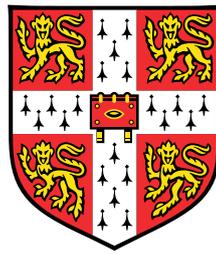


Stethoscope Acoustics



Maximilian Nussbaumer

Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Maximilian Nussbaumer

Stethoscope Acoustics

Maximilian Nussbaumer

Audible-frequency sounds from the human body are an invaluable diagnostic tool. For over 200 years, stethoscopes have been used to listen to these sounds. Despite this, the physics of how stethoscopes work remain poorly understood.

While the stethoscope itself is a simple device, its performance depends on how it is coupled to both the patient and the clinician. Existing models do not adequately address these interactions, forcing design choices to be made based on simple heuristics. The aims of this thesis are to provide a theoretical framework for understanding the acoustics of stethoscopes, propose a low order model to simulate the response, and develop an experimental methodology to validate the model.

When a stethoscope is pressed against the chest, body sounds induce small perturbations around the equilibrium position of the nonlinear chest-stethoscope system. In this thesis, a lumped-element approach is used to model these perturbations. The resulting models are validated using experiments conducted on a phantom (a laboratory model representing the human chest). Impedance measurements on the phantom and on the human chest allow differences between these systems to be accounted for.

The models presented in this thesis capture the trends associated with each of the key design parameters. Minimising the cavity volume maximises the response, while tubing significantly attenuates low frequencies and introduces distorting standing-wave resonances. Using a diaphragm attenuates the response and shifts the resonances to higher frequencies, but also allows smaller air cavities to be used. Holding a stethoscope against the chest sets the equilibrium position of the coupled system and provides a damping-dominated impedance load on the chestpiece. The strong dependence of a stethoscope's performance on external factors, such as the properties of the chest and the way it is held, makes it difficult to compare sensors objectively.

The work presented in this thesis dispels various misconceptions about how stethoscopes work and can be used to inform design choices, ultimately improving the diagnostic capabilities of future stethoscopes.

Acknowledgements

To my supervisor, Dr Anurag Agarwal, for giving me the freedom to take this research into the directions I was excited by and the guidance to ensure the quality and relevance of the work.

To the technicians at the CUED Fluids Lab, for their helpfulness and expertise. Especially to Dave Martin and Tony Lockett, and most of all to John Hazlewood, for guiding me through the nuances and practicalities of acoustics experiments, and for being great company in the acoustics lab.

To Dr Ole Nielsen, for introducing me to the electrical analogies for mechanical and acoustic systems, and for many discussions on appropriate lumped element models for the coupled chest–stethoscope system.

To Sara (Leyre) Troyas Martinez, whose Master’s project on Stethoscope Acoustics I co-supervised, for her enthusiasm and competence. The experimental results in figure 10.5 were obtained on an experimental setup built by Sara as part of her project.

To Professor Jim Woodhouse, for introducing me to experimental vibrations analysis as an undergraduate and for continuing to give me advice and assistance throughout my PhD.

To Dr Pasupathy Sivasothy at Addenbrooke’s hospital, for his immense generosity with his time, allowing us to conduct acoustic experiments during his bronchoscopies and providing invaluable medical insights.

To Dr Richard Roebuck, for encouraging me to pursue a PhD as my undergraduate director of studies and for giving me the opportunity to teach undergraduates about mechanics and vibrations.

To Nicola Cavaleri, for going above and beyond her role in providing feedback on my writing and encouragement throughout the final months of putting together this thesis.

To my colleagues at the Acousitcs Lab, especially Amélie Lamarquette, Andrew McDonald and Alastair Gregory for all the tea breaks, discussions and cake. With special thanks to Amélie Lamarquette for proof-reading much of the maths in this thesis.

To Magdalene college, for providing a home and community for over eight years since I arrived in Cambridge, and for their financial support in the final year of my PhD. To the EPSRC, for funding this research and to the Cambridge Philosophical Society for their studentship.

To my parents, for the countless ways in which they have supported me.

And most of all, to my wife Selay Pamir. I love you more than words can describe and I could fill the 300 pages of this thesis listing all of the things I want to thank you for.

Table of contents

Nomenclature	xiii
1 Introduction	1
1.1 Diagnosis	1
1.2 Auscultation	2
1.2.1 Benefits of auscultation	3
1.2.2 Issues with auscultation	4
1.3 Stethoscopes	6
1.3.1 Anatomy of the ‘modern’ analogue stethoscope	7
1.3.2 Digital stethoscope	8
1.4 Aim: a stethoscope for the 21 st century	10
1.5 Stethoscope acoustics	11
1.5.1 Problems with existing work on stethoscope acoustics	12
1.6 Objectives	13
1.7 Overview of thesis	13
2 Literature Review	17
2.1 Physical explanations	17
2.1.1 Extrapolation	17
2.1.2 Contradiction	19
2.1.3 Omission	20
2.2 Modelling approaches	21
2.2.1 Impedance-coupling and impedance-matching	22
2.2.2 Impedance matching and stethoscope performance	23
2.2.3 Methods for measuring impedance	24
2.3 Shortcomings of existing experimental work	25
2.3.1 Human experiments	25
2.3.2 Chest phantoms	27

2.4	Comparison of stethoscopes	28
2.4.1	Subjective comparison	28
2.4.2	Objective comparison	29
2.4.3	Amplitude and noise	30
2.5	Stethoscope design	31
3	Mathematical foundations	33
3.1	Potential and flow	34
3.1.1	Two-port transmission matrices	35
3.1.2	Acoustic transmission matrices	36
3.2	Lumped-element modelling of acoustical systems	38
3.2.1	Acoustic mass	38
3.2.2	Compliant cavity	39
3.2.3	LEM for a Helmholtz resonator	40
3.2.4	Application to sensors at the human chest	41
3.3	Lumped-element modelling of mechanical systems	42
3.4	Frequency response functions	44
3.4.1	Single degree of freedom damped oscillator	46
3.4.2	‘Mass locking’	48
3.4.3	‘Stiffness Shedding’	50
3.5	Electrical analogy	51
3.5.1	A note on notation	52
3.5.2	Mechanical systems	53
3.5.3	Acoustic systems	55
3.5.4	Two port formalism and transmission matrices	56
3.5.5	Boundary conditions	57
3.5.6	Coupling	57
3.5.7	Norton and Thévenin’s theorems	60
3.6	Summary	62
4	Developing a theory for stethoscope acoustics	63
4.1	Characterising stethoscope performance	64
4.1.1	Sensor impedance (‘added mass’) effects	64
4.1.2	Discrepancy between assumed and actual measurement	65
4.1.3	Non-linearity of the chest with respect to static load	66
4.1.4	Difficulty in defining a reference	66
4.2	LEM for a direct contact sensor on the chest	68

4.2.1	Sensor	68
4.2.2	Chest	70
4.2.3	Impedance analogy circuit for the coupled system	72
4.2.4	Response functions	73
4.3	Review of existing models for stethoscopes	74
4.3.1	Helmholtz resonator	74
4.3.2	Rimless contact	76
4.3.3	Rigid cavity with rim	79
4.3.4	Conclusions	83
4.4	A new theory for stethoscope acoustics	84
4.5	Derivation of a LEM for air-coupled stethoscopes	85
4.5.1	Rim	85
4.5.2	Air cavity	87
4.5.3	Chest	88
4.5.4	Holding the stethoscope	89
4.5.5	Diaphragm	90
4.5.6	Tuneable diaphragm	91
4.5.7	Complete lumped-element model	92
4.6	Circuit analogies	93
4.7	Summary	95
5	Experimental methods	97
5.1	Signal processing	97
5.1.1	Data acquisition	99
5.1.2	FRF estimators	99
5.2	Modal approach	102
5.2.1	Modal decomposition	103
5.2.2	Damping	104
5.2.3	Experimental modal analysis	105
5.2.4	Bode plot method for modal parameter extraction	106
5.2.5	Nyquist plot method for modal parameter extraction . . .	107
5.3	‘Impedance head’ for driving-point FRF measurement	108
5.3.1	Impedance head model	109
5.3.2	Model validation	112
5.3.3	Mounting stiffness	114
5.3.4	Simulated impedance head FRF measurement	115
5.4	Summary	117

6	Phantoms	119
6.1	Phantoms in the literature	120
6.1.1	Excitation	123
6.1.2	Loading	125
6.1.3	Reference signal	125
6.2	Phantom design	127
6.2.1	Materials	129
6.2.2	Geometry	129
6.2.3	Construction	130
6.2.4	Excitation	131
6.3	Phantom characterisation	131
6.3.1	Results - effect of load mass on driving point FRFs . . .	132
6.3.2	Results - effect of load contact area on driving point FRFs	136
6.3.3	Results - transfer FRF	138
6.4	2DOF LEM for the phantom	139
6.4.1	Lumped element modelling and error	142
6.4.2	Parameter extraction - varying load mass	142
6.4.3	Parameter extraction - effect of area	143
6.4.4	Performance of 2DOF LEM	144
6.5	Analysis for linear phantom behaviour	146
6.5.1	Method	146
6.5.2	Implementation	148
6.5.3	Validation - driving point impedance as a function of Z_{load}	149
6.6	Effect of sensor position on phantom behaviour	150
6.7	Summary and conclusions	152
7	Impedance measurements	153
7.1	Measuring the impedance seen by a rigid sensor on the human chest	154
7.2	Methods: balance rig setup	156
7.2.1	Data interpretation	158
7.3	Impedance seen by rigid disc on human chest	161
7.4	Impedance seen by rigid disc on phantom	164
7.4.1	Setup	164
7.4.2	Results	165
7.4.3	Effective mass trends	167
7.4.4	Effective stiffness trends	168
7.4.5	Comparison to transmission matrix for phantom	169

7.5	Attachment impedance due to holding the stethoscope	171
7.5.1	Driving-point mechanical impedance of the human hand- arm system	172
7.5.2	LEM for the hand-arm system	173
7.5.3	Experimental setup	174
7.5.4	Results	174
7.6	Summary and conclusions	176
8	Direct contact sensors	179
8.1	Direct contact sensors on the phantom	179
8.1.1	Stop-surface force reference	181
8.1.2	Thévenin impedance	184
8.2	Impedance coupling	186
8.3	Thévenin force referenced phantom response	187
8.4	Impedance coupling to the human chest	189
8.5	The effect of holding a direct contact sensor	190
8.6	Summary and conclusions	191
9	Air-coupled sensor acoustics	193
9.1	Experimental setup	194
9.2	Phantom-response of air-coupled microphones	196
9.2.1	Effect of chestpiece mass (and static load)	198
9.2.2	Effect of cavity volume	199
9.3	Interpretation of the phantom-response of air-coupled microphones	199
9.3.1	Comparison to the DCS case	200
9.3.2	Estimation of cavity-surface motion	202
9.4	Model of the ACM phantom-response	204
9.5	Stop-surface response of air-coupled microphones	209
9.5.1	Thévenin response model	211
9.6	The effect of cavity volume	214
9.7	The role of stethoscope diaphragms	215
9.8	The effect of holding an ACM	219
9.9	Summary and Conclusions	222
10	Stethoscopes with tubing	223
10.1	Coupling the mechanical and acoustic domains of an ACS	224
10.2	Attenuation in the tubing	227
10.3	Experimental setup	228

10.4 Results	229
10.5 Model	233
10.5.1 Transmission matrix between the impedance head and the air cavity entrance.	234
10.6 Application to real-world analogue stethoscopes	238
10.7 Summary and conclusions	239
11 Conclusions	241
References	245
Appendix A Derivations	259
A.1 Plane waves in ducts	259
A.2 Continuous 1D wave-equation solution for Helmholtz resonator .	262
A.3 ‘Dot Method’ for computing the dual of a circuit	263
A.4 Modal approach	264
A.4.1 FRF as modal summation	264
A.4.2 Nyquist plot - viscous damping	265
Appendix B Signal Processing	269
B.1 The discrete Fourier transform	269
B.2 Power spectral density	270
B.2.1 PSD definition	270
B.2.2 PSD estimation: Welch’s method	271
Appendix C Assumptions and details	273
C.1 Choice of FRF estimator	273
C.2 Air leaks	274
C.3 Assumption: axial microphone vibration can be neglected	276
C.4 [66] - ‘pressure response model’ for air-coupled microphones . .	279
Appendix D Simulation details	281
D.1 ABAQUS CAE Simulations	281
D.2 Figure 3.10	282
D.3 Figure 4.2	282
D.4 Figure 9.11	284
D.5 Figure 9.14	284
D.6 Transmission matrices in equation (10.1)	285
D.7 Simple duct attenuation model	285

Nomenclature

Variables

[**A**] A matrix

[**K**] Stiffness matrix

[**M**] Mass matrix

α_n Amplitude of the n th mode.

$\delta\omega/2\pi$ frequency increment between adjacent frequency bins of a DFT

$\Delta t = N\delta t$ Measurement time for a signal with $N - 1$ samples and a sampling interval of δt

δt Time-step between discrete samples

\varnothing_n Diameter of a modal circle on a Nyquist plot

η Loss factor

\hat{A} Area

\hat{a} A constant

\hat{a}_n Modal amplitude factor

\hat{B}_b Magnitude of a backwards travelling wave

\hat{B}_f Magnitude of a forwards travelling wave

\hat{p}_n	Placeholder for ω/ω_n
ι	A generic flow
ι^e	Electrical current
κ	Thermal conductivity
λ	Lumped mechanical viscous damping rate
$\tilde{\mathbf{T}}_{\text{phantom}}$	The transmission matrix between impedance-head measurements at the base of the phantom and a sensor
\mathbf{f}	Vector of external forces
\mathbf{I}	Identity matrix
\mathbf{T}	A transmission matrix
$\mathbf{T}_{\text{phantom}}$	The transmission matrix between the base of the phantom and a sensor
$\mathbf{u}^{(n)}$	Normalised modeshape for the n th mode.
\mathbf{x}	Vector of displacements
\mathcal{T}	The number of signal segments used to compute a mean
μ	Viscosity
ω	Angular frequency
ω_n	n th natural (i.e. resonance) frequency
ω_r	Damped natural frequency
ω_{-n}	n th antiresonance frequency
ϕ, ψ	Integers (typically denoting discrete points)
ρ'	Density perturbation

ρ_0	Equilibrium density
τ	A counter for segments of a signal
$z_{\phi n}$	ϕ th zero of J_n
θ	Azimuthal angle
Υ	A scaling factor for the discrete PSD.
v	A generic potential
v^e	Voltage
a	Duct radius
A', B', C', D'	Terms in a forwards transmission matrix, see equation (3.3)
A, B, C, D	Terms in a backwards transmission matrix, see equation (3.4)
b	Lumped mechanical inertance
C^e	(Electrical) capacitance
c_0	Speed of sound in air
c_n	Dimensionless modal damping ratio
c_n	Modal quality factor
d	Diameter
E	Young's modulus
e	Euler's number
F	Fourier transform of force
f	Force
f	Specific heat capacity at constant pressure

H_a	Accelerance frequency response function
H_v	Admittance frequency response function
H_x	Compliance frequency response function
i	$\sqrt{-1}$
J_n	Bessel function for the n th circumferential mode
k	Lumped mechanical stiffness
k^a	Lumped acoustical stiffness
K_d	Dynamic stiffness (inverse) frequency response function
l	Length
L^e	(Electrical) inductance
l_{\max}	Maximum spatial dimension
m	Lumped mechanical mass
m^a	Lumped acoustical mass
M_d	Dynamic mass (inverse) frequency response function
n	An integer
p	Net pressure ($p_0 + p'$)
p'	Acoustic pressure perturbation
p_0	Equilibrium pressure
q'	Volume flow perturbation
r	Radius
R^e	(Electrical) Resistance

$S(\mu)$	The DFT of discrete signal $s(n)$, where μ is the frequency bin counter
$s(n)$	A discrete signal, where n is the sample counter
t	Time
u'	Acoustic velocity perturbation
V	Volume
$w(n)$	A discrete windowing function, where n is the sample counter
X, \dot{X}, \ddot{X}	Fourier transform of displacement, velocity, acceleration
x, \dot{x}, \ddot{x}	Displacement, velocity, acceleration
y^a	Lumped acoustical admittance
$Y_{0:\hat{A}}^a$	Characteristic acoustic admittance for a duct of area \hat{A}
Y_0^s	Characteristic specific admittance
Z	Impedance frequency response function (v/l)
z	Impedance of a lumped element
$Z_{0:\hat{A}}^a$	Characteristic acoustic impedance for a duct of area \hat{A}
Z_0^s	Characteristic specific impedance ($\rho_0 c_0$)
$z_{\text{Thé}}$	Thévenin impedance
λ	Wavelength
k	Wavenumber
η	Node

Operators

\angle	phase angle of
----------	----------------

\cosh	Hyperbolic cosine function
d	The differential operator
$\ddot{g}(x)$	Second derivative of $g(x)$ with respect to its argument x
$\dot{g}(x)$	First derivative of $g(x)$ with respect to its argument x
\Im	Imaginary part
∇	Gradient operator
∇^2	Laplacian operator
∂	Partial differential operator
\Re	Real part
\sinh	Hyperbolic sine function
$\sum_{\phi=1}^{\Phi} s(\phi)$	The discrete summation of $s(\phi)$ evaluated for integer values of ϕ from 1 to Φ
$*$	Complex conjugate
t	Transpose
$f(r)$	A function of r
$g(x)$	A function of x

Superscripts

a	Acoustical
e	Electrical
m	Mechanical
s	Specific

win For a windowed signal

Subscripts

$: \tau$ for segment τ

ϕ, ψ Either ‘between ϕ and ψ ’ or ‘of ϕ with respect to ψ ’

Acronyms / Abbreviations

ACM Air-Coupled Microphone

ACS Air-Coupled Stethoscope (/Sensor)

CORSA Computerised Respiratory Sound Analysis

CPSD Cross Power Spectral Density

CT Computed Tomography

DCS Direct Contact Sensor

DFT Discrete Fourier Transform

DOF Degree Of Freedom

ECG Electrocardiogram

FFT Fast Fourier Transform (an implementation of the DFT)

FRF Frequency Response Function

LEM Lumped Element Model

PCG Phonocardiogram

PSD Power Spectral Density

SCG Seismoardiogram

SDOF Single Degree Of Freedom

SNR Signal-to-Noise Ratio

Chapter 1

Introduction

Heart and lung disease are the dominant causes of death worldwide [165]. The World Health Organisation have claimed that lung diseases “*have continuously resulted in more burden than any other diseases*” [38]. The burden is felt at both a personal level, with breathing difficulties and reduced abilities leading to a major loss in quality of life, and at a societal level, with heart and lung disease placing a significant strain on healthcare systems. Many suffers of heart and lung disease are unaware of the affliction when the disease is in its early stages. D’Arcy et al. [21] surveyed a cohort of 2500 individuals over the age of 65 in Oxfordshire, UK and found that over half of them had previously undetected (predominantly mild) valvular heart disease. Around one in twenty were found to have clinically significant valvular heart disease. In the context of an ageing population, D’Arcy et al. [21] estimate that the prevalence of clinically significant valvular heart disease will double by 2050, representing an ‘emerging epidemic’.

1.1 Diagnosis

For many of these diseases early diagnosis is crucial in enabling the appropriate medical interventions to be applied at the right time. In others (e.g. Asthma and Chronic Obstructive Pulmonary Disease) continuous monitoring of the progression of the disease is vital. ‘Diagnosis’ traditionally refers to the naming of a disease affecting a patient [75]. However, Kohn [75] points out that in modern, evidence-based medicine the aim changes from simply naming the disease to identifying the likelihood of a specific disease, while considering the quality of the

diagnostic test and assessing whether the test is worth performing. Many modern diagnostic techniques expose patients to potentially harmful levels of radiation (e.g. computed tomography (henceforth ‘CT’) scans) [74, 124]) or unpleasant invasive procedures (e.g. bronchoscopy [38]). Before performing these tests it is important to assess the pre-test likelihood of disease and to weigh up the risks and benefits of the intended procedures [75].

1.2 Auscultation

Listening to the sounds of the human body (referred to as ‘auscultation’) may be one of the oldest diagnostic techniques known to humankind. References to auscultation have been identified in the works of Hippocrates (~ 400 BC) [102], the Hindu Vedas (~ 1500-1200 BC) and the Ebers Papyrus (~ 1500 BC) [97]. For the majority of its history, auscultation has meant listening to a patient’s chest by pressing an ear against it. This practice is known as “immediate auscultation” and was never particularly widespread [89].

It was not until the invention of the stethoscope by René Laënnec in 1816, that auscultation began to become part of mainstream medical practice [132]. In the years since then, wearing a stethoscope (nowadays typically around the neck, in what Hanley and Hanley [48] refer to as the ‘cool’ position), has become a universal symbol of the medical profession [89, 97].

Laënnec did not simply invent a diagnostic device. He was at the forefront of a radical change in diagnostic philosophy: away from simply listening to patients describe their symptoms and towards the physical exam and the use of clinical correlations between observations and outcomes to diagnose disease [159]. Laënnec himself spent years correlating the sounds he heard through his stethoscope with autopsies of patients with chest disease [92, 80].

In the 21st century, objective diagnostic techniques can be split into three categories: (1) the physical exam, (2) the use of bedside diagnostic devices such as handheld-ultrasound and electrocardiogram [159], and (3) laboratory tests, high-tech imaging and invasive diagnostic procedures.

Over the past few decades a large number of advanced diagnostic techniques have become available. For the diagnosis of lung diseases, CT scans can make

use of density differences to construct a 3D map of the interior of the chest [74], while bronchoscopy allows clinicians to visually inspect the upper airways using a camera at the end of an endoscope, as well as allowing tissue samples to be taken [38]. The power and scope of these techniques has led some to claim that the stethoscope will soon be obsolete [97].

While the development of alternative diagnostic techniques has certainly diminished the once central importance of the stethoscope, the opinion that it is becoming obsolete is widely disputed (see e.g. [102, 104, 38, 36]). The clinically significant information that could be found in heart and lung sounds in the time of Laënnec is still present, and can form an important part of a diagnostic assessment. Magione [91] recalls the words of Folwer in (*Pulmonary Tuberculosis*, 1921 [35]), claiming that “*Those who advise that all stethoscopes should be ‘scrapped’ may be influenced by the fact that they do not know how to use their own*”.

1.2.1 Benefits of auscultation

In a 2011 review of techniques used for the diagnosis of lung diseases, Gautam and Pokle [38] compared four common diagnostic techniques: bronchoscopy, roentgenography (X-ray), spirometry and lung auscultation. They concluded that lung auscultation was the ‘safest’, ‘most cost effective’, ‘most accurate’, ‘simplest’, and ‘most convenient’. It appears that there is not a general consensus that the stethoscope has been superseded.

The claim that stethoscopes are the **safest** diagnostic technique is based on the fact that they are non-invasive and non-radioactive [104]. The dangers of radiation have been extensively explored and recorded (see e.g. [124]) and any diagnostic technique that makes use of radiation needs to be weighed against the radiation risk to the patient. Invasive procedures such as bronchoscopy can be very unpleasant (and potentially dangerous) for patients [63, 49]. The procedure typically involves sedation and a local anaesthetic, and also carries a risk of infection. Even spirometry (the assessment of lung function by measuring flow rates) presents a greater risk of contamination than the use of a stethoscope. This is because spirometry involves breathing into specialised medical equipment [100], while for auscultation only the external skin is in contact with the stethoscope.

It is hard to dispute that auscultation is the **cheapest** diagnostic technique [102]. Gregory [40] notes that a typical high-quality analogue stethoscope costs around £150, while a bedside ultrasound unit can cost ~ 40 times as much, and a CT scanner costs ~ 500 times as much. Furthermore, many established commercial stethoscope manufacturers now offer cheap single-use stethoscopes for disease control, and stethoscopes can be 3D-printed at low cost [127, 123]. Pavlosky et al. (2018) [123] present open-source designs for a 3D-printed stethoscope with a \$2.40 - \$3.20 material cost, although this does not factor in the cost of labour, the local cost/availability of materials, and the cost of the 3D-printer.

Whether auscultation is indeed the most **cost-effective** diagnostic technique is a more complex argument as this also depends on the utility of the outcome - this is much harder to measure, especially as several diagnostic techniques are typically used in combination. Advanced diagnostic techniques can uncover information of crucial diagnostic importance that is impossible to obtain with auscultation alone, but conversely there are important diagnostic clues that can be picked up with a stethoscope that other diagnostic tools would miss [36]. For example, Zhu et al. [172] argue that auscultation can be used to dynamically monitor the progression of pneumonia with COVID-19, while ultrasound cannot.

Finally, the argument has been put forward that stethoscopes are the **simplest** and **most convenient** diagnostic technique. Using a stethoscope does not require support from a technician, dedicated hospital space or even access to electricity. Stethoscopes are extremely portable, facilitating their use by paramedics and in remote environments. These advantages have to be weighed up against the utility of the output. The stethoscope performs one simple function: it allows clinicians to listen to the sounds produced by processes (both normal and abnormal) within the human body. This means that the utility of stethoscopes depends entirely on the user. The auditory capabilities and, equally importantly, the knowledge and experience to interpret the sounds heard, determine whether the stethoscope is a useful diagnostic device.

1.2.2 Issues with auscultation

The central issue with auscultation is the requirement for the correct interpretation of the sounds heard. This led Ertel [30] to list “the doctor” as one of the main barriers to successful auscultation. Auscultation is a skill that requires

intensive training. Recent studies have repeatably reported low auscultation proficiency among the medical profession [92, 90, 37], suggesting a lack of adequate teaching. This may be due, in part, to an increased focus on other, ‘more advanced’ diagnostic tools [102].

A lack of adequate training is, however, not the only issue. Gregory [40] notes that clinical correlations between lungs sounds and diseases are not very specific. Wheezes in particular are associated with a wide range of pathologies, which can lead to misdiagnosis [150]. Wheezes and heart murmurs can also be harmless. By some estimates 25% of the population wheezes [24], but these wheezes are not always linked to respiratory disease. Meanwhile, ‘innocent murmurs’ that are not indicative of heart disease are common in young adults [52]. Hooda et al. found that $\sim 90\%$ of a sample of 210 young adults with murmurs had no evidence of heart disease.

While the low cost of stethoscopes means that they do not represent a financial burden to clinicians in more ‘developed’ countries, cost and availability can still be a problem in poorer regions and war-zones [127]. However, it is often the access to doctors rather than stethoscopes which is the limiting factor in the availability of auscultation.

A further issue with the use of auscultation is that of hygiene. Several studies have reported on the inadequate cleaning of stethoscopes in hospitals (see e.g. [113, 86, 32]), and the recent COVID-19 pandemic in particular has led to a surge in discussion about the playoffs between auscultating patients and controlling infection [167, 172, 12] [122]. While Buonsenso et al. [12] recommend that auscultation should be replaced by the use of hand-held ultrasound devices for which infection-control is more straightforward, Zhu et al. [172] argue that despite the contamination risk, the stethoscope remains an important acoustic device.

While there are clearly many issues with auscultation as a diagnostic tool, these are certainly surmountable [159, 104]. The need for training is not unique to stethoscopes. Extensive training is also needed for the competent use of ultrasound devices [36, 140] and many advanced diagnostic techniques require entire teams of technicians [38]. The difficulty in training clinicians to interpret lung sounds can be tackled with the use of machine learning techniques to automate the detection of crucial features in the recorded sounds. These detected features can subsequently be used by clinicians to inform their diagnosis. The

development of digital stethoscopes with built-in diagnostic capabilities would also address issues associated with access to healthcare professionals in remote and disadvantaged areas. Hygiene issues are in no way unique to stethoscopes and if appropriate procedures are adopted, the risk of transmitting disease through the use of stethoscopes can be mitigated [172].

1.3 Stethoscopes

In 1816 René Laënnec was examining a young woman with symptoms of heart disease. Laënnec recounted that he was reluctant to employ immediate auscultation due to the “age and sex” of the patient. Instead, he recalled “a well known acoustic phenomenon”: that sound can be transmitted efficiently through a wooden beam. Laënnec tightly rolled a ‘quire’ (24 sheets) of paper and used this to transmit body sounds from the chest of his patient to his ear [39, 132].

Laënnec originally assumed that the sound conduction occurred predominantly through the solid body of the ‘stethoscope’, but soon came to realise that the constrained column of air played a crucial role [7]. The device Laënnec eventually settled on was a wooden cylinder with a small hole through the centre.

While the invention of the monaural stethoscope can be unambiguously attributed to Laënnec, the development of the binaural stethoscope (with sound transmission to both ears) is more controversial, with several people laying claim to the invention (see e.g. [88, 39] for details). In short, early versions of the binaural stethoscope emerged in the first half of the 19th century, often using lead pipes for the tubing [132]. In 1851, Marsh patented a stethoscope with a rubber diaphragm, and in 1855 Cammann introduced what is widely regarded as the first practical binaural stethoscope with flexible rubber tubing [132, 39, 88]. Bianchi introduced a rigid diaphragm in 1894 [132, 88], and the now classical dual mode (‘bell’ / ‘diaphragm’) stethoscope was introduced by Bowles and Sprague in 1925/6 [132, 88, 39]. Since then, the overall form of the analogue stethoscope has only marginally varied [132] although further innovations such as the tunable diaphragm (see [118]) as well as developments in manufacturing have continued to improve its performance.

At a Chinese military hospital involved in the front-line treatment of COVID-19 patients in 2020, Zhu et al. [172, 62] reported the successful use of an ‘empty

potato chip cylinder’ wrapped in a sheet of paper as a simple stethoscope that can be utilised while the ears of the clinician are covered with personal protective equipment. One could almost remark that the development of the stethoscope has ‘gone full circle’, from Laënnec’s original cylinder of rolled up paper to the ‘empty potato chip cylinders’ of 2020.

1.3.1 Anatomy of the ‘modern’ analogue stethoscope

A typical binaural analogue stethoscope is made up of a chestpiece connected by flexible tubing to a pair of earpieces as shown in figure 1.1 (see e.g. [84, 1, 67, 123]).

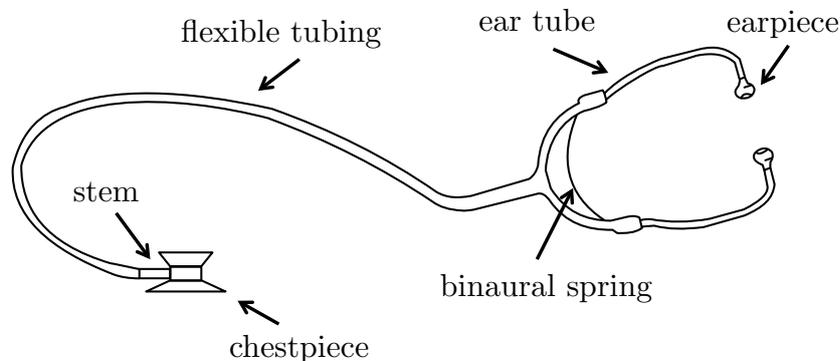


Fig. 1.1 Basic features of a typical analogue stethoscope.

The chestpiece is the part of the stethoscope that is in contact with the chest. The chestpiece can either be single or double-sided (with a rotating connector used to select which side is connected to the tubing). In the double-sided case there is typically a significant geometric difference between the two receivers: the most common design is for one side to have a relatively small contact area and a large cavity volume, while the other side has a relatively large contact area and a small cavity volume. Both, or neither, of the two sides of the chestpiece may be covered with a diaphragm. It is very unusual for the large-area low-volume side not to be covered with a diaphragm, while the low-area large-volume side is often uncovered. This has led to the popular colloquial terminology for the two sides as the ‘diaphragm side’ and the ‘bell side’ [67].

Flexible tubing connects the chestpiece to the earpieces, which are designed to form an air-tight seal at the ears. The final segment of the tubing is typically a rigid metal tube. A binaural spring is used to hold the earpieces firmly in place [7].

1.3.2 Digital stethoscope

The first electronic stethoscopes were developed in the late 19th century [88], but at the time the hardware was cumbersome and the devices were unsuitable for practical clinical use. Even though modern digital stethoscopes have evolved significantly since then (see e.g. [152]) the adoption of commercial digital stethoscopes in clinical settings has been fairly slow. This is largely due to the unnecessary additional complexity of use and the limited benefits with respect to simpler and cheaper analogue models. Typically, a stethoscope is used by clinicians as a quick ‘first-check’ in the diagnostic sequence, so increased complexity is unwelcome. Electronic stethoscopes have the advantage of allowing the storage of signals for later inspection so that changes can be observed over time and the recording can be assessed by other medical professionals. However, this is not yet common clinical practice.

Electronic recordings of heart and lung sounds also allow the waveforms to be inspected graphically. A spectrogram (a sound power plot on axes of frequency against time) of a heart sound recording is often referred to as a “*phonocardiogram*” (henceforth ‘PCG’) [69], while a spectrogram of a lung sound recording (sometimes with a simultaneous breathing trace) is referred to as a “*phonopneumogram*” [8]. Kato et al. (2009) [68] found that the visualisation of lung sounds allowed doctors to make a more accurate assessment than listening to the same sounds when assessing bilateral breath sounds to detect bronchial intubation. The core advantage of electronic stethoscopes is that they allow for computerised (and even automated) diagnosis of respiratory and cardiac disease [104, 43, 116, 69].

Unlike the analogue stethoscope, for which differences between models have been minor for the last 100 years, there are a wide variety of digital stethoscope modalities on the market and in development, with no single approach dominating the field.

Broadly speaking, digital stethoscopes can be split into two categories: ‘direct’ and ‘indirect’ contact-stethoscopes (see figure 1.2). Direct contact-stethoscopes, such as an accelerometer, measure the *absolute* motion of the chest surface [153]. In contrast, an ‘indirect’ contact-stethoscope measures the *relative* motion between two parts of the chest surface. An example of an ‘indirect’ contact-stethoscope is an analogue stethoscope [153], in which acoustic pressure perturbations arise due to differences in motion between the rim of the chestpiece and the constrained

skin surface (see figure 1.2(b)). Semmlow (2015) [135] uses a similar classification to the one we have adopted, distinguishing between sensors that use an inertial reference and sensors that use a fixed reference.

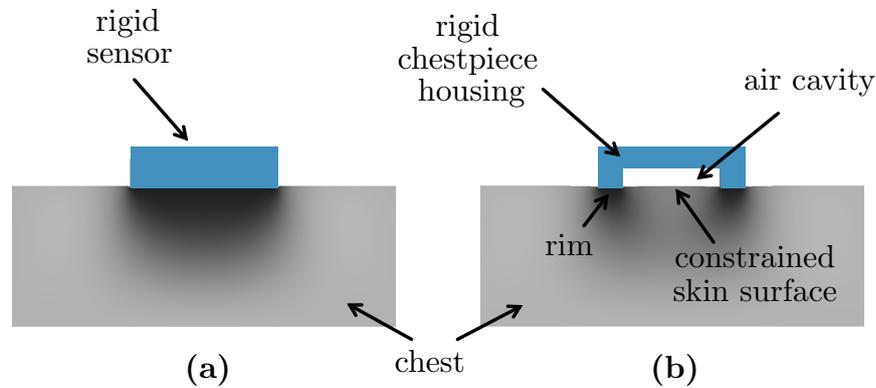


Fig. 1.2 (a) Direct contact-sensor (accelerometer) coupled to the chest. (b) Indirect contact-sensor (air-coupled stethoscope) coupled to the chest. The shading represents the vertical displacement of the chest tissue when the sensor is pressed against the chest (dark is downwards), simulated using ABAQUS CAE (see appendix D.1.)

Table 1.1 Digital stethoscope modalities.

Category	Sensor type	Examples
Direct (inertial reference ^{*1}) contact-sensor	accelerometer	[153, 119, 135]
Indirect (fixed reference ^{*1}) contact-sensor	air-coupled sensor	[117, 83, 2, 170]
	other medium-coupled sensors	[19, 50, 66]
	flexible adhesive sensor	[130, 166]
	diaphragm capacitance	[136, 137]
	piezoelectric	[106, 99]
	fixed-reference sensor	

^{*1} terminology used by Semmlow (2015) [135].

The use of lightweight accelerometers to measure the absolute motion of the chest is referred to as *seismocardiography* (henceforth ‘SCG’) (in contrast to PCG for the use of acoustic signals) [61, 146]. This area has seen a significant growth in interest over the last decade, fuelled by the availability of cheap, lightweight acceleration sensors [146]. For a detailed review of recent work in the field see e.g. Taebi et al. (2019) [146]. Jain et al. (2016) [61] compared the performance

of PCG and SCG methods by assessing the quality of recordings performed in different environments. Jain et al. concluded that air-coupled stethoscopes (PCG) perform better in a clinical setting, while lightweight accelerometers (SCG) perform better in noisy environments. Hu et al. (2014) [54] suggest that lightweight accelerometers are more suitable for wearable applications than air-coupled sensors. For recent work on wearable sensors see e.g. [22, 42, 169].

An analogue stethoscope with tubing can be converted into a digital stethoscope by using a microphone at the end of a (typically short) section of tubing [2], [170]. The commercial digital stethoscope “Eko Core” [164] also employs this method.

1.4 Aim: a stethoscope for the 21st century

Improvements in sensor-technology and the rise of machine learning can be employed to tackle many of the persisting issues with auscultation. If algorithms are able to identify abnormal features in body-sound recordings then this can help doctors to form a diagnosis without relying on extensive training and the subjective interpretation of sounds [69]. Furthermore, the doctor and the patient need not necessarily be in the same location: if the patient is able to record their own chest sounds, a doctor can access this information remotely, potentially extending the reach of healthcare to remote communities. Avoiding direct contact between patients and doctors can also reduce the risk of transmitting infectious diseases [152, 167]. In addition to one-off diagnosis, this also opens up possibilities in the monitoring of disease progression [23].

The fact that there are significantly more smartphones in the world than doctors, suggests a significant potential for tele-diagnosis and self-monitoring. In the future, doctors could even be removed from the diagnostic process entirely, although it is debatable whether this is desirable [108, 172].

Our aim is to aid the development of a stethoscope for the 21st century. This may look and function very differently from what we consider a stethoscope to be today. The key design parameters will be cost, ease of use and quality of recording, as well as compatibility with software for the interpretation of the body-sound signals obtained.

1.5 Stethoscope acoustics

Understanding stethoscope performance is a crucial aspect of automating diagnosis. Kraman et al. [79] argued that it is only by characterising the idiosyncrasies of different stethoscopes that we can compare recordings from different sensors. The same is true for the use of recordings in machine learning. Kay and Agarwal [70] argue that datasets of heart sounds recorded using different stethoscopes cannot be reliably used to train algorithms to detect abnormal body sounds. This is because if one type of stethoscope is used to record abnormal body sounds and another to record normal body sounds then it is not possible to tell whether the algorithm is in fact detecting the abnormal sounds, or just the differences between the transfer functions of the sensors. As automated auscultation begins to play an important role in the diagnostic pathway, understanding the properties of the sensors that are used to generate the data used by these algorithms is more important than ever.

Table 1.1 demonstrates that, despite the wide range of digital stethoscope sensor types described in the literature, there are two fundamental mechanisms by which body sounds can be recorded: absolute (direct) and relative (indirect) motion. In this thesis we explore the performance of a representative sensor type from each category: accelerometers and air-coupled stethoscopes. Air-coupled stethoscopes (such as the analogue stethoscope) account for the vast majority of all stethoscopes used in 2020, and form the core focus of this thesis. The experimental and theoretical techniques developed in this thesis to explore and model these two types of sensor can be extended to any of the sensor types in table 1.1.

In the words of Paul Ertel: “the stethoscope is a unique acoustical device: it is coupled to *human tissues* at both its input and output terminals”. (Ertel, 1971 [30]). Ertel identifies two couplings between the stethoscope and human tissue: the chest-chestpiece interaction and the ear-earpiece interaction. There is, in most cases, a third, which is the hand of the clinician pressing the stethoscope against the chest.

To understand the stethoscope, an investigation of the device in isolation is of little value [46]. Instead it is these three key interactions with human tissue that characterise the performance of stethoscopes.

1.5.1 Problems with existing work on stethoscope acoustics

Medical textbooks and scientific literature contain an array of contradictory and sometimes misleading claims about how stethoscopes work [108]. Findings that apply only to individual subsections of the stethoscope have been used to make statements about the performance of the entire stethoscope when coupled to the human chest [108], which is an extrapolation beyond their validity [129, 46]. Explanations found in the literature can also often contradict each other, for example on the role of the diaphragm (see e.g. [141, 67, 84, 29] versus [72, 111]).

There is little published work on modelling the entire chest-stethoscope system within a single framework. Even the most complete attempts to model air-coupled stethoscopes (such as [163] and [142, 143]) omit important aspects of the system, such as the role played by the rim of the stethoscope in constraining the chest surface [129, 67]. The most common approach to modelling stethoscopes is impedance coupling [131, 134, 82, 157, 173], but unfortunately work in this field has been plagued by several fallacies. Furthermore, Semmlow [135] argues that the common aim of ‘impedance matching’ between the stethoscope and the chest may not lead to optimal performance.

There are two main experimental approaches to stethoscope acoustics: measuring performance directly on human subjects [64, 163, 120, 41, 109, 110, 66, 129, 13, 46, 19], and using laboratory models to represent the chest [163, 157, 161, 77, 79, 19, 76, 171, 94, 51, 95, 4, 106, 105, 65, 123, 158, 73, 66, 3]. Experiments on humans are hampered by variability: both within a single subject and across subjects. This can lead to low reproducibility and uncertain conclusions. Experiments on laboratory models tackle some of these problems, but cannot capture the full nuances of a stethoscope’s use in a clinical setting [109, 110], and may thus be misleading.

The difficulties associated with experimental measurements of stethoscope performance, and the inadequacies of current modelling approaches, present difficulties when trying to optimise the design of stethoscopes for a specific application, and when comparing the performance of existing devices [46, 121].

1.6 Objectives

Inadequate explanations of how stethoscopes work hamper the ability of clinicians to choose and use stethoscopes optimally, and the ability of engineers to design the ‘best’ devices. The first objective of this thesis is therefore to:

1. Explain how stethoscopes work

As air-coupled stethoscopes are both the dominant modality and the least well understood, our primary objective is to explain how air-coupled stethoscopes work. An explanation of how direct contact stethoscopes (i.e. accelerometers on the chest) work is a secondary objective.

The second objective of this thesis is to use this physical understanding of how stethoscopes work in order to:

2. Develop a low-order model for the response of a stethoscope coupled to the chest

A complete model that takes into account all of the important influences on the response could be used in the design of new sensors, rather than following a trial and error approach.

The third objective of this thesis is to:

3. Validate the model experimentally

Our objective is to do this by developing a methodology to rigorously and reproducibly test and compare stethoscopes. This methodology has wider applications beyond validation of our modelling approach since, when designing and testing new sensors, it is crucial for the development process to have a rigorous, repeatable method for assessing performance.

1.7 Overview of thesis

The task of an analogue stethoscope is to transmit vibration signals from within the human body to the ears of a clinician as audible sound. In this thesis we break this down into a number of smaller tasks: the pressing of the chestpiece against the chest, the generation of an acoustic signal within the chestpiece, and

the transmission of this signal through the tubing. This thesis addresses these tasks in turn, culminating in a discussion of the entire analogue stethoscope in chapter 10. A benefit of this approach is that several of these ‘subtasks’ of an analogue stethoscope independently correspond to feasible electronic stethoscopes. For example, devices consisting of only a chestpiece, without tubing (so called ‘air-coupled microphones’), and devices without any air-coupling (so called ‘direct-contact sensors’). The chapters dedicated to these devices have significance both independently and as building blocks for the discussion of analogue stethoscopes.

In this thesis we introduce a new theory for air-coupled stethoscope acoustics which takes into account the intricacies of the coupling between the stethoscope and the chest. Unlike existing work, our model takes into account the key parameters that determine stethoscope performance within a single cohesive model. We also introduce a model for direct contact sensors (accelerometers) on the chest using a framework that is consistent with our model for air-coupled stethoscopes.

We validate our theory using experiments on a carefully designed phantom (a simple laboratory model representing the chest). Unlike existing work, we take into account the effect of the properties of the phantom on the transfer function between the excitation applied to the phantom and the response of the stethoscope in a rigorous manner.

In chapter 2, we review existing work in the field, highlighting issues to be addressed. We then lay out the mathematical foundations for the analysis performed in this thesis in chapter 3. Given the multi-disciplinary nature of stethoscope acoustics, which lies at the interface between practical medicine and acoustic engineering, the core mathematical techniques that we have employed are laid out in considerable detail. Readers familiar with the content may wish to skim this material and refer back to it as required.

In chapter 4, we develop new models for accelerometers and air-coupled stethoscopes coupled to the chest. The experimental techniques which we use to validate these models are outlined in chapter 5. In chapter 6, we discuss the design, construction and performance of a chest-phantom. We then explore the behaviour of the phantom when subjected to different loads in order to account for the role of the phantom in the frequency response functions obtained for different stethoscopes.

In chapter 7, we perform a series of driving-point impedance measurements on the phantom and on a representative human chest to investigate the mechanical properties of the chest ‘seen’ by sensors. In chapter 8, we investigate the performance of direct sensors on the phantom, and explore the use of impedance-coupling techniques to predict sensor performance, using the measurements presented in chapters 6 and 7. We validate our model for direct coupled stethoscopes on the chest and discuss the implications for the performance of this class of sensors.

In chapter 9, we use the phantom described in chapter 6 to characterise a set of air-coupled microphones and validate our new model for the acoustic performance of stethoscopes. We explore the effect of varying cavity volume, the use of a diaphragm, and the effect of holding the sensor. In chapter 10, we extend our analysis to include stethoscopes with tubing by coupling impedance analogy circuits for the mechanical and acoustic domains. We validate the proposed model experimentally and identify the dominant factors to consider when designing air-coupled sensors including tubing. Chapter 11 outlines the conclusions of this thesis and highlights areas of useful further work.

Chapter 2

Literature Review

2.1 Physical explanations

Shortfalls with current physical explanations of how stethoscopes work fall into three categories: extrapolation, contradiction and omission. Many explanations consider the acoustic properties of individual components of the stethoscope and *extrapolate* from this to draw conclusions about the entire coupled chest–sensor system. Explanations available in the literature also often *contradict* one another and, finally, many explanations of Stethoscope Acoustics that can be found in textbooks and journal articles on the subject *omit* some of the key factors affecting the performance of stethoscopes, such as the non-linearity of the chest. This literature review evaluates all three of these shortfalls in detail and explains why current models are problematic when considering stethoscope design.

2.1.1 Extrapolation

In his 2015 thesis on ‘Stethoscope design for auscultation in high noise environments’, Nelson [105] argues that a stethoscope can be broken down into a number of subsystems that operate independently. Nelson argues that a standalone analysis of these subsystems is “both valid and advantageous”. This claim suggests that if we can describe a transfer function from the vibration of the chest to the pressure in the air-cavity of a stethoscope’s chestpiece, then we can simply add on the transfer function between the air-cavity pressure and the ear-piece to

obtain the overall response. This approach, however, oversimplifies the notion of transfer functions in acoustics.

As early as 1941, Rappaport and Sprague [129] stated that the air-cavity in the stethoscope chestpiece “exhibits acoustic characteristics that are entirely different” depending on whether it is held in the open or against the chest. Hampton and Chaloner [46] generalised this, arguing that “the acoustic properties of any component are very different from those of a stethoscope incorporating that component”. Each component of the interconnected chest-stethoscope system provides loads and boundary conditions on the other components. These loads and boundary conditions change the transfer functions of the components. The stethoscope can be described as a *strongly coupled* system, which means that its resonant phenomena are determined primarily by the coupling rather than the resonances of the individual components. One exception to this is the tubing of the stethoscope. In the overall response the tubing-resonances are clearly distinguishable (see e.g. [29]).

Ertel et al. [29] found that the presence of a diaphragm can attenuate sounds by over 30 dB when the stethoscope is held over a speaker. It is to be expected that the air–diaphragm–air pathway significantly impedes acoustic waves, but this has little relevance for the application of the stethoscope to the chest. Nowak and Nowak [108] highlight this as an example of the perils of extrapolation in the context of stethoscope acoustics and point out that, despite the fact that any clinician who has used a stethoscope will be able to confirm that a diaphragm does not attenuate sounds by 30 dB, this result has continued to be quoted out of context over 40 years later.

Kindig et al. (1982) [72] also made a number of observations on the performance of stethoscopes based on an experimental setup in which sound was played into the chestpiece through a small air-cavity. For example, Kindig et al. [72] found that when using this experimental setup there is little difference between deep and shallow bells. The effect of coupling to the chest was not taken into account, and there is no evidence that these findings can be extrapolated to the use of stethoscopes on the chest. Nevertheless, the result that there is little difference between deep and shallow bells is presented in the medical textbook *Evidence-Based Physical Diagnosis* (McGee (2016) [98] with no reference to the experimental setup for which this claim is valid.

2.1.2 Contradiction

There are several areas within stethoscope acoustics where there are contradictions between claims made in the literature that have not been resolved. Perhaps the most prominent of these revolves around the role of the diaphragm. The diaphragm is often described as a ‘high-pass filter’ (see e.g. [141, 67, 84]). Ertel [29], for example, found that the presence of a diaphragm attenuates sounds across the entire frequency spectrum, but that the effect is most pronounced at low frequencies.

In a set of experiments in which the chest was represented by a small air-cavity excited with a speaker, Kindig (1982) [72] found that the use of a diaphragm lead to attenuation at all frequencies. This is misinterpreted in the medical textbook *Evidence-based physical diagnosis* (McGee, 2016) [98], which states that a diaphragm ‘attenuates all frequencies equally’. Equal attenuation at all frequencies is highly implausible for any physical system, and was not claimed in the original source material by Kindig.

A more serious challenge to the ‘high-pass filter’ model for stethoscope diaphragms comes from Nowak and Nowak (2019) [111]. Using a laser vibrometer to measure the vibration of diaphragms pressed against the human body, Nowak and Nowak concluded that the diaphragm does not act as an acoustic filter. However, we suggest that any component added to a mechanical system will have some effect on the frequency response of that system. While describing the diaphragm as a ‘high-pass filter’ may be too simplistic, adding a component which is stiffness-dominated in the frequency range of interest to a system tends to increase the resonant frequencies of that system (Rayleigh’s theorem). The fact that Nowak and Nowak were unable to observe this effect in a carefully designed experimental study suggests that there is a need for further investigation in this area.

The effect of application pressure has also been the subject of some dispute. Rappaport and Sprague (1941) [129] suggest that when an ‘open bell’ is pressed against the chest the skin forms a ‘natural diaphragm’. As the application pressure is increased the pre-stress of this ‘natural diaphragm’ increases, leading to filtering effects. Focussing on the ‘diaphragm’ effects of the skin can be misleading as it neglects the role of the underlying chest tissue. Nowak and Nowak [107, 111] note that the deformation of the chest tissue under the stethoscope is likely to play a dominant role in determining the effect of application pressure. Pasterkamp et

al. [120] performed a series of measurements with very low application pressures by using double-sided adhesive tape in order to attach stethoscope chestpieces to the chest. Pasterkamp et al. argue that it is unlikely that the skin acts as a pre-stressed diaphragm in this case.

2.1.3 Omission

Many factors that play an important role in the performance of stethoscopes are largely missing from the narrative, or only mentioned in passing. Examples include the non-linearity of the chest, ‘occlusion’, and the role of the stethoscope in determining the motion of the portion of the chest surface that it is in contact with.

Wodicka et al. [163] as well as Suzuki and Nakayama [142] model the interaction between a stethoscope and the human chest by considering the mechanical properties of the chest as a function of the contact area only. However, published research on the mechanical impedance of the chest shows that the properties of the chest are non-linear with respect to the load applied and the resulting equilibrium position (see e.g. [56, 148, 154, 5]).

‘Occlusion’ refers to the way in which the air-cavity fills with flesh as the stethoscope is pressed against the chest [129]. This reduces the contact area between the chest and the air [129], resulting in a lower transfer of vibrational energy. The effect is especially significant in shallow cavities, which explains why the air-cavity on the ‘bell-side’ of commercial chestpieces is typically significantly deeper (and steeper) than the cavity behind the diaphragm [67]. A simple experiment can illustrate this effect: when the diaphragm on the ‘diaphragm-side’ of a typical analogue stethoscope chestpiece is removed, and the now open cavity is pressed against the chest with moderate pressure, the entire cavity fills with chest tissue. Only the small hole that leads to the tubing is left to couple the chest-vibration to the tubing. In this case, almost no sound is heard at the earpieces. Despite this, the use of diaphragms is often assessed by comparing the performance of the same air-cavity with and without a diaphragm (see e.g. [29]).

Rappaport and Sprague [129] report that many clinicians regard the bell simply as an accumulator of sound, neglecting the role of its impedance in determining the properties of the coupled chest-stethoscope system. These comments were

made around 80 years ago, but general audience explanations on the internet still describe the role of the chestpiece in a similar way (e.g., *HowStuffWorks.com*: “vibrations hitting the chestpiece are funnelled into a narrow tube” [81]) No mention is made of the fact that the ‘vibration hitting the chestpiece’ depend, in part, on the properties of that chestpiece and the way in which it is applied to the chest [135, 153].

2.2 Modelling approaches

Beyond the general comments and qualitative descriptions of stethoscope performance discussed in section 2.1, several authors have tried to model the performance of stethoscopes quantitatively. This is typically based on one of two approaches:

- Decomposing the stethoscope into subsections and analysing these individually (without considering the interaction between subsections).
- Deriving an impedance-based model for the coupled system.

Modelling attempts based on decomposition are problematic because of the extrapolation issues raised in section 2.1.1. Once a subsection is incorporated into the full system, its boundary conditions (and thus its transfer function) will change. Rappaport and Sprague [129], as well as Kaniusas [67], derive expressions for the resonant behaviour of a clamped diaphragm that is free on both sides and use this to make general statements about the performance of stethoscopes with diaphragms. This is another version of the mistake made by applying Ertel’s experimental results for a diaphragm in air to the case where the chestpiece is pressed against the chest (discussed in section 2.1.1).

Kaniusas [67] also models the air-cavity in a stethoscope chestpiece as a Helmholtz resonator, and uses this to make claims about how parameters such as the cavity volume and tubing length affect performance. In section 4.3.1 we show that describing the air-cavity of a stethoscope as a Helmholtz resonator does not accurately capture its role.

An impedance-coupling based approach has the potential to address all of these issues and provide an accurate model of the physics of the entire coupled chest-stethoscope system. Several authors have attempted an impedance-coupling

based approach (see e.g. [131], [163] and [143]), but none of these models adequately tackle the subtleties of the chest-stethoscope interaction. A detailed review of existing impedance-based models is presented in section 4.3.2.

2.2.1 Impedance-coupling and impedance-matching

Impedance coupling refers to a mathematical technique for modelling the coupling between two systems by considering their impedances. *Impedance matching* refers to the design technique of selecting impedances in such a way as to maximise power transfer. In chapter 3 we present the mathematical background to impedance analysis in detail. For now we limit ourselves to one important clarification: the *source impedance* is the impedance of the source *as ‘seen’ by the load*, and the *load impedance* is the impedance of the load *as ‘seen’ by the source*. In the context of this clarification and the definition of the maximum power transfer theorem we outline three common misconceptions in the application of the impedance-matching method to stethoscopes.

Misconception 1: matching characteristic specific impedances

Characteristic specific impedance is purely a property of a material [10]. The impedance of a *finite* physical system does not simply equal the characteristic impedance of the material of which it is composed. The importance of the physical dimensions of a load system on the impedance ‘seen’ at the boundary depends on the frequency of interest. We can capture this in terms of the *Helmholtz number* which is the product of the wavenumber of interest and a characteristic spatial dimension. A high Helmholtz number implies a non-compact body in which the spatial dimensions are large compared to the wavelength of interest. If the chest is coupled to a system with a high Helmholtz number, then the waves propagating from the source effectively only ‘see’ the characteristic impedance of the load. If, however, the Helmholtz number is low (as is the case for stethoscopes at the frequencies of interest for auscultation), then the dimensions of the systems play an important role in the impedance coupling. Scharztz et al. [134] use characteristic specific impedances to model the transmission losses between the chest and the chestpiece cavity. A consideration of the Helmholtz number suggests that this analysis is not appropriate for typical auscultation frequencies. Similarly, Li et al. [82] argue that the proportion of the heart sound signal lost between the chest and the sensor depends on the impedance difference

between the two *mediums*. If this approach was valid then there would be no significant difference between using an air-coupled stethoscope and just standing close to the patient. Over 200 years of clinical practice suggests that this is not the case [39].

Misconception 2: neglecting contact characteristics

Given that the coupling depends on the impedances of the load and source systems ‘seen’ at the contact, one would expect the characteristics of the contact to play a significant role. The impedance of the chest ‘seen’ by an air cavity (with the surface able to deform non-uniformly) may be quite different from the impedance seen by a rigid disc. Wodicka et al. [163] and Suzuki and Nakayama [142, 143] both use the impedance ‘seen’ by a rigid disc and couple this to the impedance of the air cavity of the stethoscope. This approach also neglects the presence of the rim and its effect on the impedance ‘seen’ by the air cavity. In section 4.3.2 we discuss these models in more technical detail.

Misconception 3: comparing impedance magnitudes only

When coupling impedances, the phase of the impedances is crucial in determining the response of the coupled system. Impedance matching between a source and a load is only achieved if the imaginary components of the impedances cancel. Watrous et al. [157] and Zimmerman et al. [173] take into account the contact characteristics by measuring the impedance ‘seen’ by an air cavity surrounded by a rim; however, they only report and discuss the magnitude of the impedance measured. Watrous et al. [157] suggest that when the magnitude of the chest impedance is similar to the magnitude of the stethoscope’s impedance, power transfer to the stethoscope is efficient. However, similar magnitudes of impedance do not imply successful impedance matching, as this depends on the complex values of the impedances.

2.2.2 Impedance matching and stethoscope performance

Many papers that consider stethoscope acoustics from using an impedance approach assume that impedance matching is desirable (see e.g. [134], [157] and [82]). However, this may not be the case. Semmlow (2015) [135] argues that while impedance matching could improve the power transfer it would also reduce the signal levels at the transducer, ultimately leading to a lower output signal.

Vermairen and van Vollenhoven [153] and von Gierke and Brammer [154] show that using heavy accelerometers leads to a decrease in the high-frequency response. At frequencies far from any coupled resonances, the greatest possible acceleration level is obtained by minimising the load impedance, not by matching it to the chest. Semmlow [135] suggests that a low-mass vibration sensor may therefore outperform current stethoscopes at high frequencies.

2.2.3 Methods for measuring impedance

The mechanical and/or acoustical impedance of both the chest and the stethoscope can be measured experimentally using a variety of techniques. Mechanical impedance (the ratio of force to velocity) is typically measured by exerting a measured force on a structure and simultaneously measuring the driving point acceleration. A device designed to simultaneously measure force and driving point acceleration is called an *impedance head* ([87]). See e.g. [56], [148], [154], [153] and [5] for attempts to measure the mechanical impedance of human tissue.

Dimensionally, acoustic impedance is simply related to mechanical impedance by the contact area. However, when we refer to *acoustic impedance* we mean the acoustic impedance ‘seen’ by the air cavity. The deformable surface and the presence of the rim mean that the acoustic impedance differs significantly from the impedance seen by a rigid disc.

The acoustic impedance seen by the air cavity can be measured directly using an *impedance tube* [173, 174]. An impedance tube is a duct with at least two microphones and an excitation source [6]. For frequencies at which only plane waves propagate (see section A.1) these two microphones can be used to decompose the sound-field into waves travelling ‘forwards’ and ‘backwards’ through the duct. This technique is called ‘wave separation’. Once the sound-field has been separated into forwards and backwards travelling waves, the acoustic impedance can be calculated at any point in the duct. If a speaker is connected to one end of the tube and the other end is pressed against the chest then the setup can be used to determine the acoustical impedance of the chest as ‘seen’ by an air cavity with a contact area and rim-dimensions matching those of the tube. Watrous et al. [157] and Zimmerman et al. [173, 174] present results for impedance-tube measurement on the chest, but only consider the impedance magnitude (see section 2.2.1 for why this is an issue).

2.3 Shortcomings of existing experimental work

There are two key experimental approaches to stethoscope acoustics:

1. Measuring stethoscope performance directly on humans
2. Conducting experiments on laboratory setups designed to replicate the mechanical properties of the chest ('phantoms')

2.3.1 Human experiments

Experiments on humans subdivide into two categories: those that take human biosignals such as heart or lung sounds as the input signal, and those that use external excitation such as shakers attached to the body or sound played into the lungs.

In the context of stethoscope acoustics the term 'biosignals' is used to refer to vibrations within the human body that occur as a result of natural processes. These natural processes include the beating of the heart, breathing and digestion [67]. There have been many attempts to use biosignals as the input when evaluating the properties of stethoscopes (see e.g. [64, 163, 120, 109]). However, performing experiments on the human chest can be problematic.

Unfortunately, anatomical differences between people severely hamper the reproducibility of results based on biosignals [77] unless large representative datasets are collected that can account for variability in the population. Body sound signals do not just differ from person to person. Within a single subject they differ depending on location, breathing style and heart rate. Furthermore, the coupling of sensor to the chest is found to be very sensitive to the way in which the sensor is applied. Attaching the stethoscope with tape, resting it on the chest with a load mass to keep it in place and applying it by hand with a range of contact pressures all result in different coupled chest-stethoscope systems and therefore different output signals. Considering this extensive variability, the ability to generate data of a reproducible scientific standard (rather than anecdotal comparisons) is limited.

In the year 2000 a set of "International Computerised Respiratory Sound Analysis" (henceforth CORSA) guidelines were published [126]. The CORSA guidelines

specify information that should be included when reporting the results of lung-acoustics studies. This includes a careful description of the sensors used as well as a record of the body posture, airflow, breathing rate and environmental background conditions during the recordings. A systematic review carried out in 2014 by Oliveira and Marques [115] found that none of the reviewed papers published after the publication of the CORSA guidelines had followed the recommendations made. Reproducibility in this field of science is typically low.

Nevertheless, Nowak and Nowak [108] argue that the only truly valid approach is to test stethoscopes in their operating environment: coupled to the human chest. Their claim is that no other method can accurately replicate the stethoscope's use in a diagnostic situation. Nowak and Nowak have published a large body of experimental work on the performance of the analogue stethoscope. In order to compensate for the inherent variability of human biosignals, Nowak and Nowak simultaneously recorded synchronised Electrocardiogram (ECG) signals. They then used the ECG signals to segment the acoustic heart sound signal into individual beats which they referred to as 'acoustic events'. Nowak and Nowak selected a subset of the extracted heartbeats (using a similarity coefficient) and computed an averaged spectrum from this subset of 'acoustic events'. The resulting averaged spectra could be compared between sensors to allow the effect of various design parameters to be observed [109], [110].

This kind of attempt to compensate for variability in acoustic biosignals is not new. Wodicka [163], Pasterkamp [120] and Jones [64] all used natural lung sounds as a sound source in their studies. In all three cases, the breathing flow rate was kept within a specified range in an attempt to maintain a constant vibration input. Though basic, this still represents an improvement on research which uses natural lung sounds with no attempt to measure the flow [108] [107] [41].

Heart and breathing rates are only a limited proxy for the vibration of the chest surface. Another common technique has been to simply test two sensors side-by-side in the hope that the signal will be similar at these locations (see e.g. Joyashiki and Wada (2020) [66]). Alternatively, the sensors can be tested sequentially, with a simultaneous reference sensor at a fixed location providing a reference signal [64]. This does not take into account the non-uniformity of the chest and the fact that the vibrations transmitted to the reference sensor could differ significantly from those at the test sensor (both in an averaged sense and,

more importantly, over time). This could occur if, for example, the reference sensor picks up almost entirely heart sounds, but the test location also has a strong presence of lung sounds (or vice-versa).

As an alternative to using human biosignals, measurements can be taken with sensors coupled to the human chest but using an external, controllable source of excitation. Burch applied sensors and an excitation source to the thigh in order to minimise contamination from body sounds [13]. Rappaport and Sprague used a loudspeaker to excite the external chest surface [129]. Hampton and Chaloner [46] tried to replicate the method of Rappaport and Sprague but found that it was not possible to obtain a soundproof seal between a loudspeaker and the chestwall. Hampton and Chaloner [46] instead used a mechanical shaker to excite the chest. Dai et al. introduced a controlled sound input into the lung by playing sound through a tube into the mouth [19]. However, this study was concerned primarily with investigating mechanisms for sound transmission through the chest rather than the coupling to sensors at the chest surface.

2.3.2 Chest phantoms

The term ‘phantom’ refers to a device designed to replicate some features of the human body in order to test, calibrate and train clinicians in the use of diagnostic devices [17]. Phantoms are commonly used for a variety of medical imaging methods such as Nuclear Magnetic Resonance Imaging, computed tomography and ultrasound [156]. Depending on the application for which the phantom is used it is designed to mimic different properties of the human body. This could include electromagnetic properties, mechanical properties and ultrasound absorption [156], as well as the speed of sound, and the acoustic attenuation of the medium [17].

Imaging phantoms often simulate not only the contact of the sensor with the skin but also a range of underlying anatomical structures to be detected. Dai (2013) [19] (see also [20]) applied this technique to lung sounds by constructing a cuboid phantom made of a viscoelastic silicone elastomer (Ecoflex 00-10). The cuboid phantom was filled with a network of horizontal airways with dimensions typical for the upper airways of adults. Sound was played into the airways and measured in the tube network using a microphone. A scanning laser doppler vibrometer was used to measure the resulting surface motion.

The word ‘phantom’ has also been adopted for devices designed specifically to characterise the acoustic properties of sensors coupled to the skin rather than to mimic their use to detect vibrations from specific sources (see e.g. [19] [94] [106] [158] and [73]). Other names for such devices include ‘BioAcoustic Transducer Testing system’ [77], ‘simulator’ [161, 66] and ‘sensitivity measurement system’ [65]. The latter is potentially misleading as the word ‘sensitivity’ has a strict definition in metrology, as discussed in section 4.1.

The benefit of a phantom over direct measurements is that, provided the procedure is adequately documented, the results should be reproducible by any researcher in any laboratory. Unfortunately very little existing work on using phantoms to characterise stethoscopes is adequately documented. Another major benefit of using phantoms is that a reference input signal can be measured and used to compare sensors rigorously. A detailed overview of studies using phantoms to characterise stethoscope performance is presented in section 6.1.

2.4 Comparison of stethoscopes

2.4.1 Subjective comparison

There have been numerous investigations on the subjective performance of stethoscopes. As the subjective experience is influenced by a wide variety of non-performance related factors (such as ergonomics, appearance, familiarity and even branding) asking clinicians about their preferences is of limited scientific value. An alternative is to compare the use of different stethoscopes to diagnose disease. However, a diagnosis is ultimately qualitative, and depends on a huge variety of factors beyond the auscultation. Furthermore, if the same clinician used two different stethoscopes on the same patient it would be impossible for the information from the first auscultation not to influence the second. Liu et al. [85] solved this problem by using stethoscopes for quantitative blood pressure measurements. Liu et al. found a statistically significant (but clinically insignificant) difference in measured blood pressure values between different stethoscope designs.

Pasterkamp et al. (1997) [121] describe stethoscopes as filters acting on body sounds, applying different weightings to different frequency components [121].

Pasterkamp et al. argue that auscultation would “benefit from a more faithful representation of sound”. However, others have shown that the filtering characteristics of the stethoscope are in fact useful for human interpretation [67, 160, 129]. Semmlow (2015) [135] suggests that high sensitivity (across the relevant frequency range) is more important than a flat response when detecting faint chest sounds.

The high-pass-filter effect often attributed to the ‘diaphragm mode’ of an analogue stethoscope (see section 2.1.2 for controversies related to this claim) has been deemed useful for listening to lung sounds due to a phenomenon called masking [160, 30]. When a person hears a combination of high and low pitched sounds, the low pitched sounds mask the high pitched sounds. As lung sounds are typically slightly higher in pitch than heart sounds this means that, without a diaphragm, the masking effect can make lung sounds very hard to detect. A study into the clinical practice of clinicians and medical teachers found that a significant proportion of stethoscope users choose the bell to detect lung sounds despite the lack of unmasking [161]. Welsby suggests that this may be due to “selective hearing” – the ability of individuals to focus on sounds in which they are interested and to block out other sounds [161].

2.4.2 Objective comparison

Despite their ubiquitous use, there is no simple or standardised way to objectively compare stethoscopes or evaluate new designs. This is largely due to the fact that an evaluation of the standalone acoustic properties of a stethoscope (as performed e.g. by Abella et al. [1]) is insufficient in assessing the performance of the stethoscope when it is coupled to the human body.

To compare any two sensors, we need to know their outputs relative to a common reference input signal. This can be achieved in two ways:

1. Simultaneously measuring the sensor output signal and a reference input signal for a number of sensors, and then comparing the ratios of sensor output to measured input between sensors.
2. Ensuring that the input signal remains constant for each recording performed and then directly comparing the sensor output signals.

When applying this to stethoscopes, there is no obvious answer as to what the ‘reference input signal’ should be. The ultimate input signals we are interested in measuring are the vibrations induced by the heart and lungs. The way in which these signals are transmitted to the surface depends on sensor location, subject anatomy and even posture [53, 33]). Furthermore, the soundscape at the chest surface consists of a superposition of body sounds originating from multiple locations within the chest (as well as external sound and vibration [133].)

2.4.3 Amplitude and noise

For analogue stethoscopes the absolute amplitude of the sound at the clinician’s ears is important. Heart and lung sounds need to have amplitudes above the auditory threshold in order to be of any use [129]. For digital stethoscopes, however, the absolute amplitude of the signal is unimportant as it can be amplified to an arbitrary degree. In this case, we are interested not in the signal amplitude but rather in the signal-to-noise ratio (SNR). Li et al. [82] note that the SNR does depend in part on physical amplitude of the signal at the transducer.

Since we are interested in SNR, it would make sense to use it to characterise sensors (see e.g. [82]). For a rigorous definition of SNR we need to know the input signal, but we can also generate a crude estimate of the SNR by considering the output signal only if we can segment the output signal into segments that include the signal and segments that include only background noise.

Noise in stethoscope recordings stems from a number of different sources. The electrical circuitry in digital stethoscopes induces random electrical noise [31]. There are also several physical noise sources to consider. Nelson [105] presents a detailed overview of noise sources in stethoscope recordings. Nelson distinguished between ‘vibration noise paths’ and ‘acoustic noise paths’ (involving airborne ambient sound). The primary ‘vibration noise paths’ identified by Nelson [105] are transmission through the patient and transmission through the clinician holding the stethoscope. For a traditional stethoscope the primary ‘acoustic noise paths’ are leakage at the ears, leakage through the stethoscope and transmission of ambient noise through the patient’s chest. Suzumura and Ikegaya [144] and Schmidt et al. [133] found that the chest acts as a significant transmitter of ambient noise to a stethoscope. Zanartu et al. [171] and Joyashiki and Wada

[66] characterised sensors in terms of the ratio of their sensitivity to tissue-borne and air-borne sound. The findings of e.g. Schmidt et al. [133] suggest that it may be difficult to decouple these two pathways.

2.5 Stethoscope design

Unfortunately, the incomplete understanding of the interaction between different aspects of the stethoscope means that current theoretical work is of limited use in stethoscope design. Hampton and Chaloner [46] concluded in 1967 that “most of the previously accepted principles should be ignored and an empirical approach should be used”. We suggest that, instead of accepting the inadequacy of current models, a new theory for stethoscope acoustics is required. This thesis addresses this need with a rigorous consideration of the coupling between the stethoscope and the chest.

Chapter 3

Mathematical foundations

Sensing the vibration of the human chest and, in particular, interpreting the signals obtained is not straightforward. The signal obtained depends on the coupling between the source (the chest) and the receiver (the sensor). This coupling is affected by the anatomical and material properties of the chest, the design of the sensor, and the way in which the sensor is applied to the chest.

The interlinked complexity of the chest-stethoscope system has slowed the development of a robust understanding of the phenomena involved. Over two hundred years since the invention of the stethoscope ushered in the world of evidence-based medicine, we still do not truly understand the sensors we are using. In the past, clinicians were trained to subconsciously compensate for the idiosyncrasies of the stethoscope's acoustic response. We are now, however, in an age where digital sensors and automated diagnosis are beginning to revolutionise healthcare. As a result of this, an understanding of how bio-acoustic signals depend on the sensors we use to measure them is more important than ever.

The lumped element model (LEM) approach discretises a system to a small number of degrees of freedoms (DOFs) and describes the way in which these DOFs are linked using 'lumped' terms to represent the physics of the system. The number of degrees of freedom needed to model the behaviour of a system depends on the frequency range of interest and the physical dimensions of the system. By comparing the wavelengths of body sounds to the physical dimensions of stethoscopes, we can establish an appropriate discretisation. Apart from drastically simplifying the system we are investigating, a lumped element

formulation also allows mechanical and acoustical systems to be coupled in a straightforward manner within a single mathematical framework.

3.1 Potential and flow

The governing equations for electrical, mechanical and acoustical systems can all be expressed in terms of a *potential* (v) and a *flow* (ι). As modern digital stethoscopes stand at the cross-section of all three of these domains, developing a mathematical framework in these terms seems sensible.

The ratio of *potential* to *flow* is defined as impedance:

$$Z = v/\iota . \quad (3.1)$$

The physical dimensions of impedance vary depending on the physical properties playing the roles of *potential* and *flow*. In the electrical domain, for instance, voltage (v^e) plays the role of *potential* and current (ι^e) plays the role of *flow*. The electrical *impedance* can be expressed in terms of inductance (L^e), resistance (R^e) and capacitance (C^e). Equation 3.1 then becomes Ohm's law: $v^e = \iota^e Z_e$ [6]. Table 3.1 summarises the definitions of *potential*, *flow* and *impedance* in the acoustic, specific, mechanical and electrical domains.

Table 3.1 Potential, flow and impedance in different domains. All symbols are as defined in the nomenclature: p' , q' , f and u' denote pressure, volume flow, force and velocity, ρ_0 and c_0 denote a medium's density and speed of sound, \hat{A} represents area, and the superscripts a , s , m and e denote the acoustic, specific, mechanical and electrical domains respectively.

Domain	<i>potential</i>	<i>flow</i>	impedance	Z_0	units
Acoustic	p'	q'	Z^a	$\rho_0 c_0 / \hat{A}$	$\text{kg m}^{-4} \text{s}^{-1}$
Specific	p'	u'	Z^s	$\rho_0 c_0$	$\text{kg m}^{-2} \text{s}^{-1}$
Mechanical	$f = p' \hat{A}$	u'	Z^m	$\rho_0 c_0 \hat{A}$	kg s^{-1}
Electrical	v^e	ι^e	Z^e	$\sqrt{(L^e/C^e)}$	ohms

In general, an impedance can be defined between a potential at any point ϕ , and a flow at any point ψ . This is denoted by

$$Z(\phi, \psi, \omega) = Z_{\phi, \psi}(\omega) = v_{\phi}/\iota_{\psi}. \quad (3.2)$$

The *characteristic impedance* (Z_0) is defined as the ratio of *potential* to *flow* (at the same location) when waves are travelling in one direction only (such as in an infinite transmission line). General expressions for the characteristic impedance in different domains can be found in table 3.1.

3.1.1 Two-port transmission matrices

A *port* represents a point within a system that is associated with a *flow* and a *potential*. When describing the relationship between two ports within a system we are describing the relationship between four variables. There are $4! = 24$ different ways of expressing these relationships in terms of a 2×2 matrix. Two such matrices are the ‘forwards’ and ‘backwards’ *transmission matrices* which relate the flow and potential at one port to the flow and potential at the other port. The ‘forwards’ (or $A'B'C'D'$) transmission matrix between points 1 and 2 is given by

$$\begin{bmatrix} v_2 \\ \iota_2 \end{bmatrix} = \begin{bmatrix} A' & B' \\ C' & D' \end{bmatrix} \begin{bmatrix} v_1 \\ \iota_1 \end{bmatrix}, \quad (3.3)$$

while the ‘backwards’ (or $ABCD$) transmission matrix is given by

$$\begin{bmatrix} v_1 \\ \iota_1 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} v_2 \\ \iota_2 \end{bmatrix}. \quad (3.4)$$

The advantage of a transmission matrix formulation is that transmission matrices can be cascaded, such that a complex cascaded network can be constructed from a series of simple transmission matrices:

$$\begin{bmatrix} v_3 \\ \iota_3 \end{bmatrix} = \begin{bmatrix} A' & B' \\ C' & D' \end{bmatrix}_{3,2} \begin{bmatrix} A' & B' \\ C' & D' \end{bmatrix}_{2,1} \begin{bmatrix} v_1 \\ \iota_1 \end{bmatrix}. \quad (3.5)$$

The method of cascading transmission matrices (as in equation (3.5)) is only possible if there is a single predetermined pathway for wave-propagation [101].

More complex networks of ducts with multiple pathways cannot be modelled using two-port matrices.

A *load* or *driving point* impedance (shorthand $Z_\phi = Z_{\phi,\phi}$) is defined as the ratio of potential to flow at the same point. For a simple two-port system between ports 1 and 2, if the load impedance Z_2 at location 2 is known, then the transmission matrices given in equations (3.3) and (3.4) can be used to calculate the load impedance at location 1:

$$Z_1 = \frac{AZ_2 + B}{CZ_2 + D} \quad (3.6)$$

$$= \frac{Z_2 D' - B'}{A' - Z_2 C'}. \quad (3.7)$$

3.1.2 Acoustic transmission matrices

The solution to the 1D wave equation for sound propagating in a uniform duct (derived in appendix A.1) can be expressed in terms of forwards and backwards propagating waves of pressure and velocity perturbation:

$$p'(x, t) = \hat{B}_f e^{i(\omega t - kx)} + \hat{B}_b e^{i(\omega t + kx)}, \quad (3.8)$$

$$u'(x, t) = \frac{\hat{B}_f}{\rho_0 c_0} e^{i(\omega t - kx)} - \frac{\hat{B}_b}{\rho_0 c_0} e^{i(\omega t + kx)}. \quad (3.9)$$

\hat{B}_f and \hat{B}_b are the magnitudes of the forward and backward traveling waves, e is Euler's number, i is the square-root of minus one, ω is the angular frequency, t is time, k is the wavenumber and x is the axial coordinate. Equations (A.17) and (A.18) can be combined to give the (specific) forwards transmission matrix between points 1 and 2 separated by a distance δl :

$$\begin{bmatrix} p'_2 \\ u'_2 \end{bmatrix} = \begin{bmatrix} \cosh(i k \delta l) & -Z_0^s \sinh(i k \delta l) \\ -Y_0^s \sinh(i k \delta l) & \cosh(i k \delta l) \end{bmatrix} \begin{bmatrix} p'_1 \\ u'_1 \end{bmatrix}. \quad (3.10)$$

Y_0^s is the characteristic specific acoustic admittance, defined as $1/Z_0^s = 1/\rho_0 c_0$, while \sinh and \cosh are the hyperbolic sine and cosine functions.

For a system expressed in terms of pressure and velocity perturbations, any change in area along a duct system must be accounted for by a transmission

matrix that satisfies continuity of pressure and volume velocity. For a sudden change in area from area \hat{A}_{left} to area \hat{A}_{right} this transmission matrix is

$$\begin{bmatrix} p'_{\text{right}} \\ u'_{\text{right}} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \hat{A}_{\text{left}}/\hat{A}_{\text{right}} \end{bmatrix} \begin{bmatrix} p'_{\text{left}} \\ u'_{\text{left}} \end{bmatrix}. \quad (3.11)$$

Alternatively, the system can be expressed in terms of pressure and volume velocity (q'), in which case no separate transmission matrix is needed to account for a *sudden* discontinuity in area. Consider, for example, the duct in figure 3.1. The duct has a constant area \hat{A}_1 , but at the end of the duct there is a sudden expansion to area \hat{A}_2 . In terms of pressure (p') and velocity (u') perturbations this can be expressed as

$$\begin{bmatrix} p'_2 \\ u'_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \hat{A}_1/\hat{A}_2 \end{bmatrix} \begin{bmatrix} \cosh(ik\delta l) & -Z_0^s \sinh(ik\delta l) \\ -Y_0^s \sinh(ik\delta l) & \cosh(ik\delta l) \end{bmatrix} \begin{bmatrix} p'_1 \\ u'_1 \end{bmatrix}. \quad (3.12)$$

To convert to a formulation in terms of pressure (p') and volume velocity (q') perturbations we can apply the conversions

$$\begin{bmatrix} p' \\ q' \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \hat{A} \end{bmatrix} \begin{bmatrix} p' \\ u' \end{bmatrix}, \quad (3.13)$$

and

$$\begin{bmatrix} p' \\ u' \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1/\hat{A} \end{bmatrix} \begin{bmatrix} p' \\ q' \end{bmatrix}, \quad (3.14)$$

which lead to

$$\begin{bmatrix} p'_2 \\ q'_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & \hat{A}_2 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \hat{A}_1/\hat{A}_2 \end{bmatrix} \begin{bmatrix} \cosh(ik\delta l) & -Z_0^s \sinh(ik\delta l) \\ -Y_0^s \sinh(ik\delta l) & \cosh(ik\delta l) \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & 1/\hat{A}_1 \end{bmatrix} \begin{bmatrix} p'_1 \\ q'_1 \end{bmatrix}. \quad (3.15)$$

Equation (3.15) can be simplified to

$$\begin{bmatrix} p'_2 \\ q'_2 \end{bmatrix} = \begin{bmatrix} \cosh(ik\delta l) & -Z_{0:\hat{A}_1}^s \sinh(ik\delta l) \\ -Y_{0:\hat{A}_1}^s \sinh(ik\delta l) & \cosh(ik\delta l) \end{bmatrix} \begin{bmatrix} p'_1 \\ q'_1 \end{bmatrix}, \quad (3.16)$$

in which $Y_{0:\hat{A}_1}^a$ is the characteristic acoustic admittance for a uniform duct of area \hat{A}_1 , defined as $1/Z_{0:\hat{A}_1}^a$. $Z_{0:\hat{A}_1}^a$ is the characteristic acoustic impedance of the uniform duct segment of area \hat{A}_1 , which is defined as

$$Z_{0:\hat{A}_1}^a = \frac{\rho_0 c_0}{\hat{A}_1}. \quad (3.17)$$

3.2 Lumped-element modelling of acoustical systems

Consider a short section of duct in which only plane-waves propagate. In section 3.1.2 and appendix A.1 we discuss how to model the wave propagation through this duct segment using the 1D wave-equation. Alternatively, we can take a *lumped element* approach. In the lumped element method, a linear continuous system is discretised into a finite number of degrees of freedom connected by ‘lumped’ elements. A (sufficiently short) duct segment can be simplified to just two degrees of freedom - one at each end of the duct. We now need to capture the physics of how the pressure and velocity perturbations at the two degrees of freedom are related. The duct segment can be characterised by three effects: inertia, compressibility and dissipation. The term ‘lumped’ comes from the fact that we are taking distributed phenomena and combining (‘*lumping*’) them into a small number of elements that link the degrees of freedom.

By reducing the duct segment to two degrees of freedom, we neglect any wave-propagation phenomena between the two end points. As a result of this, the lumped element model is only valid if the physical dimensions associated with each lumped element are small compared to the shortest wavelength of interest.

3.2.1 Acoustic mass

Consider a narrow duct, in which most of the acoustic energy is stored as kinetic energy, with only a negligible amount of potential energy due to compression. We can approximate the duct as incompressible, such that the volume velocity

flowing into the duct equals the volume velocity flowing out:

$$q'_{\text{in}} = q'_{\text{out}} = q'. \quad (3.18)$$

The air in a duct of length l and area \hat{A} can then be treated as a rigid lumped mass ($\rho_0 A l$) which is accelerated by the pressure forces on its ends. Newton's second law then gives

$$\hat{A}(p'_{\text{in}} - p'_{\text{out}}) = i\omega(\rho_0 l q'), \quad (3.19)$$

in which $i\omega$ is the frequency-domain equivalent of differentiation with respect to time. Combining equations (3.18) and (3.19) in transmission matrix form yields

$$\begin{bmatrix} p'_{\text{out}} \\ q'_{\text{out}} \end{bmatrix} = \begin{bmatrix} 1 & -i\omega\rho_0 l/\hat{A} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p'_{\text{in}} \\ q'_{\text{in}} \end{bmatrix}. \quad (3.20)$$

3.2.2 Compliant cavity

Consider a short segment of duct in which compressive effects are dominant and the proportion of acoustic energy that is stored as kinetic energy is negligible. We consider only the frequency range for which the wavelengths are large compared to the length of the duct segment. This means that the acoustic pressure perturbations are approximately uniform throughout the segment. We can thus treat the short duct segment as a uniform pressure cavity with constant volume (V) and negligible kinetic energy. If the pressure is uniform in the segment then

$$p'_{\text{in}} = p'_{\text{out}} = p'. \quad (3.21)$$

Mass conservation dictates that the net mass flow into the volume must be matched by density changes within the volume, which yields

$$\rho_0(q'_{\text{in}} - q'_{\text{out}}) = V \frac{\partial \rho'}{\partial t}, \quad (3.22)$$

in which ρ' is a density perturbation and ∂ is the partial differential operator. We now assume that the density fluctuations in the cavity are isentropic (since the fluctuations are reversible, and occur so fast that there is no heat transfer to

the surroundings (adiabatic)). For an isentropic process

$$p' = c_0^2 \rho', \quad (3.23)$$

such that

$$p' = \frac{c_0^2 \rho_0 (q'_{\text{in}} - q'_{\text{out}})}{i\omega V}. \quad (3.24)$$

Combining equations (3.21) and (3.24) in transmission matrix form yields

$$\begin{bmatrix} p'_{\text{out}} \\ q'_{\text{out}} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \frac{-i\omega V}{\rho_0 c_0^2} & 1 \end{bmatrix} \begin{bmatrix} p'_{\text{in}} \\ q'_{\text{in}} \end{bmatrix}. \quad (3.25)$$

3.2.3 LEM for a Helmholtz resonator

Consider a Helmholtz resonator consisting of two segments of circular cross-section duct. Figure 3.1 shows a schematic of the system. The first segment of duct is long and narrow, while the second is short and wide. The maximum spatial dimension is the length of the long segment of duct $l_{\text{max}} = x_{2,1}$. A pressure perturbation p'_1 is applied to one end of the long duct. In this example, we consider the response of the pressure at the far end of the short duct (p'_3) to the input pressure (p'_1).

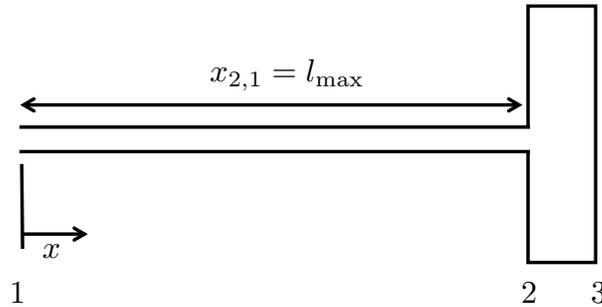


Fig. 3.1 Schematic of a simple Helmholtz resonator model with maximum spatial dimension $l_{\text{max}} = x_{2,1}$.

The long, narrow segment (1 \rightarrow 2) can be treated as an acoustic mass, while the short, wide segment (2 \rightarrow 3) can be created as an incompressible compliant cavity. Cascading the matrices in equations (3.20) and (3.25) yields

$$\begin{bmatrix} p'_3 \\ q'_3 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \frac{-i\omega V}{\rho_0 c_0^2} & 1 \end{bmatrix} \begin{bmatrix} 1 & -i\omega \rho_0 l_{\text{max}} / \hat{A} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} p'_1 \\ q'_1 \end{bmatrix}. \quad (3.26)$$

Equation (3.26) is solved by setting $q'_3 = 0$ and dividing through by p'_1 to obtain two equations for two unknowns. Figure 3.2 shows solutions for p'_3/p'_1 for the Helmholtz resonator in figure 3.1 computed by solving both the 1D wave equation (solid blue line) and the LEM (dotted grey line). The pressure ratio is shown as a function of the non-dimensional measure of frequency l_{\max}/λ , where λ is the wavelength. It is clear from the simulations that a lumped element model is valid provided the maximum spatial dimension in the problem is around ten times smaller than the longest wavelength of interest. It is possible to capture higher modes using a lumped element model by choosing an appropriate discretisation. It is ultimately the length-scale associated with the discretisation that is important.

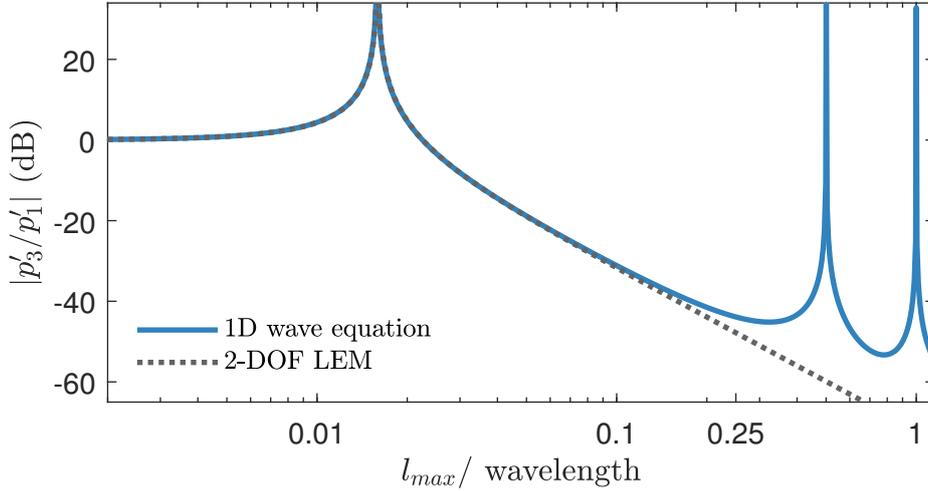


Fig. 3.2 Response of the pressure at p'_3 to the pressure at p'_1 as a function of the ratio of maximum spatial dimension (l_{\max}) to wavelength (λ) computed using both the 1D wave equation (solid line), and a 2DOF LEM (dotted line). (See appendix A.2 for details of the 1D wave equation solution.)

3.2.4 Application to sensors at the human chest

The primary relevant frequencies for auscultation are usually stated to extend from 20-60 Hz to around 600 Hz [82, 84]. Most of the energy in heart sounds can be found below 100 Hz [93], while lung sounds are typically dominant from 100 to 1000 Hz [93, 145] (-1500 Hz [163]). Some abnormal wheezes can extend up to 5 kHz [145], but in general stethoscope performance above around 3 kHz is considered unimportant.

Depending on the medium of interest, this frequency range corresponds to different wavelengths. The speed of sound in air at room temperature is around 340 m/s [28] while the speed of sound for compression waves in human body tissue is typically between 1500 and 1600 m/s [154]. At 3 kHz the wavelength of sound in air is around 11 cm, while in human soft tissue it is around 50 cm. The physical dimensions of a sensor (excluding the tubing in the case of an acoustic stethoscope) rarely exceed a few centimetres so it is reasonable to model the system using a lumped element discretisation [163].

3.3 Lumped-element modelling of mechanical systems

The lumped element method can be applied to mechanical systems by discretising the systems into a finite number of degrees of freedom which are linked by various types of lumped-element such as springs, dashpots and inerters. The combination of lumped elements chosen to model the properties of a material is referred to as a *constitutive model* [45]. The lumped elements exert forces on the degrees of freedom and, by considering the force balance at each degree of freedom, a set of equations can be written for the system.

Figure 3.3 introduces the symbols (k , λ and b) used for stiffness (springs), viscous damping (dashpots) and inertance (inerters) throughout this thesis. Dot notation is used to denote differentiation with respect to time. The figure also shows the forces exerted by each element on the degrees of freedom to which it is connected. Assuming the motion of each degree of freedom can be expressed in harmonic form as $x_i = X_i e^{i\omega t}$, the force balances can be expressed in matrix form as

$$[\mathbf{A}]\mathbf{x} = \mathbf{f}, \quad (3.27)$$

where $[\mathbf{A}]$ is a matrix containing all the information about how the degrees of freedom are linked by the lumped elements, \mathbf{x} is a vector of the displacements and \mathbf{f} is a vector of the external forces applied to each degree of freedom.

For an undamped system $\lambda = 0$ and k is purely real. In practical structures damping occurs through a variety of complicated mechanisms which typically cannot be accurately modelled in detail [31]. There are two simple models for

damping which are commonly used in vibration analysis. *Viscous damping* gives a dissipative force (f_λ) proportional to velocity and is modelled by dashpots $f_\lambda = \lambda \dot{x}$. *Structural damping* gives a dissipative force proportional to displacement and is modelled by complex stiffness $k \rightarrow k(1 + i\eta)$, where η is the ‘loss factor’.

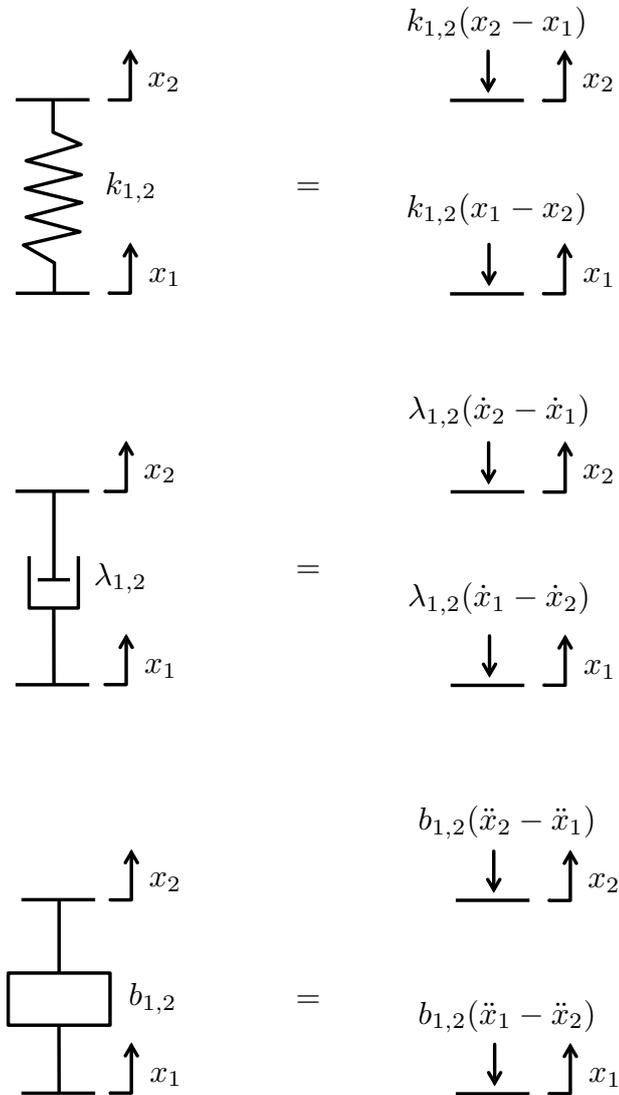


Fig. 3.3 Standard mechanical lumped elements connecting two degrees of freedom (x_1 and x_2). Stiffness ($k_{1,2}$), viscous damping ($\lambda_{1,2}$) and inertance ($b_{1,2}$). The left-hand side of the figure shows the symbols used for each element type. On the right-hand side the elements are replaced by the forces that they exert on the degrees of freedom to which they are connected.

Mass (m) is a special case of an inerter with one terminal connected to ground. This reflects the fact that it is the absolute acceleration of a mass that gives a force according to Newton’s second law of motion.

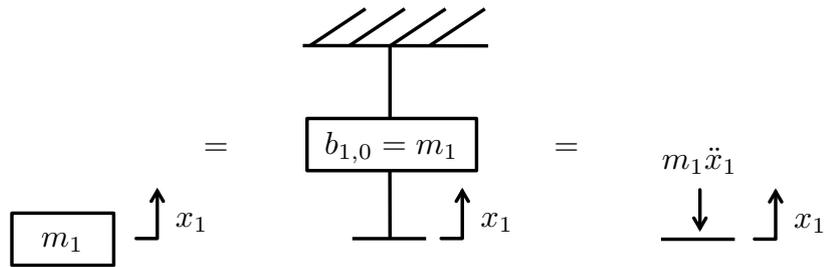


Fig. 3.4 Lumped element model for (mechanical) mass. Mass at a degree of freedom is equivalent to an inertance between that degree of freedom and (mechanical) ground.

A variety of constitutive models have been proposed to model the behaviour of viscoelastic materials. The simplest constitutive models are the Maxwell model (spring and dashpot in series), and the Kelvin-Voigt model (spring and dashpot in parallel) [45]. Both of these models struggle to capture the key properties of viscoelastic materials, so a succession of more complex models have been proposed [45]. The Standard Linear Solid Model and the Generalised Maxwell Model make use of more spring and dashpot elements to improve the match to physical behaviour [45]. More recently ‘springpots’ (fractional-order lumped elements) have been employed [18] to deal with the limitations associated with springs and dashpots and reduce the number of separate elements needed between two degrees of freedom.

3.4 Frequency response functions

In order to characterise the dynamic response of an oscillating mechanical system we consider the harmonic response of a kinematic quantity (such a displacement, velocity or acceleration) to a harmonic excitation force. The ratio of the response to the excitation as a function of frequency is called a *frequency response function* (FRF). The FRF definitions corresponding to acceleration, velocity and displacement are given in table 3.2.

Table 3.2 Terminology for frequency response functions.

Term	Symbol	Definition
Accelerance	$H_a(\phi, \psi, \omega)$	$\frac{-\omega^2 X_\phi(\omega)}{F_\psi(\omega)}$
Admittance	$H_v(\phi, \psi, \omega)$	$\frac{i\omega X_\phi(\omega)}{F_\psi(\omega)}$
Compliance	$H_x(\phi, \psi, \omega)$	$\frac{X_\phi(\omega)}{F_\psi(\omega)}$

In table 3.2, $H_a(\phi, \psi, \omega)$ refers to the accelerance evaluated for observation at the ϕ th coordinate, when subject to excitation at the ψ th coordinate, as a function of frequency (ω). $F_\psi(\omega)$ refers to the force applied at the ψ th coordinate, while $X_\phi(\omega)$ refers to the displacement at the ϕ th coordinate. By taking the ratio of force to response we can also define three *inverse* FRFs, which are shown in table 3.3.

Table 3.3 Terminology for inverse frequency response functions.

Term	Symbol	Definition
Dynamic mass	$M_d(\psi, \phi, \omega)$	$\frac{F_\psi(\omega)}{-\omega^2 X_\phi(\omega)}$
Mechanical impedance	$Z^m(\psi, \phi, \omega)$	$\frac{F_\psi(\omega)}{i\omega X_\phi(\omega)}$
Dynamic stiffness	$K_d(\psi, \phi, \omega)$	$\frac{F_\psi(\omega)}{X_\phi(\omega)}$

Tables 3.2 and 3.3 define FRFs and inverse FRFs in the mechanical domain. The same functions can be defined in the *specific* or *acoustical* domains by applying the appropriate area scaling. The nomenclature introduced in these tables is used as stated for the mechanical domain, while superscripts s and a are added for the specific and acoustical domains respectively. Thus, for example, the acoustical compliance FRF is denoted by H_x^a . The exception is impedance, for which there is no subscript, and the domain is noted by a superscript in all cases (Z^m, Z^s, Z^a).

FRFs and inverse FRFs both have their uses in modal analysis. FRFs enable simple modelling of systems as a summation of modal responses (see section

5.2.1), while *driving point* inverse FRFs can be linked directly to the spatial properties (mass, stiffness and viscous damping) of the system (see section 3.4).

3.4.1 Single degree of freedom damped oscillator

Consider a single degree of freedom system described by the displacement x_1 of the degree of freedom, as shown in figure 3.5. A force balance at x_1 yields $F_1 = (k_{1,0} + i\omega\lambda_{1,0} - \omega^2 m_1)X_1$. The (driving point) inverse FRFs can therefore be evaluated as:

$$M_d(1, 1, \omega) = -\frac{k_{1,0}}{\omega^2} + \frac{\lambda_{1,0}}{i\omega} + m_1 \quad (3.28)$$

$$Z^m(1, 1, \omega) = \frac{k_{1,0}}{i\omega} + \lambda_{1,0} + i\omega m_1 \quad (3.29)$$

$$K_d(1, 1, \omega) = k_{1,0} + i\omega\lambda_{1,0} - \omega^2 m_1 \quad (3.30)$$

In equations 3.28 to 3.30, $\lambda_{1,0}$ is non-zero for viscous damping and $k_{1,0}$ is complex for structural damping.

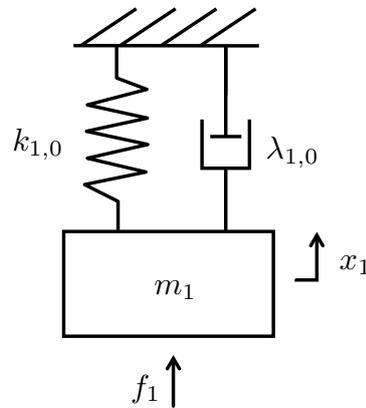


Fig. 3.5 Schematic of single degree of freedom damped oscillator.

Figure 3.6 shows the inverse FRFs as functions of (normalised) frequency for the system in figure 3.5. At low frequencies, the spring dominates the response of the system. In the dynamic mass plot this is characterised by a region with a slope of -40 dB/decade. In the mechanical impedance plot, multiplication by $i\omega$ turns this into a slope of -20 dB/decade, and in the dynamic stiffness plot this region is characterised by a constant value (equal to the stiffness of the spring).

At $\omega = \sqrt{k_{1,0}/m_1}$ the system resonates (resulting in a trough in the inverse FRFs), and we see a shift from a stiffness dominated response to a mass dominated response. At the resonance the phase of the impedance is 0 degrees, as the imaginary terms from the mass and stiffness cancel. The magnitude of the mechanical impedance at resonance therefore depends only on the damping factor.

At high frequencies the system is mass dominated. The phase of the dynamic mass tends to zero and the magnitude tends to a constant value equal to the mass (in dB) at the driving point x_1 .

The magnitude, slopes and phase angles that we see at frequencies where the system is dominated by a certain component generalise to any system. Furthermore, for any system, a trough in an inverse FRF relates to a shift from a stiffness dominated system to a mass dominated system.

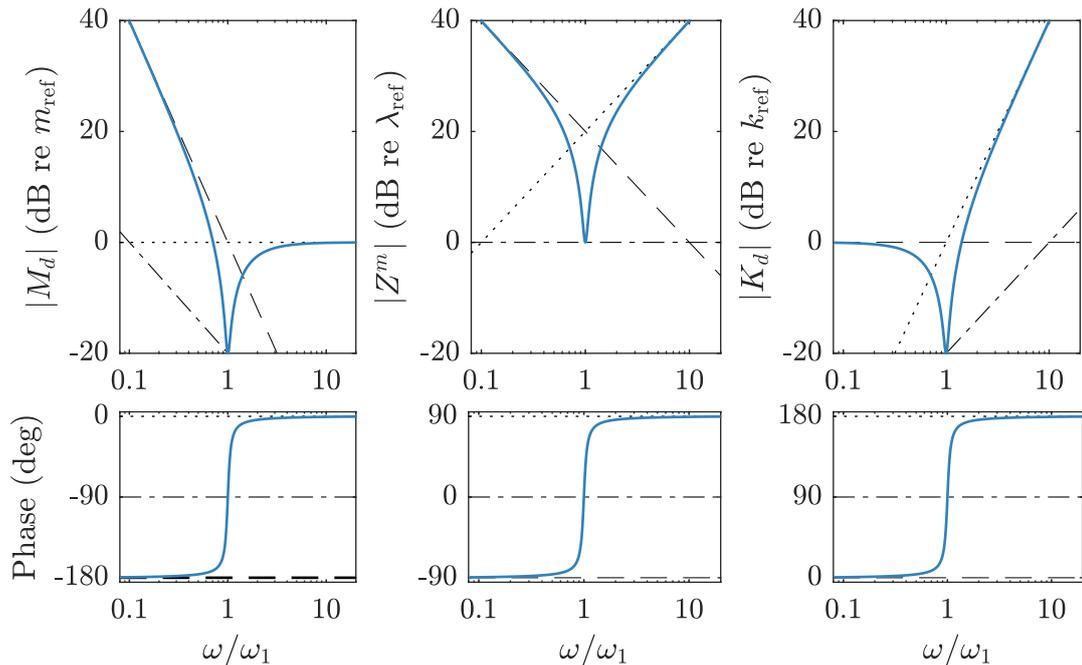


Fig. 3.6 Bode plot of the driving point dynamic mass, mechanical impedance and dynamic stiffness for the system in figure 3.5. The decibel reference values are the values of m_1 , $\lambda_{1,0}$ and $k_{1,0}$ used in the simulation. These are $m_{\text{ref}} = 1 \text{ kg}$, $\lambda_{\text{ref}} = 10 \text{ kg s}^{-1}$ and $k_{\text{ref}} = 10 \text{ kN m}^{-1}$. The grey dotted line represents a pure mass, the grey dot-dash line represents pure viscous damping and the grey dashed line represents pure stiffness.

We now illustrate some dynamic phenomena and show the signatures these leave on the driving point inverse frequency response function plots. These phenomena are evident in many of the experimentally measured FRFs in this thesis and are crucial in the coupling of sensors to the human chest.

3.4.2 ‘Mass locking’

Consider a simple two-degree-of-freedom system, in which the ground link in figure 3.5 has been replaced by a second mass m_2 that is free to move. As there is no ground connection, the first resonance is at $\omega_1 = 0$ Hz and corresponds to rigid body motion of the entire system. The system is driven by a force f_1 which is applied to m_1 . Allowing for viscous damping by including a dashpot ($\lambda_{1,2}$), the governing equations can be expressed in matrix form as

$$\begin{bmatrix} m_1 & 0 \\ 0 & m_2 \end{bmatrix} \begin{bmatrix} \ddot{x}_1 \\ \ddot{x}_2 \end{bmatrix} + \begin{bmatrix} \lambda_{1,2} & -\lambda_{1,2} \\ -\lambda_{1,2} & \lambda_{1,2} \end{bmatrix} \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} + \begin{bmatrix} k_{1,2} & -k_{1,2} \\ -k_{1,2} & k_{1,2} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} f_1 \\ 0 \end{bmatrix}. \quad (3.31)$$

At low frequency the response is dominated by the net mass ($m_1 + m_2$). As the frequency increases the response then moves through a peak in dynamic mass as m_2 resonates with $k_{1,2}$. The undamped antiresonance frequency is

$$\omega_{-1} = \sqrt{\frac{k_{1,2}}{m_2}}. \quad (3.32)$$

Above ω_{-1} , m_2 effectively stops moving and becomes ‘locked’. This means that the mass m_2 has a negligible effect on the high frequency driving point response of the system. Above ω_{-1} , the response is dominated by the stiffness ($k_{1,2}$) between the driving point and the locked mass m_2 . The peak in dynamic mass at ω_{-1} is followed by a trough at the second natural frequency of the combined system (ω_2). For an undamped system this frequency is given by:

$$\omega_2 = \sqrt{\frac{k_{1,2}(m_1 + m_2)}{(m_1 m_2)}}. \quad (3.33)$$

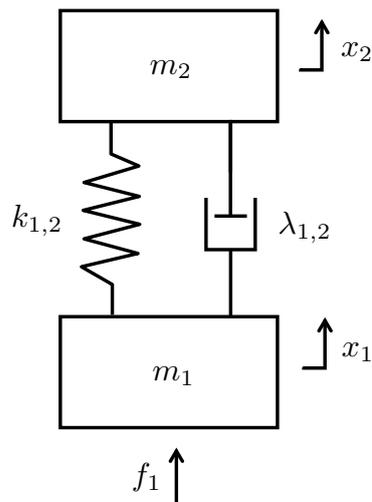


Fig. 3.7 Two degree of freedom lumped element mechanical system to demonstrate ‘mass locking’.

Beyond this resonance, the base mass dominates over the stiffness. On the dynamic mass plot (figure 3.8) the effect of these two resonances is that we move from a low frequency plateau at the combined mass of the system to a high frequency plateau related to the base mass only. Between these two mass-dominated plateaus there is an intermediate region in which stiffness effects dominate.

The distance between the two peaks is determined by the relative magnitude of the two masses. If $m_2 \ll m_1$ the gap tends to zero, while if $m_1 \ll m_2$ the gap tends to infinity. This can be verified by inspection of equations 3.32 and 3.33.

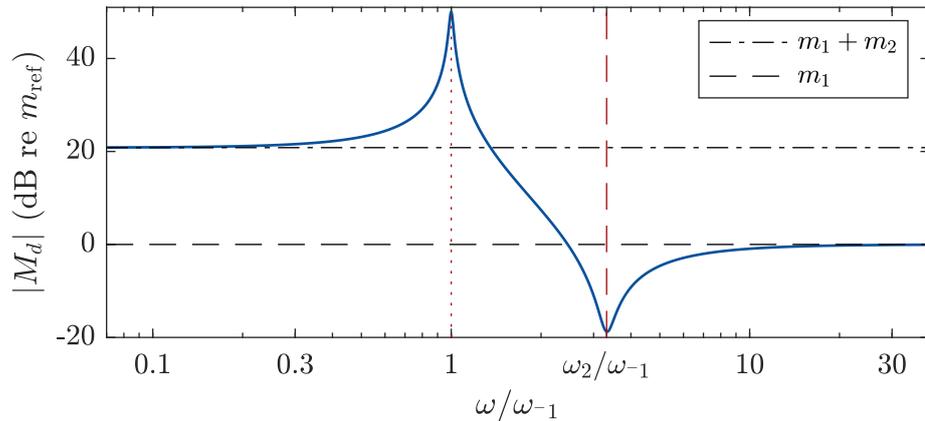


Fig. 3.8 Driving-point dynamic mass for the system in figure 3.7. Decibel reference values are the values of m , λ and k used in the simulation. These are $m_{\text{ref}} = m_1 = 0.01 \text{ kg}$, $\lambda_{\text{ref}} = \lambda_{1,2} = 1 \text{ kg s}^{-1}$, $k_{\text{ref}} = k_{1,2} = 10 \text{ kN m}^{-1}$. In this simulation $m_2 = 10m_1$. The net system mass ($m_1 + m_2$) is given by a horizontal black dot-dashed line. The driving point mass (m_1) is given by a horizontal black dashed line. The frequency axis is normalised by ω_{-1} (see equation 3.32). The vertical dotted red line corresponds to the first antiresonance, while the vertical dashed red line corresponds to the second resonance (see equation 3.33).

3.4.3 ‘Stiffness Shedding’

Mass-locking gives rise to another phenomenon which we refer to as ‘stiffness-shedding’. If a mass in the system ‘locks’, then any stiffnesses ‘beyond’ that mass (as seen from the driving point) no longer affects the driving point response of the system. In figure 3.9 the two-degree-of-freedom system in figure 3.7 is modified by adding a spring and a dashpot between x_2 and mechanical ground.

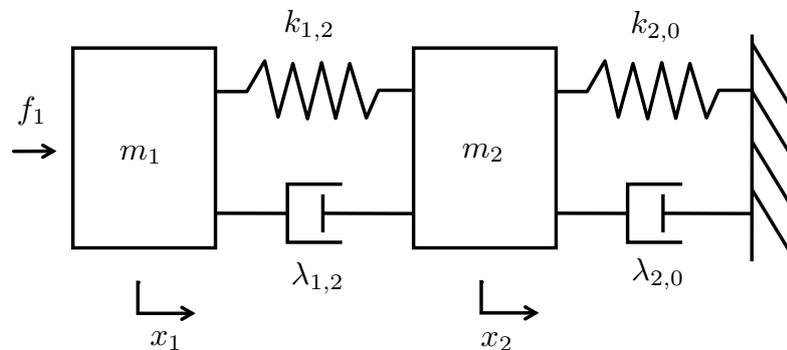


Fig. 3.9 Two degree of freedom lumped element mechanical system to demonstrate ‘stiffness shedding’.

The driving point dynamic stiffness FRF for the system in figure 3.9 is shown in figure 3.10. There is a low frequency plateau in dynamic stiffness with a magnitude that matches the net stiffness of the system $k_{\text{net}} = k_{1,2} || k_{2,0}$. The symbol $||$ denotes the parallel combination of the two terms such that

$$k_{1,2} || k_{2,0} = \frac{k_{1,2} k_{2,0}}{k_{1,2} + k_{2,0}}. \quad (3.34)$$

At intermediate frequencies ($\omega_{-1} < \omega < \omega_2$) there is a second plateau in dynamic stiffness, corresponding only to the stiffness ‘below’ the ‘locked’ mass ($k_{1,2}$). For the case shown, the plateau is cut off by the next resonance of the system (ω_2).

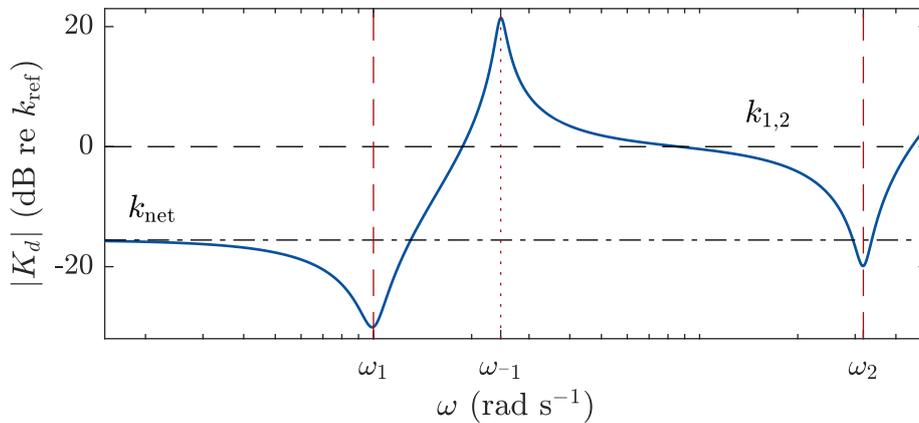


Fig. 3.10 Driving-point dynamic stiffness for the system in figure 3.9, using the parameter values given in appendix D.2. The decibel reference value is $k_{\text{ref}} = k_{1,2}$.

3.5 Electrical analogy

In section 3.1 we discuss the way in which acoustical, mechanical and electrical systems can all be described in the same manner using the concepts of potential, flow and impedance. This gives rise to an analogy: any system for which the governing equations can be expressed in terms of *flow*, *potential* and *impedance* can be converted into an equivalent electrical circuit. There are two forms of the electrical-analogy. In the *impedance analogy* impedances in the original physical domain become electrical impedances, while in the *admittance analogy* it is admittances in the original physical domain that become electrical impedances. This means that in the *impedance analogy* electrical flow and potential represent

flow and potential in the original system, while for the *admittance analogy* potential and flow are switched.

Table 3.4 Equivalence of terms in the electrical impedance analogy. The superscripts e , m and a denote terms in the electrical, mechanical and acoustic domains.

Electrical	Mechanical	Acoustic
v^e	f	p'
ι^e	\dot{x}	q'
z^e	z^m	z^a
L^e	m	$m^a = \rho_0 l / \hat{A}$
R^e	λ	R^a
C^e	$1/k$	$1/k^a * = V / \rho_0 c_0^2$

* The capacitance corresponding to an acoustic stiffness (k^a) must always have one terminal grounded.

Table 3.5 Equivalence of terms in the electrical admittance analogy.

Electrical	Mechanical	Acoustic
v^e	\dot{x}	q'
ι^e	f	p'
z^e	y^m	y^a
L^e	$1/k$	$1/k^a = V / \rho_0 c_0^2$
R^e	$1/\lambda$	$1/R^a$
C^e	b or m^*	$m^a = \rho_0 l / \hat{A}$

* The capacitance corresponding to mass (m) must always have one terminal connected to ground, the true admittance-analogue of capacitance is inertance (b).

3.5.1 A note on notation

There is a fundamental difference between frequency response functions and lumped parameter values despite the fact that they have the same physical dimensions. For clarity, let us emphasize that $M_{d:\phi,\psi}(\omega)$, $Z_{\phi,\psi}^m(\omega)$ and $K_{d:\phi,\psi}(\omega)$ are the (inverse) FRFs between an excitation force at ϕ and acceleration, velocity or displacement at ψ . On the other hand, $b_{\phi,\psi}$, $\lambda_{\phi,\psi}$ and $k_{\phi,\psi}$ are the lumped

inertance, viscous damping and stiffness between points ϕ and ψ . The lumped mass m_ϕ is equivalent to a lumped inertance between ϕ and ground ($m_\phi = b_{\phi,0}$).

We also introduce a general lumped impedance $z_{\phi,\psi}$, which is defined in the same way as $\lambda_{\phi,\psi}$, but is permitted to be complex and frequency dependant, such that the real part can capture any damping terms, while the imaginary part captures mass and stiffness terms.

3.5.2 Mechanical systems

When dealing with mechanical systems, the admittance analogy is the most intuitive as it preserves the topology of the system (i.e. the way in which the elements are connected), in contrast to the impedance analogy which does not. Circuit-manipulation techniques can be used to convert between the analogies as shown by the example in this section (see also [6]).

Consider a simple point mass under the action of two forces f_a and f_b , as shown in figure 3.11. By rearranging Newton's second law to

$$\sum \mathbf{f} - m\ddot{\mathbf{x}} = 0, \quad (3.35)$$

the acceleration of a lumped mass can be modelled as a d'Alembert (inertia) force ($-m\ddot{\mathbf{x}}$).

Now consider the admittance analogy for the system. The d'Alembert force is always related to the *absolute* acceleration of the mass, so the electrical component representing the lumped mass is a capacitance between the nodes corresponding to x_1 and ground. Mass is therefore not the true analogue of capacitance. This role is attributed to a fairly recently defined mechanical device called an 'inertor' (see [138, 139]). An 'inertor' gives a force proportional to the relative acceleration between its two ends, and can thus be represented in the admittance analogy as a capacitance between any two nodes.

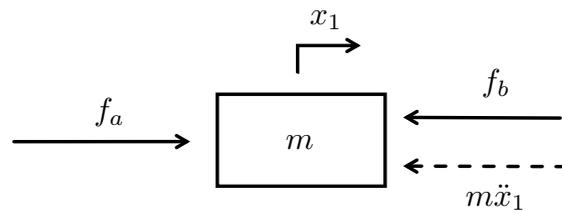


Fig. 3.11 Schematic of point mass under the action of forces. The acceleration of the point mass has been expressed as a d'Alembert (inertia) force.

Figure 3.12 shows the admittance analogy circuit for the point mass in figure 3.11. The sum of currents into node η_1 represents the force balance at coordinate x_1 . The current flowing through the capacitor ($\iota = m\ddot{x}_1$) represents the d'Alembert force on the mass.

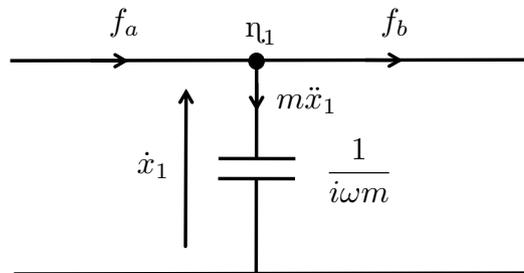


Fig. 3.12 Admittance analogy circuit for the point mass in figure 3.11.

As a more complete example of the method, figure 3.13 gives the admittance analogy for the lumped element mechanical system in figure 3.9. The topology of the system has been conserved, and the impedances in the electrical circuit are given by the mechanical admittances of the lumped elements.

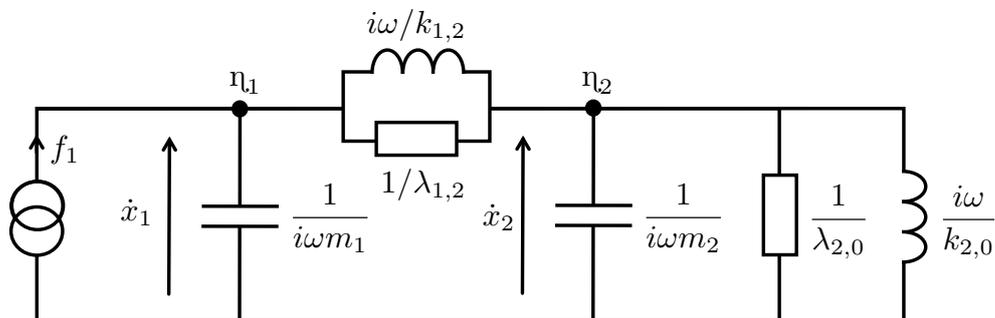


Fig. 3.13 Admittance analogy circuit for the two-degree-of-freedom lumped element system in figure 3.9.

An admittance analogy circuit can be converted into an impedance analogy circuit by taking the ‘dual’ of the circuit. The method for doing this is described in appendix A.3 (see also (Beranek, 2012) [6]). When performing this conversion, *potentials* and *flows* switch and impedances are replaced by their inverse. The topology of the circuit changes, such that the sum of flows into a node becomes the sum of potentials in a loop. For example, the forces modelled as currents ‘flowing’ into node η_1 in figure 3.13 become the voltages across the components in loop₁ in figure 3.14. In general, we see that components in parallel in the admittance analogy circuit are in series in the impedance analogy circuit and vice versa.

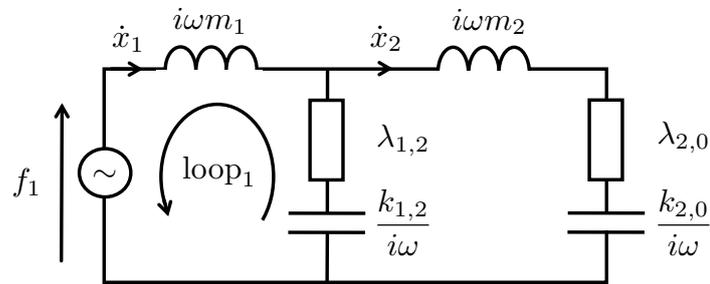


Fig. 3.14 Impedance analogy circuit for the two-degree-of-freedom lumped element system in figure 3.9.

3.5.3 Acoustic systems

For acoustic systems, it is the impedance analogy that preserves the topology of the system (see e.g. [6]). For example, the Helmholtz resonator in figure 3.1 can be converted into the impedance analogy circuit in figure 3.15, with the topology of the components conserved (the two lumped elements remain in series).

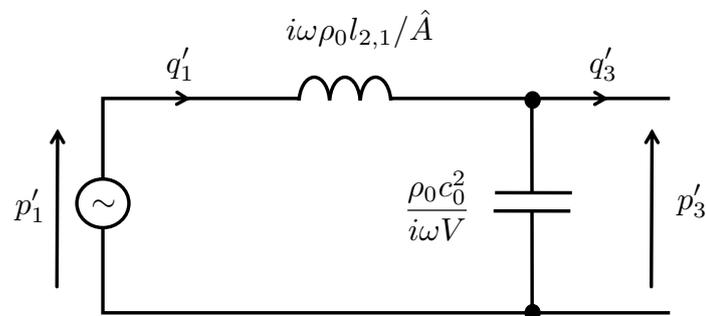


Fig. 3.15 Impedance analogy circuit for the simple Helmholtz resonator in figure 3.1.

3.5.4 Two port formalism and transmission matrices

The theory of two-ports is outlined in section 3.1. Taking v_ϕ and ι_ϕ to represent *potential* and *flow* at location ϕ , the forwards transmission matrix for a general lumped series impedance (z_{series}) can be expressed as

$$\begin{bmatrix} v_2 \\ \iota_2 \end{bmatrix} = \begin{bmatrix} 1 & -z_{\text{series}} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v_1 \\ \iota_1 \end{bmatrix}, \quad (3.36)$$

while the transmission matrix for a general lumped parallel impedance (z_{parallel}) is given by

$$\begin{bmatrix} v_2 \\ \iota_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -1/z_{\text{parallel}} & 1 \end{bmatrix} \begin{bmatrix} v_1 \\ \iota_1 \end{bmatrix}. \quad (3.37)$$

It is important to make a distinction between the electrical analogy and the two-port method. The two-port method is a mathematical technique that can only be used to provide a solution for the response of systems with a single predetermined pathway for wave propagation. The electrical analogies, on the other hand, are fundamental analogies for the governing equations and are valid for any lumped system.

In order to convert an electrical circuit into the required two-port formulation it may be necessary to rearrange the circuit *without* altering the topology of the components. For example, if we are interested in the force in spring $k_{1,2}$ in the two-degree of freedom example from figure 3.9, then the impedance analogy circuit in figure 3.14 needs to be rearranged into the circuit shown in figure 3.16.

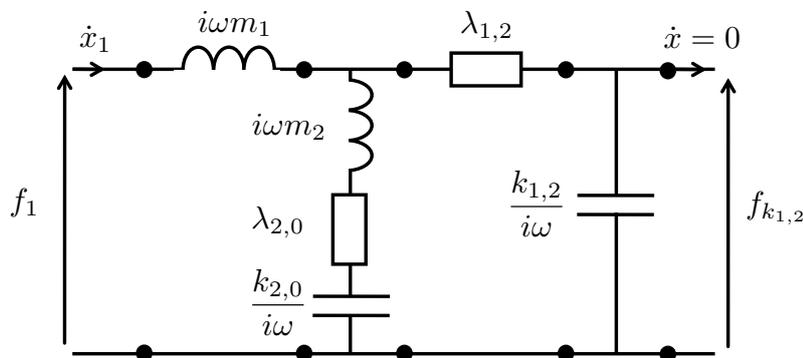


Fig. 3.16 Rearrangement of circuit in figure 3.14 into a two-port formulation for the force in spring $k_{1,2}$. The dots represent the boundaries between adjacent two-ports.

By taking the two-port matrix for each of the four two ports in figure 3.16, the circuit can be expressed as a cascaded set of multiplied transmission matrices:

$$\begin{bmatrix} f_{k_{1,2}} \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -i\omega/k_{1,2} & 1 \end{bmatrix} \begin{bmatrix} 1 & -\lambda_{1,2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/z_{2,0} & 1 \end{bmatrix} \begin{bmatrix} 1 & -i\omega m_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} f_1 \\ \dot{x}_1 \end{bmatrix}, \quad (3.38)$$

where

$$z_{2,0} = i\omega m_2 + \lambda_{2,0} + k_{2,0}/i\omega. \quad (3.39)$$

3.5.5 Boundary conditions

The boundary conditions in a physical system are represented by load impedances in equivalent circuits. For the impedance analogy, an unconstrained node in a mechanical system and an open duct (with no radiation impedance) in an acoustic system are both represented by a short-circuit (zero impedance), as there is no output *potential*. A ground-link in a mechanical system and a rigid wall in an acoustic system are both represented by an open-circuit (infinite impedance) in the impedance analogy equivalent circuit as there is no output *flow*. For the admittance analogy the opposite is true: open ducts and unconstrained ends are represented by open-circuits, while ground-links and rigid walls become short-circuits.

3.5.6 Coupling

The advantage of using electrical analogies to model physical systems lies in the ease of coupling different types of system within a single mathematical framework.

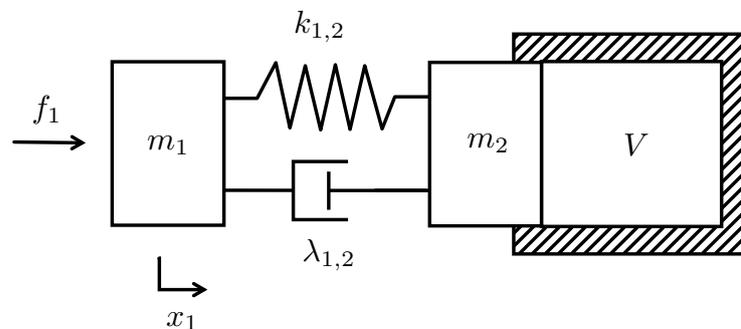


Fig. 3.17 Schematic of a coupled mechano-acoustic system.

We now consider the coupling of a mechanical system to an acoustic system. For this example, we modify the system in figure 3.9, replacing the spring and dashpot to ground with a small volume of constrained air upon which the mass m_2 acts as a piston. This setup is shown schematically in figure 3.17. Neglecting dissipation in the constrained volume of air, and assuming that the dimensions are appropriate for a lumped element model, we now have two coupled systems:

$$\begin{bmatrix} f_2 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 1 & -i\omega m_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/z_{1,2} & 1 \end{bmatrix} \begin{bmatrix} 1 & -i\omega m_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} f_1 \\ \dot{x}_1 \end{bmatrix}, \quad (3.40)$$

and

$$\begin{bmatrix} p'_2 \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -i\omega V/\rho_0 c_0^2 & 1 \end{bmatrix} \begin{bmatrix} p'_2 \\ q'_2 \end{bmatrix}, \quad (3.41)$$

where

$$z_{1,2} = \lambda_{1,2} + k_{1,2}/i\omega. \quad (3.42)$$

For convenience we represent both the mechanical and acoustic domains using electrical-impedance-analogy circuits. The uncoupled impedance analogy circuits of the two sub-systems are shown in figure 3.18.

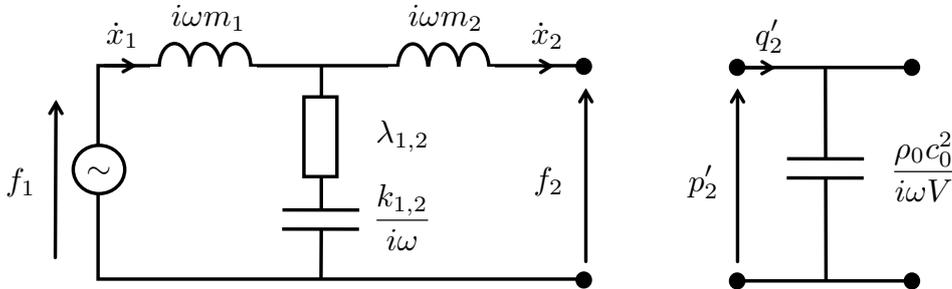


Fig. 3.18 Uncoupled impedance analogy circuits for the mechanical and acoustical sub-systems in figure 3.17.

The systems are coupled by considering the relationships between f_2 and p'_2 and between \dot{x}_2 and q'_2 . In both cases this depends on the surface area of the piston (\hat{A}):

$$p'_2 = f_2/\hat{A}, \quad (3.43)$$

$$q'_2 = \dot{x}_2\hat{A}, \quad (3.44)$$

$$\frac{f_2}{\dot{x}_2} = \hat{A}^2 \frac{p'_2}{q'_2}. \quad (3.45)$$

These relationships can be represented in the impedance analogy by a transformer with a ratio of $\hat{A} : 1$. As shown in figure 3.19, the two separate physical systems are now included within a single circuit.

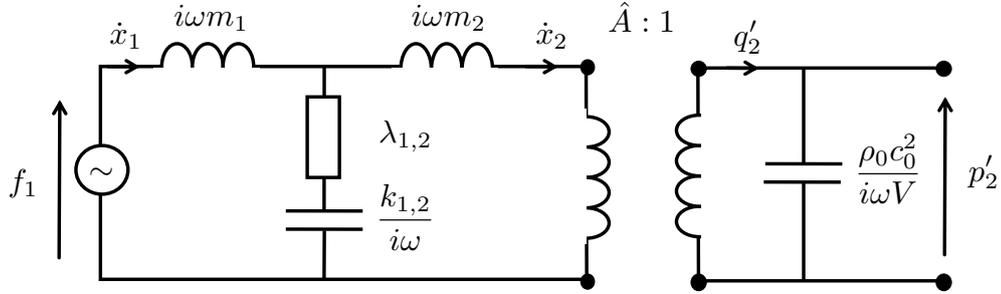


Fig. 3.19 Coupled impedance analogy circuit for system in figure 3.17. A transformer with a ratio of $\hat{A} : 1$ is used to model the coupling between the mechanical and acoustical systems.

To simplify the coupled circuit we can transfer all terms to one side of the transformer. To transfer impedance terms across a transformer they are multiplied by the square of the transformer ratio. For example, to transfer the acoustical stiffness term to the mechanical side it must be multiplied by $(\hat{A}/1)^2$. The resulting circuit is shown in figure 3.20.

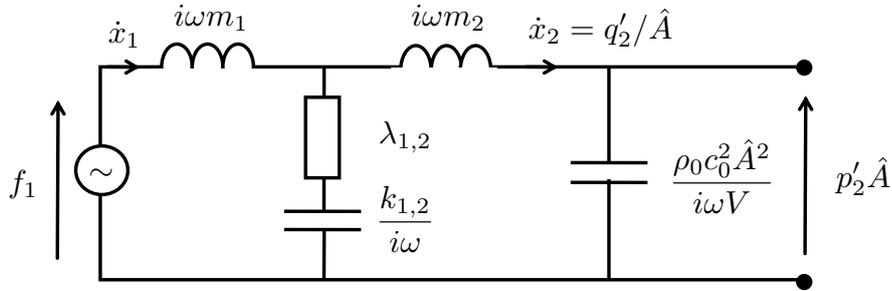


Fig. 3.20 Coupled impedance analogy circuit for system in figure 3.17, with acoustical terms converted to mechanical terms by transferring them across the transformer in figure 3.19.

In this case there is a single, predetermined pathway for wave propagation, and the system can be expressed in two-port form as

$$\begin{bmatrix} p_2' \hat{A} \\ 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \frac{-i\omega V}{\rho_0 c_0^2 \hat{A}^2} & 1 \end{bmatrix} \begin{bmatrix} 1 & -i\omega m_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/z_{1,2} & 1 \end{bmatrix} \begin{bmatrix} 1 & -i\omega m_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} f_1 \\ x_1 \end{bmatrix}. \quad (3.46)$$

When transferred to the mechanical side the stiffness of an air cavity is given by equation (3.47). For systems which are predominantly mechanical but include air cavities, this stiffness can be used directly in the mechanical model without formally working through the mathematics in this section:

$$k_a = \frac{\rho_0 c_0^2 \hat{A}^2}{V}. \quad (3.47)$$

3.5.7 Norton and Thévenin's theorems

Norton and Thévenin's theorems allow us to represent large sections of a linear circuit in terms of a single impedance term and a single source term [6]:

Norton's theorem:

Any linear circuit consisting of both energy sources and impedances can be replaced by a single current source in parallel with a single impedance.

Thévenin's theorem:

Any linear circuit consisting of both energy sources and impedances can be replaced by a single voltage source in series with a single impedance.

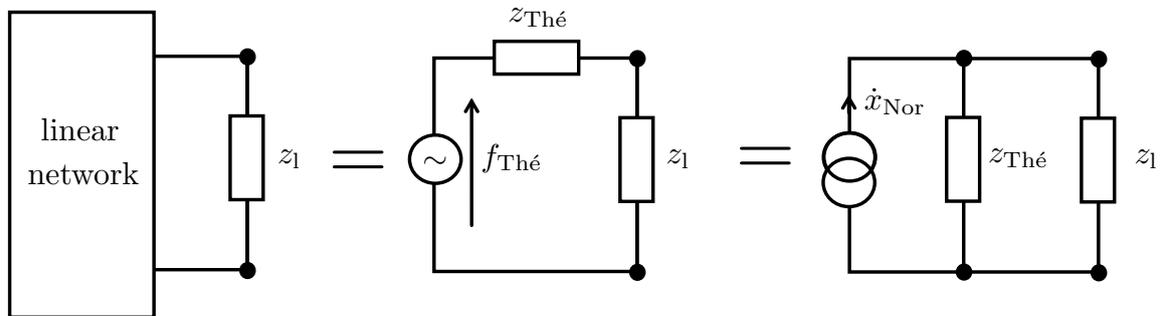


Fig. 3.21 Circuit representation of Thévenin and Norton's theorems for a mechanical impedance analogy circuit. Any linear network (left) can be replaced by a force source in series with an impedance (centre) or a velocity source in parallel with an impedance (right).

The 'Norton current' is the current that flows at the load terminals when these are connected by a short-circuit (zero impedance). The 'Thévenin voltage' is the voltage across the load terminals when the load is replaced by an open-circuit (infinite impedance). The 'Norton impedance' is equivalent to the 'Thévenin impedance' ($z_{\text{Thé}}$), and is the impedance that would be 'seen' at the load terminals

when all voltage sources are shorted and all current sources are replaced by open circuits.

Due to the analogy that exists between linear circuits and lumped mechanical/acoustical systems, Norton and Thévenin's theorems must apply to these systems too. For example, for a mechanical system the Norton velocity (\dot{x}_{Nor}) is the velocity at the output of the source system when the source is not connected to any load. The Thévenin force ($f_{\text{Thé}}$) is the force applied by the source on a completely constrained load. The Norton/Thévenin impedance is the impedance seen at the output terminals of the source when all external force and velocity inputs are inactive.

In a practical mechanical system, none of these quantities can be measured directly. If, for example, we wanted to characterise the chest using Thévenin's theorem, we would need to measure the pressure exerted on a surface that is constrained to zero motion. In practice, we cannot measure pressure without some finite displacement. Direct measurement of the Thévenin impedance is also impossible as, in a living subject, there is no way of 'switching off' all of the source terms (such as the heart beat, lung sounds and background vibration).

Despite not being amenable to direct measurement, Norton and Thévenin's theorems can provide a useful way of comparing results from different measurements on the chest. Several authors have proposed the use of the Norton surface velocity [163], or the Thévenin 'stop-surface' pressure [142] as the reference quantities to be used when comparing sensors. However, Vermarien and Vollenhoven [153] make an important point that limits the applicability of these theorems to stethoscope acoustics: the impedance of the chest is altered by the application of a sensor. This means that it is not possible to use Norton and Thévenin equivalent circuits in the traditional manner. Norton and Thévenin's theorems are, by definition, limited to *linear* systems, but the human chest is not linear. The impedance 'seen' by a load applied to the chest is a function of the load itself. In chapter 4 we show that the chest-stethoscope system can be linearised for small perturbations about an equilibrium position, which means that Norton and Thévenin's theorems can be applied to these perturbations.

3.6 Summary

In this chapter we outline the techniques required for modelling mechanical and acoustic systems using the lumped element method (LEM). The LEM is valid provided the spatial dimensions of the system are small compared to the shortest wavelengths of interest, which we have shown to be the case for stethoscopes (with the exception of the tubing).

We introduce definitions for frequency response functions (FRFs) between applied excitation and the resulting motion of a system. There are three characteristic features on FRF plots that play an important role in the interpretation of the experimental data in later chapters of this thesis: a single degree of freedom resonance, ‘mass-locking’ and ‘stiffness-shedding’. We show how these features arise using simple lumped element models and describe their characteristic features on FRF plots.

This chapter also lays out the electrical impedance and admittance analogies, which can be used to transform acoustic and mechanical systems into equivalent circuits. The benefits of this method are that it allows straightforward coupling between the acoustic, mechanical and electrical domains and that it enables a straightforward visualisation of the resulting coupled circuits. Furthermore, the impedance analogies enable the use of linear circuit analysis techniques to manipulate and interpret models. In particular, we discuss the use of the two-port method, in which transmission matrices are used to describe the propagation of waves through the chest-stethoscope system.

The methods and terminology introduced in this chapter underpin the analysis in the remainder of this thesis, and the concepts discussed herein are utilised without further introduction.

Chapter 4

Developing a theory for stethoscope acoustics

Stethoscopes are characterised by three human-sensor interactions: the interaction between the chest and the sensor, the holding of the sensor by the user and finally the interaction between the sensor and the ear.

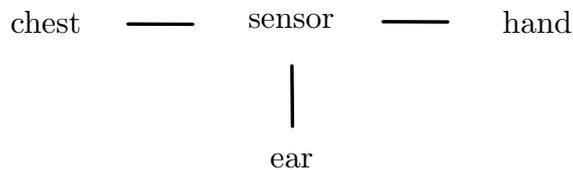


Fig. 4.1 Three human-stethoscope interactions. The nodes represent the stethoscope, the chest of the patient and the hands and ears of the clinician. The vertices represent the interactions between these nodes.

We argue that considering the nodes in figure 4.1 in isolation tells us little about the performance of stethoscopes. Instead, we need to address the vertices and develop a model for the entire coupled system.

In this chapter we consider what is meant by ‘the performance’ of a stethoscope (section 4.1), and then develop lumped-element models for direct contact sensors (section 4.2) and air-coupled sensors (sections 4.5 - 4.6) on the chest. To place this new theoretical work in context, we present (in section 4.3) a detailed technical review of the successes and shortfalls of existing models for stethoscope acoustics by Vermarien and van Vollenhoven [153], Kaniusas [67], Wodicka et al. [163],

Suzuki and Nakayama [142, 143] and Joyashiki and Wada [65, 66]. In section 4.4 we lay out a new theoretical approach to the coupled chest-stethoscope system which informs the low-order models derived in sections 4.5 - 4.6. The models proposed in this chapter are validated and explored in the remainder of this thesis.

4.1 Characterising stethoscope performance

The performance of a contact sensor can be characterised by comparing the measured variable-value to one of two references:

1. The actual value of the variable when the sensor is applied to the test piece.
2. The value that the variable would take if the sensor was not there.

The first of these is a property of the sensor itself and is typically defined as the ‘sensitivity’ of the sensor. The second is a property of the entire coupled system (consisting of both the sensor and the test-piece).

When characterising the performance of stethoscopes on the chest there are three crucial difficulties: sensor-impedance effects, non-linearity in the mechanical properties of the chest (with respect to the static load applied), and the difficulty in defining a suitable reference.

4.1.1 Sensor impedance (‘added mass’) effects

To illustrate the effect of sensor impedance on a system we consider the simple single-degree-of-freedom (henceforth ‘SDOF’) system shown in figure 3.5. The solid blue line in figure 4.2 shows the driving-point admittance FRF for this system. When we connect a sensor to this system in order to measure this FRF we alter the system. These alterations are discussed in detail in section 4.2. The green dotted line in figure 4.2 shows the driving point response of the new system. The key feature we can observe is that the original SDOF resonance shifts to a lower frequency. This is an example of a more general result: if we add mass to a system, we lower the resonance frequencies of all modes [87]. We can also observe that the amplitude of the peak in the coupled system is lower than that of the SDOF resonator. This is because we have additional sources

of damping in the sensor and at the system-sensor interface. At frequencies above the SDOF resonance we can see that the mass plateau (the level of the -20 dB/decade roll-off) differs from the unloaded case. At higher frequencies we observe a mass-locking effect and the coupled response tends towards the unloaded response (until, at much higher frequencies not shown here, the internal resonance of the sensor begins to affect the response).

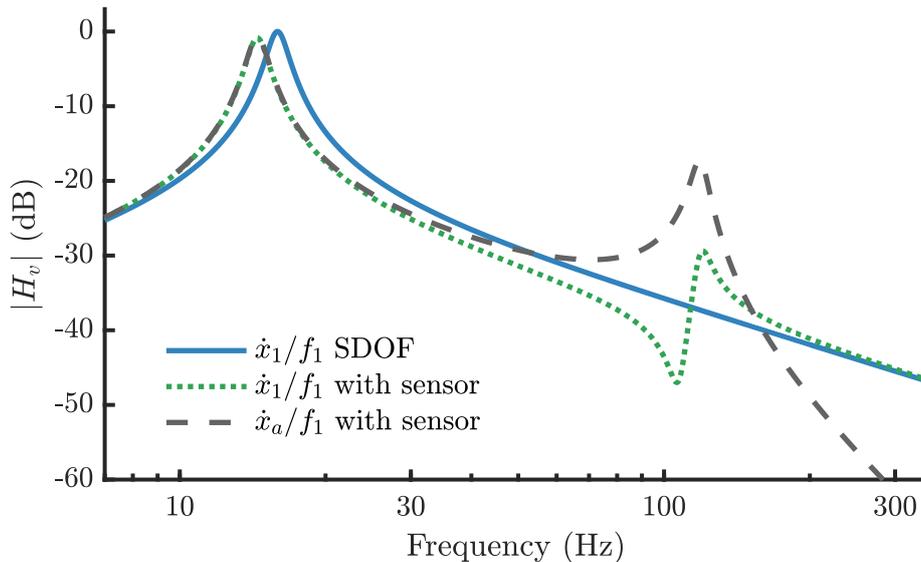


Fig. 4.2 Bode magnitude plot of the simulated admittance frequency response function for the SDOF oscillator from figure 3.5, with and without an added accelerometer. The model for the accelerometer is as described in section 4.2.1. Details of the simulation are given in appendix D.3. The decibel reference value is $1/\lambda_{1,0}$ (as defined in D.3).

The phenomenon described here is an example of what is commonly referred to as an ‘added-mass effect’. If the impedance of the sensor is significant compared to the *local* properties of the test structure, then the response deviates from the unloaded response.

4.1.2 Discrepancy between assumed and actual measurement

A further issue can arise if we think of a device as measuring a certain variable when it is, in fact, measuring a different variable. In section 4.1.1 we assume that the accelerometer gives a perfect measurement of the acceleration of the structure

on which it is mounted. In reality, the voltage output from an accelerometer is proportional to the deformation of the piezoelectric crystal within the accelerometer. In section 5.3.1 we show that this deformation is linearly proportional to the acceleration of the ‘seismic mass’ (see figure 4.3). However, the acceleration of the seismic mass differs from the acceleration of the test-structure.

The grey dashed line in figure 4.2 shows the simulated admittance FRF based on the acceleration measured by the accelerometer (\dot{X}_a/F_1). At low frequency this matches well with the loaded response of the test-structure (the green dotted line). However, at higher frequencies the two responses begin to differ. The measured response does not capture the anti-resonance seen at the driving point. After the second system resonance the measured response falls off sharply. This is because the accelerometer mass effectively ‘locks’ at these frequencies (evidence of this can also be seen in the mass-locking phenomenon observed in the driving-point response).

4.1.3 Non-linearity of the chest with respect to static load

In the example in figure 4.2 we assume that the properties of the SDOF resonator remain unchanged when the sensor is added. Previous work by other researchers (see e.g. [5, 56, 148, 153, 154]), and our measurements (in chapter 7) suggest that this is not the case for the chest. The mechanical properties of the chest exhibit a significant dependence on the static load applied. The effect of this non-linearity on the response of accelerometers on the chest has been partially investigated by Vermarien and van Vollenhoven, but (to our knowledge) the effects of chest non-linearity on air-coupled stethoscopes have not yet been investigated.

4.1.4 Difficulty in defining a reference

In the preceding sections we discuss some of the ways in which measurements made by sensors can deviate from the reference response that they aim to capture. For measurements of acoustic biosignals on the chest surface there is an additional complication: there is no clearly defined reference signal.

When investigating the acoustic signal obtained at the ear-pieces of an analogue stethoscope, it is important to consider with what these pressure perturbations should be compared. Possible examples of references include the pressure perturbations transmitted to the air if no stethoscope is present, the average velocity or volume velocity of some area of the unloaded chest surface, or a vibration defined somewhere within the body. In all of these cases there is a degree of arbitrariness in the definition, and there is no obvious ‘correct’ answer.

The literature depicts two main approaches to choosing references: the ‘*unloaded surface motion*’ (see e.g. [153, 157, 163, 142]) and the ‘*stop-surface pressure*’ (see e.g. [142]) which is the pressure perturbation that would be seen at the surface if this were loaded with an infinite impedance.

If we take the ‘*unloaded surface motion*’ as the reference, then we are faced with the issue that the motion of the chest surface is non-uniform. Non-uniform surface motion means that the value of the reference depends on the area which it is averaged over. This causes difficulties when comparing sensors with different contact areas. Even for a single stethoscope it is unclear whether we should pick the contact area between the chest and the entire chestpiece (including the rim) or just the contact area between the chest and the air-cavity to define the reference value. The fact that the reference, by nature, cannot be measured simultaneously with the actual sensor response also presents difficulties.

The ‘*stop-surface pressure*’ [142] is defined for a particular contact area, and can therefore be used to compare sensors of a given contact area. However, as for the unloaded-surface response, it is not possible to simultaneously measure the ‘*stop-surface pressure*’ and the actual response of a sensor. Furthermore, an infinite passive load impedance does not actually define a single reference, as there is still an arbitrary static load which influences the response (see section 4.2.4 for more detail on this). The ‘*stop-surface pressure*’ reference thus needs to be defined for a stated contact area and application force in order to be a valid reference. This has not been done in the literature.

If we take a reference point within the body, then, beyond the fact that this makes direct measurement of the reference very difficult, we need to consider that different body sounds have different physical origins within the chest. The location within the body that is picked as the ‘reference’ is largely arbitrary.

Concentrating on reference signals may cloud our view as to what we are actually trying to establish. The ultimate aim is to obtain a clear signal with a high signal to noise ratio and (arguably) limited distortion/filtering. For an analogue stethoscope the absolute magnitude of the acoustic signal is of importance, while for an electronic stethoscope it is only the signal-to-noise ratio (SNR) that matters. Unfortunately, a rigorous definition of signal-to-noise ratio also depends on having a known reference signal. Furthermore, while SNR can be used to compare stethoscopes, we are also interested in establishing the physical mechanisms that govern the response. For this purpose we have constructed a ‘chest phantom’: a simple laboratory model of the chest in which a reference input can be clearly defined.

4.2 LEM for a direct contact sensor on the chest

The simplest form of stethoscope is a direct contact sensor, such as an accelerometer, which is mounted on the chest and measures the motion of the loaded surface. The output of the sensor depends on the physical properties of both the sensor itself and the area of the chest to which it is coupled. We start by deriving a lumped element model for a generic direct contact sensor in section 4.2.1, and then consider how to model the impedance and excitation of the chest in section 4.2.2. Finally, we combine the two to create a full model for the coupled chest-sensor system.

4.2.1 Sensor

Figure 4.3 shows the general design of a *pressure mode* accelerometer. A typical accelerometer consists of a housing that is attached to the test-system using e.g. a bolt, wax or glue. The accelerometer-housing can also be held against the test-system by hand. Within the housing there is a ‘seismic-mass’ which is connected to the housing by a piezoelectric element. The piezoelectric element generates an electrical charge that is proportional to the relative displacement between the housing and the seismic mass. This relative displacement can be shown to be proportional to the acceleration of the seismic mass (see section 5.3.1), which is used as an estimate of the acceleration of the test-system at the measurement point. A *shear mode* accelerometer is designed with the seismic

mass and piezoelectric crystal mounted in a configuration that induces shear forces in the crystal in response to acceleration of the base. This design has several advantages (see e.g. [10]), but can be represented by the same lumped element model [10], so will not be treated separately here.

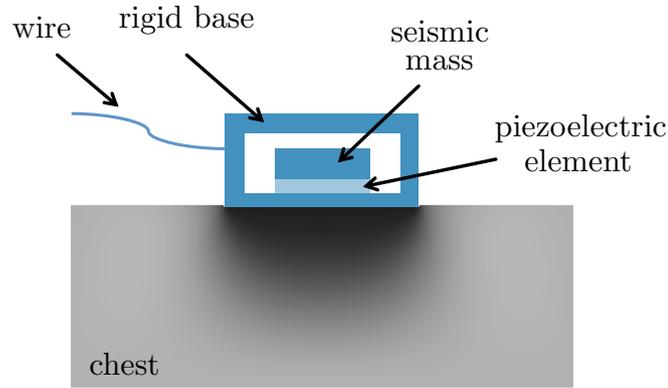


Fig. 4.3 Simplified schematic of a pressure mode accelerometer on the chest. See main-body text for an explanation of the components.

Figure 4.4 shows a lumped element model for a generic direct contact sensor. The base of the sensor (i.e. the housing) has mass of $m_{s:b}$, and is coupled to the chest surface with a spring of stiffness $k_{c,s:b}$. This stiffness is referred to as the *mounting stiffness*, and depends on the manner of application. The piezoelectric element between the base mass and the ‘seismic mass’ ($m_{s:a}$) has a stiffness of $k_{s:b,s:a}$. The base of the sensor may also be connected to mechanical-ground by a load impedance. In figure 4.4 the load impedance is represented by a stiffness $k_{s:b,0}$ but this can be replaced by any load impedance z_{load} . The load impedance could represent, for example, the sensor’s wire, or the holding of the sensor against the chest by a clinician.

In this lumped element model, damping terms have been omitted from the schematic for simplicity, but each spring in fact represents a suitable constitutive model, which may also include damping and inertance.

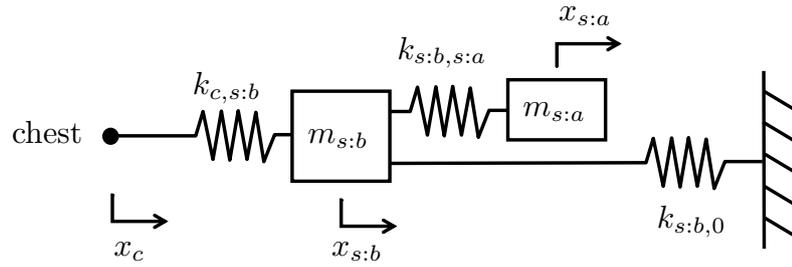


Fig. 4.4 Lumped mechanical model of a generic direct sensor. x_c is the displacement of the chest surface, $x_{s:b}$ is the displacement of the sensor base, and $x_{s:a}$ is the displacement of the sensor's seismic mass. The lumped elements are described in the main-body text. $k_{s:b,0}$ can be replaced by any complex impedance z_{load} .

4.2.2 Chest

The human chest is a complex, distributed, three-dimensional system. When we talk about including ‘the chest’ in our lumped element model of the coupled chest-sensor system we are not trying to capture the chest as a whole, but rather the chest as ‘seen’ by the sensor. Accordingly, our lumped element model for ‘the chest’ depends on the sensor used. In particular, the chest model is a function of the contact area, the contact type (rigid or deformable) and the static load from the sensor. A further complication stems from the fact that there is also a dependance on the location of the sensor on the chest, on the anatomy and posture of the person, and even on the breathing cycle. In this thesis the term ‘chest’ is used as a general placeholder for any part of the human body that a stethoscope is used on. It is impractical to take measurements on every person, in every position, at every location throughout each section of the breathing cycle. Therefore, trends need to be established with respect to contact area, contact type and static load, for a representative case, with insight into the variability introduced by other factors.

By applying Thévenin’s theorem, we assume that the chest is a linear system. The properties of the chest are non-linear with respect to the static load applied [153, 154]. However, the force fluctuations due to body sounds are small compared to the static equilibrium forces, and as such can be treated as small linearised fluctuations about an equilibrium position in a non-linear system. This means that we can use a single source term and a single impedance term to characterise ‘the chest’, but that the dependance of these terms on contact area, contact type and static load needs to be considered for each individual sensor in each

operating condition. Many researchers (e.g. Wodicka et al. [163] and Suzuki and Nakayama [142]) have made the mistake of considering only the effect of the contact area while neglecting the effects of the contact-type and the static load applied.

For a direct contact sensor the chest-sensor interface is typically rigid, which means that the relevant output impedance of the chest (as seen by the sensor) can be measured for a given contact area and load using a rigid coupler. The methodology for these measurements is described in chapter 7. Chest impedance measurements in the literature suggest that the chest (as ‘seen’ by a rigid disk in contact with the surface) can be modelled as a single-degree-of-freedom system within the frequency range of interest for auscultation [56, 148, 154, 153]. By fitting a single-degree-of-freedom LEM to the response it is possible to extract a lumped effective mass m_{eff} , stiffness k_{eff} and damping coefficient λ_{eff} .

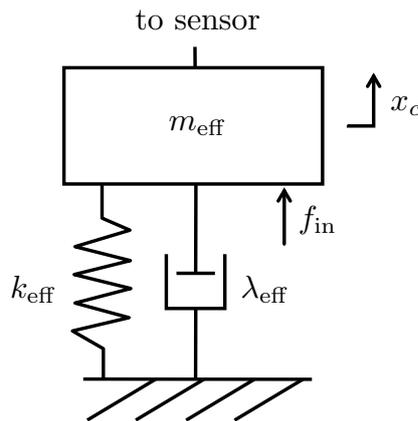


Fig. 4.5 SDOF model for the chest, as seen by a rigid sensor. A Thévenin equivalent force (f_{in}) is applied to x_c .

The effective lumped parameters naturally depend on the contact area (\hat{A}_c) and the static load (f_{static}) with which the sensor is applied against the chest. In other words, for a rigid sensor at a given location on a given chest:

$$m_{\text{eff}} = m_{\text{eff}}(\hat{A}_c, f_{\text{static}}), \quad (4.1)$$

$$k_{\text{eff}} = k_{\text{eff}}(\hat{A}_c, f_{\text{static}}), \quad (4.2)$$

$$\lambda_{\text{eff}} = \lambda_{\text{eff}}(\hat{A}_c, f_{\text{static}}). \quad (4.3)$$

To model the response of a stethoscope to body sounds we need to consider how to model the excitation source. The simplest way to do this is to assume a

Norton or Thévenin input term *for a given contact area and static load.*

$$f_{\text{in}} = f_{\text{in}}(\hat{A}_c, f_{\text{static}}). \quad (4.4)$$

Figure 4.5 shows a SDOF LEM for the chest with an input force f_{in} applied to x_c . This input force is not an external force applied to the chest-surface, but rather an ‘equivalent force’ that represents the source terms within the chest from the ‘viewpoint’ of a sensor at the chest surface. This is analogous to using a Thévenin equivalent circuit of the impedance-analogy circuit for the system as shown in figure 4.6.

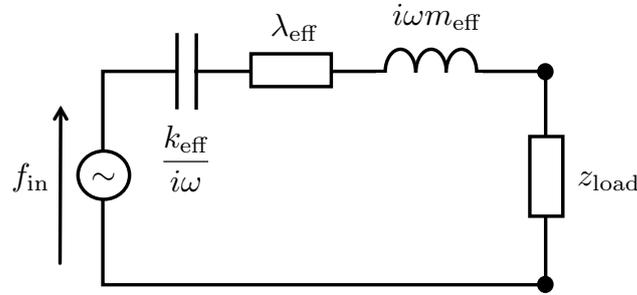


Fig. 4.6 Impedance analogy circuit for the lumped mechanical system in figure 4.5. The Thévenin input force (f_{in}), and the Thévenin impedance ($z_{\text{eff}} = i\omega m_{\text{eff}} + \lambda_{\text{eff}} + k_{\text{eff}}/i\omega$) are defined for a given contact area and static load. When a sensor is used $z_{\text{load}} = z_{\text{sensor}}$.

4.2.3 Impedance analogy circuit for the coupled system

Figure 4.7 shows an impedance analogy circuit for a direct sensor coupled to the chest. To generalise the circuit, the capacitance corresponding to the $k_{s:b,0}$ spring is replaced by a general impedance z_{load} to allow the sensor to be constrained by any passive load. We refer to this load on the sensor as the *attachment*. The attachment could simply be the wire of the accelerometer, but it could also represent a clinician’s hand pressing the sensor against the chest.

The circuit has been divided into three segments: the chest, the attachment and the sensor. The chest impedance and the excitation force are both functions of the contact area and the static load applied. To generalise this circuit we can convert it into a Thévenin equivalent circuit, replacing the three chest terms with a single impedance z_{chest} that can be measured.

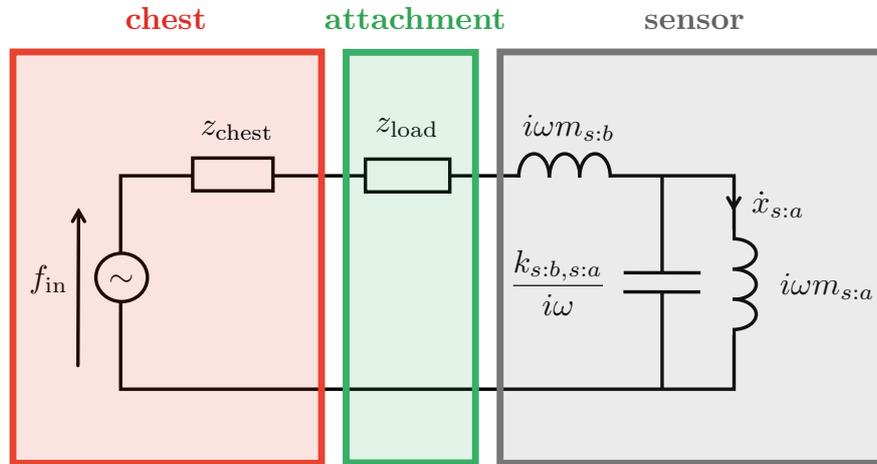


Fig. 4.7 Coupled impedance analogy circuit for a direct sensor on the chest. The contact stiffness is included in the Thévenin chest impedance (z_{chest}). The order of series terms does not effect the circuit, which allows the attachment impedance z_{load} to be positioned between the chest and the sensor rather than between the sensor and ground.

4.2.4 Response functions

Considering the coupled circuit in figure 4.7, and particularly the fact that the impedance of the chest is a function of both the sensor area and the static load applied, we can see that it is not possible to define the sensor response with respect to the unloaded response using this model. This is because if there is no load the contact area is undefined, and thus the chest impedance is also undefined. The ‘stop-surface reference pressure’ proposed by Suzuki and Nakayama [142] could be employed if we take into account that this is defined for a specific contact area and static force. In the circuit in figure 4.7, $f_{\text{in}}(\hat{A}_c, f_{\text{static}})$ is the ‘stop-surface’ reference force for the given equilibrium position.

In experiments in which a phantom with a measurable excitation signal is used, the model can also be employed to simulate the response with respect to this excitation signal. This requires a model for the phantom, as discussed in chapter 6.

4.3 Review of existing models for stethoscopes

In this section we review existing models that have been proposed to explain the acoustics of air-coupled stethoscopes. We highlight the shortcomings of these models, thereby motivating the need for a new theory for stethoscope acoustics. In section 4.3.1 we consider the claim that the chestpiece acts as a Helmholtz resonator, in section 4.3.2 we discuss models which neglect the presence of the rim of an air coupled sensor, and in section 4.3.3 we critique Suzuki and Nakayama’s [143] method of treating the air-cavity as rigid.

4.3.1 Helmholtz resonator

It is often suggested that the air-cavity of a stethoscope can be modelled as a Helmholtz resonator (see e.g. Rappaport and Sprague [129] and Kaniusas [67]). As outlined in this section, we believe this suggestion to be misleading. The term ‘Helmholtz resonator’ is used to describe a system in which an acoustic-mass and a constrained volume form a mass-spring system that resonates at a particular frequency. The mathematics of Helmholtz resonators are derived in section 3.2.3. Considering a duct of area A and length l filled with air of density ρ_0 and speed-of-sound c_0 , the lumped acoustic mass is $\rho_0 l/A$. The acoustic stiffness of a cavity of volume V and contact area A is cavity is $\rho_0 c_0^2 A^2/V$, leading to a resonance at $\omega_n = \sqrt{c_0^2 A/lV}$.

Kaniusas [67] uses this expression for the resonance of a Helmholtz resonator to argue that the resonance frequency of a stethoscope increases as the tubing area increases, the cavity volume decreases and the tubing length decreases.

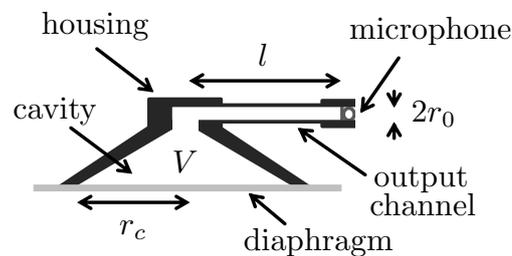


Fig. 4.8 Schematic to represent Kaniusas’s model for the ‘bell’ of an air-coupled microphone. Based on Fig. 4.28 in *Sensing by Acoustic Biosignals* (Kaniusas, 2015) [67]. V is the volume of the cavity (not including the output channel), r_c is the cavity-surface contact radius, r_0 is the radius of the output channel, such that $A = \pi r_0^2$; and l is the length of the output channel.

The problem with this model becomes evident when we consider the topology of the system that Kaniusas suggests. Figure 4.8 shows our interpretation of Kaniusas's model. The excitation is provided by the motion of the diaphragm which is in contact with the cavity volume. A short air-duct on the far side of the volume leads to a microphone. Figure 4.9 shows an impedance-analogy circuit that we have constructed for the system in figure 4.8. From this circuit it is apparent that Kaniusas's model does not, in fact, give a resonating circuit. The pressure is applied to the volume, while the lumped acoustic-mass is rigidly fixed due to the high impedance of the microphone. For this circuit it is evident that the pressure at the microphone (p'_3) simply equals the fluctuating input pressure (p'_1), with no resonant behaviour.

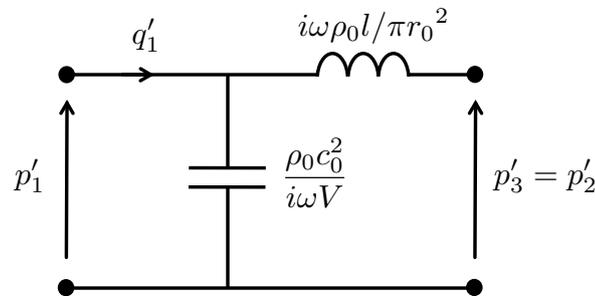


Fig. 4.9 Electrical impedance analogy circuit for Kaniusas's air-coupled microphone (shown in figure 4.8). The microphone impedance is high and is approximated here by an open circuit. This is not a resonating circuit as the mass in the output channel does not move. (For comparison see the impedance analogy circuit of a Helmholtz resonator in figure 3.15).

A further problem with this model is that it does not correctly account for the effect of the duct length between the air-cavity and the microphone. Experimental results (see e.g. [29]) show that changes to the tube length only have a small effect on the frequency of the fundamental system resonance. For a Helmholtz resonator the resonance frequency would be inversely proportional to the square-root of the tubing length. Furthermore, for typical tubing lengths the lumped-mass approach is no longer valid, and the tubing must instead be treated as a transmission line. Kaniusas uses the Helmholtz resonator model to discuss the response of stethoscopes with long tubing, despite the fact that it then becomes somewhat arbitrary as to how much mass to lump to the start of the tube.

4.3.2 Rimless contact

Many researchers who propose impedance-coupling models for the stethoscope take the problematic approach of treating an air-coupled sensor as if it were a direct contact sensor (see e.g. [163, 142]). This approach is characterised by considering the coupling of a constrained air-cavity to the chest surface, without taking into consideration the role played by the rim surrounding the cavity. When a stethoscope is pressed against the chest, the chest is loaded by both the rim and the cavity, so that the motion of the loaded chest surface will depend, in part, on the rim. The pressure in the air-cavity depends on the relative motion between the stethoscope rim and the chest surface ('seen' by the cavity). Consequently, omitting the presence of the rim entirely from a model of the coupling between the stethoscope and the chest is problematic.

Wodicka et al.'s 1994 paper [163] is a classic example of this approach. Figure 4.10 shows our interpretation of the model presented in this paper. At first inspection this seems to make sense: a Thévenin equivalent lumped element model for the chest is coupled to a T-element model for the air-cavity (see e.g. Beranek [6] for a description of the T-element model), which is in turn coupled to a SDOF model for the microphone.

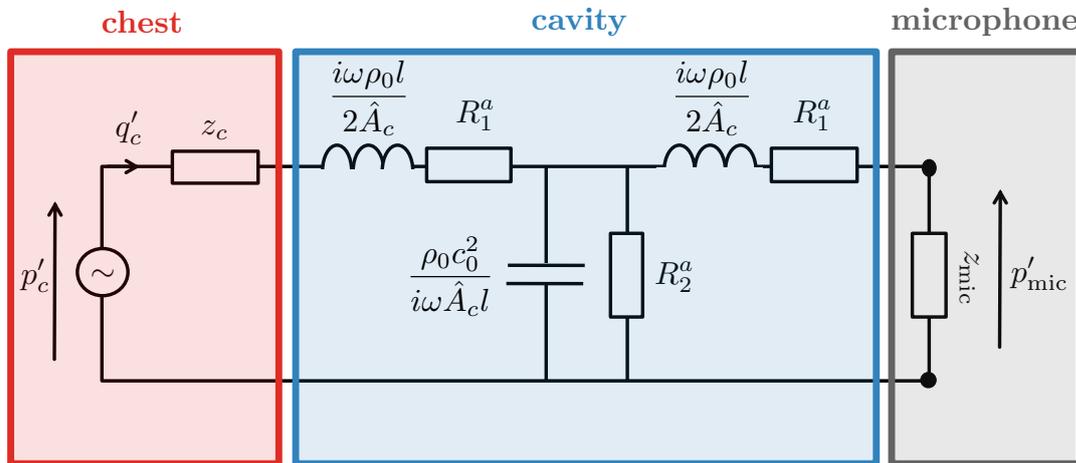


Fig. 4.10 Electrical impedance analogy circuit for an air-coupled microphone coupled to the chest according to Wodicka et al. (1994) [163]. This circuit is based on that given in Figure A1. in Measurement of Respiratory Acoustic Signals: Effect of Microphone Air Cavity Depth (Wodicka, 1994) [163]. The nomenclature has been adapted for consistency with the rest of this thesis. R_1^a and R_2^a are damping terms for the short duct segment.

On closer inspection, however, the approach breaks down. The chest impedance z_c used in this model is that ‘seen’ by a rigid disc with an area matching the cavity-contact area. By using a chest impedance model based on a measurement with a rigid contact, the model is based on the implicit assumption that the contact surface between the chest and the air-cavity is rigid, which is incorrect. Furthermore, the presence of the rim of the air-coupled sensor, which also loads the chest wall, is completely omitted. In Wodicka et al.’s model, the microphone backplate is fixed to mechanical ground, rather than to a movable housing. This would be applicable if the stethoscope was rigidly fixed to a wall and the patient leaned against it, which would be highly unusual in a clinical setting. Figure 4.11 shows a lumped mechanical system that corresponds to Wodicka et al.’s proposed circuit. Several important features of the coupled chest-stethoscope system are not captured in this circuit.

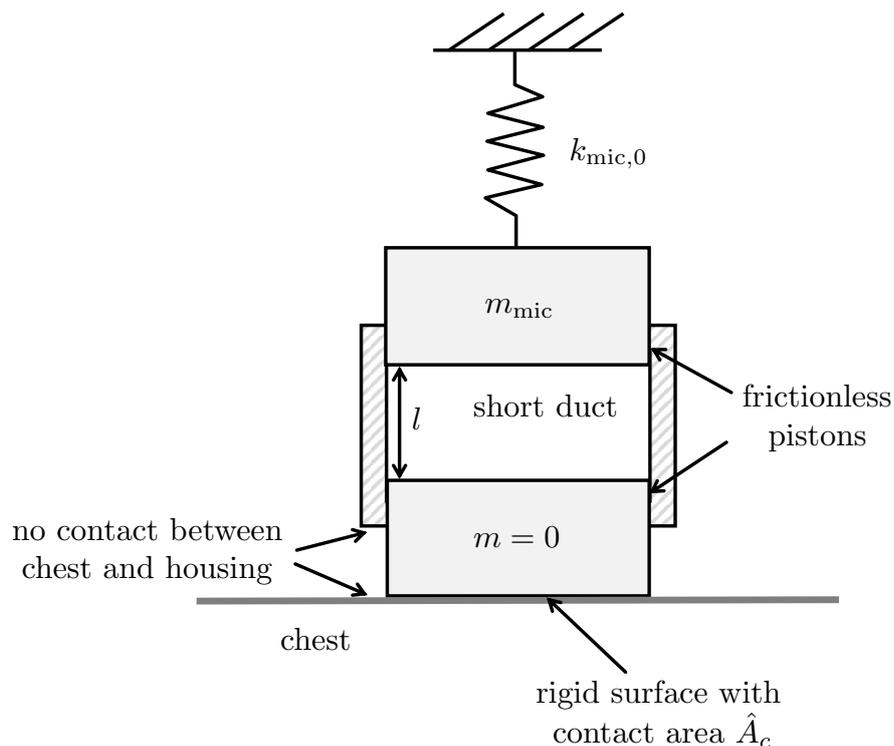


Fig. 4.11 Representation of a physical system corresponding to Wodicka et al.’s impedance analogy circuit. There is no contact between the chest and the housing of the sensor.

Suzuki and Nakayama (1998, 2000) [142, 143] take a similar approach. They make some sensible simplifications: taking the microphone impedance as infinite and modelling the cavity as a pure compliance rather than a short duct. Suzuki

and Nakayama estimate the mechanical impedance of the chest ‘seen’ by a rigid disc of a given area by analytically modelling the vibration of the disc on a semi-infinite elastic medium representing the chest. This allows them to explore the design-space more easily than relying entirely on experimental measurements of chest impedance. However, their model is still based on the assumption of a rigid contact.

Neither Wodicka et al. (1994) nor Suzuki and Nakayama (1998, 2000) include in their models the mass of the sensor or the stiffness with which the sensor is held in place. In chapter 9 we present experimental results that demonstrate that these two factors play a significant role in determining the dominant low frequency resonance in the response.

Joyashiki and Wada (2020) [66] build on the model of Suzuki and Nakayama by including the mass of the housing in their model. Joyashiki and Wada claim that they are modelling the case where the sensors are stuck to the chest wall rather than held firmly in place by a clinician. However, their model does not seem to match this case. Joyashiki and Wada’s equation for the response of an air-coupled sensor adds together the impedance due to the sensor mass ($i\omega m_s$) and the impedance due to the cavity stiffness ($k_a/i\omega$), thus implying that these terms are in series in the impedance analogy circuit (and therefore in parallel in the mechanical system) as shown in figure 4.12. This means that Joyashiki and Wada (unintentionally) model a system in which the housing mass directly loads the chest surface and there is an air cavity (of stiffness k_a) between this mass and mechanical-ground. Note that, beyond this topological error, the form of the response function presented by Joyashiki and Wada is difficult to interpret physically, as discussed in appendix C.4.

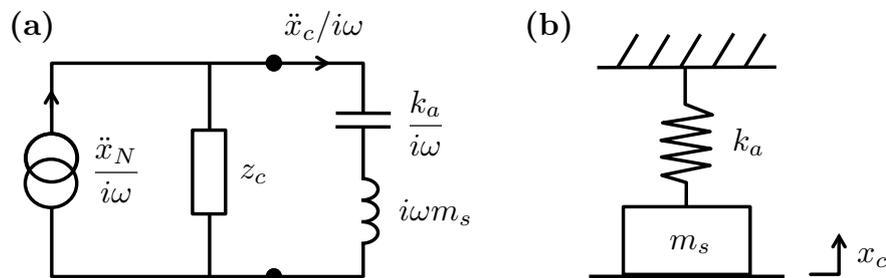


Fig. 4.12 (a) Norton equivalent electrical impedance analogy circuit that gives a response of $\ddot{x}_c/\ddot{x}_N = z_c/(z_c + (i\omega m_s + k_a/i\omega))$. (This ratio of impedances is given in the equation for the ‘pressure response’ of air-coupled microphones by Joyashiki and Wada (2020) [66].) (b) Corresponding mechanical system.

Joyashiki and Wada's model can be corrected by combining the housing-mass and cavity-stiffness impedances in parallel in the impedance analogy circuit as shown in figure 4.13. The pressure on the microphone can then be obtained by scaling the force on the housing-mass ($f_m = \ddot{x}_m m_s$) by the contact area \hat{A}_c :

$$p_m = \ddot{x}_m m_s / \hat{A}_c. \quad (4.5)$$

While this adaption makes more physical sense in terms of the topology of the lumped elements, it still does not capture the mechanics of the rim-cavity interaction. In contrast to this model, we do not expect all of the force exerted on the housing mass to be exerted through the air-cavity.

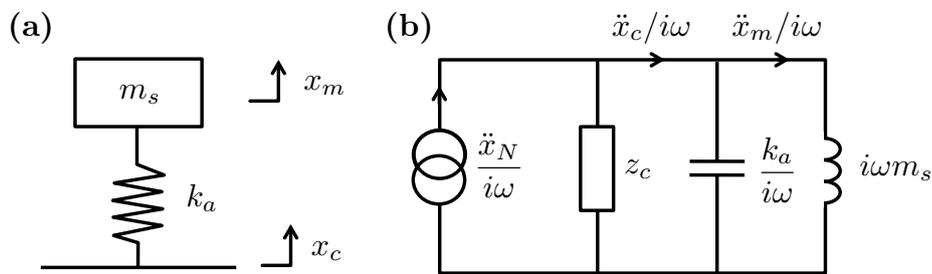


Fig. 4.13 (a) Schematic for a simple mechanical model of a (rim-less) air-coupled sensor with cavity stiffness and mass behind the microphone with no connection to mechanical ground. (b) Corresponding Norton-equivalent impedance analogy circuit.

4.3.3 Rigid cavity with rim

In section 4.3.2 we discuss models for the stethoscope that ignore the stethoscope's rim and find these to be inadequate. Suzuki and Nakayama [143] propose an alternative modelling approach in which the rim is included but the finite compliance of the air-cavity is neglected. In the models described in section 4.3.2 the finite compliance plays a dominant role in determining the response, so these two approaches are in direct contradiction. However, despite discussing both, Suzuki and Nakayama do not attempt to combine or reconcile them.

Figure 4.14 shows a cross-section of a simple air-coupled sensor coupled to the chest. The gross contact area is πr_o^2 , while the cavity-contact area is πr_i^2 .

Suzuki and Nakayama make the assumption that the entire air-coupled sensor (including the air-cavity) can be modelled as a rigid, immobile cylinder. In order

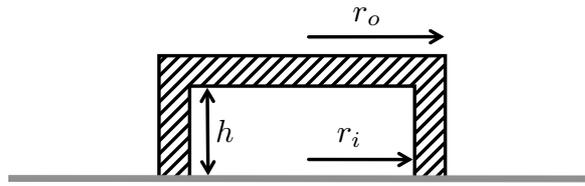


Fig. 4.14 Dimensions for a simple air coupled sensor, on a cross-section of the cylinder. The cavity volume is given by $\pi r_i^2 h$.

to estimate the pressure in the cavity, they employ the principle of superposition. Suzuki and Nakayama compute the stress profile in the ‘chest’ surface when the base of the ‘chest’ is fixed and the rigid cylinder is oscillating at a prescribed rate. They then integrate the normal stress over the cavity-contact area ($r = 0 \rightarrow r_i$) and divide this integral by the cavity-contact area (πr_i^2) to obtain an estimate for the pressure in the cavity. Suzuki and Nakayama compare this pressure estimate to the ‘stop-surface’ pressure for the same total contact area (including the rim). Since the model they use to compute the stress profile considers the stethoscope as a rigid immobile cylinder, the ‘stop-surface’ pressure can be obtained by integrating the normal stress from $r = 0 \rightarrow r_o$ and normalising by πr_o^2 .

There are numerous issues with this approach. The response ratio computed by Suzuki and Nakayama depends only on the contact-area of the cavity and the rim, while the depth of the cavity has no influence. This contradicts experimental results presented by numerous authors (see e.g. [163]). The central assumption in this method (that the air-cavity is stiff enough compared to the chest that the entire sensor can be modelled as a rigid body) is unphysical. To illustrate this we performed a series of ABAQUS CAE simulations. We were interested in the effect of the finite cavity-compliance on the surface-stress integral, not the frequency dependance of the ‘pressure-ratio’, so we used a simple static simulation. Figure 4.15 shows the surface deflection profile obtained when a stethoscope (with dimensions matching those used by Suzuki and Nakayama) is displaced downwards by 0.1 mm into an elastic medium representing the chest. Parameter values for the simulation were chosen to match those presented in the literature for the silicon elastomer Ecoflex 00-10 to match experimental work in this thesis which uses this material. The figure illustrates that, despite the high stiffness of the cavity (due to its small volume), the skin surface in contact with the air-cavity still deforms significantly (with respect to the displacement of the rim).

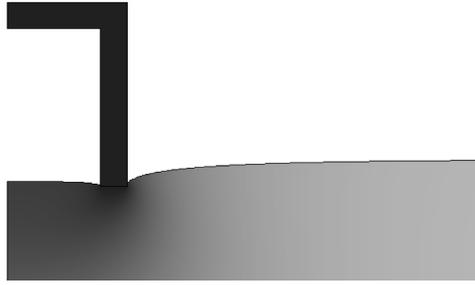


Fig. 4.15 Axisymmetric ABAQUS simulation of an air coupled sensor pushed against a viscoelastic cylinder. The sensor dimensions for this simulation are: 35 mm cavity diameter, 45 mm outer diameter of rim and 60 mm cavity depth. The shading represents downward vertical displacement.

Following the method proposed by Suzuki and Nakayama, we extracted the normal stress profile at the surface of the ‘chest’ when the sensor is displaced downwards. Figure 4.16 shows the (axisymmetric) stress distribution at the surface for two simulations: a rigid cylinder and an air-coupled sensor with a cavity. Both simulations used the same outer diameter $2r_0 = 9$ mm. The effect of the cavity on the stress profile is clearly evident: the normal stress over the cavity-contact area is significantly lower for the cavity case than it is for the rigid case.

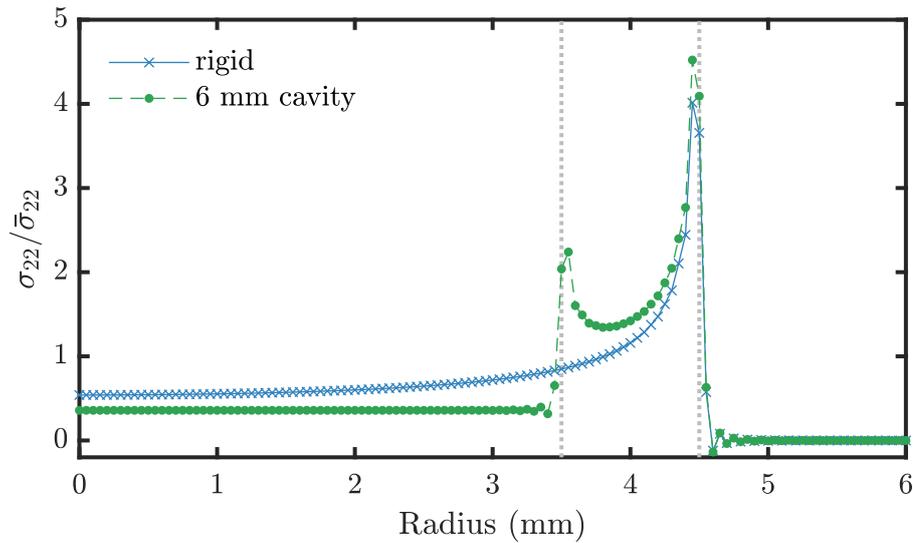


Fig. 4.16 Normalised normal stress at the ‘chest’ surface when an air-coupled sensor with an outer diameter of 9 mm, a cavity diameter of 7 mm and a (cylindrical) cavity depth of 6 mm is displaced downwards. The blue crosses show the stress distribution obtained if the entire sensor is treated as a rigid body, the green dots show the stress distribution obtained if the cavity is included in the model. The vertical grey dotted lines show the position of the sensor rim.

We then computed the pressure-ratio as defined by Suzuki and Nakayama. Figure 4.17 shows the results of simulations with a range of cavity volumes. For the geometry used by Suzuki and Nakayama, with a cavity depth of 6 mm, the rigid-contact assumption leads to an over-estimation of the pressure response by around 70%. This analysis suggests that the method presented by Suzuki and Nakayama significantly overestimates the pressure in the air-cavity of an air-coupled sensor.

The larger the cavity volume, the lower the pressure response, with a limit of $PR \simeq 0.06$ for an open cavity (for these dimensions). This suggests something is amiss: if the cavity is open then the cavity should be at atmospheric pressure (neglecting radiation impedance). The reason that we still obtain a non-zero pressure response is that we are defining the cavity pressure as the integral of the normal stress at the surface. Suzuki and Nakayama assume that the normal-stress profile at the surface matches the pressure in the cavity. However, the chest can also transmit shear forces. This means that the pressure in the cavity may be lower than the normal-stress profile suggests. Extracting the

cavity pressures from the ABAQUS simulations, we see that these are indeed lower than the averaged normal stress, as shown by the blue circles in figure 4.16.

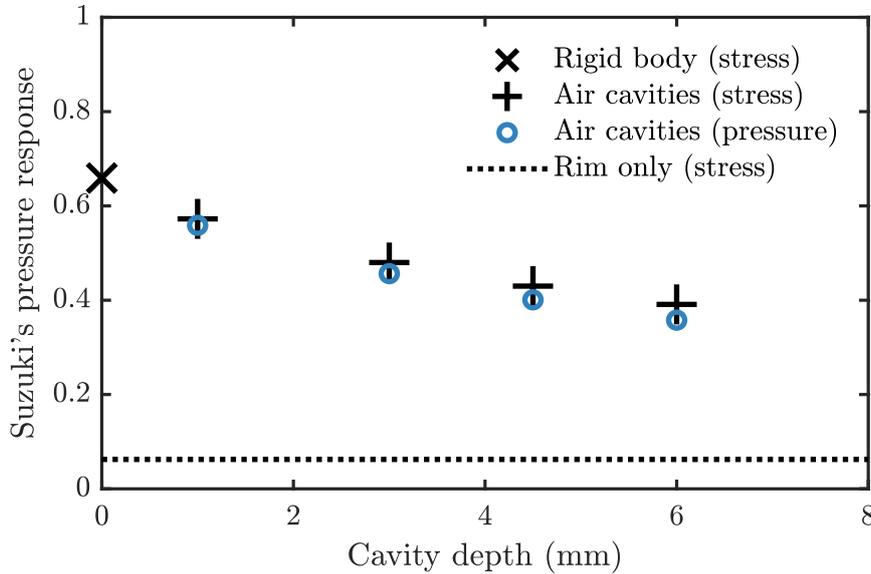


Fig. 4.17 Non-dimensional pressure response as defined by Suzuki and Nakayama [143], computed from ABAQUS CAE simulation results for an air-coupled sensor with an outer diameter of 9 mm, a cavity diameter of 7 mm and a range of (cylindrical) cavity depths. The black cross (\times) shows the pressure ratio obtained when following the method of Suzuki and Nakayama and modelling the entire contact as rigid. The black plus symbols ($+$) show the pressure ratios obtained by integrating the normal stress in simulations which included air cavities of different depths. The blue circles (\circ) show the pressure ratio evaluated for the same simulations, using the pressure in the air-cavity rather than the surface-stress integral. The black dotted line shows the case for toroid (i.e. an infinite cavity depth) with zero radiation impedance. This case is also referred to as ‘rim only’ or ‘open cavity’.

4.3.4 Conclusions

The shortfalls with existing models for stethoscope acoustics outlined in this chapter allow us to identify what is needed for a new, rigorous theory of stethoscope acoustics. It is evident that modelling a subsection of the coupled chest-stethoscope system is inadequate and misleading. The entire coupled system should be included within a single model. Furthermore, the model should take into account the contact between the rim and the chest, the finite compliance of the air cavity, the passive impedance of the stethoscope housing and its attach-

ment, the static force with which the stethoscope is applied and the non-linear properties of the chest. In section 4.4 we derive a new model for air-coupled stethoscope acoustics which takes into account all of these considerations.

4.4 A new theory for stethoscope acoustics

When a stethoscope is held against the chest a coupled system is formed. This coupled system is defined in terms of two length-scales. The displacement of the rim into the chest (and of the chest-tissue into the air-cavity) when the stethoscope is pressed against the chest has length-scales of the order of millimetres. The displacement perturbations associated with body-sounds are several orders of magnitude smaller than this [153]. We thus think of the coupled chest-stethoscope system in terms of a static equilibrium position and small perturbations about that equilibrium position.

The properties of the chest, rim, cavity and attachment (in particular the static load applied) all play a role in setting the equilibrium position for the coupled system. The system is non-linear, which plays an important role when setting the equilibrium position. However, for small perturbations the non-linearities do not play an important role and we can linearise the response about the static equilibrium position. The non-linearity arises from the fact that the mechanical properties of the system change as the stethoscope is pressed into the chest. The properties of human tissue change as it is compressed. Furthermore, when no artificial diaphragm is present, flesh can fill a significant proportion of the cavity. The extent of this filling depends on how firmly the stethoscope is pushed against the chest. The filling changes both the cavity volume and the cavity-contact area and thus alters the mechanical properties of the cavity.

For a given equilibrium position, body sounds cause small pressure fluctuations in the air-cavity of the stethoscope. The pressure perturbations in the air-cavity are related to fluctuations of the cavity volume which occur because of the relative motion between the rigid chestpiece housing and the deformable chest-surface that is in contact with the air-cavity. This relative motion is resisted by the stiffness of the air-cavity but also, importantly, by the chest tissue itself.

The presence of a diaphragm covering the cavity opening changes the system significantly. The diaphragm increases the stiffness of the cavity surface. This not

only affects the perturbations about the equilibrium position, but also means that, for the same static load applied to the chestpiece, a different static equilibrium position is set.

In the case of an air-coupled microphone (ACM) the pressure perturbations induced in the air-cavity by body sounds are measured by a microphone in the cavity wall. In the case of an analogue stethoscope the pressure perturbations are transmitted as compression waves through the tubing to the ear-pieces. In this chapter we derive a model for an ACM, which is extended to include tubing in chapter 10.

4.5 Derivation of a LEM for air-coupled stethoscopes

In this section we derive a lumped element model for an air-coupled sensor pressed against the chest. To do this, we gradually build up the system starting with a rigid toroidal rim in contact with the chest. The purpose of this derivation is to establish the topology of the LEM (the way in which the components are connected). We initially neglect damping and construct the LEM using only masses and springs. Once the topology of the system is established the springs can be replaced by appropriate constitutive models which account for stiffness, damping and inertance.

4.5.1 Rim

Figure 4.18 shows the vertical displacement profile of the chest when a rigid toroidal rim is pressed against it. The free surface of the chest within the rim is clearly influenced by the motion of the rim. When the rim is displaced downwards the average displacement of the free surface within the rim is also downwards, with an amplitude that is slightly less than that of the rim.

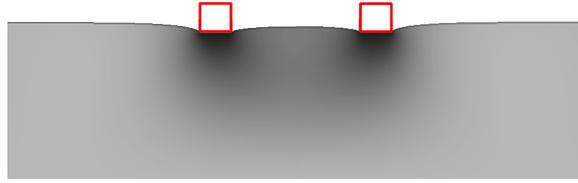


Fig. 4.18 Cross section showing the vertical displacement of the chest when a rigid rim is press down against it. The colormap shows vertical displacement with respect to the unloaded equilibrium position from an ABAQUS CAE simulation (dark is downwards).

To capture this effect, we define two degrees-of-freedom: the rim (x_s) and the surface (x_c). These DOFs are connected by a lumped stiffness ($k_{c,s}$) as shown in figure 4.19. This stiffness could be referred to as a ‘clamping stiffness’, as it is the clamping of the chest surface to the rim, that resists the motion of the proportion of the chest surface that is in contact with the chest cavity. Note that x_c is interpreted as the area-averaged vertical displacement of the free surface within the rim. Each of the DOFs we have defined has an inertia term associated with it. For the rim the inertia is given by the mass of the rim (m_s) and a small lumped-mass term from the chest tissue ($m_{ch:s}$).

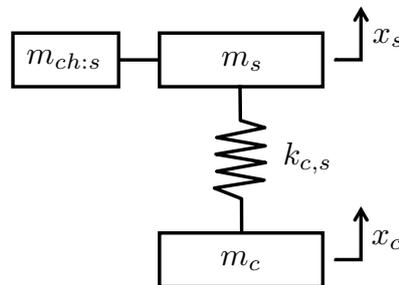


Fig. 4.19 Air coupled sensor lumped element model derivation - step one. m_s is the mass of the toroid, $m_{ch:s}$ is the part of the chest’s mass that is lumped to x_s , m_c is the part of the chest’s mass lumped to x_c and $k_{c,s}$ is the lumped stiffness between x_s and x_c . For simplicity $m_{ch:s}$ is included as part of m_s in subsequent figures.

For the free surface within the rim, the inertia term is a lumped mass (m_c) representing a proportion of the kinetic energy in the chest. A specific lumping is only valid for a particular frequency range. Using a single DOF to describe the motion of the surface will only allow a single mode of vibration to be captured [31]. Capturing the higher frequency modes with which the surface vibrates requires additional DOFs.

4.5.2 Air cavity

Next we consider the effect of a closed air cavity within the rim. We now refer to the rigid body as the ‘housing’, with ‘rim’ specifically referring to the portion of the housing that is in contact with the chest. One side of the air cavity is in contact with the chest surface, while the other side is in contact with the housing. Figure 4.20 shows a colour-map of the vertical displacement of the chest when the housing is pressed downwards. Compared to the toroidal-rim case (figure 4.18) the cavity-surface displacement matches that of the housing more closely, suggesting a greater stiffness between the two DOFs.

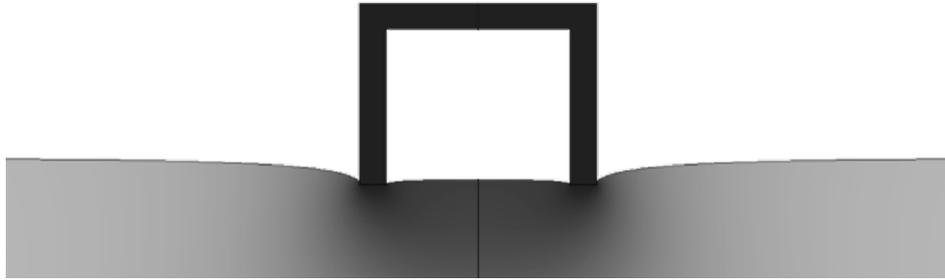


Fig. 4.20 Vertical displacement of the chest when a rigid stethoscope chestpiece (‘housing’) is pressed downwards. The colormap shows vertical displacement with respect to the unloaded equilibrium position from an ABAQUS CAE simulation (dark is downwards).

In section 3.2.2 we show that the dominant effect of a small constrained volume of air is to provide a stiffness term $k_a = \rho_0 c_0^2 \hat{A}^2 / V$. This ‘acoustical stiffness’ is in parallel with the ‘mechanical stiffness’ ($k_{c,s}$), as shown in figure 4.21. The mass of the air in the cavity is negligible and, at low frequencies, the shape of the cavity does not affect the response (see e.g. [129]).

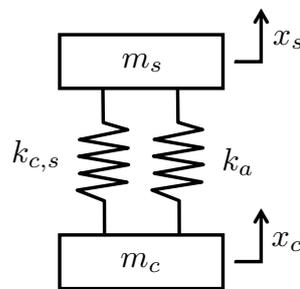


Fig. 4.21 Air coupled sensor lumped element model derivation - step two.

4.5.3 Chest

We now turn our attention to the chest. The chest tissue is already responsible for the mass lumped to the cavity-surface (m_c), and the mechanical stiffness between the surface and the sensor housing ($k_{c,s}$). The air cavity surface is also connected to the rest of the chest, but the most suitable model for this is not intuitively clear. Fortunately, we can make use of the fact that any linear system of impedance and source terms can be represented in terms of Thévenin or Norton's theorems (see section 3.5.7). Converting these theorems into the mechanical domain we obtain the LEMs show in figure 4.22, where $k_{0,c} = k_{b,c}$ is a complex term representing the impedance ($z_{0,c} = k_{0,c}/i\omega$) of the chest seen from x_c .

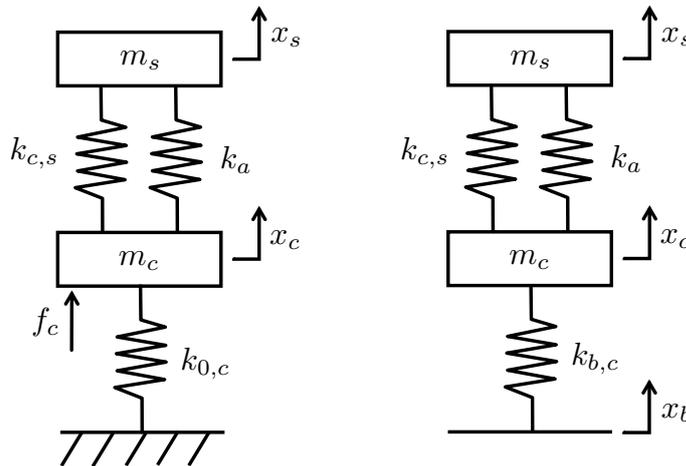


Fig. 4.22 Air-coupled sensor LEM with Thévenin and Norton representations of the chest. x_s and x_c denote the displacement of the sensor rim and the chest surface respectively. (Note that $k_{0,c} = k_{b,c}$ represents a generic Thévenin impedance and that in this representation m_c is not considered to be part of the Thévenin impedance on the chest.)

In figure 4.22 we have assumed that the rim is only connected to the inner chest via the chest tissue at the cavity surface. This makes sense for low frequencies at which a substantial portion of the mass of the chest tissue around the rim and the cavity can be lumped to the cavity-surface degree of freedom (x_c) since the motion of the rim and the cavity are in phase. At higher frequencies, the cavity and rim start to move out of phase and this simple lumping model breaks down. At these frequencies we need to include additional degrees of freedom in our LEM to accurately capture the response. The model can also be adapted to

include a link between the rim and the inner chest that bypasses the chest tissue at the cavity surface. However, we suggest that for a first approximation of the sensor's response the model proposed here should suffice.

4.5.4 Holding the stethoscope

A consideration of how the stethoscope is attached to the chest has largely been an afterthought in research on stethoscope acoustics. Indeed, Ertel (1971) [30] neglected the holding of the stethoscope in his list of the interactions between stethoscopes and human tissue (as discussed in the introduction of this chapter). In existing experimental work the pressure with which the clinician presses a stethoscope against the chest is almost universally replicated by using the weight of a load-mass with the experiment set up in a vertical configuration (see section 6.1.2). The response of any vibration sensor to any system is influenced by the manner in which the sensor is attached to the system (see section 4.1 and e.g. Ewins [31]). This suggests that replacing the hand of the clinician with a simple load mass in a model of the coupled system may be inadequate.

The hand of the clinician plays three roles in the coupled chest-stethoscope-clinician system. Firstly, the pressure with which the clinician presses the stethoscope against the chest sets the equilibrium position of the system around which there are small perturbations as body sounds are transmitted. This role can be represented by a static force $f_s(\omega = 0)$. Secondly, the hand of the clinician is a mechanical system in its own right, and thus presents an impedance between the stethoscope housing and mechanical ground, as shown in figure 4.23. The spring element $k_{s,0}$ represents a complex impedance (i.e., $k_{s,0} = i\omega z_{s,0}^m$) which may consist of any combination of mass, stiffness and damping terms.

The final role of the hand of the clinician is that it can also exert a frequency-dependent input force on the stethoscope. Nelson [105] explores this role in detail, and finds that when a stethoscope is used in an environment with high background vibration (such as a helicopter) the vibration input from the clinician's hand can significantly corrupt the stethoscope's response. However, in more common environments the pressure perturbations exerted by the clinician's hand can be expected to be small compared to the pressure perturbations at the chest of the patient throughout the frequency range of interest for auscultation. A hand

tremor, for example, typically occurs at a frequency below 10 Hz [16] and thus does not interfere with the audible-frequency auscultation signals .

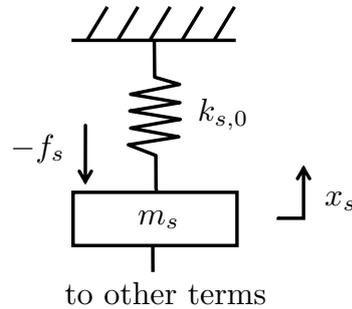


Fig. 4.23 Terms required to include the effect of holding a stethoscope in a LEM for the coupled system. $k_{s,0}$ represents the mechanical impedance of the hand-arm system and $-f_s$ represents the force exerted by the hand. The force f_s includes a static component which sets the equilibrium position of the coupled system, and a frequency-dependent component which acts as a noise source.

To our knowledge, there are no published results for the impedance due to the clinician holding the stethoscope. To address this we present experimental results for the impedance seen by the sensor ‘looking’ towards the hand in section 7.5.

4.5.5 Diaphragm

When a diaphragm (typically a thin disc of a material such as polyester, polystyrene or an epoxy resin-fibreglass laminate [118, 111]) is used on a stethoscope, this diaphragm is in contact with the rim, the air cavity, and the chest surface. For the lumped-element formulation the diaphragm does not add any new degrees of freedom, instead contributing to the existing terms. The coordinate x_c now represents the area-averaged displacement of the proportion of the diaphragm’s surface that is in contact with the air cavity. A proportion of the diaphragm’s mass ($m_{d:c}$) is lumped to this degree of freedom, along with a proportion of the chest’s mass (m_c). The remainder of the diaphragm’s mass ($m_{d:s}$) is lumped to the rim (x_s). The dominant contribution of the diaphragm, however, is the high stiffness $k_{d:c,s}$ between the rim (x_s) and the cavity-surface (x_c).

It is important to emphasise that the presence of the diaphragm does not only give additional mass, stiffness and damping terms for a given equilibrium position, but also significantly alters the equilibrium position itself. For the same static

application force, the increased stiffness means that the chest tissue fills the air cavity to a lesser extent and the rim is displaced into the chest by a smaller distance. This means that, when a diaphragm is added, the magnitudes of the other lumped terms in the model also change.

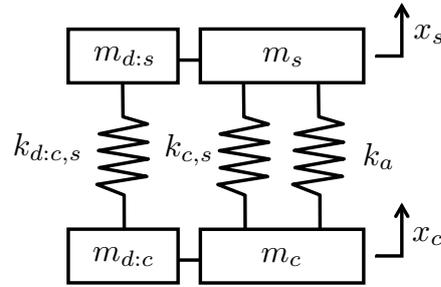


Fig. 4.24 LEM for diaphragm between the chestpiece rim x_s and the cavity surface x_c . The lumped elements are defined in the text.

4.5.6 Tuneable diaphragm

One common argument for the use of diaphragms is that the added stiffness increases the resonance frequencies of the coupled system and thus leads to a larger response at higher frequencies (such as those associated with lung sounds). Stethoscope manufacturers have found that the high stiffness of the diaphragm can make it unsuitable for auscultating low frequency sounds. The traditional solution to this issue is the dual-chestpiece: the user is able to switch between a shallow cavity covered by a diaphragm and a deeper, open cavity.

To reduce the need for switching between the two sides of the chestpiece, stethoscope manufacturers have tried to develop methods by which the effective stiffness of the diaphragm can be altered by changing the application pressure. Cefaly et al. [15] designed a chestpiece in which the diaphragm locks into new positions as the application pressure is increased. This means that as application pressure is increased the cavity-contact area and cavity volume are reduced and effective stiffness of the diaphragm increases.

Packard [118] found that the approach of Cefaly et al. still yielded a diaphragm-stiffness that was too high at low application pressure. Packard proposed a new design in which a stiff diaphragm is held by a compliant ring such that the compliant rim determines the net stiffness at low application pressures. As the application pressure is increased, the diaphragm is pressed against a rigid

inner-rim. The stiff diaphragm is then in direct contact with the sensor-housing, bypassing the compliant ring. In this new equilibrium position the effective stiffness is significantly higher (along with a reduced cavity-contact area and cavity volume). This ‘tuneable diaphragm’ design is now (2021) found in almost all stethoscopes produced by Littmann 3M. While the effect of the tuneable diaphragm of the stiffness between the housing and the cavity-surface is clear the role of the diaphragm as an acoustic filter is disputed (see e.g. [111]).

Figure 4.25 shows lumped element models for the two types of equilibrium position possible with a tuneable diaphragm. Figure 4.25(a) shows the low-application-pressure case in which the response is dominated by the low stiffness of the compliant ring ($k_{r:d,s}$). The stiffness of this compliant ring is in series with the stiffness of the diaphragm. Figure 4.25(b) shows an LEM for the low-application-pressure case. In this case the compliant ring has been bypassed and is no longer included in the LEM. The LEM now resembles that of a simple diaphragm attached directly to the rim (figure 4.24).

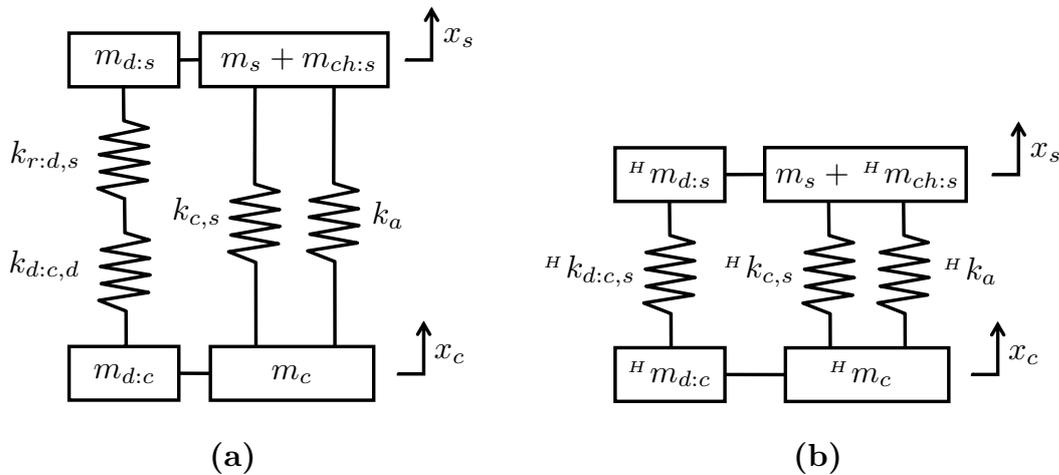


Fig. 4.25 LEM for two states of ‘tuneable diaphragm’. **(a)** small perturbations about a low-static-pressure equilibrium position. **(b)** small perturbations about a high-static-pressure equilibrium position. The values of the all of lumped parameters (except m_s) differ for the two equilibrium positions, as represented by the pre-term superscript H in subfigure **(b)**.

4.5.7 Complete lumped-element model

By combining all of the features described in sections 4.5.1 to 4.5.5 we obtain figure 4.26. We have neglected the microphone from our LEM, as the mass

and compliance of the microphone diaphragm are negligible compared to the surrounding components [163].

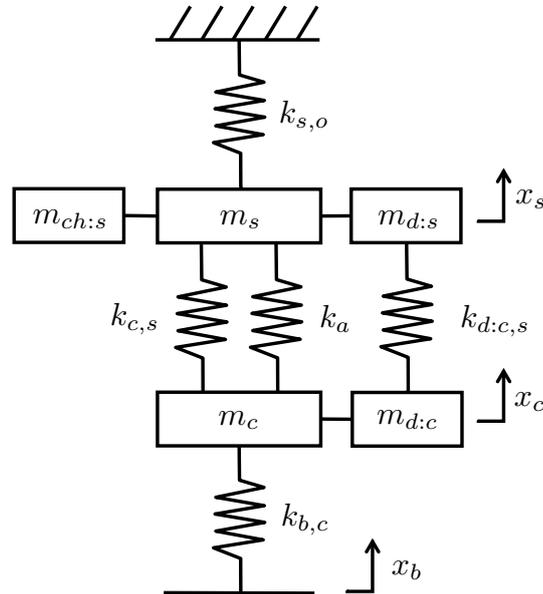


Fig. 4.26 Complete LEM for an air-coupled sensor coupled to the chest.

4.6 Circuit analogies

Figure 4.27 shows the admittance analogy circuit derived from figure 4.26.

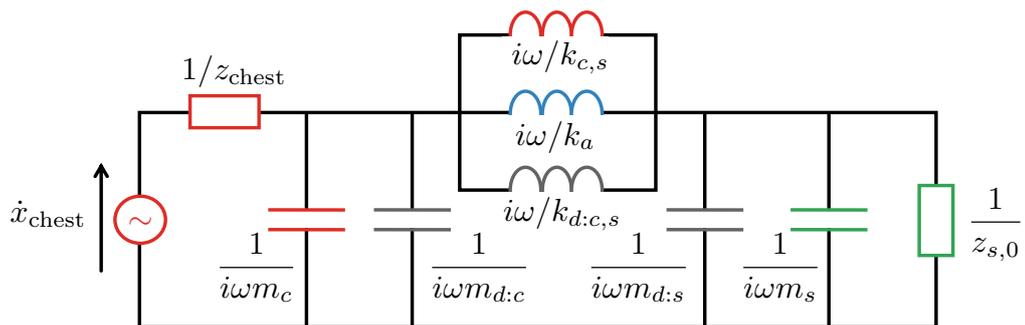


Fig. 4.27 Admittance analogy circuit for the mechanical system in figure 4.26. Red terms correspond to the chest, green terms correspond to the sensor housing and attachment, the blue term is the (acoustic) stiffness of the air cavity, and grey terms correspond to a diaphragm between the chest and the air cavity.

The corresponding impedance analogy circuit (derived by taking the dual of the admittance analogy circuit as described in appendix A.3) is shown in figure 4.28.

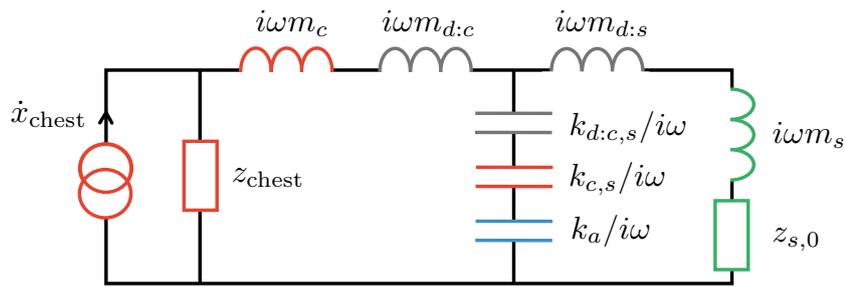


Fig. 4.28 Impedance analogy circuit for the mechanical system in figure 4.26. The colours of components match those in figure 4.27.

Rearranging this into a form amenable to two-port analysis yields figure 4.29.

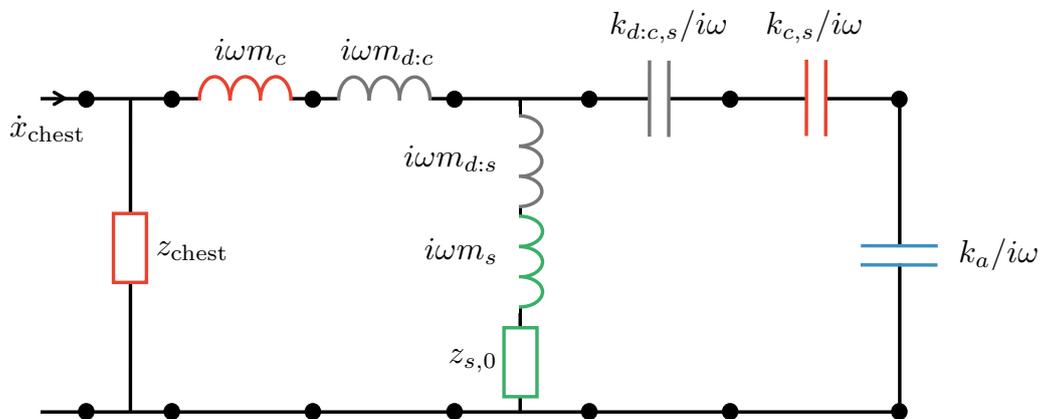


Fig. 4.29 Impedance analogy circuit from figure 4.28 rearranged to allow a two-port matrix formulation between the source and the air-cavity. The colours of components match those in figure 4.27.

The coupling of these circuits to an acoustic domain to model stethoscopes with tubing is discussed in chapter 10.

4.7 Summary

In this chapter we derive lumped element models for direct contact sensors and air-coupled sensors on the chest surface. The model presented for air-coupled sensors addresses some of the key issues with existing models which are discussed in section 4.3. In the models derived here, the value of some components can be directly measured or calculated (the mass of the sensor housing for instance). However, most of the terms need to be estimated from experimental data. In chapter 5 we outline the experimental techniques used to perform these measurements, while in chapter 6 we introduce a laboratory phantom used to represent the chest. In chapter 7 we present impedance measurements on the phantom and a human chest, and in chapter 8 we present a method for converting sensor-response measurements on the phantom into a form directly compatible with the Thévenin equivalent circuits presented in this chapter. In chapters 8 to 10, experiments on this phantom are used to validate the models derived here and to explore their implications.

Chapter 5

Experimental methods

In this chapter we lay out the mathematics underpinning the analysis of our experimental measurements, and provide a detailed model for an impedance head - a sensor used extensively in this thesis. It is essential to understand the structure and behaviour of impedance heads in order to be able to correctly interpret measurements made with these devices [11, 44].

In section 5.1 we lay out the signal processing methods used in this thesis. In section 5.2 we present an overview of the Modal approach to the analysis of vibrating systems and discuss the methods we use to extract modal parameter values from measured response data. Readers familiar with the field of experimental modal analysis may wish to skim this section. In section 5.3 we derive a model for an impedance head, exploring the relationship between the measured response and the true driving-point impedance of the test system.

5.1 Signal processing

In this thesis we present a number of experimental estimates of frequency response functions (henceforth FRFs). The methodology for this is outlined in table 5.1.

Table 5.1 Signal processing methodology.

Step	Details
(1) Identify the variables to be measured, and the parameters to be controlled. Design and validate a test setup.	This chapter outlines the fundamental testing methodologies, with further details on individual experiments given in the following chapters.
(2) Calibrate the transducers to be used.	Standard calibration procedures (pistonphone calibration for microphones and mass-calibration for force and acceleration transducers) are used for all transducers.
(3) Run a test with the desired parameter values, recording the time-domain signals from the transducers.	The data acquisition methods used are described in section 5.1.1.
(4) Convert the raw transducer signals into physical quantities using the calibration results from step (2).	Some transducers require a further frequency calibration after step (5) in accordance with data provided by the manufacturers.
(5) Convert the data into the frequency domain using the (discrete) Fast Fourier Transform (henceforth ‘FFT’).	The mathematics for the discrete Fourier transform are outlined in appendix B.1.
(6) Compute estimates of the (cross-) power spectral densities using Welch’s method.	The mathematics for Welch’s method are outlined in appendix B.2.2.
(7) Check the coherence between channels to assess the quality of the data.	Coherence is defined in section 5.1.2.
(8) Use the (cross) power spectral densities to compute estimates of the FRFs.	Estimators for FRFs are defined in section 5.1.2.

5.1.1 Data acquisition

All data acquisition is performed using a National Instruments (henceforth ‘NI’) cDAQ 9174 chassis, with an NI 9263 output module to generate a desired excitation signal, and several NI 9234 analogue input modules to sample the analogue signals from the transducers. Integrated Electronics Piezo-Electric (IEPE) transducers are used in all cases, with a constant current supplied through the input modules.

For experiments where an artificial continuous excitation source is used, this is fed from the NI 9263 output module through a power amplifier into the relevant exciter (typically a Ling Dynamic Systems V201 shaker). A sampling frequency of 51.2 kHz is used for all experiments.

5.1.2 FRF estimators

Simultaneous data-acquisition allows the FRF between two signals to be estimated. Consider two discrete signals, $s_\phi(n)$ and $s_\psi(n)$, each of length N samples, where n is the sample counter. The discrete Fourier transforms (henceforth ‘DFT’s) of the signals are denoted as $S_\phi(\mu)$ and $S_\psi(\mu)$ respectively, where μ is a counter for the frequency bins. The DFT and its inverse are defined in appendix B.1.

We define the *cross power spectral density* (henceforth ‘CPSD’) between two discrete signals (s_ϕ and s_ψ) as

$$S_{\phi,\psi} = (\Delta t/N^2)S_\phi S_\psi^* , \quad (5.1)$$

where $\Delta t = (N - 1)\delta t$ is the measurement time (δt is the time-step between samples), and $*$ denotes the complex conjugate. The CPSD is complex valued, with a phase that represents the lag between the signals. The power spectral density (henceforth ‘PSD’) is the CPSD evaluated for $\phi = \psi$. The PSD is purely real and is a measure of how the power in the signal is distributed across the frequency domain.

Measured signals always include some random measurement noise, which means that any response functions derived from them are only estimates of the response

functions for the real system. Consider, for example, a mechanical system for which we want to measure the admittance FRF ($H_{\phi,\psi}^v$). The measured force signal $s_\psi(n)$ is equal to the true force f_ψ plus random noise \mathfrak{n}_ψ , while the measured velocity signal $s_\phi(n)$ is equal to the true velocity \dot{x}_ϕ plus random noise \mathfrak{n}_ϕ . The DFTs of the measured signals are thus given by:

$$S_\phi = \dot{X}_\phi + \mathcal{N}_\phi, \quad (5.2)$$

$$S_\psi = F_\psi + \mathcal{N}_\psi, \quad (5.3)$$

where \mathcal{N} , \dot{X} and F are the DFTs of \mathfrak{n} , \dot{x} and f . We define the signal-to-noise ratio (henceforth ‘SNR’) as the power ratio between the true signal and the noise signal:

$$\text{SNR}_\psi = F_\psi F_\psi^* / \mathcal{N}_\psi \mathcal{N}_\psi^*, \quad (5.4)$$

$$\text{SNR}_\phi = \dot{X}_\phi \dot{X}_\phi^* / \mathcal{N}_\phi \mathcal{N}_\phi^*. \quad (5.5)$$

The PSD of a measured signal is given by e.g.,

$$S_{\phi\phi} = (\Delta t/N^2) (\dot{X}_\phi + \mathcal{N}_\phi)(\dot{X}_\phi^* + \mathcal{N}_\phi^*) \quad (5.6)$$

$$= (\Delta t/N^2) (\dot{X}_\phi \dot{X}_\phi^* + \mathcal{N}_\phi \dot{X}_\phi^* + \dot{X}_\phi \mathcal{N}_\phi^* + \mathcal{N}_\phi \mathcal{N}_\phi^*). \quad (5.7)$$

The CPSDs between the two measured signals are:

$$S_{\psi\phi} = (\Delta t/N^2) (F_\psi + \mathcal{N}_\psi)(\dot{X}_\phi^* + \mathcal{N}_\phi^*) \quad (5.8)$$

$$= (\Delta t/N^2) (F_\psi \dot{X}_\phi^* + \mathcal{N}_\psi \dot{X}_\phi^* + F_\psi \mathcal{N}_\phi^* + \mathcal{N}_\psi \mathcal{N}_\phi^*), \quad (5.9)$$

$$S_{\phi\psi} = (\Delta t/N^2) (\dot{X}_\phi F_\psi^* + \mathcal{N}_\phi F_\psi^* + \dot{X}_\phi \mathcal{N}_\psi^* + \mathcal{N}_\phi \mathcal{N}_\psi^*). \quad (5.10)$$

In order to estimate the PSDs and CPSDs from experimental data we use Welch’s method, which is described in appendix B.2.2 and implemented in MATLAB. Welch’s method averages the (C)PSDs from a large number of overlapping signal segments. The measurement noise is random and is therefore uncorrelated with the signal. When we average the CPSDs of a large number of uncorrelated terms

the value tends to zero. Our estimates for the PSDs and CPSDs thus become:

$$\tilde{S}_{\phi\phi} \simeq (\Delta t/N^2) (\dot{X}_\phi \dot{X}_\phi^* + \mathcal{N}_\phi \mathcal{N}_\phi^*), \quad (5.11)$$

$$\tilde{S}_{\psi\psi} \simeq (\Delta t/N^2) (F_\psi F_\psi^* + \mathcal{N}_\psi \mathcal{N}_\psi^*), \quad (5.12)$$

$$\tilde{S}_{\psi\phi} \simeq (\Delta t/N^2) (F_\psi \dot{X}_\phi^*), \quad (5.13)$$

$$\tilde{S}_{\phi\psi} \simeq (\Delta t/N^2) (\dot{X}_\phi F_\psi^*). \quad (5.14)$$

The admittance response function is given by $H_{\phi,\psi}^v = \dot{X}_\phi/F_\psi$. We define the H1 estimator (1H) and the H2 estimator (2H) for the FRF, making use of the averaged (C)PSD estimates:

$${}^1H_{\phi,\psi} = \frac{\tilde{S}_{\phi,\psi}}{\tilde{S}_{\psi,\psi}} \simeq \frac{\dot{X}_\phi F_\psi^*}{F_\psi F_\psi^* + \mathcal{N}_\psi \mathcal{N}_\psi^*}, \quad (5.15)$$

$${}^2H_{\phi,\psi} = \frac{\tilde{S}_{\phi,\phi}}{\tilde{S}_{\psi,\phi}} \simeq \frac{\dot{X}_\phi \dot{X}_\phi^* + \mathcal{N}_\phi \mathcal{N}_\phi^*}{F_\psi \dot{X}_\phi^*}. \quad (5.16)$$

By inspection of equation 5.15, and noting that $\mathcal{N}_\psi \mathcal{N}_\psi^*$ is positive by definition, it can be seen that the H1 estimator always underestimates the FRF. Likewise, inspection of equation 5.16 reveals that the H2 estimator always overestimates the FRF. This means that the H1 and H2 estimators provide upper and lower bounds on the true value of the FRF: ${}^1H < H < {}^2H$. The H1 estimator works well if the SNR on the *input* signal s_ψ is high, while the H2 estimator works well if the SNR on the *output* signal s_ϕ is high.

For a noise-less system ${}^1H = {}^2H = H$ (provided the lag is negligible compared to the FFT size). Any deviation between the H1 and H2 estimators is therefore an indication of noise in the system. We define the ratio of the H1 and H2 estimators as the coherence function γ^2 :

$$\gamma_{\psi\phi}^2(\mu) = \frac{{}^1H_{\psi\phi}(\mu)}{{}^2H_{\psi\phi}(\mu)}, \quad (5.17)$$

$$= \frac{\tilde{S}_{\phi,\psi} \tilde{S}_{\psi,\phi}}{\tilde{S}_{\psi,\psi} \tilde{S}_{\phi,\phi}}. \quad (5.18)$$

The coherence has a value between 0 and 1, with 1 meaning that the two estimators are equal and there are no external noise sources in the system. It

should be noted that the coherence value depends on the signal processing techniques used as well as on the raw measurement signals. For example, for two identical random signals with only a time lag, the coherence is zero if the FFT size is smaller than the time lag between the signals. We ensure that, wherever a coherence is given, the FFT size used in the analysis is large compared to the time lag between the signals.

When analysing experimental response data, the choice of estimator can be critical. This is especially true around resonances and antiresonances of the system. In general, the H1 estimator gives a better estimate around the troughs, while the H2 estimator gives a better estimate around the peaks of the response. The coherence function can be used to identify for which parts of the FRF the choice of estimator is important. Appendix C.1 covers these considerations in more detail.

5.2 Modal approach

Ewins [31] identifies three types of models for mechanical systems: *spatial*, *response* and *modal* models. The lumped element method, described in section 3.3, is an example of a *spatial* model, as it describes the physical properties of the system and their spatial relationship. A *response* model is a description of the system's response to excitation and can be measured experimentally. The driving point response plots in section 3.3 are examples of *response* properties derived from *spatial* properties. For simple systems we can also work in reverse and deduce *spatial* properties from measured *response* properties, but for more complex systems this becomes difficult.

Modal analysis is a powerful tool which allows us to express a system's behaviour as a summation of vibration modes. Each mode has an associated natural frequency and a modeshape. A key advantage of modal analysis is that *modal* properties can be deduced from *response* properties even for complex systems with many modes and high damping. Theoretical *spatial* models can also be converted into *modal* models, allowing direct comparisons between theory and experimental data. In this thesis, modal analysis is used extensively in the interpretation of experimental data in relation to the performance of stethoscopes.

5.2.1 Modal decomposition

An N-degree-of-freedom discrete mechanical system is characterised by N modes of vibration. The natural frequency of the n th mode is denoted by ω_n . For a discrete system, the modeshape of the n th mode is defined at a finite number of points in space and becomes a vector $\mathbf{u}^{(n)}$. The value of the n th modeshape at point ϕ is given by $u_\phi^{(n)}$. The general (harmonic) response of the system can be expressed as a sum of modal responses:

$$\mathbf{x}(t) = \sum_{n=1}^N \alpha_n \mathbf{u}^{(n)} e^{i\omega t}. \quad (5.19)$$

In equation 5.19 α_n is the amplitude and $\mathbf{u}^{(n)}$ is the normalised modeshape for the n th mode. For mathematical convenience, the modeshapes are normalised such that

$$\mathbf{u}^{(n)t} [M] \mathbf{u}^{(n)} = 1, \quad (5.20)$$

$$\mathbf{u}^{(n)t} [K] \mathbf{u}^{(n)} = \omega_n^2. \quad (5.21)$$

where $[M]$ is the mass matrix and $[K]$ is the stiffness matrix, superscript t denotes the transpose. These matrices are defined such that the undamped governing equations can be expressed in the form

$$[M]\ddot{\mathbf{x}} + [K]\mathbf{x} = \mathbf{f}. \quad (5.22)$$

Assuming harmonic excitation $\mathbf{f} = \mathbf{F}e^{i\omega t}$ and response $\mathbf{x} = \mathbf{X}e^{i\omega t}$, and further assuming that excitation is only applied at the ψ th coordinate, we can express the (undamped) FRF relating the response at point ϕ to excitation at point ψ as

$$H_x(\phi, \psi, \omega) = \sum_{n=1}^N \frac{u_k^{(n)} u_j^{(n)}}{\omega_n^2 - \omega^2}. \quad (5.23)$$

A derivation of this result can be found in appendix A.4.1. We define the ‘modal amplitude factor’ (\hat{a}_n) as

$$\hat{a}_n(\phi, \psi) = u_\phi^{(n)} u_\psi^{(n)}. \quad (5.24)$$

The sign of the modal amplitude factor determines the sign of the response at either side of the resonance. This means that if the modal amplitude factor

has the same sign for two adjacent modes then there must be a zero-crossing (antiresonance) between the resonances. The ‘driving point response’ is defined as the response when we observe and excite the same point. For example, the driving point compliance FRF is given by

$$H_x(\psi, \psi, \omega) = \sum_{n=1}^N \frac{u_\psi^{(n)} u_\psi^{(n)}}{\omega_n^2 - \omega^2} . \quad (5.25)$$

The driving point modal amplitude factor $\hat{a}_n(\psi, \psi)$ is, by definition, always positive. This means that there is a zero-crossing (anti-resonance) between each resonance in the driving point FRF.

5.2.2 Damping

Damping is the most complicated aspect of modelling vibrating systems due to the fact that it depends on a wide variety of physical mechanisms [31]. There are two simple approaches to modelling damping that are commonly used in modal analysis: viscous damping and structural (hysteresis) damping [31]. See e.g. (Maia and Silva, 1997) [87] or (Ewins, 2000) [31] for detailed discussions.

In a lumped-element model viscous damping is represented using dashpots (see figure 3.3). In a modal model we account for viscous damping by introducing a ‘modal damping ratio’ for each mode. Defining c_n as the dimensionless modal damping ratio the response function (equation (A.24)) becomes:

$$H_x(\phi, \psi, \omega) = \sum_{n=1}^N \frac{\hat{a}_n(\phi, \psi)}{\omega_n^2 + 2ic_n\omega\omega_n - \omega^2} . \quad (5.26)$$

For a single degree-of-freedom oscillator with mass m , stiffness k and viscous damping λ , such as that in figure 3.5, the relationship between the viscous damping factor and the dimensionless modal damping ratio is given by

$$c_n = \frac{\lambda}{2\sqrt{km}} . \quad (5.27)$$

The ‘quality factor’ $Q_n = 1/2c_n$ is an alternative parameter used to describe the modal damping. The relationship between the FRFs is $H_v = i\omega H_x$ and $H_a = i\omega H_v$, so evaluating the FRFs at resonance for a single degree of freedom

system with $a_n = 1/m_1$ gives:

$$H_x(\omega_n) = \frac{1/m_1}{2ic_n\omega_n^2} = \frac{1/m_1}{2i\frac{\lambda_{1,0}}{2\sqrt{m_1k_{1,0}}} \frac{k_{1,0}}{m_1}} = \frac{-i}{\lambda_{1,0}\omega_n} = \frac{-i}{2c_nk_{1,0}}, \quad (5.28)$$

$$H_v(\omega_n) = i\omega H_x(\omega_n) = \frac{1}{\lambda_{1,0}} = \frac{1}{2c_nm_1\omega_n} = \frac{Q_n a_n}{\omega_n}, \quad (5.29)$$

$$H_a(\omega_n) = i\omega H_v(\omega_n) = \omega^2 H_x(\omega_n) = \frac{i\omega_n}{\lambda_{1,0}} = \frac{i}{2c_nm_1} = iQ_n a_n. \quad (5.30)$$

It can be shown (see appendix A.4.2) that for a single mode the Nyquist plot of the admittance FRF forms a circle. The red dots in figure 5.1 show an example of this. For the compliance and accelerance response functions the Nyquist plot approximates a circle when the damping is low (see e.g., figure 5.1(a)), but deviates from a circle when the damping is increased, as shown in figure 5.1(b).

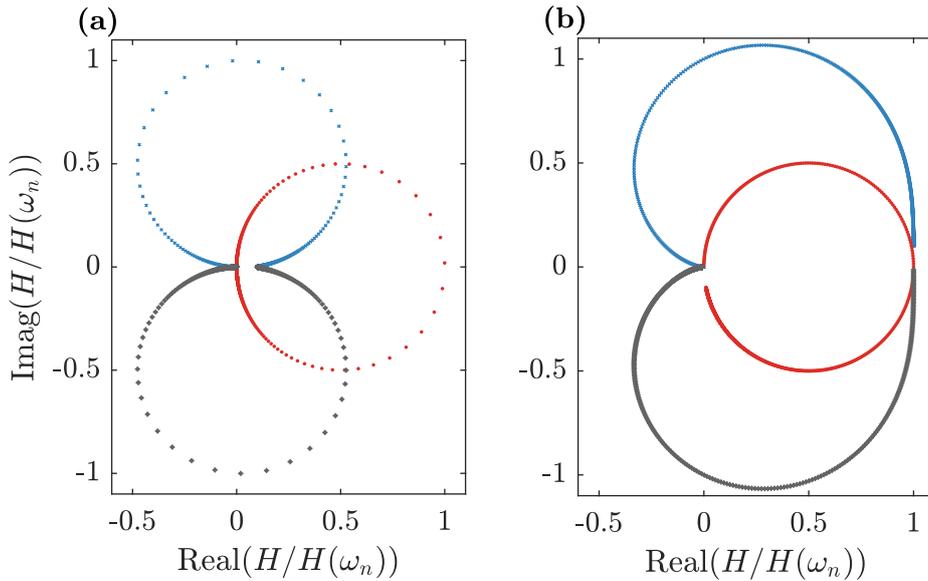


Fig. 5.1 Nyquist plots of normalised FRFs for a SDOF system with (a) $c_n = 0.05$, (b) $c_n = 0.5$. \times - H_x , \bullet - H_v , $*$ - H_a . Increasing frequency is clockwise in each case. The modal resonance coincides with the lowest point density.

5.2.3 Experimental modal analysis

The aim of experimental modal analysis is to extract the values of the three modal parameters in equation (5.26) from a measured *response model*. These modal parameters are the natural frequency ω_n , the modal amplitude factor

a_n , and a measure of the damping. For a modal model with viscous damping (e.g., equation (5.26)) the modal damping ratio c_n or the modal quality factor $Q_n = 1/2c_n$ are used.

For each individual mode, the response is at a maximum at the resonance frequency ω_n and dies away as the frequency deviates from ω_n . The sharpness of the peak is determined by the damping level. *Modal overlap* is a measure of the closeness of adjacent modes. If the modal overlap is low around mode n then the response at frequencies close to ω_n is dominated by the response of that mode, while the (*residual*) response of other modes is negligible. If modal overlap is high then the residual response of nearby modes can significantly affect the response as shown in figure 5.2.

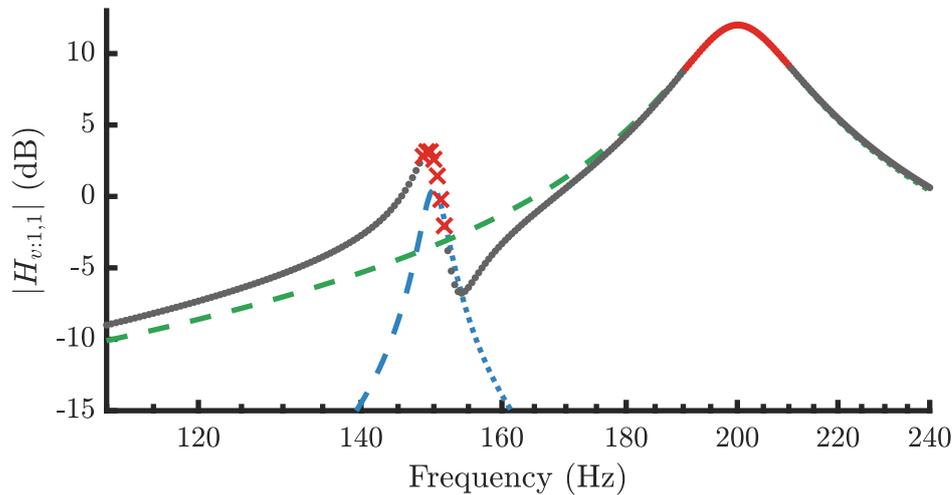


Fig. 5.2 Bode magnitude plot of the driving point admittance frequency response function for a simulated 2DOF system with significant modal overlap. First mode ($\omega_{n=1} = 150$ Hz) coloured in blue, second mode ($\omega_{n=2} = 200$ Hz) in green. In both cases regions of the FRF with positive real part are drawn as a dashed line and regions with a negative real part are drawn as a dotted line. Individual data-points of the summed system response by grey dots (\bullet). Data-points of the summed response corresponding to the half-power bandwidths of the underlying modal responses are shown using red crosses (\times).

5.2.4 Bode plot method for modal parameter extraction

The simplest method for extracting modal information from a measured FRF is to find the frequency corresponding to a peak in magnitude (ω_{peak}), the value of the response at that peak ($|H_{\text{peak}}|$) and the half-power bandwidth ($\delta\omega_{3\text{dB}}$). The

half-power bandwidth is the frequency interval between points on either sides of the peak that have magnitudes 3 dB below the peak value. For low damping we take $\omega_n \simeq \omega_{\max}$ and $Q_n \simeq \omega_n / \delta\omega_{3\text{dB},n}$. The modal amplitude factor can then be calculated by substituting $|H_{\text{peak}}|$ into equations (5.28) to (5.30).

For low modal overlap and a high frequency resolution the Bode-plot method outlined above gives reasonable results. However, if the modal overlap is high or the frequency resolution is low then the method fails to accurately capture the modal parameters. This can be seen in figure 5.2 in which the peak of the response curve for the lower-frequency mode does not provide an accurate representation of the underlying mode.

5.2.5 Nyquist plot method for modal parameter extraction

Figure 5.3 shows a Nyquist plot corresponding to the Bode plot in figure 5.2. The modal overlap has led to distortion of the modal circles and has shifted their centres away from the x -axis. The centre of the circle for the secondary mode is offset from the real axis and the peak response magnitude (maximum distance from the origin) does not occur at resonance.

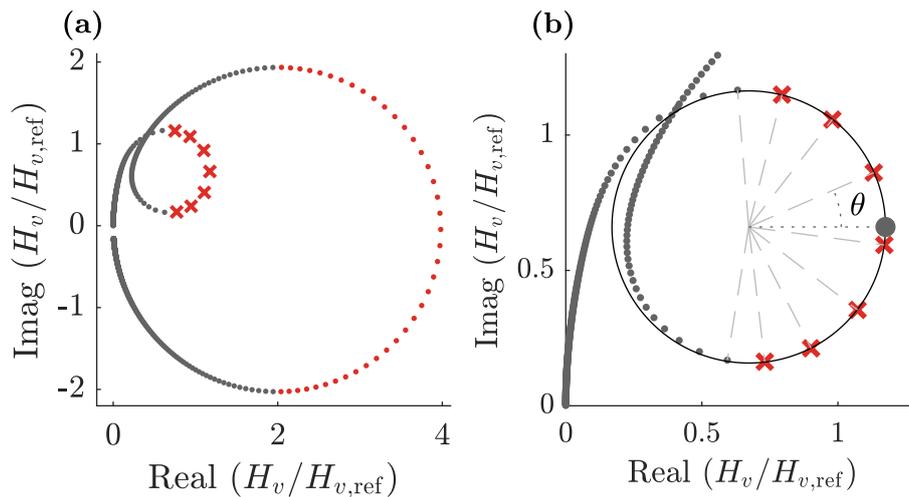


Fig. 5.3 (a) Nyquist plot corresponding to the Bode magnitude plot in figure 5.2. The datapoints shown are uniformly spaced in frequency. The red crosses and dots match those in figure 5.2. (b) Close-up of the behaviour around the low frequency resonance. θ is defined as the arc angle anti-clockwise from horizontal, and can be computed using equation 5.32.

In order to extract modal parameters we fit a circle to the data-points in the Nyquist plot (using e.g., CircleFitPrat in MATLAB). The diameter of the circle (\varnothing_n) can be expressed as:

$$\varnothing_n = \frac{Q_n a_n}{\omega_n}, \quad (5.31)$$

where a_n is the modal amplitude factor, ω_n is the resonant frequency for the mode and $Q_n = 1/2c_n$ is the modal quality factor. We define the arc angle (shown in figure 5.3(b)) as:

$$\theta = \angle(H_v(\omega) - H_c) \quad (5.32)$$

where H_c is the locus of the centre of the fitted circle, and \angle is used to denote the phase angle.

If the damping is proportional (see e.g., [31] or [87]) then ω_n can be found by interpolating to $\theta = 0$, and the half-power-points occurs at $\theta = \pm 90$. There is, however, no reason to assume that the damping in a real structure is proportional. We therefore instead make use of the fact that the angle between adjacent datapoints is largest around resonance (as seen in figure 5.3). This can be used as an interpolation condition to find ω_n . The Nyquist-plot method laid out in this section is used for all modal-parameter extraction in this thesis.

5.3 ‘Impedance head’ for driving-point FRF measurement

The impedance-coupling approaches discussed in this thesis rely on the ability to experimentally measure the driving-point response of the human chest. For the driving point response ‘seen’ by a rigid contact there are three aspects to this: forming a rigid contact with the chest, applying a known excitation force and measuring the resulting motion of the rigid contact.

We form a rigid contact by loading the chest with a metal cylinder of the appropriate mass and dimensions. The excitation force can be applied either by using impulsive excitation with an impulse hammer, or by using continuous excitation with a shaker. For systems with significant damping, it is beneficial to excite the system continuously, using a shaker. The voltage applied to the sensor cannot be directly converted into an excitation force, as the relationship between

shaker voltage and force exerted depends on the load [31]. This means that we need to measure the force exerted on the metal cylinder using a force transducer.

A device designed to simultaneously measure excitation force and driving-point acceleration is called an ‘impedance head’. One end of the impedance head is mounted directly onto the test structure, while the other is attached to the shaker via a ‘stinger’. A stinger is a thin rod that transfers axial forces from the shaker to the impedance head while limiting the transmission of transverse excitation. A stinger is required because unmeasured transverse forces can lead to errors in response measurements, particularly for compliant structures like the human chest [87].

Provided the magnitude of the shaker excitation significantly exceeds that of any other sources of excitation within the system (such as background noise and vibration) the impedance head theoretically gives a measure of the Norton / Thévenin equivalent impedance of the system ‘seen’ at the driving point.

In practice, the force measured by the impedance head is not equal to the force applied to the structure, and the acceleration measured by the impedance head is not equal to the acceleration of the driving point [11] [44]. In order to correctly interpret the data from impedance head measurements, these discrepancies must be understood. In section 5.3.1 we present a LEM for an impedance head which allows us to explain the discrepancies and their effect on the measured response.

5.3.1 Impedance head model

The impedance head is a compact device typically made mainly of titanium, with ceramic piezoelectric elements. The longitudinal wave speed in piezoelectric ceramics is typically over 4 km/s [114]. This means that, in the frequency range of interest (<3 kHz), the dimensions of the impedance head are negligible compared to the wavelengths and a lumped element model is appropriate. The titanium elements of the impedance head have very high stiffness and can thus be modelled as point masses. The piezoelectric elements have comparatively low stiffness and much lower mass and can be modelled as springs (with their mass lumped to either end of the spring).

Figure 5.4 shows a lumped element model for an impedance head. The base of the impedance head (mass m_b) is connected to the excitation source. A piezoelectric

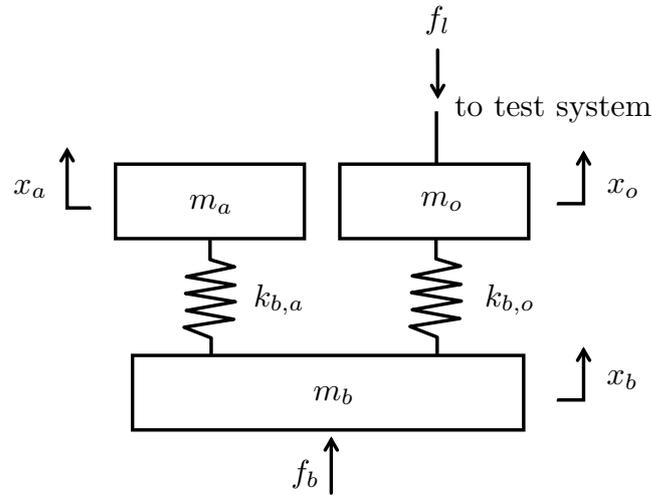


Fig. 5.4 Lumped element model of impedance head. f_l is the force exerted on the test system.

crystal (of lumped stiffness $k_{b,o}$) connects the base to the coupling mass (m_o) which is connected to the test system. The impedance head also includes an accelerometer. A seismic mass (m_a) is connected to the base mass (m_b) by a second piezoelectric crystal (of stiffness $k_{b,a}$). The matrix equation governing system is therefore given by:

$$\left(-\omega^2 \begin{bmatrix} m_b & 0 & 0 \\ 0 & m_a & 0 \\ 0 & 0 & m_o \end{bmatrix} + \begin{bmatrix} k_{b,a} + k_{b,o} & -k_{b,a} & -k_{b,o} \\ -k_{b,a} & k_{b,a} & 0 \\ -k_{b,o} & 0 & k_{b,o} \end{bmatrix} \right) \begin{bmatrix} X_b \\ X_a \\ X_o \end{bmatrix} = \begin{bmatrix} F_b \\ 0 \\ -F_l \end{bmatrix}, \quad (5.33)$$

where X and F are the frequency domain representations of displacement and force.

The deformation of the piezoelectric crystal ($x_o - x_b$) gives a voltage output that is used to estimate the force applied through the crystal ($f_{k_{b,o}} = k_{b,o}(x_o - x_b)$). This is not a direct measurement of the force applied to the load structure due to the finite mass m_o of the impedance head above the force transducer. For loads with low dynamic mass (typical for biological systems [87]) m_o can be significant compared to the effective mass of the load system. A correction for the force is given by the third row of equation (5.33):

$$-f_l = f_{k_{b,o}} + m_o \ddot{x}_o. \quad (5.34)$$

Considering the accelerometer within the impedance head, the first row in equation (5.33) gives:

$$(x_a - x_b) = \frac{-m_a}{k_{b,a}} \ddot{x}_a . \quad (5.35)$$

The piezoelectric crystal of the accelerometer gives a voltage output proportional to $(x_a - x_b)$ which is thus proportional to the acceleration of the seismic mass (\ddot{x}_a) at all frequencies. Equation (5.35) can be rearranged to give equation (5.36).

$$X_b = \left(1 - \frac{\omega^2 m_a}{k_{b,a}}\right) X_a = \left(1 - \left(\frac{\omega}{\omega_a}\right)^2\right) X_a . \quad (5.36)$$

Equation (5.36) shows that \dot{x}_a will match \dot{x}_b well provided $\omega \ll \omega_a$. This means that the accelerometer gives an accurate measurement of the base acceleration of the sensor at frequencies far below the internal resonance at $\omega_a = \sqrt{(k_{b,a}/m_a)}$. However, the base acceleration of the accelerometer (\dot{x}_b) is not equal to the acceleration of the load (\dot{x}_l).

Figure 5.5 shows a mechanical-impedance-analogy circuit for the impedance head. Thévenin’s theorem is applied to simplify the circuit, replacing terms associated with the shaker and the stinger that are not relevant to the measurement with an equivalent input force ($f_{\text{Thé}}$) and a source impedance ($z_{\text{Thé}}$). The test-structure is represented by an arbitrary load impedance z_l .

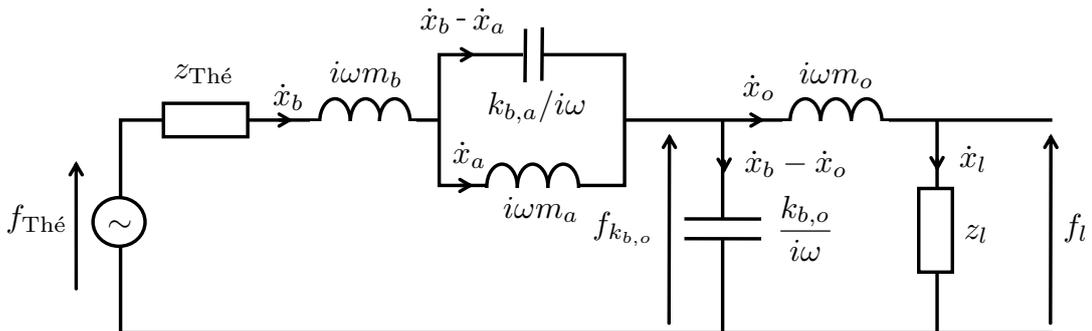


Fig. 5.5 Thévenin’s theorem applied to impedance-analogy circuit of an impedance head on a shaker, with a load impedance of z_l .

From figure 5.5 it is evident that the load impedance plays a role in determining the relationship between the measured terms (proportional to $x_b - x_a$ and $x_b - x_o$), and the terms of interest (f_l and \dot{x}_l). This means that there is no simple correction factor that can be applied to account for the bias error. In order to correctly

interpret the measurements results we must therefore make use of our LEM for the impedance head.

5.3.2 Model validation

In this thesis we present results measured using a PCB 288D01 Impedance Head. To validate our model for the impedance head and establish values for the lumped elements we use the impedance head to measure the driving point response of a set of rigid (steel and aluminium) load masses. This process can be interpreted as a ‘mass-calibration’ of the sensor (see e.g. [10]). The experimental setup for the mass-calibration is shown in figure 5.6. A range of rigid load masses were bolted to the top of the impedance head, and the system was excited with bandpass-filtered white noise for 100 s.

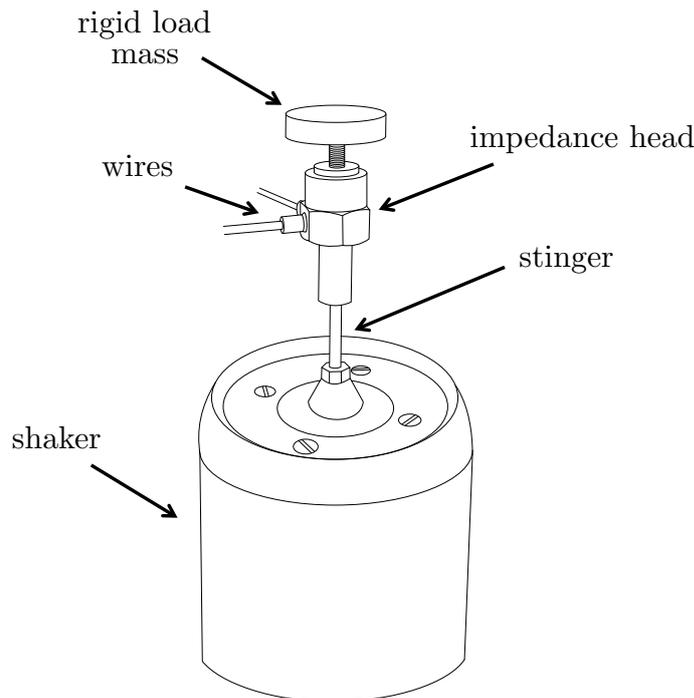


Fig. 5.6 Experimental setup for rigid mass impedance head tests.

The dynamic mass FRF, calculated using the calibration factors specified by the manufacturer, is shown for three representative cases in figure 5.7. The response is characterised by a low frequency plateau in dynamic mass, followed by an antiresonance. Our lumped element model for the impedance head suggests that

this antiresonance arises due to the mounting stiffness between the impedance head and the load. This is validated in section 5.3.3.

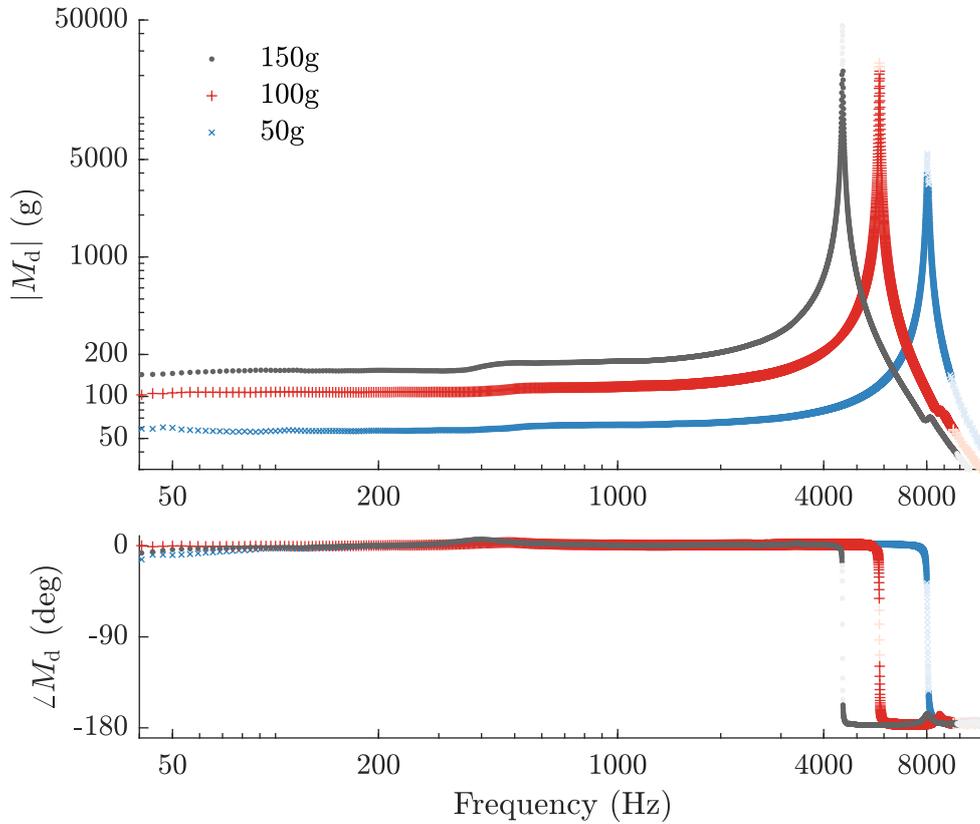


Fig. 5.7 H1 estimator for the driving point dynamic mass (M_d) measured by a PCB 288D01 impedance head using the setup shown in figure 5.6 for rigid mass loads of 50 g (blue cross \times), 100 g (red plus $+$) and 150 g (grey dot \bullet). Frequency bins with a coherence of less than 0.98 are plotted in light colour.

When we extract the value of the dynamic mass plateau from the measured FRFs for a number of known load masses, we can establish the value of the ‘mass-above-the-force-gauge’. The ‘mass-above-the-force-gauge’ is equal to the magnitude of the dynamic mass plateau when no load is applied to the impedance head ($m_o \simeq 6$ g). The manufacturers specify the ‘internal resonance’ of the accelerometer within the PCB 288D01 Impedance Head to be $\omega_{\text{accel}} \simeq 38$ kHz. This means that for the frequencies far below this resonance we can treat the $k_{b,a}/i\omega$ term in the impedance-analogy circuit in figure 5.5 as an open circuit and let $\ddot{x}_b \simeq \ddot{x}_a$. We can then simplify the impedance-analogy circuit as shown in figure 5.8.

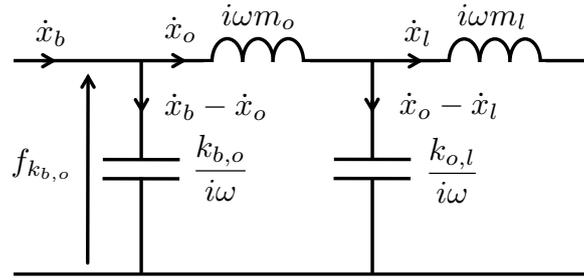


Fig. 5.8 Right hand section of the impedance-analogy circuit in figure 5.5, with the generic load impedance Z_{load} replaced by a mounting stiffness $k_{o,l}$ and a load mass m_l . For $\omega \ll \omega_{\text{accel}}$ we can approximate $\ddot{x}_b \simeq \ddot{x}_a$, which allows us to use this circuit-segment to model the FRF measured by an impedance head.

The net stiffness of the impedance head is specified by the manufacturers as ($k_{b,o} \simeq 3.5 \times 10^8$ N/m). The only remaining unknown in the model is the mounting stiffness.

5.3.3 Mounting stiffness

In this section we explore the dominant antiresonance in the driving point FRF that can be seen in figure 5.7. An antiresonance in a driving point response measurement corresponds to a natural frequency of the theoretical system that is obtained if the driven degree of freedom of the actual system is rigidly linked to mechanical ground (see e.g. [87]). This means that the frequency of the observed antiresonance is given by the lower natural frequency of:

$$\left(-\omega^2 \begin{bmatrix} m_o & 0 \\ 0 & m_l \end{bmatrix} + \begin{bmatrix} k_{b,o} + k_{o,l} & -k_{o,l} \\ -k_{o,l} & k_{o,l} \end{bmatrix} \right) \begin{bmatrix} X_o \\ X_l \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (5.37)$$

Equation 5.37 can be rearranged for the unknown mounting stiffness, giving

$$k_{o,l} = \frac{\omega_{-1}^2 m_l (k_{b,o} - \omega_{-1}^2 m_o)}{k_{b,o} - \omega_{-1}^2 (m_o + m_l)}. \quad (5.38)$$

The unknown mounting stiffness (which varies with the experimental setup) has serious repercussions for the reproducibility of these results. The location of the dominant antiresonance cannot be precisely controlled or predicted. When comparing results from tests between which the equipment has been disassembled and

reassembled we are limited to comparing results below any of the antiresonances (unless the mounting stiffness is considered part of the measurement).

To illustrate this, figure 5.9 shows two measured driving point dynamic mass curves from nominally the identical experiments. The only difference between the two cases is how firmly the rigid load mass is attached to the impedance head.

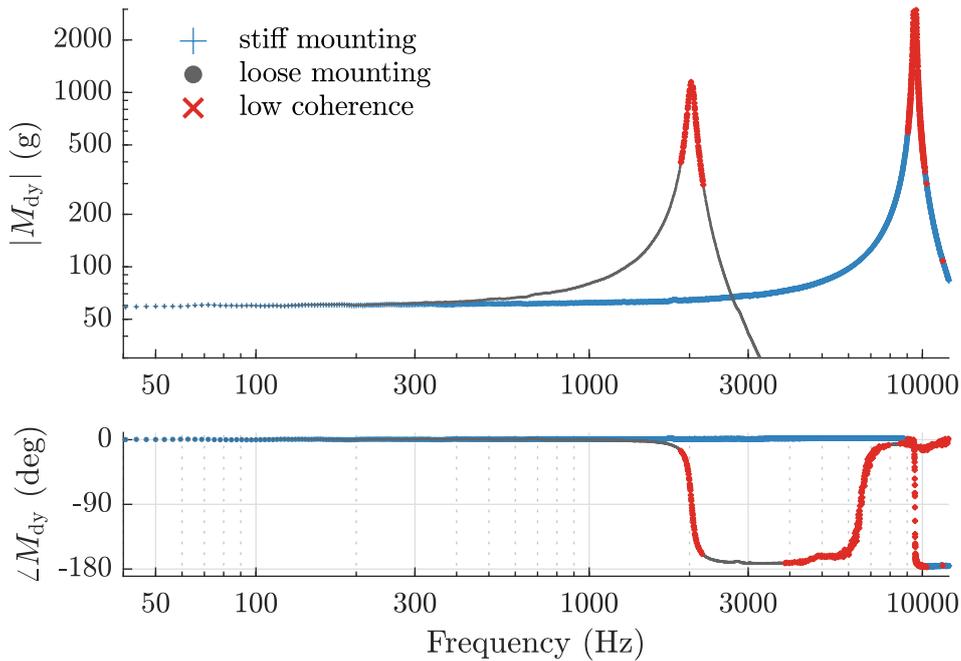


Fig. 5.9 H2 estimator of the driving point dynamic mass using an impedance head with a 51.0 g rigid load mass attached firmly (+) and loosely (•) to the impedance head. Frequency bins in which the coherence drops below 0.98 are given by red crosses (×).

In figure 5.9 the antiresonance frequency for the loose case is around 2025 Hz, while for the stiff case it is around 9500 Hz. Using equation (5.38) we can estimate the mounting stiffnesses as $8.5 \times 10^6 \text{ N m}^{-1}$ and $4.1 \times 10^8 \text{ N m}^{-1}$ respectively.

5.3.4 Simulated impedance head FRF measurement

Using the LEM parameter values derived in section 5.3.2, computing the mounting stiffness using equation (5.38) and applying appropriate modal damping factors the FRF measured by the PCB 288D01 Impedance Head can be simulated. The simulated response is superimposed onto the measured response for three

representative load masses in figure 5.10. The simulated response matches the measured response well in the frequency range up to and including the ‘mounting-antiresonance’. At frequencies beyond those shown in figure 5.10 the assumption that the impedance corresponding to the accelerometer stiffness ($k_{b,a}/i\omega$) is infinite begins to break down and the lumped elements corresponding to the accelerometer have to be included in the model. However, at these high frequencies unmeasured transverse forces begin to play a more significant role, which means that the measured force signal is no longer a valid description of the excitation of the system [31]. This means that any attempt to model this frequency region accurately is superfluous.

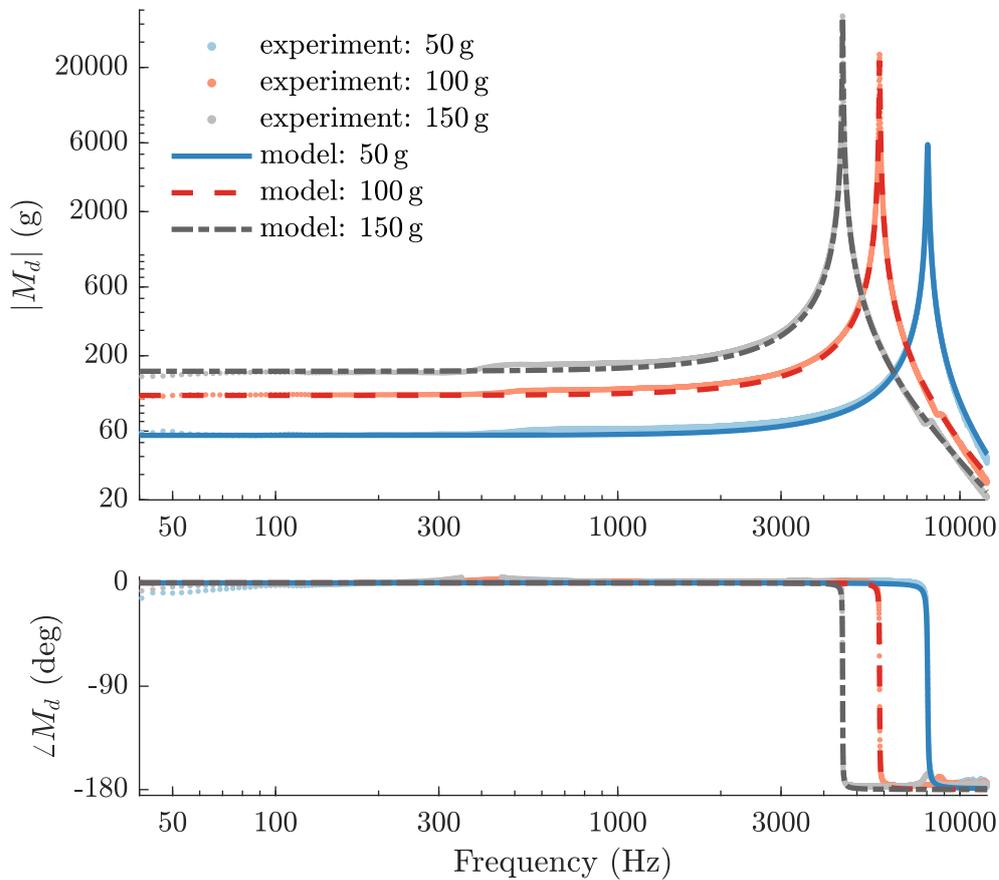


Fig. 5.10 Modelled impedance head dynamic mass FRF measurements for load masses of 50, 100 and 150 g superimposed on experimental measurements. Parameter values for the model are as discussed in sections 5.3.2 - 5.3.4. Damping values for the models are selected to match the damping observed in the measured FRFs.

We have shown that the FRFs measured using an impedance head deviate from the true driving-point FRFs of the load. The dominant deviations are an offset of

$i\omega m_o$ across all frequencies and an antiresonance associated with the mounting stiffness that dominates the response above a certain frequency.

As the mounting stiffness cannot be accurately predicted before the experiments are performed it is not possible to model the expected response without any reference to experimental data. However, the model does inform us about how to accurately interpret the FRFs measured using an impedance head. Furthermore, the model can be used to simulate changes made to the load system which do not involve altering the mounting stiffness between the impedance head and the load.

5.4 Summary

In this chapter we have outlined the fundamental techniques underlying the analysis of experimental data in this thesis. Individual experimental setups are described in the following chapters. All of these make use of the techniques we have outlined in this chapter. Experimental estimates of frequency response functions form the backbone of the results presented in this thesis. We make extensive use of an impedance head to measure both driving-point impedances and the excitation-input for our chest-phantom. The lumped element model for the impedance head presented in this chapter is used to make sense of this data.

Chapter 6

Phantoms

The human chest is a highly complex system. No two chests are the same. Significant variations are found in geometry, anatomy and material properties. These variations can be genetic, or related to age, health and physique. Adiposity (excessive body fat) in particular can have a significant effect on the propagation of vibrations through the human chest.

This degree of variability does not lend itself well to experimental reproducibility. Furthermore, the inherent variations between human chests can cloud the view of what observations on stethoscope performance are generalisable. In order to develop a model for the acoustics of the stethoscope we first require a simplified system that:

- can be represented by a simple spatial model,
- can be built and tested in a laboratory environment,
- has a clearly defined and measurable input signal,
- can be reproduced by other researchers.

As noted in section 2.3.2, a rig designed to mimic a part of the human body is often referred to as a ‘phantom’. A wide range of ‘chest phantom’ designs have been proposed in the literature and are discussed in detail in section 6.1. In section 6.2 we describe the design of our own chest-phantom, addressing some of the shortcomings in existing approaches. In section 6.3 we present a set of experiments used to characterise the behaviour of the phantom. In section 6.4 we derive a simple LEM to describe the low frequency behaviour of the phantom,

and in section 6.5 we compute a transmission matrix to describe the transmission of vibration through the phantom, allowing us to model the response of sensors over a more extended frequency range. Finally, in section 6.6 we demonstrate the robustness of our results with respect to small deviations in the positioning of the sensor on the phantom.

6.1 Phantoms in the literature

Phantoms designed to test and compare stethoscopes have been constructed using a wide variety of materials to mimic the properties of the human chest. Many early studies used water-filled structures. Wodicka et al. [163] employed a water-filled sponge sealed within a latex pouch, Warous et al. [157] (and later Pavlosky et al. [123]) utilised a water-filled ball, while Welsby et al. used a water-filled blood pressure cuff [161]. A more common technique in recent years has been to use viscoelastic elastomers to replicate the mechanical properties of the chest surface. Tables 6.1 and 6.2 give an overview of the materials used for phantoms that have been described in recent literature (2006 - 2020). Mansy et al. [96] performed a detailed investigation into the material properties of various synthetic materials to assess their suitability for soft-tissue modelling as well as their stability over time. A range of materials were found to be suitable, which explains the extensive variety observed in the literature.

The variability in phantoms can lead to difficulties in comparing results. In 2006 Kraman (et al.) set out to address this issue by designing a well documented system that could be replicated by other research groups [77]. This ‘BioAcoustic Transducer Testing’ (henceforth ‘BATT’) system consisted of a square sheet of Akton™ placed above an air cavity, with a loudspeaker in the cavity used to provide an excitation signal (see table 6.1). Kraman et al. compared the use of Akton™ with meat and observed reasonable agreement. The BATT system has been used to compare stethoscopes [79], study the effect of clothing on lung sound recording [76], and investigate the sensitivity of sensors to airborne and tissue-borne signals [171]. There has, however, been limited adoption by the wider research community. Recently (2020), Joyashiki and Wada published a comparison of three different stethoscope types in which they used a ‘respiratory sound simulator’ based on the work of Kraman et al. [77], but with several design modifications [66] (see table 6.2).

Table 6.1 Geometry, dimensions and materials used for selected phantom designs from 2006 - 2015.

Author(s)	Year	Geometry	Dimensions	Material
Dai	2006 [19]	cylinder	320 mm diameter, 150 mm height	CF 11 silicone gel (Nusil Technology, Carpinteria, CA)
Kraman et al.	2006 [77]	square sheet of material placed above an air cavity	80 mm square, 7 mm height	Akton™
Mansy et al.	2009 [94]	cylinder with rigid side and base walls, and a rigid plate embedded 10 mm below the surface	100 mm diameter, 40 mm height	Semcosil 921 (Wacker Solutions, Adreian, MI)
Hirahara et al.	2010 [51]	cylinder with rigid base and free sides	75 mm diameter, 50 mm height	urethane elastomer (not specified further)
Nelson *1	2015 [105]	cylinder with rigid base and sides	not specified	Shore 10A silicone elastomer covered with a thin polymer membrane (not specified further)
Bakhshaei et al.	2015 [4]	cylinder with rigid base and top and free sides	not specified	Ecoflex 00-10

*1 this thesis also describes the phantom in [106].

Table 6.2 Geometry, dimensions and materials used for selected phantom designs from 2018 - 2020.

Author(s)	Year	Geometry	Dimensions	Material
Joyashiki and Wada	2018 [65]	cylinder with a rigid base and free sides	75 mm diameter, 50 mm height	urethane elastomer (not specified further)
Weiss et al.	2019 [158]	rectangular geometry, with several stacked layers within a plastic container	cross sectional dimensions not specified, 5 mm of gel on top of phantom	ballistics gel (Perma-Gel Personal Ballistic Gel Test Kit; Perma-Gel, Inc.; Albany, OR, USA), castor oil and open cell foam
Klum et al.	2019 [73]	cylinder with rigid side walls and an embedded PVC structure	160 mm diameter, 90 mm height	gelatine
Keller (patent)	2020 [71]	cylinder on a rigid base, with a 87 mm diameter hole in the base leading to an air cavity, and free sides	130 mm diameter, 30 mm height	Ecoflex 00-10
Aygün and Apolskis	2020 [3]	cylinder with a rigid base and free sides	diameter not specified, 10 mm height	a porous 'open cell material fabricated from particles of plastic foam obtained from recycled car dashboards' (not specified further)
Joyashiki and Wada* ¹	2020 [66]	(geometry not specified) elastomer on a rigid base with multiple small (2mm) holes in the base leading to an air chamber	not specified	urethane elastomer (not specified further)

*¹ this paper also describes the phantom in [65].

Having investigated some of the materials available for phantom construction in detail [96], Mansy et al. also published a well documented phantom design intended for adoption by the wider research community [94, 95]. To this end, Mansy et al. published step-by-step instructions for the construction of the phantom (see [94]). Mansy et al.'s phantom consists of a cylinder of viscoelastic material, surrounded by a rigid base and walls, with a speaker embedded within the cylinder.

Despite the fact that the exact design proposed by Mansy et al. has not been adopted in the wider research community, it is worth noting that almost all of the reported phantom designs since 2009 have featured cylinders of synthetic viscoelastic material. In almost all cases the base is constrained, with the sides either left free [51, 65, 71, 3, 66] or also constrained [106, 105, 73]. As these phantoms are designed to replicate the coupling of sensors to the chest, but not the propagation of sound through the chest, they do not typically contain components designed to mimic specific anatomical features such as the bones or lungs. One notable exception is the phantom of Weiss et al. (2019) [158] which features several layers of different materials stacked inside a plastic container.

6.1.1 Excitation

Phantoms are typically excited by using either a loudspeaker or an electrodynamic shaker. If a loudspeaker is used, this is typically coupled to an air-cavity which is in contact with the phantom. If a shaker is employed, then this is typically used to vibrate a rigid element in contact with the phantom. In both cases, the excitation source can be either at the base of the phantom or embedded within the phantom material. Table 6.3 divides the phantoms described in several recent studies into these four categories.

Table 6.3 Excitation sources for transducer-testing-phantoms described in the literature.

	speaker	shaker
at base	(Wodicka, 1994) [163] (Suzuki, 1998 & 2000) [142, 143] (Kraman, 2006) [77] (Keller, 2020) [71] (Joyashiki, 2020) [66]	(Watrous, 2002) [157] (Welsby, 2003) [161] (Hirahara, 2010) [51] (Nelson, 2015) [106] (Joyashiki, 2018 & 2020) [65, 66] (Aygün, 2020) [3]
embedded	(Mansy, 2006) [94]	(Dai, 2006) [19] (Nelson, 2015) [105] (Weiss, 2019) [158] (Klum, 2019) [73]

Most of the studies we surveyed used a (white, pink or bandpass-filtered) noise signal as the excitation input [163, 143, 157, 77, 94, 106, 105, 158, 73, 71, 3, 66]. Joyashiki and Wada also played tracheal sounds (normal breathing sounds recorded at the trachea) and wheezes (abnormal lung sounds) through their phantom for demonstrative purposes [66]. Mansy et al. [94] and Klum et al. [73] used ‘spectral equalisation’ techniques to achieve an approximately flat spectrum at the *unloaded* phantom surface. As the spectrum will be modified when the surface is loaded the advantage of this procedure is limited. A wide range of frequency bands have been used for the excitation applied to phantoms, with several researchers only reporting results above 100 Hz (see e.g. [163], [77] and [51]) despite the fact that most of the energy in heart sounds is associated with frequencies below 100 Hz.

One notable exception to the use of speakers or shakers to produce the excitation signal is the work of Bakhshae et al. [4], who introduce vibrations into their cylindrical phantom using fluid flow through a constricted pipe running along the central axis of the phantom. While this may be more representative of the way certain body sounds are generated, it does not lend itself to the testing and comparison of stethoscopes as the input signal is hard to measure and control.

6.1.2 Loading

When conducting phantom-tests, one important issue to consider is how to apply the stethoscope to the chest surface. The most common approaches are sticking the sensor on with double sided tape (see e.g. [163, 79]), letting the sensor rest on top of the phantom with just its self-weight (see e.g. [161, 106]), or applying load masses to the sensor in an attempt to replicate the pressure applied by the physician (see e.g. [71, 106, 73]). In section 4.5.4 we suggest that a load-mass may not be a good model for the loading from a clinician's hand. The rigid mass provides a static force load (due to the vertical orientation of the phantom-setup), but it also provides an impedance load that affects the response of the coupled system. In clinical use, the static force with which the stethoscope is applied to the chest is not simply equal to the mass of the stethoscope chestpiece multiplied by 9.81 ms^{-2} . Furthermore, the impedance load from the clinician's hand may differ significantly from a pure mass, as is confirmed by experimental results shown in sections 7.5.4, 8.5 and 9.8. None of the studies reviewed in this thesis have taken this issue into account.

Unusually, Bakhshae et al. (2015) [4] attach stethoscopes to the curved side of their cylindrical phantom rather than to the flat top. The position of the stethoscopes is controlled using micro-manipulators, with a load-cell to measure the application pressure. Bakhshae et al. do not measure the impedance load produced by this method of attachment, which makes it hard to replicate and interpret their work.

6.1.3 Reference signal

When using a phantom to test or calibrate stethoscopes there needs to be a reference signal against which the stethoscope output can be compared. If the phantom is used purely for comparing sensors then it may be possible to not use a reference signal, but there still needs to be some way of ensuring that the input is kept constant.

Table 6.4 summarises the reference signals used for phantoms in a selection of recent studies. The table differentiates between studies in which no reference signal is used at all and studies in which the reference is 'not measured during tests'. In both cases researchers typically set the excitation voltage to the same

level in all tests and assume that this creates identical input conditions. This is not the case. Ewins [31] shows that the force output from a shaker depends not only on the voltage signal but also on the impedance of the load to which the shaker is attached, and on the shaker mounting. For studies in which no reference signal is measured at all, the resulting spectra are only valid for the particular phantom used, negating the benefits of using a phantom instead of performing measurements directly on the human chest.

Table 6.4 Examples of frequency ranges for stethoscope-phantom test results published in the literature.

Author(s)	Year	Citation	Reference signal
Wodicka et al.	1994	[163]	none
Suzuki and Nakayama	2000	[143]	microphone in the air cavity
Watrous et al.	2002	[157]	not measured during tests
Welsby	2003	[161]	none
Dai	2006	[19]	acceleration from impedance head at shaker side of stinger *1
Kraman et al.	2006	[77]	microphone in the air cavity
Mansy et al.	2009	[94]	not measured during tests
Hirahara et al.	2010	[51]	accelerometer on the phantom surface next to sensor being tested
Nelson et al.	2015	[106]	force transducer between the shaker and the phantom
Nelson	2015	[105]	accelerometer on the rigid base for ‘background vibration’, none for ‘patient noise’
Joyashiki and Wado	2018	[65]	accelerometer on the phantom surface next to the sensor being tested
Weiss et al.	2019	[158]	none*2
Klum et al.	2019	[73]	accelerometer on the phantom surface next to the sensor being tested
Keller (patent)	2020	[71]	microphone in the air cavity
Aygun and Apolskis	2020	[3]	none
Joyashiki and Wada	2020	[66]	microphone in the air cavity

*1 Positioning an impedance head on the shaker side of a stinger rather than the load side can lead to poor quality measurements (see e.g. [31]).

*2 Weiss et al. claim that they use “calibrated audio levels to ensure that each stethoscope received the same acoustic energy at the auscultation point” but provide no further information on how they do (or define) this.

Studies in which the reference is ‘not measured during the tests’ typically measure a reference signal at the unloaded phantom surface when subjected to the same nominal excitation signal that is used to test stethoscopes (see e.g. [157, 94]). As already noted there is no guarantee that applying the same electrical input signal to the shaker or loudspeaker leads to the same physical input excitation when different loads are applied to the system. Furthermore, the fact that the reference signal is not recorded simultaneously to the excitation signal precludes the computation of FRFs.

The vibration-characteristics of phantoms have rarely been fully investigated. Mansy et al. [95] found that there was a high degree of spatial uniformity in the unloaded surface motion over the central 40 mm of their 100 mm diameter phantom. However, Klum et al. [73] note that applying the stethoscope to the phantom disturbs this uniformity. Welsby et al. [161] report that their phantom “was prone to resonate at certain frequencies” but neither they nor any other published studies have, to our knowledge, undertaken a detailed experimental modal analysis of their phantom. We propose that it is not the unloaded behaviour of the phantom that should be characterised, but rather its behaviour when subjected to loading by the sensors being tested.

In this chapter we address the issues that have been raised in this literature review. We characterise the behaviour of our phantom with respect to different loads applied to the surface using a method that allows for non-linearity with respect to the loads applied. We measure the performance of sensors with respect to the force and acceleration at the base of the phantom, and use our characterisation of the phantom’s behaviour to interpret these response functions.

The response functions measured on a phantom *always* depend on the properties of the phantom. However, if we understand *how* the response functions depend on the properties of the phantom then we can use this to interpret how the same sensors would behave on the human chest. This is explored for the case of a rigid direct contact sensor in chapter 8.

6.2 Phantom design

To design a phantom for testing stethoscopes we returned to the four requirements laid out in the introduction to this chapter. The single most important element

that can be distilled from these four requirements is simplicity. The simpler the design, the easier it is to build, model and reproduce the phantom, and the easier it is to clearly define a suitable input signal.

To allow for a simple spatial model of the system, we designed the phantom to consist of only two core parts: a rigid base layer and a viscoelastic top layer, with vertical excitation. We opted for a cylindrical design, reducing the three-dimensional complexity, and also simplifying the construction. The entire phantom is freely suspended on the excitation mechanism to yield a single clearly defined input and to limit the number of unknown boundary conditions. The rigid base layer can be modelled as a single-degree-of-freedom system. The advantage of this is that it allows us to fully define the input excitation in terms of just the force on, and the velocity of, this base layer. Consequently, a shaker (rather than a speaker) is the natural choice for an excitation mechanism. An impedance head can be used between the shaker and the rigid base layer to measure the input signal during tests. The rigid base layer also allows the input excitation to be applied to the viscoelastic layer in a fairly uniform manner.

To ensure repeatability, we used readily available equipment and materials, with the main layer of the phantom constructed from a commercial viscoelastomer that has been widely used in similar work. The dimensions for the phantom were chosen by establishing a compromise between a number of requirements, while making use of widely available laboratory parts to aid reproducibility. The radius of the viscoelastic cylinder needs to be large enough to allow a range of stethoscope chestpieces to be mounted at its centre without being too close to the edge, but small enough to prevent excessive transverse motion and forces arising from small imbalances. A 150 mm diameter satisfies these requirements, and allows the phantom to be moulded using a standard 150 mm inner-diameter pipe segment. The thickness of the viscoelastic layer needs to be thin enough to not unduly attenuate the input signal, especially at higher frequencies. However, using a thinner layer increases the stiffness, which should be kept within a range that is physically plausible for modelling the human chest. A thickness of 20mm was chosen as a reasonable compromise.

This remainder of this section outlines the detailed design of the phantom used for experiments to characterise stethoscope performance in this thesis (henceforth ‘the phantom’). Sufficient detail is provided for the phantom to be replicated, to ensure the reproducibility of the results presented in this thesis.

6.2.1 Materials

The phantom consists primarily of a cylinder composed of the commercial silicone elastomer Ecoflex™ 00-10, Smooth-On Inc. Ecoflex™ 00-10 is a type of polysiloxane [168] which is easily fabricated by mixing two components and curing at room temperature [4]. Ecoflex™ 00-10 was identified by Mansy et al. as a suitable phantom material, and has been used successfully for stethoscope-testing-phantoms by e.g. Keller [71] and Bakhshae et al. [4].

Ecoflex™ has the advantage (versus other common phantom materials such as agar gel and gelatin) that its mechanical properties vary less over time, and that it is able to adhere strongly to the rigid surfaces of the phantom rig [168]. Bakhshae et al. [4] argue that *“its robustness makes it a perfect candidate for making stable phantoms”*.

The material properties of Ecoflex 00-10 have been extensively explored (see e.g. [19, 96, 34, 14, 9, 151]). A wide range of properties are reported in the literature, which may in part be due to differences in production. The exact ratio of the two components mixed together as well as moisture and air-entrapment can affect the material properties of the final product [96].

6.2.2 Geometry

The design for the phantom is shown in figure 6.1. A 75 mm diameter, 20 mm high cylinder of silicone elastomer (Ecoflex™00-10) is cast onto a 90 mm diameter, 2.5 mm thickness aluminium disc. This disc is mounted on an impedance head as shown in figure 6.2. The sides of the silicone elastomer cylinder are unconstrained in the radial direction. The total mass of the phantom is 140.3 ± 0.1 g (measured using a Sartorius Model 2255 precision balance).

In the manner adopted by e.g., Nelson et al. (2015) [106] and Joyashiki and Wada (2018) [65] we mount the entire phantom on a shaker so that the number of unknown reaction forces is minimised. Unlike Nelson et al. [106] we use an impedance head at the base of the phantom to measure both the input force and the driving point acceleration. The presence of unknown transverse forces transmitted through the impedance head at the base cannot be completely eliminated, but these are minimised by using a stinger and by ensuring a symmetrical setup [31].

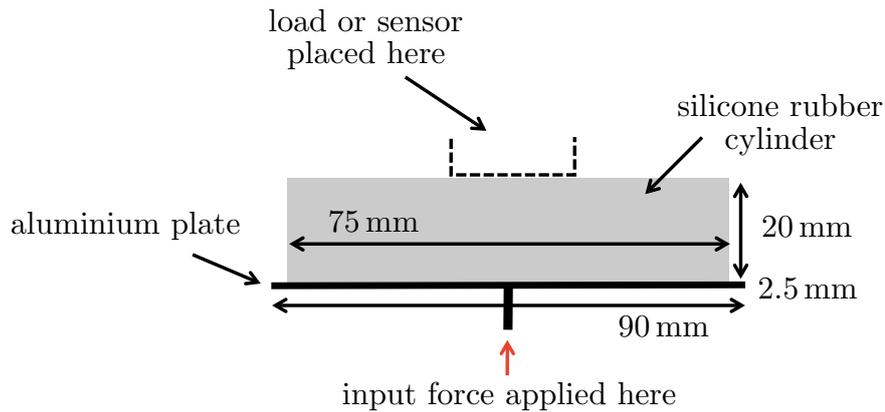


Fig. 6.1 Schematic of the phantom design.

6.2.3 Construction

The following methodology was developed for the phantom casting:

1. Assemble the phantom mould using an aluminium base plate with a tapered bolt attached to one face and a segment of perspex pipe with the appropriate diameter glued to the free face. The perspex pipe should be coaxial with the aluminium plate and the tapered bolt.
2. Line the inside of the perspex tube with paper and spray with a quick release agent.
3. Mix Ecoflex™ 00-10 with a ratio of A:B = 1:1 by volume in a container.
4. Mix vigorously to ensure a homogeneous distribution.
5. Place the Ecoflex™ 00-10 container in a vacuum chamber and degass.
6. Slowly pour the Ecoflex™ 00-10 mixture into the mould.
7. Place the mould with the Ecoflex™ 00-10 into the vacuum chamber and perform a second degassing.
8. Leave to cure for at least four hours (cure time specified by manufacturer).
9. Carefully remove the perspex pipe and the paper to leave the Ecoflex™ 00-10 cylinder unconstrained in the radial direction.

The degassing steps reduce the entrainment of air bubbles in the Ecoflex™. Air bubbles affect the mechanical properties of the phantom and thus the reproducibility of the results.

6.2.4 Excitation

To provide an excitation signal the phantom is mounted on top of a shaker (Ling Dynamic Systems V201). A voltage signal is generated in MATLAB and fed to the shaker as described in section 5.1.1. To measure the physical input signal an impedance head (see section 5.3) is mounted to the base of the phantom. The impedance head is connected to the shaker via a stinger. The fact that the weight of the phantom is supported entirely through the stinger means a fairly stiff stinger needs to be used. This reduces the extent to which the stinger is able to prevent unmeasured transverse forces (see section 5.3 for a discussion of the issues associated with this).

As the phantom directly loads the shaker without any additional external supports, the shaker armature is displaced from the centre of the coil which can cause some signal distortion [31]. Since a measured force is used as the input reference (rather than the voltage signal applied to the shaker) this is not a critical issue. The distortion mainly occurs when the voltage signal is first applied to the shaker, at which point there can be a ‘jump’ followed by a short period of transient behaviour. To account for this we truncate the first portion of the signal in the time domain.

6.3 Phantom characterisation

In section 6.1.3 we argue that there is not much to be gained by analysing the vibration behaviour of the phantom in isolation. Instead, we are interested in how the phantom behaves when it is coupled to a sensor. This section explores the vibration behaviour of the phantom described in section 6.2 when different load masses are placed in the stethoscope-testing-position on top of the phantom.

Figure 6.2 shows the experimental setup used to investigate the effect of sensor mass on the driving-point response at the base of the phantom. The phantom is loaded with either only a metal cylinder (‘rigid load mass’), or a metal cylinder

with an accelerometer (PCB Piezotronics M353B15) attached. The latter can be interpreted as a direct contact sensor. In this chapter FRFs between the acceleration of the direct sensor and the input force are utilised to explore the properties of the phantom. In chapter 8 the same experimental setup is used to explore the performance of direct contact sensors.

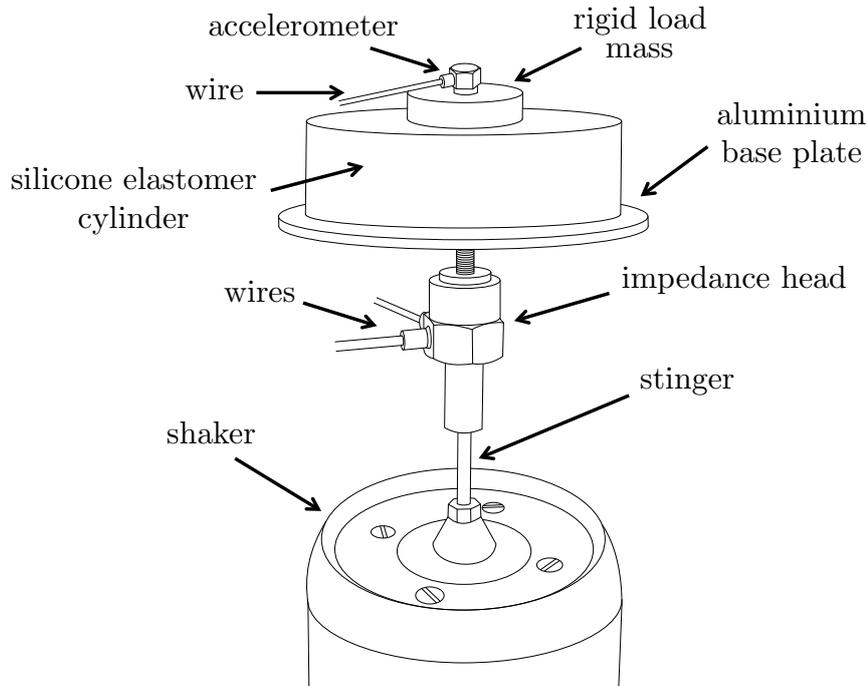


Fig. 6.2 Experimental setup for driving-point FRF measurements at the base of the phantom, with rigid load masses applied at the centre of the top surface of the phantom. An accelerometer is included for a subset of the measurements to allow a transfer FRF to be computed (see section 6.3.3 and chapter 8).

6.3.1 Results - effect of load mass on driving point FRFs

Figure 6.3 shows the driving-point admittance FRF measured at the base of the phantom when rigid load masses are applied to the top surface as shown in figure 6.2. For these measurements no accelerometer is used on the load mass. Results are shown for load masses of 50, 100 and 150 g.

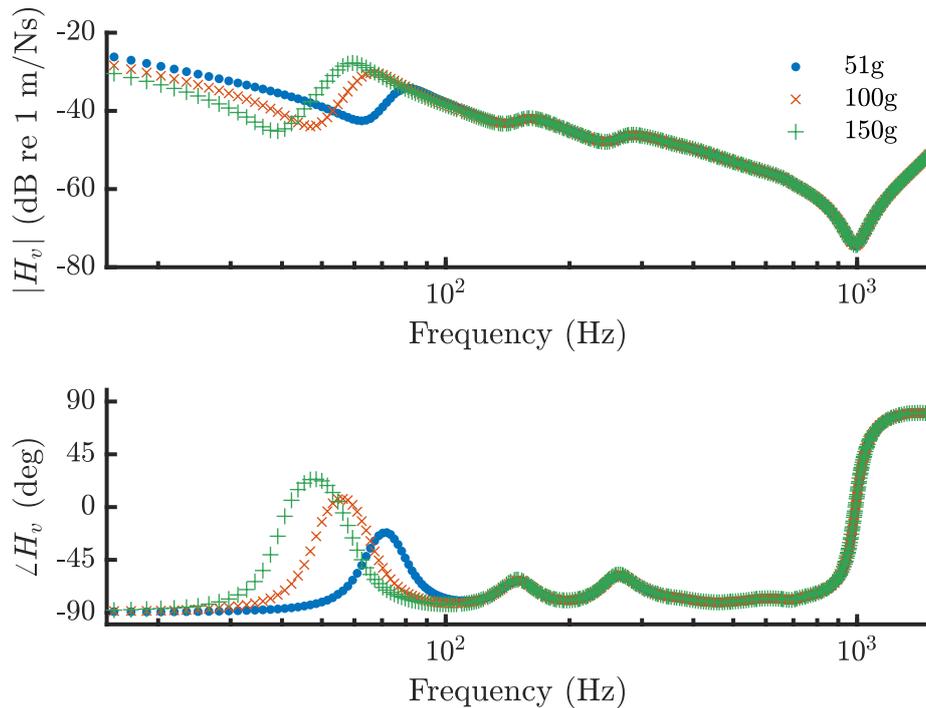


Fig. 6.3 Driving point admittance FRF measurements for the phantom loaded with 25.4 mm diameter rigid cylindrical load masses of 51 g (\bullet), 100 g (\times) and 150 g ($+$).

In all three cases the first resonance frequency of the system occurs at $\omega_1/2\pi = 0$ Hz (not directly shown on the logarithmic frequency scale), which agrees with what we expect for a ‘freely-supported’ system with no ground links (see e.g. [31]). The phase of -90 degrees tells us that the response is mass dominated at low frequency (see section 3.4 and e.g. [10]). For each case there are three clear resonances between the 0 Hz resonance and a dominant antiresonance at ~ 1 kHz.

For each load mass case, an antiresonance can be observed between each pair of resonances. This is always the case for a driving-point response and can be used as a quality control step [31]. The coherence is high (> 0.99) throughout the frequency range shown. The modes between 100 Hz and 1 kHz are quite highly damped (evident from the low peak-prominence, and the shallow phase gradient).

Comparing the admittance FRF curves for the three different load masses, it is evident that beyond ~ 100 Hz the FRFs collapse onto a single curve. This suggests

that the load mass has ‘locked’ and is no longer contributing the impedance seen at the base.

Figure 6.4 shows the resonance and antiresonance frequencies for the driving point response at the base of the phantom when this is loaded with 25.4 mm diameter loads with a range of masses. It can be seen that the trends observed in figure 6.3 hold true: only ω_{-1} and ω_2 exhibit a noticeable dependence on load mass.

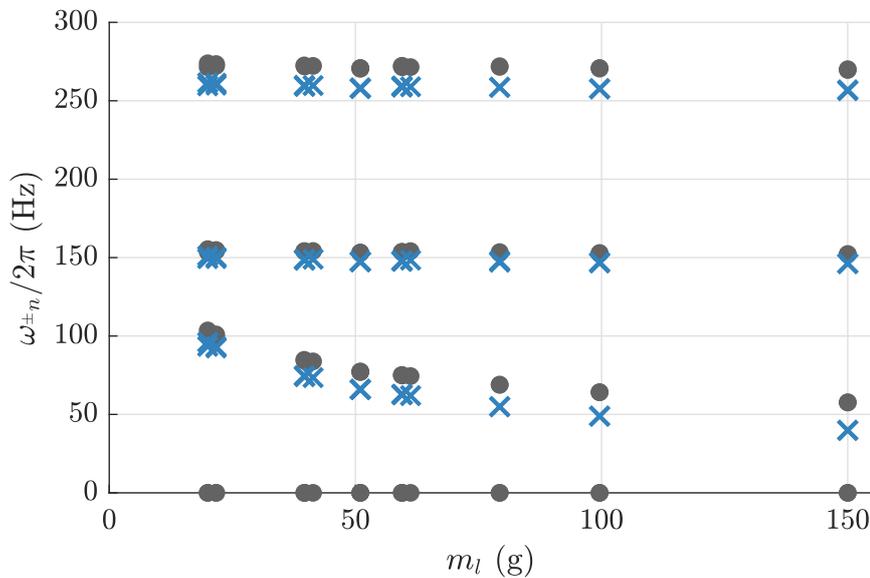


Fig. 6.4 Resonance (●) and antiresonance (×) frequencies extracted from driving point phantom response data for various load masses of diameter 25.4 mm placed centrally on the top surface of the phantom. The resonance and antiresonance frequencies are extracted by circle fitting to the Nyquist plot of the H2 estimator of the admittance and impedance FRFs respectively as discussed in section 5.2.5.

The trend in the resonances matches what we would expect from the *interlacing theorem*: the resonances of a coupled system interlace the resonances of the two subsystems that are coupled together [87]. In this case, every time mass is added to the load we are coupling the existing system with a system that has a resonance at $\omega = 0 \text{ rad s}^{-1}$. The interlacing theorem suggests that the ω_2 resonance should interlace the previous (lower load mass) ω_2 resonance and 0 rad s^{-1} . This can be observed in figure 6.4. The trend also matches the prediction made by Rayleigh’s theorem: when we add mass to the system the resonance frequencies decrease [87].

We have identified that the load mass is ‘locking’ from the perspective of the driving point at the base, and that the response is predominately mass-dominated in the frequency range of interest. This suggests that the dynamic mass FRF can yield more insight into the system’s response. Figure 6.5 shows the driving point dynamic mass FRF for 25.4 mm diameter load masses ranging from 50 - 150 g, as well as the response of the unloaded phantom.

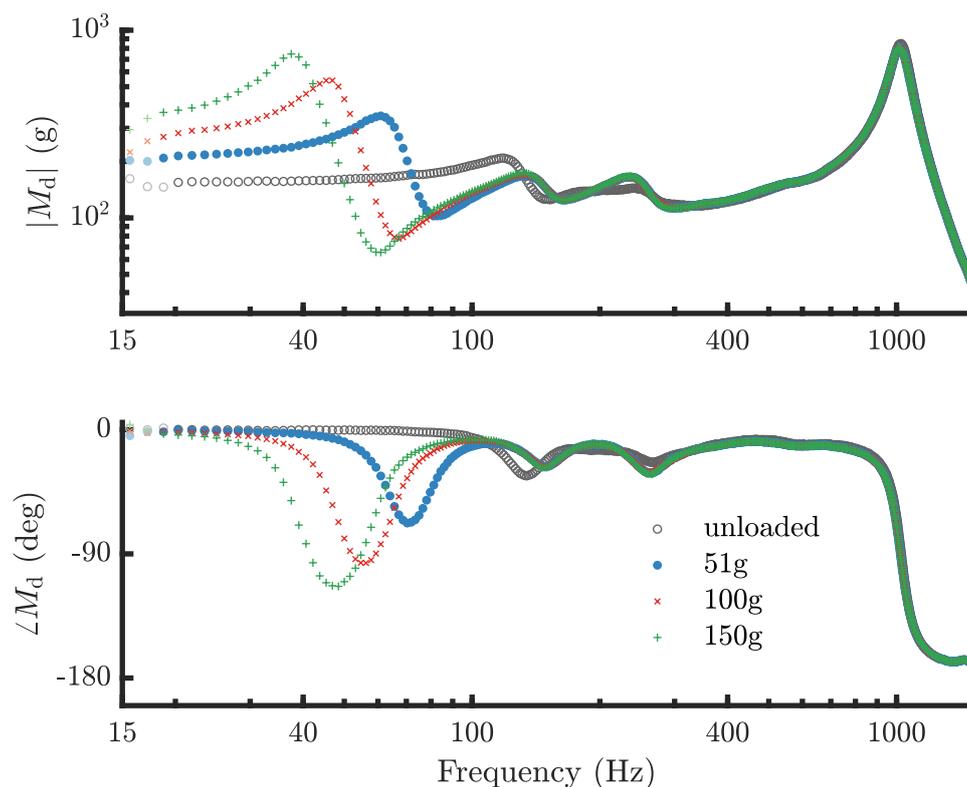


Fig. 6.5 Bode plot of driving point dynamic mass measurements for load masses on the phantom. Results are shown for the unloaded case (\circ) and for loading with 25.4 mm diameter rigid cylindrical load masses of 51 g (\bullet), 100 g (\times) and 150 g ($+$).

In all four cases, the low frequency response is characterised by a dynamic mass plateau at a magnitude approximately equal to the combined mass of the load, the phantom and the ‘mass-above-the-force-gauge’ of the impedance head (m_o). The difference between the plateau levels is approximately equal to the difference between the load masses.

In each of the loaded cases, the low-frequency dynamic mass plateau is followed by an antiresonance–resonance pair, beyond which the curves collapse and no

longer depend on the load mass. This suggests that the antiresonance - resonance pair is related to the shedding of the load mass.

In section 3.4.2 we show that after a ‘mass-locking’ antiresonance – resonance pair the dynamic mass can be expected to reach a new plateau level equal to the mass of the system remaining without the locked mass. In figure 6.5 the dynamic mass FRFs do not reach a new plateau, instead being characterised by further antiresonances and resonances. This suggests that there are resonances within the phantom at intermediate frequencies. These resonances do not depend on the load mass as shown in figure 6.4. However, in section 6.3.2 we show that the intermediate resonances do depend on the load area. This explains why the intermediate-frequency resonances of the unloaded case do not match those of the loaded cases.

While the lack of a clear ‘post mass-locking’ plateau means that we cannot simply read off the amount of mass locked, the fact that for different load masses the FRFs collapse to a single curve after the mass-locking phenomenon suggests that the difference in the amount of mass locked is equal to the difference in load mass.

6.3.2 Results - effect of load contact area on driving point FRFs

Figure 6.6 shows the driving point dynamic mass FRF measured at the base of the phantom when the phantom is loaded by ~ 50 g cylindrical loads with a range of contact areas.

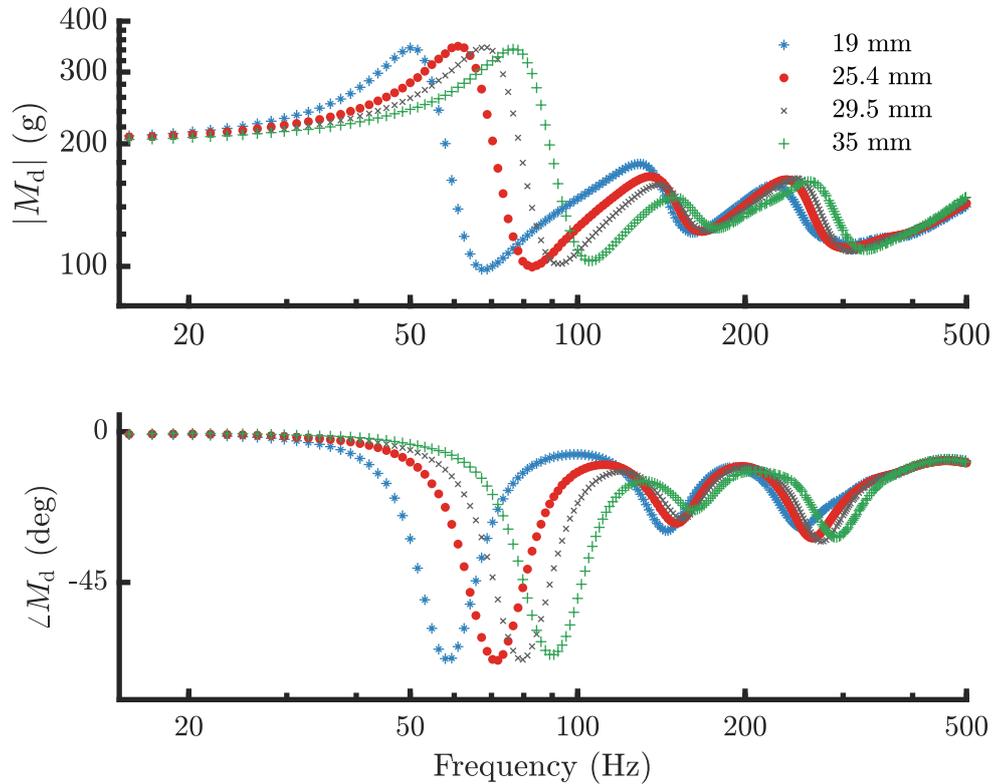


Fig. 6.6 Bode plot of the driving point dynamic mass at the base of the phantom, when the phantom is loaded with ~ 50 g rigid cylindrical load masses with different contact areas.

In figure 6.6 we observe that the response curves no longer collapse above 150 Hz, as they did for different load masses with the same contact area (see figure 6.5). In other words, when we change the contact area, the secondary resonances change frequency. This makes sense when we consider that after the load mass has locked, the proportion of the phantom's top surface that is constrained depends on the contact area of the load. The low-frequency dynamic mass plateau does not depend on load area, and the amplitude of the ω_{-1} peak and the ω_2 trough in dynamic mass are also not affected by the contact area.

Figure 6.7 shows the extracted resonance and antiresonance frequencies as a function of load area. A clear positive trend can be observed in all three cases.

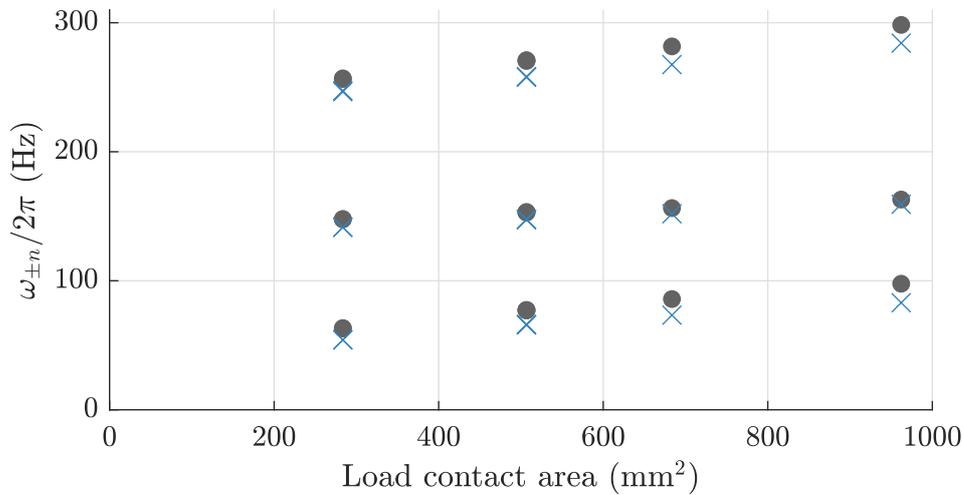


Fig. 6.7 Resonance (●) and antiresonance (×) frequencies for 50 g load masses with different contact areas on the phantom. The values are extracted from driving point response data (including that shown in figure 6.6) by circle-fitting to Nyquist plots of the admittance and impedance driving point FRFs respectively.

6.3.3 Results - transfer FRF

Figure 6.8 shows the ratios of the load and base velocities to the input force at the base of the phantom when the phantom is loaded with an accelerometer of mass ~ 45 g and a contact diameter of 25.4 mm. As expected (see e.g. [87]), the driving point antiresonances cannot be seen in the load's response. The ω_2 resonance can be clearly identified in the plot of $|H_{v:s,b}|$, but the intermediate modes are less evident. See chapter 8 for the phantom response of accelerometers with other masses.

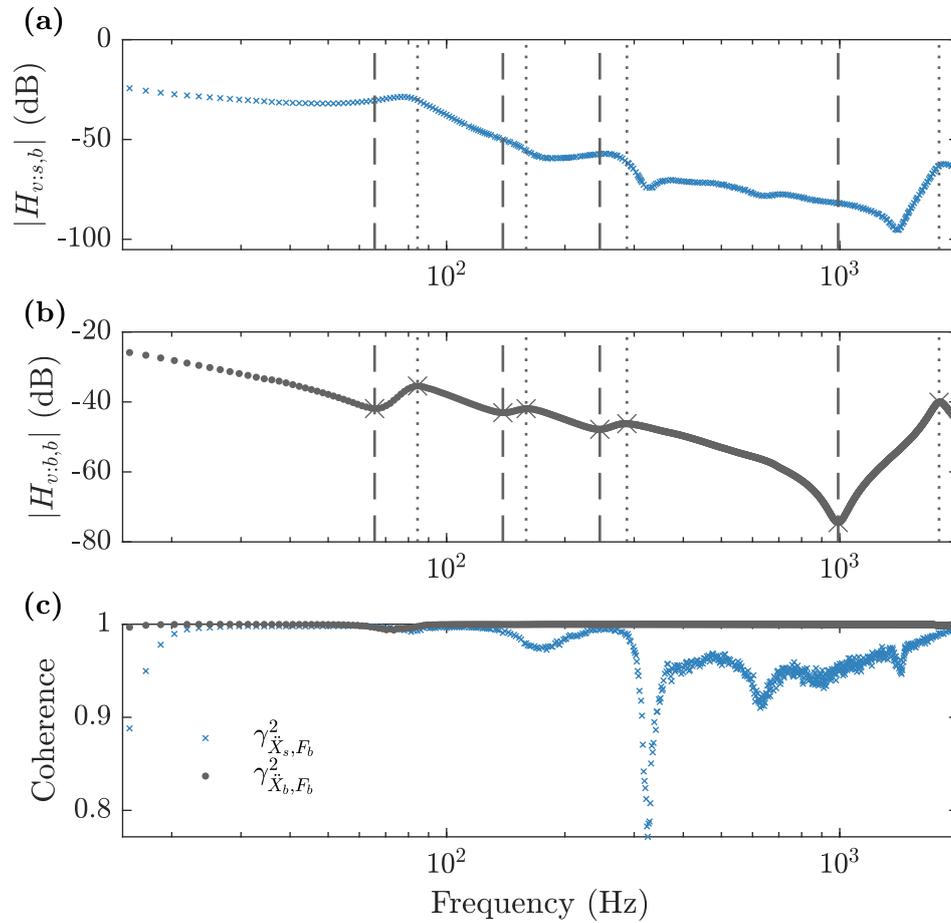


Fig. 6.8 Driving point and transfer FRF measurements for a ~ 45 g accelerometer with a contact diameter of 25.4 mm placed on the top of the phantom. **(a)**: Bode magnitude plot of the H2 estimator for the transfer admittance (load velocity with respect to input force at base). **(b)**: Bode magnitude plot of the H2 estimator of driving point admittance at the base of the phantom. Grey crosses denote the resonances and antiresonances found by circle fitting to the relevant Nyquist plots. **(c)**: coherence for both cases. The dB reference value for both admittance plots is $1 \text{ m s}^{-1} \text{ N}^{-1}$. Vertical grey dotted lines denote the frequencies at which there are resonances in the driving-point response. Vertical grey dashed lines denote the frequencies at which there are antiresonances in the driving-point response.

6.4 2DOF LEM for the phantom

Figure 6.9 shows simplified impedance analogy circuits for a simple rigid load mass on the phantom. We make use of the low-frequency approximation described

in section 5.3.2 to assume that the acceleration measured by the impedance head (\ddot{x}_a) is approximately equal to the acceleration of the base of the impedance head (\ddot{x}_b), i.e., $\ddot{x}_a \simeq \ddot{x}_b$. In figure 6.9a, a transmission matrix $\mathbf{T}_{\text{phantom}}$ is used to describe the relationship between the base of the phantom and the load.

When constructing a spatial model for the phantom we can choose how many modes of vibration to capture by selecting the number of degrees of freedom used for the phantom. Since we are mainly interested in the dominant features of the response, we employ a two degree-of-freedom (2DOF) LEM. As a result of using only two degrees of freedom to describe the phantom, this model is not be able to capture all of the modes observed in the experimental data presented in sections 6.3.1 to 6.3.3.

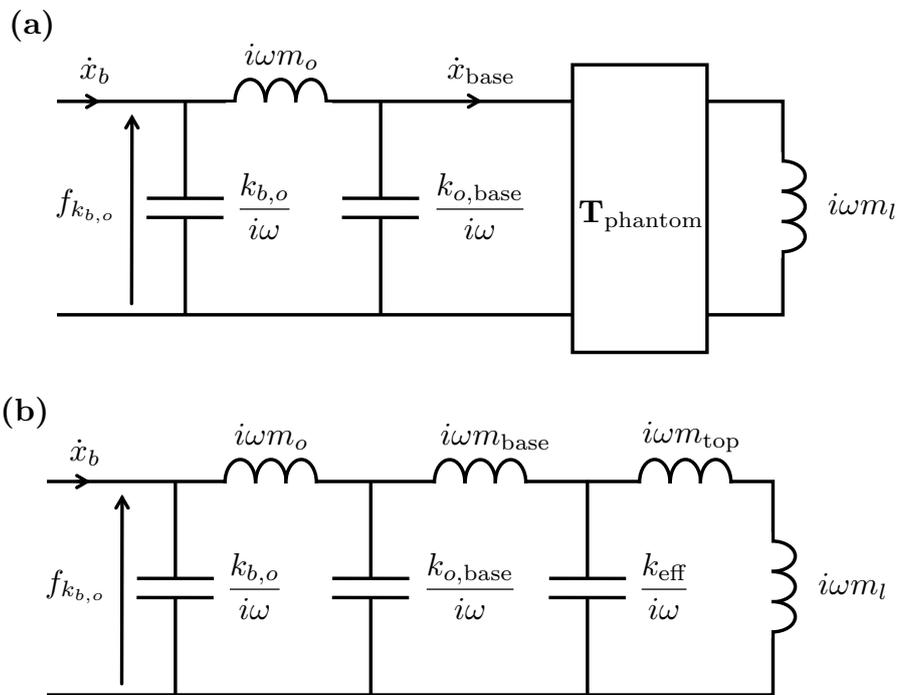


Fig. 6.9 Simplified impedance-analogy circuit for a driving-point response measurement at the base of the phantom with a rigid load applied to the top surface. **(a)** Phantom represented by a transmission matrix. **(b)** Phantom represented by a 2DOF LEM, consisting of two masses (m_{base} and m_{top}) connected by a spring (k_{eff}). Note that the shorthand form $\dot{x}_b = \dot{x}_{\text{base}}$ is also used in this thesis, and should not be confused with x_b .

In figure 6.9b a simple 2DOF LEM is used to represent the phantom. Lumped elements representing parts of the impedance head are also included to account for the difference between the measured force and acceleration and the actual driving

point force and acceleration. The mass lumped to the base of the phantom m_{base} consists of the mass of the aluminium base plate on which the Ecoflex™ 00-10 cylinder sits, as well as a proportion of the mass of the Ecoflex™ 00-10 cylinder. The mass lumped to the load (m_{top}) consists of the remaining proportion of the mass of the Ecoflex™ 00-10 cylinder. The division of the mass of the Ecoflex™ 00-10 cylinder to the two degrees of freedom in the LEM is expected to depend on the contact area of the load. The force-transducer stiffness ($k_{b,o}$) and the mounting stiffness ($k_{o,\text{base}}$) are both expected to be large compared to the effective stiffness of the Ecoflex™ 00-10 cylinder (k_{eff}).

The circuit in figure 6.9b gives a mass-locking phenomenon, which matches the experimental observations discussed in section 6.3.1. In section 3.4.2 we derive expressions for the antiresonance (ω_-) and resonance (ω_+) frequencies in the mass-locking pair. For the LEM in figure 6.9b these equations yield:

$$\omega_- = \sqrt{\frac{k_{\text{eff}}}{(m_{\text{top}} + m_l)}}, \quad (6.1)$$

$$\omega_+ = \sqrt{\frac{k_{\text{eff}}(m_{\text{base}} + m_{\text{top}} + m_l)}{m_{\text{base}}(m_{\text{top}} + m_l)}}, \quad (6.2)$$

where k_{eff} , m_{top} and m_{base} are unknown. However, we do know the net mass $m_{\text{net}} = m_{\text{base}} + m_{\text{top}} + m_l$, which allows us to express equation 6.2 as

$$\omega_+ = \sqrt{\frac{k_{\text{eff}} m_{\text{net}}}{(m_{\text{net}} - m_{tl}) m_{tl}}}, \quad (6.3)$$

where $m_{tl} = m_{\text{top}} + m_l$ for convenience.

We thus have two equations for two unknowns, which can be rearranged to give:

$$m_{\text{top}} = m_{\text{net}} \left(1 - \frac{\omega_-^2}{\omega_+^2} \right) - m_l, \quad (6.4)$$

$$k_{\text{eff}} = m_{\text{net}} \omega_-^2 \left(1 - \frac{\omega_-^2}{\omega_+^2} \right). \quad (6.5)$$

Values for m_0 and $k_{b,o}$ are provided in chapter 5. To calculate the mounting stiffness ($k_{o,\text{base}}$) from the frequency of the ‘mounting-antiresonance’ (see section 5.3.3) we use equation (5.38) with m_l replaced by m_{base} . Note the use of m_{base}

rather than the net mass of the system. This is because the mass associated with x_l has already locked by the time we reach the ‘mounting-antiresonance’.

6.4.1 Lumped element modelling and error

An important clarification needs to be made at this point: the parameters extracted from equations (6.4) and (6.5) are *not* estimates of ‘real’ physical parameters. m_{top} and k_{eff} have no meaning outside of the lumped element model in which they are defined. Any number of degrees of freedom can be chosen for a lumped element model of the phantom, and if a different number of degrees of freedom are chosen then the appropriate parameter values would not be compatible with those extracted for the 2DOF LEM. In particular, the net stiffness obtained from a (say) 3DOF LEM for the phantom would not match k_{eff} from the 2DOF LEM if the parameters in each case are chosen to capture the low frequency mass-locking feature.

For this reason it would be misleading to include error bars on any plots of extracted parameters. Error bars are typically used to give an indication of the range of values in which the true value of a measured parameter is likely to be found [147]. In practice the computation of ‘confidence intervals’ does not really answer this question (see [103]). In any case, for the selection of LEM parameters there is no ‘true’ value, so it does not make sense to talk about a region within which we expect the true value to lie. Instead, we present parameters extracted from several repeats of the experiments to demonstrate the variability in the extracted parameter values. The quality of the extracted parameters must be assessed in relation to whether they fulfil their aim, that is: whether they allow us to capture the role of the phantom in the response of sensors to excitation at the base of the phantom, for frequencies below ~ 100 Hz.

6.4.2 Parameter extraction - varying load mass

Figure 6.10 shows the values of m_{top} and k_{eff} calculated using equations (6.4) and (6.5) with the known static mass of the setup and the extracted resonance (ω_2) and antiresonance (ω_{-1}) frequencies shown in figure 6.4. Neither the lumped mass nor the lumped stiffness extracted for a 2DOF LEM of the phantom varies significantly with load mass applied for the range of masses shown (i.e., the trend

is insignificant compared to the scatter). However, the range of masses used was limited by the fact that taller cylinders would start to oscillate in transverse modes [153].

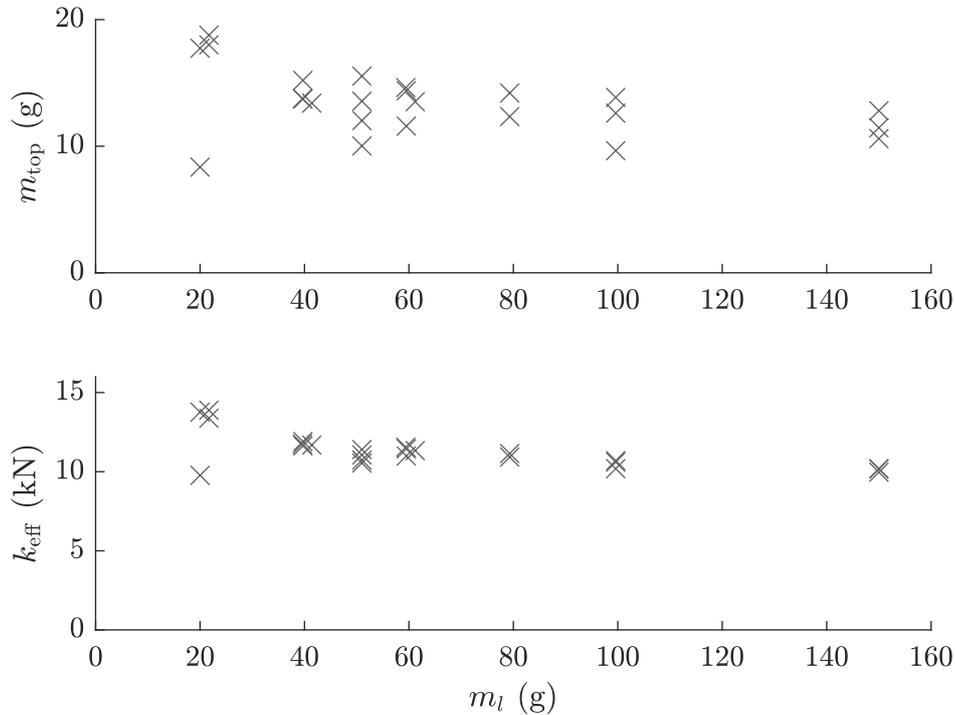


Fig. 6.10 Extracted parameters for a 2DOF LEM model of the phantom for a range of cylindrical load masses with a constant contact diameter of 25.4 mm.

6.4.3 Parameter extraction - effect of area

In figure 6.11 it can be observed that, as the contact area between the phantom and the load mass is varied, both m_{top} and k_{eff} increase with load area. The stiffness increases approximately linearly with area while the contact-area is small, but then reaches a plateau as the contact area approaches the surface area of the top of the phantom.

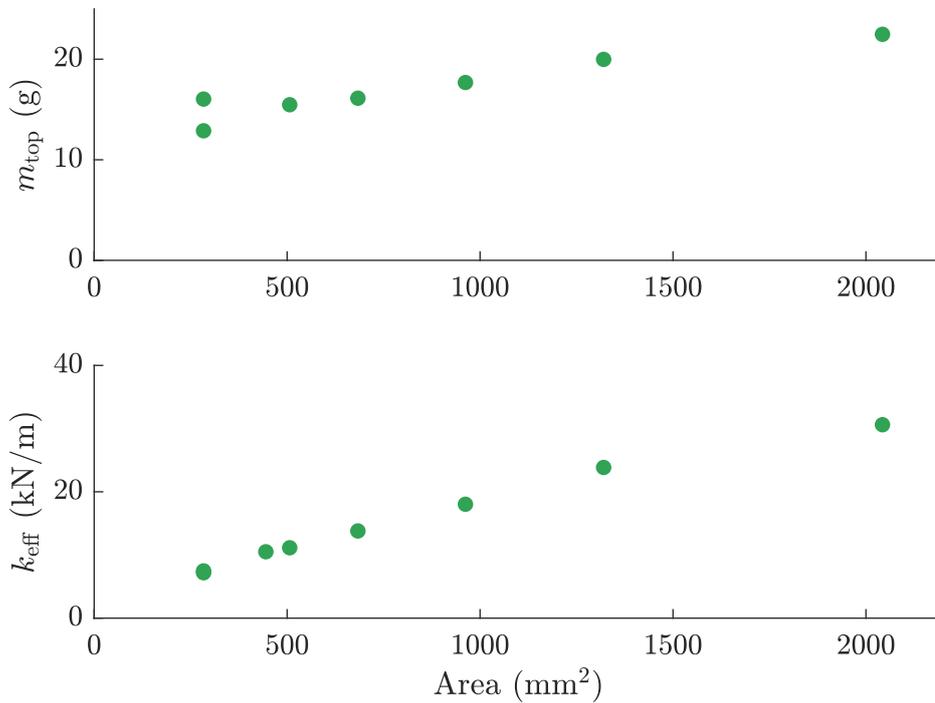


Fig. 6.11 Extracted parameters for a 2DOF LEM model of the phantom for rigid cylindrical loads of mass ~ 50 g with a range of contact areas. The surface area of the top face of the phantom is ~ 4400 mm².

6.4.4 Performance of 2DOF LEM

Figure 6.12 shows simulated results for 40 g and 60 g loads on top of the phantom. The simulation is performed using the LEM represented in figure 7.14(b). The simulated admittance FRF for each DOF in the model with respect to excitation at the base is depicted. Experimental data is shown for two of these DOFs: the base driving point and the load.

The driving point admittance simulated using a 2DOF LEM for the phantom matches the measured response well up to the second resonance frequency of the system (ω_2), and correctly captures the mounting-antiresonance, but deviates from the measured response between these two features. In particular, the model over-estimates the proportion of mass shed in the first mass-locking feature. By the time we approach the mounting antiresonance, however, the model over-estimates the mass ‘seen’ at the driving point (i.e. it underestimates the amount

of mass locked), as further mass-locking features occur that cannot be captured by a 2DOF LEM.

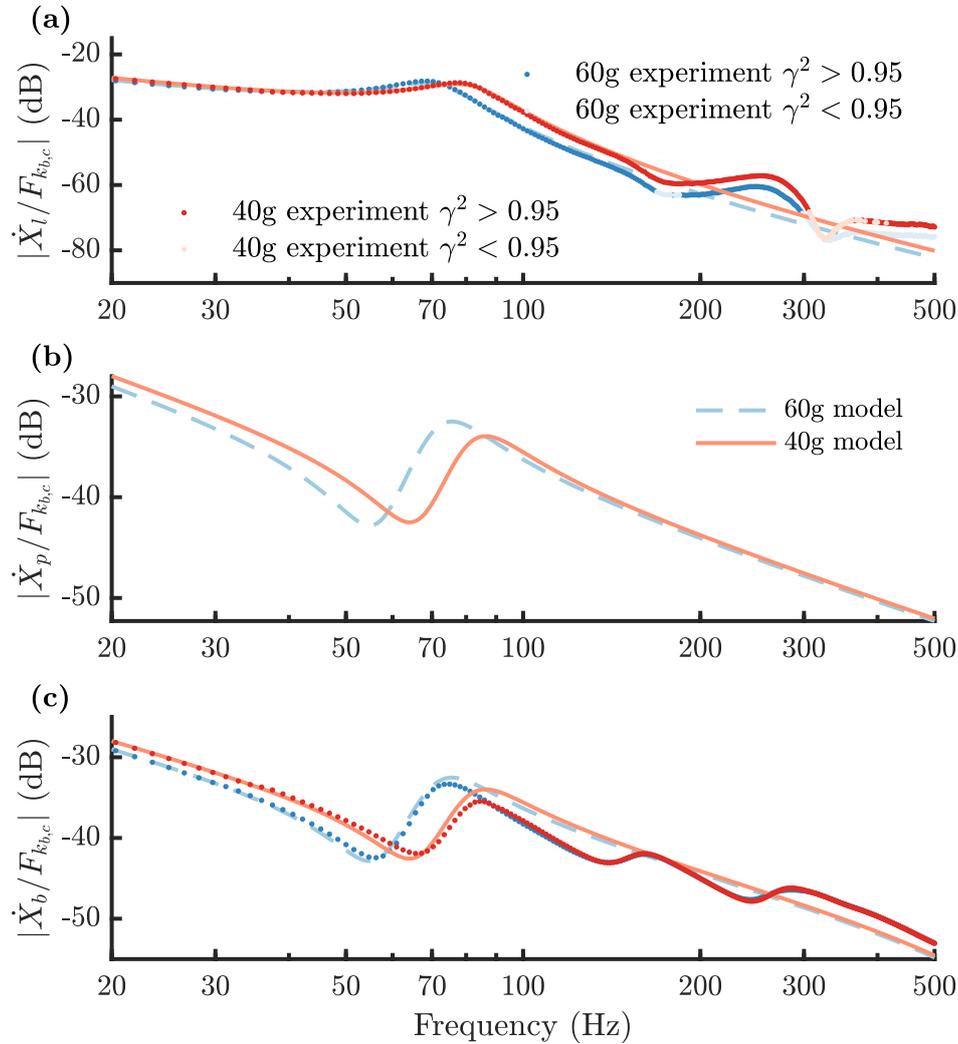


Fig. 6.12 Simulated and experimental response of two direct contact sensors on the phantom, using a 2DOF LEM to represent the phantom in the simulation. **(a)** Load velocity with respect to $F_{k_{b,c}}$. **(b)** Velocity of the base of the phantom with respect to $F_{k_{b,o}}$. **(c)** Velocity measured by the impedance head with respect to $F_{k_{b,o}}$. The dB reference is $1 \text{ m s}^{-1} \text{ N}^{-1}$ for each plot. Experimental data is shown using markers ($\bullet = 60 \text{ g}$, $\bullet = 40 \text{ g}$), with lighter shading denoting data-points with a coherence of less than 0.95. Lines in a lighter shade of the same colour represent the 2DOF-phantom simulations.

The model also provides a good fit to the measured response at the load (see figure 6.12a), despite the fact that no information about the response of the

load was used to determine the lumping used for the model. This simulation illustrates an important point: we can accurately predict the response at low frequency and predict the general shape of the overall response with only very crude data on the mechanical properties of the test-system.

6.5 Analysis for linear phantom behaviour

The phantom could, of course, be modelled using a LEM with an arbitrary number of degrees of freedom in order to better match the response. However, simply fitting a model with an arbitrary number of degrees of freedom to the measured response would not necessarily improve our understanding of the system. If the goal is simply to describe the transmission matrix between the driving point at the base of the phantom and the load on top of the phantom, then rather than using an LEM to model this transmission matrix we can determine the transmission matrix directly from experimental measurements.

6.5.1 Method

If the behaviour of the phantom can be approximated as linear (with respect to the range of static loads applied) then the analysis can be significantly simplified. For convenience we define a new transmission matrix $\tilde{\mathbf{T}}_{\text{phantom}}$ which includes the transmission matrix for the phantom itself ($\mathbf{T}_{\text{phantom}}$) in addition to terms relating the base of the phantom to the terms measured by the impedance head, such that:

$$\tilde{\mathbf{T}}_{\text{phantom}} = \begin{bmatrix} 1 & 0 \\ i\omega/k_{b,o} & 1 \end{bmatrix} \begin{bmatrix} 1 & i\omega m_o \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ i\omega/k_{o,\text{base}} & 1 \end{bmatrix} \mathbf{T}_{\text{phantom}} . \quad (6.6)$$

Figure 6.13 shows the resulting simplified impedance analogy circuit.

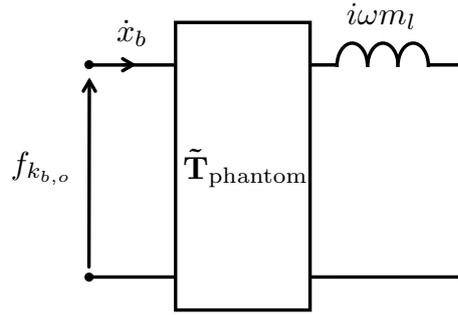


Fig. 6.13 Simplified impedance analogy circuit driving point response measurement at base of phantom with a rigid load applied. $F_{k_{b,o}}$ is the force measured by the impedance head and \dot{x}_b is the velocity of the base of the impedance head, not to be confused with the base of the phantom (\dot{x}_{base} or \dot{x}_b). For all frequencies of interest \ddot{x}_b is approximately equal to the measured acceleration \ddot{x}_a .

For the system consisting of a direct sensor of impedance Z_s mounted on a phantom which is excited through an impedance head at its base, we can write the governing equation in the form:

$$\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{phantom}}} \begin{bmatrix} Z_s \dot{X}_s \\ \dot{X}_s \end{bmatrix}, \quad (6.7)$$

where A , B , C and D are the terms in the *backwards* transmission matrix interpretation of $\tilde{\mathbf{T}}_{\text{phantom}}$ in figure 6.2(a). $F_{k_{b,o}}$ is the frequency domain representation of the force measured by the impedance head, while \dot{X}_b and \dot{X}_s are the velocities corresponding to the accelerations measured at the impedance head and at the load (related via $\dot{X} = \ddot{X}/i\omega$). Dividing equation (7.4) through by \dot{X}_s we obtain:

$$\begin{bmatrix} F_{k_{b,o}}/\dot{X}_s \\ \dot{X}_b/\dot{X}_s \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{phantom}}} \begin{bmatrix} Z_s \\ 1 \end{bmatrix}. \quad (6.8)$$

If the transmission matrix for the phantom does not depend on the load applied, then by applying different loads and measuring the response functions we can acquire a set of simultaneous equations to solve for $\tilde{\mathbf{T}}_{\text{phantom}}$. Matrix equation (6.8) consists of two simultaneous equations for four unknowns, so we need to measure FRFs for two different load impedances to solve for A , B , C and D . Using subscripts I and II for the two different loads and the associated FRFs,

this yields:

$$\underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{phantom}}} = \begin{bmatrix} (F_{kb,o}/\dot{X}_s)_{\text{I}} & (F_{kb,o}/\dot{X}_s)_{\text{II}} \\ (\dot{X}_b/\dot{X}_s)_{\text{I}} & (\dot{X}_b/\dot{X}_s)_{\text{II}} \end{bmatrix} \begin{bmatrix} (Z_s)_{\text{I}} & (Z_s)_{\text{II}} \\ 1 & 1 \end{bmatrix}^{-1}. \quad (6.9)$$

6.5.2 Implementation

The results presented in sections 6.3.1 and 6.3.2 suggest that the properties of the phantom (for transmission from the base to the load) are reasonably linear with respect to load mass (for the small range of load masses considered), but not with respect to the contact area between the phantom and the sensor. This means that *for a given contact area* we can use the method described in section 6.5.1 to compute a transmission matrix that can be used to represent the phantom.

To illustrate the method we solve equation (6.9) for A , B , C and D using FRF data for direct sensors with a contact area of 25.4 mm. For the two load-cases required we use direct contact sensors with masses of 62 g (case I) and 152 g (case II) on the top surface of the phantom.

Once the transmission matrix has been computed, the driving point response and load response can be calculated using:

$$H_{v:s,b} = \frac{1}{AZ_s + B} \quad (6.10)$$

$$H_{v:b,b} = \frac{CZ_s + D}{AZ_s + B}. \quad (6.11)$$

The use of use of $\tilde{\mathbf{T}}_{\text{phantom}}$ to simulate the response of accelerometer with different load masses is discussed in chapter 8. Figure 8.1 demonstrates the excellent match between experimental data and simulations using $\tilde{\mathbf{T}}_{\text{phantom}}$ for accelerometers with a range of load masses (within a range of masses for which the phantom behaviour can be approximated as linear with respect to load mass). This serves as a validation of the method proposed here. A further validation is presented in section 6.5.3.

6.5.3 Validation - driving point impedance as a function of Z_{load}

As a further validation of the method outlined in section 6.5.1, the driving point FRF results presented in section 6.3.1 for 25.4 mm diameter loads of various masses (without accelerometers attached) are compared to simulations using the computed transmission matrix for the phantom (with equation (6.11)). The results of this comparison are shown in figure 6.14.

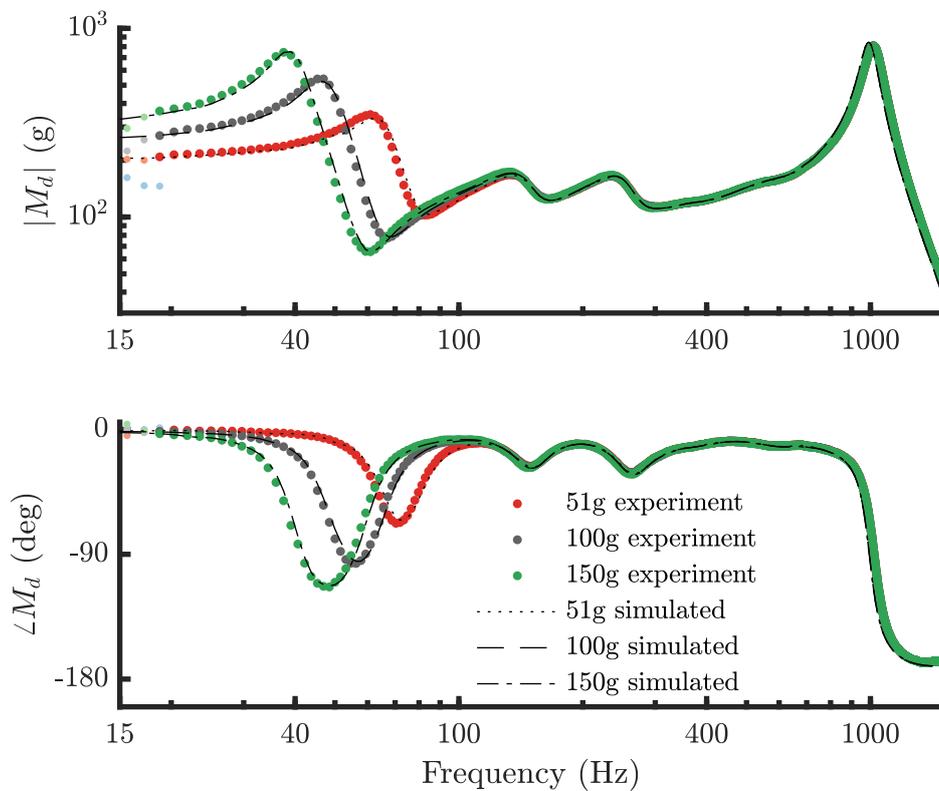


Fig. 6.14 Bode plot of the driving point dynamic mass FRFs for a range of 25.4 mm diameter load masses on the top of the phantom. Measured data-points are shown using markers, the FRFs simulated using equation (6.10) with $\tilde{\mathbf{T}}_{\text{phantom}}$ are shown using lines.

Figure 6.14 shows that the transmission matrix approach is able to accurately predict the driving point response for a range of load masses. One discrepancy that can be observed between the measured and simulated data is the fact that the anti-resonance at around 1 kHz does not quite match up. This is because this resonance depends on the mounting stiffness between the base of the phantom and the impedance head as discussed in section 5.3.3. The experimental setup

was completely disassembled and reassembled between the recording sessions for the data used to compute the transmission matrix and the recording sessions for the rigid-load-data used in this validation. The computed transmission matrix for the phantom is strictly valid only when the mounting stiffness is equal to that of the setup for the FRFs that were used to compute this matrix.

6.6 Effect of sensor position on phantom behaviour

For the data presented in this chapter (and in subsequent chapters in which this phantom is used to test stethoscopes) the load (or stethoscope) is always carefully positioned at the centre of the top surface of the phantom. In this section we explore the effect of deviations from an ideal placement. Mansy et al. [95] explored the spatial differences in the surface motion of an unloaded phantom but, as noted in section 6.3, the unloaded response of the phantom is not necessarily representative of the behaviour of the same phantom once it is loaded by sensors. We therefore wish to establish how robust the results obtained from experiments on our phantom are with respect to accurate central placement of the sensor on the phantom.

Figure 6.15 shows Bode-magnitude plots of the driving point ($H_{v:b,b}$) and transfer ($H_{v:s,b}$) admittance FRFs computed from experiments performed with the sensor placed at different locations on the surface of the phantom. The results are grouped by deviation from a central position, with small deviations ($r \leq 5$ mm) shown by blue dots and large deviations ($r \geq 8$ mm) shown by light red asterisks. The small deviations have almost no effect on the FRFs, while the larger deviations (larger than any anticipated error in positioning) have a slightly more significant effect, especially around the resonances and anti-resonances between 150-400 Hz. For both small and large deviations from a central position the low-frequency behaviour is almost unaffected by the placement, providing further justification for the adoption of a 1D-vibration assumption at these frequencies. The only case which gives visible deviation in the resonance and anti-resonance associated with the mass shedding phenomenon in the driving point response is the FRF for a 20 mm deviation from the centre. In this case the sensor is placed closer to the edge of the phantom than to its centre. Consequently, there are likely to

be significant unmeasured transverse reaction forces exerted on the structure through the force transducer which could explain the discrepancy.

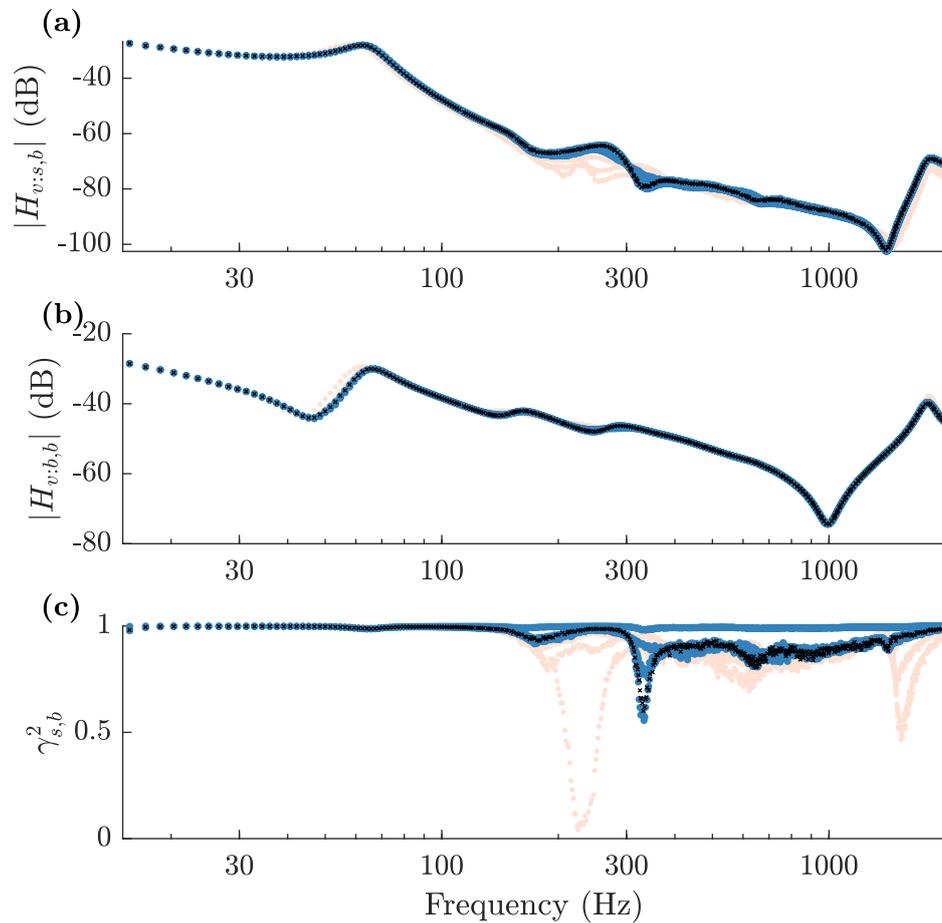


Fig. 6.15 (a) Transfer and (b) driving point admittance FRF measurements for the setup in figure 6.2 with an accelerometer on a ~ 100 g load mass, performed for a variety of radial deviations in the positioning of the load mass. (c) Coherence of the transfer response. Defining r as the distance between the centroid of the load mass and the centroid of the phantom, the results are grouped: black crosses (\times) $r = 0$, blue dots (\bullet) $2 \leq r \leq 5$ mm, light red asterisks ($*$) $8 \leq r \leq 20$ mm. Nine cases are shown.

6.7 Summary and conclusions

In this chapter we present the design of a chest-phantom for testing stethoscopes. We provide a detailed review of existing work on such phantoms in the literature and address issues found in some of the methods adopted. In particular, we highlight the need to directly measure the excitation input and the need to characterise the behaviour of the *loaded* phantom.

In this context, we present results on the driving point response measured at the base of the phantom when the top surface is loaded using rigid cylinders with a range of contact areas and masses. We show that the transmission of vibrations from the driving point at the base of the phantom to the load depends significantly on the contact area, but only weakly on the mass of the load (for a small range of light loads). A simple 2DOF LEM is able to capture the low frequency behaviour of the phantom, while a transmission matrix can be computed from measured FRF data for a given contact area in order to accurately predict the response of rigid cylindrical load on the phantom over a wide frequency range.

The phantom described in this chapter is used in chapters 8, 9 and 10 to explore the performance of direct and air-coupled sensors respectively. In chapter 7 we present driving point impedance measurements taken at the load position on the phantom, which contribute to a further characterisation of the behaviour of the phantom.

Chapter 7

Impedance measurements

In this chapter we investigate the coupled patient–stethoscope–clinician system from the perspective of the impedance seen by the sensor. This impedance has two components: the impedance of the patient’s body, and the ‘attachment impedance’ which incorporates, for example, transducer wires or the hand of the clinician. In this chapter we consider the impedance seen by a direct contact sensor (accelerometer) in detail and comment on how the method can be extended to air-coupled stethoscopes. The word ‘chest’ is used here as a placeholder for any area of the human body that a stethoscope is coupled to.

The lumped element models developed for different classes of stethoscope in chapter 4 can be used to simulate coupled chest sensor systems *provided appropriate values for the parameters are available*. The LEMs can also inform an impedance coupling approach, in which the performance of a sensor is estimated by considering the impedance of the sensor and the impedances seen by the sensor.

For direct contact stethoscopes, the impedance analogy circuit in figure 4.7 can be rearranged to yield the impedance seen by the sensor when the source term is ‘off’. Figure 7.1(a) shows this rearrangement for a generic chest impedance (z_{chest}) and a generic attachment impedance z_{load} . Figure 7.1(b) shows the same circuit with the chest modelled as a SDOF system, with an effective mass (m_{eff}), stiffness (k_{eff}) and damping rate (λ_{eff}).

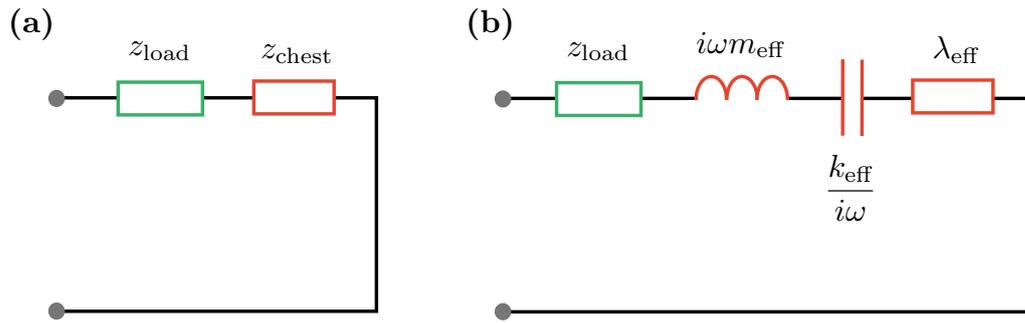


Fig. 7.1 Impedance seen by a direct sensor in contact with the human chest, and loaded by e.g. a wire or hand. (a) General model with chest and load impedances. (b) SDOF model for the chest.

The linear superposition principle [87] allows us to measure the chest impedance (z_{chest}) and the attachment-impedance (z_{load}) seen by the sensor independently, and then combine them to obtain the net impedance seen by the sensor. Sections 7.3 to 7.4 address the mechanical impedance of the chest and the phantom, while section 7.5 considers the attachment impedance associated with holding a sensor.

7.1 Measuring the impedance seen by a rigid sensor on the human chest

There is no such thing as ‘*the*’ mechanical impedance of the chest. Mechanical impedance is the ratio of a force to the velocity induced by that force. In order for this to make any sense we need to specify how the force is applied to the chest surface, and over which area the surface-velocity is averaged.

Since we are interested in the impedance seen by a direct contact sensor on the chest, the chest interface for the impedance measurement needs to match the properties of the chest-sensor interface. A typical direct contact sensor is essentially rigid (compared to the chest), so the impedance measurement can be performed by measuring the driving point impedance of a rigid coupler pressed against the chest.

This method has been adopted by several researchers to investigate the properties of human tissue (see e.g. [56, 148, 154, 153, 5]). The experimental setups used in four representative studies spanning 60 years are summarised in table 7.1.

Table 7.1 details of four representative experimental studies conducted to measure the mechanical impedance of human tissue seen by a rigid coupler.

Author(s)	Year	Location	Coupler	Preload
Bárány et al. [5]	1942	human chest	25 and 35 mm diameter rigid circular couplers	~ 5 N (not controlled)
Hussein [56]	1979	human finger	6.35 mm diameter rigid circular coupler	0.5 - 5 N (controlled)
Vermarien and van Vollenhoven [153]	1984	human chest	20, 25 and 35 mm diameter rigid circular couplers	not measured or controlled
Gierke and Brammer [154]	2002	'over soft tissue'	20 mm diameter rigid circular coupler	0 - 12 mm indentation

Bárány et al. (1941) [5] and Vermarien and van Vollenhoven (1984) [153] observed that the mechanical impedance for a given contact area varies depending on the static force ('preload') applied to the coupler. However, neither group were able to accurately measure the static preload in their experiments. Bárány observed that for low preloads (light contact) the impedance was strongly dependent on preload, while for higher preloads the influence plateaued. Hussein (1979) [56] conducted measurements on the fingertips of 16 human volunteers and found that the stiffness components of the mechanical impedance increased approximately linearly with preload from 0.5 to 5 N, while the viscosity component increased to a plateau. Inspection of the presented data suggests that the viscosity term increases approximately linearly up until an application pressure of around 80 kPa before beginning to plateau. Gierke and Brammer (2002) [154] measured the equilibrium indentation rather than the preload applied and found that both the real and imaginary parts of the mechanical impedance measured on the human chest increased with indentation.

The results presented by these authors give an indication of the trends to be expected, but are limited in scope. Only a small range of contact-areas and static loads were utilised, and the variability across different locations on the chest, as well as between subjects were only minimally explored. Gierke and

Brammer [154] and Hussein [56] both report significant inter-subject variation. This means that the impedance coupling between a direct contact sensor and the chest, and consequentially the sensor performance, is expected to be strongly subject dependent.

In chapter 6, we show that it is possible to characterise the mechanical properties of a chest phantom by using measurements at both the source and load positions. For the human chest this is not possible. In order to confirm that the measurements performed on our phantom are meaningful in terms of the performance of the same sensor on the chest we thus need to compare the properties of the two source systems as seen at the surface where the sensor is attached.

In this chapter, we measure the driving point impedance seen by a rigid sensor on the surface of our chest-phantom. We compare the features of this driving point impedance to reported chest impedance results in the literature and to our own driving point impedance measurements on the chest of a human subject.

7.2 Methods: balance rig setup

The experimental rig shown in figure 7.2 has been designed to measure the driving point impedance of the chest seen by a rigid surface. The setup achieves two important things:

1. It enables the use of a shaker to excite the chest of human participants in a supine position (i.e. lying on their backs).
2. It decouples the static load applied to the chest from the mass of the rigid coupler.

The experimental setup (henceforth referred to as the ‘balance rig’) is designed to act as a balanced lever onto which masses can be added to vary the static load applied to the chest. The design consists of a fulcrum with a rigid lever-arm pivoted at an adjustable height. A shaker is mounted at the distal end of the longer side of the lever-arm. A load pan is mounted directly above the shaker to allow mass to be added vertically above the line-of-action of the shaker-force. At the shorter side of the lever-arm a set of masses are attached to balance the torque from the shaker setup. The masses are adjusted to ensure that the static equilibrium position of the system is with the lever-arm horizontal. A similar

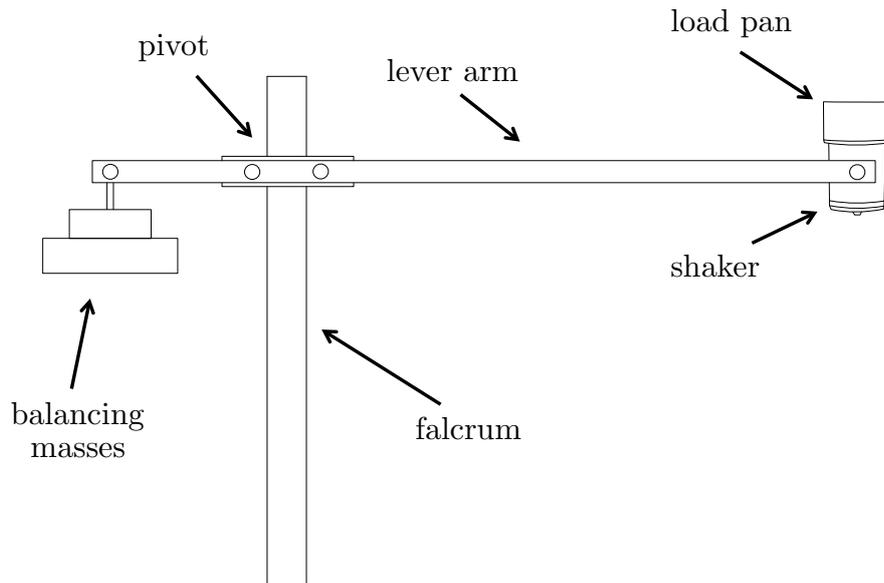


Fig. 7.2 ‘Balance rig’ for chest impedance measurements.

setup is used by Zimmerman et al. [173], albeit with an impedance tube rather than a shaker-impedance-head configuration.

The shaker is connected via a stinger to an impedance head as shown in figure 7.3. The stinger reduces the transmission of transverse reaction forces to the test structure. The presence of unknown transverse reaction forces has been shown to lead to erroneous response measurements [31]. A rigid cylindrical metal coupler is screw-mounted to the structure-side of the impedance head, so that the flat face of the coupler is perpendicular to the line-of-action of the force, and in a plane with the test-surface (chest or phantom).

In order to obtain a reasonable estimate of the Thévenin impedance, the excitation amplitude from the shaker needs to be sufficiently high so that the amplitude of body sounds and background vibration is negligible in the input signal comparison. However, the shaker amplitude should also be sufficiently small to avoid the effects of non-linearities in the mechanical properties of the chest.

Figure 7.4 shows an impedance analogy circuit for the measurement system. Terms ‘upstream’ of \ddot{x}_b ($\simeq \ddot{x}_a$) and $f_{k_{b,o}}$ have been omitted as we are only interested in establishing the relationship between these two measured terms and the chest impedance seen by the coupler (z_{chest}). The mass of the rigid coupler (m_l), the mounting stiffness ($k_{o,l}$), the ‘mass-above-the-force-gauge’ (m_o) and

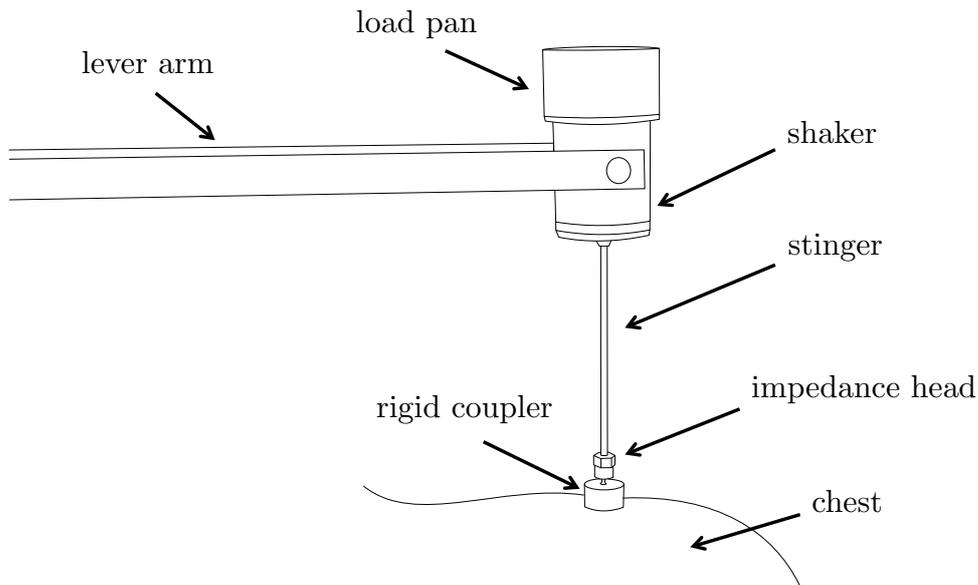


Fig. 7.3 Balance rig for chest impedance measurements: focus on shaker side.

the stiffness of the force transducer ($k_{b,o}$) all need to be taken into account to establish z_{chest} .

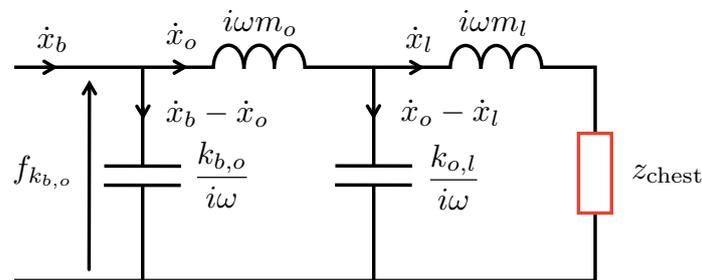


Fig. 7.4 Impedance analogy circuit for the measurement setup in figure 7.3.

7.2.1 Data interpretation

Figure 7.5 shows an example of a set of admittance FRFs measured using the balance rig with a 51 g, 24.5 mm diameter rigid coupler. When the coupler is suspended freely in air, the admittance response is a straight line at -20 dB/decade corresponding to a pure mass load within the frequency range shown. When the coupler is pressed against the human torso, the admittance response is dominated by a single resonance. Figure 7.5 shows admittance response data for two different static loads. In both cases the low frequency response is stiffness dominated, although the low frequency of the resonance means that the

anticipated +20 dB/decade plateau (see section 3.4.1) cannot be identified. For the case with the lower static load (~ 1 N) there is a clear peak, while when the static load is increased to ~ 3.5 N the peak is significantly more damped. At frequencies above the resonance frequency the responses almost collapse onto the free case, which suggests that the mass of the coupler dominates the net effective mass seen by the impedance head. Although not visible at the resolution of figure 7.5, the levels of the (-20 dB/decade) mass plateaus are lower for the loaded cases than for the free case, indicating that the chest contributes a small amount of effective mass to the driving point impedance.

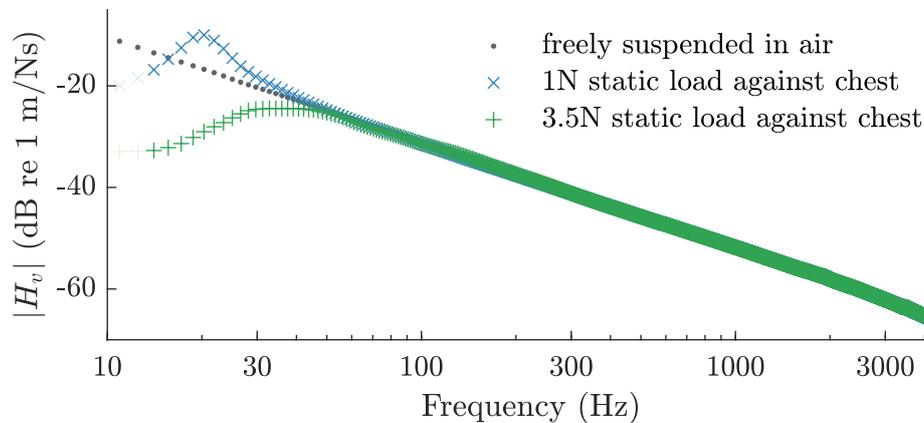


Fig. 7.5 Bode magnitude plot of the driving point admittance FRF measured using the balance rig for a 51 g, 24.5 mm diameter rigid coupler: (\bullet) suspended freely in air, (\times) pressed against the human chest with a static load of ~ 1 N, ($+$) pressed against the human chest with a static load of ~ 3.5 N. Data with a coherence of less than 0.9 are shown in light colour. Both measurements are taken on the upper back of a 26 year old male.

The fact that the response is dominated by a single resonance suggests that the system can be approximated using a SDOF lumped element model (this is consistent with e.g. the approach of Vermarien and van Vollenhoven [153]). In section 3.4.1 we show that, for the special case of a SDOF system, the real and imaginary parts of the inverse driving point FRFs relate directly to the lumped parameter values. Separating the real and imaginary parts of the mechanical impedance response function (Z^m) in equation 3.30 (for a SDOF system with mass m_1 , damping rate $\lambda_{1,0}$ and stiffness $k_{1,0}$), and allowing for a complex stiffness $k_{1,0} \rightarrow k_{1,0}(1 + i\eta)$ where η is the loss factor, the real and imaginary parts of Z^m

can be expressed as:

$$\Re(Z^m) = \frac{\eta k_{1,0}}{\omega} + \lambda_{1,0} \quad (7.1)$$

$$\Im(Z^m) = \frac{k_{1,0}}{\omega} - \omega m_1 \quad (7.2)$$

Figure 7.6 shows the magnitudes of the real and imaginary part of the mechanical impedance FRF for the ~ 1 N static load case from figure 7.5.

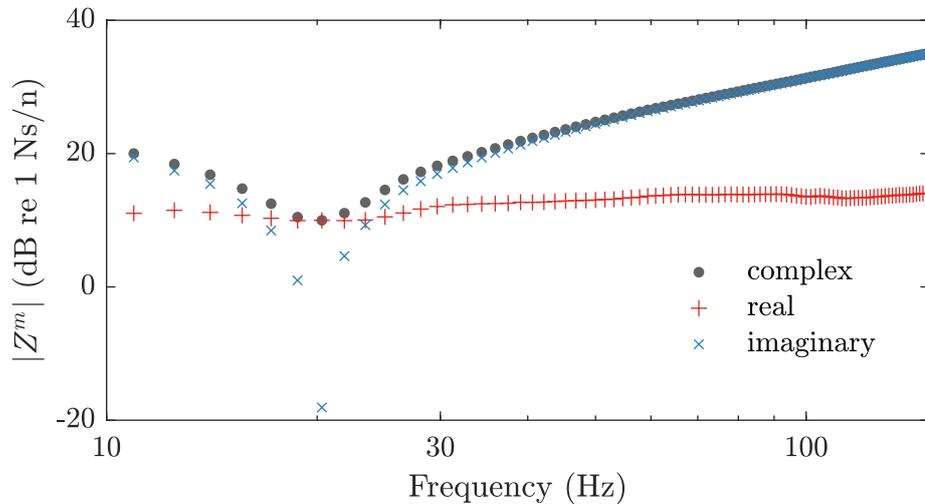


Fig. 7.6 Decibel magnitudes of the complex (\bullet), real ($+$) and imaginary (\times) components of the mechanical impedance FRF for the ~ 1 N static load case from figure 7.5.

The real part of the response is roughly constant over the frequency range shown which tells us, according to equation (7.1), that viscous damping is the dominant damping term (i.e., $\eta \simeq 0$). The viscous damping coefficient (λ_{eff}) is equal to the (constant) magnitude of the imaginary part. At low and high frequencies it is evident that the magnitude of the complex mechanical impedance FRF is dominated by the imaginary part, corresponding to stiffness and mass respectively. There is a trough in the real part of the response when the mass and stiffness terms cancel, which corresponds to the first natural frequency of the system (ω_1).

To estimate the effective mass of the chest seen by the phantom we compare the value of the mass plateau for the chest-contact case to that of the free case. The driving point dynamic mass FRFs for the ~ 1 N static load case and the freely suspended case are shown in figure 7.7. When the coupler is pressed against the

chest, the dynamic mass plateau ($|M_d|_{\text{chest:plateau}}$) is at a higher level than when it is suspended freely ($|M_d|_{\text{free:plateau}}$) which indicates that the chest contributes to the effective mass seen by the impedance head. The effective mass contributed by the chest can be computed from the difference between the two plateau levels (i.e. $m_{\text{eff}} = |M_d|_{\text{chest:plateau}} - |M_d|_{\text{free:plateau}}$).

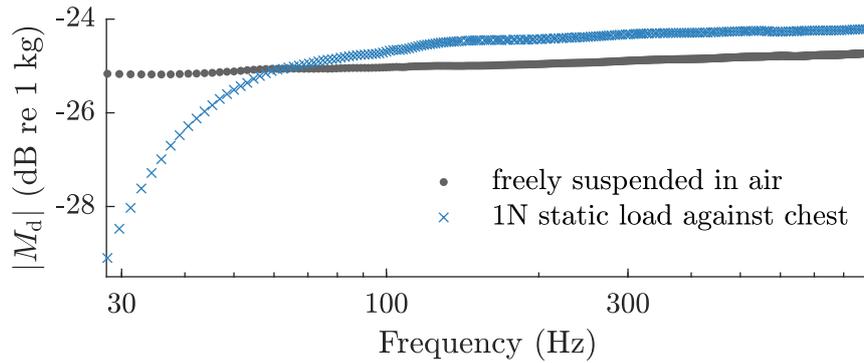


Fig. 7.7 Decibel magnitude of the real component of the dynamic mass FRF for a freely suspended coupler (\bullet) and coupler pressed against the chest with a static load of ~ 1 N (\times). $|M_d|_{\text{chest:plateau}}$ and $|M_d|_{\text{air:plateau}}$ can be extracted from the plateau levels. Data from figure 7.5.

Since the resonance (ω_1) occurs at low frequency, there are only a small number of data-points corresponding to the low-frequency stiffness plateau. Furthermore, the close vicinity of these datapoints to ω_1 suggests that the ‘true’ value of the plateau (as $\omega \rightarrow 0$) is unlikely to have been reached. Attempting to extract the effective stiffness directly from this plateau therefore leads to an underestimation. Instead, we compute the effective stiffness using $k_{\text{eff}} = \omega_1^2 |M_d|_{\text{chest:plateau}}$.

We are interested in characterising the effect of contact area and static load on the impedance seen by a rigid coupler pressed against the chest (or phantom). In order to do this we carry out a set of measurements using couplers of different contact areas and applying a range of static loads through the load pan.

7.3 Impedance seen by rigid disc on human chest

Figure 7.8 shows the cylindrical load mass in contact with skin on the upper back of a human subject. The single human subject in this study was the principle investigator. A risk assessment was carried out and informed consent was given.

Results for a single subject “ $n=1$ ” cannot provide any indication of the variability of the measured parameters among the population. The results are presented as a case study [47] and should not be used to make general inferences about the expected value of the parameters in other subjects. Nevertheless, the trends observed can be compared to results published by other researchers as a crude validation of the methodology. The purpose of these measurements is to confirm that the phantom can represent the relevant properties of the human torso to a reasonable degree. A detailed investigation of the mechanical impedance seen by sensors coupled to the chests of a range of human subjects is outside the scope of this thesis.

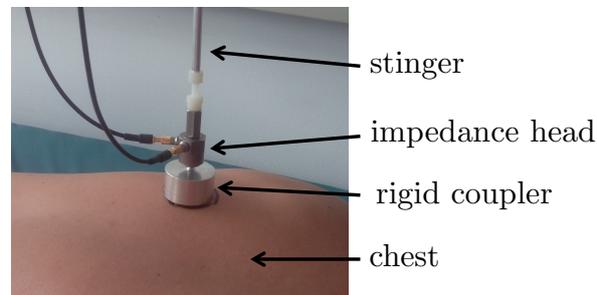


Fig. 7.8 Photograph of the ‘balance rig’ impedance head setup being used to measure the mechanical impedance seen by a rigid circular coupler on a human subject’s ‘chest’ (upper back).

Figure 7.9 shows the effective mass (m_{eff}), viscous damping rate (λ_{eff}), and stiffness (k_{eff}) extracted from measurements on the upper back of the human subject. All three terms increase with contact area. With respect to the static pressure applied, the effective mass quickly rises to a plateau, and then does not vary further with application pressure. In contrast both the damping rate and the stiffness increase continuously with static pressure over the pressure-range examined. The trends in these measurements agree with findings presented by other authors (see e.g. [56] [153]). In sections 7.4.3 and 7.4.4 we discuss the trends with respect to contact area in more depth.

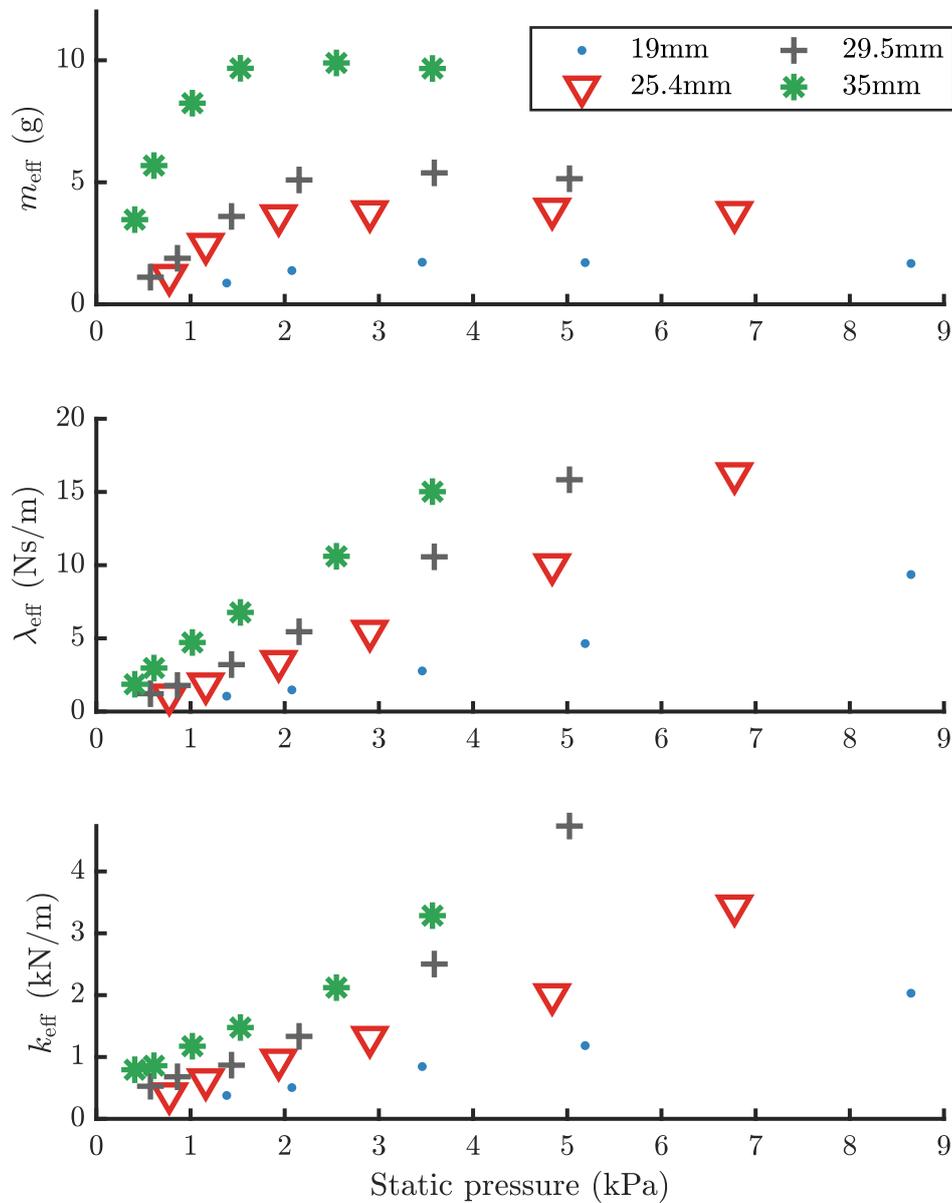


Fig. 7.9 Effective parameters for a SDOF model of the ‘chest’ impedance seen by rigid couplers with a range of contact areas, plotted against the static pressure applied through the coupler ($F_{\text{static}}/\hat{A}_c$). Data are shown for contact diameters of 19, 25.4, 29.5 and 35 mm. All measurements are conducted at the same location of the upper back of a human subject.

7.4 Impedance seen by rigid disc on phantom

In order to measure the impedance of the phantom seen by the sensor in a manner that is consistent with Thévenin and Norton's theorems the base of the phantom must be either fully constrained, or fully unconstrained. The implications of these choices are discussed in chapter 8. In this section we consider the case in which the base of the phantom is constrained. In chapter 8 to show that this is more appropriate for a comparison between the phantom and the chest.

7.4.1 Setup

Figure 7.10 shows the balance-rig (see also figures 7.2 and 7.3) used to measure the driving-point impedance seen by a rigid coupler pressed against the phantom. The phantom has been removed from the shaker-impedance-head setup at its base and has instead been firmly mounted to a fixed base. The impedance analogy circuit given in figure 7.4 can be used to represent this setup, with z_{chest} replaced by z_{phantom} . z_{phantom} is the impedance of the phantom seen by a rigid coupler of a given contact area applied with a given static load. The relationship between z_{phantom} and the descriptions of the phantom's behaviour presented in chapter 6 is discussed in section 7.4.5.

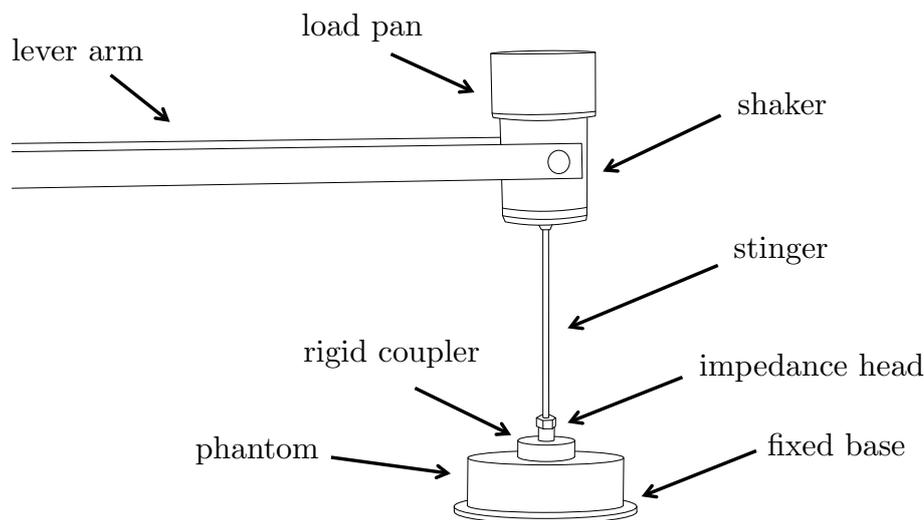


Fig. 7.10 Schematic representing the use of the balance rig for phantom impedance measurements.

7.4.2 Results

Figure 7.11 shows the effective mass, viscous damping rate and stiffness seen by rigid cylindrical couplers with three different contact areas coupled to the phantom as shown in figure 7.10. For each coupler measurements were taken with a range of static loads applied.

Comparing figure 7.11 to figure 7.9, the magnitudes of the terms for the phantom are in the appropriate range to provide a simple model for the chest. The effective mass, damping rate and stiffness all increase with increasing contact area. This matches the trends observed for the chest in both our measurements (see figure 7.9 and the literature (see e.g. [5] [56] [153] [154]).

One difference between the impedance measurements on the phantom and the chest is that the damping rate and stiffness of the phantom reach a plateau at a static pressure of around 4 kPa, while for the chest these parameters continue to increase at the same static pressure. We would expect the parameters for the chest to eventually plateau (see e.g. [5]); however, we were unable to observe this in our experiments for the range of static loads applied.

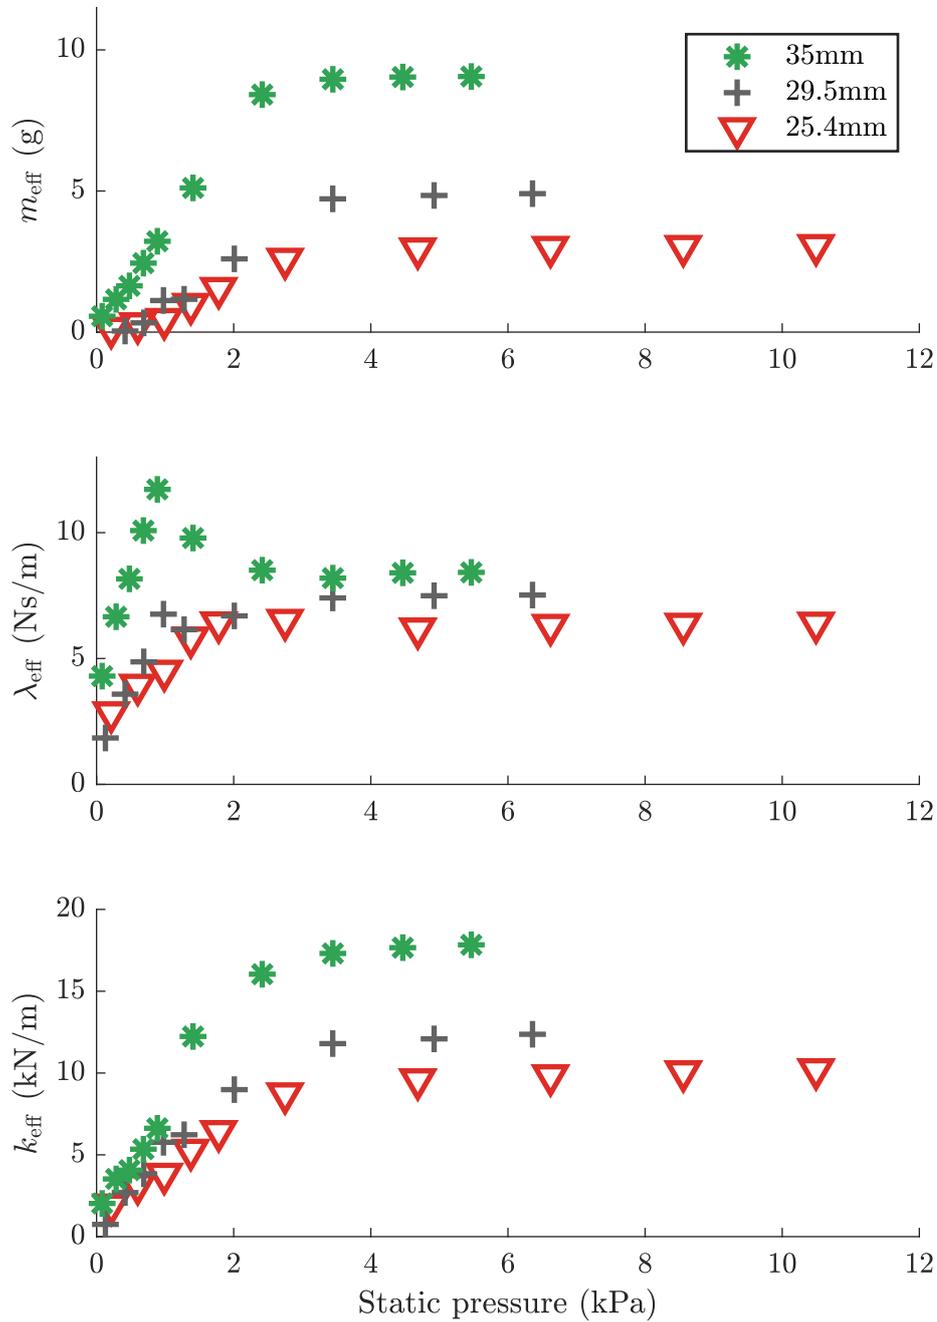


Fig. 7.11 Effective parameters for a SDOF model of the impedance of the phantom seen by rigid couplers with a range of contact areas, plotted against static pressure applied through the coupler. Data are shown for contact diameters of 25.4 mm (red triangles), 29.5 mm (grey plus signs), and 35 mm (green asterisks). All measurements are conducted at the same location of the upper back of a human subject.

7.4.3 Effective mass trends

Vermarien and van Vollenhoven [153] hypothesised that the effective mass of the chest was independent of preload for a given contact area, but were unable to verify this experimentally. In our results this is the case above a preload of around 2 kPa for the chest and around 3 kPa for the phantom. Vermarien and van Vollenhoven [153] further suggested a cubic relationship between contact diameter and effective mass:

$$m_{\text{eff}} = \hat{a}d^3 \quad (7.3)$$

where d is the coupler diameter and \hat{a} is a constant $\simeq 200 \text{ kg/m}^3$ [153].

In figure 7.12 the plateau levels of our effective mass measurements are compared with data presented by Vermarien and van Vollenhoven (1984) [153] and with equation 7.3. The effective mass seen by the sensor when it is coupled to the phantom is very close to the effective mass seen by the same sensor when it is coupled to the chest. The trend with diameter matches the d^3 trend suggested by Vermarien and van Vollenhoven [153], although more data are required to fully validate the trend. The experimental results of Vermarien and van Vollenhoven [153] yield a slightly higher effective mass for a given contact diameter (roughly corresponding to $\hat{a} = 200 \text{ kg/m}^3$), which is not surprising given the expected degree of inter-subject variability.

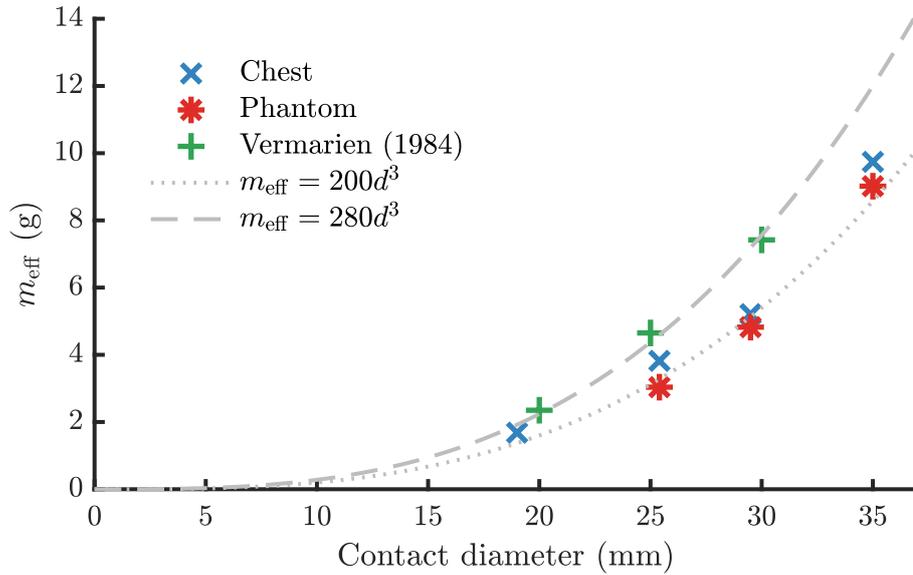


Fig. 7.12 Effective mass seen by a rigid coupler pressed against the phantom or the chest, plotted against the contact diameter. Blue crosses (\times): plateau value from chest experiments at a single location (see figure 7.9). Red asterisks ($*$): plateau value for phantom experiments (see figure 7.11). Green plus signs ($+$): data presented in table 2 of Vermarien van Vollenhoven (1984) [153] (page 173). Grey dotted line: equation 7.3 with $\hat{a} = 200 \text{ kg m}^{-3}$. Grey dashed line: equation 7.3 with $\hat{a} = 280 \text{ kg m}^{-3}$.

7.4.4 Effective stiffness trends

For a uniform elastic material with a constant cross-sectional area of \hat{A} the stiffness is given by $E\hat{A}/l$, where E is the Young's modulus of the material and l is the sample length. For the stiffness seen by a coupler of a certain area, we also expect the stiffness to scale with the contact area. In figure 7.13 the effective stiffness results from figures 7.9 and 7.11 are divided by the contact area to assess this hypothesis. For both the chest and the phantom the stiffness results from different contact areas on axes of k_{eff}/\hat{A} versus static pressure collapse onto a single curve. This confirms that for both the chest and the phantom the effective stiffness is proportional to the contact area.

The stiffness of the phantom is significantly higher than the example human-chest stiffness. Figure 7.13 also shows data from Vermarien and van Vollenhoven (1984) [153], converted into the appropriate format. For this data k_{eff}/\hat{A} lies between our values for the phantom and the chest.

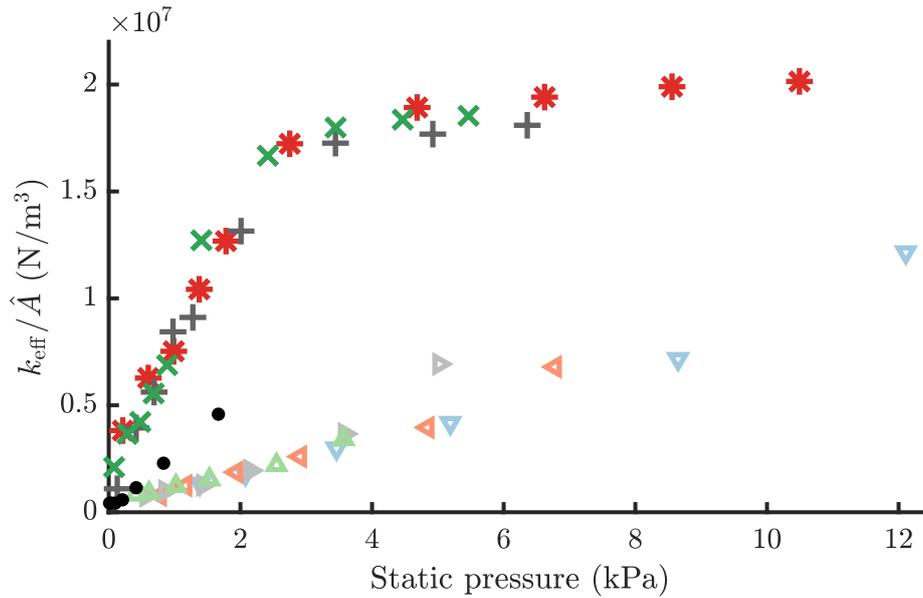


Fig. 7.13 Plot of k_{eff}/\hat{A} against the static pressure with which the coupler is applied to the surface for measurements on the upper back of a human subject (light shaded triangles) and on the phantom (dark shaded symbols). Data for a 30 mm diameter coupler from table 3 of Vermarien and van Vollenhoven (1984) [153] (page 176) is also converted into this format (black dots).

7.4.5 Comparison to transmission matrix for phantom

Figure 7.14 demonstrates the relationship between the results presented in chapter 6 and those discussed in this chapter. The impedance analogy circuit in figure 7.14a represents excitation at the base of the phantom, with the top surface of the phantom loaded by an unconstrained mass m_l . In figure 7.14b the base of the phantom is grounded (corresponding to an open circuit in the impedance analogy, see section 3.5.5). The load is excited by a test signal f_{test} , inducing a velocity \dot{x}_{test} . The ratio of F_{test} to \dot{X}_{test} is the driving point impedance approximated by the impedance head (Z_{test}). In terms of the impedance analogy circuit for the balance-rig measurement shown in figure 7.4, \dot{x}_{test} is \dot{x}_l and f_{test} is the force across “ z_{chest} ”.

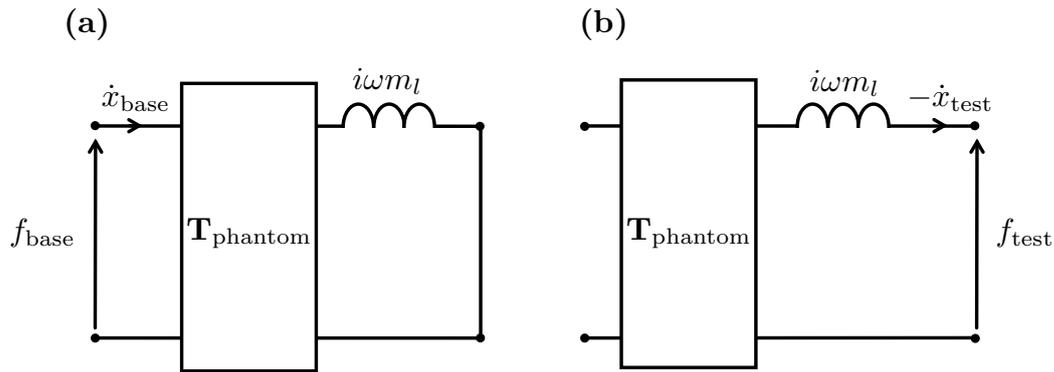


Fig. 7.14 Impedance analogy circuits for the measurement setups in (a) figure 6.1, and (b) figure 7.10. In both cases the measured signals are defined as positive *towards* the phantom.

An estimate for $\mathbf{T}_{\text{phantom}}$ is computed from measured FRF data in chapter 6. In order to compare this estimate to the impedance data in this chapter, we consider the matrix equation for the circuit in figure 7.14b:

$$\begin{bmatrix} F_{\text{base}} \\ 0 \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{phantom}}} \begin{bmatrix} F_{\text{test}} \\ -\dot{X}_{\text{test}} \end{bmatrix}. \quad (7.4)$$

The bottom line yields $0 = CF_{\text{test}} - D\dot{X}_{\text{test}}$, which allows the load impedance to be written as:

$$Z_{\text{test}} = \underbrace{D/C}_{\mathbf{T}_{\text{phantom}}} \quad (7.5)$$

Inspection of equation (7.5) makes it clear that there is a discrepancy between the phantom-properties derived in chapter 6 and the impedance data presented in this chapter. The terms in $\mathbf{T}_{\text{phantom}}$ do not depend on static load (for static loads corresponding to 60-150 g), while the measured phantom impedance does.

Figure 7.15 shows driving point impedance data from the balance-rig with loads of ~ 0.5 N and ~ 4.3 N, for a 51 g, 25.4 mm diameter rigid coupler on the phantom. The figure also shows a prediction for the value of this measurement made using the transmission matrix derived in chapter 6 with a load of 57 g (to account for the ‘mass-above-the-force-gauge’ in the measured data). The computed impedance matches the 4.3 N case from the balance-rig closely, despite the fact that the static load exerted by the coupler is ~ 0.56 N. This suggests that the low-load

data from the balance-rig should be used with caution. The characteristics of the contact may not accurately replicate the case in which a sensor lies on the phantom, loading the surface with only its self-weight. One possible explanation for this is that the coupler may not be perfectly aligned to the top surface of the phantom. If the coupler surface is not parallel to the phantom surface then at low loads the effective contact area is reduced leading to lower stiffness and damping rate and effective mass, than would be seen with an optimal contact.

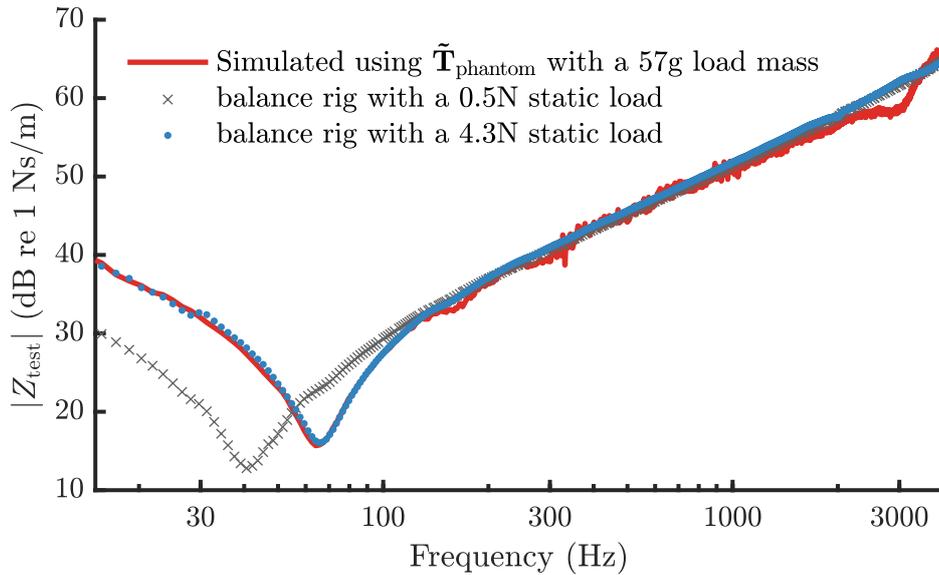


Fig. 7.15 Bode magnitude plot of the driving point mechanical impedance seen by an impedance head connected to a 25.4 mm diameter cylindrical coupler on the surface of the phantom. Balance rig measurements for a 51 g coupler are shown for two static loads (\times , \bullet). The simulated impedance computed using equation (7.5) with $\mathbf{T}_{\text{phantom}}$ and a 57 g load ($51 \text{ g} + m_o$) is shown by the solid red line.

7.5 Attachment impedance due to holding the stethoscope

Having explored the impedance of the chest seen by a sensor, we now consider the second term that contributes to the net impedance seen by the sensor in figure 7.1(a): the impedance of the attachment. While direct contact sensors are rarely attached to the chest by hand [146], air-coupled stethoscopes are almost exclusively used in this manner.

7.5.1 Driving-point mechanical impedance of the human hand-arm system

The mechanical impedance of the human hand-arm system has been extensively investigated in the context of vibration exposure when using handheld power tools [128, 27, 58]. The International Organization for Standardization (henceforth ‘ISO’) standard ‘*Mechanical vibration and shock - mechanical impedance of the human hand-arm system at the driving point*’ (ISO 10068:2012(E)) [58] presents guidelines on the range of expected values of the driving point impedance of the human hand-arm system as a function of frequency based on the results of a number of human-subject trials.

The ISO guidelines are based on a standardised experimental setup in which the hand grips a rigid bar (see ISO 8727:1997 [57]). The driving point impedance is independently defined in three orthogonal directions: the centroid of the forearm, the centroid of the rigid bar, and a third direction that is orthogonal to these. This third direction corresponds roughly to the surface-normal of the chest when holding a stethoscope. The applicability of the ISO guidelines is limited to cases where the hand grasps a handle between 19 mm and 45 mm in diameter with a grip force of between 25 N and 50 N [58].

The mechanical impedance of the hand-arm system depends on a number of factors including the geometry of the object held, the static force exerted by the hand, the posture of the hand and arm, and inter-subject variations in anatomy and tissue properties [58]. There are a number of different common techniques for gripping stethoscopes, but none closely resembles the grasping of a rigid bar. This means that the guidelines of ISO 10068:2012 cannot be directly applied.

Israr et al. consider the mechanical impedance of the human hand-arm system when gripping a stylus [59] and a ball [60] in the context of developing haptic systems. These grips are slightly more representative of how a stethoscope is held. Notably, Israr et al. found that the mechanical impedance for these two grips followed similar trends, with a resonance occurring between 160 and 320 Hz, a high frequency response dominated by mass, and a low-frequency response dominated by damping. This suggests that it is not necessary to measure the mechanical impedance of every possible stethoscope grip separately. Given the anticipated inter-subject variability, and the dependence on other factors such as

arm-posture and static load applied, we suggest that it is sufficient to characterise the mechanical impedance for a representative grip.

The results of Isar et al. [59, 60] are measured at seven discrete test frequencies between 10 and 500 Hz. In order to incorporate the effect of holding the stethoscope as an ‘attachment -impedance’ in our models a small number of discrete datapoints is insufficient. Instead we require a simple lumped element model to capture the role of the hand.

7.5.2 LEM for the hand-arm system

The ISO standard 10068:2012 [58] recommends two lumped element models for the hand-arm system. The 4DOF LEM recommended in the standard is specific to the bar-grip hold, but the 2DOF LEM presented (shown in figure 7.16) can be adapted for the stethoscope hold. At high frequencies the effect is dominated by an added-mass, while at low frequencies damping plays a significant role. The parameter values suggested by the ISO standard [58] are not expected to be applicable for the stethoscope-grip. In particular we expect the significantly smaller contact area between the hand and the object to lead to lower values of m_1 and $\lambda_{1,2}$.

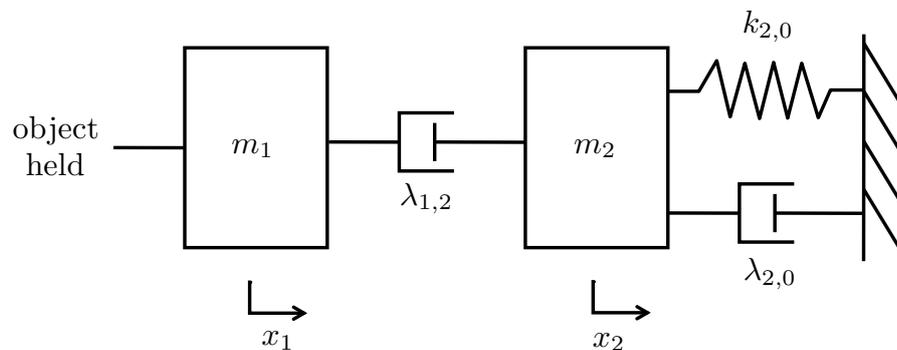


Fig. 7.16 ISO 10068:2012 2DOF LEM for the driving point mechanical impedance of the hand arm system in the x_h direction (as defined in [58] and [57]). The nomenclature for the lumped elements has been adapted for compatibility with this thesis. Values for each parameter are suggested in ISO 10068:2012 [58].

7.5.3 Experimental setup

In order to measure the attachment impedance due to the hand-arm system, we take the experimental setup described in section 5.3.2 and adapt it such that a human subject holds the load mass and applies a pressure which they subjectively describe as ‘light to moderate’. Figure 7.17 depicts a schematic of this experimental setup. The grip shown in figure 7.17 is used for practical reasons, but does not represent the most common technique for holding a stethoscope.

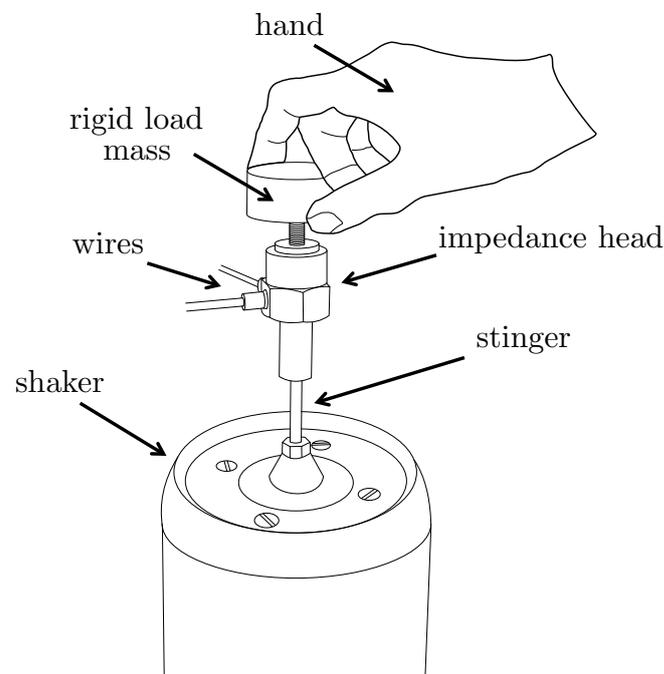


Fig. 7.17 Experimental setup for measuring the driving point mechanical impedance of the human hand-arm system when holding a coupler representing a stethoscope chestpiece.

7.5.4 Results

Figure 7.18 shows the driving point dynamic mass FRF for a 28 g coupler held as shown in figure 7.17. Seven cases are shown: five with the elbow supported (blue markers) and two with the elbow unsupported (green markers).

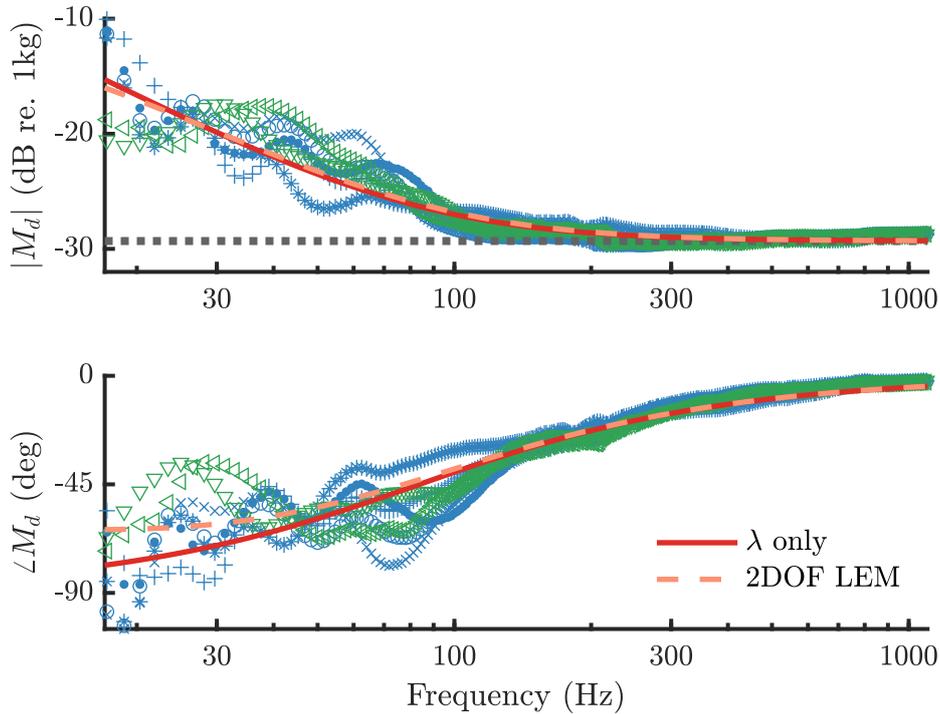


Fig. 7.18 Driving point dynamic mass FRF for a 28 g coupler held as shown in figure 7.17. Blue markers show data from measurements with the elbow resting on a support, while green markers are for an unsupported elbow. The grey dotted line represents the mass of the coupler (including the ‘mass-above-the-force-gauge’ of the impedance head). The solid red line shows a simple simulation in which the effect the hand-arm system is represented by a dashpot ($\lambda = 18$ Ns/m) to ground. The dashed light red line shows the result of a simulation in which the coupler is connected to the 2DOF LEM shown in figure 7.16, with the parameter values suggested in ISO 10068:2012 [58], except for $m_1 = 0$ and $\lambda_{1,2} = 18$ Ns/m.

At high frequencies there is little difference between FRFs for the held and free couplers. The added-mass effect due to holding the hand is negligible. The low frequency effect of holding appears to be damping dominated. The solid red line shows the results of a simulation in which the effect of holding is modelled using only a dashpot ($\lambda = 18$ Ns/m) to ground. (The value of $\lambda = 18$ Ns/m is obtained by considering the real part of the impedance FRFs, using the method described in section 7.2.1.) This simple model for the attachment impedance is able to capture the general shape of the response well. The measured responses exhibit small-scale resonances at low frequencies which are not captured by the model.

The measured FRFs for the unsupported elbow (shown by the green markers) seem to level off at low frequency, and have a phase that does not approach -90 degrees. This could be explained by considering a 2DOF LEM of the form suggested by ISO 10068:2012 [58] (see figure 7.16), with the effective contact-mass set to $m_1 = 0$ and the damping reduced to $\lambda_{1,2} = 20$ Ns/m.

7.6 Summary and conclusions

In this chapter we consider the mechanical impedance seen by a direct contact sensor on the chest. This impedance is composed of two components: the chest impedance and the attachment impedance. We develop experimental techniques to measure both of these terms.

The ‘balance rig’ presented in section 7.2 is used to measure the driving point impedance seen by a rigid coupler pressed against the chest or the phantom with a range of static loads. We demonstrate that the phantom is able to crudely replicate the mechanical properties of the chest seen by a rigid contact sensor. The effective mass, damping and stiffness of the phantom are all in the same order of magnitude as those of the chest, which further validates the use of the phantom. We show that the transmission matrix derived for the phantom in chapter 6 is able to accurately predict the impedance seen by a rigid coupler for applications pressures $\gtrsim 4$ kPa. At low static load the balance rig may not be fully aligned with the surface of the phantom leading to an lower effective contact area and consequently lower values for the extracted lumped parameters than are seen for a full contact. The data from which the transmission matrix for the phantom is computed comes from sensors balancing on top of the phantom and held in place by their self-weight. This setup is expected to yield a full contact. In chapter 8 we use the results presented in this chapter to investigate the coupling of direct sensors to the chest using an impedance coupling technique.

Considering the impedance load on a sensor from the hand and arm of the clinician holding it against the chest, we find that well established results in the literature for the the hand-arm impedance in the context of assessing the vibration load of hand-held power tools are not directly applicable to the holding of a stethoscope chestpiece. This is primarily due to significant differences in the geometry of the grip applied. In particular, we find that the stethoscope grip is

characterised by a lower effective contact-mass and a lower damping rate than the standardised bar-grip relevant to power tool usage.

We measure the driving point impedance of a rigid coupler held with light to moderate pressure and find that the dominant effect of the holding can be captured by modelling the hand as a dashpot to ground. This model can be refined by including additional degrees of freedom, but is sufficient for a first approximation of the effect. We observe significant variation in the driving point impedance for nominally identical grips, which suggests that a low order model is most appropriate.

Chapter 8

Direct contact sensors

In chapter 6, we demonstrate our ability to accurately model the response of a direct contact sensor (henceforth ‘DCS’) to excitation at the base of our chest-phantom using a transmission matrix derived from experimental FRF measurements. In this chapter, we investigate what these results can tell us about the likely performance of the same sensors on the chest. To this end, we make use of the impedance measurements presented in chapter 7.

8.1 Direct contact sensors on the phantom

Let us consider again the matrix equation relating the driving point at the base of the phantom to a sensor applied to the top surface of the phantom (as shown in figure 6.2):

$$\begin{bmatrix} F_b \\ \dot{X}_b \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{phantom}}} \begin{bmatrix} F_s \\ \dot{X}_s \end{bmatrix}. \quad (8.1)$$

F_b and \dot{X}_b are the force on and the velocity of the base of the phantom, while F_s and \dot{X}_s are the force on and the velocity of the contact point between the phantom and the sensor. For the remainder of this chapter A , B , C and D refer exclusively to the terms in $\mathbf{T}_{\text{phantom}}$, so the under-bracket notation is dropped for readability whenever the individual terms of the matrix are used in equations.

The load impedance is equal to the ratio of F_s to \dot{X}_s i.e., $z_s = Z_{s,s}^m := F_s / \dot{X}_s$. A typical direct contact sensor measures acceleration \ddot{X}_s . Equation (8.1) allows us to compute the ‘*phantom force response*’ of a DCS, which is defined as the response of the measured variable to the force at the driving point of the phantom:

$$\frac{\ddot{X}_s}{F_b} = \frac{i\omega}{Az_s + B} . \quad (8.2)$$

Likewise, the ‘*phantom acceleration response*’ is defined as the response of the measured variable to the acceleration at the driving point:

$$\frac{\ddot{X}_s}{\ddot{X}_b} = \frac{1}{Cz_s + D} . \quad (8.3)$$

Taking the ratio of the phantom force response to the phantom acceleration response yields the driving point accelerance:

$$H_{a:b,b} = \frac{\ddot{X}_b}{F_b} = \frac{\ddot{X}_s}{F_b} / \frac{\ddot{X}_s}{\ddot{X}_b} = i\omega \left(\frac{Cz_s + D}{Az_s + B} \right) . \quad (8.4)$$

All three of these FRFs are shown in figure 8.1 for four cylindrical DCSs with the same diameter of 25.4 mm, and masses ranging from 40 g to 150 g. The figure shows FRFs estimated from measurements (markers), and simulated using equations (8.2) to (8.4) (lines). The simulations are performed using a transmission matrix ($\mathbf{T}_{\text{phantom}}$) computed from experimental data for the 62 g and 152 g cases as described in section 6.5.1. To simplify the analysis we employ a slightly different form of the transmission matrix ($\tilde{\mathbf{T}}_{\text{phantom}}$) which directly relates the terms measured by the impedance head to the sensor (see section 6.5.1). The agreement between the simulation and the data is excellent between 62 and 152 g. However, for the 42 g case the simulation overestimates the second resonance frequency (ω_2), suggesting that the stiffness may be lower for this case. Using $\tilde{\mathbf{T}}_{\text{phantom}}$ to predict the 42 g case is an extrapolation, but has been included here for insight.

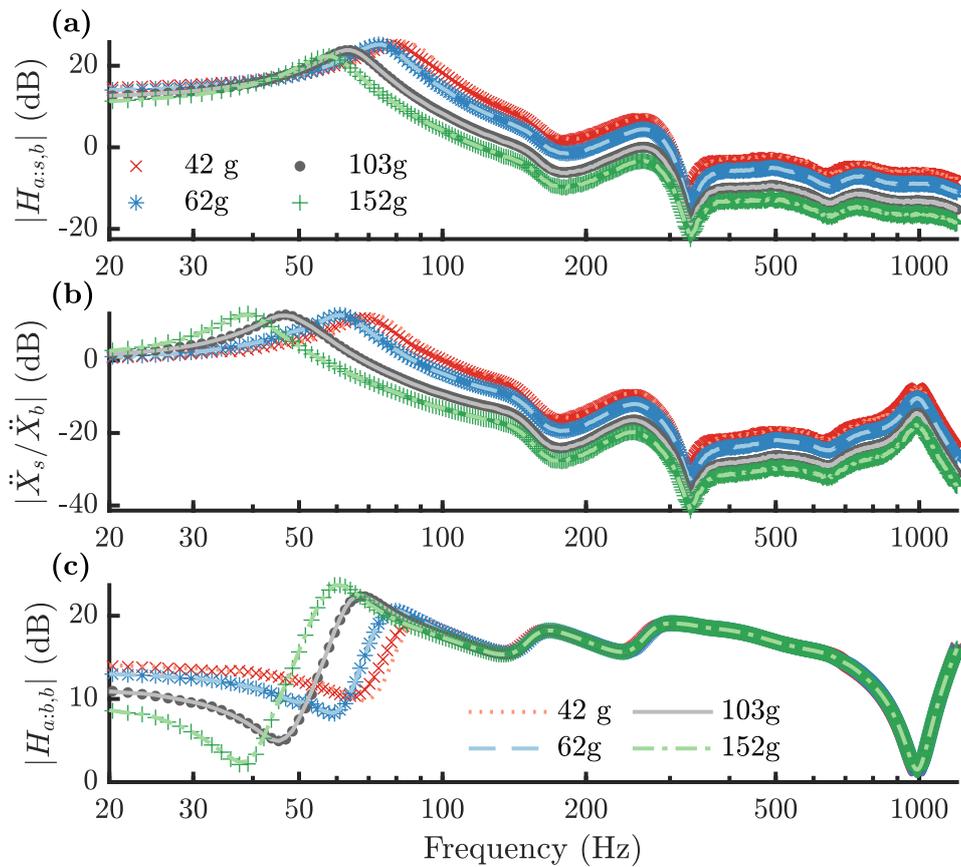


Fig. 8.1 Bode magnitude plots of FRFs for a 25.4 mm diameter rigid DCS on the phantom. Markers are used to show FRF estimates from measured data while lines show the results from simulations using equations (8.2) to (8.4) with $\tilde{\mathbf{T}}_{\text{phantom}}$. **(a)**: phantom force response ($H_{a,s,b}$). **(b)**: phantom acceleration response (\ddot{X}_s/\ddot{X}_b). **(c)**: driving point accelerance ($H_{a,b,b}$). The decibel reference for both accelerance FRFs is 1 kg^{-1} .

8.1.1 Stop-surface force reference

The response curves in figure 8.1 allow us to clearly distinguish between the performance of the sensors. However, the individual curves contain features that depend on the transmission of vibrations through the phantom and are therefore not relevant to the performance of the same sensors on the chest.

In order to assess the performance of the sensors in a more objective manner we follow the approach of Suzuki and Nakayama [142], and consider the response relative to an ‘equivalent’ *stop-surface force*. The term ‘equivalent’ needs to be

carefully defined, and we establish what we mean by this in the remainder of this section.

The *stop-surface force* is defined as the force perturbation exerted on a coupler (of a given area) with an infinite impedance. An infinite impedance load at the surface means that the motion of the surface is completely impeded, which explains the terminology used. Mathematically this ‘*stop-surface force*’ is equivalent to the Thévenin force (see section 3.5.7).

The stop-surface force cannot be measured simultaneously with the sensor-response. Instead we consider the ratio of the stop-surface force (F_{ss}) to a reference input signal at the base of the phantom (S_{ref}) and compare this to the ratio of the test-sensor response (S_{test}) to the same input signal:

$$\frac{S_{test}}{F_{ss}} = \frac{S_{test}}{S_{ref}} \bigg/ \frac{F_{ss}}{S_{ref}}. \quad (8.5)$$

The ratio F_{ss}/S_{ref} depends on both the contact area and the static load applied. As such it can either be used to compare only sensors with a given contact area and application pressure, or an arbitrary reference area and application pressure can be chosen.

Suzuki and Nakayama [142] attempt to physically measure the *stop-surface force* using a force transducer with a firm support. This is a difficult measurement to perform accurately (see section 4.1.4) and it is not clear if or how Suzuki and Nakayama [142] controlled the static load applied.

We instead determine the *stop-surface force* theoretically, using an experimentally deduced transmission matrix for the phantom. Despite the fact that we are now applying an infinite load impedance, while the transmission matrix was determined from the FRFS for two fairly low-mass couplers, we are not extrapolating. This is because the impedance load presented by the sensor can be independent of the static force applied and does not directly affect the mechanical properties of the phantom (or the chest). In our experimental setup the impedance of the coupler ($z_s = i\omega m_s$) and the static load ($F_{static} = m_s \times 9.81 \text{ m s}^{-2}$) are directly related but this does not have to be the case. A sensor of any mass can be applied with any application pressure.

Considering the matrix equation for the phantom-sensor system with the sensor velocity set to zero, i.e.:

$$\underbrace{\begin{bmatrix} F_b \\ \dot{X}_b \end{bmatrix}}_{\text{SS}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{phantom}}} \underbrace{\begin{bmatrix} F_s \\ 0 \end{bmatrix}}_{\text{SS}}, \quad (8.6)$$

we can obtain the ratio of the stop-surface force to the force at the base of the phantom as:

$$\underbrace{F_s/F_b}_{\text{SS}} = 1/A. \quad (8.7)$$

Similarly, the ratio of the stop-surface force to the base-acceleration is given by:

$$\underbrace{F_s/\ddot{X}_b}_{\text{SS}} = \frac{1}{i\omega} 1/C. \quad (8.8)$$

Depending on whether we treat the phantom base as a force-source or an acceleration-source we can define the *stop-surface force response* of the sensor using either equation (8.9) or (8.11). For a force-source we combine equations (8.2) and (8.7) to yield:

$$\underbrace{\frac{\ddot{X}_s}{F_{\text{SS}}}}_{\text{force source}} = A \frac{\ddot{X}_s}{F_b} \quad (8.9)$$

$$= \frac{i\omega}{z_s + B/A}. \quad (8.10)$$

For an acceleration-source we combine equations (8.3) and (8.8) to yield:

$$\underbrace{\frac{\ddot{X}_s}{F_{\text{SS}}}}_{\text{accel source}} = i\omega C \frac{\ddot{X}_s}{\ddot{X}_b} \quad (8.11)$$

$$= \frac{i\omega}{z_s + D/C}. \quad (8.12)$$

In practice, the base of the phantom is neither a force-source nor an acceleration source as both the driving point force and the driving point acceleration depend on the load applied at the surface of the phantom. However, this is unimportant for the analysis presented here as we are not trying to predict the response of sensors to the voltage input at the shaker in the phantom rig. Instead, we are

interested in the performance of sensors on the human chest. To determine whether a force-source or an acceleration-source is a more appropriate model we consider the Thévenin impedance ‘seen’ by a sensor on the phantom for each of these two cases and compare this to measurements of the impedance ‘seen’ by the same sensor on the human chest.

8.1.2 Thévenin impedance

The ‘*stop-surface force*’ is mathematically equivalent to a Thévenin equivalent force. This motivates the use of Thévenin and Norton equivalent circuits to model the impedance coupling of sensors to the phantom and chest. The Thévenin impedance is the impedance measured at the output terminals of a linear system when all of the source terms within that system are set to zero (see section 3.5.7). Consequently, the Thévenin impedance of the phantom depends on how the source is defined.

For a *force-source* (shorthand f.s.) at the base of the phantom, we set F_b to zero and apply test signals ($F_s = F_{\text{test}}$ and $\dot{X}_s = -\dot{X}_{\text{test}}$, as defined in figure 7.14b) at the sensor side. This yields:

$$\underbrace{\begin{bmatrix} 0 \\ \dot{X}_b \end{bmatrix}}_{\text{f.s.}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{phantom}}} \underbrace{\begin{bmatrix} F_s \\ \dot{X}_s \end{bmatrix}}_{\text{f.s.}}. \quad (8.13)$$

Defining $Z_{\text{test}}^m := F_{\text{test}}/\dot{X}_{\text{test}}$, the top line of equation (8.13) can be rearranged to give:

$$\underbrace{Z_{\text{test}}^m}_{\text{force source}} = \underbrace{B/A}_{\mathbf{T}_{\text{phantom}}}. \quad (8.14)$$

For an *acceleration-source* (shorthand a.s.) at the base of the phantom, we set \dot{X}_b to zero and apply test signals at the sensor side, yielding:

$$\underbrace{\begin{bmatrix} F_b \\ 0 \end{bmatrix}}_{\text{a.s.}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{phantom}}} \underbrace{\begin{bmatrix} F_s \\ \dot{X}_s \end{bmatrix}}_{\text{a.s.}}. \quad (8.15)$$

The top line of equation (8.13) can be rearranged to give:

$$\underbrace{Z_{\text{test}}^m}_{\text{accel source}} = \underbrace{D/C}_{\mathbf{T}_{\text{phantom}}} . \quad (8.16)$$

Figure 8.2 shows the simulated driving-point impedance seen looking at a rigid cylindrical coupler with a mass of 57 g and a contact diameter of 25.4 mm pressed against the phantom with a static force in the range of 0.4 – 1.5 N. As expected, the force-source case gives a low-frequency mass limit (+20 dB/decade), while the acceleration-source case gives a low-frequency stiffness limit (-20 dB/decade). Comparison to the form of the driving-point impedance FRF seen looking at a sensor pressed against human-chest (shown in figure 7.6 of chapter 7) tells us that an acceleration-source is a more appropriate model for the base of the phantom in order to match the chest.

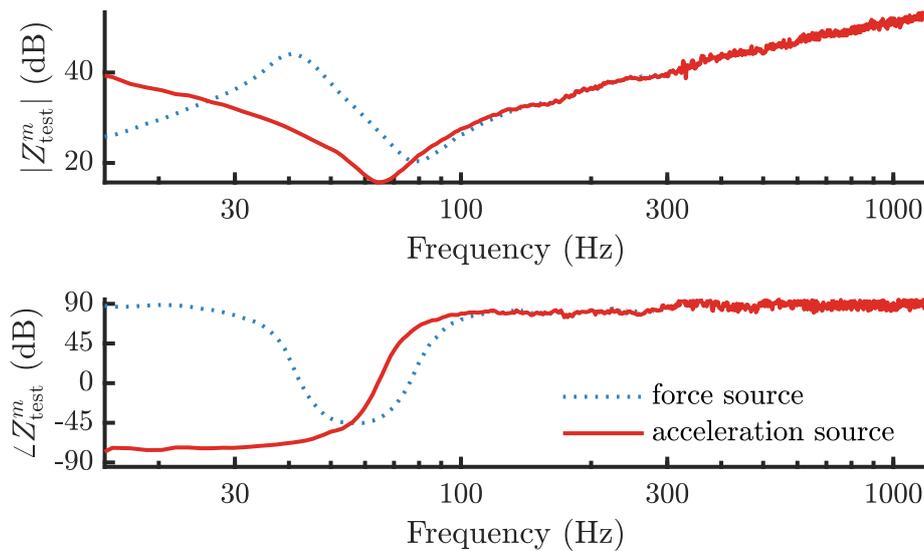


Fig. 8.2 Bode plot of the simulated impedance seen looking at a rigid cylindrical coupler with a mass of 57 g and a contact diameter of 25.4 mm pressed against the phantom with a static force in the range 0.4 – 1.5 N. Red line: force source model (equation (8.14)). Blue line: acceleration source model (equation (8.16)). The dB reference value is 1 N s m^{-1} .

It can also be seen that, despite the resonances associated with transmission through the phantom in e.g. figure 6.5, the ‘test’ impedance seen at the coupler resembles that of a two-degree-of-freedom system.

8.2 Impedance coupling

Defining the ‘test impedance’ for an acceleration-source at the base of the phantom as the Thévenin impedance, i.e.,

$$z_{\text{Thé}} = \underbrace{Z_{\text{test}}^m}_{\text{accel source}}, \quad (8.17)$$

equation (8.12) becomes:

$$\frac{\ddot{X}_s}{F_{ss}} = \frac{i\omega}{z_s + z_{\text{Thé}}}. \quad (8.18)$$

Equation 8.18 can be represented by the Thévenin-equivalent impedance analogy circuit in figure 8.3.

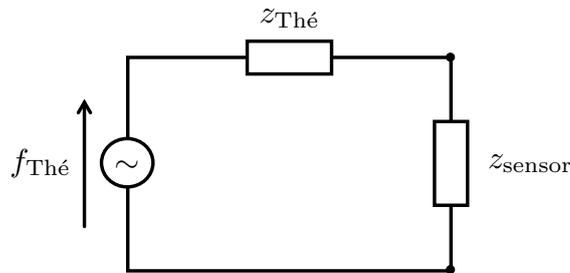


Fig. 8.3 Thévenin equivalent impedance analogy circuit for a rigid DCS on the chest. The Thévenin force ($F_{\text{Thé}}$) is equivalent to the ‘*stop-surface force*’ (F_{ss}).

The *stop-surface force response* is inversely related to the net impedance ($z_{\text{net}} = z_s + z_{\text{Thé}}$). Figure 8.4 depicts a graphical representation of the impedance coupling between the phantom and a sensor. The sensor is approximated as a pure mass, while the phantom impedance is given by a SDOF LEM with parameter values taken from the plateaus of effective mass, damping rate and stiffness in figure 7.11.

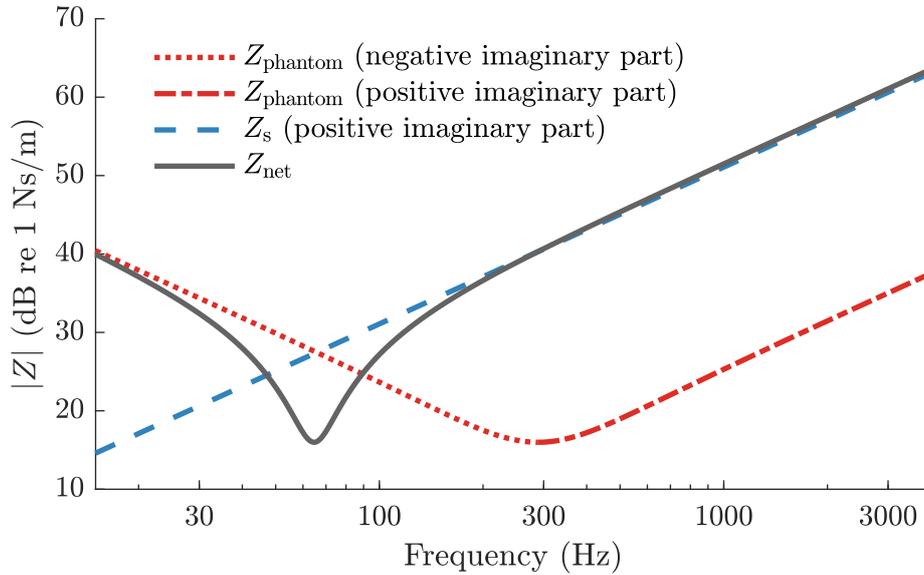


Fig. 8.4 Impedance coupling of a 57 g, 25.4 mm diameter DCS to the phantom. A SDOF LEM is used for the phantom with: $m_{\text{eff}} = 3.0$ g, $\lambda_{\text{eff}} = 6.3$ Ns/m and $k_{\text{eff}} = 9.9$ kN/m. The dB reference value is 1 Ns/m.

8.3 Thévenin force referenced phantom response

Figure 8.5 shows the *stop-surface force response* for two sensors on the phantom. The markers (+ and ×) show the response derived by converting the *phantom acceleration response* into a *stop-surface force response* using equation (8.11). The *phantom response* measurements are described in chapter 6 (see figure 6.2), as is the computation of the transmission matrix used in equation (8.11).

The lines in figure 8.5 are derived using the impedance coupling approach described in this chapter. The Thévenin impedance of the phantom is modelled using a SDOF LEM with parameter values extracted from the measurements in chapter 7 (see figure 7.10 for the experimental setup, figure 7.11 for the extracted SDOF parameters and figure 8.4 for a graphical representation of the impedance coupling).

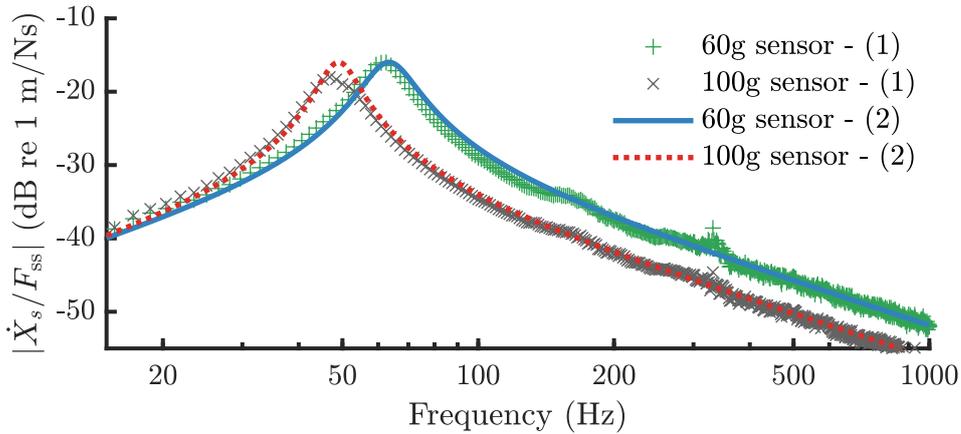


Fig. 8.5 Bode magnitude plot of the *stop-surface force response* (shown as an admittance FRF). Results are shown for two methods: (1) conversion of the phantom acceleration response using equation (8.11); (2) impedance coupling using a SDOF LEM for the Thévenin impedance of the phantom (see text). The dB reference value is $1 \text{ m s}^{-1} \text{ N}^{-1}$.

Considering the fact that the responses derived using an impedance coupling approach (+ and ×) and those derived by converting the phantom-response (••• and —) are based on measurements performed on the phantom-sensor system in an entirely different manner, the agreement between the two methods is good. The impedance coupling method predicts a slightly higher resonance frequency, and a peak height that is independent of the sensor mass, while the converted phantom acceleration response suggests that the peak amplitude decreases with sensor mass. This suggests that damping increases slightly with static load, contrary to the assumptions made in the computation of $\tilde{\mathbf{T}}_{\text{phantom}}$.

Comparing the performance of the two sensors, it can be observed that a higher sensor mass leads to a lower high-frequency response, and a lower resonance frequency. The low frequency response of both sensors is similar, while the sensor with a higher mass performs better in a small frequency range around its resonance. These observations are consistent with experimental results on the relative response of accelerometers on the chest presented by Vermarien and van Vollenhoven [153].

8.4 Impedance coupling to the human chest

Sections 8.1 to 8.3 establish a method for presenting data from the phantom in a manner that is not excessively characterised by the particularities of the measurement setup. However, we still need to interpret these measurements in light of the expected response at the chest.

In order to do this we make use of the impedance coupling approach described in section 8.2. In chapter 7 we present measurements of the chest impedance on the back of a human subject. Unlike the plateau region observed for the impedance seen by a sensor on the phantom, the chest exhibits a static-load dependence for both the effective damping rate and the effective stiffness across the entire range of loads for which measurements were performed. Studies reported in the literature observe similar trends.

To convert from the Thévenin-force response on the phantom to the Thévenin-force response on the chest, a conversion factor can be used:

$$\underbrace{\frac{\ddot{X}_s}{F_{ss}}}_{\text{chest}} = \left(\frac{1 + z_{\text{ph}}/z_s}{1 + z_{\text{ch}}/z_s} \right) \underbrace{\frac{\ddot{X}_s}{F_{ss}}}_{\text{phantom}}. \quad (8.19)$$

In equation 8.19 $z_{\text{ph}} = z_{\text{ph}}(\hat{A}, F_{\text{static}})$ is the Thévenin impedance of the phantom for the same contact area (\hat{A}) and load (F_{static}) for which $\left(\frac{\ddot{X}_s}{F_{ss}}\right)_{\text{phantom}}$ is valid and $z_{\text{ch}} = z_{\text{ch}}(\hat{A}, F_{\text{static}})$ is the Thévenin impedance of the chest for the same contact area and a desired static load.

Vermairen and van Vollenhoven [153] consider the coupling of direct contact-sensors to the chest by assuming that the static load applied is directly related to the mass of the sensor. This is only the case if the sensors are loading the chest vertically with their self-weight. This is the case for the measurements we have performed on the phantom, but is not typically the case for the practical use of direct-contact sensors on the chest. In figure 8.6 the effect of using the same DCS on the phantom and the human chest is shown for a case in which the static load is independent of the sensor mass.

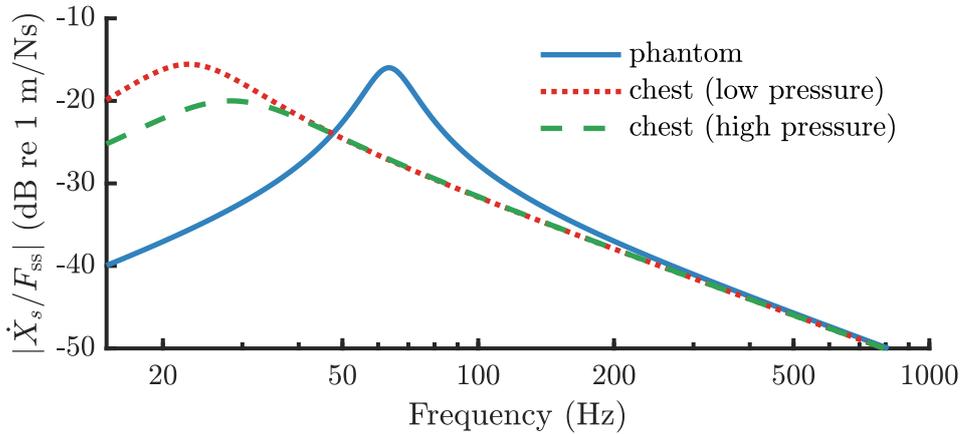


Fig. 8.6 Bode magnitude plot of the *stop-surface force response* for a 60 g, 25.4 mm diameter DCS (shown as an admittance FRF). The curves are simulated using a Thévenin-equivalent impedance coupling method with SDOF LEMs for both the phantom and the chest. The phantom case shown here matches that in figure 8.5. The response on the chest is simulated for two different contact pressures (3 kPa and 5 kPa), with parameter values extracted from the 25.4 mm diameter case in figure 7.9.

8.5 The effect of holding a direct contact sensor

Figure 8.7 shows the effect of holding the sensor on the *phantom force response*, the *phantom acceleration response* and the driving point admittance FRF. The grey dots show the free case for a 100 g sensor, while the green markers show the FRFs corresponding to three separate held cases (with the elbow resting on a support). The solid blue line shows the response of a free coupler simulated using $\tilde{\mathbf{T}}_{\text{phantom}}$ with a pure-mass load to represent the sensor. The dashed red line shows the simulated response for a coupler of the same mass linked to ground with a dashpot of $\lambda_{s,0} = 18 \text{ Ns/m}$. This is the expected attachment impedance according to the measurements and model presented in section 7.5. This simple dashpot-only model captures the effect of the loading reasonable well, although the value of $\lambda_{s,0} = 18 \text{ Ns/m}$ suggested in chapter 7 slightly overestimates the damping see for the held DCS.

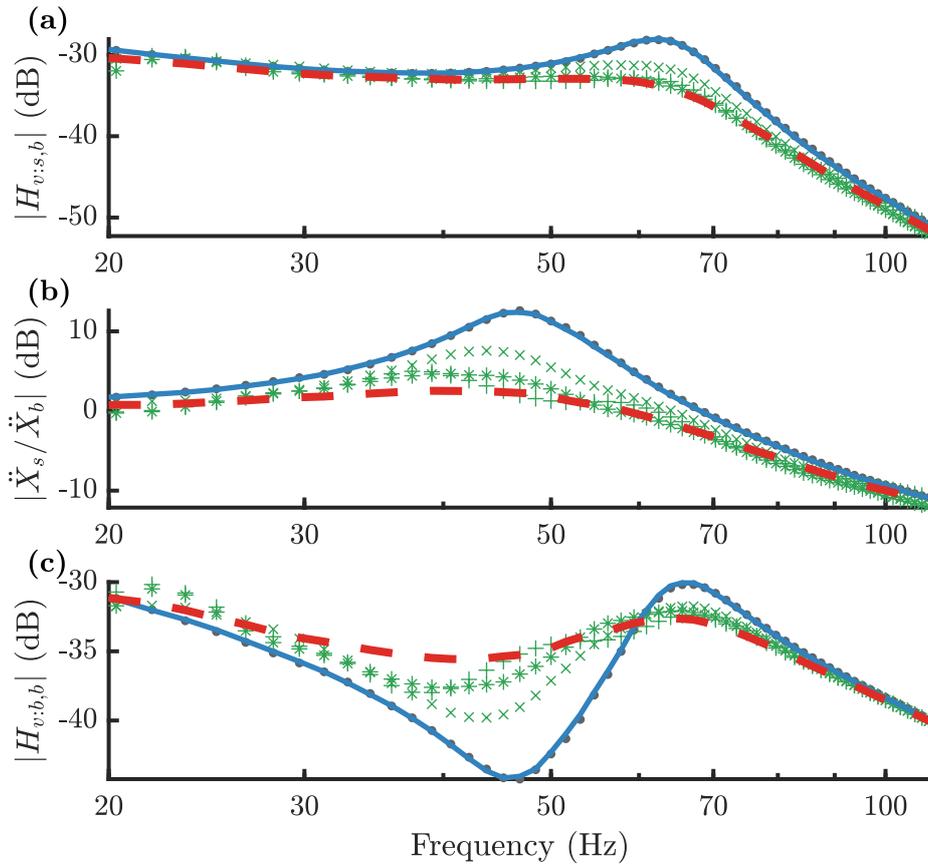


Fig. 8.7 Bode magnitude plots of H1 estimators of the FRFs for a 25.4 mm diameter rigid DCS on the phantom. **(a)**: phantom force response ($H_{a:s,b}$). **(b)**: phantom acceleration response (\ddot{X}_s/\ddot{X}_b). **(c)**: driving point accelerance ($H_{a:b,b}$). The decibel reference for both acceleration FRFs is 1 kg^{-1} . In all three plots grey dots (\bullet) denote the measured free case, green markers (\times , $+$, $*$) denote the measured held case, the blue line (—) denotes the simulated free case, and the dashed red line (-- --) denotes the simulated held case, using a λ -only model with $z_{s,0} = \lambda_{s,0} = 18 \text{ Ns/m}$. The dB reference for both admittance FRFs is 1 m/Ns .

8.6 Summary and conclusions

In this chapter, we demonstrate the use of impedance coupling techniques to predict the performance of direct contact sensors. We present a method for converting response measurements on the phantom into a format that is directly compatible with an impedance coupling approach by making use of Thévenin's

theorem. This allows us to estimate the relative performance of sensors on the chest from measurements of their performance on the phantom.

We concentrate on the physical coupling between direct contact sensors and the chest. In practice, the physics of the transducers plays an important role. For example, the surface acceleration might be maximised at the lowest possible load, but low-mass accelerometers typically have a lower sensitivity than heavier sensors. Much of the current research activity in this field is focused on the design of low-mass, high-sensitivity direct contact sensors (see e.g. [135]).

In contrast to the unloaded chest surface, a surface loaded with a DCS will move with a lower amplitude at high frequencies. As the sensor mass increases, the high-frequency performance degrades. However, the coupled chest–DCS system will have a lower first resonance frequency than the unloaded chest, predominantly due to the mass of the sensor. In the vicinity of this resonance the response is increased for the DCS case in contrast to the unloaded surface. This means that at these low frequencies an accelerometer on the chest will be moving at a higher amplitude than the unloaded surface would for the same (biological) excitation signal. The model presented in this chapter cannot be employed to directly compare the response of a free surface to a loaded one, as the model inherently includes the effect of constraining the sensor contact area to move as one - an assumption which is violated by a free surface.

Finally, we investigate the effect that holding a DCS has on the measured response. We show that the simple damping-only model for the attachment impedance due to holding proposed in chapter 7 is able to explain the changes in the measured FRFs.

Chapter 9

Air-coupled sensor acoustics

In this chapter, we present experimental measurements of the performance of air-coupled stethoscopes on the phantom described in chapter 6. The implications of these measurements are discussed in terms of the models derived in chapter 4. This chapter considers the specific case of ‘air-coupled microphones’ (henceforth ‘ACM’s). An ACM is an electronic stethoscope consisting of a chestpiece with a microphone mounted directly in the air cavity. The effect of tubing on the response is discussed in chapter 10.

This chapter starts with a description of experiments conducted on the phantom to characterise the performance of ACMs. In section 9.2, we present results for ACM performance on the phantom and discuss the effect of the chestpiece mass and the air cavity volume.

In section 9.3, we explore the ACM-performance data in more depth. We calculate the motion of the ‘skin’ surface in contact with the ACM air cavity from measured FRF data, and describe the use of a transmission matrix for the phantom in the interpretation and modelling of the ACM phantom-response.

In section 9.5, we consider the response of ACMs relative to the ‘stop-surface force’ on a rigid surface with the same outer contact area as the ACM. Modelling the transmission of vibrations from the base of the phantom to the ACM in a simplified manner allows us to compare these measured stop-surface responses with simulations of the Thévenin equivalent response according to the lumped-element ACM model presented in chapter 4. In section 9.6 we apply our model to explore the effect of cavity volume on the ACM-response, and in section 9.7 we

discuss the effects of using a diaphragm to cover the air cavity entrance. Finally, in section 9.8, we explore the effect of holding an air-coupled sensor.

9.1 Experimental setup

In order to explore how the performance of air-coupled microphones depends on cavity-volume we 3D-printed a range of chestpieces with different cavity volumes. The chestpieces were 3D-printed in Objet VeroWhitePlus RGD835 on an Objet24 3D printer to a variety of geometric specifications. For the results presented in this chapter, the outer contact area of the chestpiece and the cavity-contact area were kept constant while the cavity volume was varied. In each case, the chestpiece cavity was terminated by the flush-mounted protective grid of a 1/4 inch condenser microphone (G.R.A.S. 46BL).

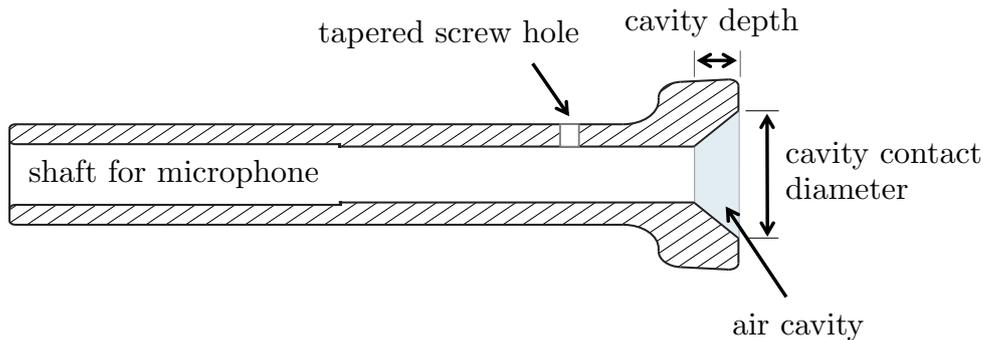


Fig. 9.1 Cross-section of 3D-printed air-coupled microphone chestpiece. The hollow shaft is filled by a microphone, with the microphone diaphragm flush to the air cavity as shown in figure 9.3. The cavity contact diameter and the cavity depth can be varied systematically to generate a set of ACMs.

Figure 9.2 shows a schematic of the experimental setup used to determine the *phantom response* of air-coupled sensors. The chestpiece rests on the top surface of the phantom and a firmly attached load mass is used to ensure the chestpiece is securely positioned and set the chestpiece mass to a desired value. An accelerometer on the load mass measures the motion of the chestpiece housing, while a microphone measures the pressure in the air cavity. Care is taken to avoid excessive air-leaks and to ensure that the setup is vertically aligned.

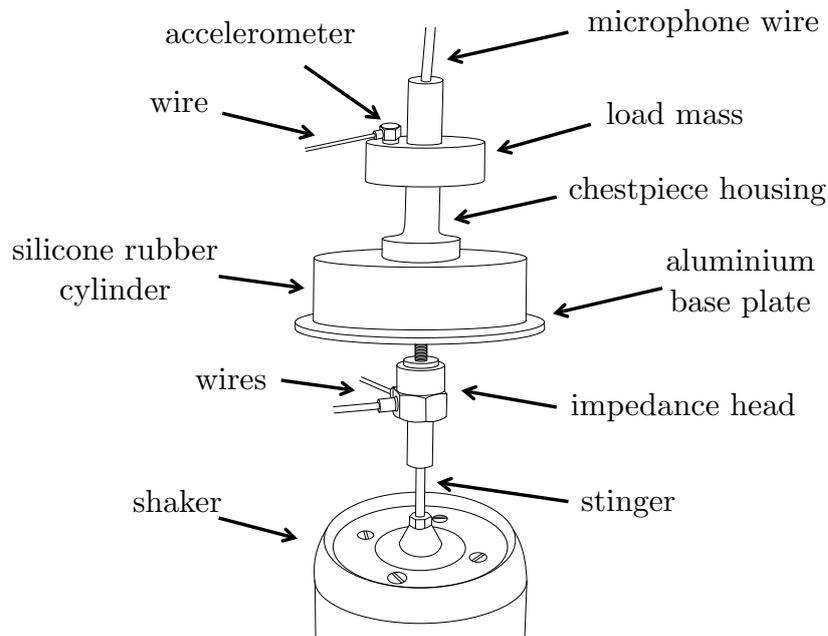


Fig. 9.2 Schematic of the experimental setup used to characterise the performance of air-coupled sensors on the phantom. Four signals are measured: the force and acceleration at the impedance head, the acceleration of the sensor housing, and the pressure in the air-cavity.

Figure 9.3 shows a cross-section of the chestpiece from figure 9.2. The shaft of the chestpiece fits tightly around the protective grid of the microphone to prevent air leakage, but the rear vent is not blocked. Any remaining air leaks would have cross-sectional dimensions several orders of magnitude smaller than the dimensions of the air cavity and are expected to only affect the response at very low frequencies (see e.g. Kraman et al. [78]), thus preventing the cavity from building up a static pressure without influencing the pressure perturbations at frequencies of interest. See appendix C.2 for more detail on the effect of air leaks. A plastic screw is used to secure the microphone in place by tightening the screw against the microphone through the tapered screw hole indicated in the schematic.

Microphones are designed to operate with a stationary base and a diaphragm that is exposed to a fluctuating pressure level. The conversion of diaphragm motion to pressure perturbation is based on the assumption that the microphone is operating in this condition. In this experiment this assumption is violated as the microphone base vibrates with the sensor housing. It can be shown (see appendix C.3) that axial vibration of the microphone-base results in an effective pressure term that is proportional to the base-acceleration. In appendix C.3 it is

shown that for the acceleration magnitudes measured in our experiments this term is negligible and does not need to be considered in the following analysis.

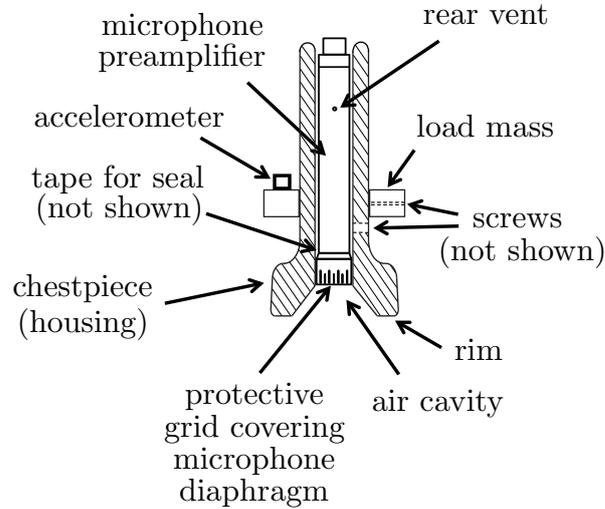


Fig. 9.3 Cross-section of the chestpiece in figure 9.2 with instrumentation.

To ensure that background noise did not corrupt our cavity-pressure measurements, we performed measurements with no excitation to determine the background noise level on all transducers. For frequencies above 20 Hz the background noise level on the microphone was found to be negligible in all measurement sessions. We also recorded the ambient sound level during the recordings with a calibrated room-microphone. This was found to be negligible when compared with sound level in the air-cavity.

9.2 Phantom-response of air-coupled microphones

Figure 9.4 shows the response of two ACMs with different cavity volumes on the phantom. For each ACM three different load masses were applied. The figure shows the response of the pressure in the air-cavity to the force (a) and acceleration (b) at the base of the phantom. The driving-point admittance FRF at the base of the phantom is also shown (c). Figure 9.4 reveals several important trends in ACM performance which are outlined in sections 9.2.1 and 9.2.2.

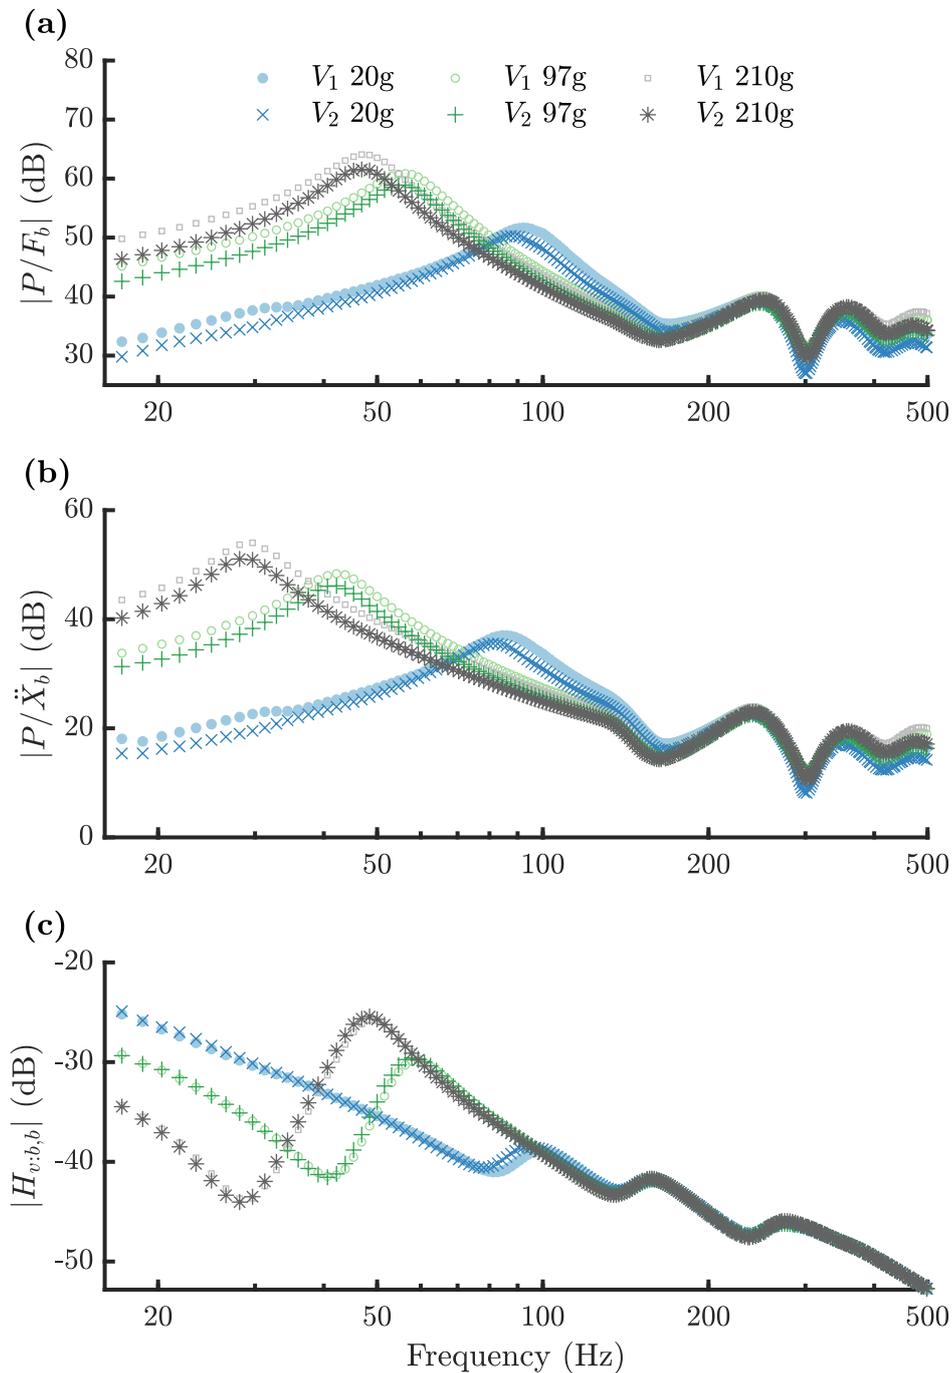


Fig. 9.4 Bode magnitude plots for FRFs from ACM tests on the phantom: (a) H1 estimator of the cavity-pressure with respect to the base-force (dB reference 1 m^{-2}), (b) H1 estimator of the cavity-pressure with respect to the base-acceleration (dB reference 1 kg m^{-2}), and (c) the driving point admittance at the base of the phantom (dB reference $1 \text{ m s}^{-1} \text{ N}^{-1}$). Results are shown for two cavity volumes: $V_1 = 1240 \text{ mm}^3$ and $V_2 = 2220 \text{ mm}^3$. For both ACMs the cavity-contact diameter is 16 mm and the outer diameter is 23.8 mm. Net chestpiece masses are as stated in the legend.

9.2.1 Effect of chestpiece mass (and static load)

The mass of the chestpiece has a significant effect on the low frequency performance of ACMs, while the high frequency response ($\gtrsim 200$ Hz) is almost unaffected by chestpiece mass. This can be explained by the fact that the large (relative) mass of the chestpiece causes it to lock at high frequencies, thus leading it to act like a grounded body regardless of the exact value of the mass. An inspection of the power spectral density (PSD) of the chestpiece acceleration (not shown here) confirms this hypothesis: the chestpiece acceleration becomes negligible at high frequencies, hitting the noise floor at around 1 kHz in most measurements (depending on the amplitude of the excitation).

At low frequencies the coupling of the chestpiece mass to the stiffness-dominated phantom plays a dominant role in setting the first antiresonance and second resonance frequencies of the coupled system. This matches our observations for direct contact sensors in chapter 6. The driving-point admittance FRFs for ACMs closely resemble those for DCSs, suggesting that the base of the phantom sees a similar load whether the surface is loaded by an ACM or a DCS.

As the chestpiece mass is increased, both the ω_2 resonance and the ω_{-1} antiresonance decrease in frequency and increase in amplitude. The ω_2 resonance determines the peak in the response of the cavity pressure to the ‘input’ force $F_{k_{b,o}}$, while the ω_{-1} antiresonance determines the peak in the response of the cavity pressure to the ‘input’ acceleration \ddot{X}_b . In both cases, increasing the sensor mass increases the level of the low-frequency response, while the mid-frequency response drops, and the high-frequency response is minimally affected.

As well as increasing the chestpiece mass, adding load masses to ACMs in our experimental setup also results in a larger static load applied to the phantom. The stiffness between the chestpiece rim and the phantom ‘tissue’ is expected to increase as the rim is pushed more firmly into the phantom (see e.g. Kaniusas’s [67] comments on the ‘natural diaphragm effect’). Increased stiffness tends to increase the frequency of resonances, but in this case that effect is outweighed by the effect of the increased mass.

As discussed in section 4.3.2, several existing theories for ACM acoustics entirely neglect the effect of the chestpiece mass on the response (see e.g. [163, 142]). The experimental results presented here confirm our hypothesis that this approach is not valid when considering the low-frequency response of ACMs.

9.2.2 Effect of cavity volume

Figure 9.4 demonstrates that increasing the cavity volume of ACMs reduces the level of the response *and* reduces the resonance frequencies. The reduction in the level of the response can be explained by the fact that when the cavity volume is larger, the same displacement perturbations at the cavity surface lead to smaller relative volume changes and therefore smaller pressure changes. The reduction in the resonance frequencies suggests that the stiffness of the air cavity plays a role in setting the resonance frequencies of the coupled system.

The volumes stated in figure 9.4 are the nominal cavity volumes. These do not include the small volume of air between the protective grid of the microphone and the diaphragm, and do not take into account the effect of occlusion. As the static load is increased we expect the actual cavity volume to decrease slightly as phantom ‘tissue’ fills the cavity.

9.3 Interpretation of the phantom-response of air-coupled microphones

Response curves of the type shown in figure 9.2 can be used to compare the performance of stethoscopes on the phantom and observe the effects of varying various design parameters. However, in order to explain these results physically, validate our model and predict the performance of the same ACMs on the chest we must extract more information from the data. In this section we present a range of semi-empirical results derived by combining measured response data with our theoretical models for the chest-sensor system. In subsection 9.3.1 we consider the forces exerted on the chestpiece of an ACM and compare these to the forces exerted on a DCS and on an infinite impedance load. In subsection 9.3.2 we explore what the pressure and housing-acceleration measurements can tell us about the motion of the constrained skin surface, by treating the small air cavity as a pure compliance.

9.3.1 Comparison to the DCS case

Figure 9.5 presents a comparison between the phantom-response of an ACM and that of a DCS with the same mass (67 g) and outer diameter (23.8 mm). Figure 9.5(a) shows the force exerted on each sensor with respect to the acceleration of the base of the phantom (F_s/\ddot{X}_b). As the force exerted on the sensors is not directly measured, these are semi-empirical results. For both the DCS (\times) and the housing of the ACM (\bullet), F_s/\ddot{X}_b is computed by multiplying the H1 estimator for \ddot{X}_s/\ddot{X}_b by the known sensor mass. The impedances of the transducer wires are neglected in this analysis. The stop-surface force ($+$) (computed as described in chapter 8) is shown for later reference.

For the ACM, figure 9.5(a) also shows an estimate of the force exerted through the air-cavity (\circ). This is obtained by multiplying the measured cavity pressure by the cavity-contact area ($F = P\hat{A}_c$). It can be seen that this accounts for only a small fraction of the total force exerted on the sensor housing. This observation suggests that Joyashiki and Wada's (2020) [66] model for ACMs, which assumes that the entire force on the housing mass is exerted through the air-cavity, is not valid (see section 4.3.2).

Figure 9.5(b) shows the phase of F_s/\ddot{X}_b . Both the DCS and the ACM housing are in phase with the base of the phantom at low frequencies. This agrees with our interpretation of the $\omega_1 = 0 \text{ rad s}^{-1}$ mode and the ω_{-1} antiresonance in the driving point response (see chapter 6). As the magnitude of F_s/\ddot{X}_b (9.5(a)) passes through the dominant peak, the phase shifts to 180° . This agrees with our interpretation of the driving point response (9.5(c)) in which we associate the second resonance (ω_2) to a mode in which the sensor moves out of phase with the base (see section 6.3). The pressure in the air cavity leads the base (and housing) acceleration by almost 90° at low frequencies. The phase difference between the housing acceleration and the cavity pressure then decreases until around 250 Hz. Above 250 Hz there is a small spike of increased phase difference before which the cavity pressure begins to lag the housing acceleration.

Figure 9.5(c) shows the driving point accelerance FRF ($H_{a,b,b}$) at the base of the phantom for the same three cases. Comparing $H_{a,b,b}$ for the DCS and ACM cases, the general features are similar, but in the ACM case the resonances occur at lower frequencies. This suggests that the overall oscillation of the coupled system is similar for the two cases, but that there are lower stiffnesses somewhere

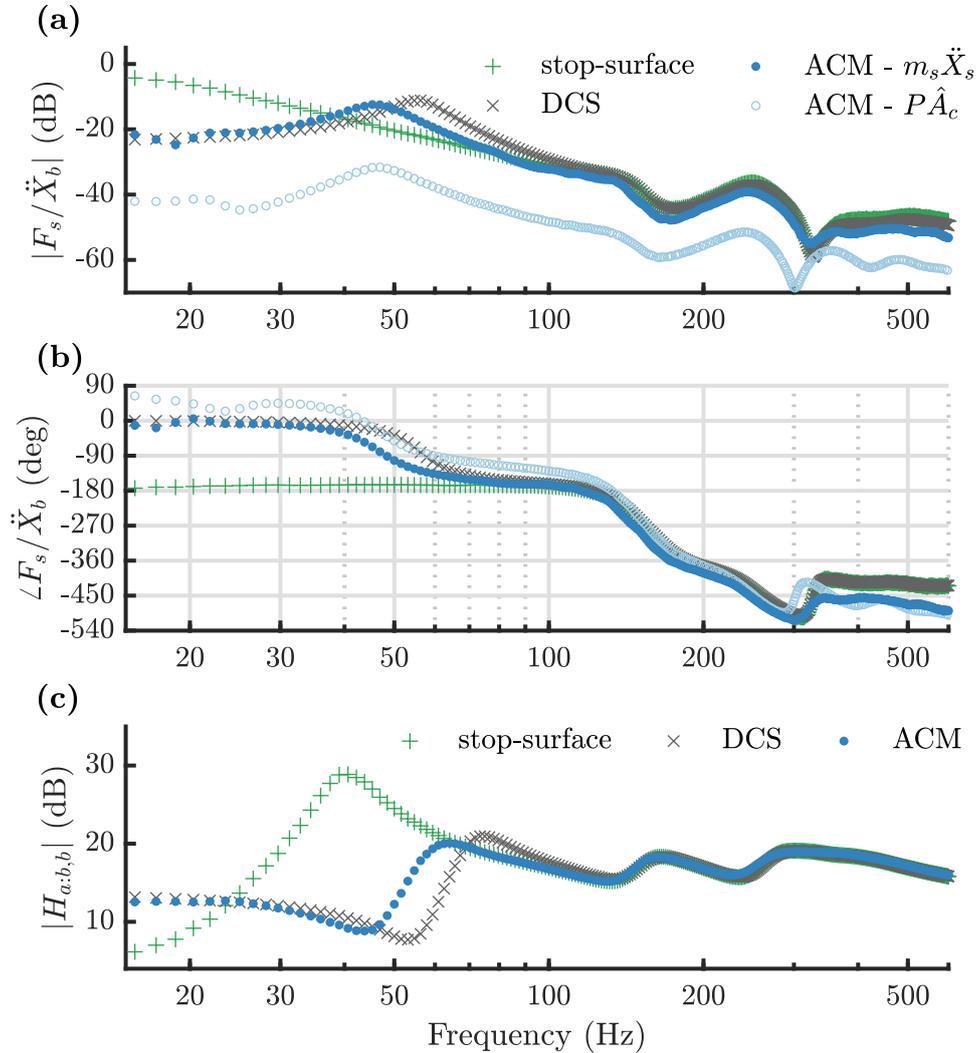


Fig. 9.5 Bode magnitude plots of the phantom-response of various sensors with 23.8 mm outer diameters. **(a)**: Semi-empirical FRF of force exerted on sensor with respect to the acceleration of the base of the phantom (F_s/\ddot{X}_b) (dB relative to 1 kg). **(b)**: Phase of F_s/\ddot{X}_b . **(c)**: Driving-point acceleration FRF (dB relative to 1 kg^{-1}). Green plus signs (+) represent the estimated stop-surface case; grey crosses (\times) represent the simulated response of a 67 g DCS (computed using $\tilde{\mathbf{T}}_{\text{phantom}}$ for the relevant contact area); and blue dots (\bullet) show the measured response of the housing of a 67 g ACM with a cavity contact area of $\sim 200 \text{ mm}^2$ and a cavity volume of $\sim 2700 \text{ mm}^3$. The hollow light blue circles (\circ) in **(a)** represent the force exerted through the air-cavity of the ACM ($F = PA_c$).

in the ACM system. This is consistent with the models presented in chapter 4, in which finite stiffness terms connect the phantom surface and the housing of the ACM, in contrast to a rigid connection in the DCS case. In section 9.4 we explore how these insights can be used to simulate the response of ACMs on the phantom.

9.3.2 Estimation of cavity-surface motion

In chapter 4, we argue that the pressure in the air-cavity (P) is dictated by volume changes due to the relative motion between the chestpiece housing (X_s) and the constrained ‘skin’ surface in contact with the air cavity (X_c). This allows us to write:

$$P = (X_c - X_s) \frac{\rho_0 c_0^2 \hat{A}_c}{V_0}, \quad (9.1)$$

where X_c is defined as in section 4.5.1 such that a change in the cavity volume $\delta V = \hat{A}_c (X_s - X_c)$. \hat{A}_c is the cavity contact area, V_0 is the equilibrium volume, ρ_0 is the density of air and c_0 is the speed of sound in air.

Equation (9.1) can be rearranged to yield a semi-empirical estimate of the surface deflection X_c :

$$X_c = \left(P \frac{V_0}{\rho_0 c_0^2 \hat{A}_c} \right) + X_s. \quad (9.2)$$

To allow comparison with experimental data we divide equation (9.2) through by a reference signal, obtaining an expression in terms of FRFs. Using the acceleration at the base of the phantom as a reference signal yields:

$$\frac{X_c}{\ddot{X}_b} = \left(\frac{P}{\ddot{X}_b} \frac{V_0}{\rho_0 c_0^2 \hat{A}_c} \right) + \frac{1}{-\omega^2} \frac{\ddot{X}_s}{\ddot{X}_b}. \quad (9.3)$$

Figure 9.6 shows a Bode plot of the semi-empirical response of the cavity surface acceleration to the phantom base acceleration (computed using equation (9.3)) superimposed on a plot of the measured response of the housing acceleration to the base-acceleration (\ddot{X}_s/\ddot{X}_b). At low frequencies, the cavity surface (X_c) and the chestpiece-housing (X_s) move in phase with both each other and the phantom-base (X_b), as can be seen by the 0° phase difference. At around 85 Hz (for this configuration) the magnitude of the surface acceleration response

intersects the magnitude of the chestpiece housing acceleration response. At around this frequency the phase difference between X_c and X_s passes through 90° .

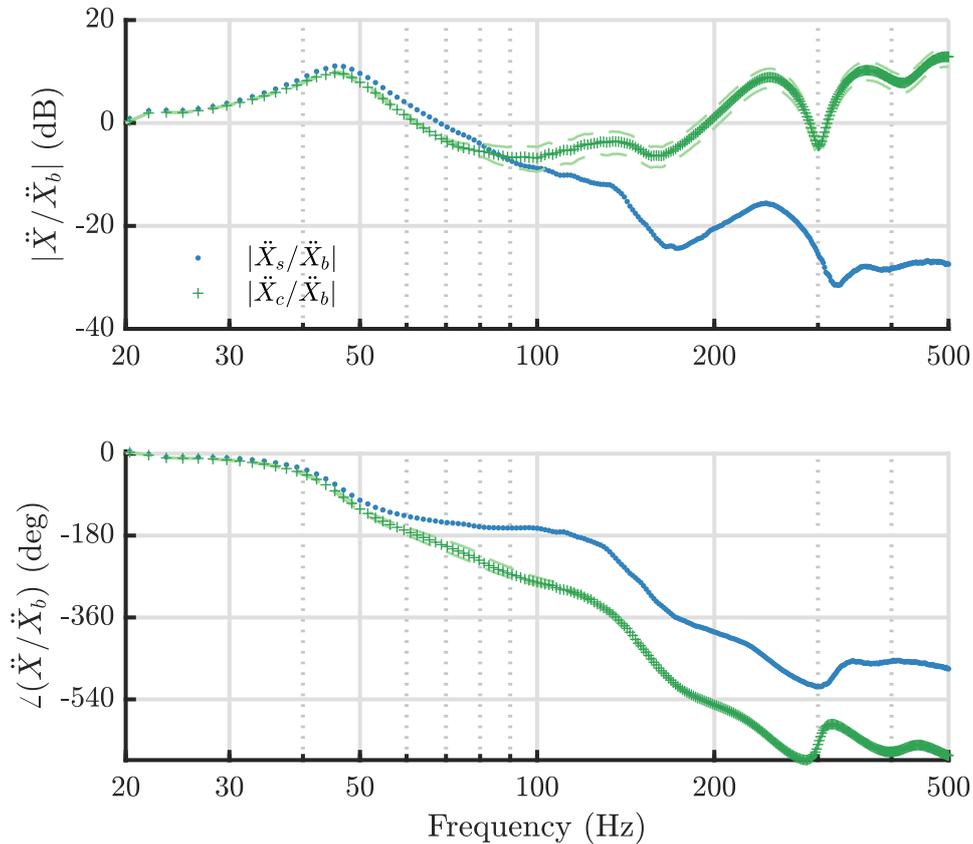


Fig. 9.6 Bode plot of the phantom-acceleration response of the housing of an ACM. Blue dots (\bullet): chestpiece housing acceleration response (H1 estimator). Green plus signs ($+$): cavity-surface acceleration response (semi-empirical: calculated from the measured \ddot{X}_s/\ddot{X}_b and P/\ddot{X}_b FRFs using equation (9.3)). Green dashed lines ($- -$): \ddot{X}_c/\ddot{X}_b calculation with 10% variation in cavity volume. The FRFs in this figure are for an ACM with a 67 g chestpiece, a nominal cavity volume of $\sim 2700 \text{ mm}^3$, an outer diameter of 23.8 mm and an inner diameter of 16 mm.

The cavity-surface motion calculation for figure 9.6 was performed using the nominal cavity-volume and cavity-contact area for the chestpiece. The green dashed lines show that there is little difference in the shape of the response for small (10%) variations in cavity volume. However, if larger static loads are applied, this could lead to a more significant variation in cavity volume, as well as a change in the effective cavity-contact area.

Furthermore, the semi-empirical estimation of \ddot{X}_c is based on the assumption that the impedance of the air cavity is a pure compliance. If a damping term is included in series with this compliance then the semi-empirical estimation of \ddot{X}_c will yield a lower response level at high frequencies.

9.4 Model of the ACM phantom-response

In this section we present a simplified approach to modelling the response of ACMs on the phantom by making use of transmission matrices computed for the rigid (DCS) case. This simplification is motivated by the consideration that the ACM case tends to the DCS case as the cavity stiffness (k_a) tends to infinity. (An infinite cavity stiffness is obtained when the cavity volume reaches zero, at which point an ACM turns into a DCS.)

Adapting the lumped element models derived in chapter 4 for DCSs and ACMs (see figures 4.7 and 4.28) to include a transmission matrix for the phantom we obtain:

$$\underbrace{\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix}}_{\text{DCS}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{ph:DCS}}} \underbrace{\begin{bmatrix} z_s \dot{X}_s \\ \dot{X}_s \end{bmatrix}}_{\text{DCS}}. \quad (9.4)$$

$$\underbrace{\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix}}_{\text{ACM}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{ph:ACM}}} \underbrace{\begin{bmatrix} z_s \dot{X}_s \\ \dot{X}_s \end{bmatrix}}_{\text{ACM}}. \quad (9.5)$$

The underbraces on the two-port vectors indicate the fact that the values of these terms depend on the sensor used (as well as on the excitation level). For the ACM case the transmission matrix between the impedance-head terms at the base of the phantom and the chestpiece housing ($\tilde{\mathbf{T}}_{\text{ph:ACM}}$) can be expressed as:

$$\underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{ph:ACM}}} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{ph:inner}}} \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{contact}}}, \quad (9.6)$$

The division between the ‘inner’ transmission matrix ($\tilde{\mathbf{T}}_{\text{ph:inner}}$) and the ‘contact’ transmission matrix ($\mathbf{T}_{\text{contact}}$) is arbitrary. In keeping with the models proposed

in chapter 4 we define the contact transmission matrix as:

$$\underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{contact}}} \simeq \underbrace{\begin{bmatrix} 1 & i\omega m_{ch:s} \\ 0 & 1 \end{bmatrix}}_{\mathbf{T}_{m_{ch:s}}^{-1}} \underbrace{\begin{bmatrix} 1 & 0 \\ \frac{i\omega}{k_{c,s}+k_a} & 1 \end{bmatrix}}_{\mathbf{T}_{k_{c,s}}} \underbrace{\begin{bmatrix} 1 & i\omega m_{ch:s} \\ 0 & 1 \end{bmatrix}}_{\mathbf{T}_{m_{ch:s}}}. \quad (9.7)$$

This can be interpreted as saying that $\tilde{\mathbf{T}}_{\text{ph:inner}}$ defines the relationship between (the force on and velocity of) the base of the phantom and (the force on and velocity of) a suitable lumped mass at the surface of the phantom. The contact transmission matrix $\mathbf{T}_{\text{contact}}$ consists primarily of a stiffness term ($k_{c,s} + k_a$) between this lumped mass and the sensor housing, but also reallocates some of the lumped mass ($m_{ch:s}$) to the housing-side of the stiffness.

As the cavity stiffness $k_a \rightarrow \infty$, $\mathbf{T}_{k_{c,s}}$ tends to the identity matrix (\mathbf{I}), yielding:

$$\tilde{\mathbf{T}}_{\text{contact}} = \mathbf{T}_{m_{ch:s}}^{-1} \mathbf{I} \mathbf{T}_{m_{ch:s}} = \mathbf{I}. \quad (9.8)$$

Consequently, as $k_a \rightarrow \infty$,

$$\tilde{\mathbf{T}}_{\text{ph:ACM}} \rightarrow \tilde{\mathbf{T}}_{\text{ph:inner}}. \quad (9.9)$$

Furthermore, we argue that as $k_a \rightarrow \infty$ the air-coupled case becomes the rigid (DCS) case, so:

$$\tilde{\mathbf{T}}_{\text{ph:inner}} \simeq \tilde{\mathbf{T}}_{\text{ph:DCS}}. \quad (9.10)$$

This allows us to use the transmission matrix for the phantom derived in chapter 6 for rigid loads to model the response of air-coupled stethoscopes:

$$\underbrace{\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix}}_{\text{ACM}} \simeq \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{\text{ph:DCS}}} \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\mathbf{T}_{\text{contact}}} \underbrace{\begin{bmatrix} z_s \dot{X}_s \\ \dot{X}_s \end{bmatrix}}_{\text{ACM}}. \quad (9.11)$$

It should be noted that this approach is a simplification which does not account for changes in the mechanical properties of the phantom due to the manner in which the phantom-surface is loaded and constrained. It is anticipated that the transmission matrix between the base of the phantom and the surface is a function of the static load applied, the way in which this static load is applied,

and the way in which the surface is constrained. However, the numerical effect of these factors cannot be easily established by experimental means, and the method presented here should serve as a reasonable approximation to illustrate the method and validate the approach.

Equation (9.11) can be used to simulate the response of the coupled ACM-phantom system to excitation at the base of the phantom. The equation is represented as an impedance analogy circuit in figure 9.7. The small mass lumped to the rim ($m_{ch:s}$ in equation (9.7)) is neglected for simplicity (the effect of including $m_{ch:s}$ in the simulation is found to be negligible).

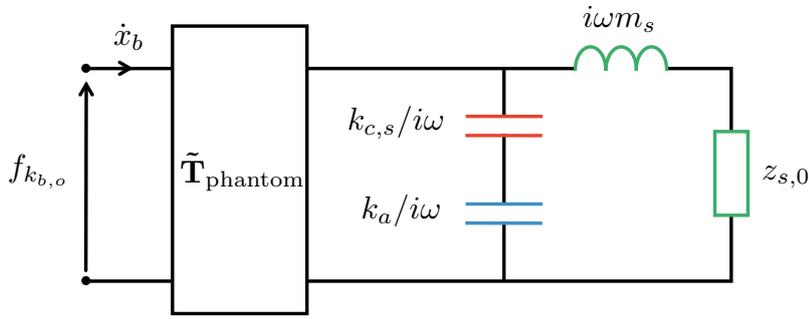


Fig. 9.7 Impedance analogy circuit representing the coupling of an ACM to the phantom. Simplified model using a transmission matrix ($\tilde{\mathbf{T}}_{\text{phantom}}$) for the phantom and neglecting $m_{ch:s}$. All symbols are as defined in section 4.5 and the colours used match those in figure 4.27.

To compute the pressure in the air cavity we first compute the force exerted on the sensor. Using, for example, the acceleration at the base of the phantom as the reference, we obtain:

$$\frac{F_s}{\ddot{X}_b} = \frac{1}{C + D/z_s}, \quad (9.12)$$

where $z_s = (i\omega m_s + z_{s,0}) \parallel (k_{c,s} + k_a)/i\omega$.

The terms k_a and $k_{c,s}$ then form a potential divider, which allows us to compute force through the air cavity as

$$F_{\text{cavity}} = \left(\frac{k_a}{k_a + k_{c,s}} \right) F_s. \quad (9.13)$$

As noted in chapter 4, the stiffness term $k_{c,s}$ represents a constitutive model such as a parallel spring-dashpot combination (giving an impedance of $z_{c,s} = \lambda_{c,s} + k_{c,s}/i\omega$ in place of $k_{c,s}/i\omega$ in the circuit). Allowing also for an arbitrary air

cavity impedance z_a , equation (9.13) then becomes:

$$F_{\text{cavity}} = \left(\frac{z_a}{z_a + z_{c,s}} \right) F_s. \quad (9.14)$$

Considering figure 9.5(a), it is evident that only a small portion of the force exerted on the sensor housing is exerted through the air cavity (compare ● with ○). This suggests that in our model $k_{c,s} > k_a$. The force exerted through the air cavity (F_{cavity}) can be converted into the cavity-pressure using the contact area:

$$P = F_{\text{cavity}} / \hat{A}_c. \quad (9.15)$$

In order to simulate the phantom-response of an ACM using equations (9.12) - (9.13), appropriate values are required for m_s , k_a , $z_{s,0}$, and $z_{c,s}$. The sensor mass m_s is measured directly on scales, while k_a is computed using equation (3.47) with nominal values for the cavity dimensions. Suitable values for $z_{s,0}$, and $z_{c,s}$ are selected by fitting them to the driving point response. For the case shown in figure 9.8 this yields values of $k_{c,s} = 16 \text{ kN/m}$, $\lambda_{c,s} = 18 \text{ Ns/m}$ and $z_{s,0} = \lambda_{s,0} = 0.5 \text{ Ns/m}$. The $\lambda_{s,0}$ term represents a small amount of damping from the transducer wires. By selecting parameter values based on the driving point response only, we are able to test our model without over-fitting to the measured response data.

Picking appropriate (and physically sensible) values for $z_{c,s}$ and $z_{s,0}$, and using these in conjunction with the known sensor mass m_s , the estimated cavity stiffness (k_a) and the transmission matrix for the DCS case ($\tilde{\mathbf{T}}_{\text{ph:DCS}}$), we are able to match the measured driving point response closely. The resulting parameter values also lead to an accurate estimation of the motion of the sensor housing (- in figures 9.8(b,c)). The simulated response of the pressure in the air cavity (- - in figure 9.8(a)) matches the general features of the measured response (in particular the dominant low frequency peak). However, the model over-predicts the cavity pressure at low frequencies and under-predicts the cavity pressure at higher frequencies.

One possible explanation for the low frequency dip in cavity-pressure is the the fact that the cavity is not completely air-tight with small leaks around the microphone allowing low-pressure equalisation to prevent a build up of static pressure within the cavity. This does not explain the recovery of the pressure level

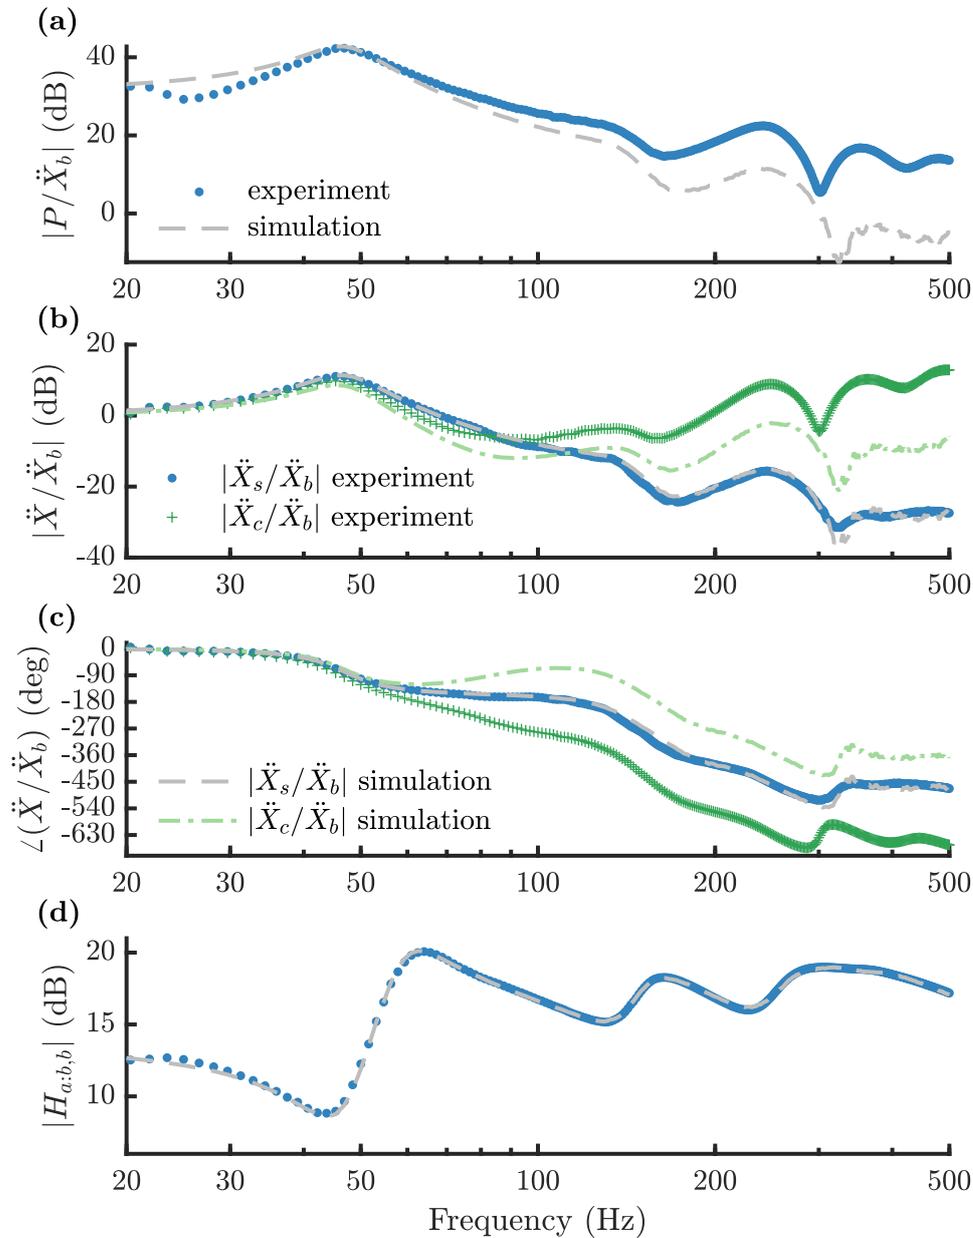


Fig. 9.8 Phantom response of an ACM with a mass of 67 g, an outer diameter of 23.8 mm, an inner diameter of 16 mm and a cavity volume of 2,700 mm³. Blue dots (•) denote experimental FRF estimates, while light grey dashed lines (—) give the simulated response for $k_a = 2.1$ kN/m, $k_{c,s} = 16$ kN/m, $\lambda_{c,s} = 18$ Ns/m and $\lambda_{s,0} = 0.5$ Ns/m, using $\tilde{\mathbf{T}}_{\text{phantom}}$ for a 23.8 mm diameter rigid contact. **(a)**: Magnitude of H1 estimator of P/\ddot{X}_b (decibel reference value 1 kg m⁻²). **(b,c)**: Magnitude and phase of H1 estimator of X_s/\ddot{X}_b (decibel reference value 1 kg m⁻²). Dark green plus signs (+): semi-empirical, light green dot-dashed lines: simulation. **(d)**: Magnitude of H1 estimator of driving point accelerance (decibel reference value 1 kg⁻¹).

at very low frequency, but it should be noted that the coherence drops rapidly below 18 Hz, and at these low frequencies the background noise is substantial.

Comparing the magnitude of the simulated cavity-surface response (\ddot{X}_c/\ddot{X}_b) with the semi-empirical estimate (the light green dot-dashed line versus the dark green plus signs in the second plot of figure 9.8), the simulation seems to underestimate the cavity-surface response, especially at high frequencies. If the semi-empirical estimation according to section 9.3.2 accurately captures the surface motion, then this suggests that the modelled cavity-surface is over-constrained with respect to the sensor housing. Considering the phase of the cavity-surface response, the simulation matches the semi-empirical estimate well at low frequency, when the entire system oscillates in phase. However, at high frequencies the phase of the semi-empirical cavity-surface response lags the housing, while the simulation suggests that the the cavity-surface response leads the housing response. We note that no value of model parameters would lead to X_c lagging X_s . The discrepancy could be due to the fact that the phantom transmission matrix for the DCS case is unable to capture the nuances of vibration-transmission when the cavity surface is able to change shape, or due to the fact that our model includes insufficient degrees of freedom to accurately capture the response at these frequencies. A further cause for discrepancy could be the fact that the semi-empirical estimate does not account for any damping in the air cavity, and may thus overestimate the high-frequency response of X_c .

9.5 Stop-surface response of air-coupled microphones

In chapter 8, we introduce a method for converting the phantom response of a sensor into an equivalent ‘stop-surface response’ by treating the base of the phantom as an acceleration source. Adapting equation (8.11) to give the response of the microphone pressure (P) with respect to the ‘stop-surface’ force (F_{ss}) yields:

$$\underbrace{\frac{P}{F_{ss}}}_{\text{accel source}} = i\omega \underbrace{C}_{\tilde{T}_{\text{phantom}}} \frac{P}{\ddot{X}_b}. \quad (9.16)$$

The ‘stop-surface’ response can be interpreted as simply being the ratio of the pressure measured in the air cavity of the ACM under test to the force exerted on an infinite-impedance load with the same outer contact area as the ACM, for the same acceleration at the base of the phantom. This conversion makes no assumptions about the physics of the ACM’s response. Figure 9.9 shows the same data presented in figure 9.4(b) converted from the ‘phantom-response’ to the ‘stop-surface’ response. The ‘stop-surface’ reference case is derived from experimental data using the method outlined in chapter 8. The response of the stop-surface force to acceleration at the base of the phantom is shown in figure 9.5.

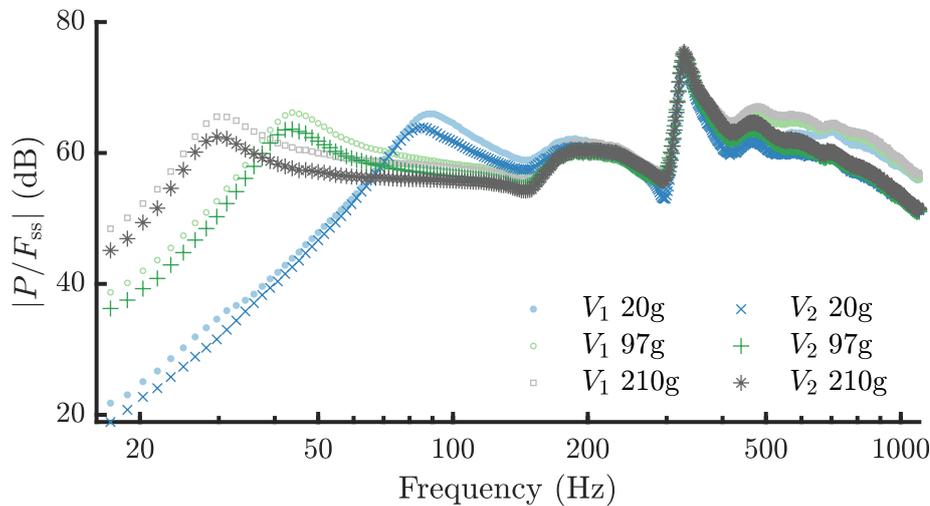


Fig. 9.9 Bode magnitude plot of the stop-surface referenced response of the pressure in the air cavity of an ACM for two cavity volumes ($V_1 = 1240 \text{ mm}^3$, $V_2 = 1.8V_1$) and three housing masses. The ‘stop-surface’ reference case is an infinite impedance rigid load with a contact-diameter equal to the outer diameter of the ACMs, excited with the same phantom-base acceleration as the ACMs (see text for details). The dB reference value is 1 m^{-2} .

For the DCS case (see chapter 8) the ‘stop-surface’ force can be interpreted as the Thévenin equivalent force in an impedance-coupling model for the response. This is because the contact conditions are the same for the DCS and the stop-surface case. As discussed in chapter 4, both the Thévenin impedance and the Thévenin equivalent force depend on the contact conditions. For an ACM the ability of the cavity-surface to change shape and move with respect to the portion of the surface that is in contact with the chestpiece rim means that the contact differs significantly from the stop-surface case. This means that the stop-surface force and the source-impedance seen by a DCS cannot be directly interpreted as the

Thévenin equivalent force and impedance for an ACM of the same outer contact area.

Nevertheless, the use of the DCS Thévenin case in the modelling of an ACM's response allows a reasonable low-order approximation of the response to be obtained. In section 9.5.1 we compare the results of modelling the ACM using this simplification to the measured stop-surface response. In sections 9.6 - 9.7 we use this simplified model to explore trends in ACM performance.

Modelling the ACM response by using the DCS stop-surface force as the Thévenin equivalent force, and using the source-impedance seen by a DCS at the phantom surface as the Thévenin impedance is mathematically equivalent to the assumption that the DCS transmission matrix for the phantom is valid for the ACM case (see section 9.4). In both cases these are simplifications to enable the ACM case to be modelled quantitatively.

9.5.1 Thévenin response model

Figure 9.10 shows a simplified version of the impedance analogy model for an ACM derived in chapter 4 (see figure 4.28). This model allows the response of the pressure in the air cavity and the motion of the housing to be computed in terms of the stop-surface force ($f_{ss} = f_{\text{Thé}}$), noting the caveats discussed in section 9.5.

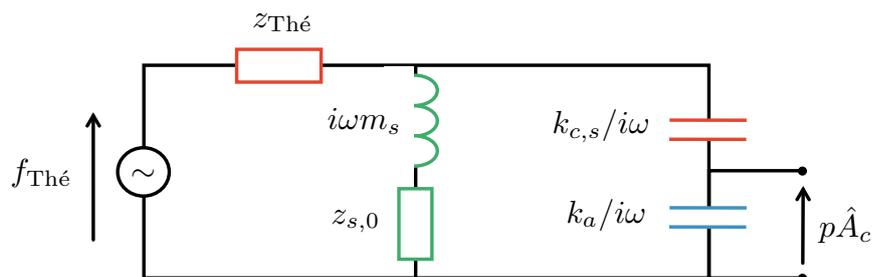


Fig. 9.10 Simplified Thévenin equivalent impedance analogy circuit for an ACM on the phantom.

Figure 9.11 contains the same information as figure 9.8(a-c). The measured and semi-empirical responses are converted to be in terms of the ‘stop-surface force’. The simulations are conducted using the circuit shown in figure 9.10. Figure 9.11(b) shows the phase of the P/F_{ss} FRF. The simulation lags the experimental

response by 90° . This could be partially explained by a significant damping term within the cavity. Small air-leaks between the microphone and the housing would lead to a resistance-dominated impedance term (z_{leak}) in parallel with $k_a/i\omega$ in figure 9.10, as discussed in appendix C.2. This could account for the low frequency discrepancy between experimental and simulated data in both magnitude and phase, but cannot account for the discrepancy at higher frequency.

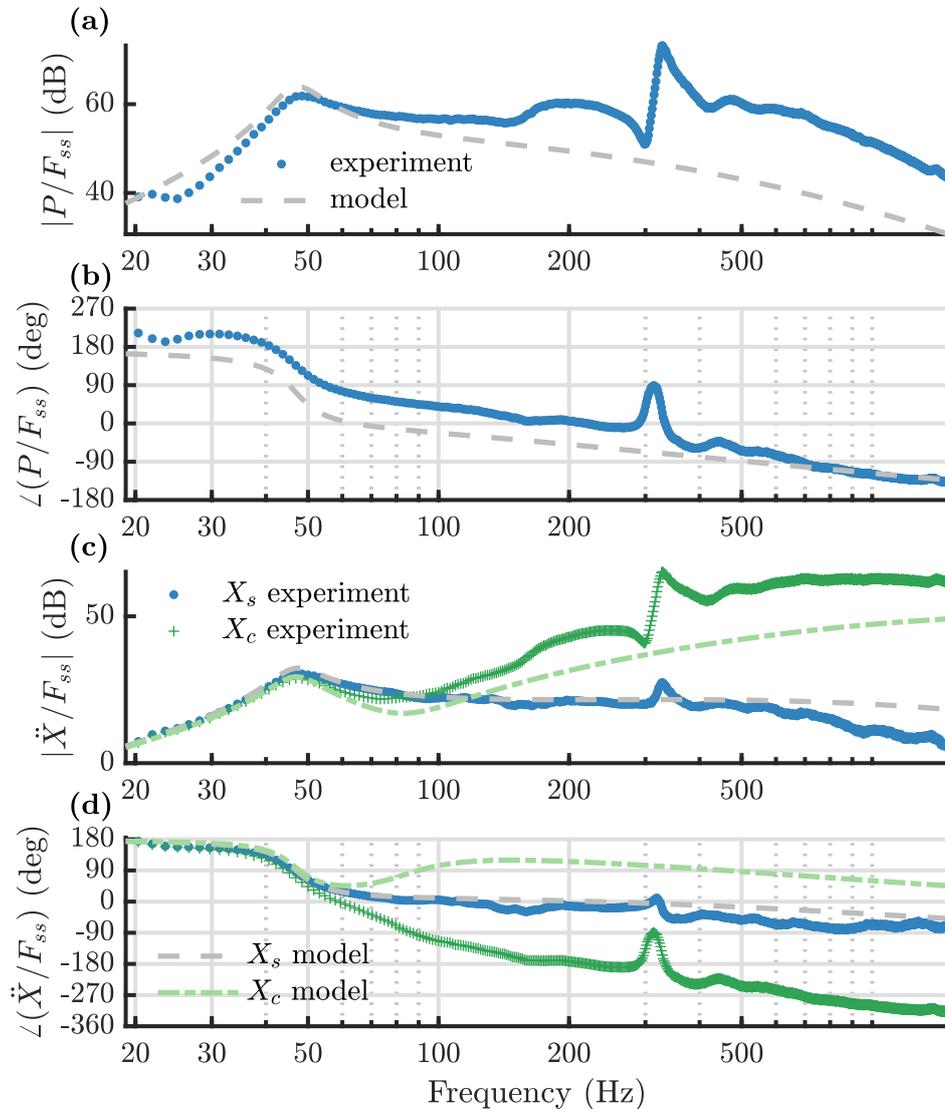


Fig. 9.11 Measured, semi-empirical and modelled stop-surface responses of an ACM and the measured and modelled phantom driving point acceleration. **(a,b)** Bode plot of the cavity-pressure with respect to the stop-surface force. **(c,d)** Bode plot of the housing and cavity-surface accelerations with respect to the stop-surface force. The blue dots denote experiment data, the green plus signs denote the semi-empirical estimation and the grey dashed and green dot-dashed lines denote the model simulations. Parameter values for the model simulations are given in appendix D.4.

The discrepancy between the modelled and semi-empirical X_c/F_{ss} is likely to be primarily due to the inadequacy of using the DCS stop-surface force and source-impedance as the Thévenin force and impedance for the chest in the ACM case. In the model, X_c/F_{ss} is a driving point accelerance, which means that the phase is limited to be between 0 and 180°. This means that, regardless of the parameter values used to model the chest-ACM contact, the model cannot match the semi-empirical phase prediction. This suggests that modelling the cavity surface as being approximately equivalent to the chest-surface in contact with the DCS is not appropriate. Nevertheless, this modelling simplification is able to capture many of the key trends in the response and, importantly, is able to accurately describe the effect of design parameters such as the cavity-volume.

9.6 The effect of cavity volume

Figure 9.12 shows the simulated effect of changing the air cavity volume of an ACM. The simulation is based on the case shown in figure 9.11, with the cavity volume varied while all other model parameters are kept constant. In this simulation we treat the cavity as a pure compliance, neglecting any damping terms.

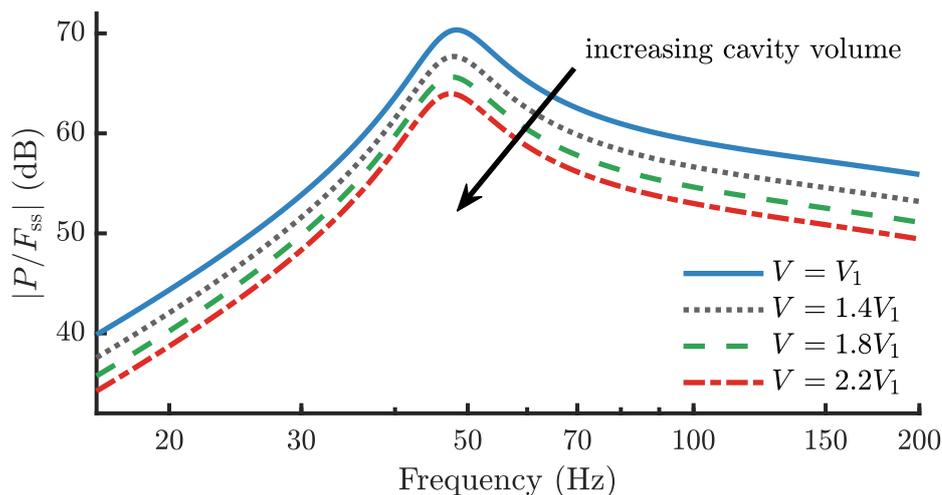


Fig. 9.12 Simulation of the effect of air cavity volume on the response of an ACM. The cavity is treated as a pure compliance. All model parameter values are as in figure 9.11, with the exception of cavity volume V . The reference volume is $V_1 = 1240 \text{ mm}^3$ and the dB reference for the response is 1 m^{-2} .

The modelled trend matches experimental observations (see section 9.2.2) which show that increasing the cavity volume reduces the level of the response and shifts the dominant low frequency resonance to lower frequency. This supports the well-established claim (see e.g. [163]) that minimising the cavity volume of an air-coupled microphone maximises its response. However, two important factors need to be considered: the effect of damping and the effect of occlusion (see section 2.1.3).

As the cavities used for the experiments presented in this thesis are designed with an inverted conical-horn profile, pressing ACMs with different cavity volumes (and thus shapes) into the chest/phantom by the same amount will lead to different changes in contact area and cavity volume. Further to changing the properties of the cavity, occlusion will also affect the stiffness between the cavity-surface and the rim that arises due to the chest/phantom tissue ($k_{c,s}$). If a greater proportion of the ‘skin’ surface is directly in contact with the chestpiece housing then we expect this stiffness to increase.

9.7 The role of stethoscope diaphragms

Figure 9.13 shows the phantom-response of an ACM with and without a diaphragm. The blue crosses and green plus signs show FRF estimates from experimental data, while the red dot-dashed line and the grey dashed line show semi-empirical estimates of the cavity-surface motion computed from the measured cavity-pressure and chestpiece-acceleration using the method described in section 9.3.2.

The addition of a diaphragm has two significant effects on the response of the pressure in the air cavity (9.13(a)). Firstly, the diaphragm leads to a significant attenuation of the pressure level and secondly, it leads to an upwards-shift in the resonance frequencies. The driving point accelerance (9.13(d)) also reveals an upwards shift in the ω_{-1} and ω_2 frequencies when the diaphragm is added, indicating increased stiffness in the system.

The observation that the inclusion of a diaphragm leads to signal attenuation at all frequencies is consistent with the claims of e.g. Ertel (1966) [29] and Kindig (1982) [72], although it should be noted that the level of attenuation is not independent of frequency as suggested by McGee (2017) [98].

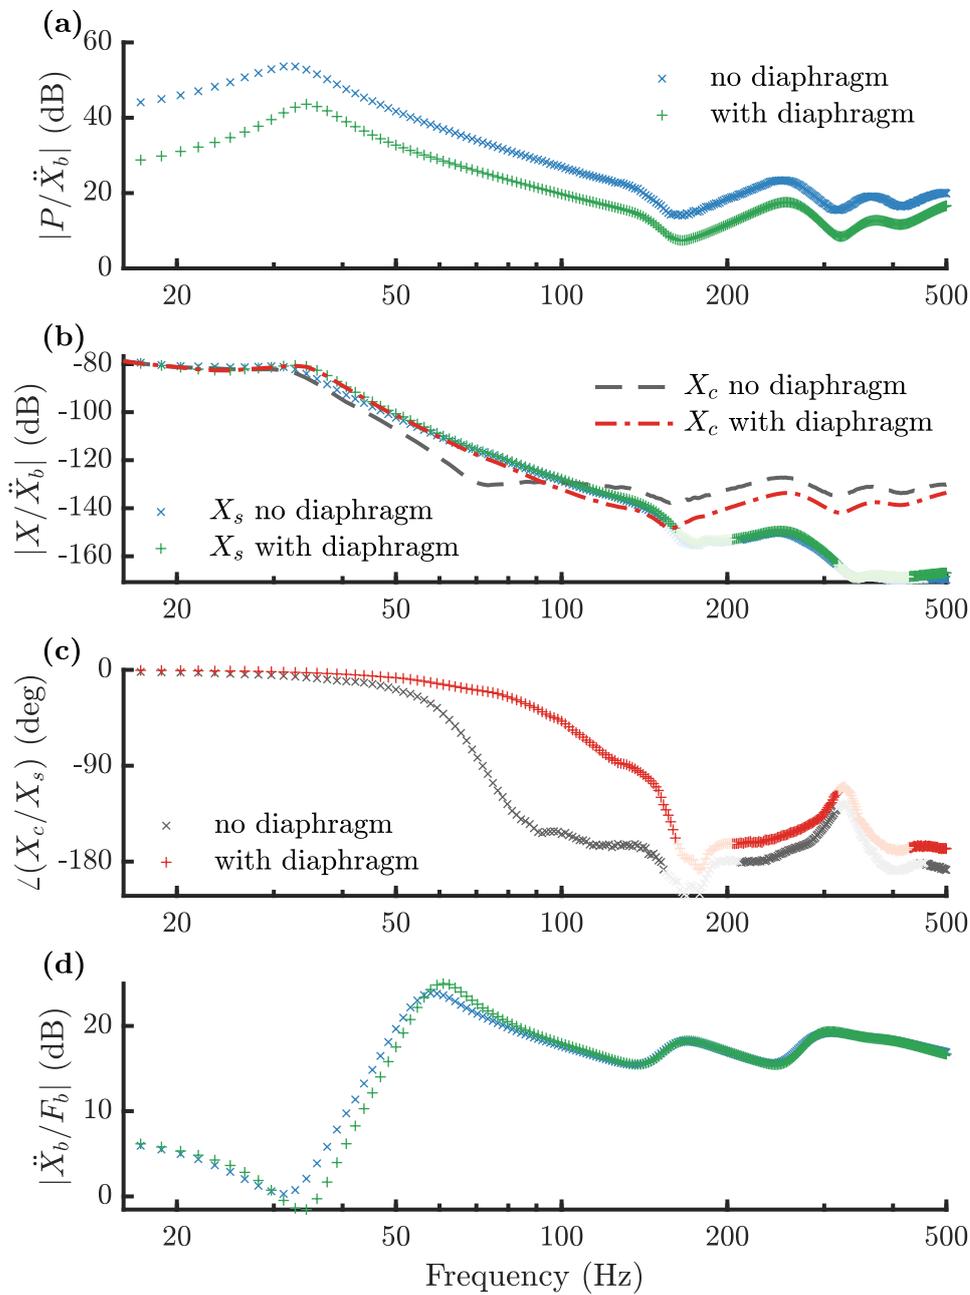


Fig. 9.13 Experimental phantom response of an ACM with and without a diaphragm. **(a)** Response of the cavity pressure with respect to the phantom-base acceleration (decibel reference value 1 kg m^{-2}). **(b)** Response of the chestpiece housing displacement (X_s) and the (semi-empirical) cavity-surface displacement (X_c) with respect to phantom-base acceleration (\ddot{X}_b) (decibel reference value 1 s^2). **(c)** Phase of the (semi-empirical) cavity-surface displacement with respect to chestpiece housing displacement. **(d)** Driving point acceleration FRF at the base of the phantom (decibel reference value 1 kg^{-1}).

Figure 9.13(b) shows the response of the housing (measured) and cavity-surface (semi-empirical) to excitation at the base of the phantom. It can be seen that the housing motion is only minimally affected by the presence of the diaphragm, with a shift in the low frequency peak to a slightly high frequency (matching observations of the cavity-pressure and driving point responses). The response of the cavity-surface, however, varies significantly with the addition of the diaphragm. In the no-diaphragm case the cavity-surface response (grey dashed line) deviates significantly from the housing response beyond the dominant low-frequency resonance, while in the diaphragm case the cavity surface response (red dot-dashed line) conforms closely to the housing response up until significantly higher frequencies (~ 150 Hz). This observation is supported by 9.13(c), which shows the phase difference between the motion of the cavity-surface and the chestpiece housing. In this figure it can be seen that the cavity-surface motion is in phase with the sensor housing motion for a considerably greater proportion of the frequency range for the diaphragm case compared to the no-diaphragm case. These observations suggest that the addition of a diaphragm significantly increases both the stiffness and the damping between the cavity-surface and the chestpiece housing. This is in agreement with the model proposed for stethoscope diaphragms in section 4.5.5.

Careful measurements of mechanical properties of different diaphragms coupled to the chest are outside the scope of this thesis. Instead, figure 9.14 demonstrates the ability of the model presented in chapter 4 to capture the key differences between the response of ACMs with and without diaphragms.

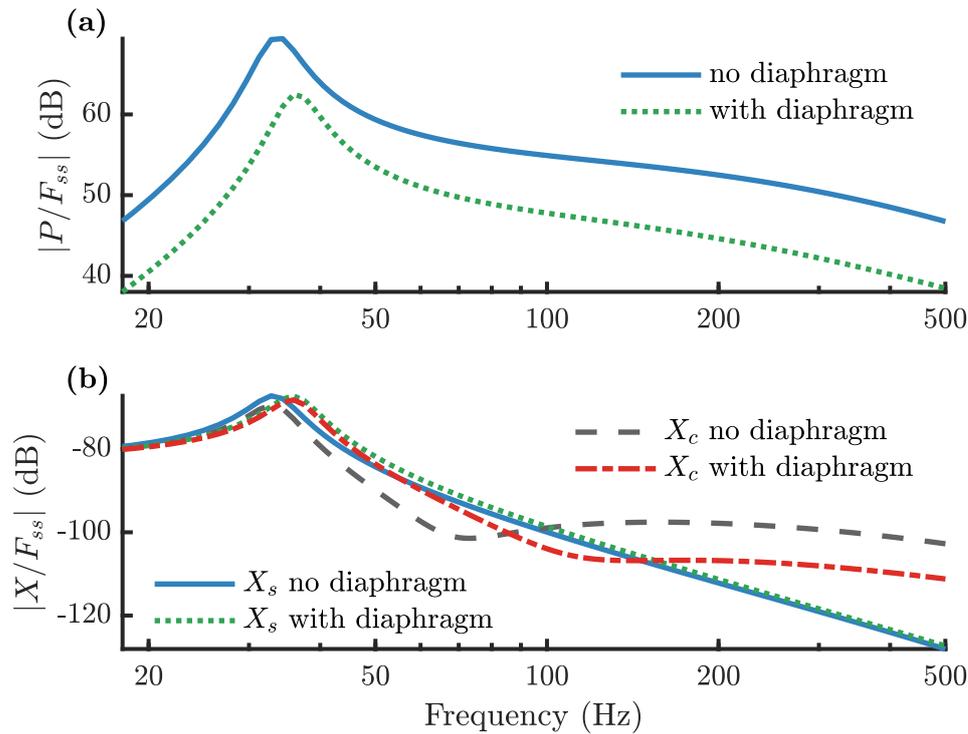


Fig. 9.14 Thévenin equivalent model simulation of the effect of adding a diaphragm to an ACM. **(a)** Response of the cavity pressure with respect to the stop-surface force (dB relative to 1 m^{-2}). **(b)** Response of the chestpiece housing and the cavity surface displacement with respect to the stop-surface force (dB relative to $1 \text{ s}^2 \text{ kg}^{-1}$). Parameter values for the model are given in appendix D.5.

The data presented in this chapter suggests that a traditional stiff diaphragm has an entirely negative effect on the performance of air-coupled sensors. However, there are several reasons why using a diaphragm may be advantageous. A stethoscope with a diaphragm can be easier to keep clean, avoiding contamination. A diaphragm can also protect the microphone in an ACM sensor. Perhaps most importantly from an acoustic perspective, the presence of a diaphragm also limits the occlusion of the air cavity by chest tissue, thus allowing air-cavities with significantly smaller volumes to be used. Decreasing the cavity volume increases the response level, and can thus make up for the attenuation caused by the diaphragm itself. Reducing the cavity volume also further increases the stiffness between the cavity-surface and the chestpiece housing (k_a) leading the resonance(s) to shift to even higher frequencies.

Figure 9.15 shows the results of a simulation comparing an ACM with a large cavity-volume and no diaphragm (commonly referred to as the ‘bell mode’)

against an ACM with the same mass and dimensions but a lower cavity volume and the inclusion of a diaphragm. While the diaphragm attenuates the response, the lower cavity volume enabled by the diaphragm amplifies the response. The net effect for this particular comparison is that the ‘bell-mode’ yields a higher response amplitude at lower frequencies, while the ‘diaphragm-mode’ yields a higher response amplitude at higher frequencies. The quantitative responses depend on the volumes used for the two ‘modes’ and on the mechanical properties of the diaphragm.

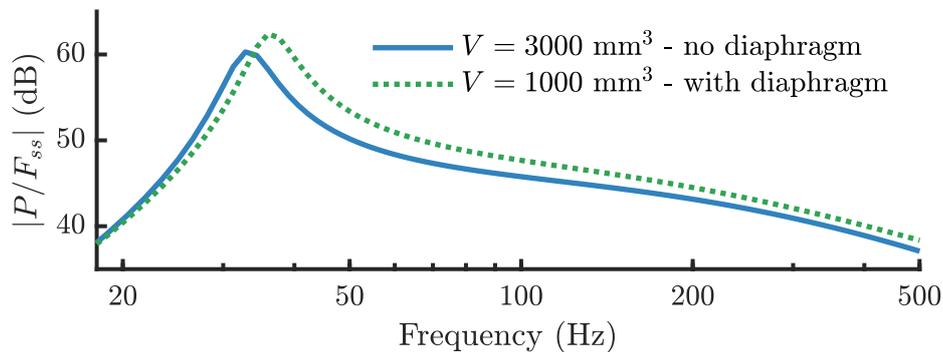


Fig. 9.15 Thévenin equivalent model simulation comparing an ACM with a large cavity volume and no diaphragm (‘bell mode’) to an ACM with a low volume with a diaphragm (‘diaphragm mode’). The decibel reference value is 1 m^{-2} . All model parameter values are as in figure 9.14, with the exception of cavity volume V .

The fact that a diaphragm allows the cavity volume to be reduced has been ignored by many studies comparing the properties of ACMs with and without diaphragms (see e.g. [67]). This can lead to the erroneous conclusion that diaphragms have an entirely negative effect on stethoscope performance.

9.8 The effect of holding an ACM

The effect of holding an ACM consists of two factors: the static load applied and the attachment impedance due to the hand-arm system. The static load applied by the holder sets the equilibrium position of the coupled system, and thus the impedance of the chest. The attachment impedance of the human hand-arm system becomes part of the coupled chest-stethoscope-clinician system. For experiments on the phantom small changes in the application pressure exerted

by the hand do not significantly effect the impedance of the phantom seen by the sensor, thus allowing us to identify the effect of the attachment impedance.

Figure 9.16 shows the H1 estimator for P/F_b for four experiments conducted on the same ACM. The blue dots show the FRF for the case in which the ACM stands freely on the top surface of the phantom. The three green markers show the FRFs estimated from three separate experiments in which the ACM was held in place by hand.

In figure 9.16 it can be seen that, in all three cases, holding the ACM decreases the amplitude of the low frequency peak without significantly changing its frequency. This suggests that the effect of holding the ACM is dominated by damping. There is a small upwards shift in the frequency of the dominant low frequency peak, which suggests that the attachment-impedance ($z_{s,0}$) also has a small stiffness component. In contrast, figure 9.4 shows that the dominant effect of adding mass to an ACM is to increase the amplitude of the low frequency peak and to shift the peak-frequency downwards significantly. This suggests that modelling the effect of application pressure on ACM performance by added load masses to the ACM chestpiece (as done in experiments by e.g. [71, 106, 73]) does not accurately capture the effect of pressing the sensor against the chest by hand.

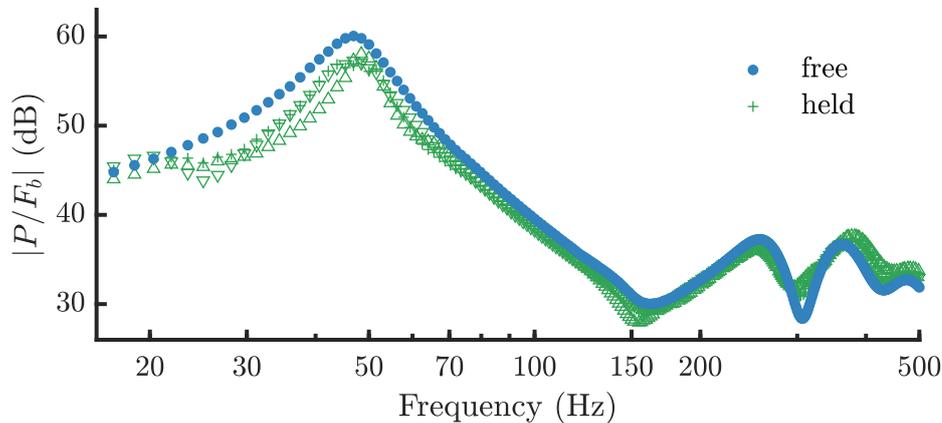


Fig. 9.16 Effect of holding the chestpiece of an ACM on the phantom response. H1 FRF estimator of the cavity pressure with respect to the input force at the base of the phantom (dB relative to 1 m^{-2}). Blue dots show the free case; green markers show three separate held cases.

Figure 9.17 shows a simulation of the phantom response of an ACM, illustrating the effect of using two different models for the holding of the ACM. The solid blue

line shows the ‘free’ case (with a chestpiece mass of ~ 250 and an attachment impedance of $z_{s,0} = \lambda_{s,0} = 0.5$ Ns/m to account for the transducer wires). The green dashed line shows the effect of modelling the hand with a spring and a dashpot in parallel ($z_{s,0} = \lambda_{s,0} + k_{s,0}/i\omega$). The dashpot is given a value of 18 Ns/m in accordance with the measurements presented in chapter 7 (section 7.5). A small stiffness of $k_{s,0} = 2$ kN/m is also included in the attachment impedance to account for the observed shift in the resonance frequency. The spring-dashpot model for the holding impedance is able to closely match the effect of holding the ACM observed in the experimental data (see figure 9.16). The grey dot-dashed line shows the effect of simulating the role of the hand using a 500 g mass added to the ACM chestpiece. This simulation does not account for the variation in the mechanical properties of the phantom due to the increased static load. However, this variation is expected to be small for the phantom, and the simulation is consistent with the experimentally observed effect of adding load mass (see figure 9.4). This simulation confirms that adding load masses to the ACM chestpiece is unable to accurately capture the effect of pressing the sensor against the chest by hand.

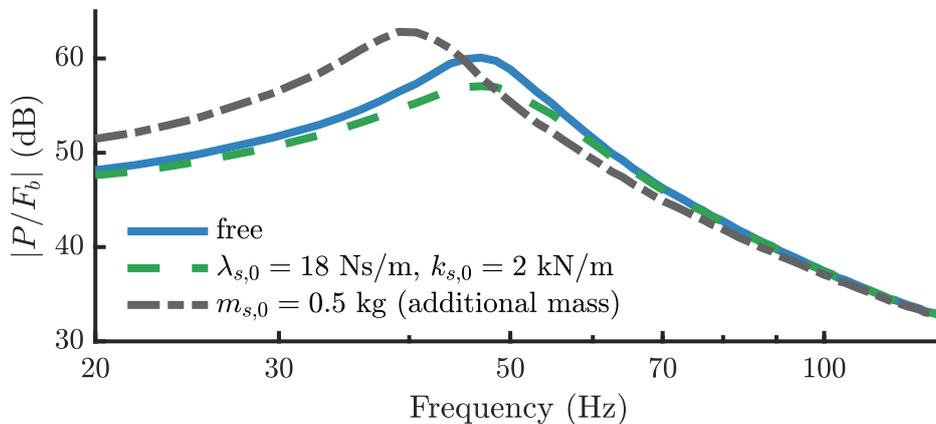


Fig. 9.17 Simulated effect of holding the chestpiece of an ACM on the response of the air cavity pressure to the excitation force at the base of the phantom. Blue solid line: free chestpiece, green dashed line: held chestpiece, with holding represented by a parallel dashpot and spring to ground ($\lambda_{s,0} = 18$ Ns/m, $k_{s,0} = 2$ kN/m), grey dot-dashed line: holding represented by a 500 g load mass.

9.9 Summary and Conclusions

In this chapter, we present results from phantom experiments on air-coupled microphones which show the effect of air cavity volume, chestpiece mass and diaphragm-use. Trends observed in the experimental data are qualitatively consistent with the ACM model derived in chapter 4. Furthermore, by making the crude simplification that the transmission matrix for the phantom obtained from the DCS case can be used to model the ACM case we are able to simulate the expected response of air-coupled sensors numerically. The simulations match the overall response reasonably well, but consistently underestimate the cavity-pressure at high frequencies. This can be largely explained by the acknowledged inadequacy of using the DCS transmission matrix for the phantom despite the significantly different contact characteristics in the ACM case. Using the DCS transmission matrix to model the ACM response has the effect of over-constraining the motion of the air cavity surface and consequently leads to an under-prediction of the volume (and thus pressure) fluctuations within the air cavity.

Importantly, our model is able to accurately capture the effects of varying cavity volume, using an artificial diaphragm between the phantom and the chest and varying the manner in which the ACM is applied to the chest. We demonstrate that the common assumption that the effect of pressing the sensor against the chest by hand can be represented in phantom-experiments by applying load masses to the chestpiece is flawed. Instead, the effect of holding the ACM is dominated by a damping term. This is consistent with the observations on the impedance of the human hand-arm system presented in chapter 7. A wider investigation of the effect of holding using a range of human subjects would strengthen these findings.

Further work is needed to characterise the transmission of vibration from the base of the phantom to the air cavity surface, as well as to determine the effect of application pressure and of the inner and outer diameters of the ACM on $k_{c,s}$. Furthermore, the additional stiffness associated with the use of an artificial diaphragm should be quantitatively characterised, and the effect of occlusion should be explored in depth. Measurements on the human chest are needed to confirm that the observed trends hold true when the same sensors are coupled to the human chest in a clinical setting.

Chapter 10

Stethoscopes with tubing

In contrast to early stethoscopes, which were constructed from rigid cylinders of wood, the flexible tubing of the modern binaural analogue stethoscope offers significant ergonomic benefits. The principle role of the tubing is to allow clinicians to auscultate patients from a distance. However, the advantages in terms of hygiene and ease-of-use are accompanied by predominantly negative acoustic effects. The use of tubing has been found to lead to a general attenuation of the signal [129], in addition to the standing wave resonances which distort the response [29, 67]. Ertel et al. [29] found that standing wave patterns differed significantly among models of analogue stethoscopes, but that in all cases these resonances dominated the response above around 250 Hz. Existing work on modelling air-coupled sensor acoustics has focussed on air-coupled microphones without tubing [163, 142, 143]. To our knowledge no complete model for the coupled chest-chestpiece-tubing system of an analogue stethoscope has been proposed and validated. In this section we address this shortcoming by proposing a model for stethoscopes with tubing. This model is based on the use of electrical analogy theory to couple a mechanical-domain model for the interaction between the chest and the chestpiece to an acoustical-domain model for the tubing.

In chapter 9, we show that the performance of an ACM depends on the local properties of the chest and the manner in which the sensor is applied. These factors differ across auscultation situations, which means that there is no single ‘stethoscope response function’. For air-coupled sensors with tubing the response also depends on the presence of leaks and the impedance of the user’s ears. Instead of exploring the acoustical properties of a single design, we introduce

general methods for the modelling and investigation of this class of sensors and discuss the salient effects associated with adding tubing to an air-coupled sensor.

In this chapter we purposefully simplify the acoustics of the tubing, neglecting e.g., the details of the damping mechanisms, the bifurcation of the tubing, and a detailed treatment the ears of the clinician. The reason for these simplifications is to ensure that the emphasis remains firmly on the core contribution of this chapter: a method for coupling our mechanical model for the interaction between the chestpiece, the chest and the hand, to an acoustic model for the propagation of sound through tubing to the ears of a clinician.

10.1 Coupling the mechanical and acoustic domains of an ACS

The air-coupled stethoscope (henceforth ‘ACS’) can be broken down into a mechanical-domain model and an acoustical-domain model. The mechanical model describes the interaction between the chest, the stethoscope rim and the air cavity surface, as discussed in section 9. The acoustical model describes the air cavity, the tubing and, in the case of an analogue stethoscope, the coupling to the human ears.

Section 3.5.6 describes a method for coupling the mechanical domain to the acoustical domain by using a transformer to connect the respective impedance-analogy circuits. Figure 10.1 shows an implementation of this coupling method to yield an impedance analogy circuit for an air-coupled sensor on the phantom. If no further acoustical terms are added, then figure 10.1(a) reduces to figure 9.7.

Terminating the circuit in figure 10.1(a) with the circuit in figure 10.1(b) allows us to model an air-coupled sensor with a segment of tubing attached to the air cavity. The narrow cross-sectional area of typical stethoscope tubes, and the low frequency range associated with auscultation, justifies the assumption of 1D plane-wave propagation along the tubing. The length of the tubing (typically around 500 - 700 mm) is not small compared to the wavelengths associated with auscultation, and thus a lumped-element model is inappropriate unless the tubing is discretised into a number of short segments. A more convenient approach is to model the tubing as a transmission line (\mathbf{T}_{tube}) using the solution to the 1D

wave equation outlined in section 3.1.2. In figure 10.1(b), the far end of the tube is terminated by an impedance z_e , which could, for instance, represent the ear of a clinician. For a microphone at the distal end of the tube the load impedance z_e can be approximated as infinite, yielding an open circuit.

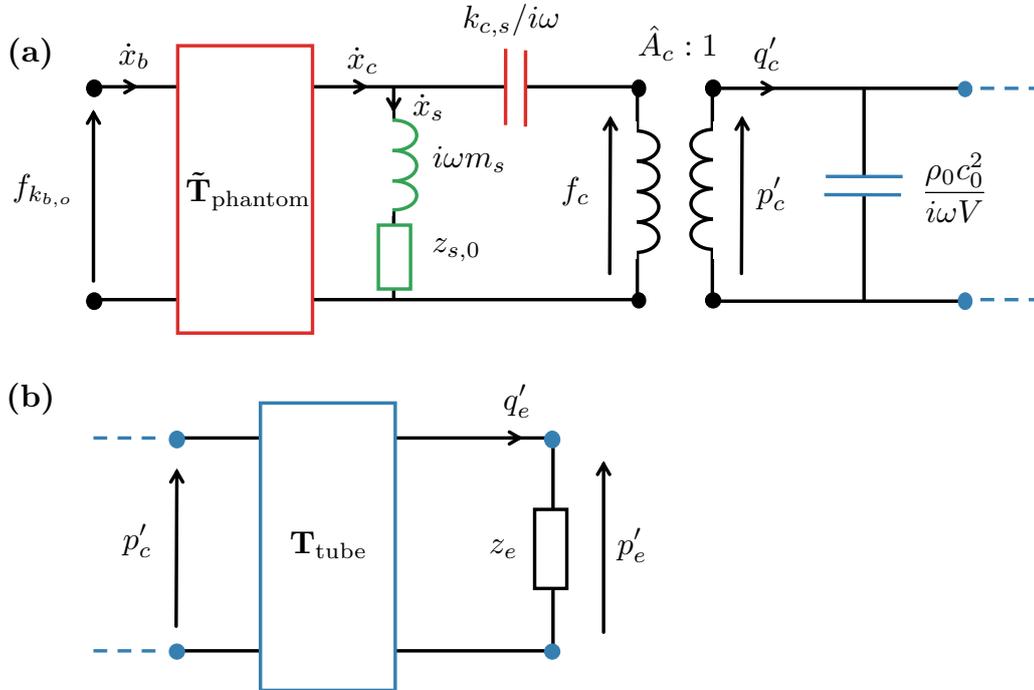


Fig. 10.1 Impedance analogy circuit for an air-coupled sensor with the mechanical and acoustic domains coupled using a transformer of ratio $\hat{A}_c : 1$. **(a)** Mechanical domain starting at the signals measured by the impedance head at the base of the phantom ($F_{k_{b,o}}$ and \dot{X}_b), coupled to an acoustic domain at the interface between the chest surface and the air cavity. The volume velocity perturbations in the air cavity are given by $q'_c = \hat{A}_c(\dot{x}_c - \dot{x}_s)$. The acoustic domain consists of a compliance term representing the air cavity to which further acoustic terms can be added. **(b)** Continuation of the acoustical domain with a transmission line (\mathbf{T}_{tube}) to represent stethoscope tubing and a load impedance z_e at the end of the tube. For a rigid termination z_e is infinite and $q'_e = 0$.

The relationship between the force and velocity measure by the impedance head ($F_{k_{b,o}}$ and \dot{X}_b) and the pressure and volume velocity at a distance x along the tube (P_x and Q_x) can be expressed using two-port notation as:

$$\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix} = \underbrace{\tilde{\mathbf{T}}_{\text{phantom}} \mathbf{T}_{z_{s,0}} \mathbf{T}_{k_{c,s}}}_{\text{mechanical}} \mathbf{T}_{\text{trans}:\hat{A}_c} \underbrace{\mathbf{T}_{k_a^a} \mathbf{T}_{\text{tube}}(x)}_{\text{acoustic}} \begin{bmatrix} P_x \\ Q_x \end{bmatrix}. \quad (10.1)$$

The transmission matrix for a transformer between a force-velocity domain and a pressure-volume velocity domain at an interface of area \hat{A}_c is given by:

$$\begin{bmatrix} F \\ \dot{X} \end{bmatrix} = \underbrace{\begin{bmatrix} \hat{A}_c & 0 \\ 0 & 1/\hat{A}_c \end{bmatrix}}_{\mathbf{T}_{\text{trans}}(\hat{A}_c)} \begin{bmatrix} P \\ Q \end{bmatrix}. \quad (10.2)$$

The transmission matrix for a tube segment of length δl is given by:

$$\mathbf{T}_{\text{tube}}(\delta l) = \begin{bmatrix} \cosh(i k \delta l) & Z_{0:\hat{A}_{\text{tube}}}^a \sinh(i k \delta l) \\ Y_{0:\hat{A}_{\text{tube}}}^a \sinh(i k \delta l) & \cosh(i k \delta l) \end{bmatrix}, \quad (10.3)$$

where $Z_{0:\hat{A}_{\text{tube}}}^a$ is the characteristic acoustic impedance of a tube of area \hat{A}_{tube} as defined in equation (3.17), and $Y_{0:\hat{A}_{\text{tube}}}^a = 1/Z_{0:\hat{A}_{\text{tube}}}^a$. The wavenumber k is permitted to be complex to allow for damping in the tubing (i.e., $ik \rightarrow ik + \hat{\alpha}$ where $\hat{\alpha}$ is the attenuation factor). The remaining transmission matrices in equation (10.1) can be inferred directly from the impedance analogy circuit in figure 10.1, using the methods outlined in chapter 3. For completeness these are given in appendix D.6.

More generally, we can define a transmission matrix $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ relating the force and velocity at the base of the phantom to the pressure and volume velocity at the cavity entrance. This yields

$$\begin{bmatrix} F_{k_{b,o}} \\ \dot{X}_b \end{bmatrix} = \tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}} \mathbf{T}_{k_a^a} \mathbf{T}_{\text{tube}}(x) \begin{bmatrix} P_x \\ Q_x \end{bmatrix}, \quad (10.4)$$

where $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}} = \tilde{\mathbf{T}}_{\text{phantom}} \mathbf{T}_{z_{s,0}} \mathbf{T}_{k_{c,s}} \mathbf{T}_{\text{trans}:\hat{A}_c}$.

Equation (10.4) can then be solved for the response of the pressure (P_x) at a distance x along the tubing as a function of the phantom input signals:

$$\frac{P_x}{F_{k_{b,o}}} = \frac{1}{\underbrace{A}_{\mathbf{T}_{\text{net}}} + \underbrace{B}_{\mathbf{T}_{\text{net}}}/Z_{x,x}^a}, \quad (10.5)$$

$$\frac{P_x}{\dot{X}_b} = \frac{1}{\underbrace{C}_{\mathbf{T}_{\text{net}}} + \underbrace{D}_{\mathbf{T}_{\text{net}}}/Z_{x,x}^a}, \quad (10.6)$$

where $\mathbf{T}_{\text{net}} = \tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}} \mathbf{T}_{k_a^a} \mathbf{T}_{\text{tube}}(x)$. $Z_{x,x}^a$ is the acoustic impedance seen at x . To compute this we first relate the volume velocity and pressure at x (Q_x and P_x) to the volume velocity and pressure at the end of the tube (Q_e and P_e) using:

$$\begin{bmatrix} P_x \\ Q_x \end{bmatrix} = \mathbf{T}_{\text{tube}}(l-x) \begin{bmatrix} P_e \\ Q_e \end{bmatrix}. \quad (10.7)$$

Given that $Q_e = P_e/z_e$, this yields:

$$Z_{x,x}^a = \frac{P_x}{Q_x} = \frac{A + B/z_e}{C + D/z_e}, \quad (10.8)$$

in which A , B , C and D all refer to the terms in $\mathbf{T}_{\text{tube}}(l-x)$.

For the rigid termination case ($z_e \rightarrow \infty$) this leads to:

$$Z_{x,x}^a = \underbrace{\frac{A}{C}}_{\mathbf{T}_{\text{tube}}(l-x)} = Z_{0:\hat{A}_{\text{tube}}}^a \coth(ik(l-x)). \quad (10.9)$$

10.2 Attenuation in the tubing

In practice, the damping of vibrations as they propagate through a stethoscope's tubing plays a significant role in determining the acoustic response of the stethoscope. In equation (3.16) we introduce transmission matrices for lossless ducts. To take into account attenuation, the wavenumber (k) and the characteristic acoustic impedance ($Z_{0:\hat{A}}$) in these matrices must be replaced by complex values [6].

For rigid-walled tubes, the attenuation is dominated by viscous and thermal losses and can be analytically modelled [6]. Approximations for these losses for a range of different tube dimensions have been proposed (see e.g., [6, 149, 26]), however, for flexible tubes the damping mechanisms are considerably more complex. Huang et al. [55] highlight three important aspects to consider for a duct system which includes flexible walls: the dissipation of energy due to damping within the flexible tubing, the break-out and break-in noise, and the sound reflection at the interfaces between rigid and flexible segments due to impedance mismatch.

A detailed consideration of the damping behaviour in a network of flexible tubes is outside the scope of this thesis. For the simulations in this chapter we employ a simple attenuation model for viscous and thermal losses in rigid ducts (see e.g.,

[149, 162]), acknowledging that this does not accurately represent the physics of our experimental setup, or of practical analogue stethoscopes. Details of the attenuation model used can be found in appendix D.7.

10.3 Experimental setup

Figure 10.2 shows the experimental setup used to determine the phantom-response of an air-coupled sensor with tubing. The setup is as described in section 9.1, except that the microphone in the ACS shaft is replaced by a segment of tubing which is terminated by a microphone. The inner and outer diameters of the tubing are 4.3 mm and 6.3 mm respectively. A 1/8 inch microphone (a G.R.A.S. 40DD prepolarised pressure microphone with a RA0082 1/8 inch - 1/4 inch adapter and a 26CB 1/4 inch preamplifier) is used to terminate the duct, with a seal around the microphone shaft.

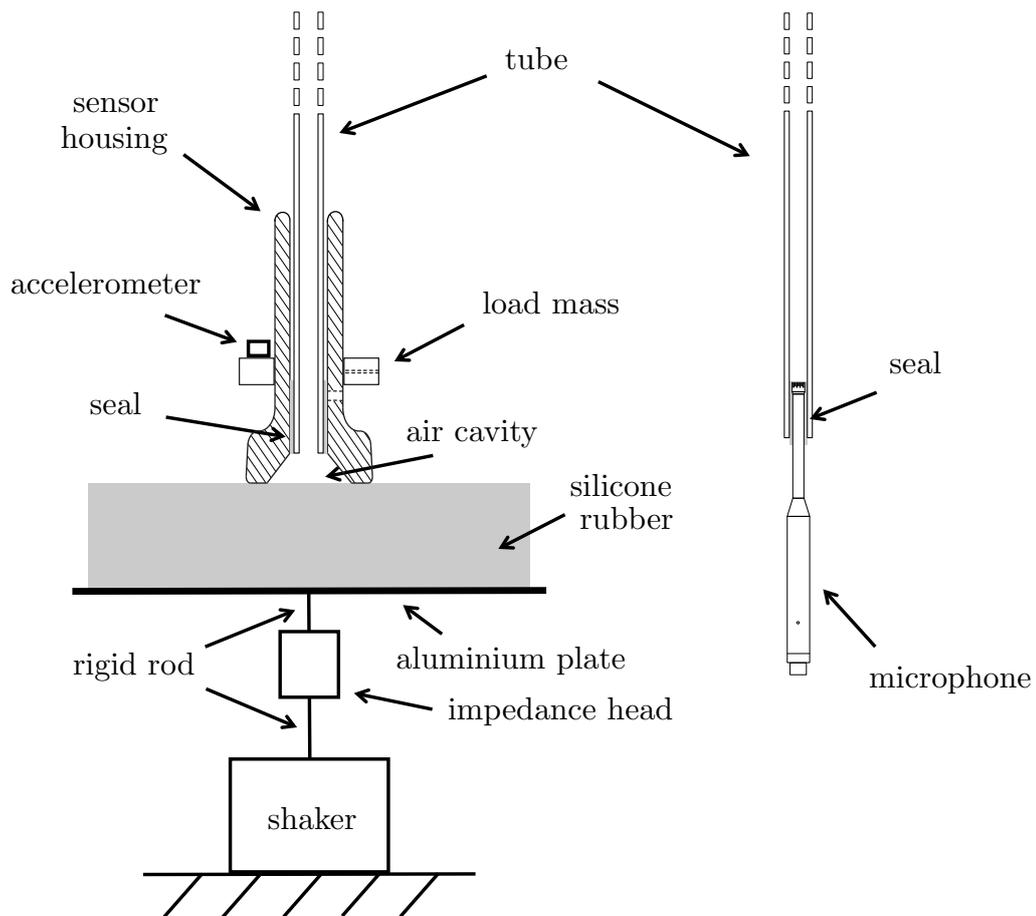


Fig. 10.2 Schematic of the experimental setup for experiments to characterise the performance of an air-coupled sensor with tubing on the phantom. Four mechanical signals are measured: the force and acceleration at the impedance head, the acceleration of the sensor housing and the pressure at the distal end of the tube.

10.4 Results

Figure 10.3 shows the response of an air-coupled stethoscope with and without tubing. For the tubing case a 395 mm long tube is used. The without-tubing case (i.e. an ACM) is shown for load masses of 63 g and 97 g to illustrate how the sensor mass only affects the low frequency response. Unlike the ACMs (which are freely resting on the phantom), the mass of the the ACS with tubing cannot be accurately determined, due to the fact that the tubing is flexible and the distal end is supported. For the case shown we estimated the effective mass seen at the phantom to be around 74 g. As illustrated by the two load-mass cases for the ACM, the chestpiece impedance has a negligible effect on the high frequency response.

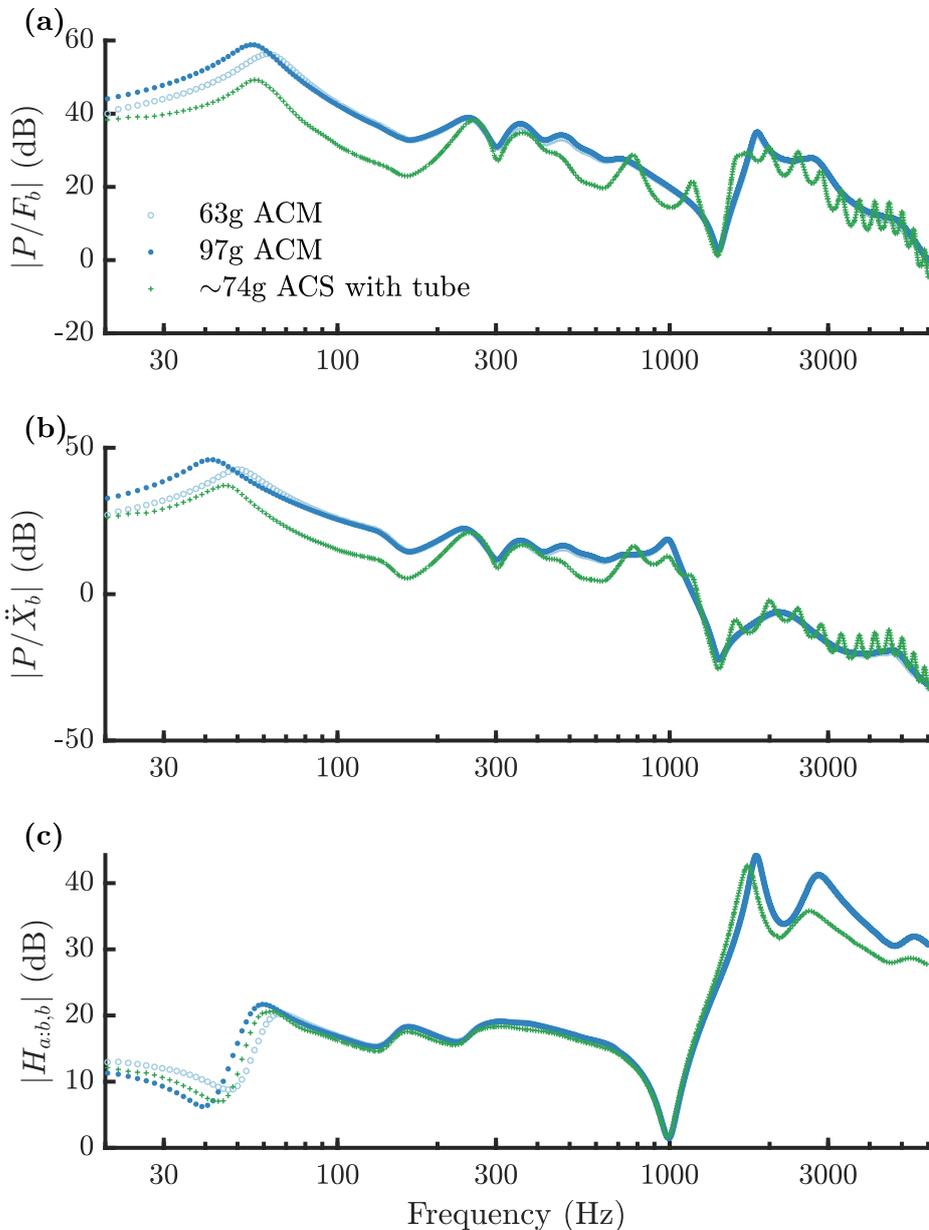


Fig. 10.3 Phantom response of an air-coupled sensor with and without tubing. **(a)** Response of the measured pressure with respect to force at the impedance head. **(b)** Response of the measured pressure with respect to acceleration at the impedance head. **(c)** Driving point acceleration FRF at the impedance head. Data are shown for an air-coupled sensor with a 2220 mm^3 chestpiece air-cavity volume, a 16 mm cavity contact diameter, and a 3.9 mm rim thickness. The no-tubing case is shown for sensor masses of 63 g and 97 g. The with-tubing case was set up as described in section 10.3, using a tube with a length of 395 mm and an inner diameter of 4.3 mm. The approximate sensor mass for the with-tubing case is $\sim 74 \text{ g}$. The decibel reference values are 1 m^{-2} , 1 kg m^{-2} and 1 kg^{-1} for figures a-c respectively. The H1 estimator of the FRF is shown in each case, with a frequency domain resolution of 1.5625 Hz.

It can be observed that the dominant effects of the tubing are to reduce the average level of the response and to introduce additional resonances. These additional resonances can be identified as standing-wave resonances arising in the tubing. The standing wave pattern can be more easily identified on a linear frequency scale, as shown in figure 10.4. The standing wave peaks are separated by an average frequency of ~ 425 Hz. By comparison, the frequency of an acoustic wave with a half-wavelength of 395 mm is ~ 435 Hz.

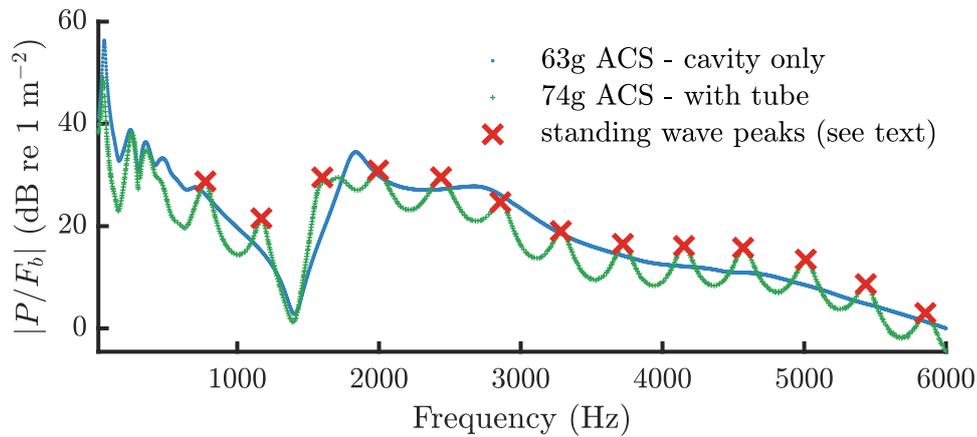


Fig. 10.4 Figure 10.3(a) redrawn on a linear frequency scale. For the ACS with tubing the peaks that can be clearly identified as standing wave resonances are shown by red crosses. In the low frequency region a standing wave resonance may be obscured by other resonances in the phantom response.

The third resonance marked in figure 10.4 falls between an antiresonance resonance pair identified as an artefact related to the mounting of the phantom on the impedance head (see section 5.3.3 and compare to figure 10.3(c)). For the ACS with tubing the ‘mounting-resonance’ occurs at a slightly lower frequency than for the ACM case and almost obscures the standing wave peak.

Inspection of the resonances in the response suggests that the lowest frequency standing wave resonance may be obscured by other resonances in the phantom-ACM response (in the 200 - 500 Hz region). To demonstrate the effect of the standing wave-resonances more clearly, figure 10.5 shows the response of an ACS with and without ~ 60 mm tubing on the alternative phantom described in Nussbaumer et al. (2019) [112]. In this alternative phantom, the base and wall of an Ecoflex 00-10 cylinder are grounded, while a shaker excites a rigid plate embedded within the Ecoflex cylinder, and an accelerometer records a reference signal on a second plate embedded just below the top surface of the phantom.

The input force is not measured and figure 10.5 shows the response of the pressure measured by the ACS to the acceleration of the embedded accelerometer. Figure 10.5 exhibits the same properties observed in figure 10.4. The presence of the tubing causes an overall attenuation of the response, which dominates at low frequencies, while at high frequencies standing wave resonances create a pattern of peaks. In figure 10.5 the lowest frequency standing-wave resonance is unobscured by other resonances and can be clearly identified.

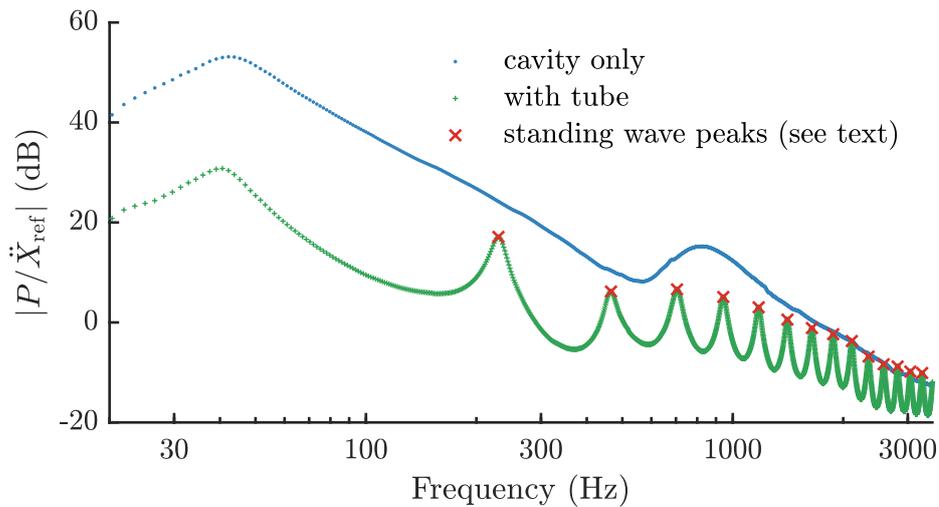


Fig. 10.5 Response of the measured ACS pressure to the acceleration of an accelerometer embedded below the surface of the alternative phantom described in [112] (see text). For the case with tubing, peaks that can be clearly identified as standing wave resonances are shown by red crosses.

In figure 10.6 the standing wave frequencies extracted from figures 10.4 and 10.5 are non-dimensionalised as $(\omega_{\text{peak}}/2\pi)(2l/c_0)$. For a closed-closed tube of length l we would expect this non-dimensional measure of the peak frequency to equal the peak number (as pressure anti-nodes at both ends of the tube yield standing waves at integer multiples of a half-wavelength, see e.g., [6]). For both cases shown in figure 10.6 the measured non-dimensional frequency is slightly lower than the closed-closed case. This matches our expectations as the finite load-impedances at each end of the tube mean that the ends of the tube are not pressure anti-nodes. A load impedance which is stiffness dominated in the frequency range of interest will lower all of the standing wave resonance-frequencies in this frequency range with respect to the infinite-load-impedance case. The deviation from the infinite-impedance case increases with frequency, as is observed in figure 10.6. This can be interpreted as an increase in the effective length of the tube. (Note that air-cavity length to the tubing length raises the

non-dimensional peak frequencies by an insignificant amount, as shown by the grey and green dots for the 14th peak in figure 10.6.)

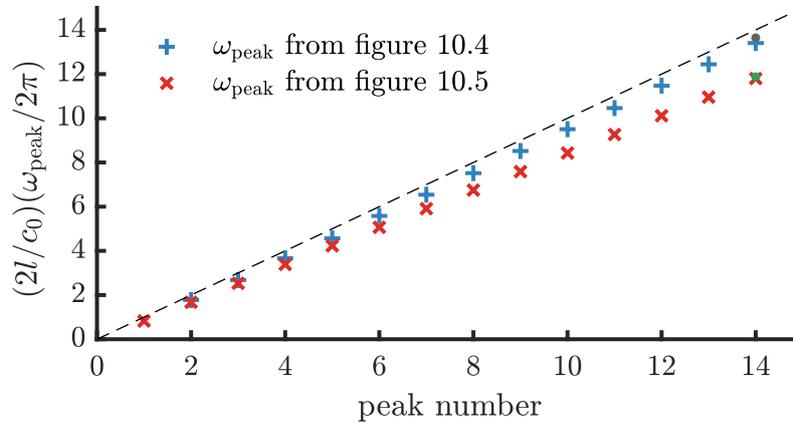


Fig. 10.6 Non-dimensionalised peak frequencies assumed to correspond to standing-wave resonances, extracted from figures 10.4 and 10.5. For l the measured tubing length between the air-cavity and the microphone is used. The grey and green dots show the effect of including the length of the air-cavity (for the 14th peak only).

10.5 Model

Using the lumped element model for the coupling of an ACM to the phantom from section 9.4, adapted to include the tubing as shown in section 10.1, we obtain the simulated response functions shown in figure 10.7.

Comparing figure 10.7 with figure 10.3, it can be seen that the model captures the dominant effects of adding tubing to an air-coupled sensor, including the low-frequency attenuation and the standing-wave resonances. There is a small discrepancy between the modelled and measured standing-wave resonance frequencies, which would be expected as these are sensitive to the impedances seen at both sides of the tubing. At the microphone end the model assumes an infinite impedance, while at the chestpiece end the model gives only a crude estimate of the impedance seen by the tube.

As the ACM (i.e. air-cavity only) model used in this analysis does not accurately simulate the measured response across the entire frequency range of interest (see chapter 9), a direct comparison to the measured data is difficult. We address this issue in section 10.5.1.

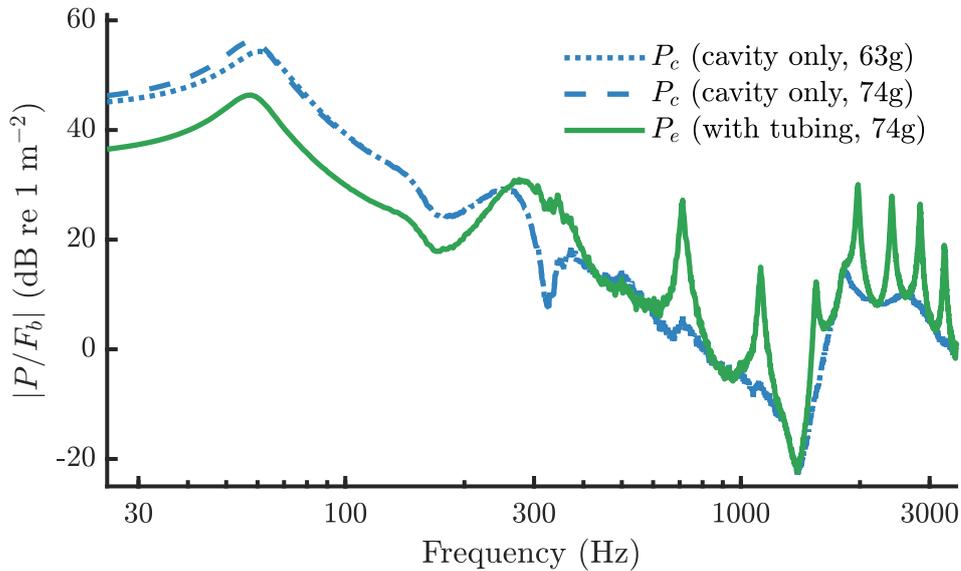


Fig. 10.7 Simulated response of ACSs on the phantom according to equation (10.1), using the attenuation model in appendix D.7. The blue dashed and dotted lines show ACMs with dimensions matching those in figure 10.3, and 63 g and 74 g chestpiece masses respectively. The green solid line shows a simulation with the same cavity volume and the addition of a 395 mm long, 4.3 mm diameter tube with the pressure measured at a rigid termination at the distal end.

10.5.1 Transmission matrix between the impedance head and the air cavity entrance.

In order to validate our model for the tubing without considering the effects of other aspects of the phantom-sensor system, we compute a transmission matrix ($\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$) between the impedance head measurements at the base of the phantom ($F_{k_{b,o}}$ and \dot{X}_b) and the pressure and volume velocity at the entrance of the air cavity ($P_{\text{cav}_{\text{in}}}$ and $Q_{\text{cav}_{\text{in}}}$).

$\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ is computed in a similar manner to the transmission matrix for the DCS described in section 6.5.1, by using ACMs with different volumes to create different impedance boundary conditions. In this case, we have a measurement of the load potential (rather than the load flow as was the case for the DCS transmission matrix computation), so the equations need to be adapted:

$$\begin{bmatrix} F_{k_b,o} \\ \dot{X}_b \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}} \begin{bmatrix} P_{\text{cav}_{\text{in}}} \\ P_{\text{cav}_{\text{in}}}/Z_{\text{cav}_{\text{in}}}^a \end{bmatrix}. \quad (10.10)$$

Dividing through by the cavity input pressure leads to:

$$\begin{bmatrix} F_{k_b,o}/P_{\text{cav}_{\text{in}}} \\ \dot{X}_b/P_{\text{cav}_{\text{in}}} \end{bmatrix} = \underbrace{\begin{bmatrix} A & B \\ C & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}} \begin{bmatrix} 1 \\ Y_{\text{cav}_{\text{in}}}^a \end{bmatrix}, \quad (10.11)$$

where $Y_{\text{cav}_{\text{in}}}^a = 1/Z_{\text{cav}_{\text{in}}}^a$. To arrive at four equations for the four unknowns we consider two different load impedances (denoted by I and II) by using two different cavity volumes, with all other parameters kept the same:

$$\begin{bmatrix} 1 & (Y_{\text{cav}_{\text{in}}}^a)_I \\ 1 & (Y_{\text{cav}_{\text{in}}}^a)_{II} \end{bmatrix} \underbrace{\begin{bmatrix} A & C \\ B & D \end{bmatrix}}_{\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}} = \begin{bmatrix} (F_{k_b,o}/P_{\text{cav}_{\text{in}}})_I & (\dot{X}_b/P_{\text{cav}_{\text{in}}})_I \\ (F_{k_b,o}/P_{\text{cav}_{\text{in}}})_{II} & (\dot{X}_b/P_{\text{cav}_{\text{in}}})_{II} \end{bmatrix}. \quad (10.12)$$

To solve for A , B , C and D we need to measure the response for two different air-cavity volumes and specify the values of $Y_{\text{cav}_{\text{in}}}^a$ for these air cavities. To do this we model the air cavities as pure compliances ($Y_{\text{cav}_{\text{in}}}^a = i\omega V/\rho_0 c_0^2$) using their nominal volumes. This simplification is expected to lead to some error as the small mass and damping terms of the acoustic impedance of the air-cavity are not taken into consideration.

Our theory suggests that $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ is a function of m_s and $z_{s,0}$. This suggests that $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ needs to be evaluated for every value of m_s and $z_{s,0}$ that we wish to model. In practice, m_s and $z_{s,0}$ mainly affect the solution at low frequencies. Consequently, using $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ for an incorrect value of m_s leads to errors in the prediction of the dominant low-frequency peak, but is unlikely to significantly affect the high-frequency response simulation.

The dashed light green line in figure 10.8 shows the simulated response of the measured pressure for a 64 g ACS with tubing with dimensions matching the ACS for which experimental data is shown in figures 10.3 and 10.4. In this example, $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ is computed using equation (10.12), using data for two 63 g ACMs with air-cavity volumes of 1240 mm³ and 2220 mm³ respectively. Both of these cases are shown in figure 10.8.

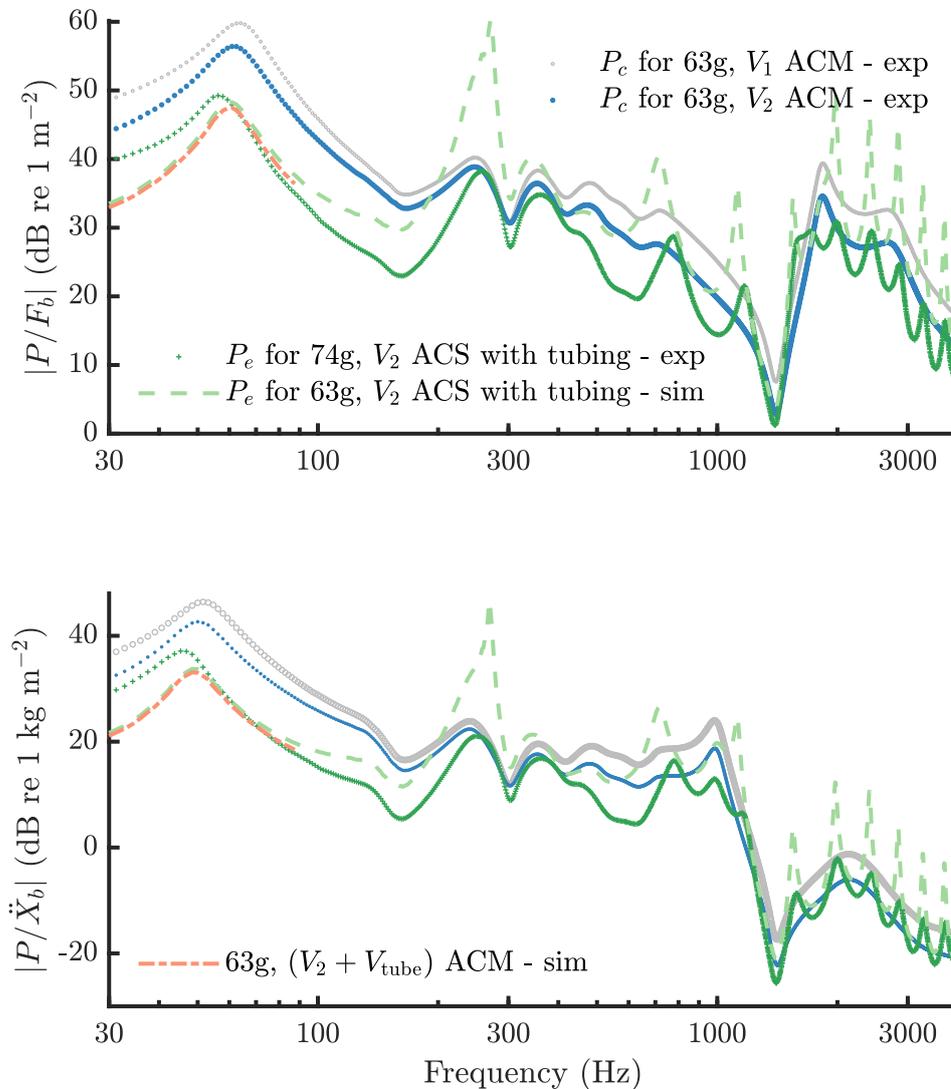


Fig. 10.8 Simulated response of a 63 g ACS with tubing according to equation (10.4), using the simple tubing attenuation model given in appendix D.7. All sensor dimensions are as specified in figure 10.7. $\tilde{\mathbf{T}}_{b,cav_{in}}$ is computed according to the method described in section 10.5.1, using FRFs measured for two ACMs with the same cavity contact area, rim thickness and chestpiece mass ($m_s = 63$ g), but different cavity volumes ($V_1 = 1240$ mm³, $V_2 = 2220$ mm³). The experimental FRFs for both of these ACMs and for a 74 g ACS with tubing are shown for comparison to the simulated response. For both the experimental and simulated cases, the tube length is 395 mm and the tube diameter is 4.3 mm. For low frequencies ($\omega < 2\pi c_0/10l$) the red dot-dashed line shows the results of a simulation in which the tubing is regarded as a lumped compliance.

By accounting for the transmission of vibrations from the base of the phantom to the interface between the phantom and the air-cavity, the use of $\tilde{\mathbf{T}}_{b,\text{cav}_{\text{in}}}$ allows us to isolate the effect of the tubing. The difference between the frequencies of the dominant low frequency peak between the measured and simulated responses can be attributed partly to the difference in m_s . Below this resonance, the simulation significantly underestimates the magnitude of the response. This could be due to the fact that in the simulation the microphone backplate is effectively connected to the sensor housing while in the experimental case it is grounded.

At higher frequencies the simulation performs reasonably but tends to overestimate the response. The simple attenuation model employed here, which assumes a rigid tube and considers only viscous and thermal losses within the duct, significantly underestimates the attenuation. This leads to an overestimation of the baseline response (at frequencies above around 80 Hz), as well as a significant overestimation of the amplitude of the standing-wave peaks. A more physical consideration of the damping mechanisms, including effect of flexible, finite impedance walls, breakout noise and bends in the duct can be expected to account for this mismatch and yield a better simulation of the response. This suggests that the attenuation properties of the tubing play a very significant role in determining the overall response of an analogue stethoscope. A further factor leading to a mismatch between the measured and simulated response levels is the assumption of an infinite load impedance in the model. A non-rigid (finite impedance) termination is expected to yield lower pressures than a rigid termination as the latter is a pressure maximum for all frequencies [6]. A complex load impedance will also have the effect of shifting the standing wave resonances frequencies, as will the complex characteristic impedance and wavenumber in the lossy tube [162].

The phantom-sensor system resonances between 200 and 500 Hz, are poorly matched by this simulation approached, which may in part be due to the fact that the computed transmission matrix between the phantom base and the air cavity is less accurate at these frequencies. Considering the 1.3 - 2 kHz band, the discrepancies around the dominant antiresonance-resonance pair in the driving point response (and in particular the frequency of the system peak between the 4th and 5th standing wave resonances) can be attributed to a slight difference in the mounting of the phantom on the impedance head between the two ACM cases and the ACS-with-tubing case (see section 5.3.3 for details).

In section 3.2.3, we argue that for $l \lesssim \lambda/10$ a lumped element approximation can provide a reasonable model of the response. The red dot-dashed line in figure 10.8 shows the effect of modelling the tube as a pure compliances for frequencies at which $l < \lambda/10$. Comparing this to the duct-simulation (green dashed line), demonstrates that, at low frequencies, the effect of the tubing is dominated by its volume.

10.6 Application to real-world analogue stethoscopes

In this chapter we illustrate how a mechanical model for the interaction between a stethoscope chestpiece and the chest can be coupled to an acoustic model for the stethoscope's tubing. Some nuances, such as the mechanical coupling of the tubing to the chestpiece, are omitted and a simplified tubing model is used for illustrative purposes.

In order to apply these principles to the design or modelling of real-world analogue stethoscopes, a detailed model for the tubing is required. This includes the geometry of the connection between the cavity and the tubing, the bifurcation of the tubing (to reach two separate earpieces) and an accurate model for the attenuation. Furthermore, real-world analogue stethoscopes are typically coupled to the ears of a clinician at the distal end of the tubing, rather than to a microphone as is the case for the results presented in this chapter.

The effect of coupling to the ear can be included by using an appropriate model for the load impedance at the distal end of the tube in the equations presented in section 10.1. The impedance of the human ear can be characterised as consisting of two main parts: the ear canal and the middle ear [155, 162]. The ear canal can be characterised as a short lossy duct with dimensions that vary between subjects while the middle ear is a complex term with multiple resonances [155, 162]. For example, Withnell and Gowdy found that the impedance of the middle ear is usually stiffness dominated below around 1 kHz and mass dominated above 4 kHz, with several small resonances in between these regions and significant inter-subject variation [162]. Combining the ear canal and the middle ear into a single load impedance yields a system that is stiffness dominated throughout the frequency range of primary interest in auscultation (i.e. to beyond 3 kHz). This

suggests that the effect of replacing a microphone at the end of the tube with a human ear will lead to a slight decrease in the overall level of the response, and a downward shift in the standing wave frequencies across the frequency range of interest.

10.7 Summary and conclusions

In this chapter we validate our model for air-coupled sensors with tubing, while demonstrating the importance of accurately capturing the attenuation coefficient and the finite load impedance.

Our model shows that at low frequencies the effect of including tubing in an ACS is dominated by the increased internal volume of the stethoscope and the attenuation along the tube. The larger the volume of the tubes, the smaller the net acoustic stiffness seen at the chest surface, reducing the response. This suggests that a compromise must be found, as narrow tubes which minimise the added volume also incur greater damping losses. At higher frequencies, the effect of including tubing is dominated by the occurrence of standing-wave resonances. We show that for a tube terminated by a microphone these resonances occur at frequencies corresponding to $\lambda = n l_{\text{eff}}/2$, where $n = 1, 2, 3, \dots$ and the effective length l_{eff} is longer than the duct length l due to the finite load impedances at both ends of the tube. Experiments in which the ears of a human act as the termination are suggested as future work.

In this chapter we focus on the methodology for modelling air-coupled sensors with tubing rather than considering individual design decisions. In general, design decisions involving tubing consist of a compromise between the effects of added volume and tube damping for low frequencies, while at high frequencies the location of the standing wave resonances depends on the tubing length and the load impedance. The prominence of standing wave resonances depends on the damping level in the tubing. When considering the design of analogue stethoscopes, the ear-impedance can be expected to be significantly lower than the impedance presented by a nominally rigid microphone. This is expected to lead to a further increase in the effective length of the tube and a consequent decrease in the frequencies of the standing wave resonances.

A mixed mode analogue-digital stethoscope which uses a side-mounted microphone in the tubing of an analogue stethoscope can be shown by our model to yield a significantly poorer performance than an ACM in which the microphone is mounted in the chestpiece-air cavity itself, or after a section of tubing that is as short as possible.

Chapter 11

Conclusions

The objectives of this thesis were to explain how stethoscopes work, to derive a low-order model for the response of stethoscopes to body sounds, and to develop experimental methods to validate the model.

The chest is a non-linear system. Existing work on stethoscope acoustics ignores this, leading to models which are unable to capture the effect of application pressure on the response of stethoscopes. Our experiments confirm that the impedance seen by a sensor on the chest varies as a function of the static load with which the sensor is applied. We argue that, while the chest is non-linear with respect to the length scales associated with pressing the stethoscope into the chest, body sounds can be considered as small, linearised perturbations around the established equilibrium position. The small size of the stethoscope chestpiece, with respect to the typical wavelengths of body sounds, justifies the use of a lumped-element approach to model the interaction between the chestpiece and the chest.

We present the design of a chest phantom which can be used to characterise the performance of stethoscopes in controlled experiments. Impedance measurements on the phantom and on a human chest verify that the phantom is a suitable model for the chest. Furthermore, these measurements allow differences between the two systems to be accounted for. We show that, for small changes in application pressure, the properties of the phantom are approximately linear. This allows us to compute transmission matrices between the input signals at the base of the phantom and the sensors tested on the phantom. These transmission

matrices allow us to directly compare measurements on the phantom with the lumped-element models presented in this thesis.

For air-coupled stethoscopes, we show that our model accurately captures the trends associated with each of the key design parameters as well as the effects of using a diaphragm and tubing. We show that the mass of the chestpiece plays a significant role in setting the dominant low frequency resonance in the response. Increasing the cavity volume lowers the measured response, and leads to a slight reduction in the resonance frequencies of the coupled system. These changes can be explained by the decreasing stiffness of the air cavity.

Our model suggests that the dominant effect of adding a diaphragm to a stethoscope is to increase the stiffness between the air cavity surface and the chestpiece rim. This leads to an increase in the resonance frequencies of the coupled system. Furthermore, the high stiffness of the diaphragm means that a lower proportion of the net force perturbations exerted on the chestpiece are exerted through the air cavity, thus reducing the level of the response. However, the use of a diaphragm also prevents the air cavity from being occluded by chest tissue. This allows stethoscopes with diaphragms to employ significantly shallower cavities. The reduction in cavity volume increases the level of the response, while also increasing the net stiffness of the system. By using both a diaphragm and a shallower cavity, the system resonance can be shifted to higher frequencies without a dramatic reduction in the level of the response.

Adding tubing to an air-coupled stethoscope is found to attenuate the overall response level, especially at low frequencies, while adding standing-wave resonances to the response. Using our model, we are able to demonstrate that at low frequencies the effect of tubing can be captured by considering two factors: the net increase in the internal volume of the stethoscope, and the energy dissipation along the tube. This leads to a compromise between thin tubes which minimise the internal volume and wider tubes which induce lower levels of damping. At high frequencies, the model presented in this thesis is able to capture the formation of standing-wave resonances. The frequencies of these resonances are shown to be sensitive to the impedances at both ends of the tube. This suggests that using the same stethoscopes on different subjects, or even at different locations on the chest of the same subject, will lead to standing wave resonances at different frequencies.

We demonstrate that holding the stethoscope provides both a static force and a damping-dominated impedance load on the chestpiece. The widely used experimental method of controlling the static load by adding mass to the chestpiece is shown to yield erroneous conclusions with regards to the effect of static load on the low-frequency response of stethoscopes.

The signal measured by a stethoscope (or received at the ears of a clinician) depends strongly on a wide variety of factors related to its use for a particular auscultation rather than just its stand-alone properties. As a result of this there is no single answer to the question “what is the response of a given stethoscope to body sounds?”. This means that the way in which a stethoscope is to be used needs to be carefully considered in the design stage and that the quantitative comparison of auscultation signals is of limited value when investigating body sounds.

When comparing stethoscopes, simply using two stethoscopes on the chest in turn is not valid unless the application pressure is controlled and the variability of the body sounds is accounted for. The comparison of two stethoscopes side by side on the chest must also be treated with care as the local mechanical properties of the chest may vary between the two sites. For digital stethoscopes, there is an additional difficulty, which is that the response can be amplified to an arbitrary level. Consequently, signal-to-noise ratios are more informative than the amplitudes of the recorded signals.

Future work

The work presented in this thesis opens up pathways for further investigation which can improve the accuracy of stethoscope-response models. In particular, we suggest that comparing measurements conducted across different phantom designs would allow a full validation of the approaches that we have introduced to account for the phantom’s behaviour.

Measurements of the impedance seen at the human chest are limited to a single human subject in this thesis. An extension of this research to a larger cohort of subjects would allow the inter-subject variability of chest impedance to be characterised. This would enable us to investigate how the performance of a stethoscope varies between individuals. A further extension to this work would

be to characterise the performance of sensors directly on the human chest by using statistical methods to account for the variabilities of the input signal (natural body sounds). This would allow us to test our hypothesis that, once impedance differences are accounted for, the results from phantom experiments can be applied to the chest.

The dependance of the stiffness between the cavity surface and the chestpiece rim on the static load applied and on the inner and outer diameters of the rim could be investigated in more detail. This would allow the effect of design changes relating to these parameters to be characterised using the models presented in this thesis. For stethoscopes with tubing, this thesis has considered only the case of electronic stethoscopes, in which a microphone is used at the distal end of the tube. To extend the analysis to analogue stethoscopes the effect of the ears has to be taken into consideration.

This thesis focuses on stethoscopes which are designed to be held by a clinician. There is a growing interest in the use of wearable devices to record heart and lung sounds. This includes utilising different sensor modalities including ultra-lightweight acceleration sensors and flexible sensors. The experimental techniques in this thesis can be used to characterise the performance of these sensors, and the modelling approaches introduced here could be extended to model the interaction between these sensors and the chest.

References

- [1] Abella, M., Formolo, J., and Penney, D. G. (1992). Comparison of the acoustic properties of six popular stethoscopes. *The Journal of the Acoustical Society of America*, 91(4):2224–2228.
- [2] Anusha, A. R., Soodi, A. L., and Kumar, S. P. (2012). Design of low-cost hardware for lung sound acquisition and determination of inspiratory-expiratory phase using respiratory waveform. In *2012 Third International Conference on Computing, Communication and Networking Technologies (ICCCNT'12)*, pages 1–5. IEEE.
- [3] Aygün, H. and Apolskis, A. (2020). The quality and reliability of the mechanical stethoscopes and Laser Doppler Vibrometer (LDV) to record tracheal sounds. *Applied Acoustics*, 161:107159.
- [4] Bakhshaei, H., Garreau, G., Tognetti, G., Shoele, K., Carrero, R., Kilmar, T., Zhu, C., Thompson, W. R., Seo, J. H., Mittal, R., and Others (2015). Mechanical design, instrumentation and measurements from a hemoacoustic cardiac phantom. In *2015 49th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE.
- [5] Bárány, E. (1942). On the mechanical impedance of the human thorax. *Acta Medica Scandinavica*, 111(3):252–260.
- [6] Beranek, L. L. and Mellow, T. (2012). *Acoustics: sound fields and transducers*. Academic Press.
- [7] Bishop, P. J. (1980). Evolution of the stethoscope. *Journal of the Royal Society of Medicine*, 73(6):448–456.
- [8] Bohadana, A., Izbicki, G., and Kraman, S. S. (2014). Fundamentals of lung auscultation. *New England Journal of Medicine*, 370(8):744–751.
- [9] Boonvisut, P. and Çavuşoğlu, M. C. (2012). Estimation of soft tissue mechanical parameters from robotic manipulation data. *IEEE/ASME Transactions on Mechatronics*, 18(5):1602–1611.
- [10] Brandt, A. (2011). *Noise and vibration analysis: signal analysis and experimental procedures*. John Wiley & Sons.
- [11] Brownjohn, J. M. W., Steele, G. H., Cawley, P., and Adams, R. D. (1980). Errors in mechanical impedance data obtained with impedance heads. *Journal of Sound and Vibration*, 73(3):461–468.

- [12] Buonsenso, D., Pata, D., and Chiaretti, A. (2020). COVID-19 outbreak: less stethoscope, more ultrasound. *The Lancet Respiratory Medicine*, 8(5):e27.
- [13] Burch, C. R. and Stock, J. P. P. (1961). A new diaphragmatic stethoscope. *British heart journal*, 23(4):447.
- [14] Cafarelli, A., Miloro, P., Verbeni, A., Carbone, M., and Menciassi, A. (2016). Speed of sound in rubber-based materials for ultrasonic phantoms. *Journal of ultrasound*, 19(4):251–256.
- [15] Cefaly, R. (1963). Acoustic stethoscope. U.S. Patent No 3,109,508. 5 Nov. 1963.
- [16] Charles, P. D., Esper, G. J., Davis, T. J., Maciunas, R. J., and Robertson, D. (1999). Classification of tremor and update on treatment. *American family physician*, 59(6):1565.
- [17] Cook, J. R., Bouchard, R. R., and Emelianov, S. Y. (2011). Tissue-mimicking phantoms for photoacoustic and ultrasonic imaging. *Biomedical optics express*, 2(11):3193–3206.
- [18] Craiem, D. and Magin, R. L. (2010). Fractional order models of viscoelasticity as an alternative in the analysis of red blood cell (RBC) membrane mechanics. *Physical biology*, 7(1):13001.
- [19] Dai, Z. (2013). *The Audible Human Project: Modeling Sound Transmission in the Lungs and Torso*. PhD thesis, University of Illinois at Chicago.
- [20] Dai, Z., Peng, Y., Mansy, H. A., and Royston, T. J. (2013). Sound Transmission in a Lung Phantom Model. In *ASME International Mechanical Engineering Congress and Exposition*, volume 56222, page V03BT03A048. American Society of Mechanical Engineers.
- [21] D’Arcy, J. L., Coffey, S., Loudon, M. A., Kennedy, A., Pearson-Stuttard, J., Birks, J., Frangou, E., Farmer, A. J., Mant, D., Wilson, J., Myerson, S. G., and Prendergast, B. D. (2016). Large-scale community echocardiographic screening reveals a major burden of undiagnosed valvular heart disease in older people: The OxVALVE Population Cohort Study. *European Heart Journal*, 37(47):3515–3522a.
- [22] Darwish, A. and Hassanien, A. E. (2011). Wearable and implantable wireless sensor network solutions for healthcare monitoring. *Sensors*, 11(6):5561–5595.
- [23] Ding, X., Clifton, D., Ji, N., Lovell, N. H., Bonato, P., Chen, W., Yu, X., Xue, Z., Xiang, T., Long, X., and Others (2020). Wearable sensing and telehealth technology with potential applications in the coronavirus pandemic. *IEEE reviews in biomedical engineering*, 14:48–70.
- [24] Dodge, R. R. and Burrows, B. (1980). The prevalence and incidence of asthma and asthma-like symptoms in a general population sample. *American Review of Respiratory Disease*, 122(4):567–575.

- [25] Dokumacı, E. (2014). On the effect of viscosity and thermal conductivity on sound propagation in ducts: A re-visit to the classical theory with extensions for higher order modes and presence of mean flow. *Journal of Sound and Vibration*, 333(21):5583–5599.
- [26] Dokumacı, E. (2021). *Duct Acoustics: Fundamentals and Applications to Mufflers and Silencers*. Cambridge University Press.
- [27] Dong, R. G., Wu, J. Z., and Welcome, D. E. (2005). Recent advances in biodynamics of human hand-arm system. *Industrial health*, 43(3):449–471.
- [28] Dowling, A. P. and Ffowcs Williams, J. E. (1983). *Sound and sources of sound*. Horwood.
- [29] Ertel, P. Y., Lawrence, M., Brown, R. K., and Stern, A. M. (1966). Stethoscope acoustics: II. Transmission and filtration patterns. *Circulation*, 34(5):899–909.
- [30] Ertel, P. Y., Lawrence, M., and Song, W. (1971). Stethoscope acoustics and the engineer: Concepts and problems. *Journal of the Audio Engineering Society*, 19(3):182–186.
- [31] Ewins, D. J. (2000). *Modal testing: theory, practice and application*. Research Studies Press Ltd.
- [32] Fenelon, L., Holcroft, L., Waters, N., and Others (2009). Contamination of stethoscopes with MRSA and current disinfection practices. *Journal of Hospital Infection*, 71(4):376–378.
- [33] Fiz, J. A., Gnitecki, J., Kraman, S. S., Wodicka, G. R., and Pasterkamp, H. (2008). Effect of body position on lung sounds in healthy young men. *Chest*, 133(3):729–736.
- [34] Fouillet, Y., Parent, C., Gropplero, G., Davoust, L., Achard, J. L., Revol-Cavalier, F., and Verplanck, N. (2017). Stretchable material for microfluidic applications. In *Multidisciplinary Digital Publishing Institute Proceedings*, volume 1, page 501.
- [35] Fowler, J. K. (1921). *Pulmonary tuberculosis*. Macmillan & Company.
- [36] Fuster, V. (2016). The stethoscope’s prognosis: very much alive and very necessary. *Journal of the American College of Cardiology*, 67(9):1118–1119.
- [37] Gardezi, S. K. M., Myerson, S. G., Chambers, J., Coffey, S., D’Arcy, J., Hobbs, F. D. R., Holt, J., Kennedy, A., Loudon, M., Prendergast, A., and Others (2018). Cardiac auscultation poorly predicts the presence of valvular heart disease in asymptomatic primary care patients. *Heart*, 104(22):1832–1835.
- [38] Gautam, N. K. and Pokle, S. B. (2011). A review and comparative analysis of techniques for diagnosis of pulmonary diseases. *International Journal of Computational Intelligence and Bioinformatics*, 3(1):43–70.
- [39] Geddes, L. A. (2005). Birth of the stethoscope. *IEEE Engineering in Medicine and Biology Magazine*, 24(1):84–86.

- [40] Gregory, A. L. (2019). *A Theory for Wheezing in Lungs*. PhD thesis, University of Cambridge.
- [41] Grundlehner, B. and Buxi, D. (2011). Methods to characterize sensors for capturing body sounds. In *2011 International Conference on Body Sensor Networks*, pages 59–64. IEEE.
- [42] Gupta, P., Moghimi, M. J., Jeong, Y., Gupta, D., Inan, O. T., and Ayazi, F. (2020). Precision wearable accelerometer contact microphones for longitudinal monitoring of mechano-acoustic cardiopulmonary signals. *NPJ digital medicine*, 3(1):1–8.
- [43] Gurung, A., Scrafford, C. G., Tielsch, J. M., Levine, O. S., and Checkley, W. (2011). Computerized lung sound analysis as diagnostic aid for the detection of abnormal lung sounds: a systematic review and meta-analysis. *Respiratory medicine*, 105(9):1396–1403.
- [44] Håkansson, B. and Carlsson, P. (1987). Bias errors in mechanical impedance data obtained with impedance heads. *Journal of Sound and Vibration*, 113(1):173–183.
- [45] Halloran, L. J. S. (2011). *A Rheological Study of Stress Relaxation in Elastomers for in situ X-Ray Diffraction Measurements*. PhD thesis, McGill University Library.
- [46] Hampton, C. S. and Chaloner, A. (1967). Which stethoscope? *British medical journal*, 4(5576):388.
- [47] Hanley, B. P., Bains, W., and Church, G. (2019). Review of scientific self-experimentation: ethics history, regulation, scenarios, and views among ethics committees and prominent scientists. *Rejuvenation research*, 22(1):31–42.
- [48] Hanley, W. B. and Hanley, A. J. G. (2000). The efficacy of stethoscope placement when not in use: traditional versus “cool”. *Cmaj*, 163(12):1562–1563.
- [49] Hatton, M. Q. F., Allen, M. B., Vathenen, A. S., Mellor, E., and Cooke, N. J. (1994). Does sedation help in fiberoptic bronchoscopy? *Bmj*, 309(6963):1206–1207.
- [50] Hirahara, T. and Ooya, T. (2010). High-sensitivity body-conducted sound sensor. In *International Conference on Natural Polymers (ICNP2010), Kottayam, India*. In Kottayam, India, pages IL141–1–IL141–4.
- [51] Hirahara, T., Otani, M., Shimizu, S., Toda, T., Nakamura, K., Nakajima, Y., and Shikano, K. (2010). Silent-speech enhancement using body-conducted vocal-tract resonance signals. *Speech Communication*, 52(4):301–313.
- [52] Hooda, A. K., Nath, C. S., and Kolhe, V. S. (1996). Systolic murmurs in apparently healthy individuals. *Medical Journal Armed Forces India*, 52(1):40–44.
- [53] Howell, W. L. and Aldridge, C. F. (1965). The Effect of Stethoscope-Applied Pressure in Auscultation: A New Instrument for Improving Discrimination. *Circulation*, 32(3):430–434.

- [54] Hu, Y., Kim, E. G., Cao, G., Liu, S., and Xu, Y. (2014). Physiological acoustic sensing based on accelerometers: A survey for mobile healthcare. *Annals of biomedical engineering*, 42(11):2264–2277.
- [55] Huang, L., Choy, Y. S., So, R., and Chong, T. (2000). Experimental study of sound propagation in a flexible duct. *the Journal of the Acoustical Society of America*, 108(2):624–631.
- [56] Hussein, H. M. G. (1979). *Measurement of the mechanical impedance of human soft tissue in vivo*. PhD thesis, Louisiana State University and Agricultural and Mechanical College.
- [57] International Organization for Standardization (1997). Mechanical vibration and shock - Human exposure - Biodynamic coordinate systems (BS ISO 827:1997+A1:2015).
- [58] International Organization for Standardization (2012). Mechanical vibration and shock - Mechanical impedance of the human hand-arm system at the driving point (ISO 10068:2012(E)).
- [59] Israr, A., Choi, S., and Tan, H. Z. (2006). Detection threshold and mechanical impedance of the hand in a pen-hold posture. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 472–477. IEEE.
- [60] Israr, A., Choi, S., and Tan, H. Z. (2007). Mechanical impedance of the hand holding a spherical tool at threshold and suprathreshold stimulation levels. In *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'07)*, pages 56–60. IEEE.
- [61] Jain, P. K., Tiwari, A. K., and Chourasia, V. S. (2016). Performance analysis of seismocardiography for heart sound signal recording in noisy scenarios. *Journal of medical engineering & technology*, 40(3):106–118.
- [62] Jiang, C., Zhao, J., Huang, B., Zhu, J., and Yu, J. (2021). A basic investigation into the optimization of cylindrical tubes used as acoustic stethoscopes for auscultation in COVID-19 diagnosis. *The Journal of the Acoustical Society of America*, 149(1):66–69.
- [63] Jin, F., Mu, D., Chu, D., Fu, E., Xie, Y., and Liu, T. (2008). Severe complications of bronchoscopy. *Respiration*, 76(4):429–433.
- [64] Jones, A., Jones, D., Kwong, K., and Siu, S. C. (1998). Acoustic performance of three stethoscope chest pieces. In *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Vol. 20 Biomedical Engineering Towards the Year 2000 and Beyond (Cat. No. 98CH36286)*, volume 6, pages 3219–3222. IEEE.
- [65] Joyashiki, T. and Wada, C. (2018). Estimation of the frequency Response of Body-conducted sound sensors via the equivalent circuit model. *Med. Biol. Eng*, 56:44–51.
- [66] Joyashiki, T. and Wada, C. (2020). Validation of a body-conducted sound sensor for respiratory sound monitoring and a comparison with several sensors. *Sensors*, 20(3):942.

- [67] Kaniusas, E. (2015). Sensing by Acoustic Biosignals. In *Biomedical Signals and Sensors II*, pages 1–90. Springer.
- [68] Kato, H., Suzuki, A., Nakajima, Y., Makino, H., Sanjo, Y., Nakai, T., Shiraishi, Y., Katoh, T., and Sato, S. (2009). A visual stethoscope to detect the position of the tracheal tube. *Anesthesia & Analgesia*, 109(6):1836–1842.
- [69] Kay, E. (2018). *Cardiac acoustics: understanding and detecting heart murmurs*. PhD thesis, University of Cambridge.
- [70] Kay, E. and Agarwal, A. (2017). DropConnected neural networks trained on time-frequency and inter-beat features for classifying heart sounds. *Physiological measurement*, 38(8):1645.
- [71] Keller, J. P., Kody, R. S., Rogers, D. J., and Watrous, R. L. (2020). Stethoscope. U.S. Patent No 10,667,782. 2 June 2020.
- [72] Kindig, J. R., Beeson, T. P., Campbell, R. W., Andries, F., and Tavel, M. E. (1982). Acoustical performance of the stethoscope: a comparative analysis. *American heart journal*, 104(2):269–275.
- [73] Klum, M., Stehling, J., Pielmus, A., Tigges, T., and Orglmeister, R. (2019). Validation System for Digital Stethoscopes. *Current Directions in Biomedical Engineering*, 5(1):261–264.
- [74] Kocher, K. E., Meurer, W. J., Fazel, R., Scott, P. A., Krumholz, H. M., and Nallamothu, B. K. (2011). National trends in use of computed tomography in the emergency department. *Annals of emergency medicine*, 58(5):452–462.
- [75] Kohn, M. A. (2014). Understanding evidence-based diagnosis. *Diagnosis*, 1(1):39–42.
- [76] Kraman, S. S. (2008). Transmission of lung sounds through light clothing. *Respiration*, 75(1):85–88.
- [77] Kraman, S. S., Pressler, G. A., Pasterkamp, H., and Wodicka, G. R. (2006a). Design, construction, and evaluation of a bioacoustic transducer testing (BATT) system for respiratory sounds. *IEEE transactions on biomedical engineering*, 53(8):1711–1715.
- [78] Kraman, S. S., Wodicka, G. R., Oh, Y., and Pasterkamp, H. (1995). Measurement of respiratory acoustic signals: effect of microphone air cavity width, shape, and venting. *Chest*, 108(4):1004–1008.
- [79] Kraman, S. S., Wodicka, G. R., Pressler, G. A., and Pasterkamp, H. (2006b). Comparison of lung sound transducers using a bioacoustic transducer testing system. *Journal of Applied Physiology*, 101(2):469–476.
- [80] Laennec, R. T. H. (2005). Medicine in Stamps René Laennec (1781-1826): inventor of the stethoscope. *Singapore Med J*, 46(3):107.
- [81] Layton, J. (2021). How Stethoscopes Work, <https://science.howstuffworks.com/innovation/everyday-innovations/stethoscopes.htm>, Accessed: 15 May 2021.

- [82] Li, H., Ren, Y., Zhang, G., Wang, R., Zhang, X., Zhang, T., Zhang, L., Cui, J., Xu, Q., and Duan, S. (2019). Design of a high SNR electronic heart sound sensor based on a MEMS bionic hydrophone. *AIP Advances*, 9(1):15005.
- [83] Li, S.-H., Lin, B.-S., Tsai, C.-H., Yang, C.-T., and Lin, B.-S. (2017). Design of wearable breathing sound monitoring system for real-time wheeze detection. *Sensors*, 17(1):171.
- [84] Littmann, D. (1961). An approach to the ideal stethoscope. *Jama*, 178(5):504–505.
- [85] Liu, C., Griffiths, C., Murray, A., and Zheng, D. (2016). Comparison of stethoscope bell and diaphragm, and of stethoscope tube length, for clinical blood pressure measurement. *Blood pressure monitoring*, 21(3):178.
- [86] Longtin, Y., Schneider, A., Tschopp, C., Renzi, G., Gayet-Ageron, A., Schrenzel, J., and Pittet, D. (2014). Contamination of stethoscopes and physicians’ hands after a physical examination. In *Mayo Clinic Proceedings*, volume 89, pages 291–299. Elsevier.
- [87] Maia, N. M. M. and e Silva, J. M. M. (1997). *Theoretical and experimental modal analysis*. Research Studies Press.
- [88] Majumdar, S. K. (2002). History of the stethoscope an overview. *Bull Indian Inst Hist Med Hyderabad*, 32(2):137–150.
- [89] Mangion, K. (2007). The stethoscope. *Malta Medical Journal*, 19(2):41–44.
- [90] Mangione, S. (2001). Cardiac auscultatory skills of physicians-in-training: a comparison of three English-speaking countries. *The American journal of medicine*, 110(3):210–216.
- [91] Mangione, S. (2012). The stethoscope as metaphor. *Cleveland Clinic journal of medicine*, 79(8):545–546.
- [92] Mangione, S. and Nieman, L. Z. (1999). Pulmonary auscultatory skills during training in internal medicine and family practice. *American journal of respiratory and critical care medicine*, 159(4):1119–1124.
- [93] Mansy, H. A. (2016). Smart Stethoscope Systems: A New Paradigm for Bedside Patient Monitoring. *J Bioeng Biomed Sci*, 6:187.
- [94] Mansy, H. A., Grahe, J., Elke, D. A., and Sandler, R. H. (2009). Design of a sound source phantom with uniform surface signal. In *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*, pages 987–990. IEEE.
- [95] Mansy, H. A., Grahe, J., Royston, T. J., and Sandler, R. H. (2011). Investigating a compact phantom and setup for testing body sound transducers. *Computers in biology and medicine*, 41(6):361–366.
- [96] Mansy, H. A., Grahe, J. R., and Sandler, R. H. (2008). Elastic properties of synthetic materials for soft tissue modeling. *Physics in Medicine & Biology*, 53(8):2115.

- [97] Markel, H. (2006). The stethoscope and the art of listening. *The New England journal of medicine*, 354(6):551.
- [98] McGee, S. (2016). *Evidence-based physical diagnosis e-book*. Elsevier Health Sciences.
- [99] McMechan, C. and So, P. (2011). Design and implementation of a low cost electronic stethoscope. In *Proceedings of 2011 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, pages 714–718. IEEE.
- [100] Miller, M. R., Hankinson, J., Brusasco, V., Burgos, F., Casaburi, R., Coates, A., Crapo, R., Enright, P., Van Der Grinten, C. P. M., Gustafsson, P., and Others (2005). Standardisation of spirometry. *European respiratory journal*, 26(2):319–338.
- [101] Mimani, A. and Munjal, M. L. (2012). Acoustical analysis of a general network of multi-port elements—an impedance matrix approach. *International Journal of Acoustics and Vibration*, 17(1):23.
- [102] Montinari, M. R. and Minelli, S. (2019). The first 200 years of cardiac auscultation and future perspectives. *Journal of multidisciplinary healthcare*, 12:183.
- [103] Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., and Wagenmakers, E.-J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic bulletin & review*, 23(1):103–123.
- [104] Murphy, R. L. H. (2008). In defense of the stethoscope. *Respiratory Care*, 53(3):355–369.
- [105] Nelson, G. (2015). *Stethoscope Design for Auscultation in High Noise Environments*. PhD thesis, University of Minnesota.
- [106] Nelson, G., Rajamani, R., and Erdman, A. (2015). Vibro-acoustic model of a piezoelectric-based stethoscope for chest sound measurements. *Measurement Science and Technology*, 26(9):95701.
- [107] Nowak, K. M. and Nowak, L. J. (2016a). On the relation between pressure applied to the chest piece of a stethoscope and parameters of the transmitted bioacoustic signals. In *Proceedings of the 22nd International Congress on Acoustics*, Buenos Aires.
- [108] Nowak, L. J. and Nowak, K. M. (2016b). Discovering new facts and revealing existing myths about the acoustic stethoscope 200 years after its invention. In *Proceedings of the 22nd International Congress on Acoustics*, Buenos Aires.
- [109] Nowak, L. J. and Nowak, K. M. (2017). Acoustic characterization of stethoscopes using auscultation sounds as test signals. *The Journal of the Acoustical Society of America*, 141(3):1940–1946.
- [110] Nowak, L. J. and Nowak, K. M. (2018). Sound differences between electronic and acoustic stethoscopes. *Biomedical engineering online*, 17(1):1–11.

- [111] Nowak, L. J. and Nowak, K. M. (2019). An experimental study on the role and function of the diaphragm in modern acoustic stethoscopes. *Applied Acoustics*, 155:24–31.
- [112] Nussbaumer, M., Troyas Martinez, L., and Agarwal, A. (2019). *A theory for stethoscope acoustics*. Universitätsbibliothek der RWTH Aachen, Aachen, Germany.
- [113] O’Flaherty, N. and Fenelon, L. (2015). The stethoscope and healthcare-associated infection: a snake in the grass or innocent bystander? *Journal of Hospital Infection*, 91(1):1–7.
- [114] Ogawa, T. (2013). Acoustic wave velocity measurement on piezoelectric ceramics. *Piezoelectric materials and devices—Practice and applications*, pages 35–55.
- [115] Oliveira, A. and Marques, A. (2014). Respiratory sounds in healthy people: a systematic review. *Respiratory medicine*, 108(4):550–570.
- [116] Oliynik, V. (2015). On potential effectiveness of integration of 3M Littmann 3200 electronic stethoscopes into the third-party diagnostic systems with auscultation signal processing. In *2015 IEEE 35th International Conference on Electronics and Nanotechnology (ELNANO)*, pages 417–421. IEEE.
- [117] Ou, D., OuYang, L., Tan, Z., Mo, H., Tian, X., and Xu, X. (2016). An electronic stethoscope for heart diseases based on micro-electro-mechanical-system microphone. In *2016 IEEE 14th International Conference on Industrial Informatics (INDIN)*, pages 882–885. IEEE.
- [118] Packard, T. J. (1984). Tunable stethoscope. U.S. Patent No 4,440,258. 3 Apr. 1984.
- [119] Padmanabhan, V., Semmlow, J. L., and Welkowitz, W. (1993). Accelerometer type cardiac transducer for detection of low-level heart sounds. *IEEE transactions on biomedical engineering*, 40(1):21–28.
- [120] Pasterkamp, H., Kraman, S. S., DeFrain, P. D., and Wodicka, G. R. (1993). Measurement of respiratory acoustical signals: comparison of sensors. *Chest*, 104(5):1518–1525.
- [121] Pasterkamp, H., Kraman, S. S., and Wodicka, G. R. (1997). Respiratory sounds: advances beyond the stethoscope. *American journal of respiratory and critical care medicine*, 156(3):974–987.
- [122] Patel, L., Gandhi, D., and Beddow, D. (2021). Controversies on the Stethoscope During COVID-19: A Necessary Tool or an Unnecessary Evil? *The American Journal of the Medical Sciences*, 361(2):278.
- [123] Pavlosky, A., Glauche, J., Chambers, S., Al-Alawi, M., Yanev, K., and Loubani, T. (2018). Validation of an effective, low cost, Free/open access 3D-printed stethoscope. *PLoS One*, 13(3):e0193087.
- [124] Picano, E. and Vano, E. (2011). The radiation issue in cardiology: the time for action is now. *Cardiovascular ultrasound*, 9(1):1–13.

- [125] Pierce, A. D. (2019). *Acoustics: an introduction to its physical principles and applications*. Springer.
- [126] Piirila, P., Sovijarvi, A. R. A., Earis, J. E., Rossi, M., Dalmasso, F., Stoneman, S. A. T., and Vanderschoot, J. (2000). Reporting results of respiratory sound analysis. *European Respiratory Review*, 10(77):636–640.
- [127] Popkin, R., Lobo, F., and Stubbs, J. (2019). Multimaterial 3D Printing for the Fabrication of Functional Stethoscopes. In *Frontiers in Biomedical Devices*, volume 41037, page V001T10A015. American Society of Mechanical Engineers.
- [128] Rakheja, S., Wu, J. Z., Dong, R. G., Schopper, A. W., and Boileau, P.-É. (2002). A comparison of biodynamic models of the human hand–arm system for applications to hand-held power tools. *Journal of Sound and Vibration*, 249(1):55–82.
- [129] Rappaport, M. B. and Sprague, H. B. (1941). Physiologic and physical laws that govern auscultation, and their clinical application: The acoustic stethoscope and the electrical amplifying stethoscope and stethograph. *American Heart Journal*, 21(3):257–318.
- [130] Reeves, W., Hilmer, C., and Miller, D. R. (1994). Disposable sensing device with contaneous conformance.
- [131] Riederer, K. and Backman, J. (1998). Frequency response of stethoscopes. In *Proc. Nordic Acoust. Meeting*.
- [132] Roguin, A. (2006). Rene Theophile Hyacinthe Laënnec (1781–1826): the man behind the stethoscope. *Clinical medicine & research*, 4(3):230–235.
- [133] Schmidt, S. E., Zimmermann, H., Hansen, J., Møller, H., Hammershøi, D., and Struijk, J. J. (2012). The chest is a significant collector of ambient noise in heart sound recordings. In *2012 Computing in Cardiology*, pages 741–744. IEEE.
- [134] Schwartz, R. S., Reeves, J. T., Sodal, I. E., and Barnes, F. S. (1980). Improved phonocardiogram system based on acoustic impedance matching. *American Journal of Physiology-Heart and Circulatory Physiology*, 238(4):H604—H609.
- [135] Semmlow, J. L. (2015). Improved heart sound detection and signal-to-noise estimation using a low-mass sensor. *IEEE Transactions on Biomedical Engineering*, 63(3):647–652.
- [136] Smith, C. (2002a). Transducer for sensing body sounds. U.S. Patent No. 6,498,854. 24 Dec. 2002.
- [137] Smith, C. (2003). Transducer for sensing body sounds. U.S. Patent No. 6,661,897. 9 Dec. 2003.
- [138] Smith, M. C. (2002b). Synthesis of mechanical networks: the inerter. *IEEE Transactions on automatic control*, 47(10):1648–1662.

- [139] Smith, M. C. (2020). The inerter: A retrospective. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:361–391.
- [140] Solomon, S. D. and Saldana, F. (2014). Point-of-care ultrasound in medical education—stop listening and look. *The New England journal of medicine*, 370(12):1083–1085.
- [141] Sprague, H. B. (1926). A new combined stethoscope chest piece. *Journal of the American Medical Association*, 86(25):1909.
- [142] Suzuki, A. and Nakayama, K. (1998). The absolute sensitivity characteristics of lung-sound transducers coupled to chest wall. In *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Vol. 20 Biomedical Engineering Towards the Year 2000 and Beyond (Cat. No. 98CH36286)*, volume 4, pages 1727–1730. IEEE.
- [143] Suzuki, A. and Nakayama, K. (2000). Characteristics of lung-sound transducers on chest wall. *Med. Electron. Biol. Eng.*, 38:298–308.
- [144] Suzumura, N. and Ikegaya, K. (1977). Characteristics of the air cavities of phonocardiographic microphones and the effects of vibration and room noise. *Medical and Biological Engineering and Computing*, 15(3):240–247.
- [145] Swarup, S. and Makaryus, A. N. (2018). Digital stethoscope: technology update. *Medical devices (Auckland, NZ)*, 11:29.
- [146] Taebi, A., Solar, B. E., Bomar, A. J., Sandler, R. H., and Mansy, H. A. (2019). Recent advances in seismocardiography. *Vibration*, 2(1):64–86.
- [147] Taylor, J. (1997). *Introduction to error analysis, the study of uncertainties in physical measurements*. University Science Books.
- [148] Thompson, D. E., Hussein, H. M., and Perritt, R. Q. (1981). Point impedance characterization of soft tissues in vivo. In *Bioengineering and the Skin*, pages 103–111. Springer.
- [149] Thompson, S. C., Gabrielson, T. B., and Warren, D. M. (2014). Analog model for thermoviscous propagation in a cylindrical tube. *The Journal of the Acoustical Society of America*, 135(2):585–590.
- [150] Traister, R. S., Fajt, M. L., Landsittel, D., and Petrov, A. A. (2014). A novel scoring system to distinguish vocal cord dysfunction from asthma. *The Journal of Allergy and Clinical Immunology: In Practice*, 2(1):65–69.
- [151] Vaicekauskaite, J., Mazurek, P., Vudayagiri, S., and Skov, A. L. (2020). Mapping the mechanical and electrical properties of commercial silicone elastomer formulations for stretchable transducers. *Journal of Materials Chemistry C*, 8(4):1273–1279.
- [152] Vasudevan, R. S., Horiuchi, Y., Torriani, F. J., Cotter, B., Maisel, S. M., Dadwal, S. S., Gaynes, R., and Maisel, A. S. (2020). Persistent Value of the Stethoscope in the Age of COVID-19. *The American journal of medicine*.

- [153] Vermarien, H. and Van Vollenhoven, E. (1984). The recording of heart vibrations: A problem of vibration measurement on soft tissue. *Medical and Biological Engineering and Computing*, 22(2):168–178.
- [154] von Gierke, H. E. and Brammer, A. J. (2002). Effects of shock and vibration on humans. *Harris' Shock and Vibration Handbook*, pages 42.1 – 42.61.
- [155] Voss, S. E., Rosowski, J. J., Merchant, S. N., and Peake, W. T. (2000). Acoustic responses of the human middle ear. *Hearing research*, 150(1-2):43–69.
- [156] Wang, K., Ho, C.-C., Zhang, C., and Wang, B. (2017). A review on the 3D printing of functional structures for medical phantoms and regenerated tissue and organ applications. *Engineering*, 3(5):653–662.
- [157] Watrous, R. L., Grove, D. M., and Bowen, D. L. (2002). Methods and results in characterizing electronic stethoscopes. In *Computers in cardiology*, pages 653–656. IEEE.
- [158] Weiss, D., Erie, C., Butera III, J., Copt, R., Yeaw, G., Harpster, M., Hughes, J., and Salem, D. N. (2019). An in vitro acoustic analysis and comparison of popular stethoscopes. *Medical Devices (Auckland, NZ)*, 12:41.
- [159] Weitz, H. H. and Mangione, S. (2000). In defense of the stethoscope and the bedside. *The American journal of medicine*, 108(8):669–671.
- [160] Welsby, P. D. and Earis, J. E. (2001). Some high pitched thoughts on chest examination. *Postgraduate Medical Journal*, 77(912):617–620.
- [161] Welsby, P. D., Parry, G., and Smith, D. (2003). The stethoscope: some preliminary investigations. *Postgraduate medical journal*, 79(938):695–698.
- [162] Withnell, R. H. and Gowdy, L. E. (2013). An analysis of the acoustic input impedance of the ear. *Journal of the Association for Research in Otolaryngology*, 14(5):611–622.
- [163] Wodicka, G. R., Kraman, S. S., Zenk, G. M., and Pasterkamp, H. (1994). Measurement of respiratory acoustic signals: effect of microphone air cavity depth. *Chest*, 106(4):1140–1144.
- [164] Wong, V. (2018). Mobile device-based stethoscope system. U.S. Patent No. 9,973,847. 15 May 2018.
- [165] World Health Organization (2020). The top 10 causes of death www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death. Accessed: 25 May 2021.
- [166] Yan, W.-L., Lopez, G., Shuzo, M., and Yamada, I. (2012). Development of a wearable acoustic respiration sensor using flexible piezoelectric film. *MIPE2012*, pages 413–415.
- [167] Yang, C., Zhang, W., Pang, Z., Zhang, J., Zou, D., Zhang, X., Guo, S., Wan, J., Wang, K., and Pang, W. (2021). A Low-Cost, Ear-Contactless Electronic Stethoscope Powered by Raspberry Pi for Auscultation of Patients With COVID-19: Prototype Development and Feasibility Study. *JMIR Medical Informatics*, 9(1):e22753.

- [168] Yasar, T. K., Royston, T. J., and Magin, R. L. (2013). Wideband MR elastography for viscoelasticity model identification. *Magnetic resonance in medicine*, 70(2):479–489.
- [169] Yilmaz, G., Rapin, M., Pessoa, D., Rocha, B. M., de Sousa, A. M., Rusconi, R., Carvalho, P., Wacker, J., Paiva, R. P., and Chételat, O. (2020). A Wearable Stethoscope for Long-Term Ambulatory Respiratory Health Monitoring. *Sensors*, 20(18):5124.
- [170] Yun, H., Kim, H., Kim, E., and Jun, M. B. G. (2020). Development of internal sound sensor using stethoscope and its applications for machine monitoring. *Procedia Manufacturing*, 48:1072–1078.
- [171] Zanartu, M., Ho, J. C., Kraman, S. S., Pasterkamp, H., Huber, J. E., and Wodicka, G. R. (2008). Air-borne and tissue-borne sensitivities of bioacoustic sensors used on the skin surface. *IEEE Transactions on Biomedical Engineering*, 56(2):443–451.
- [172] Zhu, J., Tan, Y., Huang, B., Zhu, Y., and Gao, X.-H. (2020). Don't throw the stethoscope away! The value of the stethoscope over ultrasound equipment during the COVID-19 outbreak is discussed by authors from China. *European Heart Journal*.
- [173] Zimmermann, H., Moller, H., Hansen, J., and Others (2010). Assessment of chest impedance in relation to Phonocardiography. In *Proceedings of 20th International Congress on Acoustics*.
- [174] Zimmermann, N. H., Møller, H., Hansen, J., Hammershøi, D., and Rubak, P. (2011). Effect of rib-cage structure on acoustic chest impedance. In *Proceedings of Forum Acusticum 2011*, pages 49–54. European Acoustics Association-EAA.

Appendix A

Derivations

A.1 Plane waves in ducts

Sound waves in air obey the acoustic wave equation (see e.g. Beranek (2012) [6] or Dowling (2000) [28]):

$$\frac{1}{c_0^2} \frac{\partial^2 p'}{\partial t^2} - \nabla^2 p' = 0. \quad (\text{A.1})$$

In equation (A.1) p' is the acoustic pressure perturbation such that the pressure $p = p_0 + p'$, where p_0 is the equilibrium pressure, c_0 is the speed of sound in air, ∇^2 is the laplacian operator, and ∂ is the partial differential operator.

For waves in a circular duct we adopt a cylindrical-polar coordinate system where x is coaxial with the duct, r is the radial coordinate and θ is the azimuthal angle. We can then assume a separable solution to equation (A.1) of the form

$$p(r, \theta, x, t) = f(r)g(x)e^{i(\omega t - n\theta)}. \quad (\text{A.2})$$

Here $f(r)$ is a function of r and $g(x)$ is a function of x , ω is the angular frequency, t is time and n is the azimuthal mode number. Substituting this solution into the acoustic wave equation using the cylindrical polar expression for the laplacian gives

$$\frac{1}{rf} \frac{d}{dr} (rf \dot{f}(r)) - \frac{n^2}{r^2} + \frac{\omega^2}{c_0^2} = -\frac{\ddot{g}(x)}{g}. \quad (\text{A.3})$$

To avoid confusion with the prime symbol used to denote a small perturbation, \dot{f} is used to refer to the first derivative of f with respect to its argument r (rather than with respect to t as is the case in Newton's notation). Likewise \ddot{g} refers to the second derivative of g with respect to its argument x .

The left hand side of equation (A.3) is a function of r only, while the right hand side is a function of x only. This condition can only be satisfied if both sides of the equation equal a constant. We let this constant equal k^2 . This yields two separate equations for x and r respectively:

$$\ddot{g} + k^2 g = 0, \quad (\text{A.4})$$

$$\frac{1}{rf} \frac{d}{dr} (rf \dot{f}(r)) - \frac{n^2}{r^2} + \frac{\omega^2}{c_0^2} - k^2 = 0. \quad (\text{A.5})$$

The solution for the x dependency in equation (A.4) is given by

$$g(x) = \hat{B}_f e^{-ikx} + \hat{B}_b e^{ikx} \quad (\text{A.6})$$

In equation (A.6) the constants \hat{B}_f and \hat{B}_b can be interpreted as the magnitudes of forward and backward traveling waves respectively.

Equation (A.5) is Bessel's equation. Solutions which satisfy the condition of being finite on $r = 0$ are denoted by $f(r) = J_n(\hat{\alpha}r)$ where

$$\hat{\alpha}^2 = \frac{\omega^2}{c_0^2} - k^2. \quad (\text{A.7})$$

J_n is a Bessel function and the subscript n relates to the order of the azimuthal mode. To obtain a value for $\hat{\alpha}$ and complete the dispersion relation it is necessary to consider the boundary conditions. The kinematic boundary condition is that the normal component of velocity on the wall ($r = a$) is zero. This can be expressed in terms of a condition of zero normal pressure gradient

$$\partial p' / \partial r = 0 \quad \text{on} \quad r = a, \quad (\text{A.8})$$

which leads to:

$$\dot{f}(a) = J_n(\hat{\alpha}a) = 0. \quad (\text{A.9})$$

If the ϕ th zero of J_n is denoted by $z_{\phi n}$, then the dispersion relationship (equation (A.5)) can be written as:

$$\frac{z_{\phi n}^2}{a^2} = \frac{\omega^2}{c_0^2} - k_{\phi n}^2. \quad (\text{A.10})$$

For $n = 0$ there is a root at $z_{1,0} = 0$. For this root there is no dependence of p' on either θ or r , which can be physically interpreted as a *plane* wave which is uniform across the cross-section of the duct. In this case the dispersion relationship simplifies to

$$k^2 = \frac{\omega^2}{c_0^2}. \quad (\text{A.11})$$

Equation (A.11) shows that plane waves in ducts are non-dispersive ($\partial\omega/\partial k = \omega/k = c_0$) and can propagate (k is real).

For higher order modes, if ω is below a certain value, the solution for k will be imaginary, meaning evanescent waves decaying in the x direction. This gives rise to the concept of a ‘cut-off’ frequency, below which waves of a given mode do not propagate. The mode with the lowest ‘cut-off’ frequency is the mode corresponding to $z_{1,1}$. This mode only propagates if $\omega > 1.841c_0/a$ (see e.g. [125]).

If we are interested only in frequencies below the first ‘cut-off’ frequency then only plane-waves will propagate. This amounts to setting $\dot{f}(r) = 0$ and $n = 0$ in the general solution given in equation A.2. If we then substitute $g(x)$ from equation A.6 we obtain

$$p'(x, t) = \hat{B}_f e^{i(\omega t - kx)} + \hat{B}_b e^{i(\omega t + kx)}, \quad (\text{A.12})$$

which is a general expression for plane wave propagation in a duct in terms of forwards and backwards travelling waves. Equation A.12 is the classical solution to the 1D wave equation with \hat{B}_f and \hat{B}_b representing ‘forward’ and ‘backward’ propagating waves respectively.

For acoustic perturbations conservation of momentum is given by

$$\rho_0 \frac{\partial u'}{\partial t} + \nabla p' = 0, \quad (\text{A.13})$$

where ∇ is the grad operator. For plane wave propagation $\nabla p'$ becomes $\frac{\partial p'}{\partial x}$ and the solution for velocity perturbations is:

$$u'(x, t) = \frac{\hat{B}_f}{\rho_0 c_0} e^{i(\omega t - kx)} - \frac{\hat{B}_b}{\rho_0 c_0} e^{i(\omega t + kx)}. \quad (\text{A.14})$$

From equation A.14 it is clear that the forward and backward traveling waves remain separated in terms of their velocity perturbations. For forward traveling waves the velocity perturbation is in phase with the pressure perturbation, while for backward traveling waves the velocity perturbation is 180° out of phase with the pressure perturbation. The characteristic specific impedance of a medium is defined as $Z_0^s = \rho_0 c_0$.

For a segment of duct with a constant area we can solve for $p'(l, t)$ and $u'(l, t)$ as functions of $p'(0, t)$ and $u'(0, t)$ as

$$p'(l, t) = \frac{p'(0, t)}{2}(e^{-ikl} + e^{ikl}) + \rho_0 c_0 \frac{u'(0, t)}{2}(e^{-ikl} - e^{ikl}), \quad (\text{A.15})$$

and

$$u'(l, t) = \frac{1}{\rho_0 c_0} \frac{p'(0, t)}{2}(e^{-ikl} - e^{ikl}) + \frac{u'(0, t)}{2}(e^{-ikl} + e^{ikl}). \quad (\text{A.16})$$

Equations (A.15) and (A.16) can be written in terms of the hyperbolic trigonometric functions as

$$p'(l, t) = p'(0, t) \cosh(ikl) - Z_0^s u'(0, t) \sinh(ikl), \quad (\text{A.17})$$

and

$$u'(l, t) = \frac{-1}{Z_0^s} p'(0, t) \sinh(ikl) - u'(0, t) \cosh(ikl). \quad (\text{A.18})$$

A.2 Continuous 1D wave-equation solution for Helmholtz resonator

By employing the transmission matrix formulation derived in section A.1, the 1D wave equation can be solved for the Helmholtz resonator in figure 3.1. Cascading the continuous acoustic transmission matrices for the two sections of duct we obtain

$$\begin{bmatrix} p'_3 \\ q'_3 \end{bmatrix} = \begin{bmatrix} \cosh(ikx_{3,2}) & -Z_{0:A_2}^a \sinh(ikx_{3,2}) \\ -Y_{0:A_2}^a \sinh(ikx_{3,2}) & \cosh(ikx_{3,2}) \end{bmatrix} \begin{bmatrix} \cosh(ikx_{2,1}) & -Z_{0:A_1}^a \sinh(ikx_{2,1}) \\ -Y_{0:A_1}^a \sinh(ikx_{2,1}) & \cosh(ikx_{2,1}) \end{bmatrix} \begin{bmatrix} p'_1 \\ q'_1 \end{bmatrix}. \quad (\text{A.19})$$

Applying a boundary condition of $q'_3 = 0$ for the hard wall at x_3 and dividing equation A.19 by p'_1 , we obtain two equations with two unknowns, which can be solved for p'_3/p'_1 as a function of frequency. The frequency can be non-dimensionalised as l_{\max}/λ , where the wavelength $\lambda = 2\pi c_0/\omega$, and in this example $l_{\max} = x_{2,1}$. This solution is shown by the solid blue line in figure 3.2.

A.3 ‘Dot Method’ for computing the dual of a circuit

The ‘Dot Method’ is a technique for computing the dual of a circuit. The technique is described in detail in *Acoustics Sound Fields and Transducers* (Beranek, 2012) [6]. In order to compute the dual of the circuit, a dot is placed into the centre of each mesh. One further dot is placed outside all of the meshes in the circuit. These dots will become the nodes of the dual circuit.

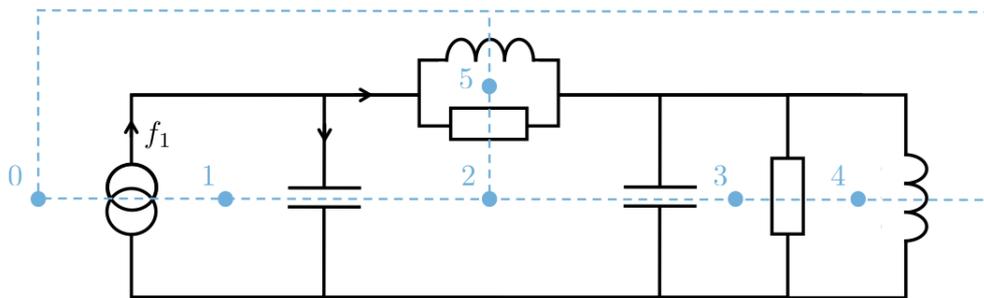


Fig. A.1 Dot method applied to the mobility analogy circuit from figure 3.13. Component values are omitted for clarity, but can be found in figure 3.13.

Lines are then drawn to connect the dots, with a line drawn through each component (impedance or source) to connect the dots on either side of it. In the dual circuit these lines give the connections between nodes, which each impedance replaced by its inverse and each source replaced by a source of the other type (i.e., *flow* versus *potential*). Figure A.1 shows the dot method applied to the mobility circuit in figure 3.13. The resulting impedance analogy circuit is shown in figure A.2.

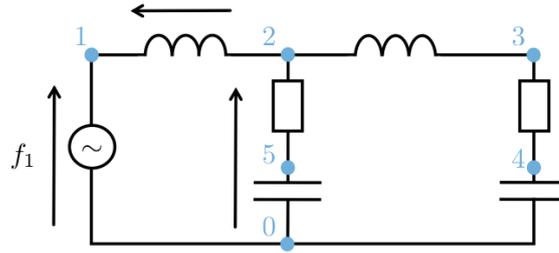


Fig. A.2 Dual of the circuit in figure A.1. Node numbering matches mesh-centre numbering in figure A.1.

A.4 Modal approach

A.4.1 FRF as modal summation

Assuming harmonic excitation and response $\mathbf{f} = \mathbf{F}e^{i\omega t}$ and $\mathbf{x} = \mathbf{X}e^{i\omega t}$, and further assuming that excitation is only applied at the ψ th coordinate, we substitute \mathbf{x} from equation (5.19) into equation (5.22) and then pre-multiply the equation by $\mathbf{u}^{(n)t}$ to obtain

$$-\omega^2\alpha_n + \alpha_n\omega_n^2 = Fu_\psi^{(n)}, \quad (\text{A.20})$$

which leads to

$$\alpha_n = \frac{Fu_\psi^{(n)}}{\omega_n^2 - \omega^2}. \quad (\text{A.21})$$

Substituting α_n from equation A.21 back into the general response (equation 5.19) we obtain

$$\mathbf{x} = F \sum_{n=1}^N \frac{u_\psi^{(n)} \mathbf{u}^{(n)}}{\omega_n^2 - \omega^2} e^{i\omega t}, \quad (\text{A.22})$$

which can be evaluated for the response at ϕ as

$$x_\phi = F \sum_{n=1}^N \frac{u_\psi^{(n)} u_\phi^{(n)}}{\omega_n^2 - \omega^2} e^{i\omega t}. \quad (\text{A.23})$$

Dividing equation A.23 by $F e^{i\omega t}$ yields the (undamped) transfer function relating the response at point ϕ to excitation at point ψ ,

$$H_x(\phi, \psi, \omega) = \sum_{n=1}^N \frac{u_k^{(n)} u_j^{(n)}}{\omega_n^2 - \omega^2}. \quad (\text{A.24})$$

A.4.2 Nyquist plot - viscous damping

The poles of equation (5.26) can be found as the solutions to $(\omega_n^2 + 2ic_n\omega\omega_n - \omega^2) = 0$, which are

$$\omega = \omega_n(\pm\sqrt{1 - c_n^2} + ic_n). \quad (\text{A.25})$$

Substituting into equation (5.19) leads to

$$\mathbf{x}(t) = \sum_{n=1}^N \alpha_n \mathbf{u}^{(n)} e^{i\omega_n(\sqrt{1-c_n^2})t} e^{-\omega_n c_n t}. \quad (\text{A.26})$$

For free vibration the damped natural frequency becomes $\omega_r = \omega_n\sqrt{1 - c_n^2}$, while $\omega_n c_n$ governs the exponential decay of the motion. For $c_n \ll 1$ the damped resonance frequency deviates negligibly from the undamped resonance frequency (i.e. $\omega_r \simeq \omega_n$).

By multiplying equation (5.26) (evaluated for a single mode in the summation) by $i\omega$ to obtain the admittance response, and rearranging the equation slightly we obtain

$$H_v(n, \phi, \psi, \omega) = \frac{a_n(\phi, \psi)}{\omega_n} \frac{i\frac{\omega}{\omega_n}}{1 + 2ic_n\frac{\omega}{\omega_n} - (\frac{\omega}{\omega_n})^2}. \quad (\text{A.27})$$

Simplifying the notation and letting $\hat{p}_n = \frac{\omega}{\omega_n}$ leads to

$$H_v = \frac{a_n}{\omega_n} \frac{i\hat{p}_n}{1 + 2ic_n\hat{p}_n - \hat{p}_n^2}. \quad (\text{A.28})$$

Multiplying the numerator and denominator of the second fraction in equation (A.28) by the complex conjugate of the denominator we obtain

$$H_v = \frac{a_n}{\omega_n} \frac{i\hat{p}_n(1 - \hat{p}_n^2 - 2i\hat{p}_n c_n)}{(1 - \hat{p}_n^2)^2 + (2\hat{p}_n c_n)^2}. \quad (\text{A.29})$$

To simplify the mathematics let $A = 1 - \hat{p}_n^2$ and $B = 2\hat{p}_n c_n$. By substituting these into equation (A.29) we obtain

$$H_v = \frac{a_n}{\omega_n} \frac{\hat{p}_n(B + Ai)}{A^2 + B^2}. \quad (\text{A.30})$$

We now define U and V such that:

$$U = \Re(H_v) - \frac{a_n}{4\omega_n c_n} = \frac{a_n}{\omega_n} \frac{\hat{p}_n B}{A^2 + B^2} - \frac{a_n}{4\omega_n c_n}, \quad (\text{A.31})$$

$$V = \Im(H_v) = \left(\frac{\hat{p}_n A}{A^2 + B^2} \right) \frac{a_n}{\omega_n}, \quad (\text{A.32})$$

where \Re and \Im denote the real and imaginary parts of the function. With some rearrangement it can be shown that

$$U^2 + V^2 = \left(\frac{a_n}{4\omega_n c_n} \right)^2, \quad (\text{A.33})$$

which means that

$$\left(\Re(H_v) - \frac{a_n}{4\omega_n c_n} \right)^2 + \left(\Im(H_v) \right)^2 = \left(\frac{a_n}{4\omega_n c_n} \right)^2. \quad (\text{A.34})$$

equation (A.34) is the equation for a circle in the form $(x - x_c)^2 + (y - y_c)^2 = r^2$. The radius of the circle is $\frac{a_n}{4\omega_n c_n}$ and the centre is at $(\frac{a_n}{4\omega_n c_n}, 0)$. This result is independent of the value of c_n .

A similar analysis can be performed on the accelerance response function, yielding:

$$\left(\Re(H_a) \right)^2 + \left(\Im(H_a) - \frac{\hat{p}_n a_n}{4c_n} \right)^2 = \left(\frac{\hat{p}_n a_n}{4c_n} \right)^2. \quad (\text{A.35})$$

Equation (A.35) differs from equation (A.34) in that \hat{p}_n is a function of frequency, so equation (A.35) is not, in general, an equation for a circle. However, if $\hat{p}_n = \frac{\omega}{\omega_n} \simeq 1$ then equation (A.35) approximates the equation for a circle with radius $\frac{a_n}{4c_n}$ and centre $(0, \frac{a_n}{4c_n})$. This occurs if $c_n \ll 1$, since in this case we have a sharp response peak and all points with non-negligible magnitude occur at frequencies very close to ω_n . If c_n is large, equation (A.35) no longer approximates a circle. Regardless of the value of c_n , the location of $H_a(\omega_n)$ on the Nyquist plot is always $(0, \frac{a_n}{2c_n})$.

For the compliance response the same method yields:

$$\left(\Re(H_x)\right)^2 + \left(\Im(H_x) + \frac{a_n}{4c_n\omega_n^2\hat{p}_n}\right)^2 = \left(\frac{a_n}{4c_n\omega_n^2\hat{p}_n}\right)^2 \quad (\text{A.36})$$

Using the same reasoning as above, if $c_n \ll 1$ then $\hat{p}_n \simeq 1$ and equation (A.36) approximates the equation for a circle. For the compliance response this is a circle of radius $\frac{a_n}{4c_n\omega_n^2}$ and centre $(0, \frac{a_n}{4c_n\omega_n^2})$. However, if c_n is large equation (A.36) no longer results in a circle. Regardless of the value of c_n , the location of $H_a(\omega_n)$ on the Nyquist plot is always $(0, -\frac{a_n}{2c_n\omega_n^2})$.

Appendix B

Signal Processing

B.1 The discrete Fourier transform

Consider a discrete signal $s(n)$, which is periodic, such that

$$s(n + N) = s(n) . \quad (\text{B.1})$$

In equation (B.1), n is the sample counter and N is the number of samples per period.

We define the discrete Fourier transform (DFT) of the signal $s(n)$ as

$$S(\mu) = \sum_{n=0}^{N-1} s(n) e^{-i2\pi\mu n/N} , \quad (\text{B.2})$$

where μ is a counter for the discrete frequency bins. The DFT is periodic with $S(\mu + N) = S(\mu)$.

The inverse DFT of $S(\mu)$ is then given by:

$$s(n) = \frac{1}{N} \sum_{\mu=0}^{N-1} S(\mu) e^{i2\pi\mu n/N} . \quad (\text{B.3})$$

The frequency increment between adjacent frequency bins ($\delta\omega/2\pi$) can be expressed as:

$$\frac{\delta\omega}{2\pi} = \frac{1}{\Delta t} = \frac{1}{N\delta t} . \quad (\text{B.4})$$

where $\Delta t = N\delta t$ is the *measurement time* and δt is the time-step between samples.

The discrete form of Parseval's theorem [10] for two discrete signals $s_\phi(n)$ and $s_\psi(n)$ is given by:

$$\sum_{n=0}^{N-1} s_\phi(n)s_\psi^*(n) = \frac{1}{N} \sum_{\mu=0}^{N-1} S_\phi(\mu)S_\psi^*(\mu) , \quad (\text{B.5})$$

where $*$ denotes the complex conjugate. For $s_\phi = s_\psi$ this yields

$$\sum_{n=0}^{N-1} |s_\phi|^2 = \frac{1}{N} \sum_{\mu=0}^{N-1} |S_\phi|^2 , \quad (\text{B.6})$$

where $| |$ is the modulus operator.

B.2 Power spectral density

B.2.1 PSD definition

Using the DFT definition from section B.1 , with the discrete form of Parseval's theorem, the signal power of discrete signal $s(n)$ is given by

$$\text{Power} = \frac{1}{\Delta t} \sum_{n=0}^{N-1} |s_n|^2 \delta t \quad (\text{B.7})$$

$$= \frac{1}{N} \sum_{n=0}^{N-1} |s_n|^2 = \frac{\Delta t}{N^2} \sum_{\mu=0}^{N-1} |S_\phi|^2 \frac{\delta \omega}{2\pi} . \quad (\text{B.8})$$

The *power spectral density* (PSD) of the signal is then defined such that:

$$\text{Power} = \sum_{\mu=0}^{N-1} \text{PSD} \frac{\delta \omega}{2\pi} , \quad (\text{B.9})$$

yielding

$$\text{PSD} = \frac{\Delta t}{N^2} |S_\phi|^2 . \quad (\text{B.10})$$

In direct analogy to the PSD defined for a signal, we can also define the cross-power spectral density (CPSD) between two signals as

$$\text{CPSD} = S_{\phi,\psi} = \frac{\Delta t}{N^2} S_{\phi} S_{\psi}^* . \quad (\text{B.11})$$

Unlike the PSD, the CPSD can be complex-valued. The phase angle of the CPSD represents the phase-lag between the signals.

B.2.2 PSD estimation: Welch's method

To estimate (C)PSDs from raw time-series data we employ Welch's method (see e.g. [10] for a good introduction). Welch's method is to segment the time-domain signal into \mathcal{T} short overlapping segments, which are each multiplied by a windowing function (to enforce periodicity). If each segment is of length N , and we define the discrete windowing function as $w(n)$, then the DFT of the windowed signal is given by

$$S_{\phi}^{\text{win}}(\mu) = \sum_{n=0}^{N-1} x_{\phi}(n) w(n) e^{-i2\pi\mu n/N} . \quad (\text{B.12})$$

We then compute the mean of the squared magnitude of the windowed DFT and multiply this by a correction factor Υ to arrive at an estimate for the double sided PSD:

$$\tilde{S}_{\phi\phi}(\mu) = \frac{\Upsilon}{\mathcal{T}} \sum_{\tau=1}^{\mathcal{T}} |S_{\phi:\tau}^{\text{win}}(\mu)|^2 . \quad (\text{B.13})$$

In equation B.13 τ is a counter for the segments and the scaling factor Υ is set so that the area under the function is equal to the energy in the time-domain signal:

$$\Upsilon = \frac{1}{N \frac{\delta\omega}{2\pi} \sum_{n=0}^{N-1} w(n)^2} \quad (\text{B.14})$$

Refer to Brandt [10] for a derivation of Υ (which is S_P in Brandt's notation).

Appendix C

Assumptions and details

C.1 Choice of FRF estimator

Figure C.1(a) shows a Bode magnitude plot of a driving-point mechanical-impedance measurement on a rigid mass, as described in section 5.3. Blue dots are used where the value of the coherence function is greater than 0.98. Crosses and plusses are used when the coherence drops below 0.98. The red crosses show the H1 estimator, while the green plus signs show the H2 estimator. The true value of the FRF lies between these two. As this is a peak, the H2 estimator gives a better approximation. This is evident from the Nyquist plot of the same data in figure C.1(b). In section 5.2.2 we show that FRF datapoints around a resonance form a circle on a Nyquist plot. As the antiresonances of driving-points FRFs correspond to the resonances of the system obtained when the driving point is grounded [87], the inverse-FRF datapoints around antiresonances should also form a circle on a Nyquist plot. In figure C.1(b) it can be seen that the datapoints from the H2 estimator of the mechanical impedance response function (Z_m) form a circle that is consistent with the higher-coherence values further from the antiresonance, while those from the H1 estimator do not.

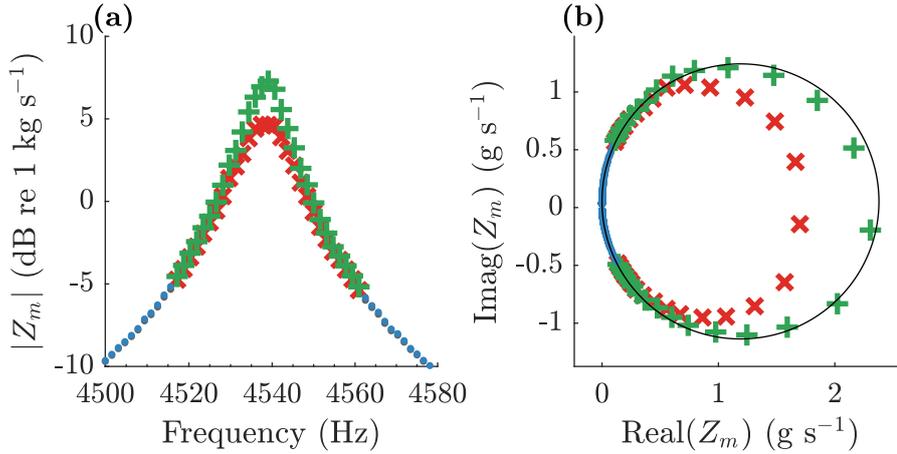


Fig. C.1 (a) Bode plot of the H1 and H2 estimators for a measured driving point impedance (150 g rigid mass on an impedance head, see section 5.3). (b) Corresponding Nyquist plot. For both plots, data with a coherence above 0.98 are shown by blue dots (\bullet). For coherence values below 0.98 the H1 estimator is shown using red crosses \times while the H2 estimator is shown using green pluses ($+$).

C.2 Air leaks

An air-leak is an acoustic-domain artefact. To simplify the analysis of its effect on the coupled chest–stethoscope system, let us consider an acoustic-domain impedance analogy circuit in which a Thévenin equivalent source is used to represent all terms except the cavity and the leak, as shown in figure C.2.

A leak between the cavity and ambient air gives an acoustic impedance (z_{leak}^a) between p'_c and ground (i.e., in parallel with the air-cavity impedance z_{cavity}^a). For a narrow leak the acoustic impedance consists of a resistance term (R_{leak}^a) and a mass term (m_{leak}^a) [6], such that:

$$z_{\text{leak}}^a = R_{\text{leak}}^a + i\omega m_{\text{leak}}^a. \quad (\text{C.1})$$

The values for R_{leak}^a and m_{leak}^a depend on the size and geometry of the leak. Analytical approximations for R_{leak}^a and m_{leak}^a exist for e.g., a narrow circular pipe and for a rectangular slit (see e.g., Beranek, 2012 [6]). However, a typical real leak will not conform to either of these geometries.

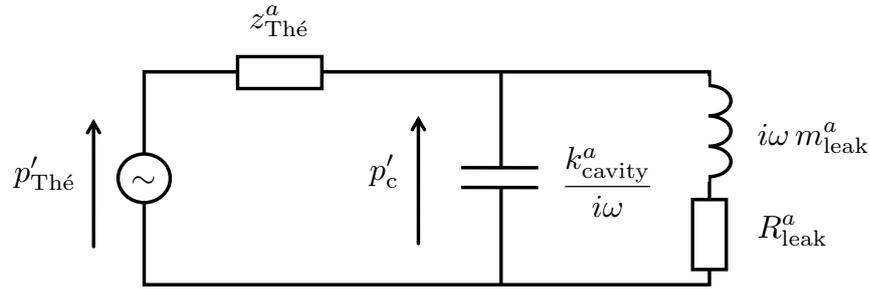


Fig. C.2 Acoustic domain impedance analogy circuit for an air-coupled microphone with a leak, with all terms except the cavity compliance and the leak replaced by a Thévenin equivalent source. The impedance of the air cavity is $z^a_{\text{cavity}} = k^a/i\omega = \rho_0 c_0^2/i\omega V$, the impedance of the leak is $z^a_{\text{leak}} = R^a_{\text{leak}} + i\omega m^a_{\text{leak}}$. The (acoustic) Thévenin impedance ($z^a_{\text{Thé}}$) is the impedance ‘seen’ by the air-cavity ‘looking’ towards the chest, with all source terms ‘off’. This term will include the effect of the sensor mass and the attachment impedance, and could be evaluated for the phantom by modifying equation (8.16) to use a transmission matrix between the measurements at the phantom’s base and the air-cavity input (as described in equation (10.10)), i.e. $z^a_{\text{Thé}} = [D/C]_{\hat{\mathbf{T}}_{b,\text{cav in}}}$.

A system with a leak will yield a Thévenin response of

$$\frac{P_c}{P_{\text{Thé}}} = \frac{z^a_{\text{cavity}} || z^a_{\text{leak}}}{z^a_{\text{Thé}} + z^a_{\text{cavity}} || z^a_{\text{leak}}}, \quad (\text{C.2})$$

where P_c is the cavity pressure and $P_{\text{Thé}}$ is a Thévenin equivalent source term. For impedances in parallel (denoted by the symbol $||$) the lowest impedance dominates. At low frequencies $z^a_{\text{cavity}} = k^a/i\omega$ is large, while z^a_{leak} is small, so the leak will have an effect on the low frequency response. At high frequencies z^a_{leak} becomes larger (tending towards an open circuit), and its effect on the cavity pressure becomes negligible. The smaller the leak area, the larger R^a_{leak} and m^a_{leak} , and thus the lower the frequency above which the effect is negligible.

The analysis here makes the assumption that there are no acoustic pressure perturbations in the ambient air. If this is untrue (i.e., in the presence of ambient noise) then the leak can also act as a source of contaminating noise for the ACM. The same considerations apply, meaning that this noise will be predominantly at low frequencies.

C.3 Assumption: axial microphone vibration can be neglected

Microphones are designed to operate with a stationary base and a diaphragm that is exposed to a fluctuating pressure level. If the base of the microphone is attached to a structure that is itself moving then the axial vibration of the microphone can lead to a false pressure reading.

A low-order model for a condenser microphone as a simple two-degree of freedom system is shown in figure C.3. A (low) diaphragm mass (m_d) is connected by a (high) stiffness ($k_{b,d}$) to a (moderate) base mass (m_b). The system is governed by the matrix equation (C.3):

$$\begin{bmatrix} k_{b,d} - \omega^2 m_d & -k_{b,d} \\ -k_{b,d} & k_{b,d} - \omega^2 m_b \end{bmatrix} \begin{bmatrix} x_d \\ x_b \end{bmatrix} = \begin{bmatrix} f_d \\ f_b \end{bmatrix}. \quad (\text{C.3})$$

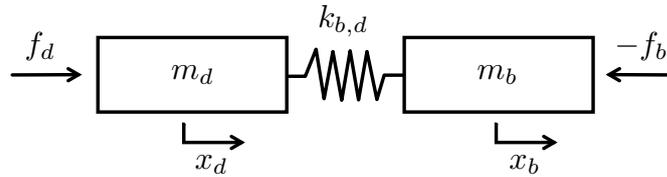


Fig. C.3 Simple lumped element model for a microphone, d denotes the diaphragm and b denotes the base.

A microphone is designed to accurately measure pressure when the base does not move. Pressure is measured (indirectly) via the capacitance between the diaphragm and the base. The measured signal is thus proportional to the displacement difference ($x_d - x_b$) between the diaphragm and the base. For the $x_b = 0$ case the top line of equation (C.3) yields:

$$f_d = (k_{b,d} - \omega^2 m_d)x_d. \quad (\text{C.4})$$

The true pressure is proportional to the force (f_d) exerted on the diaphragm. $\omega^2 m_d \ll k_{b,d}$ in the operating range of the microphone. Thus we use

$$p_{\text{true}} = \hat{C} f_d, \quad (\text{C.5})$$

$$p_{\text{meas}} = \hat{C} k_{b,d} (x_d - x_b), \quad (\text{C.6})$$

where \hat{C} is a conversion factor. Now consider the case where the microphone is not grounded and x_b is non-zero. The top line of equation (C.3) now gives

$$f_d = (k_{b,d} - \omega^2 m_d)x_d - k_{b,d}x_b, \quad (\text{C.7})$$

which can be expressed in terms of the diaphragm displacement as

$$x_d = \frac{f_d + k_{b,d}x_b}{(k_{b,d} - \omega^2 m_d)}. \quad (\text{C.8})$$

Equation (C.8) can be used to eliminate the unknown diaphragm displacement from equation (C.6), giving

$$p_{\text{meas}} = \hat{C}k_{b,d} \left(\frac{f_d + k_{b,d}x_b}{k_{b,d} - \omega^2 m_d} - x_b \right) = \hat{C}k_{b,d} \left(\frac{f_d + \omega^2 m_d x_b}{k_{b,d} - \omega^2 m_d} \right). \quad (\text{C.9})$$

Once again using the fact that $\omega^2 m_d \ll k_{b,d}$, and substituting for the true pressure from equation (C.5) we obtain:

$$\begin{aligned} p_{\text{meas}} &= p_{\text{true}} + \hat{C}\omega^2 m_d x_b \\ &= p_{\text{true}} - \hat{C}m_d \ddot{x}_b \\ &= p_{\text{true}} - \hat{C}' \ddot{x}_b \end{aligned} \quad (\text{C.10})$$

The datasheet for the G.R.A.S. 46BL 1/4 inch class of microphones states that the response to $\ddot{x}_{b\text{rms}} = 1 \text{ m/s}^2$ of root-mean-squared (rms) axial vibration is 55 dB (re. $p_{\text{ref}} = 20 \mu\text{Pa}$) of effective pressure. Assuming that this measurement is conducted in silence ($p_{\text{true}} = 0$) we obtain

$$55 = 20 \log_{10}(\hat{C}' \times 1)/p_{\text{ref}}, \quad (\text{C.11})$$

which yields

$$\hat{C}' = p_{\text{ref}} 10^{55/20} = 0.0112 \text{ Pa} / \text{m s}^{-2}. \quad (\text{C.12})$$

If there is both pressure at the diaphragm and axial vibration of the base then the measured pressure becomes:

$$p_m(t) = p_t(t) - \hat{C}' \ddot{x}_b(t) , \quad (\text{C.13})$$

$$P_m(\omega) = P_t(\omega) - \hat{C}' \ddot{X}_b(\omega) , \quad (\text{C.14})$$

$$P_m P_m^* = P_t P_t^* + \hat{C}'^2 \ddot{X}_b \ddot{X}_b^* - \hat{C}' (P_t \ddot{X}_b^* + \ddot{X}_b P_t^*) . \quad (\text{C.15})$$

In our case the pressure in the cavity and the acceleration of the housing are both in response to the same input force, so they will have a degree of coherence. This means that $P_m \ddot{X}_b^*$ and $\ddot{X}_b P_m^*$ are non zero. In order to assess the significance of the error terms we measure the microphone pressure and the housing acceleration for the setup in figure 9.2. The power spectral densities of the measured pressure and the effective pressure due to axial microphone are shown in figure C.4. The axial-vibration term is negligible compared to the measured pressure, and the estimate for the PSD of the true pressure is approximately equal to the PSD of the measured pressure.

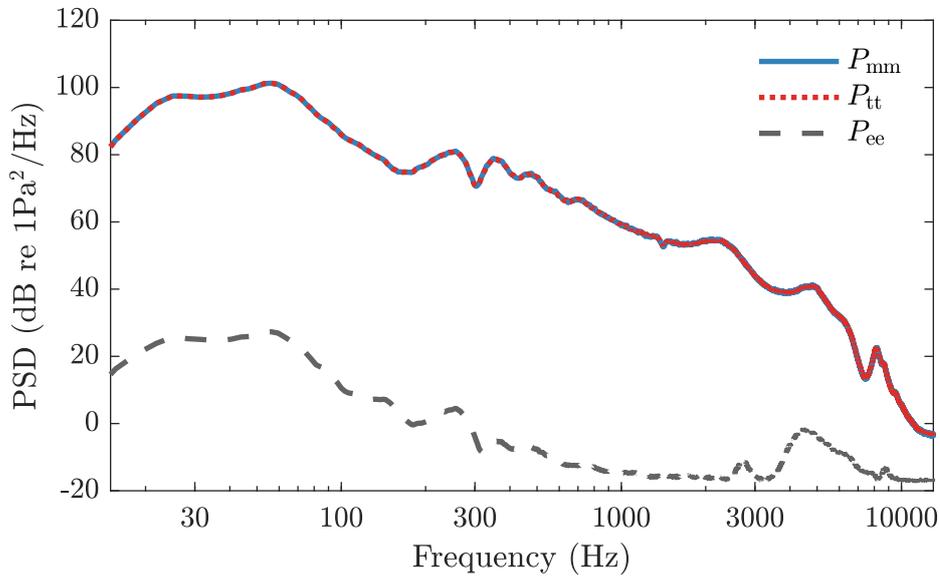


Fig. C.4 Effect of axial microphone vibration. P_{mm} is the power spectral density of the measured pressure, P_{ee} is the power spectral density of the effective pressure that would be seen due to axial microphone vibration if there was no actual pressure. P_{tt} is an estimation of the true pressure (based on a time domain combination of the signals to account for correlation).

C.4 [66] - ‘pressure response model’ for air-coupled microphones

Joyashiki and Wada [66] define a ‘pressure response model’ for air coupled sensors as:

$$\frac{z_c}{z_c + (i\omega m_s + k_a/i\omega)} \times \omega / Z_{\text{soft tissue}}^c \quad (\text{C.16})$$

The paper does not make it explicit what this expression is equal to, and the physical interpretation is therefore difficult to discern. The dimensions of this ‘pressure response’ are equivalent to acceleration over pressure. This makes them inconsistent with the dimensions of the ‘pressure response’ defined in the same paper for a different sensor type, (so-called body-coupled-sensors), for which this response is dimensionless.

The form of the impedance ratio found in this expression ($z_{\text{source}}/(z_{\text{source}} + z_{\text{load}})$) suggests that it is derived from a Norton response (see section 3.5.7). In particular, this expression could be interpreted as yielding the response of the surface velocity in the presence of a sensor of net impedance ($i\omega m_s + k_a/i\omega$), to a Norton reference velocity. Division by a specific characteristic impedance ($Z_{\text{soft tissue}}^c$) has the effect of converting the Norton reference velocity into a (loosely defined) reference pressure. Multiplication by $i\omega$ would convert the surface velocity into an acceleration. The fact that the i is omitted from the expression given suggests that the authors may only be interested in the magnitude of the response.

It is possible that there is a simple typographical error here, as multiplication by an appropriate impedance (divided by ω) would yield a pressure measurement, and thus a non-dimensional response function. However, this would not address the issue that summing $i\omega m_s$ and $k_a/i\omega$ to yield the net sensor impedance implies a topology that does not match the case description provided in the paper, as discussed in the main body of the text (section 4.3.2).

Appendix D

Simulation details

D.1 ABAQUS CAE Simulations

When discussing the physics of the interaction between a chest and a stethoscope, a visual representation of the deformation and motion of the chest is useful. Since it is impractical to capture the deformation of the chest tissue under the skin-surface experimentally, we present finite element simulations to illustrate our arguments. The role of the simulations is to illustrate theoretical arguments rather than to provide quantitative results.

All of the finite element method (FEM) simulations presented in this thesis are performed using the commercial finite element package ABAQUS CAE 6.14. The human chest is a complex system with many anatomical structures and a large degree of inter- and intra-person variability. In order to simplify the analysis and focus on the fundamental mechanisms of the coupling we consider only the case of a stethoscope coupled to a cylinder of elastic material. The geometry and material properties for the simulation are picked to match the laboratory model of the chest ('phantom') that we present in chapter 6.

The phantom is constructed of a silicone elastomer (Ecoflex™ 00-10, Smooth-On Inc). A wide range of material properties for Ecoflex™ have been reported in the literature, which may be due to the differences in air-bubble entrapment and external conditions (see section 6.2.1).

Vaicekauskaite et al. [151] suggest a Young's modulus of 0.05 MPA while Mansy et al. [96] report a Young's modulus of 0.11 MPa [96]. For the simulations

presented here, an intermediate value of 0.08 MPa is used. Ecoflex™ is almost incompressible [14], however the entrapment of air-bubble increases the compressibility. We use a Poisson's ratio of 0.495. The density of Ecoflex™ is similar to that of water. Mansy et al. report a density of 1030 kg m⁻³ while Cafarelli et al. [14] report a value of 1063 kg m⁻³. We use a value of 1050 kg m⁻³ in these simulations.

In order to capture the full modal behaviour of the system a 3D model is needed. However, when considering only the static case (or frequencies below the first non-axisymmetric mode) we can make use of the symmetry of the system to reduce the computational cost by constructing an axisymmetric model.

D.2 Figure 3.10

The lumped-element mechanical system shown in figure 3.9 is governed by the matrix equation:

$$\begin{bmatrix} m_1 & 0 \\ 0 & m_2 \end{bmatrix} \begin{bmatrix} \ddot{x}_1 \\ \ddot{x}_2 \end{bmatrix} + \begin{bmatrix} \lambda_{1,2} & -\lambda_{1,2} \\ -\lambda_{1,2} & \lambda_{1,2} + \lambda_{2,0} \end{bmatrix} \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} + \begin{bmatrix} k_{1,2} & -k_{1,2} \\ -k_{1,2} & k_{1,2} + k_{2,0} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} f_1 \\ 0 \end{bmatrix} \quad (\text{D.1})$$

This is solved for the driving point dynamic stiffness at x_1 using MATLAB R2020b for the following parameter values: $m_1 = 0.5$ g, $m_2 = 200m_1$, $\lambda_{1,2} = 0.5$ kg s⁻¹, $\lambda_{2,0} = 10\lambda_{1,2}$, $k_{1,2} = 50$ kN m⁻¹ and $k_{2,0} = 0.2k_{1,2}$.

D.3 Figure 4.2

The coupled system formed when an accelerometer is added to a SDOF oscillator is shown in figure D.1 and is governed by the matrix equation

$$\left(-\omega^2 \begin{bmatrix} m_1 & 0 & 0 \\ 0 & m_b & 0 \\ 0 & 0 & m_a \end{bmatrix} + \begin{bmatrix} k_{0,1} + k_{1,b} & -k_{1,b} & 0 \\ -k_{1,b} & k_{1,b} + k_{b,a} + k_{b,0} & -k_{b,a} \\ 0 & -k_{b,a} & k_{b,a} \end{bmatrix} \right) \begin{bmatrix} X_1 \\ X_b \\ X_a \end{bmatrix} = \begin{bmatrix} F \\ 0 \\ 0 \end{bmatrix}. \quad (\text{D.2})$$

Equation D.2 is stated for the undamped case for compactness, but can trivially be extended to include damping terms.

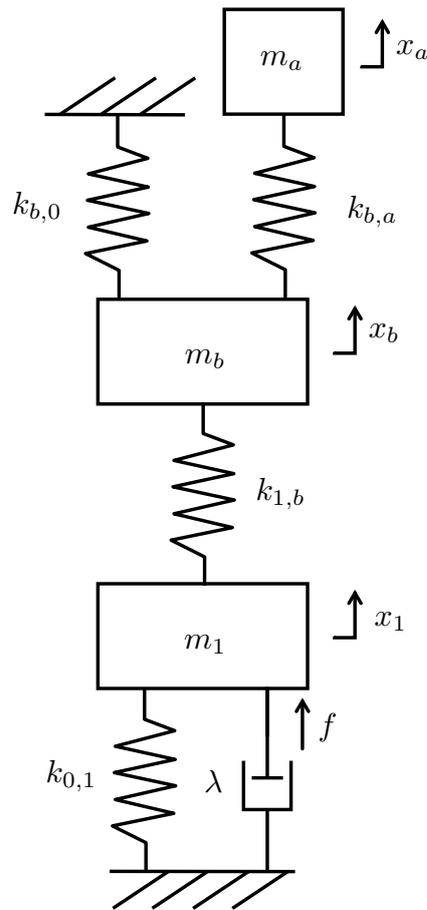


Fig. D.1 Single degree of freedom damped oscillator of mass m_1 and stiffness $k_{0,1}$, instrumented with an accelerometer, of base mass m_b and seismic mass m_a . The mounting stiffness between the test system and the accelerometer is $k_{1,b}$. The accelerometer's cable is modelled as a small stiffness between the accelerometer base and ground. The stiffness of the piezoelectric element between the seismic mass and the base mass is $k_{b,a}$. Dashpots in parallel with all springs have been omitted from the schematic for simplicity, with the exception of $\lambda_{0,1} = \lambda$

The lumped elements in the accelerometer model are defined in the legend of figure D.1. The mounting stiffness ($k_{1,b}$) and piezoelectric-element stiffness ($k_{b,a}$) are high, while the wire/load stiffness ($k_{b,0}$) and damping ($\lambda_{b,0}$ - in parallel with $k_{b,0}$) are low. This means that the dominant effect of adding the accelerometer is the added mass ($m_b + m_a$). Because of this, many researchers (e.g. Vermarien and van Vollenhoven [153]) implicitly take the mounting stiffness to be infinite and the cable stiffness to be zero in their accelerometer models.

The following parameter values are used in this simulation: $m_1 = 1$ kg, $m_b = 0.2$ kg, $m_a = 0.22$ kg, $\lambda_{0,1} = 10$ Ns/m, $\lambda_{1,b} = 10$ Ns/m, $\lambda_{b,0} = 1$ Ns/m, $\lambda_{b,a} = 0.01$ Ns/m, $k_{0,1} = 10$ kN/m, $k_{1,b} = 100$ kN/m, $k_{b,0} = 0.1$ kN/m, $k_{b,a} = 10,000$ kN/m.

D.4 Figure 9.11

Simulation of $p/f_{\text{Thé}}$ using the circuit in figure 9.10. k_a is computed using equation 3.47 with $\rho_0 = 1.225$ kg m⁻³, $c_0 = 343$ m s⁻¹, $\hat{A} = \hat{A}_c = 200$ mm² and $V = 2720$ mm³. $k_{c,s}$ is replaced by $(k_{c,s} + i\omega\lambda_{c,s})$ where $k_{c,s} = 16$ kN m⁻¹ and $\lambda_{c,s} = 18$ Ns m⁻¹. $m_s = 67$ g, and $z_{s,0} = \lambda_{s,0} = 0.5$ Ns m⁻¹. The Thévenin impedance of the chest ($z_{\text{Thé}}$) is given by a SDOF model with $m_{\text{eff}} = 2.7$ g, $\lambda_{\text{eff}} = 5.5$ Ns m⁻¹ and $k_{\text{eff}} = 8.7$ kN m⁻¹.

D.5 Figure 9.14

Simulation of $p/f_{\text{Thé}}$ using the circuit in figure 9.10, with an additional capacitor of impedance $(k_{d:c,s} + i\omega\lambda_{d:c,s})/i\omega$ as seen in figure 4.28. $k_{d:c,s} = 65$ kN m⁻¹ and $\lambda_{d:c,s} = 65$ Ns m⁻¹. k_a is computed using equation 3.47, with $\rho_0 = 1.225$ kg m⁻³, $c_0 = 343$ m s⁻¹, $\hat{A} = \hat{A}_c = 200$ mm² and $V = 990$ mm³. $k_{c,s}$ is replaced by $(k_{c,s} + i\omega\lambda_{c,s})$ where $k_{c,s} = 35$ kN m⁻¹ and $\lambda_{c,s} = 35$ Ns m⁻¹. $m_s = 222$ g and $z_{s,0} = \lambda_{s,0} = 2$ Ns m⁻¹. The Thévenin impedance of the chest ($z_{\text{Thé}}$) is given by a SDOF model with $m_{\text{eff}} = 4.9$ g, $\lambda_{\text{eff}} = 7.5$ Ns m⁻¹ and $k_{\text{eff}} = 13$ kN m⁻¹.

D.6 Transmission matrices in equation (10.1)

$$\mathbf{T}_{s,0} = \begin{bmatrix} 1 & 0 \\ 1/(i\omega m_s + z_{s,0}) & 1 \end{bmatrix} \quad (\text{D.3})$$

$$\mathbf{T}_{k_{c,s}} = \begin{bmatrix} 1 & k_{c,s}/i\omega \\ 0 & 1 \end{bmatrix} \quad (\text{D.4})$$

$$\mathbf{T}_{\text{trans}} = \begin{bmatrix} \hat{A}_c & 0 \\ 0 & 1/\hat{A}_c \end{bmatrix} \quad (\text{D.5})$$

$$\mathbf{T}_{k_a^a} = \begin{bmatrix} 1 & \rho c_0^2/i\omega V \\ 0 & 1 \end{bmatrix} \quad (\text{D.6})$$

D.7 Simple duct attenuation model

This appendix presents a simple lossy transmission line model which is used for the tubing simulations in chapter 10. Following the method presented in e.g. [149], the per-unit-length acoustic series impedance ($Z_{\text{series}}^{\prime a} = R^{\prime a} + i\omega M^{\prime a}$) and shunt admittance ($Y_{\text{shunt}}^{\prime a} = G^{\prime a} + i\omega C^{\prime a}$) of a rigid duct can be defined as

$$Z_{\text{series}}^{\prime a} = \frac{i\omega\rho_0}{\hat{A}(1 - \Psi(\delta_v))}, \quad (\text{D.7})$$

and

$$Y_{\text{shunt}}^{\prime a} = \frac{i\omega\rho_0}{\hat{A}(1 - \Psi(\delta_t))}. \quad (\text{D.8})$$

Here the superscript $^{\prime a}$ denotes a per-unit-length acoustic term and the function $\Psi(\delta)$ is defined as

$$\Psi(\delta) = \frac{2 J_1\left(\frac{(i-1)a}{\delta}\right)}{\frac{(i-1)a}{\delta} J_0\left(\frac{(i-1)a}{\delta}\right)}, \quad (\text{D.9})$$

where J denotes a Bessel function and δ_v and δ_t represent the viscous and thermal boundary layer thicknesses, which are given by

$$\delta_v = \sqrt{\frac{2\mu}{\omega\rho_0}}, \quad (\text{D.10})$$

and

$$\delta_t = \sqrt{\frac{2\kappa}{\omega\rho_0 C_p}}. \quad (\text{D.11})$$

The properties of the air in the duct are captured by its viscosity (μ), density (ρ_0), thermal conductivity (κ) and specific heat capacity at constant pressure (C_p).

For a lossy duct, the complex wavenumber then becomes

$$i\hat{k} = \hat{\alpha} + i\hat{k}_r = \sqrt{Z_{\text{series}}^{1a} Y_{\text{shunt}}^{1a}}, \quad (\text{D.12})$$

while the characteristic acoustic impedance becomes

$$Z_{0:\hat{A}}^a = \sqrt{Z_{\text{series}}^{1a} / Y_{\text{shunt}}^{1a}}. \quad (\text{D.13})$$

These complex terms are used directly in the transmission line matrices defined elsewhere in this thesis. A more detailed formulation can be found in Beranek (2012) [6], while Dokumancı (2021) [26] outlines more advanced cases such as ducts with flexible, finite-impedance walls (see also [25]).