

Acceleration sensing at the nano-g level

Development and characterisation of low-noise
microseismometers for next generation
gravitational wave detectors

Boris Boom

Doctorate committee:

Chair:

prof.dr. W.M.G. Ubachs

Vrije Universiteit Amsterdam

Promotor:

prof.dr. J.F.J. van den Brand

Universiteit Maastricht

Copromotor:

dr. A. Bertolini

Nikhef

Members:

dr. N.A. van Bakel

Nikhef

prof.dr.ir. C. Collette

Université de Liège

prof.dr. G. Hammond

University of Glasgow

prof.dr.ir. M. Kraft

Katholieke Universiteit Leuven

prof.dr. H.G. Raven

Vrije Universiteit Amsterdam

dr. H.L. Snoek

Universiteit van Amsterdam

dr.ir. R.J. Wiegerink

Universiteit Twente

Paranymphs

ing. S.P. Boom

D.J.S. van Engelen

Cover design by Boris Boom

Printed by Gildeprint - Enschede

ISBN: 978-94-6419-040-3



This work has been carried out at the National Institute for Subatomic Physics (Nikhef) and is part of the research programme *First direct detection of gravitational waves with Advanced Virgo* with project number 162, which is funded by the Dutch Research Council (NWO).



VRIJE UNIVERSITEIT

Acceleration sensing at the nano-g level

Development and characterisation of low-noise
microseismometers for next generation
gravitational wave detectors

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad Doctor
aan de Vrije Universiteit Amsterdam,
op gezag van de rector magnificus
prof.dr. V. Subramaniam,
in het openbaar te verdedigen
ten overstaan van de promotiecommissie
van de Faculteit der Bètawetenschappen
op donderdag 19 november 2020 om 13.45 uur
in de aula van de universiteit,
De Boelelaan 1105

door

Boris Anton Boom

geboren te 's-Hertogenbosch

promotor: prof.dr. J.F.J. van den Brand
copromotor: dr. A. Bertolini

to Linda



Preface

On September 14th 2015 a passing gravitational wave shook the world. Quite literally actually, for during a fraction of a second our globe was stretched and squeezed repeatedly by as much as a few femtometres, or, if you want, about the size of a single atomic nucleus. You would be excused for not having felt this wave, tiny as it is, but the tremor it created was certainly felt in the scientific community, when during a press conference on February 11th 2016 the announcement came that the two LIGO detectors had picked up this passing wave loud and clear. For the first time ever we had seen a direct proof for the existence of gravitational waves, and moreover, the observed waveform looked exactly like Einstein had predicted a century earlier.

I distinctly remember the moment I first heard the news myself. I had just started my doctoral studies two weeks before on September 1st 2015 and was visiting the Virgo gravitational wave detector in Cascina for the first time to get acquainted there. In the evening, whilst eating a pizza with some colleagues at a local restaurant not far from the detector, my doctoral advisor showed us a spectrogram of what would turn out to be the first direct detection of a gravitational wave from a pair of merging black holes. It was all quite unreal at the time. I was completely new to the field, and landed right in the middle of the culmination of decades worth of research effort. What followed was a hectic period in which the entire collaboration was fanatically checking, double-checking and then checking again all the fine details related to either the state

of the detectors at the time of the detection, or the analysis of the recorded data. All this had to be done while keeping total secrecy up to the moment of the press conference half a year later. I think it is safe to say I have learnt a lot during that intense period.

The end result of my doctoral studies now lies before you. This dissertation is focused around the development and characterisation of a novel low-noise microseismometer that uses geometric anti-springs to reach a noise performance unprecedented for its small size of roughly one square centimetre. The excellent noise performance combined with the batch-processing nature of its manufacturing make that such a seismic sensor is of interest for future generation gravitational wave detectors. A lot of the challenges that these new detectors face, especially at the lower end of their sensitive frequency band, are related to seismic motion. Arrays consisting of a large number of seismometers are likely to be required in order for these detectors to reach their design sensitivity at low frequencies.

The dissertation is structured as follows: In the first chapter, the relevant concepts in the theory of general relativity relating to gravitational waves and their interaction with physical objects will briefly be introduced. To get an idea of the principles behind gravitational wave detection and the practical limits to the sensitivity of the detectors, the basic layout of the current generation of interferometric gravitational wave detectors will be discussed. Next, the importance of a good sensitivity to gravitational waves at low frequencies is illustrated, and one of the designs for a next generation detector, Einstein Telescope, that will extend the low frequency band by a decade in frequency is highlighted. One of the major challenges in the realisation of Einstein Telescope, so-called Newtonian noise, will be discussed in Chapter 2, along with the basic principle of operation of seismic sensors in general. This chapter also contains a brief comparison of the novel microseismometer presented in this work to the current state of the art.

The main driving force behind the excellent noise performance of the microseismometer presented here, the concept of geometric anti-springs, will be discussed in Chapter 3. All details relating to the implementation of the geometric anti-spring system in the microseismometer can be found here, as well as a dimensionless design methodology for the implementation of geometric anti-springs at any desired scale. The microseismometer's most important property, its noise performance, will be analysed in Chapter 4. That chapter starts with a discussion on performing noise measurements on sensitive acceleration sensors, as they require non-trivial methods to either physically or vir-

tually suppress all input stimuli. The most important individual noise sources are subsequently modelled, and all noise levels are verified by direct measurements. Finally, the elementary fabrication process used to manufacture the sensors presented in this dissertation will be described in Chapter 5, and some of the features in the sensor's design that are related to its manufacturing are highlighted. Different approaches to making the seismometers more mechanically robust are then reviewed, and the chapter ends with a discussion on local vacuum encapsulation, an addition to the fabrication sequence that will be necessary in the future.

During the course of my doctoral studies, I've had the privilege to witness a transition in the field of gravitational wave research. It has been only a few years since the first detection, and the LIGO and Virgo interferometers are now routinely observing gravitational wave signals. In fact, we have already entered the era of publishing them in catalogues instead of presenting them as individual detections. Gravitational wave detection has been firmly established as a tool to test the theory of gravity under the most extreme circumstances, and to understand violent processes in the distant Universe. I could not have wished for a better time to work in the field.

Boris Boom, Amsterdam, February 2020

Contents

1	Gravitational wave detectors	1
1.1	Gravitational waves	2
1.1.1	Aspects of general relativity	5
1.1.2	General relativity in weak fields	7
1.1.3	The dynamics of spacetime	9
1.1.4	Experimentalist view on gravitational waves	10
1.1.5	Gravitational waves as tidal forces	11
1.1.6	A detector on Earth	13
1.2	Interferometric gravitational wave detectors	14
1.2.1	Fabry-Pérot cavity response	14
1.2.2	Gravitational waves in a single Fabry-Pérot cavity	16
1.2.3	Michelson interferometer with Fabry-Pérot arms	18
1.2.4	Quantum noise limited interferometry	20
1.2.5	Classical noise sources	27
1.3	The importance of good sensitivity at low frequencies	36
1.3.1	Low-frequency sensitivity in Advanced LIGO	37
1.3.2	Large redshift observations	38
1.3.3	Source parameters	40
1.3.4	Multimessenger astronomy	41

1.4	Einstein Telescope	44
1.4.1	Baseline layout	44
1.4.2	Sensitivity target	46
1.4.3	Site selection	47
2	Inertial sensing and Newtonian noise	49
2.1	Inertial sensors	50
2.1.1	Mechanical sensing element	50
2.1.2	The force-balance principle	52
2.2	Mitigating Newtonian noise	55
2.2.1	Passive suppression	56
2.2.2	Active suppression	58
2.3	State of the art in seismic sensors	60
3	Miniaturised geometric anti-spring technology	63
3.1	Geometric anti-spring working principle	64
3.2	Novel MEMS seismometer overview	66
3.2.1	Spring compression mechanism	66
3.2.2	Electrothermal actuators	70
3.3	Suspension spring design	75
3.3.1	Bending equations for curved springs	76
3.3.2	Solving the boundary value problem	79
3.3.3	Choosing the spring parameters	81
3.3.4	Stiffness for different working points	85
3.3.5	Rigid body modes	87
3.4	Vertical sensing	93
3.4.1	Asymmetric anti-springs	94
3.4.2	Parasitic spring for compensating gravity	95
4	Noise in MEMS seismometers	101
4.1	Noise measurements	101
4.1.1	Vibration attenuator	102
4.1.2	Three-channel correlation	102

4.2	Mechanical noise sources	105
4.2.1	Brownian noise	106
4.2.2	Internal friction	108
4.3	Electronic noise sources	114
4.3.1	Capacitive front-end	114
4.3.2	Loop-architecture	118
4.3.3	Other electrical noise sources	123
4.4	Closed-loop system performance	124
4.4.1	Measurement setup at Heimansgroeve	124
4.4.2	Correlation results	126
4.5	Modelling the gaseous damping	129
4.5.1	Damping of a body in infinite rarefied gas	129
4.5.2	Squeezed-film damping	131
4.5.3	Agreement with quality factor measurements	133
4.5.4	Damping considerations of capacitor geometry	135
5	MEMS fabrication	137
5.1	Elementary fabrication sequence	137
5.1.1	Process overview	138
5.1.2	Mitigating overetching damage	141
5.1.3	Device singulation without dicing	145
5.2	Towards a more robust fabrication process	147
5.2.1	XeF ₂ substrate etch	149
5.2.2	Wafer bonding	152
5.3	Vacuum packaging	158
5.3.1	Single die packaging	158
5.3.2	Wafer-level vacuum packaging	159
6	Conclusion and outlook	163
6.1	Conclusion	163
6.2	Outlook	166

Appendices 169

A.1 Curved spring out-of-plane stiffness 169

A.2 MEMS nominal parameters 171

A.3 Fabrication details 173

A.3.1 SOI wafer properties 173

A.3.2 Full fabrication sequence 174


A.4 Acronyms 180

Bibliography 181

Summary 191

Samenvatting 195

Acknowledgements 199



Gravitational wave detectors

Even though gravity is by far the weakest of the four presently known fundamental interactions, it is the only one that we, human beings, experience directly on a daily basis. It is therefore not surprising that it has a long and rich history in science, and written accounts on understanding the effects of gravity go back as far as the fourth century BC when Aristotle wrote his *Φυσικὴ ἀκρόασις*. About 2000 years later, Newton formulated his *law of universal gravitation*, the most widely known theory of gravity to date. It was successful in describing the orbits of the celestial bodies known at that time, and even though it is now superseded by a better theory, is still widely taught and used for its simplicity and intuitive qualities.

For just over a century now, the best theory of gravity that we have at our disposal is Einstein's theory of general relativity, or GR for short. In this theory, Einstein links gravity to his new concepts of space and time, and, in addition to fixing a few discrepancies between Newton's predictions and experimental observations, provides a plethora of predictions on new physical phenomena: the gravitational bending of light, the effect of gravity on the flow of time and the existence of black holes and gravitational waves. Over the past 100 years, scientists have subjected the theory to all possible tests, and GR passed every one of them with flying colours. Having said that, GR is known to be incompatible with quantum mechanics and to reconcile the two theories at length scales close to the Planck scale, a theory of quantum gravity is needed. String

theory is one of the approaches taken in this effort, but has of yet produced no experimentally testable predictions. The final word on gravity on small scales has yet to be spoken.

One of the most stringent tests of GR so far was only recently made possible with the first few direct detections of gravitational waves of coalescing binary systems. Whereas previous tests mainly concerned slowly moving systems with relatively weak gravitational fields, the observation of gravitational waves allowed to probe GR in the extreme gravitational fields created by black holes that were colliding at over half the speed of light.

More importantly, next to providing the most stringent test on the best theory of gravity currently available, the detection of gravitational waves provides valuable new information on our Universe. The waves carry with them information on their sources and their surroundings billions of light years away. Studying these gravitational waves is looking back in time and exploring the evolution of the Universe. The more sensitive our detectors become, the further we can look back and gain information about the early Universe.

This chapter will briefly introduce Einstein's theory of general relativity and the place of gravitational waves in it, then shift focus to the instruments used to detect these waves and finally summarise the benefits of improving their low-frequency sensitivity, and the challenges associated with doing so. It can not do justice to the decades of research and engineering that went into the currently operational gravitational wave detectors, but merely tries to convey the basic concepts used to make them work so well along with the future challenges that are relevant to this dissertation. It is important to keep in mind that by no means are these the only challenges that are to be dealt with in improving our sensitivity to gravitational waves. Gravitational wave detectors are formidable machines that push the state of the art in technology in multiple disciplines. All of them have to work if we want to succeed.

1.1 Gravitational waves

The first direct detection of gravitational waves from a pair of colliding black holes on September 14th 2015 [1] was a good start for celebrating 100 years of general relativity [2]. Not only did the waves exist, they also looked exactly as predicted by Einstein's field equations. To understand how Einstein eventually arrived at these equations, it is instructive to have a look at the two pillars of modern physics at that time: Newton's laws of motion and universal gravitation and Maxwell's laws of electrodynamics. The first was successful in describ-

ing the observed motion of the planets around the Sun¹, the latter described the interaction between magnets and charges and their associated fields, and governed the propagation of electromagnetic waves. All physical problems at hand could be described by either of the theories, or by their combination.

The combination of these two theories is exactly where Einstein saw an incompatibility that would ultimately lead him to formulate his theory of special relativity: Maxwell's laws say that an electromagnetic wave propagates with a velocity c , that is independent of the relative motion to the source. The way velocities transformed in Newton's theory meant that an observer following a beam of light at velocity v , should observe the light travelling at a velocity $c - v$, which is incompatible with Maxwell's laws. This was a first clue that one of the theories needed amending, and this is what Einstein set out to do.

A second important notion is the completely different way Maxwell's laws treat a magnet at rest or in motion with respect to the so-called ether. In the first case, there is only a static magnetic field, in the later case there is also an induced electric field. It turns out the effects of the two fields on any experiment trying to distinguish the moving magnet from the stationary one cancel exactly, but the physical treatment of the systems is different, and this bothered Einstein:

"The idea, however, that these were two, in principle different cases was unbearable for me. The difference between the two, I was convinced, could only be a difference in choice of viewpoint and not a real difference. [. . .] The phenomenon of magneto-electric induction compelled me to postulate the (special) principle of relativity." [4, Doc. 31, p. 20]

This principle of relativity is the first of two postulates that are the foundations of Einstein's theory of special relativity (or SR for short) that he published in 1905 [5]. His thoughts on the different treatment of stationary and non-stationary magnets were so important for the formulation of SR that he opens the paper with them. The resulting postulate is the following:

"[T]he phenomena of electrodynamics as well as of mechanics possess no properties corresponding to the idea of absolute rest. [. . .] the same laws of electrodynamics and optics will be valid for all frames of reference for which the equations of mechanics hold good." [5]

Einstein believed that the laws of physics should hold the same for any observer, or, more abstractly, any choice of reference frame. With SR, he for-

¹There was a slight deviation in the perihelion precession of Mercury of about 43" per century that puzzled astronomers for decades and led to the first doubts regarding the validity of Newton's gravitation laws [3].

mulated a theory that achieves this for a special set of reference frames, the inertial reference frames: in SR these are frames that have a relative motion of constant velocity. To formulate the theory, Einstein needed a second postulate, inspired by Maxwell's electrodynamics:

"Light is always propagated in empty space with a definite velocity c which is independent of the state of motion of the emitting body." [5]

The implications of this postulate reach far, because it has direct consequences for length and time intervals measured by moving observers.

Hermann Minkowski provided a useful geometric interpretation of SR in terms of a flat 4-dimensional manifold by adding time as a 4th dimension to the three spatial dimensions, resulting in what is now known as 'spacetime'. Einstein first rejected this mathematical interpretation as being "*superfluous learnedness*" [6, p. 152], but it was exactly this geometric interpretation that would provide him with the tools to later formulate his theory of general relativity. In Minkowski's geometric interpretation, the '*distance*' between two events in spacetime, the interval ds^2 , can be shown to be the same for all inertial observers. Observers can disagree about the simultaneity of events and their spacial separation, but the interval between the events is unaffected by their relative motion. The interval is given by

$$ds^2 = \eta_{\mu\nu} dx^\mu dx^\nu, \quad (1.1)$$

where the Einstein summation convention [2] is used for brevity. The four-vector x^μ lists the four coordinates (ct, x, y, z) and the components of the so-called Minkowski metric of flat spacetime in Cartesian coordinates are given by

$$\eta_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (1.2)$$

Because of the light postulate that Einstein used to build his theory, special relativity changed the views related to time and space completely. In the Galilean ideas about time and space (the ideas in which Newton's laws are embedded), the two entities are decoupled from each other and the concepts past and future mean everything before and after "*now*" respectively, independent of the event's location. Because of the finite speed c at which information can travel in SR, the concepts of universal time, universal past and universal future are lost.

Einstein realised that his new theory of special relativity required Newton's law of universal gravitation to be modified. Gravity in Newton's framework is represented as an instantaneous force between massive bodies separated in space. Because of the finite speed that information can travel at in Einstein's views, this instantaneous action at a distance did not fit the new framework. It took Einstein 10 years to reconcile gravity with his new ideas about space and time. The resulting theory of general relativity (or GR for short), was published in 1915 and 1916 in a series of papers [2]. Although the theory is called general relativity, it really is a theory of gravitation. It does not necessarily make motion more relative than it already was in special relativity, and the name is considered somewhat of a misnomer [7], but stuck with the theory.

1.1.1 Aspects of general relativity

In general relativity, Einstein adopts Minkowski's geometric interpretation and identifies spacetime as a 4-dimensional manifold with a metric. This metric, $g_{\mu\nu}$, contains all the information required to define distances between events on the manifold. In general the metric describes a globally curved spacetime, but when looking at a specific event, it can locally always be reduced to the flat metric of Eq. (1.2) by a proper choice of coordinates. This may sound somewhat counterintuitive, but we do it every day: even though we know that Earth is a sphere, locally we are used to thinking about it as a flat surface. This is an excellent approximation as long as we are only considering our immediate surroundings. The exact same thing holds in general relativity. As long as we consider small regions of spacetime only, we can model it to be locally flat, with some corrections of arbitrary precision when going further away from our origin.

Having chosen this way of representing spacetime, there are still two ingredients that are missing for obtaining a complete theory: firstly, an analogue of Newton's second law is needed to define how physical objects move in such a curved metric, and secondly, the source of the curvature should be specified in terms of the objects living in the manifold. The first ingredient is provided by the equivalence principle that gives information on the motion of particles in the curved manifold: free-falling particles move on time-like geodesics in spacetime. A geodesic is the generalisation of a straight line to curved spacetime, the shortest distance between two events. The equations of motion of particles in free fall can then be expressed in the geodesic equation

$$\frac{d^2 x^\alpha}{d\tau^2} + \Gamma^\alpha_{\mu\beta} \frac{dx^\mu}{d\tau} \frac{dx^\beta}{d\tau} = 0, \quad (1.3)$$

which reduces to Newton's second law of motion for nonrelativistic objects moving in a weak gravitational field. Here, the parameter τ denotes the proper time and $\Gamma_{\mu\beta}^{\alpha}$ is a so-called Christoffel symbol that can be written in terms of the metric $g_{\mu\nu}$ as

$$\Gamma_{\beta\mu}^{\gamma} = \frac{1}{2} g^{\alpha\gamma} (\partial_{\mu} g_{\alpha\beta} + \partial_{\beta} g_{\alpha\mu} - \partial_{\alpha} g_{\beta\mu}). \quad (1.4)$$

The second ingredient required to complete the theory is found in a relativistic analogue of Newton's law of universal gravitation $\nabla^2 \Phi = 4\pi\rho G$, which specifies that the gravitational potential Φ is generated by a mass distribution ρ . The relativistic version for ρ is the stress-energy tensor $T_{\mu\nu}$, that encodes the distribution of mass and all other forms of energy and momentum. The role of the gravitational potential in general relativity is played by the curvature of spacetime that is encoded in the metric $g_{\mu\nu}$. Curvature can be quantified in terms of the Riemann tensor

$$R^{\alpha}{}_{\beta\mu\nu} \equiv \partial_{\mu} \Gamma^{\alpha}{}_{\beta\nu} - \partial_{\nu} \Gamma^{\alpha}{}_{\beta\mu} + \Gamma^{\alpha}{}_{\sigma\mu} \Gamma^{\sigma}{}_{\beta\nu} - \Gamma^{\alpha}{}_{\sigma\nu} \Gamma^{\sigma}{}_{\beta\mu}. \quad (1.5)$$

Recalling the definition of the Christoffel symbols in Eq. (1.4) this result tells us that the curvature of spacetime, and with it also the gravitational field, is fully contained in the metric $g_{\mu\nu}$. The Riemann tensor has some useful contractions called the Ricci tensor and the Ricci scalar, defined respectively as

$$R_{\alpha\beta} \equiv R^{\mu}{}_{\alpha\mu\beta} \quad \text{and} \quad R \equiv R^{\mu}{}_{\mu} = g^{\mu\nu} R_{\mu\nu}, \quad (1.6)$$

that allow us to construct the Einstein tensor

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R. \quad (1.7)$$

Tracing back all the definitions, it is clear that the Einstein tensor is still fully determined by the metric $g_{\mu\nu}$ and it allows to write the field equations of general relativity in an elegantly compact form

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (1.8)$$

These equations provide the relativistic analogue to $\nabla^2 \Phi = 4\pi\rho G$ and together with the geodesic equation (1.3) it provides a full theory of gravity. The factor $8\pi G/c^4$ is found by requiring that GR matches observational data, *i.e.* that the Newtonian laws are recovered in the weak gravity limit.

The Einstein field equations look deceptively simple, but are in fact a compact way of writing 10 non-linear differential equations (not 16, since both the $G_{\mu\nu}$ and $T_{\mu\nu}$ tensors are symmetric). Four additional constraints come from the Bianchi identities, one for each value of μ ,

$$\nabla_\nu G^{\mu\nu} = 0, \quad (1.9)$$

where ∇_ν denotes the covariant derivative. These identities provide four relations between the ten different components of $G_{\mu\nu}$, indicating that only six of them are independent. The Einstein field equations therefore provide a total of six independent differential equations for the ten independent components of the metric tensor. The four remaining degrees of freedom reflect the liberty that general relativity provides for choosing a convenient coordinate system, determining what the ten metric components look like only after doing so.

1.1.2 General relativity in weak fields

The Einstein field equations shown above are non-linear, but can be linearised locally for weak gravitational fields by realising that since gravity is linked to the curvature encoded in the metric, regions with weak gravitational fields correspond to regions in spacetime where curvature is small. In that case, the metric can be viewed as the flat Minkowski metric with small corrections added:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad \text{with} \quad |h_{\mu\nu}| \ll 1. \quad (1.10)$$

Using this in Eq. (1.5) gives the Riemann tensor to first order in $h_{\mu\nu}$,

$$R_{\alpha\beta\mu\nu} = \frac{1}{2} (\partial_\beta \partial_\mu h_{\alpha\nu} + \partial_\alpha \partial_\nu h_{\beta\mu} - \partial_\beta \partial_\nu h_{\alpha\mu} - \partial_\alpha \partial_\mu h_{\beta\nu}). \quad (1.11)$$

To get a compact result, it is useful to define the trace reverse of $h_{\mu\nu}$,

$$\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} h \quad \text{with} \quad h \equiv \eta^{\mu\nu} h_{\mu\nu}. \quad (1.12)$$

Now by making use of the freedom that GR provides for imposing the Lorentz gauge $\partial_\nu \bar{h}^{\mu\nu} = 0$, we obtain the linearised Einstein field equations

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}, \quad (1.13)$$

where \square denotes the four-dimensional equivalent of the Laplace operator and is called the wave operator or *D'Alembertian*. When working in Cartesian co-

ordinates in flat spacetime this operator can simply be written in terms of the Laplace operator as

$$\square \equiv \partial_\alpha \partial^\alpha = \left(-\frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \nabla^2\right). \quad (1.14)$$

The Einstein field equations in Eq. (1.13) will be used to describe the behaviour of experiments on the surface of Earth, but it is not a priori obvious that the gravitational field there is weak enough for the approximation in Eq. (1.10) to hold. To find out, we can compare the result we obtained with Newtonian gravitation, for we know that in the weak field limit general relativity should reduce to Newton's theory.

Consider nonrelativistic processes in the Earth's static gravitational field and assume for the moment that there are no gravitational waves present; an excellent assumption based on our daily experiences. In that case the components in the stress-energy tensor $T_{\mu\nu}$ are ordered as $|T_{00}| \ll |T_{0j}| \ll |T_{ij}|$ [8], so the dominant contribution to the gravitational field in this frame will come from the component $T_{00} \approx \rho c^2 + O(\rho v^2)$. Moreover, since we are considering a static field, time derivatives vanish and the d'Alembertian operator in Eq. (1.14) simply reduces to the Laplace operator. The dominant field equation then reads

$$\nabla^2 \bar{h}_{00} = -\frac{16\pi G}{c^4} T_{00} = -\frac{16\pi G \rho}{c^2}. \quad (1.15)$$

Comparing this with the Newtonian gravity's field equation $\nabla^2 \Phi = 4\pi \rho G$, we see that these field equations are identical when $\bar{h}_{00} = -4\Phi/c^2$. Since this is the dominant component in $\bar{h}_{\mu\nu}$, Eq. (1.12) implies that $h \approx \bar{h}_{00}$ and thus that the metric corrections $h_{\mu\nu}$ are

$$h_{\mu\nu} \approx \frac{\bar{h}_{00}}{2} \delta_{\mu\nu} = -\frac{2\Phi}{c^2} \delta_{\mu\nu}, \quad (1.16)$$

where $\delta_{\mu\nu}$ denotes the Kronecker delta function. In other words, when choosing a metric $g_{\mu\nu} = \eta_{\mu\nu} - (2\Phi/c^2)\delta_{\mu\nu}$, GR locally reduces to Newton's universal gravitation. Since the gravitational potential near the Earth is known, the deviation from the Minkowski metric induced by the Earth's presence can be estimated:

$$h_{\mu\nu} = -\frac{2\Phi}{c^2} \delta_{\mu\nu} = -\frac{2GM}{c^2 R} \delta_{\mu\nu} \approx 10^{-9} \delta_{\mu\nu} \ll 1, \quad (1.17)$$

with M and R the mass and radius of the Earth respectively. This means the condition in Eq. (1.10) is met, and the weak-field approximation is an excellent one indeed for experiments performed on Earth.

1.1.3 The dynamics of spacetime

Where in the previous section the gravitational field was assumed to be static, we will now dive into the dynamics of general relativity. When looking at the linearised field equations given by Eq. (1.13) in a region of spacetime far away from any sources of curvature, we know that all components of the stress-energy tensor are zero, *i.e.* $T_{\mu\nu} = 0$. The resulting field equations are $\square \bar{h}_{\mu\nu} = 0$, or, more explicitly

$$\left(-\frac{1}{c^2} \frac{\partial^2}{\partial t^2} + \nabla^2\right) \bar{h}_{\mu\nu} = 0. \quad (1.18)$$

This result represents a wave equation and has wave-like solutions for $\bar{h}_{\mu\nu}$ that are of the form

$$\bar{h}_{\mu\nu} = A_{\mu\nu} \cos(k_\alpha x^\alpha) = A_{\mu\nu} \cos(\omega_{gw} t - k_i x^i), \quad \text{with } k^\alpha = \left(\frac{\omega_{gw}}{c}, k^i\right). \quad (1.19)$$

These solutions represent plane waves traveling in the direction defined by the spatial wave vector k^i . Because Eq. (1.19) is only a solution to the wave equation when $k^\alpha k_\alpha = 0$, these gravitational waves propagate at the speed of light and their frequency is given by $\omega_{gw} = c|k^i|$. The information on the wave's amplitude and polarisation is encoded in the amplitude $A_{\mu\nu}$.

The Lorentz gauge imposed on $\bar{h}_{\mu\nu}$ earlier to get to the compact form of Eq. (1.13) also poses a condition on $A_{\mu\nu}$. More specifically, it implies $A_{\mu\nu} k^\nu = 0$, or in other words, that $A_{\mu\nu}$ is orthogonal to the wave vector k^ν , such that the wave solution in Eq. (1.19) represents a transverse wave.

Einstein recognised these wavelike solutions already in 1916 and published a paper on them. When considering a general system of accelerating masses that would emit such gravitational waves, however, he concluded:

"[I]t is obvious that [the radiation of the system] has, in all imaginable cases, a practically vanishing value." [9]

We now know that in violent astrophysical events, gravitational radiation is emitted in copious amounts, but his statement still holds some truth. The gravitational waves detected by the Virgo and LIGO detectors up to now originate from the violent last moments in the lives of a pair of black holes or neutron stars, astrophysical objects whose existence was not known at the time. Moreover, by the time the emitted gravitational waves from these sources arrive on Earth, their amplitude is only of $O(10^{-21})$, a value so small that it was inconceivable to be measured with early 20th century technology.

1.1.4 Experimentalist view on gravitational waves

To understand how a passing gravitational wave can be detected, we first need to see what the physical influence is that a passing gravitational wave has on a detector in our laboratory. In a lab, we are used to measuring times with a local clock and distances with a rigid ruler as if we are in flat Euclidean space. Coordinates are defined by how many tick marks on a ruler you are away from your origin of choice. Choosing this view, even in the presence of gravitational waves, the metric is locally flat when considering a region of space significantly smaller than the typical scale on which the curvature changes. For the sake of simplicity, we set up our lab in a free-falling satellite. Such a free-falling frame has some nice properties that can be used to simplify the analysis. Firstly, the metric in such a frame is flat up to second order in x^i and the corrections only show up as

$$g_{\mu\nu} = \eta_{\mu\nu} + O\left(\frac{x^i x^j}{L_C^2}\right), \quad (1.20)$$

where L_C is the typical length scale on which the curvature changes, which in the case of gravitational waves is their reduced wavelength λ_{gw} . Because all first order corrections are zero, the Christoffel symbols as defined in Eq. (1.4) all vanish around the expansion point, *i.e.* the origin chosen inside our free-falling laboratory. Now for the assumption at the basis of linearised theory ($|h_{\mu\nu}| \ll 1$) to hold, Eq. (1.20) tells us we can only consider a lab that is smaller than the reduced wavelength of the gravitational wave we are considering. Everything that follows is only valid in this limit.

We now put two test particles in our lab, one at the origin and one a distance ξ^μ away from the origin. Denoting the geodesic of the origin with $x^\mu(\tau)$, both the curves $x^\mu(\tau)$ and $x^\mu(\tau) + \xi^\mu(\tau)$ satisfy the geodesic equation (1.3). Now for $|\xi^\mu| \ll L_C$, taking the difference of the geodesic equation at these two curves gives the equation of geodesic deviation,

$$\frac{d^2 \xi^\mu}{d\tau^2} + 2\Gamma^\mu_{\nu\rho}(x^\mu) \frac{dx^\nu}{d\tau} \frac{d\xi^\rho}{d\tau} + \xi^\sigma \partial_\sigma \Gamma^\mu_{\nu\rho}(x^\mu) \frac{dx^\nu}{d\tau} \frac{dx^\rho}{d\tau} = 0, \quad (1.21)$$

which tells us something about the movement of the particle at $x^\mu + \xi^\mu$ with respect to our origin. Now, because the Christoffel symbols vanish around x^μ , the second term in Eq. (1.21) vanishes. Moreover, if we assume the test particles to have non-relativistic velocities in our lab, we have $dx^i/d\tau \ll dx^0/d\tau$ and the

previous expression significantly simplifies to

$$\frac{d^2 \xi^i}{d\tau^2} + \xi^\sigma \partial_\sigma \Gamma_{00}^i (x^i) \left(\frac{dx^0}{d\tau} \right)^2 = 0. \quad (1.22)$$

From the form of the metric in Eq. (1.20) we can see that $\partial_0 \Gamma_{00}^i = 0$, so we have $\xi^\sigma \partial_\sigma \Gamma_{00}^i = \xi^j \partial_j \Gamma_{00}^i$. Moreover, because velocities are low, we have $dx^0/d\tau \approx c$ and $\tau \approx t$, with t being the local coordinate time indicated by stationary clocks in our lab. Finally, from the definition of the Riemann tensor in Eq. (1.5), we see that in this frame the last two terms vanish because the Christoffel symbols are all zero. Then $\partial_0 \Gamma_{00}^i = 0$ also means that

$$R_{0j0}^i = \partial_j \Gamma_{00}^i - \partial_0 \Gamma_{0j}^i = \partial_j \Gamma_{00}^i, \quad (1.23)$$

such that we can finally write Eq. (1.22) as

$$\ddot{\xi}^i = -c^2 \xi^j R_{0j0}^i. \quad (1.24)$$

1.1.5 Gravitational waves as tidal forces

Equation (1.24) tells us that in this particular frame the curvature corrections $h_{\mu\nu}$ contained in the Riemann tensor manifest themselves as forces that accelerate particles proportional to their separation from the origin ξ^i . In principle, the Riemann tensor includes corrections for all gravitational effects, but for now we will assume that the only curvature is caused by a passing gravitational wave as described in Eq. (1.19). To see what the resulting forces look like, all we have to do is calculate the components of the Riemann tensor R_{0j0}^i from $\bar{h}_{\mu\nu}$. Because in linearised GR the Riemann tensor is invariant under gauge transformations [8], we can choose whatever gauge we please to express $\bar{h}_{\mu\nu}$. The most compact form is obtained when using the so-called *Transverse Traceless* gauge, or TT gauge, corresponding to the gauge conditions

$$\bar{h}^{0\mu} = 0, \quad \bar{h}^i_i = 0 \quad \text{and} \quad \partial^j \bar{h}_{ij} = 0, \quad (1.25)$$

in which the wave amplitude $A_{\mu\nu}$ takes a conveniently simple form. Subject to these gauge conditions, Eq. (1.11) tells us that the Riemann component of Eq. (1.24) can be expressed in terms of $h_{\mu\nu}$ as²

$$R_{0j0}^i = \frac{1}{2} (\partial_0 \partial_j h_{i0} + \partial_i \partial_0 h_{0j} - \partial_0 \partial_0 h_{ij} - \partial_i \partial_j h_{00}) = -\frac{1}{2c^2} \ddot{h}_{ij}. \quad (1.26)$$

²Looking at the definition of $\bar{h}_{\mu\nu}$ in Eq. (1.12), Eq. (1.25) implies $\bar{h}_{\mu\nu} = h_{\mu\nu}$, such that when working in the TT gauge we can use $h_{\mu\nu}$ instead of $\bar{h}_{\mu\nu}$ to avoid cluttering the equations.

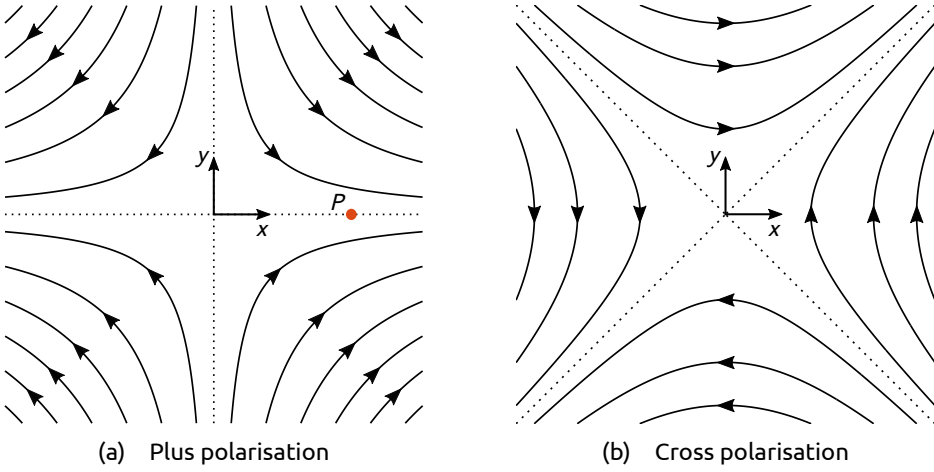


Figure 1.1 – Force line representation of the two different allowed gravitational wave polarisations for a wave travelling along the z -axis (into the page). The two polarisation states change into one another under a rotation of 45° .

When we align the z -axis in our lab along the wave's propagation vector k^i , we can rewrite Eq. (1.19) in the TT gauge as

$$h_{ij} = \begin{bmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{bmatrix} \cos\left(\omega_{gw} \left(t - \frac{z}{c}\right)\right), \quad (1.27)$$

where h_+ and h_\times indicate the amplitudes of the two independent polarisations.

Why these amplitudes are called '*h-plus*' and '*h-cross*' becomes clear when inspecting the shape of the tidal force fields they introduce when passing our laboratory. Combining Eq. (1.24) and (1.26) and interpreting the result in terms of Newtonian forces, we get

$$F^i = \frac{m}{2} \ddot{h}_{ij} \xi^j. \quad (1.28)$$

These tidal force patterns are separately plotted for each of the two polarisations in Fig. 1.1 for $t = z = 0$. Looking at the *plus-polarisation* in Fig. 1.1a, we see that particles along the x -axis get pushed away from the origin while particles along the y -axis get pushed towards the origin. The *cross-polarisation* in Fig. 1.1b shows exactly the same picture rotated by 45° .

Still looking at the plane $z = 0$ but letting the time evolve (*i.e.* letting the wave propagate through our lab), we would simply see the same force patterns modulated by $\cos(\omega_{gw} t)$. A test particle somewhere in this force field will ex-

perience a force oscillating with ω_{gw} , with an amplitude and direction that is determined by its position ξ^j . For example, a test particle at point P located along the x -axis in Fig. 1.1a will be repeatedly pushed away and pulled towards the origin with a frequency ω_{gw} . The amplitude of this oscillating force grows with the separation from the origin. This means it pays off to make a gravitational wave detector as large as possible, which is why the current interferometric detectors are multiple kilometres in size.

1.1.6 A detector on Earth

What we have learnt from the analysis of a free-falling laboratory so far can be briefly summarised as follows:

1. By choosing a Cartesian coordinates of space in our lab, we can make use of our usual Newtonian intuition in the description of our detector.
2. A passing gravitational wave shows up as a simple oscillating Newtonian force as long as the size of our detector is smaller than λ_{gw} .
3. The magnitude of this force grows linearly with the separation from the origin, so it makes sense to build large scale detectors.

Being allowed to use our Newtonian intuition, we can translate the satellite lab to a more practical lab on Earth by introducing gravity as a static force. As we have seen in Section 1.1.2, at the surface of the Earth the weak field approximation of general relativity holds and the description of gravitational waves as derived above is perfectly valid. To keep our test particles of choice from falling to the ground, we have to counteract gravity by introducing another force from some type of suspension mechanism. When a gravitational wave passes these suspended test particles, its tidal influence will just add to these static forces and shake them with respect to each other at a frequency ω_{gw} .

At this point the objective is clear; if we want to detect the effects of a gravitational wave, we need two things:

- The ability to accurately monitor the relative position of several suspended objects. This is what changes when a gravitational wave shakes them.
- In the frequency band of interest, all other forces that shake the suspended objects must be suppressed to a level below that of the tidal forces introduced by the gravitational wave as given in Eq. (1.28).

Because the amplitude of the tidal forces grows with the separation of the suspended objects, the latter is most easily achieved when making this separation large. We will now describe the techniques used in large-scale interferometric gravitational wave detectors to monitor the distance between objects.




Figure 1.2 – Schematic representation of a Fabry-Pérot cavity formed by two suspended mirrors. The cavity is illuminated from the left by a monochromatic laser beam of wavelength λ . The resulting electric fields around the input mirror are indicated. The total cavity length L is perturbed by a time dependent quantity $\delta L(t)$. Image adapted from [10].

1.2 Interferometric gravitational wave detectors

In the previous section we have determined that to detect gravitational waves, we should monitor the distance between suspended objects. The most accurate way to do this is by using measurements of the phase of light by means of interferometry. Since anything with mass will be affected by the gravitational waves, it is convenient to suspend mirrors and use them as test masses in a gravitational wave detector.

1.2.1 Fabry-Pérot cavity response

Consider the two mirrors in Fig. 1.2 that are suspended as simple pendula and together form a Fabry-Pérot cavity. The cavity is illuminated from the left by a laser of wavelength λ . Such a cavity forms an optical resonator where the field amplitude inside the cavity builds up when the cavity length L is an integer number of half wavelengths. In that case, the light that reflects off the input mirror after having completed a round trip ($r_1 E_{ii}$) interferes constructively with the light coming from the laser (E_{in}). When the cavity is on one of its resonances, the phase of the reflected light (E_r) is highly sensitive to the relative position of the two mirrors, which is exactly what we want.

To see how the phase of the reflected light changes with the mirror separation, we need to first see what happens to the field E_c inside the cavity. E_c can be written as a superposition of the input field and E_c itself, but one round trip earlier,

$$E_c(t) = t_1 E_{in}(t) + r_1 r_2 e^{2ik(L+\delta L(t))} E_c(t - 2T), \quad (1.29)$$

where $T = L/c$ is the time it takes for a photon to travel from one mirror to the

other.³ We consider the cavity locked on resonance while we move the end mirror by $\delta L(t)$ with $k\delta L(t) \ll 1$. In that case we have

$$e^{2ikL} \approx 1 \quad \text{and} \quad e^{2ik\delta L(t)} \approx 1 + 2ik\delta L(t). \quad (1.30)$$

Moreover, the output mirror typically is highly reflective, so $r_2 = 1$ (in Advanced Virgo, $r_2 \approx 0.99996$ [11]). Finally, for the purpose of this analysis we choose a constant amplitude input field $E_{in}(t) = E_0$. When we explicitly write the fields $E(t)$ as their average \bar{E} and a small time dependent perturbation $\delta E(t)$ separately, we get

$$\bar{E}_c + \delta E_c(t) = t_1 E_0 + r_1 (\bar{E}_c + \delta E_c(t - 2T) - 2ik\bar{E}_c\delta L(t)), \quad (1.31)$$

where the small cross term $\delta L(t)\delta E_c(t - 2T)$ has been dropped. In the absence of perturbations ($\delta L(t) = \delta E_c(t) = 0$), we can easily find the solution for the average static field inside the cavity on resonance,

$$\bar{E}_c = \frac{t_1}{1 - r_1} E_0, \quad (1.32)$$

and because this is a solution to Eq. (1.31), we can subtract it and solve for the small perturbations separately [10]. Transforming to the Laplace domain and solving for $\delta E_c(s)$ we get

$$\delta E_c(s) = i \frac{2kr_1\bar{E}_c\delta L(s)}{1 - r_1e^{-2Ts}} \approx i \frac{2kt_1\mathcal{F}^2E_0}{\pi^2} \frac{\delta L(s)}{1 + s/\omega_c}, \quad (1.33)$$

where in the last step the exponent was expanded as $e^{-2Ts} \approx 1 - 2Ts$, which is valid for $Ts \ll 1$.⁴ To get a compact result, we also defined the cavity finesse and the cavity pole as

$$\mathcal{F} = \frac{\pi\sqrt{r_1r_2}}{1 - r_1r_2} \quad \text{and} \quad \omega_c = \frac{\pi}{2\mathcal{F}T} = \frac{\pi c}{2\mathcal{F}L}. \quad (1.34)$$

We now have a solution for the field inside the cavity, E_c , but we are interested in the reflected field, so we use $E_r = r_1E_0 + t_1E_{cc}$ and $E_c = t_1E_0 - r_1E_{cc}$

³Actually, because the light bounced off the end mirror halfway its trip, we would have to consider the position $\delta L(t - T)$ instead of $\delta L(t)$ [10]. Since this will only introduce a small overall phase delay, I will omit it for the sake of clarity.

⁴Identifying s with the gravitational wave frequency, $Ts \ll 1$ translates to $\omega_{gw} \ll c/L$, or equivalently $L \ll \lambda_{gw}$, exactly the limit in which the tidal force description of Eq. (1.28) holds.

to get the cavity reflectivity around resonance,

$$r_c = \bar{r}_c + \delta r_c(s) = \frac{\bar{E}_r + \delta E_r(s)}{E_0} = -1 - i \frac{8\mathcal{F}}{\lambda} \frac{\delta L(s)}{1 + s/\omega_c}, \quad (1.35)$$

where we have assumed lossless mirrors ($r_1^2 + t_1^2 = 1$). Noting that $|r_c(s)| \approx 1$ for small $\delta L(s)$, this can be written in polar form as $r_c(s) = -e^{i\phi(s)}$, with

$$\frac{\phi(s)}{\delta L(s)} = \frac{8\mathcal{F}}{\lambda} \frac{1}{1 + s/\omega_c}. \quad (1.36)$$

This result tells us a few things. First of all, around one of its resonances the Fabry-Pérot cavity can effectively be modelled as a single mirror that fully reflects the incoming field amplitude, while imprinting a phase shift on it that depends on the separation of the mirrors that make up the cavity. The amount of phase change per unit of end mirror displacement imprinted on the reflected beam is given by Eq. (1.36). Secondly, it shows that above the cavity pole frequency ω_c , the sensitivity to mirror motion drops as $1/s$. This can be understood by realising that a photon on average spends a time $\tau_c = 1/2\omega_c$ travelling up and down the cavity, accumulating phase on each round trip. For frequencies above ω_c , the mirror displacement $\delta L(t)$ changes sign during the average storage time of the photons, causing them to lose already accumulated phase before exiting the cavity.

1.2.2 Gravitational waves in a single Fabry-Pérot cavity

Now assume we place the input mirror on the origin of the coordinate system as described in Section 1.1.5, and we position the end mirror along the x-axis at a position $\xi^i = (L, 0, 0)$. When a *plus-polarised* gravitational wave aligned as in Fig. 1.1a passes the cavity, its tidal forces will shake the end mirror along the x-axis, generating a phase shift in the laser beam that is reflected back from the cavity. In this situation, the only relevant component of h_{ij} as defined in Eq. (1.27) is $h_{11}(t) = h_+ \cos(\omega_{gw}t)$. Plugging this into Eq. (1.28) and taking the Laplace transform, we see that the passing gravitational wave induces a tidal force pointing along the x-axis that is given by $F^1(s) = m/2 \cdot s^2 h_{11}(s)L$. As it is suspended, the end mirror will respond to the tidal forces with the typical transfer function of a simple pendulum, which above the pendulum's natural frequency ω_0 ($s \gg \omega_0$) is given by $\delta L(s)/F^1(s) \simeq 1/ms^2$. Combining these two relations, we can get an expression for the relative length change of the cavity

in terms of the gravitational wave amplitude as

$$\frac{\delta L(s)}{L} = \frac{h_{11}(s)}{2}. \quad (1.37)$$

The result above nicely illustrates why h_{ij} is often called the gravitational wave *strain*, as it is a unitless quantity that indicates by how much the distance between two objects changes per unit length. This interpretation gives a useful means for comparing the detection noise floor (often simply called *sensitivity*) of gravitational wave detectors of different sizes by referring all the noise sources in the system to the detector input, expressing them in terms of this unitless strain amplitude.

By using the result obtained in Eq. (1.37) together with the Fabry-Pérot cavity transfer function found in Eq. (1.36), we get the cavity's phase response to a passing gravitational wave

$$\frac{\phi_x(s)}{h_{11}(s)} = \frac{4\mathcal{F}L}{\lambda} \frac{1}{1 + s/\omega_c}. \quad (1.38)$$

If the gravitational wave would be the only thing that induces a phase shift in the reflected light, we could use this single cavity to detect them. To see why this does not work in practice, recall the relation for the field inside the Fabry-Pérot in Eq. (1.29). What shows up in the exponent is the product of the cavity length with the laser wave number (kL). We chose to vary the length by $\delta L(t)$, but we might have just as well chosen to vary k , or equivalently, the laser frequency ν . Any laser will exhibit frequency fluctuations and in a single cavity they are indistinguishable from length variations. This will limit the smallest gravitational wave we can detect. Just by symmetry, a similar transfer function can be obtained for laser frequency fluctuations,

$$\frac{\phi_x(s)}{\delta\nu(s)} = \frac{8\mathcal{F}L}{c} \frac{1}{1 + s/\omega_c}. \quad (1.39)$$

Treating the phase shift generated by a gravitational wave in Eq. (1.38) as a signal and the phase shift from the laser frequency fluctuations as noise, we get the *signal to noise ratio*

$$SNR = \frac{\phi_x(f)|_{gw}}{\phi_x(f)|_\nu} = \frac{c}{2\lambda} \frac{h_{11}(f)}{\delta\nu(f)}. \quad (1.40)$$

The detector strain sensitivity $S_n(f)$ is defined as the locus of points in a single-sided power spectral density of the strain amplitude $h_{11}(f)$ where $SNR = 1$.

For the laser frequency noise that was discussed above, we can write this strain sensitivity as

$$\sqrt{S_v(f)} = \frac{2\lambda}{c} \delta\nu(f). \quad (1.41)$$

The laser source that is used in the Advanced Virgo gravitational wave detector is at the state of the art and has a typical free-running frequency noise of $100 \text{ Hz}/\sqrt{\text{Hz}}$ at a frequency of 100 Hz [12] away from the carrier. This frequency noise would result in a strain sensitivity $\sqrt{S_v(f)}$ of the order of $10^{-12}/\sqrt{\text{Hz}}$. Even when using a separate optical cavity for stabilising the laser, for example the 150 m long Advanced Virgo input mode cleaner cavity, $\sqrt{S_v(f)}$ is still in the order of $10^{-2} \text{ Hz}/\sqrt{\text{Hz}}$ at a frequency of 100 Hz [12] away from the carrier. This would result in a strain sensitivity of $10^{-16}/\sqrt{\text{Hz}}$, which is not good enough to detect gravitational waves by a long shot.

1.2.3 Michelson interferometer with Fabry-Pérot arms

What makes that gravitational waves passing Earth *can* in fact be detected, is the anti-symmetric nature of their polarisation signature. When looking at Fig. 1.1a again, we can see that if we lay out another Fabry-Pérot cavity along the y -axis, its end mirror will be shaken exactly in anti-phase with the cavity that was laid out along the x -axis. As a result, the phase transfer to $h_{11}(s)$ from this cavity is exactly the same as Eq. (1.38), but with opposite sign. However, any laser frequency fluctuations going into this cavity will not see this sign flip, and will be common in both cavities. When we now measure the *phase difference* between the light coming from the two cavities, the effect of the fluctuations coming from the laser will cancel out, while the sensitivity to gravitational waves doubles.

The interferometer layout that was used by Michelson and Morley [13, 14] in their famous null result when looking for the Earth's relative motion with respect to the so-called luminiferous ether, was designed to do exactly this: accurately measure phase differences between two beams of light that have travelled along different paths, but are coming from the same light source. Using light coming from the same source ensures that the relative phase of the two interfering beams is well defined, a necessary condition to measure a change in their relative phase. It also makes sure that any frequency fluctuations in the light will be correlated between the two beams, such that they can cancel upon detecting the phase difference at the output. The layout of two Fabry-Pérot cavity configured to operate as a Michelson interferometer is schematically shown in Fig. 1.3a.

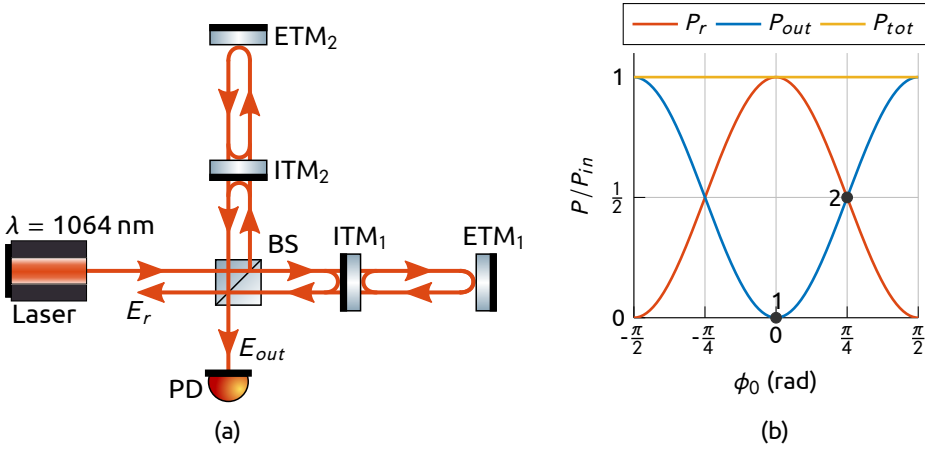


Figure 1.3 – (a) Schematic depiction of a Michelson interferometer with Fabry-Pérot arm cavities (PD: photodiode, BS: beam splitter, ITM: input test mass, ETM: end test mass). (b) Reflected and detected optical powers as a function of the setpoint ϕ_0 .

A laser sends a beam towards a 50:50 % beam-splitter that splits the laser beam into two beams of equal power going towards the two orthogonally laid out Fabry-Pérot cavities. The reflected light picks up a phase from a passing gravitational wave and interferes on the beam splitter. Taking all the reflections and phases into account, we can write the output field as

$$E_{out}(t) = -\frac{1}{2}E_{in} \left(e^{i\phi_y(t)-i\phi_0} - e^{i\phi_x(t)+i\phi_0} \right) = iE_{in} \sin(\phi_x(t) + \phi_0), \quad (1.42)$$

and the reflected field as

$$E_r(t) = -\frac{1}{2}E_{in} \left(e^{i\phi_y(t)-i\phi_0} + e^{i\phi_x(t)+i\phi_0} \right) = -E_{in} \cos(\phi_x(t) + \phi_0). \quad (1.43)$$

Here we have used the fact that $\phi_y(t) = -\phi_x(t)$, because of the anti-symmetric tidal signature of the gravitational wave for h_{11} . The setpoint phase ϕ_0 reflects our freedom to choose the interferometer's setpoint by moving the relative cavity positions with respect to the beam-splitter.

What we measure with the photodiode is the output power, given by

$$P_{out}(t) = E_{out}(t)E_{out}^*(t) = P_{in} \sin^2(\phi_x(t) + \phi_0). \quad (1.44)$$

The average output power along with the reflected power as a function of the setpoint phase ϕ_0 is shown in Fig. 1.3b. The small phase fluctuation induced by a gravitational wave $\phi_x(t)$ will modulate the output power around the chosen

setpoint. The highest sensitivity to phase fluctuations is at $\phi_0 = \pi/4$, but this is not a workable setpoint for several reasons:

- The photo-diode would need to cope with a DC offset of half the input power, which is generally high.
- Power recycling (Section 1.2.4) cannot be used in this configuration.
- The output is sensitive to fluctuations in the laser input power.

The latter of these is the most important. To see why, we need to realise that we are not interested in maximising the signal amplitude, but instead we are interested in maximising the signal-to-noise *ratio* (SNR). Writing the output power as $P_{out}(t) = \bar{P}_{out} + \delta P_{out}(t)$, Eq. (1.44) gives us the change in output power due to both $\delta P_{in}(s)$ and $\phi_x(s)$ as

$$\frac{\delta P_{out}(s)}{\delta P_{in}(s)} = \sin^2(\phi_0) \quad \text{and} \quad \frac{\delta P_{out}(s)}{\phi_x(s)} = P_{in} \sin(2\phi_0). \quad (1.45)$$

The signal to noise ratio is then given by

$$SNR = \frac{\delta P_{out}(f)|_{gw}}{\delta P_{out}(f)|_{P_{in}}} = \frac{\sin(2\phi_0)}{\sin^2(\phi_0)} \left(\frac{\delta P_{in}(f)}{P_{in}} \right)^{-1} \phi_x(f), \quad (1.46)$$

which for a given relative intensity noise in the laser $\delta P_{in}/P_{in}$, can be seen to be optimised for $\phi_0 \rightarrow 0^+$, *i.e.* a by using a small offset from the dark fringe at $\phi_0 = 0$. Again using the Advanced Virgo laser system as a reference, working at the naive working point of $\phi_0 = \pi/4$ would result in a minimum detectable phase $\phi_x(f)$ (at $SNR = 1$) of

$$\phi_{x,min}(f) = \frac{\sin^2(\pi/4)}{\sin(\pi/2)} \frac{\delta P_{in}(f)}{P_{in}} \approx \frac{1}{2} \cdot 5 \times 10^{-9} \approx 3 \times 10^{-9} \text{ rad}/\sqrt{\text{Hz}}, \quad (1.47)$$

which, using the Advanced Virgo design parameters in Table 1.1, would result in a strain sensitivity of $O(10^{-21})/\sqrt{\text{Hz}}$ through the transfer function in Eq. (1.38). This, again, is not good enough to detect gravitational waves. By instead using a more typical small offset from the dark fringe $\phi_0 \approx O(1)$ mrad, we could detect gravitational waves with strain amplitudes of $O(10^{-24})/\sqrt{\text{Hz}}$, which is about an order of magnitude below what the current generation of detectors achieves today.

1.2.4 Quantum noise limited interferometry

Now that we have established that a Michelson interferometer with two Fabry-Pérot arm cavities effectively acts as a transducer for gravitational waves, a

Table 1.1 – Advanced Virgo design parameters [11]

Parameter	Symbol	Design value
Arm cavity finesse	\mathcal{F}	443
Arm length	L	3 km
Laser wavelength	λ	1064 nm
Mirror mass	m	42 kg
Signal recycling mirror transmittance	T_{SRM}	0.2

natural question to ask ourselves is: “*how well can it perform?*” Recall from the discussion in Section 1.1.6 that there are in fact two important conditions that need to be fulfilled in order to be able to do a successful gravitational wave detection. Translated to the context of an interferometric detector these conditions read:

- We need an accurate measurement of the relative position of the interferometer’s end mirrors. This relative position will be modulated by a passing gravitational wave.
- Any parasitic forces acting on each of the interferometer’s mirrors should be reduced to a level below that of the gravitational wave’s tidal forces that we want to measure.

In practice, most of the effects that limit the strain sensitivity of interferometric gravitational wave detectors are associated with the latter category. However, there is one important bound on the accuracy of the mirror position measurement itself that will limit the interferometer’s sensitivity to gravitational waves even if the mirrors themselves remain perfectly still. This is a quantum noise effect, as it stems directly from the fact that light is quantised into photons. Because the laser beam that is used to perform the measurement consists of many individual photons, the action of measuring the interferometer’s output power with a photodiode effectively amounts to doing a counting experiment, and as with any counting experiment there are statistical fluctuations involved.

The shot noise limit

The error that is introduced by the statistical fluctuations in the amount of photons arriving at the photodiode per second was coined *shot noise*. To see to what extent this source of noise limits the sensitivity of our detectors, we have to compare its effect at the photodiode to the power fluctuations that are induced in the detector by the gravitational waves themselves. Let us take a closer look at the quantity that we measure to obtain the gravitational wave

signal, the interferometer's output power $P_{out}(t)$. Since each individual photon carries an energy $2\pi\hbar\nu$, the average value of $P_{out}(t)$ is given by

$$\bar{P}_{out} = 2\pi\hbar\nu\bar{N}_\gamma, \quad (1.48)$$

where \bar{N}_γ denotes the average number of photons arriving at the photodiode every second. For a coherent laser state, $N_\gamma(t)$ obeys Poissonian statistics, implying $\sigma(N_\gamma(t)) = \sqrt{\bar{N}_\gamma}$ for $\bar{N}_\gamma \gg 1$, so

$$\sigma(P_{out}(t)) = 2\pi\hbar\nu\sqrt{\bar{N}_\gamma} = \sqrt{2\pi\hbar\nu\bar{P}_{out}}, \quad (1.49)$$

where in the last step we have used Eq. (1.48) to eliminate \bar{N}_γ . From Eq. (1.44) we know that the average output power is given by $\bar{P}_{out} = P_{in} \sin^2(\phi_0)$, such that Eq. (1.49) represents a white noise with a single-sided spectral density of $\sqrt{4\pi\hbar\nu P_{in}} |\sin(\phi_0)|$. Comparing this to the spectral density of the power fluctuations generated by a gravitational wave in Eq. (1.45) and using the transfer function in Eq. (1.38), we obtain the shot noise limit of our detector for $\phi_0 \approx 0$,

$$\sqrt{S_{shot}(f)} = \frac{1}{8\mathcal{F}L} \sqrt{\frac{4\pi\hbar\lambda c}{P_{in}}} \sqrt{1 + \left(\frac{f}{f_c}\right)^2}. \quad (1.50)$$

This result shows us that above the cavity pole f_c the shot noise limit increases with frequency. This is a direct consequence of the detector's decreasing response to gravitational waves above the cavity pole. Additionally, we can see that the shot noise limit can be pushed down by increasing the optical input power incident on the beam-splitter. Given that the output power of continuous lasers is limited to $O(100)$ W, high input power is usually achieved by introducing another resonant optical cavity.

Recall from Fig. 1.3b, that around our dark fringe setpoint $\phi_0 \approx 0$, almost all the optical power is reflected back towards the laser, where it is lost. In analogy with the Fabry-Pérot arm cavities seen earlier, the power that impinges on the beam-splitter can be boosted significantly by putting another partially reflective mirror between the laser and the beam-splitter. When positioned correctly, the light reflected from this so-called *power recycling mirror* constructively interferes with the fresh light coming from the laser and the power on the beam-splitter builds up. The power recycling mirror in the Advanced Virgo detector, for example, is set up to provide a power recycling gain of 37.5, raising the $O(100)$ W of optical power that is coming from the laser to $O(5)$ kW impinging on the beam-splitter [11], significantly lowering the shot noise.

Radiation pressure noise

From the shot noise analysis above, it might seem that a gravitational wave detector can reach arbitrary sensitivities by simply increasing the optical input power. As always, however, there is a trade-off to be made here. The quantised nature of light does not only play a role in detecting the output power on a photodiode, it also manifests itself inside the detector's optical cavities themselves. On reflection at one of the mirrors, a photon changes its momentum from $+2\pi\hbar\nu/c$ to $-2\pi\hbar\nu/c$. This momentum is transferred to the mirror, creating a radiation force

$$F(t) = \frac{dp}{dt} = \frac{4\pi\hbar\nu}{c} N_Y(t) = \frac{2P_c(t)}{c}, \quad (1.51)$$

where we have used Eq. (1.48) to remove $N_Y(t)$ from the expression. This radiation pressure statically displaces the mirrors, but more importantly, since it is generated by numerous individually reflecting photons, it will show similar counting statistics as encountered before. This introduces an additional noise in the detector.

When we repeat the analysis in Section 1.2.1 for fluctuations in the input laser field, we find expressions for the fields inside one of the cavities:

$$\bar{E}_c = \frac{t_1}{1-r_1} \bar{E}_0 \approx \sqrt{\frac{2\mathcal{F}}{\pi}} \bar{E}_0 \quad \text{and} \quad \delta E_c(s) = \sqrt{\frac{2\mathcal{F}}{\pi}} \frac{1}{1+s/\omega_c} \delta E_0(s). \quad (1.52)$$

This shows that the fields impinging on the cavity are amplified by a factor $\sqrt{2\mathcal{F}/\pi}$, and for frequencies above ω_c the fluctuations are filtered by the cavity. The fluctuations in power resulting from $\delta E(t)$ can be obtained from

$$P(t) = (\bar{E} + \delta E(t)) (\bar{E} + \delta E(t)^*) \approx \bar{P} + \bar{E} (\delta E(t) + \delta E^*(t)) + O(\delta E(t)^2), \quad (1.53)$$

such that we have

$$\delta P(t) \approx \bar{E} (\delta E(t) + \delta E^*(t)). \quad (1.54)$$

Taking the Laplace transform and combining this with Eq. (1.52), we get a transfer function for the power fluctuations inside the cavity caused by fluctuations in the power impinging on it,

$$\frac{\delta P_c(s)}{\delta P_0(s)} = \frac{2\mathcal{F}}{\pi} \frac{1}{1+s/\omega_c}. \quad (1.55)$$

The relation in Eq. (1.49) also holds for the input power, so $\delta P_0(f)$ represents a white noise with a single-sided spectral density of $\sqrt{4\pi\hbar\nu P_0}$. From Eq. (1.51),

we then obtain the noise force acting on a mirror. Comparing this to a force induced by a passing gravitational wave gives the radiation pressure in terms of the strain sensitivity as

$$\sqrt{S_{rp}(f)} = \frac{16\sqrt{2}\mathcal{F}}{mL(2\pi f)^2} \sqrt{\frac{\hbar}{2\pi}} \frac{P_{in}}{c\lambda} \frac{1}{\sqrt{1 + (f/f_c)^2}}, \quad (1.56)$$

where an additional factor of 2 is picked up because of the fact that the gravitational wave force only acts on the cavity's end mirror, while the radiation pressure acts on both. Moreover, since only half the input power travels to each of the cavities, $P_{in}/2$ was substituted for P_0 . Equation (1.56) grows towards lower frequencies through the $1/f^2$ dependence. Moreover, it grows with increasing input power, putting it in direct conflict with reducing the shot noise.

The standard quantum limit

The opposite scaling with power for the shot noise and the radiation pressure noise makes that there is no straightforward solution to improve the detector's sensitivity. Increasing the power pays off at higher frequencies where the sensitivity is dominated by shot noise, but it makes matters worse at low frequencies where the radiation pressure contribution grows with the power. It turns out that for a given interferometer geometry, the optimum sensitivity is reached at the frequency where the two noise contributions are equal. At this specific frequency, neither decreasing nor increasing the power will improve the sensitivity. Such an optimum exists for all frequencies for a specific input power. The locus of these points for a power recycled Fabry-Pérot interferometer operated close to dark fringe is called the quantum limit and is given by

$$\sqrt{S_{QL}(f)} = \frac{1}{2\pi f L} \sqrt{\frac{8\hbar}{m}}. \quad (1.57)$$

Figure 1.4a illustrates the existence of a quantum limit by plotting the quantum noises for different input powers added in quadrature. The Advanced Virgo detector parameters listed in Table 1.1 are used to set the scale. Note that for each power, the quantum limit is reached at one frequency only and the total quantum noise never goes below this quantum limit.

Beating the standard quantum limit

The standard quantum limit of interferometry in Eq. (1.57) seems to put an ultimate limit to the sensitivity of a detector once the arm lengths and the mirror masses are given. However, it has been known since the eighties that there

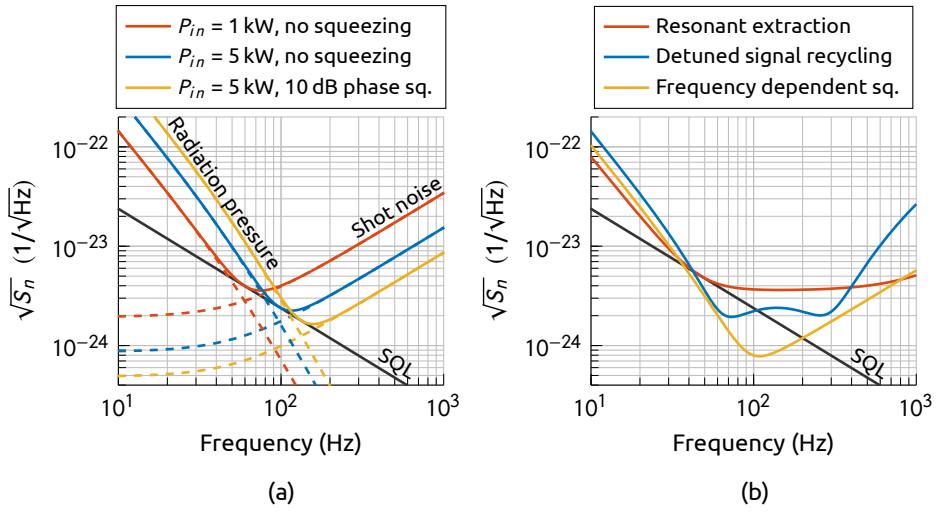


Figure 1.4 – (a) Quantum noise of a power recycled Michelson interferometer with Fabry-Pérot arms for different values of P_{in} . Frequency independent phase squeezing has the same effect as raising P_{in} . (b) The interferometer from (a) with $P_{in} = 5$ kW and three different quantum noise manipulation techniques. Note that the quantum noise can be brought below the standard quantum limit.

are several different techniques that can be exploited to obtain sensitivities beyond this quantum limit, some of which will briefly be discussed below.

The first solution is to change the interferometer response to gravitational waves, which can be done by adding yet another resonant cavity to the system. This was first proposed in 1981 [15] and the technique is usually called *signal recycling* or alternatively *resonant signal extraction*, depending on the resonant condition settings of the new cavity. The cavity is formed by adding a semi-transparent mirror between the beam-splitter and the photodiode. This often raises a rather intuitive question: “How can you detect more light on a photodetector by putting a highly reflective mirror directly in front of it?” In general, cavities can be used to enhance light fields inside the cavity, not outside. The key point with signal recycling is that we do not need to generate the detected light by sending it into the cavities ourselves. There already is a high energy reservoir stored inside the arms that gets modulated by a passing gravitational wave. Adding a signal recycling mirror then works as a resonant extractor, amplifying the amount of light that travels to the photodiode.

The resonant conditions of the signal recycling cavity can be optimised to achieve different results. This mirror’s position can be set such that signals with a specific frequency are reinjected into the interferometer with the right

phase to interfere constructively with “fresh” signals at that frequency, amplifying them before they propagate to the output. This configuration is called *detuned signal recycling*, and the resulting quantum strain noise surpassing the standard quantum limit is shown in Fig. 1.4b. The location and the depth of the typical noise minima observed in this configuration can be chosen by changing the reflectivity and position of the signal recycling mirror. This can be useful for looking for known sources at a specific frequency.

Alternatively, the signal recycling mirror can be positioned such that all signal frequencies are anti-resonant in the cavities. This lowers the storage time for gravitational wave induced signals inside the arm cavities and effectively raises the cavity pole for gravitational wave signals at the cost of slightly deteriorated peak sensitivity. The increase in the overall bandwidth of the detector in this *resonant extraction* configuration is obvious from the corresponding curve in Fig. 1.4b. Signal recycling is implemented in Advanced LIGO and planned to be implemented in Advanced Virgo, and will primarily be used in the resonant extraction configuration.

A second solution to beating the standard quantum limit is based on altering the quantum noise itself. The classical description of quantum noise introduced above tacitly assumes that there are no correlations between the contributions to shot noise and radiation pressure noise. This happens to be true for a coherent laser state, but is certainly not true in general. Shot noise originates from fluctuations in the phase quadrature of the light field, while radiation pressure noise stems from fluctuations in the amplitude quadrature [16]. These two quadratures are related by the well known *Heisenberg uncertainty relation* that states that the product of their uncertainties can never be smaller than a certain number. For the coherent laser state treated above, the fluctuations in the two quadratures are equal and minimal, resulting in the standard quantum limit of Eq. (1.57) [16].

The Heisenberg uncertainty relation only puts a restriction on the product of the shot noise and radiation pressure noise contributions, not on the individual noises themselves. By introducing correlations between the amplitude and phase quadratures, noise in one quadrature can be decreased at the expense of noise in the other quadrature, such that the uncertainty relation is still satisfied. Light that has these properties is called *squeezed light*, because noise from one quadrature is “squeezed” into the other. Using *amplitude squeezing* reduces the radiation pressure noise at the cost of increased shot noise, while *phase squeezing* does the opposite. These coupled effects on the two different quantum noises are exactly the same as what happens when the optical in-

put power is changed. Indeed, as shown in Fig. 1.4a, injecting phase squeezed light into the interferometer has the exact same effect as increasing the input power, and in itself does not allow sensitivities below the standard quantum limit. In present day detectors phase squeezing is used to improve the sensitivity. This makes sense because of two reasons:

- The low frequency sensitivity is not yet radiation pressure noise limited.
- It allows for lower shot noise without the thermal difficulties related to an increase in input power.

To beat the standard quantum limit by using squeezed light injection, *phase squeezed* light should be injected at high frequencies in order to reduce the shot noise that is dominant there, while at low-frequencies *amplitude squeezed* light should be injected to reduce the radiation pressure noise. Injecting light with such frequency dependent squeezing properties is possible by using the dispersion relation of optical filter cavities [16]. Doing this results in a broadband sensitivity improvement as shown in Fig. 1.4b. This frequency dependent squeezing is planned to be implemented in future upgrades of the current detector facilities, as well as in future generation detectors.

After all these improvements to a given power recycled Fabry-Pérot interferometer, the resulting total quantum strain noise sets the ultimate limit for its sensitivity. It reflects the accuracy limit to which the mirror positions can be measured. Recall, however, that this was only the first of two requirements to be fulfilled for successful gravitational wave detection. The remaining requirement is that we reduce all other forces that act on the mirrors to levels below the gravitational wave force we want to measure, which ideally is given by the quantum noise. With the planned broadband improvement in quantum noise provided by frequency dependent squeezing, this becomes increasingly harder to do. The most important noise sources will be treated below.

1.2.5 Classical noise sources

Aside from the noise arising from the quantum nature of the light used to measure the relative positions of all the mirrors, there is a multitude of different noise sources of classical origin that will mask a gravitational wave signal. For starters, the mirrors used as test masses are not perfect. Mechanical losses in the fused silica substrates and high-reflectivity coatings will inevitably lead to a random motion of the mirror surfaces. Moreover, building the detectors on Earth means having to deal with seismic motion some ten orders of magnitude higher than the mirror displacements to be measured. Figure 1.5 shows the Advanced Virgo design sensitivity. At high frequencies the sensitivity is set

by quantum shot noise, but below a few hundred Hertz there are many other significant noise sources that limit the sensitivity. Since this is the frequency band where a lot of interesting physics happens, it is important to understand all these noise sources to be able to minimise them.

Thermal noise

When building a gravitational wave detector, we can not beat thermodynamics. The well-known *fluctuation-dissipation theorem* [17] dictates that in all the places in the detector where some form of energy is dissipated into heat, an inverse path exists that couples statistical thermal fluctuations to the other energy domain. There are several different relevant loss mechanisms associated to the interferometer's main mirrors that all introduce their own noise in this way. All these forms of thermal noise are completely analogous to the more familiar Johnson noise in electronic circuitry, where noise enters the circuit through resistors exactly because that is where electric energy is dissipated into heat. Generally, resistor thermal noise is modelled by a voltage source that is driving an electrical circuit with a white noise of single-sided power spectral density of $v_n^2(f) = 4k_B T \Re(Z_{el}(f))$. Here k_B is the Boltzmann constant, T denotes temperature and the real part of the circuit's impedance $\Re(Z_{el}(f))$ is just the resistance value R by definition. This notion nicely carries over to other energy domains. For example, when considering losses in the mechanical domain the resulting noise can be modelled by a driving force with a single-sided power spectral density of $F_n^2(f) = 4k_B T \Re(Z_m(f))$, when we define the *mechanical impedance* as $Z_m(f) = F/v$ [18]. Here F denotes the force driving the system and v is its velocity.

Each of the mirrors in the interferometer is suspended by a monolithic system of four thin fused silica fibres to minimise the mechanical losses, but nevertheless the losses are there. Modelling the loss as an imaginary part of the pendulum's effective spring constant k_{eff} , the transfer function in response to a force acting on the mirror is

$$\frac{X(s)}{F(s)} = \frac{1}{ms^2 + k_{eff}(1 + i\phi_{eff}(s))}. \quad (1.58)$$

Here $\phi_{eff}(s)$ denotes a so-called loss angle that in general can depend on frequency. The mechanical impedance of this system can easily be obtained from the transfer function as

$$Z_m(s) = \frac{F(s)}{V(s)} = \frac{F(s)}{sX(s)} = ms + \frac{k_{eff}}{s} + i \frac{k_{eff}\phi_{eff}(s)}{s}. \quad (1.59)$$




Figure 1.5 – Design sensitivity for the Advanced Virgo detector at 125 W laser power, no squeezing and using the signal recycling mirror in the resonant extraction mode. The total noise is the incoherent quadrature sum of the shown noise sources. This result corresponds to the tuned signal recycling noise budget from [19].

To obtain the effective noise acting on the mirror, we need the real part of this, which is given by $\Re(Z_m(f)) = k_{eff}\phi_{eff}(f)/2\pi f$ because $s = 2i\pi f$. This results in a noise force with a single-sided amplitude spectral density of

$$F_n(f) = \sqrt{4k_B T \frac{k_{eff}\phi_{eff}(f)}{2\pi f}}. \quad (1.60)$$

Comparing this with a force induced by a passing gravitational wave we get

$$\sqrt{S_{susp}(f)} = \frac{f_0}{f^{5/2}L} \sqrt{2k_B T \frac{\phi_{eff}(f)}{\pi^3 m}}, \quad (1.61)$$

where f_0 denotes the natural frequency of the main pendulum mode which in Advanced Virgo is about 0.6 Hz [20]. Since the mirrors are suspended in a high vacuum environment, viscous damping through the residual gas is negligible and the losses are dominated by internal losses in the fused silica fibres. Away from any loss peaks, $\phi(f)$ is typically constant over a large range of frequencies for that type of loss. For the fibre system used in Advanced Virgo, $\phi \simeq O(5 \times 10^{-7})$ [21]. The effective loss angle is a lot lower because in a pendulum, only a small part of the energy is stored in the bending of the fibre, the

majority of the restoring force is generated by gravity, which is lossless. For this reason, ϕ_{eff} is typically of the order of $O(10^{-10})$. The resulting noise curve in Fig. 1.5 is labelled *suspension thermal noise*, and due to the $1/f^{5/2}$ behaviour, is mostly important at low frequencies. The resonant feature around 8 Hz is the vertical mirror bouncing mode. This vertical motion couples to the cavity length due to the fact that the detector is so big that the local vertical directions at both mirrors are not parallel, but due to the curvature of the Earth are misaligned by about 0.4 mrad. The resonant peaks above 400 Hz are caused by the violin modes of the suspension wires.

A second major source of thermal noise in the detector is caused by mechanical losses in the mirrors themselves. These are a combination of losses in the fused silica bulk with $\phi_{bulk} \simeq O(10^{-8})$ and losses in the high-reflectivity coatings with $\phi_{coat} \simeq O(10^{-4})$. Even though the coatings are thin, they dominate the total loss because they are four orders of magnitude more lossy than the mirror bulk. The resulting effective loss angle is $\phi_{eff} \simeq O(10^{-7})$ [22]. Following a similar argumentation as above, these losses cause fluctuations in the position of the mirror surfaces. Comparing these to the mirror displacements induced by a passing gravitational wave we get the effective strain noise. Below the first internal mirror resonance, which is typically around 7.8 kHz for Advanced Virgo [23], this noise contribution can be found to be [22]

$$\sqrt{S_{coating}(f)} = \sqrt{\frac{8k_B T (1 - \nu^2) \phi_{eff}}{\pi^{3/2} L^2 f r_0 E}}, \quad (1.62)$$

where the effect was averaged over a Gaussian laser beam shape of radius r_0 , and E and ν are the mirror's Young's modulus and Poisson ratio respectively. As can be seen in Fig. 1.5, this contribution only falls off with frequency as $1/f^{1/2}$ so it is relevant over a wider frequency range. In fact, right where the detector is most sensitive, the so-called '*bucket*', this is the limiting noise source. Significant effort is being spent on finding coating materials that improve this thermal noise contribution, making this an active research field.

The mirrors do not only experience mechanical losses, also optical energy is lost through absorption. According to the fluctuation-dissipation theorem, this gives rise to another thermal noise contribution. The noise coupling occurs through two different phenomena, both related to fluctuations in the temperature of the coating through optical power dissipation. The first contribution is the so-called *thermorefractive noise*, that enters through the temperature dependence of the index of refraction of the coating materials, altering the reflection phase when temperature fluctuations occur. The second contribution




Figure 1.6 – Measured seismic activity at the three Advanced detector sites along with Peterson’s high- and low-noise models [26]. Solid curves correspond to the 50th percentile, shaded regions indicate the 10th to 90th percentiles. Virgo data were measured with a Guralp 40T-30s in the Central Building close the beam-splitter from 4 June 2011 to 3 September 2011 [27]. Both LIGO site spectra used an STS2 seismometer in a similar location for the full month of July 2010 [28].

stems from the linear expansion coefficient of the coating material that will slightly displace the mirror surface upon the temperature fluctuating. Their collective effect on the strain noise is dubbed *thermo-optic noise* and is shown in Fig. 1.5 not to be limiting for the Advanced Virgo design.

The coupling mechanisms of all the mentioned thermal noises are diverse, but what all of them have in common is that they scale with the temperature like \sqrt{T} . This suggests that cooling down the mirrors to cryogenic temperatures would help in mitigating these effects, and indeed this is the case. However, there are significant challenges involved in cooling down mirrors to cryogenic temperatures in a high vacuum environment, especially when they are illuminated by $O(1)$ MW optical power. Moreover, the mechanical loss properties of fused silica and the tantala used in the coatings significantly deteriorate at these temperatures [24], and alternative materials are needed, involving significant R&D work. Pioneering work on cryogenic interferometry is currently being performed at the Japanese KAGRA detector, that operates sapphire mirrors, cryogenically cooled to 20 K [25].

Seismic noise

The surface of the Earth is moving constantly. The oceans hitting the shore, wind blowing on trees and buildings and human activity, they all transfer energy into the Earth that drives seismic motion. On human scales, these ground vibrations are small and we do not feel them. However, on the scale of gravitational waves they are huge and they pose an enormous challenge for gravitational wave detectors. Recall from Fig. 1.5 that above a frequency of 10 Hz, Advanced Virgo's strain sensitivity is better than $O(10^{-22})/\sqrt{\text{Hz}}$. With an arm length of 3 km, this translates to an effective mirror motion of $O(10^{-19}) \text{ m}/\sqrt{\text{Hz}}$. Figure 1.6 shows typical horizontal motion spectra measured at the three advanced detector sites of LIGO and Virgo. The seismic motion at 10 Hz for all sites is typically $O(10^{-10}) \text{ m}/\sqrt{\text{Hz}}$. The anchors of the suspension mechanisms that hold the mirrors are fixed to the ground, so for the detector to work, they need to provide over 9 orders of magnitude of vibration suppression above 10 Hz.

The required vibration attenuation is provided by using the passive vibration isolation properties of simple pendula. Above its natural frequency, a pendulum rejects the horizontal motion of the suspension point with a typical $(f_0/f)^2$ characteristic, indicating that long pendula are beneficial. A 10 m long pendulum has a natural frequency of 160 mHz and provides a factor of $O(10^3)$ attenuation at 10 Hz. This falls quite a lot short of the required 9 orders of magnitude, but luckily there is a more efficient way to provide passive vibration attenuation in the same 10 m space. By cascading multiple shorter pendula, the individual natural frequencies go up and the performance at low frequencies worsens. However, above all the system's eigenmodes the attenuation provided improves drastically because it scales like $1/f^{2N}$, with N the number of cascaded pendula. Figure 1.7a illustrates this by comparing the ground to mirror transfer functions of different pendulum chains all having the same total length of 10 m.

Figure 1.7c shows a schematic representation one of the *superattenuators* that provide the seismic attenuation in the Advanced Virgo detector. The horizontal attenuation is provided by a chain of 6 pendula suspended from an inverted pendulum pre-isolator stage. Figure 1.7b shows its performance, indicating a staggering ground motion suppression of over 14 orders of magnitude at 10 Hz [11]. The seismic noise in Fig. 1.6 gets filtered by this transfer function, resulting in the steep $1/f^{14}$ seismic noise curve in Fig. 1.5. This noise is often referred to as the *seismic wall*, for its steep slope makes it quite challenging to move it down significantly in frequency.

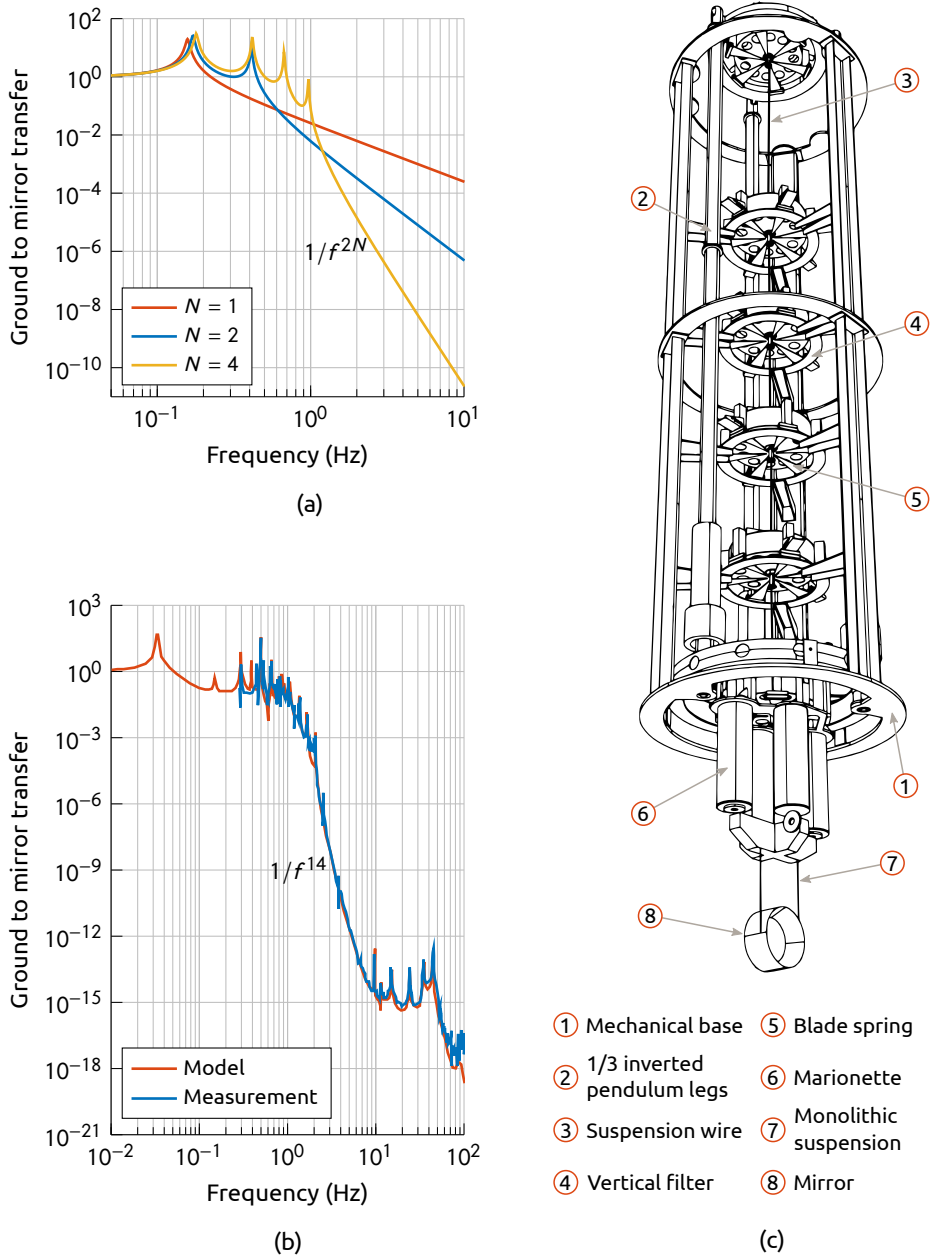


Figure 1.7 – (a) Vibration attenuation provided by a chain of N pendula with a total length of 10 m. The superattenuator providing the seismic attenuation for the main optics in Advanced Virgo is shown in (c) and consists of a 8.66 m chain of 6 pendula hung from a 6 m inverted pendulum attenuator [11] (image adapted from [29]). Its horizontal performance is shown in (b). (Figure reproduced from [11]).

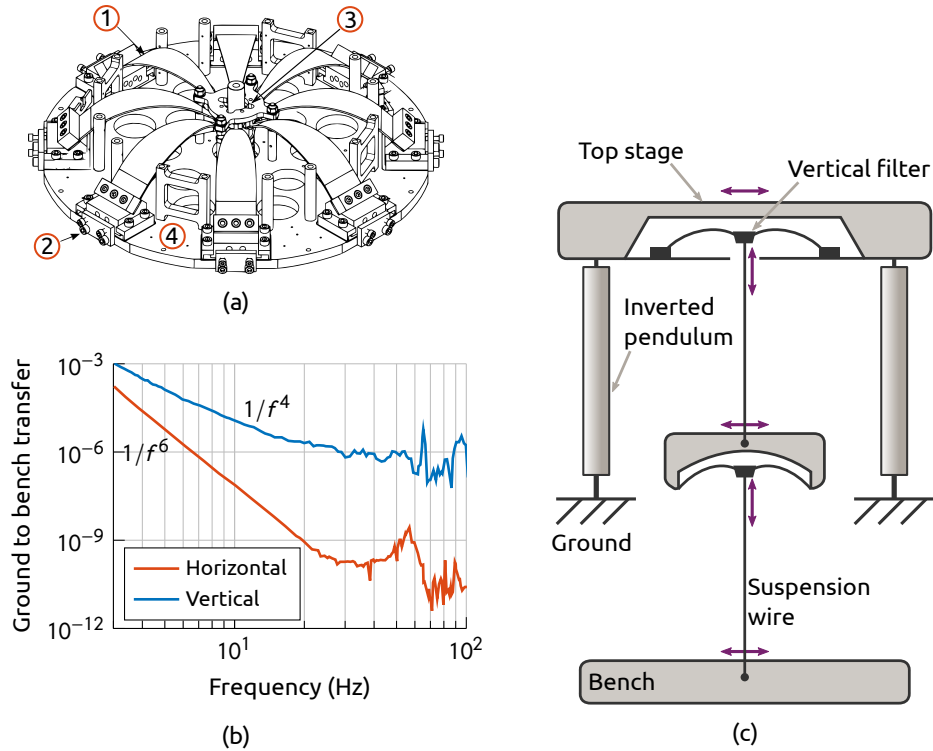


Figure 1.8 – (a) A rendering of the vertical geometric anti-spring filter used in Advanced Virgo’s auxiliary suspensions. The load is connected to the keystone (3), symmetrically suspended by a set of blade springs (1) connected to a base plate (4). The compression forces can be tuned with the screws in the compression mechanism (2). A schematic overview of the multiSAS suspension system containing two of these filters is shown in (c), with its performance indicated in (b). Images adapted from [31, 32].

The pendula themselves do not provide any suppression for vertical ground motion. Because of several mechanisms coupling vertical to horizontal motion at the mirror, significant vertical attenuation is required as well. The intermediate suspended masses contain sets of blade springs that suspend the following pendulum stage, forming a series of cascaded mass-spring-systems providing vertical seismic suppression analogous to the cascade of pendula. Since lower natural frequencies provide more attenuation, the natural frequencies of the individual blade spring systems is brought down from 1.5 Hz to about 500 mHz by means of magnetic anti-springs added around the suspension point [30].

An alternative to using magnetic anti-springs for the natural frequency reduction of vertical seismic filters is to use a purely geometric solution: the *geometric anti-spring*. This technique is used in the seismic suspensions of the

auxiliary optics in Advanced Virgo. A schematic representation of a filter implementing this solution is shown in Fig. 1.8a. The load to be suspended is carried by the central *keystone* that is surrounded by a symmetrically arranged set of blade springs. These springs are laterally loaded, such that the resultant force that is generated when the keystone is displaced from its equilibrium position effectively introduces a negative stiffness [33]. By carefully tuning the compression forces, almost all of the blade spring's internal stiffness can be cancelled, resulting in a highly compliant system. Since the operating principle of the microseismometer presented in this work is largely based on the geometric anti-spring principle used in this type of filters, it will be detailed upon in Chapter 3.

Newtonian Noise

Although challenging, through structures as the superattenuator, the interferometer's mirrors can be decoupled from the direct seismic motion at the suspension point for all frequencies above a few Hz. There is, however, a second effect related to seismic motion that can not be shielded. The local mass distribution surrounding the gravitational wave detector exerts a direct gravitational pull on its mirrors, and since seismic motion slightly perturbs this mass distribution, it also perturbs this gravitational field. The resulting noise force is called *Newtonian noise* or alternatively, *gravity gradient noise*. Figure 1.5 shows that above a few Hz, this effect is larger than the seismic motion itself. For Advanced Virgo, Newtonian noise is not expected to be a limiting noise source, but when moving to future infrastructures with improved sensitivity, this has to be dealt with.

A second source of Newtonian noise is related to density fluctuations in the atmosphere surrounding the detector's test masses. These fluctuations can be caused by infrasound waves, temperature changes or turbulent atmospheric phenomena, and cause so-called *atmospheric Newtonian noise* in the detector. A recent study shows that this noise can be rather high, and for the Advanced Virgo detector it is close to the design sensitivity [34].

As Newtonian noise couples directly to the interferometer's main mirrors, there is no way of shielding it. A future detector must therefore be built at a site with a low background seismic activity. In addition, both atmospheric and seismic Newtonian noise contributions are significantly suppressed when building the detector in an underground facility⁵ [34]. Indeed, finding a suitable location for such a detector is actively pursued, and seismic measurement campaigns have been, and are being performed to assess their suitability [35].

⁵ The explanation for the seismic contribution is clear from Fig. 1.16 in Section 1.4.3.

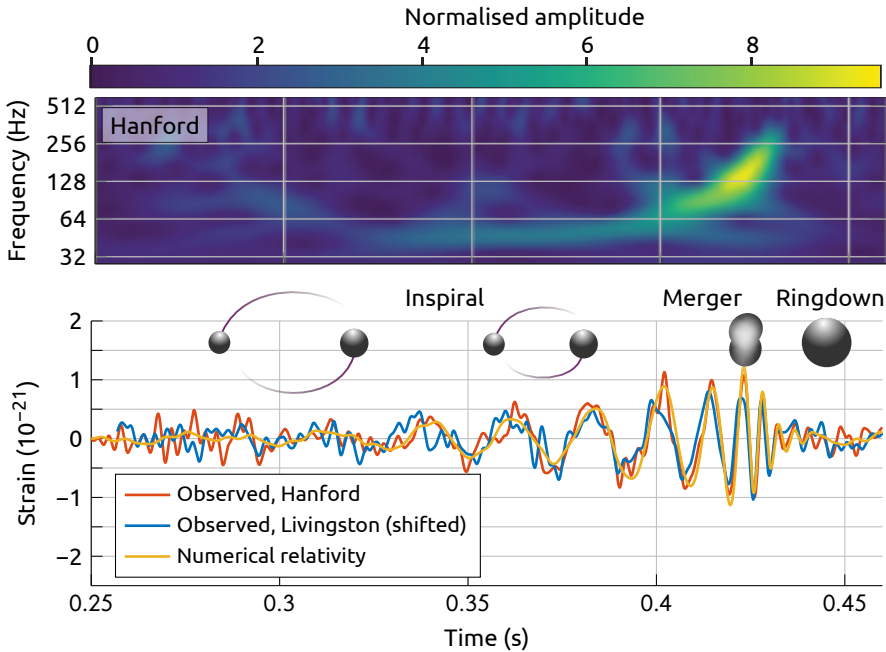


Figure 1.9 – Bottom panel: Time domain data as measured in both the LIGO detectors for the GW150914 event, along with the best fit numerical relativity waveform. The Livingston data were inverted and shifted to correct for the different arrival time. Top panel: A spectrogram made using the Hanford data showing the characteristic chirp signature. Figure adapted from [1], data source [36].

That Newtonian noise can not be shielded from affecting the mirrors does not mean that it can not be mitigated. The effect on the mirrors themselves is so small that it can only be measured with the interferometer itself. However, the effect on the mirrors will be correlated with the movement of the Earth’s crust around the mirrors, which is much larger. Monitoring the ground motion in the vicinity of the interferometer with a large array of seismic sensors is the most promising candidate mitigation technique for seismic Newtonian noise. The noise requirements on these seismometers are strict, for the detector must be built in a location with low seismic activity. More details on these Newtonian subtraction arrays will be given in Section 2.2.

1.3 The importance of good sensitivity at low frequencies

On September 14 2015 at 09:50:45 UTC the two LIGO detectors in the United States recorded their first gravitational wave signal: *GW150914* [1]. It origin-

ated from the violent merger of two black holes each with a mass roughly 30 times that of our Sun that happened over a billion light years away. The signal measured in both LIGO detectors is plotted in the bottom panel in Fig. 1.9 with only minimal bandpass filtering applied. The signal was so strong it is visible in the data by eye, something that was clearly unexpected. It has the characteristic shape of a gravitational wave generated by two compact objects colliding, where the amplitude grows while the two objects spiral towards each other, losing more and more energy to the emitted gravitational waves. As their orbit shrinks the objects speed up and start revolving around each other faster and faster, causing the signal frequency to grow, as indicated by the spectrogram in the top panel of Fig. 1.9. When the two black holes finally merge, they form a single highly excited black hole that sheds its last gravitational waves at a characteristic ringdown frequency, forever to remain silent after.

1.3.1 Low-frequency sensitivity in Advanced LIGO

The two LIGO detectors had already been looking for gravitational wave signals for extended periods of time between 2005 and 2010, but had never seen a real signal [37, 38]. Yet, within the first week of turning them on in their upgraded Advanced detector configuration, GW150914 shows up in the data with an incredibly high SNR of 23.7 [39]; *how is that possible?*

To answer that question, one should look at what has changed with the upgrade from Enhanced LIGO to Advanced LIGO. After the upgrade, the shot noise at high frequencies did improve by roughly a factor 3 to 4, but the main reason for the high SNR of GW150914 is the vastly improved sensitivity at the lower end of the frequency band. At 50 Hz, the sensitivity of Advanced LIGO is almost a factor 100 better than that of Enhanced LIGO [40]. Figure 1.10 shows the best Enhanced LIGO sensitivity at the Hanford site obtained during the S6 observation run in 2010, along with the Advanced LIGO design target and the actual sensitivity during the O1 run when GW150914 was observed. Along with the sensitivity curves, a normalised version of the Fourier transform of GW150914, $h(f)$, is shown. Realising that the expected SNR of a gravitational wave signal can be written as [39]

$$SNR^2 = \int_0^\infty \left(\frac{2|h(f)|\sqrt{f}}{\sqrt{S_n(f)}} \right)^2 d\ln(f), \quad (1.63)$$

we can interpret the area between the signal and noise curves as a measure for the expected SNR. The total area of the three shaded regions corresponds




Figure 1.10 – Sensitivities for LIGO in different configurations [40] along with the signal spectrum of GW150914 [39]. The area between the noise and the signal curves is a measure for the expected SNR. The total shaded area corresponds to the measured SNR of 23.7. The yellow shaded area represents a sub-threshold SNR < 8 that would be measured with Enhanced LIGO. Purple denotes the additional SNR obtained when scaling the Enhanced LIGO sensitivity to the O1 sensitivity around 200 Hz. As the red area is larger, most additional SNR was gained by improving low-frequency sensitivity.

to the measured SNR of 23.7 in the case of GW150914. In the Enhanced LIGO configuration, what would have been left of the SNR for the same event is represented by the yellow shaded region. With an SNR < 8, even this loud event would have been sub-threshold [40]. The purple shaded region is the SNR that would have been gained by scaling the Enhanced LIGO sensitivity to the O1 level around 200 Hz. As the red shaded region is much larger, it is clear that most of the sensitivity gain is due to the changes that were made to improve the sensitivity at low frequencies.

1.3.2 Large redshift observations

For electromagnetic observations, the wave frequencies tend to be so high that their amplitude cannot be directly measured. Therefore, most electromagnetic detectors are power-detectors. Because of energy conservation, the power arriving on Earth falls off as $1/r^2$ with r the distance to the source. For the energy of gravitational waves arriving on Earth, the same scaling relation holds, but there is one key difference: GW detectors are sensitive to the *amp-*

litute of the wave, not the power. The amplitude falls off as the square root of the power, so the measured signals only fall off as $1/r$. This scaling behaviour means that a sensitivity improvement of a factor α allows to observe the same event a factor α further out in the Universe. An interferometric gravitational wave detector is in essence omnidirectional, so this means that we can probe a volume of the Universe α^3 times larger, increasing the rate of measured events by the same factor.

One of the challenges associated to looking at events increasingly further away is related to the expansion of the Universe. All signals travelling to Earth from distant sources obtain a cosmological redshift z by the fact that the Universe expands during their travel time. This redshift value is often used as an equivalent measure for the distance to astronomical objects. Signals that had a frequency f_s upon their emission from a source at redshift z , will have their frequency lowered to $f_s/(1+z)$ by the time they reach Earth. Low-frequency sensitivity therefore becomes increasingly important as the sensitivity of gravitational wave detectors, and with it their detection range, improves.

GW150914 was detected at a distance of over a billion light years, which sounds like far away, but with an equivalent redshift of $0.09^{+0.03}_{-0.04}$ [1], was actually quite close on a cosmological scale. The measured signal frequencies were therefore only about 10 % lower than at emission. The total source frame mass for GW150914 was measured to be $65.3^{+4.1}_{-3.4} M_\odot$. Figure 1.11 shows that at design sensitivity, Advanced LIGO would be able to see a similar event out to a redshift of $O(1)$. At such a distance, the detected merger frequency would shift from 250 Hz to 136 Hz, where the sensitivity is still good. Next to the curve for Advanced LIGO, the projected range for binary coalescence events for different future detector scenarios is shown as well. Einstein Telescope (ET), for example, would potentially see a GW150914-like event all the way out to $z = O(50)$.⁶ The only reason Einstein Telescope could see this far, is that next to an order of magnitude increase in sensitivity compared to Advanced LIGO, the sensitive detection band is opened up to significantly lower frequencies. The merger of a GW150914-like event at $z = O(50)$ would be measured at 5.5 Hz, a frequency where a LIGO-like detector would have no significant sensitivity at all. More details on Einstein Telescope can be found in Section 1.4.

⁶Actually, no galaxies would have formed for redshifts this high, so Einstein Telescope would see *all* binary black hole mergers of stellar origin in the entire Universe. The potential for events to observe here would be mergers of primordial black holes created shortly after the Big Bang itself [41].




Figure 1.11 – Astrophysical reach for equal-mass non-spinning compact binary inspiral systems of future detector scenarios compared to Advanced LIGO. The maximum observable distance is shown as a function of the total intrinsic mass of the system. A Hubble constant of $67.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and a Λ CDM model of expansion were assumed. Figure adapted from [42].

1.3.3 Source parameters

In essence, black holes are exceptionally clean objects. They are extremely massive, yet in GR they are made up out of pure vacuum, not matter. The dynamics of uncharged test particles in their proximity can be fully described by two quantities only: the black hole's *mass* and its *angular momentum* (or equivalently, its *spin*). When two black holes merge, the shape of the resulting emitted gravitational waveform is also set by these parameters. Measuring such a waveform therefore gives detailed information about the properties of the black holes that collided. A gravitational waveform as measured by a detector on Earth needs additional parameters related to the relative orientation of the detector to the binary, making source parameter estimation a computationally expensive endeavour. Moreover, the dependence on the relative orientation introduces some degeneracies between the parameters. This makes constraining some of the source parameters a difficult task. Some degeneracies can be broken by using multiple detectors (*e.g.* the sky location can be independently constrained by triangulation), others, such as the relation between the binary's inclination and the observed distance will remain [43].⁷

⁷ For a binary neutron star merger this degeneracy can be broken when a host galaxy can be identified through an optical counterpart, providing an independent measure of the redshift, as was the case with GW170817 mentioned in Section 1.3.4 [44].

The wave's frequency evolution is mainly set by the system's *chirp mass*,

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad (1.64)$$

with m_1 and m_2 the binary component masses. The chirp mass can accurately be derived from the gravitational waveforms measured by the Advanced detectors [45]. The individual component masses are significantly less well constrained, because of a correlation with the black hole spins that are poorly measured. Misaligned black hole spins cause the binary orbit to precess, which in turn causes amplitude modulations in the measured waveform, as illustrated in Fig. 1.12 [46]. The waveforms measured by the Advanced detectors are at present simply too short (in band time $\ll 1$ s) to significantly constrain these effects. Extending the low-frequency band to 1 Hz would stretch the in band time of a signal from a pair of merging $10 M_\odot$ black holes to 2 days, allowing precision measurements to be performed on the source parameters. This would be of significant benefit for population studies [47].

Gravitational wave measurements only capture the last moments in the existence of a merging binary system. Doing precision measurement of the source parameters, however, also provides useful information about the binary's history and the environment it was formed in. Measuring the relative spin orientation of black hole binaries, for example, can shed light on the main formation channel of such binaries. Did the binary already exist as binary of massive stars that both became black holes during the binary's lifetime, or was it formed by dynamical capture in a dense globular cluster? The former would leave a binary with the spins being aligned with the orbital angular momentum during a so-called *common envelope* phase, and the latter would produce a binary with random relative spin orientations [48]. Distinguishing these formation channels, we will learn about the environments of these black holes and gain new insights on the conditions in the distant Universe.

1.3.4 Multimessenger astronomy

During the last cycles before a binary merges, the constituent black holes can orbit each other tens to hundreds of times per second, emitting enormous amounts of energy in the form of gravitational waves. What they do not emit, however, is light. Black holes are compact and extremely massive objects but they do not consist of matter, such that there is nothing present to emit any electromagnetic radiation. This is different when neutron stars are involved. A neutron star is the stellar remnant of a massive star that packs the total mass




Figure 1.12 – Numerically generated gravitational waveform for a coalescing pair of black holes with component masses $m_1 = 36 M_\odot$ and $m_2 = 29 M_\odot$ at a distance of 500 Mpc. Both black holes have a dimensionless spin magnitude of approximately 0.8 with the spins aligned in (a) and strongly misaligned in (b). The advanced detectors only measure the last fraction of these waveforms, making the precession signature in (b) difficult to observe.

of the sun in a sphere with a radius of only $O(10)$ km. In contrast to a black hole, a neutron star does consist of matter. When two of them merge, next to a gravitational wave signature, they emit electromagnetic radiation in a large range of different wavelengths.

The first detection of gravitational waves from a binary neutron star merger, GW170817, instantly showed the power of multimessenger astronomy, where information from multiple detection channels is combined to do unprecedented measurements. The gravitational wave signal shown in Fig. 1.13a was picked up by both the Advanced LIGO detectors as well as the Advanced Virgo detector. Having an observation in 3 detectors allowed an online analysis to triangulate the source of the gravitational wave to an area on the sky of 31 deg^2 with 90 % probability shown in Fig. 1.13b. An offline analysis eventually improved the source localisation to 16 deg^2 [45]. Within half an hour the first of these sky maps was distributed among astronomers, triggering a large electromagnetic follow-up campaign in wavelengths ranging from γ -rays to radio waves.

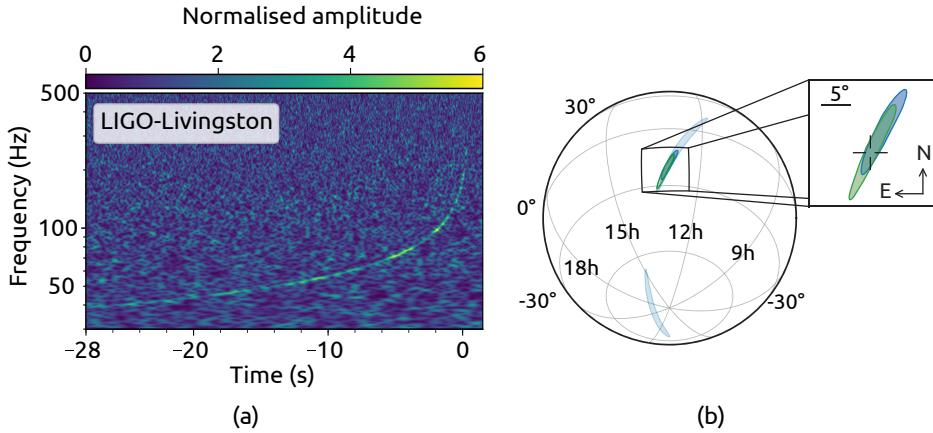


Figure 1.13 – (a) Spectrogram showing characteristic gravitational wave chirp signal for a coalescing pair of neutron stars. Times are shown relative to August 17, 2017 12:41:04 UTC. The signal is in band for almost 30 s. (b) Sky location reconstructed for GW170817 from a Hanford-Livingston (190 deg^2 , light blue contours) and Hanford-Livingston-Virgo (31 deg^2 , dark blue contour) analysis. A higher latency Hanford-Livingston-Virgo analysis improved the localisation (28 deg^2 , green contour), while the offline parameter estimation improves this to 16 deg^2 [45]. The reticle in the inset marks the position of the apparent host galaxy NGC 4993. Figure adapted from [44].

About 1.7 s after the merger time determined from the gravitational wave signal, the Fermi space based detector had already picked up a short γ -rays burst that generated an automated trigger message [49]. This was followed by the alert and the sky map from the LIGO-Virgo observation. The coincidence of these two observations prompted astronomers worldwide to follow-up on these triggers. About 10 h later, an optical counterpart was identified by multiple telescopes, identifying the host galaxy to be NCG 4993, roughly 40 Mpc away from Earth. The optical transient was identified as a so-called kilonova [50] and was traced through different wavelength ranges for several weeks, piling up a vast amount of relevant information. The joint measurement campaign of this single event produced an independent measurement of the Hubble constant of $70.0^{+12.0}_{-8.0} \text{ km s}^{-1} \text{ Mpc}^{-1}$ [51], bounded the speed of gravitational waves to be equal to the speed of light with a relative accuracy of $+7 \times 10^{-16} / -3 \times 10^{-15}$ [52], and linked merging neutron stars as the driving force behind short γ -ray bursts [49].

The rich scientific output from a single event exemplifies the potential of multi-messenger astronomy as a new research field. Also in this context, improving the low-frequency sensitivity of gravitational wave detectors brings

significant benefits. Figure 1.13a shows that GW170817 enters the sensitive frequency band of the current Advanced detectors approximately 30 s before the merger. When this frequency band would be extended down to 1 Hz, such a signal would be in the detection band for almost a week, and even a reduction to 5 Hz would allow detecting the signal for about 2 h [47]. The time of the forthcoming merger, along with an estimate of its sky location would be known before it even took place, allowing the implementation of an early-warning system that would greatly benefit multi-messenger observations.

1.4 Einstein Telescope

The era of gravitational wave astronomy has only just begun, and, as we have seen above, there is a lot to gain in improving the sensitivity of the instruments. The current Advanced LIGO and Virgo facilities will be brought to their designed sensitivity in the coming years, and new detectors in Japan (KAGRA) and India (LIGO-India) will join the network. There are projects to implement new technology in the existing facilities to improve the sensitivity even further. There will, however, be an end to how far things can be improved, and building new detector infrastructure becomes a necessity for further improvements.

There are different initiatives for realising such a third generation detector infrastructure. In the US, plans are made for a successor of Advanced LIGO called *Cosmic Explorer* [42]. The European initiative is called *Einstein Telescope*, or often just ET for short. This section focuses on Einstein Telescope because it is the most mature of the existing initiatives. An extensive conceptual design study was performed [47], and several pathfinder experiments for testing new technologies are currently being started up. In early 2018, the first steps were taken to form the ET Scientific Collaboration, and at the time of writing, over 600 people signed the ET Letter of Intent [53].

1.4.1 Baseline layout

At the core of Einstein Telescope will still be the dual recycled Michelson interferometer with Fabry-Pérot arm cavities, this time with arm lengths of 10 km to make the first step in towards increased sensitivity. Conceptually, however, there are significant changes with respect to detectors such as Advanced Virgo and Advanced LIGO. The detector will be built underground for reasons considered in Section 1.4.3. To make maximum use of the available tunnel space, ET will not consist of a single detector, but of three separate detectors in a triangular configuration as shown in Fig. 1.14a. Each vertex will be used as a




Figure 1.14 – (a) Einstein Telescope baseline layout. Three subterranean caverns hosting the main optical components will be connected by 10 km tunnels that host the vacuum vessels for the laser beams. A tunnel cross-section is shown in (b). ET will consist of three detectors (different colours in (a)) to maximise the use of available tunnel real estate. Each detector consists of two interferometers, one for high, and one for low frequencies. Figures adapted from [54, 55].

corner station for one of the detectors. With three detectors, ET will be able to disentangle the two different polarisation states of a gravitational wave, something that was not possible before.

Each of the three detectors that make up ET consists of two individual interferometers in a so-called *xylophone configuration*. One of them (ET-HF) is dedicated to measuring high frequencies, the other (ET-LF) is optimised to measure the low end of the frequency spectrum down to about 2 Hz. To implement the frequency-dependent squeezing necessary for ET-LF, two additional 1 km long filter cavities are needed. The main 10 km tunnels will contain a total of 4 long optical cavities in their own vacuum pipes, as shown in Fig. 1.14b.

Using the xylophone design the quantum noise trade-off with power as described in Section 1.2.4 can be avoided. The power in each interferometer can be individually optimised for its intended frequency band. ET-HF will boast 3 MW of optical power in the arm cavities to drive the shot noise down, while ET-LF will be operated with 18 kW only to keep the radiation pressure noise at an acceptable level. An additional benefit of this configuration is that the thermal issues related to cryogenically cooling an interferometer operating with high

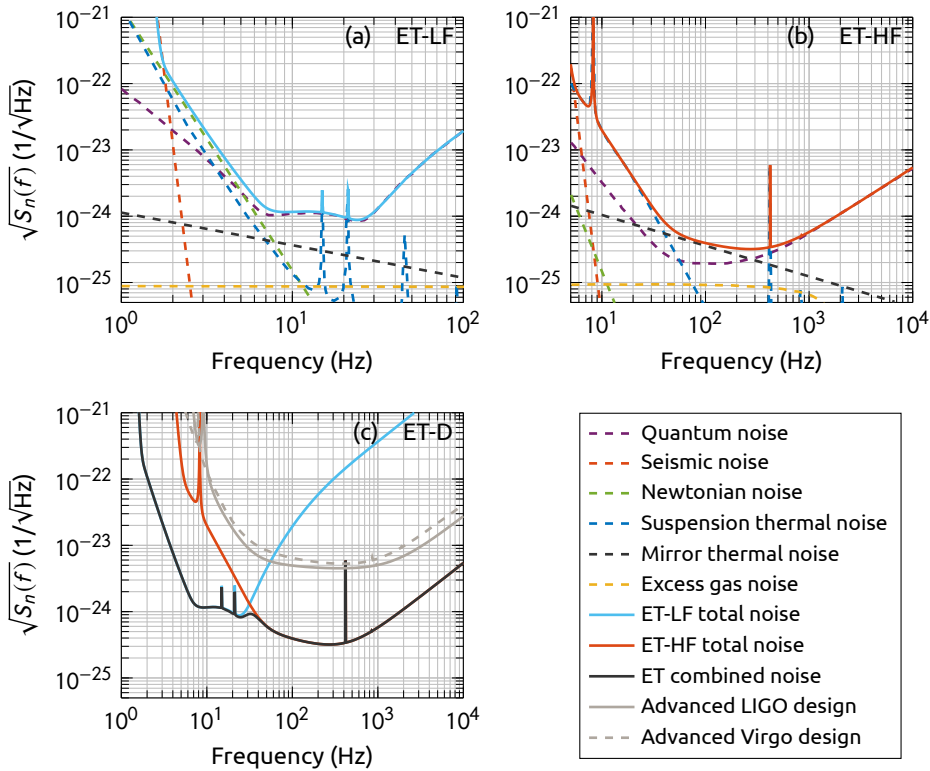


Figure 1.15 – Noise budget for Einstein Telescope’s low-frequency (a) and high-frequency (b) interferometers. (c) The total sensitivity of Einstein Telescope compound from its two component interferometers as per the design study (ET-D). Compared to the Advanced detectors, ET will be an order of magnitude more sensitive and will extend the sensitive band down to 2 Hz [47].

optical power can be avoided. The thermal noise described in Section 1.2.5 is typically more predominant at lower frequencies, such that it is only necessary to cool the ET-LF interferometer.

1.4.2 Sensitivity target

The total baseline design sensitivity for a single detector in ET is shown in Fig. 1.15c. It is the combined from its high- and low-frequency interferometer noise curve, also shown in the figure. Over the entire band, ET will be an order of magnitude more sensitive to gravitational wave strains than Advanced LIGO or Virgo. Moreover, the low-frequency cut-off will be moved almost an order of magnitude down to about 2 Hz. This incredible increase in sensitivity will allow ET to detect $O(10^5)$ binary coalescence events per year, allowing

population statistics on astrophysical objects to accumulate quickly [56].

At frequencies above 300 Hz the sensitivity is limited by the quantum noise of the ET-HF interferometer shown in Fig. 1.15b, which is minimised by both using high optical power and frequency dependent squeezing. Between 40 and 300 Hz rather the thermal noise from the mirror coatings is the limiting noise source, similar to the case in the current Advanced detectors. Below 30 Hz, ET-LF is more sensitive, and its quantum noise is dominant down to about 6 Hz, as shown in Fig. 1.15a. The radiation pressure is kept in check by a combination of using a lower optical power, heavy mirrors of about 200 kg, and frequency dependent squeezing. Below 6 Hz the sensitivity is fully limited by seismic effects. An improved version of the Advanced Virgo superattenuator shown in Fig. 1.7c will push the seismic noise limit just below 2 Hz, but even then, without reduction techniques, the Newtonian noise will be limiting up to 6 Hz. This means that a lot of the actual achievable performance at low frequencies will be set by the seismic properties of the site selected to build the infrastructure. Any improvement in sensitivity beyond this relies on Newtonian noise subtraction techniques based on input from auxiliary inertial sensors, as will be described in Section 2.2.2.

1.4.3 Site selection

The seismic noise curves in Fig. 1.15 were generated with a typical horizontal ground motion of $5 \times 10^{-10}/f^2$ m/ $\sqrt{\text{Hz}}$, a value typically observed at the underground Black Forest Observatory in Germany. This value is now the proposed target seismic background level between 1 and 10 Hz for a potential host site [35]. In this frequency band, most of the seismic motion is of anthropogenic origin. It is therefore beneficial to choose a host site in a region with low population density. However, a trade-off has to be made there, as decent road infrastructure is needed for the construction phase, and there should be scientific institutions in the area for operating the detector.

Sites with the required seismic background spectrum exist in Europe, but usually one needs to go underground to find them. Since most noise in the 1 to 10 Hz frequency band is anthropogenic, it is generated at the Earth's surface. Going underground helps to suppress this, as illustrated by a measurement in a Hungarian mine shown in Fig. 1.16. How far one needs to go underground strongly depends on the soil type and on the local geology. Currently, ET is foreseen to be built at a depth of 100 to 200 m. A borehole study at 250 m depth is under way in an area near Maastricht in the Southern part of The Netherlands to verify the suitability of the region to host ET.




Figure 1.16 – Seismic spectral variations as a function of depth in the GyöngyöSOROZI mine in Hungary. Data taken during the same period (3 - 6 April 2010) at the surface and at 70 m and 400 m depths. Above 1 Hz the seismic power is more than an order of magnitude weaker at a depth of 400 m than at the surface. Figure reproduced from [35].

Of course, choosing a site with low background seismic does not guarantee that the seismic activity is still low after the infrastructure has been built and the detector is in operation. Traffic, machinery and a bunch of people present at the detector site all produce vibrations. Care has to be taken in constructing the facility to avoid excessive man-made vibrations that spoil the location's seismic characteristics. For example, a recent study shows that the main source of low frequency seismic and acoustic noise inside the Advanced Virgo buildings is the air conditioning system [57]. There are things that can be done to improve this (*e.g.* connecting pumps and fans through spring suspensions to decrease their coupling to the buildings), but there is only so much you can do without a complete redesign of the air conditioning system. These seemingly trivial issues will become increasingly more important when moving the detector facility to a seismically quiet environment, as it will be much harder not to spoil it.

Inertial sensing and Newtonian noise

As seen in Chapter 1, several of the challenges to be overcome for gravitational wave detectors to work are related to seismic motion. Properly understanding and monitoring this motion around the test masses is therefore important. The Advanced Virgo detector for example, currently has about 175 inertial sensor channels for recording the site's seismic background and related vibrations of the detector mechanics, and to generate control signals for the feedback loops of all the seismic attenuation systems. An additional 120 inertial sensors are planned to be installed around its masses to implement and test Newtonian noise reduction techniques [58].

When moving to increasingly lower corner frequencies in future gravitational wave detectors such as Einstein Telescope, seismic effects become even more important. Moreover, the strict requirements on the seismic activity for a site to host Einstein Telescope have implications for the noise specifications of the sensors that can be used there. This chapter briefly introduces the basic operation principle of different types of seismic inertial sensors, then goes on to discuss their application in active Newtonian noise mitigation strategies related to the low-frequency performance of gravitational wave detectors, and ends with a comparison of the noise performance of the MEMS based seismic sensor presented in this work to the state of the art in seismic sensing.

2.1 Inertial sensors

Seismic sensors are traditionally inertial sensors. Depending on the application they come in many shapes and sizes, and their system architectures can vary wildly. However, at the heart of any inertial sensor is the same basic mechanical sensing element: *a mass on a spring*.

2.1.1 Mechanical sensing element

Figure 2.1a shows a schematic representation of the mechanical sensing element in a typical inertial sensor. It shows a proof mass m that is connected connected to a frame through a spring with stiffness k and a dashpot with damping constant γ . The equation of motion of this system is given by

$$m\ddot{x}_m + \gamma(\dot{x}_m - \dot{x}_f) + k(x_m - x_f) = 0, \quad (2.1)$$

where x_m and x_f denote the displacement of the proof mass and the frame, respectively. Because of the inertia of the proof mass and the fact that the connection to the frame is not rigid, the motion of the proof mass will lag that of the frame. The relative motion between the two, $x = x_f - x_m$, can be used as a proxy for the motion of the frame x_f , and therefore for the motion of anything the frame is attached to. Rewriting Eq. (2.1) in terms of the relative displacement x and converting to the Laplace domain, we get

$$\ddot{X} + \frac{\omega_0}{Q}\dot{X} + \omega_0^2 X = \ddot{X}_f, \quad (2.2)$$

where the usual substitutions $\omega_0^2 = k/m$ and $Q = m\omega_0/\gamma$ have been made to denote the system's natural frequency and quality factor, respectively.

The relevant transfer function to consider for the sensor's frequency response depends on the type of transducer employed to monitor the relative motion. In inertial sensors, both displacement transducers and velocity transducers are commonly used. The first type produces an output signal proportional to X , while the latter outputs a signal proportional to \dot{X} . The amplitude responses of sensing elements with both types of transducers and a natural frequency of 1 Hz are shown in Fig. 2.1b for comparison, with respect to frame displacement as well as frame velocity and acceleration. All these transfer functions can easily be obtained from Eq. (2.2). The displacement transfer with respect to a frame acceleration is given by the well-known transfer of a harmonic oscillator

$$\frac{X(s)}{\ddot{X}_f(s)} = \frac{1}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2}, \quad (2.3)$$

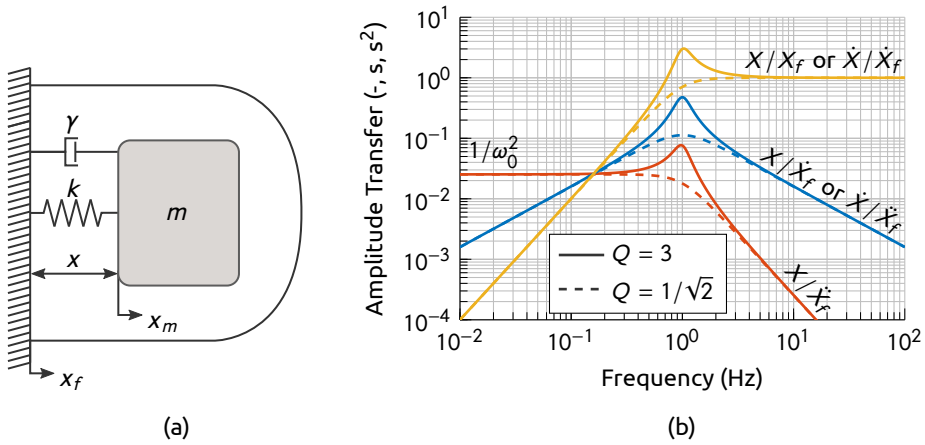


Figure 2.1 – (a) A schematic representation of a simple mechanical harmonic oscillator that can be used as an inertial sensing element. The amplitude responses of either the relative displacement X or velocity \dot{X} with respect to a frame displacement, velocity or acceleration are shown in (b). The system response is maximally flat in the passband for $Q = 1/\sqrt{2}$. Figure adapted from [59].

and since differentiation in the Laplace domain corresponds to a simple multiplication with the frequency parameter s , the other transfers differ only by factors of s .

For an inertial sensor that measures proof mass displacement, the transfer X/\ddot{X}_f is flat *below* the natural frequency, as shown in Fig. 2.1b. A signal in this frequency band can therefore be immediately interpreted as an input acceleration, and hence a sensor operating in this configuration is called an *accelerometer*. Instead, for an inertial sensor measuring the proof mass velocity such as a geophone, the relevant response \dot{X}/\dot{X}_f is flat *above* the natural frequency, and the sensor output is proportional to the frame velocity in this frequency band. The amount of damping, parametrised by Q , determines the response around the natural frequency. For $Q = 1/\sqrt{2}$, the transfer function in Eq. (2.3) reduces to a 2nd order Butterworth filter characteristic which has a maximally flat response in the passband [60].

When comparing the transfer functions for both transducer types to the same input stimulus, for example frame acceleration, it is clear that using a displacement transducer is favourable for measurements at low frequencies. The accelerometer transfer X/\ddot{X}_f remains flat all the way down to DC, while the geophone response \dot{X}/\dot{X}_f drops with 20 dB per decade with decreasing frequency. The response of an accelerometer at low frequencies is simply given by $1/\omega_0^2$, indicating that for the same transducer resolution, a lower natural

frequency results in a better noise performance. Indeed, modern broadband seismometers use a high resolution displacement transducer in combination with a low natural frequency mechanical oscillator. However, the geophone remains a popular sensor due to its simplicity and robustness. Its transducer consists of a coil that moves along with the proof mass through the static magnetic field of a permanent magnet attached to the frame, directly providing a voltage that is proportional to their relative velocity. As a result, the geophone acts as a simple voltage source that outputs an electrical signal which can easily be processed further. Moreover, the ideal quality factor of $Q = 1/\sqrt{2}$ can be set by simply shunting the sensing coil through a resistor of the correct value [60]. Displacement transducers such as LVDTs [61], capacitance bridges [62, 63] or even interferometers [64] often provide a higher resolution, but also require increasingly complex signal processing electronics.

2.1.2 The force-balance principle

Because its transfer function below ω_0 is proportional to $1/\omega_0^2$, an accelerometer needs to have a low natural frequency to have a good sensitivity at low frequencies. To be able to manufacture a low frequency mechanical oscillator with a compact form factor, a *LaCoste suspension* can be used that combines a special zero-length spring with a specific pendulum geometry to obtain theoretically infinite periods [65]. Since the required zero-length spring is relatively difficult to manufacture, compact seismometers typically employ a leaf spring suspension that approximates the LaCoste suspension only around the sensor's working point [59, 61].

Due to the intricate equilibrium of forces in these types of suspensions, they are sensitive to all kinds of external influences such as tilt and temperature changes. This makes them almost impossible to operate without an active feedback system that stabilises the mass position, and therefore the most sensitive seismic inertial sensors are typically built using the force-balance principle [60]. This type of sensor uses an actuator that generates a force to balance the inertial acceleration, such that the proof mass is forced to move along with the accelerating frame. A schematic representation of the general architecture of such a system is shown in Fig. 2.2. The proof mass reacts to an acceleration of the frame a_f with a transfer function $M(s)$ as given by Eq. (2.3), and a displacement transducer with a certain gain G_{el} converts the resulting displacement into a voltage for further processing. The displacement transducer output is then filtered by a loop filter $C(s)$ to produce the sensor's output voltage V_{out} . Successively, this signal is processed by a feedback filter $F(s)$ to gener-

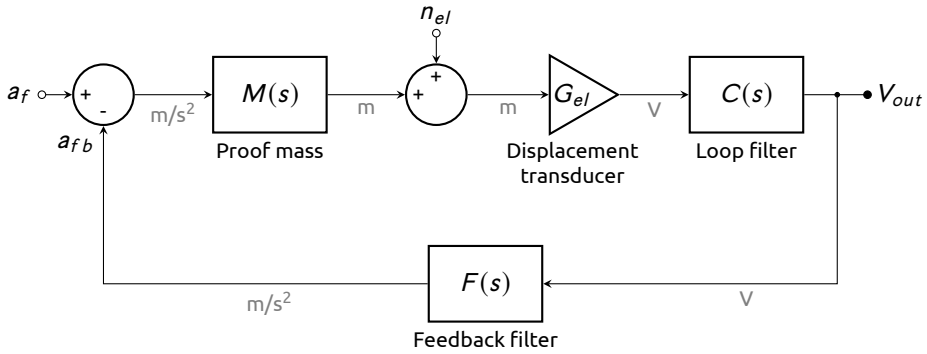


Figure 2.2 – Schematic representation of an inertial sensor built on the force-balance principle. Within the feedback system’s bandwidth, the generated feedback acceleration a_{fb} is approximately equal to the input acceleration a_f , and the feedback filter $F(s)$ determines the sensor response. The proof mass response to acceleration $M(s)$ is given by Eq. (2.3).

ate a feedback force, or equivalently, a feedback acceleration a_{fb} . Within the loop’s bandwidth, this feedback acceleration compensates the inertial acceleration a_f and minimises the relative motion between the proof mass and the frame. As a result, the requirements on the linearity of the mechanical suspension and the displacement transducer are significantly relaxed, and are now shifted towards the force transducer, which can be made to be accurate due to the limited range of motion [60].

Because of the feedback action in the loop, the sensor’s transfer function is modified. Tracing the effect of a frame acceleration through the loop and solving for V_{out}/a_f , we get

$$\frac{V_{out}(s)}{a_f(s)} = \frac{MG_{el}C}{1 + FMG_{el}C} \approx \frac{1}{F(s)}, \quad (2.4)$$

in which the last approximation holds only within the loop bandwidth where the loop gain is high, *i.e.* $|FMG_{el}C| \gg 1$. Since the transfer function in Eq. (2.4) does not depend on $M(s)$ within the loop’s bandwidth, the frame acceleration is converted to an output voltage without relying on the accuracy of the suspension system. Instead, the sensitivity is determined by the feedback filter $F(s)$ only, which is both relatively straightforward to calibrate, and more importantly, stable over time. The feedback, however, does not decrease the displacement transducer noise, and even though Eq. (2.4) shows no direct dependence on ω_0 , a low natural frequency is still necessary for a good noise performance [60]. To illustrate this, consider the noise introduced by the displace-

ment transducer (n_{el}) in Fig. 2.2. It enters the loop at a different location different from that of a_f , and therefore also has a different transfer function to the output. Tracing the effect of this noise source through the loop and solving for V_{out}/n_{el} , we get

$$\frac{V_{out}(s)}{n_{el}(s)} = \frac{G_{el}C}{1 + FMG_{el}C} \approx \frac{1}{M(s)F(s)}, \quad (2.5)$$

where again this last approximation holds only within the loop bandwidth. To refer the noise to the input of the sensor, this result still has to be divided by the closed-loop transfer function in Eq. (2.4). This yields the equivalent acceleration noise at the input caused by the displacement transducer as

$$n_{el, input} = \frac{n_{el}}{M(s)} \approx \omega_0^2 n_{el} \quad \text{for } \omega \ll \omega_0. \quad (2.6)$$

revealing that a mechanical sensing element with a low natural frequency is still beneficial for a good noise performance at low frequencies. Of course, enclosing the mechanical oscillator in a feedback loop introduces additional noise sources into the system. In general, each of these noise sources will have a different transfer function to the output, and all these should be carefully taken into account when modelling the noise in a feedback system. A detailed treatment of different noise sources and their impact on the total system performance can be found in Chapter 4.

As Eq. (2.4) shows, the response of a force-balance system is inversely proportional to the gain of the feedback path within the bandwidth of the loop. The response can therefore be tailored to specific applications by changing the frequency dependence of the feedback gain $F(s)$. In the case that the feedback filter only consists of a linear force transducer we have $F(s) = G_{fb}$, and Eq. (2.4) tells us that V_{out} is directly proportional to the frame acceleration. A sensor operating in this configuration is called a force balance accelerometer, or FBA. By choosing an appropriate loop filter $C(s)$, the sensor's useful bandwidth can be extended well beyond ω_0 , the usual bandwidth limit for a passive accelerometer, and the dynamic range can be up to 50 dB higher than that of a passive accelerometer [59]. The MEMS seismometer presented in this work is an example of an FBA. Details on its feedback architecture can be found in Section 4.3.2.

An FBA is sensitive to acceleration all the way down to DC and therefore measures its own tilt and thermal drift. To avoid the resulting reduction in dynamic range of the digitiser, sensitive broadband seismometers usually provide

an output voltage that is proportional to the frame's velocity instead of its acceleration. The proof mass in these sensors is still kept stationary with a feedback acceleration through the force-balance principle, but they achieve an output proportional to velocity through a frequency dependent feedback filter $F(s)$. This filter is typically implemented in the form of a standard PID controller [59, 63] that has a transfer function given by

$$F(s) = G_p + \frac{G_i}{s} + G_d s, \quad (2.7)$$

where the gains G_p , G_i and G_d represent the gain of the proportional, integral and derivative terms of the feedback filter, respectively. These gains are chosen such that the differential feedback term dominates above a corner frequency f_c of typically a few mHz. For frequencies above f_c , we therefore have $F(s) \approx G_d s$ and through Eq. (2.4) we see that within the control loop's bandwidth, the output is now proportional to velocity, because

$$\frac{V_{out}(s)}{a_f(s)} = \frac{V_{out}(s)}{s v_f(s)} = \frac{1}{G_d s} \quad \Rightarrow \quad \frac{V_{out}(s)}{v_f(s)} = \frac{1}{G_d}. \quad (2.8)$$

For frequencies below f_c , the integral feedback takes over and the closed-loop response with respect to $v_f(s)$ drops with 40 dB per decade towards lower frequencies [60]. The result is that the shape of the transfer function of such a broadband seismometer is described by the \ddot{X}/\dot{X}_f response in Fig. 2.1b. It therefore behaves exactly like a regular passive geophone, but with an effective natural frequency that has been electronically reduced to the mHz-range, extending the useful lower frequency range by 3 orders of magnitude.

2.2 Mitigating Newtonian noise

All different types of inertial sensors described above find applications in the Advanced Virgo gravitational wave detector. For example, the auxiliary seismic suspensions in Fig. 1.8 are inertially damped by using geophone signals [31], while the controls for the main superattenuators use custom FBAs [66] in order to have sufficient sensitivity around the rather low inverted pendulum mode of about 30 mHz. Broadband seismic motion is monitored by a combination of FBAs and broadband seismic sensors [27].

A relatively new application of inertial sensors in the context of gravitational wave detectors is the subtraction of so-called Newtonian noise, as described in Section 1.2.5. As can be seen in Fig. 1.5, the strain sensitivity of the

Advanced Virgo detector is not expected to be limited by Newtonian noise under normal seismic conditions. For next generation detectors such as Einstein Telescope, however, this is another story altogether. Newtonian noise mitigation strategies are an important part of the design of such a next generation infrastructure.

2.2.1 Passive suppression

The fact that the interferometer's test masses can not be shielded from the effects of Newtonian noise does not mean that these effects cannot be avoided to a large extent. Newtonian noise scales directly with the seismic motion around the interferometer, and the first most important thing to do then, is to make sure to build a new detector in a place with low seismic background motion. A comprehensive study of the typical seismic activity of different low-noise sites around the world [35] shows that most of the seismic motion in the relevant frequency band of 1 to 10 Hz is created by sources at the surface related to human activity, and it is therefore beneficial to build the detector in an underground location. The level of seismic activity at a site is important, but it is not the only thing that determines the resulting noise in the detector. The strength of the coupling to the detector also depends on the soil type and its dispersive properties, and even on the local topology around the test masses. While the soil properties are an integral part of the selection of a candidate site, the latter dependence may be used to minimise the coupling to the detector by optimising the size and shape of the caverns that host the test masses. Simulations show that trenches and excavations in the vicinity of the test masses can change the Newtonian noise coupling by a factor of a few [67].

Even when choosing a host site with a low background seismic activity, the Einstein Telescope design sensitivity in Fig. 1.15 is still expected to be limited by Newtonian noise from 2 to 6 Hz. At the time of writing, there are two serious candidate sites to host this future gravitational wave detector facility: one in the south of The Netherlands close to both the German and Belgian borders, and another one on the island of Sardinia in Italy. In Sardinia, an old mine allows for subterranean seismic measurements [68], while near Terziet in The Netherlands, a 250 m deep borehole seismometer was deployed for that purpose [69].

Although the surface seismic activity in the Newtonian noise band in the south of The Netherlands is far higher than in Sardinia because it is not as sparsely populated, the subterranean spectra are not as far apart as one might expect. Most seismic motion in the Newtonian noise band is anthropogenic,




Figure 2.3 – Design sensitivity for Einstein Telescope with different horizontal Newtonian noise (NN) projections. The projection used in the design study (dashed line) uses the analytical Saulson model and seismic data from the Black Forest Observatory [70]. The projection at the Terziet site uses a numerical model calibrated to both surface and borehole seismic data. The shaded band represents the 10th – 90th percentiles as well as modelling uncertainties. Figure adapted from [69].

and is therefore predominantly generated at the surface. Like any other type of wave, a seismic wave reflects off surfaces that represent an impedance contrast in the medium. In the region of the Dutch site, there is a soft soil layer roughly 40 m thick on top of hard rock which confines the seismic waves generated at the surface mostly to this soft soil layer, away from any underground detector. The reduction ratio of motion at the surface to that at 250 m depth can be as high as 40 dB during day time [69]. This layered geology has the added benefit that the gravitational wave detector is automatically protected against unforeseen anthropogenic surface noise sources that will exist in the future. This is a big advantage, since with the construction of Einstein Telescope we do not only build a detector, but a gravitational wave observatory that should operate for many decades.

Preliminary Newtonian noise estimates for the Terziet site have recently been presented. Figure 2.3 shows the Einstein Telescope design sensitivity from Fig. 1.15 along with the Newtonian noise projection as published in its conceptual design study [47]. That specific Newtonian noise projection was made with Saulson’s analytic half space model [71] using a reference seismic

acceleration of $2 \text{ ng}/\sqrt{\text{Hz}}$ as observed at the Black Forest Observatory in Germany [69]. Being a half space model, Saulson's analytical way of calculating the Newtonian noise was devised specifically for surface detectors and oversimplifies the situation for a subterranean one. To get a realistic estimate of the Newtonian noise at the Terziet site, a more sophisticated numerical model was set up to solve the full elastic wave equation for a layered geology matching the one seen around the borehole. Data from both the array of seismic sensors at the surface and the borehole seismometer were used to calibrate the model. The resulting Newtonian noise projection is shown Fig. 2.3, where the shaded band represents the 10th - 90th percentiles as well as modelling uncertainties [69]. This result shows that the expected Newtonian noise is up to a factor 3 above the design sensitivity, and additional measures have to be taken to mitigate its effects when Einstein Telescope is to reach design sensitivity at the Terziet site.

2.2.2 Active suppression

As mentioned before, the gravitational coupling between the soil and the mirrors in the detector is direct, and the mirrors can therefore not be shielded from the effects of Newtonian noise. The noise itself can only be measured directly by the interferometric gravitational wave detector, but is expected to be correlated with the seismic motion in the vicinity of the test masses in the detector. Accurately monitoring this seismic field will therefore in principle allow coherent subtraction through for example Wiener filtering [72] to mitigate the effect of Newtonian noise on the detector's output. The Newtonian noise projection in Fig. 2.3 shows that if Einstein Telescope is to reach design sensitivity at the Terziet site, such an active Newtonian noise subtraction technique is quite necessary and should suppress the effects of Newtonian noise on the detector output by a factor 3 – 5. In fact, it is now recognised that planning for Einstein Telescope must include active suppression techniques [73], as it can benefit the sensitivity even for quieter sites.

The exact implementation of the seismic sensor arrays that are required to do Newtonian noise subtraction is still under active investigation. Several detailed numerical studies have been performed on cancelling Newtonian noise arising from isotropic wave fields at the surface [35, 74], and only more recently the cancellation of the effects of body waves on subterranean detectors like Einstein Telescope were investigated [73]. Estimates for the required number of seismic sensors and their distribution vary strongly, from only a few tens of seismometers per test mass for a fully optimised array [73], to several hun-

dreds for a regularly spaced grid [35]. Moreover, it is still a challenge in the modelling of subterranean Newtonian noise to make the correct assumptions on the relative contributions from the different wave types present in a seismic field [75], so the last word on the design of these arrays has yet to be spoken.

The only instruments that are sensitive enough to detect Newtonian noise are the interferometric gravitational wave detectors themselves. Therefore, there is no hope for experimentally testing and verifying any active subtraction technique without having a detector whose sensitivity is limited by Newtonian noise [76]. Under normal seismic conditions, the sensitivity of the Advanced Virgo detector is not expected to be limited by Newtonian noise. However, seismic activity in general is far from stationary, and especially during bad weather conditions the seismic background can be elevated such that the associated Newtonian noise *is* expected to limit the sensitivity at low frequencies. Currently, an additional 120 geophones are planned to be installed around the Advanced Virgo test masses to subtract the effect of Newtonian noise during high seismic activity [58]. More importantly, however, this array should provide an experimental verification of the devised subtraction techniques using seismic sensor arrays.

The subtraction efficiency has been shown to depend on the signal to noise ratio in the seismic sensors that constitute the subtraction array [35, 76]. The requirements on the level of seismic background activity of a potential host site for a detector such as Einstein Telescope therefore directly translate to strict requirements on the noise performance of the seismic sensors to be used in the subtraction array. The seismometer self noise in the Newtonian noise frequency band of roughly 1 – 10 Hz should at least be better than the $2 \text{ ng}/\sqrt{\text{Hz}}$ ground acceleration used to generate the Einstein Telescope design sensitivity.

The potentially large number of sensors with good noise performance required for these future Newtonian noise arrays has sparked the interest for investigating low-noise MEMS sensors, because of their ability to be mass produced at relatively low cost. No commercial MEMS sensor is currently available that meets the low-noise requirement of $2 \text{ ng}/\sqrt{\text{Hz}}$ for use in future Newtonian noise cancellation arrays, but the novel MEMS seismometer presented in this work demonstrated a performance below this target and shows potential for additional improvement. The next section will put the new sensor in perspective by comparing its noise performance to the state of the art in seismic sensors. The remainder of this work is dedicated to describing the design and performance of the novel MEMS seismometer.




Figure 2.4 – A comparison of the noise performance of the novel MEMS seismic sensor presented in this work (horizontal channel) with that of a broadband seismometer (Trillium 240), a 5 Hz geophone (INOVA Quantum, vertical channel), the commercially available MEMS seismometer with the lowest acceleration noise (Sercel QuietSeis, horizontal channel), and of the MEMS based seismometer with the lowest noise performance ever recorded (Mars seismometer, vertical channel). A breakdown of the noise performance of the novel MEMS seismometer can be found in Fig. 4.18.

2.3 State of the art in seismic sensors

Broadband seismometers such as for example the Streckeisen STS-2 [63] or the Nanometrics Trillium 240 [77] have been the golden standard for seismic sensing for decades. They offer an output that is flat to ground velocity between roughly 10 mHz and 10 Hz, with a self-noise that is below the Earth’s low noise model over that entire frequency band. Especially for frequencies below 1 Hz their extremely low self-noise below 50 pg/√Hz is hard to beat. The self-noise for the Trillium 240 is shown in Fig. 2.4. Although portable, at 14 kg for the seismometer only it is a bit bulky. Its hefty price tag also makes it unsuitable for applications in arrays of hundreds of sensors, but the noise performance, especially at lower frequencies, serves as a benchmark for portable seismic sensors.

Large scale seismic reflection surveys usually use vast arrays of geophones. The noise performance of a state-of-the-art battery powered seismic node typically used for these surveys is also shown in Fig. 2.4 [78]. The Innoseis node, marketed as INOVA quantum [82], contains a 5 Hz geophone and reaches a vertical acceleration noise performance of $1 \text{ ng}/\sqrt{\text{Hz}}$ around its natural frequency in the highest gain setting. The geophone is the only sensor in the comparison that uses a velocity transducer, and the superiority of displacement transducers in terms of noise performance at low frequencies is evident. In applications that need both a large number of sensors and good noise performance at low frequencies, MEMS accelerometers are a good candidate.

The MEMS accelerometer with the lowest noise performance commercially available is QuietSeis made by Sercel [80]. It operates as an FBA in a vacuum sealed package with both capacitive position readout and capacitive feedback actuation. Its horizontal noise performance is roughly one order of magnitude above the seismic background requirement for the Einstein Telescope host site, also indicated in Fig. 2.4. Hewlett Packard has developed a MEMS accelerometer with a slightly lower, but similar noise performance of $10 \text{ ng}/\sqrt{\text{Hz}}$ above 5 Hz [83]. That sensor operates as an open-loop accelerometer with capacitive position readout and is encapsulated in vacuum on the wafer scale, but never made it to the market. The best noise performance for a MEMS based seismometer to date comes from a sensor especially designed for the InSight mission to Mars [79]. The Insight spacecraft landed on Mars in 2018, and the MEMS seismometer that was on board is currently taking data there [84]. With almost 1 gram its proof mass is made extraordinarily large by attaching small golden bars to the silicon. Also this sensor operates as an FBA with capacitive readout, but because of the large mass and the limited strength of capacitive actuators, it uses a current driven actuator based on Lorentz forces.

The MEMS seismometer presented in this work has a proof mass of 12.7 mg and therefore its noise performance is no match for the Mars seismometer that uses a proof mass that is two orders of magnitudes larger. On the other hand, the novel anti-spring suspension does allow for a noise performance that beats the best commercially available MEMS sensor by over an order of magnitude over at least three decades of frequency. In Chapter 3 we will introduce the anti-spring suspension that is responsible for the excellent noise performance and discuss its design methodology. In Chapter 4 we then go on to present models for all the relevant noise sources in the MEMS sensor and verify them through measurements, and finally in Chapter 5 we discuss the MEMS fabrication process and related issues.

Miniaturised geometric anti-spring technology

The sensing and attenuation of seismic motion are really just two sides of the same coin. As we have seen, both applications utilise the properties of the same basic structure: a mechanical harmonic oscillator. It is therefore not surprising that there are many parallels in the design principles of vibration attenuation systems and seismic sensors. For example, stiffness reduction techniques built around inverted pendula [29, 64], geometric anti-springs [31, 85] or a LaCoste suspension [65, 86] have been used in both applications to improve performance.

Not all of these techniques are suitable for application in MEMS seismometers. A microscopic inverted pendulum will generally not be heavy enough to compensate the stiffness of the flexures, and the zero-length spring that is required for a LaCoste suspension has not successfully been implemented in MEMS to date. The geometric anti-spring principle, however, *does* scale down to the micro-scale, and in what follows we describe its application in a low-noise MEMS seismometer. First we briefly outline the geometric anti-spring principle, and then present a system overview of its implementation in a seismic sensor discussing the associated challenges and design methodology. To conclude, the challenges related to the design of a sensor that senses vertical acceleration are discussed, and the feasibility of designing such a sensor is shown through measurements on a prototype.

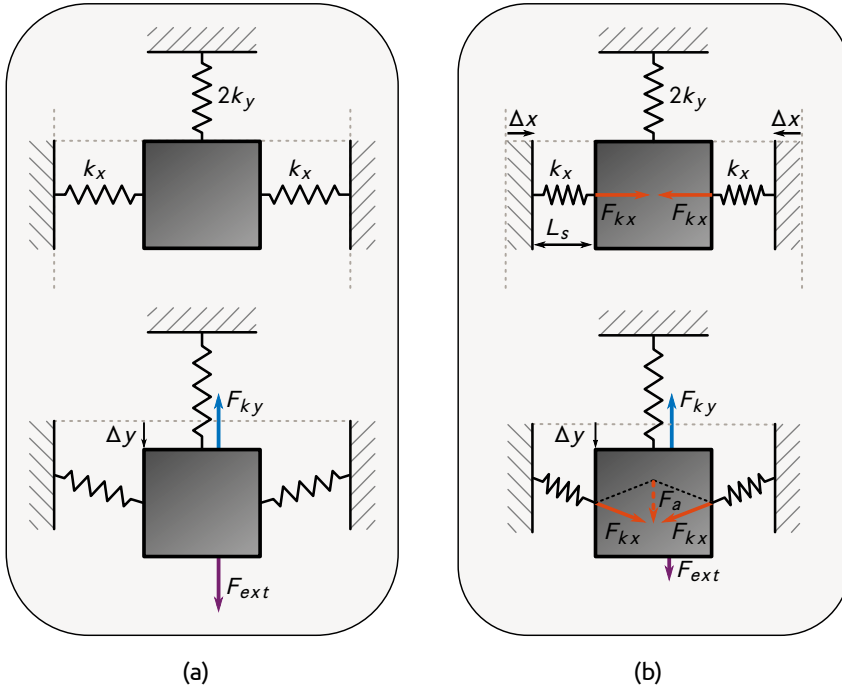


Figure 3.1 – Schematic representation of the geometric anti-spring effect. (a) A body is suspended in the horizontal plane by several uniaxial springs. Displacing this body from its equilibrium position over a distance Δy requires an external force $F_{ext} = -F_{ky}$. When compressing both the springs k_x over a distance Δx such as in (b), the external force required to obtain the same displacement Δy is reduced and consequently the effective stiffness in the y-direction is lowered.

3.1 Geometric anti-spring working principle

Conceptually, the geometric anti-spring effect can be most easily understood by examining the equilibrium of forces on a body that is suspended in the horizontal plane by several uniaxial springs. Figure 3.1 shows such a suspended body with its suspension springs aligned with the x- and y-axes. In practical implementations of geometric anti-spring systems, the stiffness in both the x- and y-directions can often be attributed to the same flexible structures that have a certain stiffness along both axes [87]. In those situations, the spring of stiffness $2k_y$ oriented along the y-direction in Fig. 3.1 does not represent a physical spring, but rather it represents the component of the stiffness in the y-direction of all the flexures in the system combined. Whether this spring is physically there or not, the underlying principle illustrated here remains exactly the same.

Consider the suspended body shown in Fig. 3.1a. Displacing this body away from its equilibrium position over a distance Δy requires an external force of $F_{ext} = 2k_y\Delta y$ that is equal and opposite to the spring's restoring force F_{ky} . In this case, the two springs oriented along the x-direction do not contribute any significant forces as long as the displacement Δy can be considered to be small compared to the length of the springs. This changes when both these springs are compressed over a distance Δx such as shown in Fig. 3.1b. The equilibrium position of the suspended body does not change, because the forces F_{kx} resulting from the compression of the two springs are anti-aligned and cancel each other exactly. What *does* change in comparison to the uncompressed case, however, is the external force required to move the body away from its equilibrium position. In this situation, when the central body is displaced over the same distance Δy as before, the forces generated by the two compressed springs misalign and give rise to an additional force $F_a = 2F_{kx}\Delta y/L_s$ pointing in the same direction as the displacement. Note that as for the restoring force of a regular spring, this force is directly proportional to the displacement Δy . However, F_a points in the direction directly opposite to that of the restoring force of a regular spring such as F_{ky} , and hence we call a structure like this a *geometric anti-spring*.

Because of the additional anti-spring effect introduced to the system, the total effective stiffness changes. Taking the additional force contributed by the compressed springs F_a into account, we get

$$k_{eff} \equiv \frac{F_{ext}}{\Delta y} = -\frac{F_{ky} + F_a}{\Delta y} = 2k_y - \frac{2F_{kx}}{L_s} \quad (3.1)$$

By changing the lateral compression force F_{kx} , the total effective stiffness k_{eff} can, in principle, be tuned down to any value, even a negative one. In practice, the stiffness is tuned as close as possible above zero, such that the system stays in a well-defined equilibrium and can correctly be operated. Since the force F_{kx} is generated by compressing a spring of stiffness k_x , tuning the system's total stiffness comes down to carefully tuning the spring anchor displacement Δx . In the macroscopic vertical geometric anti-spring filters that are used in the auxiliary vibration attenuation systems for the Advanced Virgo gravitational wave detector (shown in Fig. 1.8a), stiffness tuning is done by turning screws that displace the spring anchors. On the micro-scale, however, this step is not at all trivial. A compression system especially designed to accomplish this specific task will be treated in Section 3.2.1, but first we will present an overview of the seismic sensor as a whole.

3.2 Novel MEMS seismometer overview

An overview of the low-noise MEMS seismometer [88] is shown in Fig. 3.2. It is fabricated through the processing sequence as outlined in Section 5.1.1 from a silicon-on-insulator (SOI) wafer with a 50 μm thick device layer, a 4 μm thick buried oxide and a 400 μm thick handle layer. All intricate structures are etched into the device layer and are visible from the top as in Fig. 3.2c. The proof mass is not perforated, because it is released by an etch into the handle layer from the back side of the wafer, leaving 200 μm of the handle layer to increase the total mass. This results in a single-axis device with a 12.7 mg proof mass that can move along the direction indicated in the figure. The mass position is sensed by two sets of variable gap type capacitors as shown in Fig. 3.2e. The capacitance between the stationary and moving sets of capacitor plates changes differentially with the mass position. Comb drive actuator capacitors as shown in Fig. 3.2a can be used to apply forces to the mass in the sensing direction, either for exciting different modes of oscillation, or for applying a feedback force. More details on both these capacitive structures can be found in Figs. 4.8 and 4.10, and an overview of all the important design parameters and their nominal values is provided in Appendix A.2.

The sensor's proof mass is suspended by four sets of curved prismatic beam springs, one set of four springs at each of its corners. As shown in Fig. 3.2b, one side of the springs is connected to the proof mass, while the other side is connected to a linearly guided block. The x-position of this block can be controlled by a special actuation system, compressing the curved springs to a stressed state as shown in Fig. 3.2d. This actuation system in combination with the curved spring shape represents the unique novelty of this sensor and will be detailed upon below.

3.2.1 Spring compression mechanism

The spring elements that are most commonly used in MEMS devices usually consist of some form of prismatic beam that is clamped at both of its ends. In general, the bending stiffness for such a beam depends strongly on its dimensions. For a straight and narrow beam of length L_s that is oriented along the x-axis, the bending stiffness in the y-direction is given by [89]

$$k_y = \frac{12EI_z}{L_s^3}, \quad (3.2)$$

where E is the Young's modulus of the spring material, and I_z denotes the second moment of area of its cross-section with respect to the z-axis. For a

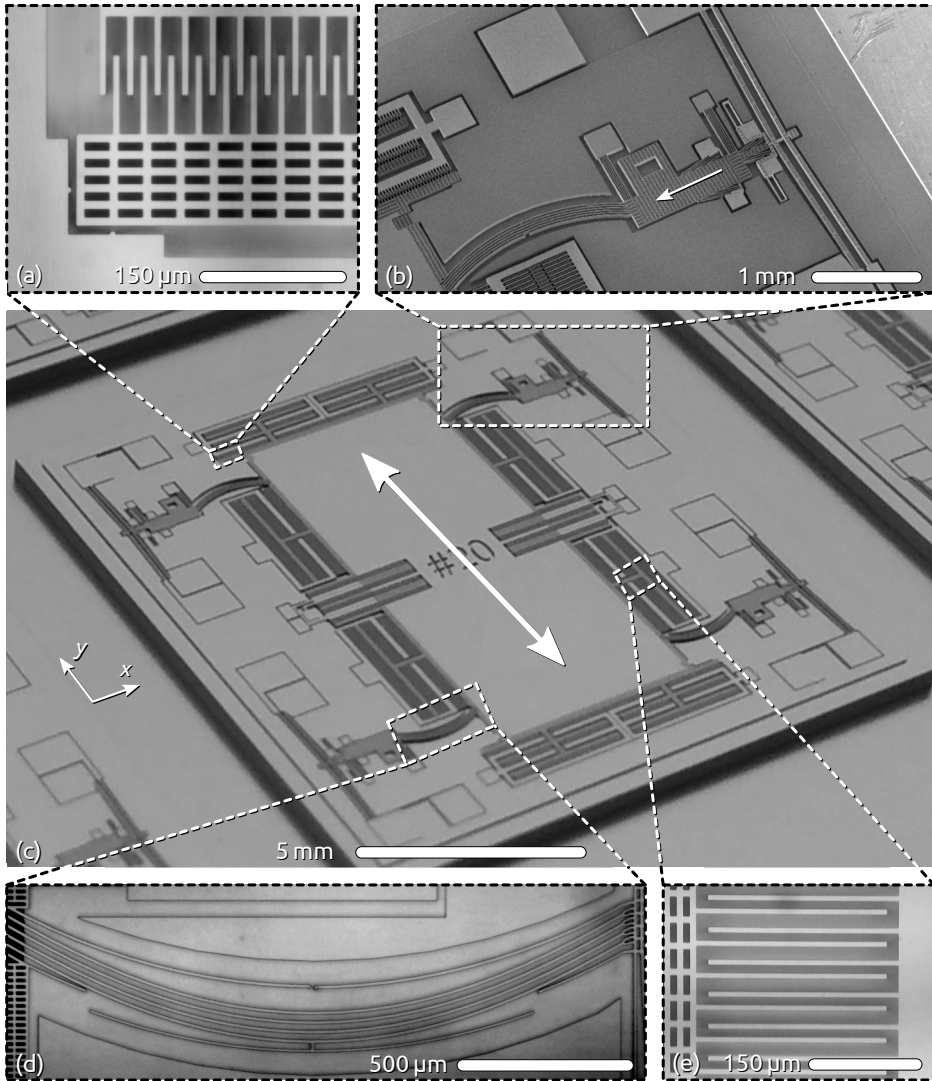


Figure 3.2 – A photograph of a bare sensor die is shown in (c). The central proof mass is suspended by curved prismatic springs (d) at its corners and can move along the y-direction as indicated by the arrow. Variable gap capacitors (e) are used to sense the proof mass position, and variable overlap capacitors (a) are used to actuate on the mass. The suspension springs can be compressed by a thermally actuated mechanism (b) for which a detailed view is provided in Fig. 3.3.

prismatic beam with a rectangular cross-section of thickness t and width w , we have $I_z = tw^3/12$ [90, Tab. A.1]. Combining this expression with Eq. (3.2), we get an expression for the bending stiffness of the spring as a function of its linear dimensions as

$$k_y = \frac{Etw^3}{L_s^3}. \quad (3.3)$$

If this bending stiffness is to be cancelled entirely by the geometric anti-spring effect, we can get an estimate for the compressive force F_{kx} that is required for this by solving Eq. (3.1) for $k_{eff} = 0$ as

$$F_{kx} = k_y L_s = \frac{Etw^3}{L_s^2}. \quad (3.4)$$

The suspension springs in the seismic sensor that is shown in Fig. 3.2 are made out of silicon through a deep reactive ion etching process¹. For the 50 μm thick silicon layer used in this case, the minimum practical spring width w that can be realised is about 6 μm . Similarly, because of out-of-plane stiffness requirements the springs can typically not be much longer than a few millimetres. The Young's modulus of silicon is crystal orientation dependent in general, but if we consider the most common [110] orientation in a (100) wafer, we have $E = 169 \text{ GPa}$ for in-plane deformations [91]. Putting these numbers together means that, according to Eq. (3.4), the compressive force that needs to be generated to cancel the MEMS flexure stiffness is at least of the order of a mN. The electrostatic actuators that are typically used in MEMS devices are not strong enough to generate such a large force, but thermal expansion actuators are [92]. One of the drawbacks of this type of actuators is that they consume a large amount of power, since they need to be Joule-heated to several hundreds of degrees Celsius above the ambient temperature. To avoid burning excessive power during sensor operation, these strong thermal actuators are used only once, just long enough to latch a compliant structure into a fixed position. Now, the compressive force required to reduce the total stiffness through the anti-spring effect has been generated and is maintained mechanically without consuming any power over the entire lifetime of the sensor.

The mechanism designed to achieve this latching action is shown in Fig. 3.3. All visible structures that are either narrow or perforated are released from the substrate below and are able to move, while structures that are not perforated are designed to remain attached to the substrate and serve as mechanical anchors. A pair of V-beam type electrothermal actuators [93] (ETAs) is

¹ Details on the fabrication sequence used to make this device can be found in Section 5.1.1.




Figure 3.3 – (a) Photograph of the anti-spring compression mechanism. The thermal actuators (ETAs) on the right can be used to repeatedly move the bucket to the left and thereby displacing the linearly guided block. This block can be latched into several fixed positions through the anti-reverse pawls indicated by the large dashed box. All the pawls are surrounded by dummy structures that serve no mechanical purpose, but protect these critical parts during the etching process. Figures (b) to (g) correspond to the view inside the large dashed box in (a) for the different available latched positions. A zoom of the anti-reverse teeth in the smaller dashed box is provided in (h).

used to generate the relatively large compression forces required. These actuators consist of long and slender silicon beams that are anchored under a small angle at their ends and are joined at the centre of the actuator. In total, the ETAs are several millimetres wide and they are visible along their full length in Fig. 3.2b. Figure 3.3a only shows their central parts at the right hand side of the photograph. The actuator beams are Joule-heated by feeding a current through them, upon which they expand and their central parts move to the left. Directly attached to the leftmost ETA is the so-called bucket that is linearly guided by a set of folded springs of which the ends are visible both at the top and bottom edges of the photograph. When the bucket is pushed to the left by the ETAs, it engages a linearly guided block that serves as the movable

anchor of the springs that suspend the proof mass. This block is in turn pushed to the left by the bucket, thereby compressing the suspension springs into a stressed state.

The large dashed box in Fig. 3.3a indicates a set of anti-reverse pawls [94] that enables latching the guided block at certain fixed positions such that it can not move back. The initial position of the anti-reverse pawls is shown in Fig. 3.3b. The thermal actuators push the linearly guided block to the left until the pawls reach the state that is shown in Fig. 3.3c, where one pair of pawls gets hooked behind the opposing anchored teeth. Subsequently, the ETAs are retracted, but the guided block will not move back thanks to the anti-reverse pawls engaging their matching teeth on the fixed anchor. Upon retracting the ETAs, another set of anti-reverse pawls located inside the bucket (Fig. 3.3h) moves to the next pair of teeth, thereby extending the total available actuation range of the thermal actuators. With the linearly guided block in its first locked position, the ETAs are heated once more until the block reaches the position that is shown in Fig. 3.3d. Note that compared to Fig. 3.3c, a different set of pawls is latched behind an anchored tooth. Using two sets of pawls enables a step size that is half the minimum tooth size of about $10\text{ }\mu\text{m}$ as defined by the etching process. The first locked position is at $15\text{ }\mu\text{m}$ block displacement and four additional locked positions are available at $5\text{ }\mu\text{m}$ increments that can be reached by repeatedly heating the ETAs for a maximum displacement of $35\text{ }\mu\text{m}$. Displacing the guided block only needs to be done once in all four corners of the device. After the ETAs have latched all four blocks to their desired operation point, the compressive force required for minimising the system's stiffness is now mechanically stored in the springs that are attached to them. After this the ETAs no longer need to be heated, consuming no power during sensor operation.

3.2.2 Electrothermal actuators

The stepped locking system described above allows for displacing the suspension spring anchors up to $35\text{ }\mu\text{m}$ in discrete steps. To be able to actually design the suspension springs that are to be compressed by this system, we also need to know the force that can be generated in this way. To get an estimate for the strength of the thermal actuators, we can use the model developed by Enikov [93]. Consider one half of a thermal actuator with its loads defined as in Fig. 3.4. Through symmetry considerations, the reaction loads at the anchor




Figure 3.4 – Schematic representation of the geometry of a V-beam type electro-thermal actuator (top left). The actuator consists of two beams of length L that are joined at the centre and are anchored under a small angle θ at their far ends. The main drawing shows one half of the actuator with its deflections and reaction loads as induced by an average temperature increase $\overline{\Delta T}$ and an external load F . Implemented values for the length L , the cross-sectional area A , the flexural rigidity EI_z and the angle θ are indicated in the figure. Image adapted from [93].

can be written as

$$\begin{aligned} P_0 &= P \cos(\theta) + \frac{F}{2} \sin(\theta), \\ T_0 &= P \sin(\theta) - \frac{F}{2} \cos(\theta) \quad \text{and} \\ M_0 &= M_1 - v(L)P_0 - LT_0, \end{aligned} \quad (3.5)$$

where F is the external load applied to the centre of the actuator, and P and M_1 denote the loads transferred by the missing half of the actuator. The latter two loads are unknown in principle, but will be fixed by applying additional symmetry conditions. Under the applied thermal and mechanical loads the actuator beam shown in Fig. 3.4 will both be bent away from its unloaded position ($v(x)$), and be stretched or compressed along its length ($u(x)$). The mechanical behaviour of the system is described by two coupled non-linear differential equations as [93]

$$\begin{aligned} \frac{du}{dx} &= \alpha \overline{\Delta T} - \frac{1}{2} \left(\frac{dv}{dx} \right)^2 - \frac{P_0}{EA} \quad \text{and} \\ \frac{d^2v}{dx^2} &= -\frac{T_0}{EI_z} x - \frac{M_0}{EI_z} - \frac{P_0}{EI_z} v, \end{aligned} \quad (3.6)$$

where α is the linear coefficient of thermal expansion of the actuator material, A is the cross-sectional area of the actuator beam and $\overline{\Delta T}$ denotes the average temperature increase of the beam. The linear coefficient of thermal expansion α for silicon is strongly temperature dependent and can be approximated by [95]

$$\alpha(T) = 3.725 \times 10^{-6} \times \left(1 - e^{-5.88 \times 10^{-3}(T-124)}\right) + 5.548 \times 10^{-10} T, \quad (3.7)$$

with T the absolute temperature of the material in Kelvin. Since the full temperature profile along the beam is not modelled here, the average beam temperature is used to estimate α .

Considering M_1 , P and $\overline{\Delta T}$ as unknown parameters in the problem, six independent boundary conditions on the displacements $u(x)$ and $v(x)$ have to be specified to be able to find a solution. Three of them immediately follow from the clamping conditions at the anchor

$$v(0) = u(0) = \left. \frac{dv}{dx} \right|_{x=0} = 0. \quad (3.8)$$

An additional two boundary conditions can be derived from the fact that the actuator as a whole is symmetric, and actually consists of two beams of length L that are joined together at the centre of the actuator. These two conditions can be written as

$$\left. \frac{dv}{dx} \right|_{x=L} = 0 \quad \text{and} \quad u(L) = v(L) \tan(\theta), \quad (3.9)$$

where the first one indicates that the two halves of the actuator are connected under a fixed angle, and the latter reflects the fact that the centre of the actuator can only move in the plane of symmetry. For the last boundary condition, we choose to pick a specific value for the actuator displacement d_{ETA} such that the final condition reads

$$v(L) = d_{ETA} \cos(\theta). \quad (3.10)$$

Numerically solving the boundary value problem in Eq. (3.6) subject to these six boundary conditions will then yield the average temperature increase $\overline{\Delta T}$ that is required to reach a given actuator displacement d_{ETA} at a specific external load F .

Doing this for a range of different loads F and displacements d_{ETA} with the ETA's parameters indicated in Fig. 3.4 gives the result shown in Fig. 3.5. From this we can see that V-beam type thermal actuators can theoretically reach

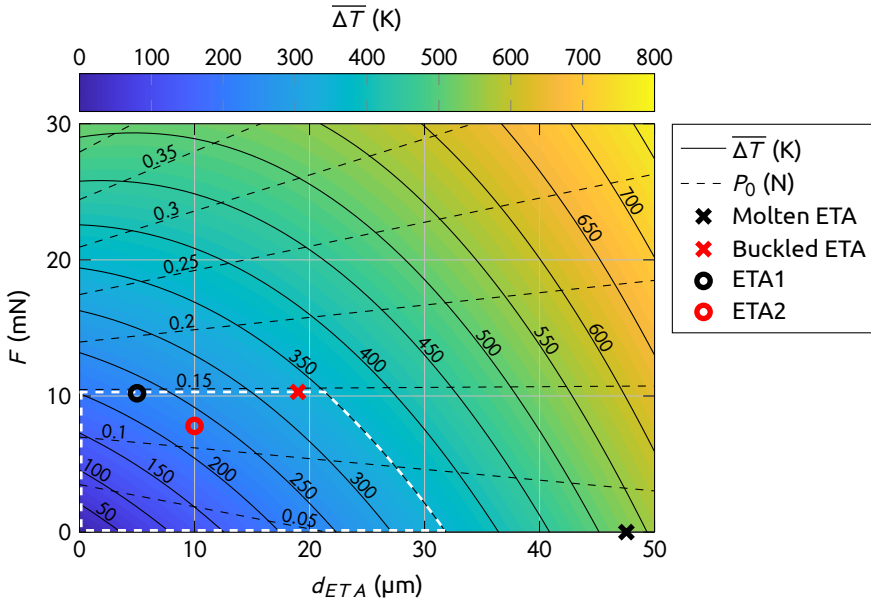


Figure 3.5 – The average temperature increase $\overline{\Delta T}$ required for the thermal actuator shown in Fig. 3.4 to reach a displacement d_{ETA} at a specific external load F . Contours of constant $\overline{\Delta T}$ are shown as solid curves, and contours of constant induced longitudinal reaction load P_0 are shown as dashed curves. Specific empirical limiting values are indicated with markers (see text). Operation the ETA within the region indicated by the white dashed curve is considered safe.

both relatively large strokes as well as large forces, but *not at the same time*. For example, heating the actuator to $\overline{\Delta T} = 350$ K can generate a displacement of just over $30 \mu\text{m}$, but only when no external load is applied. Conversely, a theoretical load of almost 20 mN is predicted, but only at limited displacements. The stepped compression system described in Section 3.2.1 can help overcome this limitation by decoupling the ETA stroke from the total spring anchor displacement. In this way, the thermal actuators can operate with a limited displacement d_{ETA} where the available compression force is large, while the total spring anchor displacement that can be achieved is extended by the latches inside the bucket.

By pushing the actuators beyond their limits, we can experimentally obtain some practical limits to the applicability of the results shown in Fig. 3.5. One actuator was heated under atmospheric conditions way beyond its practically useful limits until it melted and the electrical connection was broken, such as in the situation shown in Fig. 3.6a. Without any externally applied load the actuator reached a displacement of $48 \mu\text{m}$, and just before it was destroyed it sunk




Figure 3.6 – Some examples of ETA damage encountered after operation at too high temperatures. An ETA that has melted over its complete thickness after heating with powers in excess of 1.5 W per actuator is shown in (a). The bottom ETA in (b) shows local melting marks and the bottom ETA in (c) is severely plastically deformed.

a power of about 1.5 W [96]. This working point is indicated in Fig. 3.5 and corresponds to a $\overline{\Delta T}$ of about 530 K. The 530 K contour line therefore represents the ultimate operation limit of this actuator, although in practice operating temperatures should be kept significantly below this. Local melting damage and plastic deformation (Figs. 3.6b and 3.6c) are observed before complete failure. Silicon melts at 1687 K [97, Chap. 1], so a temperature increase of only 530 K should not be enough to melt the actuator. However, an average temperature increase $\overline{\Delta T}$ of 530 K implies a significantly larger increase near the centre of the actuator since the anchored ends remain at ambient temperature [96, 98]. A $\overline{\Delta T}$ of about 350 K gives repeatable results under atmospheric conditions without any noticeable plastic deformation and is considered to be a safe limit.

The result in Fig. 3.5 suggests that at $\overline{\Delta T} = 350$ K a single actuator can generate a force of up to 19 mN. However, the longitudinal load P_0 in these actuators grows quite rapidly with the externally applied load F , which makes them prone to buckling. Once the actuator buckles it loses its strength, so buckling sets an upper limit for F . One of the working points indicated in Fig 3.5 corresponds to the observed onset of buckling at a d_{ETA} of 19 μm . The corresponding load of 10.3 mN was estimated from the nominally designed stiffness of the

actuated spring structure. What causes the buckling is the longitudinal load P_0 for which dashed contour lines are indicated. The onset of buckling corresponds to $P_0 \approx 0.15$ N excluding all working points above this contour, thereby limiting the actuation force to roughly 10 mN. Combined with the maximum $\overline{\Delta T}$ of about 350 K, buckling limits the operating region for the ETA to within the white dashed curve.

To increase the total compression capability, the system contains a second thermal actuator behind the main one as can be seen in Fig. 3.6. The two actuators are not rigidly connected, so the second ETA has to displace about 5 μm more than the main actuator before it starts contributing to the actuation force. Optimal placement of all the anti-reverse pawls allows for locking the last step at a stroke of only 5 μm . The maximum force it can deliver is limited by the $P_0 \approx 0.15$ N contour line to 10.2 mN and is indicated in Fig. 3.5. Assuming the second ETA operates at the same temperature, it contributes an additional force of 7.8 mN for a total of 18 mN. Approximately 4 mN is required to displace all the structures visible in Fig. 3.3, leaving 14 mN to compress the suspension springs.

3.3 Suspension spring design

Now the capabilities of the compression system that will be used to stress the geometric anti-spring system are known, we can design the spring elements themselves. Recall from Eq. (3.4) that the compression force F_{kx} required to bring the effective stiffness in the y-direction down to zero is simply given by $F_{kx} = k_y L_s$. Since this compressive force is generated by displacing the anchor of the suspension springs over a distance Δx with the latching mechanism presented in Section 3.2.1, we also have $F_{kx} = k_x \Delta x$. If Δx_{min} denotes the minimum anchor displacement that can be latched by the compression mechanism, we have $F_{kx} > k_x \Delta x_{min}$. Combining these expressions directly leads to a requirement on the suspension spring stiffness ratio k_x/k_y as

$$k_y L_s > k_x \Delta x_{min} \quad \Rightarrow \quad \frac{k_x}{k_y} < \frac{L_s}{\Delta x_{min}}. \quad (3.11)$$

The minimum anchor displacement Δx_{min} that can be latched into place by the compression system is limited by the smallest feature size that can be reliably etched during the fabrication of the devices. As can be seen in Fig. 3.3, this sets Δx_{min} to 15 μm in this case, and Eq. (3.11) now tells us that for a $O(1)$ mm long spring, the stiffness ratio k_x/k_y should be smaller than $1000/15 \approx 67$. For a




Figure 3.7 – Geometry of a curved suspension spring. The left side is clamped under an angle θ_0 at the origin, while the right side connects to the proof mass under an angle $-\theta_0$. The proof mass side is loaded with forces F_x and F_y , and a torque T , such that it translates in plane and rotates over an angle α with respect to the unloaded case. The curvilinear coordinate λ is measured from the origin O and runs along the entire spring length from 0 to L .

straight prismatic beam spring, this ratio is typically $O(10^4)$, so that can not be used. Giving such a spring a slightly curved shape drastically reduces the stiffness in the x -direction while minimally affecting k_y , allowing to get the correct stiffness ratio k_x/k_y . In the following, we consider curved springs that describe part of a circle of radius R_0 and explore the parameter space for springs of this type.

3.3.1 Bending equations for curved springs

A schematic representation of one of the curved suspension springs is shown in Fig. 3.7. The spring is clamped on the left under an angle θ_0 and symmetrically connects to the proof mass on the right under an angle $-\theta_0$. The unloaded shape of the spring describes an angle $2\theta_0$ of a circular arc,

$$\theta_i(\lambda) = \theta_0 - \frac{\lambda}{R_0}, \quad (3.12)$$

where R_0 denotes the arc's constant radius of curvature, and λ is the curvilinear coordinate that runs along the neutral line of the spring beam across its entire length from 0 to L . The spring tip is loaded with forces F_x and F_y , and a torque T , such that it translates in plane and rotates over an angle α with respect to the unloaded case.

In order to find the internal bending moment $M(\lambda)$, consider the infinitesimally small part of the beam shown in Fig. 3.8. This small section can be ap-




Figure 3.8 – Free body diagram of an infinitesimally small part of the curved beam spring shown in Fig. 3.7. Such a small part can be considered as a straight line of length $d\lambda$ where the torque at both ends is caused by the internal bending moment $M(\lambda)$.

proximated as a straight line of length $d\lambda$, such that the torque balance reads

$$M(\lambda + d\lambda) - M(\lambda) - F_x d\lambda \sin(\theta) + F_y d\lambda \cos(\theta) = 0. \quad (3.13)$$

This expression can easily be rearranged to find the change of the internal bending moment along the beam,

$$\frac{M(\lambda + d\lambda) - M(\lambda)}{d\lambda} \equiv \frac{dM}{d\lambda} = F_x \sin(\theta) - F_y \cos(\theta). \quad (3.14)$$

The internal bending moment is related to the spring beam's curvature as

$$M = EI_z \frac{d(\theta - \theta_i)}{d\lambda}, \quad (3.15)$$

with EI_z the spring's flexural rigidity. For ease of notation and to make the results generally applicable, we first define the normalised coordinate $p = \lambda/L$, and the normalised loads [99]

$$\tau = \frac{L}{EI_z} T, \quad \text{and} \quad G_i = \frac{L^2}{EI_z} F_i \quad \text{for } i = x, y. \quad (3.16)$$

Using these dimensionless parameters, the differential equation governing the deformed spring shape can be found by combining Eq. (3.14) and Eq. (3.15) as

$$\frac{d^2(\theta - \theta_i)}{dp^2} - G_x \sin(\theta(p)) + G_y \cos(\theta(p)) = 0. \quad (3.17)$$

Because the unstressed beam describes a circular arc, we can use Eq. (3.12) to

show that $d^2\theta_i/dp^2 = 0$ and θ_i drops from the equation completely. This finally leads to the dimensionless non-linear differential equation

$$\frac{d^2\theta}{dp^2} = G_x \sin(\theta(p)) - G_y \cos(\theta(p)). \quad (3.18)$$

A system description in terms of these dimensionless quantities is valuable, as it allows to obtain general results independent of the scale of the beam geometry. These general results can then be mapped to the specific design scale under consideration later, by using the appropriate scaling factors from Eq. (3.16).

As boundary conditions to this problem we would like to use the Cartesian coordinates of the tip of the spring, because those are the parameters we control with the compression system. The normalised Cartesian coordinates are given by the integrals

$$\frac{x(p)}{L} = \int_0^p \cos(\theta(\xi)) d\xi, \quad \text{and} \quad \frac{y(p)}{L} = \int_0^p \sin(\theta(\xi)) d\xi. \quad (3.19)$$

To use the anchor position as a boundary condition to the problem, it is more convenient to express the relations in Eq. (3.19) in their differential form as

$$\frac{1}{L} \frac{dx}{dp} = \cos(\theta(p)), \quad \text{and} \quad \frac{1}{L} \frac{dy}{dp} = \sin(\theta(p)), \quad (3.20)$$

and solve the system of Eq. (3.20) and Eq. (3.18) simultaneously.

Considering G_x and G_y as unknown parameters in the problem, we need six boundary conditions to solve the system of equations. The clamped condition at the anchor gives the first three:

$$\frac{x(0)}{L} = \frac{y(0)}{L} = 0, \quad \text{and} \quad \theta(0) = \theta_0. \quad (3.21)$$

We choose to take the coordinates of the spring tip as an input, such that it provides three additional boundary conditions as

$$\frac{x(1)}{L} = \frac{x_i(1) + \Delta x}{L}, \quad \frac{y(1)}{L} = \frac{y_i(1) + \Delta y}{L}, \quad \text{and} \quad \theta(1) = \alpha - \theta_0, \quad (3.22)$$

where Δx , Δy and α respectively are the tip displacements and rotation that we choose, and $x_i(1)$ and $y_i(1)$ are the initial spring coordinates at the tip as

$$\frac{x_i(1)}{L} = \int_0^1 \cos(\theta_i(\xi)) d\xi, \quad \text{and} \quad \frac{y_i(1)}{L} = \int_0^1 \sin(\theta_i(\xi)) d\xi. \quad (3.23)$$




Figure 3.9 – The dimensionless load G_y as a function of the displacement of the spring tip for $\alpha = 0$ and $2\theta_0 = 59^\circ$. The contour of interest is $G_y = 0$ and is indicated in bold. The orange dashed contour indicates the locus of working points that have vanishing stiffness. The system becomes bistable for $\Delta x < \Delta x^*$. At the working point marked with a cross at $\Delta x/L = -0.020$ and $\Delta y/L = 0$, 95 % of $k_{y,0}$ is cancelled.

3.3.2 Solving the boundary value problem

Numerically solving the non-linear system of differential equations with the boundary conditions as defined by Eqs. (3.18) to (3.22) yields the dimensionless loads G_x , G_y and τ as a function of the spring tip displacement Δx , Δy and its rotation α . Figure 3.9 shows the results for G_y as a function of the tip displacement in the case that the spring tip is not allowed to rotate, *i.e.* $\alpha = 0$. The data shown were calculated for $2\theta_0 = 59^\circ$ for reasons that will become apparent later. Starting at the coordinates $\Delta y/L = \Delta x/L = 0$ and moving left along the x-axis, notice that all G_y contour lines move away from the axis, indicating that the stiffness in the y -direction decreases as the spring is compressed. Because the system is fully symmetric, we happen to be following the $G_y = 0$ contour here, indicated as the bold curve in the figure. For compressions beyond the critical compression $\Delta x^*/L$ the contour bifurcates into two stable branches and an unstable one at $\Delta y/L = 0$ indicating bistability for $\Delta x < \Delta x^*$.




Figure 3.10 – The stiffness in the y-direction normalised to the dimensionless stiffness at $\Delta x/L = \Delta y/L = 0$ ($k_{y,0} = 12.66$) corresponding to the data shown in Fig. 3.9. The stiffness k_y vanishes for all the working points on the $k_y/k_{y,0} = 0$ contour that coincides with the locus of vertical slopes in Fig. 3.9 indicated by the dashed orange curve. At the working point marked with a cross at $\Delta x/L = -0.020$ and $\Delta y/L = 0$, 95 % of $k_{y,0}$ is cancelled.

In the diagram in Fig. 3.9, a vertical slope on one of the contour lines indicates a locally vanishing y-stiffness. The locus of points where this happens is indicated with the orange dashed curve. This curve is also drawn in Fig. 3.10, which shows the stiffness in the y-direction corresponding to the data shown in Fig. 3.9. To get an idea of the relative stiffness reduction at different tip displacements, the stiffness data are normalised to the dimensionless stiffness in the unloaded situation, *i.e.* at $\Delta x/L = \Delta y/L = 0$. The dashed orange curve indeed coincides with the $k_y/k_{y,0} = 0$ contour in Fig. 3.10 that indicates all working points with vanishing stiffness. All working points to the right of this contour have a positive stiffness and are therefore stable, while all points to the left of it are unstable. In practice, even though stable working points can be found for $\Delta x < \Delta x^*$, it is better to choose one with $\Delta x > \Delta x^*$ to avoid problems related to the bistability of the system. In Fig. 3.9, the dashed curve that indic-




Figure 3.11 – Values for (a) $\Delta x^*/L$ and (b) G_x^* corresponding to the critical working point for a symmetrically curved beam spring, obtained by finding the bifurcation point as indicated in Fig. 3.9 as a function of $2\theta_0$. The corresponding bifurcation point from Fig. 3.9 is indicated in the figures.

ates all points with vanishing stiffness intersects the $G_y = 0$ contour exactly at the bifurcation, making this the most interesting working point. Approaching the bifurcation from the right, we still have only one stable solution, but the stiffness can be made arbitrarily small in principle. In practice, we choose a working point at which 95 % of the initial stiffness is cancelled. This working point is marked with a cross in both Fig. 3.9 and Fig. 3.10.

3.3.3 Choosing the spring parameters

The results shown in Figs. 3.9 and 3.10 were generated for $2\theta_0 = 59^\circ$. For this specific value of θ_0 , the compression distance that cancels the y-stiffness completely was found to be $\Delta x^*/L = -0.0218$ for $G_y = 0$. In general, such a critical compression exists for a range of different values of θ_0 , and can be found for any of them by repeating the procedure in Section 3.3.2. Doing this results in the values for $\Delta x^*/L$ and the corresponding compression load G_x^* as a function of $2\theta_0$ as shown in Fig. 3.11. The coordinate $\Delta y^*/L$ is zero in all cases due to the symmetry of the system. Given a value for θ_0 , the working point with vanishing k_y and the corresponding normalised compression load G_x^* can simply be read from these plots. The corresponding bifurcation point from Fig. 3.9 is indicated in the figure. Note that these results were obtained without knowing any properties or dimensions of the springs themselves. The critical points are determined by the initial angle θ_0 only. As we will see, this makes this type of anti-spring structure highly robust against processing tolerances.

To design the suspension springs, we should choose specific values for their length L and initial angle θ_0 . Due to constraints of different origin, not all combinations of L and θ_0 will work. To begin with, the result in Fig. 3.11a tells us that for each value of θ_0 the ratio $\Delta x^*/L$ is fixed. Since the compression system can only lock the spring in place for $15 \mu\text{m} < \Delta x^* < 35 \mu\text{m}$, this immediately results in both an upper and a lower limit on L that varies as a function of θ_0 . This is visualised in the parameter space shown in Fig. 3.12a, where the shaded region in the plot indicates the parameter combinations that have been excluded by this constraint only. This shows that a significant part of the available design space can already be excluded without any assumptions on the spring beam rigidity EI_z .

Additional constraints come from the maximum force available from the compression mechanism, the maximum allowable stress in the springs, and the out-of-plane sag induced by gravity. To map those constraints onto the parameter space shown in Fig. 3.12, however, additional assumptions on the spring rigidity EI_z have to be made first. The spring thickness is set by the device layer of the SOI wafer used to fabricate the devices, which in this case is $50 \mu\text{m}$, and since the springs are made out of silicon, we have $E \approx 169 \text{ GPa}$. To obtain a low stiffness, a small spring width w is beneficial. In principle, the smallest width that can be obtained in the fabrication process is $6 \mu\text{m}$, but to reduce the relative effect of etching tolerances, the spring width was set to a slightly higher value of $8.7 \mu\text{m}$, resulting in $EI_z \approx 4.6 \times 10^{-10} \text{ Nm}^2$.

With EI_z known, the normalisation factors in Eq. (3.16) can be used to convert the normalised compression loads G_x^* in Fig. 3.11b to the force that the compression mechanism needs to deliver for each combination of L and θ_0 . Since we know from Section 3.2.2 that this force is limited to 14 mN , this excludes a new region of the parameter space indicated in Fig. 3.12b. This second constraint excludes all the parameter combinations below the blue line, but only the combinations that were previously still allowed are shaded.

Compressing the springs will induce a stress in the springs that must be maintained at acceptable levels. The maximum stress in the spring will occur in the extreme fibres at the spring ends and is given by $\sigma = M(L)w/2I_z$ [90, Eq. 8.1-1]. Realising that we have $M(L) = EI_z/L \cdot \tau$ and that the dimensionless torque τ is known for all working points in Fig. 3.11, we get an additional constraint on the spring length

$$L > \frac{Ew}{2\sigma_{\max}} \cdot \tau(\theta_0), \quad (3.24)$$

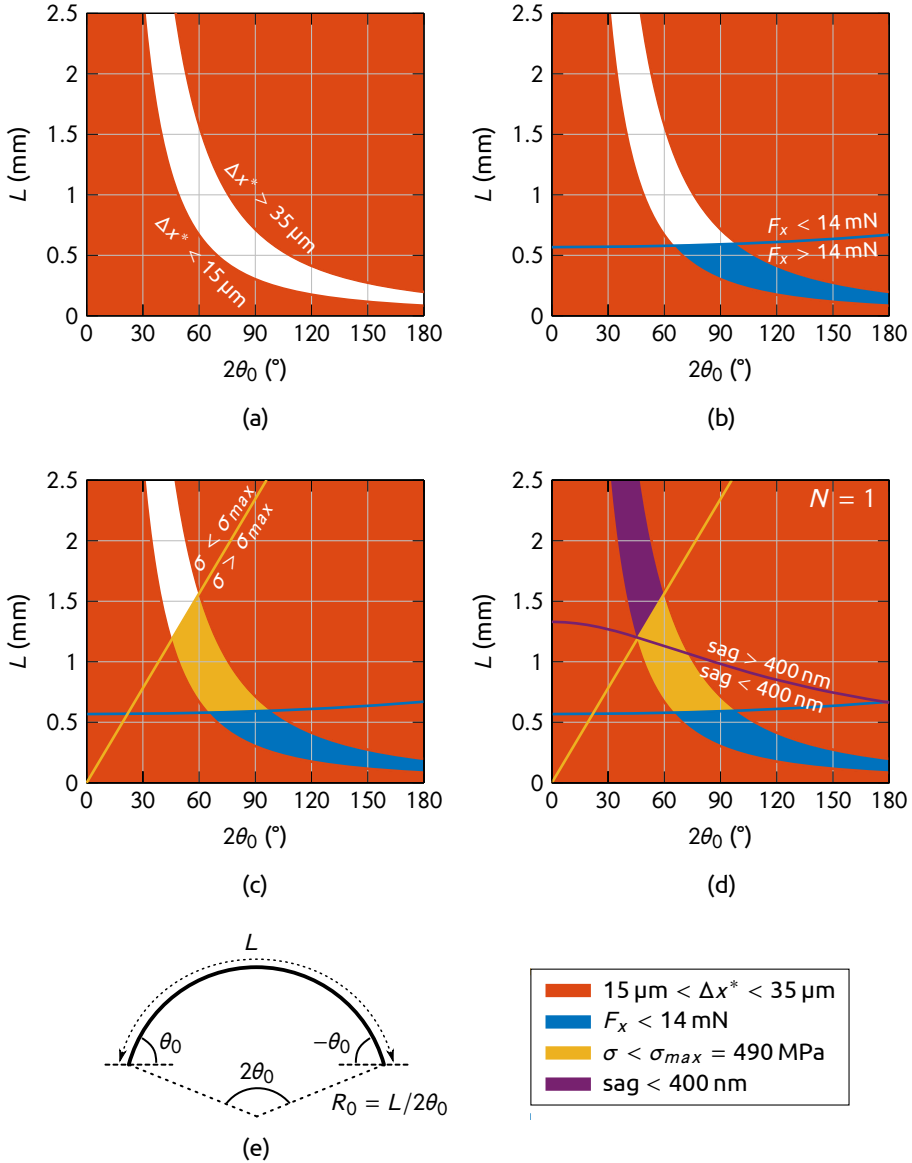


Figure 3.12 – The available parameter space for the curved springs with parameters as defined in (e). New regions of the parameter space excluded by applying the constraints consecutively in the listed order are shaded in the corresponding colour in (a)–(d). The constraints assume that a single curved spring supports a 12.7 mg proof mass at each of its four corners ($N = 1$).

where σ_{max} is the maximum allowable stress in the spring beams. The fracture strength of crystalline silicon is highly dependent on sample dimensions and reported values differ by over an order of magnitude [100]. Here, we use a fracture stress value of 1.47 GPa as reported for structures of comparable cross section as our springs [101] made through a similar fabrication process. To account for the large uncertainty in the strength of silicon, we apply a relatively large safety factor of 3, resulting in $\sigma_{max} \approx 490$ MPa. The lower bound in Eq. (3.24) then excludes yet another large piece of the parameter space as shown in Fig. 3.12c.

The final constraint is set by the fact that our proof mass of 12.7 mg is rather large, and thus gravity will cause it to sag out of plane. If we require this sag to be less than 10 % of the typical gap below the proof mass, or about 400 nm, we immediately obtain the requirement $k_z > 3.1 \times 10^2$ N/m for all the springs combined. This requirement on the out-of-plane stiffness puts an upper bound on the length of the spring beams. The compliance matrix in Appendix A.1 was used to calculate the out of plane stiffness for all combinations of L and θ_0 . The requirement on k_z formulated above then excludes yet another region of the parameter space as indicated in Fig. 3.12d.

The constraints that are outlined above seem to exclude the entire available design parameter space for this miniaturised geometric anti-spring system. However, a closer examination of Fig. 3.12d reveals that the compression mechanism is not optimally used when it compresses a single spring. The entire region of the parameter space that the requirement on the maximum compression force excludes is already unavailable due to the maximum stress allowed in the springs. In other words, when using the actuation mechanism to compress a single curved spring beam in the yellow region, the spring will break before the actuation system reaches its full potential.

To open up some of the available parameter space, we can use multiple springs beams in each corner of the seismometer. Figure 3.13 shows the parameter space subject to the same constraints as outlined above, but for four springs suspending the proof mass in each corner instead of only one. The maximum 14 mN compression force now has to be shared between four springs, raising the lower bound it sets on the spring length. However, since this region was already excluded from a stress point of view, this has no consequences. As the proof mass is now supported by four times as many springs, the out-of-plane stiffness requirement on a single spring relaxes. The upper bound it sets on L is raised, and now there is a small window of parameters that can be used while meeting all constraints.




Figure 3.13 – Constraints on the design parameter space for the curved springs constructed similar to those shown in Fig. 3.12. The constraints assume that four curved springs ($N = 4$) support a 12.7 mg proof mass at each of its corners. The black marker corresponds to $L = 1778 \mu\text{m}$ and $2\theta_0 = 59^\circ$, the spring parameters chosen for the design presented here.

The parameters chosen for this design ($L = 1778 \mu\text{m}$ and $2\theta_0 = 59^\circ$) are indicated with the marker in Fig. 3.13. To fully utilise the flexibility of the compression mechanism as presented in Section 3.2.1, we chose a working point with $\Delta x = 35 \mu\text{m}$ with a maximum spring length as to minimise $k_{y,0}$. The marker falls slightly outside the range of allowed parameters in the figure, because the chosen working point does not correspond exactly to a bifurcation point. Instead, we chose a working point that cancels 95 % of the stiffness, to ensure the stability of the system under process variation. The corresponding working point is also indicated with a marker in Fig. 3.9 and Fig. 3.10.

3.3.4 Stiffness for different working points

To test the spring design, the natural frequency of the sensing mode f_y of the seismometer presented in Fig. 3.2 was measured for all the six different compression states available. The results for three identical devices with the same 12.7 mg nominal mass coming from different locations on different wafers are shown in Fig. 3.14a, along with the behaviour predicted for device #2 from the stiffness data in Fig. 3.10. More details on the numerical data can be found in Section 3.3.5.

For an increasing spring compression the natural frequency of the sensing mode decreases as expected, indicating that the stiffness of the system is drastically reduced in the process. The data for device #2 closely follow the numerical prediction, verifying the design methodology as laid out in this sec-

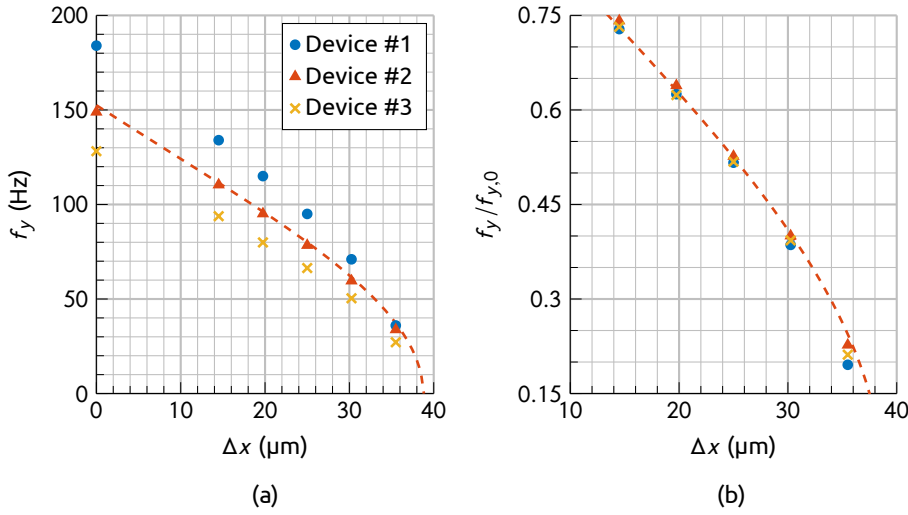


Figure 3.14 – (a) Natural frequency of the sensing mode of three MEMS seismometers of 12.7 mg nominal mass from different wafers as function of suspension spring compression. The dashed curve is generated from the numerical stiffness data in Fig. 3.10 for the spring parameters as measured on device #2 (see Section 3.3.5 for more details). (b) Data as in (a) normalised to the respective natural frequencies in the uncompressed state $f_{y,0}$.

tion. What stands out from Fig. 3.14a, though, is that the spread in the stiffness between different devices is rather large. Assuming the devices have the same mass, the ratio of natural frequencies in the uncompressed state indicates a difference in stiffness of approximately a factor two. This difference is most likely caused by variations in the spring width and the etching profile around their nominal values, which severely impact the total stiffness of the suspension springs. A spread in the etching side wall profiles of only $\pm 1^\circ$ would already cause this factor two change in stiffness and is therefore hard to avoid.

In spite of the large spread in stiffness, the anti-spring system works as anticipated, because the stiffness cancellation is purely a geometrical effect. After all, the stiffness reduction as a function of spring displacement shown in Fig. 3.10 could be calculated without any knowledge of the spring rigidity EI_z , and as such it should therefore not have an impact on the stiffness cancellation². Indeed, when normalising all the data to the uncompressed natural frequency $f_{y,0}$ as shown in Fig. 3.14b, the data for all devices line up and fol-

² However, the spring rigidity can have a slight impact since the implicit assumption made in Eq. (3.15) by moving EI_z outside the differentiation is that EI_z remains constant along the entire length of the spring, something that is only approximately true for a curved beam made out of crystalline silicon since E depends on the crystal orientation.

low the numerical expectation to within 5 % for Δx up to 30 μm . At the last compression step of 35 μm , between 94.8 % and 96.2 % of the initial stiffness is cancelled, raising the seismometer's sensitivity with a factor 20 to 25.

Alternative methods for altering the total system stiffness after device fabrication are commonly implemented in MEMS devices [89]. Typically these implementations use electrostatic effects, exploiting the non-linear dependence of the capacitance between electrodes on their relative position. Biasing such a capacitor to a constant voltage introduces a negative stiffness to the system [102, Sec. 5.3] that can be conveniently tuned by changing this voltage. Typically, however, the negative stiffness introduced in this way changes significantly over the range of motion of the mobile structures in the device. This is largely due to the fact that this range of motion is normally comparable to the separation distance of the capacitive electrodes upon which the introduced negative stiffness strongly depends. For example, the separation between the sensing electrodes in the MEMS sensor presented here is 8 μm , while the full range of motion is $\pm 5 \mu\text{m}$. Because of this variation in stiffness over the range of motion, a so-called *pull-in instability* [103] is likely to occur if a significant part of the mechanical spring stiffness is compensated electrostatically. When cancelling the mechanical stiffness through the geometric anti-spring effect, however, the full range of motion is several orders of magnitude smaller than the total length of the spring ($\pm 5 \mu\text{m}$ vs 1778 μm). This means that the stiffness is practically constant over the full range of motion of the device, and instabilities of this type do not occur. Moreover, choosing a fully mechanical implementation to minimise the system stiffness avoids the additional noise that is usually associated to the electronic circuitry required to bias the capacitor electrodes, especially at low frequencies. A recent MEMS gravimeter that was designed to monitor slow changes in the Earth's gravity also takes advantage of this property of geometric anti-springs [104]. That specific design uses a differently implemented version of the anti-spring system, and uses gravity itself for spring compression, limiting its applicability to vertical sensors only.

3.3.5 Rigid body modes

The previous section shows that compressing the curved spring beams drastically reduces the natural frequency of the y-mode. The system, however, has five more rigid body modes whose frequencies should be kept as high as possible, because they may interfere with the operation of the control loop. In order to measure them, the modes were electrostatically excited using the capacitive structures as shown in Fig. 3.2 at reduced pressures. To be able to





Figure 3.15 – Rigid body modes for all different compression states as measured on a single MEMS device. The solid curves represent the expected behaviour for the in-plane modes as derived from the numerical data presented in Section 3.3.2. The dashed curves represent estimates for the out-of-plane modes based on the compliance matrix presented in Appendix A.1 and the nominal out-of-plane compliance of the compression system.

correctly identify the different modes, all out-of-plane mode frequencies were measured with a laser doppler vibrometer scan and all in-plane mode frequencies were determined with a stroboscopic motion analyser.

Figure 3.15 shows all rigid body modes of the MEMS seismometer as a function of compression state. The data for the y-mode are also shown as device #2 in Fig. 3.14a. For the y-mode, the dimensionless stiffness results as shown in Fig. 3.10 can be used to obtain its expected behaviour as a function of spring compression. Similar dimensionless stiffness data for the other two in-plane modes, x and Tz (rotation around the z-axis), can be generated in a similar way by applying the appropriate boundary conditions for each mode. In order to map these results to mode frequencies, the dimensionless stiffness data then only need to be multiplied with the correct scaling factor EI_z/L^3 .

After inspection with an SEM, the spring side wall profile turned out not to be perfectly perpendicular to the surface. The spring as measured at the top surface was significantly wider than at the bottom. When assuming that the side walls are approximately straight and symmetric, the spring cross-section



(a)

Device parameters			
w_t	8.3 μm	I_z	$2.07 \times 10^{-21} \text{ m}^4$
w_b	7.7 μm	m	12.7 mg
t	48.8 μm	I_{xx}	$4.92 \times 10^{-11} \text{ kgm}^2$
E	156 GPa	I_{yy}	$1.24 \times 10^{-11} \text{ kgm}^2$
G	50.9 GPa	I_{zz}	$6.15 \times 10^{-11} \text{ kgm}^2$

(b)

Figure 3.16 – (a) Schematic representation of a curved spring cross-section (not to scale). (b) Device parameters used to calculate the mode frequencies in Fig. 3.15. The average values for the spring parameters in (a), w_t , w_b and t , are measured with an SEM at different locations on the device, and deviate slightly from the nominal values listed in Appendix. A.2. The inertias (m , I_{xx} , I_{yy} , I_{zz}) are numerically determined from the lithographic mask layout using the nominal process parameters.

looks like the isosceles trapezoid shown in Fig. 3.16a. The relevant second moment of area for this cross-section is [90, Tab. A.1]

$$I_z = \frac{t(w_b + w_t)(w_b^2 + w_t^2)}{48}, \quad (3.25)$$

with parameters as defined in Fig. 3.16a and their measured values in Fig. 3.16b. Because E depends on the crystal orientation, the average value along the length of the beam of 156 GPa was used [91]. The resulting modelled modes are shown as the solid curves in Fig. 3.15. They only take into account the measured spring parameters in Fig. 3.16b and are not fitted to the data. Nevertheless, they match the measured in-plane modes to within a few percent with the exception of the y-mode in the fully compressed state which is about 3.4 Hz below the modelled value of 37.2 Hz. At $\Delta x = 0 \mu\text{m}$, the x- and Tz-modes are significantly lower than expected from the model. This can be understood by realising that in this state the anchored side of the curved springs is still free to move along the x-direction, because the anti-reverse pawls in the compression system have not yet been latched into place. In this situation, it is the compliance of its the compression system that dominates and this is not incorporated in the model. The good match with the rest of the measured modes demonstrates the usefulness of the simple dimensionless model presented in this section for the design of curved geometric anti-springs.

As Fig. 3.15 shows, the frequencies of all the in-plane modes decrease with the compression of the curved springs. Note that, however, only the change

in the y-mode frequency can be attributed to the geometric anti-spring effect. The change in x- and Tz-mode frequencies is predominantly caused by the change in shape the springs undergo during compression. In the fully compressed state, the centre of the suspension spring has displaced approximately 50 μm from its initial circular shape. This is nicely illustrated in Fig. 3.2d that shows the springs in their fully compressed state, while their initial radius can be inferred from the outer dummy structures (see Section 5.1.2) that do not deform when compressing the springs. This large spring deformation presents a challenge for the device fabrication, as both the narrow trenches for the capacitors as well as wide trenches around the curved springs have to be etched. The narrower trenches suffer from RIE lag [105, Chap. 21] causing the critical spring structures to be damaged by overetching. More details on solutions to this can be found in Section 5.1.2.

The three out-of-plane modes shown in Fig. 3.15 are virtually constant over the entire compression range. They can not be obtained from the dimensionless model, since that only includes the in-plane motion of the springs. Because they hardly change, however, they can simply be estimated from the uncompressed circular case. Combining the compliance matrix in Appendix A.1 with the nominal out-of-plane compliance of the compression system yields the dashed horizontal lines in Fig. 3.15. These seem to overestimate the modes by approximately 5 %, which is likely caused by an underestimation of the compression system compliance. No SEM measurements were done on the support springs in this system and therefore their nominal design width was used for estimating its compliance. Since SEM measurements on the main springs revealed that they were narrower than their nominal values, this is likely to be the case for the support springs as well and their calculated stiffness might well be overestimated.

Stiffness ratios

In the fully compressed state the y-mode frequency is 33.8 Hz, so the effective k_y per spring is as small as $4 \times 10^{-2} \text{ N/m}$. If such a stiffness were to be achieved by a straight prismatic spring beam without compression mechanism, the device size quickly becomes impractical at a required spring length of over 5 mm. Moreover, out-of-plane stiffness considerations usually prohibit the use of such long springs. The stiffness ratio R_i can be expressed in terms of the different mode frequencies of the device as

$$R_i \equiv \frac{k_i}{k_y} = \left(\frac{f_i}{f_y} \right)^2 \quad \text{with } i = x, z. \quad (3.26)$$




Figure 3.17 – Stiffness ratios as defined in Eq. (3.26) normalised to the respective stiffness ratio in the uncompressed state. The measured values for R_x are normalised to the numerical expectation at $\Delta x = 0$, since the measured mode there is dominated by the large compliance in the compression system and not the stiffness of interest. The numerical data fit the mode measurements to within a few percent over the rest of the compression range.

The mode measurements in Fig. 3.15 can then be converted to stiffness ratios, yielding the results shown in Fig. 3.17. For a straight prismatic spring of thickness t and width w oriented along the x -axis the ratio R_z is determined by its aspect ratio, because we can write it as

$$R_z \equiv \frac{k_z}{k_y} = \frac{12EI_y/L^3}{12EI_z/L^3} = \frac{I_y}{I_z} = \frac{wt^3/12}{w^3t/12} = \left(\frac{t}{w}\right)^2. \quad (3.27)$$

The suspension spring aspect ratio used for the curved springs would result in $(t/w)^2 \approx 33$. Because the springs are curved, the uncompressed stiffness ratio is actually almost two times lower than for a straight spring at $R_z(0) = 18.0$. However, as the springs get compressed, this ratio improves significantly. In the fully compressed state, at $R_z \approx 347$ the ratio is an order of magnitude higher than what can be achieved with a simple prismatic spring. Using the anti-spring system presented here therefore allows for a relatively compact low stiffness suspension with significantly improved out-of-plane stiffness ratio. The other ratio, R_x , for a straight spring is usually much higher than R_z , limited by the ratio $(L/w)^2$ which is approximately 4×10^4 for the spring dimensions used here. This ratio decreases by almost two orders of magnitude because of the curved spring shape, but is still slightly higher than R_z .

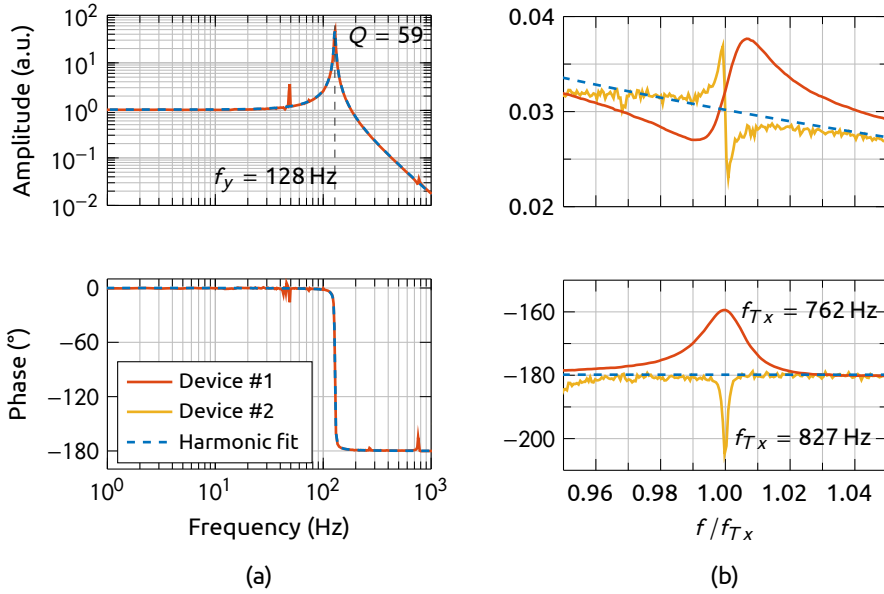


Figure 3.18 – (a) Transfer function measurement for a MEMS seismometer (device #1) at 4 mbar along with a fitted second order transfer function. The structure around 50 Hz is an artefact from the vacuum tank and the Tx-mode is visible around 762 Hz. A zoom of (a) around the Tx-mode is shown in (b) with the frequency axis normalised by $f_{Tx} = 762$ Hz. A measurement of the transfer function of a device with mirrored sensing capacitors (device #2) at 1 mbar shows an opposite phase rotation at the Tx-mode.

Mechanical transfer function

Maintaining high stiffness ratios is important in order to keep the frequencies of the modes aside from the main sensing mode as high as possible. In theory, because of the symmetry of the capacitive sensing and actuation structures, these modes should not couple to the electronic readout system. However, inevitable asymmetries introduced by processing tolerances, residual stress and sagging under gravity result in them showing up in the transfer function. The mode that couples most strongly to a differential change in both the actuation and sensing capacitors is the Tx-mode, because of the way the capacitive electrodes are laid out. Any out-of-plane offset of the proof mass causes an asymmetry that introduces a coupling to the Tx-mode only, while the other modes are unaffected.

As an example Fig. 3.18a shows a transfer function measured on an uncompressed MEMS seismometer with a 12.7 mg proof mass at a pressure of 4 mbar. The main sensing mode with $f_y = 128$ Hz is clearly visible and the transfer seems to be accurately described by a second order function. There are 2 nar-

row frequency bands in which the measured transfer deviates from the second order fit. Just below 50 Hz there is a mode of the vacuum tank that shows up in the data but is not part of the MEMS transfer function, and at 762 Hz the Tx-mode is visible. In Fig. 3.18b, where the frequency region around f_{Tx} is shown in some more detail, we can see that for this device it manifests itself as a relatively benign notch-peak structure. Such a structure has an associated *positive phase change* that can be easily handled when designing the control loop for this sensor, even when it is operated at significantly lower pressures and the gain peak increases accordingly. It turns out that this mode has to be taken into account in the design of the capacitor geometry, though. An older, virtually identical sensor prototype used a mirrored set of sensing capacitors. These work equally well for sensing the proof mass position, but in this case the Tx-mode shows up as a peak-notch structure, also indicated in Fig. 3.18b. Its associated *negative phase change* is much harder to deal with in the design of the control loop, and it made closed-loop operation impossible at decent vacuum levels for that particular sensor design.

3.4 Vertical sensing

The anti-spring implementation presented up to this point has focused primarily on a seismic sensor designed to sense acceleration in the horizontal plane. Sensing the vertical acceleration component of a seismic field is inherently much more difficult to do because of the enormous offset introduced by gravity. Including this offset as part of the output signal is not an option in this case, as the required readout dynamic range would be in excess of 150 dB at a noise level of $1 \text{ ng}/\sqrt{\text{Hz}}$ in a 100 Hz bandwidth.

One way around this is to allow the proof mass to sag under gravity over a distance g/ω_0^2 and make the readout structure periodic, such that it can accurately measure displacements around multiple equivalent working points [106, 107]. Because of the coplanar capacitor structures available in our fabrication process, the proof mass motion is restricted to $\pm 5 \mu\text{m}$. Since the sag under gravity is roughly $400 \mu\text{m}$ for a sensor with a natural frequency of 25 Hz, we can not allow the mass to simply sag under gravity. Instead, the proof mass has to be kept centred between the electrodes where the differential capacitance is nulled. To achieve this, a constant force should be generated that counteracts the effect of gravity on the proof mass. This can not be done by the feedback actuators used for the control loop because of the excessive dynamic range mentioned above. Instead, the compression force in the anti-spring system can be used to lift the proof mass by introducing a slight asymmetry to the

spring geometry. With gravity mechanically compensated, the readout system can operate as if it was a horizontal sensor with the same performance and requirements.

3.4.1 Asymmetric anti-springs

If the horizontal sensor presented in Fig. 3.2 were to be used for vertical acceleration sensing, each of the 16 suspension springs should carry a weight equivalent to $12.7/16 \approx 0.79$ mg in its compressed state. This can be achieved by rotating all the springs over a small angle ϕ in-plane, intentionally misaligning the compression forces such that the resulting component $G_x \sin(\phi)$ now carries the required mass load. Considering the nominal spring parameters ($w = 8.7 \mu\text{m}$, $t = 50 \mu\text{m}$, $L = 1778 \mu\text{m}$) and the average Young's modulus as in Fig. 3.16b, this would correspond to a dimensionless load $G_y \approx -5.8 \times 10^{-2}$. Because Fig. 3.11b shows that the stiffness along y vanishes for a critical compressive load of $G_x = -10.2$, this immediately leads to a first guess for the angle ϕ of roughly 0.33° .

Such a rotation over an angle ϕ corresponds to a simple coordinate transformation only, and does not alter the boundary conditions on the prismatic spring beams. The numerical data for G_y as presented in Fig. 3.9 along with the corresponding data for G_x can therefore be reused to study the behaviour of the rotated spring. Rotating the data over 0.33° yields a result that is similar to that shown in Fig. 3.9, but there are some subtle differences. Figure 3.19 shows the results for the rotated G_y data, where all quantities (Δx , Δy , G_y) are evaluated in the coordinate system aligned to the proof mass. A bifurcation still exists, but because of the introduced asymmetry, it is no longer on the $G_y = 0$ contour. Instead, it now resides on the contour with $G_y \approx -5.8 \times 10^{-2}$ indicating that as per the initial guess, a spring rotated over 0.33° will indeed carry a dimensionless mass load of $G_y \approx -5.8 \times 10^{-2}$ at the critical point $(\Delta x^*, \Delta y^*)$. Additionally, we now have $\Delta y^* \neq 0$, indicating that ideally the spring should be compressed along the dotted line in Fig. 3.19, at an angle with the x-axis. Note that this angle is not equal to ϕ , but instead is quite a bit larger at 11.9° .

The above shows that in theory, this approach will indeed work to compensate gravity mechanically. Even when choosing a working point where 95 % of the stiffness is cancelled as before, an angle ϕ can be found that provides the correct G_y using a similar approach. The bigger challenge now lies in the fact that due to the introduction of the small rotation of all the springs the symmetry of the system, and with it the immunity to fabrication tolerances is lost completely. More concretely, the fact that the system has to be designed for a




Figure 3.19 – The dimensionless load G_y as a function of the displacement of the spring tip for $\alpha = 0$, $2\theta_0 = 59^\circ$ and $\phi = 0.33^\circ$. The contour containing the bifurcation corresponds to $G_y = -0.058$ and is indicated in bold, while the dashed contour represents $G_y = 0$. Because of the small spring rotation ϕ , the critical point no longer resides on $\Delta y/L = 0$ and the spring is ideally compressed along the dotted line which is under an angle of $\tan^{-1}(\Delta y^*/\Delta x^*) = 11.9^\circ$ with the x-axis.

specific value of $G_y \neq 0$ introduces a direct dependence on the prismatic beam spring rigidity EI_z that was not there in the horizontal case. As witnessed by Fig. 3.14, this is a parameter that is difficult to control precisely because it is heavily influenced by fabrication process tolerances, and the data show that it can vary by at least a factor two. This means that the implementation as presented can not be expected to compensate gravity with an accuracy better than about 50 %, and the readout system still needs to cope with a clipping level of $O(1)$ g, facing the same excessive dynamic range requirement as before.

3.4.2 Parasitic spring for compensating gravity

The geometric anti-spring concept implemented here originates from the vertical seismic filters in the auxiliary suspension systems in the Virgo gravitational wave detector (Fig. 1.8). Since these filters also work in the vertical dir-

action, their pre-stressed blades similarly need to carry the exact weight of the payload. Fabrication tolerances are less of an issue here, but the force suspending the payload does change with temperature. To bring the system back into the reach of the feedback actuators after a change in temperature, a parasitic spring is connected to the filter [31]. Its anchor can be moved by means of a linear actuator, such that the static force it introduces can be tuned to cope with a range of temperature drifts.

Since the relative strength of the electrostatic actuators in the MEMS seismometer presented here is significantly larger than that of the actuators in the seismic filter, temperature fluctuations will not be a major concern. The additional parasitic spring, however, can be added to fine tune the force required to lift the proof mass against gravity and overcome the accuracy limit set by the fabrication tolerances.

Figure 3.20 shows an overview of a prototype implementation of a vertical seismic sensor with an implementation of such a parasitic compensation spring. Most functional structures can be readily recognised from a comparison with the horizontal sensor shown in Fig. 3.2. The proof mass in this case only includes the 50 μm device layer and is significantly smaller at 1.57 mg. The lighter mass allows suspending it by only one curved spring in each corner, resulting in an uncompressed natural frequency of 210 Hz. Following the same approach as above, we find that the suspension springs should all be rotated by a angle $\phi = 0.16^\circ$ to suspend the proof mass in a centred position when they are fully compressed. Instead, the angle was intentionally set to a somewhat smaller value of 0.11° to be sure that the compression force would be too small to fully lift the proof mass. The rest of the required force will then be provided by the compensation spring that is connected to the top of the proof mass as shown in Fig. 3.20.

To not defeat the purpose of the anti-spring system, the parasitic spring can only introduce minimal additional stiffness to the system. Folding several sets of 2.0 mm long and 6 μm wide springs results in a nominal stiffness of approximately 0.11 N/m. The anchored end of the spring connects to a thermally actuated anti-reverse structure similar to the one presented in Fig. 3.3. Because this system does not need to carry significant load, the actuation range could be extended to 75 μm and all 4 pawls are offset with respect to each other, allowing a finer step size of roughly 2.5 μm after an initial step of 15 μm . This actuation range together with the nominal stiffness allows for lifting between approximately 10 % and 50 % of the proof mass weight with a resolution better than 2 %.




Figure 3.20 – (a) Photograph of the prototype vertical sensor. The central 1.57 mg proof mass is suspended by one spring in each corner and can move in the direction indicated by the arrow. Sensing (d) and actuator (e) capacitors are similar to the ones in the horizontal sensor. Each of the springs is rotated over an angle $\phi = 0.11^\circ$ (b) to partly lift the mass against gravity. The compensation spring connected at the top along with its actuation mechanism (c) can be used to provide the rest of the required force.




Figure 3.21 – The additional compensation of gravity on the proof mass by the parasitic spring as a function of the displacement of its anchor Δy_c . All four curved suspension springs are fully compressed which compensates 69 % of gravity. A linear fit through the data provides a value for the parasitic stiffness of 9.0×10^{-2} N/m. In the first step that overcompensates the proof mass weight at 54 μm , 100.8 % of gravity is compensated.

Fully compressing all the curved springs by 35 μm both reduces the sensor's natural frequency from 210 Hz to 58 Hz and compensates 69 % of the proof mass weight. The latter can be determined by rotating the sensor until its DC output is nulled and then measuring its angle to the horizontal plane β . In that orientation, the net force produced by compressing the springs compensates exactly a fraction $\sin(\beta)$ of gravity. The operation of the compensation spring is illustrated in Fig. 3.21. When the compensation spring anchor is not displaced ($\Delta y_c = 0$), we only have the net force generated by compressing the curved springs lifting 69 % of the total weight of the proof mass. By then moving the compensation spring anchor upwards we introduce an additional force such that the sensor output is nulled at an orientation increasingly closer to the vertical axis, indicating that we are getting closer to compensating gravity completely. A linear fit through the data indicates that the actual value for the parasitic stiffness at 9.0×10^{-2} N/m is about 20 % lower than its nominally designed value. Moreover, it provides an estimate for the compensation force at an anchor displacement $\Delta y_c = 54 \mu\text{m}$. For this displacement, gravity is fully compensated, so no angle α exists that nulls the sensor output. We know from the estimate in Fig. 3.21, however, that at this working point gravity is mechanically compensated to within 1 %.

The proof of principle in Fig. 3.21 shows that by putting the curved suspension springs under a small angle ϕ in combination with an additional compensation spring, the offset introduced by gravity can be accurately compensated mechanically. Compensating the offset mechanically offloads this burden from the readout system, avoiding any excessive dynamic range requirements there. Chapter 4 presents a readout system architecture for a horizontal sensor and the performance that can be attained with it. The fact that gravity can be compensated mechanically, however, also means that such an electronic readout system cannot distinguish between a horizontal and a vertical sensor channel, and can be the same for all sensors.

Noise in MEMS seismometers

When designing low-noise sensors, it is important to understand the different noise mechanisms that exist. Especially in highly integrated systems such as MEMS sensors, there are non-trivial interdependencies between different parts of the system that deserve attention. This chapter will start with a description of the techniques that are used to measure the different noise contributions, and then goes on to describe where they originate. As we shall see, the sensor needs to be operated in a vacuum environment to suppress the Brownian noise to sufficiently low levels. As this is likely to be one of the cost drivers for producing a seismic sensor like this, insight into what exactly drives the vacuum requirements is essential. To this end, the chapter concludes with a detailed study of the dependence of Brownian noise on package pressure, gas composition and sensor geometry.

4.1 Noise measurements

When dealing with sensitive inertial sensors, measuring their noise performance is far from trivial. A measurement of the sensor output will just yield a measurement of all the vibrations picked up by the sensor, not giving any handle on its ultimate performance. The nicest way of measuring the self-noise is by removing all the external stimuli. Moving the setup to a seismically quiet environment can help, but for low-noise sensors there is no natural place on Earth that has sufficiently small vibrations to directly see their noise floor.




Figure 4.1 – Three sensors with generally different transfer functions $H_i(s)$ measure a common input signal x and all add their own uncorrelated noise n_i . From the three recorded outputs y_i , the sensor noise can be estimated without knowledge of the sensor transfer functions $H_i(s)$.

Luckily, depending on the frequency band of interest, there are a few different ways to obtain a measurement of the self-noise. All inertial noise spectra in this chapter were obtained by using one of the two following methods.

4.1.1 Vibration attenuator

LIGO and Virgo feature the world's best displacement noise attenuators. A prototype of one of the attenuators designed to suspend the auxiliary optics in the Advanced Virgo gravitational wave detector (as shown in Fig. 1.8) is available at Nikhef as a test platform for inertial sensors. The horizontal ground acceleration in its lab in Amsterdam is $\mathcal{O}(10^{-6})$ g/ $\sqrt{\text{Hz}}$ from 0.1 Hz to 100 Hz. This means that a horizontal attenuation of well over 10^3 is required to suppress the input to a sensor on the attenuator to well below a self-noise of $\mathcal{O}(10^{-9})$ g/ $\sqrt{\text{Hz}}$. The platform's horizontal attenuation reaches 10^4 at about 3 Hz (see Fig. 1.8b), so this platform conveniently enables a direct measurement of such a self-noise for frequencies larger than about 3 Hz.

4.1.2 Three-channel correlation

Frequencies below 3 Hz are not accessible when using the vibration attenuator to measure sensor self-noise. To be able to obtain a measurement of the sensor noise at lower frequencies, a correlation method can be used [108]. Instead of physically removing the input stimuli as with the vibration attenuator, this method uses information from three different sensor channels to remove the common mode signal from the measured spectra through correlation techniques in post-processing.

Figure 4.1 schematically depicts a situation where three sensors with generally different transfer functions $H_1(s)$, $H_2(s)$ and $H_3(s)$ measure a common input signal $x(t)$. The three different sensor output streams $y_i(t)$ ($i = 1, 2, 3$), all contain both a signal component that is convoluted with the impulse response of the sensor, $h_i(t) * x(t)$, and a noise component $n_i(t)$ [108]. When analysing the data in the frequency domain, we can write the output signals as

$$Y_i = X \cdot H_i + N_i. \quad (4.1)$$

The noise that is introduced by one of the sensors can be assumed to be uncorrelated to the other sensor noises as well as the input signal. Mathematically, this notion translates to the conditions

$$N_{ij} \equiv \langle N_i N_j^* \rangle = 0 \quad (\text{for } i \neq j), \quad \text{and} \quad \langle N_i X^* \rangle = 0. \quad (4.2)$$

From the measured data streams $y_i(t)$, we can estimate the six quantities P_{ij} that denote the cross-power spectra between the sensor channels for $i \neq j$, and the channel auto-power spectra for $i = j$ (written as P_{ii}). Using Eq. (4.1), we can write this as

$$P_{ij} \equiv \langle Y_i Y_j^* \rangle = \langle X X^* \rangle H_i H_j^* + \langle X N_j^* \rangle H_i + \langle X^* N_i \rangle H_j^* + N_{ij}. \quad (4.3)$$

The middle two terms vanish through the conditions stated in Eq. (4.2) and the last term is only non-zero for the auto-power spectra ($i = j$). Writing the auto-power spectrum P_{ii} and solving for N_{ii} , we get an estimate for the quantity we are ultimately interested in,

$$N_{ii} = P_{ii} - \langle X X^* \rangle H_i H_i^*. \quad (4.4)$$

An independent estimate for $\langle X X^* \rangle$ can be obtained from the cross-power spectrum between different sensors:

$$\langle X X^* \rangle = \frac{P_{ji}}{H_j H_i^*}. \quad (4.5)$$

Inserting this result into Eq. (4.4), we get an estimate that is independent of the common input signal X ,

$$N_{ii} = P_{ii} - P_{ji} \frac{H_i}{H_j}. \quad (4.6)$$

Finally, the transfer function ratio H_i/H_j can be estimated from the other two cross-spectra

$$\frac{P_{ik}}{P_{jk}} = \frac{\langle XX^* \rangle H_i H_k^*}{\langle XX^* \rangle H_j H_k^*} = \frac{H_i}{H_j}, \quad (4.7)$$

such that we can write an estimate for each of the sensor self-noises as

$$N_{ii} = P_{ii} - P_{ji} \frac{P_{ik}}{P_{jk}} \quad \text{for } i, j, k \in [1, 2, 3] \quad \text{and } i \neq j \neq k. \quad (4.8)$$

This estimate is completely independent of the measured common input signal X , as well as the individual sensor transfer functions H_i . The beneficial consequence of this is that it does not matter what the output quantity and frequency response of the different sensors are: as long as they measure the same input signal, this method will give an estimate for the self-noise of all the sensors. Concretely, this means that a MEMS accelerometer that outputs ground acceleration, can be compared to a broadband seismometer that outputs the ground velocity, and the noise estimation will still work.

Common mode input signals can be strongly suppressed in this way, but signals that are not common between the channels will not be removed and can generally limit the accuracy of the noise estimate. In the case of broadband seismometers, such a non-common signal is easily introduced by misalignment between the sensors. The vertical axes can usually be accurately aligned through a bubble level, but the horizontal axes can not be expected to be aligned better than about one degree. If one of the sensors is misaligned with respect to the others, the North output channel of that sensor will contain some signal from the East channel. In general, this signal will show similar spectral features, but is not correlated to the signal seen in the North axis. After the coherent subtraction, this signal will remain and mask the misaligned sensor's noise floor. Figure 4.2a illustrates this for the North output of one out of three Trillium Compact 120s broadband seismometers. After coherent subtraction, a residual remains that still contains spectral features similar to the measured acceleration spectrum. The signal suppression ratio as plotted in Fig. 4.2b, is limited to about 24 dB. This would correspond to a sensor misalignment of about 3.7° , which is not unrealistic.

Luckily, because the Trillium compact is a triaxial device, the relative orientation between the sensors can be corrected in software. By virtually rotating two sensors with respect to the target channel and optimising for the RMS of the estimated noise in the 0.1 Hz to 3 Hz frequency band, the actual alignment of the sensors can be found. When this is done, the residual spectral




Figure 4.2 – Data taken with three Trillium Compact 120s seismometers at the Heimansgroeve seismic station (HGN) were used to verify the correlation procedure. Panel (a) shows for one of the seismometers the measured seismic spectrum (P_{11}), the noise floor estimate from the raw data ($N_{11,A}$), and the noise floor estimate after virtual sensor alignment ($N_{11,B}$), compared to the specified noise floor from the data sheet [109]. The corresponding common mode suppression ratios (P_{11}/N_{11}) are shown in (b). A physical misalignment of about 3.7° limits the common mode suppression to 24 dB without virtual sensor alignment.

features vanish from the noise estimate in Fig. 4.2a, and the noise estimate corresponds to the Trillium’s expected performance. The signal suppression ratio plotted in Fig. 4.2b increases to a peak value of 48 dB, limited in this case by the noise floor of the Trillium sensor. Out of the 6 correction angles (3 for each of the two virtually rotated sensors), the dominant correction is a 3.9° rotation around the vertical axis, corresponding nicely to the 3.7° misalignment predicted from the 24 dB signal suppression limit.

4.2 Mechanical noise sources

Like the mirror suspensions encountered in Section 1.2.5, also the mechanics in a MEMS seismometer suffer from thermal noise. Explicitly writing the mechanical transfer function to an input force $F(s)$ with all the negative stiffness

contributions lumped into k_n , we get

$$H_m(s) = \frac{X(s)}{F(s)} = \frac{1/m}{s^2 + \omega_m^2 \left(1 - \frac{k_n}{k_m} + i\phi\right) + \frac{\gamma}{m}s}, \quad (4.9)$$

where k_m is the uncompressed mechanical stiffness and ω_m denotes the uncompressed MEMS natural frequency in absence of negative stiffness, $\omega_m^2 = k_m/m$. The loss angle ϕ and damping coefficient γ model the different damping mechanisms, where the first is related to losses inside the spring material and the latter to viscous damping mostly dominated by interaction with gases. Through the fluctuation-dissipation theorem [17], these losses can be associated with mechanical noises in the system. In this case, the mechanical impedance can be obtained from Eq. (4.9) as

$$Z_m(s) = \frac{F(s)}{sX(s)} = ms + \gamma + \frac{m\omega_m^2}{s} \left(1 - \frac{k_n}{k_m} + i\phi\right), \quad (4.10)$$

which results in an equivalent acceleration noise spectrum given by

$$\sqrt{S_{mech}(f)} = \sqrt{\frac{4k_B T}{m^2} \Re(Z_m(f))} = \sqrt{4k_B T \left(\frac{\gamma}{m^2} + \frac{2\pi f_m^2}{mf} \phi\right)}. \quad (4.11)$$

Note that this resulting mechanical noise floor does not depend on the amount of negative stiffness k_n . Lowering the natural frequency by means of the anti-spring system therefore does not help in lowering this noise contribution. For a given set of damping parameters ϕ and γ , the resulting thermal noise in the MEMS mechanics sets the ultimate performance that can be achieved.

4.2.1 Brownian noise

The part of the thermal acceleration noise in Eq. (4.11) associated to the damping coefficient γ is usually called *Brownian noise*. It is a frequency independent noise source that for our MEMS seismometer is usually related to gaseous damping. Because of this, the Brownian noise depends on the pressure of the gas surrounding the proof mass. Expressing γ in terms of a quality factor Q_γ , we get an estimate for the level of the Brownian noise as

$$\sqrt{S_{brown}(f)} = \sqrt{\frac{4k_B T \omega_0}{m Q_\gamma}} \quad \text{with} \quad Q_\gamma = \frac{m\omega_0}{\gamma}. \quad (4.12)$$

Here ω_0 denotes the system natural frequency including all the negative stiffness contributions ($\omega_0^2 = \omega_m^2(1 - k_n/k_m)$) and Q_γ is the quality factor resulting from *viscous damping effects* only. As we shall see in Section 4.5, the gases at reduced pressures hardly behave as viscous fluids at the length scale considered, and the damping is caused by collisions of individual gas molecules with the proof mass. However, in this context the damping is still referred to as 'viscous', because the resulting damping force is proportional to the proof mass velocity and can therefore be parametrised by γ .

In the pressure range where the total damping is dominated by gaseous damping, we have $Q \approx Q_\gamma$, so an estimate for the Brownian noise can immediately be obtained from a measurement of the system quality factor. Note that although Eq. (4.12) seems to reveal a dependence on ω_0 , the Brownian noise does not improve by lowering the natural frequency. As we will see in Section 4.5, the damping coefficient γ is set by a combination of the gas species, the pressure and the MEMS geometry, and does not change with ω_0 . Therefore, the quality factor Q_γ is expected to drop linearly with ω_0 and Eq. (4.12) represents a noise floor that is independent of ω_0 .

Measuring the output spectrum of a horizontal MEMS seismometer on the vibration attenuation platform of Section 4.1.1, the Brownian noise can be directly observed above 3 Hz. Figure 4.3 shows the input referred spectrum of a MEMS seismometer at different pressures. To correct for the frequency response of the open-loop MEMS transfer function, the measured voltage spectrum was divided by a second order transfer with $f_0 = 173.3$ Hz and Q as indicated in Fig. 4.3. A spectrum from an L22E geophone aligned with the MEMS sensor was used to get a calibrated output. The Brownian noise at 300 K estimated from the measured quality factors through Eq. (4.12) is indicated in the figure. At atmospheric pressure and a Q of 5.4, the resulting $52 \text{ ng}/\sqrt{\text{Hz}}$ of Brownian noise is the dominant noise source and limits the sensitivity over the entire measured band up to about 400 Hz. The estimate obtained from the Q measurements corresponds well to the observed noise spectrum.

Lowering the pressure raises the system's Q factor and in effect decreases the Brownian noise, revealing other noise sources. At 10 Hz, for example, a decrease in pressure from 1 mbar to 0.1 mbar barely changes the measured spectrum, as the noise there is now dominated by the resolution of the capacitive readout electronics. Around the resonance at 173.3 Hz the electronic noise is suppressed by the resonant response of the MEMS and the Brownian noise remains dominant. For all pressures, the Brownian noise estimated from the Q measurements agrees with the measured noise spectrum at f_0 to within 10 %.




Figure 4.3 – Noise spectra as measured on the vibration attenuation platform with an integrated capacitive front-end prototype ASIC using an uncompressed MEMS sensor (open-loop) with a 12.7 mg proof mass and a natural frequency of 173.3 Hz. The Brownian noise goes down with decreasing pressure, and away from resonance quickly starts being dominated by the resolution of the readout electronics. The structures in the data at 10^3 mbar are caused by acoustic coupling and only the power line harmonics remain when the pressure is reduced.

The data in Fig. 4.3 show that to be able to obtain a noise floor better than $1 \text{ ng}/\sqrt{\text{Hz}}$, the sensor should be operated at a pressure significantly lower than 10^{-1} mbar. In fact, reducing the Brownian noise is the only reason the sensor should be vacuum packaged at all. Since this could be one of the cost drivers of the fabrication process, it is important to understand how this noise scales both with pressure and with changes in the sensor geometry. A more detailed study on this is presented in Section 4.5.

4.2.2 Internal friction

The second term in the mechanical noise from Eq. (4.11) is related to losses in the material of the suspension system itself. These losses are expected to be much lower than those caused by gaseous damping at typical operating pressures. However, they are still relevant due to the frequency dependent nature

of the resulting thermal noise. The second term of Eq. (4.11) separately reads

$$\sqrt{S_{int}(f)} = \sqrt{\frac{8k_B T \pi f_m^2}{mf}} \phi, \quad (4.13)$$

and while ϕ is generally small away from loss peaks, the $1/\sqrt{f}$ dependence will ensure that the thermal noise induced by internal friction will dominate the Brownian noise at sufficiently low frequencies. This behaviour is analogous to the noise situation often encountered in electronics, and noise induced by internal friction can be considered as a kind of *mechanical flicker noise*.

In the same way as with gaseous damping, a quality factor can be associated to the internal friction losses. Using the transfer function in Eq. (4.9) and setting $\gamma = 0$, we can get an estimate for this quality factor as

$$Q_\phi = \left| \frac{H_m(\omega_0)}{H_m(0)} \right| = \left(\frac{f_0}{f_m} \right)^2 \left| \frac{1 + i(f_m/f_0)^2 \phi}{i\phi} \right| \approx \left(\frac{f_0}{f_m} \right)^2 \frac{1}{\phi}, \quad (4.14)$$

where the last result is valid in the limit $(f_m/f_0)^2 \phi \ll 1$. In absence of negative stiffness we have $f_0 = f_m$ and this reduces to $Q_\phi = 1/\phi$ as expected from this kind of modelling [18].

Obtaining an estimate for this quality factor is hard in general, because even when operating the sensor at such a low pressure that the damping from residual gas in the vacuum tank is expected to be negligible, there can be other viscous effects limiting the quality factor. For example, damping from gas molecules desorbing from the proof mass surface or through coupling to losses in the readout electronics are hard to exclude. These effects would contribute to the Brownian noise, and including their effects in an estimate of ϕ would overestimate the internal friction. Measuring the quality factor of an uncompressed MEMS sensor at 10^{-5} mbar yields a Q of 2.2×10^5 , resulting in a first estimate $\phi \approx 4.5 \times 10^{-6}$. Now this is the correct order of magnitude, but as it turns out it overestimates ϕ by about a factor four.

There is a more accurate way to get an estimate for ϕ that is based on how Q_γ and Q_ϕ scale with f_0 . Inspecting Eq. (4.14) we see that Q_ϕ scales like f_0^2 , while from Eq. (4.12) we get that Q_γ scales like f_0 . By its very nature, this MEMS device was designed to be able to change its natural frequency over a wide range of values. By measuring the quality factor as a function of f_0 , the difference in scaling can be used to disentangle the two contributions. The meas-




Figure 4.4 – Quality factor measurements on a fully compressed MEMS device with a 12.7 mg proof mass, and an uncompressed natural frequency of 185.2 Hz. Measurements were performed on a vibration isolation platform at different pressures, and the natural frequency of the device was changed by tuning the electrostatic negative stiffness introduced to the system by the voltage on the readout capacitors. All three fitted values for ϕ match to within 5 %, and the aggregate value from the three fits is $\phi = (1.18 \pm 0.05) \cdot 10^{-6}$, where the errors represent 95 % confidence bounds.

ured Q will be a combination of Q_γ and Q_ϕ as

$$\frac{1}{Q} = \frac{1}{Q_\gamma} + \frac{1}{Q_\phi} = \frac{\gamma}{2m\pi f_0} + \left(\frac{f_m}{f_0}\right)^2 \phi. \quad (4.15)$$

Slightly rearranging this relation, we get to a somewhat more useful form,

$$\frac{f_0^2}{f_m^2 Q} = \frac{\gamma}{2m\pi f_m^2} f_0 + \phi. \quad (4.16)$$

In Fig. 4.4 the quantity on the left hand side of Eq. (4.16) is plotted as a function of f_0 for Q factors measured from free decays at three different pressures. The MEMS anti-springs were fully compressed, resulting in a natural frequency of 35.7 Hz. This natural frequency was lowered further by introducing additional negative stiffness by means of increasing the driving voltage of the capacitive sensing bridge. The lowest natural frequency that could stably be operated was around 10 Hz. Representing the Q measurements as in Eq. (4.16), an estimate for ϕ can be obtained by fitting a straight line through the data. The intercept at the y-axis then is a measure for ϕ . The estimates from all three




Figure 4.5 – Noise spectrum as measured on the vibration attenuation platform on the same device as in Fig. 4.4. The electronically decreased natural frequency lowers the electronic noise, exposing directly the noise contributed by internal friction in the spring system between 5 Hz and 20 Hz. Below 5 Hz the sensor picks up residual motion from the vibration attenuator. The feature at 9.6 Hz is a residual from the MEMS resonance, and some coupling from the AC power grid is visible at 50 Hz and 100 Hz. The remaining two features at 53 Hz and 84 Hz are most likely mechanical modes of the clamped PCB.

measurement series agree to within 5 %, which means that with this method an accurate estimate for ϕ can be obtained even when Q is not limited by internal friction at all. For example, for the data at 9.5×10^{-3} mbar, the contribution from viscous damping is 3 to 10 times higher than that from internal friction over the entire measurement range. Nevertheless, the estimate for ϕ obtained from these data has a relative 95 % confidence limit of about ± 10 %. This technique for disentangling different types of damping proves to be quite powerful and could be of more general use in studying material losses.

Combining the estimates from all three fit results gives an aggregate value $\phi = (1.18 \pm 0.05) \cdot 10^{-6}$, where the error corresponds to the 95 % confidence limits. Inserting this value into Eq. (4.11) together with the result for γ obtained from the fit at 4.7×10^{-6} mbar in Fig. 4.4, we see that the internal friction noise contribution is higher than the Brownian noise for all frequencies up to about 650 Hz. This means that if we can get the capacitive readout resolution to be good enough, a similar measurement as shown in Fig. 4.3 would reveal the in-

Table 4.1 – Silicon properties at 300 K.

Parameter	Value	Source
E , Young's modulus	169 GPa	[91]
α , linear thermal expansion coefficient	$2.6 \times 10^{-6} \text{ K}^{-1}$	[95]
ρ , density	2330 kg/m^3	[91]
C , specific heat capacity	$7.1 \times 10^2 \text{ J kg}^{-1} \text{ K}^{-1}$	[110]
κ , thermal conductivity	$156 \text{ W m}^{-1} \text{ K}^{-1}$	[111]

ternal friction contribution to the thermal noise directly. Indeed, as shown in Fig. 4.5, the additional natural frequency reduction to 10 Hz has decreased the electronic noise to below the internal friction limit, and between 5 and 20 Hz the sensor's sensitivity is limited by noise resulting from internal friction. Below 5 Hz the sensor picks up residual motion from the vibration isolator, and above 20 Hz the decreasing mechanical response of the MEMS sensor results in the noise from the electronics being the limiting factor.

The internal friction and Brownian noise levels in Fig. 4.5 were plotted using the values for ϕ and γ obtained from the fits in Fig. 4.4 and the electronic noise was assumed to be white at the output. The average noise level between 100 and 200 Hz where the sensitivity is known to be limited by the capacitive readout, was referred to the sensor input to draw the electronic noise curve.

Calibrating the MEMS sensor output in this extremely high gain open-loop configuration could not be done by using the L22E geophone as before, because its output would clip whenever the motion that is required for this was introduced to the suspended platform. Instead, the calibration was done in two steps. Firstly, the actuator capacitors were calibrated in a lower sensitivity setting with $f_0 = 35.7 \text{ Hz}$ by using a motorised tilt stage and by using gravity as a reference. Secondly, the now calibrated actuator was used to inject a small calibration signal in the high sensitivity configuration with $f_0 = 9.6 \text{ Hz}$, resulting in a total output sensitivity of $1.1 \times 10^6 \text{ V/g}$ for $f \ll f_0$. The correspondence between the internal friction noise predicted from the measurement of ϕ and the calibrated spectrum is quite good, providing an independent confirmation that the observed noise indeed results from losses in the spring system. Note that from 3 Hz to 30 Hz, the noise performance is better than $1 \text{ ng}/\sqrt{\text{Hz}}$.

The bulk loss of single crystalline silicon at 300 K is of $O(10^{-8})$ [112], so that is not what limits the quality factors observed here. The fact that the estimated ϕ matches the observed thermal flicker noise floor as seen in Fig. 4.5 also excludes significant recoil losses and energy leaking away through the anchors. These effects would lower the effective Q of the system, but would not intro-




Figure 4.6 – SEM image of the surface that is created as a result of the deep reactive ion etching in the silicon device layer. On top the thermal oxide that is used as a hard mask during the etching process is still present. In the inset, the undulations caused by the consecutive etching cycles are clearly visible.

duce additional thermal noise. Another potentially significant loss mechanism, the so-called *thermoelastic damping*, can be estimated as [110]

$$\phi_{TED}(\omega) = \frac{E\alpha^2 T}{\rho C} \frac{\omega\tau}{1 + (\omega\tau)^2} \quad \text{with} \quad \tau = \frac{\rho C t^2}{\pi^2 \kappa}, \quad (4.17)$$

where t is the thickness of a single spring and the material parameters of silicon are listed in Table 4.1. Since the springs are only $8.6 \mu\text{m}$ wide, the maximum ϕ_{TED} of 1.0×10^{-4} at 300 K is reached only at a frequency of 2.0 MHz. For the frequencies below 100 Hz that are of interest here, $\phi_{TED} < O(10^{-8})$ and is therefore completely negligible.

For our springs, the surface to volume ratio is $O(10^5) \text{ m}$, and surface effects that can normally be neglected in macroscopic systems therefore become important. Quality factor measurements of thin silicon structures that have a similar surface to volume ratio indeed indicate that the dominant losses originate at the surface [113, 114]. They can, amongst other things, be caused by the native oxide present on a silicon surface that has been exposed to air, plasma etching damage, or surface roughness [110, 114]. The springs under consideration here will generally have all three. For example, Fig. 4.6 shows a SEM image of the typical rough side wall that results from deep reactive ion etching. The available data is not sufficient to conclude whether the measured losses indeed originate at the surface, and which of the surface effects would be the root cause. However, the measured value of $\phi = (1.18 \pm 0.05) \cdot 10^{-6}$ is comparable to the surface loss limited results obtained elsewhere [110, 114].

4.3 Electronic noise sources

The mechanical noise described above is associated with the mechanics of the MEMS sensor itself. For a given device geometry and package atmosphere, it sets the ultimate noise floor the sensor can reach. To get a useful signal from the sensor, however, it is embedded in an electrical readout system that injects noise of its own in several locations. It is this electronic noise that is strongly influenced by the mechanical transfer function of the MEMS sensor, and is the main motivation to use the anti-spring system to lower its natural frequency.

4.3.1 Capacitive front-end

The position of the proof-mass is read capacitively. A 100 kHz sine wave generator drives a transformer to get a balanced signal. This signal differentially drives the capacitive bridge formed by the sensing capacitors in the MEMS sensor, as shown in Fig. 4.7. The resulting residual current is sunk by a charge-amplifier that converts it to a voltage. The advantage over a voltage amplifier lies in the fact that the input of a charge-amplifier is a virtual ground node. This prevents signal loss through the parasitic capacitance at the input because it is never charged, and provides the stable reference voltage that the actuator capacitors need. The charge-amplifier feeds an inverting amplifier of which the gain can easily be changed, which in turn drives the input of a lock-in amplifier that synchronously demodulates the 100 kHz carrier. After removal of the high frequency mixing components, this circuit produces a signal proportional to the capacitive mismatch of the sensing bridge. When the proof mass is close to its centred position, this signal is also directly proportional to its displacement, giving us the information we need.

The resolution of this readout system is limited by thermal noise generated in the components of the charge-amplifier and by the parasitic resistance of the MEMS sensor R_p . Although signal loss through C_p is avoided by using charge-based readout, its value *is* important because the charge-amplifier's noise gain is proportional to the total input capacitance, which is dominated by C_p . The total input capacitance from the MEMS sensor and its connections on the PCB is approximately 100 pF with a value for R_p of about 200 Ω , as measured with an LCR meter. Using these values, a SPICE simulation predicts a noise level of 148 nV/ $\sqrt{\text{Hz}}$ at the input of the mixer at 100 kHz. Table 4.2 lists the five biggest noise contributors in the charge-amplifier that account for 98.9 % of its noise power. The 148 nV/ $\sqrt{\text{Hz}}$ around 100 kHz at the input of the lock-in amplifier translates to a white noise level of 7.4 $\mu\text{V}/\sqrt{\text{Hz}}$ around DC at its output. A measurement at that point yields 8.4 $\mu\text{V}/\sqrt{\text{Hz}}$, a difference of only about 14 %.

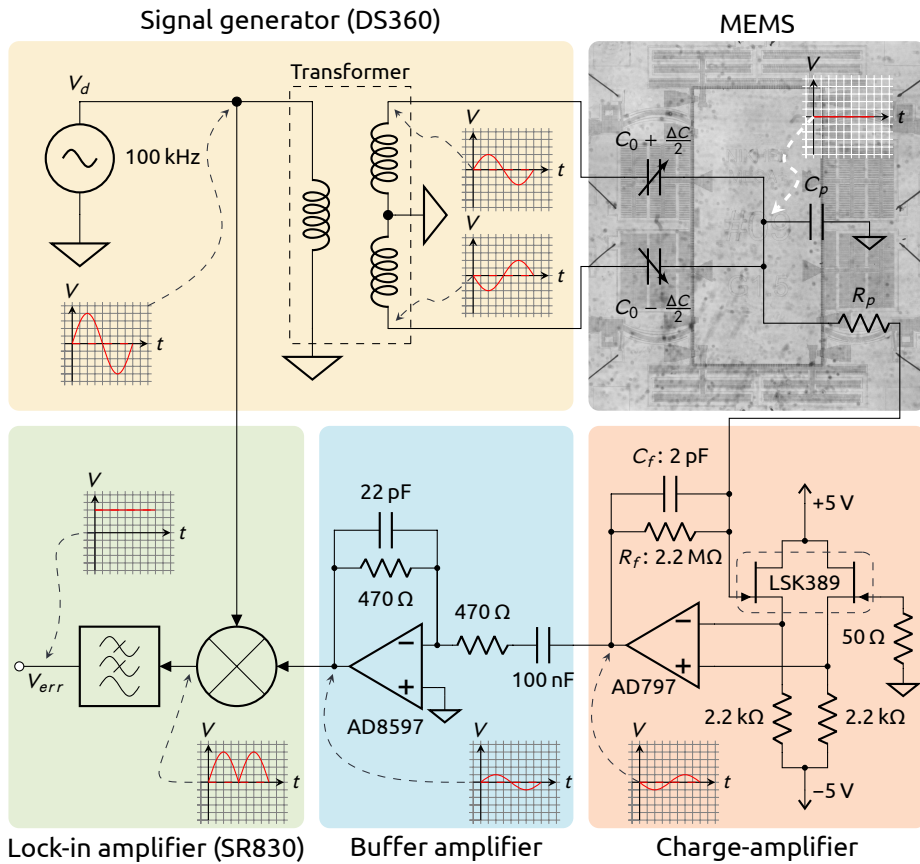



Figure 4.7 – Schematic overview of the electronic system used to read the MEMS proof mass position. A transformer converts externally generated 100 kHz sine to a balanced signal driving the MEMS capacitor bridge. The residual bridge current is integrated by a charge-amplifier, and its output voltage is buffered and sent to an external lock-in amplifier for synchronous demodulation. The circuit produces a low-frequency signal V_{err} proportional to ΔC . The signals at different parts of the circuit are indicated.

Table 4.2 – The five largest noise contributors in the readout circuit shown in Fig. 4.7.

Source	Noise at mixer input @ 100 kHz (nV/√Hz)	Relative contribution (power %)
R_p	86.5	34.0
LSK389	71.9	23.5
R_f	65.7	19.6
AD797	50.8	11.7
50 Ω	47.2	10.1
	148	98.9



(a)

Parameter	Design value
d_0	$8 \mu\text{m}$
d_1	$20 \mu\text{m}$
L_c	$232 \mu\text{m}$
t	$50 \mu\text{m}$
N_c	410

(b)

Figure 4.8 – (a) SEM image of part of the MEMS sensing capacitors. The edge of the proof mass is visible at the top of the image, and can move along the indicated y -direction. The corresponding dimensions and parameters are listed in (b).

The circuit in Fig. 4.7 produces an error voltage that is proportional to the capacitive bridge mismatch ΔC . This voltage is approximately given by

$$V_{err} \approx 50 \cdot \frac{\Delta C}{C_f} \frac{V_d}{2}, \quad (4.18)$$

where the factor 50 is the lock-in amplifier gain. Since the output noise does not depend on V_d , the SNR can be increased by simply raising V_d . Due to the low stiffness nature of our MEMS device, however, V_d is limited to a few volts RMS only, above which the negative stiffness it introduces renders the system unstable. With a typical setting of $V_d = 2.4 V_{\text{RMS}}$, the capacitive bridge resolution can be determined using Eq. (4.18) to be $2.8 \times 10^{-19} \text{ F}/\sqrt{\text{Hz}}$ at a total input capacitance of 100 pF. A recently published integrated capacitance to voltage converter circuit reached a resolution of $4.7 \times 10^{-20} \text{ F}/\sqrt{\text{Hz}}$, about a factor 6 better [115]. However, this performance was measured with a total input capacitance of about 50 pF and a bridge excitation voltage V_d of 5 V. Normalising this result to our sensor parameters would lead to a resolution of $2.0 \times 10^{-19} \text{ F}/\sqrt{\text{Hz}}$, only marginally better than the discrete circuit presented here.

A scanning electron microscope (SEM) image of a section of the position sensing capacitors is shown in Fig. 4.8a. Half of the fingers move with the proof mass in the y -direction, and the other half is fixed to the substrate. When the proof mass moves, the gaps d_0 and d_1 are modulated and the capacitance between the different electrodes changes. Using the infinite parallel plate approximation $C = \epsilon_0 A/d$, the capacitance per side can be estimated as

$$C_s(y) \approx \epsilon_0 L_c t N_c \left(\frac{1}{d_0 - y} + \frac{1}{d_1 + y} \right). \quad (4.19)$$

Using the nominal parameter values from the table in Fig. 4.8b, we can get an estimate for the total sensing capacitance per side when the proof mass is at its centred position (*i.e.* at $y = 0$) as

$$C_0 \equiv C_s(0) = \epsilon_0 L_c t N_c \left(\frac{1}{d_0} + \frac{1}{d_1} \right) \approx 7.4 \text{ pF}. \quad (4.20)$$

By taking the derivative of the expression in Eq. (4.19) with respect to the proof mass position y , the total capacitive bridge sensitivity can be estimated to be

$$2 \frac{\partial C_s}{\partial y} \approx \frac{2C_0}{d_0} \left(1 - \frac{d_0}{d_1} \right) \approx 5.5 \times 10^{-7} \text{ F/m} \quad \text{for } |y| \ll d_0, \quad (4.21)$$

with which we obtain the effective sensing resolution in terms of proof mass displacement as $n_{disp} = 5.0 \times 10^{-13} \text{ m}/\sqrt{\text{Hz}}$.

To refer the electronic noise to the input of the MEMS sensor, the position resolution should be divided by the sensor's mechanical transfer. Doing that, we end up with the effective acceleration noise caused by the charge-amplifier as

$$n_{ca} = \frac{n_{disp}}{|H_{MEMS}(s)|} = n_{disp} \sqrt{(\omega_0^2 - \omega^2)^2 + \frac{\omega_0 \omega}{Q}}. \quad (4.22)$$

It is insightful to inspect this expression in both the high and low frequency limits. For $\omega \ll \omega_0$ Eq. (4.22) reduces to $n_{ca} \approx n_{disp} \cdot \omega_0^2$, so for a given capacitive bridge resolution, lowering the MEMS natural frequency helps to reduce the noise contributed by the charge-amplifier. For $\omega \gg \omega_0$ we have $n_{ca} \approx n_{disp} \cdot \omega^2$. The acceleration noise grows with ω^2 because of the diminishing mechanical response of the MEMS above ω_0 . The high frequency limit is set by the bridge resolution only, and changing the MEMS natural frequency does not help here.

The noise scaling described above can be found in Fig. 4.9, which shows the input referred acceleration noise as measured on the vibration attenuation platform of Section 4.1.1, with the MEMS sensor operating in open-loop. The pressure was lowered to be sure that the sensor's noise floor was not dominated by Brownian noise. As the geometric anti-springs are compressed and the sensor's natural frequency is lowered, the input referred noise below f_0 indeed scales down, while the noise above all the different resonances converges to the same f^2 -sloped curve as expected from Eq. (4.22).





Figure 4.9 – Open-loop noise spectra for three different compression states of a MEMS sensor with a 12.7 mg proof mass and f_m of 128.2 Hz as measured on the vibration attenuation platform. Measurements were performed at about 3×10^{-4} mbar to suppress the Brownian noise. Lowering the sensor’s natural frequency reduces the acceleration noise floor caused by the electronics below f_0 . Above f_0 , the noise floors converge to the same f^2 -sloped curve. Values for f_0 include the negative stiffness introduced by the electronic readout system, which has the most impact on the fully compressed state, lowering f_0 from 27.1 Hz further to 17.1 Hz. All spectra were separately calibrated through comparison with an L22E geophone below 3 Hz and divided by the MEMS transfer function to refer the noise spectra to the input over the entire frequency band. Some residual of the high-Q resonance is visible at f_0 , and the feature at 4.3 Hz is most likely some electronic pick-up from the environment.

4.3.2 Loop-architecture

The MEMS sensor together with the capacitive bridge readout already constitutes a low-noise acceleration sensor. Because of the limited mechanical motion ($\pm 5 \mu\text{m}$) and the poor linearity of the capacitive bridge over this range, the useful dynamic range will be severely limited. Moreover, as the sensor needs to be operated at high Q to push the Brownian noise down, the sensor transfer function is all but flat in the frequency band of interest, and the sensor is constantly ringing at f_0 . Operating the sensor in a closed-loop force-feedback configuration can solve these issues. The control loop can damp the resonance at f_0 without raising the Brownian noise, and both the linearity and dynamic range requirement are moved from the sensing bridge to the feedback actuator.



(a)

Parameter	Design value
d_a	$7 \mu\text{m}$
w_a	$7 \mu\text{m}$
L_a	$40 \mu\text{m}$
t	$50 \mu\text{m}$
N_a	490

(b)

Figure 4.10 – (a) SEM image of part of the MEMS actuation capacitors. The edge of the proof mass is visible at the left of the image, and can move along the indicated y-direction. The corresponding dimensions and parameters are listed in (b).

The feedback force on the MEMS sensor is generated through a second set of capacitors, of which a part is shown in Fig. 4.10a. Again, half the fingers are connected to the proof mass on the left that can move in the y-direction, while the right half of the fingers is fixed to the substrate. This time when the proof mass moves, the overlap between the fingers is modulated, changing the total capacitance. Using the same parallel plate approximation as before together with the parameters listed in the table in Fig. 4.10b, the capacitance per side as a function of proof mass displacement can be estimated as

$$C_a(y) \approx \frac{2\epsilon_0(L_a + y)tN_a}{d_a}, \quad (4.23)$$

which is approximately 2.5 pF when the proof mass is centred (*i.e.* at $y = 0$). Taking the derivative of this expression with respect to the proof mass displacement y gives

$$\frac{\partial C_a}{\partial y} \approx \frac{2\epsilon_0 t N_a}{d_a} \approx 6.2 \times 10^{-8} \text{ F/m}. \quad (4.24)$$

The total potential energy of a system that consists of a capacitor $C(y)$ driven by a voltage source that puts a voltage V across its terminals is given by $U(y) = -C(y)V^2/2$. Because this is a function of y , the applied voltage gives rise to a force between the capacitor electrodes of

$$F_{el} = -\frac{\partial U}{\partial y} = \frac{\partial C}{\partial y} \frac{V^2}{2}. \quad (4.25)$$

Note that because the force depends on V^2 , the sign of F_{el} never changes with V , *i.e.* capacitive actuators can only pull. To overcome this limitation, there are two actuation capacitors, one on each side of the proof mass.

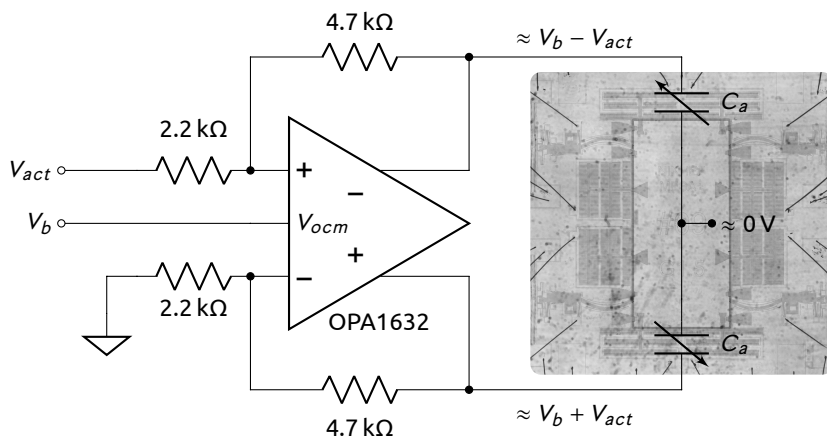


Figure 4.11 – Basic circuit to linearise the MEMS actuation force. The force generated on the proof mass by this circuit is linear in V_{act} . The actuation strength can be set by changing the bias V_b between 0 V and 5 V.

As mentioned above, operating the MEMS sensor in force-feedback mode moves the linearity requirement to the feedback actuator. The result given in Eq. (4.25) seems to be incompatible with that, as the force exerted on the proof mass through an actuation capacitor is quadratic in V , and thus far from linear. Luckily, the force response can readily be linearised by biasing both the actuation capacitors to a voltage V_b . The basic circuit used to do this is shown in Fig. 4.11. The proof mass is kept at 0 V by the charge-amplifier, so the top capacitor will be pulling *up* with $\partial C_a / \partial y \cdot (V_b + V_{act})^2 / 2$, and similarly, the bottom capacitor will be pulling *down* with $\partial C_a / \partial y \cdot (V_b - V_{act})^2 / 2$. This gives a resulting force of

$$F_{fb} = \frac{1}{2} \frac{\partial C_a}{\partial y} \left((V_b + V_{act})^2 - (V_b - V_{act})^2 \right) = 2V_b \frac{\partial C_a}{\partial y} \cdot V_{act}, \quad (4.26)$$

which is linear in the actuation voltage V_{act} supplied to the circuit. Combining this with the estimate from Eq. (4.24), the closed-loop clipping level can be estimated to be ± 0.025 g for a device with a 12.7 mg proof mass with V_b set to its maximum value of 5 V.

Having both an error signal and an actuation port available through the circuits described in Figs. 4.7 and 4.11, the MEMS sensor can be embedded in a control loop. Fig. 4.12 schematically depicts the system architecture. An input acceleration displaces the proof mass, and the capacitive readout bridge, mod-

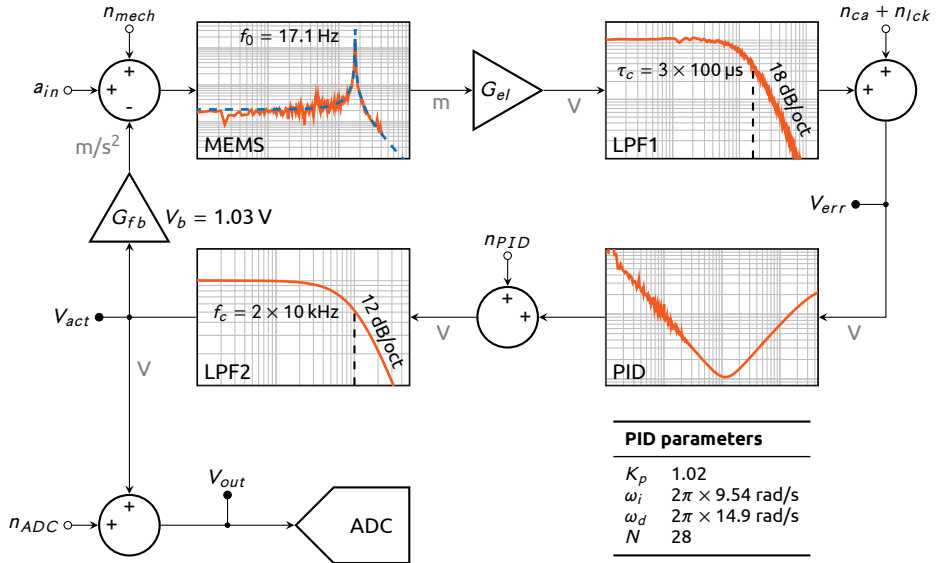


Figure 4.12 – Schematic representation of the control loop around the MEMS sensor. The entire readout circuit from Fig. 4.7 that produces V_{err} is modelled here as a gain G_{el} with a lowpass filter LPF1. The actuation circuit from Fig. 4.11 is represented by the gain G_{fb} . The error voltage drives an analogue PID controller of which the parameters as defined in Eq. (4.27) are listed. After another external lowpass filter LPF2 (SR560), the resulting V_{act} is used to generate a feedback force that keeps the proof mass centred. The voltage V_{act} is used as the sensor output, and is further digitised.

elled by a simple gain G_{el} followed by a lowpass filter¹, converts this to the error voltage V_{err} . The error signal is fed into an analogue PID controller that has a transfer function approximated by

$$H_{PID}(s) \approx K_p \left(1 + \frac{\omega_i}{s} + \frac{s}{\omega_d + s/N} \right), \quad (4.27)$$

with the implemented PID parameters given in Fig. 4.12 as measured from the transfer function in Fig. 4.13. After another lowpass filter, the resulting control voltage is sent to the actuator circuit, creating a force that counteracts the effects of the input acceleration and keeps the proof mass centred. Since the effect of V_{act} is to generate an acceleration equal to that at the input, its value can be directly used as the sensor output, and is recorded by a digitiser.

The transfer functions of the two different halves of the loop between V_{err} and V_{act} are shown in Fig. 4.13. From the data available, the MEMS transfer

¹The lowpass filter LPF1 is internal to the lock-in amplifier used for the capacitive readout, but is explicitly drawn because its phase response has significant impact on the loop transfer.

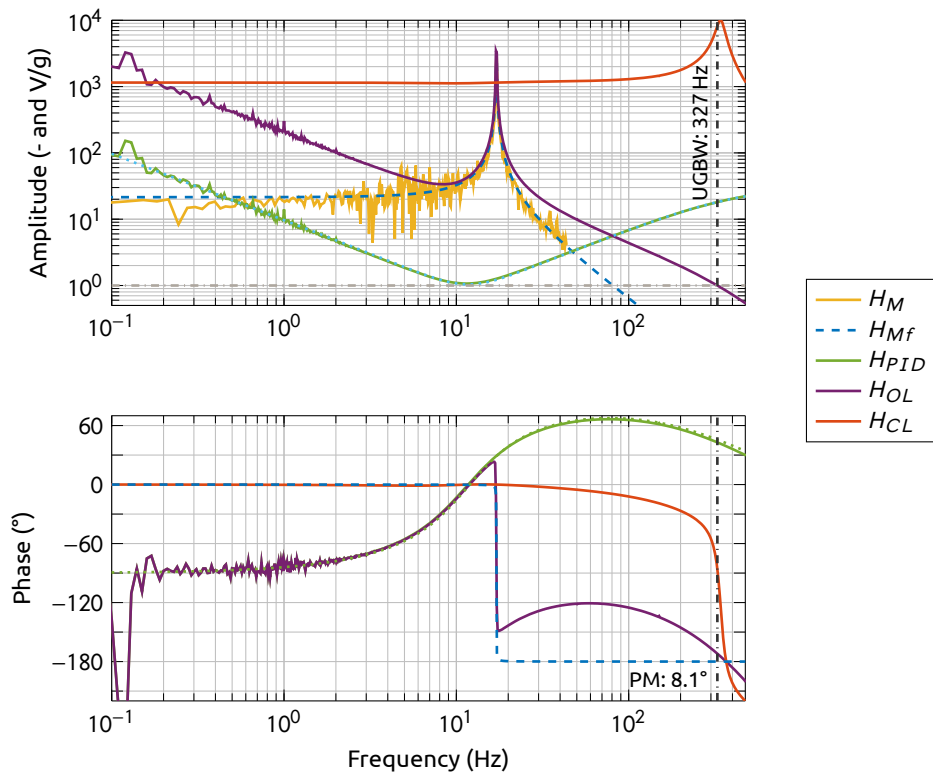


Figure 4.13 – Transfer functions for the control loop shown in Fig. 4.12. H_M denotes the MEMS transfer functions as measured from V_{act} to V_{err} in open-loop, and H_{PID} was measured from V_{err} to V_{act} in closed-loop. H_{Mf} is a second order transfer function fit to H_M , and is used to obtain the open-loop gain, $H_{OL} = H_{Mf} \cdot H_{PID}$. The closed-loop gain H_{CL} is obtained as in Eq. (4.28). The loop gain was optimised for noise performance, resulting in a relatively poor phase margin of 8.1° at a unity gain bandwidth of 327 Hz, with a gain margin of 1.1 dB.

could not be obtained above 40 Hz, so a standard second order function fit to these data, H_{Mf} , was multiplied with the measured PID transfer to calculate the open-loop gain, H_{OL} . The closed-loop gain is then obtained as

$$H_{CL}(s) = 1150 \cdot \frac{H_{OL}(s)}{1 + H_{OL}(s)}, \quad (4.28)$$

where the factor 1150 is a calibration factor obtained from comparison with a signal from a Trillium T120s seismometer. The control loop was heavily optimised for noise performance and not for stability, so the resulting phase margin at 8.1° leads to severe gain peaking around the unity gain frequency of




Figure 4.14 – Voltage noise spectra that enter the control loop in Fig. 4.12. The PID controller noise was obtained from SPICE simulations, the other noise spectra were measured with the inputs terminated with 50 Ω .

327 Hz. Operating the sensor in this way, however, has removed the high Q peak around the MEMS resonance at 17.1 Hz, and for the frequency band of interest to the noise measurement at frequencies below 50 Hz, the closed-loop transfer is flat to better than 5 % with a maximum phase rotation of 5°.

4.3.3 Other electrical noise sources

All the additional electronics in the control loop in principle also add noise to the system. For this specific setup, loop stability requirements dictate that the lock-in amplifier gain be set to a low value, causing it to introduce a noise n_{lock} comparable to that generated by the charge-amplifier. Most noises of electronic origin are modelled to enter the loop in Fig. 4.12 at the output of their contributing block, because for most sources that is where their spectrum can readily be measured. Figure 4.14 shows spectral data for the different noise sources as indicated in Fig. 4.12 that were measured in open-loop. The PID data shown in Fig. 4.14 (n_{PID}) were obtained from a SPICE simulation, since the integrator inside the PID circuitry immediately clips when the control loop is opened, and its noise can therefore not be measured directly.

To get the effective contribution in terms of input acceleration, all noise contributors should first be referred to the output by their individual closed-

loop transfer functions, and the output contributions can then be referred to the input by dividing by H_{CL} . Because of this, all noise sources further downstream become less and less significant, leaving n_{ca} and n_{lck} as the dominant electrical in-loop noise sources. The digitiser noise n_{ADC} is of course the furthest downstream, and furthermore it is over 2 orders of magnitude lower than n_{ca} . However, since it enters outside of the loop, its value at the output is not suppressed by the loop gain and is of similar order as n_{ca} noise at low frequencies. The input referred noise spectra can be found in Fig. 4.18 in the following section.

4.4 Closed-loop system performance

The mechanical noise sources have been accurately measured on the vibration attenuation platform for frequencies above 3 Hz, but the band of interest stretches to well below 100 mHz. To get a measurement of the noise at lower frequencies requires a correlation measurement as described in Section 4.1.2. The vibration isolator lab in Amsterdam has a seismic background spectrum that is about 3 orders of magnitude above the projected noise floor, exceeding the common mode rejection the correlation method can handle. To reduce common mode input to the sensors, the measurement setup was moved to the subterranean lab at Heimansgroeve. This lab, a permanent seismic station operated by the Dutch meteorological institute KNMI, is one of the seismically quietest places in The Netherlands. Its seismic background spectrum is 1 to 2 orders of magnitude lower than that of the lab in Amsterdam, ideal for measuring the low-frequency sensor noise through correlation methods.

4.4.1 Measurement setup at Heimansgroeve

A picture of the measurement setup at Heimansgroeve is shown in Fig. 4.15. The same MEMS sensor as was used for the measurements in Fig. 4.9 is placed inside a portable vacuum tank connected to a small turbomolecular pump. This sensor is measuring acceleration along the North axis in its fully compressed state. The vacuum pump was turned off during all measurements because its vibrations made low-noise measurements impossible. This does mean, however, that during a measurement the pressure inside the vacuum tank slowly rises due to outgassing. This, in turn, limits the maximum measurement time in low-noise condition to $O(1)$ h before the Brownian noise starts masking the internal friction noise. The MEMS sensor was operated in a closed-loop force-feedback configuration with the parameters as described in Fig. 4.12 and 4.13.




Figure 4.15 – Measurement setup at the Heimansgroeve (HGN) seismic station about 10 m below ground. A MEMS sensor with a 12.7 mg proof mass and f_m of 128.2 Hz inside the vacuum tank measures acceleration along the North axis. During measurements, the vacuum pump was turned off, and the Trillium Compact 120s seismometers were covered by their thermal shields. The pressure was measured using a Pfeiffer TPR265 Pirani gauge.




Figure 4.16 – Closed-loop correlation results for a MEMS sensor and one of the Trilliums after virtual sensor alignment. The Trillium and the MEMS measure the same background spectrum between 0.1 Hz and 20 Hz. The estimated Trillium noise corresponds well to its performance as specified in the data sheet [109]. A breakdown of the MEMS noise floor is shown in Fig. 4.18.

Next to the MEMS sensor, there are three Trillium Compact 120s seismometers, two of which are used to do the three-channel correlation with the MEMS output. During measurements, the Trilliums are covered by thermal shields to improve the stability at low frequencies. All signals are acquired by two Nanometrics Centaur digitisers, which have their timing synchronised through GPS. The Trillium channels are virtually aligned to the MEMS channel as described in Section 4.1.2.

4.4.2 Correlation results

The correlation measurement results are shown in Fig. 4.16 for the MEMS channel and one of the Trilliums. The noise estimate for the Trillium seismometer nicely lines up with the noise curve from its data sheet. Between 0.1 Hz and 20 Hz, the MEMS sensor and the Trilliums measure the same background spectrum. Below this frequency range, the two spectra start deviating because the signal drops below the MEMS noise floor. Above 30 Hz, the vacuum tank has several mechanical modes that will be picked up by the MEMS sensor only.




Figure 4.17 – The magnitude-squared coherence between different sensors. The MEMS loses coherence with the Trillium below 0.15 Hz because the background spectrum drops below the MEMS noise floor. The Trilliums remain coherent down to 80 mHz, below which the background spectrum drops below the Trillium noise floor as well. Above 20 Hz, both the Trilliums and the MEMS lose coherence, and the sensors cannot be assumed to be measuring the same signal.

Moreover, at these frequencies, the assumption that the sensors are measuring the same signal breaks down in general. As shown in Fig. 4.17, the coherence between sensors drops sharply above 20 Hz. This is not only true for the coherence between the MEMS and a Trillium, but also between the two Trillium sensors. Noise estimates from this analysis can therefore only be assumed to be accurate for frequencies below 20 Hz. The estimated MEMS noise floor is below $2 \text{ ng}/\sqrt{\text{Hz}}$ between 400 mHz and 20 Hz and remains below $10 \text{ ng}/\sqrt{\text{Hz}}$ down to 50 mHz. To the best of the author's knowledge, such a good noise performance is unprecedented with a proof mass of this order of magnitude.

The MEMS noise floor from Fig. 4.16 is also shown in Fig. 4.18 to compare it to all the modelled noise contributions. During the measurement, the pressure slowly rose from below $5 \times 10^{-4} \text{ mbar}$ to $4.7 \times 10^{-2} \text{ mbar}$ over the course of 2 h. Assuming a linear increase, the average pressure is $1.9 \times 10^{-2} \text{ mbar}$ for the data stretch used in the correlations. This average value is used to estimate the Brownian noise. The internal friction noise is calculated using Eq. (4.13) with the value for ϕ of 1.18×10^{-6} measured before. The charge-amplifier and the lock-in noise is shaped by the MEMS transfer and becomes dominant only above 40 Hz, where the noise estimate is no longer accurate. Mechanical noise dominates in the entire measured band from 10^{-2} Hz to 20 Hz, matching the models put forward in this chapter to better than 20 % from 10^{-1} Hz to 20 Hz.




Figure 4.18 – The MEMS noise floor as obtained from the measurement in Fig. 4.16, compared to the input referred noise entering the system as in Fig. 4.12. The measured noise floor follows the model in the entire measured band below 20 Hz and is completely dominated by mechanical noise due to the aggressive optimisation of the loop to suppress electronic noise. This MEMS sensor has a proof mass of 12.7 mg and an uncompressed natural frequency of 128.2 Hz.

4.5 Modelling the gaseous damping

As seen in Section 4.2.1, the MEMS seismometer should be operated at a pressure well below 10^{-1} mbar to bring the Brownian noise below $1 \text{ ng}/\sqrt{\text{Hz}}$. This section presents a series of Q factor measurements from which a more accurate number for the required pressure can be obtained. Next to this, we will have a closer look at which parts of the device contribute most of the Brownian noise. Usually, electronic noise requirements dictate the design of intricate structures like the actuation and sensing capacitors, but as we will see, their geometry can also significantly impact Brownian noise. A model is constructed from which the Brownian noise as a function of device geometry can be predicted.

4.5.1 Damping of a body in infinite rarefied gas

In a gas, the average distance travelled by individual molecules before colliding with another molecule is given by the mean free path,

$$\lambda_{free} = \frac{k_B T}{\sqrt{2} \pi p d_{mol}^2}, \quad (4.29)$$

where d_{mol} is the kinetic diameter of the individual gas molecules and p denotes the pressure. We already know the MEMS seismometer should be operated at pressures well below 10^{-1} mbar. In this pressure range, the mean free path of nitrogen molecules ($d_{mol} = 364 \text{ pm}$ [116, p. 6]) is expected to be longer than $O(1) \text{ mm}$, much larger than the typical cavity surrounding the proof mass. The *Knudsen number* is defined as the ratio $\text{Kn} = \lambda_{free}/L_c$, with L_c the characteristic length of the flow problem at hand. Situations such as this one with $\text{Kn} > 10$ are said to be in the *molecular flow* regime. In this flow regime, interactions between gas molecules can be neglected. When modelling the interaction between the proof mass and the gas that ultimately leads to damping, we can therefore focus on individual collisions with the proof mass surfaces.

Assuming the velocities of the gas molecules follow a Maxwell-Boltzmann distribution, the force on a unit area of a moving proof mass dA can be calculated from integrating the momentum transferred by molecules colliding with this surface. Assuming the interactions are diffuse (*i.e.* the outgoing velocity vector is uncorrelated to the incoming velocity vector), and the velocity of the proof mass, \dot{y} , is much smaller than the typical velocity of the gas molecules,

Table 4.3 – Relevant surface areas in the MEMS device used to model Q factors.

Parameter	Value	Description
A_x	$6.42 \times 10^{-6} \text{ m}^2$	Total area with unit vector along \hat{x}
A_y	$1.26 \times 10^{-5} \text{ m}^2$	Total area with unit vector along \hat{y}
A_z	$2.50 \times 10^{-5} \text{ m}^2$	Total area with unit vector along \hat{z}
A_8	$9.51 \times 10^{-6} \text{ m}^2$	Proof mass area facing 8 μm gap
A_{20}	$9.51 \times 10^{-6} \text{ m}^2$	Proof mass area facing 20 μm gap
A_{bulk}	$1.23 \times 10^{-6} \text{ m}^2$	Proof mass bulk area facing 52 μm gap

this force can be calculated to be [117]

$$\langle F_{\perp} \rangle = -p \, dA \left(1 + \dot{y}_{\perp} \sqrt{\frac{2m_0}{\pi k_B T}} \left(1 + \frac{\pi}{4} \right) \right), \quad (4.30)$$

where F_{\perp} denotes the force normal to the surface dA , m_0 is the mass of a gas molecule, and \dot{y}_{\perp} is the component of the proof mass velocity normal to the surface dA . The first term is just the force resulting from the average pressure on the unit area, but the second term is proportional to the proof mass velocity and thus corresponds to a viscous drag force. Similarly, a drag force also results from the force parallel to the unit area [117]

$$\langle F_{\parallel} \rangle = -p \, dA \dot{y}_{\parallel} \sqrt{\frac{m_0}{2\pi k_B T}} \quad (4.31)$$

where \dot{y}_{\parallel} denotes the component of the motion parallel to the unit area.

Because in our case all the surfaces are oriented either parallel or perpendicular to the direction of motion, the net force on the proof mass can be found by adding the right contribution (perpendicular or parallel) for all its different surfaces. If we define A_x as the total area of all surfaces that have their normal vector \hat{n} aligned with the \hat{x} unit vector ($\hat{x} \cdot \hat{n} = 1$), and define A_y and A_z similarly, the net force on the proof mass moving along the y-axis is given by

$$\langle F_{net} \rangle = -\dot{y} \frac{8p}{\pi \bar{v}} \left(\frac{A_z + A_x}{2} + A_y \left(1 + \frac{\pi}{4} \right) \right) = -\dot{y} \gamma_f \quad \text{with} \quad \bar{v} = \sqrt{\frac{8k_B T}{\pi m_0}}, \quad (4.32)$$

where a factor 2 was added to account for all the matching surfaces that have their normals aligned with the negative axes. The contribution from the first term in Eq. (4.30) cancels, because it is equal but oppositely oriented between the upwind and downwind surfaces. The resulting force is proportional to

and in antiphase with the proof mass velocity and therefore corresponds to a viscous drag force. Calculating the resulting quality factor $Q = m\omega_0/\gamma_f$ with the parameters in Table 4.3 for nitrogen at 10^{-1} mbar and 300 K we find $Q = 6.3 \times 10^3$ for the device in Fig. 4.3. The value actually measured is about a factor 3.5 lower, at $Q = 1.8 \times 10^3$. This discrepancy is due to the many narrow gaps around the proof mass that contribute additional damping that we have ignored so far. This additional damping actually dominates the free drag described in Eq. (4.32) and will be detailed upon in the next section.

4.5.2 Squeezed-film damping

The many narrow gaps surrounding the proof mass, mainly in the sensing capacitors, contribute additional damping. Because of the finite time it takes for an average molecule to diffuse out of the gap, the proof mass motion modulates the pressure inside the gap. To quantify this effect, we can use a model proposed by Suijlen *et al* [118] that considers the density of molecules in such a gap,

$$n = \frac{N}{V} = \frac{N}{A(d_0 - y)}, \quad (4.33)$$

where N is the total number of molecules present in a gap of total volume V , and d_0 and A are the equilibrium separation and the area of the gap walls, respectively. Because the gap is open, the total number of particles is not conserved, and a change in molecular density can generally be written as

$$\Delta n = \frac{\partial n}{\partial N} \Delta N + \frac{\partial n}{\partial V} \Delta V = \frac{\Delta N}{Ad_0} + \frac{n_0}{d_0} y, \quad (4.34)$$

where n_0 is the equilibrium density of molecules inside the gap. If such a fluctuation in molecule density exists, particles will diffuse into or out of the gap. Defining τ to be the average time it takes for a molecule to diffuse out of the gap, the net number of particles leaving the gap per second is given by

$$\frac{d}{dt}(\Delta N) = -\frac{\Delta N}{\tau}, \quad (4.35)$$

Using this result together with the ideal gas law $p = nk_B T$, we can take the time derivative of Eq (4.34) to write a differential equation for the pressure fluctuation inside the gap, Δp , around the equilibrium pressure, p_0 , as

$$\frac{d}{dt}(\Delta p) = -\frac{\Delta p}{\tau} + \frac{p_0}{d_0} \frac{dy}{dt}. \quad (4.36)$$

This equation is readily solved in the Laplace domain and gives an expression for the squeezed-film damping force on the moving wall as

$$\Delta F(s) = -A\Delta P(s) = -\frac{p_0 A}{d_0} \frac{s\tau}{1 + s\tau} Y(s) \approx -\frac{p_0 A\tau}{d_0} (sY(s)) \quad \text{for } s\tau \ll 1. \quad (4.37)$$

This approximation only holds in the $s\tau \ll 1$ limit, but this is justified as we shall shortly see that for this particular MEMS geometry, τ is of $O(10^{-6})$ s. Equation (4.37) shows that the force acting on the wall due to the modulated pressure in the gap is both proportional to and in antiphase with the wall velocity $sY(s)$, and therefore corresponds to a viscous drag force with a coefficient $\gamma_{sq} = p_0 A\tau/d_0$. All surface areas in the MEMS sensor and their corresponding gap widths are known from the mask design and can be found listed in Table 4.3. What is left to get an estimate for γ_{sq} is to find the correct value for τ .

MolFlow+ simulations

In the molecular flow regime all interactions between different gas molecules can be neglected, hence problems in this regime are well suited for Monte-Carlo simulation. A simulator tool developed at CERN, *Molflow+* [119], is available that makes these simulations quick and convenient. Using this simulation tool, the diffusion time τ can readily be obtained for different geometries.

Simulations were set up for a single gap as seen in the sensing capacitor, and the narrowest gap after that which occurs between the silicon bulk part of the proof mass and the device's handle layer. A schematic top view of the simulated sensing capacitor geometry is shown in Fig. 4.19b, where particles are emitted from the bottom plate (blue dots). The molecules are traced along their random walk bouncing from all the different surfaces until they leave the cavity (red dots). All surfaces were set to reflect diffusely with full thermal accommodation, compatible with the model used [118]. A molecule was considered to have left the gap if it either moved out of the gap from the top, or moved around or underneath the capacitor plate. As can be seen in Fig. 4.19b, most particles leave the gap through the top opening.

A distribution for the flight time of all particles is recorded, from which the average diffusion time τ can be calculated. Simulations were performed for a range of gap widths, both in air ($m_0 = 28.97$ u) and helium ($m_0 = 4.00$ u) at 300 K, the results are summarised in Fig. 4.19a. Each data point was obtained from a Monte-Carlo simulation with $O(10^6 - 10^7)$ simulated particles.

Since the expected path a gas molecule takes does not depend on the gas species, the average flight distance for all particles is the same. The diffusion times for air and helium are therefore expected to differ only due to their dif-

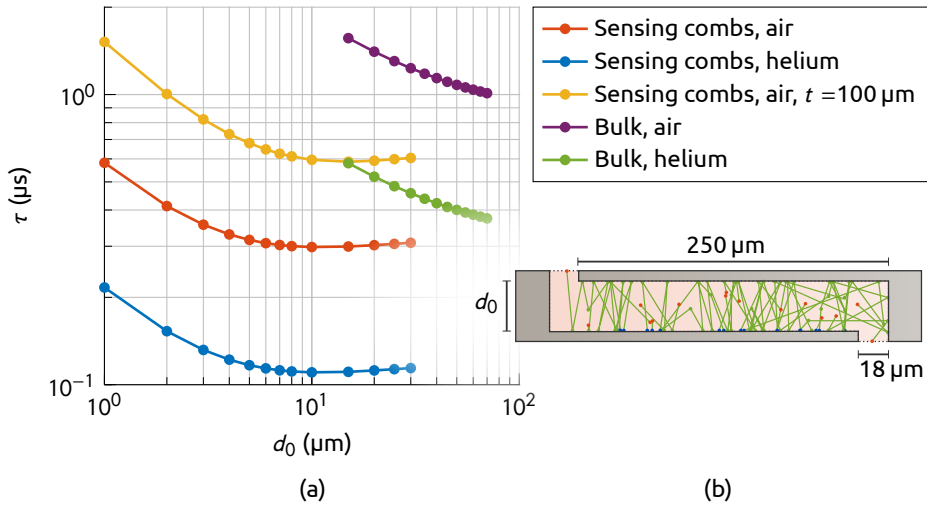


Figure 4.19 – (a) MolFlow+ simulation results for the diffusion time τ as a function of gap size d_0 for air and helium molecules at 300 K, escaping from different regions of the MEMS device. A top view of a typical simulation result inside a capacitive sensing gap is shown in (b). Particles spawned at one of the capacitive plate surfaces (blue dots) and are tracked until they leave the cavity (red dots). The average τ in (a) can be calculated from the stored flight time distribution.

ference in average particle velocity \bar{v} as defined in Eq. (4.32). The ratio of these average velocities is simply given by $\sqrt{m_{air}/m_{He}} \approx 2.69$, and indeed the results shown in Fig. 4.19a for helium and air differ by this factor to well within 1 %.

4.5.3 Agreement with quality factor measurements

Combining the contributions from free drag in Eq. (4.32) and squeezed-film damping in Eq. (4.37) gives an estimate for the total viscous damping in the MEMS device as a function of pressure. To see how well this predicts the damping in an actual MEMS sensor, this estimate is plotted in Fig. 4.20 along with measured Q values of an uncompressed MEMS device with $f_0 = 183.3$ Hz, both in air and helium atmospheres at different pressures. Q values were obtained from measuring free exponential decays, and pressures were measured with two capacitive membrane gauges (10^{-5} mbar to 10^{-1} mbar and 10^{-1} mbar to 10^3 mbar) to minimise the gas species dependence of the measurement.

At atmospheric pressure, we have $\text{Kn} \ll 10$ and the molecular flow assumption used to construct the damping model is invalid. Lowering the pressure to below $O(1)$ mbar, we have $\text{Kn} > 10$, and the measured Q values converge to the




Figure 4.20 – Q factors for an uncompressed MEMS device with a proof mass of 12.7 mg and a natural frequency of 183.3 Hz as measured on the vibration attenuation platform in air and helium atmospheres at different pressures. The expected Q factor from the modelled free drag (γ_f) and squeezed-film damping (γ_{sq}) using the parameters in Table 4.3 is plotted as a solid line. For $Kn > 10$ the measured Q values correspond to the model to within 15 %. From 10^{-4} mbar to 10^3 mbar, the pressure is measured by capacitive gauges (Balzers ACR261 and Pfeiffer CMR365) to minimise the effect of the gas species. Below 10^{-4} mbar a Pfeiffer IKR251 cold cathode gauge was used.

modelled viscous damping within about 15 %.² Below $O(10^{-3})$ mbar, gaseous damping is no longer dominant and Q converges to 2.3×10^5 , a value set by a combination of internal friction (30 %) and unmodelled viscous terms (70 %).

We can draw important conclusions from Fig. 4.20 regarding the required packaging atmosphere for the MEMS seismometers. If a noise level better than $1 \text{ ng}/\sqrt{\text{Hz}}$ is required, inverting Eq. (4.12) immediately gives the requirement $Q_\gamma > 1.6 \times 10^4$ for the device in Fig. 4.20. Assuming an air equivalent residual atmosphere, Fig. 4.20 tells us the package pressure should be below 1.2×10^{-2} mbar. Cavity pressures in this range have been previously demonstrated in wafer level scale encapsulation processes, both with and without a getter material [120, 121, 122].

²Because the characteristic lengths differ between different regions in the MEMS device, they all have their own slightly different Kn . All Kn values are calculated for the $8 \mu\text{m}$ wide gap that contributes the dominant damping term.




Figure 4.21 – (a) Q factors as a function of proof mass displacement for an uncompressed MEMS device with a proof mass of 12.7 mg and a natural frequency of 164.0 Hz as measured on the vibration attenuation platform. The proof mass displacement was modulated by tilting the MEMS on a motorised stage as shown in (b). The modelled viscous damping (free drag and squeezed-film), corresponds to the measured values to better than 5 %.

The factor $\sqrt{m_{\text{air}}/m_{\text{He}}} \approx 2.69$ difference in damping behaviour between air and helium tells us that it is not enough to know the residual cavity pressure to predict the Brownian noise level. The average mass of the gas molecules also plays a role due to the difference in velocity distribution between different gas species. Especially when using a getter material, the average gas mass will not resemble air, as the getter is mainly ineffective for noble gases. However, the influence of gas species on Brownian noise is secondary, as the total noise level only scales as $m_0^{1/4}$. Nevertheless, it is important to avoid common micro-fabrication processes that involve sputtering with noble gas ions, as these will eventually outgas and not be removed by the getter.

4.5.4 Damping considerations of capacitor geometry

Next to a dependence on pressure, the squeezed-film damping as modelled in Eq. (4.37) also shows a direct dependence on the capacitor gap. Since the proof mass in our MEMS sensor has a movement range of $\pm 5 \mu\text{m}$, all gaps can be modulated by this amount. Measuring the device's quality factor as a func-

tion of proof mass displacement, we can verify this modelled dependence. The MEMS sensor was mounted on a rotation stage as shown in Fig. 4.21b, allowing modulation of the proof mass DC position by tilting the sensor. The proof mass displacement can then directly be obtained from $\Delta y = g \sin(\theta)/\omega_0^2$, with $\theta = 0$ indicating the position in which the sensitive axis is oriented orthogonal to local gravity.

Figure 4.21a shows three series of quality factor measurements as a function of proof mass displacement for different vacuum levels of which Fig. 4.20 tells us that the damping is dominated by gaseous effects in the free molecular regime. Additionally, the expected Q_γ from the total modelled viscous damping, γ_f and γ_{sq} , is plotted as a function of the proof mass displacement Δy . The model curve for Q_γ was obtained by replacing d_0 with $d_0 \pm \Delta y$ for the relevant gaps. The geometry dependence of τ from the simulations in Fig. 4.19a is taken into account when plotting the model, and the resulting curve agrees with the measured data to within 5 %.

The results in Fig. 4.21a verify that the modelled squeezed-film damping, which is the dominant source of Brownian noise in the MEMS device, scales like A/d_0 . This happens to be the same scaling as the total sensing capacitance $C_s = \epsilon_0 A/d_0$. Normalising by C_s , we get a squeezed-film damping *per unit of sensing capacitance* as

$$\frac{\gamma_{sq}}{C_s} = \frac{p_0 A \tau}{d_0} \frac{d_0}{\epsilon_0 A} = \frac{p_0}{\epsilon_0} \tau. \quad (4.38)$$

This quantity depends on the capacitor geometry only through τ , meaning that for a given capacitance, the capacitor geometry can be optimised for electrical performance with minimal impact on Brownian noise as long as τ remains constant. The curves in Fig. 4.19a show that in the standard sensing combs, τ changes by less than 10 % for values of d_0 ranging from 4 μm to 30 μm . Using any value for d_0 in this range will provide approximately the same Brownian noise level for a given sensing capacitance. The result in Eq. (4.38) also exposes one of the downsides of using a thicker device layer. The simulations in Fig. 4.19a show that the values for τ are 1.95 to 2.61 times larger for a 100 μm thick device layer as compared to a 50 μm thick one, meaning that for the same sensing capacitance, the former will have roughly twice as much squeezed-film damping.

MEMS fabrication

Manufacturability is an important aspect of any MEMS sensor. Here we outline the basic fabrication sequence that was used to produce the MEMS seismic sensor presented in this work, and highlight the sensor's main design features related to its fabrication. This elementary fabrication process developed at the MESA+ NanoLab at the University of Twente currently works well in terms of sensor functionality, but there still are some challenges to be overcome related to processing yield and device robustness. Different implementations of improvements to the current fabrication process that address both the robustness and yield issues have been investigated and are discussed in Section 5.2. We end with a discussion on the different possibilities for the vacuum packaging of the MEMS seismometers, in order to obtain a self-contained sensor that can be used without external vacuum equipment.

5.1 Elementary fabrication sequence

To obtain the high aspect ratio features such as the capacitor structures and the suspension springs in our MEMS seismometer, highly anisotropic silicon etching is necessary. While there are multiple ways to etch silicon with high directionality, the most convenient way to etch the intricate curved features in our sensor is through deep reactive ion etching (DRIE) using the Bosch process. This process is commonly used in the fabrication of many different types of MEMS devices [105, Chap. 21], and was named after the company that de-




Figure 5.1 – Schematic representation of the mechanism used in deep reactive ion etching (Bosch process). A pulse of SF_6 gas etches silicon isotropically (a), and a subsequent pulse of C_4F_8 gas deposits a passivation layer (b). Directional ion bombardment from the reactor's plasma removes the passivation layer from horizontal surfaces where etching continues during the next etching cycle, while the side walls remain protected (c).

veloped it [123]. It is a high-density plasma etching process performed at reduced pressures, and the underlying principle is schematically illustrated in Fig. 5.1. During the etching phase, a pulse of gaseous SF_6 is fed into the reactor that etches all exposed silicon (Fig. 5.1a). This etching action is nearly isotropic but lasts only a few seconds, such that the etching distance in a single cycle is typically limited to well below $1\text{ }\mu\text{m}$, depending on the desired silicon etch rate. The high directionality is achieved by alternating the isotropic etching phase with a passivation cycle, in which a pulse of gaseous C_4F_8 is fed into the reactor that deposits a protective fluorocarbon film on all surfaces (Fig. 5.1b). During the subsequent etching cycle, this protective layer is removed on all horizontal surfaces due to directional ion bombardment from the high-density plasma, and a new pulse of SF_6 can continue to etch there (Fig. 5.1c). The silicon side walls created during the etching process, however, are not bombarded by ions and remain protected from further etching through the fluorocarbon coating. The fact that the process uses many short isotropic etching cycles results in side walls that are not smooth, but rather show a typical scalloped profile. This can be seen in the schematic representation of Fig. 5.1c, but it is also clearly visible in the electron microscopic image in Fig. 4.6.

5.1.1 Process overview

The process used to manufacture the MEMS seismometer as shown in Fig. 3.2 starts from a single 100 mm diameter silicon-on-insulator (SOI) wafer. A standard SOI wafer consists of a stack of three layers, in this case a $400\text{ }\mu\text{m}$ thick silicon handle layer, a $4\text{ }\mu\text{m}$ thick buried oxide (BOX) layer and a $50\text{ }\mu\text{m}$ thick monocrystalline silicon device layer¹. All the functional structures as described in Chapter 3, such as the readout and actuation capacitors, the suspension

¹ Details on the SOI wafers used for manufacturing can be found in Table A.3 in Appendix A.3.1.




Figure 5.2 – Schematic overview of the fabrication sequence for the prototype MEMS seismometer. An SOI wafer (a) is thermally oxidised with a $2\text{ }\mu\text{m}$ thick oxide layer (b). Where the proof mass will be, a $200\text{ }\mu\text{m}$ recess is etched with KOH (c). Subsequently, the thermal oxide is stripped and the wafer gets a new $2\text{ }\mu\text{m}$ thermal oxide layer (d). The oxide is patterned by a dry etch on both the front and back sides of the wafer (e), and two consecutive DRIE steps transfer these patterns to the silicon layers (f – g). The majority of the oxide is stripped in a 50 % HF solution (h), and finally the proof mass is released by a vapour HF under etch (i).

springs, and their compression systems are etched in the device layer through a DRIE step. The relatively large 12.7 mg proof mass is realised by separating a piece of the silicon in the handle layer from the wafer and suspending it below the device layer. All silicon is chosen to be highly P-doped to provide a low resistivity for the electrical signals without depositing a metal layer. In total, the process uses 3 photomasks.

The essential steps in the elementary microfabrication sequence are shown schematically in Fig. 5.2². The SOI wafer first receives a 2 μm thick thermal oxide layer (Fig. 5.2b). After patterning the bottom oxide layer with the first photomask, an anisotropic KOH etch is used to create a 200 μm deep recess where the proof mass will be (Fig. 5.2c). This etch ensures that the proof mass will be able to move freely when the device is mounted in a package, and moreover, the etching depth determines the total size of the proof mass. After etching the recess for the proof mass, the remaining oxide on the wafer surface is stripped in a 50 % HF solution, and a new uniform 2 μm thick thermal oxide layer is deposited (Fig. 5.2d). Thermal oxidation consumes silicon, and 45 % of the resulting silicon dioxide layer lies below the original silicon surface [105, Chap. 13]. The remaining device layer thickness of the finished devices will therefore always be slightly lower than the nominal 50 μm , which is consistent with the measurement listed in the table in Fig. 3.16b.

After thermal oxidation, the new oxide layers are patterned with two additional photomasks (Fig. 5.2e), and serve as the hard masks in the subsequent DRIE etching steps. The DRIE step on the device layer transfers all the sensor's functional features as defined in the oxide hard mask to the underlying silicon layer (Fig. 5.2f), while the etch from the bottom separates the part of the proof mass that resides in the handle layer from the frame (Fig. 5.2g). Both the DRIE steps on the front side and on the back side of the wafer naturally stop on the buried silicon oxide layer.

After DRIE, all functional structures in the devices are defined, and to finish the fabrication, part of the buried oxide layer should be removed to release the proof mass from the frame. If this layer would be completely removed by wet etching, the subsequent drying process causes so-called stiction, where the capillary forces from the drying liquid pull the moving structures together. These structures remain adhered after drying and as a result the devices will not be functional. Instead, the bulk of the oxide is removed by a timed wet etch in 50 % HF (Fig. 5.2h), and the final part of the BOX layer is removed by a vapour-

² A more complete overview with details on the parameters of the individual processing steps can be found in Appendix A.3.2.

HF etch (Fig. 5.2i) [97, Chap. 25]. The HF vapour etches all exposed oxide and undercuts all silicon features in the device layer. This way, all the narrow and perforated structures are released from the handle layer without the need for drying afterwards. Features in the device layer that are wider than twice the underetch length remain attached to the handle layer and serve either as anchors for the suspended proof mass and the spring compression system, or as fixed electrodes for the sensing and actuation capacitors. After the vapour-HF etch the proof mass is released from the frame and is free to move. Mechanical stops are incorporated in the device layer to limit the proof mass motion to approximately $5\text{ }\mu\text{m}$ in all directions. Apart from releasing the proof mass, the vapour-HF etch also separates all the individual dies from the wafer without the need for dicing. More details on this can be found in Section 5.1.3.

The most critical operation in the processing sequence is the DRIE step in the device layer. Several dummy wafers (both regular silicon and SOI) are used to tune the process parameters to obtain a clean etching profile with vertical side walls. The etching results of several critical features in the device are shown in Fig. 5.3. These SEM images show that it is possible to obtain satisfactory side wall quality in all critical features, although due to several non-uniformities typically encountered in a DRIE reactor such as local variations in plasma density or wafer temperature [124], it is challenging to get uniform and clean profiles like these all across the wafer.

5.1.2 Mitigating overetching damage

The difference in the dimensions of the features to be etched in this design is necessarily quite large. The trenches in the capacitive sensing and actuation structures have to be narrow to obtain a decent sensitivity, and with 7 to $8\text{ }\mu\text{m}$ they are set close to the minimum feature size that can be etched comfortably in a $50\text{ }\mu\text{m}$ thick device layer. At the same time, the suspension springs need enough room to deform when they are compressed to lower the suspension's stiffness. To get to the fully compressed state, the maximum spring deflection is over $50\text{ }\mu\text{m}$. This means that a trench has to be etched along a significant part of the length of the suspension spring that is more than $50\text{ }\mu\text{m}$ wide. Etching these wide open areas required for the deformation of the springs together with the narrow capacitor features is a challenge in DRIE. The majority of the problems arises from the fact that wide features etch faster than narrow ones, an effect called *reactive ion etching lag* (RIE lag) [105, Chap. 21]. RIE lag exists because the gas conductance of narrow and deep trenches is limited. The narrower the trench, the harder it is for etchant to reach the bottom of




Figure 5.3 – SEM images of the DRIE profiles in several critical locations in the design. (a) The actuation capacitors before vapour-HF etching. (b) The anchored electrode of the sensing capacitors after HF-vapour etching with the proof mass removed. The narrow capacitor plates can be seen to be underetched completely. (c) Overview of the compression mechanism with the anchored anti-reverse teeth in (d) where the moving parts have been removed.

the trench, and for the reaction products to exit it again. The result is that etch rates depend on the trench width, something that can clearly be seen in the wafer cross-sections shown in Fig. 5.4. The sensing capacitors have alternating $8\text{ }\mu\text{m}$ and $20\text{ }\mu\text{m}$ trenches, and Fig. 5.4a shows that the wider trenches are deeper than the $8\text{ }\mu\text{m}$ gaps. The wide open areas around the suspension springs in Fig. 5.4b are deeper than the region between the springs, and those are in turn deeper than the holes in the perforated area on the left.

When etching a regular silicon wafer such as the one shown in Fig. 5.4, RIE lag simply leads to different etching depths as a function of feature size. When etching an SOI wafer, however, the situation is different. To properly separate all the different functional structures both electrically and mechanically, *all* features should etch completely through the device layer. On first glance,




Figure 5.4 – SEM images of the cross-section of a regular silicon wafer used to tune the DRIE process for the device layer pattern. The difference in etch rate for differently sized features is clearly visible in both the sensing capacitors (a) and the area around the suspension springs (b).

this is not a problem, because etching naturally stops on the buried oxide layer everywhere. The wider features, however, reach this buried oxide layer *before* the narrowest features complete etching due to RIE lag, and are exposed to etching longer than necessary. This so-called overetching can have some detrimental side effects, because as well the non-conductive buried oxide layer as any electrically isolated silicon structure can develop a net charge due to the directional ion bombardment from the plasma. Due to this charging effect, ions from the plasma will be deflected and start attacking the fluorocarbon passivation layer on the side walls, causing typical notches close to the buried oxide interface [125]. The stability and deposition of the passivation layer is also strongly dependent on the local temperature, which can change in the narrow silicon features after they are released. When the side wall passivation is compromised, overetching will result in severe damage to the silicon.

Part of the problems related to overetching can usually be mitigated by enclosing the most critical features in the device by the narrowest trenches, such that they finish etching last and are not exposed to extensive overetching. However, even in the narrowest trenches a small amount of overetch-




Figure 5.5 – (a) SEM image of the cross-section of a regular silicon wafer close to the centre of the suspension springs. The inner 4 beams are the suspension springs, the outer structures only serve to protect the springs during overetching. A clean etching result on an SOI wafer (b) shows the gap in the right dummy structure. A picture of the same region in a device showing severe side wall erosion (c) shows that the majority of the damage arises in the protective structures. All scale bars correspond to 50 μm.

ing can not be avoided. The etching rate typically slightly varies across the wafer [124] such that the regions that have a higher etch rate have to be slightly overetched in order for all the features on the wafer to finish etching on the BOX layer. Therefore, optimising the etching process for uniform DRIE results all across the wafer is challenging in general.

Wherever possible, a uniform trench width is employed in the mask design and large open areas are avoided. Unfortunately, because of the large range of motion required for their compression, the large open area around the suspension springs can not be avoided. The overetching problems there, however, can be mitigated to a large extent by adding dummy structures around the springs. Figure 5.5a shows a cross-section of a regular silicon test wafer close to the centre of the suspension springs, where a total of 6 silicon beams is visible. Only the central 4 act as springs, the outermost two beams are interrupted somewhere along their length such that they do not deform when compressing the springs, and therefore add no stiffness to the system. This can be clearly seen in Fig. 3.2d, where the springs are shown in their compressed state. Although the outer beams serve no mechanical purpose, they do take most of the damage during the overetching period as witnessed by the extremely eroded side walls shown in Fig. 5.5c. They effectively shield the springs from the large open area, and improve the etching results for these critical components.

5.1.3 Device singulation without dicing

Singulating a wafer into individual devices is typically done as the final step of a microprocessing sequence. In the case of inertial MEMS devices, at this point in the process the proof mass is already released, and the device is sensitive to mechanical shock. Moreover, if no proper measures are taken to protect the sensor mechanics, the debris created by a standard dicing process can end up inside the narrow capacitor gaps and cause short circuits or block the proof mass motion altogether. To prevent the devices from failing during this final processing step, a singulation method without dicing is employed [126, 127]. Instead of dividing the wafer with a dicing saw, both the DRIE steps on the front and back sides of the wafer are used to define the individual device perimeters. During the vapour-HF etch that releases all the moveable structures in the sensor, the individual devices are released from the wafer as well.

This dicing free release process is illustrated in Fig. 5.6. At the end of the processing sequence, the patterned SOI wafer is placed on top of a regular silicon wafer that acts as a support for the released dies and ensures that the proof mass is protected during the release process. The stack of wafers is then placed on a heated wafer holder to be placed in the vapour-HF etcher (Fig. 5.6a).

At this point, all silicon structures in the wafer are still mechanically connected through the buried oxide layer (Fig. 5.6f). After enough of the buried oxide layer has been removed in the vapour HF etch, the devices are released from the wafer. A convenient visible cue for this is shown in Fig. 5.6d, where the individual devices lie slightly below the wafer surface and rest on the carrier wafer, in contrast with Fig. 5.6c where all surfaces lie in the same plane. When all devices are released, the remaining wafer frame can simply be lifted out (Fig. 5.6e), and the devices stay behind on the carrier wafer (Fig. 5.6b). A device cross-section after the release (Fig. 5.6g) shows that this single vapour-HF etch releases both the proof mass and the individual devices from their respective frames. The device's perimeters are designed such that the individual devices detach from the wafer only after all the functional structures inside the device are completely released, *i.e.* the largest features to be underetched are on the device's perimeter. This means that once all devices are seen to be released from the wafer as in Fig. 5.6d, all functional structures should be released too and the etching can be stopped. Such a visible cue in vapour-HF etching is valuable, since the etching rates in this process are difficult to control precisely.

The device cross-section in Fig. 5.6f shows that the DRIE step on the back side of the wafer requires different etching depths for different trenches. The




Figure 5.6 – (a) A photograph of the SOI wafer on the heated vapour-HF wafer holder prior to release etching (see Fig. 5.2h) along with a zoomed image (c). A schematic cross-section along AA' (f) shows that the individual devices are only held in place by the buried oxide layer. During vapour-HF etching the exposed oxide is removed and the individual devices disconnect from the wafer. A visible cue for this is shown in (d) where the devices lie slightly below the wafer surface. This is schematically illustrated by the cross-section along BB' shown in (g). When all devices are released, the remaining SOI wafer frame can be simply taken out (e) and all the singulated devices stay behind on the silicon support wafer (b).




Figure 5.7 – (a) Arcing defects between electrically isolated areas in a single device. The five affected locations are marked by dashed circles. (b) A zoomed image of one of the defects in (a), where a discharge has taken place between the proof mass and an out-of-plane mechanical stop at the location indicated by the arrow.

outer trenches that define the device's perimeter have to go through the full 400 μm thickness of the handle layer, while the trench that defines the proof mass only goes through 200 μm of silicon. To avoid excessive overetching in the proof mass trenches, these trenches are made narrower to offset part of the difference in etching depth through a slower etch rate resulting from RIE lag. The ratio of 1:5 was tuned to have both trenches reach the BOX layer at approximately the same time.

5.2 Towards a more robust fabrication process

During the tuning of the DRIE process parameters on dummy SOI wafers, defects were observed in the device layer at several locations all over the wafer. These defects occurred between areas that are separated by an etched trench. Figure 5.7 shows an overview image with 5 of the affected locations indicated, and a zoomed image of one of the defects in the device layer that occurred between the proof mass and one of the out-of-plane mechanical stops. The fact that this type of defect occurs exclusively across trenches and is only observed when etching SOI wafers and not regular silicon wafers suggests that these defects could be caused by arcing between electrically isolated areas. These areas can in principle develop a difference in electrical potential that leads to an electrical discharge, although it is not completely clear why this difference in potential builds up. Good etching results were obtained by slightly reducing the capacitively coupled plasma power in the reactor from 110 W to 90 W, but the mechanism that leads to these arcing defects is not properly understood. A fabrication sequence that keeps all silicon features electrically con-

nected during DRIE processing and only breaks the electrical connections afterwards might therefore be beneficial to avoid this effect altogether. A possible implementation of such a fabrication process will be described in Section 5.2.2.

Overall, the yield of fabrication was quite low. Even on the SOI wafer with the best DRIE etching results that did not suffer any arcing defects, the yield was only about 45 %. The vast majority of the device failures were caused by stiction, as witnessed by the fact that upon inspection with a profilometer, almost 40 % of the released devices showed a sag close to $4\text{ }\mu\text{m}$, *i.e.* the thickness of the buried oxide layer. This is surprising, because the release etch was done with HF vapour precisely to avoid the stiction problems that are usually associated to the capillary forces induced by a drying liquid in the narrow release channels. The vapour-HF etching process, however, is difficult to control, and additionally, water is released as one of the reaction products that can form a film on the walls of the release channels and possibly cause stiction issues there. The precise mechanism that contacts the proof mass to the handle layer is not well understood, but once contact is established, the adhesion forces between the two smooth surfaces can be stronger than the elastic restoring forces and the proof mass will continue to stick to the handle layer [128].

Both increasing the elastic out-of-plane restoring forces and reducing the adhesion forces between the handle and device layers would help to alleviate the stiction-related issues. The first can be achieved by using SOI wafers with a thicker device layer, and by increasing the distance between the silicon surfaces. Doubling the device layer thickness from $50\text{ }\mu\text{m}$ to $100\text{ }\mu\text{m}$, for example, would increase the restoring forces by a factor 8 for the same suspension springs. Apart from the benefits related to stiction, using a thicker device layer would also significantly increase the unwanted rigid body mode frequencies described in Section 3.3.5, and therefore make the control loop design for the sensor more straightforward. It also increases the mass and both the sensing and actuation capacitances per unit area, allowing for a reduced wafer footprint for the complete sensor. Moreover, reducing the proof mass area decreases the parasitic capacitance, which has benefits for the noise performance, as can be seen in Section 4.3.1. Moving from a $50\text{ }\mu\text{m}$ to a $100\text{ }\mu\text{m}$ thick device layer does make the DRIE process harder to tune for the same feature sizes, because the aspect ratio of all features doubles. Additionally, as was described in Section 4.5.4, doubling the device layer thickness from $50\text{ }\mu\text{m}$ to $100\text{ }\mu\text{m}$ also roughly doubles the squeezed-film damping in the sensing capacitors per unit of capacitance, increasing the resulting Brownian noise by approximately 40 % for a given package pressure.




Figure 5.8 – SEM images of wafer cross-sections after a gaseous XeF_2 etch intended to remove silicon at the bottom of the DRIE trenches only. This works well on a regular silicon wafer (a), but the removal of the thick BOX layer at the bottom of the trenches damages the hard mask on top of an SOI wafer (b) such that the XeF_2 gas also attacks the wafer surface.

The adhesion forces between the handle and device layers are proportional to their contacting surface areas, and can therefore be reduced by physically limiting the possible contact area between the layers, for example by introducing small bumps below the proof mass. Moveable features in the device layer can now only touch the bumps when they move towards the handle layer, instead of its entire parallel surface. This way the total contacting area can be reduced by several orders of magnitude, significantly reducing the adhesion forces between the silicon layers in the device. Several possible implementations of this type of bumps are discussed below.

5.2.1 XeF_2 substrate etch

A first attempt to implement the anti-stiction bumps in the handle layer was made by using a processing sequence similar to the one described in Fig. 5.2. It still uses a single SOI wafer with DRIE steps from both the front and back, but has a few additional steps after the DRIE step on the front side of the wafer. By growing a thin thermal oxide after patterning the device layer with DRIE, all freshly made silicon side walls are covered by silicon dioxide. A subsequent highly directional dry oxide etch opens the oxide at the bottom of the trenches, exposing silicon only at the bottom of all etching trenches. A gaseous XeF_2 etch can then be used to remove silicon at the bottom of all etching trenches.

Figure 5.8a shows an SEM image of the cross-section of a regular silicon wafer that was used to test this concept. A thick oxide hard mask on top is

used to define 10 μm slits that are etched into the silicon by DRIE to a depth of approximately 60 μm . A few hundred nanometre thick thermal oxide is grown on all exposed silicon surfaces and subsequently removed at the bottom of the trenches by a directional dry oxide etch. When XeF_2 gas is introduced, it starts isotropically etching silicon at the bottom of the trenches, while the rest of the silicon is nicely protected. For a regular silicon wafer the oxide skin that protects the silicon side walls is etched free by the XeF_2 gas at the bottom and is visible in Fig. 5.8a. For an SOI wafer, this would allow for selectively removing silicon below the device layer and realising anti-stiction bumps in the handle layer.

Unfortunately, the presence of the 4 μm thick buried oxide layer in an SOI wafer complicates the opening of the oxide at the bottom of the DRIE trenches. The sensor design can not tolerate a significantly thinner BOX layer because of parasitic capacitance considerations. Combined with the fact that the etch rate at the bottom of narrow trenches is significantly reduced, this means that the wafer has to be exposed to oxide dry etching for a relatively long time. Although the oxide hard mask on top of the wafer is protected by photoresist during this etch, the prolonged etching does damage the oxide hard mask to the extent that it becomes porous to XeF_2 gas. This is clearly visible in the SOI wafer cross-section in Fig. 5.8b, where next to nicely removing the silicon below the BOX layer as intended, the XeF_2 gas has also attacked the silicon at the top of the device layer through the damaged oxide mask.

This problem has been partially solved by not opening the BOX layer with a single etch, but depositing a fresh layer of PECVD oxide³ after every etched micrometre or so. The deposition rate of PECVD at the surface is significantly higher than at the bottom of the trenches, such that the protection at the surface is refreshed after each deposition step, but the BOX layer does not grow in thickness significantly. This greatly decreases the XeF_2 etching damage observed at the surface of the wafer as can be seen by comparing the top oxide silicon interface in Fig. 5.9a with that in Fig. 5.8b. The damage is not avoided entirely, however, as the sharp corners at the top of the trench have still been attacked by the etchant. A close-up of a typical edge profile is shown in Fig. 5.9b.

As can be seen around the sensing capacitors in the cross-section in Fig. 5.9a, the XeF_2 has etched away silicon from the handle layer below all trenches in the device layer to a depth of about 20 μm . Because the proof mass is perforated with a regular grid of holes in this process, the handle layer surface below it is also recessed by this 20 μm . This increases the gap between the silicon lay-

³ PECVD: Plasma Enhanced Chemical Vapour Deposition, a thin film deposition method.




Figure 5.9 – (a) SEM image of an SOI wafer cross-section at the sensing capacitors with the proof mass removed and the thick PECVD oxide still present. The 20 μm XeF₂ under-etch is clearly visible below the structures device layer. (b) Close-up of the XeF₂ etching damage to sharp corners in the device layer after all oxide has been removed. An interruption of the perforated grid shown in (c) causes an anti-stiction bump to form in the handle layer as shown in (d).

ers from 4 μm to roughly 25 μm and therefore drastically reduces the parasitic capacitance between them. The biggest advantage, though, is the possibility to introduce anti-stiction bumps in this way, by interrupting the regular grid of holes by a solid circle as shown in Fig. 5.9c. Because the XeF₂ gas can only remove the silicon below holes in the device layer, a bump such as the one shown in Fig. 5.9d will form below every circle placed in the proof mass grid. This reduces the possible contact surface area between the device and the handle layer by about 3 orders of magnitude, significantly reducing the risk for stiction.

After the XeF₂ etch, the proof mass is almost completely underetched. At the locations of the anti-stiction bumps, however, it is still connected to the handle layer through the buried oxide, and this layer should therefore be removed to release the proof mass. The bulk of the oxide present on the devices

was to be removed by a 50 % HF wet etch followed by a vapour-HF etch to do the final release from the bumps. However, a successful release could not be obtained, presumably due to a combination of the large exposed surface area of the BOX layer below the proof mass and wetting issues during the 50 % HF etch. The unsuccessful release meant that no functioning devices were produced from this processing run. In combination with the fact that the iterative etching and PECVD for opening the BOX layer became quite involved, this led to the decision not to pursue this option any further. Instead, efforts were focused on an alternative approach to implement the anti-stiction bumps that will be discussed below.

5.2.2 Wafer bonding

The majority of the problems with the process proposed in Section 5.2.1 is related to the fact that selective etching needs to be performed at the bottom of deep and narrow trenches. This etching at the bottom of trenches can be avoided by creating the cavity and the bumps on the surface of a regular silicon wafer and bonding the device layer to it afterwards. This approach gives full control over the exact shape and depth of the cavity independent of the features etched in the device layer, and gives additional freedom to choose the position and height of the anti-stiction bumps.

Process overview

A schematic overview of the processing steps to obtain such a cavity wafer is shown in Fig. 5.10. A silicon wafer is covered with a 4 μm thick thermal oxide. This oxide has to be chosen thick to minimise parasitic capacitances in the finalised device. A first lithography step and a subsequent dry oxide etch open the top oxide layer where the recess should be etched to its full 15 μm depth (Fig. 5.10b). A second lithography step then defines the full extent of the recess (Fig. 5.10c). A first 10 μm DRIE step etches the exposed silicon, leaving small bumps where oxide still covers the surface (Fig. 5.10d). This etch defines the height of the resulting anti-stiction bumps. After removing the exposed oxide by another dry oxide etch, a second 5 μm DRIE step deepens the entire recess, and together with the 4 μm thick oxide defines the final out-of plane motion range of the proof mass (Fig. 5.10e). After the photoresist has been stripped, a 100 nm thick thermal oxide is grown to cover all exposed silicon in the recess (Fig. 5.10f). Because this wafer is to be bonded to another wafer containing the proof mass and the rest of the functional structures, a stack of 10 nm tantalum, 30 nm platinum and 250 nm gold is evaporated onto the wafer and patterned by lift-off lithography to define the bond interface for thermo-




Figure 5.10 – Schematic overview of the fabrication sequence for the anti-stiction bumps in a separate silicon wafer. A $4\text{ }\mu\text{m}$ thick thermally grown oxide (a) is patterned by a dry oxide etch (b). A second lithography step (c) defines the recess that is etched by using DRIE (d). After removing the exposed oxide, a second DRIE step finishes the recess with anti-stiction bumps (e). A 100 nm thick thermal oxide is grown after photoresist stripping, followed by lift-off lithography of an evaporated Ta/Pt/Au stack (f).

compression bonding.

As before, all the sensor's functional parts are realised in the device layer of an SOI wafer, which was chosen $100\text{ }\mu\text{m}$ thick for this process. The larger thickness increases the out-of-plane stiffness for even higher resistance against stiction, but makes DRIE harder to tune. To facilitate the DRIE process, a uniform etching trench width of $8\text{ }\mu\text{m}$ was chosen everywhere in the design⁴. The required space for the large motion range of the suspension springs and their compression system is made by etching selected silicon features with XeF_2 at a later stage. Figure 5.11 shows a schematic overview of the processing steps required for patterning the device layer. It is highly similar to the process outlined in Section 5.1.1, with the difference that the DRIE trench width is uniform across the entire wafer. As with the silicon wafer in Fig. 5.10, the processing ends with growing a thin oxide to cover any exposed silicon and a lift-off pro-

⁴ The trench width on the photomask is set to $5.5\text{ }\mu\text{m}$, as etching tests revealed that the resulting $100\text{ }\mu\text{m}$ deep trenches were approximately $2.5\text{ }\mu\text{m}$ wider than defined by lithography.




Figure 5.11 – Schematic overview of the fabrication sequence for patterning the device layer with all the functional structures. An SOI wafer (a) is thermally oxidised with a 2 μm thick oxide layer (b). The oxide is patterned by a dry etch (c) and a DRIE step transfers the pattern to the silicon device layer (d). The oxide hard mask is stripped using HF (e), and a new 200 nm thick thermal oxide is grown followed by lift-off lithography of an evaporated Ta/Pt/Au stack (f).

cess to pattern a Ta/Pt/Au stack that defines the bonding interface (Fig. 5.11f).

Both processed wafers are heated to around 400 $^{\circ}\text{C}$ and pressed together with a force of approximately 20 kN to form a thermocompression bond where the gold interfaces on both wafers touch. This process forms a strong hermetic bond that can withstand temperatures much higher than the process temperature [120]. The process to get finished devices from the bonded wafer stack is schematically depicted in Fig. 5.12. The SOI handle layer that is now on top is completely removed by a 400 μm DRIE step that naturally stops on the BOX layer (Fig. 5.12b). Access holes to the silicon in the device layer that has to be removed are patterned in the now exposed BOX layer by a dry oxide etch (Fig. 5.12c). A subsequent XeF_2 etch removes the silicon that has been exposed by the access holes, and creates the space necessary for the suspension springs and their compression system to operate (Fig. 5.12d). All exposed silicon was oxidised before bonding the wafer together, so the rest of the structures should be protected even if the XeF_2 gas enters the cavity below the




Figure 5.12 – Schematic overview of the fabrication sequence for combining the wafers. After a thermocompression bond (a), the exposed SOI handle layer is completely removed by DRIE (b). Selected regions in the silicon device layers are removed by opening the buried oxide (c), and etching the exposed silicon with XeF_2 (d). A vapour-HF etch finally releases the structures (e).

proof mass. A final vapour-HF etch removes all the exposed oxide and releases the proof mass (Fig. 5.12e). This fabrication sequence uses a total of 7 photo-masks, and fits 47 devices on a single 100 mm diameter wafer.

First run results

The uniform trench width used across the design led to a good DRIE result in the 100 μm thick device layer without excessive overetching and the associated side wall damage. Moreover, all critical features remain electrically connected during DRIE, such that charging effects are avoided. As can be seen in the wafer cross-section at two of the thermal actuators in Fig. 5.13a, the etching profiles are nicely vertical and uniform over the thickness of the device layer. The zoomed image in Fig. 5.13b shows that there is only minimal notching at the bottom few micrometres of the etched trench, where both the side walls retract by approximately 1 μm . The 8.6 μm nominal trench width in this location is slightly higher than the anticipated average 8 μm projected from DRIE tests, but slightly reducing the slit widths on the photomask from 5.5 μm to




Figure 5.13 – (a) SEM image of an SOI wafer cross-section showing the DRIE profiles at two of the thermal actuators for a 100 μm thick device layer. The zoom in (b) shows a nominal trench width of 8.6 μm resulting from a 5.5 μm wide slit on the photomask, and only minimal notching in the bottom few micrometres of the etched trench.




Figure 5.14 – Gaseous XeF_2 only etches silicon that has been exposed by small circular holes in the oxide layer on top. (a) A strip of silicon is removed from the actuation capacitors. XeF_2 etching is not fully completed at this point. (b) Electrical connections between different silicon parts in the device layer are broken by the XeF_2 etch.

5 μm should bring all feature dimensions close enough to their design values. By using the uniform trench width approach to DRIE, moving to a 100 μm thick device layer is feasible.

The XeF_2 etch of the silicon device layer (Fig. 5.12d) removes selected silicon structures only where small holes are etched in the oxide layer on top. This can be seen in Fig. 5.14, where the structures below the access holes have been etched, but the rest of the silicon remains intact. This etching step serves several distinct purposes. First of all it opens up space required for the large deformations of the suspension springs, and increases the distance between certain electrodes to alter their mutual capacitance. The latter is shown in Fig. 5.14a, where a strip of silicon is removed from the actuation capacitors to

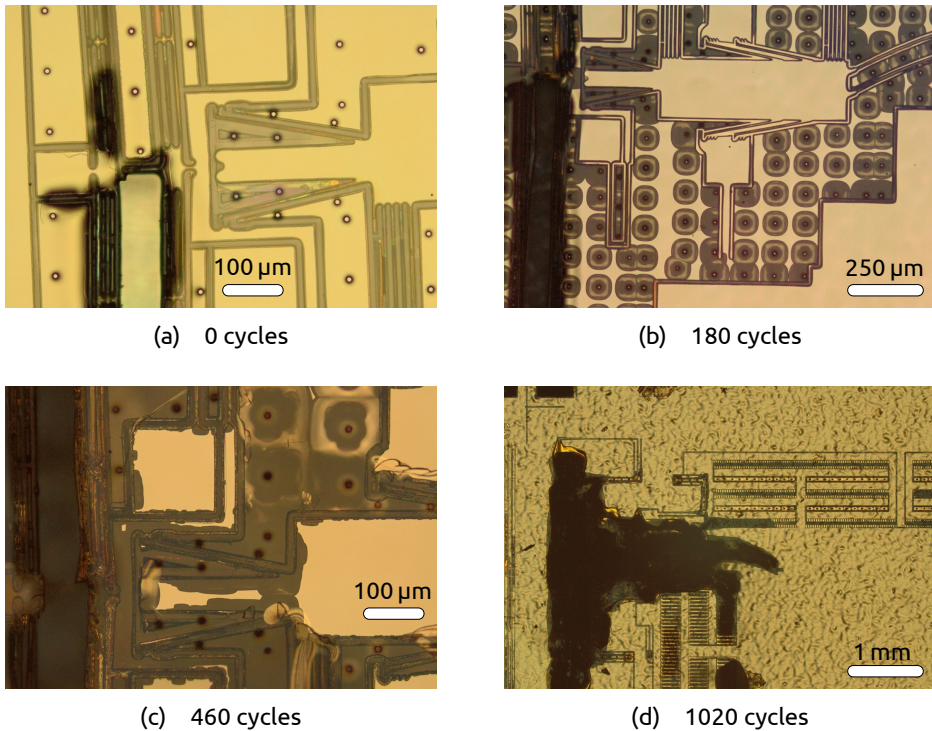


Figure 5.15 – Progress of the XeF₂ etch around one of the damaged thermal actuator areas in the wafer. One of the filler structures between the thermal actuators has fallen out of the wafer (a). Initially, the etching starts only below the intended access holes (b), but the etchant also attacks unintended regions as etching proceeds (c). When the process completes after 1020 etching cycles, the compression system and the suspension springs are severely damaged (d).

reduce the unwanted capacitance between the moving and static electrodes. Note that moveable parts do not need to be perforated in this process, as the cavity and bonding interface can be defined independent of the vapour-HF etch. Another important function of the XeF₂ etch is to break the electrical connection between all individual electrodes, as shown in Fig. 5.14b. Removing the central silicon part here breaks the connection between the three narrow strips of silicon that run to different electrical parts in the device.

After removing the silicon handle layer from the SOI wafer (Fig. 5.12b) the freshly exposed BOX layer turned out to be damaged, particularly around the silicon filler structures between the two thermal actuators. The damage is presumably caused by a combination of damage to the BOX layer from rework done on the wafers, and a differential pressure across the oxide membrane when venting the plasma etcher. Figure 5.15a shows the central part of two

thermal actuators, where the damage to the oxide layer caused the bottom silicon filler structure to fall out of the wafer completely. The large resulting gap causes water to enter the cavity below the proof mass during the lithography process to pattern the XeF_2 access holes. XeF_2 gas is highly reactive with water, forming HF that can attack the thin oxide protecting the silicon features in the device layer [105, Chap. 11]. Indeed, the XeF_2 starts etching mainly the silicon below the intended access holes (Fig. 5.15b), but as etching continues, also functional structures in the sensor start being etched from the side (Fig. 5.15c). When the XeF_2 etching is finished, the entire region around the thermal actuators and the suspension springs is severely damaged (Fig. 5.15d). This first run therefore did not produce any functional devices. At the time of writing a second run is in preparation, that will slightly alter the sequence depicted in Fig. 5.12 by leaving part of the silicon handle from the SOI wafer and performing the last lithography step on top of this layer. This avoids the oxide membrane from breaking and prevents water from entering into the cavity below the proof mass.

5.3 Vacuum packaging

As seen in Section 4.5.3, our seismometer should operate at a pressure below approximately 10^{-2} mbar to have an acceptable Brownian noise level below $1 \text{ ng}/\sqrt{\text{Hz}}$. Up to now, all low-noise testing has been performed with the MEMS seismometer placed inside vacuum tanks that were being evacuated with external pumps, but this quickly becomes impractical for testing the performance of small sensor arrays in the field. To avoid the use of separate vacuum equipment, there are two different approaches to packaging MEMS devices hermetically at package pressures better than 10^{-2} mbar that will be briefly discussed here.

5.3.1 Single die packaging

Ceramic packages are frequently used for housing both MEMS sensors and integrated circuits. They can be hermetically sealed by bonding a lid on top with a eutectic Au/Sn solder. When the devices are sealed under vacuum conditions, the package contains its own vacuum and external tanks and pumps are no longer required. However, outgassing of the surfaces inside the package can never be completely avoided, and the quality of the vacuum inside the package will deteriorate over time. When the packaged vacuum quality is important, a getter material is often included inside the package that effectively




Figure 5.16 – The Sercel QuietSeis MEMS seismometers (see Fig. 2.4) are individually packaged in a ceramic package by bonding a lid on top that contains a getter material which maintains a low package pressure over the lifetime of the device. Figure reproduced from [80].

acts as a small vacuum pump. The getter that is typically used for these applications is a patented zirconium alloy that absorbs gas molecules after it has been thermally activated at a temperature between 300 °C and 350 °C [129]. A getter inside the package will both reduce the package pressure, and extend the lifetime of the device by counteracting most of the outgassing inside the package.

Sercel’s QuietSeis MEMS based seismometer mentioned in Section 2.3 is packaged in such a ceramic package with a getter, as shown in Fig. 5.16. A commercial platform exists that provides this type of packaging for custom devices at reported pressures better than 10^{-4} mbar [130]. To get a sensor prototype suited for field testing, the devices that will be fabricated in the second run of the process described in Section 5.2.2 are planned to be packaged using this platform.

5.3.2 Wafer-level vacuum packaging

Another approach to creating a vacuum environment around the relevant parts of a MEMS device is to incorporate the hermetic sealing as part of the wafer-scale processing. This has certain advantages for the scalability of the manufacturing, but also poses significant challenges that have to be addressed. The hermetic encapsulation is typically obtained by bonding another wafer on top of the one containing the finished sensor mechanics. A typical finished device cross-section could look like the one shown in Fig. 5.17a. A process like this has the advantage that now robust mechanical range limiters in the form of the anti-stiction bumps exist on both sides of the proof mass, making the devices less vulnerable to out-of-plane shocks.

One of the challenges related to wafer-level encapsulated devices is making the proper electrical connections. The proof mass and all capacitive elec-

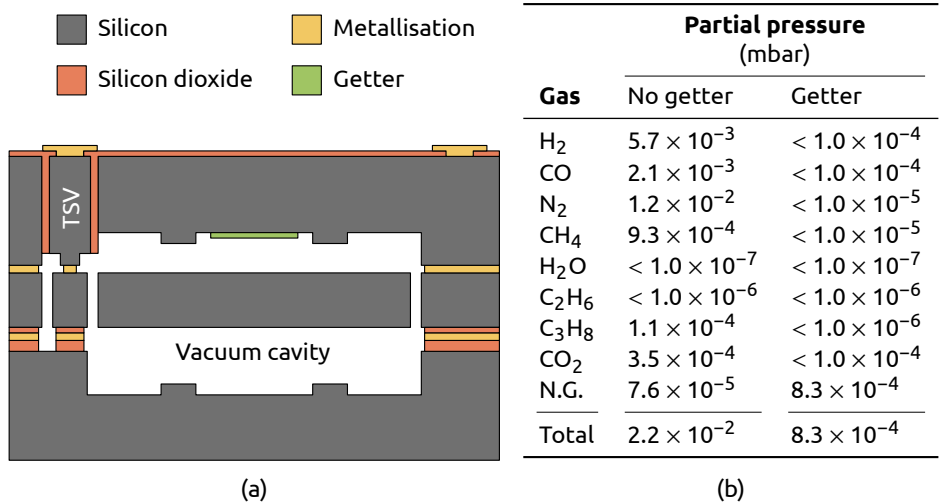


Figure 5.17 – (a) A device cross-section of a possible implementation of a wafer-level vacuum packaged seismometer. (b) Residual gas analysis of the 4.8 mm³ hermetically sealed cavity of a MEMS pressure sensor with and without getter. N.G. stands for noble gases. Data reproduced from [131].

trodes reside in a hermetically sealed cavity, and routing an electrical connection to individual bond pads outside that cavity is not straightforward. One of the possible solutions is to employ through-silicon vias (TSV) to provide an electrical connection through the thickness of the capping wafer [97, Chap. 38]. An example of a connection made using such a via is shown in Fig. 5.17a.

Wafer-level encapsulation typically results in cavities with a significantly lower volume than those in the vacuum ceramic packages, making the effect of outgassing from the surfaces of the MEMS die on the vacuum level more pronounced. Including a getter material inside the cavity helps to counteract outgassing, but it does not absorb all gases. For example, the table in Fig. 5.17b shows a comparison of a residual gas analysis for a hermetically sealed pressure sensor with and without a getter material [131]. The device with the getter inside the cavity has a cavity pressure that is 25 times lower than in the device without the getter. The getter reduces the partial pressures of all measured gases to below the detection limit, with the exception of the noble gases (N.G.) that are inert and do not interact with the getter material. A residual noble gas background is sometimes introduced on purpose to create a package with a well defined pressure [97, Chap. 39], but when the goal is to obtain a cavity with as low as possible pressure, outgassing of noble gases should be avoided. This has consequences for the design of the fabrication sequence,

as most processes that involve ion milling or sputtering use argon ions [97, Chap. 6] that will be embedded into the devices to some extent. Indeed, it has been shown that the outgassing of argon from sputtered metal layers is over 3 orders of magnitude larger than that of evaporated layers [129]. Care has been taken in designing the fabrication process described in Section 5.2.2 to only use processes and materials that are compatible with wafer-level vacuum packaging. For example, all metals are deposited through evaporation instead of the more conventional sputtering, and the thermocompression bond provides a hermetic seal around the device perimeter.

Conclusion and outlook

6.1 Conclusion

Large-scale interferometric gravitational wave detectors are formidable machines. Through a continuous battle against noise of vastly different origin over the past decades, they are now able to measure relative length changes of the order of $10^{-23}/\sqrt{\text{Hz}}$. As their sensitivity improves and their detection range increases, it becomes more and more important to also expand their detection band towards lower frequencies, because the signals from more distant sources will be red shifted towards the lower end of the frequency spectrum. While both the Advanced LIGO and Advanced Virgo detectors continue to improve their sensitivities towards the design goals, efforts also intensify to realise a next generation of gravitational wave detectors. Initiatives such as Einstein Telescope propose to improve the strain sensitivity by an order of magnitude compared to the current generation of detectors, and at the same time reduce the lower frequency cut-off to approximately 2 Hz. This combination of improvements will lead to the ability to observe mergers of stellar mass binary black holes throughout the entire Universe.

Many of the sensitivity limitations for Earth-based detectors at low frequencies are related to seismic motion. The direct coupling of ground motion to the detector's test masses can be sufficiently suppressed by state-of-the-art vibration isolation suspensions. Unfortunately, however, this does not hold for the Newtonian noise that is inevitably associated with this ground motion. Even

though Einstein Telescope is planned to be built underground to reduce many of the detrimental effects of seismic motion, developing Newtonian noise mitigation techniques is a priority if this next generation gravitational wave detector is to reach its design sensitivity. Next to a meticulous site selection and passive measures, active suppression techniques requiring arrays of seismic sensors around each of the detector's test masses have to be pursued.

This dissertation is focused around a novel design for a MEMS based seismic sensor and has demonstrated a self-noise performance below the $2 \text{ ng}/\sqrt{\text{Hz}}$ seismic background requirement set for the future Einstein Telescope host site. Moreover, the MEMS-based sensor shows the potential to be mass produced through exploiting the general batch processing nature of common microfabrication techniques. The sensor has been designed as a force balance accelerometer, which has large advantages in terms of sensitivity over the more traditional geophone that is commonly used in larger seismic arrays, especially at the lower end of the frequency spectrum.

The MEMS seismic sensor has demonstrated unprecedented self-noise performance for its limited proof mass of 12.7 mg. The main driver for this is the miniaturised geometric anti-spring system inspired by those used in the auxiliary vibration isolation suspensions in the Advanced Virgo gravitational wave detector. A novel thermally actuated latching mechanism was designed to accurately stress the spring structures that suspend the sensor's proof mass at the microscale. Doing this allows for a passive reduction of the suspension stiffness of 20 to 25 times, increasing the device's sensitivity by the same factor. A general dimensionless design methodology for the anti-springs was presented that can be used both for assessing the feasibility of a geometric anti-spring implementation at any desired scale, and for designing the spring elements themselves. Compared to conventional MEMS suspension springs, the geometric anti-spring suspension presented in this work can attain a low mechanical in-plane stiffness with a significantly improved out-of-plane stiffness.

Manufacturing tolerances have proved to pose a significant challenge for adapting the presented design for measuring vertical accelerations. To avoid any excessive dynamic range requirements on the readout system, the sensor mechanics should accurately balance the large offset that is introduced by gravity. This balance, however, is largely compromised by a poor control of the absolute stiffness of the suspension springs in the device. It has been shown that an additional compensation spring with a moveable anchor is able to cope with this spread in absolute spring stiffness, and as a result makes a vertical sensor design feasible.

Because of the low seismic background requirement for the future host site of Einstein Telescope, the novel seismometer's self-noise performance is its most important property. An extensive study of all the different major noise sources of both mechanical and electronic origin was performed, supported by direct measurements of all the respective noise spectra. Noise measurements on a vibration isolation platform showed that the stiffness reduction introduced by the geometric anti-spring system effectively reduced the electronic front-end noise as intended. As a result, the noise performance below approximately 20 Hz was completely limited by mechanical noise sources in the MEMS sensor itself. Results from the three-channel correlation setup at the Heimansgroeve seismic station confirmed this, and additionally showed that the noise performance is dominated by the sensor's mechanical noise down to at least 10 mHz. The resulting MEMS acceleration noise floor lies below $2 \text{ ng}/\sqrt{\text{Hz}}$ between 400 mHz and 20 Hz and remains below $10 \text{ ng}/\sqrt{\text{Hz}}$ all the way down to approximately 50 mHz. This result matches the noise models put forward in this dissertation to better than 20 % over the entire measured frequency range.

The seismic sensor is produced in an elementary microfabrication process that starts from a single silicon-on-insulator wafer and uses a total of three photomasks. The process currently uses wafers with a 50 μm thick device layer, but deep reactive ion etching tests have shown that it is feasible to move to wafers with a 100 μm thick device layer. This brings significant advantages in terms of out-of-plane stiffness, the device's wafer footprint, parasitic resonance modes and the device robustness in general. Additional efforts to improve the robustness of both the fabrication process and the resulting devices have converged to adapting a wafer bonding approach, where the interface between the device layer and the mechanical handle can be tailored to alleviate stiction related issues. Eventually, reaching the reported noise performance in a practical sensor will also require packaging the MEMS mechanics in a local vacuum environment. The extensive measurements and modelling of the effects of gaseous damping on the sensor mechanics tell us that the pressure inside such a hermetic package should be better than approximately 10^{-2} mbar when assuming an air equivalent atmosphere. In the design of the new fabrication process, care was taken for all processing steps and materials to be compatible with any future vacuum encapsulation solution, be it at the single-die-level or at the wafer-level.

6.2 Outlook

To be able to use the seismic sensors for their intended purpose, some development is still required, especially related to the sensor's manufacturing process. Both local vacuum encapsulation and a fabrication process that yields more robust devices are required to make further field testing convenient. At the time of writing, a fabrication run is under way that implements a slight variation of the process described in Section 5.2.2 that should solve the issues encountered in the previous run. The devices that the current run yield are to be encapsulated in a vacuum at the single-die-level inside a ceramic package containing a getter material. Contacts with the relevant manufacturing companies are in place, and we will soon have a small array of sensors that can be tested in the field.

A second important ingredient for operating multiple sensors in the field is the electronic readout system. For testing in small arrays we are currently developing a readout front-end that consists of discrete components on a PCB and implements all the functionality of the readout circuit described in Fig. 4.7. This approach is not necessarily scalable to larger arrays, so in parallel we are developing a switched capacitor front-end in a custom *application specific integrated circuit* (ASIC), that will house the same functionality on a small CMOS chip. In the longer term, this ASIC will also include all the required loop filters implemented in the form of a $\Delta\Sigma$ -modulator on chip [62, 132]. Together with the ASIC, the MEMS seismic sensor would provide an acceleration sensor with good noise performance in a compact form factor that can be of interest for applications outside of science as well. For example, these sensors could be a viable option for large scale seismic reflection surveys for geophysical studies or in the seismic service industry. Especially at low frequencies, they will provide a better noise performance than the more conventionally used geophones. In general, the batch fabrication nature of MEMS devices makes them attractive for any application that needs a good noise performance at large device volumes.

Several of the techniques developed in the context of this work have a broader application range. Firstly, the technique presented in Section 4.2.2 for accurately measuring internal losses in the suspension springs can be put in a broader perspective, as the characterisation of mechanical losses plays a role in many research fields. By measuring the system's quality factor for a range of different natural frequencies, internal friction losses can be distinguished from viscous damping effects. In fact, an accurate estimate for the internal friction loss angle ϕ can be obtained even when viscous losses domin-

ate. This is impossible to do with measurements at a single natural frequency only. Secondly, when using three-channel correlations to obtain an estimate for the noise floor of a specific sensor, an accurate common mode subtraction is beneficial. The three-channel correlation technique itself is well-known, but the data processing typically does not include an extra alignment procedure while this is generally possible for any triaxial sensor. As demonstrated in this dissertation, the additional alignment step significantly improves the common mode subtraction performance of the algorithm, and can be useful in all situations where the estimated noise floor lies significantly below the common sensor stimulus.

The pursuit of the first direct detection of gravitational waves over the past few decades has given impulses to technological innovations in several different fields. Clearly the novel MEMS seismometer presented in this work was developed in the context of gravitational wave detectors, but these detectors also feature laser technology, optical coatings, and vibration attenuators with unprecedented performance. Apart from general technological advancement, the first direct detection of gravitational waves from a pair of merging black holes also marked the beginning of a completely new field of research: *gravitational wave astronomy*. With the recently acquired ability to directly observe these waves came an abundance of new information on processes in the distant Universe, exotic objects living inside of it, and their properties. The rich scientific output after the detection of one single event, the merger of a binary neutron star, excellently illustrated how powerful the addition of the gravitational wave spectrum to the arsenal of observable messengers from the cosmos really is. This single detection started the field of multimessenger astronomy with gravitational waves. The realisation of Einstein Telescope in the future will significantly expand the current detection range, and will firmly establish gravitational wave detection as a tool for understanding the distant Universe. Undoubtedly, this will lead to numerous new discoveries and will advance our understanding of the early Universe. These are exciting times.



Appendices

A.1 Curved spring out-of-plane stiffness

The full out-of-plane compliance matrix for a curved suspension spring spanning an angle $2\theta_0$ of an arc with constant radius R_0 as shown in Fig. A.1 can be derived from tabulated expressions [90, Tab. 9.4]. The relevant second moment of area for the spring is given by $I = wt^3/12$, and E and G denote the silicon Young's modulus and shear modulus, respectively. The parameter β is defined as $\beta = EI/GK$ where K is the torsional stiffness constant of the spring cross section. For a rectangular cross section, K can be approximated as [90, Tab. 10.1]

$$K \approx tw^3 \left(\frac{1}{3} - \frac{3.36w}{16t} \left(1 - \frac{w^4}{12t^4} \right) \right). \quad (\text{A.1})$$

For the typical values of $t = 50 \mu\text{m}$ and $w = 8.7 \mu\text{m}$, we have $K \approx 9.6 \times 10^{-21} \text{ m}^4$, and when using $E = 169 \text{ GPa}$ and $G = 50.9 \text{ GPa}$ we have $\beta \approx 31$.

Denoting the rotations of the beam's tip around the x- and y-axes as θ_x and θ_y respectively, the general out-of-plane compliance matrix for any spring beam element can be written as

$$\begin{bmatrix} z \\ \theta_x \\ \theta_y \end{bmatrix} = \begin{bmatrix} c_{11} & c_{12} & c_{13} \\ c_{21} & c_{22} & c_{23} \\ c_{31} & c_{32} & c_{33} \end{bmatrix} \cdot \begin{bmatrix} F_z \\ T_x \\ T_y \end{bmatrix}. \quad (\text{A.2})$$




Figure A.1 – Schematic representation of a curved suspension spring spanning an angle $2\theta_0$ of an arc in the xy -plane with constant radius R_0 . The left side is clamped and the right side loaded with a force F_z and torques T_x and T_y , such that it translates out of plane and rotates with angles θ_x and θ_y around the x - and y -axes respectively.

In the case of a curved spring clamped at one side and loaded at its tip, the elements that make up the compliance matrix are defined as

$$\begin{aligned}
 c_{11} &= \frac{R_0^3}{2EI} \left(\frac{1}{2}(\beta - 1) \sin(4\theta_0) - 4\beta \sin(2\theta_0) + (6\beta + 2)\theta_0 \right) \\
 c_{12} &= \frac{R_0^2}{2EI} \left(((\beta - 1) \sin(2\theta_0) + (2\beta + 2)\theta_0) \cos(\theta_0) - 4\beta \sin(\theta_0) \right) \\
 c_{13} &= \frac{R_0^2}{2EI} \sin(\theta_0) ((\beta - 1) \sin(2\theta_0) - (2\beta + 2)\theta_0) \\
 c_{21} &= c_{12} \\
 c_{22} &= \frac{R_0}{2EI} ((2\beta + 2)\theta_0 + (\beta - 1) \sin(2\theta_0)) \\
 c_{23} &= 0 \\
 c_{31} &= c_{13} \\
 c_{32} &= 0 \\
 c_{33} &= \frac{R_0}{2EI} ((2\beta + 2)\theta_0 - (\beta - 1) \sin(2\theta_0))
 \end{aligned} \tag{A.3}$$

This compliance matrix can be inverted numerically to obtain the out-of-plane stiffness matrix for a curved beam with constant radius R_0 .

A.2 MEMS nominal parameters

Tables A.1 and A.2 list the nominal values for the most important design parameters of both the horizontal and vertical MEMS seismometers presented in this work. Manufacturing tolerances may cause some of the parameters to deviate from their nominal values. Wherever relevant this is indicated in the text.

Table A.1 – Nominal values for the most important design parameters of the horizontal MEMS seismometer as shown in Fig. 3.2.

Part	Parameter	Symbol	Value	Unit
general	die size	-	12×12.8	mm^2
	device layer thickness	t	50	μm
	y-mode freq. @ $\Delta x = 0 \mu\text{m}$	$f_{y,0}$	172	Hz
	y-mode freq. @ $\Delta x = 35 \mu\text{m}$	f_y	38	Hz
proof mass	total mass	m	12.7	mg
	moment of inertia x-axis	I_{xx}	4.92×10^{-11}	kgm^2
	moment of inertia y-axis	I_{yy}	1.24×10^{-11}	kgm^2
	moment of inertia z-axis	I_{zz}	6.15×10^{-11}	kgm^2
springs	total beam length	L	1778	μm
	beam width	w	8.6	μm
	spanning angle	$2\theta_0$	59	$^\circ$
	springs per corner	N	4	-
	max. compression length	Δx_{\max}	35	μm
sensing capacitors	finger separation	d_0	8	μm
	secondary gap	d_1	20	μm
	finger overlap	L_c	232	μm
	fingers per side	N_c	410	-
	capacitance per side	C_0	7.4	pF
actuation capacitors	finger separation	d_a	7	μm
	finger width	w_a	7	μm
	finger overlap	L_a	40	μm
	fingers per side	N_a	490	-
	capacitance per side	C_a	2.5	pF
ETA	half length	L	1450	μm
	cross-sectional area	A	1.2×10^{-9}	m^2
	second moment of area	I_z	1.86×10^{-19}	m^4
	clamping angle	θ	2	$^\circ$
	ETAs per corner	-	2	-

Table A.2 – Nominal values for the most important design parameters of the vertical MEMS seismometer as shown in Fig. 3.20.

Part	Parameter	Symbol	Value	Unit
general	die size	-	8.1×8.1	mm^2
	device layer thickness	t	50	μm
	y-mode freq. @ $\Delta x = 0 \mu\text{m}$	$f_{y,0}$	248	Hz
	y-mode freq. @ $\Delta x = 35 \mu\text{m}$	f_y	69	Hz
proof mass	total mass	m	1.57	mg
springs	total beam length	L	1778	μm
	beam width	w	8.6	μm
	spanning angle	$2\theta_0$	59	$^\circ$
	spring rotation	ϕ	0.11	$^\circ$
	springs per corner	N	1	-
	max. compression length	Δx_{\max}	35	μm
compensation spring	stiffness	-	0.11	N/m
	initial step	$\Delta y_{c,\min}$	15	μm
	step resolution	-	2.5	μm
	max. anchor displacement	$\Delta y_{c,\max}$	75	μm
sensing capacitors	finger separation	d_0	8	μm
	secondary gap	d_1	20	μm
	finger overlap	L_c	232	μm
	fingers per side	N_c	76	-
	capacitance per side	C_0	1.4	pF
actuation capacitors	finger separation	d_a	7	μm
	finger width	w_a	7	μm
	finger overlap	L_a	40	μm
	fingers per side	N_a	148	-
	capacitance per side	C_a	0.75	pF
ETA	half length	L	1450	μm
	cross-sectional area	A	1.2×10^{-9}	m^2
	second moment of area	I_z	1.86×10^{-19}	m^4
	clamping angle	θ	2	$^\circ$
	ETAs per corner	-	2	-

A.3 Fabrication details

Section 5.1.1 presents a brief overview of the microfabrication sequence used to manufacture the MEMS seismometer presented in this work. More details on the individual processing steps, the properties of the substrate material and the relevant deep reactive ion etching recipes are included here for reference.

A.3.1 SOI wafer properties

The MEMS seismic sensor is produced in a process that uses a single silicon-on-insulator (SOI) wafer with properties as listed in Table A.3. All silicon is highly doped to obtain a low resistivity substrate material that can serve all the required electrical functions without the need for metallisation. The standard (100) wafer orientation allows for easy etching of the proof mass recess with KOH (step 4 in Sec. A.3.2). Subsequent photolithography on a wafer with such deep recesses (step 7 in Sec. A.3.2) is generally hard, but this anisotropic etch conveniently leaves sloped side walls that can be covered with a photoresist using a spray coater. The buried oxide layer was chosen to be relatively thick, both to allow enough motion for the proof mass and to minimise parasitic capacitance to the handle layer.

Table A.3 – SOI substrate parameters

Property	Value
Orientation	(100)
Diameter	100.0 ± 0.2 mm
Silicon doping	P++, Boron
Device layer thickness	50.0 ± 0.5 μ m
Device layer resistivity	$5 \times 10^{-5} - 2 \times 10^{-4}$ Ω m
Handle layer thickness	400 ± 10 μ m
Handle layer resistivity	$5 \times 10^{-5} - 2 \times 10^{-4}$ Ω m
Buried oxide layer thickness	4μ m \pm 5 %
Supplier	SI-Mat

A.3.2 Full fabrication sequence

The series of images below shows a more detailed overview of the processing sequence used to manufacture the MEMS seismometer presented in this work, as described in Fig. 5.2. The relevant colour coding for the materials used during fabrication is indicated at the top. Microfabrication was done in the MESA+ nanolab, and details on specific machinery or standard cleaning procedures can be found in their equipment database. The process uses a total of 3 photomasks:

1. *KOH*: Defines the oxide mask for etching the proof mass recess with KOH.
2. *HANDLE*: Defines the oxide hard mask for the DRIE step on the back side of the wafer that separates the proof mass from the frame and singulates all devices without dicing.
3. *DEVICE*: Defines the hard mask for the DRIE step on the front side of the wafer that contains all intricate features in the design, such as the capacitors and the suspension springs.




Silicon



Silicon dioxide




Photoresist




1. Substrate selection

- Wafer type: silicon-on-insulator (SOI)
- Device layer: 50 μm
- Buried oxide layer: 4 μm
- Handle layer: 400 μm
- Also see Table A.3




2. Thermal oxidation

- Wafer cleaning
- Wet thermal oxidation in furnace
- Temperature: 1150 $^{\circ}\text{C}$
- Time: 12 h
- Thickness: 2 μm
- Choosing thick oxide layer here to minimise defects resulting from subsequent KOH etch




3. Photolithography

- Spin coating handle layer with photoresist
- Photoresist type: Olin 908-35
- Mask: *KOH*
- Dry etching silicon dioxide
- Machine: Adixen DE
- Depth: 2 μm
- Time: 10 min




4. Etching backside recess

- Stripping photoresist
- Wet etching exposed silicon in handle layer anisotropically in *KOH*
- Depth: 200 μm
- Time: ≈ 6 h
- Doped silicon etches slower than undoped silicon, so use P++ dummy wafers for tuning
- Measuring final recess depth




5. Stripping silicon dioxide

- RCA cleaning
- Oxide removal with wet 50 % HF etch
- Pre-furnace wafer cleaning




6. Thermal oxidation

- Wet thermal oxidation in furnace
- Temperature: 1150 $^{\circ}\text{C}$
- Time: 12 h
- Thickness: 2 μm




7. Photolithography

- Spray coating of photoresist
- Layer thickness: $\approx 5 \mu\text{m}$
- Mask: *HANDLE*
- Exposure time: 16 s
- Resist exposure was optimised for the development of the recessed features
- Features on photomask are biased to obtain correct linewidths in the recess




8. Pattern hard mask

- Dry etching silicon dioxide
- Machine: Adixen DE
- Depth: $2 \mu\text{m}$
- Time: $\approx 10 \text{ min}$




9. Wafer cleaning

- Stripping photoresist




10. Photolithography

- Spin coating device layer with photoresist
- Photoresist type: Olin 907-35
- Mask: *DEVICE*
- Use the multi-zone vacuum chuck with glass window in the mask aligner
- Take special care with the alignment marks
- Vacuum contact: pressure should be better than -0.8 bar




11. Pattern hard mask

- Dry etching of silicon dioxide
- Machine: Adixen DE
- Depth: 2 μm
- Time: ≈ 10 min




12. Wafer cleaning

- Stripping photoresist
- Plasma cleaning in TePla 300E: 45 min
- Stripping in HNO_3
- Standard cleaning




13. Deep reactive ion etching (front side)

- DRIE of silicon on front side
- Machine: Adixen SE
- Recipe: See Table A.5
- Depth: 50 μm
- Time: ≈ 19 min




14. Deep reactive ion etching (back side)

- DRIE etching of silicon on back side
- Machine: Adixen SE
- Recipe: See Table A.5
- Depth: 400 μm
- Time: ≈ 33 min




15. Wafer cleaning

- O_2 plasma cleaning with fluorocarbon removal
- Machine: TePla 360
- Piranha cleaning: 30 min



16. Silicon dioxide removal

- Stripping silicon dioxide in 50 % HF
- Time: ≈ 2 min
- Etch until oxide at both the top and bottom surfaces is removed



17. Vapour-HF release

- Vapour phase HF etch to release narrow silicon structures in the device layer from the handle
- Target underetch: $17\text{ }\mu\text{m}$
- Time: ≈ 60 min
- Temperature: 37°C
- This step singulates all devices without dicing

Table A.5 – DRIE etching recipes on the Adixen AMS 100 SE machine.

Process parameter	Front side	Back side
SF ₆ flow	350 sccm	500 sccm
SF ₆ cycle	3.5 s	6.0 s
C ₄ F ₈ flow	200 sccm	100 sccm
C ₄ F ₈ cycle	1.0 s	1.5 s
Inductively coupled plasma power	1500 W	2500 W
Capacitively coupled plasma power	90 W	60 W
Wafer stage position	200 mm	110 mm
Vacuum valve	100 %	16.5 %
Helium backing pressure	10 mbar	10 mbar
Electrode temperature	−10 °C	−40 °C

A.4 Acronyms

AC	alternating current (in general, signals with $f \neq 0$ Hz)
ADC	analogue-to-digital converter
ASD	amplitude spectral density
ASIC	application-specific integrated circuit
BOX	buried oxide
CMOS	complementary metal-oxide-semiconductor
DC	direct current (in general, signals with $f = 0$ Hz)
DRIE	deep reactive ion etching
ET(-HF/LF)	Einstein Telescope (high-frequency/low-frequency)
ETA	electrothermal actuator
FBA	force balance accelerometer
GPS	global positioning system
GR	general relativity
GW	gravitational wave
KAGRA	Kamioka gravitational wave detector
LIGO	laser interferometer gravitational-wave observatory
LPF	lowpass filter
LVDT	linear variable differential transformer
MEMS	microelectromechanical system
NN	Newtonian noise
PCB	printed circuit board
PECVD	plasma-enhanced chemical vapour deposition
PID	proportional-integral-derivative (controller)
PSD	power spectral density
RIE	reactive ion etching
RMS	root mean square
SEM	scanning electron microscope
SNR	signal-to-noise ratio
SOI	silicon-on-insulator
SPICE	simulation program with integrated circuit emphasis
SQL	standard quantum limit
SR	special relativity
TSV	through-silicon via
TT	transverse traceless
UTC	coordinated universal time

Bibliography

- [1] B. P. Abbott et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". In: *Phys. Rev. Lett.* 116 (6 Feb. 2016), page 061102.
- [2] A. Einstein. "Die Grundlage der allgemeinen Relativitätstheorie". In: *Annalen der Physik* 49.7 (1916), pages 769–822.
- [3] B. Cohen and G. E. Smith, editors. *The Cambridge Companion to Newton (Cambridge Companions to Philosophy)*. Cambridge University Press, 2002. ISBN: 0521656966.
- [4] A. Einstein. *The collected papers of Albert Einstein, Vol. 7: The Berlin Years: Writings, 1918 - 1921*. Edited by M. Janssen et al. Princeton University Press, 2002.
- [5] A. Einstein. "On the electrodynamics of moving bodies". In: *Annalen der Physik* 17.891 (1905), page 50.
- [6] A. Pais. *Subtle is the Lord: The Science and the Life of Albert Einstein*. Oxford University Press, USA, 1982. ISBN: 9780191524028.
- [7] M. Janssen and C. Lehner, editors. *The Cambridge Companion to Einstein. Volume 1*. Cambridge University Press, 2014. ISBN: 9780521828345.
- [8] B. Schutz. *A first course in general relativity*. Cambridge university press, 2009. ISBN: 9780521887052.
- [9] A. Einstein. "Approximative integration of the field equations of gravitation". In: *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* (1916), pages 688–696.

- [10] K. Izumi. *Cavity response in transmission*. <https://gwdoc.icrr.u-tokyo.ac.jp/DocDB/0011/T1201121/002/cavHPF.pdf>. [Online; accessed February 20, 2020]. 2012.
- [11] T. Accadia et al. *Advanced Virgo technical design report*. Virgo note VIR-0128A-12. The Virgo Collaboration, 2012.
- [12] E. Capocasa et al. *Injection system frequency noise budget*. Virgo note VIR-0369A-15. The Virgo Collaboration, 2015.
- [13] A. A. Michelson. "The relative motion of the Earth and the Luminiferous ether". In: *American Journal of Science* 22.128 (1881), page 120.
- [14] A. A. Michelson and E. W. Morley. "On the relative motion of the earth and the luminiferous ether". In: *American Journal of Science* 34.203 (1887), page 333.
- [15] B. J. Meers. "Recycling in laser-interferometric gravitational-wave detectors". In: *Physical Review D* 38.8 (1988), page 2317.
- [16] M. Heurs. "Gravitational wave detection using laser interferometry beyond the standard quantum limit". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376.2120 (2018), page 20170289.
- [17] H. B. Callen and R. F. Greene. "On a theorem of irreversible thermodynamics". In: *Physical Review* 86.5 (1952), page 702.
- [18] P. R. Saulson. "Thermal noise in mechanical experiments". In: *Physical Review D* 42.8 (1990), page 2437.
- [19] F. Acernese et al. "Advanced Virgo: a second-generation interferometric gravitational wave detector". In: *Classical and Quantum Gravity* 32.2 (Dec. 2014), page 024001.
- [20] F. Piergiovanni, M. Punturo, and P. Puppo. *The thermal noise of the Virgo+ and Virgo Advanced Last Stage Suspension (The PPP effect)*. Virgo note VIR-0015E-09. The Virgo Collaboration, 2009.
- [21] A. Heptonstall et al. "Enhanced characteristics of fused silica fibers using laser polishing". In: *Classical and Quantum Gravity* 31.10 (2014), page 105006.
- [22] G. M. Harry et al. "Thermal noise in interferometric gravitational wave detectors due to dielectric optical coatings". In: *Classical and Quantum Gravity* 19.5 (2002), page 897.
- [23] P. Puppo. *Q measurements on payloads with monolithic suspensions*. Virgo note VIR-0796A-18. The Virgo Collaboration, 2018.
- [24] J. Steinlechner. "Development of mirror coatings for gravitational-wave detectors". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376.2120 (2018), page 20170282.

- [25] Y. Aso et al. "Interferometer design of the KAGRA gravitational wave detector". In: *Physical Review D* 88.4 (2013), page 043007.
- [26] J. Peterson. *Observations and modeling of seismic background noise*. Open-file report 93-322. US Geological Survey, 1993.
- [27] I. Fiori. *Reference Seismic Data for Virgo*. Virgo note VIR-0390A-15. The Virgo Collaboration, 2015.
- [28] J. Harms. *Seismic spectral analysis, LHO/LLO LVEA 2009/2010*. LIGO document T1500224-v1. LIGO Scientific Collaboration, 2015.
- [29] S. Braccini et al. "Measurement of the seismic attenuation performance of the VIRGO Superattenuator". In: *Astroparticle Physics* 23.6 (2005), pages 557–565.
- [30] M. Beccaria et al. "Extending the VIRGO gravitational wave detection band down to a few Hz: metal blade springs and magnetic antisprings". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 394.3 (1997), pages 397–408.
- [31] J. van Heijningen et al. "A multistage vibration isolation system for Advanced Virgo suspended optical benches". In: *Classical and Quantum Gravity* (2019).
- [32] M. Blom. "Design and performance of the external injection bench seismic attenuation system EIB-SAS". PhD dissertation. VU University Amsterdam, 2015.
- [33] A. Bertolini et al. "Seismic noise filters, vertical resonance frequency reduction with geometric anti-springs: a feasibility study". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 435.3 (1999), pages 475–483.
- [34] D. Fiorucci et al. "Impact of infrasound atmospheric noise on gravity detectors used for astrophysical and geophysical applications". In: *Physical Review D* 97.6 (2018), page 062003.
- [35] M. G. Beker. "Low-frequency sensitivity of next generation gravitational wave detectors". PhD thesis. PhD dissertation, VU University Amsterdam, 2013.
- [36] M. Vallisneri et al. "The LIGO open science center". In: *Journal of Physics: Conference Series*. Volume 610. 1. IOP Publishing. 2015, page 012021.
- [37] J. Abadie et al. "Search for gravitational waves from low mass compact binary coalescence in LIGO's sixth science run and Virgo's science runs 2 and 3". In: *Physical Review D* 85.8 (2012), page 082002.
- [38] J. Abadie et al. "Search for gravitational waves from compact binary coalescence in LIGO and Virgo data from S5 and VSR1". In: *Physical Review D* 82.10 (2010), page 102001.
- [39] B. P. Abbott et al. "Binary black hole mergers in the first advanced LIGO observing run". In: *Physical Review X* 6.4 (2016), page 041015.

- [40] B. P. Abbott et al. "GW150914: The Advanced LIGO detectors in the era of first discoveries". In: *Physical review letters* 116.13 (2016), page 131103.
- [41] S. M. Koushiappas and A. Loeb. "Maximum redshift of gravitational wave merger events". In: *Physical review letters* 119.22 (2017), page 221104.
- [42] D. McClelland et al. *Instrument science white paper*. Technical note T1600119-v4. LIGO Scientific Collaboration, 2016.
- [43] B. P. Abbott et al. "Properties of the binary black hole merger GW150914". In: *Physical Review Letters* 116.24 (2016), page 241102.
- [44] B. P. Abbott et al. "GW170817: observation of gravitational waves from a binary neutron star inspiral". In: *Physical Review Letters* 119.16 (2017), page 161101.
- [45] B. Abbott et al. "GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs". In: *Physical Review X* 9.3 (2019), page 031040.
- [46] J. Aasi et al. "Parameter estimation for compact binary coalescence signals with the first generation gravitational-wave detector network". In: *Physical Review D* 88.6 (2013), page 062001.
- [47] M. Abernathy et al. *Einstein gravitational wave Telescope conceptual design study*. Technical report ET-0106C-10. EGO, 2011.
- [48] S. Vitale et al. "Use of gravitational waves to probe the formation channels of compact binaries". In: *Classical and Quantum Gravity* 34.3 (Jan. 2017), 03LT01.
- [49] B. P. Abbott et al. "Multi-messenger observations of a binary neutron star merger". In: *Astrophys. J. Lett* 848.2 (2017), page L12.
- [50] A. Kathirgamaraju, D. Giannios, and P. Beniamini. "Observable features of GW170817 kilonova afterglow". In: *Monthly Notices of the Royal Astronomical Society* 487.3 (June 2019), pages 3914–3921. ISSN: 0035-8711.
- [51] L. S. Collaboration et al. "A gravitational-wave standard siren measurement of the Hubble constant". In: *Nature* 551.7678 (2017), pages 85–88.
- [52] B. P. Abbott et al. "Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A". in: *The Astrophysical Journal Letters* 848.2 (2017), page L13.
- [53] M. Punturo and H. Lueck. *Einstein Telescope: letter of intent*. <http://www.et-gw.eu/index.php/letter-of-intent>. [Online; accessed May 10, 2019]. 2018.
- [54] A. Freise. *GWOptics: Einstein Telescope*. <http://www.gwoptics.org/research/et/layout/>. [Online; accessed May 15, 2019]. 2009.
- [55] M. Punturo. *Einstein Telescope: Image Gallery*. <http://www.et-gw.eu/index.php/etimages>. [Online; accessed May 15, 2019]. 2011.

- [56] V. Baibhav et al. "Gravitational-wave detection rates for compact binaries formed in isolation: LIGO/Virgo O3 and beyond". In: *Phys. Rev. D* 100 (6 Sept. 2019), page 064060.
- [57] F. Paoletti et al. *Virgo buildings acoustic and seismic noise status and mitigation ideas*. Virgo note VIR-0674A-18. The Virgo Collaboration, 2018.
- [58] J. Harms. *Newtonian Noise Cancellation*. Virgo note VIR-0333A-19. The Virgo Collaboration, 2019.
- [59] C. Collette et al. "Inertial sensors for low-frequency seismic vibration measurement". In: *Bulletin of the seismological society of America* 102.4 (2012), pages 1289–1300.
- [60] E. Wielandt and P. Bormann. "Seismic sensors and their calibration". In: *New Manual of Seismological Observatory Practices* 1 (2002), page 46.
- [61] E. Wielandt and G. Streckeisen. "The leaf-spring seismometer: Design and performance". In: *Bulletin of the Seismological Society of America* 72.6A (1982), pages 2349–2367.
- [62] H. Xu, X. Liu, and L. Yin. "A Closed-Loop $\Sigma\Delta$ Interface for a High-Q Micromechanical Capacitive Accelerometer With 200 ng/ $\sqrt{\text{Hz}}$ Input Noise Density". In: *IEEE Journal of Solid-State Circuits* 50.9 (2015), pages 2101–2112.
- [63] *Portable very-broad-band tri-axial seismometer: STS-2 low-power manual*. G. Streckeisen AG, Messgeraete. 1995.
- [64] J. van Heijningen, A. Bertolini, and J. van den Brand. "Interferometric readout of a monolithic accelerometer, towards the fm/ $\sqrt{\text{Hz}}$ resolution". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 824 (2016), pages 665–669.
- [65] L. J. LaCoste Jr. "A new type long period vertical seismograph". In: *Physics* 5.7 (1934), pages 178–180.
- [66] S. Braccini et al. "Low noise wideband accelerometer using an inductive displacement sensor". In: *Review of scientific instruments* 66.3 (1995), pages 2672–2676.
- [67] J. Harms and S. Hild. "Passive Newtonian noise suppression for gravitational-wave observatories based on shaping of the local topography". In: *Classical and Quantum Gravity* 31.18 (2014), page 185011.
- [68] L. Naticchioni on behalf of the ET-Sardinia characterization team. *Updates on ET site qualification in Sardinia*. Virgo note VIR-1185A-19. The Virgo Collaboration, 2019.
- [69] S. Koley et al. *First results of seismic studies of the Belgian-Dutch-German site for Einstein Telescope*. Virgo note VIR-1181A-19. The Virgo Collaboration, 2019.

- [70] S. Hild et al. "Sensitivity studies for third-generation gravitational wave observatories". In: *Classical and Quantum Gravity* 28.9 (2011), page 094013.
- [71] P. R. Saulson. "Terrestrial gravitational noise on a gravitational wave antenna". In: *Physical Review D* 30.4 (1984), page 732.
- [72] M. C. Tringali et al. "Seismic array measurements at Virgo's West End Building for the configuration of a Newtonian-noise cancellation system". In: *Classical and Quantum Gravity* (2019).
- [73] F. Badaracco and J. Harms. "Optimization of seismometer arrays for the cancellation of Newtonian noise from seismic body waves". In: *Classical and Quantum Gravity* (2019).
- [74] J. C. Driggers, J. Harms, and R. X. Adhikari. "Subtraction of Newtonian noise using optimized sensor arrays". In: *Physical Review D* 86.10 (2012), page 102001.
- [75] J. van den Brand. "Seismic noise and gravity-gradient noise". In: *Advanced Interferometric Gravitational-Wave Detectors*. Chapter 6, pages 161–184. ISBN: 9789813146075.
- [76] M. Beker et al. "Improving the sensitivity of future GW observatories in the 1–10 Hz band: Newtonian and seismic noise". In: *General Relativity and Gravitation* 43.2 (2011), pages 623–656.
- [77] *Trillium 240 User Guide*. Nanometrics. 2013.
- [78] M. Beker et al. "Innovations in seismic sensors driven by the search for gravitational waves". In: *The Leading Edge* 35.7 (2016), pages 590–593.
- [79] W. Pike et al. "A broad-band silicon microseismometer with 0.25 NG/rtHz performance". In: *2018 IEEE Micro Electro Mechanical Systems (MEMS)*. IEEE. 2018, pages 113–116.
- [80] J. Laine and D. Mougenot. "A high-sensitivity MEMS-based accelerometer". In: *The Leading Edge* 33.11 (2014), pages 1234–1242.
- [81] A. Fougerat, L. Guérineau, and N. Tellier. "High-quality signal recording down to 0.001 Hz with standard MEMS accelerometers". In: *SEG Technical Program Expanded Abstracts 2018*. Society of Exploration Geophysicists, 2018, pages 196–200.
- [82] *INOVA quantum product page*. <https://www.inovageo.com/products/quantum>. Accessed: 2020-01-09.
- [83] D. J. Milligan, B. D. Homeijer, and R. G. Walmsley. "An ultra-low noise MEMS accelerometer for seismic imaging". In: *Sensors, 2011 IEEE*. IEEE. 2011, pages 1281–1284.
- [84] P. Lognonné et al. "SEIS: Overview, Deployment, and First Science on the Ground". In: *Lunar and Planetary Science Conference*. Volume 50. 2019.

- [85] A. Bertolini et al. "Geometric anti-spring vertical accelerometers for seismic monitoring". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 518.1-2 (2004), pages 233–235.
- [86] P. Barriga et al. "Compact vibration isolation and suspension for Australian International Gravitational Observatory: Performance in a 72 m Fabry Perot cavity". In: *Review of Scientific Instruments* 80.11 (2009), page 114501.
- [87] H. Soemers. *Design Principles for precision mechanisms*. T-Pointprint, 2010. ISBN: 9789036531030.
- [88] B. A. Boom et al. "Nano-g accelerometer using geometric anti-springs". In: *Micro Electro Mechanical Systems (MEMS), 2017 IEEE 30th International Conference on*. IEEE, 2017, pages 33–36.
- [89] M. De Laat et al. "A review on in situ stiffness adjustment methods in MEMS". in: *Journal of Micromechanics and Microengineering* 26.6 (2016), page 063001.
- [90] W. C. Young, R. G. Budynas, A. M. Sadegh, et al. *Roark's formulas for stress and strain*. Volume 7. McGraw-Hill New York, 2002. ISBN: 007072542X.
- [91] M. A. Hopcroft, W. D. Nix, and T. W. Kenny. "What is the Young's Modulus of Silicon?" In: *Journal of microelectromechanical systems* 19.2 (2010), pages 229–238.
- [92] D. J. Bell et al. "MEMS actuators and sensors: observations on their performance and selection for purpose". In: *Journal of Micromechanics and Microengineering* 15.7 (2005), S153.
- [93] E. T. Enikov, S. S. Kedar, and K. V. Lazarov. "Analytical model for analysis and design of V-shaped thermal microactuators". In: *Journal of Microelectromechanical Systems* 14.4 (2005), pages 788–798.
- [94] P. H. Pham et al. "Single mask, simple structure micro rotational motor driven by electrostatic comb-drive actuators". In: *Journal of Micromechanics and Microengineering* 22.1 (2011), page 015008.
- [95] Y. Okada and Y. Tokumaru. "Precise determination of lattice parameter and thermal expansion coefficient of silicon between 300 and 1500 K". in: *Journal of applied physics* 56.2 (1984), pages 314–320.
- [96] J. Kamer. *Characterizing the electrical and mechanical behavior of the Electro-Thermal Actuators of the Nikhef accelerometer MEMS design*. Technical report. www.nikhef.nl/pub/services/newbiblio/other/JerryKamer_NikhefInternshipReport_2013.pdf. De Haagse Hogeschool, 2017.
- [97] M. Tilli et al. *Handbook of silicon based MEMS materials and technologies*. William Andrew, 2015. ISBN: 9780323299657.

- [98] D. van Wees. *Proof of principle measurements of a MEMS accelerometer design including anti-spring technology*. Technical report. www.nikhef.nl/pub/services/newbiblio/other/DaveVanWees_BSc_graduation_thesis_2012-12-20.pdf. De Haagse Hogeschool, 2012.
- [99] G. Cella et al. "Monolithic geometric anti-spring blades". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 540.2-3 (2005), pages 502–519.
- [100] T. Namazu and Y. Isono. "Fatigue life prediction criterion for micro–nanoscale single-crystal silicon structures". In: *Journal of Microelectromechanical Systems* 18.1 (2008), pages 129–137.
- [101] D. C. Miller et al. "Characteristics of a commercially available silicon-on-insulator MEMS material". In: *Sensors and Actuators A: Physical* 138.1 (2007), pages 130–144.
- [102] M. Elwenspoek and R. Wierink. *Mechanical microsensors*. Springer-Verlag Berlin Heidelberg GmbH, 2001. ISBN: 9783642087066.
- [103] W.-M. Zhang et al. "Electrostatic pull-in instability in MEMS/NEMS: A review". In: *Sensors and Actuators A: Physical* 214 (2014), pages 187–218.
- [104] R. Middlemiss et al. "Measurement of the Earth tides with a MEMS gravimeter". In: *Nature* 531.7596 (2016), pages 614–617.
- [105] S. Franssila. *Introduction to microfabrication*. John Wiley & Sons, 2010. ISBN: 9780470749838.
- [106] Z. Li et al. "Novel capacitive sensing system design of a microelectromechanical systems accelerometer for gravity measurement applications". In: *Micromachines* 7.9 (2016), page 167.
- [107] W. Pike et al. "Design, fabrication and testing of a micromachined seismometer with NANO-G resolution". In: *Solid-State Sensors, Actuators and Microsystems Conference, 2009. TRANSDUCERS 2009. International*. IEEE. 2009, pages 668–671.
- [108] R. Sleeman, A. Van Wettum, and J. Trampert. "Three-channel correlation analysis: A new technique to measure instrumental noise of digitizers and seismic sensors". In: *Bulletin of the Seismological Society of America* 96.1 (2006), pages 258–271.
- [109] *Trillium Compact User Guide*. Nanometrics. 2019.
- [110] S. Reid et al. "Mechanical dissipation in silicon flexures". In: *Physics Letters A* 351.4-5 (2006), pages 205–211.
- [111] C. J. Glassbrenner and G. A. Slack. "Thermal Conductivity of Silicon and Germanium from 3°K to the Melting Point". In: *Phys. Rev.* 134 (4A May 1964), A1058–A1069.

- [112] D. McGuigan et al. "Measurements of the mechanical Q of single-crystal silicon at low temperatures". In: *Journal of Low Temperature Physics* 30.5 (1978), pages 621–629.
- [113] J. Yang, T. Ono, and M. Esashi. "Energy dissipation in submicrometer thick single-crystal silicon cantilevers". In: *Journal of Microelectromechanical Systems* 11.6 (Dec. 2002), pages 775–783. ISSN: 1057-7157.
- [114] X. Liu et al. "Understanding the internal friction of a silicon micro-mechanical oscillator". In: *Materials Science and Engineering: A* 521 (2009), pages 389–392.
- [115] A. Utz et al. "An ultra-low noise capacitance to voltage converter for sensor applications in 0.35 μm CMOS". in: *Journal of Sensors and Sensor Systems* 6.2 (2017), page 285.
- [116] S. Matteucci et al. "Transport of gases and vapors in glassy and rubbery polymers". In: *Materials science of membranes for gas and vapor separation* 1 (2006), pages 1–2.
- [117] A. Cavalleri et al. "Gas damping force noise on a macroscopic test body in an infinite gas reservoir". In: *Physics Letters A* 374.34 (2010), pages 3365–3369.
- [118] M. Suijlen et al. "Squeeze film damping in the free molecular flow regime with full thermal accommodation". In: *Sensors and Actuators A: Physical* 156.1 (2009), pages 171–179.
- [119] R. Kersevan and J.-L. Pons. "Introduction to MOLFLOW+: New graphical processing unit-based Monte Carlo code for simulating molecular flows and for calculating angular coefficients in the compute unified device architecture environment". In: *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films* 27.4 (2009), pages 1017–1023.
- [120] T. Schimert et al. *Vacuum Packaging for Microelectromechanical Systems (MEMS)*. technical report. RAYTHEON CO DALLAS TX, 2002.
- [121] E. S. Topalli et al. "Pirani vacuum gauges using silicon-on-glass and dissolved-wafer processes for the characterization of MEMS vacuum packaging". In: *IEEE Sensors Journal* 9.3 (2009), pages 263–270.
- [122] D. R. Sparks, S. Massoud-Ansari, and N. Najafi. "Chip-level vacuum packaging of micromachines using nanogetters". In: *IEEE Transactions on Advanced Packaging* 26.3 (2003), pages 277–282.
- [123] F. Laermer and A. Schilp. *Method of anisotropically etching silicon*. Patent 5501893. Bosch GmbH, Mar. 1996.
- [124] K.-S. Chen et al. "Effect of process parameters on the surface morphology and mechanical performance of silicon structures after deep reactive ion etching (DRIE)". in: *Journal of Microelectromechanical Systems* 11.3 (2002), pages 264–275.

- [125] J. Li et al. "Technique for preventing stiction and notching effect on silicon-on-insulator microstructure". In: *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 21.6 (2003), pages 2530–2539.
- [126] T. Overstolz et al. "A clean wafer-scale chip-release process without dicing based on vapor phase etching". In: *17th IEEE International Conference on Micro Electro Mechanical Systems. Maastricht MEMS 2004 Technical Digest*. IEEE. 2004, pages 717–720.
- [127] I. Sari, I. Zeimpekis, and M. Kraft. "A full wafer dicing free dry release process for MEMS devices". In: *Procedia Engineering* 5 (2010), pages 850–853.
- [128] N. Tas et al. "Stiction in surface micromachining". In: *Journal of Micromechanics and Microengineering* 6.4 (1999), page 385.
- [129] A. Bonucci et al. "A new model for vacuum quality and lifetime prediction in hermetic vacuum bonded MEMS". in: *Reliability, Packaging, Testing, and Characterization of MEMS/MOEMS VII*. volume 6884. International Society for Optics and Photonics. 2008, page 68840M.
- [130] M. H. Asadian, S. Askari, and A. M. Shkel. "An ultrahigh vacuum packaging process demonstrating over 2 million Q-factor in MEMS vibratory gyroscopes". In: *IEEE sensors letters* 1.6 (2017), pages 1–4.
- [131] G. Longoni et al. "Stable and reliable Q-factor in resonant MEMS with getter film". In: *2006 IEEE International Reliability Physics Symposium Proceedings*. IEEE. 2006, pages 416–420.
- [132] F. Chen, X. Li, and M. Kraft. "Electromechanical Sigma-Delta Modulators ($\Sigma\Delta$) Force Feedback Interfaces for Capacitive MEMS Inertial Sensors: A Review". In: *IEEE Sensors Journal* 16.17 (2016), pages 6476–6495.



Summary

The first direct detection of gravitational waves in September 2015 signified both the culmination of a century long quest, and the beginning of an entirely new chapter in our exploration of the Universe. Mankind has relied solely on electromagnetic radiation for his understanding of the Universe for centuries, and with the addition of gravitational waves to the arsenal of observable messengers from the cosmos, there are numerous new things to be learned. On their own, the waves tell us a great deal about the properties of the source that emitted them and its surroundings, but they can also augment electromagnetic observations by providing additional information on the sources of particular cosmic events. With gravitational waves we can directly probe the most dense and dynamic regions of the Universe, and finally look straight into the sources that are driving the violent events in the Universe we have been observing electromagnetically for centuries.

Over the coming few years, the currently operating gravitational wave detectors will be brought to their design sensitivities, and new detectors with comparable detection characteristics will come online in both Japan and India which will increase the sky coverage of the Advanced detector network. This detector network would be able to observe gravitational waves coming from sources with a total mass similar to the first one detected in 2015 up to a redshift of approximately 1. In order to study the early Universe, this observation horizon needs to be extended significantly. One of the proposed next genera-

tion detector infrastructures, Einstein Telescope, will expand the observation horizon for sources of similar total mass to a redshift of approximately 50. This increase in detection range is only possible when the detector's bandwidth is extended to lower frequencies. Einstein Telescope will be sensitive to gravitational wave signals down to approximately 2 Hz. For Earth-based detectors such as Einstein Telescope, the detector noise at these low frequencies will be limited by so-called Newtonian noise, which is caused by fluctuations in the local gravitational field induced by seismic motion in the proximity of the detector. In order to reach the design sensitivity, active cancellation of this Newtonian noise needs to be implemented, requiring large arrays of low-noise seismic sensors.

MEMS sensors are interesting for applications that require large numbers of devices, because of the inherent batch-processing nature of their manufacturing. MEMS based seismometers are usually not considered for low-noise applications, because of two main reasons. Firstly, the seismometer's proof mass is necessarily small because of the limited size of MEMS devices. This makes that viscous damping effects caused by gases surrounding the proof mass introduce a relatively large Brownian noise. Secondly, the stiffness of the suspension springs can usually not be made small enough to obtain sufficient sensitivity to acceleration in these tiny devices. The first problem can be tackled by packaging the sensor mechanics inside a vacuum, which significantly reduces the Brownian noise level. The relatively high suspension spring stiffness, however, usually remains.

This dissertation presents a novel technique for increasing the sensitivity of MEMS seismometers to acceleration by reducing their suspension spring stiffness. This stiffness reduction is achieved by stressing the springs with a specially designed compression mechanism. The idea was inspired by the geometric anti-spring filters that are used in the auxiliary seismic isolation suspensions of the Advanced Virgo gravitational wave detector. There, the concept is used to obtain efficient vertical seismic filters with a low natural frequency. Its application to MEMS seismometers is completely new, and the implementation presented here has successfully demonstrated the ability to reduce the initial suspension spring stiffness by a factor 20 to 25. This large stiffness reduction occurs in the sensing direction of the single-axis seismometer, while only marginally affecting the stiffness in the other directions. Because a dimensionless design method was used, the results presented in this work can be used to assess the feasibility of an implementation of such a geometric anti-spring system at any scale.

The reduction of the suspension spring stiffness increases the attainable sensitivity of MEMS seismometers, and a noise performance below the seismic background requirement of $2 \text{ ng}/\sqrt{\text{Hz}}$ set for the future Einstein Telescope host site was demonstrated using this novel technique. For a sensor that only measures roughly a square centimetre, this level of noise performance is unprecedented. The acceleration noise floor of the entire sensor system operating as a force balance accelerometer lies below $2 \text{ ng}/\sqrt{\text{Hz}}$ between 400 mHz and 20 Hz, and remains below $10 \text{ ng}/\sqrt{\text{Hz}}$ all the way down to approximately 50 mHz. This measurement matches the noise models developed in this dissertation to better than 20 % over the entire measured frequency range. Due to the reduction of the suspension spring stiffness through the geometric anti-spring effect, the noise contribution of the capacitive readout and feedback circuitry is completely negligible below 20 Hz, and the noise floor there is fully dominated by thermomechanical noise in the MEMS mechanics itself.

The mechanical design of a MEMS device is always closely interwoven with its fabrication process. Although many of the individual processing steps can be considered standard practice in a semiconductor fabrication environment, the process as a whole is unique for almost every MEMS device and requires careful design. The seismometer presented in this dissertation was made using an elementary fabrication process that uses a single silicon-on-insulator wafer with a $50 \text{ }\mu\text{m}$ thick device layer. Etching tests show that it is feasible to move to wafers with a $100 \text{ }\mu\text{m}$ thick device layer, which has significant advantages for the mechanical robustness of the device. A more elaborate fabrication process was designed to overcome some of the shortcomings of the elementary prototype process, mostly related to process yield and device robustness. It uses a wafer bonding step to be able to tailor the cavity surrounding the seismometer's proof mass and minimise the potential contact area to all the moving parts in the device. Measurements show that in order to obtain its good noise performance, the MEMS seismometer needs to be operated in a vacuum environment better than 10^{-2} mbar . To obtain a practically useful sensor, this vacuum encapsulation needs to be incorporated as part of the fabrication sequence, either on the single-die-level or the wafer-level. The new process is fully compatible with either of those vacuum encapsulation methods that will have to be implemented in the future.



Samenvatting

De eerste directe detectie van zwaartekrachtsgolven in september 2015 markeerde niet alleen het hoogtepunt van een honderd jaar lange zoektocht, maar ook het begin van een geheel nieuw hoofdstuk in onze studie van het universum. De mensheid was eeuwenlang volledig afhankelijk van waarnemingen van elektromagnetische aard voor haar begrip van het universum, en met de toevoeging van zwaartekrachtsgolven aan het arsenaal van waarneembare signalen uit het heelal liggen er talrijke nieuwe ontdekkingen in het verschiet. Op zichzelf kunnen deze zwaartekrachtsgolven ons al veel vertellen over de eigenschappen van de bron die ze uitzond, maar ze kunnen ook elektromagnetische waarnemingen verrijken door extra inzicht te verschaffen in de achterliggende fysische processen die aan bepaalde gebeurtenissen ten grondslag liggen. Met zwaartekrachtsgolven kunnen we de meest compacte en dynamische regio's van het heelal onderzoeken, en eindelijk direct in de bronnen kijken die de drijvende kracht zijn achter de meest destructieve gebeurtenissen in het heelal die we al eeuwenlang elektromagnetisch bestuderen.

In de loop van de komende paar jaren zullen de momenteel operationele detectoren naar hun beoogde gevoeligheid worden gebracht en zullen er twee nieuwe detectoren met vergelijkbare eigenschappen aan het netwerk worden toegevoegd, één in India en één in Japan. Deze nieuwe detectoren zullen de dekkingsgraad van het netwerk verbeteren, zodat zwaartekrachtsgolven uit alle richtingen beter kunnen worden waargenomen. Dit detectornetwerk kan

signalen van bronnen vergelijkbaar met de eerste uit 2015 waarnemen tot een roodverschuiving van ongeveer 1. Om ook zwaartekrachtsgolven uit het vroege Universum op te kunnen vangen, moet dit bereik nog significant worden vergroot. Een van de geplande toekomstige detectoren, Einstein Telescope, kan een dergelijke bron van zwaartekrachtsgolven waarnemen tot een roodverschuiving van ongeveer 50. Deze vergroting van het detectiebereik is alleen mogelijk als de detectieband wordt uitgebreid naar lagere frequenties. Einstein Telescope zal dan ook gevoelig zijn voor zwaartekrachtsgolven tot een frequentie van ongeveer 2 Hz. Voor detectoren op aarde zoals Einstein Telescope wordt de detectorruis bij zulke lage frequenties gedomineerd door zogenaamde Newtoniaanse ruis, afkomstig van kleine fluctuaties in het lokale zwaartekrachtsveld veroorzaakt door seismische trillingen in de buurt van de detector. Om de beoogde gevoeligheid te bereiken, moet deze vorm van ruis actief worden onderdrukt en hiervoor zal een netwerk van een groot aantal seismische sensoren nodig zijn.

MEMS sensoren zijn interessant voor toepassingen waarbij grote aantallen sensoren nodig zijn, omdat de mogelijkheid tot massafabricage inherent is aan de technieken die worden gebruikt ze te produceren. Toch worden seismische MEMS sensoren meestal niet gebruikt voor toepassingen waarbij een lage ruisvloer van belang is. Hiervoor zijn twee voorname redenen te noemen. In de eerste plaats is de beschikbare massa in dit soort sensoren extreem beperkt door hun kleine formaat. Hierdoor is de mechanische ruis die wordt veroorzaakt door visceuze dampingeffecten in het omringende gas relatief groot. Ten tweede kan de stijfheid van de veerophanging op deze kleine schaal gewoonlijk niet laag genoeg worden gemaakt om de sensor voldoende gevoeligheid voor versnelling te geven. Het eerste probleem kan worden opgelost door de bewegende delen van de sensor in een vacuüm te gebruiken, waarbij de bovengenoemde mechanische ruis sterk kan worden onderdrukt. De relatief hoge stijfheid van de veerophanging blijft doorgaans echter een probleem.

Deze dissertatie presenteert een nieuwe techniek om de gevoeligheid van seismische MEMS sensoren voor versnelling te verhogen door de stijfheid van de veerophanging te verlagen. Deze verlaging wordt bewerkstelligd door de veerophanging voor te spannen met een speciaal daarvoor ontwikkeld compressiemechanisme. De oorsprong van dit idee kan worden teruggevoerd op de geometrische antiveerophangingen die worden gebruikt in de secundaire seismische isolatoren van de zwaartekrachtsgolfdetector Advanced Virgo. In die isolatoren wordt het concept van de geometrische antiveer gebruikt om een efficiënt en compact verticaal seismisch filter met een lage eigenfrien-

tie te realiseren. De toepassing van dit principe op de microschaal in seismische MEMS sensoren is volledig nieuw, en de implementatie die hier wordt beschreven heeft voor het eerst op deze schaal succesvol een mechanische stijfheidsreductie van de veerophanging met een factor 20 tot 25 gedemonstreerd. Deze grote stijfheidsreductie treedt alleen op in de gevoelige richting van de seismische sensor, in de andere richtingen verandert de stijfheid marginaal. Omdat er een dimensieloze ontwerpmethode is gebruikt, kunnen de resultaten die hier zijn gepresenteerd algemeen worden gebruikt om een geometrisch anti-veersysteem te ontwerpen op elke gewenste schaal.

De sterk gereduceerde stijfheid van de veerophanging in de sensor vergroot de gevoeligheid die met seismische MEMS sensoren te behalen is. Met gebruik van deze nieuwe techniek is een ruisniveau van onder de $2 \text{ ng}/\sqrt{\text{Hz}}$ gedemonstreerd, de grens die is bepaald voor het achtergrondniveau van seismische activiteit voor de toekomstige locatie van Einstein Telescope. Een dergelijk ruisniveau is ongeëvenaard voor een sensor die slechts een formaat heeft van ruwweg een vierkante centimeter. De gemeten ruisvloer voor het gehele sensorsysteem ligt beneden de $2 \text{ ng}/\sqrt{\text{Hz}}$ tussen 400 mHz en 20 Hz, en blijft beneden de $10 \text{ ng}/\sqrt{\text{Hz}}$ voor frequenties tot ongeveer 50 mHz. Dit komt tot binnen 20 % overeen met de ruismodellen die in deze dissertatie zijn ontwikkeld over het gehele gemeten frequentiebereik. Door de stijfheidsverlaging geïntroduceerd door het geometrische anti-veersysteem is beneden de 20 Hz de ruisbijdrage van de capacitieve uitlezings- en terugkoppelingselektronica volledig verwaarloosbaar. De ruisvloer bij die frequenties wordt volledig gedomineerd door thermische ruis in de mechanica van de MEMS sensor.

Het mechanische ontwerp van een MEMS sensor is altijd sterk verweven met zijn fabricageproces. Hoewel veel van de individuele processtappen als standaard kunnen worden beschouwd in de halfgeleiderindustrie, is dat proces als geheel vrijwel altijd uniek voor een bepaald sensorontwerp. De seismische sensor die in deze dissertatie wordt gepresenteerd is geproduceerd in een proces dat één enkele SOI wafer gebruikt met een toplaag van $50 \mu\text{m}$ dik. Resultaten van etsexperimenten tonen aan dat het mogelijk is om de sensoren ook in een $100 \mu\text{m}$ dikke toplaag te realiseren, iets dat grote voordelen heeft voor de mechanische robuustheid van de mechanische structuren. Om een aantal tekortkomingen van het elementaire proces te overkomen, voornamelijk gerelateerd aan de procesopbrengst en de mechanische robuustheid van de sensoren, is er een meer uitgebreid fabricageproces ontwikkeld. Dit nieuwe proces volgt een andere aanpak, waarbij twee aparte wafers aan elkaar worden verbonden. Dit stelt ons in staat het vlak waar de twee wafers samenkomen te be-

werken en de mogelijke raakvlakken met de bewegende delen in de sensor te minimaliseren. Metingen aan de dempingseigenschappen van de sensorstructuren laten zien dat de sensor moet opereren in een vacuüm dat beter is dan 10^{-2} mbar om het lage ruisniveau te behalen. Wil de sensor praktisch zijn in gebruik, dan moet dit vacuüm als onderdeel van het fabricageproces worden gerealiseerd, dan wel door het vacuüm verpakken van de individuele sensoren, dan wel door op waferniveau afgesloten holtes in vacuüm te realiseren. Het nieuwe procesontwerp is compatibel met beide aanpakken, waarvan er in de toekomst een aan het proces zal moeten worden toegevoegd.



Acknowledgements

My work at Nikhef during the past four and a half years has led to the result you are currently holding in your hands, and although I will be the one defending this dissertation, there are a lot of people that have either directly or indirectly contributed to it and without whom it could not have been written. First of all, I would like to extend my gratitude to all my colleagues in the gravitational wave and detector R&D groups at Nikhef for making my time there more than enjoyable. Then, fully aware that the list of people I am indebted to is a lot longer than the one I am about to enumerate, I would like to take the opportunity to mention some specific individuals anyway.

First and foremost I would like to thank my doctoral advisor Jo van den Brand for making me a part of the LIGO Virgo collaboration, and for giving me the opportunity to work in the field of gravitational wave research during what I think must be the most exciting period in its existence. Somehow you always manage to ask the right (and hardest) questions to push me to get to the bottom of something and do things right.

Alessandro, you seem to know everything there is to know about experimental practices, seismic sensors, instrumentation in general, measurement techniques, electronics, noise... and the list goes on for a while. Not only that, but you were always available to share that knowledge with me. Whatever your schedule, you always took the time to discuss ideas and experiments, and I probably would not have been able to do half the things presented in this

dissertation without your seemingly endless practical knowledge and patient advice. I also appreciate you (together with Matteo and Soumen) lugging what must have been well over 100 kg worth of measurement equipment to the Heimansgroeve through muddy, hilly fields in between a herd of cows (and only telling me you had accidentally dropped my delicate sensors on the way *after* everything turned out to still function properly...)

Eric, you are a humble man by nature, but you deserve a lot of credit for coming up with many of the ideas that ultimately led to this dissertation. I have really enjoyed our trip to a conference in the beautiful Lake District in the UK and your input on both numerical and analytical modelling has been invaluable. Thanks for your critical scrutiny of Chapter 3, the final version has benefited significantly from it.

Remco, I probably would not have pursued a PhD if you had not introduced me to this specific position during my master's project at the University of Twente. In that sense, you are the one responsible for my doctoral title, thanks! Your input made the 2017 IEEE MEMS conference a pleasant experience. Not only by commenting on both my manuscript and presentation, but more importantly by the fact that one can count on you to find the only affordable craft beer bar in the entirety of Las Vegas, *and* be aware of their happy hour schedule.

Niels, thank you for always being available for a chat or advice. You are a great listener and a pleasure to work with in general, I am glad to have you in my doctorate committee.

To my office mates Maria and Laura: thanks for all the good times both in and outside the office. Your presence made the trips to Elba, Hanover, Pisa and Benasque a lot more fun. Ka Wa, thanks for the good times during the summer school in Spain, I hope I will be able to beat you in a game of table tennis just once. Thanks as well for lending me your copy of Maggiore (which, of course, I returned much too late, sorry for that). Anuradha, thanks for generating the example gravitational waveform data for Fig. 1.12 (and for putting up with my ignorance regarding data analysis and its related terminology). Joris, fellow native of the most beautiful city in the world, thanks for acquainting me with the MultiSAS prototype at Nikhef. It has been an invaluable tool for a lot of my measurements as well as a ticket to be able to work on both the MultiSAS systems and the external injection bench at the Virgo site.

Kees, Robert and Albert, thanks for all your microfabrication work. Without you there would not have been any MEMS devices for me to work with, and certainly a lot less nice microscopic pictures as well. Moreover, thanks for all

the discussions related to MEMS fabrication: most of the things I know on the subject, I know from you. Pino, thanks for your help with operating the laser doppler vibrometer, and Tim, thanks for taking the time for our debugging sessions at the MultiSAS. I would also like to thank KNMI in general and Reinoud Sleeman in particular for kindly granting me access to the seismic station in the Heimansgroeve, without which I would have never been able to measure the low-frequency performance of our sensors.

Another heartfelt thank you goes out to my parents, Dirk and Mariëtte for encouraging and enabling my curious nature, and for supporting me through what has in fact been twenty-three years of education. I literally would not be here without you.

Then, last but certainly not least, I would like to express my gratitude and love towards the person I hold most dear. Linda, thank you for your never-ending love and support. You are the one that makes me happy. Also, you seem to have the special talent to be able to explain my work to others better than I can. For this reason, the book you are currently holding is dedicated to you, hoping my attempt of explaining it to the world will be as good as yours.

Boris Boom, Amsterdam, March 2020

