



ETH zürich

Pascal Engeler

A silicon nitride based synthetic metamaterial

DISS. ETH NO. 30231

DISS. ETH NO. 30231

**A SILICON NITRIDE BASED
SYNTHETIC METAMATERIAL**

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES

(Dr. sc. ETH Zurich)

Presented by

PASCAL ENGELER

M. sc. ETH Physics

born on 21.04.1992

accepted on the recommendation of

Prof. Dr. Sebastian D. Huber, examiner

Prof. Dr. Manfred Sigrist, co-examiner

Dr. Marc Serra-Garcia, co-examiner

Prof. Dr. Ewold Verhagen, co-examiner

On the cover: Microscope photograph of a silicon nitride resonator with gold electrodes, covered in a layer of silicon dioxide. ©Pascal Engeler

Abstract

The recent transfer of topological states from the quantum realm to classical physics has sparked ideas for new approaches in designer materials. These unusually robust configurations were no longer out of reach for classical material design. However, while it is relatively easy to identify a topological state when it appears, the inverse statement is not true. It is hard to find a material structure that will host a specific topological state, which complicates the development of novel materials.

What is more, there are experimental challenges to current metamaterial development platforms. These implementations are typically untunable, and require a complete redesign and remanufacturing for minor changes to be applied. While there exist implementations that are reconfigurable to some extent, adjustments are typically cumbersome and require for the device to be taken offline for an extended period of time. These limitations lead to long iteration times and manufacturing overhead that keeps stacking up.

The above problems are addressed in this thesis in a two pronged approach from the theoretical and the experimental side. To overcome the experimental challenges, a synthetic metamaterial is devised. The goal is a material development platform that can host large hopping models, allows real-time control over all system parameters and is fully reprogrammable by the press of a button via software. This platform only becomes a material with specific functionality when it is powered up and programmed, hence the nomenclature “synthetic”. The system as a whole consists of a sample manufactured using standard microfabrication techniques, and custom designed surrounding infrastructure, with 11 FPGAs at its core, that provides the desired level of control.

The theoretical design issues are approached via the development of a

material design framework. The resulting software stack is able to automatically discover 2 dimensional structures that host topological states. It is also platform agnostic, and works whether the target system is governed by the Schrödinger (electrons, ultra-cold atoms), the Poisson (vibrations) or the Maxwell equations (photonic crystals). The theoretical platform works, and is successfully cranking out structures.

The results found for the experimental platform are very promising. Although it has not quite arrived at the desired final state yet, there exists a clear and short path towards a fully functional system. Furthermore, being a fully parametric system, it is then also suited for the studies of other fields, such as physical neural networks, signal processing and nonlinear phenomena.

Kurzfassung

Die jüngste Übertragung topologischer Zustände aus dem Quantenbereich in die klassische Physik hat Ideen für neue Ansätze bei Designermaterialien hervorgebracht. Diese ungewöhnlich robusten Konfigurationen waren für das klassische Materialdesign nicht länger unerreichbar. Während es jedoch relativ einfach ist, einen topologischen Zustand zu identifizieren, wenn dieser auftritt, trifft die umgekehrte Aussage nicht zu. Es ist schwierig, eine Materialstruktur zu finden, die einen bestimmten topologischen Zustand beherbergt, was die Entwicklung neuartiger Materialien erschwert.

Darüber hinaus gibt es experimentelle Herausforderungen bei den derzeitigen Entwicklungsplattformen für Metamaterialien. Diese Implementierungen sind in der Regel nicht konfigurierbar und erfordern ein komplettes Redesign und eine Neuanfertigung, um kleinere Änderungen vornehmen zu können. Es gibt zwar Implementierungen, die bis zu einem gewissen Grad rekonfigurierbar sind, doch sind Anpassungen in der Regel umständlich und erfordern, dass das Gerät für einen längeren Zeitraum offline genommen wird. Diese Einschränkungen führen zu langen Iterationszeiten und zu einem immer grösser werdenden Fertigungsaufwand.

Die oben genannten Probleme werden in dieser Arbeit in einem zweigleisigen Ansatz von theoretischer und experimenteller Seite angegangen. Um die experimentellen Herausforderungen zu überwinden, wird ein synthetisches Metamaterial entwickelt. Das Ziel ist eine Materialentwicklungsplattform, die grosse Hüpfmodelle beherbergen kann, Echtzeitkontrolle über alle Systemparameter ermöglicht und die per Knopfdruck über Software vollständig umprogrammierbar ist.

Diese Plattform wird erst dann zu einem Material mit spezifischen Funktionen, wenn sie mit Strom versorgt und programmiert wird, daher

die Bezeichnung "synthetisch". Das System als Ganzes besteht aus einem Probenkörper, der mit Standard-Mikrofabrikationstechniken hergestellt wird, und einer massgeschneiderten Infrastruktur mit 11 FPGAs als Herzstück, die das gewünschte Mass an Kontrolle bietet.

Die theoretischen Designfragen werden durch die Entwicklung eines Frameworks für das Materialdesign angegangen. Der resultierende Softwarestack ist in der Lage, automatisch zweidimensionale Strukturen zu entdecken, die topologische Zustände beherbergen. Ausserdem ist sie plattformunabhängig und funktioniert unabhängig davon, ob das Zielsystem durch die Schrödinger-Gleichung (Elektronen, ultrakalte Atome), die Poisson-Gleichung (Schwingungen) oder die Maxwell-Gleichungen (photonische Kristalle) beherrscht wird. Die theoretische Plattform funktioniert und bringt erfolgreich Strukturen hervor.

Die Ergebnisse der experimentellen Plattform sind sehr vielversprechend. Obwohl sie noch nicht ganz den gewünschten Endzustand erreicht hat, ist der Weg zu einem voll funktionsfähigen System klar und kurz. Da es sich um ein voll parametrisches System handelt, eignet es sich auch für die Untersuchung anderer Bereiche, wie z. B. physikalische neuronale Netze, Signalverarbeitung und nichtlineare Phänomene.

Contents

List of Figures	vii
1 Introduction	1
1.1 Historical background	2
1.2 The second topological revolution	4
1.3 Going classical: mechanical degrees of freedom	7
1.4 Towards scalability	8
1.5 The future: tunable metamaterials	11
1.6 Target applications	12
1.7 Author's contributions	14
2 Sample	17
2.1 Large scale: geometry	17
2.2 Small scale: resonators	21
2.3 Electrostatic considerations	29
3 Fabrication	39
3.1 General Remarks	39
3.2 Fabrication Process	42
4 Setup	47
4.1 Physical support components	51
4.1.1 Vacuum chamber	51
4.1.2 Vacuum pump	53
4.2 Measurement System	53
4.2.1 Displacement Measuring Interferometer IDS3010	53
4.2.2 Linear Stages Zaber LSQ150	54

4.2.3 Digital Microscopes	54
4.3 Breakoutboard	55
4.4 RBComb API	57
4.4.1 Class SerialBridge	58
4.4.2 Class LaserInterface	59
4.4.3 Class RBComb576	59
4.5 Bridge	62
4.5.1 Hardware	63
4.5.2 Gateware	65
4.6 Minicircuit	75
4.6.1 Hardware	75
4.6.2 Gateware	78
4.7 Environment control system	85
4.7.1 Grafana	86
4.7.2 Pressure	87
4.7.3 Temperature	87
5 Measurements	97
5.1 Single drum	98
5.2 Interactions	103
6 Structure Search Design Suite	113
7 Conclusion and outlook	123
7.1 Conclusion	123
7.2 Outlook	124
Bibliography	127
List of Publications	137
Curriculum Vitae	139
Acknowledgments	141

List of Figures

2.1	Render of the final sample design	19
2.2	Initial design of the adhesive structuring mask.	20
2.3	Stress release simulation results.	25
2.4	Wafer bonded sample	26
2.5	Transverse deformation of different resonator designs.	28
2.6	Mode shapes and frequencies.	29
2.7	Closeup render of single resonator.	30
2.8	Closeup sample render of fishing holes.	32
2.9	Frequency tuning and displacement computation results.	34
2.10	Render of sample design, view over surface.	37
3.2	Microfabrication process sequence.	42
3.1	Plasma damage on gold layer.	43
3.3	Impression of finished sample.	44
3.4	Impression of finished sample.	45
3.5	Impression of finished sample.	46
4.1	Diagram of the experimental system with different subsystems coded in different colours.	49
4.2	Front view of the setup.	52
4.3	Photograph of the Breakoutboard.	56
4.4	Photograph of the Bridge board.	64
4.5	Register transfer level schematic of the top level design pro- grammed on the Bridge FPGA.	67
4.6	Register-transfer level schematic of the LaserInterfaceWith- Carli component.	71

4.7	Finite state machine diagram for SLD3.	74
4.8	Top view of a minicircuit board.	76
4.9	Top level design of minicircuit gateway.	78
4.10	Register-transfer level schematic of the Datapath block.	82
4.11	Register-transfer level schematic of the Transfer Engine block.	85
4.12	Screenshot of the Grafana dashboard visualizing parameters logged in the InfluxDB database.	86
4.13	Connectivity diagram of the temperature stabilization unit.	88
4.14	Photo of a S I C C board.	89
4.15	Photos of TSU components.	91
4.16	TSU circuit schematics.	92
5.1	Interferometer signal quality scan.	99
5.2	Resonance with phase shift.	100
5.3	Frequency and quality factor histograms.	101
5.4	Frequency map of bottom layer.	102
5.5	Resonator response to resonant coupling drive.	104
5.6	Calculated coupled resonator spectra as function of coupling detuning and strength.	106
5.7	Simulated lock-in measurements of driven two-resonator sys- tem.	108
5.8	Modulated coupling experiment results.	110
6.1	Complete catalog of samples with fragile bands.	114
6.2	Statistical analysis of fragile bands.	117
6.3	Gap optimization of a $p4mm$ sample.	120
6.4	Trivializing distance.	122

1 | Introduction

This dissertation's main thread develops experimental and software means for the design of complex novel materials. Interesting, new material properties are often found through the discovery of unconventional phases of matter. With the evolution of topological states from quantum to classical systems, a whole new class of properties has become available for macroscopic materials. However, finding a specific structure that hosts the targeted properties is hard in general: the search space is intractably large, the problem underconstrained, and conventional optimization algorithms often inappropriate. Furthermore, current experimental platforms for the implementation of complex materials typically provide no flexibility, which results in large design and manufacturing overhead. This chapter introduces the above situation historically, contextualizes the software suite, and motivates the design of the newly developed experimental platform. At the end, several examples of experiments to be conducted in the new system are provided.

1.1 Historical background

The objects of study of condensed matter physics are properties of matter. Matter can be constructed from a variety of constituents, which in turn can be arranged in a multitude of ways, and the resulting system can be subjected to an infinite variety of distinct external conditions. Each combination of these will potentially yield objects with distinct properties, leading to a seemingly intractably complicated zoo of objects. The ingenious idea that tames this monster is the concept of equivalence [1]. Instead of trying to precisely identify every possible property of each of the infinite number of objects, one asks more general questions: Can the constituents move freely, or are they bound to each other? Are the constituents organized in a specific pattern? Does the system generate a finite magnetic field? Does the system conduct electric current? The answers to such questions allow the classification of condensed matter systems into different equivalence classes called phases.

One can learn about properties of matter by studying a finite set of phases. If two systems are in the same phase, for example they both don't conduct electric current, they are in some way equivalent to each other and will have similar properties. The relevant equivalence relation at play is driven by the concept of adiabaticity. Systems that are in the same phase can be adiabatically deformed into each other, and this deformation is bijective at every step. Two systems in different phases can still be transformed into one another, but during this process their low lying states will encounter a violent rearrangement called quantum phase transition, and bijectivity breaks down [1]. The task of condensed matter physics is thus reduced to the identification of phases, and the classification of different states of matter into those phases.

Some phases seem trivial to identify. A solid is obviously different from a gas once we hold it in our hand, an insulator is obviously different from a metal once we look at the bandstructure. For every solid, whether it's a block of ice or a slab of gold, the relative arrangement of the constituents is fixed. For every crystalline insulator, whether it's silicon nitride or carbon, there will be a finite gap between valence and conduction band. More precisely, all – potentially very different – systems that are in the same phase will share some property equal for all of them. This property remains the same as long as the system does not undergo a phase transition.

In general, the above perspicuity of a phase does not exist. It is experience that separates the obvious from the obscure. Discoveries of new phases of matter are typically led by the observation of curious phenomena, which can only be satisfactorily explained by identifying the right parameters to describe the situation — such as the bandgap when distinguishing crystalline insulators from metals.

The most prominent example of phase discovery was led by the observation of the integer quantum Hall effect. Quantum Hall experiments consist of a two dimensional electron gas (2DEG) penetrated by a magnetic field, across which a voltage is applied to drive a current. In its essence, this is the study of Landau levels in a confining potential. Within a classical picture, the flowing electrons are subject to a Lorentz force, which is balanced by charge accumulation on the sides of the 2DEG, giving rise to an associated Hall resistivity. In 1980, von Klitzing et al. [2] performed experiments where they found Hall states with conductivity exactly quantized at integer multiples of $\sigma_{xy} = e^2/h$, while the longitudinal conductivity was measured to be zero. Transitions between different Hall states showed finite longitudinal conductivity.

The curious feature of this result is its universality. The quantization value of σ_{xy} is largely independent of experimental details such as sample materials, electrode placement and impurities. In all experiments, it appears to take the same value to a precision exceeding 1 part per billion. It is this universality that suggests there is something more robust at play than merely the accidental combination of constituent properties. This observation was clarified by Thouless et al. [3] expressing the Hall conductivity $\sigma_{xy} = \nu e^2/h$ in terms of a topological invariant, the TKNN integer ν , that classifies Hall states as a new phase of matter: the topological insulator.

The physics of the quantum Hall states is neither that of a metal, nor that of an insulator. In the bulk, the band structure is gapped and we observe an insulator. Along the edge of the system there are chiral channels, states that cross the gap and support the unidirectional transport of electrons. The TKNN integer describes how many such channels exist, and each of them contributes e^2/h to the conductivity of the edge.

The robustness of the effect is due to its relation to topology, which was later clarified by Simon [4] who identified the TKNN integer as a Chern number by connecting the findings of Thouless to Berry's discovery of

geometric phases [5]. In topology, the genus, the "number of holes", of a closed surface is a topologically stable property unchanged by continuous deformations, and can be calculated as the integral of the Gauss curvature. In close analogy, the Chern number is a topological invariant calculated as the integral of the Berry curvature over the first Brillouin Zone (BZ). While changing a sphere into a torus requires ripping a hole into the surface, the analog statement for quantum Hall states is that a change in Chern number necessarily requires the band gap in the bulk to close, which results in the above described metallic transition. Conversely, any perturbation that does not close the bulk gap will preserve the quantum Hall state.

So far, the quantum Hall effect appears to hinge on the magnetic field generating Landau levels. Haldane realized in 1988 [6] that an external magnetic field is not necessary to achieve quantum Hall states. He wrote down a lattice model hosting a staggered magnetic field with zero net flux, commensurate with the lattice and thus realizable in a magnetic material, that still exhibits quantum Hall states. The essential property giving rise to quantum Hall states was then no longer the magnetic field, but broken time reversal symmetry. There indeed is no way to achieve chiral edge states without breaking time reversal symmetry, because the edge states themselves break it.

1.2 The second topological revolution

In 2005, Kane and Mele [7] found a topological insulator under the name of Quantum Spin Hall Effect (QSHE), that is protected by time reversal symmetry rather than breaking it. While it still produces the general phenomenology of an insulating bulk with a conducting edge, the physics of the edge is very different from Chern insulators. The QSHE hinges on strong spin orbit coupling discriminating the motion of different spins, which gaps the bulk and spawns helical channels along the edge: up spins flow one way, down spins flow the other way. In contrast to the QHE, these channels are not spatially separated, which allows them to respect time reversal symmetry. In the bandstructure of the edge, these helical channels show up as two crossing bands. As the states present in these different gap crossing bands represent Kramers' pairs [8], no avoided crossing is generated by interactions or disorder, and the effect is stable. This topological insulator is of a different type than the original quantum Hall Chern insu-

lator, with a topological invariant $n \in \mathbb{Z}_2$ rather than \mathbb{Z} . Indeed there are only two classes of topological insulators with intact time reversal symmetry in 2D: a trivial one, and a topological one.

The concept of a symmetry of the Hamiltonian facilitating nontrivial states by stabilizing gapless edge modes can be generalized to all local symmetries of hermitian matrices: time reversal, particle-hole and chiral. The resulting periodic table of topological insulators, known as the tenfold way, lists the possible (strong) topological indices as a function of respected symmetries and spatial dimension [9–11].

The inclusion of additional (non-local) discrete symmetries further enriches the landscape of topological insulators [12]. A state might be adiabatically connected to a trivial insulator in a global sense, but restricting the allowed deformations by requiring the system to obey specific symmetries might obstruct this connection. In the presence of these symmetries the relevant Hilbertspace is reduced, and within this new configuration the state can acquire a nontrivial topological index, which – within the subspace respecting the symmetry – can not be changed without closing the bulk gap. Depending on how symmetries align with the structure, this methodology also allows the realization of higher order topological insulators (HOTIs) [13–15]. In such d -dimensional systems, it is not the $(d - 1)$ -dimensional boundary that hosts a protected gapless mode, but instead a $(d - 2)$ -dimensional corner. Drawing a bridge to the QSHE, we wouldn't have time reversal symmetry that protects the gapless edge states via Kramers' pair formation, but we would have a symmetry that only protects those states in the corners of the samples while the rest of the boundary is allowed to gap out.

The idea of adding non-local symmetries to achieve topological phases can be formalized using the concept of Wannierizability and leads to a classification known as topological quantum chemistry [16]. Starting from an atomic insulator, that is, a periodic structure with sufficiently large lattice constant for the orbitals to not hybridize, a conventional trivial insulator can be constructed by reducing the lattice constant. In fact, every trivial insulator can be adiabatically constructed in this way. The resulting bandstructure is then spanned by symmetry respecting Wannier functions located at the atomic positions: they are wannierizable. Conversely, non-trivial insulators are not adiabatically connected to atomic insulators, and their bandstructures cannot be decomposed into such orbitals [16].

This observation opens a strategy for the identification of topological insulators from their bandstructure [16]. For any given space group, one can generate all bands that are constructable out of suitably placed symmetric Wannier functions. The resulting bands are called elementary band representations (EBRs). The topological classification of a band under test can then be investigated by comparing it to the available EBRs. If the band can be written as a combination of EBRs, it must be trivial, and if it can't it must be topological. It is noteworthy that this classification is a pure symmetry argument, EBRs are characterized solely by their symmetries at distinguished points in the BZ, and the specifics of a Hamiltonian are irrelevant.

Upon close inspection, the topological band class of the topological quantum chemistry classification splits into two subcategories [17]. Firstly, there are topological bands that represent the combination of fractions of available EBRs. Examples for this type occur in systems where s and p orbitals hybridize and invert in some regions in the Brillouin zone — a phenomenon known as band inversion —, such that the resulting bands pick up parts of different EBRs. This type of band is topological in the conventional sense, and can only be trivialized by combination with other conventionally topological bands. Secondly, there are bands that can be written as a combination of available EBRs with integral weights, but the weights are not all positive. Such bands are lacking an integral number of EBRs to be trivial, so they are topological. However, their trivialization doesn't require a topological band, the addition of a trivial band with the "missing" EBRs is enough. Due to their weaker protection, this class of topological bands is called fragile. Further properties of fragile topological phases, especially pertaining their experimentally accessible signatures, were developed and measured by Peri et al. [18].

Being able to classify bands by their topological properties by itself is not yet useful for the design of new materials with interesting properties. While the above quantum chemistry classification allows the recognition of topological bands once they arise, design problems are inverse: One prescribes a functionality, such as a target topological property to be present in a bandstructure, and then desires to obtain the recipe for a material that hosts such a bandstructure. For 2 dimensional materials, this requires a lookup table listing the possible topological states for each wallpaper group. Given this database, a design algorithm can then start from a spe-

cific symmetry group and optimize the structure, with a guarantee that the target properties can be recovered. Precisely such a design suite was developed by Bird et al. [19] and will be introduced in [chapter 6](#).

1.3 Going classical: mechanical degrees of freedom

While the theory of topological insulators was spawned within the context of condensed matter physics, it is by no means limited to quantum mechanical systems. Already Berry himself realized when he discovered his geometric Berry phase, that this concept generally applies to wave physics [5]. This observation fell into obscurity, and attributing quantumness to topological insulators was a common misconception until a quantum spin hall state was realized in an irrefutably classical system [20].

The ability of classical systems to host topological states is more than a mathematical curiosity. Typical condensed matter systems consist of atoms, with properties and arrangement largely dictated by chemistry. Engineering a material arranged according to a specific space group with desired properties such as strong spin-orbit coupling then becomes a difficult, potentially impossible task. Classical systems are largely free of such restrictions, which makes them good candidates for realizing and investigating topological phases. Periodic classical materials engineered to host specific wave controlling properties are called metamaterials.

A metamaterial is a periodic structure with properties arising from structures much larger than its constituents. While the properties of the base materials do transfer to the metamaterial, its primary properties are given by structures several orders of magnitude larger than the chemical scale. Unshackled from the restrictions of chemistry, the artificial atoms and lattice can be freely engineered using subwavelength features to yield properties unobtainable in naturally occurring materials. Achieved properties include doubly negative materials – materials with negative density and elastic modulus – which allow the construction of superlenses able to resolve features much smaller than the observing wavelength [21], soft clamping, where a 2D material's engineered bandgap suppresses phonon propagation at targeted frequencies to isolate a centrally localised mode, which led to record breaking quality (Q) factors in silicon nitride res-

onators [22], and cloaking, where a metamaterial guides waves in a way that renders an enclosed object invisible at specific wavelengths [23, 24].

In 2013, Kane and Lubensky [25] realized that metamaterial design principles can be used to implement topological insulator states. Applying principles of topological bandstructures to achieve properties in metamaterials yields what is called topological metamaterials. While investigations have also been carried out in optical [26, 27], acoustic [28] and electronic systems [29–31], we will focus on the developments in the mechanical and elastic domains.

The first realization of a topological state in a mechanical system was performed by Süsstrunk and Huber in 2015 [20]. They transferred a doubled Hofstadter model [32], which exhibits helical QSH modes, to a 2D lattice of coupled macroscopic pendula. Each artificial atom consists of two pendula, and it is this subwavelength feature that spawns two distinct onsite excitations and facilitates their topologically protected helical transport along the edges of the system, in (mechanical) analogy to the spin polarized edge modes in QSH states. This experiment represents a landmark in the history of topological states, because it irrefutably demonstrated the existence of topologically nontrivial states in classical waves. Foregoing experiments had either used systems governed by quantum mechanics or quantum optics, where the quantumness of the latter was often not clear. It is this proof of principle using classical pendula that gave rise to explorations in other classical systems.

1.4 Towards scalability

While the importance of the topological pendula demonstration is obvious, the platform is not without limitations. These mainly concern manufacturability, scalability and flexibility. Large systems of macroscopic pendula are not widespread in technology, such that their manufacturing is cumbersome and expensive, requiring specialized machinery and personnel. Even for a moderate number of 270 pendula [20], the system is on the scale of meters, and scaling up the number of sites will only lead to bulkier, more inconvenient setups without any hope of integration into other systems. Finally, while there is some tunability in coupling strengths, more radical changes to the model would require a complete reboot of the system design.

To address the above limitations, one can leverage decades of progress in manufacturing technology by transitioning to microelectromechanical systems (MEMS). MEMS appeared first half a century ago, they are a well developed standard platform facilitating diverse applications including sensors, actuators and oscillators [33]. They leverage cleanroom manufacturing techniques based on lithography, a widely available, highly optimized pipeline driven by the microchip industry [34]. Given these properties, the natural path of evolution for pendula is towards miniaturization, towards MEMS.

With all the stated issues the pendulum system has, there is one big advantage over general MEMS structures: It separates well into a finite number of discrete degrees of freedom. Every pendulum has one eigenmode, and each coupling is represented by one spring. This similarity to lattice models makes transferring parameters from these models to the pendula all but trivial. In MEMS this is typically not the case. These systems live in a continuous design space and have infinitely many eigenmodes, such that specialized methods are necessary to transfer properties of lattice models into this realm.

There are different approaches that have been used to translate lattice models exhibiting topological states into continuous MEMS designs. One robust, widely used approach is based on an observation by Wu and Hu in 2015 [35]. As mentioned previously, the gapless surface states in the QSHE are stabilized by time reversal symmetry as Kramers' degeneracy protects the band crossing [7]. However, time reversal symmetry only effects Kramers' pairs if it squares to -1 , which is true for fermions, but not for bosons. There is no reason to expect a classical system to exhibit this behaviour, and the QSHE does not transfer in this form. The idea due to Wu and Hu is to generate a synthetic symmetry with Kramers-esque properties, by concatenating time reversal symmetry with a spatial symmetry that has two invariant sectors mapping into each other under time reversal, such that the combined symmetry squares to -1 . One then obtains a fermionic pseudospin living in these invariant sectors, and the synthetic symmetry effects Kramers-like pairs and can protect gapless edge states. As long as a MEMS design respects the synthetic symmetry, a QSHE can in principle be implemented. Typically this technique is applied to hexagonal lattices by choosing a 6-atom unit cell. This choice folds the Dirac cones at K, K' to a degenerate pair of cones at Γ , giving rise to a C_6 symmetry with

two irreducible representations related to inversion, thereby enabling the above mechanism. An example experimental realization in a continuous MEMS was provided by Cha et al. in 2018 [36]. Starting from a layer of silicon nitride on silicon dioxide, they partially liberated the nitride layer in a hexagonal pattern deformed to open a bulk gap while preserving the C_6 symmetry. As expected, they observed helical edge states, with the helicity given by the pseudospin state.

While the Wu and Hu approach always works to obtain helical edge states, its severe limitations are lacking generality and flexibility. It can facilitate QSH states, but only when a suitable synthetic symmetry can be produced, and any design not respecting that symmetry is impossible. For any other type of topological state or samples with restrictions prohibiting the required symmetry, the mechanism gives no design principle.

A different, more general approach for transferring properties from discrete models to continuous systems is provided by a technique known as perturbative metamaterial design [37]. It hinges on identifying lumped elements in a metamaterial design that correspond to specific functionalities of a corresponding lattice model, along with methods facilitating the efficient calculation of how small changes to the design affect their model counterparts. The crucial assumption of linearity then allows the combination of different elements to achieve a metamaterial with the target functionality. An example application demonstrating a novel topological state was performed by Serra-Garcia et al in 2018 [14]. Starting from a lattice model hosting a HOTI state, they applied the framework of perturbative metamaterials on silicon structures to obtain a corresponding metamaterial. The result is a planar silicon design that reproduces the properties predicted in the initial lattice model. In contrast to the Wu and Hu approach, perturbative metamaterial design transcends topological states and can be used to target more general properties. In 2021, Dubcek et al [38] designed a non-periodic perturbative metamaterial able to classify spoken digits by their numerical value. This application highlights the flexibility of the method, intertwining nontrivial topology with further target functionality is no longer out of reach.

1.5 The future: tunable metamaterials

Design strategies transferring models from discrete to continuous platforms represent an important step in developing the ability of realizing topological states in metamaterials. However, both of the above approaches are not satisfactory results for the efficient exploration of novel material properties. The Wu and Hu approach allows the certain transfer of QSHE states to the mechanical domain, but nothing more. There is no obvious path that allows the incorporation of further functionality into the structure. Furthermore, the resulting structure is inflexible, any tuning of the design – such as exploring deformations of the lattice – requires the design and manufacturing of a new sample. Similar issues plague the perturbative metamaterial approach. While generality is greatly expanded and general discrete models can be incorporated into metamaterials, changes to the design still necessitate a complete sample overhaul. It is this slow feedback loop that limits development throughput.

Representing the next step in development, a material development platform that overcomes the above issues should optimally have the following three properties: i) it is manufactured using standard processes, ii) it is able to incorporate lattice models with a reasonably large number of sites, iii) it allows dynamic control of all system parameters, such as neighbour couplings, preferably in real-time.

Having formulated the above requirements, we now want to translate them into a more concrete design. Leveraging the highly developed MEMS ecosystem, individual lattice sites are represented by micromachined resonators. Facilitating dynamic couplings between neighbouring resonators is a more involved undertaking, especially when a large lattice is to be generated. The two main interaction pathways are mechanical and electrostatic. Direct mechanical interaction, as used in the aforementioned applications [14, 36, 38], is robust, but severely limited with respect to dynamic tunability. As a middle ground between electrical and mechanical, piezoelectric transducers allow the electric transmission of mechanical deformations, but require electronic components in each coupling's path to facilitate control. Finally, direct electrostatic interaction via electrodes allows full scale tuning of couplings by changing a single DC voltage per coupling. Providing full control with minimal impact on scalability and complexity, this is the favoured implementation of coupling between sites.

Also implementing resonator actuation electrostatically, then allows for uniform control of the available degrees of freedom.

Above, the ideas behind a material simulation platform based on a lattice of MEMS resonators with electrostatic coupling and actuation have been outlined. The detailed implementation of this system, including the control infrastructure, is more involved and will be presented in [chapters 2](#) thru [4](#). To conclude this motivational chapter, next will be a brief sketch of four different applications facilitated by this platform. In contrast to previous metamaterial designs, the system described above provides a significant increase in flexibility and can be programmed to incorporate a variety of experiments.

1.6 Target applications

Nonabelian braiding: The first application pertains nonabelian braiding, a topological effect predicted in the context of condensed matter physics, but later transferred to the classical domain [\[39\]](#). A lattice with couplings patterned according to a Kekulé distortion with superimposed vortex will localize a zero mode pinned at the vortex core. Analogous to Majorana bound states in chiral superconductors, such modes exhibit nonabelian exchange statistics, such that the phase accumulated upon adiabatic interchange of several such modes depends on the swapping order. While the initial proposal for classical implementation of nonabelian braiding [\[39\]](#) was constructed for optical systems, this implementation is readily translated to the new MEMS platform. While a measurement of effective nonabelian braiding in optics [\[40\]](#) has already been achieved, and similar experiments have been carried out in analogous acoustic setups [\[41, 42\]](#), a mechanical demonstration is still outstanding.

Floquet and parametric phenomena: Full control over all couplings and onsite excitations allows the implementation of a large class of static lattice models. However, as these parameters can be modulated in real-time, also time dependent systems are within reach. The second application thus broadly consists of arbitrary parametric phenomena. One example from condensed matter physics are Floquet topological insulators [\[43–46\]](#), where statically trivial materials can suddenly spawn chiral edge modes when their couplings are subjected to a suitable periodic modulation, even though all Chern numbers remain trivial. Other examples are

found through parametric drive of the coupling between two resonator modes, which can yield an oscillator. Coupling such parametric oscillators then allows the study of their dynamic effects, like synchronization [47] and self-organization [48], with relevant implications for the study of brain activity during epileptic seizures [49].

Signal processing and machine learning: The third application lies in the field of machine learning. With neural networks becoming more prevalent, efficiency considerations are an important aspect of their scaling usage. As conventional electronic application specific integrated circuits (ASICs) consume many orders of magnitude more power than electromechanical counterparts [50–52], transferring machine learning applications to MEMS platforms is a promising step forwards. Leveraging the nonlinear nature of the electrostatic interaction, the system allows the implementation of arbitrary functions via neural network topologies [53, 54]. Several different machine learning models can be implemented in the system, such as model-free reinforcement learning [55], reservoir computing [56–58] and physical neural networks [59]. The ability to approximate arbitrary functions enables further signal processing applications, in similar spirit to the classifier by Dubcek et al [38].

Nonlinear phenomena: The fourth application concerns the coupling of nonlinear dynamics and condensed matter lattice models. When the resonators of the system are operated at anharmonic amplitudes, further investigations from condensed matter physics become accessible. One example is the influence of nonlinearity in the presence of disorder or flat bands. While strong disorder is well known to lead to Anderson localisation [60] and isolated flat bands typically represent atomic insulators with localized atomic orbitals, coupling shades of these two properties can effect the emergence of interesting phenomena. The additional influence of nonlinearity on this interplay is currently largely unknown.

The rest of this thesis is organized in the following way: [Chapters 2](#) thru [5](#) concern the newly developed experimental system. [Chapter 2](#) discusses the sample forming the core of the new system, describing in depth why specific design decisions were taken. In [chapter 3](#) the fabrication process of this sample, onto which most of the development effort was spent, is described. The hardware infrastructure surrounding the sample, used to drive the system, is explained in [chapter 4](#), and [chapter 5](#) highlights measured properties of the system. [Chapter 6](#) introduces the software suite for

the design of two dimensional topological materials. Finally, conclusion and outlook are provided in [chapter 7](#).

Before delving into sample design, the contributions of the author and other involved people will be declared in the next section.

1.7 Author's contributions

The goal of this section is declaring the author's and other contributor's contributions to the different chapters part of this thesis. The main contributors to the synthetic metamaterial are the author (PE), Marc Serra-Garcia (MSG), Sebastian Huber (SDH) and Ute Drechsler (UD). In addition to what's specified below, throughout the process, SDH and MSG were involved in supervisory capacity.

[Chapter 2](#) is a description of the process used to design a sample with the desired properties. PE was involved in all stages of this process, from initial platform selection to final structural optimizations. Initially, the development process was led by MSG, with PE in a more subordinate role. Roughly from the point where fabrication first started, PE took over as the hands-on lead, in close collaboration with UD. PE programmed and ran simulations, fabricated and characterized prototypes, optimized the structural design using custom software (written by PE) to address observed issues, and identified new pathways to enhance design verification.

[Chapter 3](#) describes the fabrication process and explains the reasoning behind different steps and decisions. The initial fabrication protocol was written down by MSG in collaboration with UD, while PE was involved in a more tangential role. All further development was carried out and led by PE. The initial process was then greatly changed in collaboration with UD to address observed issues, mainly pertaining manufacturability and yield, and implement changes on the sample design. The whole optimization of the fabrication process was led and done by PE, as were all fabrication runs.

[Chapter 4](#) describes the experimental system used to drive and evaluate the sample. PE was involved in the all aspects of the system design, implementation and verification. Design of the initial versions of the Minicircuits and the Bridge were led by MSG, with PE involved in an advisory, knowledge discovery and prototyping role. PE later developed extensions and updates to these two components. PE also developed the

vacuum chamber, with advisory input from Walter Bachmann, the Break-outboards, procured the vacuum pump and all off the shelf parts of the environmental control system. PE developed and implemented all aspects of the temperature stabilization unit. Julia Rupprecht integrated the camera module into the setup under supervision of PE.

[Chapter 5](#) describes measurements carried out on sample F1F4. From blank wafer to functional installed sample, F1F4 was fabricated by PE, with UD carrying out the BHF fabrication step. All measurements on F1F4, including their analysis, were carried out by PE, with SDH in a supervisory role. MSG provided input on coupling enhancement via difference modulation, which was then formalized by SDH and PE and implemented by PE.

[Chapter 6](#) describes the structure search design suite, a software stack for the discovery of topological designer materials. PE was involved in this project from the start, and developed the symmetry preserving mesher. This component, given the generators of a wallpaper group, automatically produces a mesh respecting the specified symmetries, which is then used for the finite elements computations. In the initial stages of the project PE kick started the development by implementing the first functional prototype that solved the Poisson equation using a finite elements solver on a symmetry preserving mesh, calculated band structures along specific high symmetry lines, and provided eigenstates in an interpretable basis. After this proof of concept, in the later stages of the project, PE was mainly involved in a consulting role pertaining system design, software design and performance optimizations. The text and figures that are presented in [chapter 6](#) were largely contributed by SDH for the purpose of publication, and have been adapted by PE for presentation in this thesis.

2 | Sample

Here we explain the sample design laid out in the last section, highlighting the design decisions that were taken, discussing available alternatives and pointing out problems that were encountered. The chapter is organized in the following way: First the large scale sample geometry is introduced, and it is explained in detail how the design implements the wishlist laid out in [section 1.5](#). Next, focus is placed on the design of the individual resonators. Finally, the electric design of the sample is discussed. In summary, the sample consists of 2015 silicon nitride resonators arranged in a hexagonal lattice by stacking two triangular lattices, with nearest neighbour coupling facilitated by capacitive interactions between lattice neighbours.

2.1 Large scale: geometry

As outlined in [section 1.5](#), the goal is the implementation of a nearest neighbour coupled lattice model with many sites in a MEMS, using electrostatic interactions to facilitate coupling and excitation. Each lattice site is implemented as a resonator, with site occupation corresponding to excitations placed in a specific eigenmode of that resonator. A rendering of the final sample design, made up of two 4 inch wafers, is shown in [figure 2.1](#). In this chapter, the different parts of this design will be discussed

in detail. This section focuses on large scale components, with sections 2.2 and 2.3 discussing the design of individual resonators and the electric routing, respectively.

The two wafer nature of the design is a consequence of the implementation chosen for the electrostatic coupling. The choice falls on a capacitive design, which mimics a parallel plate capacitor [61]. Capacitive coupling is commonly used in applications at scale, such as for capacitive micromachined ultrasonic transducers (CMUTs) [62, 63], meaning one can profit off existing manufacturing processes. This choice directly implies that neighbouring resonators must be placed in an overlapping stacked geometry, such that capacitors with appreciable surface area between nearest neighbours are formed.

An alternative electrostatic coupling mechanism allowing a planar implementation could be a comb-drive finger design [61, 64]. However, comb-drive interactions typically couple displacement components aligned along a specific direction in the sample plane. As the object at hand is a 2D lattice, it is then not obvious what a suitable resonator mode, coupling isotropically in all bond directions, could look like. One could switch to a stacked implementation much like the capacitive design, by using combs perpendicular to the resonator surface, but this would come at the cost of exploding manufacturing complexity and degraded stability for questionable gains. Overall, the capacitive design has been found to be preferable for the system at hand.

The sample splitting necessitated by the choice of capacitive coupling restricts the implementable lattice geometries. In order to allow a separation into two parts with nearest neighbour interactions only taking place between sites in different sublattices, the lattice must be bipartite. While there is a number of bipartite lattices in 2D, the nonabelian braiding application outlined in section 1.6 was formulated for a hexagonal lattice [39]. For this reason, the lattice is chosen to be hexagonal, built up from two triangular sublattices.

The two sublattices of resonators are manufactured on two 4 inch wafers. There are differences in the design of these two wafers. While the top wafer only contains a central array of resonators which is diced out, the bottom wafer stays whole and contains additional electrical routing, which will be discussed in depth in section 2.3. These two wafers are then mounted on top of each other, such that the triangular sublattices have a




Figure 2.1: Render of the final sample design. The sample is split into two pieces, the bottom piece consisting of a whole wafer, and a smaller diced top piece. The top part only contains resonators, and their electrodes are all connected to the same continuous region of gold. On the bottom part, there is a central region of resonators interspersed and surrounded with adhesive pillars. Each electrode on the bottom layer is routed out of the resonator region to individual bond pads placed next to fishing holes. The two layers are assembled by adhesive wafer bonding, as indicated by the two black arrows. Further information about the design is provided in the text.

suitable planar offset and overlap as required to form a hexagonal lattice. The assembly of the two layers is done via adhesive waferbonding, which uses an intermediary layer of structured adhesive material between the wafers to achieve fusion. While the precise discussion of this technique, including the selection of suitable adhesive materials, is deferred to [chapter 3](#), there are two important points concerning the design side of this joining process that should be addressed here.

Firstly, for reasons that will become clear in [section 2.3](#), the gap between the two wafers must be kept as constant as possible across the whole lattice. This requirement is not automatically fulfilled, because wafers with nontrivial stackup are typically subject to warping effected by irregular stress distributions due coating asymmetry. To counteract this bow, adhesive must be placed at regular intervals throughout the lattice, as shown in [figure 2.1](#). Surrounding the resonator region as a whole with adhesive material might seem like the cleaner solution, but has no ability to counter bow in any capacity, which yields not only unworkable gap fluctuations, but might even completely inhibit a bond.




Figure 2.2: Initial design of the adhesive structuring mask. The central resonator region is surrounded by a continuous ring of adhesive.

Secondly, the adhesive's structure should be designed with local failure tolerance as a goal. As will become clear in the process discussion of [chapter 3](#), the surface onto which the adhesive material is deposited is strongly processed and structured. Surface adhesion is then not automatic, and varies throughout the sample. In effect, the adhesive might delaminate in some regions, while it is more stable in others. Design steps can be taken to allow some robustness against these issues. In early itera-

tions, pillars of adhesive were scattered throughout the lattice, as required to counter wafer bow, with the whole resonator region surrounded by a continuous ring of adhesive. This design is shown in [figure 2.2](#). Regional delamination can then lift the ring locally, upon which it will fold onto itself and re-adhere to form a doubled or even tripled layer of adhesive potentially covering large numbers of resonators. This type of failure was observed, and immediately renders the device unbondable with no realistic chance of recovery. To make matters worse, adhesive delamination typically takes place in the final steps of the manufacturing process, yielding maximum cost on failure. The mitigation of the issue is twofold. On the fabrication side, there are process optimizations to improve adhesion, which will be explained in [chapter 3](#). On the design side, the solution, as shown in [figure 2.1](#), is replacing the surrounding continuous ring with a region of densely placed discrete pillars of adhesive. While individual pillars may then still delaminate, they are now likely to be fully flushed off the wafer rather than relaminating and destroying the sample. Even if readhesion were to occur, which has so far not been observed, there is a chance for successful recovery by manual mechanical removal of a few affected pillars.

In this section, it has been argued that the sample shall be split into two triangular parts that are stacked via gapped waferbonding to yield a hexagonal lattice of resonators. With the large scale buildup of the sample discussed, the next section will deal with the design of individual resonators.

2.2 Small scale: resonators

In this section the design of individual resonators is discussed in two parts. In the first part, a wishlist of desired properties is formulated, and in the second part the development of a design that implements these points is described.

While making drum top resonators in the MEMS ecosystem is not a specifically challenging task per se, constraining the properties of the resonators has favourable effects on the system at hand. However, additional restrictions make design and fabrication more involved, and the goal is to only apply constraints where this complexity tradeoff represents an overall simplification. In this spirit, we formulate four local requirements to

be met by the resonator design. Some of the required properties are mission critical, while others significantly reduce the complexity and cost of surrounding infrastructure. They are:

- i. Resonators have one central electrode for driving, and three electrodes close to the boundary that overlap with corresponding electrodes of neighbouring drums on the other layer for coupling.
- ii. The resonators' surface conforms with gap uniformity requirements found from electrostatic coupling considerations.
- iii. Resonators have a suitable quality (Q) factor.
- iv. Resonators have an SO(2) symmetric mode with eigenfrequency close to the audible range.

The first requirement is an obvious direct consequence of the initial goals to be reached with this system, but the other points need a bit of discussion. Requirement (ii) is effected by the nonlinearity of capacitive electrostatic interaction. The bottom line is that in the given configuration, appreciable coupling strengths are only achieved when the gap between electrode pairs is around $1\ \mu\text{m}$, such that the resonators need to be flatter than this. The derivation of this restriction is deferred to the next [section 2.3](#).

The third requirement corresponds to a lower and upper bound on the quality factor of the resonators. The quality factor Q describes the number of periods a resonator will oscillate until the oscillation energy has decayed by a factor of $1/e$. A lower bound on Q is necessary, because a number of potential experiments follow protocols that do not drive resonators throughout. Excitations are injected into the system at the start of the experiment and subsequently the evolution is observed under coupling modulations, with modulation speed restricted by adiabaticity. For this to work, the initially deposited energy must remain in the system for longer than the protocol duration. An upper bound on Q is necessary, because experiments typically need the resonators to initially be at rest. Unreasonably large quality factors then effect long ringdown wait times in sequenced experiments, which inhibits progress. For example, quality factors exceeding 10^8 at 1 MHz have been reported in MEMS resonators [22], which corresponds to a half-life time on the scale of minutes. However,

should large quality factors be found to be an issue, there are three possible mitigating steps. In order of increasing invasiveness, these are: active braking using the central electrode, increasing external friction by raising atmospheric pressure, and increasing internal friction in a specific eigenmode by depositing metal on the resonators at points where the mode shape has strong bending ¹.

Finally, the fourth requirement has two components: Firstly, the mode needs to be (approximately) SO(2) symmetric, because driving the resonator happens at its center and energy can not be efficiently transferred to inversion antisymmetric modes. Secondly, the resonance frequency directly determines speed requirements on the control system driving this sample. A lower resonance frequency implies a longer time per sample applied to the resonator electrodes, which removes strain from the processing chain generating the signals. Staying close to the audible range additionally allows leveraging the well developed and widely available audio processing ecosystem when designing the surrounding infrastructure.

There is one more restriction, effected by global sample geometry, that has to be kept in mind when designing the resonators. As shown in [figure 2.1](#), each resonator has four electrodes that are routed out of the lattice region on the bottom wafer through the gaps between resonators. The spacing between resonators on the bottom wafer, determined by the resonator diameter via the requirement that neighbouring resonators overlap, bounds the number of electrodes that can be routed, thus restricting the number of resonators. For large diameters, the number of resonators is restricted by space availability, and for small diameters by routability. The spatial extent of drums should thus be chosen consciously.

With all restrictions in place, the next part of this section lays out the design process to achieve resonators with the target properties. First, the resonator material needs to be selected. There are several popular choices, including silicon nitride (Si₃N₄) and silicon (Si) [69–71]. While both of these materials have been used in many applications to implement reasonably good resonators, there are important differences. Silicon nitride is a great insulator, with its resistivity in excess of 10¹⁵ Ω cm being over ten orders of magnitude greater than what is typically achievable in silicon, which is favourable in a design where thousands of metal traces are to be routed as densely as possible. To facilitate the typical release via

¹A more in-depth discussion of dissipation in MEMS resonators can be found in [65–68]

buffered hydrofluoric acid (BHF), silicon nitride is often stacked on silicon dioxide (SiO_2), with its resonance frequency then typically dominated by tensile prestress [34]. On the other hand, silicon for MEMS resonators is available on silicon on insulator (SOI) stackups with significantly lower prestress [72], such that its frequency is stiffness dominated at smaller thicknesses than silicon nitride. All things considered, silicon nitride is chosen for the implementation of the resonators, primarily to leverage pre-existing experience of local colleagues and an interest in developing know-how handling this stack. This doesn't mean silicon couldn't be used to implement the system, the differing properties might even improve certain issues inherent to silicon nitride.

A final comment on the consequences of the outlined material selection is in order. Choosing to work in the stress dominated regime implies strong sensitivity to temperature fluctuations. A change in temperature expands or shrinks the size of the resonator, which in turn has an inverse effect on its internal stress, which then detunes the resonance frequency. As will be reported at a later stage, the magnitude of this effect is such that active temperature stabilization needs to be employed for the resonance frequencies to remain constant over experimental timescales.

With the material selected, the electrical interaction shall be designed in next. This task is rather straightforward, there is one electrode per nearest neighbour, and a fourth electrode placed centrally for onsite excitation. Whichever sublattice a resonator belongs to, the electrode design is identical for all of them. Upon bonding, the top lattice is flipped as indicated in [figure 2.1](#), such that the electrodes align correctly. Note that this property is not automatic, but enabled by careful design. Each of the four electrodes is routed out of the drum, and to keep the drum as rotationally symmetric as possible, the central electrode is routed out in a C3 symmetric pattern to retain overall C3 symmetry. This ensures the lowest lying eigenmode will follow the same symmetry and exhibit finite deformation in the drum center, such that excitation becomes possible.

Next, the frequency design is discussed. As it directly impacts resonance frequencies, it shall be noted that the diameter of the resonators is chosen to be roughly 1 mm. The reasoning for this is a combination of measurability, which will be further discussed in [chapter 4](#), and the projected minimum number of resonators necessary to implement the most challenging application, nonabelian braiding. This choice yields a lattice

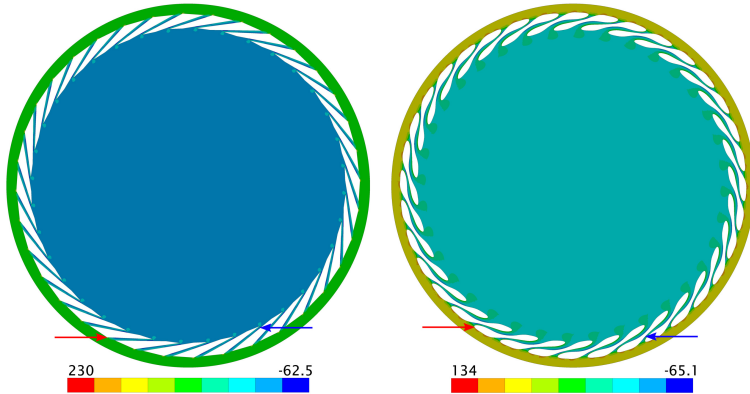


Figure 2.3: Stress release simulation results for two different resonator designs, starting out at 100 MPa prestress. Simulations were carried out in Ansys, signed von Mises stress is shown in units of MPa. Coloured arrows indicate locations with maximal and minimal stress. Left: Canted straight arms with sharp corners. Stress accumulates in hinge points, which becomes a leading factor in yield issues. Right: Smoothed out version of the design on the left. Stress still accumulates around the hinges, but to a significantly lower value. This new design improves yield by several hundred percent over the left version.

of 2015 resonators.

As mentioned above, silicon nitride membranes on silicon dioxide typically have high prestress, which leads to resonance frequencies far above the audio range. As a starting point, we begin with a material variety known as super low stress LPCVD silicon nitride ($100 \text{ MPa} \pm 50 \text{ MPa}$), which puts the lowest lying mode at roughly 100 kHz. The pathway to further frequency reduction in the tension dominated regime is allowing the resonator to relax a portion of its prestress by cutting carefully designed holes into the material. This structural optimization was first carried out in COMSOL Multiphysics and later in Ansys Mechanical. Initial designs followed the philosophy of attaching the drums to the bulk material with a few spring like suspension arms, which would elongate and allow the membrane to contract upon release. This principle was discarded as the resulting designs proved either ineffective or hard to fabricate and unsta-

ble. The brilliant idea leading to a working guiding principle is torsional stress release. Instead of targeting linear relaxation where suspending elements change their length, the membrane is suspended on canted arms. A rotational motion of the drum then allows it to contract, without requiring an appreciable change in the arm length. The task does not finish here. When cutting holes into materials under large prestress, caution is advised. With new motional freedom, the structure will move and the prestress will redistribute, and can accumulate in neuralgic points to break the resonator. This is addressed via the evolution from the initial canted design to the smoothed out design shown in [figure 2.3](#). The latter reduces the resonance frequency to roughly 15 kHz, nicely into the audio range.

Unfortunately, the fulfillment of the frequency requirement leads to a violation of the gap uniformity requirement. The discovery of this issue was led by the unscheduled observation of colourful patterns, known as Newton rings [73, 74], on all resonators in a fully bonded sample, as shown in [figure 2.4](#). Theoretical investigations involving a specially developed high performance wavelength resolved interference ray tracer confirmed the suspicion that the membrane in [figure 2.4](#) is most likely in




Figure 2.4: Wafer bonded sample, view onto the top resonator layer. The overlap with the bottom drums is clearly visible. The coloured Newton rings indicate contact.

contact with the opposing layer. The reasons effecting this touching were finally tracked down using white light interferometry to map out the topography of individual resonators. Indeed, the cutting of holes leads to the membrane deforming out of plane by roughly $3 \mu\text{m}$, in excess of the space provided by the $1 \mu\text{m}$ gap and at odds with resonator requirement (ii). Further simulations using Ansys Mechanical, yielding results consistent with measurements, suggest that there is a stress gradient of roughly 100 MPa through the thickness of the silicon nitride layer. Once holes are cut into the membrane boundary, the resonator doesn't only relax the in

plane tension via rotation, but also balances the perpendicular stress differential by deforming correspondingly. Other suspects, like a potential layer of silicon dioxide on the resonator surface, effected by exposure to oxygen plasma, couldn't be confirmed experimentally nor in simulations.

Mitigation of this effect is a nontrivial undertaking. It is not obvious how motion perpendicular to the surface can be inhibited when no structuring along this direction is possible. Guided by simulations and intuition, an array of different resonator designs was generated using purposely designed software, transferred to a mask design and fabricated, and the best performing structure selected. The leading principle for the designs was a tradeoff between less in plane relaxation and more out of plane stability, implemented by reinforcing the arms in the existing hole design in various ways. [Figure 2.5](#) shows the predicted and measured out of plane deformations of membranes following the initial optimized design, and the selected best mitigating design. A rendering of the final resonator design obtained in this process is shown in [figure 2.7](#). Calculated mode profiles and frequencies of the lowest few resonances of this design are shown in [figure 2.6](#). The lowest eigenmode is at roughly 25 kHz, still close to the audio range but clearly highlighting the frequency-flatness tradeoff at play.

Before concluding the discussion of resonator design, some comments on the quality factor are in order. Choosing silicon nitride as resonator material and roughly following known designs, reasonable Q factors are almost automatic. For this reason, the thickness of the resonators is chosen to be 500 nm. Thinner resonators have been made, but a thicker resonator is a heavier resonator, and in the tension dominated regime this implies lower frequency, a desired property.

As a brief aside it is emphasized that thinner designs should be explored in the future, especially pertaining the mitigation of lateral stress gradients effecting out of plane deformations. These gradients are hypothesized to signify structural relaxation from the silicon dioxide lattice towards stress free bulk silicon nitride, as more layers of Si_3N_4 are stacked on top. If this is the case, thinner layers of silicon nitride could exhibit far lower stress differentials, leading to less deformation. Factors reducing the Q factor from its ideal value are internal friction due to electrodes, leakage of phonons from the resonator into the wafer bulk, and friction with the surrounding medium. As the resonators at 1 mm diameter and

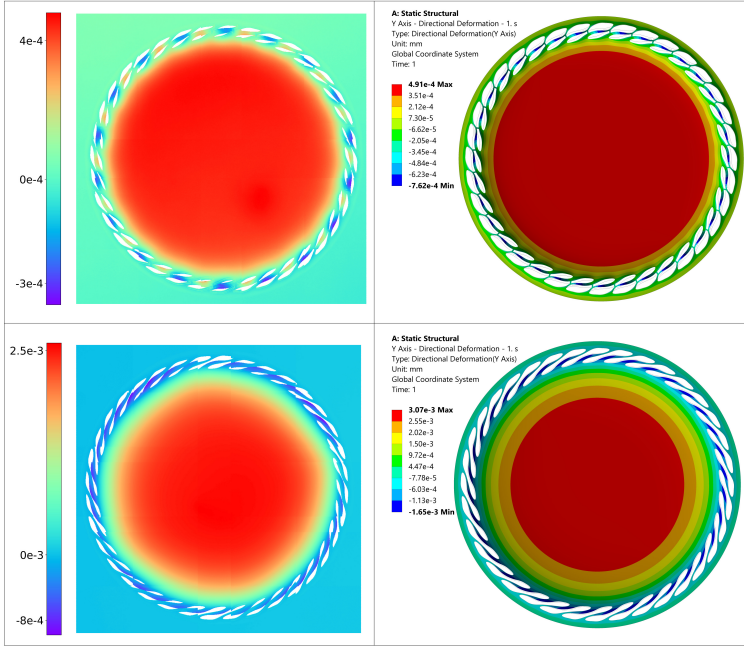


Figure 2.5: Transverse deformation of different resonator designs, comparison of measurement (left column) and simulation (right column). The results are shown in millimeters. Predicted deformations are consistently larger than observed, this could be due to a reflective gold coating applied to the drums to increase measurement contrast. Measurements were carried out in a white light interferometer (Wycko NT1100). Simulations were carried out using Ansys Mechanical, assuming a stress gradient in transverse direction through the nitride layer. Top row: Most recent design with stabilizing bridges over holes. In reality, a period doubling on the arm deformation is observed which is uncaptured by simulation. Bottom row: Early-stage design with a deformation roughly five times larger than the latest design.

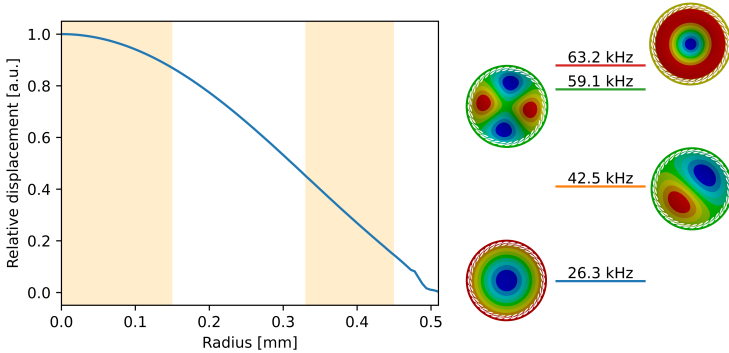


Figure 2.6: Properties of the lowest lying modes of the final resonator design. Left: Relative displacement of the resonator surface as a function of distance from the center in the lowest lying mode at 26.3 kHz. The shaded regions indicate radii where electrodes are present. Right: Frequencies of the four lowest modes with corresponding mode profiles. All shown results were calculated using Ansys Mechanical.

500 nm thickness have a huge surface to mass ratio, friction with air at atmospheric pressure will immediately force the quality factor to zero. To suppress this type of dissipation, the resonators are operated in vacuum where reasonable Q factors exceeding 10^4 are achieved with the laid out design.

In summary, the design of resonators optimized for various properties is a nontrivial undertaking, but has been achieved by a combination of various targeted optimizations and tradeoffs. In the next section, the electrical side of the sample design will be discussed.

2.3 Electrostatic considerations

The electrical design of the sample can be split into three subsections: connection to the outside world, signal routing, and interaction strength. In this section these three categories are discussed, in this order.

As has been outlined in the previous sections, electrostatic interaction is employed to facilitate resonator coupling and actuation. This involves electrode pairs forming capacitors, with each pair having one electrode on




Figure 2.7: Render of the final sample design, closeup of a single drumtop resonator with optimized stress release cutouts, electrode routing and adjacent adhesive pillars. The membrane is free standing with a hole cut through the wafer from the backside. For fabrication reasons explained in [chapter 3](#), the metal routing passing over the resonator arms is designed to be wider than the silicon nitride suspension. The membrane is 1 mm across, metal traces are 4 μm wide and have equal minimum clearance.

each sublattice, which form capacitors. Coupling capacitors are formed automatically between electrodes on neighbouring drums, the central actuation electrodes are paired with an additional electrode placed in the interstitial regions of the opposing wafer. This setup is shown in [figure 2.1](#).

Pertaining control on the couplings, the current design allows the voltage on one electrode per pair to be set. While this implementation only allows for attractive interactions, this choice strongly reduces fabrication complexity. Future developments may tackle repulsive couplings, should the need arise. For the current design this means all electrodes on the top wafer are connected to the same net, which is exposed to the outside world via a single pad. Typically this net would be tied to ground. On the bottom wafer, every electrode is exposed to the outside world individually.

Connecting 5000 electrical nets from a wafer scale sample to the outside is a nontrivial undertaking, and options are available. The main approaches are wirebonding and flip chip bonding. In the given sample, flip chip bonding requires all electrical nets to be terminated in solderable pads on the surface of the bottom wafer. Installation on a printed circuit board (PCB) then means placing the sample upside down, such that the pads are in contact with corresponding solder paste covered pads on the PCB, and reflowing the stack. This method has some concerning properties. Firstly, bond inspection requires specialized equipment, namely an X-ray machine, which is not readily available. Secondly, if the process fails, the device is likely left in a state beyond repair. Thirdly, reflow processes typically require peak temperatures of around 250° C and take place in a special atmosphere, both of which might damage the sample. And finally, solder paste typically excretes flux, which will terminally affect any resonator it comes in contact with, and can provide a conductive layer on top of the routed traces if it isn't cleaned off — which it can't be. While flip chip bonding is a very efficient process by parallelisation, the above properties disqualify it for the application at hand, at least without game changing modifications.

Wire bonding is a less efficient, but safer option. It requires each electrical net to be routed to a bond pad. These pads are then connected to corresponding bond pads on a PCB by welding a metallic wire of microscopic width. As this process is sequential, its time requirement scales linearly with the number of connections. The most prominent issue with wire bonding is low bonding reliability due to poor process optimization




Figure 2.8: Render of the final sample design, closeup showing fishing holes with adjacent bonding pads to facilitate wire bonding to a subjacent PCB, and transiting electrode routing. Each electric net contains two bonding pads, $120\ \mu\text{m}$ in diameter, for redundancy in case of bond failure.

or bond pad design. These types of problems are addressed in [chapter 3](#).

There is one further point to be made pertaining connection to the outside world. A naive approach to bond pad placement would set all of them along the boundary of the bottom wafer. However, at a diameter of 10 cm, not even half the required bond pads fit on the circumference if they are to have reasonable size and spacing. Space needs to be made by enlarging the surface length of the wafer. One idea could be to extend the existing surface by cutting trenches into it, but this complicates handling and might lead to disturbing resonance modes in the wafer. Furthermore, micromachining processes are typically unreliable close to the sample edge, where finite size effects affect resist spinning and plasma uniformity. The implementation applied here and shown in [figure 2.8](#) works similarly to ice fishing, and generates more space for bond pads by producing new wafer boundaries in the interior. Rectangular holes are cut into the wafer surface, and through each of these holes up to 10 wirebonds are threaded to the PCB below. Relevant design tradeoffs here are excessive embrittlement of the sample versus more space for wirebonds. Uniform distribution

of bonding holes ensures maximum stability is retained.

With connection to the outside world sorted out, next each electrode needs to be routed to a corresponding bond pad. The provident design discussed in [section 2.2](#) guarantees routability of all 5000 electrode nets, because space availability is guaranteed and each electrode can connect to an arbitrary bond pad. Detangling this assignment is left to external infrastructure. Performing the routing still requires the identification of a pad-electrode pairing which results in a routable sample. This is done with a custom developed trace router, employing topological sorting of the electrodes and bond pads. The result is a full routing, with traces $4\ \mu\text{m}$ in width and a minimum $4\ \mu\text{m}$ clearance in between different nets, as shown in [figures 2.1](#) and [2.8](#).

Observing the long parallel traces with tiny clearance shown in [figure 2.8](#), it is obvious that crosstalk must be considered. Relevant couplings could be of capacitive and inductive nature. The magnitude of capacitive crosstalk is driven by stray capacitance between two interacting traces, which is maximized for long neighbouring traces. At a trace thickness of $50\ \text{nm}$, this capacitance works out to roughly a few femtofarads. Using typical values for the design at hand, the maximum noise voltage coupled capacitively into a victim trace is then around $1\ \text{nV}$ [[75](#)].

Inductive crosstalk is driven by changing currents coupling to victim traces via stray inductance. The electrical nets in question are open, and the only currents present are for charging and discharging capacitors. Owing to conscious resonator frequency design, the maximum expected current slew rate is roughly $5\ \text{mA/s}$, and the mutual inductance of neighbouring traces on the sample is roughly $35\ \text{nH}$. In the worst realistic case, this leads to an inductively coupled noise voltage of several $100\ \text{pV}$ [[75](#)].

To arrive at the above estimates, the properties of the sample were used, along with the specifications of surrounding infrastructure, explained in [chapter 4](#). The selection of silicon nitride as substrate guarantees that direct leakage between traces won't be a problem either. To put the obtained estimates into context, one considers that the electrodes will be operated in the range of $0 - 5\ \text{V}$, as will be discussed shortly. For a driving digital to analog converter (DAC) to then be limited by crosstalk, its effective resolution on the given voltage range would have to exceed 32-bit, which yields a LSB corresponding to just above $1\ \text{nV}$.

In summary, the 5000 electrodes are routed to their corresponding

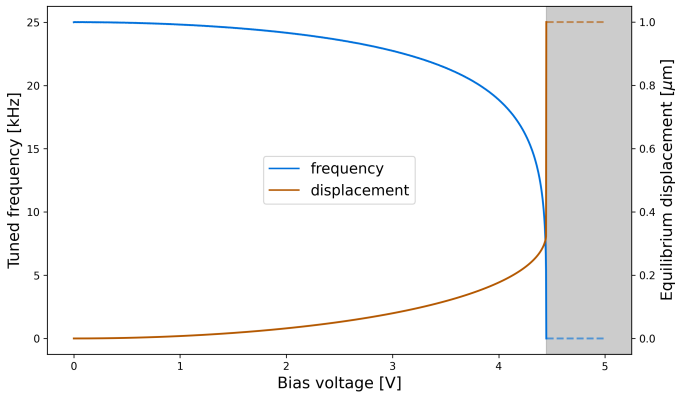


Figure 2.9: Dependence of resonance frequency (blue) and equilibrium displacement (orange) on bias voltage applied to the central electrode, obtained using [equation \(2.1\)](#) for the designed resonators at $1\ \mu\text{m}$ gap. The grey region marks the regime where the membrane is collapsed on the opposing electrode. Biases in the range of $0 - 3.8\ \text{V}$ suffice to tune the drum within the range $25 - 20\ \text{kHz}$, and at a bias of $4.45\ \text{V}$ the drum becomes infinitely soft.

bond pads using custom software, and crosstalk is not predicted to be an issue thanks to provident design on the material and resonator level. What is left on the electrical side is the determination of parameters pertaining the electrostatic interaction.

In order to determine acceptable ranges for wafer gap, electrode size and applied wafer voltage, one has to consider the target functionality. The central electrode is to provide two distinct services. When actuated with an AC signal, it can excite one of the resonator's $\text{SO}(2)$ symmetric modes. This doesn't require a strong force to be applied to the electrode, as driving modes in resonance is a very efficient process. On the other hand, the resonator can be tuned by application of a DC offset to the central electrode. This phenomenon is known as electrostatic softening [62], and deserves a brief exposition. Before delving into this effect, an argument is in order for why one desires tunability. While the design in principle yields identical resonators with identical spectra, nature is not ideal. Real

wafers have nontrivial stress distributions across their surface, and micromachining processes act (more or less) nonuniformly across the wafer surface. The fabricated resonators are then not identical and their spectra differ. To facilitate efficient energy transfer between two modes of different resonators via coupling, one then has to tune these modes to a common frequency. In general this calibration is an involved process, and will be further discussed in [chapter 4](#).

In a harmonic parallel plate approximation, the forces acting on a membrane in the design at hand are given by a spring term and a parallel plate term,

$$F_{\text{tot}}(x) = F_s(x) + F_{\text{pp}}(x) = -kx + \frac{\varepsilon_0 AV^2}{2(d-x)^2}. \quad (2.1)$$

Here we only consider a finite voltage V between the central electrode of area A and its partner at a capacitor gap d . F_s is Hooke's law [76] with spring constant k , and as usual $\varepsilon_0 \approx 8.85 \cdot 10^{-12}$ F/m denotes the vacuum permittivity and the resonator displacement x is taken positive towards the opposing electrode.

The scaling behaviours of F_s and F_{pp} guarantee the existence of a pull-in voltage V_{PI} , at which the effective spring constant $k' := -\partial_x F_{\text{tot}} = k - \partial_x F_{\text{pp}}$ vanishes. The membrane then becomes infinitely soft, and its resonance frequency is zero. As long as biases up to the pull-in voltage can be applied to the central electrode, the resonance can be tuned to any value below its nominal frequency.

The interwafer gap d can now be optimized to obtain a tractable tuning voltage range. Expanding F_{pp} to first order yields the linearized pull-in voltage $V_{\text{PI}}^{(\ell)} = \sqrt{d^3 k / \varepsilon_0 A} \propto d^{3/2}$. Whenever the gap doubles, the tuning range triples, such that gaps stable over all resonators are necessary. From a system complexity point of view, the voltage range 0 – 5 V is preferable, as it can be directly generated by many commonly available DAC chips. In the current resonator design, a gap of $1 \mu\text{m}$ yields a linearized pull-in voltage $V_{\text{PI}}^{(\ell)} \approx 5.3$ V. A more quantitative exploration was done using the full expression given in [equation \(2.1\)](#), the results are shown in [figure 2.9](#). The additional nonlinearities further add to the capacitive force, lowering the pull-in voltage into the desired range at $d = 1 \mu\text{m}$.

If samples with gaps as small as $1 \mu\text{m}$ cannot be fabricated, amplifying to larger voltages is the solution. However, there are limits to this mitiga-

tion set by breakdowns. While the dielectric strength of vacuum strongly depends on the electrode shape, breakdown field strengths probably lie in the range of $20 - 97 \text{ V}/\mu\text{m}$ [77–80], such that $V_{\text{PI}}^{(\ell)}(20 \mu\text{m}) = 475 \text{ V}$ likely leads to discharges between adjacent traces with a spacing of $4 \mu\text{m}$. A gap of roughly $1 \mu\text{m}$ shall be the goal.

Putting everything discussed in this chapter together results in the final sample design shown in [figure 2.1](#), and translates to 9 lithography masks. The predominant factor leading to successful sample design is the provident identification of interdependencies between different desired features, such that targeted decisions can be taken. With all restrictions and requirements clearly laid out, all facets of the sample can be designed to contribute to the overall goal. To give a better idea of the different scales present on the designed sample, a perspective view of the bottom layer is shown in [figure 2.10](#). The adhesive pillars being rendered in 20-fold height magnification emphasizes how stringent the requirements placed by a stable $1 \mu\text{m}$ gap, over an area 10 cm in linear extent, are.

With the sample design complete, one can move on towards manufacturing. Given the various requirements placed on the sample, this is not a trivial undertaking. In the next chapter, the micromachining process developed to fabricate this sample will be discussed, including in depth discussion on why specific steps are necessary.




Figure 2.10: Render of the final sample design, closeup highlighting the different length scales involved. Hexagons of the final lattice are visible as combination of three resonators and three central counterelectrodes that couple to the top layer. In the center of each hexagon sits a pillar of adhesive. For better visibility, the height of adhesive pillars is exaggerated by a factor of 20, the height of the metal layer by a factor of 40.

3 | Fabrication

Much like the sample design laid out in [chapter 2](#), development of the fabrication process for this sample has undergone a number of optimization iterations. While care is taken to stay close to standard processing pipelines, some peculiarities of the sample, such as yield requirements, place restrictions on the processing steps. In this chapter, the final fabrication process is discussed, with detailed explanations pertaining why and how different steps were optimized.

3.1 General Remarks

Manufacturing the outlined platform requires optimisations of standard techniques used for MEMS fabrication. Conventional fabrication protocols often employ massively parallel manufacturing followed by post selection, which can enable tolerance for almost arbitrarily low yields. With the system at hand being implemented in a wafer scale sample, defects can not be selected away a posteriori, instead they degrade the functionality of the system. Failed resonators will stay failed in the final system, and the complex wiring necessary implies a few unfortunately placed small scale defects can electronically affect hundreds of resonators. The solution is optimising the fabrication process to achieve as close to perfect yield as possible.

The fabrication process has undergone several optimization iterations,

with the goal of maximizing resonator yield. As explained above, no post selection of functional resonators is possible, all of them are part of the device, such that every failing drum degrades the overall functionality and should be avoided. Before walking through the individual processing steps, we mention four general factors that have been observed to impact yield.

1.) Masks should be cleaned before every contact lithography exposure. If a close to perfect yield is the goal, any contamination on the mask is unacceptable.

2.) Photoresist spun on top of gold should always be overexposed by 20%. Substrate reflectivity strongly impacts the dose deposited in the resist, and underexposed resist can lead to unexpected results, especially in thick layers.

3.) Developed resist should always be subjected to a short low intensity plasma ashing step before wet etching. Roughening up the surface critically improves wetting and thereby etching uniformity.

4.) Long Cr/Au traces are susceptible to destruction in plasma ashers. While the precise pathway is not known to the author, the traces have a chance to develop black spots and lose continuity. A sample that has been affected by this effect is shown in [figure 3.1](#). An RIE step with high O₂ pressure (isotropic etch) was found to be a safer alternative for dry resist cleaning.

Before a fabrication process can be developed in detail, the substrate stackup needs to be settled. In [chapter 2](#) it was argued why the resonator material is to be silicon nitride. To provide optical access to the resonators, necessary for drum readout via interferometry, the material below the drums is to be completely removed. A convenient stackup is given by silicon nitride on silicon dioxide on a silicon substrate. Silicon is convenient as the Bosch process allows the etching of vertical high aspect ratio holes, and selectivity allows silicon dioxide to act as stop layer. The final resonator release is facilitated by BHF (buffered hydrofluoric acid), an etchant removing silicon dioxide at a much higher rate than silicon nitride.

The thicknesses of the individual layers yield more degrees of freedom. As was argued in [chapter 2](#), silicon nitride resonators of various thicknesses can be manufactured, and a thickness of 500 nm was identified as a conventional value. The silicon dioxide stopping layer needs to provide enough material to absorb the inevitable Bosch overetch, but

should be thin enough to allow convenient removal with BHF. Initially, a 300 nm thick silicon dioxide layer was chosen, later this was found to be insufficient for the observed Bosch process inhomogeneity. It was then upped to a final value of 400 nm, which is still on the rather thin side, but just sufficient. Thickness of the silicon substrate must be such that it provides structural stability without degrading Bosch verticality. Conventional thicknesses range from roughly 300-500 μm , the chosen layer thickness is roughly 380 μm .

Another selection that must be settled beforehand is the adhesive used to waferbond the two device layers. Three candidate materials were considered for this process. The initial favourite was polyimide, as collaborators had worked with it before and knew how to achieve good bonds. Unfortunately, the specific polyimide product uses NMP as solvent, which was added to REACH via ANNEX XVII [81, 82]. While the ban of NMP might vary on a per cleanroom basis, for the author the desired polyimide product was unavailable.

As alternative adhesives, benzocyclobutene (BCB) and SU-8 were evaluated. Several earlier studies [83–91] involving waferbonding with BCB, SU-8 or more generally for CMUT devices were used as reference points to develop waferbonding processes for a direct comparison. Using BCB as bonding agent yielded the highest bond strength (wafers break before bond), but it requires a hard bake at temperatures in excess of 200° C, which puts stress on the other sample components. Performing hard bakes at lower temperatures might be a possibility, but this might well yield uncontrolled, unrepeatably and potentially unstable results as the BCB doesn't fully settle. As a further downside, photo-BCB was not easily available, such that dry structuring (RIE) was necessary. Using SU-8, a negative photoresist, on the other hand resulted in weaker, but still sufficiently stable bonds. The lithographic structuring and lower hardbake temperature (120° C) along with good availability led to SU-8 being the chosen adhesive compound.

With all decisions pertaining material selections completed, a fabrication process can be developed. In the next section, this process will be introduced, and various pitfalls and optimizations discussed.




Figure 3.2: Process sequence for the fabrication of the RBComb device. Photoresist masking layers have been omitted. The topmost panel shows the initial wafer stackup with 500 nm super low stress LPCVD silicon nitride on 400 nm wet thermal silicon dioxide on a $380\ \mu\text{m}$ silicon substrate.

3.2 Fabrication Process

A simplified fabrication process sequence is shown in [figure 3.2](#) and is based on surface micromachining techniques. The starting point are two 4 inch wafers with stackup shown in the top panel of [figure 3.2](#), procured from Si-Mat. One of those wafers, from here on referred to as "bottom wafer", is used to manufacture the bottom layer of the sample, while the other wafer, the "top wafer", will be turned into the top layer. See also [figure 3.5](#). As the bottom layer includes more functionality than the top

layer, some manufacturing steps are only carried out on the bottom wafer and will be indicated as such.




Figure 3.1: Damage on the gold layer caused by exposure to a 600 W oxygen plasma for 5 minutes.

In the first step, [figure 3.2 \(a\)](#), the metal layers are deposited and structured. First Cr/Au (3nm/50nm) wiring is deposited by evaporation and structured by a wet etch (AZ ECI 3012). To produce the wire bonding pads on top of the already deposited gold on the bottom wafer, 1 μm Au is sputtered and patterned in a lift-off process (LOR 10B, AZ4533). Different deposition techniques were chosen for these steps due to the combination of high deposition efficiency in sputtering, and convenient availability of the evaporator. The Cr layer is necessary to improve the weak adhesion of Au to silicon nitride. The large number of wirebonds to be placed require a high reliability during bonding, which can be guaranteed by 1 μm thick gold pads. An impression of gold traces is shown in [figure 3.3](#)

efficiency in sputtering, and convenient availability of the evaporator. The Cr layer is necessary to improve the weak adhesion of Au to silicon nitride. The large number of wirebonds to be placed require a high reliability during bonding, which can be guaranteed by 1 μm thick gold pads. An impression of gold traces is shown in [figure 3.3](#)

In the second step, [figure 3.2 \(b\)](#), the nitride layer is structured. This is done with reactive ion etching (AZ1512) stopped on the silicon dioxide layer. As the metal traces run very close to the drum holes and accidental masking should be avoided, another Au/Cr wet etch is carried out once the nitride mask has been applied.

In the third step, [figure 3.2 \(c\)](#), SU-8 photoresist is deposited on the bottom wafer and structured via lithography. It is this layer that will form the bonding agent connecting the two wafers. The properties of SU-8 change dramatically with baking times and temperatures, and failure to follow a well optimized treatment process will prevent a successful bond down the line. Firstly, the surface is treated with an Ar clean and a dehydration bake at 180° C is carried out. This step eliminates adhesion issues that were observed on the substrate at hand. Once SU-8 is applied, a two step pre exposure bake is done at 65° C for 1 minute, then 95° C for 3 minutes.

After exposure, another two step bake is carried out, at 65° C for 1 minute, then 95° C for 2 minutes. The two step nature of the bakes is critical, and after each bake the SU-8 must be left to cool off slowly. Failing to do so introduces visible inhomogeneities in the structures, which can lead to delamination and failed bonds. Finally, SU-8 is developed, making sure no residues are left on the wafer. This can be achieved by using two different beakers of developer.

In the fourth step, [figure 3.2 \(d\)](#), a protective coating is applied to the front side of both wafers, as processing will continue with a backside deep reactive etching (DRIE) step. The predominant gains on yield achieved via process optimizations originate from using this coating, instead of a more conventional DRIE preparation involving a carrier wafer fixed to the photoresist protected frontside using wax. Following the conventional approach means a carrier wafer will need to be removed from

(almost) liberated resonators. It also leads to (potentially undetected) lower cooling efficiency during DRIE processing, which can degrade the process control and lead to reactions at the wax resist interface, making the resulting mixture insoluble. The alternative coating consists of a thin sputtered silicon dioxide layer (100 nm), a thicker sputtered Al layer (500 nm) and a final layer of photoresist (AZ4562). The reasoning behind this stackup is the following. Optimally, only Al would be used, covered in a layer of photoresist for DRIE chuck protection. The inconvenient mobility at Au/Al interfaces means within short amounts of time Al_xAu_y forms on top of the traces, electrodes and bond pads, which can no longer be cleaned without damaging the sample. This problem is solved by depositing a thin silicon dioxide diffusion barrier between the aluminium coating and the device. Once all processing is finished, the protective coating can be fully removed using exclusively wet processes, crucially precluding any




Figure 3.3: Impression of gold traces.

mechanical steps and solubility issues.

In the sixth step, [figure 3.2 \(e\)](#), the silicon is patterned from the backside with deep reactive ion etching. The etch is stopped on the oxide layer. For this step, having at least a 400 nm silicon dioxide stopping layer is crucial. In earlier process iterations, a 300 nm stopping layer was used and routinely penetrated due to inhomogeneous etch rates across the wafer. Even within single drum holes this can lead to issues, as Bosch processes typically form bathtub profiles, in the case at hand attacking the drum in the center while the sides are still covered in silicon. Once all silicon has been removed, it is crucial that all remaining teflon is removed, else wet etchants in the next processing step will not enter the drum holes.




Figure 3.4: Finished sample, view through a top layer drum onto the bottom layer.

In the seventh step, [figure 3.2 \(f\)](#), the protective coating is stripped and the silicon dioxide stopping layer removed. First the stopping layer is removed with BHF, then we strip the protective coating with a succession of DMSO (120° C), Aluminium etch and BHF. Finally, the sample is rinsed with water, placed in IPA and subsequently dried. The liquid dried from the sample should have as small a surface tension as possible to lower chances of damaging the resonators. Now the two individual wafers are finished. An impression of liberated resonators is shown in [figure 3.4](#). As waferbonding is the next step, care should be taken to keep the wafer surfaces as clean as possible.

In the eighth step, [figure 3.2 \(g\)](#), the two wafers are bonded. Top and bottom wafer are placed on top of one another, and the bonding is done by applying 3 bar on the stack for 40 minutes at 120 °C. The process is carried out in vacuum. Different steps in this process do not commute, for the best results first the vacuum should be pumped, then the pressure applied, and finally the temperature raised. SU-8 being a thermosetting polymer

means the success of this process depends largely on the deposition, development and hardening procedures carried out in step (c). Alignment between the wafers is critical to ensure sufficient overlap between corresponding electrodes.

In the ninth step, [figure 3.2 \(h\)](#), the bonded sample is fixed to a PCB and electrical connections are established by wire bonding. Semi-automatic aluminium wedge bonding is employed. The available alternative, gold ball bonding, tends to exhibit lower process reliability, which would make the bonding of 5000 pads intractable. While several contacts had reported persistent issues with wirebonding reliability, these difficulties were avoided through precise steps. Careful initial tool setup, bond shape optimization, and bond pad design pertaining shape and material led to a process that works well and requires almost no manual per-bond fine tuning. This allowed operating the wire bonder in automatic mode, placing 10 bonds per trigger, which greatly increased processing speed. In [figure 3.5](#) the sample is shown before wire bonding.




Figure 3.5: Impression of the finished sample, before wire bonding.

In this chapter, the process developed for the fabrication of the sample introduced in [chapter 2](#) has been explained. Due to the monolithic design of the sample, special emphasis was placed on optimizing the process for maximized yield. This chapter concludes the discussion of the sample. In the next chapter, focus shifts to the experimental system designed to drive and probe the sample.

4 | Setup

Even the best, most interesting sample has no use without a corresponding experimental setup able to excite and measure it. To harness the full potential of the sample developed in the previous sections, precise and well synchronized control of 5000 analog channels and a system for optical measurement of resonator motion are a basic requirement. In this chapter, the custom developed experimental system facilitating control of the sample will be introduced.

In order to support the sample that was introduced and fabricated in [chapters 2](#) and [3](#), certain functionality must be provided by setup infrastructure. The sample exposes roughly 5000 independent analog nets, to which voltages in the range 0-5 V are to be applied. These 5000 voltage sequences are to be independently programmable, and must be applied to the different channels synchronously. As the resonators have resonance frequencies around 25 kHz, the Nyquist frequency of the voltage generators should lie safely above this value. The surface to weight ratio of the resonators means the sample must be operated in vacuum for finite Q factors to arise. A measurement system tracking the motion of individual drums on both sample layers is required to observe system dynamics. Finally, as the resonators are tension dominated, a temperature stabilization mechanism is desirable.

As no commercial setup fulfilling the above points is available, a custom system implementing all of these functionalities has been developed. A diagram of the system is shown in [figure 4.1](#). The components making

up the experimental system can be categorized into two 4 groups: The signal processing chain (red), the measurement system (green), the environmental control system (blue), and physical support components (orange). Before discussing each individual component in detail, an overview over these four groups shall be provided.

On the digital side, the signal processing chain can be split into two parts: real-time and non-real-time. The real-time subsystem handles tasks that need precise synchronization, such as applying an arbitrary sequence of electrode voltages and simultaneously recording a measurement. The asynchronous subsystem programs the real-time subsystem and schedules and triggers tasks on it. The synchronous subsystem is implemented on an optical star network of 11 field-programmable gate arrays (FPGAs). The hub node of this network is called Bridge, while each peripheral node corresponds to one Minicircuit, which is where the analog voltage sequences are generated. The asynchronous subsystem consists of a general purpose PC and communicates with the synchronous hub via UART. A python API has been developed, which facilitates effective experiment development without in-depth knowledge of the underlying communication protocol or hardware intricacies. There are several custom developed PCBs that are part of the signal processing chain, namely the substrates of the Bridge and Minicircuit components, and the Breakoutboard which routes the 5000 electrical nets. A final, but important, consideration concerning the signal chain is the grounding concept. Ground loops in the voltage generation system must be avoided, or inductively coupled noise will become a problem. This is achieved by having Bridge and Minicircuit components communicate optically, such that they can be electrically disconnected. There are then two ways in which the Minicircuits and the sample can be grounded. One option is to power all Minicircuits from common rails, and connect their common ground net to the Breakoutboard at a single point. The other option is to power every Minicircuit from an individual ungrounded power supply, and connect all Minicircuit grounds to the Breakoutboard. The Breakoutboard provides one ground pad for each Minicircuit, and all of the associated nets are internally connected in a single point. For either of those options, the Breakoutboard ground connects to the sample ground with a single cable.

The heart of the measurement system is an interferometric distance measurement device (IDS 3010). Two independent axes of this device are

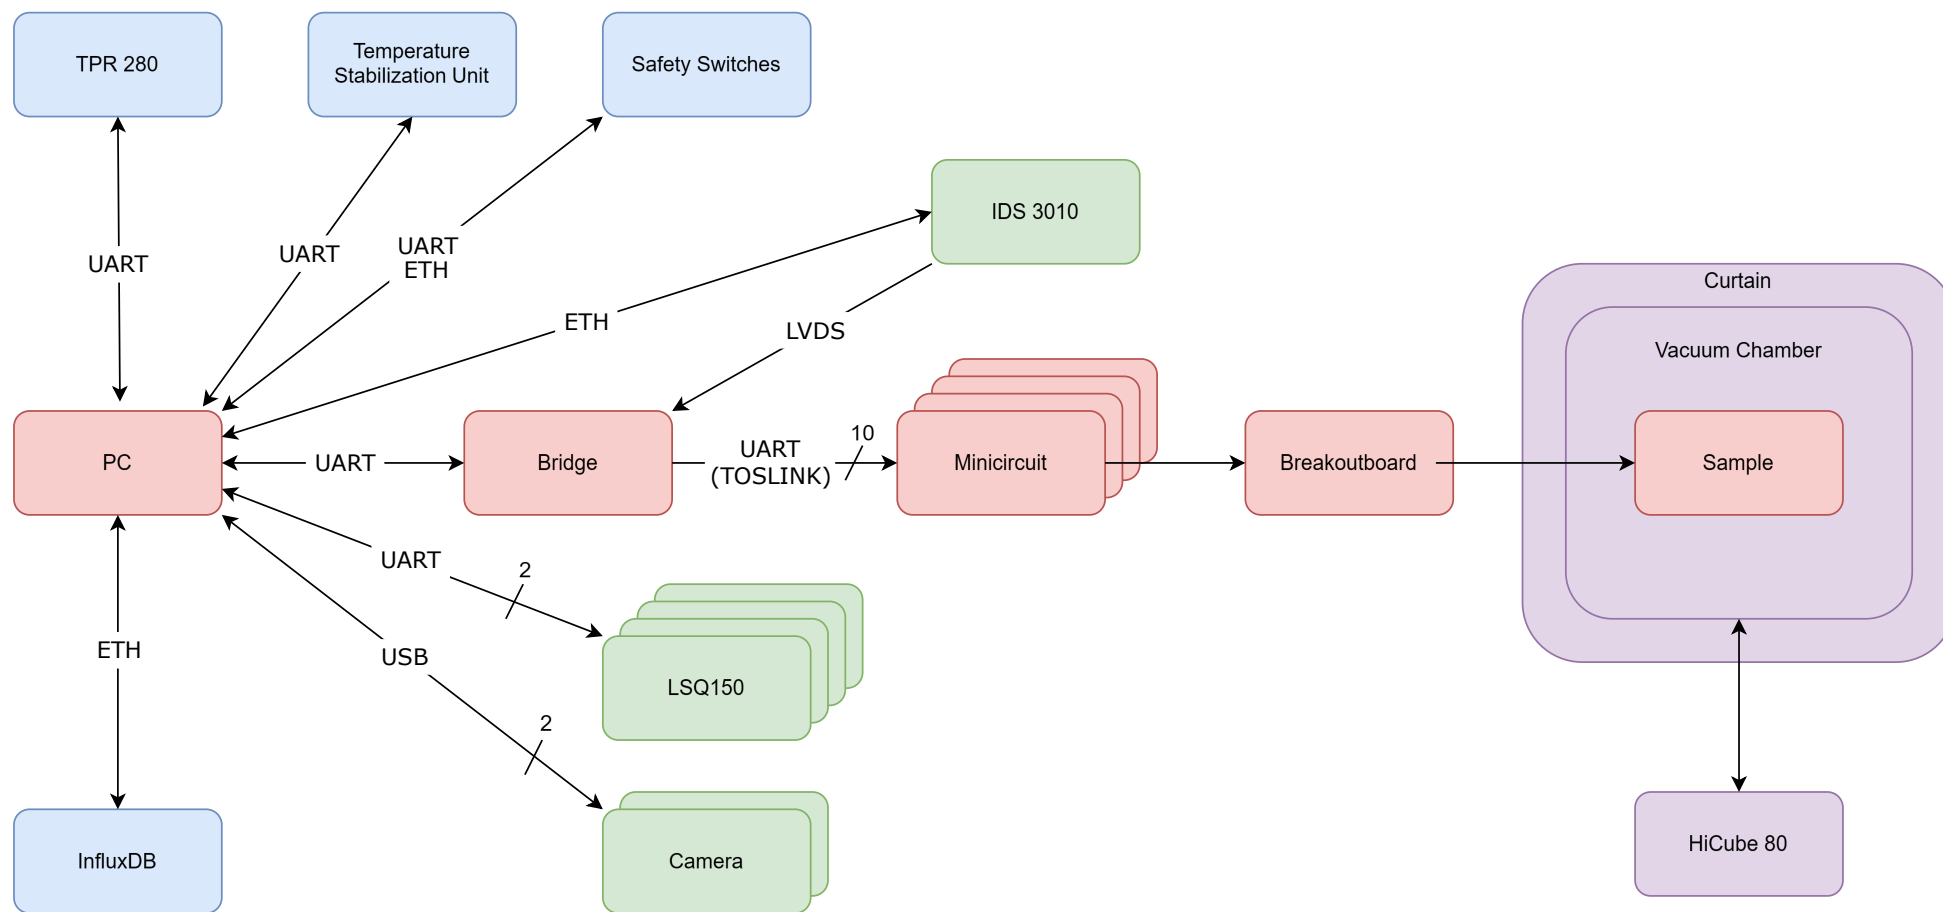


Figure 4.1: Diagram of the experimental system with different subsystems coded in different colours. Red: signal processing chain, green: measurement system, violet: physical support components, blue: environmental control system. Arrows indicate flow of information or directional influence, with labels describing the used technology where applicable.

mounted on two sets of linear x-y stages (LSQ150) together with 2 digital microscopes (Camera). One set of stages is placed on each side of the sample, such that both sample layers can be measured.

The environmental control system consists of a pressure measurement device (TPR 280) and a temperature measurement and stabilization unit (Temperature Stabilization Unit) with safety switches. This is complemented by an InfluxDB database, to which measured environmental data, along with further dynamic setup information, is uploaded and then visualized on a Grafana dashboard.

The physical support components include the vacuum chamber in which the sample is placed, a vacuum pump, isolating installations (Curtain) and an optical table, on which the setup is mounted.

With a high level overview of the system presented, the remaining sections of this chapter will introduce all of these components in detail, starting with the physical support components.

4.1 Physical support components

The major supporting apparatuses are a vacuum pump and a vacuum chamber mounted on an optical table using a custom stand. The setup is then surrounded with a curtain decoupling the experiment region from outside atmospheric fluctuations to a certain extent. Should external influences be found to be a problem, additional layers of insulation could be added. A downside to operating the sample in vacuum is reduced thermal equilibration. Silicon nitride isn't just a great electrical insulator, it also barely conducts heat, making the suppression of convection all the more significant. As a mitigating factor, there is a lot of metal on the sample, but the possibility of local temperature differences should be kept in mind.

A photograph of the front side of the setup is shown in [figure 4.2](#). In the following sections, some notable features of the vacuum setup will be discussed.

4.1.1 Vacuum chamber

The vacuum chamber was custom designed for the target application and machined out of aluminium. A front view is shown in the center of [figure 4.2](#). The chamber consists of two almost identical parts that are joined




Figure 4.2: Front side view of the setup. Notable features include: Vacuum chamber (center) with attached SUN boards (white PCBs) and TPR 280, Breakoutboard with mounted sample (visible through window), ten attached Minicircuits and thermistors connected with twisted triplet cables, measurement stages mounted on the stand with attached laser fiber in front of the window.

with 14 M5 screws, sandwiching the Breakoutboard PCB in between. Sealing is achieved with o-rings pressing onto the PCB surface from both sides. This geometry was chosen because it allows electrical connections to transcend the vacuum boundary via traces on the PCB, rather than requiring dedicated feedthroughs for 5000 electrical nets.

Each chamber part contains a flange to allow connection of a pump on one side and a pressure measurement device on the other. Equilibration between the two parts is enabled by holes cut into the Breakoutboard. Above the sample placement region on the Breakoutboard, each chamber part contains a 139 mm diameter window to allow optical measurements on both resonator layers. To avoid multipath superposition issues in interferometric measurements, the windows are slanted at an angle of 2 de-

gress relative to the plane of the sample. The quartz windows are 8 mm in thickness, which is the minimum admissible value still guaranteeing stability [92].

4.1.2 Vacuum pump

The vacuum is generated by a combined turbo and backing pump HiCube 80 by Pfeiffer Vacuum. It was initially connected to the chamber with an aluminium flex pipe, but this was later swapped with a PVC pipe to reduce the transfer of vibrations and heat into the experimental region. If further phononic decoupling should be necessary, the pipe could be routed through a granular material like sand. As the pump is located close to the chamber, even without direct conduction the heat produced by it drives convection and must be handled consciously.

This concludes the discussion of the physical support components. In the next section, the devices making up the measurement system will be introduced.

4.2 Measurement System

The Measurement System consists of three components: an interferometer, digital microscopes and a setup of linear stages. Measurement of individual resonator motion is done by recording changes in the distance between a laser fiber and a resonator. To allow measurement of all resonators, the interferometer fibers are mounted on linear stages, which facilitate programmatic repositioning.

4.2.1 Displacement Measuring Interferometer IDS3010

The Attocube IDS3010 measures displacements with a picometer resolution. It supports up to three measurement axes, different sensor heads for various applications, and streams measured data in HSSL (high speed serial link) via LVDS (low-voltage differential signaling). The precise properties of the interface can be adjusted, the device as used transmits 48-bit data packets at 430 kSPS. Each data packet contains the current distance between sensor head and reflecting object in picometers.

Several different laser heads are available for this system. The one used in the setup has a focal length of 17 mm, which is the longest focal length available. While the vacuum chamber was designed to minimizing the laser head-sample distance, the minimal achievable distance is still in excess of 20 mm. Thus the drum will always be out of the laser focus, making the spotsize finite. As was mentioned in [chapter 2](#), the resonators are designed with a diameter of 1 mm, such that targeting drums in a way that lands a large portion of the incoming beam on its surface is easy enough. Resonators with smaller diameters could suffer from interference between on- and off resonator paths.

System setup and laser alignment are carried out using a webinterface accessed by connecting to a webserver run on the device. The interface provides real-time signal strength readings, and the user adjusts sensor head positioning accordingly. To facilitate the necessary precise corrections via tilting, the sensor heads are fastened to kinematic mounts. This is shown in [figure 4.2](#), the yellow fiber passing vertically through the center corresponds to one measurement axis of the IDS3010, and terminates on a mount in front of the sample. Further positioning degrees of freedom are added by the stages described in the next section.

4.2.2 Linear Stages Zaber LSQ150

The kinematic mounts with fastened interferometer sensor heads are mounted to linear stages, as shown in [figure 4.2](#). Two pairs of stages are used, one in front of each vacuum window and each stacked to allow access to every resonator by x-y motion. The stages have a microstep size $< 0.1 \mu\text{m}$ and a repeatability $< 2 \mu\text{m}$, such that targeting individual resonator centers in reproducible fashion is easily possible.

Control of the stages is handle by the PC. A wrapper around the Zaber supplied API has been developed, that provides homing, positioning and position interrogation functionality. Using the stages, automatic experiments can be programmed that scan many different resonators in sequence.

4.2.3 Digital Microscopes

A mounting system developed in the context of a semester thesis [\[93\]](#) allows attaching digital microscopes to the stages, next to the sensor heads.

These microscopes allow in-situ introspection of individual resonators, such that instantaneous integrity of drums can be ensured and anomalous observations can potentially be traced back to damage on the sample.

The used microscope is a DinoLite AM73515MZTL, which connects to the PC via USB 3.0 and provides long working distance, up to $140\times$ magnification, 2560×1920 resolution, 45 FPS. These properties are sufficient to resolve individual metal traces on the sample from the mounting distance, which is several centimeters. While the microscopes are currently only used for live introspection, in a later iteration they could be programmed to record complete images of the sample through automatic stitching.

This concludes the discussion of the measurement system. It has now been outlined how the bones of the setup are constructed and what measurement capabilities exist. In the next section, focus shifts towards the components part of the signal chain, beginning with the Breakoutboard to which the sample is attached.

4.3 Breakoutboard

The Breakoutboard is a signal distribution PCB that distributes electrode nets from wirebonding pads to Minicircuits. A photograph of a Breakoutboard is shown in [figure 4.3](#).

The individual features on the PCB are the following. Within the sample attachment region delimited by a white circle, 5000 wirebonding pads are available connect sample electrode nets to the PCB. A rectangular hole within the sample region allows optical access to resonators on the bottom sample layer from the backside. Small circular holes around the sample region are used to fasten the PCB to the wirebonder. The rectangular hole on the right side of the PCB allows pressure equilibration between the two vacuum chamber parts, and the large circular holes spread along an oval path allow vacuum chamber screw feedthrough. Along the edges of the PCB, ten SEAM-50-03.0-S-10-2-A-K-TR zero insertion force connectors allow the reversible electrical and mechanical connection of Minicircuits. A pad part of the Breakoutboard ground net is exposed next to each zif connector.

Due to the high net count, the PCB was generated with a custom developed python script. Routing 5000 nets interactively is not an efficient process, especially when designs are fluid and several iterations need to be




Figure 4.3: Photograph of the Breakoutboard. The sample is to be mounted in the white circle, wire bonding pads are clearly visible. Zif connectors are placed along the boundary for Minicircuit attachment. Circular holes are chamber screw feedthroughs, square holes facilitate optical inspection and pressure equilibration.

performed. The resulting file was then loaded into Kicad for final manual adjustments, design rule checks and gerber generation.

The demanding functional requirements imposed on the PCB necessitate high precision manufacturing technology. 5000 densely placed wire-bonding pads need to be routed out of the sample region towards zif connectors. The resulting signal density is such that a 16 layer PCB only becomes routable when trace width and spacing are set to $100\ \mu\text{m}$, which is close to the cutoff for what's reasonably manufacturable in standard PCB technology. In fact, due to issues encountered during manufacturing, trace width and spacing are relaxed to larger values as soon as traces leave the high density area. The PCB being just about routable is not a coincidence, the sample was designed with a corresponding manufacturable Breakout-board as conscious constraint. Similarly, the Breakoutboard was routed with external requirements in mind. While all 16 layers are required to lead all nets out of the wirebonding region, outer traces are via-ed to inte-

rior layers before crossing the vacuum boundary, such that the o-ring seal is not degraded by unnecessary surface roughness.

The Breakoutboard was manufactured by the Advanced PCB track of PCBWay. Wirebonding is facilitated by an ENIG (electroless nickel immersion gold) surface finish. If bondability is not guaranteed by the manufacturer, wirebonding tests should be performed before ordering more batches of Breakoutboards.

The Breakoutboard represents a purely passive object in the signal chain. In the following sections, the digital control systems of the signal chain, which are themselves controlled by the PC via an API, will be introduced. To ease understanding, this exposition will start with a description of this API in the next section.

4.4 RBComb API

The RBComb API is written in python and allows the user to send useful commands from the PC down the signal chain without in-depth understanding of the system specifics. The goal is to facilitate the programming of full experimental protocols. An experiment consists of 5000 independently programmed sequences of analog voltages, each applied to one electrode and all played in lockstep. In parallel to this driving process, with defined temporal relation to the voltage sequences, a measurement of resonator motion shall be recorded.

As a brief aside, it be emphasized that the API merely provides the capability to apply arbitrary voltages and record motion. The process of figuring out *what* voltages should be applied, is left to the user. While this fact might seem obvious, the implied task is hard in general. From a static point of view, variations across the sample, like a position dependent interlayer gap, mean the force effected by an applied voltage depends on the position of the electrode. What is more, properties of the sample might change dynamically. As voltages are applied to electrodes, gaps and capacitances of the associated capacitors change, effecting electrostatic tuning and dynamic changes to the voltage-force relations. If precise control of forces is necessary, careful calibration is imperative.

In order to understand the API architecture, some specifics of the hardware must be known. The digital components of the signal chain consist of three logical devices: the Bridge, the laser interface, and ten identical

Minicircuits. All communication enters the FPGA network via the bridge, which constitutes the network hub, and is then distributed to the other devices. These three logical components fulfill different roles, which will be briefly laid out in their corresponding sections.

Corresponding to the signal chain, the API contains three classes `SerialBridge`, `LaserInterface`, and `RBComb576`, each of which communicates with one logical device. Furthermore, the API provides helper functions that aid the user in correctly calculating data to be transmitted to the FPGA network. In the following sections, the communication classes will be described, along with high level information on how the system works. As the goal is merely to give an idea of the capabilities of the system, the function description shall be limited to signatures and accompanying brief comments. More in-depth information about the system will be provided in [sections 4.5](#) and [4.6](#).

4.4.1 Class `SerialBridge`

The `SerialBridge` class communicates with the Bridge. As central hub, the bridge receives and forwards communication, and it synchronizes the Minicircuit devices while a measurement is running. This synchronization is done by having the bridge send periodic clock pulses to all Minicircuits in parallel, to signal that the next voltage sample is to be generated. More information about this procedure is provided in [sections 4.5](#) and [4.6](#).

The methods provided by the `SerialBridge`:

- `__init__(serialPort)`: The constructor requires the serial port to which the Bridge is connected.
- `selectBoard(boardID)`, `selectAllBoards()`, `selectLaser()`: Select either a specific Minicircuit by its ID (see [section 4.6](#)), all Minicircuits, or the laser interface for subsequent command addressing. Before a device can be talked to, it must be selected.
- `clockOn()`, `clockOff()`: Enable or disable the experiment synchronization clock. When the clock is on, the Minicircuits produce voltages and the laser interface records data, but all commands other than `clockOff()` are ignored.

4.4.2 Class LaserInterface

The `LaserInterface` class communicates with the laser interface, which handles all communication with the IDS3010. It records and stores measurement data as programmed in the experiment setup, and sends that data to the PC upon request. The provided methods are:

- `__init__(bridgeLink)`: The constructor requires a `SerialBridge` object
- `scheduleAcquisitionCarli(axis, delay, length)`: Schedule an acquisition of a trace of `length` measurement data samples from axis `axis` to begin after `delay` experiment clock ticks have passed
- `requestTimeTraceCarli()`: Fetch the last recorded measurement trace

4.4.3 Class RBComb576

The `RBComb576` class communicates with the Minicircuits. As it handles all programming of signal generation sequences, it is the most intricate of the three API communication classes.

The functionality of a Minicircuits is roughly as follows. On the general level, the DACs on a Minicircuit provide two different gain settings, and the output of the DACs can be ramped up or down to avoid hazards. Per Minicircuit, 5 signal generators are available to produce the desired output voltages: 4 sine generators and 1 interpolating wavefunction generator (IWG). For each sine generator, frequency ν and phase ϕ can be set individually. Frequency sweeping is also available on a per-generator basis.

The IWG provides configurability on a per-output channel basis. For each channel, it can store two sets of Taylor coefficients that define two voltage functions, termed A and B, which can then be played forwards and backwards in a programmable braiding sequence. To put it differently, a braiding sequence is a sequence of As and Bs, together with the information in which direction, forwards or backwards, each of the corresponding functions should be played. An example braiding sequence of length 3 is [A, A, B] with [forwards, backwards, forwards].

Each Minicircuit generates voltages for 576 analog channels. The global outputs of the 4 sine generators are mixed with amplitudes defined

per channel and added to the channel specific output of the IWG, to yield the voltage output on channel i :

$$V_i(t) = \sum_{k=0}^3 A_{k,i} \cos(2\pi\nu_k t + \phi_k) + f_i(t), \quad (4.1)$$

where f_i is the currently selected wavefunction (out of the two, and played forwards or backwards) on channel i . Apart from choosing a per-channel amplitude, each of the 5 generators can also be globally enabled or disabled in a scheduled way, which allows different stages in an experiment. For example, in a first stage the resonators could be pumped by a sine generator, and in a second stage the couplings could be modulated using the IWG.

To expose the above functionality to the user, the RBComb576 class defines several methods. To prevent confusion in the coming exposition, the reader be reminded that Minicircuits only change their voltage output when an experiment clock is supplied by the Bridge. In general, many functions depend on an application of this external clocking signal. The methods for global Minicircuit setup are:

- `setGain(highGain)`: If `highGain` is `True`, high gain mode is enabled.
- `outputRampUp()`, `outputRampDown()`: Within the next 5 seconds worth of experiment clock, the outputs of all channels will be ramped to their assigned values.
- `setCurrentState(state)`: `state` is a list of 5 binary integers `[singen0, ..., singen3, iwg]`, which enable (1) or disable (0) the corresponding generator.
- `setNextState(state, delay)`: The state is changed to `state` after `delay` experiment clock ticks.

Setup of the 4 sine generators is possible with the following methods:

- `setPhase(sineGenerator, phase)`: Set sine generator phase
- `setFrequency(sineGenerator, newFrequency)`: Set sine generator frequency

- `setAmplitude(sineGenerator, channelNum, targetAmplitude)`: Set sine generator mixing amplitude for a specific channel
- `setSweepFrequencyStep(sineGenerator, newFrequencyStep)`: Set sine generator sweeping speed
- `setSweepFrequencyMax(sineGenerator, newFrequencyMax)`: Set sine generator sweeping stop frequency

In the above signatures, `sineGenerator` is always an integer $\in \{0, \dots, 3\}$ enumerating the target sine generator. For the general setup of the IWG, the following methods are provided:

- `setTaylorCoefficients(channelNum, funSel, coeffs)`: upload Taylor coefficients for function $A(B)$ (`funSel=False(True)`) for channel number `channelNum`
- `setBraidingLength(newBraidingLength)`: set the number of elements to be played in the braiding sequence
- `setBraidingFrame(newBraidingFrame)`: start the braiding sequence at a phase offset
- `setBraidingSpeed(newBraidingSpeed)`: set the speed at which the braiding sequence runs

The following functions set properties of individual elements of the braiding sequence. They all take a list of zeros and ones as argument, and the i th entry affects the i th element in the sequence.

- `setBraidingAB(braidingAB)`: element will be function A for list entry 0, and function B for list entry 1.
- `setBraidingOrder(braidingOrder)`: element will be played forwards for list entry 0, and backwards for list entry 1.
- `setBraidingHold(braidingHold)`: element will be played normally for list entry 0, and replaced with its initial voltage for list entry 1.

All necessary functionality is provided in the API methods introduced so far. However, there are a few rough edges, which are rounded off by additional helper functions.

Many of the introduced API methods take arguments dimensioned in board units. For example, time is measured in ticks of the system clock, signal amplitude is measured in DAC LSBs. For all required obscure units, a conversion method is available in the helper class `BoardUnits`. When defining functions for the IWG, an intricate subinterval splitting and expansion procedure is necessary, as will be described in [section 4.6.2](#). The helper class `FunctionInterpolator` takes as input a python function and converts it into corresponding uploadable Taylor coefficients. It also provides predefined interpolation functions like a constant voltage offset, which can be directly uploaded to the FPGA. Finally, helper class `BoardPinout` converts between pin matrix coordinates and output integer indices. This functionality is mainly useful for debugging, taming the complexity of the system requires direct electrode-pin correspondence lookuptables.

The introduction of the python API has developed a grand overview of the capabilities of the signal chain. With this background information, the individual components of the FPGA network are now more accessible. They will be introduced in the following sections, starting with the Bridge.

4.5 Bridge

The bridge represents the hub of the FPGA network. It acts as the interface for communication between the synchronous and asynchronous parts of the signal chain. Incoming commands are interpreted and depending on their type either handled internally, or forwarded to downstream devices. As network center, the bridge programs, triggers and synchronizes the minicircuits via optical interconnects and stores relevant measurement data incoming from the interferometer in a time coherent manner.

In a system controlled by a nontrivial FPGA network, synchronization is of major importance. As was shown in the API discussion of [section 4.4](#), an experiment consists of 5000 sequences of voltages to be applied to the resonator electrodes in lockstep. Each of the 10 minicircuit FPGAs controls roughly 500 of these electrodes, and each minicircuit FPGA is clocked independently. As a set of clocks is typically affected by skew, the 10 mini-

circuit FPGAs must be explicitly synchronized via an external reference to ensure the experiment is carried out as intended and there's no growing dephasing between the application of voltages to different electrodes. In the system at hand, where FPGAs are clocked by Renesas XLH536100 CMOS oscillators with a ± 25 ppm frequency stability [94] and DAC samples are generated at 140 kHz, two 25 kHz sine waves could accumulate a 180° phase difference within less than a second¹.

This issue is prevented by the bridge in the following way. When an experiment is running, the bridge sends periodic commands to the minicircuits to signal that the next set of samples in the sequence should be generated. Once triggered, each minicircuit then takes the same roughly 650 (local) clock cycles until its DACs are fully switched to the new voltages. With this setup, the maximum voltage application delay is independent of sequence length and bounded by roughly one minicircuit clock cycle, or 10 ns, which is the maximum delay of trigger reception between different minicircuits. Other effects like propagation delay and clock skew are negligible in comparison.

As a component, the bridge consists of hardware placed on a PCB, and firmware programmed into the FPGA, which will henceforth be referred to as *gateway*. In the following section the hardware components will be explained. These consist of a custom designed PCB, which connects to a Numato Saturn device containing the controlling FPGA chip. In the section after the next, the gateway running on that FPGA is discussed.

4.5.1 Hardware

The bridge is designed as a 2 layer PCB, as shown in [figure 4.4](#). Components are placed on both, the front- and backside. The PCB was designed in Kicad and manufactured by Eurocircuits.

On the backside, there are 10 TOTX1350 TOSLINK transmitters used to send data to the 10 downstream minicircuit boards. Optical interconnection was chosen over electrical to isolate the analog signal generation stages from upstream devices. An LVDS connector facilitates a digital connection to the IDS 3010 measurement system. Three status LEDs can be connected to enable feedback to the user.

¹The hidden assumption here is that the central frequency tolerance is at best comparable to the thermal frequency stability.




Figure 4.4: Photograph of the Bridge board. Left: Backside with TOSLINK transmitters. Right: Frontside with Numato Saturn board (red) attached.

On the front side, audible user feedback can be enabled by connecting a beeper (below the expansion card), and a header allows the mounting of a Numato Saturn [95] FPGA development board. The Numato Saturn hosts a Xilinx Spartan 6 XC6SLX45 FPGA, clocked by a Renesas XLH536100 CMOS oscillator (XO). The FPGA is connected to a Micron MT46H32M16LF low-power double data rate synchronous dynamic random access memory (LPDDR SDRAM) chip providing an additional 512 Mb of memory. A FTDI FT2232H chip facilitates serial communication between an external computer and the Spartan 6 via a micro USB connector. Gateware is uploaded to the FPGA using a JTAG header. The Numato Saturn can either be powered through a dedicated barrel jack (5 V, tip positive), or via USB, with the source selected with the power selection header (PSH). On the bridge, the Numato Saturn is powered via USB, and this power is then forwarded to the other active components.

The Numato Saturn is a great fit for the application at hand. It is affordable, comes with adequate documentation, doesn't provide unnecessary functionality, and its 136 broken out IO pins are more than enough to address the target number of DAC channels. Optimal synergy is achieved

when the same device is used for all nodes in the FPGA network, which makes sense with the Numato Saturn.

Next, the gateway configured into the bridge FPGA will be discussed.

4.5.2 Gateway

As has been discussed, the bridge's brain is represented by a Xilinx Spartan 6 FPGA. In this section the gateway flashed onto this FPGA is introduced in detail.

Gateway development and programming file generation is done in Xilinx ISE 14.7 using VHDL. In the big picture, the bridge takes commands from the PC, and either handles them internally (when they concern scheduling, running and recording experiments), or forwards them to downstream devices (minicircuits) via optical interconnects. While a scheduled measurement is performed, the minicircuits are clocked and data received from the IDS 3010 via LVDS is streamed into the LPDDR DRAM chip, which can later be streamed to the PC via UART. The initial bridge design used internal BRAM, the switch to the harder to use SDRAM increased the recordable trace length by a factor of 100.

Receiving and storing measured distance data is not as trivial as it might seem. While the IDS 3010 continuously streams data at a well defined frequency (450 kSPS), we are only interested in the data during a scheduled period. When data is to be stored, the first sample will come at a temporal offset as the experiment clock and the laser clock are not synchronized. Thus the initial time offset must be stored, or an unknown phase jitter of up to 20° at 25 kHz is generated for every measurement. For similar arguments, it makes sense to store sample timestamps (in system ticks since previous sample) for all datapoints.

A register transfer level (RTL) schematic of the top level system design is shown in [figure 4.5](#). Data generally flows from the left to the right of the diagram, correspondingly blocks have input ports on their left side and output ports on their right side. The devices labeled `fd`, `fds`, `fdr` and `fdrs` are variants of D flip-flops, with optional synchronous set (`s`) and synchronous reset (`r`) ports. Further general purpose infrastructure includes differential input buffers `ibufds`, inverters `inv`, as well as or gates and and gates with various numbers of inputs.

The top level input signals to the FPGA are serial communication incoming from the mini USB port `RX_IN`, a 100 MHz clock `c3_sys_clk`,

an active-low reset signal `c3_sys_rst_n`, and the high speed serial link (HSSL) signals received via LVDS from the IDS 3010, namely the differential distance data pairs `LASER_SIGNAL_i` and the corresponding differential clock pairs `LASER_CLK_i`, one set per channel.

The top level output signals are three LED signals, which are routed to the user indicator LEDs on the PCB and represent different states of the system (device powered, clock on, command ignored), a `BEEPER_OUT` signal that carries the driving signal for a notification loudspeaker, the serial port transmit signal `STX` that is routed to transmit via the mini USB port, ten `OPT_OUT` signals that represent serial communication to be sent to the minicircuits, and various `mcb3` signals that are routed to the LPDDR SDRAM chip or, in the case of `mcb3_rzq`, connected to a resistor for calibration purposes.

The general flow through the bridge gateway is as follows. Serial input is deserialized by `LLSerialInterfaceR`, and subsequently decoded by a succession of logic gates. Depending on the decoding result, different modules are activated and perform actions corresponding to the input opcode and data. In the following portion of this chapter, the different custom blocks of the bridge gateway will be explained, and their functionality contextualised.

LLSerialInterfaceR

This device receives serial communication from the PC, and puts it on a parallel bus. It is a deserializer implemented as a shift register, designed to operate on a 5 MBaud input. Its inputs are a 100 MHz clock on `CLK100` and the serial input on `RX`. Once a complete packet has been received, `DR` goes high, signalling the data is available on the 8-wide `DATA` bus.

ClockGenerator

While a measurement is running, this device generates the sample clock that synchronizes the ten minicircuits while a measurement is running. It takes a 100 MHz clock on `CLK100`, and an active high reset signal on `Reset` that inhibits output pulse generation. When `Reset` is low, it divides the input clock down to 140 kHz, which coincides with the maximum sample rate of the AD5676 DAC chips on the minicircuits, and presents the resulting signal on the `Tick` port. The decode logic is such that the `Reset`

port is controlled correspondingly by opcodes that enable and disable the clock.

Beeper

The Beeper is used to generate audio signals to inform the user of invalid device usage, namely sending commands to the bridge while the sample clock is enabled, which leads to the commands being ignored. When Beep is detected high, a 3.3 kHz square wave is output on Speakerout for 150 ms. While the Speakerout output is active, Beeping is high, and low else.

SerialTransmit

This module outputs serial data that is then sent to minicircuits via TOSLINK. It is a serializer implemented as a shift register, operating at 5 MBaud output. The data to be sent is applied on the 8 wide Data bus. When DR is detected high, the input data is latched. When DR is next detected low, the latched data is shifted out via the TX port at 5 MBaud.

LaserInterfaceWithCarli

This component handles all commands relating to the data received from the IDS 3010. When a `selectLaser()` call is issued in the API, this is the endpoint that is selected. `LaserInterfaceWithCarli` takes as inputs a 100 MHz clock on `c3_sys_clk`, the system wide reset signal on `c3_sys_rst_n`, IDS 3010 data and clock signals for all three channels on `LCi` and `LSi`, the serial communication originating from the PC on `RX_SR` and the clock signal used to synchronize minicircuits during measurements on `sampleClock`.

The clock input port `c3_sys_clk` has a peculiarity, it needs a clock line that is trivially routed on the top level and not fanned out. This is effected by an internal clock buffer that prohibits external routing configurations. In turn, the component outputs a 100 MHz clock that can be used throughout the design, on port `CLKPLL100`. Other outputs from the component are `CTS`, which is high when the laser is currently selected as serial target and can be used to inhibit the beeper logic, and a `TX` signal which is the serial communication to be sent to the PC. The remaining outputs from

the component concern signals to be routed to the LPDDR SDRAM chip and originate from a wrapped `s6_lpddr` memory controller block (`mcb`) generated using the Xilinx Memory Interface Generator (MIG). The interior workings of `LaserInterfaceWithCarli` are intricate enough to warrant a more thorough discussion. A RTL schematic of the interior, with discrete logic absorbed into blocks, is shown in [figure 4.6](#).

The interface to the PC is represented by **SerialControl**, which decodes the incoming serial communication and feeds the result to other components. A data recording process goes along the following signal chain. For each IDS 3010 channel there is one **ReceivingEngine** that deserializes the incoming HSSL data and presents it on a 48 wide bus.

These three buses are fed to channel selection logic bunched into the block **chSelLogic**, which is essentially a multiplexer controlled by commands incoming via the serial connection. The selected data line is forwarded to a **RecordingControl** block, which reacts to measurement scheduling requests and stores the requested recording offset. The `sampleClock` signal allows **RecordingControl** to detect when a measurement starts, upon which it will skip the programmed number of initial samples, then has the memory controller **carli64** (controller amenity for `rbcomb lpddr` interfacing) write the specified number of data packets to RAM.

Each data packet is 64 bits in length, with the following format. The bits (39 : 0) are distance data, bits (53 : 40) are the number of 100 MHz ticks since either the last sample or, for the first sample, the scheduled measurement start. The bits (63 : 54) are padding set to zero. With 14 bits of timing information, sample delays of up to 163 μ s can be represented. This isn't only sufficient for the maximum ADS 3010 sample rate of 450 kSPS, there is ample headroom for sensible reductions in sampling rate in favour of increased measurement length. The high byte of laser data is dropped, as it codes for distances exceeding a meter, while a typical experiment deals with variations on the micrometer scale. This optimization was introduced in an earlier version of the laser interface, when BRAM was employed and space was severely limited, and was kept for API compatibility. It has no effect in the current design version.

For readback operations, the signal path is as follows. The **ReadControl** logic detects the request, and instructs **carli64** to transfer all non-trivial memory content. Each 64-bit data packet it receives is then split




Figure 4.6: Register-transfer level schematic of the LaserInterfaceWithCarli component. The memory controller is represented by the carli64 block.

into distance and temporal data, and sent to the PC via the **SerialSend** block as two successive serial packets. There is some headroom for optimizations here, especially pertaining packet size adjustments, which could reduce the amount of data transferred by up to 45%. While this transfer protocol was designed with short measurement traces in mind, where the data transfer time was negligible, the long traces facilitated by SDRAM make serial transfer the most significant contributor to total experiment time, such that this optimization is now worthwhile.

Communication with SDRAM is facilitated by the memory controller **carli64**. To the outside gateway, it presents an interface similar to what's typically seen in SRAM devices. This was a design target of the component, as it was to be as close to a drop in replacement for the previously used BRAM IP block as possible. Internally, it contains a Xilinx MIG generated **s6_lpddr** mcb IP core with custom reconfigured internal PLL for additional clock generation. Using the **s6_lpddr** doesn't only hide refreshing, precharging and other DRAM specific complexities, it also has the advantage that dedicated hardware on the Spartan 6 fabric is instantiated. While the difficulty of using DRAM is then reduced, the interface of a **s6_lpddr** IP core is very different from that presented by a BRAM IP core. To translate between the given SDRAM interface provided by the **s6_lpddr** mcb and the desired SRAM inspired interface, **carli64** leverages a purpose designed **SynchronousLaserDataDynamicDispatcher64** (SLD3) block.

The **SLD3** component is made for exactly the application of recording measurements, and subsequently reading that data back in full. Its state transition diagram is shown in [figure 4.7](#).

Upon power up, the component starts in the **INIT** state. Once the **s6_lpddr** communicates calibration has finished, it transitions to the **IDLE_S** state, where it waits for user input. When a low to high transition is detected on the WE port and not all data has been written yet, as defined by the input latched from the MAXADDR_IN port when memory was last empty, the state enters the write loop by transitioning to **WS_0**. In this state, the input data to be written is latched. In states **WS_1** thru **WS_3**, the lower 4 bytes of the input data is loaded into the **s6_lpddr** write queue, and in states **WS_4** thru **WS_6** the process is repeated for the high bytes. In states **WS_7** thru **WS_9**, a write command with correct address and burst length is loaded into the **s6_lpddr** command queue, and after a




Figure 4.7: Finite state machine transition diagram for the SLD3 interface translator block.

subsequent buffer state **WS_10** the device returns to idle.

The reading cycle functions slightly differently, and only allows extraction of all data at once. Starting from idle, the component enters the state **RS_0** when the RE port is detected to transition high, while the memory is not devoid of contents. In this state, the DR output is pulled low and the device waits for a read strobe pulse to be applied. Once a low to high transition is detected on the port RS, it transitions to **RS_1**, and from this state thru state **RS_3** loads a read command with correct address and burst length to the `s6_ipddr` command queue. In states **RS_4** and **RS_5** it catches the two returning 32-bit data packets and successively updates the corresponding bits on the DOUT output. The component then transitions via a buffer state to **RS_7**, where DR is set. From **RS_7** there are two possibilities. If there is still data to be read, the next state is **RS_0**. If all available data has been read, the device instead returns to idle.

The finite state machine implemented by **SLD3** was designed to facilitate swapping limited BRAM for abundant SDRAM. The resulting write throughput has been measured to be 44MB/s, and while the performance of the read side has not been measured, it is expected to exceed that of writing operations. Even though the designed finite state machine likely doesn't operate the LPDDR SDRAM chip at peak performance, care was taken that the system does not bottleneck on the DRAM read speed, and that the write speed exceeds the sample rate of the IDS 3010. If bridge-PC transfer speeds are to be increased by more than tenfold, to a bitrate exceeding 50 MBaud, the bottleneck might shift to DRAM performance. The system should then be redesigned to contain a BRAM buffer that facilitates full DRAM speed through long r/w bursts.

4.6 Minicircuit

The minicircuit represents a spoke of the FPGA network. Its purpose is controlling 72 DACs in such a way that their 576 associated analog channels carry preprogrammed voltage sequences in lockstep. A minicircuit receives single ended serial communication via an optical link from the bridge. Experiments are typically performed in a two step fashion. First, commands are sent to the bridge to setup the configuration, define settings of the DAC chips and upload data to define the voltage sequences. Once setup is complete, the experiment is then run by the bridge sending periodic triggers to each minicircuit in parallel to signal it's time to generate the next sample on the DACs.

There are ten identical minicircuits present in the experimental setup, for a total of 5760 analog channels. The prevalent challenges present in such distributed systems, namely synchronization and grounding, have already been addressed in earlier sections. Here the focus is on the buildup of individual minicircuit devices. In the next section, the hardware buildup of a minicircuit is discussed. Thereafter, the gateway flashed onto the controlling FPGA will be introduced.

4.6.1 Hardware

The minicircuit is built on a 12 layer PCB shown in [figure 4.8](#). It was designed in Kicad and manufactured by 7pcb. On the left hand side, a Nu-




Figure 4.8: Top view of a minicircuit board with important devices highlighted. Digital control is facilitated by the Numato Saturn expansion card on the left, analog signals are generated in the center by an array of DACs, and output at the header on the right side.

Numato Saturn is plugged into a header array. This is the same device that has already been introduced in an earlier section as brain of the bridge. As a minor difference, the Numato Saturn used on the minicircuit is configured to be powered via the dedicated barrel jack, the mini USB port is to be left unconnected. In order to reduce noise as much as possible, a linear powersupply is used to power the minicircuits. Next to the minicircuit, there is an array of four jumpers (the ID header), which are used to assign a four bit identifier to a minicircuit. A placed jumper represents a cleared bit, and positions closer to the PCB edge are of lower significance. For example, the minicircuit shown in [figure 4.8](#) has ID 3. As was introduced in an earlier section, API calls can address individual minicircuits by their ID.

Next to the identifier header is the TORX1355 toslink receiver (TRx) for reception of serial communication from the bridge sent via an optical link. There is no infrastructure to send data back to the bridge, the channel is inherently single ended.

Further to the right, in the center of the PCB, there are 72 AD5676 DACs arranged in a 9×8 lattice. The AD5676 [96] is an octal, 16-bit, buffered voltage output string DAC with programmable gain and a 50 MHz

SPI interface, operating in the range 0–5 V at sample rates up to 140 kSPS, and able to source up to 15 mA on its output pins. Its frame synchronization pin (SYNC) is used as chip select, and has special importance for daisy chained configurations. After a falling edge on the SYNC pin, data on the SDI pin is read in on the following 24 falling edges of SCLK, and the input is stored upon a successive rising edge on SYNC. Updating the voltages output on the eight channels is done by first loading new voltage values to the input registers of all channels via SPI, then pulsing the dedicated Load DAC (LDAC) pin low to push the input registers to the DAC registers. In this way, all channels are updated simultaneously.

Most communication lines are shared between several DAC chips. Only the Serial Data Input (SDI) net is individually connected to a unique FPGA pin for each of the 72 AD5676 devices. New data is loaded to all DAC chips in parallel, because their SPI clocks are shared, and all 576 input registers are pushed to the corresponding DAC registers simultaneously, because their Load DAC nets are shared. This level of complexity suffices to achieve all targeted functionality, as the data sent to each device is individual albeit being transferred at the same time for all devices.

On the far right of [figure 4.8](#) is the output header. Each of the 8 outputs of each of the 72 AD5676 devices is routed to one pin on this header. While individual header pins are easy to connect to and disconnect from, connecting 5760 such pins individually in a predefined pattern is an error prone task. The error potential can be removed by instead plugging the pin array into a socket, but for the given number of pins the required forces are dangerously high. Soldering minicircuits directly to breakoutboards is not a viable solution, because a swap from one sample to the next implies a swap of breakoutboards, and desoldering a large pin grid is often destructive. The solution applied here is soldering the output pin grid to an adapter PCB, which routes the signals from the given header to a Samtec zero insertion force (zif) connector SEAF-50-05.0-S-10-2-A-K-TR, which then plugs into the breakout board.

As was laid out in previous chapters, a situation might arise where higher output voltages are required. In this case, an amplifier board could be developed, which is then plugged between the minicircuit and the breakoutboard.

With the hardware laid out, the gateway flashed onto the Spartan 6 FPGA controlling the minicircuit will be discussed in the next section.




Figure 4.9: Top level design of minicircuit gateway.

4.6.2 Gateway

The top level gateway design of the minicircuit FPGA is shown in [figure 4.9](#). Inputs to the FPGA are the 4-bit boardID originating from the IDH header, a 100 MHz clock on clk, and the serial communication incoming from the bridge on rxn. The outputs are the DAC communication lines dac_data (connecting to SDI pins), dac_sync (connecting to SYNC pins), dac_clk (connecting to SCLK pins), and dac_ldac (connecting to LDAC pins), and a SerialEcho used for debugging. By design of the minicircuit PCB, all SYNC, SCLK and LDAC nets are shared between 4 AD5676 devices. However, in the FPGA gateway all of these nets are tied together to just three distinct nets, as no more finegrained communication is required.

For output generation, there are five signal generators available on each minicircuit: 4 sine generators and 1 interpolating wavefunction generator. All of these generators calculate their output from a Taylor series. Each of the 576 outputs is then a linear combination of the 4 sine generators plus the interpolated wavefunction. While each sine generator produces a global signal depending only on a programmable frequency and phase, and only the real linear combination coefficients can be programmed per channel, the interpolating wavefunction generator's Taylor

coefficients can be setup individually for each channel. More precisely, for each channel two target functions can be defined. Each target function's domain is then split into 16 intervals, and on each of these intervals the function is expanded to 5th order. As was explained in a previous section, quality-of-life python functions are available to facilitate this translation process. The resulting Taylor coefficients are then uploaded to the interpolating wavefunction generator for each channel, and the two target functions per channel can both be generated forwards and backwards, at configurable speed.

Global multipliers can be applied to all outputs, which enables ramping up/down functionality to avoid hazards. It is also possible to schedule the selective enabling/disabling of signal sources during experiment sequences, for example to separate a pumping stage from a coupling modulation stage.

As signal generation for all channels has to be done in a well synchronized way within a tight temporal budget, a brief discussion of the overall design is in order. Updating all DACs to new data is done on a per-channel basis. First, the new values for all 72 channel 0s is calculated, and then uploaded to the input registers of the corresponding devices in parallel via the SPI interface. Then this process is repeated for all other channels, until all 8 input registers of all AD5676 chips contain new data. Finally, the DAC outputs are updated simultaneously by a pulse on the LDAC net.

In order to calculate 576 voltage values from high order polynomials at 140 kSPS on a 100 MHz clock and transfer them at 50 MHz, a pipelined architecture is necessary. This pipeline has two tracks. One track consists of 5 parallel concatenations of digital signal processing (DSP) slices, that calculate the outputs of the 5 signal generators. Care was taken that these five pipelines have the same delay, such that corresponding results are produced in lockstep. The results of these 5 pipelines are then linearly superposed as programmed.

The other track is used to keep track of data propagation. Whenever the generation for the next channel starts, a bubble pulse signal is generated at the beginning of the pipeline. This bubble rises through a pipeline of equal length as the computation pipeline, such that the arrival of the bubble at various stages signals the arrival of the first datum for a new channel, and 575 clock cycles later the last datum for this channel will arrive. Similarly, when data for the last channel 7 is generated, an LDAC

bubble is sent into the pipeline. Reception of this LDAC bubble at the end of the pipeline signals the DAC input registers are fully loaded and an LDAC pulse is in order.

The above design results in a pipeline with a delay of 25 clockcycles, meaning after an initial waiting period of 250 ns, new data is output every 10 ns. All things considered, the presented design can fully update all DAC outputs at a frequency of up to 160 kHz, in excess of the 140 kHz AD5676 sample rate. With the overall architecture laid out, in the following sections the described pipeline will be traversed from input to output, and the individual blocks shown in [figure 4.9](#) will be discussed in more detail.

SerialController

This block deserializes and decodes relevant incoming serial traffic, selected using the 4-bit board identifier `BoardID`, and forwards the results to other components. It uses 5 registers to communicate input: A 6-bit `Command` register, to which opcodes are written, a 18-bit `Data` register, into which incoming data is loaded before it is either assigned to other registers or directly fetched by other components, a 15-bit `ADDR_BRAID` register, which stores memory addresses for interpolating wavefunction generator Taylor coefficient writes, a 10-bit `ADDR_SINE` register, which stores memory addresses for sine generator amplitude writes, and a 32-bit `NX_SINGEN` register, which is used for general purpose data transfer.

Incoming serial data packets are 1 Byte in length, with the 6 LSBs being the payload, and the 2 MSBs labeling how the payload is to be interpreted — as opcode, or as lower, middle or upper 6 bits of the `Data` register. In a typical application, the controlling entity would first bring the registers to a suitable state, for example load the desired target generator frequency into the `NX_SINGEN` register, and then issue a command that instructs a generator to adjust its frequency.

A notable opcode to mention is `0x3F`, which is the synchronization tick sent by the bridge and triggers the generation of the next set of voltage values. This command is interpreted by the **ClockReceiver**, which is considered next.

ClockReceiver

This component serves the single purpose of listening for a synchronization tick from the bridge, and starting the pipeline upon detection. More precisely, when an opcode $0x3F$ is detected, the component will pulse its `PipelineStartSignal` port high for one clock cycle. This pulse is registered by the **Datapath** component, which performs most of the heavy lifting for data generation, and is considered next.

Datapath

The **Datapath** component performs the calculations that yield the digital voltage values for each channel. Its internal design has some structure to it and is shown in [figure 4.10](#).

An incoming `PIPELINE_START` pulse is detected by the **OUTPUT_ENUMERATOR**. It then sends a bubble pulse out its `BUBBLE_OUT` port, and starts generating data targets on its 10-bit `OUTPUT_NUM` port, where the high 7 bits enumerate the AD5676 device, and the low 3 bits encode the output channel. Generation is done at a rate of one target per clock cycle, and starts with channel 0 and device 0, then channel 0 and device 1, and so on until channel 0 on device 71. At this point the channel is incremented, the device reset to 0, and another bubble is output. When channel 7 on device 0 is reached, an additional LDAC BUBBLE is released into the pipeline via the `LDAC_BUBBLE_OUT` port. To signal the start of a generation sequence, a pulse is output at the `START_SEQUENCE` port when a `PIPELINE_START` signal is received.

The `PIPELINE_START` signal is also detected by the **SelectiveEnabling** unit. This component listens on the serial bus for commands concerning generator state setup and scheduling, and outputs corresponding signals to the generators on its `Enable` bus. When an experiment is running, it counts the number of samples generated, and according to the configured schedule switches the generator states.

The target device and channel enumeration is received by **BraidInterpolator** and the four **SineGenerator** instances. The **BraidInterpolator** calculates the output of the interpolating wavefunction generator for the given device and channel. Input to its `PIPELINE_START` port is used to update the phase at which the wavefunctions are to be evaluated. **BraidInterpolator** listens on the serial bus for commands pertaining any

properties of the sequence playback and the writing of Taylor coefficients, which it stores in BRAM. The calculation of the output data is carried out in a pipeline of 5 LogiCORE IP DSP48 Macros [97], each of which introduce a delay of 3 clock cycles to perform a $(A + D) \times B$ operation. At each stage of the pipeline the correct coefficients are fetched from BRAM according to the corresponding input on `OutputNUM`. To synchronize the calculation, all computationally relevant values, such as the phase and Taylor coefficient memory addresses, propagate in pipelines parallel to the DSP48s, such that at any stage the correct values are trivially obtained. The result of the calculation is output at the `DATA` port.

The **SineGenerator** component works very similarly to the **BraidInterpolator** and calculates the output of a sine generator. Internally it has a pipeline of 6 LogiCORE IP DSP48 Macros, identical to those used in the **BraidInterpolator**. This device is slightly less complex, as the Taylor coefficients are hardcoded, and parameters are limited to a global frequency, a global phase offset, and an amplitude defined per channel. Just like in the **BraidInterpolator** all computationally relevant values are passed through parallel pipelines of equal length, such that the calculations are automatically synchronized. The result of the calculation is output at the `DATA` port. When a `PIPELINE_START` signal comes in, the phase and, if sweep is enabled, frequency are updated. On the serial bus, it listens for commands involving changes to frequency, phase, sweep configuration, and per channel amplitudes, which it stores in BRAM.

The pipelines of the **BraidInterpolator** and the 4 **SineGenerator** components each introduce a delay of 19 clock cycles. The output of the 5 generators is then automatically synchronized, and fed into an **AddScale** device. This component sums these 5 outputs, multiplies the result with a factor ≤ 1 to implement ramping functionality, and presents the result at its `Output` port. To perform ramping, it listens on the serial bus for change requests to the ramping setup.

While the calculated data passes through the computation pipeline comprising the 5 generators and the **AddScale** component, the bubble, the LDAC bubble, and the channel number pass through a delay line called **PipelineDelay**. This component introduces the same delay as the whole computational pipeline, such that the output bubble, LDAC bubble and channel number always correspond to the concurrent output of the **AddScale** component. The computational output from the **AddScale**

component is then extended to 24-bit by concatenating it with an AD5676 write to input register command and the target channel number. The resulting data packet can be directly sent to the target AD5676 device, and is put on the **Datapath** DATA output port, along with the correctly synchronized bubble on BUBBLE_0 and LDAC bubble on LDAC_BUBBLE_0.

The output data of the **Datapath** can not yet be transferred to the AD5676 devices, there is one more pipeline stage it must pass through first.

DirectInjection

In normal operation, the **DirectInjection** component merely performs a passthrough, connecting the inputs STREAM_IN, BUBBLE_IN, LDAC_BUBBLE_IN to the outputs STREAM_OUT, BUBBLE_OUT, LDAC_BUBBLE_OUT. It only performs nontrivial acts when the serial bus exhibits a request for direct data transfer to the AD5676 devices, such as a gain change request. When such a command is received, the corresponding 24-bit string is inserted into the pipeline and handed on to the next stage, the **TRANSFER_ENGINE**, for transferring.

TRANSFER_ENGINE

The **TRANSFER_ENGINE** performs the final data transfer to the AD5676 devices. The strategy is to wait for a full data frame (one data packet for each AD5676) to arrive, and then clock out all of that data in parallel. Its internal architecture is shown in [figure 4.11](#), and consists of 72 slave devices (one slave per AD5676), and one master. The master controls the timing on the SPI communication via the SCLK, SYNC and LDAC signals, while each slave controls the SDI line to one AD5676 device.

An incoming bubble enters at the last slave, on the next clock cycle that slave forwards the bubble to the next slave, and so on until the first slave hands the bubble to the master. The incoming data packet is input to all slaves. Whenever a slave holds the bubble, it will latch the current data packet into its input register. Consequently, when the bubble finally arrives at the master, every slave will have latched data, and the master signals the slaves to copy their input registers to their output registers via the TRANSFER signal. Next the master starts cycling the SCLK at 50 MHz, and the slaves shift data out their DATA_OUT port whenever they detect




Figure 4.11: Register-transfer level schematic of the Transfer Engine block.

SCLK low. Note that this only works because the SPI clock is exactly half the frequency of the FPGA clock. If the master received an LDAC bubble, it pulses the LDAC net low after the data transfer is complete, such that the DAC outputs are updated.

It is crucial that the slaves have two registers, one for input and one for transferring. While data is being clocked out to the AD5676 devices, the next data packets might already be coming in and a new bubble propagating up the slave chain. In a single register configuration, this would lead to data corruption. The SPI transfer duration is roughly 50 clock cycles, while the next bubble reaches the master only after 72 clock cycles, so a double register setup suffices to prevent interference between successive packets.

4.7 Environment control system

The Environment control system has a monitoring, and an active control side. On the monitoring side, parameters of the system are continuously



Figure 4.12: Screenshot of the Grafana dashboard visualizing parameters logged in the InfluxDB database.

recorded and logged to an InfluxDB database. A Grafana dashboard is used to visualize this data and enable introspection. The monitored quantities are pressure, temperature, and the power drawn by specific devices. While only one pressure is recorded, namely the pressure inside the vacuum chamber, there are several temperature sensors spread throughout the setup.

On the active control side, a PID feedback loop is deployed to stabilize the temperature recorded by one specific temperature sensor placed within the vacuum chamber. As was mentioned in [chapter 2](#), the sample is sensitive to fluctuations in temperature, and this stabilization mechanism is used to inhibit adverse influences.

In the following sections, the components of the environmental control system will be outlined in more detail. First, the dashboard will be explained, before the discussion moves on to pressure monitoring. Finally, the temperature measuring and stabilization system will be discussed.

4.7.1 Grafana

A screenshot of the Grafana dashboard is shown in [figure 4.12](#). The bottom row of panels visualizes the performance of the temperature stabi-

lization loop. The central panel shows time series data of 5 temperature sensors placed at various locations in the vicinity of the vacuum chamber. Two time series in the top right corner visualize data related to heating and minicircuit power supply units, above them are dials showing instantaneous values. On the top left are air temperature measured by the ADS3010 and the chamber pressure.

In the following sections the devices recording most of this data will be introduced.

4.7.2 Pressure

The vacuum pressure is measured using a Pfeiffer TPR 280 pirani gauge connected directly to the vacuum chamber. Initially, a cold cathode sensor was employed, but after that sensor had combusted contaminants present in the chamber and splattered the resulting products onto a sample, it was exchanged for a plasmaless device. While a pirani gauge's admissible range is reduced compared to cold cathode devices, the lower threshold of 500 nbar is still well below the typical equilibrium pressure at roughly $1 \mu\text{bar}$. The measurement result is reported to the PC several times per minute via a serial connection.

In future developments, this sensor could potentially be integrated into a PID loop to achieve active pressure control, enabling dynamic Q factor tuning. This would require the addition of an actuated valve, a suitable device could be the Pfeiffer EVR 116.

4.7.3 Temperature

The temperature control system consists of several distinct subsystems. On the one hand, there are sensors used for pure monitoring and logging purposes, such as the sensor present in the ADS3010's ECU and several sensors placed throughout the setup to track temperatures at various points. On the other hand, there is a temperature stabilization unit (TSU) keeping the temperature of the sample in a narrow window. As was mentioned in [chapter 2](#), the sample's properties are susceptible to changes in temperature, such that those fluctuations should be kept as small as possible. The goal of the TSU is to lock the temperature of the sample in a region of width $\pm 5 \text{ mK}$. This roughly constrains drum frequencies to a 1 Q window.

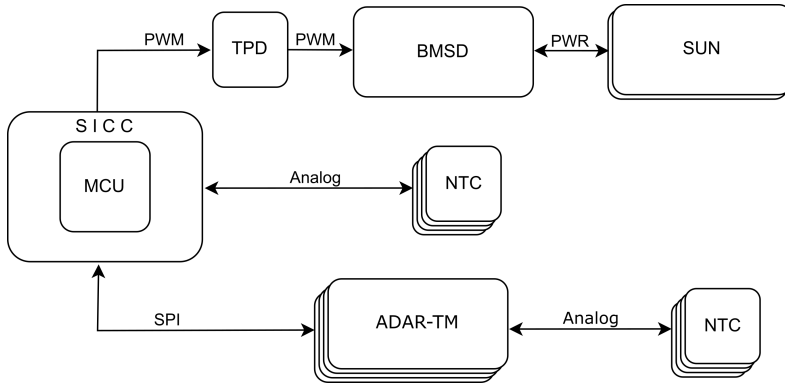


Figure 4.13: Connectivity diagram of the temperature stabilization unit. The individual components are further discussed in the text.

The TSU is constructed from a set of modules, the buildup of the resulting system is shown schematically in [figure 4.13](#). The system is controlled by an Atmel SAM3X8E ARM Cortex-M3 [98] microcontroller unit (MCU), as broken out on the Arduino Due [99] microcontroller board, which is connected to the PC via UART. Temperature measurements are performed by negative temperature coefficient thermistors (NTCs) driven by custom developed ADAR-TM modules. The MCU controls ADAR-TM modules via Serial Peripheral Interface (SPI), and modulates the current flowing through heating filaments (SUN) attached to the vacuum chamber via a voltage-controlled pulse width modulator (PWM) connected to a buck-mode step down converter (BMSD) through a totem pole MOSFET driver (TPD). Interconnection and PWM generation is enabled by a custom developed S I C C board. To achieve stabilization, the MCU then runs a PID feedback loop: Read the temperature reported by the ADAR-TM, and adjust the heating current accordingly. There is no active cooling capability present in the system. Instead, the idea is to stabilize the temperature at a value slightly higher than room temperature, namely 31°C , such that reasonably potent passive cooling can be achieved by simply not heating. The absolute value of the temperature is largely irrelevant, only stability matters.

In the following sections, the individual components making up the




Figure 4.14: Photo of a S I C C board, with functional components annotated.

system will be described in more detail. The exposition starts with the interconnect board and devices on the measurement side of the system (lower branch of [figure 4.13](#)). It then moves on to modules enabling the modulated heating of the vacuum chamber (the upper branch in [figure 4.13](#)), and concludes with a discussion of the MCU firmware.

S I C C (SPI Interconnect and PWM Circuit)

The S I C C board plugs directly into the microcontroller board, and all downstream communication originating from the MCU are routed through the S I C C. An annotated picture of the S I C C module is shown in [figure 4.14](#). It roughly provides three different functionalities. Firstly, it fans out the MCU's SPI bus to four terminals SPI0 thru SPI3 for connection to ADAR-TM modules. This facilitates the simultaneous monitoring of up to 16 high precision temperature measurement channels.

Secondly, in the PWM region, it conditions the MCU's 12-bit DAC output and routes it to control the duty cycle of an LTC6992 [100] voltage-controlled pulse width modulator chip configured to run at 50 kHz. The PWM output of the LTC6992 is exposed for further connection to the BMSD board for heating power control. The S I C C also includes the pos-

sibility of a dedicated linear 5.5 V LDO (low drop-out regulator) to extend the PWM operating rails, which can be relevant for downstream driving stages.

Thirdly, it includes four analog input stages (AIS), suitable for the direct connection of thermistors. These stages contain precision reference resistors (20 k Ω , 0.05 %, 10 ppm), and the input voltages are routed through signal conditioning circuits (active LM358DT low-pass filters) and then fed to the MCU's internal 12-bit ADC for low precision temperature monitoring. In total, 16 high precision (24-bit) and 4 low precision (12-bit) temperature measurement channels are available per S I C C device, along with one PWM output.

ADAR TM (Analog to Digital Resistive Temperature Measurement)

An ADAR TM module is essentially a PCB containing an analog to digital converter (ADC) along with its required supporting infrastructure. It is shown in [figure 4.15a](#). The core of this module is an Analog Devices AD7124-4 [101] (denoted AFE). The AD7124-4 is a completely integrated SPI controlled analog frontend (AFE) optimized for DC operation. It contains a low-noise 24-bit Σ - Δ ADC with four differential inputs with programmable gain arrays, buffers and sequencer. Programmable digital filters allow for suitable noise rejection. Further integrated infrastructure include a clock and a buffered high precision bandgap reference that can be buffered and made available at a pin of the device. All in all, this chip was selected as the best fit amongst 11 AFE/ADC devices that were compared for the application at hand.

As ADCs always bridge the analog and digital domains, the hosting PCB necessarily is a mixed-signal design and care must be taken to prevent digital noise from affecting sensitive analog nets. The principle applied here is using a continuous ground plane without slits or splits, and carefully placing digital and analog components as well separated as possible to prevent interference. Identifying and shaping return current paths (which depend on frequency) is much more important and effective than cutting the ground plane, which in many cases deteriorates noise performance. An in-depth discussion of mixed-signal design goes beyond the scope of this work, further information can be found in various application notes and blogs [102–105]. In this spirit, all components that belong to




Figure 4.15: Photos of two components used in the temperature stabilization unit. a: ADAR-TM, b: BMSD.

the digital domain are placed in the top left corner of the PCB, as shown in [figure 4.15](#). This includes the SPI header for connection to a S I C C board and an LDO to generate the digital supply. A corresponding LDO for the analog side is placed on the right hand side of the PCB.

The PCB contains four differential input stages for NTC thermistor connection (for example IS1 in [figure 4.15a](#)), with voltage supplied by the AD7124 internal reference for ratiometric measurements, high precision reference resistors and suitable differential anti aliasing filters. This measurement circuit is shown in [figure 4.16a](#). Each connection channel consists of three pins. One pin is driven by the AD7124 internal reference V_{ref} , and is to be connected to lead A of the NTC thermistor for excitation. Another pin is the positive input and is also to be connected to thermistor lead A. the third pin is to be connected to thermistor lead B and corresponds to the negative input. This input is divided by a precision reference resistor ($20\text{ k}\Omega$), with the division result routed towards the negative ADC input. Positive and negative inputs are routed through a differential anti aliasing filter with a corner frequency of 800 Hz [[106](#), [107](#)]², as required by the ADC architecture, and their differential voltage V_{diff} can be used to

²While the AD7124 datasheet does not provide details concerning the ADC modulator, the external AAF design can be found in the cited user guide and reference design.




Figure 4.16: Schematics of two circuits used in the temperature stabilization unit. a: temperature measurement circuit from NTC to ADC input as used in ADAR-TM modules. b: Totem pole MOSFET driver used for BMSD charging and discharging.

calculate the thermistor resistance via

$$R_{th} = \frac{V_{diff}}{V_{ref} - V_{diff}} \cdot 20 \text{ k}\Omega. \quad (4.2)$$

The corresponding temperature is then readily found from a temperature-resistance lookup table. The thermistor used in this application is a Measurement Specialties 44008RC precision epoxy NTC thermistor, which has a thermally conductive coating and low mass allowing efficient thermal equilibration.

On the designed PCB, the AFE chip performs to specification. In operation, at a sample rate of 9.4 SPS with 1 m long twisted triplet probe cables, RMS noise has been measured to be in the range 10-100 μ K. With the devices on the measurement side of the TSU discussed, in the following sections the focus will shift to the modules that make up the heating side.

BMSD (Buck Mode Step Down Voltage Converter)

As has been discussed in preceding sections, the MCU controls the duty cycle of an LTC6992's PWM output to modulate the heating current. This PWM signal is routed to the BMSD, and it is here where it's transformed into a corresponding heating current.

The BMSD PCB is shown in [figure 4.15b](#). In the bottom right corner, the PWM signal enters the board, and one can see the cut off TPD from which this signal originates. Just above the PWM input, a 24 V, 75 W power supply is connected, the 10 μ F tantalum decoupling capacitor was

added as an afterthought when strong switching noise was detected in other devices. The output terminals, located in the top right corner, are connected up such that their current passes through the heating elements (SUN). The Schottky diode (D), MOSFET (Q), capacitor (C) and inductivity (L) are routed to form a buck mode down converter, with the PWM signal driving the MOSFET gate, such that the heating current is modulated according to the PWM duty cycle.

It should be noted that the MOSFET in this circuit represents a lowside driver with its source terminal connected to ground, meaning one must make sure the heating circuit is well isolated. If the heating current finds ground anywhere else than through the MOSFET, control is lost and the heaters will continuously run at full power.

Next, the utility and circuit of the totem pole driving circuit will be discussed.

TPD (Bipolar Totem Pole MOSFET Driver Stage)

This module is necessary to overcome the inadequate drive strength provided by the pulse width modulator. As it turns out, the IRF1405 MOSFET used on the BMSD module has an input capacitance of 5.5 nF, which is too much to be driven directly by the LTC6992 at 50 kHz. Instead, the totem pole driving circuit shown in [figure 4.16b](#) is used. The shown circuit allows for charging and discharging currents to be directly sourced from and sunk to the rails via complementary bipolar junction transistors switched by the input PWM signal. While this yields an output signal with critically increased slew rates, there is an additional $V_{CE} \approx 0.5 V$ voltage drop to be dealt with. This is addressed by adjusting the totem pole rails, and increasing the PWM amplitude by using the auxiliary 5.5 V supply on the S I C C. More information pertaining the driving of MOSFET gates can be found in [\[108\]](#).

Using a dedicated MOSFET driver, the heating current can adequately be controlled using the BMSD. In the next section, the custom heating elements, through which the modulated heating current is routed, will be discussed.

SUN (Strongly uniform Netheater) Boards

The SUN boards are heating elements used to increase the temperature of the vacuum chamber. They are single side PCBs with aluminium back-plane for efficient thermal conduction. PCBs are well suited for custom heater design implementation, as they can be manufactured cheaply in arbitrarily designed shapes. Care must be taken however, that the involved temperatures do not exceed material tolerances.

The electrical design of a heater is rather straightforward. A continuous wide track is routed to cover the whole surface area of the PCB, with care taken to minimize the area enclosed by the current to keep induced noise to a minimum. Wide tracks are important because low impedance implies high heating power. An applied voltage U gives rise to a current $I = U/R$, such that higher heating power $P = U \cdot I = U^2/R$ is achieved with lower trace resistance R . As safety measure, each SUN PCB contains an attachment point for a thermal fuse that serves as final safeguard in case of runaway heating. In operation, two SUN boards are attached directly to the vacuum chamber and connected up in series, with input and output provided by the BMSD module. This setup can be seen in [figure 4.2](#). The resistance added to the heating net is on the order of a few Ω per SUN board, and the maximal achievable heating power at 24 V is around 60 W.

With all hardware components discussed, the next section will highlight some key points of the firmware running on the MCU.

MCU Firmware

As mentioned before, the MCU is the brain of the TSU, and runs a PID loop to stabilize the sample temperature. Setup and control of the AFEs is facilitated by a custom developed SPI driver for the AD7124. Measurements from a specific channel are used as process variable, and a suitably tuned PID expression generates a corresponding control voltage applied to the S I C C LTC6992 to modulate the heating current. The sensor providing the process variable is located inside a screw hole of the vacuum chamber.

Tuning of the proportional, integral and derivative coefficients was attempted with various methods including Cohen-Coon [109] and Ziegler-Nichols [110]. While variations of the Ziegler-Nichols method led to the best results in terms of equilibration and stability, none of the applied automatic methods was able to provide satisfactory results. Their main utility

was to obtain initial coefficient guesses. Optimal performance was only obtained when the parameters were further tuned manually.

Several different firmwares have been developed for the TSU MCU, and two are currently deployed on different devices. Firmware `temperature_monitor` is used to operate the TSU in a pure monitoring function. In terms of [figure 4.13](#), the whole top branch is discarded, and the MCU only responds to UART requests pertaining the measured temperatures.

Firmware `pid_controller` on the other hand controls the full diagram shown in [figure 4.13](#) and is deployed to stabilize the setup. The communication interface enables the user to read and write all PID parameters dynamically, and request currently measured temperatures. While this device could in principle operate autonomously, it is programmed to shut down if no communication via UART is recorded over a specific time interval. Only supervised operation is enabled, with the assumption being that a device reading the measured temperatures will detect dangerous states and react accordingly.

In the current setup, a PC is scheduled to periodically read the measured temperatures on all connected TSUs, and log them in the InfluxDB. For the TSU controlling the setup temperature, it also compares measured temperatures to set thresholds, informs users of anomalous states via text messages, and opens the heating current circuit via two distinct switches when dangerous temperatures are recorded.

This concludes the discussion of the custom developed experimental setup. While it consists of a lot of components, it still doesn't provide all functionality one could ever wish for. However, it is constructed in a modular fashion and is extensible. Ideas for potential future developments will be provided in the conclusion [chapter 7](#). In the next chapter, measurements taken with this setup to characterise a fabricated sample will be described.

5 | Measurements

In this chapter, the characterization of a fully bonded double layer sample is presented. Initially, the discussion focusses on properties of individual resonators actuated electrostatically via their central electrodes, before moving on to the coupling electrodes.

A sample called F1F4, manufactured according to the design described in [chapters 2 and 3](#), shall be investigated in this chapter. An impression of F1F4, mounted on its breakoutboard prior to wire bonding, is shown in [figure 3.5](#). While many single layer and partial samples were fabricated and characterized during the design and process development stages, sample F1F4 represents the first complete implementation.

The level of functionality of sample F1F4 shall be explored. Earlier partial samples, lacking the double layer nature, were globally actuated using a single piezo speaker glued onto the wafer surface. While this approach trivially guarantees successful excitation, selectivity is disabled. In contrast, sample F1F4 has all its electrodes paired up, but as discussed in [section 2.3](#) many things need to go right for a useful force to arise. Achieving electrostatic resonator excitation is thus the first goal. If an influence of applied electrode voltage on resonator motion is observed, the question then shifts towards the interaction strength. Depending on this magnitude, different properties of the system can be evaluated and different protocols implemented.

In the next section, the experimental approach for characterization of sample F1F4 is described, and the observed single resonator properties are

presented. The section thereafter deals with the interaction electrodes. Throughout those two sections, experimental details and challenges will be explained when fitting.

5.1 Single drum

In order to perform interferometric measurements on individual resonators, their lattice must first be identified in the coordinates of the stages described in [section 4.2.2](#). The naive approach is using a pilot laser and visually aligning it to the position of the resonators. This method is not only cumbersome manual, but also hard to perform successfully. As the resonators are read out from the backside, no specular reflections are visible from a targeted resonator, and positioning becomes a guessing game. Furthermore, spatial constraints make visual access to aimed-at resonators near impossible.

A more elegant method is found in signal contrast measurements. Scanning the interferometer over a small region of the sample while simultaneously reading out its observed signal contrast makes lattice identification near trivial. While the extended spotsize explained in [section 4.2.1](#) complicates the prediction of a specific contrast measurement, the results should exhibit a periodic pattern when the device scans over a periodic structure. The results of such a contrast measurement are shown in [figure 5.1](#). The resonators can be clearly identified, and scanning a few corners of the sample are enough to identify a lattice seed and calculate precise lattice vectors.

Before moving on to actual results, it should be noted that all resonator measurements involve a constant struggle for signal quality. Observing small metallic structures out of focus and through a curved layer of silicon nitride – namely the resonator itself – using an interferometer that is not placed perfectly perpendicularly to the wafer surface yields signal contrast well below what’s recommended by the device manufacturer. This situation is further exacerbated by a very unfortunately tuned cavity formed by the outer wafer layers [93]. These effects lead to a setup that is highly sensitive to the precise positioning of the interferometer fiber, and careful micro adjustments to the measurement position are necessary for each drum.

With the above complications identified and addressed through provi-




Figure 5.1: Result of an interferometer signal quality scan over the bottom left corner of sample F1F4. The resulting pattern allows the easy identification of the resonators (circular objects).

dent setup operation, central electrodes of resonators on sample F1F4 were excited, and corresponding motional responses of the associated drums were successfully detected in interferometer traces. As an initial characterization, a full frequency response of one resonator was recorded via the following protocol: First, a step response was taken to obtain an initial guess for the resonance frequency. Starting from this rough estimate, the resonance frequency was then zeroed in on via repetition of the following PLL technique. In the first stage, the resonator is driven at the resonance estimate. The ensuing ringdown, during which the resonator oscillates in resonance, is recorded and compared to the drive. From the dephasing between these two signals, a more precise guess for the actual resonance frequency can be calculated. Once a satisfactory resonance fre-




Figure 5.2: Measured resonance of a single resonator with phase shift, extracted via lock-in amplification. Stars show measured values, the solid lines represent simultaneous fits of lorentzian and phase function. The fit parameters are a quality factor of 15992 and a resonance frequency of 26051 Hz.

quency has been found, the resonator is driven to saturation at a number of frequencies surrounding the suspected resonance, and the resulting amplitudes and drive-resonator phases are calculated via offline lock-in amplification. The result of this measurement, along with a simultaneous lorentzian-phase fit, is shown in [figure 5.2](#). Both, the quality factor and the resonance frequency are in the expected range.

The above described measurement doesn't only prove successful resonator manufacturing, but also underlines the achieved temperature stability. The total duration of the experiment is on the scale of several dozen minutes. With the temperature dependence of the resonance frequency found to be around -61 Hz K^{-1} for the given resonators [93], the resonator temperature was likely confined to a region narrower than $\pm 5 \text{ mK}$ throughout the experiment.

While individual resonator properties are important in their own right, an understanding of the distribution of these properties over the lattice is also of interest. To this end, the resonance frequencies and quality factors of all drums on the bottom layer were measured. The results of

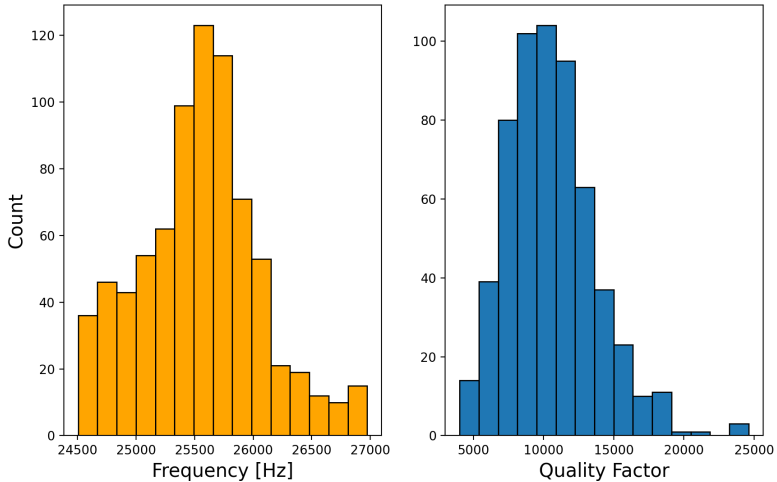


Figure 5.3: Properties of F1F4 bottom layer resonators. Left: histogram of resonance frequencies, determined via step response. Right: histogram of quality factors, determined via automated ringdown fitting.

these measurements are shown in [figure 5.3](#). While measuring the resonance frequencies is relatively straight forwards via a step response, determining quality factors is more involved. While the frequency scanning method applied to obtain [figure 5.2](#) works for individual resonators when everything is well functioning, applying this technique to hundreds of resonators would yield an experiment lasting weeks. The additional danger that many of these experiments could fail for a variety of reasons, leading to several stages of re-runs, solidly places this approach in the infeasible domain.

Instead, the more efficient and robust ring down approach for quality factor extraction is used. Fitting an exponential $\exp -t/\tau$ to the decaying envelope of a ring down process, the quality factor can be calculated via $Q = \pi\tau f_{\text{res}}$ [93], where f_{res} is the resonance frequency. This protocol only needs a single trace to be recorded, and the resulting shorter duration makes it less susceptible to a drifting environment.

While the histogram of quality factors shown in [figure 5.3](#) somewhat resembles a gaussian distribution, this is not the case for the resonance




Figure 5.4: Spatial distribution of resonance frequencies on bottom layer of F1F4. Each coloured disk corresponds to one resonator. Drums that are not represented could not be automatically measured.

frequencies, suggesting there might be more at play than random drum to drum variations. A potential explanation to this curiosity is found when looking at the spatial distribution of resonance frequencies, as shown in [figure 5.4](#). While there are defects to the pattern, nearby drums tend to have similar frequencies. This pattern could be explained by spatial variations in the residual nitride stress, which would then have a strong influence on the resonance frequencies of the stress dominated resonators.

The voids in [figure 5.4](#) don't necessarily correspond to broken resonators. As was described earlier, the success of measurements depends strongly on positioning details. Even though algorithms were developed and applied to automatically detect and correct measurement issues of various kinds, they don't come close to what can be achieved in manual operation. There are two main regions of the sample where voids do closely correspond to sample defects. One is the region on the upper left side, where a large number of drums were destroyed by an unscheduled tweezer impact during fabrication. The other region is the bottom half on the right side, but in this case the damage is not structural. Instead, drums in this region are controlled by electrodes routed via the first few wire bonds that were placed. During wire bonding, the process was continuously being optimized, and those initial wire bonds were suboptimal to a degree where it's just a question of when, not if, they break. Even without electrical connection, the resonance frequency of some drums can still be identified in the spectrum due to environmental excitation.

In order to address the detuning of different resonators, electrostatic softening was outlined in [section 2.3](#). Applying a positive static offset to

a resonator's central electrode should tune its resonance downwards by shifting its equilibrium position to lower the capacitor gap. Unfortunately, this behaviour could not be observed in sample F1F4. More generally, slow off-resonant electrode voltage modulations seemed to have no sizeable effect on the resonator's motion. As drums can reliably be driven close to their resonance frequencies, and the effect of off-resonant voltage changes is strongly dependent on the interwafer gap, the issue is likely caused by this gap being excessively large. During manufacturing a conscious decision was taken to rather err on the large gap side than risk having all drums stick to the opposing wafer. While the latter situation would have rendered all resonators completely inoperable, the former case, which has likely materialized, inhibits static voltage effects.

With single resonator properties discussed, in the next section interactions will be treated.

5.2 Interactions

In this section, interactions between neighbouring drums are the topic of interest. To verify the electrical functioning of coupling electrodes, pumping experiments are performed. More precisely, a resonator is driven to saturation using one of its coupling electrodes, and the achieved amplitude is extracted via lock-in amplification. This process is performed for both drums sharing that coupling electrode, and for various driving amplitudes.

The results of such a coupling electrode excitation experiment are shown in figure [figure 5.5](#). The achieved amplitudes for top and bottom drum are very similar, in accordance with the fact that the coupling capacitor acts equally on the two resonators. Additionally, the saturation amplitude of both resonators in a resonant central electrode drive of amplitude 0.25 V was determined. For the top resonator, this yielded an amplitude of $0.9 \text{ nm} \pm 0.1 \text{ nm}$, while for the bottom resonator a much larger amplitude of $9.8 \text{ nm} \pm 0.2 \text{ nm}$ was achieved. These results are consistent with the metal design. Coupling electrodes form less overlap area than their central counterparts, and consequently they exhibit smaller forces from equal voltages. Furthermore, the central electrode of a resonator on the bottom layer faces a solid layer of gold as its counterpart, while the central electrode on a top drum is paired with a circular pad of smaller




Figure 5.5: Saturation response of resonators to resonant drive on one of their coupling electrodes, with varying peak to peak amplitude. The measured drums are nearest neighbours, and both driven via the same coupling electrode. Blue: Response of the bottom drum to resonant coupling electrode excitation (25.954 kHz). Orange: Response of the top drum to resonant coupling electrode excitation (24.593 kHz).

diameter than its own.

With the coupling electrodes verified, the object of study becomes actual coupling of nearest neighbour drums. In the most straightforward implementation, neighbouring resonators would be tuned to a common resonance frequency, and subsequently coupled via a static voltage applied to their coupling electrode. However, as described in the previous section, tuning is not functional in sample F1F4. Instead, a different approach to coupling two resonators with detuned resonance frequencies must be explored. Considering two neighbours with different resonance frequencies $\omega_1 > \omega_2$, the idea is to modulate the voltage on the coupling electrode with the difference frequency, that is $\propto \cos(\omega_1 - \omega_2)t$. The combination of this modulation with the motion of the driven resonator will then produce frequency components at the resonance frequency of the neighbour.

In the following, this approach will be formalized more concretely. The goal is to derive an effective time-independent equation of motion for a set of nearly degenerate resonators, which are coupled with a detuning compensating drive,

$$\begin{aligned}\ddot{x}(t) &= -\omega_1^2 x(t) - \lambda_1 \dot{x}(t) + 2f \cos[(\omega_1 - \omega_2)t + \varphi] y(t), \\ \ddot{y}(t) &= -\omega_2^2 y(t) - \lambda_2 \dot{y}(t) + 2f \cos[(\omega_1 - \omega_2)t + \varphi] x(t).\end{aligned}\quad (5.1)$$

The separation of scales is assumed to be

$$\omega_1, \omega_2 \gg \omega_1 - \omega_2 \gg \lambda, \sqrt{f}. \quad (5.2)$$

Noting that $\lambda_i = \omega_i/Q_i$, this assumption implies the linewidth of an individual resonator is much smaller than the detuning. Using this, the solutions are written as the ansatz

$$x(t) = e^{i\omega_1 t} X(t), \quad y(t) = e^{i\omega_2 t} Y(t), \quad (5.3)$$

where $X(t), Y(t)$ only contain frequencies much slower than $\omega_{1,2}, \omega_1 - \omega_2$. Inserting this into [equation \(5.1\)](#), one obtains for $X(t)$

$$\underbrace{\ddot{X}(t)}_{\approx 0} = \underbrace{(-2i\omega_1 + \lambda) \dot{X}(t)}_{\approx -2i\omega_1} - \lambda\omega_1 X(t) + 2f \underbrace{[e^{i\varphi} + e^{2it(\omega_2 - \omega_1)}]}_{\approx e^{i\varphi}} Y(t) \quad (5.4)$$

Compared to ω_1 , higher order derivatives of $X(t)$, as well as λ can be neglected. Furthermore, the term $\exp[2it(\omega_2 - \omega_1)]$ is oscillating fast with respect to the time scales of X and Y and can also be dropped. This leaves the final equations for the slow scales

$$\begin{aligned}\dot{X}(t) &= -\frac{\lambda_1}{2} X(t) - i \frac{f}{2\omega_1} e^{i\varphi} Y(t), \\ \dot{Y}(t) &= -\frac{\lambda_2}{2} Y(t) - i \frac{f}{2\omega_2} e^{-i\varphi} X(t),\end{aligned}\quad (5.5)$$

which is the sought after result. A noteworthy difference to unmodulated coupled resonators is the asymmetry in coupling, $f/2\omega_1$ vs $f/2\omega_2$.

Next, the obtained effective equations of motion [\(5.5\)](#) shall be solved. First, the eigenvalues of the above problem are

$$\xi_{\pm} = \frac{\lambda_1 + \lambda_2}{4} \pm \frac{1}{2} \sqrt{\frac{(\lambda_1 - \lambda_2)^2}{4} - \frac{f^2}{\omega_1 \omega_2}} = \begin{cases} \frac{\lambda_1}{2}, \frac{\lambda_2}{2} & f \rightarrow 0 \\ \frac{\lambda_1 + \lambda_2}{4} \pm i \frac{f}{\sqrt{\omega_1 \omega_2}} & \lambda_1 \rightarrow \lambda_2 \end{cases} \quad (5.6)$$




Figure 5.6: Logarithm of the sum of the Fourier transforms of resonators $x(t)$ and $y(t)$, computed by numerical integration of [equation \(5.1\)](#) with initial conditions $x(0) = \dot{x}(0) = \dot{y}(0) = 0$ and $y(0) = 1$. The parameters are $\omega_1 = 1$, $\omega_2 = 1.2$, $Q = 500$. Top panel: Dependence of the spectrum on coupler frequency, for $f = 0.1$. The coupler frequency is given by $s(\omega_2 - \omega_1)$ with $s \in [0, 1.5]$. The analytical treatment above is found at $s = 1$ and indicated by orange arrows. Bottom panel: Dependence on the coupler strength f . The coupling modulation frequency is set at exactly $\omega_2 - \omega_1$, while the coupler strength f is scanned in the range $[0, 0.3]$.

Assuming the second case where $f \gg |\lambda_1 - \lambda_2|$ and setting $\lambda = \lambda_1 = \lambda_2$, one obtains for the two modes

$$\Xi_+(t) = \frac{e^{it\frac{f}{2\sqrt{\omega_1\omega_2}} - t\frac{\lambda}{2}}}{\sqrt{|\omega_1| + |\omega_2|}} \begin{bmatrix} e^{i(t\omega_1 + \varphi)}\sqrt{\omega_1} \\ e^{it\omega_2}\sqrt{\omega_2} \end{bmatrix}, \quad (5.7)$$

$$\Xi_-(t) = \frac{e^{-it\frac{f}{2\sqrt{\omega_1\omega_2}} - t\frac{\lambda}{2}}}{\sqrt{|\omega_1| + |\omega_2|}} \begin{bmatrix} e^{i(t\omega_1 + \varphi)}\sqrt{\omega_1} \\ -e^{it\omega_2}\sqrt{\omega_2} \end{bmatrix}. \quad (5.8)$$

$$(5.9)$$

As expected, there is a splitting of the two coupled modes. The individual resonators oscillate with frequencies

$$\omega_1 \rightarrow \omega_1 \pm \frac{f}{2\sqrt{\omega_1\omega_2}} \quad \text{and} \quad \omega_2 \rightarrow \omega_2 \pm \frac{f}{2\sqrt{\omega_1\omega_2}}. \quad (5.10)$$

The relative weight of the modes depends on the detuning, and the relative phase is imprinted by the drive. The above calculations are checked against a full numerical simulation of the original [equation \(5.1\)](#) in [figure 5.6](#), displaying the spectral dependence of the coupled system as a function of coupling detuning and strength, and [figure 5.7](#), which shows simulated lock-in measurements for various coupler phases and detunings.

Before moving on to the experimental implementation of the above theory, it should be noted that these coupling modulation results generalize to arbitrary numbers of resonators. As such, this machinery can potentially be used to completely circumvent tuning of the resonators, while still allowing the implementation of arbitrary hopping models.

The idea is to write [equation \(5.5\)](#) as

$$\Omega \dot{\vec{X}} = \mathcal{D} \vec{X}, \quad (5.11)$$

where $\Omega = \text{diag}(\omega_1, \omega_2, \dots, \omega_N)$ and $\mathcal{D}_{ij} = -\frac{\lambda_i}{2} \delta_{ij} - i(1 - \delta_{ij})\frac{f_{ij}}{2}$. Now write $\Omega = LL^\dagger$ with $L = \text{diag}(\sqrt{\omega_1}, \sqrt{\omega_2}, \dots, \sqrt{\omega_N})$ and use it as

$$LL^\dagger \dot{\vec{X}} = \mathcal{D}(L^\dagger)^{-1}L^\dagger \vec{X} \quad (5.12)$$

$$L^\dagger \dot{\vec{X}} = (L^\dagger)^{-1} \mathcal{D}(L^\dagger)^{-1} L^\dagger \vec{X} \quad (5.13)$$

$$\dot{\vec{Z}} = \mathcal{A} \vec{Z}, \quad (5.14)$$




Figure 5.7: (a) Simulated lock-in measurements of the driven two-resonator system as a function of coupler phase ϕ . Bottom panel: energy in resonator #1 at the drive frequency ω . Middle panel: energy in resonator #2 at the frequency $\omega + (\omega_2 - \omega_1)$. Top panel: the phase difference between the two resonators. One can see that the energy is not dependent on the coupler phase ϕ . The phase difference between the resonators follows the coupler phase ϕ , as predicted by the analytical treatment. (b) Same as in (a) with a constant $\phi \equiv 0$ but as a function of coupler detuning s . The system is found to be stable with respect to a small detuning. The eigenstates remain bonding and anti-bonding states, the relative weight follows the detuning.

with the definitions

$$\vec{Z} := L^\dagger \vec{X} \quad \text{and} \quad \mathcal{A} = (L^\dagger)^{-1} \mathcal{D} (L^\dagger)^{-1} \quad \text{and} \quad \mathcal{A}^\dagger = \mathcal{A}, \quad \text{if} \quad \mathcal{D}^\dagger = \mathcal{D}. \quad (5.15)$$

Generically, this new matrix takes the form

$$\mathcal{A} = \frac{-i}{2} \begin{bmatrix} -i\lambda_1 & \frac{f_{12}e^{i\varphi_{12}}}{\sqrt{\omega_1\omega_2}} & \dots & \dots \\ \frac{f_{12}e^{-i\varphi_{12}}}{\sqrt{\omega_1\omega_2}} & -i\lambda_2 & \frac{f_{23}e^{i\varphi_{23}}}{\sqrt{\omega_2\omega_3}} & \dots \\ \vdots & \frac{f_{23}e^{-i\varphi_{23}}}{\sqrt{\omega_1\omega_3}} & -i\lambda_3 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}. \quad (5.16)$$

In other words, to implement an arbitrary hopping model with an arbitrary phase pattern, all one needs to do is the following: (1) Determine all single-resonator frequencies ω_i . (2) Modulate all coupling electrodes in sequence according to $V_{ij}(t) = V_{ij}^0 \cos[(\omega_i - \omega_j)t + \varphi]$, and sweep the amplitude V_{ij}^0 to determine the effective coupling strength $t_{ij} = e^{i\varphi_{ij}} f_{ij}^{\text{eff}}(V_{ij}^0)$ by measuring the mode splitting. Having determined the t_{ij} , one can then build an arbitrary honeycomb hopping model with arbitrary phases by choosing corresponding coupler amplitudes and phases.

With the theory of modulated coupling developed, the experimental implementation will be presented next. As the results of [section 5.1](#) have shown no measurable response of the resonators to static voltages, the expectation is for the coupling strength to be extremely weak. Guided by the computations shown in [figure 5.6](#), the resonator pair from [figure 5.5](#) is further investigated in the following protocol.

In a first, the coupling voltage is constantly modulated at high amplitude and precisely determined difference frequency. Either of the two resonators is then resonantly driven via its central electrode at increasing amplitudes, and the response on the other resonator at its resonance frequency is determined via a lock-in measurement. The experiment is repeated several times for each driving amplitude, and the driven resonator's amplitude is measured as well, to ensure actual excitation. Once all sensible amplitudes have been scanned, the role of the resonators is inverted and the experiment repeated. This approach is legitimized by the assumption of very weak coupling, and the results obtained in [figure 5.6](#). The expectation is to see a growing response in the undriven resonator, as the drive on the other scales up.

In the second step, one of the resonators is constantly being excited to a high amplitude, while the coupling amplitude is scanned. Again, both resonators are measured in this experiment, and it is performed with both

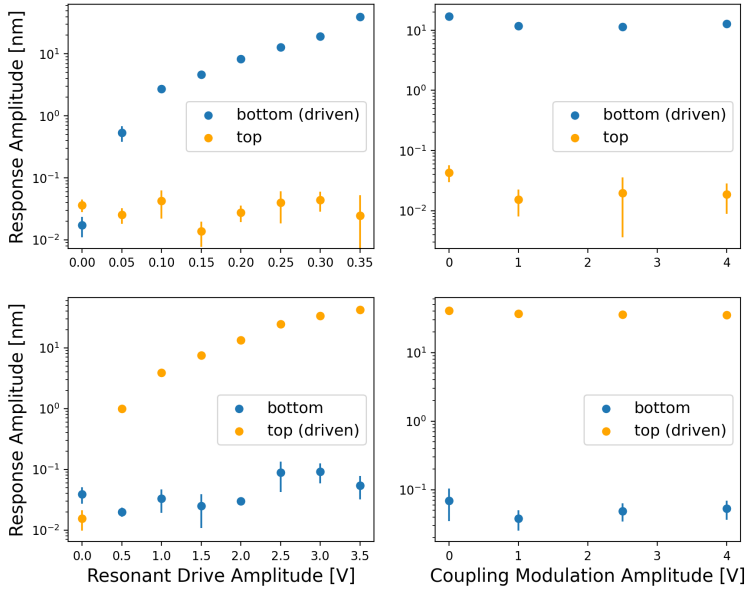


Figure 5.8: Modulated coupling experiments with detuned nearest neighbour drums. Left column: One resonator is driven at resonance with increasing amplitudes (top panel: top resonator driven at 24.597 kHz, bottom panel: top resonator driven at 25.945 kHz), while the coupling voltage is modulated at $\Delta\nu = 1.3475$ kHz with $4.75 V_{pp}$ amplitude. Right column: One resonator is driven at resonance to a high amplitude (top panel: bottom resonator driven with $0.25 V_{pp}$, bottom panel: top resonator driven at $3.5 V_{pp}$), while the coupling voltage is modulated at $\Delta\nu$ with increasing amplitudes. In all cases, the response amplitudes of the individual resonators are extracted with a lock-in measurement at their individual free resonance frequencies. No signatures of coupling have been observed.

resonators acting as driver. The growing coupling modulation amplitude corresponds to a growing coupler strength f , which, while increasing the splitting for larger values, should initially also lead to an increasing response.

The results of these two measurements are shown in [figure 5.8](#). No signature of any finite coupling is detected. While the modulated coupling approach in principle facilitates the interaction between detuned res-

onators, the involved orders of magnitude are still important. The relevant contribution proportional to $\cos \omega_2 t$ on the coupling electrode arises from the sum of the oscillation of resonator 1 proportional to $\cos \omega_1 t$ and the applied coupling modulation voltage proportional to $\cos(\omega_2 - \omega_1)t$. While the magnitude of the latter summand is voltage driven, the former acts by changing the gap of a statically charged capacitor. Using [equation \(2.1\)](#) one can see that these two situations lead to forces of vastly different magnitude. Changing an electrode voltage from 1 V to 5 V increases the force by 2400%. In contrast, reducing the capacitor gap from an initial $1 \mu\text{m}$ by 100 nm leads to an increase in force of just 24%, or 1% of what was found for the voltage case. And this is only the case where a resonator is driven at a barely safe amplitude, and the interwafer gap is $1 \mu\text{m}$. With sample F1F4, the situation is likely less optimal, and the gap modulation contribution is several orders of magnitude smaller than what's achievable with voltage modulation. The signal at the target frequency ω_2 is then likely orders of magnitude too small to effect any measurable influence on the target resonator.

Having presented the measurements, the discussion of the synthetic metamaterial platform concludes. In the next chapter, the focus changes to a more theoretical topic, namely a computational platform for the design of new materials with interesting properties.

6 | Structure Search Design Suite

Topological states are interesting objects from the point of view of material design. They offer indices that can be used as verification flags to ascertain their presence, and they can exhibit a degree of stability unavailable in trivial states. While topological quantum chemistry offers ways to recognize topological states when they arise, designing a structure to host a predefined topological state is a difficult inverse problem. Most designer materials leveraging topology rely heavily on manual input during the design stage. In this chapter, a design stack for the automatic retrieval of designer topological materials is introduced to lift these current limitations. The content presented here is in the process of being published [19].

With the recent advent of *designer materials*, a new application of topology in physics emerged. In photonic [27] and phononic metamaterials [111, 112], nano-structured electronic systems [113, 114], as well as in ensembles of cold atoms [115–119], periodic structures are created to induce artificial Bloch bands. The properties of these Bloch bands are then used to either achieve a sought after functionality of the architected material, e.g., a topological laser [120], or to observe new physical phenomena [14, 18, 121, 122].

Encoding a functionality or a new physical phenomenon in terms of a topological index is a powerful tool for two reasons. First, such an index serves as a simple and stable objective function in the *design* of a periodic




Figure 6.1: Complete catalog of samples with fragile bands. Inset: Illustration of a topological set of bands that cannot be written in terms of a tight-binding model. The yellow dots indicate the high-symmetry points in the Brillouin zone where the eigenvalues of the space group symmetries determine the topology. Main panel: For each wallpaper group the unit cell of a two-dimensional structure is presented, which describes a mass profile $\sigma(x, y)$ or a periodic dielectric constant $\epsilon(x, y)$ that leads to fragile topological bands for the phonons or photons, respectively. Structures for all the 79 distinct phases that have been theoretically predicted in [123] are found.

structure, may it be a specific optical lattice for cold atoms, an arrangement of gates defining a periodic potential for electrons in a two-dimensional electron gas, or an architected photonic or phononic crystal. Second, the achieved functionality has the chance to *display some protection* against fabrication or implementation imperfections.

As was alluded to in [chapter 1](#), finding a periodic structure whose bands have a prescribed topological index is a complicated inverse problem with no generic solution. Despite this complication, the field of designer topological materials enjoyed remarkable success over the years

[14, 18, 20, 113, 124–127]. However, in all the above examples a periodic structure had to be conceived and optimized by hand, typically guided by a simple discrete model. This approach is, however, unsuitable for large-scale, high-throughput explorations, or for problems where either no simple discrete model is known or fabrication constraints severely obstruct the intuition based workflow. Here, a systematic approach to this challenge is described and the methodology is employed on one specific example and an exhaustive characterization of a class of topological bands known as “fragile topological” is presented in all the aforementioned platforms, an impossible endeavor without a fully automated algorithm.

How are such topological bands typically found? For microscopic quantum materials databases such as the Inorganic Crystal Structure Database [128] served as a shopping list to identify new topological systems [129]. Here, focus is placed on the intermediate to large scale designer materials described above. The Bloch bands in all of these systems are described by a partial differential equation where one of their coefficients $f(\vec{r})$ is a periodic function of space. This can be a periodic potential $V(\vec{r})$ for the Schrödinger equation for electrons or ultra-cold atoms [130], a dielectric function $\epsilon(\vec{r})$ in the Maxwell equations for photonic crystals [130], or a mass distribution $\sigma(\vec{r})$ for the Poisson equation describing the vibrations in thin membranes [130].

These designer materials come with their advantages and challenges. The designed is not constrained by combinations of atoms that form chemically stable compounds like in Ref. [129], but the design space is given by a set of reasonably well-behaved functions of the spatial coordinates. This extreme design flexibility comes at a steep algorithmic price: Optimizing a structure $f(\vec{r})$ with the goal to achieve some topological nature of the elementary excitations is suffering from all the challenges inherent to high-dimensional optimization routines. Here, this challenge is overcome by using a modern evolution strategy in the form of the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [131, 132].

The approach is based on topological phases that can be detected by eigenvalues of crystalline symmetries. There are a number of reasons for this choice. First, almost all known topological insulators have a crystalline counterpart, including Chern insulators [133] or \mathbb{Z}_2 [16] insulators, to name just two examples. Second, for many designer materials, crystalline symmetries are a natural choice, as one typically replicates a lo-

cal pattern periodically. Finally, eigenvalues of crystalline symmetries are straightforward to compute and hence serve as a computationally cheap input for a large scale search of topological bands.

Concretely, the framework of topological quantum chemistry [16, 134] is used: Sets of isolated bands are deemed topological if one cannot write them in terms of a basis of exponentially localized symmetric orbitals, see Fig. 6.1. It turns out, that there are two distinct ways to fail this test: Either, bands are topological and can only be trivialized by adding another set of topological bands. All well established topological insulators having a Chern number or a \mathbb{Z}_2 index fall into this category of “stable topology”. However, there arises the possibility that bands are topological, but can be trivialized by a band that is itself trivial. This constitutes the class of “fragile topology”. It turns out by simple inspection [135], that in two dimensions without spin-orbit interaction, which for the chosen designer-materials platforms is the standard rather than the exception, fragile topology is the only possibility. Moreover, the treatment will be constrained to systems that do not break time reversal symmetry for simplicity.

All two-dimensional fragile phases have been tabulated according to their symmetry eigenvalues in Ref. [123]. Each phase can be labelled by the irreducible representations (irreps) of the wallpaper groups at high symmetry points in reciprocal space. For example, for the wallpaper group $p4mm$ with a four-fold rotation axis and two mirror planes, the possible irreps are $\Omega_1, \dots, \Omega_5$ for $\Omega = \Gamma = (0, 0)$ and $\Omega = M = (1/2, 1/2)$, respectively, as well as X_1, X_2 for the point $X = (1/2, 0)$. The points in the Brillouin zone are written with respect to the reciprocal lattice vectors. An example for a set of connected fragile bands in the table of Ref. [123] is

$$\Gamma_1 \Gamma_4, M_5, X_1 X_2; \tag{6.1}$$

where at the points Γ and X two singly degenerate bands are realized which are joined in the two-dimensional irrep M_5 at the M point [130]. There are 79 different distinct sets of irreps, or roots, which describe fragile topological bands in 11 of the 17 wallpaper groups [136]. While some of these have been experimentally observed [18], it is not known if all of them can occur in a realistic material. To find concrete material examples for all of them, to characterize those structures, and to optimize them for scientific investigations or technological applications is the concrete challenge to be met with the design algorithm proposed here.

The outline of the algorithm goes as follows. The design space is given by two-dimensional periodic functions $f(\vec{r})$. These functions are parametrized via their Fourier-coefficients $f(\vec{r}) = \sum_n u_n e^{i\vec{k}_n \cdot \vec{r}}$. The \vec{k}_n are chosen to reflect the specific lattice in a given wallpaper group and the expansion coefficients $\{u_n\}$ fulfill the necessary relations for different n 's, such that all symmetries of the wallpaper group are realized [136, 137]. For some platforms the profile is also discretized by using $\tilde{f}(\vec{r}) = f_0 + \frac{f_1 - f_0}{2} \{1 + \tanh[f(\vec{r})/\xi]\}$ with $\xi \rightarrow 0$ to encode a system where \tilde{f} only takes the two values f_0 and f_1 . The set of independent $\{u_n\}$ make up the high-dimensional continuous search space for the algorithm.

The algorithm starts with a random set of $\{u_n\}$ and solves the partial differential equations at the high-symmetry points in the Brillouin zone using a finite elements solver. In order to exactly respect the symmetries of the wallpaper group in question, the structure is meshed using a custom developed symmetry preserving mesher. For each point the solutions form irreps of the respective little group, which can be labelled by the standard symbol of the high-symmetry point and a numerical label [130]. When ordering the symbols of a given solution by eigenvalues, these form words of the form (again on the example of $p4mm$)

$$\Gamma_1 \Gamma_1 \Gamma_4 \Gamma_4 \dots, M_1 M_1 M_5 M_4 \dots, X_1 X_2 X_2 X_1 \dots \quad (6.2)$$

In a first step, one can *bundle these bands* into set of bands where symmetry compatibility relations require them to be *connected*

$$\Gamma_1^a \Gamma_1^b \Gamma_2^b \Gamma_4^c \dots, M_1^a M_5^b M_4^c \dots, X_1^a X_2^b X_2^b X_1^c \dots, \quad (6.3)$$

where those irreps with the same superscript belong to a given bundle.

In a next step it is checked if any of the bundles conform with the sought after roots of Ref. [123]. If this is the case [in the example above the bundle labelled b realizes the root of Eq. (6.1)], $\{u_n\}$ can be stored for later optimization. If none of the bundles corresponds to the target, one can either move on to the next random set $\{u_n\}$, or one can use a suitable string distance between (6.3) and a string that contains the sought after root as a cost function for CMA-ES. In CMA-ES, the parameters $\{u_n\}$ are drawn from a normal distribution $\mathcal{N}(\{\bar{u}_n\}; C_{nm})$, characterized by the means \bar{u}_n and the covariance matrices $C_{nm} = \overline{(u_n - \bar{u}_n)(u_m - \bar{u}_m)}$. The coefficients \bar{u}_n and C_{nm} are then adapted in a evolution strategy to optimize a given cost function [131, 132].

Finally, when a topological band is found, one can optimize further properties by using a refined cost function. A typical example would be to push for a maximally possible band gap without destroying the topological character. Such mixed cost functions are straightforward to implement, and this point will be come back to below.

Now the results achieved with this algorithm searching for all possible fragile bands are presented. For each of the 79 roots, the corresponding structure is displayed in Fig. 6.1. Shown are structures $f(\vec{r})$ that lead to topological bands for electromagnetic TM modes or phonons in thin membranes. Examples for the other platforms and the associated Bloch bands are shown in [130]. These structures represent the main result of this work where it is established that all fragile bands can in principle be realized.

To highlight the power of the algorithm a detailed statistical analysis of structures containing fragile roots is performed next. The most pressing questions one may have are: how likely is it to find a topological band? Does it depend on the wallpaper group? Are there classes of fragile bands that are harder to realize than others? And finally, can one target the lowest Bloch bands only, or does one have to consider highly excited bands? In Fig. 6.2 all the above questions are answered.

In each wallpaper group random structures $\{u_n\}$ are drawn until 10'000 samples with at least one fragile root amongst the 15 lowest bands have been obtained. For each of the wallpaper groups and all the three studied platforms the fraction of topological samples are then reported. It is observed that this depends strongly on the wallpaper group, and to a lesser extent, on the studied platform. One important observation is that for the group $p3$ for all platforms and for $p6$ for the Schrödinger equation, almost half of the random samples contain a fragile root. Hence, these groups may serve as an excellent starting point if one wants to optimize further properties beyond the presence of topological bands.

Next, it is showcases for the two groups $p3$ and $p4mm$ how the relative abundance is distributed amongst the different roots. While for $p3$ all roots are equally likely, in $p4mm$ there is a significant variation between the different roots. This difference between $p3$ and $p4mm$ can be understood by the different structure of the specific roots [130]. Note, that there are also extreme outliers. For example, root #3 in the group $p4gm$ is extremely hard to find. Only one out of 10^5 random structures turn out to realize this root. It is finally observed, that overall fragile bands are more or less

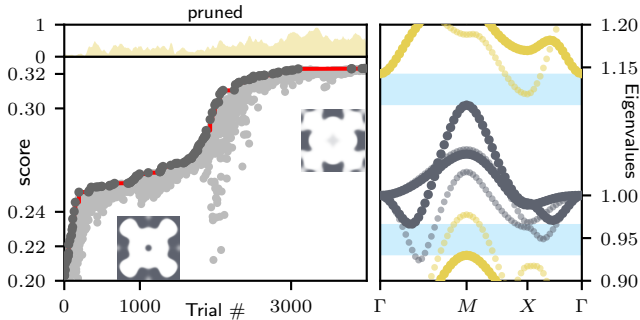


Figure 6.3: Gap optimization of a $p4mm$ sample. Left panel: The score $1/C$ indicating the size of the gap to the adjacent bands as a function of the number of optimization steps. The insets show the initial and the final real-space structure. Also shown is the fraction of samples which are pruned as they lost the initial topological bands. Right panel: Bloch bands along the high symmetry lines of $p4mm$. The small gray dots show the bands of the initial structure where the topological bands largely overlap with the bands below shown in yellow. The big dots represent the final, optimized spectrum with a full band gap to the bands above and below indicated by the blue areas.

equally likely to occur in all the studied platforms.

It is difficult to find fragile bands in the lowest bundle. In Fig. 6.2 the relative abundance of fragile roots as function of band index of the fragile bundle per wallpaper group and platform is shown. One can clearly observe that for some groups one needs to go to rather highly excited bands. This is mostly implied by the relatively complex irreps involved (the counterpart of d -, vs. s -wave orbitals).

In the presented statistical analysis focus was not placed on the size of the band gaps separating the topological from its adjacent bands. This is, however, for most practical applications one of the most important quality measures. It is shown on the example of a specific root in group $p4mm$ how CMA-ES with an appropriate cost function can address this issue. Once the sought after topological band is found in the form of a set of Fourier coefficients $\{u_n^0\}$, one can further optimize these coefficients for a large band gap. The evolution strategy of CMA-ES is optimally suited for this task: One generates additional samples by drawing new ones from $\mathcal{N}(\{u_n^0\}; \lambda_n \delta_{nm})$, with $\lambda_n \ll u_n^0$. The cost function for the update strategy

of \bar{u}_n and C_{nm} can be chosen to take the form

$$C = \frac{1}{\sum_{\vec{k}_i} \epsilon_n(\vec{k}_i) - \epsilon_{n-1}(\vec{k}_i) + \epsilon_{m+1}(\vec{k}_i) - \epsilon_m(\vec{k}_i)}, \quad (6.4)$$

where n (m) label the eigenvalues in lowest (highest) band in the bundle and the \vec{k}_i 's are chosen from a suitable set of reciprocal vectors. In order not to spoil the topology of the bundle in question, one can simply prune those structures $\{u_n^\alpha\}$ which do not conform with the root one optimizes from the population that is used in CMA-ES. In Fig. 6.3 it is shown how a topological bundle that is initially overlapping with its neighboring band below, can be brought into a fully gapped system with a sizeable gap of almost ten percent.

Two observations are worth pointing out. First, the fraction of pruned samples is growing at later stages of the optimization: At this point, improving the gap at all seems to be only possible by introducing unwanted band-inversion. Second, the initial real space structure $f_0(\vec{r})$ and the final $f_{\text{final}}(\vec{r})$ have very little in common with each other. This fact underlines why optimizations by hand or those parametrized by simple geometric motifs often do not yield satisfactory results.

Another way to protect the topology is by separating the fragile bands as much as possible from the trivializing one. Remember that fragile bundles can be trivialized by trivial bands. However, not any trivial band can do that job. By post-selecting those structures where the trivializing band is far away from the bundle in question, one can achieve the same goal as above without the need to optimize any gap. Fig. 6.4 shows one example where the fragile bundle is surrounded by bundles that do not trivialize it, essentially inducing a protecting gap to the one that does trivialize the bundle. One may ask how likely such a situation is. Fig. 6.4.b shows the number of samples as a function of the distance to their trivializing band. While one can see that this distance is exponentially suppressed, distances of up to five bundles are found.

In summary, an algorithm has been presented that can find topological bands by using the covariance matrix adaptation evolution strategy that optimizes the Fourier-representation of periodic functions entering a range of partial differential equations. The power of the approach has been demonstrated by presenting a full catalog of all possible topological bands in two dimensions in the absence of spin-orbit coupling for systems

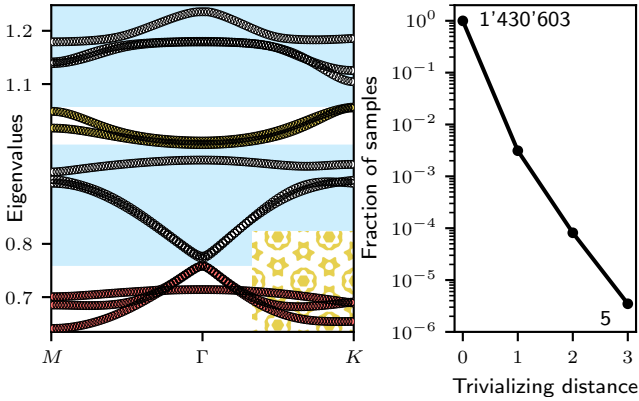


Figure 6.4: Trivializing distance. Left panel: Dispersion relation of a sample (shown in the inset) hosting a topological bundle (yellow) where the trivializing bundle is separated by two “spectator” bundles shown in white. For the purpose of topology, the effective gaps shown in blue are relevant. Right panel: Relative likelihood for a fragile topological bundle to be separated by 1, 2, or 3 bundles from its trivializing one. The graph is based on more than 1.4 million samples, however the overwhelming majority has its trivializing band direct adjacent.

of photons, phonons, electrons and ultra-cold atoms. The taken approach is by no means constrained to the presented application. Generalizations to higher-dimensions, systems that break time-reversal symmetry, or further platforms, including those of non-linear systems, should be straightforward to implement.

7 | Conclusion and outlook

7.1 Conclusion

In this thesis, the history of topological states and how they influenced the development of metamaterials was summarized. It was then explained what the current state of metamaterial design is, and what predominant issues are to be overcome. From this exposition, crucial requirements were formulated, which are to be fulfilled by an adaptive metamaterial platform: manufacturability using standard techniques, ability to host large lattice models, and dynamic – preferably real-time – control of all system parameters.

The main goal of this thesis is the development of a dynamically programmable metamaterial platform that implements these requirements and overcomes limitations of its static predecessors. The targeted platform consists of an intricate micromachined sample, in conjunction with a similarly complex experimental setup. The experimental setup is programmed using a computer, and real-time aspects are then controlled by an FPGA network. The sample consists of 2012 nearest-neighbour coupled silicon nitride resonators, and was designed and optimized over several iterations, along with the corresponding fabrication protocol. While standard MEMS fabrication methods were employed, special focus was placed on optimizing the process to achieve high resonator yield.

A sample manufactured according to the developed designs was manufactured and its properties were explored using the developed setup. Precisely programmable electrostatic excitation of arbitrary resonators in the

lattice was demonstrated. Similar control on the coupling electrodes was also shown. Using this ability, single resonator properties were measured and found to be within the targeted range. However, the sample was found to be lacking in off-resonant operation, likely due to a slightly too large gap between the constituent wafers. Even in the light of this shortcoming, operating the developed setup to reliably control arbitrary resonators represents a successful demonstration of the system as a whole.

With a functional prototype achieved, there are a number of future developments that should be explored. They will be discussed in the next section.

7.2 Outlook

There is a number of interesting pathways the system can be developed along in the future. As the lowest hanging fruit, the issues inherent to the initial version of the sample can be addressed rather easily. In order to improve the off-resonant performance of the sample, the interwafer gap should be reduced. This can be done either by manufacturing samples with various adhesive thicknesses and post-selecting, or by developing a method to reliably predict the adhesive thickness-gap size relation. The wire bond stability issues are already solved, as optimized bondshapes have been developed and tested, they merely need to be applied.

Apart from correcting current imperfections, there are various ways in which the system can be extended.

Firstly, the resonator deformation can potentially be reduced by lowering the stress differential between the resonator's surfaces. This could potentially be achieved by reducing the resonator thickness, and could bring better measurability, simpler design and easier gap control.

Secondly, other wafer stackups with lower residual stress, such as SOI, could be explored. Such configurations could bring lower temperature dependence and better behaviour pertaining deformation.

Thirdly, online lock-in measurements could be implemented on the bridge FPGA. This would facilitate much more efficient experiments by requiring many orders of magnitude less data throughput. A possible further extension could then add automatic frequency sweep recording with arbitrarily long dwell times.

Fourth, the number of sinusoidal generators on the minicircuit FPGAs

could be increased dramatically. This would enable the possibility of precisely modulated couplings between many sites in parallel, which might well be worthwhile once a suitable interwafer gap is achieved.

Fifth, the measurement system could be enhanced by adding motorized control over rotational degrees of freedom of the interferometer fibers. The struggle for signal quality is a limiting factor in the current system, and any enhancement of control on this front could greatly enhance the possibility of automated measurements.

Finally, active pressure control could be integrated in the environmental control system. This would allow dynamic quality factor tuning, thereby enabling experiments leveraging and exploring another degree of freedom.

Bibliography

- [1] Laughlin, R. B. “Nobel Lecture: Fractional quantization”. *Rev. Mod. Phys.* **71**, 863–874 (1999) (cit. on p. 2).
- [2] Klitzing, K. v., Dorda, G. & Pepper, M. “New Method for High-Accuracy Determination of the Fine-Structure Constant Based on Quantized Hall Resistance”. *Phys. Rev. Lett.* **45**, 494–497 (1980) (cit. on p. 3).
- [3] Thouless, D. J., Kohmoto, M., Nightingale, M. P. & den Nijs, M. “Quantized Hall Conductance in a Two-Dimensional Periodic Potential”. *Phys. Rev. Lett.* **49**, 405–408 (1982) (cit. on p. 3).
- [4] Simon, B. “Holonomy, the Quantum Adiabatic Theorem, and Berry’s Phase”. *Phys. Rev. Lett.* **51**, 2167–2170 (1983) (cit. on p. 3).
- [5] Berry, M. V. “Quantal phase factors accompanying adiabatic changes”. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **392**, 45–57 (1984) (cit. on pp. 4, 7).
- [6] Haldane, F. D. M. “Model for a Quantum Hall Effect without Landau Levels: Condensed-Matter Realization of the “Parity Anomaly””. *Phys. Rev. Lett.* **61**, 2015–2018 (1988) (cit. on p. 4).
- [7] Kane, C. L. & Mele, E. J. “Quantum Spin Hall Effect in Graphene”. *Phys. Rev. Lett.* **95**, 226801 (2005) (cit. on pp. 4, 9).
- [8] Kramers, H. A. “Théorie générale de la rotation paramagnétique dans les cristaux”. *Proceedings Koninklijke Akademie van Wetenschappen*, 959–972 (1930) (cit. on p. 4).
- [9] Kitaev, A. “Periodic table for topological insulators and superconductors”. *AIP Conference Proceedings* **1134**, 22–30 (2009) (cit. on p. 5).
- [10] Ryu, S., Schnyder, A. P., Furusaki, A. & Ludwig, A. W. W. “Topological insulators and superconductors: tenfold way and dimensional hierarchy”. *New Journal of Physics* **12**, 065010 (2010) (cit. on p. 5).
- [11] Chiu, C.-K., Teo, J. C. Y., Schnyder, A. P. & Ryu, S. “Classification of topological quantum matter with symmetries”. *Rev. Mod. Phys.* **88**, 035005 (2016) (cit. on p. 5).

- [12] Fu, L. “Topological Crystalline Insulators”. *Phys. Rev. Lett.* **106**, 106802 (2011) (cit. on p. 5).
- [13] Benalcazar, W. A., Bernevig, B. A. & Hughes, T. L. “Quantized electric multipole insulators”. *Science* **357**, 61–66 (2017) (cit. on p. 5).
- [14] Serra-Garcia, M., Peri, V., Süsstrunk, R., Bilal, O. R., Larsen, T., Villanueva, L. G. & Huber, S. D. “Observation of a phononic quadrupole topological insulator”. *Nature* **555**, 342–345 (2018) (cit. on pp. 5, 10, 11, 113, 115).
- [15] Schindler, F., Cook, A. M., Vergniory, M. G., Wang, Z., Parkin, S. S. P., Bernevig, B. A. & Neupert, T. “Higher-order topological insulators”. *Science Advances* **4**, eaat0346 (2018) (cit. on p. 5).
- [16] Bradlyn, B., Elcoro, L., Cano, J., Vergniory, M. G., Wang, Z., Felser, C., Aroyo, M. I. & Bernevig, B. A. “Topological quantum chemistry”. *Nature* **547**, 298–305 (2017) (cit. on pp. 5, 6, 115, 116).
- [17] Po, H. C., Watanabe, H. & Vishwanath, A. “Fragile Topology and Wannier Obstructions”. *Phys. Rev. Lett.* **121**, 126402 (2018) (cit. on p. 6).
- [18] Peri, V., Song, Z.-D., Serra-Garcia, M., Engeler, P., Queiroz, R., Huang, X., Deng, W., Liu, Z., Bernevig, B. A. & Huber, S. D. “Experimental characterization of fragile topology in an acoustic metamaterial”. *Science* **367**, 797–800 (2020) (cit. on pp. 6, 113, 115, 116).
- [19] Bird, S., Devescovi, C., Engeler, P., Valenti, A., Gökmen, D. E., Worreby, R., Peri, V. & Huber, S. D. “Design and characterization of all two-dimensional fragile bands”. *To be published* (2024) (cit. on pp. 7, 113).
- [20] Süsstrunk, R. & Huber, S. D. “Observation of phononic helical edge states in a mechanical topological insulator”. *Science* **349**, 47–50 (2015) (cit. on pp. 7, 8, 115).
- [21] Pendry, J. B. “Negative Refraction Makes a Perfect Lens”. *Phys. Rev. Lett.* **85**, 3966–3969 (2000) (cit. on p. 7).
- [22] Tsaturyan, Y., Barg, A., Polzik, E. S. & Schliesser, A. “Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution”. *Nature Nanotechnology* **12**, 776–783 (2017) (cit. on pp. 8, 22).
- [23] Bückmann, T., Thiel, M., Kadic, M., Schittny, R. & Wegener, M. “An elastomechanical unfeelability cloak made of pentamode metamaterials”. *Nature Communications* **5**, 4130 (2014) (cit. on p. 8).
- [24] Schurig, D., Mock, J. J., Justice, B. J., Cummer, S. A., Pendry, J. B., Starr, A. F. & Smith, D. R. “Metamaterial Electromagnetic Cloak at Microwave Frequencies”. *Science* **314**, 977–980 (2006) (cit. on p. 8).
- [25] Kane, C. L. & Lubensky, T. C. “Topological boundary modes in isostatic lattices”. *Nature Physics* **10**, 39–45 (2014) (cit. on p. 8).
- [26] Wang, Z., Chong, Y., Joannopoulos, J. D. & Soljačić, M. “Observation of unidirectional backscattering-immune topological electromagnetic states”. *Nature* **461**, 772–775 (2009) (cit. on p. 8).

- [27] Ozawa, T. *et al.* “Topological photonics”. *Rev. Mod. Phys.* **91**, 015006 (2019) (cit. on pp. 8, 113).
- [28] Xue, H., Yang, Y. & Zhang, B. “Topological acoustics”. *Nature Reviews Materials* **7**, 974–990 (2022) (cit. on p. 8).
- [29] Lee, C. H., Imhof, S., Berger, C., Bayer, F., Brehm, J., Molenkamp, L. W., Kiessling, T. & Thomale, R. “Topoelectrical Circuits”. *Communications Physics* **1**, 39 (2018) (cit. on p. 8).
- [30] Imhof, S. *et al.* “Topoelectrical-circuit realization of topological corner modes”. *Nature Physics* **14**, 925–929 (2018) (cit. on p. 8).
- [31] Serra-Garcia, M., Süsstrunk, R. & Huber, S. D. “Observation of quadrupole transitions and edge mode topology in an LC circuit network”. *Phys. Rev. B* **99**, 020304 (2019) (cit. on p. 8).
- [32] Hofstadter, D. R. “Energy levels and wave functions of Bloch electrons in rational and irrational magnetic fields”. *Phys. Rev. B* **14**, 2239–2249 (1976) (cit. on p. 8).
- [33] Mishra, M. K., Dubey, V., Mishra, P. M. & Khan, I. “MEMS Technology: A Review”. *Journal of Engineering Research and Reports* **4**, 1–24 (2019) (cit. on p. 9).
- [34] Dutta, S. & Pandey, A. “Overview of residual stress in MEMS structures: Its origin, measurement, and control”. *Journal of Materials Science: Materials in Electronics* **32**, 6705–6741 (2021) (cit. on pp. 9, 24).
- [35] Wu, L.-H. & Hu, X. “Scheme for Achieving a Topological Photonic Crystal by Using Dielectric Material”. *Phys. Rev. Lett.* **114**, 223901 (2015) (cit. on p. 9).
- [36] Cha, J., Kim, K. W. & Daraio, C. “Experimental realization of on-chip topological nanoelectromechanical metamaterials”. *Nature* **564**, 229–233 (2018) (cit. on pp. 10, 11).
- [37] Matlack, K. H., Serra-Garcia, M., Palermo, A., Huber, S. D. & Daraio, C. “Designing perturbative metamaterials from discrete models”. *Nature Materials* **17**, 323–328 (2018) (cit. on p. 10).
- [38] Dubček, T. *et al.* *Binary classification of spoken words with passive phononic metamaterials*. 2023 (cit. on pp. 10, 11, 13).
- [39] Iadecola, T., Schuster, T. & Chamon, C. “Non-Abelian Braiding of Light”. *Phys. Rev. Lett.* **117**, 073901 (2016) (cit. on pp. 12, 18).
- [40] Noh, J., Schuster, T., Iadecola, T., Huang, S., Wang, M., Chen, K. P., Chamon, C. & Rechtsman, M. C. “Braiding photonic topological zero modes”. *Nature Physics* **16**, 989–993 (2020) (cit. on p. 12).
- [41] Chen, Z.-G., Zhang, R.-Y., Chan, C. T. & Ma, G. “Classical non-Abelian braiding of acoustic modes”. *Nature Physics* **18**, 179–184 (2022) (cit. on p. 12).

- [42] Qiu, H., Zhang, Q., Liu, T., Fan, X., Zhang, F. & Qiu, C. “Minimal non-abelian nodal braiding in ideal metamaterials”. *Nature Communications* **14**, 1261 (2023) (cit. on p. 12).
- [43] Kitagawa, T., Berg, E., Rudner, M. & Demler, E. “Topological characterization of periodically driven quantum systems”. *Physical Review B* **82** (2010) (cit. on p. 12).
- [44] Rudner, M. S., Lindner, N. H., Berg, E. & Levin, M. *Anomalous edge states and the bulk-edge correspondence for periodically-driven two dimensional systems*. 2013 (cit. on p. 12).
- [45] Nathan, F., Abanin, D., Berg, E., Lindner, N. H. & Rudner, M. S. “Anomalous Floquet insulators”. *Physical Review B* **99** (2019) (cit. on p. 12).
- [46] Lindner, N. H., Refael, G. & Galitski, V. “Floquet topological insulator in semiconductor quantum wells”. *Nature Physics* **7**, 490–495 (2011) (cit. on p. 12).
- [47] Matheny, M. H., Grau, M., Villanueva, L. G., Karabalin, R. B., Cross, M. C. & Roukes, M. L. “Phase Synchronization of Two Anharmonic Nanomechanical Oscillators”. *PRL* **112** (2014) (cit. on p. 13).
- [48] Kotwal, T., Ronellenfitsch, H., Fischer, M., Stegmaier, A., Thomale, R. & Dunkel, J. “Active topoelectrical circuits”. *arXiv:1903.10130* (2019) (cit. on p. 13).
- [49] Dubcek, T., Ledergerber, D., Thomann, J., Aiello, G., Serra-Garcia, M., Imbach, L. & Polania, R. *Personalized identification, prediction, and stimulation of neural oscillations via data-driven models of epileptic network dynamics*. 2023 (cit. on p. 13).
- [50] Barois, T., Ayari, A., Vincent, P., Perisanu, S., Poncharal, P. & Purcell, S. T. “Ultra Low Power Consumption for Self-Oscillating Nanoelectromechanical Systems Constructed by Contacting Two Nanowires”. *Nano Letters* **13**. PMID: 23528158, 1451–1456 (2013) (cit. on p. 13).
- [51] Tracy, L. A., Luhman, D. R., Carr, S. M., Bishop, N. C., Ten Eyck, G. A., Pluym, T., Wendt, J. R., Lilly, M. P. & Carroll, M. S. “Single shot spin readout using a cryogenic high-electron-mobility transistor amplifier at sub-Kelvin temperatures”. *Applied Physics Letters* **108**, 063101 (2016) (cit. on p. 13).
- [52] Serra-Garcia, M. “Turing-complete mechanical processor via automated nonlinear system design”. *Phys. Rev. E* **100**, 042202 (2019) (cit. on p. 13).
- [53] Hecht-Nielsen, R. *Kolmogorov’s Mapping Neural Network Existence Theorem*. in (1987) (cit. on p. 13).
- [54] Kreinovich, V. Y. “Arbitrary nonlinearity is sufficient to represent all functions by neural networks: A theorem”. *Neural Networks* **4**, 381–383 (1991) (cit. on p. 13).
- [55] Sutton, R. S. & Barto, A. G. *Reinforcement Learning, Second Edition*. ISBN: 9780262039246 (The MIT Press, London, 2018) (cit. on p. 13).

- [56] Appeltant, L., Soriano, M. C., Van der Sande, G., Danckaert, J., Massar, S., Dambre, J., Schrauwen, B., Mirasso, C. R. & Fischer, I. “Information processing using a single dynamical node as complex system”. *Nature Communications* **2**, 468 (2011) (cit. on p. 13).
- [57] Tanaka, G., Yamane, T., Héroux, J. B., Nakane, R., Kanazawa, N., Takeda, S., Numata, H., Nakano, D. & Hirose, A. “Recent advances in physical reservoir computing: A review”. *Neural Networks* **115**, 100–123 (2019) (cit. on p. 13).
- [58] Zheng, T., Yang, W., Sun, J., Xiong, X., Wang, Z., Li, Z. & Zou, X. “Enhancing Performance of Reservoir Computing System Based on Coupled MEMS Resonators”. *Sensors* **21** (2021) (cit. on p. 13).
- [59] Furuhashi, G., Niiyama, T. & Sunada, S. “Physical Deep Learning Based on Optimal Control of Dynamical Systems”. *Phys. Rev. Appl.* **15**, 034092 (2021) (cit. on p. 13).
- [60] Anderson, P. W. “Absence of Diffusion in Certain Random Lattices”. *Phys. Rev.* **109**, 1492–1505 (1958) (cit. on p. 13).
- [61] Younis, M. I. *MEMS Linear and Nonlinear Statics and Dynamics*. ISBN: 978-1-4419-6020-7 (Springer New York, New York, 2011) (cit. on p. 18).
- [62] Ladabaum, I., Jin, X., Soh, H., Atalar, A. & Khuri-Yakub, B. “Surface micromachined capacitive ultrasonic transducers”. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **45**, 678–690 (1998) (cit. on pp. 18, 34).
- [63] Bozkurt, A., Ladabaum, I., Atalar, A. & Khuri-Yakub, B. “Theory and analysis of electrode size optimization for capacitive microfabricated ultrasonic transducers”. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **46**, 1364–1374 (1999) (cit. on p. 18).
- [64] Liu, C. *Foundations of MEMS*. ISBN: 978-0-13-249736-7 (Prentice Hall Press, One Lake Street Upper Saddle River, NJ, United States, 2011) (cit. on p. 18).
- [65] Yasumura, K., Stowe, T., Chow, E., Pfafman, T., Kenny, T., Stipe, B. & Rugar, D. “Quality factors in micron- and submicron-thick cantilevers”. *Journal of Microelectromechanical Systems* **9**, 117–125 (2000) (cit. on p. 23).
- [66] Lifshitz, R. & Roukes, M. L. “Thermoelastic damping in micro- and nanomechanical systems”. *Phys. Rev. B* **61**, 5600–5609 (2000) (cit. on p. 23).
- [67] Duwel, A., Candler, R. N., Kenny, T. W. & Varghese, M. “Engineering MEMS Resonators With Low Thermoelastic Damping”. *Journal of Microelectromechanical Systems* **15**, 1437–1445 (2006) (cit. on p. 23).
- [68] Rodriguez, J., Chandorkar, S. A., Watson, C. A., Glaze, G. M., Ahn, C. H., Ng, E. J., Yang, Y. & Kenny, T. W. “Direct Detection of Akhiezer Damping in a Silicon MEMS Resonator”. *Scientific Reports* **9**, 2244 (2019) (cit. on p. 23).

- [69] Mehmood, Z., Haneef, I. & Udrea, F. “Material selection for Micro-Electro-Mechanical-Systems (MEMS) using Ashby’s approach”. *Materials & Design* **157**, 412–430 (2018) (cit. on p. 23).
- [70] Zawawi, S. A., Hamzah, A. A., Majlis, B. Y. & Mohd-Yasin, F. “A Review of MEMS Capacitive Microphones”. *Micromachines* **11** (2020) (cit. on p. 23).
- [71] Verma, G., Mondal, K. & Gupta, A. “Si-based MEMS resonant sensor: A review from microfabrication perspective”. *Microelectronics Journal* **118**, 105210 (2021) (cit. on p. 23).
- [72] Yang, H, Liu, M, Zhu, Y, Wang, W, Qin, X, He, L & Jiang, K. “Characterization of Residual Stress in SOI Wafers by Using MEMS Cantilever Beams”. *Micromachines (Basel)* **14** (2023) (cit. on p. 24).
- [73] Hooke, R. *Micrographia: or Some Physiological Descriptions of Minute Bodies Made by Magnifying Glasses. With Observations and Inquiries Thereupon* (John Martyn and James Allestry, London, 1665) (cit. on p. 26).
- [74] Newton, I. *Opticks: Or, A treatise of the reflections, refractions, inflections, and colours of light. Also two treatises of the species and magnitude of curvilinear figures* (Printed for Sam. Smith and Benj. Walford, printers to the Royal Society, London, 1704) (cit. on p. 26).
- [75] Ott, H. W. *Cabling*, 44–105. ISBN: 9780470508510 (John Wiley & Sons, Ltd, 2009) (cit. on p. 33).
- [76] Hooke, R. *Lectures de Potentia Restitutiva, Or of Spring Explaining the Power of Springing Bodies* (John Martyn, London, 1678) (cit. on p. 35).
- [77] Torres, J.-M. & Dhariwal, R. S. “Electric field breakdown at micrometre separations”. *Nanotechnology* **10**, 102 (1999) (cit. on p. 36).
- [78] Slade, P. & Taylor, E. “Electrical breakdown in atmospheric air between closely spaced (0.2 /spl mu/m-40 /spl mu/m) electrical contacts”. *IEEE Transactions on Components and Packaging Technologies* **25**, 390–396 (2002) (cit. on p. 36).
- [79] Giere, S., Kurrat, M. & Schumann, U. *HV dielectric strength of shielding electrodes in vacuum circuit-breakers*. in *20th International Symposium on Discharges and Electrical Insulation in Vacuum* (2002), 119–122 (cit. on p. 36).
- [80] Babrauskas, V. *Arc Breakdown in Air over Very Small Gap Distances*. in *Proceedings of Interflam 2013* **2** (2013), 1489–1498 (cit. on p. 36).
- [81] *Substances restricted under REACH*. Accessed: March 2024. European Chemicals Agency (cit. on p. 41).
- [82] *ANNEX XVII TO REACH - Conditions of restriction - Entry 71*. Accessed: March 2024. European Chemicals Agency (cit. on p. 41).
- [83] Wiemer, M., Jia, C., Toepper, M. & Hauck, K. “Wafer Bonding with BCB and SU-8 for MEMS Packaging”. **2**, 1401–1405 (2006) (cit. on p. 41).

- [84] Oberhammer, J., Niklaus, F. & Stemme, G. “Selective wafer-level adhesive bonding with benzocyclobutene for fabrication of cavities”. *Sensors and Actuators A: Physical* **105**, 297–304 (2003) (cit. on p. 41).
- [85] Li, S., Freidhoff, C. B., Young, R. M. & Ghodssi, R. “Fabrication of micronozzles using low-temperature wafer-level bonding with SU-8”. *Journal of Micromechanics and Microengineering* **13**, 732 (2003) (cit. on p. 41).
- [86] Li, Z., Wong, L. L. P., Chen, A. I. H., Na, S., Sun, J. & Yeow, J. T. W. “Fabrication of capacitive micromachined ultrasonic transducers based on adhesive wafer bonding technique”. *Journal of Micromechanics and Microengineering* **26**, 115019 (2016) (cit. on p. 41).
- [87] Bleiker, S., Dubois, V., Schröder, S., Stemme, G. & Niklaus, F. “Adhesive wafer bonding with ultra-thin intermediate polymer layers”. *Sensors and Actuators A: Physical* **260**, 16–23 (2017) (cit. on p. 41).
- [88] Niklaus, F. “Adhesive Wafer Bonding for Microelectronic and Microelectromechanical Systems”. *Signaler, sensorer och system*, 65 (2002) (cit. on p. 41).
- [89] Niklaus, F., Stemme, G., Lu, J. Q. & Gutmann, R. J. “Adhesive wafer bonding”. *Journal of Applied Physics* **99**, 031101 (2006) (cit. on p. 41).
- [90] Niklaus, F., Kumar, R. J., McMahon, J. J., Yu, J., Lu, J.-Q., Cale, T. S. & Gutmann, R. J. “Adhesive Wafer Bonding Using Partially Cured Benzocyclobutene for Three-Dimensional Integration”. *Journal of The Electrochemical Society* **153**, G291 (2006) (cit. on p. 41).
- [91] Yamaner, F. Y., Zhang, X. & Oralkan, O. “A three-mask process for fabricating vacuum-sealed capacitive micromachined ultrasonic transducers using anodic bonding”. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **62**, 972–982 (2015) (cit. on p. 41).
- [92] Meier, E. Private Communication. 2019 (cit. on p. 53).
- [93] Rupprecht, J. *Investigation of Parametric Phenomena of Silicon Nitride Oscillators in Vacuum via Interferometry*. 2023 (cit. on pp. 54, 98, 100, 101).
- [94] *XL Family of Low Phase Noise Quartz-based PLL Oscillators*. Renesas Electronics Corporation (2022) (cit. on p. 63).
- [95] *Numato Systems Numato Saturn*. Numato Systems (cit. on p. 64).
- [96] *Octal, 16-Bit nanoDAC+ with SPI Interface*. Rev. E. Analog Devices (2023) (cit. on p. 76).
- [97] *LogiCORE IP DSP48 Macro v2.1*. DS754. v2.1. Xilinx (2011) (cit. on p. 83).
- [98] *SAM3X / SAM3A Series*. 11057C. Rev. Atmel-11057C-ATARM-SAM3X-SAM3A-Datasheet_23-Mar-15. Atmel (2015) (cit. on p. 88).
- [99] *Arduino Due Homepage*. Arduino (cit. on p. 88).
- [100] *LTC6992-1: TimerBlox: Voltage-Controlled Pulse Width Modulator (PWM)*. Rev. D. Analog Devices (2012) (cit. on p. 89).

- [101] AD7124-4: 4-Channel, Low Noise, Low Power, 24-Bit, Sigma-Delta ADC with PGA and Reference Data Sheet. Rev. E. Analog Devices (2015) (cit. on p. 90).
- [102] Peterson, Z. *Follow Your Multilayer Ground Return Path to Prevent EMI*. Altium, 2020 (cit. on p. 90).
- [103] Peterson, Z. *How to Properly Ground ADCs*. Altium, 2024 (cit. on p. 90).
- [104] Fortunato, M. *Successful PCB Grounding with Mixed-Signal Chips - Follow the Path of Least Impedance*. Maxim Integrated / Analog Devices, 2012 (cit. on p. 90).
- [105] Pithadia, S. & More, S. *Grounding in mixed-signal systems demystified, Part 1*. Texas Instruments, 2013 (cit. on p. 90).
- [106] EVAL-AD7124-4SDZ User Guide. Rev. 0. Analog Devices, 2015 (cit. on p. 91).
- [107] *Completely Integrated, 0.1°C Thermistor Measurement System Using a Low Power, Precision, 24-Bit Sigma-Delta ADC*. Analog Devices, 2020 (cit. on p. 91).
- [108] Balogh, L. *Application Report: Fundamentals of MOSFET and IGBT Gate Driver Circuits*. Rev. A. Texas Instruments, 2018 (cit. on p. 93).
- [109] Cohen, G. H. & Coon, G. A. "Theoretical Consideration of Retarded Control". *Transactions of the American Society of Mechanical Engineers* **75**. Originally published July 1, 1953, 827–834 (2022) (cit. on p. 94).
- [110] Ziegler, J. G. & Nichols, N. B. "Optimum Settings for Automatic Controllers". *Journal of Dynamic Systems, Measurement, and Control* **115**. Originally published 1942, 220–222 (1993) (cit. on p. 94).
- [111] Zhu, W., Deng, W., Liu, Y., Lu, J., Lin, Z.-K., Wang, H.-X., Huang, X., Jiang, J.-H. & Liu, Z. "Topological phononic metamaterials". *arXiv:2303.01426* (2023) (cit. on p. 113).
- [112] Huber, S. D. "Topological mechanics". *Nature Phys.* **12**, 621 (2016) (cit. on p. 113).
- [113] Gomes, K. K., Mar, W., Ko, W., Guinea, F. & Manoharan, C. H. "Designer Dirac fermions and topological phases in molecular graphene". *Nature* **483**, 306 (2012) (cit. on pp. 113, 115).
- [114] Liu, W., Ji, Z. L., Wang, Y., Modi, G., Hwang, M., Zheng, B., Sorger, V. J., Pan, A. & Agarwal, R. "Generation of helical topological exciton-polaritons". *Science* **370**, 600 (2020) (cit. on p. 113).
- [115] Mancini, M. *et al.* "Observation of chiral edge states with neutral fermions in synthetic Hall ribbons". *Science* **349**, 1510 (2015) (cit. on p. 113).
- [116] Stuhl, B. K., Lu, H.-I., Ayccock, L. M., Genkina, D. & Spielman, I. B. "Visualizing edge states with an atomic Bose gas in the quantum Hall regime". *Science* **349**, 1514 (2015) (cit. on p. 113).
- [117] Chalopin, T., Satoor, T., Evrard, A., Makhlov, V., Dalibard, J., Lopes, R. & Nascimbene, S. "Probing chiral edge dynamics and bulk topology of a synthetic Hall system". *Nature Phys.* **16**, 1017 (2020) (cit. on p. 113).

- [118] Braun, C., Saint-Jalm, R., Hesse, A., Arceri, J., Bloch, I. & Aidelsburger, M. “Real-space detection and manipulation of topological edge modes with ultracold atoms”. *arXiv:2304.01980* (2023) (cit. on p. 113).
- [119] Viebahn, K., Walter, A.-S., Bertok, E., Zhu, Z., Gächter, M., Aligia, A. A., Heidrich-Meisner, F. & Esslinger, T. “Interaction-induced charge pumping in a topological many-body system”. *arXiv:2308.03756* (2023) (cit. on p. 113).
- [120] Bandres, M. A., Wittek, S., Harari, G., Parto, M., Ren, J., Segev, M., Christodoulides, D. N. & Khajavikhan, M. “Topological insulator laser: Experiments”. *Science* **359**, 1231 (2018) (cit. on p. 113).
- [121] Lu, L., Wang, Z., Ye, D., Ran, L., Fu, L., Joannopoulos, J. D. & Soljačić, M. “Experimental observation of Weyl points”. *Science* **349**, 6248 (2015) (cit. on p. 113).
- [122] Nash, L. M., Kleckner, D., Read, A., Vitelli, V., Turner, A. M. & Irvine, W. T. M. “Topological mechanics of gyroscopic metamaterials”. *Proc. Natl. Acad. Sci. USA* **112**, 14495 (2015) (cit. on p. 113).
- [123] Song, Z.-D., Elcoro, L. & Bernevig, B. A. “Twisted bulk-boundary correspondence of fragile topology”. *Science* **367**, 794 (2020) (cit. on pp. 114, 116, 118).
- [124] Jotzu, G., Messer, M., Desbuquois, R., Lebrat, M., Uehlinger, T., Greif, D. & Esslinger, T. “Experimental realisation of the topological Haldane model”. *Nature* **515**, 237 (2014) (cit. on p. 115).
- [125] Hafezi, M., Mittal, S., Fan, J., Migdall, A. & Taylor, J. M. “Imaging topological edge states in silicon photonics”. *Nature Photon.* **7**, 1001 (2013) (cit. on p. 115).
- [126] Peterson, C., Benalcazar, W. A., Hughes, T. L. & Bahl, G. “A quantized microwave quadrupole insulator with topologically protected corner states”. *Nature* **555**, 346 (2018) (cit. on p. 115).
- [127] He, H., Qiu, C., Ye, L., Cai, X., Fan, X., Ke, M., Zhang, F. & Liu, Z. “Topological negative refraction of surface acoustic waves in a Weyl phononic crystal”. *Nature* **560**, 61 (2018) (cit. on p. 115).
- [128] *Inorganic Crystal Structure Database* (cit. on p. 115).
- [129] Vergniory, M. G., Elcoro, L., Felser, C., Regnault, N., Bernevig, B. A. & Wang, Z. “A complete catalogue of high-quality topological materials”. *Nature* **566**, 480 (2019) (cit. on p. 115).
- [130] Bird, S., Devescovi, C., Engeler, P., Valenti, A., Gökmen, D. E., Worreby, R., Peri, V. & Huber, S. D. *See supplementary information*. 2024 (cit. on pp. 115, 116, 118, 119).
- [131] Hansen, N. in (eds Lozano, J. A., Larrañaga, P., Inza, I. & Bengoetxea, E.) (Springer, Berlin, 2006) (cit. on pp. 115, 118).

-
- [132] Hansen, N. “The CMA Evolution Strategy: A Tutorial”. *arXiv:1604.00772* (2016) (cit. on pp. 115, 118).
- [133] Fu, L. & Kane, C. L. “Topological insulators with inversion symmetry”. *Phys. Rev. Lett.* **76**, 045302 (2007) (cit. on p. 115).
- [134] Bradlyn, B., Elcoro, L., Vergniory, M. G., Cano, J., Wang, Z., Felser, C., Aroyo, M. & Bernevig, B. “Band connectivity for topological quantum chemistry: Band structures as a graph theory problem”. *Phys. Rev. B* **97**, 035138 (2018) (cit. on p. 116).
- [135] *Bilbao Crystallographic Server* (cit. on p. 116).
- [136] *International Tables for Crystallography* (ed Hahn, T.) (The International Union of Crystallography, Chester, UK, 2005) (cit. on pp. 116, 118).
- [137] Verberck, B. “Symmetry-Adapted Fourier Series for the Wallpaper Groups”. *Symmetry* **4**, 379 (2012) (cit. on p. 118).

List of Publications

- [1] Peri, V., Song, Z.-D., Serra-Garcia, M., Engeler, P., Queiroz, R., Huang, X., Deng, W., Liu, Z., Bernevig, B. A. & Huber, S. D. “[Experimental characterization of fragile topology in an acoustic metamaterial](#)”. *Science* **367**, 797–800 (2020).
- [2] Neugebauer, M. J., Juraschek, D. M., Savoini, M., Engeler, P., Boie, L., Abreu, E., Spaldin, N. A. & Johnson, S. L. “[Comparison of coherent phonon generation by electronic and ionic Raman scattering in \$\text{LaAlO}_3\$](#) ”. *Phys. Rev. Res.* **3**, 013126 (2021).
- [3] Bird, S., Devescovi, C., Engeler, P., Valenti, A., Gökmen, D. E., Worreby, R., Peri, V. & Huber, S. D. “Design and characterization of all two-dimensional fragile bands”. *To be published* (2024).

Curriculum Vitae

Personal data

Name	Pascal Engeler
Date of birth	21st of April 1992
Nationality	Swiss

Education

2018 – 2024	Doctoral Studies in Physics ETH Zurich, Switzerland Advisor: Prof. Dr. Sebastian D. Huber
2015 – 2017	Master of Science ETH in Physics with distinction ETH Zurich, Switzerland. Thesis: “Nonlinear Phononics in LaAlO ₃ – A first principles study” Supervisor: Prof. Dr. Nicola A. Spaldin
2011 – 2015	Bachelor of Science in Physics ETH Zurich, Switzerland.

Acknowledgments

First and foremost, I would like to thank Sebastian Huber. His continuous support and guidance from the very start were an essential factor in the success of this thesis. Wherever the shoe pinched, however catastrophic I perceived an issue to be, I could always count on honest, calm, constructive and informed input conceived with my best interests in mind. I was trusted with a great deal of responsibility in the operation and further development of the laboratory for theoretical physics, and I am grateful for the opportunity.

Next, I would like to thank Marc Serra-Garcia, without whose immense technical experience, intuition and skill no part of the synthetic metamaterial would ever have flown. Never have I learned so much in so little time as while working with Marc, and I am grateful for the instructive and entertaining hours we spent in the lab.

Another person whose indispensable impact on this work shall not be forgotten is Ute Drechsler. Starting out as a complete greenhorn in microfabrication, Ute's guidance allowed me to tap into decades worth of experience and tricks, without which the sample fabrication could never have reached a level anywhere close to where it is now. Without her, working at BRNC would have been a lot less interesting.

I am very thankful to Manfred Sigrist and Ewold Verhagen for agreeing to act as co-examiners for my dissertation.

Further I would like to thank Fabian Kaufmann, whose input on various topics sparked ideas here and there, Philipp Bernhard, Mario Gerber and Elia Dietler. Our shared years of struggle were successful after all.

I would also like to specially mention Yves Acremann, whose input during his lecture and after have sparked my deep interest in, and a fair bit of understanding of, digital electronics.

I owe the past and current members of the CMT-MM group a great deal: Valerio, Tena, Sina, Maciek, Agnes, Eliska, Julia, Efe, Sam, Kukka, Chiara and Caroline. Unforgotten are the experiences shared in Arosa, Israel, the cliffs of Pilatus, the depths of Kunsthaus and on the racetrack. The heated, but always fair, discussions will not be forgotten.

Further thanks are in order to Nicola Spaldin and Dominik Juraschek, whose Master's thesis supervision has shaped me to this day, and Jonathan Home, Vlad Negnevitsky, Matt Grau and Cagri Oenal. You gave me a chance to dive into experimental physics with an internship and took the time to help me get up to speed, and I am grateful for your support.

I would like to thank those, whose input in one way or another shaped my thesis: Steffen Reidt, Richard Stutz, Diana Davila, Petra Burkard, Peter Märki, Albert Schliesser, Richard Norte, Alex Eichler, Mark Fischer.

Last but far from least, I would like to thank my family for their unweary support. Carla, thank you for everything.

Pascal Engeler
Zurich, April 2024

