



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Proof of concept of a hybrid MEMS-magnonic processing unit

TESI DI LAUREA MAGISTRALE IN
PHYSICS ENGINEERING - INGEGNERIA FISICA

Author: **Fabrizio Travagnin**

Student ID: 249819

Advisor: Prof. Riccardo Bertacco

Co-advisors: Antonio Angotti

Academic Year: 2024-25

Abstract

In this master thesis, conducted at Polifab in the framework of the M&MEMS European project, a proof-of-concept device for a hybrid MEMS–magnonic processor has been realized. The work addresses the growing need for alternative computing paradigms beyond CMOS technology, particularly for energy-efficient, high-frequency signal processing and hardware-level artificial intelligence.

The proposed platform exploits the propagation of backward volume spin-waves in a YIG thin film as a physical substrate for wave-based computation. Spin-waves, operating at microwave frequencies with sub-micrometer wavelengths, enable compact and low-power analog signal processing; reconfigurability is achieved by integrating micromagnets on a 4×4 array of piezoelectric MEMS membranes: by electrically actuating the membranes, the magnetic landscape governing spin-wave propagation is dynamically reshaped, enabling programmable control of wave scattering and interference with negligible static power consumption.

Two devices were designed and fabricated: N0, a characterization platform for spin-wave modulation studies, and N1, a reconfigurable processor enabling experimental inverse design. Hardware-level training was demonstrated by iteratively tuning the MEMS control voltages using two different optimization algorithms. Experimental results show effective amplitude modulation, successful objective-function maximization for different tasks, and stable convergence of the control parameters.

Beyond programmable signal processing, the demonstrated architecture provides a scalable route toward the implementation of a physical neural network, where computation emerges from wave dynamics and learning is performed directly at the hardware level through reconfigurable magnetic control.

Keywords: Magnonics, spin-waves, MEMS, Physical Neural Network, Reconfigurable Inverse Design, Wave-based Computing, Beyond-CMOS, Neuromorphic Processor

Abstract in lingua italiana

In questa tesi di laurea magistrale, svolta presso Polifab nell'ambito del progetto europeo M&MEMS, è stato realizzato un dispositivo proof-of-concept per un processore ibrido MEMS-magnonico. Il lavoro si inserisce nel contesto della ricerca di paradigmi computazionali alternativi al CMOS, con particolare riferimento all'elaborazione di segnali ad alta frequenza ed all'intelligenza artificiale a basso consumo energetico.

La piattaforma proposta sfrutta la propagazione di onde di spin di tipo backward volume in un film sottile di YIG come substrato fisico per l'elaborazione di segnale basata su onde. Le onde di spin, operanti nei gigahertz e caratterizzate da lunghezze d'onda sub-micrometriche, consentono un'elaborazione analogica compatta ed energeticamente efficiente; la riconfigurabilità del sistema è ottenuta integrando micromagneti su una matrice 4×4 di membrane MEMS piezoelettriche: l'attuazione elettrica delle membrane permette di modulare dinamicamente il campo magnetico che governa la propagazione delle onde di spin, realizzando un controllo programmabile dei fenomeni di scattering e interferenza con ad alta efficienza energetica.

Durante il corso della tesi, sono stati progettati e fabbricati due dispositivi: N0, una piattaforma di caratterizzazione per lo studio della modulazione delle onde di spin, e N1, un processore riconfigurabile che consente di essere programmato sperimentalmente. Il training a livello hardware è stato dimostrato ottimizzando iterativamente le tensioni di controllo dei MEMS tramite due diversi algoritmi. I risultati sperimentali evidenziano un'efficace modulazione dell'ampiezza, la massimizzazione della funzione obiettivo per diversi task e una convergenza stabile dei parametri di controllo.

Oltre all'elaborazione programmabile del segnale, l'architettura dimostrata rappresenta una piattaforma scalabile per l'implementazione di una rete neurale fisica, nella quale il calcolo emerge direttamente dalla dinamica delle onde e l'apprendimento avviene a livello hardware mediante controllo magnetico riconfigurabile.

Parole chiave: Magnonica, Onde di Spin, MEMS, Physical Neural Network, Hardware Riconfigurabile, Inverse Design, Wave-based Computing, Tecnologie Beyond-CMOS, Processore Neuromorfico

Contents

Abstract	i
Abstract in lingua italiana	iii
Contents	v
Introduction	1
1 Theoretical background	9
1.1 Micromagnetic model	9
1.1.1 Exchange energy	10
1.1.2 Magneto-crystalline anisotropy	11
1.1.3 Demagnetizing energy and shape anisotropy	12
1.1.4 Landau free energy	13
1.2 Magnetization dynamics	13
1.2.1 Small-signal response of a ferromagnetic medium	15
1.2.2 Ferromagnetic resonance	15
1.3 Magnetostatic spin-waves	18
1.3.1 Polder tensor	19
1.3.2 Forward volume spin-waves	20
1.3.3 Backward volume spin-waves	21
1.3.4 Magnetostatic surface spin-waves (Damon–Eshbach)	23
1.3.5 Spin-waves manipulation	24
1.4 Spin-wave excitation via RF antennas	24
1.4.1 Coplanar waveguides	25
1.4.2 Microstrip lines	25
1.5 Wave physics as an analog Recurrent Neural Network	26
1.5.1 Recurrent Neural Networks	26
1.5.2 Wave dynamics as a recurrent process	27

1.5.3	State-space representation and equivalence to an RNN	28
1.5.4	Implications for spin-wave-based implementations	29
1.6	MEMS - PMUTs	30
2	Experimental methods	35
2.1	Fabrication techniques	35
2.1.1	Optical lithography	35
2.1.2	Sputtering	38
2.1.3	Flip-chip	40
2.1.4	White light interferometry	41
2.2	Experimental setup	42
2.2.1	Vector network analyzer	43
2.2.2	RF custom PCBs	46
2.3	Software	47
2.3.1	SPSA algorithm	48
2.3.2	DS algorithm	50
3	Device design and fabrication	53
3.1	Design	53
3.1.1	N0 design	54
3.1.2	N1 design	56
3.2	N0 and N1 fabrication	57
3.2.1	Optical lithography	58
3.2.2	Sputtering and lift-off	60
3.2.3	Flip-chip	61
3.3	Fabrication results	62
3.3.1	N0	62
3.3.2	N1	63
4	N0 - Investigation of a single tunable scatterer	65
4.1	Micromagnetic simulations	65
4.1.1	Perturbation field and magnetization	66
4.1.2	Magnet height effect	69
4.1.3	Frequency dependency	71
4.1.4	Magnet shape	73
4.2	N0 characterization	75
4.2.1	MEMS displacement	75
4.2.2	Raw signal analysis	76

4.2.3	Time gating	78
4.2.4	Voltage modulation	82
4.2.5	Comparison with the simulations	84
5	N1 - Proof of concept processor	87
5.1	Preliminary characterization	87
5.2	Training results	89
5.2.1	Objective function	89
5.2.2	Preliminary device validation	91
5.2.3	DS maximization	93
5.2.4	SPSA maximization	96
5.2.5	DS - SPSA comparison	97
5.2.6	DS minimization	100
5.2.7	Maximize and minimize	101
6	Conclusions and future developments	105
	Bibliography	109
	List of Figures	115

Introduction

For more than half a century, the steady scaling of CMOS technology, commonly described by Moore's law, has driven exponential growth in computational performance. Continuous miniaturization of transistors enabled higher integration density, increased processing speed, and reduced cost per operation. However, as device dimensions approach fundamental physical and technological limits, further scaling no longer guarantees proportional improvements in performance or energy efficiency. Power density, heat dissipation, leakage currents, and interconnect complexity have emerged as dominant constraints, causing a slowdown in traditional performance scaling.

These limitations become particularly critical in the context of artificial intelligence. Classical computing architectures, where memory and processing units are physically separated, are inherently constrained by the Von Neumann bottleneck: the need for continuous data transfer between memory and processor significantly increases latency and energy consumption [1]. Modern artificial neural networks, often comprising billions of trainable parameters, require extensive data movement during both training and inference, and memory access frequently dominates the overall power budget. The scale of this inefficiency is striking when compared to biological systems: while the human brain performs massively parallel, adaptive computation with a power consumption of approximately 20 W, training a large-scale language model can require on the order of 1 MW, an energy budget sufficient to power the human brain for several years [2]. In addition, conventional CMOS technology struggles to replicate the dense three-dimensional interconnectivity characteristic of biological neural networks, leading to architectures that are both physically bulky and energetically inefficient.

To overcome these constraints, significant research efforts are directed toward neuromorphic and beyond-CMOS hardware platforms that aim to co-localize memory and computation and to exploit intrinsic physical dynamics for information processing. Proposed approaches include memristive and resistive switching devices [3] that emulate synaptic behavior at the nanoscale, photonic networks enabling high-speed parallel processing through optical interference [4], spintronic systems leveraging the electron spin degree of freedom for multifunctional computation [1, 2], as well as superconducting and organic

electronic technologies.

Among the emerging beyond-CMOS paradigms, spintronics exploits the electron spin degree of freedom, in addition to its charge, to encode and process information. Rather than replacing conventional electronics, spintronic concepts aim to complement CMOS technology by enabling more efficient execution of specific computational tasks, particularly those dominated by signal processing and parallel operations.

Within spintronics, the field of magnonics focuses on the use of spin-waves and their quanta, magnons, as information carriers [5]. Spin-waves are collective precessional excitations of the electron spin system in magnetically ordered materials such as ferromagnets, ferrimagnets, or antiferromagnets [6]. Unlike charge currents, spin-waves propagate without the motion of real particles, enabling the transfer of spin information with strongly reduced Joule heating.

spin-waves exhibit several properties that make them attractive for signal processing and computing applications. They can reach sub-100 nm wavelengths, while operating at microwave-to-THz frequencies, allowing device miniaturization beyond conventional electromagnetic wave limits. They possess long coherence lengths, up to hundreds of micrometers at room temperature, and exhibit intrinsic nonlinear behavior [7], a key ingredient for complex signal processing and neuromorphic computation. Moreover, their propagation characteristics can be tuned through material choice, magnetic field strength and orientation, and device geometry.

Information in spin-wave systems is typically encoded in the amplitude and/or phase of the wave [8]. In amplitude encoding, logic states correspond to the presence or absence of a wave, whereas in phase encoding, binary states are often assigned to phase 0 and phase π . Because spin-waves are coherent excitations, their interaction is governed by interference. Constructive interference results in amplitude enhancement, while destructive interference leads to cancellation when the phase difference is π . This property enables direct implementation of Boolean operations: for example, amplitude encoding can realize OR and XOR gates, while phase encoding naturally implements majority logic, where the output phase corresponds to the majority of input phases [9].

Beyond Boolean logic, spin-waves provide a powerful platform for analog signal processing. Linear filtering and Fourier transform operations, central to image processing and neural networks and often responsible for the majority of power consumption in digital implementations, can be realized directly through wave propagation, interference, and diffraction in engineered magnetic media [10]. In this wave-computing paradigm, information is encoded in the amplitude and phase of propagating excitations, while computation emerges

naturally from their interaction with a structured medium. As a result, mathematical operations that require extensive multiply–accumulate steps in conventional processors can be performed intrinsically by the physics of the system, without sequential instruction execution.




Figure 1: A micromagnetic simulation showing the Fourier transform property of a spin-wave lens. Image taken from [11].

To harness this capability in a systematic way, inverse-design methodologies have been introduced: instead of manually engineering device geometries based on intuition and trial-and-error, the desired input–output functionality is specified first, and an optimization algorithm iteratively tunes the free parameters of the system to reproduce the target response. In the context of spin-wave systems, Wang et al. [12] demonstrated this approach using a simulated rectangular yttrium iron garnet (YIG) region subdivided into a matrix of $100 \text{ nm} \times 100 \text{ nm}$ elements. Each element could be toggled between material presence or absence, with an effect on a pre-specified objective function evaluated via MuMax3 micromagnetic simulations. Using a custom Direct Binary Search (DBS) algorithm, the authors designed multiple device functionalities within the same framework: a linear frequency multiplexer/demultiplexer, a nonlinear switch that routes signals depending on their power, and a nonreciprocal circulator.

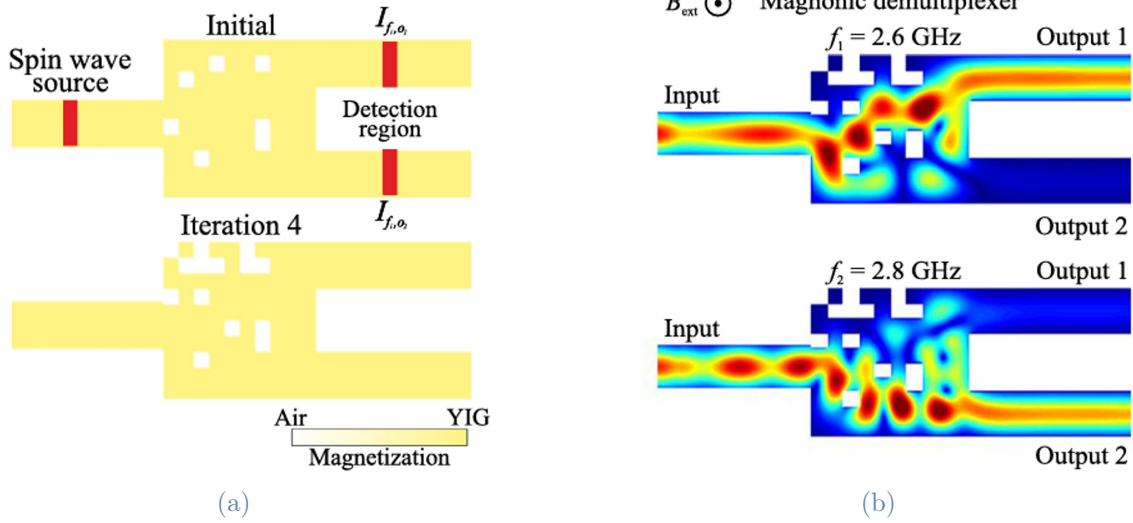


Figure 2: Simulated magnonic inverse design demultiplexer. Image taken from [12].

Additionally, Hughes et al. [13] demonstrated that physical wave systems can function as analog Recurrent Neural Networks (RNNs). By mapping the dynamics of waves in an inhomogeneous medium to the recurrent network equations, the authors showed that the wave field naturally encodes the network’s hidden state and memory through finite-velocity propagation. Using an inverse-design procedure, they optimized the spatial wave-speed distribution to classify spoken vowels directly from raw audio waveforms. The system achieved a mean testing accuracy of 86.3%, comparable to standard digital RNNs, demonstrating that wave-based analog hardware can perform complex temporal signal processing efficiently, in the signal’s native domain, with inherently low energy consumption.

Extending these concepts to fully spin-wave-based neuromorphic hardware, Papp et al. [14] demonstrated a nanoscale neural network where all computational functions, including signal routing and nonlinear activation, are realized through spin-wave propagation and interference. The network’s weights and interconnections are encoded in a spatially patterned magnetic field that scatters the spin-waves across the substrate, with interference creating the desired input–output mapping. A custom GPU-accelerated micromagnetic solver integrated with PyTorch is used to inverse-design the field pattern, effectively training the device. This work illustrates the potential for compact, low-power neural networks that operate entirely within the spin-wave domain, performing complex neuromorphic tasks directly in hardware.

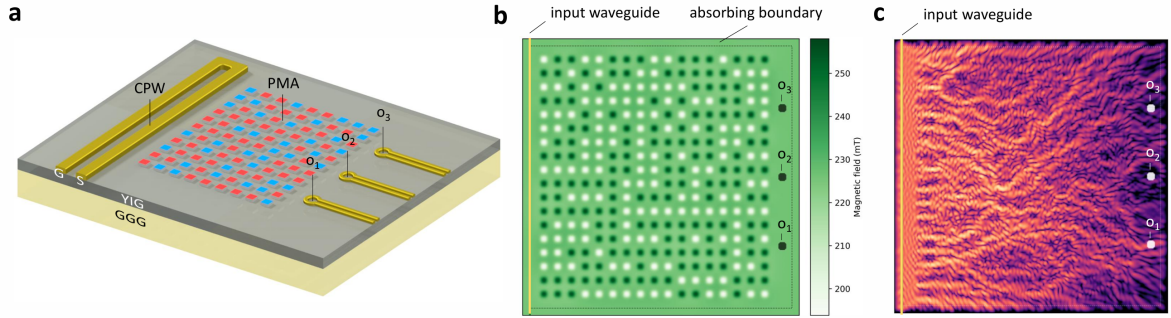


Figure 3: Device schematics and simulated results for a spin-wave-based neural network. Image taken from [14].

A key limitation of traditional inverse-design magnonic devices is their lack of reconfigurability: each device is designed and fabricated to perform a specific function, so that implementing a different functionality requires a complete redesign and fabrication process. Furthermore, simulation-based training is computationally intensive and time-consuming, restricting rapid prototyping and experimental exploration. Performing the optimization directly on a physical device could overcome these constraints, but this approach requires the device itself to be tunable.

To address these challenges, Chumak et al. [15] developed a reconfigurable spin-wave processor based on a 7×7 array of direct-current loops positioned above a YIG thin film, with multiple input and output patterned microstrip antennas for spin-wave injection and detection. By adjusting the currents in each loop, a complex and reconfigurable magnetic field landscape is created, which perturbs spin-wave propagation through interference and scattering. Using experimental feedback-loop algorithms, including direct search and genetic optimization, the system can converge on a configuration that implements a desired function, such as RF notch filters or frequency demultiplexers, directly on the hardware. This methodology enables inverse-design optimization without relying on time-consuming simulations or repeated device fabrication.

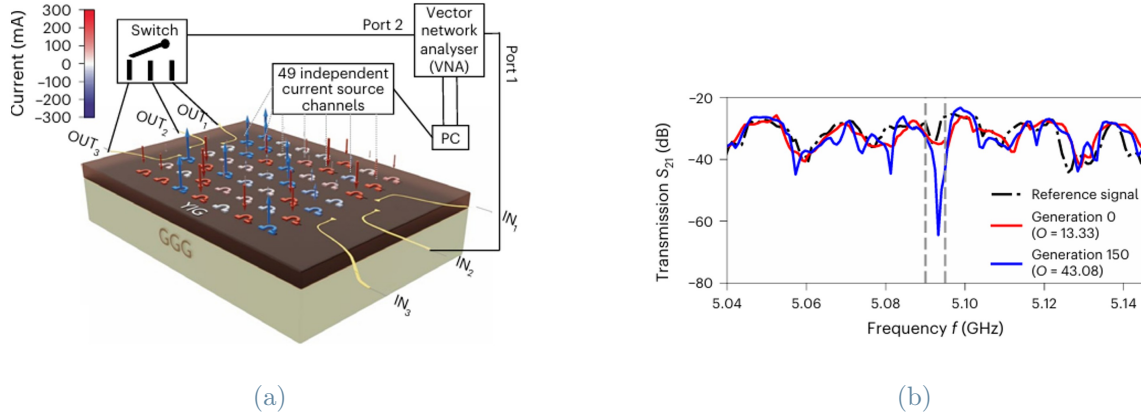


Figure 4: Device scheme and experimental results for the device implemented in [15]. The implemented functionality is a notch filter centered at 5.09 GHz with a bandwidth of 5 MHz.

Despite its reconfigurability and experimental versatility, this approach is energetically inefficient: each loop requires currents of up to 400 mA to generate magnetic fields of approximately 2 mT, resulting in substantial overall power consumption.

This limitation motivates the development of an alternative reconfigurable platform capable of preserving on-device inverse-design optimization while drastically reducing energy requirements. For this purpose, the main device developed in this thesis, referred to as N1, builds on the same physical principle of magnetic landscape engineering but replaces current-driven coils with micromagnets integrated on a MEMS array. In this architecture, the perturbation field is generated by 200 μm in side permalloy micromagnets grown on a 4×4 array of piezoelectric MEMS membranes. The MEMS array is subsequently integrated with the magnonic layer, a YIG thin film, through flip-chip assembly, ensuring close proximity between the micromagnets and the spin-wave medium. The device layout features a single input antenna positioned on one side of the structure and four output antennas on the opposite side, approximately 5 mm apart. By electrically actuating the MEMS membranes, the distance of each micromagnet from the YIG substrate can be tuned, creating a reconfigurable magnetic landscape that perturbs spin-wave propagation through scattering and interference.

The key advantage of this approach lies in its power efficiency. Unlike current loops, which require continuous electrical power to sustain a magnetic field, permanent micromagnets consume virtually no power when held in a static configuration; energy is required only during actuation of the piezoelectric membranes, resulting in negligible static power dissipation. This architecture therefore retains the tunability required for experimental inverse

design while fundamentally addressing the energy-efficiency bottleneck of coil-based systems.

In addition to N1, a secondary device, referred to as N0, was fabricated using the same technological process. Rather than implementing a reconfigurable processor, N0 serves as a dedicated characterization platform to systematically investigate spin-wave propagation in the YIG substrate and to study the effects of micromagnets with different sizes and geometries. This device enables controlled analysis of scattering, diffraction, and phase shifts induced by localized magnetic perturbations, providing the physical insight necessary to guide the design and optimization of N1.

1 | Theoretical background

This chapter provides the theoretical framework necessary for understanding the experimental work presented in this thesis. It begins with the basic concepts of micromagnetism in Section 1.1, introducing the fundamental energy contributions governing ferromagnetic systems and the resulting equilibrium magnetization states. The dynamics of magnetization are then discussed through the Landau–Lifshitz formalism, leading to the description of small-signal response and ferromagnetic resonance in Section 1.2. Building on this foundation, Section 1.3 focuses on magnetostatic spin-waves, their classification into different propagation modes, and the role of the Polder tensor in describing their electromagnetic behavior. The mechanisms for spin-wave excitation using RF antennas are subsequently presented in Section 1.4, with particular attention to coplanar waveguides and microstrip lines. Section 1.5 then presents an interesting parallelism between wave-based dynamical systems and recurrent neural networks, providing the motivation for their use in neuro-morphic computing. Finally, the chapter is concluded with Section 1.6, an introduction to MEMS PMUT devices, which form the basis of the tunable magnetic control architecture employed in this work.

1.1. Micromagnetic model

This section provides a concise overview of the general framework employed to analyze magnetic systems, namely the micromagnetic model. Within this approach, the volume of the magnetic material is discretized into subvolumes of micrometer-scale dimensions. These characteristic length scales are chosen to be sufficiently small to resolve magnetic features such as domain walls, while still remaining large enough for the continuum approximation to be valid. Under this assumption, the magnetization $\mathbf{M}(\mathbf{r})$ is treated as a smooth function of position, with its magnitude varying slowly throughout the volume under consideration.

From this perspective, the key quantity governing the system’s behavior is the spatial orientation of the magnetization rather than its magnitude. Consequently, the micromag-

netic model introduces the reduced magnetization

$$\mathbf{m}(\mathbf{r}) = \frac{\mathbf{M}(\mathbf{r})}{M_s}, \quad (1.1)$$

where M_s denotes the saturation magnetization of the magnetic material.

The model then adopts a thermodynamic formulation to identify a suitable potential capable of describing the magnetic system while accounting for all contributions that determine its energy and equilibrium state. This description is achieved through the definition of an appropriate energy functional, whose minimization yields the equilibrium magnetization configurations. The relevant functional is the Landau free energy, expressed as

$$G_L(\mathbf{H}_a, \mathbf{m}, T) = F(\mathbf{m}, T) - \mu_0 M_s \int_V \mathbf{H}_a \cdot \mathbf{m} \, d\tau. \quad (1.2)$$

In Eq. (1.2), F represents the Helmholtz free energy of the system, while the second term corresponds to the Zeeman energy, accounting for the interaction between the magnetic body and an externally applied magnetic field \mathbf{H}_a .

To fully characterize the energetic state of the system, it is necessary to specify the individual contributions that comprise the Helmholtz free energy. This energy can be formally written as

$$F(\mathbf{m}) = U(\mathbf{m}) = E_{\text{ex}} + E_{\text{an}} + E_{\text{dem}} + E_{\text{in}}. \quad (1.3)$$

Here, E_{ex} denotes the exchange energy, E_{an} the magnetic anisotropy energy, E_{dem} the demagnetizing (magnetostatic) energy, and E_{in} the contribution arising from interfacial effects. The most relevant of these energy terms will be briefly discussed in the following subsections.

1.1.1. Exchange energy

The direct exchange interaction arises from the overlap of electronic wave functions within a magnetic material. Its magnitude depends on both the crystal structure of the system and the spin character of the electrons involved. Fundamentally, the origin of the exchange interaction can be traced back to the requirement that fermionic particles, such as electrons, must exhibit antisymmetric behavior under particle exchange, as dictated by the Pauli exclusion principle. Other forms of exchange, like indirect exchange and double exchange, rely on more subtle quantum mechanical mechanisms, but they can

all be described using a similar Hamiltonian and constitute the primary mechanisms for ferromagnetic and antiferromagnetic order as they energetically favor either parallel or antiparallel alignment of neighboring magnetic moments. In a general description of a crystalline solid, the exchange interaction is commonly modeled using the Heisenberg Hamiltonian,

$$\hat{H} = - \sum_{i,j} J \mathbf{S}_i \cdot \mathbf{S}_j, \quad (1.4)$$

where the summation is typically restricted to nearest-neighbor spin pairs, J denotes the exchange constant, and \mathbf{S}_i and \mathbf{S}_j represent the spin moments associated with the interacting magnetic sites. The strength of this interaction decreases rapidly with increasing separation between magnetic moments.

Within the micromagnetic framework, the discrete Heisenberg Hamiltonian can be reformulated in a continuum description, leading to an expression for the exchange energy of the entire magnetic body. This contribution takes the form

$$E_{\text{ex}} = \frac{A}{2} \int_V [(\nabla m_x)^2 + (\nabla m_y)^2 + (\nabla m_z)^2] d\tau, \quad (1.5)$$

where A is the exchange stiffness constant and m_x , m_y , and m_z are the cartesian components of the reduced magnetization vector \mathbf{m} .

The exchange energy is nonzero only when neighboring magnetic moments exhibit relative angular deviations, thereby it favors parallel (when J is positive) alignment of spins throughout the material. As a result, this interaction promotes the formation of uniformly magnetized, low-energy states, commonly referred to as single-domain configurations. Conversely, magnetization states characterized by rapid spatial variations in orientation, such as those present in multi-domain structures or across domain walls, incur an energetic penalty due to the exchange contribution.

1.1.2. Magneto-crystalline anisotropy

Each magnetic system exhibits an energy contribution that is intrinsically linked to its underlying crystal structure, known as the magneto-crystalline anisotropy energy. This term describes the coupling between the crystallographic lattice and the orientation of the magnetization, giving rise to an anisotropic energy landscape. As a consequence, specific directions within the crystal, referred to as easy axes, are energetically favored, and along them the magnetic moments tend to align in order to minimize the total energy of the system.

Formally, the magneto-crystalline anisotropy energy can be expressed as

$$E_{\text{an}} = \int_V f_{\text{an}}(\mathbf{m}, \mathbf{n}(\mathbf{r})) d\tau, \quad (1.6)$$

where f_{an} denotes the anisotropy energy density, \mathbf{m} is the reduced magnetization, and $\mathbf{n}(\mathbf{r})$ represents the local easy-axis direction, which may in general depend on position.

This energy contribution is nonzero whenever the magnetization deviates from the easy-axis directions. The magnitude and symmetry of the magneto-crystalline anisotropy play a fundamental role in determining the magnetic properties of a material, including its coercive field, remanent magnetization, and overall magnetic hardness or softness.

1.1.3. Demagnetizing energy and shape anisotropy

To account for the long-range dipolar interactions among the magnetic moments within a sample, an additional contribution to the free energy must be introduced. This term is commonly referred to as the magnetostatic free energy, demagnetizing energy, or shape anisotropy energy. As implied by its name, this contribution depends strongly on the geometry of the magnetic body and arises from the magnetic charges that appear at its surfaces and interfaces as a consequence of the divergence of the magnetization field.

These effective magnetic charges generate an internal demagnetizing field \mathbf{H}_M , which generally acts in opposition to the magnetization and gives rise to an additional energetic cost. The demagnetizing energy can be written as

$$E_{\text{dem}} = -\frac{\mu_0}{2} \int_V M_s \mathbf{m} \cdot \mathbf{H}_M d\tau = \frac{\mu_0}{2} \int_{\text{all space}} \mathbf{H}_M^2 d\tau. \quad (1.7)$$

In contrast to the exchange interaction, which favors uniform magnetization, the demagnetizing energy promotes magnetic configurations that reduce the stray field through flux closure. This effect often leads to the formation of multi-domain states, particularly in larger or elongated samples. Consequently, in a body with negligible magneto-crystalline anisotropy, the equilibrium magnetic configuration of a system results from a competition between the exchange energy, which penalizes rapid spatial variations of the magnetization, and the demagnetizing energy, which favors flux-closed domain structures.

1.1.4. Landau free energy

By combining the energetic contributions introduced in the previous sections, it is now possible to write the complete expression of the Landau free energy within the micromagnetic framework. The total functional reads

$$G_L(\mathbf{m}, \mathbf{H}_a) = \int_V \left[\frac{A}{2} ((\nabla m_x)^2 + (\nabla m_y)^2 + (\nabla m_z)^2) + f_{\text{an}}(\mathbf{m}, \mathbf{n}(\mathbf{r})) - \frac{\mu_0}{2} M_s \mathbf{m} \cdot \mathbf{H}_M - \mu_0 M_s \mathbf{m} \cdot \mathbf{H}_a \right] d\tau + \int_S f_{\text{in}}(\mathbf{n}) dS. \quad (1.8)$$

Here, $A = \mu_0 M_s^2 l_{\text{ex}}^2$ denotes the exchange stiffness constant, while l_{ex} is the exchange length. The exchange length represents a fundamental characteristic scale of the system, as it defines the shortest distance over which the magnetization can vary significantly under the combined action of exchange and dipolar interactions. As such, it provides a natural criterion for determining the appropriate discretization length when applying the micromagnetic model to a magnetic body.

The equilibrium configurations of the reduced magnetization field $\mathbf{m}(\mathbf{r})$ are obtained by minimizing the Landau free energy functional. This condition leads to the well-known Brown equations, which express the equilibrium state as

$$\mathbf{m} \times \mathbf{H}_{\text{eff}} = 0 \quad (1.9)$$

where the effective magnetic field \mathbf{H}_{eff} is given by

$$\mathbf{H}_{\text{eff}} = \frac{A}{\mu_0 M_s} \nabla^2 \mathbf{m} + \mathbf{H}_M + \mathbf{H}_a - \frac{1}{\mu_0 M_s} \frac{\partial f_{\text{an}}}{\partial \mathbf{m}}. \quad (1.10)$$

These equations show that equilibrium is achieved when the magnetization aligns locally with the effective field, which incorporates all the energetic contributions discussed previously. The resulting effective-field distribution determines the stable magnetic configuration of the system and serves as the reference state for subsequent analyses, such as the study of magnetization dynamics.

1.2. Magnetization dynamics

The starting point for the study of magnetization dynamics is provided by the Brown equations, which describe static equilibrium configurations of the magnetization. When

a magnetic body is subjected to a magnetic field \mathbf{B} , its magnetic dipoles experience a torque that induces a rotational motion around the direction of the applied field. This phenomenon is known as Larmor precession.

Within a classical description, the precessing magnetic moments can be associated with quantum-mechanical spins through the relation

$$\boldsymbol{\mu} = -\gamma\mathbf{S},$$

where $\gamma = -ge/(2m_e)$ is the gyromagnetic ratio. Under these assumptions, the time evolution of a magnetic moment is governed by

$$\frac{d\boldsymbol{\mu}}{dt} = -\gamma\boldsymbol{\mu} \times \mathbf{B}. \quad (1.11)$$

This expression can be reformulated in terms of the magnetization vector by noting that $\mathbf{M} = n\boldsymbol{\mu}$, where n is the density of magnetic dipoles. Moreover, for a saturated magnetic material, the magnetic induction can be written as $\mathbf{B} = \mu_0\mathbf{H}_{\text{eff}}$. Substituting these relations into Eq. 1.11 yields the equation describing the undamped precessional motion of the magnetization around the effective field,

$$\frac{d\mathbf{M}}{dt} = -\gamma_0\mathbf{M} \times \mathbf{H}_{\text{eff}}, \quad (1.12)$$

where the modified gyromagnetic ratio $\gamma_0 = \mu_0\gamma$ has been introduced. The corresponding angular frequency of precession is given by $\omega_0 = \gamma_0 H_{\text{eff}}$.

Equation 1.12, which can be derived from both classical and quantum-mechanical considerations, represents an idealized description of the magnetization dynamics, as it neglects dissipative effects. In real magnetic systems, the precessional motion is subject to damping mechanisms that progressively reduce the amplitude of the precession, ultimately driving the magnetization to align with the effective field, as required by the equilibrium condition.

A phenomenological description of this dissipative behavior is provided by the Landau–Lifshitz–Gilbert (LLG) equation, which governs the damped dynamics of the magnetization:

$$\frac{d\mathbf{M}}{dt} = -\gamma_0\mathbf{M} \times \mathbf{H}_{\text{eff}} + \frac{\alpha}{M_s}\mathbf{M} \times \frac{d\mathbf{M}}{dt}. \quad (1.13)$$

Here, α is a dimensionless phenomenological parameter known as the Gilbert damping constant, which determines the rate at which the magnetization relaxes toward its equi-

librium orientation. The value of α is material-dependent and reflects the efficiency of the underlying microscopic dissipation mechanisms.

The damping term can be interpreted as contributing to both longitudinal and transverse relaxation processes with respect to the effective field. Together, the precessional and dissipative components of the LLG equation provide a complete description of the time evolution of the magnetization in micromagnetic systems.

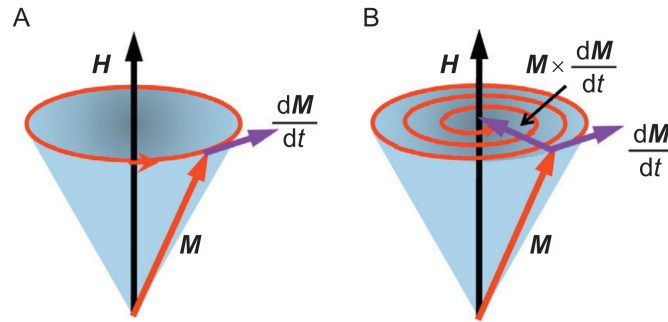


Figure 1.1: Schematic illustration of the precession of the magnetization vector \mathbf{M} around the magnetic field \mathbf{H} : (A) without damping and (B) with damping. Image taken from [16].

1.2.1. Small-signal response of a ferromagnetic medium

The objective of this section is to describe how a magnetic system responds to a small perturbation of its equilibrium configuration. In particular, we focus on the relationship between an applied radio-frequency (RF) magnetic field and the resulting small-signal magnetization, which represents a deviation from the equilibrium state. This linear response framework forms the basis for the description of two fundamental phenomena in magnetization dynamics, namely ferromagnetic resonance (FMR) and spin-wave propagation, which will be addressed in the following sections.

1.2.2. Ferromagnetic resonance

In the linear regime, the response of a magnetic medium to a small RF magnetic field can be described through the small-signal susceptibility tensor, commonly referred to as the Polder tensor. This tensor relates the magnetization and the magnetic field according to

$$\mathbf{m} = \bar{\chi} \mathbf{h}, \quad (1.14)$$

which holds for linear magnetic materials.

To determine the Polder tensor, which depends on the material properties, the Landau–Lifshitz–Gilbert equation is rewritten by introducing small perturbations around the equilibrium configuration:

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\text{eff},0} + \mathbf{h}_{\text{eff}}(t), \quad \mathbf{M} = \mathbf{M}_0 + \mathbf{m}(t). \quad (1.15)$$

Neglecting damping effects, the LLG equation reduces to the undamped precessional equation

$$\frac{d\mathbf{m}}{dt} = -\gamma_0 [(\mathbf{M}_0 + \mathbf{m}(t)) \times (\mathbf{H}_0 + \mathbf{h}(t))]. \quad (1.16)$$

By linearizing the equation and retaining only first-order terms in \mathbf{m} and \mathbf{h} , it is possible to derive the small-signal susceptibility tensor in the form

$$\bar{\chi} = \begin{pmatrix} \chi & -i\kappa \\ i\kappa & \chi \end{pmatrix}, \quad (1.17)$$

where

$$\chi = \frac{\omega_0 \omega_M}{\omega_0^2 - \omega^2}, \quad \kappa = \frac{\omega \omega_M}{\omega_0^2 - \omega^2}, \quad (1.18)$$

with $\omega_0 = \gamma_0 H_{\text{eff}}$ and $\omega_M = \gamma_0 M_s$.

The system is said to be in resonance when its response is maximized. In the lossless limit, this corresponds to the condition $\chi \rightarrow \infty$, which occurs when

$$\omega = \omega_0.$$

Under this condition, the magnetic dipoles undergo a collective and coherent precession around the effective field, giving rise to ferromagnetic resonance (FMR), which corresponds to the maximum absorption of energy from the applied RF signal.

When damping is taken into account, the same linearization procedure can be repeated by introducing the Gilbert damping parameter α . After linearization, the equation of motion becomes

$$i\omega \mathbf{m} = \hat{\mathbf{z}} \times [\omega_M \mathbf{h} - (\omega_0 - i\alpha\omega) \mathbf{m}]. \quad (1.19)$$

The resulting susceptibility tensor retains the same structure as in Eq. 1.17, with the substitution $\omega_0 \rightarrow \omega_0 - i\alpha\omega$:

$$\bar{\chi} = \begin{pmatrix} \chi & -i\kappa \\ i\kappa & \chi \end{pmatrix}, \quad \chi = \frac{\omega_M(\omega_0 - i\alpha\omega)}{(\omega_0 - i\alpha\omega)^2 - \omega^2}, \quad \kappa = \frac{\omega \omega_M}{(\omega_0 - i\alpha\omega)^2 - \omega^2}. \quad (1.20)$$

In realistic systems, the effects of finite magnetic volume and shape anisotropy cannot be neglected, as they significantly influence the resonance condition. For this reason, two thin-film magnetic configurations of particular relevance are considered.

Normally magnetized thin film In this configuration, the film thickness is assumed to be negligible compared to the in-plane dimensions, and the static magnetic field H_{DC} is applied perpendicular to the film plane. As a result, opposite magnetic surface charges appear on the film interfaces, generating a demagnetizing field

$$\mathbf{H}_M = -M_s \hat{\mathbf{z}}.$$

The total static field acting on the film is therefore

$$\mathbf{H} = (H_{\text{DC}} - M_s) \hat{\mathbf{z}}.$$

A small time-dependent perturbation applied perpendicular to the static field does not generate a significant demagnetizing field in this configuration, since the induced magnetic charges are spatially separated. Hence, $\mathbf{h} = 0$, and the FMR condition can be obtained from

$$\mathbf{h} = \bar{\chi}^{-1} \mathbf{m} = \frac{1}{\omega_M} \begin{pmatrix} \omega_0 & i\omega \\ -i\omega & \omega_0 \end{pmatrix} \begin{pmatrix} m_x \\ m_y \end{pmatrix} = 0. \quad (1.21)$$

Non-trivial solutions of this system exist only if the determinant of the matrix is equal to zero, yielding the resonance frequency

$$\omega_R = \gamma_0 (H_{\text{DC}} - M_s). \quad (1.22)$$

Tangentially magnetized thin Film In the second configuration, the static magnetic field is applied in the plane of the thin film. To maintain the static field along the z-direction, the reference frame is rotated accordingly. In this case, unlike the normally magnetized configuration, the demagnetizing field arises in response to the small-signal perturbation and acts along the y-direction.

As a result, the static field is simply $\mathbf{H} = H_{\text{DC}} \hat{\mathbf{z}}$, while the RF demagnetizing field is given by $\mathbf{h} = -\mathbf{N}_a \mathbf{m}$, where the demagnetizing tensor reads

$$\mathbf{N}_a = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}. \quad (1.23)$$

The linearized equation becomes

$$\mathbf{h} = -\mathbf{N}_a \mathbf{m} = \bar{\chi}^{-1} \mathbf{m}, \quad (1.24)$$

which can be rewritten as

$$(\bar{\chi}^{-1} + \mathbf{N}_a) \mathbf{m} = 0. \quad (1.25)$$

Once again, non-trivial solutions exist only if the determinant of the matrix vanishes, leading to the resonance condition

$$\omega_R = \gamma_0 \sqrt{H_{\text{DC}}(H_{\text{DC}} + M_s)}.$$

1.3. Magnetostatic spin-waves

Let us consider uniform plane waves propagating in an arbitrary direction within a magnetized medium. Under the conditions of sufficiently large wave vector magnitude $|\mathbf{k}|$ and $\mathbf{k} \cdot \mathbf{m} \neq 0$, it can be shown that the curl of the magnetic field \mathbf{h} and the electric field \mathbf{e} vanish as $1/|\mathbf{k}|$. As a result, Maxwell's equations reduce to

$$\nabla \times \mathbf{h} = 0, \quad (1.26)$$

$$\nabla \cdot \mathbf{b} = 0, \quad (1.27)$$

$$\nabla \times \mathbf{e} = i\omega \mathbf{b}, \quad (1.28)$$

These equations constitute the *magnetostatic approximation* to Maxwell's equations. Their solutions describe the propagation of electromagnetic disturbances in a magnetized medium, which manifest physically as collective spin precessions within the ferromagnet. In contrast to ferromagnetic resonance, where all spins precess coherently and in phase, magnetostatic waves correspond to neighboring spins precessing with a spatially varying phase. This phase difference increases linearly with position and is responsible for the wave-like propagation.




Figure 1.2: Schematic representation of a spin-wave in a 1D chain of magnetic moments.

The resulting excitations can therefore be interpreted as a particular class of spin-waves, commonly referred to as *magnetostatic spin-waves*. In this regime, dipolar interactions dominate the coupling between spins, and the characteristic wavelengths are typically on the order of micrometers for frequencies in the GHz range. Due to these properties, magnetostatic spin-waves are a good candidate for high frequency signal-processing applications.

It is worth noting that for much shorter wavelengths, in the nanometer range at comparable frequencies, exchange interactions become dominant and must be explicitly included in the theoretical description. Such exchange-dominated spin-waves, however, are currently of more limited technological relevance and will not be considered further here.

1.3.1. Polder tensor

Let us now consider an infinite magnetic medium uniformly magnetized along the z-direction and assume that the magnetostatic approximation remains valid. Since $\nabla \times \mathbf{h} = 0$, the magnetic field can be expressed in terms of a scalar magnetostatic potential ψ as

$$\mathbf{h} = -\nabla\psi.$$

For a magnetized material, the magnetic induction \mathbf{b} can be written as

$$\mathbf{b} = \boldsymbol{\mu} \mathbf{h}, \tag{1.29}$$

where $\boldsymbol{\mu}$ is the permeability tensor of the medium. This tensor can be expressed as

$$\boldsymbol{\mu} = \mathbf{I} + \bar{\chi},$$

with $\bar{\chi}$ denoting the Polder susceptibility tensor introduced previously.

In the following derivation, we adopt the simplest form of the Polder tensor, neglecting both exchange interaction and magnetic anisotropy. Under these assumptions, the susceptibility tensor takes the form

$$\bar{\chi} = \begin{pmatrix} \chi & -i\kappa & 0 \\ i\kappa & \chi & 0 \\ 0 & 0 & 0 \end{pmatrix}, \tag{1.30}$$

where

$$\chi = \frac{(\omega_0 - i\alpha\omega)\omega_M}{(\omega_0 - i\alpha\omega)^2 - \omega^2}, \quad \kappa = \frac{\omega\omega_M}{(\omega_0 - i\alpha\omega)^2 - \omega^2}.$$

Substituting the expression $\mathbf{h} = -\nabla\psi$ into Eq. 1.29, and inserting the resulting expression for \mathbf{b} into Eq. 1.27, one obtains

$$\nabla \cdot (\boldsymbol{\mu} \nabla \psi) = 0. \quad (1.31)$$

Writing the permeability tensor explicitly and expanding the above equation yields

$$(1 + \chi) \left(\frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2} \right) + \frac{\partial^2 \psi}{\partial z^2} = 0, \quad (1.32)$$

which is commonly referred to as *Walker's equation*.

The solutions of Walker's equation describe the allowed magnetostatic spin-wave modes in an infinite, uniformly magnetized medium.

1.3.2. Forward volume spin-waves

We now consider thin-film geometries, which are more technologically relevant than the infinite-medium approximation. In particular, we focus on magnetostatic spin-waves propagating in a normally magnetized thin film.

For spin-waves propagating in the plane of the film, perpendicular to the bias field, we choose the trial potential

$$\psi = \psi_0 \cos(k_z z) e^{i\mathbf{k}_t \cdot \mathbf{r}}, \quad (1.33)$$

where \mathbf{k}_t is the in-plane wave vector and k_z is the wave vector component along the film thickness. This form represents a wave propagating in the plane of the film while forming a standing wave along the z-direction. Outside the magnetic medium, the potential vanishes, ensuring localization.

Applying the boundary conditions for the RF fields \mathbf{h} and \mathbf{b} , and solving Walker's equation, yields the dispersion relation

$$\tan \left[\frac{k_z d}{2} - (1 + \chi) - \frac{n\pi}{2} \right] = \sqrt{\frac{1}{-(1 + \chi)}}, \quad (1.34)$$

where n denotes the quantized mode number arising from the finite thickness of the film.

Kalinikos [17] derived a simplified explicit expression for $n = 0$:

$$\omega^2 = \omega_0 \left[\omega_0 + \omega_M \left(1 - \frac{1 - e^{-k_t d}}{k_t d} \right) \right]. \quad (1.35)$$

Key observations regarding forward volume (FV) spin-waves include:

- All modes share the same cutoff frequency; exchange interaction is needed to lift this degeneracy.
- The dispersion depends only on $|\mathbf{k}_t|$, making propagation isotropic in-plane.
- Phase and group velocities point in the same direction, characterizing these as *forward waves*.
- The wave amplitude varies sinusoidally through the film thickness.

1.3.3. Backward volume spin-waves

For a tangentially magnetized thin film, the propagation characteristics of spin-waves depend on the angle between the wave vector \mathbf{k} and the effective magnetic field \mathbf{H}_{eff} . Backward volume (BV) spin-waves correspond to the configuration $\mathbf{k} \parallel \mathbf{H}_{\text{eff}}$, meaning that the waves propagate parallel to the in-plane magnetization direction.

In this geometry, the magnetostatic potential can be written as

$$\psi = \psi_0 \sin(k_y y) e^{i\nu k_z z}, \quad (1.36)$$

where $\nu = \pm 1$ denotes propagation along the $\pm z$ direction. By applying the appropriate boundary conditions at the film surfaces and following the standard magnetostatic formalism, the dispersion relation can be derived in the form

$$\tan \left[\frac{k_z d}{2\sqrt{-(1+\chi)}} - \frac{(n-1)\pi}{2} \right] = \sqrt{-(1+\chi)}, \quad (2.47)$$

where d is the film thickness, χ is the magnetic susceptibility, and n is the mode index ($n = 1$ corresponds to the fundamental mode).

For the fundamental mode ($n = 1$), a simplified analytical expression for the dispersion relation was derived by Kalinikos [17]:

$$\omega^2 = \omega_0 \left[\omega_0 + \omega_M \left(\frac{1 - e^{-k_z d}}{k_z d} \right) \right], \quad (1.37)$$

where $\omega_0 = \gamma\mu_0 H_{\text{eff}}$ and $\omega_M = \gamma\mu_0 M_s$.

Figure 1.3 shows the dispersion relation of the fundamental backward volume mode together with the first few higher-order quantized modes. The slope of each dispersion curve provides the group velocity, defined as

$$v_g = \frac{\partial\omega}{\partial k}. \quad (1.38)$$

From the plotted curves, it is evident that for small values of k , the fundamental mode exhibits a significantly higher group velocity than the higher-order modes. As a consequence, it reaches the receiving antenna earlier and carries most of the transmitted signal, which explains why, in the experimental spectra shown in these thesis, the dominant contribution is associated with the fundamental mode.




Figure 1.3: Dispersion diagram for backward volume spin-waves with $\omega_0/\omega_M = 0.5$.

Several important properties characterize backward volume spin-waves:

- For $k_z = 0$, the oscillation frequency reduces to the ferromagnetic resonance (FMR) frequency of the film under the given bias field.
- The group velocity is negative, while the phase velocity remains positive. This opposite sign between phase and energy propagation defines backward volume spin-waves as *backward waves*.
- The dynamic magnetization profile is quantized across the film thickness, with a sinusoidal distribution for each mode order n .

1.3.4. Magnetostatic surface spin-waves (Damon–Eshbach)

Finally, for tangentially magnetized films with $\mathbf{k} \perp \mathbf{H}_{\text{eff}}$, the modes are known as *magnetostatic surface spin-waves* or Damon–Eshbach (DE), whose dispersion relation is described by the following relation

$$\omega^2 = \omega_0(\omega_0 + \omega_M) + \frac{\omega_M^2}{4} (1 - e^{-2kd}). \quad (1.39)$$

Properties of DE modes are:

- The dispersion relation is independent of ν , but the mode fields are not; they shift from one surface to the other when the propagation direction is reversed (*field displacement non-reciprocity*).
- Phase and group velocities are parallel, making these waves *forward waves*.
- The wave amplitude decays exponentially from the surfaces, rather than varying sinusoidally across the thickness, making these true surface modes.




Figure 1.4: Dispersion relations for the three types of magnetostatic spin-waves, along with the required geometries to excite them. For each type of wave, only the first quantized mode is shown.

1.3.5. Spin-waves manipulation

As a final remark, it is important to highlight that equations 1.35, 1.37, 1.39 show that the spin-wave mode dispersion depends on ω_0 , which we previously defined as $-\gamma\mu_0 H_0$, but more generally is $-\gamma\mu_0 \mathbf{H}_{\text{eff},0}$. This implies that modifying the effective field $\mathbf{H}_{\text{eff},0}$, either in magnitude or orientation, directly affects key propagation characteristics of spin-waves, including their group velocity, wave vector \mathbf{k} , and phase.

Tuning the equilibrium magnetization landscape is therefore an effective way to manipulate spin-wave propagation. In particular, spatial variations in the equilibrium magnetization can induce reflection, scattering, and diffraction of spin-waves, analogous to wave phenomena in other physical systems. These effects can be deliberately engineered to control the propagation path, interference patterns, and phase of spin-waves, making them a powerful tool for analog signal processing applications, which is the focus of the present thesis.

1.4. Spin-wave excitation via RF antennas

A primary method to excite spin-waves in a magnetized medium is through the application of a small, oscillating magnetic field \mathbf{h}_{RF} , generated by RF antennas. When applied, this field perturbs the equilibrium magnetization, inducing a precession of the spins and launching propagating spin-waves. Conductive antennas are widely used for this purpose due to their versatility, selectivity, and ease of integration, despite a relatively low excitation efficiency.

A simple excitation scheme consists of placing a conductive strip on top of the magnetic film and driving an AC current through it. The oscillating current generates an AC magnetic field in the plane perpendicular to the conductor, with a spatial distribution determined by the geometry of the antenna. The component of the field perpendicular to the static magnetization exerts a torque, bringing the magnetization out of equilibrium and exciting spin-waves along the propagation direction.

The spatial profile of the RF field plays a crucial role in determining which spin-wave wave vectors \mathbf{k} are excited. The Fourier transform of the field distribution, $F(\mathbf{h}_{\text{RF}})$, dictates the excitation efficiency as a function of \mathbf{k} . Antennas with a well-defined spatial periodicity, such as meander-shaped antennas or coplanar waveguides, produce a field with strong Fourier components at specific wave vectors, allowing selective excitation. In contrast, uniform microstrip lines generate a broad Fourier spectrum, exciting a wide range of \mathbf{k} values, including $\mathbf{k} = 0$ (FMR).

1.4.1. Coplanar waveguides

A coplanar waveguide (CPW) consists of a central signal line (S) of width w separated by gaps s from two ground lines (G), all on top of a dielectric layer above the magnetic film. The current distribution in the conductors, including skin and proximity effects, can be approximated as four line currents at the edges of the gaps [18]:

$$J_x(y, z) = \frac{I}{2} \left[\delta\left(y + \frac{w}{2}\right) + \delta\left(y - \frac{w}{2}\right) - \delta\left(y - \frac{w}{2} - s\right) - \delta\left(y + \frac{w}{2} + s\right) \right] \delta(z), \quad (1.40)$$

where I is the current amplitude. The resulting magnetic field can be calculated via Ampère's law, and its Fourier transform has the first maximum and the first minimum respectively at

$$k_{\max} = \frac{\pi}{w + s}, \quad k_{\min} = \frac{2\pi}{w + s}, \quad (1.41)$$

allowing precise control over which spin-wave wave vectors are efficiently excited.

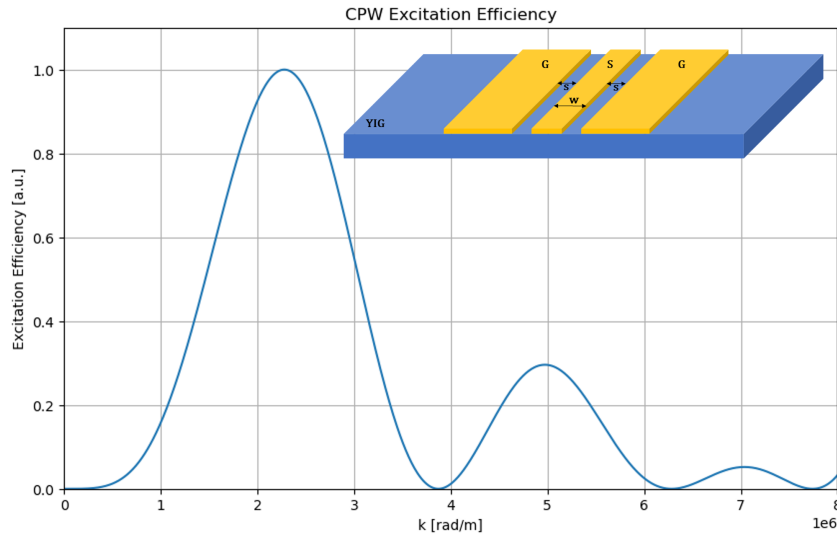


Figure 1.5: Excitation efficiency and diagram for a CPW antenna ($w = 2.5 \mu\text{m}$, $s = 1 \mu\text{m}$).

1.4.2. Microstrip lines

Microstrip antennas consist of a single conductive strip above a dielectric layer. The AC current generates an approximately rectangular in-plane field $h_{\text{RF},y}(y)$ at the magnetic layer, which can be expressed using the Karlqvist[19] approximation for thin strips:

$$h_{\text{RF},y}(y) \propto \frac{I}{2w} \text{rect}_w(y), \quad (1.42)$$

with Fourier transform

$$F(h_{\text{RF},y}(y)) \propto \frac{I}{2} \text{sinc}\left(\frac{kw}{2}\right). \quad (1.43)$$

This shows that microstrip lines provide broadband excitation, including $k = 0$ (FMR), and that the excitation bandwidth can be tuned by varying the strip width.

In summary, the key principle behind RF antenna excitation of spin-waves is that the spacial frequency components of the generated magnetic field determine which spin-wave wave vectors are efficiently excited. By tailoring the antenna geometry, one can selectively excite narrow or broad ranges of \mathbf{k} , making this method highly versatile for spin-wave experiments and devices.

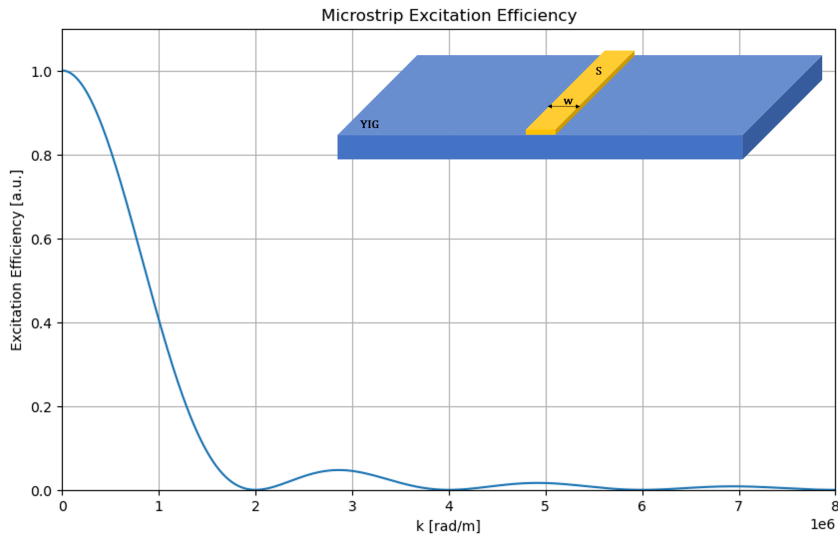


Figure 1.6: Excitation efficiency and diagram for a microstrip antenna ($w = 2.5 \mu\text{m}$).

1.5. Wave physics as an analog Recurrent Neural Network

1.5.1. Recurrent Neural Networks

Recurrent Neural Networks (RNNs) are a class of neural architectures designed to model sequential and temporal data by incorporating feedback connections within the network structure. Unlike feed-forward networks, RNNs maintain an internal state that evolves over time, enabling the processing of input sequences of arbitrary length. This internal state acts as a memory that encodes information from past inputs and influences future network responses. The following discussion has been adapted from [13].

Formally, an RNN operates on a discrete-time sequence $\{\mathbf{x}_t\}$, where the hidden state \mathbf{h}_t at time step t is computed as a function of the previous hidden state \mathbf{h}_{t-1} and the current input \mathbf{x}_t . The standard RNN formulation is given by [20]

$$\mathbf{h}_t = \sigma^{(h)} (\mathbf{W}^{(h)}\mathbf{h}_{t-1} + \mathbf{W}^{(x)}\mathbf{x}_t), \quad (1.44)$$

$$\mathbf{y}_t = \sigma^{(y)} (\mathbf{W}^{(y)}\mathbf{h}_t), \quad (1.45)$$

where $\mathbf{W}^{(h)}$, $\mathbf{W}^{(x)}$, and $\mathbf{W}^{(y)}$ are trainable weight matrices, and $\sigma^{(h)}$ and $\sigma^{(y)}$ are nonlinear activation functions. This recursive structure allows RNNs to represent temporal dependencies by repeatedly applying the same transformation at each time step, rendering them particularly suited to learn complex temporal structures and long-range dependencies in data, making it the industry standard for tasks such as natural language processing [21], machine translation, and time-series prediction[22].

While RNNs are expressive models for sequential data, their training is known to be sensitive to numerical instabilities arising from repeated matrix multiplications through time. These issues are commonly associated with vanishing or exploding gradients during backpropagation through time.




Figure 1.7: Schematic diagram of a Recurrent Neural Network.

1.5.2. Wave dynamics as a recurrent process

The propagation of waves in continuous media is governed by partial differential equations that inherently encode temporal recurrence. In particular, the scalar wave equation describes the evolution of a wave field $u(\mathbf{r}, t)$ as

$$\frac{\partial^2 u}{\partial t^2} - c^2(\mathbf{r})\nabla^2 u = f(\mathbf{r}, t), \quad (1.46)$$

where $c(\mathbf{r})$ denotes the spatially dependent wave speed, ∇^2 is the Laplacian operator, and $f(\mathbf{r}, t)$ represents an external source term. This equation arises in a wide range of physical systems, including acoustics, electromagnetism, and elastic media.

To analyze the temporal structure of the wave dynamics, the equation is discretized in time using centered finite differences with time step Δt . The second-order time derivative is approximated as

$$\frac{\partial^2 u}{\partial t^2} \approx \frac{u_{t+1} - 2u_t + u_{t-1}}{\Delta t^2}, \quad (1.47)$$

which yields the discrete-time update equation

$$u_{t+1} = (2 + \Delta t^2 c^2 \nabla^2) u_t - u_{t-1} + \Delta t^2 f_t. \quad (1.48)$$

This formulation explicitly shows that the state of the system at time $t + 1$ depends on the wave field at the two previous time steps, thereby exhibiting an intrinsic recurrent structure.

1.5.3. State-space representation and equivalence to an RNN

To express the discretized wave equation in a form analogous to a recurrent neural network, the system state is defined as the concatenation of the current and previous wave fields:

$$\mathbf{h}_t \equiv \begin{bmatrix} \mathbf{u}_t \\ \mathbf{u}_{t-1} \end{bmatrix}. \quad (1.49)$$

Where \mathbf{u}_t and \mathbf{u}_{t-1} are vectors given by the flattened fields u_t and u_{t-1} sampled on a discretized grid over the spatial domain. With this definition, the second-order wave equation can be rewritten as a first-order recurrence in an augmented state space

$$\begin{bmatrix} \mathbf{u}_{t+1} \\ \mathbf{u}_t \end{bmatrix} = \begin{bmatrix} 2\mathbf{I} + \Delta t^2 c^2 \nabla^2 & -\mathbf{I} \\ \mathbf{I} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{u}_{t-1} \end{bmatrix} + \Delta t^2 \begin{bmatrix} \mathbf{f}_t \\ \mathbf{0} \end{bmatrix} \quad (1.50)$$

that can be cast in the simpler form

$$\mathbf{h}_t = \mathbf{A}(\mathbf{h}_{t-1})\mathbf{h}_{t-1} + \mathbf{P}^{(i)}\mathbf{x}_t, \quad (1.51)$$

where \mathbf{A} is a sparse matrix encoding the discretized Laplacian operator and the spatial distribution of the wave speed, and $\mathbf{P}^{(i)}$ maps the external input signal \mathbf{x}_t to the system

state. Finally, the system output is obtained through a measurement operator

$$\mathbf{y}_t = (\mathbf{P}^{(o)} \mathbf{h}_t)^2, \quad (1.52)$$

where $\mathbf{P}^{(o)}$ selects spatial locations at which the wave intensity is measured.

This formulation is mathematically equivalent to a recurrent neural network, with the wave field history playing the role of the hidden state and the discretized wave operator acting as the recurrent transformation. Unlike conventional RNNs, the recurrent matrix \mathbf{A} is not freely parameterized but is fully determined by the physical properties of the medium, most notably the spatial distribution of the wave speed $c(\mathbf{r})$. The nonlinearity given by the activation function in Eq. 1.44 is provided in Eq. 1.51 by the dependence on \mathbf{h}_{t-1} of \mathbf{A} .

Within this framework, learning corresponds to modifying the material properties of the medium rather than adjusting abstract synaptic weights. Spatial variations in the wave speed $c(\mathbf{r})$ determine the coupling between neighboring spatial points and thereby shape the temporal evolution of the wave field. This enforces locality and finite propagation speed, which contrasts with the dense and instantaneous interactions characteristic of standard digital RNNs.

Nonlinear behavior, which is essential for emulating the activation function's behavior, arises naturally from physical mechanisms. These can be for example intensity-dependent material responses, such as the Kerr effect in nonlinear optics, as well as nonlinear dynamics in spin-waves. As a result, the role of activation functions in standard RNNs is effectively replaced by intrinsic nonlinearities governed by the underlying physical system.

1.5.4. Implications for spin-wave-based implementations

The wave-based formulation of recurrent neural networks naturally extends to a variety of physical platforms beyond optics and acoustics. In this context, spin-waves emerge as a particularly promising candidate for implementing these recurrent dynamics since they are inherently governed by nonlinear equations of motion, such as the Landau–Lifshitz–Gilbert equation (Eq. 1.13). In [14] a spin-wave-based RNN, where signal propagation, routing, and nonlinear activation arise directly from spin-wave propagation and interference, is simulated. The network's weights and interconnections are encoded in a spatially varying magnetic field pattern applied to the spin-wave medium, which acts as a scatterer shaping wave interference between sources and detectors. Training the network corresponds to solving an inverse problem to determine the magnetic field pattern that pro-

duces a desired input–output mapping, achieved using a micromagnetic solver integrated with the PyTorch machine-learning framework. The study shows that while spin-waves exhibit linear interference at low intensities, their behavior becomes strongly nonlinear at higher amplitudes, leading to a significant increase in computational capability.

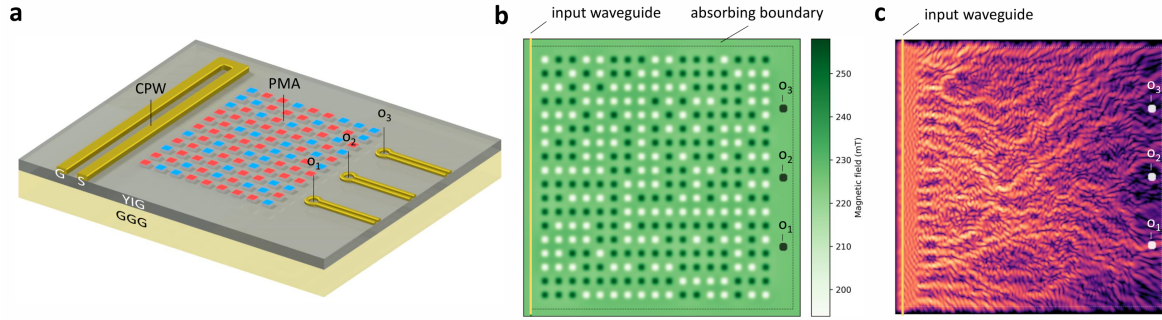


Figure 1.8: Simulation of a RNN implementation using forward volume spin-waves. Image taken from [14].

1.6. MEMS - PMUTs

Micro-electromechanical systems (MEMS) constitute a class of miniaturized devices that integrate mechanical structures, sensors, and actuators on a common substrate, typically silicon, using standard microfabrication techniques derived from the semiconductor industry. Depending on their functionality, MEMS can be broadly classified into sensors, actuators, and resonant structures, and are widely employed in applications ranging from inertial navigation [23][24] and mobile communications to biomedical fields [25] and environmental monitoring [26]. Their success is largely driven by their small footprint, low power consumption, batch fabrication capability, and high level of integration with electronic circuitry [27].

Several actuation principles are commonly employed, including electrostatic, electrothermal, electromagnetic, and piezoelectric approaches. Electrostatic actuation is among the most widely used techniques, exploiting the attractive force generated by an electric field between electrodes; it offers fast response and low power consumption, but typically requires relatively high driving voltages and exhibits nonlinear behavior. Electrothermal actuation relies on Joule heating and differential thermal expansion between material layers, enabling large displacements at low voltages, albeit at the expense of slower dynamics and increased power consumption. Electromagnetic actuation exploits Lorentz forces generated by currents flowing in magnetic fields, providing large and linear displacements,

but often involves complex fabrication and higher energy requirements. Finally, piezoelectric actuation converts an applied electric field directly into mechanical strain in thin-film piezoelectric materials, offering high forces, fast response times, and linearity at relatively low driving voltages, while remaining compatible with CMOS technologies and requiring very little power [28].

In piezoelectric MEMS, mechanical and electrical domains are coupled through the piezoelectric effect, which can be described by the linear constitutive relations

$$\mathbf{S} = \mathbf{s}^E \mathbf{T} + \mathbf{d}^t \mathbf{E},$$

$$\mathbf{D} = \mathbf{d} \mathbf{T} + \boldsymbol{\varepsilon}^T \mathbf{E},$$

where \mathbf{S} is the mechanical strain, \mathbf{T} the applied stress, \mathbf{E} the electric field, and \mathbf{D} the electric displacement. The tensors \mathbf{s}^E , \mathbf{d} , and $\boldsymbol{\varepsilon}^T$ represent the elastic compliance, piezoelectric coefficients, and permittivity, respectively; \mathbf{d}^t is the transpose of \mathbf{d} . These relations explicitly show how electrical and mechanical quantities are coupled in piezoelectric materials: by applying an electric field \mathbf{E} across the piezoelectric layer, a mechanical stress \mathbf{T} is induced, which in turn generates a strain \mathbf{S} , corresponding to a physical deformation and displacement of the membrane. Conversely, applied mechanical stress produces an electrical response. In practical terms, this electromechanical coupling allows direct control of the membrane deflection through an externally applied voltage, forming the basis of piezoelectric MEMS actuation. In most thin-film piezoelectric MEMS devices, actuation is achieved through the so-called 31-mode [29], in which an out-of-plane electric field induces in-plane strain, resulting in bending of the supporting membrane.

Piezoelectric Micro-machined Ultrasonic Transducers (PMUTs) represent a specialized subclass of piezoelectric MEMS designed to generate and detect acoustic waves through the flexural vibration of a multilayer diaphragm. A typical PMUT consists of a thin piezoelectric film, commonly lead zirconate titanate (PZT) for high electromechanical coupling or aluminum nitride (AlN) for CMOS compatibility, deposited on a passive elastic membrane made of materials such as silicon, silicon dioxide, or silicon nitride. This stack is sandwiched between top and bottom electrodes and suspended over a micro-machined cavity, with the membrane clamped at its edges. When an electric field is applied across the piezoelectric layer, in-plane strain is generated, causing the diaphragm to deflect and radiate ultrasonic waves into the surrounding medium.




Figure 1.9: 3D schematic of a PMUT device. Image taken from [30].

Compared to capacitive micro-machined ultrasonic transducers (cMUTs), PMUTs do not require a vacuum gap or high DC bias voltages, resulting in simpler operation, reduced power consumption, and improved suitability for portable systems. PMUT arrays are therefore widely employed in applications such as ultrasonic imaging, proximity sensing, haptic feedback, and acoustic communication [29].

In this work, a 4×4 array of PMUTs fabricated by STMicroelectronics is employed not for ultrasonic transduction, but as a versatile MEMS platform enabling the controlled out-of-plane displacement of integrated micro-structures. In particular, the regular geometry and independent electrical addressability of the PMUT array provide a convenient starting point for implementing a grid of micromagnets, whose distance from an underlying YIG substrate can be precisely tuned by electrically actuating the individual membranes.




Figure 1.10: Microscope image of the top side of the PMUT array chip from STMicroelectronics.

2 | Experimental methods

This chapter describes the experimental techniques, instrumentation, and software tools used for device fabrication, measurement, and optimization. Section 3.2 presents the cleanroom microfabrication techniques used to realize the system, including optical lithography, magnetron sputtering, and flip-chip bonding. Subsequently, the experimental setup developed to control and characterize the device is presented in Section 2.2. In particular, the Vector Network Analyzer (VNA), which constitutes the primary instrument for spin-wave measurements, is introduced together with the main acquisition modalities adopted in this work. The hardware architecture of the setup is then detailed, including the custom-designed RF printed circuit board and the multichannel DAC (digital to analog converter) used to control the device. Finally, the software framework developed to perform the measurements and implement the training procedures is described in Section 2.3, outlining both the device control logic and the optimization algorithms employed to tune the system response.

2.1. Fabrication techniques

2.1.1. Optical lithography

Optical lithography, or photolithography, is a fundamental technique in micro and nanofabrication and constitutes a key step in cleanroom processes used to define device geometries on thin-film substrates. The method relies on the use of photosensitive polymeric materials, known as photoresists, whose chemical properties are modified upon exposure to ultraviolet (UV) radiation, enabling the transfer of a desired pattern onto the substrate surface. Prior to exposure, the substrate is spin-coated with the photoresist using optimized recipes to obtain a controlled thickness, which must be carefully selected depending on the target structure height and on whether additive or subtractive fabrication processes are required. Photoresists can be classified according to their tonality: in positive-tone resists, UV exposure weakens the polymeric chains, increasing the solubility of the exposed regions, whereas in negative-tone resists the exposure induces cross-linking reactions that

strengthen the polymeric bonds, rendering the exposed areas less soluble during development.

In this work, optical lithography was primarily employed to perform additive fabrication steps based on the lift-off technique. In a lift-off process, the patterned photoresist acts as a temporary sacrificial layer that defines the regions where material deposition is desired. After exposure and development, the substrate is left uncovered only in the areas corresponding to the target structures. A thin film is then deposited over the entire sample surface, typically by physical vapor deposition techniques. Subsequently, the sample is immersed in a dedicated solvent, known as stripper, which dissolves the photoresist and removes the film deposited on top of it, while leaving intact the material directly deposited on the substrate. For this process to be successful, a discontinuity between the deposited film on the substrate and that on the resist sidewalls is required; this condition is achieved by engineering an undercut profile in the resist.

To obtain such an undercut profile, both positive and negative-tone chemically amplified photoresists were used in combination with a lift-off resist (LOR 15A) layer. In particular, AZ5214E was employed as a positive-tone resist, while AZ15nxt was used as a negative-tone resist. A bilayer resist stack was realized by depositing the LOR directly on the substrate and spin-coating the photoresist on top. The LOR is not sensitive to UV radiation but is chemically attacked during the development step, resulting in lateral dissolution beneath the patterned resist and the formation of a well-defined undercut. While AZ5214E was processed in its standard positive-tone configuration, AZ15nxt required an additional post-exposure bake to activate the cross-linking mechanisms characteristic of negative-tone resists.




Figure 2.1: Lithography, deposition and lift-off technique with positive resist and LOR schematic.

Pattern transfer during optical lithography was performed using a maskless aligner system. In particular, a Maskless Aligner–Heidelberg MLA100 was employed to directly expose the photoresist without the use of physical photomasks. In this approach, the desired pattern is defined electronically using computer-aided design (CAD) software and transferred to the lithography system, where it is projected onto the photoresist-coated substrate by means of dynamically controlled optics and UV illumination. The maskless configuration provides a high degree of flexibility, allowing rapid design modifications and efficient prototyping, while maintaining a minimum feature size of approximately $1\ \mu\text{m}$ and an exposure speed of up to $50\ \text{mm}^2/\text{min}$.



Figure 2.2: Picture of the MLA100 maskless aligner in Polifab’s yellow room.

2.1.2. Sputtering

Magnetron sputtering is a physical vapor deposition (PVD) technique widely used for the growth of thin films, as it provides good adhesion to the substrate together with a high degree of control over the thickness, uniformity, and composition of the deposited material. The process is based on the generation of a gaseous plasma inside a high-vacuum chamber and on the subsequent acceleration of ions from this plasma towards a source material, referred to as the target. Due to momentum and energy transfer, the impinging ions erode the target surface, causing the ejection of neutral particles in the form of individual atoms or small clusters. These particles propagate through the chamber following ballistic trajectories and, if a substrate is placed in their path, condense on its surface, forming a thin film.

The plasma is typically generated by introducing an inert gas, most commonly argon, into the deposition chamber and applying a negative bias voltage to the target. Free electrons present near the target are accelerated by the electric field and collide with the argon atoms, leading to ionizing collisions and the formation of Ar^+ ions. This cascade mechanism sustains the plasma discharge. The positively charged ions are then accelerated towards the negatively biased target, striking its surface and releasing both target material and additional free electrons. In magnetron sputtering, permanent magnets are positioned behind the target to confine the free electrons within a magnetic field localized near the target surface. This configuration both limits electron bombardment of the substrate, reducing heating and potential damage, and significantly increases the ionization efficiency of the plasma by forcing electrons to follow closed trajectories, thereby enhancing the sputtering rate.




Figure 2.3: Magnetron sputtering schematic. Image taken from [31].

During this thesis work, magnetron sputtering was carried out using a Leybold LH Z400 system to fabricate supermalloy (a Nickel, Iron and Molybdenum alloy) micromagnets on top of PMUT devices. The supermalloy films were deposited as multilayer stacks, with the magnetic layers interleaved by thin chromium layers in order to improve the magnetic properties of the system [32]. Chromium was also employed as a capping layer to protect the magnetic film from oxidation. For this reason, the possibility to switch between different targets during a single deposition process was a critical feature of the sputtering system, allowing precise control over the multilayer composition without breaking vacuum.



Figure 2.4: Picture of the Leybold LH Z400 magnetron sputtering system in Polifab's deposition area.

2.1.3. Flip-chip

flip-chip, also known as controlled collapse chip connection, is an advanced interconnection and packaging technique used to electrically and mechanically connect semiconductor devices, MEMS, and other patterned chips to external circuitry using an array of conductive bumps deposited on the device's contact pads. In contrast with conventional wire bonding, in flip-chip bonding the active side of the chip is flipped face down so that its solder bumps align with corresponding pads on a circuit board or another chip, and the connections are created by reflowing the solder to form robust metallic joints with the substrate pads. The technique begins with the formation of small solder bumps on the top-side pads of the chip during its final processing; these bumps serve as the primary electrical and mechanical bridges between the chip and the mating circuitry. The chip is then placed face down and accurately aligned with the external contact pads using an optical system and micro-positioners, which allow correction of the relative positions down to tens of micrometers. Once aligned, the solder is reflowed to complete the controlled collapse chip connection, producing compact, low-inductance electrical interconnects with improved performance and reduced profile compared to traditional approaches. Because flip-chip assembly offers high pad density and direct attachment across the die area, it is particularly suited for integrated systems requiring both electrical connection and mechanical integration.

In the present work, flip-chip bonding was used to interface the YIG chip, which contains the RF antennas for spin-wave excitation, with the MEMS PMUT devices. Gold solder bumps were used as the interconnect medium, and their bonding to the opposing pads is also referred to as ball bonding due to the shape of the gold bumps. A FINEPLACE pico 1 flip-chip bonder was employed to pick up the MEMS chip with a vacuum nozzle, align it to the YIG chip using a dedicated optical alignment system, and bring the two chips into contact under controlled force and temperature to effect the ball bonding. This approach not only establishes the required electrical connections, linking the MEMS DC actuation pads to the YIG chip's bias and RF tracks, but also enables exploitation of the out-of-plane MEMS motion to modulate the spacing between the magnets and the YIG, the key degree of freedom in the experimental platform.




Figure 2.5: Picture of the flip-chip bonder available in electronics department’s cleanroom, at Politecnico di Milano.

2.1.4. White light interferometry

White light interferometry (WLI), also known as coherence scanning interferometry, is a non-contact optical technique used to measure surface topography and vertical displacements of three-dimensional microstructures with nanometer-scale resolution over height ranges extending up to several millimeters. The method relies on the short coherence length of spectrally broad (white) light: interference fringes are produced only when the optical path lengths of the reference and measurement beams are nearly identical.

In a typical WLI setup, light from a broadband source is split into a reference beam, reflected by a flat mirror, and a measurement beam, reflected by the sample surface. The two beams are recombined on a CCD camera, where interference occurs only for surface points whose optical path difference lies within approximately half the coherence length. Because real samples exhibit varying heights, interference fringes appear locally and only at specific axial positions. By vertically scanning the sample (or equivalently the optical head) and recording the fringe contrast at each pixel, the system determines the height coordinate corresponding to maximum coherence. Repeating this process over the entire field of view enables reconstruction of a full three-dimensional surface profile.

The measurements presented in this work were performed using a Filmetrics Profil3D optical profiler, which implements WLI for thin-film and MEMS characterization. In this instrument, the vertical scan is achieved by translating the optical system rather than the sample stage. The accompanying software provides tools for surface leveling, filtering, and quantitative analysis, allowing height differences between selected points to be extracted

with high precision.




Figure 2.6: Picture of the Filmetrics Profilm3D optical profilometer, available at Polifab.

This technique was used to characterize the vertical displacement of the 4×4 PMUT array under applied DC bias. The device was mounted on a custom PCB placed on the profiler stage, while external voltage sources were employed for actuation. Surface maps acquired by WLI were processed by first leveling the surrounding substrate and then measuring the height variation at the center of each membrane, providing a direct, non-contact measurement of the PMUT deflection as a function of applied voltage.

2.2. Experimental setup

The experimental setup developed for this work is built on an anti-vibration optical table and integrates RF, magnetic field generation, and electronic control subsystems to enable the characterization and operation of the spin-wave-based devices. RF measurements are performed using a four-port Rohde & Schwarz ZNA43 vector network analyzer, while the static in-plane magnetic field required for spin-wave excitation and propagation is generated by a DXFP-4015 quadrupole clamp electromagnet. The MEMS actuators embedded in the device require analog driving voltages of up to 60 V; for this reason, a 32-channel NI-6738 PCI digital-to-analog converter is employed in combination with a 32-channel analog amplifier (PD32). The entire setup is controlled by a dedicated computer through custom Python scripts, which manage the communication with the VNA, the DAC, and the current power supply of the electromagnet. The software architecture follows an object-oriented design, allowing modular control of the different instruments and facilitating future extensions or reuse for other experimental configurations. The implemented

code governs both the actuation of the MEMS and the acquisition of RF data through the VNA. Additionally, custom PCBs have been designed and printed to serve as a signal adapter between the DAC shield and the amplifier and to allow the mounting and proper wiring of the device. Finally, two training algorithms for the N1 neuromorphic device have been implemented, namely stochastic perturbation successive approximation (SPSA) and a direct search method (DS); they are discussed in more detail in the following sections.

2.2.1. Vector network analyzer

A vector network analyzer (VNA) is an instrument designed to characterize the frequency-dependent behavior of electrical networks and devices by directly measuring the propagation of high-frequency signals. The VNA generates a sinusoidal stimulus at a precisely controlled frequency and injects it into the device under test (DUT), simultaneously measuring both the amplitude and phase of the waves that are reflected back to the input port and transmitted to the output port. The response of the system is conveniently described using scattering parameters (S-parameters), which quantify the ratios of reflected and transmitted waves relative to the input. More formally, if a_i and b_i represent the complex amplitudes of the incident and reflected waves at port i , respectively, the scattering parameters are defined as:

$$S_{11} = \frac{b_1}{a_1}, \quad S_{21} = \frac{b_2}{a_1}, \quad S_{12} = \frac{b_1}{a_2}, \quad S_{22} = \frac{b_2}{a_2},$$

where S_{11} and S_{22} correspond to reflections at ports 1 and 2, and S_{21} and S_{12} correspond to transmission from one port to the other. From these definitions, it follows that the modulus of each scattering parameter represents the amplitude ratio between the output and input waves, while the complex phase of the parameter corresponds to the phase difference accumulated between the incident and reflected or transmitted waves. By sweeping the input frequency across a desired range, the VNA captures the complete frequency response of the network, providing both magnitude and phase information necessary for detailed analysis of wave propagation phenomena. In addition to frequency-domain measurements, the vector network analyzer also provides a time-domain measurement mode, in which wave quantities or scattering parameters are acquired as a function of time. In this operating mode, the VNA monitors the temporal evolution of the transmitted and reflected signals, enabling the characterization of devices whose RF response is dynamically modulated. This capability is particularly relevant for the system investigated in this work, as the MEMS-based architecture allows for time-dependent modulation of the spin-wave propagation conditions, resulting in a corresponding temporal variation of the

measured scattering parameters.

An important aspect of the VNA-based measurement is the presence of a direct electromagnetic coupling between the transmitting and receiving antennas, in addition to the desired spin-wave propagation channel. As a consequence, the signal detected at the receiving antenna is the superposition of two contributions: a prompt EM wave directly coupled through free space or the substrate and a delayed signal associated with spin-wave propagation through the magnetic medium. The electromagnetic contribution exhibits an almost frequency-independent phase, since the RF wavelength is on the order of several centimeters, much larger than the antenna separation. In contrast, spin-waves have much shorter wavelengths and therefore accumulate a strongly frequency-dependent phase during propagation. The interference between these two signals produces a characteristic modulation of the measured transmission parameter S_{21} , resulting in oscillations of its amplitude as a function of frequency, with regularly spaced maxima typically occurring every few MHz. Additionally, interference can also occur between different spin-wave modes belonging to distinct branches of the dispersion relation, which propagate with different wavelengths and group velocities. For a correct interpretation of the measurements, it is therefore essential to disentangle these various contributions and this can be achieved through a time-gating technique.

Time gating

According to linear systems and signal theory, the frequency response of a system and its impulse response form a Fourier transform pair. In particular, the complex transmission parameter $S_{21}(\omega)$ measured by the VNA represents the frequency-domain transfer function of the DUT, while the corresponding time-domain impulse response $h(t)$ is given by the inverse Fourier transform

$$h(t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} S_{21}(\omega) e^{i\omega t} d\omega.$$

Since the frequency response is acquired over a finite bandwidth, the retrieved time-domain signal is not the true impulse response of the system. Instead, it corresponds to the convolution of the actual impulse response with a sinc function (the Fourier transform of the finite frequency window). In other words, the measured response represents the system reaction to a sinc-like pulse rather than to an ideal Dirac delta excitation.

By applying the inverse Fourier transform to the measured spectra, distinct signal packets become visible in the time domain. These packets correspond to different physical con-

tributions, such as direct electromagnetic coupling and various spin-wave modes, which arrive at different times due to their different propagation velocities. Exploiting this temporal separation, a time-gating procedure can be applied directly within the VNA: a suitable time window is selected to retain only the portion of the impulse response associated with spin-wave propagation, while suppressing the other contributions. The gated time-domain signal is then transformed back into the frequency domain, yielding a filtered S_{21} spectrum in which the spin-wave contribution is effectively isolated from the noise.

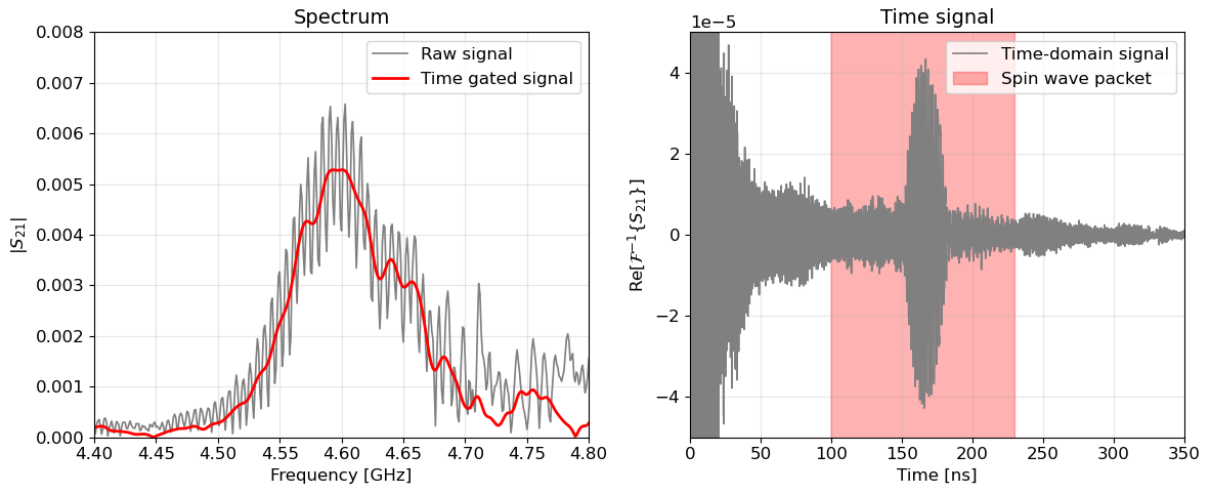


Figure 2.7: Time-gating procedure applied to a transmission measurement on N1. The raw transmission amplitude is shown in the left panel (gray curve). By inverse Fourier transforming the spectrum and considering the real part, the corresponding time-domain response is obtained (right panel), where the spin-wave packet is clearly identified and selected using the temporal window highlighted in red. Fourier transforming the gated portion back to the frequency domain yields the filtered spectrum (red curve), in which the rapid oscillations caused by interference with unwanted contributions, including directly coupled electromagnetic signals, are suppressed.

Figure 2.7 illustrates the time-gating procedure applied to a transmission measurement performed on N1. The right panel shows the time-domain response obtained via inverse Fourier transform of the measured spectrum. The spin-wave packet is identified by the red window, which defines the selected gating region. By isolating this portion of the signal and transforming it back to the frequency domain (on the left), a filtered spectrum is obtained in which the rapid oscillations caused by interference with unwanted contributions are largely suppressed. This procedure preserves the spectral envelope associated with spin-wave propagation while removing spurious components such as directly

coupled electromagnetic signals and noise. This example demonstrates how time gating enables selective analysis of the spin-wave signal, yielding a smoother spectrum that more faithfully reflects the underlying physics.

2.2.2. RF custom PCBs

To ensure a stable and reproducible platform for radio-frequency measurements and device control, a custom RF printed circuit board (PCB) was designed during this thesis work. The PCB was developed using KiCad and fabricated by an online PCB manufacturing service. The layout was specifically tailored to accommodate the hybrid YIG–MEMS sample and to enable its electrical characterization and control. The sample is mounted at the center of the PCB, where a dedicated opening was introduced to provide optical access to the YIG substrate, allowing optical investigation of spin-wave propagation to be performed without interference from the supporting structure.

The PCB provides a total of sixteen DC input lines used for MEMS actuation and control, as well as eight RF ports that route high-frequency signals from coaxial connectors to bonding pads located in close proximity to the chip. These pads are connected to the on-chip antennas and circuitry via wire bonding, minimizing parasitic effects and signal degradation. The board was designed following standard RF PCB design guidelines in order to ensure signal integrity and impedance control. It consists of a four-layer stack-up: the top layer carries the RF signal traces, implemented as coplanar transmission lines and referenced to an underlying continuous ground plane; the two inner and bottom layers are dedicated to DC routing and are largely filled with ground copper to provide shielding and reduce noise coupling. The RF traces are impedance matched and enclosed by arrays of ground vias, forming quasi-coaxial structures that suppress crosstalk and electromagnetic interference. The dielectric material used for the PCB is FR4, which represents an economical and widely available choice and offers reliable performance for RF applications up to approximately 10 GHz.

In parallel with the main PCB, a dedicated calibration board was also developed in order to properly calibrate the VNA for measurements performed on the custom platform. The calibration PCB was designed to closely replicate the RF layout, stack-up, and transmission line geometry of the main board, thereby ensuring an accurate reference plane during calibration. Standard TOSM [33] calibration structures were implemented on the board: a through, an open, a short, and a matched load. This approach allowed the removal of systematic errors introduced by cables, connectors, and PCB traces, enabling reliable and reproducible RF measurements of the devices under test.




Figure 2.8: N0 (a) and N1 (b) mounted on their custom designed PCBs.

2.3. Software

The developed software allows the vector network analyzer to perform simultaneous measurements of selected scattering parameters, while providing the possibility to choose the measurement modality, either frequency-sweep mode or time-domain mode. In addition, several acquisition parameters can be configured, including start and stop frequencies, number of sampled points, intermediate frequency bandwidth, and the optional application of the time-gating procedure. In parallel with RF acquisition, the MEMS actuation voltages can be set to fixed values or dynamically modulated using sinusoidal or square-wave waveforms. A triggering mechanism between the DAC and the VNA has been implemented to synchronize the voltage generation with the VNA time-domain measurements, enabling time-resolved characterization of the device under controlled actuation conditions. This integrated hardware–software platform has ultimately been exploited to implement and run two training algorithms on the N1 neuromorphic device, which are better explained in the following sections.

A training algorithm is employed to tune the device so as to perform a prescribed task, such as maximizing the signal received by a specific output antenna. To enable this process, a quantitative metric must be defined to assess how well a given device configuration fulfills the desired task. This is the role of the objective function, a scalar function that maps the experimentally measured parameters of the device onto a single numerical value. By reducing the system response to a scalar quantity, the comparison between different configurations becomes straightforward, allowing the identification of the most effective operating point. On these premises, the goal of a training algorithm is to explore the device parameter space and, starting from an initial configuration, identify the set of control

parameters that maximizes the objective function within a finite number of iterations. In the present work, the control parameters correspond to the actuation voltages applied to the MEMS elements of the N1 device, while the objective function takes as inputs the measured scattering parameters associated with the three output antennas. These quantities are combined into a single scalar value, expressed in arbitrary units, which represents the performance of the device for the selected task. At each iteration of the training process, the software updates the MEMS configuration, triggers the acquisition of the relevant scattering parameters through the VNA, evaluates the objective function, and determines the next set of control parameters according to the selected optimization strategy.

2.3.1. SPSA algorithm

The Simultaneous Perturbation Stochastic Approximation (SPSA) algorithm is a recursive optimization method designed for multivariate problems in which analytical gradients are unavailable or prohibitively expensive to compute. SPSA is particularly well suited for complex experimental systems, where the objective function can only be accessed through noisy measurements. Its defining feature is an efficient gradient approximation scheme that requires only two evaluations of the objective function per iteration, independent of the dimensionality p of the parameter space. This represents a significant advantage over classical finite-difference stochastic approximation methods, which require $2p$ measurements per iteration and therefore scale poorly with increasing system complexity.

At iteration k , the algorithm starts from a current estimate of the parameter vector $\hat{\mathbf{u}}^k$ and generates a random perturbation vector

$$\Delta^k = \left[\Delta_1^k, \Delta_2^k, \dots, \Delta_p^k \right],$$

whose components are independent and identically distributed, typically following a symmetric Bernoulli distribution

$$\Delta_i^k \in \{+1, -1\}.$$

Two evaluations of the objective function $= O(\mathbf{u})$ are then performed at symmetrically perturbed points:

$$y_k^+ = O(\hat{\mathbf{u}}^k + c_k \Delta^k), \quad y_k^- = O(\hat{\mathbf{u}}^k - c_k \Delta^k),$$

where c_k is the perturbation gain at iteration k .

The gradient of the objective function is approximated simultaneously for all components

using the difference between these two measurements. The i -th component of the gradient estimate is given by

$$\hat{g}_{k,i}(\hat{\mathbf{u}}^k) = \frac{y_k^+ - y_k^-}{2c_k\Delta_i^k},$$

which highlights the key property of SPSA: the same pair of objective function evaluations is reused to estimate all components of the gradient vector $\hat{\mathbf{g}}^k$. The parameter vector is then updated according to the stochastic approximation rule

$$\hat{\mathbf{u}}^{k+1} = \hat{\mathbf{u}}^k + a_k \hat{\mathbf{g}}^k(\hat{\mathbf{u}}^k),$$

where a_k is the step-size gain controlling the magnitude of the update.

The gain sequences a_k and c_k are defined as

$$a_k = \frac{a}{(k + A + 1)^\alpha}, \quad c_k = \frac{c}{(k + 1)^\gamma},$$

with a , c , A , α , and γ being user-defined non-negative coefficients. These parameters regulate the convergence behavior of the algorithm: a_k determines the learning rate, c_k sets the perturbation amplitude, and the exponents α and γ control the decay of these gains over time[34][35]. Appropriate tuning of these parameters ensures stable convergence in the presence of noise, while preventing large initial updates that could drive the system outside its physically admissible operating range.

The combination of a low measurement cost, robustness to noisy objective functions, and scalability to high-dimensional parameter spaces makes SPSA particularly well suited for the present optimization task, under the assumption that the objective function landscape is relatively smooth and does not exhibit a large number of competing local maxima. To complement this approach, a second optimization strategy based on direct search has also been implemented. While direct search methods are less susceptible to trapping in local maxima and do not rely on gradient approximations, they typically require a significantly larger number of objective function evaluations, resulting in longer optimization times compared to SPSA.




Figure 2.9: A pictorial representation of the SPSA algorithm. Image taken from [35].

2.3.2. DS algorithm

The Direct Search (DS) algorithm is a gradient-free optimization strategy adapted from the method introduced in [15]. This approach is particularly suitable for systems with discrete control parameters and complex, potentially non-convex objective landscapes. In this method, the device configuration is described by a vector of parameters \mathbf{V} where each element corresponds to a controllable voltage or setting that can be adjusted to optimize the system performance. Each parameter is constrained to a discrete set of allowed values

$$S = \{v_1, v_2, \dots, v_k\}.$$

The algorithm begins with an initial configuration \mathbf{V}_0 , which can either be selected randomly or chosen deterministically based on prior knowledge. The corresponding objective function

$$O_0 = O(\mathbf{V}_0)$$

is evaluated to establish a starting point. The optimization proceeds parameter-by-parameter: for each one of these, a trial value is randomly selected from the allowed set (excluding its current value) to create a trial configuration $\mathbf{V}_{\text{trial}}$. The objective function

$$O_{\text{trial}} = O(\mathbf{V}_{\text{trial}})$$

is evaluated, and the trial configuration is accepted if it improves the objective ($O_{\text{trial}} > O_{\text{current}}$); otherwise, the parameter remains unchanged. After all parameters have been tested, the resulting configuration defines the next iteration, and the process is repeated

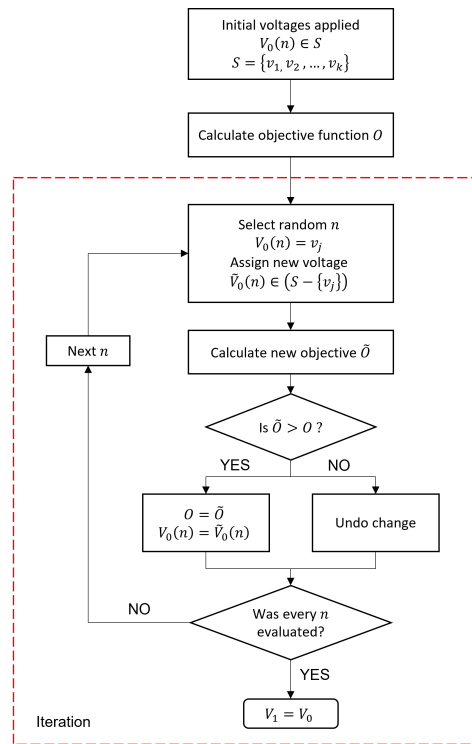


Figure 2.10: The flowchart of the implemented Direct Search algorithm.

until a convergence criterion is met, such as no further improvement in the objective or a maximum number of iterations.

The primary advantage of the DS algorithm lies in its robustness with respect to local extrema and measurement noise, as it does not rely on gradient estimates and explores the parameter space through explicit trial evaluations. This makes it particularly suitable for highly non-linear, discrete, or rugged objective landscapes, where gradient-based methods may fail or become trapped in local maxima. However, this robustness comes at the cost of a substantially higher number of objective function evaluations, since each parameter is tested individually and each trial requires a full objective function evaluation. As a result, the DS algorithm is generally slower than SPSA and less scalable to very high-dimensional parameter spaces, but provides a valuable complementary optimization strategy when reliability and global exploration are prioritized over convergence speed. Additionally, unlike SPSA, it doesn't require fine tuning to precisely converge on the objective function's maximum.

3 | Device design and fabrication

This chapter presents the design considerations and fabrication processes used to realize the two devices developed in this thesis: the N0 characterization platform and the N1 reconfigurable processor. Section 3.1 describes the design of the two devices, highlighting their geometrical layout, constraints and functional objectives. Section 3.2 details the fabrication process common to both devices, including optical lithography, sputtering and lift-off procedures for material deposition, and the flip-chip assembly used to integrate the MEMS and magnonic components. The chapter concludes with Section 3.3, which presents the fabrication results and provides important information to understand the subsequent experimental characterization.

3.1. Design

Two different devices were designed and fabricated in collaboration with the partners of the M&MEMS consortium [36]. The first device, N0, was conceived as a test platform to investigate the influence of micromagnet geometry (shape and size) on spin-wave propagation, while simultaneously enabling the characterization of antenna excitation efficiency and long-distance spin-wave transmission in YIG. The RF layout of N0 consists of eight pairs of coplanar waveguide (CPW) antennas, each associated with one MEMS actuator located in the two outermost rows of the 4×4 array. This configuration was dictated by RF routing constraints, which limited the number of addressable antennas; as a result, only eight micromagnets could be experimentally tested, despite the presence of sixteen actuators in the array.

The second device, N1, represents a first prototype of a spin-wave-based neuromorphic processor. It features a single input antenna implemented as a microstrip line spanning the full width of the chip, enabling uniform excitation of spin-waves across the device; on the opposite side, four compact CPW output antennas are equally spaced. Between the input and output antennas, a 4×4 array of micromagnets integrated on a PMUT chip provides the tunable source of the magnetic perturbation, allowing control over the spin-wave propagation dynamics.

The design and fabrication of the devices were distributed among the project partners: while Pázmány Peter Catholic University in Budapest provided the computer simulations, Technische Universität Kaiserslautern developed the RF components of the chips, i.e. the YIG substrate and the patterned antennas used for spin-wave excitation and detection. Polifab fabricated the magnetic part of the chip, carried out the hybrid assembly of the RF and MEMS parts, and provided the experimental infrastructure for device characterization. The PMUT arrays were supplied by STMicroelectronics.

For both devices, spin-wave propagation was implemented in the backward volume (BV) configuration. This choice was primarily motivated by practical considerations. BV modes require the YIG film to be magnetized in plane, which can be achieved with significantly lower magnetic fields than the out-of-plane magnetization required for forward volume (FV) modes. Moreover, the quadrupole electromagnets available in the experimental setup are designed to generate only in-plane magnetic fields, making the BV geometry the most compatible with the existing hardware. A further limitation arises from the micromagnets themselves: while the YIG film can be magnetized out of plane using moderately high magnetic fields, this is not feasible for the permalloy micromagnets due to the much higher saturation magnetization of the magnetic material. In addition, the coercive field of permalloy is very low, making it impossible to impose and maintain a stable out-of-plane remanent magnetization prior to the experiments. Achieving out-of-plane magnetization would instead require a magnetic material with both high anisotropy and high saturation magnetization, such as samarium–cobalt (SmCo). However, reliable growth of this material on MEMS membranes while preserving both the magnetic properties and the mechanical and electrical integrity of the devices has not yet been achieved. Micro-magnetic simulations further indicated promising performance in the backward volume configuration for the intended device operation. Nevertheless, the forward volume geometry remains attractive from a fundamental perspective, as it provides in-plane isotropic spin-wave propagation, which is expected to enhance magnet-induced scattering effects.

3.1.1. N0 design

As stated previously, the primary objective of the N0 device was to gain insight into how different micromagnet geometries influence spin-wave propagation. To this end, the chip layout features eight pairs of CPW antennas aligned with the outer rows of the PMUT array, each pair enclosing a magnet with distinct shape and size. Specifically, among the eight positions, two magnets are square with a side length of 200 μm , two are square with a side length of 100 μm , two are circular with a radius of 100 μm , one is circular with a radius of 200 μm , and one MEMS actuator is left without a magnet to serve as a reference.

All magnets were deposited simultaneously, so they all share the same thickness of about $1\ \mu\text{m}$.

The CPW antennas were designed to exhibit peak excitation efficiency for spin-wave wavelengths around $100\ \mu\text{m}$, a value comparable to the characteristic dimensions of the magnets, thereby maximizing their interaction with the propagating waves. Furthermore, the antennas were impedance-matched to $50\ \Omega$ in order to optimize RF power injection into the system while minimizing reflections. Additionally, 32 DC traces (two for each device) were integrated on the RF side of the chip. After the flip-chip assembly and subsequent wire-bonding to the support PCB, these traces provide a convenient routing path for delivering the DC control signals to the MEMS actuators.

The main design constraint arose from the use of prefabricated PMUT arrays provided by STMicroelectronics: while these devices offer a convenient starting point for rapid integration of micromagnets, they also impose a fixed, non-customizable geometry. As a consequence, the overall chip layout was dictated by the PMUT array architecture. The MEMS chips have a square footprint with a side length of approximately $4.3\ \text{mm}$ and host 16 circular membranes arranged in a 4×4 grid with a pitch of about $960\ \mu\text{m}$. Prior to magnet integration, the maximum out-of-plane membrane displacement was approximately $7\ \mu\text{m}$; this value decreased after fabrication of the micromagnets on top of the membranes to about $5\ \mu\text{m}$.

The RF side of the device is fabricated on a $5.6\ \mu\text{m}$ thick YIG film grown on a gadolinium gallium garnet (GGG) substrate. One side of the YIG chip was mechanically polished to achieve optical transparency, facilitating the flip-chip assembly process and enabling the investigation of spin-waves dynamics through optical techniques.




Figure 3.1: Design of the N0 device, RF side (back side) perspective.

3.1.2. N1 design

The design of the N1 device follows the same conceptual approach as N0, but with modifications tailored to its role as a neuromorphic processor. Unlike N0, all micromagnets in N1 are identical, adopting the $200\ \mu\text{m}$ square shape. On the input side, a single long stripline antenna $30\ \mu\text{m}$ wide spans the full width of the YIG chip, providing a uniform and broad source of spin-wave excitation across the entire device. On the output side, four individual CPW antennas are positioned opposite the input antenna, aligned with the rows of magnets, to collect the propagated signals from different regions of the chip. As with N0, all antennas are carefully impedance-matched to $50\ \Omega$. This is particularly important in N1 because the antennas are wire-bonded directly to a custom RF PCB, a configuration that greatly simplifies the experimental setup and allows for simultaneous signal acquisition from all output ports but introduces potential drawbacks, such as signal attenuation arising from the sub-optimal performance of the PCB traces and the wire-bond connections.




Figure 3.2: Design of the N1 device, RF side (back side) perspective.

3.2. N0 and N1 fabrication

The fabrication process is largely identical for the two devices; therefore, this section provides a general description of the main steps leading to the completed samples. The RF side of the chips was fabricated and supplied by Technische Universität Kaiserslautern. It consists of a $5.6\ \mu\text{m}$ -thick YIG film grown on a GGG substrate, which was mechanically polished on one side to ensure optical access to the YIG layer for flip-chip assembly and Brillouin light scattering measurements. Both the RF antennas and the DC routing traces were deposited simultaneously and therefore share the same material stack, consisting of SiO_2 (10 nm) / Cr (10 nm) / Cu (1 μm) / Cr (10 nm) / Au (100 nm).

At Polifab, the magnetic part of the devices was fabricated, starting from the basis provided by the 4×4 PMUT arrays. The MEMS chips were first spin-coated with photoresist and patterned via photolithography to expose the regions designated for the micromagnets. Subsequently, the magnetic multilayer stack was deposited by sputtering, and the magnets were defined through a lift-off process. After thorough electrical and mechanical characterization of the MEMS devices, the RF and MEMS chips were assembled using

flip-chip bonding, completing the fabrication process. The following subsections describe in detail the fabrication steps performed at Polifab.

3.2.1. Optical lithography

The small size, surface non-uniformity, and intrinsic fragility of the MEMS chips, which were delivered already diced, posed the main challenges during fabrication. Because the individual MEMS dies were too small and delicate to be handled directly by standard cleanroom tools, they were first mounted onto larger carrier substrates consisting of $2\text{ cm} \times 2\text{ cm}$ silicon coupons. This was achieved by spin-coating the carrier substrate with AZ15nxt resist, carefully placing the MEMS chip on top, and subsequently baking the resist. In this way, the chip was firmly bonded to the carrier and could be processed more reliably.

The primary objective of the lithography step was to coat the entire sample with photoresist while leaving openings only in the regions designated for magnet deposition. Moreover, to ensure a successful lift-off process, it was essential to obtain a pronounced undercut in the resist profile. To this end, a bilayer resist stack was employed: a first layer of LOR15A was used to facilitate undercut formation, followed by a second layer of conventional chemically amplified photoresist, either positive (AZ5214E) or negative (AZ15nxt-115cps).

The main bottleneck of the photolithography process was the spin coating of the first LOR layer. Achieving a uniform and smooth resist film across the entire sample proved challenging due to the abrupt height discontinuity introduced by mounting the MEMS chip onto the carrier substrate. The effects were particularly pronounced at the chip corners and, combined with the intrinsic surface unevenness of the piezoelectric membranes, frequently resulted in incomplete resist coverage in critical regions where the outermost membranes and DC traces are located. In several cases, air bubbles became trapped within the resist near one or more corners of the chip; these expanded during the subsequent baking step, rendering further processing impossible. In other instances, even when satisfactory coverage of both the LOR and AZ5214E layers was achieved, the resist exposure failed, revealing partially patterned features after development. This behavior was attributed to the non-uniform resist thickness, which locally exceeded the optimal focal depth of the maskless aligner, leading to insufficient exposure in thicker regions.

After numerous unsuccessful attempts and iterative optimization, a more reliable process flow was established. The final photolithography procedure adopted for devices fabrication is summarized below.




Figure 3.3: Pictures of the main problems encountered during lithography. On the left: the formation of a bubble in the LOR resist at one corner of the sample. On the right: result after development of a partially exposed feature.

The LOR 15A layer was spun at 6000 rpm, exceeding the maximum value of 4000 rpm reported in the datasheet. This modification significantly reduced the probability of air bubbles formation at the chip corners. The increased spin speed resulted only in a slight reduction of the LOR thickness, as the thickness–speed curves typically saturate at high rotation rates. The final LOR thickness was verified by SEM inspection and found to be adequate for lift-off, for which the sacrificial layer should be approximately 1.5 times thicker than the deposited metal stack.

Following spin coating, the LOR layer was baked at 190 °C for 3 minutes, after which the photoresist was applied. The selected resist was the negative-tone AZ15nxt-115cps, chosen for its higher viscosity and mechanical robustness compared to its positive counterpart. This layer was spun at 4000 rpm and soft-baked at 110 °C for 3 minutes. Patterning was performed using the MLA with an exposure dose of 400 and a defocus of 0. A post-exposure bake was then carried out at 120 °C for 1 minute, as required by the resist chemistry. Development in AZ726 MIF for 5 minutes yielded an undercut of approximately 3.5 μm , sufficient to ensure a clean lift-off of the subsequently deposited magnetic stack.




Figure 3.4: Overview and zoomed in pictures of the N1 chip after development. On the left: proper resist coverage of all the features. On the right: detail of one of the magnet features: the undercut is clearly visible.

3.2.2. Sputtering and lift-off

For the deposition of the magnets, the Leybold sputtering system was employed. The micromagnets grown on top of the MEMS are multilayered structures, consisting of alternating regions of supermalloy ($\text{Ni}_{75}\text{Fe}_{20}\text{Mo}_5$) and chromium (Cr). A 10 nm Cr layer was first deposited to serve as an adhesion promoter, followed by 80 nm of supermalloy and 5 nm of Cr forming the basic repeating unit of the stack. This multilayer sequence was repeated eleven times; finally a last supermalloy layer and 15 nm of chromium as capping were added, resulting in a total magnetic stack thickness of approximately 1030 nm. This multilayer design was chosen to prevent the formation of stripe domains in the out-of-plane direction: beyond a critical thickness, supermalloy films tend to transition from in-plane to out-of-plane anisotropy, leading to the emergence of alternating magnetic domains along the perpendicular direction, as reported by Cuccurullo et al [32]. Additionally, several breaks during deposition were needed to prevent samples overheating, that could worsen the magnetic properties of the film as well as the shape of the magnets due to resist reflow.

After deposition, the lift-off was successfully performed using MLO-07 stripper at 70 °C over a period of a few hours, yielding cleanly defined micromagnets on the MEMS membranes.




Figure 3.5: N1 pictures after lift-off. On the left: overview of the MEMS chip with the fabricated micromagnets. The bottom right membrane is missing due to fabrication issues. On the right: microscope picture of 4 membranes of the N1 chip with their patterned square magnets.

3.2.3. Flip-chip

The final step of the fabrication process consists of assembling the RF part (YIG chip) and the magnetic part (MEMS chip) into a single device. The two chips are aligned and bonded together, with one chip flipped on top of the other, as described in detail in Section 2.1.3. Small gold balls are deposited on the contacts of the MEMS chip, which is then picked up by a vacuum nozzle and brought into contact with the YIG chip. The two chips are pressed together with a force of 15 N while being heated to 400 °C for 5 minutes to allow soldering to occur.

This flip-chip step is the most critical phase of the fabrication process, as any failure at this stage would result in the loss of both chips. Several factors determine its success: the alignment of the two chips must be precise, ensuring proper registration of the antennas, the magnets, and the electrical contacts on the MEMS DC traces; any misalignment could prevent the MEMS from working. The formation of reliable electrical contacts through the gold balls is also crucial: with 32 contacts in total, variations in balls size, insufficient bonding force, or suboptimal temperature could lead to incomplete connections, particularly for the smaller bumps. Finally, the relative tilt of the chips must be carefully controlled. Any tilt can compromise electrical contact quality and cause non-uniform

magnet heights relative to the YIG surface, potentially affecting the device performance in unforeseen ways.

3.3. Fabrication results

3.3.1. N0

The fabrication of the N0 device was overall successful, despite a few imperfections arising from the challenging MEMS lithography step. In particular, the membrane intended to serve as a reference without a magnet was partially covered by magnetic material, because of bad resist coverage during lithography (Figure 3.6). For the same reason, some DC traces were unintentionally coated with a thin magnetic film during sputtering, presumably shorting them and thus preventing the MEMS actuation. These deposits were manually removed using a probe-station needle, restoring MEMS functionality. After flip-chip assembly, the membrane hosting the large circular magnet (200 μm radius) was not working, likely because of a bad electrical contact. Unfortunately, this was the only magnets featuring this design.

Despite these issues, the overall alignment between the YIG chip and the MEMS array was satisfactory, and no significant tilt was observed after bonding, indicating good mechanical assembly quality; the distance between the two chips was of about 20 μm . For characterization, the sample was mounted on a support PCB to facilitate access to the MEMS DC control lines, while the RF antennas were contacted directly using RF probes, as a custom RF PCB was not implemented for this first test device.




Figure 3.6: Microscope pictures of N0 after flip-chip. Backside view through the YIG film. On the left: overview of the bottom part of the device. On the right: zoomed view on the MEMS partially covered with magnetic material.

3.3.2. N1

The fabrication of the N1 device proved more challenging than that of N0, mainly due to the lithography issues discussed previously. The final chip selected for experimental characterization already featured one broken membrane at a corner, originating from earlier fabrication attempts. Nevertheless, this device was retained under the assumption that a corner magnet would have a limited impact on the overall spin-wave dynamics. Unfortunately, an additional MEMS element failed during post-deposition testing as a result of a defective power supply that exposed the device to an overvoltage condition. Furthermore, after the flip-chip assembly, one more MEMS was found to be non-functional, reducing the number of operational actuators to 13 out of 16 at the time of device completion. After several days of characterization, two additional MEMS elements ceased functioning. Fortunately, these were located near the output antennas, where their influence on spin-wave modulation is expected to be comparatively weaker. As a final remark, the flip-chip process resulted in a noticeable tilt between the MEMS chip and the YIG substrate: at the input side the gap was about $10\ \mu\text{m}$, while the output side exhibited a larger gap of about $25\ \mu\text{m}$. This non-uniform spacing introduces an additional asymmetry in the magnetic coupling across the device. A summary of the fabrication outcome is reported in Figure 3.7




Figure 3.7: N1 after fabrication. On the left: the chip mounted and bonded on the custom PCB. On the right: schematic view of the fabrication outcome with the broken MEMS marked with a red cross and the measured distance from the YIG film on the input and output sides.

4 | N0 - Investigation of a single tunable scatterer

This chapter investigates the physical behavior of a single tunable magnetic scatterer using the N0 device and combining micromagnetic simulations and experimental measurements. Section 4.1 presents the executed micromagnetic simulations, beginning with the analysis of the perturbation field and the magnetization distribution generated by a micromagnet above the YIG substrate, followed by a systematic study of key design parameters, including the magnet height, frequency dependence, and magnet shape, in order to understand their influence on spin-wave propagation. Section 4.2 then presents the experimental characterization of the fabricated N0 device, starting with the measurement of MEMS displacement and proceeding with the analysis of the raw and time gated RF signal. Finally, the effect of voltage-controlled MEMS actuation on spin-wave modulation is investigated and compared with the simulation results, validating the physical mechanisms and providing insight into the effect of the magnetic scatterers.

4.1. Micromagnetic simulations

The goal of the N0 device was to understand the effects of the magnets on spin-wave propagation, considering different geometries and sizes. To this purpose the measurements realized on the device have been integrated with some micromagnetic simulations performed with Mumax+ [37] on the Orfeo HPC cluster hosted in Trieste [38]. All the measurements have been performed with an in plane applied field (BV configuration) of 92 mT to keep the spin-wave signal in the 3 GHz to 4 GHz bandwidth. Of course, the simulations have been performed on the same frequency band.

The simulated system closely reproduces the N0 device configuration: it consists of a $1000\ \mu\text{m} \times 1000\ \mu\text{m}$ YIG thin-film area. Spin-waves are excited on the left side of the domain by applying a dynamic magnetic field along a straight line, with a $20\ \mu\text{m}$ thickness, spanning the full width of the simulation window. This excitation geometry is chosen to approximate the plane-wave regime as closely as possible, minimizing edge ef-

fects and lateral wavefront distortions. A static external bias magnetic field is applied in the simulations to reproduce the experimental conditions. The stray field generated by the perturbing permalloy (Py) micromagnet placed above the YIG is computed separately and then imported into the YIG film simulation as an additional static contribution.

The main assumption is that the Py micromagnet is fully saturated under the applied bias field, since the preliminary characterization conducted with a vibrating sample magnetometer (VSM) revealed a very low coercive field, of the order of 1 Oe: its stray field is therefore calculated by modeling it as a permanent magnet with the same geometry and a uniform magnetization equal to the measured saturation magnetization of 0.64 MA/m. Moreover, the ferromagnetic resonance frequency of Py at around 100 mT is significantly higher than that of YIG (approximately 8 GHz), so dynamic coupling between the spin-waves propagating in YIG and the Py magnet can be neglected. For this reason, the Py magnet is excluded from the dynamic part of the simulation, reducing the number of computational cells and significantly shortening simulation time.

To emulate different vertical positions of the magnet, the three-dimensional stray field distribution is computed once and then sliced at the desired height. The resulting two-dimensional field map is superimposed onto the static bias field applied to the YIG film. The system is driven by a sinusoidal RF excitation for a defined number of time steps. After the propagating wave reaches the opposite side of the domain and a stationary regime is established, the final oscillation periods are recorded. From these steady-state oscillations, the local amplitude and phase of the spin-wave are reconstructed at each spatial point using a lock-in-like technique. This method enables spatially resolved extraction of both intensity and phase modulation induced by the magnetic perturbation.

Simulations were performed for excitation frequencies spanning the 3 GHz to 4 GHz bandwidth in steps of 100 MHz, roughly matching the experimental frequency window. For each frequency, multiple magnet heights between 0 μm to 25 μm were considered. Two magnet geometries were investigated: a 100 μm in side square and a 100 μm in diameter circular magnet, allowing direct comparison with the experimentally studied configurations.

4.1.1. Perturbation field and magnetization

To better understand the origin of the spin-wave-magnet interaction, it is useful to analyze the perturbation induced in the YIG magnetization by the stray field of the micromagnet. Figure 4.1 shows the simulated stray field generated by a 100 μm fully saturated square magnet, together with the corresponding equilibrium magnetization configuration of the YIG film under the combined action of the external bias field and the perturbation field.

The results indicate that in regions where the transverse components of the stray field (along the y and z directions) are strongest, the local magnetization of the YIG film is slightly tilted away from its predominant alignment along the x -direction (defined by the external bias field). As a consequence, the x -component of the magnetization is locally reduced, while finite mainly out of plane components develop. This produces a small but spatially structured inhomogeneity in the magnetic landscape of the film.




Figure 4.1: (a) Simulated stray field distribution on the YIG substrate generated by a $100\ \mu\text{m}$ -side square permalloy magnet positioned $15\ \mu\text{m}$ above the film. (b) Corresponding equilibrium magnetization configuration in the YIG layer, showing the perturbation induced by the magnet stray field.

Notably, the perturbation exhibits a characteristic two-lobe structure elongated in the y -direction, reflecting the symmetry of the magnet stray field distribution. Although the deviation from uniform magnetization is small in magnitude, this localized inhomogeneity acts as an effective scattering center for propagating spin-waves. This behavior is illustrated in Figure 4.2, which presents the result of the interaction of spin-waves in the planar wave regime, excited on the left, and a $100\ \mu\text{m}$ side Py scatterer, placed in the center of a simulation at $3.2\ \text{GHz}$ with a magnet–film distance of $5\ \mu\text{m}$, chosen to enhance the interaction.

In panel A, the spatial distribution of the dynamic magnetization component $m_z(x, y, t)$ is shown at the final time step of the simulation, once a stationary regime has been reached. Panels B and C show the reconstructed amplitude and phase of the propagating wave, respectively. These quantities were extracted from the time-dependent simulation data using a lock-in-like technique, explained in the following paragraph.

Assuming a harmonic response at angular frequency ω , the local magnetization oscillation can be written as

$$m_z(x, y, t) = I(x, y) \cos(\omega t) + Q(x, y) \sin(\omega t), \quad (4.1)$$

where $I(x, y)$ and $Q(x, y)$ are the in-phase and quadrature components, respectively. These components can be obtained by multiplying the sampled signal by the corresponding reference functions and averaging over an integer number T of excitation periods:

$$I(x, y) = \frac{2}{T} \int_{t_0}^{t_0+T} m_z(x, y, t) \cos(\omega t) dt, \quad (4.2)$$

$$Q(x, y) = \frac{2}{T} \int_{t_0}^{t_0+T} m_z(x, y, t) \sin(\omega t) dt. \quad (4.3)$$

In practice, these integrals are approximated by discrete sums using eight samples per period collected over the last ten periods of the simulation, ensuring accurate reconstruction of the steady-state response.

From the in-phase and quadrature components, the local amplitude and phase are reconstructed as

$$A(x, y) = \sqrt{I^2(x, y) + Q^2(x, y)}, \quad (4.4)$$

$$\phi(x, y) = \arctan\left(-\frac{I(x, y)}{Q(x, y)}\right). \quad (4.5)$$

For clarity, the amplitude is presented as the relative variation with respect to the unperturbed case (no magnet),

$$\Delta A(x, y) = \frac{A(x, y) - A_{\text{no magnet}}(x, y)}{A_{\text{no magnet}}(x, y)}, \quad (4.6)$$

while the phase is shown as the phase shift relative to the unperturbed propagation,

$$\Delta\phi(x, y) = \phi(x, y) - \phi_{\text{no magnet}}(x, y). \quad (4.7)$$

This representation highlights the spatial regions where the magnet induces amplitude attenuation or enhancement, as well as localized phase distortions and it is adopted in all the following figures.




Figure 4.2: Results of a micromagnetic simulation performed with a 100 μm side square magnet positioned 5 μm above the YIG film, with an excitation frequency of 3.2 GHz. Panel A shows the z-component of the reduced magnetization at the last frame of the simulation. Panels B and C respectively show the spatial map of the percentage amplitude variation and the phase shift of the spin-wave induced by the presence of the micromagnet.

4.1.2. Magnet height effect

Crucially, the height of the magnet above the YIG substrate primarily affects the magnitude of the induced perturbation, while leaving its lateral spatial extent largely unchanged. Consequently, the ratio between the spin-wave wavelength and the characteristic size of the perturbation remains fixed. This ensures that, for a given excitation frequency, the effect of the magnet can be continuously tuned in strength without significant changes in the mechanism of scattering which is supposed to be largely influenced by the ratio between the wavelength and the size of the scatterer. This behavior is clearly illustrated in Figure 4.3: panels B, E, H and C, F, I show, respectively, color maps of the percentage amplitude and phase variations induced by the magnet proximity with respect to the unperturbed case, for different values of the magnet height; in panels J and K the same values are plotted along a virtual antenna (indicated by the dashed lines) positioned in the y-direction at a distance of 260 μm from the magnet.

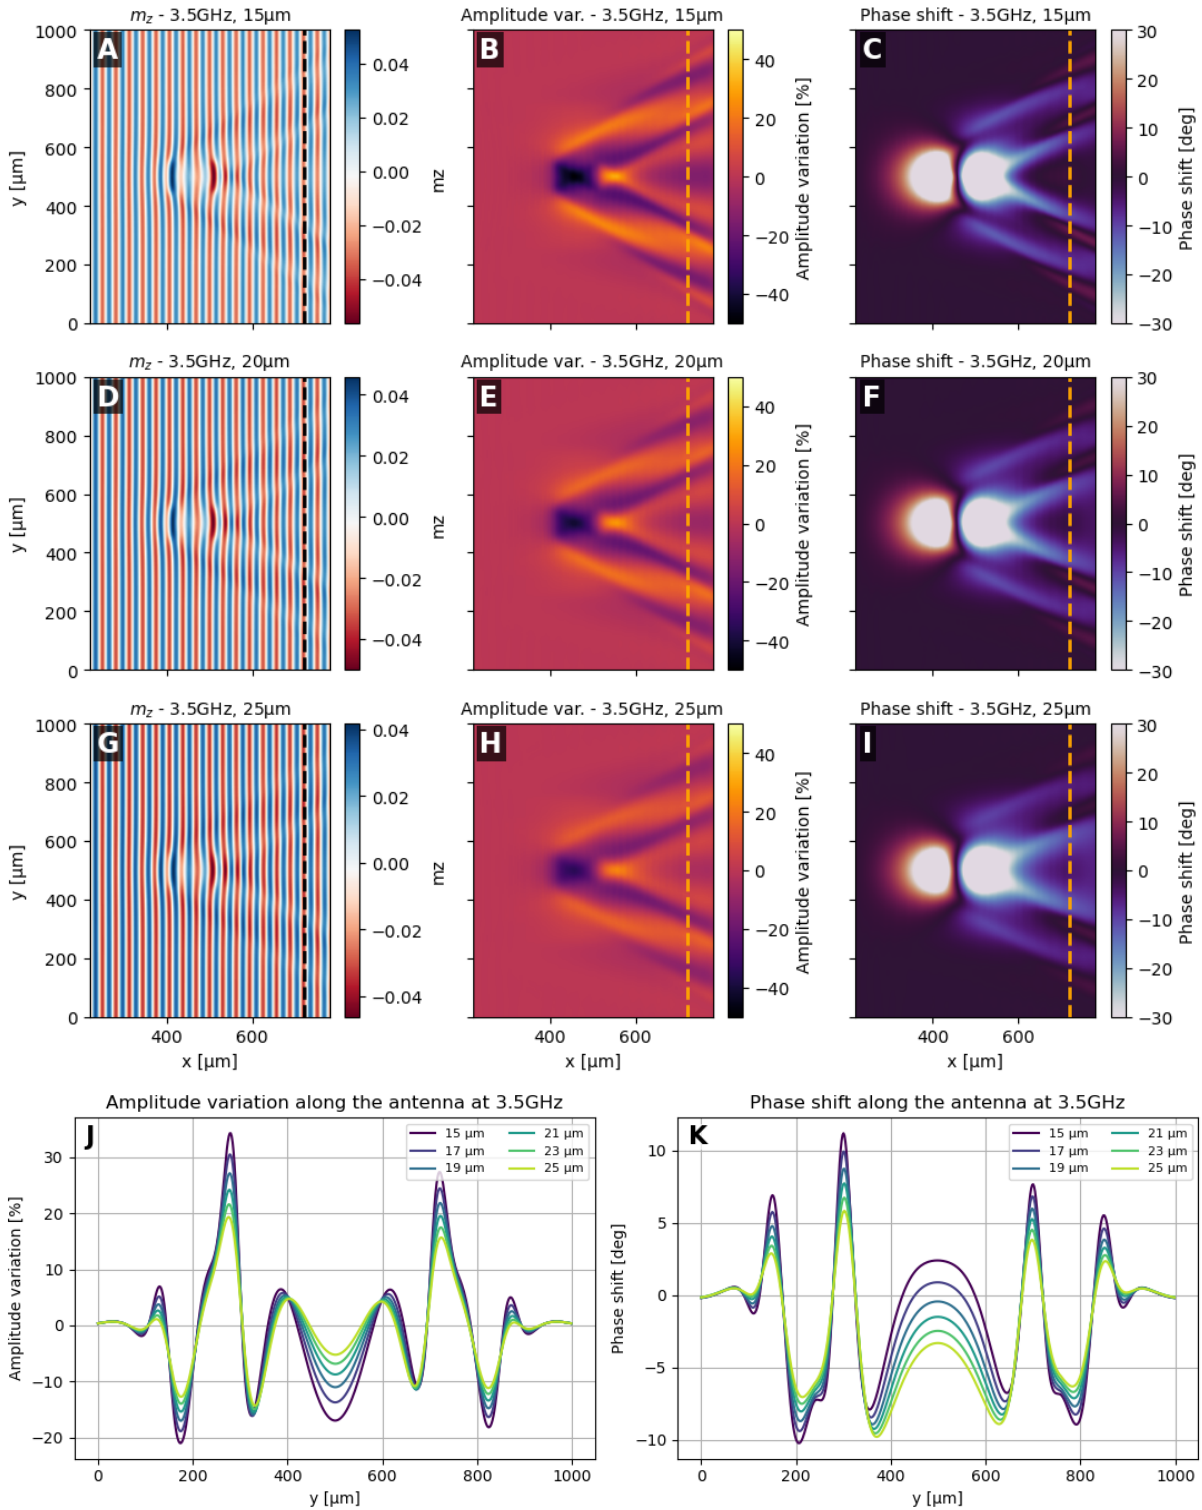


Figure 4.3: Effect of magnet height on spin-waves propagation. Panels A, D, G show the z-component of the reduced magnetization at the last simulation frame; panels B, E, H and C, F, I show, respectively, 2D maps of the percentage amplitude variations and phase shifts at 3.5 GHz for different magnet heights. Panels J and K show the corresponding amplitude and phase variations extracted along a virtual antenna (dashed lines) at a distance of 260 μm from the magnet.

A regular and monotonic trend is visible in both the intensity of the received signal and the associated phase shift as the magnet is progressively lifted away from the YIG substrate. At the same time, the global interference and scattering pattern preserves its overall spatial structure, with only its contrast being reduced. This further confirms that varying the magnet height primarily modulates the strength of the interaction, while leaving its characteristic spatial features essentially unchanged. In contrast, for the range of heights considered in the simulations (roughly corresponding to the experimental range of distances achievable from applied voltages up to 60 V) the phase variation at the center of the antenna is just on the order of 5° .

4.1.3. Frequency dependency

The frequency-dependent behavior of the system is analyzed by varying the excitation frequency. This allows the generation of spin-waves with wavelengths ranging approximately from $10\ \mu\text{m}$ to $200\ \mu\text{m}$, thereby modifying the ratio between the wavelength and the characteristic size of the magnetic perturbation.

In Figure 4.4, it is reported the case of three characteristic excitation frequencies in the spin-wave band, namely 3.0 GHz, 3.5 GHz and 4.0 GHz, corresponding to wavelengths of $13\ \mu\text{m}$, $69\ \mu\text{m}$ and $185\ \mu\text{m}$, as resulting from the dispersion relation and from the spatial periodicity observed in the unperturbed region of the simulations. As clearly visible in panels A through I, when the wavelength becomes comparable to the $100\ \mu\text{m}$ lateral dimension of the square magnet, the projected *shadow* region broadens significantly. This alters the spatial interference pattern and leads to a pronounced modulation of the amplitude detected at the receiving antenna. Additionally, the stronger interaction with the magnet at longer wavelengths, causes part of the incident wave to be reflected, leading to the formation of a standing-wave pattern between the input antenna and the magnet, as clearly visible in panel H, clearly showing the alternating nodes and antinodes in the intensity profile. This effect is further enhanced as the magnet is moved closer to the substrate.

It is also worth noting that the most substantial phase modulations occur in relatively narrow lateral regions at the sides of the antenna. In contrast, particularly for magnet-film distances exceeding $15\ \mu\text{m}$, the central portion of the antenna, corresponding to the effective sensing region in the fabricated devices, experiences predominantly amplitude modulation, with comparatively weaker phase variations.

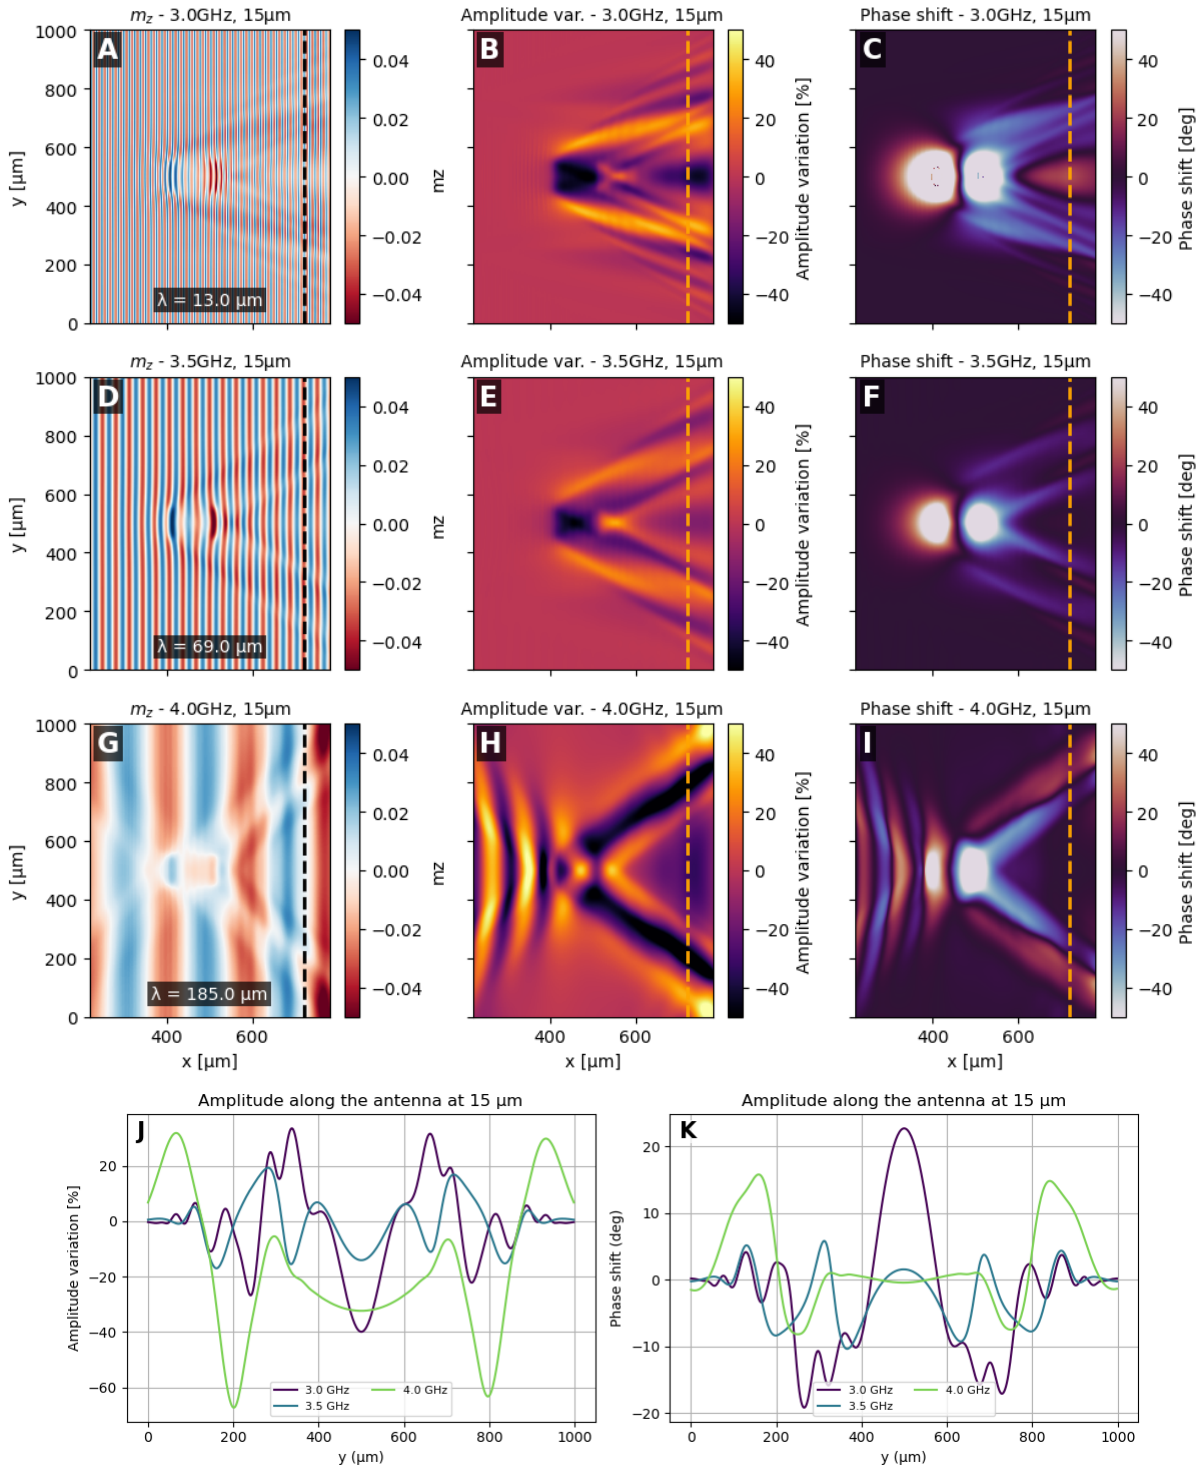


Figure 4.4: Effect of excitation frequency on spin-waves propagation. Panels A, D, G show the z -component of the reduced magnetization at the last simulation frame varying the frequency and keeping the magnet–YIG separation fixed; panels B, E, H and C, F, I show, respectively, 2D maps of the percentage amplitude variations and phase shifts with a magnet–YIG distance of $15 \mu\text{m}$ for different frequencies. Panels J and K show the corresponding amplitude and phase variations extracted along a virtual antenna (dashed lines) at a distance of $260 \mu\text{m}$ from the magnet.

In all the simulations, the formation of two narrow beams originating from the lobes of the perturbation is observed. These features can be attributed to the excitation of spin-wave caustic beams [39], which arise as a direct consequence of the strongly anisotropic dispersion relation of spin-waves in the YIG film: in the current configuration, the nature of the spin-wave depends on the orientation of the wave vector with respect to the bias magnetic field: when the wave vector is parallel to the bias field, backward volume spin-waves are excited, whereas when it is perpendicular, Damon–Eshbach modes are generated. Between these two limiting cases, the dispersion relation is highly anisotropic, and there exist specific directions along which the angle of the group velocity becomes stationary with respect to the wave vector. Under these conditions, a broad range of wave vectors propagates along nearly the same direction, concentrating the energy into narrow, intense, and weakly diffracting beams known as caustics.

In homogeneous films, caustic beams are generally excited by an omnidirectional source capable of generating a wide angular spectrum of wave vectors; this role can be fulfilled by a point like scatterer or by a narrow slit that bends the wave fronts [40]. In the present case, the localized perturbation of the magnetization induced by the magnet acts as an effective scattering center and redistributes the wave vectors into the caustic directions defined by the anisotropic dispersion.

4.1.4. Magnet shape

To investigate the role of magnet geometry in the spin-wave interaction, simulations were performed comparing a 100 μm -side square magnet with a 100 μm -diameter circular magnet. The results show a strong correspondence between the two geometries across all investigated frequencies and magnet–film separations within the 20 μm to 25 μm , which matches the operational range of the fabricated device.

Figure 4.5 illustrates this comparison for an excitation frequency of 3.5 GHz and a magnet height of 20 μm . Panels A, B, C illustrate the effect of the square magnet, showing respectively the last frame of the simulation, the amplitude percentage variation and the phase shift; panels D, E, F do the same for the round magnet. Finally, panels G and H depict the amplitude and phase variations extracted at the receiving antenna for the same frequency and varying magnet heights, in the case of the round magnet only. The similarity between the two geometries is evident, particularly when comparing these results with the corresponding data previously shown for the square magnet.

This indicates that, within the considered parameter range, the interaction is primarily governed by the lateral scale of the perturbation rather than its precise shape: as long as

the characteristic size of the magnet is comparable, the resulting scattering, diffraction, and phase modulation effects remain similar.

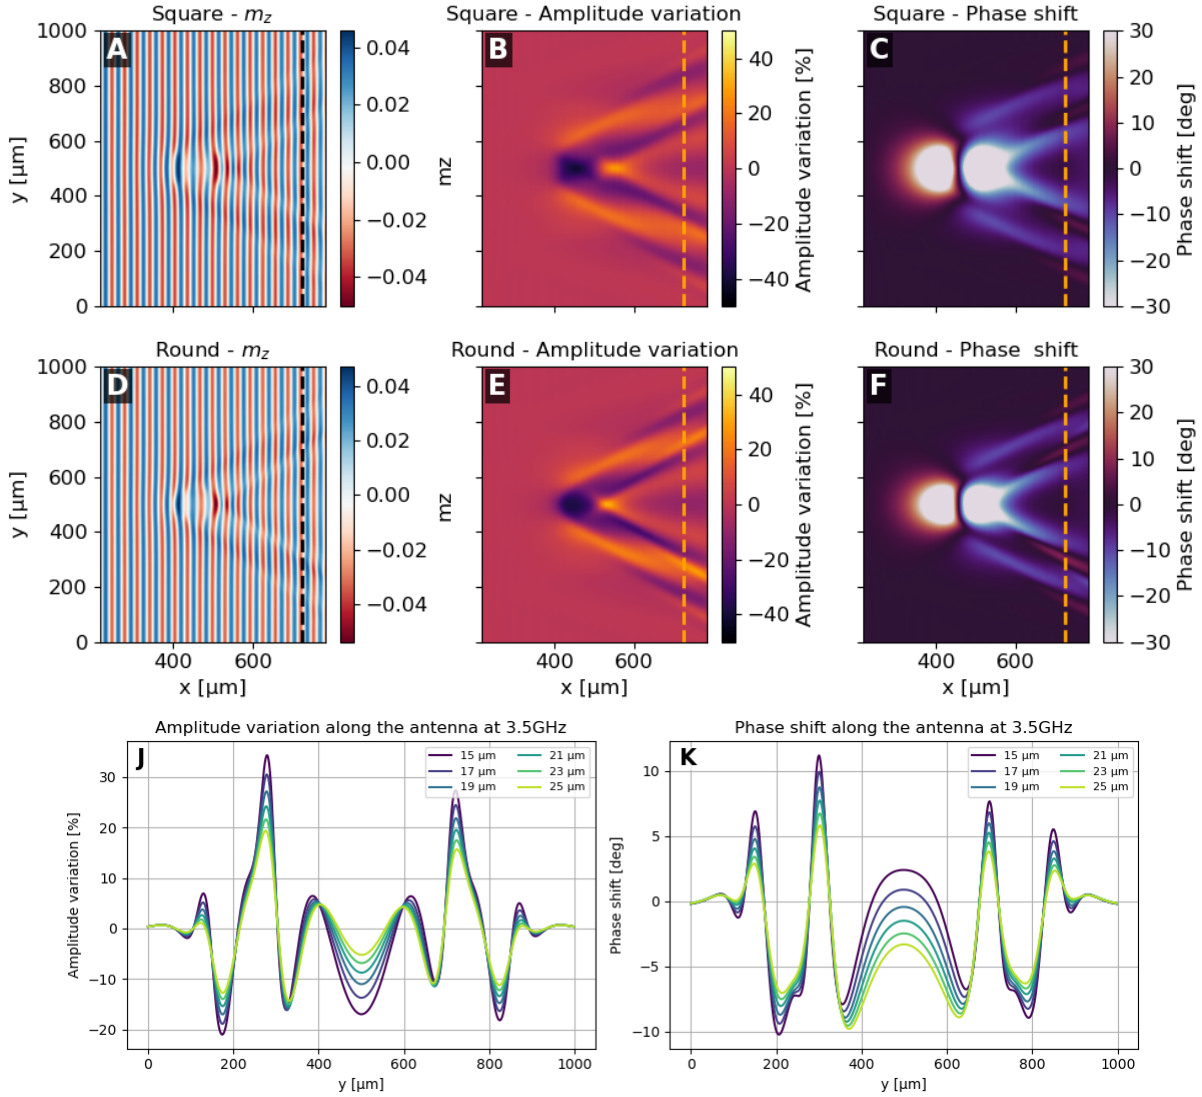


Figure 4.5: Comparison of the effect of magnet geometry on spin-waves propagation for a magnet–substrate distance of 20 μm and an excitation frequency of 3.5GHz. Panels A and D show the z-component of the reduced magnetization at the final simulation frame for square and circular magnets, respectively, both with a lateral size of 100 μm . Panels B and C display the corresponding two-dimensional maps of the percentage amplitude variation and phase shift for the square magnet, while panels E and F show the same quantities for the circular magnet. Panels G and H present the simulated amplitude and phase variations along the receiving antenna for the circular magnet case only, allowing direct comparison with Figure 4.3, panels J and K.

4.2. N0 characterization

4.2.1. MEMS displacement

The voltage-displacement characteristic of the MEMS membranes was reconstructed using an optical profilometer on the fully assembled devices. The vertical displacement, relative to the 0 V resting position, was measured from the backside of the MEMS chip for each functioning membrane in the array. Although these measurements were performed on the N1 device, they can also be considered representative of N0, since both devices were fabricated using nearly identical processes and the profilometry technique provides only moderate accuracy.

The measured heights, initially referenced to the instrument frame, have been shifted in Figure 4.6 to express the relative displacement of each membrane with respect to its resting position at 0 V. The figure shows that the total achievable displacement ranges between approximately 4 μm to 5 μm . Notably, an applied voltage of 30 V is sufficient to reach about 80% of the maximum displacement, indicating a nonlinear response with a tendency toward saturation at higher voltages.




Figure 4.6: Measured MEMS displacement as a function of the applied voltage for the N1 device. Only data from the functioning MEMS membranes are shown, together with a global fit. The maximum displacement is on the order of 5 μm , and the response begins to saturate at approximately 30 V.

4.2.2. Raw signal analysis

The first characterization of the device was performed without applying any voltage to the MEMS. The raw transmission spectra were acquired over a wide frequency range, with a nominal applied magnetic field of 92 mT. The amplitude of the transmission scattering parameter for one of the magnets of N0 is shown in Figure 4.7, together with the calculated combined excitation–detection efficiency of the coplanar waveguide antennas integrated on the device. The background contribution to the signal was removed by subtracting a reference spectrum acquired under an applied field of 165 mT, where spin-wave excitation is suppressed in the investigated frequency range.




Figure 4.7: Reference-subtracted transmission spectrum measured for one of the N0 magnets at a nominal applied field of 92 mT, compared with the calculated excitation efficiency of the coplanar waveguide antenna. The efficiency is expressed as a function of frequency using the fundamental mode dispersion at an effective field of 85 mT and reproduces the main and satellite peaks in the measured spectrum. The raw spectrum exhibits continuous oscillations due to interference between spin-wave modes and the directly coupled electromagnetic signal; the *peaks* correspond to the regions where the oscillations amplitude is enhanced.

The excitation efficiency of the antenna is usually defined as a function of the wave vector k , however, since the wave vector k and the frequency f are related through the dispersion relation, it can be equivalently expressed as a function of frequency by using the dispersion of the fundamental mode to map k onto f . This representation

allows the excitation efficiency profile to be overlaid onto the measured transmission spectrum, enabling a direct comparison between the calculated antenna response and the experimentally observed spectral features.

Importantly, the dispersion relation depends on the externally applied magnetic field, and therefore the bias magnetic field can be treated as a free parameter in the calculation of the excitation efficiency, allowing its profile to be shifted along the frequency axis to achieve optimal alignment with the experimental spectrum. This procedure is justified by the limited accuracy in the control of the applied magnetic field in the experimental setup, as well as by possible spatial non-uniformities of the field across the sample.

In the present case, the best agreement between the calculated excitation efficiency and the measured spectrum is obtained for an effective field of 85 mT, which is reasonably close to the nominal experimental value. It is important to note that the raw transmission spectrum does not show clearly defined peaks, but rather continuous oscillations resulting from the interference between different spin-wave modes (spectral beating) and the signal coming from the direct electromagnetic coupling between the antennas. Consequently, the peaks of the envelope of the measured spectrum correspond to regions where the amplitude of these oscillations increases, reflecting an enhanced spin-wave signal intensity. With the effective field adjustment, the main peaks and several satellite peaks in the transmission spectrum are well reproduced by the calculated excitation efficiency profile. However, the additional peak at higher frequencies, around 4.2 GHz, cannot be explained solely by the antenna excitation efficiency; a plausible explanation for this feature will be provided later based on the time-gating analysis.

The primary tool for understanding the spectral features is the analysis of the dispersion relations and their corresponding group velocities. Figure 4.8 reports the first three modes of the calculated dispersion relation (left panel) at an applied field of 85 mT, and their corresponding group velocities (right panel), plotted as a function of the wave vector (lower axis) and of the wavelength (upper axis); the profile of the excitation efficiency is also reported in both graphs.

As expected from the antenna design, the excitation efficiency has its maximum at a wavelength of 100 μm , while the wavelengths associated with the main excitation peak span approximately the 50 μm to 200 μm range, corresponding to values comparable with the characteristic dimensions of the magnets. Within this range, the fundamental mode ($n = 1$) exhibits the highest group velocity compared to the higher-order modes. As a consequence, spin-wave packets associated with the fundamental mode are expected to reach the detection antenna earlier in the time-domain response and with a stronger

amplitude, since slower propagation is typically associated with increased attenuation.

Another important aspect for the subsequent discussion that can be noted here is that the group velocity of the fundamental mode increases as the wave vector decreases, corresponding to higher frequencies in the dispersion relation. This behavior has direct implications for the interpretation of both the measured transmission spectra and the time-resolved signals, as it affects the relative arrival times and attenuation of the propagating spin-wave packets.




Figure 4.8: Calculated dispersion relation (left) and corresponding group velocities (right) for the first three volume modes at an applied field of 85 mT. The lower axes show the wave vector k , while the upper axes indicate the corresponding wavelength λ . The excitation efficiency of the coplanar waveguide antenna is overlaid in both panels as a function of k and λ .

4.2.3. Time gating

Although the overall spectral profile is generally consistent with the calculated antenna excitation efficiency and the fundamental branch of the dispersion relation, the measured spectra exhibit a more complex structure, indicating the presence of multiple overlapping contributions. One such contribution originates from the tapered geometry of the antennas, which is introduced to electrically connect the narrow antenna section to the larger contact pads. As shown in Figure 4.9, this tapering results in a gradual widening of the signal, the ground lines and their gap toward the edges of the sample, effectively enlarging the antenna footprint in real space. In reciprocal space, this corresponds to an excitation efficiency shifted toward lower wave vectors k and characterized by a different spectral profile compared to that of the nominal antenna section. As a result, the tapered regions act as an additional excitation source, capable of generating spin-waves

with lower k values and thus, according to the dispersion relation of the fundamental mode, higher frequencies. This means that the tapered sections can excite spin-waves in a frequency range extending beyond that of the main antenna, explaining the origin of the peak at 4.2 GHz shown in Figure 4.7. Furthermore, each excitation region, including both the nominal antenna and the tapered sections, can couple to multiple quantized volume modes of the magnetic structure. In addition to these effects, reflections occurring at the antennas and at the edges of the YIG film introduce further contributions to the detected signal. When the system is probed at a fixed frequency, spin-waves associated to all of these contributions propagate with different group velocities and accumulate distinct phase shifts. On the detection antenna, all of these components interfere in a complex, frequency dependent pattern that gives rise to the intricate morphology of the raw transmission spectra.




Figure 4.9: Optical microscope image of the two antennas on N0 enclosing a 200 μm magnet. The tapered sections of the antennas progressively widen toward the contact pads, covering a region where the magnet cannot significantly affect spin-wave propagation.

To disentangle these contributions, a time-gating analysis was performed using the impulse response of the system. Since the different packets generally propagate at different speeds, they arrive at the receiving antenna at distinct times and can therefore be separated in the time domain. Figure 2.7 illustrates this procedure for the four geometries implemented on the device. The lower panels show the time-domain signals with several gating windows highlighted, while the upper panels display the corresponding gated spectra compared with the raw spectra.

The first window, highlighted in red and spanning approximately 0.7 ns to 11 ns, isolates

the fastest propagating packet, originating from the tapered antenna region exciting the first volume mode. Consistent with the dispersion relation, this contribution is associated with smaller k values and higher group velocities, and its spectral content is primarily located in the higher-frequency region, with a shape differing from the expected excitation efficiency of the main antenna. Its very early arrival time, considering its predicted group velocity, is likely due to the tapering of the input and output antennas that effectively shifts the excitation and detection regions closer to each other, reducing the propagation distance and consequently the measured time of flight.

The second window, highlighted in green and spanning 13 ns to 50 ns, primarily contains the contribution of the first volume mode excited by the proper antenna section. In the corresponding spectrum, the excitation efficiency profile of the main antenna is clearly visible. However, additional components at higher frequencies are also present and can be attributed to spin-waves excited by the tapered sections of the antenna. As a result, a complete temporal isolation of the fundamental mode excited exclusively by the proper antenna is not possible, and the measured spectrum reflects a superposition of the main antenna response and higher-frequency components originating from the tapered geometry.

The third window, highlighted in yellow and spanning 51 ns to 82 ns, isolates a packet whose arrival time cannot be attributed to the much lower group velocity of a second volume mode. Moreover, its spectral content shows a significant overlap with the first mode associated with the proper antenna section (green spectrum). A plausible explanation is that this packet originates from two successive reflections of the fundamental mode, first at the receiving antenna and subsequently at the input antenna, before being detected again. This additional propagation path effectively increases the total travel distance, resulting in a longer time of flight. If this is the case, the packet should arrive after three times the delay of the directly transmitted one (observed in the green window), which is consistent with the measured arrival times of approximately 20 ns and 60 ns, respectively.

Finally, the direct electromagnetic coupling between the antennas is not explicitly shown, as it occurs within the first 0.7 ns but it has been excluded from the first gating window.

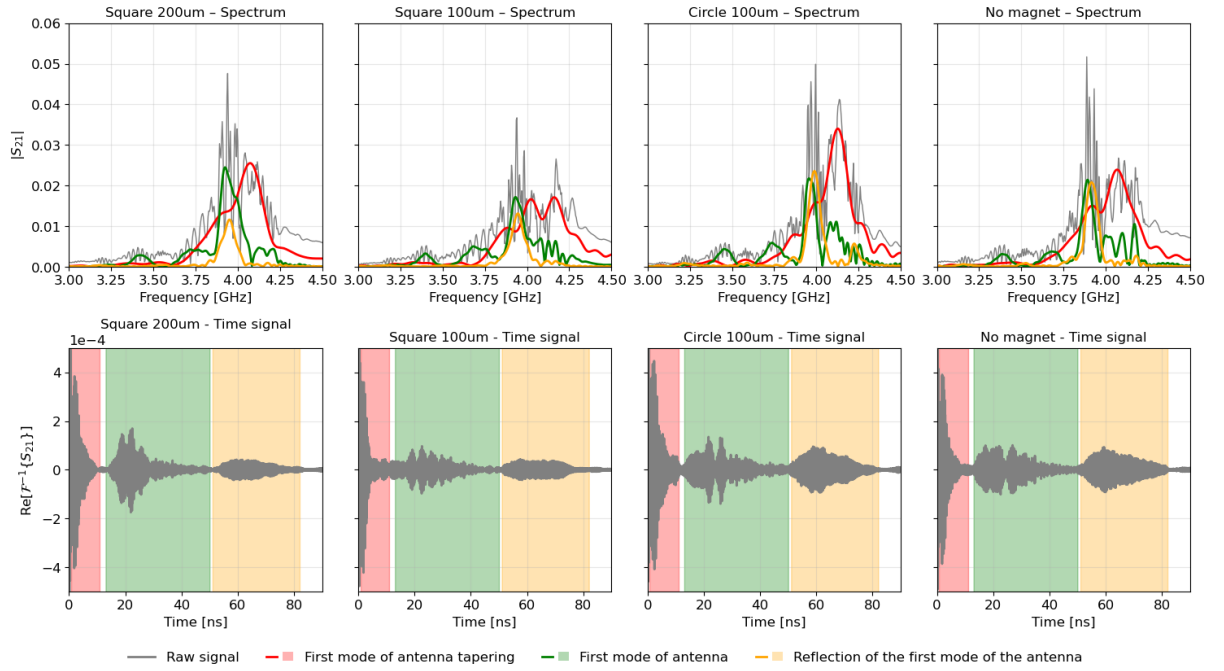


Figure 4.10: Time-gating analysis of N0 spectra for the four device geometries. Lower panels: time-domain signals obtained via inverse Fourier transform, with gating windows highlighted in red, green, and yellow to isolate different contributions. Upper panels: corresponding frequency-domain spectra for each gating window compared to the raw spectra. The red window (0.7 ns to 11 ns) isolates the fastest packet originating from the tapered antenna, exciting the first volume mode. The green window (13 ns to 50 ns) selects the first mode from the main antenna and some other contributions from the tapered region. The yellow window (51 ns to 82 ns) isolates the reflection of the first quantized volume mode of the main antenna on the receiving antenna and then back from the input antenna. The direct electromagnetically coupled contribution occurs earlier than 0.7 ns and has been excluded from the first window.

Even though a complete separation of all contributions is not achievable, a time-gated signal should still be used to better visualize the effect of magnet height modulation in the following section. The chosen time gating window is naturally the one that includes the first mode of the proper antenna (13 ns to 50 ns), which exhibits the strongest modulation with height. Contributions from the antenna tapering show negligible modulation, since the excited spin-waves propagate in regions of the film unaffected by the magnet, further confirming their origin.

4.2.4. Voltage modulation

To evaluate the effect of the micromagnets, measurements were performed by applying voltages in the 0 V to 60 V range, with a step of 2.5 V, to the MEMS actuators. Figure 4.11 shows the results in terms of amplitude for the four cases considered using the time gating window selected in the previous section. A significant modulation of the transmission is observed only in the spectral region corresponding to the main excitation-efficiency peak, where the antenna most efficiently injects energy into spin-waves, which also corresponds to the higher-frequency portion of the spectrum. In this region, according to the dispersion relation, the spin-wave wavelength becomes comparable to the lateral dimensions of the magnetic perturbation. Since the magnets are positioned relatively far from the substrate (more than $15\ \mu\text{m}$), their influence is appreciable mainly for wavelengths that interact more strongly with the perturbation field.

Importantly, the sign of the modulation is consistent with physical expectations: increasing the applied voltage raises the MEMS membrane, thereby increasing the magnet-film separation and reducing the perturbation strength, which results in higher spin-wave transmission. The modulation gradually saturates around 60 V, corresponding to the maximum displacement of the MEMS membranes.

Further confirmation of the magnetic origin of the effect is provided by measurements performed on the membrane without a magnet on top. In this case, the spin-wave transmission is only weakly affected by actuation. The small residual modulation observed is likely due to the unintended deposition of magnetic material on the sides of the piezoelectric membrane during fabrication.

The maximum transmission modulation achieved is on the order of 50% and, in agreement with the simulations, no substantial differences are observed between the two magnet geometries. The slightly lower transmitted signal recorded for the $100\ \mu\text{m}$ square magnet and for the no-magnet configuration is attributable to the time-gating procedure: as explained in the previous section, the raw signals show similar oscillating amplitudes, but the filtering operation isolates only the selected spin-wave packets, allowing the effect of the magnet on the propagation to be observed more clearly, at the expense of excluding part of the signal energy and therefore reducing the measured amplitude in some cases. Moreover, since the detected signal arises from a complex interference pattern between multiple propagating contributions, even small variations in their relative amplitude or phase can lead to significant differences in the reconstructed amplitude after time gating. As a result, the gated signal can exhibit noticeable variations across the different geometries considered, even when the corresponding raw signals have comparable intensities.

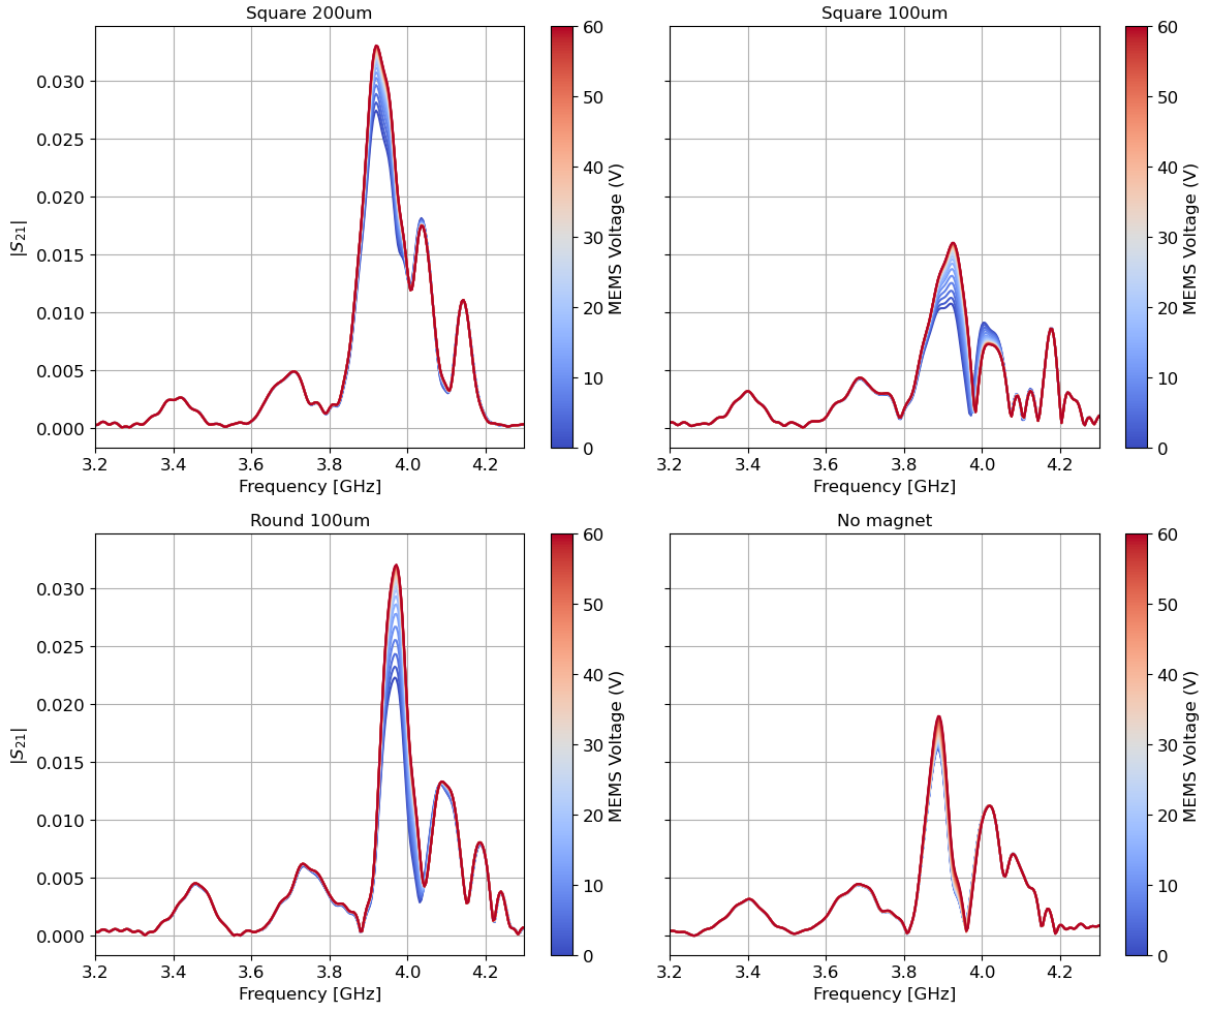


Figure 4.11: Measured transmission spectra for the four N0 geometries for different applied voltages of 0 V to 60 V and a bias magnetic field of 92 mT. Only the magnitude of the transmission scattering parameter, $|S_{12}|$, is shown.

The phase analysis has been conducted on the raw spectra, after having subtracted a reference signal, acquired with an applied bias field of 165 mT to isolate the spin-wave contribution. The effect of the magnet height modulation on the phase of the transmitted signal is illustrated in Figure 4.12, which shows the measured raw phase (on the left) alongside the time gated amplitude (on the right) for different voltages applied to the MEMS, for the square magnet 200 μm in side. As expected, each frequency corresponds to a different wavelength and therefore to a different accumulated propagation phase. The phase is not wrapped exactly within the interval $[-180^\circ, 180^\circ]$ because the measured signal in this region of the spectrum results from the superposition of two main contributions: the spin-wave signal, which can be represented as a phasor rotating with frequency, and the direct electromagnetic coupling, whose phase remains approximately constant over

the considered frequency range, as the wavelength of the electromagnetic signal traveling through the air is much larger than that of the spin-wave. The modulation induced by the applied MEMS voltage is relatively small, reaching at most about 20° in specific regions. This behavior is consistent across all investigated magnet geometries, indicating that the magnets primarily modulates the amplitude of the transmitted spin-waves rather than acting as efficient phase shifters.




Figure 4.12: Measured phase (left) and amplitude (right) spectra for the $200\ \mu\text{m}$ square magnet at applied voltages in the $0\ \text{V}$ to $60\ \text{V}$ range and a bias magnetic field of $92\ \text{mT}$. The phase spectrum was obtained by subtracting a reference spectrum acquired with a bias field of $165\ \text{mT}$. A time-gating procedure was applied to the amplitude spectrum using the same time window used previously. In the investigated frequency region, corresponding to the maximum transmission of the spin-wave signal, the phase modulation reaches a maximum value on the order of 20° .

4.2.5. Comparison with the simulations

To enable a direct comparison between experimental measurements and simulations, the computed spin-wave signal was processed by integrating the response along the central portion of the simulated receiving antenna (a $200\ \mu\text{m}$ -long region), yielding a single amplitude and phase value for each combination of excitation frequency and magnet height.

Figure 4.13 presents this comparison, showing both simulated and experimental results as a function of magnet–film distance at a fixed frequency, for the case of the round $100\ \mu\text{m}$ magnet. For the experimental data, a frequency of $3.95\ \text{GHz}$ was selected, corresponding to the maximum of the transmission signal, where the modulation is most pronounced. In the simulations, the closest available frequency point, $4\ \text{GHz}$, was used. The simulated

magnet heights were chosen in the 20 μm to 25 μm range, consistent with the experimentally accessible displacements derived from the device characterization presented earlier. The first column of the figure shows the relative amplitude variation, expressed as a percentage with respect to the lowest voltage (smallest magnet–film distance) for both simulation and experiment. This normalization highlights the modulation induced by increasing magnet height while removing absolute scaling differences between the two cases. The second column presents the corresponding phase values, shown as absolute quantities, allowing a direct comparison of the phase evolution as a function of magnet distance between simulated and experimental data.

Although the comparison is necessarily coarse due to the limited number of simulated frequency points and the indirect correspondence between voltage and magnet height, a clear qualitative agreement between simulations and experimental results can be observed. In both cases, the amplitude exhibits a substantial modulation at the selected frequency, while the phase remains only weakly affected by variations in magnet height. Most importantly, the overall trend is consistent between simulations and measurements, confirming that the model captures the essential physical behavior of the system.

In the experimental amplitude plot, a saturation of the modulation is visible at higher voltages, reflecting the limited displacement range of the MEMS actuators, as discussed in the device characterization section. This saturation is not present in the simulations, since the magnet height is directly imposed as a parameter and is not constrained by the mechanical response of the MEMS.

This general behavior is consistently observed across all investigated magnet geometries and frequencies, confirming that varying the magnet–film distance primarily modulates the strength of the interaction without significantly affecting the phase measured at the receiving antenna. This further supports the interpretation that, within the explored distance range, the magnets mainly act as scatterers, modifying the transmitted amplitude while leaving the overall accumulated phase largely unchanged.

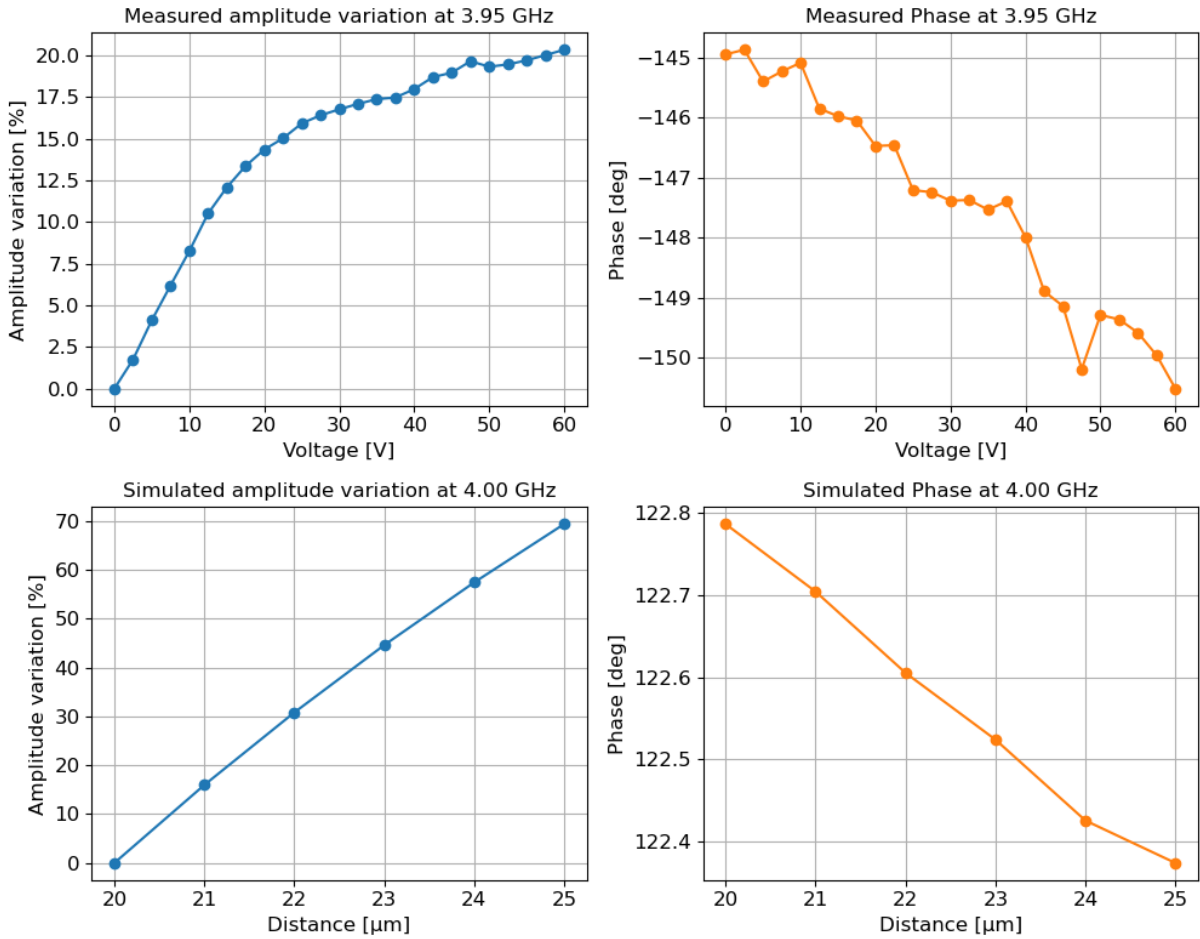


Figure 4.13: Comparison between experimental and simulated spin-wave transmission as a function of magnet–film distance, for the 100 μm round magnet, at fixed frequency of 3.95 GHz. On the simulation side, the closest available simulated frequency, 4 GHz, was used. The first column reports the amplitude variation, expressed as a percentage relative to the smallest magnet–film distance, and the second column shows the corresponding phase. Both simulations and experiments show a consistent trend, characterized by substantial amplitude modulation and only minor phase variations as the magnet–film distance increases.

5 | N1 - Proof of concept processor

This chapter presents the experimental investigation of the N1 device as a reconfigurable MEMS–magnonic processor. After a preliminary characterization, detailed in Section 5.1, the results of several training experiments are presented in Section 5.2. The discussion begins by verifying the fundamental physical mechanism enabling device operation: the magnet-mediated scattering of spin-waves and its electrical tunability via MEMS actuation. Building on this physical validation, the chapter then presents the experimental implementation of hardware-level training using the Direct Search (DS) optimization algorithm to maximize signal transmission between the input antenna and a selected output antenna. The convergence properties of the algorithm are examined through multiple optimization runs, both from identical initial configurations and from randomly initialized ones, to assess the robustness and reproducibility of the training process. The performance of DS is subsequently compared with the Simultaneous Perturbation Stochastic Approximation (SPSA) algorithm, highlighting their respective strengths in convergence speed, stability, and optimization effectiveness. Finally, the capabilities of the device are further explored by applying DS to additional optimization tasks, demonstrating the platform’s potential for programmable, reconfigurable wave-based signal processing while also identifying its current limitations.

5.1. Preliminary characterization

Before performing the training experiments, a preliminary characterization of the device was carried out to determine the optimal operating conditions and evaluate the baseline spin-wave transmission response. By tuning the bias magnetic field, it is possible to shift the spin-wave dispersion relation along the frequency axis, as explained in Section 1.3.5. This allows the selection of a frequency range in which the overall system response, comprising both the device under test, the RF PCB, and the signal transmission chain (cables and connectors), is sufficiently strong and clearly distinguishable from the background noise. By scanning different magnetic field values, it was determined that a field of approximately 100 mT enables operation in the 3 GHz to 5 GHz frequency range,

where the system response is optimal.

The first step in the characterization of the N1 device consisted in measuring the transmitted signal at the four output antennas without applying a voltage to the MEMS. The signals received by the central antennas, RF2 and RF3, were found to be weaker than those measured at the outer antennas, RF1 and RF4. The spectra were measured under a bias magnetic field of 100 mT, while the corresponding reference spectra were acquired at 180 mT, a value high enough to shift the spin-wave response outside the measured frequency band.

The four spectra are shown in Figure 5.1 with their time-gated counterparts. The lower panel of the figure reports the time-domain response of the system, obtained via inverse Fourier transform, along with the selected gating window. For antennas RF1 and RF4, the spin-wave signal can be clearly distinguished from the electromagnetic background and isolated through time gating. In contrast, for RF2 and RF3 the spin-wave contribution is significantly weaker. Considering also that most of the non-functional MEMS are located near antenna RF4, subsequent measurements were primarily focused on antenna RF1.

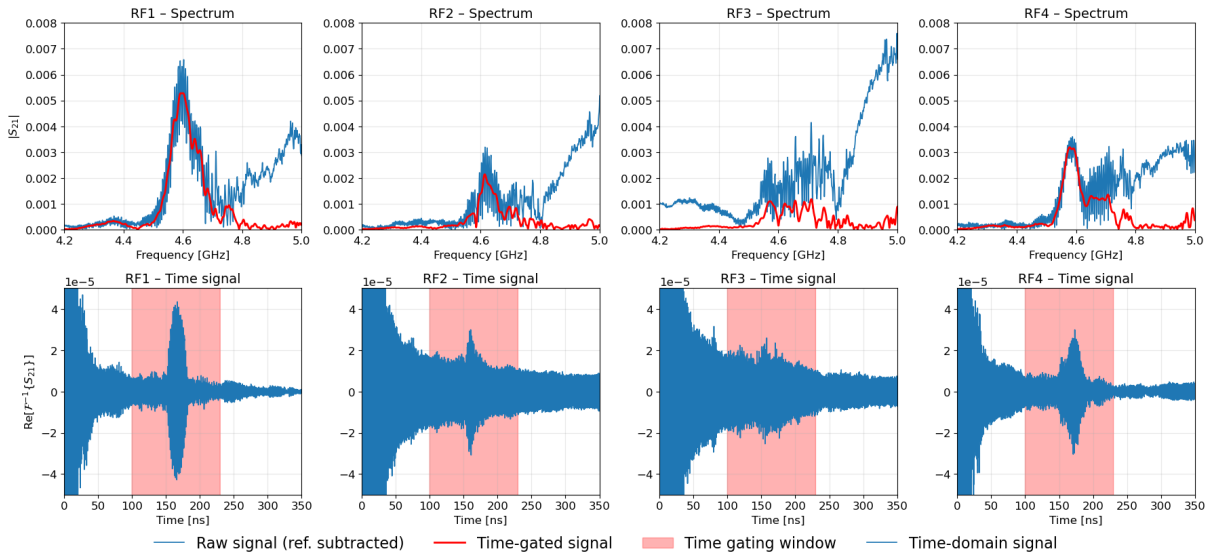


Figure 5.1: Transmission spectra measured at the four output antennas (RF1–RF4) under a bias magnetic field of 100 mT and their corresponding time-gated spectra. The lower panel shows the time-domain response obtained via inverse Fourier transform, with the selected gating window (100 ns to 230 ns) highlighted. For each antenna, a reference spectrum acquired at 100 mT has been subtracted from the raw data to remove background contributions.

Based on these preliminary observations, the measurement parameters for the detailed

characterization of the device were selected as reported in table 5.1. These parameters were used for all subsequent measurements.

Parameter	Value
Bias magnetic field	96 mT
Frequency range	4.2 GHz to 4.6 GHz
Time-gating window	100 ns to 230 ns
VNA bandwidth	50 Hz
VNA samples	401
VNA power	0 dBm

Table 5.1: Experimental parameters used in all the subsequent measurements.

5.2. Training results

5.2.1. Objective function

The objective function provides a quantitative measure of the device performance with respect to a given task by mapping the system response onto a single scalar value. During training, the function is evaluated from the measured transmission spectrum $S_{21}(f)$ and guides the optimization algorithms toward the desired modification of the spin-wave response.

The objective is defined starting from the difference between the magnitude of the measured spectrum and that of a reference spectrum,

$$\Delta|S_{21}(f)| = |S_{21}^{\text{meas}}(f)| - |S_{21}^{\text{ref}}(f)|, \quad (5.1)$$

and is obtained by integrating this quantity over selected frequency intervals.

As a first example, the maximization of the transmitted signal over a frequency band $[f_1, f_2]$ is achieved by defining the objective function as

$$O_{\text{max}} = w \int_{f_1}^{f_2} \Delta|S_{21}(f)| df. \quad (5.2)$$

Maximizing this objective corresponds to increasing the transmission in the selected frequency range with respect to the reference configuration.

More complex objectives can be constructed by combining multiple frequency intervals

with different weights. For instance, to maximize the signal in a band $[f_1, f_2]$ while simultaneously minimizing it in another band $[f_3, f_4]$, the objective function can be written as

$$O_{\text{mix}} = w_1 \int_{f_1}^{f_2} \Delta |S_{21}(f)| df - w_2 \int_{f_3}^{f_4} \Delta |S_{21}(f)| df, \quad (5.3)$$

where w_1 and w_2 are positive weighting factors that set the relative importance of the two contributions. The sign convention is chosen such that the optimization algorithms, which always maximize the objective function, are encouraged to enhance the signal in the first band and suppress it in the second.

In practice, the integrals appearing in the objective function are evaluated numerically and implemented as the mean value of the discrete spectrum difference within each frequency band. The use of a reference spectrum serves multiple purposes: it emphasizes the spectral changes induced during training, simplifies the interpretation of the sign of each contribution to the objective, and provides a common baseline for comparing the outcomes of different training runs. The objective function is not normalized; instead, the frequency integrals and weighting factors are chosen such that the resulting objective values lie in the range of a few tens to a few hundreds. This scaling is sufficiently large to avoid numerical issues, particularly for the SPSA algorithm, which relies on objective-function differences to estimate gradients and can be sensitive to excessively small values. Figure 5.2 illustrates the mathematical definition of the objective function used to maximize the transmitted spectrum over a wide frequency band.




Figure 5.2: Objective function definition to maximize the received S_{21} signal in the frequency band from 4.3 GHz to 4.5 GHz.

5.2.2. Preliminary device validation

As a first step, it was necessary to assess whether the magnets have an impact on spin-wave propagation and, consequently, on the spectra measured at antenna RF1, and whether the algorithm is able to correctly drive the system toward the desired configuration. To this end, the result of a training session aimed at maximizing the received power at antenna RF1 over a broad frequency range was compared with a physically motivated configuration.

The physically reasonable configuration consists in actuating all MEMS at their maximum voltage (60 V), thereby maximizing the distance between the magnets and the YIG film over the entire propagation area. This configuration is expected to reduce the overall perturbation introduced by the magnetic elements, mainly in terms of scattering and reflection, and therefore to increase the transmitted signal in a broad wavelength and frequency range.




Figure 5.3: Spectra comparison between the maximization training result and the configuration with 60 V applied to all MEMS. The two measured spectra have to be compared with the reference spectrum plotted in gray.

Figure 5.3 compares the transmission spectra obtained with this uniform 60 V configuration and with the final configuration reached by the training algorithm. As expected, the uniform actuation results in a general increase of the signal over the entire spectral range, however, while for frequencies higher than approximately 4.4 GHz the two spectra almost overlap, for lower frequencies the signal obtained after training is significantly higher than that achieved with the uniform configuration.

This behavior is reflected in the objective function values: the uniform 60 V configuration yields an objective value $O_{60V} = 22.21$, whereas the final trained configuration reaches $O_{tr} = 34.26$. This result shows that the magnets inside the device are able to influence spin-wave propagation and that their effect is more complex than a uniform reduction of scattering obtained by retracting all magnets. It also demonstrates that the device and the training algorithm are able, at least partially, to control this scattering process and drive the system toward the assigned objective.

Figure 5.4 reports the voltages applied to the MEMS for both the uniform 60 V configuration and the final configuration reached by the training algorithm. The actuation voltages range from 0 V to 60 V, where 0 V corresponds to the magnets being closest to the YIG film and 60 V to the maximum retraction.

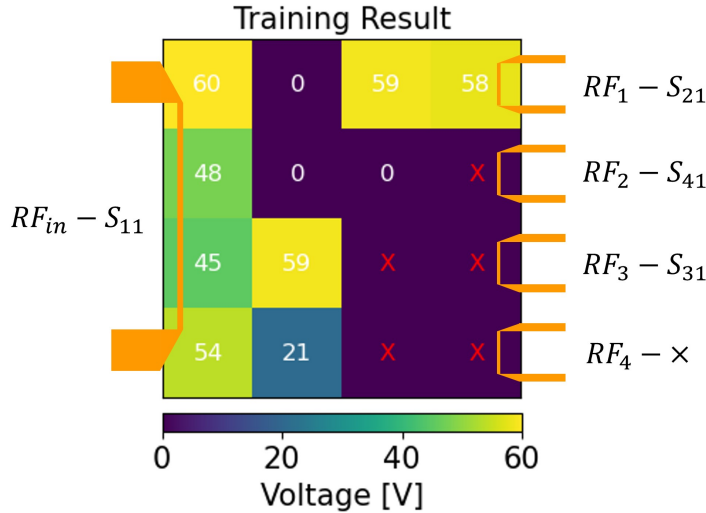


Figure 5.4: Final voltage configuration of the maximization training. The broken MEMS are marked with a red cross and are to be considered stuck at 0 V. The antennas positions and labeling are included for the sake of clarity; they will be omitted in the next figures.

Notably, the MEMS at position $[1, 2]$ (first row, second column) is set by the algorithm to 0 V. This provides further evidence that the device is maximizing the signal by exploiting more complex effects than simply creating a favorable direct propagation path. Raising this magnet would, in principle, form a *channel* between the input antenna and RF1, resulting in a signal increase; however, apparently the algorithm achieves an even higher enhancement by leveraging more complex interactions with the magnets. As a final remark, the many non-functional MEMS, stuck in the 0 V position close to the YIG film, have an effect on the spin-wave propagation that cannot be changed but that the training algorithm has to exploit to reach its goal; this means that some training tasks are

made harder or even impossible by this circumstances and it is one of the reasons why the efforts have been primarily focused on antenna RF1.

5.2.3. DS maximization

After verifying the basic functionality of the device, the two training algorithms, Direct Search and Simultaneous Perturbation Stochastic Approximation, were investigated more systematically. The objective was to evaluate their performance in terms of convergence, repeatability, and robustness with respect to the initial conditions. In particular, the ability of each algorithm to reach configurations associated with high objective-function values was assessed by initializing the training either from the same predefined voltage configuration or from different randomly generated ones.

Using the DS algorithm, the analysis initially focused on the same task discussed in the previous section, namely the maximization of the signal measured at antenna RF1 within the frequency range from 4.3 GHz to 4.5 GHz. The training sessions were divided into two subsets: one starting from a uniform initial configuration in which all MEMS actuators were set to 0 V, and another in which each training run was initialized with a different randomly generated voltage configuration. The final transmission spectra obtained for these two cases are shown in Figure 5.5.




Figure 5.5: Reference and final spectra for the maximization DS trainings starting from (a) 0 V configurations and (b) random configurations.

Within each subset, a good agreement among the training outcomes is observed. In the case of the 0 V initial configuration, both the reference spectra and the final spectra

overlap almost perfectly. For the subset starting from random initial configurations, one of the spectra appears shifted toward lower frequencies, together with its corresponding reference spectrum. This shift is most likely due to a slight inconsistency in the bias magnetic field generated by the electromagnets, possibly caused by hysteresis effects and by the fact that the measurements were performed at different times.




Figure 5.6: Comparison between the results of the DS training started from a 0 V configuration and from a random configuration.

A direct comparison between one spectrum obtained from the 0 V initial configuration and one from the random-initialization subset is shown in Figure 5.6, where a clear overlap can be observed. This behavior is further confirmed by the final voltage configurations, shown in Figure 5.7 for the two cases. It is evident that the algorithm converges toward essentially the same final configuration regardless of the initial conditions, demonstrating the robustness and repeatability of the DS optimization process.

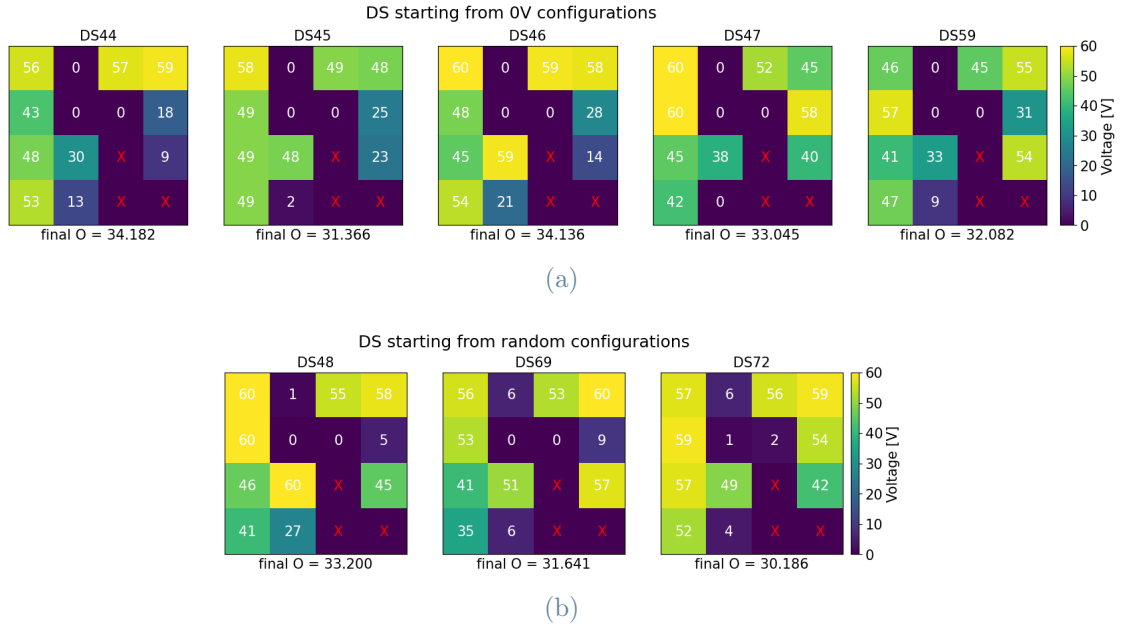


Figure 5.7: Final voltage maps for the maximization DS trainings starting from (a) 0 V configurations and (b) random configurations.

Further insight into the behavior of the Direct Search optimization is provided by the standard deviation map of the final MEMS voltages, obtained by combining the results presented above and shown in Figure 5.8. This representation allows to directly visualize which MEMS have a stronger impact on the optimization outcome and which ones play a marginal role, by highlighting how consistently each voltage is selected across different training runs.

The voltages associated with the two MEMS located close to the inner output antennas are included in the map, as at the time of the training sessions it had not yet been established that these devices were non-functional. However, the standard deviation plot clearly shows that variations in their voltage values have no measurable impact on the optimization.

Overall, the standard deviation is low for most MEMS, indicating that the algorithm converges to well reproducible voltage values. Slightly higher deviations are observed for the two MEMS located farther from both the input antenna and the receiving antenna. This behavior is consistent with expectations, as MEMS closer to the input antenna are expected to have a stronger influence on spin-wave excitation and propagation. This effect is further reinforced by the non-uniform tilt introduced during the flip-chip assembly, which results in the MEMS on the input side being positioned closer to the YIG film and therefore exerting a stronger influence on the spin-wave dynamics.




Figure 5.8: Standard deviation map of the final voltages for all DS maximization trainings, starting from both 0 V and random configurations.

5.2.4. SPSA maximization

After evaluating the performance of the Direct Search algorithm, SPSA was assessed on the same maximization task, starting from a uniform initial configuration with all MEMS set to 0 V. The training runs were executed sequentially, ensuring a very good overlap of the reference spectra, as shown in Figure 5.9. In contrast to the DS case, the final spectra exhibit a weaker agreement among different training runs, a behavior that is also reflected in the final voltage configurations shown in Figure 5.10 and in the corresponding standard deviation map.




Figure 5.9: Reference and final spectra for the SPSA maximization training started from a 0 V configuration.

Despite the increased variability, clear trends can still be identified in the final configu-

rations. The MEMS in the first row generally converge to high voltage values, with the notable exception of the membrane at position [1, 2]. In the second row, the two central MEMS are consistently set to low voltages, while the MEMS in the first column tend to assume higher voltage values. The standard deviation map further indicates that MEMS located farther from the receiving antenna have a reduced influence on the optimization outcome. However, as shown by the objective values reached in the different training runs and reported in Figure 5.10, this increased variability in the voltages configurations leads to small differences in the final achieved performance across the trainings.

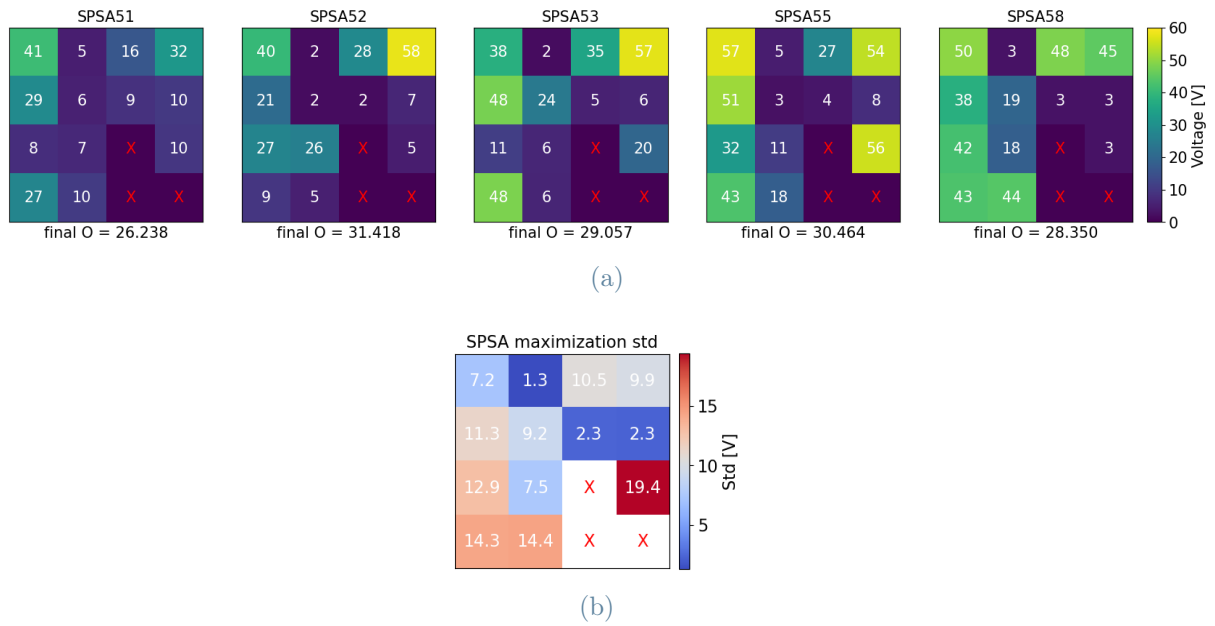


Figure 5.10: Final voltage maps for the SPSA maximization trainings (a) and their standard deviation (b).

This behavior points to a non-optimal convergence of the SPSA algorithm toward the maximum of the objective function, which is strongly dependent on an accurate tuning of the algorithm parameters. This aspect is discussed in more detail in the following paragraph, where a direct comparison between the DS and SPSA algorithms is presented.

5.2.5. DS - SPSA comparison

With the data gathered so far, it is possible to directly compare the two training algorithms when assigned the same optimization task. As already observed, SPSA consistently performs slightly worse than its DS counterpart. The best results achieved by the two algorithms are shown in Figure 5.11, where the corresponding final spectra are plotted together with their reference spectra, and the final objective values are reported.




Figure 5.11: Comparison between the results of the DS and SPSA algorithms when tasked with the maximization of the signal.

Further insight into this behavior is provided by Figure 5.12, which shows the evolution of both the objective value and the voltage configurations as a function of the iteration number. While both algorithms are able to reach relatively high objective values within a limited number of iterations, SPSA exhibits a higher degree of intrinsic instability and tends to oscillate around the optimum when not properly tuned. This behavior is clearly visible in training SPSA56, whose objective evolution is shown in the upper-right panel: in this case, an excessively large value of the gain parameter a (introduced in Section 2.3.1) leads to too large steps that prevent accurate convergence.

Additionally, an inspection of the voltage configuration plots (bottom panels of Figure 5.12) reveals that SPSA is unable to fully explore the entire 0 V to 60 V parameter range, as a finite margin is always required to compute the gradient through simultaneous perturbations. In contrast, the DS algorithm is free to directly sample the full admissible voltage range.

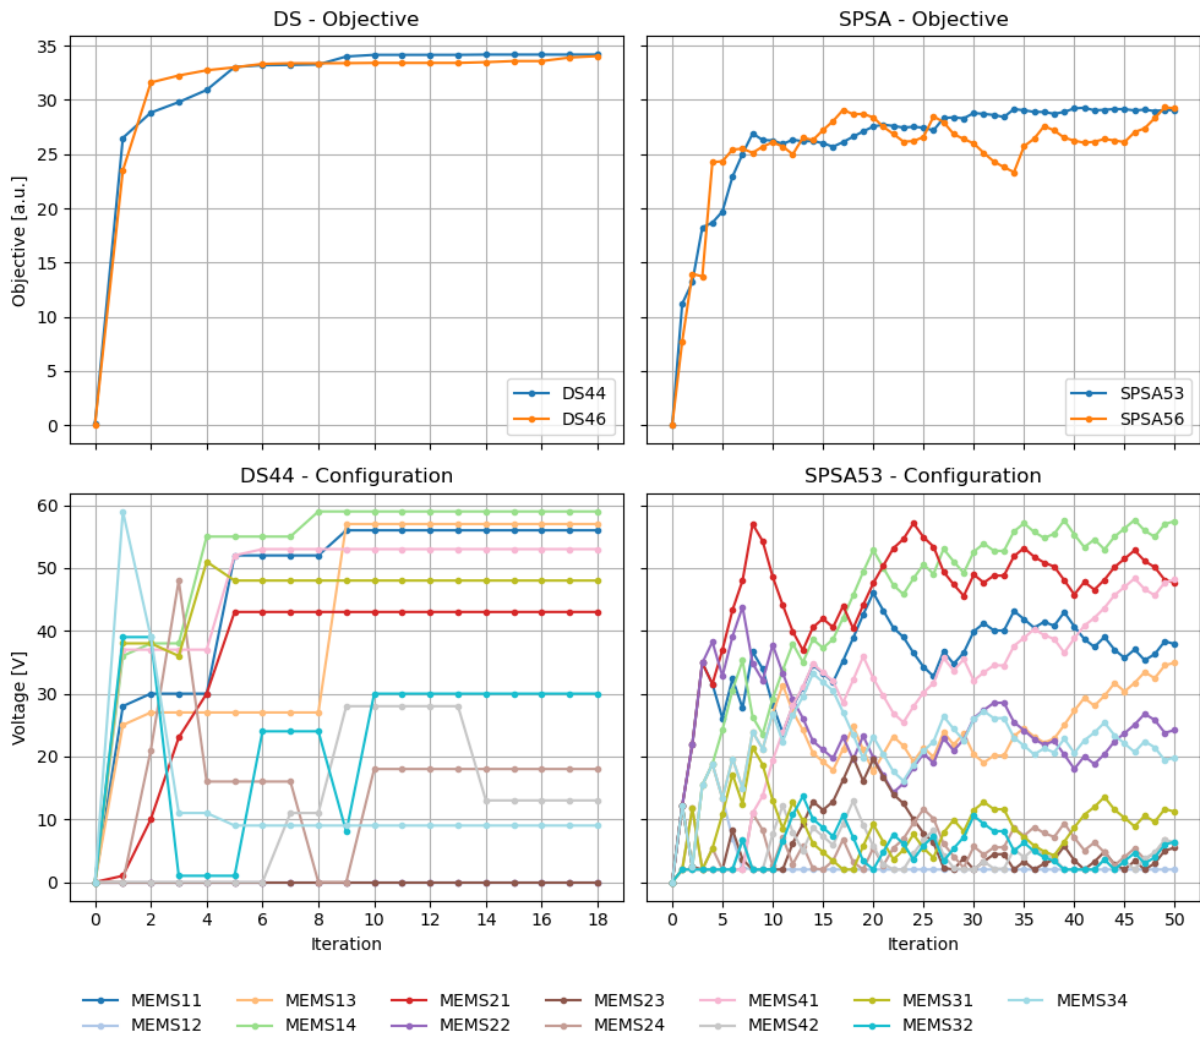


Figure 5.12: Objectives and voltage configurations evolutions during the trainings, on the left using the DS algorithm and on the right using SPSA.

Finally, although the DS algorithm typically reaches convergence in a significantly smaller number of iterations, SPSA remains computationally faster in practice. Each SPSA iteration requires only two evaluations of the objective function, whereas a single DS iteration involves one evaluation per parameter. With the VNA settings used in these experiments, one DS iteration takes approximately 2 minutes, while a single SPSA iteration requires only about 20 seconds. More importantly, the execution time of DS scales linearly with the dimensionality of the parameter space, whereas it remains constant for SPSA, making the latter particularly attractive for larger-scale systems.

5.2.6. DS minimization

After comparing the two algorithms and establishing that DS is slightly more precise than SPSA, despite requiring longer convergence times, additional objectives were investigated using the DS algorithm. In particular, the specular task with respect to the previous one has been considered, namely minimizing the signal received at antenna RF1 in the frequency range between 4.3 GHz to 4.5 GHz.

Figure 5.13 shows the spectra resulting from these training sessions and demonstrates that the device is able to reduce the transmission from input to output only for frequencies below approximately 4.4 GHz, when compared to the reference spectra. The reference spectra were consistently measured in the 0 V configuration, which also served as the initial condition for all training runs.




Figure 5.13: Reference and final spectra for the minimization training using DS starting from a 0 V configuration.

An interesting observation emerges when comparing the final voltage configurations obtained for this minimization task with those corresponding to the maximization task. Figure 5.14 reports the average voltage values over the two training sets, displayed side by side. The two configurations are clearly specular: MEMS actuators set to high voltages in the maximization case are now driven to low voltages, and vice versa. The only exception is represented by the two non-functional MEMS located closer to the output antennas.




Figure 5.14: Average voltage maps for the (a) maximization and (b) minimization trainings using Direct Search. The values are clearly specular with the exception of the ones assigned to broken MEMS.

5.2.7. Maximize and minimize

The final task tested, again using the DS algorithm, was to decrease the signal received by antenna RF1 in one frequency band while increasing it in another, with the aim of achieving a filter-like behavior. Specifically, the objective was to decrease the signal (with respect to the reference) in the 4.3 GHz to 4.4 GHz band and to increase it in the 4.4 GHz to 4.5 GHz band.

In the left panel of Figure 5.15, the resulting spectra from the trainings are shown together with the best result obtained for the previous maximization task. Two trainings were performed: the first with equal weights assigned to the minimization and maximization bands, and the second with a tenfold increase of the minimization weight, since the first training resulted in a general increase of the signal, very similar to the maximization case. The second attempt instead led to a spectrum that is essentially unchanged with respect to the reference.

By analyzing the different configurations tested and stored by the DS algorithm during the various trainings, since the algorithm effectively explores different regions of the parameter space, a configuration was identified that performs well for the opposite task, namely decreasing the signal at lower frequencies while increasing it at higher frequencies. This configuration is shown in the right panel of Figure 5.15.




Figure 5.15: Reference and final spectra obtained from the mixed minimization-maximization trainings using the DS algorithm. (a) Results of two training runs: *min1*, using equal weights for the minimization (4.3 GHz to 4.4 GHz) and maximization (4.4 GHz to 4.5 GHz) bands, and *min2*, using a tenfold higher weight for the minimization band. Both are compared with the best signal maximization result obtained in training DS46, *max best*. (b) Spectrum corresponding to a found configuration that decreases the signal at lower frequencies while increasing it at higher frequencies, relative to the 0 V reference.

These results indicate that the device is not able to manipulate the spectrum independently at each frequency, but instead exhibits a tendency to modify the response over the entire frequency band. This behavior likely originates from the fact that, with the current geometry, the interference pattern established within the device cannot be fundamentally reshaped, but only modulated in intensity across different regions. Given the limited displacement range of the MEMS and the in-plane bias field configuration, as also confirmed by the simulations, the device cannot reach a true zero-perturbation state. Instead, it can only adjust the strength of the magnetic perturbation produced by the micromagnets, without qualitatively altering the overall interference pattern.

Achieving improved tunability would therefore require the existence of a configuration corresponding to a vanishing or negligible perturbation. This could be realized, for example, by employing a hard magnetic material, allowing the magnets to be magnetized along a direction different from that of the bias field, or by adopting a forward volume configuration in which the magnets are magnetized out of plane. In such a configuration, the stray field would be more localized and would decay more rapidly with distance, enabling a true off state when the magnets are sufficiently far from the film.

Furthermore, these results highlight the importance of identifying a suitable neutral configuration to serve as a reference for the training procedures. To this end, the parameter space should first be explored in a coarse manner, for instance using the DS algorithm with a sparse voltage mesh, in order to determine the achievable range of transmission spectra. This preliminary exploration would make it possible to identify a working point corresponding to an intermediate or minimally perturbed state, which could then be used as a well-defined reference for more selective and effective optimization tasks.

6 | Conclusions and future developments

During the course of this thesis work, two devices, N0 and N1, were realized to investigate the effect of micromagnets on spin-wave propagation and to exploit this interaction for the development of a hybrid MEMS-magnonic processor.

Both devices consist of two main components: a magnonic substrate and a MEMS-based actuation system. The magnonic part is based on a YIG thin film with patterned antennas for spin-wave excitation and detection. The MEMS section comprises a 4×4 array of piezoelectric membranes, on top of which micromagnets were fabricated. The two subsystems were integrated using a flip-chip assembly technique, enabling the magnets to be positioned close to the YIG substrate at a voltage-controlled distance while consuming very little power.

To characterize the devices, custom PCBs were designed, allowing integration within a dedicated measurement setup that includes a vector network analyzer, an NI DAC signal generator and a multichannel high-voltage amplifier to independently drive the MEMS membranes. Dedicated control software was developed to coordinate the VNA acquisition, DAC output, and external bias magnetic field, enabling automated measurements and systematic voltage sweeps. In addition, two optimization algorithms were implemented to program the N1 device according to the inverse-design approach, directly training the hardware to achieve specific transmission responses.

The N0 device served as a test structure to study the interaction between spin-waves and localized magnetic perturbations as a function of magnet geometry and vertical displacement. The experimental results show a clear amplitude modulation of the transmission signal, consistent with physical expectations and with the simulation results, without evidence of strong shape-dependent effects. The simulations further indicate that the magnet dimensions must be properly selected according to the spin-wave wavelength range used in the device and that being able to move them closer to the YIG would increase significantly their effect on spin-waves.

Building on these results, N1 represents a first proof of concept of a hybrid MEMS-magnonic processor based on mechanically reconfigurable magnetic landscape. Despite fabrication limitations, the device demonstrated that spin-wave transmission and routing between input and output antennas can be modulated in a controlled and non-trivial way by adjusting the micromagnet positions. These results validate both the experimental platform and the implemented training algorithms, confirming the feasibility of this low-power, reconfigurable approach to magnonic information processing.

Several directions can be pursued to further improve and extend the results obtained in this work. A first important step concerns the implementation of hard permanent magnets, such as SmCo, on the MEMS membranes. Using a hard magnetic material would allow the magnets to be independently magnetized in a different direction with respect to the external bias field, enabling more flexible and complex magnetic landscapes creation. Initial attempts to integrate SmCo on prefabricated MEMS chips have highlighted fabrication challenges: its hard magnetic phase requires high-temperature growth, which exceeds the thermal budget of completed MEMS devices and poses an issue for standard patterning techniques. A possible solution would be to move toward a custom-fabricated MEMS array, where SmCo deposition and patterning are performed in the early stages of the process. This approach would also allow improved geometric control, denser membrane arrays, and smaller magnets, thereby increasing the number of tunable parameters in the system. However, more densely packed membranes would likely provide smaller mechanical displacements due to their reduced lateral dimensions.

Another important improvement concerns the flip-chip assembly. Reducing the separation between the micromagnets and the YIG film would significantly enhance the interaction strength. Techniques such as conductive inks or alternative bonding strategies could help bring the magnets closer to the substrate while maintaining reliable electrical connections.

From the RF perspective, the system performance could benefit from improved antenna and PCB design. Optimizing the spin-wave injection efficiency and reducing insertion losses would enhance the signal with respect to the direct electromagnetic coupling. At the same time, reducing this latter, for example through alternative antenna geometries, would minimize the need for time-gating procedures and simplify data interpretation.

On the magnonic side, transitioning to a forward-volume configuration, where the YIG film is magnetized out of plane, would enable isotropic spin-wave propagation. Such a configuration would enhance the scattering effects of the micromagnets, potentially improving the effectiveness of the programmable magnetic landscape.

Finally, increasing the injected spin-wave power to access the nonlinear regime repre-

sents an important long-term objective. Nonlinear spin-wave dynamics are essential for implementing neuromorphic and recurrent neural network functionalities, as they naturally provide activation-like behavior and richer signal transformations. Reaching this goal would unlock the full potentiality for information-processing based on the hybrid MEMS-magnonic platform demonstrated in this thesis.

Bibliography

- [1] Catherine Schuman, Thomas Potok, Robert Patton, J. Birdwell, Mark Dean, Garrett Rose, and James Plank. A survey of neuromorphic computing and neural networks in hardware. 05 2017. doi: 10.48550/arXiv.1705.06963.
- [2] Danijela Marković, Alice Mizrahi, Damien Querlioz, and Julie Grollier. Physics for neuromorphic computing. *Nature Reviews Physics*, 2(9):499–510, 2020. doi: 10.1038/s42254-020-0208-2. URL <https://doi.org/10.1038/s42254-020-0208-2>.
- [3] Robert Nawrocki, Richard Voyles, and Sean Shaheen. A mini review of neuromorphic architectures and implementations. *IEEE Transactions on Electron Devices*, 63:1–11, 10 2016. doi: 10.1109/TED.2016.2598413.
- [4] Renjie Li, Yuanhao Gong, Hai Huang, Yuze Zhou, Sixuan Mao, Zhijian Wei, and Zhaoyu Zhang. Photonics for neuromorphic computing: Fundamentals, devices, and opportunities. *Advanced Materials*, 37(2):2312825, 2025. doi: <https://doi.org/10.1002/adma.202312825>. URL <https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/adma.202312825>.
- [5] A Barman, Gianluca Gubbiotti, Sam Ladak, Adekunle Olusola Adeyeye, Maciej Krawczyk, Joachim Gräfe, Christoph Adelman, Sorin Cotofana, Azad Naeemi, Vitaliy Vasyuchka, Burkard Hillebrands, S Nikitov, Haiming Yu, Dirk Grundler, Alexandr Sadovnikov, Andrew Grachev, S. Sheshukova, Jean-Yves Duquesne, Massimiliano Marangolo, Csaba Gyorgy, Wolfgang Porod, V Demidov, Sergei Urazhdin, Sergej Demokritov, Edoardo Albisetti, Daniela Petti, Riccardo Bertacco, Helmut Schulteiss, Volodymyr Kruglyak, Vlad Poimanov, Ashok Kumar Sahoo, Jaivardhan Sinha, Hyunsoo Yang, Markus Muenzenberg, Takahiro Moriyama, Shigemi Mizukami, Pedro Landeros, Rodolfo Andrés Gallardo, Giovanni Carlotti, Joo-Von Kim, Robert Stamps, Robert Camley, Bivas Rana, Y Otani, Weichao Yu, Tao Yu, Gerrit Bauer, Christian Back, Goetz Uhrig, Oleksandr Dobrovolskiy, Sebastiaan van Dijken, Barbora Budinska, Huajun Qin, Andrii Chumak, Aleksandr Khitun, Dmitri Nikonov, Ian Young, Benjamin Zingsem, and Michael Winklhofer. The 2021 Magnon-

- ics Roadmap. *Journal of Physics: Condensed Matter*, 33(41):413001, 2021. doi: 10.1088/1361-648X/abec1a. URL <https://hal.science/hal-03229528>.
- [6] Andrii Chumak, Vitaliy Vasyuchka, Alexander Serga, and Burkard Hillebrands. Magnon spintronics. *Nature Physics*, 11:453–461, 06 2015. doi: 10.1038/nphys3347.
- [7] Abdulqader Mahmoud, Florin Ciubotaru, Frederic Vanderveken, Andrii Chumak, Said Hamdioui, Christoph Adelman, and Sorin Cotofana. Introduction to spin wave computing, 09 2021.
- [8] Ramazan Yeniceri, Wolfgang Porod, György Csaba, and Adam Papp. Non-boolean computing based on linear waves and oscillators. 09 2015. doi: 10.1109/ESSDERC.2015.7324723.
- [9] Rainer Lucas, Marc Fossorier, Yu Kou, and Shu Lin. Iterative decoding of one-step majority logic deductible codes based on belief propagation. *IEEE Transactions on Communications*, 48:931–937, 01 2000.
- [10] György Csaba, Ádám Papp, and Wolfgang Porod. Perspectives of using spin waves for computing and signal processing. *Physics Letters A*, 381(17):1471–1476, 2017. ISSN 0375-9601. doi: <https://doi.org/10.1016/j.physleta.2017.02.042>. URL <https://www.sciencedirect.com/science/article/pii/S0375960116316486>.
- [11] G. Csaba, A. Papp, and W. Porod. Spin-wave based realization of optical computing primitives. *Journal of Applied Physics*, 115(17):17C741, 03 2014. ISSN 0021-8979. doi: 10.1063/1.4868921. URL <https://doi.org/10.1063/1.4868921>.
- [12] Qi Wang, Andrii V. Chumak, and Philipp Pirro. Inverse-design magnonic devices. *Nature Communications*, 12(1):2636, 2021. doi: 10.1038/s41467-021-22897-4. URL <https://doi.org/10.1038/s41467-021-22897-4>.
- [13] Tyler W. Hughes, Ian A. D. Williamson, Momchil Minkov, and Shanhui Fan. Wave physics as an analog recurrent neural network. *Science Advances*, 5(12):eaay6946, 2019. doi: 10.1126/sciadv.aay6946. URL <https://www.science.org/doi/abs/10.1126/sciadv.aay6946>.
- [14] Csaba Gyorgy Papp Ádám, Porod Wolfgang. Nanoscale neural network using non-linear spin-wave interference. *Nature Communications*, 11 2021. doi: 10.1038/s41467-021-26711-z. URL <https://doi.org/10.1038/s41467-021-26711-z>.
- [15] Noura Zenbaa, Claas Abert, Fabian Majcen, Michael Kerber, Rostyslav Serha, Sebastian Knauer, Qi Wang, Thomas Schrefl, D. Suess, and Andrii Chumak. A uni-

- versal inverse-design magnonic device. *Nature Electronics*, 8:106–115, 01 2025. doi: 10.1038/s41928-024-01333-7.
- [16] Anjan Barman and Arabinda Haldar. Chapter one - time-domain study of magnetization dynamics in magnetic thin films and micro- and nanostructures. volume 65 of *Solid State Physics*, pages 1–108. Academic Press, 2014. doi: <https://doi.org/10.1016/B978-0-12-800175-2.00001-7>. URL <https://www.sciencedirect.com/science/article/pii/B9780128001752000017>.
- [17] B.A. Kalinikos. Excitation of propagating spin waves in ferromagnetic films. *IEE Proceedings H (Microwaves, Optics and Antennas)*, 127:4–10, 1980. doi: 10.1049/ip-h-1.1980.0002. URL <https://digital-library.theiet.org/doi/abs/10.1049/ip-h-1.1980.0002>.
- [18] Pawel Gruszecki, Maciej Kasprzak, Andriy Serebryannikov, Maciej Krawczyk, and W. Śmigaj. Microwave excitation of spin wave beams in thin ferromagnetic films. *Scientific Reports*, 6:22367, 03 2016. doi: 10.1038/srep22367.
- [19] Olle Karlqvist. *Calculation of the magnetic field in the ferromagnetic layer of a magnetic drum*. 1954.
- [20] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>.
- [21] Baolin Peng and Kaisheng Yao. Recurrent neural networks with external memory for language understanding. *Lecture Notes in Computer Science*, 05 2015. doi: 10.1007/978-3-319-25207-0_3.
- [22] Sven Crone. Artificial neural networks for time series prediction - a novel approach to inventory management using asymmetric cost functions. volume 1, pages 193–199, 01 2003.
- [23] Zhanshe Guo, Fucheng Cheng, Boyu Li, Le Cao, Chao Lu, and Ke Song. Research development of silicon mems gyroscopes: a review. *Microsystem Technologies*, 21(10): 2053–2066, 2015. ISSN 1432-1858. doi: 10.1007/s00542-015-2645-x. URL <https://doi.org/10.1007/s00542-015-2645-x>.
- [24] Honglong Chang, Liang Xue, Wei Qin, Guangmin Yuan, and Weizheng Yuan. An integrated mems gyroscope array with higher accuracy output. *Sensors*, 8(4):2886–2899, 2008. ISSN 1424-8220. doi: 10.3390/s8042886. URL <https://www.mdpi.com/1424-8220/8/4/2886>.
- [25] S. Sathya, S. Muruganand, N. Manikandan, and K. Karuppasamy. Design of capaci-

- tance based on interdigitated electrode for biomems sensor application. *Materials Science in Semiconductor Processing*, 101:206–213, 2019. ISSN 1369-8001. doi: <https://doi.org/10.1016/j.mssp.2019.06.005>. URL <https://www.sciencedirect.com/science/article/pii/S1369800119304913>.
- [26] Jiushuai Xu, Maik Bertke, Hutomo Wasisto, and Erwin Peiner. Piezoresistive micro-cantilevers for humidity sensing. *Journal of Micromechanics and Microengineering*, 29, 04 2019. doi: 10.1088/1361-6439/ab0cf5.
- [27] Abdullah Algamili, Mohd Khir, John Dennis, Abdelaziz Ahmed, Sami Alabsi, Saeed Ba Hashwan, and Muhammad Junaid. A review of actuation and sensing mechanisms in mems-based sensor devices. *Nanoscale Research Letters*, 16, 01 2021. doi: 10.1186/s11671-021-03481-7.
- [28] Chang-Beom Eom and S. Trolier-McKinstry. Thin-film piezoelectric mems. *MRS Bulletin*, 37, 11 2012. doi: 10.1557/mrs.2012.273.
- [29] Joontaek Jung, Wonjun Lee, Woojin Kang, Eunjung Shin, Jungho Ryu, and Hongsoo Choi. Review of piezoelectric micromachined ultrasonic transducers and their applications. *Journal of Micromechanics and Microengineering*, 27, 09 2017. doi: 10.1088/1361-6439/aa851b.
- [30] Sina Akhbari, Firas Sammoura, Stefon Shelton, Chen Yang, David Horsley, and Liwei Lin. Highly responsive curved aluminum nitride pmut. In *2014 IEEE 27th International Conference on Micro Electro Mechanical Systems (MEMS)*, pages 124–127, 2014. doi: 10.1109/MEMSYS.2014.6765589.
- [31] Philipus Hishimone, Hiroki Nagai, and Mitsunobu Sato. *Methods of Fabricating Thin Films for Energy Materials and Devices*. 07 2020. ISBN 978-1-78985-463-3. doi: 10.5772/intechopen.85912.
- [32] S. Cuccurullo, F. Maspero, O. Koplak, G. Pavese, E. Albisetti, M. Cantoni, and R. Bertacco. Impact of minor hysteresis loops in integrated inductors with ferromagnetic films. *Applied Physics Letters*, 122(11):113503, 03 2023. ISSN 0003-6951. doi: 10.1063/5.0127390. URL <https://doi.org/10.1063/5.0127390>.
- [33] Vector network analyzer (vna) calibration: The basics. White Paper, Rohde & Schwarz / Michael Hiebel. URL <https://pchene.wordpress.com/wp-content/uploads/2019/08/8e60b68a24a4169eecbd0763d4a9287037e4.pdf>. Accessed: 2026-02-26.
- [34] J.C. Spall. Implementation of the simultaneous perturbation algorithm for stochastic

- optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3):817–823, 1998. doi: 10.1109/7.705889.
- [35] Antal Szava and David Wierichs. PennyLane - optimization using spsa, 2025. URL https://pennylane.ai/qml/demos/tutorial_spsa.
- [36] URL <https://mandmems.eu>.
- [37] URL <https://mumax.github.io/plus/>.
- [38] URL <https://orfeo-doc.areasciencepark.it/>.
- [39] Alexis Wartelle, Franz Vilsmeier, Takuya Taniguchi, and Christian H. Back. Caustic spin wave beams in soft thin films: Properties and classification. *Phys. Rev. B*, 107:144431, Apr 2023. doi: 10.1103/PhysRevB.107.144431. URL <https://link.aps.org/doi/10.1103/PhysRevB.107.144431>.
- [40] Yifan Wang, Weizhi Yan, Nikolai Kuznetsov, Lukáš Flajšman, Huajun Qin, and Sebastiaan van Dijken. Spin-wave diffraction, caustic beam emission, and talbot carpets in a yttrium iron garnet film with magnonic fabry-perot resonator gratings. *Phys. Rev. Appl.*, 22:014038, Jul 2024. doi: 10.1103/PhysRevApplied.22.014038. URL <https://link.aps.org/doi/10.1103/PhysRevApplied.22.014038>.

List of Figures

1	A micromagnetic simulation showing the Fourier transform property of a spin-wave.	3
2	Simulated magnonic inverse design demultiplexer.	4
3	Device schematics and simulated results for a spin-wave-based neural network.	5
4	Device scheme and experimental results for the device implemented in [15].	6
1.1	Schematic illustration of the precession of the magnetization vector \mathbf{M}	15
1.2	Schematic representation of a spin-wave in a 1D chain of magnetic moments.	18
1.3	Dispersion diagram for backward volume spin-waves with $\omega_0/\omega_M = 0.5$	22
1.4	Dispersion relations for the three types of magnetostatic spin-waves and the required geometries to excite them.	23
1.5	Excitation efficiency and diagram for a CPW antenna.	25
1.6	Excitation efficiency and diagram for a microstrip antenna.	26
1.7	Schematic diagram of a Recurrent Neural Network.	27
1.8	Simulation of a RNN implementation using forward volume spin-waves.	30
1.9	3D schematic of a PMUT device.	32
1.10	Microscope image of the top side of the PMUT array chip from STMicroelectronics.	33
2.1	Lithography, deposition and lift-off technique with positive resist and LOR schematic.	37
2.2	Picture of the MLA100 maskless aligner in Polifab's yellow room.	37
2.3	Magnetron sputtering schematic.	39
2.4	Picture of the Leybold LH Z400 magnetron sputtering system in Polifab's deposition area.	39
2.5	Picture of the flip-chip bonder available in electronics department's cleanroom, at Politecnico di Milano.	41
2.6	Picture of the Filmetrics Profilm3D optical profilometer, available at Polifab.	42
2.7	Example of time-gating procedure applied on a N1 transmission spectrum.	45
2.8	N0 and N1 mounted on their respective PCBs.	47

2.9	A pictorial representation of the SPSA algorithm.	50
2.10	The flowchart of the implemented Direct Search algorithm.	51
3.1	Design of the N0 device, RF side (back side) perspective.	56
3.2	Design of the N1 device, RF side (back side) perspective.	57
3.3	Pictures of the main problems encountered during lithography.	59
3.4	Overview and zoomed in pictures of the N1 chip after development.	60
3.5	N1 pictures after lift-off.	61
3.6	Microscope pictures of N0 after flip-chip.	62
3.7	N1 after fabrication.	63
4.1	Simulated stray field distribution and YIG magnetization.	67
4.2	Results of a micromagnetic simulation performed with a 100 μm side square magnet positioned 5 μm above the YIG film.	69
4.3	Simulated effect of magnet height on spin-waves propagation.	70
4.4	Simulated effect of excitation frequency on spin-waves propagation.	72
4.5	Comparison of the simulated square and round magnet geometries.	74
4.6	Measured MEMS displacement as a function of the applied voltage for the N1 device.	75
4.7	Raw transmission spectrum of N0 compared with the calculated antenna excitation efficiency.	76
4.8	Calculated dispersion relation and group velocities for the first three volume modes at an applied field of 85 mT.	78
4.9	Optical microscope image of the two antennas on N0 enclosing a 200 μm magnet.	79
4.10	Time-gating analysis of N0 spectra for the four device geometries.	81
4.11	Measured transmission spectra for the four N0 geometries for different applied voltages of 0 V to 60 V and a bias magnetic field of 92 mT.	83
4.12	Measured phase and amplitude spectra for the 200 μm square magnet of N0.	84
4.13	Comparison between experimental and simulated spin-wave transmission as a function of magnet–film distance on N0.	86
5.1	Time-gating analysis of N1 spectra for the four output antennas.	88
5.2	Objective function definition to maximize the received S_{21} signal in the frequency band from 4.3 GHz to 4.5 GHz.	90
5.3	Spectra comparison between the maximization training result and the configuration with 60 V applied to all MEMS.	91
5.4	Final voltage configuration of the maximization training.	92

5.5	Reference and final spectra for the maximization DS trainings starting from 0 V and random configurations.	93
5.6	Comparison between the results of the DS training started from a 0 V configuration and from a random configuration.	94
5.7	Final voltage maps for the maximization DS trainings starting from 0 V and random configurations.	95
5.8	Standard deviation map of the final voltages for all DS maximization trainings.	96
5.9	Reference and final spectra for the SPSA maximization training started from a 0 V configuration.	96
5.10	Final voltage and standard deviation maps for the SPSA maximization trainings.	97
5.11	Comparison between the results of the DS and SPSA algorithms when tasked with the maximization of the signal.	98
5.12	Comparison between objectives and voltage configurations evolutions using DS and SPSA.	99
5.13	Reference and final spectra for the minimization training using DS starting from a 0 V configuration.	100
5.14	Average voltage maps for the maximization and minimization trainings using Direct Search.	101
5.15	Reference and final spectra obtained from the mixed minimization–maximization trainings using the DS algorithm.	102

