LTH
FACULTY OF
ENGINEERING

Design and Measurement of a 2 – 4.4 GHz Tunable Bandstop Filter for Interference Rejection

Harald Förare

Master's Thesis, 30 hp
Master of Science in Engineering, Electrical Engineering
Course Code: EITM01
Department of Electrical and Information Technology
Faculty of Engineering, LTH
Lund University

Supervisor: Prof. Henrik Sjöland

Examiner: Assoc. Prof. Pietro Andreani

March 2026

Series of Master's theses
LU/LTH-EIT 2026-1110
www.eit.lth.se

Printed in Sweden
Tryckeriet i E-huset, Lund

Abstract

Strong narrowband interference can significantly degrade the sensitivity of wideband receivers by driving front-end circuitry into nonlinear operation. This thesis presents the design, fabrication, and measurement of a 2 – 4.4 GHz electrically tunable varactor-based bandstop filter that can follow and block such narrowband interference before it causes problems in the receiver. It was developed using circuit and electromagnetic co-simulation with parametric optimization before fabrication on a PCB. The measured notch depth is at least 24 dB across the tuning range and exceeds 60 dB at some frequencies. Anti-series varactor placement was used to reduce even-order nonlinearities, achieving harmonic distortion $\mathrm{IIP}_2 > 41$ dBm. Cross-modulation is -70 dBc for a desired signal at 3500 MHz with a 0 dBm in-notch interferer at 2800 MHz. For the same notch, IMD3 $\mathrm{IIP}_3 = 47.3$ dBm at 2600 MHz, and $\mathrm{IIP}_3 = 23.3$ dBm at 3000 MHz. Measured results suggest that the presented tunable bandstop filter can preserve receiver sensitivity in the presence of strong narrowband interference.

Populärvetenskaplig Sammanfattning

Föreställ dig att någon viskar till dig mitt under en rockkoncert. Det är omöjligt att förstå vad de säger, eftersom örat blir överbelastat av den starka musiken och dämpar sin känslighet för att försöka undvika skada. Trots att du lätt hade förstått vad de säger i en lugnare miljö, förlorar du den förmågan på rockkoncerten.

Radiomottagare upplever ett mycket liknande problem, där de förlorar förmågan att lyssna på svaga signaler samtidigt som det finns starka störningar.

På grund av begränsningar i elektroniken i en radiomottagare kan den bara förstärka signaler proportionellt upp till en viss nivå. Om en signal är starkare än vad mottagaren klarar av börjar topparna av signalen plattas till. Detta leder till en rad problem som slutligen resulterar i att informationen i signalen går förlorad. Det är dessutom inte bara den starka signalen på en frekvens som uppfattas sämre, utan många andra signaler på andra frekvenser påverkas också.

För att undvika att detta sker behöver mottagaren dämpa sin känslighet så att ingen signal plattas till och orsakar problem. Nackdelen med detta är att redan svaga signaler på alla frekvenser går förlorade under mottagarens brusnivå när känsligheten sjunker. Det blir sällan lättare att förstå viskningen på rockkoncerten bara för att du stoppar i öronproppar.

En stark störning på bara en frekvens har alltså förmågan att försämra känsligheten för alla andra frekvenser som radiomottagaren annars hade kunnat ta emot. Detta är ett nuvarande problem som dessutom kommer att förvärras i framtiden då radiomottagare blir allt mer bredbandiga, och störningar blir allt mer vanliga.

För att lösa detta har jag byggt ett bandstoppfilter. Det är en typ av filter som blockerar ett litet intervall av frekvenser samtidigt som det släpper igenom alla

andra. Syftet med detta är att dämpa styrkan på den starka störningen så att mottagaren kan undvika att behöva sänka sin förstärkning. Då kan mottagaren bevara sin känslighet för svaga men önskade signaler på alla andra frekvenser.

Mitt filter är också stämbart vilket innebär att det är möjligt att elektroniskt styra vilken frekvens som blockeras. Detta är mycket användbart i t.ex. militära sammanhang, där fienden absolut inte meddelar i förväg vilken frekvens de kommer att använda för att störa ut dina radiomottagare.

För samhällskritiska funktioner, såsom räddningstjänst och försvar, är tillförlitlig radiokommunikation avgörande. Uppmätta resultat från detta arbetet tyder på att denna typ av stämbara bandstoppfilter kan bidra till mer robusta radiosystem till störningsfyllda miljöer.

Acknowledgements

I want to thank Prof. Henrik Sjöland, who has been my main supervisor during this project, for excellent technical discussions and helpful guidance.

The people at Evolved Aerospace funded my prototype and provided relevant context for my work. Dr. Mikael Taveniku and Arthur Zargarian contributed their time and expertise to discussions on the scope and context of the project, and Arthur also served as company supervisor. Marianna Taveniku offered to help with logistics and handled the formal parts. Pia Arvstrand kindly and expediently ordered PCBs and components.

Sirvan Abdollah Poor was very helpful during the lab phase and granted access to a desk, instruments, and equipment in the EIT research lab at LTH. He also kindly volunteered his time to discuss measurement techniques. Prof. Fredrik Tufvesson loaned me an SDR on short notice, which made it possible for me to perform the final measurements in the lab. Finally, I want to thank Assoc. Prof. Pietro Andreani for being my examiner.

Urban Carlson has been a good mentor and a friend for many years. He has allowed me to develop both as a person and professionally in his companies. Urban introduced me to Evolved Aerospace, which sparked this entire collaboration.

I am ever grateful for the support my family has given, especially my father and girlfriend, in becoming an engineer.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

As technology advances, the electromagnetic spectrum becomes increasingly contested [1]. Modern radio systems are often expected to work in environments where strong intentional or unintentional electromagnetic interference is present [2]. Such situations can occur both in the civilian world, where receivers could be located close to, for example, radar systems, cell towers or wireless networks, and also in military situations where the enemy deliberately tries to jam your systems [3].

All radio receivers have limited dynamic range. In the presence of strong interference, the receiver has to reduce its gain to prevent overload [4]. This reduces sensitivity at all frequencies, even though the interference might only occupy a small part of the total bandwidth.

This presents an increasing challenge as receivers become more wideband, and are required to listen to ever weaker signals. Such radio systems are more easily impacted by interference in only a small part of their operating bandwidth. Furthermore, as platforms shrink, there is greater emphasis put on reducing the size, weight, and cost of RF components [5]. It is therefore of significant interest to do some early adaptive filtering in an effort to prevent strong interference from ever reaching the receiver in the first place, in a cheap and compact way.

This work implements an electrically tunable bandstop filter that can be used early in a wideband receiver front-end chain to lower the power of narrowband interference, removing the need for gain reduction, and thus preserving sensitivity for all other frequencies in the operating bandwidth. Figure 1.1 illustrates the problem and how this tunable bandstop filter would solve it. Figure 1.2 shows where the tunable bandstop filter could be integrated as a part of a bigger receiver system.

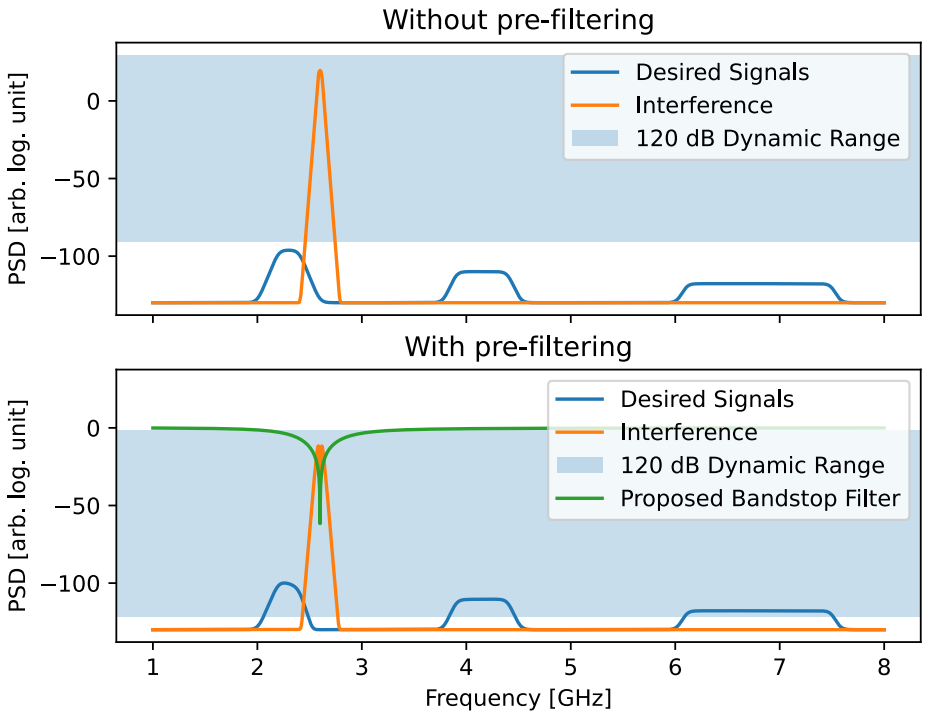

Figure 1.1: Problem illustration

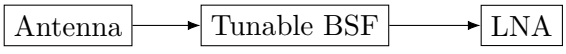

Figure 1.2: Location of tunable bandstop filter in the greater receiver system signal path

1.1 Contributions

- Designed and measured a varactor-based tunable bandstop filter based on a parallel-resonator architecture, achieving a 2 – 4.4 GHz tuning range with notch depth of at least 24 dB for all tunable frequencies, and more than 60 dB for some.
- Minimized even-order nonlinear distortion by using an anti-series varactor placement, resulting in a cross-compression that was too low to measure, and harmonic distortion $\mathrm{IIP}_2 > 41$ dBm for an in-notch interferer at 2.8 GHz.
- Achieved a cross-modulation of -70 dBc at 3.5 GHz with a 0 dBm in-notch interferer at 2.8 GHz, and connected this to third-order distortion analysis via IMD3 measurements. Discussed the trade-off between tuning range and linearity for varactor based designs, and suggested possible options for optimizing the solution towards either end.
- Investigated deviations in transfer characteristics between simulations and reality by reproducing measured deviations in a second round of simulations, narrowing down focus to specifics in the decoupling network.
- Completed an entire RF development cycle, from idea to measurement, demonstrating a solution to the problem of strong narrowband interference degrading receiver performance at all frequencies even when occupying only a small part of the total bandwidth.
- Suggested changes in the design to make an improved next revision based on simulated and measured results.

1.2 Scope & Limitations

This thesis exercises a complete RF hardware development cycle. It focuses on design, optimization, manufacture and verification of a single system block, and due to scope constraints does not go deeply into how this type of solution will be controlled, calibrated or integrated into a wider system.

Chapter 2

Theory & Background

This chapter will give the necessary theoretical background needed to understand why certain design decisions were made and how the results and discussion should be interpreted.

2.1 RF Fundamentals

Radio-frequency (RF) electronics is the study and practice of electronics when frequencies get high enough that the wavelength of the signal starts to be comparable in size to that of the circuit. In RF-electronics, due to the finite propagation speed of electromagnetic waves, circuit nodes often cannot be assumed to have the same voltage everywhere across their physical length. In addition to tools and concepts relating to low frequency analysis, new ones adapted to the wave nature of electromagnetism are introduced.

2.1.1 Characteristic impedance Z_0

For any technological device to do work, it has to transmit power. When electromagnetic waves as used in electronics transmit power, they do so along the Poynting vector $\boldsymbol{S} = \boldsymbol{E} \times \boldsymbol{H}$ [$\mathrm{W\,m^{-2}}$], where $\boldsymbol{E}$ [$\mathrm{V\,m^{-1}}$] is the electric field and $\boldsymbol{H}$ [$\mathrm{A\,m^{-1}}$] is the magnetic field intensity. The ratio between the power-propagating components of the $\boldsymbol{E}$- and $\boldsymbol{H}$ fields is the characteristic impedance Z_0. In general electrodynamics, the definition is quite technical and dependent on field structure. See [6]. For RF design, it is more practical to think of transmission lines as carrying a forward-travelling voltage wave V^+ and current wave I^+. Characteristic impedance Z_0 is then [6]:

$$Z_0 = \frac{V^+}{I^+} \tag{2.1}$$

2.1.2 Reflection coefficient Γ & S-parameters

To practically work with the varying impedances throughout a circuit and over different frequencies, it is useful to transform the impedance Z with the characteristic impedance Z_0 in the following way, to instead work with the new quantity reflection coefficient Γ:

$$\Gamma = \frac{Z - Z_0}{Z + Z_0} \tag{2.2}$$

The reflection coefficient Γ equivalently represents the complex ratio between the reflected voltage wave and the incoming voltage wave at any impedance boundary, capturing both amplitude ratio and phase difference. Note that since the impedance often varies with frequency, so does the reflection coefficient.

It is useful to extend the concept of complex ratios between voltage waves to model bigger linear RF devices. Given a device with n ports, each port p will have one incident wave a_p travelling into the device, and one reflected wave b_p travelling out of the device. Each combination of ports can be given an S-parameter, such that $S_{ij} = \frac{b_i}{a_j}$. The S-parameters for this n-port can then be assembled into an S-parameter matrix $\boldsymbol{S}$:

$$\boldsymbol{S} = \begin{bmatrix} S_{11} & S_{12} & S_{13} & \cdots & S_{1n} \\ S_{21} & S_{22} & S_{23} & \cdots & S_{2n} \\ S_{31} & S_{32} & S_{33} & \cdots & S_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ S_{n1} & S_{n2} & S_{n3} & \cdots & S_{nn} \end{bmatrix} \tag{2.3}$$

$\boldsymbol{S}$ completely captures how a linear RF device behaves. Some intuition can be gained from the S-parameter matrix. To read this matrix, it is assumed that every port is terminated with the characteristic impedance Z_0, otherwise some linear transformation has to be done to make it so, which is always possible.

For the very common case of a two-port, such as the filter we are designing, each S-parameter has a concrete meaning. S_{11} and S_{22} are the reflection coefficients for the input and output respectively. They describe what the connected components, such as antenna and LNA, will "see" when looking into the filter (assuming its interfaces are impedance matched). S_{21} might be the most interesting, as it is the

forward-transmission of the device. For our filter, ideally $|S_{21}|$ should become very small at the rejected frequency, and remain close to 1 everywhere else. S_{12} is the reverse-transmission.

2.1.3 Logarithmic units

Logarithmic units are very convenient for many tasks in RF-electronics. They simplify calculations by reducing multiplication to addition and improve intuition when dealing with very large dynamic ranges. The common logarithmic units that are being used in this thesis are [7]:

dB: General logarithmic ratio.

$$P\ [\mathrm{dB}] = 10\log_{10}\left(\frac{P}{P_{\mathrm{ref}}}\right) \tag{2.4}$$

$$= 20\log_{10}\left(\frac{V}{V_{\mathrm{ref}}}\right), \quad \left(\begin{smallmatrix}\text{for a matched system,}\\ \text{since } P\propto V^2\end{smallmatrix}\right) \tag{2.5}$$

dBm: Logarithmic power ratio compared to 1 mW.

$$P\ [\mathrm{dBm}] = 10\log_{10}\left(\frac{P\ [\mathrm{mW}]}{1\ [\mathrm{mW}]}\right) \tag{2.6}$$

dBc: Logarithmic power ratio between some tone and some reference carrier.

$$P\ [\mathrm{dBc}] = 10\log_{10}\left(\frac{P_{\mathrm{tone}}\ [\text{any linear unit}]}{P_{\mathrm{carrier}}\ [\text{same linear unit}]}\right) \tag{2.7}$$

2.1.4 Transmission lines

Since electromagnetic waves cannot propagate instantly, a new passive component is available in RF-electronics: The transmission line. A transmission line can be any type of structure that supports the propagation of electromagnetic waves. The wave propagating in the transmission line does so with some characteristic impedance, phase velocity and loss. A transmission line is defined by its characteristic impedance Z_0 [Ω], length l [m] and propagation constant $\gamma = \alpha + j\beta$, where α [$\mathrm{Np\,m^{-1}}$] is the attenuation constant and β [$\mathrm{rad\,m^{-1}}$] is the phase constant. Given a load Z_L is connected to via such a transmission line, the impedance looking into the line Z_{in} is the load impedance Z_L transformed:

$$Z_{in} = Z_0\frac{Z_L + Z_0\tanh(\gamma l)}{Z_0 + Z_L\tanh(\gamma l)} \tag{2.8}$$

A lossless transmission line of length x can be analyzed as an infinite series of lumped components [7]. This lumped component model is shown in Figure 2.1:

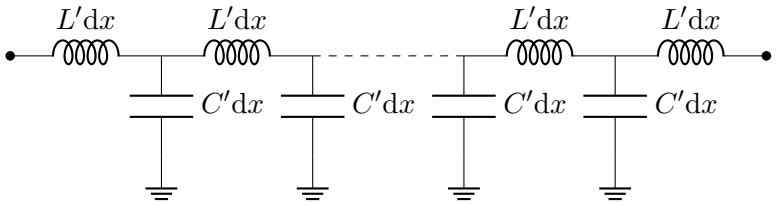

Figure 2.1: Lumped component model of a lossless transmission line

It has the following characteristic impedance Z_0 and phase velocity v_p [7]:

$$Z_0 = \sqrt{\frac{L'}{C'}} \tag{2.9}$$

$$v_p = \frac{1}{\sqrt{L'C'}} \tag{2.10}$$

Phase velocity is related to the speed of light c and the effective substrate permittivity- ϵ_r and permeability μ_r [7]:

$$v_p = \frac{c}{\sqrt{\mu_r \epsilon_r}} \tag{2.11}$$

For non-magnetic materials, $\mu_r = 1$:

$$v_p = \frac{c}{\sqrt{\epsilon_r}} \tag{2.12}$$

$$\implies \begin{aligned} L' &= \frac{Z_0\sqrt{\epsilon_r}}{c} \\ C' &= \frac{\sqrt{\epsilon_r}}{cZ_0} \end{aligned} \tag{2.13}$$

Depending on terminating impedance, short transmission lines can be used to implement approximate inductors and capacitors. Inductance adds in series, and capacitance in parallel. If the transmission line is shorted on one end, any extra shunt capacitance barely makes a difference to the total impedance since the voltage is already kept low. Conversely, for a transmission line terminated in an open

circuit, extra series inductance does little as there is no path for current to flow. These approximations hold for small electrical lengths, as the phase shift when it gets longer causes impedance to vary along the length. Short transmission lines of length x can approximate the following [6]:

$$L = x\frac{Z_0\sqrt{\epsilon_r}}{c} \quad \text{Inductor when terminated with a low impedance} \tag{2.14}$$

$$C = x\frac{\sqrt{\epsilon_r}}{cZ_0} \quad \text{Capacitor when terminated with a high impedance} \tag{2.15}$$

2.2 Simple Filters

The simplest types of band-discriminatory filters require at least one resonator. A resonator presents a fairly constant impedance at most frequencies, with a widely different impedance for one small band. Since impedance differences cause reflections, as presented in Section 2.1.2, resonators can be used in circuits to route power to different locations for different frequencies. This thesis implements a bandstop filter, for which this kind of frequency discrimination is essential. The two simplest types of bandstop filters in an electrical context are shown in Figure 2.2.

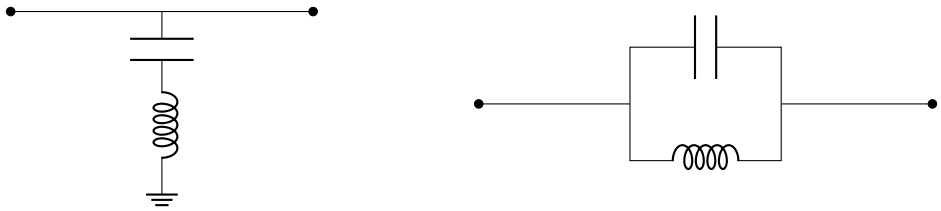

(a) Shunt-connected series LC resonator **(b)** Series-connected parallel LC resonator

Figure 2.2: Fundamental bandstop filters

The resonance frequency f_c of both a parallel- and series resonator is given by Eq. 2.16 [7]. It is apparent that the resonance frequency can be tuned by changing either the inductance or the capacitance.

$$f_c = \frac{1}{2\pi\sqrt{LC}} \tag{2.16}$$

In the ideal case, both filters in Figure 2.2 cause no power to be transmitted at the resonance frequency. The shunt-connection shorts out the RF path, and the

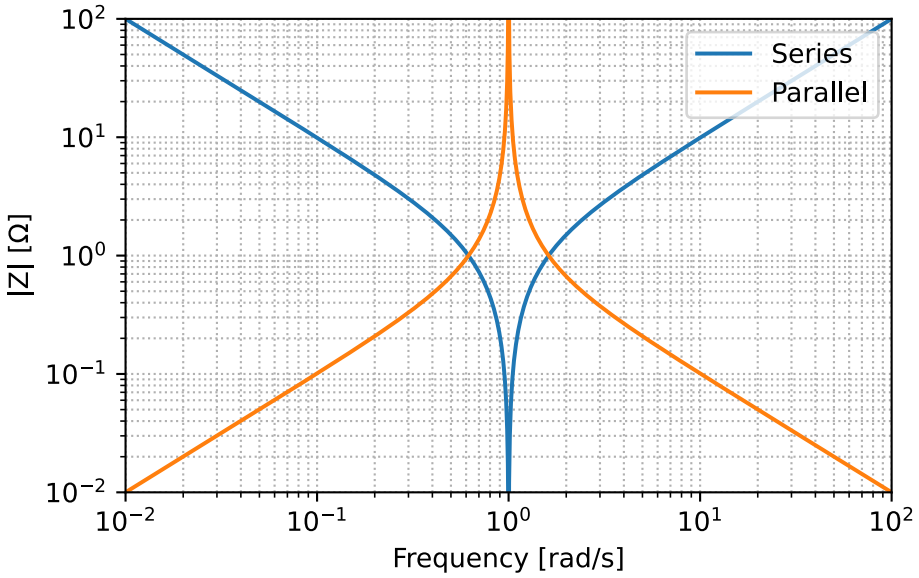

Figure 2.3: Impedance of normalized series- and parallel LC resonators over frequency

series connection presents an open circuit [7]. Figure 2.3 shows the impedance magnitude for a series- and parallel LC resonator, with $L = 1\,\mathrm{H}$ and $C = 1\,\mathrm{F}$.

Focusing on the parallel resonator seen in Figure 2.2b, its transmission with a normalized resonance frequency and different ratios of inductance to capacitance is shown in Figure 2.4.

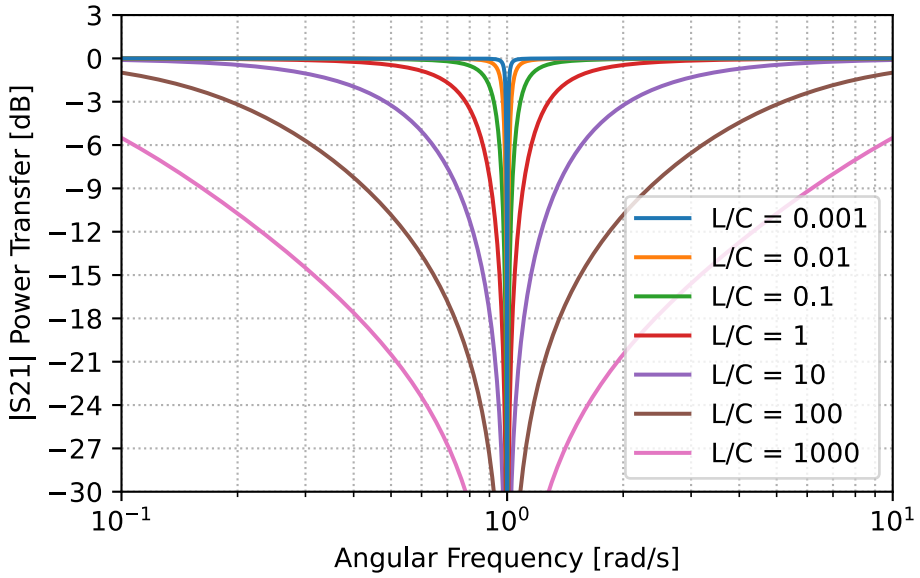

Figure 2.4: Transmission for different ratios between inductance and capacitance for a device like that in Figure 2.2b

2.2.1 Voltage & current gain

To support the boundary condition of zero or infinite impedance at resonance for a series- or parallel resonator respectively at the same time as the resonator is storing energy, a standing wave must exist in the resonator. In the case of the series resonator, the solution is a voltage wave with nodes at the boundaries of the device, since impedance is zero at resonance. By this setup, there is an anti-node somewhere inside the resonator with a maximum of voltage. It becomes apparent that even though the voltage is zero across the entire resonator, it can become very large over individual components inside the resonator. An analogous argument holds for the parallel resonator but with a standing current wave instead due to its very high impedance at resonance.

2.2.2 Lossy behavior

An ideal resonator only stores power, but a real resonator always has some loss each cycle. This loss can be due to many different reasons such as resistance or radiation. Loss in a resonator is usually quantified as the Q-factor, where a higher Q-factor is lower loss. Higher Q implies a lower bandwidth. See [7] for a thorough derivation. Half-power bandwidth Δf and Q-factor are related via Eq. 2.17.

$$Q = \frac{f_0}{\Delta f} \tag{2.17}$$

The internal waves described in Section 2.2.1 are limited by Q when driven at resonance by some external excitation [7]:

$$V_{\text{internal}} \approx Q\, V_{\text{excitation}} \tag{2.18}$$
$$I_{\text{internal}} \approx Q\, I_{\text{excitation}} \tag{2.19}$$

2.3 Varactor Diodes

As is apparent from Eq. 2.16, to get a tunable resonator, either inductance or capacitance need to be variable. One way to get tunable capacitance is to use varactor diodes.

Varactor diodes, sometimes called varicap diodes, are diodes optimized to be used as variable capacitors at RF frequencies. Their capacitance depends on the voltage over the PN-junction.

Figure 2.5 shows an equivalent model for a packaged varactor diode [8].

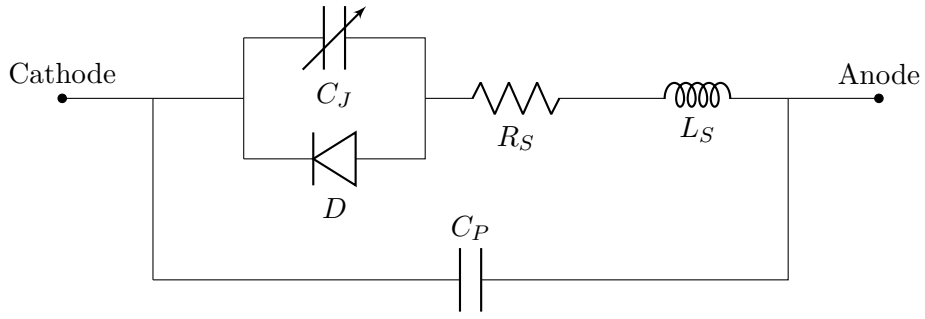

Figure 2.5: Equivalent model for a packaged varactor diode. Redrawn from [8]

The instantaneous junction capacitance C_J of a varactor diode is dependent on the applied reverse junction voltage V_R and is modelled by Eq. 2.20, where C_{JO} is the zero-bias junction capacitance, V_J is the zero-bias junction voltage, and M is a grading coefficient, typically $M = 0.5$ [8].

$$C_J\left(V_R\right) = \frac{C_{JO}}{\left(1 + \frac{V_R}{V_J}\right)^M} \tag{2.20}$$

Out of the modelled series resistance R_S, most of it is due to the semiconductor bulk resistance, that is the portion of semiconductor material that is not in the depletion region [9]. This part of R_S varies with reverse-bias voltage V_R as it changes the width of the depletion region, which in turn determines the length of resistive semiconductor bulk-material that is put in series. A higher reverse bias V_R reduces the resistive loss inside the varactor.

2.4 Microstrip Lines

A microstrip line is one way to implement a transmission line on a PCB. It consists of a lone PCB-trace on top of a dielectric over a ground plane. A widely accepted model for its characteristic impedance as a function of geometry and dielectric properties is the Hammerstad-Jensen model [7,10]. Simplifying yields the following partial asymptotic relationships, with smooth blending in between:

$$\begin{cases} Z_0 \propto \ln\left(\frac{h}{W}\right) & \text{when } h \gg W, \\ Z_0 \propto \frac{h}{W} & \text{when } h \ll W \end{cases} \tag{2.21}$$

The most basic design intuition gained from Eq. 2.21 is that in order to get a higher characteristic impedance, either the dielectric thickness has to increase or the trace width has to decrease.

2.4.1 Skin depth

As frequency increases, currents get confined to the surface of a conductor. This is quantified with skin depth δ, where at a depth δ from the surface, the current density will decay to $e^{-1} \approx 36.8\,\%$ of its value at the surface in an exponential fashion. For a conductor at frequency f with permeability μ and conductivity σ, skin depth δ is [7]:

$$\delta = \sqrt{\frac{2}{2\pi f \mu \sigma}} \tag{2.22}$$

2.5 Nonlinear Systems Theory

2.5.1 LTI systems

Linear time-invariant (LTI) systems are widely taught in undergraduate programs in electrical engineering. They are useful because their linear- and time-invariant properties allow for great flexibility in mathematical analysis. Unfortunately, due to their expressive constraints, they cannot ever model any system where signals of different frequencies interact with each other. This makes them unable to capture any type of distortion of the signal, only allowing amplitude- and phase changes per frequency, never new frequency generation. S-parameters as presented in Section 2.1.2 represent electrical LTI-systems completely in the frequency domain [7].

2.5.2 New frequencies from nonlinearities

Consider a simple univariate memoryless system with the output being a function of the current input only: $y(t) = f(x(t))$. Given the Taylor series theorem, $f(x)$ can be rewritten as a polynomial near a point x given f is sufficiently smooth [4, 11, 12]:

$$y(t) = \underbrace{a_1 x(t)}_{\substack{\text{linear} \\ \text{term}}} + \underbrace{a_0 + a_2 x(t)^2 + a_3 x(t)^3 + \cdots + a_n x(t)^n}_{\text{nonlinear terms}} \tag{2.23}$$

Nonlinear systems, like those in Eq. 2.23 can give rise to output frequencies that are not present in the input. Let:

$$x(t) = \cos(\omega t) \tag{2.24}$$

Consider one term in Eq. 2.23 when the system is excited by a pure tone:

$$x^n(t) = \cos(\omega t)^n \tag{2.25}$$

Due to Euler's formula:

$$\cos(\omega t) = \frac{e^{j\omega t} + e^{-j\omega t}}{2} \tag{2.26}$$

Substituting (2.26) into (2.25):

$$x^n(t) = \left[\frac{e^{j\omega t} + e^{-j\omega t}}{2}\right]^n$$
$$x^n(t) = \frac{1}{2^n}\left(e^{j\omega t} + e^{-j\omega t}\right)^n \tag{2.27}$$

Expand using binomial coefficients:

$$\left(e^{j\omega t} + e^{-j\omega t}\right)^n = \sum_{m=0}^{n} \binom{n}{m} \left(e^{j\omega t}\right)^m \left(e^{-j\omega t}\right)^{n-m} \tag{2.28}$$

Combine exponents:

$$\begin{aligned}\left(e^{j\omega t}\right)^m \left(e^{-j\omega t}\right)^{n-m} &= e^{j(m-(n-m))\omega t} \\ &= e^{j(2m-n)\omega t}\end{aligned} \tag{2.29}$$

Eq. 2.25 is re-written using results from (2.27), (2.28) and (2.29):

$$x^n(t) = \frac{1}{2^n} \sum_{m=0}^{n} \binom{n}{m} e^{j(2m-n)\omega t} \tag{2.30}$$

The general term shown in Eq. 2.30 contains possibly multiple complex sinusoids with frequencies dependent on order and shown in Table 2.1.

Term Order	**Harmonic Frequency**
1	$1f_0$
2	$2f_0$
3	$1f_0$, $3f_0$
4	$2f_0$, $4f_0$
5	$1f_0$, $3f_0$, $5f_0$
6	$2f_0$, $4f_0$, $6f_0$
n	$(2m-n)f_0$

Table 2.1: Possible harmonics for each term order

The system output $y(t)$ (Eq. 2.23) is the sum of such terms as shown in Eq. 2.30. Since any term with $n > 1$ gives rise to output frequencies not present in the input, any system with nonlinear terms will create frequencies in the output that were not present in the input.

In general, the above shown property of new frequency generation does not disappear when system complexity increases, such as with multiple dimensions or memory.

2.5.3 Multiple tones

Using a very similar derivation as in Section 2.5.2, but with a multi-tone excitation such as, for example, the following two tones:

$$x(t) = \cos(\omega_1 t) + \cos(\omega_2 t) \tag{2.31}$$

It can be shown that intermodulation (IM) products are created at frequencies ω, where m and n are integers corresponding to the orders of nonlinearities for each tone [12]:

$$\omega = |m\omega_1 \pm n\omega_2| \tag{2.32}$$

This phenomenon extends to any number of tones, but the analysis becomes convoluted quickly. In general [12]:

$$\begin{aligned} x(t) &= \cos(\omega_1 t) + \cos(\omega_2 t) + \cdots + \cos(\omega_k t) \\ \implies \omega &= |m_1\omega_1 \pm m_2\omega_2 \pm \cdots \pm m_k\omega_k| \end{aligned} \tag{2.33}$$

2.5.4 Nonlinear terms' rate of growth

Many measures of nonlinear behavior in RF-systems are defined largely by the rate-of-growth of certain nonlinear output terms in relation to the growth of some input excitation. The terms are convenient to distinguish on a spectrum plot with a logarithmic y-axis since an exponent inside a logarithm simply extracts out in front as a multiplicative factor:

$$\ln(x^n) = n\ln(x) \tag{2.34}$$

A term of order n thus grows at a rate n times that of the excitation.

2.5.5 Memory effects

Electrical systems are in general not memoryless since they very often contain elements that can store energy, and thus save information about the past. Components that cause a system to have memory are for example capacitors, inductors and transmission lines. Memory can also arise from thermal effects, hysteresis and many more mechanisms [12].

Memory effects make nonlinear distortion more difficult to analyze, as outputs depend on more than the present input. In practice this often shows up as for example spacing-dependent IMD3 measurements, where the beat frequency between tones gets stored in the biasing network, or even as heat in the components for low beat frequencies and higher power densities [12].

2.6 Receiver Limitations

When a receiver is subject to strong narrowband interference (NBI), the following performance degradations can occur:

1. **Gain compression:** Strong narrowband interference can drive RF- or IF amplifiers into a nonlinear operating point, causing the small-signal gain to be lower, reducing gain for the desired weak signal and thus lowering receiver sensitivity.
2. **Desensitization (blocking):** The interference causes the automatic gain control (AGC) to reduce receiver gain, making weak desired signals harder to detect.
3. **Intermodulation distortion:** Nonlinear components generate spurious signals from mixing between the interference and other signals, which can fall into the desired channel.
4. **Cross-modulation:** Modulation in the interference transfers to the desired signals, reducing the ability to extract the intended modulation of the desired signal.
5. **Reciprocal mixing:** Phase noise from the local oscillator mixes with the strong interference, raising the effective noise floor near the desired signal. This effect often combines unfavorably with cross-modulation.

6. **Noise floor elevation:** Nonlinear effects can spread interference energy across the spectrum, increasing the apparent noise floor and reducing signal-to-noise ratio.
7. **Finite filter selectivity:** Imperfect RF and IF filters allow some interference energy to leak into the desired signal path.

   ... and more: See [4, 12] for additional effects.

In order to avoid experiencing any of the above issues, the receiver generally has to reduce its gain. This impacts sensitivity and causes it to not be able to successfully receive as weak signals as were possible without the interference present.

2.7 Tools & Methods

This section seeks to briefly explain each tool and method used throughout the thesis, so it becomes apparent what function it serves when encountered in the text below. See citations for more details on any specific one.

2.7.1 Simulators

1. **SciKit-RF:** Python library for frequency-domain network analysis. Only LTI-systems. Good for gluing together different S-parameter models and is very convenient to work with in the Python ecosystem. [13]
2. **QUCS-S:** RF-focused GUI front-end for different simulators using a circuit representation. Can run harmonic-balance, extract S-parameter models, run ngspice and more. [14]
3. **Ansys HFSS:** Full-wave electromagnetic field simulator. Solves Maxwell's equations for arbitrary 3D-geometry. Can extract S-parameter models from geometry. [15]
4. **NGspice:** Numerical circuit simulator. Widely supported and implements multiple simulation domains. Handles nonlinear components defined using industry standard SPICE-models. [16]

2.7.2 Exploration & optimization

1. **Monte-Carlo:** A technique used to understand the behavior of a system by running a large number of randomized experiments. The results of these experiments are then analyzed using statistical methods to gain an understanding of the system. Widely used due to its conceptual simplicity and ability to handle systems with a large number of parameters. [17]

2. **Pareto Fronts:** A technique used to narrow down a large selection of candidates by removing those that are not Pareto-optimal. For each candidate, multiple figure of merits (FOMs) are calculated. Only candidates for which no other candidate is better than or equal to in all FOMs, and strictly better in at least one FOM are kept. Using pareto fronts is helpful because it greatly reduces decision space, only keeping good trade-offs, while also not requiring subjective decisions on the importance of each FOM. [18]

2.7.3 Lab equipment

1. **Vector Network Analyzer (VNA):** Instrument used for measuring S-parameters of a multiport (usually 2-port) device. It works by generating a known stimulus on one of the ports at a time, while simultaneously measuring the incident and reflected waves at all ports. Provides frequency dependent S-parameter matrices for the DUT. [19, 20]
2. **Spectrum Analyzer:** Instrument used for measuring the power of a signal as a function of frequency. In theory, it implements a subset of the functionality of a VNA, but in reality, it does so with much greater sensitivity, accuracy and frequency range for a given cost. It works by mixing the incoming signal with a swept intermediate frequency (IF) tone. The resulting downconverted signal is passed through a selection filter and finally a detector, measuring power within some bandwidth window swept over a bigger range. By using a small bandwidth in the selection filter, called resolution bandwidth (RBW), the power content within a very small range of frequencies is measured. This reduces the total amount of noise present at the cost of measurement time. [19, 20]
3. **Signal Generator:** Instrument used to generate a high quality reference tone, meaning with very low bandwidth and a well controlled power level. Often used in conjunction with a spectrum analyzer to measure how a DUT affects a tone passing through. Some signal generators also support modulating the tone for more complex measurements. [19, 20]
4. **Software Defined Radio (SDR):** A flexible radio system where the main functions such as different gains, local-oscillator frequency (f_{LO}), filters, analog-to-digital- (ADC) and digital-to-analog (DAC) converters can be interfaced to using software. SDRs can contain both transmitter- and receiver hardware. They are useful for prototyping and custom measurement setups due to their flexibility and low cost. Performance is often not as good as that of dedicated lab equipment with a narrower functional focus. [21]
5. **Attenuator:** A passive component that reduces the power of a signal passing through without causing reflections, provided it is terminated with the correct impedances. Good attenuators are very wideband and reduce power by an accurate and equal amount for all frequencies. They should have a constant group delay. Attenuators can be used to isolate pieces of equipment from each other's impedances, presenting a more predictable interface

for components on either end of the attenuator. They can also be used to lower the signal power to reduce nonlinearities in a receiving stage. Since they dissipate signal power, they typically increase system noise. [7, 19, 20]

2.7.4 Lab measurements

1. **S-parameters:** The settings of the VNA are specified. It is then calibrated up to the interface of the DUT (between the ports of the VNA, connected by cables and possibly extra equipment). The VNA then measures the S-parameters of the DUT as a function of frequency. The S-parameters give insights into the impedance behavior of the device. See Section 2.1.2. [19,20]
2. **Harmonic Distortion (HD):** The input of the DUT is excited by a signal generator set to a pure tone with frequency f_0, and the output is measured with a spectrum analyzer. The spectrum analyzer measures the power level at the fundamental and each expected harmonic $\{f_0, 2f_0, 3f_0, \cdots, nf_0\}$. Input tone power is swept. The data can be visualized in multiple different ways. The power of each harmonic compared to the fundamental/carrier, expressed in decibel relative to carrier (dBc), is common. Since harmonics often grow at the rate of their order, as presented in Section 2.5.4, it is also common to present intercept points, where a certain harmonic would overtake the fundamental in power if extrapolated. By looking at the fundamental, compression can also be measured from the same dataset. [4, 19, 20]
3. **Third Order Intermodulation (IMD3):** Two closely spaced- and equal power tones are injected into the DUT, often using two signal generators and a power combiner. As presented in Section 2.5.3, this will excite intermodulation products at $2f_1 - f_2$ and $2f_2 - f_1$ close to the input tones at f_1 and f_2. The level of each tone and product is measured with a spectrum analyzer as input power is swept. IMD3 is commonly reported as the power level where the third order intermodulation products surpass that of the fundamental tones, given as OIP_3 when referring to output power, and IIP_3 when referring to input power. [4, 19, 20]
4. **Cross-Modulation:** A modulated signal and an unmodulated tone are both injected into the DUT at frequencies depending on desired use-case. Using a spectrum analyzer it is evaluated at which level the modulated signal sidebands are transferred to the unmodulated tone on the output. Given as a figure in relation to the unmodulated tone output fundamental (dBc). [4, 19]
5. **Cross-Compression:** One variable power tone and one fixed power weak tone are both injected into the DUT. While the power of the variable power tone is increased, the output power of the weak tone is observed using a spectrum analyzer, looking for a decrease in power indicating cross-compression. [19]

Chapter 3

Method

3.1 Design

This section describes the design methods used to arrive at a manufacturable prototype, and includes motivations on why each specific design choice was made.

3.1.1 Technology & architecture choices

Due to complexity limitations, the design is constrained to a single resonator. It was decided to implement the variable-frequency resonator with varactor diodes as the tunable element due to their following advantages:

1. Tiny size: Allowing a compact implementation which is easier to integrate in small platforms.
2. Rapid tuning speed (frequency agility): Able to change notch frequency quickly to follow erratically moving interference.
3. Simple drive circuitry: A voltage biased device which simplifies bias-network and surrounding circuit design. This reduces project risk.
4. Low cost.
5. Compatibility with RFIC-processes: Future possibility of integration and of interest to current research [22–24] .

There are also integrated resonators such as YIG-devices. They have decade-sized tuning ranges and excellent performance, but are big, heavy and expensive. [25].

MEMS tunable capacitors are often superior to varactors in terms of Q-factor and linearity [26], but they require more complex integration and are less robust due to their mechanical nature.

3.1.2 Circuit development

Since series-resonators (Figure 2.2a) exhibit voltage gain at resonance, as outlined in Section 2.2.1, and varactors are voltage biased, series-resonators do not have a suitable topology for this type of active device. The large standing voltage wave inside the resonator at resonance would greatly modulate the junction capacitance over the waveform, resulting in large nonlinearities from signals in the notch. This is precisely what we want to avoid.

The other alternative is the parallel-resonator topology, seen in Figure 2.2b. Instead of a voltage gain at resonance, the parallel-resonator topology exhibits a current gain. This also has the potential to modulate capacitance, as the current needs to flow through the semiconductor bulk resistance, causing a voltage. This voltage is however much lower than the voltage due to voltage gain in the series-resonator, and limited to the available exctiting voltage.

Due to the above reason, it was decided to use a series-connected parallel resonator topology for the filter. Figure 3.1 shows the filter with an abstract variable capacitor. Figure 3.2 shows the filter where the abstract variable capacitor has been replaced with a varactor diode.

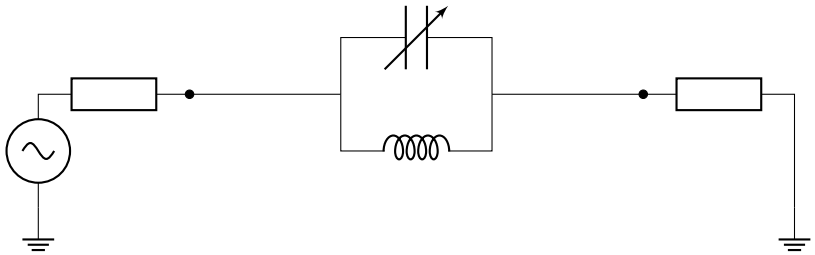

Figure 3.1: Most abstract bandstop filter

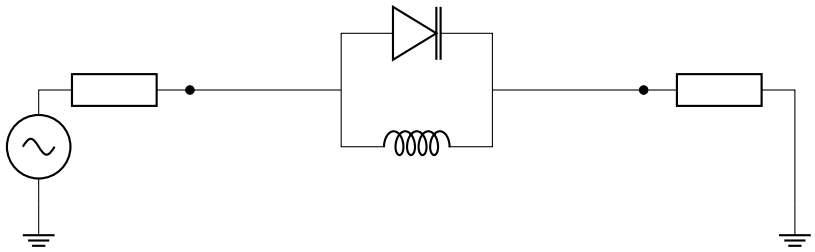

Figure 3.2: Replace abstract variable capacitor with real varactor diode

The voltage-biased nature of the varactor diode, though convenient, also presents challenges. Any voltage over the junction will modulate its capacitance, as described by Eq. 2.20. The varactor can not distinguish DC bias voltage from RF signal voltage, so any RF voltage over the varactor will to some extent modulate the capacitance throughout each signal cycle. This causes nonlinear behavior. Since a primary goal of this design is to reduce nonlinear behavior in a receiver system, and this block appears early in the signal chain, particular care must be taken to minimize the RF-voltage over any varactor diode junction. If this filter would be less linear than the front-end circuitry it precedes, it would most likely limit total system linearity.

Varactor diodes are intended to work in a reverse-biased state. The instantaneous junction capacitance as a function of reverse voltage is described by Eq. 2.20. It is a nonlinear relationship. By using two closely matched varactor diodes, connected in an anti-series configuration. Nonlinear distortion is greatly reduced. The first reason why is that they will share signal voltage quite evenly, but with opposite polarities. This causes the capacitance curves to cancel each other out for smaller signals, as they are locally similar. The second reason is that the symmetry will theoretically fully remove even-order nonlinear terms in the system response, thus removing even-order nonlinear distortion.

Refer to the capacitance as a function of reverse bias voltage in the varactor, shown in Eq. 2.20, as $C(v)$. Consider two identical varactors that are connected in anti-series (series but in opposite directions). They are reverse-biased with the same voltage V_b and the whole group is excited by some signal voltage v over it. Due to symmetry, they will share the voltage v evenly but with opposite polarities. Each varactor's instantaneous capacitance is then:

$$C_1(v) = C\left(V_b + \frac{v}{2}\right)$$
$$C_2(v) = C\left(V_b - \frac{v}{2}\right)$$

The total capacitance is, since they are connected in series, described by the fol-

lowing function:

$$C_{\text{total}}(v) = \frac{C_1(v)C_2(v)}{C_1(v) + C_2(v)} \tag{3.1}$$

We can now handle C_{total} as a single capacitor due to symmetry. To show that its capacitance as a function of voltage is an even function, $C_{\text{total}}(v) = C_{\text{total}}(-v)$ must hold [11]:

$$\begin{aligned} \frac{C_1(v)C_2(v)}{C_1(v) + C_2(v)} &= \frac{C_1(-v)C_2(-v)}{C_1(-v) + C_2(-v)} \\ \implies \frac{C(V_b + \frac{v}{2})C(V_b - \frac{v}{2})}{C(V_b + \frac{v}{2}) + C(V_b - \frac{v}{2})} &= \frac{C(V_b - \frac{v}{2})C(V_b + \frac{v}{2})}{C(V_b - \frac{v}{2}) + C(V_b + \frac{v}{2})} \\ \implies \text{LHS} &= \text{RHS} \end{aligned}$$

$\implies C_{\text{total}}(v)$ is even.

From the definition of capacitance [27]:

$$C_{\text{total}}(v) = \frac{\mathrm{d}q}{\mathrm{d}v} \tag{3.2}$$

Express the current i through the capacitor using the definition of current [27]:

$$i = \frac{\mathrm{d}q}{\mathrm{d}t} \tag{3.3}$$

$$= \frac{\mathrm{d}q}{\mathrm{d}v}\frac{\mathrm{d}v}{\mathrm{d}t} \quad \text{(using the derivative chain rule)} \tag{3.4}$$

$$= C_{\text{total}}(v)\frac{\mathrm{d}v}{\mathrm{d}t} \quad \text{(substituting in Eq. 3.2)} \tag{3.5}$$

Now evaluate the current i when voltage v polarity is reversed:

$$i(-v) = C_{\text{total}}(-v)\frac{\mathrm{d}(-v)}{\mathrm{d}t} \tag{3.6}$$

$$= -C_{\text{total}}(-v)\frac{\mathrm{d}v}{\mathrm{d}t} \quad \text{(extracting sign out of derivative)} \tag{3.7}$$

$$= -C_{\text{total}}(v)\frac{\mathrm{d}v}{\mathrm{d}t} \quad (C_{\text{total}}(v) \text{ is even}) \tag{3.8}$$

$$i(-v) = -i(v) \implies i(v) \text{ is odd!} \tag{3.9}$$

This proves that the current as a function of voltage (admittance) of the anti-series varactor pair is an odd function. If impedance is preferred, the inverse of an odd function is also odd, provided it exists [11].

Since the electrical transfer of the anti-series varactors under a signal is an odd function, its Maclaurin expansion will only have odd-order terms. The orders of terms in the polynomial expansion has a great effect on the spectral content of the transmitted signal, as presented in Section 2.5.2. Nonlinear distortion, especially even-order, is greatly reduced by using two closely matched varactor diodes in an anti-series configuration. Figure 3.3 shows the circuit after changing the single varactor to dual varactors in an anti-series topology.

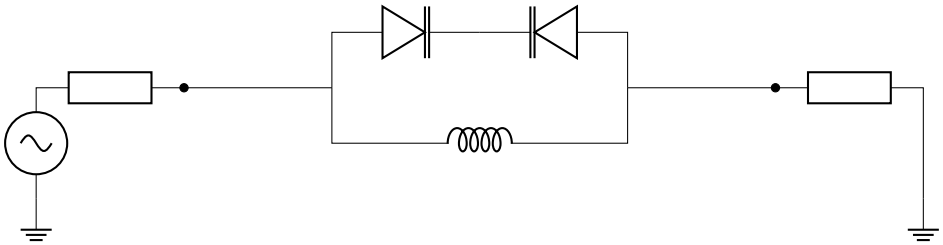

Figure 3.3: Anti-series varactors

As narrowband of a response as possible is desired, and therefore the ratio of inductance to capacitance needs to be minimized, as is evident in Figure 2.4. Above a few hundred MHz, discrete coil inductors become less useful due to various losses and inter-winding capacitance [7]. Loss is bad as it increases the width of the notch. The capacitance between the windings dominates the inductance at our operating frequencies, and causes the coil to not even appear inductive above its self-resonant frequency.

At the operating frequencies of this design it is more practical, or even essential, to realize the inductance using a short transmission line instead of a coil, as presented in Section 2.1.4. This circuit transformation is shown in Figure 3.4.

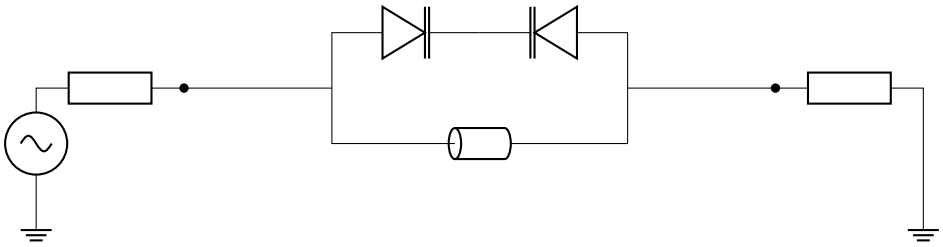

Figure 3.4: Replace inductor with short inductive transmission line

Using anti-series varactor diodes, as introduced in the step illustrated in Figure 3.3, biasing can be elegantly addressed. Biasing is implemented with resistors as they are very wideband, and will have minimal DC current through them due to the voltage-biased nature of the varactors. This circuit eliminates the need for DC-

block capacitors in the RF-path, or any complex bias-tees.

Figure 3.5 shows the addition of biasing resistors, which is the final schematic circuit.

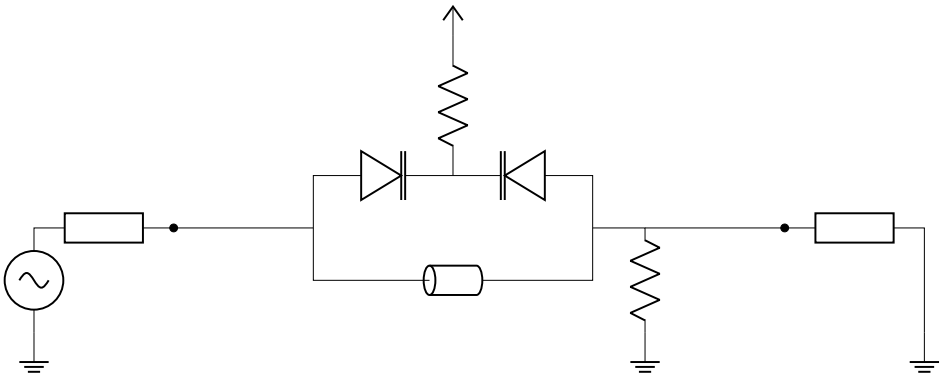

Figure 3.5: Add biasing resistors

3.1.3 Linear circuit-level simulations

Using QUCS-S, the circuit was developed in the linear circuit domain. This is an appropriate level of detail in the early phase of a design project as it enables fast iteration. These simulations are very quick to run, and support interactive tuning of design variables. The basic functionality was investigated using simple linear models of the varactors. The substitution of an inductor to short transmission line, as shown in Figures 3.3 and 3.4, was verified. Figure 3.6 shows a screenshot of QUCS-S taken during this phase of development.

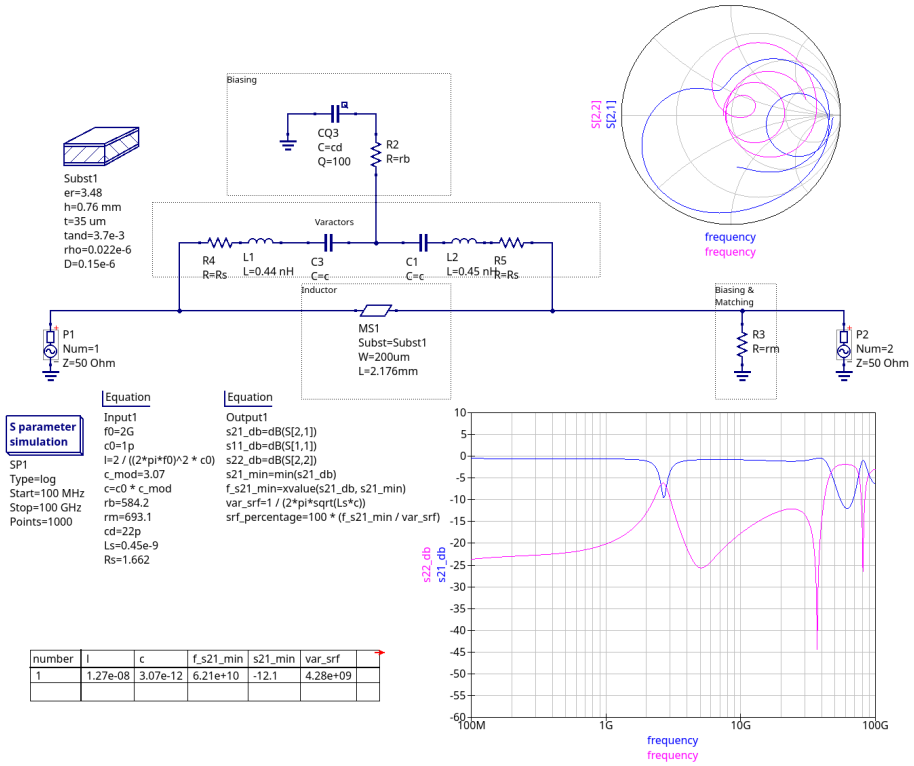

Figure 3.6: QUCS-S Screenshot

3.1.4 EM-simulations for PCB geometry

The device needs to be manufactured on a real PCB. Since certain parts, most importantly the inductive transmission line, depend completely on the actual geometry, an electromagnetic simulation is very useful. Ansys HFSS was used to perform electromagnetic simulations. The RF-part of the PCB was modelled parametrically, so that all geometry is defined by mathematical functions with input parameters, and can be programmatically varied. It was determined that the smallest components suitable for hand soldering are of the 0402-footprint (1 mm $\times$ 0.5 mm), so the design was planned around that size.

The inductive transmission line was modelled as a circle swept along an ellipse. It has uniform width all the way around. An ellipse was chosen because it has very soft curvature and thus likely won't have large impedance discontinuities. Two parameters were varied during EM-simulations: Ellipse ratio, and trace width. Ellipse ratio serves as a monotonic proxy to the length of the transmis-

sion line. Width is interesting mainly because it affects the propagation constant, both impedance and loss. Linear tapers from the microstrip lines to component pads were used due to their modelling simplicity, despite the existence of better-performing tapers, such as Klopfenstein [28].

It is very useful to have fast simulation iterations because it allows for more design options to be tested in any given time, increasing the likelihood of developing a good solution during the project. Ansys HFSS provides EM-simulations that can be very accurate, but they are also rather slow to run. Each simulation of the 3D-geometry takes me about one hour to compute for the full frequency sweep. Ansys artificially restricts the number of processor cores to two for my student license. The number of possible configurations is very large so this slow iteration time presents a challenge. Fortunately, many of the possible configurations use the same PCB-geometry, with variations only in lumped components. Since lumped components are electrically small, they can be extracted from the EM-simulation.

HFSS was set-up to expose not only the input- and output ports, but also lumped ports where the lumped components would be placed. Figure 3.7 shows a screenshot of HFSS during EM-simulations, where each magenta-colored rectangle is an exposed port. The input- and output ports are wave ports, and the pads where components would be located are lumped ports.

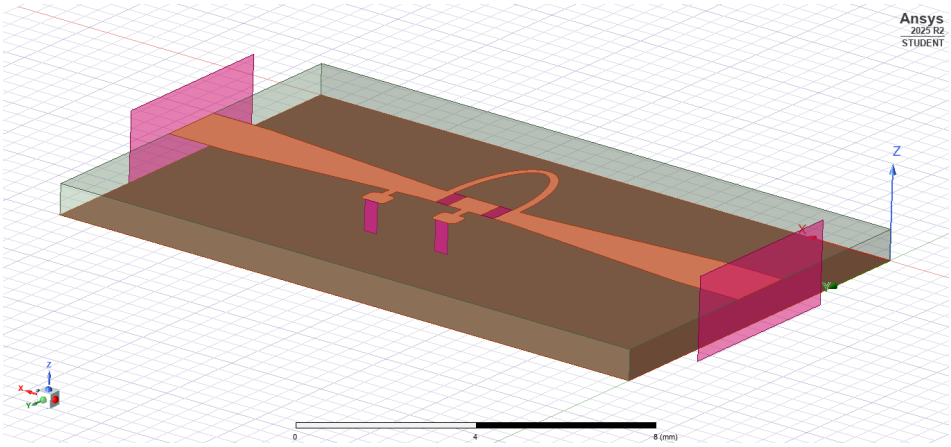

Figure 3.7: HFSS Screenshot

3.1.5 Co-simulation & optimization

With the setup described in Section 3.1.4, each HFSS-simulation generates a 6-port S-parameter model. These 6-ports were terminated inside SciKit-RF with linear models for the lumped components. Each linearly terminated model takes only a fraction of a second to evaluate the performance of, so iteration time is

improved by a factor of around 10 000 compared to doing it directly in HFSS. It is also possible to do co-simulation inside the Ansys suite, but it was decided to do it in Python instead due to tool familiarity and flexibility.

HFSS was used to generate S-parameter models of around 100 different candidate geometries. These were combined with many different variants of varactors and biasing resistors. The resistors for biasing were modelled as a resistance in series with a 300 pH inductance to approximate the package parasitics. The varactors were modelled using the circuit shown in Figure 2.5, but linearized at an operating point. Parameters of many different varactor options were tried. The resistors were sampled from a logarithmic distribution between 10 Ω and 10 kΩ. For each combination of choices for varactor, PCB and biasing resistors, the two-port was calculated for the entire capacitance sweep available in the varactor. This data was then spline-interpolated in the S-parameter domain to get continuous measures of notch frequency. The resulting S-parameters were then analyzed and given a few figures-of-merit (FOMs):

1. Worst notch depth across tuning range minus the average pass-band insertion loss
2. Notch width as average fractional bandwidth throughout the tuning range at a common threshold transmission ratio
3. Tuning range as the ratio between maximum and minimum notch bottom frequency

The dataset was split by varactor choice, and for each varactor choice, every pair of FOMs had its Pareto front calculated. Candidate configurations on the Pareto front were kept. A small software tool was written to explore these Pareto fronts interactively. First the user is presented with scatter-plots for all pairs of FOMs, then they can select one candidate to view the S-parameters. When selected, it is marked on all scatter-plots. Using interactivity makes it easier to gain an understanding of system trade-offs. Figure 3.8 shows a screenshot of this software interface.

From Figure 3.8 it is apparent that some FOMs have uncorrelated relationships. Members of these groups show up along horizontal or vertical lines. The angled- or curved groups show interesting relationships. Each categorical choice of varactor has a clear clustering. Figure 3.9 shows the most interesting relationships discovered at this phase of the project.

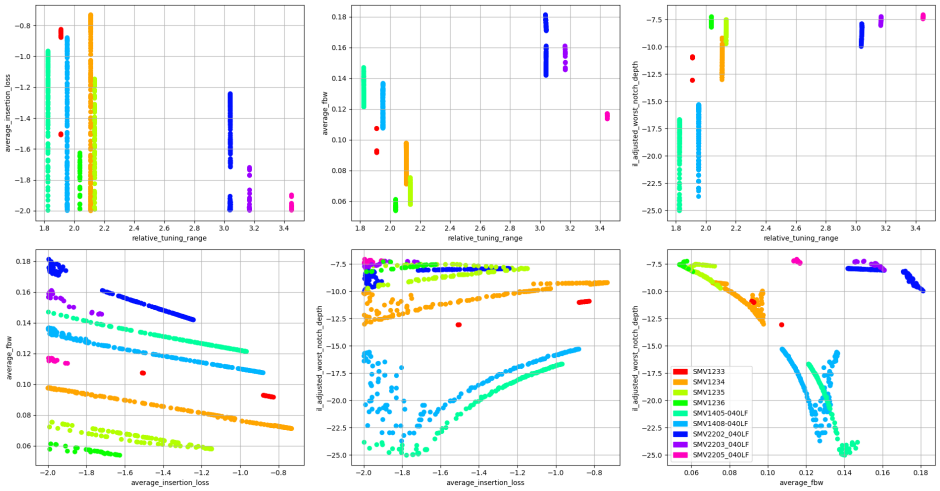

Figure 3.8: Pareto fronts in optimization

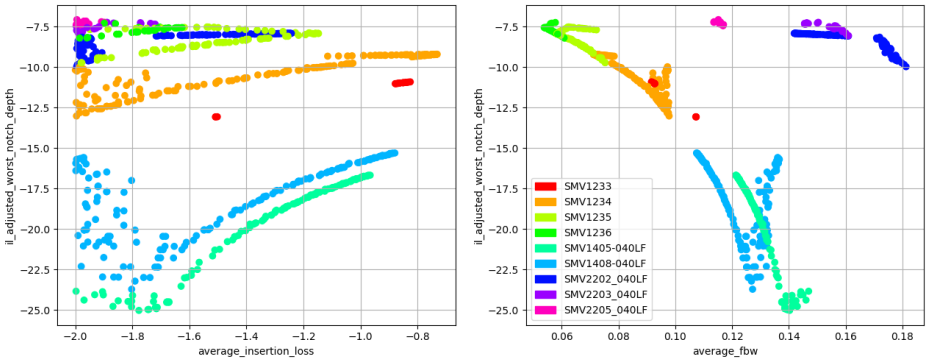

Figure 3.9: Interesting Pareto fronts from Figure 3.8

From Figure 3.9 it is seen that abrupt varactors (SMV14xx-series) achieve a much deeper notch than hyperabrupt varactors (SMV22xx-series): Hyperabrupt varactors optimize for a greater capacitance ratio. By doing so the semiconductor has to allow for more size change in the depletion region, meaning on average that the current has to flow through more resistive semiconductor bulk material. This causes losses and lowers the Q-factor of a filter, worsening the notch depth. They however have a much better tuning range. Unfortunately, the tuning range scales with the square-root of capacitance ratio, as seen in Eq. 2.16, making this an expensive trade-off for larger tuning ranges.

To select a configuration, especially in regards to the varactor choice, some foresight was put into future linearity considerations. More sensitive varactors will

have worse linearity. This is not something that can be entirely avoided by choosing a parallel-resonance topology, as the current through the semiconductor bulk resistance still would cause some voltage. Because of this, in combination with the square-root diminishing returns from increasing capacitance ratio, focus was put on the SMV1405- and SMV1408 abrupt varactors from Skyworks. The SMV1405 was finally chosen due to it having a slightly narrower notch width at the same notch depth in these studies. A favorite candidate was selected and its parameters are shown in Table 3.1.

Varactor	SMV1405	
Cathode bias resistor	261	Ω
Ground bias resistor	6.9	kΩ
Ellipse ratio	2.6	
Elliptical TL width	0.2	mm

Table 3.1: Selected favorite candidate parameters

3.1.6 PCB design

A PCB was designed containing a filter with the parameters from Table 3.1. Rogers RO4350B dielectric was chosen due to its lower loss and more controlled properties compared to something like the cheaper and extremely common FR-4 [29]. The thickness of 0.762 mm was chosen early during EM-simulations due to its higher amount of available inductance per unit length compared to thinner dielectrics, enabling inductor realization in a shorter TL. In hindsight this might not have been the optimal decision, see sections 5.1 and 5.7.

Gold and copper thicknesses were increased slightly from default values as it didn't increase costs very much. Despite thicker copper mostly being overkill loss-wise due to small ($\approx 1 - 10\,\mu m$) skin depth at GHz-frequencies, as presented in Section 2.4.1, thicker metal still produces a nicer and more robust board. Increased gold plating thickness is beneficial, as it is at the surface and has a higher conductivity than copper [7]. Vias were left untented to allow for good soldering with vias placed directly in pads, which decreases parasitic inductance.

A small decoupling network was designed for the biasing feed, with a decoupling capacitor placed after the feed to the resistor. Their common node is fed through a series resistor, which was put in to isolate the circuit slightly from possible differential-mode interference from the bias source. On the other side of that resistor, provisions for another decoupling capacitor as well as a bleed-down resistor were made.

In addition to the filter specified in Table 3.1, one with a lower ellipse ratio and one with a higher ellipse ratio were also added to the PCB. It did not increase costs

and it was unclear during this phase of the project if it would be interesting to look at those in comparison to the selected favorite, so that option was left open. Two through-connected microstrips of identical width as the traces interfacing to the filter, and using identical connectors were also added. One is the same length as the filter structure, and the other is slightly shorter. This was done to enable some characterization of the substrate material, to help determine if later measured behavior depends on the filter or its surroundings.

Figure 3.10 shows the final PCB design. Table 3.2 shows its specifications.

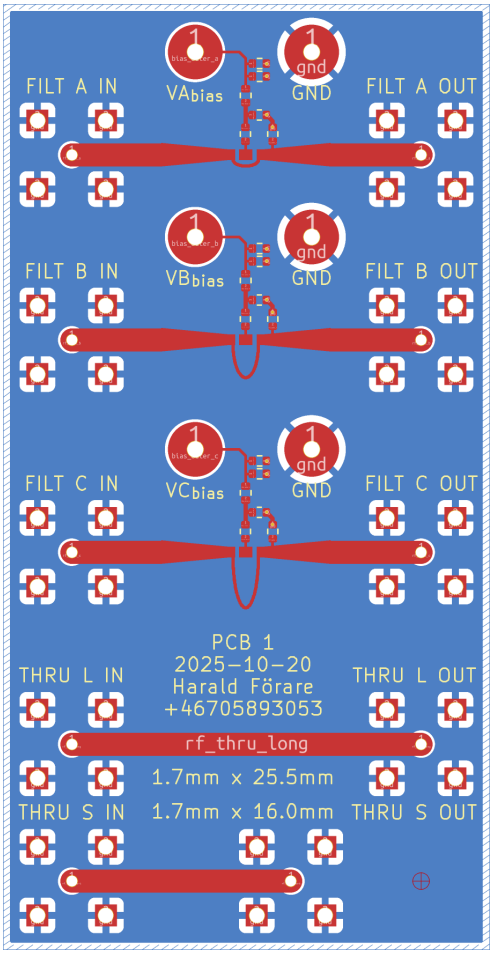

Figure 3.10: Final PCB design

Dielectric	0.762 mm Rogers RO4350B
Board Size	35.5 mm × 69.0 mm
Copper Layers	2
Surface Finish	ENIG
Gold Thickness	1U”
Copper Weight	1 oz.
Via Covering	Untented
Manufacturer	JLCPCB

Table 3.2: PCB specifications

3.1.7 Component orders

SMA connectors with a maximum frequency rating of 12.4 GHz were selected. Thin film resistors were chosen due to better accuracy, lower temperature coefficient and crucially less parasitic inductance compared to other types. The exact values for the decoupling capacitors was unknown at this stage, so a couple different values were ordered. Zero temperature- and voltage coefficient dielectric, high-Q and 50 V-rated ceramic capacitors were chosen.

The original capacitors selected were from a well known Japanese brand. Unfortunately they declined the order, most likely due to Evolved Aerospace being in the defense/military sector. A new order had to be placed for capacitors from a different manufacturer. This added about a week to the lead time.

3.1.8 Biasing network optimization

The validity of the design solutions above depend strongly on how well the lumped components match theory in reality. The biasing network should be quite forgiving in this type of architecture, since the varactor diodes are voltage-biased and there is no requirement to deliver a strong current quickly. This means that the biasing network mainly needs to present a sufficiently high impedance to avoid loading the RF-path. It was however found in Section 3.1.5 that certain resistances cause a good match with the resonator and therefore a very deep notch at one frequency in the sweep. It would be nice to realize this by closely realizing the specified impedance shown in Table 3.1, and not just some arbitrarily large value.

Using a co-simulation method very similar to the one described in sections 3.1.4 and 3.1.5, the decoupling network was optimized. Figure 3.11 shows the EM-model used in Ansys HFSS. Figure 3.12 shows the expected impedance for the best solution. It was found that using the largest available capacitors in the placed

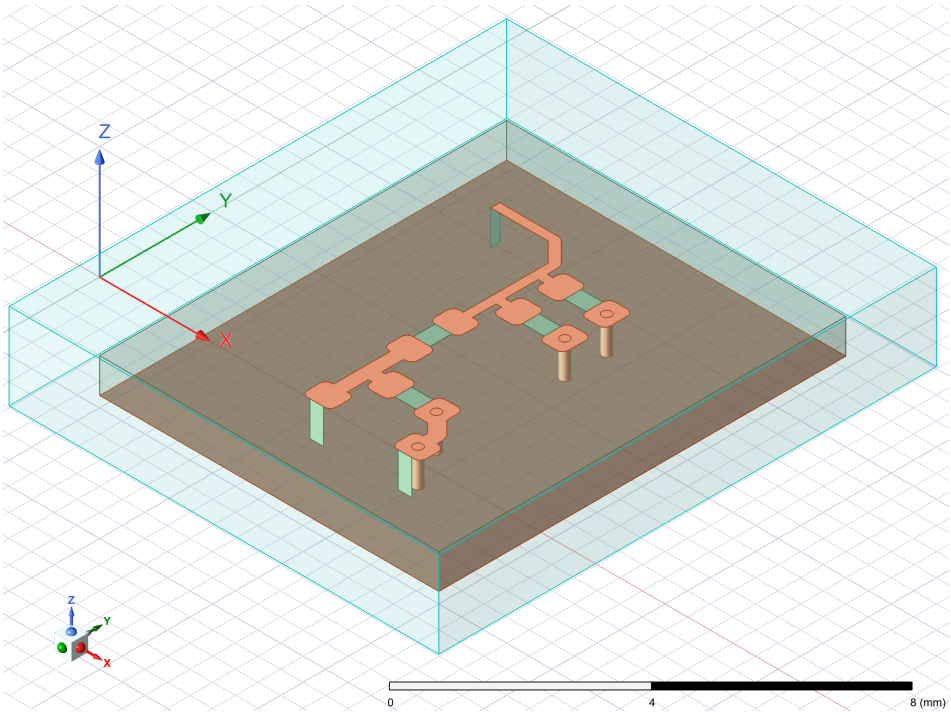

Figure 3.11: HFSS bias network screenshot

component order yields the realization where impedance can be kept low up to the highest frequency. Figure 3.13 and Table 3.3 show the selected components' locations and values.

R_{b1}	261	Ω
R_{b2}	6.9	kΩ
R_g	6.9	kΩ
R_d	10	kΩ
C_{d1}	1000	pF
C_{d2}	1000	pF

Table 3.3: Bias network component values

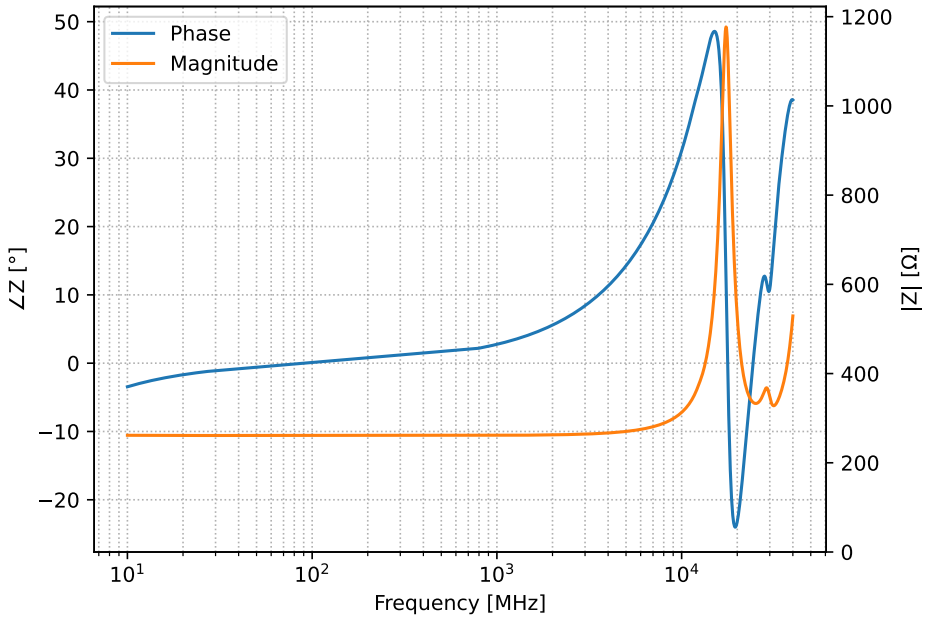

Figure 3.12: Bias network impedance

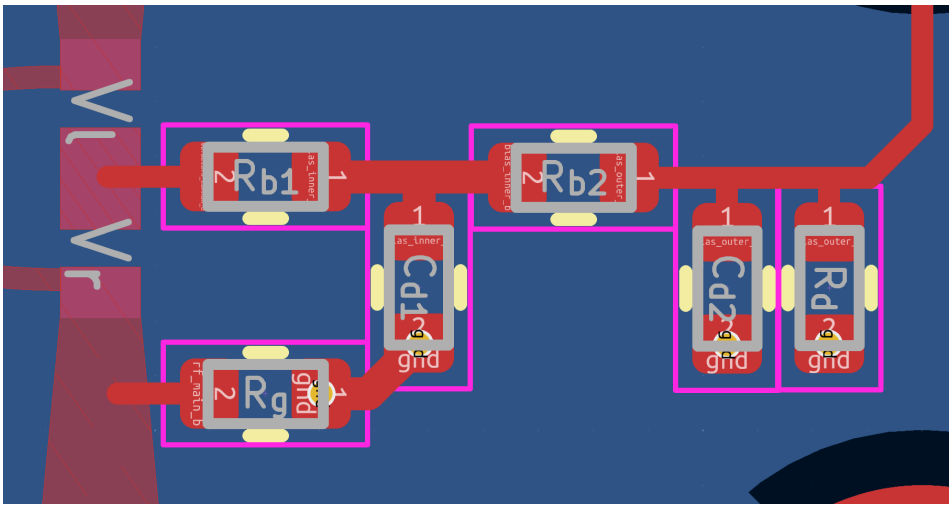

Figure 3.13: Bias network schematic

3.1.9 SPICE-simulatons for nonlinear behavior

The extra time while waiting for the PCB and components was spent on exploring nonlinear behavior using simulation.

I have the EM-simulations of the PCB done in HFSS. I have the lumped varactor model defined in SPICE from the manufacturer. It would be nice to combine these to simulate nonlinear behavior in the full circuit. There are multiple different ways to model nonlinear behavior, the most general and low-risk-of-error being direct time-domain SPICE simulatons. There are other common techniques in RF, such as harmonic balance, which make some assumptions such as signals being periodic that greatly speeds up simulation time. Due to lack of access to HB-simulators, the desire to have simple and likely correct simulations, and familiarity with SPICE, it was chosen to attempt time-series simulations in ngspice to explore nonlinearities.

The first problem is combining the inherently frequency-domain S-parameter model for the PCB with the SPICE-models for lumped nonlinear components, and run everything in a time-domain simulator. This is commonly done by fitting a lumped-network to the S-parameter data, to get an equivalent circuit. I had little success in doing this. Multiple tools were tried, namely lumped-circuit export in HFSS, and similar methods in SciKit-RF. Neither produced SPICE-models that yielded converging time-series simulations in ngspice. Some time was spent getting this to work, but the effort was abandoned due to having no progress. A compromise was made by simply fitting a simple LC-resonator model, like that in Figure 2.2b, to the linearly terminated HFSS model, extracting the inductance. This model was then combined with the SPICE-model for the varactor, shown in Figure 3.14.

Time-series data from ngspice was fourier-transformed in Python, yielding spectra. Simulations were ran to integer multiples of the least common divisor of the periods of input tones. This removes the need of windowing as complete periods of the signal are captured. Initial cycles were discarded to allow for transients to dissipate.

By looking closer at the varactor SPICE-model used (Figure 3.14), especially noting the fact that all series resistance is put outside the diode semiconductor material, it became clear that nonlinearities likely won't be well captured. Since our solution utilizes a parallel-resonator, the strong current at resonance causing a voltage over the junction will likely be the dominant mechanism for harmonic generation. This won't be captured at all by the model in Figure 3.14 since the resistance is separate. The amount of current gain is also determined by the circuit Q-factor, as presented in Section 2.2.2, which is uncertain given the fitted model.

At this point in time, the ordered hardware arrived. A decision was made to abandon these simulations and focus the effort on real world measurements.

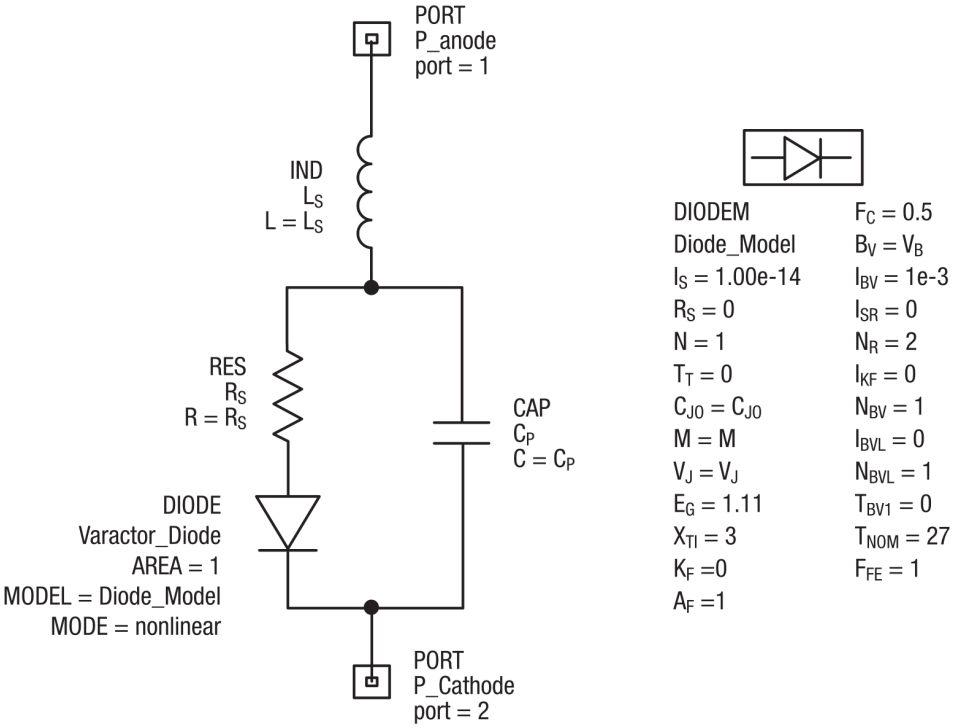

Figure 3.14: SPICE-model used for nonlinear simulations. Taken from a Skyworks application note [8]

3.2 Manufacture

The PCB was hand soldered by me. It was first cleaned with isopropanol alcohol and cotton swabs. Solder paste was then applied under microscope magnification to the pads with a small sewing needle. Components were placed on the pads using tweezers and held in place by the wet solder paste. The view through the microscope at this step is shown in Figure 3.15. The solder paste and flux used is shown in Table 3.4. The board with components was then transferred to a ceramic pizza stone that had been preheated in an oven to 150 °C. Its temperature was verified with an IR-thermometer. The hot-air gun was kept on standby and hot air was applied to the components on the board with a small nozzle expediently after placing it on the preheat surface. The specified reflow temperature for the solder paste is 219 °C, so initially the hot-air gun was set to 300 °C. Keeping in mind that the varactors and dielectric material both are rated for 260 °C only. The solder paste remained solid. After about 15 s, realizing that nothing happens, the hot-air temperature was increased to 400 °C. The solder paste then became shiny and components were "sucked" into alignment due to surface tension, indicating successful soldering.

Even though the hot-air temperature is known, it is difficult to know the actual temperature of the components and solder paste on the board surface. Solder paste becomes harder to reflow if heated for too long as the flux evaporates, which limits the minimum hot-air temperature. The maximum hot-air temperature is limited by the fact that components and dielectric material can burn and be destroyed. One also strives to minimize temperature gradients across the board during soldering, as it can cause warping, cracking and delamination. Preheating and having a lower hot-air temperature are both helpful to reduce the likelihood of that happening. It is useful to constantly move the hot-air nozzle and vary the distance between it and the board to control the process, always looking for a change from dull to shiny appearance in the solder paste, indicating successful reflow.

After all surface-mount components had been soldered, through-hole components such as connectors and bias clamp-eyes were soldered with a normal soldering iron. After all soldering was done, the board was allowed to air-cool and was finally cleaned a couple of times in isopropanol alcohol to remove all flux residue, as it can cause corrosion and electrical parasitics (I have even heard of flux-residue behaving nonlinearly at RF, causing distortion problems). Figure 3.16 shows a close-up of the finished filter.

Solder paste	Sn: 96.5%, Ag: 3%, Cu: 0.5%
Flux	Only that in solder paste

Table 3.4: Solder paste and flux used

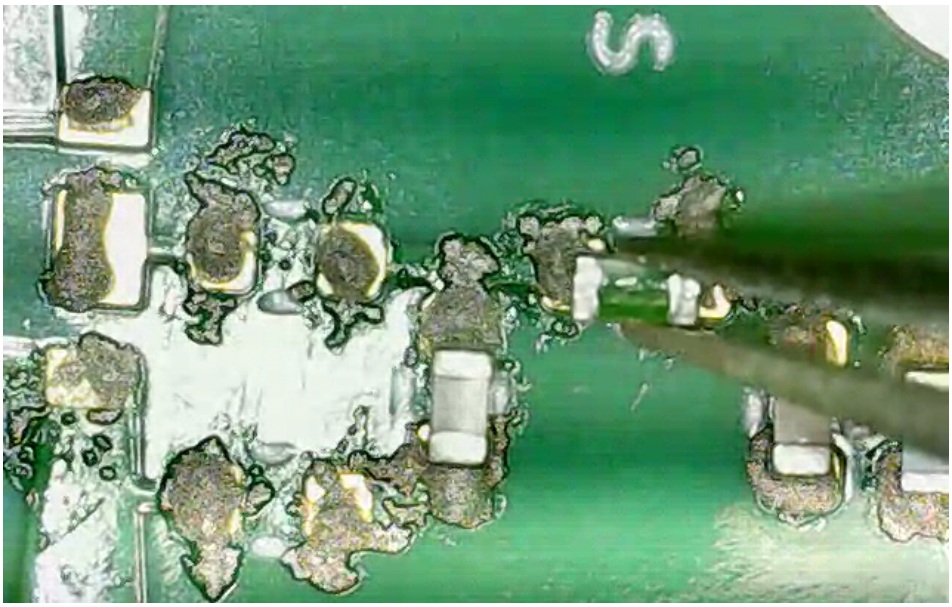

Figure 3.15: Microscope view when placing components on pads in solder paste

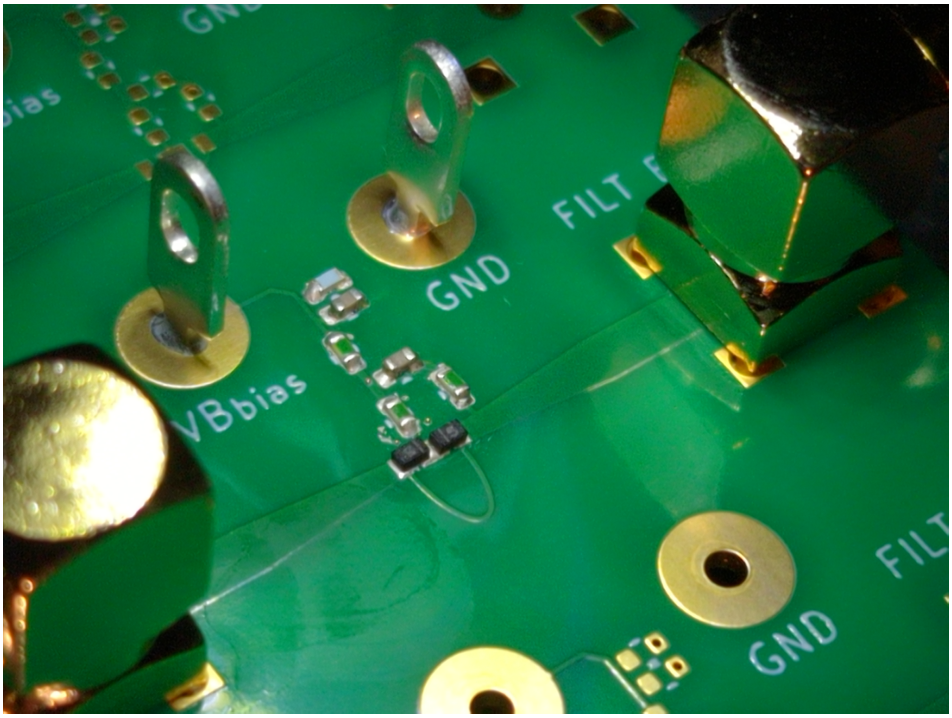

Figure 3.16: Finished PCB close-up

3.3 Measurement

The prototype was measured and verified in the EIT research lab at LTH. Table 3.5 shows all equipment except cables, fixed attenuators and computers used for lab measurements.

VNA	R&S ZVL	9 kHz	...	13.6 GHz
Spectrum Analyzer	R&S FSU	20 Hz	...	50 GHz
Signal Generator 1	Agilent E8257D PSG	250 kHz	...	40 GHz
Signal Generator 2	R&S SME 03	5 kHz	...	3 GHz
VNA Cal. Kit	R&S ZV-Z135 (female)	0 Hz	...	15 GHz
Power Combiner	MCL ZAPD-4 9330 02	2 GHz	...	4.2 GHz
SDR	A. D. ADALM-PLUTO	325 MHz	...	3.8 GHz
Multimeter	BST BS1704			
Bias Power Supply	Powerbox 3000 Lin. Reg.			
SMA Torque Wrench	Cinch 141-0000-930			0.57 N m

Table 3.5: Instruments and equipment used

3.3.1 VNA calibration

Initially there were fine undulations in the S-parameter data, like what is visible in Figure 4.2 above roughly 6.4 GHz, but with much higher amplitude. This was speculated to be due to bad calibration. It was discovered that the calibration standard specified in the VNA did not match the electrical characteristics of the calibration kit (Table 3.5) used. A new calibration standard was inputted into the VNA using the electrical model presented in the datasheet for the calibration kit. This remedied the vast majority of these abnormal S-parameter undulations. Calibration kits are not perfect electrical references. They have some parasitics. It is important to input these so that the VNA can properly compensate for them and put the reference plane directly at the boundaries of the DUT. After calibration, some fine undulations remained at higher frequencies. This could be due to cables changing impedance when being bent, connectors not being ideal, or something else. Despite this, the calibration was deemed good enough as these minor remaining undulations would not hinder characterization of the prototype in any meaningful way.

3.3.2 S-parameter measurements

The filter was connected to a VNA. S-parameters were recorded using multi-measurement averaging and a wide logarithmic frequency sweep. This was done

for multiple bias points between 0 V and 25 V. Besides the filter on the PCB, there is a microstrip line with identical substrate and connectors as for the filter. This was also measured in the exact same way as the filter. The S-parameter data was analyzed in Python where the notch center frequencies were extracted to sub-sample accuracy using spline interpolation. Table 3.6 shows the specific VNA settings used.

Number of points	2001	Num. sweeps averaged	4
Start frequency	100 MHz	Total measurement time	21 s
Stop frequency	13.6 GHz	Measurement bandwidth	1 kHz
Sweep type	Logarithmic	Power	−10 dBm

Table 3.6: VNA settings

3.3.3 Harmonic distortion

The input of the filter was conneted to a signal generator, and the output of the filter was connected to a spectrum analyzer. Code was written to take spectrum-analyzer measurements at the expected locations of harmonics at varying attenuator levels automatically. The swept variables were tone power, tone frequency and filter bias point. A RMS detector was used in the spectrum analyzer to capture power levels correctly. The complete spectrum analyzer settings are visible in Table 3.7.

Resolution bandwidth	100 Hz	Sweep time	240 ms
Video bandwidth	10 MHz	Detector	RMS
Frequency span	2 kHz	Number of points	30001

Table 3.7: Spectrum analyzer settings for harmonic distortion measurements

3.3.4 IMD3

Two signal generators were connected through a power combiner to the input of the filter. The output of the filter was connected to a spectrum analyzer. IMD3 distortion was excited by injecting two closely spaced tones. Code was written to take multiple spectrum measurements in the vicinity of the tones and expected third order intermodulation product. These were combined later in Python to gain a lower-noise result. Each measurement was taken at multiple different spectrum analyzer input attenuator values to be able to detect systematic errors in the measurement setup later on. A tone spacing of 20 kHz was used.

Since the spectrum analyzer measurements are concerning tone power levels, an RMS-detector is appropriate and was used. Furthermore each tone is well predicted from theory (Section 2.5.3) and is narrowband, which makes it practical to use a narrowband resolution filter without drastically increasing measurement time, since the sweep is small. This gives a lower noise floor [7, 19]. Since data was to be post-processed the video bandwidth was set to the maximum value, to not skew locations of peaks and hide fine details in the spectrum. Table 3.8 shows the settings used on the spectrum analyzer during IMD3 measurements.

Resolution bandwidth	200 Hz	Sweep time	250 ms
Video bandwidth	10 MHz	Detector	RMS
Frequency span	10 kHz	Number of points	30001
Attenuator	30 dB		

Table 3.8: Spectrum analyzer settings for IMD3 measurements

3.3.5 Cross-compression & cross-modulation

The filter was biased to have the notch at 2.8 GHz. Its input was connected to two signal generators through a power combiner. One supplying a low-power −20 dBm tone at 3.5 GHz, and the other supplying an AM-modulated tone centered at different frequencies for different experiments. The AM-modulated tone was centered at 2.6 GHz, 2.8 GHz and 3.0 GHz to test device behavior with interference below, at and above the notch respectively. For each experiment the AM-tone carrier power was swept. The modulation index used was 30 % and the modulation bandwidth was 1 kHz.

For cross-modulation, sidebands to the pure carrier tone were measured in relation to the carrier at 3.5 GHz ± 1 kHz. It was visualized as a function of AM-modulated tone power.

Cross-compression can be extracted from the same dataset by measuring the decrease in tone power at the unmodulated signal as a function of swept tone power. This is possible as the two measurements measure different and non-overlapping features of the output. Furthermore the cross-modulation sidebands are so small in comparison to the fundamental that even if they would somehow be subtracted from or by some other mechanism change the fundamental, it would make negligible impact on the level of the fundamental.

It was difficult to make out the modulation sidebands from the noise floor, so a very small resolution bandwidth was required in the spectrum analyzer. Similar rationale as described in Section 3.3.4 was employed when configuring the spectrum analyzer, and the exact settings are shown in Table 3.9.

Resolution bandwidth	10 Hz	Sweep time	25 s
Video bandwidth	10 MHz	Detector	RMS
Frequency span	2.5 kHz	Number of points	30001

Table 3.9: Spectrum analyzer settings for cross-modulation and cross-compression measurements

3.3.6 End-to-end SDR test

If this filter is implemented as a part of a real-world signal reconnaissance system, most of the signals-of-interest would be digitally modulated. Modern modulation schemes use both instantaneous amplitude and phase to encode information. To evaluate end-to-end performance, a test was devised where actual digitally modulated signals are sent through the device, to capture all performance effects.

With a software-defined radio (SDR), like the one specified in Table 3.5, radio signals with both amplitude- and phase modulation can be transmitted and received. Code was written to generate IQ-samples uniformly distributed on a unit disc. These were packaged into a transmission with additional features, such as a preamble and a repeating training sequence, designed to enable robust synchronization between the transmitted and received signal. The structure of a transmission is shown in Table 3.10.

Description	**Length [symbols]**
Leading silence	100
Preamble	100
Repeating training sequence ($\times 5$)	$5 \cdot 30 = 150$
Data payload	2000
Trailing silence	50

Table 3.10: Baseband SDR transmission

To limit bandwidth, reduce inter-symbol-interference (ISI), and allow for matched-filter reception, the baseband signal needs to be pulse-shaped before being sent to the radio hardware. This was done with a raise-root-cosine (RRC) filter with roll-of factor $\alpha = 0.35$ and 4 samples-per-symbol. The truncation window was 24 symbols.

The baseband transmission was then digially mixed by 100 kHz to avoid potential hardware problems with putting the baseband directly on DC. This buffer was then uploaded to the SDR.

Interference tone power and frequency were manually swept, and for each point,

the SDR was automatically controlled to record receive buffers at different LO-frequencies and combinations of transmit- and receive gains. All transmit- and receive buffers were saved in a compressed state, along with all metadata and the baseband symbols, to allow for measurements of link quality later. A small program was written to quickly correlate transmit- and receive buffers and compute a simple EVM score estimate on the live data, to give feedback on if the whole setup was working. Figure 3.17 shows a screenshot of this interface. Figure 3.18 illustrates the physical measurement setup. Some attenuation was used to create more ideal terminations for the signal generator and filter input, but none was used on the SDR receiver input, as the whole goal is to approximate a real-world scenario, and hopefully excite problems in the SDR receiver-chain.

The receive-buffer was correlated with the transmit buffer to find phase offset. Sampling instances were chosen that maximized correlation. The signals were matched in amplitude to maximize correlation and then normalized so that the average transmit power is 1. Then a final EVM-score was calculated for each.

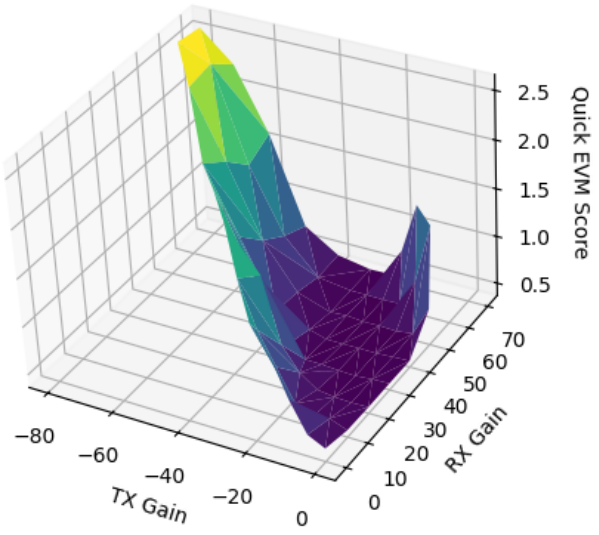

Figure 3.17: Rapid prototype live view EVM estimation used during measurements

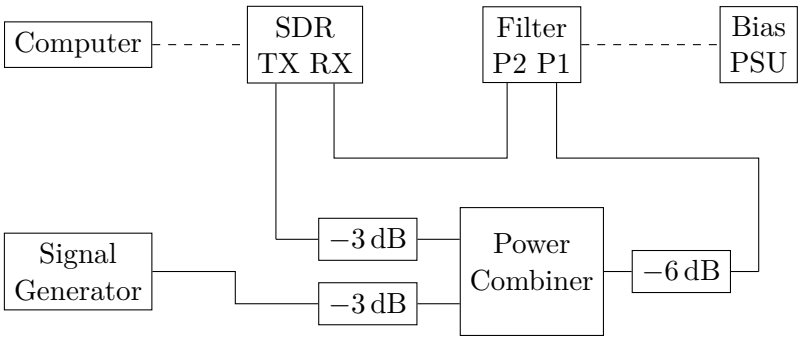

Figure 3.18: End-to-end test setup

3.3.7 Comments on organization & programming tricks

Throughout the duration of this project, it has not always been clear what the next step will be, and solutions to future problems cannot be known in advance. By using version control, in this particular case git, I can freely experiment while never losing information or the ability to go back to a working state. This is very nice as the project grows. Version control also automatically provides a kind of diary over what changes have been done. Not only code can be tracked, but also 3D simulation models, circuits, documents, lab notebooks, measurements data and general notes. Almost any type of file.

In science and engineering, reproducability and traceability are very important. Reproducability means that results can be replicated at a later date. To be able to do this, it has to be clear what was done in the first place. This is possible with traceability. Traceability makes it much easier to confidently answer questions later on about what was done in the past, which is essential to gain maximum value out of a project. In practice, these two are achieved by a combination of version control, continuous documentation and certain programming patterns.

During the lab phase of this project, there were quite a few experiments that had to be carried out within a limited amount of time. Since much of it was new to me, I didn't always know exactly what data I would need for each different analysis. With limited lab time, a balance has to be struck between doing enough analysis to know that the measurements are useful, but not so much that it takes too much time away from doing the actual measurements. Some amount of qualified guessing is required.

The negatives of this trade-off were eased by having the ability to capture more data quickly. This was done with automation of lab equipment. Measurement data and crucially metadata was dumped into simple .json-files. This is not the most efficient format, but it is extremely simple to interface with, and very flexible.

Each raw-data file was written to once and is later on not allowed to be changed. Much of the inefficiency of the raw data files was avoided by placing them inside .zstd-compressed archives.

Later on, during the analysis, these raw-data files were not read directly, but rather abstractions in code were written over them, with layers of caching to speed up raw-data access. There are mistakes in the raw-data files, but they are corrected in the code that loads them, rather than the files themselves. This maintains traceability and reproducibility. All of this is version controlled and it becomes very explicit what has been done. After caching, all of the cleaned data is available almost instantly for analysis.

Chapter 4

Results

4.1 Small-Signal S-parameters

Using the laboratory setup described in Section 3.3.2, the small-signal S-parameters of the filter were measured. The tuning range- and sensitivity of the filter is shown in Figure 4.1, and Figure 4.2 shows the forward transmission S-parameter for a few bias points. Figure 4.3 shows the notch depths measured over center frequency, but these are slightly unreliable due to too sparse frequency sampling in the VNA.

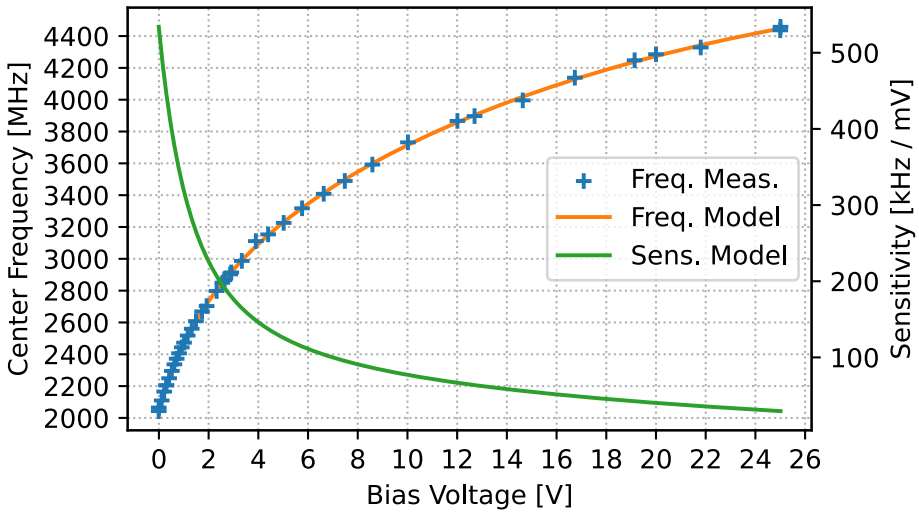

Figure 4.1: Tuning range

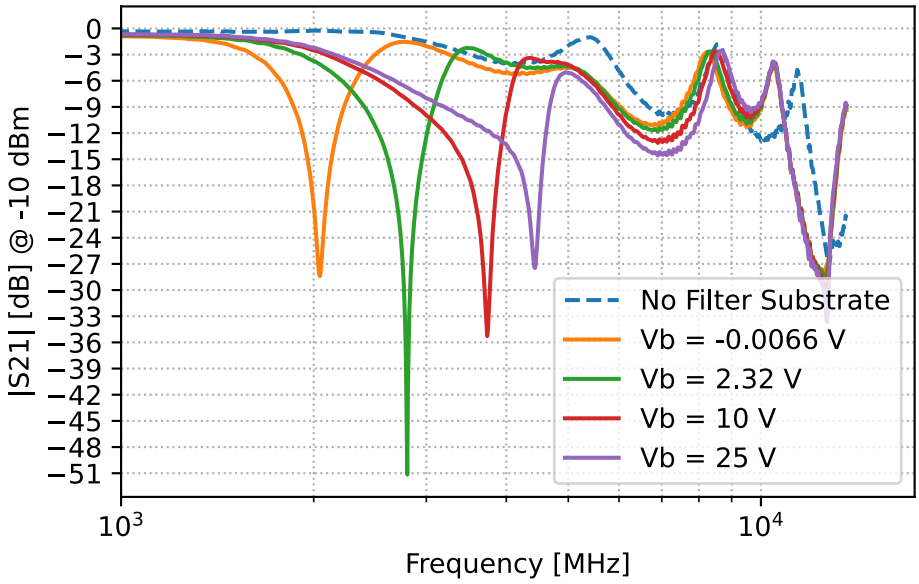

Figure 4.2: Small-signal forward transmission for a few bias points along with reference

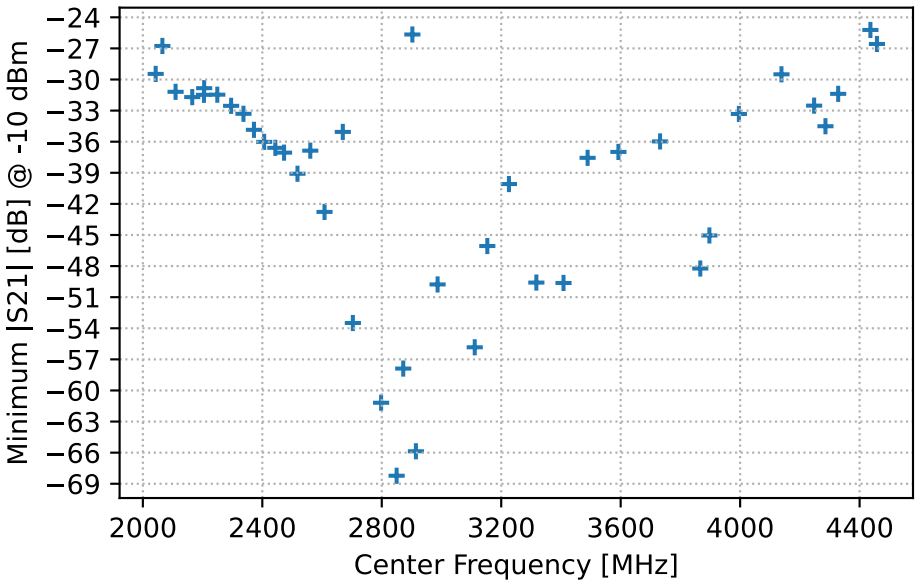

Figure 4.3: Notch depth for different bias points

4.2 Compression

For the low-voltage bias point with notch frequency $f_n = 2.8\,\text{GHz}$, the $-1\,\text{dB}$ compression point occurs at $-3\,\text{dBm}$. For the higher bias point at $f_n = 3.7\,\text{GHz}$, compression is not observed. Tones injected 80 MHz away from the notch at 2.8 GHz show a very slight deviation from a straight line over about 5 dBm. This is all visible in Figure 4.4.

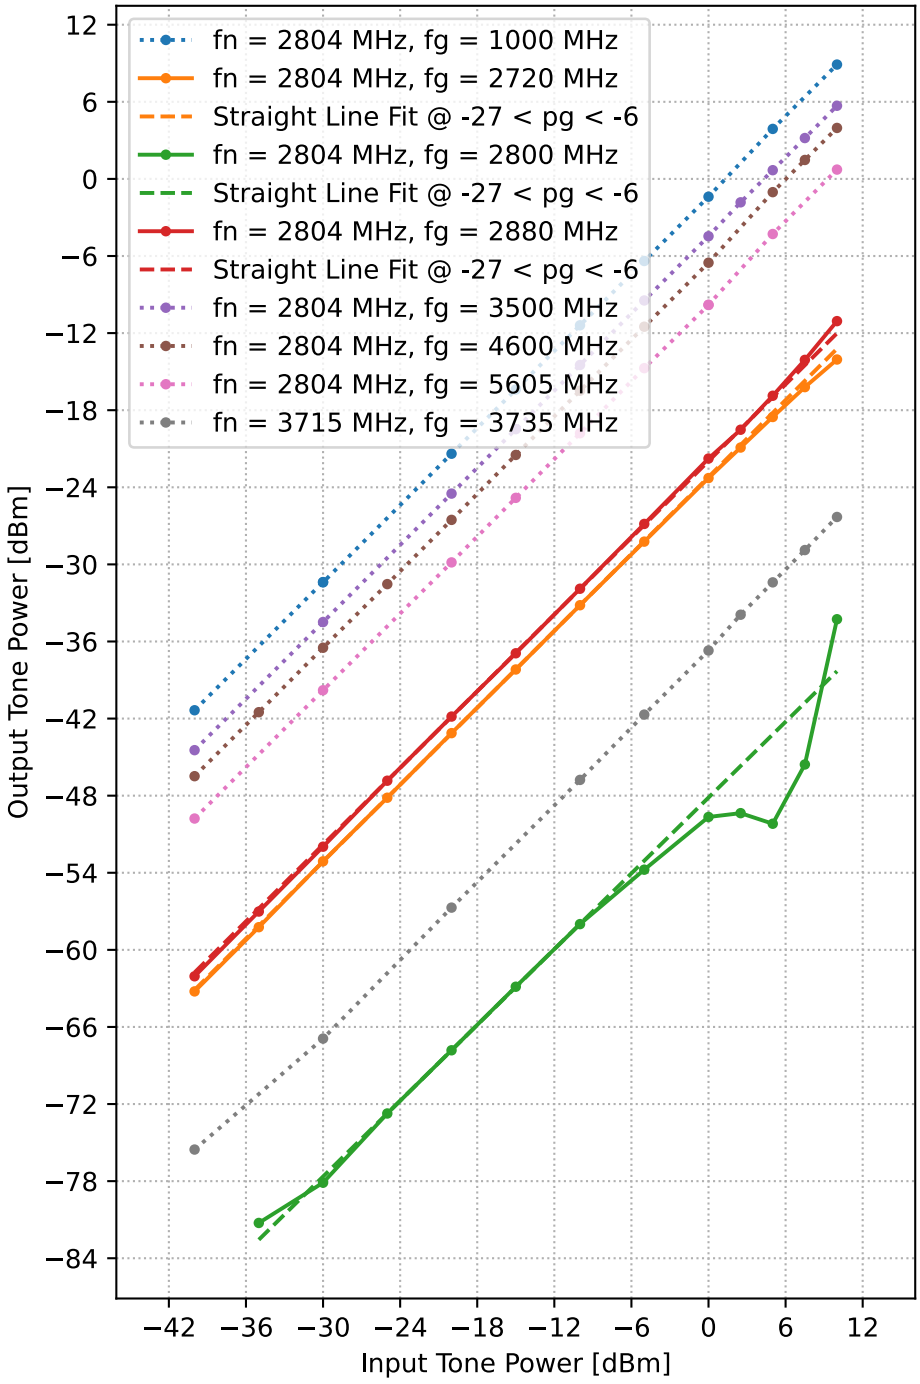

Figure 4.4: Compression

4.3 Harmonic Distortion

Throughout all measurements, only the fundamental tone, second- and third order harmonics were ever observed. No higher order harmonics than those were measurable above the noise floor. Figures 4.5, 4.6, and 4.7 show the crossings in power between a desired tone and the harmonics from an in-notch interference tone. Table 4.1 shows these intercept points in table form.

f_{notch}	$f_{\text{interference}}$	f_{desired}	IIP_2	IIP_3	OIP_2	OIP_3
2804	2800	5600	41.2	38.6	31.9	29.3
2804	2720	5600	51.5	38.0	42.3	28.7
2804	2880	5600	46.5	38.6	37.3	29.4

Table 4.1: Harmonic distortion intercept points from figures 4.5, 4.6, and 4.7. Stated here for compactness, units are MHz for frequencies and dBm for powers.

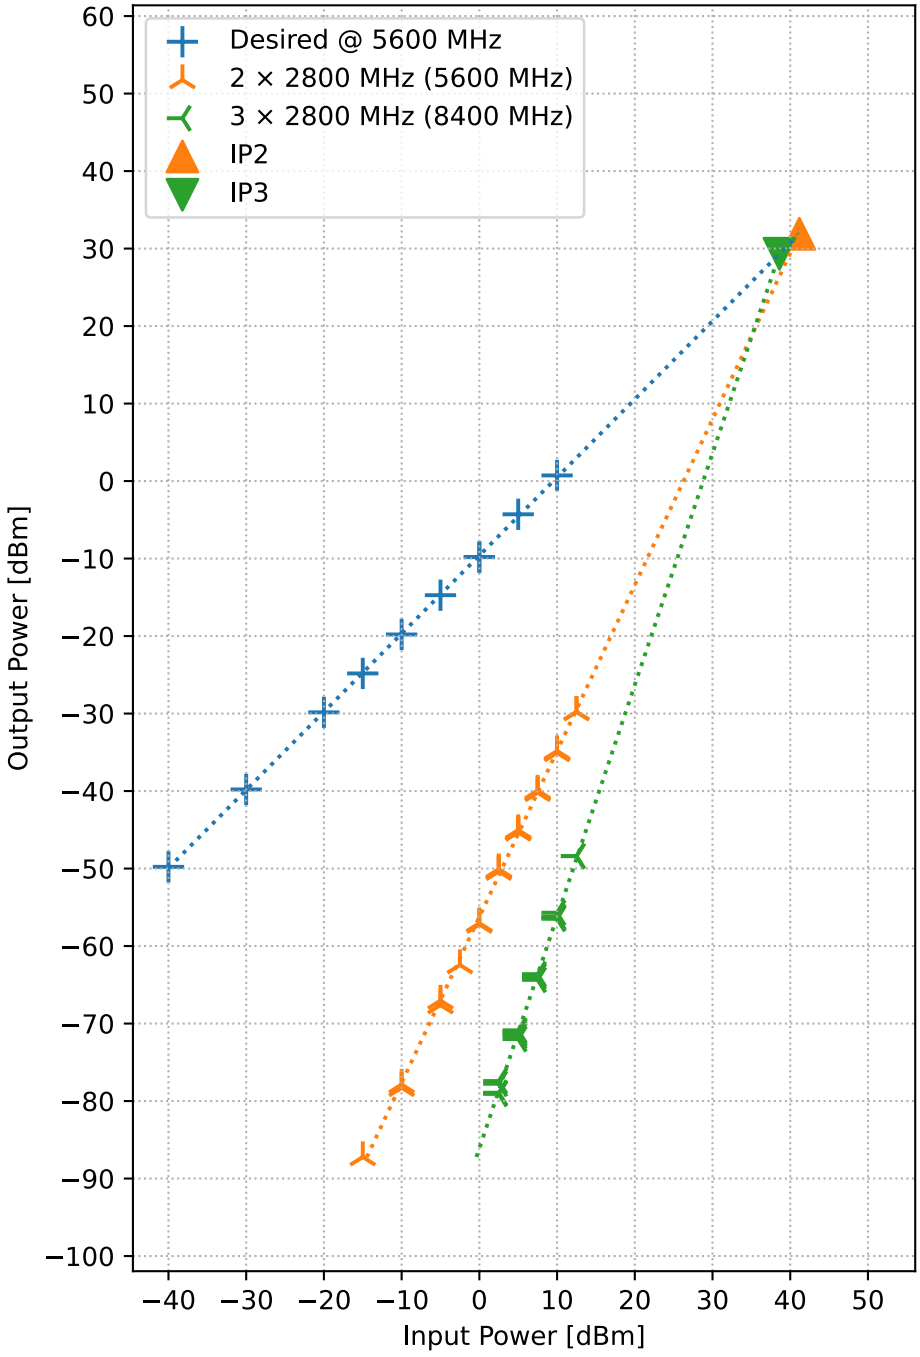

Figure 4.5: Effect of harmonic distortion from an in-notch interferer at $2800\,\mathrm{MHz}$ showing up at a desired signal at $5600\,\mathrm{MHz}$, along with the third harmonic at $8400\,\mathrm{MHz}$ for completeness

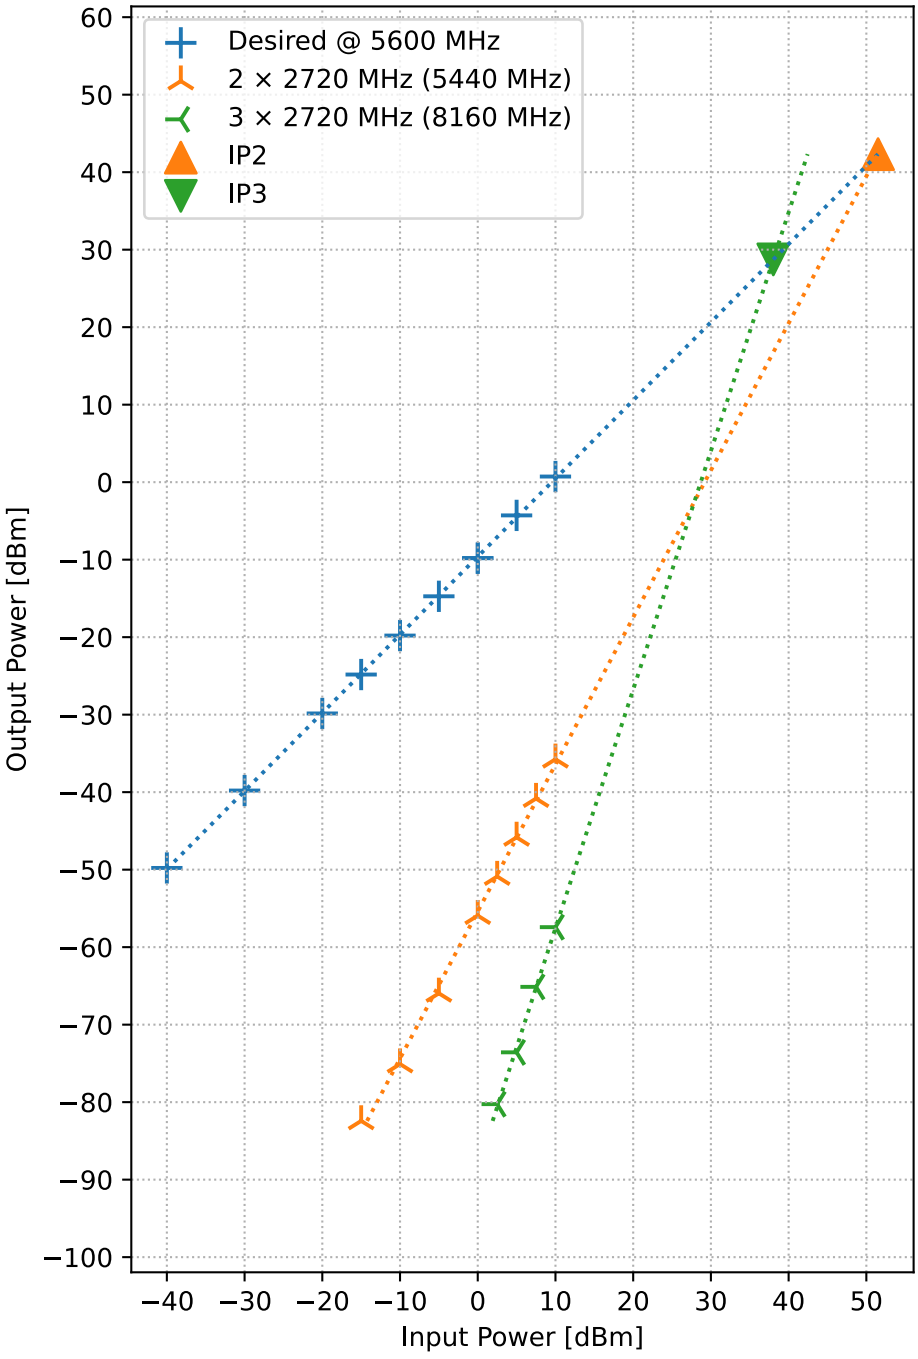

Figure 4.6: Effect of harmonic distortion from a slightly off notch center in-notch interferer at $2720\,\mathrm{MHz}$ showing up close to a desired signal at $5600\,\mathrm{MHz}$, along with the third harmonic at $8160\,\mathrm{MHz}$ for completeness

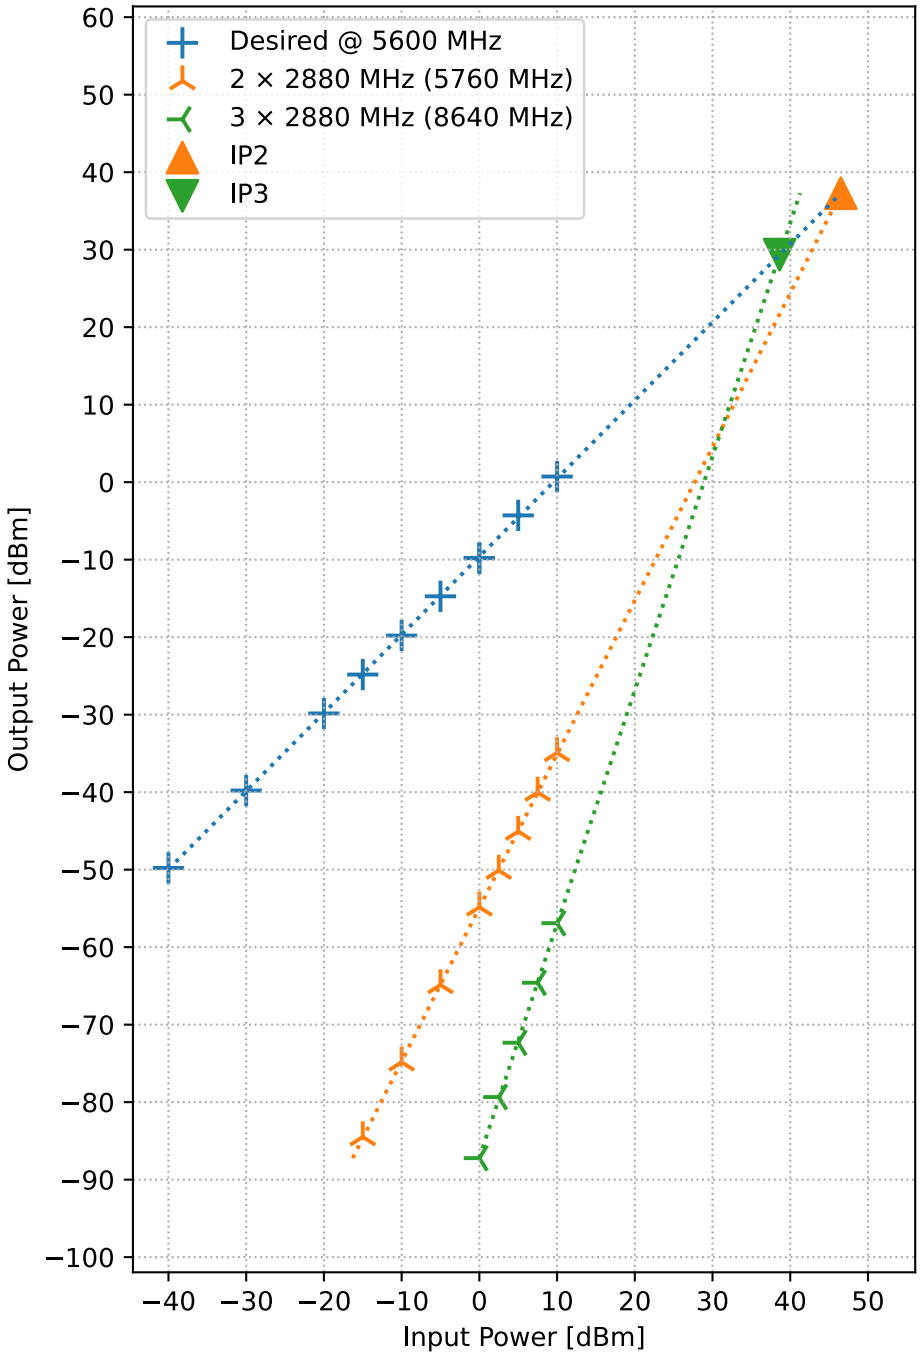

Figure 4.7: Effect of harmonic distortion from a slightly off notch center in-notch interferer at $2880\,\mathrm{MHz}$ showing up close to a desired signal at $5600\,\mathrm{MHz}$, along with the third harmonic at $8640\,\mathrm{MHz}$ for completeness

4.4 IMD3

Using the setup described in Section 3.3.4, IMD3 was measured. Figure 4.8 shows the IMD3 characteristics of the filter with the notch frequency $f_n = 2803\,\mathrm{MHz}$ and the tone spacing $\Delta f = 20\,\mathrm{kHz}$.

Tones Center Frequency [MHz]	IIP3 [dBm tone^{-1}]	OIP3 [dBm]
1000	53.9	47.7
2600	47.3	27.8
2800	13.7	−37.0
3000	23.3	6.81

Table 4.2: Third order order intercept points also seen in Figure 4.8

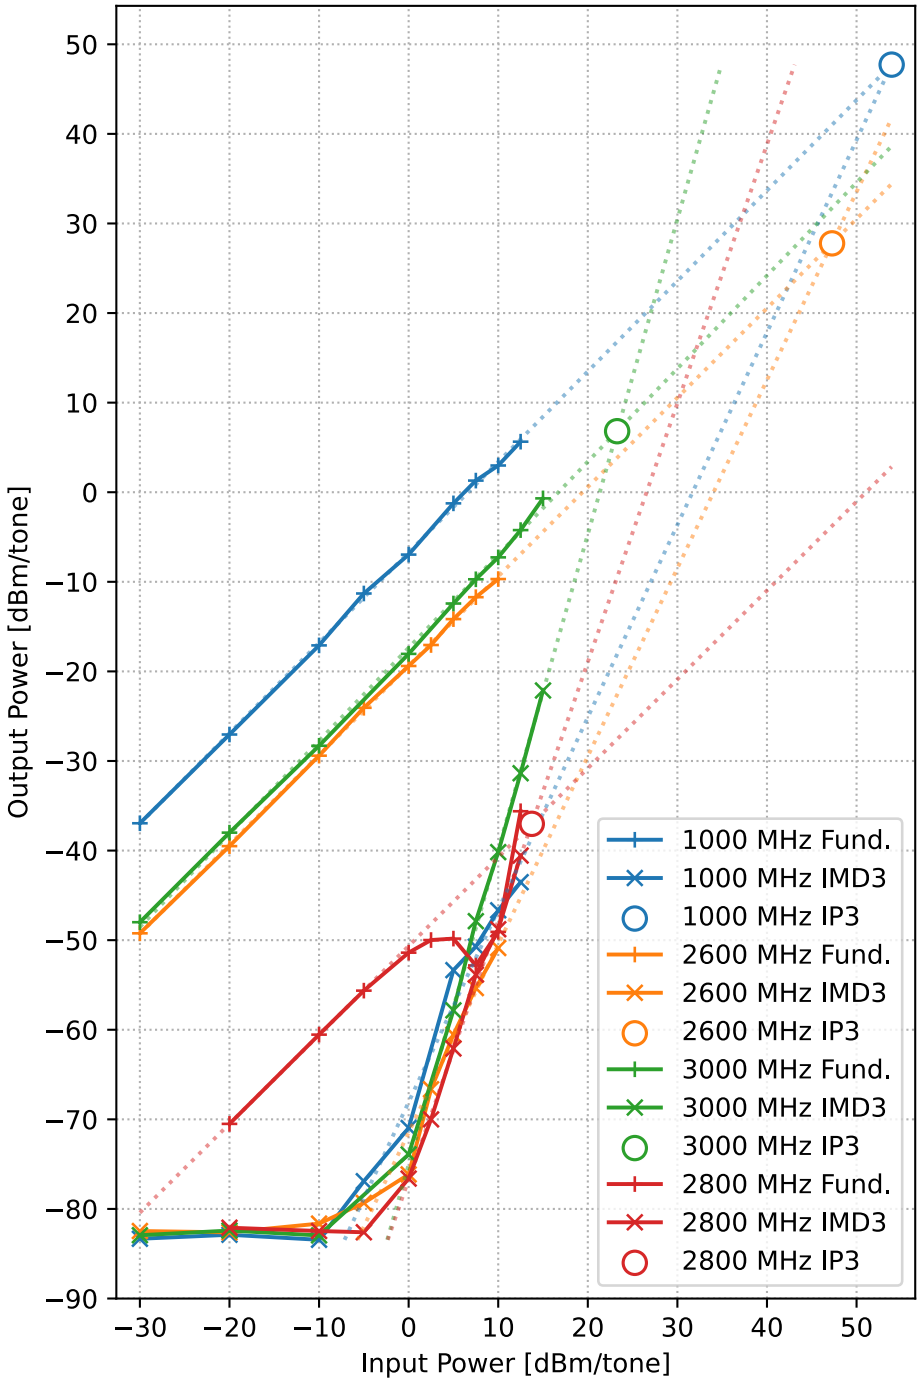

Figure 4.8: IMD3 for tones at different frequencies with notch frequency $f_{\text{notch}} = 2803\,\text{MHz}$ and tone spacing $\Delta f_{\text{tone}} = 20\,\text{kHz}$

4.5 Cross-Modulation

Cross-modulation is shown in Figure 4.9.

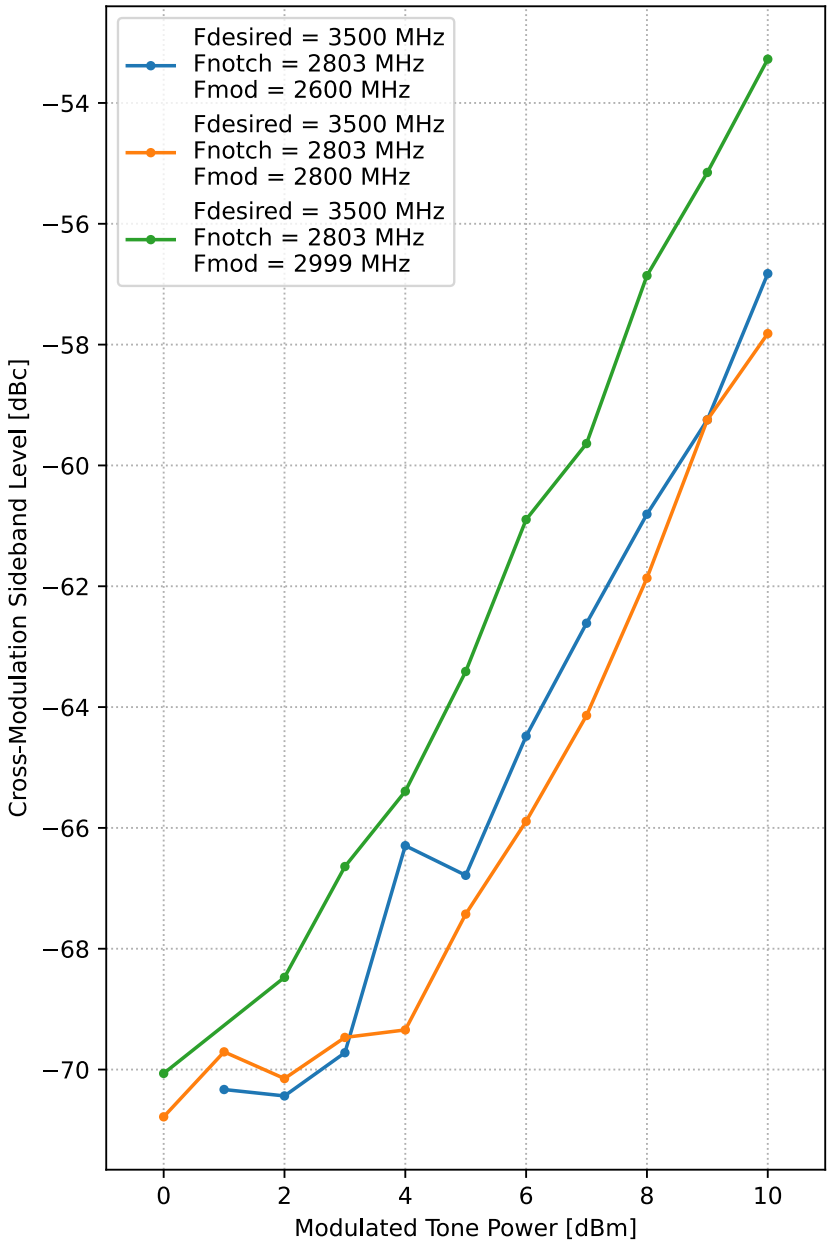

Figure 4.9: Cross-Modulation of a $-20\,\mathrm{dBm}$ signal

4.6 Cross-Compression

No discernable cross-compression could be measured. This is visible in Figure 4.10.

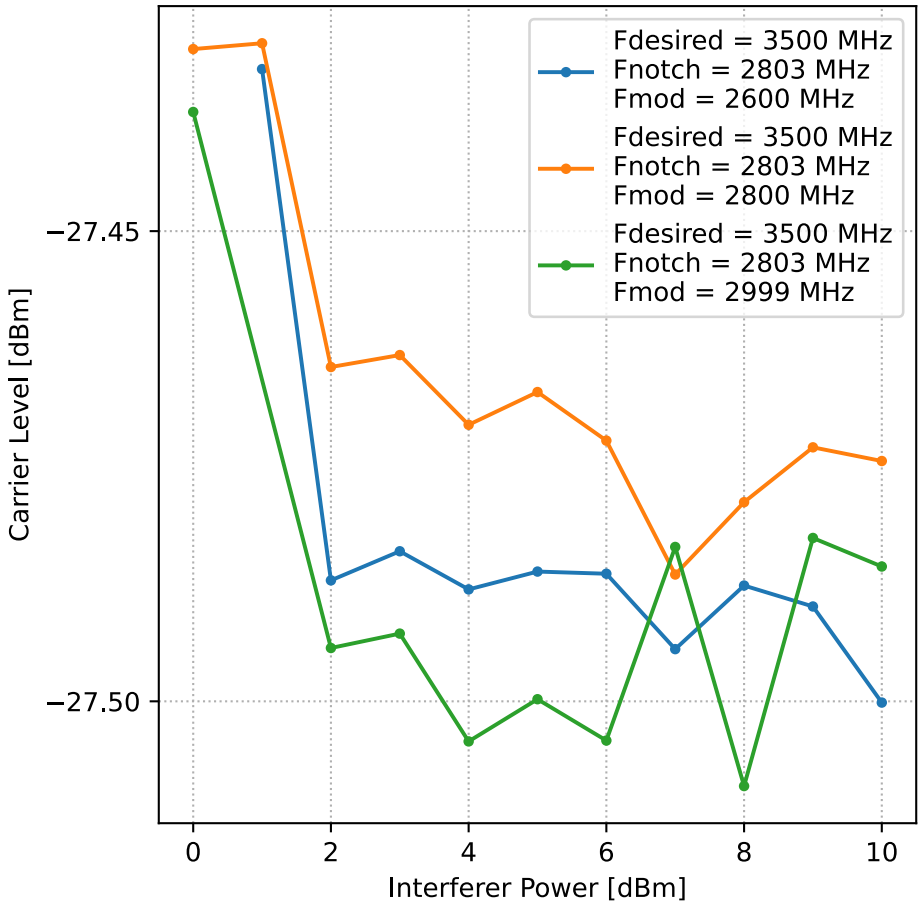

Figure 4.10: Cross-Compression

4.7 End-to-End SDR Test

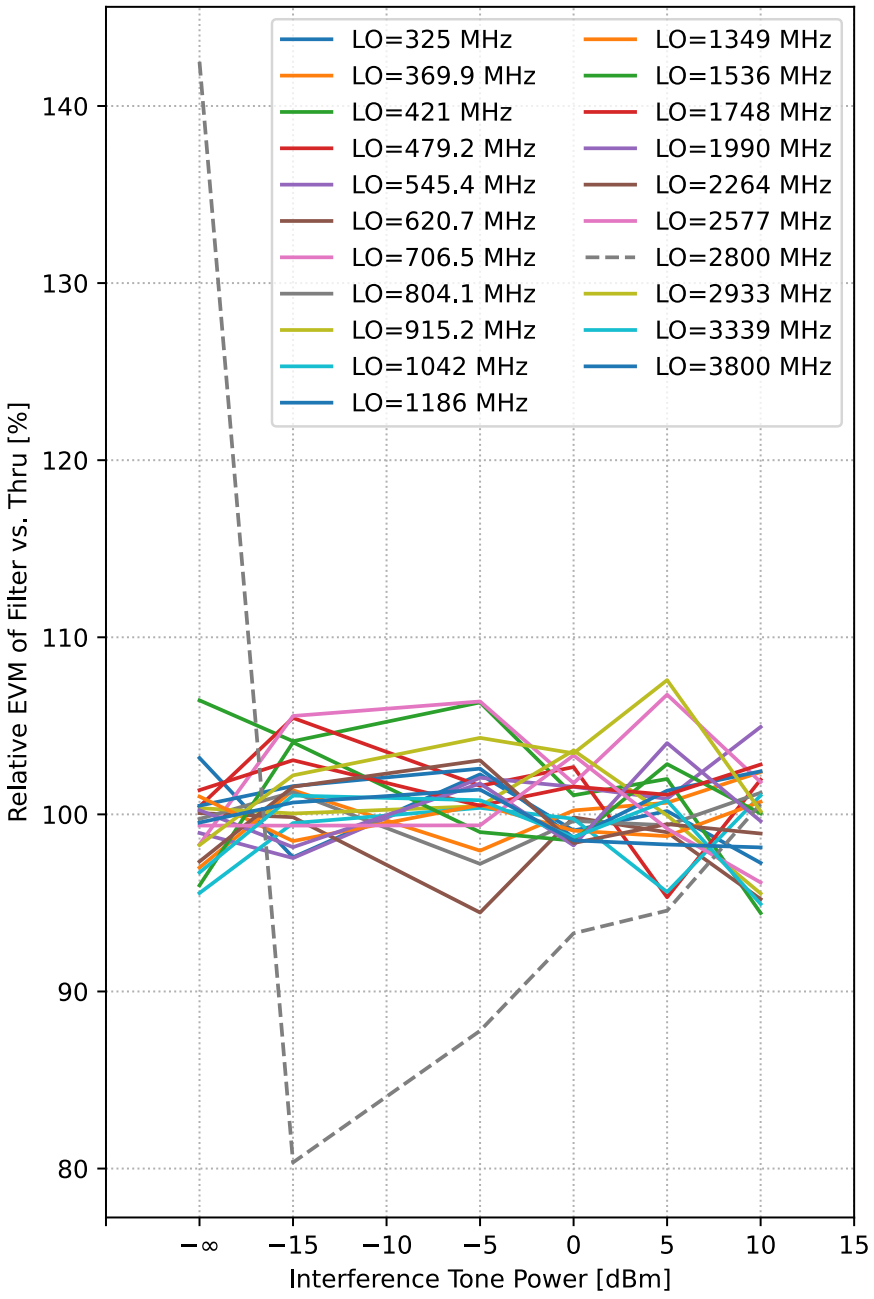

Figure 4.11: Relative EVM with filter notch frequency $f_{\text{notch}} = 2803\,\text{MHz}$ and interference tone frequency $f_{\text{tone}} = 2800\,\text{MHz}$

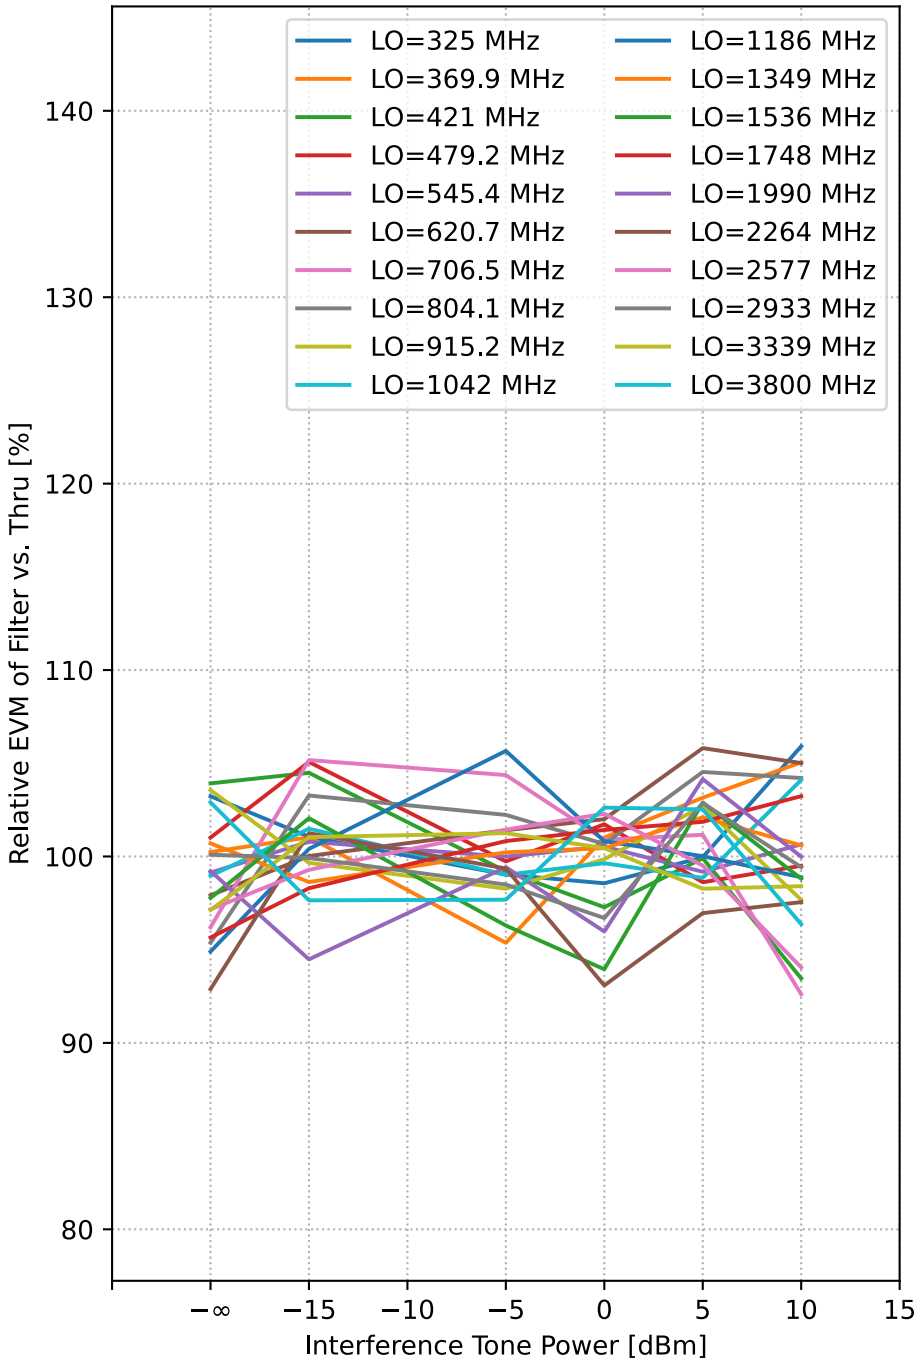

Figure 4.12: Relative EVM with filter notch frequency $f_{\text{notch}} = 3716\,\text{MHz}$ and interference tone frequency $f_{\text{tone}} = 3735\,\text{MHz}$

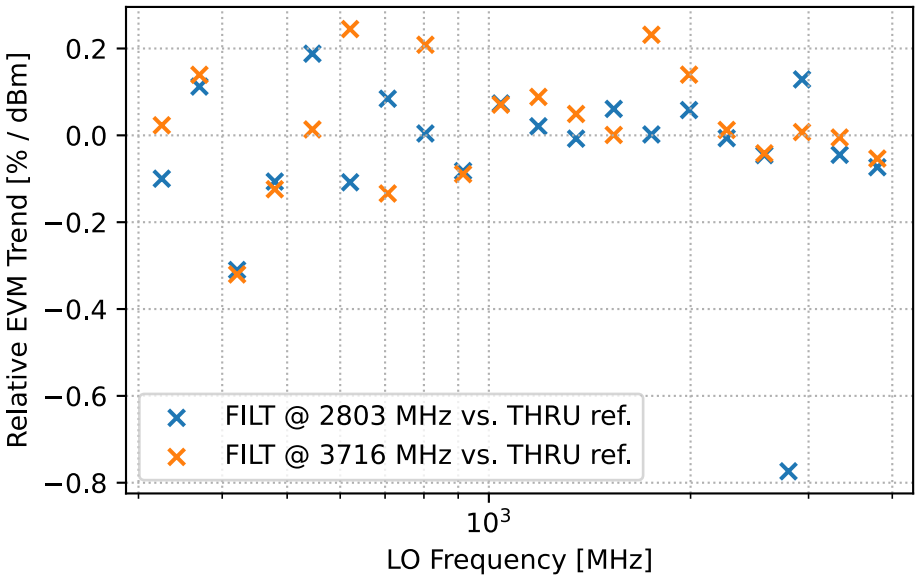

Figure 4.13: Trends of the series' drawn in figures 4.11 and 4.12

Chapter 5

Discussion

This chapter will discuss the results arrived at with the goal of evaluating design decisions and putting it into context of performance in realistic use-cases.

As the intent of this device is to remedy problems in a real receiver system, as outlined in Section 2.6, the most important tests are those that approximate real-world scenarios. Mainly scenarios where the filter would be put at the frequency of an interference source. In a perfectly ideal world, the filter would simply attenuate the power of the interferer as much as possible, leaving other frequencies untouched. In reality, the notch has some slope, limited depth, and there is some nonlinear distortion.

5.1 Small-Signal Behavior

Looking at Figure 4.1, the tuning range is rather good, covering a span of 2 GHz to 4.4 GHz. The notch depth is always at least 24 dB, and sometimes exceeds 60 dB, as seen in Figure 4.3. There is an asymmetry in the notch, which is investigated more in Section 5.4. Filter transfer is shown in Figure 4.2. It is likely that the higher-frequency deviation (above about 6 GHz) is caused by PCB and connector nonidealities, rather than the filter itself, as they are also present for the no filter substrate.

Figure 5.1 shows the simulated vs. measured forward transmission for the filter. The measured notch is generally wider. This indicates that the real filter has more loss (lower Q), and/or a higher inductance to capacitance ratio (Figure 2.4). Note that some of the perceived notch width difference is due to the higher insertion loss of the real filter pushing the entire curve downwards.

Extracting precise Q-values of the filter from Figure 5.1 is slightly complicated by the asymmetric notch and insertion loss, but it is approximately in the $2-5$ range, depending on the bias point. Higher tuning points have higher Q in the real filter.

As outlined in Section 2.3, when the varactor is biased with a larger reverse voltage, it has lower losses because less of the resistive semiconductor bulk material is put in series. This probably explains most of the variation in Q for the real filter for different bias points, compared to the simulated filter. The simulated varactor has a constant series resistance, which would be appropriate to make variable in future work.

The Q-factor of the parallel resonator is proportional to $\sqrt{\frac{C}{L}}$ [7], so at higher bias points, the lower relative capacitance decreases Q. This mechanism is dominant for the simulated filter, but is counteracted by the decreased resistance in the real filter.

A good focus for future work would be to maximize Q to minimize the notch width. The first instinct is to increase the $\frac{C}{L}$-ratio, which is probably correct. One could also look at different varactors with lower series resistance, though these often have worse tuning range. Also keep in mind that the current gain in the resonator scales with Q, as presented in Eq. 2.19, which might have implications for linearity depending on the voltage it causes in the varactor junction.

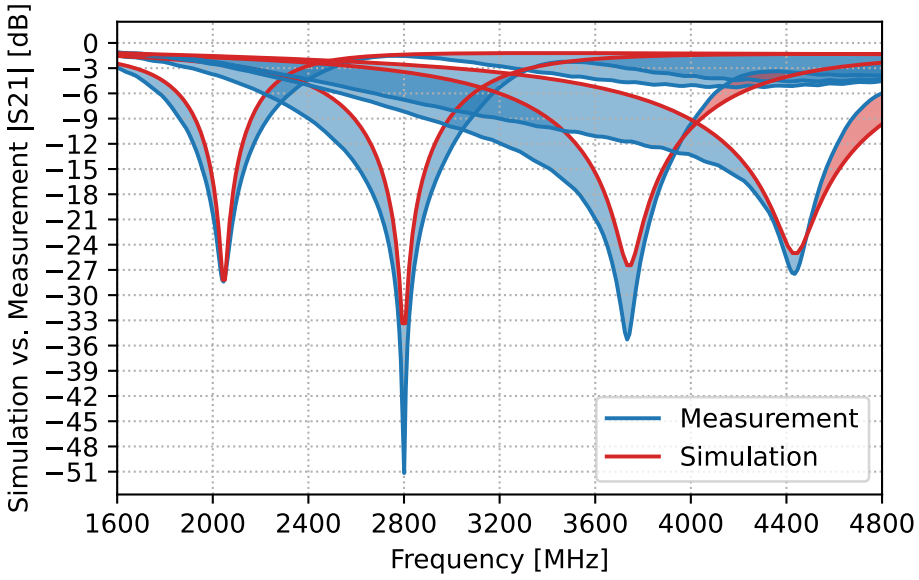

Figure 5.1: Simulated vs. measured forward transmission for the bias points seen in Figure 4.2

5.2 Distortion & Linearity

Looking at the harmonic distortion results in Figure 4.5 and the corresponding Table 4.1, the 2nd harmonic of an in-notch interferer is compared against the fundamental of a desired tone at the same frequency (5600 MHz). Their crossing is not until input power is $\mathrm{IIP}_2 = 41.2\,\mathrm{dBm}$ for the centered notch.

For completeness, the 3rd harmonic at 8400 MHz is also included. It is at a much higher frequency and is thus not completely comparable, but fortunately transmission is high at 8400 MHz, as seen in Figure 4.2, making it worth mentioning: $\mathrm{IIP}_3 = 38.6\,\mathrm{dBm}$.

The $f_{\mathrm{interference}} = 2800 \pm 80\,\mathrm{MHz}$-series', seen in figures 4.6 and 4.7, representing a slightly misaligned notch show similar values, see Table 4.1. These are all very good harmonic distortion results, and suggest that harmonics at higher frequencies won't meaningfully get through the filter and cause problems in the receiver until input powers are very high.

In a real application, the 80 MHz offset from the notch center could be due to bad control of the filter, but it could also happen if the interferer is modulated. Looking at cross-modulation in Figure 4.9, it is overall very weak for tones both at, below and above the notch, with suppression of about $-70\,\mathrm{dBc}$ at an interference power of 0 dBm (growing at twice the rate of the fundamental). Since the purpose of this device is to allow for listening to weak signals, these cross-modulation sidebands on those would be way below the noise-floor, not interfering with the intended modulation. Furthermore, cross-compression was not even measurable, as shown in Figure 4.10.

Based on the results of very good harmonic distortion IIP_2-figures and the fact that cross-compression was not measurable, it seems that the anti-series topology of the varactor diodes has been very effective in suppressing even-order distortion.

As mentioned above, cross-modulation is very weak, and the harmonic distortion analysis shows very good IIP_3-values (Table 4.1). This gives some hints on 3rd order distortion behavior. These can be further extended by analyzing the IMD3 results. Looking at Figure 4.8 and the corresponding Table 4.2, we first address the in-notch behavior. IIP_3 is rather low at 13.7 dBm, but there is also quite a lot of suppression due to the notch, with OIP_3 being $-37\,\mathrm{dBm}$. This is actually not a very interesting case, as the application doesn't require listening to closely separated channels within the notch. It is much more indicative of real-world performance to look at what happens to intermodulation of signals outside of the notch. Below the notch at 1000 MHz, the IIP_3 is excellent at 53.9 dBm. 200 MHz below the notch, IIP_3 is still excellent at 47.3 dBm. 200 MHz above the notch however, IIP_3 has dropped noticeably to 23.3 dBm. Unfortunately the signal generator capability was limited at 3 GHz, otherwise higher frequencies would have

been tested.

IMD3 is worse above the notch than below it. Above resonance, the impedance of the inductive element (short transmission line) becomes large. This confines the current to travel through the varactors instead. Due to their semiconductor bulk-resistance which this current has to flow through, there will be an RF-voltage across the PN-junction. Due to its nonlinear relationship between voltage and current, as described in Section 3.1.2, containing odd-order terms, third-order intermodulation products are created. IMD3 is expected to improve further away from the notch, as the current wave at resonance causing distortion lessens. The very good cross-modulation results lessen the concerns of bad 3rd order distortion. They are taken on a desired tone at 3500 MHz, which is a bit higher up in frequency (see Figure 4.9). These cross-modulation results also show that there is no catastrophic higher odd-order nonlinear distortion term.

If the IMD3-figure becomes too limiting, one could look at other varactor choices. For this specific design, mostly due to little familiarity early on with nonlinear dynamics, emphasis was put on having a varactor with a large capacitance ratio, giving a large tuning range. Intuitively a large capacitance ratio very likely means a higher series resistance, as there has to be more semiconductor length for the depletion-region to grow or shrink into. This is a hard trade-off and could steer a project towards using something like MEMS capacitor banks instead. One could also use a bulk capacitor (make sure it has no voltage coefficient!) in parallel with the varactor. This would reduce the current, and implicitly voltage, over the varactor junction, at the expense of tuning range. Maybe the best solution would be something like a MEMS capacitor bank for fixed steps, with fine-resolution tuning coming from a parallel varactor serving as a much smaller part of the total capacitance.

For resonators operating at frequencies pushing a couple of GHz, parasitic inductance starts to become challenging on a PCB. For the parallel resonator, it is desirable to have lots of capacitance and very little inductance to get narrow notches, as shown in Figure 2.4. Layout is difficult and made even harder with more components, such as for example a design with a varactor in parallel to a MEMS-bank. To actually address the parasitic inductance at these higher frequencies, more integration is probably the best choice. There are multi-varactor chips and flip-chip varactors available.

It is not ultimately possible to use symmetry to get rid of odd-order nonlinearities, as was possible for the even-order nonlinearities. This contributes to making third order intermodulation a classically difficult spec. The only real solution is to increase linearity, but it can be difficult and bring its own trade-offs. There is however potential for small actions to make big improvements, as third-order products grow at 3 times the rate of the fundamental, presented in Section 2.5.4.

5.3 End-to-End SDR Test

Unfortunately, the end-to-end SDR test failed to demonstrate that the filter is either helping or hurting system performance. The results are mostly within the noise floor. There is probably way too much leakage from transmitter to receiver inside the SDR for the DUT in the setup to make any difference. The only slightly promising result is in Figure 4.11 where the LO of the SDR was set to the same frequency as the interference. In that case the notch of my filter probably reduced the amount of interference reaching the receiver, giving my filter the edge. Otherwise there is no clear trend, as seen in Figure 4.13.

If attempted again, physically separated transmitters and receivers would be necessary to remove the internal SDR leakage making results useless. It would also be interesting to look at specifically placing the LO where harmonics of an interference source would show up. This could not be done with the current equipment because the filter is only tunable to a minimum notch frequency of 2 GHz, which if aligned on an interference tone would put the lowest harmonic at 4 GHz, which is higher than the maximum LO frequency of the SDR used at 3.8 GHz.

5.4 Asymmetric notch edges at higher bias voltages

In Figure 4.2, at higher bias voltages, the low-frequency edge of the notch develops a sloped section. The entire notch becomes asymmetric and there is a decrease of transmission at frequencies much lower than desired.

By coupling the node between the varactors to signal ground at the correct frequency, the behavior seen in Figure 4.2 can be replicated in QUCS-S. Figures 5.2 and 5.3 show replication of the asymmetric notch phenomenon by loading the node between the varactors with a series resonator. The only variable changed between figures 5.2 and 5.3 is the varactor junction capacitance C_j. At a lower bias point, such as in Figure 5.2, it is obscured by the desired notch. Figure 5.3 shows the asymmetry at the higher bias point.

It is not immediately apparent how coupling on the capacitive side of the resonator impacts transmission below the resonant frequency. Capacitors are thought to have high impedance at lower frequencies, and should thus prevent currents from going through that branch below resonance. By investigating the impedance of the varactor model used both throughout the project and in the QUCS-S model seen in figures 5.2 and 5.3, it becomes clear that it is actually not that high at less than 80 Ω for each varactor at 2.3 GHz where the bias-loading resonates, seen in Figure 5.4. This impedance is roughly reduced by half since there are two varactors in parallel, meaning that there is an about 40 Ω coupling to the unintended series-resonant bias feed even below the notch. It is reasonable that this causes the notch

asymmetry.

The precise values and topology of the coupling in this model is not highly important, and it depends on exact layout parasitics. These results still reasonably suggest that more care needs to be taken in the bias network, so that it does not have a series-resonant low impedance loading in-band. There is some discussion in Section 5.7 on how this might be achived with a different PCB stack-up, lowering parasitic inductance. The series-resonant frequency could also be increased a bit by reducing the trace lengths and impedances to decoupling components. For a second revision, some more EM-simulations are probably necessary to find the layout-dependent root cause. Maybe the bias topology needs to change.

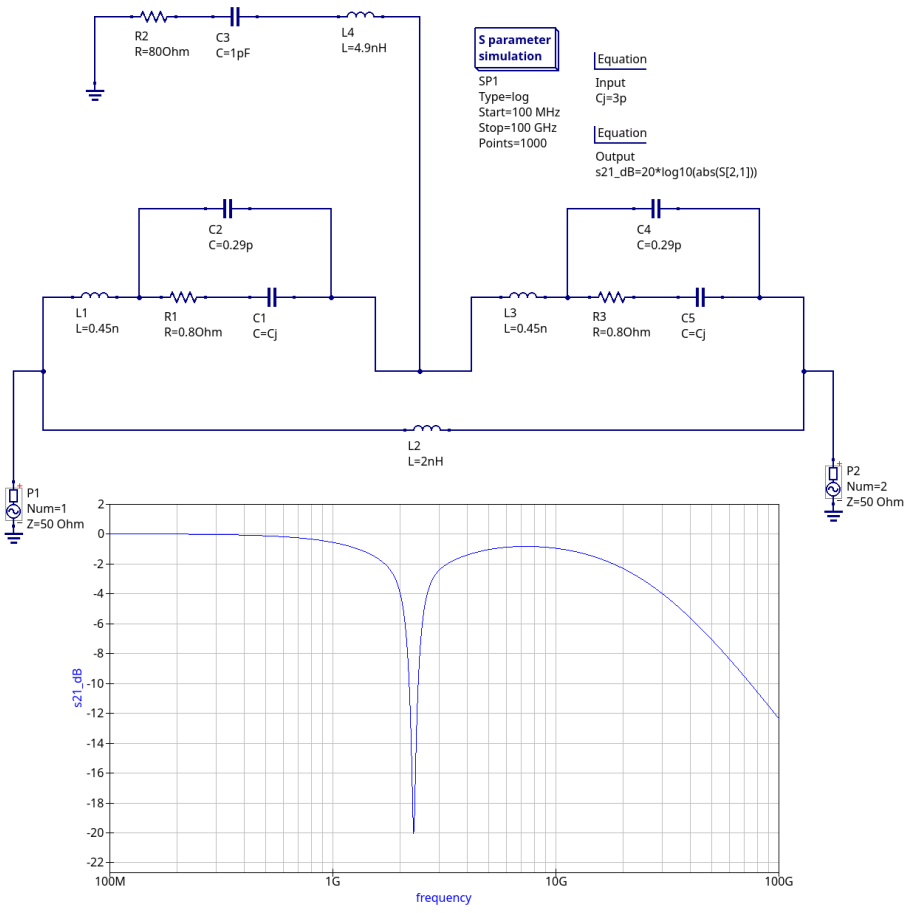

Figure 5.2: QUCS-S replication of asymmetric notch behavior at a low bias point, showing symmetric behavior of the notch here.

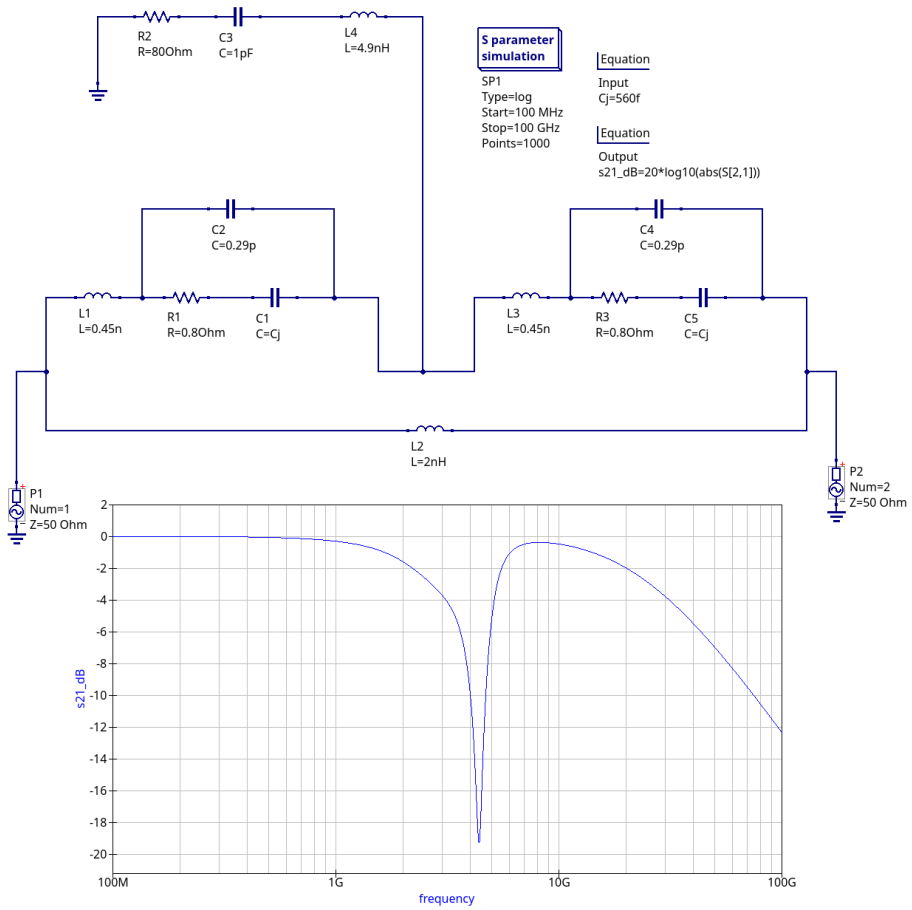

Figure 5.3: QUCS-S replication of asymmetric notch behavior at a high bias point, showing asymmetry.

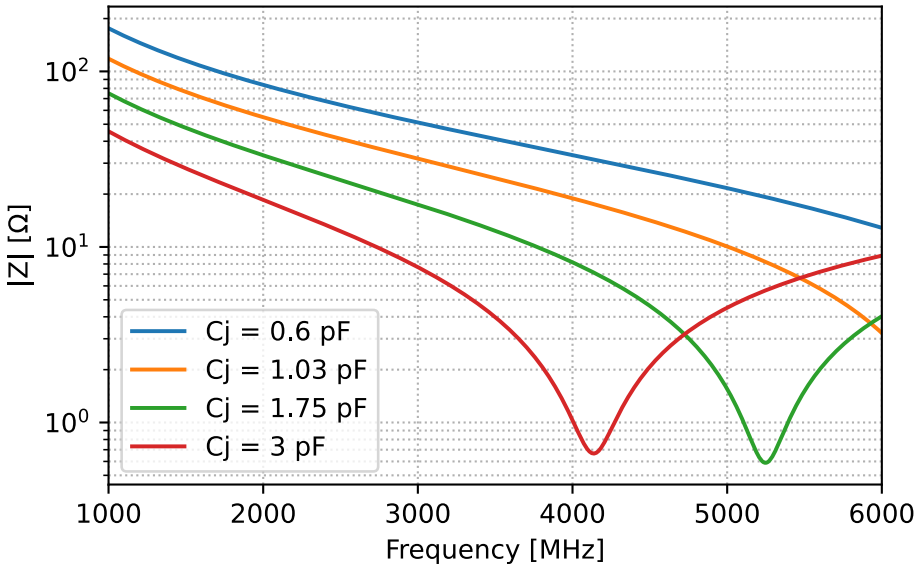

Figure 5.4: Impedance of one of the varactor models used both throughout the project and in the QUCS-S model shown in figures 5.2 and 5.3

5.5 Results at resonance are sensitive to spectrum analyzer attenuator

Outside of resonance, the level of spectrum analyzer attenuation has almost no impact on measured tone levels. At resonance, it has quite a bit of impact. If, for example, the IMD3-analysis seen in Figure 4.8 is done with many different attenuator levels, Figure 5.5 is produced. Looking at the 2800 MHz series for the fundamental per-tone power, the lines stack. The lines that are below correspond to lower attenuation, with the lowest being 5 dB. The lines on top bunch together, and the highest one corresponds to the max used attenuation of 30 dB. For the measurement series' outside of resonance, the level of attenuation has almost no effect. Two possible explainations are:

1. The automatic gain control (AGC) of the signal generator is disrupted by reflections from the device, causing it to think the output power is too high, lowering it erroneously.
2. The spectrum analyzer is not measuring power levels correctly because its impedance match to the device is bad at filter resonance.

Theory 2 likely causes most of the observed effect. Figure 5.6 shows what different loads look like through different attenuators, much like what the spectrum analyzer front-end would experience. The filter at resonance looks like a very large series impedance, which has a big reflection coefficient through small attenuators.

The front-end of the spectrum analyzer seems to be a bit sensitive to the source impedance, and when a small attenuator is used it doesn't interface with a good enough impedance match to take a correct reading. When more attenuation is used, the active front end experiences something much closer to a good match, and readings become accurate. Since the relative difference in power level (dB difference) is the same for different input power levels, the phenomenon is due to linear mismatch, not some nonlinear behavior, as that would vary with input tone power swept.

Theory 1 is less likely to contribute to the phenomenon, though since the variability of attenuator was constrained to the spectrum analyzer side, we have no way of controlling for this with present data. At resonance there is quite good isolation between signal generator and spectrum analyzer, making the AGC interaction unlikely. It could however cause a systematic error in all measurements, depending on how well isolated the output stage of the signal generator is. The stacked lines phenomenon seen in Figure 5.5 is also visible when doing the same attenuation sweep but for harmonic distortion, which has a different input-side setup without power combiner etc, causing me to not suspect that it is due to some kind of standing wave throughout the measurement setup passing through the filter.

In the future it would be best practice to by default use some level of attenuation to control impedance interfaces, especially when working with reflective devices. One would have to keep in mind the loss of noise headroom as well as higher power requirements of signal sources, pushing them towards nonlinearity. It is good that different attenuator levels were tested so that this problem could be discovered. This is very much a benefit of test automation, allowing more measurements to be taken.

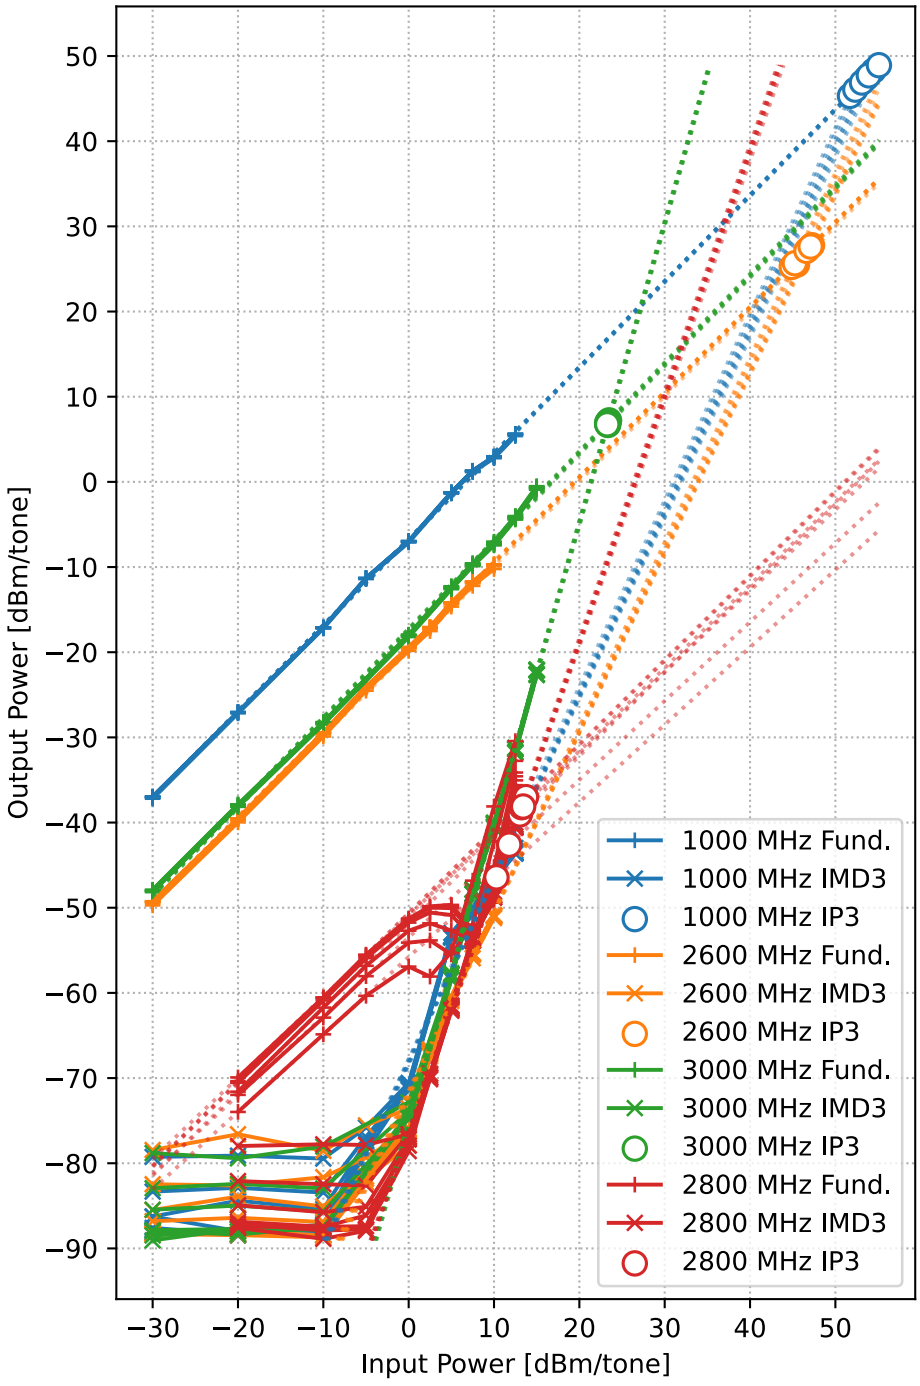

Figure 5.5: IMD3 analysis like in Figure 4.8 but with results from all attenuator values

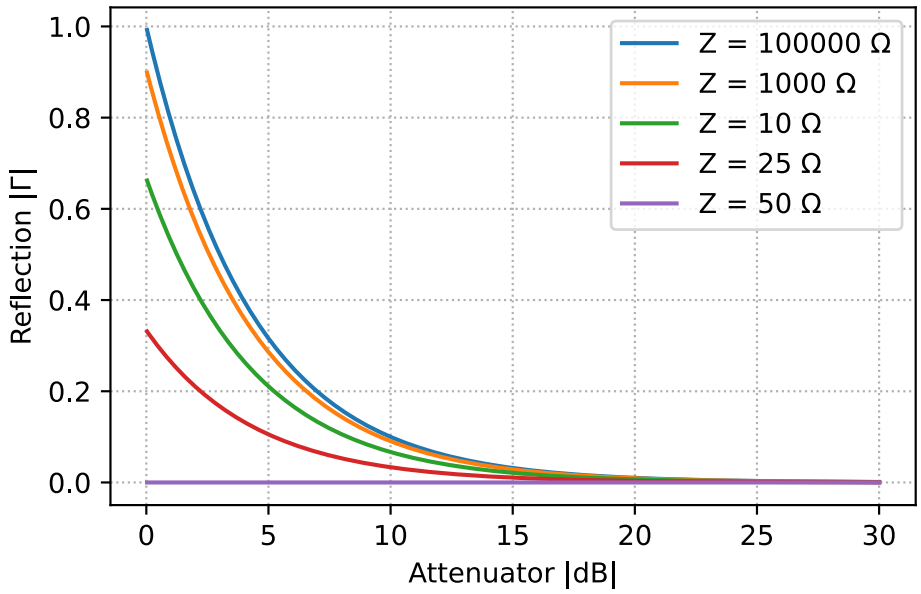

Figure 5.6: Loads seen through different attenuators

5.6 Faster-than-unity increase of fundamental power level after compression

Over about 6 dBm for an in-notch tone at 2.8 GHz, which is beyond the onset of compression, the transmitted tone suddenly starts increasing faster than unity. This is seen in figures 4.4, 4.5, 4.6, 4.7 and 4.8.

The most likely explaination for this is that the notch either detunes, broadens or both at these higher input powers. The results which were discussed in Section 5.2 suggest that the odd-order nonlinear terms are more pronounced than the even-order ones. This would direct the explaination towards notch broadening as it predominantly is caused by 3rd order nonlinearities. Intuitively 3rd order terms cause energy to be spread across the spectrum, resulting in a shallower notch. Symmetry is kept for the positive and negative half-cycles of the signal, causing no net shift in operating point.

Even-order terms would cause the average operating point to shift depending on input power, as the positive and negative half-cycles of the signal are propagated differently when power increases, causing detuning due to a shift in the average bias point from asymmetry. Even-order terms are thought to be quite suppressed, as discussed in Section 5.2. This suggests that detuning is weak, and that most of

the effect is due to notch broadening.

There is on the other hand some evidence of detuning when looking at the 2800 ± 80 MHz-series in Figure 4.4. As input power increases, the output power of a tone at 2720 MHz decreases slightly, while the output power of a tone at 2880 MHz increases slightly. In this case the notch is in the middle between them. This suggests that the notch detunes slightly downwards at higher power levels.

Both effects are likely present, with most dominance being reasonably attributed to notch broadening. This could be further investigated by power-swept VNA measurements or more elaborate pilot-tone or notch-tracking techniques.

5.7 Dielectric Thickness

A thicker dielectric was chosen early in the project to get more inductance per unit length for the inductive transmission line, as described in Section 3.1.6. After the PCB was ordered and design of the decoupling network was optimized, it became clear that excess series inductance makes it harder to design a good decoupling network. The design would probably have benefitted from having a thinner dielectric to get less via inductance, as well as the inductance from all traces being lower for the same width and length. Furthermore, a thinner dielectric would decrease the taper ratio needed to enter the pads, likely reducing non-filter reflections.

The short transmission line assumption holds for quite a bit higher frequencies than strictly necessary in the current design, so it would probably not be too bad to make the inductive transmission line longer to get the required inductance in this case, if as much inductance is desirable at all. The inductance of the inductive transmission line should probably also be minimized more to get a higher Q-factor, as discussed in Section 5.1. For a future iteration, a thinner dielectric would for these reasons likely be very beneficial.

For more advanced PCBs operating at higher frequencies, it is common to have an extremely thin dielectric (around 50 µm), most often between power- and ground layers, and sometimes for signal layers too [30]. This gives significantly thinner microstrips for a given characteristic impedance, as shown in Section 2.4. Would a too thin dielectric make the microstrip lines so small that they start to get lossy due to resistance? Yes, but it would not be the limiting factor on Q for this design. The KiCad stripline calculator gives a loss of around 0.24 dB in such an inductive transmission line, limiting Q to about 100. This is much higher than the limitation from varactor series resistance, as discussed in Section 5.1, and would thus not be the limiting factor.

This kind of thin dielectric between power- and ground planes would be very ben-

eficial for the decoupling performance, as it would remove the need for inductive microstrips to access the highest-frequency decoupling capacitance by providing plate capacitance directly between copper planes. A thin dielectric could probably go far towards creating a biasing network that solves the asymmetric notch problem discussed in Section 5.4. The main downsides would be higher cost, lower breakdown voltage and more risk of damage during soldering and handling.

5.8 Notch Frequency Control

When intending to put the notch and the tone on the same frequency, an error of around 4 MHz was usually present. This is seen in, for example, figures 4.4 and 4.8. There was no closed-loop control of the bias voltage for notch frequency, and a rather cheap multimeter was used with only 4 significant digits at the relevant range. The more accurate model of notch frequency as a function of bias voltage seen in Figure 4.1 was not created until after the lab measurements. During measurements, the expected frequency of the notch was taken from single measurements with the VNA. Thought was put into letting everything warm up and settle, but over the course of a few weeks the setup could definitely drift. In combination with the limited multimeter resolution, and probably also absolute accuracy, an error is present.

5.9 Asymmetry in IMD3 tone levels

Figure 5.7 shows one raw spectrum from IMD3 measurements. There is typically up to a 5 dB difference in input tone power, despite the signal generators being set to the same output power. This deviates from the usual assumption when measuring IMD3 that the tones should be of equal power. If the sources are calibrated well, then the difference of tone power is likely due to bad isolation between them. This would cause power from one to go backwards into the output stage of the other, possibly disturbing its internal leveling loop. More attenuation going into the power combiner would probably have helped, or even dedicated isolators. No extra isolation or attenuation was used. More care should be taken when setting up the measurements to ensure that the tone powers are equal.

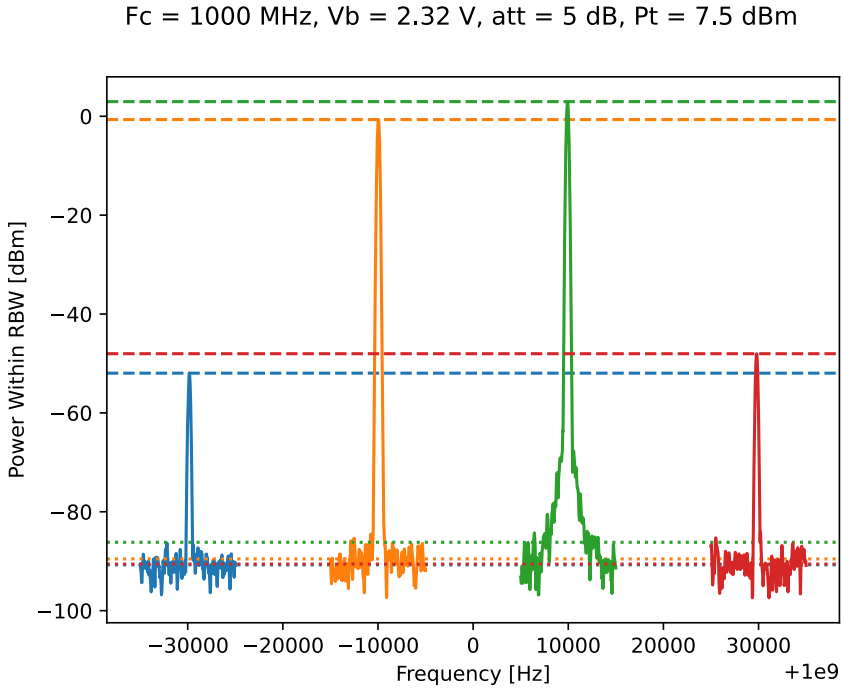

Figure 5.7: Raw IMD3 spectrum

5.10 Memory Effects

A small tone-spacing of 20 kHz was used during IMD3 measurements. This was chosen due to the desire to have similar small-signal behavior of the tones even at the quickly changing impedance at and around the notch. Memory effects, as presented in Section 2.5.5, could be excited at the beat frequency and it would be interesting to investigate whether these are present. The tone spacing could be varied, and if memory effects are present this would change IMD3-behavior. Asymmetric intermodulation products could also reveal memory effects, but those are obscured by the input tone asymmetry discussed in Section 5.9.

For continued development, further investigations of memory effects would be appropriate since 3rd order distortion likely is a limiting factor, as discussed in Section 5.2. This combined with the clear coupling to the bias network, a possible source of memory, as discussed in Section 5.4, makes this interesting.

5.11 Effects of uncontrolled PCB impedance at higher frequencies

As seen in Figure 4.2, above about 6 GHz the PCB and connectors etc. no longer transmit signals in a reliable impedance-matched way. This has implications for all measurements taken at these higher frequencies. The most potentially impacted results are harmonic distortion, seen in figures 4.5, 4.6, and 4.7, as they include measurements up to 8640 MHz. It is however possible that favorable transmission happend to occur around 8.4 GHz, as seen in Figure 4.2.

Chapter 6

Conclusion

The objective of this thesis was to solve the problem of strong narrowband interference reducing performance at all other frequencies in a wideband receiver. This was attempted by implementing a tunable bandstop filter to be placed early in the signal chain, to lower the power of interference before it ever reaches the receiver.

The filter was based on a series-connected parallel-resonator architecture, using varactor diodes as the tunable capacitance element. It was developed mainly using QUCS-S, HFSS and SciKit-RF, and then fabricated on a PCB.

A measured tuning range of 2 – 4.4 GHz was achieved with a notch depth of at least 24 dB across the entire tuning range, with some frequencies exceeding 60 dB. Unfortunately, the notch in this implementation is wider than desired, and has asymmetric edges. By reproducing this deviation in simulations, evidence points to more care being needed in the decoupling network design, as well as in the modelling of parasitics to get an optimized solution.

At 2600 MHz, which is 200 MHz below the notch tuned at 2800 MHz, IIP_3 is excellent at 47.3 dBm for IMD3, suggesting practically no intermodulation impact for signals below the notch. 200 MHz above the notch, at 3000 MHz, IIP_3 is lower at 23.3 dBm. This is still good, but not excellent, and demonstrates a fundamental trade-off between linearity and tuning range for varactor-based tunable filters. Unfortunately due to equipment limitations, it was not possible to test IMD3 at higher frequencies. IMD3 is however expected to improve further away from the notch.

Cross-modulation is very low at -70 dBc for a 0 dBm AM-modulated interferer in the notch, alleviating many of the third-order distortion concerns from the limited IMD3 measurements. It also directly provides confidence that the filter will not transfer interference modulation to other signals, preserving sensitivity to

the intended signal.

The fact that cross-compression was not observed combined with the excellent second-order harmonic distortion results of $\mathrm{IIP}_2 > 41\,\mathrm{dBm}$ suggests that the anti-series varactor topology was very effective in reducing even-order nonlinearities.

In conclusion, results of this work show that this type of tunable bandstop filter can be useful in preserving receiver sensitivity in the presence of strong narrowband interference.

Future work may include reducing the notch width by decreasing the real inductance to capacitance ratio or lowering loss, as well as doing proper integration to ensure a wideband and low insertion loss implementation.

References

[1] International Telecommunication Union, "IMT Traffic Estimates for the Years 2020 to 2030," ITU, Tech. Rep. Report ITU-R M.2370-0, 2015. [Online]. Available: https://www.itu.int/pub/R-REP-M.2370

[2] Michael Aygur and Sithamparanathan Kandeepan and Andrea Giorgetti and Akram Al-Hourani and Edward Arbon and Mark Bowyer, "Narrowband Interference Mitigation Techniques: A Survey," *IEEE Communications Surveys & Tutorials*, 2025.

[3] Michael S. Ricciardi and Pierre Souque, "The Electromagnetic Spectrum: A Strategic Domain in Modern Warfare," *PRISM, National Defense University*, vol. 9, no. 3, pp. 122–139, 2021. [Online]. Available: https://ndupress.ndu.edu/Portals/68/Documents/prism/prism_9-3/prism_9-3_122-139_Ricciardi-Souque.pdf

[4] Behzad Razavi, *RF Microelectronics*, 2nd ed. Upper Saddle River, NJ: Prentice Hall, 2012.

[5] Chippalkatti, Vinod S and Biradar, Rajashekhar C and Rana, S S, "Technology Trends in Transmit Receive Modules for Synthetic Aperture Radar Satellites," in *2022 IEEE Fourth International Conference on Advances in Electronics, Computers and Communications (ICAECC)*, 2022, pp. 1–6.

[6] Robert E. Collin, *Foundations for Microwave Engineering*, 2nd ed. Piscataway, NJ: IEEE Press and Wiley, 2001.

[7] David M. Pozar, *Microwave Engineering*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2012.

[8] Skyworks Solutions, Inc., "Varactor SPICE Models for RF VCO Applications," Tech. Rep. 200315C, Aug. 2015, application Note.

[9] Lee, T.P., "Evaluation of voltage dependent series resistance of epitaxial varactor diodes at microwave frequencies," *IEEE Transactions on Electron Devices*, vol. 12, no. 8, pp. 457–470, 1965.

[10] Hammerstad, E. and Jensen, O., "Accurate Models for Microstrip Computer-Aided Design," in *1980 IEEE MTT-S International Microwave symposium Digest*, 1980, pp. 407–409.

[11] Walter Rudin, *Principles of Mathematical Analysis*, 3rd ed. New York: McGraw-Hill, 1976.

[12] Stephen A. Maas, *Nonlinear Microwave and RF Circuits*, 2nd ed. Norwood, MA: Artech House, 2003.

[13] SciKit-RF Developers, "SciKit-RF: Object Oriented Microwave Engineering," Python package, 2025. [Online]. Available: https://scikit-rf.org

[14] Qucs-S development team, "Qucs-S: Quite Universal Circuit Simulator with SPICE," Open-source software, 2025. [Online]. Available: https://ra3xdh.github.io

[15] Ansys Inc., "Ansys HFSS," Canonsburg, PA, USA, 2025, electromagnetic simulation software. [Online]. Available: https://www.ansys.com

[16] NGspice development team, "ngspice circuit simulator," Open-source software, 2024, 42. [Online]. Available: https://ngspice.sourceforge.io

[17] Liu, Jun S., *Monte Carlo Strategies in Scientific Computing*. New York: Springer, 2004.

[18] Deb, Kalyanmoy, *Multi-Objective Optimization Using Evolutionary Algorithms*. Chichester, UK: John Wiley & Sons, 2001.

[19] D. Schreurs and N. B. Carvalho, Eds., *Microwave and Wireless Measurement Techniques*. Cambridge, UK: Cambridge University Press, 2013.

[20] V. Teppati, A. Ferrero, and M. Sayed, Eds., *Modern RF and Microwave Measurement Techniques*, ser. Cambridge RF and Microwave Engineering Series. Cambridge, United Kingdom: Cambridge University Press, 2013.

[21] Travis F. Collins and Robin Getz and Di Pu and Alexander M. Wyglinski, *Software-Defined Radio for Engineers*. Artech House, 2018.

[22] Mohammadi, Laya and Koh, Kwang-Jin, "Low power highly linear bandpass/band-stop filter for 2-4 GHz with less than 1% of fractional bandwidth in 0.13 µm CMOS technology," in *2017 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2017, pp. 292–295.

[23] Quan, Guohai and Lin, Yinhan and Xu, Taotao and Li, Zhuming and Liao, Shaowei and Yi, Xiang and Qin, Pei and Zhu, Haoshen and Xue, Quan, "A 28/39-GHz Dual-Band Multi-Mode LNA Supporting Wideband/Concurrent/Reconfigurable Operation in 40-nm CMOS," *IEEE Transactions on Circuits and Systems I: Regular Papers*, pp. 1–13, 2025.

[24] Pietro Andreani and Sven Mattisson, "A 2.4-GHz CMOS monolithic VCO with an MOS varactor," *Analog Integrated Circuits and Signal Processing*, vol. 22, no. 1, pp. 17–24, 2000.

[25] Tanbakuchi, H. and Nicholson, D. and Kunz, B. and Ishak, W., "Magnetically tunable oscillators and filters," *IEEE Transactions on Magnetics*, vol. 25, no. 5, pp. 3248–3253, 1989.

[26] Iannacci, Jacopo, "RF-MEMS for high-performance and widely reconfigurable passive components – A review with focus on future telecommunications, Internet of Things (IoT) and 5G applications," *Journal of King Saud University - Science*, vol. 29, no. 4, pp. 436–443, 10 2017.

[27] David J. Griffiths, *Introduction to Electrodynamics*, 5th ed. Cambridge, UK: Cambridge University Press, 2024.

[28] Klopfenstein, R. W., "A Transmission Line Taper of Improved Design," *Proceedings of the IRE*, vol. 44, no. 1, pp. 31–35, 1956.

[29] *RO4003C and RO4350B Laminates with TICER TCR Thin Film Resistor Foils Data Sheet*, Rogers Corporation, Chandler, AZ, USA, Aug. 2019, publication #92-134, Revised 1446 080819.

[30] Eric Bogatin, *Signal and Power Integrity—Simplified*, 3rd ed. Upper Saddle River, New Jersey: Prentice Hall, 2018, includes signal integrity, power distribution network design, and decoupling techniques.