Hybrid RF-Digital Feed-Forward Filter for High-Order Frequency Agile Filtering

by

Mark Wyville, B.Eng., M.A.Sc.

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in Electrical Engineering

Ottawa-Carleton Institute for Electrical and Computer Engineering

Department of Electronics

Carleton University

Ottawa, Ontario

June, 2011

Your file Votre référence
ISBN: 978-0-494-83234-9
Our file Notre référence
ISBN: 978-0-494-83234-9

Abstract

Wireless network operators, spectrum managers and wireless infrastructure vendors have generated demand for frequency agile RF filters. The current state of the art frequency agile RF filters are limited to low filter order.

This thesis introduces a frequency agile RF filter that is fundamentally different from the current state of the art. Digital and RF systems are combined in a feed-forward architecture. This architecture provides a frequency response with multiple tunable transfer function zeros. This feed-forward architecture is called the hybrid RF-DSP feed-forward filter.

The digital subsystem is placed in one of the feed-forward paths, and includes a digital filter. The digital filter permits the system to have a high filter order. The digital system is bookended with frequency down and up conversion stages to permit the digital filter to be used at any RF frequency. The second feed-forward path consists of an RF system, and in the simplest case is just a transmission line. This path provides a high dynamic range route through the filter, specifically for passband signals.

Noise and attenuation analyses are performed that incorporate the non-linear effects of the digital components into conventional RF performance metrics. The noise analysis demonstrates the potential for the filter to exhibit less than 1 dB of noise figure.

RF measurements are used to demonstrate proper operation of a hardware prototype. The prototype is targeted towards the RF front-end of a receiver with a dynamic interference environment at its input. Measurements demonstrate interferer suppression for operation at several bands between 800 and 1800 MHz. The group delay mismatch between the feed-forward paths is reduced with a delay line to achieve between 32 to 24 dB of

attenuation for interferers with bandwidths from 1 to 8 MHz. The largest digital filter used within the prototype has 64 complex taps, thereby providing the hybrid RF-DSP FF filter with 64 tunable transfer function zeros.

The hybrid RF-DSP feed-forward filter represents a fundamental change to frequency agile RF filter design. This fundamental change is based on harnessing the ever increasing processing power of DSP instead of the more traditional approach of incrementally improving the quality factor of tunable resonators.

Acknowledgements

I am grateful to my academic supervisor, Jim Wight, who many years ago advised me to pursue graduate studies on what I thought was my last day at university. I would like to thank him for all of the support he has made available in a number of ways.

Also from Carleton University I would like to express my gratitude towards Nagui Mikhail, the department's 'hidden hero' for all of the time and advice he provided to help me with a diverse range of issues.

At Ericsson I would like to acknowledge Russell Smiley for his leadership and for volunteering for the workload associated with overseeing a graduate student outside of the university.

At home I would like to thank my wife for her patience and loving support.

Table of Contents

List of Figures

List of Abbreviations

ADC	analog to digital converter
AR	auto regressive
ARMA	auto regressive moving average
BLMS	block least mean squares
BPF	bandpass filter
BSF	bandstop filter
BTS	base station transceiver
BW	bandwidth
CW	continuous wave
DAC	digital to analog converter
dBFS	decibel full-scale
DC	direct current, or zero frequency
DEMOD	demodulator
DSP	digital signal processing
FCC	Federal Communications Commission
FDD	frequency division duplex
FF	feed-forward
FIR	finite impulse response
FPGA	field programmable gate array
FS	full-scale
FXLMS	filtered-x least mean squares
FXBLMS	filtered-x block least mean squares
HPF	highpass filter
IIP3	input third-order intercept point
IIR	infinite impulse response
IL	insertion loss
IMD	intermodulation distortion
LCMV	linear constrained minimum variance
LO	local oscillator
LMS	least mean squares
LNA	low noise amplifier
LPF	lowpass filter
MEMS	micro-electro-mechanical systems
MMSE	minimum mean square error
MSB	most significant bit
NGD	negative group delay
PA	power amplifier
PAPR	peak to average power ratio
PSD	power spectral density
Q	quality factor
RLC	resistor inductor capacitor
RLS	recursive least squares
RX	receive
SAW	surface acoustic wave

SNR	signal to noise ratio
T-BSF	tunable bandstop filter
TX	transmit
VGA	variable gain amplifier

1 Introduction

Filtering in the RF front-end of a base station transceiver has demanding requirements. These requirements include high-order filtering, high dynamic range, and low loss. These requirements had not included frequency agility or bandwidth reconfiguration due to limitations of the state of the art filtering technologies. Frequency agility and bandwidth reconfiguration give an additional degree of freedom to spectrum regulators, wireless service providers and base station equipment vendors, which facilitate wireless systems to evolve along a new dimension.

A novel RF filter has been developed that meets the traditional demanding requirements for the RF front-end of a base station, plus it has the distinguishing characteristics of frequency agility and bandwidth reconfiguration. This new filter is classified as a hybrid RF-DSP FF filter since it has an RF and a DSP path in parallel in a feed-forward configuration. This thesis details the analyses, development, and prototyping of the hybrid RF-DSP FF filter.

1.1 Motivation

Conventional RF front-end filters for base station transceivers (BTS) are high-order bandpass filters (BPF) constructed with multiple high quality poles. These filters are not electronically reconfigurable and are relatively bulky. With the adoption of multiple antenna and transceiver architectures in the wireless community, these filters need to

reduce in size. Furthermore, demand for these filters to exhibit frequency agility exists, and has many drivers. These drivers come from BTS vendors, wireless network operators and spectrum managers.

Vendors of BTS's must offer front-end filtering systems that operate in the bands that their customers are using. This results in a large ensemble of potential filter specifications. If the filter was reconfigurable, then the same device would satisfy the specifications of multiple customers. The vendors would be able to get the radios to market faster.

Wireless network operators must adopt a specific wireless standard and frequency band before building their network since parts of the radios are band specific, including the front-end filtering. Investing in equipment that works for a single wireless standard, in a specific frequency band has inherent risk. This risk would be reduced if the radio equipment were reconfigurable for different standards and frequencies.

Spectrum managers are constrained by many factors, one being existing technology. Cognitive and software defined radio concepts have evolved, but the lack of a reconfigurable radio remains to be the Achilles' heel of these concepts. Software defined radio represents the technology that is required in the radios to implement frequency agility. Cognitive radio is a wireless concept that makes the radio aware of its interference environment, and allows the radio to adapt based on the environment [*1*]. This concept increases spectral efficiency by allowing radios to use spectrum that is unoccupied by the primary user. This task requires the radio to sense potential channels for interference, and then operate in some of those channels. Cognitive radio requires the

transmitter and receiver to be frequency agile and the receiver to have sufficient sensitivity for reliable interference sensing.

In November 2010, the FCC released a notice of inquiry to seek comments on how dynamic spectrum access radios and techniques can increase spectral efficiency, and how to effectively manage the spectrum [*2*]. This notice of inquiry also reviewed previous instances where the FCC has adopted rules to implement cognitive radio. One instance was initiated in 2003, where 255 MHz of unlicensed bandwidth was allocated near 5 GHz in a detect-and-avoid protocol. There were military and weather radar operating within these bands that were not to be interfered with. In 2008 access was given to 'white spaces', which is the unused spectrum in the broadcast TV band. Initially, interference avoidance in white spaces was performed using two methods simultaneously. The first required access to an FCC administered database that tracked frequency and geo-location of each spectrum user. The second method was spectrum sensing. Current technology was deemed inadequate for the spectrum sensing method, so in September 2010 the requirement for spectrum sensing was removed.

Spectrum refarming is a technique used in spectrum management. Spectrum refarming recovers spectrum from existing wireless operators and reallocates the operators to another band. Refarming also includes phasing out a wireless standard in a specific band and reallocating that same band with a more spectrally efficient radio service. Historically, the migration to new frequency bands has been voluntary and the transition times have been based on expiration of spectrum licenses. Fast spectrum refarming would be feasible if radios had reconfigurable front-ends. This technology would allow

spectrum regulators to allocate spectrum to innovative and risky services without the same risk associated with fixed radios. If the new service is unsuccessful, then the regulator can refarm the spectrum and the operators can reconfigure their radios for a different service.

Another topic related to spectrum management is secondary spectrum trading. This is the process where a spectrum licensee allows another operator to use their spectrum. This situation may be beneficial when the licensee's spectrum is being underutilized and another operator cannot meet their customer demand. This practice is also called spectrum leasing, and may be the only affordable method for small wireless service providers to access spectrum. The providers would benefit from equipment that can operate over different leased bands, thereby requiring frequency agility.

Demand exists for BTS filtering where the centre frequency and bandwidth are reconfigurable. These new constraints are in addition to the long standing constraints of high-order, low loss and high dynamic range.

1.2 Current State of the Art

There is an abundance of reconfigurable RF filter designs in the literature. Continuous-time filters employ RLC resonators with tunable reactive components. Many tunable continuous-time RF filters rely on low-Q tuning components, due to limitations in current technology. Other tunable continuous-time RF filters integrate high-Q tuning components into high-Q resonators. The integration process is complex and degrades the Q of the overall resonator. The maximum allowable insertion loss specification of a continuous-time filter limits the minimum Q factor of each resonator for a specific filter order and

bandwidth. It is shown in the following chapter that the current state of the art in reconfigurable continuous-time RF filtering does not meet the minimum Q factor requirements for a BTS front-end. Other tunable continuous-time filters use switches to connect or disconnect high-Q components to the filter. Filters based on this technique are limited in reconfiguration. In addition, insertion loss increases as the range of reconfiguration is expanded.

Discrete-time filters consist of multiple paths, each with a different time delay and complex gain – like a digital filter. RF discrete-time filters use tunable phase shifters, attenuators or vector modulators in each path. High-order filtering requires a large number of paths. But, each path has components that have frequency dependent behaviour and drift with environment and age. Furthermore, high-order filtering would require many wideband splitters and combiners. The monitoring circuitry for compensating the drift, the RF delays for each path, and the splitters and combiners inherently occupies a lot of space. Consequently, these types of configurations are limited to low-order filters.

The current state of the art reconfigurable filters cannot perform high-order filtering while meeting typical bandwidth and the demanding insertion loss specifications in a BTS front-end. Chapter 2 reviews these filters in more detail. Chapter 3 and beyond describes the filter developed in this work, named the hybrid RF-DSP FF filter, which performs high-order filtering while meeting typical bandwidth and stringent insertion loss specifications.

1.3 Thesis Objectives

The overall objective of this thesis is to develop a frequency agile RF filter with a high filter order. Furthermore, the noise figure and dynamic range performance should be on the same order of magnitude as typical fixed-frequency RF front-end filters. To this end, the following thesis objectives are met:

1. Develop the filter architecture.
2. Perform noise and attenuation analyses of the architecture. The resulting expressions determine the limitations of the filter, and develop into design equations.
3. Develop an adaptive control system for the filter.
4. Build a hardware prototype to validate the architecture, performance analyses, and adaptive control system.

1.4 The Hybrid RF-DSP FF Filter

For the first time ever a hybrid RF-DSP FF filter has been designed for operation at RF. One configuration of this filter is shown in Fig. 1.4.1. The FF configuration consists of a DSP path in parallel with an RF path. The FF configuration permits high power handling and large attenuation over desired stop-bands. The DSP path provides a means for high-order adaptive digital filtering, while the RF path provides a low distortion path over a large dynamic range. The down and up frequency conversion stages that bookend the digital filter allow any frequency band to be processed by the digital filter, hence permit frequency agility. The bandstop filter (BSF) increases the dynamic range of the lower

path, which would otherwise be limited by the dynamic range of the digital components. The dynamic range of the entire filter is only as large as that of the lower path.

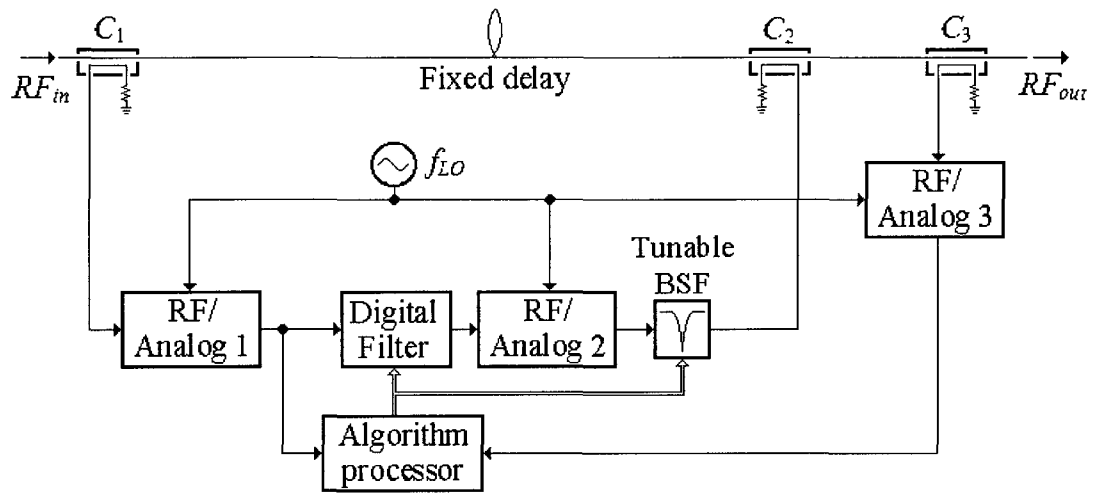

Figure 1.4.1 – System-level schematic of hybrid FF filter

The following list is a brief description of the blocks in Fig. 1.4.1:

- RF/Analog 1 – A direct down-conversion receiver consisting of a gain block, a down converter, anti-aliasing LPF's and ADC's.
- Digital Filter – Adaptive FIR filter in an FPGA.
- RF/Analog 2 – A direct up-conversion transmitter consisting of DAC's, reconstruction LPF's, an up converter and a gain block.
- Tunable BSF – Low order tunable bandstop filter.
- RF/Analog 3 – Identical to RF/Analog 1. This block is used to feedback the output signal for adaptation of the digital filter.
- Main path – The upper path in this FF system provides a low-loss path at RF.

- Algorithm processor – Adaptive control algorithms have been developed which are modified versions of standard adaptive filtering algorithms. These algorithms are run on an FPGA, and adapt the digital filter and tune the BSF.

The hybrid RF-DSP FF filter developed in this work is not the first system that uses an RF and DSP path in parallel. Parallel hybrid RF-DSP systems have been used in feedback configurations for particle accelerator apparatuses [*3*]. The bandwidth of these systems has been limited to a few MHz. The configuration in Fig. 1.4.1 bears some resemblance to a typical single channel active noise control system [*4*]. In an active noise control system the signal combining and tapping in the upper FF path is done in the acoustic domain with speakers and microphones.

1.4.1 Overview of Main Results

The hybrid RF-DSP FF filter is fundamentally different than other RF filters. Theoretical noise and attenuation analyses are reported with the purpose of relating the design variables to the filter performance. Due to the inclusion of a digital system, the RF noise figure and attenuation performance depend on quantization effects in the data converters and digital filter. The noise analysis combines these digital quantization effects into an RF noise figure. The derived noise figure expression is verified with simulations. The noise figure expression is used to determine the conditions required for the hybrid RF-DSP FF filter to exhibit a noise figure less than 1 dB.

A hardware prototype was built with evaluation boards and connectorized components. The prototype is targeted towards the RF front-end of a receiver in the presence of a dynamic interference environment. Measurements demonstrate interferer suppression for

operation at several bands between 800 and 1800 MHz. The largest digital filter used within the prototype had 64 complex taps, thereby providing the hybrid RF-DSP FF filter with 64 tunable transfer function zeros, at RF. A delay line was added to the RF feed-forward path which helped to achieve between 32 to 24 dB of attenuation for interferers with bandwidths from 1 to 8 MHz. Measurements of the hardware prototype were also used to validate the output noise expressions from the noise analysis. The noise expressions permit a designer to choose the design variables without requiring lengthy simulations for different filter realizations.

The hybrid RF-DSP FF filter represents the state of the art in high-order frequency agile RF filtering. Measurements over the cellular bands have demonstrated this filter's utility, and derived design equations isolate the relationships between the design variables and filter performance.

1.5 List of Contributions

The contributions made during this research are listed below:

1. A frequency agile RF feed-forward architecture with a digital filter in one of the feed-forward paths with applications for frequency agile RF filtering and duplexer isolation improvement.
2. Noise and attenuation analyses of this architecture that combines the effects of the RF and digital systems. Also, first-order approximation expressions for dynamic range and attenuation bandwidth performance.
3. Addition of an IIR stage within the digital filter to improve near-band attenuation performance, and provide the means to reconfigure the passband bandwidth.

4. Proposed addition of a negative group delay circuit in the feed-forward path containing the digital filter to improve wideband attenuation performance.
5. Modification to active noise control adaptive algorithm for the hybrid RF-DSP FF filter. The modification allows for a simpler relationship between the stability criteria and the allowable estimation errors of the $h_{23}[n]$ channel.
6. Addition of notch constraints in the digital filter to reduce the output noise due to the *RF/Analog* 1 chain and primary ADC's.

Two patent applications were filed by Nortel Networks with the Canadian Intellectual Property Office based on the contributions. The author of this thesis is the single author on both of these patent applications [*5,6*].

1.6 Thesis Organization

This thesis is divided into 9 chapters.

- Chapter 2 – is a survey of the state of the art technologies that are relevant to reconfigurable filtering.
- Chapter 3 – introduces the system down to the component level. This chapter includes different applications for this new filter.
- Chapter 4 – is a presentation of the relationships between the design parameters and the noise and attenuation performance. These analyses are verified by simulation.
- Chapter 5 – describes the adaptive algorithms used to control the digital portion of the filter to handle a dynamic input signal scenario, and drifting of the RF and analog components.

- Chapter 6 – introduces the hardware prototype, and reports preliminary measurements, which characterize the bottom FF path, and demonstrate proper operation of the adaptive algorithms.
- Chapter 7 – reports measurements of the hardware prototype with different input signal scenarios and filter realizations.
- Chapter 8 – includes discussions of the results and recommendations for future work. Also summarizes the thesis and lists contributions.

2 Literature Review

High-order filters are required in the RF front-end of base station radios with low loss and distortion across the passband and large attenuation over the stop-bands. The filtering can be performed by a bandstop filter (BSF) with multiple stopbands or a bandpass filter (BPF). Reconfigurable filtering requires the components of these filters be tunable. This chapter reports the state of the art in reconfigurable RF filtering, and demonstrates that a tunable BPF approach is inadequate for a BTS front-end.

2.1 Reconfigurable BPF's

The most popular BTS duplexer filter is based on a waveguide filter with cavity resonators. These cavities are coupled together to achieve the desired pole-zero placement to produce a BPF frequency response. Each resonator can have a quality factor (Q) on the order of several 1000's [7]. Multiple resonators are required for a filter to exhibit: large attenuation outside the passband, low loss across the entire passband and adequate return loss performance. In an all-pole BPF, additional poles increase the stopband attenuation by 20 dB per decade per pole. Practical designs incorporate transfer function zeros which significantly increase attenuation near certain frequencies, and forbid a generalization between filter order and stopband performance.

The insertion loss, resonator Q, filter order and bandwidth are all related for a BPF. These relationships are used to demonstrate the limitations of the all-pole BPF, and why it is not

a suitable configuration for tunable RF filtering in a BTS RF front-end. A Chebyshev BPF with 0.1 dB of ripple and N resonators is used as the baseline filter in this section. The insertion loss of this type of filter is close to optimal and the passband magnitude response exhibits regular behaviour [*8*].

The design equations associated with the Chebyshev filter are used to derive the LPF configuration. The BPF configuration is calculated using a lowpass to bandpass transformation. The insertion loss of the LPF filter with a maxima at $\omega=0$ is related to the LPF element quality factor (q) by: $IL=H_L(s)|_{s=1/q}$ [*8*]. $H_L(s)$ is the transfer function of the N^{th} order 0.1 dB ripple Chebyshev LPF with the cut-off frequency normalized to unity. The relationship between q and the insertion loss is shown in Fig. 2.1.1 for various LPF orders (N).

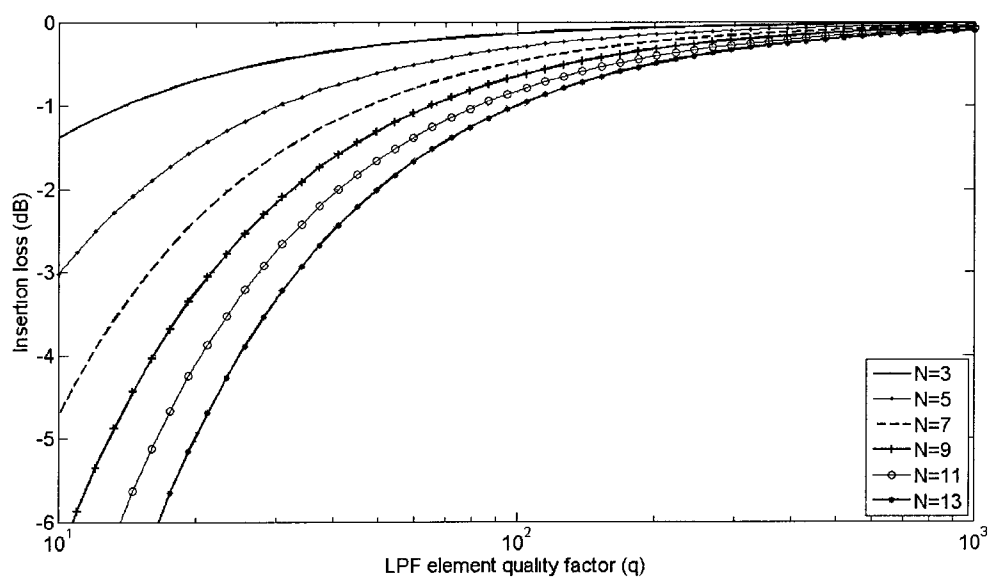

Figure 2.1.1- Insertion loss of 0.1 dB ripple Chebyshev LPF with unity cut-off frequency

The curves in Fig. 2.1.1 can be used to constrain q based on some insertion loss threshold. The relationship between the LPF element q, and the BPF resonator Q is: $Q=q/BW$ [*8*]. The fractional BW is in terms of a percentage. In Fig. 2.1.2 the insertion loss is constrained to 1 dB, and the relationship between Q and BW is plotted for various numbers of resonators (N).

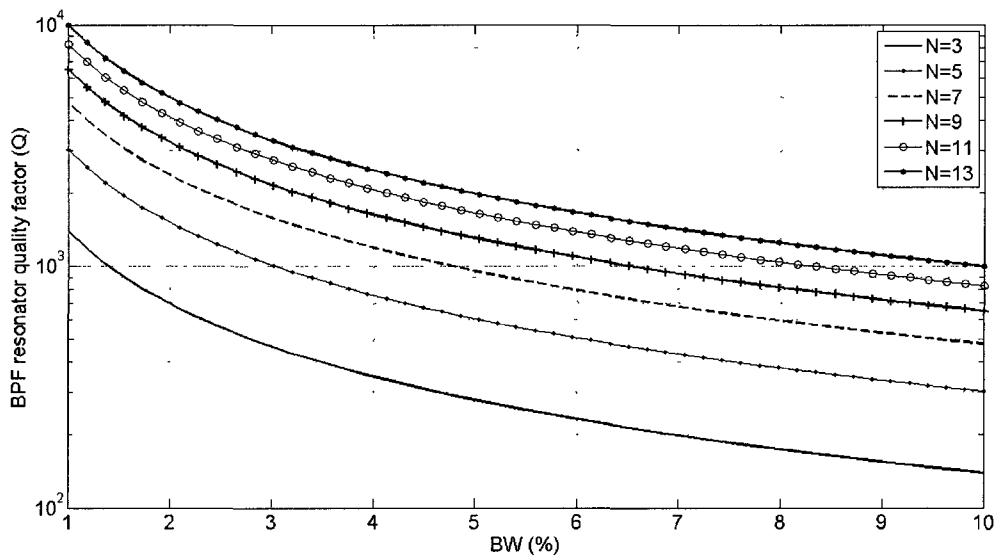

Figure 2.1.2 - Required Q for 1 dB insertion loss of 0.1 dB ripple Chebyshev BPF

The limitations of a BPF approach to RF filtering are illustrated in Fig. 2.1.2. Low insertion loss, high order and small bandwidth can only be performed by filters that consist of resonators with quality factors in the order of 1000-10000. These values are achievable with waveguide based cavity filters, but not with planar or tunable resonators, with YIG technology as the exception. The procedure that was used to generate Fig. 2.1.2 was repeated for a 1 dB ripple Chebyshev BPF with a 2 dB insertion loss threshold. The required Q for this less constrained filter is illustrated in Fig. 2.1.3. The resonator Q for

this less constrained filter is still too high for planar or tunable resonators when the fractional bandwidth is less than 5% and the number of resonators is greater than 5.

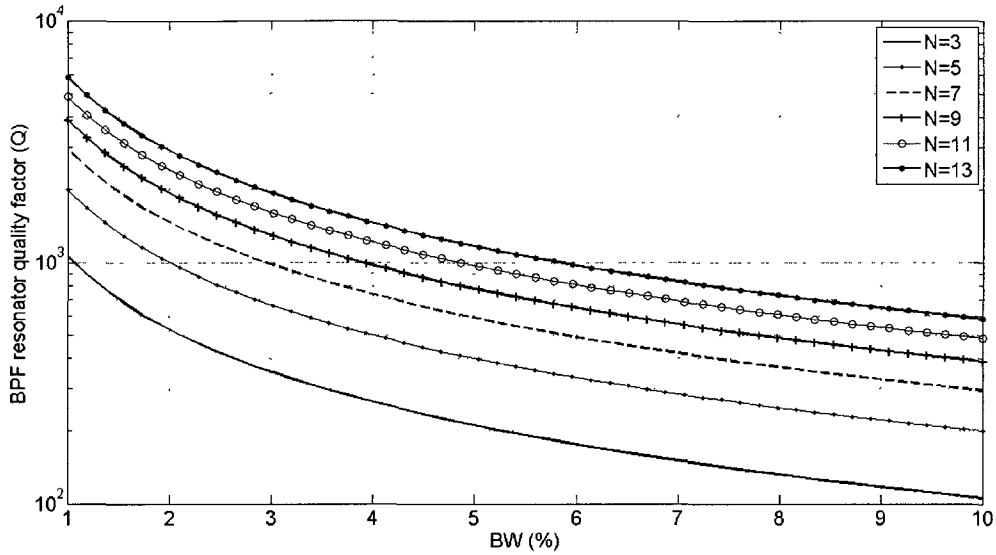

Figure 2.1.3 - Required Q for 2 dB insertion loss with 1 dB ripple Chebyshev BPF

To further demonstrate that tunable BPF's are unsuitable in a BTS RF front-end, the current state of the art in tunable BPF's is reported. Tunable BPF's are categorized as either tuning based or switch based. Tuning based resonators employ technologies such as MEMS capacitors, ferroelectric materials, semiconductor varactors, ferroelectric or ferrite materials. Switch based tunable resonators use on/off switches to dynamically change the interconnections between fixed components, such as the capacitors in a capacitor network. The switches are implemented using technologies such as MEMS switches or PIN diodes.

2.1.1 MEMS

A resonator can be tuned by mechanically changing some parameter - such as the separation between the plates of a capacitor. Also, fixed components can be dynamically

added or removed from a resonator by using metal contact switches. Both of these approaches to tuning involve mechanical tuning. MEMS is a popular technology that is used for mechanically tuning RF components by adjusting an applied DC voltage. Three types of MEMS components are reviewed: metal-contact switches, capacitive switches and varactors.

MEMS metal-contact switches and capacitive switches are useful for dynamically connecting fixed components, like capacitors, to a filter. The quality factor of MEMS metal-contact switches can be negligible compared to the adjacent fixed components [*9*]. Capacitive switches are capacitors whose on/off states give a large capacitance ratio such that when in the off state the component is effectively an open circuit. A large bias voltage is required to switch to the on state, and fortunately this large voltage prevents the components from self-switching when the RF input is large. MEMS capacitive switches can be designed with Q above 100 in the cellular bands [*9*]. Unlike MEMS capacitive switches, MEMS varactors can take on a discrete set or continuous range of capacitance values [*9*].

A tunable two-pole filter with capacitive switches was reported in [*10*]. The tuning range of the centre frequency spanned a discrete set of frequencies between 850 and 1750 MHz. The passband fractional bandwidth was also tuned from 7 to 42%. Both resonators had an off-chip inductor with Q greater than 100. Measurements demonstrated a 1 dB insertion loss across the centre frequency tuning range with the fractional bandwidth fixed at 15%. The estimated filter Q was 60-90 [*9*].

The Q factor of resonators employing tunable MEMS components is fundamentally limited by the Q of the surrounding fixed components, like the filter above. This limitation is most prevalent for planar and discrete components. To realize the potential of high-Q MEMS, they must be incorporated into resonators with fixed components that have high Q. Several papers have reported tunable resonator Q factors greater than 200 using MEMS with non-planar resonators. In these papers the resonator consists of a waveguide cavity with a capacitive post. The resonators are made tunable by using MEMS components to vary the capacitance of the post. The term 'evanescent mode waveguide filter' is used in the literature to imply the cross-section dimensions of the waveguide cavity are significantly smaller than half a wavelength.

A tunable dual cavity evanescent mode waveguide filter was reported in [*11*] with a center frequency range that spanned 2.7 to 4.03 GHz. A flexible metalized substrate was used in place of rigid metal as the top waveguide wall. A piezoelectric actuator deflected the flexible substrate to vary the spacing between the top waveguide wall and the top of the fixed capacitive posts inside the cavity. Varying this spacing varies the capacitance of the post, thereby varying the resonant frequency of the resonator. Measured results of this 2 pole filter showed a fairly consistent 1.3 dB insertion loss with a 2% fractional bandwidth across the tuning range. A second filter was fabricated that measured a 4.46 dB insertion loss for a fractional bandwidth of 0.5%. Only one tuning point was reported for the latter filter in [*11*]. The quality factor of a single resonator was measured, and varied from 360 at 2.3 GHz to 702 at 4.6 GHz.

A different flexible diaphragm evanescent mode waveguide filter with a capacitive post was reported in [*12*]. The diaphragm was deflected using electrostatic actuation. The diaphragm and a second fixed parallel plate had a potential difference applied which caused the diaphragm to move relative to the fixed plate. A single resonator was fabricated. The resonator was tuned from 3.42 to 6.16 GHz, and the corresponding measured quality factor varied between 460 and 530.

A more recent high-Q non-planar tunable filter based on evanescent waveguide resonators was reported in [*13*]. The capacitance of the post was varied in this filter by incorporating RF MEMS switch capacitors between the top of the posts and the top waveguide wall. A 2 pole filter was fabricated with a nominal fractional bandwidth of 0.5%. The center frequency range of the filter was tuned from 4.07 to 5.58 GHz. Over this range the fractional bandwidth increased from 0.44 to 0.74%, and the insertion loss increased from 3.18 to 4.91 dB. The quality factor of the filter was reported to vary from 300 to 500 over the tuning range. The primary limiting factor on Q was the effect of the bias lines running through the resonator that are required to control the MEMS components.

The state of the art in high Q MEMS tunable filters has been reported. The Q of planar filters incorporating MEMS metal contact switches and capacitive switches is limited by the fixed components. Fixed planar or discrete components do not have sufficient Q for a BTS RF front-end. The cavity resonators incorporating MEMS have Q up to 700 at frequencies well above the cellular bands. This optimum Q factor degrades as the frequency tunes away from the optimum. It is expected that the Q factor of cavity

resonators with MEMS will continue to increase in the cellular bands through research efforts. But like fixed cavity resonators, great efforts are required for relatively small improvements.

2.1.2 Varactor Diodes

PN junction diodes exhibit a voltage variable capacitance and resistance when reverse biased. A single varactor has a non-linear C-V curve which limits a varactor's linearity. An anti-series configuration is useful for extending the linearity performance of varactors [*14*]. This configuration is shown in Fig. 2.1.4a. In theory, this configuration can lead to no second and third-order distortion when the C-V exponent of both diodes is exactly 0.5, and the impedance of the centre tap, used for biasing, is orders of magnitude larger than the reactance of the varactors [*14*]. Furthermore, the diodes must be identical, and the sum of the RF and bias signal must not exceed the breakdown voltage or forward bias the varactors. This anti-series configuration has been incorporated into several tunable BPF's [*15,16*].

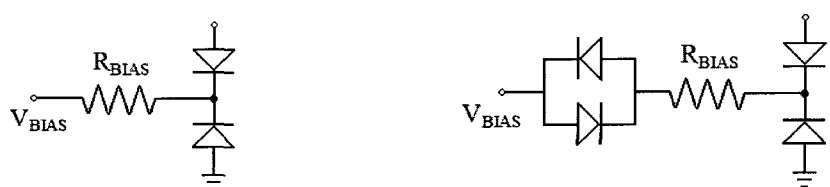

Figure 2.1.4 – Varactor configurations for improved linearity a) Anti-series b) Anti-series with anti-parallel in bias

A two-pole, two-zero tuned BPF with varactor diodes was reported in [*16*]. Two filters were built, one with single diode configurations, and the other with anti-series diode configurations. The BPF was tuned from 1.4 to 2.0 GHz. The IIP3 with a Δf of 1 MHz

ranged from 22 to 41 dBm across the range of bias voltage for the anti-series configuration filter. These IIP3 values were 13-15 dB higher than the single diode configuration filter. The IIP3 improvement reduced to 6 and 7 dB when Δf was less than the filter's 3-dB bandwidth. The diodes used for the two filters were not identical. The resonator Q for the single diode configuration filter varied from 72 to 95 over the tuning range. That for the anti-series diode configuration filter varied from 56 to 125. These resonator Q factors are close enough to expect the IIP3 improvement to be almost completely due to the difference in diode configurations. The finite third order distortion was attributed to the finite impedance of the bias resistor (10kΩ) and the C-V exponent of the diodes being 0.47 and not the ideal 0.5.

A monolithic silicon-on-glass varactor-based BPF was reported in [*15*] that was tuned from 2.5 to 3.5 GHz with a bias from 1 to 13.5 V. One of the constraints for linearity is the bias network impedance must be significantly larger than the varactor reactance. This condition must hold over frequencies containing linear and non-linear components. The reactance of the varactor is $(2\pi Cf)^{-1}$, where C is the varactor capacitance. In a two tone test there is a 2nd order frequency component at $f=f_1-f_2$, which can encounter a very large varactor reactance when the difference between f_1 and f_2 is small. In this scenario the distortion would not cancel between the two varactors, and was the reason the IIP3 dropped in [*16*] when Δf of the 2 tone test was reduced. This limitation was overcome in [*15*] by cascading the bias resistor (R_{BI}) in Fig. 2.1.4a with an anti-parallel varactor configuration, as shown in Fig. 2.1.4b. The filter had one pole and one zero with a pole-zero spacing of 400 MHz. The varactor Q varied from 100 to 600 for 5 to 20 pF at 2

GHz. An IIP3 of 46 dBm was measured for two signals in the stopband with a Δf of 2 MHz. The same IIP3 was also measured for a two-tone input in the passband. This IIP3 was for the lowest DC bias voltage of 1 V. The Q of the resonators was not reported. The maximum variation in capacitance for a diode varactor occurs for low bias voltages, but it cannot support a large AC voltage swing at low bias voltages. If the voltage swing is larger than the bias voltage, then the diode can transition into the forward bias region. In the front-end of a BTS the potential for this undesirable large signal behaviour is problematic. Furthermore, the Q factors from the papers reviewed did not meet the criteria for a BTS front-end.

2.1.3 PIN Diodes

Like MEMS, semiconductor components can be used in tunable filters for variable capacitance, like varactor diodes, or for switching, like PIN diodes. PIN diodes are used for dynamically changing connectivity between fixed components. Unlike MEMS switches, PIN diodes need to be biased with a DC current to provide a short circuit through the diode [*17*]. The DC current increases with linearity demands. The advantage of PIN diodes over MEMS is their compatibility with standard fabrication techniques. Switches limit the number of configurations that a tunable filter can realize. A large number of switches can be used to effectively realize any practical configuration. But PIN diodes are not the optimal technology for a large switching matrix due to their power demands.

2.1.4 Ferroelectric

Ferroelectric varactors are fabricated with ferroelectric materials using thin-film or thick-film processes. The permittivity of ferroelectric materials is dependent upon an applied external electric field. The most popular material for ferroelectric varactors in the literature is barium strontium titanate (BST). Thick-film processes have higher linearity than thin-film, but require a higher tuning voltage – up to several hundred volts [*18*]. Thick-film processes are limited to planar structures such as coplanar waveguide or interdigital capacitor, and the material Q is generally lower than thin-film.

Ferroelectric varactors are commercially available from Paratek and Agile RF. Two tunable BPF's were reported in [*19*], which used tunable thin-film capacitors from Paratek. These varactors had a tuning range of 4.15:1 with a control voltage of 0 to 27 V, and a Q factor over 100 from 100 to 1000 MHz [*19*]. A two-pole BPF was tuned from 300 to 450 MHz with a control voltage of 0 to 22 V. The worst case insertion loss was 1.63 dB. An IIP3 of 40 dBm was measured with two input signals in the passband separated by 2 MHz. A three-pole BPF was also fabricated and measured with the same voltage control range. The three-pole BPF was tuned from 230 to 400 MHz and had a worst case insertion loss of 2.5 dB.

More recently a thin-film BST varactor demonstrated a Q factor from 100 to 350 over the tuning range of 0 to 8 V at 1 GHz [*20*]. This high Q was obtained by using a nano-structured thin film BST. Large area pulsed laser deposition was used for depositing the thin-film. The largest Q factor for BST varactors near 1 GHz is the 100-350 range

reported in [*20*]. This state of the art performance for BST is still too low for a BTS front-end.

2.1.5 Ferrite

Ferrite materials have some properties at microwave frequencies that have put them into the front-end of radios as circulators, isolators, absorption type tunable BPF's and BSF's, and frequency selective power limiters [*21*]. Ferrites exhibit ferrimagnetic resonance in the presence of a saturating DC magnetic field. Essentially, a majority of electrons in a ferrite material tend to align their spin axes when biased by a DC magnetic field of sufficient magnitude. Ferrites have very low conductivity at microwave frequencies; therefore they exhibit low conductor loss.

Ferrite BPF's take advantage of the highly selective nature of a ferrite at the material's resonant frequency. The linewidth of a ferrite can be very narrow such that the unloaded Q factor can be in the range of 2000 to 10 000 [*21*]. YIG is the ferrite with the narrowest linewidth, which is the reason it is popular in ferrite filters requiring high Q.

In ferrite resonators signals are coupled into the fundamental mode, which is called the uniform precession mode. The resonant frequency of this mode depends primarily on the external DC magnetic bias and the geometry. A sphere shape resonator with a high surface polish is required to limit the number of natural resonant modes over the operating frequency range. The sphere must be small to ensure the magnitude of the input RF signal is constant across the YIG, otherwise undesired resonator modes are excited [*22*]. A ferrite sphere resonator with wire loops used as the coupling mechanisms is shown in Fig. 2.1.5. The input and output coupling networks are placed perpendicular to

one another to use the resonator as a bandpass filter. The input signal coupled into the uniform precession mode is subsequently coupled to an output coupling network.

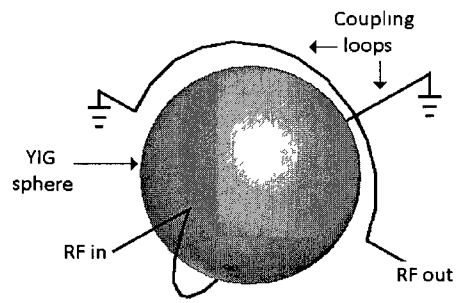

Figure 2.1.5 - Ferrite sphere BPF

The high Q property of ferrites breaks down when used at low frequencies. Small DC bias fields are required for resonance at a low frequency, but if the DC bias field is too small then the ferrite will not be completely magnetized. If not completely magnetized then domain-wall motion causes magnetic losses. The Q factor of a YIG sphere reduces to 0 as the frequency reduces to 1670 MHz [*21*]. The low-frequency limit can be reduced by doping the YIG or using a different shape like a disc. The low-frequency limit for Ga-doped YIG is 1000 MHz [*21*]. Using a disc can reduce this cutoff frequency to almost 0 [*21*]. Doping YIG and using non-spherical geometries can degrade the Q factor due to excitation of higher order modes. Essentially for operation over the 1-2 GHz range the Q factor of the resonator will not be multiple 1000's, so it will not meet the minimum Q required in a BTS front-end. Other drawbacks of ferrites are temperature dependency, IMD, fabrication complexity, and the power required to bias the material with an electromagnet.

The preceding review represents the state of the art in tunable RF filtering with a BPF approach. These approaches have demonstrated large tuning ranges, but are limited to

low order filters due to the relatively low Q-factors of the tunable resonators. These characteristics do not meet the filtering requirements of a BTS RF front-end.

2.2 Reconfigurable Multiple Stopband Filters

BPF's have become the filter configuration of choice for fixed-frequency filtering in the RF front-end of a BTS. But the lack of a suitable high Q and highly linear tunable resonator thwart the possibility of a frequency agile BPF with acceptable performance. A different approach for tunable RF filtering is based on creating multiple stopbands or notches. Multiple stopband or notch filters must adapt to the input RF environment to keep the notches collocated with undesired signals outside the passband. Multiple notch filters are implemented as all-zero filters. The locations of the zeros correspond to the notches in the frequency response. If tunable RF resonators are used to implement the zeros, then resonator Q will limit the maximum filter order as with the BPF configuration. A different all-zero technique is based on feed-forward (FF). FF systems attempt to perfectly cancel part of an RF signal using one or more copies of that signal. In this section the state of the art in RF FF attenuation is presented. RF FF has more heritage in the field of power amplifier (PA) linearization than in RF filtering. FF amplifier linearization is reviewed prior to the state of the art of FF RF filtering.

2.2.1 FF Amplifier Linearization

FF amplifier linearization is a technique that reduces the output distortion of a non-linear PA. A typical FF amplifier linearization system is shown in Fig. 2.2.1a. This system consists of two FF cancellation loops: the signal cancellation loop and the error

cancellation loop. Each loop contains a fixed delay element in one path, an adjustable complex gain in the other path, and a combiner. Successful operation of the signal cancellation loop prevents the components in the error cancellation loop from being pushed into non-linear operation. Successful operation of the error cancellation loop linearizes the PA output signal.

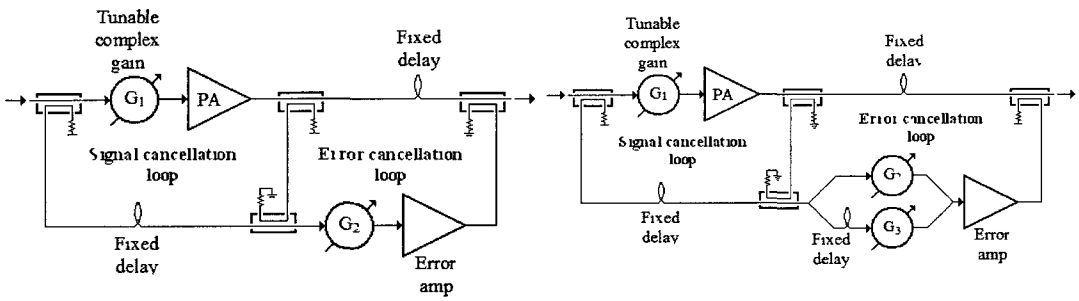

Figure 2.2.1 - FF amplifier linearization a) typical linearizer b) extra path in error cancellation loop

The primary performance metrics of the FF cancellation loops are attenuation level and attenuation bandwidth. The attenuation level is very sensitive to phase and amplitude mismatch between the FF paths. Adaptive algorithms are required to maintain adequate attenuation during environmental changes and component drift [*23*]. The attenuation bandwidth is limited by group delay mismatch between the two FF paths, which can include time delay mismatch and frequency dependent components.

The attenuation bandwidth of the conventional FF amplifier linearizer was increased by adding more FF paths in the cancellation loops in [*24*]. Each path has a different delay element. The outputs of the loops are recombined before being recombined with the output of the PA. A cancellation loop with one additional path is shown in Fig. 2.2.1b. The frequency response of this error cancellation loop has two nulls separated by a

frequency offset. Both of the complex gain elements drift; therefore they need to be tracked by an adaptive algorithm. The adaptive algorithms used in [*24*] decorrelate the signals within each branch before independently correlating the results with the FF output signal.

A technique involving a digital filter and auxiliary transmitter was used to increase the attenuation bandwidth in the signal cancellation loop in [*25*]. This hybrid RF-DSP linearizer is shown in Fig. 2.2.2. The digital FIR filter models the frequency dependent behaviour of the RF components including the PA, and the required group delay to ensure wideband signal attenuation. Conceptually, this is an extension of Cavers and Smith's multiple FF path concept from [*24*]. A digital FIR filter contains many FF paths, each can be adjusted in gain and phase to realize a desired frequency response over the attenuation bandwidth. Typically the transmitter passband is small enough to not require bandwidth enhancement in the signal cancellation loop, but this technique was reviewed due to some similarity to the research presented in this thesis.

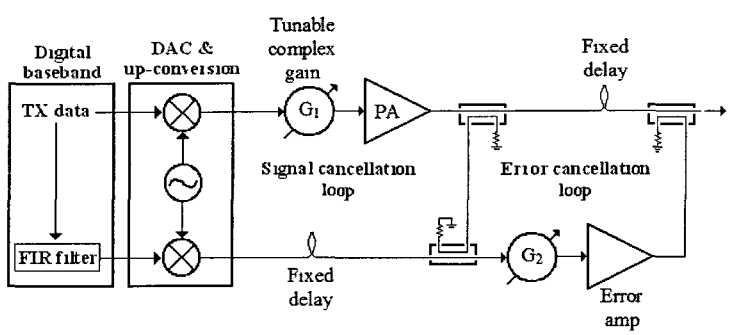

Figure 2.2.2 - DSP path in signal cancellation loop [*25*]

The bandwidth enhancement techniques in [*24*] and [*25*] offer more degrees of freedom to the control algorithm than the standard FF amplifier in Fig. 2.2.1a. These additional

degrees of freedom require more complex adaptive algorithms to maintain adequate attenuation performance. The remainder of this section is a survey of RF FF attenuation used in filtering applications.

2.2.2 Interferer Attenuation in Receiver

An RF discrete-time FF filter for a receiver was reported in [*26*] that taps the signal before the LNA, which is then used to perform signal attenuation at the LNA output. A configuration with one auxiliary FF path is shown in Fig. 2.2.3, which is capable of producing one pair of frequency response nulls that are symmetric about the centre frequency of the passband. In this system, the signal is direct down-converted in the auxiliary path, then passed through HPF's, LPF's, and a complex gain stage. The signal is then up-converted and recombined with the LNA output. The HPF's remove the receiver's passband component and noise generated in the auxiliary path's direct down-conversion mixers. The LPF's attenuate noise outside of the operating BW of the system. Large interferers that traverse the auxiliary path are adjusted in phase and amplitude to result in FF attenuation at the LNA output. Multiple interferers can be cancelled by adding more auxiliary FF paths.

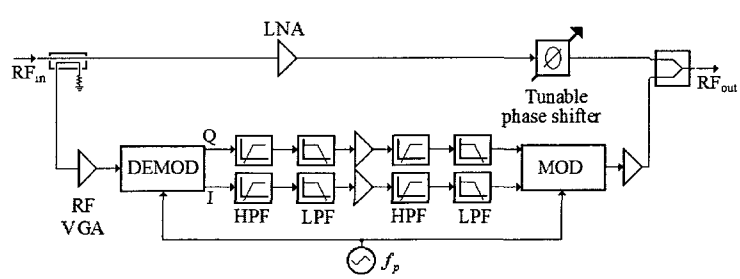

Figure 2.2.3 – Blocker attenuation in receiver prior to mixer [*26*]

A hardware prototype was reported in [*26*] with discrete components and one auxiliary path, like Fig. 2.2.3. A 20 dB coupler was used to extract the input signal for the auxiliary path. A phase shifter was placed in the main path and an RF VGA in the auxiliary path. The HPF's had a corner frequency of 2 MHz, and the LPF's had a corner frequency of 16 MHz. The LPF's were found to improve the output noise floor by 2 dB. One realization of the VGA and phase shifter settings resulted in a 26.2 dB null at 5.9 MHz below the passband centre frequency (880 MHz), and a 23.5 dB null at 4.4 MHz above. The 20 dB attenuation BW was approximately 0.27 and 0.17 MHz for the upper and lower sideband notches, respectively.

A second hardware prototype was built from discrete components that had two auxiliary paths to permit attenuation of two interferers. A phase shifter was placed into each of the auxiliary paths, instead of the main path. One realization of the VGA and phase shifter settings resulted in a 25 dB null at 10.35 MHz above the passband centre frequency (870 MHz), and a 15 dB null at 20 MHz above. Results in the lower sideband were not reported. The 20 dB attenuation BW of the deeper notch was approximately 1 MHz. Simulations were used to demonstrate proper operation of a system with multiple auxiliary paths, but only required one demodulator and modulator. Several copies of the IQ outputs of the demodulator are sent to different baseband paths. These paths are uncorrelated by putting a different non-overlapping BPF in each path. The amplitude and phase adjustment in one frequency sub-band does not impact the frequency response outside this sub-band.

A BTS requires a high order filter; therefore would require multiple auxiliary paths. As the number of auxiliary paths increase, so does the number of analog components which drift with age and environmental changes. A control algorithm was not reported in [*26*], but would need to track this drift. The frequency spacing between sub-band filters should be no more than the expected frequency separation between interferers. This constraint could result in an unmanageable number of sub-band filters. Furthermore, these sub-band filters would need to partially overlap to avoid dead spot frequencies that the FF system cannot attenuate. The FF system in [*26*] is limited to relatively low order, which is useful for attenuation of a few interferers that are adequately separated in frequency.

2.2.3 Feedforward Duplexer Isolation Improvement

RF FF systems have been designed to improve duplexer isolation. Duplexer isolation is most important in the transmitter and receiver passbands.

A FF active attenuation architecture was proposed in [*27*] to augment a SAW duplexer. This system attenuates RX band noise that is generated by the transmitter. The schematic is shown in Fig. 2.2.4. The sampled signal from the PA output is passed though a notch filter centred in the TX band. This notch filter is required to prevent the high power TX band signal from pushing the subsequent components into non-linear operation. After the notch filter, the RX band signal component is amplified, and then adjusted in amplitude and phase before being coupled back into the main path. The amplitude and phase adjustment is tunable to permit the FF attenuation to occur at any channel in the RX band. This adjustment is performed by a tunable phase shifter and a variable attenuator.

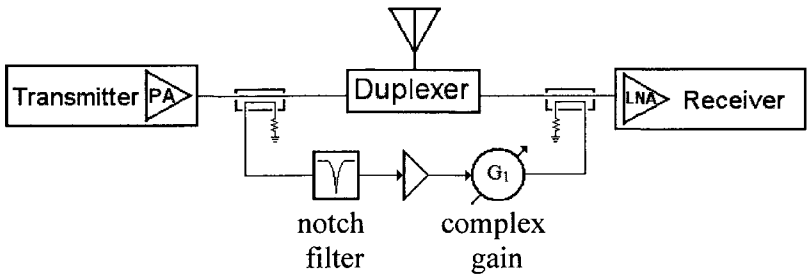

Figure 2.2.4 - FF attenuation to receiver [*27*]

The SAW duplexer contributed to a relatively large group delay mismatch between the two FF paths. In FF systems, the attenuation bandwidth narrows as the group delay mismatch increases [*23*]. Multiple FF paths were proposed in a parallel configuration for attenuation at multiple frequencies, or for a wider attenuation bandwidth.

The system was fabricated with a single FF path in [*27*]. The TX band was 824-849 MHz and the RX band was 869-894 MHz. The SAW duplexer had 40 dB of TX-RX isolation in the RX band. The active cancellation system increased the isolation by more than 20 dB over a 2 MHz channel bandwidth. This performance was reported for each channel in the RX band. A system with dual error paths was also built and tested. The results for two different cases were reported. The first case placed the two frequency response nulls 9 MHz apart, and the resulting improved isolation was 9 dB over 16 MHz. The second case had a null spacing of 4 MHz, whereby the isolation increased by 20 dB over 4.5 MHz. The insertion loss for both the transmit and receive chain was 0.27 dB. The control signals in the gain and phase adjuster were manually tuned.

Tunability of the RX channel within the RX band was demonstrated. The TX band was not tunable, due to the fixed notch filter. The proposed system did not isolate the receiver from the portion of the transmit signal that is reflected back from the antenna. A second

system was proposed that is identical to that in Fig. 2.2.4, except the signal is re-injected into the antenna port of the duplexer to cancel antenna reflections. This system would need the same components as that in Fig. 2.2.4, thereby doubling the total number of components.

A system with two FF cancellation loops was reported in [*28*] to improve duplexer isolation in a BTS. This system is located at the output of the TX chain, and is designed to cancel noise generated by the transmitter in the RX band. A block diagram is shown in Fig. 2.2.5. At the output of the transmitter, the TX band signal is much larger than the RX band noise. The TX band signal must be attenuated prior to passing it through components that operate on the RX band signal; otherwise the components would be pushed into non-linear operation. An additional FF cancellation loop is used to perform this TX band attenuation. The first loop in Fig. 2.2.5 is the signal cancellation loop. Since the loop uses FF attenuation, it can be tuned to permit a frequency agile TX band. The error cancellation loop in Fig. 2.2.5 performs the RX band attenuation. Both cancellation loops have a complex multiplier, which is implemented with a vector modulator cascaded with an amplifier. Since both loops are independently tuned, the FDD separation can vary. A reference signal is injected near the RX channel, which is measured after the single path cancellation. This measurement is used to manually tune the control signals. The band used in this work was at 1900 MHz. Simulation results showed 32 dB of attenuation in the error cancellation loop over a 5 MHz RX channel with an FDD separation of 5 MHz [*28*]. The signal cancellation loop attenuation was greater than 20

dB over the TX channel. The error cancellation loop attenuation corresponds to the improvement in duplexer isolation over the RX channel.

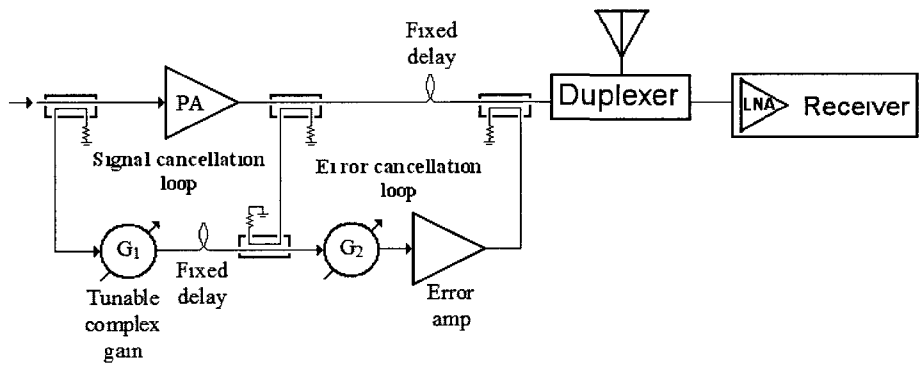

Figure 2.2.5 - Frequency agile BSF schematic [*28*]

The two preceding systems reviewed improve duplexer isolation in the RX band. Inadequate isolation over the TX band will result in the high power signal leaking into the receiver and potentially desensitizing the receiver. In [***29***] an auxiliary transmit chain is used to generate an RF signal that cancels the transmitter leakage in the receiver. This system improves isolation only in the TX band. The attenuation occurs between the duplexer and the LNA, as shown in Fig. 2.2.6. The input to the auxiliary transmitter is the same digital signal that is sent to the main transmitter, except it goes through an additional digital FIR filter. This filter is designed to have a transfer function that models the transmitter chain components including antenna reflections. A wideband training sequence is sent through the transmitter, which is used to calculate the filter coefficients.

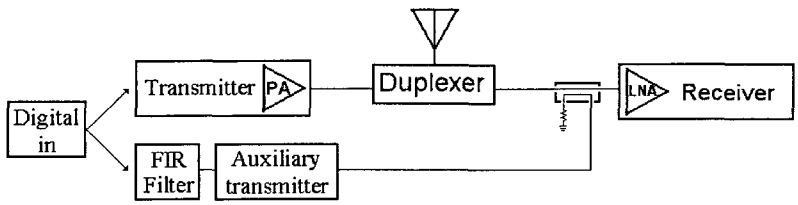

Figure 2.2.6 - Auxiliary transmitter attenuation [*29*]

The test setup included a fixed wideband duplexer that provided 25 dB of isolation between the main transmitter and the receiver. For a TX signal with a 5 MHz bandwidth centered at 1.96 GHz, the additional isolation provided by the auxiliary transmit chain was 35 dB. This result was similar for a narrow band signal centered at 2.02 GHz. This paper demonstrated wideband attenuation without the use of multiple RF cancellation paths. Adaptive control of the FIR was recommended for optimal performance, but was not implemented in the demonstrator.

The cancellation signal is a linear function of the baseband signal. The cancellation signal cannot cancel the IMD or noise generated by any of the transmitter components. Furthermore, any noise or IMD generated by the auxiliary transmitter in the RX band decreases the sensitivity of the receiver. Large attenuation performance is limited to transmitters with low noise and highly linear components.

In [*30*] a FF attenuation system was developed to enhance a fixed duplexer by improving duplexer isolation in both the TX and RX bands. The fixed duplexer used for the measurements provided at least 20 dB of isolation in both bands. The dual feed-forward path topology is shown in Fig. 2.2.7. Vector attenuators were used in the FF paths to adjust the phase and amplitude.

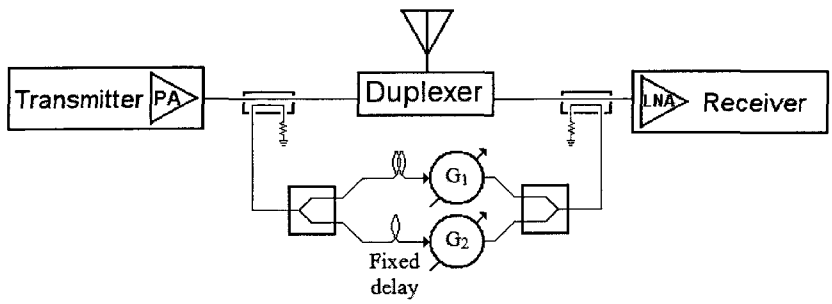

Figure 2.2.7 - Double loop attenuation [*30*]

In [*31*], measured results were published with a TX band centered at 1955 MHz, and a RX band at 2145 MHz for the system in Fig. 2.2.7. The attenuation circuit increased the TX band isolation by 47 dB and the RX band by 38 dB. The attenuation was measured over 5 MHz channel bandwidths.

A pilot tone near the centre of the RX band was coupled into the main path in the transmitter, then down converted and sampled in the receiver and used as an error signal. A control algorithm automatically tuned the signal adjusters in the FF paths.

This system was designed for a mobile terminal, not a BTS. The dynamic range of the BTS transmit signal between the TX band and RX band is much larger than mobile equipment. This would require vector modulators with the same dynamic range in the FF paths, which is infeasible for typical BTS requirements.

The state of the art in RF FF attenuation systems has been reviewed. A large attenuation bandwidth is achievable by minimizing the delay mismatch between the FF paths, or by increasing the number of FF paths. A large number of FF paths is undesirable at RF due to the complexity involved with splitting and combining RF signals, and compensating for all of the frequency dependent components and drift. These limitations have limited the FF systems reviewed to two paths.

2.3 Summary of Literature Review

The Q-factor of tunable resonator based filters results in unacceptable insertion loss in high-order configurations. RF FF filters provide low insertion loss, and high order filtering is possible by increasing the number of FF paths. But, each additional FF path requires an additional signal splitter, tunable complex gain components, and signal combiner. The increase in components, power, space, and control circuitry may not be justifiable for each increase in the filter order for a high-order filter.

3 System Description

In this research a hybrid RF-DSP FF filter has been developed that exhibits low loss in high-order configurations. Increasing the order of this filter only requires an additional tap in a digital FIR filter. The purpose of this chapter is to substantiate the configuration and each module down to the component level. The latter part of this chapter specifies how this filter integrates into an RF front-end.

3.1 System Architecture

The hybrid RF-DSP FF filter consists of two paths in a FF configuration, and a feedback path for filter adaptation. One of the FF paths contains only RF components, and the other has a digital filter along with frequency up and down conversion circuits. The path with the digital filter also contains a tunable BSF. One configuration of the hybrid RF-DSP FF filter is shown in Fig. 3.1.1. Both FF paths start at the coupler denoted C_1, and end at C_2. The feedback path starts at the coupler denoted C_3, and ends at the algorithm processor.

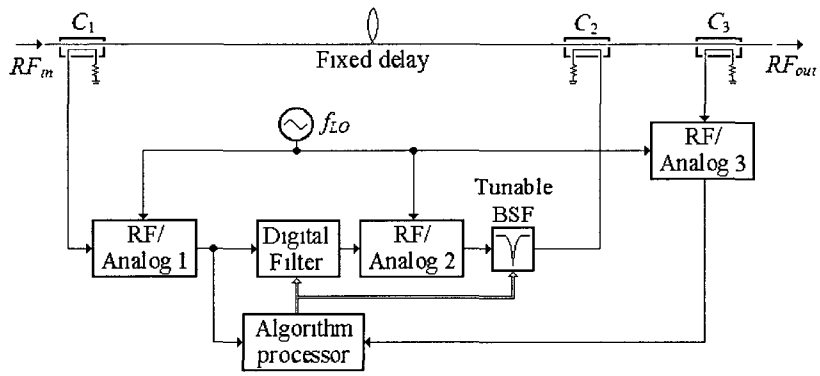

Figure 3.1.1 - System-level schematic of hybrid RF-DSP FF filter

FF configurations provide transfer function zeros. A transfer function zero can be constrained to provide a deep notch in the frequency response. A deep notch is realized when the signals in the two paths have equal magnitudes and a 180° phase difference at C_2. Each additional path in a FF configuration provides an additional transfer function zero. The digital filter in Fig. 3.1.1 is of the FIR type, and each filter tap provides a transfer function zero to the overall hybrid RF-DSP FF filter. The FIR filter can have a large number of taps, which means the entire FF filter can have a high-order all-zero frequency response. Essentially, the digital filter adjusts the magnitude and phase of signals in different bands to ensure adequate attenuation at the RF output of the hybrid RF-DSP FF filter.

The remaining components in the FF path containing the digital filter permit the digital filter to be used in an RF system. The RF/Analog modules in Fig. 3.1.1 that bookend the digital filter provide frequency conversion between RF and baseband. Varying the LO frequency to these modules varies the RF centre frequency of the system. The BSF increases the lower FF path's dynamic range, which is necessary due to the quantization noise inherent to the digital components.

The upper FF path provides the means for the desired passband signal to pass through the filter with minimal loss and distortion. In the simplest configuration the RF path is a transmission line. The insertion loss of this path is that of the thru ports of the three directional couplers and any transmission line loss. The insertion loss of the hybrid RF-DSP FF filter is independent of filter order and can be less than 1 dB by using 10 or 20 dB directional couplers for C_1, C_2 and C_3. Other RF components can be placed in this path, and some examples are shown later in this chapter. The dynamic range, power handling, passband distortion and noise of the hybrid RF-DSP FF filter depend primarily on the performance of the lower FF path.

The block diagram from Fig. 3.1.1 is shown in more detail in Fig. 3.1.2. The algorithm processor has been omitted from this figure for clarity. The lower FF path starts at the coupled port of coupler C_1, and ends at the coupled port of coupler C_2. The 'RF/Analog 1' block consists of an amplifier, demodulator, LPF's, and two ADC's. The ADC's feed a digital FIR filter. The digital filter adjusts the phase and magnitude of signals within the bands where FF attenuation is desired. The output of the digital filter feeds the DAC's in the 'RF/Analog 2' block. This block also consists of LPF's, a modulator and an amplification stage. Before being re-injected back into the upper FF path, the signal is filtered by a tunable BSF. The BSF exhibits a notch over the band where the noise figure of the hybrid RF-DSP FF filter must be minimal and no attenuation is required. The coupled port of the third directional coupler sends the hybrid RF-DSP FF filter's output signal through an amplification and frequency conversion stage. The signal is then digitized and used by adaptive algorithms in the algorithm

processor to adjust the filter taps in the digital filter. The algorithm processor also controls the tunable BSF.

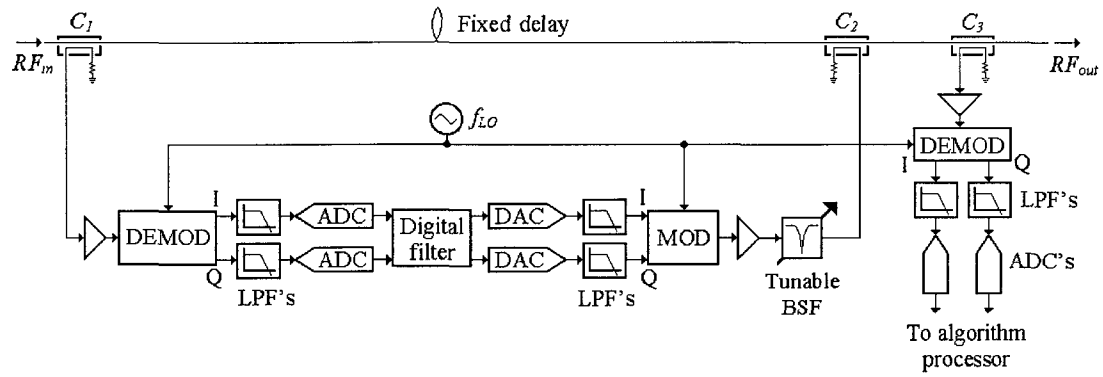

Figure 3.1.2 - Component-level schematic of hybrid RF-DSP FF filter

3.1.1 Digital Filter

DSP enables the use of high-order filtering without the number of RLC components and space demands that are required by RF or analog filtering. DSP is less sensitive to environmental factors, the frequency response is deterministic over a large bandwidth, and filter tap reconfiguration is a straightforward procedure. The major disadvantage of digital filtering in real-time applications is the inherent latency and quantization noise. Excessive latency in the DSP contributes to a large group delay mismatch between the FF paths. FPGA's are the digital technology of choice for minimal latency and rapid prototyping, hence FPGA was chosen for this research.

RF amplifiers, frequency converters, LPF's and data converters are required to reap the benefits of digital filtering in an RF system. An RF filter incorporating a digital filter along with its accompaniment of extra components is shown in Fig. 3.1.3. The demodulator (DEMOD) converts the RF signal to analog baseband, and the adjacent

LPF's are used for anti-aliasing. The ADC's and DAC's are used to interface the analog baseband and digital domains. Two ADC's and two DAC's are required to transfer the signals I and Q components. The LPF's following the DAC's are reconstruction filters, and the proceeding modulator (MOD) converts the analog baseband signal back to RF. A digital clock signal with frequency F_s is distributed to the data converters and FPGA. A digital clock frequency of F_s permits the overall RF system to operate over an RF bandwidth equal to F_s.

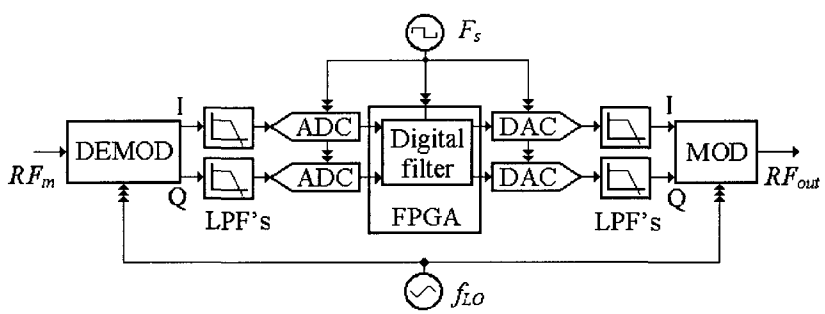

Figure 3.1.3 - DSP filtering configuration

The quantization noise generated by the digital components degrades the output noise floor. This degraded noise floor results in a significant drop in dynamic range between RF_{in} and RF_{out} in Fig. 3.1.3. The output noise floor is most problematic over the filter's passband. In low noise applications a BSF, tuned to the passband, is placed after the modulator to minimize this degradation to the noise floor. Consequently, the passband frequency response of the FF path containing the BSF has a notch over the passband. With this characteristic, ideally none of the signals in the passband at the output of the hybrid RF-DSP FF filter have come from the lower FF path.

The relationships between the design parameters of the digital components and the noise and attenuation performance of the FF filter are derived in Chapter 4. The adaptive algorithms that control the filter taps are derived in Chapter 5.

3.1.2 Direct Up and Down Conversion

The frequency up and down converters that bookend the baseband components enable the hybrid RF-DSP FF filter to be frequency agile. The RF frequency range of these components limits the RF frequency range of the entire filter. The LO source for these components is denoted by f_{LO} in Fig. 3.1.3. The following discussion justifies direct frequency conversion, which inherently has adverse effects near DC.

The LO frequency used by the frequency converters is equal to the centre frequency of the FF filter's passband. This frequency relationship means the centre frequency of the passband is mixed down to DC, where flicker noise and DC impairments are at their worst. After traversing the digital filter, the signals near DC are then up-converted to the centre frequency of the passband again, and then they are significantly attenuated by the BSF. Ideally, the BSF removes any signals from the lower FF path that were caused by flicker noise or DC impairments in the frequency converters. The BSF also removes any passband signal components that appear at its input. The purpose of the lower FF path is to adjust the phase and magnitude of signals that need to be attenuated at the FF filter's output, which does not include the passband. The only contribution to the passband at the hybrid RF-DSP FF filter's output should be from the upper FF path.

3.1.3 Tunable BSF

The BSF attenuates signals and noise within the hybrid RF-DSP FF filter's passband at RF. In doing so, the BSF increases the dynamic range of the lower FF path. In low noise applications the BSF is placed after the modulator or after the adjacent RF gain stage. The centre frequency and bandwidth of the BSF must be reconfigurable to permit the total filter to have a reconfigurable centre frequency and passband bandwidth. The first criterion is met by using an RF notch filter with a tunable centre frequency. A conventional RLC resonator or two-path FF configuration is appropriate to realize the notch. These filters must be tunable; meaning the former must have a varactor, and the latter an adjustable complex gain component. The second criterion of the BSF requires the attenuation bandwidth of the notch response to be tunable. Essentially, signals that are significantly attenuated by the notch, but fall outside the receiver's passband need to be boosted. Typically a tunable pole at RF would do the trick. But, these extra RF components can be avoided by performing the gain compensation in the digital domain – using an IIR filter.

The attenuation bandwidth of the RF notch filter is narrowed by using gain compensation in the digital domain with a technique similar to de-emphasis. The goal of gain compensation is to provide the inverse magnitude response of the RF notch, except within the passband. A digital IIR filter with a pair of complex conjugate poles and a zero at z=1 can be used to create this desired magnitude response.

The notch portion of the BSF cannot be moved to the digital domain. The RF-BSF must notch noise generated by all of the lower FF path components, including the quantization noise in the digital components and flicker noise from the modulator.

Integration of the digital de-emphasis and RF notch (RF-BSF) within the lower FF path is shown in Fig. 3.1.4 and Fig 3.1.5. These two figures show two single order configurations of the RF notch: the RLC resonator, and the two path FF.

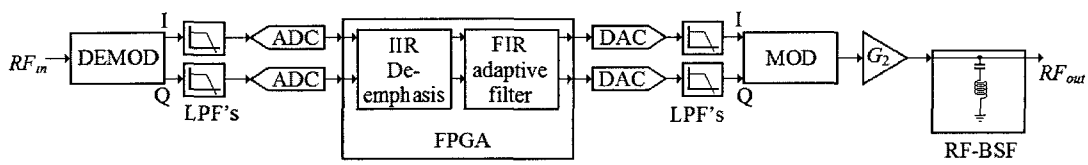

Figure 3.1.4 – RLC BSF combined digital and RF

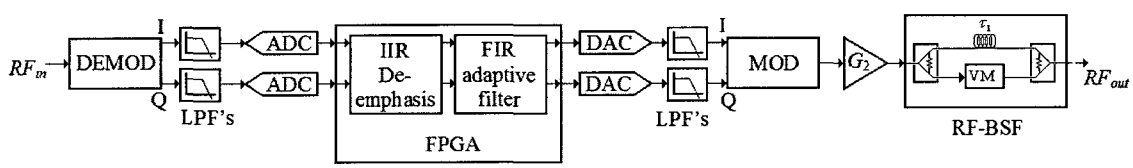

Figure 3.1.5 - FF BSF combined digital and RF

The drawback of the IIR de-emphasis is the additional group delay associated with this filter, due to the additional registers required for implementation. If the de-emphasis filter is omitted to minimize the group delay, then the adaptive FIR filter will need to perform the gain compensation. If the adaptive filter is constrained to perform gain compensation near the passband edges, then it will have fewer degrees of freedom to cause interferer attenuation away from the passband.

3.2 Applications

The hybrid RF-DSP FF filter is fundamentally different than current state of the art tunable RF filters. This section illustrates how this novel filter is integrated into the RF front-end. The filter can be used in the receive chain, in the transmit chain, and across the terminals of a duplexer.

3.2.1 Applications in the Receive Chain

The hybrid RF-DSP FF filter can be used in the receive chain to attenuate large interferers to prevent receiver components from being pushed into non-linear operation. The interference includes self-interference due to leakage of the transmit signal into the receiver as well as external interferers. The digital filter is used to adjust the amplitude and phase of the interference signals to result in FF attenuation at the output of the hybrid RF-DSP FF filter. The high-order filtering characteristic of the filter permits interference to be cancelled relatively close to the passband, which is required for adjacent channel rejection and near-band blockers. High-order filtering in a discrete-time RF filter also permits attenuation of multiple interferers.

The BSF is placed after the frequency up-conversion stage for receive chain applications. The BSF is used to notch any noise generated by the lower FF path in the receiver's passband. The tunable BSF also prevents any filtered passband signal from being re-injected back into the main path; otherwise it might combine with the desired receive signal and result in loss or distortion.

In one configuration the hybrid RF-DSP FF filter is placed between the duplexer and LNA as shown in Fig. 3.2.1. In this embodiment the upper FF path is a delay line. This

configuration is useful to prevent a wideband LNA from being pushed into non-linear operation. The noise figure of the filter is a critical performance metric since the filter is placed before the LNA. The noise figure expression is derived in the following chapter. The interference scenario at the input port is dynamic, so the filter must be adaptive.

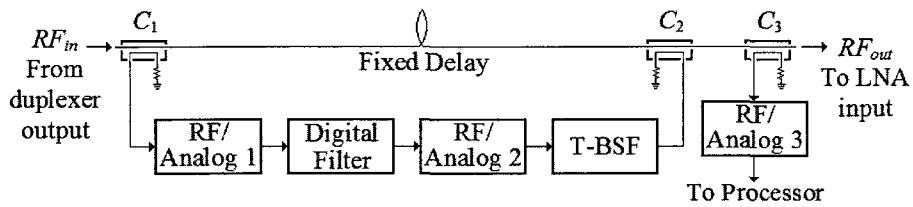

Figure 3.2.1 - Pre-LNA interference filtering application

The PSD's at different points around the filter in Fig. 3.2.1 are shown in Fig. 3.2.2. These PSD graphs are neither from simulation nor measurement. They were generated to illustrate nominal operation in a receiver application. In these inset graphs, 'fp' is the RF centre frequency of the receiver's passband, and 'Fs' is the digital clock frequency. The input PSD to the filter in Fig. 3.2.2 contains five narrowband interferers spread across a span that is 'Fs' MHz wide. This PSD is frequency down-converted to baseband in the bottom FF path before traversing the digital components. The dynamic range of the digital components is generally smaller for data converters than the surrounding analog components. This drop in dynamic range is the cause of the increased noise floor for the PSD at the output of the digital components. This PSD is up-converted to 'fp', then processed by the tunable BSF, which removes the noise within the passband that is generated by the bottom FF path. The PSD's at the outputs of the two FF paths are combined to result in FF attenuation of the five interferers. The output PSD has a relatively large noise floor outside the passband. This noise floor increase is due to the

noise floor from the bottom FF path. The filter's output signal within the passband at 'fp' ideally only consists of any signal at 'fp' at the filter's input. This application is useful when the desired passband signal is low-level.

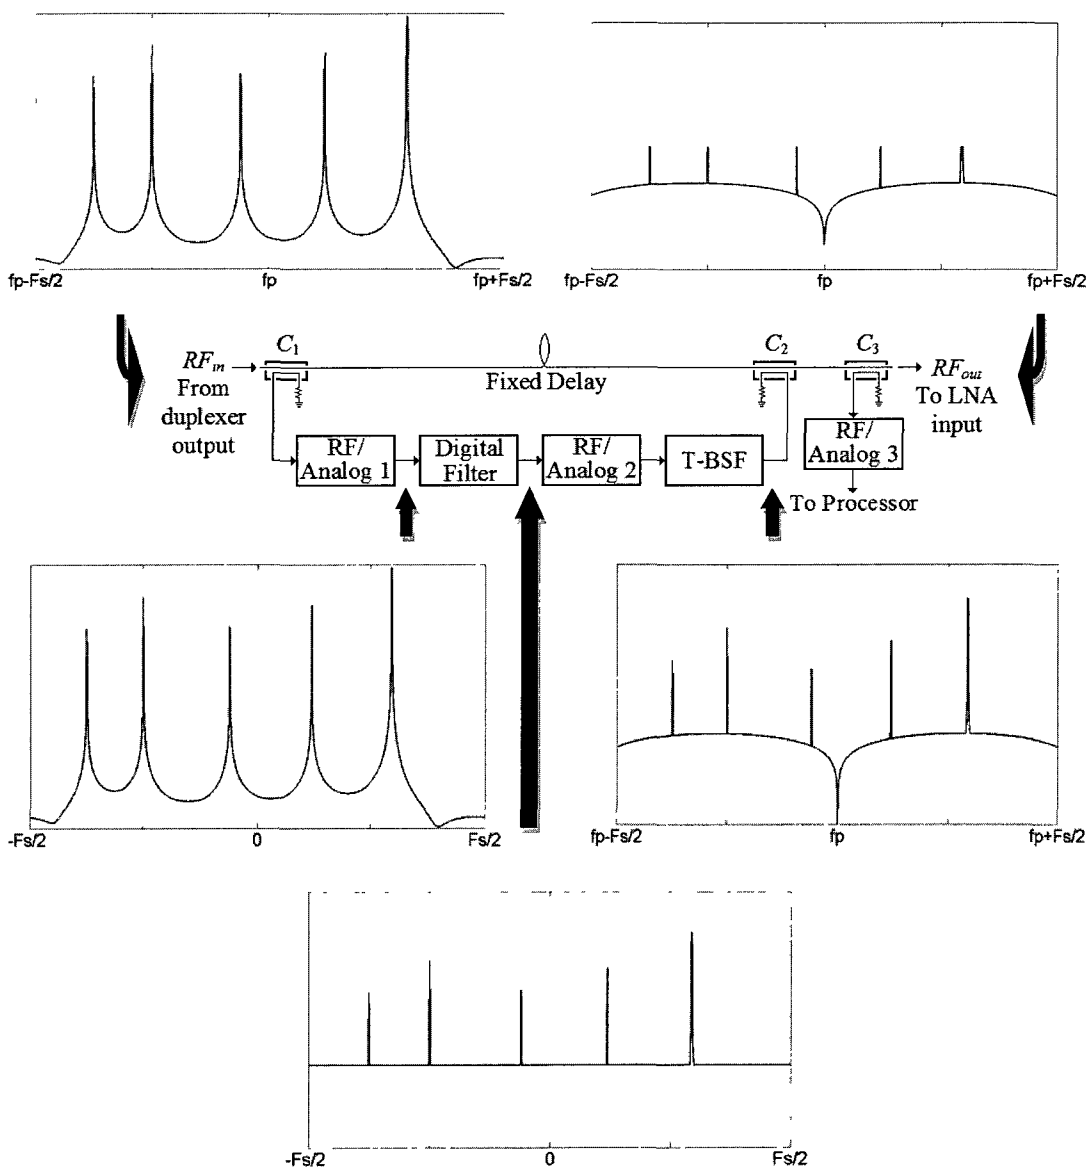

Figure 3.2.2 - PSD at different locations around filter

In another configuration the upper FF path could include the LNA, as shown in Fig. 3.2.3. This system provides additional protection to the mixers in the presence of large interferers, but not the LNA. In this configuration the noise figure of the filter is not as important as Fig. 3.2.1 since the noise is injected after the LNA. The group delay of the LNA will reduce the delay mismatch between the upper FF path and lower FF branch, which improves performance in a FF system.

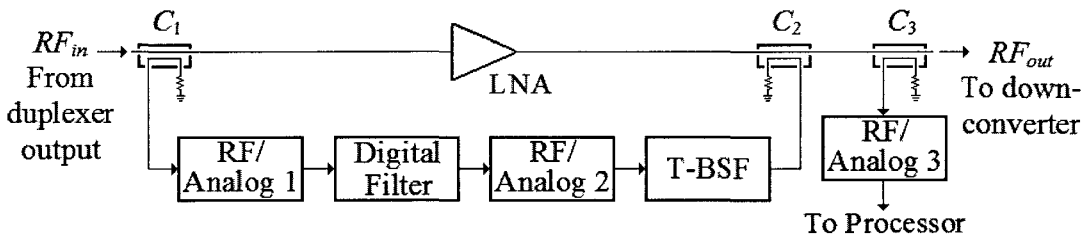

Figure 3.2.3 – System-level schematic of hybrid FF filter

Typically BTS receivers need highly linear down-conversion mixers due to large interferers. Highly linear mixers are implemented using passive components and are generally noisier than active mixers due to a large insertion loss. When the configuration in Fig. 3.2.3 is integrated into a fixed-frequency radio, then the linearity requirements on the down-conversion mixer can be relaxed, and potentially permit lower noise active mixers. This configuration would reduce the system noise figure of the receiver.

3.2.2 Application in the Transmit Chain

Typically the signals that need to be attenuated on the receive side of the radio are much stronger than the signals in the passband. On the transmit side of the radio, typically the signals to be attenuated are much weaker than the signals in the passband. The weak signals that are to be attenuated could be IMD, spurious components and noise generated

by the preceding transmitter components. The attenuation at the output of the PA can be performed anywhere outside the transmit passband, including the radio's receive passband.

The BSF is placed before the down-conversion stage for transmit chain applications and is tuned to remove the large passband signal. If this signal is not adequately attenuated, then the down-conversion and digital components will need to adjust the maximum expected input signal accordingly, which degrades the sensitivity level. Fig. 3.2.4 shows the configuration for the hybrid RF-DSP FF filter placed between the PA and the duplexer. A BSF is also placed in the feedback path to prevent the passband signal from affecting adaptation.

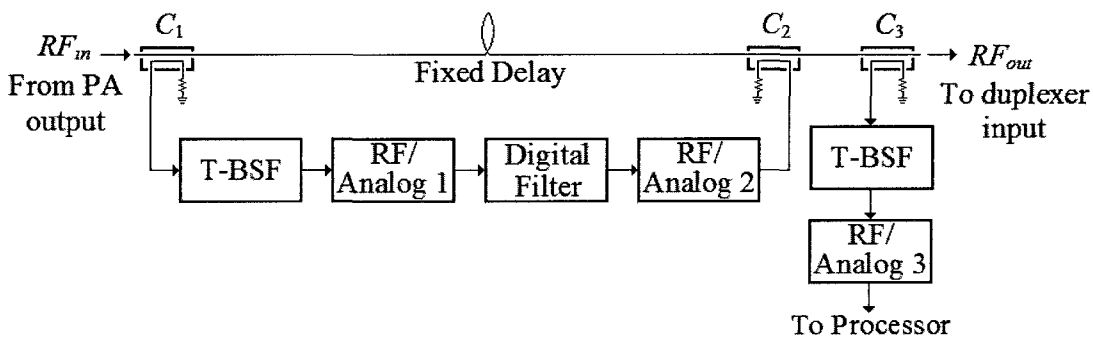

Figure 3.2.4 – Post-PA IMD and spur filtering application

The configuration in Fig. 3.2.4 bears some resemblance to FF amplifier linearization, where the BSF is implemented using the PA input signal. FF amplifier linearization was reviewed in Section 2.2.1. There has been no attempt to use a digital filter for the error cancellation loop in a FF amplifier linearization as shown in Fig. 3.2.4. FF amplifier linearization performs wideband attenuation, whereas the hybrid FF filter is more useful for attenuation in multiple narrower bands. The difference in the attenuation bandwidths

is due to a small delay mismatch between the two FF paths in conventional FF amplifier linearization, versus the larger mismatch for the hybrid RF-DSP FF filter.

The hybrid RF-DSP FF filter would be ideally suited for multiple reconfigurable stopband filtering. These stopbands would correspond to bands with strict emission constraints, including the paired receiver's passband, or other closely located receivers' passbands. Furthermore, the narrowband nature of spurs makes them a well suited candidate for attenuation by the hybrid RF-DSP FF filter.

3.2.3 Application for Duplexer Isolation Improvement

A fixed duplexer is located in the upper FF path in Fig. 3.2.5. In this configuration the hybrid RF-DSP FF filter is used to attenuate the portion of the transmitter output signal that leaks through the duplexer, including reflections from the antenna. In other words the hybrid filter is used to improve the duplexer's transmitter-receiver isolation in the transmit band.

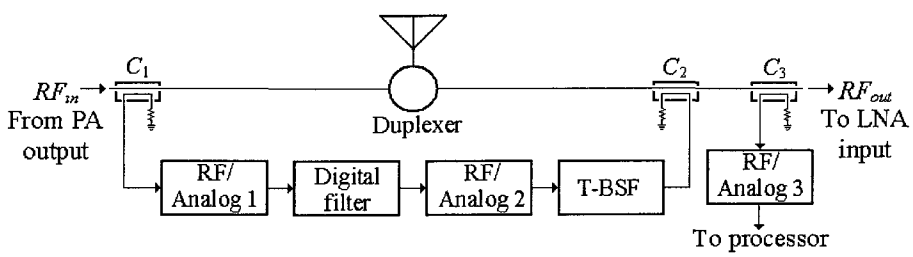

Figure 3.2.5 - Improved isolation application

The different hybrid RF-DSP FF filters described in this section varied with regards to the location of the BSF, and the RF component in the upper FF path. The BSF is located at the beginning of the lower FF path when the input signal has a much stronger passband component than stopband. The BSF is located at the end of the path in the opposite

situation. A BSF is also required in the feedback path if the passband signal is much larger than the stopband signals at the filter's output.

3.3 Summary of System Description

The hybrid RF-DSP FF filter is a frequency agile high-order filter with low insertion loss. The current state of the art of frequency agile filters can achieve low insertion loss but only in low order configurations. The benefits of high-order filtering include larger attenuation closer to the passband, and multiple notches in the frequency response. The latter benefit applies to filters with transfer function zeros.

4 Theoretical Analyses

Expressions are derived in this chapter that predict the performance of the hybrid RF-DSP FF filter. The goal of the hybrid RF-DSP FF filter in receiver applications is to attenuate strong interferes in the presence of a weak passband signal. To this end, attenuation in the interference bands, and the noise factor in the passband are of primary concern.

This chapter contains two analyses. The first analysis, Section 4.1, relates the design variables to the noise added by the hybrid RF-DSP FF filter, and includes the noise factor expression within the passband. The second analysis relates the design variables to the level of interference attenuation achieved through FF cancellation.

Finite-precision effects are introduced by the ADC's, DAC's, and anywhere that rounding or truncating occurs within the digital filter. These finite-precision effects are incorporated into the noise analysis. The digital filter tap coefficients are also represented with finite precision, which alters the digital filter's transfer function from the infinite-precision case. This effect is incorporated into the FF attenuation analysis.

4.1 Noise Analysis

The purpose of this noise analysis is to derive the noise PSD at the output of the FF filter. The output noise PSD is then used to derive the noise factor in the receiver's passband. The noise depends on the noise figures of the RF components. The noise also depends on

quantization, jitter, thermal and non-linear characteristics of the digital components. The non-linear characteristics of the RF components are not included in this analysis. The noise PSD at the output of the FF filter is denoted by $kT_{FF,out}(f)$. Throughout this noise analysis the frequency dependence may not explicitly be shown for many of the terms, but is implied (ie. $kT_{FF,out} = kT_{FF,out}(f)$).

4.1.1 Noise Model

A component level schematic of the FF filter without the monitoring feedback path is shown in Fig. 4.1.1. The RF and analog components preceding the ADC's are denoted *RF1* and those immediately following the DAC's, *RF2*. The transfer functions of the through ports of the couplers are represented by H_{C1t} and H_{C2t}. The transfer functions of the coupled ports are represented by H_{C1c} and H_{C2c}.

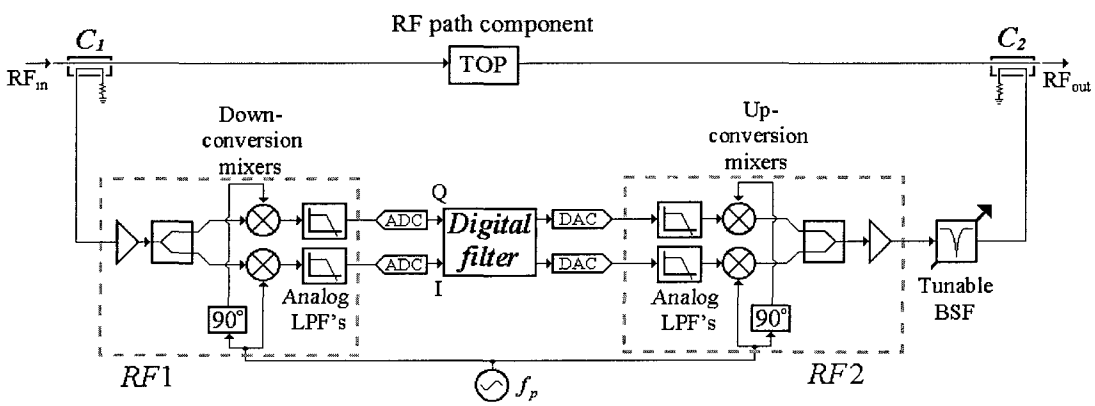

Figure 4.1.1 - Setup for noise analysis

The noise model of the hybrid RF-DSP FF filter is shown in Fig. 4.1.2. The ADC's, digital filter and DAC's are combined into one block with gain and noise factor: H_{DIG}, and F_{DIG}. The terms F_{TOP} and H_{TOP} are the noise factor and frequency response of the top path component. Potential candidates for this component are illustrated in Section 3.2.

The added noise from dissipative losses in the couplers is not explicitly shown in the figure.

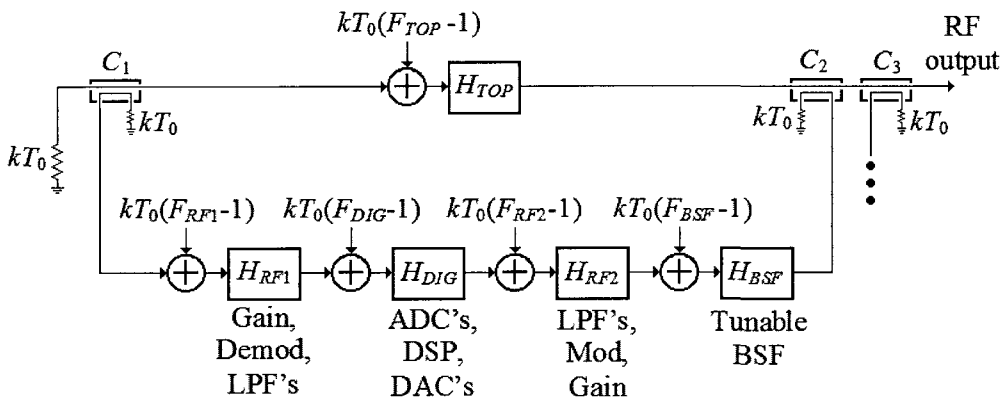

Figure 4.1.2 – Block diagram for noise analysis

The noise from the input noise source and C_1's termination resistor traverse both FF paths before being recombined at C_2. The FF paths are assumed to decorrelate these noise signals due to the large group delay mismatch between the paths. With this assumption the noise model can be simplified to that in Fig. 4.1.3. In this figure the total insertion losses of the directional couplers thru ports are represented by attenuators.

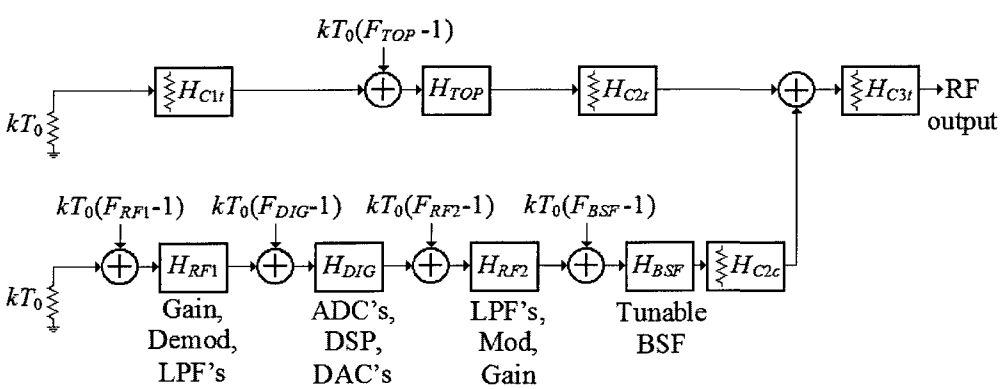

Figure 4.1.3 – Simplified block diagram for noise analysis

The noise power spectral density (PSD) at the output of the FF filter in Fig. 4.1.3 is:

$$
\begin{aligned}
kT_{FF,out} &= kT_0 + kT_0(F_{TOP}-1)\left|H_{TOP}H_{C2t}H_{C3t}\right|^2 \\
&+kT_0F_{RF1}\left|H_{RF1}H_{DIG}H_{RF2}H_{BSF}H_{C2c}H_{C3t}\right|^2 \\
&+kT_0(F_{DIG}-1)\left|H_{DIG}H_{RF2}H_{BSF}H_{C2c}H_{C3t}\right|^2 \\
&+kT_0(F_{RF2}-1)\left|H_{RF2}H_{BSF}H_{C2c}H_{C3t}\right|^2 \\
&+kT_0(F_{BSF}-1)\left|H_{BSF}H_{C2c}H_{C3t}\right|^2
\end{aligned}
\tag{4.1.1}
$$

The next steps in the noise analysis involve specifying each of the noise factor terms in (4.1.1) in terms of circuit-level parameters.

4.1.2 Noise of RF and Analog Components

The RF/analog blocks perform direct-down and direct-up frequency conversion. These modules are shown in Fig. 4.1.4.

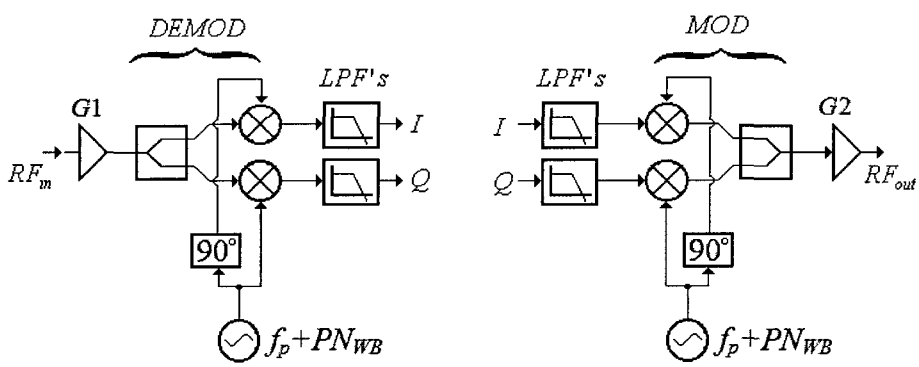

Figure 4.1.4 – Direct-conversion implementation of DEMOD and MOD blocks

The noise factors are calculated for these modules using the Friis equation. When the phase noise from the oscillator is neglected:

$$
F_{RF1} = F_{G1} + \frac{F_{DEMOD}-1}{\left|H_{G1}\right|^2} + \frac{F_{LPF}-1}{\left|H_{G1}H_{DEMOD}\right|^2}
\tag{4.1.2}
$$

$$H_{RF1} = H_{G1} H_{DEMOD} H_{LPF}$$

$$F_{RF2} = F_{LPF} + \frac{F_{MOD} - 1}{\left|H_{LPF}\right|^2} + \frac{F_{G2} - 1}{\left|H_{LPF} H_{MOD}\right|^2} \tag{4.1.3}$$

$$H_{RF2} = H_{LPF} H_{MOD} H_{G2}$$

Reciprocal mixing takes place between the wideband phase noise of the local oscillator and the interferers. The result of this mixing is additional white noise at the output of the modulator or demodulator. The additional noise PSD at the output of the modulator due to M discrete interferers is [*32*]:

$$kT_{pn,MOD,out} = PN_{WB} \left|H_{MOD}\right|^2 \sum_{i=1}^{M} P_{\text{int},MOD,in}(i) \tag{4.1.4}$$

The term $P_{int,MOD,in}(i)$ is the interferer power of the i^{th} interferer at the input to the MOD block. The wideband phase noise of the LO is represented by PN_{WB}. Wideband phase noise is not a function of frequency, and has units of dBc/Hz. To incorporate the additional noise into the noise factor, it must be referred to the input then summed with the existing input-referred noise. These two steps are shown in (4.1.5). The noise factor of the modulator block incorporating the phase noise is denoted by F'_{MOD}. This term can be substituted for F_{MOD} in (4.1.3) to include wideband phase noise in the F_{RF2} expression.

$$kT_{pn,MOD,in} = kT_{pn,MOD,out} \Big/ \left|H_{MOD}\right|^2 \tag{4.1.5}$$

$$F'_{MOD} = F_{MOD} + \frac{PN_{WB}\sum_{i=1}^{M} P_{\text{int},MOD,in}(i)}{kT_0}$$

The additive noise in (4.1.4) is for the MOD block, but also applies to the DEMOD block.

The last RF noise contributor in the lower FF path is the BSF. The gain and noise figure of the RF portion of the BSF are denoted by H_{BSF} and F_{BSF}, respectively. These parameters vary depending on the BSF implementation, two of which are shown in Section 3.1.3.

4.1.3 Noise of Digital Components

The unconventional part of this noise analysis is the noise factor of the digital components, F_{DIG}. This noise factor is a measure of the noise due to the ADC's, DAC's and digital filter (*DSP*). Noise in these components arises from:

- quantization
- truncating or rounding
- scaling values without increasing the number of bits proportionally

Quantization occurs in the ADC's. Truncating or rounding is performed after arithmetic operations. Without truncation each subsequent arithmetic result would have an increasing number of bits, and eventually more bits than the DAC can use. Scaling is required to prevent digital values from overflowing after an arithmetic operation. Scaling involves shifting the decimal location in a fixed-point system.

The magnitude response of the cascaded digital components is

$$|H_{DIG}| = |H_{ADC} H_{DSP} H_{DAC}| \tag{4.1.6}$$

The magnitude responses of the ADC's and DAC's are:

$$|H_{ADC}|^2 = \frac{\frac{1}{2}|v_{ADC,fs}|^2}{P_{ADC,fs}}$$
$$|H_{DAC}|^2 = \frac{P_{DAC,fs}}{\frac{1}{2}|v_{DAC,fs}|^2} \tag{4.1.7}$$

In this equation $P_{ADC,fs}$ is the input analog power to an ADC that results in a full-scale digital output equal to $\pm v_{ADC,fs}$. The notation is similar for the DAC's. If the peak power of an analog signal is larger than $2P_{ADC,fs}$ at the ADC input, then the digital output will saturate. Similarly, if the magnitude of a digital signal is larger than $v_{DAC,fs}$ at the DAC input, then the analog output will exhibit some type of non-linear behaviour. The units of $|H_{ADC}|$ are [digital word/volt], and those of $|H_{DAC}|$ are [volt/digital word].

The noise factor definition requires the added noise be referenced to the input of the cascade of digital components; therefore the noise is referenced to the ADC input. The input-referred noise is derived separately for the ADC's, digital filter and DAC's. The input-referred noise PSD of the ADC's is:

$$kT_{e,ADC} = \frac{P_{ADC,fs}}{SNR_{ADC,fs} B} \tag{4.1.8}$$

In (4.1.8) B is half of the digital clock frequency, and $SNR_{ADC,fs}$ is the full-scale SNR of the ADC's. This SNR value is available from component data sheets. The $SNR_{ADC,fs}$

specification does not include jitter from the clock. The theoretical limit on SNR due to jitter is based on a full-scale sinusoid at frequency f_{in} at the ADC input:

$$SNR_{ADC,jitter} = \frac{1}{(2\pi f_{in}\sigma)^2} \tag{4.1.9}$$

The jitter (σ) corresponds to the jitter of the clock cascaded with other components that exacerbate the jitter. The jitter degrades the SNR of the ADC's. To account for the jitter, the $SNR_{ADC,fs}$ term in (4.1.8) is replaced by SNR_{ADC} from (4.1.10). The result of this substitution is the input-referred noise PSD of the ADC's.

$$\frac{1}{SNR_{ADC}} = \frac{1}{SNR_{ADC,fs}} + \frac{1}{SNR_{ADC,jitter}} \tag{4.1.10}$$

A similar analysis is performed for the DAC's as for the ADC's. The F_{DIG} term requires the noise added by the DAC's be referenced to the input of the ADC's. This involves inverting the output referred noise PSD by the transfer functions of the DAC, DSP and ADC:

$$kT_{e,DAC} = \frac{P_{ADC,fs}\alpha_{DAC}^2}{SNR_{DAC,fs}B|H_{DSP}|^2} \tag{4.1.11}$$

The scaling between the ADC's and DAC's is characterized by:

$$\alpha_{DAC} = \left|\frac{v_{DAC,fs}}{v_{ADC,fs}}\right| \tag{4.1.12}$$

Equation (4.1.11) does not include jitter. Incorporating jitter involves replacing $SNR_{DAC,fs}$ in (4.1.11) with SNR_{DAC} from (4.1.13). The definition of $SNR_{DAC,jitter}$ is the same as the ADC, which is given by (4.1.9).

$$\frac{1}{SNR_{DAC}} = \frac{1}{SNR_{DAC,fs}} + \frac{1}{SNR_{DAC,jitter}} \tag{4.1.13}$$

Quantization noise in the digital filter is added at each location where a loss of precision occurs. A loss of precision is caused by truncating, rounding or scaling without increasing the number of bits accordingly. At each of these locations a quantization noise source is added to the noise analysis model. Each quantization noise source within the filter is assumed to be uniformly distributed between $-Q/2$ and $Q/2$, where Q is the quantization step-size. A real signed number represented with b bits has a step-size:

$$Q = \frac{v_{fs}}{2^{b-1}} \tag{4.1.14}$$

In (4.1.14), v_{fs} is the full-scale digital value. The noise is modeled as white within the Nyquist bandwidth with variance $Q^2/12$ [*33*]. In a digital filter there are multiple sources of quantization noise, but each source is assumed to be uncorrelated with the others, and with the input signal. This noise model has proven to be adequate for wide bandwidth signals with a dynamic range that spans many quantization levels [*34*].

A direct-form FIR filter with two complex tap coefficients and complex input data is shown in Fig. 4.1.5. The real and imaginary products of the complex multipliers are sent to parallel adders. The truncated outputs of the adders are sent to the DAC's. The noise sources at the outputs of the parallel adders in Fig. 4.1.5 correspond to the quantization

noise of the DAC's. A complex multiplier block is shown in detail in Fig. 4.1.6. In this figure a noise source is placed after each multiplier, and after the last adder stage. The quantization step-size at the output of each multiplier is Q_{MULT}, and that of each adder stage within the multiplier is Q_{ADD}.

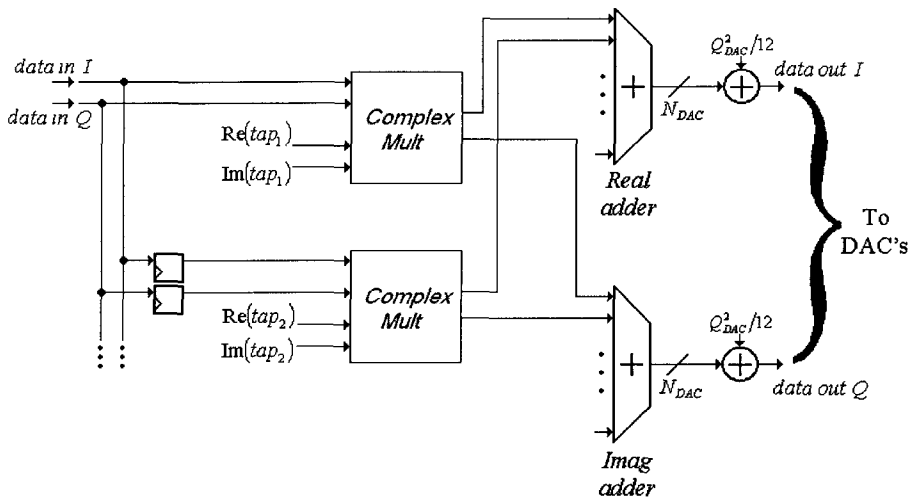

Figure 4.1.5 – Direct-form FIR filter with 2 complex taps showing quantization noise

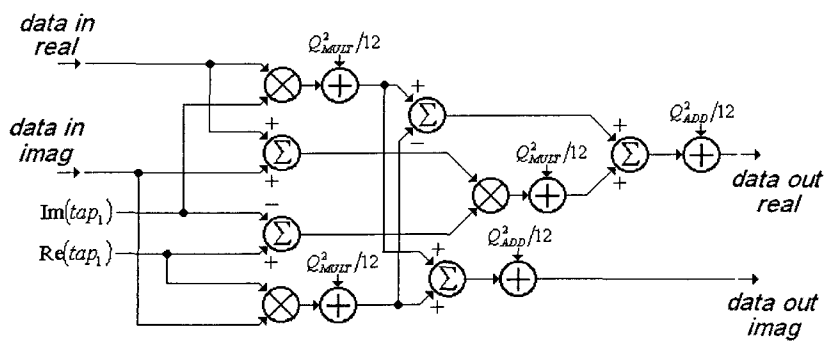

Figure 4.1.6 - Complex multiplier implementation with three multipliers

The quantization noise sources in Fig. 4.1.5 and 4.1.6 are accurate models only if a loss of precision occurs at the location of each noise source. The conditions for a loss of precision at the output of the multipliers and adders are, respectively:

$$Q_{MULT} > v_{ADC,fs} v_{TAPS,fs} / 2^{b_{ADC}+b_{TAPS}-2} \tag{4.1.15}$$

$$Q_{ADD} > Q_{MULT} \tag{4.1.16}$$

To simplify this analysis, assume a loss of precision occurs after the multipliers in Fig. 4.1.6, but not after the adders. Since no loss of precision occurs at the output of the adders, then the Q_{ADD} quantization noise sources in Fig. 4.1.6 are removed. There are three remaining noise sources in each complex multiplier. All three sources contribute noise to the real output, and two sources contribute noise to the imaginary output. The output noise of each complex multiplier is:

$$n_{out} = (n_1 - n_2) + n_3 + j(n_1 + n_2) \tag{4.1.17}$$

The noise sources n_1, n_2 and n_3 can be matched to the three sources at the outputs of the three multipliers in Fig. 4.1.6. The noise power at the output of each complex multiplier corresponds to the variance of n_{out}. The desired signal is not included in this variance estimate since it is assumed uncorrelated with the noise.

$$E\left[|n_{out}|^2\right] = E\{[(n_1 - n_2) + n_3 + j(n_1 + n_2)][(n_1 - n_2) + n_3 - j(n_1 + n_2)]\} \tag{4.1.18}$$

All three noise sources are assumed to be uncorrelated, so any cross-correlation terms resulting from the expansion of (4.1.18) are zero. The remaining terms in the variance expression are:

$$E\left[\left|n_{out}\right|^2\right] = E\left\{n_1^2 + n_2^2 + n_3^2 + n_1^2 + n_2^2\right\} = \frac{5Q_{MULT}^2}{12} \tag{4.1.19}$$

The resulting variance of each multiplier's complex output is $5Q_{MULT}^2/12$. There are N_{TAPS} filter taps, so the total noise entering the parallel adders in Fig. 4.1.5 is $5N_{TAPS}\ Q_{MULT}^2/12$. This is the variance of a complex random variable, so the digital filter's output-referred added noise PSD is spread over $2B$:

$$kT_{DSP,out} = \frac{5Q_{MULT}^2}{24B} N_{TAPS} \tag{4.1.20}$$

The Q at the output of each multiplier is $v_{MULT,fs}/2^{b_{MULT}-1}$. The scaling between the ADC's and the multiplier output's is characterized by:

$$\alpha_{MULT} = \left|\frac{v_{MULT,fs}}{v_{ADC,fs}}\right| \tag{4.1.21}$$

The F_{DIG} definition requires the noise PSD to be referred to the ADC input. This involves inverting the noise PSD by the transfer functions of the digital filter and ADC. Performing this inversion by using (4.1.7) results in:

$$kT_{e,DSP} = \frac{5P_{ADC,fs}\alpha_{MULT}^2 N_{TAPS}}{3B\left|H_{DSP}\right|^2 2^{2b_{MULT}}} \tag{4.1.22}$$

The $|H_{DSP}|$ term in (4.1.22) includes the magnitude response of the digital BSF compensation stage and the main digital filter. The digital BSF compensation stage is introduced in Section 3.1.3. One purpose of this stage is to provide a pair of complex conjugate poles. This pole configuration is accomplished using two feedback paths, each with two real multipliers. The other purpose of this stage is to provide a perfect notch at

f=0, which is accomplished using one FF path. No multiplier is needed in the FF path since a perfect notch at f=0 is obtained when the delayed path is multiplied by -1. The pole configuration and noise sources are shown in Fig. 4.1.7 and 4.1.8.

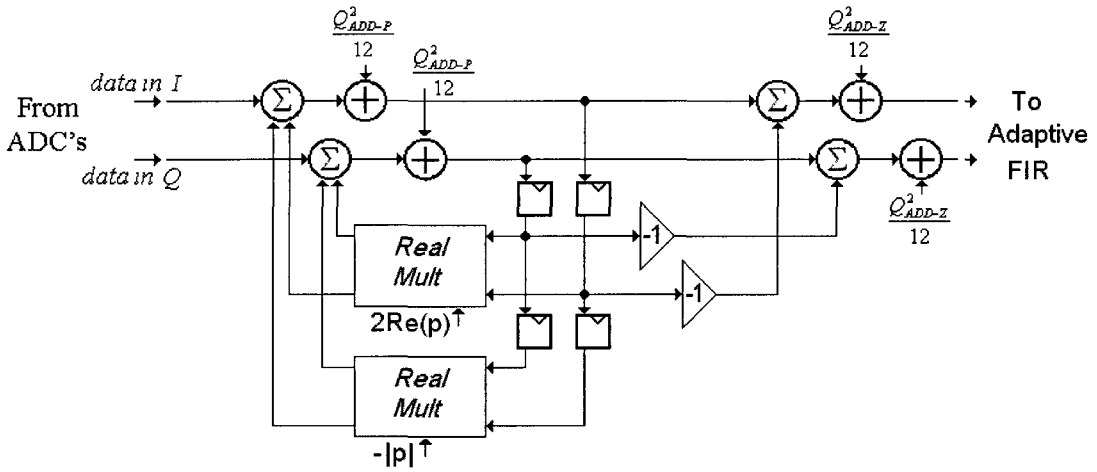

Figure 4.1.7 - Digital BSF compensation with noise sources

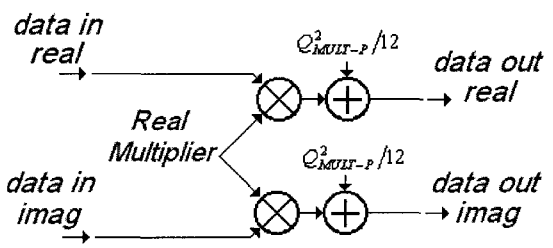

Figure 4.1.8 - Details of 'Real Mult' in Fig. 4.1.7

The output of the IIR filter in Fig. 4.1.7 has a larger dynamic range than its input signal. The IIR filter's maximum output voltage doubles for every 6 dB increase in gain compensation, thereby requiring an extra bit at the output. The scaling at the output of the digital portion of the BSF is represented by:

$$\alpha_{ADD-Z} = \frac{v_{ADD-Z,fs}}{v_{ADC,fs}} \tag{4.1.23}$$

Similarly, the scaling at the output of the multipliers in the feedback section of Fig. 4.1.7 is:

$$\alpha_{MULT-P} = \frac{v_{MULT-P,fs}}{v_{ADC,fs}} \tag{4.1.24}$$

The model is simplified by neglecting the noise sources denoted $Q_{ADD\text{-}P}$. Full precision at the output of the pole's summer can be maintained by increasing the number of bits accordingly.

The quantization noise due to rounding after the multipliers in Fig. 4.1.8 is processed by the feedback loop [*35*]. The noise variance at the output of the digital BSF module due to these multipliers is $Q^2_{MULT-P}|H_{POLES}|^2/12$. The noise PSD generated by the BSF module, denoted 'POLES', referred to the output is:

$$kT_{POLES,out} = \frac{Q^2_{ADD-Z}}{12B} + \frac{Q^2_{MULT-P}}{12B}\left|H_{POLES}\right|^2 \tag{4.1.25}$$

Referring this noise PSD to the input of the ADC requires inversion by the digital BSF and ADC transfer functions:

$$kT_{e,POLES} = \frac{2P_{ADC,fs}}{3B}\left(\frac{\alpha^2_{ADD-Z}}{2^{2b_{ADD-Z}}\left|H_{POLES}\right|^2} + \frac{\alpha^2_{MULT-P}}{2^{2b_{MULT-P}}}\right) \tag{4.1.26}$$

The term $|H_{POLES}|$ is the magnitude response of the digital BSF module, and is very small in the passband due to the notch. These characteristics mean the first term in (4.1.26) dominates the expression:

$$kT_{e,POLES} \approx \frac{2P_{ADC,fs}}{3B} \frac{\alpha_{ADD-Z}^2}{2^{2b_{ADD-Z}} |H_{POLES}|^2} \tag{4.1.27}$$

The effective input-referred noise PSD of the ADC's (4.1.8), DAC's (4.1.11) and digital filter (4.1.22) & (4.1.27) have been derived. These terms are combined to get an expression for the total digital noise factor:

$$F_{DIG} = 1 + \frac{kT_{e,ADC} + kT_{e,DAC} + kT_{e,DSP} + kT_{e,POLES}}{kT_0}$$

$$F_{DIG} = 1 + \frac{P_{ADC,fs}}{kT_0 B} \left(\frac{1}{SNR_{ADC}} + \frac{\alpha_{DAC}^2}{SNR_{DAC} |H_{DSP}|^2} + \frac{5N_{TAPS}\alpha_{MULT}^2}{3|H_{DSP}|^2 2^{2b_{MULT}}} + \frac{2\alpha_{ADD-Z}^2}{3|H_{POLES}|^2 2^{2b_{ADD-Z}}} \right) \tag{4.1.28}$$

Equation (4.1.28) is the derived noise factor for the digital components. This noise factor corresponds to F_{DIG} in Fig. 4.1.2.

4.1.4 Noise Expression Discussion

The goal of the above analyses is to derive the total noise PSD at the hybrid RF-DSP FF filter output. The total output noise PSD is given in (4.1.1), and the expressions for F_{RF1}, F_{DIG}, F_{RF2} in (4.1.2), (4.1.28), and (4.1.3), respectively. In this section the output noise floor expressions are manipulated and simplified with assumptions to answer the following two questions:

1. *Is a sub 1 dB noise figure feasible?*
2. *Under what conditions does the output noise floor degrade attenuation performance?*

The following section reviews the simulations used to verify the noise expression.

1. Is a sub 1 dB noise figure feasible?

A small noise figure is necessary for a receiver application, especially if the entire filter is placed before the receiver's LNA. In this application the RF FF path is a low-loss transmission line. The noise factor of the hybrid RF-DSP FF filter is a function of the output noise PSD from (4.1.1), and the magnitude response of the filter, $|H_{FF}|$:

$$\begin{aligned} F_{FF}(f) &= \frac{kT_{FF,out}}{kT_0|H_{FF}|^2} \\ &= \frac{1}{|H_{FF}|^2} + (F_{TOP}-1)\left|\frac{H_{TOP}H_{C2t}H_{C3t}}{H_{FF}}\right|^2 \\ &+ F_{RF1}\left|\frac{H_{RF1}H_{DIG}H_{RF2}H_{BSF}H_{C2c}H_{C3t}}{H_{FF}}\right|^2 \\ &+ (F_{DIG}-1)\left|\frac{H_{DIG}H_{RF2}H_{BSF}H_{C2c}H_{C3t}}{H_{FF}}\right|^2 \\ &+ (F_{RF2}-1)\left|\frac{H_{RF2}H_{BSF}H_{C2c}H_{C3t}}{H_{FF}}\right|^2 + (F_{BSF}-1)\left|\frac{H_{BSF}H_{C2c}H_{C3t}}{H_{FF}}\right|^2 \end{aligned} \tag{4.1.29}$$

The following assumptions are made to simplify this noise factor expression:

- Top path contains a passive component: $F_{TOP}=|H_{TOP}|^{-2}$
- The RF BSF is RLC based (ie. passive): $F_{BSF}=|H_{BSF}|^{-2}$
- The RF FF path dominates the transfer function of the total filter in the passband $|H_{FF}(f_p)|\approx|H_{C1t}H_{TOP}(f_p)\,H_{C2t}\,H_{C3t}\,|$

Under these assumptions the noise factor expression from (4.1.29) reduces to (4.1.30) in the passband.

$$F_{FF}(f_p)=\frac{1}{\left|H_{C1t}H_{TOP}H_{C2t}H_{C3t}\right|^2}+\left|\frac{H_{C2c}}{H_{C1t}H_{TOP}H_{C2t}}\right|^2+\left(\frac{1}{\left|H_{TOP}\right|^2}-1\right)\frac{1}{\left|H_{C1t}\right|^2}$$
$$+\left|\frac{H_{BSF}\,H_{C2c}}{H_{C1t}\,H_{TOP}\,H_{C2t}}\right|^2\left\{F_{RF1}\left|H_{RF1}H_{DIG}H_{RF2}\right|^2\right. \qquad (4.1.30)$$
$$\left.+(F_{DIG}-1)\left|H_{DIG}H_{RF2}\right|^2+(F_{RF2}-1)\left|H_{RF2}\right|^2-1\right\}$$

The thru port magnitude responses of the couplers account for power transferred to the coupled port and power dissipated within the coupler. Wideband directional couplers from e-MECA exhibit a maximum of 0.1 dB dissipative loss with 10, 20 and 30 dB coupling factors [*36*]. A dual-directional coupler provides a low loss combined solution for $C1$ and $C2$. One such coupler is available from e-MECA with 0.1 dB of dissipative loss [*37*], which is equivalent to $C1$ and $C2$ each exhibiting 0.05 dB of dissipative loss, and $|H_{TOP}|=1$. These insertion loss values are used to replace the thru port magnitude responses of the couplers with:

$$\begin{aligned}\left|H_{C1t}\right|^2&=10^{-0.05/10}\left(1-\left|H_{C1c}\right|^2\right)\\ \left|H_{C2t}\right|^2&=10^{-0.05/10}\left(1-\left|H_{C2c}\right|^2\right)\\ \left|H_{C3t}\right|^2&=10^{-0.1/10}\left(1-\left|H_{C3c}\right|^2\right)\end{aligned} \qquad (4.1.31)$$

The sum of the first two terms in (4.1.30) is shown in Fig. 4.1.9 for various directional coupler configurations using the relationships in (4.1.31). Since a dual-directional coupler configuration is assumed for the implementation of $C1$ and $C2$, their coupling factors are equal: $|H_{C1c}|=|H_{C2c}|$. The horizontal axis is the noise factor contribution for the indicated coupler configuration. As an example, the configuration consisting of $C1$, $C2$ and $C3$

having coupling factors of 20, 20 and 10 dB, respectively, is labelled 20,20,10, and has a noise factor contribution equal to 1.198.

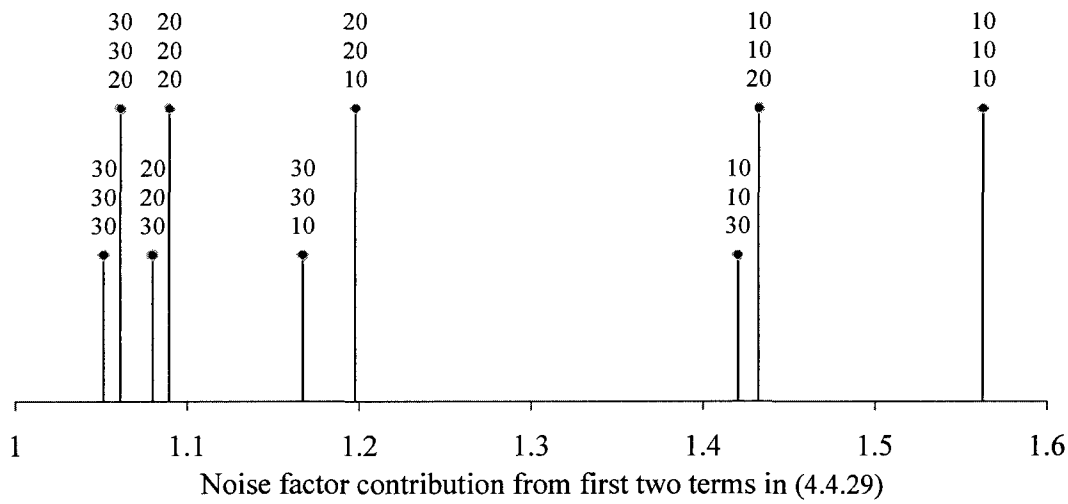

Figure 4.1.9 – Noise factor contribution from first two terms in (4.1.30) for various couplers

A sub 1 dB noise figure requires: $F_{FF}(f_p)<1.26$, which means the contribution from just the first two terms in (4.1.30) should be much less than this value. There are several four coupler configurations in Fig. 4.1.9 with a noise factor contribution between 1.05 and 1.09. Using one of these four configurations limits the remainder of the terms in (4.1.30) to between 0.17 and 0.21 to achieve a sub 1 dB noise figure. The remainder of the terms are manipulated to determine what conditions are necessary for them to equal between 0.17 and 0.21. The third term in (4.1.30) is neglected since $|H_{TOP}|^2=1$ for a dual directional coupler.

The combined response of the digital filter and RF BSF within the passband relative to an interferer occupied frequency (f_{int}) is defined as:

$$\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p) = \frac{H_{DIG}(f_p)H_{BSF}(f_p)}{H_{DIG}(f_{\text{int}})H_{BSF}(f_{\text{int}})} \tag{4.1.32}$$

The magnitude response of the digital filter and RF BSF at f_{int} must result in FF attenuation; therefore can be expressed in terms of the other components.

$$\begin{aligned} &|H_{C1t}H_{TOP}H_{C2t}| = |H_{C1c}H_{RF1}H_{DIG}(f_{\text{int}})H_{RF2}H_{BSF}(f_{\text{int}})H_{C2c}| \\ &\therefore |H_{DIG}(f_{\text{int}})H_{BSF}(f_{\text{int}})| = \left|\frac{H_{C1t}H_{TOP}H_{C2t}}{H_{C1c}H_{RF1}H_{RF2}H_{C2c}}\right| \end{aligned} \tag{4.1.33}$$

The purpose of introducing a new term in (4.1.32) is to replace the $|H_{ADC}|$ and $|H_{DAC}|$ terms, since they contain a multitude of design variables. These terms are replaced with the ratio of the digital filter and RF BSF responses in the passband to that at f_p. The magnitude response of the digital components and RF BSF in the passband becomes:

$$|H_{DIG}(f_p)H_{BSF}(f_p)| = \left|\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)\frac{H_{C1t}H_{TOP}H_{C2t}}{H_{C1c}H_{RF1}H_{RF2}H_{C2c}}\right| \tag{4.1.34}$$

The ratio of the BSF attenuation within the passband to that at the nearest attenuation frequency outside the passband, $\Delta H_{BSF}(f_p)$, is assumed to be -10 dB. This assumption is based on 20 dB of BSF attenuation in the passband, and the first stopbands starting where the BSF exhibits 10 dB of attenuation.

The terms from the last two lines of (4.1.30) are shown below, but with the substitution from (4.1.34):

$$F_{FF|BOT}(f_p) = F_{RF1}\left|\frac{\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)}{H_{C1c}}\right|^2 + (F_{DIG}-1)\left|\frac{\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)}{H_{C1c}H_{RF1}}\right|^2$$
$$+(F_{RF2}-1)\left|\frac{H_{RF2}H_{BSF}(f_p)H_{C2c}}{H_{C1t}H_{TOP}H_{C2t}}\right|^2 - \left|\frac{H_{BSF}H_{C2c}}{H_{C1t}H_{TOP}H_{C2t}}\right|^2 \quad (4.1.35)$$

The full-scale analog power levels of the ADC's and DAC's are related to the maximum RF input power of the hybrid RF-DSP FF filter. The input power to the hybrid RF-DSP FF filter that causes a full-scale digital output from the ADC's is denoted by $P_{FF,fs}$. The DAC full-scale power should be large enough to attenuate a signal in the RF path with power $P_{FF,fs}$ outside of the passband. The relationships between these full-scale power levels are given by (4.1.36). The $H_{BSF,min}(f_{int})$ term in this equation represents the minimum gain of the RF BSF outside the passband.

$$P_{ADC,fs} = P_{FF,fs}\left|H_{C1c}H_{RF1}\right|^2$$
$$P_{FF,fs} = P_{DAC,fs}\left|\frac{H_{C2c}H_{RF2}H_{BSF,\min}(f_{\text{int}})}{H_{C1t}H_{TOP}H_{C2t}}\right|^2 \quad (4.1.36)$$

Substituting these equations into (4.1.35) to remove H_{RF1} and H_{RF2} yields:

$$F_{FF|BOT}(f_p) = F_{RF1}\left|\frac{\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)}{H_{C1c}}\right|^2$$
$$+(F_{DIG}-1)\frac{P_{FF,fs}}{P_{ADC,fs}}\left|\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)\right|^2 \quad (4.1.37)$$
$$+(F_{RF2}-1)\frac{P_{FF,fs}}{P_{DAC,fs}}\left|\frac{H_{BSF}(f_p)}{H_{BSF,\min}(f_{\text{int}})}\right|^2 - \left|\frac{H_{BSF}H_{C2c}}{H_{C1t}H_{TOP}H_{C2t}}\right|^2$$

There are four terms in (4.1.37) that must sum to 0.17 or less to achieve a sub 1 dB noise figure. The maximum input interferer power for this analysis is set to -15 dBm, which corresponds to the maximum blocker power in the WCDMA standard[*38*]. Each of the

four terms is discussed individually. The goal of the noise factor contributions from the first three terms in (4.1.37) are 0.01, 0.15 and 0.01, respectively.

The first term contains the coupling factor of $C1$ in the denominator, which removes the possibility of using a 30 dB coupling factor. A trade-off between this term and the results from Fig. 4.1.9 indicate that $C1$ should have a 20 dB coupling factor. This coupling reduces the range of values for the terms in (4.1.37) to between 0.17 and 0.18, depending on whether $C3$ is a 20 or 30 dB coupler, respectively. The value of F_{RF1} corresponds to a direct-down conversion receiver, which can be on the order of 2 dB. For the first term in (4.1.37) to be less than 0.01, with F_{RF1}= 2 dB, and $|H_{C1c}|$= -20 dB, then $|\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)|$< -42 dB. As previously stated, the minimum value of $\Delta H_{BSF}(f_p)$ is assumed to be -10 dB, so the digital filter needs to provide some notches for the additional 32 dB over the passband.

The second term in (4.1.37) is the noise factor contribution from the digital components, which includes the noise factor of the digital components from (4.1.28). The digital noise terms due to rounding after the multipliers and after the IIR stage are dropped from (4.1.28). In a low noise implementation these noise sources can be insignificant if no loss of precision occurs until the data reaches the DAC's. Substituting the digital noise factor into the second term in (4.1.37), and applying the digital precision assumptions yields:

$$F_{FF|DIG}(f_p) = \frac{P_{FF,fs}}{kT_0B}\left\{\frac{\left|\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)\right|^2}{SNR_{ADC}} + \frac{\left|\alpha_{DAC}\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)\right|^2}{SNR_{DAC}\left|H_{DSP}(f_p)\right|^2}\right\} \quad (4.1.38)$$

The value of $|\Delta H_{DIG}(f_p)\Delta H_{BSF}(f_p)|$ in (4.1.38) has already been constrained to be less than -42 dB, due to the first term in (4.1.37), and the full-scale power at the RF input is -15

dBm. If the digital clock frequency is 200 MHz, then the noise factor contribution from the first term in (4.1.38) is 0.001 when SNR_{ADC}>67 dB. The second addend in (4.1.38) is due to the DAC and is simplified using (4.1.32):

$$F_{FF|DAC}(f_p) = \frac{P_{FF,fs}}{kT_0B}\left\{\frac{1}{SNR_{DAC}}\left|\frac{\alpha_{DAC}\Delta H_{BSF}(f_p)}{H_{DSP}(f_{\text{int}})}\right|^2\right\} \tag{4.1.39}$$

The ratio of the DAC full-scale value to the ADC-full scale value (α_{DAC}) should be on the same order as the maximum gain of the digital filter at the interferer occupied frequencies $|H_{DSP,max}(f_{int})|$. This minimum limit on α_{DAC} is required to avoid saturation. If this ratio is larger than $|H_{DSP,max}(f_{int})|$, then there is wasted overhead in the digital filter that will increase the noise figure. The frequencies where $H_{DSP}(f_{int})$ is maximum correspond to the nearest frequencies to the passband where attenuation is required. This maximum corresponds to the frequencies where the digital filter provides the most gain to compensate for the wide attenuation bandwidth of the RF BSF. The result of this discussion is that $\alpha_{DAC} \approx |H_{DSP,max}(f_{int})|$.

Under the assumptions presented above, the value of SNR_{DAC} must be greater than 77 dB for the expression in (4.1.39) to be less than 0.15. This value of SNR_{DAC} requires a DAC with an effective number of bits (ENOB) equal to 12.5.

The third term in (4.1.37) is due to the noise generated within the *RF*2 components. With a -15 dBm full-scale RF input to the hybrid RF-DSP FF filter, the remaining terms require $(F_{RF2}-1)/P_{DAC,FS}<3.16$ to ensure the entire third term is less than 0.01. The DAC used in the hardware prototype has $P_{DAC,FS}$=10 dBm [***39***], which would limit F_{RF2} to be

less than 31.2, which corresponds to a 14.9 dB noise figure. This noise figure specification is realistic for a cascaded modulator and amplifier .

The fourth term in (4.1.37) is negligible, and can be ignored without concern since it is negative and would only reduce the noise factor.

This analysis has laid out one set of conditions required for a sub 1 dB noise figure. The most stringent specification is for the DAC SNR. These conditions are feasible; therefore the hybrid RF-DSP FF filter can be designed to exhibit a 1 dB noise figure.

2. Under what conditions does the output noise floor degrade attenuation performance?

If the output noise floor is larger than the interferer PSD after attenuation, then the noise would appear to degrade the attenuation. The output noise PSD at interferer frequencies is obtained from (4.1.1), where f is substituted with f_{int}. Several approximations are made to simplify this expression at frequencies occupied by interferers (f_{int}):

- The magnitude response of the top and bottom paths are equal at interferer locations: $|H_{C1t}H_{TOP}H_{C2t}|=|H_{C1c}H_{RF1}H_{DIG}(f_{int})H_{RF2}H_{BSF}(f_{int})H_{C2c}|$.
- The top path is a lossless passive component: $F_{TOP}=|H_{TOP}|^{-2}=1$.
- Any terms in (4.1.1) that are expected to be near the noise floor are ignored. These include the first and last lines of (4.1.1). The justification for dropping these terms is that the filter is designed to attenuate strong interferers by 10's of decibels, which is expected to be well above the thermal noise floor.

The output noise PSD within interference bands, subject to the above approximations, is given by:

$$
\begin{aligned}
kT_{FF,out}(f_{\text{int}}) \approx\ & kT_0 F_{RF1}\left|\frac{H_{C1t}H_{C2t}H_{C3t}}{H_{C1c}}\right|^2 \\
& + kT_0(F_{DIG}-1)\left|\frac{H_{C1t}H_{C2t}H_{C3t}}{H_{C1c}H_{RF1}}\right|^2 \\
& + kT_0(F_{RF2}-1)\left|H_{RF2}H_{BSF}H_{C2c}H_{C3t}\right|^2
\end{aligned}
\tag{4.1.40}
$$

The three terms in (4.1.40) are due to the noise contribution from $RF1$, the digital components and $RF2$, respectively. The first term increases the output noise floor by the product of the coupling factor of $C1$and the noise figure of $RF1$. If $C1$ is a 20 dB coupler and F_{RF1} has a 2 dB noise figure, then this term is approximately 22 dB above the thermal noise floor. Several variables are present in the third term, including the BSF magnitude response at an interference frequency. The largest magnitude response away from the passband is 0 dB for an RLC based BSF. If $C2$ is a 20 dB coupler, then for this term to contribute as much noise as the first term, then $(F_{RF2}\text{-}1)|H_{RF2}|^2 \approx 42$ dB. This product is very large, meaning the third term will generally contribute less output noise than the first term at an interference frequency. The combined effect of the first and third terms is to raise the output noise floor by less than 25 dB at an interference frequency. It is expected that interferers will not be attenuated down this close to the noise floor, hence the noise from $RF1$ and $RF2$ does not degrade interference attenuation, under the conditions specified.

The second term in (4.1.40) is due to the digital components including the data converters. The F_{DIG} expression from (4.1.28) and the relationship in (4.1.36) are substituted into the second term in (4.1.40):

$$kT_{FF,out|DIG}(f_{\text{int}}) = \frac{P_{FF,fs}}{B}\left(\frac{1}{SNR_{ADC}} + \frac{\alpha_{DAC}^2}{SNR_{DAC}\left|H_{DSP}(f_{\text{int}})\right|^2} + \frac{5N\alpha_{MULT}^2}{3\left|H_{DSP}(f_{\text{int}})\right|^2 2^{2b_{MULT}}} + \frac{2\alpha_{ADD-Z}^2}{3\left|H_{POLES}(f_{\text{int}})\right|^2 2^{2b_{ADD-Z}}}\right) \quad (4.1.41)$$

The thru port magnitude responses of the directional couplers are omitted in (4.1.41) due to their negligible impact. Consider the case where the FF filter input power is -15 dBm and the digital clock runs at 200 MHz (B=100 MHz). The multiplication factor in front of (4.1.41) for this scenario is –95 dBm/Hz. Each addend in (4.1.41) reduces this factor by an amount that depends on the digital design variables. To ensure the noise generated by the ADC causes the noise floor to rise only 22 dB above the thermal noise floor, then: SNR_{ADC}>57 dB. The same constraint exists for SNR_{DAC} when $\alpha_{DAC} \approx |H_{DSP}(f_{int})|$. The 57 dB limit on the SNR's corresponds to an ENOB of 9.2, which is achievable. The other two terms in (4.1.41) can be negligible be ensuring adequate scaling and bit widths throughout the digital filter.

This section has shown that noise generated by the lower FF branch can be less than 30 dB above the noise floor over interferer occupied bands without requiring exotic components. This noise floor means that interferers can be attenuated down to 30 dB above the noise floor before the added noise appears to degrade the attenuation performance.

4.1.5 Simulation for Verification of F_{DIG}

Most of the RF noise analysis in the previous sections is straight from an undergraduate textbook. Digital components are not generally incorporated into an RF noise analysis.

Consequently, simulations were performed to validate the analysis using Simulink. The digital filter in Simulink includes fixed-point settings, which permits truncation, rounding or scaling after arithmetic operations. Fig. 4.1.10 and 4.1.11 show the model used in Simulink.

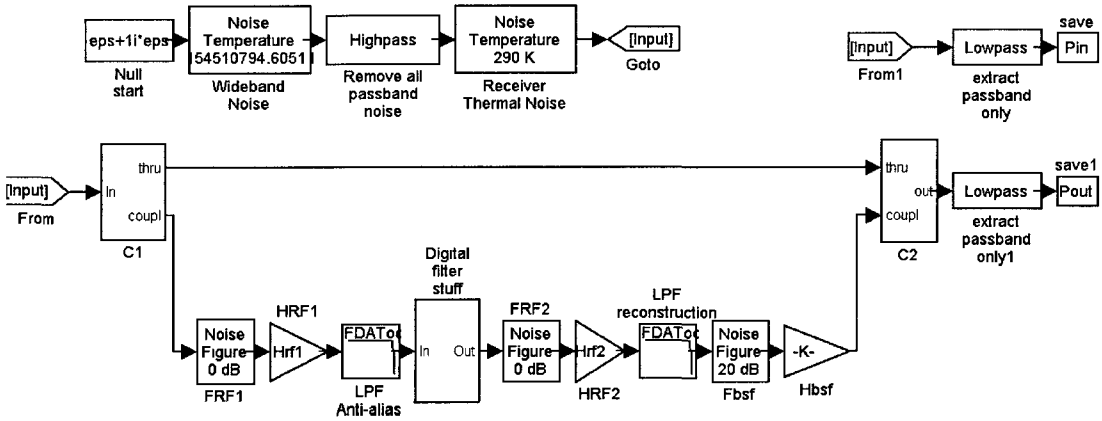

Figure 4.1.10 - Simulink model for digital noise analysis

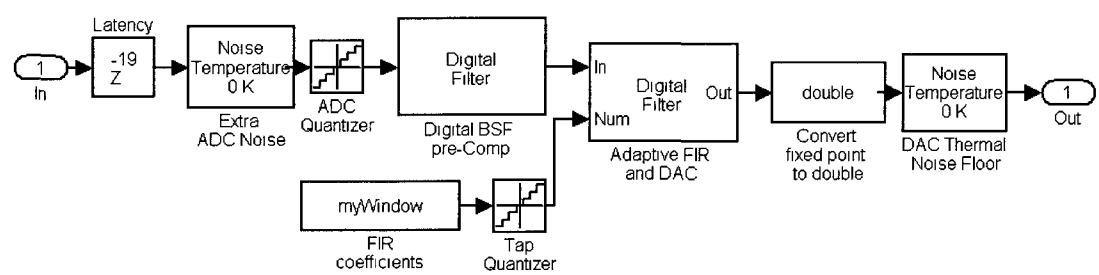

Figure 4.1.11 – Details of 'Digital filter stuff' block from Fig. 4.1.10

The cascade of blocks at the top of Fig. 4.1.10 generate the filter's input signal. This input signal has a PSD 66 dB above kT_0 outside the passband, and is equal to kT_0 within the passband. The large PSD outside the passband ensures the input signal spans many quantization levels in the digital components. The RF passband bandwidth is set to 20 MHz (B_P=10 MHz at baseband).

The following fixed component values were used in the simulation:

- H_{TOP} is a zero-length transmission line ($|H_{TOP}|=F_{TOP}=1$).
- C_1 and C_2 are ideal 20 dB couplers.
- H_{RF1} and H_{RF2} provide 20 dB of gain, and have noise factors $F_{RF1}=0$ dB, and $F_{RF2}=0$ dB (except in last row in Table 4.1). These blocks are labelled in Fig. 4.1.10.
- H_{BSF} provides 20 dB of attenuation over the passband. The BSF is implemented with a 20 dB attenuator, along with a noise figure block.
- The digital clock frequency is 200 MHz (B=100 MHz).
- Anti-alias and reconstruction filters bookend the digital filter, each with a 100 MHz bandwidth.
- The FIR part of the digital filter has 32 randomized taps. The average magnitude-squared response of $|H_{DSP}|$ in the passband, including that of the digital BSF compensation module is defined as: $\left|\widehat{H}_{DSP}\right| = \int_{-B_p}^{B_p} |H_{DSP}|^2 \, df/2B_p$.
- The DIG BSF compensation is setup to provide a maximum of 9 dB of dynamic range increase at ±15 MHz. This compensation corresponds to a maximum gain in the DIG BSF module of 6 dB, and requires poles at $z=0.7e^{\pm j0.3}$.
- The magnitude response of the ADC's is set so that an input voltage level of 1 corresponds to a digital word value of 1. The DAC's adhere to the same 1:1 relationship. These result in: $|H_{ADC}|=|H_{DAC}|=1$.
- The following digital word full-scale ratios were used throughout these simulations: $\alpha_{DAC}=2^2$, $\alpha_{MULT}=2^{-1}$, & $\alpha_{ADD\text{-}Z}=1$. These values were chosen to optimize digital resources while avoiding saturation.

The input and output noise in the passband are recorded by the blocks on the far right in Fig. 4.1.10. This data is then processed by (4.1.42) to extract the simulated noise factor (F_{sim}). The terms $S_{in}(f)$ and $S_{out}(f)$ are the PSD's of the input and output signals. These are obtained by processing the recorded data with a built-in Matlab FFT averaging function. The parameter F_{SIM} is the noise factor averaged over the passband.

$$F_{SIM} = \frac{1}{2B_p} \int_{2B_p} \frac{S_{out}(f)/|H_{FF}(f)|^2}{S_{in}(f)} df \tag{4.1.42}$$

The noise factor for the hybrid RF-DSP FF filter is defined in (4.1.30) with a passive component in the top path. This equation requires F_{DIG} from (4.1.28). The noise factor in (4.1.30) is expressed as a function of frequency, so this expression is averaged over the passband to obtain a result that is comparable with the simulation results:

$$F_{EXP} = \frac{1}{2B_p} \int_{f_0-B_p}^{f_0+B_p} F_{FF}(f) df \tag{4.1.43}$$

F_{SIM} and F_{EXP} are shown for several filter realizations in Table 4.1. The filter realizations differ in the digital components details. In the last row of the table the noise figures of the *RF*1 and *RF*2 modules are both changed from 0 to 6 dB.

Table 4.1 – Noise simulation results for varying digital resources

$\lvert\hat{H}_{DSP}\rvert$	*ADC*		*DAC*		*FIR MULT*	*Digital BSF*	F_{SIM}	F_{EXP}
(dB)	SNR(dB)	Bits	SNR(dB)	Bits	Bits	Bits		
-0.30	70	12	194	32	32	32	2.0355	1.9918
3.97	194	32	70	12	32	32	4.5340	4.5095
2.101	194	32	194	32	12	32	5.1975	5.4026
4.893	194	32	194	32	32	12	6.0154	5.9683
4.2	70	12	83	14	18	14	3.814	3.7814
-35.93	70	12	83	14	18	14	1.0866	1.0848
-1.857	70	12	83	14	18	14	3.8927	3.9182

The simulation and theoretical results are very close for these seven realizations. The derived expressions can be used to ensure efficient resource allocation of the digital components. These equations provide closed-form relationships between noise performance metrics and design parameters.

4.2 Feedforward Attenuation Analysis

FF techniques involve the cancellation of two signals that come from two different paths. Attenuation in a FF system degrades as amplitude or phase mismatches increase between the two paths. Three causes of mismatch are:

- finite-precision representation of the filter tap coefficients
- group delay mismatch between the two signal paths caused primarily by the latency inherent to the digital components,
- steady-state error associated with the adaptive algorithm that drives the digital filter's tap coefficients.

The attenuation degradation due to the first two points are derived in this chapter. The focus of the following chapter is the adaptive algorithms for the hybrid FF filter. Suitably, the degradation due to the third point is deferred to the following chapter.

4.2.1 Filter Tap Quantization

Consider a set of digital filter coefficients that result in perfect FF cancellation at several frequencies for the hybrid RF-DSP FF filter. The depth of these perfect notches will degrade when the digital filter coefficients are represented by fixed-point values. The degradation will be continually changing since the digital filter is adaptive. A statistical analysis relating FF attenuation to filter tap quantization step-size (Q_{TAPS}) is attached as Appendix A. This step-size is a function of the number of bits (b_{TAPS}) used to represent the real and imaginary components of the digital filter taps. The step-size also depends on the maximum digital value used in the fixed-point representation ($v_{TAPS,fs}$).This value is the maximum value for the real and imaginary components of the tap coefficient that can be represented without saturation or wrap-around.

$$Q_{TAPS} = \frac{v_{TAPS,fs}}{2^{b_{TAPS}-1}} \tag{4.2.1}$$

The statistical analysis follows that in [*33*] for a direct-form digital filter. The result of this analysis is an upper limit on Q_{TAPS} such that FF attenuation will be better than $|H_{FF,des}(\omega_{int})|$ for 99.97% of the time. The upper bound is given by (4.2.2). In this equation, N_{TAPS} is the number of filter taps.

$$Q_{TAPS} \le \min_{\omega_{\text{int}}} \sqrt{\frac{3}{4N_{TAPS}}} \frac{\left|H_{FF,des}(\omega_{\text{int}})\right|}{\left|H_{C1c}H_{RF1}H_{ADC}H_{DAC}H_{RF2}H_{BSF}(\omega_{\text{int}})H_{C2c}\right|} \tag{4.2.2}$$

The inequality in (4.2.2) is rearranged using (4.2) to isolate the number of tap bits:

$$b_{TAPS} \ge \log_2 \frac{\sqrt{3N_{TAPS}} \cdot \left|H_{C1c}H_{RF1}H_{ADC}H_{DAC}H_{RF2}H_{BSF}(\omega_{\text{int}})H_{C2c}\right| v_{TAPS,fs}}{\left|H_{FF,des}(\omega_{\text{int}})\right|} \tag{4.2.3}$$

The following relationships are extracted from (4.2.3):

- Every time the number of taps (N_{TAPS}) is doubled, b_{TAPS} must be increased by half a bit to maintain the same FF attenuation.
- An increase in the desired FF attenuation, $|H_{FF,des}(\omega_{\text{int}})|$, by 3 dB must be accompanied by an increase of half a bit of b_{TAPS}
- The number of bits is limited by the maximum value of the magnitude response of the BSF.

This design equation is only valid to determine the limit of FF attenuation degradation due to coefficient quantization. It is used to verify the hardware prototype has a sufficiently small tap coefficient quantization step-size to prevent attenuation degradation. Group delay mismatch also degrades FF attenuation, which is reported next.

4.2.2 Group Delay Mismatch

Group delay mismatch between the paths of a FF system is well known to limit the attenuation bandwidth. The ADC's, DAC's and DSP have inherent latency due to pipelining, and filtering. Furthermore, the anti-alias and smoothing filters have finite group delay, which increases the group delay mismatch between the top and bottom FF

paths. The effect of this group delay mismatch is derived in this section. The number of digital filter taps is not limited in this derivation.

The transfer functions of the RF components in the hybrid FF filter can be represented with their baseband equivalent functions. The hybrid FF filter block diagram in Fig. 4.2.1 shows the transfer functions in discrete-time. The discrete-time domain is used in this analysis since this derivation is focused primarily on the digital filter.

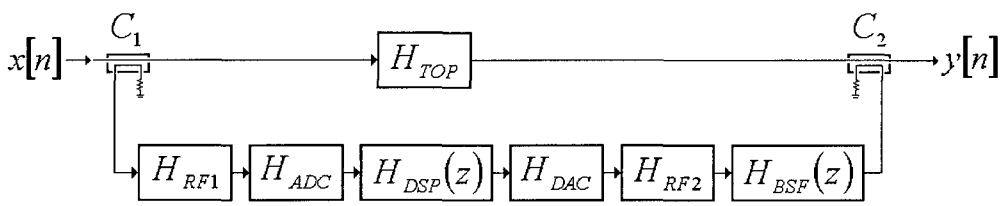

Figure 4.2.1 – Linear discrete-time model of hybrid RF-DSP FF filter

When the digital filter is in steady-state, the cascaded blocks in the bottom branch of the FF filter can be rearranged in the mathematical model without affecting the output signal. An equivalent block diagram is shown in Fig. 4.2.2.

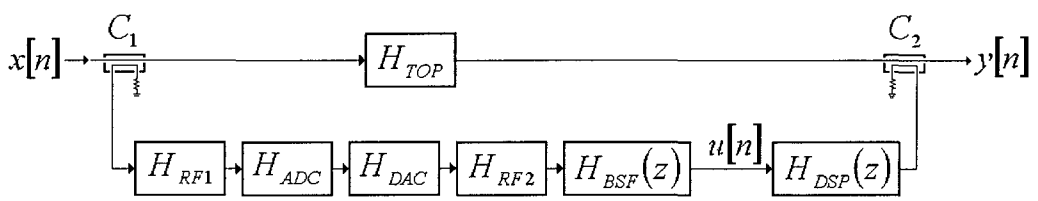

Figure 4.2.2 – Equivalent model to Fig. 4.2.1 when digital filter is in steady-state

In this formulation the coupling ports of the two directional couplers and the blocks in the bottom branch, excluding the digital filter, are combined into a block denoted $H_1(z)$. The through ports of the directional couplers and the components in the top FF path are combined into a block denoted $H_2(z)$. The compressed block diagram is shown in Fig. 4.2.3. The negative sign on the summation block is convention.

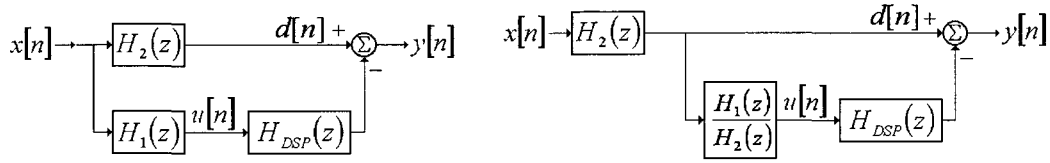

Figure 4.2.3 – Compressed block diagram a) version 1 b) version 2

When the interference environment is static then $E\{|y[n]|^2\}$ is minimized by making the digital filter equal to the Wiener filter [***40***]. Minimizing $E\{|y[n]|^2\}$ results in a whitening filter that notches interferers. The IIR causal Wiener filter is given by (4.2.4) [***40***]. A description of the terms in (4.2.4) is given after the equation.

$$H_{opt,\infty}(z) = \frac{1}{\sigma_u^2 G_u(z)}\left[\frac{S_{du}(z)}{G_u^*(1/z^*)}\right]_+ \tag{4.2.4}$$

The canonical spectral factorization of the z-spectrum of the process $u[n]$ is:

$$S_{uu}(z) = \sigma_u^2 G_u(z) G_u^*\left(\frac{1}{z^*}\right) = H_1(z) H_1^*\left(\frac{1}{z^*}\right) S_{xx}(z) \tag{4.2.5}$$

The latter relationship in (4.2.5) is obtained directly from Fig. 4.2.3a. In (4.2.5), $G_u(z)$ is the minimum phase factor, and σ_u is a real scalar. The block diagram in Fig. 4.2.3b was rearranged to aid in determining $S_{du}(z)$. In (4.2.4), $S_{du}(z)$ is the z-transform of the cross-correlation between $d[n]$ and $u[n]$:

$$S_{du}(z) = \frac{H_1^*\left(\frac{1}{z^*}\right)}{H_2^*\left(\frac{1}{z^*}\right)} S_{dd}(z) = H_1^*\left(\frac{1}{z^*}\right) H_2(z) S_{xx}(z) \tag{4.2.6}$$

The $[\square]_+$ operator in (4.2.4) is the causal component of the inverse z-transform. Equation (4.2.7) shows the proper use of this operator. This operator differs from the one-sided z-transform, since for said transform $h[n]=0$ for $n<0$.

$$\begin{aligned} \{H(z)\}_+ &= ZT\{h_+[n]\} \\ H(z) &= ZT\{h[n]\} \\ h_+[n] &= h[n]u[n] \end{aligned} \tag{4.2.7}$$

The expressions in (4.2.5) and (4.2.6) are substituted into (4.2.4) to arrive at the Wiener filter transfer function.

$$\begin{aligned} H_{opt,\infty}(z) &= \frac{1}{\sigma_u^2 G_u(z)}\left[\frac{H_2(z)H_1^*\left(\frac{1}{z^*}\right)S_{xx}(z)}{G_u^*\left(\frac{1}{z^*}\right)}\right]_+ \\ &= \frac{1}{\sigma_u^2 G_u(z)}\left[\frac{H_2(z)H_1^*\left(\frac{1}{z^*}\right)\sigma_u^2 G_u(z)G_u^*\left(\frac{1}{z^*}\right)}{G_u^*\left(\frac{1}{z^*}\right)H_1(z)H_1^*\left(\frac{1}{z^*}\right)}\right]_+ \\ &= \frac{1}{G_u(z)}\left[G_u(z)\frac{H_2(z)}{H_1(z)}\right]_+ \end{aligned} \tag{4.2.8}$$

The next step in this derivation is to decompose the function $H_2(z)/H_1(z)$ into a minimum phase component and a nominal delay term:

$$\frac{H_2(z)}{H_1(z)} = z^D \frac{H_2'(z)}{H_1'(z)} \tag{4.2.9}$$

The value of D is the nominal delay from $H_2(z)/H_1(z)$. This value includes latency and group delay. The minimum phase component is $H_2'(z)/H_1'(z)$.

The PSD of multiple interferers at the input to a receiver can appear spiky, which is characteristic of an auto-regressive (AR) process. This spikiness means that

$G_u(z)H_2'(z)/H_1'(z)$ is well modeled by an all-pole transfer function. The poles are denoted by p_i for i=1,2...M. A partial fraction expansion (PFE) of $G_u(z)H_2'(z)/H_1'(z)$ yields:

$$G_u(z)\frac{H_2'(z)}{H_1'(z)} = \sum_{i=1}^{M} A_i/\left(1 - p_i z^{-1}\right) \tag{4.2.10}$$

The A_i values are the coefficients from the PFE. The inverse z-transform of this term is:

$$IZT\left\{\left[G_u(z)\frac{H_2'(z)}{H_1'(z)}\right]\right\} = \left(g_u * h_2'/h_1'\right)[n] = \sum_{i=1}^{M} A_i p_i^n u[n] \tag{4.2.11}$$

Putting the delay term back into the transfer function gives the following inverse z-transform:

$$IZT\left\{\left[z^D G_u(z)\frac{H_2'(z)}{H_1'(z)}\right]\right\} = \left(g_u * h_2'/h_1'\right)[n] * \delta[n+D] = \sum_{i=1}^{M} A_i p_i^{n+D} u[n+D] \tag{4.2.12}$$

The last modification to this term is made by isolating the causal component of the impulse response:

$$IZT\left\{\left[z^D G_u(z)\frac{H_2'(z)}{H_1'(z)}\right]_+\right\} = \sum_{i=1}^{M} A_i p_i^{n+D} u[n] \tag{4.2.13}$$

Converting this result back to the z-domain, and using (4.2.10) gives:

$$\left[z^D G_u(z)\frac{H_2'(z)}{H_1'(z)}\right]_+ = \sum_{i=1}^{M} p_i^D A_i/\left(1 - p_i z^{-1}\right) \tag{4.2.14}$$

The LHS of this result is equal to the causal component in (4.2.8). The transfer function of the optimum Wiener filter is obtained by substituting the result from (4.2.14) back into (4.2.8).

$$H_{opt,\infty}(z)=\frac{1}{G_u(z)}\sum_{i=1}^{M}p_i^D A_i/\left(1-p_i z^{-1}\right) \tag{4.2.15}$$

When $H_{DSP}(z)=H_{opt,\infty}(z)$, the transfer function of the FF filter in Fig. 4.2.3b is:

$$H_{FF}(z)=1-\frac{\sum_{i=1}^{M}p_i^D A_i/\left(1-p_i z^{-1}\right)}{z^D\sum_{i=1}^{M}A_i/\left(1-p_i z^{-1}\right)} \tag{4.2.16}$$

In the z-domain, interference frequencies are located on the unit circle closest to the poles of the AR model. At the frequency nearest to pole i, $|1-p_i z_i^{-1}|<<|1-p_j z_i^{-1}|$, for $j\neq i$. This means that one term will dominate each of the summations in $H_{opt,\infty}(z)$ when the poles of the AR model are not densely clustered. At the frequencies nearest to pole i,

$$H_{FF}(z_i)\approx 1-\frac{p_i^D A_i/\left(1-p_i z^{-1}\right)}{z^D A_i/\left(1-p_i z^{-1}\right)}=1-\left(\frac{p_i}{z}\right)^D \tag{4.2.17}$$

This transfer function is changed into analog frequency domain, ω:

$$H_{FF}(\omega)\approx 1-|p_i|e^{j(\angle p-\omega T)D} \tag{4.2.18}$$

This transfer function is identical to that of a conventional RF FF system with a delay mismatch of T_sD. The notch of this system achieves a depth of $1-|p_i|$, and is located at the same location as the poles of the input signal's AR model. Recall that T_s is the time step

in the digital domain, and D is the group delay mismatch between $H_1(z)$ and $H_2(z)$ in Fig. 4.2.3a, in terms of clock periods. In other words, D is an integer without units.

The result in (4.2.18) is a very simple expression that relates the maximum attenuation of the FF filter to the group delay mismatch (D) and the interferer's pole amplitude in the AR model. The approximation made in (4.2.17) is valid near frequencies occupied by interferers.

4.2.3 Number of Filter Taps

The interference cancellation expression in the previous section was for an IIR digital filter, and the corresponding impulse response represents the IIR Wiener solution. The digital filter within the hybrid RF-DSP FF filter has an N-tap FIR, which nominally converges towards the N tap FIR causal Wiener solution. The MSE of the FIR filter converges to that of the IIR solution for a sufficiently large number of FIR taps. Arbitrarily increasing the number of FIR taps beyond some limit does not guarantee measurable attenuation improvement. This statement is easier to demonstrate with measurements than with theory; therefore is deferred to Section 7.3.1.

4.3 Analyses Summary

Analytical expressions have been derived relating the design variables to noise and attenuation performance metrics. These expressions can be used to choose component values based on noise and attenuation specifications. These results are valuable to trade-off hardware resources with filter performance, without needing to resort to exhaustive

simulations. The results of this chapter are used in Chapter 6 to choose the components for the hardware prototype.

A linear analysis is used throughout this chapter. This assumption is problematic if the RF or analog components are sensitive to the amplitude of their input signals, and if this amplitude is highly dynamic.

5 Algorithm Development

Adaptive algorithms are required to continuously update the digital filter coefficients in a dynamic interference environment. The digital filter must also be continuously updated to compensate for RF and analog component drift. Furthermore, the BSF contains a tunable component that is adaptively controlled to maintain a deep notch in the passband. The digital filter and adaptive algorithms are implemented in an FPGA, therefore computational resources are limited. These adaptive algorithms are presented in this chapter and verified with simulations.

5.1 Filter Optimization Strategy

The purpose of the hybrid RF-DSP FF filter in a receiver application is to attenuate out-of-band interferers and transmitter leakage. The output of this filter would ideally have a uniform PSD (white noise) except within the receiver's passband. The filter should have unity magnitude and a linear phase response in the receiver's passband. A filter that converts a non-uniform PSD to a uniform PSD is a whitening filter. A plethora of literature exists to determine the coefficients of the whitening filter, including adaptive techniques. This literature is in the field of auto-regressive moving average (ARMA) modeling, and has applications in channel equalization, system identification, spectrum estimation, linear predictive coding and active noise control.

A block diagram of a whitening filter is shown in Fig. 5.1.1. In this figure a white noise process, $w[n]$, with average power σ_w^2, is input to an ARMA model. The transfer function of this model is: $H_{ARMA}(z)=H_{num}(z)/H_{den}(z)$. The output of the ARMA model is termed an ARMA process. The PSD of this ARMA process is completely dependent upon $H_{ARMA}(z)$ and σ_w. The goal of ARMA modeling is to generate a filter to re-whiten the ARMA process. The sought after whitening filter is the inverse of the ARMA model: $H_{white}(z)=H_{den}(z)/H_{num}(z)$. The only signal that is available to aid in generating this filter is the ARMA process, $x[n]$. The values of the white noise process at the input to the ARMA model, $w[n]$, are unavailable.

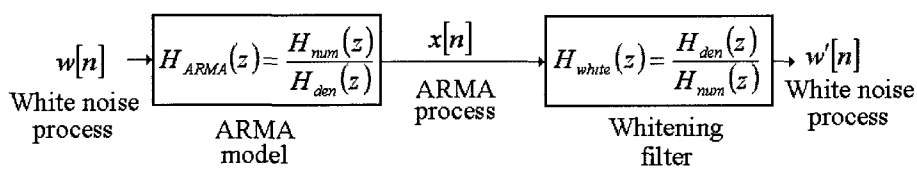

Figure 5.1.1 – ARMA process and whitening filter

The most common type of ARMA modeling is the auto-regressive (AR) case: $H_{num}(z)=1$ [*41*]. An AR process results from passing white noise through an all-pole filter. The corresponding whitening filter, $H_{white}(z)$, is then FIR. An example of an AR process is the signal at the input to a receiver. Interferers are characteristically spiky in the frequency domain, therefore are modeled well with white noise filtered by an all-pole filter. An example of this characteristic is illustrated in Fig. 5.1.2. In this figure the AR model is represented by its pole locations in the z-domain. AR techniques exist that are capable of determining $H_{white}(z)$ and σ_w based only on the AR process [*41*]. These techniques

involve choosing the coefficients that minimize the power at the output of the whitening filter [*41*].

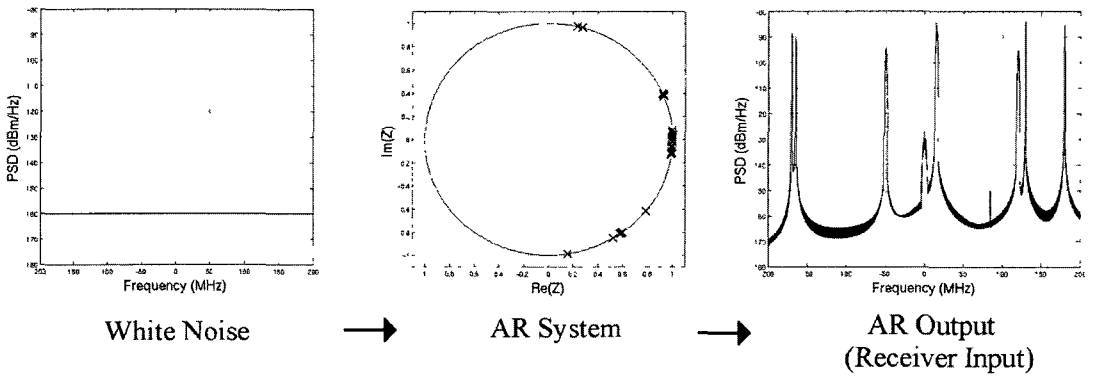

Figure 5.1.2 – AR process generated by passing white noise through an all-pole filter

The goal of the hybrid RF-DSP FF filter is to whiten the receiver's input signal outside the passband. For this application, the ARMA process in Fig. 5.1.1 corresponds to the receiver input signal, and the whitening filter corresponds to the hybrid RF-DSP FF filter. The digital filter is only one component in the whitening filter, but is the key component to permit adaptive whitening. With the AR approach the hybrid RF-DSP FF filter, including the digital filter must be FIR.

With the digital filter chosen to be FIR, the next step is to choose how the FIR coefficients are to be selected. There are two general approaches for digital filter design: deterministic and stochastic. A deterministic approach requires knowledge of the frequency bands occupied by interferers. PSD estimation is required to capture this frequency information. A deterministic digital filter is optimized by ensuring the hybrid RF-DSP FF filter has stopbands over the interferer bands. The stopbands are those of the hybrid FF filter, and not the digital filter. A benefit of the deterministic approach is the

filter taps only need to change when the interference scenario changes. A drawback is the computational burden associated with PSD estimation.

The stochastic approach was chosen for the filter design, assuming an AR process input. Standard linear adaptive filters adapt their taps in a manner that minimizes the output signal power ($w'[n]$ in Fig. 5.1.1). In both approaches to digital filter design, the transfer functions of the RF and analog components in the hybrid RF-DSP FF filter must be estimated. In this chapter three adaptive stochastic algorithms are presented for adapting the digital filter taps, estimating the RF and analog components, and tuning the BSF. The system level model is described prior to the algorithm presentations.

5.2 Hybrid RF-DSP FF Filter Model

The hybrid RF-DSP FF filter system level block diagram is shown in Fig. 5.2.1. In Fig. 5.2.2 the system level blocks have been replaced with discrete-time transfer functions. Baseband equivalent transfer functions are used in place of the RF components. The transfer function $H_1(z)$ includes the coupling port of the first coupler, down-conversion, and the ADC's that feed the digital filter. The components following the digital filter are represented by $H_2(z)$. These components include the DAC's, up-conversion, T-BSF and the coupling port of the second coupler. The coupling port of the third coupler, RF/Analog 3 and the adjacent ADC's are modelled by $H_3(z)$. It is assumed the filter operates within the linear range of the RF and analog components, so their non-linear effects are not modeled. The signals available for the adaptive algorithm are at the input ($x_{est}[n]$) and output of the digital filter, and the digital output ($y_{est}[n]$) of the ADC's that follow RF/Analog 3.

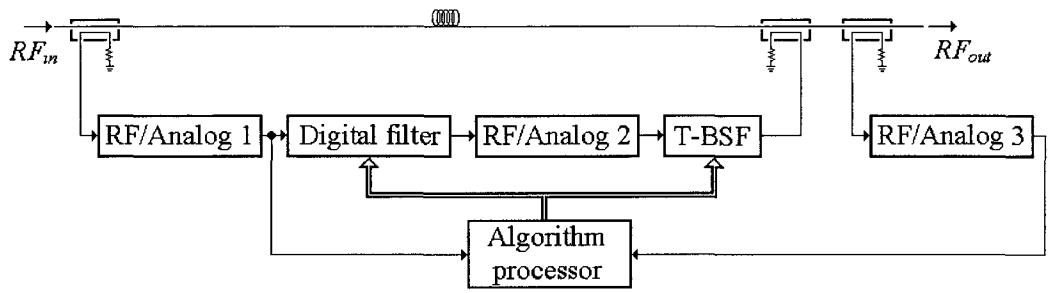

Figure 5.2.1 – Block diagram of hybrid RF-DSP FF filter

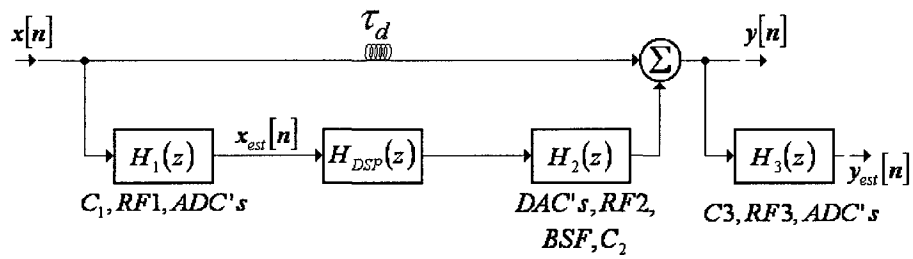

Figure 5.2.2 – Linear discrete-time model of hybrid RF-DSP FF filter

5.2.1 Model Manipulation for Adaptive Whitening

The standard form of a linear adaptive filter is shown in Fig. 5.2.3 [*42*]. The standard adaptive filter uses the input signal and desired reference signal to generate an error signal. The filter adapts in a manner that minimizes the power in the error signal. The primary advantages of this filter structure are: the goal function is convex, and existing adaptive algorithms can be used to find the goal function's global minimum. The least mean squares (LMS) and recursive least squares (RLS) adaptive algorithms can function with only the error and input signals, and do not require knowledge of the desired output signal [*42*]. The strategy used in this research is to manipulate the model in Fig. 5.2.2 so that it takes the form of Fig.5.2.3.

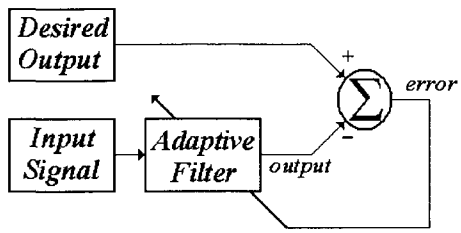

Figure 5.2.3 – Standard linear adaptive filter

The goal of the hybrid RF-DSP FF filter is to minimize the power in $y[n]$, in Fig. 5.2.2. The same goal exists for the error signal in Fig. 5.2.3. The $y[n]$ values are not available for processing since this signal is at RF. The $y_{est}[n]$ values are the baseband equivalent estimates of $y[n]$, and are available. The optimal filter coefficients that minimize the power in $y_{est}[n]$ will only be significantly different from those that minimize $y[n]$ if there are large notches or peaks in the $H_3(z)$ transfer function. Within the operating bandwidth of the hybrid RF-DSP FF filter, $H_3(z)$ consists of the passband of a LPF, an ADC, and a wideband down-converter and gain stage – none of which exhibit notches or peaks. The model in Fig. 5.2.2 is manipulated to take the form of Fig. 5.2.3 with $y_{est}[n]$ taking the place of the error signal. The $H_3(z)$ block in Fig. 5.2.2 is pushed back into both paths to form the mathematically equivalent block diagram in Fig. 5.2.4. In this figure $H_{23}(z)=H_2(z)H_3(z)$.

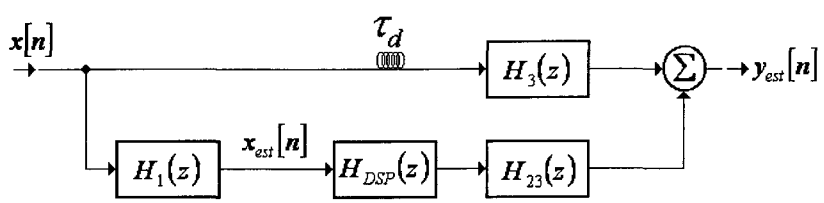

Figure 5.2.4 – Equivalent model to Fig. 5.2.2

If the digital filter coefficients do not change then $H_{DSP}(z)$ and $H_{23}(z)$ can be swapped in the model without affecting $y_{est}[n]$. This modified model is shown in Fig. 5.2.5.

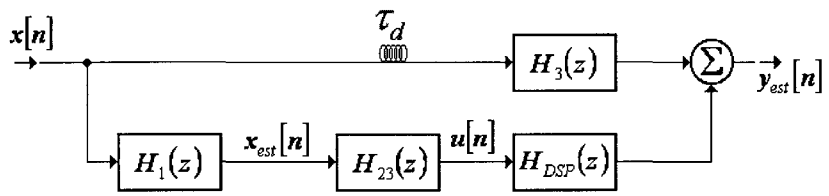

Figure 5.2.5 – Equivalent model to Fig. 5.2.4 when digital filter coefficients are fixed

The model in Fig. 5.2.5 represents the standard linear adaptive filter in Fig. 5.2.3. The desired output is the output signal of $H_3(z)$ in the top branch. The adaptive filter is the digital filter, $H_{DSP}(z)$, and the input signal is the convolution of $x_{est}[n]$ with $h_{23}[n]$. The latter term is the impulse response of $H_{23}(z)$. The error signal is $y_{est}[n]$.

There are two roadblocks associated with the model in Fig. 5.2.5. First, the result of the convolution of $x_{est}[n]$ with $h_{23}[n]$ is not readily available. This value can be estimated by passing $x_{est}[n]$ through a digital filter with a transfer function that approximates $H_{23}(z)$. The output of this filter is the input signal to the adaptive filter, and is denoted by $u_{est}[n]$. Fig. 5.2.6 is a modified version of Fig. 5.2.2, which illustrates how $u_{est}[n]$ is generated. The implementation shown in Fig. 5.2.6 is popular in active acoustic and vibration noise control systems, and is the basis of the filtered-x LMS (FXLMS) algorithm [*43*]. The FXLMS algorithm multiplies $u_{est}[n]$ and $y_{est}[n]$ together to estimate the gradient. The 'X' in the FXLMS algorithm refers to 'u' in this chapter. Due to this difference in terminology, this algorithm could be thought as the filtered-u LMS algorithm.

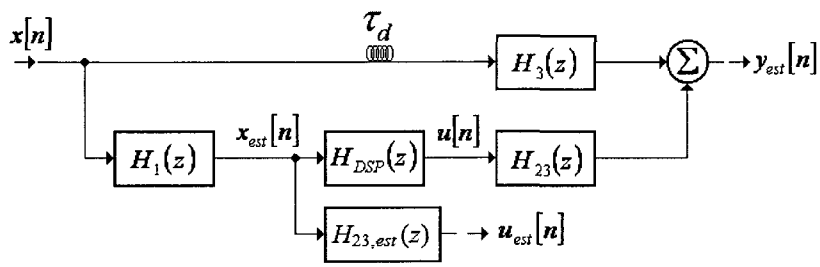

Figure 5.2.6 – Generation of $u_{est}[n]$ for gradient estimation

The second roadblock associated with the model in Fig. 5.2.5 is associated with the assumption that the digital filter coefficients do not change. An adaptive filter is not in steady-state, therefore this assumption is generally not valid, hence $H_{DSP}(z)$ and $H_{23}(z)$ cannot be swapped. This roadblock is not resolved in the FXLMS algorithm. The primary consequence of not resolving this roadblock is that the stability criteria for the standard LMS filter cannot be used to determine the allowable estimation error for $H_{23,est}(z)$. Instead, the allowable estimation error for the FXLMS is based on the location of the eigenvalues in the complex plane of the cross-correlation matrix between $u[n]$ and $u_{est}[n]$ [*4*]. Relating errors in $H_{23,est}(z)$ to the location of these eigenvalues is not straightforward [*4*].

The second roadblock is resolved below to ensure the performance metrics for the standard adaptive linear filter in Fig. 5.2.3 are valid. Fig. 5.2.7 shows two realizations of two cascaded filters, each realization differing in filter sequence. The $H_A(z)$ filter is adaptive, which means its filter coefficients change with time. The $H_B(z)$ filter has fixed coefficients, like those of $H_{23}(z)$. The time varying impulse response coefficients of $H_A(z)$ are denoted by: $a_0[n]$, $a_1[n]$, $a_2[n]$. The time invariant impulse response coefficients of $H_B(z)$ are denoted by: b_0, b_1, b_2, b_3.

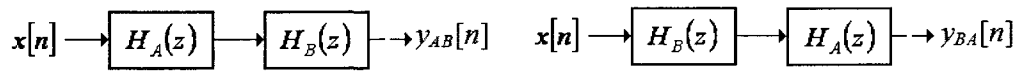

Figure 5.2.7 – Realizations of cascade of $H_A(z)$ and $H_B(z)$

The output signals of the two realizations in Fig. 5.2.7 are:

$$\begin{aligned} y_{AB}[n] = {} & (x[n]a_0[n] + x[n-1]a_1[n] + x[n-2]a_2[n])b_0 \\ & +(x[n-1]a_0[n-1] + x[n-2]a_1[n-1] + x[n-3]a_2[n-1])b_1 \\ & +(x[n-2]a_0[n-2] + x[n-3]a_1[n-2] + x[n-4]a_2[n-2])b_2 \\ & +(x[n-3]a_0[n-3] + x[n-4]a_1[n-3] + x[n-5]a_2[n-3])b_3 \end{aligned} \tag{5.2.1}$$

$$\begin{aligned} y_{BA}[n] = {} & (x[n]b_0 + x[n-1]b_1 + x[n-2]b_2 + x[n-3]b_3)\, a_0[n] \\ & +(x[n-1]b_0 + x[n-2]b_1 + x[n-3]b_2 + x[n-4]b_3)\, a_1[n] \\ & +(x[n-2]b_0 + x[n-3]b_1 + x[n-4]b_2 + x[n-5]b_3)\, a_2[n] \end{aligned} \tag{5.2.2}$$

In the above expressions, $y_{AB}[n]=y_{BA}[n]$ only when $a_i[n]=a_i[n\text{-}1]=a_i[n\text{-}2]=a_i[n\text{-}3]$ for $i=\{0,1,2\}$. This condition means the adaptive filter coefficients cannot change for a fixed period of time for the two realizations to generate the same output. This period of time corresponds to the length of the impulse response of the time invariant filter. Furthermore, after each coefficient update a specific time period must lapse before another block of data is collected for the next update. This time period is equal to the length of the impulse response of $H_B(z)$. The output data is discarded during these lapses in time since immediately after a coefficient update $a_i[n]\neq a_i[n\text{-}1]$. Under these two conditions the model of Fig. 5.2.5 is mathematically equivalent to that in Fig. 5.2.4.

To summarize, the model in Fig. 5.2.5 is valid as long as a block-based adaptive algorithm is used and $u[n]$ is generated from an estimate of $H_{23}(z)$, as shown in Fig. 5.2.6. Block-based algorithms developed for the standard adaptive filter in Fig. 5.2.3 can be employed for the hybrid RF-DSP FF filter using $u_{est}[n]$ as the input signal and $y_{est}[n]$ as

the error signal. Standard block-based adaptive algorithms must be modified to discard a block of data immediately following each coefficient update.

5.2.2 Block LMS Algorithm

A block-based adaptive algorithm is presented in this section and verified with simulation. This algorithm requires an a priori estimate of $h_{23}[n]$, denoted $h_{23,est}[n]$. The method used to obtain $h_{23,est}[n]$ is the topic of the following section.

The block LMS (BLMS) and block RLS are two adaptive block processing algorithms with significant heritage in the field of adaptive linear filtering [***41,40***]. The BLMS has a lower computational complexity, but generally converges slower than a least-squares algorithm. The adaptive algorithm is implemented on an FPGA for the hardware prototype. An FPGA has limited computational resources, therefore the BLMS is chosen for this research. Since the algorithm is based on a modified version of a block implementation of the FXLMS algorithm it is denoted FXBLMS with discard.

The LMS algorithm estimates the gradient of the goal function with respect to the filter coefficients once every clock period. The coefficients are updated at the same rate. The BLMS filter collects N_1 data samples of $u_{est}[n]$ and $y_{est}[n]$ for each coefficient update. The BLMS leaves the filter coefficients fixed for N_1 clock periods. This algorithm uses $u_{est}[n]$ and $y_{est}[n]$ as the input and error signals in Fig. 5.2.3, respectively. The update to the tap coefficients for the BLMS is given by (5.2.3) and (5.2.4). The superscript asterisk (*) indicates complex conjugation.

$$\underline{h}_{DSP}[n+1] = \underline{h}_{DSP}[n] + \frac{\mu_B}{N_1}\underline{\phi}[n] \tag{5.2.3}$$

$$\underline{\phi}[n] = \begin{bmatrix} u^*_{est}[n] & u^*_{est}[n+1] & \cdots & u^*_{est}[n+N_1-1] \\ u^*_{est}[n-1] & u^*_{est}[n] & \cdots & u^*_{est}[n+N_1-2] \\ \vdots & \vdots & \ddots & \vdots \\ u^*_{est}[n-N_{TAPS}+1] & u^*_{est}[n-N_{TAPS}+2] & \cdots & u^*_{est}[n+N_1-N_{TAPS}] \end{bmatrix} \begin{bmatrix} y_{est}[n] \\ y_{est}[n-1] \\ \vdots \\ y_{est}[n+N_1-1] \end{bmatrix} \tag{5.2.4}$$

Each coefficient update for the BLMS algorithm requires $N_1 \times N_{TAPS}$ complex multiplications, and $N_1 + N_{TAPS}$ complex additions, where the number of filter taps is denoted by N_{TAPS}. There are two design variables associated with the BLMS algorithm: N_1 and μ_B. The term N_1 is the number of data samples used for the gradient estimate, and μ_B is the step-size parameter in (5.2.3). These variables are related to the algorithms performance by the eigenvalues of the autocorrelation sequence of $u_{est}[n]$.

The square autocorrelation matrix for $u_{est}[n]$ of size N_{TAPS} has N_{TAPS} eigenvalues. These eigenvalues are denoted by $\lambda_1, \lambda_2 \ldots \lambda_{NTAPS}$. It is assumed that the largest eigenvalues correspond to individual narrowband interferers. The corresponding eigenvectors are tap weights of a filter that tend to only pass the corresponding interferer. The eigenvalue of an interferer is equal to the power of that interferer at the input to the digital filter. The smallest eigenvalues are due to the additive white noise at the filter input. In an interference environment where the interference to noise ratio is very large, the eigenvalues corresponding to the interferers will be much greater than those for the additive white noise. A large interference to noise ratio is expected at the input to a receiver, which is the target application of this algorithm.

Adaptive algorithms use feedback to adapt the filter taps. Feedback provides the potential for an unstable system. The stability constraint for the standard BLMS algorithm is [*44*]:

$$0 < \mu_B < \frac{2}{\lambda_{\max}} \tag{5.2.5}$$

Equation (5.2.5) shows the stability constraint is a function of the maximum eigenvalue of the input signal's autocorrelation sequence. The maximum eigenvalue corresponds approximately to the maximum interferer power in $u_{est}[n]$. A more conservative approach involves replacing λ_{max} with the maximum power in $u_{est}[n]$ [*42*]. This conservative upper bound on μ_B removes any ambiguity associated with the assumption that the largest interferer power is equal to λ_{max}. The effect of errors in $h_{23,est}[n]$ on stability is derived in Section 5.3.1.

The mean-square error convergence time constant along the k^{th} eigenvector is [*44*]:

$$\tau_{k,MSE} = \frac{N_1}{2\mu_B \lambda_k} \tag{5.2.6}$$

Long convergence times are associated with small eigenvalues, which correspond to weak interferers or white noise. Strong interferers have relatively large eigenvalues, so the convergence at these frequencies will be faster. Fast convergence along the associated eigenvectors requires a small N_1 and large μ_B. A large interference to noise ratio results in a large eigenvalue spread, which means that long convergence times will occur over the smallest eigenvectors. The convergence rate along these eigenvectors is insignificant to the interference attenuation of the total filter since they correspond to additive noise. The convergence time also depends on the initialization of the filter coefficients [*42*].

The standard BLMS algorithm is modified for the hybrid RF-DSP FF filter. The modification involves discarding L_{23} values following a filter coefficient update. The value of L_{23} is the effective length of the impulse response of $h_{23}[n]$. This modification increases the convergence times, so the new convergence time constants are:

$$\tau_{k,MSE} = \frac{N_1}{2\mu_B \lambda_k} \cdot \frac{N_1 + L_{23}}{N_1} = \frac{N_1 + L_{23}}{2\mu_B \lambda_k} \tag{5.2.7}$$

The numerator of (5.2.7) shows that the time constants cannot be significantly reduced when $N_1 < L_{23}$, which is contrary to the standard BLMS.

The LMS algorithm estimates the gradient for each tap update. The accuracy of this estimate depends on how many data samples are used for the estimate. Misadjustment is the ratio of the mean square of the error signal of the adaptive filter in steady-state to that of the Wiener filter. At steady-state, misadjustment is due to the noisy estimates of the gradient. The misadjustment of the BLMS algorithm is [***44***]:

$$\mathrm{M} = \frac{\mu_B}{2N_1} tr(R) \tag{5.2.8}$$

The trace operator (*tr*) of the input signal's correlation matrix (R) is the sum of the eigenvalues. The largest eigenvalues dominate the misadjustment, as evident by (5.2.8). The contribution to the misadjustment from the k^{th} eigenvector with eigenvalue σ_k, is $\mu_B \sigma_k / 2N_1$. A small misadjustment requires a large N_1 and a small μ_B. These conditions are in opposition to those for fast convergence; therefore demand a trade-off.

The adaptive algorithms are implemented in an FPGA for the hardware prototype; therefore resource usage must factor into the trade-off. At every time instant a sample

from the error signal, $y_{est}[n]$ is multiplied by a vector of current and past input signals, $\{u_{est}[n], u_{est}[n\text{-}1],\ldots,u_{est}[n\text{-}N_{TAPS}\text{+}1]\}^*$. This process requires N_{TAPS} complex multiplications every time sample. These products are accumulated for N_1 samples between updates. In terms of FPGA resources, there are N_{TAPS} complex multiplications every clock period, which is independent of N_1, therefore computational complexity is not a factor in choosing N_1.

In the literature pertinent to the BLMS algorithm [*42*,*44*], it is stated that if $N_1 > N_{TAPS}$, then redundant operations are performed as the digital filter cannot use all the input information. These comments are directed at the BLMS in a different implementation, the fast BLMS. In the fast BLMS the serialized input signals are converted to the frequency domain before being processed by the N_{TAPS} filter taps in the frequency domain. This issue of redundant operations is not relevant to the implementation in this work as the input signal is not transformed into the frequency domain. As N_1 increases beyond N_{TAPS} the misadjustment continues to improve, but comes at the cost of an increased convergence time.

The performance metrics that factor into the choice of N_1 and μ_B are stability, convergence rate, and misadjustment. These design variables are varied in the hardware prototype, and the performance metrics are reported in Chapter 7.

5.2.3 Verification with Maximum Precision Simulation

The modified BLMS algorithm is implemented in Simulink. The digital filter has 64 taps in the simulations (N_{TAPS}=64). The taps are initialized to 0. A 64 tap filter is also used to

represent $h_{23}[n]$. The algorithm used to obtain the 64 tap values of $h_{23,est}[n]$ is the topic of the next section. Finite precision effects are deferred to Section 5.4.

The hybrid RF-DSP FF filter is simulated in the baseband domain. This approach avoids the small time-steps required for RF signals. The block diagram used in Simulink is shown in Fig. 5.2.8. The noise figures of the active analog components is set to 0 dB, and the resolution of the digital components are set to maximum machine precision. These settings isolate the results to measure the operation of the BLMS algorithm without thermal or quantization noise effects. The time-step of the simulation is set to 2.5 nsec, but the time-step of the digital components is set to 5 nsec. The larger time-step in the digital components corresponds to a 200 MHz digital clock. The smaller step-size for the analog components permits the T-BSF to have a relatively narrow notch.

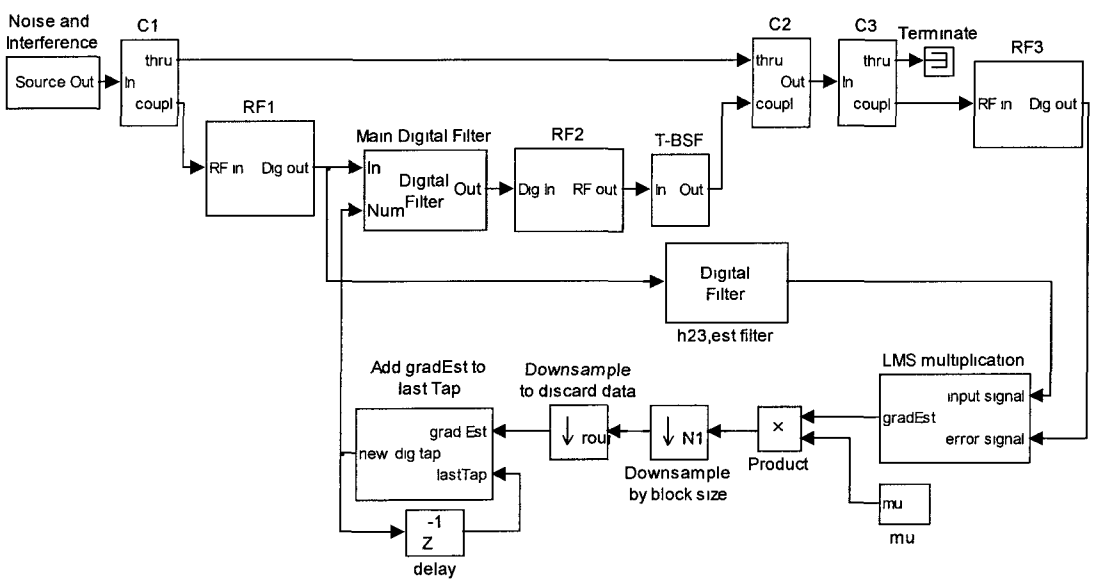

Figure 5.2.8 - Simulink model of hybrid RF-DSP FF filter with $h_{23,est}[n]$ known a priori

The blocks in Fig. 5.2.8 that represent the hybrid RF-DSP FF filter in Fig. 5.2.1 excluding the algorithm processor are: *C1*, *RF1*, *Main Digital Filter*, *RF2*, *T-BSF*, *C2*, *C3* and *RF3*. These blocks are the same as those from the noise analysis in the previous chapter with the addition of *RF3*.

The remainder of the blocks in Fig. 5.2.8 form the algorithm processor, with the exception of the *Noise and Interference* source. The *h23,est filter* block in Fig. 5.2.8 generates the $u_{est}[n]$ samples, as shown in Fig. 5.2.6. These samples are processed by the *LMS multiplication* block along with the $y_{est}[n]$ samples from the *RF3* block. This processing is the matrix multiplication in (5.2.4). The vector product from this matrix multiplication is the gradient estimate averaged over N_{TAPS} values. In this simulation the vector product is calculated every time step, but the BLMS algorithm only uses one of these results every N_{TAPS} time steps. The remainder N_{TAPS}-1 values are discarded by the first down-sample block. The second down-sample block is used to discard some of the gradient estimates immediately after a tap update.

The *Product* and *Add gradEst to last Tap* blocks perform the filter tap update from (5.2.3). The output of the latter block is the updated filter taps, which are then sent to the main digital filter.

The *Noise and Interference* block in Fig. 5.2.8 generates an interference scenario. The input PSD to the filter used in these simulations is shown in Fig. 5.2.9 with the 'PSD_{in}' curve. The input PSD represents multiple interferers, with various bandwidths and power levels. This PSD is shown at baseband since the Simulink simulations are performed in this domain.

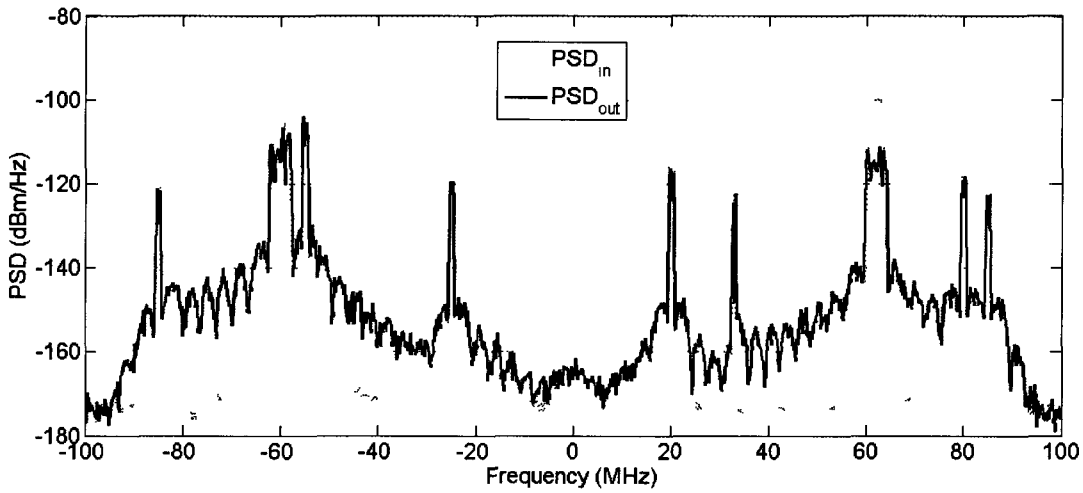

Figure 5.2.9 – PSD at input and output of hybrid RF-DSP FF filter

The algorithm is run with N_1=64, and 64 data points discarded per update. The magnitude response of the hybrid RF-DSP FF filter at two different time instants is shown in Fig. 5.2.10. These responses, denoted t_1 and t_2, are recorded during the transient phase of the filter's adaptation. The t_1 response occurs before t_2. The solid dots on the -20 dB grid line indicate the location of the interferers. Almost all of the dots have a notch going through them. The exception is the wideband notch near –60 MHz, where a notch response forms but does not reach -20 dB This figure demonstrates proper operation of the algorithm. The magnitude responses at 20 and 62 MHz are zoomed in Fig. 5 2.11. These figures show the improvement of the notch width as the filter adapts in time.

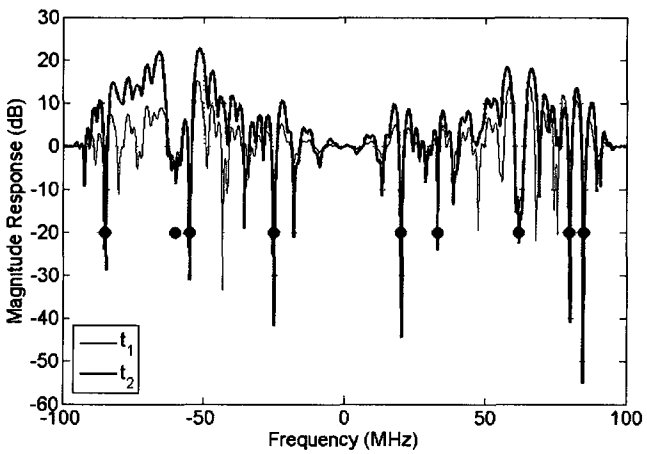

Figure 5.2.10 - Magnitude responses of hybrid RF-DSP FF filter during adaptation

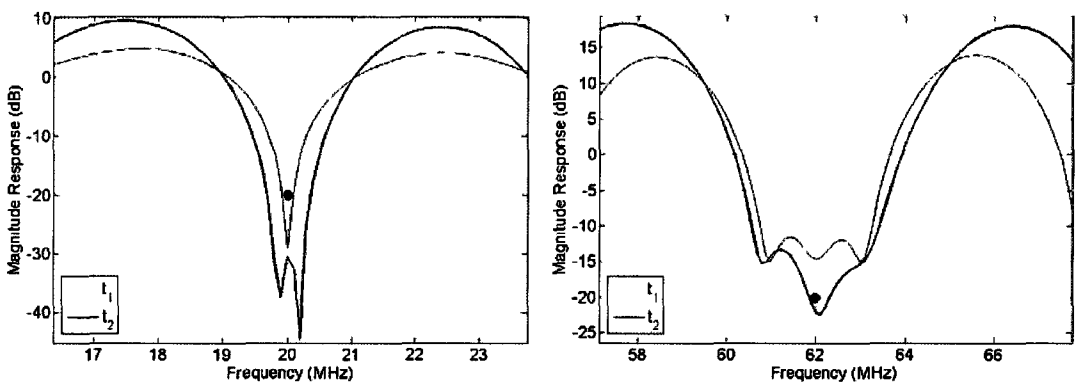

Figure 5.2.11 - Zoomed responses from Fig. 5.2.10

Starting at time t_2, the input and output samples are recorded and then processed to obtain the PSD at the input and output of the hybrid RF-DSP FF filter. The results are shown in Fig. 5.2.9. This response shows significant attenuation of the interferers at the expense of an increase in the noise floor away from the receiver's passband (0 MHz).

The above results show successful operation of the FXBLMS with discard and $h_{23,est}[n]$ known a priori. Measurements of the hardware prototype are used to relate the

algorithm's performance metrics to the input interference environment. The results are in Chapter 7. The problem of obtaining $h_{23,est}[n]$ is solved next.

5.3 Estimating $h_{23}[n]$

Approximating the impulse response of $h_{23}[n]$ is a channel estimation problem. The channel starts at the output of $H_{DSP}(z)$ and ends at the output of $H_3(z)$. Conventional linear adaptive channel estimation involves sending a training sequence through the channel and correlating the output signal with a copy of the training sequence. The training sequence only needs to be run when $h_{23,est}[n]$ needs to be updated. An update is needed when component parameters drift in $H_2(z)$ or $H_3(z)$, or the filter is reconfigured in frequency or bandwidth. The training sequence should be uncorrelated with $x_{est}[n]$ to ensure only the $h_{23}[n]$ channel is estimated. Fig. 5.3.1 shows how the training sequence, $x_{trg}[n]$, is incorporated into the hybrid RF-DSP FF filter

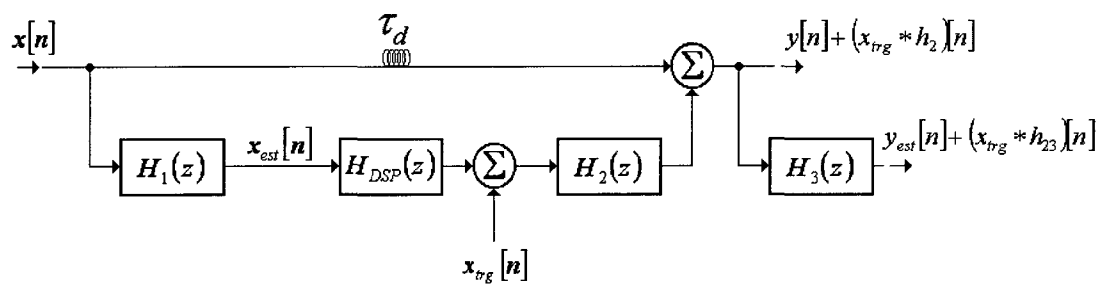

Figure 5.3.1 - Model incorporating training sequence

The BLMS algorithm is chosen for the estimation of $h_{23}[n]$ due to its relatively low computational complexity. Furthermore, hardware resources can be shared between this adaptive algorithm and that for the digital filter tap coefficients, since it also employs a version of the BLMS algorithm. The filter update equation is the same as (5.2.3), except

$u[n]$ is replaced by $x_{trg}[n]$, and $y_{est}[n]$ is replaced by $y_{est}[n]+(x_{trg}*h_{23})[n]$. Discarding data after filter updates is not required for this BLMS algorithm as the filter architecture does not deviate from the conventional form.

The design variables relevant to this algorithm relate to the design of the training sequence and the number of filter taps used in $h_{23,est}[n]$. The training sequence will appear as additive noise to the receiver. The PSD of the training sequence can be notched over the receiver's passband to ensure negligible degradation to the receiver's sensitivity performance. This training sequence characteristic is possible since an accurate estimate of $h_{23}[n]$ is not required in the passband since FF attenuation does not occur there. Signals other than the training sequence appear as noise to the $h_{23}[n]$ channel estimation algorithm. These signals include the attenuated interferers in $y_{est}[n]$, which do not have flat PSD's, therefore cannot be treated as white noise. The frequency bands occupied by these interferers correspond to the frequencies where the most accurate estimate of $h_{23}[n]$ is required. Errors in the estimate of $h_{23}[n]$ affect the performance metrics of the main digital filter's LMS algorithm, including stability. These effects are derived next.

5.3.1 Effect of Errors in $h_{23,est}[n]$

Since $h_{23,est}[n]$ does not exactly equal $h_{23}[n]$, then $u_{est}[n]$ will not exactly equal $u[n]$. The performance metrics of the LMS algorithm are based on the input signal $u[n]$ and not $u_{est}[n]$, but the latter is used by the FXBLMS with discard algorithm.

One effect of magnitude mismatch in $h_{23,est}[n]$ is to cause pre-emphasis or de-emphasis of incoming interference signals. For example, if $h_{23,est}[n]$ is 3 dB larger than $h_{23}[n]$ over a frequency band occupied by an interferer, then the FXBLMS algorithm will treat the

interferer like it is 3 dB larger and attenuate it more than is required to produce a white noise output. These effects shift the convergence point in the goal function of the LMS algorithm. In an RF environment with a large interference to noise ratio, the convergence point shift is negligible when the magnitude error is only a few decibels at interferer frequencies. The magnitude error can be even larger at frequencies occupied only by noise.

A more degrading effect of mismatch between $h_{23,est}[n]$ and $h_{23}[n]$ is to bias the gradient estimate during adaptation. The following analysis switches from a block to a sample based adaptive algorithm. This switch does not impact the solution, but allows for more concise mathematical expressions. The expected value of the estimated gradient for this modified LMS algorithm is:

$$E[\nabla J_{est}] = E[\underline{u}_{est}(n) y^*_{est}(n)] \tag{5.3.1}$$

The estimate of the input signal can be decomposed into the desired term and an error term:

$$u_{est}[n] = u[n] + u_{err}[n] \tag{5.3.2}$$

Similarly, the estimate of the impulse response of $h_{23}[n]$ can be decomposed:

$$h_{23,est}[n] = h_{23}[n] + h_{23,err}[n] \tag{5.3.3}$$

The estimate in (5.3) is expanded with (5.3.2):

$$E[\nabla J_{est}] = E[\underline{u}(n) y^*_{est}(n)] + E[\underline{u}_{err}(n) y^*_{est}(n)] \tag{5.3.4}$$

The first term of the summation in (5.3.4) is the desired expected value of the gradient estimate. The second term is the bias of the gradient estimate due to errors in $u_{est}[n]$. The remainder of this formulation is performed in the z-domain since this domain offers a more intuitive comparison of the desired and bias terms in (5.3.4). The hybrid RF-DSP FF filter can be modelled in the z-domain since it is a discrete-time filter.

Both of the expected value expressions in (5.3.4) are cross-correlation sequences, which contain a finite number of terms equal to the number of taps of the main digital filter. A finite number of terms in these expectations results in truncated versions of the cross-correlation sequences. The effects of this truncation decrease with increasing filter taps. The truncation effects are ignored, but can become significant for a digital filter with a small number of taps.

The cross-correlation terms in (5.3.4) are converted to the cross-spectral density terms in (5.3.5). In this equation $H_{FF}(z)$ is the z-domain transfer function of the entire hybrid RF-DSP FF filter.

$$\begin{aligned} ZT\{E[\nabla J_{est}]\} &= S_{u,yest}(z) + S_{uerr,yest}(z) \\ &= S_{uu}(z)\frac{H_{FF}^{*}(z)}{H_1^{*}(z)H_{23}^{*}(z)} + S_{uerr,uerr}(z)\frac{H_{FF}^{*}(z)}{H_1^{*}(z)H_{23,err}^{*}(z)} \\ &= S_{xx}(z)H_1(z)H_{23}(z)H_{FF}^{*}(z) + S_{xx}(z)H_1(z)H_{23,err}(z)H_{FF}^{*}(z) \end{aligned} \tag{5.3.5}$$

This equation represents the expected value of the gradient estimate in the z-domain. The gradient of the fast BLMS algorithm is also represented in the frequency domain. The first term of the summation in (5.3.5) is the desired term for the z-domain gradient. The second term is the bias due to the error in $h_{23,est}[n]$. The complex ratio of the bias to the desired component in the z-domain is:

$$\gamma_{grad} = \frac{H_{23,err}(z)}{H_{23}(z)} \tag{5.3.6}$$

The ratio in (5.3.6) is defined to permit the expected value of the gradient estimate in the z-domain to be expressed as:

$$ZT\{E[\nabla J_{est}]\} = ZT\{E[\nabla J]\} \times (1 + \gamma_{grad}) \tag{5.3.7}$$

The weight update equation of the adaptive algorithm is written in the z-domain with the estimated gradient substituted for the actual gradient:

$$\underline{H}_{DSP}(k+1) = \underline{H}_{DSP}(k) - \frac{\mu_B}{N_1} \times ZT\{E[\nabla J]\} \times (1 + \gamma_{grad}) \tag{5.3.8}$$

The H_{DSP} terms are the main filter's tap values in the z-domain. The step-size parameter of the weight update in (5.3.8) is effectively changed to:

$$\mu_{B,eff} = \mu_B (1 + \gamma_{grad}) \tag{5.3.9}$$

This effective step-size parameter is frequency dependent due to the frequency dependent behaviour of γ_{grad}, defined in (5.3.6). If the value of $\mu_{B,eff}$ is different than μ_B at frequencies occupied by interferers, then the stability, convergence time and misadjustment along the corresponding eigenvectors are different than if the step-size was just μ_B. The standard LMS algorithm uses real values for the step-size parameter, but here the effective value is complex, which requires re-working of the performance metrics in (5.2.5), (5.2.7) and (5.2.8).

Complex issues related to adaptive filters are typically illustrated with a two tap filter example such that the goal function can be plotted in 2-D. The adaptation of a two tap

filter with real valued taps is shown in Fig. 5.3.2 along with contours of the goal function. The initial value of both filter taps is zero. The incoming signal to the entire FF filter is the summation of a sinusoid and wideband noise, hence the ellipsoidal nature of the goal function contours. A 16 tap filter is placed before the adaptive filter, denoted $h_{23}[n]$. The input signal for the LMS algorithm is obtained by sending the incoming signal through another filter, denoted $h_{23,est}[n]$. This arrangement resembles the setup in Fig. 5.2.6. The adaptation of the coefficients is shown for an increasing estimation error in $h_{23,est}[n]$, and constant μ_B. The curves in Fig. 5.3.2 show that as the estimation error increases the trajectories change, and the convergence time changes. If the error is too large, then the filter taps become unstable.

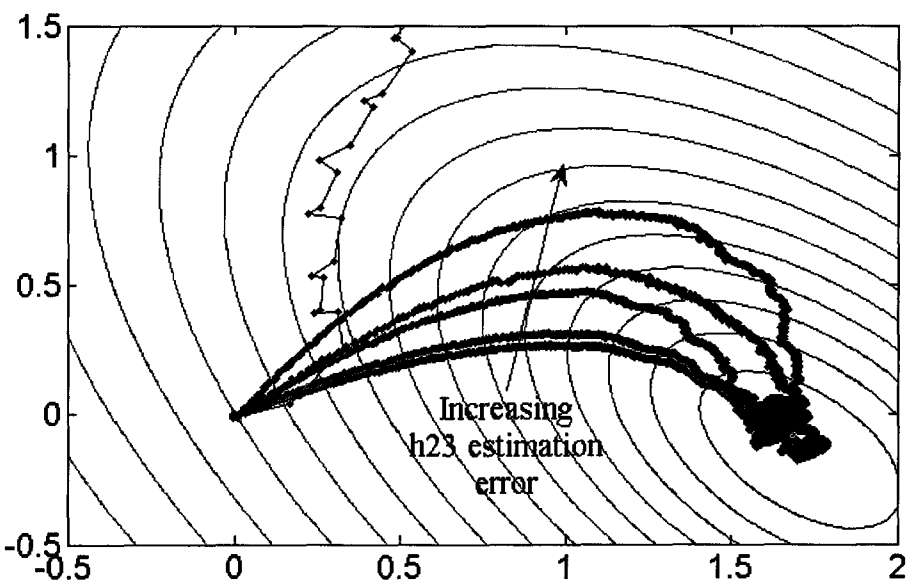

Figure 5.3.2 - Adaptation of a two tap filter where gradient estimate contains a bias

The remainder of this section is dedicated to relating the stability constraint of the main filter's adaptive algorithm and errors in $H_{23,est}(z)$. A more general version of the LMS small step-size stability constraint is [*42*]:

$$\left|1-\mu_B\lambda_i\right|<1 \tag{5.3.10}$$

To determine the effect of γ_{grad} on stability, the effective step-size from (5.3.9) is substituted into this LMS stability constraint:

$$\left|\frac{1}{\mu_B\lambda_i}-1-\gamma_{grad}\right|<\frac{1}{\mu_B\lambda_i} \tag{5.3.11}$$

In the complex plane (5.3.11) defines a circle with radius $1/\mu_B\lambda_i$ offset from the origin by a distance of $1/\mu_B\lambda_i - 1$ along the real axis. The value of this offset term can be positive or negative. The area enclosed by each circle represents a feasibility region for γ_{grad} to ensure stability. Several feasibility circles are plotted in Fig. 5.3.3, each for a different value of the original step-size, μ_B.

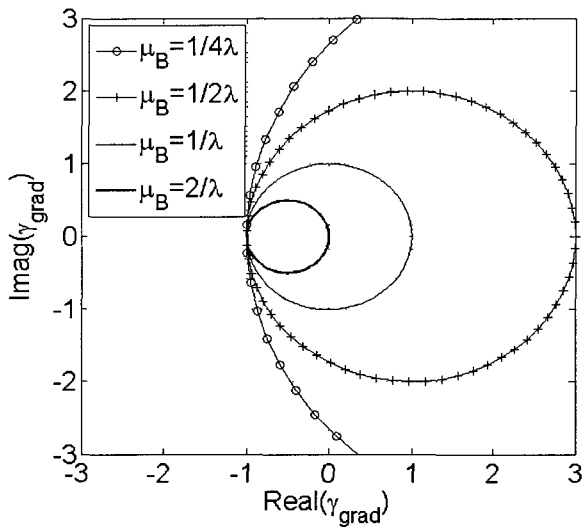

Figure 5.3.3 - Feasibility region of γ_{grad} for various values of μ_B

In the above figure the magnitude of γ_{grad} is the radial distance from the origin, and the phase of γ_{grad} is the angle from the positive real axis. Each circle in Fig. 5.3.3 encloses a region where the FXBLMS with discard algorithm will be stable, for the specified μ_B, along the eigenvector corresponding to λ. The region constrains the values of γ_{grad} to ensure stability.

The feasibility circles in Fig. 5.3.3 become larger as the value of μ_B decreases, which permits larger estimation errors in $h_{23,est}[n]$. In this biased LMS algorithm, the convergence time is still inversely proportional to μ_B, so decreasing μ_B to an arbitrarily small value is not the optimal approach. Furthermore, regardless of how large the feasibility circle is, there is only a small feasibility region when the angle of γ_{grad} is greater than $\pi/2$ and less than $3\pi/2$. Permitting large errors in $h_{23,est}[n]$ may land the angle of γ_{grad} in this range within the band of an interferer and cause the algorithm to become unstable.

Each eigenvector has its own feasibility region for γ_{grad}, for a fixed value of μ_B. In a relatively narrowband interference environment with a large interference to noise ratio, each eigenvector is fairly restricted to a frequency band occupied by an interferer. The associated eigenvalue is the interferer power, and γ_{grad} is a measure of the estimation error in $h_{23,est}[n]$ within that band. Eigenvectors with large eigenvalues correspond to small circles within Fig. 5.3.3, therefore at these frequencies the estimation error in $h_{23,est}[n]$ must be relatively small. Eigenvectors corresponding to wideband noise have small eigenvalues, which have very large feasibility circles in Fig. 5.3.3. The estimation error in $h_{23,est}[n]$ is less critical at frequency bands not occupied by interferers.

The results in this section are important because they determine how accurately the $h_{23,est}[n]$ filter must model the $h_{23}[n]$ channel. The filter's design parameters that affect this accuracy are the number of filter taps, number of samples per update, and the step-size parameter of the $h_{23}[n]$ channel estimation algorithm.

5.4 Algorithm for Tunable BSF Control

The tunable BSF is adaptive for two reasons. First, the components within the BSF may drift, thereby degrading the depth of the transfer function notch. Second, the frequency agility characteristic of the hybrid RF-DSP FF filter requires the T-BSF's notch to be frequency agile. The BSF used in the hardware prototype consists of two paths in FF, with one path containing a vector modulator and the other a delay line. The block diagram in Fig. 5.4.1 shows this BSF configuration within the hybrid RF-DSP FF filter.

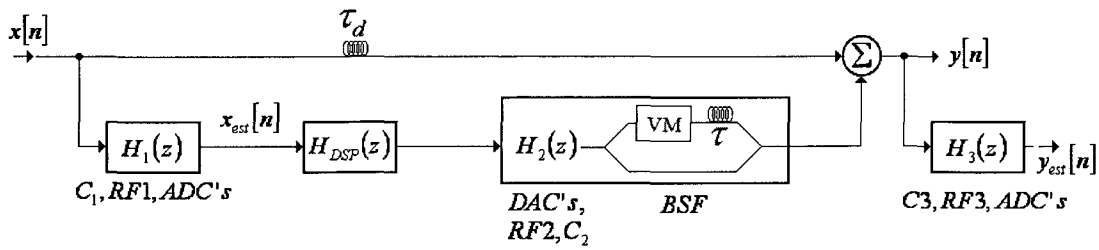

Figure 5.4.1 – BSF algorithm implementation

The notch is required in the receiver's passband, but signals in this band are not easily detected in the presence of strong interferers. The purpose of a BTS receiver is to receive signals in this passband, so the digital output of this receiver can be used as a feedback signal to adapt the notch. This feedback path would have significant latency, but drifting of the BSF components is very slow; therefore can be afforded.

A second method was used to tune the BSF without requiring a BTS receiver. This method has the FPGA inject a double sideband tone into the receiver's passband during adaptation of the BSF. The two vector modulator baseband inputs are varied over their entire ranges and the power of the tone is recorded at the output of the ADC's in the feedback path. The vector modulator inputs that result in the minimum power at these ADC's are locked in and the tone turned off. This algorithm only needs to run when the frequency is changed, or the BSF notch components significantly drift. Successful tuning of the BSF is demonstrated in the following chapter.

5.5 Finite Precision Effects

The simulation results in the previous sections used maximum floating point precision in Simulink. Finite precision effects were incorporated into the noise analysis in Chapter 4. Another effect of finite precision is to bias the tap weights of an adaptive filter. This bias

is most problematic for adaptive filters that adapt over a long period of time [*42*]. Furthermore, finite precision limits the maximum tap values.

Finite precision is incorporated into the Simulink simulations by using the fixed-point toolbox. The tap coefficients for $h_{23,est}[n]$ are known a priori. The evolution of the 32 complex taps of the main digital filter is shown in Fig. 5.5.1. The bias due to finite precision is apparent in these figures by the slowly increasing curves from between iterations 1 and 6×10^5. These signed taps are represented with 15 bits, which can take on a maximum value of 2^{14}-1 (shown in each of the figures as a thick horizontal line). The bias due to finite precision results in a consistent increase of some of the tap values. Eventually one of these taps hits the saturation level, and the filter performance becomes unpredictable. The coefficients appear to somewhat settle after iteration 6×10^5, but the bias again pushes them towards the saturation level just before iteration 12×10^5. Fortunately this bias problem is well known and the solution is to modify the LMS algorithm by using the leaky LMS algorithm [*42*].

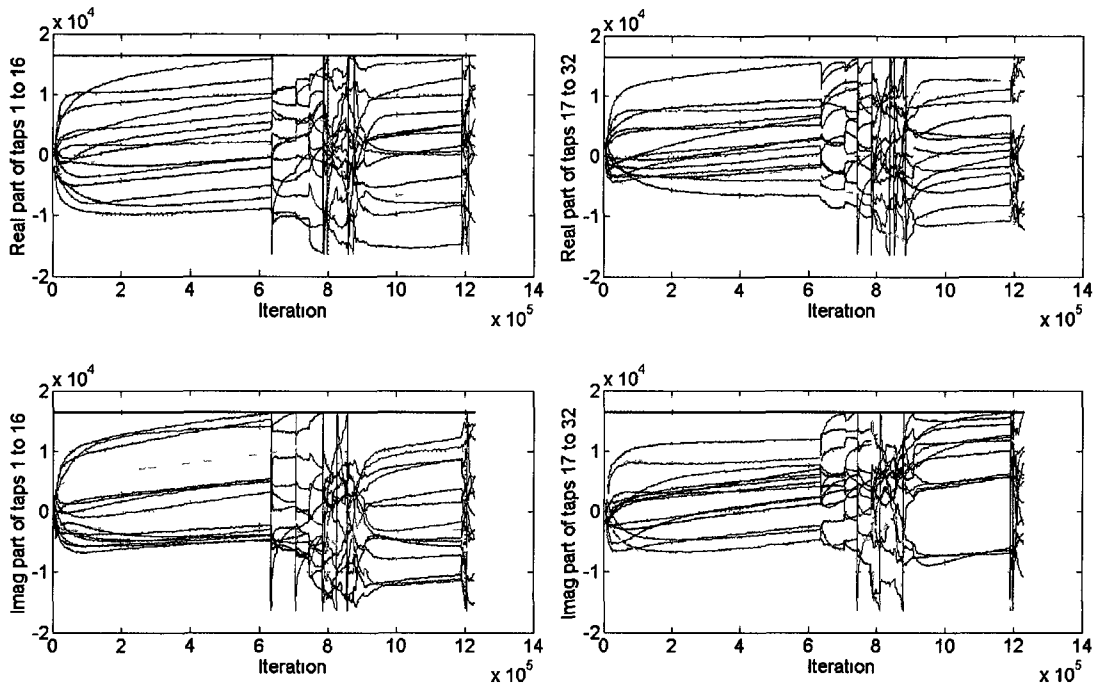

Figure 5.5.1 – Behaviour of 32 complex filter taps with finite precision

The leaky LMS algorithm modifies (5.2.3) to include a leakage term, α.

$$\underline{h}_{DSP}[n+1] = (1-\mu_B\alpha)\underline{h}_{DSP}[n] + \frac{\mu_B}{N_1}\underline{\phi}[n] \tag{5.5.1}$$

Whereas the standard LMS algorithm adapts to minimize the mean square error of the error signal in Fig. 5.2.3, the leaky LMS algorithm minimizes the goal function:

$$J = \alpha\left\| \underline{h}_{DSP}[n]\right\|^2 + E\left\{\left|y_{est}[n]\right|^2\right\} \tag{5.5.2}$$

The goal function of the standard LMS algorithm only consists of the second term in the summation of (5.5.2). With the leaky LMS algorithm the amplitude of the tap weights is not allowed to increase indefinitely. The drawback of the leaky algorithm is a degraded misadjustment [*42*].

The adaptive algorithm is modified to include a leakage factor. The adaptation of the 32 complex taps is shown in Fig. 5.5.2. In these figures there is no consistent increase in the tap values. The saturation level of the taps, $\pm 2^{14}$, corresponds to the limits of the vertical axis. These results show that incorporating the modification from (5.5.1) removes the undesirable behaviour illustrated in Fig. 5.5.1 from the FXBLMS with discard algorithm.

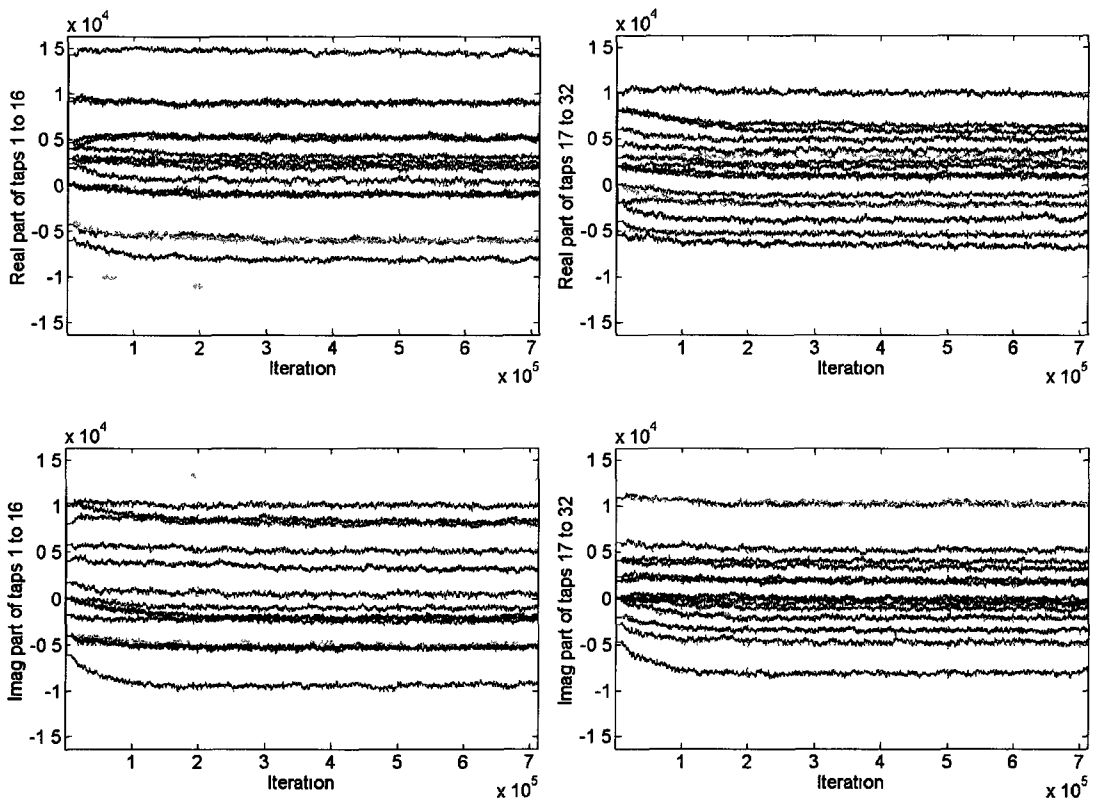

Figure 5.5.2 – Convergence of 32 complex filter taps with leaky LMS algorithm with μ=15, α=2

The PSD of the input and output signals of the hybrid RF-DSP FF filter is shown in Fig. 5.5.3. The data for the PSD's is taken after convergence, corresponding approximately to iteration 2×10^5 in Fig. 5.5.2.

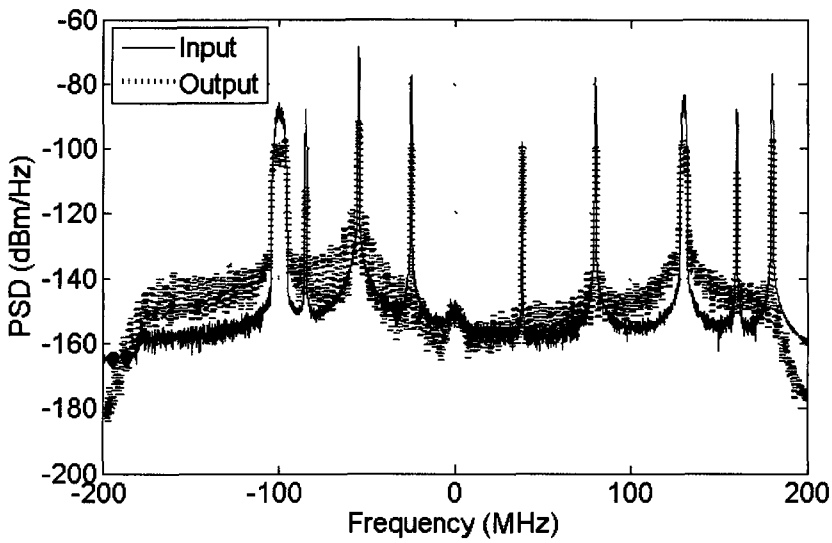

Figure 5.5.3 - PSD of input and output signals of hybrid RF-DSP FF filter

Narrowband and wideband attenuation of multiple interferers can be seen in Fig. 5.5.3, with finite precision incorporated into the simulations. The performance of the algorithms is deferred to the following chapters since the hardware prototype offers a much faster means of processing the data.

5.6 Algorithm Development Summary

This chapter has introduced the algorithms for the hybrid RF-DSP FF filter. Adaptive algorithms are required to handle a dynamic interference environment and drift of the RF and analog components. The algorithm that adapts the main digital filter is based on a block implementation of the FXLMS algorithm. Data discard is incorporated into this algorithm so the performance metrics from the standard block LMS algorithm can be used; namely stability, convergence time and misadjustment. A leakage factor is incorporated to compensate for finite precision effects. This algorithm requires

knowledge of $h_{23}[n]$. A block LMS algorithm is used to estimate this channel. Errors in this channel estimation affect the stability of the main adaptive algorithm, so a new stability criterion was derived. Lastly, a 2-D search algorithm used to tune the BSF was reviewed. The implementations of the three algorithms are shown in Fig. 5.6.1. These three algorithms are implemented on the FPGA and successful operation is reported in Chapter 6.

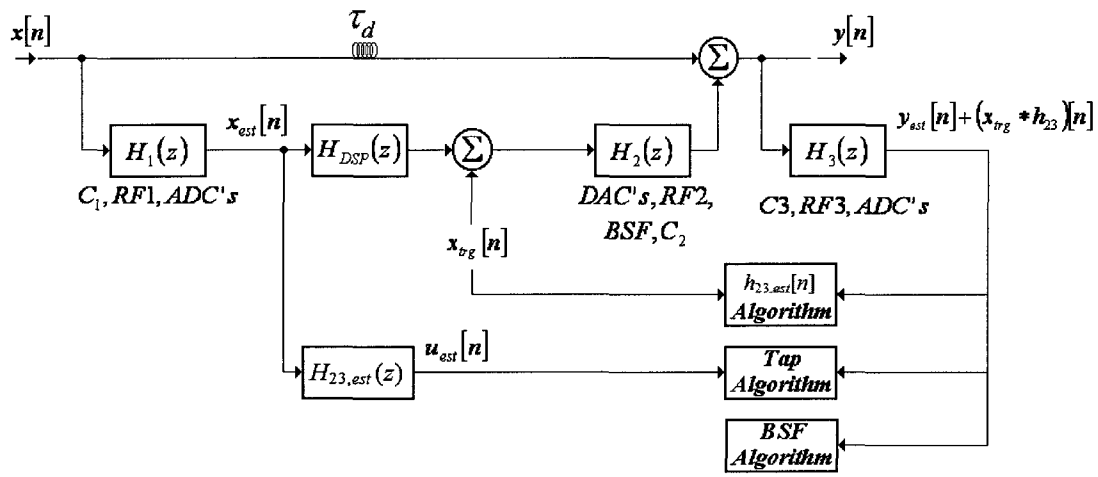

Figure 5.6.1 – Input signals to all three algorithms

6 Hardware Prototype

The analyses and simulation results in the two previous chapters showcase the potential for the hybrid RF-DSP FF filter in the RF front-end of a BTS receiver. This chapter describes the hardware prototype and demonstrates proper operation of the algorithms. The next chapter reports the performance of the hardware prototype for various input interference scenarios and filter realizations.

6.1 Description of Components

The hardware prototype consists of evaluation boards and discrete connectorized components. A block diagram of the hardware prototype is shown in Fig. 6.1.1, and a photograph in Fig. 6.1.2. Torque wrenches are standing in the photograph at the RF input and output of the hybrid RF-DSP FF filter. The evaluation boards are summarized in Table 6.1, and the connectorized components in Table 6.2. In the figure and tables the prefixes *C, G* and *D* denote couplers, amplifiers, and delay lines, respectively. The vector modulator within the BSF is denoted by VMOD, and the 3-dB splitter by SPLIT.
The LPF's adjacent to the primary ADC's and DAC's are not present in the hardware prototype since they significantly increase the group delay of the lower FF path. The additional group delay is shown in Section 6.2.7.

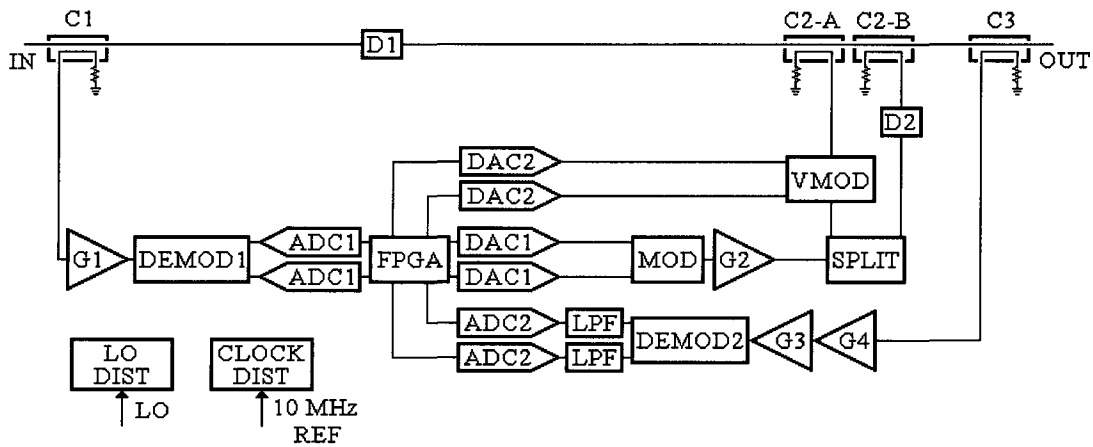

Figure 6.1.1 - Hardware prototype block diagram

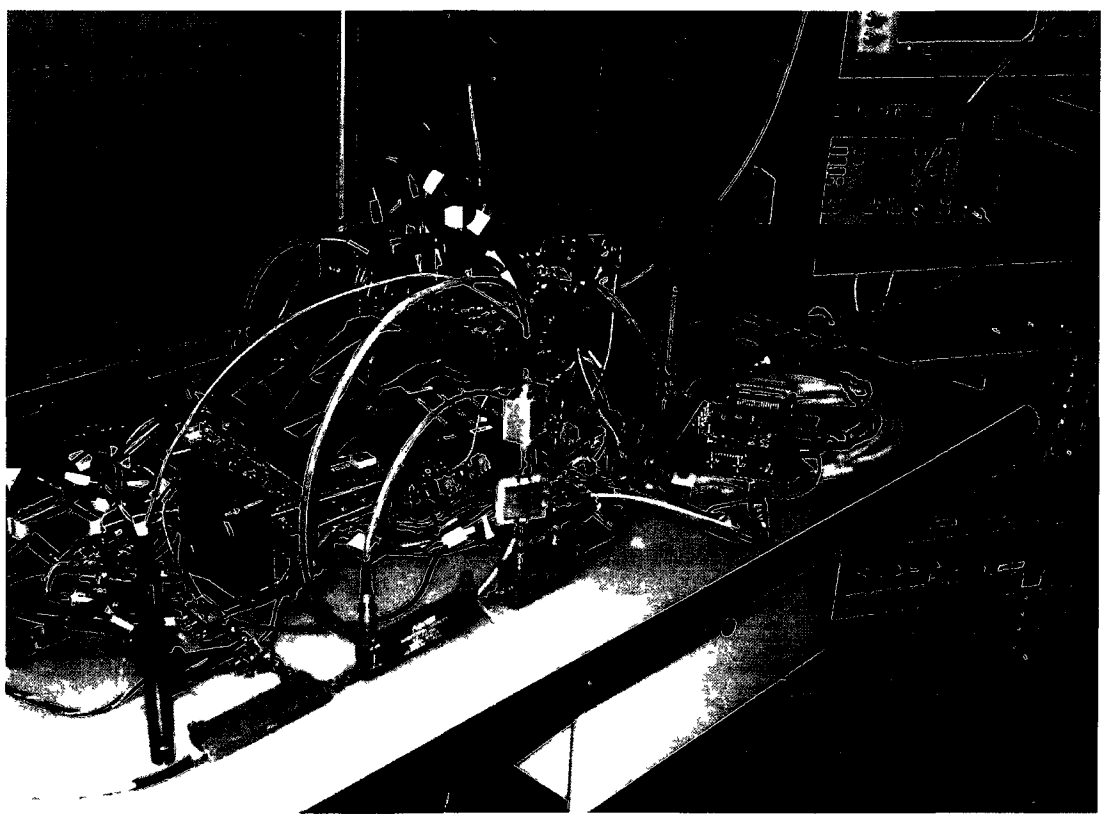

Figure 6.1.2 – Hardware prototype photograph

Table 6.1 - Evaluation boards for hardware prototype

Label	Part Number	Company	Relevant Specifications
ADC1	ADS5474	Texas Instruments	14 bit, 400 MSPS
ADC2	AD9626	Analog Devices	12 bit, 250 MSPS
DAC1	AD9744	Analog Devices	14 bit, 210 MSPS
FPGA	DE3-340	Terasic/Altera	Stratix III – 340 FPGA
DEMOD1	ADL5832	Analog Devices	RF: 0.7-2.7 GHz BB: DC-370 MHz
DEMOD2	HMC597LP4	Hittite	RF: 0.1-4 GHz, BB: DC-600 MHz
MOD	HMC697LP4	Hittite	RF: 0.45-4 GHz BB: DC-700 MHz
CLOCK DIST	AD9520-1	Analog Devices	PLL with clock dividers and multiple outputs
DAC2	AD5392	Analog Devices	14 bit, 8 channel
VMOD	HMC630LP3	Hittite	RF: 700-1000 MHz Max gain: -10 dB

Table 6.2 - Connectorized components for hardware prototype with nominal values

Label	Part Number	Company	Specifications
LPF	SLP-90+	Mini-Circuits	Cut-off: 90 MHz
G1 & G3	ZX60-33LN-S+	Mini-Circuits	50-3000 MHz Gain: 15 – 19 dB Noise figure: 1.03 dB
G2 & G4	ZX60-2411BM-S+	Mini-Circuits	800-2400 MHz Gain: 11.5 dB Noise Figure: 4.3 dB
LO DIST	2089-6406-00	MA/COM	4 way splitter 500-2000 MHz
C1	2025-6002-10	MA/COM	10 dB coupler 500-2000 MHz Thru loss: 0.81 dB
C3	4242-10	Narda	10 dB coupler 500-2000 MHz Thru loss: 0.9 dB
C2-A	4242-20	Narda	20 dB coupler 500-2000 MHz Thru loss: 0.4 dB
C2-B	CS10-1.500V	MECA	10 dB coupler 800-2200 MHz Thru loss: 0.3 dB
D1	N/A	N/A	Coaxial delay line 23 cm
D2	RG400	N/A	Coaxial delay line 110 cm
SPLIT	ZFSC-2-2500-S+	Mini-Circuits	10-2500 MHz

6.2 Bottom FF Path Characterization

The analyses in Chapter 4 result in expressions for the attenuation and noise performance based on component values. The derived noise expression is verified with measurements in Chapter 7. This verification is performed with the hardware prototype operating at 900 MHz with a 5 MHz RF passband.

Component specifications are listed in the component data sheets, but these values depend on the actual loading conditions. The demodulator and modulator in the bottom FF path do not have the nominal loading conditions, therefore their parameters are measured.

The magnitude response of a training sequence is measured at different locations around the hybrid RF-DSP FF filter to determine the magnitude responses under the actual loading conditions in the hardware prototype. Also, the output noise PSD is measured at different locations with a thermal noise input to estimate noise factors. The RF measurements are performed with the PXA signal analyzer from Agilent. The component value results are summarized in Section 6.2.9.

6.2.1 Magnitude Response *RF*1-ADC1 Chain

The magnitude response of the *RF*1 chain and ADC1's is measured by sending a band-limited signal into the *RF*1 chain input. The *RF*1 chain consists of the G1 amplifier and DEMOD1. A signal generator is used to create the input signal which has a relatively flat PSD over the 5 MHz RF passband at 900 MHz. The configuration is shown in Fig. 6.2.1. A signal analyzer is used to measure how much power is sent into the *RF*1 chain within the passband.

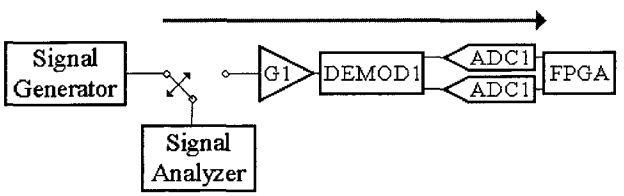

Figure 6.2.1 – Setup to measure magnitude response of *RF*1 and ADC1's

The training sequence power measured by the signal analyzer in the 5 MHz RF passband is -28.5 dBm. The power measured at the ADC1's output over a 2.5 MHz baseband passband is 53.6 dB_{dig}. The unit dB_{dig} is based on a power measurement in the fixed-point integer domain, and is referenced to 1Ω. The magnitude response of the channel is:

$$\left|H_{G1}H_{DEMOD1}H_{ADC1}\right| = 53.6\,dB_{dig} + 28.5\,dBm = 82.1\,dB_{dig-m} \qquad (6.2.1)$$

The subscript 'dig-m' is introduced since the ratio of digital power in dB_{dig} to the RF power in dBm does not have units of dB.

6.2.2 Primary ADC's

The ADC1 characterization setup has a signal generator connected to both ADC1 evaluation boards via a splitter. The input signal is a sine wave, and is increased in magnitude until it exceeds each of the ADC1's full-scale voltage ranges. This threshold is called the 0 dB full-scale (dBFS) input power, and is shown in Fig. 6.2.2a. The power values in this figure have been corrected for the attenuation due to the splitter. The magnitude responses of the ADC1 evaluation boards are proportional to the inverse of the curves in Fig. 6.2.2a (ie. multiply curves by -1). The two ADC1's are labelled ADC1-I and ADC1-Q, which correspond to the in-phase and quadrature-phase outputs of the demodulator. Some frequency variation is expected due to the output matching conditions of the amplifiers on the evaluation boards. Furthermore, these amplifiers are not designed to operate down to DC, hence the relatively high 0 dBFS power at low frequencies.

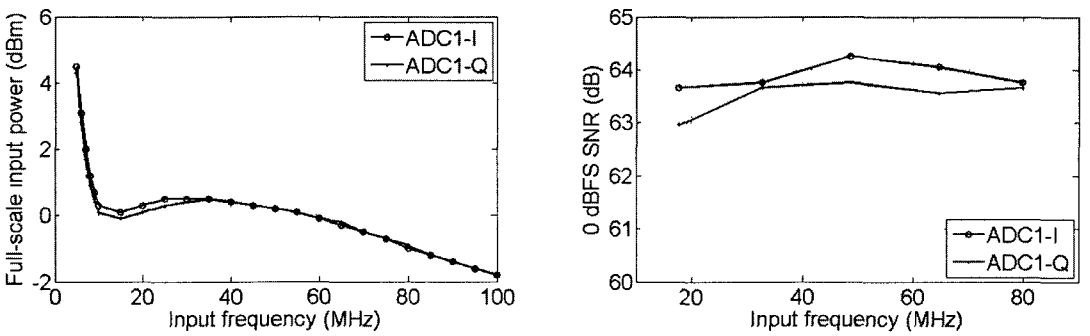

Figure 6.2.2 - ADC1's a) Full-scale input b) SNR relative to 0 dBFS input

The noise PSD is determined at the output of each ADC1 for various input frequencies and power levels. Exactly 8192 consecutive data samples are transferred from the outputs of the ADC1's to a computer and into Matlab where the PSD is processed. Normally the noise PSD is determined over the entire Nyquist bandwidth, but in this case the noise results are most important in the passband. Consequently, the average noise PSD over the 2.5 MHz baseband passband is used for the SNR calculation. The full-scale output power for each ADC1 is $\frac{1}{2}(2^{13})^2$ dB_{dig}. This full-scale power and the measured noise PSD are used to calculate the 0 dBFS SNR for each ADC1:

$$SNR_{ADC} = \frac{\frac{1}{2}\left(2^{13}\right)^2}{PSD \times B} \tag{6.2.2}$$

The 0 dBFS SNR is shown in Fig. 6.2.2b for both ADC's at various input frequencies. The measured SNR accounts for the actual ADC1 SNR, as well as digital clock jitter and noise from the other components on the ADC1 evaluation boards. The curves in Fig. 6.2.2b are used to estimate the 0 dBFS SNR for ADC1-I to be 64.5 dB, and 64.2 dB for ADC1-Q. The two ADC1's generate the same output noise as if both ADC1's had

identical SNR's of 64.3 dB. This SNR is used in place of SNR_{ADC} for the output noise expression.

6.2.3 Noise Factor of *RF*1 Chain

The noise factor of the *RF*1 chain is obtained by measuring the output noise of the ADC1's when the input PSD to the *RF*1 chain is kT_0 over the 5 MHz RF passband. A sine wave is also injected into the *RF*1 input to ensure the inputs to the data converters span several quantization levels. A directional coupler is used to inject the sine wave as shown in Fig. 6.2.3.

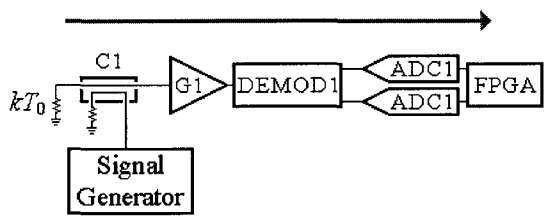

Figure 6.2.3 – Setup to measure noise factor of *RF*1

The PSD in the passband of the complex digital output signal from the ADC1's is:

$$kT_{ADC-I} + kT_{ADC-Q} = 2kT_0F_{RF1}\left|H_{G1}H_{DEMOD1}H_{ADC1}\right|^2 + 2\times\frac{\frac{1}{2}\left|v_{ADC,fs}\right|^2}{SNR_{ADC}B} \qquad (6.2.3)$$

The output noise PSD for each ADC1 is expressed in terms of its equivalent noise temperature in (6.2.3). The digital PSD on the left side of (6.2.3) is measured for different input signal levels and frequencies from the signal generator. On average this value in the passband is -61.0 dB_{dig}/Hz, The noise factor of the *RF*1 chain in (6.2.3) is solved using:

- B=100 MHz, $v_{ADC,fs}=\pm 2^{13}$ (since ADC has 14 bits)
- SNR_{ADC} = 64.3 dB (from Section 6.2.2)

- $|H_{G1}H_{DEMOD1}H_{ADC1}|$=82.1 $dB_{dig\text{-}m}$ (from Section 6.2.1)

The result of these substitutions is a noise figure estimate of: F_{RF1}=26.2 dB.

The noise figure of the *RF*1 chain is next used to isolate the noise factor of the DEMOD1 component.

$$F_{RF1} = F_{G1} + \frac{F_{DEMOD1} - 1}{|H_{G1}|^2} \Rightarrow F_{DEMOD1} = (F_{RF1} - F_{G1})|H_{G1}|^2 + 1 \qquad (6.2.4)$$

The gain and noise factor of the *G*1 amplifier are 19.2 dB and 1.03 dB, respectively. Substituting these values and F_{RF1} into (6.2.4) yields: F_{DEMOD1}= 45.3 dB. Furthermore, substituting the gain of *G*1 into (6) gives: $|H_{DEMOD1}H_{ADC1}|$=62.9 $dB_{dig\text{-}m}$. These magnitude terms do not need to be separated for the output noise expression, so the 62.9 dB is arbitrarily distributed as: $|H_{DEMOD1}|$=2.9 *dB*, and $|H_{ADC1}|$=60 $dB_{dig\text{-}m}$. The full-scale digital value of the ADC1's in this test is $\pm 2^{13}$, so according to (4.1.7) the average 0 dBFS input power of the ADC1's over the passband is: $P_{ADC,FS}$=15.3 dBm.

6.2.4 Magnitude Response DAC-*RF*2 Chain

A training sequence with a uniform PSD is generated in the FPGA and sent through the DAC's, then measured by a signal analyzer at the output of the lower FF path. The *RF*2 chain consists of the MOD and *G*2 amplifier. This setup is demonstrated in Fig. 6.2.4.

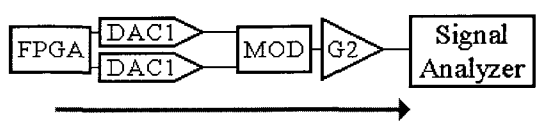

Figure 6.2.4 – Setup to measure magnitude response of the DAC's and *RF*2

The training sequence power in the digital domain has 70.5 dB_{dig} of power over the Nyquist bandwidth; which has 54.5 dB_{dig} of power in a 2.5 MHz baseband passband. The

power measured by the signal analyzer in a 5 MHz RF bandwidth at 900 MHz is -9.3 dBm. The magnitude response of the channel is:

$$\left|H_{DAC1}H_{MOD}H_{G2}\right| = -9.3\,dBm - 54.5\,dB_{dig} = -63.8\,dB_{m-dig} \tag{6.2.5}$$

The DAC's have a 10 dBm full-scale output [*39*], and the full-scale digital value within this test is $v_{DAC,fs}=\pm 2^{13}$. Also the gain of the *G*2 amplifier is 11.5 dB . These component values, along with (4.1.7) and (6.2.5) are used to calculate the gain of the modulator:

$$\left|H_{MOD}\right| = -63.8dB_{m-dig} - 11.5dB - 10dBm + 20\log_{10}\tfrac{1}{\sqrt{2}}2^{13} = -10.0\,dB \tag{6.2.6}$$

6.2.5 Noise Factor of *RF*2 Chain

The input of the *RF*2 chain is the IQ inputs of the modulator, which are high impedance ports. To determine the chain's noise factor, the DAC's are powered-off and the noise PSD at the output of the chain is measured in the passband with the signal analyzer. The measurement setup is shown in Fig. 6.2.5.

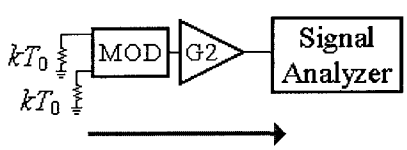

Figure 6.2.5 – Setup to measure noise factor of *RF*2

The output noise PSD of the *RF*2 chain with no inputs is:

$$kT_{RF2,out} = kT_0\left(F_{MOD} + \frac{F_{G2}-1}{\left|H_{MOD}\right|^2}\right)\left|H_{MOD}H_{G2}\right|^2 \tag{6.2.7}$$

The output noise measurement is -149.6 dBm/Hz in the 5 MHz RF passband. The modulator noise figure is obtained by applying this measurement, the modulator gain

from (6.2.6), and the *G*2 amplifier's nominal gain and noise figure to (6.2.7). The result is F_{MOD}=22.5 dB.

6.2.6 Primary DAC's

Measurements are used to incorporate the effects of the digital clock jitter and the balun on the evaluation board into the DAC1 SNR metric. The DAC's drive high impedance ports on the modulator, so they cannot be characterized using the 50Ω port on the signal analyzer.

The sine wave output of a signal generator is sent into the FPGA via the demodulator and ADC1's. The sine wave is necessary to ensure the input signals to the DAC1's span several quantization levels The signal is passed through an FIR notch filter with two transfer function zeros at z=1. This notch filter removes the noise generated outside the FPGA within the notch bandwidth. Any noise within this bandwidth at the DAC1's outputs is then due primarily to the DAC1's. The DAC1's are directly connected to the *RF*2 chain, which is connected to a signal analyzer. The measurement setup is shown in Fig. 6.2.6.

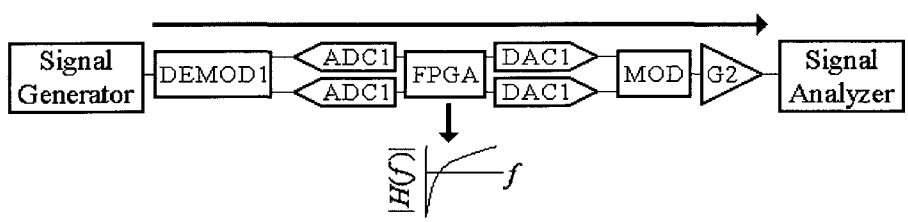

Figure 6.2.6 – Setup to measure DAC SNR

The only contribution to the noise PSD at the signal analyzer in the passband is expected to be from the DAC's and *RF*2 chain. The digital notch filter removes most of the noise from DEMOD1 and the ADC1's. The noise PSD in the passband at the signal analyzer is:

$$kT_{out} = \frac{P_{DAC,fs}}{SNR_{DAC}B}\left|H_{RF2}\right|^2 + kT_0\left(F_{RF2} - 1\right)\left|H_{RF2}\right|^2 \tag{6.2.8}$$

The SNR_{DAC} term can be isolated in (6.2.8) since the following terms are already known:

- H_{RF2} = 1.4 dB (from Section 6.2.3)
- F_{RF2} = 22.9 dB (from Section 6.2.5)

The value of kT_{out} in (6.2.8) is measured for various input frequencies from the signal generator's. The 0 dBFS SNR_{DAC} results are shown in Fig. 6.2.7.

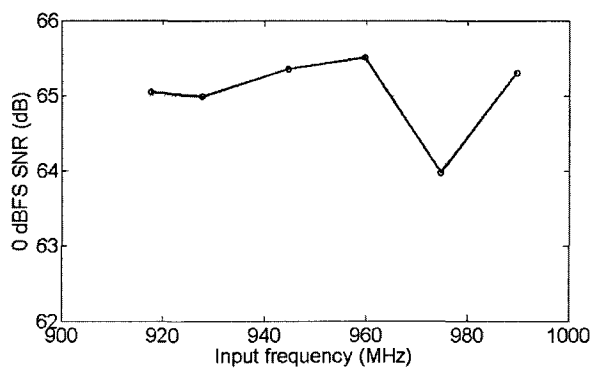

Figure 6.2.7 – 0 dBFS SNR for both DAC1's

The curves in Fig. 6.2.7 are used to estimate the 0 dBFS SNR for both DAC1's to be 65.1 dB. The SNR from the DAC data sheets is specified for a few scenarios, the most relevant being 74 dB for a clock frequency of 210 MHz, and an input frequency of 5 MHz [*39*]. The value from the data sheet does not include the evaluation board components or clock jitter.

6.2.7 Group Delay

The group delay of the FF path containing the digital components is related to the attenuation and noise performance. The best performance of the hybrid RF-DSP FF filter is obtained when the group delay mismatch between the two FF paths is minimized. The largest contribution to the group delay in this path is from the data converters and FPGA. The ADC1 data sheets specify a latency of 3.5 clock cycles [*45*]. The DAC1 data sheets specify a combined delay and settling time to 0.1% of 12 ns [*39*].

To measure the digital component latency one of the DAC1's is connected to one of the ADC1's, with a 10 dB attenuator between them. A training sequence generated in the FPGA is sent to the appropriate DAC1, and then recorded at the output of the connected ADC1. This data is processed within Matlab to determine the frequency response of the DAC1-ADC1 channel. The magnitude response and group delay of this channel are shown in Fig. 6.2.8 with and without an LPF connected between DAC1 and ADC1. The cut-off frequency of the LPF is 90 MHz.

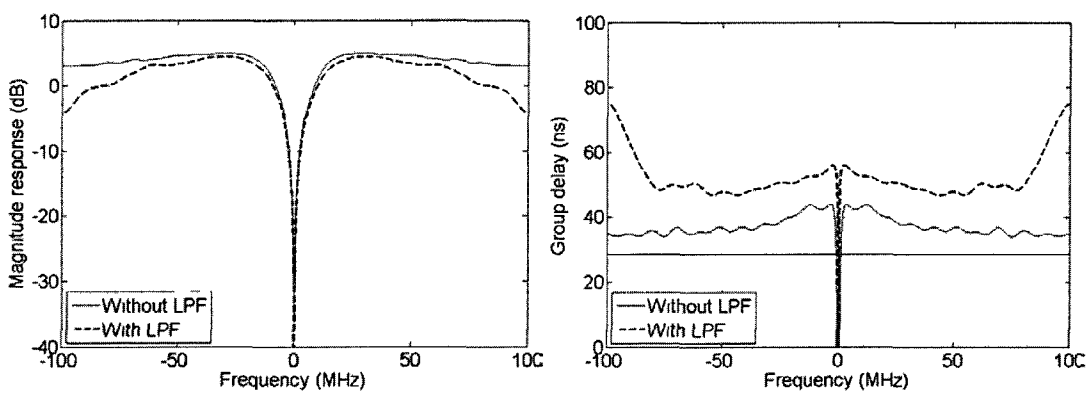

Figure 6.2.8 - Frequency response of DAC1-ADC1 channel a) Magnitude b) Group delay

The deep notch in the magnitude response is due to the transformer balun on the DAC1 evaluation board, and the low-frequency response of the amplifiers on the ADC1 evaluation board. The group delay without the LPF varies between approximately 35 and 41 ns. The combined latency from the data sheets of the ADC1 and DAC1 is 28.5 ns, which is shown by a thick horizontal line in Fig. 6.2.8b. The additional group delay is due to other components on the evaluation boards. The effect of the LPF is to add more than 10 ns of group delay.

The frequency response of the total bottom FF path is obtained by connecting the output of the *C*2 coupler to the input of the *C*1 coupler, and running the same channel estimation algorithm. The original path between these couplers is disconnected and 50Ω terminations are placed on the exposed ports. These modifications are shown in Fig. 6.2.9, where a thick arrow indicates the measured channel. The baseband equivalent frequency response of the bottom FF path is shown in Fig. 6.2.10, at four different RF frequencies.

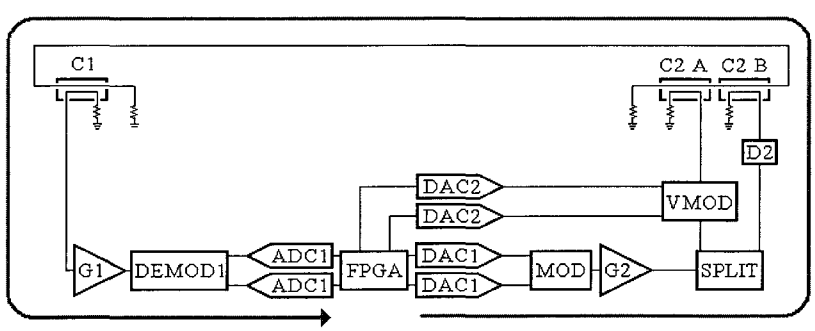

Figure 6.2.9 – Hardware setup to measure bottom FF path channel

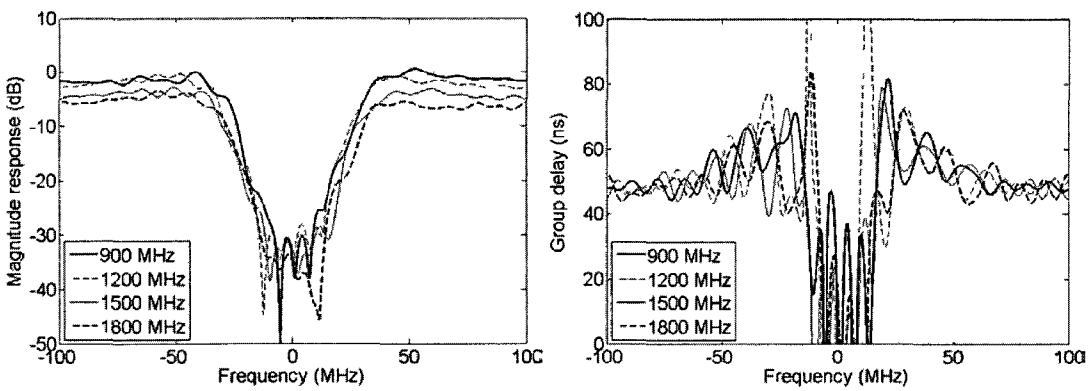

Figure 6.2.10 - Frequency response bottom path a) Magnitude b) Group delay

The magnitude response in Fig. 6.2.10 shows a wider stopband near DC than Fig. 6.2.8. This wider attenuation bandwidth is due to the tunable BSF. The group delay response has a maximum near 80 ns, which occurs near the filter's passband edges and rolls-off to below 50 ns at the edges of the operating bandwidth. During normal operation of the hybrid RF-DSP FF filter, the digital filter compensates for the magnitude response and group delay of this bottom FF path to ensure FF attenuation at the output of the filter.

6.2.8 LO Phase Noise

The wideband phase noise of the LO is required for the noise analysis. This parameter is measured using the phase noise application within the signal analyzer. The SSB phase noise of the LO is shown in Fig. 6.2.11, with the LO tuned to 900 MHz. The wideband phase noise lies somewhere between -133 and -140 dBc/Hz. The noise measurements in Chapter 7 are checked at 900 MHz, so this spot measurement of the wideband phase noise is sufficient.

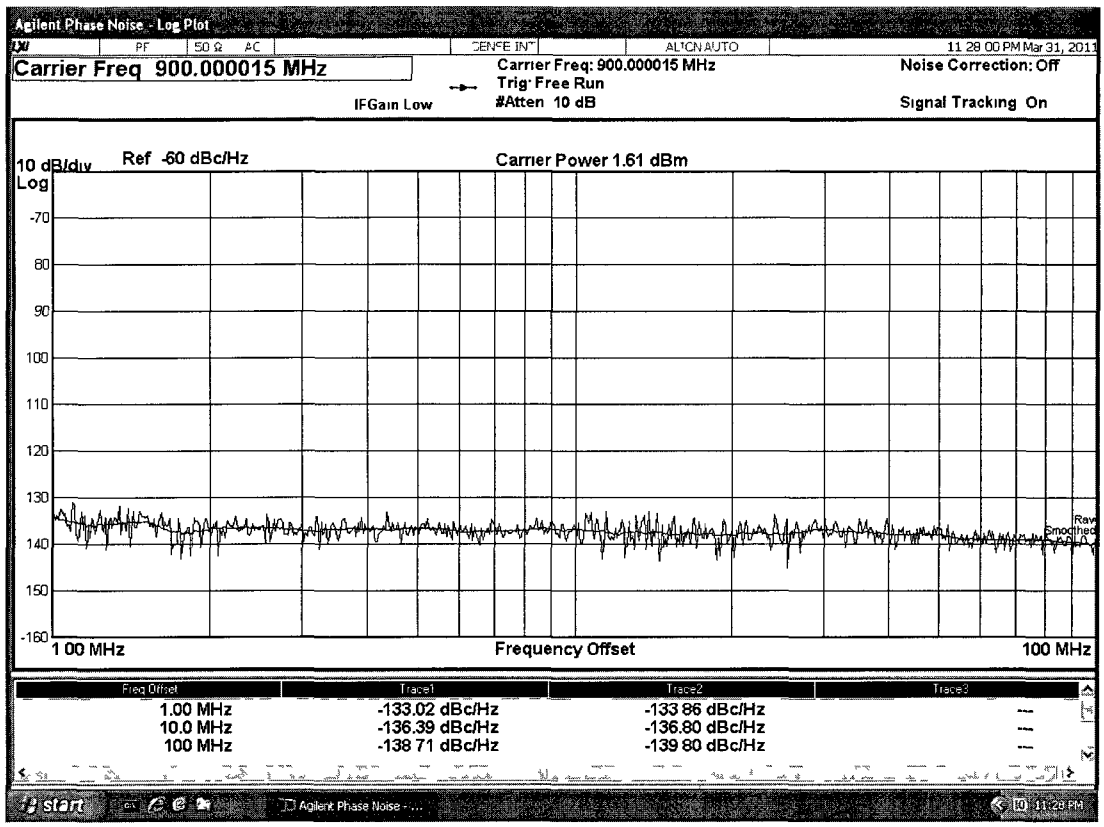

Figure 6.2.11 – SSB phase noise of LO at 900 MHz

6.2.9 Summary of Bottom FF Path Characterization

The component values obtained in this section are required for the output noise expression from Chapter 4. These component values are summarized in the following four tables. The H_{C3l} term includes 1.35 dB of attenuation from the long coaxial cable that connects the output of the filter to the signal analyzer. This cable loss was measured with a vector network analyzer.

Table 6.3 – Component values for *RF*1

Variable	Value	Comments
F_{AMP1}	1.03 dB	Table 6.2
H_{AMP1}	19.2 dB	Table 6.2
F_{DEMOD1}	45.3 dB	Section 6.2.3
H_{DEMOD1}	2.9 dB	Section 6.2.3
PN_{WB}	-133 dBc/Hz	Section 6.2.8

Table 6.4 – Component values for *RF*2

Variable	Value	Comments
F_{AMP2}	4.3 dB	Table 6.2
H_{AMP2}	11.5 dB	Table 6.2
F_{MOD}	22.5 dB	Section 6.2.5
H_{MOD}	-10.0 dB	Section 6.2.3
PN_{WB}	-133 dBc/Hz	Section 6.2.8

Table 6.5 – Values for digital terms

Variable	Value	Comments
$P_{ADC,fs}$	15.3 dBm	Section 6.2.3
B	100 MHz	Digital clock is 200 MHz
SNR_{ADC}	64.3 dB	Section 6.2.2
SNR_{DAC}	65.1 dB	6.2.6
α_{MULT}	2^{30}	Digital filter specifics
α_{DAC}	2^{16}	Digital filter specifics
N	32	Digital filter specifics
b_{MULT}	44	Digital filter specifics
$P_{DAC,fs}$	10 dBm	Data sheet [*39*]

Table 6.6 – Values for RF FF path

Variable	Value	Comments
H_{C1t}	-0.81 dB	Table 6.2
H_{C1c}	-10 dB	Table 6.2
H_{C2t}	-0.4 dB	Table 6.2
H_{C2c}	-20 dB	Table 6.2
H_{C3t}	-2.25 dB	Table 6.2
H_{TOP}	1	Small coaxial cable
F_{TOP}	1	Small coaxial cable

6.3 Algorithm Verification

Three algorithms are used within the hybrid RF-DSP FF filter. These algorithms are verified in this section.

6.3.1 Verification of BSF Algorithm

The BSF algorithm varies the I and Q inputs of a vector modulator to find the realization that causes the most attenuation in the passband. After the optimum vector modulator inputs are determined, the DAC1's send out a band-limited noise signal with most of its energy within ±5 MHz of the centre frequency. This band-limited signal is used only during BSF verification as a means to calculate the average attenuation over the passband.

A signal analyzer is connected to the output of the hybrid RF-DSP FF filter to verify proper operation of the BSF algorithm. The signal analyzer is used to measure the PSD of the band-limited signal with the BSF on ('BSF on'), and with the vector modulator tuned to maximum attenuation ('BSF off'). The output PSD's for these two cases are shown in Fig. 6.3.1 at 900 MHz. This figure shows notching due to the BSF in the passband. The output PSD data points from the analyzer are exported to Matlab, where they are subtracted (logarithmic scale) to obtain the BSF magnitude response. The exported PSD traces and resulting magnitude response are shown in Fig. 6.3.2.

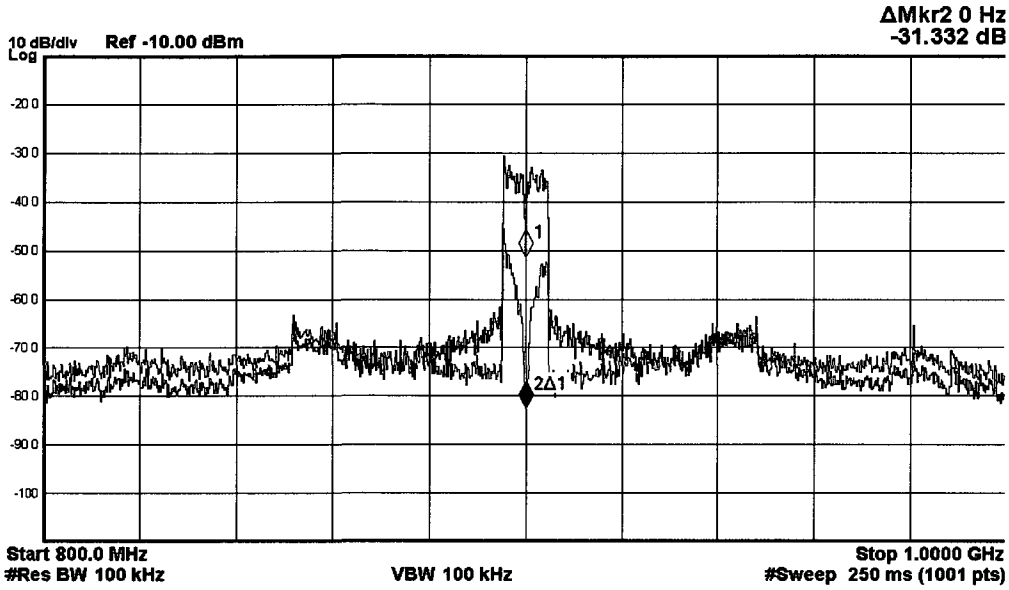

Figure 6.3.1 - PSD out of FF filter with BSF on and off

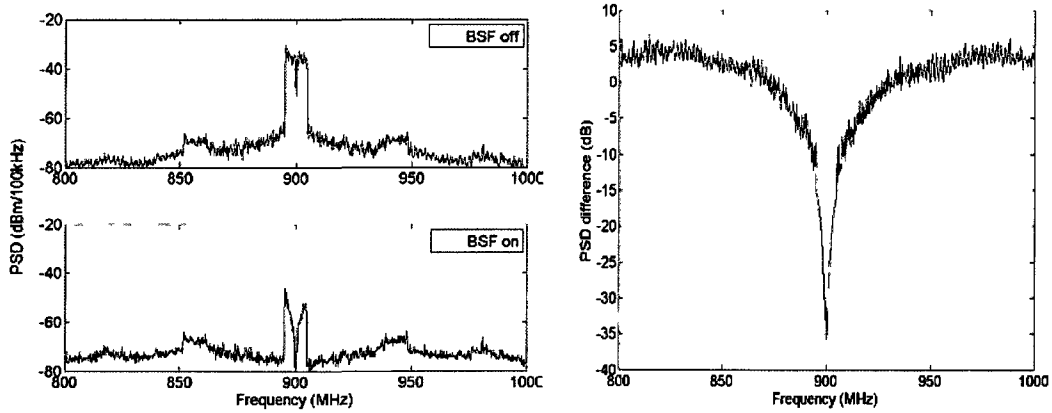

Figure 6.3.2 - PSD out of FF filter with BSF on and off exported from PXA to Matlab a) PSD with BSF in on and off states b) magnitude response

The average attenuation values over 2.5, 5 and 10 MHz RF bandwidths for the response in Fig. 6.3.2 are 30.6, 25.6 and 19.8 dB, respectively. These measurements of the BSF performance are repeated at multiple RF frequencies to arrive at the attenuation curves in Fig. 6.3.3.

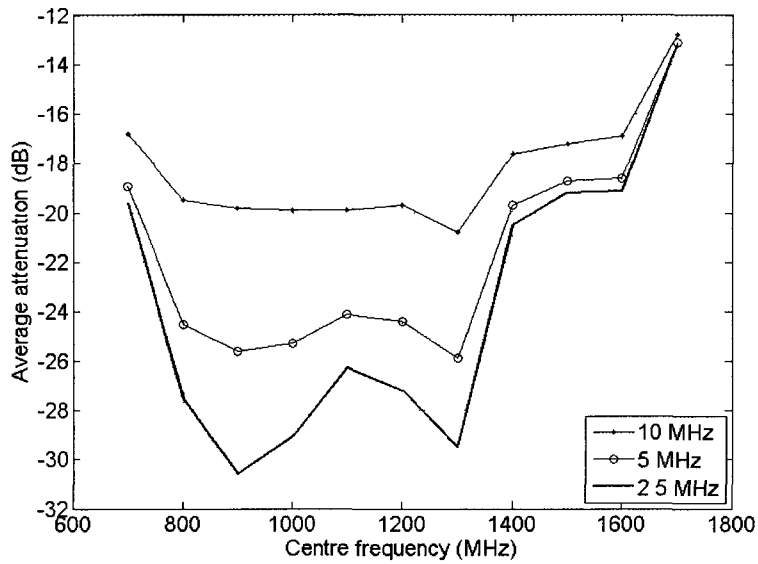

Figure 6.3.3 - BSF average attenuation for 3 different passband bandwidths

These attenuation curves show that for 800 to 1300 MHz, average attenuation values of 26, 24, and 20 dB are achievable for 2.5, 5 and 10 MHz RF passband bandwidths, respectively. The frequency limits to the BSF attenuation performance are due to the frequency range of the vector modulator. The data sheets of the vector modulator specify an RF range of 700 to 1000 MHz [*46*].

6.3.2 Verification of $h_{23,est}[n]$ Estimation Algorithm

The $h_{23,est}[n]$ channel estimation algorithm is used on the channel located between the DAC1's and the ADC2's. This channel is shown in Fig. 6.3.4. The $h_{23,est}[n]$ algorithm uses two signals. The first is the training sequence, which is the channel's input signal. The second signal is that which comes out of the ADC2's.

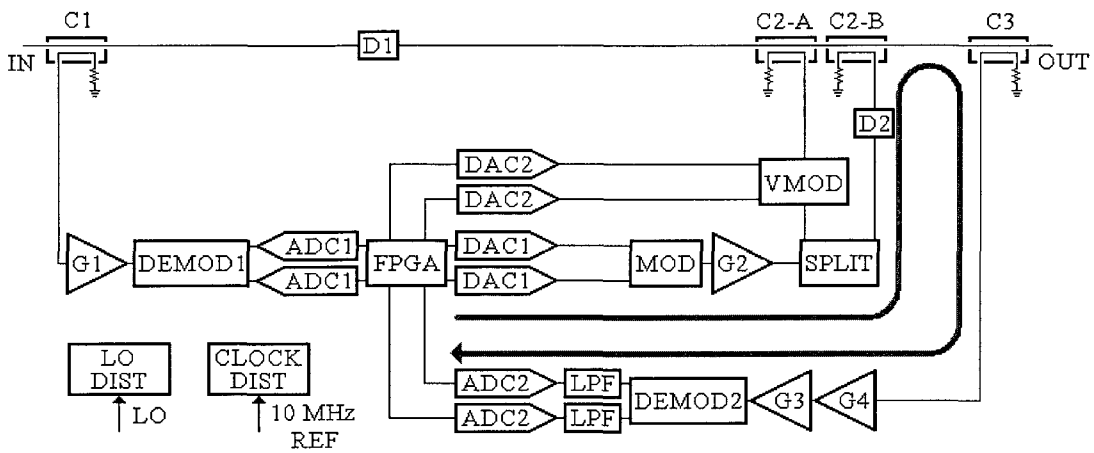

Figure 6.3.4 - Hardware prototype with $h_{23}[n]$ channel indicated

To verify proper operation of the $h_{23,est}[n]$ channel estimation algorithm, the channel's actual response is required. This response is not known, but is estimated using a software-based technique and the channel's input and output signals. These signals are picked-off the FPGA and transferred to Matlab. The channel's input signal is a training sequence with a uniform PSD over the Nyquist bandwidth. The two picked-off signals are used to find the channel response in the MMSE sense by calculating the corresponding FIR Wiener filter. Approximately 65500 sequential samples are transferred from the FPGA for the input and output signals. The FIR Wiener filter response from this software based channel estimation is the 'Simulink' curve in Fig. 6.3.5. This result is obtained with the hybrid RF-DSP FF filter operating at 1200 MHz. The steady-state solution of the FPGA algorithm is also captured. The resulting magnitude response and group delay are the 'FPGA' curves in Fig. 6.3.5.

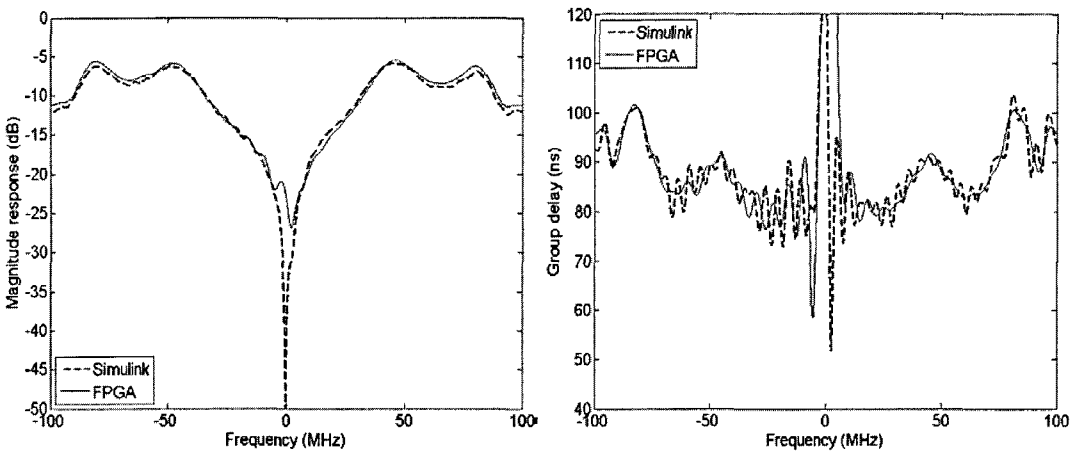

Figure 6.3.5 - Comparison of FPGA and software-based $h_{23,est}[n]$ algorithms

The curves in Fig. 6.3.5 are very similar, thereby verify proper operation of the $h_{23,est}[n]$ algorithm in the FPGA. These curves deviate in the passband, but this is due to the difference in training sequences. The FPGA based algorithm uses a training sequence that has almost no passband component to minimize degradation to the noise figure. The result is a relatively large channel estimation error within the passband. The channel estimate is not used within the passband; therefore this error is not a problem.

The curves in Fig. 6.3.5 are not sufficient to ensure the $h_{23}[n]$ channel estimation algorithm meets the stability criteria derived in Section 5.3.1. The stability criteria from Section 5.3.1 is shown again in Fig. 6.3.6.

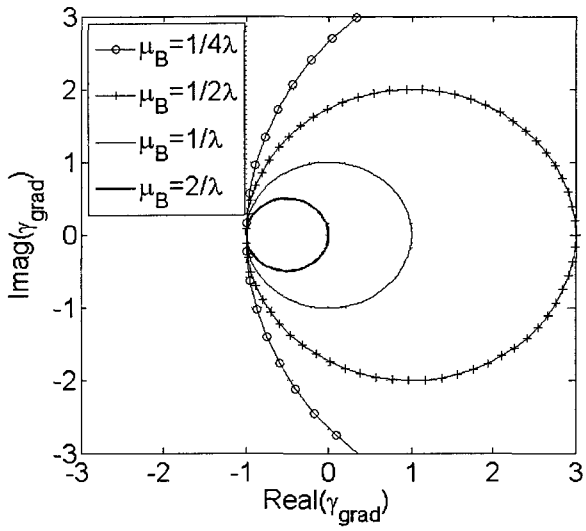

Figure 6.3.6 - Stability criteria for main algorithm (from Section 5.3.1)

The value of γ_{grad} is defined in Section 5.3.1, and is repeated in (6.3):

$$\gamma_{grad} = \frac{H_{23,err}(z)}{H_{23}(z)} \tag{6.3.1}$$

In this equation, $H_{23,err}(z)$ is the z-domain channel estimation error obtained by subtracting the channel estimate, $H_{23,est}(z)$, from the actual channel response, $H_{23}(z)$. A specific channel estimate will have a different value of γ_{grad} at each frequency point, where $z=exp(j2\pi fT_s)$. Consider plotting an individual point in Fig. 6.3.6 for γ_{grad} at each frequency in the operating bandwidth, excluding the passband. The stability criterion is then determined by the smallest circle, like those in Fig. 6.3.6, which surrounds all of the plotted points. The stability criterion can be rearranged to represent a maximum limit on $\mu_B\lambda$. The values of μ_B and λ are the adaptive algorithm step size and eigenvalue of the filtered input signal ($u_{est}[n]$) to the main digital filter. These values are introduced in Chapter 5.

The 32 tap FPGA channel estimation algorithm is run and the steady-state result is picked-off the FPGA. This result is the channel estimate, $H_{23,est}(z)$. The software based channel estimation algorithm is also run, which yields $H_{23}(z)$. The $H_{23,est}(z)$ and $H_{23}(z)$ results are then used to determine γ_{grad} from -100 to 100 MHz, in steps 0.1 MHz, excluding the passband. These values of γ_{grad} are plotted with the stability circles in Fig. 6.3.7 for passband bandwidths of 2.6 and 5 MHz. The smallest circles that surrounds the plotted points in Fig. 6.3.7a and 6.3.7b correspond to a $\mu_B\lambda$ of 0.93 and 1.12, respectively. These values define the stability criteria for the main algorithm depending on the passband bandwidth. The results in Fig. 6.3.7 are obtained with the hybrid RF-DSP FF filter operating at 900 MHz.

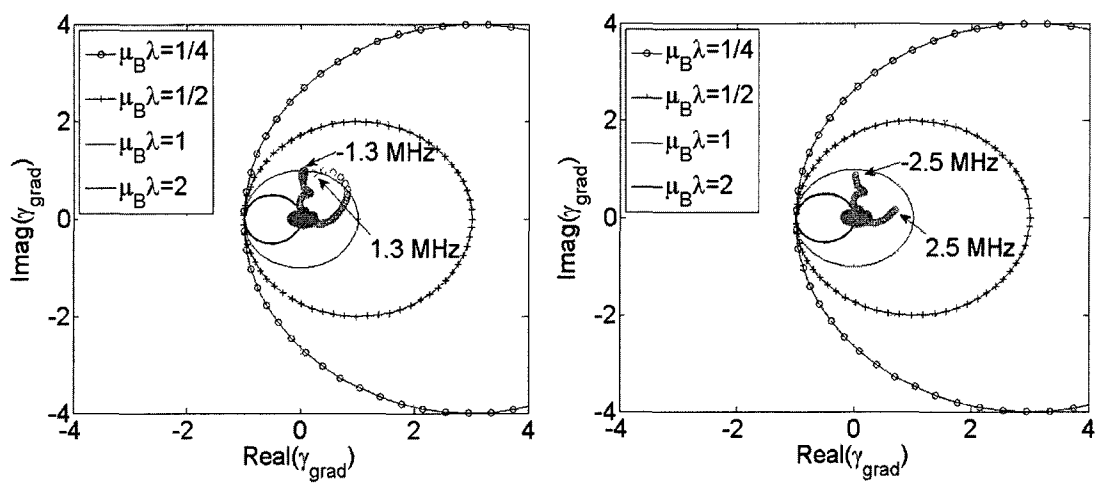

Figure 6.3.7 – Stability criteria a) 2.6 MHz passband b) 5 MHz passband

This section has demonstrated proper operation of the $h_{23,est}[n]$ channel estimation algorithm. The stability criterion was determined at one RF frequency, but is not expected to significantly vary for the other RF frequencies. The criterion corresponds to a limit on $\mu_B\lambda$ for the main algorithm.

6.3.3 Verification of Main Algorithm

The purpose of this section is to demonstrate proper operation of the main algorithm for a single realization of an input interference environment. The output PSD's with the hybrid RF-DSP FF filter turned off and on are measured by a signal analyzer and are shown in Fig. 6.3.8 and 6.3.9, respectively. Three large tones are visible in Fig. 6.3.8, each labelled with a marker. When the filter is turned on and reaches steady-state, these three sinusoids appear attenuated at the filter output, shown in Fig. 6.3.9. The attenuation values are obtained by subtracting the marker values, which are shown below the PSD's in the tables. The attenuation values are 38.2, 48.3 and 46.7 dB for the sinusoids at 926, 990 and 1015 MHz. The centre frequency of the filter is 950 MHz. These results demonstrate proper operation of the main algorithm for a single realization. The next chapter reports on more complex input interference scenarios.

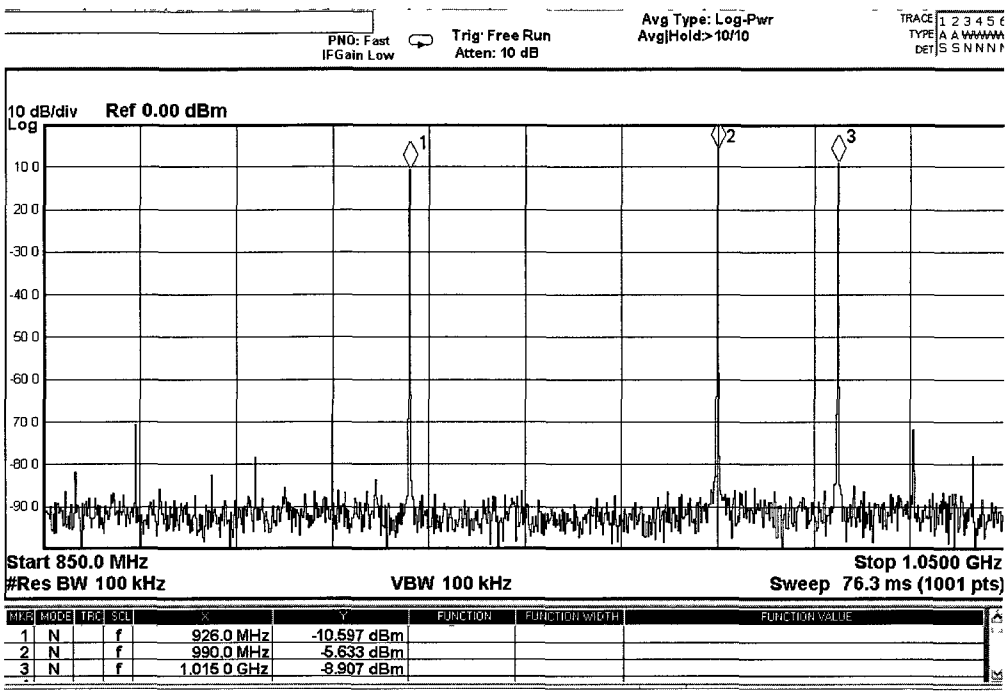

Figure 6.3.8 - Filter output with all zero filter taps

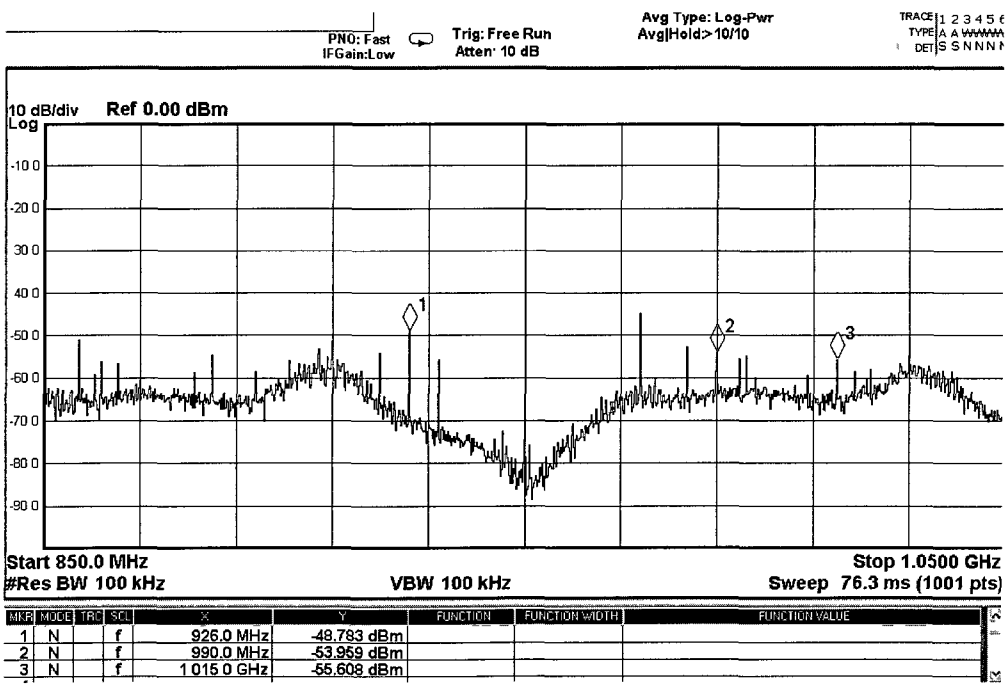

Figure 6.3.9 - Filter output with main algorithm in steady state

6.4 Hardware Prototype Summary

This chapter has introduced the hardware prototype and verified proper operation of the three adaptive algorithms. Component values have been obtained for the analytical expressions in Chapter 4 either from data sheets, or measurements. The following chapter reports measurements of the hardware prototype for different filter realizations and input interference scenarios.

7 Hardware Prototype Performance

The measured performance of the hardware prototype is reported in this chapter. Narrowband input signals are used to demonstrate the filter's: frequency agility, handling of dynamic inputs and near band attenuation performance. Wideband input signals are used to determine how the attenuation performance is related to: the number of FIR filter taps, group delay mismatch between the two FF paths, and the block size within the adaptive algorithm. Additionally, the noise analysis from Section 4.1 is verified.

7.1 Measurement Setup

The hybrid RF-DSP FF filter is tested with several different input interference scenarios. The outputs of several signal generators (Sig Gen) are combined, and then connected to the input of the filter. Four Sig Gen outputs are combined using a 4-1 combiner. A maximum of eight Sig Gen outputs are combined using two 4-1 combiners, and a 2-1 combiner to combine the outputs of the 4-1 combiners. The output of the filter is connected to an Agilent PXA signal analyzer. The measurement setup is illustrated in Fig. 7.1.1, where the details of the hybrid RF-DSP FF filter are omitted. Throughout this chapter the digital clock frequency is fixed at 200 MHz, and the LO is varied over the filter's operating range.

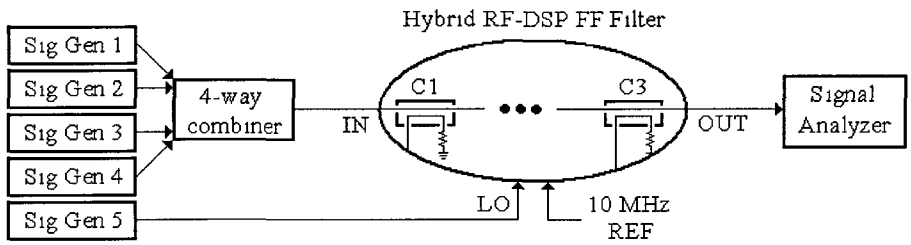

Figure 7.1.1 – Measurement setup for testing the prototype

7.2 Narrowband Inputs

Narrowband signals are used at the input to the hybrid RF-DSP FF filter in this section. This type of input signal is characteristic in an RF environment containing only narrowband interferers, or a lot of spurs. The digital filter has 32 complex taps, and the adaptive algorithm block size (N_1) is 128.

7.2.1 Frequency Agility

A multi-tone input signal is used to demonstrate frequency agility in this section. The centre frequency of the hybrid RF-DSP FF filter is varied from 800 to 1800 MHz, in 100 MHz steps. At each frequency step, 8 tones are applied to the input of the filter that lie within ±100 MHz of the filter's centre frequency. At each frequency step the FF attenuation of all 8 tones is measured. The 800 MHz output PSD of the FF filter is shown in Fig. 7.2.1 and 7.2.2 with the filter turned off and on, respectively. Eight markers are visible in Fig. 7.2.1, collocated with the eight tones. The table at the bottom of each figure lists the frequencies and magnitudes of the markers. The markers drop when the filter is turned on, as can be seen in Fig. 7.2.2. The difference between the marker heights in these two figures represents the FF attenuation inflicted by the lower FF path.

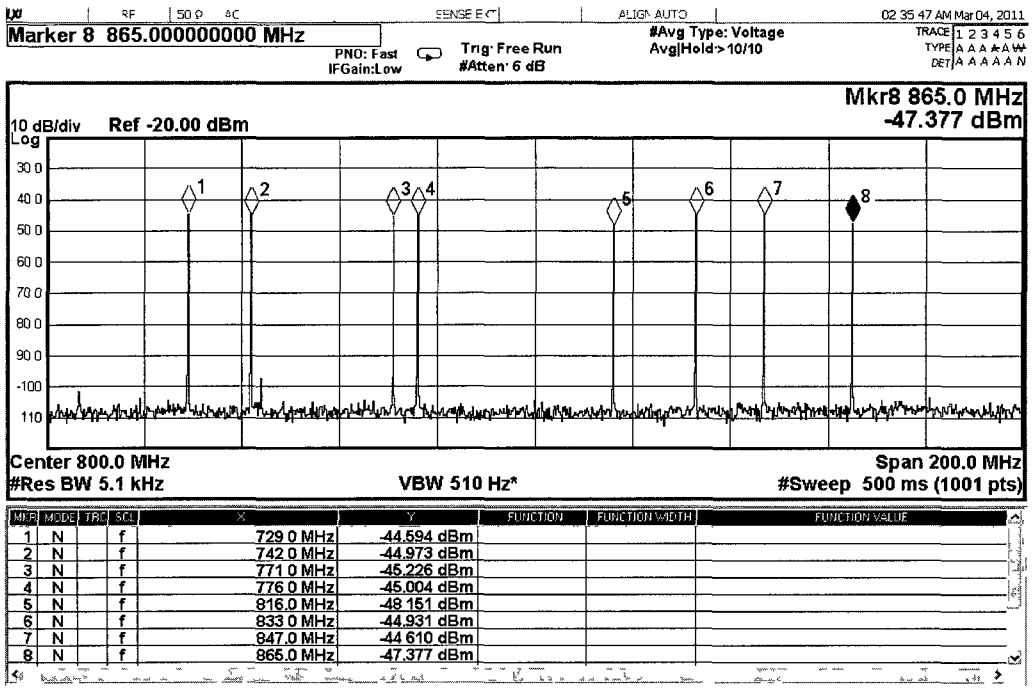

Figure 7.2.1 – Output PSD of 8 tone case with digital filter off

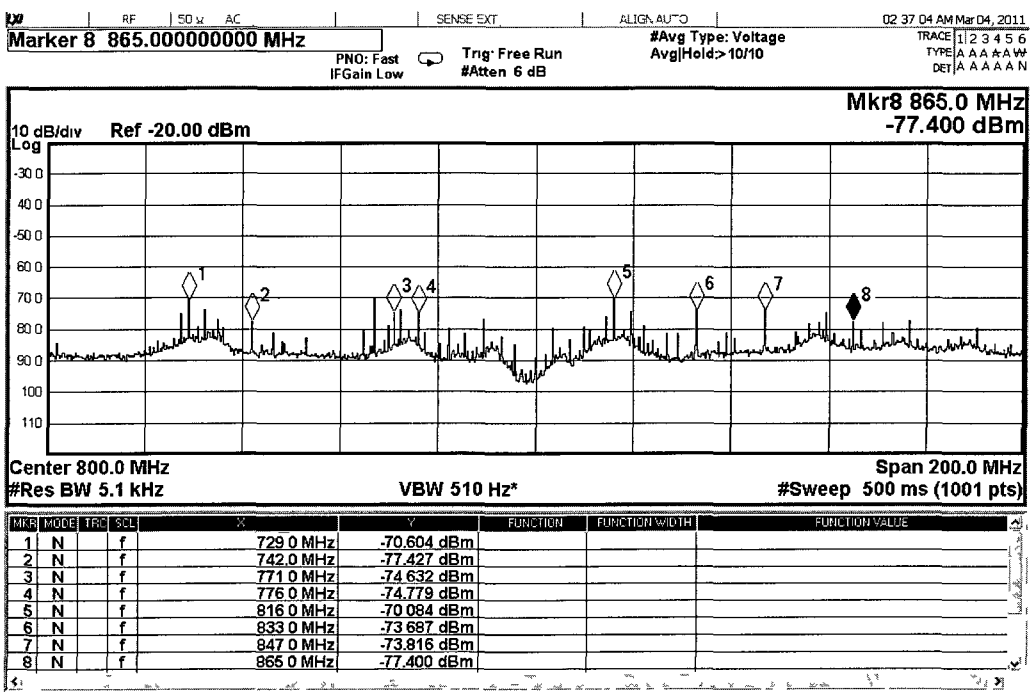

Figure 7.2.2 – Output PSD of 8 tone case with digital filter on

The differences between the marker heights are recorded at each centre frequency for each input tone. These attenuation results are tabularized in Table 7.1.

Table 7.1 – Frequency agile performance of 8 tone attenuation

Tone	Attenuation (dB) with filter operating at indicated centre frequency (MHz)										
(MHz)	800	900	1000	1100	1200	1300	1400	1500	1600	1700	1800
-71	25.8	26.7	27.0	28.3	29.0	28.1	29.7	30.6	30.0	30.5	29.7
-58	32.0	32.0	29.7	31.5	32.8	31.4	34.8	34.9	33.3	33.6	35.8
-29	29.0	32.6	28.7	33.6	34.2	30.2	38.9	39.2	33.3	33.7	39.0
-24	30.1	32.7	31.7	34.9	36.0	32.2	40.7	41.5	35.1	35.3	41.0
16	22.0	26.4	26.3	31.0	28.8	29.5	34.5	36.6	35.2	38.6	37.2
33	28.7	32.0	28.2	30.8	32.4	30.7	39.4	39.5	36.3	37.2	38.1
47	30.8	32.4	28.0	31.1	32.5	30.4	37.7	37.8	35.7	38.1	36.9
65	29.7	33.8	29.3	29.7	31.2	29.8	31.9	30.5	30.2	31.0	29.9

Excluding the shaded values, the tone attenuation is better than 28 dB for all centre frequencies. Shading indicates where attenuation is worse than 28 dB. The tone at -71 MHz is the nearest tone to the band edge of the hybrid RF-DSP FF filter's operating bandwidth. Several components cause performance to degrade near the band edges, including the LPF's in the feedback path. The attenuation is still better than 25 dB for the -71 MHz tone, so the degradation is not significant enough to avoid operating that close to the band edge. The other shaded row corresponds to the tone at 16 MHz, which is the nearest tone to the passband edge. At this location the tone is significantly attenuated by the RF BSF, hence requires gain compensation by the digital filter. More measurements are performed near the passband edge in Section 7.2.3. The results in Table 7.1 demonstrate proper operation between 800 and 1800 MHz, thereby demonstrating frequency agility.

7.2.2 Dynamically Changing Inputs

The hybrid RF-DSP FF filter handles a dynamic input interference scenario by adapting the filter's notches in order to keep them collocated with the interferers. This section demonstrates this behaviour by varying the input signal PSD in time. In order to report this adaptive behaviour, the filter's output PSD's are shown before and after each time the input signal is changed. Furthermore, the evolution of the filter's 32 taps is also plotted with respect to time. The filter tap values are extracted from the FPGA in real-time using the SignalTapII tool within the QuartusII software. This tool is a virtual logic analyzer.

The following measurements are obtained with the hybrid RF-DSP FF filter operating at a centre frequency of 1225 MHz. The input signal frequencies are specified relative to this centre frequency.

Initially the input signal consists of a tone at +73 MHz, and a 1 MHz wide signal at +28 MHz. This scenario is denoted 'Time 0'. An additional tone at -19 MHz is added to the input signal at 'Time 1'. At 'Time 2' a tone at -24 MHz is added to the input. Lastly, at 'Time 3' the 1 MHz wide signal at +28 MHz is moved to -60 MHz. The output PSD's with the digital filter turned off and on is shown in Fig. 7.2.3. This data is exported from the signal analyzer. The PSD's in the right-side column represent the output after the digital filter has converged. The purpose of this figure is to show the adaptive algorithm is operating properly with each input scenario.

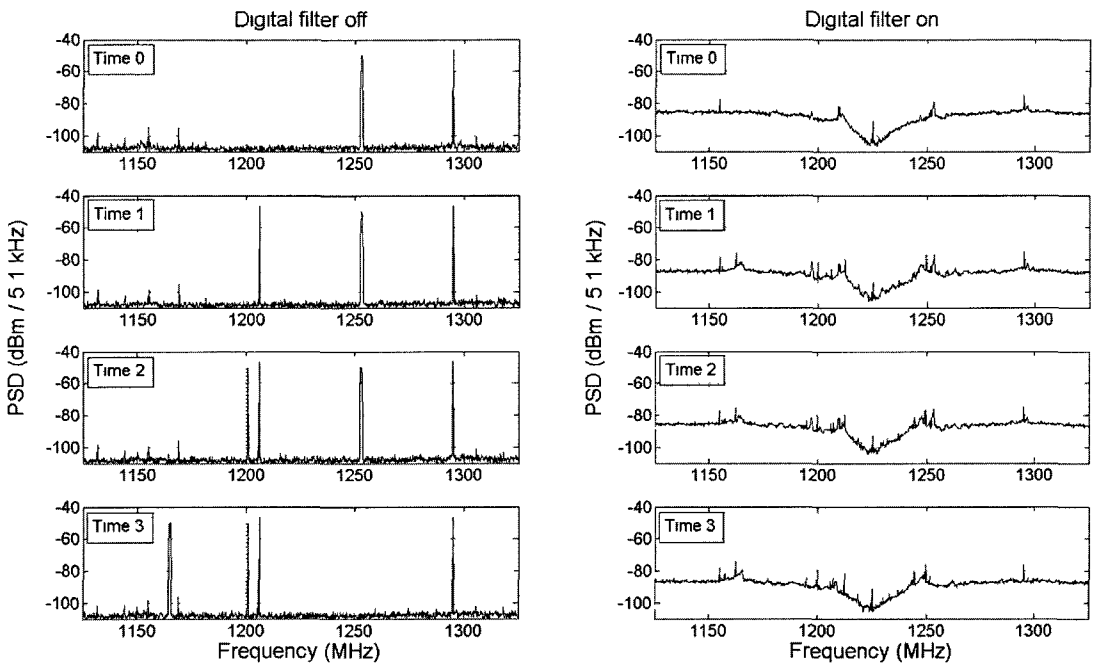

Figure 7.2.3 – Output PSD's at different times with digital filter off and on

The behaviour of the adaptive FIR filter taps during the changes to the input signal is shown in Fig. 7.2.4. The two plots show the evolution of the real and imaginary components of the 32 filter taps. The input signal changes were approximately uniformly spaced in time. The change corresponding to 'Time 1' occurs near 10 on the time scale, 'Time 2' near 23, and 'Time 3' near 35. After each change to the input signal, the taps change and converge to new values. The four regions where the taps are in steady-state correspond to the output PSD's in the right-side column of Fig. 7.2.3.

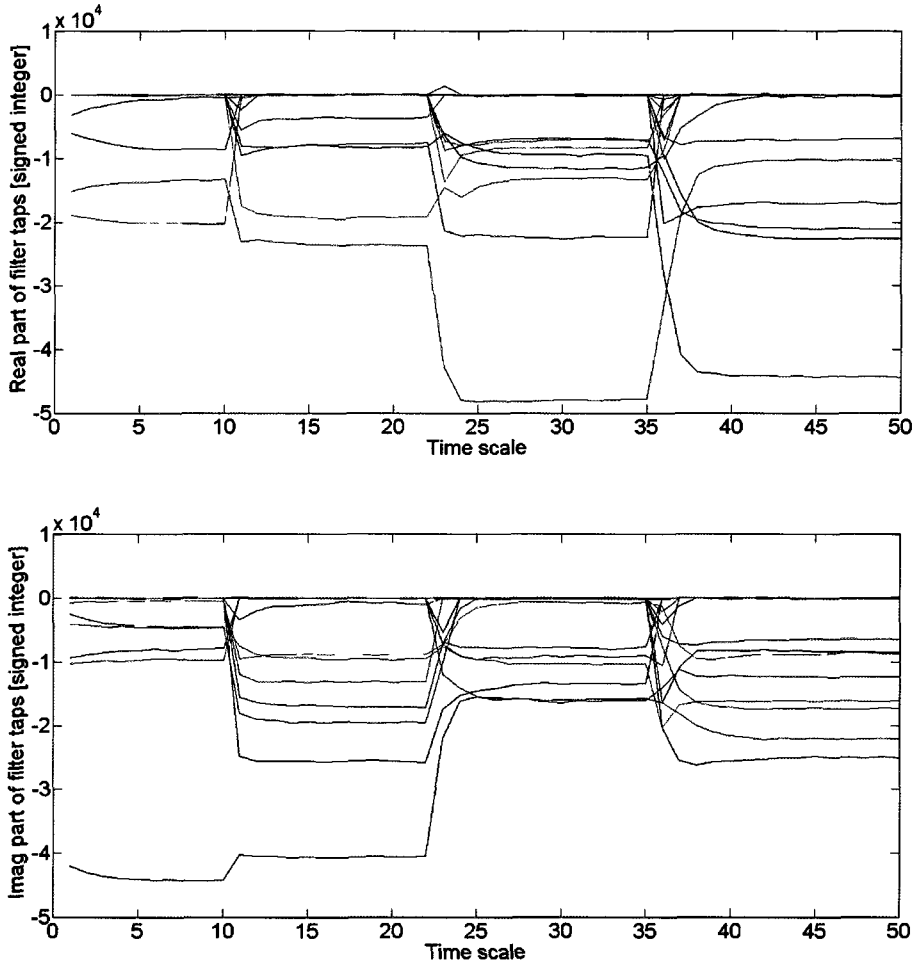

Figure 7.2.4 - Evolution of filter taps in time a) real components b) imag components

These results demonstrate successful adaptation of the hybrid RF-DSP FF filter during changes to the input signal.

7.2.3 Near Band Attenuation

The BSF at the end of the lower FF path has a large attenuation bandwidth. The large BSF attenuation can degrade the attenuation performance of the hybrid RF-DSP FF filter with regards to input signals near the filter's passband. Although these signals lie outside the passband, they are still significantly attenuated by the BSF, hence require the adaptive

filter to compensate. A digital IIR stage is introduced in Section 3.1.3 to perform this compensation, such that the adaptive FIR filter can be less constrained. The digital IIR transfer function consists of one set of complex conjugate poles and a zero at z=1 (in the z-domain). The magnitude response of the IIR stage is shown in Fig. 7.2.5 for three cases. The pole locations are different between the three cases. The digital IIR stage in the FPGA requires three extra registers in the main path between the ADC1's and DAC1's. These registers result in an additional 15 ns of group delay to the lower FF path.

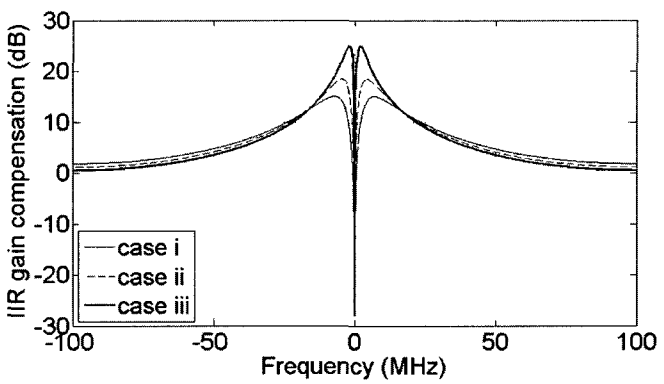

Figure 7.2.5 - Gain compensation from IIR stage in FPGA

The hybrid RF-DSP FF filter is operated at 1000 MHz for the measurements in this section. A single input signal is applied to the filter, which has a 1 MHz bandwidth and a frequency offset that is varied from +2 to +15 MHz, relative to 1000 MHz. The average power over the 1 MHz bandwidth is measured with the digital filter turned off and on. The difference between the off and on results represents the FF attenuation. The attenuation is shown in Fig. 7.2.6 for three different filter configurations:

- IIR off: IIR stage performs no filtering.
- IIR on: IIR stage has magnitude response corresponding to case iii in Fig. 7.2.5.

- RF BSF off, IIR off: Vector modulator in RF BSF is set to maximum attenuation, and the IIR stage performs no filtering.

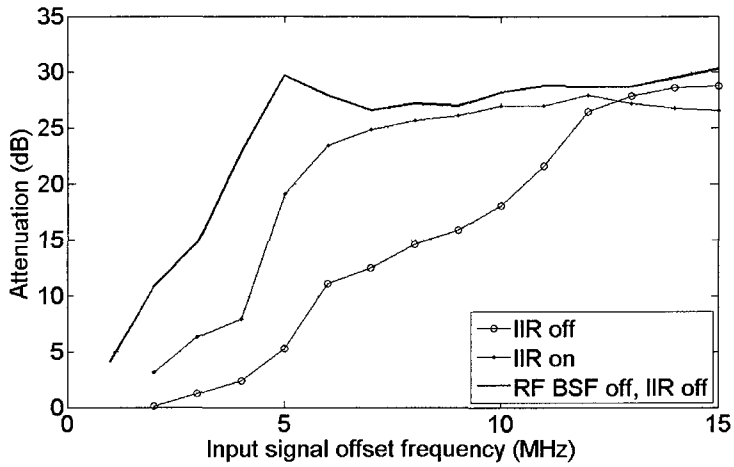

Figure 7.2.6 – Attenuation of signals near passband

The results in Fig. 7.2.6 show the IIR stage does boost attenuation between 5 and 11 MHz, with the largest improvement at 8 MHz. The large difference between the RF-BSF off curve and the IIR off curve proves that the RF-BSF causes the attenuation degradation near the passband.

7.3 Wideband Inputs

This section demonstrates the attenuation performance of a single input with different bandwidths. To keep the measurements consistent, the interferer is always located at 971 MHz, which is 21 MHz above the centre frequency of 950 MHz.

The output PSD's with the digital filter off and on are recorded by the signal analyzer for one of the test cases reported in the following sub-sections, and are shown in Fig. 7.3.1and 7.3.2. These figures show the input signal at 971 MHz has a finite bandwidth, and is successfully attenuated by over 20 dB when the filter is turned on.

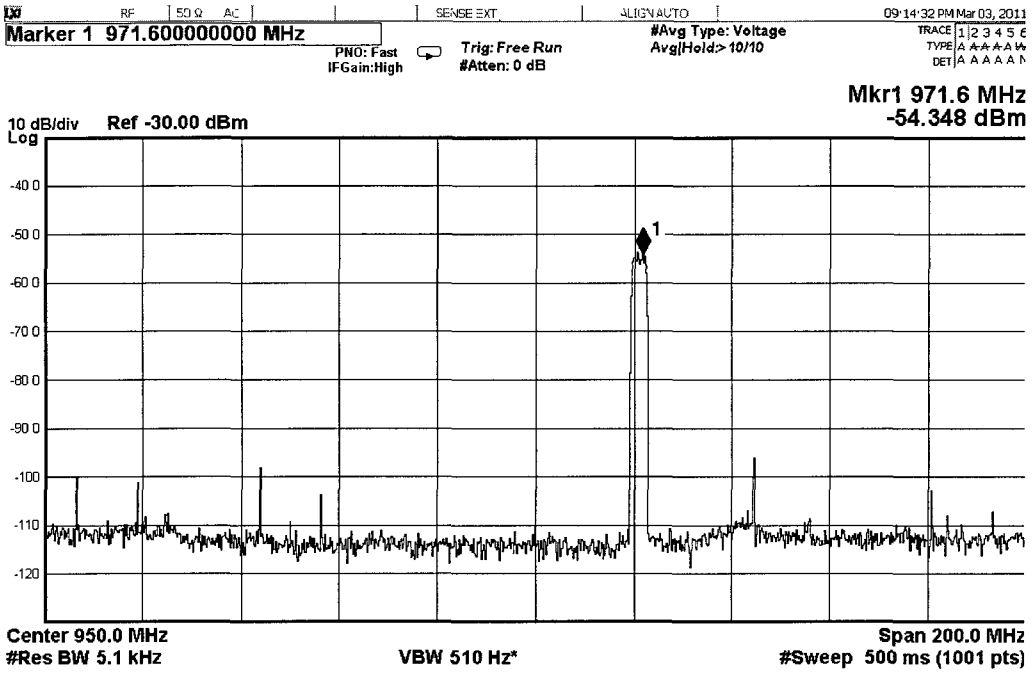

Figure 7.3.1 – Output PSD with no FPGA output

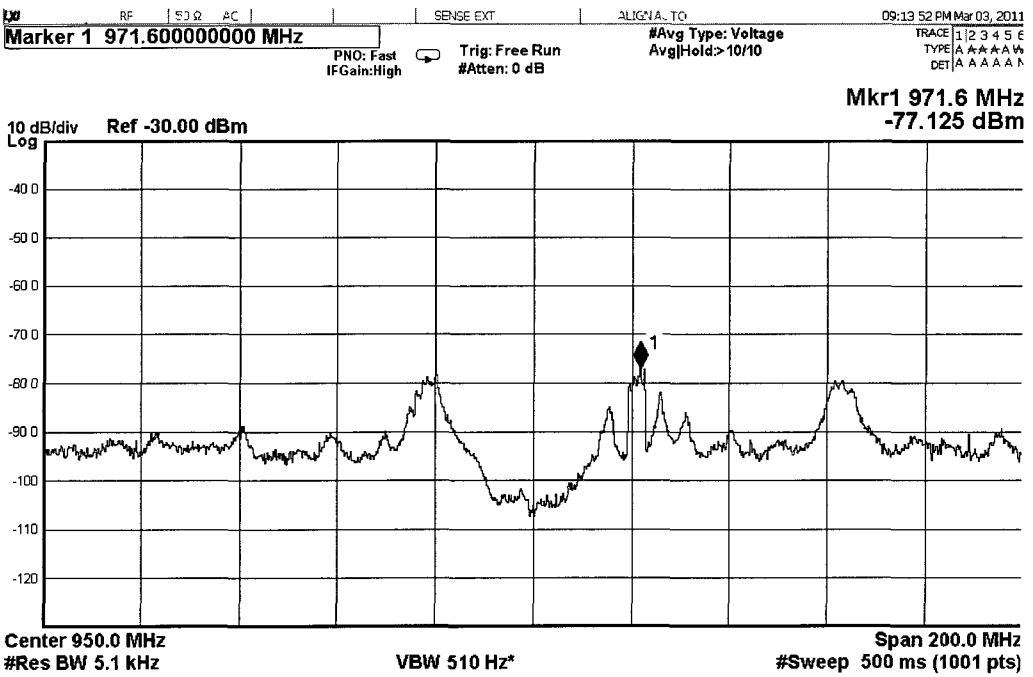

Figure 7.3.2 – Output PSD with filter in normal operating mode

The attenuation values reported in the following sub-sections are based on the average power in the specified bandwidth centered at 971 MHz. The PSD's in Fig. 7.3.3 and 7.3.4 correspond to those from Fig. 7.3.1 and 7.3.2, but are zoomed in at 971 MHz. The average power in a specified bandwidth is displayed near the top right corner of these figures, denoted 'Band Power.' The band is determined by the vertical lines that span out from marker 1. The difference between the band powers in the two figures represents the amount of FF attenuation caused by the lower FF path.

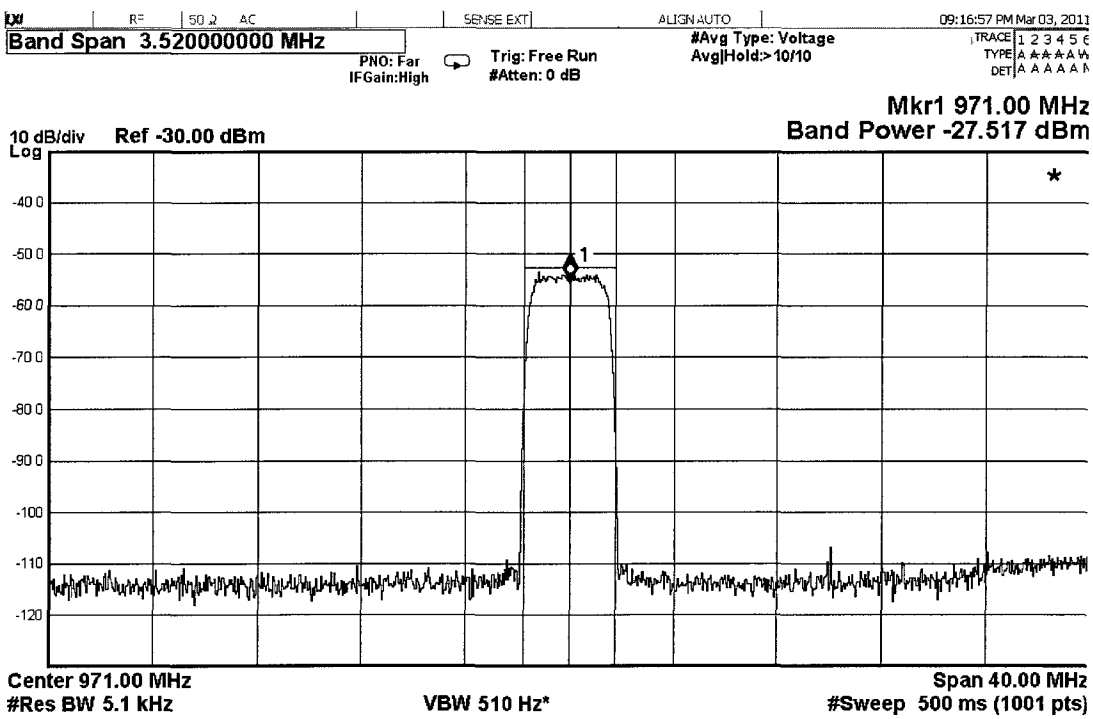

Figure 7.3.3 – Zoomed in PSD trace from Fig. 7.3.1

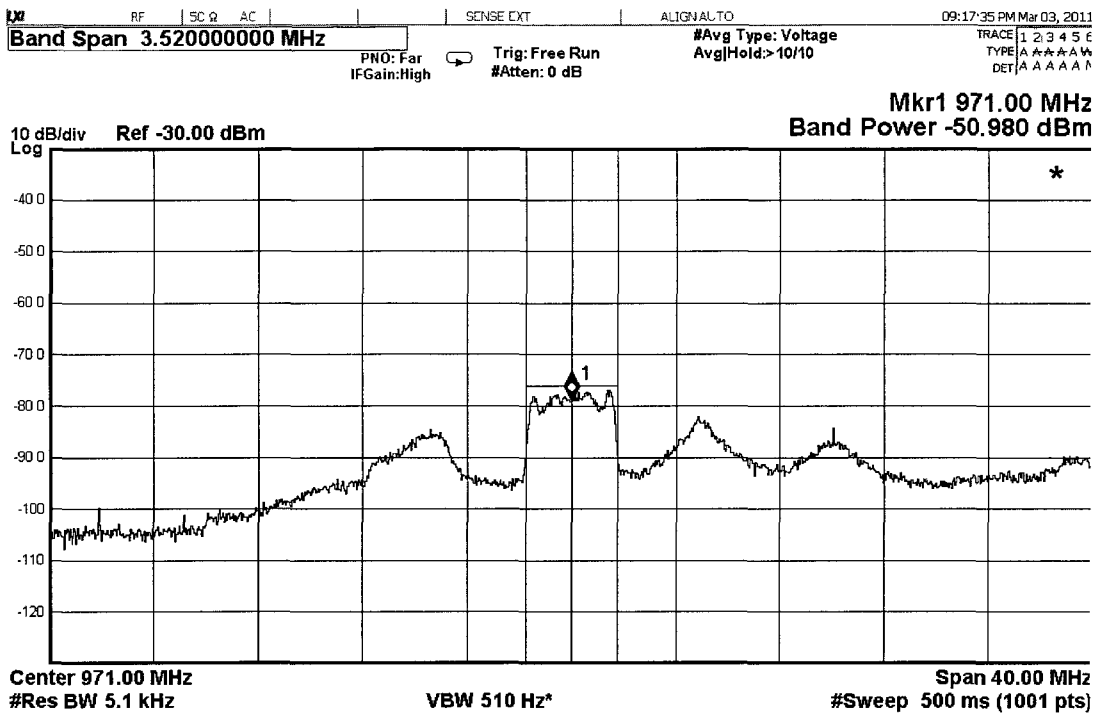

Figure 7.3.4 – Zoomed in PSD trace from Fig. 7.3.2

The FF attenuation is measured for several different filter configurations. The filter configurations differ by: the number of FIR filter taps, the block size for the FXBLMS with discard algorithm and the group delay difference between the paths. The bandwidth of the input signal is varied from 1.1 to 8.4 MHz for each configuration.

7.3.1 Vary Number of Taps

Three different digital filter configurations are compared, each with a different number of filter taps. The number of complex filter taps is 32, 48 and 64. The block size of the adaptive algorithm is kept fixed at 128. The FF attenuation caused by the lower FF path is shown in Fig. 7.3.5 for the 3 different configurations. The measurements are performed for 7 different input signals, each with a different bandwidth.

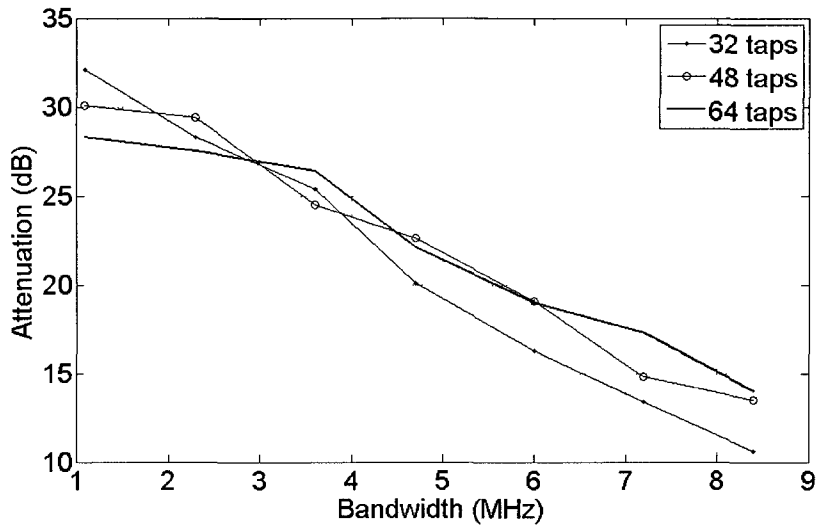

Figure 7.3.5 – FF attenuation of wideband input signals with different filter sizes

The curves in Fig. 7.3.5 show large attenuation for small bandwidths, but the attenuation degrades for large bandwidths. This characteristic is consistent with a FF system with a large group delay mismatch between the FF paths. There are only a few decibels difference between the curves in Fig. 7.3.5. This small difference shows that increasing the number of taps does not necessarily improve performance. The digital filter is supposed to converge to the Wiener filter. In general, adding more taps permits an adaptive filter to better fit the Wiener filter frequency response. But, arbitrarily increasing the number of taps may only result in a slightly better solution. This effect is seen in Fig. 7.3.5, since increasing the number of taps past 32 does not provide significantly more attenuation.

7.3.2 Vary Algorithm Block Size

The FXBLMS with discard algorithm block size determines how many data points are used to estimate the gradient for adaptation. More data points provide a more accurate

gradient estimate, which minimizes the misadjustment. The misadjustment is the difference between the steady-state solution and the optimum Wiener solution.

The block size corresponds to the number of additions performed within an accumulator per update, hence are represented by 'adds' in Fig. 7.3.6. The number of adds is varied between 128 and 512. The number of FIR filter taps is held constant at 32 for each measurement. The resulting FF attenuation versus input signal bandwidth is shown in Fig. 7.3.6.

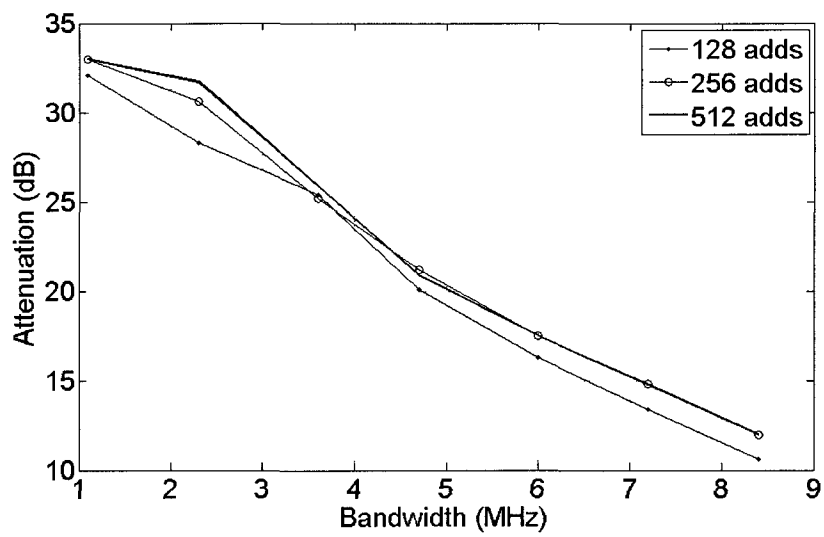

Figure 7.3.6 – FF attenuation of wideband input signals with different LMS block sizes

As with increasing the filter taps, the results in Fig. 7.3.6 show a very small improvement by increasing the LMS block size. The small improvement shows the steady-state solution of the '128 adds' case is pretty close to the Wiener solution, so increasing the block size past 128 does not bring significant improvement.

7.3.3 Vary Group Delay

The last parameter varied with a single wideband input is the group delay difference between the FF paths. The group delay of the lower FF path, excluding the registers within the FPGA, is reported in Section 6.2.7. The registers within the FPGA increase the group delay between the ADC1's and DAC1's. The number of registers within the FPGA is equal to four for the results reported in the two previous sections, excluding the IIR stage results. This is the fewest number of registers allowable for timing reasons. The digital clock runs at 200 MHz, so these registers contribute 20 ns of group delay to the lower FF path.

The wideband attenuation of three filter configurations is measured and the results are shown in Fig. 7.3.7. In this figure the 'no extra delay' curve represents the case with four registers in the FPGA. Two extra registers are added between the ADC1's and DAC1's in the FPGA for the 'Delay added to lower path' curve. The other curve is obtained by adding 7.8 metres of RG400 coax to the upper FF path. This length of cable has approximately 37 ns of group delay. Also for this configuration, the number of registers in the FPGA is equal to the minimum value of four. In all three configurations the number of filter taps is 32, and the LMS block size is 128.

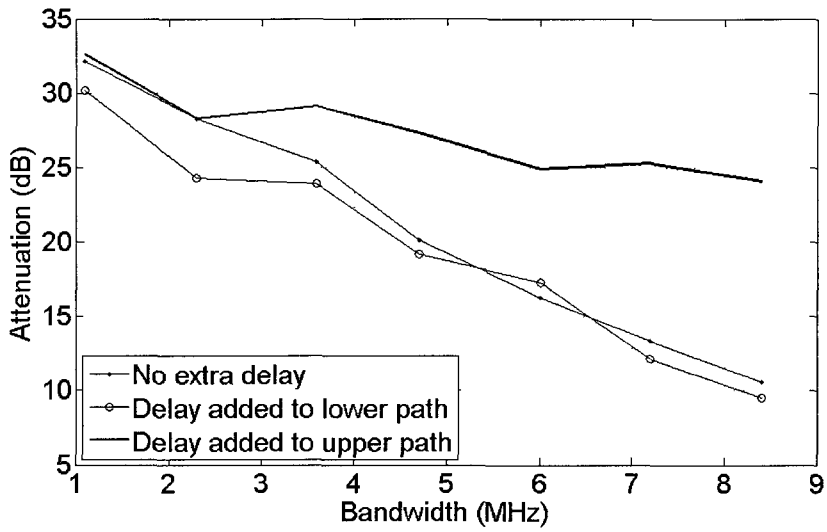

Figure 7.3.7 – FF attenuation of wideband input signals with different group delay

The results of Fig. 7.3.7 show the additional registers in the lower FF path slightly degrade the attenuation performance. The difference between these curves would be larger if the group delay mismatch between the two FF path in the 'no extra delay' case was much smaller. The delay added to the upper path shows a significant improvement in the attenuation performance for a wideband input. This result is consistent with a FF architecture where the group delay mismatch between the two paths is decreased. This improved case shows the hybrid RF-DSP FF filter can adequately attenuate wideband input signals when group delay is added to the upper FF path.

7.4 Noise Analysis Verification

The output noise PSD expression is derived in Section 4.1. This expression is checked against measured results for several filter realizations. In the following measurements the BSF is removed due to the poor linear performance of the vector modulator. Also, $C1$ is replaced with a 20 dB directional coupler. The output noise is larger without the BSF,

thereby permitting a more accurate measurement. An accurate measurement serves the greater purpose of this section, which is to validate the merger of RF and digital noise analyses.

7.4.1 Expected Output Noise

The component values required for the output noise expression are listed in the four tables in Section 6.2.9. These components values are specific to the hardware prototype. The component values from these four tables are used to determine the magnitude and noise factor expressions required by (4.1.1). The expression for each term is derived in Section 4.1. The numerical value for each term, along with the corresponding equation number from Section 4.1 is summarized below in Table 7.2. The noise factor terms for *RF*1 and *RF*2 in Table 7.2 do not include the noise due to mixing of the wideband phase noise with the input signals. This noise is included in the expected noise results in the following section.

Table 7.2 – Noise and magnitude terms for (4.1.1)

Term	Derived equation	Numerical result (linear)
F_{RF1}	(4.1.2)	408.6
$\|H_{RF1}\|$	(4.1.2)	162.2
F_{RF2}	(4.1.3)	194.7
$\|H_{RF2}\|$	(4.1.3)	1.4
F_{DIG}	(4.1.28)	$3.16\times10^{5}+1.72\times10^{9}/\|H_{DSP}\|^2$
$\|H_{DIG}\|$	(4.1.6)	$6.87\times10^{-11}\|H_{DSP}\|^2$

7.4.2 Measured Versus Expected Output Noise

The output noise PSD of the hardware prototype is measured in a 5 MHz passband for different filter realizations and input signals. The output noise PSD is averaged over the passband, excluding a small bandwidth centered at the centre frequency. The small

bandwidth at the centre frequency is omitted so the noise measurement does not include LO-RF leakage from the modulator. This detail is shown in Fig. 7.4.1, which is a zoomed-in screen capture from the signal analyzer for one of the noise PSD measurements.

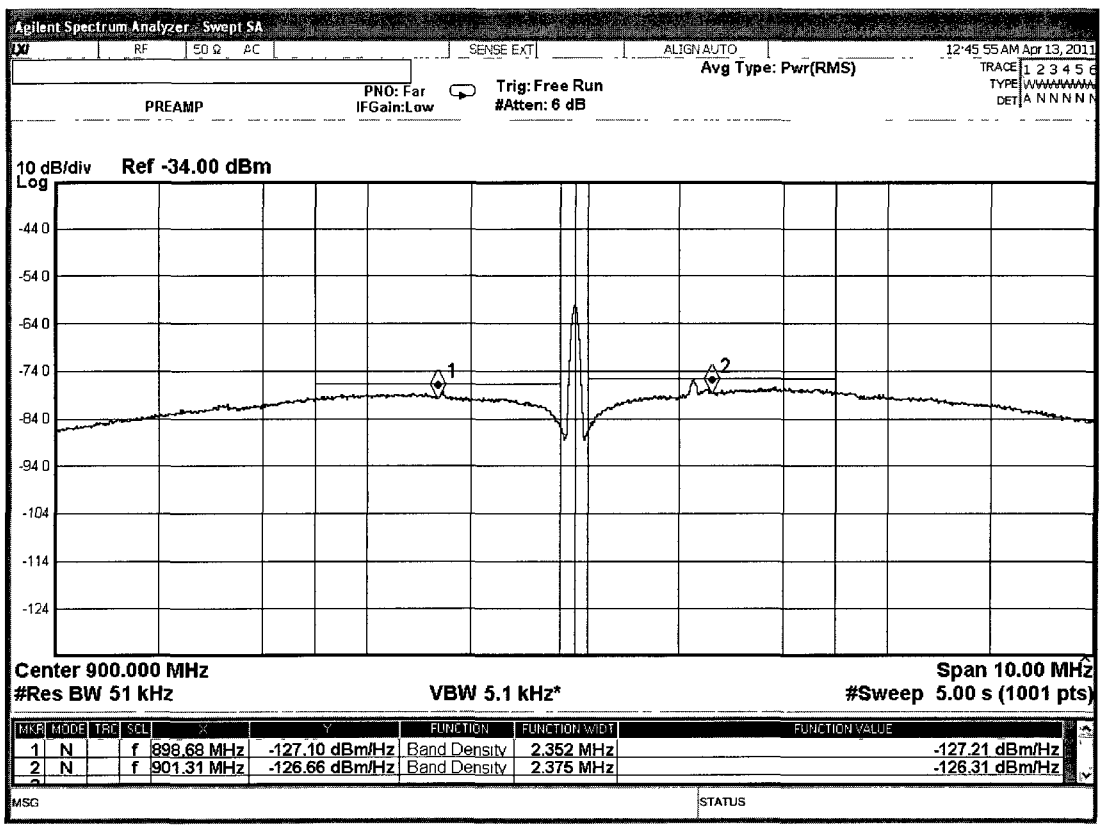

Figure 7.4.1 – Average noise PSD measurement over 5 MHz RF passband at 900 MHz

The derived noise expression is compared with the measured noise PSD in the passband for two different input interference environments. For each of these two cases the number of bits used by the ADC's and DAC's is varied. For each variation the average PSD in

the passband, and the average magnitude response of the digital filter are recorded. The centre frequency is 900 MHz for both cases.

The output PSD's for both cases with the digital components turned off are shown in the top row of Fig. 7.4.2. Significant interferer attenuation can be seen for both cases in the figures in the bottom row, which is the measured PSD's with the filter turned on. The details of the input interferers are summarized in Table 7.3.

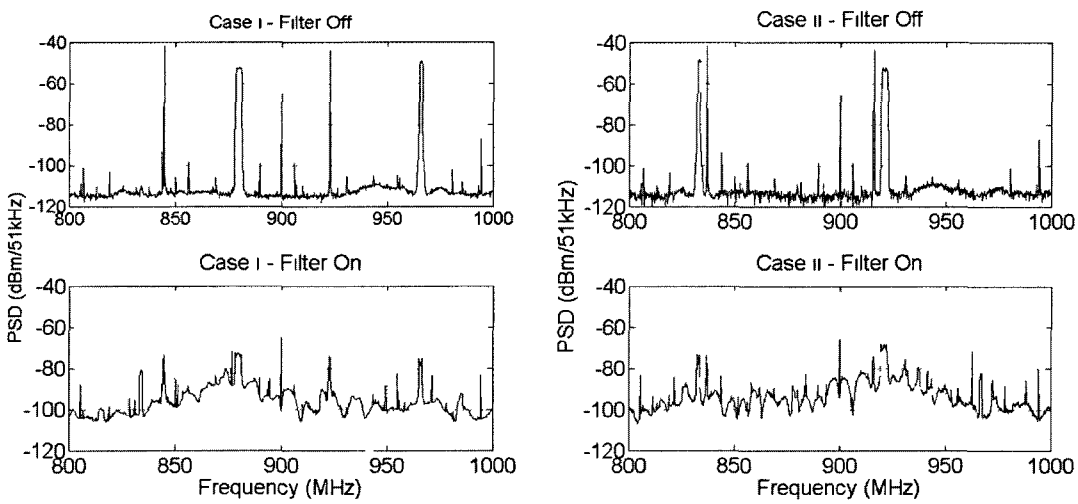

Figure 7.4.2 – Output PSD with filter on and off for both cases

Table 7.3 – Input interferers for both cases

Interferer ID	Frequency (MHz)		Bandwidth (MHz)		Power (dBm)	
	case i	case ii	case i	case ii	case i	case ii
1	845	833	CW	1.5	-32.2	-31.8
2	880	837	3.5	CW	-31.8	-32.2
3	923	916	CW	CW	-34.7	-34.5
4	966	921	1.5	3.5	-32	-32.4

The expected SNR's of the data converters are modified from those in Table 6.5 when the number of bits used is less than the nominal 14. The modification accounts for the additional quantization noise. The measured and expected noise results are shown in Fig.

7.4.3 for case i, and Fig. 7.4.4 for case ii. A null constraint is applied at DC within the digital filter for the measurements that vary the DAC bits. This constraint is discussed in Section 8.2.2, and is applied to reduce the output noise contribution from the ADC1's and *RF*1 chain.

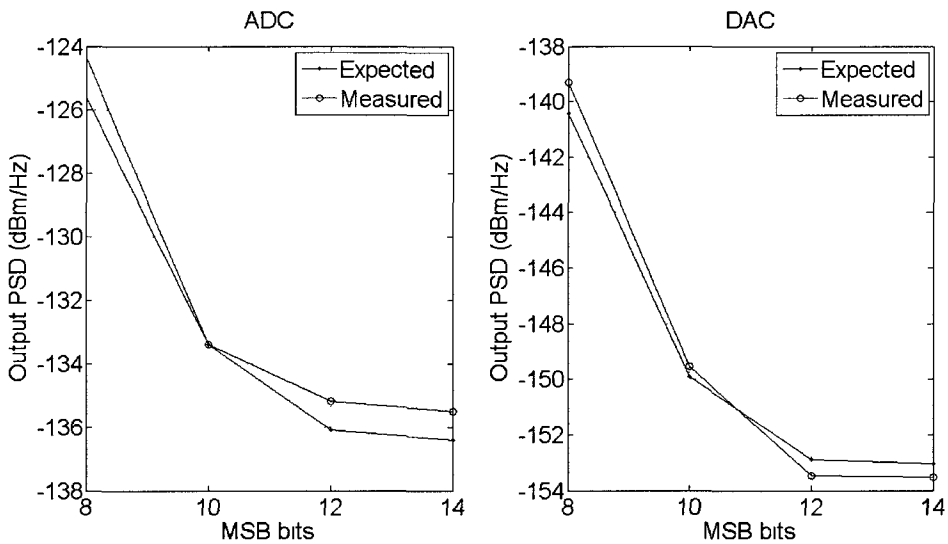

Figure 7.4.3 – Output noise PSD for case i a) various ADC1 bit widths b) various DAC1 bit widths

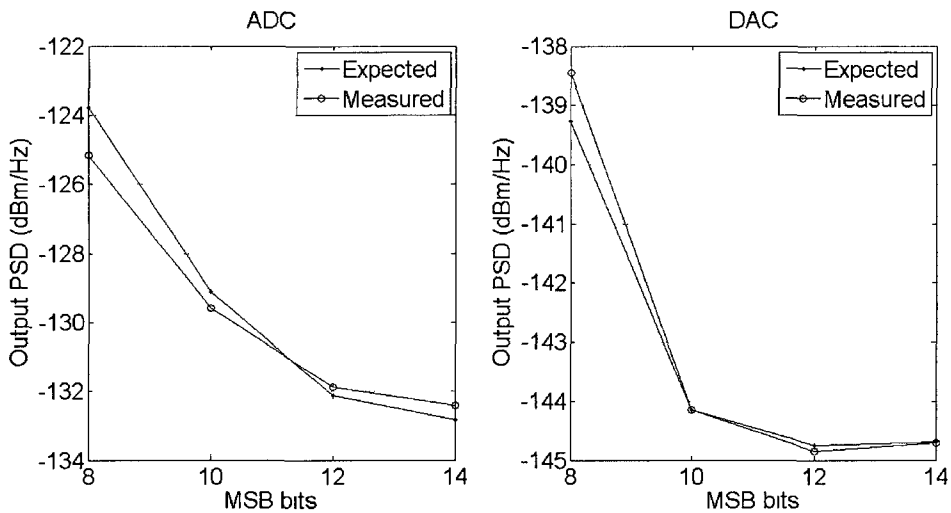

Figure 7.4.4 – Output noise PSD for case i a) various ADC1 bit widths b) various DAC1 bit widths

The output noise PSD depends on the average magnitude response of the digital filter over the passband. This response does not significantly vary when the number of bits used by the data converter changes. The average magnitude response of the digital filter in the passband is summarized in Table 7.4 for both cases. In this table $|H_{DSP}(f_p)|$ is significantly lower when the DAC bits are varied. This lower value is due to the DC null constraint mentioned above.

Table 7.4 – Average magnitude response of digital filter in noise measurements

Measurement	$\|H_{DSP}(f_p)\|$ (dB)	
	case i	case ii
Vary ADC bits	109.1	112.4
Vary DAC bits	88.7	100.1

The curves in Fig. 7.4.3 and Fig. 7.4.4 show a good match between predicted and measured performance. The output noise contribution from *RF*1, *RF*2, the ADC1's and DAC1's are shown in Fig. 7.4.5 for the same data points on the 'Expected' curve in Fig.

7.4.3. These curves are based on the derived expressions. The output noise is limited in Fig. 7.4.5a by the ADC1's and *RF*1 noise contributions. The null at DAC reduces these noise contributions in 7.4.5b, which is limited by the DAC1's. The *RF*1 and ADC1 curves almost overlap in this figure.

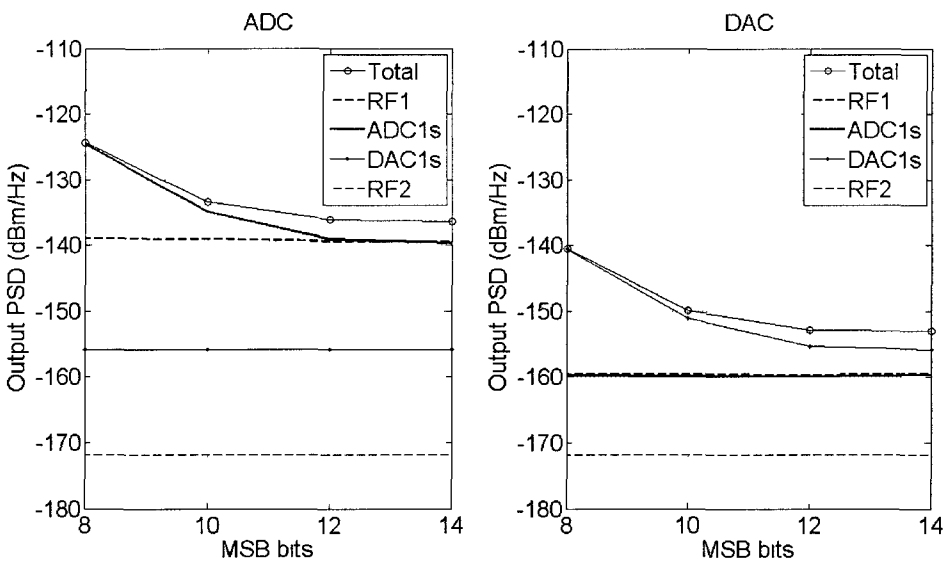

Figure 7.4.5 - Output noise PSD components for case i a) various ADC1 bit widths b) various DAC1 bit widths

This section has shown the passband noise expression derived in Section 4.1 correlates well with measurements. This correlation was observed without the BSF connected. The closed-form noise expression is useful for fixing design variables without resorting to time consuming simulations.

7.5 Prototype Performance Summary

The measured results in this chapter verify the hybrid RF-DSP FF filter hardware prototype: is frequency agile, can handle a dynamic RF environment at the input, and can

attenuate narrowband and wideband interferers. Based on the different filter configurations measured, it appears the most significant performance improvement is obtained by minimizing the group delay mismatch between the two FF paths.

8 Conclusions and Discussions

Observations made during the theoretical analyses, simulations and measurements have generated ideas for improvement of the hardware prototype and of the general analyses of the filter. These ideas are discussed below.

8.1 First Order Approximations

Some generalizations regarding the dynamic range and group delay are made in this section. The results are useful for first order approximations.

8.1.1 Dynamic Range

The dynamic range of the hybrid RF-DSP FF filter is limited to that of the FF path containing the digital components. The noise analysis from Section 4.1 showed the noise contribution of the components preceding the DAC1's can be insignificant if the digital filter is constrained to exhibit notches over the passband. One set of conditions were determined to achieve a 1 dB noise figure for the hybrid RF-DSP FF filter in Section 4.1.4. The noise due to the DAC's dominated the noise factor (F_{FF}) from the components in the lower FF path, as the noise factor allocations were: $F_{FF|RF1}$=0.01, $F_{FF|RF2}$=0.01, $F_{FF|ADC}$=0.001, $F_{FF|DAC}$=0.15. A typical PSD resulting from a single interferer is shown in Fig. 8.1.1 between the DAC's and $C2$. Only the noise generated by the DAC's is shown in the figure.

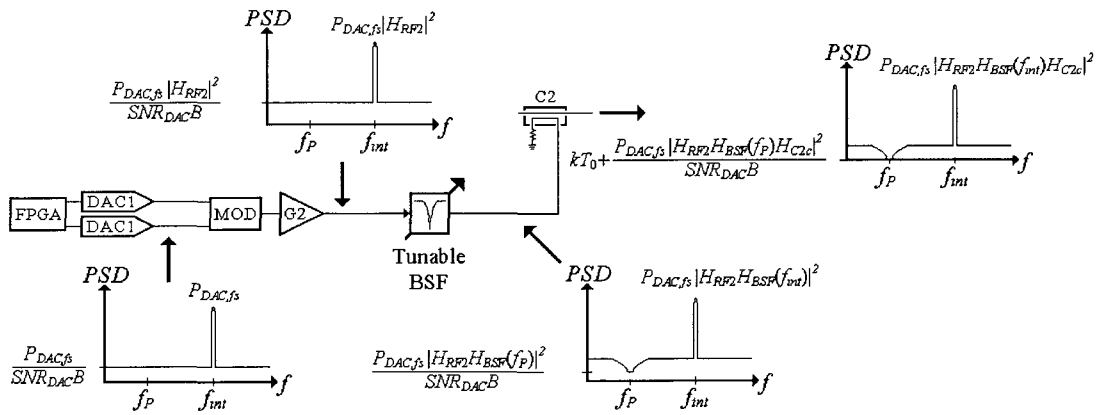

Figure 8.1.1 – Signal and noise scenario for first order approximation

A first order approximation is made relating the noise factor to the largest interferer power that can be attenuated. In Fig. 8.1.1 the filter's output signal at f_{int} shows that the largest interferer power than can be attenuated by the hybrid RF-DSP FF filter is:

$$P_{\text{int,max}} = \frac{P_{DAC,fs}}{PAPR_{\text{int}}}\left|H_{RF2}H_{BSF}(f_{\text{int}})H_{C2c}\right|^2 \tag{8.1.1}$$

The term $PAPR_{int}$ represents the peak to average power ratio of the interferer, which decreases the maximum interferer power than can be attenuated.

An approximation of the noise figure is made assuming the noise from the bottom FF path is dominated by the DAC's:

$$\begin{aligned} F_{FF} &= \frac{kT_{out}}{kT_0\left|H_{FF}(f_{PB})\right|^2} \\ &\approx \frac{1}{\left|H_{FF}(f_{PB})\right|^2} + \frac{P_{DAC,fs}}{kT_0 SNR_{DAC}B}\frac{\left|H_{RF2}H_{BSF}(f_{PB})H_{C2c}\right|^2}{\left(1-\left|H_{C1t}\right|^2\right)\left(1-\left|H_{C2t}\right|^2\right)} \end{aligned} \tag{8.1.2}$$

The maximum input power (P_{in}) that the hybrid RF-DSP FF can handle is calculated by de-embedding (7) to the filter's RF input. The expression for P_{in} is substituted into the noise factor approximation in (8.1.2) to get:

$$F_{FF}(f_p) \approx \frac{1}{\left|H_{FF}(f_p)\right|^2} + \frac{P_{in} PAPR_{\text{int}}}{kT_0 SNR_{DAC} B} \left| \frac{H_{BSF}(f_p)}{H_{BSF}(f_{\text{int}})} \right|^2 \tag{8.1.3}$$

The expression in (8.1.3) is the first order approximation that relates the noise factor (F_{FF}) to the maximum input power (P_{in}). In this equation the term $|H_{FF}(f_p)|$ is the magnitude response of the thru ports of the three directional couplers, and B is half the digital clock frequency. The magnitude of the BSF in (8.1.3) is the combined response of the RF BSF and digital gain compensation stage, which was introduced in Section 3.1.3. The approximation in (8.1.3) is useful to determine the required SNR for the DAC's, and if a digital gain compensation stage is required for the BSF given a specific P_{in}, $PAPR_{int}$, $F_{FF}(f_p)$, and RF operating bandwidth of the filter ($2B$). The approximation is valid when the RF and analog components are designed for low noise, and the digital filter is constrained to exhibit notches over the passband.

8.1.2 Group Delay

Attenuation performance of the hybrid RF-DSP FF filter was shown to strongly depend on the group delay mismatch between the two FF paths. Adding filter taps and increasing the amount of data used by the adaptive algorithm were found to only slightly improve performance. This behaviour indicates the digital filter has no problem converging to the

Wiener solution for a given group delay mismatch. But, the MMSE of the Wiener solution decreases with a decreasing group delay mismatch.

The z-spectrum at the output the hybrid RF-DSP FF filter is given for the case where the digital filter is equal to the Wiener solution in (8.1.4). This equation is derived from (4.2.5), (4.2.8), (4.2.9) and Fig. 4.2.3b. The z-spectrum of the input signal is $S_{xx}(z)$.

$$\begin{aligned}
S_{yy}(z) &= S_{xx}(z)H_{FF}(z)H_{FF}^{*}\left(\frac{1}{z^{*}}\right) \\
&= \sigma_u G_u(z)\frac{H_2(z)}{H_1(z)}\left\{1-\frac{z^{-D}}{G_u(z)}\frac{H_1'(z)}{H_2'(z)}\left[G_u(z)z^{D}\frac{H_2'(z)}{H_1'(z)}\right]_{+}\right\} \\
&\times\sigma_u G_u^{*}\left(\tfrac{1}{z^{*}}\right)\frac{H_2^{*}\left(\frac{1}{z^{*}}\right)}{H_1^{*}\left(\frac{1}{z^{*}}\right)}\left\{1-\frac{z^{-D}}{G_u^{*}\left(\frac{1}{z^{*}}\right)}\frac{H_1'^{*}\left(\frac{1}{z^{*}}\right)}{H_2'^{*}\left(\frac{1}{z^{*}}\right)}\left[G_u^{*}\left(\tfrac{1}{z^{*}}\right)z^{D}\frac{H_2'^{*}\left(\frac{1}{z^{*}}\right)}{H_1'^{*}\left(\frac{1}{z^{*}}\right)}\right]_{+}\right\} \\
&= \sigma_u\left\{G_u(z)\frac{H_2'(z)}{H_1'(z)}z^{D}-\left[G_u(z)\frac{H_2'(z)}{H_1'(z)}z^{D}\right]_{+}\right\} \\
&\times\sigma_u\left\{G_u^{*}\left(\tfrac{1}{z^{*}}\right)\frac{H_2'^{*}\left(\frac{1}{z^{*}}\right)}{H_1'^{*}\left(\frac{1}{z^{*}}\right)}z^{D}-\left[G_u^{*}\left(\tfrac{1}{z^{*}}\right)\frac{H_2'^{*}\left(\frac{1}{z^{*}}\right)}{H_1'^{*}\left(\frac{1}{z^{*}}\right)}z^{D}\right]_{+}\right\}
\end{aligned} \tag{8.1.4}$$

Recall from Section 4.2.2 that the value of D is the nominal delay from $H_2(z)/H_1(z)$. This value is the group delay mismatch between the FF paths in the discrete-time domain. A new term is introduced to simplify the expression in (8.1.4):

$$K(z)=\frac{H_2'(z)}{H_1'(z)}G_u(z) \tag{8.1.5}$$

Both the numerator and denominator terms from (8.1.5) are minimum phase, therefore $K(z)$ is minimum phase. This minimum phase property means than $k[n]$ is causal. The expression from (8.1.4) is simplified using (8.1.5):

$$S_{yy}(z) = \sigma_u^2 \left\{ K(z)z^D - \left[K(z)z^D \right]_+ \right\} \times \left\{ K^*(z)z^D - \left[K^*(z)z^D \right]_+ \right\} \tag{8.1.6}$$

The $\{K(z)z^D - [K(z)z^D]_+\}$ factor is converted into the discrete time domain in (8.1.7) to demonstrate the effect of group delay mismatch on attenuation bandwidth.

$$IZT\left\{ K(z)z^D - \left[K(z)z^D \right]_+ \right\} = k[n+D] - k[n+D]u[n] \tag{8.1.7}$$

If there was no group delay mismatch between the two FF paths, then D in (8.1.7) would be zero. If D=0, then since $k[n]$ is causal, the subtraction on the right side of (8.1.7) would also be zero (ie. $k[n]u[n]$=$k[n]$). For the D=0 case, large attenuation of the input signal occurs, and there is no limit on the attenuation bandwidth.

When D>0, then the first D terms in $k[n]$ are not subtracted in (8.1.7); therefore the subtraction on the right side yields a non-zero value. As D increases, so does the number of leading terms in $k[n]$ that are not subtracted in (8.1.7).

As a first order approximation, let $K(z)$ have a third-order low-pass Chebyshev type I response with the cut-off frequency denoted by f_c. This response is an approximation of a single input signal with an RF bandwidth of $2f_c$. The magnitude squared response from (8.1.7) is shown in Fig. 8.1.2 for different values of group delay mismatch (D), and input signal bandwidth. The frequencies in the legend refer to the RF bandwidth of the Chebyshev response, which corresponds to $2f_c$. The attenuation is proportional to the magnitude squared response, so this figure is useful in determining how much the attenuation performance will improve or degrade by changing the FF group delay mismatch.

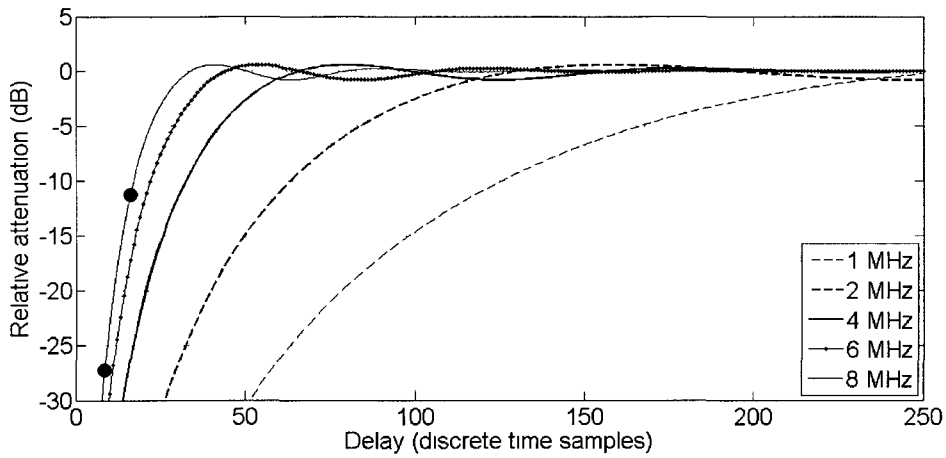

Figure 8.1.2 – Relative attenuation versus group delay mismatch with Wiener filter

The approximation made in this section was that $K(z)$ has a third-order Chebyshev type I lowpass response. The results are the curves in Fig. 8.1.2, which relate the relative attenuation to the group delay mismatch between the FF paths. In Section 7.3.3 the group delay mismatch was reduced by 37 ns, which resulted in a 15 dB improvement in attenuation for an input signal with an RF bandwidth of 8 MHz. The original group delay mismatch was approximately 80 ns (~60 ns at +21 MHz on 900 MHz curve in Fig. 6.2.10b, plus 20 ns from the four registers in the main path in the FPGA). This mismatch corresponds to D=16 for a 200 MHz clock. The mismatch reduction of 37 ns corresponds to decreasing D by ~7.5 time samples. These two points are marked on the 8 MHz curve in Fig. 8.1.2 with solid circles. These two points are separated by 16 dB. This predicted improvement is very close to the measured 15 dB improvement.

8.2 Improvements and Future Work

Several improvements are suggested in these sections that were not incorporated into the hardware prototype. These improvements can be incorporated when the hardware prototype evolves into an integrated solution.

8.2.1 Equalizing ADC's and DAC's

Since the hardware prototype is built with evaluation boards, connectorized components, and a lot of coaxial cable, there are variations between I and Q paths near the ADC1's, ADC2's and DAC1's. Amplitude and phase mismatches between these I and Q paths limit the attenuation performance of the hybrid RF-DSP FF filter. The effect of IQ mismatch is to generate an undesired image of the input signal symmetrically about DC in the complex baseband domain. The magnitude of the image decreases as the IQ mismatch decreases.

A solution to reduce the effects of IQ mismatch is to equalize the digital inputs or outputs of the data converters. The equalization corrects the mismatch, but will add some extra delay, thereby increasing the group delay mismatch between the RF and digital FF paths. Furthermore, some training technique is required to adjust the equalizer coefficients. Realistically, an integrated hardware solution can be specified to limit the IQ mismatches, and possibly eliminate the need for equalization. The effect of IQ mismatch is discussed for each of the three pairs of data converters.

ADC1's: An input signal to the hybrid RF-DSP FF filter with a frequency of $f_{LO}+f_{in}$ will have a desired signal component in the complex baseband domain at $+f_{in}$, and a smaller

undesired image at $-f_{in}$. The effect of IQ mismatch between DEMOD1 and the ADC1's is to create an additional input signal in the digital FF path that does not exist in the RF FF path. The image signal propagates through the digital filter and is sent out the DAC1's, which then shows up at the filter output. The RF output signal is used for adaptive control, so the image signal exists in the feedback path. Fortunately, the adaptive algorithm processes the images, so the digital filter is adapts to attenuate them. Attenuating the images is done by collocating them with notches in the digital filter. A problem arises when a real input signal is close to the image of another real input signal in the frequency domain.

ADC2's: Mismatch between the IQ paths between DEMOD2 and the ADC2's causes the adaptive algorithm to adapt as if there are additional signal components at the output of the filter. These additional signal components are in fact not at the RF output of the filter, and are due to the images of the ADC2's input signals. The input signals to the ADC2's decrease in magnitude as the digital filter converges; therefore the magnitudes of the image components also decrease. The image components are smaller than the actual attenuated signal components, so they do not significantly affect attenuation performance.

DAC1's: The image generated by the DAC1's shows up at the filter's RF output and cannot be attenuated. Measurements show that interferer signals can only be attenuated down to the level of this image. This limitation can be seen in the following figures:

- Fig. 6.3.9 – CW image at 974 MHz
- Fig. 7.3.2 –image at 929 MHz with 3.6 MHz bandwidth
- Fig. 7.4.2 (left side) – image at 834 MHz with 1.5 MHz bandwidth

- Fig. 7.4.2 (right side) – CW image at 963 MHz

This IQ mismatch would benefit most from equalization.

8.2.2 Notch Constraints on Digital Filter

The magnitude response of the digital filter in the passband does not impact the interferer attenuation performance, but does impact the noise figure in the passband. Notch constraints can be placed on the digital filter at multiple frequencies within the passband to reduce the output noise contributions from the *RF*1 chain and ADC1's. A technique to incorporate these constraints is to add the magnitude response of the digital filter at the desired notch locations to the algorithm's error signal. The magnitude response of the digital filter at DC is simply the summation of the filter taps. A notch constraint at DC could then be incorporated by adding the summation of the filter taps to the outputs of the ADC2's. This single notch constraint is incorporated into the hardware prototype for a single realization. The noise improvement from the signal analyzer is shown in Fig. 8.2.1.

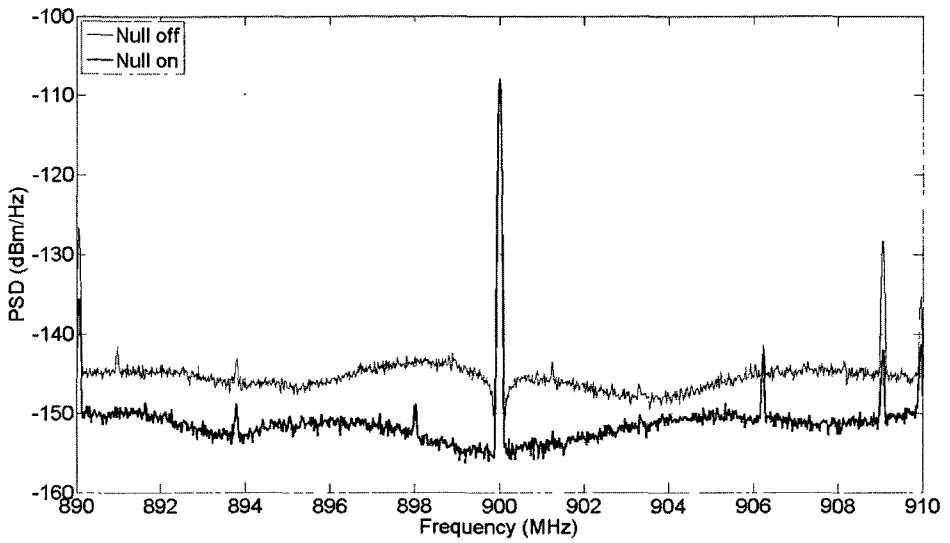

Figure 8.2.1 – Output PSD over passband with and without null constraint

The single realization includes three input signals: CW at 881 MHz, 1 MHz bandwidth at 924.28 MHz, and 1.5 MHz bandwidth at 957 MHz. The BSF is removed from the prototype for the measurements in Fig. 8.2.1. The centre frequency is 900 MHz, and the spike at 900 MHz is due to the LO-RF leakage of the modulator. Over a 5 MHz passband the output noise PSD is on average 8.5 dB lower with one null constraint. The magnitude response of the digital filter with and without the null constraint is shown in Fig. 8.2.2. The filter tap values are tapped from the FPGA using the QuartusII software. As required for FF attenuation, the two magnitude responses match at the three input signal frequencies: -19, +24.28 and + 57 MHz. The null for the 'Null on' curve can be seen at f=0.

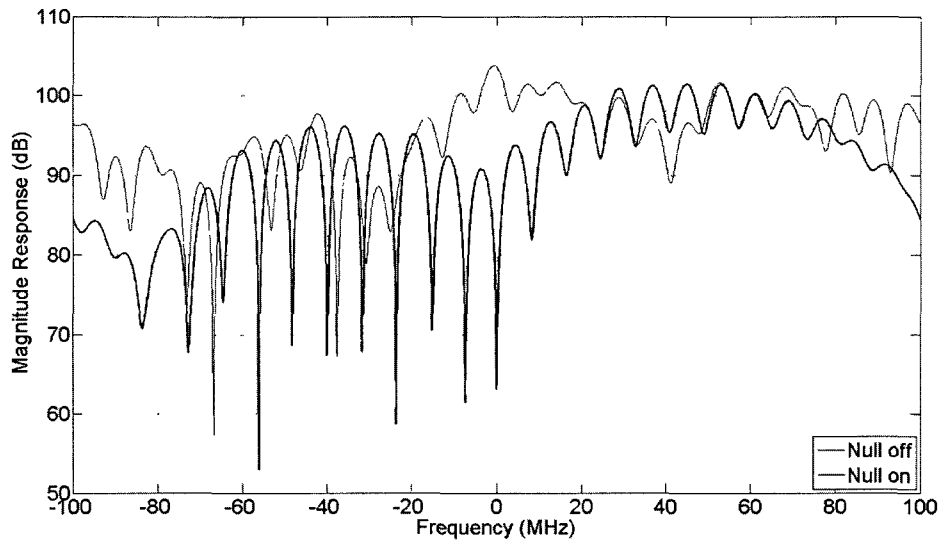

Figure 8.2.2 – Digital filter response with and without null constraint

The output PSD over the 200 MHz operating bandwidth is shown in Fig. 8.2.3 with the filter off, the filter on without the null constraint, and the filter on with the null constraint. The similarity between the bottom two PSD curves shows the attenuation performance is not significantly affected by the null constraint.

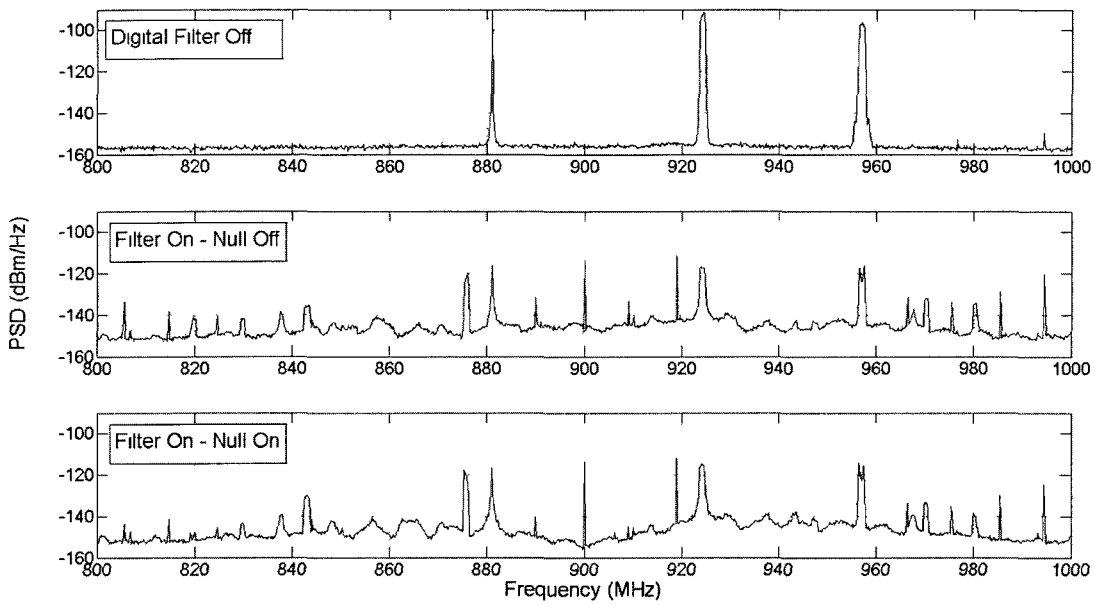

Figure 8.2.3 – Output PSD over operating bandwidth a) digital filter off b) filter on without null constraint c) filter on with null constraint

Multiple null constraints can be incorporated into the adaptive algorithm with the technique described above. These constraints improve the noise performance. But, the attenuation performance may degrade as the number of null constraints increase since each constraint removes a degree of freedom from the adaptive FIR filter.

8.2.3 Digital IIR stage

The digital BSF compensation, denoted 'IIR stage', was shown to improve near band attenuation in Section 7.2.3. The IIR stage was placed in the FPGA between the ADC1's and the adaptive FIR filter input. The $h_{23,est}[n]$ channel estimation algorithm would produce more accurate near band results if the IIR stage were placed between the adaptive FIR filter output and the DAC1's. The $h_{23}[n]$ channel would include this compensation stage, thereby making the channel's magnitude response flatter outside the

passband. A flatter frequency response could contribute towards a more accurate channel estimate, hence relaxing the stability criteria.

8.2.4 Negative Group Delay

Minimizing the group delay mismatch between the two FF paths was shown to improve attenuation bandwidth performance. Adding group delay to the RF path is one technique to minimize the group delay, but comes with the cost of a higher insertion loss in the RF path. Alternatively, a negative group delay (NGD) circuit can be placed into the FF path containing the digital components. This technique is applied to FF amplifier linearization to reduce the excess delay in the main path by 40% in [*47*]. Narrowband NGD solutions have been demonstrated in the literature at RF [*48*,*49*,*50*,*51*,*52*], analog baseband [*53*], and digital baseband [*54*]. In fact, the converged frequency response of the digital filter in the hybrid RF-DSP FF filter exhibits NGD at frequencies occupied by wideband input signals. This NGD is demonstrated in Fig. 8.2.4 and 8.2.5, which shows the group delay and magnitude responses from the hybrid RF-DSP FF filter's digital filter for a single input with a 6 MHz bandwidth. The input signal is at 942.4 MHz, which is +42.4 MHz relative to the filter's centre frequency of 900 MHz.

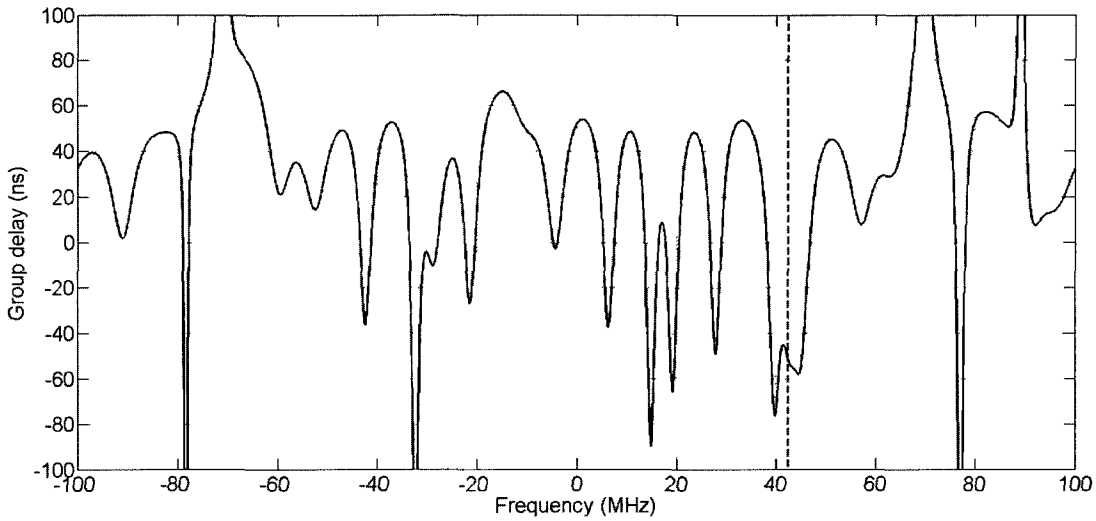

Figure 8.2.4 - Group delay response of single wideband input at +42.4 MHz

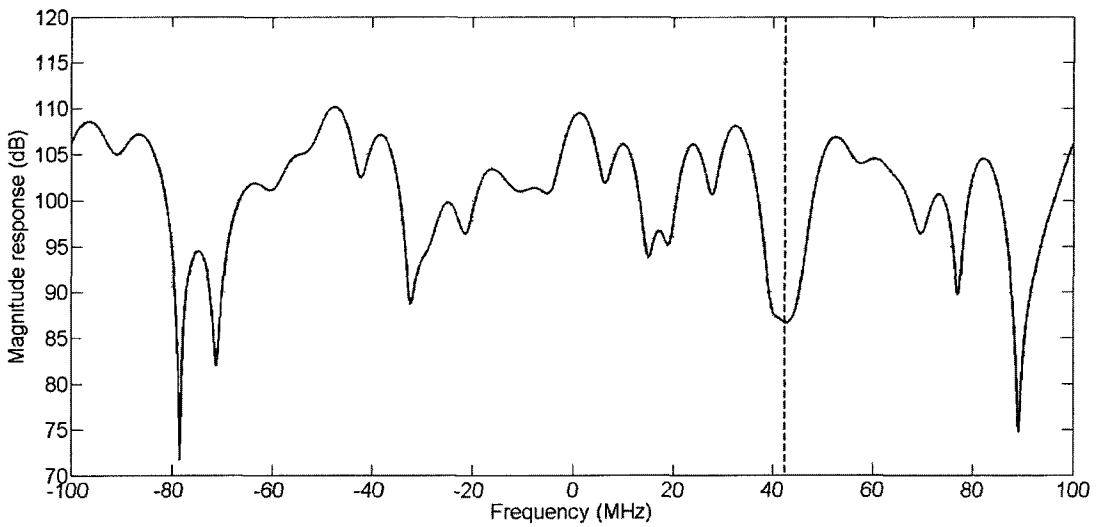

Figure 8.2.5 – Magnitude response of single wideband input at +42.4 MHz

An NGD circuit would provide some of the negative group delay that the adaptive FIR filter will be known to exhibit at frequencies occupied by wideband input signals.

Significant NGD is not attainable over the entire Nyquist bandwidth; so the NGD would need to be tuned to the wideband input's frequency, which must be known a priori.

8.3 Discussion Summary

This chapter proposed some improvements for the hybrid RF-DSP FF filter, and made some generalizations about the analysis. The proposed improvements and generalizations are based on the analyses, simulations and measured results from the previous chapters.

8.4 Conclusion

Analyses, simulations and measurements have been reported for a new type of frequency agile RF filter. Successful operation of the hardware prototype was demonstrated from 800 to 1800 MHz, and for filter orders between 32 and 64. The filter exhibits an all-zero transfer function, which is best suited for providing multiple notches or stop-bands. Wideband attenuation was also demonstrated, but required the addition of a delay line to the RF FF path to attain nearly 25 dB attenuation over bandwidths as large as 8 MHz.

8.4.1 List of Contributions

The contributions made during this research are listed below:

1. A frequency agile RF feed-forward architecture with a digital filter in one of the feed-forward paths with applications for frequency agile RF filtering and duplexer isolation improvement.

2. Noise and attenuation analyses of this architecture that combines the effects of the RF and digital systems. Also, first-order approximation expressions for dynamic range and attenuation bandwidth performance.
3. Addition of an IIR stage within the digital filter to improve near-band attenuation performance, and provide the means to reconfigure the passband bandwidth.
4. Addition of a negative group delay circuit in the feed-forward path containing the digital filter to improve wideband attenuation performance.
5. Modification to active noise control adaptive algorithm for the hybrid RF-DSP FF filter. The modification allows for a simpler relationship between the stability criteria and the allowable estimation errors of the $h_{23}[n]$ channel.
6. Addition of notch constraints in the digital filter to reduce the output noise due to the *RF/Analog* 1 chain and primary ADC's.

Two patent applications were filed by Nortel Networks with the Canadian Intellectual Property Office based on the contributions. The author of this thesis is the single author on both of these patent applications [***5***,***6***].

References

[1] S. Haykin, "Cognitive radio: brain-empowered wireless communications," *IEEE Journal on Sel. Areas in Comms.*, vol. 23, no. 2, pp. 201-220, Feb. 2005.

[2] "Promoting more efficient use of spectrum through dynamic spectrum use technologies," FCC, Washington, D.C., Notice of inquiry FCC-10-198, 2010.

[3] T. Schilcher, "RF Applications in Digital Signal Processing," *Digital Signal Processing CAS*, p. 2490283, Jun. 2007.

[4] S. Elliott, *Signal Processing for Active Control.* San Diego, CA: Academic Press, 2001.

[5] Mark Wyville, Frequency agile filter using a digital filter and bandstop filtering, filing date: 26 Nov. 2009, publication date: 10 Jun. 2010, applic. no. PCT/CA2009/001721, WIPO publication no. WO2010/063096.

[6] Mark Wyville, Frequency agile filter using a digital filter and bandstop filtering, filing date: 26 Nov. 2009, pulication date: 10 Jun. 2010, applic. no. PCT/CA2009/001722, WIPO pulic. no. WO2010/063097.

[7] G. Fischer, "Next-generation base station radio frequency architecture," *Bell Labs Technical Journal*, vol. 12, no. 2, pp. 3-18, Dec. 2007.

[8] Herman J. Blinchikoff and Anatol I. Zverev, *Filtering in the Time and Frequency Domains.* New York, NY: Wiley, 1976, ch. 6.3.

[9] G. Rebeiz et al., "Tuning in to RF MEMS," *IEEE Microwave Magazine*, vol. 10, no. 6, pp. 55-72, Oct 2009.

[10] R.M. Young et al., "Low-loss bandpass RF filter using MEMS capacitance switches to acheive a one-octave tuning range and independently variable bandwidth," *IEEE MTT-S International*, vol. 3, pp. 1781-1784, Jun. 2003.

[11] H. Joshi, H.H. Sigmarsson, D. Peroulis, and W.J. Chappell, "Highly loaded evanescent cavities for widely tunable high-Q filters," *IEEE MTT-S International*, pp. 2133-2136, Jun. 2007.

[12] X. Liu, L.P.B. Katehi, W.J. Chappell, and D. Peroulis, "A 3.4-6.2 GHz continuously tunable electrostatic MEMS resonator with quality factor of 460-30," *IEEE MTT-S International*, pp. 1149-1152, Jun. 2009.

[13] S.-J. Park, I. Reines, C. Patel, and G.M. Rebeiz, "High Q RF-MEMS 4-6 GHz tunable evanescent-mode cavity filter," *IEEE Transactions on MTT*, vol. 58, no. 2, pp. 381-389, Feb. 2010.

[14] R. G. Meyer and M. L. Stephens, "Distortion in variable-capacitance diodes," *IEEE Journal of Solid-State Circuits*, vol. 10, no. 2, pp. 47-54, Feb. 1975.

[15] K. Buisman et al., "A monolithic low-distortion low-loss silicon-on-glass varactor-tuned filter with optimized biasing," *IEEE Microwave and wireless components letters*, vol. 17, no. 1, pp. 58-60, Jan. 2007.

[16] M. A. El-Tanani and G. M. Rebeiz, "A two-pole two-zero tunable filter with improved linearity," *IEEE Trans. on MTT*, vol. 57, no. 4, pp. 830-839, Apr. 2009.

[17] Skyworks, "Design with PIN Diodes (application note)," www.skyworksinc.com,

Apr. 2009.

[18] H. Maune, M. Sazegar, Y. Zheng, A. Giere, and R. Jakoby, "Design of planar varactors based on ferroelectric thick-films," *German Microwave Conference*, pp. 102-105, Mar. 2010.

[19] M.-T. Nguyen, W. D. Yan, and E. Horne, "Broadband tunable filters using high Q passive tunable ICs," *IEEE MTT-S International*, pp. 951-954, Jun. 2008.

[20] G. Subramanyam et al., "Linearity and temperature dependence of large-area processed high-q barium strontium titanate thin-film varactors [Correspondence]," *IEEE Trans. Ultrasonics, Ferroelectrics and Frequency Control*, vol. 57, no. 7, pp. 1692-1695, Jul. 2010.

[21] T. Koryu Ishii, Ed., *Handbook of Microwave Technology: Components and Devices*. San Diego, CA: Academic Press, 1995.

[22] P.S. Carter Jr., "Magnetically-Tunable Microwave Filters Using Single-Crystal Yttrium-Iron-Garnet Resonators," *IRE Trans. on MTT*, vol. 9, no. 3, pp. 252-260, May 1961.

[23] J. K. Cavers, "Adaptation behaviour of a feedforward amplifier linearizer," *IEEE Trans. on Veh. Tech.*, vol. 44, no. 1, pp. 31-40, Feb. 1995.

[24] A. M. Smith and J. K. Cavers, "A wideband archtiecture for adaptive feedforward linearization," *IEEE VTC*, pp. 2488-2492, May 1998.

[25] R.G. Randall, J.G. McRory, and R.H. Johnston, "Broadband DSP based feedforward amplifier lineariser," *Electronics Letters*, vol. 38, no. 23, pp. 1470-1471, Nov. 2002.

[26] S. Ayazian and R. Gharpurey, "Feedforward Interference Cancellation in Radio Receiver Front-Ends," *IEEE Trans. on Circuits and Systems II: Express Briefs*, vol. 54, no. 10, pp. 901-906, Oct. 2007.

[27] T. O'Sullivan, R.A. York, B. Noren, and P.M. Asbeck, "Adaptive duplexer implemented using single-path and multipath feedforward techniques with BST phase shifters," *IEEE Trans. on MTT*, vol. 53, no. 1, pp. 106-114, Jan. 2005.

[28] A. Roussel, C.W. Nicholls, and J.S. Wight, "Frequency agile bandstop filter (FABSF)," *IEEE MTT-S International*, pp. 1099-1102, Jun. 2008.

[29] W. Schacherbauer, T. Ostertag, C.C.W. Ruppel, A. Springer, and R. Weigel, "An interference cancellation technique for the use in multiband software radio frontend design," *30th European Microwave Conference*, pp. 1-4, Oct. 2000.

[30] S. Kannangara and M. Faulkner, "Adaptive duplexer for multiband transceiver," *RAWCON Proceeedings*, pp. 381-384, Aug. 2003.

[31] S. Kannangara, "Adaptive Duplexer for Software Radio," Victoria University, Melbourne, AU, PhD dissertation 2006.

[32] Xiaopeng Li and Mohammed Ismail, *Multi-Standard CMOS Wireless Receivers: Analysis and Design.*: Springer, 2002.

[33] D. S. Chan and L. R. Rabiner, "Analysis of quantization error in the direct form for finite impulse response digital filters," *IEEE Trans. Audio & Electroacoustics*, vol. 21, no. 4, Aug. 1973.

[34] C. W. Barnes, B. N. Tran, and S. H. Leung, "On the statistics of fixed-point roundoff error," *IEEE Trans. Acoustics, Speech, and Signal Processing*, vol. 33, no. 3, pp.

595-606, Jun. 1985.

[35] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing Principles, Algorithms, and Applications*, 3rd ed. New Jersey, NY: Prentice Hall, 1996.

[36] V-line directional coupler: 7.15S-dB-1.500V, Accessed 2011, e-MECA, avail: www.e-meca.com.

[37] V-line directional coupler: 722S-dB-1.500V, Accessed 2011, e-MECA, avail: www.e-meca.com.

[38] UMTS Base Station Radio Transmission and Reception (FDD), ETSI TS 125 104 v9.6.0 (2011-01), Accessed March 2011, 3GPP, avail: www.3gpp.org.

[39] AD9744 data sheet, Accessed 2010, Analog Devices, avail: www.analog.com.

[40] S. Haykin, *Modern Filters*. NY, NY: Macmillan Publishing Company, 1989.

[41] S.L. Marple Jr., *Digital Spectral Analysis with Applications*. Englewood Cliffs, NJ: Prentice-Hall Inc., 1987.

[42] S. Haykin, *Adaptive Filter Theory*, 4th ed. New Jersey, NY: Prentice Hall, 2002.

[43] S. M. Kuo and D. R. Morgan, "Active noise control: a tutorial review," *Proceedings of the IEEE*, vol. 87, no. 6, pp. 943-973, Jun 1999.

[44] G. A. Clark, S. K. Mitra, and S. R. Parker, "Block implementation of adaptive digital filters," *IEEE Trans. on Acoustics, Speech, and Signal Processing*, vol. 29, no. 3, pp. 744-752, Jun. 1981.

[45] ADS5474 data sheet, Accessed 2010, Texas Instruments, avail: www.ti.com.

[46] HMC630LP3 data sheet, Accessed 2010, Hittite Microwave, avail: www.hittite.com.

[47] H. Noto, K. Yamauchi, M. Nakayama, and Y. Isota, "Negative group delay circuit for feed-forward amplifier," *IEEE MTT-S International*, pp. 1103-1106, Jun. 2007.

[48] S.J. Erickson, H. Khaja, and M. Mojahedi, "Time- and frequency-domain measurements for an active negative group delay circuit," *IEEE Antennas and Propagation Intern. Symp.*, vol. 3A, pp. 790-793, Jul. 2005.

[49] B. Ravelo, A. Perennec, and M. Le Roy, "Synthesis of broadband negative group delay active circuits," *IEEE MTT-S Intern.*, pp. 2177-2180, Jun. 2007.

[50] O.F. Siddiqui, M. Mojahedi, and G.V. Eleftheriades, "Periodically loaded transmission line with effective negative refractive index and negative group velocity," *IEEE Trans. Antennas and Propagation*, vol. 51, no. 10, pp. 2619-2625, Oct. 2003.

[51] C.D. Broomfield and J.K.A. Everard, "Broadband negative group delay networks for compensation of microwave oscillators and filters," *Electonics Letters*, vol. 36, no. 23, pp. 1931-1933, Nov. 2000.

[52] S. Lucyszyn, I.D. Robertson, and A.H. Aghvami, "Negative group delay synthesiser," *Electronic Letters*, vol. 29, no. 9, pp. 798-800, Apr. 1993.

[53] B. Ravelo, A. Perennec, and M. Le Roy, "Equalization of interconnect propagation delay with negative group delay active circuits," *IEEE Workshop on Signal Propagation on Interconnects*, pp. 15-18, May 2007.

[54] U. Martin and H.W. Schussler, "On digital systems with negative group delay," *Frequenz*, vol. 47, no. 5-6, pp. 106-113, May-Jun. 1993.

[55] J. B. Knowles and E. M. Olcayto, "Coefficient accuracy and digital filter response," *IEEE Trans. Circuit theory*, vol. 15, no. 1, pp. 31-41, Mar. 1968.

[56] A. Papoulis and S. Unnikrishna Pillai, *Probability, Random Variables and Stochastic Processes*, 4th ed. NY, NY: McGraw-Hill, 2002.

Appendix A – Filter Tap Quantization

This appendix presents a statistical analysis relating the digital filter tap quantization to the desired FF attenuation of the hybrid RF-DSP FF filter. The statistical analysis follows that in [*33*], for a direct-form digital filter.

The frequency response of a discrete-time FIR filter with N infinite-precision taps is given by:

$$H_{\infty}(\omega) = \sum_{n=0}^{N-1} c_{\infty}(n) e^{-jn\omega T} \tag{A.1}$$

In (A.1), ω is the continuous-time frequency variable, and T is the period of the digital clock. The n^{th} infinite-precision filter tap is denoted by $c_{\infty}(n)$. A different frequency response occurs when these taps are represented with finite-precision values, $c_{fp}(n)$.

$$H_{fp}(\omega) = \sum_{n=0}^{N-1} c_{fp}(n) e^{-jn\omega T} \tag{A.2}$$

The difference between the infinite and finite-precision frequency responses is:

$$H_e(\omega) = H_{\infty}(\omega) - H_{fp}(\omega) = \sum_{n=0}^{N-1} \left[c_{\infty}(n) - c_{fp}(n) \right] e^{-jn\omega T} \tag{A.3}$$

The effect of the transfer function error is now translated into FF attenuation. Re-arranging (A.3) gives $H_{fp}(\omega) = H_{\infty}(\omega) - H_e(\omega)$. This relationship shows that the finite-precision coefficient filter, $H_{fp}(\omega)$, can be modeled by an infinite-precision filter, $H_{\infty}(\omega)$, in parallel with $H_e(\omega)$ [*55*]. This parallel arrangement is shown in Fig. A.1.

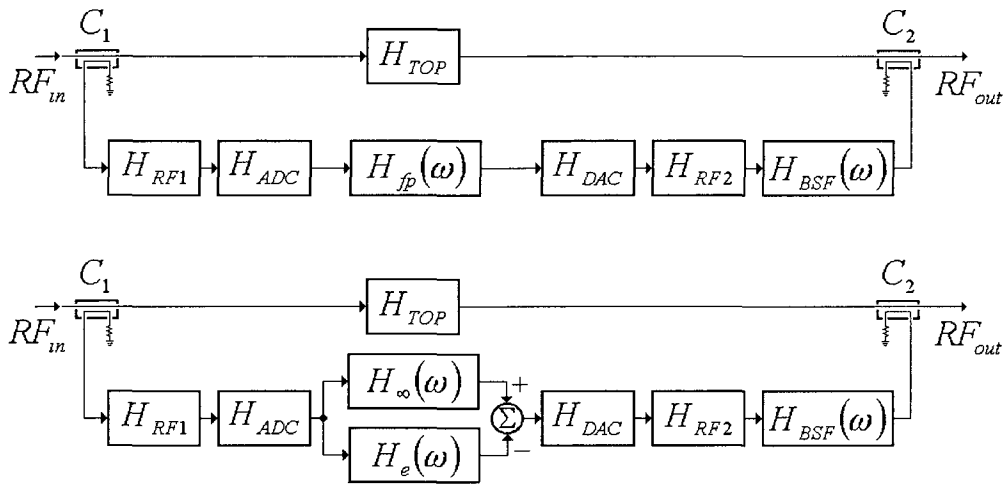

Figure A.1 – Mathematically equivalent models of hybrid FF filter with finite-precision filter taps

In Fig. A.1, the transfer functions of the RF components are expressed in terms of their baseband equivalent functions. All the component values are frequency dependent, but the '(ω)' term has only been appended to those components that intentionally exhibit a frequency dependent magnitude response.

The bottom system in Fig. A.1 consists of three parallel FF paths. If the infinite-precision filter results in perfect FF cancellation at frequencies occupied by interferers (ω_{int}), then resulting magnitude response at these frequencies is:

$$\left|H_{FF}(\omega_{\text{int}})\right| = \left|H_e(\omega_{\text{int}})\right| \cdot \left|H_{C1c} H_{RF1} H_{ADC} H_{DAC} H_{RF2} H_{BSF}(\omega_{\text{int}}) H_{C2c}\right| \qquad \text{(A.4)}$$

The above derivation has followed that from [*33*], but the remainder of this derivation does not. In [*33*], the goal was to determine the effect of coefficient quantization on the mean square error of $H_e(\omega)$, integrated over all frequencies. In this derivation the goal is to determine the effect of coefficient quantization on the magnitude of $H_e(\omega)$ only at

interferer frequencies. Furthermore, in [*33*] only real filter coefficients were considered, but complex values are required for this analysis.

The error for each coefficient due to quantization is: $c_e(n)=c_\infty(n)-c_{fp}(n)$. This definition together with (A.3) yields (A.5). The coefficient error, $c_e(n)$, is decomposed into real and imaginary parts.

$$H_e(\omega)=\sum_{n=0}^{N-1}c_e(n)e^{-jn\omega T}=\sum_{n=0}^{N-1}\left[c_{e,real}(n)+jc_{e,imag}(n)\right]e^{-jn\omega T} \tag{A.5}$$

The strategy of this analysis is to perform a statistical analysis on $H_e(\omega)$ in terms of the tap quantization step-size, Q_{TAPS}. The transfer function errors of $H_e(\omega)$ are related to the attenuation of the hybrid RF-DSP FF filter, $|H_{FF}(\omega)|$, hence can be related to Q_{TAPS}. This step size is a function of the number of bits (b_{TAPS}) used to represent the real and imaginary components of the digital filter taps. The step-size also depends on the maximum digital value used in the fixed-point representation ($v_{TAPS,fs}$).This value is the maximum value for the real and imaginary components of the tap coefficient that can be represented without saturation or wrap-around.

$$Q_{TAPS}=\frac{v_{TAPS,fs}}{2^{b_{TAPS}-1}} \tag{A.6}$$

The real part and imaginary parts of $H_e(\omega)$ are obtained using (A.5):

$$\begin{aligned}\mathrm{Re}\{H_e(\omega)\}&=\sum_{n=0}^{N-1}c_{e,real}(n)\cos(n\omega T)+c_{e,imag}(n)\sin(n\omega T)\\ \mathrm{Im}\{H_e(\omega)\}&=\sum_{n=0}^{N-1}-c_{e,real}(n)\sin(n\omega T)+c_{e,imag}(n)\cos(n\omega T)\end{aligned} \tag{A.7}$$

The coefficient quantization errors (c_e) are assumed uniformly distributed from $-Q_{TAPS}/2$ to $Q_{TAPS}/2$. These components are independent in a direct-form filter since each component is a unique input to the filter's multipliers. This independence means that when $m \neq n$: $E[c_{e,real}(m)c_{e,real}(n)]=0$, and $E[c_{e,imag}(m)c_{e,imag}(n)]=0$. Furthermore, for any m and n, $E[c_{e,real}(m)c_{e,imag}(n)]=0$.

The mean and variance of each of the N components in each of the summations in (A.7) are equal. The mean and variance of the n^{th} term in the $\mathrm{Re}\{H_e(\omega)\}$ summation are respectively:

$$\begin{aligned} E\left[c_{e,real}(n)\cos(n\omega T)+c_{e,imag}(n)\sin(n\omega T)\right] &= 0 \\ \sigma_{He}^2 = E\left[\left|c_{e,real}(n)\cos(n\omega T)+c_{e,imag}(n)\sin(n\omega T)\right|^2\right] &= \frac{Q_{TAPS}^2}{12} \end{aligned} \tag{A.8}$$

Both the real and imaginary components of $H_e(\omega)$ are the sum of N independent identically distributed random variables. For sufficiently large N, the central limit theorem states the probability distribution of the real and imaginary components of $H_e(\omega)$ will converge to a Gaussian distribution [*56*]. Applying the central limit theorem means that both $\mathrm{Re}\{H_e(\omega)\}$ and $\mathrm{Im}\{H_e(\omega)\}$ have zero mean and variance $Q_{TAPS}^2 N/12$.

Equation (A.4) requires the magnitude of $H_e(\omega)$. The magnitude of $H_e(\omega)$ is a function of the two Gaussian random variables in (A.7).

$$\left|H_e(\omega)\right| = \sqrt{\left[\mathrm{Re}\{H_e(\omega)\}\right]^2 + \left[\mathrm{Im}\{H_e(\omega)\}\right]^2} \tag{A.9}$$

The probability distribution function of the magnitude of a complex Gaussian random variable is the Rayleigh distribution [*56*]. The cumulative distribution function of the Rayleigh distribution is:

$$\Pr\left(\left|H_e(\omega)\right| < H_x\right) = 1 - \exp\left(-\frac{H_x^2}{2\sigma_{He}^2}\right) = 1 - \exp\left(-\frac{6H_x^2}{Q_{TAPS}^2 N}\right) \tag{A.10}$$

This distribution function is illustrated in Fig. A.2. In this figure the y-axis is plotted on a logarithmic axis to expand the relevant portion of the curve.

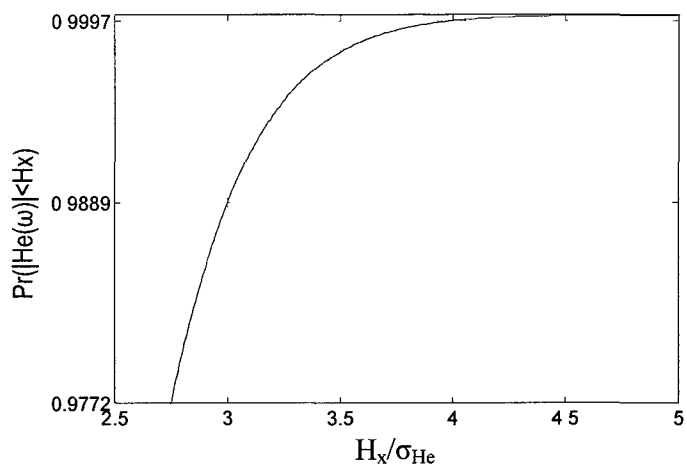

Figure A.2 – Rayleigh cumulative density function

From Fig. A.2 it can be seen that with probability of 0.9889, $|H_e(\omega)|$ will be upper bounded by three standard deviations. The standard deviation is: $\sigma_{He} = Q_{TAPS}\sqrt{N/12}$. The probability that $|H_e(\omega)|$ is bounded by four standard deviations is 0.9997. Using four standard deviations as the statistical upper bound:

$$\left|H_e(\omega)\right| \lessapprox 4\sigma_{He} = 2Q_{TAPS}\sqrt{\frac{N}{3}} \tag{A.11}$$

The result in (A.11) can be used to choose Q_{TAPS}, based on N and the desired notch degradation due to quantization. But, this limit on degradation is only valid 99.97% of the time. The $4\sigma_{He}$ statistical upper bound from (A.11) is inserted into (A.4) to arrive at an

expression relating the filter coefficient quantization step-size (Q_{TAPS}) to the desired FF attenuation, $|H_{FF,des}(\omega_{int})|$:

$$Q_{TAPS} \leq \min_{\omega_{int}} \sqrt{\frac{3}{4N}} \frac{\left|H_{FF,des}(\omega_{int})\right|}{\left|H_{C1c}H_{RF1}H_{ADC}H_{DAC}H_{RF2}H_{BSF}(\omega_{int})H_{C2c}\right|} \tag{A.12}$$

Recall from (A.6) that the quantization step-size can be represented in terms of the number of bits (b_{TAPS}) and the digital full-scale value. The inequality in (A.12) is rearranged to isolate the number of tap bits:

$$b_{TAPS} \geq \log_2 \min_{\omega_{int}} \sqrt{\frac{N}{3}} \frac{\left|H_{C1c}H_{RF1}H_{ADC}H_{DAC}H_{RF2}H_{BSF}(\omega_{int})H_{C2c}\right| v_{TAPS,fs}}{\left|H_{FF,des}(\omega_{int})\right|} \tag{A.13}$$

Simulations are run in Simulink to verify the statistical analysis. The linear constrained minimum variance (LCMV) criteria is used as the goal function for the total FF filter in order to choose the digital filter's tap coefficients. The LCMV is useful when there are more degrees of freedom in the filter optimization then there are constraints. The number of degrees of freedom is the difference between the number of digital filter taps and the number of constraints.

The magnitude response of the hybrid FF filter is constrained to exhibit notches at eight frequencies. The LCMV approach uses the additional degrees of freedom to minimize the variance of the hybrid FF filter's output. In this case the input is white noise. The notch constraints ensure the transfer function of the hybrid FF filter has perfect nulls when the tap coefficients are represented with infinite precision. These optimal digital filter coefficients are represented with fixed-point notation in the simulations with various bit widths.

In the simulation, $v_{TAPS,fs}=\pm 4$, which is chosen empirically from the maximum values required to obtain the desired notches in the transfer function. The cascaded magnitude response of the RF and analog components is unity: $|H_{C1c}H_{RF1}H_{ADC}H_{DAC}H_{RF2}H_{C2c}|=1$; therefore $|H_{FF}(\omega)|=|H_e(\omega)H_{BSF}(\omega)|$. The latter relationship is used to extract $|H_e(\omega)|$ from the simulation results. The number of digital filter taps is set to 32.

Repeated trials are required to validate a statistical analysis. In each trial eight frequency notches are randomly assigned. This setup ensures the tap coefficients are not the same values from trial to trial. The frequency response of the hybrid RF-DSP FF filter from one trial is shown in Fig. A.3. The eight constrained notches are indicated with small circles in Fig. A.3a at the -50 dB level. This figure is showing the magnitude response for the maximum precision case.

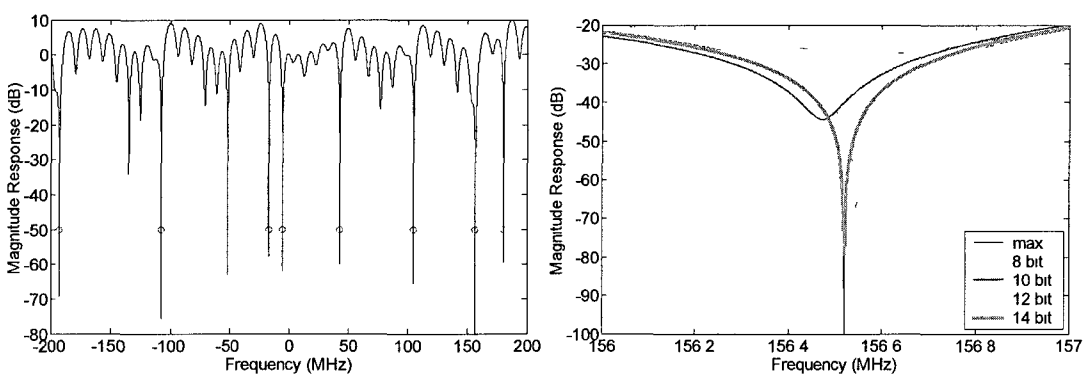

Figure A.3 a) Magnitude response of hybrid FF filter with maximum precision in the digital filter for a single trial b) Zoomed in magnitude response curves with various tap coefficient bit widths

The simulation results in Fig. A.3b show the magnitude responses at one of the notch locations for five different values of b_{TAPS}. The degradation of notch depth is illustrated for decreasing bit widths.

The magnitude response of the hybrid RF-DSP FF filter is recorded at each of the eight frequency notch locations. These eight values are recorded for 50000 trials, where each trial includes six different tap coefficient bit widths. The recorded data is used to generate a cumulative distribution function curve for each tap coefficient bit width.

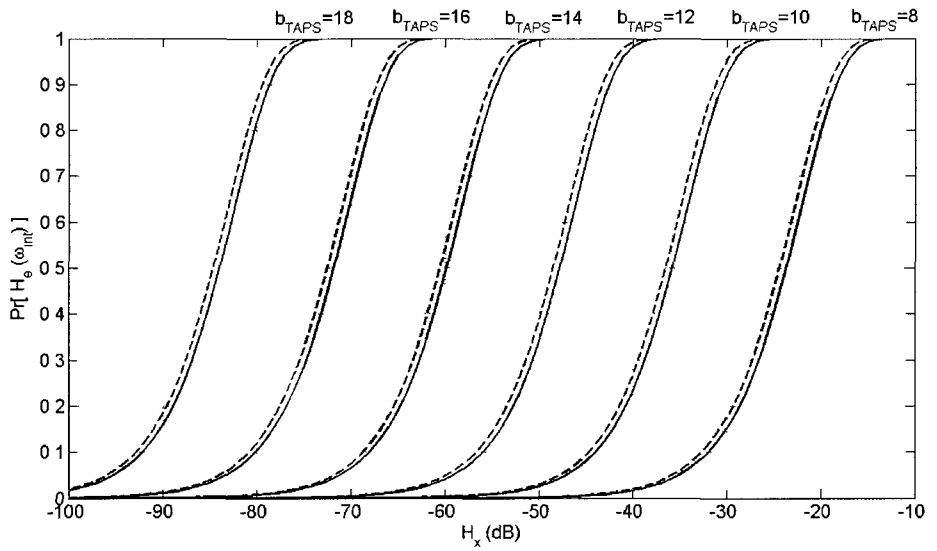

Figure A.4 - Cumulative distribution function of $|H_e(\omega_{int})|$ for varying tap coefficient bit widths (- - -) expected from (A.10), (----) simulated.

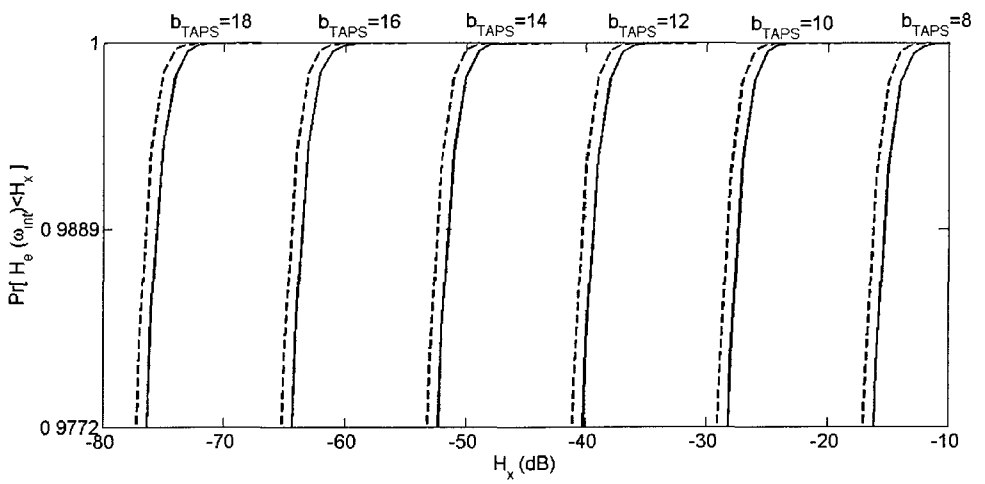

Figure A.5 – Zoomed in view of relevant region of above figure

In Fig. A.4 the cumulative distribution functions are plotted for the 6 different bit widths, based on the simulation results. Also plotted are the theoretical curves, which are based on (A.10). The theoretical and simulation curves differ by less than 2 dB for all values of the cumulative distribution function, as seen in Fig. A.4. The theoretical curves are only used in this work for large values of the distribution function. The relevant region is plotted on a logarithmic scale in Fig. A.5. There is at worst 1 to 2 dB variation between the curves in this figure.

The negligible differences between the simulation and theoretical curves in Fig. A.5 validate the results, including the application of the central limit theorem. The primary result of this section is (A.12). This equation is used to choose the tap coefficient step-size based on the number of filter taps, magnitude responses of the RF and analog components, and the desired level of FF attenuation at the interference frequencies.

The LCMV algorithm is not actually used to choose the digital filter coefficients in the hardware prototype. Constraining the filter to produce perfect notches generally does not result in the optimum output PSD. This means the infinite precision filter taps may not result in notches with infinite depth. The results of this section are useful to ensure the adaptive algorithm limits the depth of the RF-DSP FF filter's notches, and not the tap quantization.