Radio Resource Discovery for Ad-hoc Wireless Networking

M.Sc. Thesis

Niels A. Moseley

University of Twente
Department of Electrical Engineering,
Mathematics & Computer Science (EEMCS)
Signals & Systems Group (SAS)
P.O. Box 217
7500 AE Enschede
The Netherlands

Report Number:	SAS017-04
Report Date:	August 19, 2004
Period of Work:	01/05/2003 – 26/8/2004
Thesis Committee:	Prof. Dr. ir. C.H. Slump
	ir. F.W. Hoeksema
	ir. R. Schiphorst

Abstract

The radio spectrum has almost been fully allocated and the point where no more services can be added is near. This fact calls for a different spectrum management approach. Such an approach is *radio resource discovery*.

This report describes methods for determining unused or under utilized parts of the radio spectrum. These part, or bands, are identified by a radio resource discovery algorithm by means of power measurements. From the power measurements an activity map is generated.

The activity map generated by the radio resource discovery algorithm is then used by ad-hoc networking systems to position themselves in the unused or under utilized parts of the radio spectrum. This increases the effective utilization of the radio spectrum. An emergency ad-hoc networking scenario is used as a framework to derive key RF system parameters.

A low-cost radio resource discovery receiver is developed as a demonstrator and a testbed for further research into radio resource discovery algorithms. Several experiments are done to verify correct operation of the receiver and the resource discovery algorithm.

Also, general receiver concepts and radio architectures suitable for radio resource discovery are presented and discussed.

Acknowledgements

This work would not have been possible without the guidance and patience of Prof Dr. ir. Kees Slump, ir. Fokke Hoeksema and ir. Roel Schiphorst.

I would like to thank Andrzej Baranski PE1ROE for getting me interested in electronics and especially in radio. This thesis would never have been the way it is without his willingness to share his expertise, skills and even components. He suggested using the MAR6 MMIC amplifier which saved me a lot of valuable time and circuitry.

I would also like to thank my partner-in-crime Joyce, who stood by me in difficult times. You're the greatest.

Finally, my gratitude goes out to all the people at SAS who have shown interest in my work and those who have helped me.

Contents

1

Introduction

In recent years, the use of wireless communication technology has increased tremendously. Two forms of wireless communication are responsible for this growth, namely wireless telephony and wireless data networking. Although several systems, e.g. GSM, DECT, IEEE 802.11 were developed to service only one of the two types of wireless communication, there is a push towards a unified approach. With technologies like voice-over-IP (VOIP), it is possible to transmit voice communications over wireless data networking systems based on the IP protocol suite. Recently, good quality VOIP applications have become available.[1]

Wired communication systems are several orders of magnitude faster than their wireless counterparts. Because of this, there exists a demand for higher-speed wireless systems. Inevitably, this increase in speed also means an increase in transmission bandwidth.

At the moment of writing, each new wireless system is allocated a place in the radio spectrum by a regulatory body, e.g. the FCC in the USA and Agentschap Telecom[2] in The Netherlands, with high-speed systems placed at higher frequencies where bandwidth is still available. At some point in time, however, there will be no more usable spectrum available to allocate.

This allocation scheme is inefficient because the allocated space is seldom used fully all of the time. For instance, the TV bands in the Netherlands are hardly used because most of the homes have cable television. Other examples are services that are only used at certain times of the day.

[1] An example of VOIP technology using a distributed organization can be found at http://www.skype.com

[2] http://www.agentschap-telecom.nl/

This observation has led to the investigation of *opportunistic spectrum utilization* where temporarily unused parts of the spectrum are utilized by unlicensed systems. These systems are often of an *ad-hoc* nature, where the network topology changes over time as mobile nodes join or leave the network.

This work was conducted on behalf of the 'Adaptive Ad-hoc free Band Wireless Communications' project [10].

1.1 Assignment

The primary goal of this assignment was to investigate methods by which activity in the radio spectrum can be determined by producing an activity map. The activity map can then be used by an ad-hoc network to perform wireless communication in the inactive parts of the spectrum. This process is called *radio resource discovery*.

What methods can be used to detect the presence or absence of transmissions in the radio spectrum? At what frequency and time resolution should this information be available?

Some important questions to answer are what radio front-end architectures are suitable for radio resource discovery and what IF bandwidth should these radio front-ends have?

A secondary goal of this project is to develop a real-world spectrum resource discovery receiver to demonstrate resource discovery algorithms and to collect radio spectrum measurement data for further research.

An emergency ad-hoc networking scenario was used as the basis of this assignment to help define a framework and to quantify certain system parameters.

1.2 About this report

In chapter two the emergency ad-hoc network scenario is presented and relevant system parameters are derived. In chapter three possible receiver architectures for this network scenario are discussed and a suitable front-end radio architecture is chosen.

The fourth chapter introduces the concept of radio resource discovery. Two methods for measuring RF signal power are presented that allow spectrum activity to monitored. Based on one of these measurement methods a spectrum resource discovery algorithm is

Chapter five describes a real-world radio resource scanning receiver, its hardware and software setup. Some measurements performed using this receiver are presented and discussed in chapter six.

Chapter seven contains conclusions drawn from this project and makes some suggestions for future research.

2

The Ad-hoc Networking Scenario

2.1 Ad-hoc networks

An ad-hoc network is a network architecture that can be setup quickly without prior planning of layout and topology. Nodes in the network can join and leave the network as they please, often without warning. An ad-hoc network is thus very flexible in terms of number of nodes and users.

This flexibility offers an advantage to the users of the network because it allows them to forget about network infrastructure layout and planning. On the other hand, a network still requires some knowledge about where its nodes are in 'network space' and what capabilities these nodes have.

The main challenges of ad-hoc networks are routing of the network traffic and how to handle important forwarding nodes leaving the network. Luckily, solutions to these challenges exist and here they are assumed to be solved problems. A good overview of ad-hoc routing algorithms can be found in [21].

The ad-hoc network considered in this thesis has no pre-allocated radio spectrum like "classic" ad-hoc networks, e.g. Bluetooth and IEEE 802.11b.

2.2 The scenario

The ad-hoc networking scenario used here is modelled after the disaster-relief scenario described by [16].

The scenario is based on the fact that small and medium-sized urban areas may be destroyed by natural disasters like tornadoes, fires or by industrial accidents. A good real-life example is the S.E. Fireworks disaster in Enschede, the Netherlands, where a fireworks depot exploded and destroyed a large part of the city, killing 23 people and wounding more than a thousand. After the explosion most of the damage was done by fire. Allegedly, in part because of lack of communication between different disaster relief groups [4].

In such a scenario it is vital that communication between relief groups is established as quickly and as easily as possible. The fact that the land based networks in such an area will have been destroyed raises the need to communicate using wireless technologies. Access to the core telecommunications network will only be available at the periphery of the disaster area where the damage to the infrastructure is less severe.

The area of operation of the emergency network is confined to the disaster area, which is no larger than a circle with a 2.5 kilometer radius.

An ad-hoc networking approach will allow the relief groups to enter the disaster area and communicate with each other quickly. With slightly more effort the wireless ad-hoc network can be hooked up to the original core telecommunications network almost instantly. From here on, the emergency ad-hoc network will simply be called 'the network'.

The network must be able to handle a number of different modes of communication to facilitate its users. Required network services can range from low-bandwidth voice and email communication to high-bandwidth video and data traffic. Also, some traffic types behave differently from others. Video and voice traffic are a streaming types of traffic, meaning that their data is sent in a continuous fashion. On the other hand, there is *bursty* traffic where network activity is largely unpredictable. In short, the network must be able to handle a wide variety of network traffic and cope with large amounts of data throughput when required. See table 2.1 for a summary of traffic types.

In this scenario a large number of different devices can participate in the network. The most basic communication devices in the network will be voice-only handheld radios. These will send and receive low bandwidth traffic. There are also higher bandwidth devices like small computers and Personal Digital Assistants (PDAs) in the network. Finally there could be more exotic network-enabled devices like bomb disposal robots and other remotely operated vehicles (ROVs). See figure 2.1 and table 2.2 for a summary.

The network devices are either carried by a person or are mounted on or

Traffic	**Bandwidth**	**Stream/Burst**
Voice	≤64kbps	stream
Email	≤64kbps	burst
FAX	≥8kbps	stream
Inter-PDA	≥64kbps	burst
GPS location	≤8kbps	burst
Video	≥100kbps	stream
Database	≥100kbps	burst

Table 2.1: Network traffic types

Figure 2.1: Example of network enabled vehicles and units

in vehicles. This means that most network devices will move during use. It is assumed that the vehicles will not move faster than 100 km/h within the disaster area. Limiting the movement of the nodes is important because it will introduce a doppler-shift in the signals they receive and transmit.

The network devices that are carried by people must be small, lightweight and easy to carry around for extended periods of time. Also, the network devices must not interfere with the user's rescue and support tasks. They must be able to operate under extreme temperature conditions. The disaster area may contain several fires or the disaster may take place during a severe winter storm. A reasonable operating temperature range seems to be the industrial range of -40 ℃ to +85 ℃.

The network devices must be able to communicate under noisy radio conditions. Some areas may have high levels of radio interference from nearby power stations or radio transmitters. Also, the noise conditions may vary

Device category	Bandwidth
PCs and laptops	High
PDAs	Medium
Telephones	Low
Access terminals	Medium
ROVs	Medium

Table 2.2: Network device types

due to atmospheric conditions. These circumstances must not prevent the wireless communication to fail.

Finally, the nodes in the network have no pre-allocated frequency band to operate in. Therefore they will have to scan their entire frequency operating range to find parts of the spectrum that aren't used at the time the network is in operation. These unused parts of the spectrum will become the network's operating frequencies as long as the network doesn't interfere with the primary users of these frequencies.

2.3 Scenario derived system parameters

Now, it becomes possible to derive some network system parameters from the network scenario outlined earlier. The main parameters of interest are network topology, transceiver frequency range, minimum channel bandwidth, minimum transmitter output power and modulation schemes. Most of these aspects depend on other parameters like antenna size, doppler shift, signal propagation. These parameters are discussed where needed.

2.3.1 Network topology

Given that the number of nodes in the network may vary and is not known in advance, the network must be able to expand its capacity as the demand from more throughput grows. The throughput per network node is limited and there comes a point where a single network channel is no longer sufficient to sustain the response time and data throughput required. Therefore, the only solution is to have more than one network channel and allow data to be forwarded from one channel to the other in case two nodes on two different channels want to communicate with each other. This type of network is called a *multi-hop network* because each data packet is forwarded by several intermediate nodes before it reaches its final destination node.

At least some nodes in the network must be able to access multiple channels so they can forward traffic from one channel to another. These nodes either require more than one transceiver or the transceiver bandwidth must be large enough to send and receive two, or more, channels concurrently.

A frequency hopping scheme might be employed to hop between several channels using only a single transceiver. This, however, requires all the channels to be synchronized. The synchronization scheme must guarantee that no transmission destined for a node occurs when that node is listening to a different channel. This guarantee is hard to realize in an ad-hoc multi-hop network because a transmitting node may choose to transmit to any destination or forwarding node subscribed to that channel.

2.3.2 Transceiver frequency range

The transceiver frequency range is the frequency range the receiver and transmitter can use to communicate. For wireless networks this is usually a continuous range but for other systems, like AM radio broadcasting and receiving equipment, this 'range' could be several non-continuous ranges.

For the emergency ad-hoc networking scenario it is important to have a large frequency range so it is likely that the network can find an unused part within that range.

The choice of transceiver frequency range will depend on the state of transceiver electronics, maximum antenna size and signal propagation characteristics within that range.

Transceiver electronics

Technically speaking, transmitters and receiver can be made to operate on frequencies from several tens of kilohertz to a couple of hundred gigahertz. However, towards the high end of that frequency range, it becomes increasingly difficult for the designer to control tuning stability and system noise.

Antenna size

Some network devices require small to extremely small antennas to make them portable and to ensure that they do not interfere with the user's emergency support role. Also, the antennas in these devices must have an omnidirectional radiation pattern so the user can hold the device in any orientation without affecting the signal reception and transmission. And if the antennas can be made small enough, they can be fit inside the network device itself.

The size of the antenna gives a lower bound of the transceiver frequency range. This is because the size of an antenna is related to the wavelength[1] of the signal the antenna is designed to receive. Most radial antenna elements lengths are usually multiples of $\frac{1}{4}\lambda$, where λ is the wavelength. For instance, a half-wave dipole has two $\frac{1}{4}\lambda$-elements.

A single antenna can only effectively receive a certain frequency band around its design frequency. To get more bandwidth it is always possible to use more antennas designed for different frequency bands. This, however, requires the receiver to combine several antenna signals into one or to simply select the best antenna for the job at hand. Both approaches require additional frontend hardware.

Luckily, there are antenna designs that are broadband in nature. A widely used broadband antenna is the discone antenna. The ARRL Antenna handbook 2004 [18] has the following to say about discone antennas:

> The dimensions of a discone are determined by the lowest frequency of use. The antenna produces a vertically polarized signal at a low-elevation angle and it presents a good match for 50-. coax over its operating range.

To get an idea of the size of a typical discone antenna, two designs for a discone antenna can be found in appendix B. The first design is for 200 MHz and the other is for 400 MHz. These frequencies represent the lower useable frequency limit of a discone antenna. The higher frequency limit is about three times the lower limit, 600 MHz and 1200 MHz respectively.

Of the two discone designs, the 200 MHz version is quite big[2] and will interfere with the mobility requirements set by the scenario. The 400 MHz version, while still big in terms of mobile antennas, fits our goal more closely[3]. The discone's size sets the system's lower operating frequency to 400 MHz and the upper frequency to 1200 MHz.

[1] Wavelength $\lambda = \frac{v}{f}$, where v is the propagation speed and f is the frequency of the signal.

[2] It has a diameter of approximately 40 centimeters

[3] It has a diameter of approximately 20 centimeters

Other wide-band antennas include the log-periodic and multi-trap verticals [18]. These antennas tend to less broadband in nature than the discone antenna.

Signal propagation

Another reason for a upper limit of approximately 1 GHz, is that the wavelength at this frequency is in the range of 30cm. Any obstacles larger than this will block the propagation of the radio wave or have a scattering effect. This can be a problem in disaster areas with numerous large objects, e.g. earthquake zones or mountainous terrain.

A more thorough treatment of signal propagation is given in section 2.3.4.

2.3.3 Minimum channel bandwidth

The minimum channel bandwidth is the smallest unit of bandwidth that a network node must be able to receive to be able to interface with the network. The very basic nodes that require a small amount of network traffic will have a bandwidth equal to the minimum channel bandwidth. More complex nodes will have a bandwidth that is an integer multiple of the minimum channel bandwidth.

Each node has one or more wireless connections to other nodes. These connections are handled by a node's transceiver. As the bandwidth of the transceiver increases, so does it's ability to handle network traffic. Network nodes that generate or demand a large amount of traffic need more bandwidth and therefore more channels to communicate.

The types of traffic expected in the network is very diverse. This makes the amount of traffic in the network difficult to predict. Each node must be able to handle different types of traffic even though the node itself is not interested in all types of traffic. This is because the network may choose to route traffic through the node which is destined for other nodes.

Knowing the channel bandwidth and the signal-to-noise ratio of the channel, the maximum channel capacity may be calculated using the well-known Shannon channel capacity formula [15]:

$$C = B \cdot \log_2\left(1 + \frac{P}{B \cdot N_0}\right) = B \cdot \log_2\left(1 + SNR\right) \tag{2.1}$$

,where C is the channel capacity [bits/s]
B is the channel bandwidth [Hz]
N_0 is the noise power spectral density [W/Hz]
P is the received signal power [W]
SNR is the signal to noise ratio

Equation 2.1 assumes that the channel noise is additive white (spectrally flat) Gaussian noise.

Depending on the modulation scheme used, a reasonable minimum signal-to-noise ratio is 15 dB. The following table was produced by using this signal-to-noise ratio:

Bandwidth [kHz]	Capacity [kbits/s]	Capacity [koctet/s]
100	498.29	62
200	996.58	125
500	2491.45	311
1000	4982.89	623
2000	9965.78	1246
5000	24914.46	3114

A naive approach to determining the channel bandwidth would be to choose the largest bandwidth possible which is limited only by the state-of-the-art. But each unique wireless channel must be placed in an unoccupied part of the spectrum. If the channel bandwidth is too large, it will become difficult to find an unoccupied part large enough to accommodate the channel. Therefore, the channel bandwidth must not be too large.

It is the author's opinion that a single channel should not exceed 1 MHz in bandwidth. Such a channel can be accommodated easily, especially in the higher regions, e.g. above 500 MHz, of the spectrum where the more bandwidth-intensive services reside. A 1 MHz-wide channel has ample capacity for several 100 kbps video streams and data-intensive bursts.

Assuming that the modulation scheme does not offer a more efficient modulation to obtain higher data rates using the same amount of bandwidth, a network node can use multiple channels to communicate. In order to do this, a node's transceiver should be able to receive two or more channels concurrently. This can be achieved by a transceiver that has a bandwidth of two or more channels or by having multiple transceivers per node.

A single-transceiver node that can receive multiple channels will only reach it's full communicative potential if all the channels are marked unoccupied by the resource discovery algorithm. If this is the case, the full transceiver bandwidth can be used by the node to communicate.

2.3.4 Minimum transmitter output power

The minimum transmitter output power is the output power the transmitter of a network node must radiate from the antenna. More power will increase the chance of successfully communicating under adverse conditions, e.g. strong fading. But more output power will decrease battery life.

The minimum transmitter output power is determined by the battery limits and the propagation characteristics and circumstances the system is designed for. In calculating the minimum transmitter output power, it is assumed no fading or no scattering occurs, the nodes operate in a line-of-sight mode and all of the transmitter output power is radiated from isotropic radiators with 0 dB gain. The minimum output power is calculated by working out the link budget.

Battery limits

Some devices are power restricted (PDAs, Hand held telephones, Laptops). Most modern high-frequency battery operated wireless hand held equipment has a maximum transmitter output of several hundred mW. To maximize battery life, the transmitter output power should be kept as low as possible.

Path loss and propagation

The free space path loss of a carrier is related to the distance the carrier travels in free space and the wavelength of the carrier. This relationship is expressed by the following formula:

$$L = 20 \cdot \log_{10}\left(\frac{4 \cdot \pi \cdot D}{\lambda}\right) \tag{2.2}$$

,where

- L = Path loss (dB)
- D = Distance between receiver and transmitter

- λ = Wavelength of the transmitted carrier

Equation 2.2 is valid only for a single direct path from transmitter to receiver. For multi-path propagation more complex models apply because effects like fading come into play.

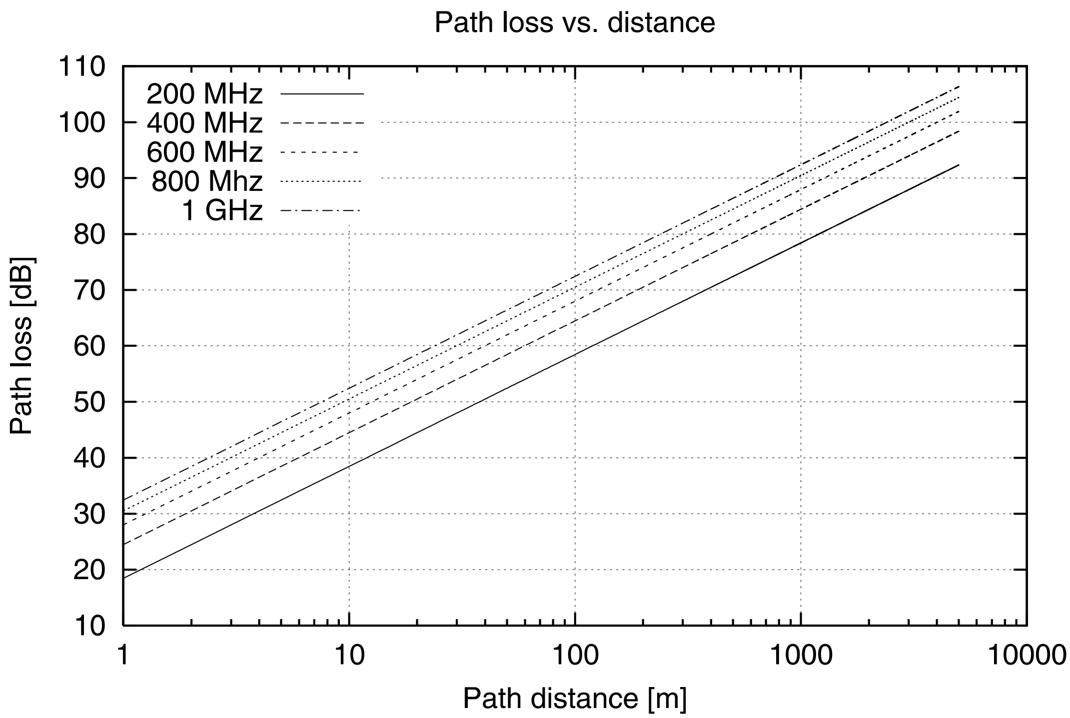

Figure 2.2: Path loss versus distance between transmitter and receiver

Plotting the path loss of a carrier against distance for different frequencies yields the plot shown in Figure 2.2. In this figure it can be seen that for low frequencies, the path loss at 5 km distance is around 92 dB. But as the frequency gets toward 1 GHz, this figure rises to about 107 dB. Path loss puts an upper limit to the frequency range, any higher than 1 GHz and the path loss will become prohibitively large for low-power communications.

Link budget

The link budget is a calculation of the transmitter power required to meet the needs of the receiver. Generally, in the case of a digital transmission system, the maximum bit error rate (BER) at the receiver is a requirement. Because the modulation method of the networking nodes is not known at this moment, it was conservatively chosen to have a minimum carrier-to-noise ratio of 15 dB at the input of the demodulator in the receiver.

The maximum distance between the transmitter and receiver is taken to be 5 km and the channel bandwidth is 1 MHz. Furthermore, it is assumed that both the transmitter and receiver have ideal isotropic antennas.

First, the channel noise is calculated, using equation A.2:

$$N = 1.38 \cdot 10^{-23}[J/K] \cdot 290[K] \cdot 1 \cdot 10^{6}[Hz] \approx 4 \cdot 10^{-12}[mW] = -114[dBm]$$

Then, the receiver's noise figure is added. This has been conservatively estimated to be 15 [dB]. Many IC based receivers have a noise figure lower than that.

$$NoiseFloor = -114 + 15 = -99[dBm]$$

To get the required power at the detector, the minimum required carrier-to-noise ratio is added.

$$P_{receiver} = -99 + 15 = -84[dBm]$$

Now, it is possible to calculate the minimum required transmitter output power by determining the path loss using the path loss equation shown in Equation 2.2. Substituting the maximum distance of 5 km for D and setting the wavelength λ to represent 200 MHz and 1 GHz respectively, the path loss on these frequencies are as follows:

$$L_{path} = 20 \cdot \log_{10}\left(\frac{4 \cdot \pi \cdot 5 \cdot 10^{3}}{1.5}\right) \approx 92[dB] \quad \text{at 200 [MHz]} \tag{2.3}$$

$$L_{path} = 20 \cdot \log_{10}\left(\frac{4 \cdot \pi \cdot 5 \cdot 10^{3}}{0.3}\right) \approx 106[dB] \quad \text{at 1 [GHz]} \tag{2.4}$$

Therefore, the required transmitter power will be:

$$P_{transmitter} = -84[dBm] + 92[dB] = 8[dBm] \quad \text{at 200 [MHz]} \tag{2.5}$$

$$P_{transmitter} = -84[dBm] + 106[dB] = 22[dBm] \quad \text{at 1 [GHz]} \tag{2.6}$$

Several loss factors have not been taken into account:

- Indoor propagation dampening
- Antenna mismatching and inefficiency
- Wave absorption and fading

The use of antenna gain allows the receiver to reclaim some of the losses listed in the table above. Unfortunately, an antenna will provide gain in some directions while in other directions it will introduce a loss with respect to an isotropic radiator.

It is advised to empirically add a couple of dBm to the minimum required transmitter power to take the aforementioned loss factors into account. A reasonable transmitter output power is therefore in the region of 30 [dBm], or about 1 [W].

Another important phenomenon that has not been taken into account is man-made noise, the noise made by electrical appliances. The man-made noise level depends on the type of environment and the time of day. During night-time most appliances are switched off.

The graph in Figure 2.3 shows the dependence of man-made noise on location and on frequency. From this figure it is clear that man-made noise is not a real problem for a system operating above 400 MHz because most of the noise is present below that frequency, even in urban areas. However, the data is from 1998 so it might be outdated, taking the recent growth of wireless systems into account.

It has to be clear that even with the most careful planning of a link budget, it can never be guaranteed that two nodes can connect successfully. The reason for this is that one can never plan for all possible propagation environments and planning for a worst-case is too expensive.

As a final note it must be said that there are more sophisticated and accurate ways of determining a link budget by using better propagation models. A couple of models are presented by Yang [33].

2.3.5 Modulation schemes

A modulation scheme takes information in the form of *bits* and modulates the carrier of the transmitter in such a way that a receiver is able to extract the information (bits) from the modulated carrier. The ideal modulation scheme can transmit an infinite number of bits in an infinitely small time using an infinitely narrow bandwidth. This, unfortunately, is not possible.

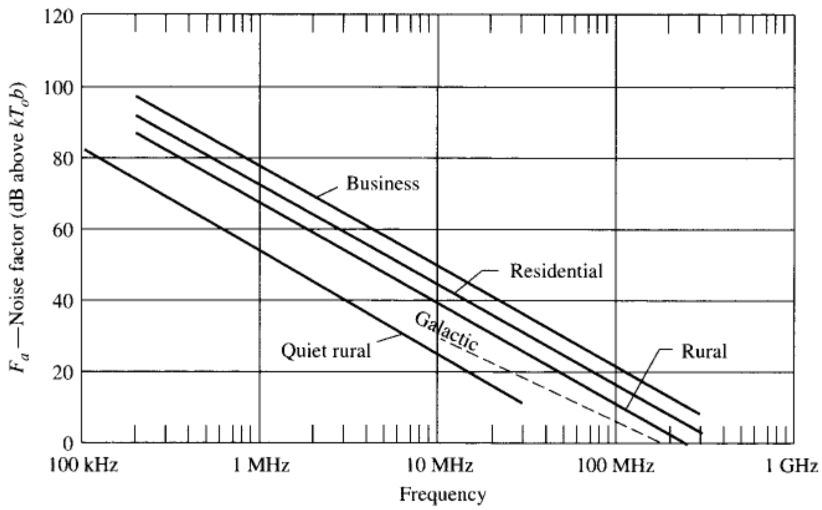

Figure 2.3: Effective antenna noise factor of man-made noise at various locations [27]

Each modulation scheme has a certain *spectral efficiency*. The spectral efficiency is a measure that expresses the modulation scheme's ability to transmit at a rate R (bits/s) within a channel of bandwidth W (Hz)[4]. The spectral efficiency, in bits/s/Hz, is determined by calculating $\frac{R}{W}$. A comparison of the spectral efficiency of different modulations schemes is shown in Figure 2.4.

The spectral efficiency is not the only property that differentiates modulation schemes. For instance, modulation schemes also differ in their ability to handle multi-path propagation effects and carrier frequency offsets.

Unfortunately, the spectral efficiency cannot be arbitrarily large because performance is limited at the receiver by noise [15]. For instance, Figure 2.4 indicates that QAM-64 is more spectrally efficient than QAM-16. Their spectral efficiency is 7 bits/s/Hz and 4 bits/s/Hz respectively. Figure 2.5 indicates that, in order to maintain the same error probability of 10^{-5}, the signal-to-noise ratio must increase from around 14 dB per bit (QAM-16) to around 19 dB per bit (QAM-64).

First, requirements of the modulation scheme are formulated after which a suitable modulation scheme is presented.

[4]In order to compare the spectral efficiency of modulation schemes in a fair way, the bit error rate at the demodulator must be kept equal.

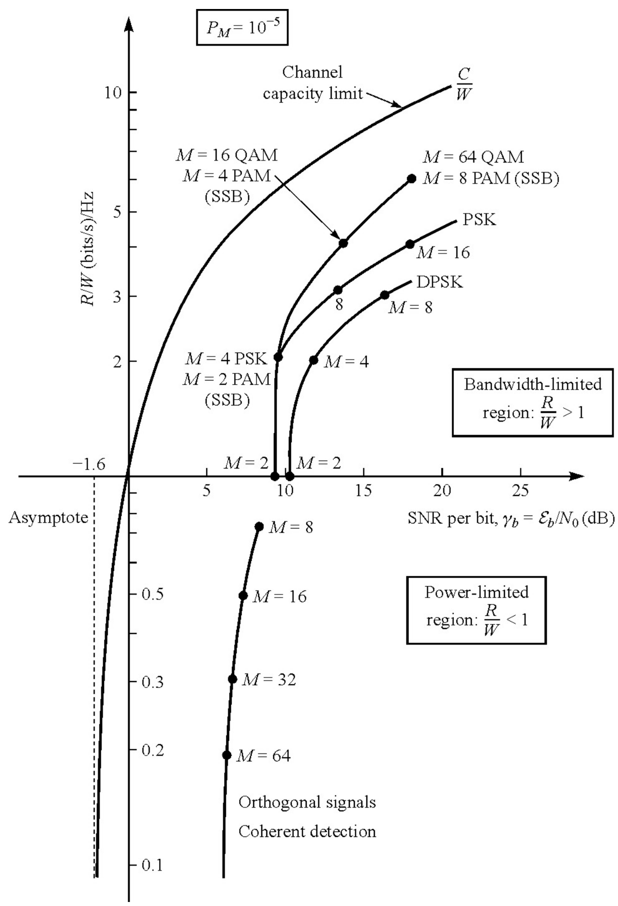

Figure 2.4: Comparison of several modulation methods at 10^{-5} symbol error probability [15]

Modulation scheme requirements

The main reason for introducing radio resource discovery is to increase the efficient usage of the radio spectrum. It would be a wasteful exercise to reclaim unoccupied radio spectrum and use it inefficiently. The more spectrally efficient a network is, the less spectrum it requires to communicate at certain data rate. Therefore, the modulation method must be as spectrally efficient as possible. This is the first requirement.

As described in Section 2.3.5, the spectral efficiency that can be attained by a modulation scheme depends on the noise and propagation conditions. The noise and propagation conditions may vary due to atmospheric circum-

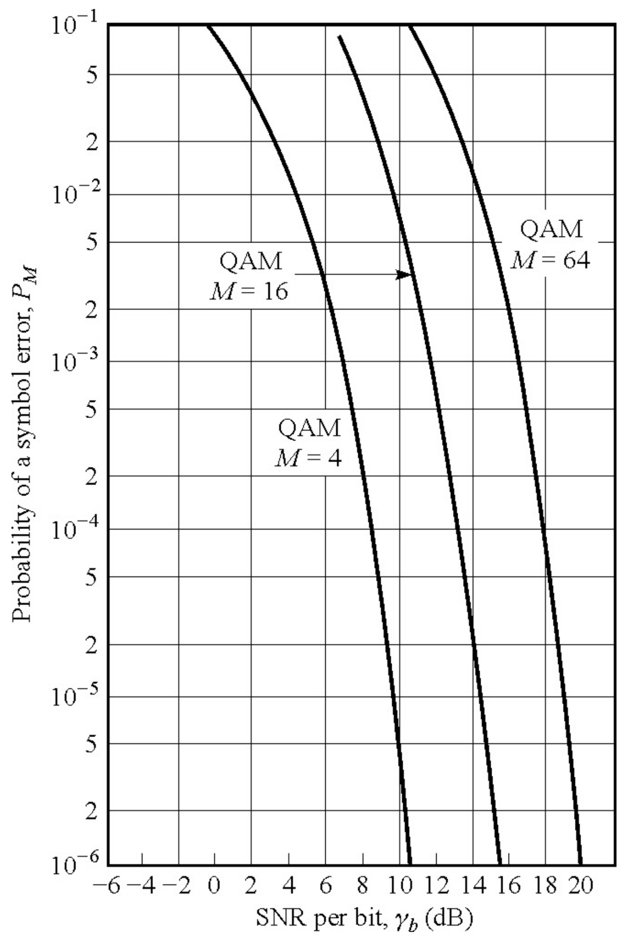

Figure 2.5: Probability of a symbol error for QAM [15]

stances. The modulation scheme must therefore be able to adapt to these varying noise and propagation conditions. Therefore, the second requirement is adaptivity.

A third requirement comes from the fact that there are multiple nodes, or users, in a network. The modulation scheme must support a multiple user model.

A fourth requirement is that the modulation scheme must be resistent to frequency offset and doppler shifting of the carrier frequency. Frequency offsets arise due to inaccuracy in the frequency determining components of the carrier oscillator in the transmitter. Technically these offset problems can be solved by using carrier offset estimation algorithms. However, some of these solutions are prohibitively expensive so some frequency offset must be tolerated.

As mentioned in Section 2.2, the network nodes move at a maximum velocity of 100 km/h. This movement will introduce a doppler shift of the carrier frequency. The maximum amount of doppler shift is calculated in the next section.

Calculating maximum doppler shift

The change in frequency, Δf, depends on the relative speed, v_r, between the transmitter and receiver and the carrier frequency, f_c.

$$\Delta f = f_c \cdot \frac{v_r}{c} \quad \text{where c equals the speed of light} \tag{2.7}$$

The doppler shift is most pronounced at the higher end of the transceiver range, which is 1 GHz in case of the scenario. At this frequency the doppler shift between two nodes travelling in opposite directions at 100 km/h is around 185 Hz.

While the maximum doppler shift of 185 Hz might not seem much, the modulation scheme must be able to handle this without problems.

A suitable modulation scheme

While there are many modulation schemes that will fit all the above requirements, a reasonable choice is Orthogonal Frequency Division Modulation, or OFDM [30]. This modulation method has recently become popular in wireless networking system, i.e. HiperLAN2.

Orthogonal Frequency Division Multiplexing belongs to the multi-carrier modulation family. The modulator divides the incoming data across multiple carriers that are modulated at a lower rate compared to a single carrier modulation system. An OFDM symbol therefore consists of multiple carriers that can be individually modulated in amplitude and phase.

The way the OFDM subcarriers are modulated can vary from BPSK for robust low rate communication to more complex QAM modulation for higher throughput. Therefore, OFDM is able to adapt to different noise conditions. It is also spectrally efficient[5] and it is able to support multiple

[5] The spectral efficiency depends on the subcarrier modulation type.

users through Frequency Domain Multiplexing (FDM) and Time Domain Multiplexing (TDM) techniques, see Lawrey [17] and Van de Beek [30].

Also, OFDM is very resistent to multi-path interference and when the doppler shift is smaller than half the distance between the OFDM subcarriers, this shift can be easily compensated. For doppler shifts larger than half the subcarrier distance, more computationally intensive techniques are needed.

2.3.6 System parameter summary

For convenience the system parameters derived in the past sections are listed here.

Parameter	Value
Topology	Multi-hop
f_{min}	400 MHz
f_{max}	1000 MHz
P_{tx}	30 dBm
$SNR_{channel,min}$	15 dB
$B_{channel,min}$	1 MHz
$f_{doppler,max}$	185 Hz

3

Radio Receiver Concepts and Architectures

This chapter first introduces some general receiver concepts and properties. These concepts will be used throughout this thesis. Then, the direct-conversion and super-heterodyne front-end architectures are presented and their characteristics discussed. Finally, an architecture is chosen to form the basis of the network nodes.

This chapter is by no means a full discussion of receiver architectures. The reader is directed to Dixon [6] or Rhode [25] for a more thorough treatment of the subject.

3.1 General receiver concepts

A receiver is made to gather energy from its antenna and convert or translate that energy into a form that the user can process. To achieve this task, the receiver performs the following subtasks [6]:

- Antenna matching
- Selection of desired signals
- Rejection of undesired signals
- Amplification by very large factors
- Demodulation

- Error detection and/or correction
- Received information conditioning and output

Every radio receiver has several parts that work together to perform all the functions mentioned above. A receiver may have one or more of the following components:

- Antennas
- Filters
- Frequency mixers
- Oscillators
- Amplifiers
- Demodulators

While antennas is a special field dealing with electromagnetic waves and their propagation, the other components are in the more mainstream category of micro electronics. Now, the question to answer is, how do the components determine the performance of a receiver.

3.2 General receiver properties

A receiver is a complex piece of equipment with many properties. When designing a receiver many desired features contradict each other, i.e. dynamic range and noise characteristics. A receiver is usually tailor made for a specific purpose. It is therefore paramount to understand the most important properties of a receiver. These properties will be discussed here.

3.2.1 Sensitivity and selectivity

Two most important properties of a receiver are sensitivity and selectivity. Both properties are explained in [6] as follows:

> Sensitivity, which is the minimum signal level that a receiver requires to produce an acceptable output signal-to-noise ratio when the input signal is modulated by a standard amount.

> Selectivity, which is the ability of a receiver to reject signals outside a given band while accepting signals that are within that band.

There are several ways of quantifying the *sensitivity* of a receiver. One method is the Minimum Discernable Signal (MDS) sensitivity which is the signal level required at the antenna to produce an output signal with a signal-to-noise ratio of 3dB. The output signal is often taken from the demodulator output.

Another way of quantifying sensitivity is the signal level that is required at the antenna to obtain a certain received uncoded bit-error-rate, usually in the range of 1×10^{-3} to 1×10^{-6}. This measurement is useful in a digital receiver where bits are received instead of continuous analog signals.

Selectivity is harder to quantify because the selectivity of a receiver depends on so many factors. One factor is the narrowest bandwidth within the receiver. Another factor is how the receiver reacts to very strong out-of-band signals. A strong out-of-band signal can desensitize a receiver such that a weak in-band signal cannot be detected.

The performance of a receiver is not only determined by its sensitivity. What good is high sensitivity when a very strong out-of-band signal is blocking the reception? It is also possible that a receiver not only responds to the frequency it is tuned to, it may also have spurious, or mirror, responses.

These mirror responses are due to imperfect frequency mixers and other nonlinearities within the receiver. Thus, when designing a receiver special attention must be paid to the non-ideal, e.g. nonlinear, properties of amplifiers, filters, mixers and other types of circuit used in receivers. A way to express the nonlinear behavior of a receiver is through its intercept or compression points [6].

3.2.2 Receiver noise figure

The noise figure of an ideal receiver is 0 dB. This means that the receiver does not add noise to the signal it receives. Unfortunately, this is not possible in reality although very expensive cryogenically cooled receivers can come close to this ideal noise figure.

The noise figure of a receiver is defined as the quotient of the signal-to-noise ratio at the input of the receiver divided by the signal-to-noise ratio at the output of the receiver:

$$NF = SNR_{i,dB} - SNR_{o,dB} \tag{3.1}$$

,where $SNR_{i,dB}$ and $SNR_{o,dB}$ are in decibels.

The noise figure must not be confused with the noise factor. The two mean the same except the noise figure is expressed in dB and the noise factor is a multiplicative constant. Often, the noise figure is denoted as 'NF' while the noise factor is written in lower case letters, 'nf' or even 'f'.

3.2.3 Gain distribution

In figure 3.1 a basic *direct-conversion*(see 3.3) receiver is shown. A signal enters the antenna and is sent through a band-pass filter to filter out unwanted frequencies before they enter the mixer. In a direct-conversion receiver, the mixer mixes the received signal down to DC. The mixed signal is then filtered using a low-pass filter to filters out unwanted products generated by the mixer. The remaining signal is then amplified for further processing.

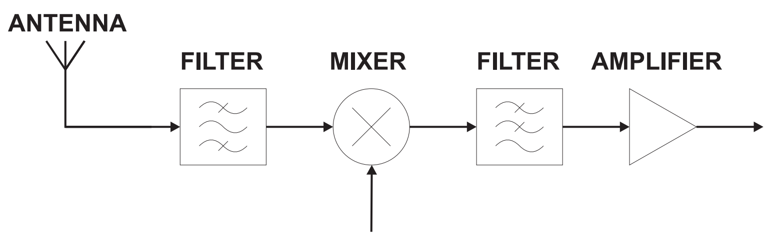

Figure 3.1: A basic direct-conversion receiver

	Filter	**Mixer**	**Filter**	**Amplifier**
Level In [dBm]	-100	-108	-84	-87
Gain [dB]	-8	24	-3	60
Cumulative gain [dB]	-8	16	13	73
Level Out [dBm]	-108	-84	-87	-27

Table 3.1: Gain distribution for the receiver shown in figure 3.1

Table 3.1 shows the gain of each stage. This is called the *gain distribution*. Also shown in the table is the cumulative gain. This is the total gain from the antenna up to and including the stage it is given for. With a gain distribution chart, each receiver stage can be optimized with respect to its input signal level to avoid nonlinear behavior under normal operating conditions[1].

3.2.4 1dB Compression point

Most stages of a receiver are designed for linear operation under normal conditions. Under strong signal conditions, however, some stages may become nonlinear due to overloading. The designer may choose to make the input signal level to these stages as small as possible. While this will certainly make the receiver less susceptive to overloading, it will increase the overall noise in the receiver under normal conditions. Generally, there is a trade-off the two factors.

A useful system parameter of a stage is its 1dB compression point. At low signal levels, the stage is considered linear, but as the input signal level is increased, the amount of signal at the output will begin to tail off. The 1dB compression point is the point where the output power level deviates 1dB from the ideal linear response.

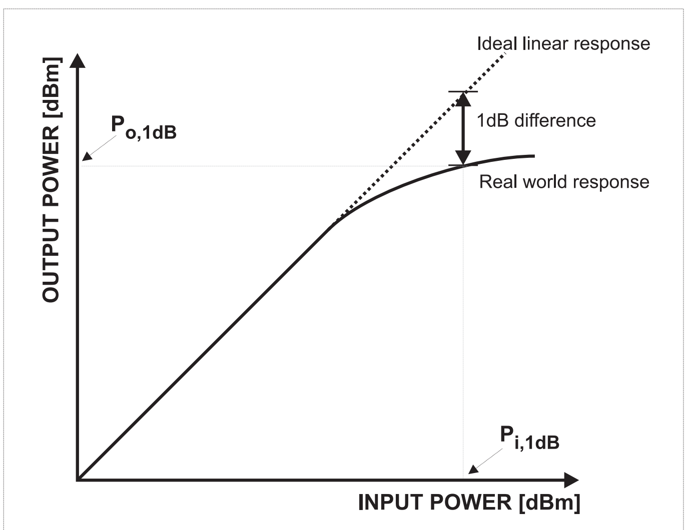

Figure 3.2: 1dB compression point

The amount of input power where the 1dB compression occurs is called

[1]Most receivers avoid this nonlinear behavior to some extent by introducing automatic gain control (AGC).

the *input 1dB point*, $P_{i,1dB}$. The amount of output power where the 1dB compression occurs is called the *output 1dB point*, $P_{o,1dB}$. This is illustrated in figure 3.2.

3.2.5 Minimum discernable signal

While the 1dB compression point determines the large signal strength signal performance of a receiver, the noise figure determines the small signal strength signal performance. This is because a very small signal can only be detected if it is stronger than the noise of the detecting system. This property is expressed as the minimum discernable signal and is defined as follows:

Definition The minimum discernable signal (MDS) is defined as 3[dB] above the noise floor of the receiver.

$$MDS = N_{floor} + 3[dB]$$

3.2.6 Dynamic range

The distance between the 1dB compression point and the minimum discernable signal is called the dynamic range. A graphical representation of dynamic range is shown in figure 3.3 and its mathematical definition is depicted by equation 3.2.

$$DR = P_{in,1dB} - MDS \tag{3.2}$$

3.2.7 Automatic gain control

It is possible to avoid nonlinear behavior of a receiver to some extent by changing the gain distribution when a strong signal is received. This is achieved by monitoring the output of a stage in the receiver. When the output reaches a certain level, the gain of the preceding amplifier stage is reduced to bring the signal within the linear region of the following stages. This process is called automatic gain control (AGC).

The gain reduction action is controlled by the AGC controller as shown in Figure 3.4.

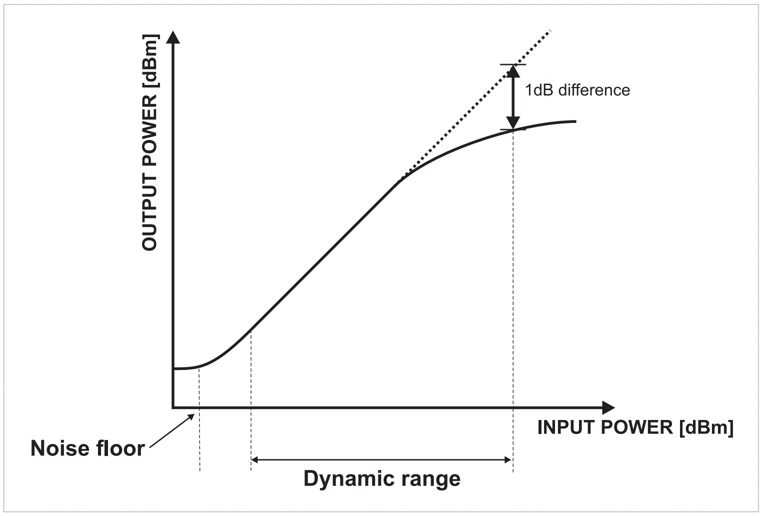

Figure 3.3: Graphical representation of dynamic range

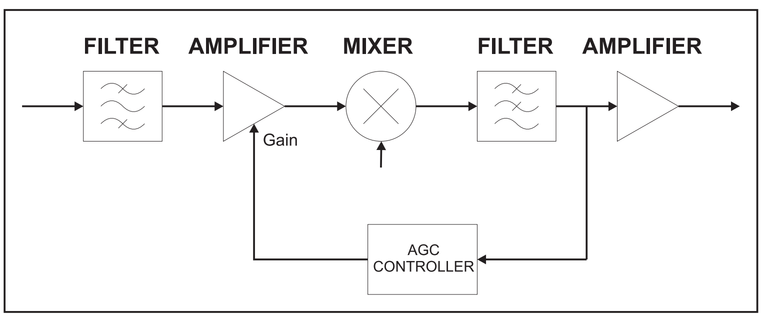

Figure 3.4: Automatic gain control is added to protect the second amplifier from overloading

3.3 The direct-conversion receiver architecture

One of the more conceptually simple receivers is the *direct-conversion* receiver (already shown in the preceding paragraphs). It is also known as the *Zero-IF* receiver because it has no *intermediate frequency* stages.

In the following sections, the receiver layout will be presented. This is followed by a review of the positive and negative aspects of the direct-conversion receiver. Then a summary is given.

3.3.1 Receiver layout

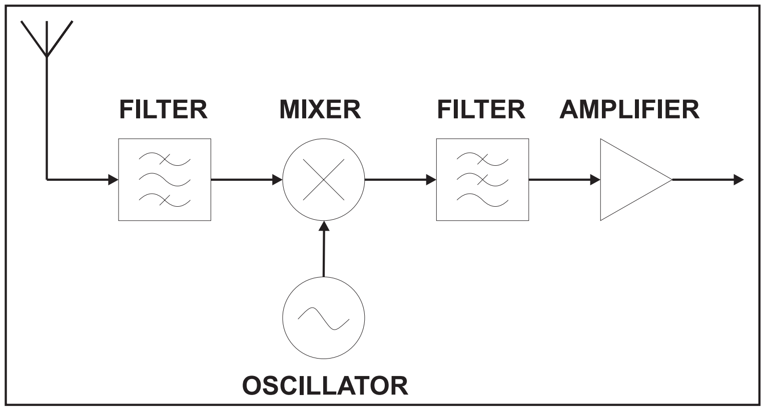

Figure 3.5: A basic direct conversion receiver

A block diagram of a direct-conversion, or DC, receiver is shown in figure 3.5. The signal from the antenna is fed into a bandpass filter to filter out most of the out-of-band signals. This filter is called a pre-selection filter. The signal at the output of the pre-selection filter is mixed down to D.C., or 0 Hz, by frequency mixer. This mixer is driven by an oscillator which is tuned exactly to the carrier frequency of the desired signal. The oscillator's output and the signal are *in beat*. The oscillator in a DC receiver is also known as a *beat-frequency oscillator*, or BFO for short.

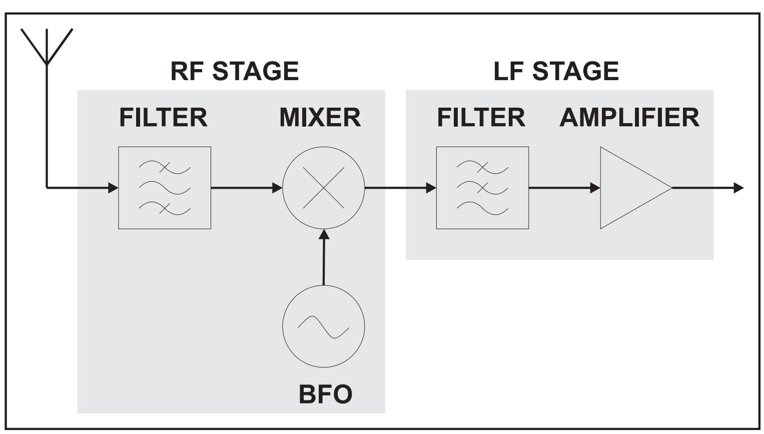

Figure 3.6: Different stages of the direct-conversion receiver

Figure 3.6 shows the two stages of a direct-conversion receiver. The first stage is the radio frequency, or RF stage. This stage operates at the frequency of the desired signal. The second stage is the low frequency, or LF, stage.

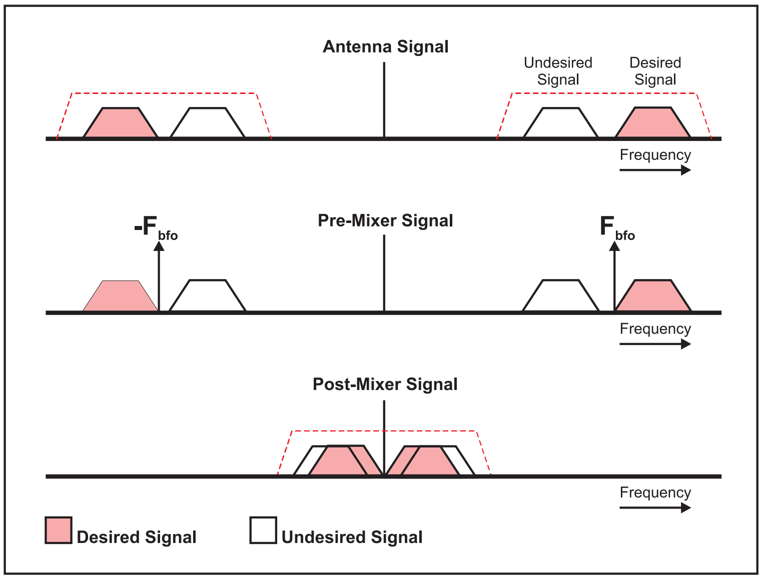

Figure 3.7: Spectrum plots

Most of the selectivity of a direct-conversion receiver is obtained in the LF stage of the receiver. The selectivity comes from the post-mixer lowpass filter because this filter is made to pass just the desired signal bandwidth. The bandpass filter at the antenna, however, does not provide as much selectivity as this filter must pass the entire band of interest.

A direct-conversion receiver has one mixer that mixes the desired signal directly to D.C., or 0 Hz. Figure 3.7 shows three spectrum charts representing the mixing process in the receiver. Two signals are present at the antenna, a desired signal and an undesired signal. The undesired signal is located very close to the desired signal in this case.

Following the signals through the receiver, the two signals are filtered by the bandpass filter but both signals survive because the filter bandwidth is too large to filter out the undesired signal. This is shown in the first spectrum chart in figure 3.7.

Then the two signals arrive at the mixer. For clarity, the frequency of the BFO and the two signals are shown depicted in the pre-mixer spectrum chart. Note that the two signals are still separate.

In the post-mixer spectrum chart, the spectrum at the mixer output is shown. The two signals are mixed together at D.C. and are no longer separable. This is because the negative frequencies of the undesired signal are mixed to the positive frequencies and vice-versa thereby overlapping the desired signal at D.C. This is one of the shortcomings of a DC receiver; it cannot differentiate between the lower and upper sideband around the BFO frequency. Luckily, the upper and lower sideband can be separated by using a quadrature mixer.

In Figure 3.8 a direct-conversion receiver employing quadrature mixing is shown. The receiver has two mixers, each driven by the same oscillator with the difference that one mixer is driven using an oscillator signal that has been phase shifted by 90 degrees. This mixer, shown in the lower branch of Figure 3.8 produces the out-of-phase Q-signal and the other mixer, in the top branch, produces the in-phase I-signal. This dual-mixer configuration is often seen as a single *complex* mixer.

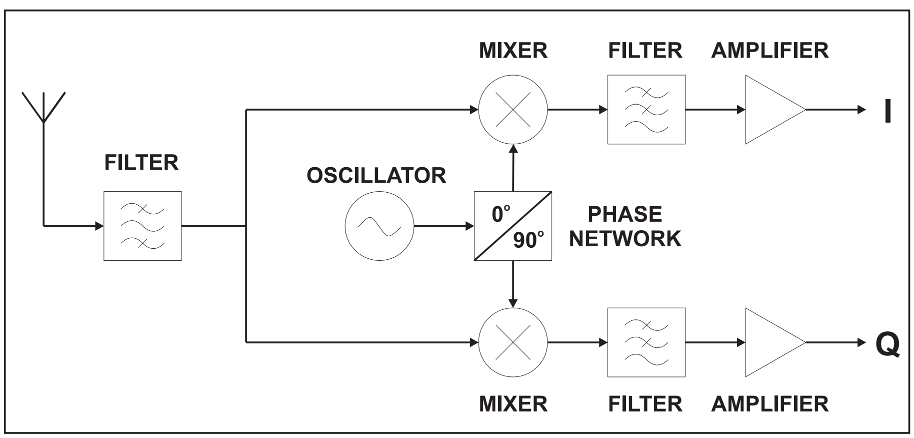

Figure 3.8: A direct-conversion receiver employing quadrature mixing

The I and Q signals have different phase relationships of the upper and lower sidebands with respect to each other. This property ensures that, although each signal has overlapped sidebands like that shown in Figure 3.7, the sidebands can be separated [11].

It has to be clear that most modern direct-conversion receivers are of this quadrature type but that the history of this receiver goes back to the early years of radio when complex mixers were unheard of.

3.3.2 Positive aspects

The direct-conversion receiver has a number of positive aspects. These qualities are:

- Low parts count.
- Low power [7].
- Highly integrated single-chip solution possible [7].
- Low number of spurious responses.

Recently, the direct-conversion receiver gained in popularity as its low parts count and low power qualities [7] make it an almost ideal architecture for implementation in integrated circuits. A good example of this is a single-chip GSM/GPRS implementation by Infineon, see Figure 3.9.

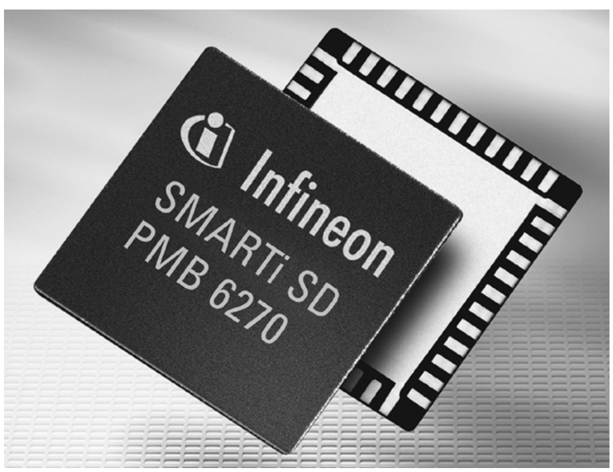

Figure 3.9: A zero-if GSM/GPRS receiver on a single chip made by Infineon

The infineon[2] website has the following to say about this chip:

> [It] integrates into a single chip all functions required of a GSM/GPRS quad-band transceiver, including an RF VCO, Zero IF receiver and Sigma-Delta transmitter chain. Its fast locking time and low

[2]http://www.infineon.com/wireless

power consumption makes it an ideal solution for GPRS class 12 implementation, as well as low-cost, voice-only phones. The SMARTi SD is entirely controllable through a 3 wire bus and is housed in a VQFN-48 standard package.

3.3.3 Negative aspects

The direct-conversion receiver has some performance problems. These are mainly caused by 1/f noise and self-mixing/dc-offset issues [32]. These two aspects are discussed in the following paragraphs.

1/f noise

In a direct-conversion receiver most of the gain is generated in the LF stage of the receiver. This stage is operated near 0 Hz where $\frac{1}{f}$-noise is very prominent[3]. The $\frac{1}{f}$-noise will be amplified by the LF stage amplifiers together with the desired signal, thereby significantly increasing the noise figure of the receiver in comparison to receivers that have a different gain distribution[4].

Self-mixing and DC-offset

Self-mixing [3] in a direct-conversion receiver occurs when energy from the BFO can leak into the antenna or the input port of the mixer. Even when the input port of the mixer is well shielded from the BFO's signal, the antenna might pickup some of the BFO's energy. This energy is not filtered out by the pre-mixer bandpass filter because the filter is centered around the BFO frequency. The BFO signal is therefore present at both inputs of the mixer as shown in Figure 3.10.

This self-mixing leads to a DC-offset voltage at the output port of the mixer. This DC-offset voltage is often much larger in amplitude than the signal of interest. So, the DC-offset voltage will swamp the desired signal.

To clarify the self-mixing process, take a look at the basic mixer formula shown in Equation 3.3.

$$s_{out}(t) = s_{in}(t) \cdot s_{bfo}(t) \tag{3.3}$$

[3] $\frac{1}{f}$ noise occurs in MOS and other semiconductor materials.

[4] i.e. the super heterodyne receiver

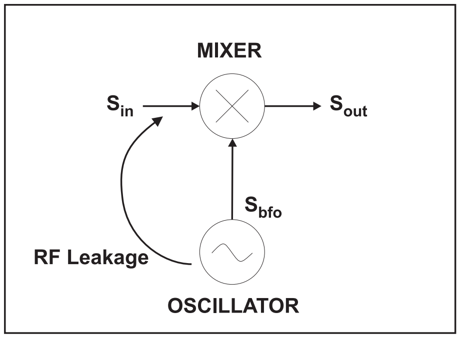

Figure 3.10: Self-mixing in direct conversion receivers

where $s_{in}(t)$ is the signal at the RF port of the mixer and $s_{bfo}(t)$ is the BFO signal at the LO port of the mixer.

The BFO signal will also be present at the RF port of the mixer through direct leakage or antenna leakage. The amplitude of this BFO signal will be smaller than that at the LO port by a factor of $c_{leakage}$.

$$s_{in}(t) = c_{leakage}\cos(\omega t) \tag{3.4}$$

$$s_{bfo}(t) = \cos(\omega t) \tag{3.5}$$

$$\begin{aligned} s_{out}(t) &= (s_{in}(t) \cdot s_{bfo}(t)) = c_{leakage}\cos^2(\omega t) \\ &= \frac{c_{leakage}}{2}(1+\cos(2\omega t)) \end{aligned} \tag{3.6}$$

$$\text{with } 0 \leq c_{leakage} < 1$$

The second harmonic, $cos(2\omega t)$, at the mixer output is usually filtered out by the post-mixer filter, leaving only the D.C. component. This D.C. component is sometimes called the DC-offset. The DC-offset cannot be filtered out by the post mixer filter if the receiver is required to have a good DC response[5].

When constructing a direct-conversion receiver it is important to minimize the $c_{leakage}$ factor. This can be achieved by employing heavy shielding around the BFO circuit. But even the most elaborate shielding may not solve the DC-offset problem because some leakage will always be present.

[5]There are solutions to this problem but all methods known to the author involve adding an extra mixer [9] [8].

When the direct-conversion receiver employs a high-level mixer, the leakage problem is made worse because the local oscillator drive level must be very high[6]. In this case, the BFO drive level can be several orders of magnitude higher than the signal being received making it easy for the BFO to swamp the desired signal.

A simple way to solve the DC-offset problem is to employ AC coupling of the mixer output. Of course, this is only possible if a DC response is not required. Many wireless communication standards provide for this by specifying a 'hole' near DC.

3.3.4 Summary

The direct-conversion receiver is a compact, low-power receiver architecture that lends itself well to implementation in integrated circuit form. Some problems of the architecture, like DC-offsets and inseparable sidebands can be solved using additional hardware or by circumventing them through clever modulation specifications, e.g. a 'hole' at DC.

3.4 The super-heterodyne receiver architecture

The super-heterodyne receiver differs from the direct-conversion receiver in that the former has one or more *intermediate frequency*, or IF, stages.

First the receiver layout is presented, followed by the positive and negative aspects of this architecture. Finally, a summary is given.

3.4.1 Receiver layout

A block diagram of a super-heterodyne receiver is shown in figure 3.11.

The example receiver uses only one IF stage but super-heterodyne receivers that employ two or even three IF stages are common.

[6]A high-level mixer usually employs more than +10dBm of LO drive level.

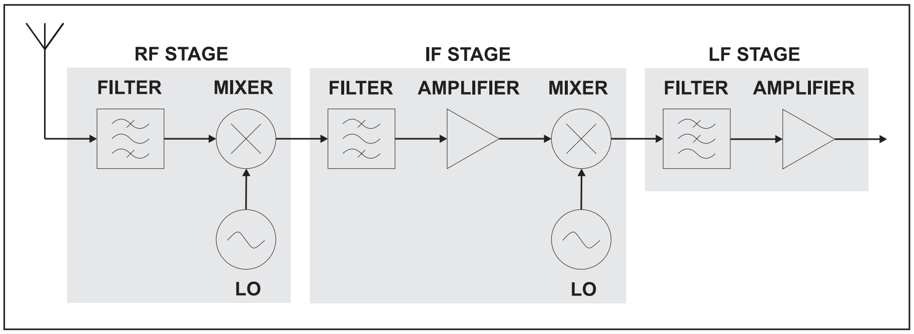

Figure 3.11: A super-heterodyne receiver

3.4.2 Positive aspects

The most compelling reason for using the super-heterodyne receiver architecture is its excellent selectivity. The selectivity is achieved by incorporating filters in every IF stage of the receiver. The super-heterodyne receiver is able to distinguish between the upper and lower sideband as it can pass one sideband while rejecting the other [25]. This feature is the main reason for its popularity.

While most of the gain in a direct-conversion receiver is concentrated the single LF stage, the super-heterodyne receiver the gain is distributed between one or more IF stages and the final LF stage. This means that the gain per stage can be lower which prevents oscillation due to RF leakage [25].

3.4.3 Negative aspects

The super-heterodyne receiver architecture has some issues associated with it. These are listed below.

- High parts count.
- Expensive, high-quality, filters needed.
- Not easy to integrate into an IC.
- Many mirror frequencies possible.
- More spurious responses.

In the following sections, these issues are discussed.

Parts count, cost and integration

The high parts count is because there are more stages compared to the direct-conversion receiver and more stages mean more parts are needed. The super-heterodyne receiver achieves its excellent selectivity through use of filters. These filters can be very expensive especially if they are required to have steep skirts. Not only are these filters expensive, they are also difficult, if not impossible, to integrate into a single chip solution. Therefore, the super-heterodyne receiver does not lend itself well to a compact single-chip solution.

Mirror frequencies and spurious responses

When designing a super-heterodyne receiver care must be taken to avoid undesired, or spurious, responses [6] [31]. These responses occur because every mixer converts two frequencies to one single frequency, in the ideal case. Usually only one of the two input frequencies is desired and the other is considered a spurious response. Real-world mixer, however, are far from ideal and will produce the same output frequency for a number of input frequencies [31].

According to Vizmuller [31], any input frequency f_{RF} satisfying the following relationship may produce a spurious response:

$$\pm m \cdot f_{RF} \pm n \cdot f_{LO} = \pm f_{IF} \tag{3.7}$$

$$\begin{aligned} \text{,where } f_{RF} &= \text{any incoming frequency into the mixer's RF port} \\ f_{LO} &= \text{the local oscillator frequency} \\ f_{IF} &= \text{the desired intermediate frequency} \\ m &= \text{any integer} \\ n &= \text{any integer} \end{aligned}$$

Each $< m, n >$ pair results in two possible spurious frequencies[31]:

$$f_{RF1} = \frac{n \cdot f_{LO} - f_{IF}}{m} \tag{3.8}$$

$$f_{RF2} = \frac{n \cdot f_{LO} + f_{IF}}{m} \tag{3.9}$$

The above equations hold for every mixer used in a receiver. The more mixers a receiver uses, the more possible spurious responses. Automated design systems are used to select intermediate and local oscillator frequencies to minimize the number of spurious responses as analysis by hand is too complex.

3.4.4 Summary

Designing a good super-heterodyne receiver is a complex task because of the complex interaction of the spurious and mirror responses. A well designed receiver, however, has excellent selectivity and can distinguish between the upper and lower sidebands of a signal if necessary. This excellent selectivity, however, is expensive to achieve and does not lend itself well to single-chip integration.

3.5 Choosing a network node front-end architecture

From the summaries of both the direct-conversion and the super-heterodyne architectures it is clear that both architectures have problems. If integration and mobility weren't an important aspect of the emergency scenario, the super-heterodyne architecture would be favored due to its excellent selectivity.

However, in the emergency scenario, integration and mobility are very important issues. The transceiver must be kept as small and as light-weight as possible. This will help reduce the total weight each rescue worker has to carry and enables network nodes to be carried for extended periods of time.

An added advantage of the direct-conversion receiver architecture is that is it cheaper to implement because it has a lower component count compared to the super-heterodyne architecture.

4

Radio Resource Discovery

4.1 Introduction

This chapter explains the concept radio resource discovery and introduces spectrum usage patterns. By detecting unused parts of the spectrum and applying models of the spectrum usage patterns an activity map is created. With this activity map, a network transceiver is able to exploit unused or under utilized parts of the radio spectrum. Two radio resource discovery architectures for generating an activity map are presented and discussed.

4.2 What is radio resource discovery?

Radio resource discovery is the process of finding unused or under utilized radio resources. In this thesis the radio resource of interest is bandwidth. Here, the object of radio resource discovery is 'recycling' bandwidth that has already been assigned, or allocated, to a primary user[1] (PU).

The 'recycling' is done by placing another user, the secondary user (SU), in the spectrum allocated to the PU. In principle, when a part of the radio spectrum has been allocated to such a PU, no other party may use this spectrum unless the PU gives permission. It is assumed that the PU has authorized the use of its allocated spectrum by a radio resource discovery enabled network.

The PU will, of course, only allow a SU access to their spectrum is the

[1]'Primary user' may also mean more than one physical user or person

SU can guarantee it will not cause interference[2].

In order to minimize the chance of interference to a primary user, the radio resource discovery system must not only find unused parts of the radio spectrum but must also predict how long these parts will remain unused. This prediction is based on the primary user's usage pattern.

4.3 Spectrum usage patterns

There are many types of PU, e.g. TV broadcast stations, WLAN systems, aeronautical communication facilities, navigational systems etc. All these types of PU use their spectrum at different times; they have different *spectrum usage patterns.*

Some PUs, like TV or radio stations, broadcast a fixed number of hours each day while other PUs broadcast short messages and have long idle periods, like aeronautical communications. Some PUs, WLAN for instance, might use their allocated spectrum fully.

A third type of usage pattern that exists is that of extremely long idle periods. This pattern may even rephrased as a not-used-at-all pattern. Although a part of the spectrum has been allocated to a PU, the PU does not occupy this spectrum fully.

In order keep the interference to the PU low, the SU must have an accurate way of predicting the usage of the PU. This prediction is based on a model of the usage pattern of the PU. Such a model is either known in advance and provided to the resource discovery system or it is formed by the resource discovery system itself.

The forming of a model by the resource discovery system is the most flexible of the two solutions because it may be an adaptive model. It is highly likely that a PU's usage pattern will change over time, for instance, when a TV station reschedules it's broadcasting times.

How to form such a model and what the model should consist of is an item for further research. From here on, it is assumed that such a model exists and that the model predicts the duration of the PU's inactivity correctly.

Now, the resource discovery system's task is to find the unused parts of the spectrum and let the model predict the amount of time the parts will

[2] In [28], it is claimed that the SU must guarantee an interference probability of less than 0.1%.

remain unused. The following section will focus on detecting unused parts of the radio spectrum.

4.4 Detecting unused parts of the spectrum

An unused part, or *band*, of the spectrum of bandwidth W contains no transmissions by definition, but it will contain noise. This noise will consist of thermal noise and man-made noise [27]. Man-made noise emanates from household appliances, cars and other equipment that uses electricity. When a transmission takes place in this band, the energy from the transmitter will add to the total power contained in this band.

The resource discovery system must make a binary decision regarding activity in the band; either there is activity within this band or there isn't. The binary decision is made by measuring the total power contained in the band and applying a threshold. If the power is below the threshold, the band is unused, if the power is above the threshold, the band is occupied.

The performance of this method depends on the accuracy of the power measurements and the threshold level. The choice where to put the threshold level is a well studied decision theory problem. The threshold level depends on the distribution functions of the 'noise-only' and 'noise+transmission' cases. Here, the distributions are considered to be unknown and a different method was used to set the threshold level.

Throughout this thesis, the power threshold level is set 3 dB higher than the noise-only case[3]. This means that a transmission must raise the total power contained in the measurement band by 3 dB before it is detected.

A band that is unused is said to be *available*. The fact that a band is available, does not mean it can actually be exploited by a secondary user. The band must be accessible as well. The following section discusses this band *accessibility*.

4.5 Band accessibility

A network transceiver is not always capable of exploiting the fact that a certain band is available. The most obvious case is where the primary user

[3] The measurement band contains only thermal and man-made noise.

refuses to give permission to access this band. For instance, the radio astronomy bands will be off limits because any transmission will interfere with the bands purpose. Also, some PU systems are upset by SU transmission during idle periods. These PU systems mostly have carrier-detect based MAC layers.

There are two cases where a band is inaccessible. The first case is when the primary user's usage pattern is unpredictable. Any transmission by the SU in the band is likely to interfere to an unacceptable degree. The second case arises when the PU's transmitters are switched on and off in succession so quickly, it is impossible for the SU's to fill the empty spaces in between the PU's transmissions. This situation might occur in bands where the PU is a high-speed WLAN.

Whether a band is accessible or not depends, in large parts, on the protocols used by the SU's network nodes. The shorter the SU's transmissions, the more likely it is an available band can be accessed.

Because of this dependency of accessibility on the MAC layers of the network, determining accessibility is not the responsibility of the resource discovery system but of the network transceiver itself. The network transceiver is simply given an *activity map* by the resource discovery system.

4.6 The activity map

Before discussing the *activity map*, the *channel* concept must be introduced. In the preceding sections, the word *band* was used to describe a part of the radio spectrum of arbitrary bandwidth and frequency location. Channels, however, have a certain bandwidth and frequency location.

The entire transceiver frequency range is divided into channels of equal bandwidth. These channels do not overlap and there are no frequencies in the transceiver frequency range that do not belong to a channel.

The *activity map*, sometimes called an *occupancy vector* [28], contains an overview of the activity in the entire transceiver frequency range. A very basic activity map contains a binary digit, or *bit*, for every channel. If the bit is set to '1', the channel it refers to is occupied. If the bit is set to '0', the channel is available.

A more complex activity map contains information on the expected duration of the current activity state. In case of an available channel, this information is used by the network node to determine the accessibility of the channel and the amount of time it may spend exploiting that channel.

4.6.1 Frequency resolution

In the above discussion of the activity map each channel has a single bit to represent the availability of the entire channel. The frequency resolution of this information is equal to the channel bandwidth. This might be considered too coarse for some situations.

Some modulation schemes, like OFDM [28], allow individual sub-carriers to be switched off to create notches in their symbol spectrum. This property could be used by the secondary user to place notches at positions in the spectrum of the channel where a primary user is transmitting thereby eliminating or reducing interference thus allowing the channel to be exploited. The SU's network nodes can only achieve this if the frequency resolution of the activity map is increased to the same level as the modulation scheme has control over.

The frequency resolution of the activity map need not be higher than the precision with which the modulation scheme can place or manipulate its own spectral emissions. A frequency resolution higher than this only leads to higher computational effort.

4.6.2 Time resolution

The time resolution of the activity map is equal to the amount of time it takes the resource discovery system to scan the complete transceiver frequency range and process the power measurements. Only after a complete scan can an updated activity map be generated.

At this time further research is needed to determine what the ideal time resolution would be. However, it is clear that the required time resolution is related to the usage patterns of the primary users that operate in the frequency range of interest.

4.7 Intermezzo

In the preceding sections the concepts behind radio resource discovery process were explained and usage patterns were discussed. Also, the use of power measurements in the discovery process was introduced.

No mention was made how the power measurements are made and what radio architectures are suitable for radio resource discovery and network operations at the same time. In the next sections these issues are discussed and a practical example of a resource discovery algorithm is given.

4.8 Network node architectures for radio resource discovery

A network node must perform radio resource discovery and take part in normal network operations. There are two network node radio architecture categories that can perform these tasks. The first category is the single transceiver architecture category and the second is the separate resource discovery receiver category. Each of the two categories can be subdivided into broadband and narrowband subcategories.

First the two main categories are presented in separate subsections after which the broad- and narrowband subcategories are discussed in a third subsection.

4.8.1 Single transceiver architecture

The first configuration is shown in figure 4.1. Here the normal network transceiver is used to determine the available spectrum for free spectrum resources. During a network operation cycle, the transceiver performs its normal network tasks. During idle periods of the network, the transceiver goes into its resource discovery mode.

This configuration is economical because it utilizes the transceiver hardware that is already present for normal network operation. This reduces both the cost of the network node and the power consumption because there is only one transceiver to operate.

A drawback of this configuration is that when the transceiver is used for normal network operation, it can't be used for spectrum resource discovery. The feasibility of this architecture depends on the required time resolution of the activity map.

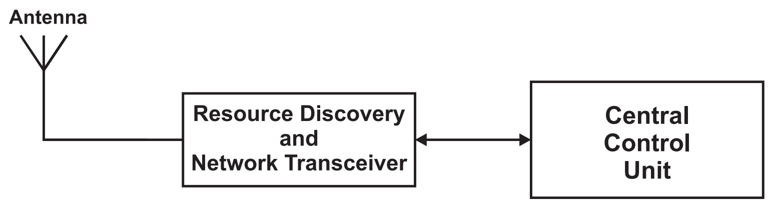

Figure 4.1: The single transceiver architecture for radio resource discovery

4.8.2 Separate resource discovery receiver architecture

The second and more elaborate configuration is depicted in figure 4.2. In this configuration an extra resource discovery receiver is added to free the normal network transceiver of resource discovery duties. A splitter/combiner is used to split the antenna signal to feed the network transceiver and the resource receiver. With this configuration it is possible to do continual spectrum resource discovery, even when the network transceiver is operational.

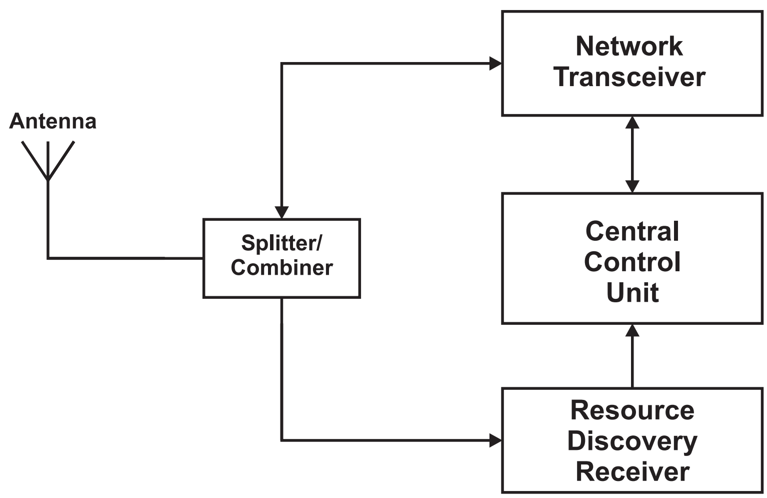

Figure 4.2: The separate resource discovery receiver architecture

A major design challenge with this configurations lies in the fact that when the normal network transceiver is transmitting, it will overload the resource discovery receiver's input stages unless some precautions are taken. The splitter/combiner must provide enough isolation between the network transceiver and the resource discovery receiver. The amount of isolation needed by this configuration depends on the network transceiver's output power and the resource discovery receiver's ability to handle high power signals. So, to avoid overloading the resource discovery receiver, its 1 dB-compression point should be very high.

In the emergency scenario, the network transceiver's output power is +30 dBm. Say that the splitter/combiner provides 20 dB of isolation between the network transceiver and the resource discovery receiver. In that case, the resource discovery receiver must be able to handle at least +10 dBm signals at the antenna, ideally without reducing the receiver's sensitivity.

4.8.3 Broadband and narrowband subcategories

Both the separate resource discovery receiver architecture and the single transceiver architecture can be placed in a broadband or narrowband subcategory.

Broadband resource discovery systems can observe the entire transceiver frequency range, while narrowband discovery systems can only observe part of the transceiver frequency range and require tuning the resource discovery receiver to cover the entire range.

So, a broadband separate resource discovery receiver is able to monitor the entire transceiver frequency range continuously. While the narrowband receiver must operate in a frequency *hopping*, or *scanning* mode so observe the complete transceiver frequency range.

While a broadband single transceiver architecture can observe the entire frequency range, it cannot observe it continuously because the resource discovery facilities are unavailable when the transceiver is required for normal network operation duties.

4.8.4 Which architecture?

The ideal architecture primarily depends on the required activity map time resolution. If continuous resource discovery is required, the separate resource discovery architecture is the only option. However, if continuous discovery is not a requirement, then the single transceiver architecture is a better alternative because of its lower power and low cost.

The choice of narrowband or broadband is determined largely by the feasibility of a broadband implementation. Mostly, a broadband solution is infeasible because of the need for high-speed, say ¿100 MSPS, A/D converters. These are expensive and power hungry and therefore not very attractive for cheap mobile network nodes.

4.9 Power measurement algorithms

Power measurements form the basis of radio resource discovery systems. In this section two methods for power measurement are introduced. Both methods sample the output of the receiver as shown in Figure 4.3.

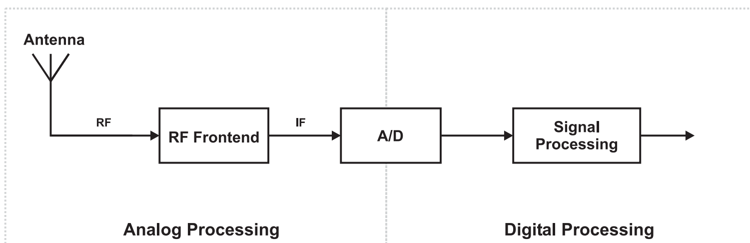

Figure 4.3: System model of a spectrum resource discovery receiver

In case of a direct-conversion receiver, the output of the receiver is already at baseband. If the receiver is a super-heterodyne, the IF signal is sampled and converted to baseband in the digital domain, see Frerking [11] for suitable methods.

The first method of power measurement measures the total power contained in the baseband signal coming from the receiver. It is referred to as the *coarse* method due to its coarse frequency resolution.

The second method uses the fourier transform to obtain power measurement data with higher frequency resolution and it therefore referred to as the *fine* method.

4.9.1 Coarse power measurement

Consider a digitized baseband signal, $x(n)$ consisting of N *real* samples. To calculate the RMS value of the signal, the sample values are first squared, then averaged and finally the square root is taken. This process is expressed by equation 4.10

$$RMS = \sqrt{\frac{1}{N} \cdot \sum_{n=0}^{N} x(n)^2} \tag{4.1}$$

The equation above can be graphically represented by the block diagram shown in figure 4.4.

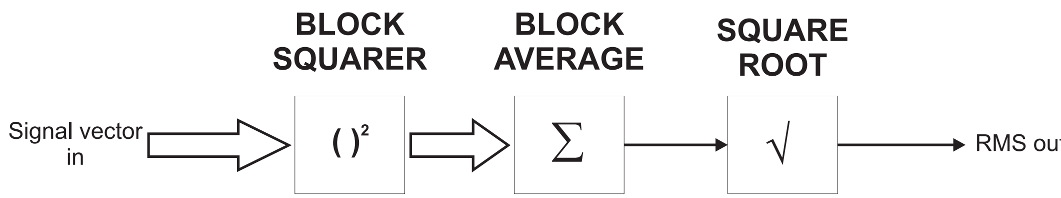

Figure 4.4: Per-sample based RMS measurement

The frequency resolution of this methods is equal to the bandwidth of the digitized baseband signal.

Analysis

To analyze this RMS measurement, consider a sinusoidal input with amplitude a, frequency θ and phase ϕ:

$$x(n) = a \cdot \sin(\theta n + \phi) \tag{4.2}$$

Then square this signal to obtain the result of the first processing step, $x(n)^2$:

$$\begin{aligned} x(n)^2 &= a^2 \cdot \sin^2(\theta n + \phi) & (4.3)\\ &= a^2 \cdot \frac{1 - \cos(2\theta n + 2\phi)}{2} \\ &= \frac{1}{2}a^2 - \frac{1}{2}a^2 \cos(2\theta n + 2\phi) & (4.4) \end{aligned}$$

Note that there are two parts, a constant part and a sinusoidal part of twice the original frequency. To analyze the averaging process, look at the two parts individually:

$$\frac{1}{N}\sum_{n=0}^{N}\frac{1}{2}a^2 = \frac{1}{2}a^2 \qquad \forall N > 0 \tag{4.5}$$

and,

$$\frac{1}{N}\sum_{n=0}^{N}\frac{1}{2}a^2\cos(2\theta n + 2\phi) \to 0 \qquad \text{as } N \to \infty \tag{4.6}$$

When the time over which the signal is averaged is long, the sinusoidal part will approach zero. Therefore it has almost no influence when N is large. The only contributing part is the part shown in equation 4.5.

Then finally, the square root of this part is taken, which gives the following result when simplified:

$$RMS\Big|_{x(n)=a\cdot\sin(\theta n+\phi)} \approx \frac{1}{2}\sqrt{2}\cdot a \qquad \text{for large } N \tag{4.7}$$

The resulting RMS measurement is thus proportional to the sinusoid's amplitude. Note that this is not a power measurement but a voltage measurement. The average power can be calculated using Equation 4.8.

$$P_{av} = \frac{V_{peak}^2}{2R} \tag{4.8}$$

where P_{av} is the average power, V_{peak} is the sinusoid's peak voltage and R is the load into which the signal's current flows.

First convert V_{RMS} into V_{peak} by multiplying V_{RMS} by $\sqrt{2}$. Then the power into a 1Ω load will be:

$$P_{av,1\Omega} = \frac{(\sqrt{2}V_{RMS})^2}{2} = V_{RMS}^2 \ [W] \tag{4.9}$$

Per-sample based algorithm

The method outlined above is *block oriented*. This means that an entire block of data must be sampled before the algorithm can do its work. With slight modification however, the above method can be applied on a per-sample basis. This modification lies in the observation that the averaging done by equations 4.5 and 4.6 is to remove the sinusoidal contribution. This can be viewed as lowpass filtering.

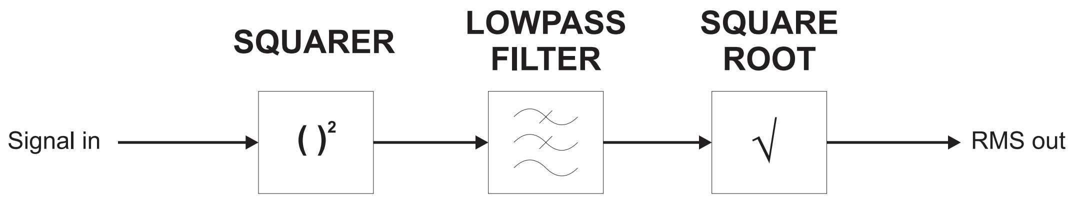

Figure 4.5: Per-sample based RMS measurement

When the block averaging is replaced by a suitable lowpass filter, a per-sample version of the algorithm is obtained, shown in figure 4.5. A suitable lowpass filter is one which has no dampening or amplification of the DC-term and which has infinite dampening of the second harmonic of the signal being measured.

The best cutoff frequency for the lowpass filter depends on the application. Making the cutoff frequency too low will result in slow convergence of the algorithm. Making the the cutoff frequency too high will result in too much harmonic leakage adding noise to the measurement.

Some comments

Both the block and the sample based algorithms take the entire input spectrum into account. If full spectrum measurement is not desired, a bandpass filter can be used in front of the algorithm to select the frequencies of interest.

Note that the block and sample based algorithms accept real valued signals only. If the sampled signal from the receiver is complex valued, the block based algorithm equation is as follows:

$$RMS = \sqrt{\frac{1}{N} \cdot \sum_{n=0}^{N} |\mathbf{x}(n)|^2} \tag{4.10}$$

where $\mathbf{x}(n)$ now denotes a complex-valued signal.

4.9.2 Fine power measurement

The coarse signal measurement outlined above produces only one RMS value for a large chunk of spectrum. This represents the complete intermediate frequency bandwidth. When a more fine grained spectral measurement is needed, the coarse method is inefficient because it would require the use of many bandpass filters, one for each subband. In this case a better solution is needed and is found in the application of the *Discrete Fourier Transform* [24], or DFT[4].

$$X(k) = \frac{1}{N}\sum_{n=0}^{N-1} x(n)e^{-\frac{j2\pi kn}{N}} \tag{4.11}$$

When the input vector x(n) consists of N points $x(0), x(1), ..., x(N-1)$, the frequency-domain representation is given by the set of N points $X(k), k = 0, 1, ..., N-1$.

The discrete fourier transform can be calculated efficiently by the fast fourier transform (FFT) algorithm which is a well-known digital signal processing primitive.

After transforming the input samples into the frequency domain, the magnitude of each frequency component is calculated and converted into an RMS value. This process is shown in figure 4.6.

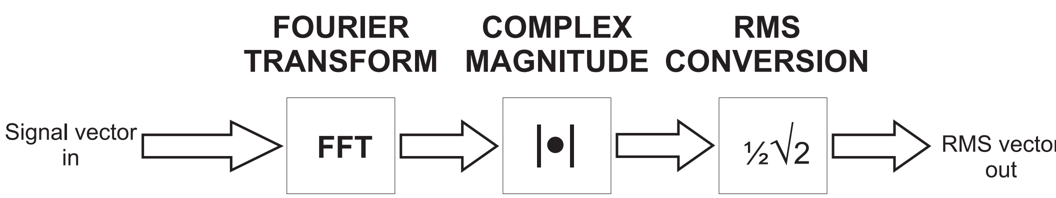

Figure 4.6: An FFT based RMS measurement

Analysis

Consider an input signal $x(n) = a \cdot e^{\frac{j2\pi n}{N}+j\theta}$ which is a complex sinusoidal signal with phase θ. This signal is transformed using the DFT:

$$\begin{aligned} X(1) &= \frac{1}{N}\sum_{n=0}^{N-1} a \cdot e^{\frac{j2\pi n}{N}+j\theta} \cdot e^{-\frac{j2\pi n}{N}} \\ &= \frac{1}{N}\sum_{n=0}^{N-1} a \cdot e^{j\theta} \\ &= a \cdot e^{j\theta} \end{aligned} \tag{4.12}$$

[4]Note that here the DFT is normalized by multiplying by $\frac{1}{N}$

$$
\begin{aligned}
X(k)|_{k\neq 1} &= \frac{1}{N}\sum_{n=0}^{N-1} a \cdot e^{\frac{j2\pi n}{N}+j\theta} \cdot e^{-\frac{j2\pi kn}{N}} \\
&= \frac{1}{N}\sum_{n=0}^{N-1} a \cdot e^{\frac{j2\pi(1-k)n}{N}} e^{j\theta} \\
&= \frac{a \cdot e^{j\theta}}{N}\sum_{n=0}^{N-1} e^{\frac{j2\pi(1-k)n}{N}} \\
&= 0 \qquad (4.13)
\end{aligned}
$$

Now, taking the magnitude of $X(1)$ the following result is produced:

$$
\begin{aligned}
|X(1)| &= |a \cdot e^{j\theta}| \\
&= a \cdot |e^{j\theta}| \qquad \text{because } a \text{ is real} \\
&= a \qquad (4.14)
\end{aligned}
$$

Thus, the DFT measures the peak amplitude of the sinusoid. To convert this amplitude to it's average power P_{av} apply the following formula:

$$P_{av} = \frac{V_{peak}^2}{2R} \qquad (4.15)$$

where P_{av} is the average power, V_{peak} is the peak amplitude and R is the load into which the signal's current flows. Thus, the power into a 1Ω load will be:

$$P_{av} = \frac{V_{peak}^2}{2}\ [W] \qquad (4.16)$$

Some comments

The resolution of the DFT depends on the number of points used. The more points used, the higher the frequency resolution but the lower the time resolution. This is a tradeoff with the DFT.

When a signal is analyzed by the DFT, the spectrum is divided into *bins*. Each bin has a fixed bandwidth which is equal to $\frac{f_s}{N}$, where f_s is the sampling frequency and N is the number of points used in the DFT. A bin with index k is centered around a frequency $f(k)$ as shown by equation 4.17.

$$f(k) = k \cdot \frac{f_s}{N} \tag{4.17}$$

Measuring power using the fine method has advantages over the coarse method besides increase frequency resolution. The noise power contained in a bin is less than the total noise power because the bandwidth is smaller[5]. If the power measurement is performed using N bins, then the noise power (per bin) is $\frac{1}{N}$ lower than in the full bandwidth case. So, this increases the sensitivity of the measurement.

A problem with the discrete fourier transform with a finite number of points is *spectral leakage*. Even when a single frequency is present in the input spectrum, the bins of the DFT react to this frequency if the DFT length isn't an exact multiple of the signal's wavelength. This is because the signal is truncated to a finite sequence using a rectangular weighting function, or *window*. By using a different weighting function to taper the finite sequence, the spectral leakage can be reduced. Well known window functions are the Hamming, Hanning and Blackman windows [24].

4.10 A resource discovery algorithm

The resource discovery algorithm presented here assumed that the resource discovery receiver is of the narrowband kind. Also, the algorithm is suitable for a single transceiver architecture because the resource receiver is not required to be operational continuously.

The algorithm is based on the coarse power measurement technique and assumes that the resource discovery receiver can only observe one channel at a time. The simplest form of activity map is generated; there is only availability information. To avoid interference problems caused by an outdated activity map, the algorithm tries to find channels that have no activity at all.

The algorithm uses four lists to classify the state of each channel. It does this by storing the *channel number* in one of the lists. Thus, a channel[6] can only be present in one list at a time. The lists are:

- Occupied list
- Candidate list

[5] Assuming that N_0 is independent of frequency.

[6] From here on, the term channel and channel number are used interchangeably.

- Longscan list
- Available list

Channels that reside in the occupied, candidate and longscan lists are marked occupied in the activity map and channels that reside in the available list are marked available in the activity map.

Each list has an procedure associated with it. When a procedure is activated, it removes the channel from the top of the list that is associated with it and performs certain operations.

The basic principe of the resource discovery algorithm is that channels that contain no activity will propagate from the occupied list to the candidate list on to the longscan list and finally to the available list. Channels that do contain activity will not propagate to the available list and get redirected to the occupied list.

The pseudo-code of each list procedure is given and discussed in the following subsections. After these four subsections a final subsection describes the calling order of the list procedures.

4.10.1 Occupied list procedure

The occupied list procedure removes the channelID at the top of the occupied list, tunes the receiver to this channel and measures the power contained in the channel using the coarse power measurement algorithm.

The power value is then compared against the threshold value. If the power value is greater than the threshold level, the channelID is placed at the bottom of the occupied list. If the power value is below the threshold, the channel is promoted by placing it at the bottom of the *candidate list*. A 'scancount' variable which is unique to every channel is reset to zero for use in the candidate procedure.

The pseudo-code of the occupied list procedure is as follows:

```
BEGIN

  channelID = PopTopOfOccupiedList();
  TuneResourceReceiver(channelID);
  power = MeasurePower();
  scancount = 0;

  IF (power > threshold)
    THEN AddAtBottomOfOccupiedList(channelID);
    ELSE AddAtBottomOfCandidateList(channelID);

END
```

4.10.2 Candidate list procedure

When a channel is present in the candidate list it means that the power in the channel was below the threshold at least once. A channel will be promoted to the longscan list when the channel survives the power thresholding test ten times. The number of times this threshold test was performed is kept in the 'scancount' variable which is unique to each channel. When the channel does not survive a threshold test, the channel is transferred to the occupied list.

A power thresholding test is only performed once per call to the candidate list procedure, so the procedure must be called at least ten times before a channel can be promoted to the longscan list.

The idea behind this procedure is to take 10 non-consecutive power measurements of the channel over a period of time.

The pseudo-code of the candidate list procedure is as follows:

```
BEGIN

  channelID = PopTopOfCandidateList();
  TuneResourceReceiver(channelID);
  power = MeasurePower();

  IF (power > threshold)
    THEN AddAtBottomOfOccupiedList(channelID);
    ELSE
    {
      IF (scancount > 9) AddAtBottomOfLongscanList(channelID);
      ELSE
      {
        scancount = scancount + 1;
        AddAtBottomOfCandidateList(channelID);
      }
    }

END
```

4.10.3 Longscan list procedure

The longscan list procedure takes consecutive channel power measurements during a fixed time period. This is to make sure that the channel doesn't contain a primary user that sends in very short bursts with long idle periods in between.

The procedure will continue measuring and thresholding the channel power until either a timer expired or until the channel power is greater than the threshold level. When the timer expires, the channel is promoted to the available list. But if the power in the channel is greater than the threshold, the channel is transferred to the occupied list.

The pseudo-code of the longscan list procedure is as follows:

```
BEGIN

  channelID = PopTopOfOccupiedList();
  TuneResourceReceiver(channelID);
  ResetTimer();

  WHILE(Timer != Expired)
  {
    power = MeasurePower();

    IF (power > threshold)
      THEN
      {
        AddAtBottomOfOccupiedList(channelID);
        EXIT;
      }
  }
  AddAtBottomOfAvailableList(channelID);

END
```

4.10.4 Available list procedure

The purpose of the available list procedure is very simple. As long as the power within a channel is below the threshold, the channel remains in the available list. If the power is above the threshold, the channel is transferred to the occupied list.

The pseudo-code of the available list procedure is as follows:

```
BEGIN

  channelID = PopTopOfAvailableList();
  TuneResourceReceiver(channelID);
  power = MeasurePower();

  IF (power > threshold)
    THEN AddAtBottomOfOccupiedList(channelID);
    ELSE AddAtBottomOfAvailableList(channelID);

END
```

4.10.5 List procedure calling order

The calling order of the list procedures is as follows:

- Occupied list procedure
- Candidate list procedure
- Longscan list procedure
- Candidate list procedure
- Available list procedure
- Candidate list procedure

This calling order puts more weight on the candidate list procedure because a single channel must be processed at least 10 times by this procedure before it gets promoted to the longscan list. This weighting scheme will allow channels to be promoted more quickly compared to a uniform weighting scheme.

4.10.6 Generating the activity map

Generating the activity map is a very simple process. All the channels in the available list are marked available while all the other channels are marked occupied.

This algorithm cannot make any guarantees on the accuracy of the activity map because it does not use a model of the primary user to predict when a channel marked available will become occupied again.

4.11 Summary

This chapter explains what radio resource discovery is and the concepts behind it. It discusses two different receiver architecture categories for radio resource discovery in mobile network nodes and introduces two bandwidth related architecture subcategories.

Two methods for power measurement were discussed, the coarse and fine methods. The coarse method forms the basis of a radio resource discovery algorithm which is also described in this chapter.

The radio resource discovery algorithm tracks the state of each channel by storing an ID for each channel in one of four lists. Each channel ID can only be present in one of there four lists. When a channel contains no activity, it gets promoted from the occupied list until it reaches the available list. To generate the activity map, all the channel with IDs in the available list are marked available while all other channels are marked occupied in the map.

4.12 Conclusions and directions

5

A Low-cost Resource Discovery Receiver

5.1 Introduction

A low-cost resource discovery receiver was designed based on a commercially available television tuner, a Tornado E67 DSP development system and some custom circuitry.

The receiver's purpose is to take power measurements across a large frequency range and store the data. The measurement data can be used for further research into spectrum usage patterns and primary user model development.

The receiver was also developed as demonstration platform for resource discovery algorithms like the one described in Section 4.10.

First, a detailed description is given of the hardware, then the software containing the resource discovery algorithm and user interface is presented.

5.2 Receiver hardware

The receiver is based on a Philips FM1246 television tuner [22] salvaged from an E-TECH PCI TV-Card that enables PCs to receive ordinary television signals.

The next sections will describe the custom hardware and tuner module in more detail.

5.2.1 Tuner module description

The tuner module is a self-contained super-heterodyne frontend with one mixing stage. The local oscillator is based on a phase locked loop (PLL) with a maximum lock-in time of 150 ms [22]. A block diagram of the tuner is shown in Figure 5.1.

The tuner module has two antenna inputs, on input is for TV signals and the other is for FM radio. In the low-cost receiver, only the TV signal input is used. The TV-input frequency range is from 45.75 to 855.25 MHz and the tuning is controlled through the I2C interface [23].

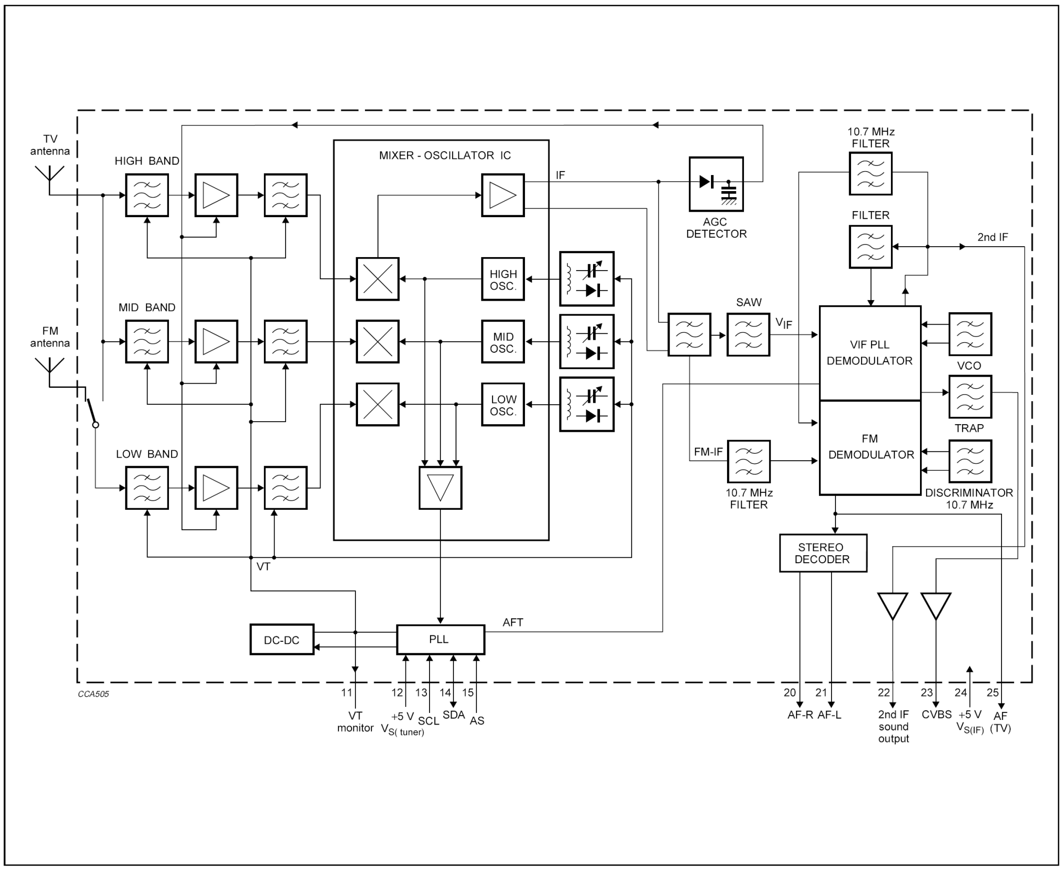

Figure 5.1: Philips FM1246 TV tuner block diagram [22]

The I2C interface of the tuner is controlled by the Tornado E67 through it's general purpose I/O port (GPIO). This GPIO port of the Tornado E67 is not I2C bus compatible and requires extra interface logic.

5.2.2 The Philips I2C bus

The I^2C or I2C bus is a two-wire serial bus developed by Philips. One wire is used as the clock and the other the data line. The I2C protocol allows multiple master and slave devices on the bus through a bus arbitration scheme based on a first-come-first-serve basis.

Each master or slave has a unique address which is eight bits long. A master device senses the bus to see if it is inactive. If the bus is inactive, the master initiates a data transfer to or from a slave by putting the unique address on the bus to inform the slave of a transfer. The slave responds to the request through an acknowledge bit. The rest of the bus conversation protocol depends on the slave device being accessed.

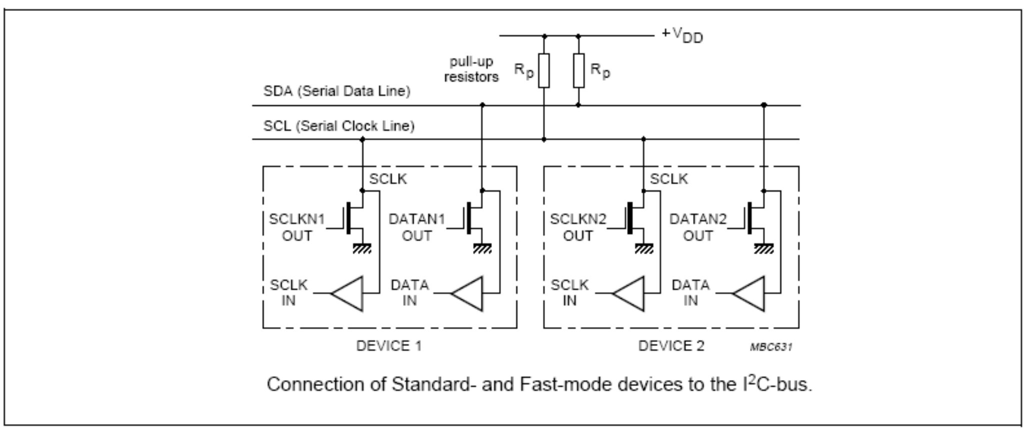

Figure 5.2: I2C bus driver interface [23]

The interface logic required for accessing the I2C bus is shown in Figure 5.2. Both the clock line (SCL) and the data line (SDA) are set to a logical "1" by the two pull-up resistors R_p. When a logical "0" is desired, the device controlling the line activates a transistor to connect the line to ground.

The currently active bus master may drive both the clock and data lines to ground. If a response is required from a slave, the slave drives the data line but the master controls the clock. For slow slave devices, the master could try to clock the slave too quickly. In such a case, the slave device will extend the low period of the clock by pulling the SCL line low. The master can detect this by sensing the clock line and comparing the result with it's own internal clock line state, any discrepancy indicates that extra wait-states are generated by the slave device.

For more information regarding the official I2C protocol, the reader is referred to the Philips website where the specification can be downloaded [23].

5.2.3 DSP to I2C glue logic

The tuner I2C pins are controlled by the DSP's general purpose 8-bit I/O port (GPIO) which is present on the serial controller. These I/O pins are not I2C bus compatible because they do not support the high-impedance state required to 'release' the lines for a logical "1".

Some glue logic based on an 74HCT244 bus driver was used to drive the I2C bus. This IC has two 4-bit bus driver sections that support a high-impedance state, see Figure 5.3. When the $\overline{OE}$ pin is a logical "1", the outputs $Y_{0...3}$ of the corresponding section are in an active drive state, that is, they follow their inputs $A_{0...3}$. When the $\overline{OE}$ pin is a logical "0", the outputs $Y_{0...3}$ enter a high-impedance state. This behavior is ideal for driving the I2C bus.

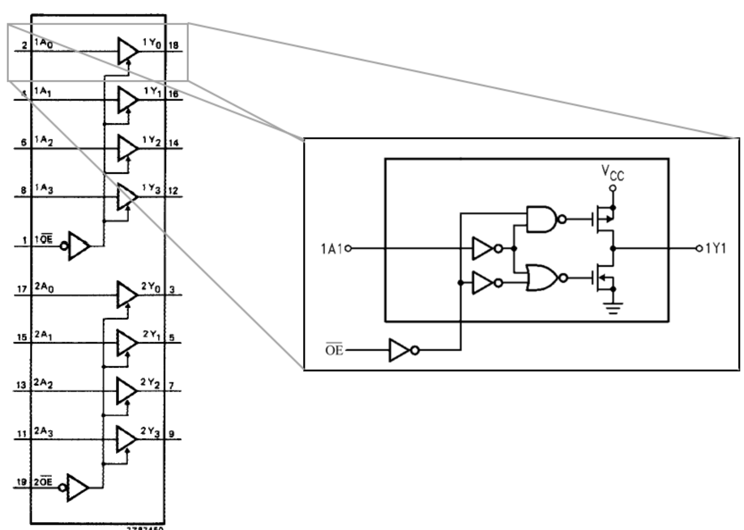

Figure 5.3: A diagram of the 74HCT244 and one output state

By connecting $1A_0$ and $2A_0$ to ground lets the $1\overline{OE}$ and $1\overline{OE}$ pins select between driving $1Y_0$ and $2Y_0$ low and the high-impedance state. In the receiver $1Y_0$ is connected to the SCL line and $2Y_0$ is connected tot the SDA line. Each line also has a 2 kΩ pull-up resistor to pull the lines high when the bus is released.

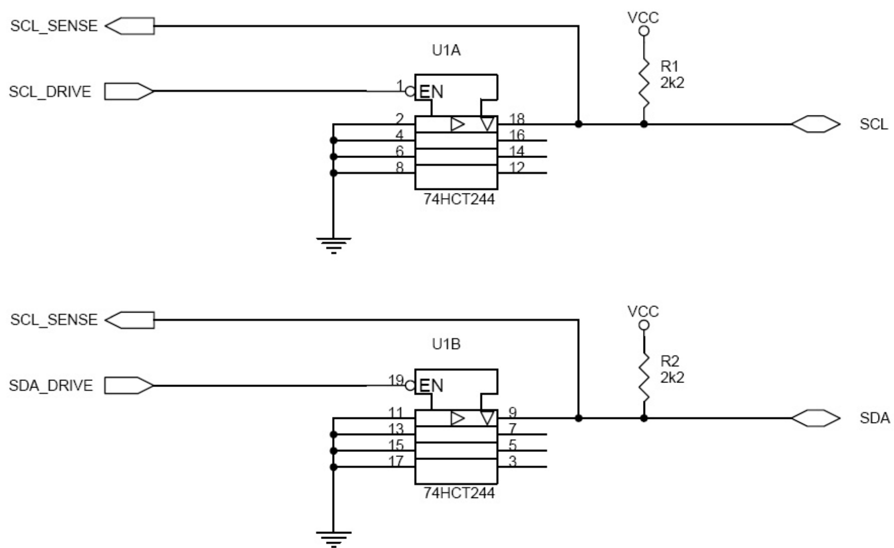

Figure 5.4: I2C DSP-to-tuner interface logic diagram

In addition to the driving the lines, the DSP must be able to sense the state of the lines. Sensing the SDA and SCL lines can be achieved by directly connecting them to 2 separate GPIO port pins configured as input. Sensing the SDA line allows data to be read from the I2C bus while sensing the SCL line allows detection of additional wait-states inserted by a slave device.

A schematic of the I2C interface logic is shown in Figure 5.4. The signals on the left hand side go to the DSP, the signals on the right go to the tuner. The signals to the DSP are connected to the development system through a Sub-D 9 connector. The pin connections are specified in the table below.

Signal name	Sub-D 9 pin
SCL_SENSE	1
SCL_DRIVE	2
SDA_SENSE	3
SDA_DRIVE	4
GROUND	9

Now the interface to the tuner has been explained, the tuner must be setup properly by the Tornado E67 using the I2C interface logic.

5.2.4 Setting up and tuning the tuner

The tuner consists of a phase locked loop based frequency synthesizer with a programmable divider that controls the tuning frequency. The desired PLL division factor is sent to the tuner through the I2C interface pins (SCL and SDA).

Through these pins other tuner parameters can be set, e.g. PLL charge-pump response and synthesizer step size. The tuner is configured as suggested by the FM1246 datasheet [22].

With the tuner under control of the Tornado E67, the IF SAW filter response is measured to see where the best place would be to sample a 1 MHz-wide chunk of spectrum. But before the actual measurement can begin, the tuner must be modified.

5.2.5 Measuring the IF SAW filter response

Unfortunately, the FM1246 tuner has an onboard IF processing section that decodes the audio subcarrier and the composite video signals from the intermediate frequency. This means that the undecoded IF signal is not available externally. A block diagram of the tuner is shown in Figure 5.1. Although a single output line is drawn from the SAW filter to the decoder chip, it is, in fact, a balanced output.

To allow access to the IF signal, the output of the IF filter was decoupled from the input of the decoder chip and attached to the two audio output pins (AF-R and AF-L). This modification is shown in Figure 5.5. The decoder chip no longer serves any function.

With the IF routed to the outside world, the frequency response of the SAW filter was measured. This was done by tuning the tuner to 50 MHz and

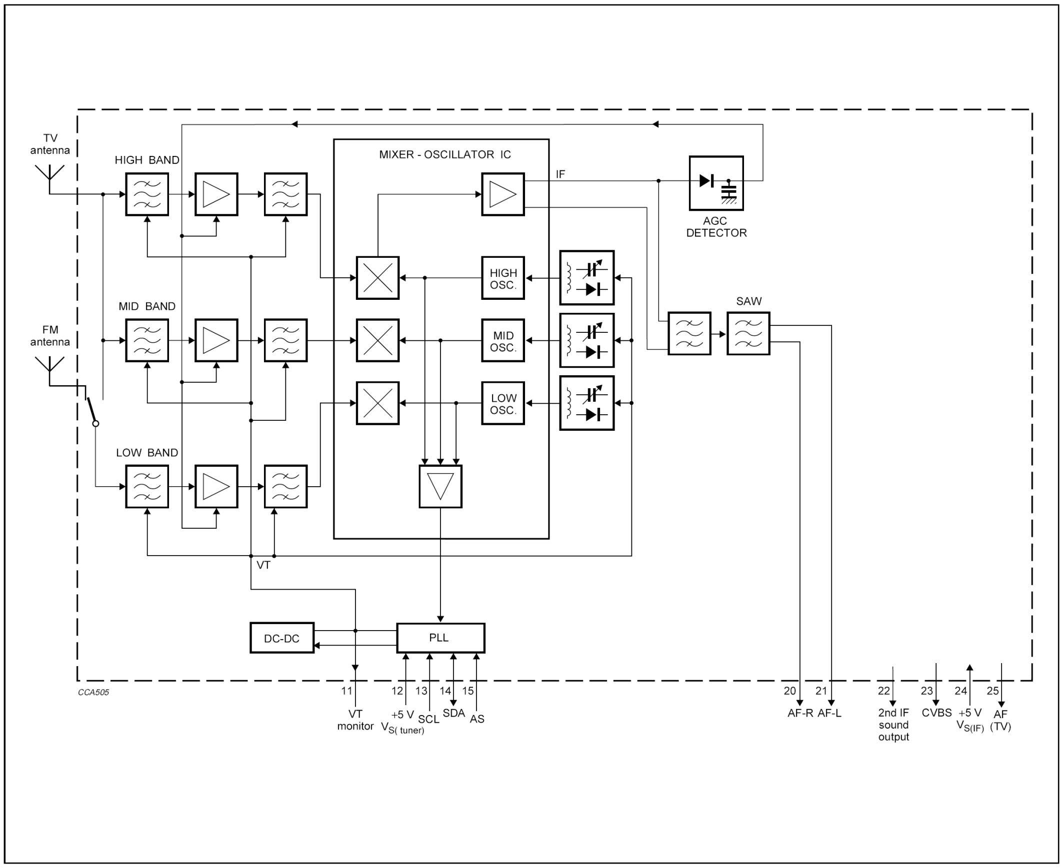

Figure 5.5: Modified Philips FM1246 TV tuner block diagram

applying a frequency sweep at the antenna. The sweep was made by an HP network analyzer. The differential output of the tuner was analyzed using a LeCroy digital oscilloscope which has a frequency analysis capability, see Figure 5.6 for a diagram.

A picture of the tuner undergoing SAW filter response testing is shown in Figure 5.7. The I2C interface, built on perfboard, can be seen mounted on the upper-left side of the ground plane. The tuner module itself is located to the right of the I2C interface.

The measured frequency response of the SAW IF filter is shown in 5.8. The results show that the IF runs from around 32 MHz to 40 MHz. The "double hump" observed on the left side of the SAW response is the audio subcarrier shelf required by the PAL television standard. The useful IF

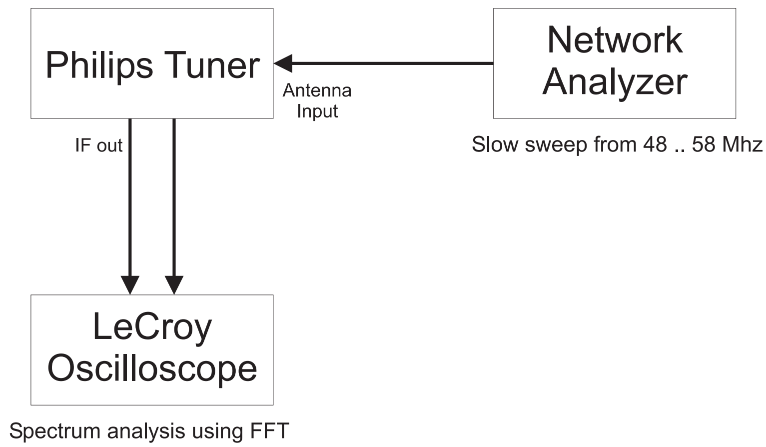

Figure 5.6: SAW Filter frequency response test setup

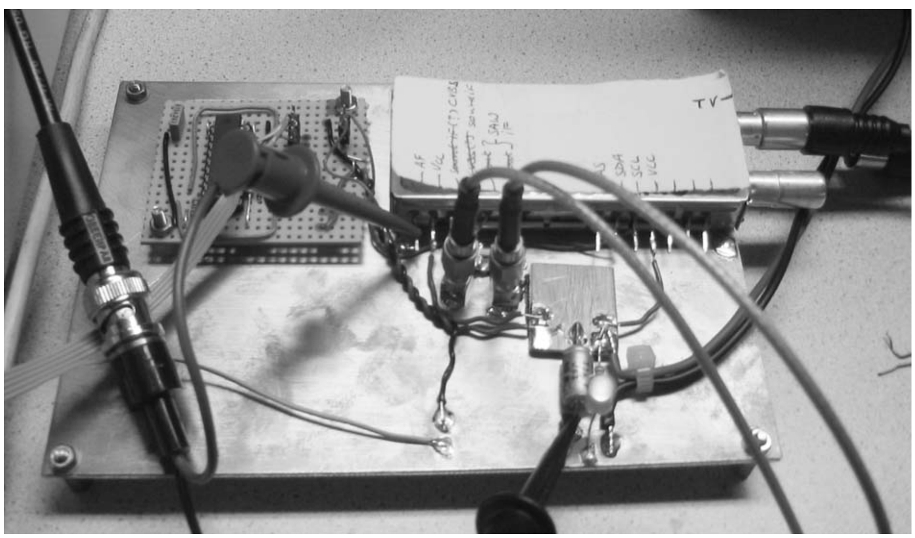

Figure 5.7: The TV Tuner setup during IF SAW filter response test

frequency range, where the filter response is approximately flat, runs from 35 to 39 MHz.

In addition to the tuner's IF spectrum, its power characteristic was also measured. This was done using the setup as shown in Figure 5.6 with the exception that the network analyzer produced a single carrier at 50 Mhz to drive the tuner. The LeCroy oscilloscope was used to measure the output power at the IF frequency by using its spectrum analysis functions. The

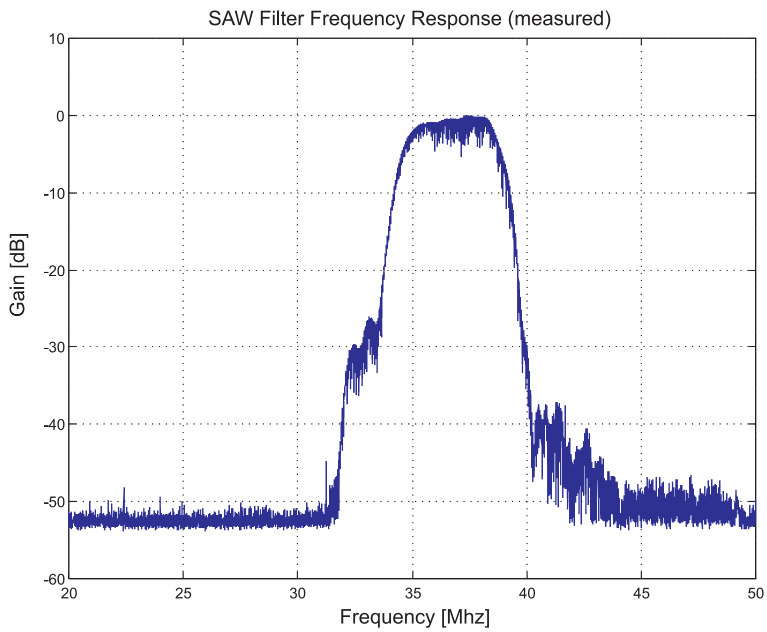

Figure 5.8: SAW filter transfer characteristic

results of these measurements are shown in Figure 5.9.

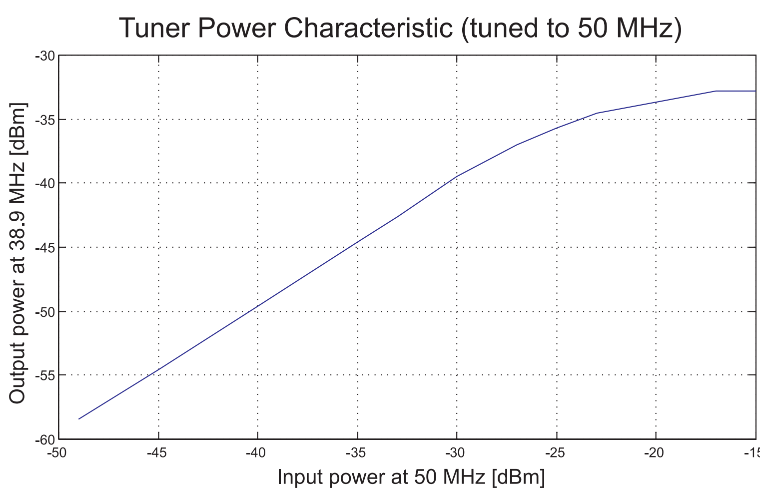

Figure 5.9: Tuner power characteristic

The output power from the tuner is weak and additional amplification is required to drive the DSP development system's A/D converter. And, according to the datasheet of the SAW filter [12], the output impedance is a high 2 kΩ and of the differential kind, see Figure 5.10.

EPCOS

SAW Components	K 2955 M
IF Filter for Intercarrier Applications	38,90 MHz

Data Sheet

Characteristics

Reference temperature:	T_A = 25 °C
Terminating source impedance:	Z_S = 50 Ω
Terminating load impedance:	Z_L = 2 kΩ || 3 pF

Figure 5.10: The EPCOS K2955M SAW filter datasheet extract

Now that the IF filter's frequency characteristic is known, the following stages can be designed. These stages are a mixer with impedance matching, a post-mixer lowpass filter and a post-mixer amplifier.

5.2.6 Down mixing the tuner IF

A further down-mixing stage is needed to bring the IF signal within correct frequency band of the A/D converter. The converter has a sampling rate of 60 MHz which allows under-sampling of the IF signal. However, this option was discarded because the SA612 mixer chip has the correct input impedance and a differential input which makes it an ideal component for directly interfacing the SAW IF output. In addition, the mixer is of the active type giving a conversion gain of around 16 dB.

By driving the mixer with an external 30 MHz local oscillator signal, the useful IF frequency range is translated to 5 to 9 MHz; ideal for the A/D converter and DSP system.

A further advantage of the SA612 mixer chip is that it allows a single-ended output. Unfortunately, the output of the mixer has a high output impedance, like the SAW filter. This poses a problem for the DSP system because the A/D converter has a low impedance of 50 Ω. Therefore, extra circuitry is needed to provide this low output impedance.

The low output impedance was achieved in two stages. The first stage is an emitter follower which has a high input impedance and a low output impedance making it ideal for this task. The voltage gain of an emitter fol-

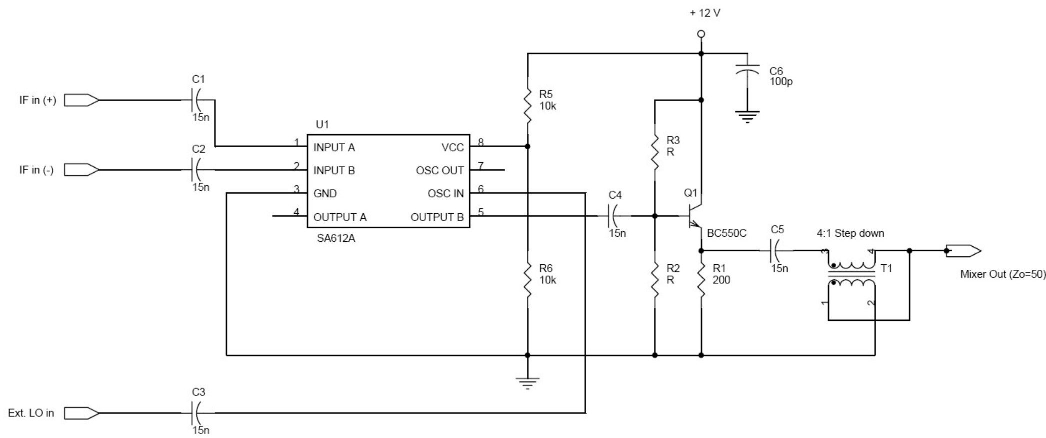

Figure 5.11: Schematic diagram of the post-tuner mixer

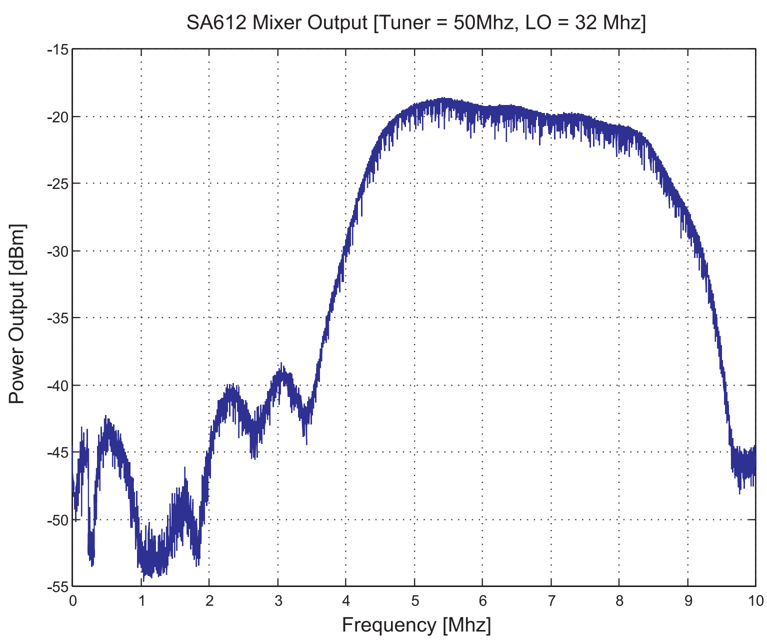

Figure 5.12: SA612A mixer output power (tuner input = -36dBm)

lower is smaller than 1 so it will introduce some loss. The output impedance of the emitter follower is approximately 200 Ω.

The second stage is a 4:1 step down transmission line transformer which converts the 200 Ω input signal into a 50 Ω output.

5.2.7 Post-mixer filter

Every mixer produces unwanted harmonics and other responses. These responses must be filtered out before they enter the following stages and especially in cases where the output of the mixer is sampled by an A/D converter. These harmonics may alias into the passband of the A/D converter causing spurious responses in the digital domain of the receiver.

The DSP system's A/D converter has an upper frequency limit of 30 MHz caused by the 60 MHz sampling rate. The dynamic range of the converter is about 72 dB owing to the 12 bits of resolution. In effect, the required filter attenuation at 30 MHz and upward is greater than 72 dB to avoid aliasing.

It is advantageous to chose the cutoff frequency of the post mixer filter as low as possible because this allows the filter order to be kept small. This reduces the number of components of the filter and its insertion loss. The lowest cutoff frequency of the filter is 6 MHz because the useful IF output from the mixer starts at 5 MHz and the minimum channel bandwidth is 1 MHz.

A Butterworth type filter was chosen because it has a maximally flat passband characteristic making it ideal for power measurement. A 5^{th}-order Butterworth with a cutoff frequency of 6 MHz has 70 dB of attenuation at 30 MHz. This was considered acceptable and it requires only two inductors. A schematic diagram of the filter is shown in Figure 5.13 and its frequency response in Figure 5.14.

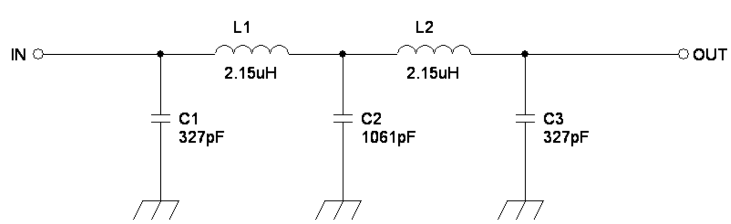

Figure 5.13: 6 MHz 5^{th}-order Butterworth lowpass filter

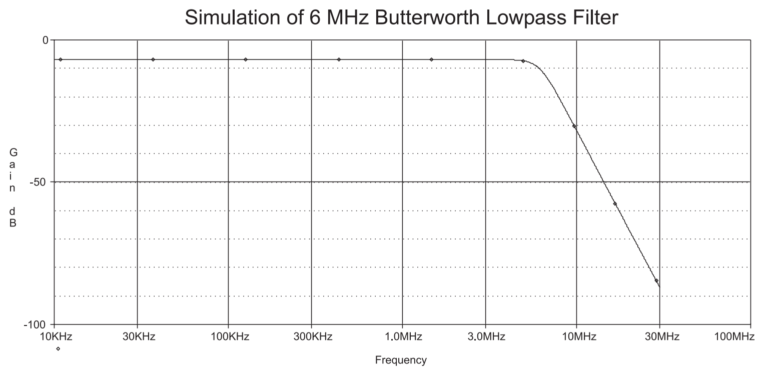

Figure 5.14: Post-mixer filter transfer characteristic

5.2.8 Post-mixer amplifier

The output of the post-mixer filter is still to weak to drive the input of the A/D converter. The converter needs a signal of 0.5 V peak into 50 Ω to exploit the full 12 bit range. This is around 4 dBm of drive power.

The output power of the post-mixer filter is about -20 dBm when the tuner is at its 1dB compression point. A post-mixer amplifier is needed to boost the signal by 24 dBm to fully exploit the range of the A/D converter.

Making good and stable amplifiers is not an easy task. This fact is corroborated by the well-known saying: "If you want to make an oscillator, try to build a good amplifier".

A quick and easy solution to the amplification problem was to build an amplifier based on a MMIC building block. These devices are small integrated circuits with bandwidths that range from 1 to several gigahertz. Most offer 10 to 20 dB of gain depending on the model.

A MAR6 MMIC [1] was selected as a suitable device because its amplification factor is 20 dB below 900 Mhz and it is unconditionally stable. The schematic diagram of the amplifier is shown in Figure 5.15.

The post-mixer amplifier allows the A/D converter to be driven over most of its linear range while still maintaining some headroom to avoid saturation even when the tuner is overloaded.

5.2.9 Receiver hardware overview

The diagram in Figure 5.16 shows how all the modules discussed earlier fit together. The heart of the receiver is the Tornado E67 DSP development

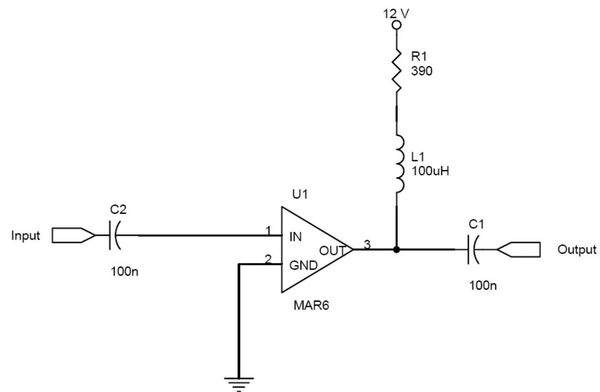

Figure 5.15: Schematic of the post-mixer amplifier

system. The Tornado E67 system controls the tuner through its general purpose I/O port and the I2C interface. The I2C bus is mainly used to set the receive frequency of the tuner.

The tuner's IF signal is connected to the mixer where it is down mixed to a lower frequency. This down-mixed signal is transformed into a 50 Ω compatible signal by an emitter follower stage and a broadband transformer. Finally, it is filtered by the 5^{th}-order Butterworth filter, to remove the unwanted mixing products, and is amplified to drive the input of the E67 system's A/D converter.

The A/D converter is connected to a HSP50214B digital programmable down-converter (PDC) from Intersil [14], which is part of the development system. The PDC extracts the 1 MHz channel from the IF signal and sends this to the DSP for further processing. So, the effective IF bandwidth of the receiver is 1 MHz.

The PC connected to the Tornado E67 sends measurement and tuning commands to the Tornado E67 over a serial RS232 interface. In this way, the Tornado E67 software is kept very simple as it only needs to measure IF power and drive the I2C bus to the tuner. The resource discovery algorithms are not part of the Tornado E67 program and reside on the PC. As most people have more experience in programming for the PC, it allows them to try out different resource discovery algorithms more easily.

The final receiver front-end for driving the Tornado E67 system is shown in Figure 5.17.

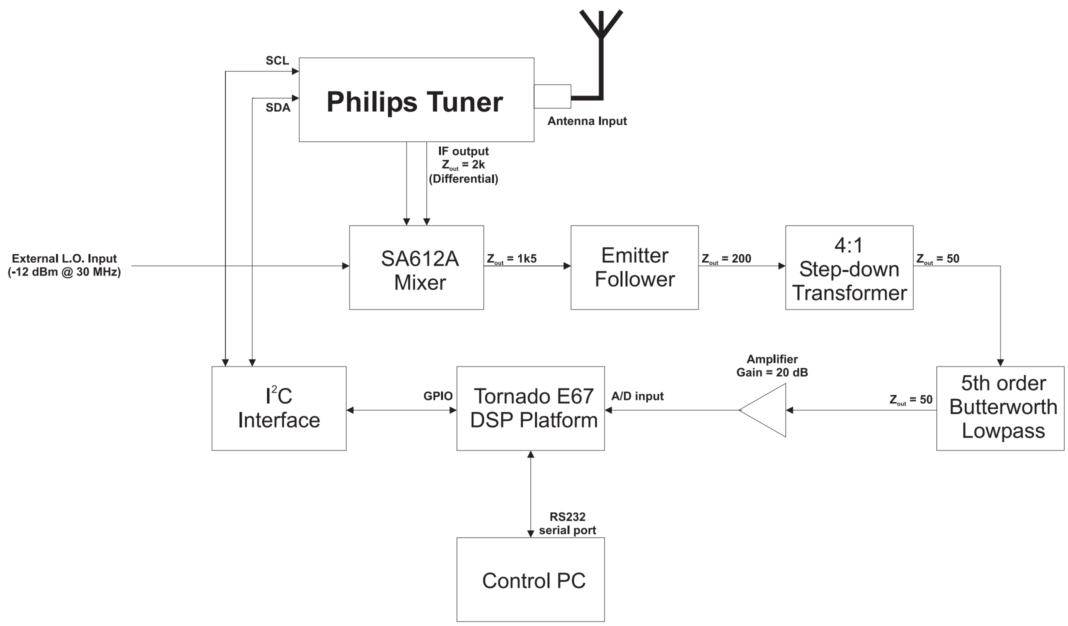

Figure 5.16: A block diagram of the complete receiver

5.3 Receiver software

The receiver software consists of two parts, the Tornado E67 software and the PC software.

The main purpose of the Tornado E67 software is capturing the IF signal and calculating its power. Other tasks of the Tornado E67 are tuning the FM1246 tuner module and various hardware initialization tasks.

The power measurement method use by the Tornado E67 is the coarse method outlined in 4.9.1. An entire 250 kHz wide IF chunk[1] is sampled into the Tornado E67. The Tornado E67 then calculates the RMS voltage and converts this into dB.

The PC software does most of the work as it sends measurement and tuning requests to the Tornado E67 system, calculates statistics of every channel, implements the resource discovery algorithm, and provides a graphical user interface. As screenshot of the user interface is shown in Figure 5.18.

The main part of the PC software is the resource discovery algorithm.

[1]Due to limitations of the Tornado E67 development system, it is impossible to sample a 1 MHz chunk.

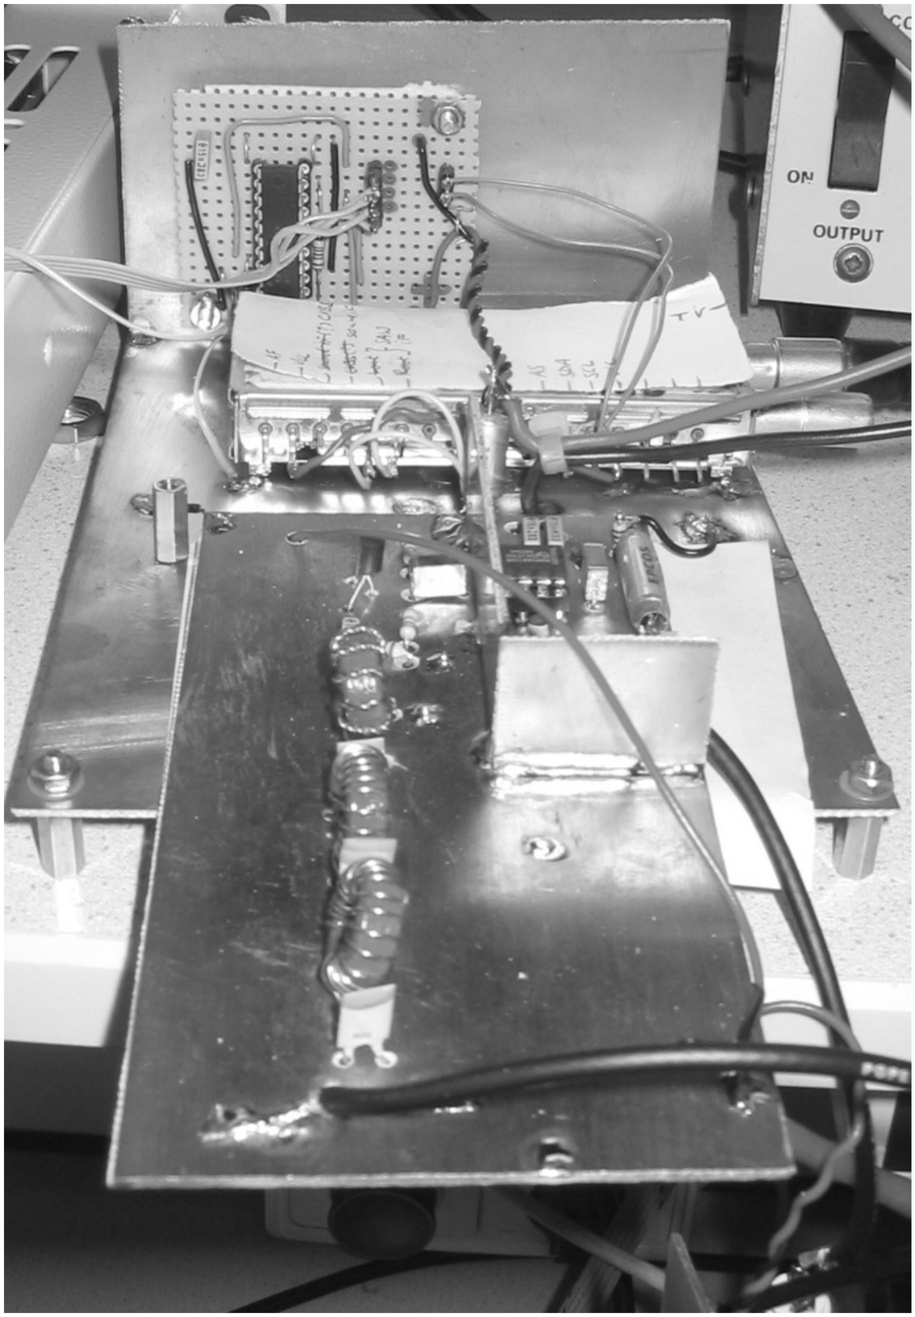

Figure 5.17: A picture of the final version of the receiver

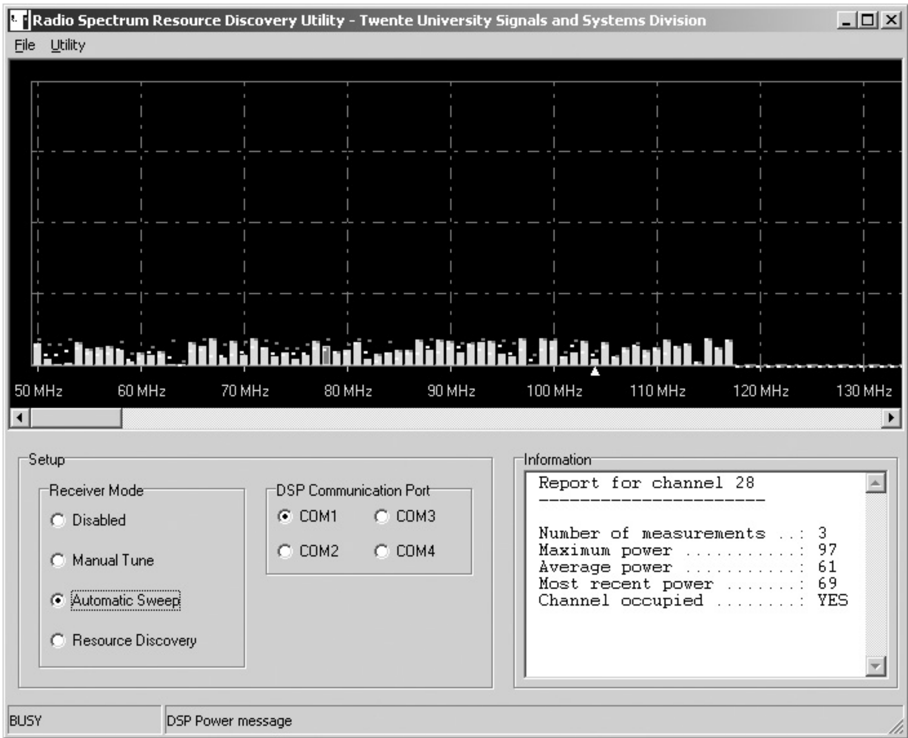

Figure 5.18: The graphical user interface of the resource discovery software

This algorithm sends measurement- and tuning-requests to the Tornado E67 through the Tornado E67 DSP command I/O subsystem. These requests are sent over the PC's RS232 serial communications port to the Tornado E67 using Microsoft Windows' File API [5].

The File API implements a blocking type interface and call to the API will not return until either a time-out occurs or data is sent/received. To avoid stalling the user interface, the resource discovery algorithm and the Tornado E67 command I/O subsystem execute in their own thread.

When the resource discovery algorithm receives a measurement result, this result is processed internally by the algorithm but it is also sent to the measurement data storage subsystem which calculates some channel statistics. The user is able to view the channel statistics by selecting the channel in the spectrum information display.

The spectrum information display is part of the graphical user interface and can clearly be seen in Figure 5.18. The most recent power result is shown for each channel by a green vertical bar. Each bar is accompanied by a red and yellow dot. The red dot shows the maximum power value in the last 32 measurements and the yellow dot shows the average power of the last 32 measurements.

When the resource discovery algorithm finds any empty channel, it notifies the spectrum display. Unoccupied channels are shown by painting the interior of each green vertical bar black making it look hollow. Occupied channels are shown by a filled green vertical bar.

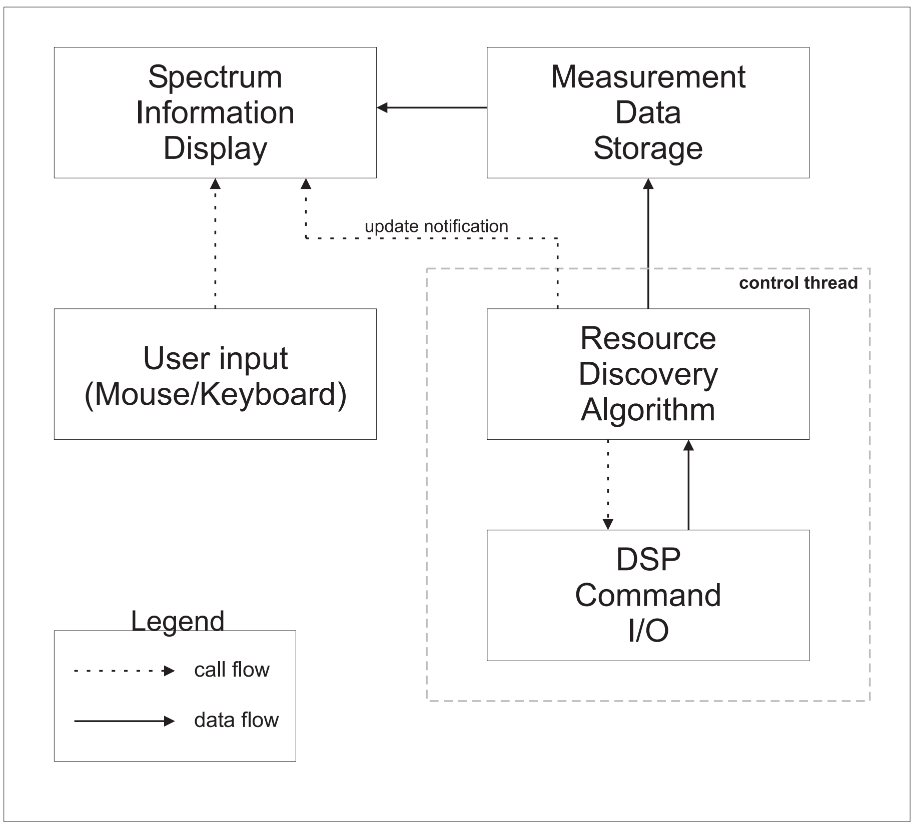

Figure 5.19: A block diagram of the PC software design

5.3.1 Duration of single channel power measurement

A power measurement of a single 1 MHz-wide channel takes a certain amount of time. First the tuner is tuned to the channel of interest. Then system must wait for the PLL of the tuner to settle. According to the tuner datasheet [22], this takes up to 150 ms. Then the Tornado E67 must sample the IF and calculate the amount of power in the channel. The data is transferred back to the PC using the Tornado E67's RS232 interface.

To measure the total time it takes to perform a power measurement on a single channel including retuning the tuner module, a full-range sweep was made across all the channels while performing power measurements. A full-range sweep takes 2 minutes and 20 seconds and sweeps over 800 channels. The total time per channel is then 175 ms.

Of this total time, only 25 ms is spent sending commands, sampling the IF signal and calculating the power. The remaining 150 ms, is spent waiting for the PLL in the tuner to settle.

5.4 Summary

In this chapter a low-cost resource discovery receiver hardware was described and some IF filter measurement results were given. This receiver is a dual conversion super-heterodyne type instead of the proposed direct-conversion receiver because it is based on a Philips TV tuner.

The software driving the receiver was described and some screenshots were shown. The resource discovery algorithm described in Section 4.10 is part of this software.

Finally, the total power measurement time was measured to be 175 ms per channel. This includes tuning the PLL in the tuner module.

The next chapter describes spectrum power measurements take with the low-cost resource discovery receiver. It also gives the resource discovery algorithm test run results.

6

Measurements using the Low-cost Resource Discovery Receiver

6.1 Introduction

The low-cost resource discovery receiver presented in Chapter 5 has two purposes. The first is to take power measurements and store them on disk for analysis. These measurements could be used to develop the usage pattern models described in Section 4.3.

The second purpose of the low-cost resource discovery receiver is to demonstrate the radio resource discovery algorithm described in Section 4.10.

The power measurements were made over the full tuning range of the low-cost resource receiver[1]. The results of the measurements are discussed in the next section.

Ideally, the low-cost resource discovery receiver should be verified for correctness. This requires all the signals at the antenna to be known. The RF signal present at the antenna input must therefore be made using test equipment. At the time when the receiver was built, our facility did not have such wide-band equipment. Instead, the UHF utility and TV bands are used for partial verification of the receiver.

[1] The full range is from 52 MHz to 852 MHz.

6.2 Spectrum power measurements

First, the a full frequency range overview measurement is shown and discussed. In order to verify the frequency correctness of the power measurements, the UHF utility band[2] and part of the UHF TV band are examined more closely.

Full spectrum power measurements were written to disk every two minutes over a two-hour period. The average power and the standard deviation [13] was calculated for each channel. This is called a *steady-state* power measurement.

6.2.1 Full frequency range overview

The full range steady-state power measurements are plotted in Figure 6.1 together with the standard deviation. The power graph is shown at the top while the deviation is shown at the bottom of the figure.

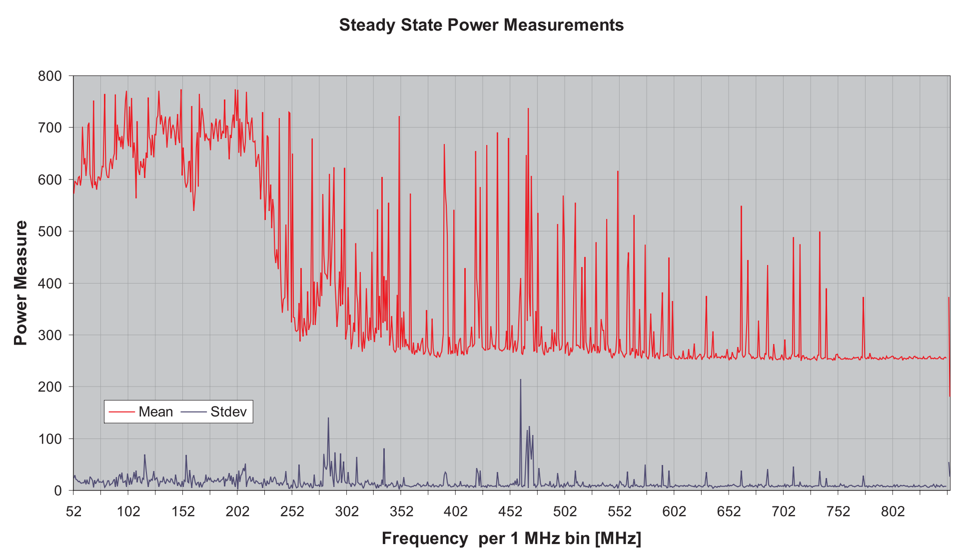

Figure 6.1: Full range steady-state power measurements

From the power measurements the man-made noise can be seen at the low end of the spectrum where it is approximately 600 units in power. It starts

[2]The UHF utility band contains police and ambulance voice communication frequencies.

to drop off up to around 250 MHz where it is no longer a major contributor. Above 250 MHz the receiver noise floor becomes the dominant noise at around 260 power units. As the frequency goes up, the activity in the spectrum becomes more and more sparse.

6.2.2 The UHF utility band

The UHF utility band runs from 460 to 470 MHz and contains police and ambulance voice transmissions. These transmissions are between a base station and police cars/ambulances. The base stations are located at a police station or a hospital. The maximum power output of these stations is 10 W effective radiated power.

An internet site that publishes police utility frequencies [29] in the vicinity lists the following frequencies:

- Enschede 466.67 MHz
- Enschede 466.91 MHz
- Boekelo 466.59 MHz
- Hengelo 467.09 MHz

Activity on these frequencies can be seen in Figure 6.2 where there is a peak in the 466 MHz channel. As an extra verification procedure, a Bearcat XLT-30 scanner/receiver was used to monitor these utility frequencies while observing the spectrum display of the resource discovery receiver. The resource discovery receiver was setup to constantly measure the power in the 466 MHz channel. The power measurement indicator could be seen moving up and down in unison with the audio from the Bearcat scanner.

6.2.3 The UHF TV band

In the Netherlands, the TV bands are not crowded because only three national TV channels are broadcast on these frequencies together with some low power local stations. The frequencies of these TV stations are published on the Internet by the Nozema [20].

For the Enschede region, the transmitter based at Markelo produces the best signal and has an effective radiated power figure of 300 kW. On the

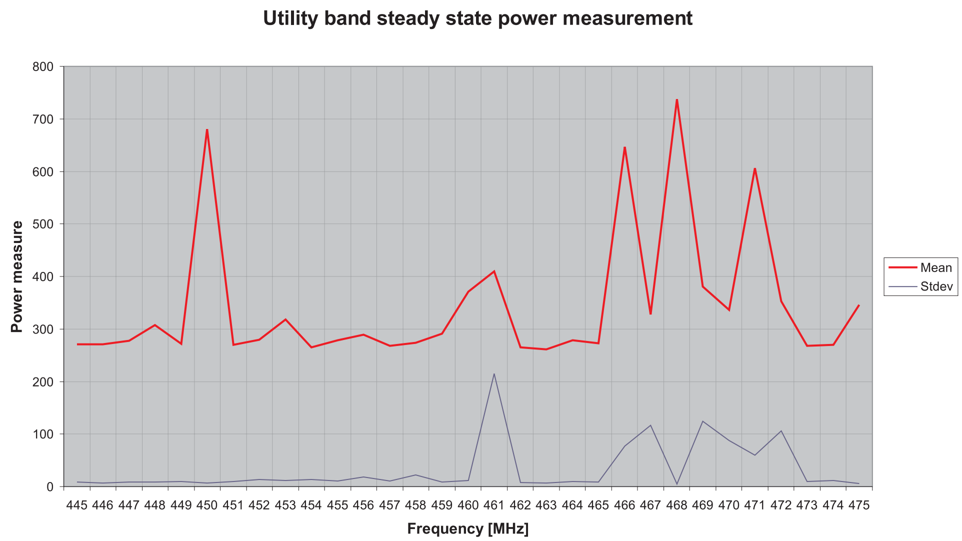

Figure 6.2: Steady-state power measurements of utility band

Nozema website, the transmitter data for Markelo [19] shows two frequencies in the UHF TV band, one for '*Nederland 2*' and the other for '*Nederland 3*'. These frequencies are:

- 711.25 MHz (Nederland 3)
- 735.25 MHz (Nederland 2)

In Figure 6.3 the steady-state power measurements of part of the UHF TV band are shown. In the graph the two TV stations are marked by an ellipse. Each ellipse contains two peaks. The left peak in the ellipse is the picture carrier while the right peak is the audio carrier.

6.3 The resource discovery algorithm test run

The software for the low-cost resource discovery receiver contains an implementation of the resource algorithm presented in Section 4.10. This algorithm was used to measure the number of available 1 MHz channels in the full tuning range of the low-cost receiver.

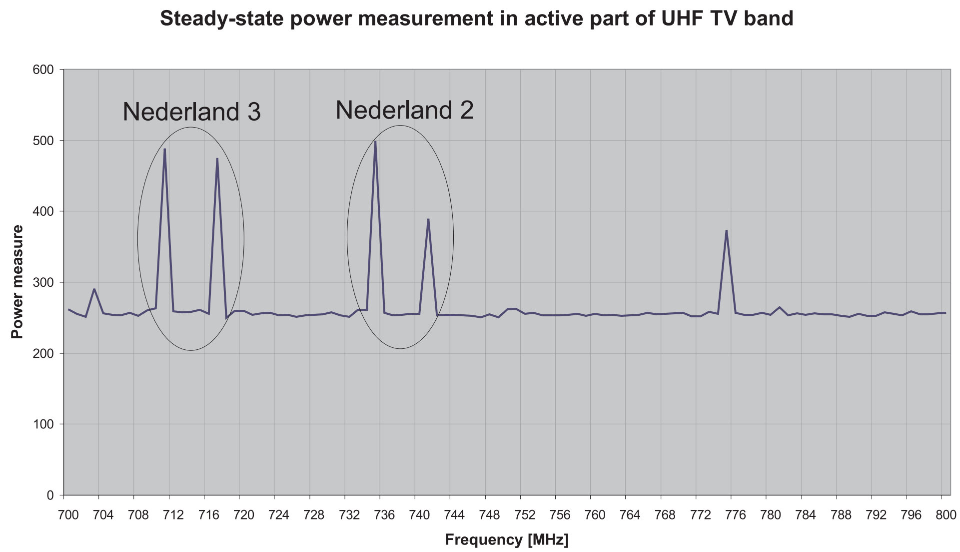

Figure 6.3: Steady-state power measurements of the UHF TV band

The number of available channels *detected* was saved to disk every 30 seconds. The data is plotted in Figure 6.4. It shows how fast the algorithm identifies available channels.

At time index 87 the resource discovery algorithm reaches a point where it has almost detected all available channels. At time index 97 it reaches the maximum number of available channels, or 544 channels. The first available channel is detected at time index 2.

The time it takes the resource discovery algorithm to detect all available channels is equal to $\frac{97-1}{2}$ minutes which equals 48 minutes. The first available channel is detected after just 30 seconds.

Most of the power in the spectrum is found at the low end of the frequency range, see Figure 6.1. This suggests that most of the available channels will be located at the high end of the spectrum. The resource discovery algorithm confirms this.

The resource discovery algorithm's spectrum display is shown in Figure 6.5 and Figure 6.6. The first figure shows the a frequency range from 785 to 851 MHz, the second a range from 95 to 126 MHz. All the channels in the high frequency figure are available (as indicated by a hollow bar) while the channels in the low frequency figure are all occupied (as indicated by a filled

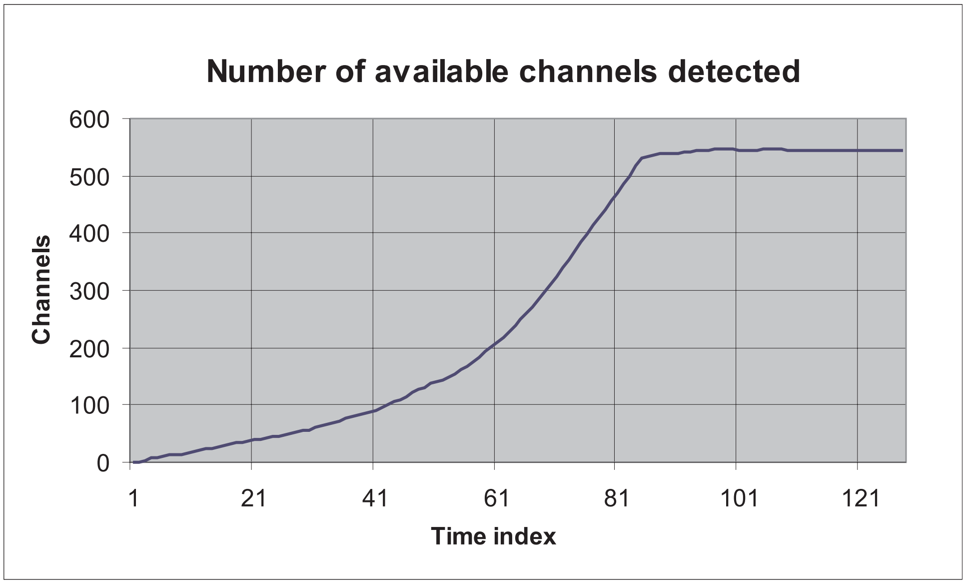

Figure 6.4: Number of detected available channels

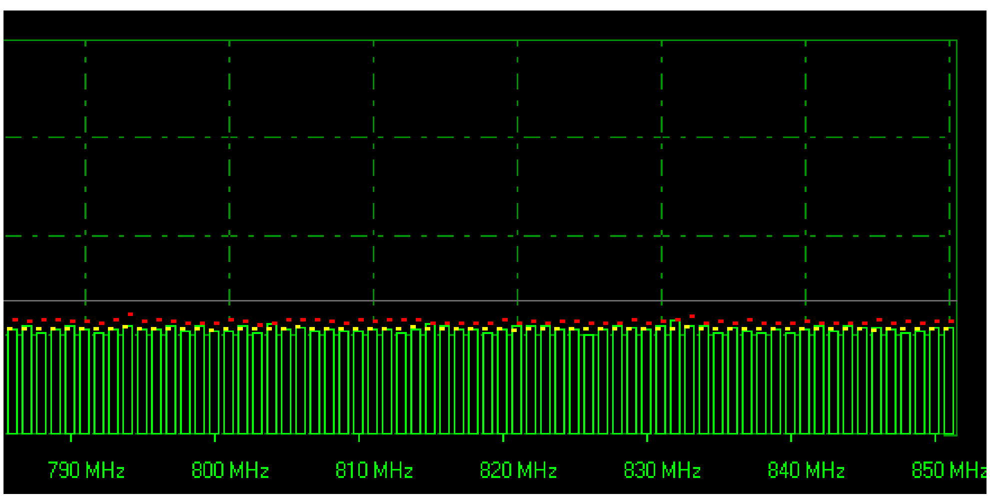

Figure 6.5: Activity at the high end of the spectrum

bar).

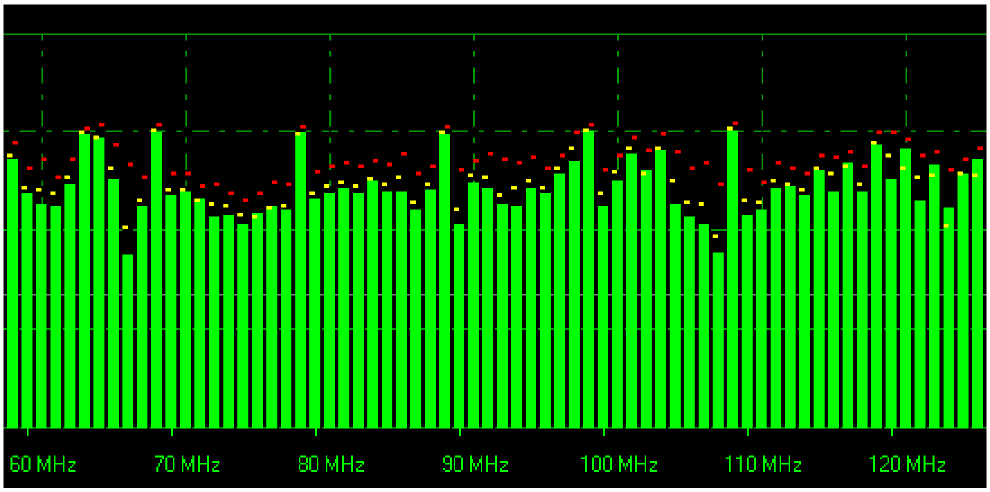

Figure 6.6: Activity at the low end of the spectrum

6.4 Summary

In this chapter the results of radio spectrum power measurements were presented. These results confirm that the low-cost resource discovery receiver is able to detect both high-power TV transmitter signals and low-power mobile radio signals. Also, the frequencies of the power measurement peaks coincide with the known frequencies of the TV transmitter and mobile radio channels. A full verification procedure was not performed.

This chapter also described the number of available channels detected by the resource discovery algorithm from Section 4.10. The total number of available 1 MHz-wide channels was found to be around 544 but no verification procedure was performed.

7

Conclusions and Recommendations

This chapter discusses the conclusions and concludes with recommendations.

7.1 Conclusions

This project had two goals, the first was to investigate methods by which activity in the radio spectrum could be determined and used to generate an activity map.

A method for determining activity in the radio spectrum based on power measurements was presented in Chapter 4 together with a resource discovery algorithm that will scan the spectrum and generate an activity map. However, no guarantee is made that the information contained in the activity map is not outdated.

The need for a primary user usage pattern model was stated and more research is needed to obtain those models. The model will require certain information about the history of the primary user's activity. As nothing is known about the frequency or time resolution requirements of these models, no assumptions can be made on the most efficient radio resource discovery node architecture. However, the most flexible architecture is the broadband separate radio resource discovery receiver architecture because it can monitor the entire transceiver frequency range continuously.

A second goal of the project was to build a low-cost radio resource discovery receiver that can be used as a demonstrator and testbed for further developments. The receiver was built and can measure the power present in 1 MHz-wide channels.

The low-cost receiver was used to test the resource algorithm from Section 4.10 and the results of this test were presented in Section 6.3. The resource discovery algorithm finds around 544 available 1 MHz-wide channels.

The low-cost receiver can be used to obtain power measurement data for use in later research projects such as primary user modelling. It can also be used to test newly developed radio resource discovery algorithms designed for the narrowband separate resource receiver architecture.

7.2 Recommendations

7.2.1 Further research

It is clear that there are still many questions to answer. One of the most important issues to be resolved is that of the primary user models. Research should be carried out in this field, particularly focussing on whether these models can achieve the necessary prediction accuracy to guarantee low interference probabilities. The frequency and time resolution requirements may also become clear from this research.

A second area that needs more research is how nodes are synchronized with respect to which channels are used and when. This becomes an important issue when the primary user causes the resource discovery system to change the activity map very often, say every second.

In this thesis each network node performs its own full frequency range radio resource discovery. But a node can communicate with its neighbor nodes. Therefore, a distributed resource discovery algorithm could be used. Research is needed to determine the feasibility of such an algorithm.

7.2.2 Use the fine method of power measurement

The fine method for power measurement not only has increased frequency resolution, it can also detect very narrow bandwidth signals more easily due to a lower noise floor. The lower noise floor allows the threshold level of the resource discovery algorithm to be lowered.

7.2.3 Sensitivity measurements of the low-cost receiver

No sensitivity measurements of the low-cost receiver were performed. It is recommended to perform these measurements over the full frequency range of the receiver to determine if the receiver is more sensitive to certain frequencies than others.

The radio resource discovery algorithm assumes that the sensitivity of the receiver is constant irrespective of frequency. If this assumption turns out to be incorrect, a frequency dependent threshold is needed for the resource discovery algorithm to compensate for this.

7.2.4 Faster tuning for the low-cost receiver

Should it prove necessary, the tuning speed of the TV Tuner could be increased by an external oscillator. The oscillator should be based on a direct digital synthesis (DDS) synthesizer. This type of synthesizer has a very fast tuning capability and COTS solutions are available, see the Analog Devices website [2] for further information.

7.2.5 Less tuning for the low-cost receiver

The low-cost radio resource receiver IF bandwidth is only 1 MHz wide. This bandwidth holds only a single channel. It is recommended to increase the IF bandwidth to 2 MHz to allow sampling two channels simultaneously. This will reduce the number of tuning operations from 800 to 400 for a full frequency range power measurement. Unfortunately, the Tornado E67 is unable to cope with more than 1 MHz bandwidth so a different platform is needed.

7.2.6 Make programming the low-cost receiver easier

Programming the low-cost receiver involves two sets of software. The first set is for the Tornado E67, the second is for the PC. Most people have experience in programming for the PC but not for the Tornado E67. The programming environment for the Tornado E67 will be unfamiliar to them and requires knowledge of embedded system programming.

To make programming the resource discovery receiver easier, the Tornado E67 could be discarded in favor of processing the IF signal directly with the PC which is already part of the receiver. This requires the PC to have an

A/D converter and 4 I/O pins to control the I2C bus interface to the tuner module of the receiver.

A

Formula

A.1 Carrier-to-Noise ratio

The following formula gives the carrier-to-noise ratio at the receiver[16]:

$$\frac{C}{N} = 20 \cdot \log_{10}\left(\frac{\lambda}{4\pi R}\right) + P_t + G_t + G_r - NF - L_t - kTB \tag{A.1}$$

,where

- C is the power of the carrier.
- N is the noise power at the location.
- λ is the wavelength of the RF carrier.
- R is the distance away from the transmitter.
- P_t is the transmitted power.
- G_t is the gain of the transmitting antenna.
- G_r is the gain of the receiving antenna.
- NF is the noise figure of the receiver.
- L_t is the total of any other losses.
- k is Boltzmann's constant

- T is the equivalent temperature of the receiver
- B is the bandwidth occupied by the signal

A.2 Channel noise power

Given a channel of bandwidth B, the channel will exhibit thermal noise which power can be calculated:

$$N = kTB \tag{A.2}$$

,where

- N = Noise power (W)
- k = Boltzmann's constant ($1.38 \cdot 10^{-23}$ J/K)
- T = System temperature (Usually 290K)
- B = channel bandwidth (Hz)

B

Two Discone Antenna Designs

Discone antennas have a receive bandwidth ratio of 10:1 and a transmit ratio of 3:1 which makes them excellent for broadband communication [18].

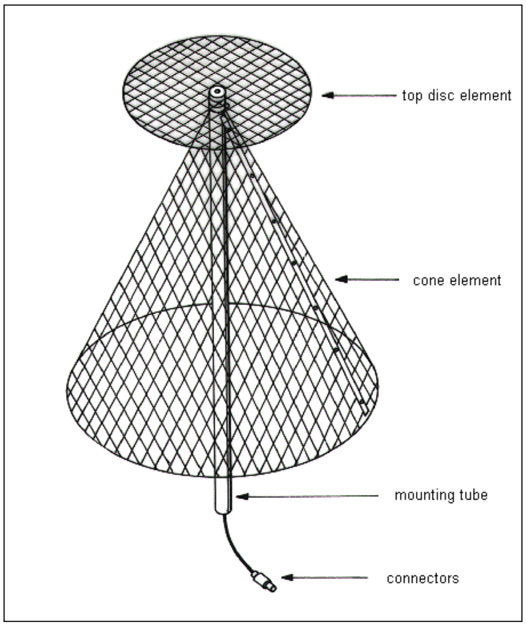

Figure B.1: Discone antenna

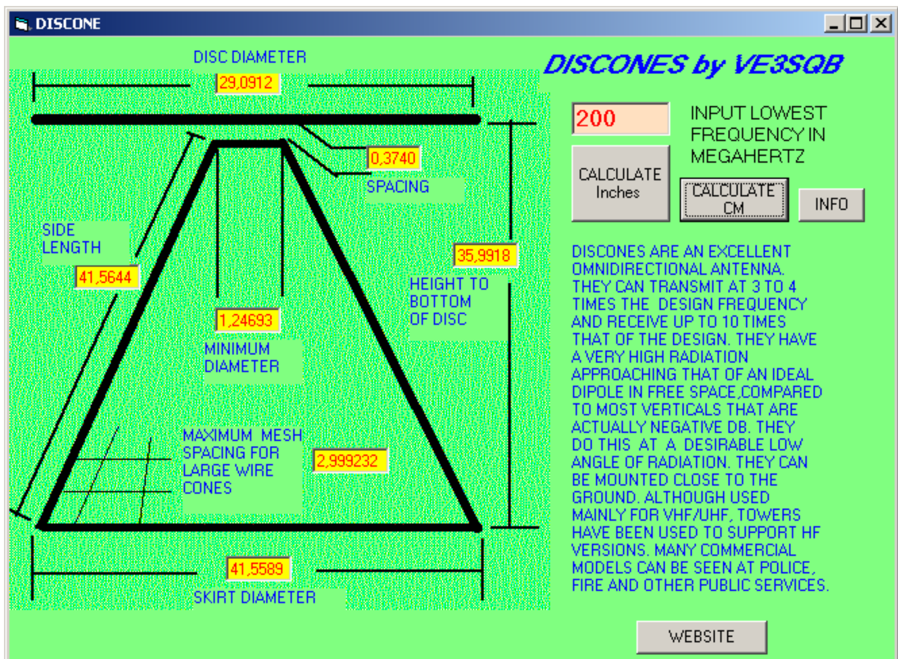

Figure B.2: Discone antenna design for 200 .. 600 MHz, dimensions are in [cm].

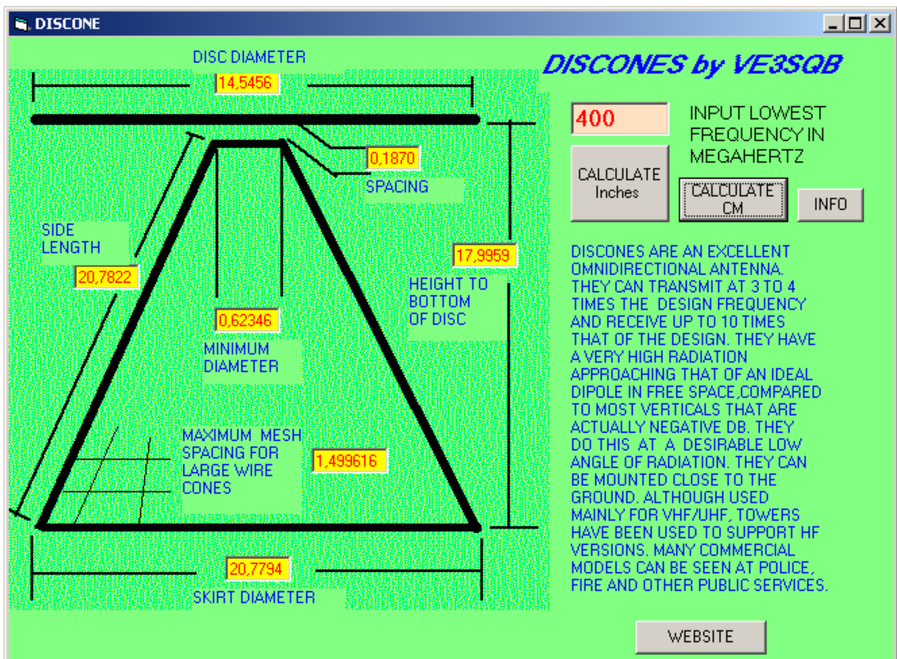

Figure B.3: Discone antenna design for 400 .. 1200 MHz, dimensions are in [cm].

C

Notes

Here are some notes that may provide ideas for further research.

C.1 Note 1

Is a carrier detection based medium access protocol suitable to 'ride' the gaps in a primary user's TDMA network? This would simplify the resource discovery subsystem's task considerably.

C.2 Note 2

Which model or set of models is accurate enough to guarantee an interference probability of less than 0.1%? Are the parameters of the models obtained or measured easily? How expensive are these models in terms of computational power?

Bibliography

[1] Agilent. MSA-0686 (MAR6) MMIC Amplifier datasheet. http://cp.literature.agilent.com/litweb/pdf/5965-9588E.pdf.

[2] Analog Devices. Analog Devices Website. http://www.analog.com/.

[3] Bengt Lindhoff and Peter Malm. BER Performance Analysis of a Direct Conversion Receiver. *IEEE Transactions on Communications*, 50(5), May 2002.

[4] Commissie Oosting. *De vuurwerkramp: Eindrapport*. Ministerie van binnenlandse zaken, Februari 2001. http://www.minbzk.nl/contents/pages/3166/eindrapport_oosting_2-01.pdf.

[5] Microsoft Corporation. WIN32 FILE API. http://msdn.microsoft.com/library/en-us/fileio/base/createfile.asp.

[6] Robert C. Dixon. *Radio Receiver Design*. Marcel Dekker Inc, 1998.

[7] Jarkko Jussila et al. A 22-mA 3.0-dB NF Direct Conversion Receiver for 3G WCDMA. *IEEE Journal of Solid-State Circuits*, 36(12), December 2001.

[8] Liwei Sheng et al. A wide-bandwidth Si/SiGe HBT direct conversion sub-harmonic mixer/downconverter. *IEEE Journal of Solid-State Circuits*, 35(9), September 2000.

[9] Sergio Sanielevici et al. A 900-MHz transceiver chipset for two-way paging applications. *IEEE Journal of Solid-State Circuits*, 33(12), December 1998.

[10] F. Brouwer, M. de Graaf, H. Nikookar and F. Hoeksema. Adaptive Ad-hoc free Band Wireless Communications - Project plan. Unpublished, May 2004.

[11] Marvin E. Frerking. *Digital Signal Processing in Communication Systems*. Kluwer Academic, 7th edition edition, 1994.

[12] EPCOS GmbH. EPCOS K2955M SAW Filter datasheet. http://www.epcos.com/inf/40/ds/mm/K2955M.pdf.

[13] Gouri K. Bhattacharyya and Richard A. Johnson. *Statistical Concepts and Methods*. Wiley and Sons, 1977.

[14] Intersil. HSP50214B Programmable Downconverter Datasheet. http://www.intersil.com/products/deviceinfo.asp?pn=HSP50214B.

[15] John G. Proakis. *Digital Communications*. McGraw Hill, 4th edition, 2000.

[16] Joseph Mitola. *Software Radio Architecture*. John Wiley Inc, 2000.

[17] Eric P. Lawrey. *Adaptive Techniques for Multiuser OFDM*. PhD thesis, James Cook University, 2001.

[18] American Radio Relay League. *ARRL Antenna Handbook 2004*. ARRL, 2004. http://www.arrl.org.

[19] Nozema. Markelo radio and TV frequencies. http://www.nozema.nl/radiotv/Opstelpunten/Markelo.html.

[20] Nozema. Nozema Website. http://www.nozema.nl/.

[21] Charles E. Perkins. *Ad-Hoc Networking*. Addison-Wesley, 2001.

[22] Philips RF Solutions. FM1246 Tuner Datasheet. http://wwww.rfsolutions.philips.com.

[23] Philips Semiconductors. The I2C-Bus Specification, January 2000. http://www.semiconductors.philips.com/acrobat/literature/9398/39340011.pdf.

[24] Richard A. Roberts and Clifford T. Mullis. *Digital Signal Processing*. Addison Wesley, 1987.

[25] Ulrich Rohde. *Communications Receivers*. McGraw-Hill, 2 edition, 1997.

[26] R. Schiphorst. *Software-Defined Radio for Wireless Local-Area Networks*. PhD thesis, Twente University, 2004. Unpublished.

[27] Albert A. Smith. *Radio Frequency Principles and Applications*. IEEE Press, 1998.

[28] Friedrich K. Jondral Timo A. Weiss. Spectrum Pooling: An Innovative Strategy for the Enhancement of Spectrum Efficiency. *IEEE Radio Communications*, March 2004.

[29] Unknown. Scanner frequencies. http://www.aardbaan.com/freqdb/nl_select.php?page=3&selectie=plaats&invoer=enschede.

[30] Jan-Jaap van de Beek et al. A Time and Frequency Synchronization Scheme for Multiuser OFDM. *IEEE Journal on Selected Areas in Communications*, 17(11), November 1999.

[31] Peter Vizmuller. *RF Design Guide*. Artech House, 1995.

[32] Won Namgoong and Teresa H. Meng. Direct-Conversion RF Receiver Design. *IEEE Transactions on Communications*, 49(3), March 2001.

[33] Samuel C. Yang. *CDMA RF System Engineering*. Artech House, 1st edition, 1998.