OFCOM

Study of current and future receiver performance

Final Report

Lewis Davies

Paul Winter

11th January 2010

ACRONYMS

ACS	Adjacent channel selectivity
ACPL	Adjacent channel power leakage
ADC	Analogue to digital convertor
AGC	Automatic Gain Control
AM	Amplitude Modulation
BER	Bit Error rate
C	Capacitance
C/N	Carrier to noise ratio
CDMA	Code Division Multiple Access
CMOS	Complementary Metal Oxide semiconductor
CW	Carrier wave
DAB	Digital Audio Broadcasting
DAC	Digital to analogue convertor
DC	Direct current, 0Hz
DSP	Digital Signal Processing
DVB-T	Digital Video Broadcasting -Terrestrial
EMC	Electromagnetic compatibility
ENOB	Effective number of bits
FDD	Frequency Division Duplex
FEM	Front End Module
FFT	Fast Fourier transform
FM	Frequency Modulation
F_s	Sampling frequency
f_t	unity gain frequency of a transistor's short circuit current gain
Gm	transconductance
GSM	Global System for Mobile communications
IC	Integrated circuit
iDTV	Integrated Digital Television
IF	Intermediate frequency
IM2	Second order intermodulation products
IM3	Third order intermodulation products
IP	Intermodulation product
IRR	Image rejection ratio
ISM	Industrial Scientific and Medical
k	Boltzmann constant (~1.38 x 10^{-23} JK^{-1})
LC	Inductor, capacitor
LNA	Low Noise Amplifier
LO	Local Oscillator
LPF	Low Pass Filter
LTE	Long term evolution (of UMTS)
MIMO	Multiple input, multiple output
MOPLL	Mixer, oscillator, PLL
NF	Noise Figure
OFDM	Orthogonal Frequency Division Multiplexing
P_{1dB}	1dB compression point
PLL	Phase lock loop
PMSE	Programme making and special events
ppm	parts per million
Q	Quality factor
QPSK	Quadrature Phase Shift Keying
RC	Resistor capacitor
RF	Radio frequency
RFIC	Radio frequency Integrated circuit

RX	Receiver
SAW	Surface acoustic wave
SFN	Single frequency network
SINR	Signal to (interference plus noise) ratio
SNR	Signal to noise ratio
SOI	Silicon On insulator
SQNR	Signal to quantisation noise ratio
T	Temperature
TDD	Time division Duplexing
TX	Transmitter
UMTS	Universal Mobile Telecommunications System
VCO	Voltage controlled oscillator
VCXO	Voltage controlled crystal oscillator
WCDMA	Wideband Code Division Multiple Access

1 EXECUTIVE SUMMARY

1.1 Introduction

This study, commissioned by Ofcom, examines issues involved in improving the performance of radio receivers in consumer equipment such as TVs and cellular devices. The study focuses on the cost / performance trade-off to make a receiver less susceptible to interference from other frequency bands, based on technologies that can be envisaged over the next ten years. An indication of how this might change over the next twenty years is also given.

For optimum use of the radio spectrum, each frequency channel is ideally independent of other frequency channels. This independence can be gained through ensuring radio transmitters transmit all their energy in the specified frequency channel and receivers are immune to energy from other frequency channels, i.e. they have good selectivity. Without full independence, the use of each frequency channel depends, to some extent, on the use of neighbouring frequency channels. This makes introducing new services into the radio spectrum very difficult.

1.2 Receiver selectivity

Receiver selectivity is the receiver's ability to receive the "wanted" signal in a frequency channel without being affected by "interfering" signals present in adjacent and other channels. In many practical cases the receiver's adjacent channel performance is critical.

Whilst many different radio receiver architectures can be used, each with their own strengths and weaknesses, selectivity in all cases is largely determined by:

- Receiver channel filter performance. In modern digital receivers this filtering is divided into two portions, the analogue filtering prior to ADC conversion of the received signal and digital filtering following the ADC. The quality of digital filtering is dictated by the dynamic range of the ADC. This is the ability to be able to receive large and small signals simultaneously.

- Reciprocal mixing. When the received RF signal is mixed with a local oscillator to convert it to typically a much lower frequency for channel filtering and demodulation, noise is added by the local oscillator. This noise, which is typically proportional to the size of unwanted interferers, can swamp small wanted received signals.

- Receiver linearity. All analogue elements of a receiver such as amplifiers, mixers, and active filters have some nonlinearity, not least because they have a maximum signal they can process. These nonlinearities introduce distortion to both the wanted and any unwanted signals. This can lead to the creation of new interfering distortion products occurring at new frequencies. If these occur at critical frequencies within the receiver, they will affect the receiver's ability to receive a small wanted signal.

- Spurious responses. Whilst ideally the receiver converts just the wanted RF signal to typically a much lower frequency for channel filtering and demodulation, in practice the receiver will have a number of spurious responses. Unwanted signals at these frequencies could block the wanted signal. The nature of the spurious responses and how they are handled is architecture dependent. One of the best known is the superhet receiver's image response. A superhet receiver reduces its susceptibility to signals at the image frequency by using an image filter.

Although not directly a receiver performance issue, transmitters do not constrain all their transmit power to their allocated transmit frequency channel. Any energy transmitted on adjacent channels is known as adjacent channel power leakage, ACPL. The ACPL sets a limit on the selectivity that can be obtained with an ideal receiver.

1.3 Receiver architectures

Receivers can be implemented using a variety of architectures. Super heterodyne receivers, capable of very good performance but needing a large number of discrete elements such as image and IF filters, were used in most applications in the 20th century. Their adjacent channel selectivity performance is gained largely through the use of an analogue IF filter. Limited digital filtering is typically used after the ADC allowing an ADC to be used with reduced dynamic range. Good far rejection is obtained as a desirable side effect of needing an image filter to minimise the receiver's image response.

More integrated approaches such as zero IF and low IF allow almost the entire receiver to be fabricated in silicon. When manufactured in large volumes the unit cost is very low, and performance, after calibration, between one receiver and another is very consistent. These receivers do not use external IF filters and typically rely more heavily on digital filtering to obtain their selectivity than a superhet receiver does. The receiver's ADC can place a key limitation on the selectivity obtainable.

1.4 Current device cost versus selectivity

In order to improve receiver selectivity, the cost of the RF receiver is likely to increase, due to the increased development effort and larger die size required. Assuming the additional costs are passed through the supply chain from the semiconductor vendor, through the product manufacturers to the consumer, the retail price of the product will increase.

Figure 1-1: Typical consumer electronics supply chain

To estimate the influence additional receiver selectivity will have on RF receiver production cost in a quantifiable way the concept of improving the ADC's resolution has been used in the report. With appropriate receiver design this additional ADC effective number of bits (ENOB) or dynamic range can be used to obtain better selectivity through improved digital filtering. This approach allows adjacent channel selectivity to be improved at all but the highest signal levels, but has little effect on far-off selectivity. To improve far-off selectivity, techniques such as tracking filters need to be used.

One additional ADC bit of resolution, assuming the same sample rate, will allow the ADC to have up to 6dB additional dynamic range. By plotting the ADC's ENOB and silicon area for a variety of ADCs reported in the literature, all using a similar silicon process, see figure 1-2, the typical relationship between an ADC's selectivity and cost can be derived.

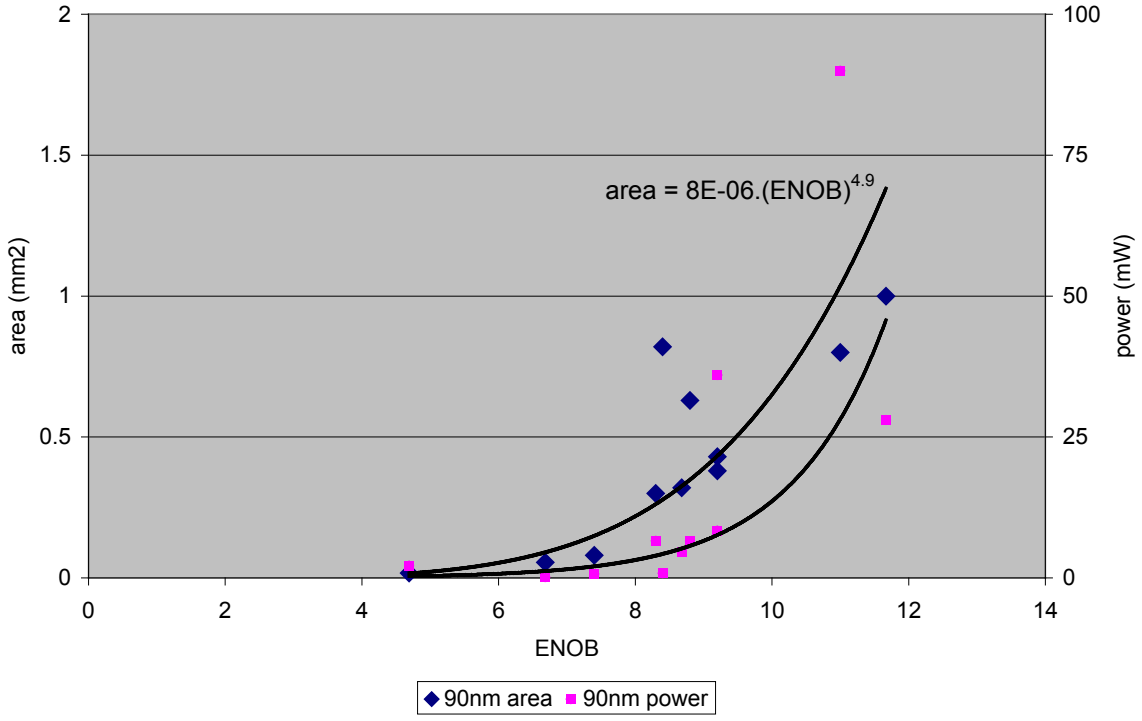

Figure 1-2: Silicon area and power consumption versus ENOB

By assuming that silicon area is proportional to the cost of the receiver and that an existing ADC accounts for typically 35% of the analogue/mixed signal area of the IC, understanding the relationship between silicon area and ADC performance allows the relationship between receiver selectivity and cost to be predicted as shown in figure 1-3. The graph is for a 9bit ADC, in a receiver costing a nominal $1, being improved to 12bits. The graph is limited to 12 bits as this is the limit of current ADC technology at this sampling rate.

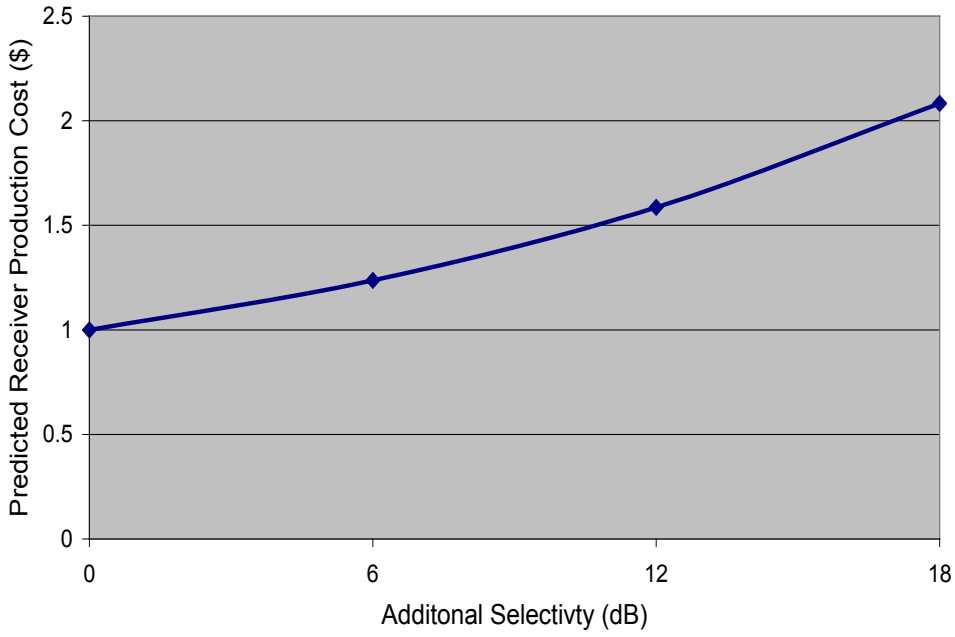

Figure 1-3: Additional selectivity effect on nominal receiver IC production cost

By assuming the receiver IC cost is passed through the supply chain the effect of additional selectivity on various consumer products can be predicted as shown in figure 1-4.

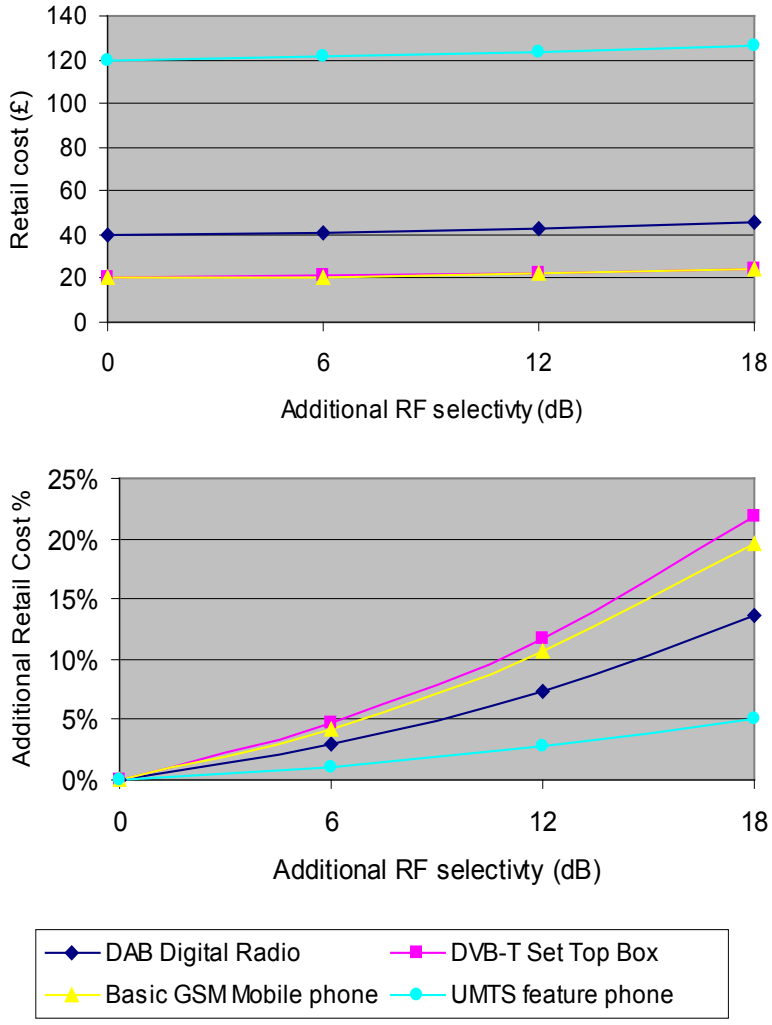

Figure 1-4: Product cost versus RF selectivity

Whilst increasing receiver selectivity is expected to lead to increased semiconductor silicon area, production cost, and research and development costs, it has been identified that semiconductor manufacturers may absorb these additional costs in some cases.

In systems where there is a clearly defined applicable standard which is seen as adequate by the industry, for example cellular, RF receiver selectivity performance beyond that required by the applicable standard the receiver is working to is generally not perceived as important by the product manufacturer or retailer. It is thought unlikely in this case that the receiver manufacturer would choose to improve receiver selectivity.

However in some systems the applicable standard is not seen as adequate by parts of the industry. One example of this is TV. Premium brands are likely to require tuners with additional selectivity in order to avoid the risk of a poor customer experience.

It is anticipated that the continual improvement of silicon processes and receiver design techniques will allow in time any additional costs to be reduced through enhanced design, reducing long term increase in semiconductor production costs below those shown in the figure above.

1.5 Anticipated Potential Future Receiver Performance

In anticipating the trends in future receiver performance the following can be identified:

- As the RF spectrum becomes more crowded and spectrum regulation tends to become more market-led, it can be anticipated that for interference to not limit wireless system performance, receiver selectivity will become more important.
- As receive bandwidths increase, receive sensitivity will reduce necessitating higher transmitter power to maintain the link budget. For a given receiver with constrained large signal handling, the selectivity of the receiver will degrade correspondingly.
- For mature standards, in the highest volume designs, the RF receiver is increasingly being merged with digital basebands.
- Potential for RFICs to cater for a reduced number of standards or standard families, but with some standards being used for a wider range of applications.
- Increased use of standardisation in wireless systems with proprietary applications migrating to standardised protocols such as Zigbee whenever possible.
- Increased semiconductor development costs needing to be offset by increased revenues through higher product volumes.
- New processes, such as SOI, which may replace bulk CMOS once the physical limitations of shrinking CMOS geometries are reached, allowing predictions such as Moore's law to possibly hold for longer into the future.
- A wide range of receiver design techniques which will allow at least the selectivity performance of RF receivers specified in current standards to be met at lower cost in more integrated receivers.

Development costs for the highest density, lowest production cost CMOS nodes are rising rapidly. Therefore the size of market needed to justify the investment in the silicon development will also rise. This suggests that there will be opportunities in medium size markets for slightly less integrated approaches, e.g. separate RF and baseband. These might use lower density CMOS processes (with reasonable f_t and supply voltage) allowing better large signal handling.

In some cases, if costs are acceptable, more specialist silicon processes such as SiGe could be considered. SiGe can operate with higher supply voltages than CMOS and therefore has better large signal handling. However its production cost is greater than CMOS and its transistor density is lower. This makes it acceptable for devices with a large proportion of RF circuitry, where passive components, which do not scale with the silicon process, dominate, but not acceptable for devices with a large proportion of digital circuitry.

1 INTRODUCTION

This study, commissioned by Ofcom, investigates current and future RF receiver performance focussing on the receiver selectivity for a range of receiver architectures used in typical consumer devices.

For optimum use of the radio spectrum, each frequency channel is ideally independent of other frequency channels. This independence can be gained through ensuring that radio transmitters transmit all their energy in the specified frequency channel and receivers are immune to energy from other frequency channels. Without full independence, the use of each frequency channel depends, to some extent, on the use of neighbouring frequency channels.

Transmitter emissions are managed through regulatory defined spectrum masks and EMC emissions limits. These masks and limits do not ensure that all the transmitter emissions are in the transmit channel but do ensure that they are constrained to a reasonable known level.

Adequate receiver selectivity ensures that transmissions in neighbouring, and other frequency channels, does not affect the reception of the wanted signal. However, management of the radio spectrum tends not to involve specification of receivers on the basis that the radio spectrum licence holders will be incentivised to optimise the performance of their receivers. If they do not, then they risk suffering interference from neighbouring users.

This study examines, in detail, the issues involved in improving the performance of RF receivers in consumer equipment such as TVs and cellular devices. Within the study, we define RF receivers as the portion of the system which receives the RF signal and converts it to a demodulated baseband signal. The study focuses on the cost / performance trade off making a receiver less susceptible to interference based on technologies that can be envisaged over the next ten years with an indication of how this might change over the next twenty years.

In practice, the product manufacturer is likely to seek a competitive advantage, through cost engineering the product so that the product can be produced at the lowest possible cost. This often means that the product only just meets any required performance specifications or standards, unless there is a competitive advantage in exceeding them.

Receiver selectivity is rarely, if ever, promoted by product manufacturers. This may be because selectivity is a complex area which is often not well understood by the consumer. In practice, whilst the consumer expects the device to work well, if a device does suffer interference, the consumer may “blame” the network or atmospheric conditions. In this context, it is difficult for the product manufacturer to use a positive marketing message to justify additional cost for selectivity.

Recent thinking in this area has suggested that there may be cases where the regulator should intervene, perhaps during the definition of standards, to ensure that receivers with appropriate selectivity are developed to make them reasonably immune to interference from neighbouring frequency channels.

This study is intended to help the regulator better understand the relationship between receiver selectivity and product costs for consumer devices. Both adjacent channel selectivity (ACS) and blocking performance associated with out-of-band signals are considered.

The study report is arranged as follows:

- Chapter two looks at what receiver selectivity is, and what the receiver requirements are for good selectivity. An appendix to this chapter puts these requirements into context with a worked example for a DVB-T receiver.

- Chapter three looks at typical receiver architectures currently in use, beginning with super heterodyne receivers, capable of very good performance but needing a large number of discrete elements such as filters, before moving to architectures which lend themselves to more integrated approaches with fewer external components.

- Chapter four examines the drivers behind the silicon process choice for modern RF receiver ICs focussing on how the silicon material properties affect the parameters of both integrated transistor and passive components.

- Chapter five examines how current product cost and retail price are influenced by receiver selectivity requirements concluding with graphs of retail price versus receiver selectivity. In determining these device costs and retail prices, factors such as receiver architecture, power consumption, and product volumes are considered as well as the underlying economic factors within the industry such as typical margins.

- Chapter six looks at potential future receiver performance by first looking at some of the drivers and trends and then looking at some of the techniques which may be used to fulfil the requirements.

- Chapter seven concludes the report by discussing how device costs and retail prices may change with device selectivity in the future.

2 RECEIVER SELECTIVITY

This chapter investigates what is receiver selectivity, what receiver requirements are needed for good selectivity and why real world receivers may struggle to achieve the selectivity required for good system performance and spectrum use efficiency.

Receiver selectivity is the receiver's ability to receive the "wanted" signal in a frequency channel without being affected by "interfering" signals present in adjacent and other channels. Figure 2-1 shows for both an ideal and a practical receiver, the receiver's response to typical signals at the receiver's input.

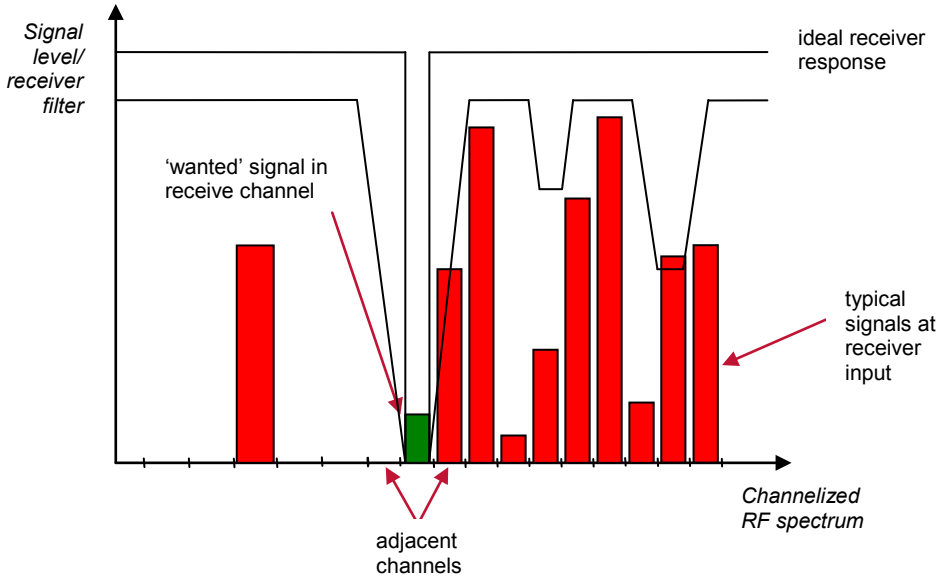

Figure 2-1: *Ideal and practical receiver response*

When just noise is considered, the receiver's sensitivity, i.e. the minimum level of radio signal that can be decoded, is defined by the receiver's noise figure and the signal to noise ratio (SNR) of the received signal required for adequate signal demodulation by the receiver as depicted in figure 2-2a. The ideal receiver is not affected by interfering signals in the adjacent and other frequency channels. Depending on how good the receiver's selectivity is, these interfering signals can dramatically affect the SNR needed for adequate signal demodulation, and the signal to (interference plus noise) ratio (SINR) becomes a more meaningful term as shown in figure 2-2b.

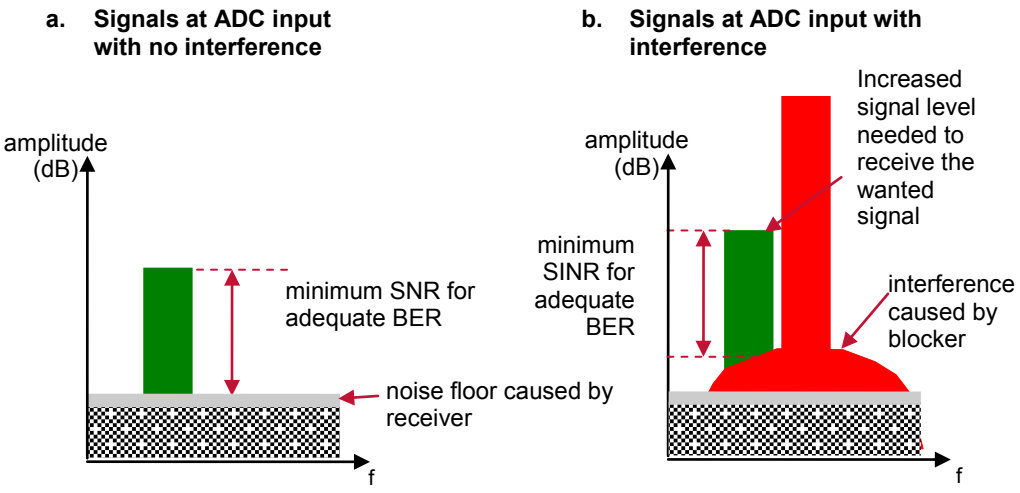

Figure 2-2: Signals reaching receiver's ADC, with and without interference

No single parameter defines a receiver's selectivity. Instead a number of parameters are used to define elements of the receiver's selectivity. These include receiver adjacent and alternate channel selectivity, far-off selectivity and intermodulation.

Receiver single channel selectivity defines the maximum unwanted interfering signal in a single channel a specified number of channels away from the wanted signal that the receiver can reject, whilst still adequately decoding the wanted signal received at a defined level. It is generally expressed as the ratio between the wanted and unwanted signal power level. For example adjacent channel receiver selectivity defines the maximum unwanted interfering signal in the adjacent channel (n ± 1) that the receiver can reject whilst still adequately decoding the wanted signal received at a defined level. Alternate channel selectivity is very similar. In this case the unwanted interfering signal is in the alternate channel (n ± 2). 'Far-off' selectivity considers interference at frequencies several channels away from the wanted signal. Intermodulation specifications typically consider large interfering signals in two channels.

2.1 How air interface standards and network planning deal with limited selectivity

Practically most receivers have poorer adjacent channel selectivity than other selectivity parameters such as alternate or far-off selectivity. Within a single air interface standard, such as GSM, this can be dealt with by specifying a receiver adjacent channel rejection which is significantly poorer than say the alternate channel specification and then using appropriate network planning techniques to avoid the limited adjacent channel rejection becoming an issue. The network planning techniques used may involve adjusting the position of base-stations, choosing appropriate frequencies of operation, and adjusting power levels.

Where two air interface standards use adjacent spectrum, for example DVB-T and the proposed use of the 790 to 862MHz band for cellular or broadband meet; cooperation between the two users is likely to be required. This may involve using a guard band to separate the two users or cooperatively planning the networks of both systems. In unplanned networks such as Wi-Fi, interference may limit performance in often undesirable ways.

As the RF spectrum becomes more crowded, if performance is to be maintained, receiver selectivity will become more important especially at the boundary between different radio systems.

2.2 Receiver requirements for good selectivity

Good selectivity demands a receiver with a wide dynamic range, i.e. with the ability to handle a large interferer whilst simultaneously receiving a small wanted signal. This requires a receiver able to handle large amplitude signals whilst also having a low noise floor.

2.2.1 Receive channel filter

With the exception of FFT based OFDM demodulation, a receiver's demodulator tends to have limited frequency selectivity. Therefore, for the demodulator to be unaffected by out of band signals, interfering signals need to be suppressed by the receiver to a level such that there is a sufficient C/N of the wanted signal at the receiver's demodulator. Ideally the wanted signal should be unaffected. The suppression of the unwanted signal can be implemented through a combination of analogue and digital filtering.

Traditionally much of the suppression needed was implemented through analogue fixed IF filters implemented using ceramic, crystal or SAW filters. These filters tend to have very high Q and accurately defined cut off frequencies. It is not practical to integrate these types of filters into an IC. Analogue filters using integrated components within the IC tend to have insufficient Q and cut off frequencies defined by the tolerance of the on chip components. These tolerances, when IC process, voltage and temperature variations are taken into account, can be as poor as 50%. Calibration techniques, either done during manufacture or within the device itself can improve the accuracy of the filter cut off frequency greatly but are often alone not adequate for good receiver performance. The dynamic range of the analogue filter is limited by on chip noise, and interference from other analogue and digital sources, and feed through of the received signals.

Instead modern receivers tend to rely much more on digital filters implemented after the analogue signal is sampled. The filter characteristics are much more repeatable and cut off frequency is accurately defined by the receiver's clock frequency. Assuming enough DSP processing power to perform the computations, the dynamic range of the digital filter is limited by the effective resolution of the analogue to digital convertor (ADC) used to convert the received radio frequency analogue signal to digital samples. The receive channel filter response is the combined response of the analogue and digital filters cascaded together.

2.2.1.i Sampling and analogue to digital conversion

All receivers using digital modulation need to convert, or quantise, the received radio frequency analogue signal to digital samples. One of the key characteristics of an ADC is its dynamic range. This is the difference between the amplitude of the largest and smallest signal it can handle simultaneously, determined, in the ideal convertor, by the number of digital bits each sample of the analogue signal is quantised to. The rounding error between the actual analogue signal and the quantised digital signal appears as quantisation noise.

The signal to quantisation noise ratio (SQNR) defines the ratio between the maximum signal that the ADC can handle and the quantisation noise produced. For an ideal analogue-to-digital converter, where the quantization error is uniformly distributed between −1/2 LSB and +1/2 LSB, and the signal has a uniform distribution covering all input levels, SQNR can be calculated from:

$$SQNR_{ADC} = 20\log10(2n) \approx 6.02n \text{ dB}$$

where n is the number of bits.

The SQNR determines the theoretical maximum ADC dynamic range. It will only be obtained when a signal with an amplitude covering the full scale input range of the ADC is sampled. If a lower amplitude signal is sampled by the ADC the SQNR of the sampled signal will be proportionally reduced.

In practice, effects such as distortion in the analogue section, jitter of the sampling clock and kT/C thermal noise introduced by the sampling switch, will create additional noise and spurious products reducing the dynamic range slightly. The ADC's SNR is a practical measure of a real ADC's maximum dynamic range. It characterises the ratio between the fundamental signal and

the noise in the sampled spectrum. Often the Effective Number Of Bits (ENOB) of the useful signal data in the ADC's output digital signal is used rather the SNR of the input signal. The ADC ENOB is always lower than the ADC's headline number of bits of resolution.

For maximum dynamic range the amplitude of the input signal must cover the full scale input range of the ADC. By using an automatic gain control (AGC) system, the gain of the receiver prior to the ADC can be continually adjusted ensuring a near full ADC scale signal is being sampled.

The ADC is not able to distinguish between frequency component in the sampled signal occurring above half of the sample clock frequency and lower frequency components of the signal. The higher frequency components are known as alias components. The ADC's Nyquist bandwidth is the frequency bandwidth over which the ADC can operate without forming alias components. It is half the sample clock frequency.

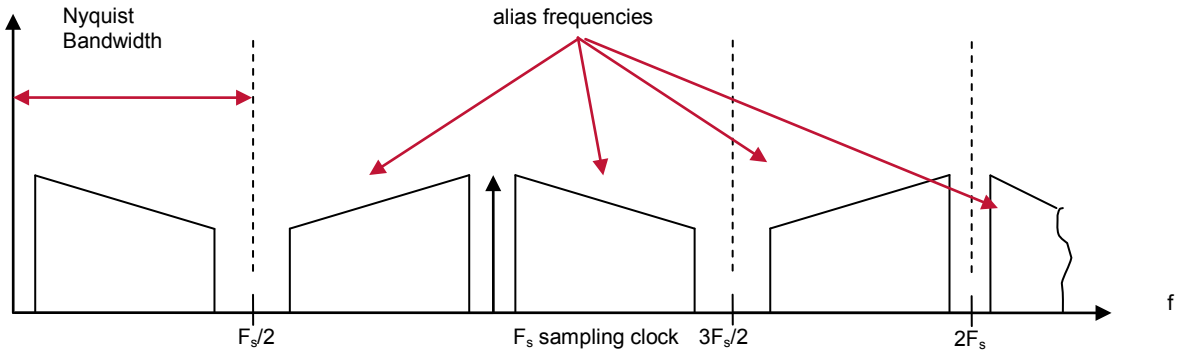

Figure 2-3: ADC's Nyquist bandwidth and alias frequencies

To avoid aliasing the signal must be filtered prior to sampling. This filter is known as an alias filter. The alias filter must remove any energy in the alias bands which, if sampled, would appear in the digital replica of the analogue signal. Although most common, it is not necessary to low pass filter the lowest set (baseband) of sampled frequencies, instead it is possible to bandpass filter an alias frequency removing signals at higher and lower alias frequencies; this is known as sub-sampling.

The analogue filter prior to the ADC in a digital receiver thus serves two purposes. The first as part of the channel filter, the second as an anti alias filter. By using a higher frequency sampling clock the anti-alias filter requirement is relaxed as the alias frequency is further from the wanted frequency allowing a filter with a shallower transition from its pass band to stop band.

The ADC's RMS quantisation noise error level is fixed by the input range and number of bits of the ADC. It is independent of bandwidth and in a simple "Nyquist" ADC the noise energy is spread evenly over the ADC's Nyquist bandwidth. When the Nyquist bandwidth is made wider by using a higher sample rate, the noise density (watts/Hz) is lower maintaining the same total noise energy per sample. Therefore within the desired signal bandwidth, the SQNR is improved. This is shown pictorially in figure 2-4.

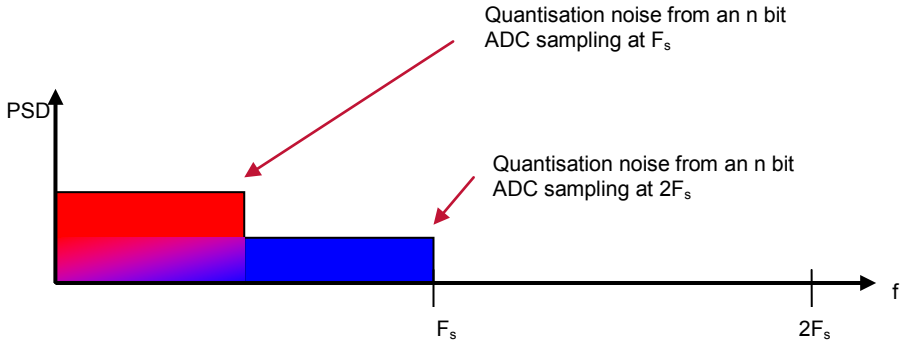

Figure 2-4: lowering the quantisation noise power density by raising the sampling rate

By doubling the sample rate it can be seen that the ADC noise density is halved giving a 3dB improvement in dynamic range. This is equivalent to adding an extra half bit to the ADC's resolution. Using this approach, known as oversampling, the designer can trade ADC resolution with the ADC sampling rate. Taken to extremes the ADC can be replaced with a single bit ADC, i.e. a single latched comparator operating at very high frequencies. However effects such as jitter of the sampling clock and KT/C noise introduced by the sampling switch increase with frequency so it is difficult to make high resolution, high speed ADCs

2.2.2 Phase noise and reciprocal mixing

Any practical local oscillator signal is not a perfectly pure sine-wave; instead it has noise sidebands known as phase noise. Several factors govern the amount of phase noise generated. For a tuned resonant oscillator, the type of oscillator typically producing the most pure waveform; the Q, or quality factor of the resonating tank circuit is an important contributing component to the oscillator phase noise.

Leeson's formula predicts a tuned resonant oscillator phase noise and is shown in figure 2-5. Its shape is governed by the following factors:

- Far from the carrier the noise is constant. This represents the broad band noise floor of the oscillator circuit.
- Closer to the carrier the Q of the resonator dictates the noise floor. The $1/f^2$ response comes from the filtering action of the resonating tank circuit.
- Close to the carrier, flicker noise with a 1/f characteristic combines with the resonator noise to produce a $1/f^3$ response.

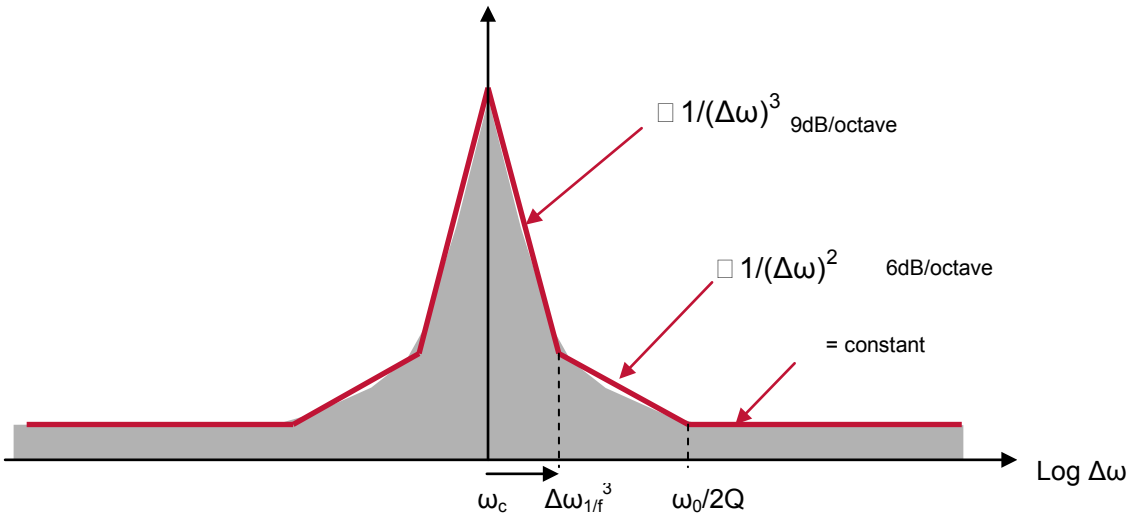

Figure 2-5: Local oscillator phase noise

In practice LO oscillators use a PLL based frequency synthesiser to lock the tuned oscillator to a frequency multiple of a reference oscillator. The synthesiser reduces the phase noise produced by the tuned oscillator close to the carrier within the loop bandwidth of the PLL to a level similar to that produced by the reference oscillator. Further from the carrier, beyond the loop bandwidth of the PLL, the synthesiser adds a small amount of additional noise.

Figure 2-6 shows the effects of mixing a wanted signal along with a larger adjacent channel signal, both depicted as a pure sine wave, to a lower IF frequency. The local oscillator noise is mixed with both the wanted signal and adjacent channel interferer. The down converted wanted signal is swamped by the adjacent channel signal, now with added phase noise. This effect is known as reciprocal mixing.

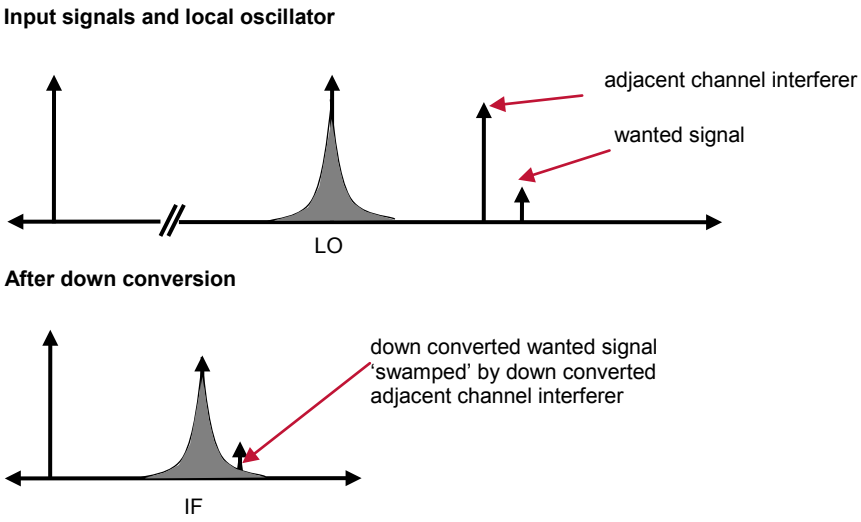

Figure 2-6: Reciprocal mixing

2.2.3 Receiver linearity

All analogue elements of a receiver such as amplifiers, mixers, and active filters have some nonlinearity, not least because they have a maximum signal they can process. These non-

linearities introduce distortion to both the wanted and any unwanted signals. This can lead to the creation of new interfering distortion products occurring at new frequencies. If these occur at critical frequencies within the receiver, this will affect the SINR needed for an adequate C/N at the demodulator.

In a real receiver multiple receiver stages are cascaded together. Intermodulation products developed in one stage are fed through to the next and therefore cumulate at each stage. As the relatively weak received signals are also generally amplified by each stage of a receiver, prior to demodulation, latter “back end” stages need to be able to better handle large signals than do the earlier “front end” stages.

2.2.4 Spurious responses

Ideally amplifiers just amplify the signals and all higher order products are minimised. Mixers on the other hand maximise the 2nd order sum (f1+f2) and difference (f1-f2) products. One of the mixer input frequencies is from a local oscillator (LO) allowing RF signals received at one frequency to be translated in frequency either up (up-conversion) or down (down-conversion) to another frequency. If the wanted signal at $\omega_c + \omega_{if}$ is mixed with a local oscillator at $-\omega_c$ the sum of the signals is at ω_{if}. At the same time any signals or noise at the “image” frequency $\omega_c - \omega_{if}$ are also translated to ω_{if}. This is known as a spurious response and is shown in figure 2-7. The mixer’s amplitude response to the image product is identical to the wanted frequency. Handling the image frequency effectively is a critical part of receiver design as any signals at the image frequency will cause interference and degrade the receiver’s selectivity. How this image response is handled is the key difference in the various receiver architectures discussed in chapter 3.

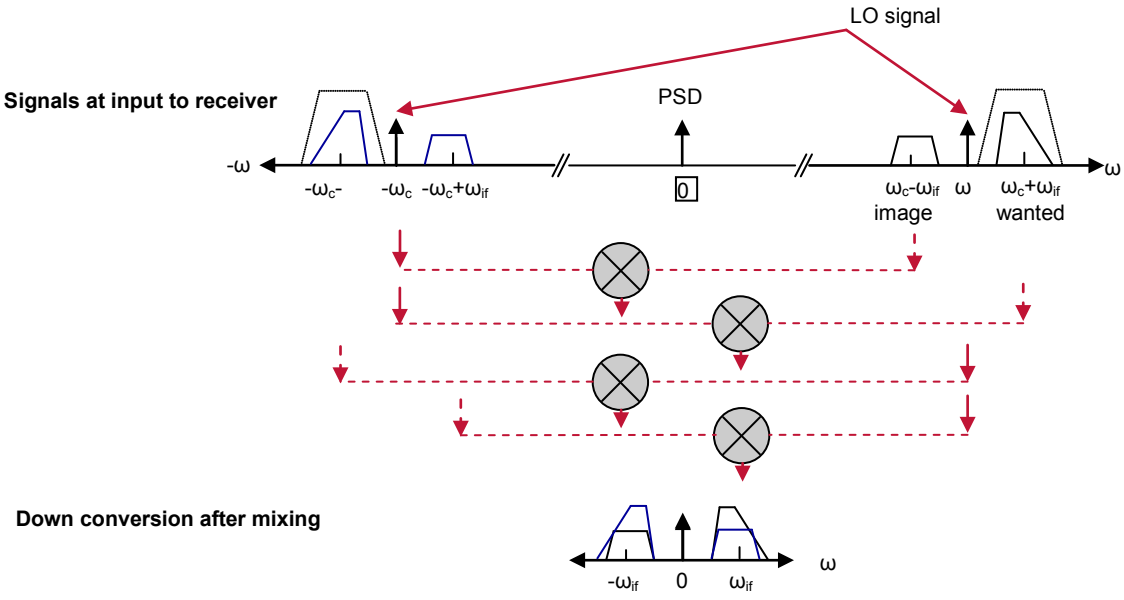

Figure 2-7; Image products after downconverson

In addition to the image response, mixers have other spurious responses at $mf_{RF} \pm nf_{LO}$. In many receiver designs it is possible to band pass filter the received signal before amplification so that the receiver only needs to deal with RF signals in a relatively narrow frequency band, limiting the number of spurious products which may be generated.

Unwanted received signals at any frequency which modulate with the local oscillator or its harmonics, or with any signal present in the receiver (e.g. digital clocks) at another frequency, to form a product at the wanted received frequency, will be an issue degrading or completely

blocking the reception of the wanted channel. Local oscillator harmonics can be a significant issue in wideband receiver designs such as TV tuners.

2.3 Transmit adjacent channel power leakage

Although not directly a receiver performance issue, transmitters do not constrain all their transmit power to their allocated transmit frequency channel. This can be due to the modulation scheme used or local oscillator phase noise. Any energy transmitted on adjacent channels is known as adjacent channel power leakage, ACPL. The ACPL sets a limit on the selectivity that can be obtained with an ideal receiver.

2.4 Summary

Table 2-1 lists the various receiver impairments discussed above and examines the number and type of interferers required to cause a selectivity issue.

Impairment	**Susceptible frequency channels**	**Number and type of interferer required to cause a selectivity issue**
Channel filtering (analogue and digital)	Mostly the adjacent channel with channels further from the wanted having a reduced susceptibility as the analogue filter transitions to its stop band. The ultimate rejection of the filter will influence the far-off blocking performance of the receiver.	At least one
Linearity	Any frequencies that are not adequately suppressed by an input filter prior to processing by the receivers analogue element, amplifiers, mixers, etc operating non-linearly.	Either one with non constant envelope modulation or two signals. The intermodulation products must fall on the wanted, IF, DC frequency etc (receiver architecture dependent).
Spurious responses	Any that are not adequately suppressed prior to processing by the receivers analogue elements, amplifiers, mixers, etc; that combine with the receiver's LO (and its harmonics) or internally generated spurious frequencies which result in a spurious product being generated that falls on the wanted, IF, DC frequency (receiver architecture dependent).	At least one
ADC aliasing	Any frequencies that are not adequately suppressed by the ADC's alias filter prior to sampling.	At least one with spacing such that interferer is aliased into the ADC's desired band
LO phase noise – reciprocal mixing	Mostly the adjacent channel with channels further from the wanted have a reduced susceptibility as the LO phase noise levels out. The ultimate LO noise floor of the filter will influence the far-off blocking performance of the receiver	At least one
Transmit adjacent channel power leakage	Mostly the adjacent channel with channels further from the "wanted" having a reduced susceptibility as the transmitter output filter transitions to its stop band.	At least one

Table 2-1: Receiver selectivity impairments

From the table it can be seen that, with the exception of receiver linearity, a receiver's selectivity can be determined using a single interferer. To comprehensively characterise a receiver's selectivity, potential interferers at any frequency which may be present at the receiver's input need to be considered.

Appendix One gives a numerical example of how these factors actually affect a practical receiver's selectivity.

3 RECEIVER ARCHITECTURES

This chapter looks at typical receiver architectures currently used, beginning with super heterodyne receivers, capable of very good performance but needing a large number of discrete elements such as filters, before moving to architectures which lend themselves to more integrated approaches with fewer external components. These more integrated approaches allow almost the entire receiver to be fabricated in silicon, so that when manufactured in large volumes the unit cost is very low, and performance, after calibration, between one receiver and another is very consistent. The low cost is due to very low material costs, and rapid, highly automated, assembly and test.

This highly integrated approach lends itself very well to high volume applications such as cellular or broadcast where total world, per annum, volumes are greater than one billion[1]. However the very high development costs; especially when using the smallest feature size, where the IC mask cost alone can be a few million dollars[2], prohibit this approach for smaller volume applications, such as PMSE.

3.1 Super heterodyne

Superhet receivers were used in the majority of receivers from when Armstrong first popularised the approach in 1917 up until around the year 2000. At this point alternative receiver architectures more suitable for complete integration into an IC, such as direct conversion, became increasingly popular.

A single stage superhet is shown in figure 3-1. The mixer mixes the received RF signal with a local oscillator (LO) signal converting the received RF signal to an intermediate or “IF” frequency. The LO signal can be higher than the RF signal (high side LO) or alternatively lower than the RF signal (low side LO).

-RF + LO = IF	high side LO
RF - LO = IF	low side LO

The image filter is required to stop the receiver from responding to the signals at the image frequency. Receiver selectivity is gained by the IF filter bandpass filtering the wanted signal suppressing power in the channels adjacent to the wanted signal. Following the IF filter the signal can then be sampled using a sub-sampling approach as discussed in section 2.2.1.i. With this approach the IF filter acts as the ADC’s anti alias filter. Alternatively it can be converted by a second mixer to baseband for digital sampling or analogue signal detection.

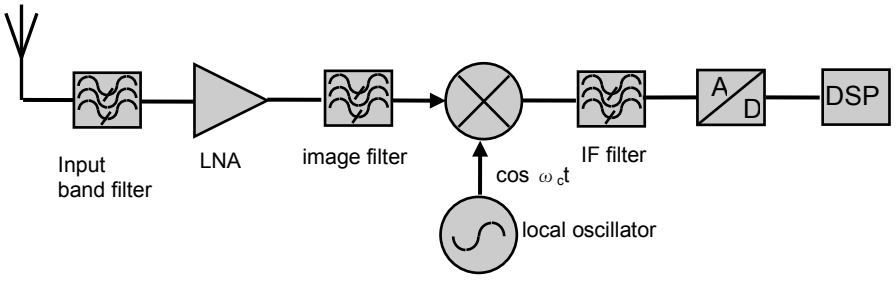

Figure 3-1: Superhet architecture

In most receivers the IF is at a fixed frequency. To allow the receiver to be tuned to a range of receive frequencies the local oscillator is varied.

[1] IDC Worldwide Quarterly Mobile Phone Tracker, February 4, 2009
[2] 20th July 2009. “Cheaper options for chip designs” IET, http://kn.theiet.org/magazine/issues/0913/cheaper-chip-designs-0913.cfm [accessed 24th July 2009]

3.1.1 IF filter

The IF filter is generally at a lower fixed frequency than the received RF signal allowing a high Q filter to be implemented to provide the required selectivity. The ideal IF filter, shown in figure 3-2, should have a flat pass band and good group delay to pass the wanted frequency channel without distortion; with very steep skirts on both sides of the pass band in order to be able to reject the adjacent channels well.

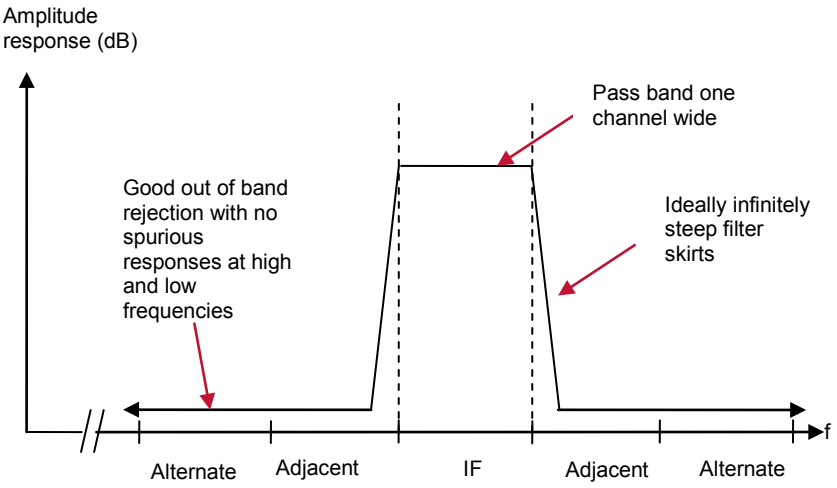

Figure 3-2: Ideal filter response

A number of typically passive technologies are used to realise IF filters including crystals, ceramic, and SAW filters. These filters tend to be available for a range of system standard IF frequencies including 455KHz - AM, 10.7MHz – FM broadcast and 36MHz – TV. Due to the high Q and stability at high frequencies required, the IF filter is rarely integrated into an IC.

3.1.2 Image filter

Assuming a low side local oscillator is used, where the local oscillator is at a lower frequency than the wanted receive frequency, i.e. the wanted frequency is at $\omega_c+\omega_{if;}$ then the input band filter and image filter together must reject the receiver's image response at $\omega_c-\omega_{if}$. For a high side local oscillator the opposite is true, the wanted frequency is at $\omega_c-\omega_{if}$ and the receiver's image response is at $\omega_c+\omega_{if}$.

These filters need to operate directly at the RF frequency and as they have limited Q don't provide any real channel filtering. Any filter placed directly at the input of receiver will directly affect, and may dominate, the receiver's noise figure and therefore the receiver's sensitivity to receive weak signals; therefore it must have low insertion loss. The insertion loss of the filter following the LNA (low noise amplifier) is less critical. However, in this case, the LNA must have adequate linearity to not distort any signals passed by the input band filter. A significant positive side effect of an image filter is that it does attenuate other potential 'far-off' interferers apart from the image frequency before the signals reach the mixer. This reduces the effects of reciprocal mixing and limited channel filtering dynamic range on the receiver's far selectivity.

The image filter may have a fixed pass band, wide enough to pass all the channels within the tuning range of the receiver. These types of filters, depending on the frequency of operation and the pass bandwidth required, can be made from various different materials. At lower frequencies wire wound inductors and ceramic capacitors are often used whilst for cellular physically small SAW filters etched on a variety of substrates are commonly used.

As the tuning range of the receiver increases, the lowest received frequency and the highest image frequency get closer together making it more difficult to get adequate rejection until eventually the lowest wanted frequency and highest image overlap. At this point it becomes impossible to realise the image filter using this approach.

Another approach is to use a much narrower filter which has a tuneable pass-band. It is generally more complicated to build a tuneable filter. However assuming the filter can "track" the receive frequency it does allow a receiver with a wide tuning range to be developed. Tracking filters are often realised by using varactor tuned diodes to act as variable capacitors in an inductor capacitor tuned circuit. These allow the filter's pass band frequency to be adjusted with a DC voltage.

3.1.3 Image reject mixers

The superhet receiver discussed so far uses the amplitude response of filters to reject the image frequency. Another approach is to use mixers operating in quadrature to cancel out the image signal. Two architectures are often used to realise this approach, Hartley and Weaver. For both the Weaver and Hartley image reject mixers, it is necessary to generate signals in quadrature to each other. For perfect image rejection these quadrature signals need to be phase and gain matched across the frequency band of interest.

3.1.3.i Hartley

The Hartley image rejection architecture is shown in figure 3-3. The RF signal is down converted by quadrature LO signals. The resulting IF signals are then low pass filtered and after one is phase shifted by 90° the IF signals are combined. This result is that, depending on which channel is subjected to the 90° phase shift, either the image or wanted channel being rejected.

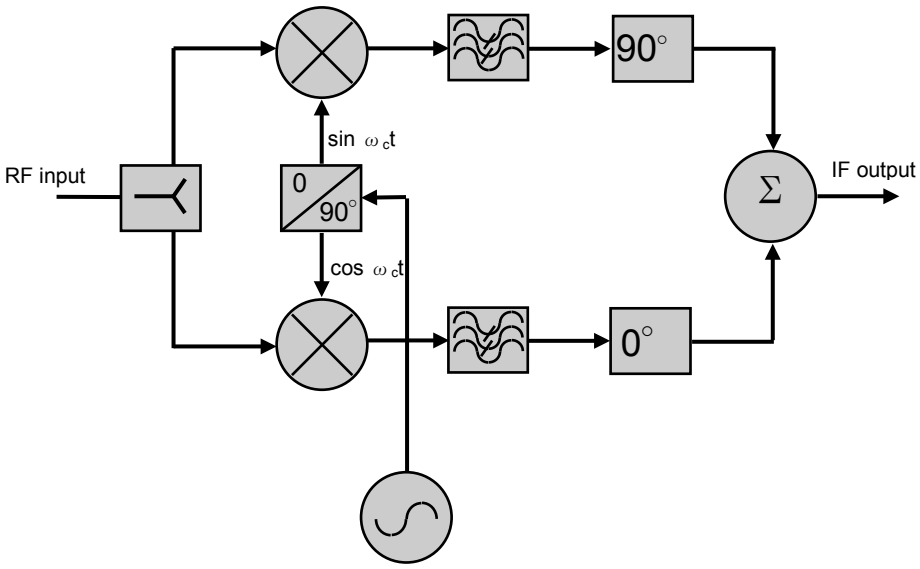

Figure 3-3: Hartley image reject architecture

The image rejection obtainable is dependent on how close to 90°the phase shift of the local oscillator is, how similar are the amplitude responses of the two arms of the mixer and how close to 90° is the final phase shift. It can be shown that an amplitude mismatch of 0.1dB and a phase mismatch of 1° yields around 41dB of IRR. To realise this degree of IRR requires careful design and possibly some form of calibration to minimise any imbalances.

3.1.3.ii Weaver

The Weaver approach overcomes the amplitude mismatch issues caused by needing to add a 90°phase shift to one arm of the quadrature mixer by adding a second pair of mixers to realise the phase shift. This second mix does create another set of image frequencies which need to be addressed. One approach is to use a second IF centred on DC (0Hz) with the sampling of the signal being done in quadrature as shown in figure 3-4.

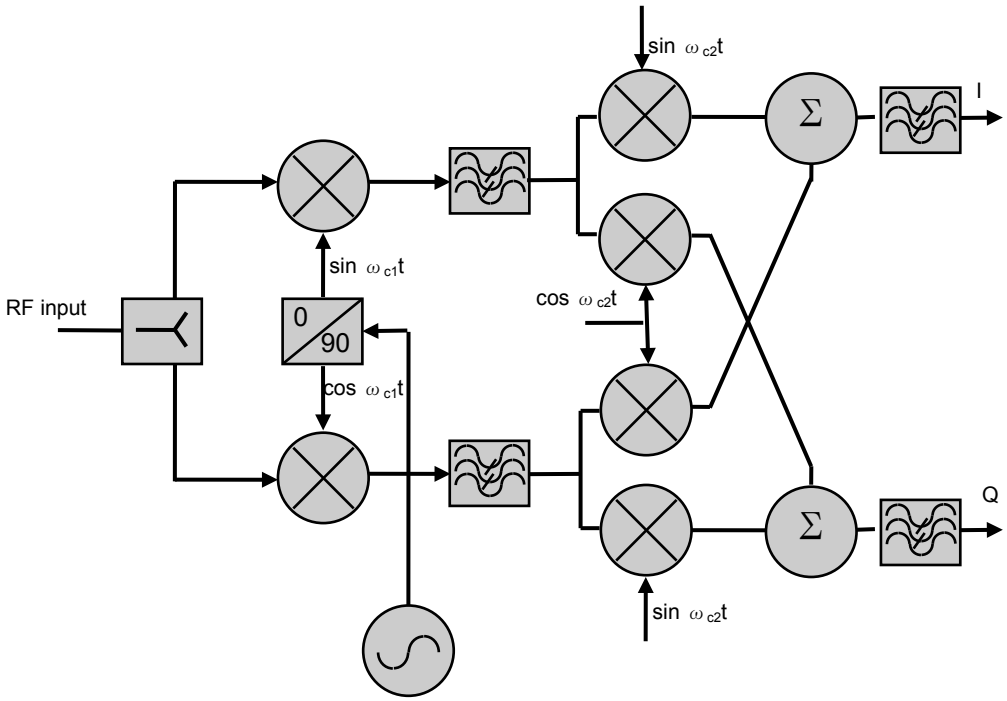

Figure 3-4: Quadrature Weaver architecture

3.2 Zero IF Receiver

The zero IF receiver overcomes the superhet IF image response issue by directly converting the signal to baseband centred on 0Hz using two mixers operating in quadrature. Its architecture is shown in figure 3-5.

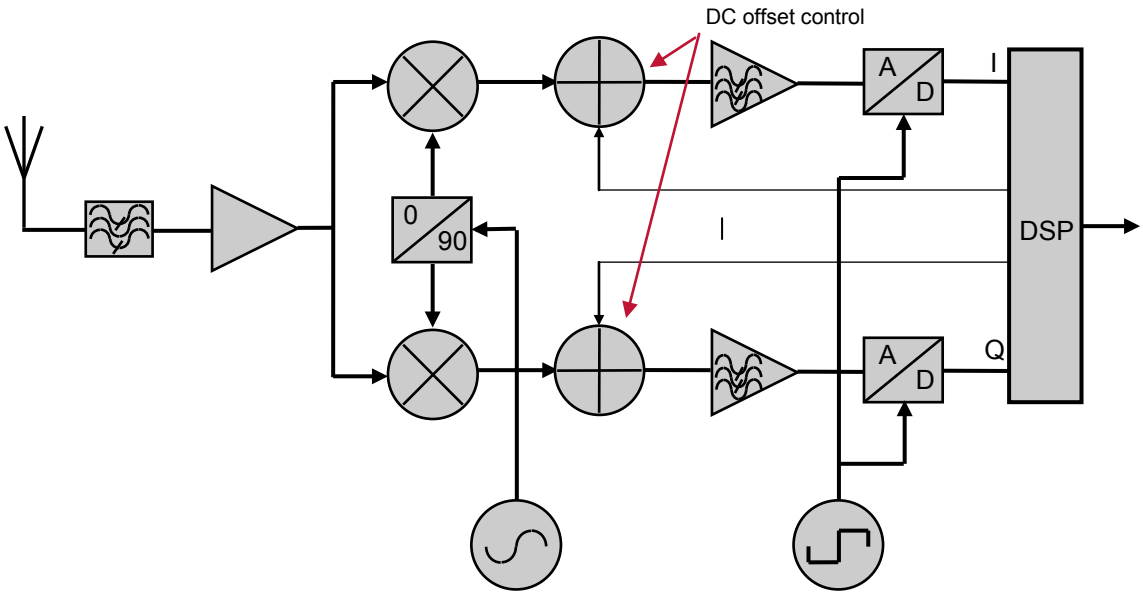

Figure 3-5: Zero IF receiver

In the zero IF receiver the LO is at the same frequency as the received RF frequency. This causes the down converted image signal to fall directly on the wanted signal so that both the wanted and image signals are mirror images of each other, with each reflected around the frequency axis.

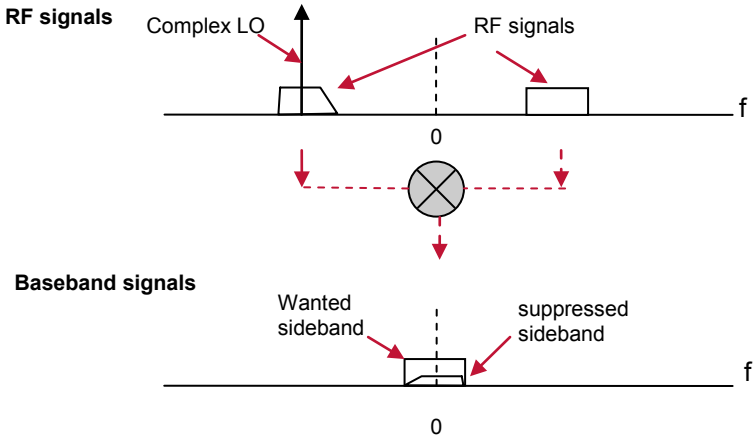

Figure 3-6: Zero IF image down conversion with image suppression

Whilst the zero IF approach deals with the image response, the receiver still has significant spurious responses at odd LO harmonics, i.e. $3f_{LO}$, $5f_{LO}$, $7f_{LO....}$ This is commonly due to the use of a square wave LO signal. In a zero IF receiver, as the local oscillator is at the RF frequency, any received signals at these frequencies will cause interference. In narrow band receivers, where frequencies at $3f_{RF}$ etc do not need to be received, fixed frequency input filters can be used. In very wide band receivers, such as cable TV receivers needing to cover 48 to 860MHz, this can be a significant problem.

The analogue low pass filters following the mixer help provide the receiver's selectivity and act as anti-alias filters to the ADC. If all the selectivity is provided by these filters, they must have a cut off frequency at half the channel bandwidth and must reject the adjacent channel and other channels further from the wanted frequency by the selectivity required.

Usually some of the receiver's selectivity requirement is realised with digital filters following the ADC. Digital filters don't suffer from many of the limitations of analogue filters such as their performance being affected by component and silicon process tolerances, cross talk and noise. This allows their performance to be closely defined and very repeatable. In addition they can be implemented in low cost digital CMOS making use of either DSP processors or custom digital circuitry. With this approach the ADC must have enough bits of resolution to sample any high level adjacent and other channels without clipping, whilst not degrading the low level wanted signal with quantisation noise.

Whilst the zero IF approach minimises the image issue, the architecture does introduce other issues. These issues are mainly centred on needing a lot of amplifier gain to amplify signals near or at DC. These include:

- Second order receiver nonlinearity causing spurious products at DC
- Local oscillator leakage causing varying levels of DC offsets
- Amplifier DC offsets
- Flicker noise reducing sensitivity

These limitations detract from the simplicity of the zero-IF approach requiring a receiver with good IP2 performance and extensive calibration to overcome the DC offset issues.

3.3 Low IF receiver

A low IF receiver attempts to overcome the DC offset and 1/f noise issues associated with zero IF receivers whilst still using an approach that lends itself to a high degree of integration. Many radio standards require less selectivity for interfering signals occurring in adjacent channels than they do for interfering signals in other channels. Low IF receivers often make use of this by choosing an IF frequency which causes the image frequency to fall into an adjacent channel.

Using a low IF frequency allows the IF channel filter to be integrated into silicon or implemented digitally. Figure 3-7 shows where the adjacent, alternate channels and image response are down converted to when a low side local oscillator is used positioned on the edge of the wanted frequency channel. After down conversion, signals in one of the channels adjacent to the wanted receive channel fall into the wanted channel and signals in one of the alternate channels fall into the receiver's adjacent channel. As the wanted signal is just above DC, DC offset and 1/f noise issues are not such a concern as they are in zero IF receivers.

It can be seen that the receiver must achieve sufficient image rejection to meet the required adjacent channel specification. In addition, it can be seen that the lower alternate channel, after down conversion, lies next to the wanted frequency. Any adjacent channel leakage power (ACPL) from the lower alternate channel originally transmitted on the high side of the transmission, after down-conversion will fall into the wanted channel. This energy cannot be suppressed by filtering after down conversion and therefore the receiver must have sufficient image rejection to adequately suppress the signal. It is found in systems such as GSM with poor ACPL performance, that ACPL sets the image rejection requirements in a low IF receiver.

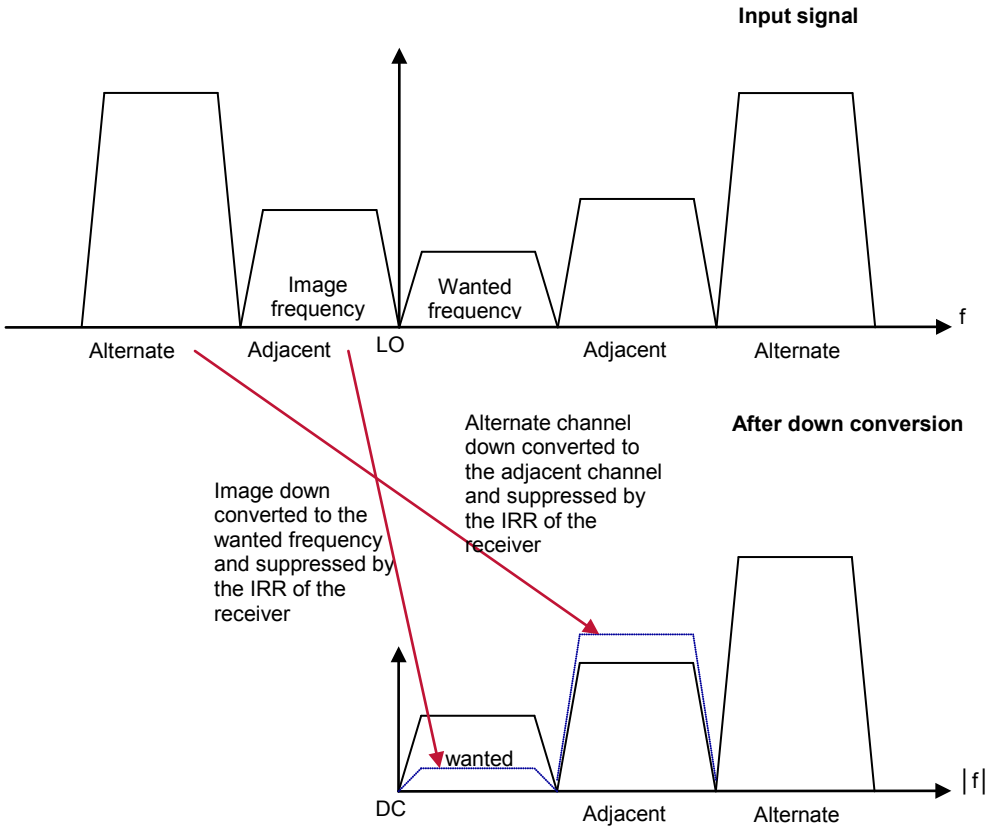

Figure 3-7: Low IF down-conversion

As the image is so close to the wanted frequency, an image filter at the receiver input can't be used, however image reject techniques can be used. One approach is by using the dual quadrature mixer Weaver architecture as shown in figure 3-8. The second set of mixers is implemented digitally, and the second digital LO is set so that the output is centred around DC.

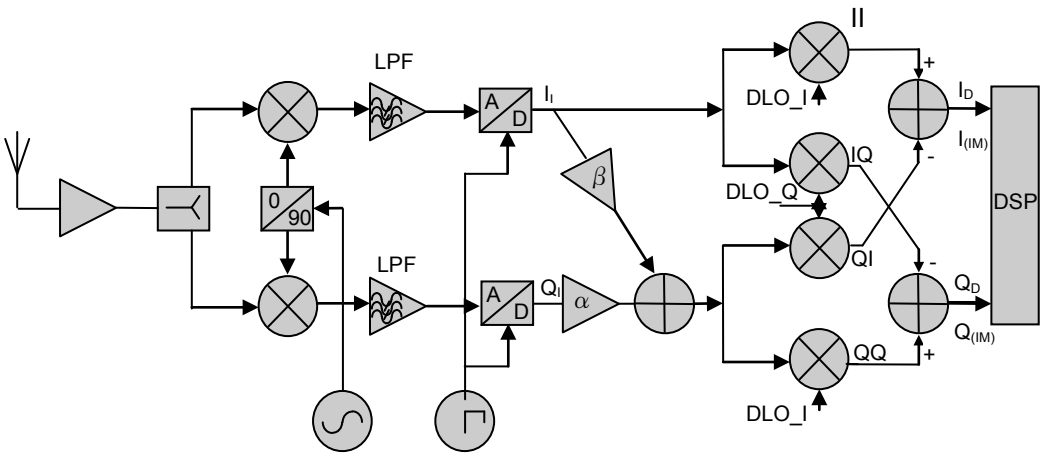

Figure 3-8: Digital low IF Weaver architecture

The signals are down converted to a low IF frequency by the first set of mixers. The I and Q signals are low pass (anti alias) filtered and sampled. As discussed in section 3.1.3, due to phase and gain errors between each arm of an analogue image reject mixer, it is very difficult to achieve greater than 25 to 35dB of image rejection without calibration. Overcoming the imbalances of the first analogue mixer can be achieved using the amplifiers just after the ADCs with gains α and β to modify the I and Q signals slightly allowing image rejection figures of typically up to 40dB. The values of α and β needs to be determined by a calibration process. The accuracy of the calibration process and final image rejection obtainable is at least in part due to the resolution of the ADCs used.

A *polyphase* band pass filter can be used instead of digital down conversion to obtain reasonable image rejection. The key attribute of a *polyphase* filter is that it provides a different filter response for positive and negative frequencies unlike most filters which just respond to the absolute frequency of the signal and not the sign of the signal. Using this approach an 'image reject' filter can be built.

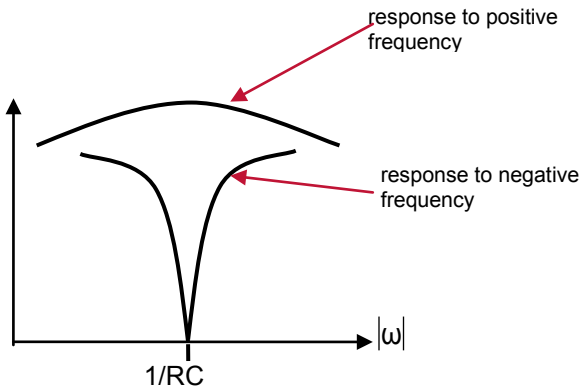

Figure 3-9: Ideal response of a single stage polyphase image reject filter

Figure 3-10 shows an image reject polyphase filter integrated with a quadrature mixer to implement a low IF receiver. Whilst this approach shows the polyphase filter which rejects the image frequency it does not show the channel filter. The channel filter may be implemented in the analogue domain prior to the ADC. Alternatively it may be implemented digitally. In either case there needs to be sufficient filtering prior to the ADC to avoid aliasing issues.

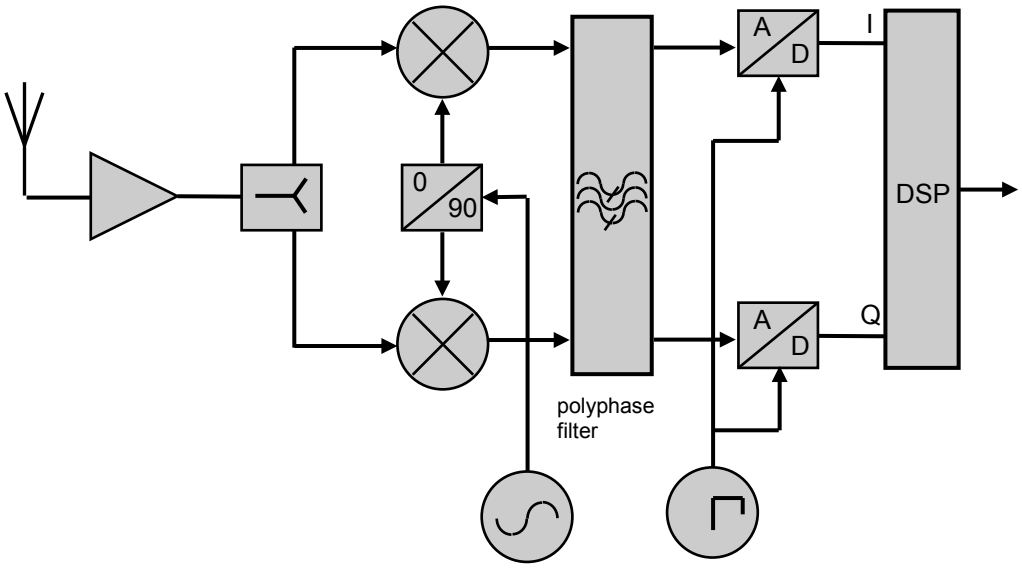

Figure 3-10: Low IF receiver with polyphase filter

Phase and amplitude variations in the quadrature mixer and imbalances in the polyphase filter all contribute to limiting the receiver's image rejection.

3.4 Architecture comparison

Table 3-1 summarises the selectivity limitations of the various receiver architectures discussed.

	Superhet	**Low IF**	**Zero IF**
Sensitivity	LNA, flicker noise not important	LNA, flicker noise may be important	LNA, flicker noise more important
Image rejection	Image filter and image reject mixer rejection. Enough image rejection needed for adequate selectivity at the image frequency, typical tens of MHz away from the wanted	Phase and amplitude matching of mixers plus: • Digital approach - enough ADC dynamic range and bandwidth for digital dual down convertor • Analogue approach - polyphase filter rejection Enough image rejection needed for adequate C/N and selectivity at frequencies close to the wanted ACPL from the alternate channel mixes into the wanted signal and may dictate the image rejection required	Not a selectivity issue Enough image rejection needed for adequate C/N for signal decoding
Spurious response rejection	Limited by: • Input filter rejection • Mixer $mF_{RF}\pm nF_{LO}$ response • Mixer RF to IF and LO to IF isolation It is likely that all the spurious frequencies will be in frequency channels well away from the wanted channel and therefore may be subject to interfering signals at frequencies much higher than the wanted signal	Limited by: • Input filter rejection • Mixer $mF_{RF}\pm nF_{LO}$ response • Mixer RF to IF and LO to IF isolation A number of the most significant spurious frequencies due to factors such as the mixers ½ IF response will be in the wanted or adjacent frequency channels and therefore not subject to interfering signals higher than the adjacent channel signal	Limited by: • Input filter rejection • Mixer $mF_{RF}\pm nF_{LO}$ response • Mixer RF to IF and LO to IF isolation A number of the most significant spurious frequencies due to factors such as the mixers ½ IF response will be in the wanted frequency channel and therefore not subject to interfering signals at frequencies higher than the wanted signal
Channel filtering	Limited by: • Discrete filter + ADC dynamic range • LO phase noise External high Q filters needed	Digital approach limited by: • ADC dynamic range • mixer image rejection • LO phase noise Analogue approach limited by: • Integrated analogue filter • mixer image rejection • polyphase filter rejection • LO phase noise Medium Q filters needed, either integrated analogue filters or digital	Limited by: • integrated analogue filter • ADC dynamic range • LO phase noise Lowest Q filters needed, either integrated analogue filters or digital
Linearity	IP3 important, IP2 not so important	IP3 and IP2 important	IP2 critical

Table 3-1: Receiver architecture selectivity limitations

Key points from this table are:

- The superhet's image filter can be eliminated in the zero IF and low IF architectures. However in a superhet, an image filter, whilst also acting as an input RF filter, can provide very significant spurious response rejection, channel filtering and enhanced linearity for interfering signals several channels away from the wanted frequency.

- A superhet's adjacent channel filtering is provided by a discrete fixed frequency IF filter with potentially very high Q, and post ADC digital filtering. The IF filter's high Q can relax the receiver's ADC dynamic range requirement compared to a zero or low IF receiver.

- A Low IF receiver's adjacent and alternate channel selectivity is provided by a combination of the receiver's image rejection and channel filtering. Selectivity of one of the adjacent channels is provided solely by the receiver's image rejection. For interfering signals several channels away from the wanted, minimal input filtering will limit receiver selectivity and linearity.

- In zero IF receivers, adjacent channel selectivity is provided by the receiver's integrated analogue and digital filtering. Zero IF receiver's performance can be limited by even order nonlinearity. Minimal input filtering will limit receiver selectivity and linearity for interfering signals several channels away from the wanted.

3.5 Typical architectures used in various applications

The table below highlights some of the architectures used for various applications. Many applications are listed under various architectures. This is because often the optimum architecture choice is not clear cut and may depend on the preferred approach of individual manufacturers due to the silicon process, intellectual property or knowledge and experience an individual manufacturer has available.

Architecture	Application	Driver for architecture choice
Super-heterodyne	'Traditional' analogue or analogue/ digital TV Tuner	Analogue video output required. Good selectivity obtained through use of SAW filter and input tracking filter.
	PMSE and other low volume applications	Can be realised without custom application specific ICs. These would need large markets for economically viable development.
	Traditional FM radio	Analogue audio output required, low cost but fairly large
Low IF (analogue polyphase filter approach)	Short range wireless	Only requires one ADC
	DAB	Allows the RF IC to interface with existing baseband ICs incorporating one ADC and no support for DC offset cancelation.
Low IF digital approach	GSM	Standard doesn't require good adjacent channel selectivity so image rejection is not to much of an issue, significantly reduced 1/f noise (especially if implemented in CMOS) and DC offset issues
	Digital TV tuners	Significantly reduced 1/f noise (especially if implemented in CMOS) and DC offset issues. Higher power consumption arising from using higher speed ADCs are not such an issue in mains powered device
Zero IF	UMTS	Good selectivity required at low power (lower ADC clock speed)
	Digital TV tuners	Good selectivity and noise figure possible especially if using a non pure CMOS process so 1/f noise is not an issue.

Table 3-2: Typical architectures used in various applications

4 DRIVERS BEHIND SILICON PROCESS CHOICE

This chapter investigates why a silicon vendor might choose a particular silicon process for an RFIC. With the increased use of the fabless chip production model, i.e. the semiconductor vendor outsources the fabrication of the IC, a wide choice of silicon process is available to the silicon vendor.

Many of the processing costs in IC manufacture are silicon area related, so minimising area is a way of reducing manufacturing cost and also allows more functionality to be integrated into a single IC. The drive for reduced cost and increased functionality has created the relentless pressure to minimise the dimensions of the transistors used in an IC process. This is exemplified by the reduction in CMOS gate size from several microns in the first 4000 series CMOS (long channel devices) in the late 1960s to gate sizes of a few tens of nanometres today (short channel devices). To date, this gate size reduction has allowed 'Moore's law' prediction to be fulfilled.

Dernard's scaling theory, first published in 1974, suggested that for proper scaling of MOSFETs, the devices physical dimensions, voltage, and doping needed to be adjusted appropriately; and with device scaling, for a digital device, would come reduced circuit delays. An analogue amplifier's equivalent parameter to circuit delay is its f_T , the frequency at which its current gain reduces to one.

As CMOS feature size has reduced high frequency performance, exemplified by the device's peak f_T, has improved considerably. This has allowed CMOS high frequency performance to overtake other more expensive approaches such as SiGe BiCMOS. For example, CMOS MOSFETs have similar high frequency performance to SiGe BiCMOS transistors but have half the feature size and hence occupy a quarter of the silicon area, i.e. a 130nm BiCMOS transistor has comparable high frequency performance to a 65nm CMOS MOSFET.

As CMOS geometries have reduced, effects due to very high electric fields being created with even moderate voltages, and other short channel effects such as oxide breakdown have reduced the device's supply voltage required to a little over 1V and made the device's behaviour much more complex. A short channel device's behaviour cannot be well defined by simple equations; instead complex modelling is required to accurately define this. The supply voltage plays a significant part in determining the large signal handling and IP3 of the device. The link between device node size, supply voltage and f_T is shown in figure 4-1.

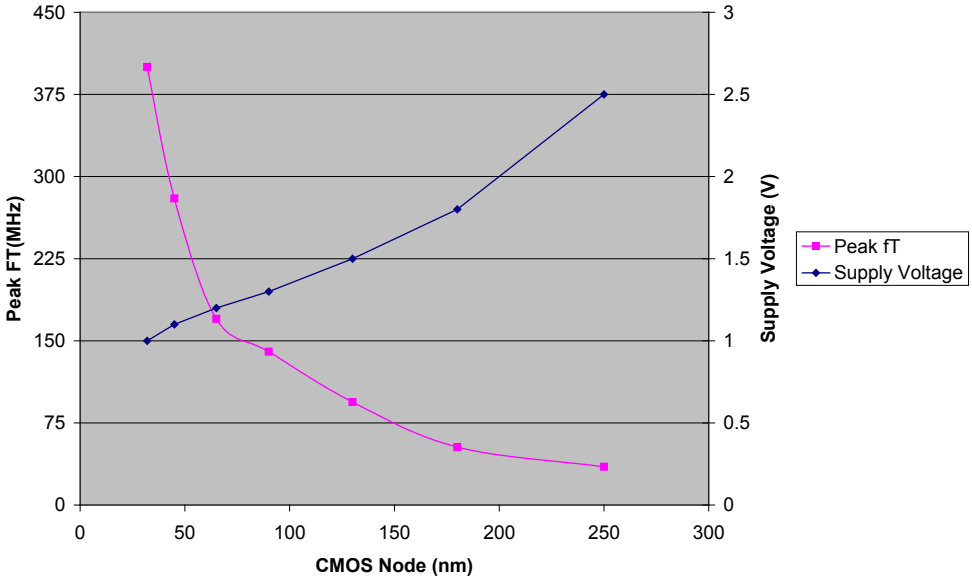

Figure 4-1: Supply voltage and f_T with CMOS node[3]

Another important limitation of deep sub micron CMOS includes a noise mechanism, occurring at low frequencies, known as flicker noise or 1/f noise. This limits receiver sensitivity when low and particularly zero IFs are used.

CMOS digital circuits in general scale. This allows a design initially implemented in one CMOS process to be easily migrated to the next node as it becomes available, lowering IC cost. In general, for stand alone RF circuits, surface area is dominated by passive devices such as capacitors, inductors and I/O pads. These do not scale in the same way so the same economies are not applicable. Instead BiCMOS often offers the lowest cost option despite an additional 20% increase in processing complexity[4].

There is a drive for the IC manufacturer to increase functionality in a shrinking form factor, and reduce 'time to market'. This has forced analogue RF components, which often occupy 30 to 40% of the board space, and take very significant design effort to successfully implement, to be integrated where possible. This has helped drive the choice of circuit topologies and receiver architectures to those which can be more easily accommodated in an IC process.

Circuit topology techniques include eliminating inductors where possible, using matched component techniques based on the using the ratio between two components rather than the absolute value of a component, using digital approaches where possible either directly or to calibrate an analogue stage. Precision is often obtained by referencing the circuit to the receiver's frequency reference, accurate to a few part per million (ppm) rather than to a precise amplitude reference. Examples of this include using an onboard oscillator to continuously calibrate a corner frequency of a slave filter whilst a similar master filter is used to filter the actual signal.

Whilst improved performance but more costly IC processes, such as SiGe, can be used for RF specific integrated circuits, a final potential integration step is to combine the RF with the baseband. As much of the die area will be taken up with digital circuits, (especially true if a relatively large node process is used) this approach demands that the most advanced digital silicon processes are used even if the process is not entirely suitable for high performance analogue RF circuits. These steer the RF architecture towards low IF designs removing the 1/f noise issues and new architecture approaches. These will be discussed in chapter 6.

[3] ITRS, ITRS 2003 to 2007
[4] Thomas H. Lee, 2004, "The Design of CMOS Radio Frequency Integrated Circuits", 2nd Edition, Cambridge University Press

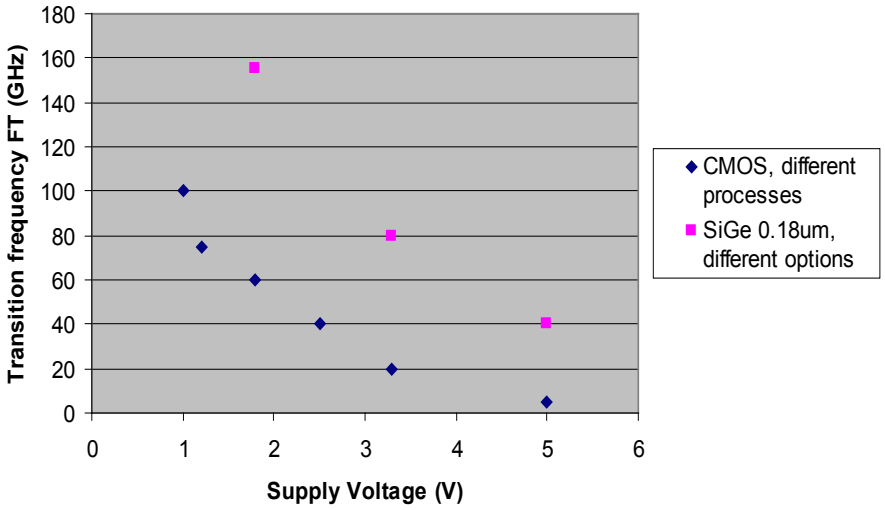

Figure 4-2: Frequency versus voltage plot for CMOS and SiGe[5]

As the combined design effort for both the baseband and RF parts of the system need to be concentrated into the one very large IC design, the designers must be very confident that the design requirement is stable and unlikely to change. The approach also limits the designer's ability to upgrade the digital design to the latest, lowest cost process. Therefore this approach is only ever likely to be seen in well defined very high volume applications. The approach was first used with Bluetooth and is now used with more demanding systems such as UMTS[6].

It can be seen that the choice of silicon process for an RFIC depends on:

- Performance requirement. Processes such as SiGe offer the best performance but are more expensive than CMOS, especially for designs with significant amounts of digital circuitry.

- Cost requirement. Deep sub micron CMOS offers the lowest cost per transistor. This is especially important for designs with significant amounts of digital circuitry and allows very high levels of integration. Some important RF parameters such as linearity are poor.

[5] Alain-Serge Porret and Alvin Wong, "Silicon- Germanium: The superior semiconductor technology for solid state TV tuners" Xceive Corporation, June 28th 2006 http://www.videsignline.com/189601591%3Bjsessionid=31QKDVGAHN0QHQE1GHOSKHWATMY32JVN?printableArticle=true [accessed 15th October 2009]

[6] Qualcomm press announcement http://www.umts-forum.org/content/view/1342/81/ [accessed 7th December 2009]

5 CURRENT DEVICE COST VERSUS SELECTIVITY

In order to improve receiver selectivity the cost of the RF receiver is likely to increase due to the increased development effort and larger die size required. Assuming the additional costs are passed through the supply chain from the semiconductor vendor, through the product manufacturers to the consumer, the retail price of the product will increase.

Figure 5-1: Typical consumer electronics supply chain

In order to determine how the product cost to the consumer will change with RF selectivity three factors need to be taken into account. These are:

- The influence additional receiver selectivity has on the cost of the RF receiver.
- How closely related the production cost of the receiver is to the price the semiconductor vendor sells it to the product manufacturer for.
- The influence the price the product manufacturer pays for the RF receiver has on the product's retail price.

5.1 RF receiver cost's influence on the product's retail price

The influence the RF receiver cost has on the final product's retail price will depend on the cost of the other components needed to make a product. For instance in a set top box the RF tuner cost has a very significant bearing on the overall product cost. However for an integrated digital LCD TV the influence the same tuner would have on the overall product cost is very much less as many more costly components, such as the LCD panel, are required.

Device retail prices are given below based on typical high street and internet prices. Ex-factory prices have been estimated based on the retail prices.

Product	**Typical Retail Price**[1]	**Estimated ex-factory price**[2]	**Notes**
DAB Digital Radio	£40	$28	
DVB-T Set Top Box	£20	$14	
Basic GSM Mobile phone	£20	$14	SIM free[3]
UMTS feature phone	£120	$85	SIM free[3]

Table 5-1: Typical products retail and estimated ex-factory price

Notes

1. Low end devices have been selected as the RF receiver cost will more heavily influence the retail cost than a more expensive device with more features
2. Allowance for VAT (15%), retailer margin, brand margin, and dollar to pound exchange rate (~1.6)
3. SIM free phone used as mobile operators significantly influence UK device costs

RF receiver volume costs, derived from a number of sources, have been used to estimate the influence the RF receiver has, as a proportion of the overall product cost.

Product	RF Receiver	Estimated Cost of RF receiver (or transceiver)	Percentage of ex-factory price
DAB Digital Radio	DAB RF IC	$1.80	6.5%
DVB-T Set Top Box	TV Tuner (canned tuner or Silicon)	$2	14%
Basic GSM Mobile phone	GSM transceiver RFIC	$1.80	13%
UMTS feature phone	UMTS/GSM transceiver RFIC	$2.80	3.3%

Table 5-2: Estimated RF receiver actual cost and percentage of ex-factory costs

Retail prices within one product category are generally directly proportional to the ex- factory price. This is seen especially at the lower end of the market where there is little room for large 'mark ups' for premium branding or styling so any changes in ex-factory prices will be directly reflected in the retail price.

The relationship between RF receiver cost to the product manufacturer and the retail price are shown in the graphs in figure 5-2.

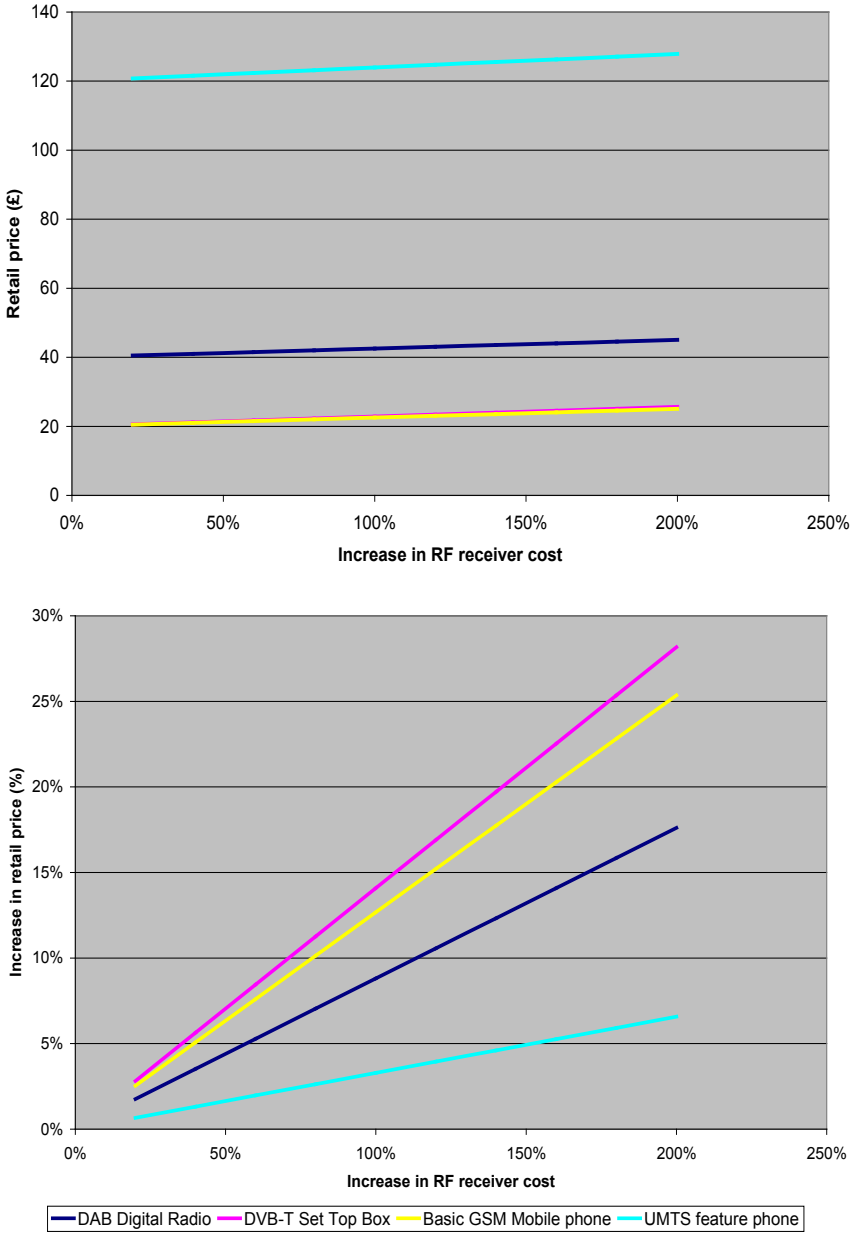

Figure 5-2: Product cost versus RF receiver cost to the product manufacturer

For example if the cost of a TV tuner IC to the product manufacturer increases by 100% the retail price can be expected to change by 14% from £20 to £22.80.

5.2 Additional receiver selectivity's influence on RF receiver production cost

As noted in chapter one there is no one single definition of receiver selectivity. To estimate the influence additional receiver selectivity will have on RF receiver production cost in quantifiable way the concept of improving the ADC's resolution has been used. With appropriate receiver design this additional ADC dynamic range can be used to obtain better selectivity through improved digital filtering.

By estimating the additional costs of incorporating a higher resolution ADC, along with the modifications to the receiver needed to realise the benefits of the additional ADC resolution, a graph of additional ADC resolution versus additional RF receiver cost can be constructed. For high volume products, using RF receivers almost entirely integrated into a single IC, the silicon vendor's RF receiver production cost is closely linked to silicon area. Therefore by estimating the change in silicon area, assuming the same silicon process is used, the change in receiver production cost can be estimated.

5.2.1 Improved ADC resolution

One additional ADC bit of resolution, assuming the same sample rate, will mean the ADC's quantisation noise is reduced by 6dB allowing the ADC to have up to 6dB additional dynamic range.

Throughout the analysis the same ADC sample rate and full scale amplitude level of the ADC will be assumed. As there are additional benefits in oversampling as much as possible, allowing a simpler anti alias filter to be used, it has been assumed that the ADC sample rate is probably set as high as it can reasonably be for the chosen silicon process. Therefore to gain additional resolution, additional ADC bits are required. This requires an ADC taking a larger silicon area. For any ADC, the full scale amplitude of the ADC is likely to be set close to the maximum voltage which can be tolerated in the silicon process as this minimises the effects of thermal and circuit noise on the ADC performance.

Figure 5-3 shows the relationship between silicon area and the effective number of bits (ENOB) for ADCs which may be used in a wideband receiver such as a digital TV tuner, Wi-Fi or LTE receiver. The parts selected have a range of architectures including the two architectures of choice at this sampling rate, pipeline and sigma delta. The data has been derived from a silicon IP vendor's website[7] and from ISSCC 2007[8] and 2008 papers. They have all been implemented in a 90nm CMOS process and characterised with a channel bandwidth of 10MHz.

[7] S3, parts SAD80M10BC90 and SAD80M13BC90, http://www.s3group.com/silicon-ip/view-all/ [accessed 31st July, 2009]

[8] Sotir Ouzounov et al, "A 1.2V, 121-Mode Continuous -Time ΣΔ Modulator for Wireless Receivers in 90nm CMOS", NXP Semiconductors, Eindhoven University of Technology, ISSCC 2007

Dan Huber et al, "A 10b 160MS/s 84mW 1V Sub-ranging ADC in 90nm CMOS" University of California, ISSCC 2007

Masato Yoshioka et al, "25.1: A 0.8V 10b 80MS/s 6.5mW Pipelined ADC with Regulated Overdrive Voltage Biasing" Fujitsu Laboratories Ltd, ISSCC 2007

G. Van der Plas, "A 150MS/s 133µW 7b ADC in 90nm Digital CMOS using a Comparator Based Asynchronous Binary Search sub-ADC", IMEC, ISSCC 2007

Pukar Malla, "A 28mW Spectrum Sensing reconfigurable 20MHz 72dB-SNR 70dB_SNDR DT ΔΣ ADC for 902.11n/WiMAX Receivers", Intel Corporation, Cornell University, Georgia Institute of Technology, ISSCC 2007

Young-Deuk Jeon et al, "A 4.7mW 0.32mm^2 10b 30MS/s Pipelined ADC Without a Front-End S/H in 90nm CMOS", ETRI, Korea ISSCC 2007

Jan Craninckx et al, "A 65fJ/Conversion-Step, 0-to-50MS/s 0-to-0.7mW 9bit Charge-Sharing SAR ADC in 90nm Digital CMOS", IMEC, Belgium, ISCC 2007

V. Giannini et al, "An 820µW 9b 40MS/s Noise Tolerant Dynamic SAR ADC in 90nmDigital CMOS", IMEC, Belgium, University of Salento, Italy ISCC 2008

Michiel van Elzakker et al, "A 1.9µW 4.4fJ/Converison-Step 10b 1MS/s charge-Redistribution ADC", University of Twente, ISCC 2008

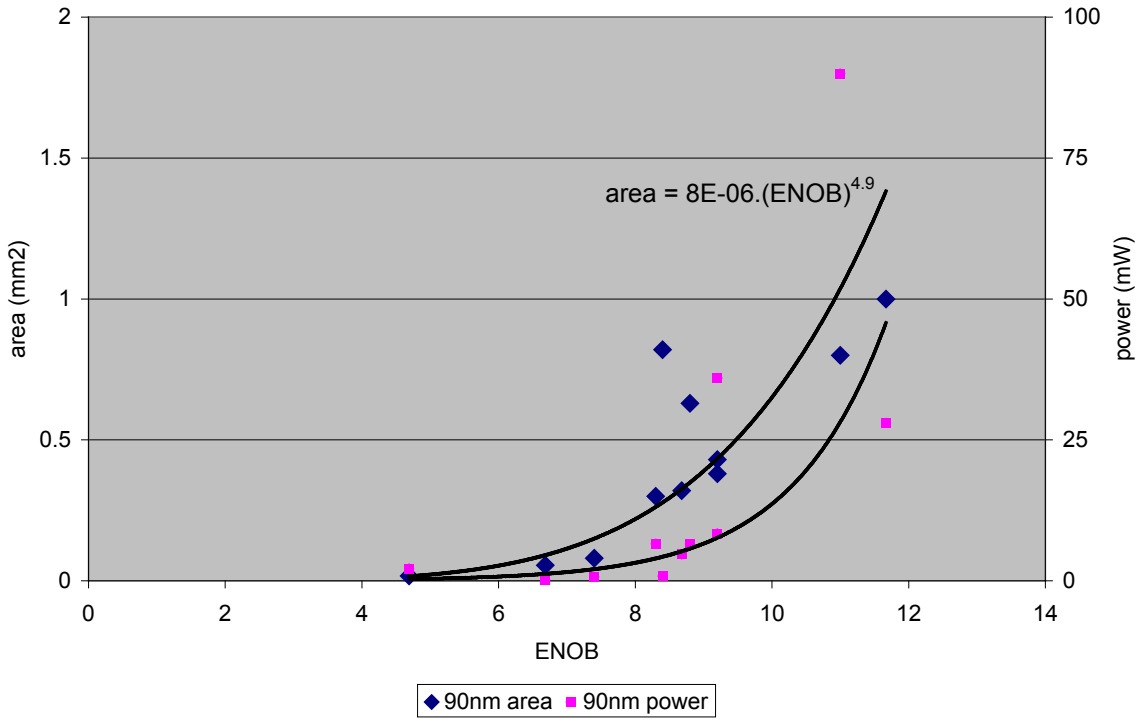

Figure 5-3: Silicon area and power consumption versus ENOB

From the graph it can be determined that obtaining additional effective bits of resolution take significant additional silicon area and power. The convertors displayed have generally been designed primarily for low power portable applications with silicon area, to some extent, a secondary limitation. There is a cluster of convertors with approximately nine effective bits as this is the resolution typically required in many current wireless systems. The plotted trend lines can be used to determine the change in silicon area as the number of bits changed. It should be noted that 12 bits of resolution appears to be the effective current limit of CMOS technology with 10MHz of bandwidth.

If the requirement for low power consumption and low cost is relaxed, allowing alternative silicon processes to be used, greater resolution can be obtained. For example ADCs with13 bits of resolution, typically taking 500mW, are often used in GSM base station applications. If a GSM base-station was not to down convert its signal first, it is predicted that 15 bits of ADC resolution at 1GHz, consuming 15W of power, would be needed.

5.2.2 Other receiver changes required

Other receiver changes which may be required to gain the benefit of the additional ADC resolution are discussed below.

5.2.2.i Anti alias filter

An ADC needs an anti alias filter to remove any energy in the alias bands which if sampled would appear in the digital replica of the analogue signal. Any energy in the first alias frequency band will be aliased directly into the wanted frequency band, where it could fall directly onto the wanted signal. If this happens filtering cannot be used to remove the unwanted signal degrading the potential SNR obtainable.

Figure 5-4 shows the effects of aliasing. Assuming the receiver noise density and gain is constant with frequency, to avoid noise in the alias band affecting the ADC performance, the anti-alias filter must reject signals at F_w-F_s and above by, at least, the dynamic range of the ADC. If this criterion is met, any signals at the full amplitude of the system reaching the anti alias filter will not affect the ADC performance.

In practice, especially if a low sampling rate is used, this criterion may not always be met. In this case the receiver performance for interfering signals around f_s will be defined by the rejection of the anti aliasing filter.

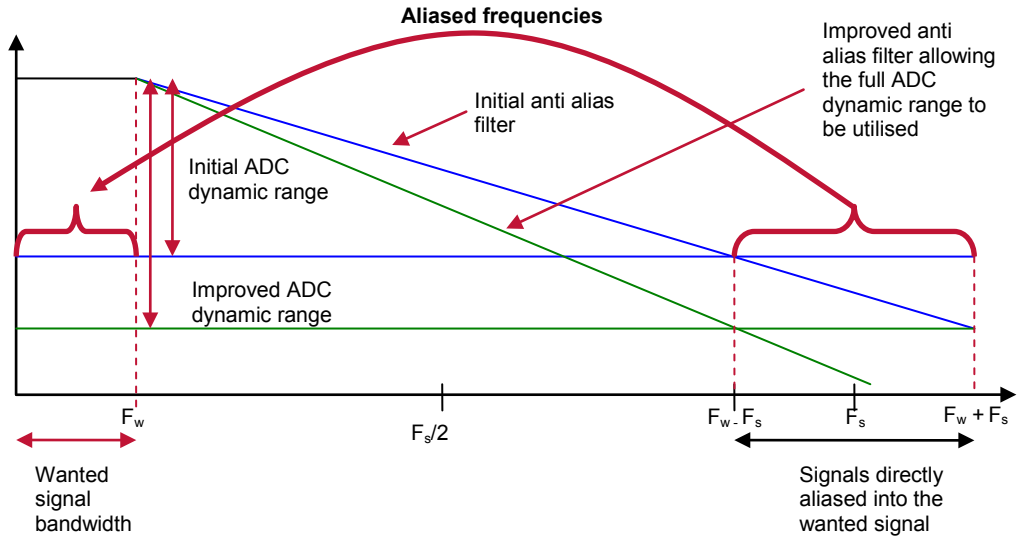

Figure 5-4: Anti-alias filtering effects

For each additional bit of ADC resolution, the ADC noise floor will be lower by 6dB, so 6dB more effective filtering will be needed at the first alias frequency (F_w-F_s). This could be implemented by using more poles of analogue filtering or possibly by lowering the filter's cut of frequency. As there is a trend for using high sampling rate sigma to delta convertors, with sample rates of well over a 100MHz allowing the anti alias filter to be implemented as a simple RC type filter, this is not expected to take significant additional die area.

5.2.2.ii Digital Filtering

For the additional ADC resolution to be used to gain better receiver selectivity, additional processing precision of the digital filters will be required. High speed digital filters are generally implemented using custom logic. In deep submicron process the area the additional gates required is expected to be small compared to the total die area.

5.2.2.iii Front end receiver changes needed for better large signal handling

The large signal handing performance of a receiver is largely defined by its third order intercept point. Intermodulation products may be caused by cross modulation of a single adjacent channel interferer or by multiple signals intermodulating. Figure 5-5 shows the improvements to the receiver's input third order intercept required to make full use of the ADC's additional dynamic range under different input signal conditions

Many receiver designs use AGC to help deal with large interferers as shown in case 1 below. However with reduced gain the noise figure of the receiver tends to increase. This can only be rectified with a receiver with improved large signal handling. As in case 1, an IP3 improvement of +9dB will allow a 6dB improvement in large signal handling.

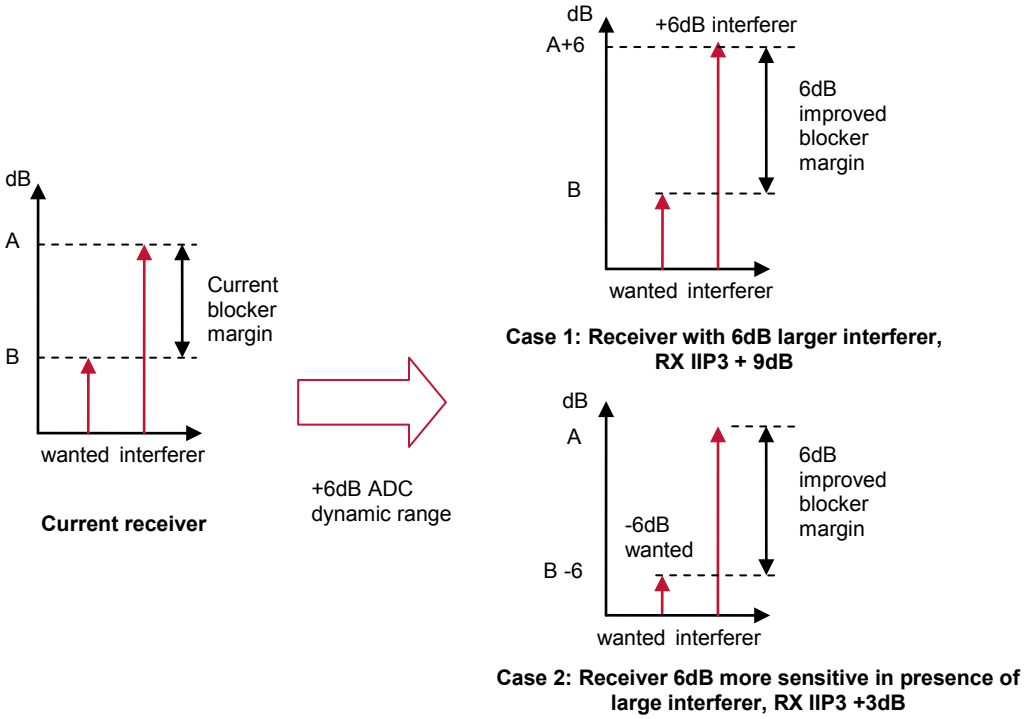

Figure 5-5: Affect of large signal handling on a receivers required IIP3

To obtain a 9dB better receiver third order input intercept point, considerably better LNA, mixer and amplifier large signal performance will be required. Steps which can be taken include:

- Increasing the stage current.
- Increasing the supply voltage.
- Using techniques such as emitter degeneration to linearise the stage. This is likely to degrade the noise figure of the stage slightly.
- An LNA's IIP3 is proportional to Q^2 of the LNA. By reducing the Q of the LNA, making it wider band, IIP3 can be improved. This is likely to degrade the noise figure of the LNA.
- Using negative feedback to linearise the stage. This is likely to degrade the noise figure of the stage slightly and increase power consumption.
- Using derivative superposition to cancel out some of the intermodulation products. This requires very well controlled silicon processes.
- Improving the noise figure of the ADC to reduce the gain required. At high sampling frequencies, the ADC noise figure is likely to be limited by kT/C sampling noise effects and clock jitter.

Some of the simpler steps are likely to have been already taken in the existing design. In battery power designs significantly increasing power consumption may not be acceptable. Whilst many techniques degrade noise figure the practical effects of this on receiver sensitivity are likely to be fairly small.

Deep submicron CMOS requires a low supply voltage of about 1V and has poor linearity. For significantly improved large signal handling it may be necessary to change the silicon process to one with a higher supply voltage with improved linearity. Supply voltage has a fairly direct effect on the size of signal a stage can handle. For example, with all else being equal, 10 dB larger signals can be handled with a silicon process which has a supply voltage allowing it to handle 3.3V peak to peak signals as to one which can handle 1V signals.

If a 3.3V supply voltage was required for large signal handling, a 350nm CMOS process would be needed and the transistor's f_T would be limited to less than 20GHz, making many circuit techniques unusable at the higher cellular and WLAN frequencies. Any digital circuits would be

very large. However by using a 180nm SiGe process, transistors with f_T of 78GHz could be used and digital circuit size would be more reasonable. This would still probably preclude the development of cost effective single chip radios.

None of the smaller steps are thought likely to increase actual silicon area significantly unless a change in silicon process is required. If a different silicon process is required the costs of producing the IC may change radically, especially if significant amount of digital circuitry is included.

5.2.2.iii.1 Benefits of additional ADC resolution

Assuming appropriate changes are made in both the analogue and digital parts of the receiver to allow the additional ADC resolution to be utilised, the benefits which can be gained are shown in figure 5-6. The bandpass filtering shown in the diagram is for a superhet receiver where an IF frequency is sub-sampled, for example a 'canned' TV tuner. With baseband sampling, used for example in low IF and zero IF receivers, low pass filtering would be used.

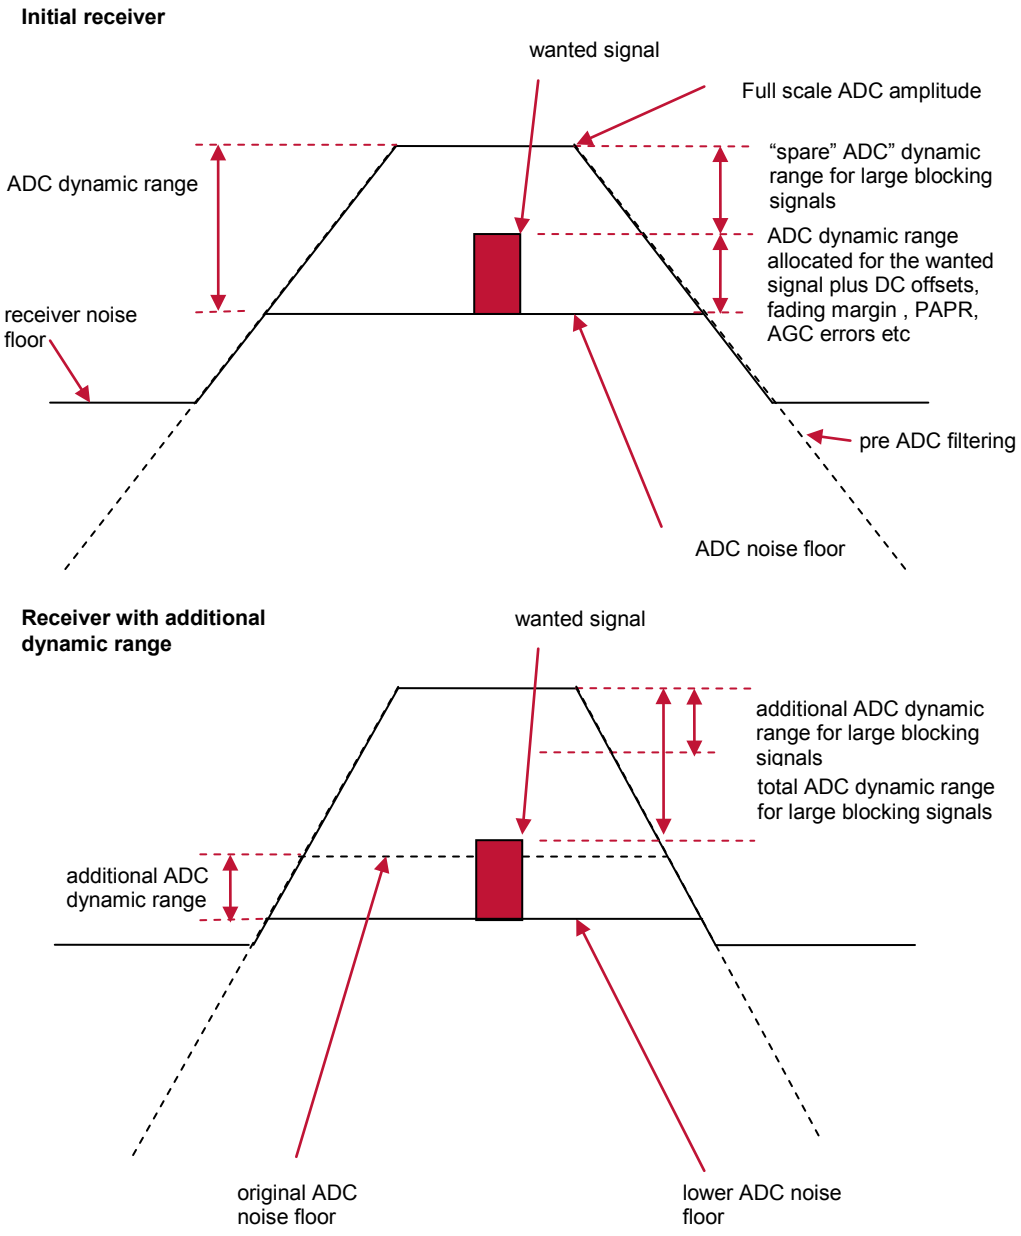

Figure 5-6: Effect of additional dynamic range

When the receiver is receiving fairly small signals and is therefore operating with high gain, the noise floor of the receiver will be low and defined largely by the 'front end' of the receiver. This should allow the full benefit of the additional ADC resolution to be realised.

Assuming the receiver dynamic range is not limited by the receiver's noise floor prior to sampling, the additional ADC resolution will allow greater receiver dynamic range. For each additional effective bit of ADC resolution:

- 6dB larger unwanted signals falling into the pass band of the pre ADC anti-alias/channel filter can be handled without the ADC overloading whilst still being able to decode the wanted signal. The unwanted signals sampled by the ADC can be filtered out by digital filtering prior to demodulation.
- Unwanted signals at a frequency above the corner frequency of the anti alias filter, but below the Nyquist bandwidth of the ADC will be attenuated by the anti alias filter. Therefore the receiver will be able to handle 6dB larger signals due to the ADC plus potentially up to 6dB larger signals due to the improved anti alias filtering. The additional improvement from the anti alias filter is only likely to be realised for signals close to the alias frequency.

.

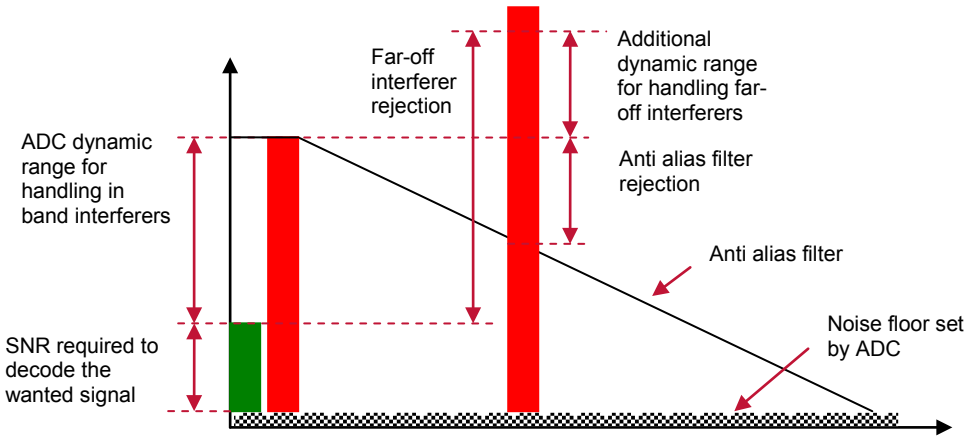

Figure 5-7: Selectivity with noise floor set by ADC

When the receiver is receiving larger signals and therefore has reduced gain, the receiver's noise floor will be raised. In this case the receiver's dynamic range is likely to be limited by the receiver's noise floor prior to sampling and also by the supply voltage. In this case additional ADC resolution will not allow greater receiver dynamic range.

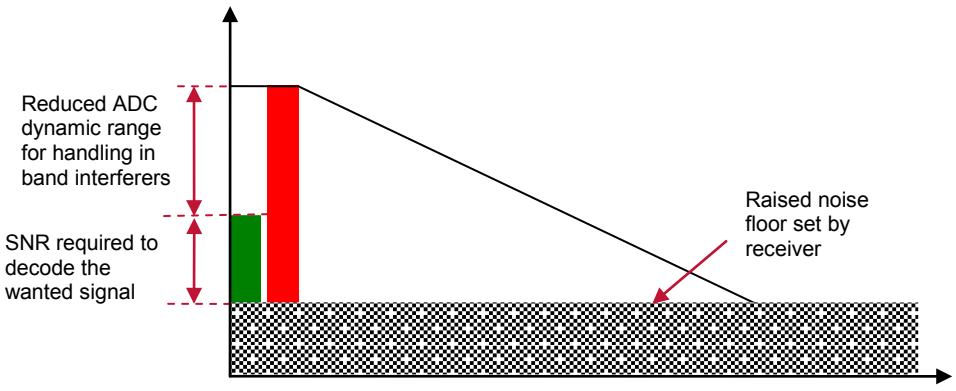

Figure 5-8: Selectivity by front end with raised noise floor

When the receiver is receiving large signals and is therefore operating with little gain, the noise floor of the receiver will be defined largely by the ‘back end’ of the receiver and is likely to be high. In this case, for effective filtering of large signals, the filter must be located as close to the front end of the receiver as possible. This ensures the filter has the maximum possible SNR before being limited by the noise floor of the receiver.

The potential improvements gained from additional ADC bits are listed in table 5-3 for the various receiver architectures. In each case ‘small’ signal conditions are assumed, i.e. the wanted signal is near the sensitivity limit of the receiver.

Parameter	Superhet	Zero IF	Digital Low IF Weaver	Analogue low IF with poly-phase filter
Small signal adjacent channel rejection (n±1)	**Up to 6dB per bit** Assumes: 1. Analogue filtering does not fully suppress the adjacent channel to below the receiver's noise floor 2. ACR is not limited by ACPL from the transmitter	**Up to 6dB per bit** Assumes: 1. Analogue filtering does not fully suppress the adjacent channel to below the receiver's noise floor 2. ACR is not limited by ACPL from the transmitter	**Up to 6dB per bit** Will allow: 1. Additional resolution to allow better gain and phase calibration of the analogue down conversion stage 2. Additional resolution for digital image suppression The improved performance allows better C/N of the wanted in the presence of the adjacent channel interferer located close to the LO and more dynamic range for the alternate channel interferer located close to the LO. Adjacent and alternate channel rejection of interferers away from the LO are less of a problem Assumes: 1. ACR is not limited by ACPL from the transmitter 2. Alternate channel rejection is not limited by ACPL from the transmitter	**No effect** on the ACR of the adjacent channel closest to the LO **Up to 6dB per bit** improved ACR for the adjacent channel furthest from the LO. However this adjacent channel is generally less of a problem than the one closest to the LO
Small signal alternate channel rejection (n±2)	**None** Assumes: 1. Analogue filtering suppresses the alternate channel to below the receiver's noise floor	**Up to 6dB per bit** Assumes: 1. Analogue filtering does not fully suppress the alternate channel to below the receiver's noise floor		**Up to 6dB per bit** Assumes: 1. The channel / anti-alias filter does not fully suppress the alternate channel to below the receiver's noise floor
Small signal far-off channel rejection (n±m)	**None** Assumes: 1. Analogue filtering suppresses the alternate channel to below the receiver's noise floor	**6dB per bit, plus some benefit due to improved anti-alias filtering** Assumes: 1. Analogue filtering does not fully suppress the unwanted channel to below the receiver's noise floor. 2. Aliasing is not an issue	**6dB per bit, plus some benefit due to improved anti-alias filtering** Assumes: 1. Analogue filtering does not fully suppress the unwanted channel to below the receiver's noise floor. 2. Aliasing is not an issue	**6dB per bit, plus some benefit due to improved anti-alias filtering** Assumes: 1. Analogue filtering does not fully suppress the unwanted channel to below the receiver's noise floor. 2. Aliasing is not an issue
Analogue domain - image rejection	**None** As the image falls directly on the wanted signal improved ADC resolution will not help	**Not applicable**	**Not applicable**	**None** As the image falls directly on the wanted signal improved ADC resolution will not help

Table 5-3: Benefits of additional ADC resolution for different receiver architectures

It can be seen that for commonly the most critical case, adjacent channel rejection, can provide up to 6dB additional rejection per ADC bit. For interferers further from the wanted, the analogue filtering prior to the ADC is likely to play a more important role, suppressing the interferer towards the receiver's noise floor limiting the practical benefits of improved ADC resolution.

For large signal conditions, receiver performance is limited by the relatively poor IP3 performance of the silicon process and the relatively high receiver noise figure. The only way of improving unwanted signal rejection is to implement filters at the front end of the receiver. These will have a relatively low Q and therefore will not help adjacent and possibly alternate channel rejection, but will significantly help far-off rejection.

Superhet receivers need good front end filtering for good image rejection. This filter plays a significant role in improving all far-off selectivity, especially with large signals. When compared to superhets, direct conversion and low IF receivers tend to not have a significant front end filter and therefore have reduced far-off selectivity.

Due to using high gain antennas mounted on roof tops to enhance the signal level compared to that found near ground level, TV tuners receive significantly stronger signals than those received by cellular or Wifi systems using relatively inefficient antennas positioned close to the ground. This allows TV broadcasters to provide adequate coverage with a relatively low number of transmitters compared to say a cellular network. Whilst the TV signal may be relatively low, requiring a high gain antenna for reception, the TV antenna, if it points towards the cellular base-station, will also amplify the cellular signal to a high level. If this situation occurs, the TV receiver will need much greater selectivity than a cellular receiver with a poor antenna.

To improve far of selectivity some of the newer 'silicon' TV tuners, using generally low IF architectures incorporate tracking filters. This allows the manufacturers to obtain similar or even better far-off selectivity performance than traditional canned tuners using superhet receivers. Both architectures rely on adjacent channel filtering implemented using a SAW filter followed by an ADC and digital filtering in the case of the superhet receiver and limited analogue channel filtering and an ADC and digital filtering in the case of the silicon tuner. Tracking filters don't help adjacent channel selectivity.

5.3 Receiver and product cost versus selectivity

For a given silicon process, silicon area is approximately proportional to the IC production cost. In trying to improve receiver selectivity the largest area factor is likely to be the improved ADC. Figure 5-9 charts, for 90nm CMOS, how the ADC's effective resolution affects ADC silicon area. The die area of the rest of the IC is not expected to rise significantly. By examination of die micrographs from recent papers covering a wide range of handheld receiver technologies including GSM[9], Bluetooth[10], TV[11] and 802.11, using similar CMOS processes, it can be seen that the ADC typically occupies between 20 and 50% of the total RF/analogue die area. The average figure has been used.

By combining these three parameters, the additional production cost of improved selectivity can be predicted. This is shown in the figure below for a nominal receiver with nine effective bits of resolution, typical for current implementations, costing a nominal $1. The improvement in selectivity has been limited to three additional bits, i.e. using a 12 bit ADC. As already noted, this is close to the limit for today's technology.

[9] Ed K Iniewski, "Wireless Technologies, Circuits, Systems and Devices," Chapter 10, page 268 by R.B. Staszewski, "Digital RF Processor (DRP™), CRC Press, 2008

[10] R B Staszewski et al, "Digital RF Processor (DRPTM) for Cellular phones", ICCAD-2005 Embedded Tutorial 2A.3 http://www2.iccad.com/data2/iccad/iccad_05acceptedpapers.nsf/9cfb1ebaaf59043587256a6a00031f78/e25be1f58da7d131872570530070afb2/$FILE/2A_3slides.PDF [28th August 2009]

[11] Supisa Lerstaveesin et al, "A 48–860 MHz CMOS Low-IF Direct-Conversion DTV Tuner", IEEE Journal Of Solid-State Circuits, September 2008.

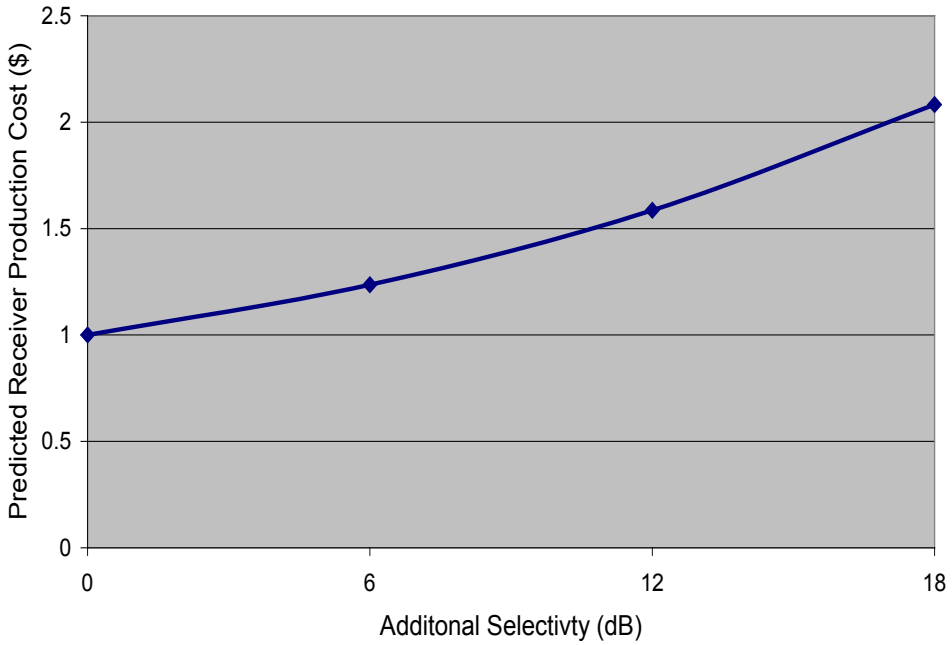

Figure 5-9: Additional selectivity effect on nominal receiver production cost

IC production costs, represented in a semiconductor vendor's annual report as 'cost of sales', is a relatively small portion of the costs a semiconductor vendor incurs. Typically a vendor's costs can be broken down as follows[12]:

Item	Portion of total costs incurred
Costs of sales, i.e. production costs	55%
Research and development	30%
Selling, marketing, general and administration	20%

Table 5-4: Typical semiconductors vendor's costs

The cost of sales figure of 55% of total costs for the semiconductor industry is very low compared to say other companies in the electronics value chain, such as contract manufacturers and retail, where costs of sales figures of up to 95% are common. This shows that the semiconductor vendor's production costs have less of an effect on their total revenues and profits compared to costs in other sectors. More important to a semiconductor vendor is securing 'design wins' into as many products as possible, as even a modest increase in sales can lead to a significant increase in profit.

We have assumed the price of the RFIC paid by the contract manufacturer increases in proportion to the RFIC production cost. With the low cost of sales this is perhaps a worst case assumption, from the consumer's point of view, as it allows the semiconductor vendor to maintain the same gross margin with increased turnover, giving them scope for increased net profits. It does allow the semiconductor vendor to finance the additional research and development the semiconductor vendor is likely to need to undertake to achieve the improved RF performance. Whether the contract manufacturer or consumer is prepared to pay the additional cost is questioned in section 5.4.

By combining the influence selectivity has on the nominal selling price with the influence the RF receiver selling price has on the overall product costs, the influence additional receiver selectivity has on overall product cost can be predicted. This is shown in figure 5-10.

[12] Derived from examination of the annual report of several major semiconductor vendors specialising in analogue and wireless

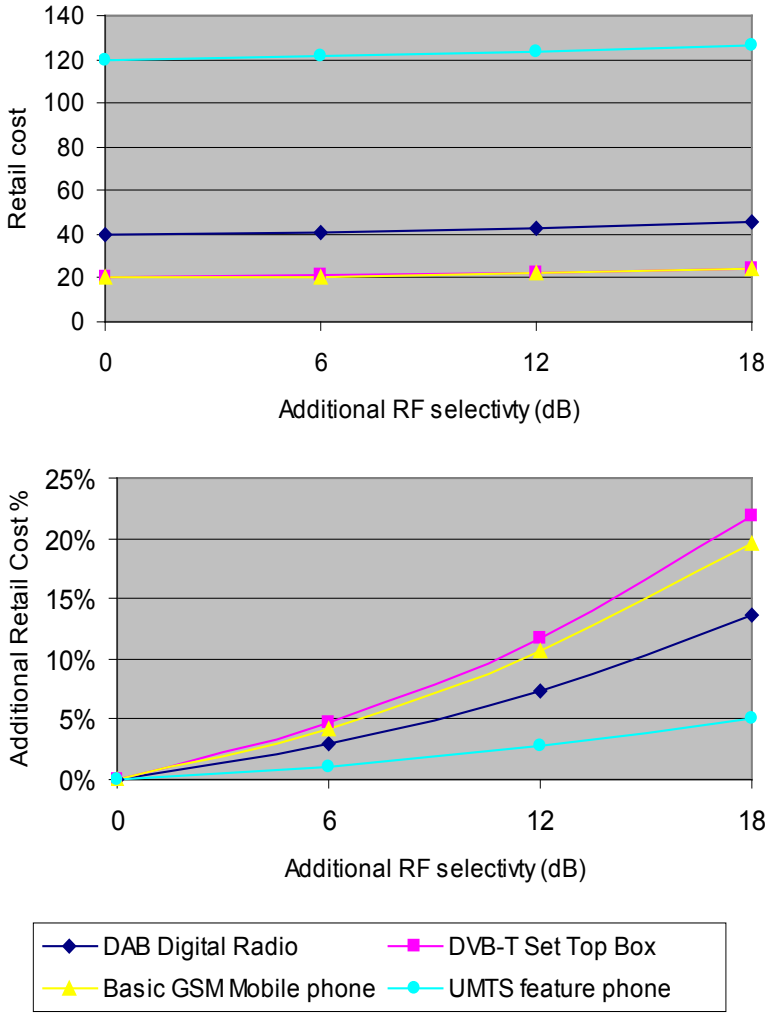

Figure 5-10: Product cost versus RF selectivity

It should be noted that this prediction is for:

- Highly integrated products with all channel filtering integrated into the IC. With this approach, multiple receiver bandwidths can be accommodated within a single design by using programmable filters. This lends itself to multi standard and world standard receivers allowing the largest possible product volumes for each design.

- Costs associated with increased power consumption are neglected, e.g. needing a larger battery. Power consumption will increase rapidly with improved selectivity. Although this has not been investigated in detail, initial indications suggest that power consumption could increase at a similar, or even faster rate, than the silicon area increase with resolution. In portable products, power consumption is potentially a more significant design driver than cost as this leads to reduced operation times, and/or larger, less attractive products as the battery size is increased. The higher selectivity levels could potentially lead to a commercially unacceptable product. Receivers, such as those used in set top boxes, are often left powered even when the unit is in standby. This is, for example, to allow EPG updates. These could lead to products with higher standby currents which may conflict with energy use guidelines and legislation.

5.4 Cost versus selectivity - the semiconductor industry view

To corroborate the findings in the report the results were discussed with a number of representatives from the semiconductor and communications industry.

A common view, best presented as an example by a representative of a major semiconductor company was:

"The price a semiconductor vendor can sell its products for is dictated by the contract manufacturer who has to meet a very strictly defined price. In practice, the price of the part is almost the only thing that matters to the contract manufacturer "

This price is dictated by the acceptable high street price. As seen in the report typical contract manufacturers and electronic products retailers' profits represent a very small part of their turnover. By placing a strict cost target on the contract manufacturer the brand or retailer can maximise their profits.

"….Given a sufficiently large opportunity in the semiconductor vendor's market area, the semiconductor vendor will develop a product for the market, almost irrespective of the requirements. An example is the US analogue TV switch off where the ATSC US digital TV standard receiver needs to meet the fairly stringent A/74 standard"

This standard has a far-off (>±6channels) selectivity protection ratio of 57dB at low levels, compared to say the MBRAI limit of 42dB, and has fairly stringent large signal handling requirement.

"…*Several silicon vendors have developed products for this market and they have been produced at a price which allows digital adaptors to be "given away" under the US government's $40 coupon scheme".*

Whilst this price is significantly less than that widely predicted only a couple of years ago for these digital adaptors, it should be noted that this price is still significantly more than an equivalent basic DVB-T set top box, which may retail at around $28 in the US, once differences in sales tax have been taken into account. Whether the cost difference is due to the increased RF specification, demodulator difference between ATSC and DVB-T, or the government coupon making it unnecessary to drop the price further is difficult to predict.

"…the cost of the ATSC tuner IC is little more than a lower specification DVB-T tuner…"

"the higher specification product will take longer to develop; this was seen as delays in the ATSC programme".

This view is probably backed by typical semiconductor vendors cost breakdowns as shown in table 5-4. The fairly low cost of sales means that increasing sales are more important than profit margins per unit. In order to maintain or increase sales volumes, the semiconductor vendor must develop new markets where possible and to try to have a superior product to help gain design wins over a competitor. This has led to very high research and development expenditure within the industry and has helped lead to the rapid growth in the performance of semiconductors whilst also allowing product prices to drop.

Putting aside short term price increases due to shortage of supply, for semiconductor prices to increase in the long term, something which rarely, if ever happens; there must be an increase in performance as perceived by the customer. A superior product may be one with improved consumer perceived performance, for instance a more powerful processor, or one which allows the product manufacturer to lower his overall costs.

For RFIC vendors the approach is usually to incorporate additional functionality into the receiver and incorporate "additional services" into the sale. This is usually done by:

- Reducing the number of additional components, for example filters, inductors and capacitors, needed for the receiver to operate allowing the semiconductor vendor to be

paid a larger proportion of the overall Bill of Material cost. In some cases this reduces the complexity of the PCB the components are mounted on, also lowering product cost.

- Reducing the production cost by reducing the assembly and test time costs. Assembly time can be reduced by reducing the number of components, test costs can be reduced by incorporating self test algorithms in the device.

- Reducing the amount of design work the product manufacturer has to undertake by making the RFIC as simple to implement as possible, supplying embedded software, reference designs and test suites.

Where the applicable selectivity standard is seen as adequate by the industry, RF receiver selectivity performance, beyond that required by the applicable standard, is not perceived as important. This leads the semiconductor vendor to design into their receiver potentially just enough selectivity to meet any mandated standards.

TV and radio receiver are perhaps slightly different as there is no mandatory minimum standard defined although in the UK manufacturers can choice to obtain accreditations such as the digital tick logo administered by the DTG which do require some selectivity conformance testing. In practice it is believed that many of the premium brands will require selectivity performance significantly exceeding that needed for the DTG tests.

This topic was discussed at a workshop attended by over 20 representatives from the industry held at TTP on the 1st December 2009.

It was recognised in the workshop that selectivity is a complex area which is often not well understood by the consumer. In practice, whilst the consumer expects the device to work well, if a device does suffer interference, the consumer may “blame” the network or atmospheric conditions. In this context, it is difficult for the product manufacturer to use a positive marketing message to justify additional cost for selectivity.

However the consensus at the workshop was that the premium manufacturers would require additional selectivity beyond that required in any standards to help ensure the product works in as many situations as possible, or perhaps as a minimum, at least in as many situations as their competitor’s product would work in. This is done in order to minimise poor consumer experiences and product returns, in turn protecting their brand value and building brand loyalty.

A premium product brand would typically implement this wish by placing a requirement on their tuner supplier (canned or silicon) to deliver a near state of the art tuner whilst also expecting it to be cost effective. Silicon vendors have responded to this. For example, Silicon Laboratories, who presented an overview of their silicon tuner at the workshop, have developed silicon tuners with a performance comparable to or better than conventional MOPLL can tuners. This is, according to Silicon Labs, to allow them to target the “*high-end OEM iDTV market, not limited to low end market by performance, e.g. PC, portable, cheap set-top box*”.

Even at the “low end” of the market the workshop recognised that retailers valued reliability as this avoided customer returns. However within this segment it was recognised that consumer electronics margins are small and any additional costs were avoided.

The workshop recognised that self incentivisation within an industry sector for improving receiver selectivity was difficult when the requirement for receiver improvement was brought about by new interferers from other new services occupying spectrum previously unused, or previously used in a way that did not cause interference to the receiver. Ways of the beneficiary of the new spectrum compensating the victim of the interferer were discussed, but no ways were identified of making this approach practically work.

It was acknowledged that the introduction of a new service could lead to a large number of existing users suffering interference. Whilst receiver manufacturer could perhaps be encouraged to improve new receivers, there would always be a very significant legacy issue as not all users would purchase new equipment. This effect could be reduced by introducing

enhancements to the existing service, for example by moving from DVB-T to DVB-T2 which would encourage users to purchase new equipment.

The workshop recognised that interference issues occur mainly at the boundary between different spectrum users. Therefore limiting the number of boundaries limits the interference issues. For instance it was questioned why the DSO spectrum had been split into two blocks creating four boundaries rather than one block with two boundaries.

The use of guard bands between users was discussed. Guard bands were seen as an effective but inefficient way of protecting one service from another. It was suggested the guard bands could be released as unprotected ISM spectrum.

5.5 Conclusions

The conclusions of this chapter, “device cost versus selectivity” are:

- Increasing receiver selectivity will lead to increased semiconductor silicon area, production cost, and research and development costs.
- In systems where there is a clearly defined applicable standard which is seen as adequate by the industry, for example cellular, RF receiver selectivity performance beyond that required by the applicable standard the receiver is working to is generally not perceived as important by the manufacturer or retailer.
- In systems where the applicable standard is not seen as adequate by parts of the industry, for example TV, premium brands may require their supplier to supply tuners with additional selectivity in order to provide as good a customer experience as possible whilst at the “low end” of the market limited margins forced manufacturers to avoid any additional cost.
- The high gross margins in the semiconductor industry may allow, if required to maintain or grow sales volume, the semiconductor vendor to absorb the additional costs of limited selectivity improvements.
- The continual improvement of silicon processes and receiver design techniques are thought likely to allow any additional costs to be ‘engineered around’, limiting any long term increase in semiconductor production costs.
- Introducing new services often places a requirement for the receivers of existing spectrum users operating near the new service to have improved selectivity. Even if new receivers have improved selectivity, this will not overcome the issue with legacy receivers.

6 ANTICIPATED POTENTIAL FUTURE RECEIVER PERFORMANCE

6.1 Drivers and trends

In anticipating the future architectures and level of integration for various applications it is worth looking backwards to identify the trends in the relevant sectors in the last ten years and using these to help to predict what might happen in the next ten years.

6.1.1 Cellular

Ten years ago GSM was well established with UMTS in early development. Currently UMTS is established with the plans for LTE being put in place. In 2019, LTE is likely to be well established. GSM is not predicted to be phased out by this date although it may well have less bandwidth allocated to it. It has been suggested that LTE may replace UMTS in the future. LTE's specifications demand similar receiver selectivity as UMTS.

In more expensive "smart phone" devices, increased functionality is being seen, primarily through more complex software and a larger display, but also with the introduction of Wi-Fi, GPS and possibly near field communication systems into the terminal. This trend is expected to continue. There will also continue to be a market, especially in developing economies, for basic phones with voice, SMS capability and simple software.

Historically most cellular systems have used FDD. In FDD all base-station transmitters transmit at one fixed set of frequencies for the downlink (base-station to mobile); and all mobile transmitters operator at another fixed set of frequencies for the uplink (mobile to base-station). For adjacent channel interference to occur in FDD, the interferer and victim must be of different equipment types, i.e. one is a mobile, whilst the other is a base-station.

Assuming base-stations always have some physical separation from mobiles; receivers can be used with a relatively low adjacent channel performance without limiting system performance. With the introduction of the femtocell (small low power cellular base stations) some cellular infrastructure equipment could be installed in the home and office, rather than on towers and building roofs, introducing new interference sources possibly very closely situated to other receivers.

In the future more cellular systems may use TDD. In TDD, time is used to divide the transmit and receive signals. This can bring spectrum efficiency benefits if there are unequal amounts of traffic in the uplink and down link. This is increasingly likely as cellular networks become more data centric. In addition TDD doesn't require 'paired' spectrum. However TDD systems can bring physically close together transmitters and receivers simultaneously operating in adjacent channels. Therefore TDD receivers are likely to require much higher levels of adjacent channel rejection than current receivers.

The table below highlights trends over the last ten years and predicts trends for the next ten years. Integration trends are shown pictorially in figure 6-1.

	Time		
	1999	2009	2019 (anticipated)
Standard	GSM	GSM, UMTS	GSM, UMTS, LTE
Architecture	Superhet with discrete hardware and SW channel filtering	Zero IF, Low IF	Discrete time, highly software defined, MIMO transceivers incorporated into some devices
Integration	Dual band with separate RF transceiver, mixed signal device, and baseband. Ancillary items included a separate VCTCXO, SAW IF filter, discrete VCO, LNAs and several voltage regulators. A large number of passive components were used. Multilayer PCB technology	Multi band (5 to 7) RF IC, with integrated PMU. The few external components include a crystal and a few passive components. The front end filter is often integrated into a FEM. There is increasing integration with baseband and multimedia processor[13] Multilayer PCB with micro-via technology.	Multi band (many), with a high level of integration with the cellular baseband and/or possibly the RF power amplifier. Possible integration with other wireless connectivity such as Bluetooth, FM radio and little or no fixed frequency filtering.
Si Process	BiCMOS	90nm CMOS, possibly SiGe	16nm CMOS[14]

Table 6-1: Cellular trends, 1999 to 2019

[13] Qualcomm press announcement http://www.umts-forum.org/content/view/1342/81/ [accessed 7th December 2009],
Ed K Iniewski, "Wireless Technologies, Circuits, Systems and Devices," Chapter 10 by R.B. Staszewski, "Digital RF Processor (DRP™), CRC Press, 2008

[14] ITRS, "Radio Frequency And Analog/Mixed-Signal Technologies For Wireless Communications", ITRS2007, http://www.itrs.net/Links/2007ITRS/2007_Chapters/2007_Wireless.pdf [accessed 8th December 2009

1800MHz GSM phone, ~1999

900 and 1800MHz low cost GSM phone, 2009

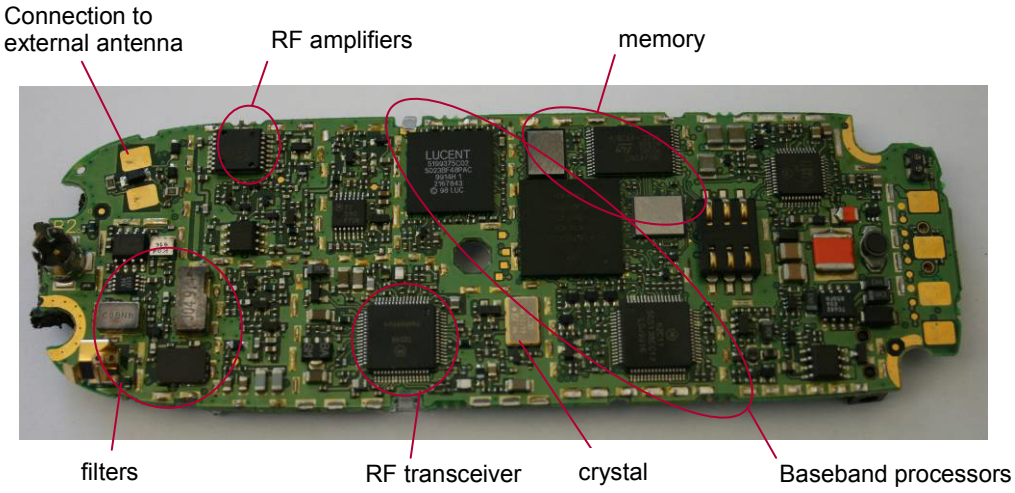

1800MHz GSM phone, ~1999

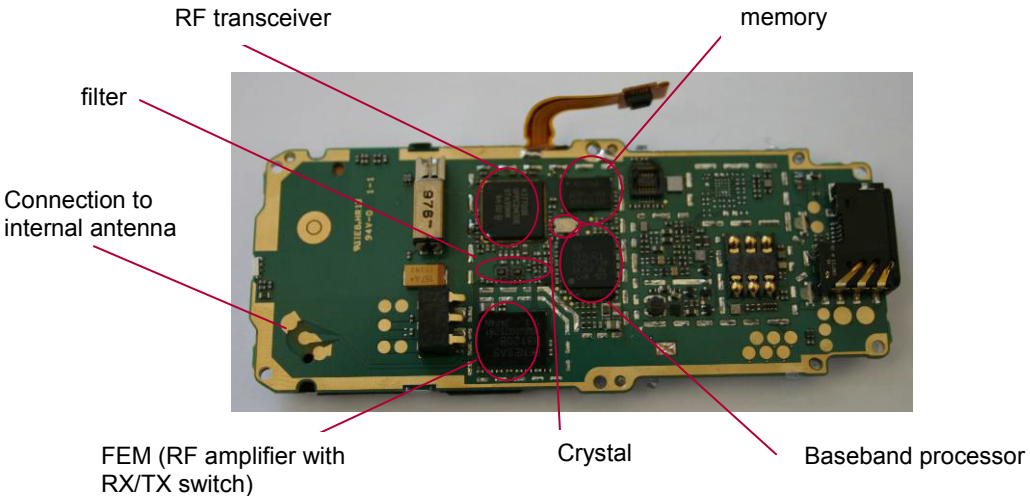

900 and 1800MHz low cost GSM phone, 2009

Figure 6-1: Cellular integration, 1999 to 2009

6.1.2 Broadcast

Terrestrial TV tuners have been used as an example in this section whilst broadcast radio receivers have some similarities. Two approaches to tuner implementation are currently common; silicon tuners and canned tuners. It is anticipated[15] that once analogue services are switched off silicon tuners will rapidly dominate. This is likely to be due to silicon tuners with outputs suitable for analogue TV signals being more complex, and therefore more costly than silicon tuners with outputs suitable for digital TV. Additional factors include:

- Although the cost of silicon tuners is currently more than canned tuners, the cost of digital tuners is dropping and is anticipated to reach the price of canned tuners in 2010 or 2011, a similar time frame to analogue switch of in many countries.

- Digital tuners are much smaller than canned tuners. This makes them more suitable for integrating into slim-line flat panel TVs and into compact digital video recorder with multiple tuners.

It has been suggested that terrestrial TV may be replaced with broadband[16] and possibly satellite, freeing up spectrum for other applications although no dates have been set for this. This is unlikely to happen by 2019 but thought possible by 2029[17].

	Time			
	1999	2009 Canned-tuner	2009 Silicon tuner	2019 (anticipated)
Standard	analogue (PAL)	Multi standard analogue and digital	DVB-T, occasionally hybrid PAL/DVB-T	DVB-T, DVB-T2
Architecture	Superhet	Superhet	Zero IF, Low IF	Zero IF, Low IF, discrete time, software defined
Integration	Canned tuner implemented on single sided PCB with separate RF and IF devices, discrete LNAs, varactors and hand adjusted inductors	Similar to 1999, although single RF IC now used, + LNAs, multiple SAW filters to cope with different standards	Single chip external inductor for tuning front end filter, crystal	Single chip integrated with baseband
Si Process	Bipolar	Bipolar (lowest cost)	SiGe or CMOS, Sony[18] suggest CMOS geometry will rapidly reduce from 250nm to less than 90nm by the beginning of 2011.	16nm CMOS[19]

Table 6-2: TV tuner trends, 1999 to 2019

[15] Nikkei Electronics Asia, " Expanded Use of Silicon Tuners Transforms TVs" 5th June 2009, http://techon.nikkeibp.co.jp/article/HONSHI/20090526/170774/ [accessed 8th December 2009]
[16] For example, Williams and Marks, "A Framework for Evaluating the Value of Next Generation Broadband" A report for the Broadband Stakeholder Group, Plum Consulting, June 2008
[17] For example Lewin et al, "Entertainment in the UK in 2028" A report for Ofcom, Plum Consulting, February 2008
[18] Sony, "CMOS Silicon Tuners for Large-Screen TV sets", http://www.sony.net/Products/SC-HP/cx_news/vol56/pdf/featuring56.pdf [accessed 8th December 2009]
[19] Predicted by ITRS, ITRS2007

6.1.3 Other applications

In 1999 there was a wide range of mainly fairly low volume RF applications including wireless LAN, PMSE, wireless telemetry and a variety of short range wireless applications. Many of these applications used specific proprietary protocols. Analogue and digital (DECT) cordless telephones were one of the few mass market applications using open standards. Standards such as Bluetooth and 802.11 were in their infancy with no significant product volumes. By today's standards, integration of devices was fairly low with a number of ICs with parameters such as the frequency range the radio could operate over set by external components. This allowed, using commonly available components, receiver designs for specific non standard applications to be implemented at almost any frequency.

By 2009 technologies such as 802.11 and Bluetooth have become mass market with Bluetooth integrated in all but the lowest cost mobile phones and both technologies integrated into many laptop computers.

Many new short range wireless applications have been developed utilising standards such as Zigbee operating in a wide range of ISM bands. Initially these used RFIC specifically developed for use in the ISM bands but capable of supporting multiple protocols. Increasingly these RFICs are being integrated with processors specifically for single standards such as Zigbee. Whilst it is still possible to cost effectively produce products using non standard protocols, the RF ICs have their operating frequency range set by the IC itself restricting their operating range to the ISM bands. Many of the components commonly available ten years ago, which allowed designs to be implemented at almost any frequency, are now not available. This makes designs at other 'non standard' frequencies such as PMSE much more difficult and costly to implement than 'standard' applications.

In the future it can be anticipated that increasingly applications will be forced down a standards driven route. Fortunately the standards are becoming more flexible allowing them to be used for a wider range of applications. Future versions of mass market technologies such as Bluetooth and Wi-Fi are likely to be integrated into a wider range of products. The distinction between Wi-Fi, and cellular are likely to blur with the introduction of standards such as WiMAX intended to operate in licensed spectrum.

	Time			
	1999	2009 - high volume	2009 - low volume	2019 – high volume (anticipated)
Standard	DECT, proprietary	DECT, 802.11, Bluetooth…	Proprietary, TETRA	DECT, 802.11, Bluetooth…
Architecture	Superhet, super regenerative for low cost short range receivers	Zero IF, low IF,	Superhet	Discrete time, highly software defined
Integration	Low density RFICs and transistors	High, often combined with multiple wireless interfaces such as GPS and FM radio	Fairly discrete	very high, combined with multiple wireless interfaces and possible cellular transceivers
Si Process	Bipolar	CMOS	BiCMOS	16nm CMOS[20]

Table 6-3: Other applications trends, 1999 to 2019

[20] Predicted by ITRS, ITRS2007

6.1.4 Semiconductor development cost

Over the last decade we have seen spectacular improvements in the price and performance of electronic products. Whilst products prices have dropped performance has increased. This is largely due to advances in semiconductors and has been obtained through semiconductor vendors spending large portions of their revenue on research and development.

With a reduction in the semiconductor node size, IC mask costs and engineering development (NRE) have increased radically as shown in figure 6-2. In order to justify undertaking a design the number of units of the IC which need to be sold has increased. This number has been estimated[21] to be in the region of 50 million units for complex designs in the smaller geometries. Therefore to justify the development, assuming a design has a competitive life of say two years and there are four major competitors in the market area, a total market size of around 100 million units per year is required.

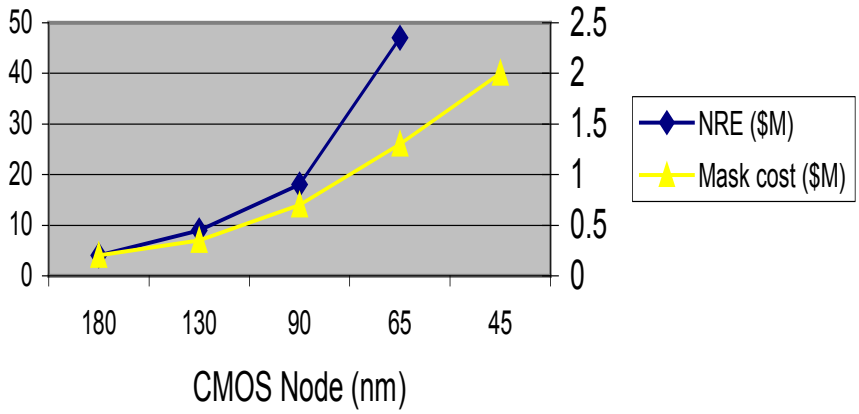

Figure 6-2: NRE and mask costs versus CMOS node[22][23]

The investment and sales required suggests that this type of development can only be undertaken by large companies operating in very large markets. This can be used to explain why there are currently fewer “design starts” in the first year of a new node becoming available compared to the historic introduction of new nodes and why, increasingly, lower volume designs with fewer gates are being implemented in less “leading edge” CMOS geometries.

The increased development costs have historically been offset by increased revenues due to greatly increased unit sales, albeit with reduced unit cost. Worldwide semiconductor revenues typically rose by 16% a year between by 1960 and 2000. This century revenue growth has slowed to around 6%.

6.1.5 Silicon evolution

SiGe has already been mentioned as a process of choice for high performance RF. However as it is bipolar based it tends to be one or two process nodes behind CMOS die sizes and therefore more costly.

Silicon on insulator (SOI) is another fairly new technology which has been used for many special applications in the past such as radiation-hardened or high voltage circuits. Compared with bulk CMOS, SOI technology is able to offer a higher maximum frequency and better

[21] Kumar, Rakkesh, “Fabless Semiconductor Implementation” The McGraw-Hill Companies 2008
[22] Kumar, Rakkesh, “Fabless Semiconductor Implementation” The McGraw-Hill Companies 2008
[23] Chang, Morris, “Foundry Future: Challenges in the 21st Century” TSMC, ISSCC 2007

linearity. Its lower parasitic capacitance, than conventional bulk CMOS, means that it can also offer lower digital power consumption. The table below makes a basic comparison between the 0.18um CMOS and 0.35um SOI CMOS technologies[24]. The ITRS, an industry body which roadmaps semiconductor performance predicts that SOI may be mature enough to be used in volume applications by around 2015.

Function	Si CMOS 0.18um	SOI CMOS 0.35um
Fmax (GHz)	30	60
Linearity	Good	Best
NFmin (@2GHz)	<0.8	<0.8
RF switches	Poor	Best
Low power digital	Yes	Yes
Passives integration	Poor	Good
A/D; D/A	Yes	Yes
3V swing (dynamic range)	No	Yes
EEPROM / Flash	No	Yes
Isolation	Poor	Good -> best
Cost	Best	Good

Table 6-4: CMOS and SOI CMOS comparison

Both SiGe and SOI are both possible technologies which could replace CMOS as CMOS scaling reaches towards its physical limits. Figure 6-3 highlights some of the silicon material a trends.

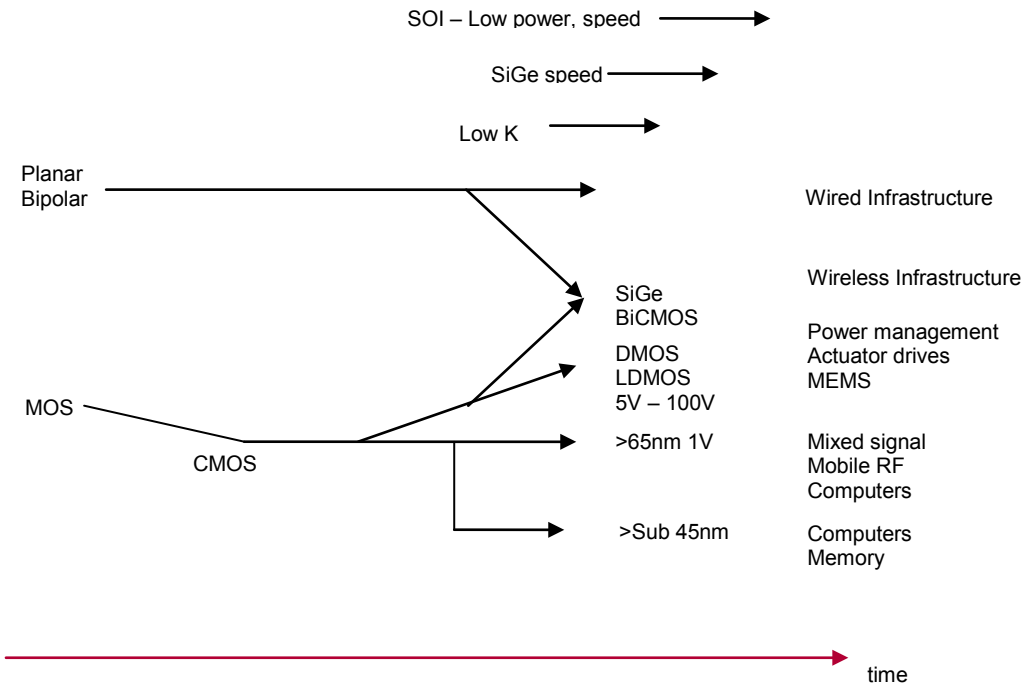

Figure 6-3: Semiconductor materials

6.2 Receiver design techniques – hot topics

Most of the newer receiver design techniques are aimed at making the best use of CMOS with relatively poor large signal handling and to minimise the number of external components required.

[24] Ma, Vivian, "SOI vs CMOS for Analog Circuit" http://www.eecg.toronto.edu/~kphang/papers/2001/ma_SOI.pdf [15th October 2009]

6.2.1 Tracking filters

Tracking bandpass filters, located before the down conversion mixer have traditionally been used in superhet wideband receivers for image rejection. As well as providing image rejection these filters can help provide good rejection of interferers several channels away from the wanted by reducing the unwanted signal amplitude prior to down conversion. They are generally not required in narrower band receivers as fixed frequency SAW filters can be used to reduce the effects of far-off (out of band) signals. Silicon tuner TV manufacturers, although not needing image rejection, have found it necessary to integrate tracking filters into silicon TV tuners so that performance comparable to canned tuners can be obtained.

The effectiveness of the tracking filter on the overall receiver's performance is largely dictated by:

- The Q of the filter.
- How accurately the filter can be aligned to the wanted channel.
- The position in the receiver. The closer the filter is to the antenna the more it will protect the entire receiver from unwanted out of band signals but it will contribute more to the noise figure of the receiver.

Traditionally implemented tracking filters, using several inductors and varactor diodes acting as variable capacitors, do not lend themselves to integration. Currently silicon tuner manufacturers are taking a number of approaches including:

- A system in package (SiP) approach where external to the IC is inductors and varactor diodes. Within the IC is a DC to DC convertor to generate the high voltage for the diodes and a synthesiser to generate the signal for automatic calibration.[25]
- An IC with integrated capacitors calibrated during manufacture which resonate with a high tolerance off board inductor[26]
- Filters implemented fully onboard. This technique uses on-board capacitors and integrators to realise a bandpass filter at the RF frequency[27]
- Feed forward blocker cancelation using a receiver cancelation loop as shown in figure 6-4. It has been suggested that this technique could replace SAW filters in a GSM receiver[28].

[25] Filatre et al, " A SiP Silicon Tuner with integrated LC Tracking filter for both Cable and Terrestrial TV Reception" NXP Semiconductors, ISSCC 2007
[26] For example MAX3580, Maxim Semiconductors
[27] Sun et al, "On chip Active RF Tracking Filter with 60dB 3rd order harmonic Rejection for Digital TV tuners", http://u-radio.kaist.ac.kr/pdf/conferences/2008/9.pdf [accessed 8th October 2009]
[28] Hooman Darabi, "A Blocker Filtering Technique for Wireless receivers", Broadcom Corporation, ISSCC 2007

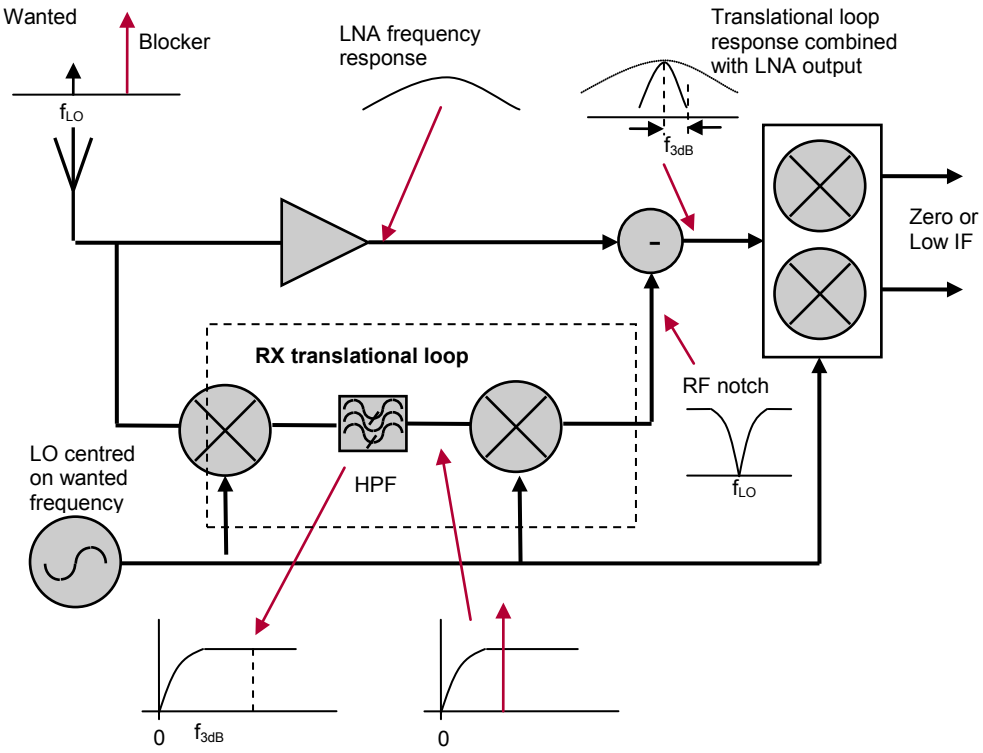

Figure 6-4: Blocker cancelation using feed-forward techniques

Of these techniques the SiP approach is thought likely to be short lived with a fully integrated approach, either using a bandpass filter or techniques such as the feed-forward translational loop being the preferred long term solution.

.

6.2.2 Discrete Time receivers

The high level of integration demanded by modern small products has driven the requirement for RF circuits used in some very high volume applications to be integrated with large digital circuits. This approach has forced engineers to apply digital techniques to analogue RF functions. This allows RF circuits to share the same CMOS process as digital circuits and benefit from CMOS scaling. However these circuits need to cope with a silicon process which does not have specialised linearized components; requires a reduced supply voltage and increases the coupled interference levels form other parts of the die.

Digital techniques applicable to CMOS receivers include discrete time sampling receivers[29]. A direct sampling mixer is shown below. The receiver's LNA is replaced by or followed by a trans-conductance amplifier which converts the received RF voltage vRF into a current iRF. The FETs mixes the signal down and the charge is stored on C_s.

[29] R.B. Staszewski., "Digital RF Processor (DRP™)" from "Wireless Technologies, Circuits, Systems and Devices" edited by K Iniewski, CRC Press 2008

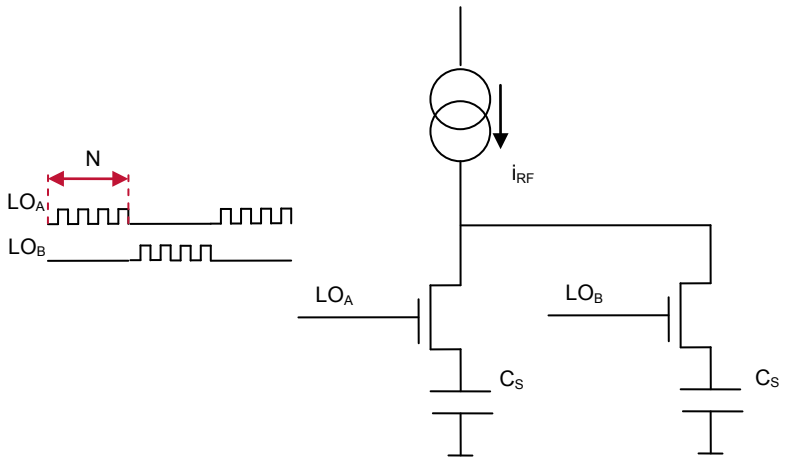

Figure 6-5: Sampling mixer

By integrating, over time, the charge on C_s over N cycles the resulting voltage ($V=Q/C_s$) increases giving rise to a discrete signal processing gain of N. In addition a Temporal FIR (Finite Impulse Response) Moving Average (TMA) filter is implemented with a sinc frequency response. An example is shown in figure 6-6 with N set to 8 and a sample rate of 2.4GHz. It can be seen that with the 2.4GHz sampling rate, e.g. for a Bluetooth signal, the first notch in the filter response occurs at 300MHz.

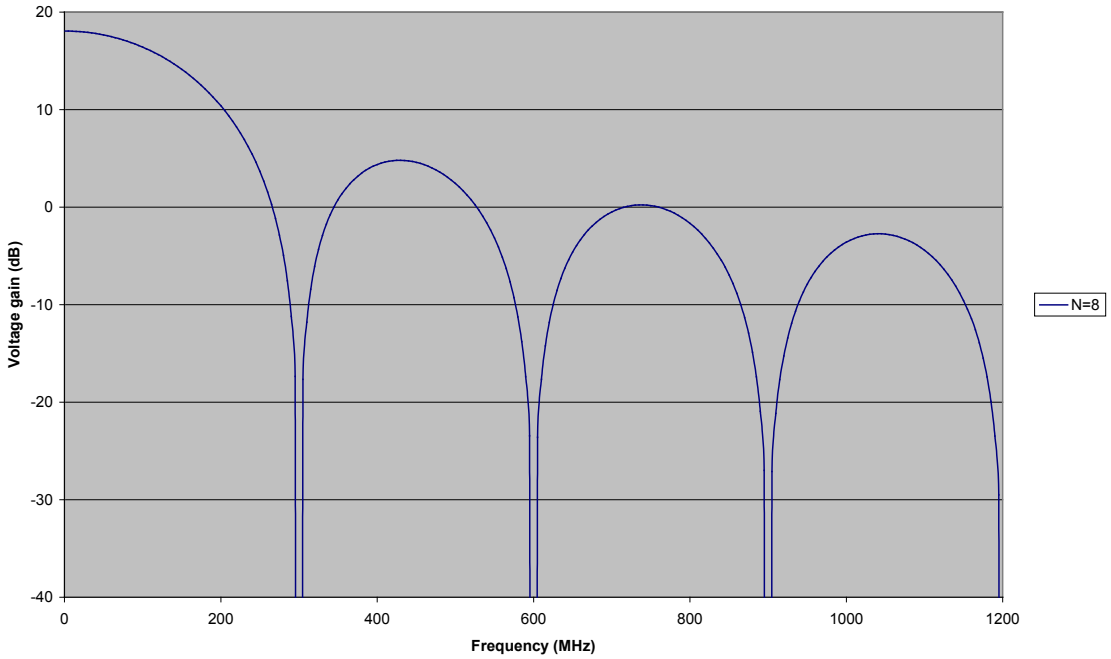

Figure 6-6: Transfer function of the sampling mixer

With C_s split into C_H and two C_R capacitors, the charge can be stored and read out continually. One C_R is used for integrating the RF voltage whilst the other C_R is used for readout. C_H continually stores a portion of the charge allowing a delay element to be created.

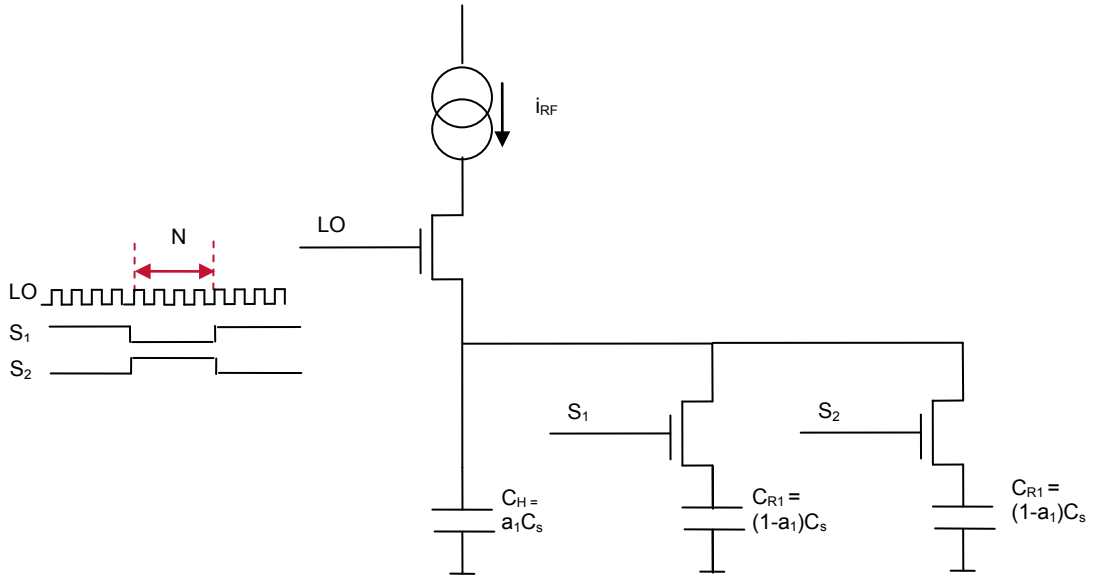

Figure 6-7: IIR filter with cyclic charge readout

With C_s split into C_H and C_R a first order IIR (Infinite Impulse Response) filter is created. This is a much stronger filter than the FIR filter. It is not practical to read out the charge at 300MHz. By adding multiple switched readout capacitors the readout time can be increased. For example by replacing each read out capacitor with four capacitors, the readout rate can drop to 75MHz. This adds a second moving average FIR filter with the first notch also occurring 75MHz. The mixers combined filter responses provide around 50dB of rejection for signals above 75MHz. A "standard" continuous time RF filter is still needed prior to the mixer to prevent interferers around harmonics of the sampling clock from folding to baseband.

A second switched capacitor IIR filter is incorporated in the readout buffer used to isolate the high impedance of the mixer from the low output impedance needed to drive the next stage. By placing the 3dB corner frequencies of the two IIR filters at the edge of the wanted frequency band significant close in filtering can be achieved. When used with an ADC with sufficient dynamic range, this approach can be used to allow a GSM receiver to pass with some margin the GSM 3MHz blocking tests.

Whilst traditional continuous time analogue filters using topologies such as Gm-C, active RC and LC depend on amplifier gains and passive component absolute values to define their cut off frequency and other parameters, the discrete time receiver's selectivity can be digitally controlled by the LO clock frequency and capacitance ratios. Gain can be controlled by changing the capacitance ratios within the circuit. Both of these are amenable to migration to digital CMOS and can potentially be migrated from one CMOS node to another without too much redesign work.

6.2.3 Direct sampling receivers

An ideal versatile radio would perhaps remove any frequency selective analogue parts, except for the antenna, allowing the radio to be fully software defined in order for it to be able to receive any frequency band that could be sampled by its ADC.

Currently cellular and Wi-Fi frequencies extend to nearly 6GHz. In order to adequately sample this bandwidth, a radio receiver would require a sample rate of 12GHz and perhaps 12 bits of dynamic range. Current ADC technology is some what behind this. Increasing the power helps,

GSM base station ADCs operate with 13bits of resolution at 100MHz but consume almost 1 Watt, far too much for a handheld device.

A figure of merit (FOM) of ADCs is[30]

$$FOM = \frac{(2^{ENOB})f_{sample}}{P}$$

For sub-sampling ADCs f_{sample} can be replaced with twice the effective resolution bandwidth of the convertor as this maintains the Nyquist criterion.

Assuming a battery powered device with a maximum allowed ADC current of 10mW, similar to today's ADCs, an ADC with a FOM of $4.9x10^6$ GHz/W would be required. For a similar performance device operating in a mains powered device, e.g. a television, consuming 1W of power a FOM of 49×10^3 GHz/W would be required.

By understanding, due to improved silicon processes, how the ADC's FOM is likely to improve with time it is possible to gain an understanding of when direct sampling of the RF signal could potentially be implemented. The ITRS have studied this. In 2007 the best convertors had a FOM of or 1500GHz/W. By 2022 the ITRS anticipate convertors with a FOM of 6 to 10×10^3GHz/W. This is still 5 times less than that required for a direct sampling ADC system described above.

It is conceivable that by 2022 a direct sampling sample DVB-T tuner could be built. It is felt unlikely that direct sampling handheld equipment would be viable within the next twenty years.

Whilst it appears impractical for full Nyquist bandwidth sampling of the received spectrum it is conceivable that a sub sampling approach could be used, allowing portions of the spectrum to be sampled, assuming appropriate filtering can be implemented.

6.2.4 Spurious response rejection techniques

A traditional superhet receiver uses the amplitude response of filters to reject the image frequency whilst more modern architectures based around quadrature mixers uses phase based cancellation techniques to remove unwanted receiver responses.

It can be shown that similar *polyphase* techniques can be used to cancel out unwanted mixer products generated due to mixer responses at harmonics of the local oscillator. This technique is often known as a harmonic reject mixer.

An ideal mixer with a sine wave local oscillator doesn't produce spurious products at harmonics of the local oscillator. However in practice, the local oscillator is a square wave with odd harmonics magnitudes of 1/3,1/5, 1/7, etc. In a complex mixer these harmonics create products at $+3f_{LO}$, $-5f_{LO}$, $+7f_{LO}$ and so on. By placing three complex mixers in parallel, each fed with one of the polyphase LOs (same frequency, offset in phase) the mixer responses due to the 3rd and 5th order LO harmonics can be cancelled. By extending the number of receiver paths and LO phases the more harmonics can be cancelled.

Similar techniques can be used to eliminate some intermodulation products. Balanced circuits, used in virtually all RF IC design are effectively a polyphase circuit as one circuit branch is fed with the inverse (rotated in phase by 180°) to the other. This allows, assuming perfect balance, 2nd order intermodulation products to be cancelled. The approach can in principle be extended to some higher order products. Unfortunately it can't be extended to cancel 3rd order intermodulation products.

[30] International Technology Roadmap for Semiconductors, 2007 Edition System Drivers, http://www.itrs.net/Links/2007ITRS/2007_Chapters/2007_SystemDrivers.pdf [3rd September 2009]

6.3 Conclusions

In anticipating the trends in future receiver performance the following can be identified:

- As the RF spectrum becomes more crowded and spectrum regulation tends to become more market-led, it can be anticipated that for interference to not limit wireless system performance, receiver selectivity will become more important.
- As receive bandwidths increase, receive sensitivity will reduce necessitating higher transmitted power to maintain the link budget. For a given receiver with constrained large signal handling the selectivity of the receiver will degrade correspondingly.
- Increased receiver integration. For mature standards, in the highest volume designs, the RF receiver is increasingly being merged with digital basebands.
- Potential for RFICs to cater for a reduced number of standards or standard families, but with some standards being used for a wider range of applications.
- Increased use of standardisation in wireless systems with proprietary applications migrating to standardised protocols such as Zigbee whenever possible.
- Increased semiconductor development costs needing to be offset by increased revenues due to increased product volumes.
- New processes, such as SOI which may replace bulk CMOS once the physical limitations of shrinking CMOS geometries are reached, allowing predictions such as Moore's law to possibly hold for longer into the future.
- A wide range of receiver design techniques which will allow at least the selectivity performance of RF receivers specified in current standards to be met at lower cost in more integrated receivers.

7 FUTURE DEVICE COST VERSUS SELECTIVITY

In looking at current product costs versus selectivity, the selectivity has some influence on the overall product cost.

In the future, in systems where there is a clearly defined applicable standard which is seen as adequate by the industry, for example cellular, RF receiver selectivity performance beyond that required by the applicable standard is generally not perceived as important. In this case there is little impetus for semiconductor designers to improve the selectivity of the receiver beyond that mandated.

In systems where the applicable selectivity standard is not seen as adequate by parts of the industry, for example TV, premium brands are thought likely in the future to continue to ask their tuner supplier to supply parts with additional selectivity in order to provide as good a customer experience as possible. As long as some TV manufacturers perceive a need for improved selectivity, their semiconductor suppliers will be encouraged to improve the tuners selectivity performance.

It can be anticipated that the semiconductor manufacturer will make steps to improve the tuner's selectivity performance, ideally without increasing the tuner cost, in order to continue to gain design wins with these TV manufacturers. Some of the additional costs of improved selectivity will be offset by new design techniques and, for the digital portions of the receiver, denser CMOS nodes. At the same time an increase in tuner sales is also likely due to multi-tuner digital video recorders and TVs supporting picture in picture functions. This will also allow some of the costs of additional selectivity to be offset.

Unless improved selectivity is mandated, it is likely that some receiver manufacturers may choose not to improve their device's selectivity.

The receiver's selectivity performance will always be limited by physically possible bounds. The practical and theoretical physical limitations are dependent on the radio standard and are highly dependent on the C/N needed to decode the signal and the bandwidth of the signal. A practical example is demonstrated in appendix 1.

Improvements in selectivity will take time. A likely scenario will involve small selectivity improvement steps being taken with each new generation of RFICs occurring typically every two to three years. For any new RFIC there must be a big enough market to justify its development cost. In relatively new areas such as silicon TV tuner design, engineering experience gained will help lead to the practical selectivity of tuners getting closer to the theoretical limitations. This has already been seen in more mature areas such as GSM receivers. Early GSM phones struggled to achieve the -102dBm reference sensitivity required by the GSM standard. Modern chipsets achieve nearly -110dBm.

Moving forward, advances in silicon will allow lower cost very highly integrated receivers to be produced. For the highest volume applications, based on mature stable applications, the RF transceiver and base band may be combined.

As development costs rise for the highest density lowest production cost CMOS nodes, the size of market needed to justify the investment in the silicon development will also rise. This suggests that there will be opportunities in medium size markets for slightly less integrated approaches, e.g. with separate RF and baseband devices. These might use lower density CMOS processes (with reasonable f_t and supply voltage) allowing better large signal handling. If required, and where costs are acceptable, more specialist silicon processes such as SiGe, allowing very good large signal handling, could be considered.

Lewis Davies

Paul Winter

TTP is the premier independent technology and product development company. We work in partnership with our clients to bring new products to market, creating new business from advances in technology. Our partners include some of the world's largest and most successful companies as well as some of the smallest and most ambitious.

The Technology Partnership plc
Melbourn Science Park
Melbourn
Hertfordshire SG8 6EE UK

T +44 1763 262626
F +44 1763 261582
www.ttp.com
enquiries@ttp.com