White Paper

Communications Systems Design

This white paper presents an overview of system level design for wireless communications equipment. It will be useful not just to those involved in the development of new communications systems, but also to designers of new products using existing systems, or anybody seeking to develop new architectures and/or components to reduce product cost and/or size. It also describes the benefits and implications of developing highly integrated System on a Chip (SoC) solutions, which dominate in today's high volume communications products.

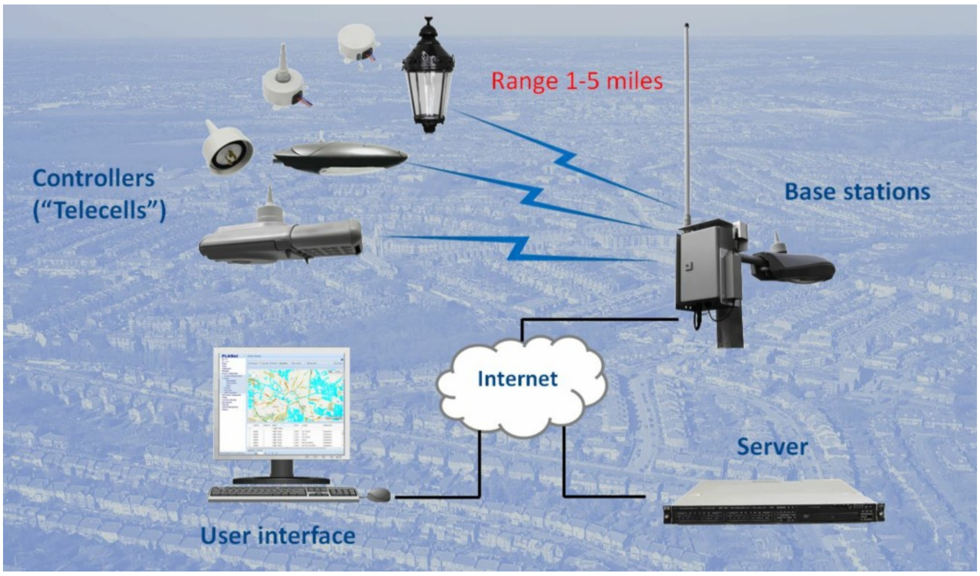

Figure 1: Wireless street lighting control system (courtesy of Telensa)

Introduction

Radio (wireless) communications systems have moved a long way in the last 100 years. From radio broadcasting, the first mass market application, to today's ubiquitous Smartphones offering an array of applications and communication at ever increasing data rates. This march of increasing functionality and complexity was enabled by the advent of the integrated circuit and the microprocessor, which together have facilitated miniaturisation and low-cost high volume manufacture.

Modern communications systems proliferate throughout the available spectrum and use a wide range of channel bandwidths, modulation schemes and data rates. The main traffic for wireless communications systems is now data rather than voice and the required data rates are ever rising. When designing a communications system the implementation choices depend on a host of requirements, some of which can be conflicting.

These include the required functionality and performance, the available technology, operational conditions and limitations, legislative restrictions, development budgets, unit cost targets and timescales. The design process is complex, requiring an in-depth understanding of technology, radio systems, digital processing and the ability to accurately simulate the effects of a practical realisation at a system level.

Applications and Requirements

Although consumer demand for wireless data is relentless, it would be wrong to think that all current developments in communications are aimed at maximising traffic throughput over wide channel bandwidths. This type of application certainly occupies an enormous amount of global development effort, but there are also applications where only a tiny amount of information needs to be communicated between sites and on an infrequent basis.

A good example of a modern low data-rate application is in the monitoring and control of street lighting, such as that depicted in Figure 1 (courtesy of Telensa).

Here each street lamp is fitted with a small two way radio link allowing monitoring and control by the local council (City authority). The council can control when the street lights are turned on, and to save energy late into the night, when they can be dimmed to a lower level or switched off. The street light can report back the state of its bulb and ballast, or more recently LEDs, so that replacement can be undertaken when near end-of-life. The amount of data that needs to be communicated to/from each light is small and infrequent. There is currently strong and growing interest in low data rate M2M (machine-to-machine) communications such as this.

Another growth area in communications is high data rate links for personal consumer applications. High data-rate WLAN and cellular communications is now the norm and is expected by users. As higher and higher wireless data rates have become available, new applications evolve to make use of the additional capacity. The volume of data that must be transferred is greatly amplified in the back-haul network where the live data from thousands of users needs to be transferred between the local cellular tower and the cellular operator regional control centre.

Microwave line of sight links are often used for this purpose, operating at high frequencies where wide channel bandwidths are available. In addition to using wide channel bandwidths, microwave links make use of higher order modulation schemes to squeeze more data through each channel.

Satellite systems allow truly global coverage. Global Navigation Satellite Systems (GNSS) provide one-way communications of navigational information worldwide. Two way satellite services are provided by other organisations and are used by emergency services in times of disasters as well as by news reporters and those who need communications beyond the range of the normal cellular services.

Satellite Satellite receivers are very sensitive. The signal levels are very low by the time they arrive on the ground and special attention is paid to minimising their noise figure. The down side is that they are more prone to interference from high power systems on nearby frequencies. Special protection is given to the satellite frequency bands by restrictions on other users nearby to ensure their radiated power was not excessive. However, the demand for growth of cellular services has recently put this protection under pressure with the allocation of cellular band frequencies close to some satellite bands.

Battery life is an important design parameter in many communications products. In some cases this simply needs to be long enough to avoid inconvenience (excessive re-charging of mobile devices) but for other applications (vehicle security devices, remote meter reading devices) a battery life of many years is essential for effective commercial deployment. Such a requirement must be considered from the outset and will be a key parameter in the system design process.

Modulation Schemes and Data Rates

The one resource that all radio applications have in common is the radio spectrum and this has to be shared by all users. As the number of applications continue to increase, it becomes more and more important that the use of this resource is undertaken efficiently.

In all cases, the modulation must be appropriate for the application, as efficient as possible and constrained to the bounds of the allocated channel to

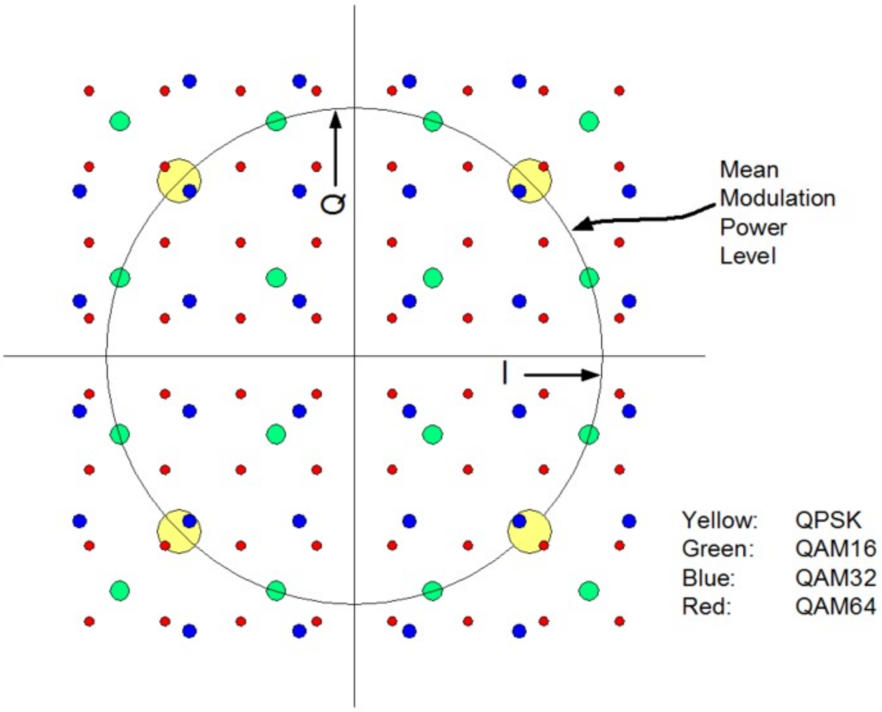

Figure 2: Combined I-Q constellation diagram

avoid spectral leakage into adjacent channels which could potentially interfere with users there.

There are many different modulation schemes operating in systems around the world today. For microwave point-to-point links, so called "high order" modulation schemes such as QAM64, QAM256 & QAM512 are commonly used. These offer high data transfer rates as each symbol of the modulation represents 6, 8 and 9 bits of data respectively. Their high data density comes at a price as it places challenging requirements on the error vector magnitude (EVM), the recovered eye diagram and the control of the symbol detection points.

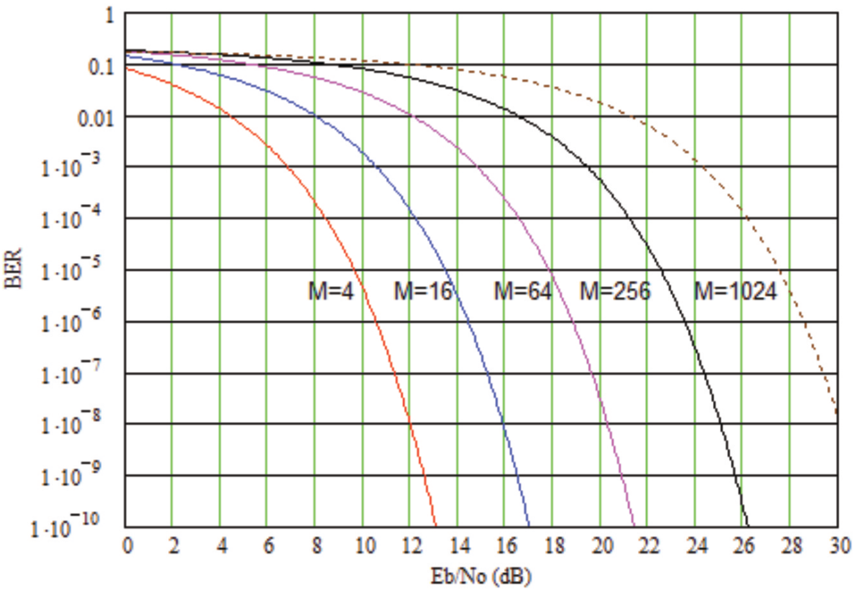

Figure 3: BER versus Eb/No for different modulation schemes,M=4: QPSK, M=16: QAM16 etc.

To illustrate this, Figure 2 overlays four separate modulation schemes on a single constellation diagram. The four schemes: QPSK, QAM16, QAM32 & QAM64 all have the same mean power level, but it is clear that the individual constellation points of the highest scheme (QAM64) are much closer to each other than are those of the lowest scheme (QPSK).

This makes the higher order scheme much more susceptible to noise and distortion since the amount of noise or distortion required to displace one particular constellation point from its region into that of another is smaller than that required for a low order scheme.

The presence of noise in the receive channel sets the fundamental limit on the Bit Error Rate (BER) that can be achieved for a given modulation scheme.

The higher the order of modulation used the higher the BER for a given noise level. Figure 3 shows a plot of BER versus Eb/No for a range of modulation schemes, where Eb is the energy per bit and No the noise spectral density. It is clear that squeezing more data through a given bandwidth by the use of higher order modulation schemes has an associated penalty in terms of the power that must be transmitted to

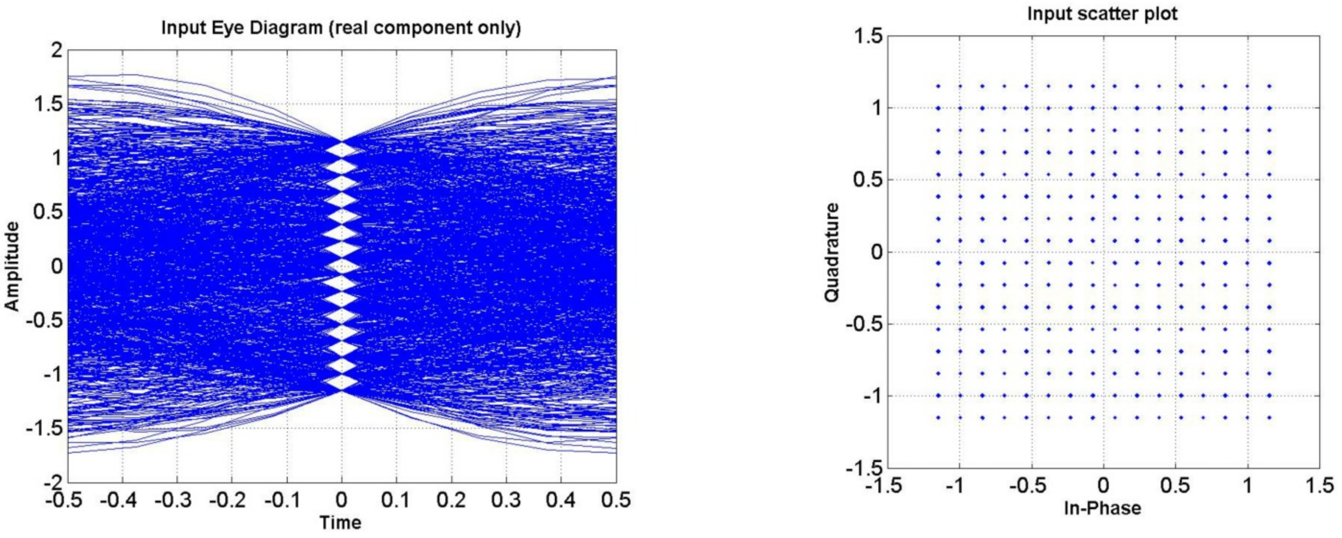

Figure 4: Constellation and eye diagram for QAM256 modulation

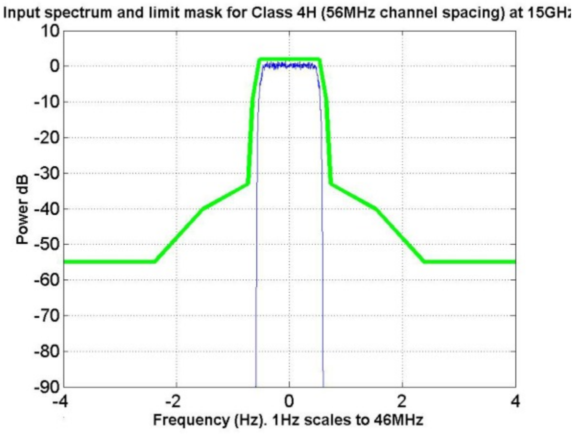

Figure 5: QAM 256 modulation spectrum (Blue) and transmit spectral mask (Green) for point-to-point links

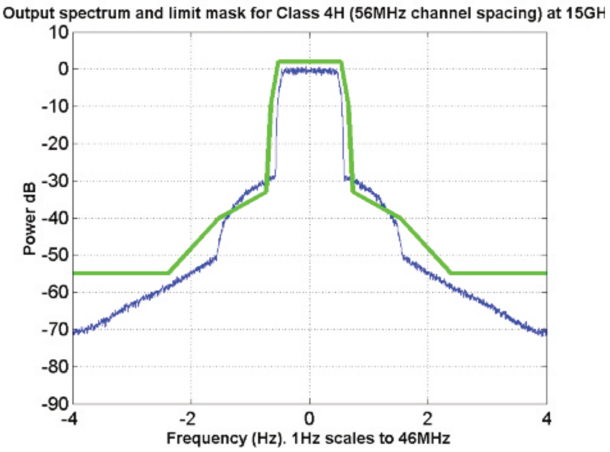

Figure 6: QAM 256 modulation suffering excessive spectral regrowth due to PA distortions

achieve the same BER (assuming a receiver with the same noise performance). There is also additional complexity in the radio design itself.

Figure 4 depicts the constellation and eye diagram for QAM256 modulation without any real-world corruption.

In a practical system imperfections in the transmitter, such as amplitude and phase distortion in the RF amplifiers and the phase noise of the local oscillator would cause distortion, the eyes would close a little and the constellation points would spread.

The frequency spectrum of a perfect QAM256 signal is depicted in Figure 5 together with the spectral mask showing the emission limits for a microwave point-to-point link (green trace). Root-raised cosine filtering of the transmitted data ensures its tight spectral control. It also introduces Inter-Symbol-Interference (ISI) away from the centre of the eye, as seen in Figure 4. The receiver must maintain symbol timing with sufficient accuracy to determine the symbol's value only at the eye opening point.

In addition to corrupting the constellation and causing eye closure, distortion in the transmitter can cause the signal to self-intermodulate and produce spectral regrowth (leakage) into adjacent channels. Figure 6 shows the spectral regrowth that PA non-linearities can cause. Here the regrowth exceeds the statutory limits imposed by the regulatory authorities.

This level of distortion also corrupts the transmitted constellation and eye diagram as depicted in Figure 7 and Figure 8.

Distortion within the receiver must also be well contained. Channel filtering, if applied too severely, can distort the modulated waveform, resulting in further eye closure and constellation distortion. This directly causes an increase in bit error rate (BER). When designing a radio system, sufficient time should be devoted to undertaking system simulations to predict performance in terms of EVM, spectral regrowth, constellation and eye distortion and BER sensitivity.

It is possible to simulate the entire communications system from bits in to bits out. This is a time consuming task but can be worthwhile particularly in cases where custom System on Chip (SoC) ICs are under development. Figure 9 compares the simulated BER to measured BER for a low-IF demodulator designed by PRFI and implemented as part of a custom ASIC now in volume production. The close agreement is a testament to the accuracy that can be achieved with careful design and simulation.

Modern cellular systems make use of Orthogonal Frequency Division Multiplexing (OFDM) to further increase data throughput. OFDM can be understood as a two-step process; first the data to be transmitted is split into a large number of lower data-rate streams. Each of these data streams is then modulated onto a sub-carrier as depicted in Figure 10, where each subcarrier is orthogonal with the others. Secondly, individual subcarriers are then modulated with

Equation 1: $$\text{Path loss (dB)} = 20.\text{Log}_{10}\left(\frac{4\pi}{\lambda}\right) + 10\text{n}.\text{Log}_{10}(\text{d}) - 10.\text{Log}_{10}(\text{Gtx}.\text{Grx})$$

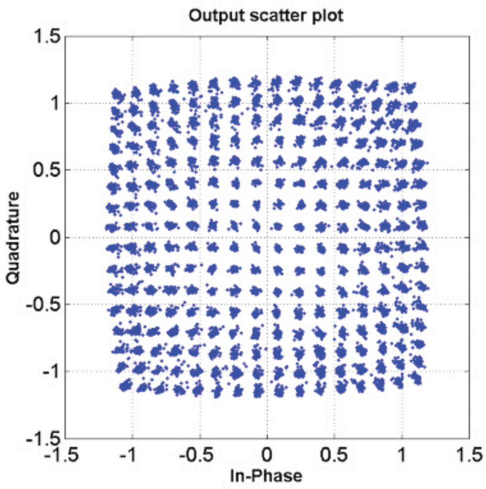

Figure 7: QAM 256 constellation corrupted by PA non-linearities

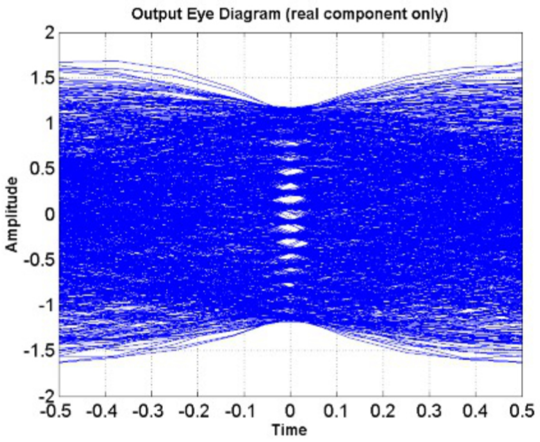

Figure 8: QAM 256 eye diagram corrupted by PA non-linearities

high order modulation to further increase capacity.

The sub-carriers are orthogonal because the carrier frequencies are chosen so that there is no cross-talk between the sub-channels meaning guard bands between the sub-carriers are not required. The use of OFDM allows a high degree of compaction in frequency giving a very high data density per unit of transmission bandwidth. The advantage of OFDM over high data-rate single carrier schemes is its ability to cope with propagation issues such as frequency selective fading, this is discussed in more detail below.

Use of spatial diversity is also increasingly common in modern mobile devices.

This introduces two or more separate transmission paths between the mobile phone and the cellular mast. These MIMO (multiple input multiple output) systems increase further the data throughput to and from the phone. Again, radio system design of these applications involves a lot of detailed simulation of the performance of MIMO OFDM modulated signals.

In free space, with a single direct line-of-sight path between transmitting and receiving antennas, the path loss follows a simple square law as defined in Equation 1, where d is the distance between the two antennas, Gtx and Grx are the antenna gains and n =2.

This propagation model would be valid, for example, for satellite to satellite communications where there is one direct propagation path and no additional paths due to reflections from objects. For ground based point-to-point links, reflections off the earth mean that the free space model is no longer valid. This effect can be accounted for by modifying Equation 1 based on antenna height and ground reflection coefficient. For low angles of reflection (long distance links) the path loss tends to a 4th law model.

Propagation and Link Budgets

At microwave frequencies additional loss is incurred due to rain. An allowance for this (rain fade margin) must

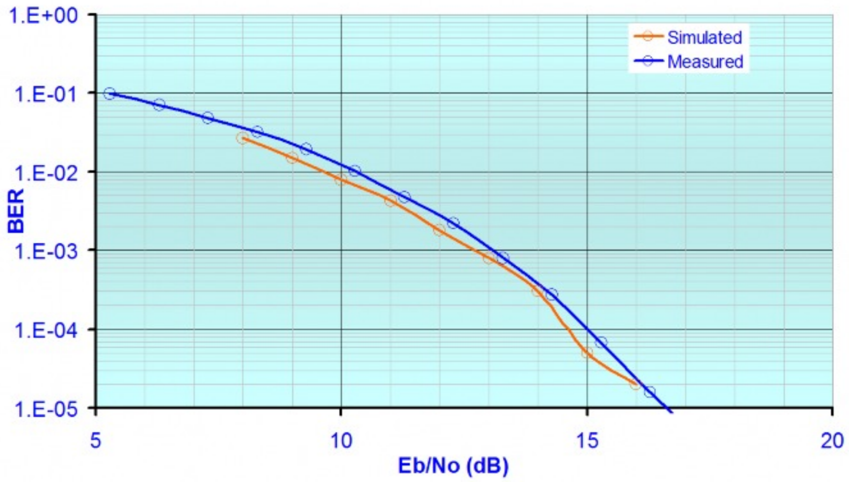

Figure 9: Comparison of simulated to measured BER

be made during the system design. The magnitude of this allowance depends on the location of the link (and so the typical weather) and the target availability (percentage of time that the specified loss must not be exceeded due to rain fall). Statistical data exists to allow appropriate rain fade margins to be determined.

There is also additional loss due to atmospheric absorption. This is both frequency dependent and altitude dependent but only really starts to become significant at mm-wave frequencies. The oxygen absorption peak at around 60GHz is an example of atmospheric absorption.

When designing a communications system the link budget must be balanced to ensure that there is adequate signal power arriving at the receiver to ensure acceptable Bit Error Rate (BER) for the demodulated data.

The link budget can be improved by increasing transmit power, increasing antenna gain (transmit and/or receive), reducing noise figure or changing modulation scheme or rate so that the required SNR at the receiver is obtained.

For non-line-of-sight propagation, additional losses due to reflection, diffraction and scattering caused by obstructions must be included. The simplest way to account for this is to increase the value of "n" in Equation 1 depending on the frequency, heights of the antennas and propagation environment. For urban environments, appropriate values of n can vary between 2, when streets can act like waveguides, and 6, when significant obstructions exist.

A more complex, but potentially more representative, approach is to model the propagation channel and include this in the system simulation. It is possible for reflections to cause two copies of the transmitted signal to arrive at the receiver with a phase difference of close to 180°. This would be a specific case of flat fading and it gives rise to a large dip in the received signal strength.

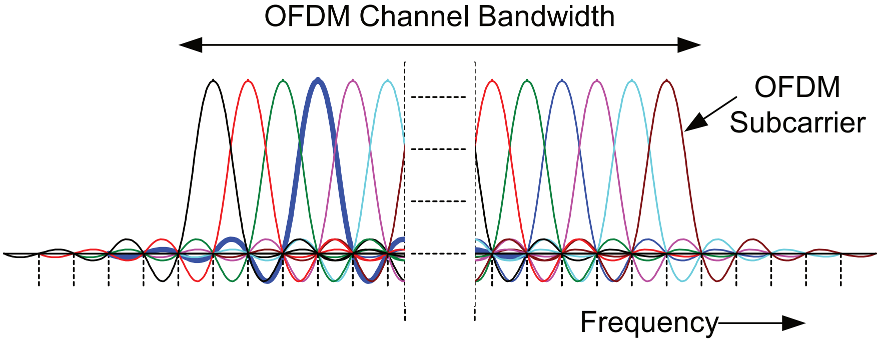

Figure 10: Simplified Image of OFDM

An approach to combat this is to implement receive antenna diversity, where an additional receive antenna positioned in a physically different location is also available and will be subject to different multi-path propagation effects.

The impact of multipath propagation on a transmission system can be assessed by undertaking a system simulation with an appropriate multipath model. A number of internationally recognised models have been developed to cover a range of different propagation conditions and Figure 11 depicts the channel amplitude response for two cases of the ITU Indoor Office Channel A model. The Indoor Office Channel A model is a 6-path static model having an overall 35nsec rms delay spread.

The difference between the two responses in Figure 11 is simply that Response No. 1 will occur at one RF frequency and Response No. 2 will occur at another RF frequency. Note that the x-axis is centred about 0Hz and this represents the centre frequency of the carrier.

The impact of this channel model on a 5MBd QPSK signal having a 6.25MHz bandwidth will now be illustrated. Note that the signal bandwidth will occupy the frequency range of ±3.125MHz about the zero Hz centre frequency. One would intuitively expect Response No.1 to have less impact on the signal when compared to Response No. 2 because the amplitude distortion of Response No.1 over the ±3.125MHz band about zero Hz is much less than that for Response No. 2 and this is indeed the case.

Figure 12 depicts the eye diagram for the two responses. As expected, the eye opening for Response No. 1 is good and shows much less distortion than that for Response No. 2; the position of the eye opening in the two responses of Figure 12 has been centred.

The corresponding constellation plot for the two responses is presented in Figure 13 and again highlights the better performance of Response 1. It shows that a channel having Response

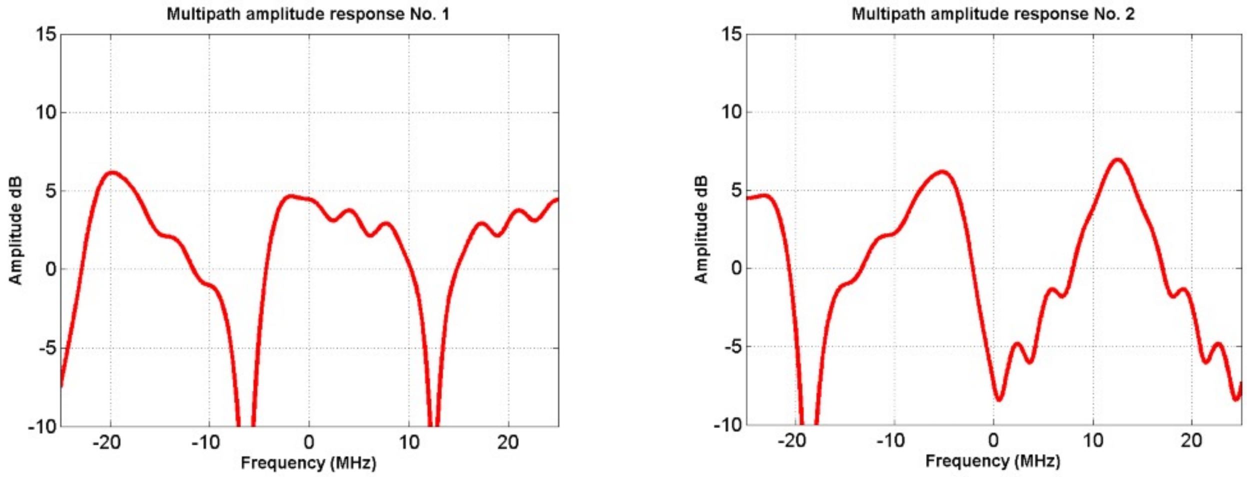

Figure 11: Two instances of the propagation path amplitude response of ITU Indoor Office Channel A

2 will have more difficulties receiving the signal. Combatting the effects of multipath must be taken into account when a system is being designed.

One effective means of combating ISI is to periodically transmit a known sequence of data (a training sequence). The receiver knows what it expects to receive when this sequence is transmitted and can remove the effects of ISI by adapting the weights of a digital filter (an adaptive equalizer) to essentially create an inverse model of the multipath interference caused by the reflections in the channel. With moving users (and / or objects moving in the local environment) the channel and its multi-path behaviour are constantly changing so the filter will be continually adaptive.

The use of OFDM, depicted in Figure 10, also helps to combat ISI. In this case the high data rate signal is split into multiple (n) sub-channels of a lower data rate. Each of these sub-channels is separated in frequency by the reciprocal of the symbol rate to minimise interference between other sub-channels. The resulting signal comprises n sub-carriers each with a data rate and modulation bandwidth reduced by a factor n. This allows each sub-channel to propagate independently along a multipath channel with very little amplitude and phase variation for that reduced bandwidth.

There will still be amplitude variation between sub-channels, but for an individual sub-channel the aim is to have a minimal variation. This means the sub-carriers can now be recovered without the need for adaptive equalisation of the channel, although it does require a more complex FFT based demodulator. It should be noted that whilst the need for equalisation may have been removed, some sub-carriers will be transmitted in propagation nulls, frequencies where there is significant attenuation, and these subcarriers may arrive at the receiver below the noise floor and therefore not recovered. The use of appropriate error correction and interleaving of data across sub-carriers can be used in mitigation against this effect.

Multiple Access Techniques

To allow multiple users access to the same communications system, a means of separating the different signals from each other must be used. This can be by separation in time (Time Division Multiple Access - TDMA), in frequency (Frequency Division Multiple Access - FDMA) or by applying a unique code (Code Division Multiple Access - CDMA).

Combinations of these techniques are also used. Separation in time and/or frequency is also used to separate receive signals from transmit signal so providing full duplex communications. These approaches all have different benefits and which is most appropriate depends on the requirements of the system.

Operating Frequency and Interfering Signals

The radio frequency spectrum is a global resource and its use is co-ordinated on a worldwide basis by the International Telecommunications Union (ITU); an agency of the United Nations based in Geneva. The ITU publishes a table of frequency allocations that represents agreement amongst the 191 member states of the ITU. It separates the globe into three radio regions on geographical grounds:

- Region 1: Europe, Middle East, Africa,the former Soviet Union, including Siberia; and Mongolia;

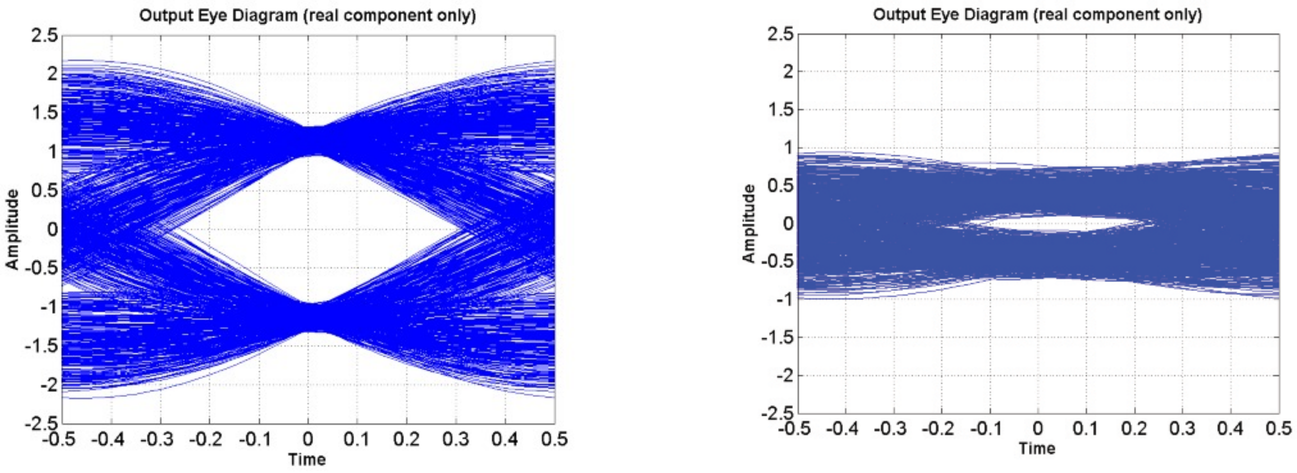

Figure 12: Eye diagram for 5MBd QPSK signal subjected to ITU Indoor Office Channel A multipath model. Response No.1 (LHS) and Response No.2 (RHS)

- Region 2: North and South America and Pacific (East of the International Date Line);
- Region 3: Asia, Australia and the Pacific Rim (West of the International Date Line).

Frequency allocations may differ to some extent between the three regions and the table is revised periodically at World Radiocommunications Conferences. Individual countries are responsible for managing the radio spectrum in their own jurisdiction and generally follow the lead given by the ITU table, but they are free to modify the table to suit their own needs. Any modification will usually be undertaken with the aim of not causing interference to other countries as these will be assumed to be abiding by the ITU plan.

Each country controls access to use the frequency spectrum in their jurisdiction by a variety of means. For some parts of the spectrum, the country regulator will set regulatory limits and issue licenses to control the interference between users; an example of this being users of private mobile radio systems. For other parts, the regulator provides less restriction and relies on users to self-control their emissions; an example of this would be cellular operators. Yet other parts of the spectrum are designated licence-exempt and here there is relatively weak regulatory control of interference levels; an example of this is the Industrial, Scientific and Medical (ISM) bands.

Any communications system must adhere to the appropriate legislation and the choice of operating frequency has a huge impact on the cost, achievable performance and design complexity.

Radio spectrum is a shared resource, so other users could be operating on nearby frequencies. Whatever the operating frequency of the system there will be strict limits on the allowed emissions at all other frequencies and this can be a significant design challenge, particularly for spectrum close to the operating band. Problems that can give rise to unwanted emissions from a transmitter include:

- Harmonic emissions
- Non-harmonic spurious emissions
- Adjacent channel leakage of carrier phase noise
- Spectral regrowth in adjacent channels caused by carrier self-intermodulation
- Spurious transient emissions caused by transient effects

It is also important that communication receivers can operate in the presence of legitimate communications devices operating at other frequencies. This requires care and attention during the design process. The main interference mechanisms experienced in a receiver are:

- Insufficient selectivity where imperfections in a receiver allow an adjacent transmission to corrupt a wanted signal. This can be combated by additional filtering within the receiver.
- Reciprocal mixing - the phase noise sidebands of the local oscillator are sufficiently high to cause an unwanted signal to be down-converted into the receive band
- Insufficient intermodulation immunity - receiver imperfections allow transmissions from two or more signals to interact within the receiver and generate an on-channel product that corrupts the wanted signal …. even though the

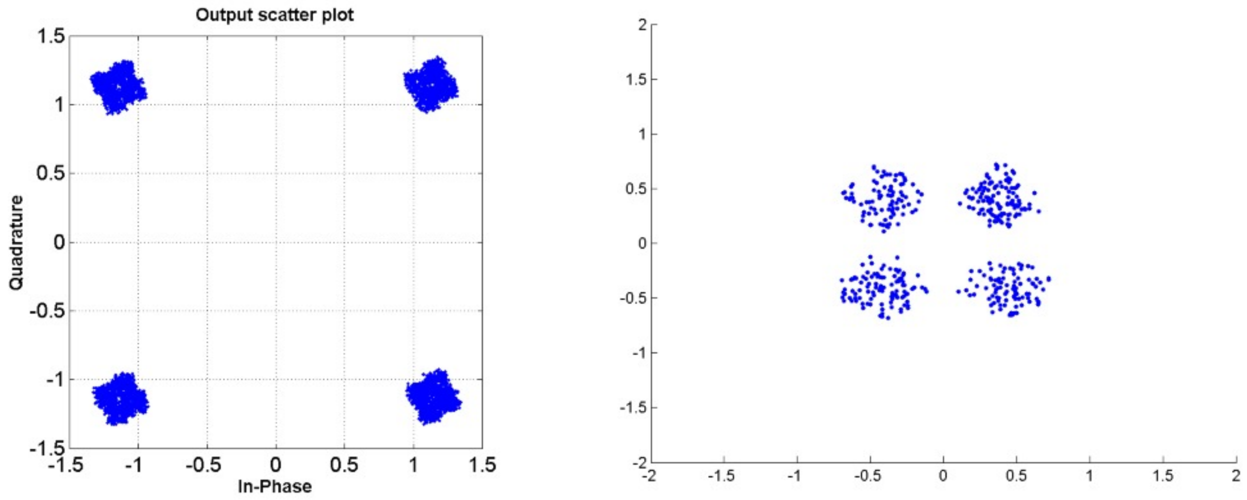

Figure 13: Constellation diagram for 5MBd QPSK signal subjected to ITU Indoor Office Channel A multipath model. Response No.1 (LHS) and Response No.2 (RHS)

interfering transmissions may be well separated from the wanted channel

- Spurious responses - unintended mixing products translate an unwanted signal into the receive band
- Blocking - where a very large signal overloads a receiver attempting to receive a weak wanted signal

Careful transceiver design is required to ensure the problems listed above are minimised for the application. Appropriate filtering must be included, the individual building blocks must have adequate linearity, LO and IF frequencies must be selected with care, local oscillators must have adequate phase noise performance, Automatic Gain Control (AGC) may be necessary.

Radio Architectures

A broad range of radio architectures have been developed. For high volume world wide applications, an aim will always be to seek the size and cost benefits of an integrated solution. A Bluetooth transceiver is one such example where the whole transceiver architecture is realised in silicon. Lower volume applications may be realised by a mixture of integrated and discrete components.

The move to full silicon integration over the last twenty or so years has been achieved by adapting transceiver architectures to make use of the advantages of integration whilst steering around the disadvantages. For instance it is possible to realise integrated analogue filters having useful selectivity at low frequencies but not at RF frequencies, so receiver architectures have been developed or adapted to make use of this. Such architectures include low-IF and zero-IF types and are the preferred approaches for highly integrated radios.

Figure 14 shows a block diagram of a single conversion heterodyne transceiver; an architecture that goes back to the early days of radio when designs were implemented using discrete components. The receiver mixer converts the signal from the transmission frequency to a standard IF such as 10.7MHz, 21.4MHz, 45MHz or 70MHz where crystal or ceramic filters provide narrow band channel filtering with high selectivity. Similar filtering in the transmitter ensures the transmitted signal only broadcasts on the wanted channel. Filtering at RF is used to reject image responses caused by the mixing process, provide

protection against blocking signals and reject transmitter harmonics Double conversion designs are used in applications requiring significant IF gain. The predominance of discrete filtering makes the architecture unsuited for integration onto silicon and so unable to fully benefit from the miniaturisation it can offer. This was the main reason that new architectures were sought, which were better suited in integration.

Figure 15 shows a block diagram of a zero-IF transceiver. This was one of the first developed with full silicon integration in mind. The received signal is mixed directly down to baseband where low-pass filters now

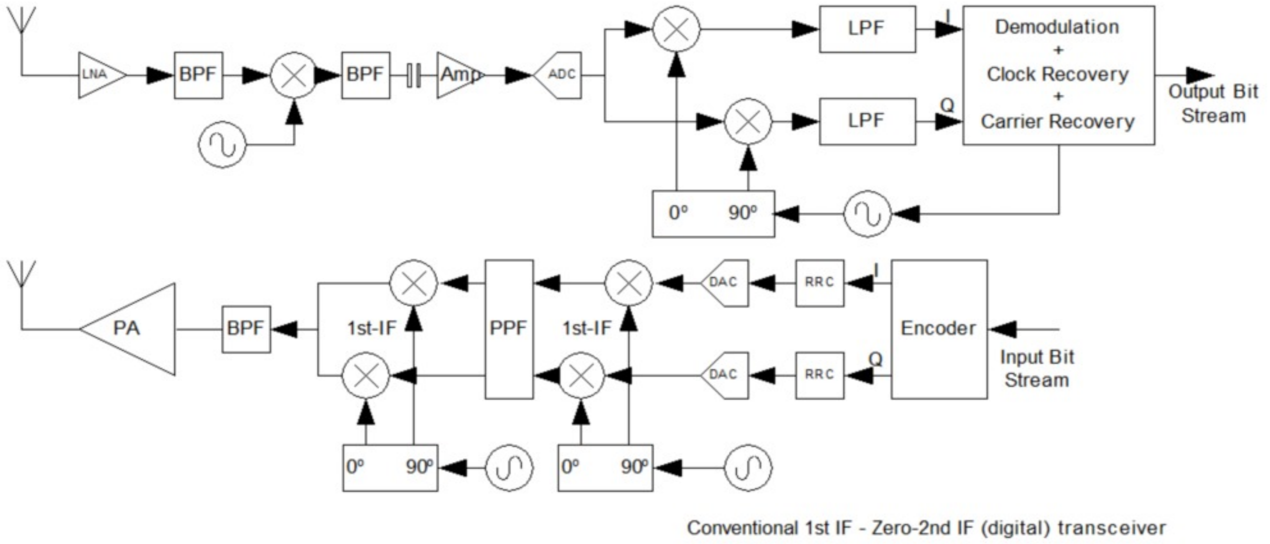

Figure 14: Single conversion heterodyne architecture

provide the channel filtering. Low pass analogue filters at baseband are suited for integration and can be tailored to a desired response. Two baseband channels (I and Q) are required to preserve both amplitude and phase information in the received signal. A drawback with the architecture is that it is susceptible to self-reception of its own local oscillator with the result that a DC component (a DC offset) appears in the baseband alongside the wanted signal.

This DC component needs to be removed and this can be done by either DC blocking with a low cut-off high pass filter or by a feedback arrangement that measures the DC offset and cancels it by adding an equal but opposite DC component via a DAC. Another disadvantage to this architecture is that it is susceptible to 2nd order distortion whereby the envelope of an unwanted interferer is translated directly to baseband.

The zero-IF transmitter is well suited to integration with the baseband modulated signal appearing immediately at final frequency. As before, two identical baseband channels are required to correctly produce the output modulated waveform. Imperfections in the mixer will give rise to leakage of the local oscillator signal appearing on the output, and to combat this, a DC offset correction (cancellation) would need to be applied.
Figure 16 shows a block diagram of a low-IF transceiver; an alternative design for full integration. In this case the received RF signal is mixed down to a very low frequency IF - close to DC - which, as with the zero-IF approach, allows the channel filters to be fully integrated.

These “*polyphase*” filters make use of the quadrature IF channel and produce a passband response at the frequency of the low-IF and a rejected response at its image frequency. The amount of image rejection achieved by this architecture is dependent on the degree of matching that can be achieved in the mixers and IF and the quality of the quadrature LO signal driving the mixers. The degree of image rejection will be finite and is likely to limit the receiver selectivity at this frequency.
An advantage to the low-IF receiver compared to the zero-IF (direct conversion) is that the self-reception that produced a problematic DC component can be overcome by the addition of a DC block. A second down conversion stage in the digital domain can be added to obtain a baseband signal output.
A low-IF transmitter converts a quadrature modulated signal at a low IF directly up to the final frequency. Matching in the low-IF strip and quadrature is essential to minimise the leakage of the local oscillator and leakage of the image signal to the output, however the degree of matching needed to meet regulatory requirements is often insufficient and an external bandpass filter may be required.

This approach can be problematic and so some solutions will use the simpler zero-IF architecture for the transmitter and adopt a mixed low-IF receive, zero-IF transmit design.

Further simplification of an integrated architecture may be possible when the

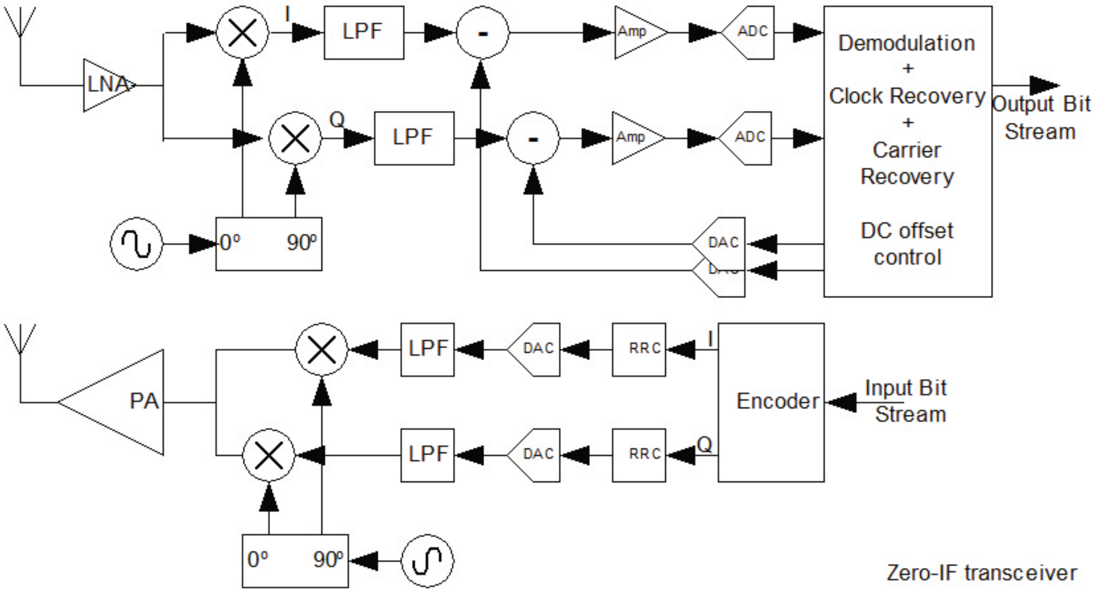

Figure 15: Direct (zero-IF) conversion architecture

operation of the radio system is considered. For example:

- A TDD system has only the transmitter or the receiver operating at any one time and this makes it possible to share one local oscillator between the two parts.
- Some protocols may have periods where it is know that there will be no signal to receive, such dead periods can be used to undertake self-calibration routines such as:
    - Cancelling a DC offset in the receive IF strip,
    - Tuning a banded VCO to the correct frequency,
    - Minimising the local oscillator feed-though in a transmitter,
    - Minimising the image response by adjustment of the local oscillator I-Q phase and amplitude balance

As in all engineering, trade-offs will form part of the design process, for example:

- The receiver selectivity that can be achieved will depend on a range of factors including the degree of filtering implemented, the quality of the local oscillator, specifically its sideband noise and for non-direct conversion architectures, the amount of image rejection that can be achieved.

- The trade-off between sensitivity and linearity will require careful specification of the individual blocks that make up a receiver chain. Too much receiver gain, whilst favouring sensitivity, will also reduce dynamic range and loss of performance when in the presence of strong interfering signals. Noise figure is often the guiding benchmark for the ultimate sensitivity that can be obtained whilst intercept point (for both second and third order distortion) is a benchmark for high level interference immunity.

- The transmitted EVM and adjacent channel leakage power will depend on the quality of the local oscillator, the modulation format and complexity, the linearity of the transmit strip amplifiers and their associated noise floors, etc.

Some, but not all, applications will be able to take advantage of a software defined radio approach - where the received signal is mixed to a low frequency and applied to an ADC for digitisation and then subsequently filtered, amplified and detected in the digital domain. The adaptability of such a software implementation provides a high level of flexibility to the application areas.

System Simulation

Designing and implementing a new communications system is time consuming and expensive. Effective simulation of the whole communications

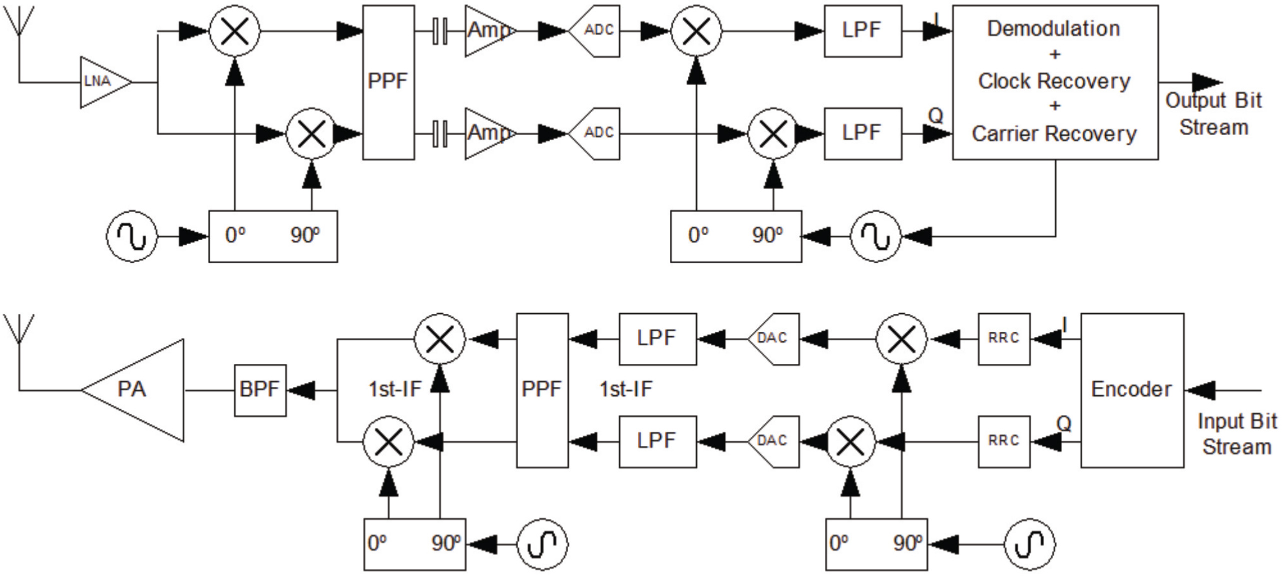

Low-1st IF (analog) - Zero-2nd IF (digital) transceiver

Figure 16: Low-IF architecture

system prior to implementation is essential. For the RF chain a cascaded spread-sheet analysis of parameters such as NF, gain, IP3, IP2 and compression was the traditional approach and this is still a valuable tool during the development process.

Various system simulator packages are now commercially available that allow this approach to be taken a step further with the incorporation of aspects such as filter responses, phase noise profiles of local oscillators and the effects of out of band interferers.

To fully optimise the performance of a communications system it is necessary to undertake a complete end to end communications system analysis from bits in to bits out, including the effects of the channel itself. This requires co-simulation of digital and analogue circuitry and is no small undertaking. For practical reasons systems are often developed with a slightly scaled back approach, where sections of the system are simulated independently.

If a highly integrated custom IC is to be developed, either for a new system or as a cost/size reduction exercise for an existing system, the need to ensure adequate performance through simulation is even more vital. The cost of committing to an integrated implementation is high, the options for implementing changes after fabrication are restricted and the timescales and costs of a second iteration cause nightmares for many project managers. In this instance simulation of the complete communications system, such as the BER simulation included in Figure 9, are highly recommended.

Concluding Remarks

The development of a new communications system is a huge undertaking. In many cases the use of an existing system and/or standard is the most sensible approach offering shortest timescales to deployment, lowest development costs and lowest risk. However, when new application scenarios or new functional requirements arise and the expected volumes justify the development costs a custom communications system can lead to performance and cost benefits. This white paper has given an overview of the communications system design process and discussed implementation choices and development requirements.