2.13.2.9 Antenna Diversity

Since propagation channels introduce multipath fading, the reception system must be designed to overcome fading in some way. The available alternatives include:

- Reduced channel symbol rates to reduce intersymbol interference (ISI)
- Structuring the data to be resilient to the effects of fading
- Diversity transmission and/or reception
- Slow FH
- Increased instantaneous bandwidth for multipath resolution and equalization

Reducing channel symbol rates may be necessary if other measures are ineffective. Interleaving and FEC reduce the impact of erasures introduced through fading. Diversity transmission and reception reduces this probability.

The impact of including diversity in an SDR includes both technical and economic challenges. Diversity antennas require parallel RF/IF conversion and ADC channels, increasing the cost of the system. They make it possible to delay and combine diversity paths more precisely and adaptively than is possible with analog approaches. The economic challenges center on minimizing the cost of such parallelism. The antenna, RF/IF processing, and ADC path can account for upwards of 60% of the procurement cost of a base transmission station.

The primary tradeoff, then, is to provide diversity in the architecture in a way that balances benefit against cost. CDMA's inherently wide bandwidth is robust in multipath, but also benefits from diversity combining, subject to receiver complexity constraints.

2.13.3 RF and IF Processing Tradeoffs

The second tradeoff concerns RF and IF conversion. Multichannel transceivers in TDD bands require an interference-suppression architecture that could include antenna isolation, programmable analog filters and active cancellation. The transmitter may require both linear operation (e.g., for W-CDMA and QAM waveforms) and nonlinear operation (e.g., class-C amplifier for high-power efficiency with FSK or PSK waveforms).

Single channel receivers may resort to nonlinear distortion of the incoming waveform (e.g., for a narrow-band direct-conversion architecture). Multichannel receivers, on the other hand, must match the RF and IF conversion parameters to the ADC, digital filtering and signal recovery algorithms of the back-end. The goal of this tradeoff is to balance the noise, spurious components, intermodulation products and artifacts. Since the direct-conversion receiver has fewer parts, its manufacturing costs may be less than the superheterodyne, which may have better performance. The thermal noise floor will be determined by the total bandwidth (e.g., in interference-limited bands below 400 MHz), or by the first LNA (e.g., in cellular and microwave bands). The thermal noise will be processed through the RF and IF conversion stages, resulting in noise shaping across the passband.

2.13.4 ADC Tradeoffs

Maximum sampling rate is obtained for a given clock technology in a quadrature-sampling ADC architecture. Such ADCs can introduce nonlinearities due to mismatching between the

in-phase (real) and quadrature (imaginary) conversion channels. Real oversampling with digital quadrature provides a lower-complexity alternative. In addition, one must match the ADC architecture to the structure of the service bands being supported. If two or three 25 MHz bands spaced hundreds of MHz apart are to be supported, one may have more total dynamic range using multiple medium-bandwidth ADCs instead of one super wideband (e.g., 500 MHz) ADC. Each medium-bandwidth ADCs access band may be programmable by tuning the final LO. Such an approach therefore complicates the RF architecture, but reduces the interconnect bandwidth and processing capacity of the next stage.

2.13.5 Digital Architecture Tradeoffs

The fourth tradeoff concerns the mix of parallelism and pipelining of the digital signal processing hardware from ASICs and FPGAs to DSPs and general-purpose processors. High-speed (gigabyte-per-second) digital interconnect is necessary to fan wideband ADC streams out to digital filtering FPGAs or ASICs. Reconfigurable processors and despreader ASICs may reduce or eliminate the need for wideband digital interconnect by either embedding the interconnect on-chip or producing baseband streams directly.

Digital filtering of high-data-rate ADC streams yields much lower data rate subscriber baseband channels. Medium bandwidth digital interconnect (hundreds of megabytes per second) then provides flexible paths among DSPs and general-purpose processors. The architecture of local and global memory among the processors also can be a significant contributor to algorithm performance. Balancing these high-speed data flows and bandwidth reduction steps against clusters of processing capacity and memory is a central concern of the digital architecture tradeoffs.

2.13.6 Software architecture tradeoffs

The fifth tradeoff concerns the organization of the radio software into appropriately packaged data structures and real-time algorithms. Interfaces among applications and services need to be radio-aware so that the radio's low data rates, high variability in data transfer times and occasional outages do not severely curtail user satisfaction with the services. In the radio applications layer, object-oriented design techniques help to group related data structures with appropriate algorithm methods. This simplifies detailed design, development, testing, deployment, and evolution of the software architecture. The terminology and approaches of object-oriented design may be applied to all the functions of the radio. This facilitates the realization of those functions in hardware, firmware, or software as a function of technology and project needs.

Projects may be implemented using conventional software techniques. With such approaches, the radio applications and infrastructure software elements are interwoven. Open-architecture approaches now favor the use of the industry-standard CORBA in radio infrastructure middleware. Such middleware reduces the coupling between radio functions and distributed processor hardware. This adds flexibility but requires processing capacity above that which is needed for a closed architecture. Accurate characterization of the processing requirements of modular collections of open-architecture software can be challenging.

2.13.7 Performance Management Tradeoffs

The final major tradeoff concerns the management of processing demand offered by the software against the resources provided by the hardware platform. Accurate characterization of processing demand requires benchmarking. A sustained measurement and instrumentation campaign to monitor performance implications of development decisions reduces development risk. Performance prediction and management steps add cost to a software radio development program. One therefore must balance the cost of performance management against the benefits of reduced development risk.

Other important tradeoffs include end-to-end tradeoffs. One of the most important in the software radio is the allocation of dynamic range among RF, IF, ADC, DSP hardware (e.g., FPGAs and ASICs) and algorithms. Another is the allocation of software objects to hardware components.

Resources critical to software radio architecture include I/O bandwidth, memory, and processing capacity. Good estimates of the demand for such resources result in a well-informed mapping of software objects to heterogeneous multiprocessing hardware. Depending on the details of the hardware, the critical resource may be the capacity of the embedded processors, memory, bus, mass storage, or some other I/O subsystem. The first step is the identification of the system's critical re-sources. The critical resource model characterizes each significant processing facility, data flow path, and control flow path in the system.

A critical resource is any computational entity (CPU, DSP unit, floating-point processor, I/O bus, etc.) in the system. Hand-coded assembly language algorithms may outperform high-order language (HOL) code (e.g., Ada or C) by an order of magnitude. Rigorous analysis of demand and capacity in terms of standards MOPS per critical resource yield useful predictions of performance. Initial estimates generated during the project-planning phase are generally not more accurate than a factor of two.

In a multithreaded DSP, there may be no explicit queues but if more than one thread, task or user is ready to run, its time spent waiting for the resource constitutes queuing delay. SDRs require three to four times the raw hardware processing capacity of ASICs and special purpose chips.

2.13.8 Cost Tradeoffs

Since cost of production electronics is nearly a linear function of parts count, the number of antennas and related RF/IF paths is critical. For each antenna, there must be at least some minimum amount of RF circuitry. And in most multiband, multimode radio designs, the parallelism of analog RF equipment extends to the ADC. As a result, the antenna and RF subsystems can account for upward of 60% of the procurement costs of a radio node. SPEAKeasy I and II, therefore, put considerable effort into developing all-band antennas, but with little success. The antenna therefore remains one of the most challenging aspects of SDR platform technology development.

2.13.9 RF/IF Conversion Segment Tradeoffs

Software radios require wideband RF/IF conversion, large dynamic range, and programmable analog signal processing parameters. In addition, a high-quality SDR architecture includes specific measures to mitigate the interference readily generated by SDR operation.

This includes a final stage of up-conversion from an IF, band-pass filtering to suppress adjacent channel interference, and final power amplification. First-generation cellular systems did not employ power control to any significant degree. CDMA systems, including 3G W-CDMA, require power control on each frame (50 to 100 times per second). SDRs may be implemented with a DAC as the interface between IF up-conversion and the RF segment. Alternatively, a high-speed DAC may directly feed the final power amplifier.

Power amplifiers have less-than-ideal performance, including amplitude ripple and phase distortion. Although these effects may be relatively small, failure to address them may have serious consequences on SDR performance. Amplitude ripple, for example, degrades the transmitted power across the band, particularly near the band edges. IF processing may compensate for this.

Since the receiver must overcome channel impairments, it may be more complex and technically demanding than the transmitter. The direct conversion receiver modulates a reference signal against the received RF (or IF) signal to yield a baseband binary analog waveform in the in-phase and quadrature (I&Q) channels. Although this kind of RF conversion has nonlinear characteristics, it is particularly effective for single-user applications such as handsets. It may not work well for multiuser applications, however.

2.13.10 Digital Processing Tradeoffs

A digital hard-ware design is a configuration of digital building blocks. These include ASICs, FPGAs, ADCs, DACs, digital interconnect, digital filters, DSPs, memory, bulk storage, I/O channels and general-purpose processors. A digital hardware architecture may be characterized via a reference platform, the minimum set of characteristics necessary to define a consistent family of designs of SDR hardware.

Part 3:

3.0 Enabling Technologies

The component technologies that form the backbone of SDR systems and set their performance limits are ADCs, DSPs, FPGAs, ASICs, DDCs/DUCs, software architectures & tools, filters and RF amplifiers.

The commercial and military interest in SDR has grown in recent years primarily because high performance DSPs are now available at reasonable prices. Radio functions are being implemented in programmable digital devices rather than hardwired analog components. The latest DSP operates at speeds up to 1.1GHZ and offer performance of nearly 9 billion instructions per second. FPGAs can now provide densities of up to 2 million gates with low power consumption. These numbers are ever improving. Yet more work needs to be done to improve the connectivity and interoperability of these components.

3.1 ADC

ADC is the most critical element of an SDR since its speed determines how close to the antenna the A to D conversion can be done. Usually ADC performance is defined based on speed (number of samples per second), resolution (how many bits each sample is coded into), and linearity (how accurately the digital output codes are related to the analog input values). For a given radio service and frequency band, the ADC's speed, resolution, and dynamic performance determine whether the digitization will occur at the RF or IF or in the baseband.

ADCs will need to sample at a rate double that of the highest frequency contained in the signal inorder to meet the Nyquist criterion and prevent aliasing. However, for bandlimited signals at very high frequencies, this leads to using a very high sampling rate.

All signals above the sampling frequency are aliased. So the ADC must be preceded by an analog antialiasing filter which passes only the band of interest and reject all other unwanted frequency components. The filter must exhibit good phase linearity and group delay performance with a minimum of insertion loss. Usually in the basic SDR architecture, the antialiasing function is effectively performed by the entire Analog Front End prior to the ADC.

It is the combined amplitude and phase response of all the components that must be considered. Using two or more stages of IF analog conversion prior to the ADC can simplify the filtering design task In the future, as the analog section in the SDR reduces in size, the antialiasing job will become more difficult. A SAW filter can be used as the primary antialiasing filter because of the design's excellent roll-off in the transition band. But the SAW filter exhibits quite high insertion loss. So it must be ensured that the SAW filter is preceded by a very low noise amplifier to ensure adequate noise figure.

Digitizing wide signal bandwidths can be challenging for data converters. In doing so, spurious performance is critical. Unlike single-carrier IF sampling, entire signal bands are digitized, thus allowing potentially hundreds of signals to be digitized. It is important that the converter should not generate spurious signals that interfere with the desired signals. These

spurious signals can be in the form of harmonics or intermodulation products. Either way, the results could be poor receiver performance. Spurious response consists of the second, third and higher-order harmonics. SFDR primarily determines how large an in-band or out-of-band interferer can be. If an interferer is sufficiently large, the A/D converter may generate a harmonic that appears as a co-channel interferer of the desired signal.

Spurious performance is very closely related to the air interface standard and is often the limiting factor in a receiver. For SDR, the distortion of the A/D converter must not impact overall system performance.

Spurious components have very narrow bandwidths and usually their amplitude is much larger than the rms quantisation noise power with in the information bandwidth. For a receiver the dynamic range will be limited by the spurious energy as the small information signals compete. The spurious performance must be such that any spurious signal generated is smaller than the desired signals of interest.

Distortion produced by digital conversion can be attributed to its static nonlinearly and dynamic nonlinearly features. Dynamic nonlinearity is the more important SDR-related parameter and this is commonly expressed through specification such as signal to noise and distortion (SINAD) and SFDR (Spurious Free Dynamic Range) SFDR.

SINAD is the ratio of the rms of required signal power to the rms of all the unwanted spectral components (including harmonies excluding dc components). A good data converter will exhibit a SINAD very close to its SNR.

Spurious-free dynamic range is defined as the range of signal amplitudes that can simultaneously be processed without distortion or be resolved by a receiver without the emergence of spurious signals above the noise floor. SFDR is expressed as the ratio of the rms signal amplitude to the rms value of the peak spurious spectral component. Spurious component may or may not be a related harmonic of the input signal caused by nonlinearity in the converter. SFDR encapsulates linearity and quantisation noise performance. It is a more important parameter than SNR or SINAD when considering the performance of either an ADC or DAC for SDR.

Effective Number Of Bits is the measured performance, in bits, of an ADC with respect to input frequency. As input frequency increases, overall noise also increases, thereby reducing the ENOB. Effective Number Of Bits specifies the dynamic performance of an ADC at a specific frequency, amplitude and sampling rate relative to an ideal ADC's quantisation noise.

For ADCs and DACs, a useful ***Figure of Merit*** is sample rate multiplied by SFDR.

Clock jitter has the effect of increasing system noise. The clocks used for sampling and holding are assumed to be perfect. Natural phenomena including bandwidth limitations, skew and noise prevents these clocks from being perfect.

Dithering is a method for randomizing the quantisation errors of an ADC by adding an uncorrelated analog signal to the desired signal at the input to the ADC. The quantisation error is not evenly distributed, it is ***correlated*** to the input signal. The function of Dithering is to ***decorrelate*** this error, to distribute the error across the entire spectrum.

The dithering signal is most often pseudorandom noise. By adding wideband pseudorandom noise to the input sine wave prior to sampling and digitization, the level of spurious components can be reduced by approximately 20dB. Dithering randomizes the dynamic nonlinearity errors, effectively reducing the harmonic content of the error signal. In other words, energy in the spurious signals is smeared out over many frequencies. Dithering can provide significant improvements in the ability of a converter to extract signals below its LSB resolution and in linearizing a converter's performance.

Dithering will improve the SFDR, but noise energy will be added to the system. This overall decrease in SNR may be acceptable if SFDR is the prime system requirement.

Dithering can be performed with a reduced impact on the noise floor by employing different techniques. The first approach uses larger amplitude but narrower band pseudorandom noise. This narrowband dithering energy can be added to a section of the Nyquist band seldom used by the communication systems. Filtering can then be used to remove the dithering energy and restore the SNR. A more complicated alternative digitally subtracts the pseudorandom dither signal, following digital conversion.

The aim of the SDR designer is to push the sampling frequency as high as possible to maximise the processable bandwidth of each processing chain. Higher IF and sampling frequencies cause a decrease in SNR. A doubling in sampling frequency will reduce the SNR by 6 dB for a fixed jitter.

3.2 DACs

High performance DACs are specifically used in the Tx signal path to reconstruct one or more carriers that have been digitally modulated. Many of the DSP functions are being integrated with the DAC itself to enhance its performance and to enable new transmitter architectures.

The quantisation noise generated by a DAC within the information band is generally not important for the transmitter in a radio system. This noise is attenuated during propagation and ends up being significantly less than the thermal noise seen at the receiver. However the out of band spurious components will be radiated and can cause unwanted intermodulation products appearing inside the information band of the receiver at the other end of the link.

In the reconstruction process of a digitally synthesized signal, it is the DAC and its nonideal characteristics, which often yield unpredictable results. In some cases, it is the performance of the DAC which actually determines whether a particular modulation scheme or system architecture can meet the specification. State-of-the-art DACs are 14 bit devices with SNR higher than 80dBc and sampling rate of 400Msamples/s. DACs are available from TI, Intersil, AD etc., that are suitable for 2G and 3G multicarrier SDR transmitters.

3.2.1 ADC/DAC status

- Mature Technology
- Key parameters- SFDR and SNR barriers cross the IF requirements 100dB/14Bits/16Bits 100MSPS and more.
- Promising Roadmap of industry leaders

- Can ENABLE the SDR Now.

The first developments of wideband software defined 2G basestations were fundamentally limited by the resolution and SFDR of digital converters. Designs with bandwidths of 5MHz or greater were limited to 12bit ADCs and the SFDRs of these devices were considered insufficient for some air interfaces like GSM. Now suitable wideband 14 bit ADCs with enough performance for GSM and better 16bit converters are available.

Today's high-performance 14-bit ADCs are enabling digitization at the first IF level. ADCs with sufficient performance are now available for software receivers using AMPS, GSM, CDMA2000 or UMTS standards.

Recent advances in converter performance are allowing the designers to move the ADC closer to the antenna, thereby simplifying the RF section and doing more and more signal processing in the digital domain. Concurrently, developers are investigating the integration of high-performance ADCs with DDCs to ease the task of interfacing these converters with follow-on baseband DSPs in the receive path of the communication chain

In the future, it is expected that ADC manufacturers will use super conducting technology that will operate at frequencies of several tens of GHz with SDFR between 120dB and 160dB. These ADCs will be able to translate microwave frequencies directly from the antenna into a digital bitstream. The high speed will enable the design of broadband SDR receivers that will cover bandwidths greater than 100 MHz. The high sensitivity will also enable the removal of any amplifier from the receiver chain.

3.3 DSPs

In an SDR, many of the radio functions are implemented in digital signal processing. The processing requirement in a modern 3G terminal is estimated to exceed 6-13 billion operations per second. Traditional approach is to place all of this processing power into a DSP. DSPs, as on today, which are rated at 9 billion operations per second also have a power dissipation that is measured in terms of many multiples of watts, with the end result of a battery life that is measured in minutes.

DSPs could theoretically handle the speed of processing required for high-speed signal processing of second generation Air Interface Standards. A DSP attempting to do signal processing for a second generation cellular handset would have to operate at a clock speed in excess of 325 MHz. In practice, because of bus delays, the need to write to buffers etc., it turns out to be in the GHz clock range. Early DSPs could not run at these speeds.

As DSP development progressed, it became clear that since power consumption varied directly with processor speed, it would not be practical to operate a DSP at these clock rates for battery powered applications. It also turned out that even for systems run on utility-provided power, there were issues of size, heat dissipation etc. which would make pure DSP solutions for high speed signal processing impractical.

Costs per MIPS of DSPs and general purpose CPUs have dropped below $10 per MIPS. The costs continue to drop by a factor of two every few years. At the same time, absolute processing capacities continue to climb into the hundreds of millions of floating-point operations per second (MFLOPS) to billions of FLOPS (GFLOPS) per chip.

3.3.1 Choice of DSP

Processing power will depend upon system factors. The choice of a DSP to obtain the required computation speed is not a straightforward matter of specifying the highest clock speed. Architecture and instruction sets greatly affect the speed of algorithm execution. MIPS (millions of instructions per second) is not a valid measure, since each manufacturer counts instructions differently. A highly useful recommended measure more closely related to algorithm execution is the peak *million-multiply-accumulates-per-second* (MMACS). This calculation is the product of the clock speed and the number of MACs the DSP is capable of executing per clock cycle.

On average, every 18 months, processing power doubles for the same volume, power consumption and cost. This equals to an order of 10x improvement every 6 to 7yrs. Investigations into the increase is algorithmic complexity experienced during the transitions from 1G to 3G suggest that the actual need for processing power is increasing rapidly. New architectural technologies are required to keep pace and the industry cannot rely on faster versions of traditional processing architectures like Super Harvard and others to solve the need for more signal processing power. Two recently introduced new classes to consider are: ***very long instruction word*** **(VLIW)** and ***static superscalar***.

VLIW describes an instruction set philosophy in which the compiler packs a number of simple, noninterdependent operations into the same instruction word. The sequencing mechanism in *VLIW* relies on an instruction format wherein every single execution unit in the chip is under direct programmer or compiler control. When fetched from cache or memory into the processor, these words are easily broken up and the operations dispatched to independent execution units. VLIW can perhaps best be described as a software- or compiler-based superscalar technology. VLIW attempts to reduce cost and increase execution speed by reducing hardware complexity.

Unfortunately, VLIW has little or no hardware support for maintaining the integrity of data dependencies or avoiding scheduling hazards associated with real-time processing. In VLIW, all operation latencies in a particular implementation are fully exposed to software. The **TMS320C6x series** from Texas Instruments is an example of a VLIW architecture.

Static superscalar architecture enforces a consistent and functionally well-defined programming model, and the schedule is determined prior to run time. It incorporates static scheduling techniques like those found in VLIW, but it retains many superscalar and RISC attributes, enabling real time systems. Consequently, code can be written directly in assembly without requiring sophisticated timing prediction. The TigerSHARC DSP from Analog Devices is an example of a *static superscalar* architecture.

The new TMS320C64x DSP core will deliver the highest performance DSPs in the world, with clock speeds of up to 1.1 GigaHertz (GHz) and performance near 9,000 million instructions per second (MIPS). This delivers 10x the DSP performance of the current industry leading devices like TMS320C62x, for key applications.

The new TMS320C55x DSP core slashes power consumption to a cool 0.05 mW per MIPS and will deliver performance up to 800 MIPS. With six times lower power than the industry-leading TMS320C54x, the C55x offers substantially longer battery life. The C55x builds on TI's C54x, the industry's most power-efficient DSPs used in 70 percent of the world's cell

phones, by enabling four times longer battery life to extend cell phone usage from days to weeks.

Both new DSP cores are software-compatible with TI's previous DSP generations for faster time to market and reduced development time. Additionally, TI's integrated development environment, eXpressDSP Real-Time Software Technology, supports the new C64x and C55x DSP cores by providing a complete, open DSP software environment with the tools and software to begin development today.

The C64x is the world's highest performance DSP core and the C55x is the lowest power DSP core to ever hit the market.

TI's next generation C64x DSP core is ideal for the needs of wireless infrastructure manufacturers, delivering the highest performance DSPs in the world, with speeds of up to 1.1 GigaHertz (GHz) and performance near 9,000 MIPS. Initial devices running at frequencies of 600 to 800 MHz provide speeds comparable to 10 times the DSP performance of the current industry leader chips. Ten of the top 13 leading wireless infrastructure manufacturers are designing their 3G basestations using TI's DSP technology.

3.4 ASICs

Although ASICs provide better performance at lower cost, their programmability declines as their level of integration increases. The use of a commercial ASIC targeted at a specific error correction algorithm would be inappropriate for any platform, and an FPGA or DSP implementation would be a better choice.

The number of functions or features required all consume silicon area, which equates to cost. Any processing function, which is placed in silicon, is totally inflexible to change. Any changes to standards, bug fixes, better algorithms, all equate to a new silicon design.

3.5 FPGAs

The signal processing requirements for military and commercial radio systems employing high data rate signals or spread spectrum modulation easily exceeds the processing speeds currently available in off-the-shelf DSP microprocessors and General Purpose Processors. Modern FPGAs are the only re-programmable components that can implement functions beyond the capabilities of today's DSP microprocessors. In recent years, the FPGA has advanced significantly in speed, size, and gate density.

These devices can provide the programmability of software, the high speed of hardware and can be re-configured on a real-time basis with no physical change to the hardware. It combines the flexibility of a programmable DSP with the performance of a pure ASIC.

New FPGA devices utilize extremely small silicon geometries allowing both very high clock speeds and low core voltages. Logic development has changed too, with new design tools that open up this area to software engineers, where once it was the realm of the hardware engineer. This means that algorithm and protocol development will be done in a more object-oriented way, enabling the creation of more complex systems.

Very large FPGA devices are difficult to manufacture and hence extremely expensive when compared with DSPs. The largest FPGAs are priced similar to DSPs when comparing dollars per theoretical MMACS. However, this situation rapidly deteriorates if actual signal processing capacity only reaches 25% of the maximum. For this case, the FPGA becomes four times more expensive.

Now it is feasible to implement processor cores and DSP-based algorithms within a single FPGA. Algorithms that can be partitioned to process multiple signals in parallel can be more efficiently implemented in an FPGA device. Compared to implementation in a processor, such partitioning enables the parallel calculation of many more MACs in an FPGA device. Modern FPGAs have sufficient capacity to fit multiple algorithms into a single device along with the interface circuitry required by the application, as a single chip solution. However, the implementation of DSP algorithms in FPGAs presents several technical challenges over DSP microprocessors.

Certain FPGAs have specific internal logic to implement efficient DSP algorithms, such as the Xilinx Virtex-II and Virtex-PRO families. All of the required processing of an SDR system can thus be accomplished in FPGA devices.

3.6 Digital Down/Up Converters

The 2G and 3G multichannel digital frequency downconversion has the following major requirements:

Down conversion requirements:-

- ?? Filter or isolate narrow band of frequencies from wideband source and reject the remainder of the band.
- ?? Translate the isolated carrier down in frequency usually from IF to baseband.
- ?? Reduce the data rate to some integer multiple of the information rate

Up conversion requirements:-

- ?? Translate one or more narrow band signal sources up in frequency from baseband to IF.
- ?? Combine the baseband sources to create one wideband signal
- ?? Increase the data rate to a digital 1F rate.

Multirate Processing

The data is processed at more than one sample rate in a system. Two key multirate DSP operations of decimation and interpolation are efficient algorithms most often used to decrease or increase the data rate at any given point in a system.

Wideband multicarrier SDR is an example of a multirate system. This is because the sampling frequency at IF is fixed for many information signals on different carrier frequencies. To move each carrier to baseband or 0 Hz, each carrier must be shifted by a different increment relative to the sampling frequency. However, the baseband bandpass characteristics must be the same for each carrier and the data rate must be an integer of the