# An Introduction to Software Defined Radios

By Jeff Blaine, ACØC – July 2009

**Rig types:  SCR vs SDR**

----------------------------------

Modern rigs are largely of the "SCR" variety – meaning software *controlled* radio.  The major functions of the rig are implemented in hardware, with a dedicated microcontroller providing rig management and interface – and (since the late 1990's) usually with a DSP section – to provide specialty signal processing (e.g. filtering, noise reduction, notch).

Rigs of the SCR type are generally fixed function – with firmware updates (in some cases) providing somewhat optimized function, bug fixes and improved DSP performance.  But the fundamental rig capabilities are limited by the manufacturer's specific hardware implementation.

The SDR – software defined radio – is a generic hardware platform – where the hardware has the function of converting the RF signal into a digital format in as direct and simple manner as possible.  Once the signal is rendered in a digital format, all further manipulation is done exclusively in software.  In a very simple block diagram, the SDR can be diagramed this way:

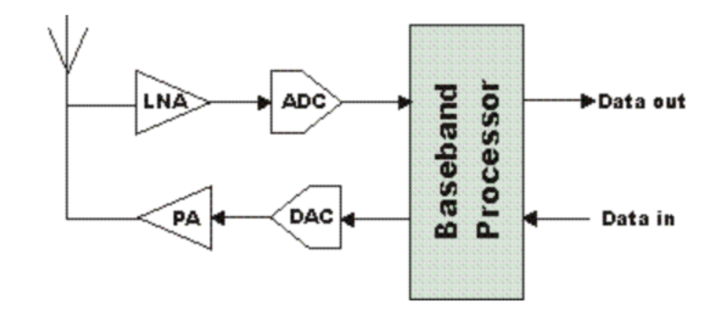


Block diagram of an 'Ideal' Software Defined Radio

Notice the four block functions – properly implemented – are generic to all modes and asks the rig will perform (within the frequency limits of the hardware).  That leaves the entire balance of the rig to be implemented completely in software.  Filtering, notch, noise reduction, mode and demod options, user interface, and control of the outside world are all under the exclusive control of the computer.  The SDR is said to be "always improving" and is "impossible to make obsolete."  In the ideal case, that is probably an accurate statement.

Having defined the world as being either the SCR or SDR variety, let's take a look at some actual examples. We want to start with what is common to the average ham.  Some of the most popular rigs on the mid-priced market are the FT-2000, K3 and the Flex 5000a.  Next we will compare these three rigs - with an eye to understanding each rig's use of the SCR and SDR technology.

**Superhet Rigs**

--------------------

Below is the block diagram of the FT-2000. A traditional triple up-conversion superhet design. The signal path is straight forward – an analog implementation through the 3rd mixer. From the 3rd mixer, the analog 30khz signal is converted to a digital signal for further processing by the DSP. Modulation detection, filtering functions and AGC are implemented by the DSP. A digital to analog conversation step follow driving the AF output.

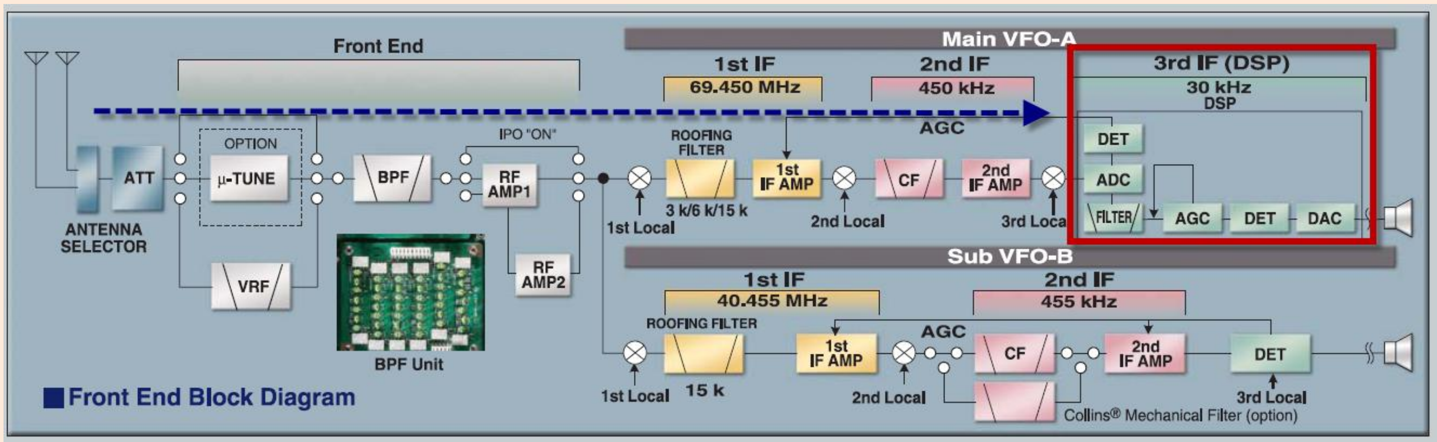


A similar architecture is provided in this diagram of the K3 receive path illustrated by the blue arrows. Again we see a traditional superhet design; in this case implemented as a dual down-conversion. The 2nd mixer output at 15khz handles the similar functions as does the FT2000.

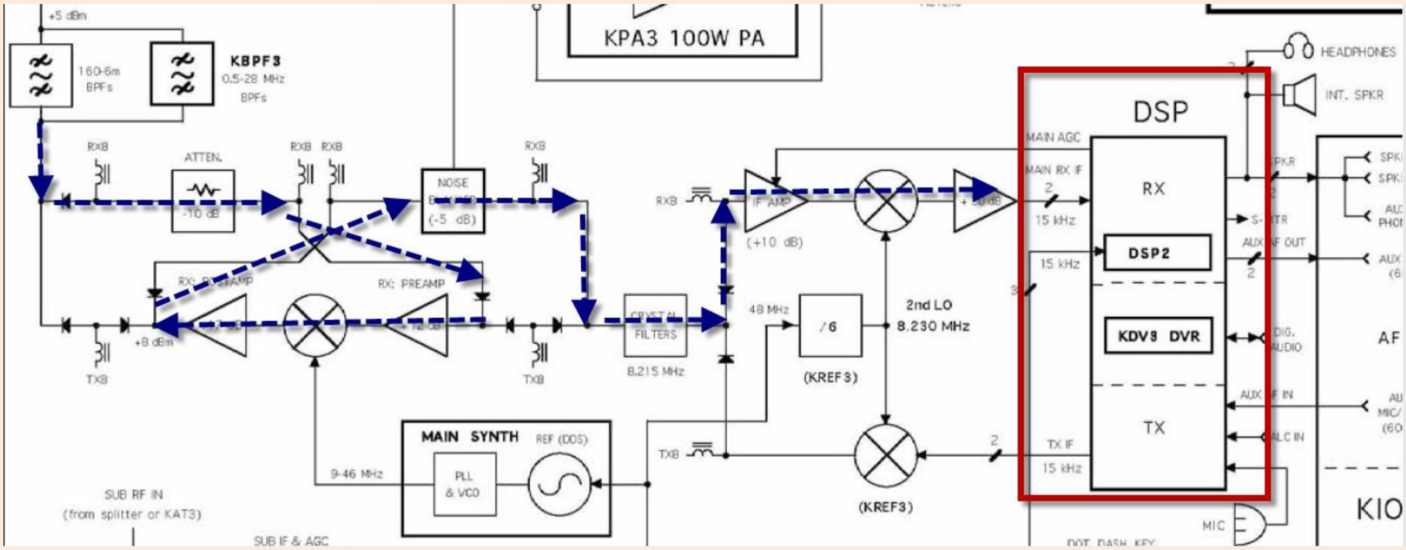


Other basic to high-performance rigs, including all of the Kenwood, Icom and Yaesu series, are variations of the same basic superhet topology although the details (and performance) vary greatly.

All traditional superhets share a common characteristic – in that the range of signals that the receiver must deal with is very large – and the chance for encountering two strong and closely spaced signals increases accordingly. This problem is illustrated here:

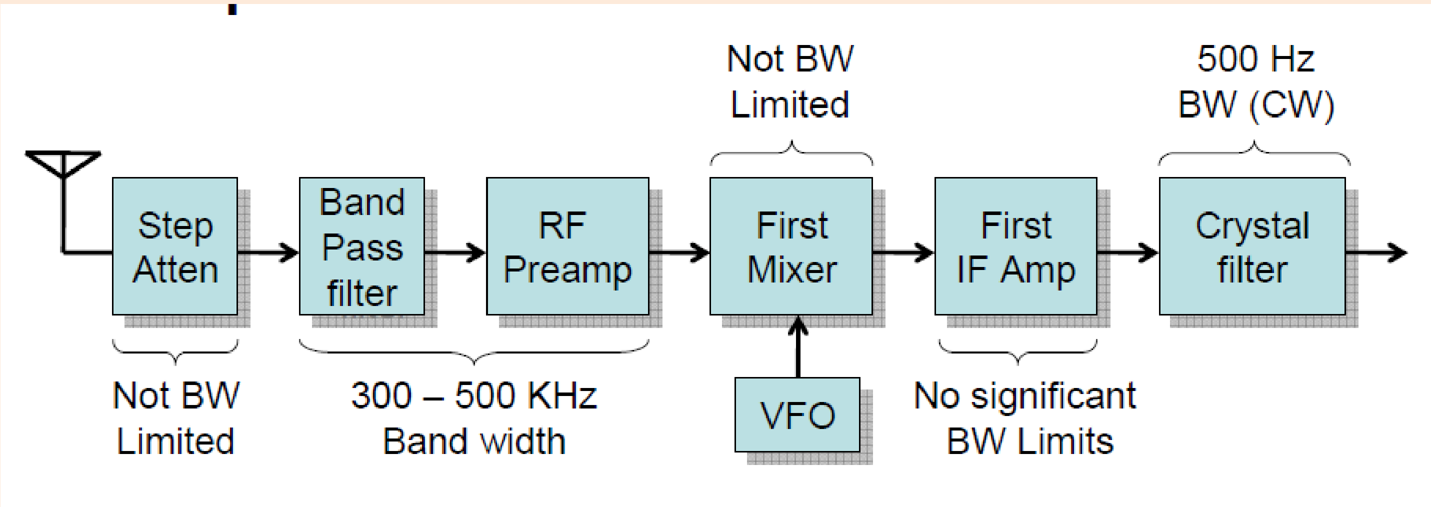


Source: [http://www.norcalqrp.org/files/AustinNC2030Presentation.pdf](http://www.norcalqrp.org/files/AustinNC2030Presentation.pdf)

The IF chain prior to the first roofing filter will see the entire spectrum of signals – limited only by the effect of the initial band-pass filtering.  So the majority of the IMDDR performance of the rig is defined by the 1st mixer and the roofing filter set.  In the worst case, strong signals present inside the roofing filter pass band create a tough environment for the 2nd mixer and subsequent stages.

IMMDR - "intermodulation dynamic range."  All rigs will generate false signals through an IMD mixing action.  The IMMDR specification indicates the level of these false signals under a given signal strength and spacing condition.

As the spacing between the two signals becomes smaller – and the signal strengths increase, the IMD-generated false signals become more likely to create problems for the operator.  Of course, receiver performance is quantified in many ways and the IMMDR metric is just one – however in the popular discussion, this has become the most common method of comparing the capabilities of two rigs in dense signal conditions.  The focus on IMMDR as a measure of fitness is really an acknowledgement that the typical rig now has all of the other essential ingredients that we have had to evaluate separately in the past – VFO stability, sensitivity, frequency coverage, transmit spectrum and mode – to name a few.

The fantastic IMMDR performance of the K3 is largely due to a 1st strong mixer and narrow filtering available in the 1st IF.  The Yaesu design suffers by comparison primary because the difficulty in making similar narrow filters at 70mhz is an expensive and difficult task.  Normalizing for IF frequency differences, a 500 Hz 1st IF filter for the Yaesu would have the same bandwidth/center-frequency ratio as a 60 Hz 1st IF filter for the K3.  While the details vary, the fundamental weakness of the design is that the rig's fundamental performance is dictated by the capability of the hardware design – filtering and mixing stages – and will remain unchanged for the life of the rig.  Regardless of subsequent improvements in the DSP and rig control code the manufacturer may provide.

**SDR Rigs**

-----------------

The magic of the SDR comes in two parts. 1) The hardware required to implement the SDR unbelievably simple. And 2), because the signals are processed entirely in the digital format inside the computer, the performance - considering the relatively sparse hardware requirement – can be as good as any superhet.

The most common and basic type of SDR implementation is the QSD or "phasing" method - made popular by Gerald Youngblood in his 2002 QEX series. The prototype highlighted in that series became the Flex SDR-1000 and was the basis for creation of Flex Radio. The QEX articles introduce the reader to the basics of SDR and is recommended reading. Links to this series, and other resources, follow at the end of the article.

In the QSD SDR implementation, RF is sampled and converted to an analog baseband AF signal in a single step. The resulting "IQ" output can be demodulated and filtered entirely within the host computer - without further hardware. The block diagram of the QSD-based SDR looks like this.

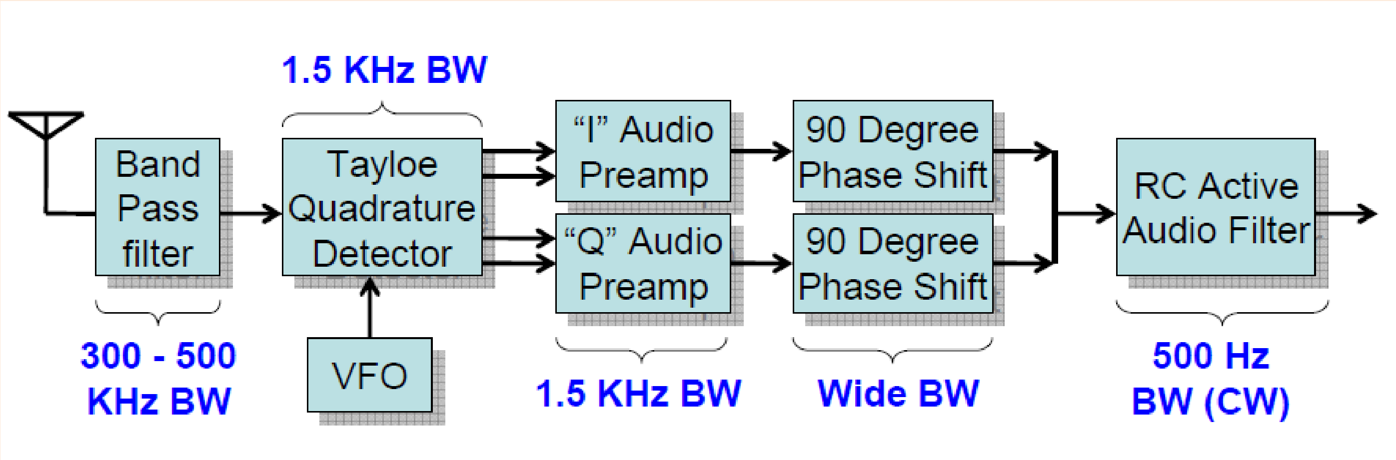


The mixer or detector (terms vary) that performs this step is frequently called the "Tayloe" mixer/detector (generically called "QSD") and was popularized by Dan Tayloe, N7VE. Key features of the QSD are:

1. The output of the QSD is the baseband in a format called IQ – where the identical signal is presented as a dual phase signal: The I or in-phase and Q or quadrature (90-degree phase shift).

2. The circuitry that performs the QSD is quite simple – and is essentially an analog sample-and-hold circuit. This diagram below is an excerpt from the QEX article (part 1) and shows the basic function of the QSD. As we will see later, the circuit can be further simplified from the 4 capacitor design here, down to just 2 with little change in performance capabilities.

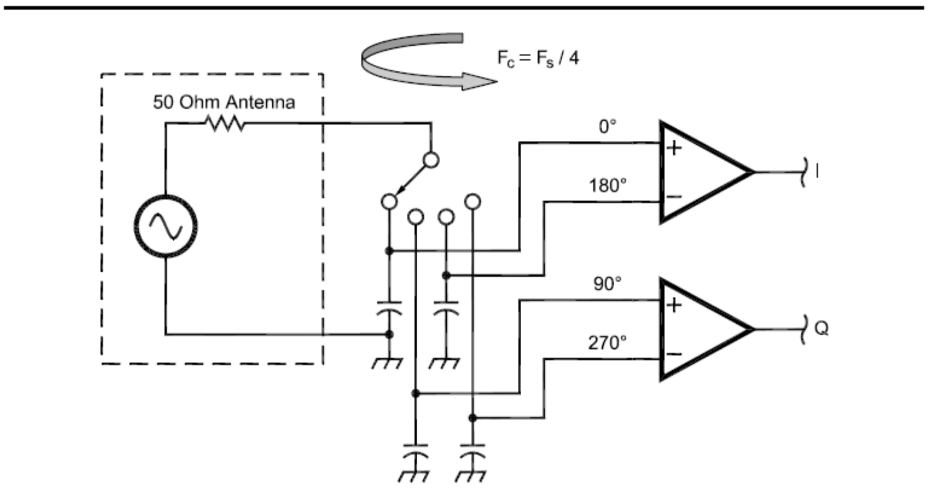


**Fig 10—Tayloe detector: The switch rotates at the carrier frequency so that each capacitor samples the signal once each revolution. The 0° and 180° capacitors differentially sum to provide the in-phase (*I*) signal and the 90° and 270° capacitors sum to provide the quadrature (*Q*) signal.**

The simplicity of the QSD contrasts to it's excellent performance. A typical high-end mixer will have a conversion loss of about 6-8db, meaning a best-case NF of about 7-9db. NF is important because it provides an estimate of how much noise is introduced by the rig alone and provides a floor limiting the rigs performance (higher NF → worse performance). The QSD, by contrast, has about a 1db conversion loss. Which means a QSD implementation may have an entire "front end" NF of about 2db. A typical NF of a multiple conversion superhet may be 20db+.

A side benefit of the QSD is a natural filtering action that occurs at the frequency of operation. With careful attention to I/Q balance, variations of this QSD circuit are capable +20 IMD3 (20khz) in practice with no input preselector filtering required. Commercially available mid-priced SDR exhibit typical IP3 points in the +20 to +50 range. Performance in the analog world beyond and IP3 of +30db is very difficult to achieve and generally is found in only the most premium of rigs.

With respect to the IMMDR metric discussed earlier, all SDR rigs exhibit a flat IMMDR response. That means the IMMDR performance is unchanged regardless of the signal spacing… For example, the Flex 5000a has an IMMDR of about 95db – and that is valid at 20khz, 5khz, 2khz or 100 Hz – any signal spacing. The traditional superhet design will see a huge drop in IMMDR performance as the two signals and their IMD products are narrow enough to fit within the passband of the roofing filter.

SDR Software

------------------

The IQ signal is fed through the ADC and is now in a format for the computer to use. Let's take a look at some of the most common SDR software programs. We are going to skip the explanation of how the SDR program works internally in producing these results – and instead focus on the programs basic capabilities - in order to understand how the SDR programs are typically used. If you are interested in learning how the software actually performs its various functions – and the mathematical theory behind the DSP processing – please see the references at the end of the article.

Probably the granddad of the modern SDR programs used in the ham world is PowerSDR. This is the software interface Gerald Youngblood built to run the SDR-1000 as detailed in the QEX series. From part 3 of the QEX article, we see this block diagram of the basic signal processing operation. We are going to skip entirely the details here other than to point out a few key points:

1. The IQ signal comes in on the left.

2. Filtering of various types, demodulation and digital AGC happen in the middle.

3. And an audio signal is output on the right side. [There physically exists a DAC just prior to the AF Out but that function is typically performed by the computer sound card.]

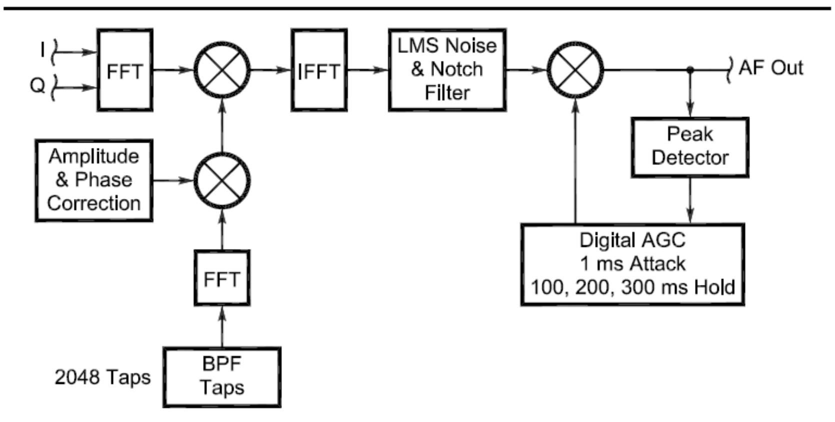


**Fig 1—DSP software architecture block diagram.**

That explanation gets us from the IQ signal to useable decoded audio. This is the basic function of the SDR software.

In practical use, the software needs to provide a lot more adjustment and control features as well. All of the control logic that exists hard-coded into a modern transceiver is handled by the computer as well. Typically by a USB interface that sits on the SDR. The hardware on the SDR related to switching and other housekeeping functions is typically “dumb” in that there is no logic provided. Think of this as a set of relays that perform various switching functions (T/R, preselector, power on/off) – all under the explicit control of the software program.

If you are familiar with the Flex 5000a software, you will recognize the screen shot below. The software is called PowerSDR and is open-sourced (free for public use). PowerSDR provides the user/computer interface for the Flex radios and for many other radios SDR as well.

A list of the more commonly used features and capabilities of PowerSDR include:

- Pan-adapter view showing a segment of the band up to 196khz in width
- IQ balance fine-tuning of the image rejection
- Bandpass filtering including near ideal shape factors - the “brick-wall” effect
- Notch filtering
- Noise reduction of both band noise (hash) and impulse noise of a repetitive nature

- Frequency selection including split operation using keypad entry, or mouse to quickly QSY
- Power on/off and some control of rig specific hardware (enable preamp, select antenna, etc)
- Volume control and audio equalizer to tailor the sound to user preferences
- Record and playback of an entire spectrum of recorded IQ-format
- Additional view options including a waterfall display of the entire band

In this screen capture, PowerSDR is configured to display a traditional bandscope view (top half) and waterfall (bottom half). The portion of bandwidth is adjustable – and here we are looking at about 20 Khz. Note the two CW signals indicated by the yellow arrows in the picture – and their accompanying traces in the waterfall below. The yellow cross-hairs are tied to the mouse – moving the mouse moves the cross hairs. So tuning in a signal can be done simply by aligning the yellow cross hair on the signal indicated on the bandscope, and clicking.


PowerSDR's use is not limited to the Flex radios. For example, I use PowerSDR to provide bandscope functions to my FT2000 – as shown in the picture above. Many of the convenience functions of PowerSDR work with the 2K - including adjustable bandwidth display, signal decoding, filtering and noise reduction, as well as mouse-click instant QSY.

Many other programs exist with differing capabilities. Spectraview, provided by RFSPACE is optimized for bandscope applications and is quite popular as well:

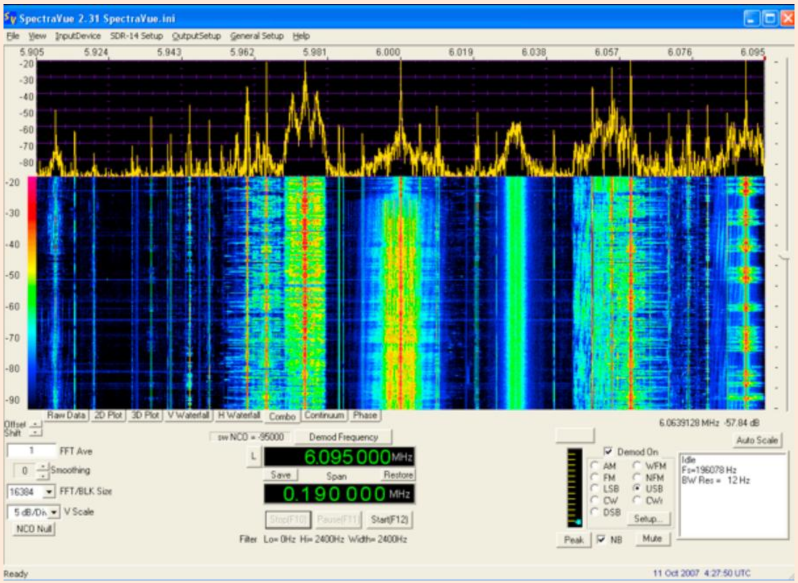


Rocky, by Alex VE3NEA of CW Skimmer fame.  Rocky is a powerful program with a simple interface suitable for SDR hardware that is receive and transmit capable as well as for bandscope applications:

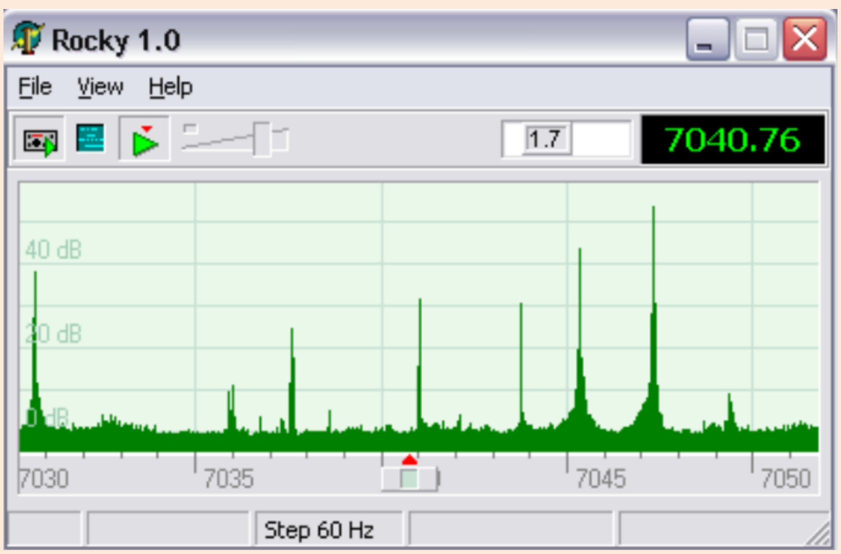


The importance of the SDR computer program cannot be overstated.  The key point is that these programs perform the similar decoding, filtering and control functions.  And many applications are functionally interchangeable.  So for a given SDR hardware platform, the user may have their choice of interfaces and is able to select the one must suited for their particular operating style.

And because the programs are in a constant state of improvement, the SDR's capabilities and functions continue to evolve on a daily basis.  Because the personality and capability of the SDR is completely defined by the software, updating the SDR program essentially gives you a "new radio" with a software download in every traditional measure. That extends beyond improvements in function, DSP capability or bug fixes – to areas far more difficult (or impossible) to alter in traditional rigs. For example, if a new mode – let's call it "FM CW" were created, SDR users would have this ability available from their hardware as soon as the computer program were updated to include this new mode.

Now let's take a look at the Flex as we move back from the topic of software and back toward hardware…

**SDR Example: $3000 Flex 5000a**

-----------------

Recall the QSD described earlier – and here notice the similar QSD circuit found at the heart of the Flex 5000a (below). The same combination of the basic QSD design and the computer software package are entirely responsible for the excellent performance of the Flex.

The essential part of the rig is the preselector and QSD enclosed in the red box. The ADC (analog to digital conversion) is marked as "A/D Converter RX1" is integrated into the Flex 5000. In most lesser-cost implementations, a computer's sound card provides this A/D conversion function.

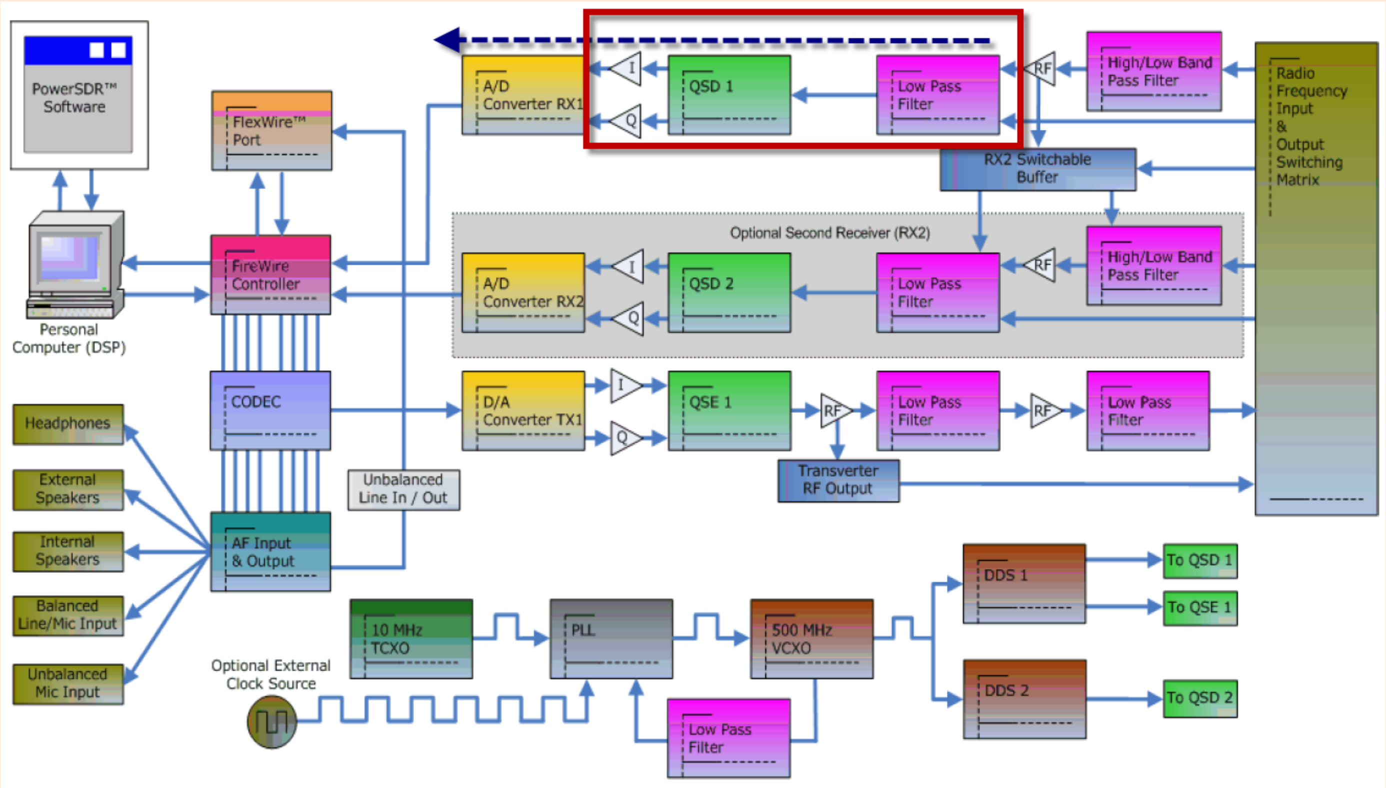

**SDR Example:  $13 Softrock SR-II**

---------------------------------------------------

As an example of what “unbelievable simple” is, we cite the $13 Softrock SR-II kit.  This single band crystal controlled receiver schematic is shown below contains 4 main functional blocks.  Construction involves mounting a few SMT parts and takes about 2 hours to complete in total.

Key circuit features include:

1. The clock (green) is a colpitts oscillator with buffer following, driving a high-speed 7474 flip-flop.  The outputs of the flip-flop provide two clock signals, separated by 90 degrees, necessary for the sampling function of the QSD.

2. The heart of the rig is the QSD – in blue color.  The QSD is implemented here with a simple CMOS switch, two capacitors and two resistors.  The RF is directly sampled – the two 0.047uf caps hold the sampled voltage level until refreshed on the next cycle.

3. The AF amp (red) provides gain and buffering for the small voltages present on the sample-hold caps.  The output drives the computer sound card.  The sound card provides the final hardware step – an analog to digital conversion.  Once in digital format, the computer software provides all the traditional receiver features.

4. The preselector (yellow color) provides a dual pole filter and impedance matching function.  The purpose of the L/C on the preselector is not for selectivity – but to suppress very far out of band, very strong signals that may be present in the environment (e.g. commercial AM broadcast service).  It provides no significant selectivity.  The rigs selectivity is a function of the QSD and the computer software.

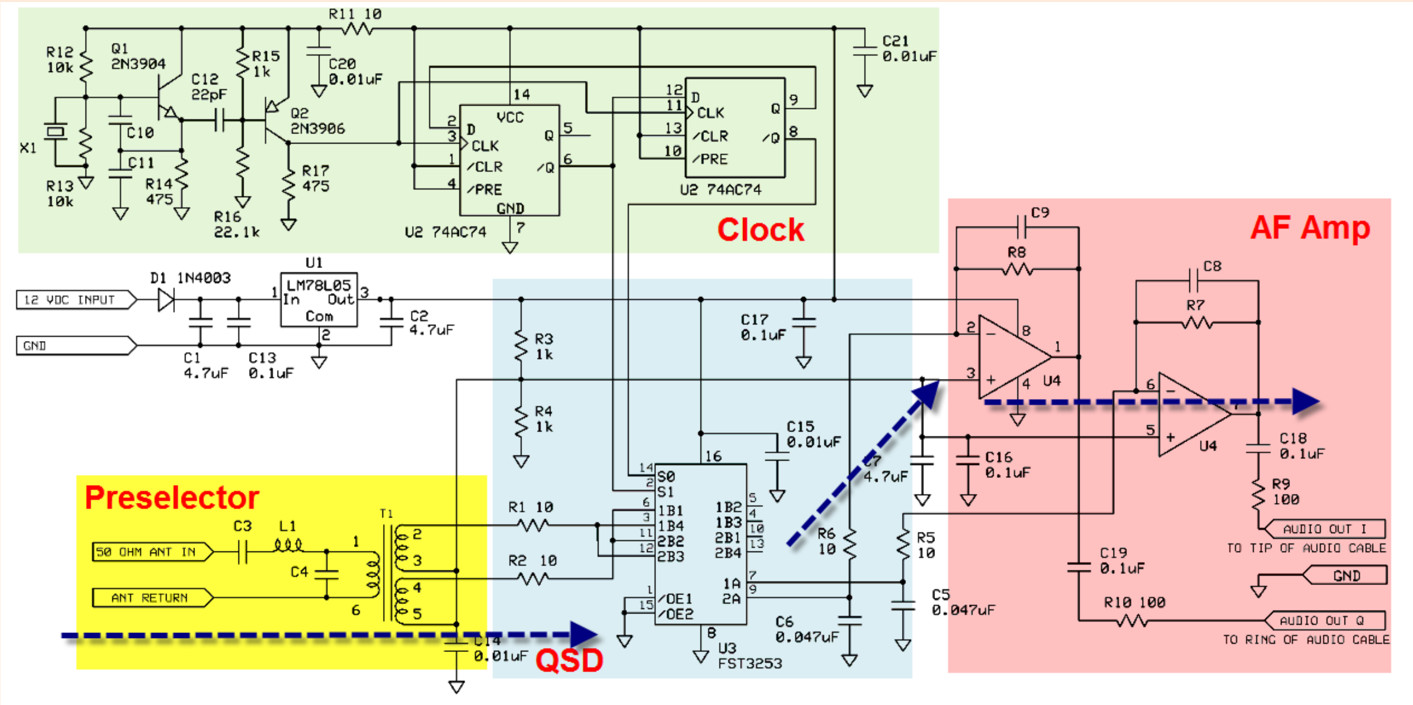

Below we see a Softrock SRII enclosed; the small square PCB on the left is a preamp.  The bottom side board photo (below) shows the SMT component locations.

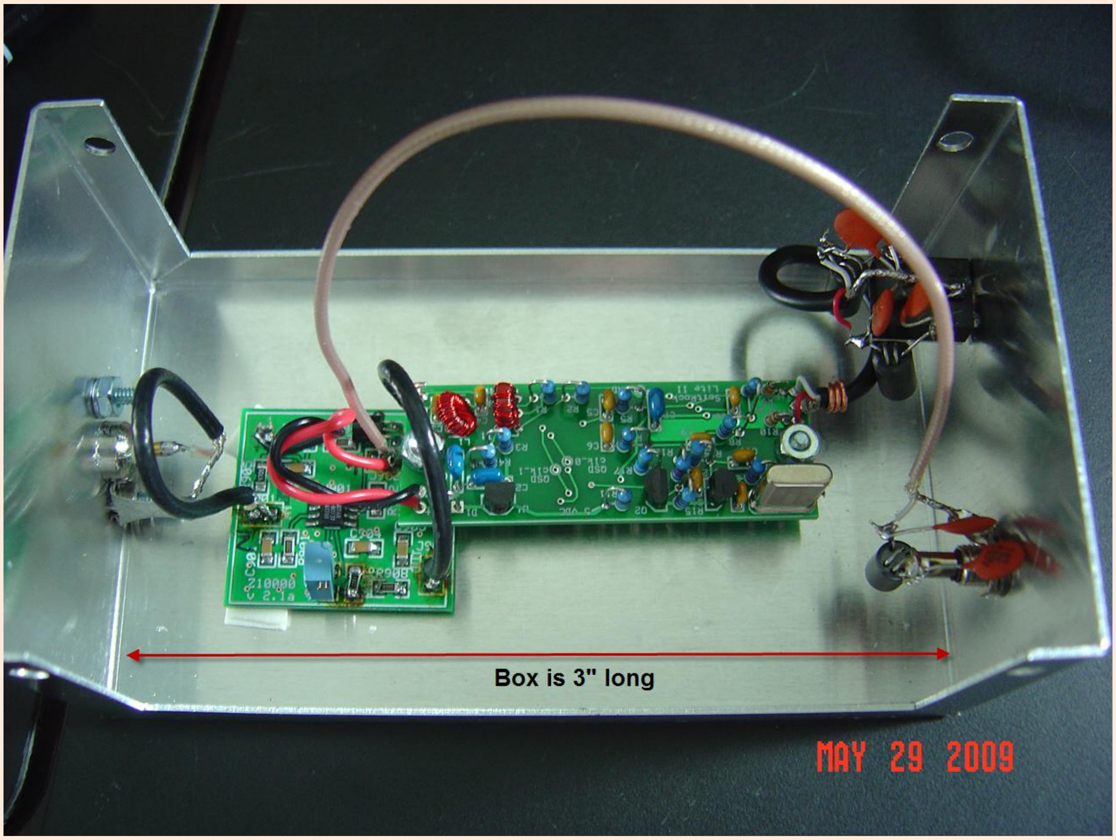


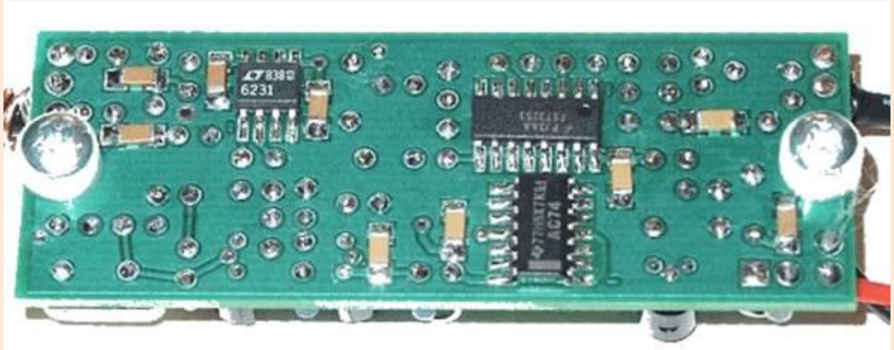

This SRII is tuned for the 10.55mhz IF of a FT-2000 and provides bandscope service.  The displayed bandwidth and sensitivity (MDS) depends on the sound card it's matched with.  In my application, the bandscope covers about a 100khz spectrum, and the minimum observable signal is better than -117dbm.

The MDS (minimum detectable signal) of the Softrock depends on the strength of the signal - and the noise level of the A/D conversion (the computer sound card and surrounding computer EMI).  It's less costly to buy or build a wide-band preamp - compared the generally higher cost of specialty sound cards or the difficulty in quieting down the EMI levels of a computer.

**SDR Example:  $85 Softrock v9 BPF**

---------------------------------------------------

The second example of the simplicity of the SDR circuitry is found in the Softrock v9. This rig is an all band, all mode receiver capable of 160-10 or 80-6m coverage.  The complete kit requires perhaps 8 hours of assembly time.

The block diagram (right) is very similar to the SRII, with added capability of an on-board VFO, switchable band-pass preselector and a USB connection for rig frequency control from the keyboard.

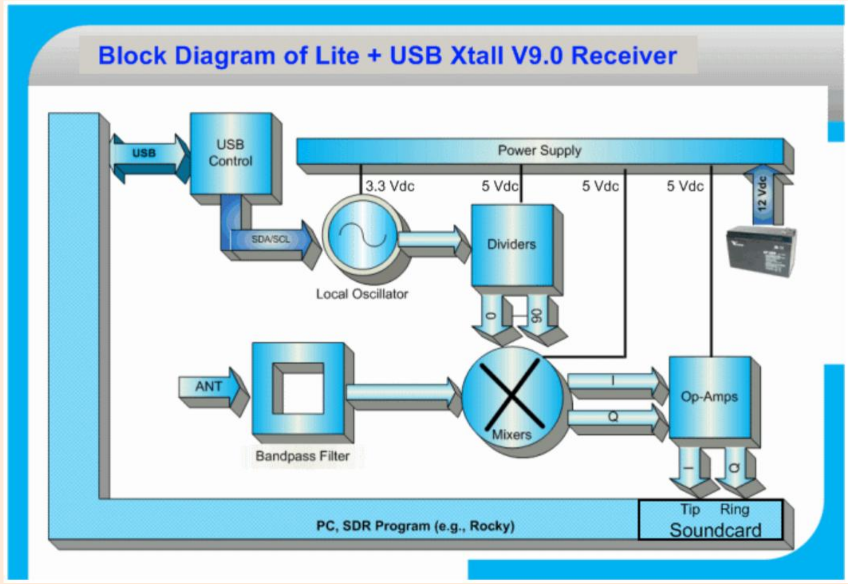


The Softrock v9 schematic here is very similar to the SR2:

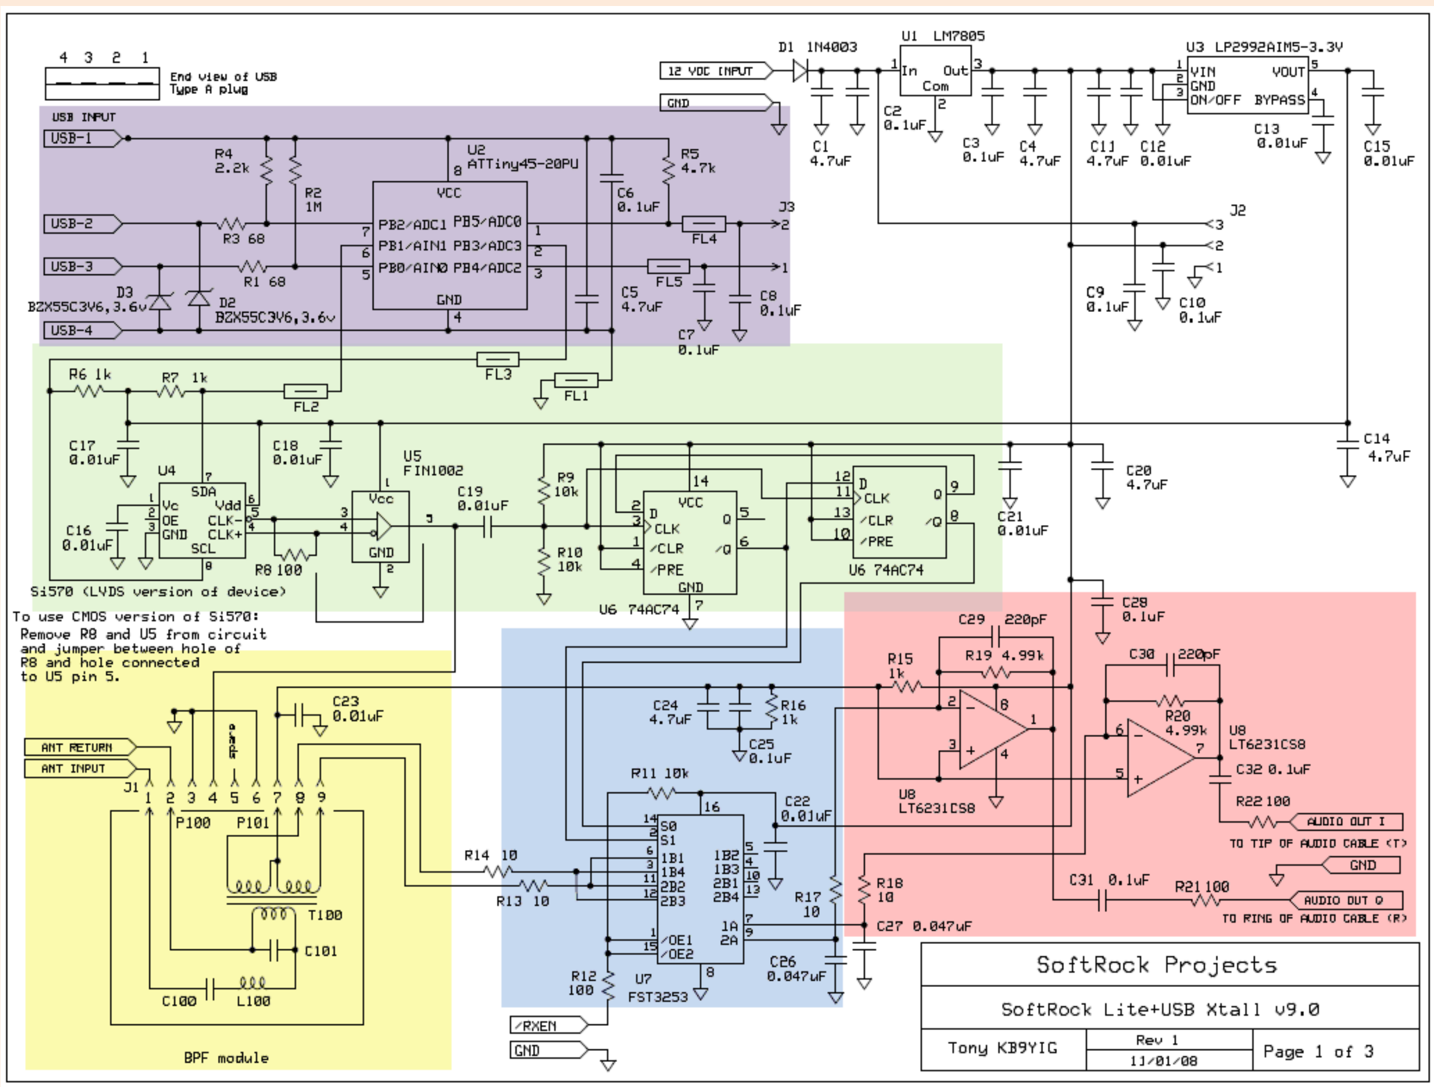

In green is the clock circuit – replacing the crystal oscillator and buffer is a single-chip direct-digital-synthesizer (DDS) chip - the Si570 - a $20 marvel of modern technology.  Just 5x7mm in size and drawing 80ma, the Si570 chip provides a programmable frequency output covering the 0.1-160mhz range - flat in amplitude to within 1db – and with a nearly perfect 50:50 output waveform duty cycle.  Jitter and phase noise spectra are nearly indistinguishable from the “benchmark” of ham signal generators – the HP8640B.

The DDS frequency is set via USB connection to the computer.  A programmable microcontroller (purple) provides the USB interface and in turn sets the frequency of the DDS and control lines for the preselector.

When matched with a small daughter card containing 5-pole L/C filtering and SMT CMOS switching, this Softrock v9 becomes a true frequency agile receiver covering all of the HF band – fully under control of the computer.  Pictured below is the Softrock v9 in an enclosure the same size as that for the SRII – about 3” in length.

Performance is very similar to the SRII.  The example shown below (with preamp mounted on the sidewall) has an MDS of about -129dbm.  It’s fully portable and self contained running entirely off the power provided by the USB connection.

I have further modified this unit (not shown in this photo) to provide a -73dbm (50uV S9 signal standard) output from the DDS output which allows me to use the rig as a frequency and signal strength standard providing a low cost substitute for the lack of a HP8640B of my own.

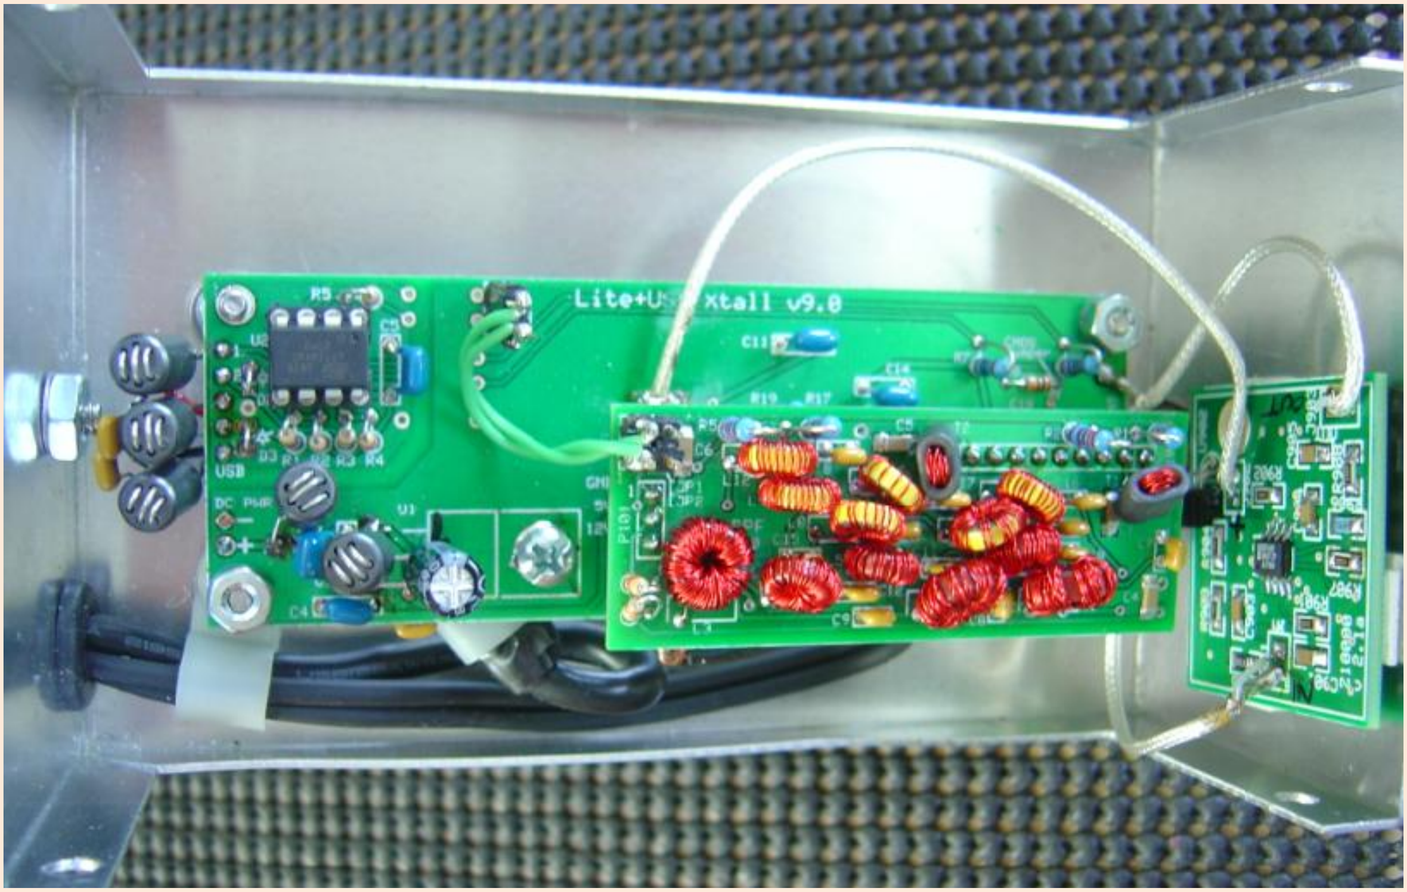

ACØC CW Skimmer Array

-------------------------------

The low cost and high performance of the SDR technology opens opportunities that are not practical from a traditional hardware-based platform standpoint. An example of this is the CW Skimmer Array.

Here, we have combined 10 of the SRII SDR receivers into a single system that provides CW Skimmer sampling of all bands, 160m-6m inclusive.

In the photo below can be seen the individual SDR modules mounted in two rows alone either side. DC supply is the wiring from the top center and down along the inside of the boards. AF output is carried along the wires on the outside of the row of boards.

RF signal input for all bands come from a single antenna terminating through the coax connector at the bottom of the picture. The incoming RF is fed into a preamp module (red arrow) based around a Mini Circuits +31db NF=2.7 single chip wide-band amp. The red square insert shows the bottom side of the preamp board, and the yellow arrow is pointing to the MC amp chip. It’s about 1.5mm in diameter!

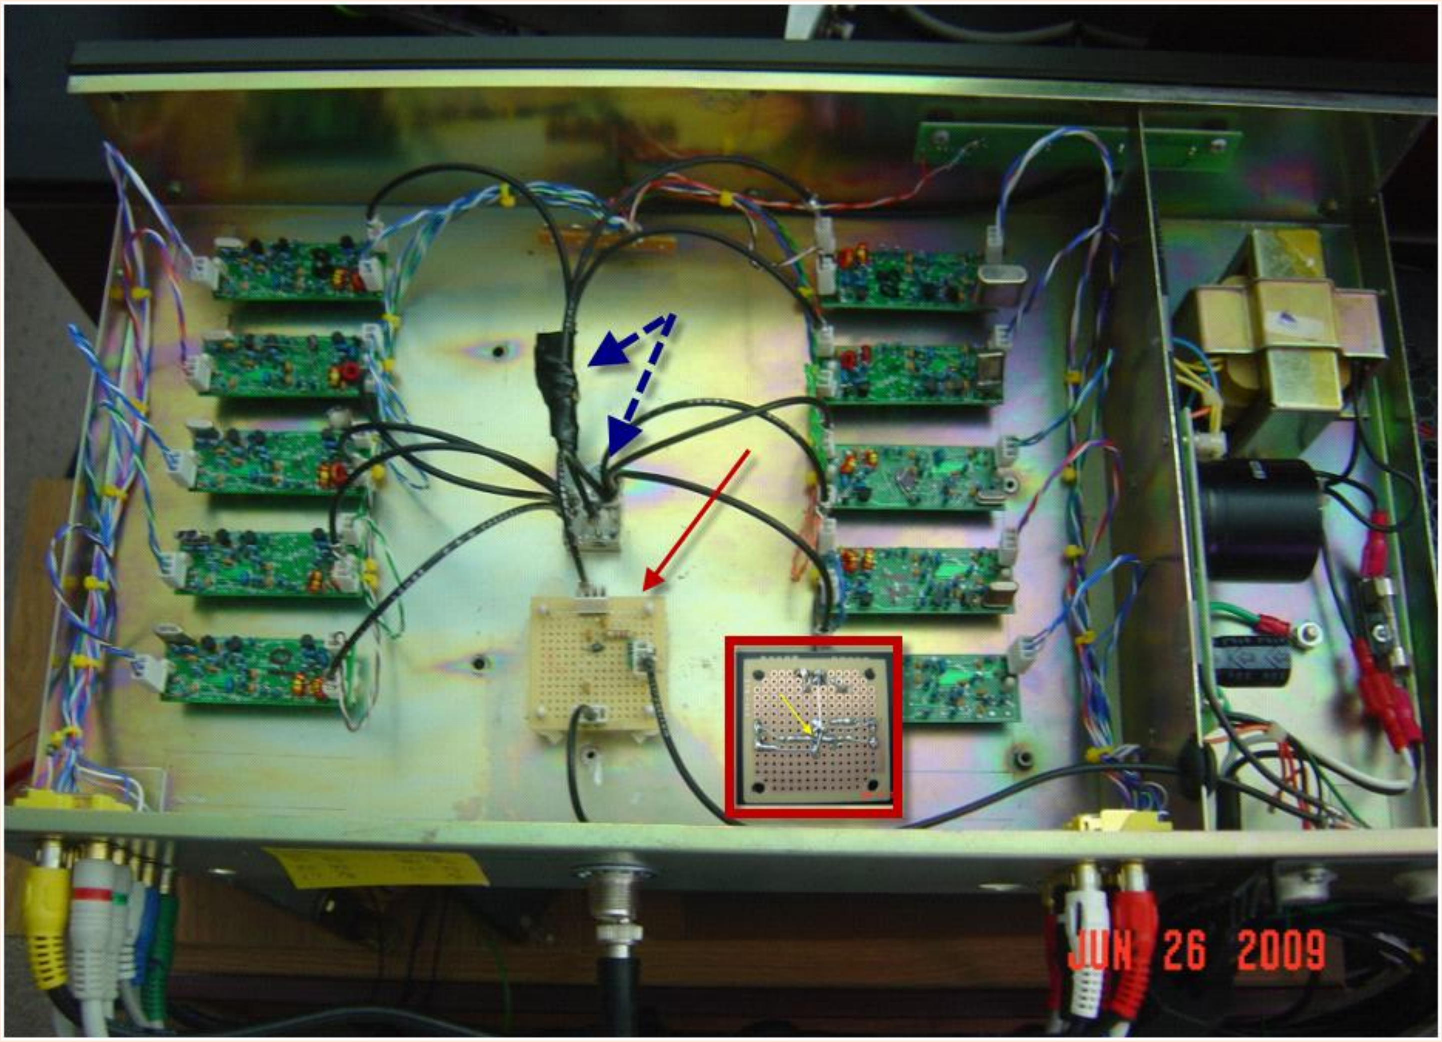

The output of the preamp is fed into an 8-way Mini Circuits splitter – one output leg of the 8-way is further is further split out into 3 additional lines serving the 160/80/40m bands. The 2nd series splitter adds an additional 6db or so of insertion loss - but given the high band noise levels encountered on these bands, the impact on sensitivity is insignificant.

The power supply (right side) was part of the original NABU enclosure. It provides 18V @ 1A and 5V @ 1A using a power zener regulation circuit. The only power supply modifications required were the addition of a 9V 3-terminal regulator mounted against the back wall to provide the proper voltage for the preamp bias. This minimizes the heating of the 78L05 mini-regulators located on each of the SDR boards.

The center frequency of each SDR is set at about 45khz from the bottom end of each band. Coupled with a bank of 96 Khz sound card inputs, the Skimmers provide nearly 100 khz of coverage per band. Crystals were sourced from Surplus Sales in some cases – and for the boards where no commercially available crystal could be found (15m, 17m and 6m), small self-contained single chip oscillators were used. These chips are cheap ($6) but have relatively high phase noise levels and that can exhibit reciprocal mixing under some strong signal conditions – especially given the high gain preamp at the head. However, the CW Skimmer software is not at all bothered by the added band noise - should I ever encounter a strong signal on these bands.

The 10 instances of CW Skimmer software is split between a Quad core 3.4ghz and a dual core 3.6 ghz machine and loads each computer at about a 50% CPU utilization level. On extreme signal population conditions like field day, the CPU load can reach 100% requiring me to take a few of the less interesting bands off-line temporarily. The magic of SDR and CW Skimmer is only made possible by a LOT of computer horsepower.

Initially I had plans to build an i7-based computer system which could accommodate the three sound cards and handle the CPU loading requirements of the entire array all in one computer system. However, the author of CW Skimmer Alex has implemented a “virtual” skimmer system hosted in the FPGA of the QS1R direct-sampling SDR which is quite interesting. All of the DSP signal processing and CW decoding is done in hardware for the 7 simultaneous skimmer sessions - and the output via USB from the QS1R is a single telnet feed – ready for tapping by your favorite spot management utility. The appeal of that solution is that it requires very little computer resources by comparison to the individual SDR/individual Skimmer session solution. And the combined cost of the QS1R and the software is about the same as the computer needed to host the 10-band skimmer.

| Band | MDS dbm |
|---|---|
| 160m | -128 |
| 80m | -114 |
| 40m | -131 |
| 30m | -135 |
| 20m | -127 |
| 17m | -119 |
| 15m | -140 |
| 12m | -138 |
| 10m | -137 |
| 6m | -118 |

The Skimmer Array is almost complete and the receiver sensitivity is quite impressive. A few changes with the canned oscillator to a low noise variety are all that remain to complete the project.

High sensitivity is a necessity in this application - especially on the high bands where the band noise and typical signal strength is quite low – because the antenna is not resonant on these bands. And in some cases, an antenna tuner may be inline adding another 10db drop in signal strength, at the worst case. However, even at these reduced signal levels, the MDS is generally below the band noise floor ensuring that the skimmer array will hear all the signals present on the band.

The 10-skimmer sessions feed into the Wintelnetx application created by Dave K1TTT.  This provides a single consolidated feed point for the spot processing program.  And the K1TTT program has dupe checking as well as the ability to filter out domestic vs. DX calls.  Here's a screenshot with 9 skimmer sessions running…

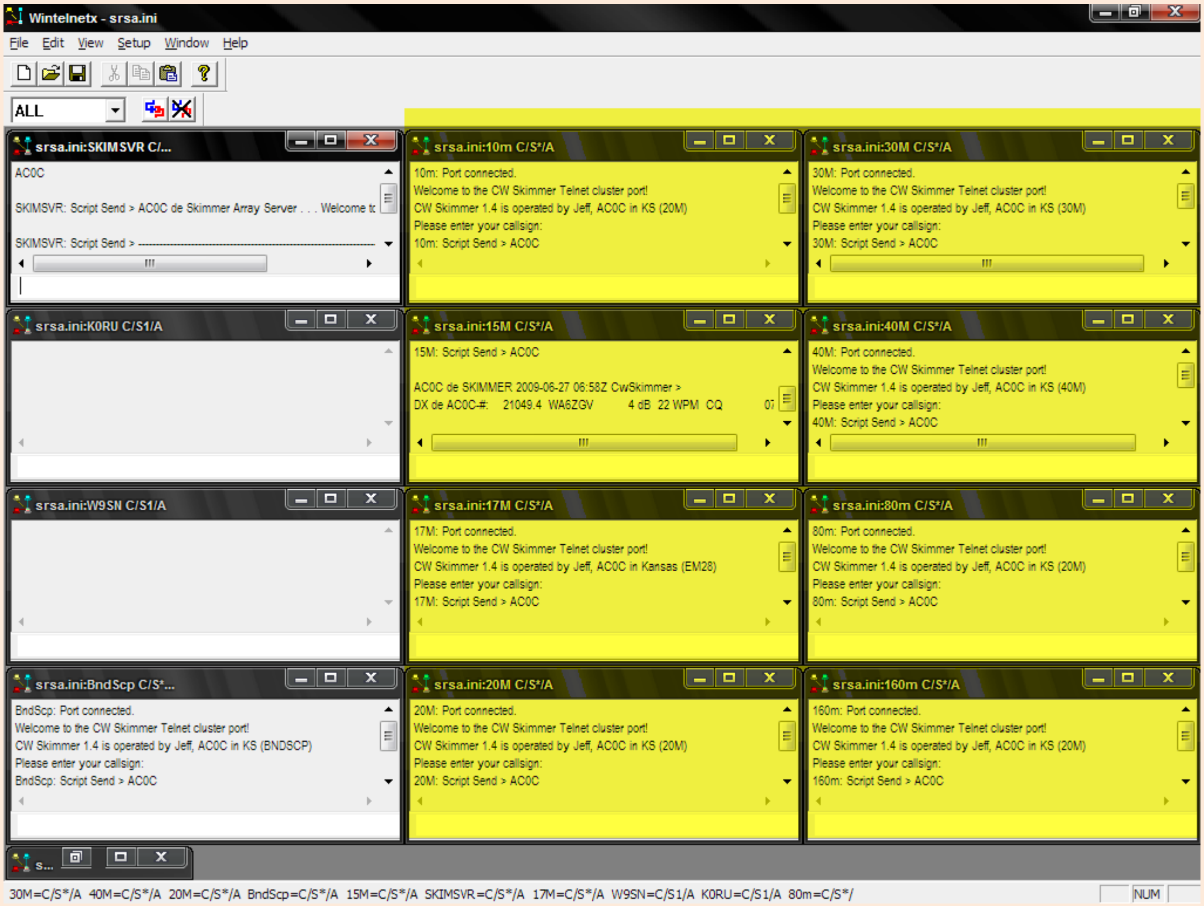


On field day the array was serving spots at up to 200 per minute - and that was with only 6 bands functioning at that time.

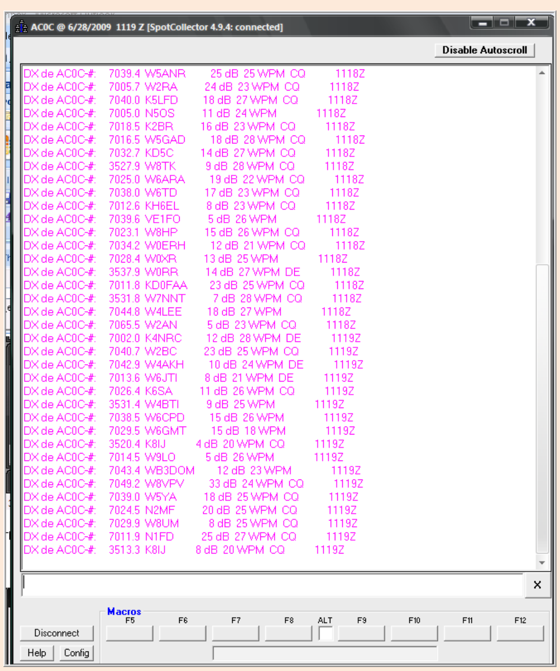


**SDR:  Is All that Glitters Gold?**

-----------------------------------------

In the practical implementation of the less complex, lower cost SDR implementations, a couple of practical issues pop up that keep us from easily realizing performance that would in fact beat a good superhet design.

The major considerations include:

1. IQ balance.  Oscillator phase noise and jitter - as well as variations in the QSD sample hold circuit - result in very minor differences in the IQ balance.  The levels are not EXACTLY the same, and the phase relationship is not EXACTLY 90 degrees.  And optimizing them in hardware is not simple.

However, in the spirit of the SDR theme, these balance differences can easily be accommodated by the computer software. So while the hardware balance may not perfect, the computer can make it nearly so giving the operator final performance that is far better than it is in reality.

For example, the typical image rejection of the Softrock is only about 25db (uncorrected by software), when measuring the hardware alone. And when corrected by software, the image rejection improves to the 70-100db range ensuring that we will not see any image on our bandscope except for the ghost of a gigantic input signal.

2. The ADC function served by the computer sound card is probably the weakest point of the entire system. Modern computers have much improved sound systems – but the bandwidth capability of these sound cards is tailored for audio – and that generally means a 20khz upper limit. A 20khz sound card upper frequency means we will only be able to see about 40khz of bandwidth spectrum.

   The good news is that the noise level of these built-in cards has improved dramatically in the last 10 years and the typical built-in sound card can provide quite adequate service for bandscope applications where it's quite common to look at a 40khz or smaller spectrum. In the bandscope application, the signals for a crowded band may be too closely spaced on the screen to allow the use of wide bandwidths in practice. Higher performing sound cards are available at lower costs now as well – and some very excellent cards are easily available on eBay for reasonable cost. The Soundblaster Live 24, for example, provides 96khz of spectrum visibility and can be had for around $15 delivered.

3. Dynamic range is dynamic range... For use in extreme signal cases, the same attention must be given to the component selection in the QSD as with the entirely hardware based rig. The signal range between the typical rig's MDS with a preamp activated - and a full-scale S9+60 indication - is well over 130db. That means the antenna may see ranges from 0.04uV (-135dbm) all the way up to perhaps 0.25v (0dbm).

   On the ADC conversion end, getting that huge dynamic range requires very high performance chips (read that as expensive chips) and extreme attention to circuit layout. The very best add-in sound cards may reach about 110 db of dynamic range capability.

   Solving this is a mater of design preference. Implementing some form of hardware AGC in the IQ chain is one option. Another is to decide in advance what the typical operating conditions will be – favoring either the strong end or the weak end of the signal range – and optimize the hardware for that. And another is to consider that MDS alone is not a very useful metric given practical band noise levels. My 80m skimmer SDR has a -117dbm sensitivity – however, with typical band noise levels of around -100dbm, the high sensitivity is not needed. For the lower bands, the practical MDS is band-noise limited rather than SDR limited.

There are subtleties and complications of the SDR implementation. However, when compared to the difficulty of traditional receiver design of similar performance, these difficulties are trivial. With all the signal processing being handled in software, the hardware builders challenge is minimal by comparison.

In the SDR world, all that glitters may not in fact be gold – but lumps of coal are few and far between - even in the worst of the SDR implementations.

**The SDR of the Future**

------------------------------

As fantastic as the QSD performs, it's now considered an old technology and is not applicable for new performance-sensitive applications. Replacing the RF to baseband conversion served previously by QSD, these new designs feature high speed ADC instead. With an ADC running at 110 Mhz or higher, the rig can directly read the entire 160-6m HF spectrum in one single step.

Commercial examples of this type of direct-sampling SDR topology include:

| | |
|---|---|
| The RFSPACE SDR-IQ and SDR-14 | http://www.rfspace.com/SDR-14.html |
| QSR1 | http://www.srl-llc.com/ |
| Perseus | http://www.microtelecom.it/perseus/ |

For more performance metrics, see http://www.hfdecoding.com/uploads/SDR_RX_measurements.pdf

The price point for these SDRs are about $1000 point (except for the SDR-IQ). And except for the SDR-IQ (200khz max), all of these receivers have the ability to serve as a 30-mhz spectrum analyzer.

And a fully functional direct-sampling SDR-based ham transceiver is now available from ADAT. http://www.adat.ch/index_e.html

The future will most certainly hold greater and greater levels of merging between the SCR and SDR. Perhaps the next steps will be to take the DSP frequency up so that one additional mixer stage can be eliminated. However its almost certain that the "standard" ham transceiver 10 years from now will be of a direct-sampling nature.

It's a challenging time for rig manufacturers, especially the SDR manufacturers – because the emphasis of the SDR move is more and more toward minimizing the hardware platform. And given most of the software programs for SDR are open-source and available at no cost, it will be difficult to differentiate the product within the market unless the computer portion is moved to a dedicated and self-contained format (e.g. Flex 5000c).

What remains is the user interface issue – the knobs vs. mouse debate. For me, the combination of the Yaesu FT-2000 equipped with a Softrock-driven PowerSDR bandscope provides the best of both worlds from a user interface standpoint. Perhaps this format will be more common in the future. The K3, for example, already provides easy access to the 1st IF output via back-panel jack - making it very easy to fit a bandscope SDR like the excellent N8LP LP-PAN.

Regardless of how the SDR revolution progresses, the ham community will be the beneficiary. The SDR has pushed the bar of performance for any price point very high. Competing rigs will need to provide more performance for less cost than the trends we have seen historically.

What About Transmitting?

----------------------------------

While the discussion above speaks of receivers exclusively, the implementation of the transmit function follows the same principles – just in a reverse order operation.

DAC from the computer provides the analog signal in IQ format. It's merged using the QSD into a RF level output signal. This is followed by traditional methods of amplification, low-pass filtering, switching and etc. See the QEX article on the SDR-1000 for a good overview of the transmit variation using the QSD

Further Reading

---

Radio Receiver Design by Kevin McClaning and Tom Vito

A more comprehensive review of the theory behind the Tayloe QSD may be found here: http://www.norcalqrp.org/files/Tayloe_mixer_x3a.pdf

Brush up on your DSP understanding with the free online book DSP for Engineers and Scientists at: http://www.dspguide.com/

For more on the basics of IQ processing, see http://members.cox.net/k9ivb/files/why_I&Q_radcom_jan07.pdf (basic) and http://www.dspguru.com/info/tutor/QuadSignals.pdf (detailed)

The article series is available for download in 4 parts from Flex at: http://www.flex-radio.com/Data/Doc/qex1.pdf through http://www.flex-radio.com/Data/Doc/qex4.pdf

Flex Radio 5000a user manual, complete with details on the operation of PowerSDR: http://support.flex-radio.com/Downloads.aspx?id=183

Tony Parks for Softrock orders: https://www.kb9yig.com/

Tony Parks 2006 Dayton presentation on Softrock: http://www.tracey.org/wjt/temp/dayton06/tony-352x240-2p.avi

Heathkit-style builder notes for all Softrock variations: http://www.wb5rvz.com/sdr/

Rocky SDR software: http://www.dxatlas.com/Rocky/

Yahoo users group for the entire Softrock series: http://groups.yahoo.com/group/softrock40

Wide Band Image Rejection coming to PowerSDR: http://w9oy-sdr.blogspot.com/2009/05/wbir.html

The HPSDR project – SDR in the extreme: http://openhpsdr.org/

A massive link list for all things SDR: http://planete.inria.fr/SoftwareRadio/