“JOURNAL OF THE CONSTRUCTORS CLUB”

Hot Iron

Winter 2017 Issue E98

Editorial

Have we reached ‘peak’ home electronic construction? Commentators talk about peak oil or the peak of all sorts of industries, but are we near that for radio related enthusiast electronics? I fancy we are! As I have argued many times before, it has never been so easy to build simple electronic projects – gone are the days of nasty metalworking to mount your valves or even stage after stage of discrete semi-conductors to realise a multi-band double conversion superhet receiver, let alone the associated single sideband phone transmitter! The ability to mount parts and integrated circuits on copper clad laminate, with or without etched connecting tracks, makes building even single band one-off transceivers relatively easy. This can now be done literally on your kitchen table with the whole project kept in place for next week on a tray that can be safely stored away ready for the next session. Can it get easier than this – maybe not!

How will the technology progress and will it allow keen experimenters to dabble? Undoubtedly, integrated circuits will become more capable – putting the whole RX into a single chip is by no means novel. At the same time, computing ‘devices’ are working even faster so that they can perform complex mathematical functions on ever higher bandwidth signals. The rise of Software Defined Radio techniques is such that very few commercial HF TCVR designs are anything other than a high performance analogue to digital converter followed by a very high speed processor, with a ‘nearly’ analogue power output stages for audio and RF! Inevitably the chips in such designs are made for the mass market in surface mount format and utilising these in an experimental form is very challenging in many physical ways. On top of that, is the challenge presented by the software which is required to make them perform. Very few of us understand the advanced mathematics behind serious digital processing; and the ability to alter or adapt that to suit ones own experimental ideas is minimal unless the original designer built in the software ‘knobs’ to permit such experiments - pretty unlikely if it was commercially developed! Hence I would argue that from now on it will become increasingly hard to use this sort of advanced technology.

So by my assessment, it has never been so easy to build moderately complex electronics and it is likely to get harder in the future – conclusion – we are at ‘peak electronics build’! I hope I will be proved wrong by one of you reading this – please write and tell me so, with a contribution for Hot Iron. Meanwhile, I say, ‘get on and enjoy it while you can’!!

Happy Christmas and a healthy New Year to you all! Tim G3PCJ

Contents

Hot Iron is published by Tim Walford G3PCJ of Walford Electronics Ltd. for members of the Construction Club. It is a quarterly newsletter, distributed by e mail, and is free to those who have asked for it. Just let me know you would like it by e mailing me at electronics@walfords.net

Kit Developments

Too much of the last three months has been spent on erecting buildings for the farm and starting the process of building a new house – it takes much effort to get through all the regulatory procedures! Most of my electronic time has been devoted to getting the prototype Beer RX (right) working – see later for some circuit snippets. It is a multiband design (in its simplest form shown right, for any 2 of the 4 bands 20 - 80m). The complexity of its analogue local oscillator scheme just shows up the advantage of the modern DDS or phase locked loop micro-processor driven designs. In this photo you can see four double tuned band pass filters for the RF input, and Local Oscillator mixer output, on both of the two bands of this basic form; when extended for the other two bands in this group, the RX needs another four filters (as well as the unavoidable two (or 4) low pass filters for the transmitter!). The days of complex analogue multi-band designs are drawing to a close – this might be my last in this style!

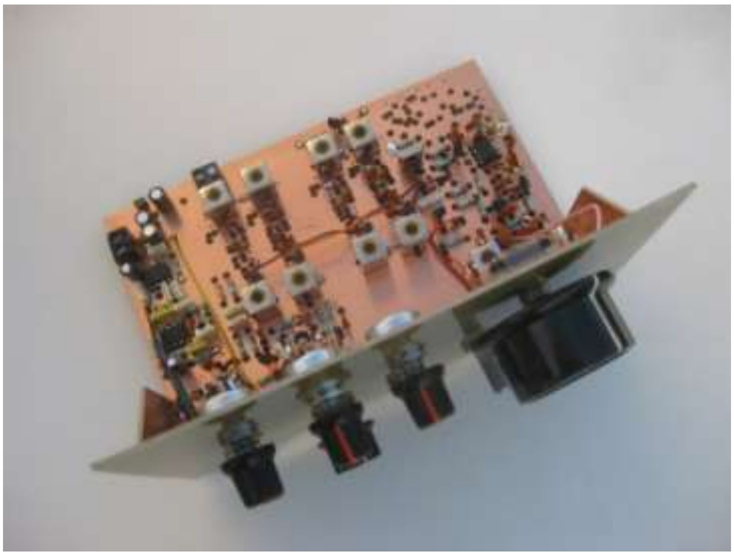

The next article **'VFOs and all that'** comparing the merits of different of techniques to produce stable RF signals (mainly in a frequency sense) is highly relevant. Pete Juliano N6QW threw down the gauntlet in the last issue of Hot Iron, so I sought a defence of the traditional HF analogue approach. Peter Thornton G6NGR took on that task but says the days of complex analogue 'VFO's are over! I don't entirely agree - the analogue approach has considerable merit for very simple single band designs such as those two ideas outlined below – mainly on a 'performance for cost/complexity' basis. This has led to the following doodling suggestions:-

The **Sheppy** might be a simple single band 1.5W CW DC transceiver based on a ceramic resonator for 80m, or a crystal for the higher bands to avoid drift/chirp. This is the sort of small simple rig that should make it possible to take out and about without trailing a large suitcase for all the extras! It should fit easily into the small upright format with a minimal PCB front panel. Another 'use' would be for demonstration of simple radio techniques across relatively short distances (half mile or so across playing field) when visiting schools. Target price is to be under £40. The name Sheppy is a river in Somerset, & continues that 'theme' for such simple CW rigs!

The **Quantock** might be a simple DSB phone TCVR; again with a DC receiver and using an 80m ceramic resonator. For the same reasons as above, crystals are needed for the higher bands. For those wishing to avoid the wonderful simplicity of phone using ancient Amplitude Modulation, the next easiest technique has to be double sideband modulation with a suppressed carrier; which is of course fully compatible with stations operating conventional single sideband phone. This would be an update of the very successful Brendon DSB phone TCVR that was used by many Buildathon groups. Again the target price is under £40 and it should fit into the small upright format easily. The Quantocks are a small range of hills (near the Brendon Hills) in north-west Somerset – which is the theme I have adopted for my DSB designs!

VFOs and all that – *Peter Thornton G6PNR*

Introduction

This is a comparison of the “Variable Frequency Oscillators” available to amateur constructors: L /C circuits and mixer VFO's; Pulled Ceramic Resonators; analogue with digital error correction (Huff and Puff and similar drift correction); pure digital (synthesizers and DDS methods). I haven't considered crystal oscillators or VXO's as they are different beasties; besides, the frequency you can pull a crystal is - - - not a lot!

My criteria for the VFO standards are: (1) Cost; (2) Construction difficulty; (3) Drift performance; (4) Purity of output note (harmonic content, spurs, noise sidebands, phase noise, micro-phony); (5) “RIT” for receiver local oscillators. I have split the Amateur RF spectrum into sections, as VFO's for each section face different demands: (1) LF, 132kHz - 1MHz; (2) Low HF, 1MHz - 5MHz; (3) Mid HF, 5MHz - 14MHz ; Hi HF, 14MHz - 30MHz; (4) Low VHF, 30MHz - 70MHz; (5) Hi VHF, 70MHz - 200MHz. Above 200MHz (realistically) only synthesizer or DDS methods fit if you want full VFO functionality; or for microwave SSB you're looking at temperature compensated crystal oscillators feeding multiplier chains - or the very esoteric DDS technologies currently appearing that have internal spur and noise cancelling. ECL gates can be used for GHz Huff-n-Puff stabilisers, but these are mostly experimental at this time.

The following tables summarise the results I've found over the years.

L/C Variable oscillators

	LF	**Lo-HF**	**Mid-HF**	**Hi-HF**	**Lo-VHF**	**Hi-VHF**
Cost	Mid-High	Mid-High	High	n/a	n/a	n/a
Construct.	High	High	High	n/a	n/a	n/a
Drift	Low(ish)	Low-Mid	Mid-High	n/a	n/a	n/a
Purity	Good	Good	Good	n/a	n/a	n/a
RIT	Possible	Possible	Possible	n/a	n/a	n/a

Cost is high, as is hassle in explaining why you've spent $'s, £'s, Euro's (delete as appropriate) on that superb Jackson variable capacitor, the oh-so-rare Oxley “Tempa-trimmers” and the beautiful analogue tuning scale and reduction gearbox you'll need to build and use the beast.

Construction is difficult in that you'll be milling solid blocks of aluminium, annealing wire, toroids, and other components; changing capacitors for the next week, using a hot-air gun and freezer spray knowing that if you take it outside you'll have to start all over again.

Drift - note that LF and Lo-HF L/C oscillators are intrinsically stable - this probably means that drift is a complex function of frequency, supply volts and temperature. They are fine for most modes below 2MHz; however, for WSPR, QRSS and similar esoteric modulation, you'll need a more stable oscillator, perhaps by digitally dividing a higher frequency oscillator. A rule of thumb: if your 5MHz L/C oscillator achieves a drift less than 20Hz / minute, it's good; if it drifts less than 5Hz / minute it's amazing.

Purity of the output depends very much on the oscillator having correct biasing, stable power supply(s), lack of vibration and being buffered properly (amongst many other variables). The big advantage of an L/C oscillator is they don't generate “spurs” or many noise sidebands.

RIT can be done by switching in additional capacitors to shift the oscillator slightly. Sudden switching causes "chirp" and harmonics. RIT applied by manually adjusting a small value trimmer capacitor is not as prone to "chirp", but the oscillator will take time to settle and this can produce "warble" notes.

Range is set by the L/C ratio, but as the oscillator components age and/or transistor gain drops, L/C VFO's can suddenly stop oscillating for no apparent reason when tuning. Some ageing L/C VFO's have a nasty habit of stopping at odd frequencies whilst tuning across the range, but starting again once the tuning stops! In this instance, try looking for dried up electrolytic decoupling capacitors.

Mixer type L/C VFO's

	LF	**Lo-HF**	**Mid-HF**	**Hi-HF**	**Lo-VHF**	**Hi-VHF**
Cost	High	High	High	High-Tricky	High-Painful	High-Painful
Construct.	Complex	Complex	Complex	Complex	Complex	Complex
Drift / min.	Low-Mid	Low-Mid	Low-Mid	Low-Mid	Low-Mid	Low-Mid
Purity	Good/Excel	Good/Excel.	Good/Excel	Good/Excel	Good/Excel	Good/Excel
RIT	Possible	Possible	Possible	Possible	Possible	Possible

Cost - The oscillator mixes a VFO with a fixed crystal oscillator, and the sum or difference frequency is filtered out. You get the adjustment range (and drift....!) of a VFO, plus stages of filtering for the wanted output without all the unwanted mixer products. This costs money, space and power, and it takes a very subtle design to minimise the oscillator signals and mixer products appearing on the desired output as phase noise and sidebands. It CAN be done, but it needs expertise and experience to get the design right (if you're lucky).

Construction is complex, with two oscillators to power, set up and mix, keeping levels within the range the mixer can work with properly. The frequency you want has to be filtered out and the filter has to be both wideband to work over the full range of the VFO, yet have sufficient Q to reject the unwanted mixer products. Not an easy job above 20MHz!

All other features are as an L/C oscillator.

Ceramic Resonators

	LF	Lo-HF	Mid-HF	Hi-HF	Lo-VHF	Hi-VHF
Cost	Low	Low	Low	Tricky	n/a	n/a
Construct.	Reasonable	Reasonable	Reasonable	Reasonable	n/a	n/a
Drift / min.	V. Good	V. Good	V. Good	Good	n/a	n/a
Purity	Good	Good	Good	Good	n/a	n/a
RIT	Possible	Possible	Possible	Possible	n/a	n/a

Cost is low - if (and it's a big "if") you can find a stock value that's the frequency you want. Custom resonators are available; be prepared for "how many thousand are you ordering, Sir?" If you can use stock frequency resonators you can make a very good oscillator very cheaply and in minimal PCB space.

Construction is as crystal oscillators; the "super-VXO" circuit with two resonators in parallel can give amazing shifts (+/- 10's of kHz in some instances) in frequency, but temperature induced drift increases dramatically as more "pull" is applied. Ceramic resonators are notoriously temperature sensitive, nowhere near as stable as a crystal.

Drift - Temperature is the biggest cause of trouble, then supply voltage and biasing - you have to avoid internal heating, too. Don't expect anything like the temperature stability of a crystal oscillator, but you'll run SSB/CW no bother on 160, 80, 60, 40m; possibly 30 and 20m. (My experience suggests not higher than 80m is best G3PCJ!).

Purity of the output depends very much on the oscillator having correct biasing , stable power supply, lack of vibration and loaded to a minimum (amongst other variables). Ceramic Resonators are distinctly micro-phonic - tap one with a pencil whilst monitoring the note on a receiver!

RIT can be done by switching in additional capacitors to shift the frequency, but this can be "chirpy". RIT applied by manually adjusting a small value trimmer capacitor is not as prone to "chirp", but the oscillator will take a finite time to settle, and this can produce "warble" notes.

Range is as per crystal oscillator methods - the "super VXO" being very effective. (I would have said the tuning range for ceramic resonators is often near 1 to 2 % of their frequency – very much higher than for a crystal oscillator, so makes them viable for low cost 80m rigs. G3PCJ)

Huff and Puff digital correction

	LF	Lo-HF	Mid-HF	Hi-HF	Lo-VHF	Hi-VHF
Cost	Low	Low	Low	Low	Low	n/a
Construct.	Reasonable	Reasonable	Reasonable	Reasonable	Reasonable	n/a
Drift / min.	V. Good	V. Good	V. Good	Good	Good	n/a
Purity	Good*	Good*	Good*	Good*	Good*	n/a
RIT	Easy	Easy	Easy	Easy	Easy	n/a

* depends on output LPF if sine wave required.

Cost Using 50MHz clock "74HCxx" logic is amazingly cheap. You'll need tuning varicaps to get the range, though the varicap diodes can be LED's, 1N4001's etc. The simplest circuit I've seen to implement Huff-n-Puff uses two 74HCxx logic chips, and a "fast" mode (statistical) H-n-P with three! PIC's and other micro-controllers can implement H-n-P's easily and cheaply - but why bother if a couple of 74HCxx gates do the job? Keep in mind you're using an L/C oscillator but tuning over LF and low/mid HF bands that are very crowded is easy.

Construction Be aware that 74HCxx logic has nS edges; high speed PCB techniques are mandatory and fast edges create hash, so watch the screening and isolation. An output LPF is needed for applications requiring a clean sine wave as the output is logic level pulses.

Drift - You can't polish muck; neither can H-n-P stabilise a poor oscillator. Start with a decent L/C oscillator and a H-n-P will render it superb. The tuning goes in steps of a few Hz depending on the timing period and counter/divider ratios; by using fixed NP0 and mica tank capacitors shunted with varicaps you can cover any of the amateur bands easily. A H-n-P will take some time to phase lock but once locked will hold the frequency (lock range > capture range).

Purity Depends on the quality of the basic oscillator the H-n-P correction is applied to. Start with a poor design and the system cannot phase lock. A decent oscillator with H-n-P applied become a superb oscillator; not perhaps to DDS or synthesizer levels but perfectly adequate for WSPR, QRSS, and demanding SSB adherents.

RIT Can be done by switching in additional tank capacitors to shift the oscillator slightly. Sudden switching causes the oscillator to jump to a new frequency; RIT, if not excessive, will not lose phase lock but the loop will take time to settle and might skip steps. Alternatively a slow (C-R) ramp voltage to the varicaps will allow the phase lock to hold.

Range Depends on the L/C oscillator and phase lock characteristics of the H-n-P; this is a function of the division and timing ratios.

Comment by G3PCJ. Implementing Huff and Puff often needs an extra tuning 'device' (eg a varactor) to implement the small capacitance corrections – in some instances, making sure this aspect of the control loop is stable can be problematical!

DDS / Synthesizers

	LF	Lo-HF	Mid-HF	Hi-HF	Lo-VHF	Hi-VHF
Cost	Low	Low	Low	Low	Low	Low
Construct.	Reasonable	Reasonable	Reasonable	Reasonable	Reasonable	Reasonable
Drift / min.	< 1Hz	< 1Hz	< 1Hz	< 1Hz	< 1Hz	< 1Hz
Purity	Good / LPF	Good / LPF	Good / LPF	Good / LPF	Good / LPF	Good / LPF
RIT	Easy	Easy	Easy	Easy	Easy	Easy

Cost It is low and falling. The biggest expense is probably the controller (Arduino, Butterfly, PIC, etc.) and the display, but the prices are dropping every day. Direct frequency readout eliminates dials and associated mechanics so what's not to like?

Construction High speed digital switching and displays create hash up to many MHz so watch the screening and isolation. An output filter is mandatory if you want anything other than a logic level digital output.

Drift - Literally "rock solid". Expect stability fully equal (if not better) to a similar frequency crystal oscillator.

Purity Depends on the quality of the output filter, as the output is logic level pulses. Some applications, like switching mixers and diode double balanced mixers are happy with logic level drive; but if you want a clean sine wave for a transmitter you'll need a good output filter. The job is easier with a DDS with active internal "spur" reduction (bless the engineers at Analog Devices!). Considering the frequency range of a modern DDS (kHz to GHz) the filtering for a clean sine wave is probably the biggest factor. You'll be looking at switched banks of filters if you're covering many bands.

RIT Can be done by a few digital manoeuvres; no bother at all.

Range Hz to Ghz in modern devices!

Comment by G3PCJ. There maybe some drawbacks! Are they as well understood as the simpler form of oscillator – definitely NO! Can they be adapted as easily as the simple LC form of oscillator? This depends very much on what has been provided in the software of the driving micro-processor and the programming skills of the builder to alter that software. You need also to be aware that the processor (being digital) could also generate hash that might interfere with the reception just like the Huff and Puff approach! You might also need to be a dab hand at combining more than one form of electronic construction if attempting this yourself!

Summary Chart

This chart shows the frequency band horizontally, and the potential drift vertically. The higher up, the higher the drift.

	LF	**Lo-HF**	**Mid-HF**	**Hi-HF**	**Lo-VHF**	**Hi-VHF**	**μ-Waves**
L/C	YES!	YES	yes/maybe	NO	NO!	NO!	NO!
Mixer	YES!	YES!	YES	yes	no	NO	NO!
Cer. Res.	YES!	YES!	YES	maybe	no	NO	NO!
Huff-n-Puff	YES!	YES!	YES!	YES!	yes	maybe*	NO**
DDS Synth.	YES!	YES!	YES!	YES!	YES!	YES!	YES!

*if a hi-VHF signal is pre-scaled then H-n-P methods will work, but stability is an issue as loop gains have to be higher, effectively multiplied by the pre-scale factor.

** new GHz pre-scalers might make H-n-P a feasible and cheaper alternative to DDS for narrow bands or spot frequencies; or ECL gates in conventional H-n-P circuits running directly at GHz.

Key:

YES! = Without a doubt, what are you waiting for?

YES = Definitely, get on with it, use good construction and don't skimp.

yes = You can, use good construction and parts.

maybe = Well, maybe...

no = Not really, perhaps a waste of time; but please feel free to prove me wrong.

NO = Don't waste your time and money, but try it if you really insist...

NO! = Don't bother. No, honestly, DON'T BOTHER.

The next note continues the theme from the last Hot Iron, of improving our understanding of the mathematical processes that occur in a modern Software Defined Radio.

DSP Fundamentals By Gary C. Sutcliffe, W9XT

To hams used to working with analogue circuits like VFOs, audio amplifiers, RF linear amplifiers and the like, the current trend towards DSP (Digital Signal Processing), SDR (Software Defined Radio) and the proliferation of the digital modes like PSK and the WSJT suite can seem a bit daunting. This article looks at the fundamentals of DSP with simple examples without going too deeply into the math.

Let's start with looking at the characteristics of an analogue signal such as the one shown right – Fig 1. It is a simple sine wave. We can pick any random point in time on the signal and measure the voltage to any arbitrary number of digits limited only by the quality of our instruments and by noise. Another key aspect of analogue signals is that you can pick any two points in time and there will be an infinite number of points between them. This is called a continuous function. There are no gaps between any two points on the signal.

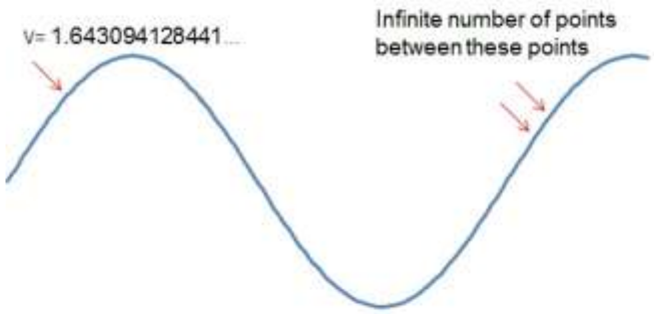

We can take an analogue signal and make changes to it. We can reduce the voltage with a resistor voltage divider or make a filter with resistors, inductors and capacitors. We can amplify the signals with transistor or tube amplifiers. The laws of physics determine how these components will affect a signal fed into the circuit. Digital systems like computers and microprocessors work in a different world, often called the digital domain. They take discrete numerical values and can perform mathematical operations on them. For example, it can take two numbers, X and Y, and multiply them together to get result Z. Microprocessors are getting faster and faster and perform math in real time if the processor is fast enough compared to the frequency of the analogue signal. The circuit will collect data on the analogue signal and convert to it to numerical values. The processor can then perform math on the data, and finally convert the result back to an analogue signal.

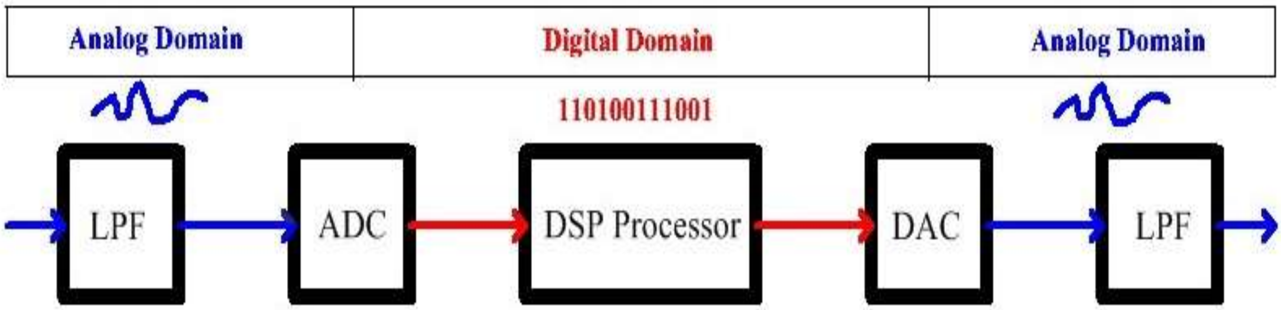

The diagram above Fig 2 shows a block diagram of a simple DSP system. For now, ignore the low pass filter blocks (LPF), the analogue signal goes into a device called an Analogue to Digital Converter (ADC). An ADC is essentially a volt meter that measures the analogue voltage at one point in time and converts it to a number in binary format. The processor takes the values provided by the ADC and performs mathematical functions on them. Finally, the processed data is fed into a Digital to Analogue Converter (DAC) which generates the analogue voltage equivalent of the binary number given to it.

Now, the processor can do some trivial math on the digital data. It can divide the signal by some number and essentially be equivalent to a resistor voltage divider. It can multiply it by some number and become an amplifier. If it multiplies the signal by -1 it simply inverts the signal. Of course, if this is all you want to do, it is much simpler to just use conventional electronic components. The real advantage is when you want to do complex DSP algorithms.
Let's back up a bit and take a closer look at some of the processes involved. As mentioned before, analogue signals are continuous functions. Computers must work with discrete values, at least until infinitely fast computers are developed. So, the data must be fed into the processor in chunks separated by time.

Figure 3 below shows a signal that is "sampled" at regular intervals, and the table of data that is created. Each dot on the waveform is the point where the signal is sampled. The table shows the values at each sample period.

0.000
0.479
0.841
0.997
0.909
0.598
0.141
-0.351
-0.757
-0.978
-0.959
-0.706
-0.279
0.215
0.657
0.938
0.989
0.798
0.412

Note that all the values of the signal between sample points are ignored. We are losing fidelity of the signal by ignoring these lost segments of the signal. However, Nyquist's Theorem states that if we sample at least twice as fast as the highest frequency present, we can recreate the original signal. In practice samples are usually done at a higher rate than specified by the Nyquist criteria. Going back to Figure 1, a low pass filter is usually placed in front of the ADC to prevent higher frequency signals getting into the system. Those signals will result in errors in the signal processing.

In selecting an ADC, two important specifications are considered (there are others but are beyond the scope of this article). First is the sample rate. That is how fast signals can be converted to the binary format. As pointed out by Nyquist, this will affect the maximum frequency signal that can be worked with. Another important specification is the resolution. The output of an ADC is in the binary number system used by computers. Analogue to digital converters are specified in the number of bits, 8 bit, 12 bit, 14 bit, etc. The number of bits indicates the number of different values the input signal can be divided into. For example, an 8 bit ADC can show a signal as a value up to 2^8 or 256 values. A 14 bit ADC can represent that same signal as 2^{14} or 16384 different values. More bits are required if your application requires a high dynamic range and/or more precise measurements. Making an ADC with both high sampling rates and high resolution is hard, and has been a factor limiting how fast DSP developed. Today High speed/high resolution ADCs are available that are capable of direct sampling at RF frequencies. This has allowed development of direct sampling SDRs which don't require mixing the desired RF signal to a lower frequency.

Before going into the processing of the signal let's take a quick look at the other end of a DSP system, the DAC as shown in Figure 2. Like the ADC, the number of bits and the maximum operating frequency of a DAC are critical in determining the capability of our DSP system. The output of a DAC will look something like shown in Figure 4 right. Because the DAC operates on discrete values that change at a regular rate, we can't represent the signal as a continuous analogue signal. The voltage will remain constant for the entire sample period. We can however clean that up by putting a low pass filter on the output of the DAC as shown in Figure 2.

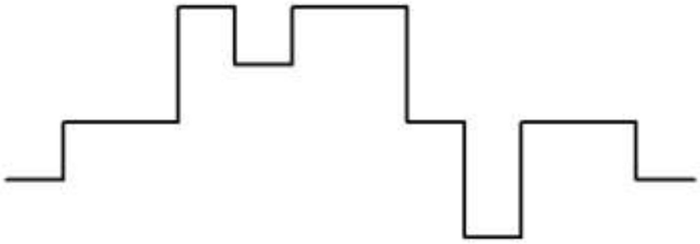

The real magic of DSP occurs in the processing and the algorithm implemented in the software. To start, let's look at a very simple implementation of a low pass filter. Figure 5 below shows a sine wave with some noise added on the left. This was created in a spread sheet where a table was created by using the SIN function. A small amount of noise was added by using the RANDOM function to cause slight variations in the signal. The PLOT function draws a graph of the resulting table. It is clearly seen as a sine wave with some noise.

To implement a very simple low pass filter we can just average the signal. In this example we take a value, add in the previous and next values, and divide by three. We do this

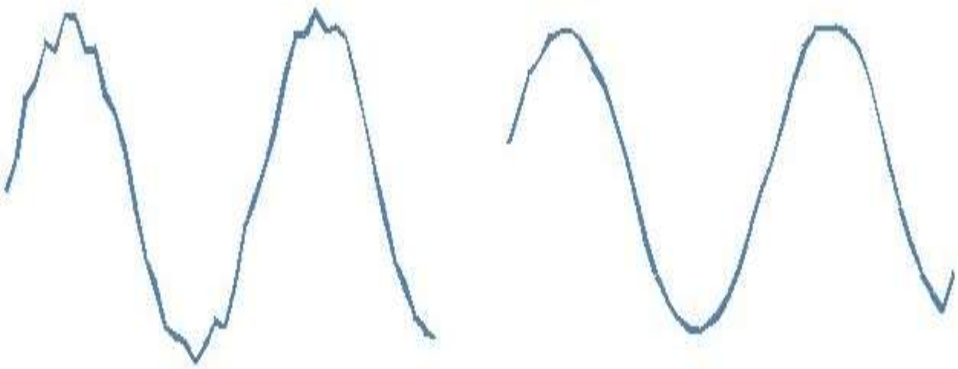

0.185424	0.244936
0.549383	0.462038
0.651308	0.696421
0.888572	0.790422
0.831386	0.924232
1.052738	0.972532
1.033471	0.972644
0.831724	0.904917
0.849558	0.755619
0.585576	0.631673
0.459885	0.430243
0.245268	0.198822
-0.10869	-0.08733
-0.39857	-0.35708

for every value in the table. The right hand waveform above shows the resulting plot after averaging. While all the noise (look closely at the peaks) has not been removed with this simple filter algorithm, it is certainly closer to the original sine wave without noise added. You can create a spreadsheet and play around with different methods of manipulating the data if you are interested.

This example is very simple and probably would be better implemented as a simple analogue RC Op amp filter, but it shows the basic method that many DSP algorithms use. These algorithms take a given sample, then take adjacent samples and multiply each sample by a specific coefficient number then add them all together. Our simple example we only used three samples and used an effective coefficient value of 1/3 for each value. Typical DSP algorithms use many more cells, usually factors of 2 like 32, 64 or 128 samples. The magic is in setting the coefficients that each cell is multiplied by to get the desired function on the signal.

Figure 6 below shows a generic form used by many DSP algorithms. The top row represents the samples taken in time. S0 is the current sample being processed. S-1 is the sample taken previous to S0. S-2 is two samples back, etc. S+1 is the sample after S1, and so on, out to some number of samples, n. The second row shows coefficients that are multiplied by each sample. In our simple averaging filter, we only had 3 samples (S-1, S0, S+1), and the same coefficient for each sample. Of course more complex algorithms will have different coefficients, and are the key to implement different functions. The results of all the samples multiplied by their respective coefficients are added together to produce a single value that will ultimately be sent to the DAC. When the calculations on a single sample are complete, the values in the top row are shifted right. A new value will be shifted into S+n, and the value in S-n will be discarded. This will produce a stream of values for our processed signal. As the number of cells increase, and the frequency of the signal increases, the processor must be more powerful to perform all the operations before the next sample of data comes in. Special DSP processor chips have been developed that do the multiply and accumulate functions very efficiently. The sound card and video chips in your PC have DSP processors to do these jobs.

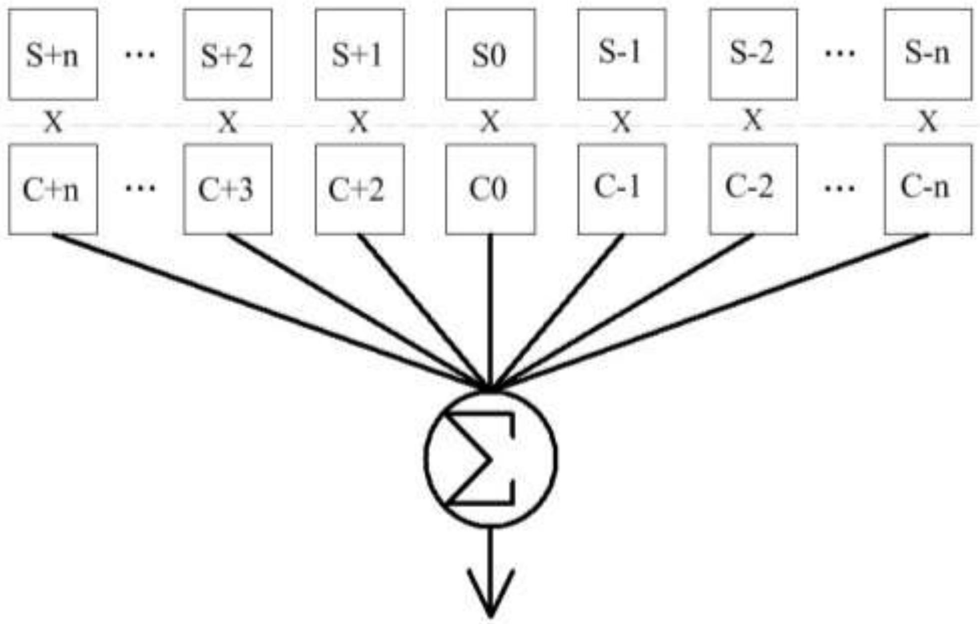

So, why is DSP much better in many applications? There are two main reasons. The first is flexibility. Say you have a conventional analogue radio and are using it to listen to SSB transmissions. If you want now to listen to an FM transmission you need to have a FM demodulator circuit. All the extra circuitry increases cost. With DSP you just change software much like you use your pc browser to surf the web then bring up your spread sheet program to work on your budget. Software is expensive to develop, but every copy is essentially zero additional production cost.

The other reason is that some functions are difficult and/or expensive to implement in hardware. One example is a very narrow filter with sharp skirts and low bandpass ripple, say for CW reception. You can make them but you need to use a lot of components. These components need to have very tight tolerances to get the desired results, and precise components are much more expensive. Furthermore, over time the components age and change value. The filter will not perform as well over time. Digital systems don't have this problem.

Hopefully this very brief introduction to DSP will give you a better feel for what is involved with the digital revolution that is hitting ham radio.

Thank you, Gary, for that excellent explanation. For those wanting a good reference book I can recommend An Introduction to Software Defined Radio by Andrew Barron ZL3DW. It is ISBN-13: 978-1500119935 and ISBN-10: 1500119938 originally published in July 2014. Thanks to Jim Gailer for lending me this more detailed explanation of SDR techniques. Tim G3PCJ

Circuit ideas to avoid RF wires to front panels – as in the Beer RX – by G3PCJ

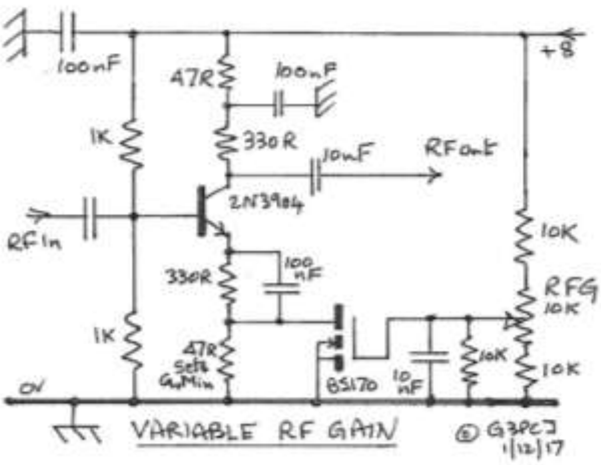

Variable gain RF amp – one scheme is to have an RF amp stage using a bipolar (2N3904) in the 'common emitter' configuration with a 'RF hot' variable resistance in the emitter lead to control the gain. Substitute a MOSFET for the variable and its On resistance can be changed by the DC voltage on its gate. It is convenient to put a fixed resistor in parallel with it to set the minimum gain. The resistors in series with the control pot are to restrict the gate voltage to near the useful range of approx 1.75 to 3 volts for OFF to hard ON. Voltage gain varies from about x 5 to x 20.

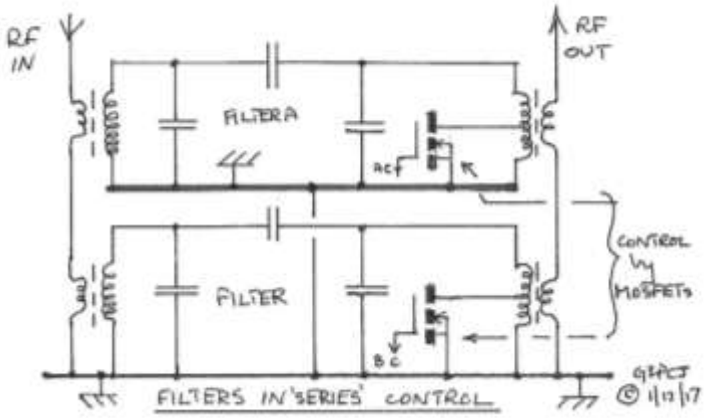

Band switching - normally this would be done by re-directing the low impedance input and outputs of the various bandpass filters. An alternative is to wire these low impedance windings in series and by some means short out the unwanted ones when not wanted. This can instead be done with the high impedance windings where the short would be reflected back across the low impedance winding. A switch across the whole of the secondary would not be good owing to the stray capacitance, but a switch across half of the high impedance winding would be enough to disable that filter! Make the switch a MOSFET and it can be controlled hard on/off by DC volts on its gate from the band switch. TOKO 333X series inductors often have a suitable tap!

Crystal switching – a Colpitts oscillator with a grounded crystal makes diode switching easy. Add some biasing resistors to turn the diodes ON when wanted, and also a reverse voltage (to minimise diode capacitance) when NOT wanted. Then add MOSFET switches, with the same hard on/off characteristics as for band switching above, across the crystals and their small off drain capacitance will not upset the crystal when it is active. The only slight drawback is potential changes in oscillator signal amplitude when changing from one crystal to another.

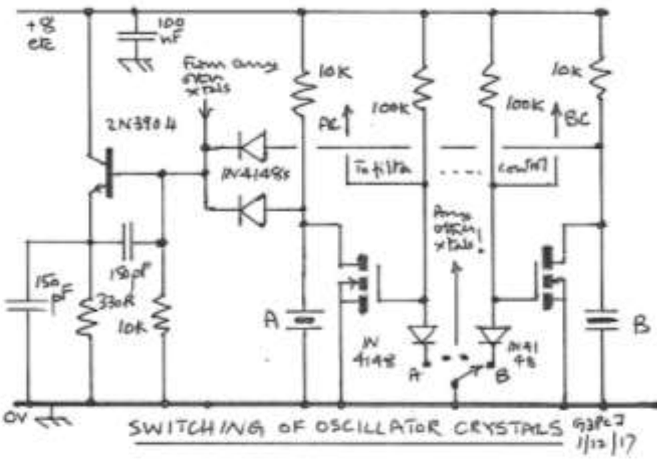

Happy Christmas to you all!