An Overview of Machine Generated Modes Used on Beacons and the Changes that may be Needed for New Modulation Types.

Andy Talbot G4JNT March 2021

A table of acronyms can be found at the end of this document.

Introduction

In February 2006 a newly upgraded GB3VHF 144MHz Beacon went on air [1]. Part of its transmission sequence included the, what was then very new, JT65B weak signal mode; part of the WSJT suite [2] for weak signal working. Previously that beacon, as had a few others around the World, carried a simple RTTY message as well as the obligatory CW identifier. But as far as can be determined GB3VHF was the first permanent beacon in the World to carry any advanced MGM modulation. The five microwave beacons on Bell Hill in Dorset, GB3SCS/F/C/X/K were subsequently modified to carry either JT4G or JT65C modulation.

Since then many beacons have adopted various modulations, nearly all Multi-Frequency Shift Keying (MFSK), taking either one of those from the WSJT suite, usually JT4G or JT65B or C or a specially developed mode, PI4 [3] developed from JT4G and customised to beacon transmission identification. Frequency bands have ranged from 5MHz on the GB3ORK beacon carrying JT9A, to several 24GHz beacons. There are probably MGM modes on beacons on even higher bands somewhere in the World.

A number of hardware solutions have been developed in the intervening fifteen years for generating these multi-frequency modes on the various bands from HF to SHF; some solutions better than others. With several new and improved WSJT-X modes coming along that adopt Gaussian smoothing of the frequency transitions to minimise signal bandwidth, some of those modulation generation techniques need to be re-thought and new ones introduced.

This paper will look at the modulation techniques that have been used, describing the pros and cons of each and suggesting ways forward offering maximum flexibility.

Direct Digital Synthesizer

GB3VHF uses an AD9852 DDS device clocked at 204.8MHz. The clock waveform is a GPS locked input at 12.8MHz, multiplied up using the internal PLL clock multiplier. The DDS generates the RF output directly at 72MHz which is subsequently doubled to 144MHz. Symbol data for the JT65B message is pre-calculated and the symbol values, numbers from 0 to 64, are stored in the PIC microcontroller. Each symbol, at approximately 370ms intervals, is converted to a frequency offset by multiplying its value by a constant. The constant is the value that would have to be programmed into the DDS to generate a frequency equal to the JT65B tone spacing of 5.4Hz. This offset is then added to the fixed value needed to generate the RF output. The result is sent to the DDS to directly generate RF at each successive tone interval. The DDS chip can provide fine tuning resolution and at VHF near-exact frequency setting is possible. As the AD9852 has a 48 bit frequency accumulator, that device could be used for accurate frequency generation up to microwave frequencies with suitable RF frequency multiplication.

However, a DDS followed by multiplication is not suitable for critical beacon usage. All DDS chips generate close in spurii that are often unpredictable in their frequency offset. When the RF is passed through a frequency multiplier these spurii increase in relative amplitude as the square of the multiplication factor so for UHF and up a DDS used as a direct frequency source needs to be used with extreme care. GB3VHF required a 72MHz crystal filter to be added to remove close-in spurii that appeared at unacceptable levels when multiplied to 144MHz. The two images show the output of the GB3VHF DDS module after the frequency doubler. The first is without the crystal filter, the second with it added in. Span is 3MHz, RBW 30kHz.

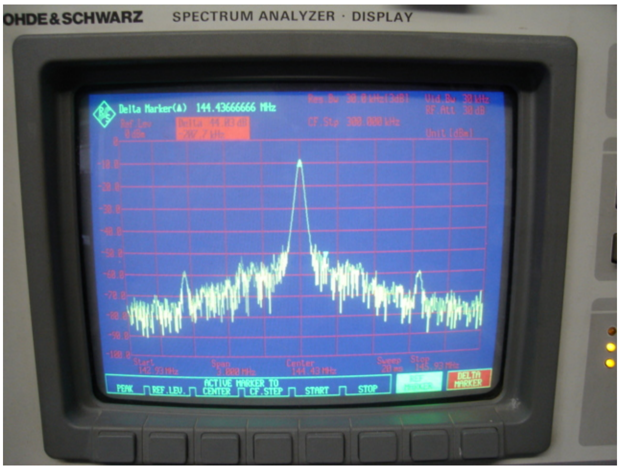

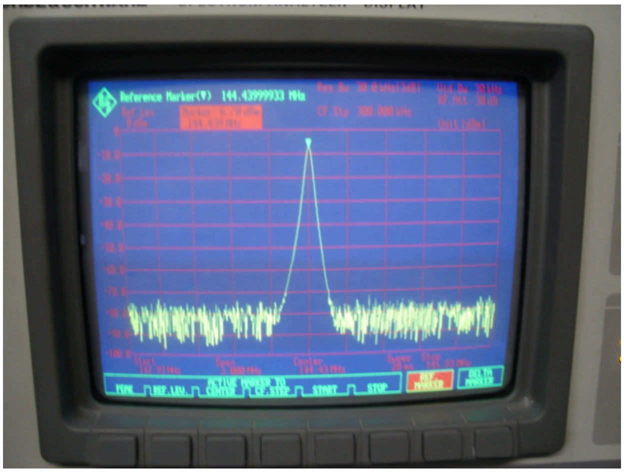

Adopting a better, more modern DDS clocked at a higher frequency, typically 1GHz generated at very low phase noise from a source locked to a master reference input, things are improved. This is the route taken on the *Next Generation Beacon Project* by OZ2M [3] using an AD9912 DDS device and ideally suited to all bands up to UHF. By using alias products of the DDS, and where output cleanliness is slightly less important this source is even useable up to low microwave bands. But it is not cheap. A good quality low phase noise reference-locked 1GHz signal is needed as a clock.

PLL Multiplication

For microwave beacons it is therefore generally unacceptable to directly multiply the output from a DDS source unless close-in spurii are filtered out or removed by design. Depending on just how close in these spurii are, some frequency multiplication techniques may be possible. A Phase Locked Loop multiplier fed with a reference input derived from a modulated DDS is one possibility. If the PLL loop bandwidth is made narrow enough, and the close in spurii from the DDS are not *that* close, the PLL may be able to filter out the unwanted products. There is usually enough flexibility with regard to the PLL multiplication factor to allow optimisation of the DDS frequency to allow a crystal filter to be added. Direct multiplication in a PLL was the technique used for the GB3SCS beacon on 2.3GHz [4] but even with a crystal filter, some spurii within a few kHz of the main signal are noticeable and have been reported by local stations.

Reverse DDS (RDDS)

Nearly all the older pre-existing microwave beacons make use of similar hardware; a crystal oscillator usually running at somewhere in the range 90 – 120MHz followed by a cascade of multipliers and filters to get up to the wanted output frequency. While this solution gives probably the cleanest, lowest phase noise output it does not lend itself to adding advanced modulation or of frequency locking to a master reference. Nowadays beacons are expected to be able to provide a frequency standard for calibration and as such really need to be locked to GPS or a similar standard.

A retro-fit for such hardware was developed to allow existing beacons to be upgraded to high stability locked operation with MFSK data mode capability [5]. The crystal oscillator source has to be modified to allow voltage tuning capability, making it a VCXO covering a range of perhaps one or two parts-per-million. This usually involves little more than adding a varicap diode in series with the crystal with a few R and C decoupling components. A direct output sample of the 100MHz also has to be provided.

The 100MHz voltage controlled crystal oscillator forms the DDS clock. The DDS itself is programmed to generate the reference frequency, usually 10MHz, from this input. The resulting output is then compared to the 10MHz reference input in a phase detector, the output of which is then used to steer the crystal onto the exact frequency that, when divided in the DDS, gives 10MHz. The DDS can be reprogrammed to impart the frequency shifts needed for MGM operation using the same code as that used for direct generation – except that by virtue of being in a feedback loop, the frequency setting codes effectively work backwards. An increased value programmed into the DDS register requires a lower clock input to generate the same 10MHz output.

The PLL bandwidth is defined mainly by the voltage tuning coefficient of the VCXO and rarely exceeds a few hundred Hz. This is much narrower than a traditional PLL used for direct multiplication could be persuaded to work at and is sufficient to clean up DDS spurii. The loop bandwidth is wide enough to not harm the pulse edges of the majority of WSJT modes and PI4. For a while the 3.4GHz beacon GB3SCF used RDDS and successfully carried a 50 baud RTTY signal with 850Hz shift. The phase locked loop was wide enough to allow this through without noticeable corruption or harmful smoothing of the FSK.

The low phase noise and general cleanliness of the crystal-oscillator multiplier route is still kept, so in a sense we have the best of both worlds. This is the reason quite a large number of UK beacons have adopted the RDDS technique, with a retrofit kit being made available at one point. But for new designs RDDS is not always the best solution. Crystals are proving more difficult and expensive to obtain nowadays; RF multipliers needing tuning and setting up and are a very hardware-intensive route to a microwave source.

A problem with RDDS is getting the frequency resolution needed. If a 48 bit DDS like the AD9852 is used, there is no problem, but the common 32 bit devices like the AD9850 can fall short of accuracy of the generated tones on the higher frequency bands. At 10GHz for example when used with a 10MHz reference, a tone setting accuracy of around 10Hz is possible. While such errors on tone spacing may be acceptable for the wide spaced JT4G and PI4 – probably – this makes generating JT65 or similar impractical. GB3SCF on 3.4GHz uses an AD9852 that allows for JT65C modulation. GB3SCX on 10GHz, by pure fluke, manages just a few Hz tone spacing discrepancy with an AD9850, [6] but it could be worse with potentially up to 15Hz error between adjacent tones if an unfortunate base frequency were selected.

Fractional-N Synthesis

A Fractional-N synthesizer is a normal PLL type whose internal divider is jittered randomly to average out at an intermediate effective division rate. By using a high comparison frequency and jitter rate in comparison with the loop bandwidth, the effects of the jitter can be mostly filtered out. The result is a synthesizer capable of tuning in small frequency steps with an 'acceptable' level of spurii and phase noise.

Fract-N devices generate an output frequency from a reference F_R equal to:

$F_{OUT} = F_R * (N + F/D)$ where N, F and D are register values programmed into the chip in real time.

The larger D can be made, so the smaller the frequency increment that can be achieved. For any specific output frequency there is appreciable leeway in selection of F and D, it is just their ratio that needs to remain constant for any given F_{OUT}. One way to generate an output with small offsets for any MFSK mode is to choose a value of D that gives exactly the offset needed. For example, JT4G requires tones, or frequencies, spaced 315Hz apart. If a 10MHz reference F_R is assumed, then by making D = 10MHz/315Hz = 31746, a unit change in the value of F results in a frequency shift of 315Hz. Generating the MGM is now simply a matter of adding the tone number to a constant then programming the result into the F register for each symbol.

Not all Fract-N devices allow a high D. The popular ADF4351 allows only a maximum of 4096, so narrow frequency shifts are not possible with that device. However the LMX2470 has a 22 bit D register so small frequency shifts of just a few Hz can be achieved. The GB3WGI transatlantic beacon uses an LMX2470 to generate at 1156MHz, subsequently divided by 8 for the 144MHz signal carrying JT65B modulation.

The LMX2541 initially showed considerable promise. With its integrated VCO and output divider that could be set in the range 1 to 64, a beacon source for narrow spaced MFSK modes at any frequency from 35MHz to 4GHz could be generated in a small low IC-count assembly. Unfortunately it was discovered, almost at the last minute, that when the F register in this chip is changed the internal VCO goes though an auto-calibration procedure that generates brief wideband glitches, making MGM via F register reprogramming unworkable. A work-around was achieved by changing the base F and D values and reprogramming the D register instead, which didn’t generate the recalibration process. This works quite well but, as for the RDDS, the N vs. Frequency response now works backward and the values are more complex to calculate. There are also some uncomfortable frequencies that just won’t work very well – generally those close to a multiple of the reference.

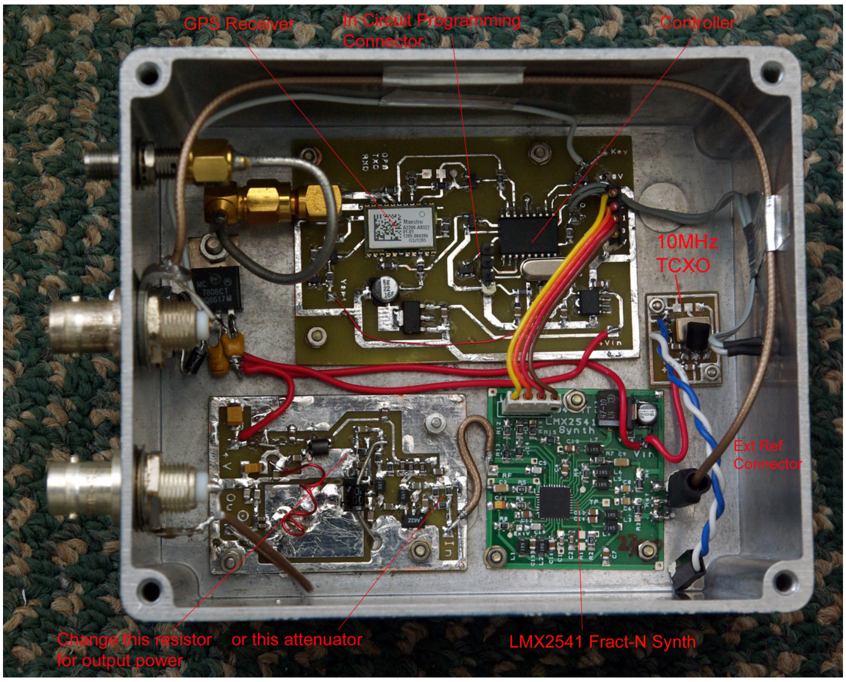

An LMX2541 based solution generating at 144MHz was used for both DL0SHF EME beacons where their outputs at 144MHz are mixed up to 10GHz and 24GHz. An early version of the hardware shown here.

In their latest incarnation the sources generate both QRA64 and Q65 modulation types at two different tone spacings/timings, all selectable via user switching. The recent 432MHz beacon programme was offered an LMX2541 solution to start with, but a few beacon keepers subsequently replaced them by *Next Generation Beacon Project* hardware to take advantage of its lower phase noise for beacons located at critical sites. A few microwave personal beacons have been built using the LMX2541, sometimes followed by multiplication. The phase noise is adequate, but generally considered a bit excessive for a high profile beacon on a good site.

The latest contender in Fract-N sources is the ADF5355 device, useable for direct generation to 6.8GHz and with its internal doubler to 13.6GHz output. Unlike other Analog Devices synthesizers, this one allows for very small frequency spacings by making use of a dual Fract-N engine and, importantly, its auto calibration and be turned off. By judicious choice of registers, it can be programmed for frequency shifts in a way more akin to a DDS than a Fract-N synthesizer. Full details for using the ADF5355 in this way can be found at [7].

Gaussian Smoothed Transitions

So far all the MFSK modes under discussion have been generated by abruptly switching from one tone output to the next at a symbol boundary. In a DDS this does result in phase continuity at symbol changes, as specified in the WSJT user Manual. But in a PLL implementation the shape of the frequency jump is entirely at the whim of the loop filter. Only the RDDS makes any attempt to smooth out the transition, and does that in a completely uncontrolled way dependent on the narrow loop bandwidth. For JT4, JT65 and PI4 modes these abrupt changes are acceptable, with the resulting 'key clicks' tolerated.

The latest FT8 and FT4 modes used at HF, as well as FST4/FST4W for the LF bands have a carefully shaped frequency transition between their symbols. Frequency 'glides' from one tone to the next with a Gaussian shaped transition. The result is a considerable reduction in bandwidth and virtually no 'key clicks', but now needs a compete rethink as to how the waveform can be generated directly in hardware. It is now no longer possible to just reprogramme a register setting at each symbol boundary and the whole symbol, for every one of them in the message, has to be constructed from a range of much smaller increments or decrements in frequency. In effect, we need to simulate what is going on inside a PC running the WSJT-X software, generating the waveform at a sampling rate many times higher than the symbol rate.

Clearly a Fractional-N synthesizer solution is completely out of the question. A DDS, however, can be programmed at sufficiently fine resolution and produces a phase-continuous waveform across every change in frequency. Full details of the exact Gaussian pulse shape is given in the QEX article at [8], although the mathematical description there is geared more towards generating the waveform from first principles for a soundcard, rather than as a series of frequency values fed to a DDS. The beacon controller software needs to be completely redesigned so it no longer generates timing based on symbol intervals. Now it will have to generate timing at a sample rate adequate to properly construct a version of the Gaussian transition. In practice the sampling interval will need to be an exact multiple of the symbol rate so the Gaussian waveform can be stored in a lookup table and merged with the value for each symbol.

No details have been worked out at this stage but a cursory working-through of the QEX equations suggests a sampling rate perhaps 20 to 100 times that of the symbol rate should suffice. This is well within the capability of the same mid-range 16F family of PIC microcontrollers as used for all WSJT beacon controllers to date. Reprogramming of an AD9852 DDS amplitude at 8kHz sampling in real time has been done using a 16F628 PIC device to generate a PSK31 waveform. The sample rate for this application can be appreciably lower than that needed for PSK31.

DDS Implementations for the Higher Frequencies.

In view of the inadvisability of trying to multiply a DDS source to get to higher frequency bands, and the fact that PLL solutions are not an option, and RDDS is not an ideal solution for new designs, we need another way to generate modulation for the VHF up to microwave frequencies.

Frequency mixing is one way forward. Generate the modulation at a relatively low frequency in a DDS and upconvert in a mixer using a fixed local oscillator. The traditional approach would be to use a single mixer and filter out the image response. This is generally feasible where the low frequency

IF is perhaps no lower than 1 - 5% of the output frequency and good image-rejection filtering is possible. If we assume that the highest frequency that is reasonable to generate in a DDS is, say, 144MHz then RF filtering is going to be commensurate with that used for transverters for the same frequency bands. In fact, any existing transverter design could be used but is not a particularly cost-effective route.

Quadrature or I/Q upconversion is one solution and here there are a number of advantages to be had. The opposite-sideband rejection possible in quadrature mixers depends on the mixer type and how well it is set up. Dedicated I/Q upconverter chips – and there are a very large number of these around, many covering up to several GHz – can manage around 40dB rejection without too much effort. A sample of these were tested and the findings can be seen at [9]. A home built mixer design using discrete diodes may only deliver 10 – 20dB. In all cases this rejection is not sufficient if the image frequency falls out of band; it still needs further filtering, although the filter requirements are now greatly reduced.

Baseband IQ Mixing

New rules apply and the image response problem mostly goes away if the waveform is generated at baseband, centred on zero frequency, or DC. Now the unwanted image lies on top of the passband of the wanted signal and, provided it is reduced by 15dB or more, will have no detrimental effect on the wanted signal. The issue now comes down to being able to generate the MFSK waveform as a dual channel IQ or quadrature signal carrying positive and negative frequencies. (Negative frequency is represented by a 180° switch in phase between I and Q channels.)

A DDS with a quadrature output is needed, but as this now only has to run at a few hundred Hz, or a kHz at most (the sampling rate we need to reconstruct the Gaussian transition), a simple solution based around a microcontroller and pair of D/A converters is all that is needed. Furthermore, in any DDS having quadrature outputs negative frequencies are generated automatically, by programming the NCO increment with a twos-complement version of the value that would otherwise be sent so it counts downwards. The values for each tone, or each sample on the Gaussian ramp, can be calculated in exactly the same way as for a standard DDS except that some of them will now be negative twos-complement numbers.

The photograph below shows two implementations of a quadrature output DDS using a PIC device. The upper one uses a D/A converter made from discrete components in an R/2R ladder, capable of a sampling rate up to around 80kHz, while the smaller unit uses a two-channel D/A converter chip controlled via SPI from the PIC which can sample at up to 16kHz. Both these units are designed to behave as if they were a dedicated DDS device and take an SPI or similar input from another controller generating the modulation. Both have been used to test-drive WSJT modulation via I/Q upconverters. Devices aimed at DSP operation, such as those in the Microchip *dsPic* family would allow more sophisticated modulation at higher bandwidths.

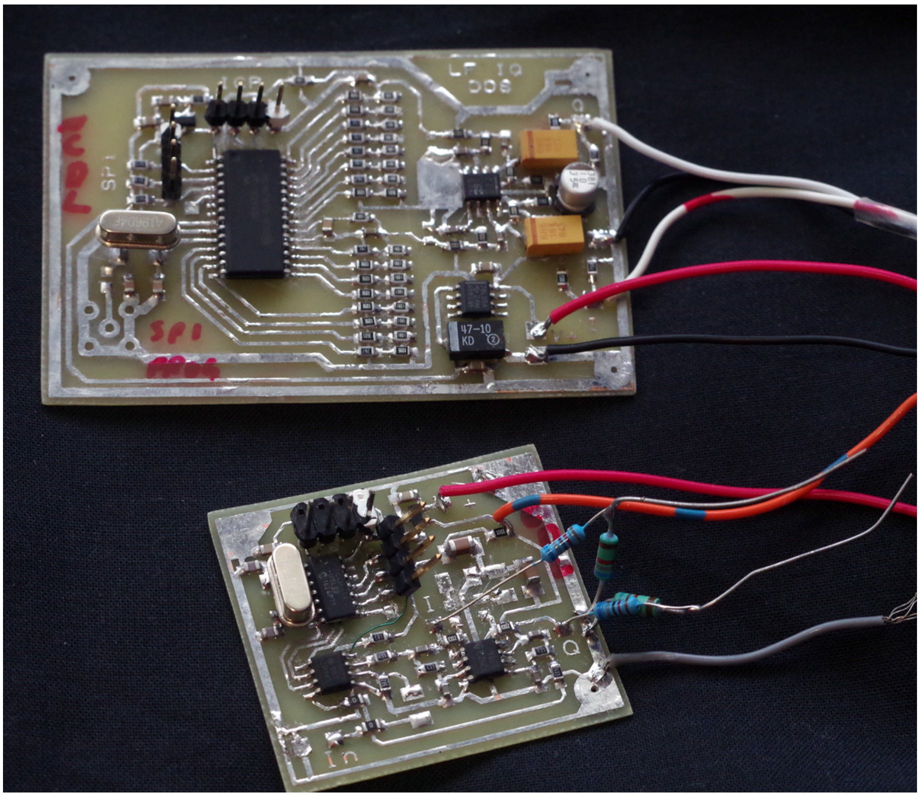

Conclusions

A number of ways for directly generating RF carrying MFSK modulation have been developed over the 15 or so years that WSJT modes have been used on beacons. New bandwidth-efficient waveforms present a problem with many of these traditional generator sources and a new approach based on I/Q upconversion has been proposed. At the time of writing demonstration hardware has been built in various forms, but only ever as test boards, nothing in any final version that can be used as a demonstrator. All tests so far suggest the technique is very viable – and straightforward.

References

[1] GB3VHF 144MHz Beacon http://www.gb3vhf.co.uk/GB3VHFtechnicaldetails.html

[2] WSJT Weak Signal Modes https://physics.princeton.edu/pulsar/k1jt/

[3] PI4 and the Next Generation Beacon Project https://rudius.net/oz2m/ngnb/pi4.htm

[4] GB3SCS 2.3GHz Beacon http://www.g4jnt.com/SCRBG/TheNewGB3SCS.htm

[5] Reverse DDS Kit https://www.qsl.net/yo4hfu/RDDS.html

[6] JT4G on GB3SCX http://www.g4jnt.com/SCRBG/JT4_GB3SCX.pdf

[7] Using the ADF5350 for MFSK http://g4jnt.com/ADF5355_Synthesizer_Control.pdf

[8] FT8 and FT4 Communication Protocols https://physics.princeton.edu/pulsar/k1jt/FT4_FT8_QEX.pdf

[9] Integrated Quadrature Upconverter Chips. *Design Notes*, RadCom Feb 2012. Also '*Design Ideas for Radio Amateurs*' RSGB, Page 41

Acronyms used in the text.

D/A	Digital to Analogue [converter]
DDS / RDDS	Direct Digital Synthesizer / Reverse DDS [a technique using a DDS backwards as the pseudo-divider in a phase locked loop].
DSP	Digital Signal Processing
FSK / MFSK	Frequency Shift Keying / Multiple Frequency Shift Keying
MGM	Machine Generated Modes. The official IARU designation for most computer generated data modes used on the amateur bands.
PLL	Phase Locked Loop
SPI	Serial Peripheral Interface. A three wire interface between a microcontroller and a peripheral device such as a synthesizer chip
VCO	Voltage Controlled Oscillator
VCXO	Voltage Controlled Crystal Oscillator
WSJT	'Weak Signals by Joe Taylor' A suite of weak signal data modes mostly using MFSK modulation such as JT65 and JT4.
PI4	An alternative modulation, derived from JT4 used specifically for beacon identifiers.