A NARROW BAND HIGH-RESOLUTION SYNTHESIZER USING A DIRECT DIGITAL SYNTHESIZER FOLLOWED BY REPEATED DIVIDING AND MIXING

Richard Karlquist, Hewlett-Packard Laboratories,
3500 Deer Creek Rd., MS 26M-3, Palo Alto, CA 94303-1392

Abstract

A synthesizer is described that generates a narrow band of frequencies around 10 MHz with extremely small step size and high spectral purity. A DDS running at 667 kHz is upconverted to 10.667 MHz and divided back to 667 kHz several times to improve the resolution and spurious suppression of the DDS. The final 10.667 MHz offset frequency is then mixed with 667 kHz resulting in a 10 MHz output. The post mixer filtering is done with inexpensive 10.7 MHz ceramic filters of the type used in FM broadcast receivers.

Applications

Metrology

It is frequently necessary to add a small, precisely controllable offset to a standard frequency signal. In some cases, the purpose is to "steer" a free running clock to agree with a known time reference such as UTC without losing the information about the error between the clock's own time and the reference. In other cases, the purpose is to get a beat note between the clock being tested and a reference clock, in order to apply resolution enhancement techniques to the frequency difference measurement [1]. A third related use would be as the "vernier" section of a synthesizer used in an atomic frequency standard.

In these applications, it is usually unsatisfactory to directly steer the clock in question. It may not be steerable at all, or may use a suboptimal method such as variations in the C field of an atomic standard. Even if none of these problems is present, it is desirable to have both a corrected and uncorrected version of the clock available simultaneously so that the uncorrected version can be compared to UTC. Hence what is called for is an external two-port frequency offset device.

Communications

A non-metrology use would be to facilitate a technique known in the paging industry as "simulcasting" where a paging receiver may receive pages from multiple base station transmitters at the same time. If two stations are in "zero beat," destructive interference will result. To get a controlled beat note, a small offset is added to the base station's frequency reference. Since this reference is frequency multiplied to 800 MHz or more, the offset synthesizer needs to have very good spectral purity. Another communications use might occur when building modulators. With some modifications to the architecture, it is possible to build frequency stable analog modulators. In many cases, analog modulation is still simpler and more accurate than digital modulation, except for the frequency drift problem.

Prior art

A popular method of generating frequency offsets with moderate performance is the phase microstepper. It makes use of the fact that a small frequency offset is equivalent to a slow phase ramp. The microstepper controls a digitally programmable phase shifter, initiating phase increments at appropriate intervals. When 360º of phase is reached, it resets to 0º, which is equivalent. It is subject to phase jumps, phase drift and spectral impurities. To improve on this performance, sophisticated multiloop phase locked VCXO synthesizers have been built [2]. While the performance of these systems is excellent, it depends on the use of high quality VCXO's. In some cases, general purpose synthesized signal generators have been pressed into service for frequency offset work, but even the best of these is only marginally useful in metrology work and in any case they are not very cost effective for this type of task. The most appropriate general purpose synthesizers for this application are ones fashioned after the repeated mix and divide architecture popularized in the Hewlett-Packard 5100A produced in the 1960's. Even these architectures tend to have a problem with temperature coefficient of phase. They are also very complex and expensive to implement.

A synthesizer will be described below that is essentially a simplified 5100 type architecture with just enough functionality to implement small frequency offsets of a 10 MHz reference. The important issue of improving the temperature induced phase drift will be covered in some detail.

Requirements

Frequency tuning range

A tuning range of ±1 Hz is quite adequate for many metrology applications and will work for simulcasting. For generating beat notes to measure frequency differences, up to ±1 kHz may be needed. In modulator work, a range up to ±5 kHz is sufficient for most FM voice and data channels. As shown below, there is a tradeoff: more resolution and spectral purity come at the price of reducing tuning range.

Spectral purity

In metrology applications, it is customary to require non-harmonically related spurious sidebands to be better than -80 dBc as a minimum, with -100 to -120 dBc being desirable. When observed with phase noise test instrumentation, spurs may be visible down to -140 dBc or even less. In communications applications, a good rule of thumb is to keep all spurs below -40 dBc *after multiplication.* In many cases, regulations require -60 to -80 dBc performance. With multiplication factors in the range of 100 to 1000, the original spurs are enhanced 40 to 60 dB. Hence the synthesizer needs to hold spurious sidebands to -80 dBc to -140 dBc as measured at the 10 MHz reference.

Resolution (step size)

For metrology, resolutions of at least 1 part in 10^{15} are desirable, which means steps measured in nanoHz at 10 MHz. Less critical applications may only need to have steps in the milliHz range. The synthesizer to be described utilizes a direct digital synthesizer (DDS) at the front end of the synthesizer to generate steps, so step size is to some extent a non-issue. By extending the DDS accumulator size, a relatively easy thing to do, it is possible to achieve any desired step size. While the divide and mix sections to be described later do add 4 bits of resolution apiece, this is not their major function, a least in the case where a DDS is used. In modulator applications where the DDS is replaced by an analog frequency/phase modulator, this 4 bits of resolution is more properly stated as 24 dB of additional dynamic range in the analog domain., which is a non-trivial improvement.

Phase stability vs temperature

By definition, frequency is the time rate of change of phase. If phase is temperature sensitive, and the ambient temperature changes with time, there will effectively be a frequency error [4]. The required phase stability depends on the assumed rate of change of temperature and the frequency error tolerance, but a good rule of thumb is to keep the overall phase drift expressed as a change in delay time to under 1 nanosecond over the full expected temperature range, such as 0 to 70°C.

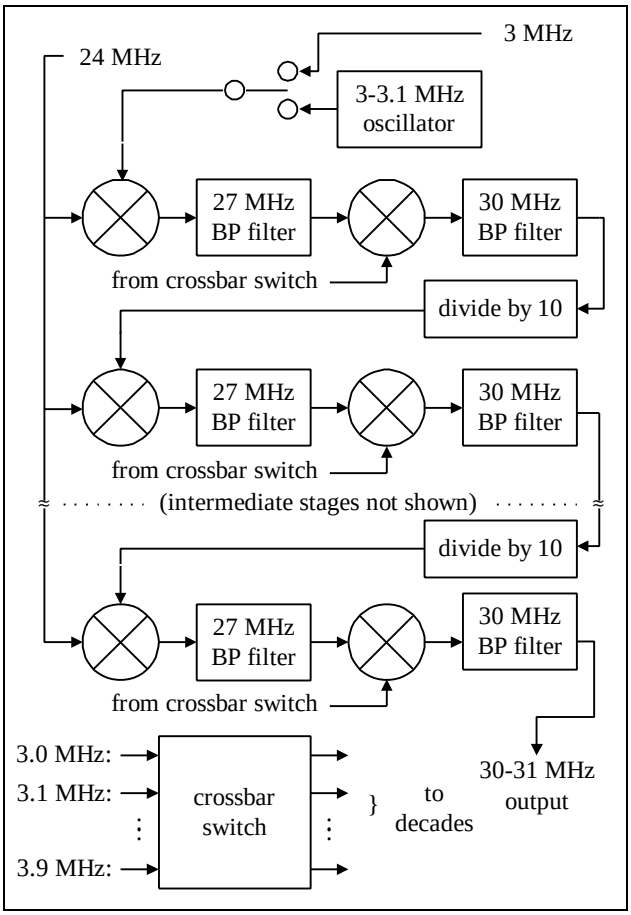

Figure 1. HP 5100 type architecture.

The classical mix and divide architecture

As an introduction, a brief review of the HP 5100 type of synthesizer will be given, which represents the mix and divide family of architectures. Figure 1 shows the 5100 block diagram. A synthesizer driver system on a separate chassis (not shown) generates reference frequencies at 24 MHz and 3.0, 3.1, 3.2, ... 3.9 MHz from the 5 MHz reference frequency input. The 24 MHz signal goes to all modules. A crossbar switch routes the 3 to 3.9 MHz signals to the appropriate divide and mix modules according to the frequency selected on the front panel. Each module is connected to a digit switch on the front panel which has a possible setting of

0 through 9. If this setting is denoted "N," then the switch will feed 3.N MHz to the corresponding module. In a two-stage upconversion process, each module converts an incoming tone in the 3 to 3.l MHz band to a tone in the 30 to 31 MHz band, then divides the resulting frequency by 10 to get an output that again falls in the 3 to 3.1 MHz band. Although the input and output frequencies fall in the same band, they are different because a decade of resolution has been added effectively. For example, if the input frequency to a module is 3.89 MHz, the output frequency will be 3.N89 MHz, where N is the setting of the digit for that module. E.g., if N=5, then the frequency would be 3.589 MHz. The modules are cascaded so that any amount of resolution can be realized by using a sufficient number of modules. The last module omits the ÷10, resulting in a 30 to 31 MHz output band. The first module's input that would normally come from the previous module is either hardwired to 3 MHz or can be connected to a "search oscillator" that tunes continuously from 3 to 3.1 MHz. This allows "infinite" resolution with good stability because any frequency drift in this oscillator is reduced by a factor of 10^M, where M is the number of mix and divide modules. Of course, the total tuning range (originally 100 kHz) is also divided by 10^M. Each module improves the spectral purity of the signal it receives from the previous module by 20 dB due to its ÷10 frequency divider.

The 5100 architecture as an offset synthesizer

Suppose a tuning range of only 1 Hz were needed for a typical metrology application. Then a search oscillator tunable over 1 MHz could be followed by six mix and divide modules, and one mix-only module, resulting in the required coverage. If it could be specified that the output frequency would be 30.000000 to 30.000001 MHz, then all modules could be hardwired to 3 MHz thereby eliminating the crossbar switch and 3.1 through 3.9 MHz generators. Having done this, the result would be that each module's first mixer would combine 3 MHz with 24 MHz to get 27 MHz. Hence these mixers could be dispensed with in all modules by instead feeding 27 MHz directly to the second mixer. The result is the reduction of the 5100 architecture to little more than the decade modules, and they are simplified to have only one stage of mixing and filtering instead of two. Finally, the search oscillator can be replaced by its modern equivalent, the direct digital synthesizer, to make the system a true synthesizer. The advantage of this arrangement compared to using the DDS output directly is that each module reduces DDS spurs by 20 dB when the signal is frequency divided and adds 4 bits of resolution to the DDS. The resulting configuration is shown in Figure 2.

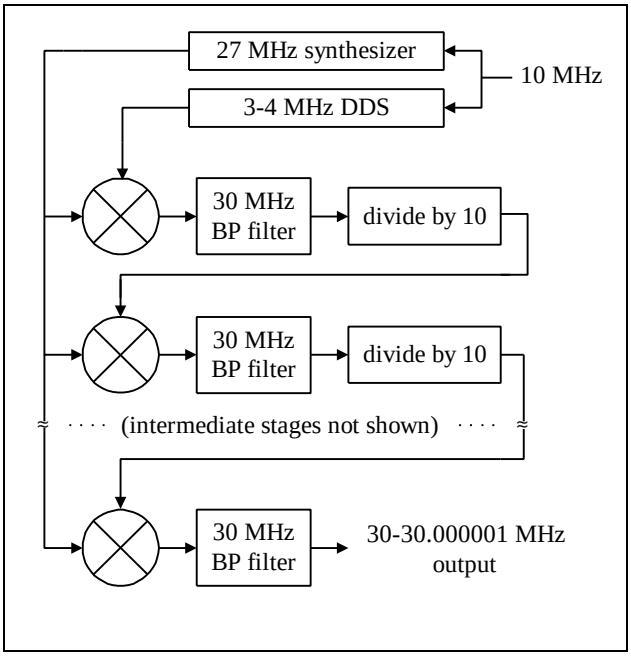

Figure 2. Narrowband architecture with DDS.

Filtering

The bandpass filter after each mixer is essential to getting high spectral purity. In the original 5100 and in the initial stages of the simplified architecture of Figure 2, the bandpass filter must pass 30 to 31 MHz, a fractional bandwidth of 3.3%. In the later stages of the simplified architecture, the passband can be much narrower, if non-identical modules are acceptable. In all cases, the filter must be capable of rejecting the 24 MHz image and the 3 and 27 MHz input frequencies. The 9:1 ratio of the mixer input frequencies results in a situation where the only mixer products that fall within the 3.3% passband are high order ones that decrease rapidly as the mixer drive level is backed off. It is very important that the out-of-band mixer products not be allowed to reach the divider. If they do, the non-linear nature of the divider will alias them back into the passband and compromise spectral purity. Hence the main factor in maintaining spectral purity is proper filtering.

The tradeoffs in filtering that apply here are bandwidth limitations, temperature drift, microphonics, phase noise, and cost. With the 3.3% bandwidth required for the 5100 architecture, crystal filters are ruled out because they are incapable of achieving such a wide bandwidth. On the other hand, 3.3% is so narrow as to

be barely feasible for LC filters, and there isn't much prospect for improvements in passive filter technology. Besides the manufacturing problems of producing such narrow filters, LC filters have a substantial problem with temperature stability. This could cause a temperature dependent phase shift that would be most damaging in the last stage. One possible approach to this problem would be to use a crystal filter for the last filter only, since the required passband at that stage would be narrow enough to implement with a crystal filter. This would not necessarily alleviate temperature stability problems because of the long delay time through the filter caused by the narrow bandwidth. A long delay time represents many radians of phase shift and hence a small change in bandwidth or center frequency could cause a large change in delay time. Other problems with crystal filters are phase noise and possible microphonics.

Ceramic filters

What is needed is a moderate bandwidth filter with good stability. A so-called "ceramic" piezoelectric filter is a good match for this requirement. It is optimal for applications such as this one requiring bandwidths too narrow for LC filters and too wide for crystal filters. Ceramic filters are more stable than LC filters but are wide enough to avoid the close in phase noise and microphonic problems of crystal filters. Additionally, ceramic filters can be much cheaper than either of the alternatives. They are used worldwide in FM broadcast receivers for IF filtering and cost a fraction of a dollar in production quantities. These filters are generally available only with a center frequency of 10.7 MHz and with various bandwidths from less than 100 kHz to 500 kHz. Although these filters have been in use for decades, their performance has improved in recent years. When receivers were changed from analog tuning to synthesized digital tuning, it became necessary to improve the tolerance and stability of the center frequency. This is because it was no longer possible to compensate for an error in the IF filter center frequency by merely offsetting the analog tuning slightly. It is now easy to buy filters that are centered very precisely at 10.7 MHz and have low temperature drift. Recently, larger bandwidths such as 400 or 500 kHz have been implemented for use in DBS receivers, where the deviation is greater than with FM broadcast.

10.7 MHz architecture

The key to using 10.7 MHz filters is to determine a technique for scaling the simplified 5100 type architecture such that the filter frequency becomes 10.7 MHz. If all the frequencies in the architecture in figure 2 were merely scaled by a factor of $^{10.7}\!/_{30}$, the filters would of course change to 10.7 MHz, but now the 27 MHz reference frequency scales to 9.63 MHz, which is quite inconvenient to synthesize from 10 MHz. (figure 3a). If instead the frequencies are scaled by a factor of $^{10}\!/_{27}$, then the 10 MHz could be used directly as a reference, but the filters would need to be changed to 11.111 MHz (figure 3b). The problem is that the reference and filter frequencies are too far apart. The mix and divide process can be generalized as shown in figure 3c. The frequencies obey the equation: $F_{ref} = (F_{IF})(N-1)/N$, where N is an integer equal to the frequency divider ratio.

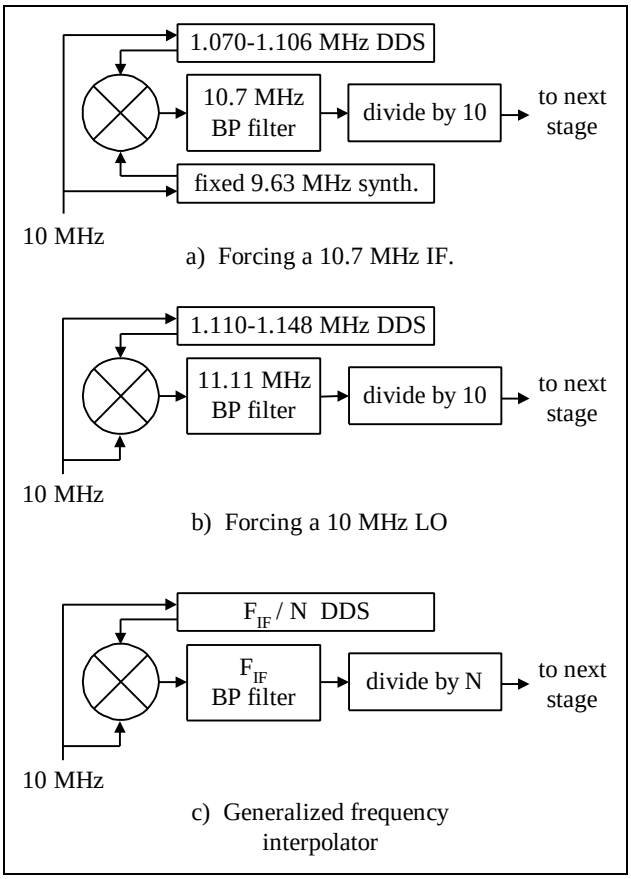

Figure 3. Transformations of 5100 architecture.

Table 1 shows the results of this equation for values of N of interest, assuming F - IF = 10.7 MHz. The values of N that give the closest reference frequency to 10 MHz are 15 and 16, which result in 9.987 and 10.031 MHz respectively. It is now possible to take advantage of the fact that it is acceptable to operate the filter slightly off center, if a small decrease in tuning range can be tolerated. If the N=15 and N=16 configurations are scaled

in frequency to get the reference frequency to exactly 10 MHz, then the resulting IF frequencies are $10\frac{5}{7}$ =10.714286... MHz and $10\frac{2}{3}$ =10.66666... MHz respectively (figures 4a and 4b). With several hundred kHz of bandwidth available, sacrificing 14 or 33 kHz is a reasonable compromise. Although the IF is closer to 10.7 MHz with N=15, it is usually preferable to use N=16 for three reasons:

N:	Ref. freq.	N:	Ref. freq.	N:	Ref. freq.
4	8.025	14	8.560	24	10.254
5	8.560	15	9.987	25	10.272
6	8.917	16	10.031	26	10.288
7	9.171	17	10.071	27	10.304
8	9.363	18	10.106	28	10.318
9	9.511	19	10.137	29	10.331
10	9.630	20	10.165	30	10.343
11	9.727	21	10.191	31	10.355
12	9.808	22	10.214	32	10.366
13	9.877	23	10.235	33	10.376

Table 1. Possible ref. frequencies for 10.7 MHz IF.

1. The divider is easier to implement in digital hardware and gives an output with very low even harmonics because the divide ratio is a power of 2.

2. It is usually easier to deal with fractional frequencies involving factors of $\frac{1}{3}$ than $\frac{1}{7}$, assuming it is necessary to convert the repeating decimal frequency to a more useful number. For example, $10\frac{2}{3}$ can be used to drive a frequency tripler to produce 32 MHz. With $10\frac{5}{7}$, a difficult-to-build frequency septupler would be required to produce 75 MHz.

3. A lower DDS frequency is possible as will be explained below.

Front and back ends for the mix and divide chain

A chain constructed of the modules shown in figure 4c requires an input near $\frac{2}{3}$ MHz ≈ 667 kHz for the first module. In order to have the capability of achieving zero offset, the input must be at exactly $\frac{2}{3}$ MHz, not merely a good approximation. Most DDS's are binary or decimal based. Neither of these can generate exactly $\frac{2}{3}$ MHz starting from 10 MHz., regardless of the number of bits of resolution. A ternary (base 3) DDS would be required, which is possible, but highly unlikely to be available. Instead, the DDS is operated in a band centered at 2 MHz and is followed by a ÷3 frequency divider (figure 5). In most cases, a DDS output of exactly 2 MHz is easily realized. If the N=15 architecture is chosen, then the DDS would need to generate 5 MHz instead of 2 MHz and a divide by 7 circuit would need to be used to obtain $\frac{5}{7}$ MHz ≈ 714.286 kHz.

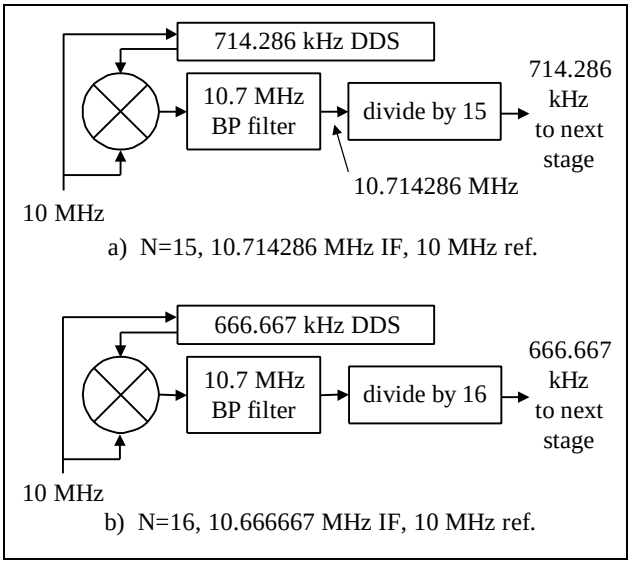

Figure 4. Usable frequency schemes for 10 MHz ref.

At the end of the chain, 10.667 MHz is available with programmable offset, but of course it needs to be converted back to 10 MHz to be useful. This is accomplished with one additional mixer and filter stage. However in this case the LO for the mixer is the 10.667 MHz from the previous stage. and the filter is tuned to 10 MHz.

The output stage raises the issue of filtering, since the standard 10.7 MHz filters are no longer applicable. At the end of the chain, the filter bandwidth required is negligible so it is possible to use a crystal filter. This filter should be as wide as possible to minimize phase noise and microphonics effects. Another possibility is to use 10 MHz ceramic resonators in combination with discrete capacitors to build a 10 MHz filter. These resonators are similar to the ones in the 10.7 MHz filters, and, fortunately, 10 MHz is one of the standard frequencies for them. They are normally used as low priced replacements for microprocessor clock crystals. Their moderate Q is a good match to this requirement since the nearest frequency to be rejected is 7% away from the

passband center. It is also possible to use a phase locked VCXO as a cleanup loop to get pure 10 MHz out.

Figure 5 shows the complete design with the blocks regrouped into iterated divide/mix/filter modules, which will be henceforth be referred to as "**interpolators.**" This organization is more logical for this design than the 5100 partitioning into mix/filter/divide modules.

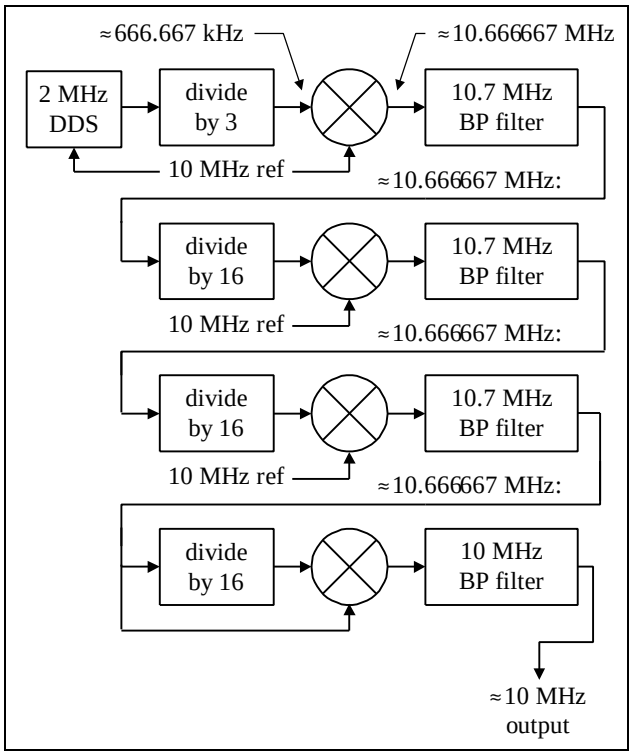

Figure 5. 10 MHz output, 10 MHz ref. synthesizer.

Design constraints

The architecture of figure 5 raises the question of how many interpolators are appropriate. Each iteration decreases step size by a factor of 16 (which may or may not be very significant) and it also improves spectral purity (i.e. phase noise, spurious sidebands, jitter, etc.) by 24 dB = 20 log 16, which usually is of great importance. After a sufficient number of stages, the total tuning range may become too small to be usable or the 24 dB per stage improvement in spectral purity may cease to occur because the noise floor of the interpolators or the reference frequency source becomes the limiting factor. The limit on tuning range depends on the number of interpolators and the bandwidth of the IF filters used. With the DBS type filters, about 400 kHz of range is available. This results in final tuning ranges of 25 kHz, 1.5 kHz, 100 Hz, 6 Hz, and 0.4 Hz for 1, 2, 3, 4, and 5 interpolations respectively. Since 5 interpolations gives a theoretical spectral improvement of 120 dB (in addition to the 10 dB from the divide by 3 after the DDS), it is unlikely that more than 5 interpolators would ever be used, even if less tuning range were acceptable.

Critical system design details

To get maximum performance for the overall design, it is important that the 10 MHz distribution amplifier (not shown in figure 5) have high isolation between outputs to keep the slightly different IF frequencies of each interpolator from intermixing. Also, it is helpful for it to have low distortion, especially even order, in order to minimize mixer distortion. As usual, it should also have low phase noise and good short term stability so as not to degrade the reference source. The various interpolators should be well isolated from each other by proper ground plane management and PC board layout, or installed in individual shielded boxes if necessary. Finally, adequate power supply filtering should be used to preclude coupling through the power distribution system.

Critical module design details

At the front end of the interpolator, a high quality sine wave to logic level converter should be used. A excellent treatment of this subject is given in reference [5]. After converting to logic levels, a low jitter logic family such as 74ACXXX should be used for the divider, which is preferably of the synchronous type, not the ripple type. The divider must be followed by a low pass filter to suppress harmonics; an elliptic function filter is probably the most appropriate type for this purpose, with 5, 7, or 9 poles depending on the required performance. A pad should be used between the filter and the mixer to properly terminate both the filter and the IF port of the mixer, which is driven with the filtered 667 kHz. The amount of attenuation in the pad can be adjusted to give the desired mixer drive level. A drive level of -30 to -10 dBm is reasonable for a +7dBm LO mixer depending on the required level of performance. For extremely high performance systems, it may be appropriate to use high drive level mixers.

The output of the mixer goes through two to four stages of filtering. It is convenient to insert amplifiers between the filters to prevent interaction between the ceramic filters since they are not designed to be cascaded directly. The amplifiers not only must make up for the filter loss of 3 to 6 dB but also provide a net gain to help

raise the low level signal coming out of the mixer to a higher amplitude to drive the clock of the next divider. The filters, when used as recommended, have a rounded ("stoop shouldered") amplitude response in order to get constant group delay, which is critical in FM demodulation. When several of them are cascaded, the resulting response has an even more pronounced rounding. Fortunately, it is possible to fix this problem by deliberately mismatching the filters. It is also possible to change the center frequency somewhat by adding capacitors (or inductors) in series with the input and output terminals. The effect of temperature on phase for a typical filter is 3 milliradians/°C. The phase does not accumulate through the interpolators, but rather is only critical in the last interpolator and the final frequency translation stage filtering.

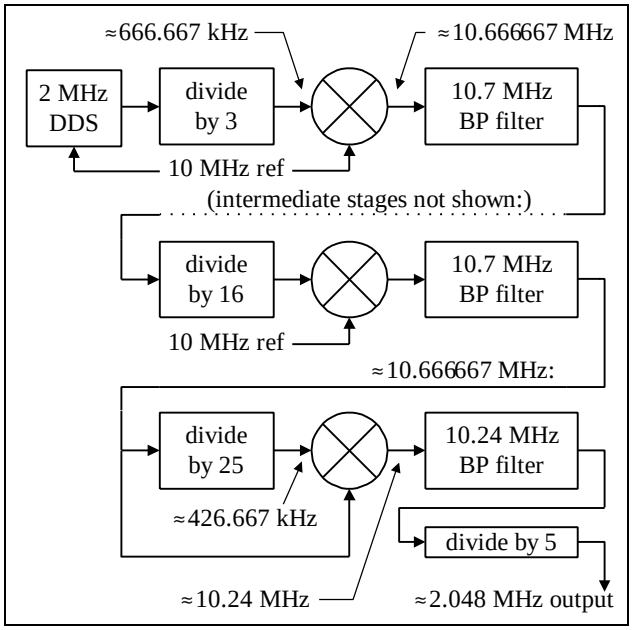

Figure 6. 2.048 MHz output, 10 MHz ref. synthesizer.

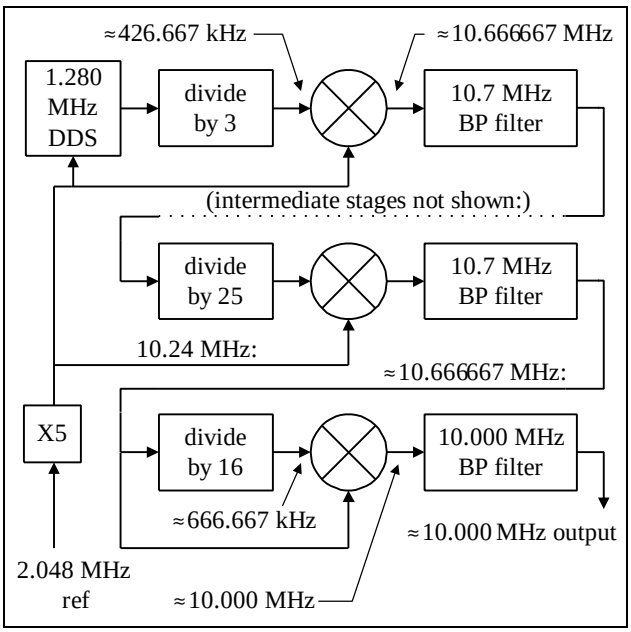

Figure 7. 10 MHz output, 2.048 MHz ref synthesizer.

Input and output frequencies other than 10 MHz

The final translation stage can be changed to produce frequencies of interest other than 10 MHz. For example, 2.048 MHz can be synthesized from a 10 MHz reference. Figure 6 shows the last divider changed to ÷25 to convert 10.667 MHz to 10.24 MHz. This can then be divided by 5 to get an offsettable 2.048 MHz output.

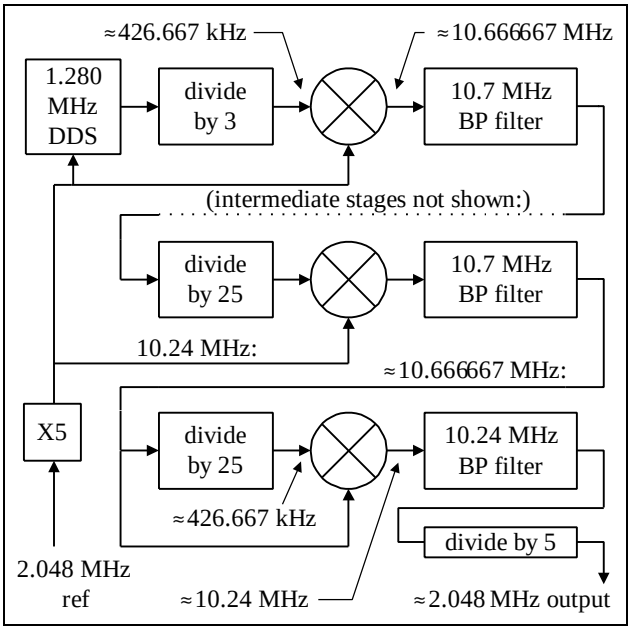

Figure 8. 2.048 MHz out, 2.048 MHz ref. synthesizer.

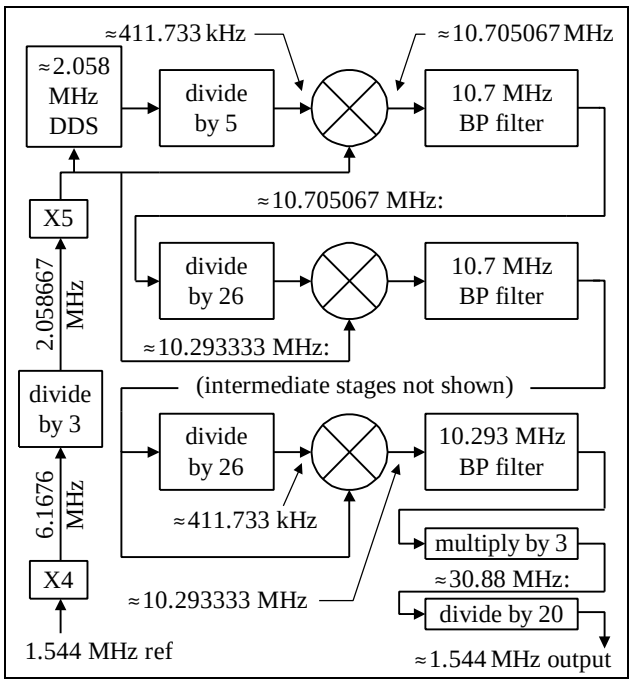

Figure 9. 1.544 MHz out, 1.544 MHz ref. synthesizer.

In figure 7, the last divider has been returned to ÷16 but the previous ones have been changed to ÷25. This makes it possible to multiply a 2.048 MHz input by 5 to get 10.24 MHz, which is used as the LO for the mixers. The IF frequency is still 10.667 MHz and the output stage of figure 5 produces offset 10 MHz. Figure 8 combines the techniques of the two previous figures to get a synthesizer with 2.048 MHz reference frequency and an offset 2.048 MHz output.

In figure 9, the North American telecom frequency of 1.544 MHz is multiplied by ⅓ to get 2.058667 MHz. This is close enough to 2.048 MHz to use the scheme of figure 8, but it is advantageous to increase the divide ratio from 25 to 26 to get the signal closer to the center of the passband. The result in figure 9 is the North American equivalent of figure 8, with 1.544 MHz in and offset 1.544 MHz out.

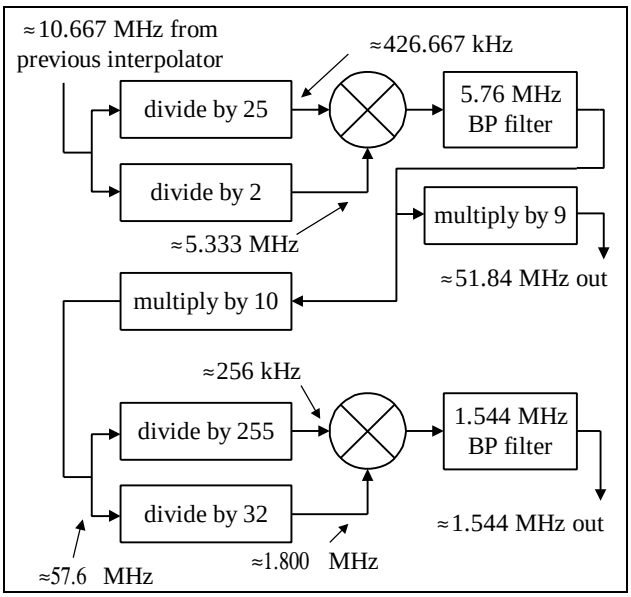

Figure 10. Output sections for 51.84 and 1.544 MHz.

In figure 10, output sections are shown for generating the Sonet/SDH frequency reference of 51.84 MHz or 1.544 MHz starting with either the 10 MHz input synthesizer shown in figure 5 or the 2.048 MHz input synthesizer shown in figure 7. Figure 11 shows a synthesizer for exciting cesium beam tubes (or absorption cells) with a direct multiply frequency scheme. The synthesizer produces 10.714 MHz which can be multiplied by 78, then multiplied by 11 resulting in the cesium transition frequency of 9.192 GHz. There is sufficient tuning range to move 50 kHz off the main line to the Zeeman line for measuring the effective C field magnitude [6]. The output of the ×78 multiplier in figure 11 is conveniently filtered by an off the shelf 836.5 MHz dielectric resonator filter, of the type used for cellular telephone transceivers.

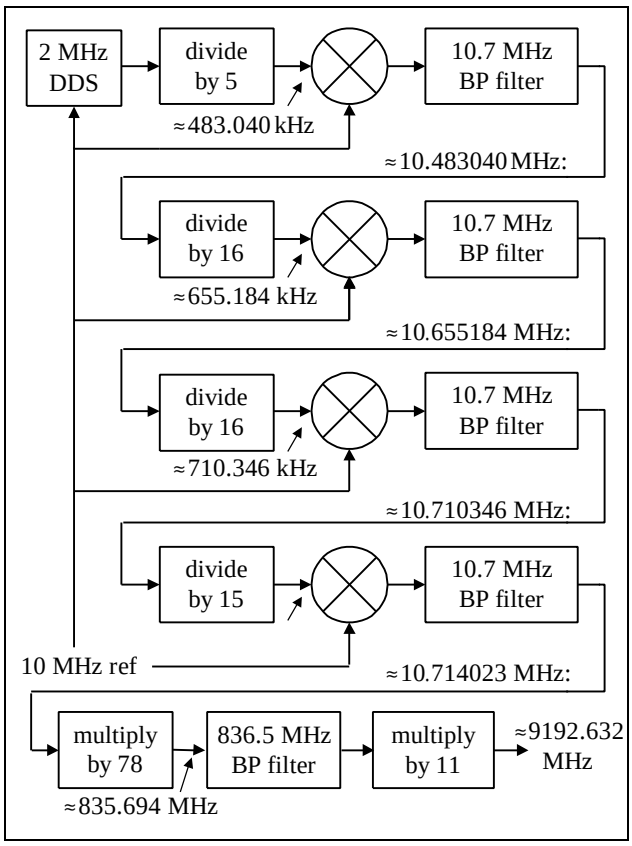

Figure 11. RF chain for cesium standard.

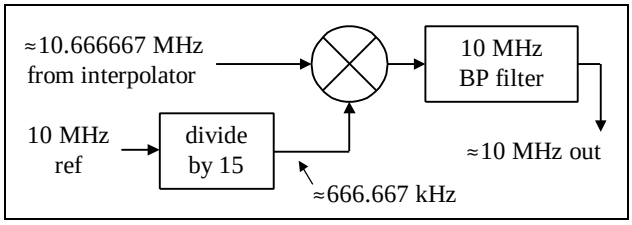

Figure 12. Alternative 10 MHz output architecture.

Figure 12 shows an alternative output architecture where the ÷N in the path from the previous stage is changed to ÷(N - 1) and relocated to the reference path. This modification can also be adapted to the other frequency schemes previously described. Figure 13 shows a 10.23 MHz output synthesizer for use as a GPS timebase locked to a 10 MHz reference. It is adapted from figure 6.

Frequency modulator applications

If the DDS is replaced by a frequency or phase modulator, then the synthesizer can be used to reduce the modulation index.

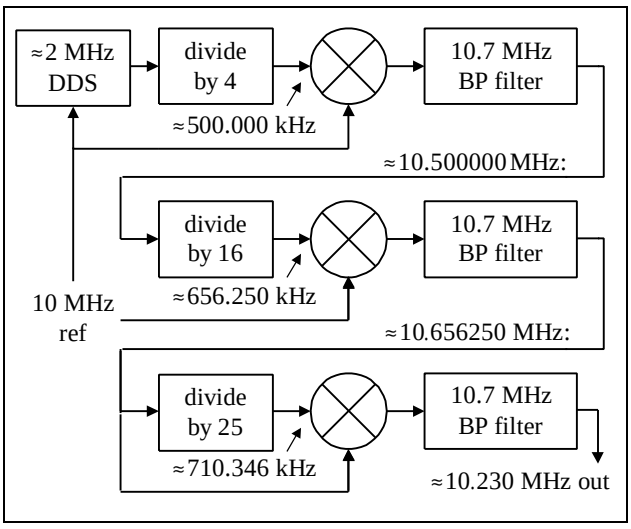

Figure 13. 10.23 MHz out, 10 MHz ref synthesizer.

This allows generating the signal with much wider deviation than the actual signal. The advantage is that any center frequency drift of the carrier is reduced correspondingly. This drift reduction permits the use of a true frequency modulator with DC coupling utilizing a VCO, rather than an AC coupled technique such as an Armstrong modulator. DC coupled modulation is important in asynchronous data transmission. Current workarounds for this problem involve questionable schemes such as putting a very narrow phase locked loop on the VCO so that the data is "almost" DC coupled or using a DDS to generate mark and space frequencies. The problem with using a DDS is that the data pulse shaping filters have to be digital filters. Some systems implement a workaround consisting of a look up table that describes the transition from mark to space or vice versa. This method is suboptimal with respect to intersymbol interference and spectrum utilization.

Experimental verification

A experimental model was built using two interpolators with 3 filters each and a final 10 MHz filter with 5 discrete ceramic resonators. A synthesized signal generator was used in place of the DDS. With a moderate amount of attention to the details enumerated above, it was easy to keep spurious sidebands down at least 120 dB. The ceramic filters did not significantly increase phase noise.

References

[1] D. W. Allan and H. Daams, "Picosecond time difference measurement system," in Proceedings of the

29th Annual Symposium on Frequency Control, 1975, pp. 404-411.

[2] R. L. Hamell, P. F. Kuhnle and R. L. Sydnor, "An Improved Offset Generator Developed for Allan Deviation Measurement of Ultra Stable Frequency Standards," in Proceedings of the 23rd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting (NASA confereence publication 3159), 3-5 Dec. 1991, pp. 209-218.

[3] V. E. Van Duzer, "A 0-50 Mc Frequency Synthesizer with Excellent Stability, Fast Switching, and Fine Resolution," Hewlett-Packard Journal, vol. 15, no. 9, May 1964, pp. 1-6.

[4] R. K. Karlquist, "A New RF Architecture for Cesium Frequency Standards," in Proceedings of the 1992 IEEE Frequency Control Symposium, 27-29 May 1992, pp. 134-142.

[5] G. J. Dick, P. F. Kuhnle and R. L. Sydnor, "Zero-Crossing Detector with Sub-Microsecond Jitter and Crosstalk," in Proceedings of the 22nd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting (NASA Conference Publication 3116), 4-6 Dec. 1990, pp. 269-282.

[6] L. S. Cutler and R. B. Giffard, "Architecture and Algorithms for New Cesium Beam Frequency Standard Electronics," in Proceedings of the 1992 IEEE Frequency Control Symposium, 27-29 May 1992, pp. 127-133.