

Almost Synthesis: An Offset PLL with a Reference Divider.

Wes Hayward, w7zoi, updated 5 Dec 2011, 14 Jan 13, 6 Jan 15 (minor edit)

This note deals with an extremely simple Phase Locked Loop (PLL) frequency synthesizer. There is nothing fundamentally new about the design. In the formal sense, it is just a so-called “offset” loop, a PLL without division where the VCO is offset in frequency from the phase detection. The design has a twist that makes it both practical and low cost. The discussion may also serve as a PLL tutorial for some readers.

Features:

I’ve encountered two seemingly isolated frequency tuning problems in recent times:

- (1) How does one build a local oscillator system that is not only stable enough for SSB and CW, but some digital modes as well?
- (2) How does one build a simple frequency control system that provides good bandwidth and smooth, backlash free tuning while avoiding the mechanical complication of traditional methods?

The scheme described in this note addresses both problems. The listed performance features are provided with some of the digital methods of the day. These include both DDS (Direct Digital Synthesis) and Si570 based designs. But these are not always simple, have medium or even high cost, and require a microprocessor for control. The DDS systems are often compromised by severe spurious responses.

Some Block Diagrams and Concepts.

A traditional PLL is shown in Fig 1 below. A voltage controlled oscillator (VCO) is applied to a phase/frequency detector. The reference for the detector is the output of a crystal oscillator at f_x . The phase detector produces a signal with a DC component that is proportional to the phase difference between the two oscillators. This is amplified in the so-called loop filter. The output is then applied to the VCO frequency control input, usually a varactor diode. The result is that the VCO has exactly the same frequency and phase as the crystal oscillator.

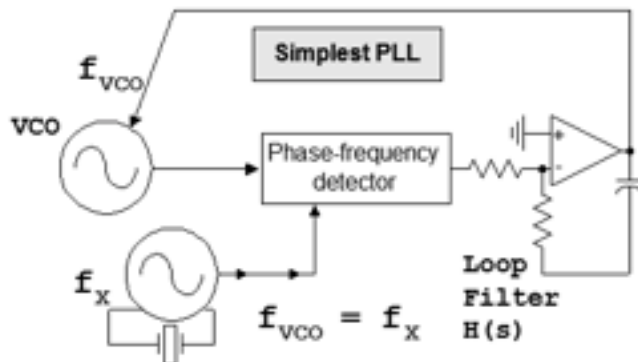


Fig. 1.

The details of the design of the PLL are found in numerous books and will not be repeated here. I recommend the now classic text by Ulrich Rohde, Digital PLL Frequency Synthesizers: Theory and Design (Prentice-Hall, 1983) although other books and various semiconductor application notes are also good. One note of caution: A PLL, in my opinion, is best designed, by the numbers and with some analysis. Many electronic circuits can be built by pulling *catalog* circuits from books, duplicating those circuits, adjusting a resistor value here or there, and applying power. This does not generally work so well with a PLL.

The traditional singleloop PLL using programmable digital dividers is shown in Fig 2. Consider an example. Assume the crystal oscillator produces an output at 1 MHz. The R value is set at 100, producing a reference at the phase detector of 10 kHz. The VCO is designed to operate in the 5 MHz range. Assume that the programming applied to the N divider is for N=503. The VCO frequency will be divided by this amount to yield an output of 10 kHz. This only happens if the VCO operates at 5.03 MHz. Changing N to 504 will cause the VCO to move to 5.04 MHz, and so forth. The range can be quite large here.

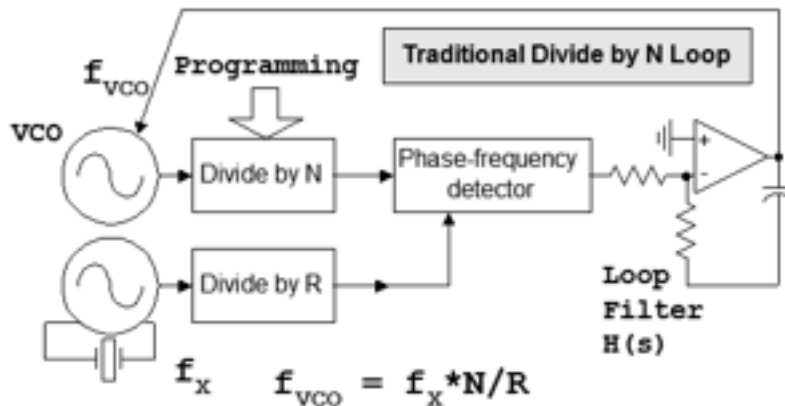


Fig 2.

The classic loop of Fig 2 is limited by many factors, but is still a useful tool. The frequency step is fixed by R and the f_x . Phase noise is a problem with this topology. Noise on the crystal oscillator is applied to the R divider where it is reduced, but still present at the phase detector. This noise, plus excess detector noise, is the reference for the phase detector. This noise eventually makes its way to the VCO where it is larger owing to the frequency multiplication action of the loop. The VCO phase noise is increased by $20\text{Log}(N)$ dB over that present at the detector. For the example discussed, the degradation is 54 dB.

Think of the noise as uncertainty in the frequency at the phase detector. Assume an uncertainty of 1 Hz. The VCO uncertainty is now 530 Hz in the example. Application of modulation analysis produces the $20\text{Log}(N)$ result.

This is not quite as bad as it might seem, for the increase in phase noise only happens within the bandwidth of the PLL. (Loop *bandwidth* is the frequency where the *closed loop* gain is down 3 dB with the measured parameter being the control voltage. This bandwidth is very close to the gain *cross over* frequency of the *open loop* response.) Assume that we have designed the loop to have a loop bandwidth of 1 kHz. Loop BW should generally be much less than the reference frequency. The 54 dB phase noise degradation mentioned above will happen next to the carrier and out to 1 kHz from the carrier. Beyond the 1 kHz separation from the carrier, the ultimate system phase noise is determined primarily by the VCO. Clearly, phase noise performance is vital in a good VCO. It is *the* driving performance parameter.

Phase noise is discussed in virtually all of the texts on PLL design as well as those dealing with transmitter and receiver design. This noise is the central difficulty encountered in virtually all communications equipment, far exceeding the other traditional performance metrics of sensitivity, selectivity, (thermal) stability, and dynamic range. Phase noise is not an esoteric curiosity that only engineers worry about. Rather, it is the root of much of the reciprocal mixing that limits what we can hear, or the quality of what we transmit.

Figure 3 below is an offset PLL with a mixer. The VCO is applied to a mixer with the other mixer input being a crystal controlled oscillator. Mixer output is processed with a low pass filter before being applied to a phase-frequency detector.




Fig 3.

An example of such a loop is included in [Experimental Methods in RF Design](#). (See the discussion beginning on page 4.22.) In the EMRFD example the stability of a 1.5 MHz manually tuned oscillator is transferred to 14 MHz. That design still required the weight, bulk, and mechanical complexity of a dial drive applied to the manually tuned oscillator.

Figure 4 shows the modification that defines the synthesizer of this note. Here, the manually tuned oscillator (MTO) is applied to a programmable frequency divider.

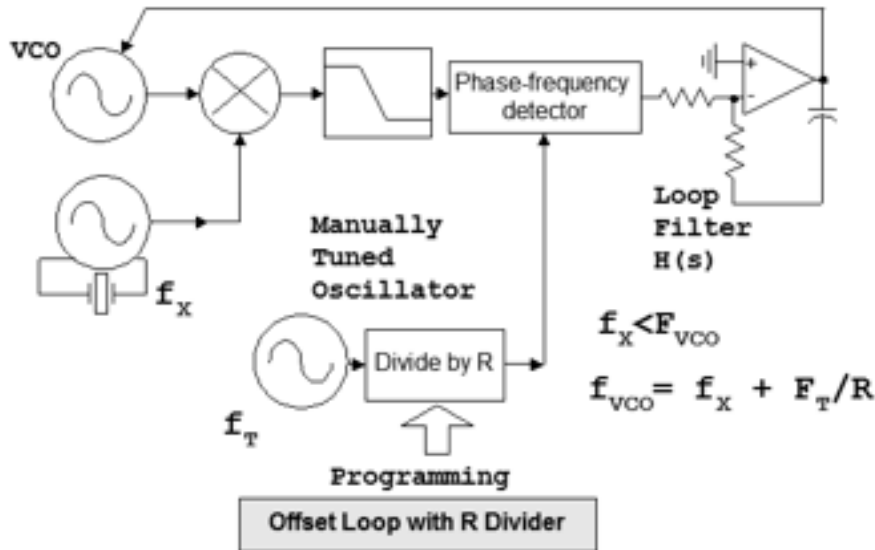


Fig 4.

The utility of this scheme is best illustrated with a numeric example, a source at 10 MHz region and above. A standard catalog microprocessor crystal is found at 9.83 MHz. The difference that emanates from the mixer and low pass filter is 170 kHz. This reference is obtained from a free running L/C *manually tuned oscillator* (MTO) at 2.89 MHz. It is divided by 17 to produce the 170 kHz reference. Assume that MTO is tuned upward to 3.19 MHz, but is still divided by 17. This yields a reference frequency of 188 kHz. The reference frequency has changed by 18 kHz. The VCO will also change by 18 kHz.

This narrow range is both a virtue and a problem. It is useful because we can tune the MTO without a vernier drive mechanism. A large knob works well, leaving no hint of backlash. But the problem is that the range is restricted.

The overall range is expanded by changing the divisor R. Assume that R drops from 17 to 16. The MTO range (2.89 to 3.19) now produces reference frequencies from 181 to 199 kHz, an 18 kHz range that overlaps the first range, but is above it. We will show detailed data below.

If R is made programmable from 10 to 17, a net range of 150 kHz is tuned by the system. The MTO produces the fine tuning, completely free of backlash, while one of eight larger tuning ranges are selected with three toggle switches. The stability is excellent, for an already stable L/C oscillator at 3 MHz is made even more stable by a factor of 10 to 17. The traditional PLL parameter N is 1 in this system. (See Fig 2 above and related discussion.) Hence, there is no phase noise compromise. The spectral purity of the MTO is enhanced by a factor of 10 or more and then transferred to the VCO.

A disadvantage is the redundancy produced by overlapping ranges. This is required in order to allow all frequencies to be achieved within a large tuning range. If there were no overlap at the low end of the overall range, there would be gaps at the high end.

Figure 5 shows the equations that set up the synthesizer and some data using the example already presented.




Fig 5

The Mathcad display shown above is fine for me, but not of much use to others. A computer program was written to do the calculations with an output shown in Fig 6.




Fig 6.

The program also includes a routine to calculate the loop filter R and C values. SPICE can be used to do the classic open-loop Bode plots. This will be shown below.

A system was built using the frequencies of the above example. Tuning curves for this system are shown in Fig 7.




Fig 7. The number on the right are binary and represent the values loaded into the divider. There is considerable frequency overlap at the lower end of the overall range.

The details:

The block diagram for the system is shown in Fig 7, immediately below.

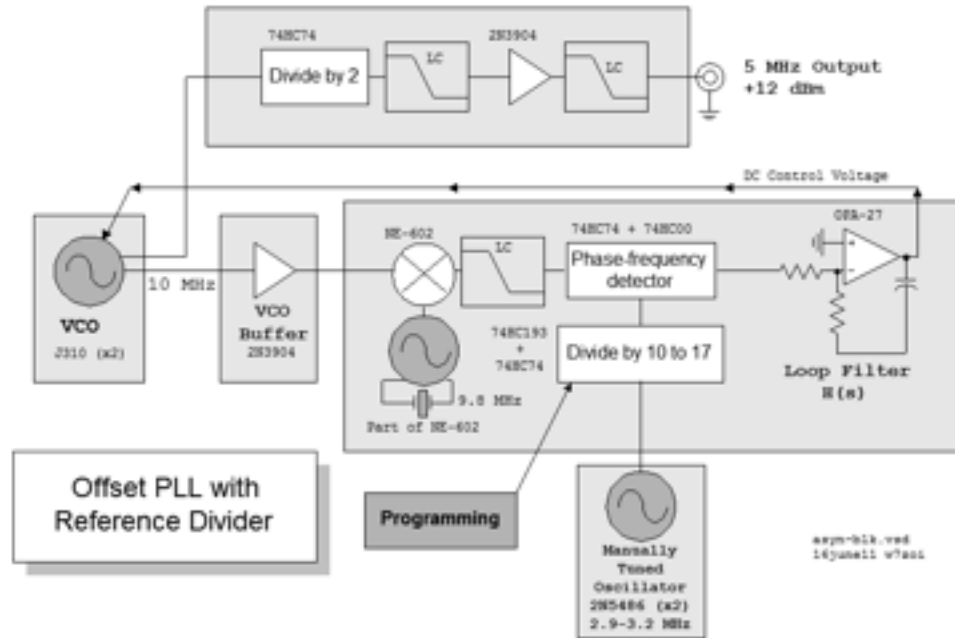


Fig 7. Synthesizer System.

The heart of the box is the VCO, shown in Fig 8. This is an outgrowth of that shown in Fig 4.34 of Experimental Methods in RF Design. Indeed, the original module was modified to get this circuit. The tuning curve is in Fig 9; Fig 10 is an inside view.




Fig 8. VCO




Fig 9. VCO Tuning



Fig 10. VCO.

The schematic for the manually tuned oscillator, MTO, is shown in Fig 11 with a photo in Fig 12.

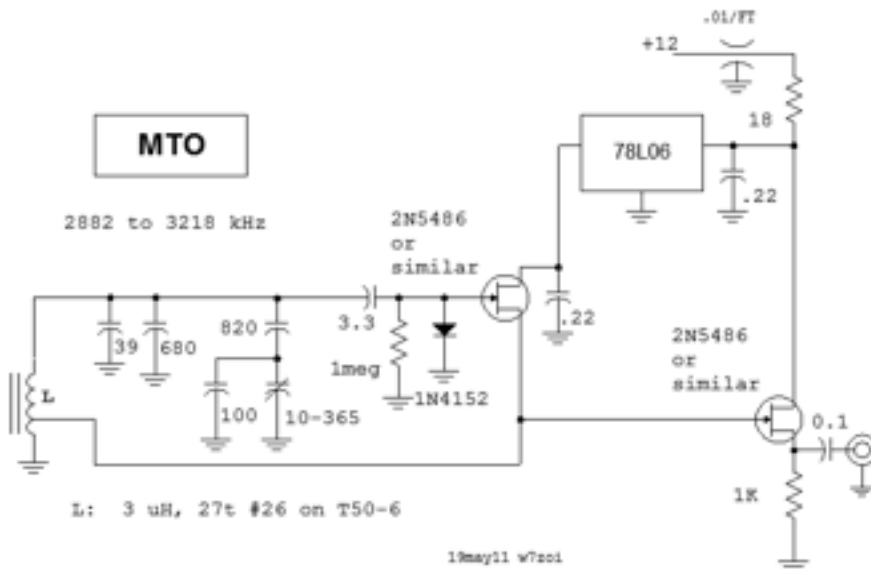


Fig 11. MTO.



Fig 12. MTO.

The details of the system are shown in Figures 13 and 14. The initial attempt to build this system used a 74HC4046 phase/frequency detector. Problems were encountered with loop oscillation. The dependable circuit with a dual flip flop and NAND gate (74HC74 and 74HC00) was built and is now used. A synchronous divider with a 74HC193 and a following 74HC74 form the programmable divider. Mixing is done with a NE602 while a high performance op-amp, the OPA-27, is the loop filter.

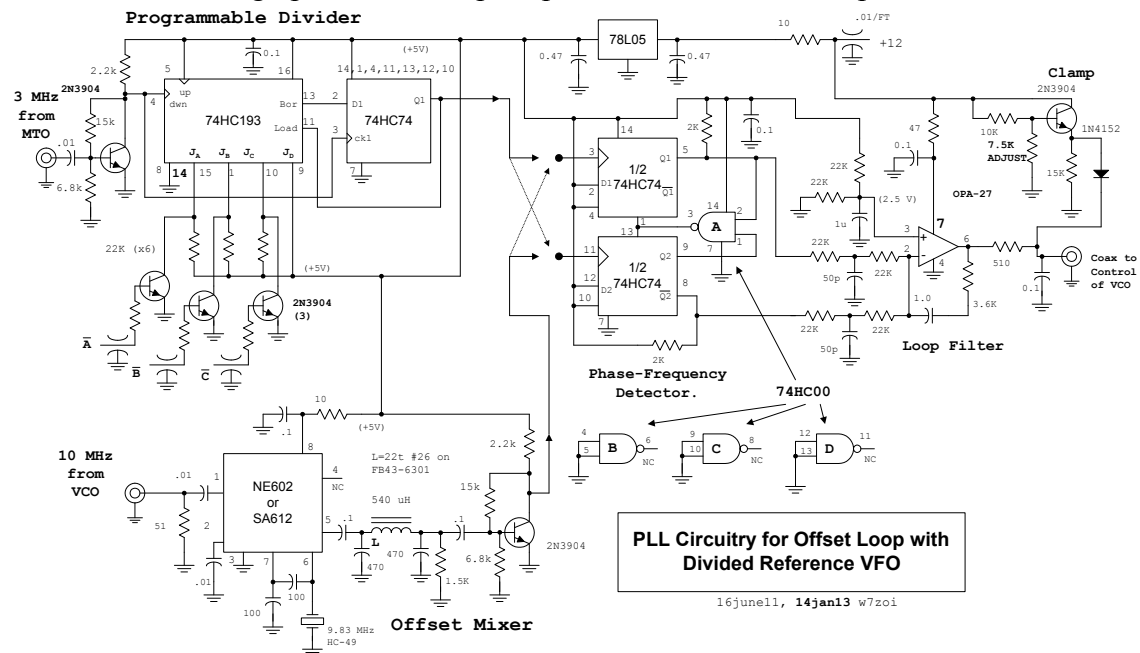


Fig 13. The programmable divider, mixer, crystal oscillator, phase detector, and loop filter are contain in one module. A Clamp circuit prevents VCO operation at a frequency

below the 9.83 MHz crystal frequency. This is important to prevent an unstable lock on the wrong side.



Fig 14. Interior view of synthesizer. The interface from the loop filter to the VCO is with coax.

Once the synthesizer system is operating and “locked,” it must be processed for use. In my application, I needed a 5 MHz oscillator for use with my home station. The 10 MHz from the PLL is divided and amplified with the circuitry shown in Fig 15.




Fig 15. The output

available from this module is +12 dBm.

Spectrum analyzer measurements of the 10 MHz signal showed that the reference sidebands separated from the output by the reference frequency at the VCO were only down by about 65 dB from the carrier. This turned out to result from some of the 200 kHz (approximately) signal that leaked into the VCO module from the synthesizer board. The VCO output was isolated from the signal that was propagated back on the coax from the synthesizer board (Fig 13.) with a common base amplifier, shown in Fig 16. After this circuit was installed, the reference sidebands were over 100 dB down.




Fig 16. This circuit was built in a

small box, but could just as well be incorporated within the VCO module. It should NOT be in the synthesizer module. The transformer can be 12 bifilar turns on an FT-37-43 core.

PLL Analysis with SPICE

The equations for the various elements of the PLL are well known and won't be repeated here. Graphic display of the Bode plot is still useful. The loop that we wish to analyze is shown below (Fig 17.) This is a linear circuit, even though we know that many components are nonlinear. Here, linear approximations are used, as is normally done with PLL analysis.




Fig 17.

The first element to be approximated is the VCO. This is specified in terms of an output frequency versus an input voltage. Also, so far as baseband operation goes, the VCO is an **integrator**. (This arises because frequency is the rate of change of instantaneous phase.) Hence, the VCO is approximated by an ideal integrator with a 1 Ohm resistance and 1 Farad capacitor, followed by an amplifier that has the proper gain. Numeric parameters for our example will be shown in a figure below.

The next element in the chain is the traditional divide by N element. $N=1$ for our application, but the element is included here for general use. The R divider that is not zero in our application does not impact PLL operation, for it is outside the loop. It merely generates the reference at the phase detector.

The phase detector is the next element in the chain. Various detectors will have different gain values. This one is 0.8 V/radian. One can derive this by inspection of the circuit when operating from 5 volts. (See [Introduction to RF Design](#), ARRL, 1994, p322.)

The next element is the loop filter. Actually, that term is something of a misnomer. The op-amp and related components do indeed form a low pass filter element. However, the cutoff frequency is not related to that of the loop. Rather, the op-amp is part of the larger low pass filter formed by all of the loop elements when glued together. The input resistance, R1, in combination with the capacitor, C of Fig 17, are the dominant elements that determine the open loop crossover frequency. R2 injects a stabilizing effect into the loop. If R2 was not there, the gain would be decreasing so fast that loop stability

might be an issue. But the use of the R2 calms the loop, which stabilizes it. R2 also dampens the loop such that transients (e.g., when R is changed) do not cause a ringing response.

There is yet another filter element in the picture presented in Fig 17. This is an output RC filter. The R is a necessary element, for it is best to isolate the op-amp from loading by a shunt capacitor. The capacitor that is present consists usually of an RF bypass part, or even coax cable capacitance. These elements are part of the loop and should be included in the analysis.

There is one final factor that is approximated with this analysis. This is the “transducer” nature of the real elements. Many of the blocks in a PLL are transducers to the extent that the energy output is not the same form as the input. The VCO has an input of Volts, but an output that is a frequency. The phase detector has an input phase, but an output of volts. All of the elements in the analysis are pure “volts in and volts out” blocks. This does not matter. By the time we get through the total loop, we again have volts versus volts. This detail is of no consequence for a mathematical analysis, but comes into play when we model the loop in SPICE. Fig 18 shows an analysis in LT-SPICE from Linear Technology Corp., downloaded on the web.




Fig 18. The input is from V1, a voltage source of strength 1. This is shown as an output in green, the “0_dB_ref.” The loop output signal is “out” from the right side of the circuit. Also shown is a “slope_ref.” This is a signal configured to cross the 0 dB line at 12 dB/octave, or 20 dB/decade. The place where it crosses is forced to equal the place where the signal crosses by adjusting the parameter “slopegain.”

In the example shown, R1 is 44K, which is the series sum of two 22K resistors. In the practical circuit, these are doubled, for there are two outputs from the phase detector. R1 is split in two to allow insertion of a small bypass cap, 50 pF in our case. This removes the “sharp edges” without really impacting the loop.

Examination of the region of the curve where the output crosses 0 tells most of the story. A stable loop should have gain crossing at a rate less than the 12 dB/octave rate of the blue trace. It is, however, suitable for the rate to increase beyond loop crossover. It is interesting in this loop to change the R and C values for the output filter and look at the response.

Results

The first thing that one does with a PLL is to apply it to a spectrum analyzer. The results are shown in Figures 19 and 20.




Fig 19. The programming is “000.”




Fig 20. Programming = "111."

The difference between the two cases is the shift in center frequency and spacing of the reference sidebands changes. (Analyzer BW is 30 kHz). These measurements were taken before the buffer amplifier was added, which pushed the reference sidebands well into the noise. The measurement was done by attenuating the strong carrier while leaving the sidebands unchanged. This was done with a crystal notch filter. (See "Oscillator Noise Evaluation with a Crystal Notch Filter," QEX, July, 2008)

Figure 21 shows the interior of the synthesizer. The photo was taken before the buffer amplifier was included in the system. Fig 22 shows the system in use as a second LO for a CW transceiver.




Fig 21.



Fig 22.

Although we have not yet measured phase noise, it appears to be excellent. Measurements were performed in recent contest situations to compare the synthesizer with an LC oscillator; no reciprocal mixing was observed.

Conclusions

This variation on the simple PLL has turned out to be a handy one. With $N=1$, the usual compromise in phase of a divide-by- N with high N is avoided. The division by R enhances the stability of what is nominally an LC oscillator based system. In this implementation, an already stable 3 MHz LC oscillator produced drift that was further reduced by a factor of 10 to 17.

The scheme has the additional advantage of enhanced bandwidth. An ancillary result of the slow tuning is the elimination of mechanical complexity, a major factor for craftsmanship challenged experimenters such as myself.

This scheme is completely traditional with no special integrated circuits or microprocessors required.

But this system is not the ultimate solution. The method is useful when a narrow tuning range is required, as was the case with my committed CW transceiver requiring a geometric LO tuning range of only 1.5%. It is not nearly as attractive when a more typical frequency range of 10% or more is sought. At a more subtle level, overlapping incremental tuning ranges are a minor problem. Some overlap between increments is

necessary, but it is a minor annoyance to reach the end of one range and need to retune the MTO, increment the digital word, and find the previous frequency.

This synthesizer system was used for frequency control. Most synthesizers that use a microprocessor also use the processor to generate a digital readout. My receiver already had a readout.

Toggle switches represent both a feature and problem. Push buttons to go up or down to control divisor R would be easy to implement. The toggle switches allowed easy scheme to test the concept and generated a system that would be especially useful for a portable set up where battery power might be critical. Note that one could, in some circumstances, get by without a digital readout so long as the R value was known

Hybrid or mixed schemes might be a better solution. That is, traditional divideby N schemes are reasonable so long as N is not excessive.