application note

Fractional-N Synthesizers

David Owen

IFR (and formerly Marconi Instruments) owns the key IPR associated with Fractional-N synthesis technology. This application note provides historical information on fractional-N and how the IFR system works.

In 1984 John Wells of Marconi Instruments invented a new solution to the problem of Fractional-N Synthesizers that required the use of no analog components to correct the noise and spurious normally introduced by fractional divider schemes. Marconi Instruments made extensive use of this Fractional-N Synthesizer following the introduction of the first signal generator, the 2030 and 2031, that exploited the technology.

In April 1995 Marconi Instruments was awarded the Queens Award for Technological Achievement for work on Fractional-N Synthesizer systems. Since the acquisition of Marconi Instruments by IFR the company has continued to exploit this synthesis scheme in its new products. The technology has been licensed to a number of companies.

This article describes Fractional-N Synthesizers and the associated systems that IFR and Marconi Instruments have used in order to implement the technology in commercially available products.

Basic Fractional N

Fractional-N Synthesizers have been used for many years to improve the performance of indirect frequency synthesizers. A simple indirect synthesizer uses a phase-locked loop (PLL) containing variable division ratio dividers. The frequency resolution of such a synthesizer is limited by the phase detector rate - if a 1 kHz phase detector rate is used, then the resolution will also be 1 kHz. However, many applications demand much finer frequency control, and for this reason, a technique known as Fractional-N Synthesis was developed.

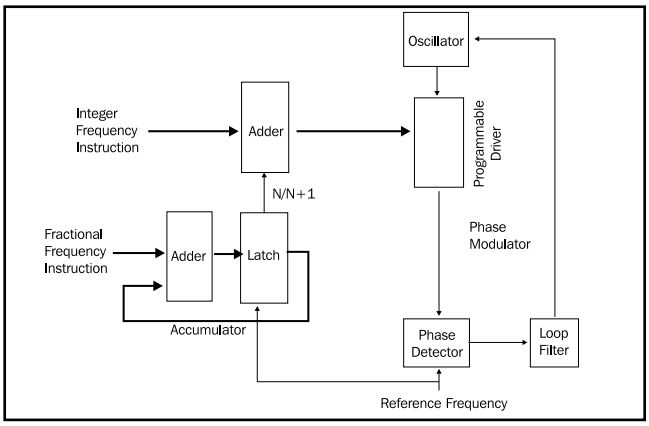

Figure 1 - A basic Fractional-N system with no jitter correction based on a single accumulator

Fig. 1 shows the principle. The division ratio of the divider is made to have a fractional component by changing the division ratio of the divider periodically, so the average value contains a fractional element. If, for instance, a fractional value of 0.1 is required then the division ratio is changed by 1 every tenth cycle. If a fractional value of 0.01 is required then the division ratio is changed by 1 on every one hundredth cycle. This offers much finer frequency control than integer based systems.

An accumulator whose digital output is incremented for each cycle of the divider by the fractional frequency instruction is a convenient method of controlling the division ratio. The accumulator uses an adder latch to add the contents of its input to its current output on each cycle of the clock. It behaves as the digital equivalent of an integrator and since the integral of the frequency is phase its output represents the relative phase of the fractional component. Everytime the accumulator reaches its capacity it produces an overflow which changes the divider division ratio.

In the above example if the accumulator has a 0.1 fractional instruction added to its output for each clock cycle then the accumulator will be full and overflow every tenth cycle. In a similar way, if the fractional instruction is .01, the accumulator will overflow every 100 cycles.

There is a price to be paid for the improvement in the frequency resolution that these systems provide. The manipulation of the divider ratio generates phase perturbations, and hence spurious signals, which have to be eliminated in a useful synthesizer design. The nature of the phase perturbations is predictable and can be cancelled using an analog correction system. Solutions that use various techniques to improve the performance of PLL's based on the Fractional-N system have been patented[1,2,3,4,5].

Marconi Instruments used a Fractional-N Synthesizer system with analog correction in the design of its very popular 2022 10 kHz to 1 GHz Signal Generator (Fig. 2) to provide 10 Hz resolution based on a phase detector rate of 40 kHz. An accumulator that changes the division ratio by 1 each time the accumulator overflows controls the divider.

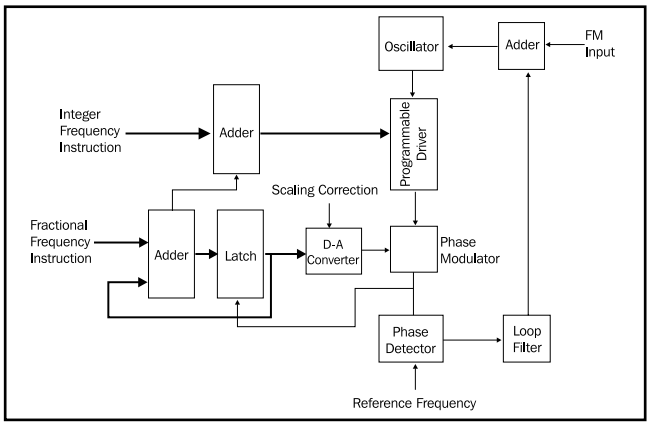

Figure 2 - An example of an early Fractional-N synthesizer used in 2022 signal generator using analog correction

The residual content of the accumulator represents the residual phase error in the PLL and, by routing this signal to a D-to-A converter and applying the analog signal to a phase modulator within the PLL, the residual errors can in principle be eliminated as shown in Figure 3.

With a small fractional frequency instruction at the input to the Fractional-N system the output of the accumulator increases its digital value on each cycle of the phase detector clock. The residual phase error in the PLL also accumulates in the same way. When the accumulator overflows (i.e. its content becomes greater than the capacity of the accumulator) it overflows and the divider division ratio is changed by one for one cycle of the phase detector. Changing the division ratio by one for one cycle of the phase detector effectively absorbs one cycle of the VCO frequency, and hence introduces a 360° phase shift. The D-A converter produces an analog representation of the PLL accumulated phase error and is used to phase modulate the PLL to cancel the analog error.

As shown in Figure 2, the amplitude of the correction system also has to be scaled as the average output frequency (and therefore the average division ratio) in the PLL changes. This arises because the phase modulator acts at a constant frequency (the phase detector rate) and a phase change (or time change) at the detector has a different phase impact on the VCO according to the division ratio of the divider.

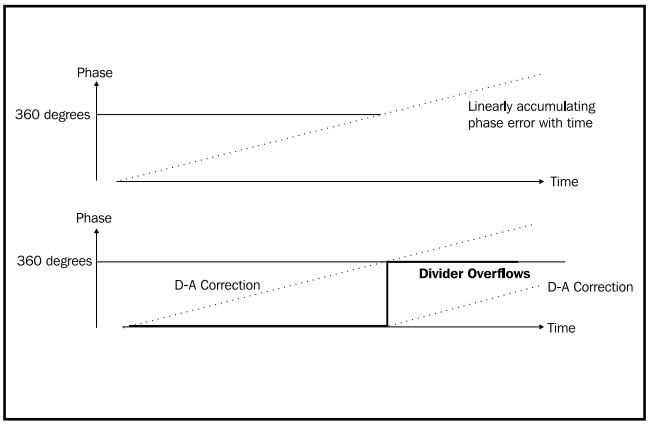

Figure 3 - Analog correction of the phase jitter in a Fractional-N system

However, there is a limit to the accuracy of the cancellation that can be achieved and obtaining cancellation of better than 1% can be troublesome. In addition, the use of analog components to correct for the error results in a design that requires careful setting up and can drift with time

The need to apply analog correction to such synthesis schemes also has the effect of restricting the phase detector rate that can be used. This stems from the need to run the correction circuitry conservatively in order to maintain good accuracy and linearity, and from the need to avoid the injection of wide-band noise from the correction system.

For optimum performance, in terms of speed and noise, a PLL should be operated at the highest practical phase detector rate to minimize divider ratios and give the highest margin between the loop bandwidth and the phase detector frequency. It is clear that there are significant design compromises needed when designing PLL's of this type.

King[6] described an improved version of a Fractional-N system. This used two accumulators, but still required analog correction in order to minimize spurious responses.

An ideal Fractional-N system would not require any analog components in order to cancel these residual phase errors. Intuition at the time suggested that this was not a practical possibility, but an all-digital solution was found. The all-digital solution results in a PLL performance that is virtually indistinguishable from the non-fractional loop operating at the same phase detector rate - and is capable of operating at much higher phase detector rates than those which relied on analog correction techniques.

John Wells' Fractional-N System

A simple implementation of this system is shown in Fig. 4.

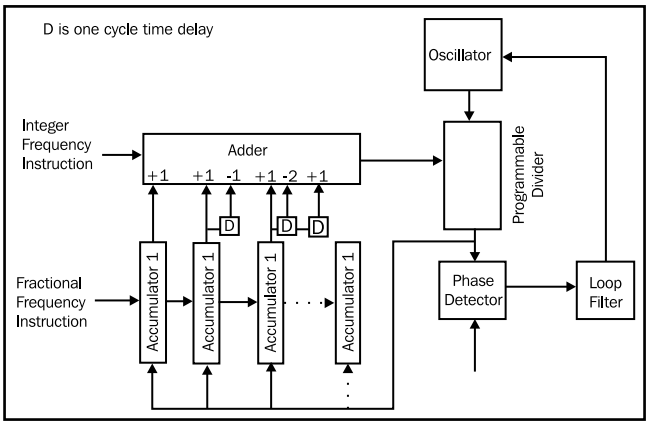

Figure 4 - Fractional-N system using digital correction

This digital Fractional-N system is based on the principle of noise shaping[7]. Rather than trying to cancel the phase errors, this system modifies their pattern so as to minimize low-frequency spectral content. The technique actively reduces the generation of these low-frequency signals and exchanges them for increased levels of higher-frequency signals.

Within the context of the PLL this is a highly satisfactory arrangement because the loop itself acts as a low-pass filter. Additional low-pass filtering can then be included within the loop filter to prevent any of the modified perturbations from significantly modulating the voltage-controlled oscillator.

In the example of Fig. 4, the single accumulator of the earliest forms of Fractional-N Synthesizer is replaced with three or more accumulators, the output of each being connected to the input of the next. The overflow from each of the accumulators manipulates the division ratio of the divider.

Effect of accumulator overflows

The first accumulator overflow acts in the same way as the accumulator in the simplest Fractional-N systems; it changes the division ratio divider from N to N+ 1 for one cycle when the accumulator overflows. The remainder output from the first accumulator represents the phase error that would result if no other correction were applied.

The second accumulator digitally integrates the output of the first accumulator and subsequent accumulators repeat this process. The question then arises as to how to manipulate the division ratio of the divider in order to reduce the low frequency phase jitter that it produces on its output. The solution is that the overflow from the second accumulator needs to manipulate the division ratio by the differential of the effect of the first accumulator. In a similar way the output from the third accumulator manipulates the division ratio by the differential of the effect that an overflow output from the second accumulator causes, and so forth.

The first accumulator produces a change of division ratio that changes from N to N+ 1 and then back to N.

The second accumulator carry overflow is required to generate a sequence that changes from N to N+1 and then to N-1 and then to N, when an overflow is initiated.

In a similar way the third accumulator overflow changes the division ratio from N to N+ 1 then to N-2, then N+ 1 and then returns to N.

A fourth accumulator would use a sequence N+ 1, N-3, N+3 and N-1.

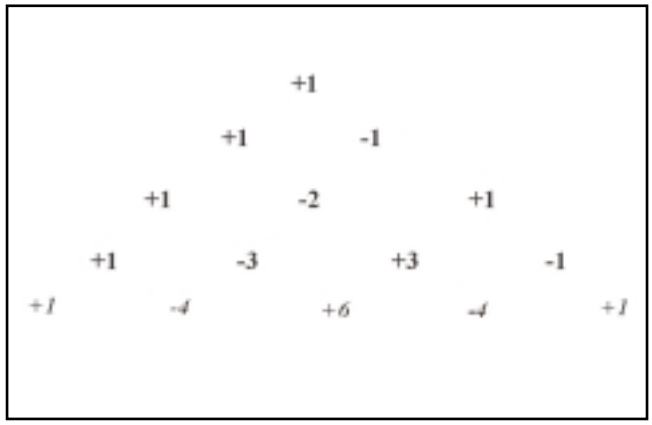

Figure 5 - Pattern of division ratio changes caused by accumulator overflows shown as Pascal's Triangle

The sequences that are generated correspond to Pascal's triangle but with the overall sum of each row being weighted to be zero, with the exception of the first which is required to correct the overall division ratio of the divider in order to obtain the required fractional frequency. Because the second and

subsequent rows of the Pascal triangles introduce an average divider change of 0, these accumulator overflows have no long-term effect on the division ratio of the divider. However, they do serve to remove low-frequency components from the divider's output spectrum and to transfer the energy to higher frequencies where they can be successfully filtered by the PLL.

As shown in Fig. 4, the accumulator overflows can all be routed via a cascade of single-clock-period delay networks (implemented by D-Type flip-flops) and an adder with weighted inputs so that they generate the required division ratio changes.

Alternative circuit implementation

The original patent was drawn in the manner described above. However, the same effect can be obtained from the redrawn version shown in Fig. 6.

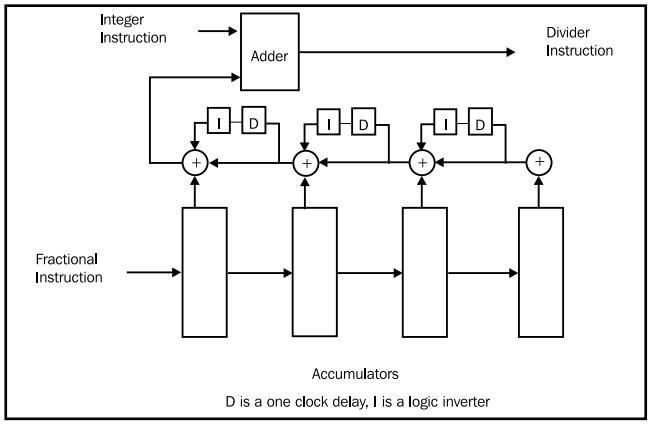

Figure 6 - A different circuit implementation of the Fractional-N system

The accumulators can be considered to be digital integrators. The weighting functions generated by the overflow from them are generated by differentiating the effect of the previous accumulator's overflow.

An overflow is differentiated by turning an accumulator overflow into an N+1, N-1 sequence. This sequence is generated by initially passing the overflow directly to an adder, and then passing an inverted version one-clock cycle later to the same adder. If you track the resulting signal from each overflow through the chain of summers, delay (D) and inverter (I) circuits you can build up the same Pascal's triangle sequence as that disclosed in Wells. Replacing the differentiating function (N converted to N+1, N-1) with an alternative function leads to sequences other than Pascal's triangles with broadly similar levels of performance, however the Pascal's triangle solution is the simplest and optimum solution requiring the least division ratio changes.

Understanding the improvement

A pictorial representation of how the phase jitter is reduced is shown in Fig. 7. In this diagram the fractional component of the frequency instruction is relatively high, but it does give an indication of the sort of mechanisms at work. The target phase is a representation of the required phase of the VCO if there were no phase jitter present. In the first plot, using one accumulator (the simplest form of Fractional-N), after a period of time the division ratio of the divider is changed from N to N+1. This introduces an abrupt 360° change in phase that corrects the accumulated phase error of the VCO. The residual phase jitter is clearly a ramp with the residual phase error being equal to the remainder output of the first accumulator, the same as was disclosed for simple Fractional-N systems.

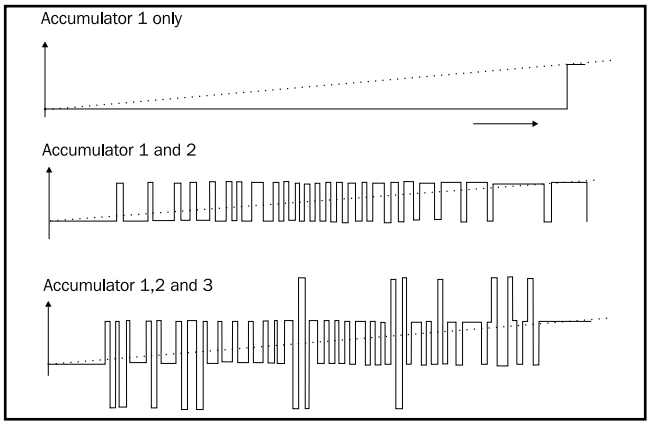

Figure 7 - Pictorial view of the effect of the accumulator overflows on residual phase error in the PLL

The second plot shows the effect of introducing a second accumulator. The second accumulator overflows much more often than the first and a phase perturbation is introduced much earlier.

When the accumulator overflows, division ratios of N+1, N-1 on successive clock cycles is introduced. As can be seen from the plot, the area under the positive phase transient is approximately equal to the area below the target phase. After this overflow the average area of accumulated phase is now much smaller than it would have been. This sequence of events is repeated through the fractional cycle. As the content of the first accumulator increases (that is, getting closer to an overflow) the second accumulator overflows more often. If you imagine what the plot would look like after passing through a low pass filter the error is clearly less than that shown in the first plot.

The removal of the low-frequency content in the output from the phase detector is not, however, perfect and needs to be further reduced in a useful synthesizer design. The third accumulator over-flows more frequently than the second accumulator and the third plot on Fig. 7 shows the effect of this third accumulator. The first N+1, N-2, N+1 sequence starts at an earlier point in time. In fact, the accumulator overflows so often that overflows tend to be superimposed on top of one another. The phase correction starts much earlier in the cycle and occurs much more often. Greater division ratio changes occur, but if the results are passed through the PLL filter, the end result is a signal whose low-frequency content is much lower, and because most of the high-frequency content can be removed by the PLL filter, the overall performance of the loop is much improved.

The impact of this process can be simulated and plotted. The impact of the noise and jitter in Fractional- N systems averaged over a large number of possible fractional division ratios is a flat (white) noise signal. This is not surprising since the peak to peak jitter is always 360° - the first accumulator limits the peak error. A similar analysis can be performed on the three accumulators using commercially available mathematical tools running on a PC. When the original simulations were done on Fractional-N it took up to one hour of CPU time on a main frame computer to do a simulation representing less than one second of operation of the system. The results of a simulation are shown in Figure 8.

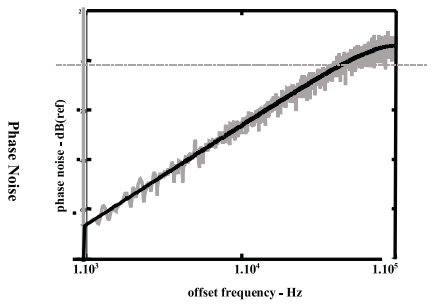

Figure 8 - Noise shaping introduced by the Fractional-N system

Instead of being flat noise the noise profile is frequency shaped, the low frequency noise being much lower than the 0 dB reference line. The higher frequency noise is actually higher then an uncorrected system (as usual you never get a gain without a penalty). The slope of the noise is dependent on the number of accumulators, the more accumulators the steeper the slope and the more the high frequency noise is lifted. It can be shown that the accumulators can be added ad infinitum to shape the resulting phase noise progressively and give less and less low frequency content. In practice, however, most systems require only three accumulators before the resulting spurious signals are virtually eliminated from a practical design.

There are good reasons for not adding more accumulators than is required because as the number of accumulators increase the division ratio excursions and high-frequency signal content will increase. This results in greater excursions in peak phase error at the input to the phase detector, and non-linearity in the phase detector can cause high frequency products to be intermodulated down to lower frequencies. For this reason the phase detectors used in the IFR (Marconi Instruments) implementation of this Fractional-N system use a highly linear, digitally based phase detector that avoids the dead zones and non-linearitys associated with the more conventional dual D-Type phase detectors. However, continued addition of accumulators creates the need for an ever-widening range of division ratios and that causes problems with the design of programmable dividers. The ultimate limit, of course, is set by the need to keep all the ratios positive. The design compromise is essentially one of using enough accumulators, but not more than is required since it would increase the division ratio required in the PLL and require greater levels of performance on the phase detector.

Comparing performance

The performance of a particular design can be assessed by operating the Fractional-N Synthesizer system with the fractional component of the frequency instructions set to zero. If the content of the first and subsequent accumulators are all set to zero, then with no fractional instruction into the first accumulator, the synthesizer will behave as a conventional integer PLL. If, however, a random number is inserted into the output of the first accumulator then all of the subsequent accumulators will change and generate accumulator overflows. This will result in the synthesizer generating sequences of division ratio changes that have no net effect on the division ratio but if the design of the PLL is imperfect it will degrade the performance of the Fractional-N Synthesizer system because of non-linear behavior. With careful design, it is possible to make a Wells' Fractional-N Synthesizer, whose performance is the same in the integer mode as in the fractional mode.

DC-Coupled FM

For synthesizer schemes used in signal generators there is a requirement that the system should enable the use of frequency modulation. Many applications, such as FSK (frequency shift keying) systems used in paging, require that the frequency modulation allows DC control of the synthesizers output frequency. Traditionally this has been a difficult technical problem, often "solved" using complex analog solutions to implement which give marginal performance.

The IFR Fractional-N system lends itself to a DC-coupled FM system, patented by IFR (Marconi Instruments)[8], based on the use of a 1-bit A-D converter as shown in Fig. 9. As with the Fractional-N system this A-D converter is substantially over-sampled so that it can be noise-shaped, in much the same way CD quality analog-to-digital converters now operate. The system is extremely effective when used with the new Fractional-N system because it virtually eliminates carrier frequency error in the DC-coupled FM mode.

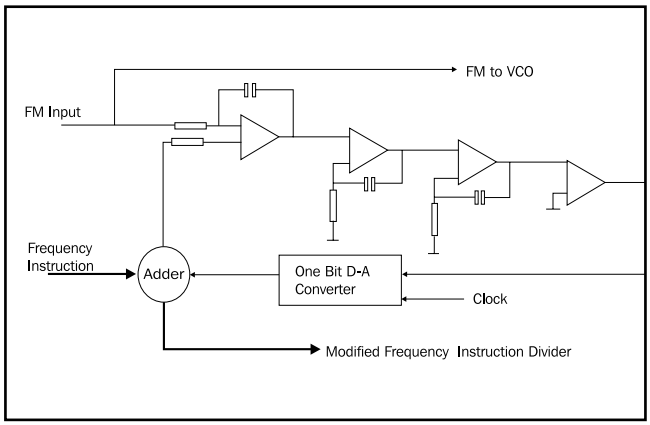

Figure 9 - DC coupled FM circuits added to the Fractional-N system using one bit A-D converter technology

The FM signal is applied directly to the VCO to produce FM outside the PLL bandwidth. Inside the PLL BW the Fractional-N divider is used to manipulate the frequency setting of the synthesizer in sympathy with the required FM signal. The Fractional-N Divider is controlled by applying the analog signal to a third order feedback system deployed around a one bit digital to analog converter - it simply decides whether the error signal is greater than or less than 0 volts for each clock cycle of the phase detector.

The analog feedback is stabilized by making two of the three integrators have unity gain above a break frequency so that the loop is stable. The gain of the feedback system increases with lower modulation frequencies, so the noise of the signal sent to the divider is much lower at low offsets than at high offset frequencies. To understand how the system operates it is not necessary to understand the detail of the pattern of logic 1's and 0's that the digital output produces. Instead consider the system to be a simple feedback circuit with a digitizer (the one bit converter) adding noise (quantization noise). The feedback system automatically forces the pattern of data to reduce the noise at low frequencies since the gain of the feedback increases with lower frequency. Just like the Fractional-N system the DC FM system is noise shaped and takes advantage of the fact the PLL rejects the higher frequency components.

The magnitude of the frequency (divider) changes introduced by the 1's and 0's from the DC FM system determines the sensitivity of the DC FM system.

The result is that the DC FM system is exceptionally well behaved. The DC error in the system is simply related to the DC errors in the analog operational amplifiers used in the system - there is no reliance on RF components in oscillators being

stable with time or temperature.

The use of two port modulation (modifying the oscillator and the divider with an FM signal) requires that the two paths are calibrated so they have the same gain and their phases are matched. A separate IFR patent[10] discloses an automatic system that achieved this with a coherent detector arrangement to match the gain of the DC FM path and the VCO path. The DC FM system based on a one-bit converter also has the great advantage of introducing very little time delay into the DC path, making it relatively easy to match the DC path and the analog (VCO) path. What errors in phase are present can be matched by using digital prediction systems.

Commercial Products

The first products to use the new Fractional-N system were the 2030 and 2031 Signal Generators that were introduced in November 1989. Covering frequency ranges of 1.35 GHz and 2.7 GHz, respectively, these generators offered 0.1 Hz frequency resolution and a DC-coupled FM performance which is still the best available on any signal generator.

The range of signal generators was further enhanced by the 2032 5.4 GHz Signal Generator (still with 0.1 Hz resolution) and a range of low-noise signal generators, the 2040 series, which relied on the Fractional-N system to implement a novel low-noise architecture[9].

The system is incorporated in the microwave synthesizer used in the 6200 series Microwave Test Sets. The inherent speed advantages and simple design has resulted in a low-cost, agile, synthesized source which can allow fast scalar and distance-to-fault measurements on microwave systems to be made.

Radio test sets 2965, 2945, 2966, 2967, 2968 and 2935 rely on the Fractional-N system to provide synthesizers for both their signal generators and their spectrum analyzer local oscillators.

The 2023, 2025, 2026 signal generators and the first ever broadband VXI signal generator, 3002, have a more modern implementation of this synthesis technique to achieve new levels of cost versus performance levels.

Exploitation of the technology has continued in many IFR products. The synthesizers and local oscillators in the 6800 range of Microwave Test Sets and the 2310 Tetra Analyzer are based on the Fractional-N system and the 2319 RF Digitizer uses Fractional-N local oscillator systems.

The Fractional-N system has also been used in radio products by Marconi Communication Systems Ltd. to provide a cost-effective, high performance local oscillator in professional radio designs.

The technology has been licensed to Hewlett Packard (now Agilent) and Rhode and Schwarz for use in test and measurement products.

References

1 BRAYMER N.B, 'Frequency Synthesizer'. US Patent 3555446, 12 January 1971.

2 GILLETTE G.C, 'Frequency Synthesizer', US Patent 3582810, 1 June 1971.

3 UNDERWOOD M.J, 'Frequency Synthesizer'. UK Patent 1447418, filed 29 March 1974, published 25 August 1976.

4 COX R.G, 'Frequency Synthesizer'. US Patent 3976945, 24 August 1976.

5 KINGSFORD-SMITH C.A., 'Device for Synthesizing Frequencies which are Rational Multiples of a Fundamental Frequency', US Patent 3928813, 23 December 1973.

6 KING N.J.R, 'Phase-Locked Loop Variable Frequency Generator'. US Patent 4204174, 20 May 1980.

7 WELLS J.N, 'Frequency Synthesizers'. European Patent EP 125790, filed 11 April 1984, published 21 November 1984.

S THOMAS. N.E, European Patent EP 322139, filed 8 December 1988, published 28 June 1989.

9 OWEN D.P, 'Frequency Synthesizers'. British Patent GB 2103035. filed 4 August 1981 published 9 February 1983.

10 OWEN D.P, and WELLS J.N, US5079522, Variable Frequency Signal Generator, priority Oct 20 1989

Copies or extracts of these patents can be obtained from the Patent Office web sites.

IFR Americas, Inc., 10200 West York Street, Wichita, Kansas, 67215-8999, USA. E-mail: info@ifrsys.com
Tel: +1 316 522 4981 Toll Free USA: 1 800 835 2352 Fax: +1 316 522 1360

IFR, Longacres House, Six Hills Way, Stevenage SG1 2AN, United Kingdom. E-mail: info@ifrsys.com
Tel: +44 (0) 1438 742200 Freephone UK: 0800 282 388 Fax: +44 (0) 1438 727601

Part no. 46891/890
Issue 1