Create signals having optimum resolution, response, and noise

Simultaneously achieving fine frequency resolution, fast switching speed, and low phase noise is the hallmark of the signal-generation technique known as direct digital synthesis.

Earl McCune Jr, *Digital RF Solutions Corp*

Direct digital synthesis (DDS), also known as direct digital frequency synthesis (DDFS), is a newcomer to the toolbox of engineers who develop hardware for generating signals and waveforms. Because of component tolerances, value variations with time, and manufacturing inconsistencies, traditional analog techniques can only approximate a desired signal. In contrast, DDS calculates the signal directly.

DDS is known as a numeric—rather than a digital—technique. Two concerns account for this categorization: one technical, the other more marketing oriented. The technical reason stems from the direct calculations; the signal is actually generated by manipulating numbers. Although "digital" can have the same meaning as "numeric," digital can also refer to signals having (usually) two fixed amplitude levels. DDS does use techniques that are digital in both senses, but the two most important aspects of DDS are the numeric means it uses to represent quantities and the inherent precision that results from its use of numbers.

The marketing-oriented reason for calling DDS a numeric technique relates to the constraints imposed by traditional analog design. To maintain waveform inaccuracies of 0.1% or less (60-dB dynamic range), designers have avoided digital circuits. Such circuits have a reputation for generating noise currents that degrade signal purity in sensitive analog circuits nearby. With DDS, the digital circuits actually generate the analog signals. Using "numeric" rather than "digital" as the descriptor helps to dissociate DDS from digital circuits' unsavory reputation.

You can view DDS as an extension of digital signal processing (DSP) in accordance with the ideas presented in **Fig 1**. DSP has been around for many years. It is often used for filtering signals after an analog-to-digital conversion. DSP is also used to transform such digitized signals. For example, DSP techniques such as the FFT can transform a time-domain signal into an equivalent frequency-domain signal. These uses are

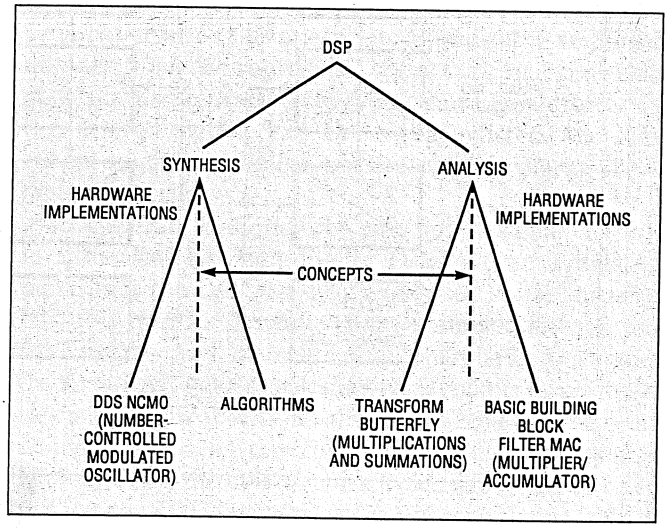

***Fig 1—You can view DDS** as closely parallel to DSP. Both subject areas have sets of mathematical and hardware tools whose functions are intimately related.*

What is really important about DDS is the numeric means it uses to represent quantities and the inherent precision that results from its use of numbers.

analytical; they take an existing signal and change it.

Signal synthesis doesn't begin with an existing signal. DDS takes a small set of parameters (numbers) that describe the desired signal and generates a number sequence that represents the signal. This number sequence usually undergoes a digital-to-analog conversion to finally produce an analog signal. (See **box**, "The sampling theorem backwards.")

A major motivation for the development of DDS is achieving high accuracy at moderate cost. Calculators selling for $9.95 are accurate to 12 digits, whereas analog systems must incur large costs to maintain 0.1% accuracy—equivalent to three digits. By maximizing the use of digital techniques, DDS generates accurate analog signals inexpensively.

There are several ways to implement DDS. Where generating high frequencies is unnecessary—for example, in the voice band—you can implement DDS with general-purpose microprocessors. Low- and medium-speed phone-line modems have been built this way for years. As the required signal frequency increases above the audio range, the computing overhead for signal generation increases proportionally. Somewhere in the low RF (radio-frequency) range, the computing requirements become prohibitive even for modern, high-speed general-purpose μPs. At these higher signal frequencies, you should consider implementing DDS with dedicated hardware.

Dedicated DDS devices are optimized for signal synthesis. Their inputs are the signal parameters; their output is the desired signal. The control processor only needs to keep up with the signal parameters, not with the complete signal. The DDS device acts as a peripheral, freeing the main processor for other tasks.

Fig 2 shows the basic block diagram of a direct digital synthesizer. The DDS has three main blocks: the digi-

The sampling theorem backwards

Direct Digital Synthesis (DDS) obviously belongs to the synthesis side of Digital Signal Processing (DSP). There is no signal to be sampled or processed; rather, there is a sampled signal to be used. In this respect, DDS operates the sampling theorem in the reverse of the usual direction.

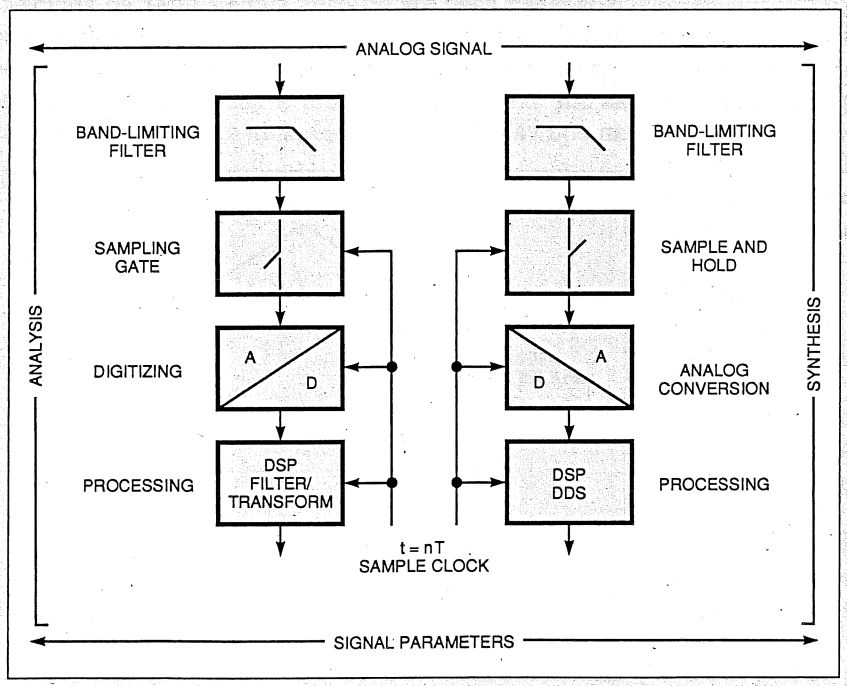

***Fig A—Waveform analysis** using DSP and waveform synthesis using DDS involve similar operations. But DSP and DDS reverse the order of the operations.*

The left side of **Fig A** shows the sampling process as you would conventionally apply it in DSP. You must first band-limit the input signal with an anti-aliasing filter, after which the signal is sampled and digitized. The digitizer provides a number sequence that can be further processed to identify characteristic parameters.

With DDS, the characteristic parameters exist first. The number sequence is generated from these parameters and then converted into an analog waveform. The alias signals characteristic of a sampled signal are then removed with a band-limiting filter. The math is the same; the order of performing the operations is different.

tal accumulator, a waveform map, and the digital-to-analog converter. The input parameters are the signal frequency, represented by a number, and the timebase clock. The whole assembly is often called a number-controlled oscillator (NCO).

The objective of the NCO is to produce a signal s(t) according to the basic signal equation

$$s(t) = A \cos(2\pi \cdot f \cdot t). \quad (1)$$

To the NCO, f is the signal-frequency-number input parameter, and t is the time reference provided by the clock. The waveform map provides the sinusoidal cosine waveform and the digital-to-analog conversion sets the output-signal amplitude, A.

The argument of the cosine is the signal phase, which, for a fixed output frequency, must be a linear ramp. The digital accumulator generates this ramp. At every cycle of the clock, a phase increment corresponding to the desired output frequency is added to the existing phase value. At a particular output frequency, this increment will be fixed, and its repeated accumulation results in the desired ramp.

The accumulator is not a counter. The step size of a counter is fixed, usually at unity. For the digital accumulator, the step size corresponds to the signal-frequency number, f. This distinction will become important shortly in deriving the DDS tuning relationship.

Waveform map

If an NCO's output waveform is fixed as a cosine, a fixed ROM can implement the waveform map. Addresses in the ROM will represent the phase position within the output-signal cycle, and the data stored at each address will be the corresponding cosine amplitude.

Note that the waveform mapping is general. Maps can be made for disk read/write-head waveforms, nonlinear-phase signals, and even noise waveforms. A particularly special case is the operation of two waveform maps in parallel, one with a cosine and the other with a sine waveform. This technique provides absolute quadrature signals, which are required by many signal-processing and DSP applications.

The waveform map must operate at the full clock speed. Because the map follows the digital accumulator, each clock pulse provides different information to the map, and the map must settle completely within each clock period. If it doesn't do so, it will incorrectly

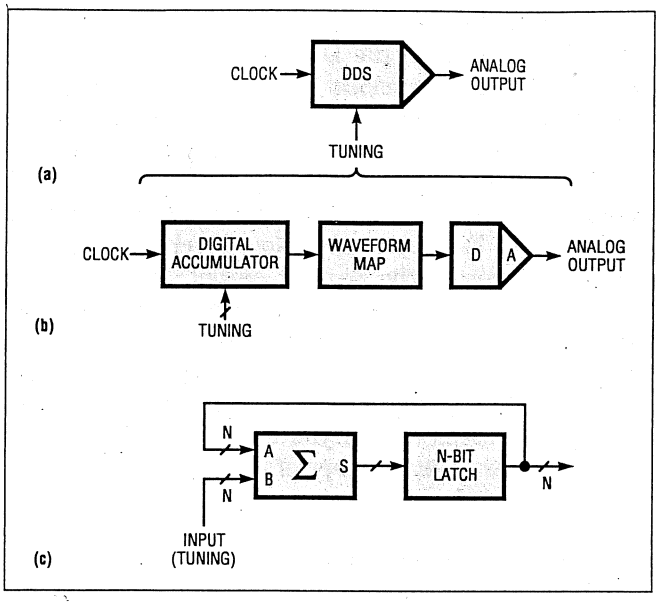

Fig 2—You can think of a DDS *as a single block* ***(a)*** *with digital and tuning inputs and an analog output. The more complex representation* ***(b)*** *more closely approximates the real nature of the DDS, and the 2-block representation of the digital accumulator in* ***c*** *suggests that the accumulator is more than a simple counter.*

convert the phase information from the digital accumulator to the corresponding amplitude value. For example, a 10-MHz clock requires the waveform map to have a cycle time of less than 100 nsec.

The analog-conversion block takes the amplitude number sequence from the waveform map and converts it into a single analog signal. Today, a single-chip DAC usually performs this operation. Because the amplitude number sequence represents the actual, real-time samples that an accurate ADC would have generated from the desired signal, had the signal existed, the DAC output signal is the (re)constructed waveform of the desired signal.

In general, DAC devices are not designed for use in DDS (**Ref 1**). High-quality DDS output signals demand that the DAC not only have good static linearity, but also that its dynamic characteristics (slewing and settling) be well matched and controlled. A common fix for a DAC with unsatisfactory dynamic characteristics is to follow it with a sample-and-hold (S/H) circuit. Doing so replaces the DAC's dynamic characteristics with those of the S/H circuit, which are generally better. As the DAC manufacturers improve their products to meet the needs of high-quality DDS outputs, the use of S/H circuits in DDS will diminish.

Discussions of DACs in DDS applications generally

DDS takes a small set of parameters (numbers) that describe the desired signal and generates a number sequence that represents the signal.

assume that the DAC settling time is less than the clock period. In fact, the DAC settling time should be *much* less than the clock period (**Ref 1**). Very fast settling produces output steps that are more nearly square, and more closely approximate the perfect rectangles assumed by the sampling theorem. This DAC requirement leads to a DDS rule of thumb: With a given set of hardware, the slowest clock frequency that can generate the desired output frequency will provide the lowest level of spurious outputs. The cleaner output is a direct result of the more rectangular shape of the output steps. In other words, you'll get better results if you use a slow clock to generate fewer high-quality steps than if you use a higher frequency clock to generate many steps of lower quality. In DDS, more is generally worse, not better.

Incorporating modulation

For communications purposes, pure sine waves are essentially useless. To pass information along, you must modulate the sine wave in some way. A sine wave has three characteristics capable of modulation: amplitude, frequency, and phase. Including them in the general signal, **Eq 1** gives the general communications signal:

$$s(t) = A(t)\cos(2\pi \cdot (f + f_m(t))t + p(t)).$$

A(t) represents amplitude modulation (AM), $f_m(t)$ represents frequency modulation (FM), and p(t) represents phase modulation (PM). A DDS device that includes modulation capabilities is called a modulated NCO, an NCMO for number-controlled modulated oscillator, or an MNCO for modulated, number-controlled oscillator.

Interest in using continuous-phase signals to conserve output bandwidth is growing. A signal that does not have phase continuity will be discontinuous at its first derivative. Rapid changes in a waveform produce high-frequency sidebands, and a signal whose first derivative is discontinuous can have high-frequency sidebands that contain significant energy. The more of the signal's high-order derivatives that are continuous, the smaller the waveform's high-frequency content will be.

DDS is inherently a continuous-phase technique; its output-waveform calculation always proceeds from the present point, whether or not any parameter changes. Therefore, DDS completely eliminates switching transients, overshoot, and undershoot. By definition, phase modulation can produce phase discontinuities, but a modulated NCO can only approximate them. The modulated NCO approximates phase discontinuities by performing a large number of small phase steps in quick succession. Ideally, one of these steps should occur in every DDS clock cycle. You can purchase devices and hardware that phase-modulate at this rate (**Refs 2, 3,** and **4**).

Developing DDS designs

Frequency resolution is one of the primary issues of any synthesized-signal design. For DDS, you find the output frequency f_o from the relationship

$$f_o = (f_{clk}/K) \cdot M \text{ Hz}, \quad (2)$$

where f_{clk} is the applied clock frequency, in Hz, M is the tuning number applied to the DDS, and K is the operating modulus of the DDS digital accumulator.

The clock frequency sets the DDS's sampling rate, which is the rate at which the DDS updates the signal-amplitude samples. In almost all cases, the sampling rate is equal to the frequency of the applied clock. The design of the DDS device determines the operating modulus, K, which equals the number of states that the accumulator can take on. Since most devices use binary circuits, K is usually a power of two such as 2^{24} or 2^{32}. When DDS devices use decimal circuits, K is a power of ten such as 10^6 or 10^8. A new technique called variable resolution (VR) lets you set K to any number between 1 and the digital accumulator's maximum intrinsic number of states.

The tuning number, M, is an integer between zero and K/2. The upper bound is called the Nyquist limit, a requirement from the sampling theorem to guarantee a unique output frequency. When M = 0, **Eq 2** shows that the output frequency will also go to zero. Therefore DDS designs include dc within their tuning bandwidth. The frequency resolution of the DDS is the derivative of **Eq 2** with respect to M. This calculation gives the DDS resolution (f_{res}) as

$$f_{res} = f_{clk}/K. \quad (3)$$

The frequency resolution is identical to the output frequency when M = 1. Because M must be an integer, all output frequencies will be harmonics of the resolution given in **Eq 3**. For this reason, this resolution is occasionally called the DDS quantization frequency.

DDS frequency resolutions can be very small. As an example, consider a 24-bit binary DDS device oper-

ating with a 10-MHz clock. The 24-bit accumulator sets K to 16,777,216 and yields a frequency resolution of 10,000,000/16,777,216 = 0.59 Hz. Ease of achieving fine frequency resolution is a fundamental characteristic of the DDS. More bits give even finer steps.

Several applications require an exact frequency resolution, and have a single, high-precision reference frequency for the DDS clock. For these designs, a simple algebraic shift of **Eq 3** would be useful:

$$K = f_{clk}/f_{res}. \quad (4)$$

Such an approach would set the DDS modulus and the desired resolution exactly. The Variable Resolution (VR) technique accomplishes these objectives. If a design requires precise resolution, such as 2.85714 Hz or 0.100000 Hz, it must either supply a special clock frequency to the DDS according to

$$f_{clk} = K \cdot f_{res},$$

or use VR technology, according to **Eq 4**.

Synthesizer output bandwidth

As mentioned before, there is a maximum output frequency at which the tuning relationship of **Eq 2** holds. This frequency (the Nyquist frequency) is equal to one half of the applied clock frequency. As with most real designs, the practical upper limit is lower than the theoretical limit. For DDS, the practical upper limit lies around 40 to 45% of the clock frequency. This limitation holds not only for the tuned output frequency, but also includes the sum of any and all modulation sidebands above the carrier.

Output settling-time performance

The first D in DDS stands for direct, which means DDS designs do not use feedback to ensure output-frequency accuracy. This approach differs from the PLL approach, which depends on feedback to achieve output-frequency accuracy. The PLL must stabilize its feedback loop to effect any frequency change. Hence, changing the frequency of a PLL takes longer than the reciprocal of the PLL's bandwidth.

As soon as you change the applied-frequency number, M, a DDS starts using the new M value in its calculations. The DDS' pipeline depth (number of calculation stages) establishes the time required for the output signal to reflect this frequency change. If, for example, a DDS has a 10-MHz clock and 32 stages, it will require 3.2 microseconds. More efficient designs, using fewer stages, switch proportionally faster: With the same clock, a 6-stage DDS will settle in 600 nanoseconds. Hardware is readily available with 5 stages of registers from the applied tuning number to the DAC output. With a 10-MHz clock, such a design exhibits the same frequency agility (settling time) as a 32-stage design operating at 64 MHz.

Dealing with output alias signals

Because it is a sampled system, a DDS has a multiple-signal output spectrum. In addition to the desired output signal at f_o, there are output signals at each harmonic of the clock plus f_o. The output spectrum is therefore f_o, $f_{clk}+f_o$, $2\cdot f_{clk}+f_o$, $3\cdot f_{clk}+f_o$, etc. The additional signals, often referred to as alias signals, are mixing products of the output signal with the clock and all of its harmonics. These signals must be removed by lowpass filtering to leave only the desired f_o signal. This filtering is the reverse of the antialias filtering used ahead of ADCs.

To make the output filter realizable, the upper output frequency is nominally 40% of the clock frequency. As the upper output frequency approaches the Nyquist frequency from below, the alias frequency at $f_{clk}-f_o$ approaches the Nyquist frequency from above. The closer the upper output and alias frequencies get to each other, the more complex the output lowpass filter must become. If you try to make the output frequency exceed 40% of the clock frequency, the output filter quickly becomes impractical. For example, achieving 60 dB of alias rejection when the output frequency is 42% of the clock requires a Chebyshev lowpass filter with more than 20 sections. Realizing any analog filter of seven sections or more requires special care, but producing a filter with 20 sections is almost impossible.

At the output of the lowpass filter, analog signal (re)construction is complete. The signal can then be used directly, or additionally processed by such conventional analog techniques as amplification, mixing, limiting, and multiplication.

Determining output-signal quality

The cosine is a transcendental function, which the waveform map cannot quantize precisely. The quantization errors that inevitably exist produce signals in addition to the main signal and the alias terms. These nonharmonic extra signals are called DDS spurs. In a good DDS design, these spurs will be relatively low in amplitude—at least 60 dB below the main output.

There are three main blocks: the digital accumulator, a waveform map, and the digital-to-analog conversion.

In the best DDS designs, the spurs are 70 to 80 dB below the main output. If required, under some conditions, you can reduce the spurs even further.

When you tune the main signal, the DDS spurs generally move around much faster than the main signal does. You can predict the location of the spurs by modifying a technique used in analog systems with mixers (**Refs 5** and **6**). Near output frequencies that are integral submultiples of the clock (⅓, ¼, 1/10, etc.), the spurs all "gather around" the main output. At exact integral submultiples of the clock, the spurs all cross over the main output. You can quickly check a DDS design's degree of spur generation by tuning the DDS to near the ⅓- and ¼-clock-frequency crossovers. You can then directly evaluate the output spectrum quality from plots similar to those of **Fig 3**.

Construction hints for success

A DDS, by nature, uses digital circuits to perform analog functions. The absence of noise generation is essential for a successful design. High-frequency construction techniques are necessary. Clock frequencies are typically 10 to 50 MHz. On a standard pc board made from G-10 or FR-4 laminates, the shortest wavelength is about 270 cm. Transmission-line effects are not important here, because a typical DDS design covers 10 to 20 cm—under 1/10 of a wavelength.

Loop currents are another matter. The DDS's digital circuits have edge speeds of less than 10 nsec. The currents these circuits produce have significant energy at 500 MHz or more. The signal current must flow from the driving IC's dc supply pin, through its output transistor(s), to the receiving IC's input pin, and back to the supply through the ground. **Fig 4** shows this current path.

Any signal transfer from one IC to another has three loop currents. The driving IC draws current from its supply pin to turn on an output-pin driver. Some of this current flows into the interconnection as signal current. The rest flows out the driving IC's ground pin to return to the supply.

The signal current flows to the receiving IC and draws a matching current from the IC as required to conserve charge. (Basic physics strikes again!) This return signal current flows from the receiving IC's ground pin, underneath the signal trace (if there is a ground plane there), and back to the driving IC to return to the supply.

A third loop current results from the effects of the signal on the receiving IC. The receiving IC will draw current from its own supply pin in response to the stimulus. This current will flow out of the IC's ground pin and will return to the power supply.

Good low-noise construction will guarantee that all of these currents flow in the smallest possible area. Radiation and other interference increases in direct proportion to the area enclosed by the conductors carrying these currents. Also, if more than one current flows in a conductor, the currents can interact. Careful layout of the pc board is essential to minimize signal-loop sizes and impedances common to several loops.

Current loops around each IC are a slightly different

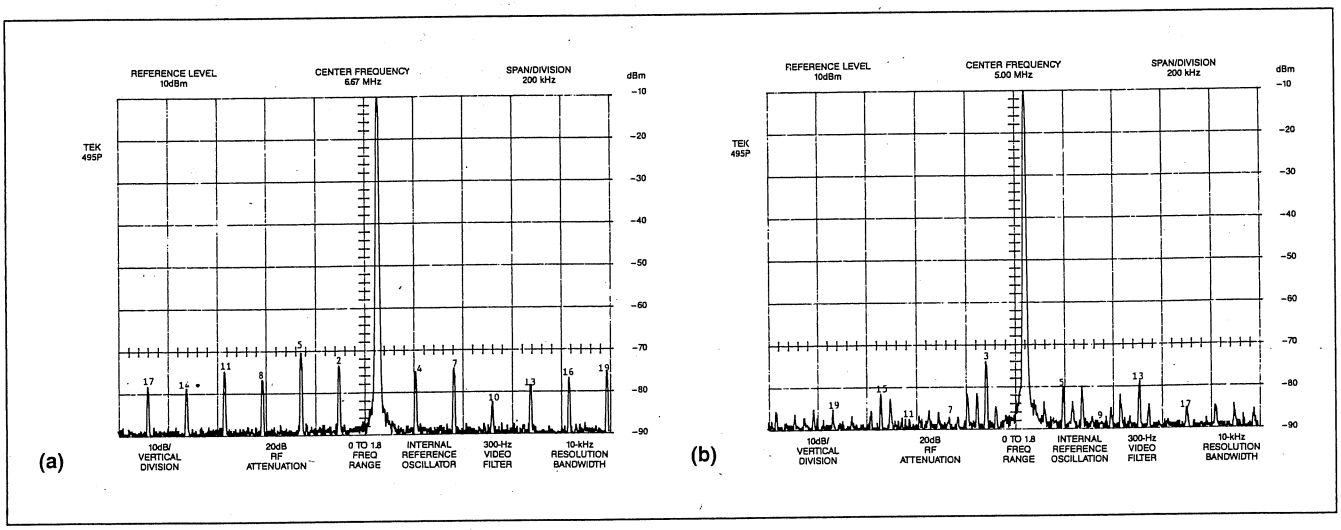

***Fig 3—You can quickly check a DDS** design's degree of spur generation by tuning the DDS to near the ⅓- **(a)** and ¼-clock-frequency **(b)** crossovers. Spectral plots such as these then readily reveal the spurs.*

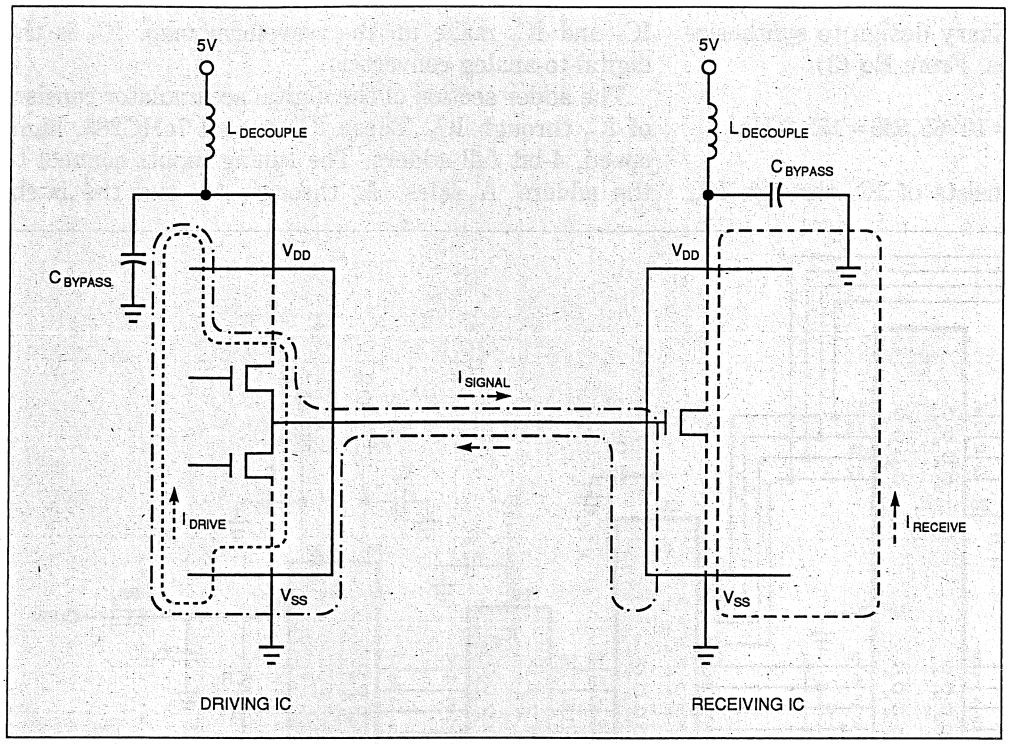

***Fig 4—When one IC sends a signal** to another, the currents that flow create electromagnetic fields. Understanding the current paths in signal leads and ground planes helps you to minimize the spurious signals generated by the changing fields.*

matter. The bias component of the switching currents should never flow far from the IC—this is the idea behind the use of bypass capacitors. If the impedance of the bypass capacitor is lower than that of the power supply as seen by the IC power pin, the switching current will come from and return to the bypass capacitor (hence the name). Between switching transients, the bypass capacitor will recharge from the higher impedance supply. If there is no bypass capacitor, the IC will be forced to draw power from the power supply during the transient. Because power supplies are rarely right next to their loads, the loop-current path will enclose a large area. Radiation and interference will be severe.

Multilayer PC boards and surface-mounted components are a significant help in controlling these currents. A ground-plane layer immediately below the signal traces minimizes the signal-current path length. Surface-mounted bypass capacitors have lower lead inductances than do through-hole-mounted capacitors. Mounting these devices on the back of a pc board further reduces the length of the current path. Nevertheless, these modern components and construction techniques are mixed blessings. Besides costing more, they can cause problems with power distribution.

A bypass capacitor will work when its impedance, as seen by the IC's power pin, is lower than that of the power supply. By design, a power plane, like a ground plane, has a very low impedance. For the bypassing to control the transient loop current, there must be some impedance between the point where the IC draws current from the power plane and the junction of the IC's power pin and the bypass capacitor. This impedance can take the form of a ferrite bead or, if the dc current is low enough, a smaller chip inductor. A 470-nH inductance is a good value for this application. Without the series inductance, the bypassing is ineffective; the ICs draw energy directly from the low-impedance power and ground planes in a manner that makes control essentially impossible.

Fig 5 shows a simple DDS design using standard components. This synthesizer, which is useful as a general-purpose signal generator covering the audio and low-RF ranges, has the following specifications:

Output Frequency Range	0 to 4 MHz
Frequency Resolution	152.6 Hz
Spurious Signal Suppression	50 dB below carrier
Frequency Switching Time	<500 nsec

A 10-MHz crystal-reference frequency drives this sample design.

With a given set of hardware, the slowest clock frequency that can generate the desired output frequency will provide the lowest level of spurious outputs.

The DDS uses a 16-bit binary design to synthesize the required frequency steps. From **Eq (3)**

$$f_{res} = f_{clk}/K = 10\ MHz/2^{16} = 10^7/65{,}536 = 152.6\ Hz.$$

The digital accumulator consists of IC_1 through IC_6. IC_7 and IC_8 make up the waveform map. IC_9 is the digital-to-analog converter.

The adder section of the digital accumulator consists of IC_1 through IC_4. These ICs are all 74HC283, high-speed, 4-bit full adders. The tuning inputs connect to the adders' A sides, A_0 through A_3, and the latch-

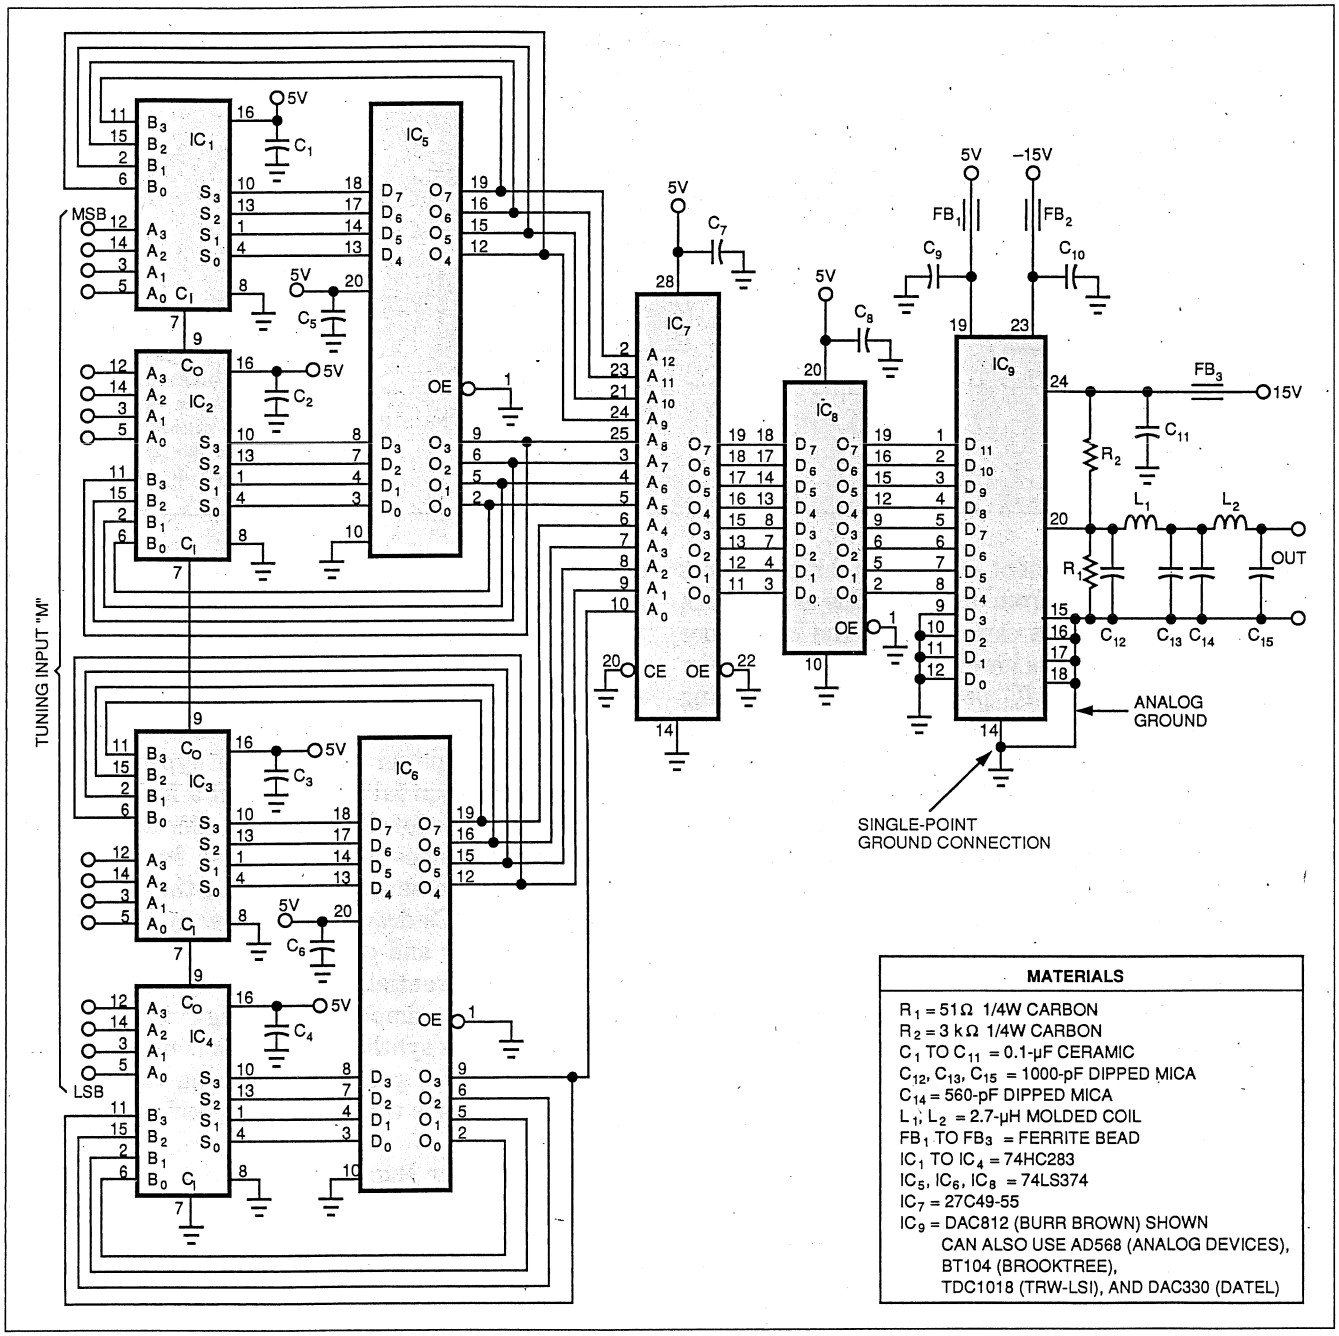

***Fig 5—This 0- to 4-MHz DDS** uses nine readily available ICs. All of the other components are passive.*

DDS designs do not use feedback to ensure output-frequency accuracy. A PLL depends on feedback to achieve output-frequency accuracy.

output feedback connects to their B sides, B_0 through B_3. This arrangement is arbitrary, but it is easy to remember. To configure the four devices as a single 16-bit adder, the carry output of one stage connects to the carry input of the next-higher stage. Latches IC_5 and IC_6 complete the digital accumulator. The adder outputs drive the latch inputs: IC_1 and IC_2 drive IC_5; IC_3 and IC_4 drive IC_6.

The waveform map consists of PROM IC_7 and latch IC_8. The PROM contains 8k (2^{13}) bytes, so there are 13 address bits. The PROM contains a full sine wave calculated from

$$\text{data} = 127 \cdot (\sin(2\pi \cdot \text{address}/8192) + 1),$$

where 0<address<8191, and the argument of the sine function is in radians. Latch IC_8 retimes the data output. You can find the required PROM-device cycle time from

$$\begin{aligned}\text{clock\_period} = {} & \text{propagation\_delay } (IC_5, IC_6) \ldots t_{pd} \\ & + \text{PROM\_cycle\_time } (IC_7) \ldots t_{cy} \\ & + \text{latch\_setup\_time } (IC_8) \ldots t_{su},\end{aligned}$$

so that, using LS series (74LS374) latch devices for IC_5, IC_6, and IC_8,

$$\begin{aligned}t_{cy} &< T_C - t_{pd} - t_{su} \\ &< 100 - 25 - 20 \text{ nsec} \\ &< 55 \text{ nsec}.\end{aligned}$$

Faster latch devices will allow correspondingly longer PROM cycle times. For this example, a PROM device of type 27C49 through 27C55, or a similar part, permits 10-MHz operation.

IC_9 performs digital-to-analog conversion. Several devices meet the requirements of this function; the Burr-Brown DAC812 is shown. This device settles to 12-bit accuracy in less than 50 nsec, about half of the clock period. R_1 converts the 0- to 10-mA DAC output current to a voltage and also sets the output impedance of the DAC. A resistance of 50Ω matches the conventional line impedance of RF circuitry. Resistor R_2 injects a 5-mA current into the DAC output node to center the DAC output range around zero.

The 5-section LC, lowpass filter removes the output-signal alias terms. The component values produce a Chebyshev filter with 0.1 dB of passband ripple and a cutoff frequency of 4 MHz in a 50Ω system.

Thus, designing the filter completes the synthesizer's design. The output signal power is nominally 0.5 mW, which is 5 dB less than 1 mW (−5 dBm). The frequency-switching time is two clock periods, that is, 200 nsec. The measured spurious-signal suppression is −52 dBc (52 dB less than carrier level).

Square-wave time-base generator

Many applications do not require sine waves. In timing applications, square waves are all you need. Because DDS is a digital technique, you would expect that a square-wave output would be natural. Alas, the DDS does not produce a square wave without some help from its designer. The reasoning, which is covered more thoroughly elsewhere (**Ref 7**), is only summarized here.

The problem stems from the fact that a DDS design is a synchronous digital system. Its state will change, along with any output, only in response to a clock edge. The resolution of any output cycle is therefore limited by the clock period: the duration of any individual output cycle must be an integer number of clock periods.

Changing the output frequency by 1 Hz involves a change in the output period far smaller than the clock period of any current DDS. Any variation from the required period is modulation, which creates sidebands and spurs. Because the time quantization is too coarse to achieve the necessary signal quality, interpolation is necessary. The DDS output lowpass filter (LPF) performs the interpolation. Because the LPF is built from fixed components, it qualifies as an LTI (linear, time-invariant) network—a concept from early circuit-analysis classes. LTI networks have a "natural frequency" which is sinusoidal; you obtain the highest quality output signals by driving the filter with sinusoidal signals. This purpose underlies the cosine waveform map and the DAC. To effectively filter the phase information, you must make the LPF as "happy" as possible. The LPF will then perform the required time interpolation and yield a clean signal. But it is a sine wave. To make it square, you must employ amplitude limiting, usually with a comparator of some sort.

Fig 6 shows a wideband clock generator built with DDS. Note that this design creates a square wave by limiting a sine wave's amplitude. From a 50-MHz clock, this generator will provide a direct square-wave output from dc to the limit of the LPF, probably around 20 MHz.

At low output frequencies, the jitter on the DDS square-wave output is less objectionable than it is at

If you try to make the output frequency exceed 40% of the clock frequency, the output filter quickly becomes impractical.

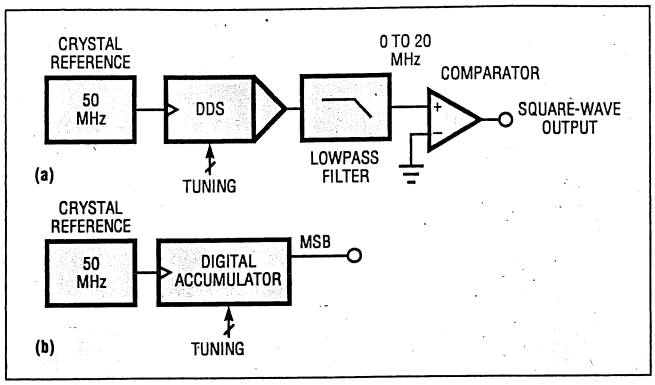

Fig 6—To obtain a high-quality *square wave over a broad range of frequencies, you must have the DDS generate a sine wave. Then you must filter out the harmonics and limit the sine wave's amplitude, as in* ***a****. But if you are interested only in low-frequency square waves and you can tolerate moderate jitter, the DDS's digital accumulator alone may suffice,* ***(b)****.*

high frequencies. If you need only a low-frequency output, you need not generate a sine wave and limit its amplitude to produce a square wave. **Fig 6b** shows an implementation that dispenses with sine-wave generation and amplitude limiting. With a strict limitation on its bandwidth, the direct output exhibits tolerable jitter, according to:

$$f_{max}, \text{ direct square wave} = f_{clk} \cdot \%\_\text{jitter}/2.$$

This example uses a 50-MHz clock. Square waves with 0.1% jitter are directly obtainable from dc to 25 kHz. If you increase the jitter tolerance to 1%, the output bandwidth can increase to 250 kHz. Though it is not very efficient as square-wave generators go, this DDS design would be small and exceedingly stable.

Arbitrary output waveforms are possible, as mentioned earlier, by changing the information in the waveform map. **Fig 7** shows two common ways to change this information; real-time map changes are possible. Through the use of a dual-port RAM, a computer can insert changes to the waveform as the DDS operates. A generator with such an architecture can produce high-quality speech and is suited to any application that requires a large number of waveforms.

If the number of different waveform types is relatively small, a single PROM can hold them all. An external processor can select the desired waveform with the upper address bits. A typical application of this type is in testing of disk-drive read circuits. The PROM holds proper waveforms for positive and negative flux transitions as well as several types of problem waveforms for each transition direction. The processor can supply the circuit under test with normal waveforms, occasionally insert a particular error, and then return to normal operation. Digital buses and communications links are testable in a similar way.

Direct Digital Synthesis is a real technology, well beyond the experimental stage. DDS devices and subassemblies are available today from several manufacturers. Because DDS does not rely on feedback, it can simultaneously realize small frequency steps, fast frequency switching, and low phase noise. Careful construction techniques have solved the spurious-signal problem that has traditionally limited DDS applications in communications systems. Hence, the doors are open

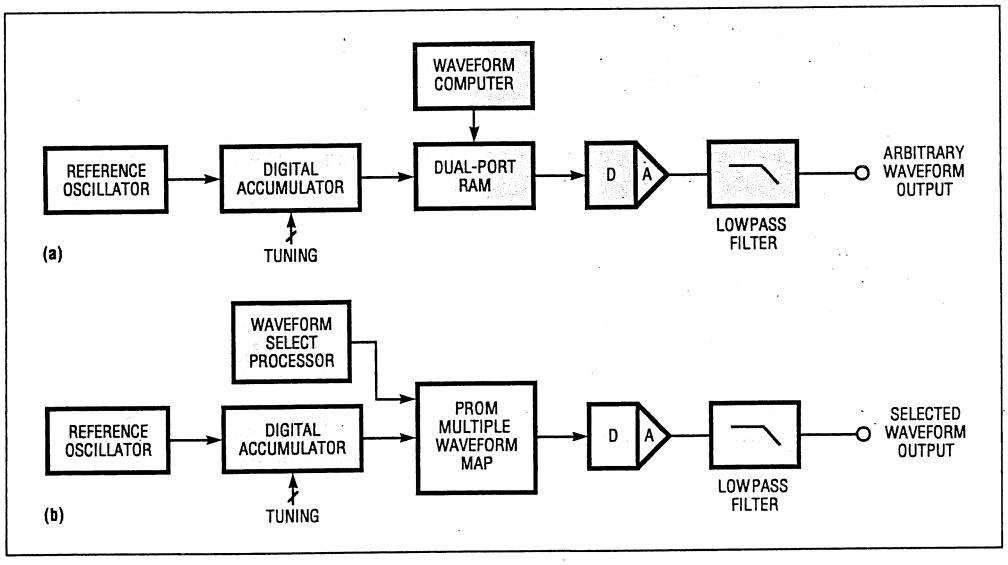

Fig 7—Obtaining arbitrary waveforms *from a DDS can involve using a dual-ported memory and feeding it new waveform maps on the fly,* ***(a)****. For less demanding applications, a ROM can store several waveform maps,* ***(b)****, and the MSBs of the ROM address can determine which waveform the DDS produces.*

to system performance that was impossible not long ago.

Dedicated hardware, both commercial and military, is in production. Standard products are used primarily in low- to medium-volume applications. Where quantities are larger, custom products can offer attractive competitive advantages. DDS is moving into the engineer's toolbox to improve product performance while keeping pricing within reason. Through DDS, the digital world's economy is moving into the RF and analog worlds. EDN

References

1. McCune, E, "Control of spurious signals in DDS," *Digital RF Solutions Corp,* Santa Clara, CA, AN1004.
2. Stokes, P, "Basic 'analog' NCMO applications," *Digital RF Solutions Corp,* AN1003.
3. McCune, E, "Data modulation using direct digital synthesis," *Proceedings of the RF technology expo,* 1990, pg 413.
4. McCune, E, "Data communications and numeric modulation," *Digital RF Solutions Corp,* AN4001.
5. McCune, E, "Quantitative prediction of DDS spurs," *Digital RF Solutions Corp,* AN1012.
6. Digital RF Solutions Corp, *"Spur predictor disk,"* (spreadsheet worksheets for Lotus 1-2-3 and Framework II/III).
7. McCune, E, "Numeric synthesis of square waves," *Digital RF Solutions Corp,* AN1008.

Author's biography

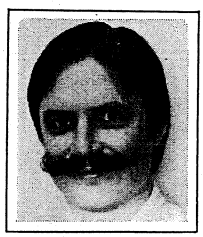

Earl McCune Jr is VP of Engineering at Digital RF Solutions Corp. He has been with the firm for nearly five years and has developed module- and board-level synthesizers using DDS and other techniques. He holds a BSEE from the University of California at Berkeley and an MSEE from Stanford University (Stanford, CA). Earl is a member of IEEE. His leisure activities include hiking, bicycling, and working with model aircraft.

Article Interest Quotient (Circle One)
High 497 Medium 498 Low 499