Telemetry Simulation Using Direct Digital Synthesis Techniques

Randall S. Pitchford
Project Manager, Telemetry Systems
Frontier Engineering, Inc.
P.O. Box 5755
Vandenberg AFB, CA 93437

Introduction

Direct digital synthesis technology has been employed in the development of a telemetry data simulator constructed for the Western Space and Missile Center (WSMC). The telemetry simulator, known as TDVS II, is briefly described to provide background; however, the principal subject is related to the development of programmable synthesizer modules employed in the TDVS II system. The programmable synthesizer modules (or PSMs) utilize direct digital synthesizer (DDS) technology to generate a variety of common telemetry signals for simulation output.

The internal behavior of DDS devices has been thoroughly examined in the literature for nearly 20 years. The author is aware of significant work in this area by every major aerospace contractor, as well as a broad range of activity by semiconductor developers, and in the universities. The purpose here is to expand awareness of the subject and its basic concepts in support of applications for the telemetry industry.

During the TDVS II application development period, new DDS devices have appeared and several advances in device technology (in terms of both speed and technique) have been effected. Many fundamental communications technologies will move into greater capacity and offer new capabilities over the next few years as a direct result of DDS technology. Among these are: cellular telephony, high-definition television and video delivery systems in general, data communications down to the general business facsimile and home modem level, and other communications systems of various types to include telemetry systems.

A recent literature search of the topic, limited only to documents available in English, indicates that some 25 articles and dissertations of significance have appeared since 1985, with over 30% of these appearing in international forums (including Germany, Japan, Great Britain, Portugal, Finland...). Product advertisements can readily be found in various publications on test instruments, amateur radio, etc., which indicate that international knowledge and product application of the technology is becoming increasingly widespread.

Brief Overview of TDVS II

One area of the TDVS II system, that of producing (on a single module) simulated telemetry signals representative of all the basic forms in IRIG 106-86 (low-megabit rate baseband PCM, as well as IRIG FM/FM, PAM/FM, and PCM/FM), was largely made possible as a direct result of the availability of low-cost direct digital synthesis components. The following brief overview of TDVS II will clarify the general implementation of the system, and the strategic application of DDS devices.

The TDVS II computer ensemble consists of two back-to-back 68020-based VME systems, where the front-end of one system consists of PC/AT compatible workstations and the back-end of the other system consists of two VXI crates known as Telemetry Synthesis Units (TSUs). Users of TDVS II may define a desired telemetry simulation using fill-in-the-blanks data entry screens on the workstations. Based on the user's definition, a bit-wise representation of the desired telemetry data stream(s) is constructed and routed to the second VME frame. The second VME frame is also capable of user interaction through an Ethernet interface to the front-end frame, such that user requested real-time changes to a simulation may be effected while the simulated telemetry data is being "streamed" to the TSUs.

Each TSU contains twelve (12) Programmable Synthesizer Modules (PSMs). Each PSM can generate either one channel of baseband PCM at data rates in the lower megabits per second ranges, or one IRIG FM subcarrier (from the basic IRIG set of 85 PBW and CBW channels) modulated with PCM, PAM, or basic waveform (e.g., sine, ramp, DAC-ID pattern, etc.) data. Altogether, TDVS II can generate up to 24 channels of synthesized telemetry (24 effective information channels). However, the two TSU subsystems together generate four unique telemetry outputs, up to three of which may be baseband PCM, and one of which is an FDM composite containing up to 24 subcarriers (if none of the PCMs are active, and depending on normal channel spectra constraints).

The FM subcarriers may be modulated such that up to three of these channels may be PCM/FM; up to four may be PAM/FM; and the remainder are modulated with simple waveforms, multi-point calibration patterns, DAC-ID sequences, etc. In addition to the three baseband PCM and one FDM output, the system also has a fifth non-unique telemetry output which may be configured to produce "special" telemetry outputs from combinations of the first four, such as PCM/FDM hybrid signals.

As an overall system, TDVS II is configured as a three-link telemetry simulator. The telemetry structure capabilities and data generation capacity of TDVS II, while not specifically the subject of this paper, tend to exceed the capabilities of most other telemetry simulators presented at this forum in recent years, and so are briefly reviewed in Table 1.

Some aspects of the general telemetry simulation issue which relate more to testing of range hardware through the receivers, bit syncs, etc., are not directly addressed by TDVS II, which is specifically oriented toward generation of substantial data structure and content. While some basic Doppler simulation capabilities exist in an associated RF transmission system (Simulated Telemetry Open Loop System II - STOLS II), the overall ensemble does not presently provide for more exotic testing techniques such as noise and/or jitter injection, PCM vanishing, random wrong bits, induced ringing, etc. TDVS II is primarily intended as a high-volume, high-accuracy data simulator: the simulation of various characteristics of real-world RF signals encountered in aerospace vehicle communications is best handled by off-the-shelf signal simulators or other techniques which can be introduced at the receivers.

The overall TDVS II ensemble also provides for distribution of its outputs to various range facilities, recording and/or playback of simulations, local quick-look analysis of simulations, or transmission of up to three S and/or L band RF signals (STOLS II) modulated with the TDVS II simulated telemetry. All TDVS II simulated IRIG telemetry waveforms are ultimately generated on the PSM modules using direct digital synthesis techniques in one way or another.

Basic Techniques of Direct Digital Synthesis (DDS)

Since the earliest days of radio, designers have employed various techniques to generate stable signals of controllable frequency. Until recent years, all techniques for synthesizing waveforms were implemented in the frequency domain; including RC and LC oscillators, crystal oscillators, PLL circuits related to crystal, RC, or LC circuits, etc. In the last 20 years, the development of high speed digital integrated circuits has made it possible to synthesize waveforms directly in the time domain (i.e., as the time-phased points, or samples, of a desired final waveform). The essential rules for the number of samples required for successful direct digital synthesis of waveforms were actually developed about 50 years ago at Bell Laboratories when Harry Nyquist defined the basic relationship between the frequency domain and samples in the time domain.

Some writers in the field incorrectly state that a basic direct digital synthesizer can be readily constructed from a counter, a ROM or EPROM, and a DAC assembled in a straight chain. This type of device is simply a clocked waveform map, and not a DDS.

The fundamental feature of a true DDS is the use of a phase accumulator; typically consisting of an input register (actually a delta phase, or “hop” register), a digital adder with bit-width corresponding to that of the registers, and a register which serves as an accumulator. The accumulator value at any point in time is used to address a look-up table (or waveform map, usually implemented in ROM) which outputs the sine (or other waveshape function value) of the phase angle represented by the current number in the accumulator. In most present day implementations, the waveform map is followed by a DAC (and a low-pass filter). A frequency is synthesized by virtue of a feedback path from the accumulator to the adder, such that the current delta phase value (hop size) is added to each subsequent accumulator output value to produce a new value. Eventually the accumulator reaches overflow, at which point a full cycle of possible phase values has been exhausted and another sequence begins. The frequency output by such a circuit is a function of the clock frequency, the length of the accumulator, and the delta phase increment. This arrangement is illustrated in Figure 1.

Another popular view of this operation (also shown in Figure 1) is based on the idea that a repetitive waveform can be visualized as rotation from point to point around a circle. The circumference of the circle (sometimes called a phase circle) is equal to the number of values which can be represented in the phase accumulator (e.g., a 24-bit wide system can represent 2^{24}, or 16,777,216 discrete phase values). Rotation around the phase circle takes place in discrete increments at a constant clock rate, such that the only variable is the size of the increment (the delta phase value, or the number of points to skip when progressing to the next phase state on the circle). Phase accumulator overflow is the equivalent of starting another trip around the circle.

The delta phase value (ΔN) is the change in accumulated phase over the time interval between clock pulses, and the rate of change of phase represents the output frequency of the circuit. As shown in the relationship (1), below, this output frequency is dependent on three parameters: f_{clk}, the frequency of the timebase clock in Hertz; n, the length in bits of the of the phase accumulator; and ΔN, the phase increment, which is effectively the tuning input value.

(1) $$dN/dt = (\Delta N/2^n) / (1/f_{clk}) = \Delta N(f_{clk}/2^n) = f_{out}$$

Thus, tuning is accomplished by changing the size of the hops (phase deltas) taken by the phase accumulator while being clocked around the phase circle. Since the delta phase (tuning) register is a separate structure from the actual accumulator portion of the circuit, the result of a change in tuning value is simply that the overall phase accumulator action continues from its current phase value at the new increment. Frequency changes therefore occur directly and in a phase continuous fashion. However, since the values output from the overall phase accumulator are used to address entries in a waveform map, the result is fundamentally a sampled data system, subject to Nyquist's theorem. Thus, the maximum output frequency which can practically be generated is something less than half of the clock frequency.

Since the DDS is a sampled data system, an obvious source of errors is the finite quantization of a sine wave as amplitude values in the waveform map. Predictably, the final output sinewave will exhibit distortion; however, the distortion of the DDS is generally not harmonic relative to the sine wave itself, but rather the result of arithmetic shifts of the original spectrum which are spaced according to the sampling frequency. This sampling effect produces image responses above and below the clock frequency, in addition to the desired fundamental. The predictable frequencies for these images may be derived from:

(2) $f_{img} = (N*f_{clk})\pm f_{out}$ where N=0,1,2,...,.

These images must be filtered using a lowpass (or bandpass) filter at the DAC output.

Properties of the waveform map relative to the DAC are another predictable source of spurious energy in a DDS system. The basic purpose of the waveform map is to translate quantized phase values into quantized amplitude values for presentation to the DAC. Typically, a waveform map contains 180º of time-phase addressed amplitude information for a cosine function. The cosine is often used instead of the sine in order to minimize amplitude errors at zero crossings, and to provide for synchronization of delta phase changes with peaks in the carrier waveform (zero slope points).

In general, DDS implementations truncate the bit-width of the overall accumulator chain prior to input to the waveform map, which in turn outputs an amplitude resolution appropriate to the input bit-width of the DAC. The truncation of the resolution of the phase value introduces jitter in the output waveform which appears as spurs in the spectrum, and the amplitude quantization of the object waveform also produces spurs. Phase truncation errors in most DDS devices contribute spurs no greater than about -70dbc, but the behavior of typical DACs is likely to result in a higher figure for spurious content derived from amplitude related factors.

It is important to note that the intrinsic quantization errors of the DDS topology do not necessarily appear as uniformly distributed noise across the Nyquist bandwidth, as conventional wisdom might imply. In general, depending upon the specific properties of the DDS, its waveform maps, and D-to-A converter (i.e., bit-width of the phase accumulator vs bit-width of the maps vs resolution and other attributes of the DAC), quantization errors will repeat periodically at certain frequencies and produce additional spurs in the output. This occurs as a result of differential non-linearity of the DAC which becomes apparent when exact phase sequence values, and hence, exact amplitude values input to the DAC are repeated.

Current wisdom regarding DDS application is that the most significant source of spurious signal content is the result of DAC quantization errors. However, it would appear that essentially all currently available DDS/DAC combinations result in spurs no greater than -60dbc. In general, the so-called 6db-per-DAC-bit rule (quantization SNR = 10.8db + 6db-per-bit of DAC resolution) cannot be used as a prediction of DAC spur levels for reasons cited in the preceding paragraph; and DDS/DAC combinations do not always appear to behave in strict compliance with the expected 6db slope.

While many other characteristics of the DDS could be presented, the scope of this portion of the paper is limited to a basic understanding of the technique, its inherent problems and general oddities, and features relating to its application in the specific IRIG telemetry environment. Therefore, one more overview topic must be considered: that of digital modulation in such systems.

Modulation basically consists of varying in the time domain one or more of the three primary signal parameters: frequency, phase, or amplitude. In order to quickly understand how direct digital modulation of these signal parameters might be accomplished, consider the following general equation for a signal:

(3) $$s(t) = A f(Tt+N),$$

where:
- A is the signal amplitude
- T is the signal frequency in radians
- N is the signal base phase
- $f(\bullet)$ is the signal waveshape function

Modulating this general signal consists of introducing variations in time upon the basic signal parameters, which in the general case produces the equation:

(4) $$s(t) = A(t)f(T(t)t+N(t))$$

where:
$A(t)$ is the amplitude modulation
$T(t)$ is the frequency modulation
$N(t)$ is the phase modulation

By considering the basic DDS circuit of Figure 1 in the light of the above equation (4), one can readily deduce where each type of modulation would have to be effected in the system to produce the desired result. Frequency modulation would have to be ahead of the delta phase register, phase modulation would have to be after the adder but ahead of the (phase) accumulator register, and amplitude modulation would have to be after the (normalized) waveform map.

DDS Math is Discrete-Time

Generally, in modern communications, reality tends to be described in one of two ways: either as continuous-time (CT) functions (familiar to all of us as analog signals), or as discrete-time (DT) functions (familiar as sampled systems, or as related to sampling theory). In commonly encountered cases, a discrete-time signal is simply a sequence of values (like the sequence of values in a waveform map); but if the values of the sequence are represented by only a finite number of bits (as in a finite state machine), then the sequence can only represent a discrete set of values (commonly 0 and 1) and the sequence is called a digital signal. Much of the theory that is applied in DSP is actually the theory of DT signals and systems, in that no amplitude quantization is assumed in the mathematics.

The mathematics of CT systems and DT systems exhibit a general and fairly consistent duality of concepts, and other parallel features. For essentially every concept applied to the analysis of CT systems, a corresponding concept can be found in the literature for DT systems (e.g., continuous convolution and discrete convolution, or differential equations and difference equations). In the telemetry industry, we have to deal in both of these mathematical worlds, and are generally familiar with the fact that despite the apparent duality of concepts, it is simply not possible to apply directly the mathematics of CT systems to DT systems, or vice versa.

The systems we deal with in the telemetry world typically consist of both analog and digital subsystems, with appropriate ADC and DAC devices at their interfaces. In fact, it has become common practice for us to put a digital computer in the control loop for a multi-million-dollar missile or aircraft whose behavior and sensors are fundamentally analog by nature. As we are all well aware, analytical difficulties frequently occur at the boundaries between the analog and digital regions of our systems, because the mathematics used on the two sides of the interface must be different.

As a simple example, we can compare a typical digital sequence to its analog equivalent. Take the case of a DT signal sequence s(n), derived from an analog signal $s_a(t)$ by ideal sampling, such that:

(5) $$s(n) = s_a(t) \mid t=nT$$

The analog or CT model for the same sampled signal is denoted by $s^*(t)$ and defined by:

(6) $$s^*(t) = \sum_{n=-\infty}^{+\infty} s_a(t) W_a(t-nT)$$

where: $W_a(t)$ is an analog impulse function.

Of course, we recognize that both s(n) and $s^*(t)$ are used throughout the literature to represent an ideal sampled signal, and that even though s(n) and $s^*(t)$ represent the same essential information:

- s(n) is a DT signal, and
- $s^*(t)$ is a CT signal.

The demonstrated reality is that the two descriptions are not mathematically identical; and although this is a very fundamental point, s(n) remains a “DT-world” model of a sampled phenomenon, and $s^*(t)$ is simultaneously the “CT-world” model of the same phenomenon.

In a DDS, the phase accumulator is described by DT mathematics, and is a sampled data system. The DAC-filter pair is also a sampled data system. However, the output signal objective is still described in terms of CT mathematics. This results in the same generic analytical difficulties which occur whenever analog (CT) and digital sampled system (DT) phenomena meet.

DDS Output - Analog Signal Recovery

There are at least twenty important parameters involved with the specification of a digital to analog conversion scheme, and a complete discussion of this subject is beyond the scope and purpose of this paper. However, a brief review of some of the key phenomena encountered in waveform reconstruction is in order.

In general, DDS output is reconstructed through a zero-order hold type digital-to-analog conversion, where each output value is a function of its binary weight value (a number) and is held until the next sample arrives. As the binary numbers change, the output varies in discrete steps. If the binary numbers were generated by an up/down counter, a triangle step function response would be observed at the output, and spectral analysis would show that a large amount of undesirable high-frequency energy is present.

To eliminate this, the DAC is usually followed by a smoothing filter having a cutoff frequency no greater than half the sampling frequency. This filter produces a smoothed version of the converter output which is actually a convolved function. What this means mathematically is that the spectrum of the resulting signal is the product of two spectra: that of the step function (sinx/x), and that of the band-limiting analog filter.

Conventional wisdom would hold that an increase in the number of samples per cycle for a particular output frequency will provide a better approximation of the desired signal, and correspondingly lower distortion. However, this is not always the result with typical real DACs. With an ideal DAC, transition time from one sample to another is zero, and a very large number of samples makes sense. However, with a real DAC fewer samples per cycle can actually improve performance because the DAC-settling-time induced waveform error is averaged over more time, thus reducing its overall contribution to distortion of the output signal. Remember, only two points are theoretically required to reconstruct a complete sinewave with the proper DAC and filter combination. More is not always better.

The TDVS II PSMs use a low-pass filter at the DAC output which corners sharply at 2 MHz, thus limiting the spectrum to the IRIG FDM passband. Well known alternatives include following the DDS with a PLL, or applying a more specific filter at the DAC output.

The smoothed output of the DDS may exhibit aliasing effects resulting from the phase and attenuation relationships of the signal recovery process (the DAC and smoothing filter combination). Instead of immediately following the DDS with a

DAC and reconstruction filter, attenuation due to the DAC (sinx/x) spectrum shape could, theoretically, be compensated for by applying a digital filter with an inverse response (x/sinx) prior to conversion, providing an overall flat magnitude signal response to be smoothed by the filter. Other alternatives exist, and some alternative to the current DAC-filter arrangement is both required and likely to occur in future DDS developments.

Spectral Containment Considerations with DDS Technology

There are four basic categories of spurious spectral energy intrinsic to the general application of a DDS device. These areas are all discussed elsewhere in this paper, but the following is a basic list of the types of DDS-related spurious energy, and their sources:

- non-harmonic discrete spurs - DAC non-linearity
- harmonic spurs - DAC non-linearity
- aliasing responses - the images.
- broadband phase noise - the DDS clock source

However, for most applications, the DDS exhibits impressive spectral purity because it is inherently phase continuous. Thus, upon receiving a new operating frequency, the device continues on from its current phase position at the new frequency. This performance is equivalent to continuous phase frequency modulation (CPFM). This characteristic operating mode of the DDS results in the consumption of less bandwidth than that of a frequency-switched FM signal generated by a conventional VCO.

Also, in cases where the waveform map is structured such that signal peaks occur when the phase accumulator overflows and starts through the phase circle sequence again (e.g., as in cosine waveform mapping), then signal frequency shifts can be synchronized with peaks in the carrier waveform by using the carry-out flag from the DDS. This mode is referred to as continuous phase-slope frequency modulation (CPSFM), because at the signal peaks the phase-slope is zero. The end result of synchronization of frequency changes with phase accumulator overflows can be additional minimization of sideband energy, depending on the relationship between the modulating frequency and the carrier frequency.

Development of a Telemetry Synthesizer using a DDS

Considerable engineering development, along with the typical experiences with most of Murphy’s Laws, has been involved in development of the PSMs. The DDS products of several manufacturers have been the subject of analysis or experimentation at various times during the TDVS II system development. No two DDS devices have the exact same topology, and all have specific proprietary features which are oriented toward one or another class of application. One particular variation on the DDS, which contains internal modulation stages for FM and PM, is known as the Number Controlled Modulated Oscillator (emphasis added) or NCMO*.

At the time of this writing, the NCMO remains the only DDS on the market with a direct digital modulation input independent of the tuning input, and ported at the fundamental bit-width of the phase accumulator. This particular characteristic of the NCMO is critical to concepts employed in the PSM which simultaneously modulate the device and deviation map the resultant signal to the IRIG channel constraints.

The basic architecture of the NCMO (Digital RF Solutions Corporation, Santa Clara, CA) is shown in Figure 2. Note that while modulation capabilities are on-package, the waveform maps are external components. In addition, a single 12-bit external DAC is used with this device.

The PSMs in TDVS II are basically required to be capable of creating any kind of telemetry allowable in IRIG 106-86, in the channel quantities cited earlier. This includes baseband PCM at accurate bit rates and any permissible IRIG encoding; as well as a mix of IRIG FM subcarriers which may be modulated with PAM, PCM, or basic waveforms. During the development process, interesting characteristics related to applying the NCMO to each of these areas were discovered.

Telemetry Data Streaming: The DDS as a Bit-Rate Clock

The circuit approach for output streaming of a constructed telemetry data stream (generated on disk by TDVS II) is illustrated in Figure 3. In this scheme, a FIFO is used as a rate buffer between the data flow behavior of the VME computer system and the data rate required by the particular PCM, PAM, or FM data simulation. The FIFO depth is presently 8K based upon a service-time analysis; however, pin-compatible FIFO ICs are available at reasonable cost with volumes from 2K through 16K.

The DDS in this case is used simply as an inexpensive, highly accurate, programmable bit-rate clock, where the inherent numerical accuracy of a DDS approach means “no tweaking.” An encapsulated crystal timebase oscillator is used as the frequency source for the DDS. The timebase frequency in the PSM is 16,777,216(2^{24}) resulting in a tuneable frequency resolution of 1 Hertz/bit with the

Digital RF Solutions 24-bit NCMO. Since DDS devices inherit both their long term and short term stabilities from the timebase used; and since crystal oscillators, being single frequency devices, are designed with excellent drift and residual noise characteristics; the DDS output also possesses these properties, and theoretically improves upon them. Improvement of the timebase noise characteristic occurs by virtue of the DDS output frequency being lower than that of the reference clock, with phase noise improvement in accordance with the "20 log N" rule.

An interesting aspect of using the DDS as a bit-rate clock is found in the fact that the clock signal needed by subsequent logic must be a square wave (having sharp corners, and therefore containing very high frequency harmonic components); but the DDS has the fundamental constraint that the output can contain no frequency component above the Nyquist rate. One intuitive solution to this dilemma would derive from the fact the DDS, representing a digital process itself, already contains squarewaves in its operations.

So, why not use the output MSB before the waveform map stage as the source for a programmable clock? The problem with this approach is that while the DDS can be viewed as a specialized divider, in reality it is a phase step counter. The MSB of the DDS output does provide output frequency information; however, it also represents more than just output frequency information, especially at higher frequencies. The MSB exhibits degrees of edge jitter in all cases, except when the phase increment is set so that the output frequency is an exact integral subharmonic of the fundamental clock. This condition is not suitable for cases in which the DDS is being applied as a frequency programmable squarewave source, since it is expected to not exhibit significant jitter at any programmed output frequency.

Squarewave synthesis with the DDS is best performed by following the DAC with a low-pass filter to obtain a sine wave with suitable zero-crossing accuracy; then, applying a comparator-to-ground to detect the zero crossings. Considerable analysis has been applied to this subject by DDS device manufacturers, with the general (and not surprising) conclusion that the low-pass filter is the essential element for ensuring that jitter in the DDS output is eliminated.

The lowest bit-rates required in TDVS II are 64 bits per second, and the highest bit rates exceed 2 megabits per second. At the lower frequencies, a higher number of phase points are applied to the waveform map by the DDS, and the resultant output exhibits very low jitter. At the higher frequencies, the filter comes into play.

Programmable generation of the various IRIG encodings (i.e., NRZ, BI*N*, etc.) is implemented on a state machine contained in the registered PROM which precedes the circuit's output driver, and is clocked by the squared NCMO output. Pre-Modulation filtering for baseband PCM is handled at a later point in the system architecture through the use of commercially available precision programmable filters.

Digitally Generated IRIG FM Subcarriers

Basic IRIG FDM telemetry is defined by Table 3-1 (Proportional-Bandwidth Subcarriers) and Table 3-2 (Constant Bandwidth Subcarriers) of IRIG 106-86 (where the enclosed area of Table 3-2 contains the basic CBW channel set). In all, there are some 85 channels defined in the basic IRIG "set" of PBW and CBW channels. The design strategy in TDVS II provides that any PSM may be a source of not only baseband PCM; but also a source for any subcarrier from the basic IRIG set, where the programmably selected subcarrier may be modulated with PCM, PAM, or other waveforms. In addition, the amplitude of the individual subcarriers may be programmably scaled (via a multiplier, as in the case for amplitude modulation in Figure 1) to simulate a pre-emphasis relationship.

The purpose of IRIG FM simulation in TDVS II is to provide an integrated and programmable signal source for range validation purposes. This includes the ability to cycle through DAC-ID sequences on each channel, the ability to place standard five and eleven point calibration sequences on each channel (which in this case is done with incredible accuracy due to the digital technique employed), and the ability to provide simulation of telemetry data on FM subcarriers.

One design challenge in applying the NCMO to this function was development of a technique to establish deviation limits which correspond to the IRIG standards for each individual channel. This aspect of the application is called "deviation mapping," and two basic approaches to achieving this feature have been explored during the TDVS II development.

One approach is to apply a memory medium between the final modulating data, and the modulation port of the NCMO. This approach is illustrated in Figure 4. In the basic example shown, an 8-bit telemetry data value and a 7-bit IRIG channel address are presented to the deviation map memory. Through the deviation map, the telemetry data value is scaled to a 24-bit NCMO modulation port value such that the 256-element range of the 8-bit data value produces a 24-bit modulation value which falls exactly in the range of a particular IRIG channel. The effective data amplitude represented in the 8-bit telemetry data value modulates the NCMO frequency excursion, while the information added by the map simply ensures that the channel's frequency boundaries are never exceeded.

The deviation map can be implemented as a fixed structure in EPROM, or can be implemented with writeable memory so that its entries can be changed under program control. In addition, a programmable single entry map could be applied instead of a larger memory representing the deviation limits for all of the available channels.

Another approach to achieving the deviation mapping effect is to use a multiplier rather than a memory. While this approach is not illustrated, the basic idea is the same: to scale a value from its intrinsic bit-width, to the bit-width of the modulation port, within limits defined by the permissible IRIG (or non-IRIG, if required) frequency excursions.

Data Resolution

An amplitude resolution of 8-bits is fine for DAC-IDs or calibration points. However, amplitude quantization effects on an 8-bit digital representation of data (for example, a sine wave), result in a highly inaccurate waveform as the slope approaches zero (due to rounding), or if the data is scaled. However, there are alternatives to simply widening all the data paths.

One way to deal with this problem is to use 8-bit amplitude values as "phase-pointers" (only the zero crossings or the peaks are really necessary), and interpolate the individual time-sequence of the values into higher accuracy data-points at some convenient memory-bearing site such as the deviation map. Remember, we're dealing with amplitude at this point.

Local Generation of Repetitive Waveforms

The PSMs use a recirculating downloadable FIFO technique. This permits the PSM to be placed in a local mode in which the module will repeatedly issue its data set through the NCMO (or as baseband PCM, clocked by the squared-NCMO) up to the depth of the FIFO (2K to 16K). In the IRIG FM case, this mode is used for DAC-ID data, five and eleven point calibration data, and for simple waveforms such as sines, ramps, triangles, etc. As indicated, a PCM frame which is structured so that its total volume falls within the FIFO depth may also be recycled, as well as a PAM data frame set meeting the same total data volume criteria.

PCM/FM and PAM/FM - Modulation with Sharp-Edged Signals

When a series of baseband PCM signal states are applied to the NCMO modulation port, the output carrier hops back and forth between the two numerically-selected frequencies (i.e., FSK). While the NCMO doesn't violate the laws of physics, and the expected result of some "spreading" is apparent in the output spectrum, the technology does offer spectral performance improvements over most prior-generation SCOs. Frequency slewing an SCO is simply not the same as the frequency hops of a continuous phase digitally generated FM device, either mathematically or in practice.

Spectral results have been observed in which NCMO generated calibration subcarriers exhibited 40-45dbm peak-to-trough levels for adjacent IRIG channels, while a conventional range calibrator didn't exceed 35dbm for the same adjacent channels. Comparative performance in the lower IRIG channels can be especially striking.

Sharp-edged waveforms can be "shaped" before presentation to the NCMO modulation port to achieve improved spectral performance, if the application warrants. The digital version of pre-modulation filtering is called "transition mapping," or "transition shaping." A low-cost implementation technique involves the use of a DSP assisted interpolator, but is rate-limited to about 40 Kilobit PCM by the performance parameters of current-generation DSPs.

One novel approach to providing frequency interpolation between synthesis steps is based on dynamically changing the modulation mode of the NCMO. A simple counter can be employed to generate modulation port sweep values. If the NCMO is in phase modulation mode, as the speed of the counter is varied, the output frequency changes accordingly because a linear phase ramp produces a frequency shift. Either a state machine or a DSP can be employed to change the modulation mode and control the counter, depending upon the parameters of the application.

It is important to note that transition shaping may not be required at all in IRIG FM telemetry applications due to the intrinsic spectral purity and precise numerical modulation of the NCMO, when considered in the light of normal performance of existing SCOs and discriminators.

Special Considerations for DDS Application

An important area of special consideration when dealing with DDS devices is the area of instrumentation problems (test bench instruments, not range instrumentation). In general, spectrum analysis test equipment can easily mask actual performance of a DDS. Since all test equipment exhibits some degree of non-linear performance, and since DDS devices by nature can generate many output signals, spectrum analyzer non-linearities which mix DDS created signals with instrument created images can result in highly confusing situations.

The NCMO is a 5 volt device, and when associated with a DAC in a typical circuit with a ±5 volt supply, can produce output products which essentially overdrive the dynamic range capability of the spectrum analyzer. In general, highly detailed analysis of NCMO spectral performance requires a precision attenuator at the analyzer input. Spurious signals displayed by the spectrum analyzer may well be artifacts of the analyzer's radio circuitry, and not a characteristic of performance of the DDS. Even the highest-quality analyzers may require 30db of input attenuation, and many units will require 40db to guarantee sufficient dynamic range. Some older analyzers simply can't be applied to work with DDS designs in scenarios where "those last few db" are significant.

Another special consideration, and one of the most challenging aspects of applying a DDS does not stem from the device itself, but is found in the realization of the extent to which engineers are educated in an analog world. The "tank circuit mentality" can be very hard to overcome, both in the developer, and among those attending development design reviews. A sign in the TDVS II development area asks: "What is the slew-rate from zero to one?"

Manufacturers in the DDS industry itself have been largely responsible for creating the impression that DDS devices are somehow identical to conventional oscillators, giving rise to the Idea that conventional wisdom regarding the behavior of various earlier generation linear systems may be directly applied to the DDS. An oscillator is a linear system, subject essentially to "inertial" effects proportional to the Q of a resonator. A digital synthesizer is just not well related to the world of LaPlace transforms and the like, and exhibits no "inertial effects."

A useful visual aid when presenting the idea of a DDS application to "linear observers" is to have an old broom handy. The idea is to hold the broom out at arm's length while visualizing that the end of the broom is a pointer to a particular frequency. In order to point to a different frequency it is necessary to swing the broom around to point in another direction, which involves the rate at which one swings the broom, and the physical instability of stopping in just the right place (i.e., ringing).

The DDS is so conceptually different that the broom is next placed on the floor. The visualization is that of being radially surrounded by a large number of "virtual brooms," where selecting a new frequency (for tuning or modulation) is simply a matter of selecting which "virtual broom" to pick up. There is no inertia, and all the brooms could be sequentially picked up in any order. The collection of "virtual brooms" represents the fundamental concept of the "signal palette" of the DDS.

A real DDS can, in fact, rapidly hop from frequency to frequency, doing so as a result of frequency modulation by various sharp-edged waveforms, while maintaining remarkable spectral purity in its output. The process is the number-controlled result of a sampled data system which is capable of changing its output pattern from the phase-value sequence for one frequency to the phase-value sequence for another frequency at a point of zero phase slope between two successive cycles of the phase circle. The principal issue in its application has to do with the waveform reconstruction process, and is centered around the DAC and the subsequent filter used for waveform reconstruction.

Possibilities for Future Telemetry Systems

Although TDVS II is a development project, and not a research project, one of the most challenging aspects of the program has been to keep abreast of the rate at which today's laboratory research turns into tomorrow's products. When the author was first presented with the idea of TDVS II as a generalized telemetry simulator, DDS technology was basically a laboratory curiosity. Today, DDS technology is reaching the status of altering the basic concepts, techniques, and economics associated with electronic communication in the both the military and commercial marketplaces.

Some papers presented at various episodes of this forum in recent years have tended to see the future in terms of ever higher PCM bit-rates. However, the real figure-of-merit for any data communications system is actually the ratio of the data architecture's bit-rate to the bandwidth applied in the information transfer. This ratio is the bandwidth efficiency of the information system.

If more signal states could be introduced into an information communication system, the bandwidth efficiency could be increased. Bandwidth is consumed every time a signal changes state. If more signal states were applied in a system, each state could represent a larger set of information, and the system would require fewer state changes to transfer a desired collection of information sets.

In any communication system, the information bearing signal must conform to the limitations of its channel. While the bit-streams to be transmitted are inherently discrete-time (DT), the physical media has traditionally been viewed as continuous-time (CT) in nature. As a result, modulation, the process of imposing the bit-stream on the media, is viewed as a means of representing the discrete-time bit stream as a continuous-time signal. The intrinsically quantized nature of DDS technology makes alternative views of the modulation process possible.

Recently, the academic community has begun to back away from the ever increasing alphabet soup of acronyms which are used to indicate different forms of modulation. In this low-letter-count view, most of what we deal with is PAM of one sub-type or another, in which a sequence of time-translates of a basic pulse is amplitude modulated by a sequence of data symbols. When the primitive idea of baseband PAM is extended to passband transmission by introduction of a sinusoidal carrier signal, the whole gamut of alphabet-soup techniques become available as special cases of passband PAM. Important examples include phase-shift keying (PSK), amplitude and phase modulation (AM-PM), and quadrature amplitude modulation (QAM).

Recall that with the NCMO, all the basic means for applying modulation are directly available in digital (hence, DT) form; meaning, in essence, that the NCMO can be considered to be polar in nature, and number-driven in technique. Those familiar with modem technology will also be aware that modern data communication carrier systems apply amplitude coding in conjunction with phase shifts in the form of passband PAM known as QAM. Conventional QAM uses two modulators simultaneously and independently to convert the desired output signal vector into its Cartesian (x,y) coordinates and perform a vector addition to create the final output.

Since the modulation inputs in a conventional QAM system are polar coordinates (i.e., each signal state exhibits a unique combination of carrier and phase), the combination of signal states is commonly presented on a polar display and referred to as a signal constellation. Well-known constellations are 16-QAM, and the higher-order types of v.29 and v.32 communications. Note that bandwidth efficiency (bits/second/hertz) increases with the use of a higher number of precision signal states.

Since the numeric synthesizer is polar in nature, it can accept digital input data directly in polar coordinates. Instead of Quadrature Amplitude Modulation, the device directly accommodates vector (or polar) data modulation formats. In the telemetry world, the acronym “PDM” used to mean Pulse Duration Modulation. However, the term PDM is now a registered trademark of the NCMO manufacturer, Digital RF Solutions. PDM now means Polar Data Modulation, the direct digital relative of QAM (i.e., digital number driven QAM).

The significance of this new-generation PDM idea is that the full digital resolution of both the amplitude and phase modulators of a DDS system can be applied to the representation of signal states. If the twelve bits of the NCMO phase modulation input (4096 signalling states) is considered in conjunction with just an 8-bit amplitude multiplier (256 states), then over a million unique output signal vectors are possible (4096x256=1,048,576). Each of these output vectors is uniquely generated from its numeric coordinates, with the particular numeric accuracy of the modulator inputs (consider a 12-bit amplitude multiplier, or 4096x4096). While the commercial applications of this line of thinking are of considerable significance (both technologically and economically), additional possibilities for future telemetry approaches also exist, some of which have been discussed at previous ITCs.

During the development of TDVS II, the author has often encountered curiosity regarding the telemetry synthesis aspect of the TDVS II project. The usual reaction to the PSM, a VXI/VME PC board with two complete VME-interfaced DDS systems on-board, is one of bewilderment and: “But, aren’t those things really just frequency sources?” Trying to explain that the NCMO (and DDS technology in general) can be used either to establish PCM data rates, or as a digital sub-carrier oscillator, or as a source of newer, more complex forms of telemetry, typically produces some unique reactions. But when the ability to do PAM with the technology is mentioned, the reaction is almost consistently: “PAM is dead.” Of course, in our world of telemetry systems, PDM is dead too... isn’t it?

Concluding Comments

Direct digital synthesis technology has been employed in a new range telemetry simulator, and is beginning to appear in various commercial telemetry products. The cost of single-chip DDS devices has fallen by over 500% in the past five years, making these devices available and cost effective for a wide spectrum of applications. Trends in digital communication techniques which are finding application in the telemetry field can often be directly implemented using DDS technology. The inherent numerical accuracy, directly accessible digital tuning and modulation features, characteristic stability, and intrinsic spectral purity of DDS devices ensure that they will find their way into many important future telemetry system innovations.

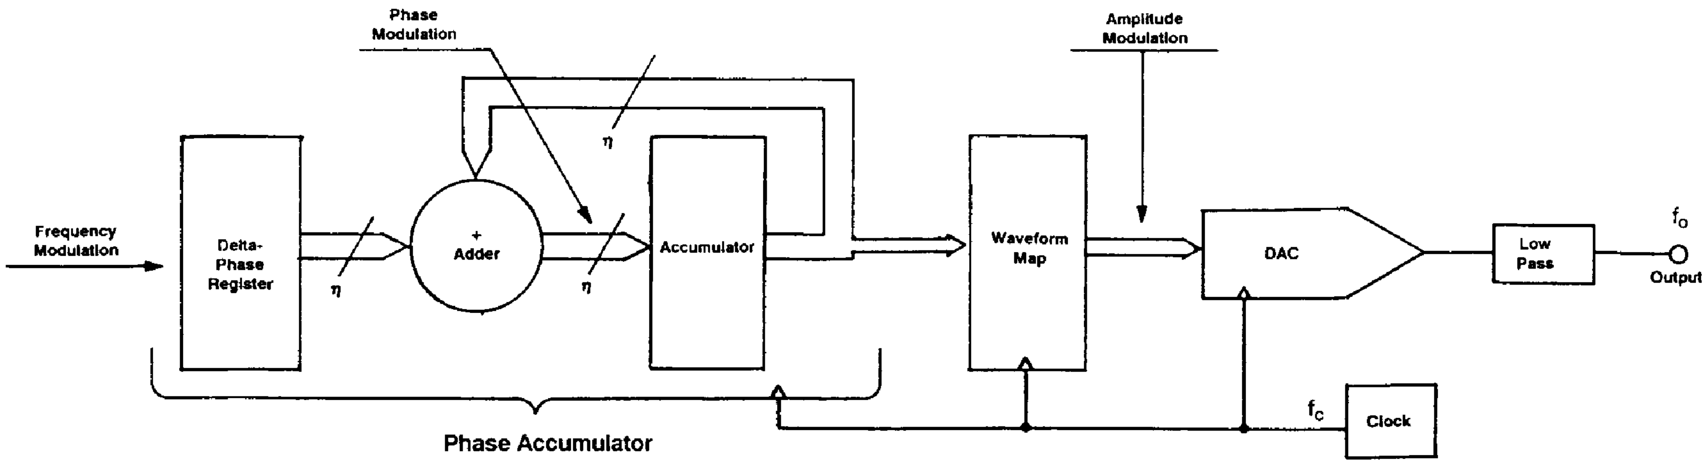

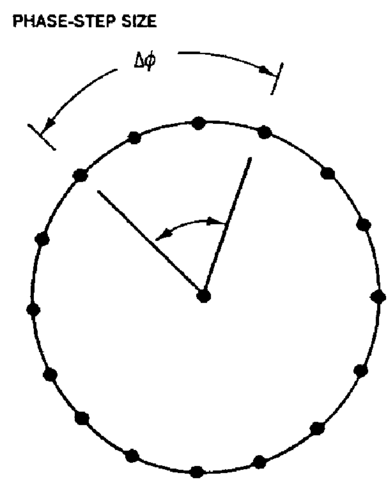

DIGITAL PHASE CIRCLE

$\underline{\eta}$	NUMBER OF PHASE POINTS
8	256
12	4096
16	65536
20	1048576
24	16777216
28	268435456
32	4294967296

$$f_o = \frac{\Delta\phi \cdot f_c}{2^{\eta}} \Bigg\} \text{ DIRECT-DIGITAL-SYNTHESIZER TUNING EQUATION}$$

FIGURE I - TOPOLOGY OF THE DDS

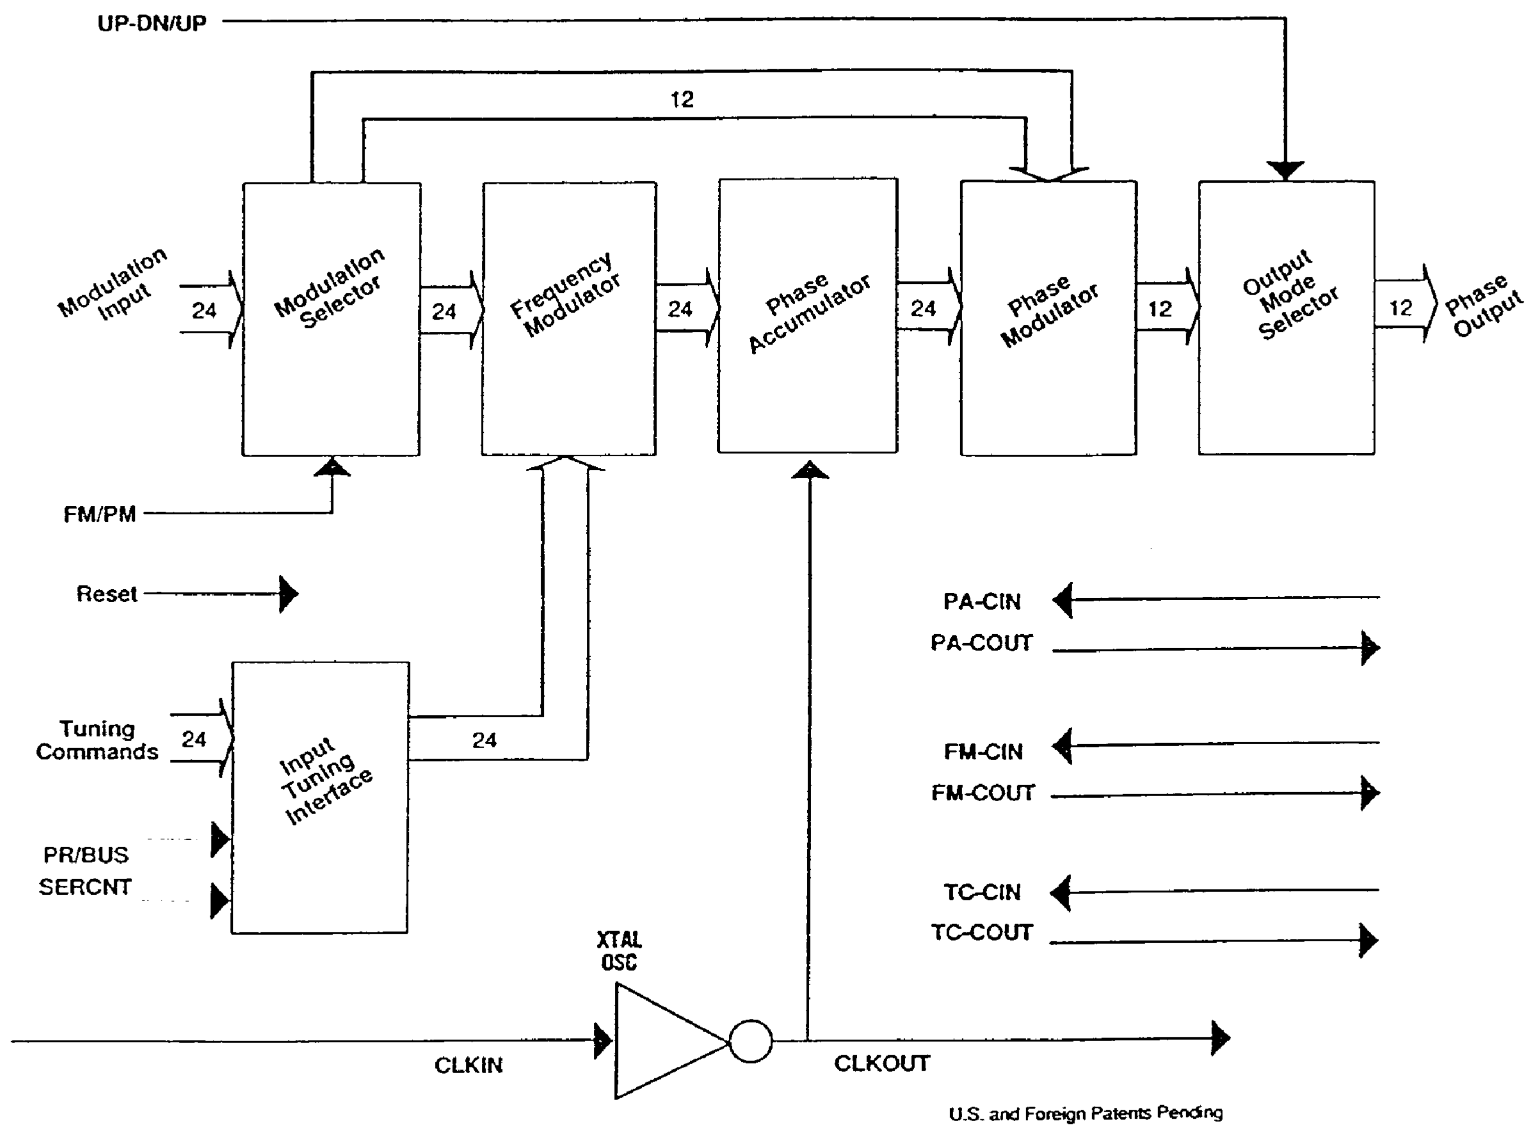

FIGURE 2 - ARCHITECTURE OF THE NCMO

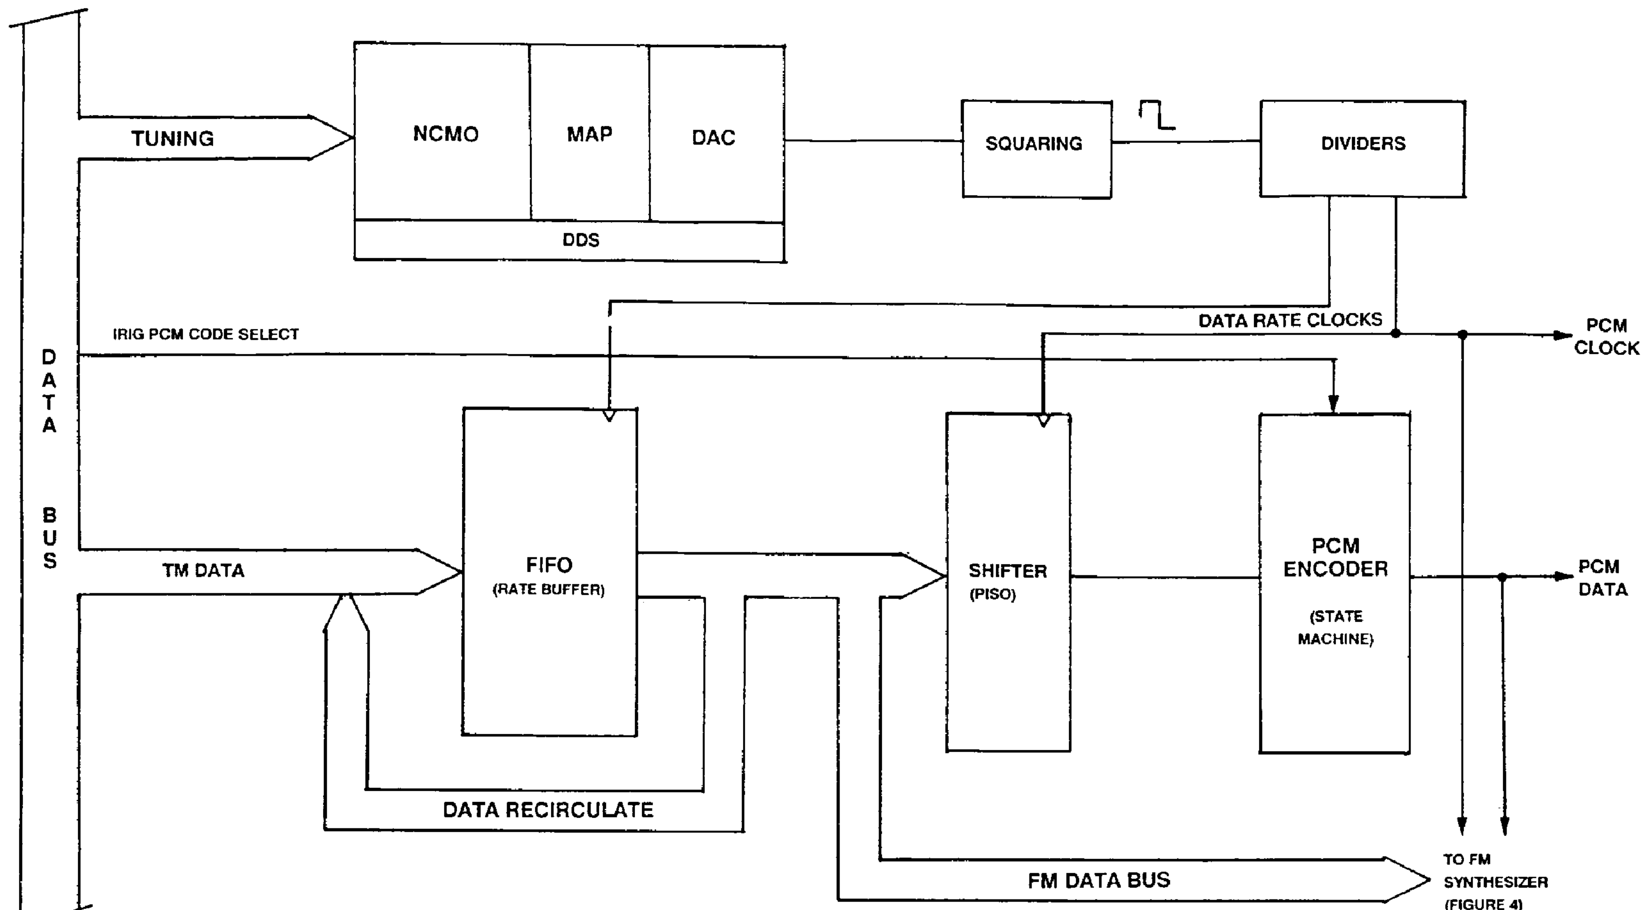

FIGURE 3 - THE DDS AS A DATA RATE CLOCK

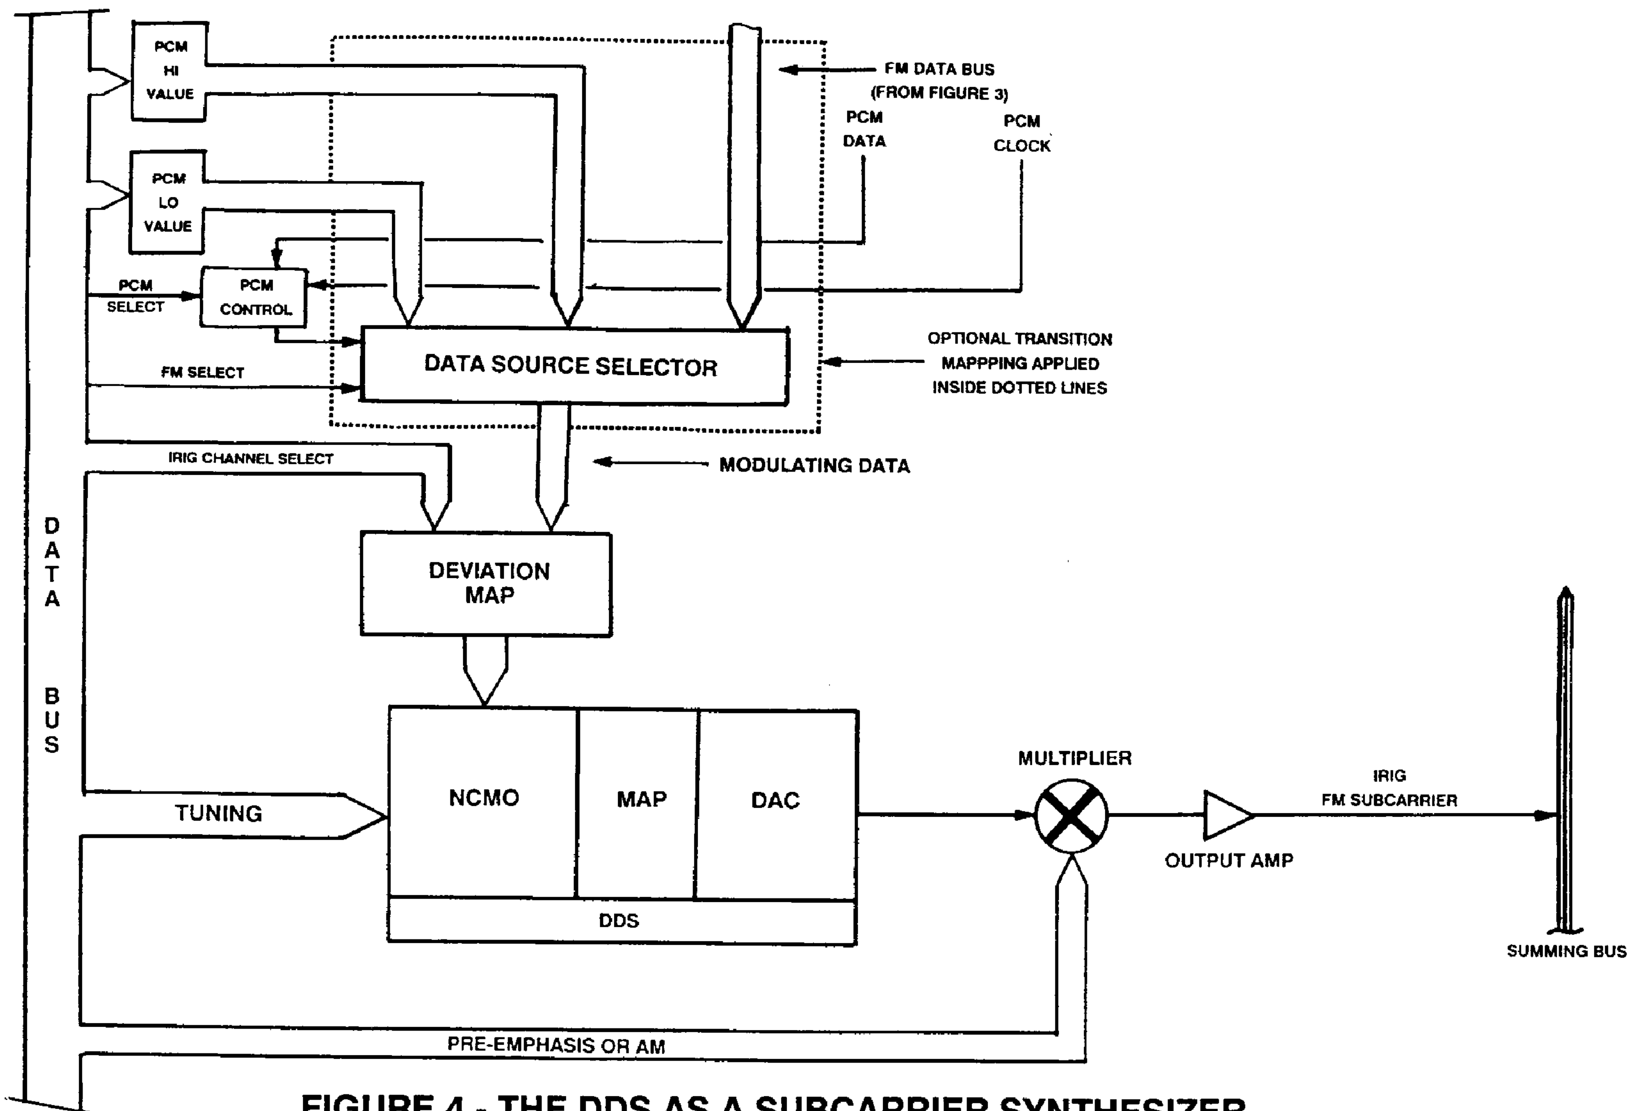

FIGURE 4 - THE DDS AS A SUBCARRIER SYNTHESIZER

TABLE 1

TDVS II Data Structuring Capabilities

TDVS II is capable of generating telemetry with the following characteristics:

From 20 to 65,535 bits per major frame.
From 1 to 99 major frame formats in one simulation.
From 1 to 256 minor frames per major frame.
From 1 to 4095 word time slots per minor frame.
From 0 to 5 submultiple frames with individual synchronization, each of which may be from 2 to 256 words long.
From 0 to 99 strapping sets per measurement.
From 4 to 64 bits per word time slot may be defined for telemetry simulation.
Odd, even, or no parity on each word time slot. Parity bit leading or trailing.
From 4 to 64 bits for sync and fill data.
From 1 to 64 bits per syllable.
From 1 to 9 syllables per word time slot.
From 1 to 4096 measurements may be defined in a major frame.
From 1 to 64 bits per measurement.
From 1 to 64 measurement formats.
From 1 to 64 bits per subfield. Each measurement may be divided into any number of subfields, up to the number of bits it contains.

Measurements may be sent MSB or LSB first.

Integer values can be expressed from 1 to 48 bits. Integer formats may be signed or unsigned. Signed integers may be in one’s complement, two’s complement, or sign and magnitude format.

Floating point words are user definable as to the number and order of bits for the mantissa and the exponent, and the location of the sign bit. The precision of the data inserted into those words is limited to 52 bits for the mantissa, and the exponent values are limited to 11 bits. Conventional or “hidden bit” normalization is supported. Biased or standard exponent notation is supported.

IRIG PCM modulation codes as defined in the IRIG Telemetry Standards are supported as specified.

IRIG standard FM subchannels as defined in the IRIG Telemetry Standards are supported as specified.

REFERENCES

Direct Digital Synthesis - Twenty Years of Development

1. Tierney, J., Rader, C., & Gold, B.; “A Digital Frequency Synthesizer,” IEEE Transactions on Audio and Electronics, March, 1971.

2. Cooper, H., “Why Complicate Frequency Synthesis?”, Electronic Design, July 19,1974.

3. Bjerede, B. & Fisher, G., “A New Phase Accumulator Approach to Frequency Synthesis,” Proceedings of IEEE NAECON 76, May, 1976. Also: “An Efficient Hardware Implementation for High Resolution Frequency Synthesis,” Proceedings of the 30th Annual Frequency Control Symposium, IEEE, June, 1976.

4. Moore, F.R., “Table Lookup Noise for Sinusoidal Digital Oscillators,” Computer Music Journal, April, 1977.

5. Tierney, J. (et.al.), “A Digital Frequency Synthesizer,” IEEE Transactions on Audio and Electroacoustics, March, 1978.

6. Nossen, E., “Digital Frequency Synthesis,” United States Patent 4,206,425, 3 June 1980.

7. Wheatly, C. & Phillips, D., “Spurious Suppression in Direct Digital Synthesizers,” Proceedings of the 35th Annual Frequency Control Symposium, IEEE, May, 1981.

8. Hoppes, P., “A Radiation Hardened Low-Power Numerically Controlled Oscillator,” Proceedings of the Custom Integrated Circuits Conference, IEEE, May, 1982.

9. Wheatley, C., “Digital Frequency Synthesizer with Random Jittering for Reducing Discrete Spectral Spurs,” United States Patent 4,4010,954, 18 October 1983.

10. Williams, F., “A Digital Frequency Synthesizer,” QST, April, 1984.

11. Williams, F., “A Microprocessor Controller for the Digital Frequency Synthesizer,” QST, February, 1985.

12. Eisenson, H., “Frequency Hopping with Direct Digital Synthesis,” R. F. Design, December, 1985.

13. Kisenwether, E. & Troxell, W., “Performance Analysis of the Numerically Controlled Oscillator,” Proceedings of the 40th Annual Frequency Control Symposium, IEEE, May, 1986.

14. Olsen, J. & Fishman, P., “Truncation Effects in Direct Digital Synthesis,” Twentieth Asilomar Conference on Signals, IEEE, November, 1986. (IEEE Pub. Jan. 1987)

15. Nicholas, H. & Samueli, H., “An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the presence of Phase-Accumulator Truncation,” Proceedings of the 41st Annual Frequency Control Symposium, IEEE, May, 1987.

16. Williams, F., “Direct Numerical Synthesis of Sinewaves,” R. F. Design, May, 1987.

17. McCune, E., “Control of Spurious Signals in Direct Digital Synthesizers,” Proceedings of RF Technology Expo, February, 1988.

18. Zavrel, R., “Digital Modulation using the NCMO,” R. F. Design, March, 1988.

19. Nicholas H., Samueli, H. & Kim B., “The Optimization of Direct Digital Synthesizer Performance in the Presence of Finite Word Length Effects,” Proceedings of the 42nd Annual Frequency Control Symposium, IEEE, June, 1988.

20. Gallant, J., “Direct Digital Synthesizers - Devices refine the art of frequency synthesis.,” EDN, November 9, 1989.

21. Edwin, A., “Direct Digital Synthesis Applications,” Microwave Journal, January, 1990.

22. McCune, E., “Digital Communications Using Direct Digital Synthesis,” R. F. Design, January, 1990.

23. McCune, E., “Data Modulation Using Direct Digital Synthesis,” Proceedings of RF Technology Expo, March, 1990.

*	NOTES:

(1). The term “NCO” means Numerically Controlled Oscillator or Number Controlled Oscillator, interchangeably. This term is used generically in the literature, and is not the subject of trademark or copyright.

(2). The term “NCMO” means Numerically Controlled Modulated Oscillator or Number Controlled Modulated Oscillator, interchangeably. This term is not used generically, and is a registered trademark of Digital RF Solutions Corporation. The terms “NCMO” and “PDM” are used herein with permission of the copyright holder.