

SECTION 4

DACS FOR DSP APPLICATIONS

- DAC Structures
- Low Distortion DAC Architectures
- DAC Logic
- Sigma-Delta DACs
- Direct Digital Synthesis (DDS)

DACs FOR DSP APPLICATIONS

SECTION 4

DACs FOR DSP APPLICATIONS

Walt Kester, James Bryant

DAC STRUCTURES

The most commonly used DAC structures (other than a simple 1-bit DAC based on a single switch used with a reference voltage) are binary weighted DACs or ladder networks, but these, though relatively simple in structure, require quite complex analysis. We will start by examining one of the simplest structures of all, the Kelvin divider shown in Figure 4.1. An N-bit version of this DAC simply consists of 2^N equal resistors in series. The output is taken from the appropriate tap by closing one of the 2^N switches by decoding 1 of 2^N switches from the N-bit data. Recent DACs using this architecture are referred to as “string DACs.”

SIMPLEST VOLTAGE OUTPUT DAC: THE KELVIN DIVIDER (“STRING DAC”)

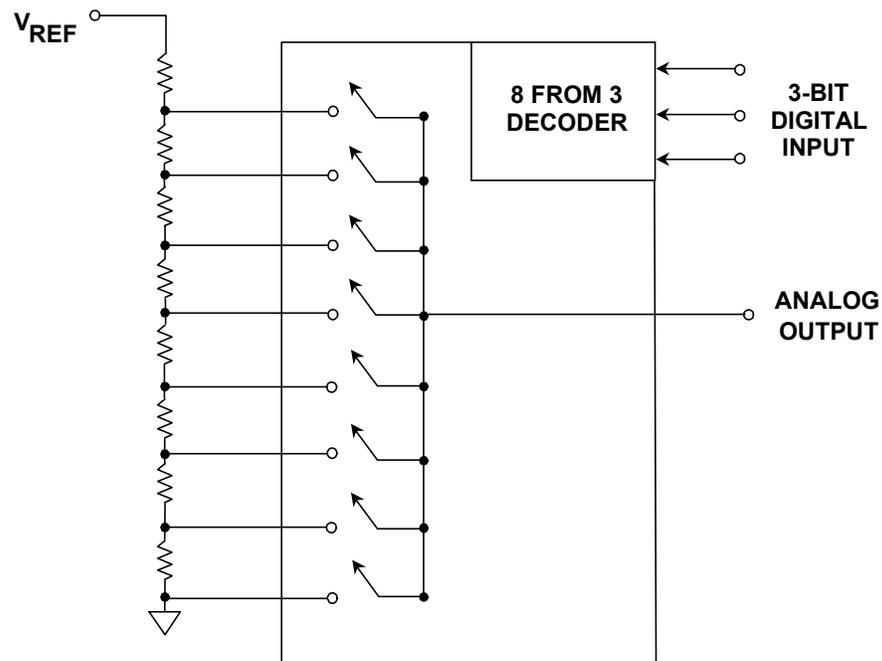


Figure 4.1

This architecture is simple, has a voltage output (but a code-varying Z_{OUT}), and is inherently monotonic (even if a resistor is zero, $OUTPUT_N$ cannot exceed $OUTPUT_{N+1}$). It is linear if all the resistors are equal, but may be made deliberately non-linear if a non-linear DAC is required. Since only two switches operate during a transition, it is a low-glitch architecture. Its major drawback is the

DACS FOR DSP APPLICATIONS

large number of resistors required for high resolution, and as a result it is not commonly used - but, as we shall see later, it is used as a component in more complex DAC structures.

There is an analogous current output DAC which consists, again, of 2^N resistors (or current sources), but in this case they are all connected in parallel between the reference voltage input and the virtual ground output (see Figure 4.2).

THE SIMPLEST CURRENT OUTPUT DAC

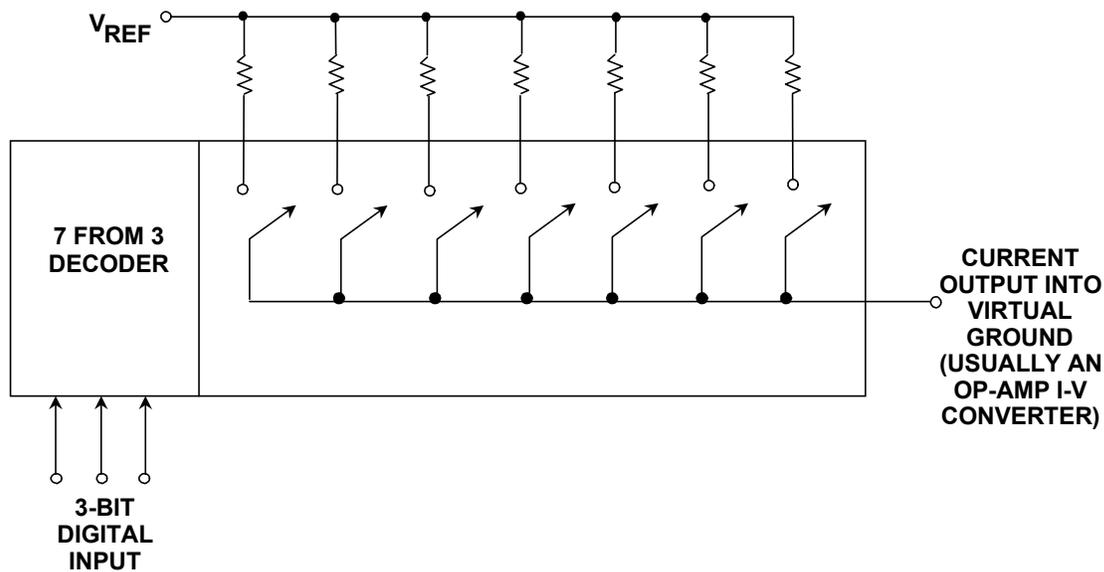


Figure 4.2

In this DAC, once a resistor is switched into circuit by increasing digital code, any further increases do not switch it out again. The structure is thus inherently monotonic, irrespective of inaccuracies in the resistors, and like the previous case, may be made intentionally non-linear where a specific non-linearity is required. Again, as in the previous case, the architecture is rarely, if ever, used to fabricate a complete DAC because of the large numbers of resistors and switches required. However, it is often used as a component in a more complex DAC structure.

Unlike the Kelvin divider, this type of DAC does not have a unique name, although both types are referred to as *fully decoded DACs* or *thermometer DACs* or *string DACs*.

Fully-decoded DACs are often used as components of more complex DACs. The most common are "segmented DACs" where part of the output of a fully decoded DAC is further subdivided. The structure is used because the fully decoded DAC is inherently monotonic, so if the subdivision is also monotonic, the whole resulting DAC is also monotonic.

A voltage segmented DAC (see Figure 4.3) works by further sub-dividing the voltage across one resistor of a Kelvin divider. The sub-division may be done with a further Kelvin divider (in which case the whole structure is known as a "Kelvin-Varley divider", or with some other DAC structure.

SEGMENTED VOLTAGE DACs

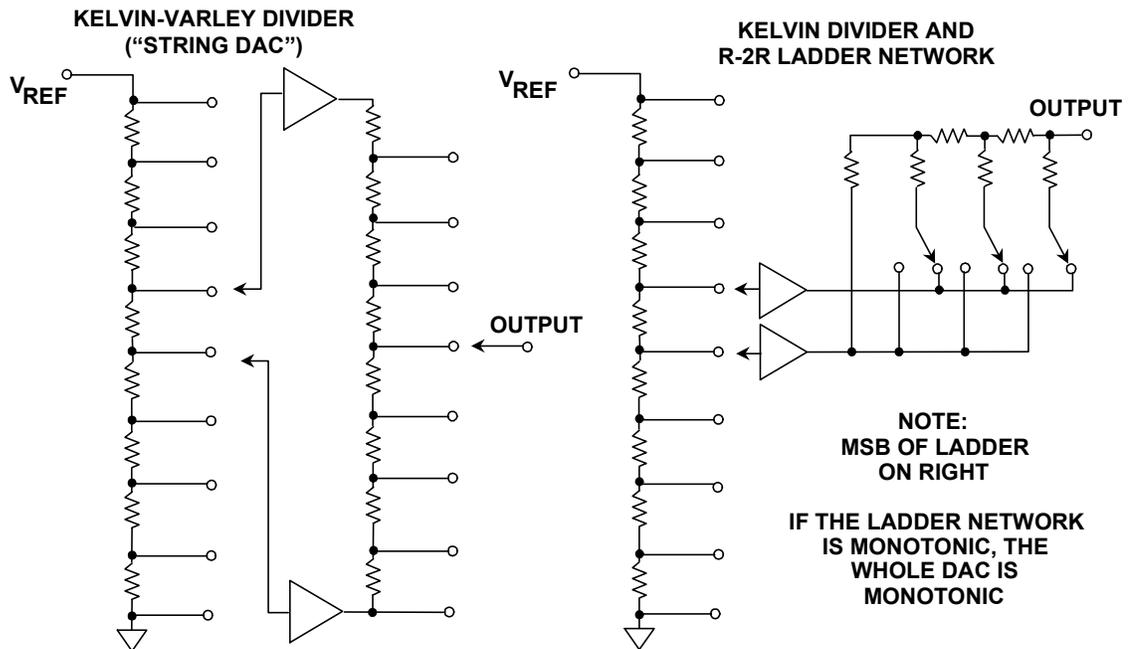


Figure 4.3

In all DACs, the output is the product of the reference voltage and the digital code, so in that sense, all DACs are multiplying DACs, but many DACs operate well only over a limited range of V_{ref} . True MDACs, however, are designed to operate over a wide range of V_{ref} . A strict definition of a multiplying DAC demands that its reference voltage range includes 0V, and many, especially current mode ladder networks with CMOS switches, permit positive, negative, and AC V_{ref} . DACs which do not work down to 0V V_{ref} are still useful, however, and types where V_{ref} can vary by 10:1 or so are often called MDACs, although a more accurate description might be "semi-multiplying" DACs.

LOW DISTORTION DAC ARCHITECTURES

Because of the emphasis in communications systems on DDS DACs with high SFDR, much effort has been placed on determining optimum DAC architectures. Practically all low distortion high speed DACs make use of some form of non-saturating current-mode switching. A straight binary DAC with one current switch per bit produces code-dependent glitches as discussed above and is certainly not the most optimum architecture (Figure 4.4). A DAC with one current source per code level can be shown not to have code-dependent glitches, but it is not practical to implement for high resolutions. However, this performance can be approached by

decoding the first few MSBs into a "thermometer" code and have one current switch per level. For example, a 5-bit thermometer DAC would have an architecture similar to that shown in Figure 4.5.

5-BIT BINARY DAC ARCHITECTURES

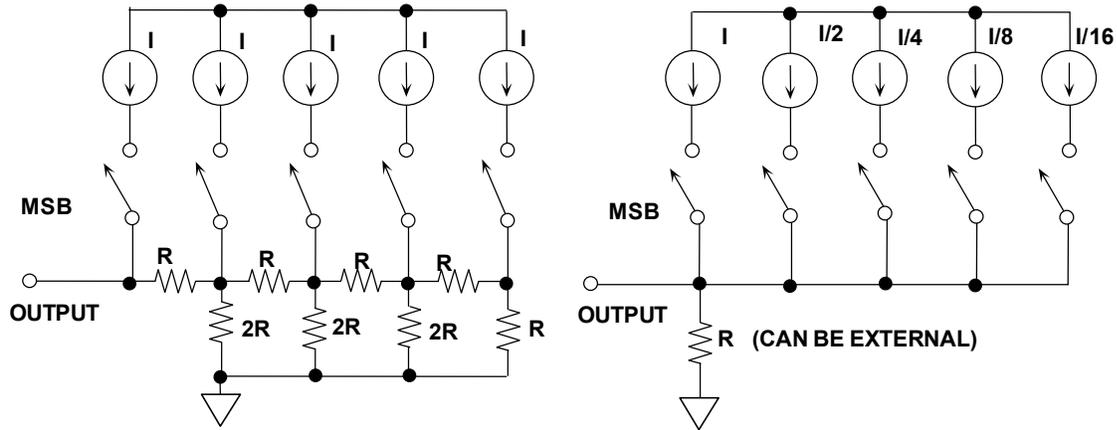


Figure 4.4

**5-BIT "THERMOMETER" OR "FULLY-DECODED" DAC
MINIMIZES CODE-DEPENDENT GLITCHES**

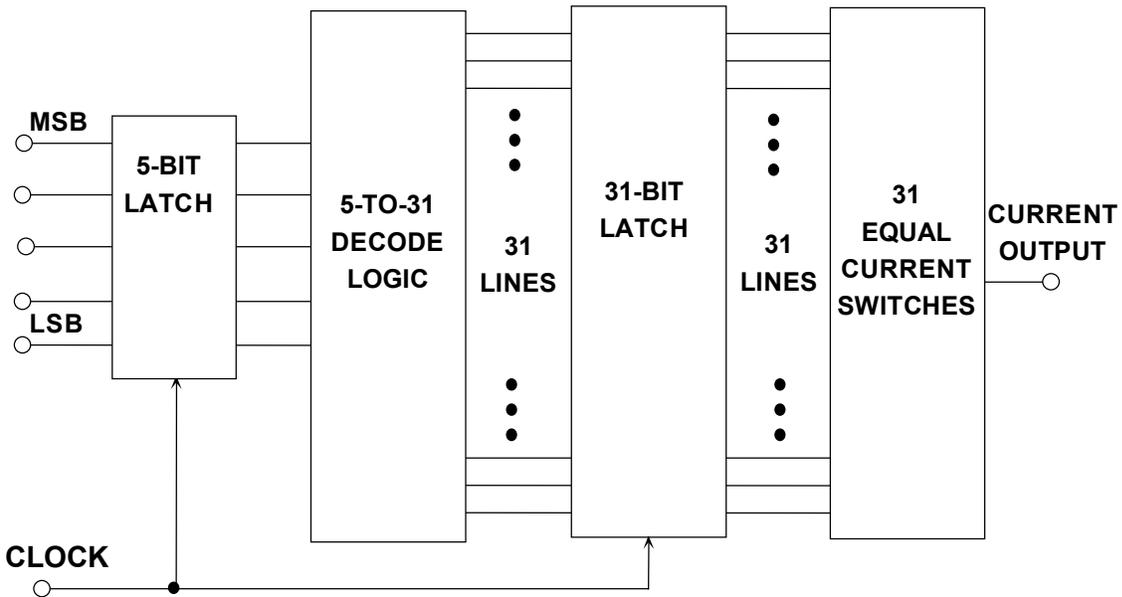


Figure 4.5

The input binary word is latched and then decoded into 31 outputs which drive a second latch. The output of the second latch drives 31 equally weighted current switches whose outputs are summed together. This scheme effectively removes nearly all the code-dependence of the output glitch. The residual glitch that does occur at the output is equal regardless of the output code change (it is code-independent) and can be filtered because it occurs at the DAC update frequency and its harmonics. The distortion mechanisms associated with the full-decoded architecture are primarily asymmetrical output slewing, finite switch turn-on and turn-off times, and integral nonlinearity.

The obvious disadvantage of this type of thermometer DAC is the large number of latches and switches required to make a 14, 12, 10, or even 8-bit DAC. However, if this technique is used on the 5 MSBs of an 8, 10, 12, or 14-bit DAC, a significant reduction in the code-dependent glitch is possible. This process is called *segmentation* and is quite common in low distortion DACs.

Figure 4.6 shows a scheme whereby the first 5 bits of a 10-bit DAC are decoded as described above and drive 31 equally weighted switches. The last 5 bits are derived from binarily weighted current sources. Equally weighted current sources driving an R/2R resistor ladder could be used to derive the LSBs, however, this approach requires thin film resistors which are not generally available on a low-cost CMOS process. Also, the use of R/2R networks lowers the DAC output impedance, thereby requiring more drive current to develop the same voltage across a fixed load resistance.

10-BIT SEGMENTED DAC

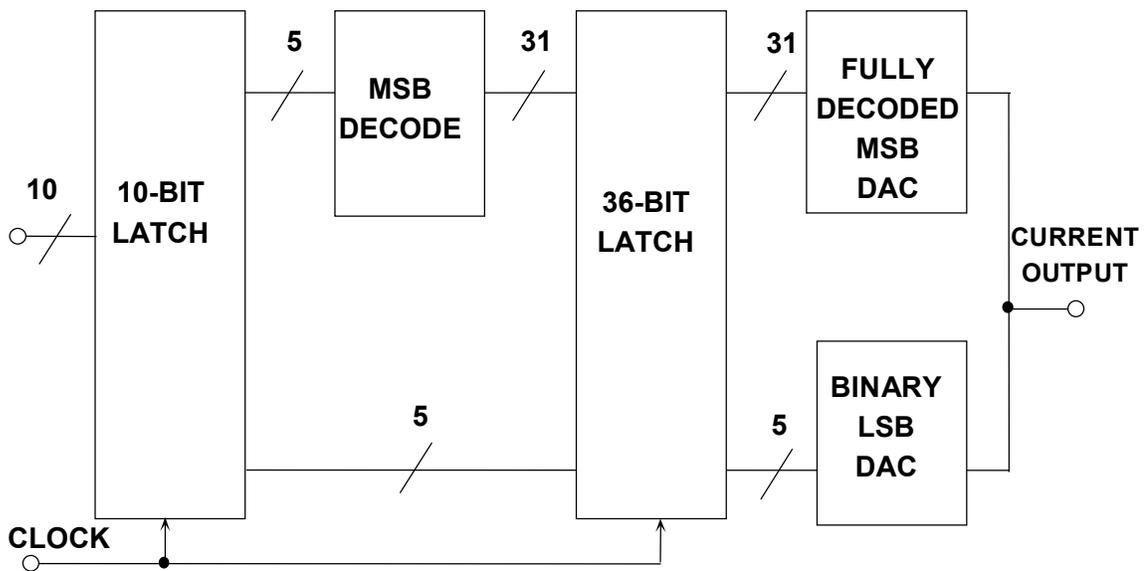


Figure 4.6

DACS FOR DSP APPLICATIONS

The AD9772 14-bit, 150MSPS TxDAC™ uses three sections of segmentation as shown in Figure 4.7. Other members of the AD977x-family and the AD985x-family also use this same core.

The first 5 bits (MSBs) are fully decoded and drive 31 equally weighted current switches, each supplying 512 LSBs of current. The next 4 bits are decoded into 15 lines which drive 15 current switches, each supplying 32 LSBs of current. The 5 LSBs are latched and drive a traditional binary weighted DAC which supplies 1 LSB per output level. A total of 51 current switches and latches are required to implement this architecture.

AD9772 TxDAC™ 14-BIT CMOS DAC CORE

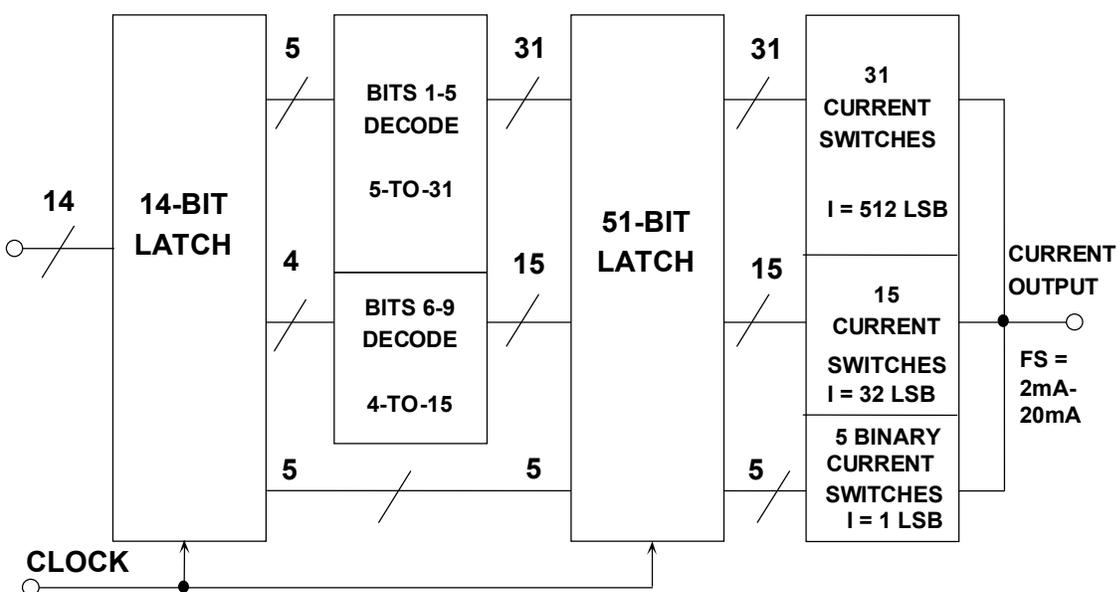


Figure 4.7

The basic current switching cell is made up of a differential PMOS transistor pair as shown in Figure 4.8. The differential pairs are driven with low-level logic to minimize switching transients and time skew. The DAC outputs are symmetrical differential currents which help to minimize even-order distortion products (especially when driving a differential output such as a transformer or an op amp differential I/V converter).

The overall architecture of the AD977x TxDAC™ family and the AD985x-DDS family is an excellent tradeoff between power/performance and allows the entire DAC function to be implemented on a standard CMOS process with no thin film resistors. Single-supply operation on +3.3V or +5V make the devices extremely attractive for portable and low power applications.

PMOS TRANSISTOR CURRENT SWITCHES

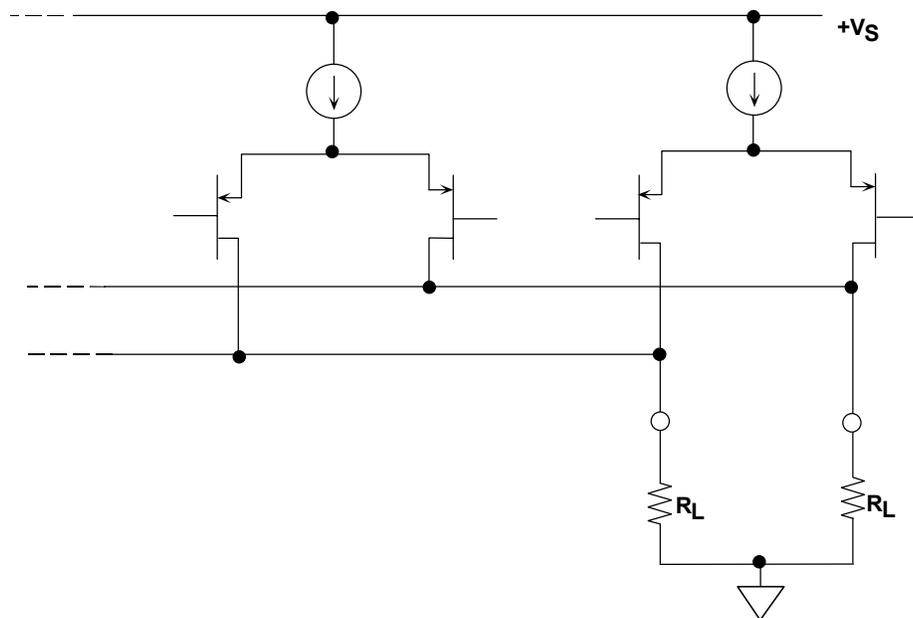


Figure 4.8

DAC LOGIC

The earliest monolithic DACs contained little, if any, logic circuitry, and parallel data had to be maintained on the digital input to maintain the digital output. Today almost all DACs have input latches, and data need only be written once, not maintained.

There are innumerable variations of DAC input structure which will not be discussed here, but the majority today are "double-buffered". A double-buffered DAC has two sets of latches. Data is initially latched in the first rank and subsequently transferred to the second as shown in Figure 4.9. There are three reasons why this arrangement is useful.

The first is that it allows data to enter the DAC in many different ways. A DAC without a latch, or with a single latch, must be loaded with all bits at once, in parallel, since otherwise its output during loading may be totally different from what it was or what it is to become. A double-buffered DAC, on the other hand, may be loaded with parallel data, serial data, or with 4-bit or 8-bit words, or whatever, and the output will be unaffected until the new data is completely loaded and the DAC receives its update instruction.

DOUBLE-BUFFERED DAC PERMITS COMPLEX INPUT STRUCTURES AND SIMULTANEOUS UPDATE

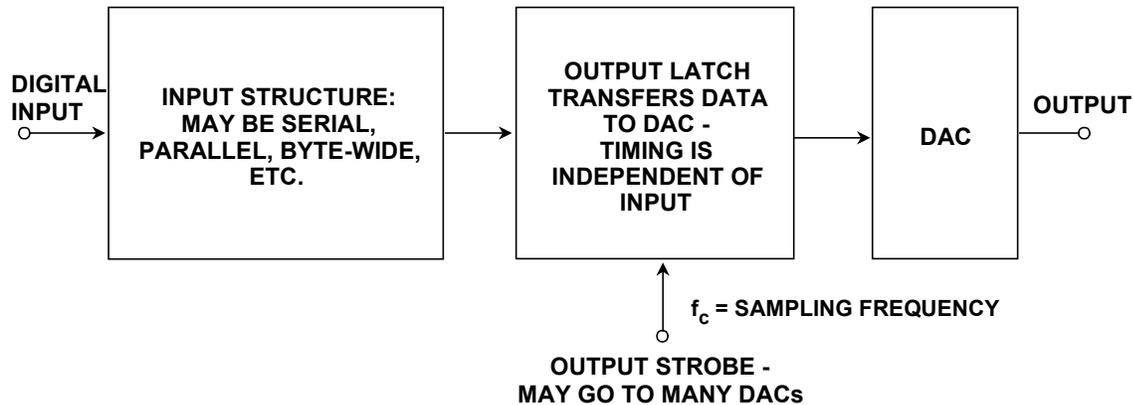


Figure 4.9

The second feature of this type of input structure is that the output clock can operate at a fixed frequency (the DAC update rate), while the input latch can be loaded asynchronously. This is useful in real-time signal reconstruction applications.

The third convenience of the double-buffered structure is that many DACs may be updated simultaneously: data is loaded into the first rank of each DAC in turn, and when all is ready, the output buffers of all DACs are updated at once. There are many DAC applications where the output of several DACs must change simultaneously, and the double-buffered structure allows this to be done very easily.

Most early monolithic high resolution DACs had parallel or byte-wide data ports and tended to be connected to parallel data buses and address decoders and addressed by microprocessors as if they were very small write-only memories (some DACs are not write-only, but can have their contents read as well - this is convenient for some applications but is not very common). A DAC connected to a data bus is vulnerable to capacitive coupling of logic noise from the bus to the analog output, and many DACs today have serial data structures. These are less vulnerable to such noise (since fewer noisy pins are involved), use fewer pins and therefore take less space, and are frequently more convenient for use with modern micro-processors, many of which have serial data ports. Some, but not all, of such serial DACs have data outputs as well as data inputs so that several DACs may be connected in series and data clocked to them all from a single serial port. The arrangement is referred to as "daisy-chaining".

Another development in DACs is the ability to make several on a single chip, which is useful to reduce PCB sizes and assembly costs. Today (2000) it is possible to buy sixteen 8-bit, eight 12-bit, four 14-bit, or two 16-/18/20/22/24-bit DACs in a single package. In the future, even higher densities are probable.

INTERPOLATING DACS

In ADC-based systems, oversampling can ease the requirements on the antialiasing filter, and a sigma-delta ADC has this inherent advantage. In a DAC-based system (such as DDS), the concept of interpolation can be used in a similar manner. This concept is common in digital audio CD players, where the basic update rate of the data from the CD is about 44kSPS. "Zeros" are inserted into the parallel data, thereby increasing the effective update rate to 4-times, 8-times, or 16-times the fundamental throughput rate. The 4x, 8x, or 16x data stream is passed through a digital interpolation filter which generates the extra data points. The high oversampling rate moves the image frequencies higher, thereby allowing a less complex filter with a wider transition band. The sigma-delta 1-bit DAC architecture represents the ultimate extension of this concept and has become popular in modern CD players.

The same concept can be applied to a high speed DAC. Assume a traditional DAC is driven at an input word rate of 30MSPS (see Figure 4.10A). Assume the DAC output frequency is 10MHz. The image frequency component at $30-10 = 20\text{MHz}$ must be attenuated by the analog antialiasing filter, and the transition band of the filter is 10 to 20MHz. Assume that the image frequency must be attenuated by 60dB. The filter must therefore go from a passband of 10MHz to 60dB stopband attenuation over the transition band lying between 10 and 20MHz (one octave). A Butterworth filter design gives 6dB attenuation per octave for each pole. Therefore, a minimum of 10 poles is required to provide the desired attenuation. Filters become even more complex as the transition band becomes narrower.

Assume that we increase the DAC update rate to 60MSPS and insert a "zero" between each original data sample. The parallel data stream is now 60MSPS, but we must now determine the value of the zero-value data points. This is done by passing the 60MSPS data stream with the added zeros through a digital interpolation filter which computes the additional data points. The response of the digital filter relative to the 2-times oversampling frequency is shown in Figure 4.10B. The analog antialiasing filter transition zone is now 10 to 50MHz (the first image occurs at $2f_c - f_o = 60 - 10 = 50\text{MHz}$). This transition zone is a little greater than 2 octaves, implying that a 5- or 6-pole Butterworth filter is sufficient.

The AD9772 is a 2-times oversampling interpolating 14-bit DAC, and a simplified block diagram is shown in Figure 4.11. The device is designed to handle 14-bit input word rates up to 150MSPS. The output word rate is 300MSPS maximum. For an output frequency of 60MHz, and an update rate of 150MHz, and an oversampling ratio of 2, the image frequency occurs at $300\text{MHz} - 60\text{MHz} = 240\text{MHz}$. The transition band for the analog filter is therefore 60MHz to 240MHz. Without oversampling, the image frequency occurs at $150\text{MHz} - 60\text{MHz} = 90\text{MHz}$, and the filter transition band is 60MHz to 90MHz.

**ANALOG FILTER REQUIREMENTS FOR $f_o = 10\text{MHz}$:
 $f_c = 30\text{MSPS}$, AND $f_c = 60\text{MSPS}$**

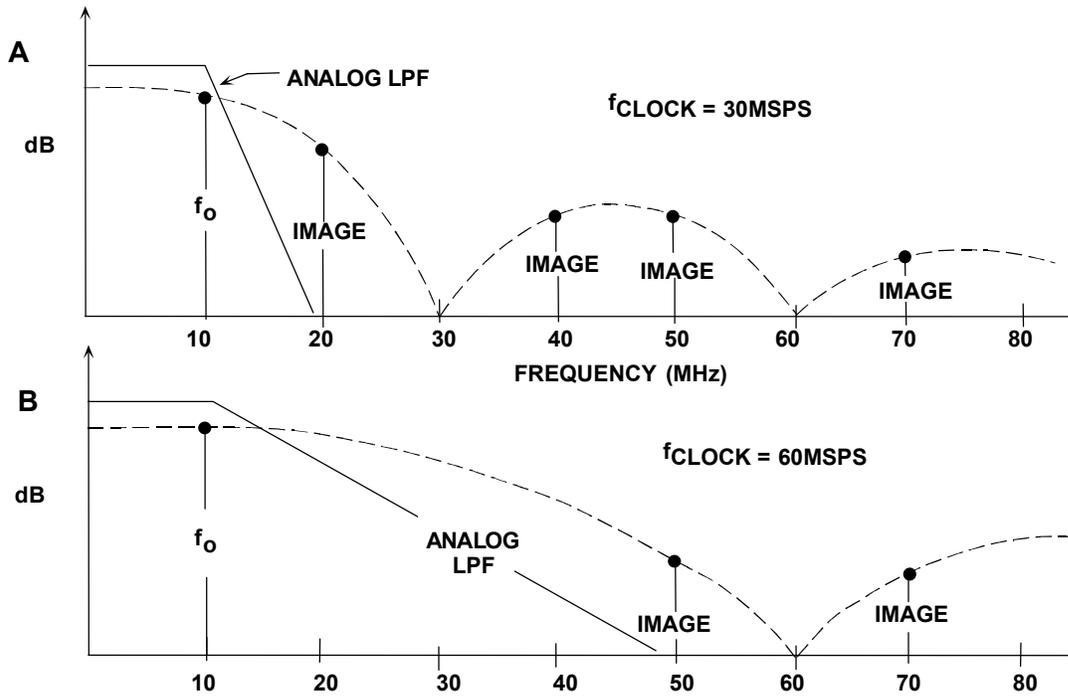


Figure 4.10

AD9772 14-BIT, 150MSPS INTERPOLATING TxDAC™

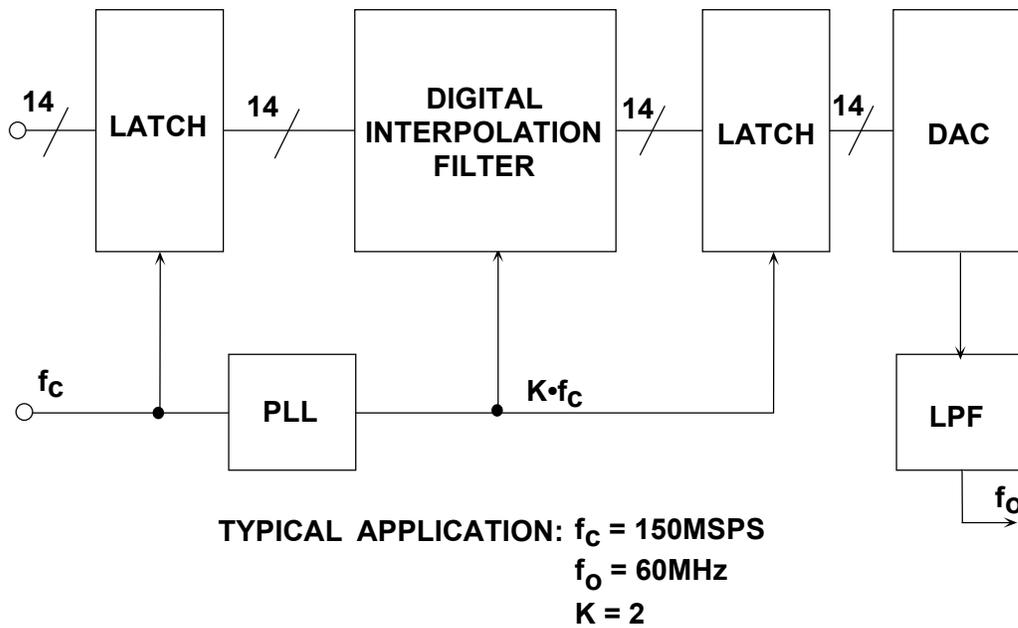


Figure 4.11

SIGMA-DELTA DACS

Another way of obtaining high resolution is to use oversampling techniques and a 1-bit DAC. The technique, known as sigma-delta (Σ - Δ), is computation intensive, so has only recently become practical for the manufacture of high resolution DACs, but since it uses a 1-bit DAC, it is intrinsically linear and monotonic.

A Σ - Δ DAC, unlike the Σ - Δ ADC, is mostly digital (see Figure 4.12A). It consists of an "interpolation filter" (a digital circuit which accepts data at a low rate, inserts zeros at a high rate, and then applies a digital filter algorithm and outputs data at a high rate), a Σ - Δ modulator (which effectively acts as a low pass filter to the signal but as a high pass filter to the quantization noise, and converts the resulting data to a high speed bit stream), and a 1-bit DAC whose output switches between equal positive and negative reference voltages. The output is filtered in an external analog LPF. Because of the high oversampling frequency, the complexity of the LPF is much less than the case of traditional Nyquist operation.

SIGMA-DELTA ($\Sigma\Delta$) DACs: 1-BIT AND MULTIBIT

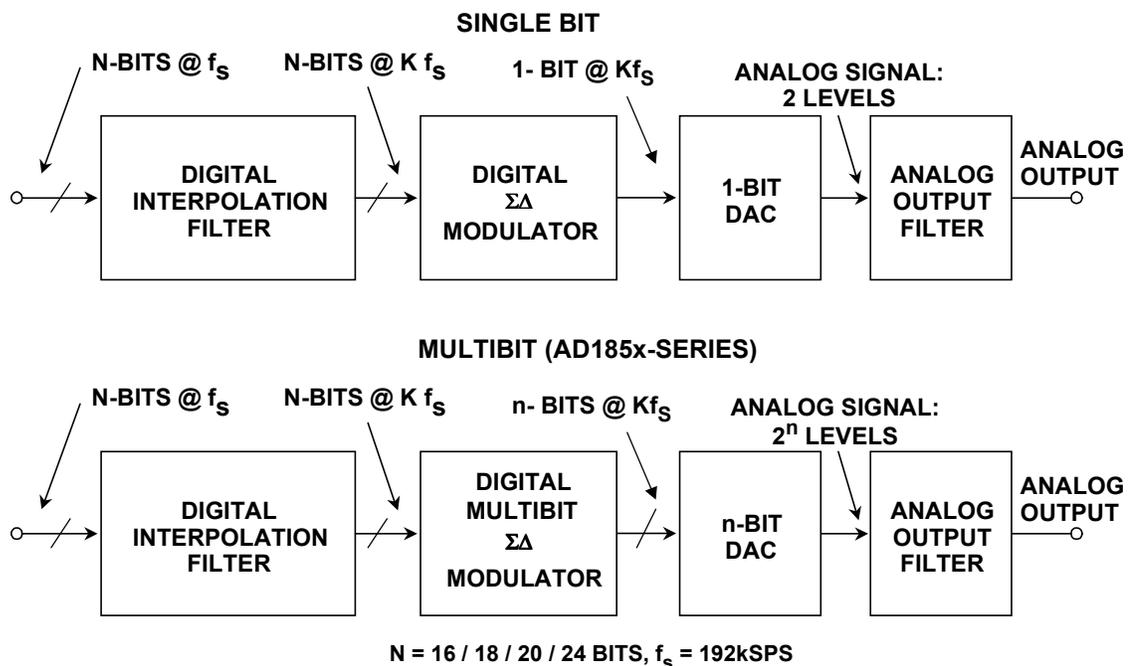


Figure 4.12

It is possible to use more than one bit in the DAC, and this leads to the *multibit* architecture shown in Figure 4.12B. The concept is similar to that of the interpolating DAC previously discussed, with the addition of the digital sigma-delta modulator. In the past, multibit DACs have been difficult to design because of the accuracy requirement on the n-bit internal DAC (this DAC, although only n-bits, must have the linearity of the final number of bits, N). The AD185x-series of audio DACs, however use a proprietary *data scrambling* technique (called *direct data*

scrambling, or D²S) which overcomes this problem and produces excellent performance with respect to THD + N. For instance, the AD1853 dual 24-bit, 192kSPS DAC has greater than 115dB THD + N at a 48kSPS sampling rate.

DIRECT DIGITAL SYNTHESIS (DDS)

A frequency synthesizer generates multiple frequencies from one or more frequency references. These devices have been used for decades, especially in communications systems. Many are based upon switching and mixing frequency outputs from a bank of crystal oscillators. Others have been based upon well understood techniques utilizing phase-locked loops (PLLs). This mature technology is illustrated in Figure 4.13. A fixed-frequency reference drives one input of the phase comparator. The other phase comparator input is driven from a divide-by-N counter which is in turn driven by a voltage-controlled-oscillator (VCO). Negative feedback forces the output of the internal loop filter to a value which makes the VCO output frequency N-times the reference frequency. The time constant of the loop is controlled by the loop filter. There are many tradeoffs in designing a PLL, such a phase noise, tuning speed, frequency resolution, etc., and there are many good references on the subject (see References 1, 2, and 3).

FREQUENCY SYNTHESIS USING OSCILLATORS AND PHASE-LOCKED LOOPS

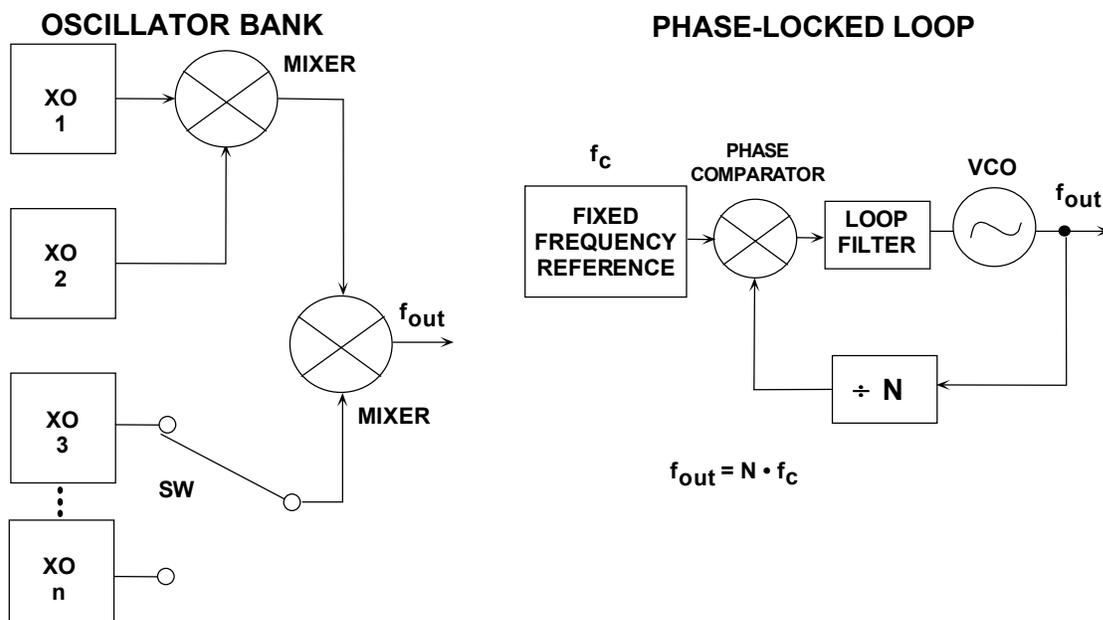


Figure 4.13

With the widespread use of digital techniques in instrumentation and communications systems, a digitally-controlled method of generating multiple frequencies from a reference frequency source has evolved called Direct Digital Synthesis (DDS). The basic architecture is shown in Figure 4.14. In this simplified

model, a stable clock drives a programmable-read-only-memory (PROM) which stores one or more integral number of cycles of a sinewave (or other arbitrary waveform, for that matter). As the address counter steps through each memory location, the corresponding digital amplitude of the signal at each location drives a DAC which in turn generates the analog output signal. The spectral purity of the final analog output signal is determined primarily by the DAC. The phase noise is basically that of the reference clock.

FUNDAMENTAL DIRECT DIGITAL SYNTHESIS SYSTEM

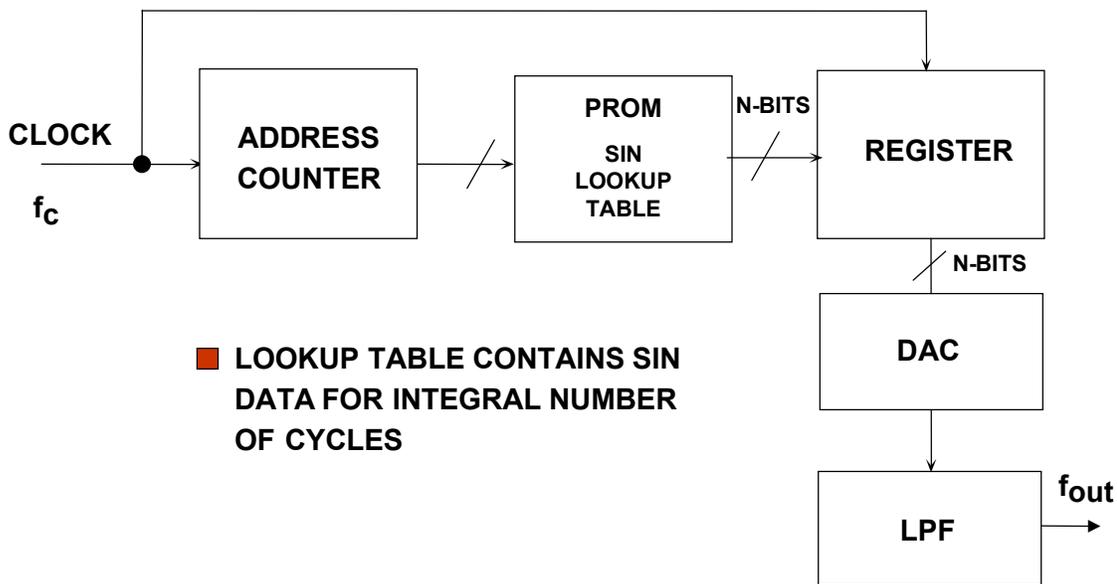


Figure 4.14

The DDS system differs from the PLL in several ways. Because a DDS system is a sampled data system, all the issues involved in sampling must be considered: quantization noise, aliasing, filtering, etc. For instance, the higher order harmonics of the DAC output frequencies fold back into the Nyquist bandwidth, making them unfilterable, whereas, the higher order harmonics of the output of PLL-based synthesizers can be filtered. There are other considerations which will be discussed shortly.

A fundamental problem with this simple DDS system is that the final output frequency can be changed only by changing the reference clock frequency or by reprogramming the PROM, making it rather inflexible. A practical DDS system implements this basic function in a much more flexible and efficient manner using digital hardware called a Numerically Controlled Oscillator (NCO). A block diagram of such a system is shown in Figure 4.15.

A FLEXIBLE DDS SYSTEM

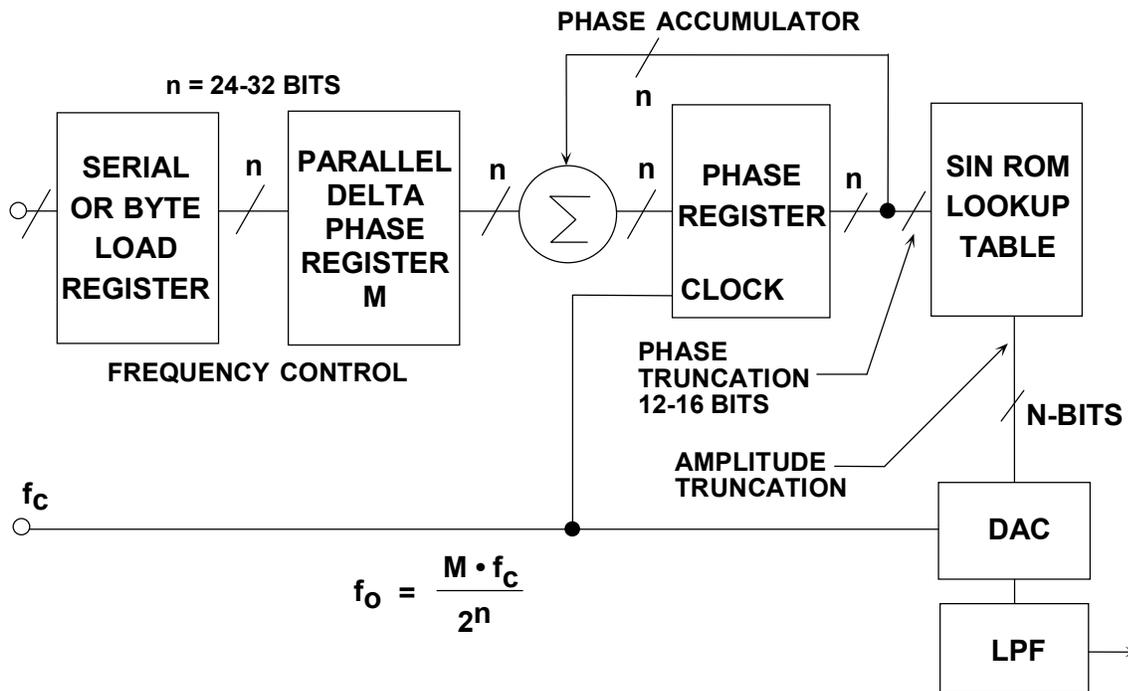


Figure 4.15

The heart of the system is the *phase accumulator* whose contents is updated once each clock cycle. Each time the phase accumulator is updated, the digital number, M , stored in the *delta phase register* is added to the number in the phase accumulator register. Assume that the number in the delta phase register is $00\dots01$ and that the initial contents of the phase accumulator is $00\dots00$. The phase accumulator is updated by $00\dots01$ on each clock cycle. If the accumulator is 32-bits wide, 2^{32} clock cycles (over 4 billion) are required before the phase accumulator returns to $00\dots00$, and the cycle repeats.

The truncated output of the phase accumulator serves as the address to a sine (or cosine) lookup table. Each address in the lookup table corresponds to a phase point on the sinewave from 0° to 360° . The lookup table contains the corresponding digital amplitude information for one complete cycle of a sinewave. (Actually, only data for 90° is required because the quadrature data is contained in the two MSBs). The lookup table therefore maps the phase information from the phase accumulator into a digital amplitude word, which in turn drives the DAC.

Consider the case for $n=32$, and $M=1$. The phase accumulator steps through each of 2^{32} possible outputs before it overflows. The corresponding output sinewave frequency is equal to the clock frequency divided by 2^{32} . If $M=2$, then the phase accumulator register "rolls over" twice as fast, and the output frequency is doubled. This can be generalized as follows.

For an n-bit phase accumulator (n generally ranges from 24 to 32 in most DDS systems), there are 2^n possible phase points. The digital word in the delta phase register, M, represents the amount the phase accumulator is incremented each clock cycle. If f_c is the clock frequency, then the frequency of the output sinewave is equal to:

$$f_o = \frac{M \cdot f_c}{2^n}.$$

This equation is known as the DDS "tuning equation." Note that the frequency resolution of the system is equal to $f_c/2^n$. For n=32, the resolution is greater than one part in four billion! In a practical DDS system, all the bits out of the phase accumulator are not passed on to the lookup table, but are truncated, leaving only the first 12 to 16 MSBs. This reduces the size of the lookup table and does not affect the frequency resolution. The phase truncation only adds a small but acceptable amount of phase noise to the final output; the bulk of the output distortion comes from the DAC itself.

The basic DDS system described above is extremely flexible and has high resolution. The frequency can be changed instantaneously with no phase discontinuity by simply changing the contents of the M-register. However, practical DDS systems first require the execution of a serial, or byte-loading sequence to get the new frequency word into an internal buffer register which precedes the parallel-output M-register. This is done to minimize package pin count. After the new word is loaded into the buffer register, the parallel-output delta phase register is clocked, thereby changing all the bits simultaneously. The number of clock cycles required to load the delta-phase buffer register determines the maximum rate at which the output frequency can be changed.

The AD9850 125MSPS DDS system (Figure 4.16) uses a 32-bit phase accumulator which is truncated to 14-bits (MSBs) before being passed to the lookup table. The final digital output is 10-bits to the internal DAC. The AD9850 allows the output phase to be modulated using an additional register and an adder placed between the output of the phase accumulator register and the input to the lookup table. The AD9850 uses a 5-bit word to control the phase which allows shifting the phase in increments of 180°, 90°, 45°, 22.5°, 11.25°, and any combination thereof. The device also contains an internal high speed comparator which can be configured to accept the (externally) filtered output of the DAC to generate a low-jitter output pulse suitable for driving the sampling clock input of an ADC. The full scale output current can be adjusted from 10 to 20mA using a single external resistor, and the output voltage compliance is +1V.

The frequency tuning (delta-phase register input word) and phase modulation words are loaded into the AD9850 via a parallel or serial loading format. The parallel load format consists of five consecutive loads of an 8-bit control word (byte). The first 8-bit byte controls phase modulation (5-bits), power-down enable (1-bit), and loading format (2-bits). Bytes 2-5 comprise the 32-bit frequency tuning word. The maximum control register update frequency is 23MHz. Serial loading of the AD9850 is accomplished via a 40-bit serial data stream on a single pin. Maximum update rate of the control register in the serial-load mode is 3MHz.

DACS FOR DSP APPLICATIONS

The AD9850 consumes only 380mW of power on a single +5V supply at a maximum 125MSPS clock rate. The device is available in a 28-pin surface mount SSOP (Shrink Small Outline Package).

AD9850 CMOS 125MSPS DDS/DAC SYNTHESIZER

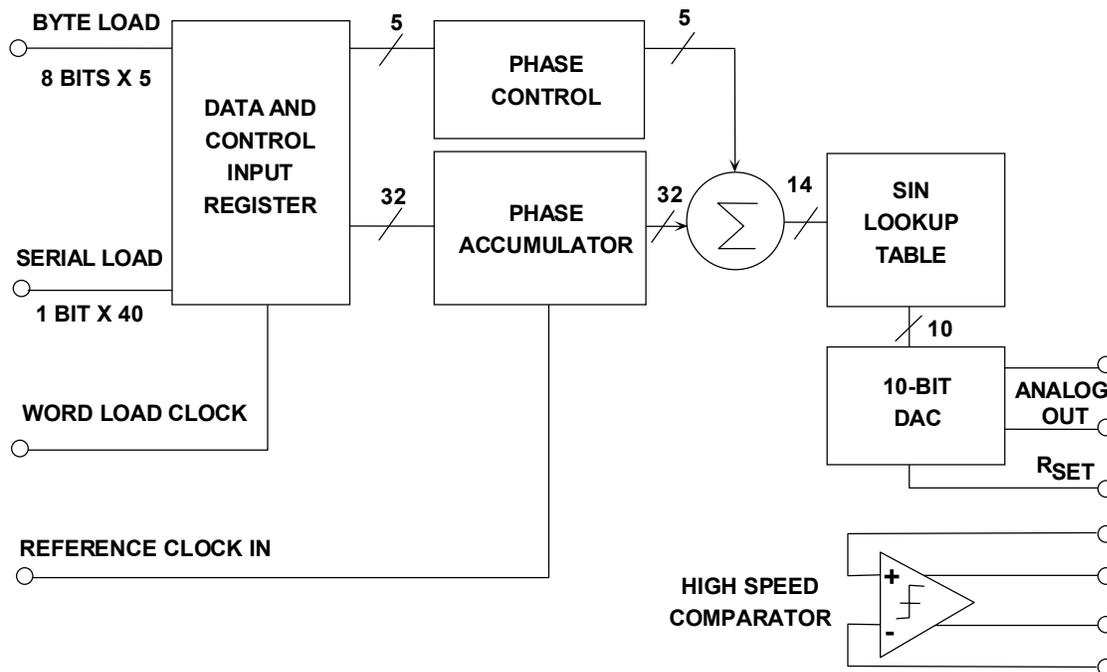


Figure 4.16

Analog Devices offers a number of DDS systems for a variety of applications. The AD983x family are low-cost 10-bit systems with clock frequencies up to 50MSPS. The AD985x-family offers 10-bit and 12-bit systems with clock frequencies up to 300MSPS and additional functions such as quadrature and phase modulation, chirp-mode capability, and programmable on-chip reference clock multipliers.

REFERENCES

1. R.E. Best, **Phase-Locked Loops**, McGraw-Hill, New York, 1984.
2. F.M. Gardner, **Phaselock Techniques**, 2nd Edition, John Wiley, New York, 1979.
3. *Phase-Locked Loop Design Fundamentals*, Applications Note AN-535, Motorola, Inc.
4. **The ARRL Handbook for Radio Amateurs**, American Radio Relay League, Newington, CT, 1992.
5. Richard J. Kerr and Lindsay A. Weaver, *Pseudorandom Dither for Frequency Synthesis Noise*, United States Patent Number 4,901,265, February 13, 1990.
6. Henry T. Nicholas, III and Henry Samuelli, *An Analysis of the Output Spectrum of Direct Digital Frequency Synthesizers in the Presence of Phase-Accumulator Truncation*, **IEEE 41st Annual Frequency Control Symposium Digest of Papers**, 1987, pp. 495-502, IEEE Publication No. CH2427-3/87/0000-495.
7. Henry T. Nicholas, III and Henry Samuelli, *The Optimization of Direct Digital Frequency Synthesizer Performance in the Presence of Finite Word Length Effects*, **IEEE 42nd Annual Frequency Control Symposium Digest of Papers**, 1988, pp. 357-363, IEEE Publication No. CH2588-2/88/0000-357.

