

Considerations for Successive Approximation A→D Converters

Jim Williams

The most popular A→D method employed today is the successive approximation register (SAR) converter (see Box, "The Successive Approximation Technique"). Numerous monolithic, hybrid and modular devices embodying the successive approximation technique are available, and monolithic devices are slowly gaining in performance. Nevertheless, hybrid and modular SAR types feature the best performance. In particular, at the 12-bit level, the fastest monolithic devices currently available require about 10 μ s to convert. Modular and hybrid units achieve

conversion speeds below 2 μ s, although they are quite expensive. Because of these factors, it is often desirable to build, rather than buy, a high speed 12-bit SAR converter. Even in cases where high speed is not required, lower cost may still mandate building the circuit instead of using a monolithic device.

Figure 1 shows a simple 12-bit, 12 μ s SAR converter. Understanding this circuit's performance limitations is useful

LT, LT, LTC, LTM, Linear Technology and the Linear logo are registered trademarks of Linear Technology Corporation. All other trademarks are the property of their respective owners.

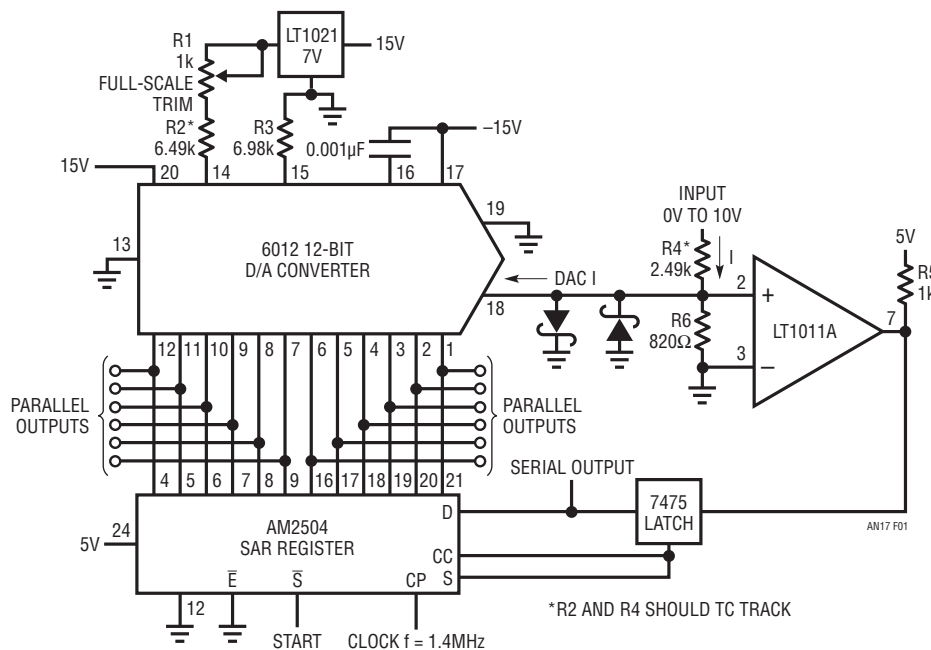


Figure 1. Basic 12-Bit, 12 μ s Successive Approximation A→D Converter

Application Note 17

in designing faster converters. Figure 2 shows waveforms of operation. Trace A is the clock, which is applied to the 2504 IC successive approximation register (SAR), while Trace B is the start pulse. On the rising edge of the start pulse, the SAR-DAC combination begins to test each bit, beginning with the MSB. This action is reflected in conditions at the LT1011's positive input (Trace C). This waveform is seen to sequentially converge towards zero as the SAR, DAC and comparator servo the node. After the LSB has been converted the "conversion complete" (CC) line (Trace D) goes high, signaling the end of the sequence. The 7475 latch prevents the comparator from responding to input noise or shifts after the conversion is complete. It is reset at the next "conversion command". The major limitations on speed in this circuit are the DAC and the comparator. Most bipolar DACs require 150ns to 200ns to settle for a worst-case (full-scale) step and the comparator's delay time must also be accounted for. The clamp diodes limit overdrive, aiding comparator response. Additionally, the 820Ω resistor to ground shunts the DACs output capacitance, helping the comparator-DAC node settle more quickly. The shunt degrades the voltage per-LSB available to the comparator, but the LT1011's high gain makes up for this.

In general, this is a fairly typical 12-bit SAR converter with good speed and low cost. To get higher conversion speed requires more sophisticated circuitry.

Figure 3 shows a circuit, which uses a clock modulation scheme to decrease conversion time. The A→D is identical to Figure 1's circuit, but the clock terminal (CP) is driven by a 2-speed oscillator. Figure 4 shows operating details. A convert command pulse (Trace A) initiates the SAR routine. Simultaneously, the 7474 flip-flop's Q output is set high (Trace C), biasing Q1. This causes the 47pF capacitor to be paralleled with the 33pF unit. These capacitors are part of the timing network of C1, which is configured as an oscillator. C1's output pulses (Trace B) drive the SAR's clock terminal (CP in the schematic). After the third MSB has been converted, the flip-flop is reset (Trace C). Q1 goes off and the clock oscillator (Trace B) speeds up. The increase in clock speed results in less dwell time per bit at the DAC-comparator junction (Trace D), allowing faster total conversion time. Trace E, the conversion complete pulse (CC), drops low 7.5μs after the conversion started.

This clock modulation approach buys significantly improved speed, but does nothing to get around the comparator's contribution to delay. Minimizing comparator delay would seem to be as simple as using a faster device. At the 8-bit or even 10-bit level this usually works, but 12-bit performance raises problems. Replacing the LT1011, a 150ns device, with a 10ns LT1016 increases speed, but decreases available gain. The LT1011 has a minimum gain of 200,000. The LT1016's high speed sacrifices gain.




Figure 2. 12μs A→D Waveforms

Application Note 17

This figure clearly means the comparator cannot do the job without some help. The input signal reduction caused by the shunt resistor at the DAC output worsens the problem. Finally, the comparator's speed degrades for such low level overdrives.

The solution to the aforementioned difficulties is to place a gain stage ahead of the comparator. While the gain stage adds some delay, it also increases gain, providing the needed overdrive to the comparator.

Figure 5 shows a simple pre-amplifier. This pre-amp-comparator combination gives adequate gain and an overall response time of 40ns to 50ns. A1 is set up as a Schottky bounded amplifier. The bound diodes prevent A1 from saturating due to excessive summing point overdrive, aiding response time. The 10pF capacitor, a typical value,

compensates DAC output capacitance and is selected for best amplifier damping. The 10k feedback resistor, also typical, is chosen for best gain-bandwidth performance. Voltage gains of 4 to 10 are common. Figure 6 shows performance. Trace A, a test input pulse, causes A1's output (Trace B) to slew through zero (screen center horizontal line). When A1 crosses zero, C1's input biases negative and it responds (C1's output is Trace C) 10ns later with a TTL output.

A simple circuit like this results in faster comparisons. Substituted for the LT1011 in Figure 3's circuit, it permits conversion times in the 3µs to 5µs range. Further reduction in conversion time is possible with a faster discrete preamplifier.




Figure 5. Simple Pre-Amplifier for the Comparator




Figure 6. Pre-Amplified Comparator Waveforms

Figure 7 shows a very fast pre-amplifier built with GHz range transistors. This cascoded differential amplifier is placed ahead of C1, an LT1016. Q4 and Q5 provide bias current compensation for Q1's base current. Figure 8

shows results for a test input signal (Trace A). C1's output (Trace B) switches in 15ns to 20ns. About 10ns of this delay is due to C1, with the pre-amplifier contributing the rest.

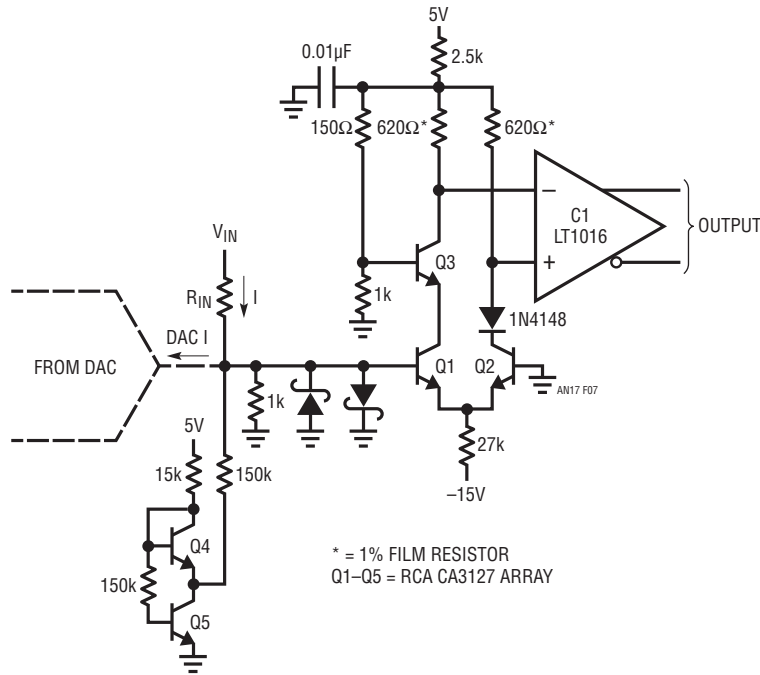


Figure 7. Fast Pre-Amplifier Comparator




Figure 8. Fast Pre-Amplifier Comparator Waveforms

Application Note 17

Figure 9 shows the discrete pre-amplifier used in a very fast 12-bit SAR converter. The design utilizes a variety of techniques to attain extremely high speed. Primary speed enhancing features include a closed loop clock control

method and active summing node clamping. The circuit achieves a full 12-bit conversion in 1.8 μ s, about the practical limit with off-the-shelf components.

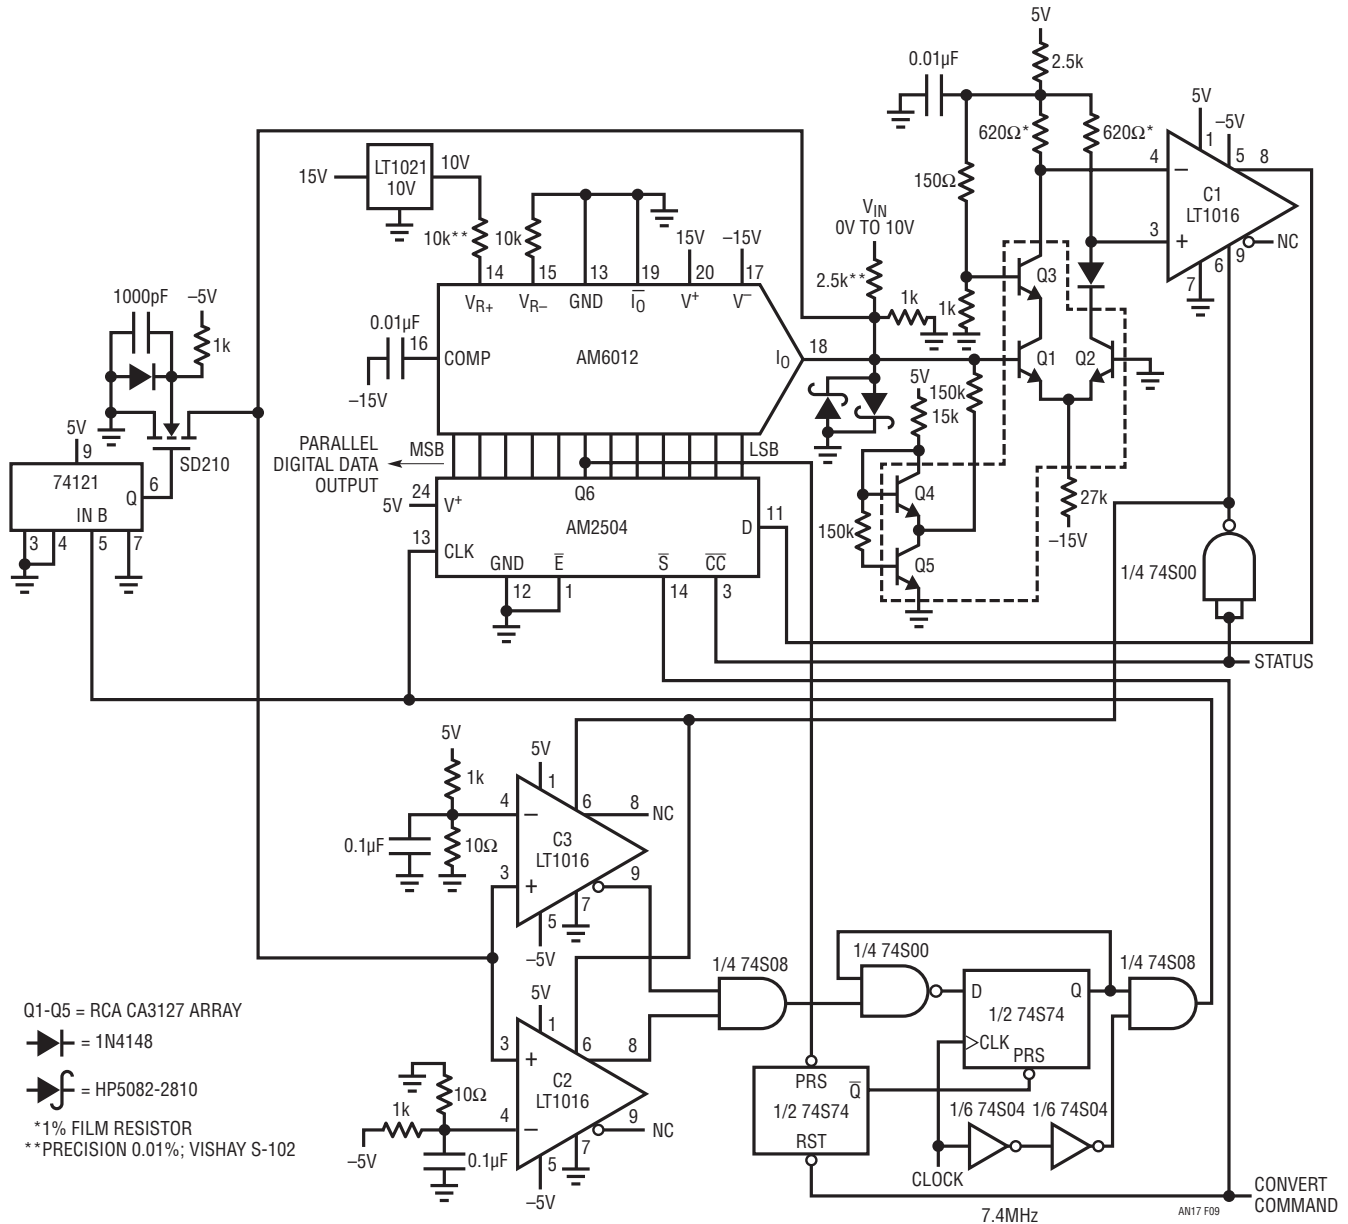


Figure 9. 12-Bit 1.8 μ s SAR A→D

The design is similar in concept to Figure 3, except that the fast pre-amplifier replaces the LT1011. Additionally, the clock speed change is implemented with the digital logic shown. Unlike before, the clock rate is accelerated after the fifth MSB is converted. During conversion of the upper four bits, the clock rate is controlled by a closed loop to maximize overall speed. The loop monitors conditions at the DAC-comparator summing node. If the node is outside $\pm 50\text{mV}$, the SAR is clocked at the maximum rate. For node responses inside $\pm 50\text{mV}$, the clock rate is retarded, giving adequate time for settling. The clock loop speeds conversion by not waiting for bits which aren't going to settle within $\pm 50\text{mV}$. C2 and C3 form a high speed window comparator, which delivers summing node information in digital form to the clock logic.

Figure 10 shows the effects of the closed loop clocking scheme. Trace A is the convert command. Trace B is the gated output of the C2-C3 window comparator. Trace B's state controls the clock line, which is Trace C. Trace D is the summing point, and the dwell time-per-bit is controlled by the window comparator's decision. Beyond the fifth bit, the SAR's Q6 line instructs the clock logic to run at maximum speed. As described to this point, the circuit achieves a $1.9\mu\text{s}$ conversion time.

If the 74121 one-shot and associated circuitry are included, conversion time is reduced to $1.8\mu\text{s}$. These components form an active clamp at the DAC-comparator summing

node. Each time the SAR clock is pulsed (Trace A, Figure 11), the 74121 puts out a 30ns FET gate pulse (Trace B). The FET comes on, shunting the summing node (Trace C) to ground. The FET's low on-resistance aids DAC settling by discharging the DACs 30pF output capacitance for 30ns. The summing node (Trace C) is reset to zero by this action at each SAR-directed step. When the one shot times out, the node settles to its final value. This active clamping results in about a 10ns-per-bit time savings.

The circuit's $1.8\mu\text{s}$ conversion time is very close to what is practically achievable for a 12-bit SAR A→D converter. The special techniques used result in an effective DAC settling time of about 100ns per bit. Comparator-pre-amp delay is about 20ns per bit and SAR chip delays are in the 25ns per bit range. Adding these together gives:

$$100\text{ns} + 20\text{ns} + 25\text{ns} \cdot 12 = 1.74\mu\text{s}.$$

The comparator and SAR delays, about 31% of the total, are not easily reduced. A discrete Schottky SAR design and a faster pre-amp can cut this figure somewhat, but the DAC settling time, 69% of the total, remains. The effective 100ns/bit DAC settling time compares favorably with published specifications for monolithic DACs, and is not readily reducible. Beyond this speed, other conversion methods are required.




Figure 10. Figure 9's Waveforms




Figure 11. Figure 9's Waveforms Using Active Clamping

The Successive Approximation Technique

The successive approximation technique is probably as old as the first crude weighing scale ever constructed. It is most easily visualized when considering the operation of a beam balance. The unknown weight, in one pan, is determined by successive trials with standard weights placed in the other pan. Overweight-underweight decisions are made by the balance as standard weights (and combinations of them) are successively tried in a logical sequence, which converges towards balancing the scale.

Successive approximation A→D converters start with the MSB and proceed toward the LSB as each under-over decision is made. The figure shows the summing node response (Trace A) as the DAC, instructed by the clock driven (Trace B) successive approximation register (SAR) logic, tries different bit weights. The comparator's decisions are also shown (Trace C). Note how the summing point sequentially converges towards zero, the analog of null in a beam balance.


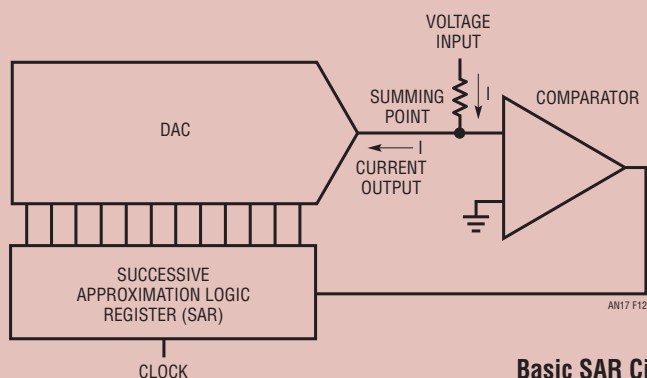
Digital to Analog Converters in SAR Applications

Selecting a DAC for use in a SAR-based A→D requires some thought. Most often, bipolar current mode DACs are employed because of their higher speed. CMOS DACs output capacitance, in the 100pF to 150pF range, causes excessive summing node settling times. Monolithic bipolar types, in the 30pF region, settle more quickly. Voltage mode output DACs are almost never used because they are not necessary to achieve summing action and they are substantially slower than current output types.

Speed is often important, and since the DAC is the slowest part of the converter, it should be carefully considered. Settling time specifications for DACs are usually stated for full-scale transitions. Smaller bit changes take less

time, so some interpretation of the full-scale settling time number can be made when considering the DAC's effective settling time-per-bit in an A→D application. Unfortunately, the complex dynamics of DAC internals prevent simple straight line calculations (e.g., 1 LSB will not settle in one-half the time of full-scale for a 12-bit unit). At moderate speeds, the simplest course is to allow the specified full-scale settling time for each bit decision. The conservative method will never get you into trouble, but almost certainly guarantees slower than necessary DAC performance. The best way to find out just how far you can push the DACs settling time specification in a SAR application is to consult the manufacturer. Additionally, it is worthwhile to actually measure the settling time under conditions appropriate to the intended use (see LTC Application Note 10, "Methods for Measuring Op Amp Settling Time", for circuits readily adaptable to DAC settling time measurements). The wide variety of DACs and individual output termination requirements make obtainable results vary considerably. However, some guidelines on what to expect are possible. For example, the popular 565A type, specified at 250ns full-scale settling time into 0Ω, can achieve 110ns to 150ns effective settling time-per-bit in SAR applications with careful design. It is also worth noting that the dynamics of DAC types can vary considerably between manufacturers of what is nominally the same part.

Speed is not the only concern. The DACs DC specifications translate directly into A→D error terms. Linearity, drift, accuracy and other DC terms contribute on a 1:1 basis to the A→D's error characteristics. One specification, monotonicity, can contribute a particularly nasty term. The effect of a non-monotonic DAC is an inability of the A→D to produce some output codes (missing codes) under any input condition.



Basic SAR Circuit and Waveforms