High Speed Comparator Techniques

Jim Williams

INTRODUCTION

Comparators may be the most underrated and underutilized monolithic linear component. This is unfortunate because comparators are one of the most flexible and universally applicable components available. In large measure the lack of recognition is due to the IC op amp, whose versatility allows it to dominate the analog design world. Comparators are frequently perceived as devices which crudely express analog signals in digital form—a 1-bit A–D converter. Strictly speaking, this viewpoint is correct. It is also wastefully constrictive in its outlook. Comparators don't "just compare" in the same way that op amps don't "just amplify".

Comparators, in particular high speed comparators, can be used to implement linear circuit functions which are as sophisticated as any op amp-based circuit. Judiciously combining a fast comparator with op amps is a key to achieving high performance results. In general, op amp-based circuits capitalize on their ability to close a feedback loop with precision. Ideally, such loops are maintained continuously over time. Conversely, comparator circuits are often based on speed and have a discontinuous output over time. While each approach has its merits, a fusion of both yields the best circuits.

This effort's initial sections are devoted to familiarizing the reader with the realities and difficulties of high speed comparator circuit work. The mechanics and subtleties of achieving precision circuit operation at DC and low frequency have been well documented. Relatively little has appeared which discusses, in practical terms, how to get fast circuitry to work. In developing such circuits, even the most veteran designers sometimes feel that nature is conspiring against them. In some measure this is true. Like all engineering endeavors, high speed circuits can only work if negotiated compromises with nature are arranged. Ignorance of, or contempt for, physical law is a direct route to frustration. In this regard, much of the text and appendices are directed at developing awareness of and respect for circuit parasitics and fundamental limitations. This approach is maintained in the applications section, where the notion of "negotiated compromises" is expressed in terms of resistor values and compensation techniques. Many of the application circuits use the LT1016's speed to improve on a standard circuit. Some utilize the speed to implement a traditional function in a non-traditional way, with attendant advantages. A (very) few operate at or near the state-of-the-art for a given circuit type, regardless of approach. Substantial effort has been expended in developing these examples and documenting their operation. The resultant level of detail is justified in the hope that it will be catalytic. The circuits should stimulate new ideas to suit particular needs, while demonstrating the LT1016's capabilities in an instructive manner.

TABLE OF CONTENTS

THE LT1016 — AN OVERVIEW

A new ultra high speed comparator, the LT1016, features TTL-compatible complementary outputs and 10ns response time. Other capabilities include a latch pin and good DC input characteristics (see Figure 1). The LT1016's outputs directly drive all TTL families, including the new higher speed ASTTL and FAST parts. Additionally, TTL outputs make the device easier to use in linear circuit applications where ECL output levels are often inconvenient.

A substantial amount of design effort has made the LT1016 relatively easy to use. It is much less prone to oscillation and other vagaries than some slower comparators, even with slow input signals. In particular, the LT1016 is stable in its linear region, a feature no other high speed comparator has. Additionally, output stage switching does not appreciably change power supply current, further enhancing stability. These features make the application of the 200GHz gain-bandwidth LT1016 considerably easier than other fast comparators. Unfortunately, laws of physics dictate that the circuit *environment* the LT1016 works in must be properly prepared. The performance limits of high speed circuitry are often determined by parasitics such as stray capacitance, ground impedance, and layout. Some of these considerations are present in digital systems where designers are comfortable describing bit patterns and memory access times in terms of nanoseconds. The LT1016 can be used in such fast digital systems and Figure 2 shows just how fast the device is. The simple test circuit allows us to see that the LT1016's (Trace B) response to the pulse generator (Trace A) is faster than a TTL inverter (Trace C)! In fact, the inverter's output never gets to a TTL "0" level. Linear circuits operating with this kind of speed make many engineers justifiably wary. Nanosecond domain linear circuits are widely associated with oscillations, mysterious shifts in circuit characteristics, unintended modes of operation and outright failure to function.

Other common problems include different measurement results using various pieces of test equipment, inability to make measurement connections to the circuit without inducing spurious responses and dissimilar operation between two "identical" circuits. If the components used

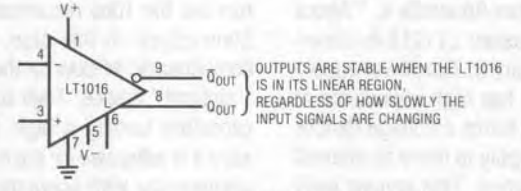

Figure 1. The LT1016 at a Glance

Figure 2. LT1016 vs a TTL Gate

in the circuit are good and the design is sound, all of the above problems can usually be traced to failure to provide a proper circuit "environment." To learn how to do this requires studying the causes of the aforementioned difficulties.

The Rogue's Gallery of High Speed Comparator Problems

By far the most common error involves power supply bypassing. Bypassing is necessary to maintain low supply impedance. DC resistance and inductance in supply wires and PC traces can quickly build up to unacceptable levels. This allows the supply line to move as internal current levels of the devices connected to it change. This will almost always cause unruly operation. In addition, several devices connected to an unbypassed supply can "communicate" through the finite supply impedances, causing erratic modes. Bypass capacitors furnish a simple way to eliminate this problem by providing a local reservoir of energy at the device. The bypass capacitor acts like an electrical flywheel to keep supply impedance low at high frequencies. The choice of what type of capacitors to use for bypassing is a critical issue and should be approached carefully (see Appendix A, "About Bypass Capacitors"). An unbypassed LT1016 is shown responding to a pulse input in Figure 3. The power supply the LT1016 sees at its terminals has high impedance at high frequency. This impedance forms a voltage divider with the LT1016, allowing the supply to move as internal conditions in the comparator change. This causes local feedback and oscillation occurs. Although the LT1016 responds to the input pulse, its output is a blur of 100MHz oscillation. *Always use bypass capacitors.*

In Figure 4 the LT1016's supplies are bypassed, but it still oscillates. In this case, the bypass units are either too far from the device or are lossy capacitors. *Use capacitors with good high frequency characteristics and mount them as close as possible to the LT1016. An inch of wire between the capacitor and the LT1016 can cause problems.*

In Figure 5 the device is properly bypassed but a new problem pops up. This photo shows both outputs of the comparator. Trace A appears normal, but Trace B shows an excursion of almost 8V — quite a trick for a device running from a +5V supply. This is a commonly reported problem in high speed circuits and can be quite confusing. It is not due to suspension of natural law, but is traceable to a grossly miscompensated or improperly selected oscilloscope probe. *Use probes which match your oscilloscope's input characteristics and compensate them properly* (for a discussion on probes, see Appendix B, "About Probes and Scopes"). Figure 6 shows another probe-induced problem. Here, the amplitude seems correct but the 10ns response time LT1016 appears to have 50ns edges! In this case, the probe used is too heavily compensated or slow for the oscilloscope. Never use 1X or "straight" probes. Their bandwidth is 20MHz or less and capacitive loading is high. *Check probe bandwidth to ensure it is adequate for the measurement. Similarly, use an oscilloscope with adequate bandwidth.*

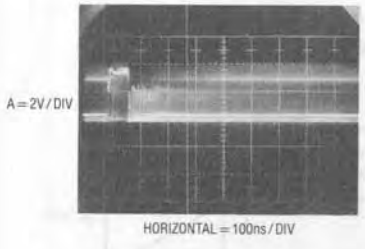

Figure 3. Unbypassed LT1016 Response

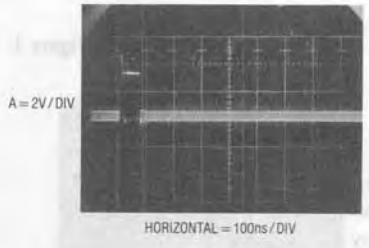

Figure 4. LT1016 Response with Poor Bypassing

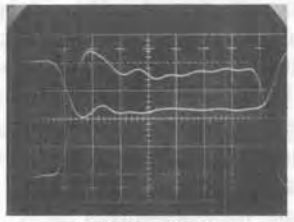

Figure 5. Improper Probe Compensation Causes Seemingly Unexplainable Amplitude Error

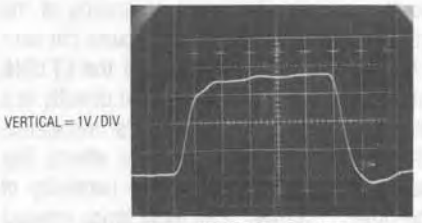

Figure 6. Overcompensated or Slow Probes Make Edges Look Too Slow

In Figure 7 the probes are properly selected and applied but the LT1016's output rings and distorts badly. In this case, the probe ground lead is too long. For general purpose work most probes come with ground leads about 6 inches long. At low frequencies this is fine. At high speed, the long ground lead looks inductive, causing the ringing shown. High quality probes are always supplied with some short ground straps to deal with this problem. Some come with very short spring clips which fix directly to the probe tip to facilitate a low impedance ground connection. For fast work, the ground connection to the probe should not exceed 1 inch in length. *Keep the probe ground connection as short as possible.*

The difficulty in Figure 8 is delay and inadequate amplitude (Trace B). A small delay on the leading edge is followed by a large delay before the falling edge begins. Additionally, a lengthy, tailing response stretches 70ns before finally settling out. The amplitude only rises to 1.5V. A common oversight is responsible for these conditions.

A FET probe monitors the LT1016 output in this example. The probe's common-mode input range has been exceeded, causing it to overload and clip the output badly. The small delay on the rising edge is characteristic of active probes and is legitimate. During the time the output is high, the probe is driven deeply into saturation. When the output falls, the probe's overload recovery is lengthy and uneven, causing the delay and tailing.

Know your FET probe. Account for the delay of its active circuitry. Avoid saturation effects due to common-mode input limitations (typically ± 1V). Use 10X and 100X attenuator heads when required.

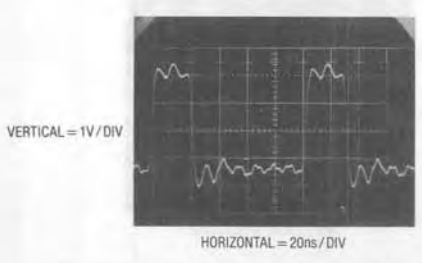

Figure 7. Typical Results Due to Poor Probe Grounding

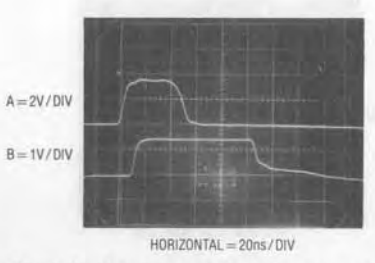

Figure 8. Overdriven FET Probe Causes Delayed, Tailing Response

Figure 9 shows the LT1016's output (Trace B) oscillating near 40MHz as it responds to an input (Trace A). Note that the input signal shows artifacts of the oscillation. This example is caused by improper grounding of the comparator. In this case, the LT1016's ground pin connection is 1 inch long. The ground lead of the LT1016 must be as short as possible and connected directly to a low impedance ground point. Any substantial impedance in the LT1016's ground path will generate effects like this. The reason for this is related to the necessity of bypassing the power supplies. The inductance created by a long device ground lead permits mixing of ground currents, causing undesired effects in the device. The solution here is simple. *Keep the LT1016's ground pin connection as short (typically 1/4 inch) as possible and run it directly to a low impedance ground. Do not use sockets.*

Figure 10 addresses the issue of the "low impedance ground," referred to previously. In this example, the output is clean except for chattering around the edges. This photograph was generated by running the LT1016 without a "ground plane." A ground plane is formed by using a continuous conductive plane over the surface of the circuit board (the theory behind ground planes is discussed in Appendix C). The only breaks in this plane are for the circuit's necessary current paths. The ground plane serves two functions. Because it is flat (AC currents travel along the surface of a conductor) and covers the entire area of the board, it provides a way to access a low inductance ground from anywhere on the board. Also, it minimizes the effects of stray capacitance in the circuit by referring them to ground. This breaks up potential unintended and harmful feedback paths. *Always use a ground plane with the LT1016.*

"Fuzz" on the edges is the difficulty in Figure 11. This condition appears similar to Figure 10, but the oscillation is more stubborn and persists well after the output has gone low. This condition is due to stray capacitive feedback from the outputs to the inputs. A 3kΩ input source impedance and 3pF of stray feedback allowed this oscillation. The solution for this condition is not too difficult. *Keep source impedances as low as possible, preferably 1kΩ or less. Route output and input pins and components away from each other.*

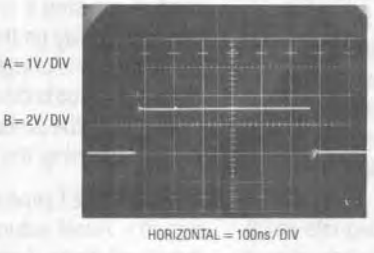

Figure 9. Excessive LT1016 Ground Path Resistance Causes Oscillation

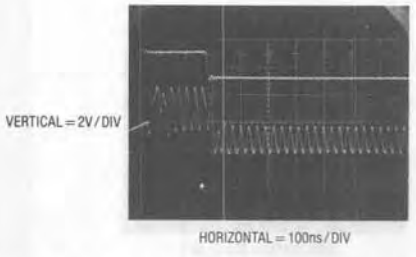

Figure 10. Transition Instabilities Due to No Ground Plane

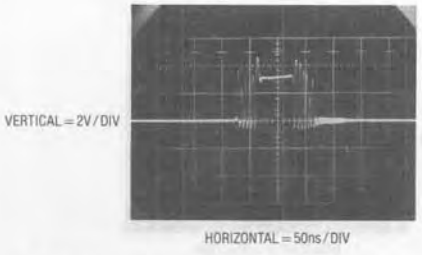

Figure 11. 3pF Stray Capacitive Feedback with 3kΩ Source Can Cause Oscillation

The opposite of stray-caused oscillations appears in Figure 12. Here, the output response (Trace B) badly lags the input (Trace A). This is due to some combination of high source impedance and stray capacitance to ground at the input. The resulting RC forces a lagged response at the input and output delay occurs. An RC combination of 2kΩ source resistance and 10pF to ground gives a 20ns time constant — significantly longer than the LT1016's response time. *Keep source impedances low and minimize stray input capacitance to ground.*

Figure 13 shows another capacitance-related problem. Here the output does not oscillate, but the transitions are discontinuous and relatively slow. The villain of this situation is a large output load capacitance. This could be caused by cable driving, excessive output lead length or the input characteristics of the circuit being driven. In most situations this is undesirable and may be eliminated by buffering heavy capacitive loads. In a few circumstances it may not affect overall circuit operation and is tolerable. *Consider the comparator's output load characteristics and their potential effect on the circuit. If necessary, buffer the load.*

Another output-caused fault is shown in Figure 14. The output transitions are initially correct but end in a ringing condition. The key to the solution here is the ringing. What is happening is caused by an output lead which is too long. The output lead looks like an unterminated transmission line at high frequencies and reflections occur. This accounts for the abrupt reversal of direction on the leading edge and the ringing. If the comparator is driving TTL this may be acceptable, but other loads may not tolerate it. In this instance, the direction reversal on the leading edge might cause trouble in a fast TTL load. *Keep output lead lengths short. If they get much longer than a few inches, terminate with a resistor (typically 250Ω–400Ω).*

A final malady is presented in Figure 15. These waveforms are reminiscent of the input RC-induced delay of Figure 12. The output waveform initially responds to the input's leading edge, but then returns to zero before going high again. When it does go high, it slews slowly. Additional odd characteristics include pronounced overshoot and pulse top aberration. The fall time is also slow and well delayed from the input. This is certainly strange

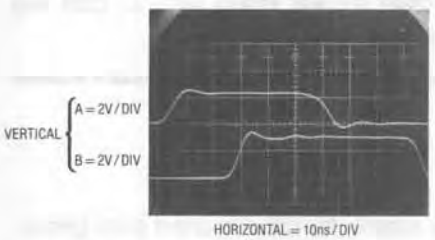

Figure 12. Stray 5pF Capacitance from Input to Ground Causes Delay

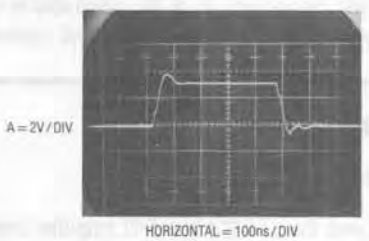

Figure 13. Excessive Load Capacitance Forces Edge Distortion

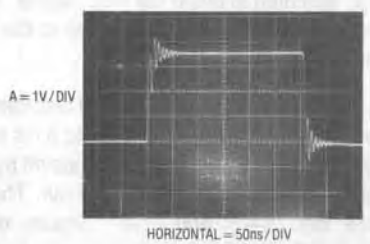

Figure 14. Lengthy, Unterminated Output Lines Ring from Reflections

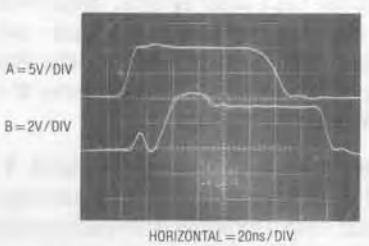

Figure 15. Input Common-Mode Overdrive Generates Odd Outputs

behavior from a TTL output. What is going on here? The input pulse is responsible for all these anomalies. Its 10V amplitude is well outside the +5V powered LT1016's common-mode input range. Internal input clamps prevent this pulse from damaging the LT1016, but an overdrive of this magnitude results in poor response. *Keep input signals inside the LT1016's common-mode range at all times.*

Oscilloscopes

A few of the examples illustrated dealt with probe-caused problems. Although it should be obvious, it is worth mentioning that the choice of oscilloscope employed is crucial. Be certain of the characteristic of the probe-oscilloscope combination you are using. Rise time, bandwidth, resistive and capacitive loading, delay, overdrive recovery and other limitations must be kept in mind. High speed linear circuitry demands a great deal from test equipment and countless hours can be saved if the characteristics of the instruments used are well known (see Appendix C, "Measuring Equipment Response"). In fact, it is possible to use seemingly inadequate equipment to get good results if the equipment's limitations are well known and respected. All of the applications which follow involve rise times and delays well above the 100MHz–200MHz region, but 90% of the development work was done with a 50MHz oscilloscope. Familiarity with equipment and thoughtful measurement technique permit useful measurements seemingly beyond instrument specifications. A 50MHz oscilloscope cannot track a 5ns rise time pulse, but it can measure a 2ns delay between two such events. Using such techniques, it is often possible to deduce the desired information. There are situations where no amount of cleverness will work and the right equipment, e.g., a faster oscilloscope, must be used.

In general, use equipment you trust and measurement techniques you understand. Keep asking questions and don't be satisfied until everything you see on the oscilloscope is accounted for and makes sense.

The LT1016, combined with the precautionary notes listed above, permits fast linear circuit functions which are difficult or impossible using other approaches. Many of the applications presented represent the state-of-the-art for a particular circuit function. Some show new and improved ways to implement standard functions by utilizing the LT1016's speed. All have been carefully (and painfully) worked out and should serve as good idea sources for potential users of the device.

APPLICATIONS SECTION

1Hz-10MHz V→F Converter

The LT1016 and the LT1012 low drift amplifier combine to form a high speed V→F converter in Figure 16. A variety of circuit techniques is used to achieve a 1Hz to 10MHz output. Overrange to 12 MHz (V_{IN} = 12V) is provided. This circuit has a wider dynamic range (140dB, or 7 decades) than any commercially available unit. The 10MHz full-scale frequency is 10 times faster than currently available monolithic V→Fs. The theory of operation is based on the identity Q = CV.

Each time the circuit produces an output pulse, it feeds back a fixed quantity of charge (Q) to a summing node (Σ). The circuit's input furnishes a comparison current at the summing node. The difference signal at the node is integrated in a monitoring amplifier's feedback capacitor. The amplifier controls the circuit's output pulse generator, completing a feedback loop around the integrating amplifier. To maintain the summing node at zero, the pulse generator runs at a frequency which permits enough charge pumping to offset the input signal. Thus, the output frequency will be linearly related to the input voltage. A1 is the integrating amplifier.

For low bias, high speed operation, a pair of discrete FETs directly drives A1's output stages, replacing A1's monolithic input circuitry. A1's input stage is turned off by connecting the input pins to the negative 15V rail. The FET gates become the "+" and "−" inputs of the amplifier. 0.2μV/°C offset drift performance is obtained by stabilizing the A1-FET combination with A2, a precision op amp. A2 measures the DC value of the negative

input, compares it to ground, and forces the positive input to maintain offset balance in the A1-FET combination. Note that A2 is configured as an integrator and cannot see high frequency signals. It functions only at DC and low frequency. The A1-FET combination is arranged as an integrator with a 100pF feedback capacitor. When a positive voltage is applied to the input, A1's output integrates in a negative direction (Trace A, Figure 17). During this period, C1's inverting output is low. A very high speed level shifter, Q1–Q2 (see Appendix D, "About Level Shifters"), inverts this output and drives the zener reference bridge. The bridge's positive output is used to charge the 33pF capacitor. The 1.2V diode string provides cancellation and temperature compensation for the diode drops in the bridge so that the 33pF unit charges to $V_Z + V_{BE}$ Q3.

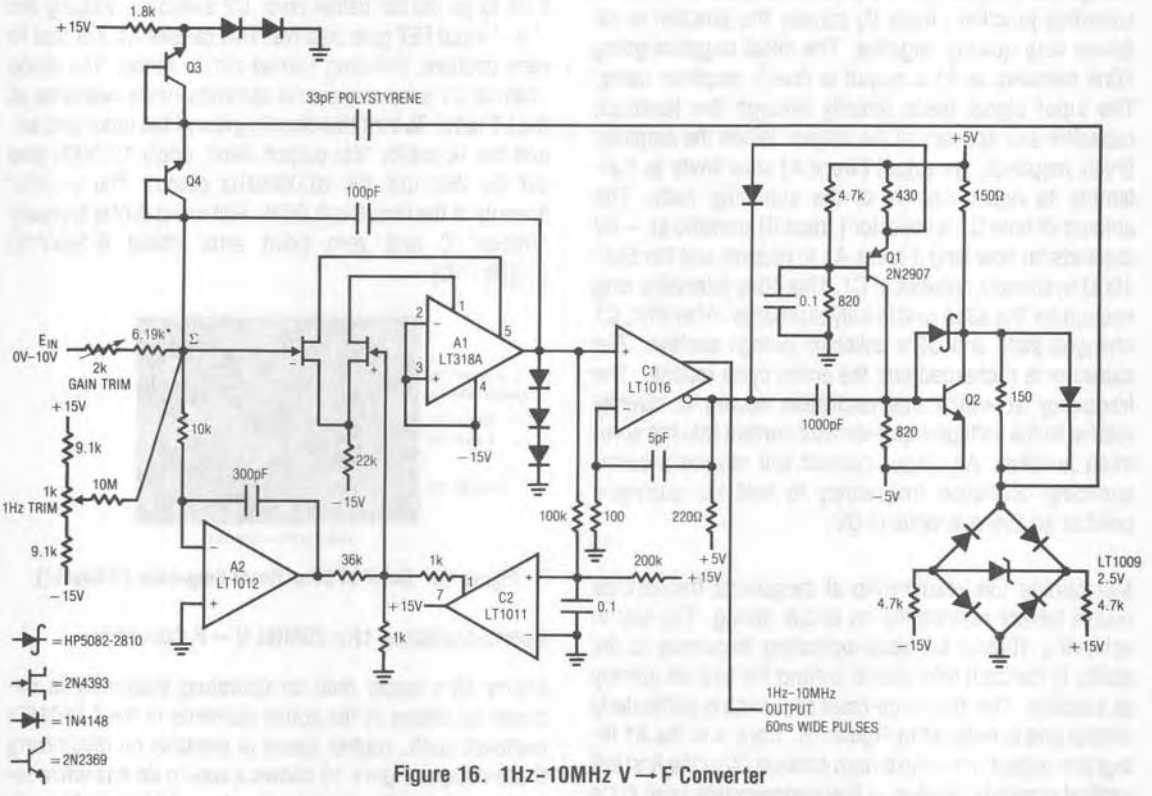

Figure 16. 1Hz–10MHz V → F Converter

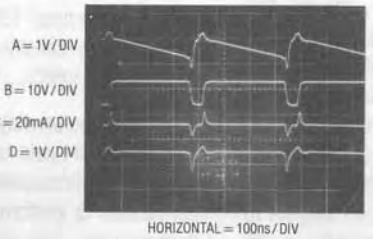

Figure 17. 10MHz V → Fs Operating Waveforms

When A1's output crosses zero, C1's inverting output goes high and Q2's (Trace B) collector goes to −5V. This causes the 33pF unit to dispense charge into the summing node via Q4's V_{BE}. The amount of charge dispensed is a direct function of the voltage that the 33pF unit was charged to (Q=CV). Q4's V_{BE} compensates the Q3 V_{BE} term in the capacitor's charge equation. The current which flows through the 33pF unit (Trace C) reflects this charge pumping action. The removal of current from A1's summing junction (Trace D) causes the junction to be driven very quickly negative. The initial negative-going 20ns transient at A1's output is due to amplifier delay. The input signal feeds directly through the feedback capacitor and appears at the output. When the amplifier finally responds, its output (Trace A) slew limits as it attempts to regain control of the summing node. The amount of time Q2's collector (Trace B) remains at −5V depends on how long it takes A1 to recover and the 5pF-100Ω hysteresis network at C1. This 60ns interval is long enough for the 33pF unit to fully discharge. After this, C1 changes state and Q2's collector swings positive. The capacitor is recharged and the entire cycle repeats. The frequency at which this oscillation occurs is directly related to the voltage-input-derived current into the summing junction. Any input current will require a corresponding oscillation frequenncy to hold the summing point at an average value of 0V.

Maintaining this relationship at megahertz frequencies places severe restrictions on circuit timing. The key to achieving 10MHz full-scale operating frequency is the ability to transmit information around the loop as quickly as possible. The discharge-reset sequence is particularly critical and is detailed in Figure 18. Trace A is the A1 integrator output. Its ramp output crosses 0V at the first left vertical graticule division. A few nanoseconds later, C1's inverting output begins to rise (Trace B), driving the Q1-Q2 level shifter output negative (Trace C). Q2's collector begins to head negative about 12ns after A1's output crosses 0V. 4ns later, the summing point (Trace D) begins to go negative as current is pulled from it through the 33pF capacitor. At 25ns, C1's inverting output is fully up, Q2's collector is at −5V, and the summing point has been pulled to its negative extreme. Now, A1 begins to take control. Its output (Trace A) slews rapidly in the positive direction, restoring the summing point. At 60ns, A1 is in control of the summing node and the integration ramp begins again.

Start-up and overdrive conditions could force A1's output to go to the negative rail and stay there. The AC-coupled nature of the charge dispensing loop can preclude normal operation and the circuit may latch. C2 provides a "watchdog" function for this condition. If A1's output tries to go too far below zero, C2 switches, forcing the "+" input FET gate positive. This causes A1's output to slew positive, initiating normal circuit action. The diode chain at C1's input prevents common-mode overdrive at the LT1016. To trim this circuit, ground the input and adjust the 1k pot for 1Hz output. Next, apply 10.000V and set the 2kΩ unit for 10.000MHz output. The transfer linearity of the circuit is 0.06%. Full-scale drift is typically 50ppm/°C and zero point error about 0.2μV/°C (0.2Hz/°C).

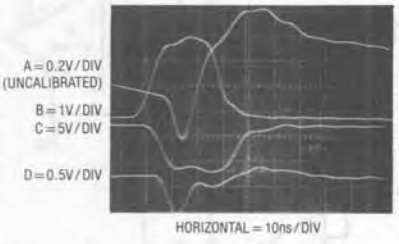

Figure 18. Detail of 60ns Reset Sequence (Whoosh!)

Quartz-Stabilized 1Hz-30MHz V→F Converter

Figure 16's upper limit on operating frequency is imposed by delays in the active elements in the LT1016's feedback path. Higher speed is possible by minimizing these delays. Figure 19 shows a way to do this while retaining good drift and linearity characteristics. The circuit's untrimmed 150dB dynamic range is 1000 times greater than commercially available V→F converters, whether monolithic, hybrid, or modular.

The technique employed allows the LT1016 to roar along at a 30MHz full-scale output frequency, substantially faster than any commercially available V→F. The actual V→F conversion is performed by the circuit shown inside the dashed lines. This circuit functions similarly to Figure 16.

The level shift and zener bridge are eliminated. Q1 charges the 200pF capacitor, which is unloaded by the Q2-Q3 buffer. When the LT1016's negative input rises above its positive input, its output goes low, pulling charge out of the capacitor via Q4, which serves as a low leakage diode. The 2.7pF capacitor provides positive feedback. If the left end of the 100k input resistor is driven from a voltage source, the LT1016 oscillates over a 1Hz to 30MHz range. Although this simple circuit is fast, its linearity is poor and drift exceeds 5000ppm/°C.

The remaining components in Figure 19 form a quartz-locked sampled-data loop to correct these terms without sacrificing speed. The loop works by counting the number of pulses at the LT1016's output during a fixed interval and converting this information to a voltage. The voltage is compared to the circuit's input by an amplifier which drives the LT1016 V→F circuit. This closed loop technique relies on the stability of the time interval and the digital-to-voltage conversion to achieve circuit stability. Frequent updating of the loop ensures long term stability. Figure 20 shows how the circuit functions.

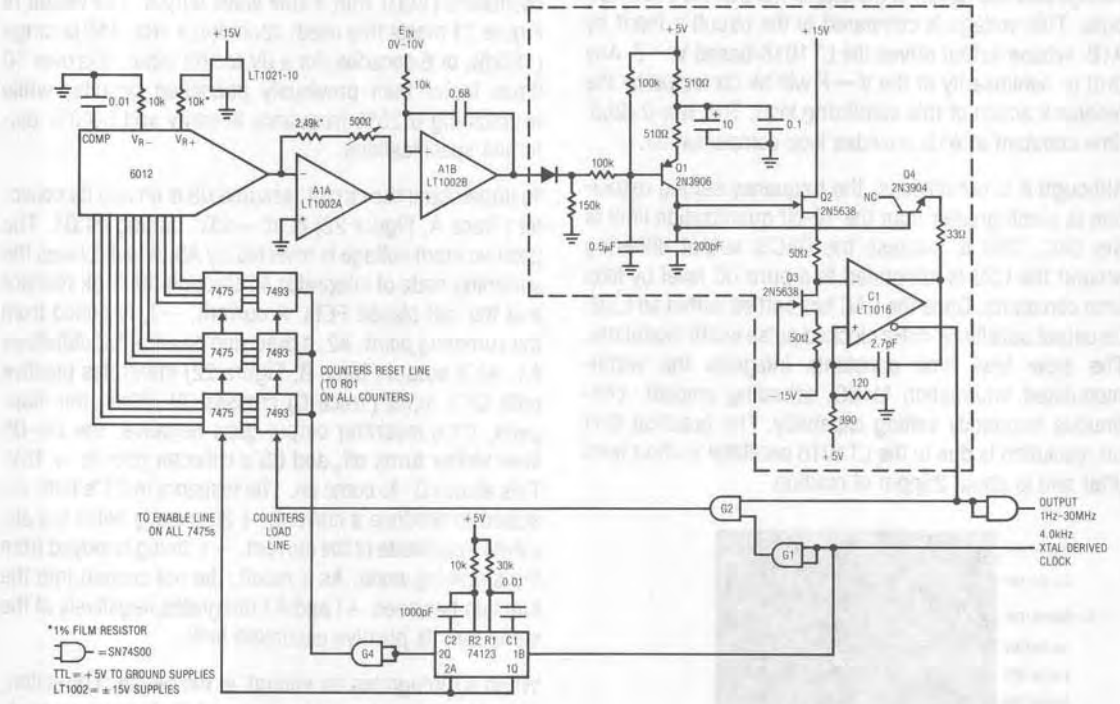

Figure 19. 30MHz V→F Utilizes Sampled Loop for High Stability and Linearity

Waveforms A, B and C are the LT1016's negative input, output and positive input, respectively. Their similarity to Figure 17 (Traces A, B and C) reflects the two circuits' commonality of operation. Trace D shows the quartz-crystal-derived 4kHz clock. During the clock's low portion, the LT1016's gated output appears at G2's output (Trace E). This data is loaded into counters which drive a 12-bit DAC via the 7475 latches. When the clock goes high, one section of the 74123 one-shot generates a pulse (Trace F), allowing the latches to acquire the counter's data. After this pulse goes low, the one-shot's second half pulses (Trace G) the counter's reset line. At the clock's next falling edge the entire cycle repeats. The DAC and its associated output amplifier (A1A) provide a voltage representation of the digital word at the 7475 outputs. This voltage is compared to the circuit's input by A1B, whose output drives the LT1016-based V→F. Any drift or nonlinearity in the V→F will be corrected by the feedback action of this stabilizing loop. The 10k-0.68μF time constant at A1B provides loop compensation.

Although it is not obvious, the frequency setting resolution is much greater than the 12-bit quantization limit of the DAC. This is because the DAC's output dithering around the LSB is integrated to a pure DC level by loop time constants. Once the DAC has settled within an LSB, its output acts like a 4kHz clocked pulse width modulator. The slow loop time constants integrate the width-modulated information to DC, affording smooth, continuous frequency setting capability. The practical limit on resolution is due to the LT1016 oscillator's short term jitter and is about 25ppm of reading.

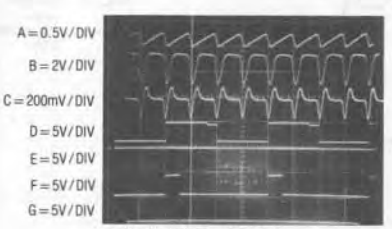

Figure 20. Waveforms for Figure 19. Sampled Data Loop (Traces D-G) Stabilizes Basic V→F (Traces A-C)

Although this approach allows higher speeds than Figure 16, there are some trade-offs. The loop's sampled nature, combined with its long time constants, limit settling time to about 100ms. Thus, although its output is faster than Figure 16's, it cannot track quickly varying inputs. Circuit linearity is DAC limited to 0.025% with full-scale drift of 50ppm/°C. Zero point drift of 1Hz/°C is due to A1B's 0.3μV/°C offset drift.

1Hz–1MHz Voltage-Controlled Sine Wave Oscillator

Both V→F converters described have pulse outputs. Many applications such as audio, shaker table driving, and automatic test equipment require voltage-controlled oscillators (VCO) with a sine wave output. The circuit of Figure 21 meets this need, spanning a 1Hz–1MHz range (120dB, or 6 decades) for a 0V to 10V input. It is over 10 times faster than previously published circuits, while maintaining 0.25% frequency linearity and 0.40% distortion specifications.

To understand the circuit, assume Q5 is on and its collector (Trace A, Figure 22) is at −15V, cutting off Q1. The positive input voltage is inverted by A3, which biases the summing node of integrator A1 through the 3.6k resistor and the self-biased FETs. A current, −I, is pulled from the summing point. A2, a precision op amp, DC stabilizes A1. A1's output (Trace B, Figure 22) integrates positive until C1's input (Trace C) crosses 0V. When this happens, C1's inverting output goes negative, the Q4–Q5 level shifter turns off, and Q5's collector goes to +15V. This allows Q1 to come on. The resistors in Q1's path are scaled to produce a current, +2I, exactly twice the absolute magnitude of the current, −I, being removed from the summing node. As a result, the net current into the junction becomes +I and A1 integrates negatively at the same rate its positive excursion took.

When A1 integrates far enough in the negative direction, C1's "+" input crosses zero and its outputs reverse. This switches the Q4-Q5 level shifter's state, Q1 goes off and the entire cycle repeats. The result is a triangle waveform at A1's output. The frequency of this triangle is dependent on the circuit's input voltage and varies from 1Hz to 1MHz with a 0V–10V input.

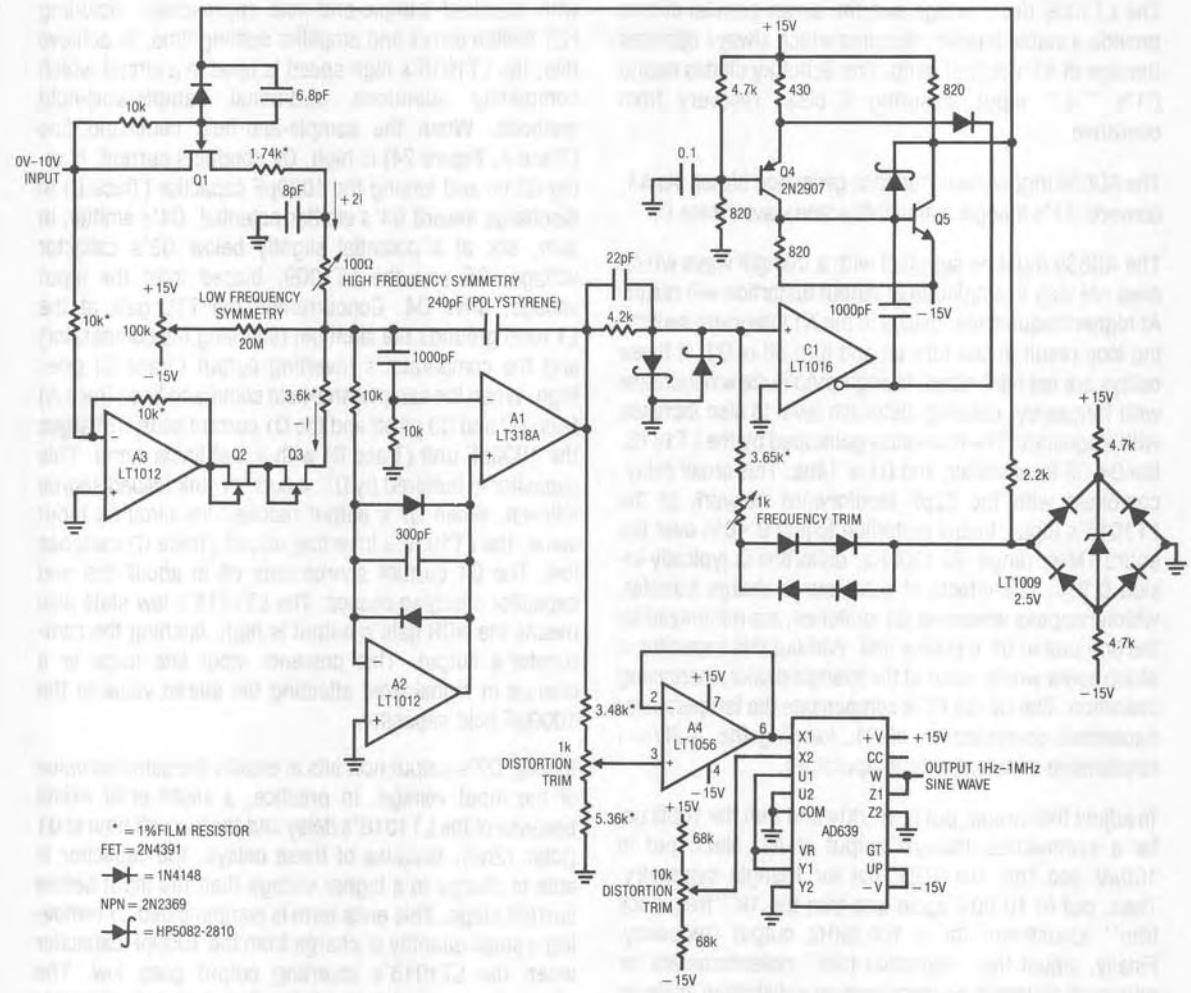

Figure 21. 1Hz–1MHz Sine Wave Output VCO

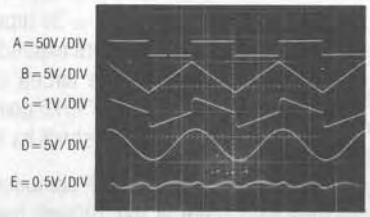

Figure 22. Sine Wave VCO Waveforms

The LT1009 diode bridge and the series-parallel diodes provide a stable bipolar reference which always opposes the sign of A1's output ramp. The Schottky diodes bound C1's "+" input, assuring it clean recovery from overdrive.

The AD639 trigonometric function generator, biased via A4, converts A1's triangle output into a sine wave (Trace D).

The AD639 must be supplied with a triangle wave which does not vary in amplitude or output distortion will result. At higher frequencies, delays in the A1 integrator switching loop result in late turn on and turn off of Q1. If these delays are not minimized, triangle amplitude will increase with frequency, causing distortion level to also increase with frequency. The total delay generated by the LT1016, the Q4-Q5 level shifter, and Q1 is 14ns. This small delay, combined with the 22pF feedforward network at the LT1016's input, keeps distortion to just 0.40% over the entire 1MHz range. At 100kHz, distortion is typically inside 0.2%. The effects of gate-source charge transfer, which happens whenever Q1 switches, are minimized by the 8pF unit in Q1's source line. Without this capacitor, a sharp spike would occur at the triangle peaks, increasing distortion. The Q2-Q3 FETs compensate the temperature-dependent on-resistance of Q1, keeping the $+2I/-I$ relationship constant with temperature.

To adjust this circuit, put in 10.00V and trim the 100Ω pot for a symmetrical triangle output at A1. Next, put in 100μV and trim the 100k pot for triangle symmetry. Then, put in 10.00V again and trim the 1k "frequency trim" adjustment for a 100.0kHz output frequency. Finally, adjust the "distortion trim" potentiometers for minimum distortion as measured on a distortion analyzer (Trace E). Slight readjustment of the other potentiometers may be required to get lowest possible distortion.

200ns-0.01% Sample-and-Hold Circuit

Figure 23's circuit uses the LT1016's high speed to improve upon a standard circuit function. The 200ns acquisition time is well beyond monolithic sample-and-hold capabilities and is matched only by hybrid and modular units selling in the $200 range. Other specifications exceed the best commercial unit's performance. This circuit also gets around many of the problems associated with standard sample-and-hold approaches, including FET switch errors and amplifier settling time. To achieve this, the LT1016's high speed is used in a circuit which completely abandons traditional sample-and-hold methods. When the sample-and-hold command line (Trace A, Figure 24) is high, Q2 conducts current, biasing Q3 on and forcing the 1000pF capacitor (Trace B) to discharge toward Q4's emitter potential. Q4's emitter, in turn, sits at a potential slightly below Q3's collector voltage. Q5 and the LT1009, biased from the input voltage, drive Q4. Concurrently, the TTL gate at the LT1016 grounds the latch pin (enabling the comparator) and the comparator's inverting output (Trace C) goes high. When the sample-and-hold command line (Trace A) falls, Q2 and Q3 go off and the Q1 current source charges the 1000pF unit (Trace B) with a fast linear ramp. This capacitor is buffered by Q7, a current sink loaded source follower. When Q7's output reaches the circuit's input value, the LT1016's inverting output (Trace C) switches low. The Q1 current source cuts off in about 2ns and capacitor charging ceases. The LT1016's low state also means the NOR gate's output is high, latching the comparator's output. This prevents input line noise or a change in signal from affecting the stored value in the 1000pF hold capacitor.

Ideally, Q7's output now sits at exactly the sampled value of the input voltage. In practice, a slight error exists because of the LT1016's delay and the turn-off time of Q1 (total 12ns). Because of these delays, the capacitor is able to charge to a higher voltage than the input before current stops. This error term is compensated by removing a small quantity of charge from the 1000pF capacitor when the LT1016's inverting output goes low. The charge is removed through the 8pF-1kΩ potentiometer network. Because the charging ramp's slope is fixed, the error term is constant and the compensation works over the circuit's ±3V input common-mode range. The lower four traces are expanded to show detail of the compensation and the circuit's critical ramp turn-off sequence. When the LT1016 goes off (Trace D), the ramp is seen to slightly overshoot its final value (Trace E).

The 1kΩ-8pF combination pulls enough charge (Trace F) out of the 1000pF hold capacitor to bring it back to the correct value. Trace G is the $\overline{\text{NOW}}$ line. It falls low 2 gate delays after the LT1016 inverting output goes low. When

this line goes low, the circuit's sampled output has settled from the correction transient and is valid data. The total time from the falling of the sample-and-hold line to the $\overline{NOW}$ output going low will always be inside 200ns.

The circuit's 200ns acquisition time is due to the high slew rate of the charging ramp and the action of Q4, Q5 and the LT1009. These components form a wideband tracking amplifier whose output is always a fixed amount below the input. Q7's current source load (Q6) ensures that its V_{GS} does not change. Thus, Q3 will always reset the capacitor a small, relatively constant amount below any circuit input. In this way, the ramp does not have to run very long before it crosses the input value, and acquisition time versus input voltage is constant. In

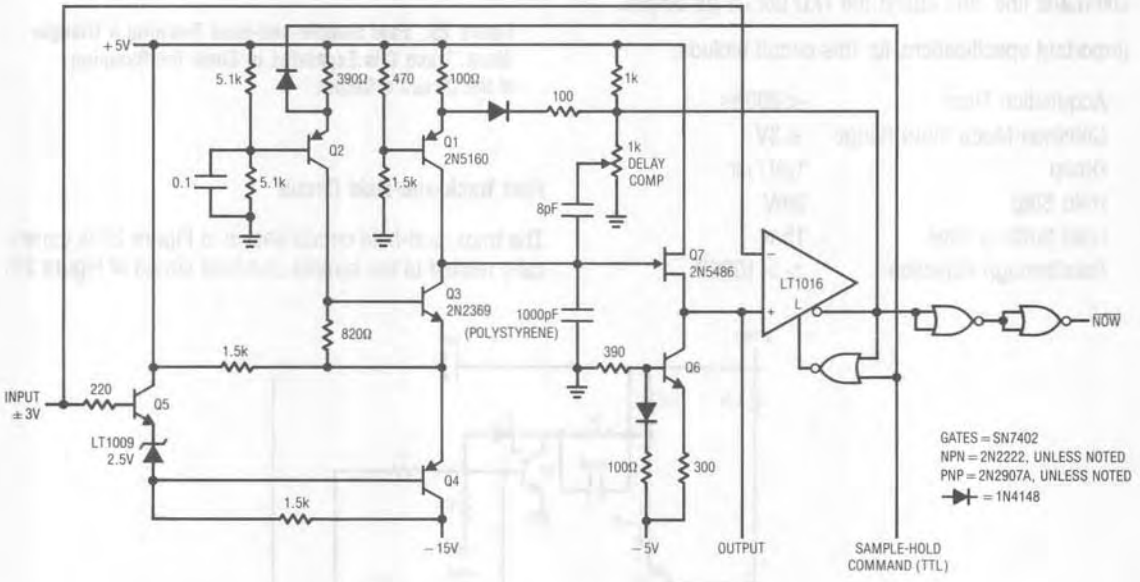

Figure 23. 200ns Sample-and-Hold

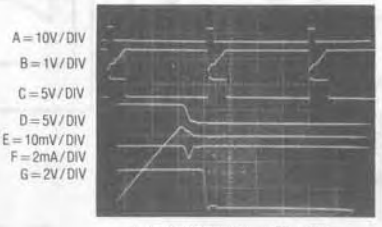

Figure 24. Fast Sample-and-Hold Waveforms. Traces A-C Show Ramp-Compare Action. Traces D-G Detail Delay Compensation

Figure 25 the circuit is shown sampling a bipolar triangle wave. Trace A is the input and Trace B is the circuit output. Trace C is an expansion of Trace B (the ''smearing'' of the sampled pedestals in Trace C is due to the repetitive asynchronous sampling of the triangle). The action of the tracking amplifier is readily apparent. It always resets the ramp to the same point below the input voltage, regardless of the common-mode level. To calibrate the circuit, ground the input, repetitively pulse the sample-and-hold command line, and adjust the 1kΩ pot for 0V output.

Important specifications for this circuit include:

Acquisition Time	<200ns
Common-Mode Input Range	± 3V
Droop	$1\mu V/\mu s$
Hold Step	2mV
Hold Settling Time	15ns
Feedthrough Rejection	$>>100$dB

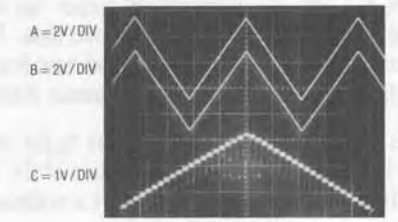

Figure 25. Fast Sample-and-Hold Tracking a Triangle Wave. Trace C is Expanded to Show the Ramping of the Circuit's Output.

Fast Track-and-Hold Circuit

The track-and-hold circuit shown in Figure 26 is generically related to the sample-and-hold circuit of Figure 23.

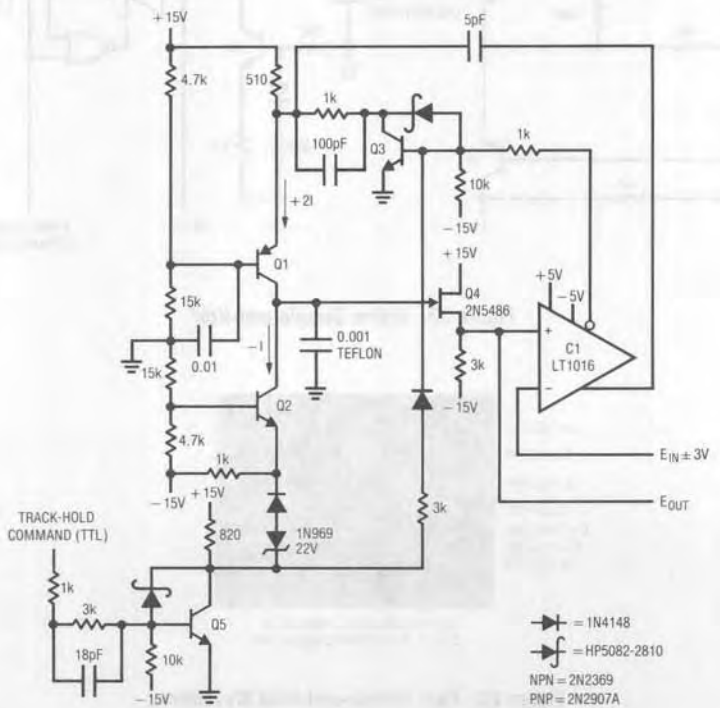

Figure 26. Comparator-Based Track-and-Hold

It also forsakes standard techniques in favor of an approach based on the LT1016's speed. This circuit's main blocks are a switched current source (Q1–Q3), a current sink (Q2), a FET follower (Q4), and the LT1016. To understand the circuit, assume the voltage stored in the 0.001μF hold capacitor is below the input potential and the track-and-hold command line (Trace A, Figure 27) is at a TTL ''1'' (track mode). Under these conditions Q5 is on and C1's output is positive. C1's inverting output is low and Q3 is off, allowing the Q1 current source to charge the hold capacitor. The Q2 current sink is also operating, but at ½ the current density of Q1. The hold capacitor charges positively. When Q4's source (Trace B, Figure 27) ramps to the input voltage's value, C1's outputs reverse state. Q3 comes on, quickly turning off the Q1 current source. The 5pF feedforward capacitor speeds up Q1's turnoff by bypassing Q3. With Q1 off, Q2's sink current discharges the hold capacitor. This causes C1's output to change state and oscillation commences (Trace B, Figure 27). This controlled, 10mV-25MHz oscillation centers itself around the input voltage's value. When the track-and-hold line (Trace A) goes low, Q5 ceases conducting, Q1 and Q2 immediately go off, oscillations cease and the circuit's output sits within $\pm$5mV of the input value at the time of turn-off. This 5mV uncertainty, caused by the nature of the circuit's operation, limits accuracy to 8 bits.

Figure 28 shows what happens when a square wave is fed into the circuit. Trace A is the input. Trace B is the output. Trace C is the track-and-hold command line and Trace D is the LT1016's output. Note that the controlled oscillation stops cleanly when the track-and-hold line goes low. If the source-sink transistors were run at higher currents, the circuit's output would slew much faster to keep up with the input's transitions. The oscillation's error band would also proportionately enlarge. The 25MHz update rate allows this circuit to track a relatively slow signal very closely with settling time under 10ns when switched into hold.

10ns Sample-and-Hold

Figure 29 shows a 10ns acquisition time sample-and-hold which can be used with repetitive signals only. Here, the LT1016 (C1) drives a differential integrator's (A1) input. Feedback from the integrator back to the LT1016 closes a loop around the circuit. Figure 30 shows what happens when a 1MHz sine wave (Trace A, Figure 30) is applied to the input. C2 generates a zero crossing signal (Trace B) and one-shot ''A'' (Trace C) provides an adjustable width. One-shot B's Q output produces a 30ns pulse (Trace D) which is fed into a logic network with the $\overline{Q}$ signal. The two inverter delays in Q's path give its associated gate a shorter duration output (Trace F) than $\overline{Q}$'s gate (Trace E). The last gate subtracts these two signals and generates a 10ns spike. This is inverted (Trace G) and fed to C1's latch pin. Each time the latch is enabled the comparator responds to the condition of the summing junction at its ''+'' input. If summing error is positive, A1 pulls current. If the error is negative, A1 sources current to the junction. After a number of input cycles, A1's output settles at a DC value which is the same as the level sampled during the time the latch is enabled. The ''delay adjust'' allows the 10ns sampling ''window'' to be positioned anywhere on the input sine wave.

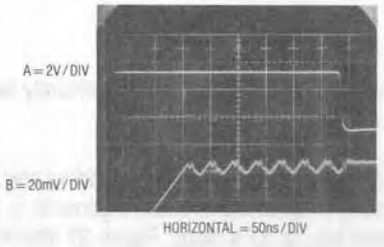

Figure 27. Track-and-Hold Circuit Acquiring an Input

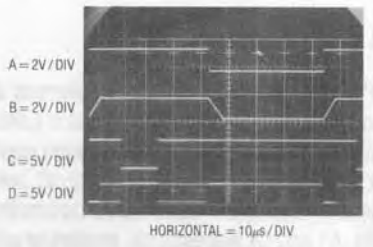

Figure 28. Track-and-Hold Responding to a Square Wave Input

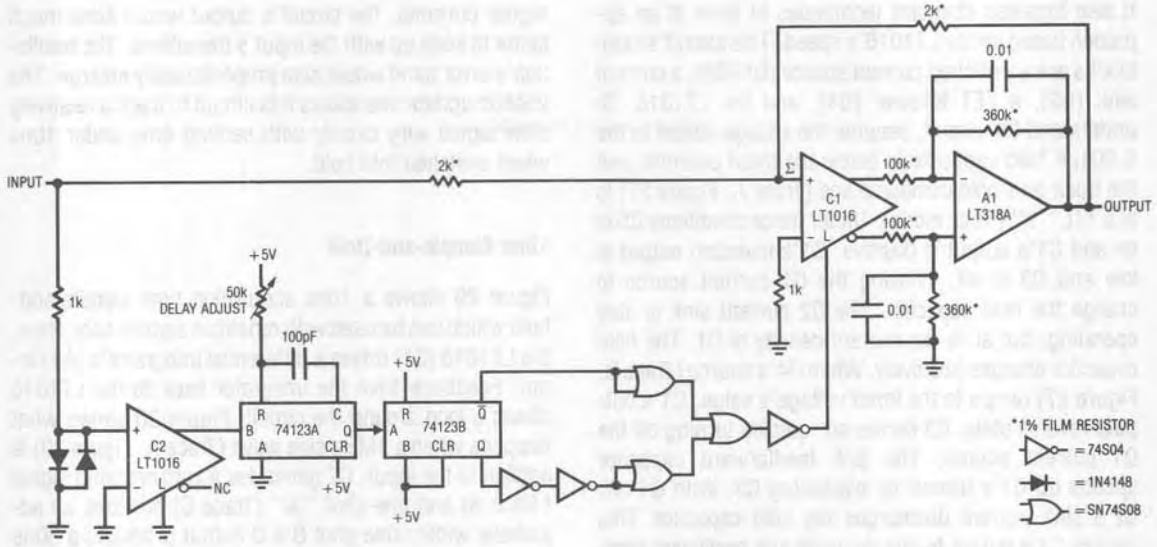

Figure 29. 10ns Sample-and-Hold for Repetitive Signals

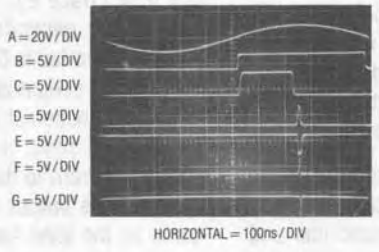

Figure 30. Waveforms for 10ns Sample-and-Hold. 10ns Sampling Window (Trace G) May be Positioned Anywhere on Input (Trace A)

5μs, 12-Bit A-D Converter

The LT1016's high speed is used to implement a fast 12-bit A–D converter in Figure 31. The circuit is a modified form of the standard successive approximation approach and is faster than most commercial SAR 12-bit units. In this arrangement the 2504 successive approximation register (SAR), A1 and C1 test each bit, beginning with the MSB, and produce a digital word representing V_{IN}'s value. To get faster conversion time, the clock (C2) is sped up after the third MSB is converted. This takes advantage of the segmented DAC used, which has significantly faster settling time for the lower 9 bits.

A1 provides preamplification for C1 while adding only 7ns delay. The preamplification allows clean response to ½ LSB (1.22mV) overdrives at A1's input. Figure 32 shows the converter at work. To aid in observing operation, A1 has been eliminated and the DAC-input node drives the LT1016 "+" input directly. A1 should be employed in normal use.

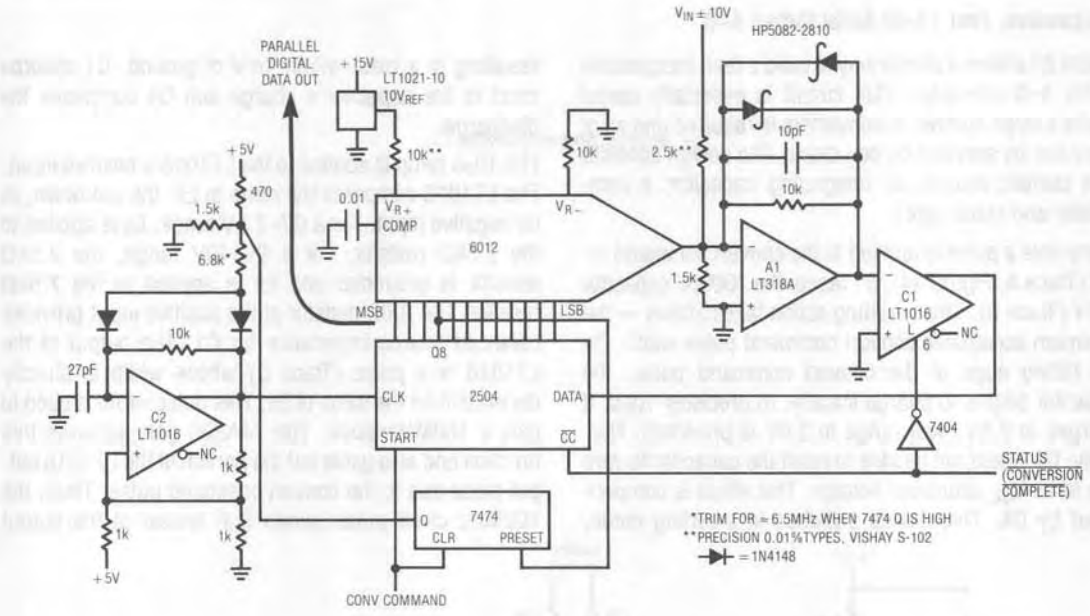

Figure 31. 5μs, 12-Bit SAR Converter. Clock is Sped up after the Third Bit, Shortening Overall Conversion Time

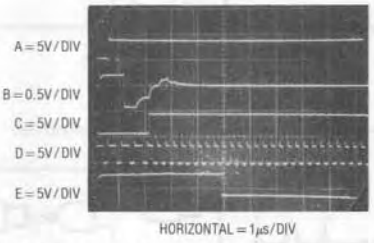

Figure 32. Fast SAR Converter Waveforms. Note Clock (Trace D) Speed-Up after 3rd Bit Conversion

The conversion begins when the ''convert command'' line (Trace A, Figure 32) drops low. When this happens the SAR begins to test each bit. The DAC output (Trace B), fed to the Schottky-clamped C1 input, sequentially converges toward final value. After the third MSB has been established, the 7474 Q line goes high (Trace C), forcing the 2.1MHz clock to shift to 3.2MHz (Trace D). This speeds up conversion of the remaining 9 bits, minimizing overall A–D time. When conversion is complete, the status line (Trace E) drops low and C1's latch is set by the TTL inverter, preventing the comparator from responding to input noise or shifts. The next convert command reinitiates the entire cycle. Note that on the lowest order bits C1 must accurately respond to small signals without sacrificing speed. The high gain-bandwidth required makes this application one of the most difficult for a comparator. This circuit's 5μs conversion time is fast for a 12-bit A–D. Faster conversion time is possible, although the design becomes more complex. A ''stretched'' version of this circuit, with 1.8μs conversion time, appears in AN17, ''Considerations for Successive Approximation A–D Converters''.

Inexpensive, Fast 10-Bit Serial Output A-D

Figure 33 shows a simple way to build a fast, inexpensive 10-bit A–D converter. This circuit is especially useful where a large number of converters is required and all of them can be serviced by one clock. The design consists of a current source, an integrating capacitor, a comparator and some gates.

Every time a pulse is applied to the convert command input (Trace A, Figure 34), Q1 resets the 1000pF capacitor to 0V (Trace B). This resetting action takes 200ns — the minimum acceptable convert command pulse width. On the falling edge of the convert command pulse, the capacitor begins to charge linearly. In precisely 10μs, it charges to 2.5V (over range to 3.0V is provided). Normally, Q1 would not be able to reset the capacitor to zero due to its V_{CE} saturation voltage. This effect is compensated by Q4. This device switches in inverting mode, resulting in a reset within 1mV of ground. Q1 absorbs most of the capacitor's charge and Q4 completes the discharge.

The 10μs ramp is applied to the LT1016's positive input. The LT1016 compares the ramp to Ex, the unknown, at its negative input. For a 0V–2.5V range, Ex is applied to the 2.5kΩ resistor. For a 0V–10V range, the 2.5kΩ resistor is grounded and Ex is applied to the 7.5kΩ resistor. The 2.0k resistor at the positive input provides balanced source impedance for C1. The output of the LT1016 is a pulse (Trace C) whose width is directly dependent on the value of Ex. This pulse width is used to gate a 100MHz clock. The 74AS00 gate achieves this function and also gates out the portion of the LT1016 output pulse due to the convert command pulse. Thus, the 100MHz clock pulse bursts that appear at the output

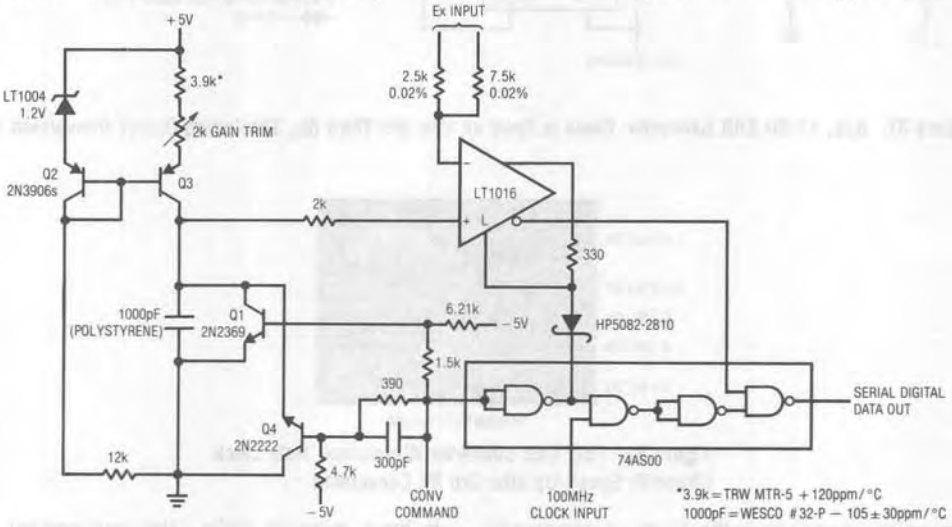

Figure 33. Simple, Fast 10-Bit A → D

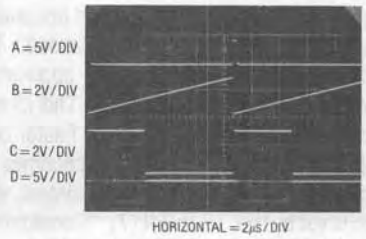

Figure 34. Waveforms for 10-Bit A → D

(Trace D) are proportional to Ex. For a 0V–10V input, 1024 pulses appear at full-scale, 512 at 5.00V, etc. The resistor-diode network at the LT1016's latch pin ensures clean comparator transitions by locking the LT1016 outputs after the conversion is completed. This latch is broken by the next convert command pulse.

The current source scaling resistor and ramp capacitor specified provide good temperature compensation because of their opposing thermal coefficients. The circuit will typically hold ±1 LSB accuracy over 0°C–70°C with an additional ±1 LSB uncertainty due to the asynchronous relationship between the clock and the conversion sequence.

Figure 35 details the most critical part of the converter's operation, the reset phase. Trace A is the convert command. Trace B is the capacitor (greatly magnified) resetting to zero. The comparator output appears in Trace C and Trace D is the gated serial output. Observe that the output pulses do not appear until the capacitor has started to ramp (just past mid-screen), even though the comparator is high.

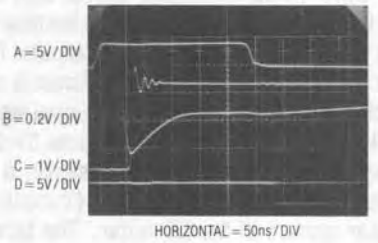

Figure 35. Figure 33's Reset Sequence. Q1-Q4 Combination Gives Quick, Low Offset Zero Reset

2.5MHz Precision Rectifier/AC Voltmeter

Most precision rectifier circuits rely on operational amplifiers to correct for diode drops. Although this scheme works well, bandwidth limitations usually restrict these circuits to operation below 100kHz. Figure 36 shows the LT1016 in an open-loop, synchronous rectifier configuration which has high accuracy out to 2.5MHz. An input

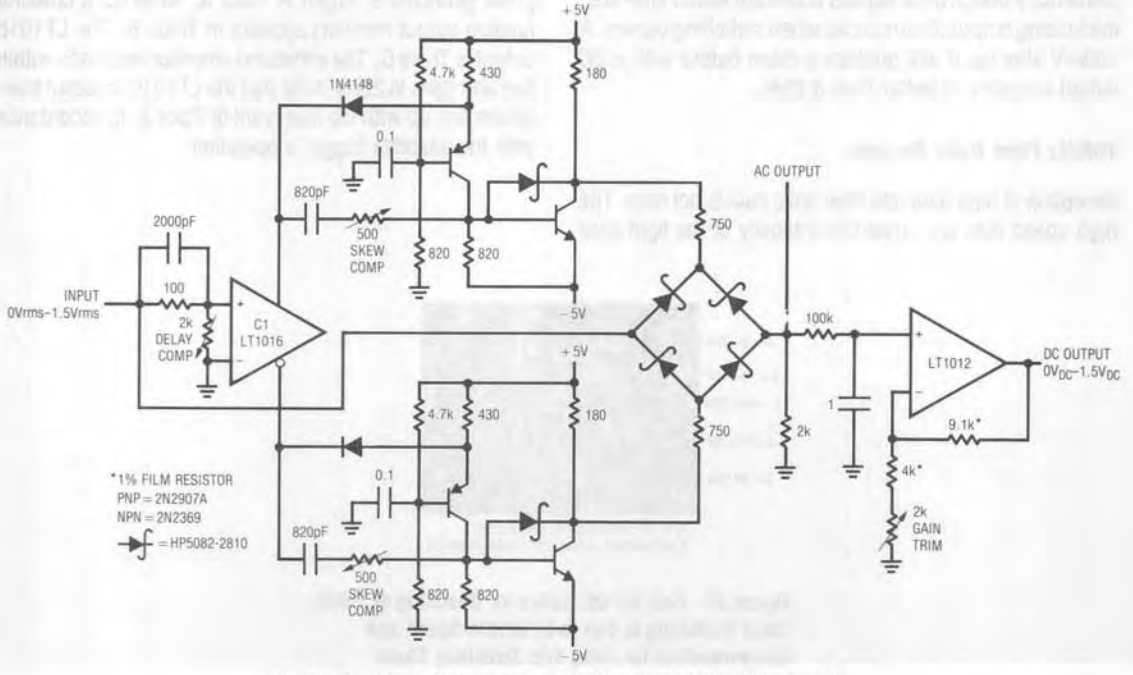

Figure 36. Fast, Synchronous Rectifier-Based AC-DC Converter

1MHz sine wave (Trace A, Figure 37) is zero cross detected by C1. Both of C1's outputs drive identical level shifters with fast (delay = 2ns–3ns), ± 5V outputs. These outputs bias a Schottky switching bridge (Traces B and C are the switched corners of the bridge). The input signal is fed to the left-midsection of the bridge. Because C1 drives the bridge synchronously with the input signal, a half-wave rectified sine appears at the AC output (Trace D). The DC RMS value appears at the DC output. The Schottky bridge gives fast switching and eliminates the charge pump-through that a FET switch would contribute. This is evident in Trace E, which is an expanded version of Trace D. The waveform is clean with the exception of very small disturbances where bridge switching occurs. To calibrate this circuit, apply a 1MHz–2MHz 1Vp-p. Sine wave and adjust the delay compensation so bridge switching occurs when the sine crosses zero. This adjustment corrects for the small delays through the LT1016 and the level shifters. Next, adjust the skew compensation potentiometers for minimum aberrations in the AC output signal. These trims slightly shift the phase of the rising output edge of their respective level shifter. This allows skew in the complementary bridge drive signals to be kept within 1ns–2ns, minimizing output disturbances when switching occurs. A 100mV sine input will produce a clean output with a DC output accuracy of better than 0.25%.

10MHz Fiber Optic Receiver

Reception of high data rate fiber optic data is not easy. The high speed data and uncertain intensity of the light level can cause erroneous results unless the receiver is carefully designed. The fiber optic receiver shown in Figure 38 will accurately condition a wide range of light inputs at up to 10MHz data rates. Its digital output features an adaptive threshold trigger which accommodates varying signal intensities due to component aging and other causes. An analog output is also available to monitor the detector output. The optical signal is detected by the PIN photodiode and amplified by a broadband fed-back stage, Q1–Q3. A second, similar, stage gives further amplification. The output of this stage (Q5's collector) biases a 2-way peak detector (Q6–Q7). The maximum peak is stored in Q6's emitter capacitor, while the minimum excursion is retained in Q7's emitter capacitor. The DC value of Q5's output signal's mid-point appears at the junction of the 0.005μF capacitor and the 22MΩ unit. This point will always sit midway between the signal's excursions, regardless of absolute amplitude. This signal-adaptive voltage is buffered by the low bias LT1012 to set the trigger voltage at the LT1016's positive input. The LT1016's negative input is biased directly from Q5's collector. Figure 39 shows the results using the test circuit indicated in Figure 38. The pulse generator's output is Trace A, while Q5's collector (analog output monitor) appears in Trace B. The LT1016 output is Trace C. The wideband amplifier responds within 5ns and rises in 25ns. Note that the LT1016's output transitions line up with the mid-point of Trace B, in accordance with the adaptive trigger's operation.

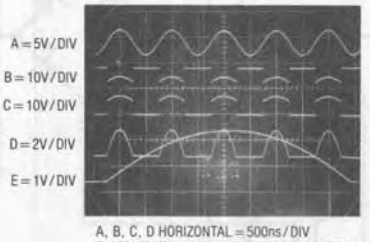

Figure 37. Fast AC-DC Converter Operating at 1MHz. Clean Switching is Due to LT1016's Speed and Compensations for Delay and Switching Skew

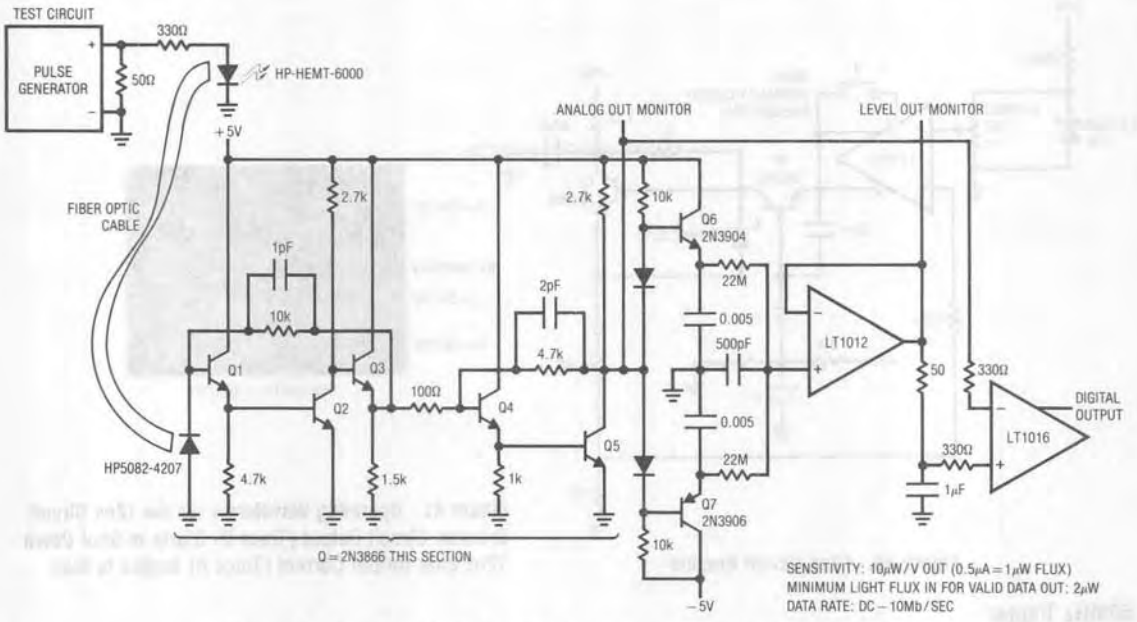

Figure 38. Fast Fiber Optic Receiver is Immune to Shifts in Operating Point

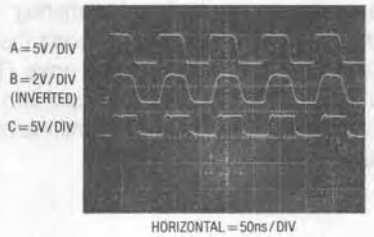

Figure 39. Fiber Optic Receiver Waveforms

12ns Circuit Breaker

Figure 40 shows a simple circuit which will turn off current in a load 12ns after it exceeds a preset value. This circuit has been used to protect integrated circuits during developmental probing and is also useful for protecting expensive loads during trimming and calibration. It is 3 times faster and less complex than previously published circuits. Under normal conditions the voltage across the 10Ω shunt is smaller than the potential at the LT1016's negative input. This keeps Q1 off and Q2 receives bias, driving the load. When an overload occurs (in this case via a test circuit, whose output is Trace A, Figure 41), the current through the 10Ω sense resistor begins to increase (Trace B, Figure 41). When this current exceeds the preset value, the LT1016's outputs (non-inverting output shown in Trace C) reverse. This provides ideal turn-on drive for Q1 and it cuts off Q2 (Q2 emitter is Trace D) in 5ns. The delay from the onset of excessive load current to complete shutdown is just 13ns. Once the circuit has triggered, the LT1016 is held in its latched state by feedback from the non-inverting output. When the load fault has been cleared the pushbutton can be used to reset the circuit.

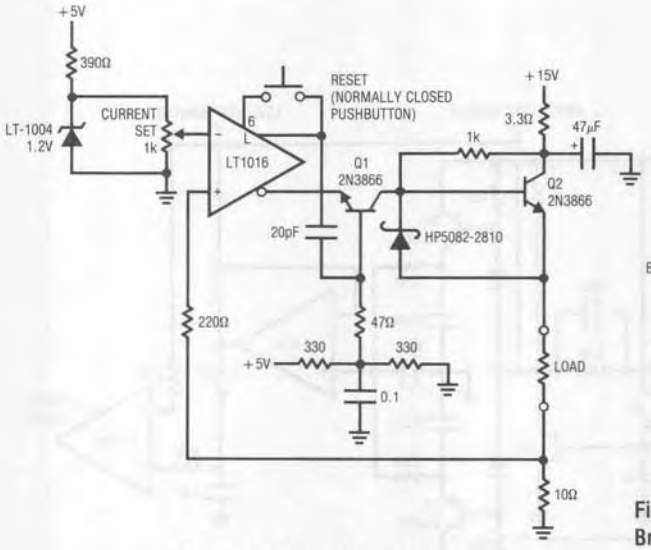

Figure 40. 12ns Circuit Breaker

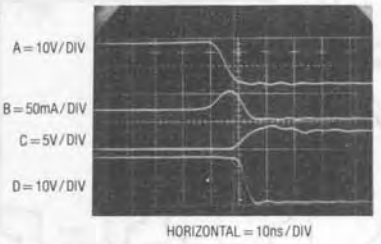

Figure 41. Operating Waveforms for the 12ns Circuit Breaker. Circuit Output (Trace D) Starts to Shut Down 12ns after Output Current (Trace B) Begins to Rise

50MHz Trigger

Counters and other instruments require a trigger circuit. Designing a fast, stable trigger is not easy, and often entails a considerable amount of discrete circuitry. Figure 42 shows a simple trigger with 100mV sensitivity at 50MHz. The FETs comprise a simple high speed buffer and the LT1016 compares the buffer's output to the potential at the ''trigger level'' potentiometer, which may be either polarity. The 10k resistor provides hysteresis, eliminating ''chattering'' caused by noisy input signals. Figure 43 shows the trigger's response (Trace B) to a 50MHz sine wave (Trace A). To calibrate this circuit, ground the input and adjust the ''input zero'' control for 0V at Q2's drain terminal.

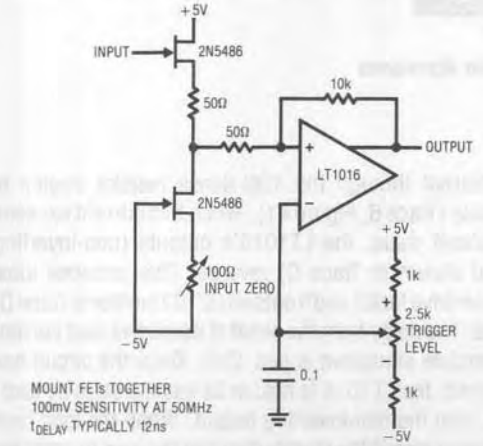

Figure 42. 50MHz Trigger

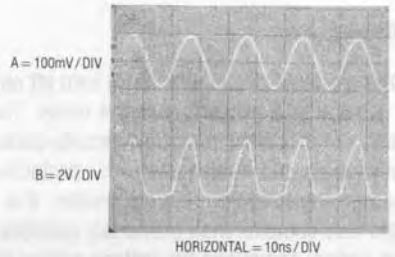

Figure 43. Trigger Responding to a 50MHz Sine Input

REFERENCES

Dendinger, S., "One IC Makes Precision Sample and Hold," *EDN,* May 20, 1977.

Pease, R. A., "Amplitude to Frequency Converter," U.S. Patent #3,746,968, Filed September, 1972.

Hewlett-Packard Application Note #915, "Threshold Detection of Visible and Infra-Red Radiation with PIN Photodiodes."

Williams, J., "A Few Proven Techniques Ease Sine-Wave-Generator Design," *EDN,* November 20, 1980, page 143.

Williams, J., "Simple Techniques Fine-Tune Sample-Hold Performance," *Electronic Design,* November 12, 1981, page 235.

Baker, R. H., "Boosting Transistor Switching Speed," *Electronics,* Vol. 30, 1957, pages 190–193.

Bunze, V., "Matching Oscilloscope and Probe for Better Measurements," *Electronics,* March 1, 1973, pages 88–93.

APPENDIX A

About Bypass Capacitors

Bypass capacitors are used to maintain low power supply impedance at the point of load. Parasitic resistance and inductance in supply lines mean that the power supply impedance can be quite high. As frequency goes up, the inductive parasitic becomes particularly troublesome. Even if these parasitic terms did not exist, or if local regulation is used, bypassing is still necessary because no power supply or regulator has zero output impedance at 100MHz. What type of bypass capacitor to use is determined by the application, frequency domain of the circuit, cost, board space and many other considerations. Some useful generalizations can be made.

All capacitors contain parasitic terms, some of which appear in Figure A1. In bypass applications, leakage and dielectric absorption are second order terms but series R and L are not. These latter terms limit the capacitor's ability to damp transients and maintain low supply impedance. Bypass capacitors must often be large values so they can absorb long transients, necessitating electrolytic types which have large series R and L.

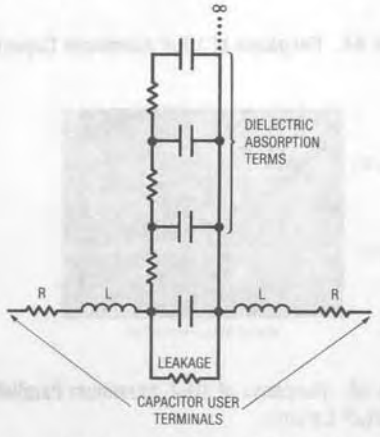

Figure A1. Parasitic Terms of a Capacitor

Different types of electrolytics and electrolytic-non-polar combinations have markedly different characteristics. Which type(s) to use is a matter of passionate debate in some circles and the test circuit (Figure A2) and accompanying photos are useful. The photos show the response of 5 bypassing methods to the transient generated by the test circuit. Figure A3 shows an unbypassed line which sags and ripples badly at large amplitudes. Figure A4 uses an aluminum 10μF electrolytic to considerably cut the disturbance, but there is still plenty of potential trouble. A tantalum 10μF unit offers cleaner response in A5 and the 10μF aluminum combined with a 0.01μF ceramic type is even better in A6. Combining electrolytics with non-polarized capacitors is a popular way to get good response but beware of picking the wrong duo. The right (wrong) combination of supply line parasitics and paralleled dissimilar capacitors can produce a resonant, ringing response, as in A7. Caveat!

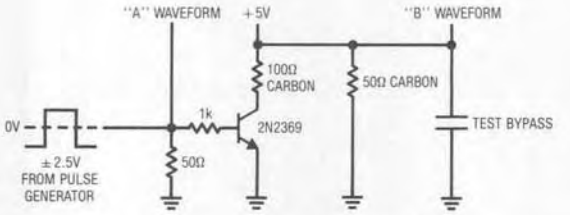

Figure A2. Bypass Capacitor Test Circuit

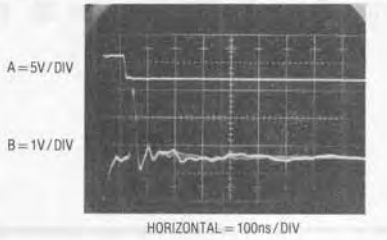

Figure A3. Response of Unbypassed Line

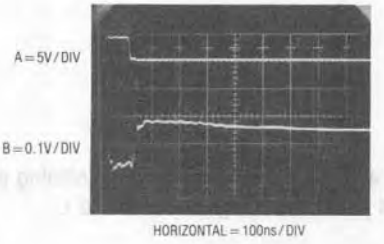

Figure A4. Response of 10μF Aluminum Capacitor

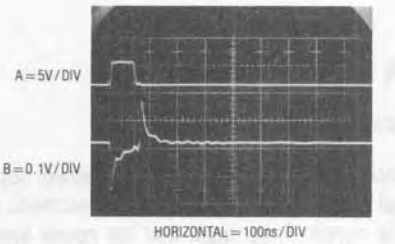

Figure A5. Response of 10μF Tantalum Capacitor

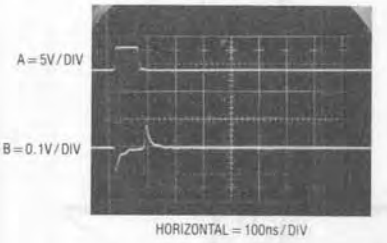

Figure A6. Response of 10μF Aluminum Paralleled by 0.01μF Ceramic

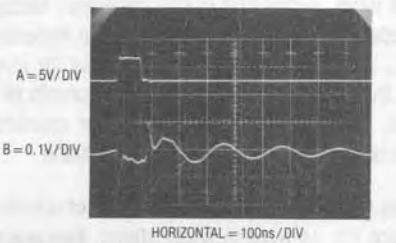

Figure A7. Some Paralleled Combinations can Ring. Try before Specifying!

APPENDIX B

About Probes and Oscilloscopes

The oscilloscope-probe combination used in high speed work is the most important equipment decision the designer must make. Ideally, the oscilloscope should have at least 150MHz bandwidth for work with the LT1016, but slower instruments are acceptable if their limitations are well understood. Be aware of your scope's behavior with respect to input impedance, noise, overdrive recovery, sweep nonlinearity, triggering, channel-to-channel feedthrough and other characteristics.

Probes are the most overlooked cause of oscilloscope mismeasurement. All probes have some effect on the point they are measuring. The most obvious is input resistance, but input capacitance usually dominates in a high speed measurement. Much time can be lost chasing circuit events which are actually due to improperly selected or applied probes. An 8pF probe looking at a 1kΩ source impedance will form an 8ns lag — close to the LT1016's response time! Low impedance probes, designed for 50Ω inputs, (with 500Ω to 1kΩ resistance) usually have input capacitance of 1pF or 2pF. They are a very good choice if you can stand the low resistance. FET probes maintain high input resistance and keep capacitance at the 1pF level but have substantially more delay than passive probes. FET probes also have limitations on input common-mode range which must be adhered to or serious measurement errors will result. Contrary to popular belief, FET probes *do not* have extremely high input resistance — some types are as low as 100kΩ.

Current probes are useful and convenient. The passive transformer-based types are fast and have less delay than the Hall effect-based versions. The Hall types, however, respond at DC and low frequency and the transformer types typically roll off around 100Hz to 1kHz. Both types have saturation limitations which, when exceeded, cause odd results on the CRT which will confuse the unwary.

When using different probes remember that they all have different delay times, meaning that apparent timing errors will occur on the CRT. Know what the individual probe delays are and account for them in interpreting the CRT display.

By far the greatest source of error in probe use is grounding. Poor probe grounding can cause ripples and discontinuities in the waveform observed. In some cases the choice and placement of a probe's ground strap will affect waveforms on another channel. In the worst case, connecting the probe's ground wire will virtually disable the circuit being measured. The cause of these problems is due to parasitic inductance in the probe's ground connection. In most oscilloscope measurements this is not a problem, but at nanosecond speeds it becomes critical. Fast probes are always supplied with a variety of spring clips and accessories designed to aid in making the lowest possible inductive connection to ground. Most of these attachments assume a ground plane is in use, which it should be. Always try to make the shortest possible connection to ground — anything longer than 1 inch may cause trouble.

The simple network of Figure B1 shows just how easy it is for poorly chosen or used probes to cause bad results. A 9pF input capacitance probe with a 4 inch long ground-strap monitors the output (Trace B, Figure B2). Although the input (Trace A) is clean, the output contains ringing. Using the same probe with a ¼ inch spring tip ground connection accessory seemingly cleans up everything (Figure B3). However, substituting a 1pF FET probe (Figure B4) reveals a 50% output amplitude error in measurement B3! The FET probe's low input capacitance allows a more accurate version of circuit action. The FET probe does, however, contribute its own form of error. Note that the probe's response is tardy by 5ns due to delay in its active circuitry. Hence, separate measurements with each probe are required to determine amplitude and timing parameters of the output.

A final form of probe is the human finger. Probing the circuit with a finger can accentuate desired or undesired effects, giving clues that may be useful. The finger can be used to introduce stray capacitance to a suspected circuit node while observing results on the CRT. Two fingers, lightly moistened, can be used to provide an experimental resistance path. Some high speed engineers are particularly adept at these techniques and can estimate the capacitive and resistive effects created with surprising accuracy.

Examples of the probes discussed, along with different forms of grounding implements, are shown in Figure B5.

Probes A, B, E, and F are standard types equipped with various forms of low impedance grounding attachments. The conventional ground lead used on G is more convenient to work with but will cause ringing and other effects at high frequencies, rendering it useless. H has a very short ground lead. This is better, but can still cause trouble at high speeds. D is a FET probe. The active circuitry in the probe and a very short ground connector ensure low parasitic capacitance and inductance. C is a separated FET probe attenuator head. Such heads allow the probe to be used at higher voltage levels (e.g., ±10V or ±100V). The miniature coaxial connector shown can be mounted on the circuit board and the probe mated with it. This technique provides the lowest possible parasitic inductance in the ground path and is especially recommended.

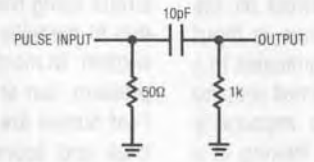

Figure B1. Probe Test Circuit

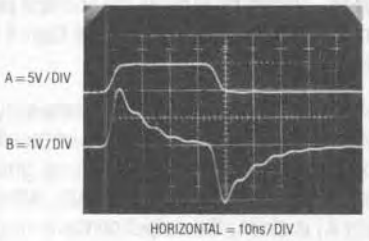

Figure B2. Test Circuit Output with 9pF Probe and 4 Inch Ground Strap

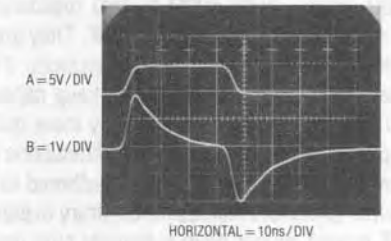

Figure B3. Test Circuit Output with 9pF Probe and 0.25 Inch Ground Strap

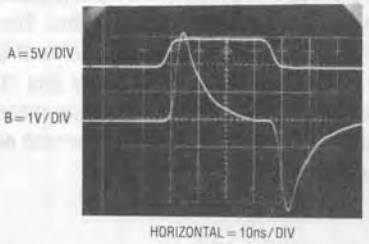

Figure B4. Test Circuit Output with FET Probe

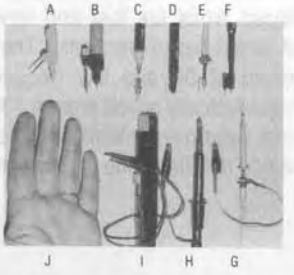

Figure B5. Various Probe-Ground Strap Configurations

I is a current probe. A ground connection is not usually required. However, at high speeds the ground connection may result in a cleaner CRT presentation. Because no current flows in the ground lead of these probes, a long strap is usually permissible. J is typical of the finger probes described in the text. Note the ground strap on the third finger.

The low inductance ground connectors shown are available from probe manufacturers and are always supplied with good quality, high frequency probes. Because most oscilloscope measurements do not require them, they invariably become lost. There is no substitute for these devices when they are needed, so it is prudent to take care of them. This is especially applicable to the ground strap on the finger probe.

APPENDIX C

About Ground Planes

Many times in high frequency circuit layout the term "ground plane" is used, most often as a mystical and ill-defined cure to spurious circuit operation. In fact, there is little mystery to the usefulness and operation of ground planes, and like many phenomena, their fundamental operational principle is surprisingly simple.

Ground planes are primarily useful for minimizing circuit inductance. They do this by utilizing basic magnetic theory. Current flowing in a wire produces an associated magnetic field. The field's strength is proportional to the current and inversely related to the distance from the conductor. Thus, we can visualize a wire carrying current (Figure C1) surrounded by radii of magnetic field. The unbounded field becomes smaller with distance. A wire's inductance is defined as the energy stored in the field set up by the wire's current. To compute the wire's inductance requires integrating the field over the wire's length and the total radial area of the field. This implies integrating on the radius from $R = R_W$ to infinity, a very large number. However, consider the case where we have two wires in space carrying the same current in either direction (Figure C2). The fields produced cancel.

In this case, the inductance is much smaller than in the sample wire case and can be made arbitrarily small by reducing the distance between the two wires. This reduction of inductance between current carrying conductors is the underlying reason for ground planes. In a normal circuit, the path current takes from the signal source, through its conductor and back to ground includes a large loop area. This produces a large inductance for this conductor which can cause ringing due to LRC effects. It is worth noting that 10nH at 100MHz has an impedance of 6Ω. At 10mA a 60mV drop results.

A ground plane provides a return path directly under the signal carrying conductor through which return current can flow. The conductor's small physical separation means the inductance is low. Return current has a direct path to ground, regardless of the number of branches associated with the conductor. Currents will always flow through the return path of lowest impedance. In a properly designed ground plane this path is directly under the signal conductor. In a practical circuit it is desirable to "ground plane" one whole side of the PC card (usually the component side for wave solder considerations) and run the signal conductors on the other side. This will give a low inductance path for all the return currents.

Aside from minimizing parasitic inductance, ground planes have additional benefits. Their flat surface minimizes resistive losses due to AC "skin effect" (AC currents travel along a conductor's surface). Additionally, they aid the circuit's high frequency stability by referring stray capacitances to ground.

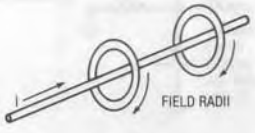

Figure C1. Single Wire Case

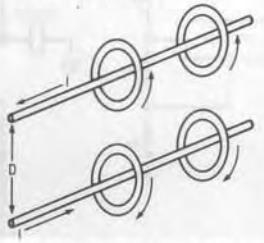

Figure C2. Two Wire Case

Some practical hints for ground planes are:

1. Ground plane as much area as possible on the component side of the board, especially under traces that operate at high frequency.

2. Mount components that conduct substantial fast rise currents (termination resistors, ICs, transistors, decoupling capacitors) as close to the board as possible.

3. Where common ground potential is important (i.e., at comparator inputs), try to single point the critical components into the ground plane to avoid voltage drops.

For example, in Figure C3's common A/D circuit, good practice would dictate that grounds 2, 3, 4 and 6 be as close to single point as possible.

Fast, large currents must flow through R1, R2, D1 and D2 during the DAC settle time. Therefore, D1, D2, R1 and R2 should be mounted close to the ground plane to minimize their inductance. R3 and C1 don't carry any current, so their inductance is less important; they could be vertically inserted to save space and to allow point 4 to be single point common with 2, 3 and 6. In critical circuits the designer must often trade off the beneficial effects of lowered inductance versus the loss of single point ground.

4. Keep trace length short. Inductance varies directly with length and no ground plane will achieve perfect cancellation.

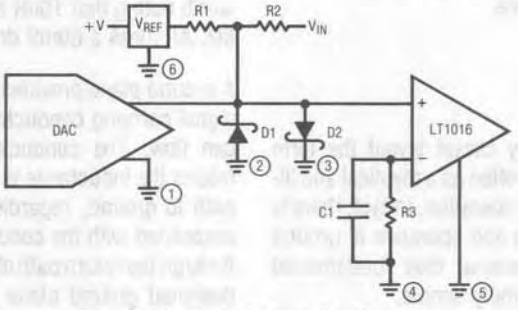

Figure C3. Typical Grounding Scheme

APPENDIX D

Measuring Equipment Response

The 10ns response time of the LT1016 and the circuitry it is used in will challenge the best test equipment. Many of the measurements made utilize equipment near the limit of its capabilities. It is a good idea to verify parameters such as probe and scope rise time and differences in delays between probes and even oscilloscope channels. To do this, a source of very fast, clean pulses is necessary. The circuit shown in Figure D1 uses a tunnel diode to generate a pulse with a rise time well under 1ns.

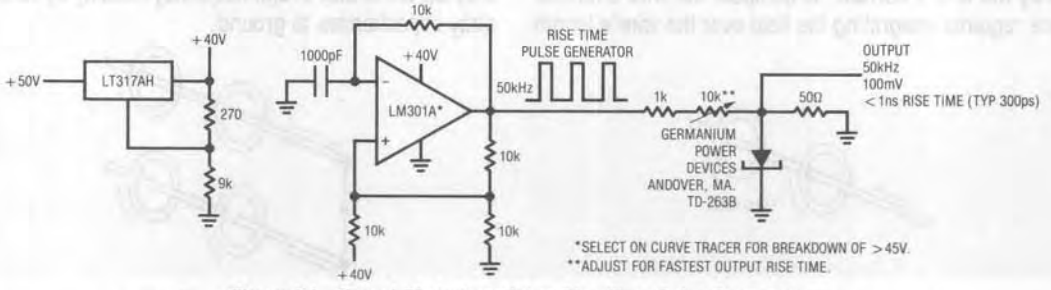

Figure D1. Tunnel Diode-Based 1ns Rise Time Pulse Generator

Figure C2 shows that the pulse is also very clean, with no attendant ringing or noise. In this photo the pulse is used to check a probe-scope combination with a specified 1.4ns rise time. The display shows that the equipment is being properly used and is in specification. Using the tunnel diode generator to perform tests such as this can save countless hours pursuing "circuit problems," which in reality are caused by misapplied or out of spec equipment.

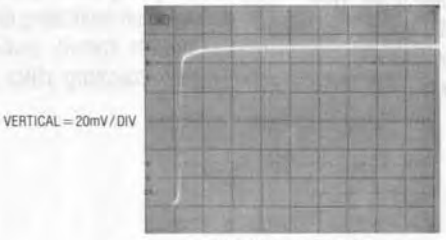

Figure D2. Figure D1's Output Monitored on a 275MHz Oscilloscope

APPENDIX E

About Level Shifts

The TTL output of the LT1016 will interface with many circuits directly. Many applications, however, require some form of level shifting of the output swing. With LT1016-based circuits this is not trivial because it is desirable to maintain very low delay in the level shifting stage. When designing level shifters, keep in mind that the TTL output of the LT1016 is a sink-source pair (Figure E1) with good ability to drive capacitance (such as feedforward capacitors).

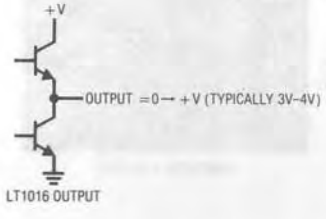

Figure E1

Figure E2 shows a non-inverting voltage gain stage with a 15V output. When the LT1016 switches, the base-emitter voltages at the 2N2369 reverse, causing it to switch very quickly. The 2N3866 emitter-follower gives a low impedance output and the Schottky diode aids current sink capability.

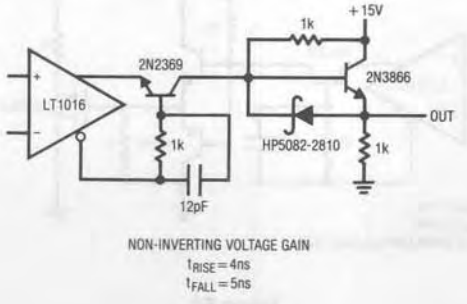

Figure E2

Figure E3 is a very versatile stage. It features a bipolar swing which may be programmed by varying the output transistor's supplies. This 3ns delay stage is ideal for driving FET switch gates. Q1, a gated current source, switches the Baker-clamped output transistor, Q2. The heavy feedforward capacitor from the LT1016 is the key to low delay, providing Q2's base with nearly ideal drive. This capacitor loads the LT1016's output transition (Trace A, Figure E5), but Q2's switching is clean (Trace B, Figure E5) with 3ns delay on the rise and fall of the pulse.

Figure E4 is similar to E2 except that a sink transistor has replaced the Schottky diode. The two emitter-followers drive a power MOSFET which switches 1A at 15V. Most of the 7ns–9ns delay in this stage occurs in the MOSFET and the 2N2369.

When designing level shifters, remember to use transistors with fast switching times and high f_T's. To get the kind of results shown, switching times in the ns range and f_T's approaching 1GHz are required.

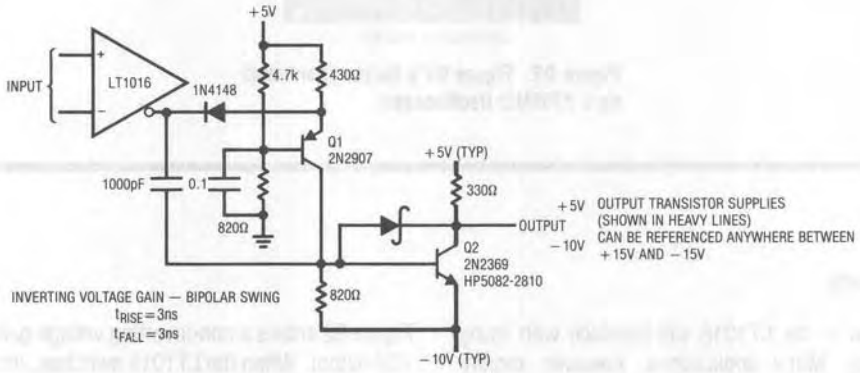

Figure E3

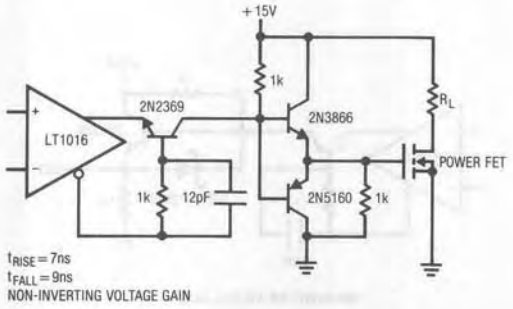

Figure E4

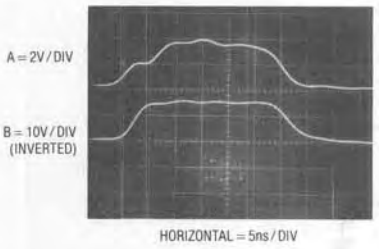

Figure E5. Figure E3's Waveforms