

THE
JIM WILLIAMS
P A P E R S

Subduing high-speed op-amp problems

The application of high-speed op amps requires special attention to a multitude of potential problems. You need to guard against noise intrusion, capacitance-loading effects, and parasitic conductive paths—without neglecting the compromises you may have to make between compensation and gain.

Jim Williams, *Linear Technology Corp.*

Compared with their low-speed siblings, high-speed op amps are subject to many problems, many of which can result in oscillation. The forte of the operational amplifier is negative feedback, which stabilizes the operating point and fixes the gain. However, positive feedback or delayed negative feedback can cause oscillation. Thus, even a properly functioning amplifier constantly lives in the shadow of oscillation.

When oscillation occurs, several major candidates for blame are present. If the power supply is unby-passed, the impedance the amplifier sees at its power terminals is high, particularly at high frequency. This impedance forms a voltage divider with the amplifier, letting the supply vary as internal conditions in the amplifier change. This variation can cause local feedback, resulting in oscillation. The obvious cure is to bypass the amplifier. Power-supply impedance must be low to ensure stable operation.

A second common cause of oscillation is positive feed-

back. In most amplifier circuits feedback is negative, although the circuit may also use controlled amounts of positive feedback. In a circuit that normally has only negative feedback, unintended positive feedback may occur with poor layout. Check for possible parasitic feedback paths and unwanted or overlooked feedback action. To the extent possible, minimize the impedances seen by the amplifier inputs. A low impedance helps to attenuate the effects of parasitic feedback paths to the inputs. Similarly, minimize exposed input-trace area. Route the amplifier outputs and other signals well away from sensitive nodes. Sometimes no amount of layout finesse will work, and shielding is required. Use shielding only when required—extensive shielding is a sloppy substitute for good layout practice.

Watch out for time delays

A third cause of oscillation is negative feedback that arrives significantly delayed in time. Under these conditions, the amplifier hopelessly tries to servo-control a feedback signal that consistently arrives “too late.” The servo action takes the form of an electronic tail chase, with oscillation centered around the ideal servo point. The most common causes of this problem are reactive loading of the amplifier (most notably capacitive loads such as cable) and circuitry (such as power amplifiers) within the amplifier’s feedback path. In many cases, isolating the reactive load from the amplifier’s output (and feedback path) with a resistor will solve the problem. Sometimes rolling off the amplifier’s frequency response will fix the problem, but in high-speed circuits this may not be an option.

Placing power-gain or other types of stages within

THE
JIM WILLIAMS
P A P E R S

the amplifier's feedback path adds time delay to the stabilizing feedback. If the delay is significant, oscillation commences. Stages operating within the amplifier's loop should have a minimum time lag compared with the amplifier's speed capability. At lower speeds, this is not too difficult, but a stage within a 100-MHz amplifier's loop must be fast. As mentioned before, rolling off the amplifier's frequency response eases the job, but is usually undesirable in a wideband circuit. You should make every effort to maximize the bandwidth of the added stages before resorting to amplifier roll-off. In this way, you can achieve the fastest overall bandwidth while maintaining stability. (See the box, "The oscillation problem and frequency compensation," which discusses power-gain stages and other types of stages operating within amplifier loops.)

In certain cases, it may appear that there is an oscillation problem when there really isn't. For example, the low-level, square-wave output in Fig 1 appears to suffer from parasitic oscillation. In actuality, the disturbance is typical of that caused by fast digital clocking or switching-regulator noise getting into critical circuit nodes. Plan for parasitic radiative or conductive paths and eliminate them with appropriate layout and shielding.

The Fig 2 display underscores the previous statement. The scope trace shows the output from a gain-of-ten inverter with a 1-k Ω input resistance. The output exhibits severe peaking induced by only 1-pF of parasitic capacitance across the 1-k Ω resistor. The 50 Ω -terminated input source provides only 20 mV of drive, but that's more than enough to cause problems, even with only 1 pF of stray coupling. In this case the solution was a ground-referred shield at a right angle to, and encircling, the 1-k Ω resistor. Plan for parasitic radiative paths and eliminate them with appropriate shielding.

Too low a gain can cause problems

A decompensated amplifier running at too low a gain produced the oscillation shown in Fig 3. The penalty for a decompensated amplifier's increased speed is a restriction on the minimum allowable gain. Decompensated amplifiers are simply not stable below some (specified) minimum gain, and no amount of wishful thinking will change this. Such oversight is common with these devices, although the amplifier never fails to remind the user. Observe gain restrictions when using decompensated amplifiers.

Text continued on pg 140




Fig 1—This square-wave output disturbance is typical of switching-regulator noise at critical circuit nodes. Poor layout is the culprit.




Fig 2—The peaking and ringing at the output of a 10 \times amplifier is the result of only 1 pF of stray capacitance across a 1-k Ω input resistor.




Fig 3—A decompensated amplifier running at low gain produced this oscillation. Such amplifiers are not stable below a specified minimum gain.

The oscillation problem and frequency compensation

All feedback systems have the propensity to oscillate. Basic theory tells us that an oscillator requires both gain and phase shift. Unfortunately, feedback systems, such as operational amplifiers, also have gain and phase shift. The close relationship between oscillators and operational amplifiers requires careful attention when designing op-amp circuits. In particular, excessive input-to-output phase shift can cause the amplifier to oscillate when you apply feedback. Furthermore, any time delay placed in the amplifier's feedback path introduces additional phase shift, increasing the likelihood of oscillation. This phase shift is why feedback-loop-enclosed stages can cause oscillation.

A large body of complex mathematics describing stability criteria is available, and is useful in predicting the stability characteristics of feedback amplifiers. For sophisticated applications, this complex approach is essential for optimum performance. However, comparatively little information is available that discusses, in practical terms, how to understand and address the issues of compensating feedback amplifiers. The following is a practical approach to stabilizing various combinations of amplifier and power-gain stages, although the considerations are also useful to other feedback systems.

Two categories exist

Oscillation problems in amplifier/power-booster combinations fall into two broad categories: local and loop oscillations. *Local* oscillations can occur in the booster stage, but should not appear in the IC op amp, which presumably was debugged prior to sale. These oscillations are the result of transistor parasitics, layout problems, or circuit-configuration

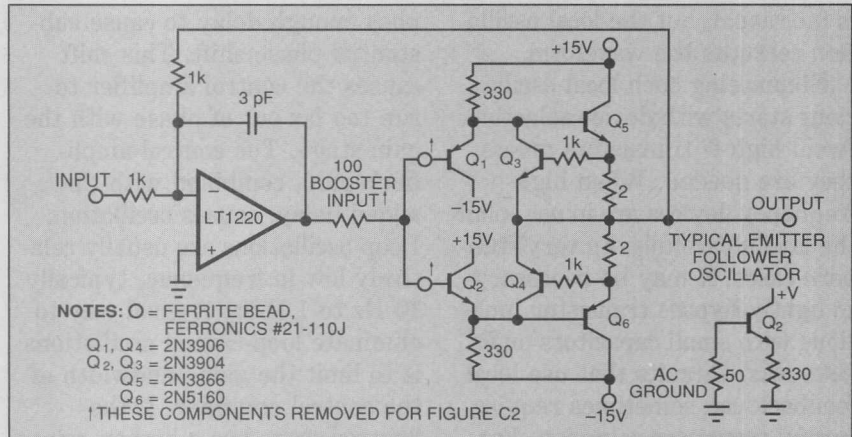


Fig A—In this booster circuit, the 100Ω resistor and the ferrite beads at the inputs of Q₁ and Q₂ play a critical role in maintaining stability.

instabilities. The oscillations are usually relatively high in frequency, typically in the 0.5- to 100-MHz range. Usually, local booster-stage oscillations do not cause loop disruption. The major loop continues to function, but contains artifacts of the local oscillation.

Fig A furnishes an instructive example. The Q₁ and Q₂ emitter-follower pair has a reasonably high f_T . These devices will oscillate if driven from a low-impedance source (Refs 1 and 2). To prevent problems, the 100Ω resistor and the ferrite beads are included to make the op amp's output look like a higher impedance. Q₅ and Q₆, also emitter followers, have an even higher f_T , but 330Ω sources drive them, eliminating the oscillation problem. The **Fig B** photo shows the

action of the **Fig A** circuit without the 100Ω resistor and the ferrite beads. Trace A is the input, and Trace B is the output. The resultant high-frequency oscillation is typical of locally caused disturbances. Note that the major loop

Text continued on pg 138




Fig B—Removing the 100Ω resistor and the ferrite beads from the **Fig A** circuit results in local oscillations, such as those shown here.

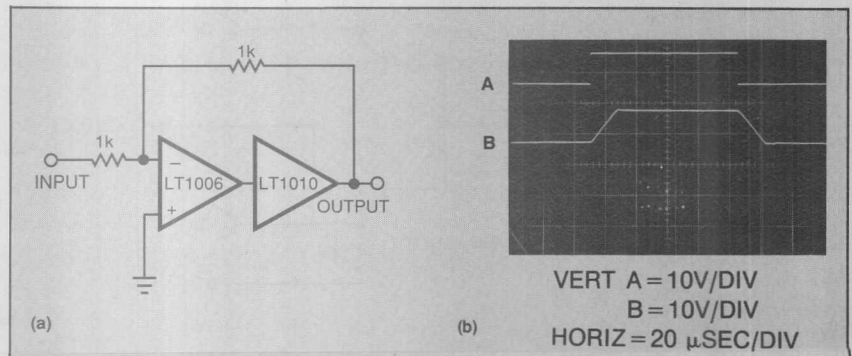


Fig C—This slow op amp and medium-speed booster (a) produced the stable output shown in b.

The oscillation problem and frequency compensation (continued)

is functional, but the local oscillation corrupts the waveform.

Eliminating such local oscillations starts with device selection. Avoid high f_T transistors unless they are needed. When high-frequency devices are in use, plan the layout carefully. In very stubborn cases, it may be necessary to lightly bypass transistor junctions with small capacitors or RC networks. Circuits that use local feedback can sometimes require careful transistor selection. For example, transistors operating in a local loop may require different f_T s to achieve stability. Emitter followers are notorious sources of oscillation, and should never be directly driven from low-impedance sources.

Loop oscillations are caused when the added gain stage sup-

plies enough delay to cause substantial phase shift. This shift causes the control amplifier to run too far out of phase with the gain stage. The control amplifier's gain, combined with the added delay, causes oscillation. Loop oscillations are usually relatively low in frequency, typically 10 Hz to 1 MHz. A good way to eliminate loop-caused oscillations is to limit the gain-bandwidth of the control amplifier. If the booster stage has a higher gain-bandwidth than the control amplifier, its phase delay is easily accommodated in the loop. When the control amplifier's gain-bandwidth dominates, oscillation is ensured. Under these conditions, the control amplifier hopelessly tries to servo-control a feedback signal that consistently arrives

too late, and the oscillation centers around the ideal servo point.

Frequency response roll-off of the control amplifier will almost always cure loop oscillations. In many situations, it is preferable to brutally force the compensation using large capacitors in the major feedback loop. As a general rule, it is wise to stabilize the loop by rolling-off the control amplifier's gain-bandwidth. The feedback capacitor serves to trim only the step response; do not rely on it to stop outright oscillation.

Fig C illustrates these issues. The LT1006 amplifier used with the LT1010 current buffer produces the output shown in Fig Cb. As before, trace A is the input, and trace B is the output. The LT1006 has less than 1-MHz gain-bandwidth. The LT1010's 20-MHz gain-bandwidth introduces negligible loop delay, and dynamics are clean. In this case, the LT1006's internal roll-off is well below that of the output stage, and stability is achieved with no external compensation components. Fig Da uses the LT1223, which has a 100-MHz bandwidth, as the control amplifier. Fig Db shows the results. Here, the control amplifier's roll-off is well beyond the output stages, causing problems. The phase shift through the LT1010




Fig D—This fast op amp and medium-speed booster (a) produced the oscillations shown in b.

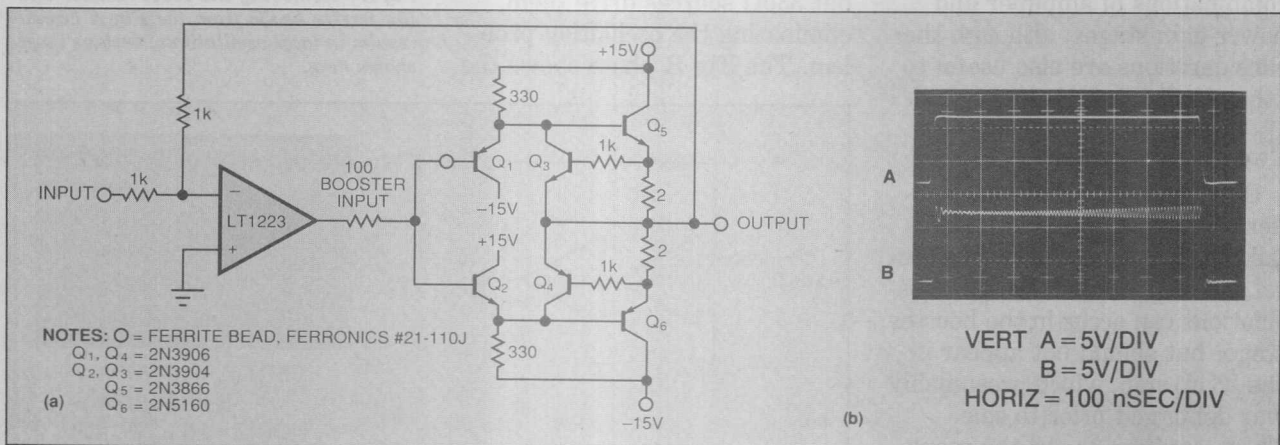


Fig E—The fast amplifier and fast booster combination in a attempts to correct the problems in the Fig Da circuit. Although the result is an improvement, the 100-MHz oscillation indicates that the booster stage is still too slow for the op amp b.

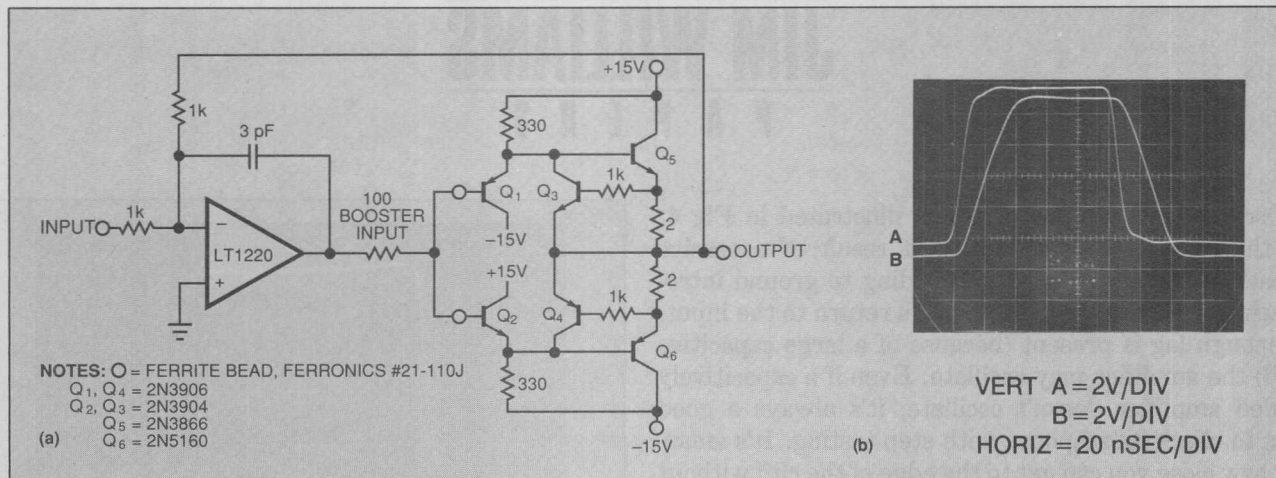


Fig F—This circuit (a) produces more pleasing results than Fig Ea's circuit. The 45-MHz LT1220 replaces the 100-MHz LT1223. The slower amplifier now works well with the booster stage in its loop. The circuit is well controlled, with no sign of oscillation (b).

is now appreciable and oscillations occur. Stabilizing this circuit requires degenerating the control amplifier's gain-bandwidth.

The fact that the slower op-amp circuit doesn't oscillate is a key to understanding how to compensate booster loops. With the slow device, compensation is free. The faster amplifier requires roll-off components for stability. Practically, the LT1223's speed is simply too fast for the LT1010. A somewhat slower amplifier is the way to go. Alternatively, you could use a faster booster. The circuit of Fig Ea attempts using this faster booster, but doesn't quite make it. Although the result (Fig Eb) is less corrupted than before, the 100-MHz oscillation indicates that the booster stage is still too slow for the LT1223.

Nearly identical to Fig Ea, the Fig Fa circuit produces more pleasant results. Here, the 45-MHz LT1220 replaces the 100-MHz LT1223. The slower amplifier, combined with minimum (3 pF) local compensation, works well with the booster stage in its loop. The result (Fig Fb) is a high-speed output that is well controlled, with no sign of oscillations.




Fig G—This current source uses a 40-MHz LT1194 with a gain of 10 and a 50-MHz LT1190. The 100-pF feedback capacitor ensures a fast, stable loop.

Power boosters are not the only things you can place within an amplifier's feedback loop. The Fig G current source is an interesting variation. There is no power booster in the loop, but rather a 40-MHz differential amplifier with a gain of 10. For stability, the circuit uses the 50-MHz LT1190. The local 100-pF feedback slows it down a bit more, and the loop is fast and stable. What happens if you remove the 100-pF feedback path? Fig H shows that the loop is no longer stable, because the LT1190 control amplifier cannot servo-control the phase-shifted feedback at the higher frequency. So, put that 100-pF capacitor back in.

Even if you broke the output-input connection between the LT1190 and the LT1194 and in-




Fig H—Removing the 100-pF feedback capacitor from the Fig G circuit subjects the op amp to phase-shifted feedback, causing the oscillation shown here. Put that capacitor back in.

serted a booster stage, the circuit would still be stable—if you retained the 100-pF feedback capacitor. This tells us that the control amplifier doesn't care what generates the causal feedback between its input and output, as long as there isn't excessive delay.

When compensating loops like these, remember to investigate the effects of various loads and operating conditions. Sometimes a compensation scheme that appears to be proper gives bad results for some conditions. For this reason, check the completed circuit over as wide a variety of operating conditions as possible.

THE
JIM WILLIAMS
P A P E R S

Oscillation is also the problem illustrated in Fig 4. In this case, the oscillation is a result of excessive capacitive loading. Capacitive loading to ground introduces a lag in the feedback signal's return to the input. If enough lag is present (because of a large capacitive load) the amplifier may oscillate. Even if a capacitively loaded amplifier doesn't oscillate, it's always a good idea to check its response with step testing. It's amazing how close you can get to the edge of the cliff without falling off, except when you build 10,000 production units. Avoid capacitive loading. If such loading is necessary, check performance margins, and isolate or buffer the load.

The Fig 5 waveform appears to be one cycle of oscillation. The output initially responds, but abruptly reverses direction, overshoots, and then heads positive again. Some overshoot again occurs, with a long tail and a small dip well before a nonlinear slew returns the waveform to zero. Ugly overshoot and tailing complete the cycle. This is certainly strange behavior, making you wonder what is going on. The input pulse is responsible for all these anomalies. The pulse's amplitude takes the amplifier outside its common-mode limits, inducing the bizarre effects shown. Keep inputs within the specified common-mode limits at all times.

Fig 6 shows an oscillation-laden output (trace B) trying to unity-gain invert the input (trace A). The input's form is distinguishable in the output, but corrupted with very high-frequency oscillation and overshoot. In this case, the amplifier includes a booster within its loop to provide increased output current. The disturbances noted are traceable to local instabili-




Fig 5—What appears to be one cycle of oscillation is actually the result of a high-amplitude pulse that exceeds the common-mode range of the amplifier.

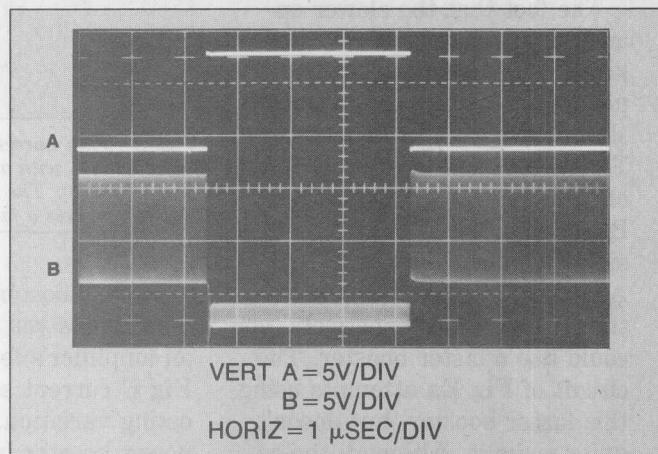


Fig 6—Local oscillation in a booster stage produced this result. Wideband booster stages are prone to high-frequency parasitic oscillation.




Fig 4—The oscillation shown here is the result of excessive capacitive loading, which causes a lag in the feedback to the input of the amplifier.

ties within the booster circuit. When using output booster stages, make sure they are inherently stable before placing them inside an amplifier's feedback loop. Wideband booster stages are particularly prone to device-level parasitic, high-frequency oscillation.

The booster-augmented, unity-gain inverting op amp in Fig 7 also oscillates, but at a much lower frequency. Overshoot and nonlinear recovery dominate the waveform's envelope. Unlike the previous example, this behavior is not caused by local oscillations within the booster stage. Instead, the booster is simply too slow for the op-amp's feedback loop. The booster introduces enough lag to force oscillation, even as it hopelessly tries to maintain loop closure. Make sure that booster




Fig 7—Loop oscillation in a booster stage produced this result. Note the lower frequency of oscillation compared with the local oscillation shown in Fig 6.




Fig 8—Excessive source impedance produced this serene—but undesired—result.

stages are fast enough to maintain stability when placed in the amplifier's feedback loop.

The serene rise and fall of the Fig 8 trace's pulse is a welcome relief from the oscillatory screaming of the previous examples. Unfortunately, such tranquilized behavior is simply too slow. This waveform is the result of excessive source impedance. The high impedance combines with the amplifier's input capacitance to band-limit the input, and the output reflects this action. Reduce the source impedance to a level that maintains the desired bandwidth, and minimize stray input capacitance.

EDN

References

1. Chessman, M and N Sokol, "Prevent Emitter-Follower Oscillation," *Electronic Design*, June 21, 1976, pg 110.
2. DeBella, G B, "Stability of Capacitively-Loaded Emitter Followers—a Simplified Approach," *Hewlett-Packard Journal* 17, April 1966, pg 15.
3. Addis, John, "Vertical Amplifiers and Engineering," *Analog Circuit Design; Art, Science and Personalities*, Butterworths, 1991.
4. Williams, Jim, "Methods for Measuring Op-Amp Settling Time," *Linear Technology Corporation, Application Note 10*, July 1985.

Acronyms used in this article

ac—Alternating current
 DAC—Digital-to-analog converter
 FET—Field-effect transistor
 UHF—Ultrahigh frequency

5. Tektronix, Inc, "Type 1S1 Sampling Plug-In Operating and Service Manual," *Tektronix, Inc*, 1965.
6. Mulvey, J, "Sampling Oscilloscope Circuits," *Tektronix, Inc, Concept Series*. 1970.
7. Addis, John, "Sampling Oscilloscopes," *Private Communication*, February, 1991.
8. Harvey, Barry, "Take the Guesswork Out of Settling-Time Measurements," *EDN*, September 19, 1985, pg 177.

Author's biography

For more information on this article's author, turn to page 163 in the October 10, 1991, issue.

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

Measuring amplifier settling time

High-resolution measurement of an amplifier's settling time is often necessary. Frequently, a DAC drives the amplifier. Of particular importance is the time required for the DAC-amplifier combination to settle to its final value after an input step. This specification lets you set a circuit's timing margins with confidence that the data produced is accurate. The settling time is the

total length of time from the input-step application until the amplifier's output remains within a specified error-band around the final value.

Fig A shows one way to measure DAC-amplifier settling time. The circuit uses the false-sum-node technique. The resistors and amplifier form a summing network. The amplifier output will step positive when the DAC

moves. During amplifier slewing, the diodes limit the voltage excursion at the oscilloscope probe. The summing node is arranged so that, when settling occurs, the oscilloscope probe voltage should be zero. Note that the resistor divider's attenuation means the probe's output will be one-half the actual settled voltage.

In theory, the Fig A circuit lets

Text continued on pg 142

overload recovery will completely dominate the observed waveform, and all measurements will be meaningless.

One way to achieve reliable settling-time measurements is to clip the incoming waveform in time, as well as amplitude. If you prevent the oscilloscope from seeing the waveform until settling is nearly complete, you can avoid overload problems. Doing this requires placing a switch at the settle circuit's output and controlling it with an input-triggered, variable delay. FET switches are not suitable because of their gate-source capacitance. This capacitance lets gate-drive artifacts corrupt the oscilloscope display, producing confusing readings. In the worst case, gate-drive transients

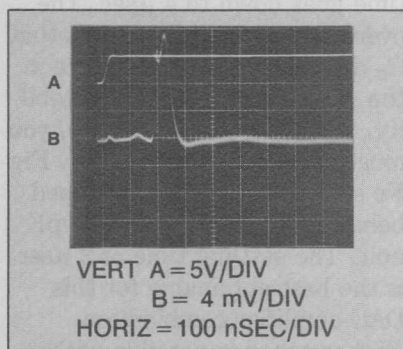


Fig C—The 280-nsec settling time shown here was measured using the Fig B circuit. The sampling switch closes just before the third vertical division, which lets you observe the settling detail without overdriving the oscilloscope.




Fig D—This 280-nsec settling-time measurement was obtained using a sampling scope at Fig B's sampling-scope output. The settling time and waveform is identical to that of Fig C.

will be large enough to induce overload, defeating the switch's purpose.

Fig B shows a way to implement a switch that largely eliminates these problems. This circuit lets you observe settling within 1 mV. The Schottky sampling bridge is the actual switch. The bridge's inherent balance, combined with matched diodes and very high-speed complementary bridge switching, yields a clean, switched output. An output-buffer stage unloads the settle node and drives the diode bridge.

The operation of the DAC-amplifier is as before. The additional circuitry provides the delayed switching function, eliminating oscilloscope overdrive. Buffering the settle node and driving the Schottky bridge is the EL2004, a unity-gain broadband FET-input buffer that has a 3-pF input capacitance and a 350-MHz bandwidth. The pulse generator's output fires the 74123 1-shot circuit. The arrangement of the 74123 produces a delayed pulse whose width sets the on-time of the diode bridge. The 20-k Ω potentiometer controls the pulse delay; the 5-k Ω potentiometer controls the pulse width. If you set the delay appropriately, the scope will not see any input until settling is nearly complete, eliminating overdrive. You adjust the width of the sample window so that all remaining settling activity is observable. In this way, the oscilloscope's output is reliable, and you can take meaningful data. Q_1 through Q_4 shift the level of the 1-shot's output, providing complementary switching drive to the bridge. The actual switching transistors, Q_1 and Q_2 , are UHF types, permitting true differential bridge-switching with less than 1 nsec of time skew (Ref 4).

Using an oscilloscope having adequate sensitivity, you can observe the output of the bridge di-

rectly or you can look at the output of the LT1222, which provides a 10 \times amplified version. A third output, taken directly from the EL2004, is also available. This output, which bypasses the entire switching circuitry, provides a monitoring point for a sampling oscilloscope. Because of their operating nature, sampling oscilloscopes are inherently immune to overload (Refs 5, 6, and 7). A good test of this settling-time test fixture (and the above statement) is to compare the signals displayed by the sampling scope and the Schottky-bridge-aided real-time scope.

As an additional test, you can employ a completely different (albeit considerably more complex) method of measuring settling time, described by Barry Harvey (Ref 8). All three approaches represent good measurement techniques, and if you use proper construction, the results should be identical. That is, the data produced by the three methods has a high probability of being valid.

Figs C, D, and E illustrate settling time details of an AD565A DAC and an LT1220 op amp. The photos represent the sampling bridge, sampling scope, and "Harvey" methods, respectively. Photos **Figs C and E** display the input step for convenience in ascertaining the elapsed time. Photo **Fig D**, taken with a single-trace sampling oscilloscope (Tektronix 1S1 with a P6032 cathode-follower probe in a 556 main-frame), uses the left-most vertical graticule line as its zero-time reference. All methods agree on a 280-nsec settling time to 0.01% (1 mV on a 10V step). Note that Harvey's method inherently adds 30 nsec, which you must subtract from the displayed 310-nsec to get the real number. Note also that the shape of the settling waveform—in every detail—is

Text continued on pg 144

Measuring amplifier settling time (continued)

identical in all three photographs. This kind of agreement provides a high degree of credibility for the measured results.

Some poorly designed amplifiers exhibit a substantial "thermal tail" after responding to an input step. This phenomenon, caused by die heating, can cause the output to wander outside desired limits long after settling has apparently occurred. After checking the settling at high speed, it is always a good idea to slow down the oscilloscope sweep and look for thermal tails. Often, you can accentuate the thermal tail's effect by loading the amplifier's output. Such a tail can make an amplifier appear to have settled in a much shorter time than it actually has.

Select the feedback capacitor

To get the best possible settling time from any amplifier, you should choose the feedback capacitor, C_F , carefully. The purpose of C_F is to roll off the amplifier gain at the frequency that permits the best dynamic response. The optimum value for C_F will depend on the feedback resistor's value and the characteristics of the source. DACs are one of the most common sources and also one of the most difficult. Usually, you must convert a

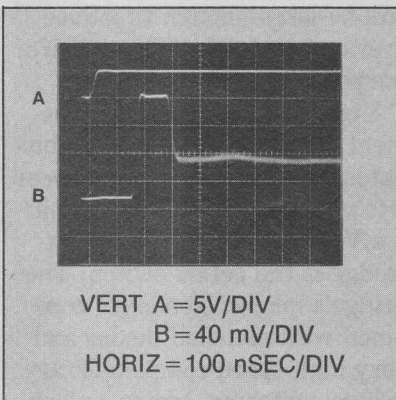


Fig E—Harvey's method was used to obtain this 280-nsec settling-time measurement. After subtracting this method's inherent 30-nsec delay, the settling time and waveform are identical to that of Figs C and D.

DAC's current output to a voltage. Although an op amp can easily do this, care is required to obtain good dynamic performance. A fast DAC can settle to 0.01% in 200 nsec or less, but its output also includes a parasitic capacitance term, making the amplifier's job more difficult. Normally, the circuit unloads the DAC's current output directly into the amplifier's summing junction, placing the parasitic capacitance from ground to the amplifier's input. The capacitance introduces feedback phase-shift at high frequencies, forcing the amplifier to hunt and ring about the final value before settling.

Different DACs have different values of output capacitance. CMOS DACs have the highest output capacitance, which varies with code. Bipolar DACs typically have 20- to 30-pF of capacitance, which is stable over all codes. Because of their output capacitance, DACs furnish an instructive example in amplifier compensation.

Fig Fa shows the response of an industry-standard DAC-80 type and a relatively slow op amp. Trace A is the input, and traces B and C are the amplifier and settle outputs, respectively. In this example, there is no feedback capacitor (C_F), and the amplifier rings badly before settling. In **Fig Fb**, an 82-pF unit stops the ringing and settling time goes down to 4 μ sec. The overdamped response means that C_F dominates the capacitance at the AUT's input, assuring stability. For the fastest response, you must reduce the value of C_F . **Fig Fc** shows the critically damped behavior obtained with a 22-pF unit. The settling time of 2 μ sec is the best obtainable for this DAC-amplifier combination. Higher speed is possible with faster amplifiers and DACs, but the compensation issues remain the same.

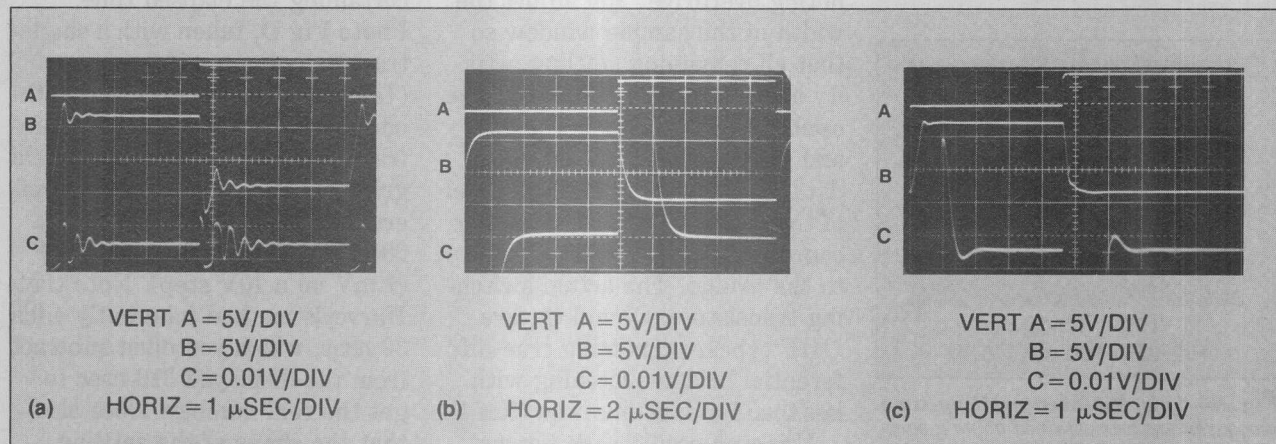


Fig F—These three scope traces show the effects of different values of feedback capacitors on DAC-amplifier combinations. The values for the three photos are no capacitance (a), 82 pF (b), and 22 pF (c).