Does Your Op Amp Oscillate?

Barry Harvey, Staff Design Engineer, Linear Technology® Corp.

Well, it shouldn't. We analog designers take great pains to make our amplifiers stable when we design them, but there are many situations that cause them to oscillate in the real world. Various types of loads can make them sing. Improperly designed feedback networks can cause instability. Insufficient supply bypassing can offend. Finally, inputs and outputs can oscillate by themselves as one-port systems. This article will address common causes of oscillation and their remedies.

Some Basics

Figure 1a shows the block diagram of a non-rail-to-rail amplifier. The inputs control the g_m block which drives the gain node and is buffered at the output. The compensation capacitor C_C is the dominant frequency response element.
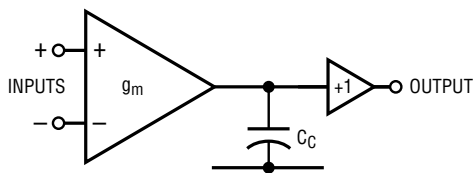

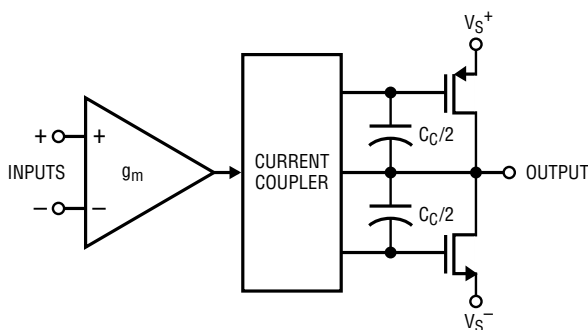**Figure 1A Typical Non-Rail-To-Rail Op Amp Topology**

The return of C_C would go to ground if there were such a pin; however op amps traditionally have no ground and the capacitor current will return to one or both supplies. Figure 1b is a block diagram of the simplest amplifier with rail-to-rail output. The input g_m's output current is sent through a current coupler that splits the drive current between the output transistors. Frequency response is dominated by the two $C_C/2$s, which are effectively in parallel. These two topologies describe the vast majority of op amps that use external feedback. Figure 1c shows the frequency responses of our ideal amplifiers, which display similar behavior although they are electrically different. The single-pole compensation created by g_m and C_C gives a unity-gain-bandwidth product frequency of GBF = $g_m / (2\pi C_C)$. The phase lag of these amplifiers drops from –180 to –270° around GBF/A_{VOL}, where A_{VOL} is the open-loop amplifier DC gain. The phase hangs in at –270° for frequencies well above this frequency. This is known as dominant pole compensation, where the C_C pole dominates the response and hides various frequency limitations of the active circuitry.
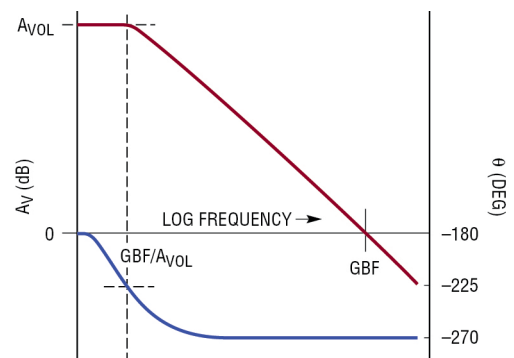
Figure 1B Typical Rail-To-Rail Op Amp Topology

Figure 1C Idealized Frequency Response of Op Amps

Figure 2 shows the open-loop gain and phase response over frequency for the LTC®6268 amplifier. This is a neat little low-noise 500MHz amplifier with rail-to-rail outputs and only 3fA bias current, and is a good example of real amplifier behavior. The dominant compensation's –90° phase lag starts at about 0.1MHz, reaches –270° at about 8MHz, but moves past –270° beyond 30MHz. In practice, all amplifiers have high-frequency phase lags additional to the basic dominant compensation lag due to extra gain stages and the output stage. Typically, the extra phase lag starts at around GBF/10.
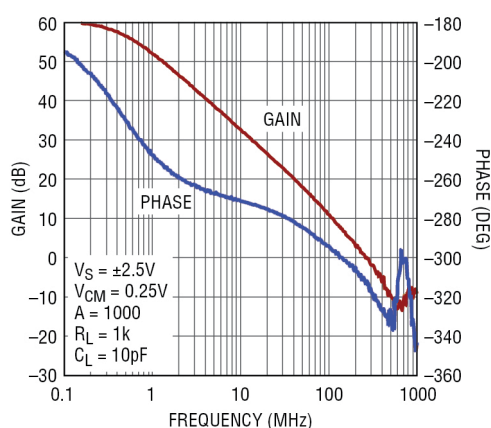

Figure 2. LTC6268 Gain and Phase vs. Frequency

Stability with feedback is a matter of 'round the loop gain and phase; or A_{VOL} times the feedback factor, loop gain in short. If we connect the LTC6268 in unity-gain configuration, then 100% of the output voltage is fed back. At very low frequencies the output is the negative of the –input, or –180° phase lag. Compensation adds another –90° lag through the amplifier, setting –270° lag from –input to output. Oscillation will occur when the loop phase lag increases to ±360°, or multiples of it, and the loop gain is at least 1V/V or 0dB. The phase margin is a measure of how far from 360° the phase lag is when the gain is 1V/V or 0dB. Figure 2 shows us that the phase margin is about 70° (10pF red curve) at 130MHz. This is a very healthy number; phase margin down to perhaps 35° is usable.

A less popular topic is gain margin, although it is just as important a parameter. When the phase descends to zero margin at some high frequency, the amplifier will oscillate if the gain is at least 1V/V or 0dB. As shown in Figure 2, when the phase drops to 0 (or multiples of 360°, or –180 as in the figure), the gain is about –24dB at approximately 1GHz. This is a very low gain; no oscillations could occur at this frequency. In practice one wants at least 4dB of gain margin.

Decompensated Amplifiers

While the LTC6268 is quite stable at unity gain, there are a few op amps that are intentionally not. By designing the amplifier compensation to be stable only at higher closed-loop gains, design trade-offs can deliver higher slew-rate, wider GBF, and lower input noise than unity-gain-compensated. Figure 3 shows the open-loop gain and phase of the LT®6230-10. This amplifier is intended to be used at a fed-back gain of 10 or higher, so the feedback network will attenuate the output by at least 10. With this feedback network we look for the frequency where the open-loop gain is 10V/V or 20dB and find a phase margin of 58° at 50MHz (±5V supplies). At unity gain, our phase margin is just about 0°, and the amplifier will oscillate.
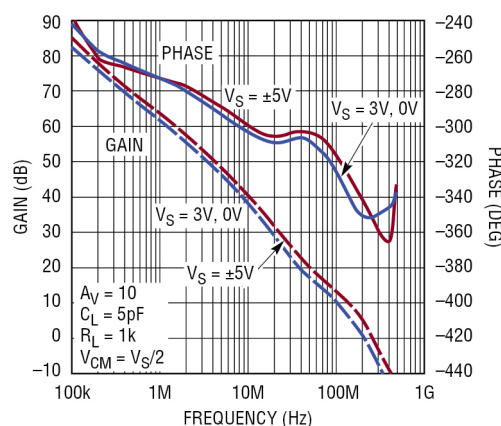

Figure 3. LT6230-10 Gain And Phase vs. Frequency

an148fa

One observation is that all amplifiers will be more stable when providing more closed-loop gain than the minimum stable gain. Even a gain of 1.5 makes a unity-gain-stable amplifier much more stable.

Feedback Networks

While on the subject, the feedback network itself can induce oscillation. Note in Figure 4 we have placed a parasitic capacitance in parallel with the feedback divider. This is inevitable; each terminal of each component on a circuit board has about 0.5pF to ground, plus that of the traces. In practice nodes have minimum capacitance of 2pF, and ~2pF per inch of trace. It is easy to rack up 5pF of parasitic. Consider the LTC6268 providing a gain of +2. Attempting to save power, we set the values of R_F and R_G to a rather high 10kΩ. With C_{PAR} = 4pF, the feedback network has a pole at $1/(2\pi \cdot R_F||R_G \cdot C_{PAR})$, or 8MHz.

Using the fact that the phase lag of the feedback network is –atan(f/8MHz), we can estimate that 360° lag around the loop will be at about 35MHz, where the amplifier has

a lag of –261° and the feedback network lags –79°. At this phase and frequency, the amplifier still has 22dB of gain while the divider gain is

$$0.5/\sqrt{1+\left(f/8MHz\right)^2} = 0.1114 \text{ or } -19dB.$$

The amplifier's 22dB times the feedback –19dB yield a loop gain of +3dB at 0° phase and the circuit oscillates. One must size down the feedback resistor value to work with parasitic capacitance so that the feedback pole is well beyond the loop's unity-gain frequency. At least a 6× pole to GBF ratio is good.

Op amp inputs themselves can be fairly capacitive, emulating C_{PAR}. In particular low-noise and low-V_{OS} amplifiers have large input transistors and may have larger input capacitance than other amplifiers, which loads their feedback networks. You need to consult the data sheet to see how much more capacitance will be in parallel with C_{PAR}. Fortunately, the LT6268 has only 0.45pF, a very low value for such a low-noise amplifier. The circuit with parasitics can be simulated using Linear Technology's macro-model running on LTspice®, which is free.
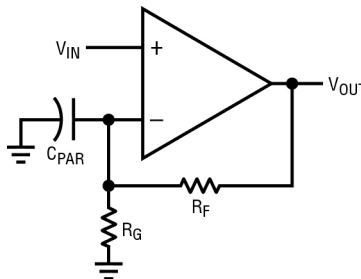

Figure 4. Parasitic Capacitance Loading the Feedback Network

Application Note 148

Figure 5 shows methods to make the divider more tolerant to capacitance. Figure 5a shows a non-inverting amplifier arrangement with R_{IN} added. Assuming V_{IN} is a low impedance source ($\ll R_{IN}$), R_{IN} will effectively attenuate the feedback signal without changing the closed loop gain. R_{IN} will also lower the impedance of the divider, increasing the feedback pole frequency, hopefully beyond GBF. The 'round the loop bandwidth is reduced by R_{IN}, and the input offset and noise are magnified by it.

Figure 5b shows an inverting configuration. R_G performs the loop attenuation again without changing the closed-loop gain. In this case, the input impedance is not disturbed by R_G, but noise, offset, and bandwidth are worsened.

Figure 5c shows the preferred method to compensate C_{PAR} in a non-inverting amplifier. If we set $C_F \cdot R_F = C_{PAR} \cdot R_G$ we have a compensated attenuator such that the feedback divider now has the same attenuation at all frequencies and the C_{PAR} problem is solved. Mismatch in the products will cause "bumps" in the pass band of the amplifier and "shelves" in the response where low frequency responses are flat but

change to another flat level around $f = 1/2\pi \cdot C_{PAR} \cdot R_G$. Figure 5d shows the equivalent C_{PAR} compensation for inverting amplifiers. The frequency response needs to be analyzed to find a proper C_F, and the bandwidth of the amplifier is part of the analysis.

Some comments on current-feedback amplifiers (CFAs) are in order here. If the amplifier in Figure 5a were a CFA, then R_{IN} would do little to modify frequency response, since the –input is a very low impedance and actively copies the +input. Noise would be somewhat worsened, and additional –input bias current would effectively occur as V_{OS}/R_{IN}. Likewise, the circuit in Figure 5b is unaltered by R_G with respect to frequency response. The inverting input is not just a virtual ground, it's a true low impedance to ground and is already tolerant of C_{PAR} (inverting mode only!). DC errors are similar to those shown in Figure 5a. Figures 5c and 5d may be preferred for voltage-input op amps, but CFAs simply cannot tolerate a direct feedback capacitor without oscillation.
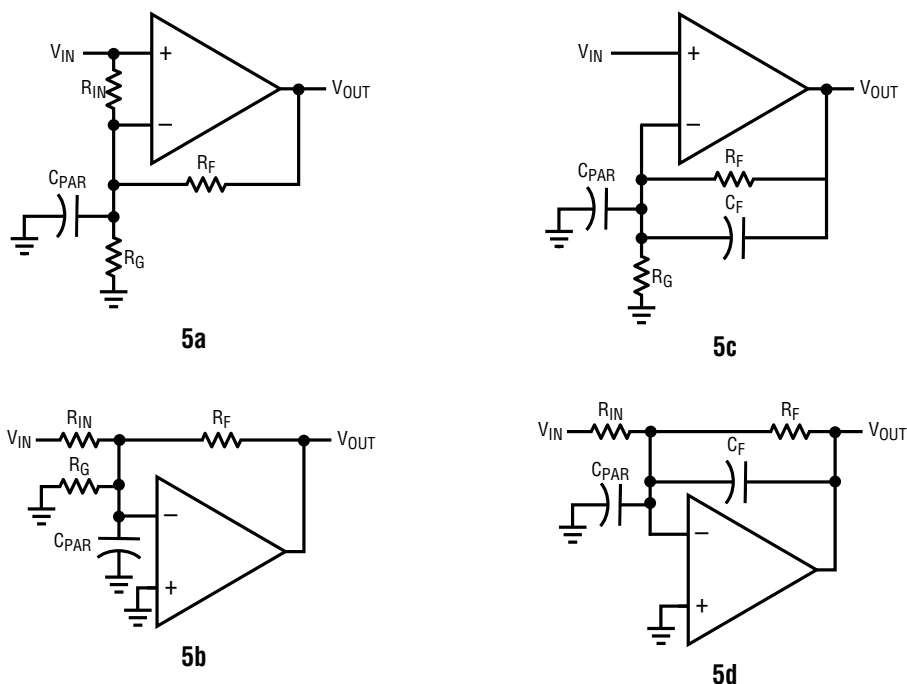

Figure 5. Methods of Reducing C_{PAR} Effect

an148fa

AN148-4

Load Issues

Just as feedback capacitance can erode phase margin, so too can load capacitance. Figure 6 shows the LTC6268 output impedance over frequency for a few gain settings. Note that the unity-gain output impedance is lower than that of higher gains. The full feedback allows the open-loop gain to reduce the inherent output impedance of the amplifier. Thus the gain of 10 output impedance in Figure 6 is generally 10× higher than the unity-gain results. There is 1/10th the 'round the loop gain due to the feedback attenuator diminishing loop gain that otherwise would reduce the closed-loop output impedance. The open-loop output impedance is about 30Ω, made obvious by the flat region of the gain-of-100 curve at high frequencies. In this region, from about gain-bandwidth-frequency/100 to gain-bandwidth-frequency, there is not enough loop gain to reduce the open-loop output impedance.
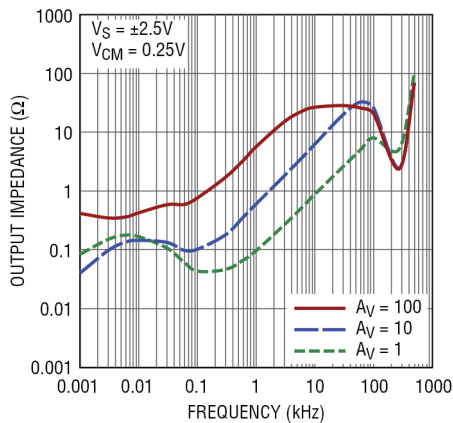

Figure 6. LTC6268 Output Impedance vs. Frequency for Three Gains

Capacitor loads will cause phase and amplitude lags with the open-loop output impedance. For instance, a 50pF load with our LTC6268 30Ω output impedance makes another pole at 106MHz, where the output has a phase lag of −45° and −3dB attenuation. At this frequency, the amplifier has a phase of −295° and gain of 10dB. Assuming a unity-gain feedback, we do not quite achieve oscillation because the phase has not made it to ±360° (at 106MHz). At 150MHz, however, the amplifier has 305° lag and 5dB gain. The out-put pole has a phase of $-a\tan(150\text{MHz}/106\text{MHz}) = -55°$ and a gain of

$$\frac{1}{\sqrt{1+\left(150\text{MHz}\Big/106\text{MHz}\right)^2}} = 0.577$$

or −4.8dB. Multiplying gains around the loop, we get 360° and +0.2dB gain, again an oscillator. 50pF appears to be the minimum load capacitance that will force the LTC6268 to oscillate.

The most common way to prevent load capacitance from causing oscillation is to simply place a small-value resistor in series with it, after the feedback connection. Values of 10Ω to 50Ω will limit the phase lag capacitive loads can induce, as well as isolate the amplifier from low capacitive impedances at very high speed. Drawbacks include DC and low frequency errors depending on resistive aspects of the load, limited frequency response at the capacitive load, and signal distortion if the load capacitance is not constant with voltage.

Oscillations caused by load capacitance can often be stopped by raising the amplifier closed-loop gain. Running the amplifier at higher closed-loop gains means that the feedback attenuator also attenuates the loop gain at frequencies where the loop phase is ±360°. For instance, if we use the LTC6268 at a closed-loop gain of +10, we see the amplifier has a gain of 10V/V or 20dB at 40MHz, where the phase lag is 285°. To achieve oscillation we would need an output pole, causing an additional 75° lag. We can solve for the output pole by using

$$-75° = -a\tan(40\text{MHz} / \text{Fpole}) \rightarrow \text{Fpole} = 10.6\text{MHz}.$$

This pole frequency comes from a load capacitance of 500pF and 30Ω output impedance. The output pole gain is

$$\frac{1}{\sqrt{1+\left(40\text{MHz}\Big/10.6\text{MHz}\right)^2}} = 0.026.$$

With an unloaded open-loop gain of 10 we have a 'round the loop gain of 0.26 at the oscillation frequency, so this

time we do not get oscillation, at least oscillation caused by a simple output pole. Thus we have increased load capacitance tolerated from 50pF to 500pF by raising the closed-loop gain.

Unterminated transmission lines are also very bad loads, since they present wild impedance and phase changes repetitive with frequency (see the impedance of an unterminated 9' cable in Figure 7). If your amplifier can drive the cable safely at one low-frequency resonance, it's likely to oscillate at some higher frequency as its own phase margin degrades. If the cable must be unterminated, then a back-match resistor in series with the output can isolate the cable's radical impedance variations. Further, even though transient reflections from the unterminated end of the cable rush right back to the amplifier, the back-match

resistor correctly absorbs the energy if its value matches the characteristic impedance of the cable. If the back-match doesn't match the cable impedance, some energy will reflect from the amplifier and termination, running back down to the unterminated end. When the energy hits the end, it efficiently reflects back to the amplifier yet again, and we have a series of pulses bouncing back and forth, but attenuated each time.

Figure 8 shows a more complete output impedance model. The term R_{OUT} is the same 30Ω we have been discussing in the LTC6268, and we have added a term L_{OUT}. This is a combination of physical inductance and an electronic equivalent of inductance. The physical package, bond-wire, and external inductances add up to 5nH to 15nH, smaller with smaller packages. In addition, there is an electroni-
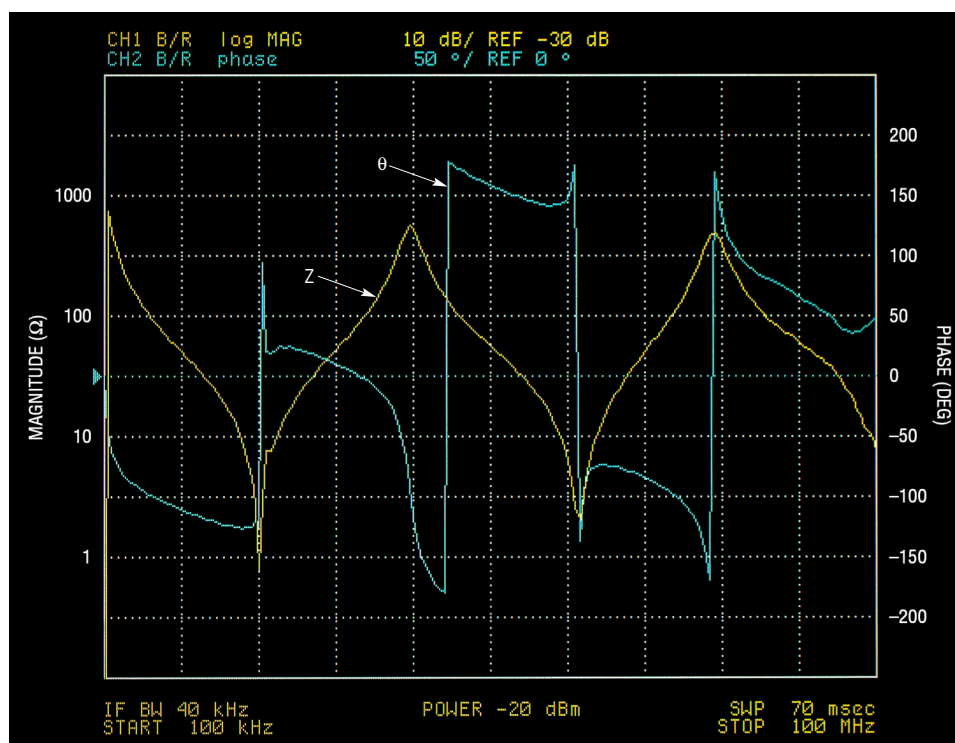

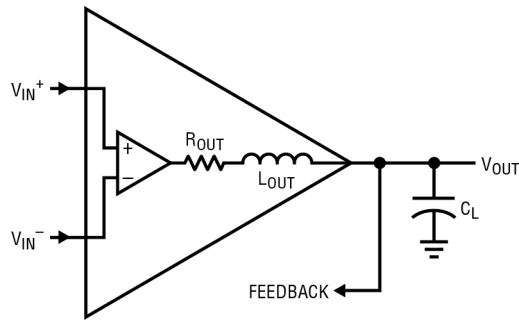**Figure 7. Impedance And Phase of an Unterminated Coaxial Cable**

Figure 8. Inductive Part of Amplifier Output Impedance

cally generated inductance range of 20nH to 70nH for most amplifiers, especially devices using bipolar output devices. The parasitic base resistance of the output transistor is transformed to inductance by the finite Ft of the devices. The hazard is that L_{OUT} can interact with C_L to form a series-resonant tuned circuit whose impedance can drop to levels R_{OUT} cannot drive without, yet again, more phase lag within the loop and potential oscillation. For example, set $L_{OUT} = 60nH$ and $C_L = 50pF$. The resonant frequency is

$$\frac{1}{\sqrt{2\pi L_{OUT} \bullet C_L}} = 92MHz \,,$$

well within the passband of the LTC6268. This series-resonant circuit effectively loads the output at resonance and severely modifies loop phase around the resonance. Unfortunately, L_{OUT} is not noted in amplifier data sheets, but one can sometimes see its effect in the open-loop output impedance graphs. In general, this effect is not important for amplifiers with bandwidth less than about 50MHz.

One solution is shown Figure 9. R_{SNUB} and C_{SNUB} create what is called a snubber whose purpose is to de-Q the resonant circuit so that is will not impose a very low resonant impedance at the amplifier output. R_{SNUB} is generally valued at the reactance of C_L at resonance, which is $-j35\Omega$ in this example, to bring the Q of the output resonance down to ~1. C_{SNUB} is sized to insert R_{SNUB} fully at the output resonant frequency, that is, the reactance of $C_{SNUB} << C_L$. $C_{SNUB} = 10 \bullet C_L$ is practical. C_{SNUB} unloads the amplifier at intermediate and low frequencies, especially at DC.

If C_{SNUB} is very large, then the amplifier will be heavily loaded by R_{SNUB} at medium or low frequencies, and gain accuracy, closed-loop bandwidth, and distortion may be compromised. Nevertheless with a bit of tweaking, the snubber is often useful for taming reactive loads, but it must be empirically sized.

A current-feedback amplifier's –input is really a buffer output, and will also exhibit the series features of Figure 8. Thus it can oscillate by itself against a C_{PAR} just like an output. C_{PAR} and any associated inductance must be minimized. Unfortunately, a snubber at the –input will modify
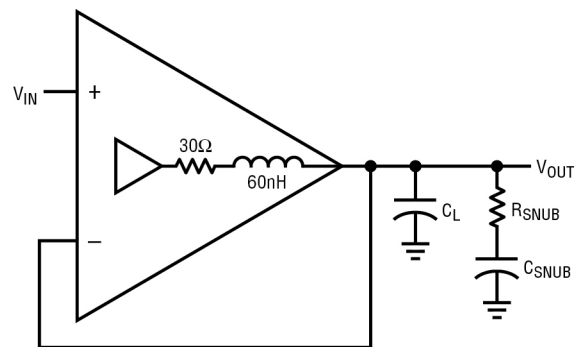

Figure 9. Using an Output Snubber

the closed-loop gain over frequency and is not useful.

Strange Impedances

Many amplifiers exhibit input impedance oddities at high frequencies. This is most true of amplifiers with two input transistors in series, as in a Darlington arrangement. Many amplifiers have an npn/pnp transistor pair at the input which behaves similarly to a Darlington over frequency. There are frequencies, generally well beyond GBF, where the real part of input impedance goes negative. An inductive source impedance will resonate with input and board capacitance, and the negative-real component fuels oscillation. When driven from an unterminated cable, this can also allow oscillation at many repetitive frequencies. If a long inductive line is unavoidable at the input, it can be broken up with a few series energy-absorbing resistors, or a medium-impedance snubber (about 300Ω) can be installed at the amplifier input lead.

Power Supplies

The last source of oscillation to consider is the power supply bypass. Figure 10 shows a fragment of output circuitry. L_{VS}^+ and L_{VS}^- are inevitable series inductances of the package, IC bond-wires, physical length of the bypass capacitor (which is also inductive like any conductor), and circuit board trace inductances. Also included are outer inductance that connects local bypasses with the rest of the power bus, if it's not a power plane. While 3nH to 10nH may not seem like much, it is 3.8Ω to j12Ω at 200MHz. If an output transistor conducts large high-frequency output current, then there will be a drop across its supply inductance.

The rest of the amplifier needs a quiet power supply, because over frequency it cannot reject the supply. In Figure 11 we see the power supply rejection ratio (PSRR) over frequency for the LTC6268. Because the compensation capacitors are associated to the power supplies in all op amps without ground pins, they couple supply noise into the amplifier that the g_m must counteract. The PSRR diminishes as 1/f due to the compensation, and past 130MHz the supply rejection actually becomes gain. With gain in PSRR at 200MHz, output currents can disturb the supply voltages inside the LV_S inductors, which through
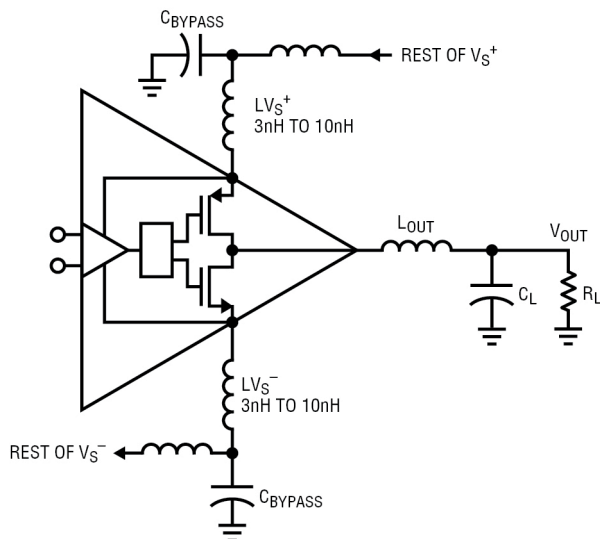
PSRR amplification become strong amplifier signals, driving output currents, creating internal supply signals, etc., causing the amplifier to oscillate. This is why all amplifiers' supplies must be carefully bypassed with low inductance traces and components. Further, the supply bypass capacitors must be much larger than any load capacitance.

If we consider frequencies around 500MHz, our 3nH to 10nH becomes j9.4Ω to j31.4Ω. This is high enough for the output transistor alone to oscillate inside its inductances and IC component capacitances, especially at larger output currents where the transistor g_m and bandwidth increase. Special attention is required since today's semiconductor manufacturing processes employ transistors whose bandwidths are very high, at least at large output currents.
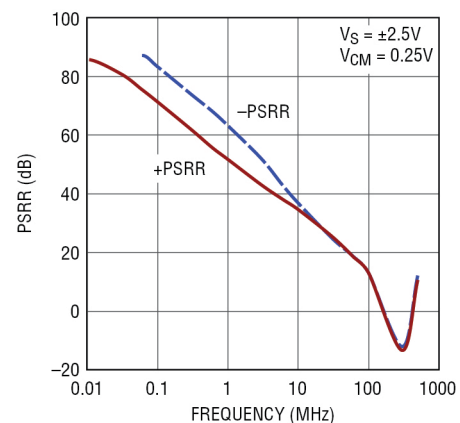
Figure 11. LTC6268 Power Supply Rejection Over Frequency

Conclusion

In summary, the designer needs to consider parasitic capacitance and inductance associated with each op amp terminal, and the nature of the load. The amplifiers are designed to be stable within a nominal environment, but each application requires its own analysis.

References

Jung, Walt, *Op Amp Applications Handbook*, pages 497-499, Newnes, 2005.

Figure 10. Power Supply Bypass Details