

Ask the Applications Engineer—28

By Eamon Nash*

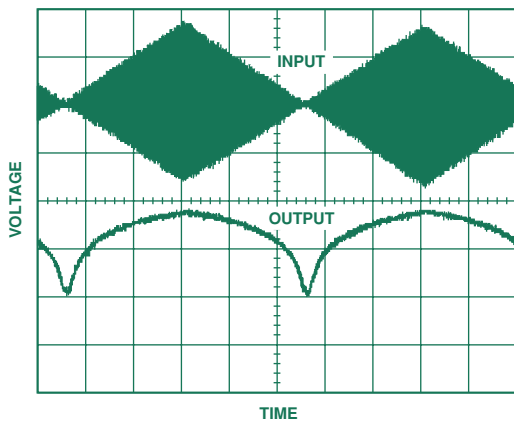
LOGARITHMIC AMPLIFIERS EXPLAINED

Q. I've just been reading data sheets of some recently released Analog Devices log amps and I'm still a little confused about what exactly a log amp does.

A. You're not alone. Over the years, I have had to deal with lots of inquiries about the changing emphasis on functions that log amps perform and radically different design concepts. Let me start by asking you, what do you expect to see at the output of a log amp?

Q. Well, I suppose that I would expect to see an output proportional to the logarithm of the input voltage or current, as you describe in the Nonlinear Circuits Handbook | <<http://www.analog.com/publications/magazines/Dialogue/Anniversary/books.html>> and the Linear Design Seminar Notes | <http://www.analog.com/publications/press/misc/press_123094.html>.

A. Well, that's a good start but we need to be more specific. The term log amp, as it is generally understood in communications technology, refers to a device which calculates the log of an input signal's envelope. What does that mean in practice? Take a look at the scope photo below. This shows a 10-MHz sine wave modulated by a 100-kHz triangular wave and the gross logarithmic response of the AD8307, a 500-MHz 90-dB log amp. Note that the input signal on the scope photo consists of many cycles of the 10-MHz signal, compressed together, using the time/div knob of the oscilloscope. We do this to show the envelope of the signal, with its much slower repetition frequency of 100 kHz. As the envelope of the signal increases linearly, we can see the characteristic "log (x)" form in the output response of the log amp. In contrast, if our measurement device were a linear envelope detector (a filtered rectified output), the output would simply be a tri-wave.




Q. So I don't see the log of the instantaneous signal?

A. That's correct, and it's the source of much of the confusion. The log amp gives an indication of the instant-by-instant low-frequency changes in the envelope, or *amplitude*, of the signal in the log domain in the same way that a digital voltmeter, set to "ac volts," gives a steady (linear) reading when the input is connected to a constant amplitude sine wave and follows any adjustments to the amplitude. A device that calculates the

instantaneous log of the input signal is quite different, especially for bipolar signals.

On that point, let's digress for a moment to consider such a device. Think about what would happen when an ac input signal crosses zero and goes negative. Remember, the mathematical function, $\log x$, is undefined for x real and less than or equal to zero, or $-x$ greater than or equal to zero (see figure).



However, as the figure shows, the inverse hyperbolic sine, $\sinh^{-1} x$, which passes symmetrically through zero, is a good approximation to the combination of $\log 2x$ and minus $\log(-2x)$, especially for large values of $|x|$. And yes, it is possible to build such a log amp; in fact, Analog Devices many years ago manufactured and sold Model 752 N & P temperature-compensated log diode modules, which—in complementary feedback pairs—performed that function. Such devices, which calculate the instantaneous log of the input signal are called *baseband log amps* (the term "true log amp" is also used). The focus of this discussion, however, is on envelope-detecting log amps, also referred to as demodulating log amps, which have interesting applications in RF and IF circuitry for communications.

Q. But, from what you have just said, I would imagine that a log amp is generally not used to demodulate signals?


A. Yes, that is correct. The term demodulating came to be applied to this type of device because a log amp recovers the log of the envelope of a signal in a process somewhat like that of demodulating AM signals.

In general, the principal application of log amps is to measure signal *strength*, as opposed to detecting signal *content*. The log amp's output signal, which can represent a many-decade dynamic range of high-frequency input signal amplitudes by a relatively narrow range, is typically used to regulate gain. The classic example of this is using a log amp in an automatic gain control loop, to regulate the gain of a variable-gain amplifier. The receiver of a cellular base station, for example, might use the signal from a log amp to regulate the receiver gain. In transmitters, log amps are also used to measure and regulate transmitted power.

However, there are some applications where a log amp is used to demodulate a signal. The figure shows a received signal that has been modulated using *amplitude shift keying* (ASK). This simple modulation scheme, similar to early transmissions of radar pulses, conveys digital information by transmitting a series of RF bursts (logic 1 = burst, logic 0 = no burst). When this

*His photo and a brief biography appear in Analog Dialogue 33, page 71.


signal is applied to a log amp, the output is a pulse train which can be applied to a comparator to give a digital output. Notice that the actual amplitude of the burst is of little importance; we only want to detect its presence or absence. Indeed, it is the log amp's ability to convert a signal which varies over a large dynamic range (10 mV to 1 V in this case) into one that varies over a much smaller range (1 V to 3 V) that makes the use of a log amp so appealing in this application.



Q. Can you explain briefly how a log amp works?

A. The figure shows a simplified block diagram of a log amp. The core of the device is a cascaded chain of amplifiers. These amplifiers have linear gain, usually somewhere between 10 dB and 20 dB. For simplicity of explanation, in this example, we have chosen a chain of 5 amplifiers, each with a gain of 20 dB, or 10 \times . Now imagine a small sine wave being fed into the first amplifier in the chain. The first amplifier will amplify the signal by a factor of 10 before it is applied to the second amplifier. So as the signal passes through each subsequent stage, it is amplified by an additional 20 dB.

Now, as the signal progresses down the gain chain, it will at some stage get so big that it will begin to clip (the term *limit* is also used) as shown. In the simplified example, this clipping level (a desired effect) has been set at 1 V peak. The amplifiers in the gain chain would be designed to limit at this same precise level.



After the signal has gone into limiting in one of the stages (this happens at the output of the third stage in the figure), the limited signal continues down the signal chain, clipping at each stage and maintaining its 1 V peak amplitude as it goes.

The signal at the output of each amplifier is also fed into a full wave rectifier. The outputs of these rectifiers are summed together as shown and applied to a low-pass filter, which removes the ripple of the full-wave rectified signal. Note that the contributions of the earliest stages are so small as to be negligible. This yields an output (often referred to as the "video" output), which will be a steady-state quasi-logarithmic dc output for a steady-state ac input signal. The actual devices contain innovations in circuit design that shape the gain and limiting functions to produce smooth and accurate logarithmic behavior between the decade breaks, with the limiter output sum comparable to the *characteristic*, and the contribution of the less-than-limited terms to the *mantissa*.

To understand how this signal transformation yields the log of the input signal's envelope, consider what happens if the input signal is reduced by 20 dB. As it stands in the figure, the unfiltered output of the summer is about 4V peak (from three stages that are limiting and a fourth that is just about to limit). If the input signal is reduced by a factor of 10, the output of one stage at the input end of the chain will become negligible, and there will be one less stage in limiting. Because of the voltage lost from this stage, the summed output will drop to approximately 3 V. If the input signal is reduced by a further 20 dB, the summed output will drop to about 2 V.

So the output is changing by 1 V for each factor-of-10 (20-dB) amplitude change at the input. We can describe the log amp then as having a slope of 50 mV/dB.

Q. O.K. I understand the logarithmic transformation. Now can you explain what the Intercept is?

A. The slope and intercept are the two specifications that define the transfer function of the log amp, that is, the relationship between output voltage and input signal level. The figure shows the transfer function at 900 MHz, and over temperature, of the AD8313, a 100-MHz-to-2.5-GHz 65-dB log amp. You can see that the output voltage changes by about 180 mV for a 10 dB change at the input. From this we can deduce that the slope of the transfer function is 18 mV/dB.


As the input signal drops down below about -65 dBm, the response begins to flatten out at the bottom of the device's range (at around 0.5 V, in this case). However, if the linear part of the transfer function is extrapolated until it crosses the horizontal axis (0 V theoretical output), it passes through a point called the *intercept* (at about -93 dBm in this case). Once the slope and intercept of a particular device are known (these will always be given in the data sheet), we can predict the nominal output voltage of the log amp for any input level within the linear range of the device (about -65 dBm to 0 dBm in this case) using the simple equation:

$$V_{OUT} = Slope \times (P_{IN} - Intercept)$$

For example, if the input signal is -40 dBm the output voltage will be equal to

$$18 \text{ mV/dB} \times (-40 \text{ dBm} - (-93 \text{ dBm})) = 0.95 \text{ V}$$

It is worth noting that an increase in the intercept's value *decreases* the output voltage.



The figure also shows plots of deviations from the ideal, i.e., *log conformance*, at -40°C , $+25^{\circ}\text{C}$, and $+85^{\circ}\text{C}$. For example, at $+25^{\circ}\text{C}$, the log conformance is to within at least ± 1 dB for an input in the range -2 dBm to -67 dBm (over a smaller range, the log conformance is even better). For this reason, we call the AD8313 a 65-dB log amp. We could just as easily say that the AD8313 has a dynamic range of 73 dB for log conformance within 3 dB.

Q. In doing some measurements, I've found that the output level at which the output voltage flattens out is higher than specified in the data sheet. This is costing me dynamic range at the low end. What is causing this?

A. I come across this quite a bit. This is usually caused by the input picking up and measuring an external noise. Remember that our log amps can have an input bandwidth of as much as 2.5 GHz! The log amp does not know the difference between the wanted signal and the noise. This happens quite a lot in laboratory environments, where multiple signal sources may be present. Remember, in the case of a wide-range log amp, a -60 -dBm noise signal, coming from your colleague who is testing his new cellular phone at the next lab bench, can wipe out the bottom 20-dB of your dynamic range.

A good test is to ground both differential inputs of the log amp. Because log amps are generally ac-coupled, you should do this by connecting the inputs to ground through coupling capacitors.

Solving the problem of noise pickup generally requires some kind of filtering. This is also achieved indirectly by using a matching network at the input. A narrow-band matching



network will have a filter characteristic and will also provide some gain for the wanted signal. Matching networks are discussed in more detail in data sheets for the AD8307, AD8309, and AD8313.

Q. What corner frequency is typically chosen for the output stage's low-pass filter?

A. There is a design trade-off here. The corner frequency of the on-chip low-pass filter must be set low enough to adequately remove the ripple of the full-wave rectified signal at the output of the summer. This ripple will be at a frequency 2 times the input signal frequency. However the RC time constant of the low-pass filter determines the maximum rise time of the output. Setting the corner frequency too low will result in the log amp having a sluggish response to a fast-changing input envelope.

The ability of a log amp to respond to fast changing signals is critical in applications where short RF bursts are being detected. In addition to the ASK example discussed earlier, another good example of this is RADAR. The figure on the left shows the response of the AD8313 to a short 100 MHz burst. In general, the log-amp's response time is characterized by the metric 10% to 90% rise time. The table below compares the rise times and other important specifications of different Analog Devices log amps.

Now take a look at the figure on the right. This shows you what will happen if the frequency of the input signal is lower than the corner frequency of the output filter. As might be expected, the full wave rectified signal appears unfiltered at the output. However this situation can easily be improved by adding additional low-pass filtering at the output.



Part Number	Input Bandwidth	10%–90% Rise Time	Dynamic Range	Log Conformance	Limiter Output
AD606	50 MHz	360 ns	80 dB	± 1.5 dB	Yes
AD640	120 MHz	6 ns	50 dB	± 1 dB	Yes
AD641	250 MHz	6 ns	44 dB	± 2 dB	Yes
AD8306	500 MHz	67 ns	95 dB	± 0.4 dB	Yes
AD8307	500 MHz	500 ns	92 dB	± 1 dB	No
AD8309	500 MHz	67 ns	100 dB	± 1 dB	Yes
AD8313	2500 MHz	45 ns	65 dB	± 1 dB	No

Q. I notice that there is an unusual tail on the output signal at the right. What is causing that?


A. That is an interesting effect that results from the nature of the log transformation that is taking place. Looking again at transfer function plot (i.e. voltage out vs. input level), we can see that at low input levels, small changes in the input signal have a significant effect on the output voltage. For example a change in the input level from 7 mV to 700 μ V (or about -30 dBm to -50 dBm) has the same effect as a change in input level from 70 mV to 7 mV. That is what is expected from a logarithmic amplifier. However, looking at the input signal (i.e., the RF burst) with the naked eye, we do not see small changes in the mV range. What's happening in the figure is that the burst does not turn off instantly but drops to some level and then decays exponentially to zero. And the log of a decaying exponential signal is a straight line similar to the tail in the plot.

Q. Is there a way to speed up the rise time of the log amp's output?

A. This is not possible if the internal low-pass filter is buffered, which is the case in most devices. However the figure shows one exception: the unbuffered output stage of the AD8307 is here represented by a current source of 2 μ A/dB, which is looking at an internal load of 12.5 k Ω . The current source and the resistance combine to give a nominal slope of 25 mV/dB. The 5-pF capacitance in parallel with the 12.5-k Ω resistance combines to yield a low-pass corner frequency of 2.5 MHz. The associated 10%-90% rise time is about 500 ns.

In the figure, an external 1.37-k Ω shunt resistor has been added. Now, the overall load resistance is reduced to around 1.25 k Ω . This will decrease the rise time ten-fold. However the overall logarithmic slope has also decreased ten-fold. As a result, external gain is required to get back to a slope of 25 mV/dB.

You may also want to take a look at the Application Note AN-405. This shows how to improve the response time of the AD606.




Q. Returning to the architecture of a typical log amp, is the heavily clipped signal at the end of the gain chain in any way useful?

A. The signal at the end of the linear gain chain has the property that its amplitude is constant for all signal levels within the dynamic range of the log amp. This type of signal is very useful in phase- or frequency demodulation applications. Remember that in a phase-modulation scheme (e.g. QPSK or broadcast FM), there is no useful information contained in the signal's amplitude; all the information is contained in the phase. Indeed, amplitude variations in the signal can make the demodulation process quite a bit more difficult. So the signal at the output

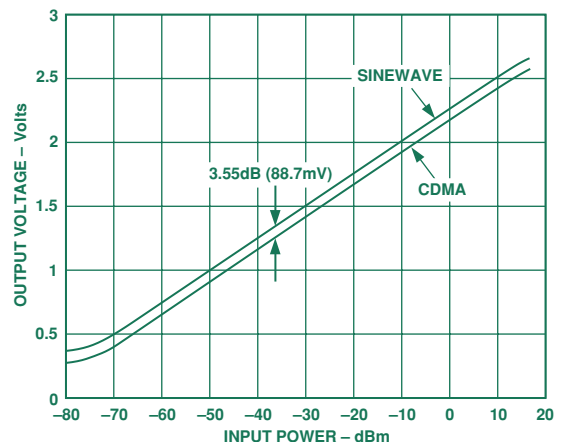
of the linear gain chain is often made available to give a limiter output. This signal can then be applied to a phase or frequency demodulator.

The degree to which the phase of the output signal changes as the input level changes is called *phase skew*. Remember, the phase between input and output is generally not important. It is more important to know that the phase from input to output stays constant as the input signal is swept over its dynamic range. The figure shows the phase skew of the AD8309's limiter output, measured at 100 MHz. As you can see, the phase varies by about 6° over the device's dynamic range and over temperature.



Q. I noticed that something strange happens when I drive the log amp with a square wave.

A. Log amps are generally specified for a sine wave input. The effect of differing signal waveforms is to shift the effective value of the log amp's intercept upwards or downwards. Graphically, this looks like a vertical shift in the log amp's transfer function (see figure), without affecting the logarithmic slope. The figure shows the transfer function of the AD8307 when alternately fed by an unmodulated sine wave and by a CDMA channel (9 channels on) of the same rms power. The output voltage will differ by the equivalent of 3.55 dB (88.7 mV) over the complete dynamic range of the device.



The table shows the correction factors that should be applied to measure the rms signal strength of various signal types with a logarithmic amplifier which has been characterized using a sine wave input. So, to measure the rms power of a squarewave, for example, the mV equivalent of the dB value given in the table (-3.01 dB, which corresponds to 75.25 mV in the case of the AD8307) should be subtracted from the output voltage of the log amp.

Signal Type	Correction Factor (Add to Output Reading)
Sine Wave	0 dB
Square Wave or DC	-3.01 dB
Triangular Wave	+0.9 dB
GSM Channel (All Time Slots On)	+0.55 dB
CDMA Forward Link (Nine Channels On)	+3.55 dB
Reverse CDMA Channel	0.5 dB
PDC Channel (All Time Slots On)	+0.58 dB
Gaussian Noise	+2.51 dB

Q. In your data sheets you sometimes give input levels in dBm and sometimes in dBV. Can you explain why?


A. Signal levels in communications applications are usually specified in dBm. The dBm unit is defined as the power in dB relative to 1 mW i.e.,

$$Power \text{ (dBm)} = 10 \log_{10} (Power/1 \text{ mW})$$

Since power in watts is equal to the rms voltage squared, divided by the load impedance, we can also write this as

$$Power \text{ (dBm)} = 10 \log_{10} ((V_{rms}^2/R)/1 \text{ mW})$$

It follows that 0 dBm occurs at 1 mW, 10 dBm corresponds to 10 mW, +30 dBm corresponds to 1 W, etc. Because impedance is a component of this equation, it is always necessary to specify load impedance when talking about dBm levels.



Log amps, however fundamentally respond to voltage, not to power. The input to a log amp is usually terminated with an external 50-Ω resistor to give an overall input impedance of

approximately 50 Ω, as shown in the figure (the log amp has a relatively high input impedance, typically in the 300 Ω to 1000 Ω range). If the log amp is driven with a 200-Ω signal and the input is terminated in 200 Ω, the output voltage of the log amp will be higher compared to the same amount of power from a 50-Ω input signal. As a result, it is more useful to work with the *voltage* at the log amp's input. An appropriate unit, therefore, would be dBV, defined as the voltage level in dB relative to 1 V, i.e.,

$$Voltage \text{ (dBV)} = 20 \log_{10} (V_{rms}/1 \text{ V})$$

However, there is disagreement in the industry as to whether the 1-V reference is 1 V peak (i.e., amplitude) or 1 V rms. Most lab instruments (e.g., signal generators, spectrum analyzers) use 1 V rms as their reference. Based upon this, dBV readings are converted to dBm by adding 13 dB. So -13 dBV is equal to 0 dBm.

As a practical matter, the industry will continue to talk about input levels to log amps in terms of dBm power levels, with the implicit assumption that it is based on a 50 Ω impedance, even if it is not completely correct to do so. As a result it is prudent to provide specifications in *both* dBm and dBV in data sheets.

The figure shows how mV, dBV, dBm and mW relate to each other for a load impedance of 50 Ω. If the load impedance were 20 Ω, for example, the V (rms), V (p-p) and dBV scales will be shifted downward relative to the dBm and mW scales. Also, the V (p-p) scale will shift relative to the V (rms) scale if the peak to rms ratio (also called crest factor) is something other than $\sqrt{2}$ (the peak to rms ratio of a sine wave).

