ON Semiconductor

Is Now

onsemi™

To learn more about onsemi™, please visit our website at
www.onsemi.com

Zero-Drift Precision Op Amps: Advantages and Limitations of the Chopper-Stabilized Architecture

ON Semiconductor®

APPLICATION NOTE

What are zero−drift precision op amps?

Zero−drift precision op amps are specialized op amps designed for applications that require high output accuracy due to small differential voltages. Not only do they feature low input offset voltage, but they also have high CMRR, high PSRR, high open loop gain, and low drift over temperature and time. These features make them ideal for applications such as low−side current sensing and sensor interface, particularly with very small differential signals.

Table 1. Key Parameters for precision op amps include all parameters that effect accuracy and precision

Key Parameter	Symbol	Units	Importance
Input Offset Voltage	V_{OS} or V_{IO}	μV	Lower offset allows lower differential voltages to be measured with accuracy
Input Offset Voltage Drift	dV_{OS}/dT or $\Delta V_{OS}/\Delta T$	μV/°C	Lower drift prevents the offset voltage changing over temperature
Common Mode Rejection Ratio	CMRR	dB	High CMRR indicates that the offset voltage is less susceptible to variation in the common mode voltage
Open Loop Voltage Gain	A_{VOL}	dB	High open loop gain results in better closed loop gain accuracy
Power Supply Rejection Ratio	PSRR	dB	High PSRR means that the offset voltage is less susceptible to variation in the supply

Why is input offset voltage so important?

Offset voltage is one of the parameters that limits the smallest signal that can be captured dependably. This defines the low level of the dynamic range.

Input offset voltage is a key parameter for all op amps. In datasheets, it is commonly referred to as V_{OS} or V_{IO}. It is the inherent differential voltage between the IN+ and IN− terminals and a measure of how well the input pairs are matched. For an ideal op amp, $V_{IN+} = V_{IN−}$ for a closed loop system. In the real world, $V_{IN−}$ can't equal V_{IN+} due to the input offset voltage.

Although there are silicon−level design techniques that can be implemented to improve input pair matching, the manufacturing process is the major contributing factor to creating an input offset voltage. Imperfections in the semiconductor material give rise to the internal voltage difference between the input pins. Different types of imperfections that arise from the manufacturing process will produce different temperature coefficients. This part−to−part variation can cause the drift (input offset voltage drift over temperature) of a particular part to be higher or lower than the typical value on the datasheet. Furthermore, the drift coefficient may be positive or negative with temperature. This makes it difficult to simply calibrate out input offset voltage in the application. In some cases, reducing offset or drift in conventional linear op amps will incur a penalty in power consumption.

The input offset voltage is multiplied by the gain and added to the output voltage, essentially adding an error factor to the output, as shown Figure 1. This parameter becomes critical when measuring small differential voltages. As the differential voltage becomes smaller, the error due to input offset voltage increases.

Figure 1. Current sensing with an op amp in a difference amplifier configuration. Low offset voltage is critical since the input offset voltage is amplified by the noise gain, creating an offset error at the output.

In the difference amplifier circuit shown in Figure 1 above, the output voltage is a sum of the signal gain term and the noise gain term:

$$V_{OUT} = \left(V_{SENSE} \cdot \frac{R_F}{R_1} \right) + \left(V_{OS} \cdot \left(1 + \frac{R_F}{R_1} \right) \right) \qquad \text{(eq. 1)}$$

As an internal op amp parameter, the input offset voltage is multiplied by the noise gain and not the signal gain. This results in an output offset error.

A precision amplifier seeks to minimize this offset as far as possible, utilizing a multitude of techniques to decrease the input offset voltage. For zero−drift amplifiers, this applies particularly to low frequency and DC signals. Table 2 compares the maximum input offset of commonly used general purpose op amps compared to chopper−stabilized zero−drift amplifiers.

Table 2. MAXIMUM OFFSET VOLTAGE COMPARISON BETWEEN COMMON GENERAL PURPOSE OP AMPS AND THE CHOPPER−STABILIZED ZERO−DRIFT OP AMPS

Part Number	Description	Max V_{OS} at 25°C
LM321	Legacy general purpose op amp	7000 μV
NCS20071	General purpose op amp	3500 μV
NCS21911	Chopper−stabilized zero−drift op amp	25 μV
NCS333A	Chopper−stabilized zero−drift op amp	10 μV

What are zero−drift op amps made of?

Precision op amps are able to achieve "zero−drift" offset voltage, maintaining low input offset voltage over temperature variation and time, through a number of techniques. One of the ways that an amplifier can achieve this is by using a design technique that periodically measures the input offset voltage and corrects the offset at the output. This type of architecture is referred to as chopper−stabilized.

Like all engineering solutions, zero−drift op amps also have their limitations. One of the less obvious is a result of the fact that the internal circuit of the chopper−stabilized amplifier contains a clocked system. A simplified block diagram of the chopper−stabilized architecture used in ON Semiconductor's NCS333 and NCS21911 is shown in Figure 2. While some might debate that this type of chopping is a real−time system, practice has shown that it is susceptible to classic sampling system problems of aliasing or heterodyning. The main artifact from chopper−stabilized op amps occurs when a signal nears the clock frequency of the chopper. In this application note, we will use the word aliasing, but that will encompass what is probably more appropriately called heterodyning.
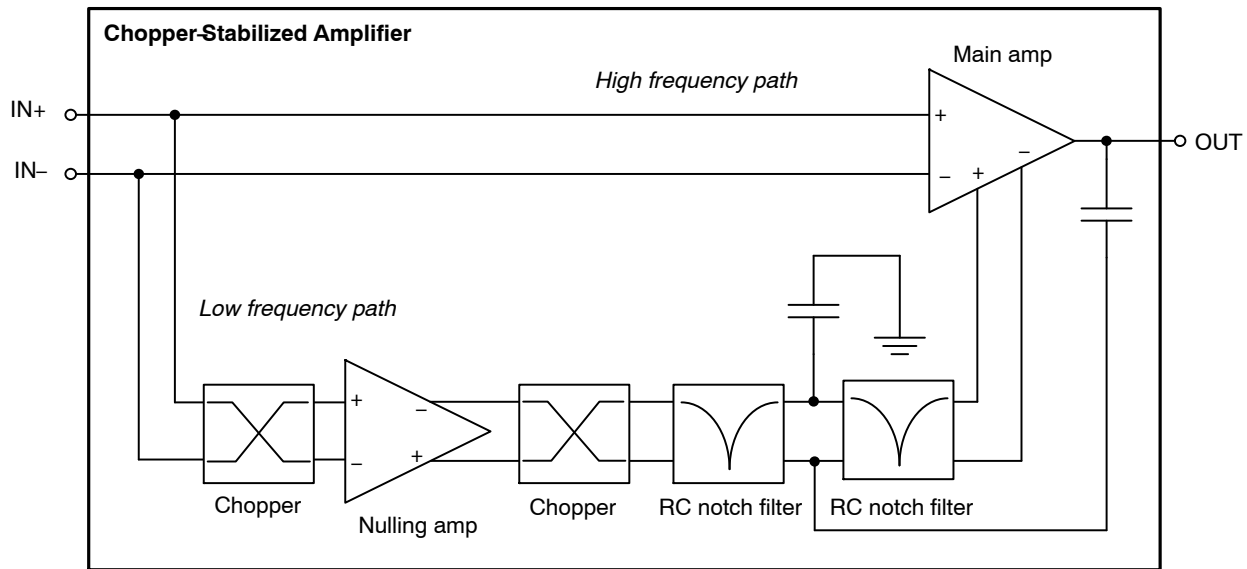
Chopper-Stabilized Amplifier

High frequency path

Main amp

IN+

IN−

OUT

Low frequency path

Chopper

Nulling amp

Chopper

RC notch filter

RC notch filter

Figure 2. Simplified block diagram of a chopper−stabilized op amp

In Figure 2, the lower signal path is where the chopper samples the input offset voltage, which is then used to correct the offset at the output. This offset correction frequency is within the overall bandwidth of the amplifier. Because this type of architecture uses a sampling method, the optimal performance occurs when the input signal frequencies remain below the related Nyquist frequency. This means that not only do the input signal frequencies need to be within the closed loop bandwidth but also within half the offset correction frequency for best performance. This allows the chopper to maintain a sampling frequency above the Nyquist rate, eliminating the likelihood of aliasing. As the signal frequency exceeds the Nyquist frequency, aliasing may occur at the output. This is an inherent limitation of all chopper and chopper−stabilized architectures due to the sampling system.

$$f_{NYQUIST} = \frac{1}{2} \cdot f_{CLOCK} \qquad \text{(eq. 2)}$$

The chopper−stabilized architecture benefits from having a feed−forward path, which is shown as the upper signal path of the block diagram in Figure 2. This is a high speed signal path that extends the gain bandwidth beyond the sampling frequency. Not only does this help retain high frequency components of the input signal, but it also improves the loop gain at low frequencies. Consider that the open loop gain of an op amp decreases by −20 dB/decade. When the unity gain bandwidth is increased, the plot shifts towards higher gain as well. An example is shown in Figure 3. When the op amp is put into a closed loop system, the increased open loop gain of the system improves the closed loop accuracy of the system. This is especially useful for low−side current sensing and sensor interface applications where the signal is low frequency and the differential voltage is relatively small.

Figure 3. Open loop gain over frequency of two chopper−stabilized amplifiers. The higher bandwidth NCS21911 demonstrates how increasing the unity gain bandwidth also increases the overall open loop gain. Increased open loop gain improves the accuracy of closed loop system, even at DC.

Nevertheless, not all zero−drift amplifiers are built the same. Different implementations of the architecture can have different results. Even with the limitations due to sampling, ON Semiconductor's NCS333 and NCS21911 series op amps have minimal aliasing and are less susceptible to aliasing effects when compared to competitor parts from other manufacturers. This is because of ON Semiconductor's patented approach using two cascaded, symmetrical, RC notch filters tuned to the chopper frequency and its fifth harmonic to reduce aliasing effects.

Another zero−drift architecture is referred to as "auto−zero". The block diagram of the auto−zero architecture, shown in Figure 4, is similar to the chopper−stabilized architecture, though the implementation is different. The auto−zero architecture has a main amplifier and a nulling amp. This method also uses a clocked system. In the first phase, the switched capacitors hold the offset error from the previous phase on the nulling amp output. In the second phase, the offset from the nulling amp output is used to correct the offset of the main amplifier. The architectural differences between auto−zero and chopper−stabilized amplifiers result in differences in noise performance and aliasing susceptibility, which will be discussed in later sections.
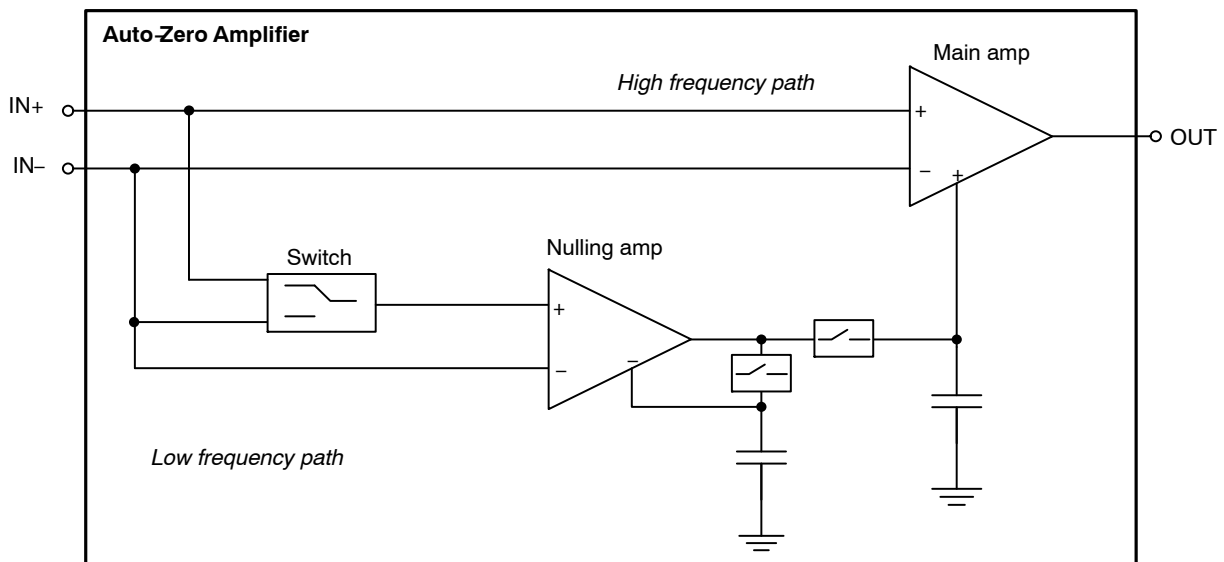

Figure 4. Simplified block diagram of an auto−zero op amp

Identifying zero−drift amplifier clock frequency

Many zero−drift amplifier datasheets provide no information on the internal clock frequency. Sometimes, it may be mentioned in a paragraph in the application section. Occasionally, it may be identified by a perturbation in the noise or bandwidth plot, alluding to what the clock frequency may be. As a consequence, it is up to the user to test the circuit for susceptibility to aliasing. The method shared here is very simple: time domain testing while observing the output on an oscilloscope. This can easily be done by sweeping the amplifier input over a range of frequencies up to the gain−bandwidth product. The internal clock frequencies of all known zero−drift amplifiers are, to the authors' knowledge, within the gain−bandwidth of the amplifier, often at roughly one−third of the gain−bandwidth. These amplifiers are going to perform their best at signal bandwidths less than half of that frequency.

Discovering and testing aliasing

Some datasheets for zero−drift amplifiers state that they are free from aliasing. One may be able to assume that those manufacturers made an effort to measure any possible aliasing and found none. In the development of zero−drift amplifiers here at ON Semiconductor, initial measurements on competitive amplifiers gave them a clean bill of health.

At the time, clock artifacts in the output of the competitor's devices were not found. Further testing has shown that aliasing can still be found on these devices using a simple technique of time−domain analysis.

Aliasing was discovered when customers reported problems with systems using zero−drift op amps from several manufacturers. In these cases, there was a common theme where the signal of interest, a low frequency or DC signal, had high amplitude, high−frequency interference or ripple signals superimposed. The result on the end system varied, including closed loop systems stabilizing at incorrect conditions and systems unable to report a correct signal.

Past efforts to discover aliasing involved the use of sophisticated spectral and network analysis systems that provided inconclusive results. To take a more fundamental approach, an oscilloscope was connected to the amplifier output for direct visual observation. For input stimulation, a generator was used to sweep the input frequency around where the clock frequency was expected to be (and beyond that if necessary) to see if a "beat frequency" could be produced on the output. This method worked surprisingly well, considering that the initial effort was with a gain of +1 configuration, shown in Figure 5, arguably one of the most linear op−amp configurations.


Figure 5. The test circuit to detect aliasing is a simple unity gain buffer. Essential to the technique is viewing the device output on an iscilloscope. It seems that spectrum and network analyzers don't always detect signals related to the internal workings of zero−drift amplifiers.

The first op−amp chosen for this test was ON Semiconductor's NCS325 which is an auto−zero technology amplifier, as opposed to chopper−stabilized like the other devices tested. In theory, an auto−zero architecture will exhibit more dramatic aliasing than a chopper−stabilized type which makes it a convenient first choice to validate the test. Figure 6 depicts the aliasing of the NCS325. Measuring a familiar amplifier first made validating these tests easy since the clock frequency is known.


Figure 6. Aliasing output of the first amplifier tested, the ON Semiconductor NCS325 in a simple +1 V/V buffer. The upper blue trace is the input signal, and lower, orange trace is the aliasing seen at the amplifier output.

It is important to remember at this point; aliasing is not a defect of sampling amplifiers, it is a behavior. Knowledge of this behavior and how to avoid it can make zero−drift amplifiers operate at their best.

After checking the NCS325, the ON Semiconductor NCS333, which is a chopper−stabilized amplifier, was tested next. An interesting result occurred here, the only noticeable aliasing that could be found occurred at twice the clock frequency. This points out that performing this test to discover aliasing may require sweeping throughout the bandwidth of the amplifier to detect these signals. Figure 7 depicts the aliasing signal from the NCS333.

Figure 7. NCS333 Chopper−stabilized zero−drift op amp aliasing. The aliasing was expected to occur near the clock frequency but despite our best efforts it didn't occur. Nonetheless aliasing did occur at the second harmonic of the clock frequency.

A competitor's zero−drift chopper−stabilized amplifier was also tested for aliasing. This popular amplifier's datasheet states it has no aliasing. However, Figure 8 depicts aliasing at approximately the fundamental frequency of the internal clock. For this amplifier, previous extensive testing with spectrum and network analyzers could not uncover evidence of the clock or its frequency.
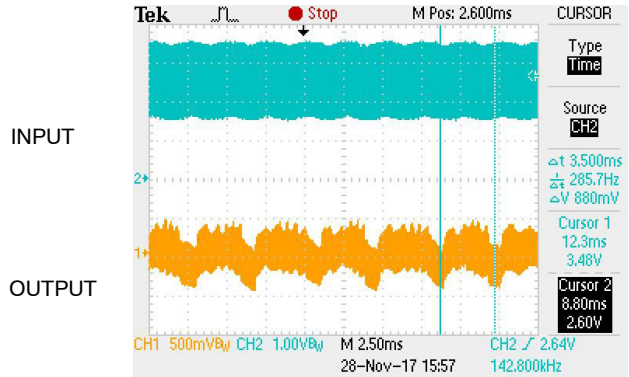

Figure 8. Brand−X chopper−stabilized zero−drift op amp aliasing. The datasheet for this device claims there is no aliasing.

Similarly, the NCS21911 precision op amp, with its 2 MHz bandwidth, shows aliasing when the input signal is 500 kHz with a gain of approximately G = −1 V/V, as shown in Figure 9.


Figure 9. Aliasing in the NCS21911, which is a 36 V, 2 MHz, precision amplifier. The aliasing is still relatively well−controlled at 500 kHz. The center, blue trace is the input signal, and the larger, magenta trace is the amplifier output exhibiting aliasing.

However, the NCS21911 aliasing is relatively well−controlled when compared to another manufacturer's counterpart under the same conditions, as shown in Figure 10.


Figure 10. Aliasing in the competitor's 36 V, 2 MHz precision amplifier shows more erratic behavior in the output at the same signal frequency of 500 kHz. The center, blue trace is the input signal, and the larger, magenta trace is the amplifier output exhibiting aliasing.

Another example is shown between the NCS21911 and the competitor's 2 MHz chopper−stabilized precision op amp. The NCS21911 shows minimal aliasing in the 1 MHz to 2 MHz range in the unity gain buffer circuit, as shown in Figure 11. In contrast, the competitor part reacts, behaving normally at 1 MHz, exhibiting aliasing at 1.5 MHz, and diminished aliasing at 2 MHz (along with bandwidth), as seen in Figure 12.

Figure 11. NCS21911 in a unity gain circuit with a small signal at (top) 1 MHz, (middle) 1.5 MHz, and (bottom) 2 Mhz. Aliasing is minimal.

Figure 12. The competitor's 2 MHz chopper–stabilized precision op amp with a small signal at (top) 1 MHz, (middle) 1.5 MHz, and (bottom) 2 MHz. Aliasing is pronounced at 1.5 MHz and diminished as the input signal increases to 2 MHz. Also note that the lower bandwidth of the competitor part as seen on the bottom waveform.

Not every chopper stabilized amplifier is built the same. That is why it is critical to test each part over the entire operating frequency range.

Systems susceptible to aliasing

A system is vulnerable to aliasing in cases where the signal of interest is accompanied by high frequency coupling of stray signals or large high frequency ripple. The result can include merely delivering incorrect or noisy values, or control loops settling on incorrect operating points.

According to the Nyquist sampling theorem, the zero−drift clock should be at least twice the maximum frequency component of the signal of interest. In other words, the maximum frequency of the input signal should be less than or equal to half of the amplifier's internal clock.

How does one adhere to the Nyquist sampling theory? It's easy enough to establish an upper limit for signal frequency ($f_{in} < f_{CLOCK}/2$), but pickup from stray signals, noise, or ripple could contain frequencies higher than the Nyquist frequency. These frequencies may then alias into the appropriate frequency range, resulting in errors or incorrect readings.

To ensure that the frequency content of the input signal is limited to the usable frequency range, a low pass filter can be added before the amplifier. This filter acts as an anti−alias filter. By attenuating the higher frequencies (beyond the Nyquist frequency), it reduces or eliminates the aliasing effect. The anti−aliasing filtering must be purely analog filtering before the inputs of the amplifier. Often a simple RC filter will suffice, as shown in Figure 13. Elaborate and fancy filter architectures should rarely be necessary. Do not configure the amplifier as part of the filter in an active filter circuit.


**Figure 13. An anti−aliasing filter can be as simple as a two−section RC filter.
The filter must be placed ahead of the amplifier inputs.**

Cascading zero−drift amplifiers can also pose a risk, since multiple clock frequencies may interact and result in aliasing.

Transient response considerations

Due to the time−based sampling of the chopper channel architecture, there is a time aspect to the zero−drift amplifier achieving its low offset, which means that offset correction doesn't occur instantaneously. A large dynamic step on the amplifier input or, even worse, an input overload can create conditions in which the loop will need time to reestablish a low offset. This essentially affects the settling time and behavior.

The use of higher clock frequencies has enabled relatively fast recovery and settling times; nevertheless, those parameters are typically in the tens of microseconds or higher for zero−drift amplifiers. As usual, this is due to design tradeoffs. At the transistor level amplifier design, opting for faster settling times can lead to higher offset voltage. Typically, a lower input offset voltage specification has the higher priority.

Turn−on time and robust design

Because zero−drift amplifiers incorporate a fair amount of logic circuitry, it is not surprising that they also include some means of ensuring a defined behavior during startup and power glitches (such as brownouts). When first powering up an offset−correcting amplifier, there is a small period of time during which the output will reflect the uncorrected offset. Once the supply voltage reaches a defined trip−point set by the Power−On Reset (POR) circuit, the offset correction mechanism needs a few clock cycles until the output of the amplifier is within the specified offset voltage limits. Typically, this amplifier startup time is not a critical item from an overall system perspective as it is usually well within the power−up time of the whole system. This may be the reason why many op amp manufacturers do not show this parameter within their zero−drift amplifier datasheets. It should be noted that the startup time also depends on the configured gain of the amplifier−larger gains can increase the overall start−up time.

In very critical systems thought should be given to the fact that linear amplifiers simply dispense with these complications, and yield much more robust power up performance. Some precision amplifiers use trim to achieve the low offset voltage instead of a chopper−stabilized or auto−zero architecture. This eliminates any clocked system with the amplifier. This is a critical consideration in many designs such as large industrial circuit breakers. The tradeoff is that these trimmed, linear amplifiers do not necessarily reach the same ultra−low input offset voltage performance of zero−drift amplifiers.

The zero−drift effect of improved rail−to−rail performance

Rail−to−rail input op amps use two inputs pairs to achieve the widened common mode input voltage range. A PMOS pair may be used as the input stage for the lower input voltage region and an NMOS pair may be used for the higher input voltage region. Each input pair has its own corresponding input offset voltage. When the common

mode voltage shifts from one region to the other, there is typically a crossover region where the offset voltage jumps from one region to the next.

Rail−to−rail input performance in zero−drift op amps show significant benefits compared to their non−zero−drift companions by dramatically reducing the effect of the input stage crossover region between the PMOS and NMOS input pairs. Offset voltage and offset voltage drift performance near the limits of the common mode input voltage is excellent, which is why zero−drift amplifiers are often used for applications such as high side current sensing as well.

The zero−drift effect on low frequency noise

Zero−drift chopper−stabilized amplifiers are especially suited for accurate, high−gain amplification at lower frequencies. In general, they do not exhibit the higher bandwidth of linear op amps, and the location of their clock frequency establishes a practical frequency limit on signal fidelity, as is discussed in the section on aliasing. This makes performance at low frequencies especially important, and the chopper−stabilized architecture further contributes to low frequency usefulness by eliminating the classic linear op amp 1/f input voltage noise. Many high gain sensor applications are at low frequencies, making zero−drift amplifiers a natural choice for this function. Despite the use of the term "low frequencies" here, these amplifiers provide excellent performance in general up to around 100 kHz.

Figure 14. Voltage noise density plots of conventional (a. NCS2005) and zero−drift (b. NCS333) amplifiers show the elimination of 1/f noise in the zero−drift amplifier. Note that the conventional amplifier plot does not show below 10 Hz, while the zero−drift plot is down to 1 Hz.

Just like voltage noise, the chopper−stabilization also eliminates 1/f current noise. However, chopper−stabilized amplifiers exhibit increased input current noise from the chopping, due to charge injection of the input switches. This increased current reduces the level at which input impedances can cause noise to equal or exceed the level of voltage noise. Using the NCS333 as an example, which has 62nV/√Hz input voltage noise at 1 kHz, the 350fA/√Hz input current noise will result in noise exceeding the input voltage noise when input impedances greater than 177 kΩ are used.

In comparison, zero−drift auto−zero amplifiers alias noise down to the baseband. This gives the auto−zero architecture a disadvantage compared to the chopper−stabilized architecture in cases where the input signals are at DC or low frequencies.

The zero−drift effect on input current

Input current spikes are present on all zero drift amplifiers due to the chopper−stabilization technique. These current spikes are caused by charge injection and clock feedthrough.

The input current is averaged out into the I_{IB} specification, but the input bias current is not truly constant. In reality, the input current spikes occur periodically with the clock frequency.

As the input current draws across input resistors, this results in input voltage spikes that can get multiplied by the gain. In order to minimize the voltage spikes, very large input resistor values are not recommended. Input current spikes can also be filtered out with a simple, low−pass RC filter as shown previously in Figure 13. The filter frequency should be set below the chopper sampling rate.

Additionally, the input current spikes make zero−drift amplifiers very poor candidates for transimpedance amplifiers, which are designed to measure input current.

The absence of the zero−drift effect in Spice models

Spice simulations do not provide any insight into zero−drift amplifier behaviors such as aliasing. All Spice models of zero−drift amplifiers are continuous−time models. They are designed to model as closely as possible the linear performance of the op−amp. The chopper is not

modeled. They are continuous time because clocked and sampled systems simulate much slower.

Conclusion

Zero−drift amplifiers provide superb DC and low frequency performance. Gain−bandwidth product is a less−than−ideal specification for determining the practical bandwidth of a zero−drift amplifier circuit, especially since their internal clocks are within this bandwidth. Achieving the best performance requires knowledge of the internal clocking frequency which is not always available, but sometimes other clues and testing will reveal it.