A 140 dB-CMRR Current-Feedback Instrumentation Amplifier Employing Ping-Pong Auto-Zeroing and Chopping

Michiel A. P. Pertijs, *Senior Member, IEEE*, and Wilko J. Kindt

***Abstract*—This paper presents a precision general-purpose current-feedback instrumentation amplifier (CFIA) that employs a combination of ping-pong auto-zeroing and chopping to cancel its offset and $1/f$ noise. A comparison of offset-cancellation techniques shows that neither chopping nor auto-zeroing is an ideal solution for general-purpose CFIAs, since chopping results in output ripple, and auto-zeroing is associated with increased low-frequency noise. The presented CFIA mitigates these unintended side effects through a combination of these techniques. A ping-pong auto-zeroed input stage with slow-settling offset-nulling loops is applied to limit the bandwidth of the increased noise to less than half of the auto-zeroing frequency. This noise is then modulated away from DC by chopping the input stage at half the auto-zeroing frequency, reducing the low-frequency noise to the 27 nV/$\sqrt{\text{Hz}}$ white-noise level, without introducing extra output ripple. The auto-zeroing is augmented with settling phases to further reduce output transients. The CFIA was realized in a 0.5 μm analog CMOS process and achieves a typical offset of 2.8 μV and a CMRR of 140 dB in a common-mode voltage range that includes the negative supply.**

***Index Terms*—Instrumentation amplifiers, current-feedback amplifiers, dynamic offset cancellation, chopping, auto-zeroing.**

I. Introduction

CURRENT-FEEDBACK instrumentation amplifiers (CFIAs) are attractive because of their ability to sense differential input voltages in a common-mode voltage range (CMVR) that includes either of the supply rails [1]. They can thus interface directly, for instance, with a current-sense resistor in series with the supply [2], or with a ground-referenced sensor in a single-supply system, as illustrated in Fig. 1(a).

Such applications often involve mV-level input voltages [3], [2], [4]. To accurately process these voltages, the input-referred errors introduced by a general-purpose precision CFIA should be at the μV level. This implies a μV-level input-referred offset and noise. Moreover, if the amplifier is exposed to input common-mode voltage variations in the order of 1 V, it should have a common-mode rejection ratio (CMRR) in the order of 120 dB.

Manuscript received March 14, 2010; revised June 03, 2010; accepted June 10, 2010. Date of current version September 24, 2010. This paper was approved by Associate Editor Andreas Kaiser.

M. Pertijs was with National Semiconductor, Delft, The Netherlands, and is now with the Delft University of Technology, Delft, The Netherlands (e-mail: pertijs@ieee.org).

W. Kindt is with National Semiconductor, Delft, The Netherlands.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2010.2060253

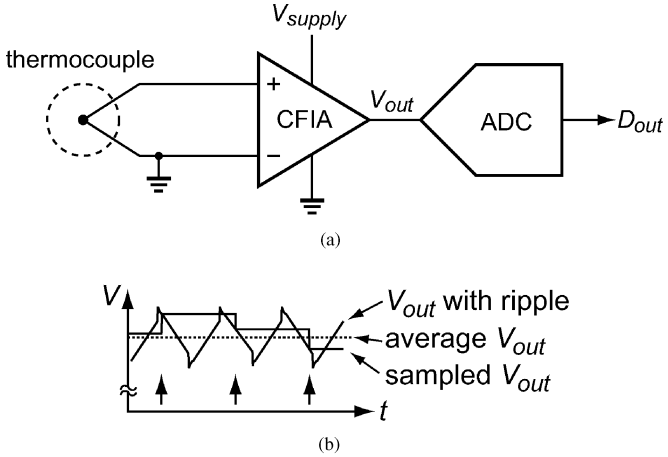

Fig. 1. (a) Current-feedback instrumentation amplifier interfacing between a grounded sensor and an ADC. (b) Error due to sampling of output ripple by the ADC.

A MOS input is desired to achieve a high input impedance and low input bias currents. To reduce the offset voltage of the MOS input transistors, which is typically at the mV-level, and their $1/f$ noise, dynamic offset cancellation (DOC) techniques such as chopping or auto-zeroing are required. A problem that complicates the application of DOC techniques in a general-purpose CFIA are the switching transients or ripple that occur at the amplifier's output. Such transients can cause large errors when they are sampled by an ADC connected to the output of the amplifier [Fig. 1(b)]. While in principle this can be prevented by synchronizing the sampling with the DOC (see, e.g., [3]), this is not possible in many practical systems, or not desired due to the additional complexity that such synchronization introduces at the system level. Therefore, a general-purpose CFIA should provide a clean, ripple-free output signal, making any internally used DOC techniques transparent to the user.

Recently, several precision CFIAs employing chopping have been reported [2], [4], [5], [6]. In a chopped amplifier, offset and $1/f$ noise are eliminated by modulating them away from DC. The modulated offset results in ripple at the output of the amplifier, which can be suppressed by incorporating the chopped amplifier as a precision low-frequency path in a multi-path amplifier [5]. Further ripple reduction can be obtained by auto-zeroing the chopped low-frequency path [2], or by including a ripple-reduction loop that senses the ripple at the amplifier's output and feeds back a correction signal to the input stage [4], [6]. A common disadvantage of these approaches, which will be reviewed in detail in Section III-B, is that they modulate the full

differential input signal (not just the offset). Depending on the implementation of the chopped input stage, this can give rise to signal-dependent switching transients at the amplifier's output.

This paper presents an alternative approach: a precision CFIA that uses ping-pong auto-zeroing to achieve μV-level offset as well as a clean output voltage [7]. Rather than modulating the offset, it is sampled and subtracted, so that no ripple is generated. Settling phases are added to this auto-zeroing process to prevent any switching transients from reaching the amplifier's output. The increased low-frequency noise associated with auto-zeroing is modulated away from DC by also chopping the amplifier. In contrast with earlier reported combinations of chopping and autozeroing, which use a chopping frequency higher than the auto-zeroing frequency [8], [2], the amplifier is chopped at half of the auto-zeroing frequency. This low chopping frequency prevents the introduction of new switching transients. Thus, the CFIA combines the clean output of an auto-zeroed amplifier with the low noise of a chopper amplifier.

This paper is organized as follows. Section II briefly reviews the properties of CFIAs. Section III describes how three commonly used offset cancellation techniques, trimming, chopping and auto-zeroing, can be applied to CFIAs. The main limitations of these techniques are discussed. Section IV then introduces two ways of combining chopping and auto-zeroing: first, chopping at twice the auto-zeroing frequency, as in [8], which is shown to remove the increased low-frequency noise associated with auto-zeroing, but at the expense of increased output transients; second, chopping at half the auto-zeroing frequency, which is enabled by a slow-settling offset-nulling loop, and does not increase the output transients. Section V describes the implementation details of the prototype CFIA employing this new combination of chopping and auto-zeroing. Section VI presents measurement results of this prototype. The paper ends with conclusions in Section VII.

II. Current-Feedback Instrumentation Amplifiers

Two important properties distinguish CFIAs from other instrumentation amplifiers, most notably from conventional 3-opamp instrumentation amplifiers: their superior CMRR, and their ability to sense input voltages in a CMVR that includes either of the supply rails [1], [5]. Both are a direct result of the fact that the differential input voltage $V_{\rm in}$ is converted into a current by an input tranconductor $g_{\rm in}$, as shown in Fig. 2. A similar transconductor $g_{\rm fb}$, incorporated in a feedback loop, senses the voltage drop $V_{\rm fb}$ across a resistive divider between the output voltage $V_{\rm out}$ and a reference voltage $V_{\rm ref}$. This feedback loop drives the output such that the output currents of the transconductors cancel. As a result, $V_{\rm fb}$ will equal $V_{\rm in}$, assuming the transconductances are equal, and the output voltage will be

$$V_{\rm out} = V_{\rm ref} + \frac{R_1 + R_2}{R_2} V_{\rm in}. \tag{1}$$

The CMRR of a CFIA is determined by a combination of isolation and balancing [1]. The high output impedance of $g_{\rm in}$ and that of the internal bias current sources in $g_{\rm in}$ (not shown in Fig. 2) isolate the input common-mode voltage from the rest

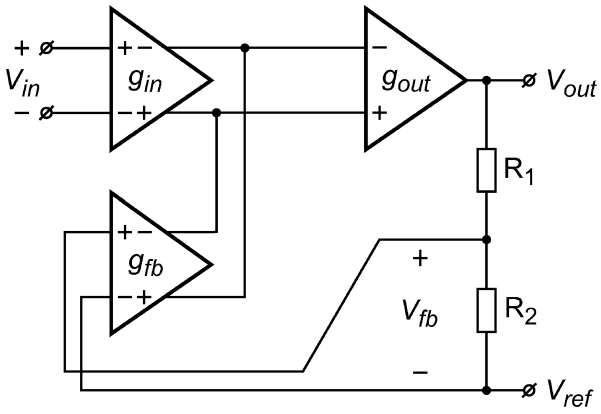

Fig. 2. A typical current-feedback instrumentation amplifier.

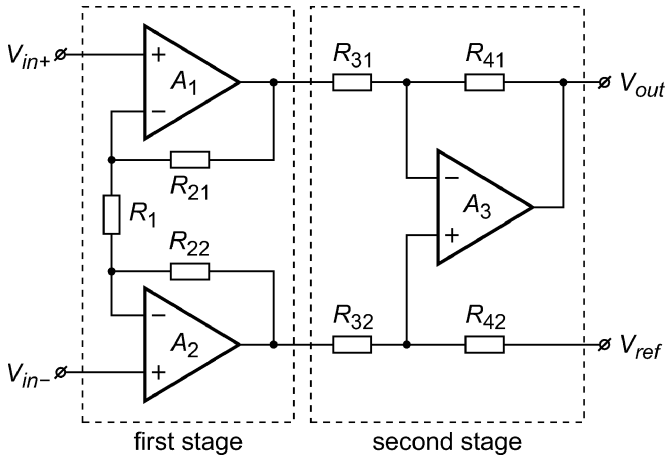

Fig. 3. Conventional 3-opamp instrumentation amplifier.

of the circuit. As a result, variations in the input common-mode voltage lead to only very small currents at the output of $g_{\rm in}$. Moreover, if the output impedances are balanced, i.e., matched between the two halves of the differential signal path, these small currents will be common-mode currents that will be absorbed by a common-mode feedback circuit at the output of $g_{\rm in}$, and will therefore not give rise to a change in the output voltage. It is thus a combination of finite isolation and finite balancing that limits the CMRR of a CFIA. In practice, values well in excess of 120 dB can be achieved [2].

The CMRR of a 3-opamp instrumentation amplifier (Fig. 3), in contrast, is determined by the product of the finite differential gain of its first stage, and the finite CMRR of its second stage. The latter is determined by balancing only, i.e., by the matching of the feedback resistors R_{31}, R_{32}, R_{41}, and R_{42}. This typically leads to much poorer CMRR values, in the order of 80 dB [9].

CFIAs can sense input levels at or slightly below the supply rails by exploiting the inherent level-shift provided by the input transistors of $g_{\rm in}$. Assuming, for instance, that $g_{\rm in}$ has pMOS input transistors, the threshold voltage of these transistors can level shift an input voltage at or slightly beyond the negative supply to a level within the supply rails, where it can be converted into a current.

A 3-opamp instrumentation amplifier, in contrast, relies on feedback at the input common-mode level: amplifiers A_1 and A_2 not only sense at this level, but also need to provide output

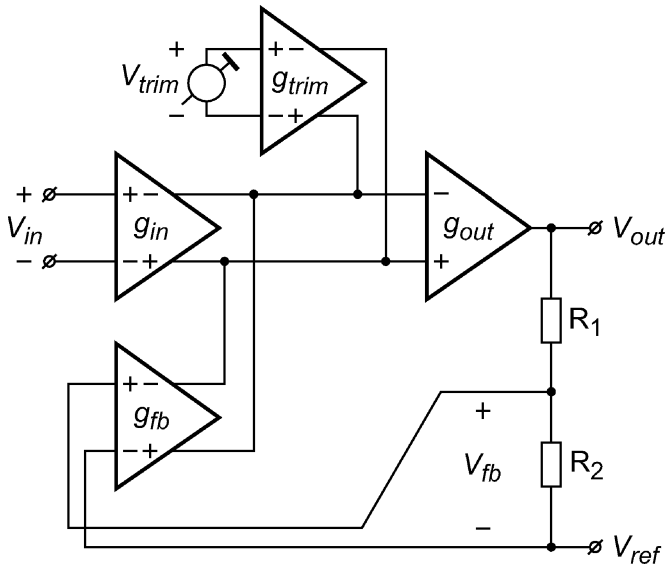

Fig. 4. A trimmed CFIA.

voltages at this common-mode level. Since they cannot drive at or beyond the supply rails, the CMVR of such a 3-opamp instrumentation amplifier cannot include the supply rails.

III. Offset Cancellation in CFIAs

In order to achieve the μV-level offset and noise required by precision applications using a MOS-input CFIA, offset cancellation is required to reduce the offset voltage and $1/f$ noise of the MOS input transistors. In this section, three classes of offset cancellation techniques are discussed: trimming, auto-zeroing, and chopping [10].

A. Trimming

Trimming consists of measuring and adjusting the offset during production. In a CFIA, such an adjustment can be conveniently implemented using an additional transconductor $g_{\rm trim}$ to which an offset-compensation voltage is applied. This transconductor then generates a current that effectively nulls the combined offset currents of the input and feedback transconductors (Fig. 4). More precisely, it nulls the offset *difference* between these transconductors. To avoid having to generate a stable mV-level compensation voltage, $g_{\rm trim}$ is made smaller than $g_{\rm in}$ and $g_{\rm fb}$, so that the compensation voltage is in fact an amplified version of the offset difference.

While this approach can be used to obtain an order-of-magnitude reduction of the offset, it is unable to reduce a mV-level initial offset to the required μV-level, because offset drift is not compensated for. Moreover, trimming does not eliminate $1/f$ noise. Dynamic offset cancellation techniques such as chopping or auto-zeroing are therefore needed to mitigate these problems.

B. Chopping

In a chopped CFIA, the input and feedback voltages are modulated away from DC by means of two polartiy-reversing switches (or "chopper" switches), as shown in Fig. 5. At the

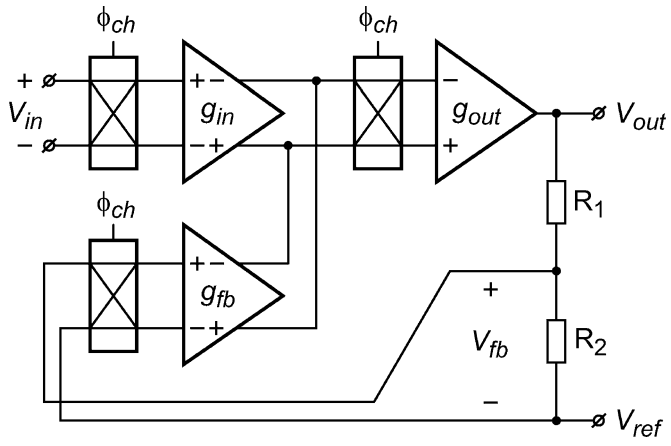

Fig. 5. A chopped CFIA.

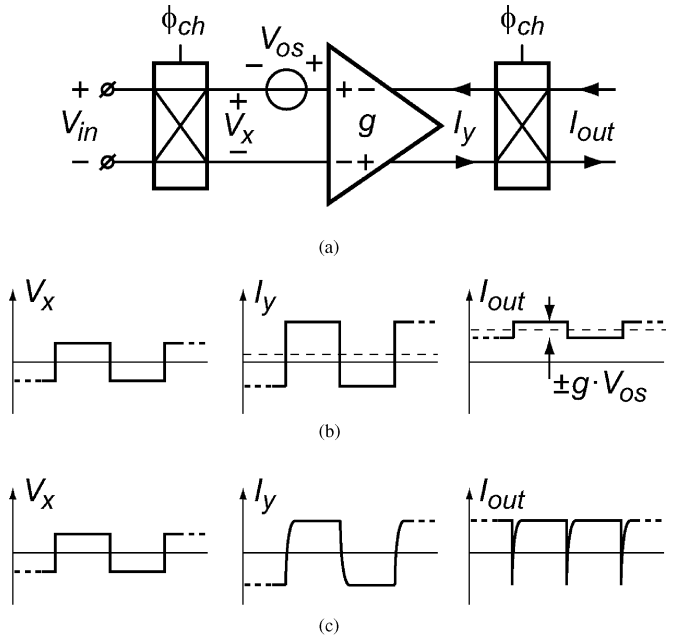

Fig. 6. (a) A chopped transconductor. (b) Offset results in a modulated output current. (c) Even without offset, finite bandwidth results in output glitches.

outputs of the transconductors, a second chopper switch modulates the output current back to DC, while modulating the offset and $1/f$ noise of the transconductors away from DC, where they can, in principle, be filtered out.

As illustrated in Fig. 6(a) and (b) for a single chopped transconductor, the modulation of the offset leads to a square-wave modulated output current. In a chopped Miller-compensated CFIA, this square-wave modulated current is integrated by the output stage, giving rise to a triangular output ripple with an amplitude proportional to the offset. This ripple gives rise to large errors when sampled by an ADC [as illustrated in Fig. 1(b)], unless the sampling is synchronized with the chopping and takes place at the zero-crossings of the ripple [11].

If such synchronization is not possible, there are several ways to reduce the ripple. One way is to include a low-pass filter at the output of the amplifier. This, however, limits the usable signal

bandwidth and in some cases requires off-chip components to realize a corner frequency well below the chopping frequency.

A second approach to reduce the output ripple is to use a multi-path amplifier topology [12], [2]. This approach is referred to as chopper stabilization (although the same term is sometimes also used for regular chopper amplifiers). The amplifier comprises a chopped low-frequency (LF) path and a non-chopped high-frequency (HF) path. The frequency at which the HF path takes over from the LF path is chosen below the chopping frequency, so that the HF path is active at the chopping frequency and suppresses the ripple injected by the LF path. The suppression that can thus be obtained, however, is limited, as it is proportional to the ratio of the take-over frequency and the chopping frequency. On the one hand, the take-over frequency cannot be reduced below the $1/f$ corner frequency of the HF path, while, on the other hand, the chopping frequency cannot be increased indefinitely because switching transients and residual offset due to charge injection by the chopper switches will increase. The ripple of a multi-path amplifier can be further reduced by the application of auto-zeroing in the LF path (see Section III-C) [2]. This, however, comes at the cost of increased low-frequency noise.

A third approach to suppress ripple is to include a notch filter in the amplifier [13]. A switched-capacitor notch filter can be synchronized to the chopper clock so as to filter out the ripple. Such a notch filter will result in notches in the transfer function of the amplifier at the chopping frequency and its odd harmonics. These notches can be eliminated by means of a multi-path topology, in which a non-chopped HF path bypasses the notch filter at higher frequencies [13].

A fourth and final approach is to use a ripple-reduction loop [4]. Such a loop synchronously detects signals at the chopping frequency at the amplifier's output and feeds back an offset-correction signal to the input stage so as to null these signals. While such a loop effectively suppresses the ripple, it also suppresses input-signal components at the chopping frequency, and thus also results in a notch in the amplifier's transfer function. Moreover, a ripple-reduction loop is likely to be sensitive to load transients at or close to the chopping frequency. Since the loop senses the ripple at the amplifier's output, it cannot distinguish between such load transients and ripple due to modulated offset. As a result, load transients may lead to an incorrect offset-correction signal, and thus effectively introduce offset. Such transients could occur, for instance, if the output of the amplifier is loaded by the sampling capacitor at the input of a subsequent ADC. Both disadvantages can be mitigated by incorporating the ripple-reduction loop in the LF path of a multi-path amplifier [6].

A common disadvantage of all these approaches, when they are applied in a CFIA, is that the full differential input signal (and feedback signal) is modulated. Even if the transconductors are offset-free, their finite bandwidth will give rise to switching transients, which, when demodulated at the output of the transconductors, turn into glitches, as illustrated for a single chopped transconductor in Fig. 6(c). This is in particular a problem if the bandwidth of the transconductors is limited due to the fact that local feedback is used to improve their linearity or common-mode rejection (see Section V-D). However, even with fast open-loop transconductors, similar transients will result from the charging and discharging of the transconductors' differential input capacitances by the input signal source and the feedback network.

In a CFIA, these glitches will occur both at the output of g_{in} and g_{fb}. Therefore, they will cancel if they are perfectly matched. In practice, however, the time constants in the input and the feedback path will not be perfectly matched, and glitches will appear at the output. Moreover, such mismatched glitches may also lead to gain errors, since they reduce the effective transconductance of g_{in} and g_{fb}. Note, incidentally, that these issues do not occur in chopped *opamps*, such as those presented in [11] and [12], because in these opamps only a small error signal is chopped rather than the full input signal.

C. Auto-Zeroing

Auto-zeroing is an alternative to chopping that does not produce output ripple. The input stage of an auto-zeroed CFIA is operated in two alternating phases: during one phase its offset is sampled, during the other the input signal is amplified while subtracting the sampled offset, leading to an offset-free and ripple-free output signal.

The easiest implementation of this approach uses capacitors at the input of the amplifier to sample and hold the offset [10]. This implies, however, that any errors in this process, for instance due to switch charge-injection or kT/C noise, add directly to the amplifier's input voltage. By combining input offset storage with chopping, μV-level offset and noise have nevertheless been achieved [14].

An alternative is to store the offset not at the input of the amplifier, but at the input of an additional transconductor g_{AZ} [10], as shown in Fig. 7 (ignoring the "pong stage" for now). In a zeroing phase ϕ_Z, the inputs of g_{in} and g_{fb} are shorted, so that an offset current appears at their outputs. The combined offset currents of g_{in} and g_{fb} are fed to an integration capacitor C_{AZ} that drives a third transconductor g_{AZ} so as to generate a compensating current that cancels the offset currents. Thus, the offset is suppressed by the loop gain of this offset-nulling loop.

In a subsequent amplification phase ϕ_A, the integration capacitor is disconnected. As a result, the voltage across this capacitor is held constant, and g_{AZ} continues to cancel the offset currents. Then, the input and feedback voltages are connected to the inputs of g_{in} and g_{fb}, while their outputs are connected to the CFIA's output stage. Now, the CFIA operates as before, but, ideally, without offset.

The offset compensation by means of an additional transconductor g_{AZ} is similar to the approach taken in a trimmed CFIA (Fig. 4), except that the offset-compensation voltage at the input of this transconductor is now updated regularly during operation of the CFIA. Thus, offset drift and $1/f$ noise are also eliminated, provided the CFIA is auto-zeroed at a rate f_{AZ} higher than the transconductors' $1/f$ corner frequency.

Like g_{trim} in a trimmed CFIA, g_{AZ} is typically smaller than g_{in} and g_{fb}, so that an amplified offset voltage appears across C_{AZ}. Thus, any errors in this voltage, which will determine the residual offset of the CFIA, will be attenuated when referred to the input. Such errors are mainly due to charge injection by

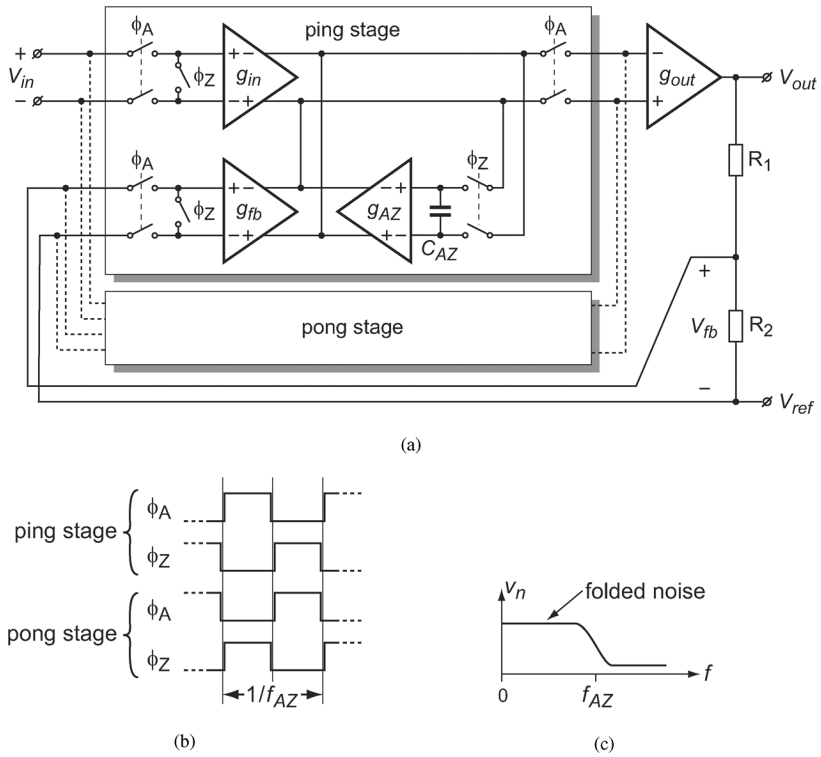

Fig. 7. A ping-pong auto-zeroed CFIA: (a) block diagram, (b) timing diagram, (c) input-referred noise PSD.

switches ϕ_Z (or rather charge-injection mismatch in a fully differential topology).

In such an auto-zeroed CFIA, the input voltage is only amplified during ϕ_A, which is typically only half of the time. This is undesirable for a general-purpose amplifier, as it will cause down-conversion of input-signal components (including noise) at harmonics of $f_{\rm AZ}$. Continuous amplification of the input signal can be obtained by using a so-called ping-pong topology [15], in which a second input stage, comprising a additional set of input and feedback transconductors with nulling circuitry, is added to the amplifier [Fig. 7(a)]. This "pong" stage is operated at opposite clock phases, so that one stage amplifies the input signal while the other is zeroed, and vice versa.

This ping-pong topology comes at the cost of the extra die area and power consumption of the additional input stage. A somewhat more efficient approach that also enables a continuous amplification of the input voltage is the ping-pong-pang topology [5]. In this approach, three transconductors are used: one amplifies the input voltage, another the feedback voltage, while the third is zeroed. These transconductors change places, so that each of them is regularly zeroed. In comparison to the ping-pong approach, this saves one transconductor. However, each individual transconductor will now require an offset-nulling loop (rather than each pair of input and feedback transconductors), so that three such loops are required, instead of two.

Auto-zeroed CFIAs, in contrast with chopped CFIAs, provide a ripple-free output voltage, because the offset is sampled and then subtracted, rather than modulated. The only output transients that remain are associated with the charge-injection of the switches, or with the step response of the transconductors when they are switched between zeroing and amplification. Like in the case of chopped CFIAs, this step response (caused by the finite bandwidth of the transconductors) can give rise to output glitches and gain errors (Fig. 8). However, as will be described in Section V-C, these effects can be prevented by including settling phases.

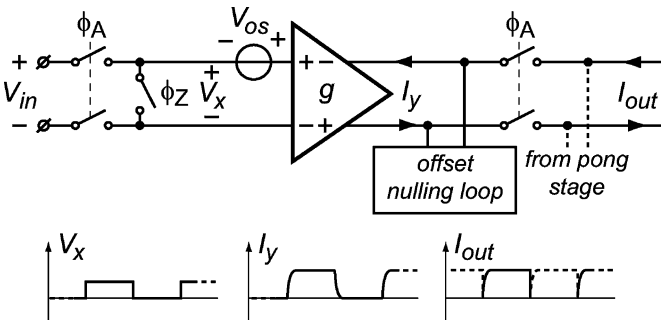

Fig. 8. Switching transients in an auto-zeroed transconductor.

A common disadvantage of all auto-zeroing approaches is the increased low-frequency noise associated with the sampling that takes place in the offset-nulling loop. Typically, the bandwidth $f_{\rm null} = g_{\rm AZ}/2\pi C_{\rm AZ}$ of this loop significantly exceeds the auto-zeroing frequency $f_{\rm AZ}$ (e.g., by a factor of 5), so as to allow the loop to fully settle within one zeroing phase. As a result, the noise within this bandwidth will alias to the bandwidth from DC to $f_{\rm AZ}$, resulting in an input-referred noise power spectral density (PSD) that is a multiple of the white-noise level of the (non-auto-zeroed) transconductors [Fig. 7(c)].

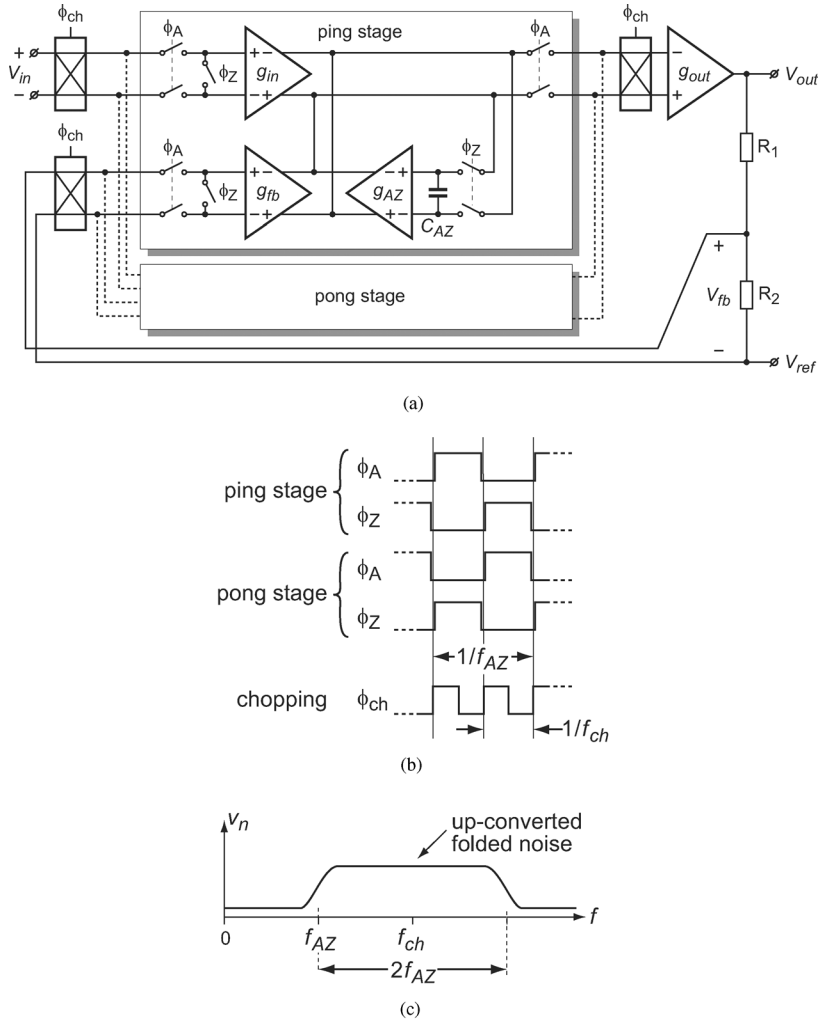

Fig. 9. A ping-pong auto-zeroed CFIA chopped at $f_{\rm ch} = 2f_{\rm AZ}$: (a) block diagram, (b) timing diagram, (c) input-referred noise PSD.

IV. Chopped Auto-Zeroed CFIAs

As described in the previous section, auto-zeroing provides a ripple-free output, but at the cost of increased low-frequency noise. Chopping, on the other hand, does not lead to increased noise, but is associated with ripple. In this section, it will be shown that the two techniques can be combined to obtain the best of both worlds.

A. Auto-Zeroed CFIA Chopped at $2f_{\rm AZ}$

One approach to remove the increased low-frequency noise due to auto-zeroing is to chop the auto-zeroed amplifier at a frequency $f_{\rm chop}$ above the auto-zeroing frequency $f_{\rm AZ}$. This technique was described in [8] for application in a ping-pong auto-zeroed opamp. It can be applied to the ping-pong auto-zeroed CFIA of Fig. 7 by equipping it with chopper switches at the input and output of the input stages [Fig. 9(a)], operated at $2f_{\rm AZ}$ [Fig. 9(b)]. Thus, the undersampled noise in the bandwidth from DC to $f_{\rm AZ}$ is up-converted to $2f_{\rm AZ}$, while the original noise PSD around $2f_{\rm AZ}$ is down-converted to DC [Fig. 9(c)]. The result is a CFIA with the same low-frequency noise as a chopped CFIA, i.e., the white-noise level of the transconductors, but without ripple, because the auto-zeroing eliminates the offset that would give rise to ripple in a regular chopped CFIA.

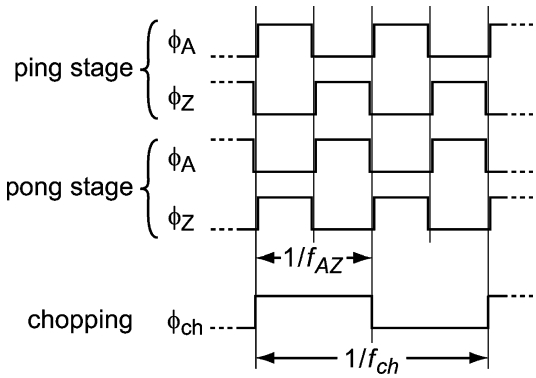

Fig. 10. Timing diagram of a ping-pong auto-zeroed CFIA chopped at $f_{\rm ch} = f_{\rm AZ}/2$.

While this approach effectively solves the noise problem associated with auto-zeroing, the added polarity reversals halfway through the amplification phase introduce extra glitches due to the finite bandwidth of the transconductors, similar to the problem described in Section III-B and illustrated in Fig. 6.

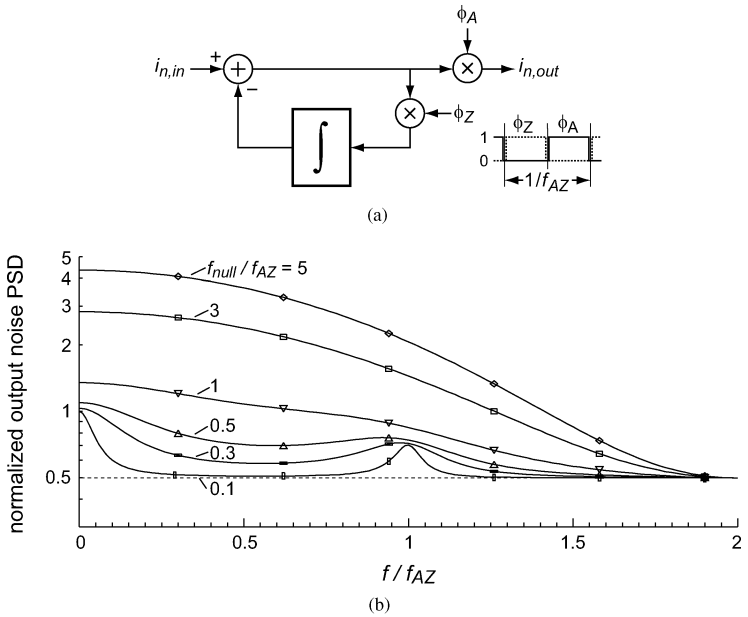

Fig. 11. Auto-zeroing with a slow-settling nulling loop: (a) model of the nulling loop, in which the bandwidth $f_{\rm null}$ is set by unity-gain frequency of the integrator; (b) simulated output noise PSD for different values of $f_{\rm null}$, normalized to the input noise PSD.

Therefore, this technique is not very suitable if a CFIA with a clean output is desired.

B. Auto-Zeroed CFIA Chopped at $f_{\rm AZ}/2$

Extra glitches can be avoided if the CFIA is chopped at a frequency $f_{\rm chop}$ *lower* than $f_{\rm AZ}$, e.g., at $f_{\rm AZ}/2$. The polarity reversals associated with the chopping then coincide with the transitions from zeroing to amplification, and merely imply that the ping-pong stages are inserted into the signal path with alternating polarity (Fig. 10).

While this slow chopping comes without extra glitches, it only improves the amplifier's low frequency noise if the noise PSD of the (unchopped) auto-zeroed amplifier at $f_{\rm AZ}/2$ is low, since noise at this frequency gets down-converted to DC by the chopping. As discussed before, this is not the case in a regular auto-zeroed amplifier, the noise PSD of which is elevated from DC to $f_{\rm AZ}$ due to aliasing [Fig. 7(c)].

By employing a *slow-settling* nulling loop, however, the bandwidth over which the noise PSD is increased due to auto-zeroing can be reduced. This can be explained using the model shown in Fig. 11(a). In this model, the combined output noise currents of $g_{\rm in}, g_{\rm fb}$ and $g_{\rm AZ}$ are represented as a noise current $i_{n,\rm in}$ that is input to the auto-zeroing circuit; $C_{\rm AZ}$ and $g_{\rm AZ}$ are represented as an ideal integrator, with a unity-gain frequency of $g_{\rm AZ}/2\pi C_{\rm AZ}$ that sets the nulling bandwidth $f_{\rm null}$; the switches ϕ_A and ϕ_Z, finally, are represented as multipliers driven by square-wave signals between 0 and 1 (with an average value of 0.5).

Fig. 11(b) shows the simulated output noise PSD, normalized to the input noise PSD, for different values of $f_{\rm null}$. These results were obtained by means of periodic steady state and periodic noise simulation using SpectreRF [16]. If $f_{\rm null}$ is significantly larger than $f_{\rm AZ}$ (which is the normal approach to allow for complete settling of the nulling loop within one auto-zeroing cycle), undersampling gives rise to an increased PSD near DC. By choosing $f_{\rm null} < f_{\rm AZ}/2$, this PSD can be reduced to the input PSD. As a consequence of this choice, the nulling loop will no longer settle within one cycle, but it will require multiple cycles to settle. This need not be a problem, as long as the state of the integrator in the nulling loop is preserved from cycle to cycle.

Irrespective of how low $f_{\rm null}$ is made, however, the output PSD around DC does not drop below the input PSD, which is twice as high as the PSD obtained around $f_{\rm AZ}/2$. This can be understood as follows. While near-DC components of $i_{n,\rm in}$ get eliminated by the auto-zeroing, components of $i_{n,\rm in}$ at or near the odd harmonics of $f_{\rm AZ}$ get down-converted to DC by the ϕ_Z multiplier. The resulting DC component at the input of the integrator is cancelled, due to the negative feedback loop, by an equally large DC component at the integrator's output, which also makes its way to the output through the ϕ_A multiplier. This multiplier *also* down-converts the components of $i_{n,\rm in}$ at or near the odd harmonics of $f_{\rm AZ}$, so that these components effectively get down-converted on both clock phases, and reach the output unattenuated, hence leading to a noise PSD around DC equal to the input noise PSD. This is twice the PSD that would be obtained without auto-zeroing, in which case an output PSD of half the input PSD would be obtained due to the 50% duty cycle.

Given the lower PSD at $f_{\rm AZ}/2$, it is now advantageous to chop at $f_{\rm AZ}/2$. The simulated effect of this is shown in Fig. 12. For small values of $f_{\rm null}$, the chopping reduces the low-frequency noise PSD to half the input noise PSD, which means that the auto-zeroing process no longer introduces additional noise. Fig. 12 also clearly shows that this slow chopping is only benificial if a slow-settling loop is used. For the typical choice of $f_{\rm null} = 5f_{\rm AZ}$, there is virtually no improvement in the low-frequency noise.

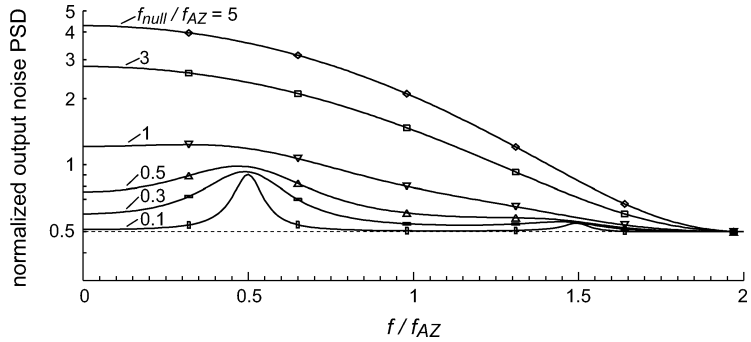

Fig. 12. Simulated noise PSD of an auto-zeroed amplifier chopped $f_{\rm ch} = f_{\rm AZ}/2$ for different values of $f_{\rm null}$.

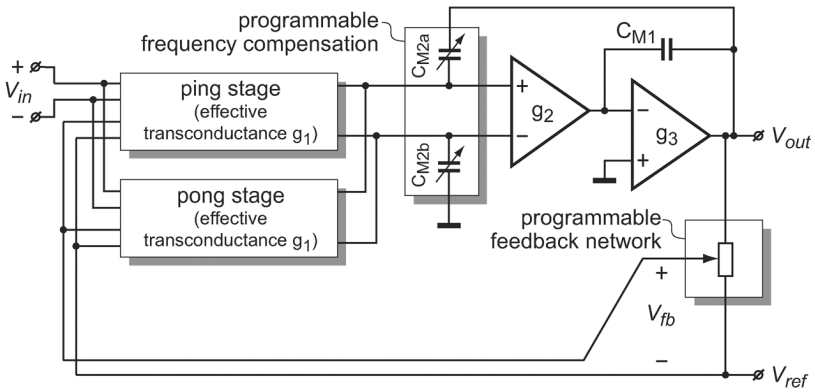

Fig. 13. Simplified circuit diagram of the instrumentation amplifier.

V. Circuit Implementation

A. Amplifier Architecture

Fig. 13 shows a simplified top-level circuit diagram of the implemented CFIA. It consists of three gain stages: a ping-pong auto-zeroed input stage chopped at half the auto-zeroing frequency, with an effective transconductance g_1; an intermediate stage g_2; and a class-AB output stage g_3. In this way, an overall open-loop DC gain in excess of 120 dB is obtained.

The closed-loop gain is defined by a precision on-chip resistive feedback network. The gain realized by this network can be programmed via a serial digital interface to be $10\times, 20\times, 50\times, 100\times, 200\times, 500\times$, or $1000\times$. Alternatively, the gain can also be defined by the user using an external resistive feedback network (not shown).

B. Programmable Frequency Compensation

The amplifier has been stabilized using nested-Miller compensation, comprising a local Miller capacitor C_{M1} around the output stage g_3, and a global Miller feedback consisting of capacitors $C_{M2a,b}$ (Fig. 13). So as to be able to maximize the closed-loop bandwidth for a given gain setting, the latter has been made programmable through the serial interface. In contrast with a fixed Miller compensation, which would result in a constant gain-bandwidth product, and hence a lower closed-loop bandwidth for higher gain settings, this programmable compensation network provides a bandwidth that is roughly independent of the gain setting. This is achieved by selecting a smaller capacitor for $C_{M2a,b}$ for higher gain settings.

C. Ping-Pong Stages

Fig. 14 shows a circuit diagram of one of the (identical) chopped ping-pong auto-zeroed input stages. The implementation differs in several aspects from the simplified version shown in Fig. 9(a). The most important difference is the addition of a settling phase ϕ_D, which prevents settling transients that occur after the transitions between the amplification phase ϕ_A and the zeroing phase ϕ_Z from reaching the output and from affecting the offset-nulling loop. During the settling phase phase ϕ_D, the output current of the stage is shorted rather than routed to the intermediate stage or to the nulling integrator. As illustrated in the timing diagram of Fig. 15, this is done for a short period of time after the rising edge of ϕ_A, so as to dissipate the transient associated with the transition from zeroing to amplification (indicated by a dotted line in Fig. 15), after which the current is routed to the intermediate stage during the output phase ϕ_O. Similarly, the transient associated with the transition from amplification to zeroing is dissipated after the rising edge of ϕ_Z, after which the current is routed to the nulling integrator during the nulling phase ϕ_N. The three signals ϕ_O, ϕ_N and ϕ_D thus form a triplet of non-overlapping clocks that steer the current at the output of the stage.

As shown in Fig. 15, the timing of the ϕ_O signal of the ping stage and that of the pong stage is arranged such that there is always one stage connected to the intermediate stage, and thus a continuous signal path is provided. To make time for the settling phases, the ϕ_A phases of the two stages become slightly overlapping. That is, just before one stage, say the ping stage, is removed from the signal path to be auto-zeroed, the input of the

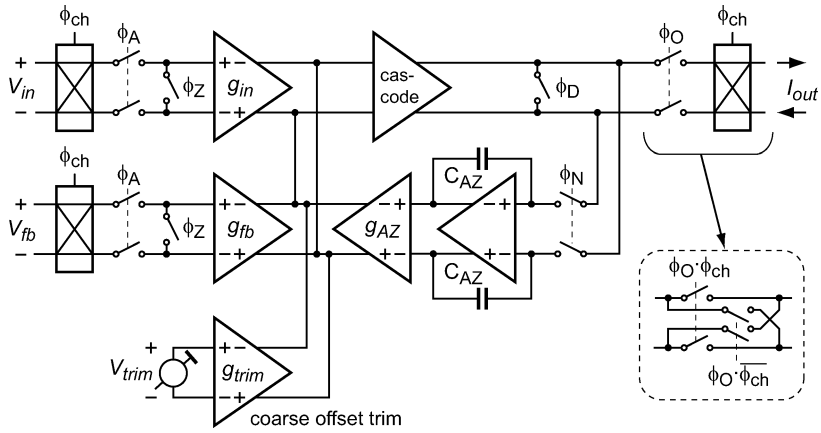

Fig. 14. Circuit diagram of the ping-pong stages.

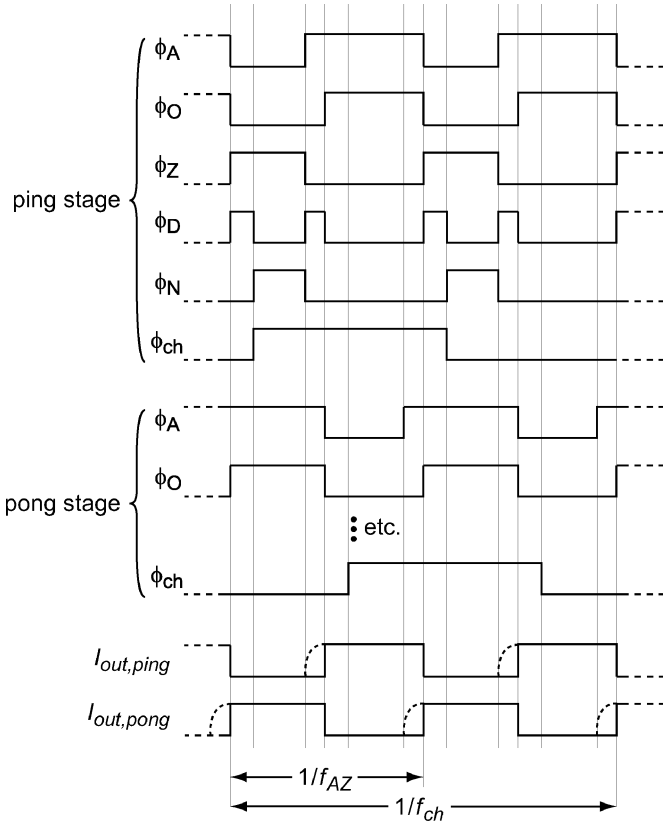

Fig. 15. Timing diagram of the ping-pong stages.

pong stage is already connected to allow its transconductors to settle. The output of the pong stage will only be connected into the signal path after this settling is completed.

The chopper switches $\phi_{\rm ch}$ that implement the chopping at $f_{\rm AZ}/2$ have been implemented inside the ping-pong stages. At the input, their timing is arranged such that they switch during the zeroing phase, when the input is not connected, so that the polarity reversal has the least effect. At the output, the chopper switches have been merged with the current-steering switches ϕ_O, as shown in the inset of Fig. 14, so that the polarity with which the output is connected simply toggles in successive amplification phases without introducing any new switching instants.

A further difference between the implementation shown in Fig. 14 and the principle shown in Fig. 9(a) is the use of an active nulling integrator. The advantage of this in comparison with the passive integrator of Fig. 9(a) is that the output current of the stage is always steered into a virtual ground, either that of the Miller-integrator formed by the intermediate stage and the output stage, or that of the nulling integrator. As a result, no differential signal is present on current-steering switches ϕ_O, ϕ_N, and ϕ_D. Thus, charge-injection mismatch is minimized, and parasitic capacitors at the output of the stage are kept at the same voltage, leading to smaller transients and lower residual offset. Common-mode transients are avoided by keeping the two virtual grounds at the same common-mode level by means of common-mode feedback loops (not shown).

Residual differential charge injection upon opening of switches ϕ_N leads to a small change in the voltage stored on the auto-zeroing capacitors $C_{\rm AZ}$. To reduce the resulting input-referred offset, $g_{\rm AZ}$ is made $200\times$ smaller than $g_{\rm in}$ and $g_{\rm fb}$. As a result, the steady-state output voltage of the nulling integrator is a $200\times$ amplified version of the initial input-referred offset. To ensure that large initial offsets do not saturate the output of this integrator, a coarse offset trim is implemented by means of an extra transconductor $g_{\rm trim}$ driven by a voltage $V_{\rm trim}$. This voltage is derived from an on-chip bandgap reference, and is programmed with a 4-bit resolution during production testing to minimize the output of the nulling integrator.

In order to increase both the open-loop DC gain of the signal path, and the DC loop gain of the nulling loop, a cascode stage has been included. The input of this stage forms the summing node at which the currents of transconductors $g_{\rm in}, g_{\rm fb}$, and $g_{\rm AZ}$ are added, while its output connects to the current steering switches ϕ_O, ϕ_N, and ϕ_D.

D. Input Transconductors

The input transconductor $g_{\rm in}$ in Fig. 14 converts the differential input voltage to a current while isolating the rest of the circuitry from the input common-mode voltage. In order to obtain an input common-mode voltage range (CMVR) that includes the negative supply rail, a simple pMOS differential pair

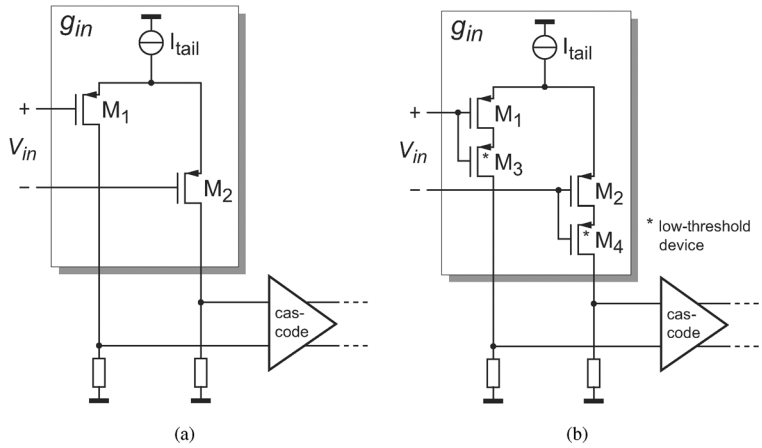

Fig. 16. Input transconductors based on (a) a simple differential pair, and (b) a differential pair cascoded using low-threshold transistors.

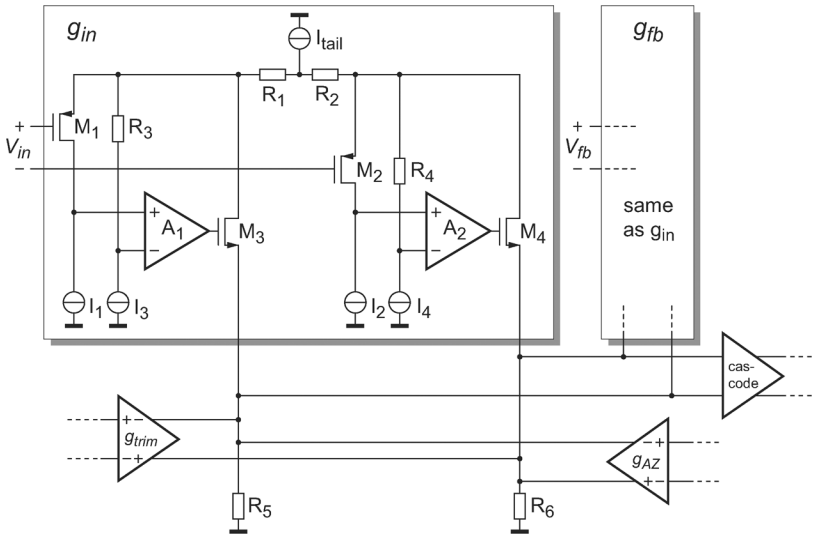

Fig. 17. Circuit diagram of the implemented input transconductors and their context.

coupled to a folded cascode stage could be used, as shown in Fig. 16(a). However, due to the finite output impedance of the pMOS transistors, the CMRR obtained using such an approach is quite poor. Moreover, it also leads to a common-mode dependent transconductance, and hence a common-mode dependent gain error.

In [4], these problems were addressed by including low-threshold cascode transistors in series with the drains of the (high-threshold) input transistors, as shown in Fig. 16(b). These cascodes increase the output impedance of the input transistors and keep them at a fixed drain-source voltage, thus making their transconductance common-mode independent. While this solution is very power-efficient (the cascodes share the bias current of the input transistors), it can be difficult (depending on the characteristics of the low-threshold devices) to guarantee a CMVR that includes the negative supply rail with good yield over process corners and temperature.

To prevent such yield issues, an alternative approach to keeping the input transistors at a fixed, common-mode independent drain-source voltage is used: they are incorporated in a local feedback loop, as shown in Fig. 17. The pMOS input transistors $M_{1,2}$ level-shift the input common-mode voltage up by a gate-source voltage, and reproduce the differential input voltage across resistors $R_{1,2}$. To make the gate-source voltages of $M_{1,2}$ both common-mode and signal independent, the drain currents of $M_{1,2}$ are set by current sources $I_{1,2}$, while their drain-source voltages are made equal to the voltage drop across resistors $R_{3,4}$ by amplifiers $A_{1,2}$. The differential current required to achieve this is drawn by $M_{3,4}$, equals $V_{\rm in}/(R_1 + R_2)$, and forms the output current of the transconductor.

The transconductance of $g_{\rm in}$ is thus determined by resistors $R_{1,2}$ rather than by the transconductance of $M_{1,2}$. By means of careful layout, these resistors can be matched to better than 0.1% with the corresponding resistors in the (identical) feedback transconductor $g_{\rm fb}$, so that the overall gain accuracy is determined by the overall resistive feedback network.

The output current of $g_{\rm in}$ is summed on load resistors $R_{5,6}$ with the output current of $g_{\rm fb}$, that of the offset-nulling loop circuitry provided by $g_{\rm AZ}$, and that of the offset trimming circuit provided by $g_{\rm trim}$. The resulting current flows through the (folded) cascode stage to the intermediate stage.

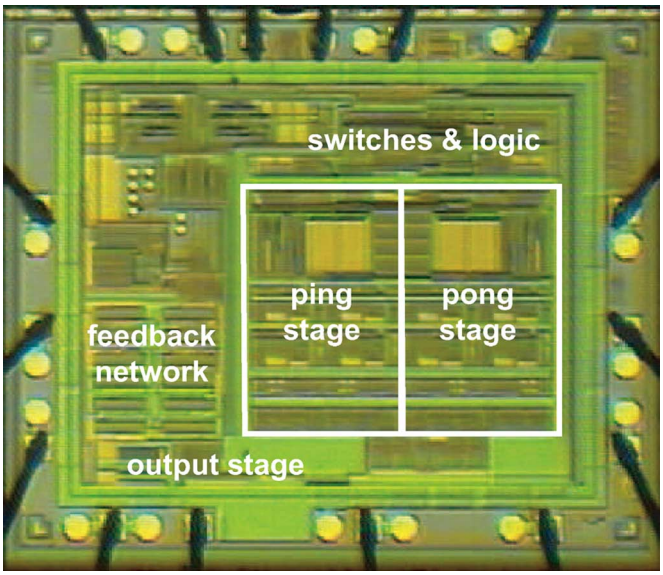

Fig. 18. Chip microphotograph of the instrumentation amplifier.

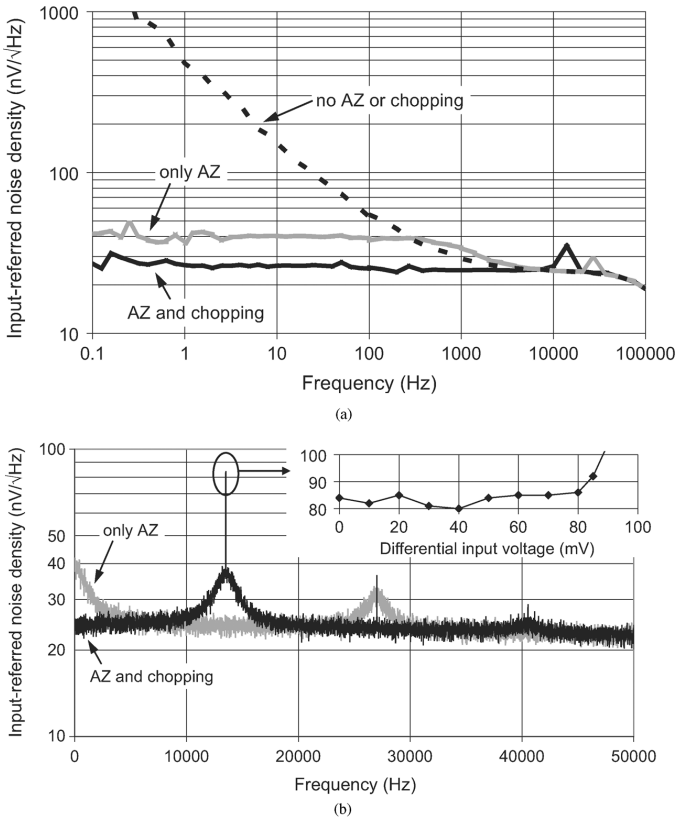

Fig. 19. Measured input-referred noise spectra: (a) on a logarithmic frequency scale; (b) on a linear frequency scale, with inset showing the PSD at the chopping frequency as a function of the differential input voltage. Measurements were performed with a resolution bandwidth of 10 Hz.

VI. Measurement Results

The CFIA was realized in an analog 0.5 μm CMOS process, and measures 2.5 mm^2, including the programmable-gain feedback network and serial interface (Fig. 18). It consumes 1.7 mA from a 3.0 V to 5.5 V supply, the majority of which is consumed in the input and feedback transconductors in the two ping-pong stages, which consume 275 μA each. The remaining

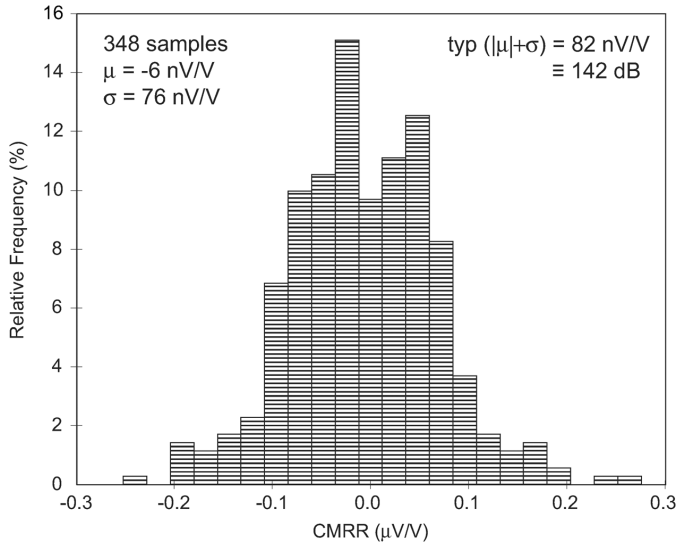

Fig. 20. Histogram of the amplifier's CMRR in a CMVR from 0.1 V below the negative supply to midsupply, measured with a gain of 1000x at a supply of 5 V.

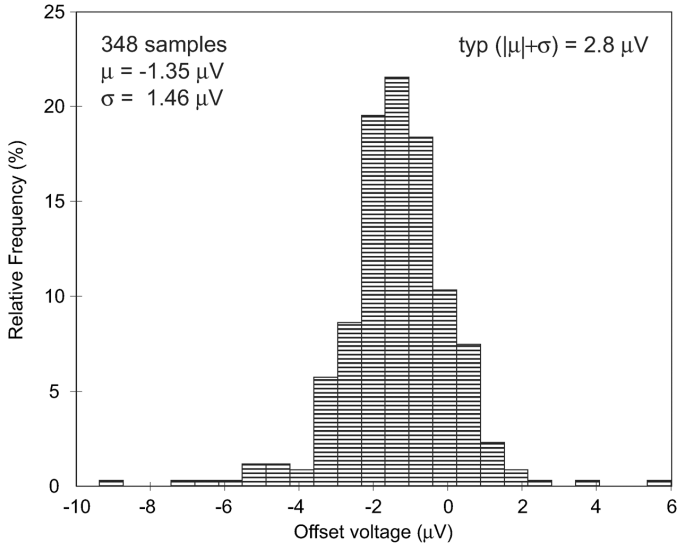

Fig. 21. Histogram of the amplifier's input-referred offset voltage measured with a gain of $1000\times$ at a supply of 5 V.

circuitry in each of the ping-pong stages, including the cascode stage, the nulling integrator and the trimming transconductor, consume about 100 μA. The intermediate stage and the output stage, finally, consume 30 μA and 290 μA, respectively.

Fig. 19 shows the input-referred noise spectrum measured at a gain of 1000. The timing of the ping-pong stages has been made programmable so as to be able to switch off the chopping and the auto-zeroing. With the ping-pong AZ and the chopping at $f_{\rm AZ}/2$ switched off, the amplifier's $1/f$ noise is clearly visible [Fig. 19(a)]. With only AZ enabled, a $\sqrt{2}$ times increase in the noise level within the bandwidth of the nulling loop is observed, as expected. Finally, when this is modulated to $f_{\rm AZ}/2$ by means of chopping, the noise level around DC equals the white noise level of 27 nV/$\sqrt{\rm Hz}$, as desired.

The inset in Fig. 19(b) shows the PSD at $f_{\rm AZ}/2$, which is a measure of the output ripple, as a function of the differential

TABLE I
PERFORMANCE SUMMARY AND COMPARISON WITH STATE-OF-THE-ART CURRENT-FEEDBACK INSTRUMENTATION AMPLIFIERS

	Witte VLSI '09 [14]	Wu JSSC '09 [4]	Fan ISSCC '10 [6]	MAX4209 [17]	This work
Process	0.7μm CMOS	0.7μm CMOS	0.7μm CMOS	–	0.5μm CMOS
Input offset voltage	2.5μV	5μV	2μV	3μV	2.8μV
Input offset drift	–	–	–	$10\text{nV}/^\circ\text{C}$	$3\text{nV}/^\circ\text{C}$
CMRR	130dB	>120dB	137dB	135dB	142dB
PSRR	114dB	>120dB	120dB	125dB	138dB
Gain error	0.6%	$\pm0.5\%$	0.53%	$\pm0.25\%$	$\pm0.1\%$
Input voltage noise	$42\text{nV}/\sqrt{\text{Hz}}$	$15\text{nV}/\sqrt{\text{Hz}}$	$21\text{nV}/\sqrt{\text{Hz}}$	$140\text{nV}/\sqrt{\text{Hz}}$	$27\text{nV}/\sqrt{\text{Hz}}$
Gain-bandwidth	640kHz	800kHz	900kHz	750kHz	800kHz
Supply current	0.325mA	0.23mA	0.143mA	0.75mA	1.7mA
Supply voltage	3.3V to 5.5V	5V	5V	2.85V to 5.5V	3.0V to 5.5V

input voltage for a gain of 10. Within the amplifier's differential input range, this PSD is independent of the input voltage, showing the effectiveness of the presented settling scheme.

Wafersort data from over 300 samples show a typical CMRR of about 140 dB (mean+sigma) in a CMVR from 100 mV below the negative rail to midsupply (Fig. 20), and a typical offset of 2.8 μV (mean+sigma, Fig. 21). The typical power-supply rejection at DC is 125 nV/V, or 138 dB. Over the temperature range of $-40\,^\circ$C to $125\,^\circ$C, the offset typically shifts by 0.5 μV, which is equivalent to a temperature drift of 3 nV/°C.

Table I summarizes and compares the performance of the CFIA, which is in line with the state of the art [4], [6], [14], [17].

VII. CONCLUSION

The first implementation of a current-feedback instrumentation amplifier has been presented that uses ping-pong auto-zeroing to reduce input-referred offset and $1/f$ noise to the μV level required by general-purpose precision applications. A new approach to combining auto-zeroing with chopping has been introduced that enables the elimination of the increased low-frequency noise associated with auto-zeroing. In contrast with earlier reported similar combinations, the amplifier employs a slow-settling offset-nulling loop to limit the bandwidth over which the auto-zeroing introduces under-sampled noise to less than half of the auto-zeroing frequency. The amplifier is then chopped at half of the auto-zeroing frequency to restore the low-frequency noise PSD to the white-noise level. This slow-chopping approach can be implemented simply by inserting the ping-pong stages of the amplifier into the signal path with alternating polarity. The key advantage of this approach compared to chopping at a frequency higher than the auto-zeroing frequency is that no extra switching transients are introduced, which is essential to obtain the transient-free output desired for a general-purpose precision instrumenation amplifier. To further reduce switching transients, the implemented CFIA employs an auto-zeroing scheme with settling phases that prevent settling transients due to the limited bandwidth of the input and feedback transconductors from reaching the output.

The CFIA was implemented in an analog 0.5 μm CMOS process. Noise measurements confirm the effectiveness of the proposed combination of auto-zeroing and chopping, as well as of the clocking scheme with settling phases. The precision performance of the CFIA is in line with the state of the art, with a typical offset of 3 μV and a CMRR of 140 dB in a common-mode voltage range that includes the negative supply.

REFERENCES

[1] B. J. van den Dool and J. H. Huijsing, "Indirect current feedback instrumentation amplifier with a common-mode input range that includes the negative rail," *IEEE J. Solid-State Circuits*, vol. 28, no. 7, pp. 743–749, Jul. 1993.

[2] J. F. Witte, J. H. Huijsing, and K. A. A. Makinwa, "A current-feedback instrumentation amplifier with 5 μV offset for bidirectional high-side current-sensing," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2769–2775, Dec. 2008.

[3] P. G. Blanken and S. E. J. Menten, "A 10 μV-offset 8 kHz bandwidth 4th-order chopped $\Sigma\Delta$ A/D converter for battery management," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 388–476.

[4] R. Wu, K. A. A. Makinwa, and J. H. Huijsing, "A chopper current-feedback instrumentation amplifier with 1 mHz $1/f$ noise corner and an AC-coupled ripple reduction loop," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, pp. 3232–3243, Dec. 2009.

[5] J. H. Huijsing, "Instrumentation amplifier developments," in *Proc. AACD*, Apr. 2008, pp. 105–136.

[6] Q. Fan, J. H. Huijsing, and K. A. A. Makinwa, "A 21 $\text{nV}/\sqrt{\text{Hz}}$ chopper-stabilized multipath current-feedback instrumentation amplifier with 2μV offset," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2010, pp. 80–81.

[7] M. A. P. Pertijs and W. J. Kindt, "A 140 dB-CMRR current-feedback instrumentation amplifier employing ping-pong auto-zeroing and chopping," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2009, pp. 324–325.

[8] A. T. K. Tang, "A 3 μV-offset operational amplifier with 20 $\text{nV}/\sqrt{\text{Hz}}$ input noise PSD at DC employing both chopping and autozeroing," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 386–387.

[9] A. T. K. Tang, "Enhanced programmable instrumentation amplifier," in *Proc. IEEE Sensors*, Nov. 2005, pp. 955–958.

[10] C. C. Enz and G. C. Temes, "Circuit techniques for reducing the effects of op-amp imperfections: autozeroing, correlated double sampling, and chopper stabilization," *Proc. IEEE*, vol. 84, no. 11, pp. 1584–1614, Nov. 1996.

[11] A. Bakker and J. H. Huijsing, "A CMOS chopper opamp with integrated low-pass filter," in *Proc. ESSCIRC*, Sep. 1997, pp. 200–203.

[12] A. Thomsen, D. Kasha, and W. Lee, "A five stage chopper stabilized instrumentation amplifier using feedforward compensation," in *Symp. VLSI Circuits Dig. Tech. Papers*, Jun. 1998, pp. 220–223.

[13] R. Burt and R. Zhang, "A micropower chopper-stabilized operational amplifier using a SC notch filter with synchronous integration inside the continuous-time signal path," *IEEE J. Solid-State Circuits*, vol. 41, no. 12, pp. 2729–2736, Dec. 2006.

[14] J. F. Witte, J. H. Huijsing, and K. A. A. Makinwa, "A chopper and auto-zero offset-stabilized CMOS instrumentation amplifier," in *Symp. VLSI Circuits Dig. Tech. Papers*, Aug. 2009, pp. 210–211.

[15] C.-G. Yu and R. L. Geiger, "An automatic offset compensation scheme with ping-pong control for CMOS operational amplifiers," *IEEE J. Solid-State Circuits*, vol. 29, no. 5, pp. 601–610, May 1994.

[16] K. Kundert, "Simulating switched-capacitors filters with SpectreRF," The Designer's Guide Community, Jul. 2006 [Online]. Available: http://www.designers-guide.org/Analysis/sc-filters.pdf

[17] Maxim Int. Prod. Rev. 0, "MAX4209 ultra-low offset/drift, precision instrumentation amplifiers with REF buffer," Sep. 2007 [Online]. Available: http://datasheets.maxim-ic.com/en/ds/MAX4208-MAX4209.pdf

Michiel A. P. Pertijs (S'99–M'06–SM'10) received the M.Sc. and Ph.D. degrees in electrical engineering (both *cum laude*) from Delft University of Technology, Delft, The Netherlands, in 2000 and 2005, respectively.

From 2000 to 2005, he was a researcher with the Electronic Instrumentation Laboratory of Delft University of Technology, working on high-accuracy CMOS smart temperature sensors. The results of his research have been applied in several commercial temperature sensors in cooperation with Philips Semiconductors. From 2005 to 2008, he was with National Semiconductor, Delft, where he designed precision operational amplifiers and instrumentation amplifiers. From 2008 to 2009, he was a Senior Researcher at IMEC/Holst Centre, Eindhoven, The Netherlands, where he worked on ultra-low-power electronics for wireless autonomous transducer systems. In 2009, he joined the Electronic Instrumentation Laboratory of Delft University of Technology as an Assistant Professor, working on sensor interface circuits and smart sensors. He has authored or coauthored one book, eight patents, and over 35 technical papers. His research interests include analog and mixed-signal electronics and smart sensors.

Dr. Pertijs is a member of the technical program committee of the European Solid-State Circuits Conference (ESSCIRC). He received the ISSCC 2005 Jack Kilby Award for Outstanding Student Paper, the IEEE JOURNAL OF SOLID-STATE CIRCUITS 2005 Best Paper Award, and the 2006 Simon Stevin Gezel Award from the Dutch Technology Foundation STW.

Wilko J. Kindt received the M.Sc. and Ph.D. degrees in electrical engineering (both *cum laude*) from Delft University of Technology, Delft, The Netherlands, in 1993 and 1999, respectively.

From 1993 to 1998, he worked as a researcher at the Electronic Instrumentation Laboratory of Delft University of Technology on Geiger mode avalanche photodiode arrays for spatially resolved single photon counting. In 1996 and 1997, he was a trainee at Exar Corporation, Fremont, CA. In 1999, he joined the imaging products group of National Semiconductor Corporation, Santa Clara, CA, where he worked on CMOS image sensors. In 2002, he transferred to the amplifiers product line, where he worked on precision operational amplifiers and current sense amplifiers. In 2004, he moved from National Semiconductor's headquarters in California to that company's design center in Delft, The Netherlands, where he is currently the design manager for the precision systems design group. He is the author of 15 technical papers and holds 15 U.S. patents.