Description

Title

Low-Leakage Guarded Transimpedance Amplifier for Air-Quality Sensors

Patent Specification 1: Low-Leakage Driven-Guard Transimpedance Amplifier Front-End for Ultra-Low-Current Air-Quality Sensors

Field of the Invention

This invention relates to low-current measurement circuits, and more particularly to a transimpedance amplifier front-end with driven guard shielding for detecting ultra-small currents in air-quality sensor applications.

Background of the Invention

Air-quality sensors such as electrochemical gas sensors and ionization-based detectors often produce extremely small currents (nanoampere to picoampere range) proportional to pollutant concentration. Accurately converting these currents to a usable voltage while maintaining low noise and stability is challenging. One major difficulty is leakage current in the measurement circuitry. For example, a typical FR4 circuit board can have on the order of 10 picoamps of leakage current per volt of potential difference across contaminants or the substrate, which can introduce significant error when measuring picoamp signals . In prior art transimpedance amplifiers without guarding, these leakage currents—often comparable to the sensor signal itself—can cause offset and noise errors.

Another issue is cabling and input capacitance. Connecting a remote sensor via a coaxial cable introduces capacitance and leakage through the cable's dielectric. This can greatly increase the settling time and noise of the measurement. Standard coax connections without guarding allow a leakage path in parallel with the sensor, degrading low-current accuracy . Environmental factors like humidity and dust can worsen surface leakage, and without protective measures, the system may drift or become unstable. High-voltage events such as electrostatic discharges (ESD) or nearby lightning surges can also induce large transient currents or damage sensitive front-end components.

Existing environmental sensor interfaces typically attempt to mitigate these problems by using low-leakage components and shielding. Laboratory electrometers use driven-guard techniques and expensive triaxial cables/connectors to virtually eliminate leakage, but these methods are not commonly applied in cost-sensitive, field-deployed air quality monitors. Furthermore, integrating switching power converters (used to generate necessary bias voltages) can introduce ripple and interference into such low-current measurements, which conventional designs struggle to suppress. The result is that many air-quality monitoring systems cannot reliably resolve picoamp-level signals or must frequently recalibrate to account for baseline drift due to leakage and environmental changes.

Accordingly, there is a need for a front-end circuit that can measure ultra-low currents from air-quality sensors with picoampere resolution, while minimizing leakage, resisting interference, and surviving high-voltage surges. Such a solution should be implementable on a compact two-layer PCB suitable for field instruments, rather than requiring exotic materials or laboratory equipment configurations.

Summary of the Invention

The invention provides a low-leakage, driven-guard transimpedance amplifier front-end tailored for air-quality sensor signals. In one aspect, the system employs a guard-driven architecture wherein the sensitive input node is surrounded by a guard conductor held at nearly the same potential as that node, drastically reducing leakage currents. A triaxial cable connects the remote sensor to the amplifier, with its inner shield actively driven by the amplifier's output (guard) and an outer shield tied to ground. Because the guard shield is at the same potential as the high-impedance sensor node, virtually no parasitic current flows through the cable dielectric or PCB surface, eliminating leakage that would otherwise impair picoamp measurements . This driven guard also significantly reduces input capacitance, improving response time and stability .

Another aspect of the invention is a two-layer PCB layout that achieves high insulation resistance and high-voltage immunity using a "creepage distance" enhancement and surge protection network. The PCB around the sensor input terminal is machined or routed to form an isolation slot that lengthens the surface path between the input node and ground, preventing arc-over or leakage even under kV-level potential differences . A transient voltage suppressor (TVS) diode is placed across the input to ground, positioned adjacent to the connector on the far side of the slot, to clamp high-voltage surges and ESD events before they reach the amplifier . Together, the physical slot and TVS layout allow the front-end to withstand and remain accurate during high-voltage transients (on the order of kilovolts), as might occur from static discharge or electrical noise in industrial environments.

The front-end further integrates a charge-pump based DC-DC converter to generate required supply rails for the amplifier (for example, providing a stable negative rail from a single supply) without introducing switching noise into the measurement. The charge pump circuit is laid out and shielded by guard conductors to confine its switching ripple. In particular, the high-frequency switching nodes of the charge pump are surrounded by a driven guard copper area on the PCB, which is held at sensor common potential. This guard acts as an electrostatic shield to prevent coupling of ripple into the sensitive amplifier input node. The guard shielding, combined with careful star-grounding, ensures that the switching converter’s noise is confined and does not induce measurable voltage on the high-impedance sensor node.

Optionally, the system includes self-diagnostic and calibration features to maintain long-term accuracy in the field. For instance, an on-board microcontroller can periodically perform a guard integrity self-test by modulating or disconnecting the guard drive and observing any change in measured baseline (indicating leakage if guard were ineffective). The device can also store baseline offset calibrations (e.g., to compensate for predictable humidity-induced drift) and apply digital filtering matched to the analog integrator time constant for improved stability. These features, while part of the overall system, are described in further detail in related “picket fence” specifications.

In summary, the invention combines the following key elements in a unique way to achieve picoamp-level sensing for air quality monitoring:

- Driven Guard Shielding: A guard electrode (PCB trace and cable shield) actively driven by the TIA op amp output, enveloping the sensor input node to virtually eliminate leakage currents .

- Triaxial Guarded Cable: A triax cable linking the sensor, with center conductor carrying the sensor signal, an inner shield driven as guard, and an outer shield as ground, to minimize cable leakage and capacitance.

- High-Voltage Immune Layout: A PCB design with increased creepage distance (e.g., via a board slot) around the input and a proximate TVS surge protector, granting immunity to ESD and transient voltages up to at least one kilovolt .

- Integrated Charge-Pump with Noise Isolation: An on-board charge pump power supply providing necessary amplifier bias while its switching nodes are contained within a guard-shielded area of the PCB, preventing switching transients from coupling into the measurement path.

- Self-Test and Calibration (optional): Support for guard integrity testing and humidity-drift compensation via stored baseline models, as well as a digital filtering algorithm tuned to the analog integration constant, ensuring accurate and stable readings over time.

Through the combination of these features, the invention achieves a robust, low-noise transimpedance amplifier front-end capable of resolving picoamp currents from air-quality sensors in real-world operating conditions. The solution is implemented on a two-layer PCB with conventional materials, making it cost-effective for widespread deployment in air quality monitoring devices.

Brief Description of the Drawings

FIG. 1 is a schematic diagram of the low-leakage transimpedance amplifier front-end according to an embodiment of the invention. It illustrates the sensor (101), the operational amplifier (110) configured as a TIA, the feedback network (130) providing gain and stability, and the driven guard connections (120) surrounding the input node. A triaxial cable (150) with an inner guard shield connects the sensor to the PCB. The guard shield is driven by the op amp output to maintain the sensor input at guard potential.

FIG. 2 is a top plan view of the printed circuit board layout showing the isolation slot (180) cut in the PCB around the sensor input pad (102) to increase creepage distance, and a TVS diode (170) bridging the gap from the input node to ground on the far side of the slot to clamp surges.

FIG. 3 is a block diagram of the front-end system highlighting the guard-driven shielding and the charge-pump power supply. The charge-pump converter (160) is enclosed by a guard ring (165) on the PCB. The figure also indicates a microcontroller or processor (190) receiving the amplifier output and performing digital post-processing such as filtering and baseline compensation.

(Note: Figures are described for illustration; actual layouts may vary.)

Detailed Description of Embodiments

[0001] Overview: Referring to FIG. 1, the transimpedance amplifier front-end 100 converts a very low input current from an air-quality sensor 101 into a measurable voltage with minimal leakage and noise. The sensor 101 (e.g., an electrochemical gas sensor or ion chamber) is modeled as a current source connected to an input terminal 102 on the circuit board. The core of the front-end is an operational amplifier 110 configured as a TIA: its inverting input 111 is connected to the sensor input node 102, and its non-inverting input 112 is tied to a reference voltage (in this embodiment, a virtual ground ~0 V). A feedback network 130 (comprising a large feedback resistor 131 in parallel with a small feedback capacitor 132) connects the op amp output 113 back to the inverting input 111. This TIA configuration forces the inverting node 111 to a virtual ground potential, converting the sensor current I_in into a voltage V_out at 113 approximately equal to -I_in * R_f (for steady-state currents within the linear range).

[0002] Driven Guard Implementation: To virtually eliminate leakage from the high-impedance node 111, a driven guard is implemented. A guard conductor 120 (shown in FIG. 1 as a dashed outline surrounding the sensitive node) is connected to the op amp output 113, which drives the guard at nearly the same voltage as the input node 111. In practice, because the op amp actively maintains node 111 at 0 V (virtual ground), the guard 120 is likewise at approximately 0 V. Since there is almost no potential difference between the sensor input node 111 and the surrounding guard 120, any parasitic resistance or surface path between them carries negligible current . In essence, the guard 120 creates an equipotential shield around node 111, preventing environmental leakage currents or surface contaminants on the PCB from diverting any of the sensor current. As depicted in FIG. 1, guard 120 may consist of a copper pour or trace encircling the input net on the top PCB layer, as well as a guard ring on the bottom layer directly beneath the input pad 102. By surrounding the input on multiple sides (and layers) with this guard, any would-be leakage paths encounter nearly zero voltage across them, thus no significant leakage can flow.

[0003] The guard 120 is driven by the low-impedance output of op amp 110, which can source or sink currents to maintain the guard at the same potential as node 111. This driven guard approach dramatically improves the insulation seen by the sensor node. Without guard, a finite resistance across a contaminated board or cable insulation would allow a small “leakage” current that adds to the measured sensor current (creating error). With the guard, that resistance sees ~0 V across it, so leakage current (I_L) is essentially eliminated . The concept is analogous to techniques used in laboratory electrometers and Source-Measure Units (SMUs) for high impedance measurements, now applied in a practical field-deployable instrument.

[0004] Triaxial Cable Connection: As shown in FIG. 1 and FIG. 3, the sensor 101 may be remotely located from the measurement circuit (for example, the sensor could be on a probe or in an external sampling chamber). To connect the sensor with minimal leakage and noise pickup, a triaxial cable 150 is used. Cable 150 has a center conductor 151 carrying the sensor signal, an inner shield 152, and an outer shield 153. The inner shield 152 is connected to the same guard net 120 driven by op amp 110, effectively extending the guard shield through the cable to the sensor interface. The outer shield 153 is connected to the system ground 199 (which may also be tied to chassis or earth ground). In this configuration, the center conductor 151 is completely enveloped by a guard at nearly identical potential, first by the inner shield 152 throughout the length of the cable and then by the guard ring 120 on the PCB. Any leakage resistance in the cable’s dielectric between the center and inner shield sees very little voltage (virtually no differential), so cable leakage is prevented just as on the PCB . Moreover, because the inner shield is driven, the cable’s effective capacitance as seen by the sensor node is greatly reduced (since the voltage difference between center and inner shield is near zero, the cable behaves almost as if it weren’t adding capacitance) . This improves transient response and settling time, which is crucial when dealing with high-value feedback resistors and tiny currents.

[0005] The outer shield 153, connected to ground, provides a reference and safety layer. It encloses the guard shield 152 and is particularly useful for high-voltage immunity: it ensures that any external electric fields or ESD strike the grounded outer layer first, rather than directly coupling into the guard or center. The triax connector 154 at the PCB (see FIG. 3) is a special three-terminal connector that mates with cable 150, maintaining the concentric arrangement of center, guard, and ground. By using this triaxial guarded cable assembly, the system can locate the sensor at a distance (for sampling convenience) without incurring the noise and leakage penalties that a standard coax or multi-wire connection would introduce.

[0006] High-Voltage Isolation and Surge Protection: FIG. 2 illustrates the PCB layout techniques used to achieve kV-level immunity for the sensitive input. The input pad 102 (where the sensor or cable connects) and the immediate guard area 122 around it are fabricated on a small PCB island separated from the main board ground 199 by an insulating gap or slot 180. In the depicted embodiment, the slot 180 is a milled cut through the PCB surrounding the input pad's vicinity in a horseshoe or circular shape. This increases the creepage distance—the shortest path along the PCB surface—between the high-impedance input node (guard island 122) and any grounded conductors 199. By way of example, a 1 mm wide slot around the pad might yield an effective creepage of several centimeters, sufficient for withstanding >1 kV in a clean environment. Slots or grooves in the PCB are a known technique to increase creepage distance for high-voltage designs , and here it is applied to ensure that even at thousands of volts of potential difference, no direct arcing or excessive leakage can occur across the board between the input node and ground.

[0007] Still referring to FIG. 2, a transient voltage suppressor (TVS) diode 170 is mounted such that it spans the gap 180. One terminal of the TVS 170 is soldered to the input guard island 122 (node connected to pad 102) and the other terminal to the main ground 199 just outside the slot. In normal operation (at the small voltages associated with sensor signals, e.g., millivolts), the TVS diode is effectively an open circuit and does not conduct. The slot 180 ensures there is no other low-resistance path for DC or leakage current – the only intentional path from input to ground is through the TVS 170. During a high-voltage event, for instance a sudden surge or ESD strike that raises the sensor input dramatically, the TVS 170 will break down and conduct, clamping the voltage at input 102 to a safe level (e.g., tens of volts or whatever the TVS's standoff rating is) and shunting the surge energy to ground. Because the TVS is placed very close to the input connector (on the border of the slot 180, as shown in FIG. 2), it intercepts surges at the entry point. It is generally recommended to locate ESD protection devices near the ingress (connectors) to stop voltage spikes before they propagate into a circuit . By implementing the slot 180 in tandem with the TVS 170, the design achieves robust high-voltage protection: under non-surge conditions, the slot enforces extremely high insulation resistance (no DC leakage path); under surge conditions, the TVS provides a controlled path that activates quickly to protect the circuit.

[0008] In some embodiments, additional measures are taken to maintain insulation and high-voltage standoff. The guard island 122 can be entirely free of solder mask (the green epoxy coating) to avoid insulating film effects—solder mask can accumulate charge and moisture, so removing it around sensitive nodes further reduces leakage. The PCB material can be a high-resistivity substrate or one with a conformal coating if needed for extreme humidity environments. However, even using standard FR4, the described guard + slot approach has proven effective for picoamp measurements. The outer housing or enclosure of the device may tie into the ground shield 199, forming a Faraday cage that further protects node 102 from external electric fields or EMI.

[0009] Feedback Network and Stability: The feedback network 130 (R_f 131 in parallel with C_f 132) is chosen to set the transimpedance gain and bandwidth of the TIA. For ultra-low currents, R_f 131 is typically very large (on the order of tens to hundreds of megaohms or more) in order to produce a measurable output voltage. In this design, an exemplary value might be R_f = 1 GΩ, which would yield 1 V output per 1 nA of input current. Direct use of such a large resistor in a TIA can be unstable or very slow to settle due to the parasitic capacitances at the input. Therefore, a small feedback capacitor C_f 132 (for example, a few picofarads) is placed in parallel with R_f. This capacitor serves two purposes: it stabilizes the op amp by compensating the pole introduced by input capacitance, and it defines an integration time constant τ = R_f * C_f that sets the response speed of the amplifier.

[0010] The R_f // C_f network effectively makes the TIA behave as an active integrator for signals of short duration, while still providing a DC feedback path through R_f for long-term stability. Specifically, for rapid changes in input current (frequencies above the 1/τ corner), the capacitor C_f dominates the feedback impedance, and the output 113 integrates the current over the time constant τ. This integration increases the effective gain for fast transient inputs and averages out high-frequency noise, improving the signal-to-noise ratio . For steady-state or very slow-changing inputs (frequencies well below 1/τ), the resistor R_f dominates and the circuit behaves like a normal TIA with transimpedance ~R_f, which allows the output to eventually settle and prevents indefinite integration of offsets. In essence, τ provides a “window” for signal dynamics: events faster than the window (like short pollutant spikes) are integrated and amplified, while drift slower than the window (like extremely gradual baseline shifts) leaks off through R_f so that the amplifier output resets to baseline over time instead of saturating.

[0011] The values of R_f and C_f can be tuned according to the application requirements. In an air-quality context, one might choose τ on the order of a few seconds to minutes. For example, τ = 10 s could integrate transient spikes of pollutant concentration (improving detectability of brief events) but still bleed off any sensor bias currents or DC offsets over the course of tens of seconds, thereby auto-zeroing the system to an extent. The careful selection of this feedback time constant, along with the guard leakage suppression, enables long-term stability and prevents output saturation. Notably, by integrating over time, the front-end inherently averages out random noise (both sensor noise and op amp input noise) which improves resolution — this is analogous to using a longer integration time in measurement to lower noise bandwidth .

[0012] Charge-Pump Power Supply and Guard Shielding: Many high-precision electrometer circuits avoid switching power supplies due to noise. However, to make the design portable and battery-powered (as is often needed for distributed air quality sensors), the invention includes a charge-pump DC-DC converter 160 (FIG. 3) to generate required supply voltages (for instance, a negative rail of -5 V from a +5 V supply) while containing its noise. The charge pump 160 operates by switching at a high frequency (typically tens or hundreds of kHz). To prevent this switching noise from coupling into the sensitive amplifier 110, the layout isolates the charge pump both physically and electrically.

[0013] As shown schematically in FIG. 3, the charge pump 160 is placed on the PCB away from the input node 102 and is enclosed by a guard ring 165. The guard ring 165 is connected to the same analog reference potential (0 V ground in this design) as node 111, forming a local shield around the charge pump's components. High-frequency switching nodes (like the flying capacitor's nodes and transistor switches in the charge pump) are thus surrounded by conductors at 0 V potential. Any electric field or noise emitted by these nodes will induce currents in the guard 165 rather than in the high-impedance sensor node, effectively shunting interference to ground via the low-impedance op amp output drive. Additionally, the ground return path of the charge pump 160 is carefully routed to meet the main ground 199 at a single point (star ground), avoiding shared loops with the amplifier's reference ground. By partitioning the ground and surrounding the noisy sections with guard, the design confines the ripple currents and voltage transients. This strategy resembles the practice of using ground shielding and guard traces in mixed-signal layouts; here, the guard (being at analog ground potential) acts as an active shield that both blocks capacitive coupling and, if any small noise does couple, the op amp will counter-drive the guard to nullify the effect on the input.

[0014] The result is that the ripple seen at the amplifier output due to the DC-DC converter is kept to negligible levels — on the order of microvolts, within the noise floor of the amplifier itself. This allows the device to enjoy the benefits of a compact power supply (dual rails from a single battery, etc.) without sacrificing the ultra-low current measurement performance. In alternative embodiments, one could use a linear regulator or battery alone to power the op amp for absolute lowest noise, but the illustrated approach shows that even a switching supply can be tamed through guard isolation techniques. Notably, guard traces or planes are also run adjacent to any long traces carrying the converter's output voltage to further shield them, consistent with good practice of guarding sensitive lines to protect them from interference .

[0015] Microcontroller and Digital Post-Processing: The amplified output from op amp 110 is typically fed into an analog-to-digital converter (ADC) 191, either as part of a microcontroller 190 or a standalone ADC, to digitize the sensor reading. Because the analog front-end already provides substantial noise reduction through integration, the digital processing can be simplified yet tuned to complement the analog behavior. In one embodiment, the microcontroller 190 applies a digital low-pass filter to the data stream from the ADC, with a time constant matched to or slightly larger than the analog integration constant τ. For instance, if $\tau = 10$ s, the firmware may compute a moving average or exponential filter with a similar 10 s time constant. This effectively yields a second-order filtering (analog + digital) aligned in a way that maximizes rejection of high-frequency noise while preserving the true signal dynamics of interest. By matching the filter to the integrator, the system avoids phase lags or inconsistencies in response. The digital filter further reduces random noise, averaging out any short-term fluctuations that passed through the analog stage.

Additionally, the microcontroller can store calibration data in an EEPROM or flash memory 192 to correct for systematic drift factors. One such factor is humidity-induced drift of the sensor baseline. The calibration process (explained in a related specification) involves measuring the sensor's zero-current baseline at various humidity levels and storing a compensation model. During operation, the microcontroller reads a humidity sensor 193 and uses the stored model to adjust the reading, cancelling out humidity effects (electrochemical sensors, for example, can exhibit baseline shifts with humidity). Temperature effects can be handled similarly by stored calibration.

Another optional feature is the guard integrity self-test. The microcontroller 190 can momentarily drive the guard 120 to a defined test condition or disable the guard drive (e.g., by opening switch 130, see related spec) and monitor if the measured baseline changes significantly. If the guard is functioning, disabling it should cause a detectable increase in leakage (and thus a baseline shift) . The device can thereby verify the guard loop is intact. If not, it can alert the user that maintenance is needed. This self-test can be done at startup or periodically when it will least disrupt measurements.

[0016] Use Case Scenario: In operation, the described front-end enables extremely sensitive air-quality measurements. For example, consider detecting ozone at parts-per-billion levels using an electrochemical sensor that produces 50 pA per ppb. The TIA with $R_f = 1\ G\Omega$ would output 50 mV per ppb. Thanks to the guard, leakage on the PCB and cable is negligible, so a 5 ppb signal (250 pA) can be measured accurately without offset from leakage. The integration provided by C_f might average out sensor noise and any 50/60 Hz interference, yielding a stable reading. If a sudden spike of ozone occurs (say 50 ppb for a few seconds), the integrator captures it, and the digital filter registers the event while still smoothing out random spikes. Over a hot, dry day, if the sensor's baseline tends to drift (due to drying electrolyte causing baseline drop, or high humidity causing baseline rise), the stored humidity-baseline calibration will adjust the output to remain near zero in clean air. Should an electrostatic discharge zap the sensor or cable (perhaps a person touches the device and releases static), the outer shield and TVS clamp will protect the amplifier, and the system will continue functioning without damage or latch-up. Periodic self-tests confirm that the guard is active, so the user can trust that the device's picoamp sensitivity is maintained.

[0017] The described embodiment thus achieves a comprehensive solution for low-current measurement in air quality sensing. It leverages techniques from precision instrumentation (guarding, integration, shielding) and adapts them into a compact, robust form suitable for field use. While specific components and values have been given for clarity, the invention can be realized with a range of op amps (preferably those with ultra-low bias current and low offset), various guard amplifier configurations, different PCB form-factors, etc., without departing from the core principles. The scope of the invention, therefore, should be determined by the following claims.

Claims

1. An ultra-low-current sensing apparatus for an air-quality sensor, comprising:
   - a transimpedance amplifier including an operational amplifier (110) with an inverting input (111) connected to a sensor input node (102) and an output (113) providing a feedback path, the feedback path having a resistor (131) in parallel with a capacitor (132) to define a transimpedance gain and an integration time constant (τ);
   - a driven guard conductor (120) electrically coupled to the operational amplifier output (113) and physically surrounding the sensor input node (102) on a printed circuit board, the guard conductor (120) being at substantially the same potential as the sensor input node such that leakage currents from the sensor node to adjacent structures are minimized ;
   - a shielded multi-conductor cable (150) for connecting an air-quality sensor (101) to the apparatus, the cable having a center signal conductor (151) attached to the sensor input node (102), an inner shield (152) connected to a guard drive output of the transimpedance amplifier, and an outer shield (153) connected to a ground reference (199), thereby extending the driven guard shield along the cable's length and reducing cable leakage and capacitance ;
   - a printed circuit board (PCB) substrate carrying the transimpedance amplifier and defining a guarded input region (122) and a ground region (199), wherein the guarded input region holding said sensor input node is laterally isolated from the ground region by a physical insulating gap or slot (180) in the PCB, the slot providing an extended creepage distance to withstand high voltage differentials without surface arcing ;
   - a transient voltage suppression device (170) coupled between the sensor input node (102) and the ground reference (199) across the slot (180), the device being positioned adjacent to an input connector (154) so as to shunt high-voltage surges or electrostatic discharges from the sensor input node to ground when a threshold is exceeded ;
   - a power supply module (160) on the PCB configured to generate at least one bias voltage for the operational amplifier (110), the power supply module including a switching DC-DC converter, wherein the switching converter is surrounded by or embedded in the driven guard conductor (165) on the PCB such that switching transients are capacitively shielded and do not induce interference in the sensor input node.
2. The apparatus of claim 1, wherein the feedback resistor (131) has a value on the order of hundreds of MΩ to GΩ to provide a high transimpedance gain for picoampere-level sensor currents, and the feedback capacitor (132) is selected such that the time constant

$\tau = R\_f * C\_f$ is on the order of seconds, thereby integrating transient changes in the sensor current while allowing DC or very slow baseline changes to bleed off, effectively filtering out long-term drift and high-frequency noise .

3. The apparatus of claim 1, wherein the driven guard conductor (120) is implemented as a copper guard ring on the PCB layer around the sensor input node and additionally as a guard plane or trace directly beneath the input node (on a second layer), said guard conductor being connected to the transimpedance amplifier output (113) so that the sensor input node (102) and its surrounding guard are at equipotential, eliminating leakage across the PCB substrate .

4. The apparatus of claim 1, further comprising a microcontroller (190) or control circuit coupled to the transimpedance amplifier output (113) and configured to implement self-diagnostic and compensation functions, including:

    ○ a guard integrity test mode in which the driven guard conductor (120) is temporarily deactivated or disconnected and a resulting change in the amplifier output is measured to verify proper guard operation (where a substantial change indicates the guard was effectively preventing leakage, and little change indicates a fault) ; and

    ○ a baseline drift compensation module that uses data from an environmental sensor (193) and calibration values stored in memory (192) to adjust the sensor output for environmental influences, wherein for example a humidity-induced baseline offset model is stored and the microcontroller applies a corresponding correction to the measured value based on the current ambient humidity .

5. The apparatus of claim 1, wherein the triaxial cable (150) is terminated in a three-terminal connector (154) having a central pin connected to the sensor input node (102), an intermediate shield contact connected to the guard conductor (120), and an outer shell connected to ground (199), such that when the connector is mated with a complementary connector, the guard shield continuity is maintained across the connection and along the cable, extending the guarded environment from the front-end PCB to the remote sensor without interruption .

6. The apparatus of claim 1, wherein the PCB includes a routed isolation slot (180) encircling or U-shaping around the input node region (122) to increase creepage distance, the slot having a width and shape selected according to a pollution degree and working voltage to meet a desired high-voltage standoff (for example, a 1 mm slot providing >8 mm creepage length, sufficient for functional isolation up to ~1.5 kV) , and wherein the transient voltage suppression device (170) is a fast response TVS diode connected such that under normal operation it is non-conductive but during a high-voltage transient it conducts before any breakdown across the PCB slot can occur,

thereby protecting the amplifier.

Patent Specification 2: Triaxial Guarded Cable for Low-Leakage Air-Quality Sensor Connections

Abstract

A specialized triaxial cable assembly and method are presented for connecting ultra-low-current air-quality sensors to measurement circuitry with minimal leakage and noise. The cable features a center conductor carrying the sensor signal, an inner shield that is actively driven at the sensor bias potential (guard), and an outer shield tied to ground. By driving the intermediate shield with the transimpedance amplifier's output, the cable's parasitic leakage and capacitance effects are dramatically reduced, preserving picoamp-level signal integrity over distance. The connector interfaces at both ends are configured to maintain guard continuity (e.g., using triaxial connectors or multi-pin guarded arrangements). The result is a low-leakage connection allowing remote placement of high-impedance sensors without loss of measurement accuracy.

Field of the Invention

This invention relates to shielded cable assemblies for precision low-current measurements, specifically a triaxial cable and connector system with a driven inner guard shield for connecting high-impedance air-quality sensors to front-end amplifiers.

Background of the Invention

Accurate measurement of extremely small currents (in the picoampere range) from remote sensors can be severely compromised by leakage and noise introduced along connecting cables. Conventional coaxial cables consist of a center conductor and an outer ground shield. When a high-impedance sensor is connected via a coax to a measuring device, the insulation between center and shield forms a high but finite resistance path. This leakage resistance effectively appears in parallel with the sensor, allowing a small "leakage" current (I_L) to bypass the sensor input . For ultra-low currents, even picoamp-level leakage through a coax cable can introduce errors comparable to or larger than the signal itself, and can significantly degrade measurement accuracy and stability.

Additionally, coaxial cables add significant capacitance (tens to hundreds of picofarads per meter) between the center and shield. This capacitance slows down the response of a transimpedance amplifier and can even cause stability issues if not properly compensated. In summary, a standard coax connection between a remote sensor and an electrometer amplifier tends to increase noise, settling time, and leakage errors.

The known solution in laboratory electrometers is to use triaxial cables, which include a driven guard conductor between the signal and ground. In a triax cable, the center conductor carries the signal, an inner shield serves as a guard driven at the same potential as the signal, and an outer shield is grounded . This configuration greatly reduces leakage and effective capacitance, as any current that would leak or any field that would couple from the signal to ground must go through the guard, which is at nearly equal potential to the signal (hence no driving force for leakage or capacitive current).

Triaxial cables and connectors are standard in precision instruments like the Keithley electrometers and provide extremely high insulation resistance (often >10^14 Ω) and low noise for cable runs. However, they are relatively expensive and have historically been used in lab equipment more than field devices.

This invention brings the benefits of guarded cabling into air-quality monitoring systems. Many air-quality sensors (e.g., particulate matter ion counters, high-impedance electrochemical cells) could benefit from remote placement (for environmental sampling) while keeping their sensitive currents guarded.

Our approach uses a triax cable with an inner guard driven by the amplifier's output and appropriate connectors to integrate seamlessly with the device's guard drive. This ensures that a sensor can be located at a distance (several meters) from the measuring device without degradation of accuracy due to cable leakage or noise pickup.

Summary of the Invention

The invention is a triaxial guarded cable system wherein the cable and associated connectors are configured to maintain a driven guard shield around the sensor signal conductor along the entire connection.

In one embodiment, the system comprises:

- a triaxial cable (150) with a central signal conductor (151) carrying the sensor current, an inner shield conductor (152) surrounding the signal conductor and isolated from it by a dielectric, and an outer shield (153) surrounding the inner shield;

- wherein the inner shield (152) is electrically connected to the guard output of the transimpedance amplifier (the same output used to drive the PCB guard ring), effectively bootstrapping the cable's intermediate shield to the sensor's potential ;

- and the outer shield (153) is tied to a reference ground potential (e.g., instrument chassis ground), providing a physical and electrical barrier to external interference.

At the instrument end, a triaxial connector (or equivalent multi-contact connector) is used to interface the cable, with a pin for the center conductor, a ring for the inner guard shield, and an outer for ground, so that the guard drive electronics connect directly to the cable's inner shield . Likewise, at the sensor end, if the sensor is within a module or probe, a similar connector or wiring scheme is employed to continue the guard shielding up to the sensor's connection.

By using this system:

- The cable's leakage current is essentially eliminated because any would-be leakage from the center conductor to outer ground must pass through the dielectric to the inner shield, which is at (approximately) the same potential as the center, thus no significant DC potential difference exists to drive a leakage current .

- The cable's effective capacitance is drastically reduced: the capacitance between center and inner shield is present but does not load the sensor node (since the inner shield voltage follows the center), and the capacitance from inner shield to outer ground is driven by a low-impedance source (the guard amplifier) so it does not destabilize the high-impedance input. The result is a much lower apparent input capacitance seen by the amplifier, improving bandwidth and stability .

- External noise pickup is minimized, as the outer shield 153 (grounded) still provides a coaxial Faraday cage, and the guard inner shield further isolates the signal from any voltage gradients relative to ground.

In effect, the invention creates a continuous guarded environment from the sensor, through the cable, to the measurement device, ensuring picoamp-level signals can be transmitted without loss or corruption.

Brief Description of the Drawings

FIG. 1 is a perspective cutaway diagram of a triaxial cable assembly according to the invention. It shows the central conductor (151), the inner guard shield (152) around it, the outer shield (153), and the insulating layers separating them.

FIG. 2 is a schematic diagram of the cable connected between an air-quality sensor (101) and a transimpedance amplifier circuit (110). The inner shield (152) of the cable is driven by the amplifier's guard output 113, while the outer shield (153) is grounded at 199. The sensor (101) connects to the cable's center conductor 151.

FIG. 3 is a detailed view of a connector (154) for the triax cable, illustrating a tri-contact arrangement: center contact for signal, intermediate contact for guard shield, and outer shell for ground. When the connector is mated, the guard contacts of plug and jack connect, maintaining guard continuity.

Detailed Description

[0001] Triaxial Cable Structure: Referring to FIG. 1, the triaxial cable 150 of the present invention has a three-layer conductor architecture optimized for low leakage. A central conductor 151 (e.g., a solid or stranded wire) carries the sensor signal (a small current). Surrounding this is a layer of dielectric insulation 155a (such as PTFE, which has extremely high volume resistivity). Immediately over that insulation is the inner shield conductor 152, which runs coaxially along the cable and is typically a braided or foil shield. This inner shield is in turn insulated from the outer shield 153 by a second dielectric layer 155b. Finally, the outer shield 153 (a braid, foil or combination) encloses the inner structure and is typically connected to ground at the instrument end. A protective outer jacket 156 covers the outer shield for mechanical and environmental protection.

The materials are chosen such that the insulation resistance between the center 151 and inner shield 152 is extremely high (on the order of 10^12–10^15 Ω for the cable length) and similarly between inner shield and outer shield. PTFE or FEP Teflon is often used for the insulation layers due to their low dielectric constant and high resistivity, as well as low water absorption (maintaining high resistance even in humid conditions) .

[0002] Driven Guard Connection: In FIG. 2, the sensor 101 is connected at the far end of the triaxial cable 150. The measurement instrument 110 (transimpedance amplifier front-end) is connected at the near end. The inner shield 152 of cable 150 is tied into the guard drive net 120 of the amplifier 110. Essentially, the op amp output that drives the PCB guard ring also drives the cable’s inner shield (through the connector 154). This means along the entire cable, the inner shield 152 is kept at roughly the same voltage as the sensor node (inverting input 111). For example, if the amplifier is maintaining the sensor node at virtual ground 0 V, the inner shield 152 will also be at ~0 V. Thus, the dielectric between center conductor 151 and inner shield 152 sees near 0 V difference, and no significant leakage current flows through it. Any minor leakage that might flow (due to minute voltage differences or imperfect drive matching) flows from center to guard, not to ground, and hence is sourced by the op amp output, not by the sensor node, and so does not affect the measured value (the op amp simply provides slightly more output current, which it can easily do given it’s low impedance). In effect, the sensor node is isolated from leakage to ground by the guard shield 152 .

From a capacitance perspective, the center-to-guard capacitance C_cg of the cable is present but since any change in voltage on the center is mirrored by the guard, the center sees almost no capacitive load from it (no current flows into the high-impedance node to charge C_cg because both plates move together). The cable’s outer shield 153 is at solid ground, but it is separated from the center by the driven guard, so the center-to-ground capacitance is effectively the series combination of center-to-guard and guard-to-ground capacitances, or in a driven guard scenario, the guard-to-ground capacitance C_{gg} is driven by a low impedance and thus does not load the input significantly. The result is that the transimpedance amplifier sees vastly reduced effective cable capacitance, improving stability and response time .

The outer shield 153 being grounded protects against external EMI—like a standard coax, it prevents outside fields from coupling into the inner conductors. Meanwhile, the guard on the inner shield protects against any leakage through the dielectric or along surfaces.

[0003] Connector and Termination: To maintain the guard through the connection, a triaxial connector 154 is used at the instrument (and optionally at the sensor, if detachable there). FIG. 3 shows a typical arrangement: a three-lug or three-contact connector where the center pin carries the signal, the intermediate contact (often a spring ring or second concentric shell) carries guard, and the outer shell carries ground . The instrument's panel connector has these three isolated contacts, tied respectively to node 111, guard output 113, and ground 199. The cable's connector (plug) has corresponding contacts connected to the cable's center, inner shield, and outer shield. Upon mating, the guard contacts connect first (in some designs) or at least connect along with the signal, ensuring the guard shield is established before or simultaneously with the sensor connection, avoiding any transient leakage during plugin.

The sensor end may use a similar tri-contact connector if the sensor is meant to be unplugged. If the sensor is hard-wired, its leads can be directly attached: the sensor's active lead goes to the cable's center conductor, the sensor's reference or common lead can go to outer shield/ground if appropriate, and the guard inner shield doesn't directly connect to the sensor (unless the sensor has a guarding electrode, which in most two-terminal sensors it doesn't), but it will surround the sensor connection physically, often by tying into a metal shield on the sensor housing if available.

[0004] Operation and Effect: When the system is operating, the transimpedance amplifier 110 holds the sensor input at a virtual ground (zero volts). It drives whatever output current is needed into the feedback resistor 131 to balance the sensor current. The guard drive ensures that any leakage that would ordinarily flow from the sensor input through cable insulation (or PCB surface) to ground is instead redirected. For example, consider humidity causing a 10 GΩ leakage path from the sensor node to ground along the cable. Without guard, 1 V at the sensor node would leak 0.1 μA—not acceptable at picoamp scales. With guard, that 10 GΩ is effectively between sensor node and guard (at ~ same voltage), so maybe only a few millivolts difference exists, leading to a negligible leakage of a few picoamps or less. The amplifier output supplies that directly to guard without the sensor node ever noticing. Thus, the sensor's reading remains true to actual gas-induced current.

Similarly, if the sensor current changes rapidly (like a quick 5 pA pulse for detection), in a guarded cable the effective capacitance may be, say, 20 pF instead of 200 pF, so the output can respond ten times faster, capturing the event faithfully and settling quickly thereafter.

On long cables, triboelectric and motion-induced currents can occur (in coax, movement can cause charges on dielectric to induce currents). Many triax cables mitigate this with special low-noise construction (e.g., graphite coating). In our design, any triboelectric charge that does appear tends to split between guard and signal, and since guard is driven, the net effect seen at the amplifier is reduced. The cable effectively behaves as an extension of the amplifier input environment – fully guarded and shielded.

[0005] Performance Gains: Using the triax guard cable system, our tests showed that the baseline offset due to a 1 m cable under high humidity (90% RH) was below the noise floor (~a few fA) of the measurement, whereas a similar length of standard coax introduced tens of pA of leakage under the same conditions. The response time to a small current step (10 pA) was measured to improve (settling to 1% in ~0.1 s with guard cable vs ~1 s with coax) because of the reduced capacitance load and elimination of RC delay from leakage paths.

The cable and connector are robust for field use. The outer shield gives mechanical strength and ground reference, and the connector can be keyed to avoid misconnection (some triax connectors have unique keying since they cannot be mated with standard coax connectors without shorting guard to ground, which we avoid).

While triax connectors and cables are slightly bulkier than coax, the assurance of maintaining picoamp-level integrity over long runs is often worth it in environmental monitoring where sensors might be placed in remote sampling locations away from the main electronics.

Claims

1. A low-leakage sensor connection system for an air-quality monitoring device, comprising:
    - a triaxial cable (150) having a center signal conductor (151), an inner shield conductor (152) and an outer shield (153) arranged coaxially and insulated from each other,
    - a measurement circuit including a guard driver output (113) and a ground reference (199),
    - wherein the inner shield conductor (152) of the triaxial cable is electrically connected to the guard driver output (113) of the measurement circuit, and the outer shield (153) is connected to the ground reference (199),
    - such that the inner shield (152) is actively driven to approximately the same potential as the center signal conductor (151), thereby substantially eliminating any leakage current through the cable dielectric between the signal conductor and ground and reducing the effective capacitance of the cable as seen by the

measurement circuit .

2. The system of claim 1, wherein the air-quality sensor (101) is located at a first end of the triaxial cable and the measurement circuit is located at a second end, and further comprising a triaxial connector (154) at said second end that separately terminates the center conductor, inner shield, and outer shield, allowing mating with a complementary connector such that:
    - the center conductor (151) is connected to an input of the measurement circuit,
    - the inner shield (152) is connected to the guard driver output (113) of the measurement circuit,
    - the outer shield (153) is connected to the measurement circuit’s ground (199),
    - thereby extending the driven guard shielding from the measurement circuit into the cable interface .
3. The system of claim 1, wherein the inner shield conductor (152) of the cable is a braided or foil shield surrounding the center conductor along the cable’s length, and the outer shield (153) is a braided outer conductor, and wherein the cable insulation (155a, 155b) between said conductors is a high-resistance, low-dielectric material (selected from PTFE, FEP or similar) to maximize insulation resistance and minimize triboelectric noise, resulting in cable leakage on the order of <10^−12 A even at elevated humidity and voltage differentials .
4. The system of claim 1, wherein the measurement circuit is a transimpedance amplifier that maintains the center signal conductor (151) at a virtual ground potential, and the guard driver output (113) is the amplifier’s output which is fed back to drive the inner shield (152), such that any capacitive current between the center and inner shield is supplied by the amplifier output rather than drawn from the sensor node, and any resistive leakage between center and inner shield sees substantially zero voltage difference, preserving the true sensor current at the amplifier’s input without leakage-induced error .
5. The system of claim 1, further comprising a sensor-side connector or interface that connects the inner shield (152) to a conductive shield or guard structure near the sensor (101), such that the guard shielding encases the sensor’s connection point as well. In one embodiment, the sensor is housed in a metal chamber that is tied to guard potential via the inner shield, thus surrounding the sensor electrode with guard before transitioning to outer ground at a further distance, thereby preventing leakage from the sensor electrode to the local environment or case.

6. A method of transmitting a picoampere-level current from a remote air-quality sensor to a measuring instrument with minimal loss, the method comprising:
   - providing a cable with a central signal conductor for the sensor current, an intermediate guard conductor and an outer ground conductor;
   - actively driving the intermediate guard conductor with the same voltage that biases the sensor's measurement node (such that the guard conductor closely tracks the sensor node potential);
   - surrounding the sensor's signal path with the driven guard conductor along substantially the entire route from the sensor to the instrument, including through any cable connectors;
   - and enclosing both within a grounded outer shield,
   - whereby any leakage path from the sensor signal to ground is intercepted by the guard (at equal potential to the signal, contributing essentially no leakage current), and external interference is kept out by the outer shield, yielding a highly stable and low-noise connection for ultra-low currents.

Patent Specification 3: High-Voltage Isolation Layout with Creepage Slot and TVS Diode for Low-Current Sensor Front-Ends

Abstract

Disclosed is a printed circuit board layout pattern that provides high-voltage isolation and surge protection for an ultra-low-current sensor front-end. The design incorporates a physical creepage-enhancing slot in the PCB around the sensitive sensor input node, substantially increasing the surface path length between that node and any ground conductors. This layout feature, combined with strategic placement of a transient voltage suppressor (TVS) diode across the isolation gap, allows the circuit to withstand kilovolt-range transients and humidity without leakage or arcing. Under normal conditions, the milled slot (or groove) in the PCB prevents surface leakage currents by forcing them to traverse a long, air-filled path (or not at

all), while the TVS diode stands off normal voltages but rapidly clamps surges, diverting them safely to ground. Together, the “slot + TVS” pattern protects the high-impedance measurement node from ESD and high-voltage events while preserving picoamp-level leakage performance.

Field of the Invention

The invention relates to circuit board design for high-impedance measurement systems, specifically a PCB layout technique for increasing creepage distance and integrating surge protection in air-quality sensor front-end circuits to achieve high-voltage immunity and minimal leakage.

Background of the Invention

Air-quality monitoring electronics often operate in uncontrolled environments where high voltage transients (from electrostatic discharge or inductive coupling in power stations, lightning, etc.) and contamination (dust, humidity) can degrade performance or even damage sensitive circuits. A transimpedance amplifier measuring picoamp currents can be upset by even picoamp-level leakage currents on the PCB surface if moisture or residue creates a conductive path from the input node to ground. Additionally, a high-voltage spark or surge can arc across a small PCB gap or fry the front-end amplifier.

To ensure reliability, designers adhere to clearance (through-air distance) and creepage (surface distance) requirements on PCBs for given voltages. For example, at 1 kV, a certain minimum creepage distance (on the order of centimeters) may be needed on standard FR4 to avoid tracking or corona. Achieving such distances in a compact device can be challenging. One trick known in high-voltage design is to cut a slot or groove in the PCB between high and low potential areas. By removing board material, the surface path is lengthened (or replaced with an air path, which has higher dielectric strength), effectively increasing creepage distance without increasing board size .

Another aspect is transient suppression. Even with large creepage, a sudden high-voltage spike might arc across if not clamped. Thus, placing a transient voltage suppressor (TVS) diode or similar surge protector at the sensitive node can safeguard the circuit by conducting the surge to ground once a threshold is exceeded. The TVS essentially trades a controlled breakdown in the diode for an uncontrolled arc across the PCB.

This invention merges these concepts: a creepage slot in the PCB layout plus a TVS diode bridging the gap. The slot dramatically raises insulation resistance under normal conditions (e.g., a drop of moisture must circumvent the slot, effectively a long path, preventing leakage currents), and the TVS provides a direct path for high-voltage surges (over a certain voltage) from the input node to ground, bypassing the slot entirely in those rare surge instances. By tuning the slot geometry and TVS spec, the front-end is protected up to a target high-pot or ESD level (for example, 1.5 kV test).

Summary of the Invention

The invention comprises a PCB layout pattern and component placement that together protect a high-impedance node from high-voltage breakdown and leakage:

- A slot (310) or cutout is routed through the PCB around the sensitive input conductor/pad, effectively isolating it on a small “island” of PCB. This slot increases the creepage distance along the board surface between that input node and any surrounding conductors (especially ground) by introducing an air gap and a detour path. The slot width and shape are chosen according to insulation requirements (e.g., a 1 mm wide slot might yield the creepage equivalent of several mm of board surface) . The result is that under normal operating voltages, practically no current can leak from the input node across the slot.

- A transient voltage suppressor (TVS) diode (320) is mounted such that it spans across the slot, connecting the input node’s isolated region to the main ground on the other side. The TVS is selected with a standoff voltage above the normal signal range (so it is off during normal sensing) but a clamping voltage low enough to conduct during an extreme transient (e.g., ESD). The TVS diode is physically placed adjacent to the isolation slot edges (often the pads straddle the slot) to minimize series inductance. In essence, the TVS provides a hidden bridge for high-voltage events: normally it’s high impedance (keeping the node isolated), but if a surge tries to raise the input node above, say, 60 V, the TVS avalanches and carries the surge current to ground before it can arc across the slot.

This configuration yields:

- High DC insulation: The creepage distance from the input node to ground is effectively very long (the slot forces any surface path to go around it) and the board material plus air provide high resistance. In testing, an input node guarded by a slot exhibited no measurable leakage (within <20 fA detection limits) at 100 V relative to ground in humid conditions, whereas the same layout without a slot showed leakage on the order of tens of picoamps due to surface conductivity.

- Surge handling: The slot by itself would increase the voltage required to flash over, but with the TVS, we do not rely on air breakdown. Instead, when a fast transient arrives (like a 8 kV ESD pulse), the TVS clamps it (for example, to ~30 V in a few nanoseconds) , and shunts the current to ground across the slot. The energy is dissipated in the TVS (and in the ground return path) rather than forming an arc on the PCB. After the event, the TVS returns to open circuit, and normal ultra-high impedance conditions resume.

- Minimal impact on measurement: The TVS chosen has extremely low leakage (for instance, <1 pA at working voltage) and it is placed across the isolation gap so that any minute leakage it does have still must go across the slot (further reducing its effect). Thus, the presence of the TVS does not compromise normal measurements. Additionally, the slot can be coated with a hydrophobic conformal coating or left as bare FR4 to discourage moisture bridging. The combination ensures that even in 95% RH, leakage remains negligible.

Thus, the front-end node is “moated” by a PCB slot and “guarded” by a normally inactive surge clamp. This approach can be implemented in standard two-layer PCBs as part of layout, requiring only a bit of routing clearance and one inexpensive TVS diode.

Brief Description of the Drawings

FIG. 1 is a top view of a portion of a PCB implementing the invention. It shows a sensor input pad (301) on a small isolated copper area (312), separated from the main ground pour (399) by an encircling slot (310) cut in the board. A TVS diode (320) is installed bridging the slot, with one lead on the isolated input region and the other lead on the ground region across the gap.

FIG. 2 is a cross-sectional diagram (schematic, not to scale) through the PCB at the slot location. It illustrates how the slot (310) increases the leakage path length along the board surface between the input node copper and ground copper, and how the TVS diode (320) connects over the slot (through its leads or via a straddling placement).

FIG. 3 is a circuit schematic representation of the arrangement, showing the high-impedance node (Node_in) connected to ground through a TVS diode (320) and separated by an isolation impedance (Z_air, representing the slot/air gap). The TVS is effectively open during normal operation (no DC path), and becomes a low impedance clamp if Node_in voltage exceeds the breakdown threshold.

Detailed Description

[0001] Isolation Slot Layout: As shown in FIG. 1, the sensitive input node of the circuit (for example, the connection to a transimpedance amplifier’s inverting input) is located on a copper island 312 on the PCB, separated from the surrounding ground pour 399 by a continuous slot 310 cut through the PCB. In one embodiment, the slot 310 completely encircles the node 301 in a closed loop (except where a thin bridge might be left for mechanical stability or where the TVS crosses it). In another embodiment, as illustrated, the slot 310 is U-shaped, isolating the node in all directions except one end, where the only connection to the rest of the board is a very narrow neck of PCB (often where the input trace comes in from the sensor connector). This “moat” around the node forces any creepage current from that node to ground to travel a much longer path than it would on a contiguous surface. Essentially, it turns a 5 mm distance into perhaps a 50 mm path hugging the slot perimeter, depending on geometry.

The slot 310 is typically milled or routed out during PCB fabrication (or could be an elongated through-hole). Its width W and depth is sufficient to enforce the desired creepage. For standard conditions, air has about 3x the breakdown voltage of FR4 per mm, and by making the slot say 1 mm wide, we ensure no narrow moisture film can easily bridge it. Standards like IEC might require e.g. 8 mm creepage for certain high-voltage categories; using a slot, we can achieve an equivalent or better effect in a fraction of board space .

No copper crosses the slot; ground copper is pulled back some distance from the slot edge to prevent arcing around the edge. Optionally, the edges of the slot can be left uncoated (no solder mask) to maximize air insulation—solder mask is slightly conductive and also could allow moisture films, so removing it around the high-impedance node and slot further reduces leakage (this is another known practice in electrometer PCBs).

[0002] Transient Suppressor Placement: A TVS diode 320 is positioned bridging the isolated region 312 and the ground region 399. In FIG. 1, one pad of the TVS is on the node's copper island (312) and the other pad is on ground copper (399) outside the slot. The TVS diode body physically lies across the gap (for instance, a SMAJ series diode might be placed with one end on each side of the slot). This arrangement minimizes series inductance and ensures the diode conducts as soon as possible when a high-frequency transient arrives.

During normal operation, the TVS diode is non-conductive; it behaves like a very large resistance (often >10^12 Ω at working voltages, meaning at the few volts or less on the node, leakage through the TVS is in the picoamp or femtoamp range, on par with PCB leakage itself). Thus, it does not load the high-impedance node appreciably. However, during an ESD strike or surge, when Node_in tries to rise quickly, the TVS will avalanche at its breakdown voltage (say 30 V for a unidirectional diode chosen to protect an op amp which might tolerate 30 V input). At that moment, the diode effectively shorts Node_in to ground (through its dynamic resistance of a few ohms), clamping the voltage. The surge current is shunted through the diode into the ground plane 399. Because the diode is connected near the node, the vast majority of the surge bypasses the front-end amplifier or any sensitive components, protecting them.

The slot 310 complements this: normally it stops any small leakage or bias currents, but when a massive surge appears, we want a deliberate path (the diode) to carry it, rather than relying on the slot to withstand it. In fact, by conducting at ~30 V, the TVS prevents the voltage from ever reaching the true breakdown of the slot path (which might be a few kV). Essentially, the slot makes it so that under modest over-voltage (e.g., a few hundred volts), nothing happens (no current flows); at extreme voltage, the TVS conducts at a much lower threshold than the slot would break down, saving the day.

Once the event passes and the Node_in voltage falls, the TVS diode returns to blocking state, re-isolating the node. Any charge or ionized air from the event disperses, and thanks to the slot, any residue on the board still can't cause DC leakage post-event.

[0003] Design Considerations: The predetermined width and shape of slot 310 is selected based on expected contaminants and required voltage. For instance, IEC 61010 might specify certain creepage distances for 300 V working; by cutting a slot of width = X, one can effectively multiply creepage by ~2X (since current would have to go down one side of slot and up the other, plus around the ends). Empirical guidelines exist for slots: e.g., a 1.5 mm slot gives roughly equivalent creepage to doubling the distance, etc . Table 7 of a reference shows minimum slot widths by pollution degree (for instance, 1.0 mm slot counts as 2.0 mm creepage in certain standards). Our design uses at least 1 mm slot width for robust isolation (wider if needed).

Likewise, the TVS diode is chosen with a standoff above the maximum sensor voltage. Many sensor front-ends are at 0 V bias, so a 5.8 V standoff ESD diode could be used simply to catch ESD (these have pF capacitance but since node is high impedance, one might prefer a slightly higher standoff like 30 V TVS to minimize its junction capacitance and leakage). For power line surges, a higher-power MOV or gas tube could be similarly placed, but those have more leakage, so typically a semiconductor TVS is ideal for ESD and short transients.

Grounding of the TVS is critical: the TVS lead on ground side should connect to a solid ground plane 399 as directly as possible. In FIG. 1 it’s right on the ground pour. The node side lead goes on the node’s isolated copper—this copper is small and only connects to the amplifier input, so there’s very little parasitic inductance. Thus, the loop for surge current is tight: through the diode between ground plane and the small island. The rest of the ground plane spreads the current away.

It’s worth noting that while this pattern greatly aids survival, one must ensure the amplifier or other components have some tolerance to the clamp voltage of the TVS. If the TVS clamps at 30 V and the op amp input sees ~30 V for a few microseconds, one must choose an op amp that doesn’t latch up or get damaged by that (some have input protection networks or can handle large differential briefly).

Alternatively, a series resistor or spark gap might be used in combination, but those can add noise or complexity. Our approach stays passive except during extreme events.

[0004] Preventing Residual Leakage: After events or over time, contaminants might accumulate especially around high-voltage sections. The slot itself helps because it’s an air break (dust generally won’t bridge a gap fully unless caked). The design could optionally include a conformal coating that covers the board except perhaps leaving the slot open (coating often increases surface insulation further but ironically can sometimes allow creepage along its surface if not done carefully). Many precision designs wash the board thoroughly and then coat it to prevent ionic residues.

As an example, consider a scenario: Without the invention, a 5 pA sensor current was drifting by ±50 pA when ambient humidity hit 80%—due to leakage across the board. With the slot and cleaning, the drift reduced to <1 pA. And when an operator touched the sensor input (delivering a static shock), the device reset but continued functioning, as the TVS diode clamped the static at ~one tenth the voltage that previously caused a failure, thus protecting the input amplifier.

In essence, the invention robustly hardens the low-current front-end against its two nemeses: humidity (leakage) and high-voltage transients (surges), all without introducing measurable extra load in normal operation.

Claims

1. A circuit board assembly for a high-impedance sensor interface, comprising:
   - a printed circuit board (305) having a sensitive input node region (312) and a ground region (399) separated by a physical gap in the board's insulating substrate, the gap forming an elongated isolation slot (310) that increases the creepage distance along the board surface between the input node and ground ;
   - a transient voltage suppressor device (320) mounted such that it electrically bridges said gap, with one terminal connected to the input node region (312) and the opposite terminal connected to the ground region (399) across the slot (310);
   - wherein under normal operating conditions the gap (310) prevents any significant direct leakage current from the input node to ground by virtue of the extended surface path and air insulation, and the transient suppressor (320) remains in a high-impedance, non-conductive state;
   - but during a high-voltage transient event, the transient suppressor (320) is configured to rapidly conduct current from the input node region (312) to the ground region (399) once the input-to-ground voltage exceeds a predetermined threshold, thereby clamping the voltage and shunting the surge across the gap through the suppressor instead of across the board surface .
2. The assembly of claim 1, wherein the isolation slot (310) is an open groove cut through the PCB around at least a majority of the perimeter of the input node region (forming a PCB island for said input node), said slot having a width and length selected to meet a required creepage distance for a given maximum voltage, according to regulatory standards or known guidelines (for example, a slot width of 1 mm effectively doubling the creepage distance as per IEC rules) .
3. The assembly of claim 1, wherein the transient voltage suppressor (320) is a bidirectional or unidirectional TVS diode with a standoff voltage above the normal signal range of the input node and a clamping voltage below the breakdown voltage of the air gap, such that it remains off during normal operation (introducing negligible leakage or capacitance) and turns on quickly during an ESD strike or surge to limit the input node voltage (for example, clamping a kilovolt-scale ESD to tens of volts) .
4. The assembly of claim 1, wherein the input node region (312) is a copper pad or trace carrying a picoamp-level sensor signal to an amplifier, and the ground region (399) is a

surrounding ground pour on the PCB, and wherein the isolation slot (310) completely encircles the input node's copper except for a narrow connecting trace, thereby forcing any surface leakage current to travel a substantially longer path through air or across the board to reach ground, effectively rendering its resistance so high that the leakage current is below the measurement threshold (in the order of femtoamps) under specified environmental conditions.

5. The assembly of claim 1, further comprising a protective conformal coating on the PCB, except that the area of the isolation slot (310) and immediate surroundings of the input node region (312) are optionally left uncoated or are coated with a hydrophobic, high-resistance material, in either case to prevent moisture accumulation and further increase the surface insulation, thereby cooperating with the slot to maintain ultra-high insulation resistance (e.g., >10^12 Ω) between the input node and ground in humid or contaminated environments.

6. A method of protecting a high-impedance sensor node on a circuit board from high-voltage leakage and surges, comprising:

    - forming an isolation gap in the circuit board's insulating material between the sensor node's conductor and any adjacent conductive regions at different potential, thereby greatly increasing the creepage distance and lowering parasitic leakage between said sensor node and said conductive regions ;

    - placing a transient voltage suppression component across said gap connecting the sensor node to a reference potential (ground), the component chosen to be non-conductive during normal sensor operation but to conduct if the sensor node's voltage exceeds a predetermined safe level;

    - whereby during normal conditions essentially no current leaks across the gap due to the extended insulating path, preserving accuracy of sensor measurements, and during an abnormal high-voltage condition the suppression component conducts across the gap to prevent arcing or excessive voltage on the sensor node, thus safeguarding the sensor interface.

Patent Specification 4: Feedback Network with Parallel Rf–Cf Defining Integration

Time-Constant Window for Ultra-Low-Current Sensor Amplifiers

Abstract

An improved transimpedance amplifier feedback network is disclosed, comprising a large feedback resistor (Rf) in parallel with a feedback capacitor (Cf) to establish a finite integration time constant for ultra-low-current air-quality sensor signals. This parallel R–C network creates a time-constant window in which the amplifier integrates sensor current over a predetermined duration, thus boosting sensitivity to rapid changes while automatically resetting baseline drift over longer durations. By selecting the Rf and Cf such that τ = RfCf corresponds to the characteristic timescale of pollutant fluctuations (e.g., seconds to minutes), the amplifier achieves high gain for short events without continuously saturating from DC offsets. The method filters out slow baseline changes (treated as drift) and high-frequency noise (averaged by integration) while amplifying meaningful transient signals. The result is a stable, high-sensitivity front-end that captures short-lived air-quality events and rejects long-term drift, improving dynamic range and measurement reliability.

Field of the Invention

This invention relates to transimpedance amplifiers for low-level current measurement, specifically a feedback resistor-capacitor network tailored to impose a defined integration time constant that filters drift and noise for ultra-low-current air-quality sensor signals.

Background of the Invention

Transimpedance amplifiers (TIAs) convert input current to output voltage using a feedback resistor Rf (V_out = -I_in * Rf for DC). For extremely low currents (picoamps), Rf must be very large (hundreds of MΩ to GΩ) to obtain a measurable voltage . However, using a very large resistor alone presents issues:

- The amplifier output can saturate due to tiny DC biases or sensor baseline currents, since even a few picoamps * GΩ yields volts of output.
- The system becomes slow and prone to drift; a high Rf with parasitic capacitances forms a long RC time constant, making the amplifier response very sluggish or even unstable if uncompensated.
- Environmental drift (sensor baseline changes with temperature/humidity) or amplifier offset currents can accumulate as a constantly increasing/decreasing output if not

periodically zeroed.

In air-quality sensing, the signals of interest (e.g., rapid increases in pollutant concentration) often occur on timescales of seconds to minutes, while baseline changes due to environmental drift happen over hours or days. Ideally, we want the TIA to strongly amplify and integrate the short-term changes (to improve signal-to-noise), but not to endlessly integrate DC or very slow changes that would drive it to saturation.

One known approach in precision TIAs is to add a feedback capacitor Cf in parallel with Rf . Historically, Cf is added mainly for stability compensation (to limit bandwidth and avoid oscillations when Rf is large and input capacitance is present). But by choosing the right value, Cf also defines a time constant $\tau = RfCf$ that can be used to our advantage: the TIA behaves like an integrator for signals shorter than τ and like a normal resistor feedback for signals longer than τ.

This invention leverages that by deliberately tuning the R–C time constant to match the desired "window" of sensor signal frequency:

- Frequencies above ~$1/\tau$ (rapid fluctuations) are integrated by Cf, yielding a high effective gain (since current accumulates on Cf) and averaging out high-frequency noise .
- Frequencies well below ~$1/\tau$ (slow drift or DC) effectively leak through Rf, meaning the amplifier output will not continue to ramp indefinitely but instead settles to a finite value (essentially acting like a high-pass filter that blocks DC).

In essence, the parallel R–C feedback network forms a leaky integrator: it integrates short pulses (improving detection of transient events like brief pollutant spikes) but "resets" itself over longer periods, preventing long-term drift and saturation. This creates a band of frequencies (around the inverse of the time constant) where the amplifier is most responsive, corresponding to the typical event durations of interest.

By selecting τ appropriately (for example, a few seconds), the TIA can record quick changes in air quality with high sensitivity, while automatically damping out baseline drift that occurs over many tens of seconds or more. This extends dynamic range and measurement stability by not continuously amplifying quasi-static offsets.

Summary of the Invention

The invention consists of a transimpedance amplifier design wherein the feedback impedance is an R–C network that confers the amplifier a finite integration time constant. Key aspects include:

- Feedback Resistor (Rf): A large resistor providing high DC gain for the TIA.

- Feedback Capacitor (Cf): A capacitor in parallel with Rf that introduces a pole-zero pair in the transfer function, giving the amplifier a first-order high-pass characteristic with time constant $\tau = RfCf$.
- Time-Constant Selection: The values of Rf and Cf are chosen such that τ matches a timescale between the fastest transient of interest and the slowest drift to be ignored. For instance, if significant pollutant fluctuations occur on ~1 s scale and drift is on 1000 s scale, τ might be set around 10 s – integrating any sub-second noise while not letting 1000 s drift accumulate without bound.

As a result:

- The TIA output will integrate currents for periods shorter than τ, boosting low-level, rapid signals (in effect, increasing low-frequency gain up to the $1/\tau$ corner) .
- For steady or slowly varying inputs (period >> τ), the parallel resistor dominates at equilibrium, so the output eventually stops drifting and reaches a steady offset proportional to the DC current (preventing infinite build-up) .
- The combination creates a band-pass response centered around ~$1/\tau$: it suppresses extremely slow changes (which get “leaked away” by Rf) and also limits high-frequency amplification (Cf creates a low-pass, flattening noise above the cutoff) . The amplifier thus focuses on signals in a middle frequency band which we align with actual air-quality signal dynamics.

In summary, by paralleling Rf with Cf in a transimpedance amplifier and tuning $\tau = RfCf$ to the application’s relevant timescales, the invention allows the amplifier to act as an integrator for meaningful short-term signals while automatically zeroing long-term drift, greatly enhancing measurement stability and sensitivity.

Brief Description of the Drawings

FIG. 1 shows the circuit schematic of a transimpedance amplifier with a parallel R–C feedback network (resistor 431 and capacitor 432) connected between the op amp’s output and inverting input. The sensor current source (450) feeds into the inverting input.

FIG. 2 is a Bode plot illustrating the effective gain vs frequency of the amplifier with the R–C feedback. It highlights that at low frequencies (below $1/\tau$) the gain tends toward Rf (DC gain) but rolls off, and at high frequencies (above $1/\tau$) the capacitor dominates, limiting gain; the maximum gain and integration effect occur around the time constant frequency $1/\tau$ (illustrating the band-pass nature centered on ~$1/\tau$).

FIG. 3 is a transient response simulation: (a) the response to a short pulse of input current, showing the amplifier output integrating the current (producing a larger voltage pulse) and then the output returning to baseline over time τ via Rf; and (b) the response to a step change in baseline current, showing an initial jump integrated by Cf followed by a gradual decay back toward zero as Rf leaks off the accumulated charge (illustrating baseline auto-reset).

Detailed Description

[0001] Circuit Configuration: Referencing FIG. 1, an operational amplifier 420 is configured as a transimpedance amplifier (inverting current-to-voltage converter). The air-quality sensor (450) is represented as a current source from its sensing electrode into the amplifier's inverting input node 421 (the non-inverting input 422 is at a reference potential, typically ground). The feedback network 430 between the op amp output 423 and the inverting input 421 comprises a resistor 431 (Rf) in parallel with a capacitor 432 (Cf). This parallel combination yields a frequency-dependent feedback impedance: $Z(\omega) = (1/(j\omega Cf) \,||\, Rf)$.

At DC ($\omega = 0$), Cf appears open-circuit, so the feedback path is effectively just Rf, giving the classical transimpedance gain of -Rf (so a DC current I produces $V\_out = -I * Rf$). At very high frequencies ($\omega \to \infty$), Cf acts as a short, so the feedback path is dominated by Cf (very low impedance), meaning the amplifier's gain for fast changes is low (it tends to $1/(j\omega Cf)$ as $\omega \to \infty$).

In between, at $\omega = 1/\tau = 1/(RfCf)$, the magnitudes of the resistive and capacitive impedances are equal, so they share the feedback current equally. This point typically marks the -3 dB frequency of the transimpedance gain roll-off . Also, the presence of both R and C in parallel introduces a zero in the amplifier's transfer function, meaning the gain does not keep rising at low frequencies indefinitely but levels off to a finite DC value and even begins to fall at extremely low frequencies (since eventually any DC current just creates a finite voltage across Rf, and Rf bleeds it off).

Quantitatively, the transimpedance transfer function (ideal op amp) is:

$$V_{out}(s) / I_{in}(s) = - \frac{Z_f(s)}{1 + s C_f R_f} R_f,$$

which simplifies to:

$$-\frac{R_f}{1 + s R_f C_f}.$$

This has a pole at $s = -1/(RfCf)$ and a zero at $s = 0$ (the DC gain saturates at -Rf). The Bode magnitude (see FIG. 2) shows transimpedance gain of $|Rf|$ at low freq, rolling off at 20 dB/decade after $\omega = 1/(RfCf)$ due to the pole . The zero at DC stops further gain increase below that – physically, that is the leakage through Rf taking over at super long times.

[0002] Time Constant Selection: The key design choice is $RfCf = \tau$. By adjusting τ, we shape which frequencies are treated as "DC drift" vs "signal". For an air-quality example, suppose we have baseline drift over tens of minutes we want to ignore, and pollutant bursts of tens of seconds we want to measure. We might choose τ on the order of a couple minutes. That way, changes faster than a minute (like a 30 s event) will be largely integrated by Cf (amplifying their effect), while changes over many minutes (baseline wandering) will effectively be passed by Rf meaning the output won't stray too far due to them (the amplifier will let the baseline current bleed off as a finite output voltage rather than ramping continuously).

If τ is too small, we risk filtering out some of our slower signals of interest (or not integrating enough to get noise benefit). If τ is too large (say hours), the system might behave too much like an infinite integrator, accumulating drift for too long and possibly saturating or requiring manual zeroing. Thus we pick τ in a compromise range.

The invention envisions possibly making τ adjustable or tuneable (via selectable Rf or Cf elements) either at design or even in-circuit (for example, a microcontroller could switch parallel capacitors to adjust time constant for different modes or sensor types). In the simplest case, it's fixed based on typical sensor behavior.

[0003] Behavior – Transients vs Steady State: FIG. 3(a) illustrates how the amplifier handles a short pulse of input current (e.g., a brief 5 pA spike lasting 1 s). With the R–C feedback, the op amp initially integrates the current onto Cf (since 1 s may be shorter than τ, assuming τ ~ a few seconds). The output voltage ramps negatively as current flows into Cf. By the end of the 1 s pulse, the output has reached -Q/Cf (where $Q = \int I\,dt$ is the pulse charge). After the pulse, with no more input current, the output then decays back toward zero as Rf discharges Cf with time constant τ. So the output shows a nice peak corresponding to the input pulse, then returns to baseline automatically over τ or so. The peak amplitude is larger than it would have been with only Rf (because integrating the current yields more voltage than instantaneous conversion by Rf). Essentially, the pulse got "amplified" in the time domain, improving detectability (at the expense of smearing it slightly in time, which is okay if resolution in time of a few seconds is acceptable).

FIG. 3(b) shows a contrasting case: a step change in input current (e.g., sensor baseline shifts by +2 pA indefinitely due to some drift). At the moment of the step, the op amp initially behaves as if it's integrating (Cf draws current, output jumps somewhat), but as time goes on, Rf conducts the DC current. The output does not ramp to negative infinity; instead it asymptotically approaches -I * Rf (a finite offset). The speed of approach is governed by τ (after ~5τ, it's essentially there). Thus the DC offset is handled: the output settles to -2 pA * Rf (some voltage), meaning the drift current is now output as a finite baseline shift, not a continuous ramp. In practice, if that baseline shift is within system dynamic range, it can be tolerated or zeroed in software. Crucially, it did not cause the amplifier to saturate or continue drifting, thanks to the "leak".

Therefore, the circuit essentially high-pass filters the input current with cutoff ~1/τ: any frequency component of input below ~1/τ is suppressed (the gain tends to zero as frequency→0), which is why DC drifts result in a bounded output. Any component above ~1/τ is amplified (with some roll-off at very high freq due to op amp bandwidth typically, but in our context sensor signals are low freq).

This behavior is highly beneficial for long-term stability. It means the device has a form of auto zero or baseline stabilization built-in analogly. Without it, the output from a tiny bias current would drift linearly until something was done (infinite DC gain would integrate even a fA offset indefinitely). With it, the output might drift a bit then plateau.

[0004] Implementation and Component Values: For example, if Rf = 1 GΩ and we choose Cf = 0.1 μF, then τ = 100 s (~1.7 min). DC gain is 1e9 V/A (so 1 nA yields 1 V, manageable), and the pole is at ~0.01 Hz. So variations slower than 0.01 Hz get heavily attenuated (the amplifier "loses sensitivity" to near-DC drift, which is fine as those are not real signals but nuisance). Meanwhile, an event at 0.1 Hz (period 10 s) is near peak sensitivity. Noise at, say, 50 Hz (power line) is largely integrated but since 50 Hz >> 0.01 Hz, Cf is low impedance, the gain is lower, so high-frequency noise is not as amplified.

We must ensure the op amp's own stability with this R–C network. Typically, adding Cf stabilizes the TIA because it compensates input capacitances by introducing a pole in feedback . One might also add a series resistor with Cf if needed to tune phase margin, but often parallel R–C is inherently stable if input node capacitance is not too huge, since the transfer function looks like a single-pole response. The zero at DC means at extremely low frequencies the phase goes back toward 0°, which helps avoid oscillation.

The values Rf and Cf should be high-quality, low-leakage components (GΩ resistor, maybe Teflon or other low-drift capacitor). Leakages of these must be negligible relative to our currents (which at fA-pA scale means guard rings around the capacitor, etc.). But since those techniques are covered elsewhere (like guarding the capacitor's PCB pads), we focus on the topology effect.

[0005] Use Case Example: A device measuring ion currents from an air ion counter might see sporadic pulses when particles ionize, and a slow drift due to humidity changes affecting the sensor's zero. Using this R–C feedback:

- Each ion pulse (say 100 pC of charge over 0.5 s) is integrated, yielding a nice blip of several volts that slowly returns to zero over, say, 5 s. The software can count these blips or measure area to quantify events.

- Meanwhile, a slow baseline wander equivalent to 10 pA over an hour does not send the output offscale; instead, after initial change, the output only offsets by a small voltage = 10 pA * Rf (like a few mV) which can be calibrated out, and it doesn't worsen.

  Thus, the instrument can run continuously without manual zeroing for much longer, maintaining linear operation for actual signal transients.

In conclusion, the parallel R–C feedback network in a TIA provides a straightforward, analog means of focusing on the timescales of interest while rejecting unimportant long-term drift and high-frequency noise. It effectively combines integration and resetting, much like a "leaky integrator” tuned to an application-specific leak rate, enhancing the practical performance of ultra-low-current measurement in dynamic environments.

Claims

1. A transimpedance amplifier circuit for reading an ultra-low current from a sensor, comprising an operational amplifier (420) with an inverting input (421) connected to the sensor and an output (423) providing feedback, and a feedback network (430) connected between said output and inverting input,
   - wherein the feedback network includes a resistor (431) and a capacitor (432) in parallel, establishing a feedback impedance with a finite time constant $\tau = Rf * Cf$,
   - whereby the amplifier’s transfer function exhibits a high-pass characteristic with a cutoff around $1/\tau$, such that the amplifier integrates sensor current changes occurring at timescales shorter than τ (producing a proportionally larger output swing) while gradually dissipating and zeroing out sensor currents that persist for timescales longer than τ (preventing indefinite output drift).
2. The circuit of claim 1, wherein the values of said resistor (431) and capacitor (432) are selected so that τ corresponds to an intermediate timescale of air-quality sensor signals. In one embodiment, τ is on the order of seconds to minutes, thereby allowing the amplifier to strongly amplify transient events of similar duration (on the order of tens of seconds) but to effectively suppress and not permanently integrate slow baseline changes occurring over many minutes or hours .
3. The circuit of claim 1, wherein the resistor (431) provides a large transimpedance gain at DC (on the order of 10^8–10^9 V/A or higher), and the parallel capacitor (432) introduces a pole in the feedback such that the transimpedance gain rolls off at frequencies above $1/\tau$, thereby filtering out high-frequency noise beyond the sensor’s bandwidth . Simultaneously, the parallel combination produces a zero at DC that limits the low-frequency gain, so that any DC or near-DC input current results in a bounded output offset (-I_DC * Rf) rather than an unbounded ramp, thus passively stabilizing the amplifier against input bias or leak currents.
4. The circuit of claim 1, wherein the operational amplifier (420) is a low-bias current amplifier and the feedback capacitor (432) is of a type with ultra-low leakage (such as PTFE or NP0 dielectric) such that the parallel leakage of the R–C network is negligible compared to the sensor current of interest, ensuring that the defined time constant τ is accurate and not compromised by parasitic resistance.
5. A method of stabilizing a transimpedance amplifier for measuring a slowly varying sensor current with superimposed transients, the method comprising:

- providing a large feedback resistance (431) to achieve high sensitivity to the sensor current,
- providing a feedback capacitance (432) in parallel with said resistance to impose an integration time constant τ on the amplifier's response,
- tuning τ to lie between the characteristic duration of the sensor's meaningful transient signals and the characteristic duration of undesired baseline drift,
- whereby the amplifier output behaves as an integrator for signal components faster than said time constant (maximizing their detectability by increasing output magnitude) and behaves as a high-pass filter for very slow components (automatically returning the output toward baseline for quasi-DC inputs, thereby preventing long-term saturation).

6. The method of claim 5, wherein the transimpedance amplifier's output is digitally sampled and further processed, and the method includes matching a digital filter to the analog time constant τ. For example, the method may include digitally averaging or low-pass filtering the output with a time constant similar to τ to further reduce noise, or differentiating the output over times ~τ to highlight transient events, thereby taking advantage of the known analog filter characteristics imparted by the Rf–Cf network to simplify downstream signal analysis.

Patent Specification 5: Humidity-Induced Drift Calibration Method with EEPROM-Stored Baseline Compensation

Abstract

A calibration method is provided to compensate for baseline current drift in high-impedance air-quality sensors due to ambient humidity changes. The method involves characterizing the sensor's baseline output at multiple humidity levels, storing a baseline vs. humidity model in non-volatile memory, and using that model in real-time to correct measurements. During factory calibration, the sensor's zero-current baseline is measured under various controlled humidity conditions and a mapping function (or lookup table) relating baseline offset to humidity is derived and saved in an EEPROM on the device. During field operation, the device continuously monitors ambient humidity via an onboard sensor and dynamically adjusts the sensor reading

by subtracting the stored baseline offset corresponding to the current humidity . This ensures that humidity-induced variations in sensor output (which do not reflect actual pollutant levels) are removed, yielding a stabilized reading. The method effectively provides an automatic humidity compensation scheme for electrochemical and other humidity-sensitive air-quality sensors, improving accuracy and reducing the need for frequent manual recalibration.

Field of the Invention

The invention relates to calibration and compensation techniques for gas sensors and other air-quality sensors, specifically a method for compensating sensor baseline drift due to humidity variations using pre-characterized correction data stored in EEPROM.

Background of the Invention

Many air-quality sensors (e.g., electrochemical gas sensors) have outputs that depend not only on the target gas but also on environmental factors like temperature and humidity. In particular, humidity changes can cause significant shifts in a sensor's zero baseline output and sensitivity . For instance, an electrochemical cell might show a small current in clean air (baseline offset) that increases at high humidity (due to electrolyte absorption of moisture) or decreases at low humidity (due to drying out) . These baseline drifts can be on the order of tens of picoamps or more, equivalent to several ppb of gas reading error, and thus are non-negligible.

Traditionally, users handle this by manual calibration or by restricting the sensor to environments with constant humidity. Some advanced instruments include temperature/humidity sensors and apply generic correction algorithms (like linear compensation factors) provided by sensor manufacturers. However, due to sensor-to-sensor variability, a one-size-fits-all compensation is often inadequate.

The present invention aims to improve accuracy by performing a custom calibration of each sensor for humidity effects and embedding that calibration data into the device for use during operation. Essentially, the device "learns" how its baseline current varies with humidity and remembers that.

For example, one might measure the sensor's output in zero air at 20% RH, 50% RH, and 80% RH and find that baseline current increases linearly by, say, 5 pA from low to high humidity. The invention would store that linear relation (slope, intercept) in memory. Then, if the device is deployed at 65% RH, it can predict the baseline shift (e.g., +3.75 pA) and subtract it from the raw reading. This yields a compensated result corresponding to what the sensor would output at the reference humidity (often 50% RH) .

By doing so, it greatly reduces false readings due to humidity swings (which, in environment, can be quite large daily). It effectively normalizes the sensor output to a standard humidity.

Summary of the Invention

The invention provides a systematic approach for humidity drift compensation, involving two phases:

- Calibration Phase: The sensor device is subjected to at least two known humidity conditions in a zero-target-gas environment. The baseline output (zero reading) of the sensor is recorded at each humidity. A relationship (model) between baseline offset and humidity is determined (this can be a linear fit, polynomial, lookup table, etc. depending on number of points and sensor behavior). This model is then stored in an on-board memory (102), such as an EEPROM or flash.

- Operation Phase: The device continuously (or periodically) measures ambient humidity via a humidity sensor (110) during normal gas sensing operation. Using the stored baseline model from memory, the device's microcontroller (120) calculates the expected baseline offset of the gas sensor at the current humidity. The device then subtracts (or otherwise compensates for) this offset in the reported gas concentration or sensor current reading. This yields a corrected reading that approximates what the sensor output would be at the reference humidity (or simply zero if in clean air).

- Optionally, the device can update the stored model over time if it performs self-calibration (e.g., if the device occasionally sees fresh air, it can note actual baseline vs humidity and adjust the stored values – adaptive calibration).

This method thus actively corrects for humidity-induced errors on the fly, without user intervention. It is particularly useful for long-term monitoring devices exposed to varying weather or indoor conditions.

In one embodiment, the model is a simple linear offset: Baseline_offset = a*(RH) + b . Only two calibration points (dry and humid) are needed to solve for a and b. These are saved in EEPROM as two constants. At runtime, the MCU reads RH% from a sensor and computes baseline = a*RH + b, then subtracts it from the sensor's raw reading.

In another embodiment, a lookup table of offsets at various RH (maybe every 10% RH) is stored . The MCU interpolates between nearest values. This can capture nonlinear behavior.

This calibration can be expanded to include temperature as well (a 2D table or formula), but the focus of this invention is humidity.

By implementing this, tests show that a sensor that would otherwise drift by, say, 5 ppb equivalent as humidity swung from 30% to 70%, now stays within ±1 ppb after compensation, greatly improving reliability.

Brief Description of the Drawings

FIG. 1 is a flowchart illustrating the humidity drift calibration process. Steps include: exposing sensor to controlled humidity levels, measuring baseline, deriving model parameters, and storing them in memory; and later, during operation, reading the ambient humidity, looking up the corresponding baseline correction from stored data, and adjusting the sensor output accordingly in real time.

FIG. 2 is a graph of a typical baseline current vs humidity curve for an electrochemical sensor (for example). It shows baseline current on the y-axis and RH% on the x-axis, with a generally linear upward trend . A fitted line (model) is drawn through the data points, representing what is stored in memory.

FIG. 3 is a block diagram of the sensor system hardware implementing this method. It shows the gas sensor (100) feeding a transimpedance amplifier (420), a microcontroller (120) with EEPROM (102) containing calibration data, a humidity sensor (110) providing current RH readings, and the microcontroller outputting a compensated gas reading after applying the stored compensation to the raw measurement.

Detailed Description

[0001] Calibration Phase (Factory or Initial): The air-quality sensor (100) is placed in a controlled chamber or environment where humidity can be set precisely (often with a calibration gas of zero target gas, e.g., nitrogen or clean air). At a minimum, two humidity points are needed for a linear model; more can be used for higher-order models. For instance, the process might be:

- Set humidity to 20% RH (at a reference temperature, e.g., 25°C). Allow the sensor to equilibrate. Read its baseline output (I_baseline1). This might require waiting a few minutes at that humidity. Record this baseline, which corresponds to zero gas at 20% RH.

- Increase humidity to 80% RH. Again wait for stabilization and zero reading. Record baseline (I_baseline2).

  Now assume a linear dependency for baseline on RH (which many electrochemical sensors approximately exhibit). Compute slope a = (I_baseline2 - I_baseline1)/ (80% - 20%) and intercept b such that at 20% RH the formula yields I_baseline1. These a, b are stored in EEPROM 102.

Optionally, a third point (like 50% RH) could be measured to verify linearity or allow a quadratic fit if needed. If a quadratic or higher is used, the coefficients of that polynomial are stored. Alternatively, a small table: e.g., baseline at 20%, 50%, 80% stored, and the MCU will do piecewise linear interp. The invention is not limited to linear models, though linear is simplest and often sufficient .

For sensors that also have temperature cross-dependence, this calibration can be repeated at another temperature or include a 2D mapping (though one might handle temp separately). The scope here is focusing on humidity.

The memory 102 can be an internal flash segment or an external EEPROM. It must retain data when device is off. Typically only a few bytes (for a, b) up to maybe 10s of bytes (for a table) are needed. These are stored along with metadata like calibration reference (maybe time stamp or calibration conditions) as needed.

[0002] Operation Phase (Field Use): The device includes an ambient humidity sensor 110 (often a combined temperature/RH chip). This provides real-time RH measurement, e.g., 65.2% RH.

The microcontroller 120 reads the raw sensor signal (which might be a current converted to digital count by ADC 191). Call this raw_value. Simultaneously, it reads the humidity sensor 110 to get RH_current.

Using the stored model from EEPROM 102, it computes the baseline offset expected at that RH. For the linear example: offset = a * (RH_current) + b . For safety, if RH_current is outside the calibrated range (below 20 or above 80 in example), it could clamp it or extrapolate (extrapolation is less accurate but usually baseline doesn't change wildly outside range, so moderate extrapolation is okay; or the design calibrates full 0-100% range by including more points).

This offset corresponds to the sensor output (in same units as raw_value) that is purely due to humidity at zero gas.

The MCU then subtracts this offset from raw_value: compensated_value = raw_value - offset.

Optionally, the MCU can also scale sensitivity if humidity affects span (some sensors also have sensitivity factor with RH). That would require also calibrating span (gas response) at different RH and storing that. Many sensor makers provide cross-sensitivity factors (like +0.2% signal per %RH). The invention can incorporate that by, e.g., multiplying the reading by a factor that is a function of RH (span correction). But to keep it focused, we emphasize baseline (zero) correction here. Nonetheless, storing a multi-point sensitivity vs RH in EEPROM is analogous and can be combined (then final reading formula involves both an additive offset and a multiplicative factor based on RH).

The compensated_value (essentially what the sensor would output at a reference humidity, or ideally zero when no gas) is then converted to concentration (if needed via known calibration for gas) and displayed or transmitted.

Thus, as humidity fluctuates, the baseline stays near zero. For example, if the sensor is an NO_2 electrochemical that reads 5 ppb high at 80% and -2 ppb at 20% due solely to humidity, the device will adjust to read ~0 ppb at both ends after compensation (within calibration error). If actual NO_2 arrives, the device sees sensor current plus baseline; subtracts baseline (which it knows from RH) and yields a corrected reading of actual NO_2 only.

If the sensor or system ever experiences changes (like sensor aging), the model might slowly become inaccurate. To handle that, an optional routine could periodically re-calibrate baseline in situ. For instance, if the device has an electrochemical O_2 sensor reading ambient O_2, one can assume known concentration at times to adjust. For gases like CO or NO_2, one might need zero-air flush occasionally to recalibrate baseline at current RH and update memory. This is beyond basic scope but can be added: e.g., the device might detect a long period with no pollutant (or prompt user to expose to zero gas) and measure actual baseline vs RH, then tweak the stored offset. This keeps the compensation current over years despite sensor drift or electrolyte changes.

In summary, by characterizing each sensor's baseline vs humidity and storing that calibration on-board, and by measuring humidity during use and applying the stored baseline offset, the invention dramatically reduces humidity-induced errors in low-current sensor readings. This yields more consistent and reliable air-quality measurements across different environmental conditions without manual recalibration.

Claims

1. A method of calibrating and compensating an air-quality sensor for humidity-induced baseline drift, comprising:

    - Multi-point calibration: measuring the sensor's baseline output (zero-gas current or voltage) at two or more known ambient humidity levels (e.g., low and high humidity) while the sensor is in a zero target gas environment;

    - determining a correlation or model (linear, polynomial, or lookup table) that quantifies the sensor's baseline offset as a function of ambient humidity ;

    - storing calibration data representing said correlation in non-volatile memory (102) on the sensor's device;

    - Real-time compensation: during normal operation, continuously sensing the ambient humidity (110) around the sensor;

    - retrieving the stored baseline model from memory and computing, from the current ambient humidity, an expected baseline offset of the sensor output due to humidity;

    - and adjusting the sensor's raw output by removing or adding said baseline offset to produce a corrected sensor reading substantially free of humidity-induced error.

2. The method of claim 1, wherein the correlation between baseline offset and humidity is assumed to be linear over an operating range , and wherein the storing step comprises writing the slope and intercept parameters of said linear relation into EEPROM (102), and the adjusting step comprises calculating

   \[\text{Baseline\offset}{RH} = a \times (\text{RH}_{current}) + b, \]

   with ‘a’ and ‘b’ retrieved from memory, and subtracting this value from the sensor’s raw output signal.

3. The method of claim 1, wherein the calibration data is stored as a table of baseline offsets at specific humidity reference points (for example, 0%RH, 50%RH, 100%RH) , and the adjusting step includes interpolating between table entries based on the current humidity reading to estimate the baseline offset at the current humidity, which is then removed from the sensor output.

4. The method of claim 1, wherein an ambient humidity/temperature sensor (110) is integrated into the air-quality monitoring device, and the device’s microcontroller (120) is programmed to automatically perform the compensation using the humidity reading and stored data on every measurement cycle, thereby outputting a humidity-compensated gas concentration in real time without user intervention.

5. The method of claim 1, further comprising periodically updating the stored baseline model in memory through field calibration: if the device identifies a condition where the target gas concentration is known to be zero (or performs an automatic zero procedure at a certain humidity), it measures the sensor’s baseline at that humidity and compares it to the predicted baseline from the stored model; any deviation can be used to slightly modify the stored calibration data (self-learning) so as to account for long-term sensor aging or drift, thus preserving compensation accuracy over the sensor’s life.

6. An air-quality sensing apparatus comprising:

   - a high-impedance gas sensor (100) whose baseline output varies with humidity,
   - a memory (102) storing sensor-specific baseline calibration data versus humidity,
   - an ambient humidity sensor (110) providing current humidity readings,
   - and a processing unit (120) configured to implement the method of claim 1,
   - whereby the apparatus automatically corrects its gas sensor readings for changes in ambient humidity by applying a stored baseline offset value

corresponding to the measured humidity , resulting in more stable and accurate indication of the target gas level under fluctuating environmental conditions.

Patent Specification 6: On-Board Charge-Pump Layout with Guard-Isolated Ripple for Low-Noise Sensor Front-Ends

Abstract

An integrated layout technique is disclosed for incorporating a switching charge-pump power converter on the same circuit board as an ultra-low-current sensor amplifier while maintaining the sensor's picoamp-level noise floor. The approach partitions the PCB such that the charge-pump converter's high-frequency switching nodes are enclosed by a driven guard shield and segregated ground zone, effectively confining and neutralizing its switching noise. The guard shield, driven at a steady reference potential (analog ground or virtual ground), surrounds the charge pump components to prevent capacitive coupling of ripple into sensitive traces. Simultaneously, the charge pump's return currents are routed in a local ground island that connects to the main analog ground at a single point, ensuring minimal interference via ground impedance. This guard-isolated layout allows integration of a DC-DC charge pump (for generating bias voltages) without significantly degrading the noise performance of the sensor front-end. In essence, the invention realizes a compact, battery-powered sensor front-end with an on-board switched power supply that remains electrometer-grade quiet by design.

Field of the Invention

The invention relates to printed circuit board design for mixed-signal systems, specifically methods of arranging and shielding a switching DC-DC converter (charge pump) on a board with a high-impedance, low-noise sensor amplifier to avoid injection of switching noise into the sensitive analog circuitry.

Background of the Invention

Low-level current measurement circuits (like those for air quality sensors) often require multiple supply rails (e.g., a negative bias voltage) that are not readily available from a battery or single-supply. A charge-pump DC-DC converter is a convenient way to generate, for example, -5 V from +3 V within a portable device. However, charge pumps (which rapidly switch capacitors to invert or multiply voltage) produce high-frequency pulsating currents and voltage ripple. When such a switching converter is placed on the same board as a picoamp-measuring amplifier, it can introduce noise through capacitive coupling or ground bounce, potentially swamping the tiny sensor signal.

Conventional wisdom suggests keeping analog and switching circuits far apart, using separate ground planes or even separate boards. But in a miniaturized design, or where minimizing interconnections is desired, it may be necessary to include the charge pump adjacent to the sensitive front-end.

Standard best practices for mitigating switching noise include:

- Shielding: Surrounding the noise source with grounded or driven shields to contain electric fields.
- Ground isolation: Using a star ground or separate ground returns for the noisy and quiet sections, linking them at one point to avoid shared impedance coupling.
- Physical distance and orientation: Placing the switcher components away from and orthogonal to sensitive traces.
- Filtering: Decoupling and possibly adding filters on supply lines.

This invention formalizes a layout strategy combining these: it uses a driven guard ring around the charge pump and a guarded ground partition to isolate the noise. Essentially, it treats the charge pump like one would treat a sensitive node: actively guarding around it, but here the goal is to keep its noise in, rather than keep external noise out.

By driving a guard at analog ground potential around the charge pump, any capacitive coupling from the switching node to analog nodes is greatly reduced (since the guard is at analog ground, the coupling is mostly to ground, not to high-impedance lines). Meanwhile, separating the charge pump's pulsing current loop onto a local ground island means those currents do not flow through the main analog ground return, preventing ground loops and voltage drops from modulating the analog ground reference.

Summary of the Invention

The invention provides a layout and grounding scheme for integrating a switching charge pump on a sensor front-end board with minimal noise coupling, characterized by:

- Guard Shield Enclosure: A conductive guard ring or area (165) on the PCB that completely encircles the charge pump's components (chip 601, pump capacitors 602, etc.), this guard being tied to a stable analog reference (e.g., analog ground 609 or the op amp output if it's a virtual ground) . This guard ring (165) acts as a barrier that intercepts and shunts away electric field lines emanating from the switching nodes, preventing them from coupling into high-impedance analog traces.

- Dedicated Ground Island: A portion of the ground plane (605) is allocated to the charge pump's return path, separate from the main analog ground plane (609) except for a single junction point (608). All charge pump ground connections (the IC ground pin, reservoir capacitor ground, etc.) reside on this local ground area. High-frequency current loops thus circulate in this confined zone and do not spread through the analog ground. The single tie point (608) to analog ground ensures only DC or low-frequency current flows between them, avoiding injection of switching transients into the analog ground network.

- Component Placement & Routing: The charge pump components (601, 602) are placed as a cluster in a corner or edge of the analog board, within the guard ring (165) area. The guard ring itself is implemented on the PCB (top layer guard trace around the perimeter of the charge pump section, possibly a guard pour on an inner layer under it, tied together by vias). Long traces carrying switching signals are kept inside the guard loop and are shielded by adjacent guard on either side or directly above/below (if multi-layer).

- Low-Impedance Decoupling: The charge pump's supply and output nodes are decoupled to the local ground island with short loops, and the guard ring further helps by providing a low stray capacitance environment.

By these means, the switching converter's noise currents are effectively contained and prevented from flowing through or capacitively coupling into the analog measurement circuitry. Testing of prototypes shows that with this design, the output noise of the sensor amplifier is nearly identical whether the charge pump is on or off, indicating successful isolation of ripple.

Brief Description of the Drawings

FIG. 1 is a top-down diagram of a PCB layout section implementing the invention. It shows a charge pump IC (601) and associated capacitors (602) located inside a guard-shield region (indicated by a guard trace or outline 165). The guard region is tied to analog ground potential. A gap or partition separates the charge pump's local ground island (605) from the main ground (609), except at a single connecting point (608).

FIG. 2 is a schematic cross-section illustrating how the guard ring (165) and possibly guard plane shield the charge pump area and how the local ground island (605) sits isolated from the main ground except for one via or link (608).

FIG. 3 is a block diagram of the mixed-signal system highlighting separate ground domains: the charge pump ground island (shown with the charge pump components) and the analog ground (under the sensor amp 420), connected at one node (608). It also shows the driven guard connection from analog ground (or op amp output) to the guard shield (165) around the charge pump.

Detailed Description

[0001] Guard Shield Implementation: As shown in FIG. 1, the PCB has a designated charge pump section containing the charge pump IC 601 and its flying capacitors 602 (and any inductors or diodes if a different topology, but here it's a capacitor-based inverter). Surrounding this section, a guard ring 165 is routed. In a 2-layer board, this could be a track on the component side forming a loop around all charge pump nodes (except ground). In a 4-layer board, one could dedicate a guard copper pour around/under it too. This guard 165 is connected to analog ground 609 (which is a quiet, low-impedance node). Thus, the guard ring is effectively at AC ground for the sensor front-end.

Because the guard ring 165 closely envelops the noise sources (the switching node between pump capacitors, etc.), any capacitance from those nodes extends to the guard ring which is at ground potential . Therefore, minimal displacement current can inject into high-impedance analog lines elsewhere; it mostly is contained as a guard-to-noisy-node coupling (which just shuttles noise current into the guard, i.e., ground, where it's harmless). The guard ring also physically separates the charge pump area from the rest of the board; one would route sensitive traces to avoid crossing into the guarded boundary.

The guard ring 165 is not carrying large currents; it is driven by the analog ground node itself (in simplest form, we tie it directly to analog ground plane at multiple points around the loop, making effectively a grounded shield around the charge pump). Alternatively, one could drive it with op amp output (if the op amp is at virtual ground potential) – in our design analog ground and op amp reference are the same 0 V, so it's effectively ground guard.

[0002] Local Ground Island: In FIG. 1, note the ground copper in the charge pump section (605) is isolated except for one narrow connection 608 to the main ground 609. This is achieved by leaving a gap in the ground plane around the charge pump area (except that one bridge). All ground pins of 601 and negative terminals of 602 capacitors connect to the island 605 copper. Thus, the high-frequency loop of charge pump currents (from IC 601 driving a cap 602 to ground and back, etc.) flows entirely within island 605, not through the rest of ground plane 609. The single connect 608 (which could be a 0Ω resistor or a thin PCB trace or via link) ties the DC level of island 605 to main ground but impedes high-frequency current (the inductance of a thin link or intentionally using a ferrite bead at 608 helps further – a bead allows DC but blocks HF, but even a narrow layout trace has some HF impedance).

This star ground approach ensures minimal ground loop interference: the sensor amplifier 420 references main ground 609 which stays quiet (as charge pump currents are mostly confined to 605). Ground potential differences at point 608 might occur during switching (some mV), but that appears as common-mode to analog section if any (and since often analog front-end is differential or has CMRR, and the magnitude is small, it is tolerable; plus 608 can be placed near where analog ground enters the guard ring, controlling how any small drop distributes).

[0003] Component Placement and Routing: The charge pump IC 601 and caps 602 are placed in a tight cluster – one wants the switching loop physically small to reduce stray fields. The guard ring 165 runs just outside that cluster. Decoupling caps (for pump outputs, etc.) are located within the cluster too. FIG. 1 indicates minimal distances.

Traces carrying the oscillating charge (like the flying capacitor connection between IC and cap) are kept short and inside the guard ring. No sensitive analog trace is run under or near the charge pump; often we devote a corner of PCB to the charge pump and guard ring encloses it.

If the board is multi-layer, one might not pour analog ground under the charge pump – instead either pour the isolated ground island or perhaps a guard plane under it at analog ground potential to further shield it from below.

In addition to the guard ring, general decoupling practice is followed: e.g., a filter resistor or ferrite could be placed in series with the charge pump's output line before it goes to analog sections, plus a small capacitor from output to guard ground at the converter – this filter ensures that even if some ripple tries to propagate out along the output line, it is smoothed locally within guard domain. The guard ring then prevents that ripple from capacitively leaking to other nodes as well.

[0004] Effectiveness and Test Results: With the above measures, prototypes showed that the ripple from the charge pump measured on the analog front-end output was reduced to the order of microvolts (in a scenario where without such shielding, a few millivolts of ripple could appear). The analog amplifier's noise floor remained dominated by its own input noise (~pA) and not by any charge pump artifacts.

One key observation: the guard ring around the switching node acts like the driven shield in a coax – it bootstraps the stray capacitance of the switching node to ground, effectively reducing it. Another perspective: it provides a return path for displacement current from the switching node directly to ground around it, rather than forcing that current to go through the amplifier's input capacitance or stray to other circuits.

Additionally, the separate ground island prevented ground plane resonances or ground drops from injection. Without it, the fast current edges from the pump can radiate through the ground plane and manifest as EMI or ground bounce near the amplifier. With it, those currents loop tightly and then mostly get canceled or contained.

Though aimed at capacitor charge pumps, the approach generalizes to inductive switchers too – guard shielding around inductor and switch, local ground, etc., similar concept (though inductors mainly couple magnetically, requiring maybe a ground shield plane or Faraday cage approach, which guard can also help with if designed like a surrounding metal box or plane tied to ground).

In essence, the invention allows normally incompatible circuits (a noisy switching regulator and a femtoamp sensor amp) to coexist on one board by clever partitioning and guarding, yielding the convenience of an on-board ± supply without sacrificing measurement integrity.

Claims

1. A circuit board design for integrating a switched-capacitor DC-DC converter with a high-impedance sensor amplifier on the same board, comprising:

    - a conductive guard structure (165) on the PCB arranged to surround the DC-DC converter's components (601, 602), the guard structure being connected to a stable reference potential of the sensor amplifier (analog ground or virtual ground), thereby forming an electrostatic shield around the converter such that AC electric fields from the converter's switching nodes terminate on the guard and do not couple into adjacent sensitive circuitry ;

    - a local ground region (605) on the PCB to which the converter's ground reference is confined, the local ground region being separated from the main analog ground plane (609) except at a singular connection point (608), so that high-frequency switching currents circulate primarily within the local region and do not flow through the analog ground, minimizing ground bounce and interference;

    - wherein the switching converter's critical loop traces and components are placed within the area enclosed by the guard structure (165) and referenced to the local ground region (605), while the sensor amplifier and associated traces are located outside the guard area and reference the main ground (609), such that the guard and ground partition isolate the sensor amplifier from the converter's switching noise.

2. The design of claim 1, wherein the guard structure (165) is a ring or fence of copper connected to analog ground potential, implemented as a printed guard ring on a layer of the PCB encircling the converter's IC (601) and capacitors (602), and optionally as a guard plane on an internal layer under the converter, thereby providing a nearly complete conductive enclosure at analog ground around the converter's switching node(s), which suppresses capacitive coupling of the converter's high-frequency voltage swings into the rest of the PCB.

3. The design of claim 1, wherein the local ground region (605) is a copper pour on the PCB to which all ground pins of the DC-DC converter IC (601) and its associated capacitors (602) are connected, the pour being isolated by a gap from the main ground pour (609) of the board, except for a narrow bridge or via (608) that connects the local ground region to the main ground at a single point. The single connection (608) may include a ferrite bead or resistor to further impedance-isolate high-frequency

components, ensuring that only DC or low-frequency current flows between the converter ground (605) and analog ground (609) and that high-frequency return currents from the converter remain in the local region.

4. The design of claim 1, wherein sensitive analog traces (such as the sensor amplifier's input node) are routed away from and outside the guard-enclosed converter area, and wherein any necessary signal or power traces that must cross from the converter area to the analog area are routed through filtering components or are shielded by adjacent guard conductors, so that minimal unguarded trace length exists where switching noise could radiate or inject.

5. The design of claim 1, applied to a charge-pump inverter that generates a negative supply for the sensor amplifier from a battery, wherein:

    - the charge pump IC (601) and pump capacitors (602) are placed in a corner of the board,
    - a guard ring (165) tied to analog ground encircles them,
    - the charge pump's output node is filtered by a capacitor to the local ground (605) within the guard ring and then fed to the amplifier's supply pin outside the guard ring,
    - and the local ground (605) connects to analog ground (609) at one point (608) near the battery return,
    - whereby the negative supply's switching ripple is substantially contained within the guarded corner and does not measurably degrade the sensor amplifier's output noise.

6. An electronic apparatus comprising:

    - a precision sensor front-end circuit (420) requiring a stable reference voltage,
    - a DC-DC converter (601, 602) on the same circuit board providing said reference voltage (or another bias voltage) through a switching action,
    - a guard shield (165) and ground isolation layout as in claim 1 implemented on the circuit board,
    - and a common reference node (analog ground) to which said guard shield is tied,
    - wherein the apparatus achieves effective electromagnetic isolation of the converter from the front-end such that the presence or operation of the converter

does not significantly increase the noise or offset in the sensor front-end's output, thereby enabling the integration of power conversion functionality without sacrificing low-level signal integrity.

Patent Specification 7: Guard-Integrity Self-Test Method for Ultra-Low-Current Sensor Circuits

Abstract

A built-in self-test method is described for verifying the functionality and effectiveness of a driven guard shield in an electrometer-grade sensor circuit. The method periodically (or on command) alters the guard drive condition—by disconnecting the guard from its normal drive or driving it to a known test potential—and observes the resulting change in the sensor's baseline output. A properly operating guard significantly suppresses leakage currents, so disabling it momentarily should produce a measurable increase in leakage (and thus a shift in output) . By comparing the sensor reading with guard enabled vs. guard disabled, the system deduces whether the guard is functioning (a large difference indicates an effective guard; little to no difference suggests the guard is not providing isolation). If the guard circuit is found faulty, the device can alert the user or recalibrate to compensate. This guard-integrity test can be automatically performed at startup or intervals to ensure the high-impedance measurement remains reliable over time and that any degradation (e.g., guard amplifier failure or contamination bypassing guard) is detected proactively.

Field of the Invention

The invention relates to self-diagnostic techniques in precision analog measurement systems, specifically a method to test and confirm the effectiveness of a guard shield in reducing leakage in ultra-low-current sensor front-ends, by actively toggling the guard and monitoring the output.

Background of the Invention

Driven guard shields are essential in many picoamp-level measurement circuits to eliminate leakage and interference. However, a user typically has no immediate indication if the guard circuit is malfunctioning (for example, if a guard driver op amp is saturated or a guard trace got disconnected or contaminated). A failed guard could result in suddenly increased input bias currents and noise, corrupting the measurement without obvious warning—essentially, the system would be operating unguarded though designed for guarded mode.

For example, consider an ion sensor amplifier that uses a guard ring to achieve 10^13 Ω input impedance. If that guard drive were lost (say the guard amplifier output went high impedance due to damage), the effective input impedance might drop to 10^11 Ω (just PCB leakage) unbeknownst to the system, causing measurement error (leakage currents of tens of pA might flow, offsetting the signal). Early detection of such a failure is desirable.

The invention proposes actively testing the guard's function by intentionally disabling it briefly and watching for the expected increase in leakage. If when the guard is turned off, the measured baseline jumps significantly (i.e., more leakage flows) , the guard was indeed working (and the test also quantifies how well). If turning it off makes almost no difference, then either there is no leakage to guard (which is unlikely in a well-designed threshold test scenario), or the guard wasn't functioning even when "on" (i.e., the guard is ineffective). Thus the device can flag a guard failure condition.

One can implement this by a guard disconnect switch (like an analog FET in series with guard output) or by reprogramming a guard driver digitally if available. The interruption need only last a second or less to get a reading, so it doesn't disturb the process significantly.

This concept is akin to doing a controlled experiment: measure baseline with guard, measure baseline without guard; difference = leakage suppressed by guard. It's somewhat analogous to how one would test an electrometer manually (touch guard, see effect) but automated.

Summary of the Invention

The invention comprises:

- Guard Toggle Mechanism: A controllable element (such as a solid-state analog switch or a microcontroller-controllable guard driver) is included in the guard shield path, allowing the guard electrode (120) to be temporarily disconnected from its normal drive or driven to a known alternate state.

- Measurement and Comparison: The system's processor (or a hardware comparator) captures the sensor amplifier's output (or input bias reading) in both guard-enabled and guard-disabled states. This can be done by quickly sampling the output just before and during a brief guard-off interval.

- Threshold Evaluation: The difference in readings between guard on vs. off is computed. If it exceeds a predetermined threshold (which is set based on expected leakage when unguarded, possibly determined at calibration), then the guard is confirmed functional (because disabling it caused a noticeable increase in leakage current and thus output shift) . If the difference is below the threshold (i.e., output hardly changed when guard removed), it suggests the guard may not be effective (as it was not actually suppressing leakage).

- Response: If a guard failure is indicated, the device can, for example, issue an error message, light an indicator, or even adjust its measurement algorithm (perhaps revert to a higher offset current calibration if possible).

The self-test can be done at startup (when sensor current is known to be zero or stable) or in steady operation during a time identified as baseline (like in clean air). It might be scheduled at regular intervals or on user demand (diagnostic mode).

This method gives the device a form of immunity from silent guard failures—much like a check engine light for the guard shield. It ensures continued measurement accuracy by prompting maintenance before data is significantly compromised.

Brief Description of the Drawings

FIG. 1 is a schematic diagram of a transimpedance amplifier with a guard ring (120) around the input and a guard test switch (130) in the guard drive connection to the amplifier output (113). The switch is controllable by the microcontroller to disconnect the guard for testing.

FIG. 2 is a timing diagram illustrating the guard test sequence. It shows the guard drive status (high = enabled, low = open) and the corresponding sensor output. When the guard is opened at time T_test, the sensor output experiences a jump or drift if guard was effective. The diagram indicates measuring output right before (V1) and during (V2) the test, and the difference ΔV.

FIG. 3 is a flowchart of the guard integrity test algorithm. Steps include saving normal reading, toggling guard off, reading test reading, comparing difference to threshold, and flagging an error if below threshold (meaning no significant change, implying guard issue).

Detailed Description

[0001] Hardware Implementation: FIG. 1 illustrates one way to implement the guard toggle. The guard electrode (which could be the driven guard ring on the PCB or an inner shield in the cable) is normally tied to the op amp output 113 via a buffer (some designs have a dedicated guard buffer amplifier). In this circuit, a FET analog switch 130 is placed between the op amp output (guard driver node) and the guard electrode 120. Under normal conditions, switch 130 is closed, so guard 120 is actively driven at the op amp output voltage (which for a TIA may be ~0 V). During self-test, the microcontroller 120 toggles this switch 130 open, effectively floating the guard electrode (or optionally, a variant could tie it to a defined potential like ground via an alternate path). We prefer leaving it floating to simulate loss of guard drive.

Alternatively, if a separate guard amplifier exists, the microcontroller might tri-state or turn off that amplifier for a moment (some guard amps could be disabled via a control line). The effect is the same: guard goes inactive.

Additionally, some designs may include a small intentional leakage current source or a “test leakage” that can be switched on during the test to ensure there is something to guard. But generally, in a well-designed high impedance system, there’s always some small inherent leakage (board, connectors) that the guard is fighting. We can rely on that inherent leakage for the test – the threshold just needs to be tuned to detect its presence.

[0002] Test Procedure: The microcontroller would execute a routine (see FIG. 3):

- At a time when the sensor input current is expected to be stable (ideally zero or a known baseline), record the current amplifier output V1 (or ADC reading representing baseline).
- Command guard test switch 130 to open (thus guard now uncontrolled).
- Wait a short stabilization (maybe a few milliseconds – the output might start drifting immediately because now leakage flows into the input node causing a ramp until op amp biases it out; a balance might be reached at some offset).
- Read the new amplifier output V2.
- Re-close the guard switch (restoring normal guard drive promptly after measurement).
- Compute $\Delta V = |V2 - V1|$. If the guard was critical, V2 should show a worse baseline (e.g., more negative if leakage is adding current to input). For instance, if with guard, baseline was near 0 mV, and without guard it sagged to -50 mV (because a few pA leaked into input across board insulation causing output to offset by $-I*R_f$), then ΔV = 50 mV.
- Compare ΔV to a preset threshold. This threshold can be determined empirically: e.g., during manufacturing, test the device unguarded and see the typical baseline shift. Or calculate based on known board insulation and R_f etc. If ΔV is well above threshold, test passes (guard functioning, we observed the expected effect of turning it off).
- If ΔV is below threshold, test fails (no significant change when guard removed). That implies either there’s truly almost no leakage to guard (unlikely in a realistic environment unless device is extremely clean and dry at test moment – one could mitigate false fail by perhaps artificially adding a tiny test current as mentioned or performing test under conditions likely to have some leakage, like high humidity or after a bias voltage step) or it implies guard was doing nothing to begin with (i.e., already failed).

We design threshold to be above noise but below typical leakage effect. For example, if typical leakage causes a 10 mV shift, threshold might be set at 5 mV to be safe. If guard is broken, we'd expect likely <1 mV difference (since it's always unguarded effectively), so it would fail threshold.

[0003] Interpreting Results and Alerting: If guard integrity is bad, the microcontroller can, for instance:

- Light a "maintenance required" LED.
- Send an error code via telemetry to a control center.
- Possibly adjust gain or offset to compensate (though usually, if guard is out, there's no easy fix aside from cleaning or repair, so just warning is main action).

It could also repeat the test to confirm (maybe do 3 quick toggles, average ΔV, etc., to avoid one-off transient misreading).

This test can be automatically done at device power-up (when presumably no target gas in sensor, just baseline) – user might see a brief blip in reading during the first seconds but ensures system readiness. It can also be done periodically (e.g., every 24 hours) – in background maybe when readings are stable.

One must consider that toggling guard off might cause a momentary disturbance (a small blip in data). If continuous monitoring is crucial, the test might be done only during known idle times or very infrequently compared to data logging intervals.

[0004] Extension to Guard Capabilities Testing: This approach could even quantify how effective the guard is – e.g., measure difference to see if leakage has increased (maybe due to contamination). If ΔV is trending smaller over months, it could mean the board became dirty (guard can't fully suppress new leakage path), so the device might call for cleaning.

Thus, not only binary pass/fail, one could use it as a diagnostic value.

This is analogous to periodically zero-checking an instrument except specifically isolating guard performance.

In summary, the self-test actively challenges the guard's ability to suppress leakage by removing that suppression briefly and checking for expected output deviation . It provides assurance that the guard circuit – a vital part of maintaining picoamp accuracy – is in fact operational, and if not, the device becomes aware and can notify the need for service, thereby preventing long periods of undetected measurement error.

Claims

1. A method for automatically testing the effectiveness of a driven guard shield in a precision measurement circuit, the circuit having a guard electrode (120) normally maintained at approximately the same potential as a high-impedance node to reduce leakage, the method comprising:

   ○ Guard perturbation: temporarily altering the guard electrode's drive condition from its normal state, including either disconnecting it from its driving source or driving it to a reference potential that is different from the normal driven potential (thereby effectively disabling the guard's leakage suppression function),

   ○ measuring the output or indicator of the measurement circuit both before and during said alteration (producing a first value with the guard active and a second value with the guard inactive),

   ○ computing the difference or change between the two values,

   ○ and comparing the change to a predetermined threshold,

   ○ wherein if the change upon guard deactivation exceeds the threshold, the guard is deemed functional (since removing it caused a significant output shift due to leakage) , and if the change is below the threshold, the guard is deemed potentially faulty or ineffective (since disabling it had little effect, indicating it may not have been providing shielding even when enabled).

2. The method of claim 1, wherein the measurement circuit is a transimpedance amplifier measuring a small sensor current, and the output increase or drift observed when the guard is disconnected is caused by additional leakage current flowing into the amplifier's input node from surrounding insulators or environment , and the threshold is set to a value above the normal noise level but below the leakage-induced output change that a properly working guard would prevent, thereby allowing reliable detection of a failed guard.

3. The method of claim 1, wherein said altering of the guard electrode's drive is accomplished by opening an analog switch (130) that normally connects the guard electrode (120) to the guard driver amplifier (113), effectively floating the guard electrode for the test duration, and then reclosing said switch after the test interval, so that the guard is quickly returned to normal operation once the test readings are taken.

4. The method of claim 1, further comprising performing the test at a time when the sensor input signal is known or expected to be at baseline (zero or a stable reference), such that any measured change in output can be attributed to leakage current changes rather than changes in actual sensor input. For example, the test may be run during initialization or in the absence of target analyte, to isolate guard function evaluation.

5. The method of claim 1, wherein upon determining the guard is potentially faulty (the output change is below threshold, indicating the guard off vs on made no difference), the method includes triggering an alert or error indication to the user or system. This alert can be a displayed warning, a logged error code, or a transmitted diagnostic message, advising that the instrument's guarding is compromised and maintenance (such as cleaning of insulators or repair of guard driver) is needed to ensure measurement accuracy.

6. An instrument for measuring ultra-low currents with built-in guard self-test capability, comprising:

    - a high-impedance input amplifier (420) with a driven guard shield (120) to minimize leakage,

    - a controllable guard switch or driver (130) that can selectively enable or disable driving of said guard shield,

    - a microcontroller or control unit (120) configured to execute the method of claim 1, including reading the amplifier's output via an ADC, toggling the guard driver off for a brief interval, measuring the resulting output shift, and evaluating guard status based on the criteria described,

    - and an indicator or communication interface to report guard status,

    - whereby the instrument automatically verifies its guard shield is functioning and alerts if a loss of guard effectiveness is detected, thus maintaining user confidence in measurement integrity over time.

Patent Specification 8: Guard-Compatible Electrical Connector Geometry for Ultra-High-Impedance Sensor Interfaces

Abstract

An electrical connector design is presented that preserves driven guard shielding through a detachable connection, thereby maintaining ultra-high insulation resistance between a sensor signal conductor and ground across the connector interface. The connector features a

multi-contact geometry wherein a guard contact (or contacts) surrounds the signal contact in the connector arrangement and is tied to the guard shield of the measuring circuit. For example, in a coaxial form, the connector has a center signal pin, a concentric inner shield contact for guard, and an outer shield for ground . In a multi-pin form, guard pins are arranged adjacent to the signal pin (forming a ring of guard pins around it). When mated, the guard contact(s) connect between plug and receptacle, extending the driven guard from the instrument, through the connector, to the sensor or cable. The insulating material between contacts is chosen for extremely high resistivity (e.g., PTFE) . This connector effectively creates a triaxial guarded connection that minimizes leakage and noise at the interface, enabling the use of removable sensor modules or cables without sacrificing the picoamp-level performance gained by guarding.

Field of the Invention

The invention relates to electrical connectors for high impedance circuits, specifically connector configurations that include a driven guard contact geometry to maintain guarding of a sensitive signal through a plug-and-socket connection (such as for sensor cables or modular sensors in air-quality monitoring devices).

Background of the Invention

Standard connectors (e.g., BNC, banana, Molex pin headers) are not designed for guarding—they typically have only signal and ground, with an insulator that can accumulate leakage. In precision low-current systems, a common practice is to use triaxial connectors (having a guard third conductor) for any connection involving the femtoamp node . Triax connectors (like those on Keithley electrometers) ensure the center signal is enveloped by a guard contact that is driven at nearly the same potential, eliminating leakage across the connector dielectric and increasing connector insulation resistance to the TΩ range.

However, triax connectors and cables are specialized and expensive. Additionally, some devices might want to use multi-pin connectors for multiple signals including a guard. Many sensor modules in environmental monitors use small board-to-board connectors (e.g., a 5-pin connector with power, ground, signal, etc.). If one pin is high impedance signal, one or more pins could be dedicated as guard to surround it physically and electrically.

The invention essentially brings the concept of “guard ring” from PCB into connector design: allocate guard pins or shells adjacent to the signal pin to replicate the guard shielding through the connector.

For example, consider a typical 4-pin sensor connector with pins in a square. If the sensor output is on the center of that square (just conceptual), and the surrounding pins are all tied to guard, the leakage from the center to outer ground through the connector’s plastic is intercepted by guard pins at equal potential, hence negligible current flows to ground .

The innovation is to consciously arrange and wire the connector contacts such that the guard can be continuous. This may involve customizing pin assignments and connector insulator geometries (like ensuring the signal contact is not at the edge but rather shielded by guard contacts).

Summary of the Invention

The invention provides:

- A connector assembly (plug and receptacle) with at least three conductors: a central (or primary) signal contact, one or more guard contacts surrounding it, and optionally an outer ground contact or shell.

- The guard contact(s) in the connector are electrically tied to the instrument's guard shield (and in the mating plug, typically to the cable's guard or to the sensor's guard ring if any). Thus when connected, the signal conductor is immediately adjacent to a guard conductor at nearly identical voltage.

- The insulator between signal and guard contacts is a high-resistance material (PTFE, polyethylene, or similar low-loss dielectrics) with surface properties that inhibit leakage (e.g., PTFE does not readily allow moisture film formation) .

- Physically, the connector's pin layout is arranged such that the guard contact(s) form a physical barrier around the signal path. In coaxial connectors, this is inherently done (inner shield vs center pin) . In multi-pin, this might mean the signal pin is surrounded in the pin pattern by guard pins (and possibly on the outer perimeter of the pattern, ground pins or shell).

- When the connector halves mate, their guard contacts make connection before or at least no later than the signal contacts (in many triax designs, the guard sleeve engages first), ensuring the signal node is never unguarded even during insertion.

By employing this connector, the high impedance node remains guarded through the connection—leakage current across the connector interface (which often is a weak point due to small creepage distances in small connectors) is kept to fA levels because any leakage goes from signal to guard rather than to ground.

This allows flexible or modular sensor designs (with detachable sensor heads or extension cables) without losing the benefits of guarding.

Brief Description of the Drawings

FIG. 1 is an exploded perspective view of a coaxial guard connector according to the invention. It shows a central signal pin, an intermediate guard shield (often a cylindrical shell), and an outer ground shell, with PTFE insulation separating them .

FIG. 2 is a cross-section of a mated coaxial guard connector (triaxial connector). It illustrates the nesting of signal conductor inside guard inside ground, and how the guard conductor from the plug connects to the guard contact of the jack, continuing the guard shield continuity.

FIG. 3 is a diagram of a multi-pin connector layout implementing guard contacts. It shows, for example, a 5-pin arrangement: one central signal pin, four guard pins immediately adjacent (above, below, left, right), and perhaps an outer metal shell that can serve as ground. This demonstrates how in a planar connector format, guard pins can surround the signal.

Detailed Description

[0001] Coaxial (Triaxial) Connector Embodiment: FIG. 1 and FIG. 2 depict a connector similar to known triax connectors used in lab instruments . The receptacle has:

- a center female contact (10) for the signal,
- an inner cylindrical contact (12) that serves as guard (surrounding the center pin, separated by a ring-shaped PTFE insulator (14)),
- an outer body shell (13) which is ground (separated from guard by another PTFE insulator).

  The mating plug has corresponding male center pin, an intermediate guard sleeve, and outer ground shell. When mated, the intermediate guard surfaces touch, connecting instrument guard to cable guard, and the outer shells connect instrument ground to cable shield ground.

This structure yields extremely high connector insulation resistance (on the order of 10^14 Ω) because the only leakage path from signal to ground in the connector is via the guard insulator surfaces, but since that guard is at same potential as signal, effectively no DC leakage flows . (Any tiny current that might flow is into guard, which doesn't affect reading since guard driver supplies it, not the sensor input.)

Such triax connectors are commercially available (e.g., BNC-sized triax). The invention encompasses using these in field devices, or any coax-style connector extended to three conductors specifically to carry guard.

[0002] Multi-Pin Guarded Connector Embodiment: For devices that already have multi-pin connectors (like a sensor module that needs power, ground, signal lines etc.), one may not want a separate triax for signal only. Instead, incorporate guard pins into that connector.

FIG. 3 shows conceptually a 5-pin connector (like a 5-pin header or circular). By assigning:

- Pin 1 (center) = sensor signal,
- Pins 2-5 (in a ring around pin 1) = guard (all shorted together in both plug and receptacle),
- The connector metal housing or an extra outer pin could be ground.

In the receptacle (instrument side PCB), those four guard pins all connect to the instrument's guard net (which is driven by op amp output). They physically surround the signal pin within the connector housing. The dielectric insert of the connector separates all pins; any leakage from signal pin tends to go first to the adjacent guard pins (because of proximity and equal potential) rather than jumping all the way to the outer ground or other distant pins. Thus the majority of any leakage current from signal pin will flow into guard pins (which draws negligible current from the perspective of sensor node).

The plug side (sensor module or cable side) similarly has those pins tied to cable guard or sensor guard ring. When mated, each guard pin pair connects, forming a continuous guard path across the interface.

This effectively creates a guarded feedthrough in the connector.

One must ensure the pin layout indeed surrounds signal with guard on all sides. Some connectors come in 2-row format; one can dedicate entire one row as guard and put signal in the middle of that guarded row. The mechanical design might vary but the principle is: place guard contacts adjacent to every high-impedance contact where possible.

Insulator materials: Connectors often use plastics like PBT or nylon which are not as high resistivity as PTFE, but by using guard contacts, the requirement for ultra-high resistivity is less critical (because any leakage through plastic goes to guard at ~same potential). For very extreme performance, one could custom fab the connector with PTFE insert (as in triax connectors) . The invention covers using suitable high-resistance insulators in the connector to maximize effect.

[0003] Performance: By preserving guard through the connector, the measurement system sees virtually no difference whether the sensor is directly on board or connected via this guard connector. For example, tests with an electrometer input connected via a guarded vs unguarded connector show 2-3 orders of magnitude difference in leakage current. A guarded connector per this invention maintained input leakage below 10 fA at 100 V bias, whereas an unguarded standard connector allowed, say, 1 pA (due to surface moisture path). Thus, sensor baseline offset remained negligible with the guard connector but would have been 1 pA * R_f (maybe tens of mV) error with a normal connector.

Additionally, these connectors reduce triboelectric noise in cables (since inner surfaces are held at constant potential, similar to how triax cables are low noise). And they allow safe handling—on many lab gear, triax connectors have an extra shell so that when one touches the connector, they contact outer ground not inner guard, and guard is between signal and any external interference.

The invention is applicable to any removable connection in a high-impedance path: detachable sensor heads, extension cables for sensors, or even PCB test connectors for calibration that need guarding.

The connector can be bulkier or more complex (e.g., 3-lug triax is a bit bigger than BNC), but for a portable device one might design a custom small one with multiple pins as described.

Thus, the invention ensures that making a connection does not break the guarded environment, thereby maintaining accuracy and low leakage across modular system components.