Review

Not peer-reviewed version

On Recent Advances in Design of Transimpedance Amplifiers in CMOS: A Taxonomy of Topological Enhancements Beyond the Transimpedance Limit

Agata Romanova * and Vaidotas Barzdenas

Posted Date: 10 June 2026

doi: 10.20944/preprints202606.0753.v1

Keywords: transimpedance amplifier (TIA); C MOS a nalog c ircuits; t opology t axonomy; transimpedance limit; assumption relaxation; bandwidth enhancement; low-noise design; optical receivers; inductive peaking; regulated cascode

 doi:10.20944/preprints202606.0753.v1

Review

On Recent Advances in Design of Transimpedance Amplifiers in CMOS: A Taxonomy of Topological Enhancements Beyond the Transimpedance Limit

Agata Romanova * and Vaidotas Barzdenas

Department of Computer Science and Communications Technologies, Vilnius Tech, Vilnius, Lithuania; Sauletekio al. 11, Vilnius 10223, Lithuania

* Correspondence: agata.romanova@vilniustech.lt

Abstract

Transimpedance amplifiers (TIAs) are the critical current-to-voltage interface in optical receivers, LiDAR front-ends, biomedical sensors, and unconventional applications such as magnetic-resonance receiver-coil arrays and wide-bandgap ultraviolet detectors, and their CMOS design is governed by a fundamental gain-bandwidth-noise trade-off whose structure is rarely made explicit. This review introduces a unifying framework rooted in three explicit assumptions underlying the classical shunt-feedback TIA limit: a single-pole core amplifier (A1), a resistive feedback element (A2), and the full input capacitance loading the feedback summing node (A3). Relaxing one or more of these is shown to be the common structural thread behind every class of bandwidth or noise enhancement in the recent literature, and all surveyed architectures are organized into a six-tier taxonomy, from Tier 0 designs operating within the classical limit to Tier 5 topologies that bypass all three assumptions simultaneously. This taxonomy is supplemented by an orthogonal configurability axis spanning single- and dual-control reconfigurable, variable-gain, and dynamic-range-extension designs. We further show that stability is not removed by these relaxations but migrates with the tier, from the global phase margin of the classical loop to a local regulating loop, a group-delay-flatness constraint, an input-passivity condition, or a multi-loop interaction, so that each architecture carries a predictable stability locus alongside its noise and bandwidth consequences. The taxonomy is cross-referenced with application domains, with closed-form noise-floor boundary plots parametrized by input capacitance and amplifier gain-bandwidth product, and with the CMOS technology landscape, where we argue that the most advanced node is not universally optimal and that node and topology act as complementary rather than competing levers. A single consistent figure of merit, applied uniformly to a representative set of CMOS realizations from 0.6 μm to 16 nm FinFET, shows no monotonic improvement with publication year or node and is presented as a diagnostic indicator rather than an absolute ranking. The review closes with an outlook on 200 Gb/s /lane links, wide-bandgap sensor integration, and the FinFET-to-gate-all-around device transition.

Keywords: transimpedance amplifier (TIA); CMOS analog circuits; topology taxonomy; transimpedance limit; assumption relaxation; bandwidth enhancement; low-noise design; optical receivers; inductive peaking; regulated cascode

1. Introduction

The transimpedance amplifier (TIA) is a small circuit with an outsized role. It is the first active stage that turns a weak input current, most often the photocurrent of a photodiode (PD), but equally the output of a capacitive, resistive, or current-mode sensor, into a voltage that the rest of the system can amplify, process, and digitize. In this sense the TIA is one of the principal interfaces through which electronics senses the physical world, and it sits at the very front of the receiver chain, where its noise, bandwidth, and sensitivity propagate directly into the system signal-to-noise ratio (SNR) and

 doi:10.20944/preprints202606.0753.v1

ultimately bound the performance of the entire receiver [1,2]. Because it is the first building block, every imperfection it introduces is amplified downstream and can no longer be recovered, which is precisely why the TIA so often decides whether a receiver meets its sensitivity target at all.

This unassuming block underpins much of the modern data economy. The relentless expansion of the global data sphere, currently driven by 5G/6G deployment, the scaling of artificial-intelligence (AI) compute clusters, hyperscale data centers, and the rise of autonomous sensing, rests on optical links whose ultimate reach and rate are gated by the receiver front-end, and therefore by the TIA [3]. As link rates climb toward 200 Gb/s per lane and modulation moves from simple on-off keying to multilevel Pulse-Amplitude Modulation (PAM) and coherent formats, the TIA ceases to be a mere gain stage and becomes a signal-conditioning engine that must reconcile noise, bandwidth, linearity, and power simultaneously within an ever-tighter supply-voltage budget.

What makes the TIA a perennial research object is that these requirements are not independent. For the canonical Resistive Shunt-Feedback (SFB) TIA, transimpedance gain, bandwidth, and input-referred noise are bound together by the so-called transimpedance limit: the feedback resistance that maximizes the gain and minimizes the noise is also the one that collapses the bandwidth, and the input capacitance couples the gain, bandwidth, and noise penalties to a single physical source [2–4]. The SFB topology is valued precisely because this trade-off can be reformulated into a clean closed form, which has made it at once the textbook reference design and the baseline against which every subsequent enhancement shall be measured.

In practice, however, two rather different design cultures have grown up around this circuit. In the first, the TIA is merely one block among many inside a larger integrated circuit. This can be a high-speed receiver co-integrated with clock-and-data recovery, Continuous-Time Linear Equalizers (CTLE), Variable-Gain Amplifiers (VGA), Analog-to-Digital Converters (ADC), and Digital Signal Processing (DSP) on a single die. Here a robust, well-understood SFB cell is frequently the pragmatic choice, with the heavy lifting of bandwidth recovery and adaptation pushed off-loop into the surrounding system and its DSP [5–9]. In the second, the TIA itself is the subject of the research, and considerable ingenuity is invested in a single, elaborate cell that pushes against the transimpedance limit: regulated-cascode input stages, capacitive feedback, inductive and active peaking, or fully distributed gain [10–15]. Historically, comparatively few works have married those advanced standalone TIA cells with full system integration as these two cultures optimize different targets. While the system robustness and yield are prioritized on one side, peak cell-level performance becomes the target objective on the other. In practice, they rarely meet in the same publication. This split is itself a useful lens, because it explains why nominally comparable circuits report incomparable Figures of Merit (FOM), and why an architecture that looks optimal in isolation may be unnecessary once a processing chain is already in place.

Although optical communication and its many variants have historically been the dominant driving force behind the development of TIAs, starting with long-haul and datacenter links to passive optical networks, Light Detection and Ranging (LiDAR) [16], optical fiber sensing [17], and Optical Time-Domain Reflectometry (OTDR) [3,18–20], the application space of the TIA is now considerably broader than optics alone. Low-power, low-noise front-ends for biomedical and biosensing instrumentation operate at bandwidths as low as a few hundred kilohertz while demanding femto- to pico-ampere noise floors [21–23]. Among other applications one shall mention radiation detectors [24], infrared spectroscopy [25], magnetic-resonance and NMR receiver-coil arrays [26], radio-over-fiber small-cell front-ends with the frequency selectivity folded into the transimpedance interface itself [27]. Finally, significant group is formed by MEMS and resonant-sensor readouts [28,29] and emerging cryogenic and qubit-readout interfaces shall be mentioned with each type imposing its own requirements, often dramatically different from those in optical applications. In the magnetic-resonance case, for instance, high linearity and dynamic range become the primary specification while the transimpedance gain is deliberately kept low so that the input node remains a virtual ground for the coil [26]. Obviously, this is an opposite emphasis to a high-gain optical receiver. These emerging domains do not merely

 doi:10.20944/preprints202606.0753.v1

reuse optical-receiver circuits, but often bring new constraints (ultra-low bandwidth, extreme dynamic range, high detector bias voltage, wide-bandgap detector capacitance) and, in turn, motivate new circuit-level design ideas. This reinforces the observation that no single TIA architecture is best for all applications, and that competing designs are most usefully understood through their complementary properties. Note that even within the optical domain, for short-distance links much higher input-referred noise currents can be tolerated than for long-distance links and, therefore, noisier TIA front-ends can be adopted [30].

Given this breadth, the recent literature is difficult to navigate. Authoritative foundations do exist such as the textbook treatment of optical-receiver TIAs and the formalization of the transimpedance limit by Sackinger [2,4], together with application-specific surveys such as the OTDR-oriented review of [19] and the CMOS MEMS resonator-oriented review of [29]. Similarly, a domain-specific comparative study of TIA topologies for biomedical applications, such as photoplethysmography and Near-Infrared Spectroscopy (NIRS), was presented by Atef et al. [23]. Their work evaluates various architectures across distinct simulation scenarios, ranging from 5 kHz bandwidths with 10 pF PD capacitance for low-power, low-noise medical instruments, up to 100 MHz bandwidths with 2 pF capacitance for higher-speed demands. Additionally, a comparative architectural study tracking the advancement of CMOS TIAs for data communication and optical receivers was compiled by Badal [31], highlighting design optimization paths across varied instrumentation frameworks. Yet the broader body of recent CMOS work still reads largely as a fragmented collection of circuit-level modifications such as ad-hoc inductive peaking, empirical compensation schemes, and design choices that appear to resist any systematic ordering. The difficulty is therefore twofold. First, there is a missing structure, in that individual innovations are seldom related to one another and one may need a significant design expertise to navigate the solution. Secondly, a clear design taxonomy is missing as there seems to be no agreed framework into which a given circuit-level decision can be placed and compared. A review that addresses neither of these points has a risk of degenerating into yet another exhaustive or not so exhaustive catalogue of publications rather than an analysis of the design principles behind them. The same navigability problem is now also being attacked from the design-automation side, where symbolic exploration frameworks enumerate and prune candidate analog topologies before committing to costly simulation [27]. Here a human-readable taxonomy and a machine-readable topology space are complementary responses to the same underlying fragmentation of TIA design approaches. Recent domain-specific surveys illustrate this tendency directly. For example, the CMOS-MEMS review of [29] enumerates roughly ten front-end topologies side by side, each accompanied by its gain, bandwidth, and noise expressions, but without an organizing principle that would relate them to one another or explain why each emerges from the same underlying limit. A similar flat enumeration is observed in the review by Badal [31], which summarizes performance trends and specific design specifications across discrete topologies but leaves the broader structural taxonomy unaddressed. This tendency is equally visible in the biomedical survey by Atef et al. [23], which systematically compares noise, gain, and power trade-offs of established TIA topologies under specific application constraints and derives their mathematical models, yet without providing a unified structural classification. Such enumerations are valuable as domain-scoped lists, yet they leave open precisely the structural question this review attempts to address.

This work seeks to close that gap by viewing the field through the lens of the transimpedance limit and the assumptions that underlie it. We show that the SFB limit rests on three explicit structural assumptions: a single-pole core amplifier (A1), a resistive feedback element (A2), and the full input capacitance loading the feedback summing node (A3). Assuming this, one may notice that essentially every reported bandwidth- or noise-enhancement technique can be read as the relaxation of one or more of these assumptions. This yields a compact, six-tier taxonomy that maps a sprawling literature onto a small, ordered space of assumption relaxations, each with predictable consequences for the gain-bandwidth-noise trade-off. We do not claim this to be the only valid way of organizing the field. It is only one suggested perspective that we have found both structurally honest and pedagogically

 doi:10.20944/preprints202606.0753.v1

useful, and other orderings are certainly possible. We also state plainly that this review does not attempt to enumerate all TIA designs. Such completeness attempt would be infeasible even for a textbook. Instead, we use a deliberately curated, representative set of publications to demonstrate that the taxonomy holds, drawing on the corresponding works as supporting evidence rather than as an exhaustive inventory.

Within that aim, the scope of the article is defined by three objectives:

- **Temporal relevance**: analysis shall include both the classical the works and those from the last several years including a period spanning the industry's transition from planar CMOS through FinFET toward gate-all-around devices.
- **Taxonomic clarity**: a move away from the exhaustive-list format toward a grouping of works by the underlying assumption they relax relative to the reference SFB TIA.
- **Application breadth**: a treatment that spans the optical-communication and sensing ecosystems together, from high-speed datacenter links to high-sensitivity LiDAR, biomedical, and emerging readout circuits, while remaining representative rather than exhaustive.

The remainder of the paper is organized as follows. Section 2 introduces the governing performance metrics and trade-offs, including the additional requirements imposed by the PAM-4 and Co-Packaged Optics (CPO) era. Section 3 reviews the technological landscape, beginning with the historical migration from III-V compound semiconductors to CMOS and continuing through the planar-to-GAAFET device evolution and the rise of CPO and Silicon Photonics (SiPh). Section 4 develops the transimpedance limit and the assumption-relaxation taxonomy based on SFB TIA architecture, while Section 5 surveys several well-known alternative topologies from CG TIA through to distributed and current-mode front-ends. Section 6 examines advanced bandwidth- and noise-optimization techniques, and Section 7 shows how the same taxonomy organizes the stability behavior of each tier. Section 8 turns to digitally-assisted and high-linearity TIAs. Section 9 provides a comparative, FOM-based benchmark of representative designs across technology nodes and publication years. Finally, Section 10 discusses the outlook toward 200 Gb/s-per-lane links, wide-bandgap sensor integration, and the FinFET-to-gate-all-around transition, and Section 11 concludes the work.

2. Fundamentals and Performance Trade-Offs

The TIA converts the PD current into a voltage and, being the first block of the receiver chain, fixes the envelope within which every subsequent stage must operate. Its design is therefore a solution to a multi-objective optimization problem in which transimpedance gain, bandwidth, input-referred noise, phase linearity, dynamic range, power, and area cannot be improved independently [14,20]. This section defines the parameters that quantify these objectives, explains why each is important for a TIA specifically, and indicates in general terms how each is improved. The concrete circuit techniques, however, and the formal gain-bandwidth-noise bound known as the *transimpedance limit*, are deferred to Section 4.

2.1. Gain, Bandwidth, and Input-Referred Noise

The transimpedance gain is the small-signal ratio of output voltage to input current,

$$Z_T(s) = \frac{V_{\text{out}}(s)}{I_{\text{in}}(s)}, \qquad s = j\omega, \tag{1}$$

where $R_T = |Z_T(0)|$ is often called the transresistance [20]. A large Z_T is desirable because it relaxes the gain and noise requirements of the following main amplifier (MA). Since the output buffer contributes little gain, it is the TIA and MA together that set the overall conversion gain. The difficulty is that gain and bandwidth are strongly coupled, so raising R_T, whether through a larger feedback resistor R_F or a higher core-amplifier gain, is bounded by both stability and the target bandwidth.

The bandwidth $\text{BW}_{-3\text{dB}}$ is the upper frequency at which $|Z_T|$ falls 3 dB below its mid-band value. It must always be quoted together with the maximum permitted in-band peaking, because

 doi:10.20944/preprints202606.0753.v1

bandwidth can be inflated by tolerating peaking that a given application may not accept [20]. Note that AC-coupled optical front-ends additionally exhibit a low-frequency cutoff. The bandwidth must span the signal spectrum (for Non-Return-to-Zero (NRZ), roughly $0.7\times$ the symbol rate), where too little of that causes intersymbol interference (ISI), while excess bandwidth needlessly integrates additional noise.

The input-referred noise current (equivalently, the equivalent input noise current) determines the receiver sensitivity, i.e. the smallest detectable signal. It is obtained by referring the output noise back to the input through the gain,

$$\overline{i^2_{n,\mathrm{TIA}}}(f) = \frac{\overline{v^2_{n,\mathrm{TIA}}}(f)}{|Z_T(f)|^2}, \tag{2}$$

and is reported as a spectral density in $\mathrm{pA}/\sqrt{\mathrm{Hz}}$. Because the density is frequency dependent, a rigorous comparison requires the full spectrum description. In practice, however, a band-averaged value is quoted, and the total RMS noise follows from integrating $|Z_T|^2\,\overline{i_n^2}$ to at least twice the bandwidth and dividing by R_T [20]. The TIA dominates the receiver noise because it sits closest to the signal source, so the noise of later stages is suppressed by the preceding gain. Crucially, a noise figure is meaningful only when quoted alongside the PD capacitance, and in practice the ESD and pad capacitances as well, since the total input capacitance C_T sets the source impedance and shapes both noise and bandwidth. Noise is lowered, in general, by maximizing the gain and input-device transconductance while minimising C_T.

2.2. Phase Linearity, Passband Flatness, and Dynamic Range

Beyond magnitude, the group delay $\tau_g(\omega) = -\,\mathrm{d}\phi/\mathrm{d}\omega$ characterizes the phase response. A linear phase yields a constant group delay, and the group-delay variation $\Delta\tau_g$, the largest deviation within the band, is the primary measure of linear (phase) distortion [20]. Excess $\Delta\tau_g$ spreads pulses and produces ISI. This metric is in direct tension with bandwidth extension: the high-Q complex poles introduced by inductive or active peaking, while widening the band, often also degrade $\Delta\tau_g$. Unfortunately, this interesting characteristic of a TIA design is frequently omitted when new TIA designs are reported [32,33].

The same poles also govern passband flatness. For a second-order response, $Q = 1/\sqrt{2}$ (Butterworth) gives maximally flat magnitude with no peaking, $Q = 1/\sqrt{3}$ (Bessel) gives maximally flat group delay, and higher Q trades a faster roll-off and an approximately 40% bandwidth extension for magnitude ripple and elevated $\Delta\tau_g$ [20]. Passband ripple, defined the gain variation across the band, manifests as overshoot and ringing in the pulse response and is a gain error that any downstream equalizer must absorb. Since magnitude flatness and phase flatness cannot be maximized simultaneously, the designer aligns the response toward Bessel-like behavior for phase-critical links or Butterworth for amplitude-critical ones. This choice is invisible to a single bandwidth number, which is precisely why peaking and $\Delta\tau_g$ must be reported next to it.

Finally, the dynamic range is the ratio of the maximum to the minimum detectable input current. The lower bound is set by the input-referred noise (sensitivity) and the upper bound by overload and gain compression [20]. A low-impedance front-end offers wide bandwidth and large dynamic range at the cost of sensitivity, whereas a high-impedance front-end maximizes sensitivity and low noise but is limited in dynamic range and bandwidth. The range is extended by varying the transimpedance with signal strength, for instance through Automatic Gain Control (AGC) or controlled compression [17,34,35].

2.3. Linearity, PAM-4, and Co-Packaged Optics

The migration from NRZ to four-level PAM-4 has expanded the criteria for a good TIA beyond the typical gain-bandwidth-noise triplet. PAM-4 encodes information in four amplitude levels, i.e. three stacked eyes, so any gain compression squeezes the levels unequally in a way a linear equalizer cannot fully undo. Linearity, quantified by the Total Harmonic Distortion (THD) and the 1-dB compression

 doi:10.20944/preprints202606.0753.v1

point (P_{1dB}), therefore becomes a primary design target rather than an afterthought as it was for compression-tolerant NRZ [5,6]. Active bandwidth extension techniques are a common offender here as the effective impedance of an active inductor varies with the instantaneous signal swing. A sufficiently linear TIA is what allows the downstream feed-forward and decision-feedback equalizers to reopen the eye.

Phase linearity, already discussed above, is correspondingly more demanding under PAM-4. Modern high-speed designs target a Group-Delay Variation (GDV) below roughly ± 5 ps across the -3 dB band to meet jitter and ISI budgets [32]. Two further constraints arise from the integrated, high-density environment. Power-supply rejection matters because the TIA shares a die with switching DSP, where poor rejection converts supply ripple into timing jitter, motivating low-dropout regulation and fully differential signaling. For CPO, the input return loss (S_{11}) becomes relevant for matching to the optical engine across the short interconnect, while burst-mode applications such as passive optical networks add a fast settling-time requirement as the TIA adapts to packet-to-packet power changes. These advanced metrics do not displace the primary triplet, but rather add important application-weighted constraints, and it is the relative weighting, rather than any single number, that distinguishes a datacom front-end from a MEMS sensing or a LiDAR front-end.

3. Technological Landscape

The technological landscape for TIA has undergone a significant shift, driven by the continuous scaling of CMOS technology and the emergence of heterogeneous integration. While historical designs relied on the high transconductance of SiGe BiCMOS, the modern requirement for SoC integration and lower cost-per-bit has positioned CMOS as the dominant platform for high-speed optical front-ends.

Historically, high-performance optical receivers were the domain of III-V compound semiconductors (InP, GaAs HBT) owing to their superior electron mobility and transit frequency f_T [3,15]. The migration of lightwave front-ends toward silicon followed a well-documented technology-substitution pattern: at 155 Mb/s and 622 Mb/s the shift from bipolar to CMOS occurred first, while GaAs HBT remained dominant at 2.5 Gb/s and 10 Gb/s well into the 2000s [36]. A landmark systematic benchmark of a 0.13 μm CMOS process for 40 Gb/s optical communication was provided by Razavi (2002) where the author showed that the simulated f_T of NMOS devices falls to approximately 62 GHz at the slow high-temperature process corner, making direct 40 GHz operation challenging, and three-stage differential ring oscillators with resistive loads were found to oscillate at only 18 GHz. This finding motivated the shift from ring-oscillator PLLs toward LC-VCO-based clock synthesis and inductive bandwidth-extension techniques that now form the core of the CMOS high-speed TIA toolkit. The driving forces behind this migration are economic and architectural rather than purely performance-driven.

On the cost side, high-volume CMOS wafer costs are substantially lower than III-V fabrication and continue to decline with each process generation. Large multi-customer foundry capacity also removes the need to own or exclusively partner with a dedicated fabrication facility, lowering the capital barrier to entry [36]. These arguments apply directly to TIA production at datacom volumes.

The dominant architectural driver is, of course, the SoC integration density. CMOS allows the TIA to be co-integrated with multi-channel DSPs, CDR circuits, and memory arrays on a single die, eliminating the power-hungry and bandwidth-limiting electrical interconnects that would otherwise connect a standalone III-V TIA to a CMOS back-end. The availability of millions of transistors also enables on-chip ESD protection structures and BIST circuitry that are impractical in III-V technologies [36]. For CPO and SiPh, the requirement that the TIA reside within the same package as the host ASIC makes a III-V front-end physically incompatible with 2.5D and 3D integration schemes.

Node scaling provides a third incentive. Transitioning from planar CMOS through FinFET to GAAFET ($\leq$5 nm) increases f_T/f_{max} and the raw transconductance density per unit area, enabling bandwidth-extension techniques (active inductors, inductive inter-stage peaking) that were impractical at older nodes [3,37]. Equally important is the proliferation of metal layers: a modern 0.13 μm

 doi:10.20944/preprints202606.0753.v1

generation already offers eight metal layers, providing the passive-device substrate (spiral inductors, stacked fringe capacitors, transmission lines) that is indispensable for inductive peaking, LC VCOs, and the fringe bypass capacitors used in TIA inter-stage coupling [3].

The low supply voltage imposed by thin gate oxides at advanced nodes restricts the voltage headroom available for cascode and Regulated-Cascode (RGC) topologies, and severe parasitic capacitances degrade bandwidth at each gain stage [37]. This headroom constraint was articulated early in the CMOS-TIA literature by Kromer et al. in [30]. As the supply voltage was scaled down with each node while the threshold voltage was not reduced proportionally, two gate-source voltages can no longer be stacked at an 80 nm/1 V operating point, which excludes conventional cascode and source-follower stages and constrains the regulating-loop bias of a plain RGC. The same scaling lowers the intrinsic device gain below ten, so that high transimpedance gain demands several cascaded stages with significant power and stability costs [30,38]. Note that this is precisely the constraint that motivates the headroom-releasing input stages discussed in Section 4. The low supply voltage imposed by thin gate oxides at advanced nodes restricts the voltage headroom available for cascode and RGC topologies, and severe parasitic capacitances degrade bandwidth at each gain stage [37]. The MOSFET transconductance-per-current ratio is lower than that of an HBT, leading to higher input-referred channel noise for a given bias current [39]. Substrate resistivity in bulk CMOS is orders of magnitude lower than in semi-insulating GaAs, increasing substrate coupling through on-chip spiral inductors and raising crosstalk in multi-channel arrays [3,36] even though the deep n-well available in modern processes can partially mitigate this by isolating sensitive analog nodes from substrate noise [3]. Finally, CMOS foundry RF compact models (BSIM and successors) have historically lagged III-V models in large-signal accuracy, creating simulation-to-silicon discrepancies that add design margin and reduce first-pass success [36].

This residual noise gap between competing technologies is still real and quantifiable. SiGe BiCMOS, combining inherently lower noise with superior $f_T/f_{\max}$, has long been the powerhouse for TIAs, and its noise advantage remains significant even at moderate speeds [5]. For example, [24,40] demonstrate that a 0.35 μm BiCMOS TIA using a Common-Emitter Common-Base (CECB) cascode achieves 2.13–2.44 pA/$\sqrt{\text{Hz}}$ at 70 dBΩ and 2.3 GHz, roughly an order of magnitude beyond CMOS designs at comparable gain [41]. The higher g_m/I_C of the BJT directly reduces input-referred channel noise, and the CB cascode fully suppresses the second-stage noise contribution [24]. Unlike GaAs HBT, SiGe BiCMOS is fabricated in a silicon foundry and can be co-integrated with CMOS logic while retaining the g_m/I_C advantage of a bipolar device [24]. The practical consequence for designers choosing bulk CMOS is that this noise disadvantage cannot be recovered by topology alone it and must be absorbed through the circuit-level innovations such as active feedback, noise cancellation, and RGC input isolation as discussed below. The justification for CMOS is therefore integration economics and SoC density, not intrinsic analog superiority.

These device-level penalties feed a more general design caution: advancing to the most aggressive available node is a lever, not an unconditional gain, and the measured power-performance-area behavior of real silicon departs from roadmap idealizations enough that the node must be chosen against the application rather than assumed from the calendar. A single-foundry benchmark spanning a low-leakage CMOS family from 180 nm to 28 nm by [42] quantifies this as although each generation yields, on average, a 57% frequency increase and a >20% power-delay-product reduction, the gains are non-uniform. The same study documents the rising process variability, model and design-rule complexity, and PVT sensitivity (including the sub-40 nm temperature-inversion effect) that accompany scaling, and concludes that real performance requires a per-process evaluation before node selection [42]. This is the technology-side counterpart to the architectural result reached later in this review (Section 11) as node and topology are complementary rather than competing levers, and a well-chosen topology at a mature node can outperform a conservative design at an advanced one.

For a large class of TIA applications this calculus favors a mature node outright. Where the PD is co-integrated on-die, the APD photo-response constrains the usable metal stack and biasing, which

 doi:10.20944/preprints202606.0753.v1

is one reason recent LiDAR receivers are realised in 180 nm CMOS despite the raw speed available at finer nodes [16]. Similarly, where the detector demands tens of volts of reverse bias, as for GaN and Ga_2O_3 APDs, a high-voltage mature process becomes mandatory and runs directly counter to the sub-1 V scaling trend (Section 10). More broadly, for the sensing, biomedical, and LiDAR front-ends surveyed here the binding specifications are noise, dynamic range, and detector compatibility at modest bandwidth and none of which scales favorably with the node. Thus the lower design complexity, reduced variability, and lower cost of 130-180 nm CMOS frequently dominate the decision.

The migration from planar MOSFETs to three-dimensional structures represents the most significant change in the TIA design environment. At technology nodes below 28 nm, traditional planar devices suffered from severe short-channel effects and high leakage currents, which limited the achievable DC gain and increased noise floors.

- **Fin Field-Effect Transistor (FinFET) Technology**: Introduced at the 22 nm node, the FinFET replaced the flat channel with a vertical silicon **fin**. This 3D geometry allowed the gate to wrap around the channel on three sides, providing superior electrostatic control and enabling higher drive currents within a smaller footprint. For TIAs, this resulted in improved f_T/f_{max} ratios, though it introduced new parasitic capacitances associated with the fin-to-gate sidewalls that must be carefully managed.
- **FD-SOI (Fully Depleted Silicon-on-Insulator)**: As a parallel path to FinFET, FD-SOI offers a planar alternative that utilizes an ultra-thin buried oxide layer. This architecture minimizes junction parasitics and allows for back-biasing, a technique where the threshold voltage V_{th} can be dynamically tuned to optimize the TIA for either high-speed performance or ultra-low power consumption depending on the link state.
- **Gate-All-Around FET (GAAFET) and Nanosheets**: For sub-3 nm nodes, the industry is transitioning to GAAFET or nanosheet transistors. By surrounding the channel on all four sides, GAAFETs offer a near-perfect control over the channel potential, further reducing leakage and allowing for lower supply voltages V_{DD}. In TIA design, this scaling provides the bandwidth necessary for 200 Gbps-per-lane targets but necessitates innovative circuit techniques to handle the extremely limited voltage headroom.

On the other side, at least on the optical domain the technological landscape is currently being redefined by CPO. In older discrete modules, the TIA is separated from the host ASIC by several centimeters of PCB traces, leading to significant signal degradation and power loss. CPO seeks to disrupt this by bringing the optical engine and the TIA directly onto the same package as the main switch or the processing unit. This shift places new demands on TIA technology:

- **Density and Thermal Management**: CPO arrays require 1D or 2D TIA banks. This necessitates inductor-less or active-inductor designs to minimize silicon area, as large spiral inductors prevent the high-density packing required for terabit-scale interconnects.
- **Optical-Electrical Co-Optimization**: With the PD placed in close proximity to the TIA via micro-bumps or monolithic SiPh, the input parasitic capacitance C_T is significantly reduced. TIA designers are leveraging this to push bandwidth boundaries without the traditional noise penalties associated with high-capacitance wire-bonded interfaces.

4. Resistive Shunt-Feedback TIA and the Transimpedance Limit

The resistive SFB TIA is the reference topology against which every other architecture is measured. This section introduces this famous circuit and its trade-offs, derives the closed-form transimpedance limit that governs it, and abstracts the three assumptions behind that limit into the assumption-relaxation taxonomy used throughout the remainder of the paper. The architectures that overcome the limit by other means are deferred to Section 5. This primacy is not specific to optical receivers: even in domain-specific surveys of CMOS-MEMS resonator front-ends, the resistive SFB TIA is identified as the simplest topology and the one offering the most favorable gain-noise-bandwidth compromise among the alternatives considered [29].

 doi:10.20944/preprints202606.0753.v1

4.1. The Resistive Shunt-Feedback TIA

Any discussion of CMOS TIAs starts with the classical resistive SFB TIA of Figure 1(a) with a voltage amplifier of gain $-A$ closed by a feedback resistor R_F. The shunt feedback simultaneously presents a low input impedance, which is desirable for a current-sensing front-end that must cope with a large external PD capacitance C_D, and a low output impedance suited to driving the subsequent stage. The topology is ubiquitous because R_F carries no large bias current and can therefore be made large, and because the feedback sets the transimpedance without the headroom penalty incurred by feed-forward configurations [1].

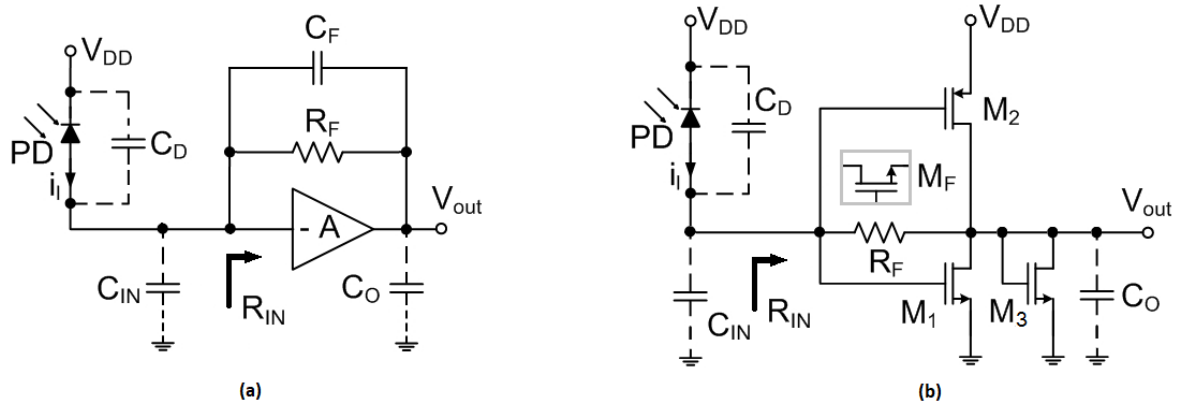

Figure 1. Basic configuration of a resistive SFB TIA: (a) version with an optional feedback capacitor C_F. (b) Push-pull inverter-based SFB TIA with an active load and triode biased NMOS device instead of the feedback resistor.

Its weaknesses follow from the same loop. The interaction of R_F with C_D sets the dominant pole, so a large or poorly-controlled C_D both limits bandwidth and degrades stability when the amplifier must serve PDs of varying capacitance [43]. Furthermore, under realistic speed, headroom and power budgets the input transistor can contribute as much input-referred noise as R_F itself [1]. A small compensation capacitor C_F placed across R_F restores a maximally flat response and reduces the sensitivity to C_D, but it also lowers the achievable transimpedance and so is of a limited use when larger gain is a major objective. A common variant replaces the feedback resistor with a triode-biased NMOS device tunable through its gate voltage [44],

$$R_F = \frac{1}{\frac{W}{L}\,\mu C_{ox}\,(V_{GS} - V_T)}, \tag{3}$$

which adds gain programmability at the cost of severe sensitivity to process variation once the overdrive becomes small, i.e. when R_F is made large [7]. Apart of the classical CS-based implementation, the inverter-based approach often accompanied by an active load is used (see Figure 1(b)). Despite all known issues, such an inverter-based SFB cell remains widely adopted for its simplicity, and a version of it attains among the highest FOM values in the surveyed set [16]. This emphasizes the ongoing necessity of optimizing conventional topologies to handle the strict trade-offs imposed by input node capacitance dependency, a core performance bottleneck detailed extensively in [31]. For the voltage amplifier itself the designer has considerable freedom. As the TIA noise is dominated by the first stage, a simpler amplifier with fewer active devices is preferable when it can supply enough gain, and a single stage remains viable if so. Many designs instead cascade push-pull inverters (three in [44], five in [45]). Recall that the inverter is a power-efficient analog gain element that delivers twice the g_m for a given drain current, tolerates the sub-1 V supplies of advanced nodes [5] and accommodates larger input currents. The earlier view that inverter front-ends are too slow for high-speed optics (see comments e.g. in [37]) seem no longer to holds as a 16 nm FinFET inverter TIA with only 10 GHz of intrinsic bandwidth reaches 32 GHz once a downstream CTLE recovers the band [5].

 doi:10.20944/preprints202606.0753.v1

4.2. The Transimpedance Limit and Its Modeling

Let C_T denote the total capacitance at the input (summing) node, comprising the PD itself, bond-pad, ESD and amplifier-gate contributions. For a single-pole core amplifier of gain-bandwidth product GBW, demanding a maximally flat (Butterworth) response ties the closed-loop bandwidth to R_F and C_T,

$$f_{-3\,\mathrm{dB}} \approx \sqrt{\frac{\mathrm{GBW}}{2\pi R_F C_T}}, \tag{4}$$

which rearranges into the maximum transimpedance attainable at a target bandwidth as follows:

$$R_{F,\max} \leq \frac{\mathrm{GBW}}{2\pi f_{-3\,\mathrm{dB}}^2 C_T}. \tag{5}$$

For multi-stage cores the effective GBW must be taken from the dominant pole of the loop gain rather than the open-loop unity-gain frequency. Three implications follow from the formulation above. The dependence is square-law and doubling the bandwidth costs a factor of four in transimpedance. Because the input-referred thermal noise is dominated by R_F, a smaller R_F raises the noise floor, so the transimpedance limit is simultaneously a sensitivity limit [4]. And the R_F-C_T interaction (only partly tamed by C_F) produces the phase shift that bounds how far the loop can be pushed before peaking turns into ringing.

While the limit is universal, CMOS makes it hard to approach. Falling supply voltages cap the DC drop allowed across R_F and, at 3-5 nm FinFET/GAAFET, even a single cascode is constrained by the minimum device stack needed for useful gain. The effective in-loop GBW is far below the raw f_T once gate resistance and parasitics are included (typically $f_T/5$-$f_T/10$) and the e.g. the inverter stages used to recover g_m at low voltage are sensitive to PVT drift away from the Butterworth point. A compounding and often-overlooked factor is the operating point at which the raw f_T is quoted as that peak f_T is specified at the maximum supply, whereas analog front-ends bias their devices at roughly half V_{DD} for signal swing, so the f_T available to the circuit is already substantially below the headline figure before parasitics are counted. Authors in [38] note that the nominal 200 GHz f_T of a 40 nm process falls below 100 GHz at a practical analog bias, which together with the $f_T/5$-$f_T/10$ explains why a nominally fast node often delivers a far more modest usable GBW [38].

Figures 2 and 3 plot the resulting noise floor. Substituting $R_{F,\max}$ from (5) into the dominant thermal term $i_n = \sqrt{4kT/R_F}$ gives:

$$i_n = \sqrt{\frac{8\pi\, kT\, C_T}{\mathrm{GBW}}}\, f_{-3\,\mathrm{dB}}, \tag{6}$$

with k the Boltzmann constant and T being the temperature. No SFB TIA at a given C_T and GBW can sit below the corresponding line as long as the three core assumptions hold. The lines are strictly parallel on log-log axes (slope 1, +20 dB/decade), so doubling the bandwidth always costs 6 dB of noise floor. Increasing C_T shifts every line up by $20\log_{10}\sqrt{C_{T,2}/C_{T,1}}$ dB, coupling the bandwidth and noise penalties to the single physical cause of input capacitance. The largest gains therefore come not from approaching the limit but from architectures (RGC, capacitive feedback, distributed gain) that structurally decouple C_T from the dominant pole. In Figure 3 the GBW values are parametric because the in-loop value is only $f_T/5$–$f_T/10$ of the transistor f_T. Advancing a technological node shifts the floor down by $20\log_{10}\sqrt{\mathrm{GBW}_2/\mathrm{GBW}_1}$ dB without changing the slope, so node scaling and architecture act as complementary rather than competing measures.

Several routes overcome the $1/\mathrm{GBW}^2$ penalty without leaving the SFB family in spirit: an RGC input lowers the effective C_T seen by the loop and admits a larger R_F, inductive peaking adds zeros that cancel the dominant pole and extend bandwidth at constant R_F (quality factors as low as $Q = 3$-4 already yield $\approx$ 50% extension in a single stage, against 82% for an ideal inductor, which justifies compact spirals over large high-Q structures [3]), while local feedback loops, as in multi-stage and

 doi:10.20944/preprints202606.0753.v1

Cherry-Hooper (CH) designs, decouple gain from bandwidth and replacing R_F with a capacitor removes the resistor's thermal noise outright. Distributed-gain, noise-matched-input and current-mode techniques sidestep the limit at the system level by never accumulating the full C_T at one node.

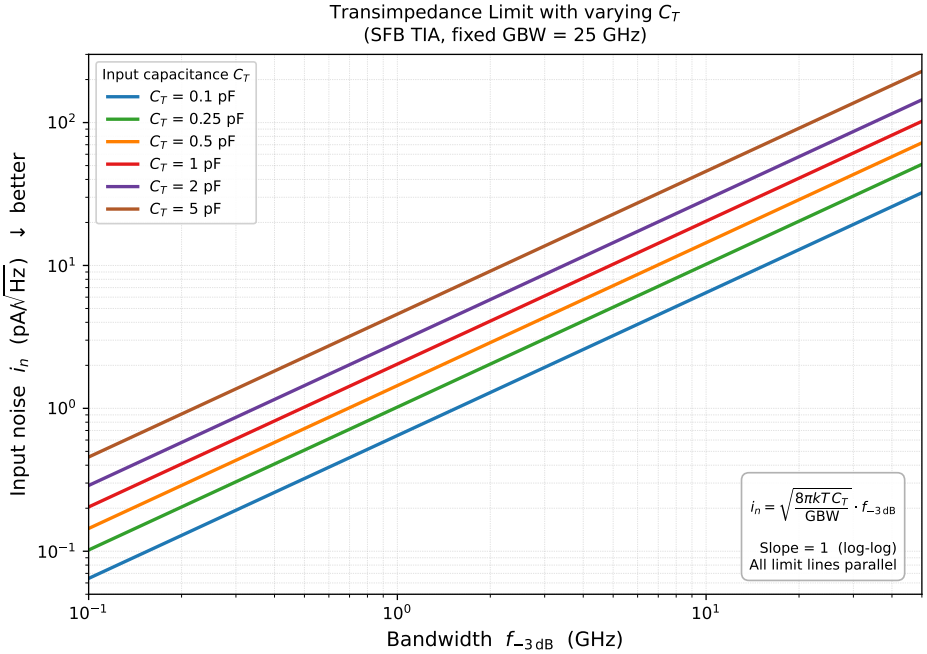

Figure 2. Input-referred noise current density i_n versus -3 dB bandwidth $f_{-3\,\mathrm{dB}}$ of the transimpedance limit for an SFB TIA at fixed GBW $= 25$ GHz, for six values of total input capacitance $C_T \in \{0.1, 0.25, 0.5, 1.0, 2.0, 5.0\}$ pF.

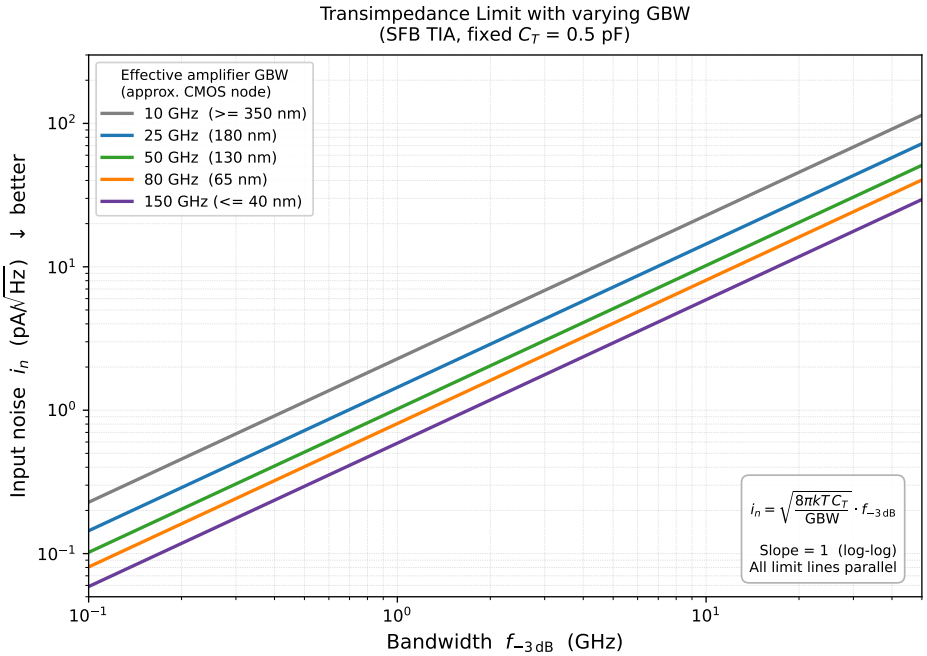

Figure 3. Input-referred noise current density i_n versus -3 dB bandwidth $f_{-3\,\mathrm{dB}}$ of the transimpedance limit for an SFB TIA at fixed $C_T = 0.5$ pF, for five values of effective GBW $\in \{10, 25, 50, 80, 150\}$ GHz, corresponding approximately to the $\geq$350, 180, 130, 65 and $\leq$40 nm CMOS nodes.

4.3. The Three Assumptions and the Assumption-Relaxation Taxonomy

The limit (5) holds strictly only for a single-loop, single-pole, resistive SFB TIA, but it matters far beyond that case because it rests on three independent assumptions:

 doi:10.20944/preprints202606.0753.v1

A1 The core amplifier has a single dominant pole, $A(s) = A_0/(1 + s/\omega_0)$, giving a well-defined GBW;

A2 The feedback element is resistive, so its thermal noise $4kT/R_F$ enters the signal path directly;

A3 The full input capacitance C_T (PD, pad, ESD, gate) sits at the summing node and forms the dominant pole with R_F.

These assumptions map onto three physical handicaps: the amplifier's finite speed, the resistor's thermal noise, and the capacitive loading of the feedback node. Every advanced architecture attacks one or more of them, which is what makes the limit a useful organizing baseline rather than a property of one circuit. When a design claims, say, a 3× better noise-bandwidth product than *a conventional TIA*, the baseline being beaten is the SFB design governed by (5). An RGC presents a low impedance to C_D and relaxes A3, inductive or active peaking injects zeros and relaxes A1 while capacitive feedback removes R_F and relaxes A2. Comparable limits exist for every topology, but the SFB case is the one where the algebra collapses to a clean closed form, making it the natural pedagogical entry point [4].

This motivates us to suggest a six-tier *assumption-relaxation taxonomy* that classifies any CMOS TIA by which assumptions its mechanism breaks: Tier 0 relaxes none and operates within the limit, i.e. by exchanging the core amplifier with a more performing one. Tier 1 relaxes A3 (decoupling C_T from the feedback pole) while Tier 2 relaxes A1 (injecting zeros to overcome single-pole roll-off). Finally Tier 3 relaxes A2 (eliminating, reducing, or noise-cancelling R_F). More advanced group formed by Tier 4 relaxes two assumptions simultaneously forming a compound design, while the ultimate Tier 5 bypasses all three assumptions through a different signal-processing paradigm. Orthogonal to this axis is a *configurability* dimension (AGC, variable-gain, dynamic-range extension and limiting designs) that traverses operating points without changing the underlying tier. With this we may formulate the survey as a consistent set of the design decisions (tuning knobs) applied on a common baseline.

4.4. The Architecture Map

Table 1 applies this taxonomy to the surveyed literature. Each row is an architecture family, where the three central columns mark with a check or a cross whether A1, A2 and A3 are left intact or relaxed. The last two columns give the essential mechanism and representative designs, with quantitative detail deferred to the benchmarking of Section 9. Rows are grouped into the six tiers, followed by the orthogonal configurability band whose A1/A2/A3 cells are left undetermined because a configurable cell may sit at any tier.

Several families are deliberately consolidated so the map shows mechanisms rather than instances. Shunt, series, T-coil and π-type (PIP) peaking are treated as one passive approach, since all inject zeros through inductors and differ mainly in component count and group-delay behavior. CH, active-inductor, capacitive-degeneration and parallel-TIA variants are merged into a single active zero injection family, as all synthesize a zero without an explicit spiral. The pseudo-resistor designs (sub-threshold MOSFET, multi-element, switched-resistor) form their own A2 family. The T-resistor network (see [29]) is kept in Tier 0 with the note that synthesizing a large effective R_F from smaller resistors does not relax any assumption. Finally, the Feedforward Common-Gate (FCG) input is listed as the endpoint of the CG → RGC → FCG input impedance reduction lineage in Tier 1, while the complete FCG-based receivers that add a peaking mechanism appear as compound Tier 4 entries.

 doi:10.20944/preprints202606.0753.v1

Table 1. Classification of CMOS TIA architectures by the three assumptions underlying the transimpedance limit. A ✓ mark indicates the assumption holds; ✗ indicates it is relaxed or bypassed by the topology. A1 = single-pole core amplifier; A2 = resistive feedback element R_F; A3 = full C_T at the feedback summing node.

Architecture family	A1	A2	A3	Core mechanism	Representative designs
Tier 0: *All three assumptions hold: transimpedance limit applies directly*					
Resistive SFB family	✓	✓	✓	Single or cascaded gain stage closed by a global resistive R_F; gain $\approx R_F$ with noise set by $4kT/R_F$. The T-resistor variant synthesizes a large effective R_F from smaller resistors at the cost of higher noise.	[44,46]; cascoded-inverter SFB [38]
High-impedance passive load	✓	✓	✓	Open-loop resistive (or active) load, $G = R_L$, BW $= 1/(R_L C_P)$; no feedback, hence poor noise, poor dynamic range, and no PVT stabilization.	[22,29]
Tier 1: *Assumption A3 broken: C_T decoupled from the dominant feedback pole*					
CG / RGC / FCG / enhanced-cascode family	✓	✓	✗	A low input impedance $Z_{in} \approx 1/[g_m(1+A)]$ absorbs C_D at the source node instead of at the summing node; transimpedance set by the load R_D. The CG $\rightarrow$ RGC $\rightarrow$ FCG lineage raises the loop gain A at each step (FCG additionally releasing headroom), lowering Z_{in} further.	CG [47]; RGC [11]; cascoded-booster RGC [14]; regulated inverter-cascode [25]; FCG input [30]
Input-impedance transformation	✓	✓	✗	A passive L-network, bondwire, or input T-coil absorbs C_D into a matching/filter network rather than letting it load the feedback pole.	Input T-coil with center-tap ESD diode [5]; series-peaked input + AGC [17]
Tier 2: *Assumption A1 broken: zeros injected into transfer function, single-pole roll-off overcome*					
Passive inductive peaking	✗	✓	✓	On-chip inductors add zeros and complex pole pairs to $H(f)$, extending bandwidth with no DC-power penalty; shunt, series/inter-stage, T-coil, π-type (PIP), and transformer variants differ mainly in component count and group-delay behavior.	Shunt/series [3,37,45]; π-PIP [33,41,48]; transformer-coupled [32]; inductive feedback [46]
Active zero injection	✗	✓	✓	Transistor networks synthesize an inductive impedance or inject a zero without spirals (gyrator-C active inductor, source-degenerated current source, or transconductor feedback), avoiding spiral area and substrate coupling.	[39,49]
Cascode Miller suppression	✗	✓	✓	A CG cascode holds the input transistor's drain at low impedance, suppressing Miller multiplication of C_{gd} and shifting the dominant pole up; realized passively (folded-cascode CFC/MFC, current-reuse inverter input) or with a gain-boosted active cascode.	CFC / MFC [50]; gain-boosted [51]
Tier 3: *Assumption A2 broken: R_F eliminated, reduced, or its noise cancelled*					
Capacitive feedback (CF)	✓	✗	✓	C_F replaces R_F; gain set by a capacitor ratio and load, contributing zero thermal noise (noise $\propto 1/g_m$ rather than $1/R_F$); requires an auxiliary DC-bias path.	[1,18,26,28]
Pseudo-resistor (PR) feedback	✓	✗	✓	A sub-threshold MOSFET or diode stack presents a GΩ-range feedback resistance whose shot noise $2qI_D \ll 4kT/R$ at femto-ampere bias; multi-element (MEPR) and switched-resistor variants add PVT tuning.	[52]
Active feedback / noise-cancellation	✓	✗	✓	Two partial-A2 routes: a MOSFET in triode replaces R_F (gain $\approx 1/g_{m,\text{triode}}$, gate-tunable), or a noise-cancellation path around a retained R_F cancels device channel noise (the R_F thermal noise itself is not cancelled).	MOSFET-triode [6,44,53]; NC-TIA [54]

doi:10.20944/preprints202606.0753.v1

Table 1. *Cont.*

Architecture family	A1	A2	A3	Core mechanism	Representative designs
Tier 4: *Two or more assumptions broken simultaneously: compound relaxation*					
RGC + peaking (A1 + A3)	✗	✓	✗	C_D isolated by an RGC input (A3) *and* zeros injected by passive or active peaking (A1); the dominant family for ≥20 GHz CMOS. At low headroom a two-stage CG+CS booster and active inductors replace the single CS and passive spirals.	[12,13,15,34,55,56]
Cross-coupled immittance converter (A1 + A3)	✗	✓	✗	A cross-coupled device injects a negative resistance that drives $R_{in} \to 0$ and neutralizes C_T (A3); extending the cross-coupling impedance to $R_c + sL_c$ adds a negative inductance resonating with C_T for a complex-pole pair (A1). Conditional stability requires $R_{in} > 0$ across PVT.	[57,58]
FCG input + A1 peaking (A1 + A3)	✗	✓	✗	The FCG input decouples C_D (A3); bandwidth is then extended either by passive shunt-peaking or, inductorlessly, by a capacitive-degeneration zero plus inter-stage negative-Miller neutralization (A1), the latter bounded by an input-passivity constraint.	[30,59]
Active CDF / inverter main-amp (A2 + A3)	✓	✗	✗	A common-drain active feedback stage replaces R_F (no resistor noise, A2) and simultaneously lowers Z_{in} by the RGC mechanism (A3); single-pole roll-off retained (A1 intact). Suited to low-V_{DD} nodes.	ICDF-TIA [60]
MOSFET-triode R_F + inductive zero (A1 + A2)	✗	✗	✓	A triode-MOSFET R_F (gate-tunable, partial A2) plus an injected inductor zero in the resulting third-order response (A1); feedback-path L_f placement yields lower noise than inter-stage placement.	[45,46]
CF + active BW extension (A1 + A2)	✗	✗	✓	Noise-free capacitive feedback (A2) combined with output-buffer inductance or active peaking to extend bandwidth beyond the bare CF response (A1).	[1,18]
Current-reuse CS–CG cascode (A1 + A3)	✗	✓	✗	CS and CG share bias current (no power penalty); the CG isolates C_D (A3) while the cascode interaction introduces a non-dominant zero (A1).	[61]
Current-mirror input + shunt peaking (A1 + A3)	✗	✓	✗	A diode-connected input transistor presents $1/g_m$ and absorbs C_D at the mirror node (A3); input/output shunt inductors (or active-inductor substitutes) inject zeros (A1), with resistive R_F retained (A2).	[62]
System-level compound: BW-limited SFB + input T-coil + post CTLE/VGA (A1 + A3)	✗	✓	✗	Each stage is simple but the chain breaks two assumptions: a deliberately BW-limited SFB front-end (large R_F for low noise), a passive input T-coil absorbing $C_{PD}/C_{pad}/C_{ESD}$ (A3 at the port), and off-loop CH-CTLE/VGA equalization recovering bandwidth (A1). The off-loop A1 recovery may be frequency-domain (CTLE/peaking) or time-domain (DFE). Hallmark of DSP-coupled FinFET PAM-4 receivers.	[5,6,9]
Configurability axis (orthogonal to Tiers 0–5)					
Reconfigurable gain / bandwidth	—	—	—	Single- or dual-control adjustment of R_F and/or open-loop gain A to traverse multiple gain/bandwidth or dynamic-range operating points without changing the underlying tier; includes AGC and switched-feedback limiting for wide dynamic range.	Dual-control [7]; programmable-gain CF [20]; AGC [17]; switched-FB limiting [16]; current-diversion VGA [35]

 doi:10.20944/preprints202606.0753.v1

Table 1. *Cont.*

Architecture family	A1	A2	A3	Core mechanism	Representative designs
Tier 5: *All three assumptions structurally bypassed: new signal-processing paradigm*					
Distributed / traveling-wave	✗	✗	✗	C_D absorbed into an artificial LC transmission line; gain distributed across N cells with no single C_T-loaded node; bandwidth set by the ladder cutoff. High power, since all cells are biased simultaneously.	[63]
Discrete-time / switched-capacitor	✗	✗	✗	Input current integrated onto a sampling capacitor; gain set by clock period and capacitor, noise $= kT/C$ rather than $4kT/R_F$, so the classical limit does not apply.	[29,64]
Pseudo-R integrator-differentiator	✗	✗	✗	Integrator with a GΩ pseudo-R DC path sets very high gain; a differentiator (C_F in series with R_F) restores flat bandwidth, with gain and BW set independently by capacitor ratios.	[65]
Log-domain / translinear / semi-digital	✗	✗	✗	Signal processed in the log / current domain for automatic dynamic-range compression, or a digital servo (ADC→LPF→DAC) cancels DC, avoiding both CF saturation and large-R_F noise.	[66], [67]

 doi:10.20944/preprints202606.0753.v1

5. Other Classical TIA Architectures

While SFB TIA remains the industry standard due to its simplicity and inherent wide-band potential, its performance is strictly limited by the feedback resistor R_F, which creates an unavoidable trade-off between thermal noise (high R_F) and bandwidth (low R_F) since the open-loop gain cannot be increased indefinitely without stability problems [1,68].

Several relatively old and well-known TIA topologies exist which try to address the known limitations of the SFB TIA. One of the best known examples is a common-gate (CG) TIA shown in Figure 4. This TIA is able to provide a relatively broad band, but has a fairly high input noise current due to R_D and M_2 are directly referred to the input. Moreover, as the DC voltage drop across R_D must be maximized to minimize its noise current and achieve a high gain, the allowable headroom for M_2 is also limited making its noise contribution significant. Finally, for large transistors and/or large C_D the noise contribution by M_1 also rises at high frequencies [1].

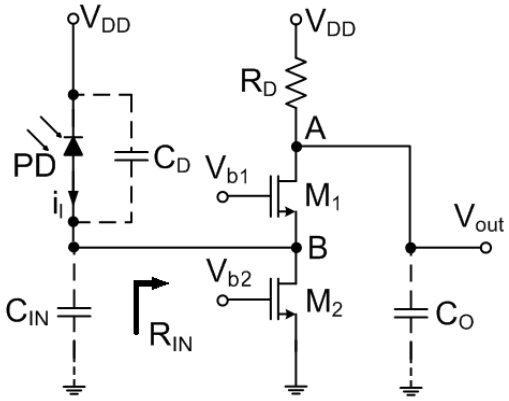

Figure 4. Basic configuration of a CG TIA.

Interesting modifications to the CG TIA were proposed in [10], which pushed C_D to a non-dominant pole so that it no longer determines the bandwidth. The authors demonstrated approximately 3.5× higher transimpedance gain at the same 1.2 GHz bandwidth compared with a classical common-source design [10], and approximately 3× compared with the conventional CG topology at 2 GHz [68].

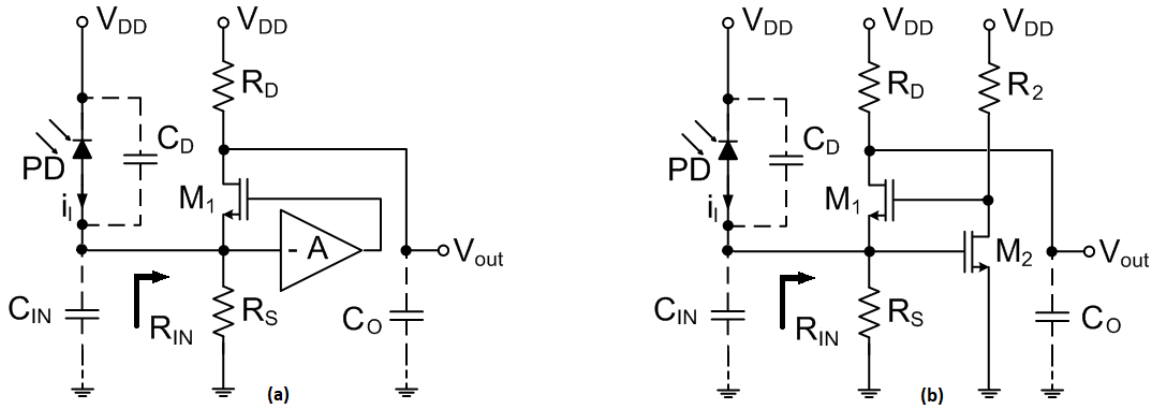

Figure 5. Regulated cascode TIA. (a) General RGC TIA concept. (b) The simplest implementation of RGC TIA with voltage amplifier as a CS stage.

The RGC topology, introduced by the same authors in the follow-up work [68] and subsequently presented as a complete TIA in [11], addresses the residual input-impedance limitation of the CG topology by lowering the input impedance through a dedicated local feedback stage (see Figure 5), effectively decoupling C_D from the input node entirely. In the RGC circuit a common-source auxiliary amplifier with gain $(1 + g_{mB}R_B)$ is wrapped around the input transistor, reducing the effective input impedance by that same factor relative to a plain CG stage. As a result, C_D sees a near virtual-ground

 doi:10.20944/preprints202606.0753.v1

termination and its contribution to the dominant closed-loop pole becomes negligible [11,68]. A further noise benefit arises from the current-gain property of the RGC stage. As the current gain of M_1 is close to unity, the current noise sources of R_1 and M_1 are referred directly back to the input node, causing the voltage noise of M_1 to become redundant and cancel in the total input-referred noise, a mechanism that lowers the equivalent input noise current below that of a conventional CG stage [11]. In the complete RGC TIA the feedback resistor R_F is still connected to the drain of M_1 rather than the output node, for two reasons: the dominant pole would otherwise be set by the large R_1 alone giving a narrow bandwidth, and connecting both ends of R_F to nodes at similar DC potentials avoids disturbing the bias [11]. While RGC enables even higher bandwidth than the CG variant, it requires additional active stages for local feedback and buffering, which increase power consumption and introduce extra poles that demand careful stability analysis in high-speed applications. Still, the RGC topology is widely used configuration for high-speed applications [12].

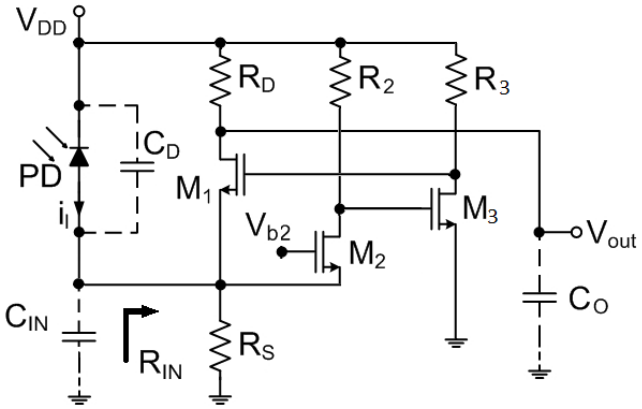

Figure 6. Basic implementation of the FCG TIA.

The natural next step in this same input impedance reduction lineage is the Feedforward Common-Gate (FCG) input stage introduced by Kromer et al. (2004). The basic configuration can be seen in Figure 6 and is a further explicit modification of the previously discussed RGC. The limitation of the plain RGC is that its CS regulating amplifier stacks on top of the input device, and the resulting series of stacked V_{DS} drops consumes the very voltage headroom that low-V_{DD} nodes cannot spare. This caps the regulating gain and hence the achievable input-impedance suppression [30]. The FCG restructures the regulation path to release this headroom while simultaneously increasing the regulation gain. The single CS regulating stage of the RGC is replaced by a two-stage regulation path consisting of a CG stage of gain $g_{m2}R_2$ followed by a common-source stage of gain $g_{m3}R_3$ so that the low-frequency input impedance becomes

$$Z_{\text{in,FCG}}(0) \;\approx\; \frac{1}{g_{m1}\left(1 + g_{m2}R_2\, g_{m3}R_3\right)}, \tag{7}$$

in direct contrast to the single-stage suppression $1/[g_{m1}(1 + g_{m2}R_2)]$ of the RGC and the unboosted $1/g_{m1}$ of the bare CG [11,68]. The three topologies therefore form a mechanistic progression in which each step attacks the same constraint of the loading of C_D on the input pole (assumption A3) with a progressively larger suppression factor. This factor is unity for CG case, the single-stage loop gain $(1 + g_{mB}R_B)$ for the RGC, and the product loop gain $(1 + g_{m2}R_2\, g_{m3}R_3)$ for the FCG. Crucially, the discriminating variable across the RGC $\rightarrow$ FCG step is not gain alone but headroom. The FCG attains its larger suppression factor precisely because its regulation path is restructured to free, rather than consume, the stacked headroom, which is what makes the input-impedance advantage realizable at the low supply voltages of deep-submicron CMOS [30].

This headroom-versus-gain framing also sharpens the relationship between the FCG and the enhanced-RGC designs discussed below. The cascoded-booster RGC of [14] and the FCG are two answers to the same question on how to raise the regulation gain beyond that of a plain single-transistor RGC booster. While the former approach stacks a cascode inside the feedback booster and

 doi:10.20944/preprints202606.0753.v1

pays for the extra gain in headroom, the FCG restructures the regulation into a feedforward path and recovers headroom in the process. Still, both remain, at the input-stage level, A3 relaxations in the sense of Table 1 (the system order and the resistive feedback element are left intact). The bandwidth-extension mechanisms that promote a complete FCG-based receiver to a compound tier (the inductorless capacitive-degeneration zero and inter-stage negative-Miller neutralization of [59] analyzed in Section 6 are layered on top of the FCG input stage rather than being intrinsic to it.

A structurally related but architecturally distinct variant is the Folded Cascode (FC) TIA, in which a CS NMOS input stage (M1) is paired with a folded CG PMOS stage (M3) in separate DC bias branches [50]. The low impedance presented by M3 at M1's drain eliminates the Miller multiplication, breaking the single-pole rolloff assumption (A1) and extending bandwidth without any passive inductive element. Critically, because the PD C_D connects to M1's gate rather than to M3's source, A3 remains intact in contrast to the RGC where local feedback absorbs C_D entirely. This places the FC-TIA in Tier 2, not Tier 1. Lima et al. [50] compared a conventional FC-TIA (CFC) with a modified variant (MFC) that replaces the NMOS CS stage and its current-source bias with a full CMOS inverter, yielding an effectively increased transconductance at the same bias current. The input pole also shifts providing a higher bandwidth at the cost of a modest increase in power and a mild increase in total input capacitance. The design's principal advantage is reduced implementation area for identical performance at the same bandwidth.

A particularly compact and instructive instantiation of the enhanced-RGC family was reported by Sadeghi et al. in [14] in 65 nm CMOS. The conventional RGC booster amplifier (a single CS transistor) is replaced by a cascoded CS stage which increases the booster gain. This also consequently lowers the RGC input impedance by the same factor. No inductor or active inductor is used and the resulting TIA contains only five transistors and consumes 760 μW at 1 V supply while delivering 60 dBΩ gain a 1.65 GHz bandwidth. The design is firmly Tier 1 of Table 1: A1 holds (the system remains second-order with the input pole dominant; the M3 cascode raises booster gain rather than injecting a zero), A2 holds, and A3 alone is relaxed via the RGC mechanism. The work is significant because it demonstrates how far the basic RGC class can be pushed by a single architectural refinement of the booster stage, without recourse to inductive peaking ([12], [13]), inverter-based feedback ([60], [15]), or active inductors ([61]). At the comparable power, the achieved transimpedance and noise are competitive with several compound Tier 4 designs, suggesting that the booster cascode is an attractive first-resort enhancement before invoking multi-assumption compound relaxations.

A somehow related route to raising the regulation gain was taken by [25] in their Regulated Inverter-Cascode TIA (RIC-TIA). Rather than cascoding the booster (as in [14]) or restructuring it into a feedforward path (as in the FCG of [30]), two local voltage amplifiers each wrap a regulating loop around one cascode device of an inverter-cascode core, multiplying its transconductance by the loop gain (A+1) and thereby raising the open-loop output resistance and lowering the input impedance. A secondary benefit is that these regulating loops replace the two external bias sources of the unregulated inverter cascode. The design is firmly Tier 1 as the resistive R_F is retained (A2 holds) and no zero is injected, so the response remains single-pole with the input pole dominant (A1 holds). Here only A3 is relaxed via the RGC mechanism. Note that the design targets frequency-domain NIRS in 130 nm CMOS and a gain-noise-power corner characteristic of the low-bandwidth biomedical-sensing regime rather than the high-speed optical cluster typical for previously discussed approaches.

Later RGC-derived designs demonstrate that the basic RGC input stage can be further compounded with active zero-injection techniques. Hosseinisharif et al. [15] replaced the standard single-transistor bias stage of the RGC with a cascoded CMOS inverter as the active feedback network. The the cascode suppresses the Miller effect of the gate-drain capacitances of the feedback transistors, effectively injecting an additional zero into the transfer function, while simultaneously boosting the feedback gain and further reducing input resistance. A second active inductor stage at the output resonates with C_{out} and moves the output pole to higher frequency. The combined results in both C_D isolation (A3) and dual active zero injection (A1) is classified in Table 1 as a Tier 4 compound design.

 doi:10.20944/preprints202606.0753.v1

The penalty is a moderately elevated input noise larger than that of a typical RGC stage since the transistors of the cascoded inverter and active inductor all refer their channel thermal noise directly to the input. Note that the active inductor in [15] is placed at the TIA output to move the output pole, whereas in [13] the active inductor replaces the load of the RGC input transistor itself, directly shaping the transfer function at the most bandwidth-critical internal node. Note also that the active inductor requires larger voltage headroom and higher supply voltage or voltage boosting techniques may be needed for advanced nodes [51].

Yet another low-voltage variant that preserves the C_D-decoupling property of the RGC while eliminating the resistive feedback entirely is the Inverter-based Common-Drain Feedback (ICDF) TIA of [60], where the CS local feedback stage of the RGC is replaced by a full CMOS inverter and the load resistors by PMOS devices, yielding a higher feedback gain A_{inv} at the same bias current and thus a lower Z_{in} than an iso-power RGC at 1.1 V supply.

Deep-submicron CMOS nodes have revitalized the interest in inverter-based TIAs. By utilizing both NMOS and PMOS devices in a complementary self-biased inverter configuration, designers can achieve higher g_m values per unit of bias current compared to traditional CS stages. This architecture is exceptionally power-efficient and provides inherent rail-to-rail output swing, making it ideal for direct driving of subsequent limiters or CDR circuits in dense SoCs. The ICDF architecture of [60] is a good example of this trend. Here a push-pull CMOS inverter serves as the main amplifying stage while a common-drain pair provides active local feedback, removing R_F from the signal path and decoupling C_D simultaneously. This design, implemented at 40 nm represents a Tier 4 with a compound relaxation of A2 and A3.

A representative example of this strategy is also the inverter-based TIA of [61], implemented in 0.18 µm CMOS for 2.5 Gb/s optical applications. The core push-pull inverter is augmented by two diode-connected transistors at the input node whose combined transconductance reduces the effective input resistance and moves the dominant PD pole to a higher frequency. This partially relaxes the A3 constraint without requiring a RGC stage. At the output, two transistors synthesize an active inductor that resonates with the load capacitance, injecting a zero into the transfer function and extending the bandwidth with a compact, area-efficient realization of the A1 relaxation that avoids the large chip area of passive spiral inductors. A resistive shunt-feedback network R_F provides DC biasing, stabilizes the operating point, and further lowers both the input and output impedance. As classified in Table 1, this compound use of an active input impedance reduction (A3) and an active zero injection (A1) with resistive feedback (A2 remains intact) positions the design firmly in Tier 4.

A distinct inverter-based refinement from the same group is the inverter-cascode TIA of [38] in 40 nm CMOS. Cascode devices are stacked in series with the push-pull input pair to raise the open-loop output resistance and voltage gain and to suppress the Miller multiplication of the input C_{gd}, while the global shunt feedback remains a physical resistor R_F. Because the cascode boosts the gain element without injecting a zero or relocating the dominant input pole, all three assumptions hold and the design is essentially Tier 0. The design is the cascoded-inverter refinement of the plain inverter-SFB cell rather than an assumption relaxation and can be seen as the direct unregulated precursor of the regulated inverter-cascode of [25]. In the latter adding the two source-regulating amplifiers of [25] around the cascode devices is precisely what promotes this Tier-0 cell to the Tier-1 (A3-relaxed) RIC-TIA discussed above.

The earliest design in this lineage by the same research group is the 0.35 µm precursor of [53], which already establishes the combined inverter and MOSFET triode feedback paradigm but lacks the inductive zero-injection mechanism that later promotes the topology to Tier 4. The design cascades three identical push-pull inverters. Here in each stage a diode-connected NMOS active load M_2 sets the stage gain as $A_{\text{OL}} = (g_{m1} + g_{m3})/g_{m2}$, replacing the high-impedance r_o-limited output of a standard CMOS inverter and reducing the Miller multiplication of the input gate-drain capacitance. A PMOS biased in triode closes the global shunt feedback eliminating the $4kTB/R_F$ thermal-noise term that an equivalent passive resistor would contribute. Within the framework of this review the design is

 doi:10.20944/preprints202606.0753.v1

firmly Tier 3 as A2 is partially relaxed by the MOSFET-triode R_F exactly as in [44], while A1 and A3 hold. Notably, the diode-connected active load does not alter the tier: a diode-connected MOSFET is purely resistive at the band of interest and injects no zero, in clear contrast with the active inductor loads of [13,39,49,61] which do break A1. The active load is a stage-level area- and noise-trade-off and is orthogonal to the assumption-relaxation paradigm.

Later the work [45] carries the same inverter and MOSFET triode feedback philosophy to 0.18 µm CMOS but reworks three stage-level details in the process. The diode-connected load is removed, the cascade is extended to five stages, and the feedback MOSFET is switched from PMOS to NMOS. Furthermore, the authors add a single series inductor between the fourth and fifth gain stages. The inter-stage L creates a third-order LC transfer function that pushes the -3 dB bandwidth from 7.77 GHz (no inductor) to 9.28 GHz (1.19× extension), simultaneously relaxing A1 (zero injection) and A2 (MOSFET-triode R_F) and placing the design in Tier 4 alongside [46].

The follow-on [46] compounds the same A1 + A2 relaxation but in a *single-stage* push-pull inverter, with the inductor L_F placed in series with the MOSFET R_F inside the feedback branch rather than between gain stages. The combined effect of dropping the cascade (which lowers the baseline input-referred noise from 27.76 pA/$\sqrt{\text{Hz}}$ to 11.4 pA/$\sqrt{\text{Hz}}$ before any inductor is added) and moving L into the feedback loop (which directly suppresses the in-band R_F-noise contribution, whereas the inter-stage placement of [45] actively raises the noise floor from 27.76 to 36.23 pA/$\sqrt{\text{Hz}}$) reduces the final input-referred noise to 7.7 pA/$\sqrt{\text{Hz}}$ at 7.9 GHz, at the cost of a modestly lower bandwidth. Together with the 2010 precursor of [53], the three works form a clean Tier 3 → Tier 4 transition in which one additional A1 mechanism (inter-stage L, then feedback-branch L) is layered onto a held-fixed MOSFET triode A2 relaxation. Note that the across-paper architectural difference is, however, broader than the inductor placement alone and also involves the per-stage load and the stage count.

A compound Tier 4 example combining three simultaneous A1-breaking mechanisms with RGC was reported in [34]. The core TIA employs a RGC input and two series inductors L1 and L2 absorbed into an LC input-matching ladder (passive A1). A subsequent capacitive-degeneration stage injects a zero to cancel the remaining dominant pole thus extending the bandwidth to the next highest pole (active A1). A parallel shunt-feedback AGC loop (a PMOS transistor controlled by a peak detector) is disabled in the weak-input regime and engages progressively under overload, reducing the transimpedance from 490 Ω to 170 Ω to prevent saturation and timing jitter. The same peak-detector controlled feedback transistor AGC principle underlies the wide dynamic range fiber-sensing TIA of [17], where it is deployed for sensitivity range extension (123.5,dB DR) rather than jitter limited overload protection. This itself is a nice illustration that the identical control loops may serve different binding specifications across application domains.

Moving away from traditional continuous-time amplification, switched-capacitor (SC) and discrete-time TIAs represent a paradigm shift for low-power sensing. By using an active integrator and periodic reset, these TIAs effectively integrate the photocurrent over a clock cycle, providing superior noise filtering and robust rejection of low-frequency interference. This approach is increasingly popular in LiDAR and OTDR systems where signal processing occurs in the digital domain. For ultra-short-reach links (e.g., chip-to-chip optical interconnects), latency and power efficiency are paramount. Current-mode TIAs bypass the traditional transimpedance conversion stage, instead amplifying the photocurrent directly in the current domain. By eliminating the high-impedance node, these architectures minimize the impact of parasitic capacitance and allow for significantly higher bandwidths at the cost of higher noise floor as a compromise often acceptable for short-reach applications with high optical power budgets.

The performance of a conventional lumped-element designed TIA is limited by the characteristics of the active devices they use in simple feedback circuit [63]. As data rates exceed 100 Gbps per lane, the physical dimensions of the TIA become comparable to the signal wavelength. Distributed and traveling-wave TIAs solve this by treating the amplifier as a transmission line segment. By distributing the input and output transconductance stages along a synthetic transmission line, these designs are in

 doi:10.20944/preprints202606.0753.v1

principle able to achieve bandwidths limited only by the technology's f_T, making them the preferred solution for the 100 GHz and above regime in advanced CMOS nodes.

One of the earlier works on the distributed TIA paradigm in standard CMOS was reported by [63]. Two chips were fabricated in a 0.18 µm CMOS process. The first cascades two identical three-stage cascode distributed amplifiers to overcome the inherently low gain of a single distributed stage. Each stage employs cascode gain cells (CS and CG pairs) with a 20 Ω damping resistor at the gate of the CG transistor to suppress the high-Q resonance introduced by the drain-source capacitance of that device, which would otherwise make the cascode cell difficult to stabilize. The second chip is a single-stage cascode distributed amplifier with only two gain cells, each consisting of a cascode device, a 10 Ω damping resistor, and a bypass capacitor. The amplifier is DC-coupled at both input and output. This topology trades gain against power and die area, achieving 48 dBΩ transimpedance and a 30 GHz bandwidth at 50 mW, at the time of publication the largest bandwidth reported for any CMOS TIA. In both designs, on-chip spiral inductors are enclosed by a CPW ground shield to reduce mutual coupling and suppress substrate-noise injection into the LC ladder. Both designs are assigned to Tier 5 of Table 1 as all three classical assumptions are bypassed simultaneously. A1 is broken because the distributed ladder absorbs the cascode input and output capacitances, and the gain-bandwidth product exceeds f_T. A2 is broken because there is no resistive shunt-feedback path; the noise floor is set by the channel noise of the cascode cells, not by a thermal feedback resistor. A3 is broken because the C_D is absorbed into the LC input transmission line and does not load a single summing node. The penalty, consistent with the distributed paradigm, is low power efficiency as all gain stages must be biased simultaneously and the distributed topology offers no feedback mechanism to stabilize the DC operating point against possible PVT drift.

In massive multi-channel arrays, crosstalk and power supply noise are critical bottlenecks. There some benefits and some disadvantages of the single-ended designs. Such architecture reduces power consumption up to half and thermal noise by up to a factor of $\sqrt{2}$ when compared to replica-based TIA architectures. Compared to other differential variants, it also saves a significant active silicon area and also gives a big relief from the design overhead dealing with mismatches in the amplitude and the phase errors in differential paths, especially when the targeted symbol rate has a UI < 20 ps [5]. However, single-ended architectures are more sensitive to power supply noise, which can elevate the power supply induced jitter (PSIJ). Moreover, they are highly susceptible to substrate noise and common-mode interference. Conversely, differential architectures provide excellent rejection of PSIJ and crosstalk, which is essential for dense CPO applications. However, they require twice the number of components and necessitate precise layout symmetry to prevent signal skew and phase mismatches in high-speed data streams. The motivation above is framed by high-speed optical arrays, but a structurally different driver for the differential form arises in resonant-sensor readout. When the TIA output is fed back to a MEMS resonator to sustain closed-loop oscillation, the resonator's parasitic feedthrough capacitance appears in parallel with the TIA transimpedance and erodes its effective bandwidth. Here implementing both the resonator and the TIA differentially cancels the feedthrough of the positive and negative branches and recovers the bandwidth a single-ended realization would lose [28]. The same application class constrains the TIA phase rather than only its magnitude as a sustained oscillator requires the front-end to contribute a phase near 0° across the operating band so that it does not perturb the loop phase condition, which forces the −3 dB bandwidth to lie well above the oscillation frequency [29]. The MEMS front-end of [28] holds its phase within roughly ±10° from a few kilohertz to 200 kHz by placing its 1.8 MHz bandwidth about an order of magnitude above the intended closed-loop oscillation range [28]. This is the low-frequency, phase-driven counterpart of the group-delay-flatness requirement that governs high-speed links, and it reinforces that the binding TIA specification is set by the application rather than by the topology alone.

Two concrete pre-FinFET realizations of the fully-differential array approach in 0.18 μm CMOS illustrate the simple-CG → ACG mechanism upgrade at a fixed process node. [47] reports a 4-channel array for DVI / HDMI cables built from a basic CG input stage (no local feedback loop), followed

 doi:10.20944/preprints202606.0753.v1

by two CS gain stages with global resistive R_F and a passive DC-balance output buffer realized by low-pass filtered feedback. The work [8] upgrades the input stage to an ACG (functionally a Park & Toumazou RGC with local-feedback gain boosting the input-impedance suppression) and replaces the passive DC-balance with an active f_T-doubler-based DC-cancellation auxiliary path so that the servo loop does not collapse the signal-path bandwidth. An active-feedback gain stage with negative-capacitance compensation is added downstream as a limiting amplifier. Both TIA cells belong to Tier 1 of Table 1: the ACG mechanism of [8] is structurally the same A3 relaxation as the simple CG of [47], only with the input impedance further suppressed by the local feedback gain, which is what allows the gain headroom and bandwidth uplift between the two designs at fixed process and fixed pseudo-differential pattern. In both, common-mode balance is obtained by tying the inverting input to an off-chip capacitor C_x matched to C_D, retaining the PSIJ / CMRR benefits of a true differential path without the area cost of a twin-PD scheme but at the cost of one external component per channel. The per-channel power overheads of 50-72 mW from 1.8 V typical of these dense differential CG arrays motivated the migration to inductor-less, lower-headroom variants ([14], [15]) in subsequent sub-100 nm generations. The downstream active-feedback LA of [8] is a post-TIA bandwidth-extension mechanism that does not affect the TIA tier, but the same active-feedback principle recurs in the post-equalizer chains of modern PAM-4 receivers such as [5].

6. Advanced Bandwidth and Noise Optimization

Modern high-speed TIA design requires techniques that go beyond simple shunt-feedback, moving toward sophisticated passive and active networks to manipulate the transfer function and suppress noise.

6.1. Passive Peaking Networks

Inductive peaking is probably the most common used technique for the bandwidth extension where the inductors are placed in strategic circuit locations, resulting in resonance phenomena with parasitic capacitances [37,50]. A basic discussion on passive devices including inductors in modern CMOS can be found in [3]. However, the usage of inductors to extend the amplifier bandwidth results in several obvious drawbacks such as increased chip area, which makes them inappropriate for high density designs, increase in substrate coupling, resulting in higher crosstalk and challenges in efficient inductor design as significant efforts may be needed on accurate modeling and parasitics extraction. Finally, design kits also provide fewer number of inductance values. Nevertheless, passive peaking networks, particularly T-coils, have become the cornerstone of bandwidth enhancement in commercial optical receiver ICs. Unlike simple shunt peaking, which often introduces significant GDV due to the sharp phase shift near the cutoff frequency, the T-coil network acts as a bridge between the input capacitance and the load. A rigorous fair comparison on a common basis shows that T-coil peaking achieves a realistic Bandwidth Extension Ratio (BWER) of 2.40 (versus the commonly cited 2.82 derived with C_D neglected) and shunt-series peaking achieves 1.83 (versus 3.46 derived under the impractical assumptions) [33]. These figures apply to single-stage or two-stage configurations and the extension factor grows substantially in multi-stage designs with inductors at every inter-stage node, reaching 3× for a 5-stage cascade [37]. While LCL-T networks provide similar benefits, T-coils offer superior area efficiency, making them the preferred solution for high-density monolithic integration. A particularly attractive property of passive inductive shunt peaking, shared by none of the active alternatives, is that the bandwidth enhancement comes without any additional power dissipation [39].

However, this advantage is conditional on the physical feasibility of the required inductance. An important practical boundary is established by the inductance requirement as discussed in [39]. For example, for $R_L = 1\,\text{k}\Omega$ and $C_{\text{eff}} = 0.2\,\text{pF}$, the optimum-group-delay factor $m = 0.32$ demands $L = 64\,\text{nH}$, which is impractical on chip. Passive shunt peaking therefore becomes ineffective for designs with high load resistance at frequencies below approximately 5 GHz, and is ruled out entirely in ultra-compact or dense VLSI integration contexts where area is the binding constraint [13]. Here the both motivations drive the active inductance approaches discussed below.

 doi:10.20944/preprints202606.0753.v1

Numerous works have been reported on TIAs with passive peaking networks. For example, [69] reported a fully-differential CG TIA with on-chip spiral shunt peaking inductors. Note that this design also partially relaxes A3, but since there is no regulated cascode feedback loop of the Park/Toumazou type, the A3 relaxation due to the CG stage is structural rather than an engineered compound addition. The most aggressive passive peaking factor reported for a single CS stage was achieved by [33,41] via the proposed π-type inductor peaking (PIP) technique. Three inductors per stage (shunting the input node, in series between stages, and shunting the output node) produce a 5th-order transfer function with two real zeros and two complex conjugate pole pairs. The design achieves a BWER of 3.31$\times$ with a gain variation within 2.0 dB, but at the cost of a GDV exceeding 45 ps above 24 GHz, a direct consequence of the complex-pole response optimized for bandwidth at the expense of phase linearity [33]. Eye diagram measurements at 40 Gb/s confirm functional operation but show broadened edge transitions attributable exactly to this GDV effect.

A complementary passive peaking strategy, targeting multi-stage inverter-configuration TIAs specifically, was introduced by [37]. Rather than concentrating multiple inductors within a single stage, inductors of equal value L_T are placed between every pair of gain stages, so that each inter-stage node forms a third-order LC-ladder section with the combined parasitic capacitances of the adjacent stages. The key system-level consequence is that the bandwidth degradation normally incurred when cascading stages (which reduces the 3-dB frequency of a conventional 5-stage inverter TIA to 2.4 GHz) is almost entirely prevented. The same 5-stage cascade with full inter-stage peaking achieves 7.4 GHz in simulation and 7.2 GHz in measurement, a 3$\times$ improvement over the unpeaked baseline at no additional power dissipation. To prevent the resulting LC ladder from degrading at the input and output terminations due to impedance mismatch, M-derived half-circuits (with series-L_z and shunt-C_z branches) are employed at both ports. These additionally absorb PD capacitance into the impedance transformation network rather than leaving it as the dominant loading capacitance at the summing node. A practically important property reported in the same work is the insensitivity of this topology to inductor quality factor as a 50% reduction in Q degrades gain by only 2 dB and reduces bandwidth by 3%, in contrast to shunt-peaking designs where spiral inductor stray capacitance directly adds to the node capacitance and collapses the peaking effect.

A more recent passive-peaking variant replaces the separate series- and shunt-peaking spirals with a transformer (a pair of positively-coupled inductors) in the forward path [32] in a 16 nm FinFET inverter-based TIA. Two such transformers act as analog passive equalizers that inject high-frequency boost, relaxing the transimpedance single-pole limit (A1) at no static-power cost, while the resistive R_F is deliberately kept out of parallel with the inverter C_{gd} so the feedback impedance is not degraded at high frequency. Stacking the two coils as a transformer rather than placing two independent spirals raises the effective quality factor and hence the achievable boost, which lets the design recover $\sim$57% of the silicon area of an equivalent combined series and shunt-peaked implementation while delivering $\sim$1.36$\times$ the bandwidth at the same R_F [32]. Crucially for PAM-4 signaling, the frequency-independent feedback factor yields a lower in-band GDV ($\sim$3 ps), addressing exactly the phase-linearity penalty that limits aggressive inductive peaking (Section 7). The same work motivates this choice by contrasting it with negative-capacitance bandwidth generation, which extends bandwidth but at a much higher noise penalty [32], the same trade-off discussed for negative-Miller neutralization in Section 6. Notably, the reported 58 dBΩ/17.4 GHz front-end supports 64 Gb/s PAM-4 *without* any explicit equalizer, with a 100 Gb/s eye recovered once a 5-tap FFE is added, illustrating how in-cell passive equalization can defer the power and complexity of an off-loop equalizer chain.

A structurally distinct inductive peaking variant places the inductor the feedback path, in series with R_F, rather than at the signal input or the output node. The feedback impedance $Z_F(s) = R_F + sL$ adds one zero and one pole to the transfer function relative to the resistive-only topology, elevating the system to third order [46]. Beyond bandwidth extension, the rising $|Z_F|$ at high frequency partially suppresses the feedback resistor noise contribution in the same band where the input-referred noise current would otherwise peak. This forms a dual role shared with the L_F inductor in the CH design

 doi:10.20944/preprints202606.0753.v1

of [12] and results in reducing the cut-off noise from 11.4 pA/$\sqrt{\text{Hz}}$ (resistive-only) to 7.8 pA/$\sqrt{\text{Hz}}$ at 7.9 GHz.

6.2. *Active Feedback & Zero-Pole Management*

To circumvent the transimpedance limit while avoiding the area penalty of on-chip inductors, active feedback and zero-pole management have emerged as critical design paths. These techniques involve the use of active loads or secondary feedback loops to introduce active zeros in the transfer function, which cancel out the dominant poles caused by the PD and transistor input capacitance. By carefully placing these zeros, designers can achieve a flat frequency response and extend the bandwidth significantly. These architectures are particularly suited for processes where high-quality passive inductors are unavailable or prohibitively expensive, effectively providing a digital equivalent of software-defined bandwidth control.

One of the most used approach is to add a so-called active inductor. This is a circuit that contains at least one active element and does not contain a physical inductor. This circuit is characterized by an impedance that increases with frequency within a specific frequency range. An active inductor circuit addresses the limitations associated with passive inductors and serves as their replacement for extending the bandwidth. Obviously, there can be different active circuits which approximate the inductor and even if they are designed for the same equivalent inductance value, due to the non-idealities their impact on the TIA performance can be very different (see comparison of true passive and two different active inductors in [62]). Apart of an obvious impact on the bandwidth and gain, active inductors may also significantly affect the THD to be changing with the magnitude of the input signal current and may negatively affect the linearity of the system. Unfortunately, this level of analysis can be seldom found in concise works on recent developments in CMOS TIAs.

One of the earliest measured CMOS demonstration of active inductance synthesis for TIA bandwidth extension was reported by [39]. A current-source transistor with inductive source degeneration L_S presents an output impedance, amplifying a physically small 15.3 nH spiral to synthesise the large effective inductance that passive shunt peaking would require but cannot practically implement at 1-2 GHz. Applied as the load of a cascode TIA with global resistive feedback (A2 intact, A3 intact), the active load extends the bandwidth by more than 90 % relative to a conventional shunt-peaked cascode, where the same 15-nH passive inductor produces zero bandwidth improvement . A further benefit, structurally identical to the dual role of L_F in [12] and L_s in [33], is that the active load suppresses load-transistor noise at high frequency and provides noise reduction without reducing the DC bias current of the input stage. Area efficiency is improved further by using 3-level 2-turn multilevel spiral inductors occupying one-third the chip area of single-level equivalents at the same L and Q.

A structurally simpler active-L implementation was reported by [49] targeting Plastic Optical Fibre (POF) receivers with large C_D. An NMOS transistor with series resistor R and shunt capacitor C forms a gyrator-C active inductor as the CS drain load, synthesizing impedance with zero and pole. The zero compensates the dominant input pole, permitting a 6 dB larger R_F at the same 2 GHz bandwidth relative to a baseline CS-TIA, which reduces the integrated input-referred noise current by 15 %. Note that no passive spiral is used and the TIA core occupies only $10\times35\,\mu\text{m}^2$.

A gain-boosting approach to the same Tier 2 (A1) classification was reported by [51] in the same 130 nm node but targeting a different corner of the design space: $C_\text{D} = 200$ fF (InGaAs PD) rather than the 2 pF of [49]. An auxiliary two-stage CS amplifier drives the gate of cascode transistor, boosting its effective transconductance and increasing the cascode output resistance. The higher open-loop gain raises the closed-loop bandwidth while simultaneously suppressing Miller multiplication of by the same factor. This mechanism is structurally identical to the FC TIA, implemented via active gate drive rather than a passive CG topology. The auxiliary amplifier adds 5% of total power and extends bandwidth from 4.8 to 7 GHz (31.4%), while leaving noise nearly unchanged. This is the first reported application of the OPAMP gain-boosting technique to TIA design.

A structurally distinct active-inductance implementation is the Hara-type active inductor of [13], where the load resistor of the RGC input transistor is replaced by transistor biased such that its

 doi:10.20944/preprints202606.0753.v1

gate-source and gate-drain capacitances, together with the bias resistor synthesize an impedance with an inductive reactance. This introduces a zero/complex-pole pair into the transfer function without any passive spiral. This contrasts with the active inductance of [39], which synthesizes the inductive effect via source degeneration of a current-source transistor. The Hara approach is better suited to deep-submicron processes where the low supply voltage ($V_{DD} = 1$ V in 28 nm) prevents the stacking required by the degenerated current-source topology. A second A1-breaking mechanism in the same design is a CH resistor inserted in the two-stage A_{RGC} loop: it trades loop gain for loop bandwidth, preventing the large parasitic capacitance collapsing the RGC feedback before it can suppress the input impedance. The compound result with RGC (A3) and dual active zero injection (A1) (no passive inductor) results the smallest reported footprint in the $\geq$20 GHz CMOS TIA class at the time of publication.

The CH mechanism, introduced as an intra-loop A1-relaxation technique in [12] and [13], has also migrated into a structurally different role in modern DSP-coupled FinFET receivers with digitally-programmable post-TIA CTLE. In [5] (16 nm FinFET, 32 GHz, 0.9 V) the Stage-2 and Stage-3 equalizers are CH transconductors with a parallel CR-based high-pass branch whose cutoff is digitally tuned through two resistor banks. Together they recover the bandwidth that the deliberately low-BW Stage-1 SFB front-end gives up to noise. The architectural significance is that the same A1-relaxation mechanism (active zero injection by a CH transconductor) is now used outside the TIA feedback loop rather than within it, so that the CH stage's own thermal noise is suppressed in the input-referred budget instead of adding directly to the TIA noise floor. This off-loop usage of the CH mechanism is the structural enabler of the system-level compound Tier 4 (A1 + A3) row introduced in Table 1.

For example, [15] used two mechanisms simultaneously in addition to RGC input: Miller-cancelling cascode in the feedback network and an active inductive peaking at the output. A similar compound approach but at significantly lower power was demonstrated by [61]), where an active inductor at the TIA output provides inductive peaking without any passive spiral. At the same time the diode-connected input transistors simultaneously reduce the input impedance and shift the PD pole. The same RGC and active-inductor combination, but with the input-impedance reduction supplied specifically by a Sadeghi-style cascoded booster rather than a Miller-canceling feedback inverter, was reported by [55] in 90 nm CMOS. Here the booster cascode purely lowers the RGC input resistance (A3, no zero injected, as in [14]) while a single active inductor output stage supplies the A1 zero, reaching 50.5 dBΩ/7.3 GHz at only 1 mW in simulation. A 65 nm variant from the same group embeds an equivalent active-RGC / active-inductor cell ahead of a replica TIA and a three-stage LA to form a complete sub-5 mW 5 Gb/s receiver with the TIA core remaining the same A1 + A3 mechanism [56].

A recent simulation-only design in the same A1 + A3 sub-family is the current-mirror TIA of [62]. The diode-connected mirror input M_2 achieves A3 isolation identically to a CG stage, while shunt inductors placed at both input and output extend the bandwidth to 2.8 GHz (A1). Three inductor implementations are compared: a passive spiral and two active-inductor substitutes based on a PMOS gyrator-C and an SSF circuit, respectively. The passive variant achieves the lowest noise and widest bandwidth, whereas both active-inductor variants degrade THD due to the additional nonlinear elements. The principal design constraint is that the small $R_F = 10\,\Omega$ required to achieve low Z_{IN} limits the closed-loop transimpedance to 30.6 dBΩ, substantially below other Tier 4 designs, confirming that the input-isolation benefit of the current-mirror topology is traded against transimpedance gain at this C_D.

A particularly efficient compound design in the same Tier 4 family was demonstrated by [12], where the conventional RGC is augmented by a local CH feedback consisting of a resistor and series inductor connecting the drain to the gate of M1. The resistive feedback converts the two real poles of the RGC into a complex conjugate pair, extending bandwidth, while the inductor simultaneously resonates with parasitic capacitance and introduces a zero for further peaking (A1). A critical and under-appreciated property of this topology is that both inductors serve a dual role. While one makes the thermal noise of the feedback resistor frequency-dependent, suppressing its contribution at high

 doi:10.20944/preprints202606.0753.v1

frequencies, the series inductor at the input extends the bandwidth and reduces the input-referred noise.

A simpler Tier 4 compound design that forgoes the RGC input stage was demonstrated in [46]. Here a CMOS push-pull inverter retains the full C_D at its input (A3 intact) but replaces the feedback resistor with a MOSFET biased in triode (A2 partially relaxed) and places an inductor in series within that feedback path (A1 relaxed via zero injection). Note that here the A2 relaxation is modest as the MOSFET channel noise nearly matches the thermal noise of the equivalent resistor, and the noise improvement relative to the $R_F + L$ confirms that the principal benefit of this topology is bandwidth extension through inductive feedback and not the noise elimination.

A compound Tier 4 (A2 + A3) design targeting the same 40 nm CMOS process was demonstrated by [60] at 10 Gb/s. The ICDF-TIA replaces both the resistive load and the feedback resistor with PMOS transistors, so that the common-drain (CD) source-follower pair acts as the feedback element. The resulting input impedance follows the same RGC form but with the inverter gain replacing the CS stage gain, which is gain-limited by headroom at 1.1 V. At matched input impedance this reduces bias current by 28 % relative to RGC.

A useful remark at this point is that the Tier 4 compound designs mentioned here pack the two assumption relaxations into a single TIA cell, accepting a more complex active feedback network in exchange for a mechanically compact stage. The 16,nm Patel design [5] illustrates the structurally opposite strategy: a Tier-0 SFB inverter at the input, a passive T-coil at the input port carrying out A3 relaxation, and a cascaded CH CTLE/ Variable-Gain Amplifier (VGA) chain carrying out A1 relaxation outside the feedback loop. Both strategies converge on the same Tier 4 row in Table 1 but operate differently. The intra-loop compounds minimize stage count and area at the cost of feedback complexity and PVT sensitivity, while system-level compounds preserve a simple TIA cell at the cost of an explicit equalizer and VGA chain. Which trade-off dominates depends primarily on whether DSP-grade calibration is already mandatory in the receiver as it is for 100,Gb/s-class PAM-4 links targeting CPO deployment.

6.3. Noise-Canceling Techniques

In the most popular SFB TIA architecture current noise of the resistor is directly added to the input-referred current noise of the TIA which dramatically degrades the overall noise performance. A solution to this problem are noise-canceling TIAs which utilize feed-forward or bridge-based architectures to subtract the noise contribution of the main input transistor. An example of such architecture is the capacitive feedback TIA initially proposed in [1] and is shown in Figure 7. Here the approach to reduce the noise of the circuit is to replace the noisy feedback resistor with a noise-free circuit such as a capacitor or their combination. Such a circuit would eliminate the disadvantages of typical resistive feedback while maintaining the advantages due to the circuit's feedback structure such as almost constant transimpedance gain in the bandwidth of interest and reduced sensitivity to process and temperature [43]. The capacitive feedback architecture is formed by a dedicated network of two capacitors C_1 and C_2. The mid-frequency value of the transimpedance gain R_T for this generic architecture can be, in theory, approximated as:

$$R_T = \left(1 + \frac{C_1}{C_2}\right) R_2. \tag{8}$$

This expression for the transimpedance gain is valid under the assumption of an infinite gain A and the overall gain of this TIA circuit is defined by the feedback circuit formed solely by the two feedback capacitors C_1 and C_2. Although it can be used for a very rough analysis of this promising architecture, it may not be extremely helpful for predicting the performance of a real circuit in CMOS and, therefore, more accurate models are needed [70].

 doi:10.20944/preprints202606.0753.v1

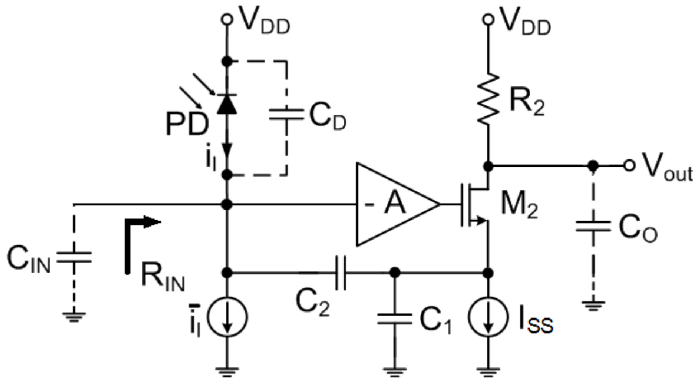

Figure 7. Basic configuration of a capacitive feedback TIA.

The original approach of Razavi was later refined in the series of works of Shahdoost Shahdoost et al. (2014, 2016) who introduced DC dark current elimination circuit, multi-stage feed-forward paths to achieve ultra-low input-referred noise densities while maintaining the gain required for high-sensitivity optical links. Further improvements have been introduced by [20] including an improved PMOS-based area-efficient biasing circuit, a detailed noise model that resolves discrepancies in earlier closed-form expressions, and a refined accurate transfer-function model that explicitly captured both bias resistors and the inter-stage coupling capacitor C_C [73]. The programmable-gain extension of the same methodology is discussed later as an instance of the dual-control configurability class.

A complementary CF-TIA extension targeting a substantially different application class is reported by [26] for MRI receiver-coil arrays. The design retains the noise-free capacitive feedback network of [1] but introduces two CF-specific refinements. First, the buffer is realized as a complementary push-pull stage, and the resistive load R_D at the buffer output is eliminated because the downstream stage is a current-mode mixer sensing the output current directly. The elimination of R_D removes the dominant residual thermal noise term that, R_F has been replaced by the capacitances, would otherwise set the noise floor in conventional CF designs. Unfortunately, the push-pull buffer stage increases the capacitive loading of the voltage amplifier and this additional loading may be undesirable for high speed optical communication applications while being acceptable for MR systems with the operation frequencies below 150 MHz. This example illustrates that the same Tier 3 mechanism (A2 broken with capacitive feedback) supports application requirements very different from those of optical receivers, and that within the CF family there is substantial room for buffer-stage and bias-stage refinement that does not alter the assumption-relaxation pattern.

In [21] as a voltage amplifier a Miller two-stage topology was proposed. According to the authors this design improves both the gain and noise rejection when compared to reference designs at a price of an increased design complexity. As a disadvantage of the capacitive feedback design one can mention its relative complexity and that the noise and the bandwidth is affected by the parasitics of the P_D. A further instance of the same Tier 3 mechanism in a non-optical domain is the differential MEMS-resonator front-end of [28]. The design casts the Razavi capacitive-feedback network as an explicit current amplifier. The capacitor-ratio current gain reaches an on-chip 56 MΩ at 65 fA/$\sqrt{\text{Hz}}$ and 436 μW in 0.18 μm CMOS thus achieving the gain impossible to realize with a physical feedback resistor, at a bandwidth (1.8 MHz) matched to resonant-sensor rather than optical use. Because the current gain $(1 + C_2/C_1)$ multiplies the load resistor while the load-resistor noise is divided by the very same factor when referred to the input, and because the $R_D C_L$ pole sits outside the feedback loop (so that stability is not governed by R_D, provided R_D stays below the drain-source resistance of M_1), the gain, noise, and bandwidth of this topology are set by a largely decoupled set of variables, in contrast to the resistive SFB TIA where R_F alone fixes gain, noise, bandwidth, and loop stability at once [28]. This is the concrete example in which breaking A2 opens a distinct trade-off space rather than merely deleting a noise term.

 doi:10.20944/preprints202606.0753.v1

While capacitive feedback (Tier 3) eliminates resistor noise by breaking (A2), an alternative approach focuses on canceling the noise of the active devices within a Tier 1 multi-path framework. Atef et al. in [54] demonstrated a 2.5 Gb/s TIA with noise cancellation that maintains the resistive SFB structure but introduces a feedforward path to subtract the thermal noise of the input CMOS inverter. In this architecture, the input device noise appears in-phase at the inverter's input and output, whereas the signal is inverted. A subsequent summation stage (acting as an analog subtractor) effectively removes the transistor noise contribution.

Taken together, selected works fom Atef's span Tiers 0-4: the inverter-cascode of [38] is a Tier-0 cascoded-inverter SFB cell, promoted to the Tier-1 regulated inverter-cascode of [25] (A3 relaxed, NIRS front-end) by adding the source-regulating amplifiers, while the remaining works target large- or moderate-capacitance optical front-ends. [49] extends bandwidth via gyrator-C active-L zero injection (A1, Tier 2; $C_T = 2\,\text{pF}$, POF target), [51] extends the bandwidth via gain-boosted cascode Miller suppression (A1, Tier 2; $C_{\text{PD}} = 200\,\text{fF}$, optical-fiber target) while both in 130 nm but addressing opposite PD capacitance regimes with orthogonal Tier 2 mechanisms. Finally [54] partially reduces device noise via NC cancellation around a retained R_F (A2 partial, Tier 3; 40 nm) and [60] simultaneously decouples C_D and eliminates resistive R_F via the ICDF mechanism (A2 + A3, Tier 4; 40 nm). The progression illustrates how a single application constraint of interfacing with a large or non-negligible PD capacitance, motivates distinct assumption relaxations depending on whether the binding bottleneck is bandwidth (A1), noise floor (A2), or input-pole dominance (A3), and that within a single tier multiple mechanistically orthogonal solutions may coexist.

6.4. Input Capacitance Neutralization

The canonical silicon-proven realization of this approach is the cross-coupled immittance-converter (CIC) TIA of [57]. A cross-coupled transistor M_2 connected between the drain of the main CG device M_1 and the cross-coupling impedance Z_C injects a negative resistance term into the input node, driving $R_{IN} \approx 0$. This decouples $C_T = C_{PD} + C_{IN}$ from the dominant input pole without requiring a RGC feedback loop, breaking assumption A3. When Z_C is extended to $R_C + sL_C$, the same node acquires a negative inductance L_{IN} that resonates with C_T, forming a complex-conjugate pole pair and breaking assumption A1. The topology therefore falls in Tier 4 of Table 1. The critical PVT constraint is that R_{IN} must remain positive across all process corners and it sets a lower bound on L_C by directly coupling the achievable stability margin to the transimpedance gain. [58] carry the same immittance-converter principle into a low-power RGC-based modified-MIC realization. Their refinement is to place an on-chip inductor in series with the cross-coupling resistor, $Z_C = R_C + sL_C$. The negative term adds to L_{IN} and cancels the source-degeneration penalty that the RGC device otherwise imposes. The R_{IN} reduction (A3) and the L_{IN} reduction no longer trade against one another. This decoupling breaks the gain-stability trade-off without the elevated bias current needed to hold R_{IN} small by transconductance alone. The authors also showed that placing L_C in series with R_C outperforms conventional shunt peaking as the former moves the dominant input pole rather than merely peaking the load.

A capacitive rather than immittance-based route to the same input-pole relief was taken by [59] in 65 nm CMOS. The input pole is first suppressed by a FCG stage, whose regulation loop lowers the input impedance below that of a plain RGC and thereby relaxes A3. The bandwidth is then extended without any inductor by injecting a zero through a capacitive degeneration CS gain stage and by neutralizing the internal inter-stage parasitics with negative-Miller capacitors (both being A1 mechanisms). The design classifies as a Tier 4 (A1 + A3) compound in Table 1. Because the bandwidth extension is purely capacitive, the active area is roughly an order of magnitude smaller than the inductively-peaked 65 nm designs at comparable bandwidth [33] at the cost of a higher input-referred noise since the degeneration and neutralization devices refer their channel noise toward the input. The price of the neutralization, as discussed in Section 7.2, is an input-passivity constraint that bounds how much capacitance may be canceled before the stage sarts to oscillate.

 doi:10.20944/preprints202606.0753.v1

7. Stability Across the Assumption-Relaxation Taxonomy

Every CMOS TIA surveyed in this review is obviously reported as a stable circuit. Each design was demonstrated either in silicon or in simulation with a well-behaved frequency and step response. Stability is, therefore, not a property that distinguishes a working design from a broken one, but a margin that some architectures preserve more comfortably, and more robustly against PVT drift, than others. The relevant question is consequently not whether a given topology is stable, but where its stability margin comes from, how much slack it leaves, and which failure mode appears first when the design is pushed toward higher gain or wider bandwidth. This section attempts to answer those questions within the same $\{A1, A2, A3\}$ vocabulary used throughout the review, and argues that stability is not an additional, orthogonal concern bolted onto the transimpedance limit but is encoded in the very same second-order constraint from which the limit is derived. The central observation is that relaxing any one of the three assumptions does not remove the stability problem, but rather shifts it into the mechanism that performs the relaxation.

7.1. The Transimpedance Limit Is a Stability Boundary

Under the two dynamic assumptions A1 (single-pole core amplifier) and A3 (full C_T at the summing node), the closed-loop transimpedance of the resistive SFB TIA is second-order and can be written in the canonical form [28]:

$$Z_T(s) = \frac{-R_F}{1 + \dfrac{s}{Q\,\omega_n} + \dfrac{s^2}{\omega_n^2}}, \qquad \omega_n = \sqrt{\frac{2\pi\,\mathrm{GBW}}{R_F\,C_T}}, \tag{9}$$

where the closed-loop quality factor Q increases monotonically with the loop time constant $R_F C_T$ relative to the available loop gain-bandwidth product. The maximally flat (Butterworth) response corresponds to $Q = 1/\sqrt{2}$, which is equivalent to a closed-loop phase margin of approximately 65° and, for a second-order system, places the -3 dB frequency exactly at ω_n. Substituting $Q = 1/\sqrt{2}$ into (9) returns the bandwidth identity used to define the transimpedance limit, $f_{-3\,\mathrm{dB}} \approx \sqrt{\mathrm{GBW}/(2\pi R_F C_T)}$, and hence the Butterworth-matched maximum feedback resistance $R_{F,\max} = \mathrm{GBW}/(2\pi f_{-3\,\mathrm{dB}}^2 C_{\mathrm{TOT}})$ that anchors the noise-floor boundary plots of Figure 2 and Figure 3.

The consequence is that the transimpedance-limit line is not merely a bandwidth boundary, but is simultaneously the maximum-gain locus, the minimum-noise locus, and the minimum-flatness ($Q = 1/\sqrt{2}$) stability locus. The three coincide because all three are consequences of the same denominator in (9). Pushing R_F above $R_{F,\max}$, the natural temptation when chasing lower noise, does not just overshoot a bandwidth target. It also drives Q above $1/\sqrt{2}$, moves the closed-loop complex-pole pair toward the imaginary axis, lowers the phase margin below 65°, and produces gain peaking with the corresponding step-response ringing. This is the precise mechanism behind the qualitative warning, already made in the limit derivation of this review, that the interaction between R_F and C_T may cause a phase shift that leads to ringing or oscillations. The same equivalence between the transimpedance limit, the maximally flat response, and the stability margin is developed in detail by Säckinger for SFB, common-gate, regulated-cascode, and distributed front-ends [2,4], and the practical compensation procedure, choosing a feedback capacitor C_F to set a target Q (and therefore a target phase margin) at the expense of bandwidth, is the standard industrial recipe [2]. The compensation capacitor across R_F introduced in the SFB discussion is therefore best understood not as a separate trick but as the designer trading a controlled amount of the limit-defined bandwidth for phase margin. Here the transimpedance limit becomes the natural reference frame for comparing the stability of competing architectures: an architecture earns a stability advantage precisely when an assumption relaxation lets it reach a target $(R_F, f_{-3\,\mathrm{dB}})$ operating point with a larger phase margin, or with weaker sensitivity of that margin to PVT, than the SFB baseline operating at the same point on its limit line.

 doi:10.20944/preprints202606.0753.v1

7.2. Tier 1 and Negative-Immittance Front-Ends (A3, optionally + A1): Stability Migrates to a Local Loop

Relaxing A3 by decoupling C_T from the feedback summing node is the most widely used bandwidth-extension strategy, and it is also the one whose stability behavior has been studied most explicitly in the recent literature. The RGC family attains its low input impedance through a local feedback loop wrapped around the common-gate input device. This local loop has its own loop gain and its own phase margin, distinct from the global transimpedance loop. The key stability consequence, analyzed for the low-voltage RGC by [74], is that the very action that extends bandwidth, raising the transconductance of the regulating stage to push the input pole to higher frequency, simultaneously erodes the phase margin of the local loop, so that beyond a critical g_m the step response rings and the amplifier ultimately becomes unstable [74]. There is, in other words, a direct bandwidth-versus-stability trade-off internal to the A3 mechanism, and the compensation technique proposed in that work exists specifically to buy back a degree of design freedom that the bandwidth push consumes (interestingly that the authors showed no increase in noise for the proposed technique). Low-power RGC variants that fold the regulating amplifier into a modified-MIC structure inherit this local-loop constraint, but because they also synthesize a negative input immittance, their dominant stability limit is the positivity (passivity) constraint discussed next [58].

The negative-immittance front-ends occupy the extreme of this behavior. The cross-coupled CIC TIA of Taghavi in [57] drives the input resistance toward zero by injecting a negative resistance, and a negative resistance element is only conditionally stable. The design is well-behaved only while the net input resistance R_{IN} remains positive across all corners. As already noted in the taxonomy, this requirement sets a lower bound on the cross-coupling inductance L_c and couples the achievable stability margin directly to the transimpedance gain [57], a coupling that is qualitatively absent in the resistive SFB baseline. Although the L_C-bearing variants of this family are classified as Tier 4 in Table 1 (the negative L_{IN} breaks A1 in addition to A3), they are grouped here by their shared stability locus, since the binding constraint in every case is the positivity of the synthesized immittance rather than the high-order pole placement that governs the Tier 2 entries. Negative-capacitance neutralization of the input pole behaves analogously, where [59] obtained broadband, high-gain operation by a negative-Miller-capacitance network, but the amount of capacitance that may be neutralized is bounded by the same conditional stability constraint, since over-neutralization presents a net negative capacitance to the input node and provokes oscillation [59]. In all three cases the A3 relaxation does not eliminate the stability question posed by C_T at the summing node, but rather converts it into a positivity (passivity) constraint on a synthesized immittance, which the designer must hold across PVT.

7.3. Tier 2 Zero Injection (A1): From Phase Margin to Group-Delay Flatness

Breaking A1 by injecting zeros, whether through passive inductive peaking, gyrator-C active inductors, or active zero-injection of the CH type, changes the nature of the stability question rather than its difficulty. A zero contributes phase lead and, if placed correctly, can improve the phase margin of the loop it sits in. This is why peaking becomes so attractive. However, every injected zero is accompanied by additional poles, raising the system order beyond two, so the single, scalar phase-margin criterion of (9) is replaced by a pole-zero-placement problem in which magnitude flatness, phase margin, and group-delay flatness can no longer be optimized independently. The dominant failure mode is therefore not theoscillation, but an excessive GDV and the ISI it produces under PAM-4 signaling. The PIP design of [33] is one of the clearest examples in the surveyed set. Its fifth-order transfer function, optimized for a $3.31\times$ bandwidth-extension ratio, exhibits a GDV in excess of 45 ps and a linear-phase range that falls several GHz short of the magnitude -3 dB bandwidth [33]. This is a textbook symptom of complex poles placed for bandwidth at the expense of a phase. Active realizations reintroduce a genuine loop-stability concern on top of the GDV concern, because a gyrator-C active inductor is itself a feedback structure with finite phase margin that can ring if its bias or loading drifts. Note that the transformer-coupled inverter front-ends similarly trade a magnetic resonance for bandwidth and their stability hinges on the damping of that resonance [32] (note that this work is

 doi:10.20944/preprints202606.0753.v1

also one of the few surveyed designs to report on explicit minimization of an in-band GDV (3,ps) as a primary objective). The recurrent engineering response is to use an explicit damping. For example, the distributed and cascode designs of [63] insert small (10-20 Ω) series damping resistors specifically to suppress the high-Q resonance of the cascode device that would otherwise make the cell difficult to stabilize [63]. In the taxonomy's terms, A1 relaxation moves the binding stability constraint from the gain margin of a second-order loop to the flatness of a high-order group-delay characteristic, and the cost of bandwidth is paid in phase linearity rather than in noise.

7.4. Tier 3 Non-Resistive Feedback (A2): Multi-Loop Stability

Relaxing A2 removes the thermal noise of R_F but replaces the frequency-independent resistive feedback factor with a frequency-dependent one, which has two stability consequences. First, with capacitive feedback the noise gain acquires its own poles and zeros, so the rate of closure between the open-loop gain and noise at the crossover frequency, rather than a single op-amp phase margin, governs stability. Here maintaining a 20 dB/decade rate of closure now requires deliberate placement of a compensating zero, and the same C_F-for-Q trade-off of Section 7.1 reappears in a capacitive guise [2]. Second, and more importantly, a capacitive or pseudo-resistive feedback network cannot set the DC operating point, so these designs invariably add a low-frequency DC-servo or pseudo-resistor bias loop. The TIA then becomes a multi-loop system, and its stability requires that the servo loop shall be slow enough to be dynamically decoupled from the signal loop. If the two unity gain frequencies approach each other, the design exhibits low-frequency peaking rather than high-frequency ringing. This is why the capacitive-feedback front-ends of [43] and the related works build their DC-cancellation paths around deliberately high-impedance pseudo-resistors, where the large servo time constant is a stability requirement, not merely a noise-bandwidth choice. The same multi-loop, slow-servo requirement appears even when the feedback element is resistive. The wide dynamic range SFB design of [35] adds a *DC-restore* loop that senses the DC voltage across R_F and equalizes its two terminals so that only the small signal is fed back, decoupling the quiescent operating point from PD DC-current drift. As with the CF DC-servos above, this loop must be low-bandwidth enough to be dynamically separated from the signal loop, so that an A2-intact resistive TIA acquires the same multi-loop stability character once a DC-cancellation path is added. A complementary illustration that the multi-loop stability concern is not confined to A2-broken designs is the dual-feedback transimpedance-limiting amplifier of [16]. Although its feedback is entirely resistive (a Tier-0 core in the taxonomy of Section 4), it superimposes a switched active-feedback limiting path on the main SFB for dynamic range extension. The authors report that the active path feedback resistor must be carefully optimized to preserve stability [16], making explicit that once a second feedback mechanism coexists with the main loop, regardless of which assumption it does or does not relax, the stability margin is set by the interaction of the loops rather than by either of them alone.

7.5. Tier 4-5 Compound and Paradigm-Shift Topologies

Compound Tier 4 designs inherit the union of the stability constraints of the assumptions they relax and those constraints interact. An RGC-plus-peaking design must simultaneously satisfy the local-loop phase margin of its A3 mechanism and the group-delay flatness of its A1 mechanism, and the inductive peaking that flattens the magnitude response can reduce the local-loop phase margin it was not designed to protect. This is why the most aggressive Tier 4 optical front-ends are also those that report the most elaborate compensation networks. The system-level compound route, exemplified by the 16 nm FinFET receiver discussed earlier, sidesteps intra-cell multi-loop interaction by keeping a Tier-0 SFB cell stable in isolation and shifting the bandwidth-recovery effort into an off-loop equalizer chain, at the cost of equalizer-chain phase budgeting and the added ESD-capacitance loading of the summing node that gets harder to stabilize as the input-referred bandwidth rises.

Tier 5 paradigm-shift topologies replace the continuous-time phase-margin question entirely. In the distributed/traveling-wave TIA the gain cells must each present the correct termination to the artificial LC ladder at every in-band frequency, so stability becomes a matter of preserving the

 doi:10.20944/preprints202606.0753.v1

ladder's characteristic impedance. No cell can be power-gated or under-biased without perturbing the impedance and provoking ripple or band-edge peaking. Critically, the absence of any global feedback path means there is no virtual-ground mechanism to stabilize the DC operating point against PVT drift [63]. The discrete-time switched-capacitor TIA shifts the question further into the sampled-data domain, where the stability is a settling-and-aliasing problem (the integrator must settle within the sampling phase and the loop must not undersample its own dynamics) rather than a phase-margin problem. This is why the continuous-time transimpedance limit and the corresponding Q-based stability criterion of Section 7.1, simply do not apply.

7.6. Synthesis: Stability Is Conserved, Not Removed

Table 2 consolidates the discussion by mapping each assumption relaxation onto the location of its stability margin, its dominant failure mode, and a representative recent analysis. The pattern that emerges is the converse of the performance story told by the rest of the review. Where each assumption relaxation removes one face of the transimpedance limit, gaining bandwidth (A1), noise (A2), or input-pole headroom (A3), it conserves the stability burden by relocating it into the new mechanism: a local feedback loop (A3), a group-delay characteristic (A1), or a second DC-servo loop together with a frequency-dependent noise gain (A2). The total amount of stability analysis required does not fall as the architecture grows more sophisticated as it migrates in one-tiers and, in the compound tiers, compounds. This forms the stability counterpart of the review's central thesis. Just as the taxonomy predicts which performance bottleneck a given topology relieves, it also predicts what form the stability analysis must be taken for that topology, and therefore which design margin should be scrutinized first.

A practical corollary concerns reporting. Very few of the surveyed works report quantitative phase margin, Q, or PVT-spread of the closed-loop response, and the only design in the curated set with measured closed-loop PVT compensation is the dynamic voltage scaling FinFET front-end of [6], whose on-chip loop holds the inverter g_m and triode R_F at their target values, and thus the closed-loop response near its intended Q, across the TT/SS/FF corners [6]. The framework of this section suggests that a more uniform reporting convention, stating the closed-loop Q or phase margin and its worst-case PVT spread alongside the conventional gain-bandwidth-noise figures, would make the stability dimension of the taxonomy as comparable across works as the FOM already makes the sensitivity dimension.

doi:10.20944/preprints202606.0753.v1

Table 2. Stability margin by assumption-relaxation tier. The "stability locus" indicates where the dominant stability constraint resides once the assumption is relaxed; the limit-based $Q = 1/\sqrt{2}$ ($\sim$65° phase-margin) criterion of Section 7.1 applies directly only to Tier 0.

Tier (relaxed assumption)	**Stability locus**	**Dominant failure mode**	**Representative analysis**
Tier 0 (none)	Global loop phase margin	Gain peaking / step ringing when $R_{\mathrm{F}} > R_{F,\max}$ ($Q > 1/\sqrt{2}$)	[2,4,16]
Tier 1 (A3)	Local (regulating) loop	Local-loop ringing as g_m is raised for bandwidth	[74]
Tier 1/4 (A3 [+ A1], neg. immittance)	Synthesized-immittance passivity	Oscillation when $R_{\mathrm{IN}} < 0$ / over-neutralization	[57–59]
Tier 2 (A1)	High-order pole-zero placement	Group-delay variation / ISI; active-L resonance	[32,33,63]
Tier 3 (A2)	Multi-loop (signal + DC servo)	Low-frequency peaking / motorboating; rate-of-closure >20 dB/dec	[43]
Tier 4 (compound)	Union of the above, interacting	Peaking-loop vs. regulating-loop margin conflict	[5,6]
Tier 5 (paradigm shift)	Ladder impedance / sampled-data	Band-edge ripple, PVT operating-point drift, settling	[63]

 doi:10.20944/preprints202606.0753.v1

8. Digitally-Assisted and High-Linearity TIAs

8.1. AGC and Wide-Dynamic-Range (WDR)

Beyond gain compression for sensitivity adaptation, recent CMOS TIAs increasingly target parametric reconfigurability across multiple operating points. Classical AGC adjusts a single control axis (usually R_F or g_m) to suppress the overload. The widest reported instance of this single-axis class is the optical fiber sensing front-end of [17], which partitions the input range into inactive, one-active and both-active regions. The solution contains a peak detector-driven MOSFET (one transistor replaces the feedback resistor while a shunt transistor diverts excess photocurrent) and yields a 123.5 dB dynamic range [17]. A mechanistically distinct single-control instance is the CMOS Transimpedance-Limiting Amplifier (CTLA) of [16] for short-range LiDAR. Rather than compressing gain with a peak detector-driven AGC, a diode-connected NMOS switch engages a second, active feedback path once the input photocurrent exceeds a fixed threshold, lowering the input resistance and driving the output rapidly into limiting. This folds the conventional power hungry multi-stage LA into the TIA itself, yielding a 56.6 dB input dynamic range at 23.6 mW in 180 nm CMOS [16]. Unlike the Liu front-end, no always-on input impedance reduction element is present, so the small signal core remains a plain Tier-0 inverter SFB design and the limiting path is purely a single-control dynamic-range mechanism on the configurability axis. A third single-control variant, predating both, is the wide-dynamic-range preamplifier of [35] in 0.13 µm CMOS. Here the gain knob is neither a peak detector AGC nor a switched limiter, but an input-stage MOSFET biased in the linear region that diverts part of the photocurrent away from R_F, manually setting the transimpedance and yielding a 5 µA-2 mA (≈52 dB) input current range. As in [16], the small-signal core stays a plain Tier-0 resistive SFB cell, confirming that the dynamic-range mechanism is orthogonal to the assumption-relaxation tier regardless of whether it is implemented as AGC, switched limiting, or current diversion. On the other hand, the dual-control variable bandwidth or gain class adjusts two control axes simultaneously based the recurring observation that a single tuning parameter cannot independently shape both the lower and the upper edge of the passband. Two designs at structurally distinct tiers illustrate the same dual-control principle. [7] implement it at Tier 0 (inverter SFB) in 65 nm CMOS. Here a 12-bit binary-weighted PMOS array placed between V_{DD} and the PMOS source of every inverter stage in the TIA and CH post-amplifier reconfigures the bias current of every stage simultaneously, while a switched-MOSFET resistor bank replaces every R_F. Switches rather than triode MOSFETs are used in the resistor bank explicitly to reduce PVT sensitivity at the high-R_F corners. The two control axes are tied in lockstep so that the front-end transimpedance remains constant, while bandwidth and power scale linearly with data rate from 1.25 Gb/s/0.32 mW to 20 Gb/s/13.5 mW. The work [73] implements the same dual-control principle at Tier 3 (CF, A2 broken) for OTDR sensitivity adaptation. Here one feedback capacitor C_1 is reconfigured via a parallel-switched-capacitor bank to address the upper passband, while a switched bias-resistance at the gate of the input transistor independently sets the lower passband through the high-pass component. The five gain configurations span 10–500 kΩ scaling linearly with R_T. Taken together, the two designs make the structural point that the configurability axis is orthogonal to the assumption relaxation tier as each operating point sits at its own fixed tier (Tier 0 for Dash, Tier 3 for Romanova), and the dual-control mechanism is the means by which the design traverses points that would otherwise demand distinct fixed implementations.

8.2. Integrated Post-Processing

When the TIA shares a die with the receiver DSP, bandwidth recovery can be delegated entirely to post-processing, and the equalizer may act in either the frequency or the time domain. The frequency-domain route (CTLE, VGA, in-cell peaking) was discussed above. The time-domain alternative is Decision Feedback Equalization (DFE), exemplified by the 64 Gb/s 14 nm FinFET receiver of [9]. Here the SFB inverter TIA is deliberately bandwidth-limited below the classical data-rate optimum so that R_F can be enlarged for higher gain and lower noise, and a 1-tap speculative DFE recovers the resulting post-cursor ISI. The structural significance for this review is that a DFE recovers bandwidth without

 doi:10.20944/preprints202606.0753.v1

the high-frequency noise amplification intrinsic to any linear equalizer. Here the CTLE boosts signal and noise alike and adds its own noise, whereas a 1-tap DFE cancels ISI with essentially no noise penalty, improving SNR by $\sim \sqrt{2}$ at the optimum [9]. This places the design in the same system-level compound family as [5,6], but with the off-loop A1 relaxation performed in the time domain rather than the frequency domain. Two cell-level techniques further distinguish the work. First the TIA reuses its own input node as the negative differential output (a self-referenced scheme) rather than instantiating a replica, which restores the full R_F transimpedance, lowers the integrated noise and converts low-frequency device noise to common mode so that flicker noise is absent from the input-referred spectrum. Finally a closed-form power-supply-rejection condition is met by sizing the inverter symmetrically and balancing the pad capacitance to both rails [9].

Equalization need not reside off-loop. A CTLE or VGA co-designed after the TIA recover bandwidth and channel loss, but at the cost of extra power and, in low-V_{DD} advanced nodes, considerable design difficulty [32]. The alternative is to fold the equalization into the TIA cell as analog passive equalization, sacrificing neither static power nor low-frequency gain. The transformer-based inverter TIA of [32] is a representative case, reaching 64 Gb/s PAM-4 without an explicit equalizer (Section 6). Where the equalization effort is placed, either off-loop in a CTLE/DSP chain versus in-cell as passive peaking, is another instance of the system-versus-cell design-culture split noted earlier.

9. Comparative Analysis & Benchmarking

9.1. The State-of-the-Art (SOTA) Landscape

A set of representative recent works on TIA in CMOS with the main parameters is shown in Table 3. Note that works with incomplete or ambiguous specification of the TIA design (i.e. missing noise current as in [34] or power ([50], [45]) are not shown in the table). It is also holds for works like [58], where, for example, the minimum and not band averaged value is given and recalculation to the band-average may be non-trivial. With this we do not provide FOM values which are not directly comparable to other entries even though the designs can be extremely promising (i.e. see [13]). On the other hand, in some of the works the reported power consumption is not for the TIA, but may include the subsequent amplification steps and the buffer as done, for example, in [7]. In this case the numbers are reported for the complete module if TIA-specific values cannot be unambiguously extracted from the publication.

The reason a single quantitative metric is needed in the first place is that direct, line-by-line comparison of the entries in Table 3 is structurally misleading. The surveyed designs span almost two decades of CMOS process evolution, from 600 nm to 16 nm FinFET, with raw f_T increasing by more than an order of magnitude across that range. Yet the analog-relevant device figures of merit do not scale uniformly with f_T. For example, the intrinsic gain $g_m r_o$ generally degrades as the channel shortens, matching worsens, and the noise corner shifts with the supply voltage collapsing toward 0.7-0.8 V at advanced nodes. A TIA in 16 nm FinFET with a nominally higher f_T than a 130 nm BiCMOS counterpart does not automatically achieve better gain-bandwidth-noise efficiency at the circuit level, because the headroom available to bias for high r_o shrinks in step with V_{DD} [38]. A meaningful comparison metric must therefore consolidate the four primary electrical observables (transimpedance gain, bandwidth, input-referred noise, and power dissipation) into a single number that rewards designs achieving high sensitivity and bandwidth at low energy-per-bit, while remaining agnostic to the specific process node used to obtain those numbers. The classical product $Z_T \cdot BW/P$ falls short of this requirement because it ignores the noise penalty paid for bandwidth extension and underweighs designs operating with large PD or pad capacitances, both of which dominate the achievable sensitivity at the system level.

To allow a quantitative, single-criterion comparison of TIA designs fabricated in different technology nodes and targeting different applications, a FOM is employed. A comprehensive FOM

 doi:10.20944/preprints202606.0753.v1

that balances the primary TIA performance parameters such as bandwidth, transimpedance gain, input-referred noise, and power consumption was proposed in [20]:

$$\mathrm{FOM} = \frac{\sqrt{BW\,[\mathrm{GHz}]} \;\cdot\; R_T\,[\Omega] \;\cdot\; C_T\,[\mathrm{pF}]}{\overline{i_n}\,[\mathrm{pA}/\sqrt{\mathrm{Hz}}] \;\cdot\; P\,[\mathrm{mW}]}, \tag{10}$$

where BW is the -3 dB bandwidth (in GHz), R_T is the transimpedance gain in linear scale (Ω; note that Table 3 lists R_T in dBΩ), C_T is the total input capacitance in pF, $\overline{i_n}$ is the average input-referred noise current spectral density in $[\mathrm{pA}/\sqrt{\mathrm{Hz}}]$, and P is the total power consumption in mW. The square root of the bandwidth is used deliberately to maintain dimensional consistency with the noise current spectral density expressed in $\mathrm{pA}/\sqrt{\mathrm{Hz}}$; a linear BW term in the numerator would otherwise introduce an undue bias toward wide-band designs and over-penalize noise-optimized, lower-speed amplifiers [20]. The capacitance factor C_T in (10) rewards designs capable of operating with larger PD (and bonding-pad, ESD) input capacitances, which is relevant for applications such as OTDR or LiDAR where the PD junction capacitance is non-negligible. In practice, however, C_T is not universally reported in the literature as many works either specify only the bare value of C_D without accounting for pad and ESD parasitics, or omit the value entirely. Retaining C_T in the FOM would therefore require adopting an assumed default value for a substantial fraction of the surveyed works which may alter the resulting ranking in a non-trivial way. To preserve comparability across all entries of Table 3 without introducing such assumptions, the following reduced FOM is suggested:

$$\mathrm{FOM} = \frac{\sqrt{BW\,[\mathrm{GHz}]} \;\cdot\; R_T\,[\Omega]}{\overline{i_n}\,[\mathrm{pA}/\sqrt{\mathrm{Hz}}] \;\cdot\; P\,[\mathrm{mW}]}. \tag{11}$$

This expression retains the essential trade-off between speed, gain, noise, and power while remaining computable from the parameters that are consistently and unambiguously reported across the surveyed literature. It should be noted that, as with any scalar performance metric, the FOM in (11) is an approximation. The expression does not capture linearity, group-delay flatness, or dynamic range, and should be regarded as a *comparative indicator* rather than an absolute measure of design quality. Other works also proposed different FOM formulations. For example the authors in [74] suggested to count the number of used inductors. This however, can be also interpreted as an implicit area penalty and makes fewer sense for modern designs which either avoid inductors whatsoever or often use active inductor techniques. Finally, we use the square root of the bandwidth, while [61] employs the direct value in GHz. An interesting suggestion had been provided in [60], where the authors suggested to add f_T to the denominator of the FOM expression. [13] proposed a FOM that additionally divides by active chip area, rewarding the inductor-less compactness that distinguishes their 28 nm design. While this captures a genuine engineering achievement, active area is not adopted here for two reasons: it is inconsistently reported across the literature (many works cite only the pad frame footprint, or omit the figure entirely), and it conflates circuit topology with the layout skills and design-rule maturity of the specific technology node, both of which vary independently of the electrical performance being compared. An interesting FOM was suggested [62] where the supply voltage was included in the denominator, while the bandwidth in the numerator was is raised to the power 2 in order to emphasize its importance in optical communications as addressed by the authors.

doi:10.20944/preprints202606.0753.v1

Table 3. Parameters and performance of selected TIAs in CMOS.

Author	Year	Process, nm (CMOS)	R_T (dBΩ)	$BW_{-3\,\mathrm{dB}}$ (GHz)	Power / VDD (mW, V)	Avg. in.-ref. noise (pA / $\sqrt{\mathrm{Hz}}$)	Topology	Application
Abdelrahman [51]	2016	130	56.65	7.0	1.95 @ 1.5 V	7.5	CS + gain-boosted cascode and resistive R_F	Optical Receiver
Abdollahi [74]	2018	180	52.0	3.0	4.32 @ 1.8 V	18.5	Low-voltage RGC (level-shifter) with R_m–C_m compensation; inductorless	Optical Receiver
Atef [60]	2013	40	47.0	8.0	2.03 @ 1.1 V	22.0	Inverter (main amp) + common-drain (PMOS) feedback	Optical Receiver
Atef [38]	2013	40	55.3	8.0	3.01 @ 1.2 V	12.3	Inverter-cascode (cascoded inverter core) + resistive R_F	Optical Receiver
Bashiri [12]	2010	65	46.7	26.1	39.9 @ 1.2 V	30.0	RGC + CH($R_f + L_f$) + L_s input	Optical Receiver
Dash [7]	2013	65	74.5	13.1	13.5 @ 1.0 V	18.0	Inverter SFB + binary-weighted PMOS array/switched-MOSFET R_F bank	Optical Receiver
Hammoudi [53]	2010	350	54.5	2.75	53.5 @ 3.3 V	12.8	3-stage push-pull inverter with diode-connected PMOS active load global PMOS-in-triode R_F	Optical Receiver
Honarmand [56]	2020	65	46.24	3.52	0.573 @ 1.0 V	15.0	Active-RGC input; active-inductor output; resistive R_F + 3-stage LA	Optical Receiver
Hosseinisharif [15]	2020	90	40.0	6.4	1.6 @ 1.2 V	25.0	Regulated cascode; cascoded inverter; active inductor output	Optical Receiver
Hu [21]	2010	65	162.35	0.00145	0.0293 @ 0.8 V	0.09	Capacitive-feedback TIA with Miller Two-Stage OpAmp	Biosensor
Jin [33]	2008	180	51.0	30.5	60.1 @ 1.8 V	55.7	Four cascaded CS stages with PIP	Optical Receiver
Kromer [30]	2004	80	52.8	13.4	2.2 @ 1.0 V	28.0	FCG input + passive shunt-peaking L	Optical Receiver (short-reach)
Lakshmikumar [6]	2019	16	78.0	27.0	60.8 @ 1.8 V	18.3	Inverter SFB + MOSFET-triode R_F + 4-stage g_m/g_m shorted-inverter PGAs with shunt-peaking L's + DVS PVT-compensation loop	Optical Receiver
Liu [63]	2004	180	58.0	15.0	200.0 @ 1.8 V	12.0	Two cascaded 3-stage cascode distributed amplifiers; DC-coupling source-follower interstage	Optical Receiver

doi:10.20944/preprints202606.0753.v1

Table 3. *Cont.*

Author	Year	Process, nm (CMOS)	R_T (dBΩ)	$BW_{-3\,\mathrm{dB}}$ (GHz)	Power / VDD (mW, V)	Avg. in.-ref. noise (pA / $\sqrt{\mathrm{Hz}}$)	Topology	Application
Liu [63]	2004	180	48.0	30.0	50.0 @ 1.8 V	16.0	Single-stage cascode distributed amplifier, 2 gain cells, DC-coupled	Optical Receiver
Liu [17]	2015	180	87.8	1.4	8.1 @ 1.8 V	2.75	Three-stage push-pull inverter SFB + input series-peaking L (decouples C_D) + classical single-control AGC (peak-detector → MOSFET R_F M_{10} + shunt M_{11}); 123.5 dB DR (10 nA–15 mA)	Optical Fiber Sensing
Oh & Lee [39]	2004	350	68.0	1.73	50.0 @ 3.3 V	3.3	Cascode + R_F + active-L load	Optical Receiver
Oh & Park [8]	2007	180	96.0	4.7	72.0 @ 1.8 V	25.0	ACG (RGC-equivalent) + CS + R_f + f_T-doubler DC cancellation, fully-differential	Optical Receiver (OPCB)
Ozkaya [9]	2017	14	56.9	14.2	3.52 @ 0.9 V	5.7	Self-referenced SFB inverter + series-peaking L + off-loop time-domain 1-tap speculative DFE (system-level compound Tier 4)	Optical Receiver (NRZ)
Park [11]	1998	600	61.0	3.5	135.0 @ 5.0 V	4.2	Regulated cascode	Optical Receiver
Park [47]	2007	180	64.0	2.1	50.0 @ 1.8 V	32.0	CG input + 2-stage CS gain + LPF-feedback DC-balance buffer; 4-ch fully-differential array	Optical Receiver (DVI/HDMI)
Patel [5]	2023	16	63.0	32.0	47.0 @ 0.9 V	16.9	BW-limited SFB inverter front-end + input T-coil/ESD + CH CTLE + VGA/CTLE + final TIS with L_1/L_2 peaking (co-packaged optics, system-level compound Tier 4)	Optical Receiver (PAM-4)
Razavi [1]	2000	600	78.8	0.55	30.0 @ 3.0 V	4.5	Capacitive-feedback TIA	Optical Receiver
Romanova [20]	2021	250	80.2	0.9	29.3 @ 2.5 V	1.6	Capacitive-feedback TIA with DC current elimination and PMOS biasing	OTDR
Romanova [20]	2022	180	80.3	1.0	21.0 @ 1.8 V	1.8	Capacitive-feedback TIA with DC current elimination and PMOS biasing	OTDR
Sadeghi [14]	2021	65	60.01	1.64	0.76 @ 1.0 V	5.79	Cascoded-booster RGC	Optical Receiver

doi:10.20944/preprints202606.0753.v1

Table 3. *Cont.*

Author	Year	Process, nm (CMOS)	R_T (dBΩ)	$BW_{-3\,\mathrm{dB}}$ (GHz)	Power / VDD (mW, V)	Avg. in.-ref. noise (pA / $\sqrt{\mathrm{Hz}}$)	Topology	Application
Salhi [46]	2017	180	50.8	7.9	7.2 @ 1.8 V	7.7	Push-pull inverter + MOSFET-in-triode + L	Optical Receiver
Shaaban [62]	2024	45	30.6	2.8	1.63 @ 1.0 V	23.6	Current-mirror + R_F + shunt passive L_1 (out) L_2 (in)	Optical Receiver
Shahdoost [43]	2014	180	75.5	1.62	26.3 @ 2.2 V	3.18	Capacitive-feedback TIA with DC current elimination	Optical Receiver
Shahdoost [72]	2016	130	76.0	1.76	13.7 @ 1.5 V	2.67	Capacitive-feedback TIA with DC current elimination	Optical Receiver
Soltanisarvestani [55]	2020	90	50.5	7.3	1.0 @ 1.2 V	13.7	Cascoded-booster RGC; active-inductor output; resistive R_F	Optical Receiver
Taghavi [57]	2015	130	55.3	6.0	2.0 @ 1.5 V	24.	Cross-coupled immittance-converter (CIC): CG input + cross-coupled M_2 (negative-R) + L_c (negative-L) + cascode M_3 (Miller suppression); first-order response (TIA1)	Optical Receiver
Taghavi [57]	2015	130	55.4	8.8	2.0 @ 1.5 V	28.	Cross-coupled immittance-converter (CIC): CG input + cross-coupled M_2 (negative-R) + L_c bondwire-extended (negative-L) + cascode M_3; second-order response (TIA2)	Optical Receiver
Wu [37]	2005	180	61.0	7.2	70.2 @ 1.8 V	8.2	5-stage inverter + multiple inter-stage series L + M-derived I/O matching	Optical Receiver
Zhang [59]	2023	65	60.1	23.7	4.8 @ 1.8 V	40.7	FCG input (A3) + capacitive-degeneration CS zero + inter-stage negative-Miller-C neutralization	Optical Receiver
Zohoori [61]	2019	180	42.24	1.96	0.97 @ 1.5 V	11.7	Modified push-pull inverter; active inductor output; diode-connected input; resistive SFB	Optical Receiver

 doi:10.20944/preprints202606.0753.v1

We first delimit the scope of the FOM and the rationale for computing it. Scalar FOMs compress the gain-bandwidth-noise-power trade-off into a single quantity, but are usually defined by the authors of the design with the functional form chosen to favor that design's strengths. Cross-paper comparison on such author-defined metrics is consequently unreliable. Applying (11) uniformly across Table 3 removes this degree of freedom, yielding a reproducible measure computed solely from consistently reported parameters. So constructed, the metric is diagnostic rather than ordinal as it resolves clusters and outliers and permits a direct test of two common expectations, namely that a higher assumption-relaxation tier and a newer technology node each raise the FOM. However, in practice, neither holds. As we can see the most complex multi-assumption designs are not FOM-optimal, and neither Figure 8 nor Figure 9 exhibits a monotonic dependence of FOM on node or publication year. This absence of trend is itself a result, consistent with node and topology acting as complementary rather than competing design levers.

Four caveats constrain interpretation. (i) The sample is representative, not curated by citation count. Here a low FOM therefore typically indicates optimization for an objective excluded from (11), such as dynamic range [16,17], group-delay flatness, linearity, or area, rather than an inferior design itself. (ii) Measured and post-layout-simulated entries are mixed in Figs. 8, 9. The simulated values are systematically optimistic and not strictly commensurate with measured ones. (iii) Reporting is heterogeneous as gain may be small-signal or compressed, bandwidth specified with or without the PD, noise averaged or peak, and power referred to the cell or the full chain (hence the cell-level figures adopted for [32]), so reporting convention alone can shift the FOM appreciably. (iv) A scalar cannot rank designs across application domains as a high-gain, sub-megahertz biosensor and a multi-GHz optical datacom front-end lie on incommensurate axes, and any mixed ranking is dominated by high-gain, low-bandwidth sensing entries irrespective of design quality. The FOM is thus valid only as a comparative indicator within a common bandwidth and application class (e.g. see grouping in [31]), and primarily identifies which designs warrant the mechanistic analysis that follows.

Nevertheless, several interesting observations can be done from the table. Switch from RGC to CF architecture for optical receivers at 600 nm resulted in enormous FOM improvement. It is also interesting to observe that relaxing two constraints simultaneously (e.g. A1 + A3 group in [15] does not automatically result in superior FOM value, although one can admit that the achieved power (1.6 mW) is fairly low compared to other designs with GHz range bandwidth. Note that both [15,61] reached rather low power consumption of around or below 1mW, while [61] also added a replica TIA to eliminate the common mode noise before the differential LA stages. A third low-power member of this A1 + A3 group, [55], reaches 1 mW at 7.3 GHz while its relatively large FOM stems from its higher gain and lower (simulated) noise, a reminder that the simulated entries sit optimistically high. The PIP design of [33,41] achieved the highest passive peaking factor (3.31×) but paid a steep FOM penalty from its elevated noise and power, illustrating that maximizing the BW alone my result in significant power consumption and noise current of 55.7 $pA/\sqrt{Hz}$. Note as the [41] is open-loop, it shows factor 4–5× worse noise than a well-designed RGC or CH TIA at similar bandwidth. This is the single biggest penalty of the open-loop approach. Another point is that four cascaded high-order peaking networks accumulate phase non-linearly. Jin reports linear phase only up to 26 GHz, already 4.5 GHz short of the 3-dB bandwidth and this is a direct symptom of the open-loop cascade topology. Finally the power can be also considered larger than necessary. With no feedback to set the operating point via a virtual ground, each stage must be individually biased for both gain and speed, leading to the relatively high 60.1 mW power consumption.

At the opposite extreme, the fiber-sensing front-end of [17] attains the highest FOM of the presented curated set, but for reasons orthogonal to its headline feature. The value is driven entirely by the simultaneous combination of high gain, low input-referred noise and relatively low power at a modest 1.4 GHz, while its 123.5 dB dynamic range, the actual contribution of the single-control AGC, is precisely the quantity that (11) omits. This design therefore exemplifies the limitation noted with (11): a scalar FOM rewards the gain-noise-power corner and is blind to the wide-dynamic-range objective

 doi:10.20944/preprints202606.0753.v1

that motivated the given topology in the first place. The same caveat applies to the LiDAR limiting-TIA of [16]. Its high FOM reflects an 88.8 dBΩ small-signal gain at low noise and modest power, whereas its actual design objective, the 56.6 dB dynamic range delivered by the switched-active-feedback limiting, is again invisible to (11).

One of the most recent FinFET entry, [5], sits at the opposite corner of the high-speed cluster. Its position is structurally informative because the TIA cell itself is a single Tier-0 SFB inverter. The compound Tier 4 (A1 + A3) behavior visible in this high-bandwidth/moderate-power point is contributed entirely by the input T-coil and the off-loop CTLE/VGA chain, which between them absorb the BW-noise trade-off that an intra-loop compound such as [15] packs into the TIA itself. The two routes to the same tier therefore populate adjacent regions of the FOM space rather than competing for the same point, with the choice between them driven primarily by whether the receiver architecture already mandates DSP-grade equalization downstream.

The [6] design sits at the same 16 nm node but at FOM ≈ 37 and is an almost $4\times$ improvement over [5] at the same node and four years earlier, contributed by a more aggressive intra-cell A2 relaxation (MOSFET-triode R_F instead of fixed R_F), PGA chain with shunt peaking in every stage, and an active DVS compensation loop.

A structurally analogous but mechanistically distinct power penalty is incurred by the distributed TIA of [63], which achieves 30 GHz bandwidth in the same 0.18 µm node with a relatively low FOM. The root cause differs from the PIP case, however. Whereas the PIP penalty arises from open-loop noise accumulation across four cascaded stages, the distributed penalty is structural. All N gain cells must be biased at full quiescent current simultaneously, because each cell must present the correct termination impedance to the LC ladder at every frequency up to the cutoff. Here no cell can be power-gated or operated at a reduced bias without perturbing the ladder's characteristic impedance and collapsing the bandwidth. Bandwidth in the distributed paradigm is therefore purchased directly with bias current rather than with an assumption relaxation in the noise budget and the FOM defined by (11), which penalizes both noise and power in the denominator, captures this trade-off cleanly. The two designs together illustrate a general principle that aggressive open-loop topologies, whether based on cascaded peaking networks or distributed gain, consistently populate the low-FOM tail of the high-speed cluster, because both paradigms decouple bandwidth from the feedback-stabilized operating point that allows SFB and RGC designs to achieve competitive noise and power figures simultaneously.

A notable secondary benefit of PIP is at the input stage, first analytically identified in [33]. The inductive reactance of the PIP network increases with frequency, partially blocking the noise currents of the matching resistors and the first transistor from reaching the input node. This results in a frequency decreasing input-referred noise current over the in-band frequency range, i.e. the opposite of the typically rising noise in capacitance-loaded stages, with the noise current falling from a relatively high value at low frequencies (11.9 dB NF at 3.5 GHz) to a minimum of near 25 GHz. A complementary demonstration of the PIP technique, targeting 10 Gb/s operation, was reported by the same group in [48]. Extending the cascade to six stages (vs. four for the 40 Gb/s design) the circuit achieves the highest Z_T among 10 Gb/s CMOS TIAs at time of publication. Interestingly, the stage count of six was not chosen to maximize the bandwidth (the bandwidth-optimal count is nine), but rather to maximise GBW/Power ratio as beyond six stages the power increases faster than bandwidth, reversing the FOM trend.

Unfortunately rather few works reported on accurate PVT analysis when presenting their designs. Except of several rather general claims on PVT performance for a particular type of architecture, no PVT analysis is shown. Some works did, however, present such results. For example [62] reports static PVT simulations and [55,56] likewise report simulated temperature and supply-voltage sweeps together with Monte-Carlo runs of gain and bandwidth, though again in simulation only. The work [6] goes further by reporting both measured PVT characterization across TT/SS/FF samples and an on-chip dynamic voltage scaling loop that actively compensates the inverter g_m and triode R_F

 doi:10.20944/preprints202606.0753.v1

against PVT drift, improving the worst-case bandwidth from 10.5 GHz to 26.3 GHz. Unfortunately, absence of the comparison basis (other works) does not allow to benchmark the designs w.r.t. this important performance metric, but [6] establishes that closed-loop PVT compensation is a viable and necessary architectural element for sub-1 V FinFET inverter SFB front-ends, conceptually distinct from the user-controlled configurability axis of Section 8.1.

Note that Tier 0 approach (inverter with R_F) reported in [7] also demonstrates relatively high FOM without any special circuit-level adjustments except of achieving low-power via current reusage by the push-pull inverter architecture. This high number, still, can be also probably attributed to rather low input capacitance of 100 fF used in the network simulation. For computing the FOM value for [32] standalone TIA-cell numbers and not the abstract's full-chain values were used and the latter would include the balun, CML buffer, and offset-cancellation loop, so using them would unfairly penalize the FOM against other works especially those which design the TIA front-end only.

9.2. Graphical Trend Analysis

To provide a more generalized insight into the evolution of the TIA in CMOS performance in time and its dependency on CMOS node we show the FOM value plots for the works in the Table 3 correspondingly in Figure 8 and Figure 9. Here we shall remark that FOM is a comparative indicator within similar application domains rather than an absolute ranking across all application classes. For example, the work [21] reports a capacitive feedback TIA with enormous FOM of approximately 1e+9. However, since the FOM scales as linear, not log, the 162 dBΩ but extremely low-bandwidth bio sensor design will always dominate any FOM plot that mixes application domains. The same caveat applies to the 155 dBΩ MEMS-resonator front-end of [28], whose 1.8 MHz bandwidth and femto-ampere noise floor place it in the same high-FOM, low-bandwidth sensing regime as [21].

Within the high-speed optical cluster of Figs. 8 and 9, the most recent FinFET entry [5] sits at the 16,nm corner and illustrates the broader trajectory of the field. The highest-bandwidth measured CMOS TIAs of the last five years have all adopted some form of compound Tier 4 architecture, but they split into two structurally distinct sub-families. Intra-loop compound designs (RGC plus active zero injection, e.g. [12,15,61]) dominate at older sub-100,nm nodes where DSP-grade equalization in the same package is not yet standard. System-level compound designs (Tier-0 cell + input T-coil + off-loop CTLE/VGA, e.g. [5]) emerge at 16,nm, where the SoC context already includes the CTLE/VGA/ADC chain and the TIA can be deliberately simplified to minimize its noise contribution. This requires us to separate the intra-loop and the system-level compounding as the latter itself is strongly correlated with the maturity of the receiver-DSP ecosystem.

doi:10.20944/preprints202606.0753.v1

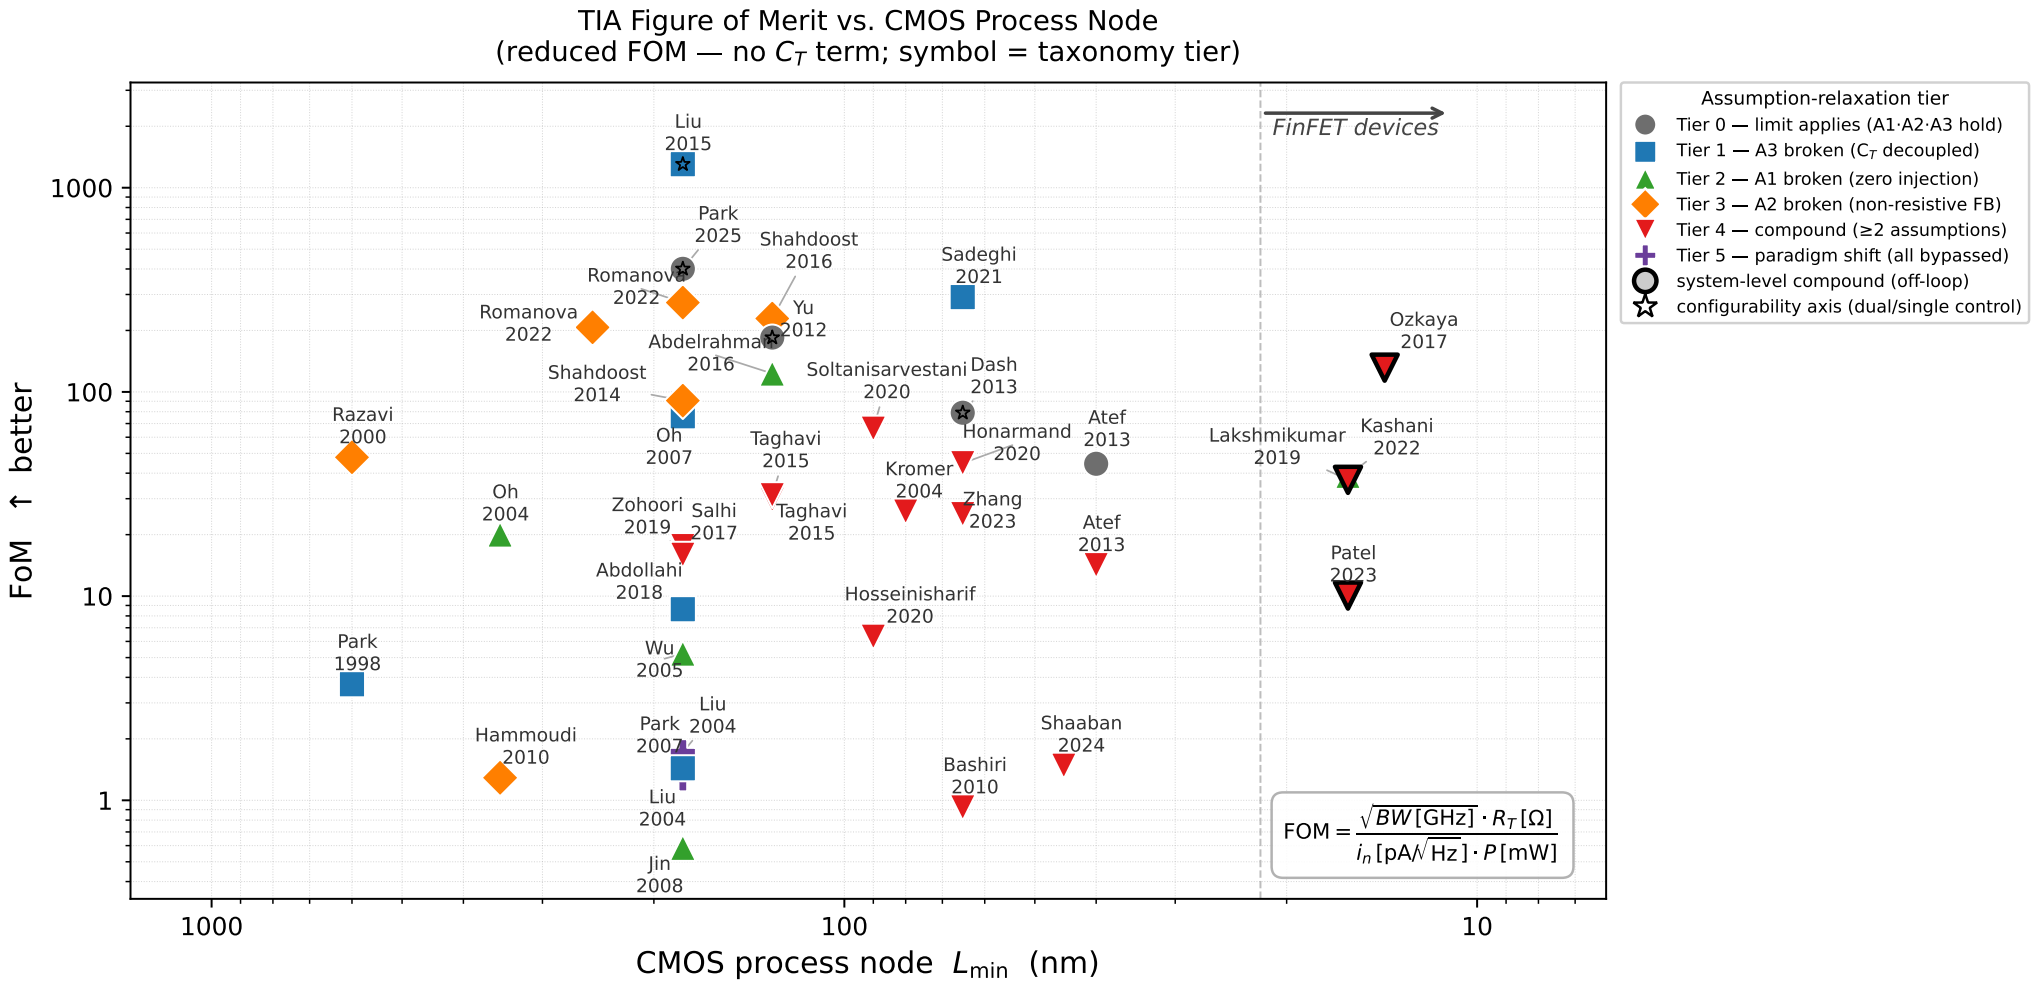

Figure 8. TIA FOM vs. CMOS process node for a set of representative designs. Entries include both measured and post-layout simulated data.

doi:10.20944/preprints202606.0753.v1

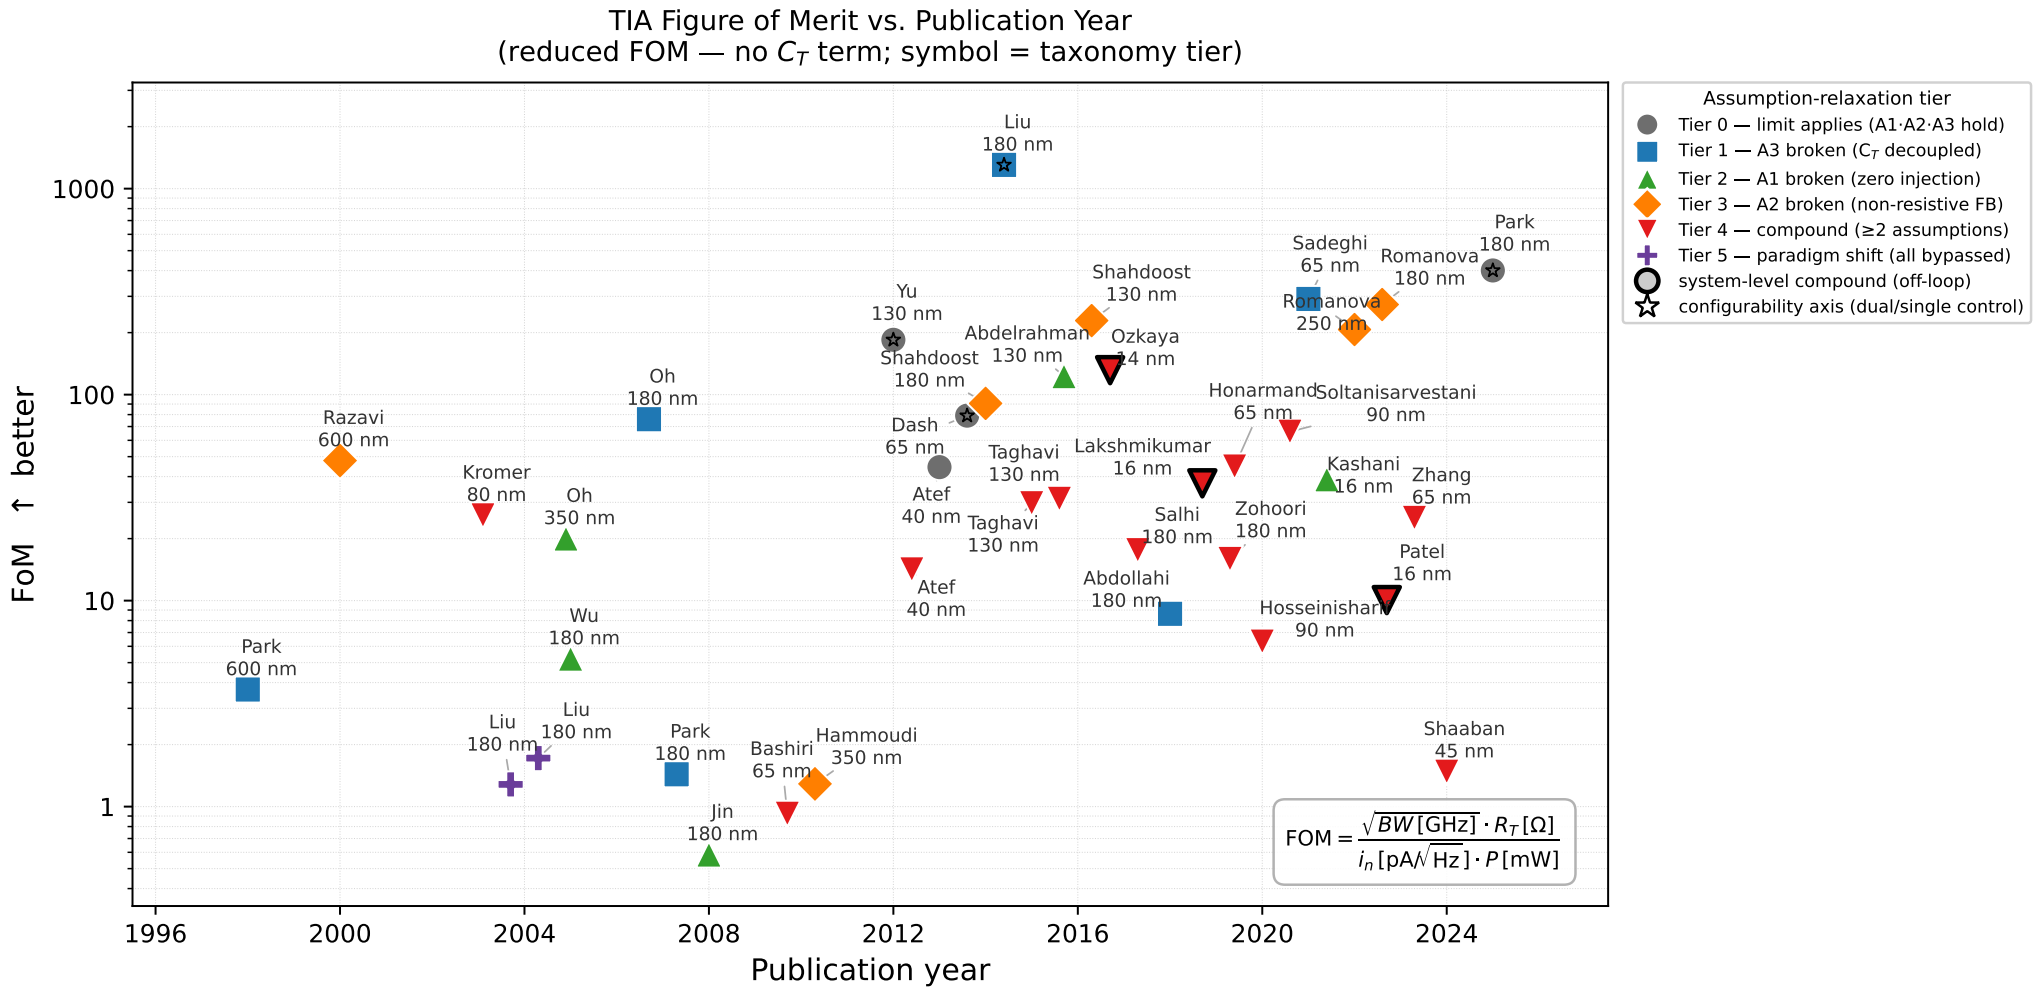

Figure 9. TIA FOM vs. publication year for a set of representative designs. Entries include both measured and post-layout simulated data. When multiple entries share the same publication year, a small horizontal jitter offset is added to prevent the markers from overlapping.

 doi:10.20944/preprints202606.0753.v1

10. Future Outlook & Challenges

The evolution of CMOS transimpedance amplifiers is continually shaped by physical device constraints and emerging system-level demands. Looking forward, the trajectory of TIA design is defined by three major frontiers: the push toward 200,Gb/s per lane in datacom, the integration of non-standard wide-bandgap sensors, and the physical transition to Gate-All-Around (GAA) devices. Across all three, the assumption-relaxation taxonomy developed in this review provides a predictive framework for understanding which topologies will remain viable.

10.1. Towards 200,Gb/s/lane and 1.6,T Links

The deployment of 800,GbE and 1.6,TbE in AI/ML clusters and hyperscale data centers has fixed 200,Gb/s,/lane (224,Gb/s including FEC overhead) as the near-term target, with 400,Gb/s,/lane research demonstrations already reported at OFC 2025. At the wireline end, fully-integrated 224,Gb/s PAM-4 transceivers have now been demonstrated in 3,nm FinFET with 3,pJ/b efficiency over a 40,dB insertion-loss channel [75], and DSP-based 212,Gb/s PAM-4 transceivers for optical direct-detect applications are operational in 5,nm FinFET [76]. For optical coherent transport, ISSCC 2024 reported a 200,GS/s 8-bit ADC-based coherent receiver with $>$60,GHz analog front-end bandwidth in 5,nm FinFET [77].

These designs reshape the TIA specification in three concrete ways. First, the linearity-bandwidth pair displaces the classical gain-bandwidth-noise triplet as the primary design driver. First, PAM-4 and PAM-6 modulations demand $<$1, THD across the Nyquist band to maintain usable equalized eye openings and group-delay variation tolerances tighten to a few picoseconds. Second, the analog front-end is no longer an isolated stage but a co-designed block with the CTLE, VGA (e.g., [5]), and ADC, blurring the boundary between the TIA and downstream DSP as the sampling rate approaches the TIA's -3,dB frequency. Third, the noise budget is increasingly dominated by quantization and CTLE-injected noise rather than the TIA's input-referred noise alone. Within the assumption-relaxation framework, these extreme requirements mandate compound Tier 4 and Tier 5 architectures. Single-relaxation designs (Tiers 0–3) are no longer sufficient at $\geq$200,Gb/s,/lane. Instead, successful front-ends must simultaneously co-deploy A3-relaxing input networks, A1-relaxing active/passive peaking, and DSP-assisted off-loop equalization.

10.2. Integration with Wide-Bandgap Sensors

Beyond traditional silicon PD-based optical receivers, a growing class of applications uses wide-bandgap (WBG) PDs such as GaN and AlGaN APDs for ultraviolet (UV) astronomy and flame detection, SiC APDs for deep-UV, and β-Ga_2O_3 PDs for high-temperature solar-blind imaging [78]. While their wide bandgaps ($\geq$ 3.4,eV for GaN, $\geq$ 4.8,eV for Ga_2O_3) suppress dark current, the authors impose readout constraints drastically different from datacom optics.

First, the PD operating bias can reach tens of volts (up to $\approx$80,V for GaN APDs), requiring the readout to either be implemented in a mature, high-voltage CMOS process (e.g., 0.35,μm HV-CMOS) or be galvanically isolated. This directly opposes the sub-1,V scaling trend of advanced optical receivers. Second, the single-photon-counting photocurrent levels (hundreds of picoamperes) strongly favor A2-relaxed capacitive TIA topologies (Tier 3) over the resistive-feedback designs dominant in high-speed datacom. Third, the dynamic-range requirements span 60-90,dB while needing to quench post-avalanche current surges, requires robust dual-control reconfigurable architectures. The hybrid GaN-APD,+,CMOS-CTIA array reported by [78] illustrates how these constraints can be met simultaneously. Importantly, this demonstrates also that the TIA assumption-relaxation framework remains universal where the architectural axis (which assumption is broken, and how) applies independently of the detector material, even as the bounding application specifications completely invert.

10.3. The Role of GAAFET in Future TIA Design

To understand the impact of the transition from FinFET to Gate-All-Around (GAA-NS) architectures at the 2-3,nm nodes, one must first recognize the structural limits reached by preceding

 doi:10.20944/preprints202606.0753.v1

generations. State-of-the-art 16,nm and 22,nm FinFET designs [5,6,79] have largely converged on system-level compound Tier 4 architectures. To overcome the finite intrinsic gain and voltage headroom limits of FinFETs, these high-speed front-ends rely on simple Tier 0 or Tier 3 TIA cores paired with off-loop A1 relaxations (such as programmable CTLEs or inductive peaking) and A3 input-matching networks (such as T-coils). While highly effective for 112–128,Gb/s PAM-4, these topologies highlight that device scaling alone is exhausted without compound architectural relaxations.

The GAA-NS transition reshapes this baseline through three intertwined device-level changes. First, the wrap-around gate yields near-ideal electrostatic control, reducing drain-induced barrier lowering and enabling sub-60,mV/decade subthreshold operation [80]. Second, unlike the discrete fin count in FinFETs, nanosheet width is continuously variable, restoring the fine analog sizing flexibility lost in the FinFET era. Third, fabricated GAA-NSH OTAs in a 22,nm-class flow exhibit roughly 4,dB higher intrinsic gain and 2.6$\times$ lower power than FinFET equivalents at the same supply [81]. For TIA design, this increased intrinsic gain is paramount as it directly mitigates the finite-A limitations that currently bottleneck modern Tier 3 (capacitive-feedback) and Tier 1 (RGC) topologies, making previously theoretical node approximations tractable in silicon.

However, these gains are counterbalanced by severe design challenges. Supply voltages continue to fall ($\leq$0.7,V at 2,nm), exacerbating the voltage-headroom problem. Traditional stacked cascodes are rendered virtually impossible, and even single-transistor common-source stages run out of headroom when biased for optimal g_m/I_D. Furthermore, larger parasitic source-drain capacitances per unit W and self-heating in the 5-6,nm thick nanosheets mean the effective in-loop TIA bandwidth will be noticeably smaller than the raw f_T suggests [81]. Consequently, the GAA transition will not permit a return to simple topologies, but rather the extreme voltage constraints will accelerate the migration toward aggressive Tier 4 and Tier 5 compound architectures, as designers are forced to relax assumptions structurally to survive the vanishing headroom.

11. Conclusions

The main result of this review is that the large variety of reported CMOS TIA designs reduces to a small number of basic options. The classical transimpedance limit rests on three assumptions, namely a single-pole core amplifier (A1), a resistive feedback element (A2), and the full input capacitance loading the feedback summing node (A3). It was found that every bandwidth- or noise-enhancement technique discussed in the literature relaxes one or more of these three assumptions, so that what is usually presented as a long list of separate circuit techniques is in fact a limited set of mechanisms and their combinations. Once the assumptions relaxed by a given topology are known, its place among all other designs is fixed. This view is consistent with the perspective established for early CMOS optical-communication circuits [3], and it is the sense in which the present work adds to the understanding of the individual designs it cites.

Two observations follow from this view that are not apparent when the designs are considered one by one. The first is that the relaxation of an assumption does not remove the stability constraint but moves it to a different place in the circuit. As designs progress from the classical feedback loop toward more advanced topologies, the stability requirement shifts from the global phase margin to a local regulating loop, then to a flat-group-delay condition, then to an input-passivity condition, and finally to the interaction between several simultaneous loops. An improvement in bandwidth or noise is therefore always obtained at the cost of a new stability condition that has to be satisfied. The second observation is that, because each of the three assumptions corresponds to a specific physical limitation, the proposed classification has a predictive value and is not only descriptive. Once the relaxed assumptions are known, the location of the main stability risk and the quantity that bounds the noise floor can be identified before a detailed circuit analysis is performed.

A further result concerns the comparison of the reported designs, which shows no clear improvement of the FOM with either the year of publication or the technology node (if one does not attempt to preselect the competing works). The technology node and the circuit topology behave as independent

 doi:10.20944/preprints202606.0753.v1

design parameters, and a well-chosen topology implemented in a mature 130-180 nm process can outperform a simpler design realized in a more advanced node. For this reason the FOM is best understood as a diagnostic indicator within a given bandwidth and application range, rather than as an absolute ranking across different applications, since a single number obtained from heterogeneously reported and partly simulated data cannot order designs intended for different purposes.

Finally, the review suggests that the basic analog structure of the CMOS TIA has remained largely unchanged over time. The 112-128 Gb/s PAM-4 front-ends implemented in 22 nm and 16 nm FinFET technologies are built from the same inverter-based shunt-feedback and inductive-peaking techniques that were already used for 0.13 μm, 40 Gb/s designs two decades earlier. The main progress has taken place around the amplifier, in the choice of the technology node, in digital equalization, and in the configurability that allows a single circuit to cover several operating points, rather than in the analog core itself. The classification proposed here provides a consistent way to describe the existing designs and to position the future ones intended for 200 Gb/s /lane links, wide-bandgap detector readouts, and the transition from FinFET to GAA devices.

11.1. Automated and Symbolic Topology Synthesis

A complementary strategy to device and architectural scaling is the automation of the topology-selection step itself. Where this review organizes TIA designs by which closed-form transimpedance limit assumption each relaxes, recent design automation work [27] from 2026 makes that closed-form structure machine-actionable on the example of TIA design task. The authors suggest the symbolic exploration framework to enumerate impedance-block permutations of a circuit template in order to derive exact transfer functions, and prune non-bandpass or unstable candidates by symbolic sign and pole analysis before any SPICE run. Then the survivors are handed to Bayesian or evolutionary sizing algorithms. The assumption-relaxation taxonomy proposed in this review, in effect, a hand-curated low-dimensional version of the same idea, where each tier corresponds to a structural transformation of the SFB transfer function, and the tier of a candidate topology is exactly the kind of closed-form property a symbolic engine can classify automatically. A natural extension, already sketched in the EDA literature, is an agentic flow in which the taxonomy supplies the prior while constraining the search to the assumption relaxation that matches a target specification.

Abbreviations

The following abbreviations are used in this manuscript:

5G/6G	5th/6th Generation
ACG	Advanced Common-Gate
ADC	Analog-to-Digital Converter
AGC	Automatic Gain Control
AI	Artificial Intelligence
APD	Avalanche Photodiode
ASIC	Application-Specific Integrated Circuit
BEOL	Back-End-of-Line
BiCMOS	Bipolar CMOS
BIST	Built-In Self-Test
BJT	Bipolar Junction Transistor
BW	Bandwidth
BWER	Bandwidth Extension Ratio
CB	Common-Base
CCII	Second-Generation Current Conveyor
CD	Common-Drain
CDM	Charged-Device Model
CDR	Clock and Data Recovery
CECB	Common-Emitter Common-Base

 doi:10.20944/preprints202606.0753.v1

CF	Capacitive Feedback
CFC	Conventional Folded-Cascode
CG	Common-Gate
CH	Cherry-Hooper
CIC	Cross-Coupled Immittance Converter
CMOS	Complementary Metal-Oxide-Semiconductor
CMRR	Common-Mode Rejection Ratio
CPO	Co-Packaged Optics
CPW	Coplanar Waveguide
CS	Common-Source
CTIA	Capacitive Transimpedance Amplifier
CTLA	CMOS Transimpedance-Limiting-Amplifier
CTLE	Continuous-Time Linear Equalizer
DAC	Digital-to-Analog Converter
DFE	Decision Feedback Equalizer
DR	Dynamic Range
DSP	Digital Signal Processing
DVI	Digital Visual Interface
DVS	Dynamic Voltage Scaling
ESD	Electrostatic Discharge
FC	Folded-Cascode
FCG	Feedforward Common-Gate
FD-SOI	Fully Depleted Silicon-on-Insulator
FEC	Forward Error Correction
FET	Field-Effect Transistor
FFE	Feed-Forward Equalizer
FinFET	Fin Field-Effect Transistor
FOM	Figure of Merit
GAA	Gate-All-Around
GAAFET	Gate-All-Around Field-Effect Transistor
GBW	Gain-Bandwidth Product
GDV	Group Delay Variation
HBM	Human-Body Model
HBT	Heterojunction Bipolar Transistor
HDMI	High-Definition Multimedia Interface
HV-CMOS	High-Voltage CMOS
IC	Integrated Circuit
ICDF	Inverter-based Common-Drain Feedback
ISI	Inter-Symbol Interference
LA	Limiting Amplifier
LDO	Low-Dropout Regulator
LiDAR	Light Detection and Ranging
MEMS	Micro-Electromechanical Systems
MEPR	Multi-Element Pseudo-Resistor
MFC	Modified Folded-Cascode
MOSFET	Metal-Oxide-Semiconductor Field-Effect Transistor
MRI	Magnetic Resonance Imaging
NC	Noise Cancellation
NF	Noise Figure
NIRS	Near Infrared Spectroscopy
NMOS	n-channel MOSFET
NMR	Nuclear Magnetic Resonance
NRZ	Non-Return-to-Zero
OPCB	Optical Printed-Circuit Boards

 doi:10.20944/preprints202606.0753.v1

OTA	Operational Transconductance Amplifier
OTDR	Optical Time-Domain Reflectometry
PAM-4	4-level Pulse Amplitude Modulation
PCB	Printed Circuit Board
PD	Photodiode
PGA	Programmable-Gain Amplifier
PIP	π-type Inductor Peaking
PLL	Phase-Locked Loop
PMOS	p-channel MOSFET
POF	Plastic Optical Fibre
PON	Passive Optical Network
PR	Pseudo-Resistor
PSIJ	Power Supply Induced Jitter
PSRR	Power Supply Rejection Ratio
PTAT	Proportional To Absolute Temperature
PVT	Process, Voltage, and Temperature
RF	Radio Frequency
RIC	Regulated Inverter Cascode
RGC	Regulated Cascode
SC	Switched-Capacitor
SFB	Shunt-Feedback
SiGe	Silicon-Germanium
SiPh	Silicon Photonics
SNR	Signal-to-Noise Ratio
SoC	System-on-Chip
SOI	Silicon-on-Insulator
SOTA	State-of-the-Art
SSF	Super Source Follower
STM	Scanning Tunneling Microscope
TCAD	Technology Computer-Aided Design
THD	Total Harmonic Distortion
TIA	Transimpedance Amplifier
TIS	Transimpedance Stage
UI	Unit Interval
UV	Ultraviolet
VCII	Second-Generation Voltage Conveyor
VCO	Voltage-Controlled Oscillator
VGA	Variable-Gain Amplifier
VLSI	Very Large-Scale Integration
WBG	Wide-Bandgap
WDR	Wide Dynamic Range

References

1. Razavi, B. A 622 Mb/s 4.5 pA/$\sqrt{(}Hz)$ CMOS transimpedance amplifier [for optical receiver front-end]. In Proceedings of the 2000 IEEE International Solid-State Circuits Conference. Digest of Technical Papers (Cat. No.00CH37056), Feb 2000, pp. 162–163. https://doi.org/10.1109/ISSCC.2000.839732.
2. Säckinger, E. *Analysis and Design of Transimpedance Amplifiers for Optical Receivers*; Wiley, 2017. https://doi.org/10.1002/9781119264422.
3. Razavi, B. Prospects of CMOS technology for high-speed optical communication circuits. *IEEE Journal of Solid-State Circuits* **2002**, *37*, 1135–1145. https://doi.org/10.1109/JSSC.2002.801195.
4. Säckinger, E. The Transimpedance Limit. *IEEE Transactions on Circuits and Systems I: Regular Papers* **2010**, *57*, 1848–1856. https://doi.org/10.1109/TCSI.2009.2037847.

 doi:10.20944/preprints202606.0753.v1

5. Patel, D.; Sharif-Bakhtiar, A.; Carusone, T.C. A 112-Gb/s -8.2-dBm Sensitivity 4-PAM Linear TIA in 16-nm CMOS With Co-Packaged Photodiodes. *IEEE Journal of Solid-State Circuits* **2023**, *58*, 771–784. https://doi.org/10.1109/JSSC.2022.3218558.
6. Lakshmikumar, K.R.; Kurylak, A.; Nagaraju, M.; Booth, R.; Nandwana, R.K.; Pampanin, J.; Boccuzzi, V. A Process and Temperature Insensitive CMOS Linear TIA for 100 Gb/s/ λ PAM-4 Optical Links. *IEEE Journal of Solid-State Circuits* **2019**, *54*, 3180–3190. https://doi.org/10.1109/JSSC.2019.2939652.
7. Dash, P.P.; Cowan, G.; Liboiron-Ladouceur, O. A variable-bandwidth, power-scalable optical receiver front-end in 65 nm. In Proceedings of the 2013 IEEE 56th International Midwest Symposium on Circuits and Systems (MWSCAS), 2013, pp. 717–720. https://doi.org/10.1109/MWSCAS.2013.6674749.
8. Oh, W.S.; Park, K. A 12-Channel 60-Gb/s Transimpedance Amplifier and Limiting Amplifier Array for OPCB Applications. In Proceedings of the 2007 14th IEEE International Conference on Electronics, Circuits and Systems, 2007, pp. 22–25. https://doi.org/10.1109/ICECS.2007.4510921.
9. Ozkaya, I.; Cevrero, A.; Francese, P.A.; Menolfi, C.; Morf, T.; Brändli, M.; Kuchta, D.M.; Kull, L.; Baks, C.W.; Proesel, J.E.; et al. A 64-Gb/s 1.4-pJ/b NRZ Optical Receiver Data-Path in 14-nm CMOS FinFET. *IEEE Journal of Solid-State Circuits* **2017**, *52*, 3458–3473. https://doi.org/10.1109/JSSC.2017.2734913.
10. Toumazou, C.; Park, S.M. Wideband low noise CMOS transimpedance amplifier for gigaHertz operation. *Electronics Letters* **1996**, *32*, 1194–1196. https://doi.org/10.1049/el:19960814.
11. Park, S.M.; Toumazou, C. Low noise current-mode CMOS transimpedance amplifier for giga-bit optical communication. In Proceedings of the ISCAS '98. Proceedings of the 1998 IEEE International Symposium on Circuits and Systems (Cat. No.98CH36187), May 1998, Vol. 1, pp. 293–296 vol.1. https://doi.org/10.1109/ISCAS.1998.704419.
12. Bashiri, S.; Plett, C.; Aguirre, J.; Schvan, P. A 40 Gb/s transimpedance amplifier in 65 nm CMOS. In Proceedings of the Proceedings of 2010 IEEE International Symposium on Circuits and Systems, 2010, pp. 757–760. https://doi.org/10.1109/ISCAS.2010.5537465.
13. Szilagyi, L.; Henker, R.; Ellinger, F. An inductor-less ultra-compact transimpedance amplifier for 30 Gbps in 28 nm CMOS with high energy-efficiency. In Proceedings of the 2014 IEEE 57th International Midwest Symposium on Circuits and Systems (MWSCAS), Aug 2014, pp. 957–960. https://doi.org/10.1109/MWSCAS.2014.6908575.
14. Sadeghi, S.; Nayeri, M.; Dolatshahi, M.; Moftakharzadeh, A. An active CMOS optical receiver employing an inductor-less, low-noise and high-gain regulated cascode transimpedance amplifier. *Microelectronics Journal* **2021**, *110*, 105015. https://doi.org/https://doi.org/10.1016/j.mejo.2021.105015.
15. Hosseinisharif, S.A.; Pourahmadi, M.; Shayesteh, M.R. Utilization of a cascoded-inverter in an RGC structure as a low-power, broadband TIA. *Microelectronics Journal* **2020**, *99*, 104749. https://doi.org/https://doi.org/10.1016/j.mejo.2020.104749.
16. Park, S.; Lee, S.; Seo, B.; Jung, D.; Choi, S.; Park, S.M. A Low-Noise CMOS Transimpedance-Limiting Amplifier for Dynamic Range Extension. *Micromachines* **2025**, *16*. https://doi.org/10.3390/mi16020153.
17. Liu, L.; Zou, J.; Ma, N.; Zhu, Z.; Yang, Y. A CMOS transimpedance amplifier with high gain and wide dynamic range for optical fiber sensing system. *Optik* **2015**, *126*, 1389 – 1393. https://doi.org/https://doi.org/10.1016/j.ijleo.2015.04.021.
18. Shahdoost, S.; Bozorgzadeh, B.; Medi, A.; Saniei, N. Low-noise transimpedance amplifier design procedure for optical communications. In Proceedings of the 22nd Austrian Workshop on Microelectronics (Austrochip), Oct 2014, pp. 1–5. https://doi.org/10.1109/Austrochip.2014.6946323.
19. Romanova, A.; Barzdenas, V. A Review of Modern CMOS Transimpedance Amplifiers for OTDR Applications. *Electronics* **2019**, *8*. https://doi.org/10.3390/electronics8101073.
20. Romanova, A. Design of Transimpedance Amplifiers for Broadband Time-Domain Optical Reflectometer Systems. Doctoral dissertation, Vilnius Gediminas Technical University, Vilnius, Lithuania, 2022. Technological Sciences, Electrical and Electronic Engineering (T 001).
21. Hu, J.; Kim, Y.B.; Ayers, J. A 65nm CMOS ultra low power and low noise 131M front-end transimpedance amplifier. In Proceedings of the Proceedings - IEEE International SOC Conference, SOCC 2010, 09 2010, pp. 281–284. https://doi.org/10.1109/SOCC.2010.5784645.
22. Kamrani, E.; Chaddad, A.; Lesage, F.; Sawan, M. Integrated Transimpedance Amplifiers Dedicated to Low-Noise and Low-Power Biomedical Applications. In Proceedings of the 2013 29th Southern Biomedical Engineering Conference, May 2013, pp. 5–6. https://doi.org/10.1109/SBEC.2013.11.

 doi:10.20944/preprints202606.0753.v1

23. Atef, A.; Atef, M.; Khaled, E.E.M.; Abbas, M. CMOS Transimpedance Amplifiers for Biomedical Applications: A Comparative Study. *IEEE Circuits and Systems Magazine* **2020**, *20*, 12–31. https://doi.org/10.1109/MCAS.2019.2961724.
24. Jiao, Y.; Li, W.; Zhang, D.; Wen, W. A 70 dBohm 2.3 GHz low noise transimpedance amplifier. In Proceedings of the 2013 IEEE Third International Conference on Consumer Electronics ¿ Berlin (ICCE-Berlin), 2013, pp. 322–325. https://doi.org/10.1109/ICCE-Berlin.2013.6697985.
25. Atef, A.; Atef, M.; Abbas, M.; Khaled, E.E.M. High-sensitivity regulated inverter cascode transimpedance amplifier for near infrared spectroscopy. In Proceedings of the 2016 Fourth International Japan-Egypt Conference on Electronics, Communications and Computers (JEC-ECC), 2016, pp. 99–102. https://doi.org/10.1109/JEC-ECC.2016.7518977.
26. Ghanad, M.; Dehollain, C. A sub 1 dB NF high dynamic range low-input impedance CMOS amplifier. In Proceedings of the 2016 IEEE 59th International Midwest Symposium on Circuits and Systems (MWSCAS), 2016, pp. 1–4. https://doi.org/10.1109/MWSCAS.2016.7870023.
27. Noori Zadeh, D.; Elamien, M.B. SymXplorer: Symbolic Analog Topology Exploration of a Tunable Common-Gate Bandpass TIA for Radio-over-Fiber Applications. *Electronics* **2026**, *15*. https://doi.org/10.3390/electronics15030515.
28. Salvia, J.; Lajevardi, P.; Hekmat, M.; Murmann, B. A 56MΩ CMOS TIA for MEMS applications. In Proceedings of the 2009 IEEE Custom Integrated Circuits Conference, Sep. 2009, pp. 199–202. https://doi.org/10.1109/CICC.2009.5280878.
29. Kumar, S.; Luhach, A.; Vanita, V. TIA Topologies used in CMOS-MEMS Resonators: A Detailed Review. *WCES-2022: Workshop on Control and Embedded Systems* **2022**, *3314*, 53–64.
30. Kromer, C.; Sialm, G.; Morf, T.; Schmatz, M.; Ellinger, F.; Erni, D.; Jackel, H. A low-power 20-GHz 52-dB/Ω transimpedance amplifier in 80-nm CMOS. *IEEE Journal of Solid-State Circuits* **2004**, *39*, 885–894. https://doi.org/10.1109/JSSC.2004.827807.
31. Badal, M.T.I.; Reaz, M.B.I.; Yeng, L.S.; Bhuiyan, M.A.S.; Haque, F. Advancement of CMOS Transimpedance Amplifier for Optical Receiver. *Transactions on Electrical and Electronic Materials* **2019**, *20*, 73–84. https://doi.org/10.1007/s42341-018-00092-5.
32. Haghi Kashani, M.; Shakiba, H.; Sheikholeslami, A. A Low-Noise High-Gain Broadband Transformer-Based Inverter-Based Transimpedance Amplifier. *IEEE Open Journal of Circuits and Systems* **2022**, *3*, 72–81. https://doi.org/10.1109/OJCAS.2022.3164396.
33. Jin, J.D.; Hsu, S.S.H. A 40-Gb/s Transimpedance Amplifier in 0.18-*mu*m CMOS Technology. *IEEE Journal of Solid-State Circuits* **2008**, *43*, 1449–1457. https://doi.org/10.1109/JSSC.2008.922735.
34. Jia Xu.; Zhenghao Lu. A 10 GB/s CMOS transimpedance amplifier with adaptive gain. In Proceedings of the 2011 International Conference on Electric Information and Control Engineering, 2011, pp. 5–7. https://doi.org/10.1109/ICEICE.2011.5777268.
35. Yu, G.; Zou, X.; Zhang, L.; Zou, Q.; zheng, M.; Zhong, J. A low-noise high-gain transimpedance amplifier with high dynamic range in 0.13um CMOS. In Proceedings of the 2012 IEEE International Symposium on Radio-Frequency Integration Technology (RFIT), 2012, pp. 37–40. https://doi.org/10.1109/RFIT.2012.6401606.
36. Zampardi, P.J. Will CMOS amplifiers ever Kick-GaAs? In Proceedings of the IEEE Custom Integrated Circuits Conference 2010, Sept 2010, pp. 1–4. https://doi.org/10.1109/CICC.2010.5617402.
37. Wu, C.H.; Lee, C.H.; Chen, W.S.; Liu, S.I. CMOS wideband amplifiers using multiple inductive-series peaking technique. *IEEE Journal of Solid-State Circuits* **2005**, *40*, 548–552. https://doi.org/10.1109/JSSC.2004.840979.
38. Atef, M.; Chen, H.; Zimmermann, H. 10Gb/s inverter based cascode transimpedance amplifier in 40 nm CMOS technology. In Proceedings of the 2013 IEEE 16th International Symposium on Design and Diagnostics of Electronic Circuits & Systems (DDECS), 2013, pp. 72–75. https://doi.org/10.1109/DDECS.2013.6549791.
39. Yong-Hun Oh.; Sang-Gug Lee. An inductance enhancement technique and its application to a shunt-peaked 2.5 Gb/s transimpedance amplifier design. *IEEE Transactions on Circuits and Systems II: Express Briefs* **2004**, *51*, 624–628. https://doi.org/10.1109/TCSII.2004.836883.
40. Jiao, Y.; Li, W.; Zhang, D.; Wen, W.; Zhang, Z. A low noise transimpedance amplifier for 2.5 Gb/s optical communications. In Proceedings of the 2013 25th International Conference on Microelectronics (ICM), 2013, pp. 1–4. https://doi.org/10.1109/ICM.2013.6734953.
41. Jin, J.D.; Hsu, S.S.H. 40-Gb/s Transimpedance Amplifier in 0.18-um CMOS Technology. In Proceedings of the 2006 Proceedings of the 32nd European Solid-State Circuits Conference, Sept 2006, pp. 520–523. https://doi.org/10.1109/ESSCIR.2006.307504.

 doi:10.20944/preprints202606.0753.v1

42. Yin, M.; Li, Z.; Zhang, W.; Liu, H.; Zhou, H.; You, Y.; Wang, C. Power, performance, and area evaluation across 180nm-28nm technology nodes based on benchmark circuits. *IEICE Electronics Express* **2024**, *21*, 20240194–20240194. https://doi.org/10.1587/elex.21.20240194.
43. Shahdoost, S.; Medi, A.; Bozorgzadeh, B.; Saniei, N. A novel design methodology for low-noise and high-gain transimpedance amplifiers. In Proceedings of the 2014 Argentine Conference on Micro-Nanoelectronics, Technology and Applications (EAMTA), July 2014, pp. 77–82. https://doi.org/10.1109/EAMTA.2014.6906083.
44. Singh, D.; Shankar, A.; Kumar, M. A Study on Transimpedance Amplifier in 0.35 um CMOS Technology. *International Journal of Computer Applications* **2012**, *51*, 4–6. Full text available.
45. Escid, H.; Salhi, S.; Slimane, A. Bandwidth enhancement for 0.18 um CMOS transimpedance amplifier circuit. In Proceedings of the 2013 25th International Conference on Microelectronics (ICM), Dec 2013, pp. 1–4. https://doi.org/10.1109/ICM.2013.6734945.
46. Salhi, S.; Escid, H.; Slimane, A. Design of high speed transimpedance amplifier for optical communication systems. In Proceedings of the 2017 Seminar on Detection Systems Architectures and Technologies (DAT), Feb 2017, pp. 1–5. https://doi.org/10.1109/DAT.2017.7889191.
47. Park, K.; Oh, W.S.; Choi, B.Y.; Han, J.W.; Park, S.M. A 4-channel 12.5 Gb/s Common-Gate Transimpedance Amplifier Array for DVI/HDMI Applications. In Proceedings of the 2007 IEEE International Symposium on Circuits and Systems (ISCAS), 2007, pp. 2192–2195. https://doi.org/10.1109/ISCAS.2007.378609.
48. Jin, J.D.; Hsu, S.S.H. A 75-dBOhm 10-Gb/s Transimpedance Amplifier in 0.18- *mu*m CMOS Technology. *IEEE Photonics Technology Letters* **2008**, *20*, 2177–2179. https://doi.org/10.1109/LPT.2008.2007812.
49. Atef, M.; Abd-elrahman, D. 2.5 Gbit/s compact transimpedance amplifier using active inductor in 130nm CMOS technology. In Proceedings of the 2014 Proceedings of the 21st International Conference Mixed Design of Integrated Circuits and Systems (MIXDES), 2014, pp. 103–107. https://doi.org/10.1109/MIXDES.2014.6872165.
50. Lima, E.N.; Cura, J.L.; Alves, L.N. Folded-cascode transimpedance amplifiers employing a CMOS inverter as input stage. In Proceedings of the 2013 European Conference on Circuit Theory and Design (ECCTD), 2013, pp. 1–4. https://doi.org/10.1109/ECCTD.2013.6662204.
51. Abd-elrahman, D.; Atef, M.; Wang, G. 10 Gb/s 1.95 mW active cascode transimpedance amplifier for high speed optical receivers. In Proceedings of the 2016 IEEE 59th International Midwest Symposium on Circuits and Systems (MWSCAS), 2016, pp. 1–4. https://doi.org/10.1109/MWSCAS.2016.7870132.
52. Rajabzadeh, M.; Djekic, D.; Haeberle, M.; Becker, J.; Anders, J.; Ortmanns, M. Comparison Study of Integrated Potentiostats: Resistive-TIA, Capacitive-TIA, CT Sigma-Delta Modulator. In Proceedings of the 2018 IEEE International Symposium on Circuits and Systems (ISCAS), 2018, pp. 1–5. https://doi.org/10.1109/ISCAS.2018.8351029.
53. Hammoudi, E.; Mokhtar, A. 2.75 GHz low noise 0.35 usingm CMOS transimpedance amplifier. In Proceedings of the Control Automation (MED), 2010 18th Mediterranean Conference on, June 2010, pp. 928–932. https://doi.org/10.1109/MED.2010.5547749.
54. Atef, M.; Zimmermann, H. 2.5Gbit/s transimpedance amplifier using noise cancelling for optical receivers. In Proceedings of the 2012 IEEE International Symposium on Circuits and Systems (ISCAS), 2012, pp. 1740–1743. https://doi.org/10.1109/ISCAS.2012.6271599.
55. Soltanisarvestani, R.; Zohoori, S.; Soltanisarvestani, A. A RGC-based, low-power, CMOS transimpedance amplifier for 10Gb/s optical receivers. *International Journal of Electronics* **2020**, *107*, 444–460, [https://doi.org/10.1080/00207217.2019.1661027]. https://doi.org/10.1080/00207217.2019.1661027.
56. Honarmand, S.; Zohoori, S.; Abbasi, K. A Low-power, CMOS Optical Communication Receiver System for 5Gbps Applications based on RGC Structure. *Majlesi Journal of Electrical Engineering* **2020**, *14*. https://doi.org/10.29252/mjee.14.4.57.
57. Taghavi, M.H.; Belostotski, L.; Haslett, J.W. A CMOS Low-Power Cross-Coupled Immittance-Converter Transimpedance Amplifier. *IEEE Microwave and Wireless Components Letters* **2015**, *25*, 403–405. https://doi.org/10.1109/LMWC.2015.2421253.
58. Hosseini, M.; Jafarnejad, R.; Sobhi, J.; Koozehkanani, Z.D. A low power wideband RGC-based modified-MIC trans-impedance amplifier in 0.18 um CMOS process. *Microelectronics Journal* **2020**, *96*, 104682. https://doi.org/https://doi.org/10.1016/j.mejo.2019.104682.
59. Zhang, J.; Wang, Z.; Ma, C. A CMOS transimpedance amplifier with broad-band and high gain based on negative Miller capacitance. *Integration* **2023**, *91*, 60–66. https://doi.org/https://doi.org/10.1016/j.vlsi.2023.03.004.

 doi:10.20944/preprints202606.0753.v1

60. Atef, M.; Zimmermann, H. Low-power 10 Gb/s inductorless inverter based common-drain active feedback transimpedance amplifier in 40 nm CMOS. *Analog Integr Circ Sig Process* **2013**.
61. Zohoori, S.; Dolatshahi, M.; Pourahmadi, M.; Hajisafari, M. An Inverter-Based, CMOS, Low-Power Optical Receiver Front-End. *Fiber and Integrated Optics* **2019**, *38*, 1–20, [https://doi.org/10.1080/01468030.2019.1567871]. https://doi.org/10.1080/01468030.2019.1567871.
62. Shaaban, A.; Sharroush, S.; Abdalla, Y.; Nafea, S. Design of a Novel CMOS Transimpedance Amplifier Based on a Current Mirror. *International Journal of Circuit Theory and Applications* **2025**, *53*, 4919–4934, [https://onlinelibrary.wiley.com/doi/pdf/10.1002/cta.4353]. https://doi.org/https://doi.org/10.1002/cta.4353.
63. Liu, R.C.; Wang, H. DC-to-15- and DC-to-30-GHz CMOS distributed transimpedance amplifiers. In Proceedings of the 2004 IEE Radio Frequency Integrated Circuits (RFIC) Systems. Digest of Papers, 2004, pp. 535–538. https://doi.org/10.1109/RFIC.2004.1320676.
64. Tang, Y.; Zhang, Y.; Fedder, G.K.; Carley, L.R. An ultra-low noise Switched Capacitor Transimpedance Amplifier for parallel Scanning Tunneling Microscopy. In Proceedings of the SENSORS, 2012 IEEE, 2012, pp. 1–4. https://doi.org/10.1109/ICSENS.2012.6411083.
65. Ferrari, G.; Gozzini, F.; Sampietro, M. A Current-Sensitive Front-End Amplifier for Nano-Biosensors with a 2MHz BW. In Proceedings of the 2007 IEEE International Solid-State Circuits Conference. Digest of Technical Papers, Feb 2007, pp. 164–165. https://doi.org/10.1109/ISSCC.2007.373345.
66. Micusik, D.; Zimmermann, H. 130dB-DR Transimpedance Amplifier with Monotonic Logarithmic Compression and High-Current Monitor. In Proceedings of the 2008 IEEE International Solid-State Circuits Conference - Digest of Technical Papers, 2008, pp. 78–597. https://doi.org/10.1109/ISSCC.2008.4523065.
67. Hsu, C.L.; Jiang, H.; Venkatesh, A.G.; Hall, D.A. A Hybrid Semi-Digital Transimpedance Amplifier With Noise Cancellation Technique for Nanopore-Based DNA Sequencing. *IEEE Transactions on Biomedical Circuits and Systems* **2015**, *9*, 652–661. https://doi.org/10.1109/TBCAS.2015.2496232.
68. Park, S.M.; Toumazou, C. Gigahertz low noise CMOS transimpedance amplifier. In Proceedings of the Circuits and Systems, 1997. ISCAS '97., Proceedings of 1997 IEEE International Symposium on, Jun 1997, Vol. 1, pp. 209–212 vol.1. https://doi.org/10.1109/ISCAS.1997.608671.
69. Hasaneen, E.S.A.M.; Wahab, M.A.A.; Okaley, N. Bandwidth Extension of CMOS Transimpedance Amplifier Using On-Chip Spiral Inductor. In Proceedings of the 2007 2nd International Design and Test Workshop, 2007, pp. 43–46. https://doi.org/10.1109/IDT.2007.4437425.
70. Romanova, A.; Barzdenas, V. On Noise Modeling of Capacitive Feedback Transimpedance Amplifiers in CMOS. *Applied Sciences* **2022**, *12*. https://doi.org/10.3390/app121910186.
71. Shahdoost, S.; Medi, A.; Saniei, N. A 1.93 pA/sqrt(Hz) transimpedance amplifier for 2.5 Gb/s optical communications. In Proceedings of the 2011 IEEE International Symposium of Circuits and Systems (ISCAS), May 2011, pp. 2889–2892. https://doi.org/10.1109/ISCAS.2011.5938235.
72. Shahdoost, S.; Medi, A.; Saniei, N. Design of low-noise transimpedance amplifiers with capacitive feedback. *Analog Integrated Circuits and Signal Processing* **2016**, *86*, 233–240. https://doi.org/https://doi.org/10.1007/s10470-015-0669-x.
73. Romanova, A.; Barzdenas, V. A design methodology for programmable-gain low-noise TIA in CMOS. *Journal of Electrical Engineering* **2021**, *72*, 147–157. https://doi.org/doi:10.2478/jee-2021-0021.
74. Abdollahi, B.; Mesgari, B.; Saeedi, S.; Nabavi, A. Stability analysis and compensation technique for low-voltage regulated cascode transimpedance amplifier. *Microelectronics Journal* **2018**, *71*, 37–46. https://doi.org/https://doi.org/10.1016/j.mejo.2017.11.009.
75. Pfaff, D.; Nummer, M.; Hai, N.; Xia, J.; Ge Yang, K.; Mohsenpour, M.M.; Leong, C.H.C.H.; LaCroix, M.A.; Zamanlooy, B.; Eeckelaert, T.; et al. A 224 Gb/s 3 pJ/bit 40 dB Insertion Loss Transceiver in 3-nm FinFET CMOS. *IEEE Journal of Solid-State Circuits* **2025**, *60*, 9–22. https://doi.org/10.1109/JSSC.2024.3466092.
76. Wang, J.Q.; Tan, A.; Iyer, A.; Fan, A.; Farhoodfar, A.; Alnabulsi, B.; Smith, B.; Loi, C.; Ho, C.R.; Cartina, D.; et al. 7.1 A 2.69pJ/b 212Gb/s DSP-Based PAM-4 Transceiver for Optical Direct-Detect Application in 5nm FinFET. In Proceedings of the 2024 IEEE International Solid-State Circuits Conference (ISSCC), 2024, Vol. 67, pp. 123–125. https://doi.org/10.1109/ISSCC49657.2024.10454275.
77. Nguyen, R.L.; Mellati, A.; Fernandez, A.; Iyer, A.; Fan, A.; Reyes, B.; Abidin, C.; Nani, C.; Albano, D.; Ahmad, F.; et al. 18.4 A 200GS/s 8b 20fJ/c-s Receiver with >60GHz AFE Bandwidth for 800Gb/s Optical Coherent Communications in 5nm FinFET. In Proceedings of the 2024 IEEE International Solid-State Circuits Conference (ISSCC), 2024, Vol. 67, pp. 344–346. https://doi.org/10.1109/ISSCC49657.2024.10454385.

 doi:10.20944/preprints202606.0753.v1

78. Padmanabhan, P.; Hancock, B.; Nikzad, S.; Bell, L.D.; Kroep, K.; Charbon, E. A Hybrid Readout Solution for GaN-Based Detectors Using CMOS Technology. *Sensors* **2018**, *18*. https://doi.org/10.3390/s18020449.
79. Daneshgar, S.; Li, H.; Kim, T.; Balamurugan, G. A 128 Gb/s, 11.2 mW Single-Ended PAM4 Linear TIA With 2.7 μA RMS Input Noise in 22 nm FinFET CMOS. *IEEE Journal of Solid-State Circuits* **2022**, *57*, 1397–1408. https://doi.org/10.1109/JSSC.2022.3147467.
80. Yeung, C.W.; Zhang, J.; Chao, R.; Kwon, O.; Vega, R.; Tsutsui, G.; Miao, X.; Zhang, C.; Sohn, C.W.; Moon, B.K.; et al. Channel Geometry Impact and Narrow Sheet Effect of Stacked Nanosheet. In Proceedings of the 2018 IEEE International Electron Devices Meeting (IEDM), 2018, pp. 28.6.1–28.6.4. https://doi.org/10.1109/IEDM.2018.8614608.
81. Sousa, J.C.; Perina, W.F.; Rangel, R.; Simoen, E.; Veloso, A.; Martino, J.A.; Agopian, P.G. Design of operational transconductance amplifier with Gate-All-Around Nanosheet MOSFET using experimental data from room temperature to 200 °C. *Solid-State Electronics* **2022**, *189*, 108238. https://doi.org/https://doi.org/10.1016/j.sse.2022.108238.