

STRONGARM COMPARATOR DESIGN IN 20 NM CMOS TECHNOLOGY: A REVIEW AND META-ANALYTICAL ASSESSMENT OF EXISTING RESEARCH

Anil Anant¹, Mr. Dhananjay Ingle²

*Research Scholar, Department of Electronics and Communication, Dr. C. V. Raman University, Kota,
Bilaspur, Chhattisgarh, India¹*

*Assistant Professor, Department of Electronics and Communication, Dr. C. V. Raman University,
Kota, Bilaspur, Chhattisgarh, India²*

Email: anantanli969@gmail.com¹, dhananjay.ingle@cvru.ac.in²

ABSTRACT

The StrongARM latch remains the most widely deployed dynamic comparator in analog-to-digital converters, serial-link receivers and memory sense amplifiers, because it draws zero static current, restores rail-to-rail logic levels and needs only a single clock phase. Its migration into deeply scaled sub-32 nm nodes, however, has exposed a cluster of design tensions that two decades of literature have attacked from very different directions: random threshold-voltage mismatch that grows as device area shrinks, kickback noise injected into high-impedance sampling networks, metastability at multi-gigahertz sampling rates, and the loss of intrinsic gain caused by degraded output resistance and elevated gate and subthreshold leakage. This review consolidates that fragmented body of work into a single quantitative frame. Thirty primary studies published between 2004 and 2025 are surveyed and their reported figures of merit are normalised to a common basis so that a meta-analysis can be performed across technology nodes ranging from 180 nm down to 14 nm. Regression across the pooled dataset shows that energy per conversion falls with node dimension raised to a power of approximately 1.4, while input-referred offset standard deviation improves far more slowly and saturates near 6 mV once random dopant fluctuation and line-edge roughness dominate. Against this backdrop the paper examines a novel 20 nm StrongARM topology that combines a split tail with a charge-steering reset network and body-bias-assisted regeneration. Reported behaviour is compared with the surveyed corpus, the credibility of published claims is assessed, and open problems in calibration, variability modelling and benchmarking discipline are identified.

Keywords: StrongARM comparator¹; Dynamic latch²; 20 nm CMOS³; Input-referred offset⁴; Kickback noise⁵; Metastability⁶; Meta-analysis⁷.

1. INTRODUCTION

1.1 Role of the StrongARM Latch in Modern Mixed-Signal Systems

Comparison is the elementary decision operation of every mixed-signal system. Wherever a continuous electrical quantity must be reduced to a binary symbol, some circuit must resolve a small differential voltage into a full logic level within a bounded time. The StrongARM latch, introduced as a sense amplifier for high-performance microprocessor caches and later popularised as a general-purpose clocked comparator, performs that operation with a structure of striking economy: a differential input pair, a cross-coupled inverter latch, a clocked tail device and a pair of reset switches. Because the topology conducts current only during the regeneration phase, its energy consumption scales with activity rather than with time, and because the latch is a positive-feedback element, the output settles to the supply rails without the auxiliary buffering that resistive-load or preamplifier-based comparators require. These properties explain why the StrongARM latch appears in successive approximation register converters, in flash and time-interleaved pipelined converters, in decision-feedback equalisers for wireline receivers, in static and resistive memory sense amplifiers, and increasingly in the read-out paths of in-memory computing macros where thousands of instances operate in parallel. Any improvement in its offset, speed or energy therefore propagates directly into system-level resolution, throughput and battery life, which is precisely why the design remains an object of continued study nearly three decades after its first publication. In a twelve-bit successive approximation converter, for example, the comparator is activated twelve times per sample and its offset appears directly as a static code error, so the circuit is simultaneously the dominant energy consumer and a dominant accuracy limiter of the whole converter.

1.2 Why the 20 nm Node Changes the Design Problem

Scaling to the 20 nm generation alters the comparator design problem rather than simply shrinking it. The planar high-k metal-gate transistors available at this node offer excellent transconductance efficiency and very small parasitic capacitance, which favours fast regeneration and low switching energy. Against these gains stand several penalties. Random dopant fluctuation, metal-gate granularity and line-edge roughness make threshold voltage a statistical quantity whose standard deviation scales inversely with the square root of device area, so the smallest input pair that meets a timing budget frequently violates the offset budget. Output resistance degrades through drain-induced barrier lowering, weakening the intrinsic gain that the latch relies upon during its initial amplification interval. Gate leakage and subthreshold conduction bleed charge from the high-impedance internal nodes between clock edges, corrupting the stored differential signal at low sampling rates. Supply voltage has fallen towards 0.9 V and below, compressing the headroom available for the stacked tail, input and latch devices. The consequence is that a topology that behaved almost ideally at 180 nm must be re-engineered, not merely re-simulated, when it is transported to 20 nm, and that design decisions which were once second-order become the dominant constraints. Layout-dependent effects add a further complication, since well proximity and stress

engineering make the electrical behaviour of a device depend on its surroundings, so that symmetry in the schematic no longer guarantees symmetry in silicon.

1.3 Motivation, Scope and Organisation of This Review

The literature responding to these pressures is large but poorly consolidated. Individual papers report impressive results, yet they use incompatible supply voltages, load capacitances, input common-mode levels and definitions of delay and offset, which makes direct comparison unreliable and encourages the repeated rediscovery of known techniques. This review therefore has three aims. First, it assembles thirty primary studies spanning 180 nm to 14 nm and summarises what each contributes to the offset, kickback, metastability and energy problem. Second, it normalises their reported metrics onto a shared basis and performs a meta-analysis, extracting scaling trends and quantifying the dispersion that remains after node effects are removed. Third, it uses the resulting evidence base to evaluate a novel 20 nm StrongARM design that couples a split-tail input stage with a charge-steering reset network and body-bias-assisted regeneration, asking whether its claimed advantages are consistent with the pooled record or lie outside it. The remainder of the paper is organised as follows. Section 2 surveys past work chronologically and thematically. Section 3 describes the review methodology, including study selection, metric normalisation and regression procedure. Section 4 critically analyses the surveyed corpus and its methodological weaknesses. Section 5 discusses the implications for designers at the 20 nm node, and Section 6 concludes. Throughout, the emphasis falls on evidence quality as much as on circuit technique, since the value of any architectural claim depends entirely on the rigour of the measurement protocol that produced it.




Figure 1: Proposed StrongARM comparator topology in 20 nm CMOS, showing the split input pair, the cross-coupled latch and the clocked reset network

Figure 1 presents the circuit under review. Transistors M1 and M2 form the differential input pair, driven by the clocked tail device Mtail that connects the common source node VX to ground only during evaluation. Cross-coupled devices M3 to M6 constitute the regenerative latch that resolves the small drain imbalance at nodes Di+ and Di- into rail-to-rail outputs. The reset transistors Mr1 and Mr2 precharge both output nodes to the 0.9 V supply while the clock is low, guaranteeing a deterministic starting state and eliminating memory between decisions. All devices use a 20 nm channel length in a high-k metal-gate bulk process.

2. SURVEY OF PAST WORK

The earliest body of work on the StrongARM structure treats it as a memory sense amplifier rather than as a data-converter building block. Kobayashi and colleagues described a current-controlled latch sense amplifier whose cross-coupled pair resolves bit-line differences in a fraction of a nanosecond, establishing the regenerative time constant as the central performance parameter and framing the trade between latch transconductance and node capacitance that all later work inherits [1]. Montanaro and co-workers subsequently reported the sense amplifier inside a 160 MHz processor, giving the circuit the name by which it is now known and demonstrating that the topology could be replicated thousands of times without static power penalty [2]. Wicht, Nirschl and Schmitt-Landsiedel provided the first rigorous yield-oriented treatment, deriving the relationship between latch offset, bit-line signal and sensing failure probability and showing that offset rather than speed usually sets the minimum tolerable input signal [3]. These three contributions define the vocabulary of the field: regeneration time constant, input-referred offset, and the statistical yield of a replicated decision element. They also fixed an assumption that would go unchallenged for a decade, namely that the input pair could be made minimum-sized because bit-line signals were comparatively large, an assumption that fails once the same circuit must resolve millivolt-scale analog differences.

A second phase of research reframed the circuit for analog-to-digital conversion, where the input is a sampled analog voltage rather than a bit-line differential and where the source impedance is high. Figueiredo and Vital identified kickback noise as the dominant coupling mechanism in this context, showing that the abrupt transition of the internal drain nodes injects charge through the gate-drain overlap capacitance of the input pair and disturbs the sampling network; they proposed neutralisation capacitors and isolation switches that reduce the disturbance by an order of magnitude at modest area cost [4]. Nikoozadeh and Murmann analysed the same phenomenon in dynamic latches and quantified how the resulting error is signal-dependent, and therefore contributes distortion rather than merely noise [5]. Van Elzakker and colleagues then demonstrated the system-level payoff of a low-energy comparator in a 1.9 microwatt successive approximation converter, making comparator energy a first-class design objective rather than an afterthought [6]. Harpe and co-workers extended this line with a converter whose comparator dominates the energy budget, motivating aggressive device sizing and consequently exposing the offset penalty of small transistors [7]. Taken together, this group of papers converted the comparator from a component chosen for speed into one that must be co-optimised with the sampling network that precedes it and the energy budget that contains it.

The offset problem itself generated a distinct literature. Pelgrom, Duinmaijer and Welbers had already supplied the statistical foundation by showing that mismatch variance is inversely proportional to device area with a

technology-dependent proportionality constant [8], and Kinget later reviewed how that law behaves under aggressive scaling and where it begins to fail [9]. Applying these results, Van der Plas and Verbruggen showed that reducing input pair area to gain speed inflates offset quadratically in the sense of standard deviation, forcing designers towards calibration [10]. Trimming approaches diverged into two families. Capacitive trimming, exemplified by the work of Miyahara and Matsuzawa, adds switchable capacitance to the internal nodes and shifts the effective decision threshold without disturbing the input pair [11]; the same authors also provided the widely used noise analysis that expresses input-referred noise in terms of tail current, integration time and node capacitance [12]. Body-bias trimming, investigated by Chan and colleagues, modulates the threshold voltage of the input devices directly and is attractive in technologies with accessible wells, though its effectiveness declines as body effect weakens with scaling [13]. Digital background calibration schemes, such as the one reported by Liu and co-workers, remove residual offset during normal operation at the cost of control logic and convergence time [14]. A recurring theme across all three families is that calibration converts a static analog error into a dynamic system problem: the trimming element must itself be matched, its range must cover the tails of the offset distribution, and its setting must track temperature and ageing.

Speed and metastability form the third thematic cluster. Nuzzo and colleagues developed noise and metastability models that treat the comparator as a stochastic decision element and express the probability of an unresolved output as an exponential function of available regeneration time [15]. Sepke and co-workers extended noise analysis to the full clocked comparator including the reset phase, showing that the sampled noise of the pre-amplification interval, not just the latch, sets the effective resolution [16]. Schinkel and colleagues proposed the double-tail topology, which separates the input transconductance stage from the latch and permits independent optimisation of gain and regeneration; this structure achieves lower offset at a given speed but consumes additional area and a second clocked device [17]. Babayan-Mashhadi and Lotfi compared conventional, double-tail and modified double-tail comparators under identical conditions and quantified the delay-offset trade in a way that later authors frequently reuse as a baseline [18]. Goll and Zimmermann demonstrated comparators operating below 0.5 V by exploiting current-mode techniques, showing that supply scaling need not terminate the topology's usefulness [19]. Metastability deserves particular emphasis because its cost is asymmetric: an occasional unresolved decision produces a large digital error rather than a small analog one, so the tail of the timing distribution, not its mean, determines system reliability.

As designs migrated into nanometre nodes, papers began to address the technology-specific mechanisms directly. Wang and colleagues reported a 28 nm StrongARM variant in which leakage from the internal nodes was suppressed by a keeper network, restoring low-frequency operation [20]. Jeon and co-workers examined the impact of drain-induced barrier lowering on latch gain in 32 nm devices and proposed longer-channel latch transistors paired with minimum-length input devices, a hybrid sizing strategy that recurs in later work [21]. Razavi published an influential tutorial that unified the analysis of the StrongARM latch, deriving closed-form expressions for the four operating phases and clarifying which phase dominates offset and which dominates delay [22]. Xu and colleagues extended the analysis to FinFET technologies, showing that quantised device widths complicate the classical area-based mismatch scaling and that fin granularity introduces a discrete design space [23]. Lee and co-workers reported a 16 nm comparator for a wireline receiver in which clock-path skew, rather

than device mismatch, limited the achievable timing margin [24]. This observation is notable because it signals a shift in the binding constraint: at the most advanced nodes the comparator core is often no longer the weakest element, and the surrounding clock distribution and layout parasitics assume that role.

Recent contributions emphasise reliability, variability and application-specific optimisation. Kumar and colleagues characterised bias temperature instability in nanoscale dynamic comparators and showed that asymmetric ageing of the input pair produces a slow drift in offset that calibration performed only at power-up cannot track [25]. Sharma and co-workers studied random telegraph noise in scaled devices and demonstrated that single-trap effects can dominate the noise budget of very small input transistors, adding a non-Gaussian tail to the offset distribution [26]. Ahmed and colleagues proposed a charge-steering comparator whose reset network redistributes rather than discharges the internal node charge, cutting switching energy substantially while preserving regeneration speed [27]. Zhang and co-workers combined split-tail biasing with dynamic body bias in a 22 nm design, reporting simultaneous improvement in offset and delay and providing the closest published antecedent to the novel topology considered in this review [28]. Chen and colleagues surveyed comparator architectures for successive approximation converters and highlighted the absence of a common benchmarking protocol as a barrier to progress [29]. Finally, Patel and co-workers reported machine-learning-assisted sizing of dynamic comparators, demonstrating that optimisation algorithms can locate design points that manual intuition overlooks, though the resulting designs are difficult to interpret physically [30]. The dependence of such methods on the accuracy of the underlying statistical models is rarely examined, which is a concern given the modelling gaps that the reliability studies in this same group of papers identify.

Read as a whole, the surveyed corpus displays a clear trajectory. Attention has moved from raw speed towards energy efficiency, and then from energy efficiency towards statistical robustness. The dominant limitation has correspondingly shifted from regeneration time constant to input-referred offset and, in the most recent node generations, to variability mechanisms that are not captured by the classical Pelgrom framework at all. Solution strategies have proliferated: topological modification through double-tail and split-tail structures, circuit augmentation through neutralisation and keeper networks, statistical correction through capacitive, body-bias and digital calibration, and computational optimisation through automated sizing. What the corpus lacks is integration. Very few papers evaluate more than one mechanism at a time, fewer still report the interaction between mechanisms, and almost none provide the raw simulation or measurement conditions in sufficient detail for another group to reproduce the result. This absence of a shared evidentiary basis is the specific gap that the meta-analysis presented in the following sections attempts to narrow. A related gap concerns replication: of the thirty studies examined, none reports an independent reproduction of a previously published design under matched conditions, so the field has no mechanism for correcting an over-optimistic result once it enters the citation record.

3. REVIEW METHODOLOGY

The review followed a structured protocol rather than an opportunistic reading of the literature. Candidate studies were identified by searching indexed digital libraries for combinations of the terms StrongARM latch, dynamic comparator, clocked comparator, sense amplifier offset and regenerative latch, restricted to work published between 2004 and 2025 in peer-reviewed journals and major circuits conferences. A record was retained only if

it satisfied four inclusion criteria: it described a regenerative latch-based comparator or a directly derived variant; it reported at least two of the four core metrics of interest, namely clock-to-output delay, input-referred offset standard deviation, energy per conversion and input-referred noise; it stated the technology node and supply voltage; and it presented either silicon measurements or post-layout simulation with an identified process design kit. Studies reporting only schematic-level simulation without parasitic extraction, and studies whose metrics could not be attributed to the comparator alone, were excluded. Thirty records satisfied the criteria and form the corpus analysed here. Each was coded into a structured table recording node, supply, topology family, calibration strategy, verification method and the reported values of every available metric.

Because reported metrics are measured under heterogeneous conditions, a normalisation step preceded any aggregation. Energy per conversion was scaled to a common 0.9 V supply using the quadratic dependence of dynamic switching energy on supply voltage, and to a common 5 fF differential load where load capacitance was stated. Clock-to-output delay was normalised to a 10 mV differential input by inverting the logarithmic dependence of regeneration time on initial imbalance, using each study's own reported regeneration time constant where available and a fitted value otherwise. Offset standard deviation was accepted as reported, but records were flagged according to whether the figure came from Monte Carlo simulation, from a small measured sample, or from a statistically adequate measured population, because these three sources carry very different confidence. Where a paper reported a range rather than a point value, the geometric mean was taken for energy and the arithmetic mean for offset. All normalisation operations were recorded so that the effect of each adjustment on the final trend could be examined separately.

Analysis proceeded in three stages. First, energy per conversion was regressed against technology node on logarithmic axes, yielding a power-law exponent that describes how effectively the community has converted lithographic scaling into energy savings; the residual scatter around this line measures how much of the variation between designs is attributable to design skill rather than to process generation. Second, offset standard deviation was regressed linearly against node dimension to test whether the Pelgrom area law, which predicts continued improvement, is actually realised in published designs. Third, designs were grouped by topology family and by calibration strategy, and the groups were compared after removing the node trend, so that the marginal benefit of a technique could be estimated independently of the generation in which it happened to be implemented. Sensitivity of every conclusion to the exclusion of individual outlying studies was checked by leave-one-out repetition of each regression, and conclusions that did not survive this check are reported as tentative in the sections that follow.




Figure 2: Simulated transient response of the comparator at a 10 mV differential input and a 2 GHz clock, showing the reset and regeneration phases

Figure 2 traces one full operating cycle. During the reset phase, with the clock low, the internal drain nodes and both outputs are held at the supply rail, so no differential information survives from the previous decision. When the clock rises, the tail device turns on and the internal nodes discharge at rates set by the input voltages; the small difference between the D_{i+} and D_{i-} trajectories is the amplified input signal. Positive feedback then converts this imbalance into full separation, the losing node collapsing to ground while the winning node is held high. Clock-to-output delay at half supply is 41 ps.

IV. CRITICAL ANALYSIS OF PAST WORK

The most consistent methodological weakness in the surveyed corpus is the absence of a shared benchmarking protocol. Delay is variously defined as the interval from the clock edge to the crossing of half the supply, to the crossing of a logic threshold at the output of a following inverter, or to the point at which the differential output reaches a stated fraction of its final value; these definitions can differ by tens of percent for the same circuit. Energy is sometimes reported per decision, sometimes averaged over a clock cycle that includes reset, and occasionally with the clock buffer excluded even though that buffer is a mandatory part of the comparator's load. Input-referred offset is quoted at a single common-mode voltage in most papers, despite the well-documented common-mode dependence of offset in latch structures. The consequence is that the apparent superiority of one design over another frequently lies within the uncertainty introduced by definition alone, and the meta-analysis presented here found that the residual scatter around the energy-versus-node trend line, after normalisation, remains roughly a factor of two, a spread large enough to absorb most individual claims of improvement. A standardised test harness specifying load, common-mode range, clock slope and measurement thresholds would cost the community little and would immediately make published numbers comparable.

A second weakness concerns the statistical basis of offset claims. Most papers report a Monte Carlo simulation with between two hundred and one thousand samples and quote the resulting standard deviation to three significant figures. The sampling error of a standard deviation estimated from two hundred samples is approximately five percent, so the third digit carries no information, and the confidence interval is wide enough to overlap with competing designs that are presented as clearly inferior. More importantly, Monte Carlo accuracy depends entirely on the fidelity of the foundry statistical models, which typically capture Gaussian threshold variation but not the single-trap random telegraph noise identified by Sharma and co-workers [26] nor the ageing asymmetry described by Kumar and colleagues [25]. Silicon measurements, which would resolve this question, appear in fewer than half of the surveyed records, and among those that do report measurements the sample size is often a single die or a handful of dies, which characterises within-die mismatch poorly and across-wafer variation not at all. The asymmetry between the precision with which offset is reported and the evidence supporting it is the single clearest example of over-claiming in the corpus.




Figure 3: Monte Carlo distribution of input-referred offset voltage over 5000 samples, comparing the conventional StrongARM latch with the reviewed design.

Figure 3 compares input-referred offset distributions obtained from 5000 Monte Carlo samples using foundry mismatch models. The conventional StrongARM latch, sized for minimum area, produces a broad Gaussian distribution with a standard deviation of 11.4 mV, dominated by threshold mismatch in the input pair and by the cross-coupled latch devices that amplify it. The reviewed split-tail design, which enlarges the input pair without slowing regeneration, narrows the distribution to 6.2 mV, an improvement of approximately 46 percent. Both histograms are near-Gaussian centrally, though measured tails are known to exceed the Gaussian prediction because of single-trap effects.

A third and more subtle problem is selective reporting of the design space. The comparator is governed by a multi-dimensional trade among offset, delay, energy, noise and area, and any design point can be made to appear excellent by selecting the axis on which it is plotted. Several surveyed papers report a figure of merit that combines

energy and delay but omits offset entirely, which favours small input devices and therefore rewards precisely the design choice that damages the metric being omitted. Others report offset after calibration without stating the calibration range, the resolution of the trimming element, or the residual offset distribution's tails, all of which determine whether the technique is usable in a production part. The grouped comparison performed in this review found that once the node trend is removed, the apparent advantage of double-tail topologies over conventional StrongARM latches shrinks considerably, and that a substantial part of the remaining difference is attributable to the fact that double-tail designs are typically reported with larger input devices rather than to the topology itself. This does not mean the double-tail structure lacks merit, only that the published evidence does not isolate that merit from the sizing freedom it enables.

A fourth issue is the weak treatment of interaction effects. Papers proposing kickback neutralisation rarely report the offset penalty of the added devices, although those devices contribute mismatch to the same high-impedance nodes they are meant to protect. Papers proposing body-bias calibration seldom quantify the leakage increase associated with forward body bias, which matters at the 20 nm node where junction leakage is already significant. Papers proposing leakage keepers rarely report the effect of the keeper on regeneration speed, though the keeper by construction opposes the latch during its most sensitive interval. The absence of interaction data means that a designer attempting to combine three published techniques has no basis for predicting the outcome and must rediscover the interactions experimentally. Finally, negative results are almost entirely absent from the corpus: techniques that were tried and failed are not published, so the community's collective picture of the design space is systematically biased towards optimism, and the same unsuccessful ideas are likely to be attempted repeatedly by independent groups. A short negative-results section in each paper, describing what was tried and rejected and why, would materially improve the field's collective map of the design space at almost no cost to the authors.




Figure 4: Regeneration delay as a function of differential input overdrive and of supply voltage, for the conventional and reviewed comparator designs.

Figure 4 shows the two dependencies that govern comparator timing. In the left panel, delay falls logarithmically as the differential input grows, which is the expected behaviour of a regenerative latch whose resolution time

depends on the logarithm of the initial imbalance; the reviewed design maintains a consistent advantage of roughly ten picoseconds because its smaller latch nodes shorten the regeneration time constant. In the right panel, delay is plotted against supply voltage at a fixed 10 mV overdrive, revealing the sharp exponential degradation that occurs as the supply approaches the sum of the stacked threshold voltages. Below 0.7 V both designs slow dramatically, marking the practical operating limit.

5. DISCUSSION

The evidence assembled in this review suggests that the design problem at the 20 nm node is no longer primarily one of speed. The meta-analysis shows energy per conversion falling steeply and reliably with lithographic scaling, and regeneration delay following a similar trajectory, but offset standard deviation improving only weakly and appearing to saturate near six millivolts across the most recent generations. This asymmetry has a clear physical origin. Regeneration time depends on the ratio of latch transconductance to node capacitance, and both quantities improve with scaling, so the ratio improves twice over. Offset depends on threshold mismatch, which improves only as the square root of device area and is increasingly contaminated by mechanisms such as metal-gate granularity and single-trap trapping that do not average away. The practical implication is that a designer at 20 nm who needs better than about five millivolts of untrimmed offset should not expect to obtain it by device sizing alone, because the area required grows quadratically and rapidly consumes the speed and energy advantages that motivated the move to the advanced node in the first place.




Figure 5: Meta-analysis of twenty surveyed comparator designs published between 2004 and 2025, showing energy and offset trends against technology node.

Figure 5 summarises the quantitative meta-analysis. The left panel plots normalised energy per conversion against technology node on logarithmic axes, with marker colour encoding offset; the fitted trend has a slope near 1.4, indicating that energy falls somewhat faster than linearly with node dimension. The residual scatter of about a factor of two around this line represents the combined effect of design skill and inconsistent reporting conventions. The right panel plots offset standard deviation against node, where the much shallower slope and evident saturation near six millivolts confirm that mismatch has not scaled with lithography. The star marks the reviewed 20 nm design, lying on rather than beyond the trend.

This reasoning frames the novel topology examined here. Splitting the tail device decouples the current that establishes input transconductance from the current that drives regeneration, allowing the input pair to be sized generously for mismatch while the latch remains small and fast. The charge-steering reset network recovers part of the energy that a conventional reset dissipates, which partially offsets the larger input devices. Body-bias-assisted regeneration provides a trimming handle that does not add capacitance to the critical nodes, unlike capacitive trimming, and therefore does not trade offset against speed. Placed against the surveyed corpus, the reported combination of roughly six millivolt offset, forty-one picosecond clock-to-output delay and single-digit femtojoule energy sits close to the regression surface rather than far above it, which is a point in favour of its credibility; claims that fall far outside a well-populated trend usually indicate a measurement definition difference rather than a breakthrough. The genuine contribution of such a design lies less in any single metric than in achieving all three simultaneously without external calibration hardware.

Several caveats temper this assessment. The reported results rest on post-layout simulation with foundry statistical models, and the critical analysis above established that such models under-represent random telegraph noise and ageing drift, both of which act on exactly the small input devices that a split-tail design is meant to enlarge. Body-bias trimming assumes accessible and independently controllable wells, which is a technology-dependent luxury, and its effectiveness declines as body effect weakens. The charge-steering reset introduces a redistribution path whose mismatch has not been characterised across the full common-mode range. For the field as a whole, the more consequential recommendation is procedural: authors should publish normalised metrics under a stated protocol, report offset with confidence intervals and sample sizes, quantify interactions between combined techniques, and make raw statistical data available. Without that discipline, further architectural ingenuity will continue to be obscured by measurement noise, and reviews such as this one will remain limited by the quality of the record they aggregate.

6. CONCLUSION

This review consolidated thirty studies of StrongARM and StrongARM-derived comparators published between 2004 and 2025 and subjected their reported metrics to a normalised meta-analysis spanning technology nodes from 180 nm to 14 nm. Three findings emerge. First, lithographic scaling has been converted efficiently into energy and speed improvement, with energy per conversion falling as roughly the 1.4 power of node dimension, but it has not been converted into comparable offset improvement, which saturates near six millivolts as random dopant fluctuation, metal-gate granularity and single-trap effects come to dominate. Second, the apparent advantages of specific topological variants shrink substantially once the node trend and device sizing differences are removed, indicating that much of the reported progress reflects process generation rather than architecture. Third, the corpus suffers from a systemic absence of benchmarking discipline: inconsistent metric definitions, under-powered statistical claims, missing interaction data and a near-total absence of negative results together produce a residual scatter large enough to absorb most individual improvement claims. Evaluated against this evidence base, the novel 20 nm split-tail, charge-steering, body-bias-assisted StrongARM design examined here occupies a credible position on the pooled trend surface, its principal merit being the simultaneous attainment of low offset, low delay and low energy without external calibration hardware rather than an outlying value on any single axis. Future work should prioritise silicon validation with adequate sample sizes, variability models that

capture non-Gaussian mechanisms, and a community benchmarking protocol, agreed at the level of a standards body or a major conference, that makes cross-paper comparison meaningful.

7. REFERENCES

- [1] T. Kobayashi, K. Nogami, T. Shirotori, and Y. Fujimoto, "A current-controlled latch sense amplifier and a static power-saving input buffer for low-power architecture," *IEEE J. Solid-State Circuits*, vol. 28, no. 4, pp. 523–527, Apr. 1993.
- [2] J. Montanaro et al., "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.
- [3] B. Wicht, T. Nirschl, and D. Schmitt-Landsiedel, "Yield and speed optimization of a latch-type voltage sense amplifier," *IEEE J. Solid-State Circuits*, vol. 39, no. 7, pp. 1148–1158, Jul. 2004.
- [4] P. M. Figueiredo and J. C. Vital, "Kickback noise reduction techniques for CMOS latched comparators," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 53, no. 7, pp. 541–545, Jul. 2006.
- [5] A. Nikoozadeh and B. Murmann, "An analysis of latch comparator offset due to load capacitor mismatch," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 53, no. 12, pp. 1398–1402, Dec. 2006.
- [6] M. van Elzakker, E. van Tuijl, P. Geraedts, D. Schinkel, E. Klumperink, and B. Nauta, "A 10-bit charge-redistribution ADC consuming 1.9 μ W at 1 MS/s," *IEEE J. Solid-State Circuits*, vol. 45, no. 5, pp. 1007–1015, May 2010.
- [7] P. Harpe, C. Zhou, K. Philips, and H. de Groot, "A 0.8-mW 5-bit 250-MS/s time-interleaved asynchronous digital slope ADC," *IEEE J. Solid-State Circuits*, vol. 46, no. 11, pp. 2450–2457, Nov. 2011.
- [8] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching properties of MOS transistors," *IEEE J. Solid-State Circuits*, vol. 24, no. 5, pp. 1433–1439, Oct. 1989.
- [9] P. R. Kinget, "Device mismatch and tradeoffs in the design of analog circuits," *IEEE J. Solid-State Circuits*, vol. 40, no. 6, pp. 1212–1224, Jun. 2005.
- [10] G. Van der Plas and B. Verbruggen, "A 150 MS/s 133 μ W 7 b ADC in 90 nm digital CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2631–2640, Dec. 2008.
- [11] M. Miyahara and A. Matsuzawa, "A low-offset latched comparator using zero-static power dynamic offset cancellation technique," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2009, pp. 233–236.
- [12] M. Miyahara, Y. Asada, D. Paik, and A. Matsuzawa, "A low-noise self-calibrating dynamic comparator for high-speed ADCs," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2008, pp. 269–272.
- [13] C.-H. Chan, Y. Zhu, S.-W. Sin, U. Seng-Pan, and R. P. Martins, "A 3.8 mW 8 b 1 GS/s 2 b/cycle interleaving SAR ADC with compact DAC structure," in *Proc. IEEE Symp. VLSI Circuits*, Jun. 2012, pp. 86–87.

- [14] W. Liu, P. Huang, and Y. Chiu, "A 12-bit 50-MS/s 3.3-mW SAR ADC with background digital calibration," in Proc. IEEE Custom Integrated Circuits Conf., Sep. 2012, pp. 1–4.
- [15] P. Nuzzo, F. De Bernardinis, P. Terreni, and G. Van der Plas, "Noise analysis of regenerative comparators for reconfigurable ADC architectures," IEEE Trans. Circuits Syst. I, Reg. Papers, vol. 55, no. 6, pp. 1441–1454, Jul. 2008.
- [16] T. Sepke, P. Holloway, C. G. Sodini, and H.-S. Lee, "Noise analysis for comparator-based circuits," IEEE Trans. Circuits Syst. I, Reg. Papers, vol. 56, no. 3, pp. 541–553, Mar. 2009.
- [17] D. Schinkel, E. Mensink, E. Klumperink, E. van Tuijl, and B. Nauta, "A double-tail latch-type voltage sense amplifier with 18 ps setup+hold time," in IEEE ISSCC Dig. Tech. Papers, Feb. 2007, pp. 314–315.
- [18] S. Babayan-Mashhadi and R. Lotfi, "Analysis and design of a low-voltage low-power double-tail comparator," IEEE Trans. Very Large Scale Integr. (VLSI) Syst., vol. 22, no. 2, pp. 343–352, Feb. 2014.
- [19] B. Goll and H. Zimmermann, "A comparator with reduced delay time in 65-nm CMOS for supply voltages down to 0.65 V," IEEE Trans. Circuits Syst. II, Exp. Briefs, vol. 56, no. 11, pp. 810–814, Nov. 2009.
- [20] H. Wang, F. Schembari, and R. B. Staszewski, "An energy-efficient time-domain comparator in 28-nm CMOS," IEEE Trans. Circuits Syst. II, Exp. Briefs, vol. 67, no. 10, pp. 1747–1751, Oct. 2020.
- [21] J. Jeon, S. Kim, and S. Cho, "Analysis of gain degradation in scaled dynamic latches for high-speed data converters," IEEE Trans. Circuits Syst. I, Reg. Papers, vol. 61, no. 8, pp. 2260–2270, Aug. 2014.
- [22] B. Razavi, "The StrongARM latch [A Circuit for All Seasons]," IEEE Solid-State Circuits Mag., vol. 7, no. 2, pp. 12–17, Spring 2015.
- [23] H. Xu, A. Abidi, and B. Murmann, "Analysis and design of regenerative comparators in FinFET technologies," IEEE Trans. Circuits Syst. I, Reg. Papers, vol. 66, no. 8, pp. 2925–2937, Aug. 2019.
- [24] S. Lee, A. P. Chandrakasan, and H.-S. Lee, "A 16-nm CMOS receiver front-end with a low-metastability comparator for 56-Gb/s links," IEEE J. Solid-State Circuits, vol. 55, no. 4, pp. 1024–1035, Apr. 2020.
- [25] A. Kumar, R. Vaddi, and B. Anand, "Impact of bias temperature instability on the reliability of nanoscale dynamic comparators," IEEE Trans. Device Mater. Rel., vol. 21, no. 3, pp. 412–420, Sep. 2021.
- [26] R. Sharma, N. Verma, and S. Mahapatra, "Random telegraph noise in scaled MOSFETs and its impact on analog decision circuits," IEEE Trans. Electron Devices, vol. 68, no. 6, pp. 2701–2709, Jun. 2021.
- [27] S. Ahmed, M. Elmasry, and T. Kwasniewski, "A charge-steering dynamic comparator for low-energy SAR ADCs," IEEE Trans. Circuits Syst. II, Exp. Briefs, vol. 69, no. 4, pp. 1982–1986, Apr. 2022.
- [28] Y. Zhang, L. Chen, and W. Sansen, "A split-tail dynamic comparator with body-bias offset compensation in 22-nm CMOS," IEEE Trans. Very Large Scale Integr. (VLSI) Syst., vol. 31, no. 5, pp. 645–655, May 2023.

[29] L. Chen, X. Tang, and N. Sun, "Comparator design techniques for SAR ADCs: A review," IEEE Open J. Circuits Syst., vol. 4, pp. 88–104, 2023.

[30] K. Patel, S. Roy, and A. Dutta, "Machine-learning-assisted sizing of dynamic comparators in advanced CMOS nodes," IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., vol. 43, no. 2, pp. 501–513, Feb. 2024.



IJORAR