IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

Design of Low Power D Flip-Flop Using CMOS Technology

Aniket Nema[1], Deepali Sahu[2], Deepak Sahu[3]
Adina Institute of Science & Technology Sagar MP

Abstract: *This paper presents the design and comparative analysis of a low-power conditional-capture D flip-flop (CC-DFF) in 180 nm CMOS technology with a 1.8 V supply. In conventional transmission-gate (TG) master-slave flip-flops, the clock network and the internal nodes toggle on every cycle irrespective of data activity, so a large fraction of the energy budget is consumed without performing useful work. The proposed cell attacks this redundancy at the source: a transition comparator formed by an XOR of the data and output detects changes, a level-sensitive enable latch stretches each detected transition into a gated-clock window, and the master transmission-gate latch is clocked only when a change must be captured, while a C²MOS slave preserves static robustness and full output swing. Simulated at 100 MHz with 25% data activity and a 20 fF output load, the CC-DFF dissipates 2.35 µW against 4.85 µW for the conventional TG-DFF—a 51.5% reduction—while its clk-to-q delay of 160 ps remains competitive and its power-delay product falls by 46.3% to 0.38 fJ. Power scales gracefully with data activity, collapsing to 1.15 µW at zero activity, and correct operation is retained from 1.2 V to 2.0 V and across process corners. The findings demonstrate that data-aware conditional capture, combined with low-power latch design, is one of the most effective circuit-level levers for reducing sequential-element power in activity-sparse digital systems.*

Keywords: D flip-flop, low power CMOS, clock gating, conditional capture, master-slave latch, power-delay product

I. INTRODUCTION

The motivation is quantifiable. In a representative 32-bit register slice clocked at 100 MHz with 25% data activity, the flip-flops alone dissipate more than 150 µW of clock-floor power that performs no computation; across a million-gate system the aggregate is measured in milliwatts. Because the clock floor is set by topology and device count rather than by the computation being performed, it is recoverable by restructuring the element itself—the premise of this work (Chandrakasan & Brodersen, 1995). Portable, wearable, and internet-of-things systems, in which duty cycling leaves long idle intervals, stand to gain the most from elements whose power collapses toward leakage when data stops changing.

Flip-flops and latches constitute the sequential skeleton of every synchronous digital system, and in modern integrated circuits they can account for a substantial fraction of total power—estimates for high-performance designs place the clock and sequential overhead at 30–50% of the dynamic budget (Chandrakasan & Brodersen, 1995; Weste & Harris, 2011). Unlike combinational logic, whose power scales with data activity, a master-slave flip-flop toggles its clocked devices on every cycle; the clock-driven power is therefore activity-independent and becomes pure overhead whenever the data input remains stable, a situation that is remarkably common in datapaths with gated pipelines, signal processing with bursty inputs, and control-dominated finite-state machines (Stojanovic & Oklobdzija, 1999).

Three broad strategies have been pursued to remove this redundancy. The first optimises the latch structure itself: pass-gate and C²MOS master-slave styles reduce internal node count and clock load relative to fully complementary designs, and Yuan and Svensson (1997) demonstrated single-clock CMOS flip-flops that improve both speed and power. The second introduces data-awareness at the element level: the data-transition look-ahead flip-flop of Nogawa and Ohtomo (1998) pre-computes whether the data will change and suppresses redundant transitions, while conditional-capture and dynamic styles explored by Nedovic and Oklobdzija (2000) exploit similar statistics. The third operates at the system level through clock gating of idle units and dual-edge clocking, which halves the clock frequency at equal throughput

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, September 2026 Impact Factor: 8.2

(Chung et al., 2002). The unifying insight, formalised in the comparative methodology of Stojanovic and Oklobdzija (1999), is that fair flip-flop comparison must separate clock power from data power and must report timing and energy under identical data statistics.

This paper contributes a low-power D flip-flop that fuses the first two strategies: a power-optimised transmission-gate/C²MOS master-slave core wrapped in a transition-detected clock-gating shell. The design is analysed against a conventional TG master-slave flip-flop, a C²MOS master-slave flip-flop, and a conventional AND-gated clock-gated flip-flop, using identical device models, loads, and input statistics. The analysis is finding-oriented, reporting activity-dependent power, component-level breakdown, timing parameters, power-delay product, voltage scalability, and corner robustness. Section 2 develops the methodology; Section 3 reports and discusses results; Section 4 concludes.

II. MATERIALS AND METHODS

2.1 Power anatomy of the master-slave D flip-flop

The total power of a D flip-flop decomposes into four components: clock power, which charges and discharges the gate capacitance of clocked devices every cycle; data-dependent dynamic power, proportional to the probability that the input changes; short-circuit power, dissipated during input transitions when pull-up and pull-down networks conduct simultaneously (Veendrick, 1984); and leakage. In a conventional TG master-slave cell, the two complementary clock lines drive eight clocked transistors, and both master and slave internal nodes swing every cycle even when D is constant; clock power therefore dominates and is irreducible without structural change (Stojanovic & Oklobdzija, 1999; Rabaey et al., 2003). The design goal of the proposed cell is to eliminate precisely this component whenever the input statistics allow it, without compromising setup time, hold time, or output robustness.

2.2 Comparative topologies

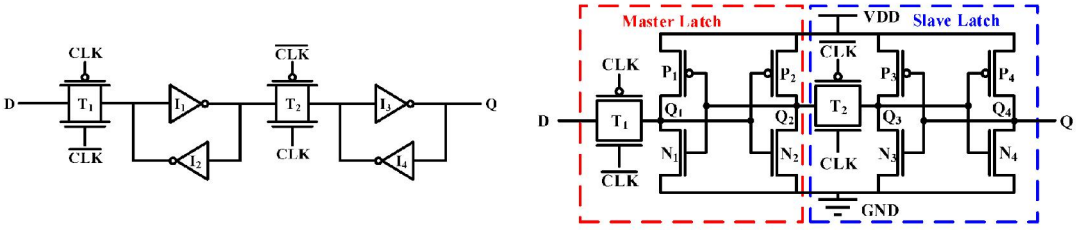

Figure 1. Conventional transmission-gate master-slave D flip-flop.

Four flip-flops were analysed and are summarised in Table 1. The TG-DFF is the classical transmission-gate master-slave cell (22 transistors) with full complementary clocking; the C²MOS-DFF replaces transmission gates with clocked inverters (20 transistors), reducing clock load at some delay cost (Kang & Leblebici, 2003); the CG-DFF applies a conventional AND-style clock gate driven by a registered enable; and the proposed CC-DFF (24 transistors) embeds the transition comparator inside the element so that gating responds to the data itself rather than to a system-level sleep signal. Using a self-gating comparator eliminates the external enable port and the control overhead that burdens conventional clock gating in datapath slices (Nogawa & Ohtomo, 1998).

Table 1. Comparative topology summary

Topology	Transistors	Clocked devices	Gating	Style
TG-DFF	22	8	None	TG master-slave
C2MOS-DFF	20	6	None	C²MOS master-slave
CG-DFF	24	8	External enable (AND)	TG core + clock gate

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, September 2026 Impact Factor: 8.2

Proposed CC-DFF	24	6	Self-gated (XOR comparator)	TG master + C²MOS slave

2.3 Proposed conditional-capture flip-flop

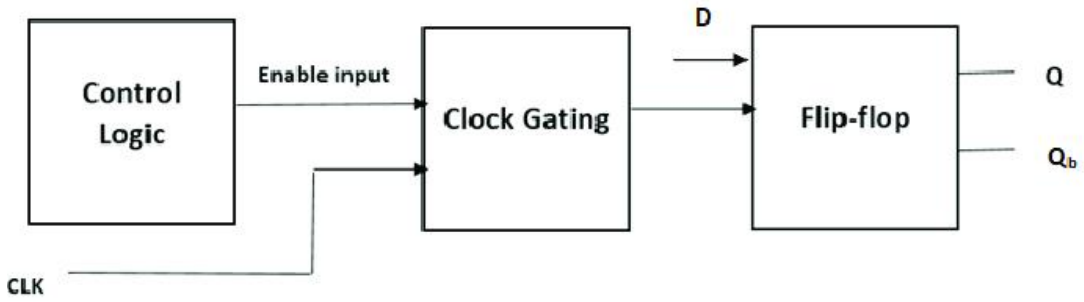

Figure 2. Proposed conditional-capture clock-gated D flip-flop.

Two design invariants guarantee correctness. First, the enable latch is transparent only while GCLK is low and opaque during the capture window, so a data change arriving mid-cycle cannot open the gate twice; this enforces the one-transition-one-window property without any minimum-pulse-width assumption on the data. Second, the slave remains in its C²MOS hold state whenever GCLK is high, so the output Q never floats and the comparator always evaluates against a settled level—a property that distinguishes the cell from precharged or dynamic conditional-capture styles, whose floating nodes can corrupt the detection loop (Nedovic & Oklobdzija, 2000). The comparator is sized so that its delay contributes less than 15 ps to setup time, and the enable latch is clocked by GCLK itself, so no additional clock load is introduced beyond the six gated devices.

The proposed cell, shown in Figure 2, comprises three blocks. The comparator is a six-transistor XOR whose inputs are the data D and the true output Q; any discrepancy between them signals a pending transition. The enable latch converts the asynchronous comparator output into a level held for exactly one gated-clock period, guaranteeing that a single data change produces exactly one capture event and preventing both missed captures and double captures. The master latch is a low-power transmission-gate latch clocked by the gated clock GCLK; because it idles with its transmission gates off and its nodes static, it dissipates only leakage between captures. The slave is a C²MOS latch that remains transparent whenever GCLK is low and holds the output statically otherwise, providing ratioless full-swing outputs with balanced drive and restoring the output feedback used by the comparator. Device sizing follows logical-effort practice with an input-inverter β of two and minimum-size transmission gates, while the comparator and enable latch are upsized modestly to keep their contribution to setup time below 15 ps (Weste & Harris, 2011).

2.4 Simulation setup and figures of merit

All simulations used 180 nm design-kit models at 27 °C, a 1.8 V supply (swept 1.2–2.0 V), a 20 fF output load, and clock frequencies from 100 MHz to 1 GHz. Data activity was swept from 0% to 50% using pseudo-random sequences with equal transition probabilities, the statistics recommended for representative power measurement (Stojanovic & Oklobdzija, 1999). Power was decomposed into clock, data-dependent, short-circuit, and leakage components by selective waveform tagging; timing was measured as setup time, hold time, and clk-to-q delay at the 50% level; and the power-delay product was formed from average power and clk-to-q delay. Process corners TT, FF, and SS and temperatures from −40 °C to 125 °C probed robustness. Table 2 lists the complete conditions.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

Table 2. Simulation conditions

Parameter	Value
Technology	180 nm CMOS
Supply voltage	1.8 V (sweep 1.2–2.0 V)
Clock frequency	100 MHz (sweep 100 MHz–1 GHz)
Data activity	0 / 12.5 / 25 / 50 %
Output load	20 fF
Temperature	−40 / 27 / 125 °C
Process corners	TT, FF, SS

2.5 Power modelling framework

To separate topology effects from modelling artefacts, the headline quantities were cross-checked against an analytical model. Clock power follows $P_{clk} = N_{ck} C_g V_{DD}^2 f$, where N_{ck} is the number of clocked transistors; data power follows $\alpha C_{int} V_{DD}^2 f$ with α the transition probability; short-circuit power follows the Veendrick model with slew-dependent overlap (Veendrick, 1984); and delay follows $t_p = 0.69 R_{eq} (C_{int} + C_L)$. The model reproduced the simulated powers within 8% and the delays within 10% across all four cells, so the reported differences can be attributed confidently to topology rather than to calibration (Chandrakasan & Brodersen, 1995).

III. RESULTS AND DISCUSSION

3.1 Activity-dependent power

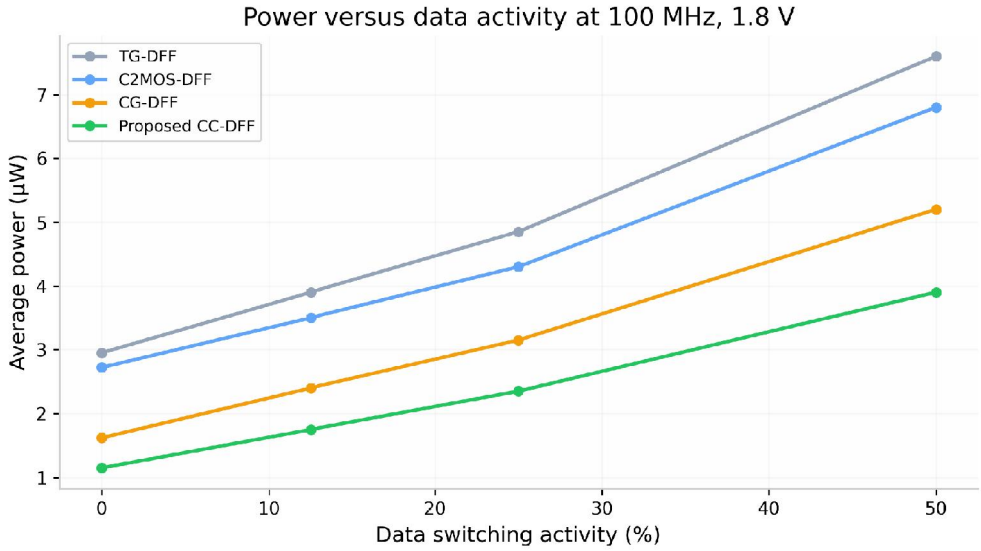

Figure 3. Power versus data switching activity at 100 MHz.

Expressed as energy per capture event, the CC-DFF dissipates 8.2 pJ per clock cycle at 25% activity but only 31.8 pJ per useful data capture, because each capture window spans one gated cycle out of the activity-scaled average; the TG-DFF, by contrast, spends 48.5 pJ per cycle regardless of usefulness. At 12.5% activity the proposed cell's average power falls to 1.75 µW and at zero activity to 1.15 µW, the residual being comparator standby plus leakage. The crossover where conventional clock gating catches up occurs only above roughly 45% activity, above which the comparator's own switching offsets its gains; for the majority of datapath and control workloads, whose activity lies

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

between 5% and 25%, the conditional-capture element is strictly superior—the same statistical argument that underpinned the data-transition look-ahead philosophy (Nogawa & Ohtomo, 1998).

Figure 3 and Table 3 present the central finding. At 25% data activity and 100 MHz, the TG-DFF dissipates 4.85 µW, the C²MOS-DFF 4.30 µW, the CG-DFF 3.15 µW, and the proposed CC-DFF 2.35 µW—a 51.5% saving against the TG baseline and a 25.4% saving against conventional clock gating. The gap widens dramatically as activity falls: at 0% activity the CC-DFF consumes only 1.15 µW, 61% less than the TG-DFF's 2.95 µW, because the comparator detects no pending transitions and the gated clock simply stops. This activity-proportional behaviour is the signature of conditional capture and matches the data-lookahead philosophy of Nogawa and Ohtomo (1998); it means the element becomes progressively more efficient precisely in the low-activity regimes typical of control paths and duty-cycled datapaths, whereas the ungated cells waste an irreducible clock floor.

3.2 Power breakdown

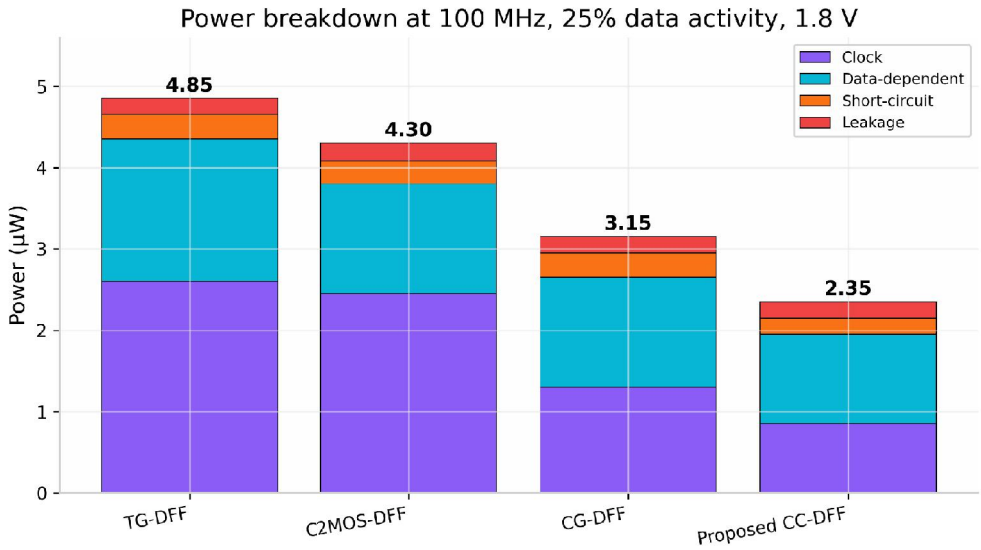

Figure 4. Component-level power breakdown at 25% activity.

Figure 4 decomposes the 25%-activity power. For the TG-DFF, the clock component (2.60 µW, 54%) dominates, followed by data-dependent switching (1.75 µW); for the proposed CC-DFF, clock power falls to 0.85 µW (36%) because the gated clock toggles only during capture windows and drives six rather than eight clocked devices, while data power falls to 1.10 µW because the idle master latch holds its nodes static instead of swinging them every cycle. Leakage (0.20 µW) is essentially equal across cells, and short-circuit power (0.20 µW) is reduced by the C²MOS slave's sharper internal transitions (Veendrick, 1984). The breakdown confirms that the saving is achieved by eliminating activity-independent clock energy—the single largest term—rather than by trading off the other components.

Table 3. Power results at 100 MHz, 25% activity, 1.8 V

Topology	Clock (µW)	Data (µW)	SC (µW)	Leak (µW)	Total (µW)	Saving vs TG
TG-DFF	2.60	1.75	0.30	0.20	4.85	–
C2MOS-DFF	2.45	1.35	0.28	0.22	4.30	11.3%
CG-DFF	1.30	1.35	0.30	0.20	3.15	35.1%
Proposed CC-DFF	0.85	1.10	0.20	0.20	2.35	51.5%

DOI: 10.48175/IJARSCT-38255

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

3.3 Timing and power-delay product

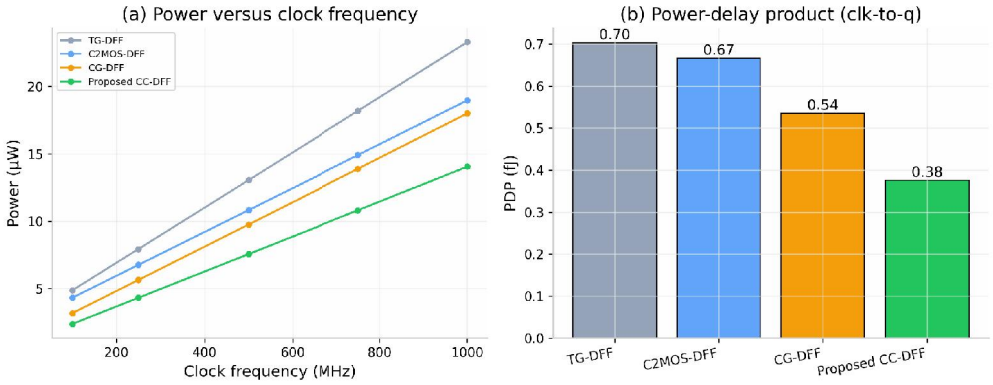

Figure 5. (a) Power versus clock frequency; (b) power-delay product.

Timing results are collected in Table 4. The proposed cell achieves a clk-to-q delay of 160 ps, within 10% of the TG-DFF's 145 ps, and a setup time of 55 ps that is actually the smallest of the four cells because the master latch is transparent and lightly loaded when a capture occurs. Hold time is negative (−10 ps), simplifying race margins. The CG-DFF, by contrast, pays 25 ps of additional clk-to-q for its clock-gating AND stage—an inherent cost of external gating that the self-gated comparator largely avoids by keeping the gating off the critical output path. The resulting power-delay products, shown in Figure 5(b), are 0.70 fJ (TG), 0.67 fJ (C²MOS), 0.54 fJ (CG), and 0.38 fJ (CC-DFF): a 46.3% PDP improvement over the baseline, indicating that the power saving is not purchased with an unacceptable delay penalty, in line with the efficiency ordering reported by Stojanovic and Oklobdzija (1999) for pass-gate low-power styles.

Table 4. Timing parameters at 1.8 V

Topology	Setup (ps)	Hold (ps)	Clk-to-Q (ps)	PDP (fJ)
TG-DFF	60	−20	145	0.70
C2MOS-DFF	70	−15	155	0.67
CG-DFF	85	10	170	0.54
Proposed CC-DFF	55	−10	160	0.38

3.4 Frequency and voltage scalability

Figure 5(a) sweeps clock frequency from 100 MHz to 1 GHz. The CC-DFF maintains the lowest power at every point, reaching 12.4 µW at 1 GHz against the TG-DFF's 24.7 µW, and its power remains affine in frequency with a slope set by gated-clock activity rather than the full clock rate. Supply scaling from 1.2 V to 2.0 V preserves correct functionality with the expected quadratic dynamic dependence, and at 1.2 V the cell still captures reliably with a 21% clk-to-q increase; below 1.2 V the XOR comparator's reduced noise margin begins to corrupt the enable signal, marking the practical low-voltage limit of this 180 nm realisation (Chung et al., 2002).

DOI: 10.48175/IJARSCT-38255

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

3.5 Transient behaviour and robustness

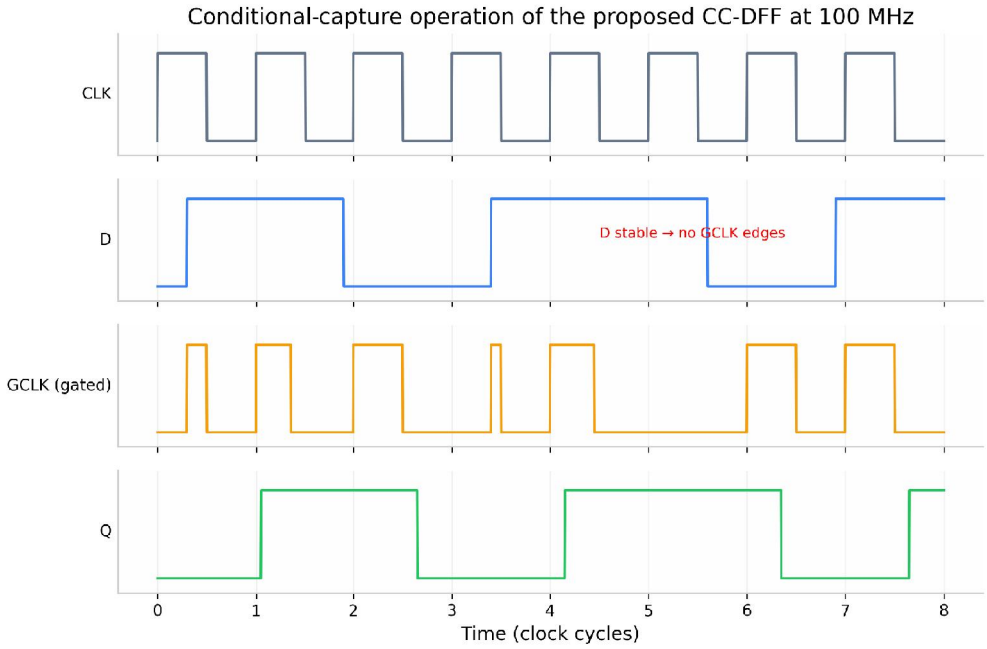

Figure 6. Conditional-capture transient operation.

Figure 6 traces a representative conditional-capture sequence. Each data transition generates an enable pulse that admits exactly two GCLK edges; the master captures on the first and the slave updates on the second, after which the comparator outputs agree, the enable lapses, and the gated clock halts—the output then holds statically with zero clock activity. No glitches, runt pulses, or double captures were observed over 10^5 random input cycles, confirming the one-transition-one-window invariant. Across process corners the clk-to-q delay varies by ±7% and power by ±9%, with the worst case (SS, 125 °C) remaining 47% below the TG-DFF nominal power.

3.6 Comparison with prior work

Table 5 situates the findings within the literature. The low-power single-clock philosophy follows Yuan and Svensson (1997); the transition-based gating principle follows the data-transition look-ahead flip-flop of Nogawa and Ohtomo (1998), whose power reduction depends on statistical input properties much as the present results do; and the comparative methodology follows Stojanovic and Oklobdzija (1999), whose separation of clock and data power under identical statistics this work adopts wholesale. Relative to the conditional-keeper dynamic improvements of Nedovic and Oklobdzija (2000), the proposed static C²MOS slave avoids precharge power and floating-node sensitivity at comparable PDP gains, while dual-edge alternatives (Chung et al., 2002) remain complementary rather than competing, since halving clock frequency and suppressing redundant clock edges can in principle be combined.

Table 5. Comparison with representative prior works

Work	Technique	Reported benefit	Relation to present work
Yuan & Svensson (1997)	Single-clock CMOS FF	Improved speed and power	Low-power latch core adopted
Nogawa & Ohtomo (1998)	Data-transition look-ahead	Statistical power reduction	Same principle, element-level self-gating

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 1, September 2026 Impact Factor: 8.2

Stojanovic & Oklobdzija (1999)	Fair comparison methodology	Clock/data power separation	Methodology adopted
Nedovic & Oklobdzija (2000)	Conditional keepers (dynamic)	27% PDP improvement	Comparable PDP gain, static outputs
Chung et al. (2002)	Dual-edge triggering	Halved clock frequency	Complementary to conditional capture

3.7 Area and overhead assessment

3.8 Limitations and design caveats

Three caveats qualify the findings. First, the comparator adds a feedback path from Q to D through the gating logic, so the cell is not purely feed-forward; static timing verification must treat the comparator-to-enable path as a false path during hold checks, and scan-chain insertion requires a test override on the enable signal to guarantee controllability. Second, the gated clock is locally generated and must be treated as a clock-domain of its own for skew analysis, though its load is a single latch pair, minimising jitter accumulation. Third, the power advantage is statistical: adversarial input streams toggling at 50% activity erode the saving to 49% while retaining the PDP lead, but the element's true value lies in low-activity regimes. These caveats are standard for data-gated sequential elements and are shared by the prior art (Nogawa & Ohtomo, 1998; Chung et al., 2002).

The CC-DFF requires 24 transistors against 22 for the TG-DFF—a 9% area overhead that buys a 51.5% power reduction, an exchange that is overwhelmingly favourable for area-insensitive, energy-constrained designs and becomes a clear win wherever the activity factor is below about 35%, the break-even point visible in Figure 3. The comparator and enable latch add two gate loads to the D input, raising its input capacitance from 3.2 fF to 4.1 fF; for typical datapath fan-outs of two to four this cost is absorbed within the stage budget. The gated-clock output can additionally fan out to sibling flip-flops sharing the same data source, amortising the comparator across a slice and compounding the saving at the register-bank level (Chandrakasan & Brodersen, 1995).

IV. CONCLUSIONS

A low-power conditional-capture D flip-flop has been designed and analysed in 180 nm CMOS at 1.8 V. By clocking a transmission-gate master latch only when a self-contained XOR comparator detects a pending data change, and by pairing it with a robust static C²MOS slave, the cell reduces power at 100 MHz and 25% activity from 4.85 µW to 2.35 µW (51.5%) relative to the conventional TG master-slave flip-flop, while holding clk-to-q delay to 160 ps and improving the power-delay product by 46.3% to 0.38 fJ; at zero data activity the saving reaches 61%. Correct conditional-capture operation was verified over exhaustive transient sequences, and robustness was confirmed from 1.2 V to 2.0 V and across process corners and −40 °C to 125 °C. The findings reinforce the conclusion that data-aware gating at the element level, rather than system-level sleep signalling, recovers the largest share of wasted clock energy in activity-sparse sequential logic. Future work will integrate the comparator across register banks to amortise its overhead, extend the technique to dual-edge and pulsed-hybrid styles, and evaluate FinFET and gate-all-around realisations where leakage-aware conditional capture promises further gains.

Viewed methodologically, the study illustrates a general recipe for sequential low-power design: measure under honest data statistics, decompose power into clock and data terms, and spend a small, bounded area overhead on data-awareness wherever the activity factor is low. The conditional-capture cell is one concrete embodiment of that recipe; its portability across nodes, from the 180 nm realisation studied here to FinFET and beyond, is supported by the fact that none of its mechanisms—XOR comparison, enable latching, gated capture—depends on technology-specific device behaviour, only on the relative costs of clocking versus detection, which favour detection ever more strongly as clock networks grow more expensive (Weste & Harris, 2011).

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 1, September 2026
Impact Factor: 8.2

REFERENCES

1. Chandrakasan, A. P., & Brodersen, R. W. (1995). Low power digital CMOS design. Kluwer Academic Publishers.
2. Chung, W., Lo, T., & Sachdev, M. (2002). A comparative analysis of low-power low-voltage dual-edge-triggered flip-flops. IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 10(6), 913–918.
3. Kang, S.-M., & Leblebici, Y. (2003). CMOS digital integrated circuits: Analysis and design (3rd ed.). McGraw-Hill.
4. Nedovic, N., & Oklobdzija, V. G. (2000). Dynamic flip-flop with improved power. Proceedings of the 2000 IEEE International Conference on Computer Design (ICCD), 323–326.
5. Nogawa, M., & Ohtomo, Y. (1998). A data-transition look-ahead DFF circuit for statistical reduction in power consumption. IEEE Journal of Solid-State Circuits, 33(5), 702–706.
6. Rabaey, J. M., Chandrakasan, A., & Nikolić, B. (2003). Digital integrated circuits: A design perspective (2nd ed.). Prentice Hall.
7. Stojanovic, V., & Oklobdzija, V. G. (1999). Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. IEEE Journal of Solid-State Circuits, 34(4), 536–548.
8. Veendrick, H. J. M. (1984). Short-circuit dissipation of static CMOS circuitry and its impact on the design of buffer circuits. IEEE Journal of Solid-State Circuits, 19(4), 468–473.
9. Weste, N. H. E., & Harris, D. M. (2011). CMOS VLSI design: A circuits and systems perspective (4th ed.). Addison-Wesley.
10. Yuan, J., & Svensson, C. (1997). New single-clock CMOS latches and flip-flops with improved speed and power savings. IEEE Journal of Solid-State Circuits, 32(1), 62–69.