

RX AMS Blind CDR with Oversampling and Elastic FIFO

For USB 2.0 Hi-Speed Receiver – 480 Mbps on SKY130

Tushar Goyal

tushar.goyal@berkeley.edu

Yash Kodali

yash.kodali@berkeley.edu

Abstract—This paper outlines a RX AMS Blind Clock and Data Recovery (CDR) circuit for USB 2.0 Full-Speed and High-Speed PHY, targeting both 12 and 480 Mbps operation using the SKY130 process. Traditional PLL- and DLL-based CDRs are vulnerable to long lock times, jitter, and PVT variation, which makes them less suitable for digitally dominant, portable PHY implementations. We implement a fully digital 5x oversampling-based CDR approach with a novel Add-Drop FIFO (AD-FIFO) that dynamically compensates for frequency offsets and input jitter by inserting or removing bits as needed. The recovered serial data stream is handed off to the RX Logic block. The design is synthesized in Chisel using the Chipyard SoC framework and evaluated through extensive simulations. Results demonstrate correct bit recovery under ± 1000 ppm frequency drift, validating the architecture’s robustness and suitability for USB PHYs.

The total area of the design is under $10,000 \mu\text{m}^2$, with the AD-FIFO consuming $6674 \mu\text{m}^2$, CRD consuming $2730 \mu\text{m}^2$, and oversampler consuming $328 \mu\text{m}^2$. Simulations confirm correct operation under ± 1000 ppm frequency drift, demonstrating the design’s robustness and suitability for integration in low-power digital USB PHYs.

Index Terms—blind clock and data recovery (CDR), digital oversampling, low power design, SKY130

I. INTRODUCTION

Reliable clock and data recovery (CDR) is essential in USB 2.0 Hi-Speed (480 Mbps) and Full-Speed (12 Mbps) receivers, where only the data signal is transmitted and no dedicated clock is embedded. This creates fundamental challenges in recovering timing information from the data stream, particularly in the presence of clock drift, jitter, and long runs of consecutive identical digits (CIDs) that degrade the timing margin.

Traditionally, USB receivers typically employ analog Phase-Locked Loops (PLLs) or Delay-Locked Loops (DLLs) for CDR. PLLs track the phase and frequency of incoming signals to align a recovered clock, while DLLs offer better jitter performance by manipulating delay stages. [4] However, both approaches require sensitive analog circuitry that is vulnerable to PVT (process, voltage, and temperature) variations, exhibit long lock times, and introduce significant design complexity. These limitations are problematic for highly integrated or portable systems that prioritize digital design, fast startup, and power efficiency. [2], [3]

An alternative approach is digital oversampling-based CDR, often referred to as "blind" CDR. This method oversamples

the data stream by a factor (e.g., 3x, 5x, or higher) without requiring phase-locked synchronization. The receiver can digitally select an optimal sample within each bit period, avoiding analog loops and reducing design risk.

This project presents a fully digital RX AMS CDR system based on 5x oversampling at 2.4 GHz to support both USB 2.0 Hi-Speed and Full-Speed modes. To improve robustness, an Add-Drop FIFO (AD-FIFO) is introduced to dynamically handle missing or duplicate bits caused by jitter and frequency offset. The design is implemented in Chisel and synthesized using the Hammer flow for the SKY130 process from Cadence, targeting low area, low power, and robust digital integration.

II. EXISTING SOLUTIONS

A. PLL Based CDR




Fig. 1. PLL-Based CDR

Phase-Locked Loop (PLL)-based CDRs are the standard architecture in many high-speed serial protocols. These systems rely on analog phase detectors, charge pumps, loop filters, and voltage-controlled oscillators (VCOs) to track the frequency and phase of the incoming signal. When locked, the recovered clock is used to sample data near the center of the data eye, minimizing bit errors. PLLs can achieve excellent jitter performance, but their reliance on precise analog tuning makes them sensitive to PVT (process, voltage, temperature) variation. Furthermore, the lock time of PLLs is non-negligible, which poses challenges for bursty or packetized protocols such as USB 2.0. [4], [5]

B. DLL Based CDR




Fig. 2. DLL-Based CDR

Delay-Locked Loop (DLL)-based CDRs address some limitations of PLLs by adjusting discrete delay lines rather than oscillator frequency. They provide finer control of phase resolution and are often used in interfaces requiring tight jitter margins, such as DDR. However, DLLs still require analog phase detection and delay-calibration circuitry and do not eliminate the need for multiphase clock generation. These circuits also remain susceptible to supply and temperature variations, and their complexity makes them less feasible for area-constrained USB PHY integration. [4], [5]

C. Digital Overampling CDR with AD-FIFO




Fig. 3. BDR-Based CDR

In blind CDR, the incoming serial data stream is sampled at a fixed multiple (e.g. 3x-8x) of the target bit rate using a free-running high-speed clock. At each bit period, a digital phase selector chooses the most reliable sample, ideally near the center of the data eye. Since this method is feedforward and fully digital, it offers fast lock times (zero or one bit period) and is highly robust to PVT variation.

In this work, we implement a blind CDR design with 5x oversampling. A key enhancement is the addition of an Add-Drop FIFO (AD-FIFO) that dynamically inserts or removes bits in response to accumulated timing mismatch caused by frequency offset or jitter. This prevents the accumulation of timing errors and ensures accurate data recovery. [1]

The key benefits of this approach include:

- Jitter Tolerance: Oversampling inherently provides robustness against timing jitter, as multiple samples per bit period allow for accurate edge detection.
- Simplified Clocking: Utilizing a high-speed, free-running clock for sampling reduces the complexity associated with phase-locked clock recovery.
- Digital Processing: The oversampled data can be processed using digital logic, simplifying the overall system design and potentially reducing power consumption.

- Error Correction: The AD-FIFO resolves errors due to frequency mismatch, improving long-run stability.

However, this method necessitates the generation and management of a high-frequency sampling clock (e.g., 2.4 GHz), which can pose challenges in terms of clock distribution and power consumption in other parts of the device.

III. DESIGN SPACE EXPLORATIONS

A. 5x vs 8x Oversampling

Oversampling rate is a central design tradeoff. Higher oversampling (e.g., 8x) improves temporal resolution for edge detection, providing finer phase granularity and wider sampling margin. However, this improvement comes at a cost: logic toggle activity increases substantially, impacting dynamic power and timing closure. Higher oversampling also necessitates a faster PLL or reference clock source, and deeper datapaths to store and process multiple samples per bit. For example, at 8x oversampling, a 480 Mbps data stream demands a 3.84 GHz clock and significant pipeline depth.

Importantly, diminishing returns are observed beyond 5x oversampling—particularly when using a center-picking method for bit recovery, where 5x is often sufficient for reliable performance. This configuration balances detection accuracy with system power and area constraints, while enabling robust bit-level recovery across moderate jitter and drift.

B. Majority Voting vs Center-Picking Method

Two main strategies exist for selecting the correct sample from the oversampled window: majority voting and center-picking. Majority voting selects the most common value among samples, but cannot track phase drift across bit periods. This makes it vulnerable to frequency offset.

In contrast, center-picking detects transition edges and aligns sampling to the middle of each bit eye, dynamically adjusting phase. This approach is more resilient to drift and jitter.

C. Depth of Add/Drop FIFO

The depth of the AD-FIFO is another critical dimension. A shallow FIFO (e.g., 8 cells) offers minimal area and latency but limited drift tolerance, risking underflow or overflow. A deeper FIFO (e.g., 25 cells, as used in our design) accommodates large packet sizes and drift margins (e.g., ± 1000 ppm), and ensures that the token does not reach FIFO boundaries under worst-case conditions. However deeper FIFOs also incur higher costs in register count, routing complexity, and longer warm-up durations.

We decide the FIFO depth based on:

- Maximum expected packet length (e.g., 8255 bits in USB2.0)
- Worst-case allowable frequency offset (± 500 ppm)
- Bit error tolerance due to jitter and transition ambiguity

$$N = (\text{Max packet size}) \times \Delta f_{\text{offset}} + N_{\text{margin}} \quad (1)$$

Here, Δf_{offset} represents the maximum allowed frequency offset (e.g., ± 500 ppm for USB 2.0) [3], and N_{margin} provides additional slack to absorb jitter and guard against insert/drop overshoot. For a USB packet length of 8255 bits and $\Delta f = 500 \times 10^{-6}$, this yields:

$$N = 8255 \times 0.0005 + 8 \approx 12 \text{ cells} \quad (2)$$

In our implementation, we conservatively select a depth of 25 cells to provide capability to handle N drops of inserts at once. [1]

D. Clamping vs. Soft Recovery in AD-FIFO Control

When the AD-FIFO token approaches its buffer limits, the control logic needs to decide whether to halt corrective operations or attempt speculative recovery. A clamping strategy halts insert/drop activity at buffer boundaries, preventing further corruption but potentially allowing drift-induced bit skew to propagate. However, a soft recovery strategy instead could continue operation by injecting fixed or inferred bits. This minimizes protocol interruption, but risks transient bit errors. Our implementation currently assumes clamping behavior as the depth of the AD-FIFO can already tolerate the max USB 2.0 frequency offset and to not introduce any incorrect data into the bitstream passed off to RX Logic.

E. Hysteresis in Insert/Drop Control

One potential strategy to improve noise resilience in insert/drop signaling is to incorporate hysteresis into the phase selection logic. This would require that a phase drift be sustained over multiple sampling windows before triggering a correction. It would help avoid overreaction to transient jitter that might not represent true frequency offset.

However, in our implementation, hysteresis is intentionally not applied. Blind CDR for USB 2.0 must respond rapidly to clock drift and jitter, particularly in short packets with limited timing margin. Introducing hysteresis would delay correction and allow drift-induced misalignment to accumulate, potentially corrupting multiple bits before recovery begins. As such, we opt for immediate response to transition-detected insert/drop events. This helps ensure fast tracking even at the cost of potential overcorrection in marginal cases.

F. Bit-Level vs. Byte-Level Recovery

Conventional oversampling CDRs often recover data in byte-wide chunks using wide FIFOs and demux logic. This increases silicon area and power. Our design recovers data at the bit level using a 5x sampling scheme and token-controlled bit-level AD-FIFO, which significantly reduces register count and simplifies integration with serial RX logic.

IV. SUMMARY OF SPECS

- **Data Rates Supported:** 12 Mbps (Full-Speed), 480 Mbps (Hi-Speed)
- **Oversampling Rate:** 5x (2.4 GHz sampling clock)
- **Sampling Clock Domain:** 2.4 GHz (oversampler), 480 MHz (CRD and AD-FIFO)

- **Jitter Tolerance:** Up to ± 1000 ppm frequency offset correction
- **Bit Error Rate (BER):** $< 10^{-12}$
- **FIFO Depth:** 25 cells
- **Total Core Area:** $9969 \mu\text{m}^2$

- AD-FIFO: $6674 \mu\text{m}^2$
- CRD: $2730 \mu\text{m}^2$
- Oversampler: $328 \mu\text{m}^2$

- **Technology:** SKY130

V. DETAILED DESIGN METHODOLOGY




Fig. 4. Top Level Design

The final architecture was selected after exploring various oversampling ratios, FIFO depths, and sampling strategies, as discussed in Section IV. The resulting design consists of three primary blocks arranged in a feed-forward pipeline:

- 1) **Oversampler** (5x at 2.4 GHz): Captures a 5-sample window of each incoming bit.
- 2) **Coarse Data Recovery (CDR)**: Identifies transitions, selects the best phase, and detects timing anomalies.
- 3) **Add-Drop FIFO (AD-FIFO)**: Dynamically corrects for frequency drift by adding or removing bits in the recovered stream.

These blocks operate across two clock domains: the oversampler is driven at 2.4 GHz, while the CDR and AD-FIFO are clocked at 480 MHz. The recovered serial data output is synchronized for downstream RX Logic processing.

A. Oversampler Design

The oversampler captures 5 samples per bit period using a 2.4 GHz clock. A 5-stage D flip-flop (DFF) shift register is used, with each DFF sampling the input in sequence due to propagation delay.

Implementation Details:

- Instantiate five DFFs in series, forming a 5-stage shift register.
- Clock all flip-flops with a 2.4 GHz sampling clock.
- Output samples as `samples[0]` to `samples[4]`, with `samples[0]` being the newest sample

B. Coarse Data Recovery Design




Fig. 5. Coarse Data Recovery Design

- 1) **Edge Detector:** Detects bit transitions.
- 2) **Phase Selector:** Analyzes bit transitions and decides which phase to sample. Moreover, it outputs the ADD/DROP signals for the AD-FIFO.
- 3) **Data Sampler:** Selects the appropriate bit according to the information of the phase.

1) *Edge Detector:* The purpose of this block is to detect signal transitions and track whether the current sampling point (determined by the phase_pointer) is early or late. This allows us to determine the optimal sampling point within the 5-sample window.

A transition is identified using XOR operations between the last sample of the previous window and all samples from the current window. The transition information gets stored in a flag register.

Implementation Details:

```
if samples = prev then flag = flag_prev
else f = Fill(5, prev[4])
flag = samples  $\oplus$  f
```

2) *Phase Selector:* The selector's goal is to track these transitions and reposition the sampling point toward the most stable region of the bit period—typically the center of the eye diagram—where the likelihood of bit errors is minimal.

The implementation follows a center-picking heuristic that maps observed transition locations to empirically chosen sampling positions. Specifically, the module inspects a 5-bit vector flag that encodes where a transition occurred in the most recent sample window. If an edge is detected at position 0 (the earliest sample), the selector chooses position 2 for subsequent sampling. Similarly, an edge at position 1 maps to position 3, and so on. This mapping is designed to realign the sampling point as close as possible to the midpoint of the stable region following the transition.

The mapping is implemented using a priority-encoded conditional logic structure that checks each bit of flag in order and updates the selector output (sel) accordingly. If no edge is detected, the previous sampling phase is retained to ensure temporal consistency.

Implementation Details:

```
if edge at position 0 then sel = 2
else if edge at position 1 then sel = 3
else if edge at position 2 then sel = 4
else if edge at position 3 then sel = 0
else if edge at position 4 then sel = 1
else sel = previous_sel
```

Edge Cases: In ambiguous patterns such as 11011 or 00100, transitions appear on both ends, causing multiple detected edges. To mitigate this, we apply a voting mechanism across adjacent windows. If there occur two edge transitions in the same sampling window. The CDR chooses the bit with greater than or equal to 3 occurrences in the window.

3) *Data Sampler:* Once the optimal phase is tracked, a 5-to-1 multiplexer selects the corresponding sample from the samples[0:4] array. By selecting the sample closest to the ideal phase, the data sampler minimizes timing distortions, such as jitter, that could cause errors at the bit boundaries. This process significantly improves the reliability of data recovery, ensuring the integrity of the transmitted signal.

Implementation Details:

- Use a 5-input MUX with the select line driven by phase_pointer.
- The output is the recovered serial data bit.
- This MUX and phase pointer are clocked at 480 MHz (bit rate), not at 2.4 GHz.
- The output is latched in a DFF to ensure clean transitions for downstream logic.

C. Add-Drop FIFO Design




Fig. 6. AD-FIFO State Transition Diagram [1]

To handle frequency offset, the coarse data stream may contain duplicated or missing bits. The AD-FIFO corrects this by dynamically adjusting the data flow based on the insert/drop signals issued by the CDR. This functionality is essential in blind CDR systems where oversampling occurs using a free-running clock without tight phase tracking.

In USB 2.0 high-speed mode, packet lengths can reach up to 8255 bits, and the receiver must tolerate a frequency offset of ± 500 ppm relative to the transmitter. Without correction, errors accumulate and corrupt the decoded data stream.

The AD-FIFO mitigates these effects using a circular bit-level buffer and a one-hot token pointer. The token determines

which FIFO cell is currently output to downstream RX logic. Depending on the insert or drop request, the token either advances or retreats through the buffer, effectively adding or removing a bit without disrupting the FIFO ordering.

1) *FIFO Architecture:*

- **Depth:** 25 entries
- **Data Structure:** A shift-register FIFO
- **Token:** An index (register) indicating the output cell.
- **Clock Domain:** Entire AD-FIFO operates at 480 MHz, synchronized to the bit rate.

2) *Data Flow and Token Behavior:* At each clock cycle, the FIFO performs a left-shift operation to accept a new bit from the Coarse Data Recovery (CDR) stage. Three cases determine how the shift and token update behave:

1) **Normal Operation:**

- Data is shifted in from the left.
- Token remains unchanged.
- FIFO acts like a classic shift register.

2) **Drop Case:** Triggered when the same bit is detected twice (i.e., transmitter is too fast).

- The input bit is shifted in normally.
- The token is decremented by 1 (unless already at zero).
- This skips the output of a duplicated bit.

3) **Insert Case:** Triggered when a bit is missed (i.e., receiver too fast).

- Enter a multi-stage state machine (Insert1, Insert2).
- Insert1: The inverse of the upcoming bit is inserted. The token is incremented (unless at max).
- Insert2: Resume normal shifting with the next input bit.




Fig. 7. AD-FIFO State Transition Diagram

3) *Insert-Driven FSM Design:* The insert operation requires two sequential clock cycles to fully handle the data injection. The FSM consists of the following states:

- **sIdle:** Normal state. Handles drop or nominal shift. Transitions to **sInsert1** if insert is asserted.
- **sInsert1:** Inserts a dummy bit (computed as the inverse of the upcoming bit). Token is incremented.

- **sInsert2:** Shifts in the true incoming bit. Returns to **sIdle**.

Token Clamping: To avoid out-of-bounds behavior, the token is clamped at both ends:

- $\text{token} = 0 \rightarrow \text{Drop}$ does not decrement.
- $\text{token} = \text{depth}-1 \rightarrow \text{Insert}$ does not increment.

This ensures stability even under sustained drift.

VI. RESULTS

A. Verification

To validate correctness and robustness, we developed an extensive set of testbenches targeting both unit-level and full-system behavior. Each component of the blind CDR pipeline—oversampler, coarse data recovery (CDR), and the Add-Drop FIFO (AD-FIFO)—was tested under nominal and worst-case conditions, including simulated clock drift, jitter, and protocol edge cases.

• Oversampler

– The oversampler was validated by injecting known bit patterns and checking the resulting 5-sample output window. The window updated every cycle, and each sample within the window was verified to be temporally ordered and logically correct with respect to the injected serial stream. Tests confirmed that a correct oversampled window was retrieved for each bit

• CDR

The CDR block was tested using oversampled input windows that simulate phase drift and transition ambiguity. Three main categories of tests were applied:

– **Transition Detection Accuracy:** Inputs like 0b11110, 0b00001, and 0b10001 were used to check that transitions are correctly flagged and mapped to sampling phases.

– **Phase Selector Logic:** Validated that the selector updates its sampling index based on the edge position and properly retains the prior index when no transition is present.

– **Edge Case Handling:** Inputs like 0b11011 and 0b00100 were used to stress-test multi-edge and symmetric transitions. In these cases, the design correctly applies a soft voting mechanism across adjacent windows to stabilize sampling.

– **Insert/Drop Signal Generation:** We verified that the CDR module emits correct insert or drop signals when drift is emulated by introducing early or delayed transitions over multiple windows.

• Add-Drop FIFO (AD-FIFO):

The AD-FIFO was validated using several targeted test cases in Chisel to ensure correct functionality and boundary safety.

– **Normal Operation Test:** Sequential bits were fed into the FIFO and observed from the midpoint token to confirm basic shift and pass-through behavior.

- **Drop Case:** Simulated a duplicated bit by asserting the drop signal mid-stream. Verified that the FIFO correctly skipped a value and realigned the output.
- **Insert Case:** Triggered an insert cycle and validated that a dummy bit (inverse of upcoming bit) was inserted, and the token advanced as expected.
- **Mixed Flow:** Combined insertions, drops, and nominal shifts to confirm multi-cycle transitions across FSM states. Output was compared against ground-truth bit streams.
- **Clamp Behavior - Insert Upper Bound:** Drove the token to the end of the buffer and triggered inserts. Verified that the token was clamped and no overflow occurred.
- **Clamp Behavior - Drop Lower Bound:** Forced token to buffer start and asserted drops. Verified clamping and safe behavior under extreme frequency offset.
- **Output Initialization Delay:** Confirmed that the output does not reflect valid recovered bits until the FIFO is sufficiently primed, to ensure proper warm-up latency.

- **Integration Testing**

- Simulate end-to-end flow from AFE input to RX Logic handoff with a PLL reference clock.
- Perform error injection (bit errors, jitter, clock glitches) and validate system recovery, similar to CRD-level testing.
- Evaluate overall system resilience under worst-case timing and jitter conditions (e.g., ± 1000 ppm).
 - * Three scenarios are tested:
 - 1) Input clock is slower than expected (-1000 ppm)
 - 2) Input clock is exactly matching (0 ppm)
 - 3) Input clock is faster than expected ($+1000$ ppm)
 - * For each scenario, the system generates 20,000 random bytes and verifies correctness.

VII. PPA SUMMARY

- **Data Rate Support:** Verified operation at 480 Mbps (Hi-Speed USB) and 12 Mbps (Full-Speed), meeting full USB 2.0 PHY requirements.
- **Lock Time:** The architecture achieves a lock time of zero cycles, because of its feed-forward design. Bit recovery begins immediately but there is a warm-up latency of 12 cycles for the Add-Drop FIFO.
- **Bit Error Rate (BER):** Verified $BER < 10^{-12}$ in simulation across 20,000+ bytes with ± 1000 ppm frequency offset. No bit errors were detected.
- **Jitter Tolerance:** Full system operation was confirmed with drift up to ± 1000 ppm and oversampled windows containing multiple transitions or CIDs.

Synthesis using the SKY130 process produced the following post-elaboration area estimates:

Module	Cell Count	Total Area (μm^2)
Oversampler	11	328
Coarse Data Recovery (CRD)	77	2,730
Add-Drop FIFO	188	6,674
Total	277	9,969

TABLE I
AREA BREAKDOWN OF SYNTHESIZED MODULES

Full post-synthesis and gate-level power analysis was not completed due to time constraints and issues with the HAMMER process and Volutus. However, based on related literature [1], we would expect an estimated power consumption of less than 10mW at 480 Mbps.

VIII. CONCLUSION

This report presents a robust digital CDR solution for USB 2.0 using blind oversampling. The addition of an AD-FIFO improves reliability by dynamically handling timing errors. Implemented in SKY130, the design is optimized for low power and efficient area utilization. Future work will focus on full-chip integration and further power optimizations.

REFERENCES

- [1] S.-H. Park, K.-H. Choi, J.-B. Shin, J.-Y. Sim, and H.-J. Park, “A single-data-bit blind oversampling data-recovery circuit with an add-drop FIFO for USB2.0 high-speed interface,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 55, no. 2, pp. 156–160, Feb. 2008, doi: 10.1109/TCSIL.2007.911791.
- [2] J.-J. Nam, Y.-J. Kim, K.-H. Choi, and H.-J. Park, “A UTMI-compatible physical-layer USB2.0 transceiver chip,” in *Proc. IEEE Int. SOC Conf.*, Portland, OR, USA, 2003, pp. 309–312, doi: 10.1109/SOC.2003.1241532.
- [3] Microchip Technology Inc., “USB3300 Hi-Speed USB Transceiver with ULPI Interface,” Datasheet, Doc. No. 00002142A, [Online]. Available: <https://ww1.microchip.com/downloads/en/DeviceDoc/00002142A.pdf>
- [4] S. Palermo, “Clock and Data Recovery (CDR) Techniques,” Lecture Notes, Dept. of Electrical and Computer Engineering, Texas A&M University, [Online]. Available: https://people.engr.tamu.edu/spalermo/eece689/lecture12_ee720_cdrs.pdf
- [5] D. K. Jeong, “Topics in IC Design – Clock and Data Recovery (CDR),” Lecture Notes, Department of Electrical and Computer Engineering, Seoul National University, 2020. [Online]. Available: <https://ocw.snu.ac.kr/sites/default/files/NODE/Lec%206%20-%20Clock%20and%20Data%20Recovery.pdf>