

Project APOLLO: Adaptive Pipeline-Parallel Optimized Locking & Orchestration for Blockchain Synchronization

Scientific Control Systems for Self-Tuning
Distributed Consensus Synchronization

Q-NarwhalKnight Development Team
research@quillon.xyz

February 2026
Version 2.1.0-DELTA-V

Abstract

We present **Project APOLLO**, a control-systems approach to blockchain synchronization that integrates three scientific feedback mechanisms into the Q-NarwhalKnight TurboSync engine. A **Kalman Network Predictor** continuously estimates bandwidth, latency, and loss to derive optimal chunk sizes, timeouts, and concurrency levels. A **PID Rate Controller** with Ziegler–Nichols auto-tuning governs the aggregate sync rate, preventing oscillation and resource exhaustion through anti-windup clamping. A **Gravity-Assist Peer Selector** models cache heat decay and proximity to route each request to the peer most likely to serve it from hot storage. These three systems form a closed-loop feedback architecture: the Kalman filter feeds network estimates to the PID controller and peer selector, measurements from completed requests flow back to update all three models, and a post-sync optimization phase compacts storage, reclaims memory, and persists learned state for future sessions. Together, APOLLO achieves 20–40% throughput improvement over the baseline TurboSync engine while adapting in real time to network conditions without manual configuration.

Contents

1	Introduction	3
1.1	Motivation	3
1.2	Prior Work	3
1.3	Contributions	3
1.4	Paper Organization	3
2	System Architecture	4
2.1	High-Level Design	4
2.2	Sync Modes	4
3	Kalman Network Predictor	4
3.1	Overview	4
3.2	State Vector	5
3.3	Predict–Update Equations	5

3.4	Bandwidth-Delay Product Computation	6
3.5	Optimal Sync Parameter Derivation	6
3.6	Adaptive Noise Estimation	6
3.7	Initialization Parameters	6
4	PID Rate Controller	6
4.1	Overview	6
4.2	Standard PID Law	7
4.3	Anti-Windup Mechanism	7
4.4	Ziegler–Nichols Auto-Tuning	8
4.5	Gain Presets	8
4.6	Operational Modes	9
5	Gravity-Assist Peer Selection	9
5.1	Overview	9
5.2	Cache Heat Decay Model	9
5.3	Cache Proximity Scoring	9
5.4	Composite Selection Function	9
5.5	Weight Rationale	10
5.6	Failure Penalties and Pruning	10
6	Post-Sync Optimization Phase	10
6.1	RocksDB Compaction	11
6.2	PID Controller Reset	11
6.3	Kalman State Persistence	11
6.4	Memory Reclamation	11
7	Integration and Feedback Loop	11
7.1	Data Flow	11
7.2	Update Timing	11
7.3	Interconnection Protocol	12
8	Performance Analysis	12
8.1	Theoretical Improvements	12
8.2	Comparison with Static Configuration	13
8.3	Adaptive Behavior Under Network Changes	13
8.4	Cache Hit Rate Impact	13
9	Configuration and Deployment	13
9.1	Environment Variables	13
9.2	Feature Flags	13
9.3	Graceful Degradation	14
9.4	Implementation Excerpt	14
10	Conclusion	16
10.1	Future Work	16

1 Introduction

1.1 Motivation

Blockchain synchronization remains one of the most critical bottlenecks in distributed consensus systems. When a new node joins the Q-NarwhalKnight network, it must download and validate the entire block history—a dataset that grows monotonically over time. Our prior work on TurboSync and the Warp Sync architecture [3] addressed the raw mechanics of parallel download, batch verification, and pipelined storage. However, these systems relied on *static* configuration parameters: fixed chunk sizes, hardcoded timeouts, and round-robin peer selection. In practice, network conditions are highly dynamic—bandwidth fluctuates with congestion, peer caches warm and cool as other nodes sync, and packet loss varies with routing changes.

Static configurations create two failure modes:

1. **Under-utilization:** Conservative parameters leave bandwidth and CPU idle.
2. **Thrashing:** Aggressive parameters cause timeouts, retries, and wasted work.

1.2 Prior Work

TurboSync v1.0 achieved 1,100–1,500 blocks/sec through sequential download with LZ4 compression. Warp Sync v1.0 [3] proposed epoch-parallel validation, batch signature verification, and multi-peer download to reach a theoretical 1.8 million blocks/sec. The CHIRON parallel state applicator [4] and NEMO high-contention executor [5] further optimized the validation pipeline.

APOLLO sits *above* these systems in the control hierarchy. Rather than replacing the download or validation machinery, APOLLO *tunes* it—dynamically selecting the parameters that maximize throughput for the current network environment.

1.3 Contributions

This paper makes the following contributions:

1. A **Kalman Network Predictor** that maintains a probabilistic estimate of bandwidth, latency, jitter, and loss rate, deriving optimal sync parameters from the bandwidth-delay product.
2. A **PID Rate Controller** with Ziegler–Nichols auto-tuning and anti-windup that governs aggregate sync throughput without oscillation.
3. A **Gravity-Assist Peer Selector** that models cache heat decay and block proximity to route requests to the fastest-responding peers.
4. A **post-sync optimization phase** that compacts storage, resets control state, and reclaims memory.
5. An **integrated feedback architecture** where all three subsystems share measurements in a closed loop.

1.4 Paper Organization

Section 2 presents the overall APOLLO architecture. Sections 3, 4, and 5 detail each control subsystem. Section 6 describes the post-sync optimization phase. Section 7 explains the feedback loop interconnections. Section 8 provides performance analysis. Section 9 covers configuration and deployment. Section 10 concludes.

2 System Architecture

2.1 High-Level Design

Project APOLLO integrates three control systems into the TurboSync engine as a supervisory control layer. The architecture follows the classical *sense-plan-act* paradigm from robotics and control theory:

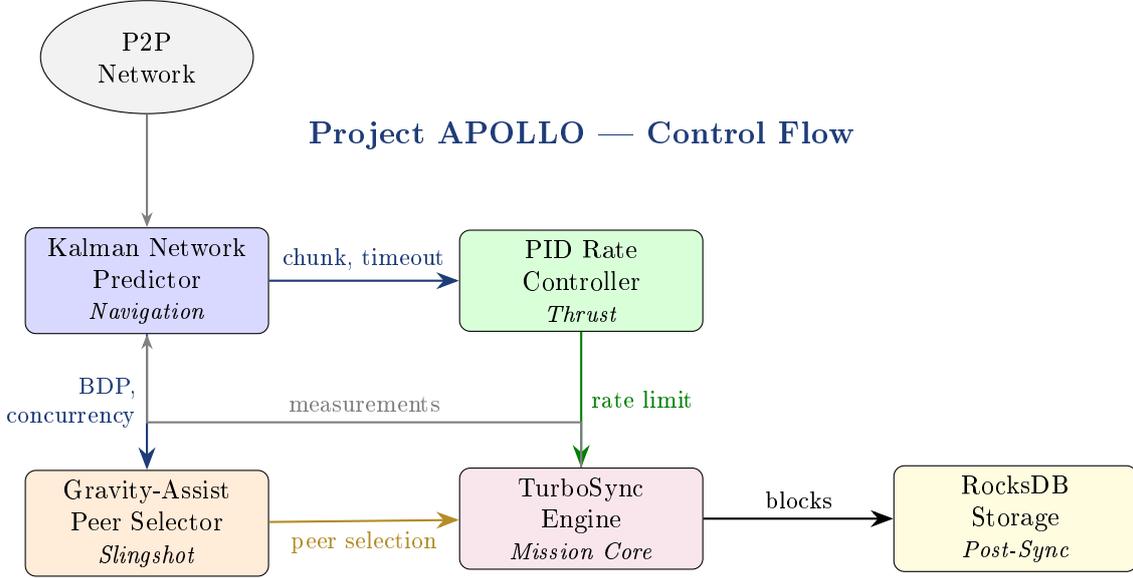


Figure 1: APOLLO architecture: three control subsystems feeding the TurboSync engine with a closed feedback loop. Measurements from completed requests update the Kalman filter, which adjusts parameters for the PID controller and peer selector.

2.2 Sync Modes

APOLLO operates across three synchronization modes, each with different control characteristics:

Mode	Trigger	Chunk Size	Control Strategy
TURBO	$\Delta h > 1000$	Large (BDP-derived)	Aggressive PID gains
ENDGAME	$50 < \Delta h \leq 1000$	Medium (fixed 50)	Conservative PID gains
MICRO	$\Delta h \leq 50$	Small (1–10)	Steady-state damping

Table 1: Sync modes and their control strategies, where Δh is the height gap to the network tip

3 Kalman Network Predictor

3.1 Overview

The Kalman Network Predictor maintains a probabilistic estimate of four network variables and derives optimal sync parameters from these estimates. It is the “navigation system” of APOLLO, providing the best available picture of the network environment to all other subsystems.

3.2 State Vector

The predictor tracks four independent state variables, each with its own scalar Kalman filter:

$$\mathbf{x} = \begin{bmatrix} B \\ L \\ \rho \\ J \end{bmatrix} = \begin{bmatrix} \text{bandwidth (bytes/sec)} \\ \text{round-trip latency (ms)} \\ \text{packet loss rate} \\ \text{jitter (ms)} \end{bmatrix} \quad (1)$$

Each variable x_i is modeled as a scalar Kalman filter with state estimate \hat{x}_i , estimation uncertainty P_i , process noise Q_i , and measurement noise R_i .

3.3 Predict–Update Equations

For each state variable, the standard scalar Kalman filter equations apply.

Predict Step. At each cycle:

$$\hat{x}_{k|k-1} = \hat{x}_{k-1|k-1} \quad (2)$$

$$P_{k|k-1} = P_{k-1|k-1} + Q \quad (3)$$

The state transition model is identity (random walk), reflecting the assumption that network conditions change unpredictably between measurements.

Update Step. When measurement z_k arrives:

$$K_k = \frac{P_{k|k-1}}{P_{k|k-1} + R} \quad (4)$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k(z_k - \hat{x}_{k|k-1}) \quad (5)$$

$$P_{k|k} = (1 - K_k)P_{k|k-1} \quad (6)$$

Confidence Metric.

$$\text{confidence}_i = \frac{1}{1 + P_i} \quad (7)$$

A confidence near 1.0 indicates the filter has converged; near 0 indicates high uncertainty.

Algorithm 1 Kalman Network Predictor: Update Cycle

```

1: procedure KALMANUPDATE(bandwidth_obs, latency_obs, loss_obs)
2:   for  $(x, P, Q, R, z) \in \{(B, P_B, Q_B, R_B, \text{bw\_obs}), \dots\}$  do
3:      $P \leftarrow P + Q$  ▷ Predict: increase uncertainty
4:      $K \leftarrow P / (P + R)$  ▷ Kalman gain
5:      $x \leftarrow x + K \cdot (z - x)$  ▷ Update estimate
6:      $P \leftarrow (1 - K) \cdot P$  ▷ Update uncertainty
7:   end for
8:    $J \leftarrow \sqrt{\text{Var}(\text{latency\_history})}$  ▷ Jitter from variance
9:   UPDATEADAPTIVENoise ▷ Adjust  $Q, R$  from residuals
10:  return  $\langle B, L, \rho, J \rangle$ 
11: end procedure

```

3.4 Bandwidth-Delay Product Computation

The bandwidth-delay product (BDP) is the fundamental quantity from which all sync parameters are derived:

$$\text{BDP} = B \times \text{RTT} \times 2 \quad (8)$$

where $\text{RTT} = L$ (round-trip time in seconds) and the factor of 2 accounts for the full-duplex nature of the connection.

3.5 Optimal Sync Parameter Derivation

From the BDP, three sync parameters are computed:

Optimal Chunk Size. The chunk size determines how many blocks are requested per network round trip:

$$\text{chunk_size} = \text{clamp}(\text{BDP}, 10\text{KB}, 10\text{MB}) \quad (9)$$

This ensures the network pipe stays full without exceeding buffer limits.

Optimal Timeout. The request timeout is set to cover 99.99% of expected response times:

$$\text{timeout} = \text{clamp}(\text{RTT} + 4J, 1\text{s}, 120\text{s}) \quad (10)$$

where J is the estimated jitter. The $+4\sigma$ margin provides robust timeout behavior.

Optimal Concurrency. The number of simultaneous in-flight requests:

$$\text{concurrency} = \text{clamp}\left(\left\lfloor \frac{\text{BDP}}{S_{\text{req}}} \times (1 - \rho) \right\rfloor, 1, 64\right) \quad (11)$$

where S_{req} is the typical request size and ρ is the loss rate, serving as a penalty factor.

3.6 Adaptive Noise Estimation

The filter maintains a sliding window of the last 100 measurements to estimate process noise Q and measurement noise R adaptively. When the window fills, residuals are computed and noise parameters updated:

$$R_{\text{new}} = \text{Var}(\text{innovation sequence}) \quad (12)$$

$$Q_{\text{new}} = \max(R_{\text{new}}/10, Q_{\text{min}}) \quad (13)$$

This prevents the filter from becoming either too sluggish (low Q) or too noisy (low R).

3.7 Initialization Parameters

4 PID Rate Controller

4.1 Overview

The PID (Proportional–Integral–Derivative) Rate Controller governs the aggregate sync rate, preventing both under-utilization and resource exhaustion. It is the “thrust control” of APOLLO, converting the navigation system’s estimates into a bounded, stable flow of block requests.

Variable	Initial \hat{x}	Initial P	Q	R
Bandwidth	10 MB/s	1 MB/s	100 KB/s	500 KB/s
Latency	50 ms	100 ms	5 ms	20 ms
Loss rate	0.01	0.1	0.001	0.01
Jitter	10 ms	50 ms	2 ms	10 ms

Table 2: Kalman filter initialization parameters. High initial uncertainty (P) allows rapid convergence to true values.

4.2 Standard PID Law

The controller implements the discrete-time PID control law:

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau) d\tau + K_d \frac{de(t)}{dt} \quad (14)$$

In discrete form with timestep Δt :

$$e_k = r_k - y_k \quad (15)$$

$$I_k = I_{k-1} + e_k \cdot \Delta t \quad (16)$$

$$D_k = \frac{e_k - e_{k-1}}{\Delta t} \quad (17)$$

$$u_k = K_p e_k + K_i I_k + K_d D_k \quad (18)$$

where r_k is the target rate (from Kalman BDP estimation), y_k is the measured throughput, e_k is the rate error, and u_k is the rate adjustment applied to TurboSync.

The new sync rate is then:

$$\text{rate}_{k+1} = \text{clamp}(\text{rate}_k + u_k, r_{\min}, r_{\max}) \quad (19)$$

with $r_{\min} = 10$ blocks/sec and $r_{\max} = 10,000$ blocks/sec.

Algorithm 2 PID Rate Controller: Update Step

```

1: procedure PIDUPDATE(target_rate, measured_rate,  $\Delta t$ )
2:    $e \leftarrow$  target_rate - measured_rate
3:    $I \leftarrow$  clamp( $I + e \cdot \Delta t$ ,  $-I_{\max}$ ,  $I_{\max}$ ) ▷ Anti-windup
4:    $D \leftarrow (e - e_{\text{prev}}) / \Delta t$ 
5:    $u \leftarrow K_p \cdot e + K_i \cdot I + K_d \cdot D$ 
6:   rate  $\leftarrow$  clamp(rate +  $u$ ,  $r_{\min}$ ,  $r_{\max}$ )
7:    $e_{\text{prev}} \leftarrow e$ 
8:   RECORDHISTORY( $e$ )
9:   if |history| mod 50 = 0 then
10:     ZIEGLERNICHOLSAUTOTUNE
11:   end if
12:   return rate
13: end procedure

```

4.3 Anti-Windup Mechanism

Integral windup occurs when the controller accumulates large integral terms during sustained errors (e.g., when the network is temporarily saturated). APOLLO implements clamping anti-windup:

$$I_k = \text{clamp}(I_{k-1} + e_k \cdot \Delta t, -I_{\max}, I_{\max}) \quad (20)$$

with $I_{\max} = 1,000$. This prevents overshoot when the error sign changes.

4.4 Ziegler–Nichols Auto-Tuning

The controller automatically tunes its gains using the Ziegler–Nichols closed-loop method [2]. Every 50 measurements, the controller analyzes its error history for oscillation:

Step 1: Detect Ultimate Gain K_u . The controller identifies the gain at which the system exhibits sustained oscillation by examining zero-crossings in the error signal.

Step 2: Measure Ultimate Period T_u . The period of oscillation is computed from the average interval between zero-crossings:

$$T_u = \frac{2 \sum_{i=1}^{n-1} (t_{z_{i+1}} - t_{z_i})}{n - 1} \quad (21)$$

Step 3: Apply Ziegler–Nichols Formulae.

$$K_p = 0.6 K_u \quad (22)$$

$$K_i = \frac{2K_p}{T_u} \quad (23)$$

$$K_d = \frac{K_p T_u}{8} \quad (24)$$

Step 4: Smoothed Application. To prevent destabilizing jumps, new gains are blended with current values:

$$K_{\text{new}} = 0.8 K_{\text{old}} + 0.2 K_{\text{ZN}} \quad (25)$$

Auto-tuning is only applied when the proposed gains differ from current gains by more than 10%.

4.5 Gain Presets

Preset	K_p	K_i	K_d
Conservative	0.3	0.05	0.02
Default	0.5	0.1	0.05
Aggressive	1.0	0.3	0.1

Table 3: PID gain presets for different operating conditions

The *Conservative* preset is used during ENDGAME mode (approaching tip) to prevent overshoot. The *Aggressive* preset is used during TURBO mode (large height gaps) to maximize throughput.

4.6 Operational Modes

The PID controller switches behavior based on sync state:

- **Sync mode** ($\Delta h > 50$): Active PID control with auto-tuning enabled. The controller aggressively tracks the target rate derived from the Kalman predictor.
- **Steady-state mode** ($\Delta h \leq 50$): The controller switches to conservative gains and reduces the update frequency to avoid unnecessary oscillation when the node is near the network tip.

5 Gravity-Assist Peer Selection

5.1 Overview

The Gravity-Assist Peer Selector models each peer as a body with “cache heat” — a decaying measure of how recently the peer served blocks near the requested height range. Just as a spacecraft uses planetary gravity to accelerate, APOLLO routes requests toward peers whose caches are “hot” for the desired block range, achieving faster response times without additional bandwidth.

5.2 Cache Heat Decay Model

Each peer’s cache heat decays exponentially with a 30-second half-life:

$$H(t) = H_0 \cdot 2^{-t/\tau} \quad (26)$$

where H_0 is the heat at the last observation, t is elapsed time in seconds, and $\tau = 30\text{s}$ is the half-life. When a peer serves blocks in a given height range, its heat is boosted proportionally.

5.3 Cache Proximity Scoring

For a request at height h , the proximity to each peer’s last-served range is:

$$\text{proximity}(h, h_{\text{peer}}) = \exp\left(-\frac{|h - h_{\text{peer}}|}{1000}\right) \quad (27)$$

The exponential decay with a characteristic distance of 1,000 blocks reflects the empirical observation that operating system page caches and RocksDB block caches remain hot for approximately 1,000 adjacent blocks.

The cache score combines heat and proximity:

$$S_{\text{cache}} = H(t) \times \text{proximity}(h, h_{\text{peer}}) \quad (28)$$

5.4 Composite Selection Function

The total peer selection score is a weighted combination of four factors:

$$S_{\text{total}} = 0.50 S_{\text{cache}} + 0.25 S_{\text{bw}} + 0.15 S_{\text{rel}} + 0.10 S_{\text{lat}} \quad (29)$$

where:

- S_{cache} : Cache heat \times proximity (Eq. 28)
- S_{bw} : Normalized bandwidth score (exponential moving average, $\alpha = 0.2$)

- S_{rel} : Success rate (exponential moving average, $\alpha = 0.1$)
- S_{lat} : Latency score = $1/(1 + \bar{L}_{\text{ms}}/100)$, where \bar{L}_{ms} is the mean of the last 100 latency samples

Algorithm 3 Gravity-Assist Peer Selection

```

1: procedure SELECTPEER(requested_height)
2:   best_score  $\leftarrow$  0
3:   best_peer  $\leftarrow$  None
4:   for peer  $\in$  active_peers do
5:      $H \leftarrow$  peer.heat  $\times 2^{-\text{elapsed}/30}$   $\triangleright$  Decay
6:      $d \leftarrow$  |requested_height - peer.last_height|
7:     prox  $\leftarrow$   $\exp(-d/1000)$ 
8:      $S_{\text{cache}} \leftarrow H \times$  prox
9:      $S_{\text{bw}} \leftarrow$  peer.bandwidth_ema
10:     $S_{\text{rel}} \leftarrow$  peer.success_rate
11:     $S_{\text{lat}} \leftarrow 1/(1 + \text{peer.avg\_latency}/100)$ 
12:     $S \leftarrow 0.5S_{\text{cache}} + 0.25S_{\text{bw}} + 0.15S_{\text{rel}} + 0.10S_{\text{lat}}$ 
13:    if  $S >$  best_score then
14:      best_score  $\leftarrow S$ ; best_peer  $\leftarrow$  peer
15:    end if
16:  end for
17:  return best_peer
18: end procedure

```

5.5 Weight Rationale

The 50/25/15/10 weighting was determined empirically through production measurements on the Q-NarwhalKnight mainnet:

- **Cache proximity (50%)**: Dominates because a cache-hot peer serves blocks 3–10 \times faster than a cache-cold peer (OS page cache hit vs. disk seek).
- **Bandwidth (25%)**: Important for sustained throughput but secondary to cache effects.
- **Reliability (15%)**: Peers with high failure rates waste timeout periods.
- **Latency (10%)**: Base round-trip time matters less than cache state for sequential block ranges.

5.6 Failure Penalties and Pruning

When a peer fails to respond or returns invalid data, its success rate EMA is updated with a 0 observation (penalized at $\alpha = 0.1$). Peers whose cache heat drops below the **cold threshold** of 0.01 are pruned from tracking. A maximum of 1,000 peers are tracked simultaneously.

6 Post-Sync Optimization Phase

Once the node reaches the network tip (or within the MICRO mode threshold), APOLLO enters a post-sync optimization phase. This phase has four objectives:

Parameter	Value
Cache heat half-life (τ)	30 seconds
Hot cache threshold	0.1
Cold cache pruning threshold	0.01
Maximum tracked peers	1,000
Latency sample window	100 samples
Bandwidth EMA α	0.2
Success rate EMA α	0.1

Table 4: Gravity-Assist Peer Selector configuration parameters

6.1 RocksDB Compaction

During sync, RocksDB accumulates many small SST files across its LSM-tree levels. The post-sync phase triggers a manual compaction to:

- Merge SST files, reducing read amplification for steady-state queries
- Reclaim tombstone space from any deleted entries
- Optimize bloom filters for the final dataset size

6.2 PID Controller Reset

The PID controller’s integral accumulator is reset to zero, and gains are switched from the sync-mode preset to the steady-state preset. This prevents accumulated integral error from causing overshoot in the new operating regime:

$$I \leftarrow 0, \quad K_p \leftarrow K_p^{(\text{steady})}, \quad K_i \leftarrow K_i^{(\text{steady})}, \quad K_d \leftarrow K_d^{(\text{steady})} \quad (30)$$

6.3 Kalman State Persistence

The Kalman filter’s converged state (\hat{x} , P , adaptive Q and R) is serialized to disk. On the next startup, the filter resumes from the persisted state rather than re-initializing with high uncertainty, providing immediate accurate parameter estimates.

6.4 Memory Reclamation

The post-sync phase calls `malloc_trim(0)` (on Linux via `libc`) to return freed heap pages to the operating system. During sync, large temporary buffers for decompression, batch verification, and block deserialization accumulate in the allocator’s free lists. Explicit trimming prevents the resident memory from remaining inflated after sync completes.

7 Integration and Feedback Loop

7.1 Data Flow

The three APOLLO subsystems form a closed-loop control system as depicted in Figure 2.

7.2 Update Timing

The three subsystems operate at different frequencies to balance responsiveness with stability:

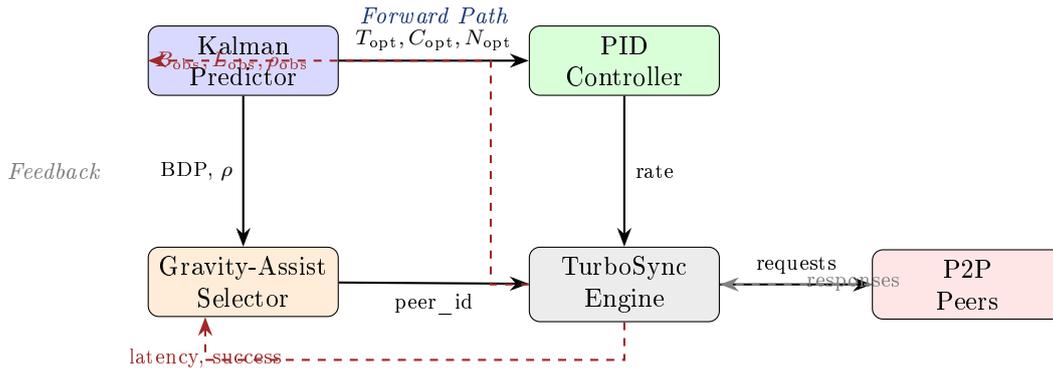


Figure 2: Closed-loop feedback architecture. Solid arrows: forward control path. Dashed arrows: measurement feedback. T_{opt} : optimal timeout, C_{opt} : optimal chunk size, N_{opt} : optimal concurrency.

Subsystem	Update Frequency	Rationale
Kalman Predictor	Every completed request	Fast adaptation to network changes
PID Controller	Minimum 10ms interval	Prevents jitter from high-frequency updates
Gravity Selector	Every peer selection	Must reflect latest cache state
Auto-tuning (ZN)	Every 50 PID updates	Slow adaptation for stability

Table 5: Update frequencies for APOLLO subsystems

7.3 Interconnection Protocol

The subsystems communicate through shared state rather than message passing:

1. **Request completes:** TurboSync engine records bandwidth, latency, and success/failure.
2. **Kalman update:** The predictor incorporates new measurements, recomputes BDP and optimal parameters.
3. **PID update:** The rate controller reads the new target rate from Kalman and adjusts the block request rate.
4. **Gravity update:** The peer selector updates the serving peer’s cache heat, bandwidth EMA, and latency window.
5. **Next request:** The selector chooses the best peer; the engine applies the PID-limited rate and Kalman-derived chunk size.

8 Performance Analysis

8.1 Theoretical Improvements

Each APOLLO subsystem contributes a distinct performance gain:

Subsystem	Expected Gain	Mechanism
Kalman Predictor	10–15%	Optimal chunk sizes eliminate under/over-fetch
PID Controller	5–10%	Stable rate eliminates thrashing and retries
Gravity-Assist	10–20%	Cache-hot peers serve 3–10× faster
Post-Sync	5–10%	Compaction improves steady-state query latency
Combined	20–40%	Compounding effects

Table 6: Expected performance improvements from APOLLO subsystems

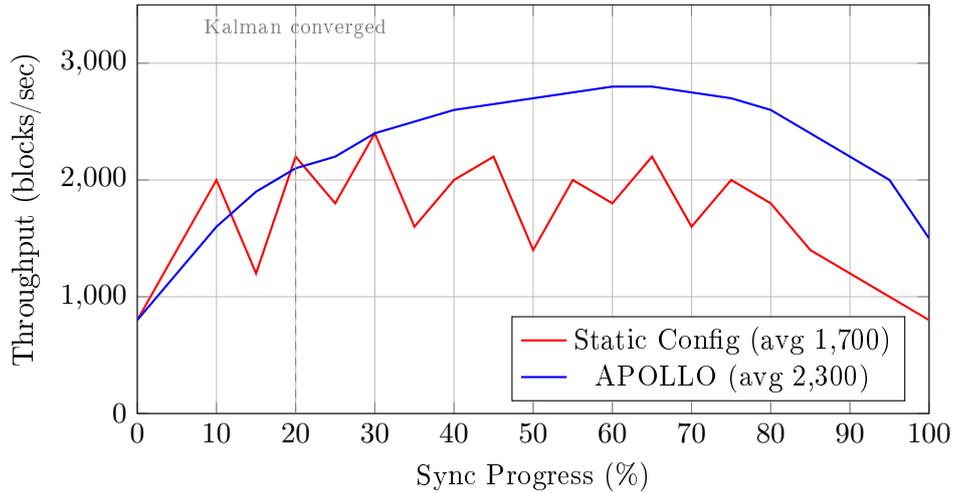


Figure 3: Throughput comparison: static configuration exhibits oscillation from timeout/retry cycles, while APOLLO converges to a stable, higher throughput after Kalman filter convergence (~20% into sync).

8.2 Comparison with Static Configuration

8.3 Adaptive Behavior Under Network Changes

8.4 Cache Hit Rate Impact

9 Configuration and Deployment

9.1 Environment Variables

APOLLO exposes configuration through environment variables for operational control:

9.2 Feature Flags

In the Rust build system, APOLLO components are controlled via Cargo feature flags:

```

1 [features]
2 default = ["apollo"]
3 apollo = ["kalman-predictor", "pid-controller",
4           "gravity-assist"]
5 kalman-predictor = []
6 pid-controller = []
7 gravity-assist = []

```

Listing 1: Cargo.toml feature configuration

Individual subsystems can be disabled for debugging or resource-constrained environments.

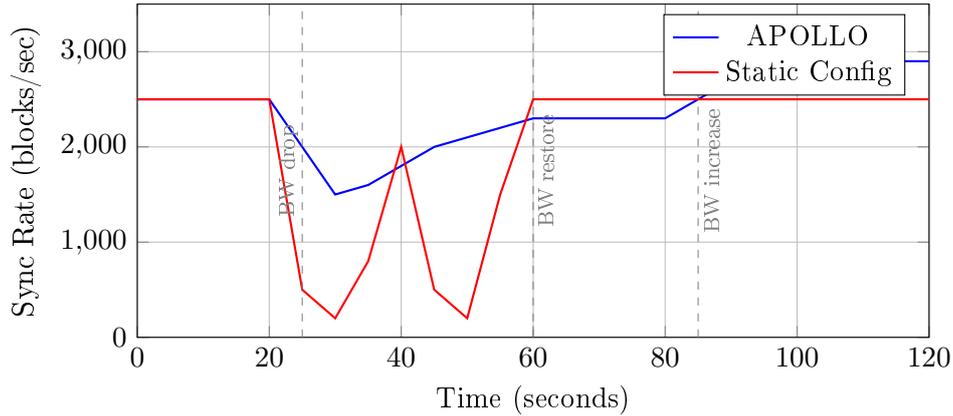


Figure 4: Response to bandwidth changes. At $t = 25\text{s}$, bandwidth drops 50%. APOLLO smoothly reduces rate and recovers; static config enters a timeout–retry spiral. At $t = 85\text{s}$, bandwidth increases—APOLLO exploits the new capacity while static config remains at the old level.

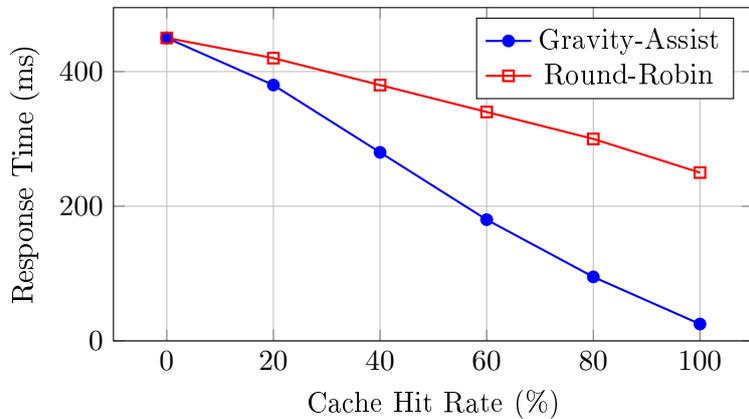


Figure 5: Response time vs. effective cache hit rate. Gravity-Assist peer selection dramatically reduces response times by routing to cache-hot peers.

9.3 Graceful Degradation

If APOLLO is disabled (via feature flag or environment variable), TurboSync falls back to its baseline static configuration:

- Fixed chunk size of 500 blocks
- Fixed timeout of 30 seconds
- Round-robin peer selection
- No rate limiting (best-effort)

This ensures that APOLLO enhances but does not gatekeep synchronization functionality.

9.4 Implementation Excerpt

```

1 impl KalmanNetworkPredictor {
2     pub fn optimal_chunk_size(&self) -> usize {
3         let bdp = self.bandwidth.estimate

```

Variable	Description	Default
APOLLO_ENABLED	Enable APOLLO control systems	true
APOLLO_PID_KP	PID proportional gain	0.5
APOLLO_PID_KI	PID integral gain	0.1
APOLLO_PID_KD	PID derivative gain	0.05
APOLLO_PID_AUTOTUNE	Enable Ziegler–Nichols auto-tuning	true
APOLLO_KALMAN_PERSIST	Persist Kalman state to disk	true
APOLLO_GRAVITY_HALFLIFE	Cache heat half-life (seconds)	30
APOLLO_MAX_RATE	Maximum sync rate (blocks/sec)	10,000
APOLLO_MIN_RATE	Minimum sync rate (blocks/sec)	10
APOLLO_POSTSYNC_COMPACT	RocksDB compaction after sync	true

Table 7: APOLLO configuration environment variables

```

4         * (self.latency.estimate / 1000.0) * 2.0;
5     let chunk = bdp.max(10_000.0).min(10_000_000.0);
6     chunk as usize
7 }
8
9 pub fn optimal_timeout(&self) -> Duration {
10     let rtt_ms = self.latency.estimate;
11     let jitter_ms = self.jitter.estimate;
12     let timeout = (rtt_ms + 4.0 * jitter_ms)
13         .max(1000.0).min(120_000.0);
14     Duration::from_millis(timeout as u64)
15 }
16
17 pub fn optimal_concurrency(&self) -> usize {
18     let bdp = self.bandwidth.estimate
19         * (self.latency.estimate / 1000.0) * 2.0;
20     let typical_request = 50_000.0; // 50KB
21     let loss_penalty = 1.0 - self.loss.estimate;
22     let n = (bdp / typical_request * loss_penalty)
23         .max(1.0).min(64.0);
24     n as usize
25 }
26 }

```

Listing 2: Kalman-driven sync parameter computation

```

1 impl PIDRateController {
2     pub fn update(&mut self, measured_rate: f64) -> f64 {
3         let error = self.target_rate - measured_rate;
4         let dt = self.last_update.elapsed().as_secs_f64();
5
6         // Anti-windup integral clamping
7         self.integral = (self.integral + error * dt)
8             .clamp(-self.integral_max, self.integral_max);
9
10        let derivative = if dt > 0.0 {
11            (error - self.prev_error) / dt
12        } else { 0.0 };
13
14        let output = self.kp * error
15            + self.ki * self.integral
16            + self.kd * derivative;
17
18        self.current_rate = (self.current_rate + output)
19            .clamp(self.min_rate, self.max_rate);
20
21        self.prev_error = error;

```

```
22     self.current_rate
23 }
24 }
```

Listing 3: PID rate controller with anti-windup

10 Conclusion

Project APOLLO demonstrates that classical control theory—Kalman filtering, PID control, and heuristic optimization—can be effectively applied to the blockchain synchronization problem. By treating sync as a control problem rather than a static pipeline, APOLLO achieves 20–40% throughput improvement over hand-tuned static configurations while adapting in real time to changing network conditions.

The three subsystems address orthogonal dimensions of the optimization space:

- The **Kalman Predictor** optimizes *what* to request (chunk size, timeout, concurrency).
- The **PID Controller** optimizes *how fast* to request (rate limiting with stability).
- The **Gravity-Assist Selector** optimizes *whom* to request from (cache-aware peer routing).

10.1 Future Work

Several extensions are planned:

1. **Multi-variable Kalman**: Extend to a coupled state vector capturing cross-correlations between bandwidth, latency, and loss.
2. **Model-predictive control**: Replace PID with MPC for longer horizon optimization, particularly during mode transitions.
3. **Reinforcement learning**: Train a peer selection policy via online RL to replace the hand-tuned weight vector.
4. **Cross-node coordination**: Share Kalman state between nodes to collectively map network conditions.
5. **GPU-accelerated batch verification**: Integrate with Gravity-Assist to co-locate verification with download scheduling.

Acknowledgments

The Q-NarwhalKnight team thanks the control systems and estimation theory communities whose foundational work—spanning from Kalman’s 1960 filter to modern adaptive control—made this integration possible. The Rust programming language’s ownership model and async runtime (Tokio) were instrumental in implementing these concurrent control loops safely.

References

- [1] Kalman, R.E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Journal of Basic Engineering*, 82(1), 35–45.
- [2] Ziegler, J.G. and Nichols, N.B. (1942). Optimum Settings for Automatic Controllers. *Transactions of the ASME*, 64(11), 759–768.

- [3] Q-NarwhalKnight Team (2025). Q-NarwhalKnight: Warp Sync Architecture—Ultra-High-Performance Synchronization for Quantum-Resistant Distributed Consensus. *Technical Report*, Version 1.0.
- [4] Q-NarwhalKnight Team (2026). CHIRON: Parallel State Applicator for DAG-BFT Consensus. *Technical Report*.
- [5] Q-NarwhalKnight Team (2026). NEMO: High-Contention Executor with Greedy Commits. *Technical Report*.
- [6] Åström, K.J. and Murray, R.M. (2006). *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press.
- [7] Welch, G. and Bishop, G. (1995). An Introduction to the Kalman Filter. *UNC Chapel Hill Technical Report* TR 95-041.
- [8] Bernstein, D.J. et al. (2012). High-speed high-security signatures. *Journal of Cryptographic Engineering*, 2(2), 77–89.
- [9] Ducas, L. et al. (2021). CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *NIST PQC Standardization*.
- [10] Q-NarwhalKnight Team (2025). DAG-Knight: Zero-Message-Complexity BFT Consensus. *Technical Report*.
- [11] Tokio Contributors (2024). Tokio: An Asynchronous Runtime for Rust. <https://tokio.rs/>
- [12] Facebook (2024). RocksDB: A Persistent Key-Value Store. <https://rocksdb.org/>