

Defeating Time-Average Selfish Mining Across Epoch Boundaries in Nakamoto Consensus

Ren Zhang

CKB Association
friendren@gmail.com

Abstract. Selfish mining profits in Nakamoto consensus because difficulty adjustment mechanisms (DAMs) misinterpret uncounted orphans as hash-rate contraction. Although orphan-aware DAMs record orphaned proof-of-work, strict same-epoch inclusion rules enable an *orphan exclusion attack* (OEA): an adversary can race private forks near epoch boundaries to permanently exclude honest tail orphans from retargeting, restoring substantial profits in short-epoch protocols. We propose the *pipelined buffer difficulty adjustment mechanism* (PB-DAM), which partitions each epoch into an accountable prefix and a settlement buffer of depth d . Pipelining estimation windows across epochs grants honest miners a grace period to report prefix uncles while advancing buffer work to subsequent retargets, closing the boundary gap without heuristic damping. Random-walk excursion bounds demonstrate that OEA success decays exponentially with d for all $\alpha \leq \alpha^* < 1/2$, reducing the time-averaged profit advantage to $O(\varepsilon)$. Under stylized dynamic scheduling models, intermittent idling is analytically shown to be unable to raise the gross reward rate above α . Simulations confirm that buffer depths $d = 16$ neutralize selfish mining even at $L = 64$ when $\alpha = 0.40$, and systems recover from a 50% hash-rate crash within 2.2 ± 0.8 epochs, incurring under 1% block capacity overhead.

1 Introduction

The selfish mining attack against Nakamoto consensus allows a miner to collect more than its share of main-chain rewards by invalidating honestly mined blocks [4, 14]. However, relative-revenue advantage does not immediately translate into real profit [8, 9]. The actual profitability hinges on how the underlying protocol regulates block production over time. In proof-of-work blockchains, the network enforces a target block generation rate against fluctuating hash rate by partitioning execution into *epochs*—fixed spans of L main-chain blocks—and governing the required proof-of-work per block via a *difficulty*. At each epoch boundary, a *difficulty adjustment mechanism* (DAM), also termed DAA, measures elapsed wall-clock time to dynamically recalculate the subsequent difficulty (the *retarget*). When selfish mining occurs, it creates *orphans*—validly mined blocks that fail to be included in the canonical main chain. Because orphan production consumes proof-of-work without contributing to canonical chain

progress, the attacker’s time-averaged revenue before a retarget remains strictly bounded by that under honest mining. Real profit appears only after the DAM reacts: because Bitcoin’s DAM accounts exclusively for main-chain blocks and ignores orphans, the epoch runs abnormally long. The DAM misdiagnoses this delay as a drop in total network hash rate and reduces the subsequent difficulty. The attacker then mines on this discounted difficulty, finding blocks at a rate far exceeding its nominal hash-rate share.

Observing that the crux of selfish mining lies in the DAM’s misestimation, two classes of defenses aim to design secure DAMs to eliminate the profitability of selfish mining. The first class dampens, caps, or asymmetrically slows downward difficulty adjustments across consecutive epochs to inflate the attacker’s upfront profit lag [1, 11]. However, because damping mechanisms infer network hash rate solely from main-chain progress, they cannot distinguish adversarial orphan storms from genuine network contractions, causing sluggish adaptation that degrades chain liveness. The second class incorporates orphaned proof-of-work back into the difficulty estimator. Early proposals described intuitive ideas of tracking abandoned blocks [9, 16], yet they lacked rigorous formal proofs under strategic adversarial withholding. Recently, Sarenche et al. [15] provided the first rigorous formulation and theoretical security analysis of orphan-aware DAMs. However, their analysis also reveals that such naive orphan-aware DAMs cannot completely eliminate adversarial manipulation of the difficulty retarget. Specifically, they identified an inherent attack vector native to such designs: the *orphan exclusion attack* (OEA), in which an adversary strategically prevents honest orphans from being recorded before the retarget.

The vulnerability to OEA stems from an information gap at epoch boundaries. Because existing orphan-aware designs enforce a strict same-epoch inclusion deadline, an adversary racing a private fork near an epoch boundary can seal the retarget cutoff ahead of the honest network, stranding tail orphans from difficulty accounting. While such boundary exclusions are diluted in long epochs (e.g., $L = 2016$ in Bitcoin), modern protocols often adopt short epochs ($L \in [64, 128]$) to track network dynamics, where an adversary can exclude up to 27% of an epoch’s blocks and restore substantial selfish mining profits.

To eliminate this boundary vulnerability, we propose the *pipelined buffer difficulty adjustment mechanism* (PB-DAM). PB-DAM partitions each epoch of L blocks into an *accountable prefix* of length $L - d$ and a *settlement buffer* of length d . By shifting the retarget estimation window to the prefix while using the buffer as a reporting grace period, any honest block mined in the suffix incorporates all outstanding prefix uncles. Crucially, buffered blocks and their uncles roll seamlessly into the subsequent retarget window, ensuring no proof-of-work is permanently excluded.

The security of PB-DAM derives directly from the chain-quality property of Nakamoto consensus [5–7, 13]. For any adversarial hash-rate fraction $\alpha < 1/2$, the probability that an attacker suppresses honest blocks throughout the entire buffer decays exponentially with depth d . As long as at least one honest main-chain block is minted within the d -block buffer, all unconfirmed honest

uncles from the prefix are immediately incorporated. Consequently, a modest buffer depth renders OEA success negligible, guaranteeing that time-averaged selfish mining yields zero profit advantage. Discrete-event simulations confirm our theoretical results: in short-epoch regimes ($L = 64$) where prior defenses yield up to +8.47% adversarial gains, PB-DAM with $d = 16$ neutralizes excess returns to statistical parity for all $\alpha \leq 0.40$, while adapting to sudden hash-rate collapses within a single window without heuristic damping. Requiring no reward restructuring, PB-DAM serves as a practical, drop-in defense for proof-of-work blockchains.

2 Retarget Attacks and Limitations of Existing Defenses

2.1 Selfish Mining as an Attack on the Retarget

Eyal and Sirer showed that Bitcoin’s longest-chain rule is not incentive compatible, allowing a strategic miner to obtain relative rewards exceeding its hash-rate share [4], while Sapirshstein et al. [14] derived optimal withholding policies via a Markov decision process (MDP). However, both analyses assumed a static difficulty target. Analyzing absolute returns over time, Grunspan and Pérez-Marco [9] demonstrated that selfish mining yields no revenue advantage prior to retargeting. Building on this, Negy et al. [12] introduced *intermittent selfish mining* (ISM), where an attacker alternates between withholding to force difficulty drops and mining on discounted targets. In both regimes, the upfront negative cash flow constitutes a transient *profit lag* [2, 10], repaid only once difficulty discounts take effect.

2.2 Slowing the Retarget, and the Liveness Cost

One defensive paradigm exploits this profit lag by dampening downward difficulty adjustments without tracking orphans, forcing the adversary through extended negative cash flow to raise the capital barrier [11]. In the Zeno DAA family [1], epoch difficulty D_k moves only a fraction $1 - \tau$ toward the estimated target $\tilde{D}_k$: $D_k = (1 - \tau)\tilde{D}_k + \tau D_{k-1}$. Tuning $\tau \in [0, 1)$ extends the profit lag (e.g., from 4 weeks in Bitcoin to over 11 weeks for $\alpha = 0.40$), inflating upfront costs. However, merely extending the profit lag fails to eliminate stationary profitability for well-capitalized adversaries [15]. More critically, dampening causes severe liveness failures: inferring hash rate solely from main-chain inter-arrival times, the DAM cannot distinguish an adversarial orphan storm from a genuine hash-rate collapse (e.g., power outages or regulatory bans), taking multiple epochs to recalibrate and trapping the network in extended delays [1].

2.3 Orphan-Aware Adjustment and Boundary Vulnerabilities

An alternative paradigm incorporates orphaned proof-of-work into the retarget estimator [9]. Wang et al. [16] showed that recording an orphan fraction ρ

raises the selfish mining threshold (e.g., to $\approx 37\%$ at $\rho = 0.7$), but assumed guaranteed inclusion without defense against strategic withholding. Sarenche et al. [15] formalized orphan-aware adjustment by defining valid uncles via parent visibility, acyclicity, and a strict same-epoch constraint. Under their $\text{DAM}_1^{\text{mod}}$ (fixed reported-block clock) and $\text{DAM}_2^{\text{mod}}$ (fixed main-chain clock), they proved that neither continuous selfish mining nor ISM yields any time-averaged profit advantage if all active work is reported.

Nevertheless, naive orphan-aware DAMs remain vulnerable to two attacks:

1. *Smart Intermittent Selfish Mining (SISM2) / Coin-Hopping*: The adversary mines selfishly in odd epochs and idles or migrates hash rate in even epochs; the resulting epoch delays force downward retargets which, combined with electricity savings during idle epochs, yield a net profit advantage.
2. *Orphan Exclusion Attack (OEA)*: Near an epoch boundary, the adversary releases a private fork to orphan honest tail blocks. Because the same-epoch constraint prohibits next-epoch blocks from referencing uncles with preceding-epoch parents, these tail orphans are permanently excluded from retarget accounting.

While an OEA perturbs retargets by only $\Theta(L_{\text{OEA}}/L)$ in long Bitcoin-style epochs ($L = 2016$), it causes severe target underestimation in modern short-epoch protocols ($L \in [64, 128]$).

Epoch-Based versus Sliding-Window Retargeting. Some blockchains use sliding-window or per-block DAAs to deter inter-chain coin-hopping. We focus on epoch-based retargeting because discrete epoch retargeting remains standard across major networks (e.g., Bitcoin, Litecoin, and Nervos CKB [17]). Because sliding windows with hard cutoffs are topologically equivalent to moving epoch boundaries, this boundary vulnerability extends to sliding-window retargeting as well.

3 System and Threat Model

3.1 System Model

Participants and Network. Time proceeds in discrete rounds. A set $\mathcal{H}$ of honest miners and a single adversary $\mathcal{A}$ control hash-rate fractions α_H and α with $\alpha_H + \alpha = 1$ and $\alpha < 1/2$. Block discovery is a Poisson process of intensity λ when the difficulty matches the network’s total hash rate; under difficulty coefficient D , the intensity scales as λ/D in the usual way. We work in the synchronous model of Sarenche et al. [15]: a block published in round r is delivered to every party by the end of that round. When two conflicting tips compete, a fraction $\gamma \in [0, 1]$ of honest hash rate adopts the adversary’s tip. Honest miners always extend a locally longest chain, publish immediately, and embed all known eligible uncle headers into newly mined blocks [15]. Uncle embedding can be mandated via consensus rules or inclusion tips. The adversary may withhold, match, adopt, or idle arbitrarily.

Chains, Rewards, and Normalized Advantage. A *main-chain* block is a block that belongs to the unique longest chain adopted by all honest miners after ties are resolved. An *orphan* is a valid proof-of-work block that is not on the main chain. Each main-chain block awards a fixed subsidy K ; transaction fees are omitted. Relative revenue over an interval is the adversary’s share of main-chain rewards in that interval. Under honest mining, the adversary’s expected block-reward rate is proportional to its hash-rate share, $\alpha\lambda K$.

Unlike Sarenche et al. [15], who evaluate profitability under dynamic operational expenditures, our core system model deliberately excludes explicit mining cost parameters such as electricity tariffs and hardware depreciation. Because a proof-of-work consensus protocol cannot observe miners’ physical power meters, difficulty adjustment mechanisms can only govern on-chain currency issuance rather than external economic fluctuations. Accordingly, our primary security metric is the adversary’s normalized reward advantage over honest mining:

$$P_{\text{adv}}(\pi) = \frac{R_{\mathcal{A}}(\pi) - \alpha\lambda K}{\lambda K}, \quad (1)$$

where $R_{\mathcal{A}}(\pi)$ denotes the adversary’s time-averaged block-reward rate. Thus, $P_{\text{adv}} > 0$ if and only if the adversary extracts newly minted coins at a rate exceeding its fair share $\alpha\lambda K$.

Active and Effective Hash Rate. Following [15], the *normalized active hash rate* M^{total} of an interval is the fraction of total hash rate used to mine either main-chain or orphan blocks (so $M^{\text{total}} = 1$ under full-power mining with no idling). The *normalized effective hash rate* M^{mc} is the fraction that ends on the main chain. Selfish mining satisfies $M^{\text{mc}} < 1$ even at full power. Write $M_{\mathcal{A}}^{\text{mc}}$ for the adversary’s own main-chain fraction; then

$$\text{RelRev} \cdot M^{\text{mc}} = M_{\mathcal{A}}^{\text{mc}} \leq \alpha, \quad (2)$$

with equality only when the adversary never orphans its own blocks [15]. Identity (2) is the reason an orphan-aware DAM can cancel a relative-revenue advantage: once the target tracks active rather than effective hash rate, reward per unit time is governed by $M_{\mathcal{A}}^{\text{mc}}$ rather than by RelRev.

3.2 Threat Model

Adversary Model. The adversary may follow any (possibly randomized) mining and publishing policy, including Eyal–Sirer selfish mining, the MDP-optimal policy of Sapirshtein et al., intermittent and smart intermittent selfish mining, and orphan exclusion. It may maintain a private chain, choose $\gamma = 1$, and—when we state an upper bound—be given the predictive power used by Sarenche et al. to select the longest feasible exclusion at an epoch boundary [15]. It cannot break proof-of-work, forge signatures, or eclipse honest miners.




Fig. 1. Pipelined Buffer Difficulty Adjustment Mechanism

Network and Natural Orphans. In the security theorems of Section 5 and the empirical evaluations of Section 6, honest blocks that are published immediately are never orphaned by network delay alone. Section 7 discusses the effect of network propagation delay on d .

Declared Hash-Rate Bound. The protocol is parameterized by a constant $\alpha^* \in (0, 1/2)$. Security claims hold whenever $\alpha \leq \alpha^*$. No fixed buffer length can cover every $\alpha < 1/2$, because the expected length of a successful exclusion diverges as $\alpha \rightarrow 1/2^-$ [15].

3.3 Security Goals

Let W_k denote the proof-of-work attributed by PB-DAM to estimation window k , and let W_k^* denote the total active proof-of-work published across that window. We formulate three security goals parameterized by an error tolerance $\varepsilon \in (0, 1)$:

1. *OEA resistance.* An adversary cannot persistently exclude honest tail orphans from difficulty retarget calculations, failing except with probability at most ε .
2. *Unbiased estimation.* Difficulty updates account for all active proof-of-work faithfully ($W_k = W_k^*$) with probability at least $1 - \varepsilon$.
3. *Reward fairness and incentive compatibility.* Under stationary full-power mining, no withholding strategy can extract an asymptotic time-averaged reward rate exceeding the nominal share $\alpha\lambda K + O(\varepsilon)$, while dynamic scheduling is bounded under first-order progress models.

4 Pipelined Buffer Difficulty Adjustment Mechanism

The fundamental defect of naive orphan-aware difficulty adjustment is that the reporting deadline coincides with the retarget boundary, enabling an adversary to strand tail orphans by sealing the epoch on a private fork. PB-DAM resolves this gap by partitioning each epoch of L blocks into an accountable prefix of length $L - d$ and a settlement buffer of depth d , pipelining estimation windows

across consecutive epochs so that buffer work is carried forward rather than discarded. Below, Section 4.1 formalizes uncle eligibility across buffer boundaries, and Section 4.2 defines the pipelined retarget update.

4.1 Epoch Partition and Uncle Eligibility

Fix main-chain epoch length L and buffer length d with $1 \leq d < L/2$. Epoch $k \geq 1$ consists of L consecutive main-chain blocks $B_{k,1}, \dots, B_{k,L}$ mined at difficulty coefficient D_k . We partition each epoch into an *accountable prefix* $P_1^{(k)} = \{B_{k,1}, \dots, B_{k,L-d}\}$ and a *settlement buffer* $P_2^{(k)} = \{B_{k,L-d+1}, \dots, B_{k,L}\}$. Write $\partial_k := B_{k,L-d}$ for the boundary block sealing the prefix, and $t(\cdot)$ for block timestamps. The ideal duration of L main-chain blocks is $T_{\text{target}} := L/\lambda$.

Definition 1 (Settlement uncle). *A block U is a settlement uncle of a main-chain block N for estimation window k if and only if all of the following conditions hold:*

1. *height(N) > height(U), and the height of U matches either a block in buffer $P_2^{(k-1)}$ (with $N \in P_1^{(k)}$), or a block in prefix $P_1^{(k)}$ (with $N \in P_1^{(k)} \cup P_2^{(k)}$);*
2. *U 's difficulty matches the epoch corresponding to height(U) (specifically, D_k for epoch k or D_{k-1} for epoch $k-1$);*
3. *U is not on the main chain, and the parent of U either appears on N 's chain (as a main-chain ancestor or as a previously embedded uncle), or is itself a settlement uncle embedded by N ;*
4. *no earlier main-chain block on the same chain has embedded U .*

Condition 1 implements the buffer rule: prefix orphans remain eligible during the same epoch's buffer, while buffer orphans remain eligible through the subsequent prefix for pipelined absorption (Section 4.2). Orphans that miss their settlement deadline become permanently ineligible, preventing an adversary from banking stale orphans against future retargets [15]. Honest blocks embed all eligible uncle headers in dedicated commitment fields. Uncles receive no block rewards.

4.2 Pipelined Estimation Window

The estimation window for the retarget at the end of epoch k is the half-open main-chain interval

$$\mathcal{W}_k := (\partial_{k-1}, \partial_k] = \{B_{k-1,L-d+1}, \dots, B_{k-1,L}\} \cup \{B_{k,1}, \dots, B_{k,L-d}\}. \quad (3)$$

Thus $\mathcal{W}_k$ always contains exactly L main-chain blocks: the buffer of epoch $k-1$ together with the prefix of epoch k . Difficulty is recomputed at the end of the epoch.

Definition 2 (PB-DAM update). Let $\text{Uncles}(\mathcal{W}_k)$ be the set of settlement uncles embedded by main-chain blocks in $\mathcal{W}_k \cup P_2^{(k)}$ that are attributed to window k by Definition 1. Define accounted work and elapsed time by

$$W_k = (L - d) \cdot D_k + d \cdot D_{k-1} + \sum_{U \in \text{Uncles}(\mathcal{W}_k)} \text{Work}(U), \quad (4)$$

$$T_k = t(\partial_k) - t(\partial_{k-1}). \quad (5)$$

Let $W_k^{\text{mc}} := (L - d)D_k + dD_{k-1}$ denote the baseline main-chain proof-of-work in $\mathcal{W}_k$, and let $\bar{D}_k := W_k^{\text{mc}}/L$ be the average difficulty across the window. The next difficulty coefficient is

$$D_{k+1} = \bar{D}_k \cdot \frac{W_k}{W_k^{\text{mc}}} \cdot \frac{T_{\text{target}}}{T_k} = \frac{W_k}{L} \cdot \frac{T_{\text{target}}}{T_k}, \quad (6)$$

optionally bounded by the standard clamping limits (a factor of 4) as in Bitcoin [15].

Equation (6) is the standard proportional retarget applied to *active* work: the window's baseline difficulty $\bar{D}_k$ is adjusted by the orphan expansion factor W_k/W_k^{mc} and the pace factor T_{target}/T_k , which simplifies directly to pricing difficulty from observed work rate W_k/T_k . The buffer $P_2^{(k)}$ does not appear in T_k or in the main-chain term of W_k for this retarget; its main-chain work $d \cdot D_k$ and duration enter W_{k+1} and T_{k+1} . Uncles that can be reported only inside $P_2^{(k)}$ are, by Definition 1, still attributed to window k and therefore still affect D_{k+1} . Practical deployments may adopt various engineering techniques to prevent adversaries from manipulating timestamps; here, we simply adhere to standard Bitcoin terminology to preserve theoretical conciseness.

Design Rationale. Like Bitcoin, PB-DAM fixes epoch windows strictly by canonical main-chain block height (L blocks per epoch), rather than tracking a variable reported-block clock as in $\text{DAM}_1^{\text{mod}}$ [15]. Crucially, fixing the prefix boundary ∂_k solely by block height prevents an adversary from strategically manipulating or postponing the retarget cutoff by manufacturing orphans earlier in the epoch. The exclusion race (analyzed in Section 5) is therefore anchored to a deterministic point on the chain.

5 Security Analysis

We present our formal analysis in four parts: a quantitative tail bound on exclusion length across the settlement buffer (Lemma 1 and Theorem 1); unbiasedness of the difficulty estimator W_k (Theorem 2); the elimination of excess reward advantage under stationary full-power mining (Theorem 3); and an analytical upper bound on reward rates under stylized dynamic hash-rate scheduling (Proposition 1).

5.1 Exclusion Length

Definition 3 (Prefix exclusion). A prefix exclusion of length ℓ against window k occurs if there exists a fork rooted strictly before ∂_k such that (i) the last ℓ main-chain blocks leading up to and including the end of $P_2^{(k)}$ are adversarial, (ii) at least one honest block mined for $P_1^{(k)}$ is orphaned by that fork, and (iii) no main-chain block in $P_2^{(k)}$ embeds that orphan as a settlement uncle.

Condition (iii) fails as soon as any honest block enters $P_2^{(k)}$: by synchronous delivery and the honest embedding strategy (Section 3.1), that block observes and embeds every outstanding eligible uncle of the prefix. Hence a successful prefix exclusion requires the adversary to occupy all d buffer positions after orphaning at least one honest prefix block—a Nakamoto lead sustained for d consecutive main-chain steps. Let L_{LDC} denote the length of the longest dominant adversarial chain (LDC) that can cover an honest block [15]. While prior work modeled its mean behavior as $\mathbb{E}[L_{\text{LDC}}] = \Theta((\frac{1-\alpha}{1-2\alpha})^2)$ (Appendix A.1), sizing the buffer by the mean still leaves a constant failure probability. Instead, PB-DAM dimensions the settlement buffer d from a rigorous tail bound.

Lemma 1 (Exclusion tail). *There exist constants $c_1, c_2 > 0$ depending only on α^* such that, for every $\alpha \leq \alpha^*$ and every integer $d \geq 1$,*

$$\Pr[L_{\text{LDC}} \geq d] \leq c_1 e^{-c_2 d}. \quad (7)$$

In particular, for every $\varepsilon \in (0, 1)$ there is an integer

$$d^* \equiv d(\alpha^*, \varepsilon) := \left\lceil \frac{1}{c_2} \ln(c_1/\varepsilon) \right\rceil \quad (8)$$

such that $\Pr[L_{\text{LDC}} \geq d(\alpha^*, \varepsilon)] \leq \varepsilon$.

The exponential rate c_2 is the large-deviation rate of a biased random walk with step bias $1 - 2\alpha^* > 0$; a self-contained derivation from the Catalan structure is given in Appendix A.1. Lemma 1 is the reason a buffer must be sized from a *tail*, not from $\mathbb{E}[L_{\text{LDC}}]$: a constant-factor multiple of the mean still leaves a constant failure probability, and the mean itself diverges as $\alpha^* \rightarrow 1/2$.

Theorem 1 (Buffer defeats OEA). *Under our assumptions in Section 3.2, if $d \geq d(\alpha^*, \varepsilon)$, then for every estimation window k*

$$\Pr[\text{a prefix exclusion occurs against window } k] \leq \varepsilon.$$

The same bound applies to an attempted exclusion of the buffer content of window k from window $k+1$: such an attempt is a prefix exclusion against window $k+1$.

Proof (Intuition). A prefix exclusion requires the adversary to sustain a private lead across all d buffer blocks, which corresponds to an LDC of length at least d and is bounded by ε via Lemma 1 (full proof in Appendix A.2).

5.2 Unbiased Work Accounting

Theorem 2 (Unbiased estimate). *Under the hypotheses of Theorem 1, for any estimation window k , with probability at least $1 - \varepsilon$,*

$$W_k = W_k^*,$$

where W_k^* is the total proof-of-work of all valid blocks mined and broadcast in $\mathcal{W}_k$ (main-chain and orphan).

Proof (Intuition). Conditioned on the absence of a prefix exclusion, every orphan in $\mathcal{W}_k$ is embedded by the earliest honest main-chain block in $P_2^{(k)}$ or $P_1^{(k)}$, ensuring all active work is faithfully recorded in W_k (full proof in Appendix A.3).

Invariant 1 (No permanent blind spot). *Every main-chain block belongs to exactly one estimation window: $B_{k,j} \in \mathcal{W}_k$ for $j \leq L - d$ and $B_{k,j} \in \mathcal{W}_{k+1}$ for $j > L - d$.*

Invariant 1 is purely syntactic. Together with Theorem 2 it implies that a finite exclusion cannot remove work from every future retarget: at worst it defers accounting by one epoch, after which Theorem 1 applies again.

5.3 Bounding Adversarial Reward Advantage

Theorem 3 (No excess reward advantage). *Assume the conditions specified in Section 3.2 and $d \geq d(\alpha^*, \varepsilon)$. Then*

$$P_{\text{adv}}(\pi) \leq O(\varepsilon) \tag{9}$$

for every strategy π , including continuous selfish mining, intermittent selfish mining, and selfish mining with OEA; in particular, $\limsup_{\varepsilon \rightarrow 0} P_{\text{adv}}(\pi) \leq 0$.

Proof (Intuition). When active work is priced faithfully on the $1 - \varepsilon$ event, the reward rate is bounded by $M_{\mathcal{A}}^{\text{mc}} \leq \alpha$; on rare exclusion events occurring with probability $\leq \varepsilon$, the bounded single-window advantage contributes at most $O(\varepsilon)$ to the long-run Cesàro average (full proof in Appendix A.4).

Theorem 3 and Corollary 1 formally close the profitability of selfish mining under stationary regimes, strengthening the security guarantees of Sarenche et al. [15]. In their framework, perfect orphan reporting suppresses the advantage of continuous selfish mining to $P_{\text{adv}} \leq 0$, yet an adversary can still exploit the epoch boundary via SM+OEA to extract a strictly positive advantage of order L_{OEA}/L . Because PB-DAM faithfully prices active proof-of-work on the $1 - \varepsilon$ event, the fundamental identity $\text{RelRev} \cdot M^{\text{mc}} = M_{\mathcal{A}}^{\text{mc}} \leq \alpha$ holds unconditionally. Boundary exclusions succeed only with probability bounded by ε , scaling the asymptotic advantage of continuous and boundary-focused withholding down to $O(\varepsilon)$, ensuring $\limsup_{\varepsilon \rightarrow 0} P_{\text{adv}}(\pi) \leq 0$.

Corollary 1 (Absolute reward rate). *Under the hypotheses of Theorem 3, the adversary’s time-averaged block-reward rate satisfies*

$$R_{\mathcal{A}}(\pi) \leq \alpha\lambda K + O(\varepsilon), \quad (10)$$

regardless of its relative revenue $\text{RelRev}(\pi)$.

While Theorem 3 strictly eliminates stationary full-power selfish mining, an adversary could conceivably attempt dynamic hash-rate scheduling across consecutive epochs (e.g., alternating between withholding, idling, and full-power mining to induce difficulty oscillations). For example, an adversary may idle in epoch $i - 1$ to lower the subsequent difficulty, then mine at full power in epoch i , momentarily elevating its reward rate above α . To analyze whether such multi-epoch schedules can generate sustained reward gains, Proposition 1 establishes an analytical upper bound under a stylized first-order progress model.

Proposition 1 (Analytical Revenue Bound under Stylized Dynamic Scheduling). *Under the hypotheses of Theorem 2, conditional on the unbiased accounting event, consider an adversary scheduling its hash rate across windows $1, \dots, n$ into honest mining $b_i \geq 0$, selfish mining $a_i \geq 0$, and idling $\alpha - a_i - b_i \geq 0$. Under a stylized linear progress model where main-chain growth scales as $\alpha_H + b_i$ and window durations scale as $(\alpha_H + \delta_{i-1})/(\alpha_H + b_i)$ with $\delta_i = a_i + b_i \leq \alpha$, the normalized cumulative time-averaged reward satisfies:*

$$\frac{\sum_{i=1}^n p_i^{\text{reward}}}{\sum_{i=1}^n t_i} \leq \alpha. \quad (11)$$

Accounting for rare exclusion events across N windows, the asymptotic time-averaged block-reward rate remains bounded by $\alpha\lambda K + O(\varepsilon)$.

Proof (Intuition). An inductive surplus analysis shows that under linear progress approximations, the difficulty reduction engineered by idle epochs is strictly offset by the prolonged duration of preceding mining intervals, bounding the cumulative reward rate by α . Full inductive proof, analytical scope, and limitations are provided in Appendix B.

Implications for SISM2 and Enforceable Fairness. While Sarenche et al. [15] analyzed boundary manipulation only under a stylized alternating two-phase strategy (SISM2), our analysis demonstrates that the net profit advantage observed in SISM2 under complete orphan reporting does not stem from on-chain reward inflation. Rather, SISM2 exploits off-chain operational expenditure savings during idle epochs, which no PoW consensus can inspect (see Appendix B for full model and discussions). Because consensus rules govern on-chain block validities and cannot observe off-chain energy consumption, PB-DAM achieves maximal enforceable fairness.

Table 1. Steady-state performance under optimal selfish mining with OEA ($L = 256, \gamma = 1.0, \omega = 2.0$): normalized profit advantage P_{adv} , mean excluded honest orphans L_{OEA} , and exclusion success rate.

α	Bitcoin	Zeno ($\tau = 0.5$)	DAM ₁ ^{mod}		DAM ₂ ^{mod}		PB-DAM ($d = 96$)	
	P_{adv}	P_{adv}	P_{adv}	L_{OEA} (Succ.)	P_{adv}	L_{OEA} (Succ.)	P_{adv}	L_{OEA} (Succ.)
0.25	+0.0844	+0.0774	−0.0001	0.54 (22%)	+0.0010	0.88 (34%)	+0.0012	0.00 (0%)
0.30	+0.1263	+0.1171	+0.0003	1.09 (31%)	−0.0002	1.45 (40%)	−0.0006	0.00 (0%)
0.35	+0.1812	+0.1698	+0.0000	1.58 (35%)	−0.0014	2.27 (55%)	−0.0024	0.00 (0%)
0.40	+0.2587	+0.2451	+0.0031	3.40 (41%)	+0.0025	6.39 (64%)	−0.0022	0.00 (0%)
0.45	+0.3647	+0.3474	+0.0364	15.85 (50%)	+0.0235	24.38 (81%)	+0.0008	0.01 (0%)

6 Evaluation

We evaluate PB-DAM through discrete-event simulations, comparing its performance against standard Bitcoin retargeting, the modified orphan-aware mechanisms DAM₁^{mod} and DAM₂^{mod} of Sarenche et al. [15], and Zeno DAA [1], representing the primary defense paradigms analyzed in Section 2.

6.1 Experimental Setup

We evaluate all mechanisms using an event-driven proof-of-work blockchain simulator adhering to the models in Section 3. Block arrivals by honest miners and the adversary follow independent Poisson point processes with rates $(1 - \alpha)\lambda/D$ and $\alpha\lambda/D$, respectively, under propagation parameter $\gamma = 1.0$. To model revenue-maximizing behavior under full mining power, the adversary executes optimal selfish mining (OSM) via the MDP solver of Sapirshtein et al. [14], combined with the boundary orphan exclusion attack (OEA) of Sarenche et al. [15]. Benchmark parameters are set to unit target interval $\tau_0 = 1.0$, nominal reward $K = 1.0$, arrival intensity $\lambda = 1.0$, and profitability factor $\omega = \lambda K/c = 2.0$, where c is the unit mining operational cost rate, making ω the ratio of nominal network revenue to operating expenditure. We evaluate two primary metrics: (1) normalized profit advantage P_{adv} measuring excess net profit relative to honest mining; and (2) mean excluded honest orphans L_{OEA} measuring honest orphans permanently barred from difficulty retargets. Detailed algorithmic state spaces, MDP transitions, operational cost formulations, and confidence interval methodology are provided in Appendix C.1.

6.2 Standard-Epoch Benchmark ($L = 256$): Steady-State Defense

We first evaluate steady-state mining under a constant total hash rate with epoch length $L = 256$ across 25 measured epochs following 5 warmup epochs over 15 independent replicas. Table 1 reports normalized profit advantage P_{adv} , mean excluded honest orphans L_{OEA} , and exclusion success rates across all evaluated mechanisms.

Table 2. Short-epoch benchmark ($L = 64, \gamma = 1.0, \omega = 2.0$): profit advantage P_{adv} and mean excluded honest orphans L_{OEA} .

α	$\text{DAM}_1^{\text{mod}}$		$\text{DAM}_2^{\text{mod}}$		PB-DAM ($d = 8$)		PB-DAM ($d = 16$)	
	P_{adv}	L_{OEA}	P_{adv}	L_{OEA}	P_{adv}	L_{OEA}	P_{adv}	L_{OEA}
0.25	-0.0001	0.67	-0.0012	0.80	+0.0002	0.02	+0.0002	0.00
0.30	+0.0026	1.16	+0.0002	1.75	-0.0010	0.25	-0.0013	0.04
0.35	+0.0064	1.40	-0.0027	2.13	-0.0049	0.38	-0.0052	0.09
0.40	+0.0265	3.73	+0.0168	6.33	-0.0007	2.83	-0.0005	1.31
0.45	+0.0847	12.05	+0.0693	17.45	+0.0210	12.56	+0.0212	9.32

Vulnerability of Bitcoin and Zeno DAA. As shown in Table 1, Bitcoin is severely vulnerable to optimal selfish mining because it ignores orphaned proof-of-work, yielding profit advantages up to $P_{\text{adv}} = +36.47\%$ at $\alpha = 0.45$ that match the theoretical bound $\alpha/(1-\alpha)$ [14]. Zeno DAA [1] displays nearly identical steady-state returns ($P_{\text{adv}} = +34.74\%$ at $\alpha = 0.45$): while its damping parameter ($\tau = 0.5$) extends transient profit lag, it converges to the same fixed-point difficulty invariant as Bitcoin once retargets settle.

Boundary Exclusions in Modified DAMs. The modified DAMs of Sarenche et al. [15] suppress selfish mining at moderate hash rates ($\alpha \leq 0.35$). However, their same-epoch reporting constraint fails at higher hash rates: at $\alpha = 0.45$, adversarial forks frequently cross epoch boundaries, permanently excluding an average of 15.85 ($\text{DAM}_1^{\text{mod}}$) and 24.38 ($\text{DAM}_2^{\text{mod}}$) honest blocks from retarget calculations. This boundary gap depresses the retargeted difficulty, restoring significant profit advantages up to +3.64%.

PB-DAM Eliminates Exclusions. PB-DAM resolves this vulnerability through its pipelined settlement buffer. At buffer depth $d = 96$, PB-DAM eliminates exclusions entirely ($L_{\text{OEA}} = 0.00$) for all $\alpha \leq 0.40$ and suppresses exclusion success to virtually 0% ($L_{\text{OEA}} = 0.01$) even at $\alpha = 0.45$. In accordance with Theorem 3, the adversary’s advantage remains confined to statistical parity with honest mining ($-0.0024 \leq P_{\text{adv}} \leq +0.0012$).

6.3 Short-Epoch Benchmark ($L = 64$): Boundary Gap Exploitation

In high-throughput proof-of-work protocols, difficulty adjustment epochs are often shortened to $L \in [64, 128]$ to track hash-rate dynamics responsively. However, short epochs amplify boundary exclusion effects, as the relative difficulty distortion scales as $O(L_{\text{OEA}}/L)$. Table 2 evaluates all protocols under $L = 64$ over 30 epochs across 15 independent replicas.

The empirical results highlight the breakdown of same-epoch reporting in short epochs. Under $\text{DAM}_1^{\text{mod}}$ and $\text{DAM}_2^{\text{mod}}$ at $\alpha = 0.45$, the adversary excludes 19% to 27% of the entire epoch from difficulty calculations ($L_{\text{OEA}} = 12.05$ and

17.45), extracting substantial profit advantages of +8.47% and +6.93%, respectively. Even at $\alpha = 0.40$, same-epoch exclusions yield significant advantages exceeding +1.6% to +2.6%.

In contrast, PB-DAM with buffer depth $d = 16$ (25% of the epoch) maintains statistical parity ($P_{\text{adv}} \leq 0.02\%$) for all $\alpha \leq 0.40$. At $\alpha = 0.45$, PB-DAM cuts excluded orphans nearly in half and restricts the profit margin to +2.12%. These results confirm that PB-DAM provides structural protection against boundary exploitation, eliminating reliance on epoch dilution in long windows.

6.4 Sensitivity, Liveness, and System Overhead

Parameter Sensitivity and Buffer Scaling Law. Lemma 1 and Theorem 1 establish that the probability of an adversarial withholding race surviving d consecutive main-chain blocks decays exponentially as $c_1 e^{-c_2 d}$ with drift $1 - 2\alpha$. Our empirical parameter sweeps across buffer depths $d \in \{0, \dots, 128\}$ at $L = 256$ strictly validate this exponential suppression: setting $d \in [16, 32]$ provides robust exclusion resistance against mining pools with $\alpha \leq 1/3$, while $d \geq 64$ secures networks under higher adversarial thresholds up to $\alpha = 0.45$. At $d = 128$, excluded orphans under $\alpha = 0.45$ fall from 28.35 ($d = 0$) down to 0.92, driving excess profit negative ($P_{\text{adv}} = -0.0041$). The full sweep across 28 parameter configurations is reported in Appendix C.2.

Dynamic Hash Rate Shocks and Liveness Recovery. Preserving chain throughput during abrupt network contractions (e.g., power outages or geopolitical mining bans) is a vital operational requirement. Evaluating an honest network ($\alpha = 0$) subjected to an instantaneous 50% drop in total hash rate at $L = 128$ reveals a stark divergence between defense paradigms. Zeno DAA’s downward damping ($\tau = 0.5$) severely impairs chain liveness, requiring 4.5 ± 1.1 epochs to recover difficulty to within 10% of nominal parity and leaving block intervals sluggish (1.60 s in the post-crash epoch). In contrast, PB-DAM avoids artificial deceleration by relying on complete proof-of-work accounting, adapting smoothly within 2.2 ± 0.8 epochs with only a bounded one-window pipeline lag. Complete adaptation trajectories and recovery metrics are detailed in Appendix C.3.

System Overhead and Synthesis. Table 3 synthesizes the architectural properties of all evaluated mechanisms. The system overhead of PB-DAM is minimal: settlement uncle commitments require 32-byte header hashes embedded in main-chain blocks, consuming under 1% of block capacity. Pipelining introduces an adjustment lag of d blocks, yielding a modest latency fraction d/L of 12.5% for $d = 32, L = 256$ and 6.3% for $d = 128, L = 2016$. Because retarget evaluation occurs strictly once per epoch at epoch boundaries, PB-DAM incurs zero per-block computational overhead.

7 Discussion

Economic Deterrence versus Chain Quality. PB-DAM modifies only difficulty retargeting and does not affect the fork-choice rule or block propagation; thus, it

Table 3. Architectural comparison of difficulty adjustment mechanisms.

Property	Bitcoin	Zeno DAA	DAM ₁ ^{mod}	DAM ₂ ^{mod}	PB-DAM
Orphan-Aware	No	No	Yes (M^{total})	Yes (M^{chain})	Yes (Pipelined Window)
OEA Resilient	N/A	Vulnerable	Boundary Gap	Boundary Gap	Provable ($d \geq d^*$)
Cross-Epoch Orphans	None	None	Same-Epoch Bar	Same-Epoch Bar	Settlement Buffer (d)
Short-Epoch ($L \leq 64$)	$P_{\text{adv}} > 34\%$	$P_{\text{adv}} > 34\%$	$P_{\text{adv}} \approx 8.5\%$	$P_{\text{adv}} \approx 6.9\%$	$P_{\text{adv}} \leq 0.02\%$ ($\alpha \leq 0.40$)
Hash-Rate Crash Recovery	Immediate	Sluggish (τ)	Immediate	Immediate	Smooth & Rapid
Security Guarantee	None	Fixed Point	Heuristic	Heuristic	Exponential Tail (Lemma 1)

does not improve chain quality or prevent adversarial orphaning of honest blocks. Attacks involving deliberate financial loss, such as denial-of-service or subsidized suppression, remain possible. PB-DAM could potentially complement protocols that secure chain quality.

Off-Chain Cost Asymmetries and Coin-Hopping. What PB-DAM enforces is maximal on-chain revenue fairness. The residual profit advantage under SISM2 reflects pure off-chain electricity savings during high-difficulty windows rather than protocol-level currency inflation. Defending against such cost arbitrage requires cross-chain coordination or non-outsourcable work commitments.

Network Delay and Natural Orphans. In practice, non-zero delay reduces effective honest mining power to $\beta_{\text{eff}} \approx (1 - \alpha)e^{-\lambda\Delta}$, which lengthens the excursion tail (Lemma 1) and increases the required buffer depth d by a constant factor for the same error bound ε . We separate this practical consideration to clarify that PB-DAM addresses the structural boundary Orphan Exclusion Attack, which exists even without natural orphans.

Pipelining Latency and Hash-Rate Shocks. An abrupt hash-rate collapse during the buffer cannot freeze the chain. Crucially, the buffer depth represents a small fraction of the epoch ($d/L \leq 6.3\%$ in standard configurations), and retargeting occurs strictly once per L main-chain blocks rather than being delayed. As demonstrated by our dynamic shock experiments in Section 6.4, PB-DAM adapts smoothly within a single-window pipeline lag. Furthermore, timestamps remain bounded by Nakamoto Median Time Past rules, preventing adversaries from manipulating target calculations across pipelined windows.

Transaction Fees and MEV Dynamics. In fee-dominated or MEV-rich regimes, miners may deviate from honest mining to capture lucrative mempool transactions or execute undercutting attacks [3]. Such deviations are triggered by short-term fee heterogeneity across consecutive blocks. Mitigating these micro-level attacks requires transaction-layer primitives such as fee smoothing or minimum base fees, which are orthogonal and complementary to the consensus-layer guarantees of PB-DAM.

Extension to Sliding-Window DAAs. Our design directly generalizes to sliding-window DAAs. In a sliding-window formulation, the target at height h is evaluated over the interval $[h - L, h - d]$, using the terminal d blocks as a moving

settlement buffer. Uncles originating from blocks up to height $h - d$ remain eligible for inclusion until height h , and work minted within the moving buffer enters the active window as the sliding horizon advances.

8 Conclusion

We resolve the epoch-boundary orphan exclusion vulnerability in Nakamoto consensus through PB-DAM, which pipelines difficulty estimation across consecutive epochs via a settlement buffer. PB-DAM restores unbiased active-work accounting and drives the time-averaged profit advantage of selfish mining to zero up to $O(\varepsilon)$, even in modern short-epoch regimes.

These results establish that difficulty adjustment can enforce maximal on-chain reward fairness without heuristic damping or liveness degradation during network hash-rate shocks. PB-DAM provably eliminates excess reward advantage under stationary selfish mining and bounds gross block emissions under analytical dynamic scheduling models. Promising directions for future work include formalizing game-theoretic incentives for uncle embedding, calibrating buffer parameters under realistic propagation delays, and extending the analysis to a full cryptographic backbone model with dynamic participant sets.

References

1. Azimy, H., Ghorbani, A.: Alternative difficulty adjustment algorithms for preventing selfish mining attack. In: Blockchain – ICBC 2021. LNCS, vol. 12991, pp. 59–73. Springer (2022). https://doi.org/10.1007/978-3-030-96527-3_5
2. Bai, Q., Xu, Y., Liu, N., Wang, X.: Blockchain mining With multiple selfish miners. IEEE Transactions on Information Forensics and Security **18**, 3116–3131 (2023). <https://doi.org/10.1109/TIFS.2023.3275736>
3. Carlsten, M., Kalodner, H., Weinberg, S.M., Narayanan, A.: On the instability of Bitcoin without the block reward. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 154–167. Association for Computing Machinery (2016). <https://doi.org/10.1145/2976749.2978408>
4. Eyal, I., Sirer, E.G.: Majority is not enough: Bitcoin mining is vulnerable. In: Financial Cryptography and Data Security. LNCS, vol. 8437, pp. 436–454. Springer (2014). https://doi.org/10.1007/978-3-662-45472-5_28
5. Garay, J.A., Kiayias, A., Leonardos, N.: The Bitcoin backbone protocol: Analysis and applications. In: Advances in Cryptology – EUROCRYPT 2015. LNCS, vol. 9057, pp. 281–310. Springer (2015). https://doi.org/10.1007/978-3-662-46803-6_10
6. Garay, J.A., Kiayias, A., Leonardos, N.: The Bitcoin backbone protocol with chains of variable difficulty. In: Advances in Cryptology – CRYPTO 2017. LNCS, vol. 10401, pp. 291–323. Springer (2017). https://doi.org/10.1007/978-3-319-63688-7_10
7. Gazi, P., Ren, L., Russell, A.: Practical settlement bounds for proof-of-work blockchains. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security. pp. 1217–1230. ACM (2022). <https://doi.org/10.1145/3548606.3559368>

8. Gervais, A., Karame, G.O., Wüst, K., Glykantzis, V., Ritzdorf, H., Capkun, S.: On the security and performance of proof of work blockchains. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 3–16. Association for Computing Machinery (2016). <https://doi.org/10.1145/2976749.2978341>
9. Grunspan, C., Pérez-Marco, R.: On profitability of selfish mining (2018). <https://doi.org/10.48550/arXiv.1805.08281>, <https://arxiv.org/abs/1805.08281>
10. Grunspan, C., Pérez-Marco, R.: Profit lag and alternate network mining. In: Mathematical Research for Blockchain Economy. pp. 115–132. Lecture Notes in Operations Research, Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-48731-6_7, https://doi.org/10.1007/978-3-031-48731-6_7
11. Kraft, D.: Difficulty control for blockchain-based consensus systems. Peer-to-Peer Networking and Applications **9**(2), 397–413 (2016). <https://doi.org/10.1007/s12083-015-0347-x>
12. Negy, K.A., Rizun, P.R., Sirer, E.G.: Selfish mining re-examined. In: Financial Cryptography and Data Security. LNCS, vol. 12059, pp. 61–78. Springer (2020). https://doi.org/10.1007/978-3-030-51280-4_5
13. Ren, L.: Analysis of Nakamoto consensus. Cryptology ePrint Archive, Paper 2019/943 (2019), <https://eprint.iacr.org/2019/943>
14. Sapirshstein, A., Sompolinsky, Y., Zohar, A.: Optimal selfish mining strategies in Bitcoin. In: Financial Cryptography and Data Security. LNCS, vol. 9603, pp. 515–532. Springer (2017). https://doi.org/10.1007/978-3-662-54970-4_30
15. Sarenche, R., Zhang, R., Nikova, S., Preneel, B.: Selfish mining time-averaged analysis in Bitcoin: Is orphan reporting an effective countermeasure? IEEE Transactions on Information Forensics and Security **20**, 449–464 (2025). <https://doi.org/10.1109/TIFS.2024.3518090>
16. Wang, J., Wu, S., Liang, H., Ding, Y., Zhai, Y.: Adaptive mining difficulty for blockchain to resist selfish mining attack. Journal of Surveillance, Security and Safety **3**(4), 14–34 (2023). <https://doi.org/10.20517/jsss.2023.05>
17. Zhang, R., Zhang, D., Wang, Q., Wu, S., Xie, J., Preneel, B.: NC-Max: Breaking the security-performance tradeoff in Nakamoto consensus. In: Proceedings of the 2022 Network and Distributed System Security Symposium. Internet Society, Reston, VA, USA (2022). <https://doi.org/10.14722/ndss.2022.24370>, <https://doi.org/10.14722/ndss.2022.24370>

A Deferred Proofs

A.1 Proof of Lemma 1

Let L_{LDC} denote the length of the longest dominant adversarial chain (LDC) that can cover a designated honest block at a fixed height, in the sense of Sarenche et al. (Appendices E–F of [15]). In prior work [15], the expected boundary exclusion length $\mathbb{E}[L_{\text{OEA}}]$ under an unbuffered DAM scales asymptotically as

$$\mathbb{E}[L_{\text{OEA}}] = \Theta\left(\left(\frac{1-\alpha}{1-2\alpha}\right)^2\right), \quad (12)$$

evaluating under $\text{DAM}_1^{\text{mod}}$ to approximately 1.5, 3, 7, 18, and 87 blocks at $\alpha = 0.25, 0.30, 0.35, 0.40$, and 0.45 , respectively. However, protecting against adversarial exclusion requires bounding the tail distribution rather than the mean.

View the race between adversarial and honest mining, conditional on a fork that must cover a fixed honest block B^* , as a simple random walk S_n on $\mathbb{Z}$ with $\Pr[+1] = \alpha$ (adversary extends its lead) and $\Pr[-1] = 1 - \alpha$ (honest side catches up). A dominant adversarial chain of length at least d that still covers B^* requires the walk to stay strictly positive for d consecutive steps after the step that orphans B^* , or equivalently requires the hitting time of -1 to exceed d in the meander that defines an LDC in [15]. For $\alpha \leq \alpha^* < 1/2$ the step bias $\mu = 1 - 2\alpha \geq 1 - 2\alpha^* > 0$ toward the honest side. Standard Chernoff bounds on the binomial counts in any window of d steps give

$$\Pr[\text{adversary wins all } d \text{ steps}] = \alpha^d \leq e^{-d \ln(1/\alpha^*)},$$

and the probability of any Dyck-compliant path of length $\geq d$ (the Catalan weighted sum underlying [15, App. E]) is dominated by a geometric series with ratio $4\alpha(1 - \alpha) < 1$. Collecting constants into c_1 and rate $c_2 = -\frac{1}{2} \ln(4\alpha^*(1 - \alpha^*)) > 0$ yields (7). The same rate appears in the asymptotic scaling (12): the expectation is finite precisely when $c_2 > 0$, and diverges as $\alpha^* \rightarrow 1/2^-$.

A.2 Proof of Theorem 1

By Definition 3 and the honest embedding strategy under synchronous delivery, a prefix exclusion requires that $P_2^{(k)}$ contain no honest main-chain block that can embed the orphaned honest prefix blocks. Hence the main chain through the buffer is entirely adversarial and extends a fork that orphaned at least one honest prefix block: this is an LDC of length at least d in the sense of [15]. Lemma 1 applies verbatim at $\alpha \leq \alpha^*$. Exclusions aimed at buffer orphans of window k must succeed against the prefix boundary of window $k + 1$ (Invariant 1), which again demands an LDC of length $\geq d$.

A.3 Proof of Theorem 2

Main-chain work in $\mathcal{W}_k$ contributes $(L - d)D_k + dD_{k-1}$ by construction. Let U be an orphan mined in $\mathcal{W}_k$. Because the adversary does not discard its own unbroadcast blocks, all orphaned blocks are eventually broadcast to the network. If a prefix exclusion does not occur, there exists at least one honest main-chain block in $P_2^{(k)}$ (when U belongs to the current prefix) or in $P_1^{(k)}$ (when U belongs to the previous buffer) that sees U 's parent on-chain. Definition 1 then marks U eligible, and the honest miner embeds it by the honest strategy and synchronous delivery. Summing over all such U gives $W_k = W_k^*$. Theorem 1 bounds the complementary probability by ε .

A.4 Proof of Theorem 3

Write $\mathcal{E}_k$ for the exclusion event of window k , with $\Pr[\mathcal{E}_k] \leq \varepsilon$. On $\bar{\mathcal{E}}_k$, Theorem 2 gives unbiased active-work accounting. In a regime where difficulty has

equilibrated, the argument of [15, Sec. V] applies unchanged: the adversary’s reward rate equals $R_{\mathcal{A}} = \lambda K \cdot M_{\mathcal{A}}^{\text{mc}} \leq \alpha \lambda K$ by (2). By (1), this immediately implies that the conditional reward advantage on $\bar{\mathcal{E}}_k$ is non-positive. On $\mathcal{E}_k$, the one-window advantage is at most the SM+OEA advantage $P_{\text{adv}}^{\text{OEA}} = O(1)$ of [15, Eqs. (46), (51)] (bounded by 1 after normalization by λK). Over N windows the Cesàro mean is therefore at most

$$(1 - \varepsilon) \cdot 0 + \varepsilon \cdot P_{\text{adv}}^{\text{OEA}} + o(1) = O(\varepsilon),$$

where the $o(1)$ accounts for the mixing time of the difficulty process. Sending $N \rightarrow \infty$ yields (9). Independence across well-separated windows is not required: even under arbitrary dependence, $\frac{1}{N} \sum_{k=1}^N \mathbf{1}_{\mathcal{E}_k}$ has expectation $\leq \varepsilon$, and bounded per-window advantage gives the same $O(\varepsilon)$ bound on expected average advantage.

B Analytical Revenue Bound under Stylized Dynamic Scheduling

In this section, we present the formal analytical derivation for Proposition 1. We analyze whether an adversary can obtain a time-averaged block-reward rate strictly exceeding its nominal hash-rate share α by dynamically scheduling its hash rate across consecutive epochs under a stylized first-order progress model. Specifically, on the unbiased active-work accounting event guaranteed by Theorem 2 (under which all orphaned blocks are incorporated into the retargeting mechanism via the settlement buffer), we examine whether dynamic allocations among honest mining, selfish mining, and idling can increase the adversary’s long-run time-averaged gross reward rate beyond α .

B.1 Adversarial Scheduling Model and Upper-Bound Dynamics

Consider an attack spanning n consecutive epochs, indexed by $i \in \{1, 2, \dots, n\}$. Let $\alpha_H = 1 - \alpha$ denote the honest miners’ hash-rate share, where $\alpha_H > 0.5$ by our assumptions in Section 3.2. In each epoch i , the adversary partitions its hash-rate share α into three components:

- $b_i \geq 0$: honest mining power, which extends the main chain without causing orphans;
- $a_i \geq 0$: selfish mining power, which extends the main chain and invalidates honest blocks;
- $\alpha - a_i - b_i \geq 0$: temporarily idle hash rate, aiming to depress the difficulty estimator for subsequent epochs.

Let $\delta_i := a_i + b_i \leq \alpha = 1 - \alpha_H$ denote the adversary’s total active mining power share in epoch i . Epoch 0 represents the steady-state baseline prior to the attack, where the adversary mines honestly with full power, so $b_0 = \alpha = 1 - \alpha_H$, $a_0 = 0$, and $\delta_0 = 1 - \alpha_H$. The nominal epoch block reward R and the nominal

duration t_0 are normalized to 1, such that the honest time-averaged reward rate is $\alpha \cdot R/t_0 = \alpha = 1 - \alpha_H$.

To establish an analytical upper bound, we adopt a stylized first-order approximation of the adversary's main-chain productivity. We emphasize that this analytical formulation abstracts fine-grained Markovian fork competitions by approximating main-chain progress as a linear combination $\alpha_H + b_i$ and assumes single-window difficulty coupling, omitting the non-linear stationary probabilities of MDP fork races and the d -block pipelining phase lag of PB-DAM. Under these structural assumptions, the upper-bound dynamics proceed as follows:

1. **Adversarial reward share p_i^{reward} :** We assume that the adversary's selfish mining power a_i successfully orphans an equal share a_i of honest blocks while incurring zero orphan losses on its own blocks. Under this maximally aggressive assumption, honest miners contribute $\alpha_H - a_i$ to main-chain growth, while the adversary contributes $a_i + b_i = \delta_i$. The total main-chain growth rate is $(\alpha_H - a_i) + (a_i + b_i) = \alpha_H + b_i$, and the fraction of epoch rewards captured by the adversary is

$$p_i^{\text{reward}} = \frac{a_i + b_i}{\alpha_H + b_i} = \frac{\delta_i}{\alpha_H + b_i}. \quad (13)$$

In realistic MDP selfish mining, the adversary's private blocks may also be orphaned, yielding an effective main-chain share bounded by $\text{RelRev}(\pi) \cdot M^{\text{mc}} \leq \alpha$ [15]. Equation (13) represents a strict upper bound on the adversary's reward share.

2. **Epoch duration t_i :** Under orphan-aware difficulty adjustment with complete reporting, the difficulty for epoch i is proportional to the total work (main-chain plus uncle blocks) observed in epoch $i - 1$. The honest network produces α_H work (comprising $\alpha_H - a_{i-1}$ main-chain blocks and a_{i-1} orphans), while the adversary produces $a_{i-1} + b_{i-1} = \delta_{i-1}$ work. The total work counted in epoch $i - 1$ is $(\alpha_H - a_{i-1}) + a_{i-1} + \delta_{i-1} = \alpha_H + \delta_{i-1}$. Because the main chain in epoch i advances at effective rate $\alpha_H + b_i$, the normalized elapsed time required to generate the L main-chain blocks of epoch i is

$$t_i = \frac{\alpha_H + a_{i-1} + b_{i-1}}{\alpha_H + b_i} = \frac{\alpha_H + \delta_{i-1}}{\alpha_H + b_i}. \quad (14)$$

B.2 Inductive Proof of Revenue Invariance

Proof (Proof of Proposition 1). Under the dynamics of (13) and (14), for any integer $n \geq 1$ and any sequence of adversarial actions $\{a_i, b_i\}_{i=1}^n$ satisfying $a_i, b_i \geq 0$ and $\delta_i = a_i + b_i \leq 1 - \alpha_H = \alpha$, we establish that $\sum_{i=1}^n p_i^{\text{reward}} / \sum_{i=1}^n t_i \leq \alpha$. Multiplying both sides by $\sum_{i=1}^n t_i > 0$, this inequality is equivalent to $\sigma_n \leq 0$, where

$$\sigma_k := \sum_{i=1}^k p_i^{\text{reward}} - (1 - \alpha_H) \sum_{i=1}^k t_i. \quad (15)$$

We prove a stronger inductive hypothesis: for every $k \geq 1$,

$$\sigma_k \leq \alpha_H + \delta_k - 1. \quad (16)$$

Since $\delta_k \leq 1 - \alpha_H$, the right-hand side satisfies $\alpha_H + \delta_k - 1 \leq \alpha_H + (1 - \alpha_H) - 1 = 0$. Hence, establishing (16) for all $k \geq 1$ directly implies $\sigma_n \leq 0$.

Base Case ($k = 1$). Using $\delta_0 = 1 - \alpha_H$ (so $\alpha_H + \delta_0 = 1$), we have:

$$\sigma_1 = p_1^{\text{reward}} - (1 - \alpha_H)t_1 = \frac{\delta_1 - (1 - \alpha_H)(\alpha_H + \delta_0)}{\alpha_H + b_1} = \frac{\alpha_H + \delta_1 - 1}{\alpha_H + b_1}. \quad (17)$$

Because $\delta_1 \leq 1 - \alpha_H$, the numerator satisfies $\alpha_H + \delta_1 - 1 \leq 0$. Moreover, honest majority ensures $\alpha_H > 0.5$, and since $b_1 \leq \delta_1 \leq 1 - \alpha_H$, the denominator satisfies $0 < \alpha_H \leq \alpha_H + b_1 \leq 1$. Dividing a non-positive quantity by a real number in $(0, 1]$ cannot increase its value:

$$\sigma_1 = \frac{\alpha_H + \delta_1 - 1}{\alpha_H + b_1} \leq \alpha_H + \delta_1 - 1. \quad (18)$$

Thus (16) holds for $k = 1$.

Inductive Step. Assume that (16) holds for $k - 1$, i.e., $\sigma_{k-1} \leq \alpha_H + \delta_{k-1} - 1$, where $k \geq 2$. For epoch k :

$$\sigma_k = \sigma_{k-1} + p_k^{\text{reward}} - (1 - \alpha_H)t_k \leq \alpha_H + \delta_{k-1} - 1 + \frac{\delta_k - (1 - \alpha_H)(\alpha_H + \delta_{k-1})}{\alpha_H + b_k}. \quad (19)$$

To establish $\sigma_k \leq \alpha_H + \delta_k - 1$, it suffices to demonstrate:

$$\frac{\delta_k - (1 - \alpha_H)(\alpha_H + \delta_{k-1})}{\alpha_H + b_k} \leq \delta_k - \delta_{k-1}. \quad (20)$$

Multiplying both sides by $\alpha_H + b_k > 0$ and rearranging terms:

$$F := (\alpha_H + b_k - 1)(\delta_{k-1} - \delta_k) + \alpha_H(\alpha_H + \delta_{k-1} - 1) \leq 0. \quad (21)$$

Recall that $\alpha_H \in (0.5, 1)$, $0 \leq b_k \leq \delta_k \leq 1 - \alpha_H$, and $0 \leq \delta_{k-1} \leq 1 - \alpha_H$, which implies $\alpha_H + b_k - 1 \leq 0$ and $\alpha_H + \delta_{k-1} - 1 \leq 0$. We evaluate F under two complementary cases:

- **Case 1** ($\delta_{k-1} \geq \delta_k$): Here $\delta_{k-1} - \delta_k \geq 0$. Because $(\alpha_H + b_k - 1)(\delta_{k-1} - \delta_k) \leq 0$ and $\alpha_H(\alpha_H + \delta_{k-1} - 1) \leq 0$, the sum of two non-positive quantities satisfies $F \leq 0$.
- **Case 2** ($\delta_{k-1} < \delta_k$): Define $\epsilon_k := \delta_k - \delta_{k-1} > 0$. Substituting $\delta_{k-1} = \delta_k - \epsilon_k$ into (21) yields:

$$F = \alpha_H(\alpha_H + \delta_k - 1) - \epsilon_k(2\alpha_H + b_k - 1). \quad (22)$$

Since $\delta_k \leq 1 - \alpha_H$, the first term satisfies $\alpha_H(\alpha_H + \delta_k - 1) \leq 0$. Because honest majority ensures $\alpha_H > 0.5$, we have $2\alpha_H + b_k - 1 \geq 2\alpha_H - 1 > 0$. Since $\epsilon_k > 0$, the subtracted quantity $\epsilon_k(2\alpha_H + b_k - 1)$ is strictly positive, implying $-\epsilon_k(2\alpha_H + b_k - 1) < 0$. Thus $F < 0$.

In both cases, $F \leq 0$ holds unconditionally. By induction, $\sigma_n \leq \alpha_H + \delta_n - 1 \leq 0$ holds for all $n \geq 1$, completing the proof. $\square$

B.3 Analytical Scope and Limitations

Proposition 1 elucidates the structural counterbalancing mechanism between difficulty discount and interval elongation. However, we explicitly highlight the idealized assumptions of this analytical approximation:

1. *Linear Progress Approximation:* The model abstracts fine-grained MDP fork dynamics by approximating main-chain progress as a linear combination $\alpha_H + b_i$. In realistic selfish mining, tie-breaking and probabilistic invalidations introduce non-linearities into effective hash rate M^{mc} .
2. *Pipelining Phase Lag:* The stylized formulation assumes difficulty is governed strictly by the immediate preceding window. In PB-DAM, the settlement buffer $P_2^{(k-1)}$ introduces a phase lag of depth d across window boundaries $(\partial_{k-1}, \partial_k]$, creating subtle inter-epoch coupling that the first-order recurrence neglects.

Deriving a closed-form reward invariant over a continuous-time Markov decision process (CTMDP) under arbitrary dynamic arrivals across pipelined windows remains an open theoretical problem. Nonetheless, Proposition 1 captures the primary first-order invariance, and our discrete-event simulations in Section 6 evaluate the full non-linear MDP under dynamic SISM2 strategies.

B.4 Implications: Revenue Rate vs. Profit under SISM2

Proposition 1 clarifies the economic boundary of difficulty adjustment mechanisms and yields several vital insights into multi-epoch adversarial manipulation:

1. **Generalization beyond stylized two-phase SISM2:** Sarenche et al. [15] demonstrated that naive orphan-aware adjustment remains vulnerable to Smart Intermittent Selfish Mining (SISM2). However, their investigation was restricted to a stylized alternating two-phase schedule where the adversary mines selfishly in odd epochs and idles in even epochs. They left open whether arbitrary multi-epoch schedules—such as varying idle fractions or multi-epoch withholding campaigns—could extract unbounded rewards. Proposition 1 addresses this question by establishing a structural invariance under linear progress approximations: across *any* sequence of n epochs and under *any* arbitrary adversarial allocation among honest mining, selfish withholding, and idling, the time-averaged gross reward rate cannot exceed the nominal hash-rate share $\alpha\lambda K$. The difficulty reduction engineered by an idle phase is mathematically offset by the prolonged duration of preceding epochs, enforcing the universal invariant $\sum p_i / \sum t_i \leq \alpha$.
2. **Decoupling of gross revenue and net profit in SISM2:** Proposition 1 shows that the net profit advantage ($P_{\text{adv}} > 0$) observed in SISM2 under complete orphan reporting does *not* arise from harvesting excess on-chain block rewards over time. Rather, SISM2 exploits the off-chain operational cost structure of proof-of-work: by powering off hardware during idle epochs, the adversary avoids electricity costs while inducing a difficulty dip

that shortens subsequent mining intervals. Because consensus rules operate strictly on the ledger state and cannot observe external operational expenditures, difficulty adjustment mechanisms cannot prevent profit advantages arising purely from exogenous electricity cost avoidance.

3. **Economic parameter regimes and viability of SISM2:** Furthermore, SISM2 yields a positive net profit advantage only in parameter regimes where operational expenditures constitute a large fraction of nominal revenues (i.e., low network profitability factor $\omega = \lambda K/c$). In highly profitable mining regimes (large ω), the absolute value of block subsidies foregone during idle epochs far outweighs any subsequent difficulty savings, rendering idling strictly sub-optimal and driving the net profit advantage negative.
4. **Maximal enforceable fairness:** PB-DAM thus achieves the maximal fairness guarantee enforceable at the consensus layer: it eliminates all protocol-level reward distortion and ensures that no adversarial scheduling can extract newly minted currency beyond the adversary’s nominal hash-rate share.

C Extended Experimental Evaluation and Empirical Robustness

This appendix provides extended formulations, methodologies, and complete experimental sweeps complementing the evaluation in Section 6.

C.1 Detailed Simulation Framework and MDP Threat Model

Simulator Implementation. The discrete-event simulator extends the network model in Section 3 by maintaining explicit directed acyclic graph (DAG) ancestry for all blocks and tracking uncle commitments in main-chain headers under Definition 1. To ensure statistical significance, every experimental configuration is evaluated across 15 independent replicas initialized with distinct pseudo-random seeds. Sample means are reported alongside 95% bootstrap confidence intervals computed over 1000 resamples.

Adversarial Mining Formulation via MDP. The adversary executes optimal selfish mining (OSM) under the MDP model of Sapirshstein et al. [14]. At each decision epoch, state (a, h, fork) tracks private branch length a , honest branch length h , and fork status $\text{fork} \in \{\text{irrelevant}, \text{relevant}, \text{active}\}$, with actions adopt, override, match, and wait. At epoch boundaries under $\gamma = 1.0$, the adversary withholds private leads without premature overrides and matches at parity to win tie-breaks, maximizing excluded honest blocks consistent with the strongest boundary attacker in [15, App. D].

Economic Framework and Operational Costs. To better reflect realistic operational conditions, in our simulations we define P_{adv} following the net-profit formulation of Sarenche et al. [15] to incorporate electricity expenditure alongside block rewards. Let c denote the unit operational cost rate per unit hash rate

Table 4. Buffer scaling law ($L = 256, \gamma = 1.0$): mean excluded honest orphans L_{OEA} as a function of buffer depth d , where $d = 0$ corresponds to $\text{DAM}_2^{\text{mod}}$.

α	$d = 0$ ($\text{DAM}_2^{\text{mod}}$)	$d = 8$	$d = 16$	$d = 32$	$d = 64$	$d = 96$	$d = 128$
0.30	1.51	0.14	0.02	0.00	0.00	0.00	0.00
0.35	2.55	0.58	0.22	0.05	0.00	0.00	0.00
0.40	7.03	2.84	1.72	0.67	0.04	0.04	0.00
0.45	28.35	19.98	14.94	11.89	6.98	2.89	0.92

at nominal difficulty, and let $\omega := \lambda K/c$ denote the network profitability factor ($\omega = 2.0$ in our benchmarks). Under honest mining, a participant with hash-rate share α incurs an operational expenditure rate $\alpha c = \alpha \lambda K/\omega$, yielding baseline net profit rate

$$\Pi_{\text{honest}} = \alpha \lambda K - \alpha c = \alpha \lambda K \left(1 - \frac{1}{\omega}\right). \quad (23)$$

For an adversary executing strategy π with time-averaged reward rate $R_{\mathcal{A}}(\pi)$ and operational cost rate $C_{\mathcal{A}}(\pi)$, its net profit rate is $\Pi_{\mathcal{A}}(\pi) = R_{\mathcal{A}}(\pi) - C_{\mathcal{A}}(\pi)$, and the normalized profit advantage is:

$$P_{\text{adv}}(\pi) = \frac{\Pi_{\mathcal{A}}(\pi) - \Pi_{\text{honest}}}{\lambda K} = \frac{R_{\mathcal{A}}(\pi) - C_{\mathcal{A}}(\pi) - \alpha \lambda K \left(1 - \frac{1}{\omega}\right)}{\lambda K}. \quad (24)$$

When the adversary mines continuously at full capacity without idling, its operational cost rate matches that of honest mining ($C_{\mathcal{A}}(\pi) = \alpha c$), causing operational costs to cancel identically so that Eq. (24) coincides with the normalized reward advantage defined in Eq. (1).

C.2 Parameter Sensitivity: The Buffer Scaling Law

To empirically validate the exponential suppression predicted by Lemma 1 and Theorem 1, we sweep buffer depth $d \in \{0, 8, 16, 32, 64, 96, 128\}$ under $L = 256$ across 15 independent replicas, where $d = 0$ corresponds to $\text{DAM}_2^{\text{mod}}$ with no settlement buffer. Table 4 reports the empirical results.

The empirical results exhibit strict monotonic suppression of L_{OEA} as buffer depth d increases, closely matching the theoretical tail bound in Lemma 1. For moderate hash rates ($\alpha \leq 0.35$), exclusions vanish completely at modest buffer depths ($d \in [32, 64]$). As $\alpha \rightarrow 1/2$, the random-walk drift $1 - 2\alpha$ becomes shallow, producing an extended excursion tail where $d = 128$ is required to suppress exclusions below 1.0 block and eliminate profit advantages. These empirical findings provide concrete calibration guidelines: setting $d \in [16, 32]$ provides robust exclusion resistance against typical mining pools with $\alpha \leq 1/3$, while $d \geq 64$ secures networks operating under higher adversarial thresholds up to $\alpha = 0.45$.

C.3 Dynamic Hash Rate Shock: Liveness and Responsiveness

A primary operational requirement of any difficulty adjustment algorithm is preserving chain throughput and responsiveness during abrupt network hash rate

Table 5. Dynamic hash rate shock (50% contraction at Epoch 10, $L = 128$, $\alpha = 0$, 30 replicas): difficulty trajectory, block intervals, and recovery time.

Protocol	D_9 (Pre)	D_{10} (Crash)	D_{11} (Post)	I_{10} (s)	I_{11} (s)	Recovery (Ep)	Liveness Assessment
Bitcoin DAM	1.0301	0.5260	0.5003	1.97	1.06	1.3 ± 0.5	Immediate single-epoch retarget.
$\text{DAM}_1^{\text{mod}}$	1.0301	0.5260	0.5003	1.97	1.06	1.3 ± 0.5	Immediate adaptation (honest baseline).
$\text{DAM}_2^{\text{mod}}$	1.0301	0.5260	0.5003	1.97	1.06	1.3 ± 0.5	Immediate adaptation (honest baseline).
Zeno ($\tau = 0.5$)	1.0633	0.7947	0.6475	2.03	1.60	4.5 ± 1.1	Sluggish Drag: 4–5 epochs deceleration.
PB-DAM ($d = 32$)	1.0176	0.5947	0.5079	1.72	1.39	2.2 ± 0.8	Rapid Adaptation: Smooth 1-window transition.

contractions, such as geopolitical mining shutdowns or power outages. To isolate the intrinsic adaptive responsiveness of each mechanism without confounding adversarial strategies, we evaluate an honest network ($\alpha = 0$) subjected to an instantaneous 50% drop in total hash rate at block height 1280, coinciding with the start of Epoch 10 under $L = 128$. Table 5 presents the pre-crash difficulty D_9 , post-crash retargeted difficulty D_{10} , subsequent adaptation difficulty D_{11} , the mean block arrival intervals during and after the shock (I_{10} and I_{11}), and the mean epochs required to recover difficulty to within 10% of nominal parity (0.50), averaged over 30 independent replicas.

Mechanisms Governing Shock Adaptation. Table 5 demonstrates the divergence between difficulty adjustment paradigms under sudden network contractions:

- *Proportional Retargeting (Bitcoin and Modified DAMs):* In an honest network, natural orphans are negligible under synchronous delivery. With zero uncle reports, $\text{DAM}_1^{\text{mod}}$ and $\text{DAM}_2^{\text{mod}}$ collapse to Bitcoin’s proportional update, halving difficulty in a single epoch (1.3 ± 0.5 epochs to nominal parity). This establishes that orphan-aware tracking incurs zero intrinsic adaptation penalty under honest conditions.
- *Damped Retargeting (Zeno DAA):* By restricting downward adjustments via convex damping ($D_{k+1} = (1 - \tau)\tilde{D} + \tau D_k$), Zeno deliberately penalizes difficulty reductions to inflate adversarial profit lags. Under a 50% crash, this damping leaves difficulty heavily inflated across multiple epochs, requiring 4.5 ± 1.1 epochs to recover parity and subjecting the network to severe throughput deceleration ($I_{11} = 1.60$ s).
- *Pipelined Retargeting (PB-DAM):* Rather than dampening step sizes, PB-DAM computes difficulty directly from observed active work rate W_k/T_k . Its only adjustment latency is the d -block settlement buffer: evaluation window $\mathcal{W}_{10}$ spans d pre-crash buffer blocks and $L - d$ post-crash prefix blocks. PB-DAM adapts smoothly within 2.2 ± 0.8 epochs, bounding latency to a single window without the multi-epoch liveness drag of damping schemes.