International Journal of Wireless and Ad Hoc Communication

♦ ISSN: 2692-4056 Vol: 11, No. 01, 2027 ♦ PP. 34–45

Article DOI: https://doi.org/10.54216/IJWAC.110104



FULL LENGTH ARTICLE

Freshness–Transmission Trade-offs in Wireless Ad Hoc Communication: An Analysis of Age-Threshold Random Access

Wadhah Ahmed Muthanna Abdullah[1,*]

[1] Saint Petersburg State University, Saint Petersburg, Russia

Email: st.082532@student.spbu.ru

Received: June 20, 2026 Revised: August 29, 2026 Accepted: November 02, 2026 ⋆ Corresponding author

ABSTRACT

Wireless ad hoc monitoring requires access rules that keep received information fresh without excessive contention. Age-threshold random access suppresses recently successful transmitters, but its performance depends on the transmission probability used after the silent interval. This study examines that dependence through renewal analysis and 1,280 reproducible simulation runs covering network size, independent erasures, correlated outages, and unequal link quality. Twenty independent replications are used per configuration. Exact slotted ALOHA and round-robin expressions provide implementation checks; a decoupling approximation links age, throughput, and transmission attempts for threshold access. At 50 peers and an independent erasure probability of 0.2, a silent interval of 50 slots with access probability 0.04 lowers mean age from 168.59 to 140.02 slots, while increasing attempts from 1.00 to 1.45 per slot. A longer silent interval of 100 slots combined with probability 0.08 raises mean age to 398.94 slots. Reducing that probability to 0.04 instead produces 101.46 slots and 0.84 attempts per slot. The results locate a collision-driven reversal that disappears from comparisons reporting only a preferred configuration. Correlated outages and link heterogeneity further show why mean age, deadline violations, and fairness require separate evaluation. The findings support joint threshold–probability selection within an explicit transmission budget. They apply to a synchronized, single collision domain with ideal acknowledgments; spatial reuse and multi-hop forwarding require further analysis.

Keywords: Wireless ad hoc communication ▪ Age of information ▪ Slotted ALOHA ▪ Random access ▪ Packet erasures ▪ Transmission budget

1. INTRODUCTION

A wireless peer that has just delivered a status update has little immediate reason to compete with a peer whose last delivery occurred hundreds of slots earlier. Equal access probabilities ignore that difference: in a decentralized monitoring network, the resulting collisions can leave the receiver with stale information even when many transmissions are attempted, which makes the age of information (AoI) a useful complement to throughput. AoI measures the time since generation of the most recently received update. Surveys by Abbas et al. [1] and Kahraman et al. [2] describe its relationship with sampling, buffering, and medium access across Internet of Things applications.

Age-threshold access introduces a simple local restriction: after success, a transmitter remains silent for a prescribed interval, then enters probabilistic contention. The rule can be implemented with a counter and an acknowledgment. Its simplicity is attractive for ad hoc communication, where peers may share a channel without a scheduler allocating every transmission, but the silent interval and the subsequent access probability act together. Increasing the threshold can

© 2027 The Author(s). Published by ASPG. This article is licensed under CC BY 4.0.

reduce unnecessary contention; increasing the probability can undo that reduction once enough peers become eligible. Both settings matter.

Recent analyses have established the value of age-sensitive access under several network models. Zhao et al. [3] study threshold access with spatial interference, while Moradian et al. [4] adjust frame length and age-gain thresholds in machine-type communication. Packet errors change the preferred threshold–probability pair [5]. These findings motivate a narrower engineering question: under a fixed and fully specified collision model, how much of an apparent freshness gain survives when transmission attempts, age violations, and unequal service are reported alongside the mean? The answer matters whenever the same access rule will be used under more than one loss condition.

A second question concerns temporal structure. Two channels can have the same long-run erasure probability and still differ in the sequence of opportunities they offer. If an outage lasts long enough for recently successful peers to become eligible, its end can release a larger contending population; the consequence depends on both the access probability and the current age distribution. An independent-erasure experiment cannot capture that interaction. We therefore compare independent losses with a stationary two-state outage process while holding the marginal loss probability at 0.2.

This paper provides a controlled computational analysis of an established protocol family. It does not introduce threshold access as a new medium-access mechanism. The contribution is a reproducible account of its freshness–transmission trade-off, including an unfavorable aggressive configuration, a finite-grid resource comparison, a check of the decoupling approximation, and separate stress tests for burstiness and link heterogeneity. The central outcome is conditional: a threshold can improve freshness and reduce attempts when paired with a suitable probability, yet a larger probability can drive the same network toward persistent contention. All numerical results are generated by the stated simulator, and the accompanying source code and aggregate replication records permit the calculations to be repeated.

2. RELATED WORK

2.1 Age analysis and access control

Exact age calculations provide the natural starting point for a simulation study. Fiems and Vinel [6] derive mean-age results for slotted ALOHA, and Fiems and Vinel [7] extend the analysis to distributional properties. Their work makes the distinction between a mean and an age tail explicit. Our generate-at-will model admits a particularly simple geometric baseline, which we use as an implementation check before interpreting threshold results.

Frame-based access offers another way to organize contention. Wang et al. [8] analyze age under frame slotted Aloha, and Yue et al. [9] study a frame-based status-updating protocol. Both address structured access opportunities; the present experiment holds slot duration and packet length fixed and removes frame adaptation so that the effects of threshold and probability remain identifiable. Moradian et al. [4] go further by adjusting access decisions using estimated age-gain distributions. That approach supplies a relevant design direction, but a faithful implementation would require the estimator and frame-control logic to be reproduced, which falls outside this comparison.

Backoff also changes the spacing between successful deliveries. Buratto et al. [10] compare reactive and proactive strategies through inter-delivery statistics. Their distinction is relevant because the same average delivery rate can be associated with different delivery variability. Here, the silent interval follows a successful reception and the terminal generates a fresh packet whenever it attempts access; there is no retransmission queue retaining an old payload. This choice isolates the renewal timing mechanism.

2.2 Channel errors, spatial interference, and energy

The spatial model in Zhao et al. [3] permits interference effects that a single collision domain cannot reproduce. The present abstraction instead assumes that simultaneous attempts prevent all receptions, making it suitable for a deliberately restricted comparison of local contention rules. Results for density in a spatial network cannot be equated directly with our results for the number of peers in one shared domain. The latter has no spatial reuse.

Tzeng et al. [5] explicitly include packet errors in threshold selection. We adopt the same general premise that access settings should depend on reliability, while testing transparent fixed settings and a small parameter grid. Energy harvesting introduces an additional state variable because access also depends on available stored energy [11]. Our transmission count is a resource proxy: multiplying it by an assumed per-attempt energy gives transmit energy, but it excludes sensing, reception, acknowledgment processing, and radio wake-up costs. Battery lifetime remains unmeasured.

Physical-layer reception capability can change the contention penalty. Fan et al. [12] study NOMA-assisted ALOHA throughput, while Zhu et al. [13] develop age-gain-dependent random access for event-driven periodic updating. NOMA can make some simultaneous transmissions useful. The collision model used here assigns them zero useful throughput. This is a substantive modeling boundary, particularly for the aggressive setting whose failure is driven by concurrent attempts.

2.3 Position of the present analysis

The relevant literature already supplies threshold protocols, frame adaptation, backoff analysis, and richer reception models. Our study asks what can be learned from a smaller model whose assumptions and every reported number can be inspected. The comparisons use fixed rules, including a poor rule, rather than selecting a winner from a large unreported tuning process. A separate grid then examines how that poor result changes when the access probability is reduced.

The evidence is organized around three questions. How does threshold access alter the mean and tail of age when the loss probability increases? Which tested configurations remain competitive after attempts per slot are counted? How do temporal correlation and link heterogeneity change those conclusions? These questions concern mechanism and parameter sensitivity. They do not establish superiority over current adaptive protocols.

3. SYSTEM MODEL AND MATHEMATICAL ANALYSIS

3.1 Wireless contention domain and age convention

Consider N source–destination links sharing a synchronized wireless collision domain. Each source generates a new status packet at the beginning of any slot in which it transmits. A packet occupies one slot. Every destination receives its intended packet only if that source is the sole transmitter and the channel does not erase the packet. A successful reception is acknowledged without delay or error before the next access decision.

Let $A_i(t)$ denote the age at the end of slot t. With $Y_i(t) = 1$ for a successful delivery in that slot, the state evolves as

$$A_i(t) = \begin{cases} 1, & Y_i(t) = 1, \\ A_i(t-1)+1, & Y_i(t) = 0. \end{cases} \tag{1}$$

A successfully delivered packet has age one because transmission occupies one complete slot. This convention fixes the additive constants in the renewal expressions below. All ages are measured in slots; conversion to seconds requires multiplication by a specified slot duration.

Each source uses the local age from the previous slot. For threshold $T \geq 0$ and eligible-node access probability p, its attempt indicator is

$$\begin{aligned} X_i(t) &= \mathbf{1}\{A_i(t-1) > T\}\mathbf{1}\{U_i(t) < p\}, \\ U_i(t) &\sim \mathrm{Unif}(0,1). \end{aligned} \tag{2}$$

The uniforms are independent across peers and slots. The strict inequality in Equation (2) suppresses exactly T slots after success. Ordinary slotted ALOHA (SA) sets $T = 0$. We use $p = c/N$, where c is a dimensionless access multiplier. Table 1 records the notation used throughout.

3.2 Exact independent-access baseline

Under SA, a particular peer succeeds with probability

$$q_0 = p(1-p)^{N-1}(1-\varepsilon). \tag{3}$$

Independent loss across slots makes the peer's inter-delivery interval geometric on $\{1,2,\ldots\}$. The exact stationary mean age and age-violation probability are

$$\bar{A}_{\mathrm{SA}} = \frac{1}{q_0}, \qquad V_D = (1-q_0)^D. \tag{4}$$

The second expression counts ages strictly greater than the integer deadline D. For the same model, throughput is $S = Nq_0$, attempts are $G = Np$, idle probability is $I = (1-p)^N$, and collision probability is

$$C = 1-(1-p)^N - Np(1-p)^{N-1}. \tag{5}$$

A singleton erasure belongs to neither the idle nor the collision category. This distinction is used in the simulation counters.

For fixed N and $\varepsilon < 1$, differentiating $\log q_0$ with respect to p gives $1/p-(N-1)/(1-p)$. Its zero occurs at $p = 1/N$, which minimizes mean age within the fixed-probability SA family. The independent erasure probability multiplies the success rate and therefore changes the achieved age without moving this optimum. At large N, $\bar{A}_{\mathrm{SA}}/N$ approaches $e/(1-\varepsilon)$ for that setting.

Table 1. Model notation and units.

Symbol	Unit	Definition
N	peers	Number of contending source–destination links
$A_i(t)$	slots	Receiver age at the end of slot t
T	slots	Number of silent slots after a successful update
p, c	dimensionless	Eligible-node attempt probability and multiplier $c = Np$
ε, ρ	dimensionless	Marginal erasure probability and outage persistence
G, S	per slot	Network attempts and successful updates
C, I	probability	Fraction of collision slots and idle slots
V_D	probability	Fraction of peer–slot ages greater than D
J	dimensionless	Jain index calculated from inverse per-peer mean ages
q, τ	probability	Eligible-node success probability and stationary attempt probability
L, m	slots	Inter-delivery interval and its expectation

3.3 Renewal representation of threshold access

A threshold cycle consists of a deterministic silent interval and a random eligible period. To obtain a tractable approximation, assume that an eligible peer faces a constant success probability q and that other peers attempt independently with stationary probability τ. Then

$$L = T + Z, \qquad Z \sim \mathrm{Geom}(q), \tag{6}$$

$$m = \mathbb{E}[L] = T + q^{-1}, \tag{7}$$

$$\mathbb{E}[L^2] = T^2 + 2Tq^{-1} + (2-q)q^{-2}. \tag{8}$$

A cycle of length L accumulates ages $1,2,\ldots,L$. Renewal reward theory consequently gives

$$\bar{A} \simeq \frac{\mathbb{E}[L^2]+\mathbb{E}[L]}{2\mathbb{E}[L]} = \frac{m+1}{2} + \frac{1-q}{2q^2 m}. \tag{9}$$

The second term records variability in the eligible period. It explains why improving mean age need not require increasing throughput: a narrower distribution of intervals can reduce the accumulated age even if the mean interval becomes longer.

The eligible fraction is $1/(1+Tq)$, so the decoupling equations are

$$\tau = \frac{p}{1+Tq}, \qquad q = p(1-\tau)^{N-1}(1-\varepsilon). \tag{10}$$

We solve the scalar equation numerically, using a bracket scan over $0 < q \leq p$ and a root solver. The implementation records

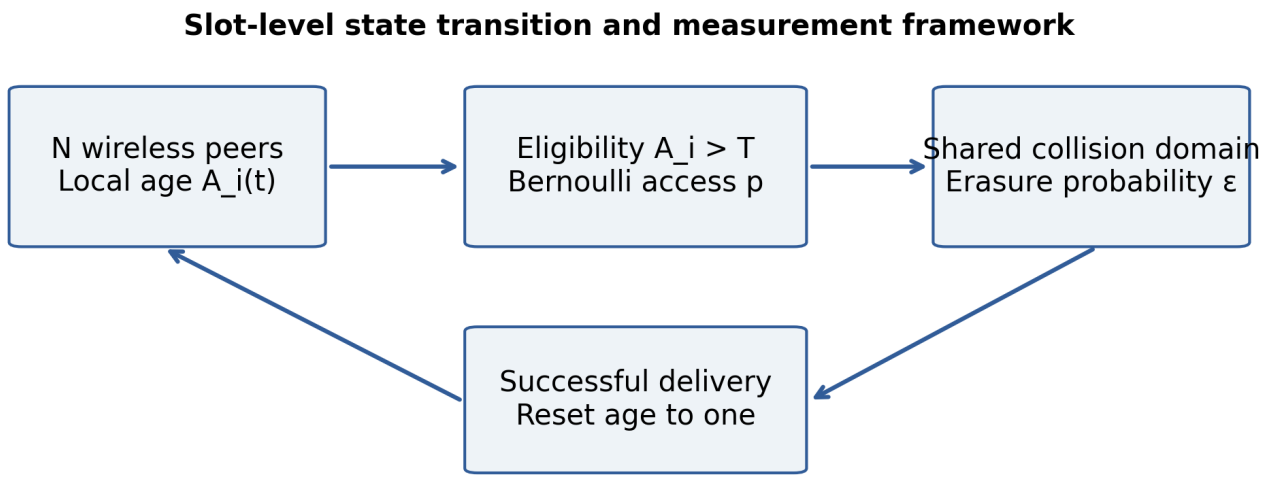


Figure 1. State and measurement framework for age-threshold access. A failed or suppressed transmission leaves the receiver age increasing by one slot. Acknowledgments are ideal.

the largest root if more than one is located. Root selection would require an additional stability analysis in a model with multiple admissible branches. For the twelve reported grid points, the numerical scan found one root each.

Equation (10) is a decoupling approximation for threshold access. Eligibility histories share the same collision process, so their independence is not exact. The simulation retains that dependence. Agreement between the two is assessed numerically, and the approximation is not used to generate the reported simulation means.

3.4 Tail probability and transmission resource

The expected number of ages above D in a renewal cycle is $\mathbb{E}[(L-D)^+]$. With the same geometric approximation,

$$V_D \simeq \begin{cases} 1-D/m, & D < T, \\ (1-q)^{D-T}/(qm), & D \geq T. \end{cases} \tag{11}$$

A threshold close to the deadline leaves fewer opportunities to recover from an unsuccessful eligible period. Raising T is therefore not a direct prescription for reducing the tail, even when it removes many early attempts. Both q and the remaining recovery window change.

Network transmission load and throughput satisfy

$$G \simeq \frac{Np}{1+Tq}, \qquad S \simeq \frac{Nq}{1+Tq}, \qquad \frac{G}{S} \simeq \frac{p}{q}. \tag{12}$$

The ratio G/S is attempts per delivered update. If each attempt consumes E_{tx} joules, the transmit-energy component per slot is $E_{\mathrm{tx}}G$. A complete device model would add idle listening, sampling, and switching terms, which are not measured here.

For a finite tested set $\mathscr{P}$, a resource-constrained comparison can be written as

$$\begin{aligned} \min_{(T,p)\in\mathscr{P}} \quad & \bar{A}(T,p) \\ \text{subject to} \quad & G(T,p) \leq G_{\max}, \\ & V_D(T,p) \leq v_{\max}. \end{aligned} \tag{13}$$

This expression defines a selection criterion. The experiment evaluates its feasible points on a small grid; it does not solve a continuous global optimization problem.

3.5 Scheduled reference and heterogeneous service

A collision-free round-robin (RR) schedule assigns one peer to each slot, repeating every N slots. Under independent erasures, a peer succeeds once every NK slots, with $K \sim \mathrm{Geom}(1-\varepsilon)$. Substitution into the renewal formula yields

$$\bar{A}_{\mathrm{RR}} = \frac{N(1+\varepsilon)}{2(1-\varepsilon)} + \frac{1}{2}. \tag{14}$$

RR has one attempt per slot and requires a known population, synchronized ordering, and a maintained schedule. It is an idealized scheduled comparator, not an implementation-equivalent decentralized competitor or a universal lower bound.

For unequal link quality, we calculate each peer's mean age $\bar{A}_i$ and define $f_i = 1/\bar{A}_i$. Fairness is measured by

$$J = \frac{(\sum_{i=1}^{N} f_i)^2}{N\sum_{i=1}^{N} f_i^2}. \tag{15}$$

A high J indicates similar inverse mean ages across peers. Low age is a separate requirement. A network in which every peer has equally stale information can attain J close to one.

4. METHODOLOGY

4.1 Experimental design

The core experiment crosses $N \in \{20,50,100\}$ with $\varepsilon \in \{0,0.2,0.4\}$ and four rules. SA uses $(T,c) = (0,1)$; AT1 uses $(N,2)$; AT2 uses $(2N,4)$; RR follows the scheduled reference. AT1 and AT2 are labels for fixed parameter settings of the same age-threshold rule. AT2 is retained even where it performs poorly. Table 2 gives the full experiment.

All simulations start with age one for every peer. The first 10,000 slots are discarded from measurement, while the age state continues into the observed interval. Each ordinary run records 80,000 slots. A longer check increases that interval to 320,000 slots for the four core rules at $N = 50$ and $\varepsilon = 0.2$. No physical measurement dataset is used: the numerical records are generated observations from the specified

Table 2. Simulation design. Every configuration has 20 independent replications.

Experiment	Settings and run length
Core	$N = 20, 50, 100$; $\varepsilon = 0, 0.2, 0.4$; 36 configurations; 10,000-slot burn-in; 80,000 measured slots
Parameter grid	$N = 50$; $\varepsilon = 0.2$; 12 configurations; 10,000-slot burn-in; 80,000 measured slots
Correlated outage	$N = 50$; $\varepsilon = 0.2$; 8 configurations; 10,000-slot burn-in; 80,000 measured slots
Unequal links	$N = 50$; $\varepsilon = 0.05/0.35$; 4 configurations; 10,000-slot burn-in; 80,000 measured slots
Long-run check	$N = 50$; $\varepsilon = 0.2$; 4 configurations; 10,000-slot burn-in; 320,000 measured slots
Total	64 configurations; 1,280 runs

The grid uses $T \in \{0, 50, 100\}$ and $c \in \{0.5, 1, 2, 4\}$. Correlated outages use $\rho \in \{0.8, 0.95\}$. Unequal links divide the peers into equal groups. Some grid points duplicate core settings intentionally.

computational model. Literature available by 31 May 2026 defines the research context; the simulation does not claim a historical collection date.

4.2 Correlated outage and unequal-link experiments

Let $B_t = 1$ denote a bad slot that erases every potential singleton delivery. A two-state chain is specified by

$$\begin{aligned} \Pr(B_t = 1 \mid B_{t-1} = 0) &= \varepsilon(1-\rho), \\ \Pr(B_t = 1 \mid B_{t-1} = 1) &= \varepsilon + \rho(1-\varepsilon). \end{aligned} \tag{16}$$

Its stationary bad-state probability is ε and its lag-one correlation is ρ. Initialization draws B_0 from that stationary distribution. At $\rho = 0$, the model reduces to independent slot erasures. Since there can be at most one successful candidate in a slot, a common independent slot erasure and a link-specific independent erasure have the same reception distribution in the homogeneous experiment.

The mean duration of a bad run is

$$\mathbb{E}[L_B] = \frac{1}{(1-\varepsilon)(1-\rho)}. \tag{17}$$

At $\varepsilon = 0.2$, it equals 1.25, 6.25, and 25 slots for $\rho = 0$, 0.8, and 0.95. These are explicit outage scenarios. They are not fitted to a propagation trace.

The heterogeneous experiment gives half the peers erasure probability 0.05 and half probability 0.35, independently across slots. Their average link erasure probability remains 0.2. The access rule is identical for both groups, allowing differences in mean age to arise from reliability and the resulting eligibility histories. This experiment does not combine heterogeneous links with correlated outages; each stressor is changed separately to preserve interpretability.

4.3 Measurement and uncertainty

For replication r, measured length K, and deadline $D = 5N$, the primary estimators are

$$\widehat{A}_r = \frac{1}{NK} \sum_{t=1}^{K} \sum_{i=1}^{N} A_i(t), \tag{18}$$

$$\widehat{V}_{D,r} = \frac{1}{NK} \sum_{t=1}^{K} \sum_{i=1}^{N} \mathbf{1}\{A_i(t) > D\}, \tag{19}$$

$$\widehat{G}_r = \frac{1}{K} \sum_{t=1}^{K} \sum_{i=1}^{N} X_i(t). \tag{20}$$

Throughput counts successful slots divided by K. Collision and idle fractions count slots with at least two attempts and with zero attempts, respectively. Group ages are averages of per-peer means, and J is computed within each replication before averaging across replications.

The reported estimate is the mean of 20 replication estimates. Its 95% Monte Carlo confidence interval uses

$$\widehat{\mu} \pm t_{0.975,19} \frac{s_r}{\sqrt{20}}, \qquad t_{0.975,19} \approx 2.093. \tag{21}$$

The replication is the independent unit. Because ages are strongly serially dependent and a delivery changes several subsequent age observations, treating all peer–slot observations as independent would produce an apparent sample size far larger than the number of independently generated runs. Intervals describe simulation variability conditional on the model, not uncertainty about real wireless deployments. They are pointwise intervals and do not provide simultaneous coverage for every comparison in the parameter grid.

4.4 Implementation and verification

The simulator is written in C++17. A specified SplitMix64 generator produces uniforms from the upper 53 bits of each output, with seed $20260531 + 1000003r + 1009N$ for replication index $r = 0, \ldots, 19$. Every slot consumes N access uniforms followed by one channel uniform, even when a threshold suppresses an attempt. This fixed draw order supports matched random streams between policies; replications have separate seeds. The analysis uses marginal replication intervals and makes no unreported significance claim from this pairing.

The update order is access decision, channel state, delivery decision, age update, and measurement. A successful packet resets age to one in the same measured slot. Assertions in the verification script check row counts, parameter ranges, nonnegative rates, success bounded by one, and the relation $I + C + S \leq 1$, whose remainder consists of erased singleton attempts. Exact SA and RR means provide an independent numerical target. The complete code, CSV summaries, and figure-generation script accompany the manuscript.

5. RESULTS

5.1 Agreement with exact baselines

Table 3 compares simulated and exact ages for all 18 baseline cases. The largest absolute relative discrepancy is 0.408%. RR without loss gives exactly $(N+1)/2$ slots, as expected from a deterministic cycle. SA converges to the geometric

expression despite collisions among peers because the fixed access decisions are independent from one slot to the next.

Table 3. Exact and simulated mean ages for independent erasures.

N	ε	Rule	Exact / simulation / CI half-width / error (%)
20	0.0	SA	53.001 / 53.025 / 0.129 / 0.045%
20	0.0	RR	10.500 / 10.500 / 0.000 / 0.000%
20	0.2	SA	66.251 / 66.224 / 0.226 / −0.041%
20	0.2	RR	15.500 / 15.486 / 0.017 / −0.088%
20	0.4	SA	88.334 / 88.217 / 0.359 / −0.133%
20	0.4	RR	23.833 / 23.841 / 0.056 / 0.032%
50	0.0	SA	134.553 / 134.549 / 0.339 / −0.003%
50	0.0	RR	25.500 / 25.500 / 0.000 / 0.000%
50	0.2	SA	168.191 / 168.588 / 0.476 / 0.236%
50	0.2	RR	38.000 / 38.015 / 0.065 / 0.038%
50	0.4	SA	224.254 / 225.170 / 0.934 / 0.408%
50	0.4	RR	58.833 / 58.874 / 0.110 / 0.070%
100	0.0	SA	270.468 / 270.385 / 1.002 / −0.031%
100	0.0	RR	50.500 / 50.500 / 0.000 / 0.000%
100	0.2	SA	338.085 / 337.930 / 1.454 / −0.046%
100	0.2	RR	75.500 / 75.495 / 0.144 / −0.007%
100	0.4	SA	450.780 / 452.085 / 2.524 / 0.290%
100	0.4	RR	117.167 / 117.060 / 0.278 / −0.091%

CI denotes the half-width of a 95% Monte Carlo confidence interval. Relative error is $100(\widehat{A} - A_{\text{exact}})/A_{\text{exact}}$. Ages are in slots.

These checks test the age convention, slot accounting, erasure handling, and measurement accumulation. They do not verify the realism of the collision model. A simulator can agree with an exact formula for its own abstraction while omitting a mechanism that matters in an actual network, such as capture or delayed acknowledgments. The approximation comparison later in this section addresses a different question: how much error is introduced by decoupling threshold-dependent peer states.

5.2 Loss sensitivity and the reversal under aggressive access

At $N = 50$, AT1 reduces mean age relative to SA under each tested independent loss level, but the size of the reduction changes markedly. With no erasures, mean age falls from 134.55 to 98.72 slots, a reduction of 26.63%. At $\varepsilon = 0.2$, it falls from 168.59 to 140.02 slots, or 16.94%. At $\varepsilon = 0.4$, the reduction is only 5.05%. The attempt load simultaneously increases because fewer successful deliveries return peers to the silent state.

AT2 behaves differently. Its mean age is 100.10 slots without erasures, close to AT1, but rises to 398.94 slots at loss probability 0.2 and 760.67 slots at 0.4. The respective attempt loads are 1.43, 3.15, and 3.53 per slot. The longer silent interval does not prevent congestion once the high eligible-node probability leaves many peers competing through repeated unsuccessful slots.

Figure 2 separates that deterioration from the mean alone. At loss probability 0.2, AT2 exceeds the 250-slot age deadline in 53.35% of peer–slot observations, compared with 22.58% for SA and 16.09% for AT1. Its throughput is only 0.106 updates per slot. That combination identifies a collision-dominated configuration: more attempts purchase fewer deliveries and leave a larger fraction of observations stale.

Table 4. Joint freshness and transmission results at $N = 50$.

ε	Rule	$\bar{A}$ / CI / V_{250} / G / S / J
0.0	SA	134.55 / 0.34 / 0.1549 / 1.000 / 0.372 / 0.9968
0.0	AT1	98.72 / 0.37 / 0.0684 / 1.281 / 0.359 / 0.9982
0.0	AT2	100.10 / 2.93 / 0.0625 / 1.428 / 0.321 / 0.9981
0.0	RR	25.50 / 0.00 / 0.0000 / 1.000 / 1.000 / 1.0000
0.2	SA	168.59 / 0.48 / 0.2258 / 1.000 / 0.297 / 0.9959
0.2	AT1	140.02 / 0.63 / 0.1609 / 1.452 / 0.274 / 0.9972
0.2	AT2	398.94 / 5.90 / 0.5335 / 3.151 / 0.106 / 0.9923
0.2	RR	38.01 / 0.07 / 0.0003 / 1.000 / 0.800 / 0.9995
0.4	SA	225.17 / 0.93 / 0.3285 / 1.000 / 0.223 / 0.9944
0.4	AT1	213.80 / 0.94 / 0.3083 / 1.612 / 0.194 / 0.9948
0.4	AT2	760.67 / 10.56 / 0.7240 / 3.532 / 0.059 / 0.9853
0.4	RR	58.87 / 0.11 / 0.0103 / 1.000 / 0.600 / 0.9989

The deadline is 250 slots. G is attempts per slot, S is successes per slot, and J is the fairness index. The listed CI applies to mean age.

AT1 also illustrates the distinction between age and throughput. At loss probability 0.2, its throughput of 0.274 is below the SA value of 0.297, yet its mean age is lower. Equation (9) permits this outcome because age depends on the second moment of delivery intervals as well as their mean. Recently successful peers are temporarily removed from contention, changing the allocation of access opportunities across the age distribution.

5.3 Network size and normalized age

Table 5 and Figure 3 hold $\varepsilon = 0.2$ while increasing N. Dividing age by N makes the capacity constraint visible. SA remains close to its large-population limit $e/0.8 \approx 3.398$, whereas the threshold settings have different normalized trajectories because the combination of silent duration and finite-population collision probability changes their effective load.

Scaling the deadline with N asks whether service deteriorates relative to the population sharing the channel. It does not imply that an application can tolerate an arbitrarily longer deadline as the network grows. For a fixed physical deadline, population growth would demand a separate capacity decision, such as channel partitioning or a more capable receiver, since a favorable normalized-age curve allows the absolute age to grow and therefore cannot justify adding peers to an unchanged application.

The scheduled reference maintains a much smaller age because it avoids collisions and gives each peer a known opportunity every N slots. At $N = 50$ and $\varepsilon = 0.2$, its 38.01-slot simulated mean agrees with the exact 38-slot value. The result quantifies the cost of the tested uncoordinated access rules under ideal scheduling conditions; because the experiment assigns no maintenance overhead to RR, the observed difference covers collision avoidance and service timing but leaves the cost of establishing and maintaining that schedule unmeasured.

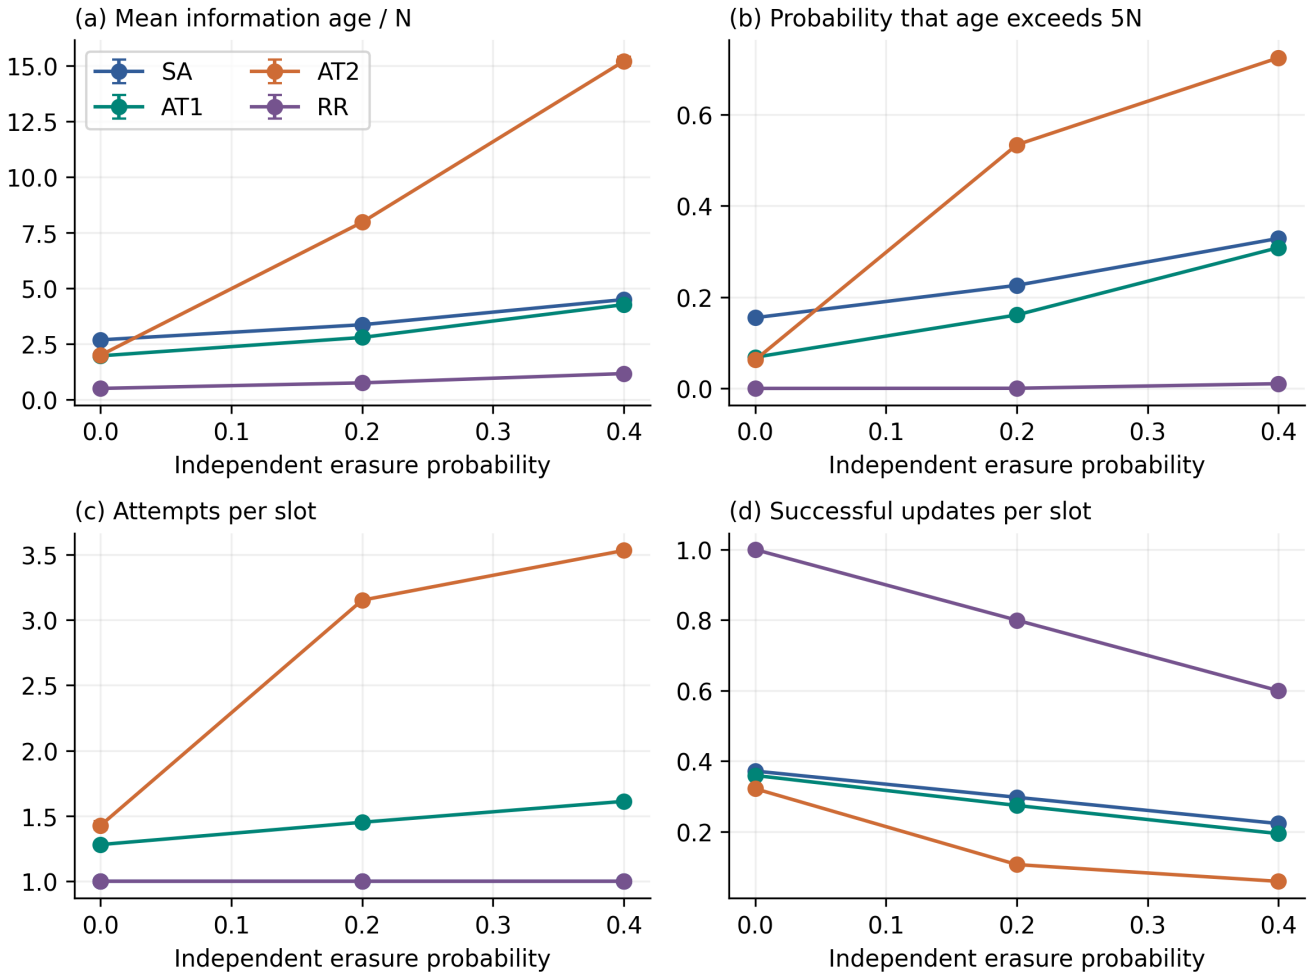


Figure 2. Loss sensitivity at $N = 50$: mean age, age violations, attempts, and throughput. Error bars are 95% Monte Carlo confidence intervals over 20 replications and may be smaller than the plotting symbols. AT1 uses $(T, c) = (50, 2)$; AT2 uses $(100, 4)$.

Table 5. Network-size sensitivity at independent erasure probability 0.2.

N	Rule	$\bar{A}$ / V_{5N} / G / C / I
20	SA	66.22 / 0.2182 / 1.000 / 0.264 / 0.359
20	AT1	55.42 / 0.1562 / 1.446 / 0.429 / 0.224
20	AT2	193.45 / 0.5830 / 3.222 / 0.838 / 0.041
20	RR	15.49 / 0.0003 / 1.000 / 0.000 / 0.000
50	SA	168.59 / 0.2258 / 1.000 / 0.264 / 0.364
50	AT1	140.02 / 0.1609 / 1.452 / 0.428 / 0.229
50	AT2	398.94 / 0.5335 / 3.151 / 0.825 / 0.042
50	RR	38.01 / 0.0003 / 1.000 / 0.000 / 0.000
100	SA	337.93 / 0.2273 / 0.999 / 0.264 / 0.366
100	AT1	280.85 / 0.1621 / 1.455 / 0.428 / 0.231
100	AT2	766.63 / 0.5228 / 3.140 / 0.823 / 0.042
100	RR	75.49 / 0.0003 / 1.000 / 0.000 / 0.000

The deadline scales as $D = 5N$. C is the collision-slot fraction and I is the idle-slot fraction. Erased singleton slots account for the remaining share after idle, collision, and successful slots are counted.

5.4 Threshold–probability grid and feasible resource choices

The parameter grid changes T and c separately at $N = 50$ and $\varepsilon = 0.2$. This is the key test of the aggressive configuration. With $T = 100$, reducing c from 4 to 2 lowers mean age from 398.94 to 101.46 slots and lowers attempts from 3.151 to 0.835 per slot. The corresponding changes are 74.57% and 73.50%. The threshold is unchanged. Access intensity accounts for the reversal.

Within the twelve tested settings, $(T, c) = (100, 2)$ has the lowest mean age. It also meets a budget of one attempt per slot. SA at $(0, 1)$ has mean age 168.59 and approximately one attempt per slot, so the grid winner improves mean age by 39.82% while using 16.49% fewer attempts. The search is finite. Smaller probability increments or different thresholds could produce other feasible points.

Figure 5 gives the resource view. The curve for $T = 100$ initially improves as c rises from 0.5 to 2, then turns sharply upward at 4. Attempts grow faster than useful service after the turn. For $T = 0$, increasing c from 1 to 4 raises mean age from 168.59 to 940.41 slots, consistent with the decline in singleton probability once the fixed-probability optimum is exceeded.

The grid also separates an age optimum from a low-load setting. At $(T, c) = (100, 1)$, attempts fall to 0.508 per slot and mean age is 128.38 slots. A stricter budget of 0.6 attempts

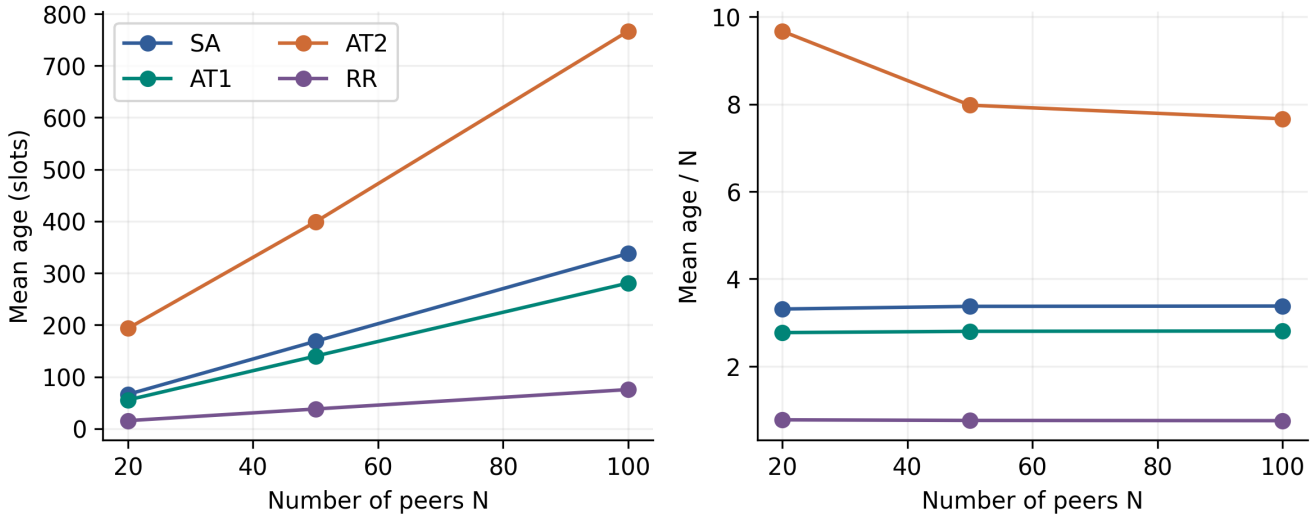


Figure 3. Absolute and normalized mean age under increasing peer count at $\varepsilon = 0.2$. The reference RR schedule retains one assigned transmission per slot.

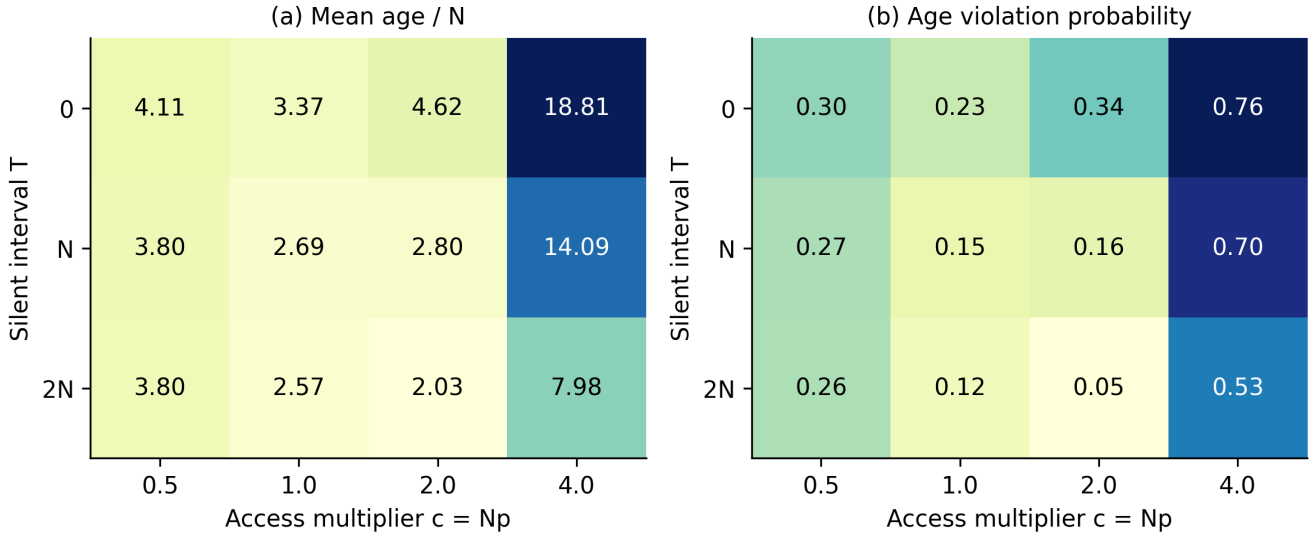


Figure 4. Joint parameter sensitivity at $N = 50$ and $\varepsilon = 0.2$. Every cell is a simulation mean. Panel (a) divides mean age by N; panel (b) reports the probability of age exceeding 250 slots. The color scale is independent in each panel.

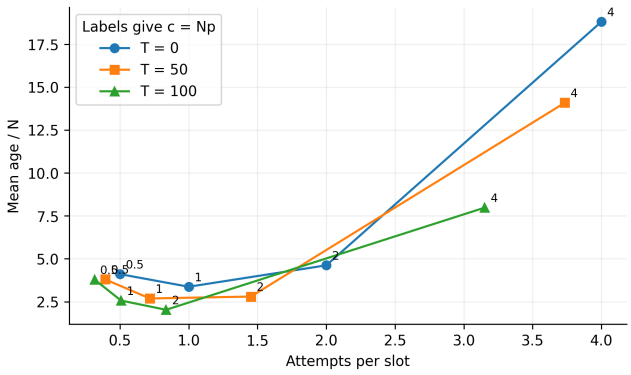


Figure 5. Freshness–transmission trade-off over the finite parameter grid. Each marker is a tested configuration; lines connect settings with the same silent interval for visual guidance. The lines do not represent evaluated intermediate configurations or a continuous Pareto frontier.

per slot would exclude the lowest-age point and admit this one. Equation (13) makes that choice explicit, avoiding a performance ranking that hides the available transmission resource.

5.5 Accuracy and limits of the decoupling approximation

The mean-age approximation tracks the parameter-grid pattern, including the high-age region. Across the twelve cases, its largest absolute relative error is 2.28%, occurring at $(T, c) = (100, 4)$. At $(100, 2)$, the error is -0.92%. Table 6 gives every discrepancy, and Figure 6 contrasts the absolute fit with relative error.

Equation (10) captures the homogeneous independent-loss results closely enough to explain their main pattern, although agreement over a finite set of parameters cannot establish the assumed independence of peers whose eligibility histories share the same collision process. In particular, the test does not evaluate approximation accuracy for correlated outages or unequal links, where conditioning on channel state or peer class would be necessary. Extending a successful homogeneous fit by simply substituting the same average loss

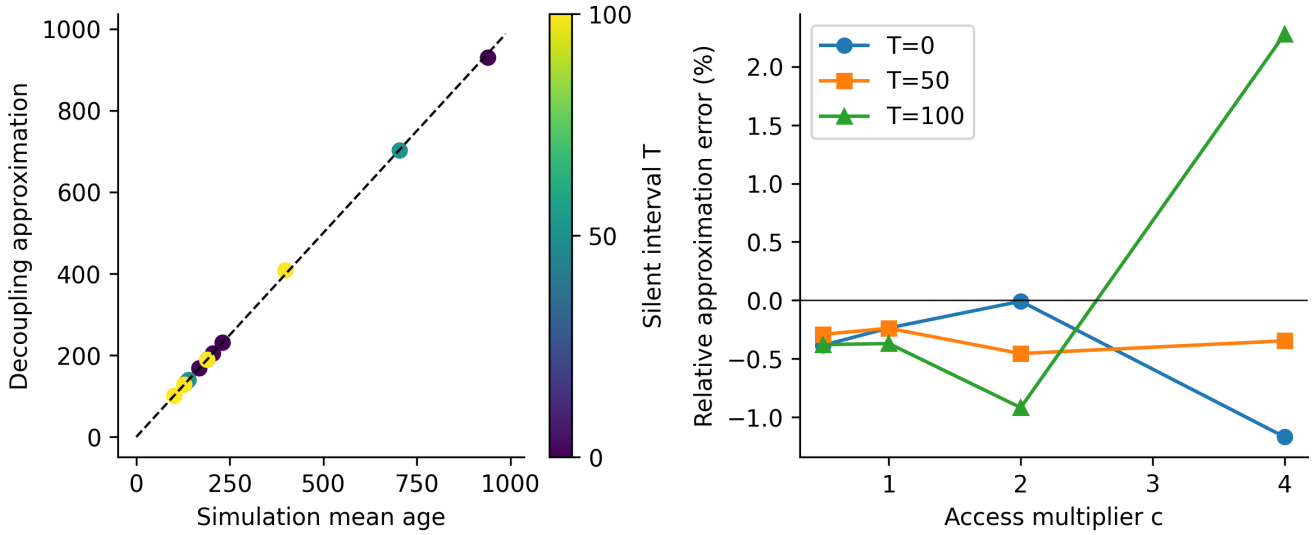


Figure 6. Decoupling accuracy for the twelve grid settings. The dashed diagonal is exact agreement; the relative-error panel uses the simulation mean as denominator. Colors in the left panel distinguish thresholds.

Table 6. Full parameter-grid results and decoupling approximation.

T	c	Simulation: $\bar{A}$ / V_{250} / G / S	Approximation: $\bar{A}$ / error (%)
0	0.5	205.33 / 0.2951 / 0.499 / 0.244	204.54 / −0.38%
0	1.0	168.59 / 0.2258 / 1.000 / 0.297	168.19 / −0.24%
0	2.0	230.99 / 0.3381 / 2.000 / 0.216	230.97 / −0.01%
0	4.0	940.41 / 0.7648 / 4.000 / 0.054	929.43 / −1.17%
50	0.5	190.04 / 0.2656 / 0.393 / 0.213	189.49 / −0.29%
50	1.0	134.41 / 0.1480 / 0.717 / 0.282	134.09 / −0.24%
50	2.0	140.02 / 0.1609 / 1.452 / 0.274	139.38 / −0.45%
50	4.0	704.38 / 0.7016 / 3.733 / 0.067	701.94 / −0.35%
100	0.5	189.78 / 0.2621 / 0.315 / 0.184	189.06 / −0.38%
100	1.0	128.38 / 0.1187 / 0.508 / 0.246	127.90 / −0.37%
100	2.0	101.46 / 0.0524 / 0.835 / 0.291	100.53 / −0.92%
100	4.0	398.94 / 0.5335 / 3.151 / 0.106	408.03 / 2.28%

Approx. is the mean age from Equations (8)–(10). Error is $100(A_{\text{approx}} - \widehat{A})/\widehat{A}$. The simulation means are independent of the approximation.

probability would discard the very dependence introduced in those stress tests.

5.6 Correlated outages at a fixed marginal loss rate

At the same marginal loss probability of 0.2, increasing ρ to 0.95 raises SA mean age from 168.59 to 172.89 slots. RR rises from 38.01 to 39.12 slots. AT1 changes little, from 140.02 to 140.41 slots, while AT2 decreases from 398.94 to 357.46 slots. The direction depends on the rule.

Table 7. Correlated-outage sensitivity at $N = 50$ and marginal loss probability 0.2.

ρ	Rule	$\bar{A}$ / CI / V_{250} / G / S
0.00	SA	168.59 / 0.48 / 0.2258 / 1.000 / 0.297
0.00	AT1	140.02 / 0.63 / 0.1609 / 1.452 / 0.274
0.00	AT2	398.94 / 5.90 / 0.5335 / 3.151 / 0.106
0.00	RR	38.01 / 0.07 / 0.0003 / 1.000 / 0.800
0.80	SA	169.25 / 0.72 / 0.2269 / 1.000 / 0.297
0.80	AT1	139.88 / 1.00 / 0.1608 / 1.451 / 0.274
0.80	AT2	389.13 / 7.37 / 0.5203 / 3.112 / 0.111
0.80	RR	38.10 / 0.19 / 0.0003 / 1.000 / 0.799
0.95	SA	172.89 / 1.20 / 0.2349 / 1.000 / 0.297
0.95	AT1	140.41 / 1.45 / 0.1628 / 1.447 / 0.276
0.95	AT2	357.46 / 10.15 / 0.4834 / 2.996 / 0.125
0.95	RR	39.12 / 0.46 / 0.0010 / 1.000 / 0.799

The $\rho = 0$ rows reuse the core independent-loss results. CI half-widths apply to mean age. Every peer experiences the same channel state in a slot; access uniforms remain independent.

The AT2 decrease is consistent with an interaction between outage timing and the evolving eligible population, but the aggregate counters cannot establish a unique causal pathway inside that interaction. Conditional age distributions at transitions from bad to good states would be needed to separate temporary synchronization from changes in the effective collision process. We retain the result without interpreting it as a general benefit of bursty channels.

The tail gives a further qualification. RR's violation probabil-

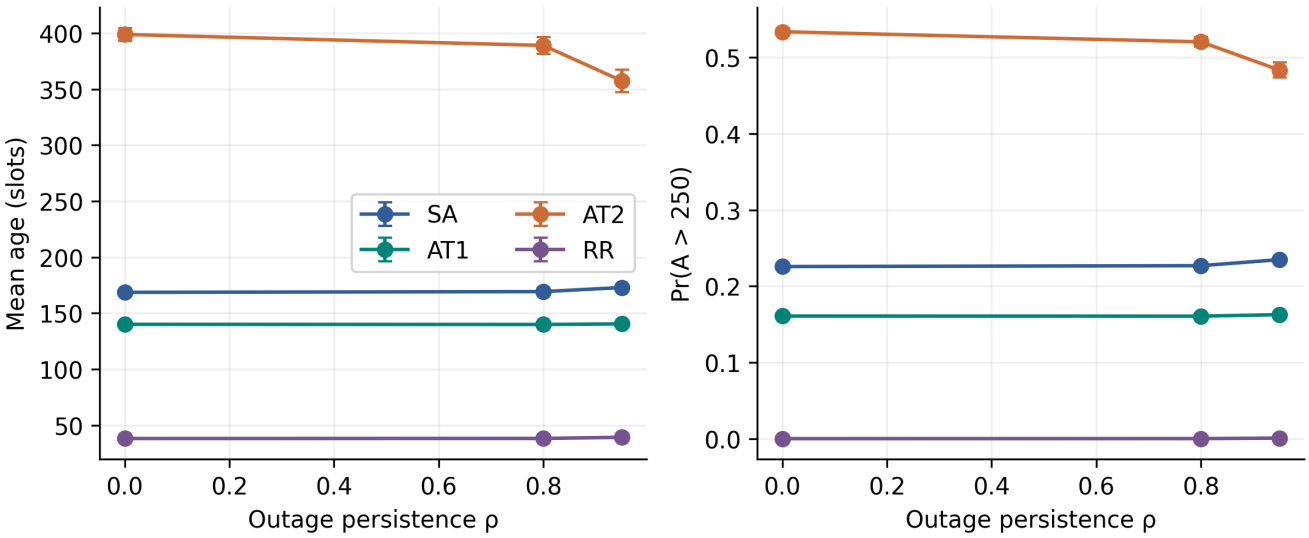


Figure 7. Effect of outage persistence on mean age and age violations. The marginal erasure probability is fixed at 0.2. The decline for AT2 is specific to the aggressive parameter setting and the common two-state outage model.

ity at 250 slots rises from 0.000317 to 0.000996 as ρ changes from zero to 0.95, although its mean remains much smaller than the random-access means. A low mean can coexist with a material relative increase in a rare-event probability. The absolute probabilities remain small in this particular experiment; applications with stricter deadlines should evaluate their own D directly.

5.7 Unequal links and fairness

Equal-size link classes with erasure probabilities 0.05 and 0.35 produce a mean link loss of 0.2, but different service outcomes. Under SA, the strong-link class has mean age 142.01 slots and the weak-link class 207.77. AT1 lowers those means to 120.56 and 170.52. AT2 leaves both classes substantially staler. Table 8 reports the group means together with aggregate age and fairness.

Table 8. Unequal-link results for 25 strong and 25 weak peers.

Rule	$\bar{A}_s$ / $\bar{A}_w$ / $\bar{A}$ / V_{250} / J / G
SA	142.01 / 207.77 / 174.89 / 0.2348 / 0.9622 / 1.000
AT1	120.56 / 170.52 / 145.54 / 0.1718 / 0.9688 / 1.458
AT2	349.56 / 490.27 / 419.91 / 0.5455 / 0.9650 / 3.177
RR	28.15 / 52.31 / 40.23 / 0.0026 / 0.9172 / 1.000

Subscripts s and w denote erasure probabilities 0.05 and 0.35, respectively. J is calculated within each replication from per-peer inverse mean ages, then averaged.

RR produces the lowest ages in both groups, 28.15 and 52.31 slots, yet its fairness index of 0.917 is lower than the random-access indices. This is not a contradiction. Equation (15) measures relative equality of inverse ages, while the left panel of Figure 8 measures their absolute level. Reporting J alone would obscure RR's better service to every group in this scenario.

AT1's fairness index is 0.969, compared with 0.962 for SA. Weak links remain eligible for longer after failures, giving them more chances to contend, while recently successful peers spend time silent. That mechanism is compatible with the smaller relative group gap. Its magnitude should not be generalized to unequal traffic rates or heterogeneous transmission power, neither of which is included in this experiment.

5.8 Longer observation windows

The longer-run check preserves the model and seeds while increasing measurement length by a factor of four. Table 9 reports mean age, its Monte Carlo interval, the age tail, and attempt load. The purpose is to determine whether the principal ordering depends on the chosen 80,000-slot observation interval.

Table 9. Observation-length check at $N = 50$ and $\varepsilon = 0.2$.

Rule	Slots	$\bar{A}$ / CI_A / V_{250} / CI_V / G
SA	80000	168.588 / 0.476 / 0.22581 / 0.00117 / 1.0001
AT1	80000	140.021 / 0.632 / 0.16093 / 0.00136 / 1.4522
AT2	80000	398.943 / 5.899 / 0.53352 / 0.00512 / 3.1512
RR	80000	38.015 / 0.065 / 0.00032 / 0.00004 / 1.0000
SA	320000	168.451 / 0.246 / 0.22556 / 0.00049 / 0.9996
AT1	320000	139.823 / 0.373 / 0.16065 / 0.00082 / 1.4517
AT2	320000	396.696 / 2.390 / 0.53128 / 0.00239 / 3.1448
RR	320000	38.006 / 0.022 / 0.00032 / 0.00002 / 1.0000

Each row summarizes 20 replications after a 10,000-slot burn-in. The longer runs extend the same seeded paths, so the two lengths are not independent samples. CIs are computed across replications at each length.

The poor performance of AT2 remains visible at the longer horizon, as do the lower mean age of AT1 relative to SA and the much smaller RR mean. A longer run narrows Monte Carlo uncertainty but cannot correct a missing physical mechanism. It is also insufficient to establish convergence for every possible threshold, because larger silent intervals or rarer successful receptions could require substantially longer transients and observation windows.

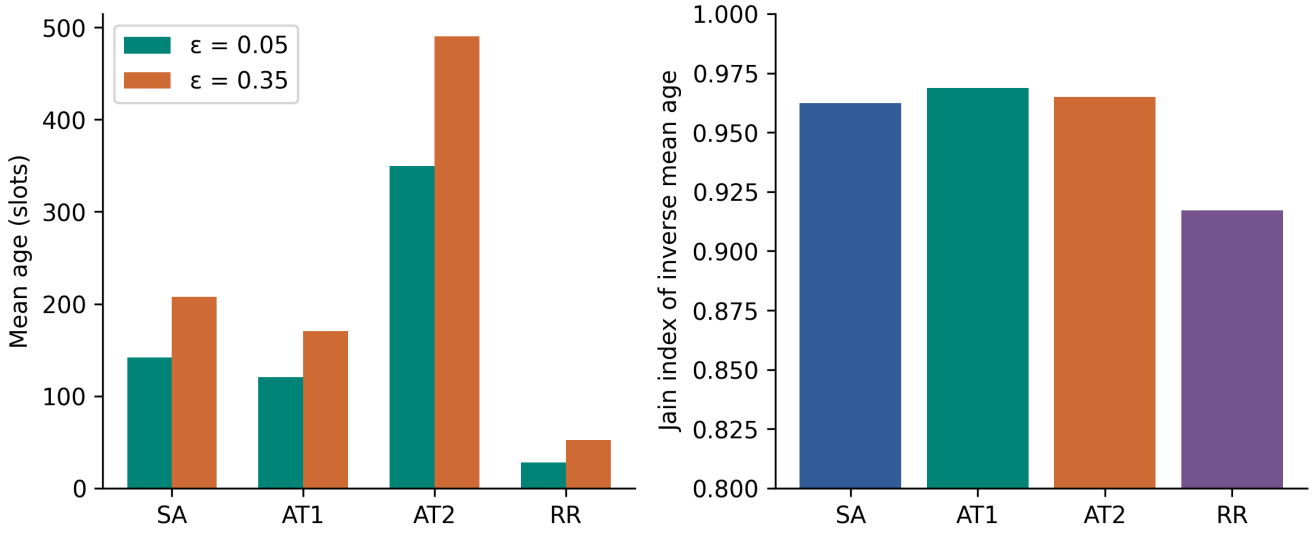


Figure 8. Group-specific age and fairness under unequal independent erasure probabilities. The fairness-axis range is 0.8–1.0; the group-age bars preserve the absolute performance differences.

6. DISCUSSION

6.1 Joint tuning is a resource decision

The practical lesson is to choose a silent interval and access probability together, then evaluate the choice against a transmission budget and an application deadline. The contrast between $(T,c) = (100,4)$ and $(100,2)$ is particularly informative because the former is poor despite having the same threshold as the best point in the tested grid. A protocol description that reports only the threshold would conceal a decisive parameter.

The renewal approximation explains why this dependence is nonlinear. A smaller success probability q lengthens the eligible period, increasing the fraction of peers that can attempt. That raises τ, which lowers the singleton term $(1-\tau)^{N-1}$ and can further reduce q. Under the tested aggressive setting, additional losses push the system into a high-attempt, low-success region. The numerical results establish the operating pattern; they do not claim a formal bifurcation or prove multiple equilibria.

A resource constraint changes the preferred setting even within the small grid. At a budget of one attempt per slot, $(100,2)$ is the lowest-age tested point. At 0.6 attempts per slot, $(100,1)$ becomes preferable among the feasible grid points. This comparison gives a concrete reason to report G next to $\bar{A}$. It also avoids calling every extra attempt an energy loss: whether that attempt is acceptable depends on radio costs and the application's value of freshness.

6.2 Mean, tail, and fairness answer different questions

A receiver that occasionally spends a long interval without a fresh update may be unsuitable for a deadline-sensitive task even when its long-run average age is acceptable. The tail estimator V_D therefore deserves its own constraint. Figure 2 shows that the aggressive setting worsens both mean and tail at moderate loss, while the outage experiment shows a different case in which a comparatively small mean change accompanies a larger relative change in a rare tail event.

The fairness result requires equal care. RR gives both link classes fresher information yet a smaller Jain index of inverse mean age. The metric is behaving as defined: it compares relative equality. A design review should examine the class means before treating a larger J as an improvement, and should specify whether it values equality of opportunities, equality of ages, or a minimum freshness guarantee. Those objectives are mathematically different.

For unequal links, one possible extension is class-specific p_i and T_i. Such a controller would need reliable estimates of per-link performance and safeguards against measurement noise. The present experiment does not optimize these quantities, so it supports their investigation without offering a validated heterogeneous control rule.

6.3 Implications for wireless ad hoc implementation

A peer implementing Equation (2) stores one age counter, a threshold, and a transmission probability. The local decision has constant computational cost. Although that makes the rule straightforward to implement, distributed access still requires synchronized slots and knowledge of successful delivery, so a firmware implementation must maintain the timing and feedback assumptions on which the counter update depends. Clock drift, asymmetric reception, or missed acknowledgments can cause the local counter to differ from the receiver's true age, changing eligibility decisions.

Population estimation is another practical requirement. The rule $p = c/N$ assumes a known number of active peers. If the deployed population is N but a peer uses $\widehat{N}$, the actual aggregate access multiplier is approximately $cN/\widehat{N}$ before threshold suppression. Underestimating population therefore moves operation toward the aggressive side of the grid. The direction follows directly from the parameterization; its quantitative cost under arrivals and departures remains unmeasured.

Spatial reuse would also alter the interpretation of G. In this study, every simultaneous pair conflicts. In a larger ad hoc network, sufficiently separated links may transmit successfully at the same time, so restricting network-wide attempts to one could be unnecessarily conservative. A conflict graph or signal-to-interference model would be required before using these numerical settings outside a single contention domain.

6.4 Limitations and next experiments

The study assumes generate-at-will updates, constant packet duration, perfect feedback, and a fixed peer population. It excludes queueing of exogenous traffic, mobility, hidden terminals, capture, multiple-packet reception, and multi-hop routing. These restrictions isolate access timing, but they prevent claims about end-to-end application delay or standards-compliant wireless performance. The 13 journal references establish the research context; the reported comparisons are against transparent baseline rules, not reproduced implementations of every cited method.

The correlated-outage result for AT2 warrants targeted follow-up. Recording the eligible population and collision rate conditional on channel transitions would help explain its improvement under persistent outages. A broader sensitivity analysis should also vary the deadline, the burn-in duration, the population estimate, and acknowledgment reliability. No parameter in the present model was fitted to field observations.

A hardware or packet-level validation would need to account for synchronization and acknowledgment airtime in the slot duration, measure per-attempt and listening energy, and compare the resulting age distribution with the abstract prediction. Those measurements would determine whether the favorable grid points retain their resource advantage once implementation costs are included. Until then, the numerical settings are reproducible experimental choices rather than deployment recommendations.

7. CONCLUSION

Age-threshold random access offers a measurable freshness benefit only when its silent interval and eligible-node transmission probability are considered jointly. Across the tested single-domain model, a moderate threshold rule reduced mean age relative to ordinary slotted ALOHA, but its advantage narrowed as independent losses increased. A more aggressive rule became collision dominated: at 50 peers and loss probability 0.2, its mean age was 398.94 slots, compared with 168.59 for ordinary access.

Keeping the 100-slot silent interval and reducing the access probability from 0.08 to 0.04 changed that result to 101.46 slots while cutting attempts to 0.835 per slot. This was the lowest-age point in the tested grid and satisfied a one-attempt-per-slot budget. The exact baseline checks, approximation comparison, and longer runs support the internal consistency of the analysis. Correlated outages and unequal links show why age tails and class-specific service must accompany the mean. Extension to spatially reused, mobile, or multi-hop ad hoc networks requires a richer channel model and implementation measurements.

REFERENCES

[1] Q. Abbas, S. A. Hassan, H. K. Qureshi, K. Dev, and H. Jung, "A comprehensive survey on age of informa-tion in massive IoT networks," *Computer Communica-tions*, vol. 197, pp. 199–213, 2023.

[2] I. Kahraman, A. Köse, M. Koca, and E. Anarim, "Age of information in Internet of Things: A survey," *IEEE Internet of Things Journal*, vol. 11, no. 6, pp. 9896–9914, 2024.

[3] F. Zhao, N. Pappas, C. Ma, X. Sun, T. Q. S. Quek, and H. H. Yang, "Age-threshold slotted ALOHA for optimizing information freshness in mobile networks," *IEEE Transactions on Wireless Communications*, vol. 23, no. 11, pp. 17 236–17 251, 2024.

[4] M. Moradian, A. Dadlani, A. Khonsari, and H. Tabassum, "Age-aware dynamic frame slotted ALOHA for machine-type communications," *IEEE Transactions on Communications*, vol. 72, no. 5, pp. 2639–2654, 2024.

[5] S.-S. Tzeng, Y.-J. Lin, and S.-W. Wang, "Packet error-aware threshold ALOHA protocol in wireless networks," *Wireless Personal Communications*, vol. 142, pp. 21–38, 2025.

[6] D. Fiems and A. Vinel, "Mean value analysis of the age of information in slotted ALOHA," *IEEE Communications Letters*, vol. 27, no. 10, pp. 2588–2592, 2023.

[7] D. Fiems and A. Vinel, "The age-of-information distribution in slotted ALOHA," *IEEE Access*, vol. 12, pp. 159 462–159 472, 2024.

[8] J. Wang, J. Yu, X. Chen, L. Chen, C. Qiu, and J. An, "Age of information for frame slotted ALOHA," *IEEE Transactions on Communications*, vol. 71, no. 4, pp. 2121–2135, 2023.

[9] Z. Yue, H. H. Yang, M. Zhang, and N. Pappas, "Age of information under frame slotted ALOHA-based status updating protocol," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 7, pp. 2071–2089, 2023.

[10] A. Buratto, A. Munari, and L. Badia, "Strategic backoff of slotted ALOHA for minimal age of information," *IEEE Communications Letters*, vol. 29, no. 1, pp. 155–159, 2025.

[11] F. Zhao, N. Pappas, M. Zhang, and H. H. Yang, "Age of information in random access networks with energy harvesting," *IEEE Journal on Selected Areas in Communications*, vol. 43, no. 11, pp. 3813–3829, 2025.

[12] W. Fan, P. Fan, and Z. Ding, "On the throughput of NOMA-ALOHA in massive IoT with sparse active users," *IEEE Wireless Communications Letters*, vol. 13, no. 3, pp. 582–586, 2024.

[13] Y. Zhu, Y. Zhu, A. Gong, Y. Lin, Y.-H. Lo, and Y. Zhang, "Age-gain-dependent random access for event-driven periodic updating," *IEEE Transactions on Communications*, vol. 73, no. 8, pp. 5774–5787, 2025.