Age-Optimal Target Wake Time: Provably Good Wake Schedules for Energy-Constrained Wi-Fi Status Updating

Haoyu Wang
Department of Computer Science
University of Massachusetts Boston
Boston, MA, USA
haoyu.wang001@umb.edu

Bo Sheng
Department of Computer Science
University of Massachusetts Boston
Boston, MA, USA
bo.sheng@umb.edu

Xiaoqian Zhang
Department of Computer Science
University of Nebraska Omaha
Omaha, NE, USA
xiaoqianzhang@unomaha.edu

***Abstract*—Target Wake Time (TWT), introduced in IEEE 802.11ax, lets an access point schedule exactly when each station wakes, transmits, and dozes. Existing TWT schedulers optimize energy or throughput, treating information freshness at best as a constraint and offering no performance guarantees. We design the wake schedule itself for freshness: minimize the weighted average Age of Information (AoI) over stations subject to per-station energy budgets, where the decision variables are the TWT triples (wake interval, offset, service period duration). We derive a renewal-exact AoI model for TWT under per-SP block fading and validate it against packet-level 802.11ax simulation with $\sim$1% mean error. We show that, unlike preemptive scheduling, non-preemptive TWT packing can be infeasible at schedule density 1, and identify the granularity condition under which a small-first best-fit packer provably succeeds. Around this we build HARMONIC-GREEDY, a scheduler combining a convex relaxation, anchor-optimized power-of-two rounding, and a best-of-uniform safeguard, and prove it is a constant-factor approximation: $4/\ln 2 \approx 5.77$ under a mild granularity assumption and $6/\ln 2 \approx 8.66$ unconditionally. We implement the complete system in ns-3 – a TWT wake/doze mechanism integrated with the power-save architecture, plus the scheduler – and show that it is the only scheduler that stays near a relaxation lower bound across all regimes: against a strong energy-greedy baseline it ties when per-station energy floors already pin the periods, and wins by 4–36% exactly where the schedule density is binding and must be redistributed by AoI weight or channel quality rather than by energy budget – the regime our analysis identifies.**

***Index Terms*—Age of Information; Target Wake Time; IEEE 802.11ax; scheduling; energy efficiency; ns-3**

I. INTRODUCTION

Real-time monitoring applications – industrial sensing, health wearables, smart-building telemetry – send small periodic status updates over Wi-Fi while running on batteries. Two trends collide in such systems. On one hand, the relevant performance metric is not throughput or even delay, but the *Age of Information* (AoI) [1], [2]: the time elapsed since the generation of the freshest delivered update. On the other, the dominant energy cost is simply being awake: a station that wakes rarely is efficient but stale; one that wakes often is fresh but dies quickly.

IEEE 802.11ax [3] introduced Target Wake Time (TWT) precisely to manage this trade-off mechanically: the AP and a station agree on a periodic schedule of *service periods* (SPs) – a wake interval T, an offset ϕ, and an SP duration d – and the station dozes outside its SPs. TWT thus exposes, for the first time in mainstream Wi-Fi, the exact control surface that AoI theory wants to optimize: *when* each source is allowed to deliver. Yet existing TWT schedulers optimize energy or throughput, treat freshness at best as a constraint, and offer no guarantees (Section II).

This paper designs the TWT schedule itself for freshness, with guarantees. Our contributions:

- **Model** (Sections III–IV): a renewal-exact expression for time-average and peak AoI of a TWT station under per-SP block fading, including the per-wake overhead c_w. A structural lemma shows that batching retransmission attempts within an SP is *never* beneficial without wake overhead; the per-wake cost c_w is thus exactly what separates TWT schedule design from abstract sampling models.
- **Hardness structure** (Section V): a three-station example with schedule density (the fraction of air time the SPs reserve, $\sum_i d_i/T_i$) exactly 1 that is infeasible under *every* placement – non-preemptive TWT packing fundamentally differs from preemptive utilization bounds (harmonic rate-monotonic is optimal at density 1 only because preemption can split service across gap copies).
- **Algorithm and guarantees** (Sections VI–VII): HARMONIC-GREEDY, combining a water-filling convex relaxation, power-of-two period rounding with an anchor chosen by a candidate-argmin rule (evaluate a small, provably sufficient set of grid anchors and keep the best), small-first best-fit packing of power-of-two-sized reservations, and a best-of-uniform safeguard. We prove a deterministic $4/\ln 2 \approx 5.77$ approximation under a mild granularity assumption, and $6/\ln 2 \approx 8.66$ unconditionally. The packing analysis gives a tight granularity threshold $d \le t_0(1-U)$.
- **System** (Section VIII): a complete ns-3 implementation – a TWT wake/doze mechanism integrated with the

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

802.11 power-save architecture, with SP-boundary queue gating and TSF-anchored offsets – validated against the analytical model with $\sim 1\%$ mean error, plus five protocol findings about deploying TWT (e.g., doze announcements cannot be acknowledged over a faded channel; pending retransmissions must be gated or they erase the energy savings).
- **Evaluation** (Section IX): against both a naive (equal-interval) and a strong (energy-greedy) baseline, plus a relaxation lower bound. We separate two effects – period diversity, which the energy-greedy baseline already captures, and AoI-weighting, which we isolate in uniform-budget binding-density regimes where only our scheduler wins (4–6.5% over both baselines, up to 36% at scale with mixed modulation rates), exactly where the theory predicts.

II. Related Work

Age of Information. Since its introduction as a freshness metric [1], AoI has grown into a broad field spanning queueing, sampling, and estimation [2], [4], [5]. Foundational results characterize the freshness-optimal update policy for a single source [6], the effect of packet management and queue discipline [7]–[9], and AoI's role in IoT [10]. AoI now also serves as a system-level metric in edge-AI applications – federated unlearning in mobile edge computing [11], digital-twin healthcare [12] – which presuppose, rather than design, the delivery schedule beneath them. Variants such as the Age of Incorrect Information [13] refine the metric for estimation. Our renewal model (Sec. IV) specializes this machinery to the TWT service-period structure.

AoI scheduling. Minimizing weighted AoI across many sources over a shared channel is the closest body of work. Centralized per-slot policies and their optimality were established for broadcast networks with reliable and unreliable links [14], [15] and with stochastic arrivals [16], [17]; the restless-bandit structure yields Whittle-index policies that are near-optimal and, in cases, asymptotically optimal [18]. These works schedule *which source transmits in each slot* under a centralized controller; we instead design *periodic wake agreements* fixed in advance, where a sleeping station is unreachable between its service periods – a constraint absent from slot-by-slot models. A parallel line constrains energy explicitly – energy-harvesting sources [19], [20], retransmission/ARQ budgets [21], and distributed freshness optimization [22] – but optimizes transmission timing rather than a negotiated periodic wake schedule with a hard duty-cycle budget. Closest in spirit are sleep-wake AoI optimization for CSMA sensors [23], which tunes continuous wake *rates* under battery constraints but has no negotiated periodic structure, and age-agnostic cyclic schedulers [24], which optimize deterministic transmission patterns but model neither energy nor sleeping. A pull-based line measures freshness only when the monitor asks: the Query Age of Information [25] and Whittle-index policies that optimize age at query instants [26]. TWT is structurally dual: the delivery opportunities, not the queries, occur at negotiated instants, and our objective remains the continuous-time average.

TWT and Wi-Fi. TWT was introduced in IEEE 802.11ax to manage energy and contention [3], [27], [28], and is extended in Wi-Fi 7 (802.11be) – notably restricted TWT for latency-sensitive traffic and multi-link operation [29], [30] – making schedule design increasingly central [31]. TASPER [32] schedules TWT SPs to maximize throughput and energy efficiency subject to AoI *constraints*, proves NP-hardness, and gives a heuristic without guarantees; timely-throughput TWT grouping maximizes deadline-met throughput via drift-plus-penalty [33]; experimental IIoT schedulers [34] respect AoI deadlines deterministically; and TWT mechanism implementations and measurement studies [35]–[37] provide the platform context. Single-device closed-form AoI analyses exist for the cellular sleep-mode analogue of TWT (DRX) [38] and, most recently, for the TWT mechanism itself [39] – performance analysis of given configurations, not schedule design across a station *population*, which is what we do here. To our knowledge no prior work makes AoI the *objective* of TWT schedule design or provides approximation guarantees.

Periodic real-time scheduling. Our packing analysis connects to classical real-time theory. Rate-monotonic scheduling is optimal for preemptive periodic tasks [40], and harmonic task sets are schedulable up to utilization 1 – but only because preemption can split a job across the gaps left by higher-rate tasks. TWT service periods are *contiguous and non-preemptive*, which is exactly why our Proposition 2 shows density 1 can be infeasible. The relevant non-preemptive analogue is the pinwheel problem and distance-constrained scheduling [41]–[43], whose density thresholds and harmonic constructions inform our small-first best-fit packer; we adapt them to the AoI-weighted, energy-budgeted TWT setting.

III. System Model

N stations are associated with one AP. Station i holds a TWT agreement (T_i, ϕ_i, d_i), where T_i is its *wake interval* (the time between consecutive wake-ups), ϕ_i its *offset* (the start of its first SP on the shared clock), and d_i its *SP duration* (the contiguous awake time per wake): the station wakes at times $\phi_i + kT_i$, $k = 0, 1, \ldots$, stays awake for d_i, and dozes otherwise (Fig. 1; Table I collects notation). SPs of different stations must not overlap: keeping them disjoint is the scheduler's responsibility, and it is precisely what designs contention away. Station i samples its source at the start of each SP (generate-at-will) and transmits the update uplink within the SP.

Channel. We model per-SP block fading: the channel coherence time exceeds the SP duration, so all transmission attempts within one SP share fate; an SP succeeds with probability $p_i \in (0, 1]$, independently across SPs. This is the regime in which the retransmission policy genuinely matters for AoI. At the opposite extreme – per-attempt i.i.d. loss – a few in-SP retries drive the per-SP success probability to 1 and the channel dimension of the problem disappears; we verify this empirically in Section IX.

TABLE I
NOTATION.

(T_i, ϕ_i, d_i)	TWT triple: wake interval, offset, SP duration
ρ_i; $T_i^{\min}$	duty-cycle budget; energy floor d_i/ρ_i
p_i	per-SP success probability (block fading)
q_i, k_i, s_i	per-attempt success, attempts per SP, attempt airtime
$c_{w,i}$	per-wake overhead (ramp-up, synchronization)
δ_i	mean in-SP delivery delay
w_i; a_i	AoI weight; cost coefficient $w_i(1/p_i - 1/2)$
$\bar{A}_i$, $\bar{A}_i^{\text{peak}}$	time-average and peak AoI
U	schedule density $\sum_i d_i/T_i$
$d_{\max}$	largest SP duration $\max_k d_k$
t_0, level j	harmonic grid anchor; grid period $t_0 2^j$
$\rho^\star$; λ	scheduler density target; water-filling multiplier

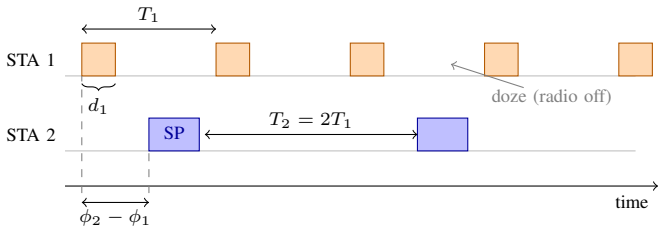


Fig. 1. TWT wake schedules: station i wakes for d_i every T_i at offset ϕ_i on a common (TSF-anchored) grid; the scheduler chooses the triples so that service periods are disjoint.

Energy. The station's energy cost is its awake time: each wake costs a fixed overhead $c_{w,i}$ (radio ramp-up, synchronization) plus the SP airtime. Station i has a duty-cycle budget ρ_i: $d_i/T_i \le \rho_i$, equivalently $T_i \ge T_i^{\min} := d_i/\rho_i$.

Objective. Let $A_i(t)$ be the age of the freshest update from station i delivered to the AP. We minimize the weighted time-average AoI $\sum_i w_i \bar{A}_i$; Section VII notes that all results carry to weighted peak AoI.

Scope. We deliberately study the cleanest setting that isolates schedule design: uplink status updates, generate-at-will sampling, scheduled non-overlapping SPs (so contention is designed away), and per-SP block fading. These assumptions are standard in AoI scheduling [14], [24] and matched to monitoring IoT, the dominant TWT use case. Three violations deserve comment. *Stochastic arrivals:* if updates arrive randomly rather than on demand, the wake schedule itself stays feasible unchanged, but the sampled update may already be stale at wake time; the objective acquires an arrival-age component that couples the stations only through the density constraint, turning the design into a constrained-MDP problem (Section X). *Contention:* traffic from non-TWT stations can collide with SPs; such losses enter the model through p_i as long as they are independent across SPs, and where that cannot be assumed, 802.11be restricted TWT provides protected SPs. *Intra-SP channel variation:* block fading is the boundary case in which retransmission policy matters at all – under per-attempt i.i.d. loss, in-SP retries drive the SP success probability to 1 and the problem degenerates gracefully (Section IV; verified empirically in Section IX). Downlink AoI (broadcast TWT, AP-side delivery) is likewise deferred to Section X.

IV. Single-Station Structure

Successful SPs form a renewal process; the standard renewal-reward AoI argument [1], [2] then gives the following, which is the linchpin that turns TWT scheduling into a linear objective (Section V).

Proposition 1 (TWT AoI)**.** *With wake interval T, per-SP success probability p, and mean in-SP delivery delay δ,*

$$\bar{A} = \delta + T\Big(\frac{1}{p} - \frac{1}{2}\Big), \qquad \bar{A}^{\text{peak}} = \delta + \frac{T}{p}.$$

Sketch. Let G be the number of wake intervals between successive successful SPs. Since each SP succeeds independently with probability p, G is geometric on $\{1, 2, \dots\}$, so $\mathbb{E}[G] = 1/p$ and $\mathbb{E}[G^2] = (2-p)/p^2$. The inter-delivery time is $X = GT$. After a successful delivery the AoI drops to the in-SP delay (mean δ) and then grows linearly until the next success, so each renewal cycle contributes area $\delta X + X^2/2$. Dividing expected area by expected cycle length gives $\bar{A} = \delta + \mathbb{E}[X^2]/(2\mathbb{E}[X]) = \delta + T(1/p - 1/2)$, and the per-cycle peak $\delta + X$ averages to $\bar{A}^{\text{peak}} = \delta + \mathbb{E}[X] = \delta + T/p$. □

When the channel is per-attempt i.i.d. with success probability q and the SP holds k attempts, $p = 1 - (1-q)^k$ and $d = ks + c_w$ with s the per-attempt airtime. Choosing k under the energy-tight period $T = d/\rho$ reveals a clean dichotomy:

Lemma 1 (batching needs overhead)**.** *With $c_w = 0$, the single attempt $k^* = 1$ is optimal for every q: writing $f(k) = (k + c_w/s)(1/p(k) - 1/2)$, one computes $f(2) - f(1) = q/(2(2-q)) > 0$. With $c_w > 0$, $k^* > 1$ for sufficiently large c_w/s at any fixed q.*

The per-wake overhead is thus load-bearing: it is exactly what distinguishes TWT schedule design from abstract sampling models, and we keep it in the model throughout.

We validated Proposition 1 against packet-level 802.11ax simulation (Section IX): measured time-average AoI matches within 0.02% ($p \approx 1$, two parameter sets) and within 0.6–1.2% (seed-averaged) under per-SP block fading with $p \in \{0.8, 0.6, 0.4\}$, with measured duty cycles equal to d/T to four decimals.

V. The Multi-Station Problem

By Proposition 1 the schedule cost is $J(T) = \sum_i a_i T_i + C$ with $a_i := w_i(1/p_i - 1/2)$ and C a T-independent constant; the design problem is

$$\min_{T,\phi} \ \sum_i a_i T_i \qquad \text{s.t. SPs pairwise disjoint}, \quad T_i \ge T_i^{\min}. \tag{1}$$

Disjointness implies the *density* bound $U = \sum_i d_i/T_i \le 1$, but the converse fails in a structural way:

Proposition 2 (density 1 can be unpackable)**.** *Let t_0 be a base period and consider stations A_1: $(T{=}t_0, d{=}t_0/2)$, A_2: $(T{=}t_0, d{=}t_0/4)$, B: $(T{=}2t_0, d{=}t_0/2)$, with $U = 1$. No placement of offsets makes the SPs disjoint.*

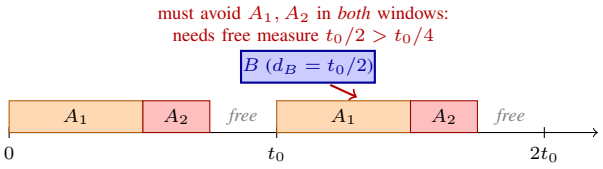


Fig. 2. Proposition 2: at density 1, the per-window free set has measure $t_0/4$, but B's contiguous SP needs $t_0/2$ of it – in every window, under every placement.

Proof. Both A_1 and A_2 have period t_0, so their SPs trace the same pattern in every base window $[kt_0, (k+1)t_0)$, and the free set S left in each window, of measure $t_0/4$, is likewise identical from window to window. Now place B: its SP is a contiguous interval of length $t_0/2 \le t_0$, so it meets at most two consecutive windows, splitting into (at worst) a suffix of one window and a prefix of the next. Each piece must avoid A_1 and A_2, i.e. lie inside that window's copy of S; and since S repeats identically, the two pieces occupy *disjoint* parts of the same per-window set. Their lengths sum to $t_0/2$, forcing $|S| \ge t_0/2 > t_0/4$ – a contradiction (Fig. 2). □

Remark 1. Under *preemptive* scheduling the same instance is feasible (harmonic rate-monotonic is optimal at utilization 1): preemption splits B's service across the two gap copies. TWT SPs are contiguous by design, so utilization intuition from real-time scheduling does not transfer. Any packing guarantee must involve granularity, which our analysis makes precise (Lemma 5).

A second structural fact bounds every period from below:

Lemma 2 (long SPs dominate periods)**.** *In any feasible schedule, $T_i \ge d_{\max} := \max_k d_k$ for all i.*

Proof. Suppose some station A had $T_A < d_B$ for a station B. The wake starts of A are spaced T_A apart, so at least one falls strictly inside B's SP of length d_B, and the A-SP beginning there overlaps it. Hence every period is at least the longest SP duration. □

The following small instance is carried through Sections VI–VII to make each stage of the design concrete.

Example 1 (running)**.** Three stations with equal SPs $d_i = 4$ ms, uniform duty budgets $\rho_i = 0.5$ (energy floors $T_i^{\min} = 8$ ms), reliable channels ($p_i = 1$, so $a_i = w_i/2$), and weights $w = (8, 2, 1)$, i.e. $a = (4, 1, \frac{1}{2})$. At the floors the density would be $3 \cdot 4/8 = 1.5 > 1$: the air time the stations *want* exceeds what exists, so the density constraint binds and the periods must be differentiated – by weight, not by budget, since the budgets are identical: exactly the regime our evaluation isolates (Section IX).

VI. The Harmonic-Greedy Scheduler

Algorithm 1 lists the scheduler. Line 1 sizes each station's SP by the single-station dichotomy of Section IV and precomputes the cost coefficients a_i. Line 2 solves the convex relaxation below; its water-filling solution shortens the periods of high-weight, poor-channel stations first. Line 3 snaps the ideal periods onto a power-of-two grid – this is what makes the later packing tractable – with the grid anchor chosen by the candidate-argmin rule analyzed in Corollary 1. Line 4 reinvests the slack that rounding up creates into extra transmission attempts. Line 5 places offsets, and line 6 guards against the octave-quantization pathology discussed at the end of this section. The relaxation of line 2 replaces disjointness by its density consequence:

Algorithm 1 Harmonic-Greedy

Require: per station i: weight w_i, duty budget ρ_i, per-attempt success q_i, attempt airtime s_i, wake overhead $c_{w,i}$; density target $\rho^\star$

Ensure: disjoint TWT triples (T_i, ϕ_i, d_i) with $d_i/T_i \le \rho_i$

1: *size SPs:* $k_i \leftarrow$ energy-tight optimal attempts (Lemma 1); $d_i \leftarrow k_i s_i + c_{w,i}$; $p_i \leftarrow 1-(1-q_i)^{k_i}$; $a_i \leftarrow w_i(1/p_i - \frac{1}{2})$; $T_i^{\min} \leftarrow d_i/\rho_i$

2: *relax:* $T_i \leftarrow \max(T_i^{\min}, d_{\max}, \sqrt{\lambda d_i/a_i})$, bisecting $\lambda \ge 0$ until $\sum_i d_i/T_i \le \rho^\star$ (Lemma 3)

3: *anchor and round:* fold each T_i by powers of two into $(T_{\min}/2, T_{\min}]$, $T_{\min} = \min_i T_i$; among these candidate anchors pick the t_0 minimizing $\sum_i a_i \lceil T_i \rceil_{t_0}$, where $\lceil \cdot \rceil_{t_0}$ rounds up on the grid $\{t_0 2^j\}$ (Corollary 1); set $T_i \leftarrow \lceil T_i \rceil_{t_0}$

4: *re-expand:* for stations with $p_i < 1$, grow k_i (hence d_i) into the rounding slack, capped by ρ_i – extra attempts at fixed T_i are free for AoI

5: *pack:* in increasing-T_i order, place each ϕ_i by small-first best-fit (certified variant: dyadic reservations; practical variant: first-fit of the true d_i; Section VII)

6: *safeguard:* $T_u \leftarrow \max(\max_i T_i^{\min}, \sum_i d_i)$; if $\sum_i a_i T_u$ beats the rounded cost, or packing failed, return the uniform schedule (offsets staggered back-to-back) instead

$$J^* = \min \Big\{ \sum_i a_i T_i : \sum_i \frac{d_i}{T_i} \le 1,\ T_i \ge \max(T_i^{\min}, d_{\max}) \Big\}, \tag{2}$$

Every feasible schedule satisfies the constraints of (2) – disjoint SPs occupy at most the full timeline, and Lemma 2 supplies the $d_{\max}$ floor – so J^* lower-bounds (1).

Lemma 3 (water-filling)**.** (2) *is convex; its solution is $T_i(\lambda) = \max\left(T_i^{\min}, d_{\max}, \sqrt{\lambda d_i/a_i}\right)$ with $\lambda \ge 0$ the density multiplier (found by bisection).*

Sketch. The objective is linear and each d_i/T_i is convex on $T_i > 0$, so (2) is a convex program and KKT conditions are sufficient. For fixed λ the Lagrangian separates across stations; its i-th term $a_i T_i + \lambda d_i/T_i$ is strictly convex with unconstrained minimizer $\sqrt{\lambda d_i/a_i}$, so the coordinatewise minimizer over the feasible ray is that value clipped at the floor $\max(T_i^{\min}, d_{\max})$ – exactly $T_i(\lambda)$. The density $\sum_i d_i/T_i(\lambda)$ is continuous and nonincreasing in λ, so bisection finds the smallest λ at which it meets the density budget ($\lambda = 0$ if the floors alone satisfy

it), which is complementary slackness. Full proof: technical report. □

Two implementation-driven design elements deserve emphasis, both discovered by simulation. First, the *anchor matters*: rounding to powers of two of an arbitrary base can double a period. Writing $T_{\min} = \min_i T_i$ for the smallest relaxed period, we instead choose the anchor among the ideal periods folded into $(T_{\min}/2, T_{\min}]$ – each candidate makes one station's rounding lossless – which is provably no worse than the average over random anchors (Corollary 1) and empirically much better. Second, the *safeguard matters*: when the relaxation's period spread is sub-octave (all ideal periods within a factor 2), the harmonic grid cannot express the differentiation and rounding losses dominate. For example, ideal periods of 30 and 50 ms (ratio $1.67 < 2$) round to the *same* power-of-two grid level, so the harmonic schedule cannot give the more important station a shorter period and merely pays the rounding cost; the safeguard then returns the uniform schedule instead. We call this the *octave quantization* effect and quantify it in Section IX.

Example 2 (continued)**.** On the instance of Example 1, water-filling clamps station 1 to its floor and spreads the remaining density by weight: $T^* = (8, 13.7, 19.3)$ ms, density exactly 1, relaxed cost $J^* = \sum_i a_i T_i^* = 55.3$ (a lower bound). Folding T^* into $(4, 8]$ gives the anchor candidates $\{8, 6.83, 4.83\}$; rounding up on each candidate grid costs $\{64.0,\ 81.9,\ 67.6\}$, so line 3 picks $t_0 = 8$ and $T = (8, 16, 32)$ ms – station 1's rounding is lossless and the cost is $1.16\times$ the relaxation bound. The uniform alternative $T_u = \max(8, 3{\cdot}4) = 12$ ms would cost $66.0 > 64.0$, so the safeguard keeps the harmonic schedule. Packing in increasing-period order places $\phi = (0, 4, 12)$ ms, disjoint over the 32 ms hyperperiod.

VII. Approximation Guarantees

The practical scheduler packs true SP durations at density 0.9. A worst-case certificate for that exact configuration would have to survive adversarial duration mixes; we instead analyze a deliberately conservative variant of Algorithm 1 and charge the difference to the constants. The analyzed variant differs in two ways. First, it runs the relaxation at a density target $\rho^\star$ (fixed to $1/4$ in Theorem 1), leaving the packing stage guaranteed headroom. Second, it packs dyadic *reservations*: each station reserves the smallest grid divisor $t_0/2^m \geq d_i$ and transmits only for d_i inside its reservation, so reservations consume schedule space but no energy. Packing is small-first best-fit (SF-BF): process stations by increasing period; place each into a minimum-size free block of sufficient size, carving from the left and returning the staircase remainder. The analysis then proceeds in three steps: a structural invariant on the free blocks SF-BF maintains (Lemma 4), a sufficient packing threshold (Lemma 5), and the end-to-end guarantees (Theorems 1 and 2, sharpened by the anchor rule of Corollary 1).

Lemma 4 (count invariant)**.** *When SF-BF serves a request of size q from a chosen free block of size b, the staircase blocks it returns have sizes in $[q, b)$ – sizes at which, by best-fit minimality of b, no free block currently exists. Consequently new free blocks are only created at sizes whose count is zero, and at the start of level-j processing every size has at most 2^j free blocks. Both bounds are tight.*

Lemma 5 (packing threshold)**.** *If every reservation satisfies $d'_i \leq t_0(1-U)$ with $U = \sum_i d'_i/T_i$, SF-BF places all stations. Conversely (cf. Proposition 2), a station of level ≥ 1 must satisfy $d' \leq t_0(1-U_0)$ where U_0 is the lower-level utilization: the threshold is tight up to the requesting station's own density share.*

Proof sketch. Suppose a request of size q fails at level j. Every free block is then smaller than q, and since each of the sizes $q/2, q/4, \ldots$ carries at most 2^j blocks (Lemma 4), the total free measure is below $2^j q$. On the other hand, the stations placed so far occupy at most a fraction U of the level-j window of length $t_0 2^j$, leaving free measure at least $t_0 2^j(1-U) \geq 2^j q$ under the stated granularity – a contradiction. Full proofs are in the technical report, and all lemmas are additionally machine-checked on ∼3,000 randomized harmonic instances. □

Theorem 1. *Under the granularity assumption $d_i \leq t_0/4$,* HARMONIC-GREEDY *(analysis variant, $\rho^\star = 1/4$) returns a feasible schedule of cost at most $8 \cdot \mathrm{OPT}$.*

Corollary 1 (candidate-argmin anchor)**.** *The rounded cost, as a function of the anchor phase $\theta \in [0, 1)$ ($t_0 = T_{\min} 2^{-\theta}$), is piecewise strictly decreasing with minima exactly at the folded candidates; hence the candidate-argmin anchor satisfies $\min_{\mathrm{cand}} C \leq \mathbb{E}_\theta[C] = (1/\ln 2)\sum_i a_i T_i$, improving Theorem 1 deterministically to $4/\ln 2 \approx 5.77$.*

Theorem 2 (unconditional)**.** *Without any granularity assumption, scaling the relaxation solution by $s = 6$ (which simultaneously enforces the density budget and, via Lemma 2, the long-SP margin $t_0 > 3d_{\max}$) yields a 12-approximation, and $6/\ln 2 \approx 8.66$ with the candidate-argmin anchor. The constant 6 is optimal for this accounting.*

Example 3 (concluded)**.** On Example 1 after rounding, $U = 4/8 + 4/16 + 4/32 = 0.875$, so Lemma 5 would require $d \leq t_0(1-U) = 1$ ms; the instance violates the threshold ($d = 4$ ms) yet packs anyway – the condition is sufficient, not necessary, and the certified variant's $\rho^\star = 1/4$ exists precisely to guarantee it in the worst case. The achieved cost, $1.16\times$ the relaxation bound, is typical of the gap between certificate and behavior: across every regime of Table III the practical variant lands within 1.05–$1.25\times$ the relaxation lower bound – itself unachievable in general – against the $5.77\times$ worst-case certificate. The bound's role is to exclude pathological instances, not to predict typical loss.

Proposition 3 (peak AoI)**.** *All of the above holds verbatim for weighted peak AoI: by Proposition 1 the objective is again linear in T with coefficients $a_i^{\mathrm{peak}} = w_i/p_i$.*

For $N = 2$ the problem admits an *exact* solution, with a structural twist:

Theorem 3 (exact N=2)**.** *Feasibility forces a rational period ratio $T_2/T_1 = p/q$ with $T_1 \geq q(d_1 + d_2)$. For each q the cost $J(r, q) = (a_1 + a_2 r)\max(T_1^{\min}, T_2^{\min}/r, q(d_1+d_2))$ is quasiconvex in $r = p/q$ – strictly decreasing while the $T_2^{\min}/r$ term is active, then strictly increasing. Its integer-numerator optimum is therefore one of the two grid neighbors of the continuous minimizer, and q ranges over a finite, explicit set, yielding an exact $O(Q)$ algorithm (verified against exhaustive search on* $4{,}000$ *instances). Notably,* non-integer *ratios are strictly optimal on an open set of instances – both energy floors can bind only at a fractional ratio – so the harmonic restriction has a real price already at $N = 2$, precisely the rounding loss our analysis charges for. Full proof: technical report.*

Full proofs of all results appear in the appendix of the technical-report version.

VIII. Implementation in ns-3

We implement the full system in ns-3.[1] It comprises a `Twt PowerSaveManager` realizing individual implicit TWT as a subclass of the 802.11 power-save architecture, and the scheduler as a pluggable policy module. Three mechanism elements proved essential, each corresponding to a measured failure mode: (i) *deferral, not wake-on-traffic*: channel access requested outside an SP must not wake the radio – queued frames contend automatically at the next SP start; (ii) *SP-boundary queue gating*: without it, pending retransmissions keep the station awake until the fade lifts (we measured a retry storm of 1.5×10^5 transmissions), erasing the energy savings. We gate by *freezing* rather than dropping: at SP end we mark the station's per-access-category EDCA queues blocked through ns-3's existing power-save block-reason mechanism (the same hook the standard PS manager uses), and release it at the next SP start. Blocked frames are not dequeued for transmission, so their sequence numbers, retry counts, and Block-Ack state are preserved untouched; the deferred updates simply ride the next SP. This restores the exact k-attempts-per-SP structure the model assumes and keeps the PHY in sleep between SPs; (iii) *TSF-anchored offsets*: per-station offsets must live on a common absolute time grid – otherwise the computed disjointness has no shared reference and is silently lost. We further document protocol findings on TWT deployment: beacon-loss monitoring disassociates dozing stations; downlink frames (including ADDBA responses) buffered for a dozing station deadlock the uplink through asymmetric Block Ack state; and doze announcements are best-effort by nature – a PM indication cannot be acknowledged over a channel the AP cannot hear, which is precisely why TWT is schedule-based.

[1]Code, data, and a reproduction guide are available at https://github.com/Haoyu2/age-optimal-twt – an ns-3 `contrib` module plus the `src/wifi` mechanism patch, with the raw sweep results, released to support artifact evaluation. The extended version with full proofs is available as arXiv:2608.21596.

Scheduler cost. The scheduler runs on the AP once per (re)association or renegotiation epoch, not per frame. Lines 1–2 of Algorithm 1 are $O(N)$ (the bisection runs a fixed number of iterations to machine precision), anchor selection evaluates N candidate grids at $O(N)$ each, and packing touches each SP instance of the hyperperiod. Measured on a single commodity-VM core (median of 20 runs, heterogeneous synthetic workloads): 0.7 ms at $N = 64$ stations, 9 ms at $N = 256$, and 43 ms at $N = 512$ – negligible against association signaling. The runtime state is one triple per station plus the per-AC gating flags, and the data path adds only the two $O(1)$ queue gate/ungate operations per SP boundary, reusing the existing power-save block-reason hook.

IX. Evaluation

Setup. All experiments use 802.11ax at 2.4 GHz over a 20 MHz channel (HeMcs3 unless stated otherwise), with 10–50 stations of heterogeneous weights, duty budgets, and per-SP fading. Each station delivers one status update per wake interval, and the AP integrates AoI exactly, event by event. Runs last 300 s over 10 seeds; we report 95% confidence intervals.

Baselines. We compare Harmonic-Greedy against two baselines. *Equal-interval* gives all stations a single common period (the naive TWT schedule). *Energy-greedy* is a strong heuristic in the spirit of energy-efficiency-first TWT schedulers such as TASPER [32]: it gives each station its shortest budget-feasible period $\max(T_i^{\min}, d_{\max})$, uniformly stretches to feasibility if overbooked, and reuses our rounding, packing and safeguard. It is period-*diverse* but AoI-weight-*unaware*, so it isolates the value of our weighted period assignment: any margin of Harmonic-Greedy over Energy-greedy is attributable to that step alone. The two baselines let us separate two questions – does period diversity help (Energy-greedy vs. Equal-interval), and does AoI-weighting help beyond it (Harmonic-Greedy vs. Energy-greedy). Slot-based AoI policies (max-weight, Whittle index [14], [18]) are not directly comparable: they assume the scheduler can reach any station in any slot, whereas a dozing TWT station is unreachable between its negotiated SPs, and re-deriving index policies over wake agreements is the stochastic-arrivals extension of Section X. The relaxation lower bound plays the comparison role instead: it bounds the cost of *every* feasible TWT schedule – hence of any AoI policy adapted to TWT – and Harmonic-Greedy stays within 5–25% of it below.

Energy–freshness frontier. Figure 3 sweeps the duty-cycle budgets (8× weight skew): the schedulers trace the AoI–energy trade-off, with Harmonic-Greedy dominating in the energy-constrained region where ideal-period spread exceeds an octave and coinciding with the baseline where the safeguard detects collapse.

Model validation. (Fig. 4, Table II) Across the loss grid $e \in \{0, 0.1, \ldots, 0.6\}$, measured AoI tracks $\delta + T(1/p - 1/2)$ with mean error ∼1%: 0.003% at $p = 1$ on every seed, and a worst single seed of 4.4% at p=0.4 – consistent with the geometric-tail variance of a 300 s run. The measured renewal

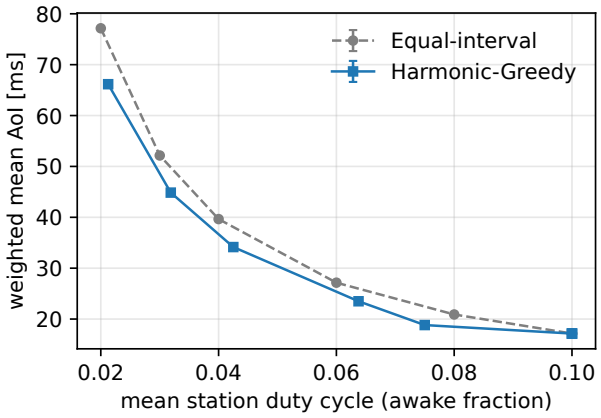


Fig. 3. Weighted mean AoI vs. mean duty cycle across energy budgets (8× weight skew, 10 stations, 5 seeds).

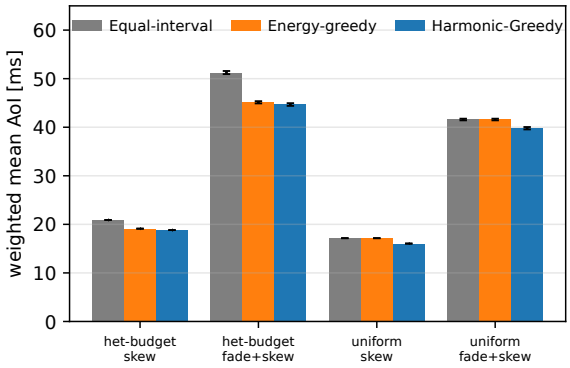


Fig. 5. Weighted mean AoI by regime (10 seeds, 95% CI). With heterogeneous budgets (left) the strong energy-greedy baseline already matches HARMONIC-GREEDY; with uniform budgets and binding density (right) it collapses to equal-interval and only AoI-weighting wins.

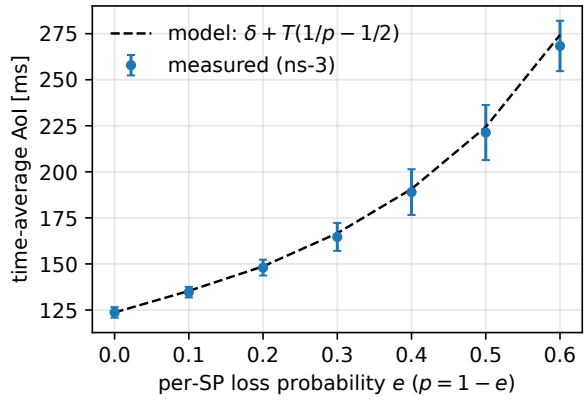


Fig. 4. Single-station validation: measured time-average AoI (3 seeds, 95% CI) vs the renewal model, across per-SP loss probabilities (T=100 ms, d=8 ms).

components $\mathbb{E}[G]$ and $\mathbb{E}[G^2]/2\mathbb{E}[G]$ are within 1% of their predicted values T/p and $T(1/p - 1/2)$.

Period diversity vs. AoI-weighting. (Fig. 5, Table III) The two effects separate cleanly by regime. *When budgets are heterogeneous*, the per-station energy floors $T_i^{\min} = d_i/\rho_i$ already pin the periods, so "minimize every period then pack" is near-optimal: Energy-greedy captures essentially the whole gain over equal-interval and HARMONIC-GREEDY only matches it (het-budget skew: 18.8 vs 19.1 ms, both vs 20.9 for equal-interval; fading+skew similar). We report this plainly – period diversity, not AoI-weighting, drives those gains, and the LB column confirms both schedulers already sit near the optimum there. *When budgets are uniform*, Energy-greedy collapses to equal-interval (identical floors give identical periods), and only the AoI-weighted assignment redistributes the binding schedule density toward the stations that matter: under uniform budgets with binding density, HARMONIC-GREEDY wins 6.5% with weight skew (16.0 vs 17.2 ms) and 4.3% with skew plus fading – the latter via the channel-aware coefficient $a_i \propto 1/p_i$. This is the regime that isolates the paper's contribution. Weighted *peak* AoI tracks the same pattern (Proposition 3); duty cycles equal d_i/T_i to four decimals and SPs never overlap in all runs.

Near-optimality. The LB column of Table III is the relaxation optimum – a lower bound on *any* feasible schedule. HARMONIC-GREEDY is the only scheduler that stays close to LB in *every* regime; Energy-greedy is near-LB when energy floors pin the periods but drifts away once the binding density must be redistributed by weight or channel rather than by budget. The residual gap to LB is the price of harmonic rounding and contiguous (non-preemptive) packing, which Proposition 2 shows is unavoidable in general.

Scaling and decomposition. (Fig. 7, Fig. 6) The pipeline runs cleanly to $N = 50$ stations (exact duty cycles, zero SP overlap, predictions within 2%). Heterogeneous SP sizes – alternating HeMcs0/HeMcs7, 700-byte updates, airtimes differing ∼3× – give the cleanest decomposition of the two effects at scale (3 seeds, CIs within 0.1 ms). At N=50: equal-interval 38.5 ms; Energy-greedy 29.5 ms (a 23% gain from period diversity alone, as fast stations need not wait for slow ones); and HARMONIC-GREEDY 24.6 ms (a further 17% from AoI-weighting, for 36% over equal-interval). Both ingredients contribute and HARMONIC-GREEDY beats the strong baseline outright here, because diverse airtimes create both period diversity *and* a binding density that AoI-weighting can exploit. (N=20 shows the same split: 16.8 → 13.3 → 11.2 ms.)

Certified vs. practical variant. HARMONIC-GREEDY has two configurations that share the relaxation, anchor rule, and safeguard but differ in the packing stage (Sec. VI). The *certified* variant (dyadic reservations, $\rho^\star$=1/4) is the one Theorems 1–2 bound; the *practical* variant ($\rho^\star$=0.9, first-fit of true SP durations) is the one that produces the gains above. The conservative $\rho^\star$=1/4 is forced by the two factor-of-two losses the worst-case packing analysis must absorb in series: rounding each SP duration up to a dyadic reservation can double the schedule density, and the packing threshold

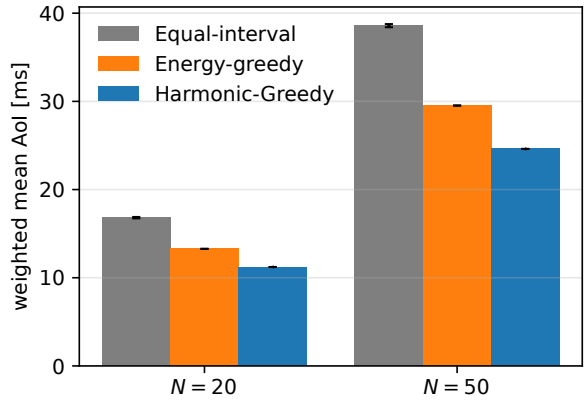


Fig. 6. Heterogeneous modulation rates (HeMcs0/HeMcs7 alternating, 700-byte updates, 8× weight skew, 3 seeds): the gain decomposes into period diversity (energy-greedy over equal-interval) and AoI-weighting (HARMONIC-GREEDY over energy-greedy); 36% total at N=50.

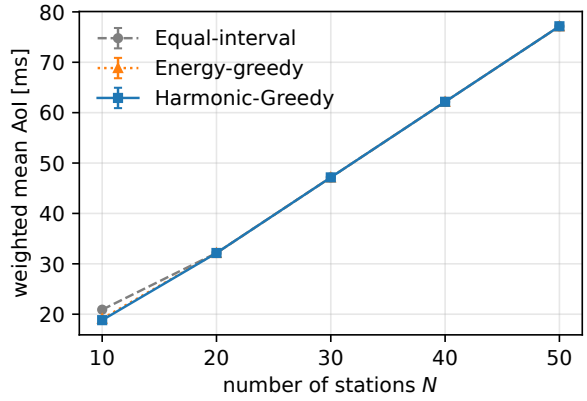


Fig. 7. Scaling with the number of stations (8× weight skew, heterogeneous budgets, 3 seeds): all three schedulers track closely once the schedule is capacity-bound (energy floors pin the periods).

(Lemma 5) admits a request only when $d' \leq t_0(1-U)$, which the count invariant satisfies for $U \leq 1/2$; starting from $\rho^\star$=1/4 leaves exactly the headroom for both. The practical variant instead packs the true SP durations greedily, so neither loss applies and density 0.9 is feasible – but it has no worst-case certificate. The two are thus not interchangeable: the theorem certifies the *template* (relaxation, anchor rule, packing, safeguard), while the high-density practical variant is what we evaluate throughout. Under its conservative budget the certified variant's safeguard returns the uniform schedule in the heterogeneous-budget regimes, so it tracks equal-interval there rather than the practical variant; closing this gap – a tighter packing analysis that certifies a density target near 0.9 – is open (Sec. X).

X. CONCLUSION

TWT turns freshness-versus-energy into a schedule design problem. We formulated it with AoI as the objective, derived a renewal-exact model validated at the packet level, exposed the non-preemptive structure that separates TWT packing from classical utilization bounds, and built HARMONIC-GREEDY – to our knowledge the first TWT scheduler with approximation guarantees – together with an exact N=2 solution and an open ns-3 implementation.

For practitioners, the evaluation reads as a regime map. When per-station energy floors pin the periods, energy-first scheduling is already near-optimal and AoI-weighting adds nothing. When the schedule density binds under uniform budgets or mixed modulation rates, AoI-weighted assignment is the only policy that redistributes air time toward the stations that matter, worth 4–36%.

Four directions remain open. *Tighter packing certificates*: the certified $\rho^\star$=1/4 is conservative only because the worst-case analysis absorbs two compounding factor-of-two losses in series – dyadic duration rounding can double the density, and the non-preemptive threshold (a consequence of the density-1 infeasibility of Proposition 2) admits requests only up to half the window; the practical $\rho^\star$=0.9 variant never triggers either worst case, so an amortized packing argument would likely certify a density target near 0.9 – the single highest-leverage next step. *Stochastic arrivals*: index policies over wake agreements, with asymptotic-optimality machinery. *Exact solutions beyond* N=2. *Downlink AoI*: broadcast TWT with AP-side schedule-aware delivery.

TABLE II
SINGLE-STATION MODEL VALIDATION (SUMMARY).

Regime	Predicted	Error
$p \approx 1$, T=100/8 ms	$\delta + T/2$	0.016%
$p \approx 1$, T=50/4 ms	$\delta + T/2$	0.065%
block fading p=0.8/0.6/0.4	$\delta + T(1/p - 1/2)$	$\leq$ 1.2% (mean)

TABLE III
WEIGHTED MEAN AOI (MS; MEAN ± 95% CI OVER 10 SEEDS), 10 STATIONS, 300 S. EQUAL = EQUAL-INTERVAL, EG = ENERGY-GREEDY (STRONG BASELINE), H-G = OURS, LB = RELAXATION LOWER BOUND (UNACHIEVABLE). TOP: HETEROGENEOUS BUDGETS – EG ALREADY NEAR-OPTIMAL, H-G MATCHES. BOTTOM: UNIFORM BUDGETS, BINDING DENSITY – EG COLLAPSES TO EQUAL, ONLY H-G (AOI-WEIGHTING) WINS.

Regime	Equal	EG	H-G	LB
het. budget, skew	20.9	19.1	18.8	15.8
het. budget, fade+skew	51.3	45.1	44.7	35.7
unif. budget, skew	17.2	17.2	**16.0**	15.2
unif. budget, fade+skew	41.6	41.6	**39.8**	35.9

ACKNOWLEDGMENT

This material is based upon work supported in part by the National Science Foundation under Award No. CNS-2450832. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*, 2012.

[2] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, 2021.

[3] IEEE, "IEEE standard 802.11ax-2021: Enhancements for high-efficiency WLAN," IEEE Std 802.11ax-2021, 2021.

[4] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool," *Foundations and Trends in Networking*, vol. 12, no. 3, 2017.

[5] B. Wu, J. Huang, and S. Yu, ""X of Information" continuum: A survey on AI-driven multi-dimensional metrics for next-generation networked systems," *IEEE Communications Surveys & Tutorials*, vol. 28, pp. 5307–5344, 2026.

[6] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksal, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Transactions on Information Theory*, vol. 63, no. 11, 2017.

[7] M. Costa, M. Codreanu, and A. Ephremides, "On the age of information in status update systems with packet management," *IEEE Transactions on Information Theory*, vol. 62, no. 4, 2016.

[8] A. M. Bedewy, Y. Sun, and N. B. Shroff, "Minimizing the age of information through queues," *IEEE Transactions on Information Theory*, vol. 65, no. 8, 2019.

[9] R. D. Yates and S. K. Kaul, "The age of information: Real-time status updating by multiple sources," *IEEE Transactions on Information Theory*, vol. 65, no. 3, 2019.

[10] M. A. Abd-Elmagid, N. Pappas, and H. S. Dhillon, "On the role of age of information in the internet of things," *IEEE Communications Magazine*, vol. 57, no. 12, 2019.

[11] Z. Ding, B. Wu, and J. Huang, "SCALE: Sensitivity-aware federated unlearning with information freshness optimization for mobile edge computing," in *Proc. IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2026, pp. 1476–1486.

[12] B. Wu, J. Huang, and Q. Duan, "Real-time intelligent healthcare enabled by federated digital twins with AoI optimization," *IEEE Network*, vol. 40, no. 2, pp. 184–191, 2026.

[13] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, 2020.

[14] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, 2018.

[15] R. Talak, S. Karaman, and E. Modiano, "Optimizing information freshness in wireless networks under general interference constraints," *IEEE/ACM Transactions on Networking*, vol. 28, no. 1, 2020.

[16] Y.-P. Hsu, E. Modiano, and L. Duan, "Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals," *IEEE Transactions on Mobile Computing*, vol. 19, no. 12, 2020.

[17] I. Kadota and E. Modiano, "Minimizing the age of information in wireless networks with stochastic arrivals," *IEEE Transactions on Mobile Computing*, vol. 20, no. 3, 2021.

[18] V. Tripathi and E. Modiano, "A whittle index approach to minimizing functions of age of information," *IEEE/ACM Transactions on Networking*, 2024, dOI 10.1109/TNET.2024.3452006.

[19] R. D. Yates, "Lazy is timely: Status updates by an energy harvesting source," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, 2015.

[20] M. A. Abd-Elmagid and H. S. Dhillon, "Average peak age-of-information minimization in UAV-assisted IoT networks," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, 2019.

[21] E. T. Ceran, D. Gündüz, and A. György, "Average age of information with hybrid ARQ under a resource constraint," *IEEE Transactions on Wireless Communications*, vol. 18, no. 3, 2019.

[22] R. Talak, S. Karaman, and E. Modiano, "Distributed scheduling algorithms for optimizing information freshness in wireless networks," in *Proc. IEEE SPAWC*, 2018.

[23] A. M. Bedewy, Y. Sun, R. Singh, and N. B. Shroff, "Optimizing information freshness using low-power status updates via sleep-wake scheduling," in *Proc. ACM MobiHoc*, 2020.

[24] E. O. Gamgam, N. Akar, and S. Ulukus, "Minimizing age of information with generate-at-will status updates and age-agnostic cyclic scheduling," *arXiv preprint arXiv:2311.18791*, 2023.

[25] F. Chiariotti, J. Holm, A. E. Kalør, B. Soret, S. K. Jensen, T. B. Pedersen, and P. Popovski, "Query age of information: Freshness in pull-based communication," *IEEE Transactions on Communications*, vol. 70, no. 3, pp. 1606–1622, 2022.

[26] J. Liu and H. Chen, "Optimizing AoI at query in multiuser wireless uplink networks: A Whittle index approach," *arXiv preprint arXiv:2411.02108*, 2024.

[27] B. Bellalta, "IEEE 802.11ax: High-efficiency WLANs," *IEEE Wireless Communications*, vol. 23, no. 1, 2016.

[28] E. Khorov, A. Kiryanov, A. Lyakhov, and G. Bianchi, "A tutorial on IEEE 802.11ax high efficiency WLANs," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, 2019.

[29] D. López-Pérez, A. Garcia-Rodriguez, L. Galati-Giordano, M. Kasslin, and K. Doppler, "IEEE 802.11be extremely high throughput: The next generation of Wi-Fi technology beyond 802.11ax," *IEEE Communications Magazine*, vol. 57, no. 9, 2019.

[30] A. Garcia-Rodriguez, D. López-Pérez, L. Galati-Giordano, and G. Geraci, "IEEE 802.11be: Wi-Fi 7 strikes back," *IEEE Communications Magazine*, vol. 59, no. 4, 2021.

[31] M. Yang, B. Li, and Z. Yan, "Target wake time in IEEE 802.11 WLANs: Survey, challenges, and opportunities," *Computer Communications*, 2025.

[32] F. Busacca, C. Puligheddu, F. Raviglione, R. Rusca, C. Casetti, C. F. Chiasserini, and S. Palazzo, "Target wake time scheduling for time-sensitive and energy-efficient Wi-Fi networks," *IEEE Transactions on Mobile Computing*, vol. 25, no. 3, pp. 3469–3487, 2026.

[33] R. Roy, R. V. Bhat, P. Hathi, N. Akhtar, and N. M. Balasubramanya, "Maximization of timely throughput with target wake time in IEEE 802.11ax," in *Proc. IEEE International Conference on Communications (ICC)*, 2023, pp. 647–652.

[34] G. Rajendran and S. Agnihotri, "Deterministic scheduling over Wi-Fi 6 using target wake time: An experimental approach," *arXiv preprint arXiv:2505.00447*, 2025.

[35] S. K. Venkateswaran *et al.*, "IEEE 802.11ax target wake time: Design and performance analysis in ns-3," in *Proc. Workshop on ns-3 (WNS3)*, 2024, doi:10.1145/3659111.3659115.

[36] E. Mozaffariahrar, M. Menth, and S. Avallone, "Implementation and evaluation of IEEE 802.11ax target wake time feature in ns-3," in *Proc. Workshop on ns-3 (WNS3)*, 2024.

[37] M. Nurchis and B. Bellalta, "Target wake time: Scheduled access in IEEE 802.11ax WLANs," *IEEE Wireless Communications*, vol. 26, no. 2, 2019.

[38] H.-C. Lin, K.-H. Lin, and H.-Y. Wei, "Age of information for power-saving devices with DRX mechanism," in *Proc. IEEE International Conference on Communications (ICC)*, 2023, pp. 3234–3239.

[39] P.-H. Ting, C.-W. Tsai, W.-D. Shen, and H.-Y. Wei, "Information freshness and energy efficiency tradeoffs in IEEE 802.11: Age of information performance analysis of target wake time mechanism," *IEEE Transactions on Green Communications and Networking*, vol. 10, pp. 2778–2788, 2026.

[40] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the ACM*, vol. 20, no. 1, pp. 46–61, 1973.

[41] R. Holte, A. Mok, L. Rosier, I. Tulchinsky, and D. Varvel, "The pinwheel: A real-time scheduling problem," in *Proc. Hawaii Int. Conf. System Sciences (HICSS)*, 1989.

[42] C.-C. Han and K.-J. Lin, "Scheduling distance-constrained real-time tasks," in *Proc. IEEE Real-Time Systems Symposium (RTSS)*, 1992.

[43] M. Y. Chan and F. Chin, "General schedulers for the pinwheel problem based on double-integer reduction," *IEEE Transactions on Computers*, vol. 41, no. 6, 1992.