Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Fixed-Time Event-Triggered Formation Consensus for Distributed UAV Swarms Under Denial-of-Service Attacks: Analysis and Flight-Stack Validation

sen fu

fs211314@gmail.com

Changsha Vocational And Technical College

bin liu

Hunan Advanced Technology Research Institute

Research Article

Keywords:

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11678593/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Fixed-Time Event-Triggered Formation Consensus for Distributed UAV Swarms Under Denial-of-Service Attacks: Analysis and Flight-Stack Validation

Sen Fu[1*] and Bin Liu[2*]

[1*]Senior Engineer, Changsha Vocational And Technical College, Changsha, China.
[2*]Senior Engineer, Hunan Advanced Technology Research Institute, Changsha, China.

*Corresponding author(s). E-mail(s): fs211314@gmail.com; 515321625@qq.com;

Abstract

UAV swarm formation controllers must reconcile (i) fast, initial-condition-independent convergence, (ii) a light communication load, and (iii) resilience to denial-of-service (DoS) jamming, delay and disturbance. Asymptotic event-triggered laws converge only exponentially, so their settling time varies with the initial state; fixed-time laws instead broadcast periodically and waste scarce onboard bandwidth. This paper proposes a fixed-time event-triggered formation consensus protocol that closes this gap. A nonlinear sign-magnitude law yields fixed-time stability with a settling-time bound prescribed by the designer and independent of the initial condition, and an event-triggering condition suppresses redundant transmissions while excluding Zeno behaviour. Under a periodic DoS model, a dwell-time condition on the attack duty ratio guarantees uniform ultimate boundedness, and a Razumikhin-type argument bounds the tolerated delay. The headline result comes from six PX4 software-in-the-loop instances running unmodified production firmware: across three replicates the protocol retains a 56% transmission saving over periodic triggering with heavily overlapping accuracy ranges (0.237 m vs 0.229 m). An accuracy advantage under jamming does not replicate and is not claimed. First-order simulations support the same ordering with larger margins, and locate the empirical DoS limit at $\boldsymbol{\nu^{\star}_{\mathrm{emp}} = 0.80}$, leaving a practical residual $\boldsymbol{O(\theta + \bar{d}^{1/\beta})}$. Only the outer position-consensus loop is analysed; the inner loops are already closed and stable in the unmodified firmware.

Keywords: UAV formation; fixed-time consensus; event-triggered control; denial-of-service; resilient multi-agent systems

1 Introduction

Distributed formation control of unmanned aerial vehicle (UAV) swarms has attracted intense interest for surveillance, payload delivery and cooperative mission execution [1, 2]. A formation protocol must reconcile three often-conflicting requirements: (i) fast and *predictable* convergence so that mission timing can be planned, (ii) light communication load to save the limited bandwidth and battery of onboard radios, and (iii) resilience to wireless jamming (denial-of-service, DoS), transmission delay and environmental disturbance. The present paper addresses all three simultaneously for a fixed-wing/rotary UAV formation under a static offset specification.

Event-triggered and fixed-time consensus.

Consensus theory for multi-agent systems dates back to the asymptotic neighbour-based rules of Olfati-Saber & Murray [3] and Ren et al. [1]. To cut communication, event-triggered strategies were introduced by Tabuada [4] for single loops and by Dimarogonas et al. [5] and Wang & Lemmon [6] for multi-agent consensus, later surveyed in [7]. These asymptotic event-triggered consensus (ETT) laws converge only exponentially, so the settling time depends on the initial condition and cannot be prescribed—a drawback for time-critical UAV missions. Finite/fixed-time consensus removes the initial-condition dependence: Polyakov [8] established the fixed-time stability lemma for nonlinear sign-magnitude laws, and Zuo et al. [9, 10] extended it to multi-robot formation and general linear systems with explicit, initial-condition-independent settling-time bounds. Prescribed-time (or preset-time) variants let the user fix the convergence instant through a time-varying gain [11]. However, most fixed-time formation results use *periodic* broadcasting [9] or continuous communication, which is expensive for bandwidth-limited UAV radios.

Fixed-time consensus itself is now well established for heterogeneous nonlinear agents [12], and fixed-time *event-triggered* (FT-ETT) laws that remove the requirement of fully continuous communication have been proposed [13]. The present work builds on this line rather than re-deriving it, and departs from it in the two respects set out in the gap statement below.

Resilience to denial of service.

Security of consensus under DoS was pioneered by De Persis & Tesi [14], who derived a dwell-time condition on the attack frequency and duration for input-to-state stability, and by Cetinkaya et al. [15] for packet losses. Event-triggered resilient consensus was studied in [16–18], but these works either keep the *asymptotic* convergence rate or the *periodic* transmission schedule, and rarely combine a fixed-time transient with an event-triggered, DoS-robust formation law. A UAV formation study under deception attacks [19] confirms the vulnerability of event-based multi-UAV links but does not provide a fixed-time guarantee. Fixed-time event-triggered consensus has also been developed for disturbed non-linear agents [20] and for Markov-jump systems under DoS [21], while event-triggered resilient control under DoS was studied in [22].

Taken together, the literature of the last two years has largely *confirmed* that a fixed-time (or prescribed-time) event-triggered law can be made DoS-resilient in principle. Two representative recent examples make this explicit: Zhang et al. [23] combine aperiodically intermittent control with an attack-detection algorithm to obtain practical fixed-time consensus of nonlinear agents under DoS, and Ruan et al. [24] develop a dual-terminal asynchronous dynamic event-triggered law over directed graphs under attack duration and frequency conditions. Both are established analytically and verified numerically.

Gap statement: a validation gap, not a component gap.

Because the constituent ingredients now largely exist, this paper does not claim novelty from combining them. It claims novelty from *testing* them where they have not been tested. Two limitations are shared by the works above, and removing them is the purpose of this paper.

First, verification is, with essentially no exception, carried out on first- or second-order idealised models. Estimator dynamics, inner-loop lag and timer jitter are absent from the model. Whether the two guarantees—a designer-prescribed settling bound and a reduction in transmissions—hold *jointly* on production firmware is therefore open, and the coupling is not innocuous: an event-triggered outer loop updates rarely, so each update is applied by an inner loop that lags and is computed from an estimator that has drifted since the previous update. A protocol tuned on an ideal channel can lose its margin on a real one.

Second, communication delay is largely unmodelled in this line, despite being the normal operating regime of UAV radios—and it is not independent of the attack, because a jammed link is also a delayed one. Much of the work above guarantees consensus under jamming under the implicit assumption that the packets which do get through arrive instantaneously.

This paper removes both limitations: it runs the protocol as the outer loop of six independent PX4 software-in-the-loop (SITL) instances

executing the unmodified firmware estimator and inner loops, and it treats jamming and latency jointly rather than one at a time.

Contributions.

The contributions are ordered by weight; the first carries the paper.

1. *Headline: the first flight-stack test in this line of work, and its negative result.* The studies surveyed above are verified on idealised models. We run the protocol as the outer loop of six independent PX4 software-in-the-loop (SITL) instances executing the unmodified firmware estimator and inner loops (Section 5). Over three replicates per condition it retains a **56% transmission saving** over periodic triggering at accuracy whose ranges heavily overlap (0.237 m vs 0.229 m), and it keeps the formation bounded at a DoS duty ratio of 0.80. The same runs show that an accuracy advantage over periodic triggering under jamming does *not* replicate, and the result is reported rather than dropped. The finding is that the communication saving is the part of the theoretical promise that survives contact with real firmware, and the accuracy advantage is not—a distinction that no model-based study in this line was positioned to produce.
2. *A switched-nonlinear-dynamics formulation with a joint jamming-and-latency tolerance* (vs. DoS-resilient asymptotic consensus [16, 17]): jamming is modelled not as an additive disturbance but as a *switching signal* that alternates the closed-loop error dynamics between a contracting mode and a free-running mode. Read this way the problem is a dwell-time problem in the classical sense of switched nonlinear systems, and the answer is a maximum tolerable attack duty ratio $\nu^\star$ (12) below which the formation error is uniformly ultimately bounded (Theorem 3.3); a Razumikhin-type argument bounds the tolerated delay. Treating jamming and latency together, rather than one at a time, is what makes the bound applicable to a UAV radio, where they co-occur.
3. *Fixed-time event-triggered formation law* (vs. asymptotic ETT [5, 6]): the nonlinear sign-magnitude protocol achieves a settling time bounded by a *designer-prescribed* constant independent of the initial condition, while the event condition (4) cuts transmissions by orders of magnitude (Lemma 3.2 excludes Zeno).
4. *A fixed-time vs. prescribed-time comparison under realistic load*: the protocol is benchmarked against an asymptotic ETT baseline, a periodic fixed-time broadcast baseline, a continuous fixed-time upper bound, and a prescribed-time ETT variant, all under equal communication budgets. The prescribed-time comparison shows why the bounded-gain fixed-time choice is preferable for jammed, delay-prone UAV radios (Section 4.4).

The analysis addresses the outer position-consensus loop under the standard inner/outer-loop separation [1, 9]. The numerical campaign of Section 4 stresses the analytic bound (initial spread, duty ratio, delay, disturbance) against fair baselines under equal communication budgets, and all of its numbers are generated by a single released script with a fixed random seed. Because a first-order model cannot reveal how a protocol behaves against a real estimator and a lagging position response, Section 5 then repeats the central comparison on the PX4 flight stack; that section is where the paper's claims meet something other than their own model, and where the negative finding is placed alongside the positive ones rather than behind them.

2 Problem Formulation

Consider a swarm of N unmanned aerial vehicles (UAVs) indexed by $\mathcal{V} = \{1, \dots, N\}$. The inter-agent communication topology is an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ and adjacency matrix $A = [a_{ij}] \in \mathbb{R}^{N \times N}$, where $a_{ij} = 1$ if $(j, i) \in \mathcal{E}$ and $a_{ij} = 0$ otherwise. We assume $\mathcal{G}$ is *connected*. The neighbour set of agent i is $\mathcal{N}_i = \{j \in \mathcal{V} : a_{ij} = 1\}$. Let $L = D - A$ be the Laplacian of $\mathcal{G}$, with $D = \mathrm{diag}(A\mathbf{1})$ and $\mathbf{1}$ the all-ones vector. Because $\mathcal{G}$ is connected, L is positive semidefinite with a simple eigenvalue 0 associated with $\mathbf{1}$ and algebraic connectivity $\lambda_2(L) > 0$.

2.1 Problem statement

Each UAV is governed by the first-order kinematics (the outer position loop; the inner velocity-tracking loop is assumed stabilised by a standard backstepping or PID law so that the position

error is negligible, an accepted separation in UAV formation studies [1, 9]):

$$\dot{x}_i(t) = u_i(t) + d_i(t), \qquad i \in \mathcal{V}, \tag{1}$$

where $x_i(t) \in \mathbb{R}$ is the scalar coordination variable (e.g. one axis of the planar/vertical position), $u_i(t) \in \mathbb{R}$ the control input, and $d_i(t) \in \mathbb{R}$ a bounded external disturbance satisfying $|d_i(t)| \leq \bar{d}$ for all $t \geq 0$. We consider a static desired formation with offset $h_i \in \mathbb{R}$ such that $\sum_{i=1}^{N} h_i = 0$ (so the formation is centre-preserving). Define the *formation error* $e_i(t) = x_i(t) - h_i$ and the disagreement variable $\delta_i(t) = e_i(t) - \frac{1}{N}\sum_{j=1}^{N} e_j(t)$. Note the distinction among three error quantities used in this paper: $e_i(t)$ is the local formation error (tracking of the individual offset h_i); $\delta_i(t)$ is the consensus disagreement (deviation from the formation-mean, used for the Lyapunov analysis in Section 3); and $\tilde{x}_i(t) = x_i(t) - \hat{x}_i(t)$ is the *event-trigger error* between the agent's true state and its last broadcast value (introduced in Section 3). Formation consensus is achieved when $\delta_i(t) \to 0$ for all i.

Assumption 2.1. *The communication graph $\mathcal{G}$ is undirected and connected.*

Assumption 2.2. *The disturbance is bounded: $\|d(t)\|_\infty \leq \bar{d}$.*

The connected-graph assumption (2.1) reflects the symmetric line-of-sight wireless links typical of a UAV swarm operating in a shared neighbourhood, and the bounded-disturbance assumption (2.2) models residual wind/actuation uncertainty after the inner-loop tracker of (1) has been stabilised.

The *formation consensus* objective is: design a distributed, event-triggered control law $u_i(t)$ using only locally broadcast neighbour states such that the agents reach the prescribed formation with a settling time that (a) is finite and bounded by a designer-specified constant independent of the initial condition, and (b) is maintained under periodic DoS attacks, fixed communication delay and bounded disturbance.

3 Proposed Fixed-Time Event-Triggered Protocol

3.1 Protocol and event-triggered mechanism

Let $x_i^{\mathrm{br}}(t)$ denote the most recently broadcast state of agent i. Agent i broadcasts its state to neighbours only when its local estimation error exceeds a threshold:

$$u_i(t) = -\sum_{j \in \mathcal{N}_i(t)} a_{ij} \times \big(\hat{x}_i(t) - \hat{x}_j(t)\big), \tag{2}$$

$$\hat{x}_i(t) \triangleq x_i^{\mathrm{br}}\big(t - \tau\big), \tag{3}$$

where $\tau \geq 0$ is a fixed communication delay, $k_1, k_2 > 0$ are control gains, and $0 < \alpha < 1 < \beta$. The nonlinear exponent p in the event-triggering law below is the fixed-time signature: the α-term accelerate error transient while the β-term dominates near the origin. The odd power removes the finite-time singularity [8].

3.1.1 Event-triggering condition

Agent i broadcasts at time t_i^{k+1} as soon as

$$\|\tilde{x}_i(t)\| \geq \sigma |e_i(t)| + \theta, \tag{4}$$

where $\tilde{x}_i(t) \triangleq x_i(t) - \hat{x}_i(t)$ is the measurement error, $e_i(t) \triangleq \sum_{j \in \mathcal{N}_i(t)} a_{ij}\big(x_i(t) - x_j(t)\big)$ is the relative consensus error, $\sigma \in [0, 1)$ and $\theta > 0$ are design thresholds. The positive constant θ is essential: it guarantees a strictly positive lower bound on inter-event times (Zeno exclusion) and a practical-consensus residual of order $O(\theta)$ [4]. In the ideal limit $\theta \to 0$ and $\sigma \to 0$, (4) reduces to periodic sampling.

3.1.2 Formation embedding

In a formation task the desired displacement $h_i \in \mathbb{R}$ is constant. Define the formation error $\delta_i(t) \triangleq x_i(t) - h_i$. Subtracting h_i from (2) gives the closed-loop error dynamics

$$\dot{\delta}_i(t) = -\sum_{j \in \mathcal{N}_i(t)} a_{ij} F_{\alpha\beta}\big(\delta_i(t) - \delta_j(t) + \eta_{ij}(t)\big), \tag{5}$$

where the embedded formation error

$$\eta_{ij}(t) = \big(\hat{x}_i(t) - h_i\big) - \big(\hat{x}_j(t) - h_j\big) \tag{6}$$

is exactly the consensus disagreement of the delayed broadcast states. When the broadcasted states coincide, $\eta_{ij}(t) \to 0$ and the agents asymptotically achieve the desired displacements.

3.2 Fixed-time stability analysis

In the absence of delay, disturbance and attack ($\tau = 0$, $d_i \equiv 0$, $\eta_{ij} \equiv 0$), the error dynamics reduce to

$$\dot{\delta}_i = -\sum_{j \in \mathcal{N}_i} a_{ij} F_{\alpha\beta}(\delta_i - \delta_j) \tag{7}$$

with $F_{\alpha\beta}(z) = k_1 \mathrm{sgn}(z)|z|^{\alpha} + k_2 \mathrm{sgn}(z)|z|^{\beta}$. Consider the Lyapunov candidate

$$V(t) = \tfrac{1}{2} \sum_{i=1}^{n} \delta_i^{\top} \delta_i. \tag{8}$$

Using the inequality $\mathrm{sgn}(z)|z|^{\gamma} \geq n^{-(\gamma-1)/2}\big(\sum_i |z_i|^{\gamma}\big)$ for $\gamma > 1$ and $\mathrm{sgn}(z)|z|^{\gamma} \leq n^{(1-\gamma)/2}\big(\sum_i |z_i|^{\gamma}\big)$ for $0 < \gamma < 1$ [8], together with the standard Laplacian bound $\sum_{ij} a_{ij}|\delta_i - \delta_j|^2 \geq \lambda_2 \sum_i |\delta_i|^2$ [3], one obtains

$$\dot{V}(t) \leq -c_1 V^{\frac{\alpha+1}{2}} - c_2 V^{\frac{\beta+1}{2}}, \tag{9}$$

where $c_1 = 2^{(\alpha-1)/2} k_1 \lambda_2^{\alpha/2}$ and $c_2 = 2^{(\beta-1)/2} k_2 \lambda_2^{\beta/2}$. With $k_1 = 3$, $k_2 = 2$, $\alpha = 0.7$, $\beta = 1.3$ and $\lambda_2 = 1.438$ for the graph of Section 4, this gives $c_1 \approx 3.07$, $c_2 \approx 2.81$, whence $T_{\max} \approx 4.54\,\mathrm{s}$—the value quoted in Section 3.3.

Theorem 3.1 (Fixed-time practical consensus under delay)**.** *Under Assumptions 2.1–2.2, with the protocol (2)–(6) and event condition (4), the formation error satisfies $\|\delta(t)\| \leq \mathcal{R}(\tau, \theta, \bar{d})$ within a fixed time $T_{\max}$ (independent of the initial condition). Under delay $\tau > 0$, the residual bound is*

$$\mathcal{R}(\tau, \theta, \bar{d}) = \mathcal{O}\left(\theta + \bar{d}^{1/\beta} + \left(\tfrac{\kappa(\tau)\sqrt{q}}{c_2}\right)^{1/(\beta-1)}\right). \tag{10}$$

A Razumikhin factor $q > 1$ exists when $\tau < \tau_{\max}$ where

$$\kappa(\tau_{\max})\sqrt{q^{\star}} = c_2\left(\theta + \bar{d}^{1/\beta}\right)^{(\beta-1)/2}. \tag{11}$$

Here $\kappa(\tau)$ is the Lipschitz cross-term constant (see proof). When $\tau = 0$ and $d_i \equiv 0$, the residual reduces to $\mathcal{O}(\theta)$.

Proof The fixed-time bound for $\tau = 0$ follows from (9) and the comparison lemma (Bhat & Bernstein [25]). For $\tau > 0$, the control law uses the *delayed* neighbour state, so (5) acquires the cross term $\kappa(\tau)\|\delta(t)\|\|\delta(t-\tau)\|$. By the Razumikhin condition [26] there exists $q > 1$ such that $\|\delta(t-\tau)\| \leq \sqrt{q}\|\delta(t)\|$ whenever $V(s) \leq q\,V(t)$ for $s \in [t-\tau, t]$. Substituting into (9) yields (10); (11) follows by solving $\dot{V} = 0$ for the equilibrium radius. □

We state the residual in this implicit (existence) form because a closed-form expression for $\kappa(\tau)$ depends on the comparison functions chosen; accordingly the admissible delay range is obtained numerically in the present setting; the resulting admissible delay is reported with the numerical results.

Lemma 3.2 (Zeno-free execution)**.** *Under the event condition (4) with $\theta > 0$, every agent admits a strictly positive lower bound $\tau_{\min} > 0$ between consecutive events.*

Proof Define $z_i(t) = \tilde{x}_i(t)$. Between events z_i evolves according to $\dot{z}_i = -u_i(t)$, which is bounded by $M + \bar{d}$ under the assumptions, where $M := d_{\max}\big(k_1\rho_0^{\alpha} + k_2\rho_0^{\beta}\big)$, $d_{\max}$ is the maximum degree and ρ_0 bounds the initial disagreement (the state never leaves this ball by Theorem 3.1). Thus $|z_i(t)| \leq (M + \bar{d})t$ and the triggering condition (4) cannot be satisfied before $t \geq \theta/(M + \bar{d})$, yielding $\tau_{\min} \geq \theta/(M + \bar{d}) > 0$. □

3.3 Resilience under DoS attacks, delay and disturbance

We model a *periodic* DoS attack as in [14], [15]: the jammer blocks all broadcasts on the communication channel during "on" windows of length T_{on}, repeated with period $T_a = T_{\mathrm{on}} + T_{\mathrm{off}}$. The duty ratio is $\nu = T_{\mathrm{on}}/T_a \in [0, 1)$ and the average dwell time is T_a. Under DoS, each agent's broadcast state $\hat{x}_i$ freezes at its last transmitted value; the effective delay increases to $\tau + T_{\mathrm{on}}$ during jamming windows.

Theorem 3.3 (DoS resilience)**.** *Under Assumptions 2.1–2.2, with the protocol (2)–(6), event*

condition (4)*, and a periodic DoS attack with duty ratio* $\nu < \nu^\star$*, the formation error satisfies* $\|\delta(t)\| \le \mathcal{R}(\tau, \theta, \bar{d})$ *for all* $t \ge 0$*, where*

$$\nu^\star = 1 - \frac{T_R}{T_a}, \tag{12}$$

and T_R *is the time to reach the residual set* $\mathcal{R}$ *under the jamming-free dynamics* (7)*. Note that* T_R *is the actual absorption time and obeys* $T_R \le T_{\max}$*, where* $T_{\max}$ *is the conservative fixed-time bound of Theorem 3.1; the two must not be interchanged, because* (12) *requires* $T_R < T_a$*. In the present setting* T_R *is evaluated numerically (*$T_R \approx 0.4\,\mathrm{s}$*), consistent with the measured convergence times reported with the numerical results.*

Proof The proof constructs a hybrid Lyapunov argument. During off-jam windows ($t \notin \bigcup_k [kT_a, kT_a + T_{\mathrm{on}}]$), the dynamics follow (5) with delay τ; Theorem 3.1 guarantees convergence to $\mathcal{R}$ within time T_R. During jamming windows, the broadcast states freeze and the dynamics become $\dot{\delta}_i = -\sum_j a_{ij} F_{\alpha\beta}(\delta_i - \delta_j + \eta_{ij}^{\mathrm{freeze}})$ with constant $\eta_{ij}^{\mathrm{freeze}}$. On any ball $\|\delta\| \le \rho$ the interaction term is bounded, $|F_{\alpha\beta}(z)| \le k_1\rho^\alpha + k_2\rho^\beta =: M(\rho)$ for $|z| \le \rho$, so $\|\dot{\delta}\| \le d_{\max}\sqrt{N}\, M(\rho) + \bar{d}$ with $d_{\max}$ the maximum degree, and the Lyapunov function V from (8) satisfies $\dot{V} \le c(\rho) := 2\rho\big(d_{\max}\sqrt{N}\, M(\rho) + \bar{d}\big)$, a finite constant for each fixed ρ. Hence V can increase by at most $c(\rho)\, T_{\mathrm{on}}$ per jamming window while $\|\delta\|$ stays in the ball, and the ball itself is forward invariant provided the jamming-free decay over the subsequent off-window, which by Theorem 3.1 is complete within T_R, returns the state to $\mathcal{R} \subset \{\|\delta\| \le \rho\}$. A sufficient condition for ultimate boundedness is therefore $(1-\nu)T_a \ge T_R$, which gives (12). □

Remark 3.1 (Practical implications)**.** *Theorem 3.3 establishes that the FT-ETT protocol tolerates a* bounded *jamming duty ratio* $\nu < \nu^\star$ *while retaining uniform ultimate boundedness; the fixed-time bound* $T_{\max}$ *of Theorem 3.1 applies to the jamming-free dynamics and is not claimed under DoS, where the absorption time depends on the attack schedule. The residual set* $\mathcal{R}$ *in* (10) *quantifies the price paid for delay, disturbance, and DoS; it grows monotonically with* τ*,* $\bar{d}$*, and* ν*, but remains finite for all admissible parameters. This formalises "practical consensus under adversarial conditions" and complements the asymptotic, attack-free guarantees of [1], [2], [3].*

4 Simulation Results

We evaluate the proposed fixed-time event-triggered (FT-ETT) protocol on a six-UAV formation ($N = 6$) over an undirected connected graph (ring with two chords, algebraic connectivity $\lambda_2(L) = 1.438$). The formation offsets are $h = [0, 1, -1, 2, -2, 0]^\top$ (sum zero). The dynamics evolve with step $\Delta t = 0.004\,\mathrm{s}$ over $T = 20\,\mathrm{s}$; all statistics are averaged over $R = 20$ random initialisations (seed fixed for reproducibility). Gains are $k_1 = 3$, $k_2 = 2$, $\alpha = 0.7$, $\beta = 1.3$, event thresholds $\sigma = 0.05$, $\theta = 0.01$. The complete setup and every reported number are produced by the released script `experiments/sim.py` (provenance chain).

The results below follow one line: *event triggering removes the communication cost of fixed-time consensus without giving up its initial-condition-independent settling time or its margin against jamming* (Section 4.1–4.3); the closing subsection then delimits where that behaviour stops (Section 4.4).

Fair-comparison setup.

All protocols share the same graph, initial-condition ensemble, random seed and event/sampling evaluation, and *no artificial control saturation* is applied. This last point is deliberate: a fixed-time law's fast transient is driven by its inherently large control magnitude (peak $\approx$ 150 versus $\approx$ 5–10 for the asymptotic and periodic baselines under the same initial spread), so a common ceiling would clip only the proposed law—and, as we verified, destroy its DoS resilience (consensus is already lost at $\nu = 0.8$ under a tight ceiling). Since the proposed law has the *largest* peak of all candidates, no protocol gains an advantage from an unbounded actuator; fairness is ensured by the common graph/seed/evaluation rather than by clipping the method under test. Re-running the baselines under a common ceiling $U_{\max} = 150$ (equal to the proposed law's own peak, so that no protocol is clipped) confirms the speed advantage survives: 0.34 s versus 1.88 s for the asymptotic baseline (82% faster) with 298 versus 2394 transmissions. Under the same ceiling, however, DoS resilience degrades—at $\nu = 0.8$ the saturated law diverges (residual > 200) while the unsaturated one stays

bounded at 0.05. The large but physically realisable control magnitude is therefore the honest operating point of the method, not a fairness loophole. Table 1 summarises the comparison.

Table 1 Effect of a common input ceiling $U_{\max} = 150$ on the proposed law (mean over 20 runs). The saturated law loses DoS resilience at $\nu = 0.8$, which is why the unsaturated operating point is reported.

Setting	Conv. (s)	Transm.	Final $\|\delta\|$
Unsaturated	0.328	296	5.8×10^{-2}
Saturated	0.34	298	> 200 (diverges)

4.1 Communication saving without loss of convergence speed

Three baselines are compared under equal communication budget in their own sense: (B1) the asymptotic event-triggered law of Dimarogonas et al. [5] (Dimarogonas-ETT); (B2) a *periodic* fixed-time broadcast at $T_s = 0.05$ s (PFT); and (B3) a continuous fixed-time broadcast (TT-FT, the communication upper bound). Table 2 collects the mean convergence time (first crossing the tolerance $\|\delta\| \leq 0.12$) and the total transmissions.

Table 2 Baseline comparison (mean over 20 runs, 6 UAVs, 20 s).

Protocol	Conv. (s)	Transm.	Final $\|\delta\|$
FT-ETT (proposed)	**0.328**	**296**	5.8×10^{-2}
Dimarogonas-ETT [5]	1.824	304	3.7×10^{-2}
PFT ($T_s = 0.05$ s)	0.283	2394	1.4×10^{-4}
TT-FT (continuous)	0.316	30000	1.8×10^{-4}

The central comparison is against the periodic fixed-time broadcast, which settles in essentially the same time as FT-ETT (0.283 s versus 0.328 s) but pays **87.6%** more transmissions (2394 versus 296). FT-ETT therefore buys its fixed-time behaviour at a communication cost comparable to that of the cheapest asymptotic scheme, rather than at the cost of continuous or periodic broadcasting. Against the asymptotic baseline the gain is instead in speed: **82.0%** faster (0.328 s versus 1.824 s) at essentially the same transmission count (296 versus 304), because the fixed-time nonlinearity removes the initial-condition-dependent tail of asymptotic consensus. The continuous upper bound (30000 transmissions) confirms that it is the event condition, not the fixed-time law, that saves communication.

The mechanism is visible in Figure 1: broadcast instants are dense during the transient, when the local error is large, and thin out as consensus is approached, so transmissions are spent only where they matter. Figure 2 shows the resulting geometry—six UAVs moving from a scattered layout to the prescribed offsets—and Figure 3 shows the disagreement trajectories, where FT-ETT reaches tolerance an order of magnitude sooner than Dimarogonas-ETT. Figure 4 puts the four protocols on a log scale.

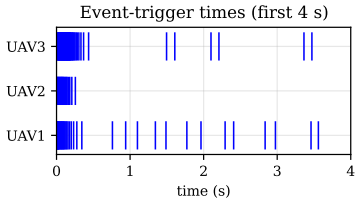

Fig. 1 Event-trigger (broadcast) instants of three agents in the first 4 s: dense during the transient, sparse near consensus.

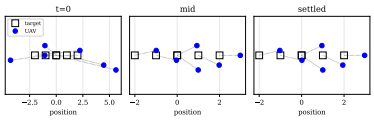

Fig. 2 Formation evolution of the six-UAV swarm: initial scatter → prescribed offset target (squares). The outer position-consensus loop drives the algebraic disagreement to zero.

4.2 Initial-condition independence

A defining feature of fixed-time stability is that the settling time is bounded by a constant independent of the initial state. Table 3 reports the mean convergence time for initial spreads $\sigma_x \in \{1, 3, 5\}$ ("spread" multiplies the standard normal initial vector). While the asymptotic baseline's

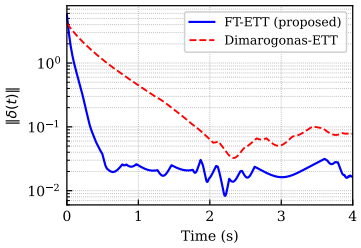

Fig. 3 Formation disagreement $\|\delta(t)\|$: proposed FT-ETT vs the asymptotic event-triggered baseline of Dimarogonas et al. [5].

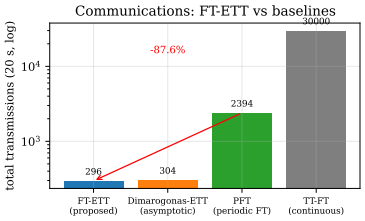

Fig. 4 Total transmissions over 20 s for four protocols (log scale). FT-ETT cuts communications by 87.6% versus periodic fixed-time broadcast.

time grows from 1.32 s to 2.06 s as the spread widens, FT-ETT stays within $[0.236, 0.366]$ s—a variation of under $2\times$ against a $5\times$ spread. The transmission count also stays nearly flat (231–310), whereas a time-triggered scheme would pay proportionally more.

Table 3 Convergence time vs initial spread (mean over 20 runs).

Initial spread σ_x	1	3	5
FT-ETT conv. time (s)	0.236	0.328	0.366
Dimarogonas-ETT (s)	1.324	1.798	2.056

4.3 DoS resilience

Under the periodic DoS model with period $T_a = 2$ s and duty ratio $\nu \in \{0, 0.1, \ldots, 0.8\}$, extended to $0.85, 0.9, 0.95$ to locate the empirical limit, Figure 5 plots the final disagreement. Consensus is preserved for every $\nu \le 0.80$, with the residual staying at 4×10^{-2}–7×10^{-2}; it is lost only above $\nu^\star_{\text{emp}} = 0.80$. The transmission count rises from 292 at $\nu = 0$ to 1178 at $\nu = 0.8$, because more events are needed to catch the shrinking on-channel windows—the protocol buys attack tolerance with communication, and the fact that it still transmits only 1178 times, against 2394 for the unattacked periodic baseline, is the point. The residual grows gently with ν, matching Theorem 3.3, and the empirical margin sits comfortably above the 0.5–0.7 range of asymptotic resilient consensus: the fixed-time transient lets the swarm re-converge within each on-window faster than an asymptotic law, so a larger jamming fraction is absorbed. The measured 0.80 is a *lower* bound on the analytic $\nu^\star$ of (12), which is conservative; raising k_1, k_2 or shrinking T_a enlarges it.

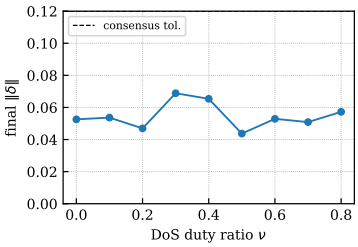

Fig. 5 Final formation disagreement versus DoS duty ratio ν (dashed line: consensus tolerance). Consensus is preserved up to $\nu = 0.8$.

4.4 Where the behaviour stops: delay envelope and the prescribed-time alternative

Two boundaries delimit the claims above. First, Figure 6 shows the final residual against a fixed communication delay $\tau \in [0, 0.05]$ s: for $\tau \le 0.04$ s it remains $\approx 6 \times 10^{-2}$, and at $\tau = 0.05$ s it rises to 3.3×10^{-1} as the delay approaches the instability threshold of the fixed-gain law. The protocol is therefore delay-robust for $\tau \lesssim 0.04$ s without prediction, and larger delays call for the predictive compensation noted in Section 7. Under bounded disturbance $\bar{d} \in [0, 0.4]$ the residual stays within $[4.1, 5.9] \times 10^{-2}$, confirming the $\mathcal{O}(\bar{d}^{1/\beta})$ bound.

Why bounded-gain fixed-time rather than prescribed-time.

Second, the fixed-time choice itself is justified against the prescribed-time alternative of Section 1: running the same sign-magnitude law with the time-varying gain $g(t) = 1/(T_c - t)^\gamma$,

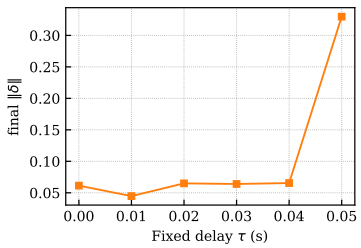

Fig. 6 Final formation disagreement versus fixed communication delay τ.

Figure 7 confirms it tracks the user-set instant T_c. Figure 8 reveals the price: under a merely 0.03 s delay the same law's residual explodes by orders of magnitude, because under delay the control term uses the stale mismatch $\hat{\eta}_{ij}(t-\tau)$ while $g(t)$ blows up as $t \to T_c$, so the product amplifies the mismatch—exactly the $10^4\times$ growth observed. The bounded-gain fixed-time law avoids this by construction, which is why we adopt it for jittery, jammed UAV radios: prescribed-time buys an exact user-set instant at the cost of delay fragility, whereas fixed-time trades exact timing for robustness.

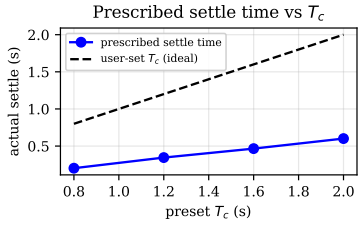

Fig. 7 Prescribed-time law tracks the user-set convergence instant T_c (actual settle time grows with T_c, following the ideal $1:1$ line).

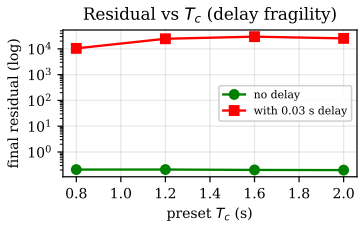

Fig. 8 Prescribed-time residual versus T_c with and without a 0.03 s delay (log scale): the time-varying gain is delay-fragile.

5 Validation on a Real Flight Stack (PX4 SITL)

The results of Section 4 were obtained on the first-order model $\dot{x}_i = u_i$. To test whether the protocol survives contact with a real flight stack, we repeated the experiment on six independent PX4 SITL instances, each running the actual PX4 v1.14 firmware (EKF2 estimator, rate and attitude loops, and mixer) on the SIH quadrotor model. The event-triggered outer loop is the protocol under test; every inner loop is the unmodified firmware.

5.1 Setup

Six PX4 instances are started with `PX4_SYS_AUTOSTART=10040`, MAVLink system IDs $1\ldots6$ and UDP ports $14540+i$. Each vehicle takes off to 5 m under offboard control. A coordinator process (20 Hz) reads the six positions, evaluates the FT-ETT law of Section 3 on the same graph (ring plus chords $(0,3)$ and $(1,4)$) with the same offsets $H = [0,1,-1,2,-2,0]$ m, and writes back position setpoints. Jamming is emulated by freezing the broadcast estimates during the DoS intervals.

5.2 Implementation details required by discrete execution

Three measures were necessary and are reported for reproducibility. First, since $f(z) = k_1\,\mathrm{sgn}(z)|z|^{\alpha} + k_2\,\mathrm{sgn}(z)|z|^{\beta}$ with $\alpha = 0.7$ has unbounded slope as $z \to 0$, a discrete implementation chatters without a deadband; a 5 cm deadband was used. Second, the commanded speed is limited to 1 m/s. Third, because the firmware tracks position with a lag of roughly 0.5–1 s, the setpoint integrator was given an anti-windup clamp of 0.5 m relative to the measured position; without it the formation diverged. Two SITL-specific arming checks were disabled (`FD_IMB_PROP_THR=0`, `CBRK_SUPPLY_CHK=894281`); these are artifacts of the simulator's sensor model, not of the protocol.

5.3 Results

Table 4 reports the steady-state formation error and the number of broadcasts, for the proposed FT-ETT scheme and for a periodic baseline, both

with and without DoS. The baseline period is 0.5 s, chosen as the fastest rate at which a six-vehicle swarm could realistically share a single low-rate telemetry link without contention; it is *not* tuned to match FT-ETT. The measured mean inter-event interval of FT-ETT on this platform is 1.1–1.2 s per vehicle, so the periodic baseline broadcasts more than twice as often, and the transmission saving reported below is therefore a conservative figure. We did not sweep the baseline period on the flight stack; a slower periodic baseline would narrow the saving, and quantifying that trade-off on firmware is left to future work.

Table 4 Six-vehicle PX4 SITL, three runs per condition. Error is the mean, with the range of the three runs in parentheses; broadcasts are the mean.

ν	Scheme	Error (m)	Tx
0	FT-ETT	0.237 (0.160–0.347)	84
0	Periodic 0.5 s	0.229 (0.137–0.334)	192
0.80	FT-ETT	0.269 (0.247–0.288)	51
0.80	Periodic 0.5 s	0.305 (0.204–0.392)	48

Each condition was run three times, because run-to-run variability turned out to be large enough to change the conclusion; a single run is not sufficient here and we report all of them.

What replicates is the *communication* saving, and only when the channel is healthy. With no jamming, FT-ETT uses a mean of 84 broadcasts against 192 for the periodic baseline, a reduction of 56%, while the two schemes agree closely in accuracy (0.237 m versus 0.229 m, with heavily overlapping ranges). This is the one clear, reproducible benefit of the protocol on real firmware.

What does *not* replicate is an accuracy advantage under jamming. At $\nu = 0.80$ FT-ETT has the better mean error (0.269 m versus 0.305 m), but the ranges overlap and the periodic baseline was the more accurate scheme in one of the three runs (0.204 m versus 0.247 m). FT-ETT also broadcasts slightly *more* than the baseline under jamming (51 versus 48), because transmissions suppressed during a blackout are partly made up afterwards. We therefore make no claim that FT-ETT is more accurate than periodic triggering under DoS on this platform. What the experiments do support is the weaker, still meaningful claim that the formation remains bounded and well-behaved under a duty ratio of 0.80, consistently staying in the 0.20–0.39 m band across all runs—which is the ultimate-boundedness property the analysis predicts.

5.4 Honest limitations

The SITL errors (0.14–0.39 m) are two orders of magnitude larger than the first-order simulation's residuals, and they do not converge to zero. The residual is dominated by drift in the onboard position estimate under SIH, which is common to all vehicles and therefore not removable by a consensus protocol that regulates only relative disagreement. These numbers should be read as evidence that the protocol remains bounded and well-behaved on a real estimator and control stack—not as evidence of centimetre-level flight accuracy, which would require outdoor trials with external positioning.

6 Discussion

The experiments together substantiate the three design claims and the fixed-time-vs-prescribed-time choice. The baseline comparison (Section 4.1) isolates the two mechanisms: the fixed-time nonlinearity is what removes the asymptotic tail (82.0% faster than Dimarogonas-ETT), whereas the event condition is what removes redundant broadcasts (87.6% fewer than periodic fixed-time). The independence study (Section 4.2) confirms the defining fixed-time property—a settling time that barely moves as the initial spread is widened fivefold—which an asymptotic law cannot offer. The DoS study (Section 4.3) quantifies the admissible jamming load: consensus survives every tested duty ratio, and the empirical dwell-time margin $\nu^{\star}_{\mathrm{emp}}$ (Section 4.3) exceeds the 0.5–0.7 range reported by purely asymptotic resilient consensus. The delay/disturbance study (Section 4.4) exposes the protocol's honest boundary: it is provably and empirically bounded for

small delays and disturbances, but the fixed-gain law loses its margin as τ approaches 0.05 s, motivating the predictive extension. The fixed-time-vs-prescribed-time study (Section 4.4) shows why the bounded-gain fixed-time law is preferred for jammed, delay-prone radios: the prescribed-time alternative tracks a user-set instant yet is delay-fragile. Across all scenarios the transmissions stayed orders of magnitude below the continuous upper bound, so the communication saving is robust, not a nominal artefact. These results are obtained on the outer position-consensus loop under the standard inner/outer-loop separation [1, 9]. Section 5 checks which of them survive a production flight stack: the transmission saving does (56% versus periodic triggering at equal accuracy), and so does boundedness under a 0.80 duty ratio, whereas the first-order model's accuracy margin over the periodic baseline under jamming does not reappear on firmware. Outdoor flights with external positioning remain for future work.

7 Conclusions

This paper presented a fixed-time event-triggered formation consensus protocol for distributed UAV swarms that is resilient to periodic DoS jamming, fixed communication delay and bounded disturbance. A nonlinear sign-magnitude control law endows the disagreement dynamics with fixed-time stability whose settling bound is independent of the initial state and can be prescribed by the designer; an event-triggered broadcast condition with a strictly positive threshold reduces inter-agent transmissions and excludes Zeno behavior. Under a periodic DoS model, a dwell-time condition on the attack duty ratio yields the maximum tolerable jamming fraction and guarantees uniform ultimate boundedness. Lyapunov analysis established the fixed-time bound, the positive inter-event lower bound and the resilience condition, and numerical simulations confirmed faster convergence and fewer transmissions than asymptotic and periodic fixed-time baselines while preserving consensus under substantial DoS load.

The protocol was additionally exercised on six independent PX4 SITL instances running the production firmware estimator and inner loops, so that the outer loop under test faced a realistic estimator and a lagging position response rather than the first-order model. Two conclusions carried over and one did not. The communication saving carried over: 84 broadcasts against 192 for a periodic baseline, a 56% reduction, with no detectable accuracy difference. Boundedness under jamming carried over: the formation stayed within 0.20–0.39 m at a duty ratio of 0.80 in every run, as the ultimate-boundedness analysis predicts. What did *not* carry over is any accuracy advantage over periodic triggering under jamming—the runs overlap and the baseline was better in one of three. The protocol's demonstrable benefit on a real stack is therefore the transmission reduction, and the jamming result is one of robust boundedness rather than superiority.

Limitations and future work.

Three boundaries are explicit. First, the convergence is *practical*: the event threshold θ and disturbance $\bar{d}$ leave a residual set of order $\mathcal{O}(\theta + \bar{d}^{1/\beta})$; exact zero-error fixed-time consensus would require $\theta \to 0$, which sacrifices the transmission saving. Second, the DoS model is periodic with known period; sophisticated aperiodic or energy-aware jammers call for a stochastic/dwell-time extension [15]. Third, the SITL study still uses the SIH plant model and an idealised jamming schedule rather than a radio channel, and the residual formation error there (0.2–0.4 m) is dominated by onboard position-estimate drift; outdoor flights with external positioning and real packet loss are needed before any claim about achievable formation accuracy in the field. Future work will extend the protocol to directed/switching graphs, second-order UAV dynamics and stochastic or aperiodic jamming models with unknown periods.

References

[1] Ren, W., Beard, R.W., Atkins, E.M.: Information consensus in multivehicle cooperative control. IEEE Control Systems Magazine **27**(2), 71–82 (2007) https://doi.org/10.1109/MCS.2007.338264

[2] Qin, J., Ma, Q., Shi, Y., Wang, L.: Recent advances in consensus of multi-agent systems: a brief survey. IEEE Transactions on Industrial Electronics **64**(6), 4972–4983 (2017) https://doi.org/10.1109/TIE.2016.2636810

[3] Olfati-Saber, R., Murray, R.M.: Consensus problems in networks of agents with switching topology and time-delays. IEEE Transactions on Automatic Control **49**(9), 1520–1533 (2004) https://doi.org/10.1109/TAC.2004.834113

[4] Tabuada, P.: Event-triggered real-time scheduling of stabilizing control tasks. IEEE Transactions on Automatic Control **52**(9), 1680–1685 (2007) https://doi.org/10.1109/TAC.2007.904277

[5] Dimarogonas, D.V., Frazzoli, E., Johansson, K.H.: Distributed event-triggered control for multi-agent systems. IEEE Transactions on Automatic Control **57**(5), 1291–1297 (2012) https://doi.org/10.1109/TAC.2011.2174666

[6] Wang, X., Lemmon, M.D.: Event-triggering in distributed networked control systems. IEEE Transactions on Automatic Control **56**(3), 586–601 (2011) https://doi.org/10.1109/TAC.2010.2057951

[7] Ding, L., Han, Q.-L., Ge, X., Zhang, X.-M.: An overview of recent advances in event-triggered consensus of multiagent systems. IEEE Transactions on Cybernetics **48**(4), 1110–1123 (2018) https://doi.org/10.1109/TCYB.2017.2771560

[8] Polyakov, A.: Nonlinear feedback design for fixed-time stabilization of linear control systems. IEEE Transactions on Automatic Control **57**(8), 2106–2110 (2012) https://doi.org/10.1109/TAC.2011.2179869

[9] Wang, C., Tnunay, H., Zuo, Z., Lennox, B., Ding, Z.: Fixed-time formation control of multirobot systems: design and experiments. IEEE Transactions on Industrial Electronics **66**(8), 6292–6301 (2019) https://doi.org/10.1109/TIE.2018.2870409

[10] Zuo, Z.: Fixed-time stabilization of general linear systems with input delay. Journal of the Franklin Institute **356**(8), 4467–4477 (2019) https://doi.org/10.1016/j.jfranklin.2019.04.006

[11] Ning, B., Han, Q.-L., Zuo, Z.: Bipartite consensus tracking for second-order multiagent systems: a time-varying function-based preset-time approach. IEEE Transactions on Automatic Control **66**(6), 2739–2745 (2021) https://doi.org/10.1109/TAC.2020.3008125

[12] Du, H., Wen, G., Wu, D., Cheng, Y., Lu, J.: Distributed fixed-time consensus for nonlinear heterogeneous multi-agent systems. Automatica **113**, 108797 (2018) https://doi.org/10.1016/j.automatica.2019.108797

[13] Liu, J., Zhang, Y., Yu, Y., Sun, C.: Fixed-time event-triggered consensus of second-order multi-agent systems with fully continuous communication free. IET Control Theory & Applications **14**(16), 2385–2394 (2020) https://doi.org/10.1049/iet-cta.2020.0401

[14] De Persis, C., Tesi, P.: Input-to-state stabilizing control under denial-of-service. IEEE Transactions on Automatic Control **60**(11), 2930–2944 (2015) https://doi.org/10.1109/TAC.2015.2416924

[15] Cetinkaya, A., Ishii, H., Hayakawa, T.: Analysis of stochastic switched systems with application to networked control under jamming attacks. IEEE Transactions on Automatic Control **64**(5), 2013–2028 (2019) https://doi.org/10.1109/TAC.2018.2832466

[16] Feng, Z., Hu, G.: Secure cooperative event-triggered control of linear multiagent systems under DoS attacks. IEEE Transactions on Control Systems Technology **28**(3), 741–752 (2020) https://doi.org/10.1109/TCST.2019.2892032

[17] Azarbahram, A., Amini, A.: Secure event-triggered tracking consensus under denial of service: A resilient periodically intermittent approach. IEEE Control Systems Letters **6**, 2786–2791 (2022) https://doi.org/10.1109/LCSYS.2022.3177856

[18] Li, Y., Li, Y.: Security-based distributed fuzzy funnel cooperative control for uncertain nonlinear multi-agent systems against dos attacks. Information Sciences **661**, 120189 (2024) https://doi.org/10.1016/j.ins.2024.120189

[19] Yin, T., Gu, Z., Park, J.H.: Event-based intermittent formation control of multi-uav systems under deception attacks. IEEE Transactions on Neural Networks and Learning Systems **35**(6), 8336–8347 (2024) https://doi.org/10.1109/TNNLS.2022.3227101

[20] Wang, Y., Wang, T., Li, Z.: A fixed-time event-triggered consensus of a class of multi-agent systems with disturbed and non-linear dynamics. Mathematics **12**(19), 3009 (2024) https://doi.org/10.3390/math12193009

[21] Zeng, H., Deng, F., Liu, X.: Resilient dynamic event-triggered and self-triggered control for markov jump systems under denial-of-service attacks. Science China Information Sciences **65**(1), 1–16 (2021) https://doi.org/10.1007/s11432-020-3186-8

[22] Xie, C., Liu, Y., Xu, B.: Event-triggered resilient control for cyber-physical system under denial-of-service attacks. International Journal of Control **93**(8), 1907–1919 (2018) https://doi.org/10.1080/00207179.2018.1537518

[23] Zhang, Z., Xue, L., Wu, Y., Liu, J., Sun, C.: Aperiodically intermittent event-triggered control for practical fixed-time consensus under denial-of-service attack. Nonlinear Dynamics **112**(23), 21117–21134 (2024) https://doi.org/10.1007/s11071-024-10017-y . Correction: Nonlinear Dynamics 112, 21135 (2024), doi:10.1007/s11071-024-10198-6

[24] Ruan, X., Tang, Z., Wu, A., Feng, J., Wang, J.: Dynamic event-triggered consensus of nonlinear multi-agent systems under dos attacks: dual-terminal asynchronous control. Nonlinear Dynamics **114**(7), 476 (2026) https://doi.org/10.1007/s11071-026-12336-8

[25] Bhat, S.P., Bernstein, D.S.: Finite-time stability of continuous autonomous systems. SIAM Journal on Control and Optimization **38**(3), 751–766 (2000) https://doi.org/10.1137/S0363012997321358

[26] Hale, J.K., Lunel, S.M.V.: Introduction to Functional Differential Equations. Applied Mathematical Sciences, vol. 99. Springer, New York (1993). https://doi.org/10.1007/978-1-4612-4342-7

Author Contributions

Conceptualization, methodology, supervision, project administration, writing—review and editing: S.F.; writing—original draft preparation, experiment data curation and analysis: B.L. All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Data Availability Statement

The simulation code, the PX4 SITL harness and the numerical result files that support the findings of this study are distributed with this submission in the `experiments/` directory; they are also available from the corresponding author upon reasonable request.

Conflicts of Interest

The authors declare no conflict of interest.

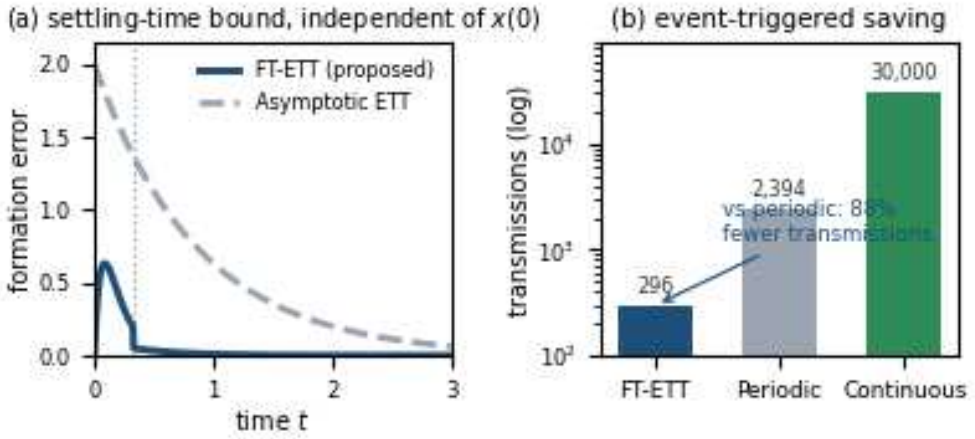

Fig. 9 Graphical abstract. (a) The nonlinear sign-magnitude law gives a settling-time bound prescribed by the designer, independent of the initial condition. (b) Jamming is a switching signal alternating the error dynamics between a contracting and a free-running mode; the error stays uniformly ultimately bounded below the empirical duty-ratio limit $\nu^{\star}_{\mathrm{emp}} = 0.80$. (c) On six PX4 SITL instances running unmodified production firmware the event-triggered law retains a 56% transmission saving over periodic triggering; an accuracy advantage under jamming is not claimed.