Aerospace Engineering Communications
http://dx.doi.org/10.62762/AEC.2026.732691

RESEARCH ARTICLE

Event-Triggered Attitude Control of Underactuated Spacecraft

Chengxi Zhang [1,*]

[1] School of Internet of Things Engineering, Jiangnan University, Wuxi 214122, China

Abstract

This paper presents an event-triggered attitude control method for an underactuated spacecraft with only two bounded body-axis torques. Gyroscopic coupling supplies the missing control action: two periodic angular-rate references generate the average gyroscopic rate product required to regulate the unactuated axis. A relative-plus-absolute event condition updates the bounded torque command, while a zero-order hold and a timeout provide implementable inter-event intervals. An elementary energy argument establishes asymptotic stability of the explicitly defined nominal averaged local error model, while a nonlinear simulation assesses the complete implementation. Under a 2 N·m torque limit, an initial attitude error of 32.6° is reduced to 0.108° at 200 s. The event implementation performs 295 full control-law evaluations, and its completed inter-update intervals range from 0.2 to 1.0 s. This corresponds to 70.5% fewer full control-law evaluations and torque-command updates than a 5 Hz periodic implementation.

Keywords: underactuated spacecraft, attitude control, event-triggered control.

Submitted: 01 January 2026
Revised: 26 August 2026
Accepted: 15 September 2026
Published: 21 September 2026

Vol. 1, **No.** 3, 2026.
10.62762/AEC.2026.732691

***Corresponding author:**
Chengxi Zhang
dongfangxy@163.com

1 Introduction

Spacecraft attitude control is required whenever a vehicle slews toward a commanded orientation, points an antenna, or maintains an instrument line of sight; an overview of the underlying determination and control foundations is given in [1]. Most attitude controllers assume three independent body-axis torques. This assumption can fail after the loss of a reaction wheel, control-moment gyro, or thruster channel. If no redundant actuator is available, the spacecraft becomes underactuated: two commands must control three rotational degrees of freedom. A conventional three-axis proportional–derivative controller can no longer be applied directly.

The absence of one torque does not mean that all motion about that axis is impossible. Euler's rigid-body equations contain gyroscopic coupling. When the principal moments of inertia are different, simultaneous angular motion about two axes creates acceleration about the third axis. Krishnan, Reyhanoglu, and McClamroch established an early nonlinear two-torque attitude control result [2]; more recent work has considered velocity-free two-torque feedback [3]. Standard rigid-spacecraft dynamics and the physical origin of this coupling are detailed in [4]. These studies show why the nonlinear model

Citation
Zhang, C. (2026). Event-Triggered Attitude Control of Underactuated Spacecraft. *Aerospace Engineering Communications*, 1(3), 110–118.

matters: the linearized model at rest hides precisely the coupling needed to recover the missing control direction. The present paper uses the same physical mechanism but packages it as a simple pair of periodic rate references that can be explained using the average of $\sin^2(\nu t)$.

Constant-command attitude tracking also requires a representation of orientation error. Unit quaternions are convenient because they avoid the singularities of Euler angles and are inexpensive to propagate onboard. Their double-cover property must still be handled by selecting the shorter quaternion representative.

More sophisticated rotation-matrix and hybrid quaternion methods can provide strong almost-global or global stability properties, respectively [5, 6]; here we intentionally restrict attention to the local command-tracking region so that the controller and its stability explanation remain accessible.

Actuator limits are equally important. Fault-tolerant spacecraft control and actuator-loss accommodation have been surveyed in [7]. A controller that works only by asking for tens of newton-meters may be mathematically interesting but unusable on a small or medium spacecraft. Practical studies therefore consider saturation, uncertainty, actuator failures, and quantized torque [8–11]. Finite-time convergence designs provide an alternative route that achieves fast attitude response under disturbances and inertia uncertainty at the cost of greater analytical complexity [12]. Attitude tracking of underactuated spacecraft using internal torques has also been studied via gyroscopic coupling on SO(3) [13].

Event-triggered control addresses a different onboard limitation. Instead of recomputing and transmitting a command at every clock tick, it updates the command only when a measurable error has changed enough. The basic idea and the distinction between event-triggered and self-triggered control are well established [14, 15]. Spacecraft applications have used event conditions and learning observers to reduce communication or controller updates while preserving a small error set [16, 17]. Related work has also addressed two-torque spacecraft attitude control under various constraints [18], and event-triggered adaptive attitude tracking under unknown actuator faults [19]. For the underactuated controller considered here, however, a purely state-based event is insufficient: the virtual missing-axis input is generated by a time-dependent carrier. If its torque sample is held indefinitely, the phase information becomes stale. We therefore combine a state-dependent condition with a maximum holding time.

1.1 Contributions and positioning

This paper combines underactuated attitude control, two bounded spacecraft torques, and event-triggered updates. The desired quaternion is constant after the initial command step. Thus, we consider constant-command tracking, which is equivalently local practical stabilization of the attitude-error dynamics, rather than tracking of a general time-varying trajectory. The main contributions are as follows:

1. We present a transparent carrier construction that converts the desired unactuated-axis acceleration into two periodic rate references whose frozen-parameter product has the required average value.

2. We develop a practical event-triggered implementation that combines a 2 N·m torque limit, a zero-order hold, and a digital event rule that enforces engineering-scale intervals of 0.2–1.0 s while excluding infinitely fast updates by construction.

3. We provide a nominal averaged-system analysis, a reproducible nonlinear simulation, and a threshold sweep to explain the local mechanism, assess the complete implementation, and quantify how the event threshold changes the update frequency.

The remainder of this paper is organized as follows. Section 2 introduces the model and constant-command objective; Sections 3 and 4 present the gyroscopic controller and event implementation; Section 5 gives the numerical study; and Section 6 concludes the paper. Appendix A derives the carrier references in (9)–(12).

Notation

For $x \in \mathbb{R}^3$, $x^\times$ is the skew-symmetric matrix satisfying $x^\times y = x \times y$. The Euclidean norm is $\|\cdot\|$. For a periodic signal f, $\overline{f}$ denotes its average over one carrier period. Likewise, a bar over a state, e.g., $\bar{q}_{ve,1}$ or $\bar{\omega}_1$, denotes the corresponding slow state of the nominal averaged model in Section 3.3.

2 Model and Constant-Command Tracking Objective

2.1 Rigid-body model

Let $\omega = \mathrm{col}(\omega_1, \omega_2, \omega_3)$ be the body angular velocity. Body axes are selected along the principal axes, so that $J = \mathrm{diag}(J_1, J_2, J_3) > 0$. The rigid-body dynamics are

$$J\dot{\omega} = -\omega^{\times} J\omega + \tau, \qquad \tau = \begin{bmatrix} 0 \\ \tau_2 \\ \tau_3 \end{bmatrix}. \tag{1}$$

Equivalently,

$$J_1\dot{\omega}_1 = (J_2 - J_3)\omega_2\omega_3, \tag{2}$$
$$J_2\dot{\omega}_2 = (J_3 - J_1)\omega_3\omega_1 + \tau_2, \tag{3}$$
$$J_3\dot{\omega}_3 = (J_1 - J_2)\omega_1\omega_2 + \tau_3. \tag{4}$$

The structural coupling condition required by the virtual-input construction used below is

$$J_2 \neq J_3. \tag{5}$$

If $J_2 = J_3$, (2) gives $\dot{\omega}_1 = 0$ and the proposed mechanism cannot regulate a nonzero first-axis rate. This is a condition on the proposed coupling channel, not a claim of global controllability for the full nonlinear spacecraft.

2.2 Quaternion kinematics

Let the commanded attitude after the initial step be the constant quaternion q_d and let $q_e = q_d^{-1} \otimes q = \mathrm{col}(q_{0e}, q_{ve}) \in \mathbb{S}^3$ be the scalar-first error quaternion. In general the angular-rate error is $\omega_e = \omega - R_e^{\top}\omega_d$, where R_e is the error rotation matrix. For the constant command used here, $\omega_d = 0$ and hence $\omega_e = \omega$. The error kinematics are

$$\dot{q}_{0e} = -\frac{1}{2}\, q_{ve}^{\top}\omega, \tag{6}$$
$$\dot{q}_{ve} = \frac{1}{2}\left(q_{0e}\mathrm{I} + q_{ve}^{\times}\right)\omega. \tag{7}$$

The quaternion sign is selected so that $q_{0e} \geq 0$; the degenerate case $q_{0e} = 0$ corresponds to a $180°$ rotation and is excluded by restricting attention to the local chart $q_{0e} > 0$. The objective is to drive (q_{ve}, ω) to a small neighborhood of zero using only τ_2 and τ_3, subject to

$$|\tau_2| \leq \tau_{\max}, \qquad |\tau_3| \leq \tau_{\max}, \qquad \tau_{\max} = 2\ \mathrm{N{\cdot}m}. \tag{8}$$

Figure 1 summarizes the physical control situation and the event-triggered signal path.

3 Gyroscopic Virtual-Input Controller

3.1 Carrier construction

Near the desired attitude, (7) gives $\dot{q}_{ve,1} \simeq \omega_1/2$. Select a desired unactuated-axis acceleration

$$v_1 = -k_r q_{ve,1} - d_r\omega_1, \qquad k_r, d_r > 0, \tag{9}$$

and define the product required by (2) as

$$p = \frac{J_1 v_1}{J_2 - J_3}. \tag{10}$$

The step-by-step derivation of (9)–(12) is given in the Appendix. For a frozen p, two zero-mean angular-rate references are constructed as

$$r_2(t) = A(p)\sin(\nu t), \qquad r_3(t) = B(p)\sin(\nu t), \tag{11}$$
$$A(p) = [2(p^2 + \epsilon_p^2)]^{1/4}, \quad B(p) = \frac{2p}{A(p)},$$

where $\nu > 0$ is the carrier frequency. The regularization parameter $\epsilon_p > 0$ has the same units as p and keeps $A(p)$ strictly positive, which avoids a singular derivative and division by zero at $p = 0$. For a frozen p over one carrier period,

$$\overline{r_2 r_3} = \frac{A(p)B(p)}{2} = p. \tag{12}$$

Thus, although both references have zero mean, their product creates a nonzero average input in (2). This is the key step: replacing either reference by zero would remove control authority about axis 1. The factorization $A(p)B(p) = 2p$ is not unique; the selected smooth factorization is one convenient implementation. It also implies $A(0) > 0$ and $B(0) = 0$, so a small residual carrier remains at the target (see Section 3.3).

3.2 Two-axis torque law

Let $q_{ve,a} = \mathrm{col}(q_{ve,2}, q_{ve,3})$, $\omega_a = \mathrm{col}(\omega_2, \omega_3)$, and $r = \mathrm{col}(r_2, r_3)$. Choose the commanded actuated-axis acceleration

$$a = \dot{r} - k_\omega(\omega_a - r) - k_q q_{ve,a}, \qquad k_\omega, k_q > 0. \tag{13}$$

The available torques are

$$\tau_2^c = J_2 a_2 - (J_3 - J_1)\omega_3\omega_1, \tag{14}$$
$$\tau_3^c = J_3 a_3 - (J_1 - J_2)\omega_1\omega_2. \tag{15}$$

The actuator command applies componentwise saturation,

$$\tau_i^{\star} = \mathrm{sat}_{\tau_{\max}}(\tau_i^c) = \mathrm{sgn}(\tau_i^c)\min(|\tau_i^c|, \tau_{\max}), \qquad i = 2, 3. \tag{16}$$

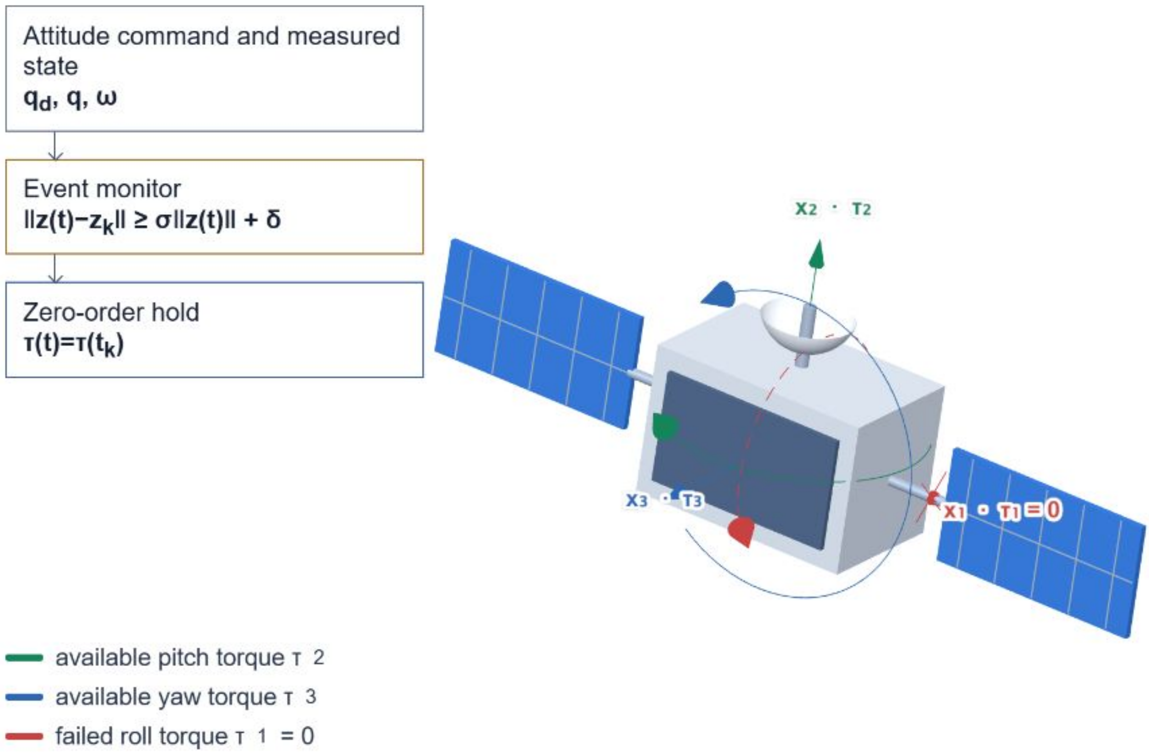

Figure 1. Underactuated spacecraft attitude-command tracking architecture. The first body-axis actuator is unavailable, and the event monitor updates a zero-order-held command for the remaining two axes.

When saturation is inactive, substitution into (3)–(4) yields $\dot{\omega}_a = a$. Consequently, the actuated rates track the carrier while the $-k_q q_{ve,a}$ term removes lateral attitude offsets. Saturation preserves the specified torque ceiling but can slow this inner tracking loop. The derivatives in (13) are analytical. For example, $A' = Ap/[2(p^2 + \epsilon_p^2)]$ and $B' = 2/A - 2pA'/A^2$, where the primes denote differentiation with respect to p. Because $p = p(q_e, \omega)$, the reference derivatives contain both amplitude and phase terms:

$$\begin{aligned}
\dot{r}_2 &= A'(p)\dot{p}\sin(\nu t) + \nu A(p)\cos(\nu t), \\
\dot{r}_3 &= B'(p)\dot{p}\sin(\nu t) + \nu B(p)\cos(\nu t), \\
\dot{p} &= \frac{J_1}{J_2 - J_3}\left(-k_r \dot{q}_{ve,1} - d_r \dot{\omega}_1\right) \\
&= -\frac{J_1 k_r}{J_2 - J_3}\dot{q}_{ve,1} - d_r \omega_2 \omega_3.
\end{aligned}$$

The exact quantities $\dot{q}_{ve,1}$ and $\dot{\omega}_1$ are evaluated from (7) and (2) at each control evaluation; no numerical differentiation or angular-acceleration measurement is required.

3.3 Nominal averaged local behavior

The stability idea can be understood without advanced nonlinear theory, but the model to which it applies must be stated explicitly. First consider the nominal averaged slow subsystem in the local quaternion chart. Assume that p changes slowly during one carrier period, the actuated rates track r_2 and r_3, the torque is continuously recomputed, and saturation is inactive. The following calculation concerns only this explicitly defined reduced model; it is not an averaging-theorem result for the complete closed loop. Equation (12) then makes the average of (2) equal to the requested virtual acceleration v_1. Near the commanded attitude, $\dot{q}_{ve,1} \approx \omega_1/2$, so the slow first-axis tracking errors obey

$$\dot{\bar{q}}_{ve,1} = \tfrac{1}{2}\bar{\omega}_1, \qquad \dot{\bar{\omega}}_1 = -k_r \bar{q}_{ve,1} - d_r \bar{\omega}_1. \tag{17}$$

This is simply a damped second-order system: k_r pulls the attitude error toward zero and d_r removes angular motion.

For completeness, define the nonnegative error energy

$$V_r = k_r \bar{q}_{ve,1}^2 + \tfrac{1}{2}\bar{\omega}_1^2. \tag{18}$$

Using (17) gives

$$\dot{V}_r = -d_r \bar{\omega}_1^2 \le 0. \tag{19}$$

Thus the average energy cannot increase. Within the set $\dot{V}_r = 0$, a nonzero attitude error immediately creates a nonzero rate through the second equation in (17). Hence the largest invariant subset of $\dot{V}_r = 0$ is the origin, which is asymptotically stable for this nominal averaged first-axis subsystem.

The directly actuated axes admit the same simple interpretation. Let $e_a = \omega_a - r$. Under the same frozen-p assumption, the carriers have zero mean and the local nominal average equations are

$$\begin{aligned}\dot{\bar{q}}_{ve,a} &\simeq \tfrac{1}{2}(\bar{e}_a + \bar{r}) = \tfrac{1}{2}\bar{e}_a, \qquad \bar{r} = 0,\\ \dot{\bar{e}}_a &= -k_\omega \bar{e}_a - k_q \bar{q}_{ve,a}.\end{aligned}$$

For $V_a = k_q\|\bar{q}_{ve,a}\|^2 + \|\bar{e}_a\|^2/2$, direct differentiation gives $\dot{V}_a = -k_\omega\|\bar{e}_a\|^2 \le 0$. Within $\dot{V}_a = 0$, $\bar{e}_a = 0$, and the second average equation then requires $\bar{q}_{ve,a} = 0$ for invariance. This completes the spring–damper explanation for the nominal averaged three-axis error model.

This calculation does not establish asymptotic stability of the full time-periodic, saturated, event-triggered system. In particular, $A(0) = (2\epsilon_p^2)^{1/4} > 0$, so the regularized carrier does not vanish exactly and the error origin is generally not an equilibrium of the complete closed loop. Finite carrier frequency, inner-loop tracking error, saturation, and zero-order holding add further differences from the ideal average model. The implementable controller is therefore evaluated in terms of local practical tracking to a carrier- and implementation-dependent neighborhood. The nonlinear simulation checks this behavior for the selected initial condition; it is not a proof of global convergence.

The parameters also have a clear physical interpretation. A faster carrier better reproduces the desired average product but requires more feedforward torque. A slower carrier respects the actuator ceiling more easily but slows the maneuver. Decreasing ϵ_p reduces the residual carrier near the target but sharpens the amplitude derivatives; increasing it has the opposite effect.

4 Event-Triggered Implementation

Define the dimensionless weighted local state

$$z = \mathrm{col}(2q_{ve}, \omega/\omega_s), \qquad \omega_s = 1\ \mathrm{rad/s}, \tag{20}$$

and let $z_k = z(t_k)$. At an event time, (9)–(16) are evaluated and the bounded torque is held:

$$\tau(t) = \mathrm{col}(0, \tau_2^\star(t_k), \tau_3^\star(t_k)), \quad t \in [t_k, t_{k+1}). \tag{21}$$

Starting at $t_k + h_{\min}$, a digital monitor checks every $h_{\min}$ seconds whether

$$\|z(t) - z_k\| \ge \sigma\|z(t)\| + \delta, \tag{22}$$

where $\sigma > 0$ and $\delta > 0$. The first grid point satisfying (22) becomes t_{k+1}. If no such point occurs, a forced update is made at $t_k + h_{\max}$; here $h_{\max}$ is selected as an integer multiple of $h_{\min}$. Hence

$$h_{\min} \le t_{k+1} - t_k \le h_{\max}. \tag{23}$$

The lower bound excludes Zeno behavior by construction. The upper bound is important here because the nominal law is explicitly time-varying: even a nearly constant state must not hold one carrier sample indefinitely.

The meaning of (22) is straightforward. The left-hand side is how much the measured tracking state has changed since the last update. The first term on the right allows a larger mismatch when the tracking error is large; the absolute term δ prevents unnecessary updates near zero. Increasing δ usually lengthens the event interval, but it can also increase the final tracking neighborhood. At a monitoring instant, the event test either resets the sampled-state mismatch or verifies that it is below the trigger threshold. The timeout separately limits the holding time to $h_{\max}$ and the carrier-phase age to $\nu h_{\max}$. No continuous-time mismatch or analytical ultimate bound is claimed between monitoring instants; the resulting practical tracking neighborhood is assessed numerically in Section 5.

4.1 Practical tuning sequence

The parameters can be selected without an optimization routine. First, choose k_r and d_r from the desired slow first-axis poles in (17). Next, set k_ω above the resulting slow bandwidth and choose k_q for a well-damped lateral response. Increase ν until the carrier average is accurate, while checking torque and rate peaks. For event implementation, select $h_{\min}$ from the sensing and processor limit, take $h_{\max}$ as a small fraction of the carrier period, and then increase σ and δ toward the target update frequency while checking the allowed tracking error. This sequence separates control design from scheduling.

At each event, implementation requires only the following operations: select the quaternion sign; evaluate v_1, p, A, B and their derivatives; compute $r, \dot{r}$; evaluate (13)–(16); and store the two bounded torque commands and z_k. Between events, the measured state is mapped to z at each scheduled monitoring instant, and the trigger norm and timeout are evaluated. The full control law and torque command are not recomputed unless an event or timeout occurs. No online optimization is involved.

5 Numerical Example

The complete nonlinear equations (1)–(7) were integrated by fixed-step fourth-order Runge–Kutta with a 0.01 s step for 200 s. No random quantities are used. The commanded attitude is the identity quaternion after a step command at $t = 0$. The initial 3–2–1 roll–pitch–yaw angles are $(25, -12, 15)^\circ$, corresponding to an initial attitude-command error of 32.6°, and the initial angular rate is $(3, -2, 2.5)^\circ$/s. Quaternions are normalized after every integration step. Table 1 lists every controller and event parameter.

Table 1. Simulation parameters.

Quantity	Value
J	$\mathrm{diag}(8, 5, 3)$ kg·m^2
$\tau_{\max}$	2 N·m
k_r, d_r	0.05 s^{-2}, 0.45 s^{-1}
k_q, k_ω	0.25 s^{-2}, 1.2 s^{-1}
ν, ϵ_p	0.8 rad/s, 10^{-6} s^{-2}
ω_s	1 rad/s
σ, δ	$0.08,\ 0.010$
$h_{\min}, h_{\max}$	$0.20,\ 1.00$ s

Figure 2 compares the attitude-angle error $\theta_e = 2\cos^{-1}|q_{0e}|$ with a baseline recomputed at every 0.01 s integration step (100 Hz), and shows the event-triggered angular rates and torques. The first torque remains identically zero, while the dashed horizontal lines show the ± 2 N·m limits. The limit is reached during the initial transient but is never exceeded. The carrier temporarily increases the total attitude-angle error because it must create motion on the available axes before gyroscopic coupling can correct the missing axis. This behavior is an expected cost of the simple two-torque strategy rather than hidden overshoot. The peak attitude error is 56.8°, and the peak angular-rate norm is 34.6°/s.

The event-triggered response and the 100 Hz baseline remain close. At 200 s, their attitude errors are 0.108° and 0.0897°, respectively. For the event-triggered case, the maximum attitude error over the last 20 s is 0.110° and the maximum angular-rate norm is 0.0937°/s. The response enters and thereafter remains in the band $\theta_e < 1^\circ$ and $\|\omega\| < 0.2^\circ$/s at 131.7 s.

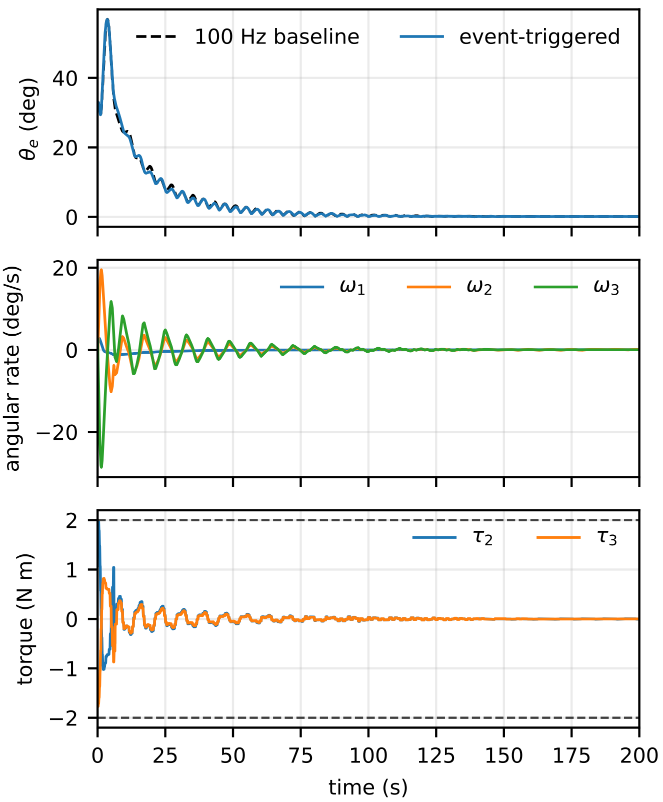

Figure 2. Nonlinear closed-loop response. Dashed: baseline recomputed at every 0.01 s integration step (100 Hz). Solid: event-triggered zero-order-hold control.

At least one available-axis torque equals the 2 N·m limit in only 0.100% of the recorded samples; after the initial slew, the required torque is much smaller. The small nonzero terminal neighborhood is consistent with the combined effects of the residual regularized carrier, finite-frequency averaging, and event holding; it should not be interpreted as exact convergence of the full time-periodic loop.

Figure 3 shows the completed inter-update intervals. The reported 295 full control-law evaluations include the initial evaluation at $t_0 = 0$ and 294 subsequent event- or timeout-induced evaluations; no terminal evaluation at $t = 200$ is counted. The 294 differences between successive update times have minimum, mean, and maximum values of 0.20, 0.679, and 1.00 s. The last update occurs at 199.6 s, and the final truncated 0.4 s hold is not included in these interval statistics. With the same endpoint convention, a 5 Hz periodic implementation uses 1000 applied updates over $[0, 200)$ s. The event rule therefore removes 70.5% of the full control-law evaluations and torque-command updates. The event monitor still samples the state every 0.20 s, and the comparison does not by itself measure processor energy.

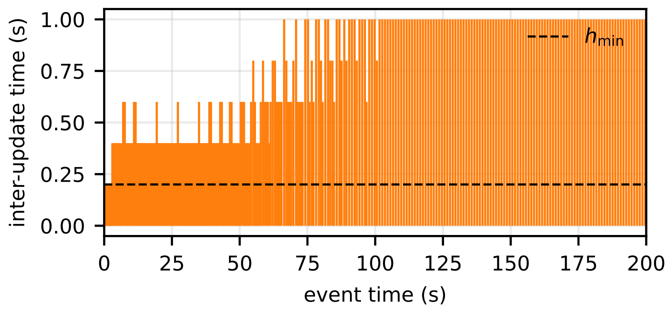

Figure 3. Completed inter-update intervals; the dashed line marks the imposed minimum dwell time.

To examine sensitivity to the absolute event threshold, Table 2 repeats the deterministic simulation for three absolute thresholds, with every other parameter unchanged. Every row uses the same initial/terminal counting convention and averages only completed inter-update intervals; the final column is the maximum attitude error over the last 20 s. Increasing δ from 0.005 to 0.020 reduces the update count from 360 to 256 and increases the mean interval from 0.555 to 0.780 s. The last-20-s attitude bound remains close to 0.11° because the 1 s timeout dominates near the command. Thus, this sweep demonstrates update-frequency sensitivity to δ, while the tested range does not show a pronounced accuracy tradeoff. It is not presented as a robustness campaign.

Table 2. Absolute-threshold sensitivity.

δ	Updates	Reduction	Mean (s)	$\max \theta_e$
0.005	360	64.0%	0.555	0.1100°
0.010	295	70.5%	0.679	0.1098°
0.020	256	74.4%	0.780	0.1102°

6 Conclusion

This paper presented an event-triggered attitude control method for an underactuated spacecraft with one failed actuator. Two periodic rate references use ordinary gyroscopic coupling to create the missing average control action. A simple error-energy calculation explains the asymptotic behavior of the explicitly defined nominal averaged subsystem. In the reported nonlinear case, the residual carrier, finite-frequency averaging, saturation, and event holding produce a small tracking neighborhood rather than exact convergence. The actual commands are limited to ± 2 N·m and held between events separated by 0.2–1.0 s. For the selected initial condition, the nonlinear quaternion simulation gives a final attitude error of 0.108°, while the event rule uses 70.5% fewer full control-law evaluations and torque-command updates than a 5 Hz periodic implementation.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported by the National Natural Science Foundation of China under Grant 62573211.

Conflicts of Interest

Chengxi Zhang served as an Editor-in-Chief of the *Aerospace Engineering Communications* at the time of manuscript submission. To ensure the integrity of the peer-review process, Chengxi Zhang was not involved in the editorial handling, peer review, or decision-making process for this manuscript, which was handled independently by another editor. The author declares no other conflicts of interest.

AI Use Statement

The author declares that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] Forbes, J. R. (2015). Fundamentals of spacecraft attitude determination and control [bookshelf]. *IEEE Control Systems Magazine*, *35*(4), 56-58. [CrossRef]

[2] Krishnan, H., Reyhanoglu, M., & McClamroch, H. (1994). Attitude stabilization of a rigid spacecraft using two control torques: A nonlinear control approach based on the spacecraft attitude dynamics. *Automatica*, *30*(6), 1023-1027. [CrossRef]

[3] Frias, A., de Ruiter, A. H., & Kumar, K. D. (2019). Velocity-free spacecraft attitude stabilization using two control torques. *Automatica*, *109*, 108553. [CrossRef]

[4] Hughes, P. C. (1986). *Spacecraft Attitude Dynamics*. New York: Wiley. (ISBN 0-471-81842-9; reprinted by Dover Publications, Mineola, NY, 2004.) https://archive.org/details/spacecraftattitu0000hugh

[5] Sanyal, A., Fosbury, A., Chaturvedi, N., & Bernstein, D. S. (2009). Inertia-free spacecraft attitude tracking with disturbance rejection and almost global stabilization. *Journal of guidance, control, and dynamics*, *32*(4), 1167-1178. [CrossRef]

[6] Mayhew, C. G., Sanfelice, R. G., & Teel, A. R. (2011). Quaternion-based hybrid control for robust global attitude tracking. *IEEE Transactions on Automatic control*, *56*(11), 2555-2566. [CrossRef]

[7] Yin, S., Xiao, B., Ding, S. X., & Zhou, D. (2016). A review on recent development of spacecraft attitude fault tolerant control system. *IEEE Transactions on Industrial Electronics*, *63*(5), 3311-3320. [CrossRef]

[8] BoÏskovic, J. D., Li, S. M., & Mehra, R. K. (2001). Robust adaptive variable structure control of spacecraft under control input saturation. *Journal of Guidance, Control, and Dynamics*, *24*(1), 14-22. [CrossRef]

[9] Hu, Q., Xiao, B., Wang, D., & Poh, E. K. (2013). Attitude control of spacecraft with actuator uncertainty. *Journal of Guidance, Control, and Dynamics*, *36*(6), 1771-1776. [CrossRef]

[10] Shao, X., Hu, Q., Shi, Y., & Jiang, B. (2018). Fault-tolerant prescribed performance attitude tracking control for spacecraft under input saturation. *IEEE Transactions on Control Systems Technology*, *28*(2), 574-582. [CrossRef]

[11] Wu, B., & Cao, X. (2017). Robust attitude tracking control for spacecraft with quantized torques. *IEEE transactions on Aerospace and Electronic Systems*, *54*(2), 1020-1028. [CrossRef]

[12] Zhu, Z., Xia, Y., & Fu, M. (2011). Attitude stabilization of rigid spacecraft with finite-time convergence. *International Journal of Robust and Nonlinear Control*, *21*(6), 686-702. [CrossRef]

[13] Gui, H., Vukovich, G., & Xu, S. (2015). Attitude tracking of a rigid spacecraft using two internal torques. *IEEE Transactions on Aerospace and Electronic Systems*, *51*(4), 2900-2913. [CrossRef]

[14] Tabuada, P. (2007). Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Transactions on Automatic control*, *52*(9), 1680-1685. [CrossRef]

[15] Heemels, W. P., Johansson, K. H., & Tabuada, P. (2012, December). An introduction to event-triggered and self-triggered control. In *2012 ieee 51st ieee conference on decision and control (cdc)* (pp. 3270-3285). IEEE. [CrossRef]

[16] Wu, B., Shen, Q., & Cao, X. (2018). Event-triggered attitude control of spacecraft. *Advances in Space Research*, *61*(3), 927-934. [CrossRef]

[17] Zhang, C., Wang, J., Zhang, D., & Shao, X. (2018). Learning observer based and event-triggered control to spacecraft against actuator faults. *Aerospace Science and Technology*, *78*, 522-530. [CrossRef]

[18] Zou, A. M., Kumar, K. D., & de Ruiter, A. H. (2017). Spacecraft attitude control using two control torques. *Information Sciences*, *408*, 23-40. [CrossRef]

[19] Wang, C., Guo, L., Wen, C., Hu, Q., & Qiao, J. (2019). Event-triggered adaptive attitude tracking control for spacecraft with unknown actuator faults. *IEEE Transactions on Industrial Electronics*, *67*(3), 2241-2250. [CrossRef]

Appendix

A Derivation and Scope of the Carrier Construction

The carrier in (9)–(12) is obtained by treating the gyroscopic rate product as a virtual input. The derivation is given below in three steps.

A.1 Desired motion of the unactuated axis

From (2),

$$\dot{\omega}_1 = c_J\omega_2\omega_3, \qquad c_J = \frac{J_2 - J_3}{J_1} \neq 0. \tag{A1}$$

Thus, the product $\omega_2\omega_3$, rather than either available torque alone, determines the acceleration about the unactuated axis. The exact first component of (7) is

$$\dot{q}_{ve,1} = \tfrac{1}{2}(q_{0e}\omega_1 + q_{ve,2}\omega_3 - q_{ve,3}\omega_2). \tag{A2}$$

Near the constant command, $q_{0e} \simeq 1$ and the lateral cross terms are small, giving $\dot{q}_{ve,1} \simeq \omega_1/2$. We therefore select the desired first-axis acceleration as (9). If $\dot{\omega}_1 = v_1$ could be imposed directly, the local error would satisfy

$$\ddot{q}_{ve,1} + d_r\dot{q}_{ve,1} + \frac{k_r}{2}q_{ve,1} \simeq 0, \tag{A3}$$

which is a damped second-order equation for $k_r, d_r > 0$.

A.2 Required average rate product

Because axis 1 has no actuator, (A1) shows that the required average product is

$$\overline{\omega_2\omega_3} = \frac{v_1}{c_J} = \frac{J_1 v_1}{J_2 - J_3} = p, \tag{A4}$$

which gives (10). Let $T_c = 2\pi/\nu$, temporarily freeze p over one carrier period, and define

$$\langle f \rangle_{t_0} := \frac{1}{T_c}\int_{t_0}^{t_0+T_c} f(t)\,dt.$$

For the references in (11),

$$\langle r_2 \rangle_{t_0} = 0, \qquad \langle r_3 \rangle_{t_0} = 0,$$
$$\langle r_2 r_3 \rangle_{t_0} = \frac{A(p)B(p)}{2}. \tag{A5}$$

Consequently, only the amplitude product is prescribed:

$$A(p)B(p) = 2p.$$

For $p > 0$ the two carriers are in phase; for $p < 0$, the signed amplitude $B(p)$ makes them opposite in phase. This produces either sign of the required virtual input without a nonzero mean in either frozen-parameter reference.

A.3 Regularization and approximation scope

The amplitude condition does not determine a unique pair. The choice in (11) keeps $A(p) > 0$, gives $B(p)$ the sign of p, and satisfies $A(p)B(p) = 2p$ exactly for every p. It is one convenient smooth allocation between axes 2 and 3; no optimality of this allocation is claimed. Treating radians as dimensionless, v_1, p, and ϵ_p have units s^{-2}, while A, B, r_2, and r_3 have units s^{-1}. Accordingly, k_r and k_q have units s^{-2}, and d_r and k_ω have units s^{-1}.

In the implemented controller, $p = p(t)$ is not exactly frozen. The identity $A(p)B(p) = 2p$ still holds pointwise, but the one-period product average is

$$\langle r_2 r_3 \rangle_{t_0} = \frac{2}{T_c} \int_{t_0}^{t_0+T_c} p(t) \sin^2(\nu t)\, dt, \tag{A6}$$

which need not equal $p(t_0)$. If $|\dot{p}| \le L_p$ during that period, then

$$\left| \langle r_2 r_3 \rangle_{t_0} - p(t_0) \right| \le L_p T_c. \tag{A7}$$

Thus, the frozen-p relation becomes accurate when p changes slowly relative to the carrier. Moreover, if $e_i = \omega_i - r_i$, then

$$\omega_2\omega_3 - r_2 r_3 = r_2 e_3 + r_3 e_2 + e_2 e_3, \tag{A8}$$

so inner-loop error, saturation, and event-held torque further perturb the requested average virtual input.

Finally, $\epsilon_p > 0$ makes the amplitude derivatives finite, but $A(0) = 2^{1/4}\sqrt{\epsilon_p} > 0$ and $B(0) = 0$. A small axis-2 carrier therefore remains at $p = 0$, so the full time-periodic closed loop generally does not have the error origin as an equilibrium. This explains why the implemented result is local practical tracking, whereas asymptotic stability applies only to the nominal averaged subsystem. This tradeoff cannot be removed by making both factors smooth and zero at the origin while preserving the exact identity $A(p)B(p) = 2p$: if differentiable A and B both vanished at zero, then $(AB)'(0) = 0$, in contradiction with $(2p)' = 2$.

Chengxi Zhang received B.S. and M.S. degrees from Harbin Institute of Technology, China, in 2012 and 2015; and Ph.D. degree from Shanghai Jiao Tong University, China, in 2019. He is now an Associate Professor at Jiangnan University, China. His interests are space engineering, robotic systems & control. He is the Editor-in-Chief of *Aerospace Engineering Communications*. He is an Editor of *Frontiers in Aerospace Engineering*, *Frontiers in Mechanical Engineering*, and *Scientific Reports*, on the Editorial Board of *Symmetry* and *Aerospace Systems*. He is a Committee Member of the Youth Working Committee, Chinese Association of Automation since 2025.