Aerodynamic Roll Control on Project Zephyrus

Connor Sterling

June 2026

1 Introduction

The MIT Rocket Team's Project Zephyrus is a high altitude amateur rocket, targeting an apogee of 100,000 ft. Comprising a single stage with a P class motor, Zephyrus is arguably the team's most ambitious high-powered project to date. In addition to the high impulse motor, Zephyrus features aerodynamic roll control. Using two "tabs" and PD control, the rocket is stabilized to prevent roll and yield pristine camera footage.

This report aims to document the methods and results of roll control on the project to date. This engineer has authored similar reports for Project Xanthus, the demonstrator and precursor to Zephyrus [1, 2], as well as documentation on servo motor system identification experiments for both Xanthus [3] and Zephyrus [4]. The reader is encouraged to review these, as they provide valuable background on the project. However, this document is intended to stand alone, and does not require familiarity with the aforementioned reports.

Section 2 briefly explains the rocket's roll dynamics. Section 3 describes the control system design process. Zephyrus' test launch is recounted in Section 4, and an analysis thereof is presented in Section 5. Section 6 concludes the report.

2 Roll Dynamics

A similar explanation of roll dynamics can be found in [1]. The physics are identical as those described here, though there have been some changes in notation.

A rocket in roll can be treated as a rigid body acted on by the sum of external moments. The governing equation for this motion is

$$J_{xx}\ddot{\phi} = M_x. \tag{1}$$

The roll angle is represented as ϕ. J_{xx} is the moment of inertia about the roll (x) axis and M_x is the sum of moments about the axis. This neglects damping due to roll angular velocity (i.e. it does not include a term proportional to $\dot{\phi}$). This is justifiable if $\dot{\phi} \ll V$ [5].

There are two main sources of roll moment on the rocket: steady forcing due to manufacturing defects, M_x^{mfg}, and forcing due to control surface actuation, M_x^{ctrl}. Zephyrus has two aileron-like control surfaces referred to as "tabs". They can be seen in Fig. 1, signed by one Olivia Garcia and a Connor Sterling.

The actuation angle of the tabs has been designated α since for a rocket flying parallel to the wind this corresponds to the tabs' angle of attack. The moment due to the tabs is given by

$$M_x^{\text{ctrl}} = \frac{1}{2}\rho V^2 \; S_{\text{ref}} \; c_{\text{ref}} \; C_{M_x\alpha} \; \alpha. \tag{2}$$

The first group of terms, $\rho V^2/2$, is the dynamic pressure. S_{ref} and c_{ref} are the reference area and distance, respectively. For this project, the area of one tab is used as reference and the reference distance is the lever arm from the tab's center to the centerline of the rocket. $C_{M_x\alpha}$ is the derivative of tab-induced moment with respect to deflection angle at $\alpha = 0$ ($C_{M_x\alpha} = N_{tabs}C_{L\alpha}$ from [1]). For small angles of attack, tab-induced moment is linear with α. N.B. Equation 2 only holds in this linear regime.

Let q denote the dynamic pressure. The transfer function from α to ϕ is then

$$\frac{\phi}{\alpha} \equiv G = \frac{C_{M_x\alpha} \; S_{\text{ref}} \; c_{\text{ref}} \; q}{J_{xx}s^2}. \tag{3}$$

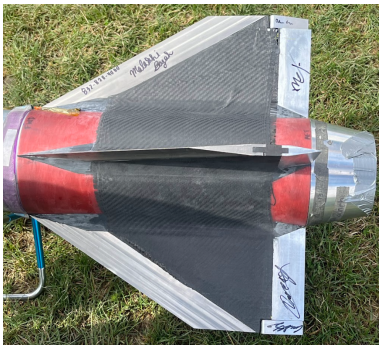

Figure 1: Zephyrus test launch fin can. The fins are "Dusty Doughnut, Jr." (top), "Rusty Radish, Jr." (center), "Crusty Cranberry, Jr.", (bottom), and "Trusty Turnip, Jr." (not pictured). The tabs are at the tips of Dusty Doughnut, Jr., and Crusty Cranberry, Jr.

The tabs are actuated by KST A12-T servo motors, commanded by the flight computer. There is no guarantee that the commanded angle (α^*) is the same as that actually achieved in flight. In fact, this is all but guaranteed to be untrue. For one, the servo motor is not an instantaneous device. Additionally, deflecting the tabs in the high-speed freestream creates an adverse torque on the servo. This actually prevents the tab from reaching the commanded angle in steady state. The servo transfer function is modeled as

$$\frac{\alpha}{\alpha^*} \equiv G_{servo}^{CL} = k_{att} \frac{b_2 s^2 + b_1 s + b_0}{a_3 s^3 + a_2 s^2 + a_1 s + a_0} \tag{4}$$

with a certain set of parameters k_{att}, a_i, and b_j. The term k_{att} is the attenuation factor. It decreases linearly with the product $C_{M_x\alpha} \cdot q$ and is unity for $q = 0$. For more information on the servo model and how its parameters are obtained, see [4].

When the commanded angle is determined by a feedback law such that

$$\alpha^* = K(\phi^* - \phi) \tag{5}$$

the system's closed loop transfer function is

$$\frac{\phi}{\phi^*} = \frac{K \cdot G_{servo}^{CL} \cdot G}{1 + K \cdot G_{servo}^{CL} \cdot G}. \tag{6}$$

Fig. 2 is a simple block diagram of the system. A more detailed one can be found in [1] or [3].

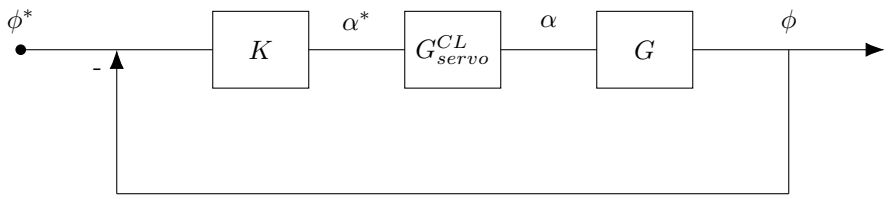

Figure 2: Simplified block diagram of system

Above is presented the linearized system dynamics for a rocket in roll. Note that this is a time-varying system. The moment of inertia will decrease as fuel is expelled. The velocity of the rocket changes drastically throughout the flight. Density decreases to less than 2% of its sea level value by the target altitude. $C_{M_x\alpha}$ is dependent on the Mach number, which varies with velocity and air temperature. And, as mentioned above, k_{att} depends on these time-varying quantities as well. The only constants in the transfer function are S_{ref}, c_{ref}, and the servo parameters a_i and b_j. This results in the need for a time-varying controller, the design of which is discussed in the following section.

3 Controller Design

As on Project Xanthus, Zephyrus uses proportional-derivative (PD) control. The controller takes the form:

$$K = K_P\ (T_D s + 1). \tag{7}$$

K_P is the dimensionless proportional gain and T_D is the derivative time constant in seconds.

To account for the time-varying quantities in the dynamics, the controller is designed at some reference condition. Let $\bar{L}(t)$ denote the time-varying component of the loop gain:

$$\bar{L}(t) = \frac{C_{M_x\alpha}(t)\ \rho(t)\ V(t)^2}{J_{xx}(t)}. \tag{8}$$

For reasons that will become apparent, k_{att} is not included in $\bar{L}(t)$ even though it is time varying. A reference $\bar{L}_0$ is chosen such that

$$\bar{L}_0 = \bar{L}(t_0), \tag{9}$$

where t_0, the reference time, corresponds to predicted max q. It is not necessary for t_0 to be the time when the rocket experiences the highest dynamic pressure, nor does L_0 even need to be a product of terms whose values arise at the same time. However, the choice of t_0 at max q is intuitive.

The PD controller is then designed at the reference conditions, yielding some K_0. For the purposes of choosing K_0, k_{att} is set to unity. The time-varying controller is then scaled by $\bar{L}_0/\bar{L}(t)$ to place the closed loop poles at the same points for any loop gain. To prevent aerodynamic stall over the control surfaces and mechanical stall of the servo motor, the tab command angle α^* saturates at $\pm 10°$. Finally, because the achieved angle will only be a fraction of the command, the latter is scaled by $1/k_{att}$. Thus, the controller is implemented as:

$$K(t) = \text{sat}\left[K_0 \cdot \frac{\bar{L}_0}{\bar{L}(t)}\right] \cdot \frac{1}{k_{att}}. \tag{10}$$

(Credit for the analytic work behind the extension of the simple $q_0/q(t)$ used in [1, 2] to the more general $\bar{L}_0/\bar{L}(t)$ must be given to this project's Chief Engineer.)

The PD controller comprises two parameters: K_P and T_D. The location of the controller's zero is solely a function of T_D. The position of this zero determines the phase of the frequency response. For a given T_D, altering K_P changes the magnitude of the response, without affecting the phase. If one enforces the constraint that K_P be chosen to maximize phase margin, this collapses the problem to a single degree of freedom (T_D). Fig. 3 plots the phase and gain margins as a function of T_D.

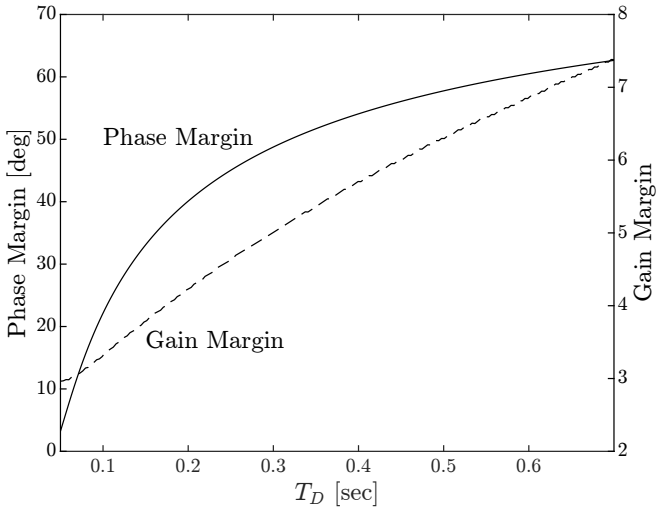

Figure 3: Phase and gain margin vs. T_D. As T_D is increased, phase margin asymptotes to 90° and gain margin becomes arbitrarily large.

Increasing the derivative time constant raises the margins. The phase margin asymptotes to 90°, and the gain margin can be made arbitrarily large. However, increasing T_D also amplifies noise in the gyroscope's signal. Thus, a balance must be struck between stability/damping and noise sensitivity.

A time constant of 0.25 seconds was chosen based on this engineer's judgment, corresponding to $K_P = 0.084$ This resulted in a gain margin of 4.65 and phase margin of 45°. Gain crossover is at 1.44 Hz and phase crossover is at 5.54 Hz. Root locus and Bode plots of the system with the controller are shown in Figs. 4 and 5.

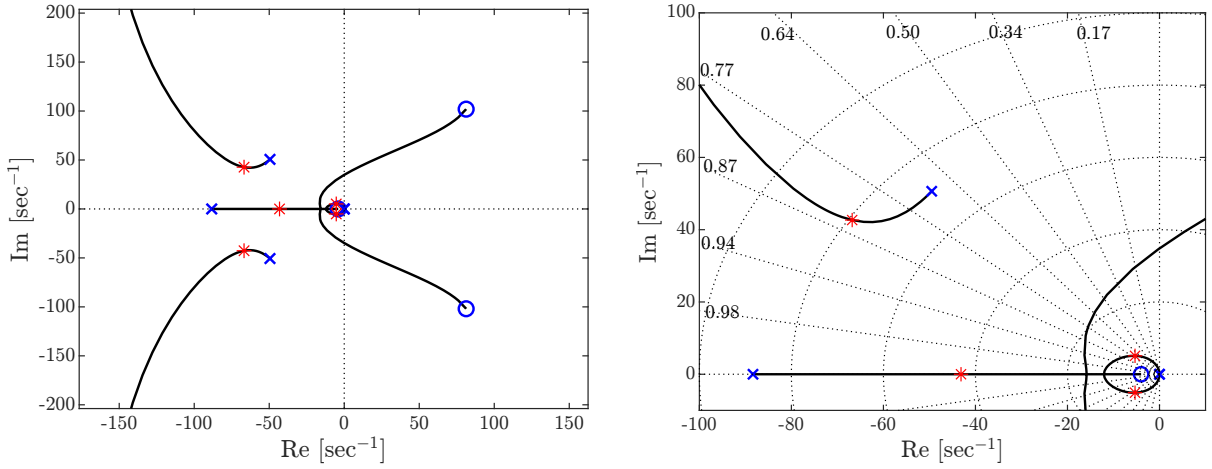

Figure 4: Root locus plot indicating chosen poles as red stars. The two open loop poles at (or approximately at) the origin are the rigid body roll dynamics. The complex zeros and all other poles are from the servo dynamics. The real zero is due to the controller.

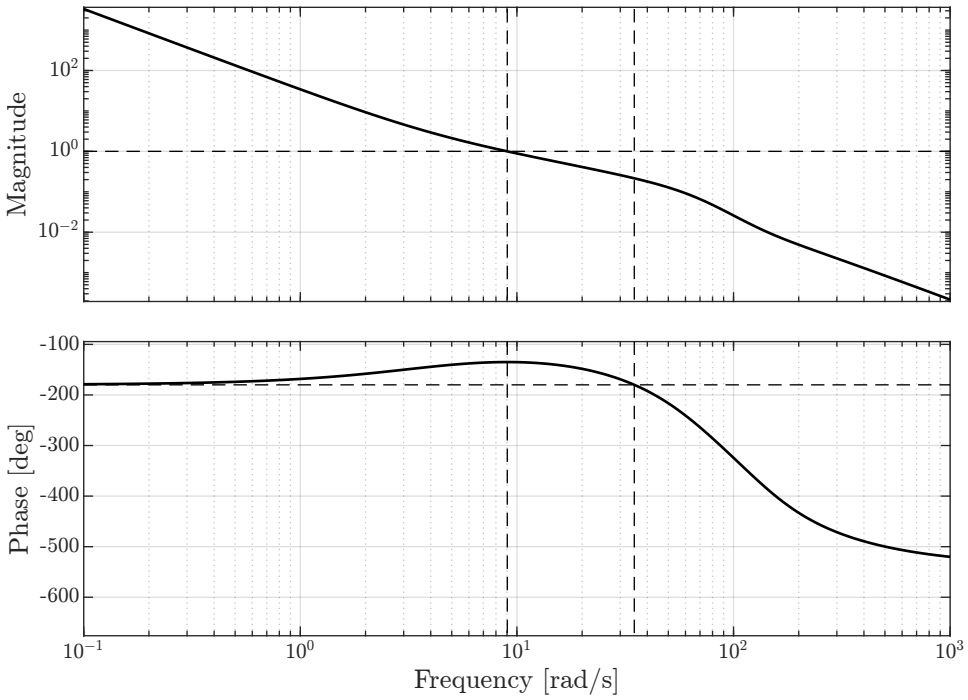

Figure 5: Bode plot indicating crossover frequencies and margins. Gain margin is 4.65 with $|L| = 1$ at 1.44 Hz. Phase margin is 45.04° with $\angle L = -180°$ at 5.54 Hz.

4 Launch Summary

The Zephyrus test launch took place on April 11th, 2026, at 14:47 EDT. The launch site was at the Upstate Rocketry Research Group (URRG) in Penn Yan, New York. According to accelerometer data from the commercial flight computer (a Multitronix Kate-3 system [6]), apogee was at 4785 m (15,698 ft). Max q occurred at a velocity of 409 m/s (1342 ft/s), corresponding to a Mach number of 1.2. Fig. 6 presents the flight profile until deployment.

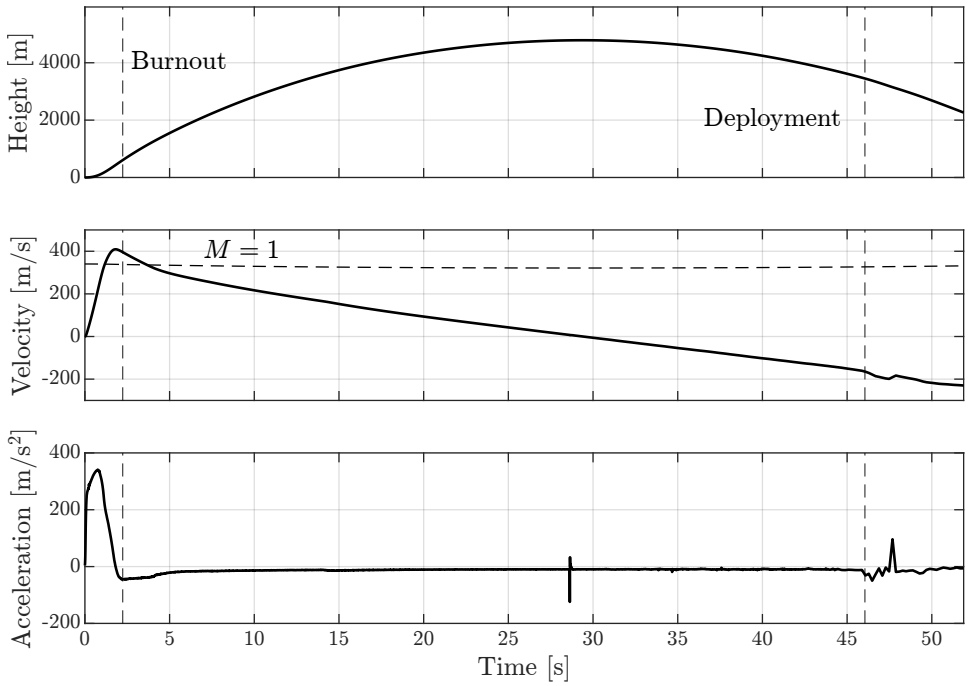

Figure 6: Zephyrus test launch profile until deployment. Apogee: 4785 m. Max velocity: 409 m/s (Mach 1.20).

Three critical system failures occurred during Zephyrus' test launch. First, the gyroscope data were processed incorrectly, scaling the true values by a factor of four. Thus, for every degree of rotation of the rocket, the flight computer (and controller) believed the vehicle had rolled four degrees. This effectively increased K_P by four times. The system appeared to remain stable, due to both the robust gain margin and benign effects of saturation. Second, the wires used to route power and signal from the flight computer to both servo motors were severed approximately three seconds into the flight. The period of (albeit violent) stability therefore only lasted three seconds, after which the rocket entered an unstable spin. These two problems in themselves were a tragedy for the Aerodynamics and Control subteam. To add salt to the wound, the last major system failure occurred on the recovery side. Upon deployment, the main parachute line was severed from the rocket, and Zephyrus descended solely on the drogue. It struck a tree at roughly 30 m/s (100 ft/s). The majority of the rocket was unscathed—but only because the fin can took one for the team and was obliterated on impact. Fig. 7 shows the damage.

Overall, test launch was a dark day for Aerodynamics and Control. Not only did roll control fail, but the subteam now faced the prospect of rebuilding the fin can from scratch. On the bright side, the GPS tracking system allowed the rocket to be located after it landed (unlike Project Aurora). Moreover, the propulsion systems worked like a charm.

A more complete version of this report would document the second launch of Zephyrus to its full altitude target of 100,000 ft and analyze the successful roll control. Unfortunately, two anomalous static fires of the full, P-class motor prevented the completion of the project on schedule. As of this writing, there are tentative plans to pursue a full launch of Zephyrus in January, but this is contingent on successfully static firing the P-motor.

For now, a brief analysis of the first three seconds of flight is presented in the following section.

(a) (b) (c)

Figure 7: The fin can post-launch. In (a), Dusty Doughnut, Jr., and Rusty Radish, Jr., remain attached to the can. In (b), Crusty Cranberry, Jr., is seen detached from the rocket with tab deflected. Trusty Turnip, Jr., buried in leaves in (c), sports a large dent in her trailing edge.

5 Analysis

The three seconds of active control do not provide a great indicator on the efficacy of the controller. The majority of this period is during the burn, and command values were four times greater than they should have been because of the gyroscope readout bug. Fig. 8 depicts the roll angle, commanded tab angle with predicted actual angle, and normalized loop gain over the seconds when roll control was active.

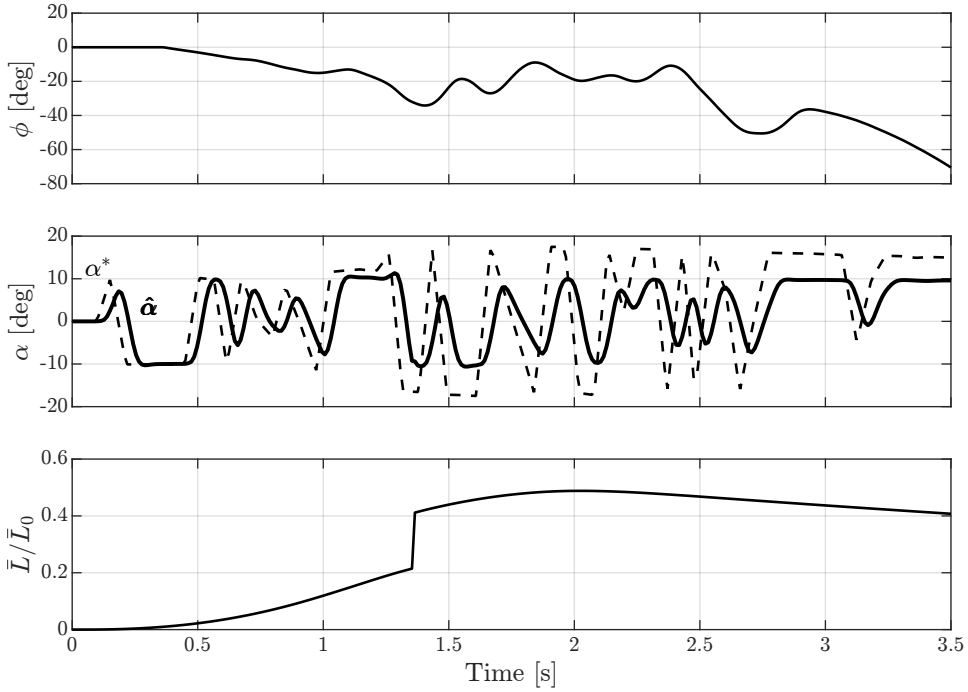

Figure 8: Roll angle, commanded tab angle, and normalized loop gain for the first 3.5 seconds of Zephyrus test launch. After about 2.9 seconds, ϕ diverges because of the servo disconnection. The commanded angle α^* is four times greater than it should be. It is plotted with $\hat{\alpha}$, the predicted tab angle. The jump in $\bar{L}/\bar{L}_0$ is due to a discontinuity in $C_{M_\infty\alpha}$ at Mach 1.

After about 2.9 seconds, the roll angle diverges. This is when the servos were disconnected. The second graph displays the commanded and estimated actual tab angles. The latter is calculated using the transfer

function G^{CL}_{servo}:

$$\hat{\alpha}(t) = g^{CL}_{servo}(t) \ * \ \alpha^*(t). \tag{11}$$

As can be seen, $\hat{\alpha}$ lags behind α^*, and does not exceed $\pm 10^\circ$, the saturation point. The lowest graph is a reference for the relative authority of and aerodynamic loading on the tabs, the ratio of current to reference time-varying loop gain components described in Section 3. The jump in the plot is due to a discontinuity in the graph of $C_{M_x\alpha}$ vs. Mach number. While the coefficient of moment is roughly constant for subsonic speeds (with some odd behavior in the transonic regime that is not included in the model), it sharply increases once the speed is supersonic, then diminishes as Mach number increases further.

As in [2], attempts were made to compare the actual and expected responses via error modeling techniques. The measured angular acceleration is compared to the predicted acceleration based on the estimated tab angle. The latter is calculated via the second derivative of Equation 3. The model takes three parameters: an efficacy constant, C_{eff}, manufacturing error, e_{mfg}, and time lag, t_{lag}. Its form is:

$$\ddot{\phi}(t) = C_{\text{eff}} \cdot \ddot{g}\left(\hat{\alpha}(t - t_{\text{lag}}) \ + \ e_{\text{mfg}}\right), \tag{12}$$

where

$$\ddot{g} = \frac{C_{M_x\alpha} \ S_{\text{ref}} \ c_{\text{ref}} \ q}{J_{xx}}. \tag{13}$$

The efficacy constant is a measure of how large the magnitude of the response is relative to prediction. One would therefore expect $C_{\text{eff}} = 1$. The manufacturing error simulates a constant roll forcing (ostensibly due to asymmetries from the manufacturing process). It has units of *degrees of tab deflection.* So $e_{\text{mfg}} = 1^\circ$ corresponds to a continuous moment equivalent to both tabs being actuated to 1°. The last parameter, time lag, reflects the possibility that the phase of the servo model is not completely accurate. Positive t_{lag} corresponds to the servo actuating slower than predicted.

One complication in this analysis is that the servos were controlled by the custom flight computer (distinct from the commercial Kate-3 device), which did not save a hard log of the flight data. Thus, the only record of the gyroscope readings and servo commands from the flight computer is what was transmitted via radio to the ground stations during flight. While Kate logs data at 100 Hz, the ground stations received transmissions variably at around 15 Hz. Since the flight computer operates at close to 50 Hz, this is only a partial record of the gyroscope readings and commands. The data analyzed here are from "Ground Station 4", where this engineer was positioned during the launch.

The optimal parameters were found to be $C_{\text{eff}} = 0.94$, $e_{\text{mfg}} = 0.9^\circ$, and $t_{\text{lag}} = 0.024$ seconds. The fitted curve is shown against the measured in Fig. 9.

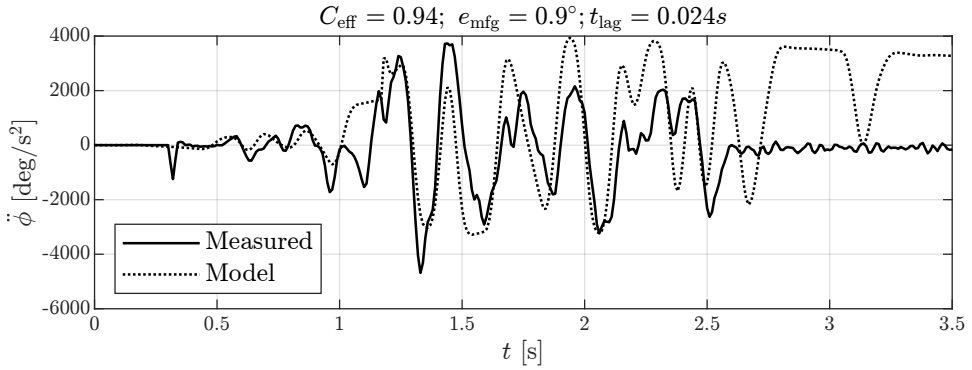

Figure 9: Error model fit to angular acceleration in roll. Optimal parameters are $C_{\text{eff}} = 0.94$, $e_{\text{mfg}} = 0.9^\circ$, and $t_{\text{lag}} = 0.024$ seconds. The servos lose power at roughly 2.5 seconds, so the model no longer aligns with measurement.

The fit is mediocre at best, but one must bear in mind that there are only three seconds of less-than-ideal data. With C_{eff} close to one, there is not convincing evidence that significant errors are present in the model. Given the quality of the data, however, perhaps the best conclusion to be drawn is none at all.

6 Conclusion

Project Zephyrus represents a great deal of progress within the team's high-powered rocketry division. The team has developed new telemetry hardware and software, accurate GPS tracking, and fired a P-class motor. Within Aerodynamics and Control, a whole new architecture was developed and deployed. This was accompanied by new manufacturing techniques, experimental investigations, and advances in the theoretical know-how of the subteam.

That being said, the work is far from over. Concerning control in particular, there are numerous avenues open for exploration, both on Project Zephyrus if the team aims to complete that mission and on future projects. This report will close with a selected list of such opportunities:

- **New mechanical architectures.** While trailing edge control surfaces on the fins are effective, other architectures possess advantages as well. Forward positioned canards (possibly with a "spin can") could be employed, or the fins themselves could be made to rotate. Such methodologies would provide much more authority and motivate advances in the subteam's mechanical design capabilities.

- **Thinner fins.** The half-inch-thick fins on Zephyrus come with a substantial drag penalty. The limiting factor is the width of the servo, so finding a smaller high-torque servo or adjusting the control surface to allow for a weaker motor would be beneficial for a high-altitude launch.

- **Lead compensation.** One concern with any derivative controller is noise amplification. This was the motivation for choosing such a low T_D, leaving potential phase margin still "on the table." Such risk can be mitigated by implementing a lead controller instead of derivative feedback.

- **Velocity control.** The roll stabilization problem can be framed in two ways: controlling position or velocity. The former approach has been adopted so far, with the setpoint being $\phi = 0$. Because the actual roll angle doesn't matter, the same result can be achieved by setting $\dot{\phi} = 0$. A PID controller could be used in this problem, but would require differencing the gyroscope readings.

- **Three axis control.** With two tabs, only the roll mode can be easily controlled. With three or four control surfaces, pitch and yaw can be stabilized as well. Control of the latter two degrees of freedom allows for trajectory control or regulation.

- **Modern control techniques.** To date, solely classical methods have been used in designing a control system for this project. Particularly if three-axis control is pursued, adopting state-space techniques may be prudent. Such a move could be accompanied by the use of an LQR, or more classical pole-placement methods.

References

[1] C. Sterling, "Aerodynamic Roll Control on Project Xanthus," 2026.

[2] ——, "Addendum to Aerodynamic Roll Control on Project Xanthus," 2026.

[3] ——, "Servo Motor Characterization for Project Xanthus," 2025.

[4] ——, "Servo Motor Characterization for Project Zephyrus," 2026.

[5] S. Niskanen, "OpenRocket technical documentation," 2013.

[6] Multitronix, "Kate-3 Transmitter Brochure," 2021. [Online]. Available: https://www.multitronix.com/kate-3-transmitter.html