Addendum to Aerodynamic Roll Control on Project Xanthus

Connor Sterling

May 2026

1 Introduction

In November 2025, the MIT Rocket Team launched the flight vehicle Xanthus, a high powered rocket flying on an I class motor, with the goal of achieving active roll control. The development of the control architecture and analysis of Xanthus' two flights can be found in [1], the precursor to this report. In brief, Xanthus was launched and recovered twice. The first flight did not achieve roll control. The second achieved roll control, but not stabilization. With the aim of finally stabilizing Xanthus in roll, the rocket was launched once more in January. Roll stabilization was successful, but Xanthus was destroyed due to recovery failure. This report includes further analysis on the second flight of Xanthus, as well as a description and analysis of the third.

Section 2 contains the additional analysis of Xanthus flight II. Section 3 briefly describes the controller for flight III. The launch is recounted in Section 4 and analyzed in Section 5. Conclusions about Project Xanthus are found in Section 6.

2 Further Analysis of Xanthus Flight II

In [1], analysis of the estimated instantaneous natural frequency was used to compare the actual roll response of Xanthus to predictions. The error is defined as

$$\text{Error} = \frac{|L_{obs.}| - |L_{theor.}|}{|L_{theor.}|}, \tag{1}$$

and was found to be -50.1%. Thus, the observed loop gain was half that predicted.

Before Xanthus' third launch, further analysis was conducted. A method of regression was used to match the observed results to a model incorporating error terms. The model takes the form

$$\ddot{\phi} = C_{\text{eff}} \frac{G_\alpha}{J_{xx}} (\alpha + e_{\text{mfg}}). \tag{2}$$

The leading term C_{eff} is the "efficacy constant", a measure of how much predicted tab authority is manifested as acceleration. This captures the accuracy of the theoretical dynamics model. The term e_{mfg} is the manufacturing error. It captures the net steady state roll forcing in the absence of tab actuation, likely due to asymmetries in the manufactured rocket. C_{eff} is dimensionless, and e_{mfg} has units of degrees, representing the equivalent forcing of the control surfaces when actuated to that angle (thus, if e_{mfg} were greater than 10°, the saturation value, the control surfaces would be unable to stop the rocket from rolling).

The nature of the error model allows the regression to be done in two stages. First, for a given C_{eff}, the optimal e_{mfg} is that which matches the means of observed and theoretical roll. This collapses the problem to one-dimensional optimization, and the best efficacy constant is chosen by minimizing the squared error. The first six seconds of the flight are used to construct the error model. The model with best fitting parameters is plotted against the measured roll in Fig. 1.

The best fitting C_{eff} is 0.59, corroborating the findings in [1]. The manufacturing error is -2.9°. This is non-negligible, but can be overcome by the control surfaces.

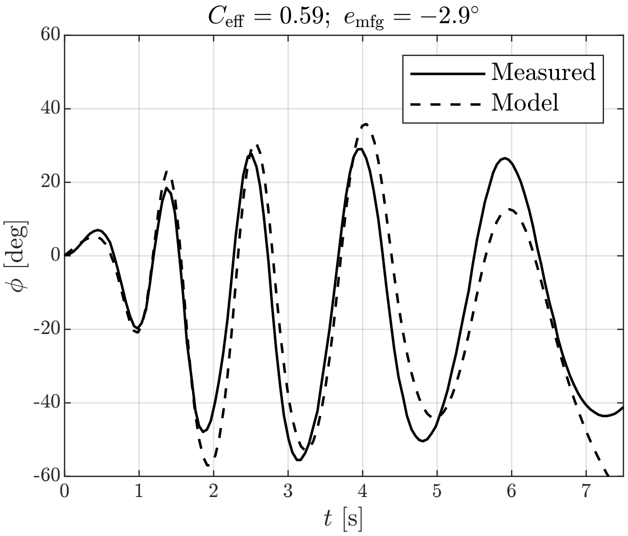

Figure 1: Measured vs. modeled roll for C_{eff}=0.59 and e_{mfg}=-2.9°.

The initial report on Xanthus' roll control declined to provide an explanation for the discrepancy between prediction and observation. It was later determined that the primary cause of this issue was the calculation of $C_{L\alpha}$. The initial calculation of the tabs' lift coefficient was done using CFD on a single fin in a freestream. This is not a good idea. Refined estimates using CFD on the rocket in its entirety provide $C_{L\alpha}$ that are closer to the observed values.

3 Controller Design

Detailed methodology behind the controller design for Project Xanthus is covered in [1]. Figs. 2–3 are the relevant stability plots. There were minor changes in the rocket's inertia. An efficacy constant of 0.5 was used to adjust the plant to reflect the findings of Xanthus flights I and II. K_p was set to 0.260 and T_d was kept at 0.20. This yielded a healthy gain margin of 4.59 and phase margin of 41.8°.

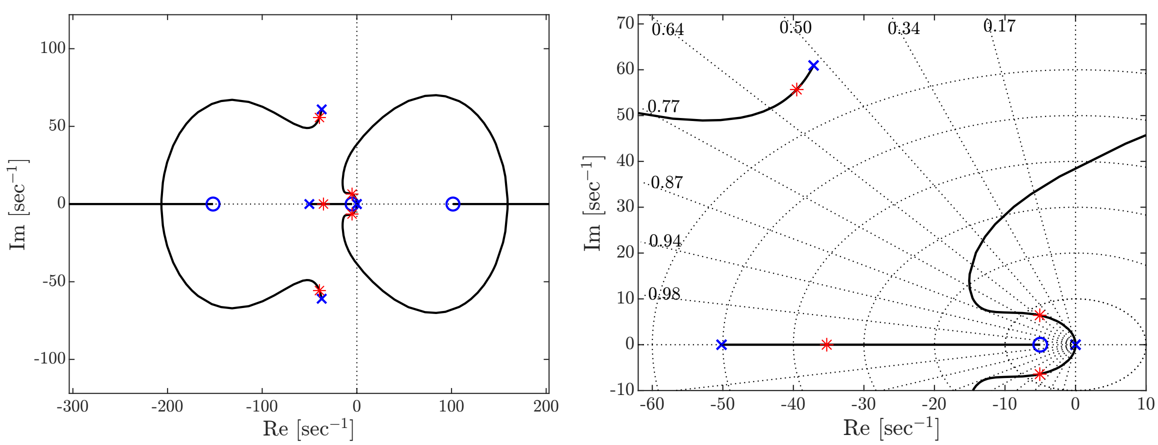

(a) **(b)**

Figure 2: 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 left and right zeros and all other poles are from the servo dynamics. The center zero is due to the controller.

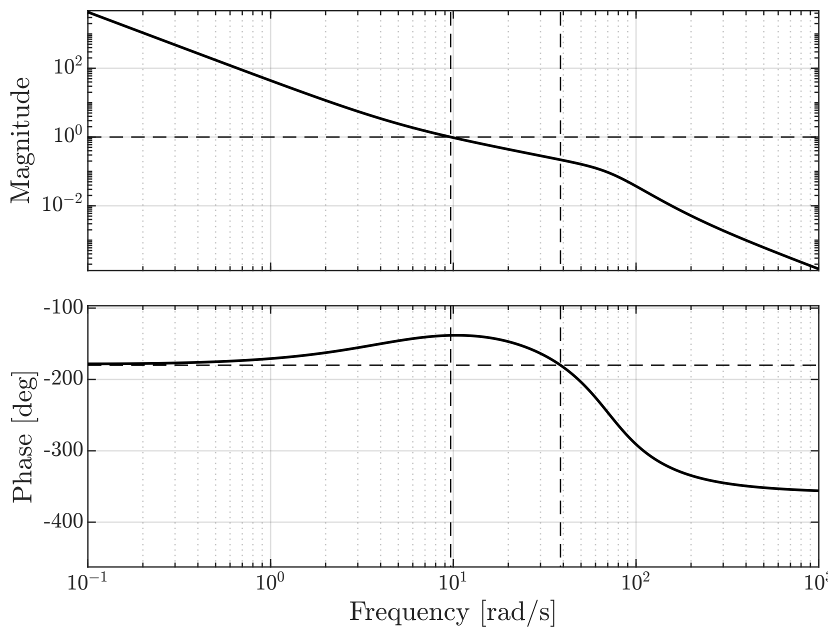

Figure 3: Bode plot indicating crossover frequencies and margins. Gain margin is 4.59 with $|L| = 1$ at 1.54 Hz. Phase margin is 41.79° with $\angle L = -180°$ at 6.14 Hz.

4 Launch Summary

Xanthus' final launch took place on January 17th, 2026 at approximately 11:00 a.m. at URRG. Accelerometer readings indicate an apogee of 460 m (1,501 ft) and maximum speed of 98 m/s (321 ft/s). Unfortunately, a short in the pyrotechnics caused the flight computer to restart after appogee. The parachute failed to deploy as a result, and Xanthus came down ballistically. Although the recovery and avionics bays were demolished, the fin can was mostly unharmed. Fig. 4 depicts the flight profile.

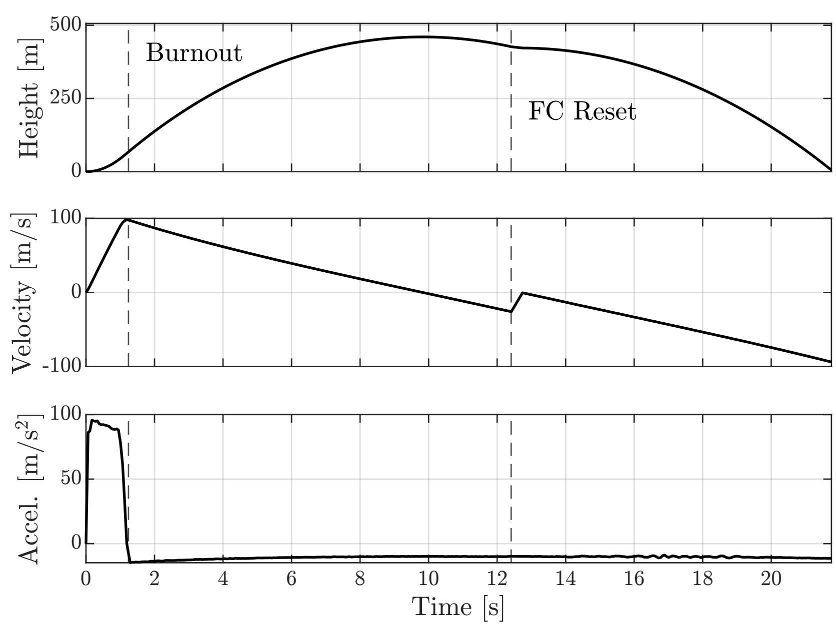

Figure 4: Flight 3 profile until estimated landing. Max height: 460 m. Max velocity: 98 m/s.

Despite Xanthus' tragic end, her final flight saw the realization of the project's aim: roll stabilization. The following section will explore this achievement.

5 Analysis

Flight III provides the first dataset depicting what successful roll control looks like. The roll time series, along with tab command angle and dynamic pressure, is plotted in Fig. 5. As can be seen, the roll angle is kept to within a 25° range, despite a sizable initial disturbance. The reader is encouraged to view camera shell footage from the flight; it will no doubt make more of an impression than the plot in Fig. 5. Unfortunately, only footage from the upwards pointing camera is available, as the downward-facing footage was destroyed in the crash.

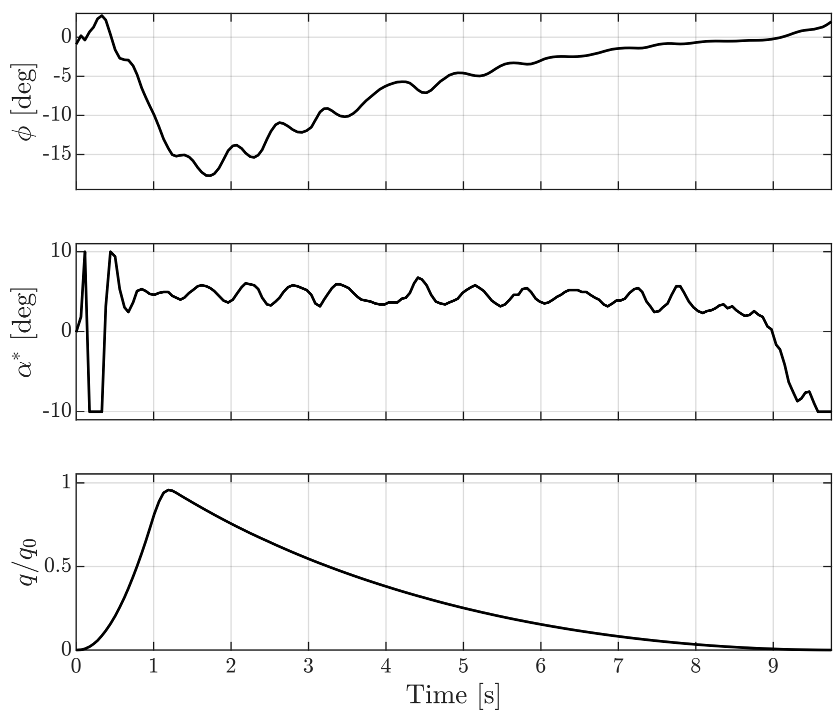

Figure 5: Roll angle, commanded tab angle, and normalized dynamic pressure. After an initial disturbance, the controller brings ϕ back to the setpoint.

The same regression methods described in Section 2 were used to analyze flight III. The aperiodic nature of the third flight (compared to the friendly sinusoid of flight II) made the fitting more difficult. Thus, instead of an objective function based on roll, angular acceleration $\ddot{\phi}$ was used. The best fitting parameters are $C_{\text{eff}} = 0.45$ and $e_{\text{mfg}} = -4.7°$. This efficacy constant is in the ballpark of those from the second flight analyses. Perhaps the larger manufacturing error can be attributed to damages from weathering two previous launches. The best fitting model is plotted against actual angular acceleration in Fig. 6.

6 Conclusion

Overall, Xanthus was the first thrice-launched flight vehicle in the MIT Rocket Team's history, the first with successful active control, and one of many to crash land. The project was invaluable in developing expertise in controls, system identification, modeling, mechanical design, and analysis. Beyond Aerodynamics and Control, Project Xanthus exposed the Avionics and Structures subteams to several new problems and some new solutions. And it left the team in a good position to continue pursuing Project Zephyrus. The details

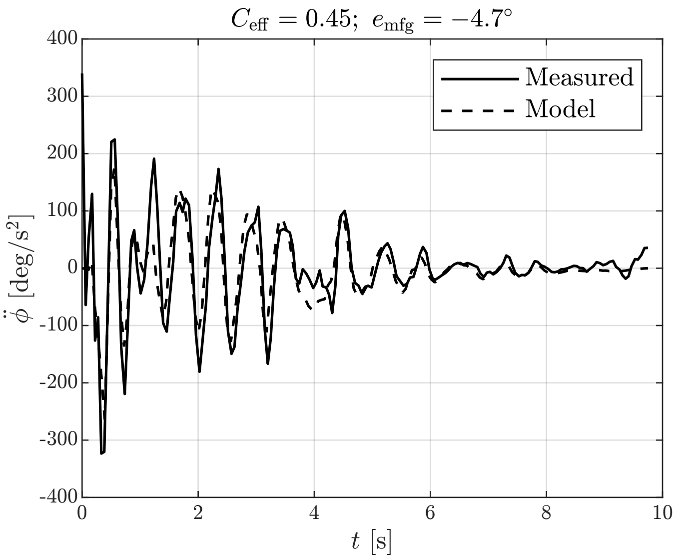

Figure 6: Measured vs. modeled roll for C_{eff}=0.45 and e_{mfg}=-4.7°.

(and shortcomings) of the latter can be found in another report. This engineer looks forward to a future where every MIT rocket has pristine, spin-free onboard footage.

References

[1] C. Sterling, “Aerodynamic Roll Control on Project Xanthus,” 2025.