An Architecture for Trusted Long-Duration Satellite Autonomy

Jeremy Ludwig[1] and Daniel Tuohy[2]
Stottler Henke Associates, San Mateo, CA, 94402

Hanspeter Schaub[3] and Anaïs Cheval[4]
University of Colorado Boulder, Boulder, CO 80309

Sean Phillips[5] and Kristina Miller[6]
Air Force Research Laboratory, Kirtland AFB, NM, 87117

Traditional satellite operations, which rely heavily on ground monitoring and command and control, face growing limitations. Onboard satellite autonomy offers a promising solution to these challenges. By enabling spacecraft to navigate, execute tasks, monitor health, and analyze data independently, autonomous systems can enhance mission effectiveness, improve resilience, and reduce operational costs. Recent progress has advanced individual capabilities such as autonomous planning, fault management, and decision-making. However, integrating these elements into a cohesive, adaptable framework for long-duration satellite operations remains an open challenge. This paper introduces preliminary work on the Lifetime Autonomy Software for Spacecraft Operations (LASSO) autonomy architecture, specifically designed to address this gap. The contribution of this paper is to present work from the initial development phase, outline future enhancements, and demonstrate LASSO's potential to impact satellite operations management.

I. Introduction[*]

Traditional satellite operations, heavily reliant on ground monitoring and command streams, face increasing limitations in the context of modern space missions [1]. The scale, complexity, and distance inherent in contemporary Earth observation, global communication constellations, and deep space exploration challenge this operational model. Significant communication delays, particularly for interplanetary missions, preclude effective real-time control, while the management of large, dynamic constellations necessitates rapid, localized decision-making far exceeding ground capabilities. Onboard satellite autonomy emerges as a solution to these challenges [2]. By enabling spacecraft to navigate, monitor health, execute tasks, and potentially analyze data independently, autonomy enhances mission effectiveness, increases resilience, and reduces operational costs.

Recent work in satellite autonomy reflects a growing trend towards enhancing spacecraft independence and resilience. This encompasses advancements across several critical domains. In autonomous mission planning, the focus is on enabling satellites to independently schedule and execute tasks, optimize resource allocation, and

[1] Director, AIAA Member.
[2] Senior AI Software Engineer.
[3] Distinguished Professor and Department Chair, Aerospace Engineering Sciences Department, University of Colorado, AIAA Fellow
[4] Graduate Research Assistant, Aerospace Engineering Sciences Department, University of Colorado, Student Member
[5] Technology Advisor, Space Control Branch, AIAA Member
[6] Research Aerospace Engineer, Space Control Branch
[*] Approved for public release; distribution is unlimited. Public Affairs release approval # AFRL-2025-2269

coordinate with other spacecraft to achieve assigned objectives (e.g., [2], [3], [4], [5]). Another critical area is autonomous fault detection and recovery, which aims to equip satellites with the ability to identify, diagnose, and resolve anomalies without ground intervention. Recent efforts in this field leverage machine learning and model-based approaches to enhance the reliability and longevity of satellite missions (e.g., [6], [7], [8], [9]). Furthermore, onboard autonomous decision-making is driven by the need for real-time responses and efficient data handling, especially in deep space missions and large constellations. This includes developing software capable of processing data onboard and making intelligent choices when unexpected issues occur (e.g., [1], [2], [10], [11]). Finally, with the increasing complexity of autonomous systems overall, explainability is emerging as a vital aspect, aiming to provide transparency and understanding into the decision-making processes of AI-driven satellite operations [12].

While recent work has significantly advanced all these areas (autonomous planning, fault management, decision-making, and explainability), effectively integrating these capabilities into a cohesive, adaptable framework for long-duration satellite operations remains an open challenge. While there are existing frameworks such as Modular Autonomous Systems Technology (MAST) [13], there are no existing adaptable framework that brings together these diverse research threads to support satellite operations over long durations without operator intervention and that are designed to be applicable across a variety of satellite busses.

This paper introduces preliminary work on the Lifetime Autonomy Software for Spacecraft Operations (LASSO) autonomy architecture, specifically designed to address this integration gap. To achieve this, LASSO aims to leverage the interaction between key onboard autonomy capabilities. Given high-level mission goals, it automatically schedules spacecraft bus and payload tasks while continuously monitoring telemetry for anomalous conditions. Upon detecting fault conditions, LASSO takes corrective actions intended to prevent unnecessary entry into safe mode, thereby enabling continued payload operation where possible. This functionality relies on coordinating mission planning (selecting and sequencing operations), fault detection (identifying anomalies), and decision-making (reacting to real-time conditions and determining appropriate responses). Key design goals guiding LASSO's development include broad applicability across diverse satellite platforms with minimal customization and adaptability to evolving satellite health and performance characteristics over its lifetime. The contribution of this paper is to present work from the initial development phase, outline future enhancements, and demonstrate LASSO's potential to impact satellite operations management.

II. Problem Statement

This paper adresses the challenge of designing, implementing, and demonstrating specific onboard autonomous decision-making capabilities towards an autonomy architecture that increases payload operational time and/or reduces the need for continuous operator monitoring and intervention. The decisions expected to be made by the autonomy architecture fall into several categories:

- Optimizing Routine Tasks: Automating nominal operations such as determining optimal battery charging schedules based on predicted power generation and consumption, or planning attitude adjustments to maximize payload data acquisition opportunities while respecting spacecraft constraints.
- Detecting and Responding to Slow Degradations: Identifying gradual performance declines that might otherwise go unnoticed until they cause significant issues. An example includes monitoring battery cell voltage trends over extended periods to recognize declining capacity and potentially adjust operational strategies proactively.
- Anomaly Detection and Response: Rapidly identifying and reacting to sudden off-nominal events. Examples include responding to an unexpected drop in solar panel charging capability, detecting components operating outside their specified temperature limits, and recognizing degraded reaction wheel performance.

A critical operational requirement is the avoidance of unnecessary entry into spacecraft Safe Mode. While Safe Mode protects the vehicle, it typically halts payload operations, potentially leading to days or weeks of downtime for recovery and reconfiguration. Therefore, a primary objective is to mitigate issues and maintain operational capability without resorting to Safe Mode unless absolutely necessary for spacecraft survival.

Autonomous decisions must operate within constraints acceptable to human operators, particularly concerning irreversible actions or those with significant, unpredictable side effects. For instance, autonomously deciding to permanently switch to a backup reaction wheel is deemed outside the scope of available decisions. Only actions that are reversible or have well-understood, limited consequences are allowable. Beyond the core decision-making capabilities, the project requirements emphasized the development of robust visualization tools and explainable AI methods. These elements are crucial for enabling human operators to understand the autonomous system's state, reasoning, and actions, thereby calibrating trust and facilitating the eventual transition to operational use. The autonomy engine is required to provide clear justifications for its decisions. Finally, the architecture should be applicable to diverse satellite platforms and payloads with minimal customization.

III. Methods

As an initial step towards addressing this problem statement, a LASSO proof of concept was developed that focused on simulating the core interactions between spacecraft dynamics, mission planner, and anomaly detection modules, as well as design work towards an autonomy executive for high-level decision making as shown in Fig. 1.

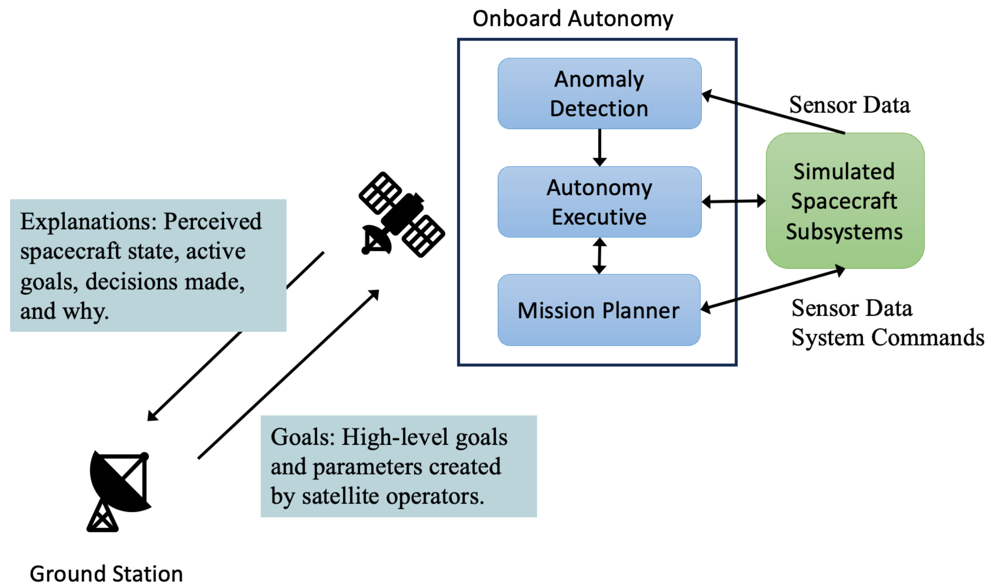

Fig. 1 Proof of Concept Overview.

The spacecraft subsystems and orbital environment were simulated using Basilisk (BSK), a high-fidelity, physics-based simulation framework [14]. Key bus components relevant to autonomy decisions were modeled, including solar panels for power generation, a battery model tracking state of charge, and reaction wheels for attitude control. The simulated satellite bus characteristics (mass, size, power) were representative of a small satellite (330kg). The simulated payload was an Earth-observing optical instrument performing "striping" operations – sequential imaging along a defined ground track. Both the instrument apertures and the solar panels were modeled as fixed to the spacecraft body, meaning attitude changes were required for both pointing the payload and orienting the solar panels, creating inherent resource conflicts for the scheduler to manage.

To test the system's responsiveness, several types of anomalies were introduced into the BSK simulation environment: Increased friction in the reaction wheels, affecting attitude control performance and power draw; reduced power draw limit, simulating a fault where the maximum power available to actuators (like reaction wheels) is less than expected; encoder errors, introducing inaccuracies in the measurement of reaction wheel speeds; and reduced battery capacity, simulating cell degradation over time. Of these, only the first was tested during the initial work.

Typical simulation runs involved initializing the spacecraft state and then propagating its dynamics under the control of the Mission Planner for a complete orbit (approximately 2-3 hours of simulated time). Telemetry data streams (e.g., battery voltage, wheel speeds, power levels, attitude errors) were generated throughout the simulation for both nominal runs and runs where specific anomalies were injected. These telemetry data files served as the primary integration point with the anomaly detection system, which trained on some initial data and then analyzed the remaining runs looking for anomalies. The generated dataset contained 50 nominal orbital scenarios and 3 off-nominal cases featuring elevated reaction wheel friction. The off-nominal cases included a friction coefficient of 0.05, approximately 10x higher than manufacturer specifications for new components. This level of friction is still relatively modest in terms of its impact but interferes with the efficiency of attitude adjustments sufficiently to cause some imaging tasks to fail. Each dataset captured a complete orbit with various imaging tasks, including comprehensive telemetry on reaction wheels, batteries, translational state, attitude state, and task-specific attitude errors.

Mission Planning was implemented using a policy learned offline via reinforcement learning, leveraging the BSK-RL toolkit [15]. This policy aimed to maximize the completion of payload striping tasks, given the current spacecraft state (attitude, battery charge, wheel saturation). The available actions for the scheduler included pointing towards the Sun for battery charging, performing reaction wheel desaturation maneuvers, initiating data downlink, and commanding an imaging strip. The learned policy implicitly balanced task completion with maintaining the spacecraft's health and safety (e.g., sufficient battery charge, manageable wheel speeds).

The **Anomaly Detection** component analyzes the generated telemetry streams, looking for anomalous conditions based on a learned pattern of life that updates over time. The initial anomaly detection component relied on two existing anomaly detection algorithms: TRIAD (statistics based) [16] and ADTM (machine-learning based) [17]. A key objective of this initial work was to evaluate the performance of these distinct anomaly detection approaches on the simulated data, using the results to guide the design of the integrated anomaly detection subsystem in future work.

The **Autonomy Executive** component for decision making was design-only for this initial work. Its envisioned functions include decision making and communication and explainability. Decision making involves synthesizing inputs from the scheduler (task success/failure), anomaly detection (identified anomalies), and internal satellite models (model-based reasoning) to make higher-level decisions aimed at prolonging payload operation in spite of anomalies. For example, if the scheduler reports task failures concurrently with an anomaly detection alert in the attitude control system, the executive might decide to switch to a less demanding operational mode or attempt a mitigation action if available and safe. Communication & explainability involves interfacing with operators, likely through existing ground system software (e.g., OpenC3 Cosmos), to provide situational awareness. This includes reporting planned tasks, current system health, detected anomalies, and crucially, the decisions made by the autonomy and the rationale behind them.

IV. Results

Results from the initial work demonstrate the feasibility of autonomous scheduling and anomaly detection. First, the Mission Planner based on BSK-RL effectively managed spacecraft actions to optimize task completion under varying conditions. Second, as part of the anomaly detection component, TRIAD exhibited strong performance in distinguishing nominal from anomalous conditions. While ADTM did not detect the same anomaly, it provided valuable insights into multivariate anomaly detection and the value of ensemble methods. Third, a design of an autonomy executive for high-level decision making and communication was created, employing model-based reasoning (MBR). The design shows promise in providing directives to the mission planner for mitigating identified issues while at the same time explaining the reasons behind these directives to the human operator.

A. Striping Task and Simulation

The simulation models a large, agile, high-resolution Earth-observing satellite designed for precise imaging tasks, rather than a small constellation satellite. The specific task simulated is "striping," which involves continuous imaging along a defined ground path using a Time-Delay Integration (TDI) pushbroom scan line camera. Achieving acceptable image quality with this type of camera requires strict attitude control to keep the sensor's projection perpendicular to the scanning direction and maintain a constant velocity matching the desired acquisition speed relative to the Earth's surface.

To implement this in BSK, specific attitude guidance modules were developed. These modules calculate the necessary time-varying attitude and rate commands (references) for the satellite to effectively "track" the target strip on the ground, similar to tracking a moving target. This process includes an essential pre-imaging slew maneuver to ensure the satellite achieves the required pointing accuracy before starting the image capture. These guidance references are then used by BSK's feedback control system to command the satellite's attitude control actuators (like reaction wheels). Fig. 2 illustrates the control and guidance system as the spacecraft attempts to image a strip from Boulder to Phoenix.

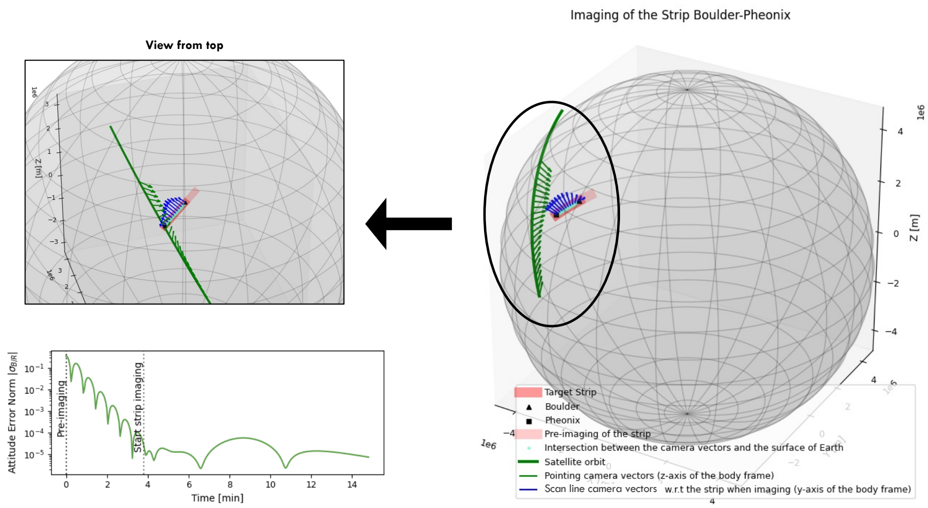

Fig. 2 Illustration of the Striping Task.

B. Mission Planner

The mission planning agent applies a policy, learned offline using BSK-RL, that maximizes completion of striping tasks given the current state. The image striping problem is formulated as a Markov Decision Process (MDP). The action space of the agent is pointing to the Sun to charge the battery, desaturating the reaction wheels, downloading data, and imaging a strip. The policy attempts to maximize completion of striping tasks while maintaining the health and safety of the spacecraft.

The agent was trained using the Proximal Policy Optimization (PPO) algorithm over simulated multi-orbit episodes. In each episode, the agent was presented with 500 randomly generated potential imaging strips (varying in length from 100-1500km and priority 0-1). The pre-imaging time and acquisition speed are set to 3 minutes and 3 km/s, respectively, for all strips. Observing the next twenty accessible strips, the agent's goal was to choose which strip to image next to maximize the total priority value accumulated during the episode or to perform one of the non-imaging actions.

Preliminary results indicated successful learning as shown in Fig. 3. The trained agent developed a policy that prioritized imaging strips with higher assigned priority values. Interestingly, it also learned to avoid selecting longer strips, a behavior attributed to the reward being based solely on priority, making longer strips less efficient for maximizing cumulative reward in this training setup.

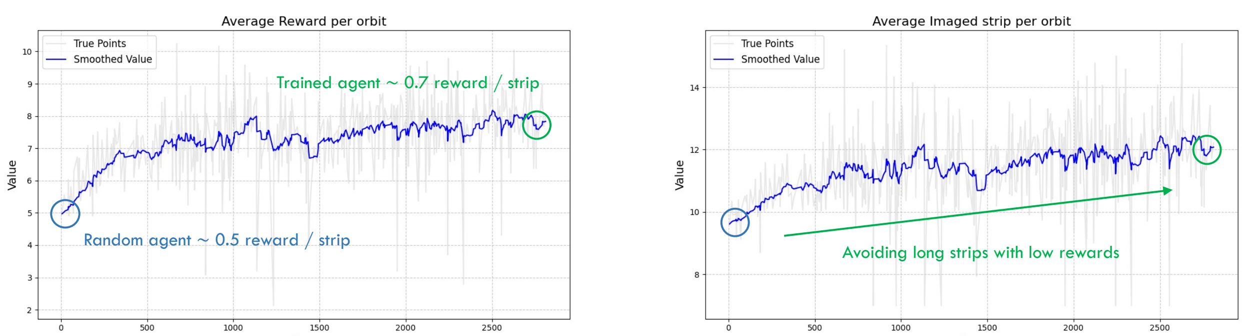

Fig. 3 Preliminary Training Results.

C. Anomaly Detection

Our initial investigation focused on using TRIAD and ADTM to analyze simulated satellite telemetry data generated while a version of the mission planner carried out the striping task. The dataset comprised fifty nominal orbital scenarios and three off-nominal cases featuring elevated reaction wheel friction. The off-nominal cases included a friction coefficient of 0.05, approximately 10x higher than manufacturer specifications for new components. This level of friction is still relatively modest in terms of its impact, but it does interfere with the efficiency of attitude adjustments sufficiently to cause some imaging tasks to fail. Each dataset captured a complete orbit with various imaging tasks, including comprehensive telemetry on reaction wheels, batteries, translational state, attitude state, and task-specific attitude errors.

TRIAD generates *thermodynamic variables*, which are mathematical functions calculated over recent data windows that summarize expected system behavior. Fault-free training data is used to establish normal operating ranges for detecting anomalies, and online training can be used to continually update the definition of expected behavior. TRIAD demonstrated positive performance in anomaly detection for this case, reliably identifying true positives while generating only sparse, low-density false positives as shown in Fig. 4. In this figure, the X-axis is time and the Y-axis is the number of standard deviations above/below the maximum recorded statistic learned from the nominal training data. False positives were rare in nominal scenarios and, when detected, were always sparse (see left graph). True positive anomalies were consistently detected across the examples, and these detections were characterized by dense signatures primarily involving 'variance within a small window' and 'average frequency on the Fourier spectrum' statistics (see right graph). Further sensitivity tuning eliminated false positives entirely while maintaining detection capability, suggesting separability between nominal and off-nominal cases.

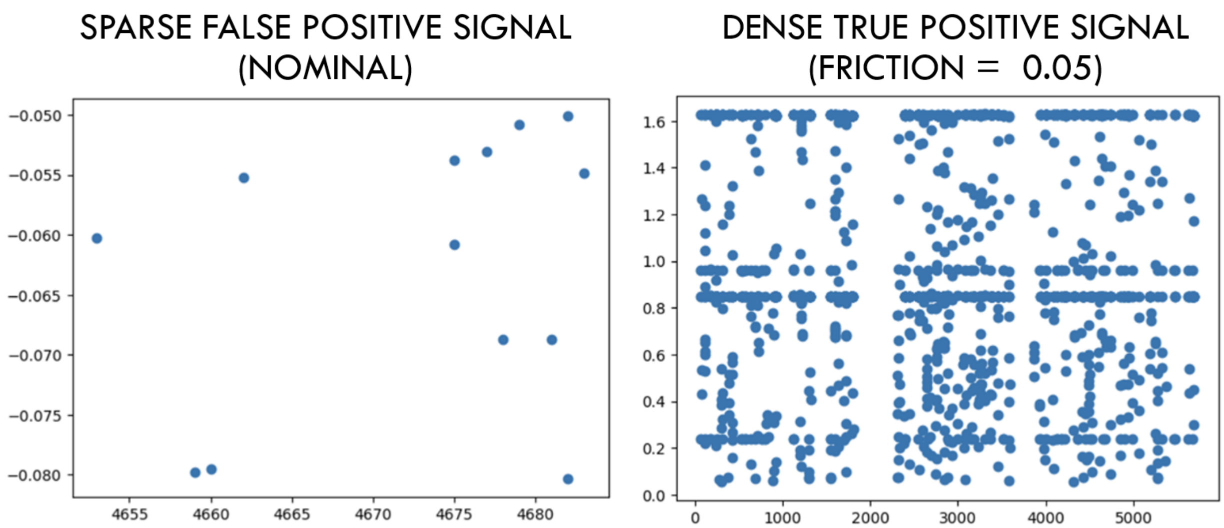

Fig. 4 Anomaly Detection Results, where X-axis is Time and Y-axis is a Measure Related to Standard Deviations from Expected Nominal Distributions.

ADTM leverages Self-Organizing Map (SOM) Neural Networks trained on nominal subsystem data, which do not require faulty or hand-labeled data. ADTM's multivariate analysis, while detecting some anomalies, proved less reliable than TRIAD for this specific use case. We attribute this performance difference to the temporal nature of friction-induced anomalies, which manifest as patterns over time rather than as instantaneous multivariate deviations. While ADTM did not successfully detect this particular anomaly, we still believe it has a role to play in future work as part of an ensemble of anomaly detection algorithms within the anomaly detection component.

D. Autonomy Executive

The envisioned Autonomy Executive employs model-based reasoning (MBR) to make critical decisions regarding spacecraft bus operations, aiming to prolong payload functionality despite off-nominal conditions. Specifically, the executive is responsible for interpreting findings from the anomaly detection system, using MBR to predict potential effects, and subsequently selecting appropriate mitigation strategies. These strategies are then communicated to the mission planner, typically as revised operational constraints or activated policies, which guide subsequent planning. Analogously, this is like selecting an all-wheel drive mode in a car, where the system enacts significant internal policy changes (e.g., power distribution, traction control) to keep the vehicle moving safely and effectively in adverse conditions, allowing the driver (representing the higher-level mission objectives) to proceed with minimal disruption.

The executive relies on several types of data in addition to the anomaly detection results. It receives data from the satellite bus and payload, including outcome performance metrics (e.g., task success/failure), behavioral performance metrics (e.g., attitude error), and raw telemetry. The executive is built around satellite anomaly mitigation model like the notional one shown in Fig. 5, which relates specific causes to generic possible symptoms and identifies policies/modes of operation that can mitigate each. In this case, the anomaly detected in the reaction wheel points to an issue with attitude control in general. Given this, there are two policies to choose from that would not have irreversible side effects: either to switch to a backup wheel or pursue a task-reduction policy where the payload selects tasks that do not require such agile attitude adjustments. These policies include a priority number, which is used to select a single action when multiple apply. It is envisioned that the model would be enhanced with actions that could be taken to help narrow down the cause, which would be of higher priority than mitigation actions.

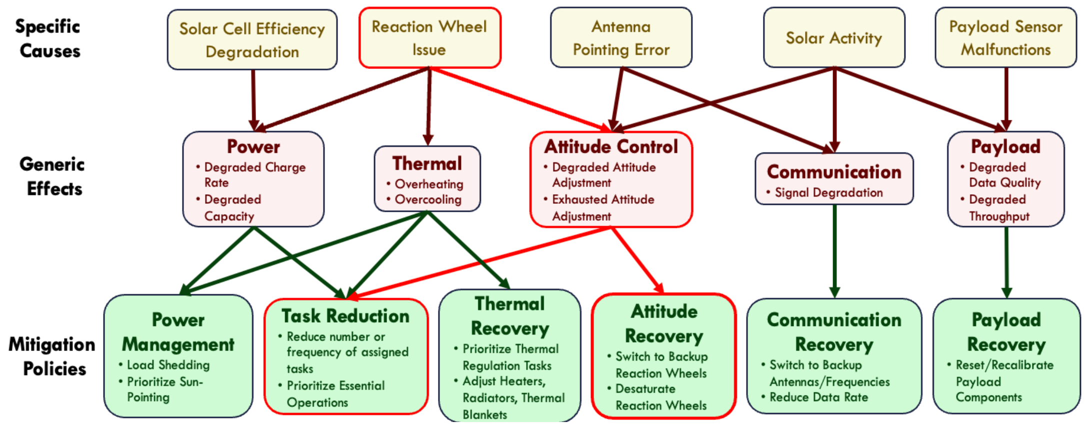

Fig. 5 Satellite Anomaly Mitigation Model.

Activation of a specific mitigation policy directly influences the mission planner's subsequent operations. The planner must then adapt its strategy, potentially by generating a revised set of actions and schedule or operating under adjusted constraints dictated by the policy. This adaptation might involve switching between specialized planning agents pre-trained for different mitigation states, or by utilizing the active policy information as input to a single adaptive planning agent. Regardless of the implementation, such mitigation actions are expected to substantially impact the resulting mission plan, while the system maintains the capability to explain these autonomous adjustments to the human operator as shown in Fig. 6.

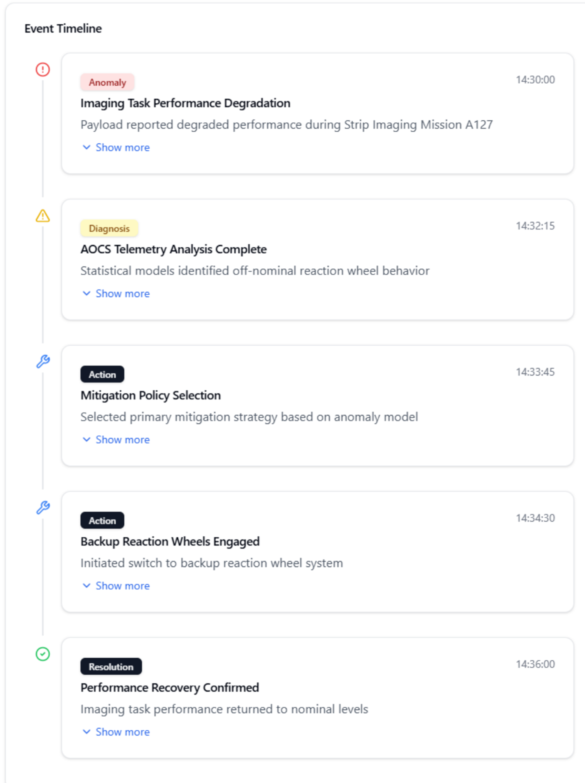

Fig. 6 UI Mockup for Reaction Wheel Event

V. Conclusion and Future Work

The LASSO autonomy architecture holds significant potential to address the challenges of long-duration, operator-free satellite missions. By automating routine tasks, proactively responding to anomalies, and adapting to changing satellite characteristics, an architecture such as LASSO can enhance mission life, improve operational efficiency, and enable operators to focus on broader strategic objectives. Future research, development, and validation efforts will pave the way for eventual transition into operational use.

Building on the initial success, future development aims to create a complete prototype of the LASSO architecture. Key enhancements include: (1) improved simulation capabilities to generate more realistic telemetry data, model general bus and payload interfaces, and incorporate additional failure types; (2) enhanced mission planner for continuous imaging tasks and decentralized constellation scheduling; (3) expanded anomaly detection research, including the evaluation of open-source libraries and the integration of temporal behavior patterns into ADTM; (4) development of the autonomy executive using model-based reasoning to make high-level decisions for prolonging payload operation; (5) design and implementation of visualization tools to effectively communicate decisions and rationale to operators, fostering trust in the autonomous system; and (6) comprehensive system integration and testing, including long-term performance characterization through simulated lifespan excerpts (1–10 years).

References

[1] Joseph Melville, Andrew Narvaez, and Lee Jasper, “A Review of Automation in Small Satellite Operations,” presented at the IEEE Aerospace Conference, Big Sky, MT, 2025.

[2] R. Ravishankar and O. L. de Weck, “Onboard Computing for Autonomy and Decision Making in Satellite Systems: An Overview,” in *AIAA SCITECH 2025 Forum*, in AIAA SciTech Forum. , American Institute of Aeronautics and Astronautics, 2025. doi: 10.2514/6.2025-1433.

[3] M. Stephenson, L. Mantovani, and H. Schaub, “Intent Sharing for Emergent Collaboration in Autonomous Earth Observing Constellations,” in *AAS Guidance and Control Conference, Breckenridge, CO*, 2024, pp. 24–192.

[4] M. Stephenson, L. Mantovani, S. Phillips, and H. Schaub, “Using Enhanced Simulation Environments to Accelerate Reinforcement Learning for Long-Duration Satellite Autonomy,” in *AIAA SCITECH 2024 Forum*, Orlando, FL: American Institute of Aeronautics and Astronautics, Jan. 2024. doi: 10.2514/6.2024-0990.

[5] S. Hilton, K. Thangavel, A. Gardi, and R. Sabatini, “Intelligent Mission Planning for Autonomous Distributed Satellite Systems,” *Acta Astronautica*, vol. 225, Sep. 2024, doi: 10.1016/j.actaastro.2024.08.050.

[6] G. Danke and R. Iyer, “An Autonomous Reinforcement Learning Framework for Fault Recovery and Mission Replanning on CubeSats,” *Small Satellite Conference*, Aug. 2024, [Online]. Available: https://digitalcommons.usu.edu/smallsat/2024/all2024/108

[7] M. Hedayati, A. Barzegar, and A. Rahimi, “Fault Diagnosis and Prognosis of Satellites and Unmanned Aerial Vehicles: A Review,” *Applied Sciences*, vol. 14, no. 20, Art. no. 20, Jan. 2024, doi: 10.3390/app14209487.

[8] J. Murphy, J. E. Ward, and B. Mac Namee, “An Overview of Machine Learning Techniques for Onboard Anomaly Detection in Satellite Telemetry*,” in *2023 European Data Handling & Data Processing Conference (EDHPC)*, Oct. 2023, pp. 1–6. doi: 10.23919/EDHPC59100.2023.10396403.

[9] S. Cuéllar, M. Santos, F. Alonso, E. Fabregas, and G. Farias, “Explainable anomaly detection in spacecraft telemetry,” *Engineering Applications of Artificial Intelligence*, vol. 133, p. 108083, Jul. 2024, doi: 10.1016/j.engappai.2024.108083.

[10] K. Thangavel *et al.*, “Artificial Intelligence for Trusted Autonomous Satellite Operations,” *Progress in Aerospace Sciences*, vol. 144, p. 100960, Jan. 2024, doi: 10.1016/j.paerosci.2023.100960.

[11] C. Adams, B. Kempa, M. Iatauro, J. Frank, and W. Vaughan, “An overview of distributed spacecraft autonomy at nasa ames,” presented at the 37th Annual Small Satellite Conference, 2023.

[12] K. L. Hobbs *et al.*, “The Safe Trusted Autonomy for Responsible Space Program,” Jan. 10, 2025, *arXiv*: arXiv:2501.05984. doi: 10.48550/arXiv.2501.05984.

[13] J. M. Badger, P. Strawser, and C. Claunch, “A distributed hierarchical framework for autonomous spacecraft control,” in *2019 IEEE Aerospace Conference*, IEEE, 2019, pp. 1–8.

[14] P. W. Kenneally, S. Piggott, and H. Schaub, “Basilisk: A Flexible, Scalable and Modular Astrodynamics Simulation Framework,” *Journal of Aerospace Information Systems*, vol. 17, no. 9, pp. 496–507, Sep. 2020, doi: 10.2514/1.I010762.

[15] M. A. Stephenson and H. Schaub, “BSK-RL: Modular, High-Fidelity Reinforcement Learning Environments for Spacecraft Tasking,” in *75th International Astronautical Congress, Milan, Italy, IAF*, 2024.

[16] R. Stottler *et al.*, “Modular AI for Faults: Local Watch and Efficient Response,” presented at the IEEE Aerospace Conference, Mar. 2025.

[17] R. Stottler, S. Ramachandran, C. Healy, A. Singhal, and E. Finnigan, “On-board, autonomous, hybrid spacecraft subsystem fault and anomaly detection, diagnosis, root cause determination, and recovery,” in *Proceedings of Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, 2022.