Autonomous, Hybrid Space System Fault and Anomaly Detection, Diagnosis, Root Cause Determination, and Recovery

Dick Stottler, Abhimanyu Singhal, Christopher Healy, and Dr. Sowmya Ramachandran, PhD
Stottler Henke Associates, Inc.
Kerry Quinn, Joseph Palmieri, and Shawn Logan
Astrobotic Technology

ABSTRACT

Bottom line up front (BLUF): our MBR and ML techniques correctly identified critical spacecraft subsystem anomalies in milliseconds in a wide variety of scenarios, while avoiding false alarms. An important component of Space Situational Awareness (SSA) / Space Domain Awareness (SDA) is knowledge of the true status of friendly assets and whether any assets are under attack. Therefore, it is important to be able to detect faults and other anomalies, and determine the components involved and the root cause as well as whether that root cause is likely an external attack. During space conflict, communications to satellites may be disrupted, requiring them to intelligently and autonomously "take care of themselves," i.e., effectively detect faults, diagnose their root causes, and develop and execute recovery plans, autonomously, without necessarily being able to communicate with ground controllers. This lack of communication is analogous to lunar rovers and power systems where communication can be disrupted by terrain and other factors or take too long for some catastrophic anomalies. Astrobotic, for NASA, is developing a rover, Vertical Solar Array Technology (VSAT), that traverses over the lunar surface to an advantageous position, then unfurls its 60' high Roll Out Solar Array (ROSA) photovoltaic mast to provide power for other lunar systems. Given the height of the deployed/deploying ROSA, the rover is very unstable and must adhere to very strict leveling requirements to avoid tipping over, with catastrophic loss of mission, even just a few degrees of lean would be disastrous. Prior to ROSA deployment a gimbal levels the base of the ROSA and then locks. As the solar array is unfurled inertial measurement units (IMUs) continuously monitor the array's movement, including any lean, force sensors monitor the force on each of 4 wheels, and several side-facing and upward facing cameras observe the events. A problem during or after ROSA deployment may be very dynamic, denying ground controllers enough time to correct any problem, given the round trip communication delays. It is therefore important that the VSAT be equipped with the means to quickly detect problems, perform diagnosis and root cause determination, and quickly safe the system.

Traditionally, Fault Detection, Isolation, and Recovery (FDIR) systems have utilized Model Based Reasoning (MBR), which requires knowledge of the subsystem design and the behavior of components down to the desired level of diagnosis. To the degree this information is readily available, it is important to make good use of it. However, the field of machine learning (ML) has shown that systems can also learn, offline, the normal behavior of complex systems in many different environments and states, and then detect abnormal behavior in real time. These systems can also be trained with known abnormal states, and recognize these more specifically when they occur. With the new types of VSAT subsystems (such as mechanical components and related sensors) came new challenges to be overcome. Some concerns included quick reaction times needed to avoid tipping or buckling during mast deployment and, at the opposite end of the spectrum, detecting very gradual changes, hard to discern in sensor noise (the mast moves very, very slowly while tracking the Sun). In some cases, data is severely limited, reducing the applicability of a pure ML approach. In our previous work, we outlined our modular approach to fault detection and diagnosis utilizing MBR and ML as well as a third independent method called the Thermodynamic Reasoning and Intelligent Anomaly Detection (TRIAD) system. Similarly to how aggregate variables such as thermodynamic variables such as pressure and temperature can summarize microstate variables (e.g. the speeds of individual molecules), TRIAD utilizes aggregate quantities such as mean, minimum, maximum, and Fourier Transforms to characterize anomalies. We also described how this hybridization enables additional confidence in diagnosis, as the advantages of each approach are emphasized while the disadvantages are mitigated, and summarized how we planned to apply these methods on the VSAT platform and subsystems. This paper first quickly reviews the concepts then describes progress on this work since our last paper, presented at AMOS 2023. This includes validation of the hybrid approach to fault detection, diagnosis, and recovery via a physical simulation of the VSAT

platform as well as results from multiple fault detection modules. We enumerate many relevant scenarios, developed in conjunction with Astrobotic to best capture realistic, critical faults, as well as metrics from our approach. We show that the previously discussed methods are capable of both detecting and characterizing mechanical anomalies from simulated VSAT telemetry data within tens of milliseconds of the faults occurring, well below the allotted "reaction time" of 100 milliseconds. The paper will present quantitative results for a large range of fault scenarios, including soil collapse, soil slippage, ROSA levelling errors, and a wide variety of sensor faults. Both MBR and TRIAD were effective at detection and diagnosis and, as mentioned in our previous work, we identified several areas where hybridization of both techniques provides a significant advantage over the use of just one or the other. We will also present data gathered during slope testing of the actual ground VSAT physical prototype with some fore-shadowing of what we plan to present next year, based on testing with this (and more) actual data from the actual hardware prototype. We conclude with a discussion on the direction this work will take in the future. Based on these results, Astrobotic plans to include MAIFLOWER on the actual lunar VSAT, with integration beginning this Fall.

1. HIGH-LEVEL MAIFLOWER ARCHITECTURE

Stottler Henke's VSAT anomaly detection and diagnosis system is called MAIFLOWER (Modular Artificial Intelligence for Faults: Local Online Watch and Efficient Response). During normal operations, MAIFLOWER monitors onboard sensor values to automatically characterize subsystem components and to be prepared to detect failures. During a failure scenario, MAIFLOWER very quickly detects and diagnoses the problem, then immediately sends the root cause to VSAT's control system so that it can take immediate actions to stabilize the situation.

Any intelligent, adaptive system must inherently be a closed loop system (i.e., in basic terms: the system must sense what is occurring and make appropriate decisions to take suitable actions). The first part of this "sense-decide-act" loop involves perception, understanding the situation from the raw sensor values. Over long time scales, this in turn involves characterization. The other perceptual function is detecting faults and diagnosing their causes. MAIFLOWER is a specialization of our MEASTRO close loop architecture, shown at right, though MAIFLOWER is focused on the characterization, anomaly detection, and diagnosis, relying on VSAT's control system to "close the loop"

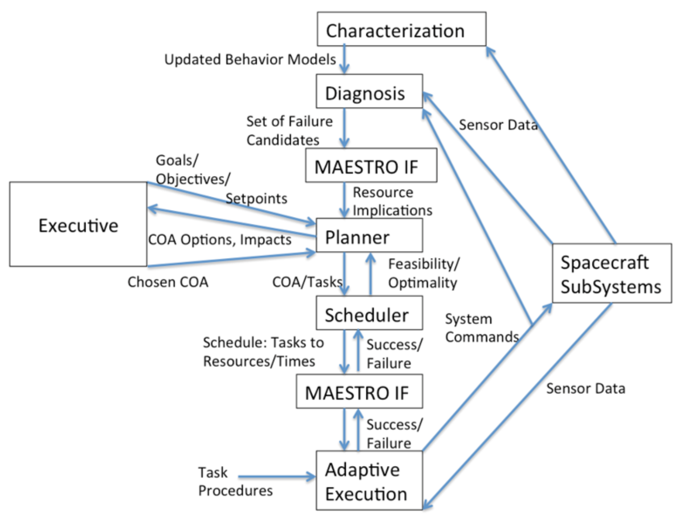

Fig. 1. MAESTRO/MAIFLOWER High-Level Architecture.

Model-based reasoning (MBR) systems are often used to detect and diagnose faults. These systems encode the schematic information of subsystems, which includes the components (including sensors), their normal behavior and known abnormal modes of behavior, and the connections between components. During normal operations, the model is used to simulate the current behavior and compare the simulated sensor output values to the actual sensor outputs. Significant deviations are used to detect some kind of fault. Then the model is used to reason which component faults are most likely to lead to the currently deviating sensor values. The set of possible faults (possibly including sensor faults) which explain the sensor values is the MBR diagnosis engine's output. The process of using the model to diagnose failures is considered somewhat analogous to the reasoning an electrical engineer would employ when using a schematic to try to diagnose the fault. The process can be made more efficient by various heuristics used by spacecraft engineers to quickly diagnose problems and include knowledge of which components are most likely to fail (and how, e.g., mechanical relays tend to fail open while solid state relays tend to fail closed) and/or are the most likely explanation for certain types of sensor values. MAIFLOWER takes advantage of these heuristics. MBR engines may identify one specific fault and/or a set of possible faults. For the most important scenarios (during and after ROSA deployment) the VSAT rover body should be stationary. Therefore, an anomaly is anything that indicates movement in the data of any sensor (including any optical flow) in the cameras. And because the rover is a rigid body with its wheels and gimbal locked, any movement will be seen in multiple sensors and with given mathematical relationships between the deviations across sensors, based on the type of movement (linear versus angular displacement) occurring.

In addition to MBR, MAIFLOWER leverages an expansion of the TRIAD system developed within MAESTRO to synthesize an extensive model-free module to detect and diagnose faults. TRIAD has shown perfect performance in fault detection and classification across fault types in our project to develop for NASA, an intelligent Astronaut Agent based on the MAESTRO system. Beyond this, the TRIAD system constitutes an overarching framework that can intelligently incorporate and synthesize any set of additional feature-based anomaly detection systems. Feature-based anomaly detection systems utilize deterministic functions to extract time series of low-dimensional features from incoming streams of sensor data and perform simple threshold-based anomaly detection on the extracted series. Nearly every prominent anomaly detection algorithm can be classified as feature-based anomaly detection, from state-of-the-art deep learning algorithms using Transformer and Convolutional Architectures [4][5][11][16] to more traditional algorithms making use of Hidden Markov Models or Self-Organizing Maps [6][10]. TRIAD is designed to intelligently synthesize these algorithms instead of running them in parallel.

The advantages of TRIAD's synthesis are twofold. First, TRIAD's synthesis is significantly more accurate than any parallel implementation, for purposes of both detection and characterization. Second, TRIAD enables more wholistic modelling of the utility of each method's individual contribution to the accuracy of the system. For example, a computationally slow/expensive anomaly detection method may outcompete cheaper methods, run individually or in parallel, but find itself redundant when incorporated into TRIAD's system. By contrast, a cheaper method may fail to provide any novel detection when compared against already-implemented methods but significantly enhance expressivity when incorporated in intelligent concert with these methods within TRIAD.

MAIFLOWER extends TRIAD with feature functions from cutting-edge time-series anomaly detection algorithms, including the state-of-the-art methods that utilize transformer and CNN architectures. Transformer and CNN architectures have broadly dominated the state-of-the-art for anomaly detection in recent years. Deep learning encoders are trained on a variety of downstream tasks, including detection of synthetic faults [11], the prediction of future data [5], and reconstruction accuracy [4]. The feature encodings are then subject to threshold-based detection. We are investigating a "mix-and-match" method of encoding architectures (transformers, CNNs) and downstream tasks, incorporating each into TRIAD in order to comprehensively assess and implement the most robust and efficient combinations of state-of-the-art techniques.

2. MAIFLOWER BENEFITS

MAIFLOWER is beneficial for the following reasons:

- MAIFLOWER provides a generalized and modular fault management architecture that can be quickly spun up for any number of subsystems.
- Autonomous, high-speed anomaly detection along with "root cause" analysis by correlating time-series data across subsystems, thereby capturing cascading impacts of single faults on a spacecraft as a whole.
- Introduction of transformer-based anomaly detection to fault detection and space domain.
- Astrobotic's VSAT provides a real platform to prove MAIFLOWER's feasibility for adaptation to other spacecraft.
- Some anomalies, such as soil collapse and slippage, during or after ROSA deployment, can only be recovered from if an autonomous detection and diagnosis capability exists onboard. There is just insufficient time for ground controllers to save the VSAT mission under these circumstances, which is why Astrobotic considers MAIFLOWER critical to the success of their mission.
- MAIFLOWER's flexible, general, open architecture enables easy interfacing with other algorithms and other 3rd-party software.
- As described below, the hybrid approach combines the strengths and mitigates the weaknesses of BMR or ML when applied alone.

3. BENEFITS/DISADVANTAGES OF DIFFERENT APPROACHES COMPARED TO HYBRID APPROACH

MBR for Detection and Diagnosis has several benefits and a few disadvantages:

- Benefits:
    - Detect and diagnose anomalies never before encountered
    - Effectively utilize existing design knowledge of subsystem engineers
    - Do not require large amounts of data
    - Can explain their reasoning and are human understandable
    - Have models and code that can be validated
    - Can diagnose down to the lowest modelled component level
    - Can handle rare (but modelled) operating conditions
    - Can reconfigure, recalculate resources, and replan
- Disadvantages:
    - Time-consuming to develop
    - May have too-loose detection thresholds
    - Are limited, to known, designed, components and behaviors

TRIAD and Machine Learning for Detection and Diagnosis have benefits and a few disadvantages:

- Benefits:
    - Little development time required
    - Can learn new, unknown relationships and phenomena
- Disadvantages:
    - Requires a lot of data for training
    - Cannot reliably extrapolate
    - Flight certification may be challenging
    - Susceptible to false alarms
    - Only useful for diagnosis if explicitly trained on the faults to be indicated

4. HYBRIDIZING MBR AND TRIAD

A hybrid fault detection system uses both TRIAD/ML and MBR to monitor sensor values in an effort to identify anomalous behavior. Hybridizing MBR and TRIAD emphasizes the benefits of each while minimizing their disadvantages. Generally, the benefits of MBR coincide with the disadvantages of TRIAD, and vice versa, so the two complement each other well [3]. When an anomaly is detected, a few possibilities present themselves: both systems detect the anomaly, only TRIAD detects it, or it is only detected by the MBR system. TRIAD detects all data streams that are "different" from the data on which it was trained, and as such it generates reports when encountering novel outside contexts, even when all systems are functioning effectively. MBR is a good check on these false alarms, especially if a method exists to determine if TRIAD training data does not include data from the current operating state. However, TRIAD's model-free anomaly detection can also detect faults that escape the analytic limits of MBR. Because MBR and TRIAD are very different technologies (especially in that one is model-based while the other is model-free), when they agree, this consensus instills extra confidence in the result (compared to using MBR or TRIAD alone).

The hybrid system to anomaly detection and diagnosis effectively combines the benefits and mitigates the weaknesses of the individual technologies. Specifically the hybrid system has the following benefits:

- Utilizes a priori design knowledge
- Utilizes real or simulation sensor data
- Very independent technologies
- Can use its MBR model to diagnose, reconfigure, recalculate, and replan in reaction to newly detected faults and can explain its reasoning
- Can discover unknown or unmodelled relationships
- Can detect and diagnose anomalies never before encountered or trained for
- Can use the ML component to disambiguate overlapping MBR states

5. ASTROBOTIC'S VSAT SYSTEM

Aa NASA intends to dramatically increase the number of manmade objects and systems on the Moon, it is essential to develop a grid to generate and distribute electrical power. This grid will consist of several elements, from the generators to the intermediate distributors to the final recipients, and a multitude of faults could occur at any stage in this process. Astrobotic's Vertical Solar Array Technology (VSAT) is an example of a generator platform. At a high level, it resembles a rover with a large vertical solar array that can unfurl once the rover has reached a suitable area for solar power generation and that rotates to track the motion of the Sun. The VSAT is likely to be deployed near the lunar south pole, which is why it is oriented vertically. A graphic of the array is shown lower right, while the figure, left, provides a close-up view of the rover.

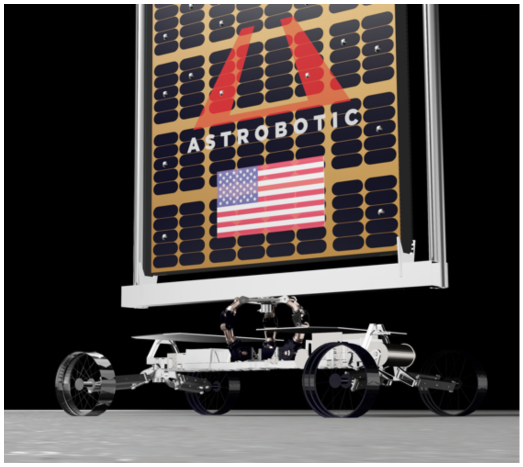

The VSAT platform will face several challenges. First, the solar array is significantly larger than the base of the rover; as the array unfurls, the center of gravity eventually sits far above the rover's base. This means that if the array is not perpendicular (relative to the plum vector on the Moon), it may buckle, and worst of all, the entire system is susceptible to tipping over. The operational window for the deployment of the array is thus very precise and any amount of leaning must be carefully tracked and adjusted for (or the deployment process must be cancelled and tried again). A requirement for the unfurling to even begin is the settling of the wheels of the main rover body into the lunar regolith. The VSAT platform will do this by wiggling its wheels (small motions in opposition to each other and opposite to that of the opposing wheel) to dig into the lunar regolith and make the platform more stable for the solar array deployment.

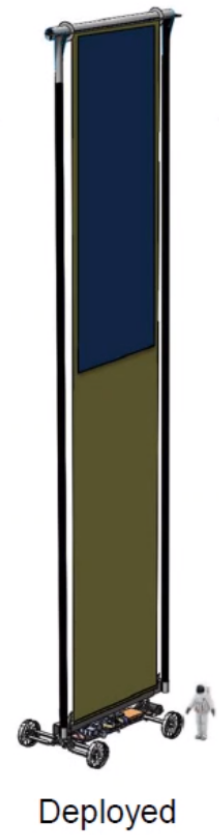

6. FAULT SCENARIOS

The most critical and catastrophic result of a fault is, of course, tipping. In order to be rolled up when stowed, the booms on either side are similar to a plastic straw with a long vertical slit where the "straw" can be rolled up by flattening out the curvature of the straw but which snaps back to cylinder as it is unwound. During stowing operations, the opposite effect occurs, the "straw" is flattened as it is rolled up. Given the height of the center of mass of the deployed ROSA and the width of the wheel base, the maximum allowable tilt can be calculated and it is 3 degrees. Most of the weight (because that is where the solar cells actually are) is in the upper half of the 60' mast. Combined with the fact that the booms act like a spring, it causes the mast to sag further in the direction out of vertical. There is also a potential for an inverted pendulum effect. This could be quite complex because the mast is less stiff in the direction normal to the surface of the mast and more stiff in the direction along the surface of the mast, from one beam toward the other. The inverted pendulum can therefore have an elliptical motion. The inverted pendulum effect, if it occurs, decreases the effective safety tolerance to less than the static value of 3 degrees.

These failure scenarios can be a result of several causes. The sensors used to determine "verticalness" may have too much error. Energy could be added to the system, causing the inverted pendulum, possibly from rough gears or motors or other damage due to launch vibrations. If the soil gives way during or after deployment, it would introduce both an off-angle *and* energy for the inverted pendulum. The soil could give way as a result of vibrations during mast deployment, the very slow Sun-tracking process, during unwinding

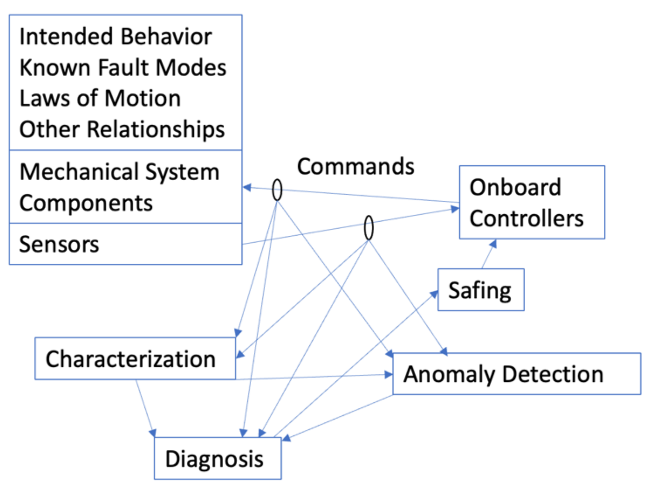

Fig, 2. High-Level Architecture of MAIFLOWER Applied to VSAT every 28 days (because the solar array cannot continue tracking in one direction more than one full revolution), or moonquakes.

7. VSAT MODEL

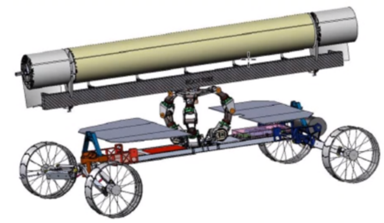

Driving Configuration

Several kinds of sensors are onboard the VSAT to help detect anomalies before they become catastrophic. Unfortunately given the round-trip communications delays, it would not always be possible for humans on the ground to take effective action in time, such that the VSAT needs an onboard, autonomous, sense and react capability. The types of sensors include load sensors at each of the wheels, inclinometers, gimbal joint angle sensors, mast deployment motor sensors, and up-facing and side-facing cameras. The data these sensors collect are processed with both the MBR system and TRIAD, which also includes the ML techniques discussed earlier. A high-level architecture of MAIFLOWER applied to VSAT is shown to the right, followed by the model of the mechanical system components.

Relationships are used to both sanity-check sensor values (i.e., identify faulty sensors) and to diagnose causes of anomalies. An example of a relationship is that using the angle of the body from vertical, the angle of the platform and solar array base, and the state of the solar array (how far deployed, deployment motion, inverted pendulum motion, etc.), the wheel loads can be calculated and cross-checked against the wheel load sensors. Another example is how the vertical motion of the IMU at the top of the mast can be checked against both the up-facing camera and motor sensors which can be used to also calculate the degree of deployment. Cross-vertical motions of the top IMU (i.e., inverted pendulum motions) can be cross-checked against the up-facing cameras. Angular velocities detected by the IMUs can be cross-checked against optical flows in the side facing cameras. Small, very slow changes in attitude, below the noise thresholds of the IMUs can be found by attitude calculations with star-tracking from the up-facing cameras. By detecting and diagnosing problems early, before they become so severe that disaster strikes, MAIFLOWER can greatly decrease mission risk. The primary safing mechanism is to stop the current operation and retract the mast (to improve its center of gravity). For example, if during deployment, inverted pendulum motion is detected, stopping all motion would tend to decrease the problem. If there is a danger of tipping, retracting the mast will improve the problem. If the soil starts collapsing or sliding during deployment, solar tracking, or unwinding, halting the motion will stop the vibration. For significant angular deviations of the rover, the gimbal will need to be quickly releveled.

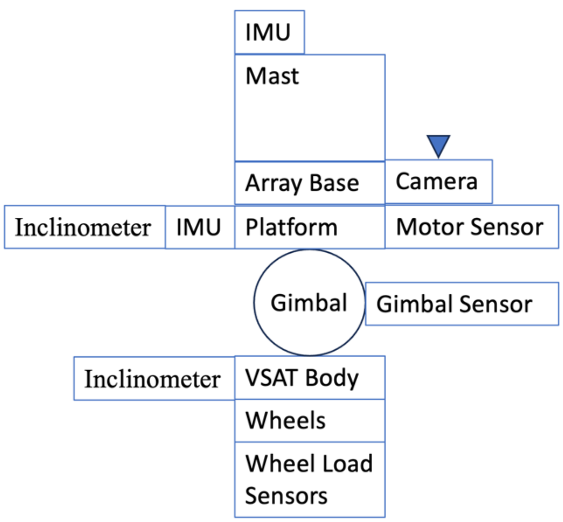

Fig. 3. Model of Mechanical System Components

8. SIMULATION STUDY RESULTS

The initial MAIFLOWER prototype comprises three independent sections: the physics simulation and two modules for anomaly detection: a model-based reasoning module and a data-driven module called TRIAD. We ensured a firewall between the development of the simulation component and that of the anomaly detection components so that the latter would not be unrealistically influenced or informed by the former. Specifically, the MBR module uses a model with lower fidelity than the simulation. Initially we focused on the mechanical aspects of the VSAT mission, in particular the scenario where the VSAT has navigated to the location of choice, settled into the lunar regolith, and is beginning to deploy the onboard Roll Out Solar Array (ROSA). After discussing this scenario with Astrobotic, it became clear that their primary concern was the danger of the entire platform tipping over due to one of several possible fault scenarios:

- Soil collapse underneath the wheels. In this scenario, the rover has dug in but a sudden shift in the regolith causes the entire platform to become unlevel with the surface. This sudden collapse has the potential to tilt the rover past the tipping point.
- Soil slippage. Soil slippage involves translational movement of the VSAT itself on some slope. Due to the heightened center of gravity, this movement also poses a tipping threat.
- Sensor error or mechanical launch damage that prevents the ROSA from leveling initially. This is a scenario where the gimbal that orients the ROSA such that its plane is parallel to the gravitational vector of the Moon is incorrectly positioned to be slightly off from level (due to launch damage or other issue(s)). In this configuration, deploying the ROSA may cause the weight of the VSAT to be unevenly shifted, which can become a tipping risk.

We developed a low-fidelity physical simulation of the VSAT to model its kinematic behavior on the lunar surface. Our simulation fully models:

- The slope of the lunar surface.
- Equations of motion for the rover and attached ROSA. This includes the normal force from the surface, associated friction, and gravity.
- Stabilizing and attachment interactions between the VSAT's Mobility Platform and the ROSA.
- Failures in sensors, soil, and leveling mechanism.
- Realistic noise for all sensors.
- ROSA solar array deployment process.

Our simulation generated telemetry data for a list of sensors, which was selected after conversations with Astrobotic regarding which sensors were planned to be supported. These onboard sensors included inertial measurement units (IMUs), horizontal cameras to extract information on angular velocity, force sensors on the wheels, and upward-facing cameras to determine the orientation of the ROSA as well as the overall attitude of the rover itself. For additional realism, we included sensor noise as part of the output telemetry as noise is a problem any real system would have to properly manage. Once the physics simulation was complete, we added a capacity to induce faults into the system at specific times, such as:

- Soil collapse under specific wheels.
- Soil slippage in a specific direction for a specified amount of time.
- Errors in the ROSA leveling process.
- Sensor faults of various kinds for a specified set of sensors.

The MBR module was designed to ingest telemetry data from the physics simulation and was an implementation of the MBR techniques discussed above. We designed this module as an enhancement of our previous work and specifically for the VSAT's mechanical system. At a high level, MBR used a basic model of the VSAT's mechanical components (explicitly including the underlying soil as a component) to compare incoming sensor data with expected telemetry data to determine if any sensors were reading anomalous values. Of those anomalous sensors, MBR performed crosschecking to determine if the issue was due to a sensor fault (reading random values, frozen at a single value, reading only zero values, etc.). In cases where no sensor fault was detected, MBR then invoked a diagnostic procedure to determine which of the components was the root cause of the anomaly and attempt to characterize the anomaly as one of the above known fault cases. MBR was able to effectively and quickly diagnose the above faults as well as faults in the sensors themselves. In actual operation, the model used by MBR will never reflect with perfect accuracy the reality of the device and its environment. To prove feasibility, therefore, it was important that we made sure that the model used by MBR was at a lower fidelity than the model that the simulation uses, which we did; furthermore, we showed that even with this realistic handicap, the model could always detect and diagnose the anomalies correctly and quickly.

We also designed a data-driven module using a method developed by Stottler Henke called Thermodynamic Reasoning for Intelligent Anomaly Detection (TRIAD). TRIAD takes inspiration from thermodynamics to learn patterns in higher-level features that are emergent from the underlying sensor values. These features include standard statistical measures, such as mean and standard deviation, as well as Fourier transforms of the sensor data to leverage patterns in the frequency domain to detect anomalies that might be opaque to standard analysis, especially in the presence of noisy data. The online system consisted of two parts: the first was trained to detect anomalies while the second attempted to classify those anomalies by matching them to known specific types of faults. TRIAD learned from both nominal and anomalous data offline, but the anomaly *detection* framework is only trained on nominal data and can detect faults on which it was not explicitly trained. TRIAD proved effective in finding anomalies in complex sensor data and matching them to specific faults, both when the faults originated due to a component and when the fault was due to sensor failure. Of specific note was TRIAD's performance in anomaly detection: across a variety of faults, it overwhelmingly detected anomalies within several sensor readings of the fault origination, typically in under a fifth of a second.

We demonstrated several capabilities of MAIFLOWER. These included:

- <u>Detection and characterization of soil collapse.</u> The initial prototype was able to detect and characterize both symmetric and asymmetric soil collapses from the onboard sensors. Both MBR and TRIAD were efficient in finding faults, taking just 3.6 milliseconds and 30 milliseconds of processing time respectively. Even in catastrophic scenarios where the soil collapsed in a way that caused the rover to fall over, MAIFLOWER detected collapse several seconds before the point of no return.
- <u>Detection and characterization of soil slippage.</u> Similar to soil collapse, MAIFLOWER was able to detect this fault in 3.5 milliseconds (MBR) and 30 milliseconds (TRIAD).
- <u>ROSA leveling errors.</u> As the ROSA was to be leveled at the beginning of the simulation start, MBR and TRIAD were able to find this fault immediately.
- A large variety of sensor faults:
    - Reading random values
    - Frozen at a singular value
    - Reading only zeros as output
    - Combinations of the above errors
    - Detecting soil collapse in spite of a sensor error
    - Finding anomalies despite a sudden increase in noise from a relevant sensor

9. DETAILED SCENARIO RESULTS

9.1 Prototype Scenario 1: Symmetric Soil Collapse

In this scenario, soil collapses 1cm under both wheels on one side of the rover at 30 seconds. This induces oscillation in the ROSA but does not result in tipping.

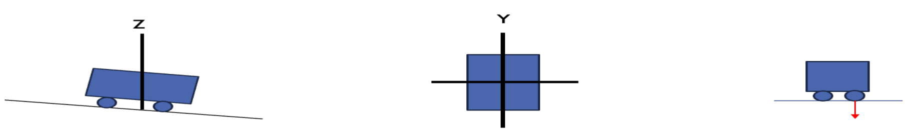

Fig. 4. Sketch of Soil Collapse

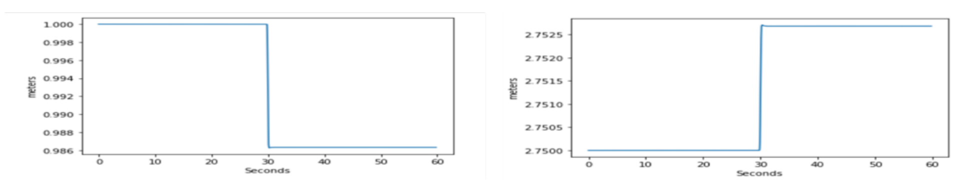

Fig. 5. X and Z Position on Non-Collapsed Corner

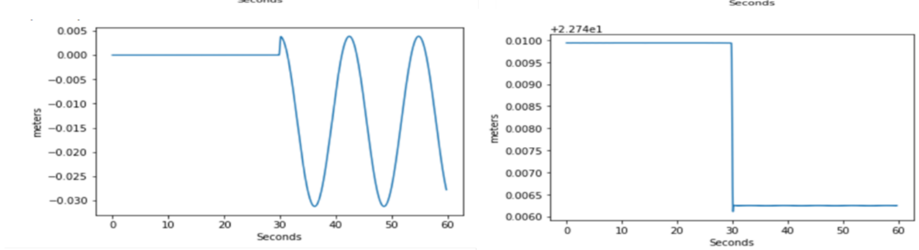

Fig. 6. X and Z Positions of Top Of Panel

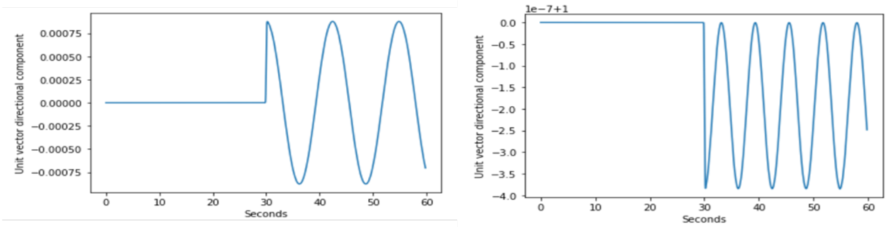

Fig. 7. X and Z Components of Panel Orientation

9.1.1 Model-Based Reasoning Results

Using the methodology described above, MBR first noticed anomalous values in each of the four leg sensors and proceeded to crosscheck these sensors against each other. Once determining that all four leg sensors were reading consistent values despite deviating from the expected values, MBR proceeded to component diagnosis. Here, MBR first checked the soil component for a fault. As mentioned above, we explicitly chose to model the soil as a component because, even though it is physically separate from the VSAT itself, faults such as collapse and slippage can still be attributed to it. During diagnosis, MBR inspected the list of anomalous sensors and found all four leg sensors were reading anomalous values (*the cause of which cannot be attributed to some kind of sensor malfunction)*, which it recognized was indicative of a soil collapse.

To date, MBR and TRIAD are decoupled from any downstream processing, so we routed all output to a file in place of this. Figure 8, left, shows this entire process, from sensor crosschecking (first red box) to the final determination of a fault (second red box).

```
starting diagnosis
forward pass through components
anomaly detected! leg 1 sensor is off by 7.182228618419921%
anomaly detected! leg 2 sensor is off by -18.882300726303477%
anomaly detected! leg 3 sensor is off by 7.182228618419921%
anomaly detected! leg 4 sensor is off by -18.88230072630349%
sensors {'leg1': 1, 'leg2': 1, 'leg3': 1, 'leg4': 1} are reading off nominal values
now running anomaly check for sensor leg1. crosschecking against the following: ['leg1', 'leg2',
'leg3', 'leg4']
sensor anomaly detection complete, results: ('sensor_consistent', None)
now running anomaly check for sensor leg2. crosschecking against the following: ['leg1', 'leg2',
'leg3', 'leg4']
sensor anomaly detection complete, results: ('sensor_consistent', None)
now running anomaly check for sensor leg3. crosschecking against the following: ['leg1', 'leg2',
'leg3', 'leg4']
sensor anomaly detection complete, results: ('sensor_consistent', None)
now running anomaly check for sensor leg4. crosschecking against the following: ['leg1', 'leg2',
'leg3', 'leg4']
sensor anomaly detection complete, results: ('sensor_consistent', None)
one of the components is faulty
checking if soil component is the cause of the anomaly
found 4 leg force sensor anomalies: soil is likely undergoing a collapse
the cause of the anomaly was from soil component
diagnosis complete
```

Fig. 8. From Sensor Crosschecking (1st Red Box) to Final Fault Determination (2nd Red Box)

9.1.2 TRIAD Results

```
TIME 5.12
TIME 5.15
frequency domain anomaly detected
                                                 suspected soil collapse
TIME 5.180000000000001
frequency domain anomaly detected
                                                 suspected soil collapse
TIME 5.21
frequency domain anomaly detected
                                                 suspected soil collapse
TIME 5.24
frequency domain anomaly detected
                                                 suspected soil collapse
TIME 5.270000000000005
boundary value anomaly detected
                                                 soil collapse
TIME 5.300000000000001
boundary value anomaly detected
                                                 soil collapse
TIME 5.33
boundary value anomaly detected
                                                 soil collapse
TIME 5.36
boundary value anomaly detected
                                                 soil collapse
TIME 5.390000000000001
```

As shown in Fig. 9, TRIAD notices an anomaly .15 seconds after the collapse, guesses the nature of the anomaly correctly instantly, and confirms with confidence .27 seconds after collapse

Fig. 9. TRIAD Notices Anomaly .15 Seconds After Collapse, Correctly Guesses Nature of Anomaly Instantly, and Confirms with Confidence .27 Seconds After Collapse

9.2 Prototype Scenario 2: Asymmetric Soil Collapse

In this scenario, soil collapses beneath all wheels barring the front left wheel, triggering oscillations in the ROSA in both the X and Y directions. These oscillations are of varying frequencies, as the spring constants holding the ROSA in place along its stable and unstable axes are different. The back right wheel hits the ground first, and the VSAT briefly sways, held by two wheels, until the back left wheel makes contact with the ground at around 7.5 seconds, stabilizing the system. Note that the normal force induced on each wheel post collapse oscillates—this is due to the swaying of the ROSA.

9.2.1 Simulation Results

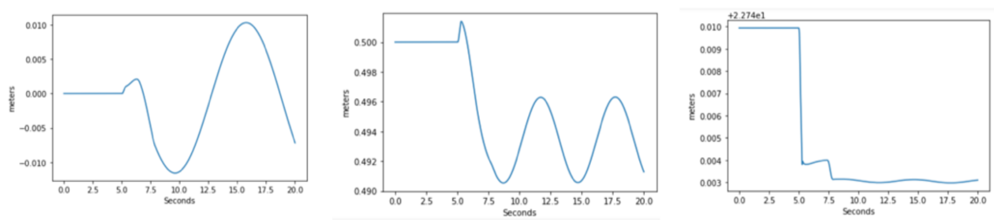

Fig. 10. X, Y, Z Positions of ROSA Tip

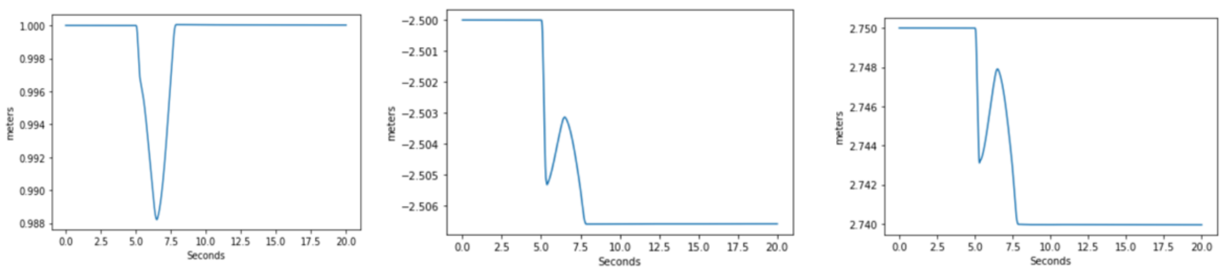

Fig. 11. X, Y, Z Position of the Back Left Corner

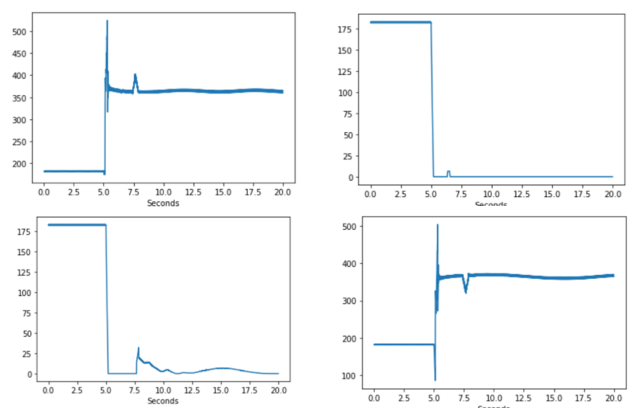

Fig. 12. Normal Force (Newtons) on Each Wheel (Top Left Chart Corresponds to Front Left Wheel, Etc.)

9.2.2 MBR Results

Similar to the symmetric soil collapse scenario, MBR inspects the list of anomalous sensors (once crosschecking is complete) for any patterns it has been programmed to find. Here, it finds that four wheel force sensors and two gimbal accelerations are anomalous and determines that the soil must be undergoing an asymmetric collapse.

Note that the final line of this output indicates how much time MBR took to run: 2.7 milliseconds. This time is measured from when the telemetry data is received to when the diagnosis for the current step has been made.

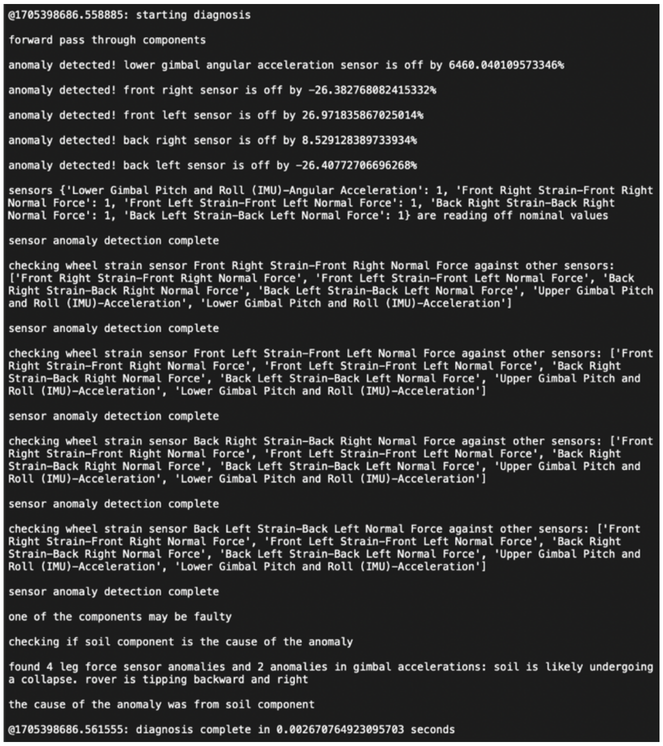

```
@1705398686.558885: starting diagnosis

forward pass through components

anomaly detected! lower gimbal angular acceleration sensor is off by 6460.040109573346%

anomaly detected! front right sensor is off by -26.382768082415332%

anomaly detected! front left sensor is off by 26.971835867025014%

anomaly detected! back right sensor is off by 8.529128389733934%

anomaly detected! back left sensor is off by -26.40772706696268%

sensors {'Lower Gimbal Pitch and Roll (IMU)-Angular Acceleration': 1, 'Front Right Strain-Front Right
Normal Force': 1, 'Front Left Strain-Front Left Normal Force': 1, 'Back Right Strain-Back Right
Normal Force': 1, 'Back Left Strain-Back Left Normal Force': 1} are reading off nominal values

sensor anomaly detection complete

checking wheel strain sensor Front Right Strain-Front Right Normal Force against other sensors:
['Front Right Strain-Front Right Normal Force', 'Front Left Strain-Front Left Normal Force', 'Back
Right Strain-Back Right Normal Force', 'Back Left Strain-Back Left Normal Force', 'Upper Gimbal Pitch
and Roll (IMU)-Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-Acceleration']

sensor anomaly detection complete

checking wheel strain sensor Front Left Strain-Front Left Normal Force against other sensors: ['Front
Right Strain-Front Right Normal Force', 'Front Left Strain-Front Left Normal Force', 'Back Right
Strain-Back Right Normal Force', 'Back Left Strain-Back Left Normal Force', 'Upper Gimbal Pitch and
Roll (IMU)-Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-Acceleration']

sensor anomaly detection complete

checking wheel strain sensor Back Right Strain-Back Right Normal Force against other sensors: ['Front
Right Strain-Front Right Normal Force', 'Front Left Strain-Front Left Normal Force', 'Back Right
Strain-Back Right Normal Force', 'Back Left Strain-Back Left Normal Force', 'Upper Gimbal Pitch and
Roll (IMU)-Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-Acceleration']

sensor anomaly detection complete

checking wheel strain sensor Back Left Strain-Back Left Normal Force against other sensors: ['Front
Right Strain-Front Right Normal Force', 'Front Left Strain-Front Left Normal Force', 'Back Right
Strain-Back Right Normal Force', 'Back Left Strain-Back Left Normal Force', 'Upper Gimbal Pitch and
Roll (IMU)-Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-Acceleration']

sensor anomaly detection complete

one of the components may be faulty

checking if soil component is the cause of the anomaly

found 4 leg force sensor anomalies and 2 anomalies in gimbal accelerations: soil is likely undergoing
a collapse. rover is tipping backward and right

the cause of the anomaly was from soil component

@1705398686.561555: diagnosis complete in 0.002670764923095703 seconds
```

Fig. 13. Within 2.7 Milliseconds, MBR Finds Four Wheel Force Sensors and Two Gimbal Accelerations are Anomalous and Determines Soil Must be Undergoing Asymmetric Collapse

9.2.3 *TRIAD Results*

TRIAD notices an anomaly .15 seconds after the collapse, guesses the nature of the anomaly correctly instantly, and confirms with confidence .3 seconds after collapse.

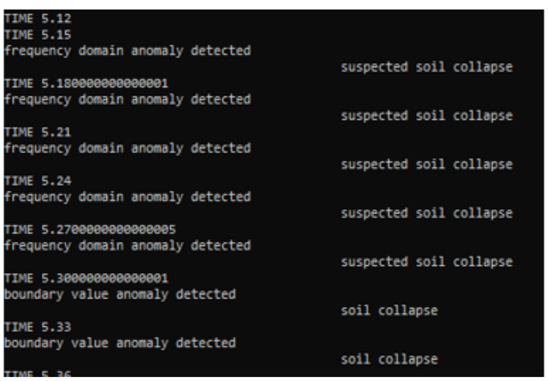

```
TIME 5.12
TIME 5.15
frequency domain anomaly detected
                                          suspected soil collapse
TIME 5.180000000000001
frequency domain anomaly detected
                                          suspected soil collapse
TIME 5.21
frequency domain anomaly detected
                                          suspected soil collapse
TIME 5.24
frequency domain anomaly detected
                                          suspected soil collapse
TIME 5.2700000000000005
frequency domain anomaly detected
                                          suspected soil collapse
TIME 5.300000000000001
boundary value anomaly detected
                                          soil collapse
TIME 5.33
boundary value anomaly detected
                                          soil collapse
```

Fig. 14. TRIAD Notices Anomaly .15 Seconds After Collapse, Correctly Guesses Nature of Anomaly Instantly, and Confirms with Confidence .3 Seconds After Collapse

9.3 Asymmetric Soil Slipping

In this scenario, the rover slips backwards and to the left, at a 6-degree incline. In this scenario, friction is returned to the system faster (creating more dramatic oscillations in the ROSA), but is never fully recovered. This lower friction creates subtle slipping and more erratic normal force distributions even after the VSAT is no longer sliding, as the torque normally exerted by friction helps to counteract the oscillation of the ROSA. As a result of this, the VSAT itself briefly tilts at around 18 seconds—this corresponds two nearly simultaneous peaks along the two axes of the ROSA's oscillation, which is enough to lift two of the legs up. Upon recovering, the VSAT briefly tips in the opposite direction as well. Note that the normal forces of each leg briefly reach 0 after these events—however, due to increased friction by this point, the VSAT no longer tips. Also note that the varying distributions of normal force on each leg as the rover slips lead to varying levels of friction. This causes the rover to rotate as it descends down the slope. This can be seen from the below charts in Figure 15, 16, and 17 in several ways. First, note that the changes in X and Y positions of the corner vary by about a factor of 2; as the slope has an even incline between the

X and the Y directions, in the absence of rotation, the X and Y positions would change by roughly the same amount. Secondly, note how the X and Y positions of the panel both contain both the high and low frequency oscillations corresponding to both steadying springs; this is because the steadying springs are no longer aligned with the global X and Y directions.

9.3.1 Simulation Results

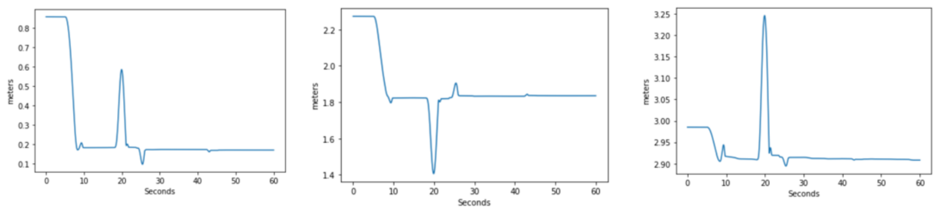

Fig. 15. X, Y, Z Positions of Back Right Corner of VSAT

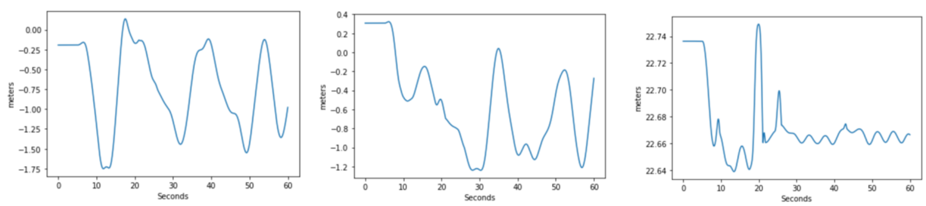

Fig. 16. X, Y, Z Positions of the Panel Tip

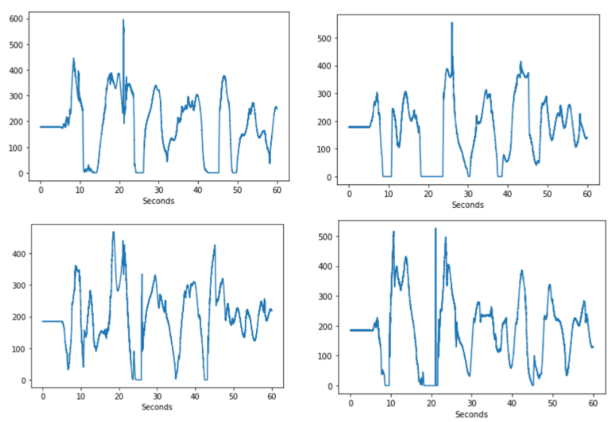

Fig. 17. Normal Force (Newtons) on Each of the Four Wheels (Upward Front Corresponds to Front Left, etc.)

9.3.2 MBR Results

In this scenario, MBR finds a long list of anomalous sensors. Of these sensors, it notices that four are sideways cameras, two are in upward facing cameras, and four are in the gimbal linear and angular accelerations. Because this specific set of sensors are showing anomalies, MBR diagnoses this problem as soil slippage in roughly 3.5 milliseconds.

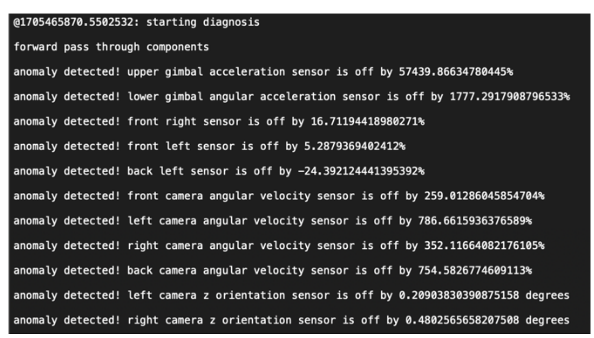

```
@1705465870.5502532: starting diagnosis
forward pass through components
anomaly detected! upper gimbal acceleration sensor is off by 57439.86634780445%
anomaly detected! lower gimbal angular acceleration sensor is off by 1777.2917908796533%
anomaly detected! front right sensor is off by 16.71194418980271%
anomaly detected! front left sensor is off by 5.2879369402412%
anomaly detected! back left sensor is off by -24.392124441395392%
anomaly detected! front camera angular velocity sensor is off by 259.01286045854704%
anomaly detected! left camera angular velocity sensor is off by 786.6615936376589%
anomaly detected! right camera angular velocity sensor is off by 352.11664082176105%
anomaly detected! back camera angular velocity sensor is off by 754.5826774609113%
anomaly detected! left camera z orientation sensor is off by 0.20903830390875158 degrees
anomaly detected! right camera z orientation sensor is off by 0.4802565658207508 degrees
```

Figure 18. MBR Finds a Long List of Anomalous Sensors

```
one of the components may be faulty

checking if soil component is the cause of the anomaly

found 4 anomalies in sideways cameras, 2 anomalies in upwards cameras, and 4 anomalies in gimbal
angular and linear accelerations: rover is likely slipping on the soil. rover is tilted backward and
left

the cause of the anomaly was from soil component

@1705465870.553728: diagnosis complete in 0.0034749507904052734 seconds
```

Fig. 19. MBR Diagnoses Problem as Soil Slippage in Roughly 3.5 Milliseconds

9.3.3 TRIAD Results

TRIAD detects the anomaly at .15 seconds, accurately guesses its nature, and classifies it successfully at .36 seconds. Note that this is TRIAD's feed in "Verbose Mode," which lists out the anomalous sensor values associated with its diagnostics, which increase over time as the downstream effects of the slippage become more noticeable. The volume and relevance of anomalous sensors can serve as an additional metric of TRIAD's confidence in its diagnoses.

```
TIME 5.15
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll
 (IMU)-Acceleration_y']
TIME 5.180000000000001
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll
 (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z']
TIME 5.21
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll
 (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z']
TIME 5.24
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll
 (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z', 'Panel Pitch and Roll (IMU)-Angular Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_x', 'Upper Gimbal
 Pitch and Roll (IMU)-Velocity_y']
TIME 5.2700000000000005
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll
 (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_x', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_y', 'Panel Pitch and R
oll (IMU)-Acceleration_z', 'Panel Pitch and Roll (IMU)-Angular Acceleration_x']
TIME 5.300000000000001
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Panel Pitch and Roll (IMU)-Acceleration_z', 'Panel Pitch and Roll (IMU)-Angular
 Acceleration_x', 'Panel Pitch and Roll (IMU)-Angular Acceleration_y', 'Panel Pitch and Roll (IMU)-Velocity_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Ro
ll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_x', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_y', 'Panel Pitch and
 Roll (IMU)-Velocity_x']
TIME 5.33
frequency domain anomaly detected
                                             suspected slipping
relevant sensors:  ['Panel Pitch and Roll (IMU)-Acceleration_x', 'Panel Pitch and Roll (IMU)-Acceleration_y', 'Panel Pitch and Roll (IMU)-Angular Acceleration_y', 'Panel Pitch and Roll (IMU)
-Velocity_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_y', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_z', 'Upper Gimbal Pitch
and Roll (IMU)-Velocity_x', 'Upper Gimbal Pitch and Roll (IMU)-Velocity_y', 'Panel Pitch and Roll (IMU)-Angular Acceleration_x', 'Panel Pitch and Roll (IMU)-Velocity_x']
TIME 5.36
boundary value anomaly detected
                                             slipping
relevant sensors:  [{'name': 'Panel Pitch and Roll (IMU)-Angular Acceleration_x', 'fault_type': 'anomaly', 'elaboration': 'high', 'confidence': 0.8333333333333334}, {'name': 'Upper Gimbal Pi
tch and Roll (IMU)-Acceleration_x', 'fault_type': 'anomaly', 'elaboration': 'low', 'confidence': 0.8333333333333334}, {'name': 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_y', 'fault_type
': 'anomaly', 'elaboration': 'low', 'confidence': 0.8333333333333334}]
TIME 5.390000000000001
boundary value anomaly detected
                                             slipping
relevant sensors:  [{'name': 'Panel Pitch and Roll (IMU)-Acceleration_y', 'fault_type': 'anomaly', 'elaboration': 'low', 'confidence': 0.8333333333333334}, {'name': 'Panel Pitch and Roll (IM
U)-Angular Acceleration_x', 'fault_type': 'anomaly', 'elaboration': 'high', 'confidence': 0.8333333333333334}, {'name': 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_x', 'fault_type': 'ano
maly', 'elaboration': 'low', 'confidence': 0.8333333333333334}, {'name': 'Upper Gimbal Pitch and Roll (IMU)-Acceleration_y', 'fault_type': 'anomaly', 'elaboration': 'low', 'confidence': 0.83
33333333333334}, {'name': 'Upper Gimbal Pitch and Roll (IMU)-Velocity_y', 'fault_type': 'anomaly', 'elaboration': 'low', 'confidence': 0.8333333333333334}]
TIME 5.420000000000001
boundary value anomaly detected
```

Fig. 20. TRIAD's Feed in "Verbose Mode," Listing Out the Anomalous Sensor Values Associated with Diagnostics; in this Scenario, TRIAD Detects Anomaly at .15 Seconds, Accurately Guesses its Nature, and Successfully Classifies Anomaly at .36 Seconds

9.4 Random Sensor Error

In this scenario, the portion of the upper gimbal IMU that measures position freezes. This type of random sensor scenario is maximally adversarial, as the random behavior is bounded by past nominal values, and as such it can be difficult to distinguish from simple noisy readings.

9.4.1 MBR Results

MBR processes and catches sensor errors as part of its crosschecking step. Here, MBR recognizes that a sensor is being inconsistent against the list of sensors that should be correlated. MBR is not particularly sure what has happened, so a general message is sent stating that the sensor is broken in an unknown way and that more data is needed.

```
checking position sensor Upper Gimbal Pitch and Roll (IMU)-Position against other sensors: ['Panel
Pitch and Roll (IMU)-Position', 'Panel Pitch and Roll (IMU)-Velocity', 'Panel Pitch and Roll (IMU)-
Acceleration', 'Panel Pitch and Roll (IMU)-Angular Velocity', 'Panel Pitch and Roll (IMU)-Angular
Acceleration', 'Upper Gimbal Pitch and Roll (IMU)-Position', 'Upper Gimbal Pitch and Roll (IMU)-
Velocity', 'Upper Gimbal Pitch and Roll (IMU)-Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-
Position', 'Lower Gimbal Pitch and Roll (IMU)-Velocity', 'Lower Gimbal Pitch and Roll (IMU)-
Acceleration', 'Lower Gimbal Pitch and Roll (IMU)-Angular Velocity', 'Lower Gimbal Pitch and Roll
(IMU)-Angular Acceleration']

sensor anomaly detection complete

sensor is inconsistent with other sensors in an unknown way.
there may be a fault, this sensor may be broken, or this anomaly may be transitory
more data is needed

one of the components may be faulty

checking if soil component is the cause of the anomaly

checking if ROSA component is the cause of the anomaly

@1705459222.296879: diagnosis complete in 0.003480195999145508 seconds
```

Fig. 21. A General Message is Sent Stating that Sensor is Broken in Unknown Way and More Data Needed

9.4.2 TRIAD Results

TRIAD detects an anomaly in the frequency domain at .15 seconds after the error and pattern-matches the anomalous readings to soil slippage. At .27 seconds after the error, TRIAD's high confidence prediction diverges from its original guess and correctly diagnoses that the error is sourced from a random sensor reading. The information TRIAD leverages is twofold—first, the anomalous sensor's behavior loses consistency with slippage in any particular direction, and second, other sensor anomalies that would indicate slipping are not triggered in time.

```
TIME 5.09
TIME 5.12
frequency domain anomaly detected
                                            suspected slipping
TIME 5.15
frequency domain anomaly detected
                                            suspected slipping
TIME 5.180000000000001
frequency domain anomaly detected
                                            suspected slipping
TIME 5.21
frequency domain anomaly detected
                                            suspected slipping
TIME 5.24
frequency domain anomaly detected
                                            suspected slipping
TIME 5.2700000000000005
BROKEN SENSOR:  Upper Gimbal Pitch and Roll (IMU)-Position_x . BEHAVIOR:  random
BROKEN SENSOR:  Upper Gimbal Pitch and Roll (IMU)-Position_z . BEHAVIOR:  random
TIME 5.300000000000001
BROKEN SENSOR:  Upper Gimbal Pitch and Roll (IMU)-Position_x . BEHAVIOR:  random
BROKEN SENSOR:  Upper Gimbal Pitch and Roll (IMU)-Position_y . BEHAVIOR:  random
BROKEN SENSOR:  Upper Gimbal Pitch and Roll (IMU)-Position_z . BEHAVIOR:  random
```

Fig. 22. TRIAD Detects Anomaly in the Frequency Domain at .15 Seconds After Error, Pattern-Matches Anomalous Readings to Soil Slippage; at .27 Seconds After Error, TRIAD's High Confidence Prediction Diverges from Original Guess, Correctly Diagnosing that Error is Sourced from Random Sensor Reading

9.5 Frozen Sensor Error

In this scenario, the strain sensor on the back left leg freezes.

9.5.1 MBR Results

While MBR can catch sensor errors, this specific error is actually undetectable due to the scenarios we have selected. The Phase I MBR system only performs crosschecks on sensors it has found to be anomalous; in other words, if a sensor continues to read nominal values, MBR will not be able to tell if it has been tampered with in any way. For the scenarios that were selected, the VSAT platform was expected to be static and there were no commands to send it into motion of any kind. This means that a frozen sensor would always be reading nominal values as it is within the range of what MBR is expecting to see. A future iteration may improve this and instead monitor the standard deviation of the last several values to compare against the overall expected deviation. This, however, is already performed by TRIAD, demonstrating how MBR and TRIAD complement each other and can cover for each other's weaknesses.

9.5.2 TRIAD Results

```
TIME 5.180000000000001
TIME 5.21
TIME 5.24
TIME 5.2700000000000005
TIME 5.300000000000001
TIME 5.33
BROKEN SENSOR:  Back Left Strain-Back Left Normal Force . BEHAVIOR:  frozen
TIME 5.36
BROKEN SENSOR:  Back Left Strain-Back Left Normal Force . BEHAVIOR:  frozen
TIME 5.390000000000001
BROKEN SENSOR:  Back Left Strain-Back Left Normal Force . BEHAVIOR:  frozen
```

TRIAD detects and correctly classifies the anomaly .33 seconds after it is triggered.

Fig. 23. Anomaly Detection and Classification Within .33 Seconds

10. DATA FROM THE ACTUAL PHYSICAL PROTOTYPE DURING SLOPE LAB TESTING

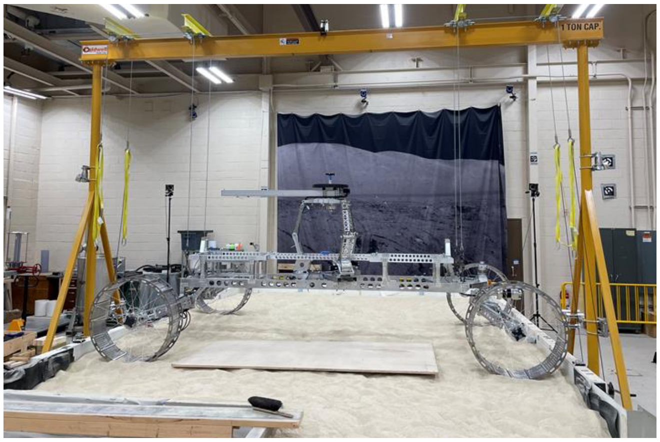

Under their contract with NASA, Astrobotic built a ground test article of their VSAT design and tested in NASA's slope lab, with simulated regolith to verify it met the stated requirements. This picture shows the roughly 500 pound prototype in the lab with a 0 degree slope. The arm and weight in the center of the picture represent the weight and horizontal location of a ROSA leaning at the edge of its 3 degree limit. That arm is mover to different positions and force and position readings are recorded. The next four graphs show the recorded weight on the two rear wheels, the simulated ROSA mass's Z (forward/backward) location, ROSA mass's X (port/starboard) location and, finally, the weight on the two forward wheels.

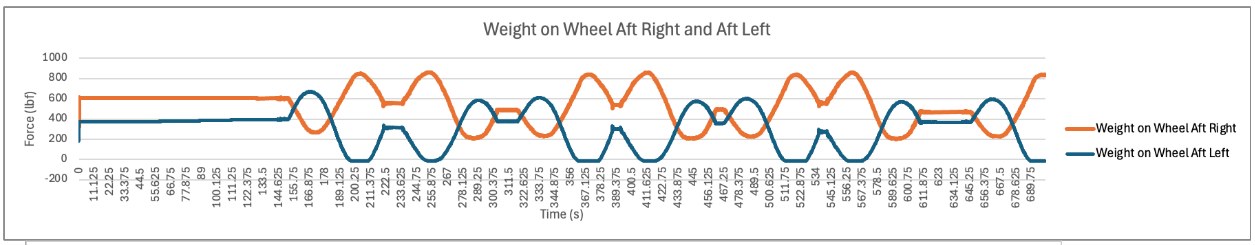

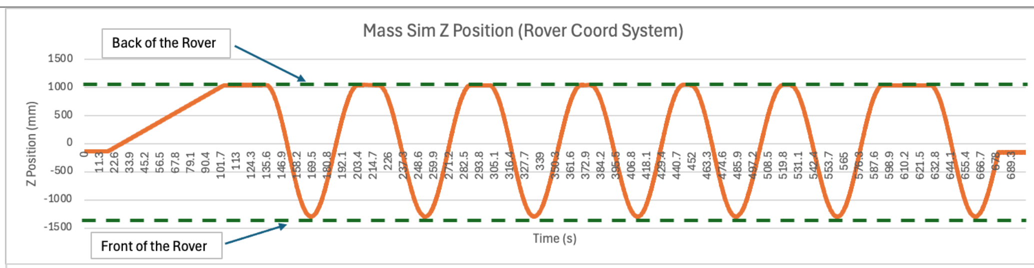

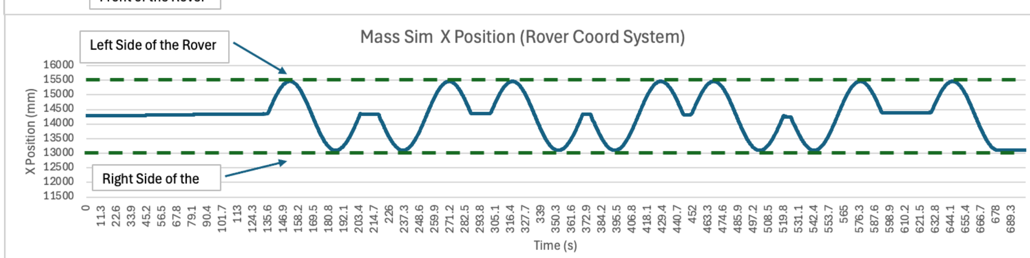

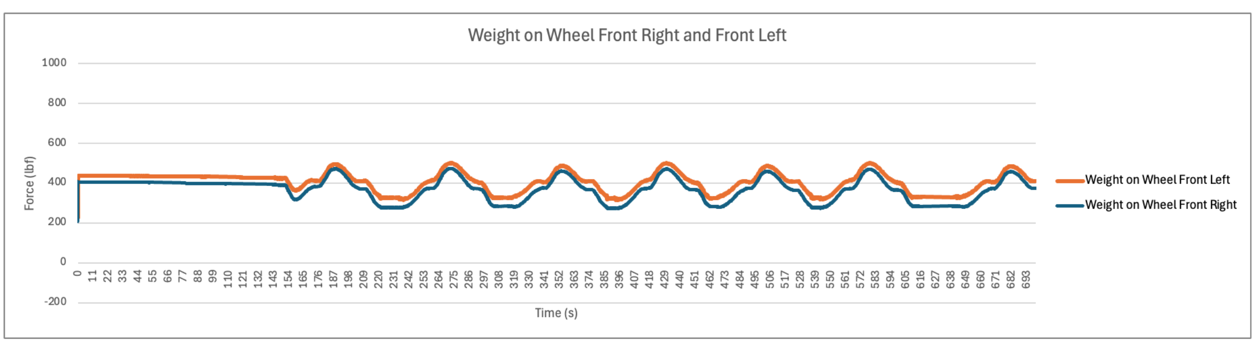

The front two wheels are articulated in the roll direction (to cover rough terrain) so their forces, in a fairly static scenario, are synchronized.

11. FUTURE WORK

The ultimate plan is to field MAIFLOWER on the VSAT that lands on the moon. The next steps toward accomplishing that goal are to receive the full set of slope lab data, build out a high-fidelity simulation for testing and to generate training data, develop improved MBR models and software, test the system against the slope lab data, perform additional experiments where the actual physical rover undergoes soil collapse, soil slippage, ROSA leveling errors, and various sensor errors to confirm MAIFLOWER always behaves appropriately, and integrate MAIFLOWER with VSAT's onboard software.

12. CONCLUSIONS

We have developed a prototype of the MAIFLOWER architecture and applied it to simulated data where the simulation was at a higher fidelity than the MBR models and showed, that even with this limitation, the MBR system could correctly responds (detects and diagnoses so the VSAT control system could compensate) to the most important scenarios in literally a few milliseconds. This was Java code running on a modest desktop machine. Similarly, TRIAD was shown to correctly respond within 10s of milliseconds. The project is therefore proceeding toward its ultimate use on the lunar surface.

13. References

[1] Gao, H. et al. (2022). Tsmae: A novel anomaly detection approach for Internet of Things time series data using memory-augmented autoencoder. *IEEE Transactions on Network Science and Engineering*.
[2] Geneva, N., & Zabaras, N. (2022). Transformers for modeling physical systems. *Neural Networks*, *146*, 272-289.
[3] Li, Z. et al. (2018). Data-driven bearing fault identification using improved hidden Markov model and self-organizing map. *Computers & Industrial Engineering*, *116*(2018), 37-46.
[4] Ramachandran, S., Rosengarten, M., & Belardi, C. (2020). Semi-supervised machine learning for spacecraft anomaly detection & diagnosis. In *2020 IEEE Aerospace Conference*. IEEE.
[5] Ren, H. et al. (2019). Time-series anomaly detection service at Microsoft. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
[6] Stottler, R., Singhal, A., Healy, C., Ramachandran, S., Quinn, K., Palmieri, J., & Logan, S. (2023) Autonomous, Hybrid Space System Fault and Anomaly Detection, Diagnosis, Root Cause Determination, and Recovery. In *Proceedings of Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), 2023*.
[7] Stottler, R., Ramachandran, S., Healy, C., Singhal, A., & Finnigan, E. (2022). 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*.
[8] Stottler, R. (2020). *On-board, autonomous, hybrid spacecraft subsystem fault and anomaly detection, diagnosis, and recovery*. Space Domain Awareness Data Science Working Group, November 2020.
[9] Stottler, R., Ramachandran, S., Belardi, C., & Mandayam, R. (2020). On-board, autonomous, hybrid spacecraft subsystem fault and anomaly detection, diagnosis, and recovery. In *Proceedings of Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), 2020*.
[10] Xu, J. et al. (2021). *Anomaly transformer: Time series anomaly detection with association discrepancy.* arXiv preprint arXiv:2110.02642.