

AIRBUS	ADAP Project	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 1 of 12
---------------	---------------------	---

Project: ADAP Project

Executive Summary Report

	Name	Function	Date	Signature
Prepared by	ADAP Team	-		
Checked by	Alisa Krstova	Technical Manager		
Approved by	Filippo Ales	Study Manager		

Approval and release information:

This document has been approved either by manual signatures as shown above, or via electronic workflow as shown on the cover page preceding this page.

Security information:

Security information applies either by manual signatures in the box below or by electronic workflow as shown in the box on the cover page preceding this page.

Export Controlled:	n.a.
Document	n.a.
Rater: Filippo Ales	Rating Date: 23/06/2023

In collaboration with:



TABLE OF CONTENTS

1	PROJECT OBJECTIVES.....	5
2	PROJECT ACHIEVEMENTS.....	7
2.1	Use-case 1 - AOCS Anomaly on Sensor	7
2.1.1	Pitch and Roll Sensor equipment management & FDIR	7
2.1.2	ADAP system results for UC1.....	8
2.2	Use-Case 2 – Payload Antenna Anomaly.....	9
2.2.1	Payload antenna thermal control management & FDIR	9
2.2.2	ADAP system results for UC2.....	10
2.3	Use-Case 3 - Solar Array Degradation	11
2.3.1	Electrical power system equipment management & FDIR	11
2.3.2	ADAP system results for UC3.....	11

	ADAP Project Executive Summary Report	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 3 of 12
---	---	---

LIST OF FIGURES

Figure 1 Close-up of the anomaly8

LIST OF TABLES

No table of figures entries found.

Change Log

Issue	Date	Description of Change
1	23/06/2023	first issue

1 PROJECT OBJECTIVES

The aim of a Spacecraft's Failure Detection Isolation and Recovery (FDIR) is to guarantee the target of missions' reliability, availability, maintainability and operational autonomy (failure management) in order to ensure its success despite the potential occurrence of failures. The implementation principle of the Spacecraft's FDIR relies on an incremental safeguarding approach where, in the first instance, use is made of the redundant resources when the primary resource is suspected to have caused the failure. However, in case the redundancy switching has not led to resolving the problem, the satellite SAFE mode is entered for preventing a mission loss. The FDIR implementation strategy in terms of failure tolerance technique is also linked to mission phase. During nominal operations, the satellite FDIR strategy is mainly driven by the availability requirement, i.e. is designed to maintain the mission whenever it is possible or to minimize the mission outage duration when the mission cannot be maintained. This means that for the majority of failures it must be possible to autonomously recover nominal operational conditions within a short time. The FDIR strategy used in non-operational phases of the mission (e.g. launch, early orbit phase - which are limited to "lower" satellite modes not used in nominal operations), is mainly driven by the safety of the satellite. It aims at placing the satellite in a safe state in which it can remain without time constraints allowing the Ground to perform the necessary investigations and recovery operations (Fail-Safe approach).

An on-board FDIR design (for example, based on the Packet Utilization Standard) requires that all feared events are identified, made observable and that proper monitorings are associated to the observable. Hence, this classical approach implies that all the possible failures are identified at a very early stage of a project and that all the thresholds are set correctly. For simple architectures, it is possible to easily implement a failure detection system since the failure modes and observables at play are few. But for more advanced systems like a spacecraft, the FDIR implementation requires high development costs and very often also high operational costs, associated to the resolution of in-orbit anomalies which were not foreseen during the spacecraft's design phase.

It can turn out to be more (cost-)effective to have an on-board system which is learning from the data telemetry so that it can carry out the on-board monitoring activities without prior knowledge of expected failures. In this case, the use of Machine Learning algorithms can significantly enhance the performance of the on-board FDIR, notably in identifying and isolating failures (especially unforeseen ones) at the lowest level possible (equipment and subsystem level) thus fostering mission availability and autonomy. In particular, following an unsupervised approach to anomaly detection, the machine learning model is only fed with "nominal" data and should be able to predict when something is not "as expected". In other words there is no need to inject failure examples or failure data but these will be autonomously detected by the system. The advantages of such an approach are therefore reflected in failure detectability and system reactivity.

Although a plethora of failure/anomaly detection strategies have been developed for time series analysis (case on of on-board telemetry monitoring) and used across missions, due to the ever-growing complexity of the spacecraft being built each year, the field of smart anomaly detection both on ground and on board remains challenging. An important limitation relates to hardware and computational resources (a LEON IV processor or space-qualified FPGAs have much less computational power compared to modern GPUs) and therefore adaptation of these techniques is necessary.

For this study, a class of neural networks especially suitable for solving time series problems, such as recurrent neural networks (RNN), convolutional neural networks (CNNs) and their variations, are mostly considered for on-board implementation. These networks can capture the temporal context in a dataset by efficiently processing and prioritizing historical information to perform inference on current data. Specifically, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) improve the memory of the network of long-term dependencies in the data. Moreover, LSTMs and GRUs can be used as components (layers) in complex autoencoder architectures as well as part of unsupervised methods.

The algorithms developed are deployed on hardware using a coprocessor architecture design as preferred solutions since it would lead to the earliest adoption (for example, as payload of opportunity). The data processing unit in which the AI-based FDIR algorithms are deployed is thus a self-standing coprocessor able to handle and decode autonomously telemetries and telecommands coming from OBC and potentially also directly from equipment. Based on a trade-off performed on hardware equipment available on the market, the Zynq US+ based CHICS computer was selected in a self-standing coprocessor configuration as it gives more flexibility in terms of inference options.

	<p style="text-align: center;">ADAP Project Executive Summary Report</p>	<p>ADAP-ADSF-RP-23000051149</p> <p>Issue: 1</p> <p>Date: 23/06/2023</p> <p>Page: 6 of 12</p>
---	---	--

In order to test the AI-based FDIR implementation, the in-orbit telemetry of a satellite constellation (labelled as SAT1 to SAT4 in this study) has been used. For this study, three use-cases have been prioritized, each one with a different and very specific anomaly signature. In particular, an AOCS anomaly (labelled as use-case 1 in the study), a failure in the thermal sub-system (labelled as use-case 2 in the study), and an anomaly on the power sub-system (labelled as use-case 3 in the study) have been explored. All use cases have in common that the detection and isolation with a classical FDIR has been problematic in orbit, either due to complex relations between satellite units, or due to the fact that the standard approach of assigning fixed high and low limits to single parameters with sometimes very large ranges (in order to cope with very different orbital constraints) is not sufficient to safely detect an anomaly in a timely manner.

2 PROJECT ACHIEVEMENTS

The Anomaly Detection – Anomaly Prognosis (ADAP) is an AI-based FDIR system which implements ML with the aim of increasing satellite failure detection capabilities. It has been conceived following two main design drives:

- Compatibility with PUS-based FDIR and/or classical hard-coded FDIR designs
- Capability to perform anomaly detection and isolation as well as prediction of potential failure cases independently from one another

The use-cases selected to demonstrate the ADAP capabilities have been selected based on the data considerations which are listed hereafter:

- Data available from the selected use-cases shall allow to mimic the training of a ML-based FDIR solution for future missions. Hence, the selected use-cases shall allow:
 - to retrieve, produce and mix synthetic and real TM in order to train the AI and check the expected behaviour “in-orbit”
 - to test that the AI solution can be generalized such that the deployment of the ADAP solution to the same sub-systems (for example, thermal, power etc...) of different satellites shall have the same behaviour

In order to fulfill these objectives, Airbus has taken advantage of the telemetry of a constellation of 4 satellites that, from platform side, are the same and are at couples in the same orbital plane. For such constellation, Engineering Test Benches as well as constellation simulators are also available. The use-cases analyzed are:

1. to train the models with the TM of one satellite and afterwards use the ADAP system to detect and predict anomalies on the same satellite
2. to train the model with the TM of one satellite and use the ADAP system to detect anomalies on a twin satellite (generalization of the solution)

For the scope of the project, the use-cases described hereafter are considered.

2.1 USE-CASE 1 - AOCS ANOMALY ON SENSOR

use-case 1 is an AOCS Pitch and Roll Sensor misbehavior, which is leading first to a subsystem and afterwards to a system critical anomaly with consequent transition to satellite SAFE Mode.

This anomaly was observed on all spacecrafts of the constellation as the sensors were impacted by an external perturbation that was unknown until it appeared. The perturbation phenomena was created by very low temperatures on the upper atmospheric layers during winter which don't allow the sensor to clearly detect the transition between Earth (hot side) and Deep Space (cold side). Indeed, the sensor was never before used in this configuration and as it turned out the design was not fully suitable for the new operating conditions. The perturbation phenomena was leading to an increased noise in the pitch and roll measurements and ultimately to an erroneous estimation of pitch and roll pointing.

This noise in the pitch and roll measurement was caused by a seasonal perturbation with a re-occurrence under the same orbital conditions and with the same latitudes of the earth in the sensor field of view. This behavior was observed on all the satellites during the same period of the year.

2.1.1 Pitch and Roll Sensor equipment management & FDIR

The AOCS sensors are subject to the following FDIR level 1 monitorings:

- Sensor power status (failures of Power Distributor LCL and sensor DC/DC converter)
- Sensor temperature (Sensor dissipative DC/DC failure cases)
- Sensor current (Sensor dissipative DC/DC failure cases, low consumption failure cases, failure of Power Distributor LCL)
- Sensor house-keeping bits
- Correct Pointing signal (detector failure cases)

The triggering of these monitorings (which are activated when the corresponding unit is on) will set to failed the corresponding sensor in the failure status of the equipment handler and initiate a sensor switch-over if the unit is used for attitude control or a switch-off when the unit is not used for attitude control.

The reproduction, distribution and utilization of this document as well as the communication of its contents to others without express authorization is prohibited. Offenders will be held liable for the payment of damages.

All rights reserved in the event of the grant of a patent, utility model or design.

In addition, in order to cover failures of the remote interface unit to the sensor interface function, specific monitorings of the house-keeping bits are introduced which are activated after a first sensor failure (detected via the house-keeping bit monitoring). In such a case a remote interface switch-over is performed. The triggering of the sensor house-keeping monitoring could indeed be due to a remote interface unit's failure rather than to a sensor failure. In such a case these additional monitorings (which have a lower filter value for more reactivity) will trigger and perform a remote interface switch-over.

The monitoring of the pointing signal for the AOCS-modes allows sensor detector failure cases to be recovered. This monitoring is indeed classified FDIR level 1 since it is activated only in the modes where correct pointing is expected. The FDIR monitoring on roll/pitch attitude deviations are implemented separately to protect the AOCS function against real attitude deviation.

The AOCS higher-level FDIR level monitorings include, among others, the monitoring of the attitude deviation. This monitoring is triggered when Roll/Pitch/Yaw deviations and angular momentum deviations are above programmable thresholds for a programmable time interval. In case this monitor triggers, the satellite transitions to SAFE Mode.

Based on the implementation of the PUS-based FDIR, there is no possibility to detect this failure at a lower level by just taking into account the observables of the feared events for the sensor. Indeed, the anomaly is not triggered by a failure of the sensor itself, but by an external disturbance, leading to a wrong estimation of the Pitch and Yaw calculation for the specific orbit position. Therefore, despite this being a systematic anomaly, there is no possibility to avoid SAFE Mode triggering by means of classical FDIR implementations without a specific (and complex) software patch specifically designed to tackle this issue.

2.1.2 ADAP system results for UC1

Anomaly Detection

A ML-based FDIR solution has been capable to identify that the readouts of the sensor are not correct by correlating these with the current orbital position and thus pin-point that the failure is actually coming from the sensor and not from a more general attitude control problem. Indeed from qualitative point of view the model performs very well despite, from the classical quantitative point of view, this seems not to be the case. By performing a graphical comparison of the predicted labels (marked in red) with the original true FDIR labels (the telemetry data points that the operators labelled as the start of the anomaly - marked in pink) it is possible to see that the ADAP system was able to detect the anomaly before the classical FDIR actually triggered.

Operators that have performed the original anomaly investigation confirmed that the peaks flagged by the ADAP system are not in the nominal range. hence, the ADAP system was capable of detecting the anomaly at equipment level several time steps before the classical FDIR triggered.

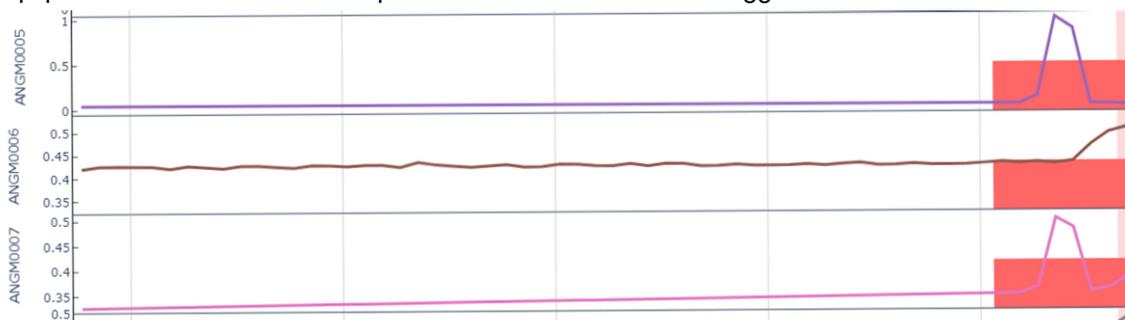


Figure 1 Close-up of the anomaly

Anomaly Prognosis

For the anomaly prognosis part of the ADAP system, three different model types and architectures have been trained:

- 2D convolution classifier
- LSTM classifier
- XGBoost classifier

The reproduction, distribution and utilization of this document as well as the communication of its contents to others without express authorization is prohibited. Offenders will be held liable for the payment of damages.

All rights reserved in the event of the grant of a patent, utility model or design.

AO/1-10612/20/NL/AS

© - Airbus Defence and Space – 2022

	ADAP Project Executive Summary Report	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 9 of 12
---	---	---

All three classifiers work very well and have similar performances. The best model is the XGBoost classifier, but it has the disadvantage of needing a dedicated deployment framework for generating the model application which comes at a considerable license price. The convolution and LSTM classifiers are neural networks like the detection models and therefore can use standard frameworks (for example, Vitis AI) that come without a license or with smaller license fees (in the case of Matlab). The performances of the convolution and LSTM classifiers are almost identical although the LSTM classifier has a slightly better performance, but more important, it is a much smaller model with only 22 081 parameters compared to the convolution classifier with 465 121 parameters. Hence, the LSTM classifier has lower impact in terms of HW.

Generalization

For this use-case, the overall performance of the model once tested with data from a different spacecraft is poor. In general, a performance drop >50% has been observed. This might be due to:

- **Different scales of the channels:** We use the same scaler from the original satellite, but in case the channels are not in the same scale this will lead to bad performance eventually.
- **Incorrect or inconsistent labelling:** This is a general problem for use-case 1, which already affected the performance during training. It is also possible that the model just doesn't generalise well.

Note: Generally, when models perform well on validation data but have poor performance on testing data it might be an indicator of over fitting. However, this assumes that the training, validation and testing data comes from the same distribution and/or population. Since the data comes from different spacecrafts, this assumption is not necessarily true (e.g. different scales of the channels). Over fitting cannot completely be excluded as a reason for the bad performance of the models on other spacecrafts but it is more likely that different scales and especially incorrect labelling play a major role compared to over fitting.

Hardware Porting & Performance

For the AOCS pitch and roll sensor anomaly, a convolutional autoencoder model was deployed on the ADAP system with Vitis AI. Unfortunately, the quantized model for both the anomaly detection and the anomaly prognosis showed a drop in performance of about 30% compared to the non quantized model. Quantized Aware Training (with QAT the model is trained after being quantized to improve on the accuracy loss) has been employed to try and improve the model performance. QAT allowed to reduce the delta-drop in performance to about 25%. Hence, for this use-case, the results once the model was deployed on the edge are poor.

2.2 USE-CASE 2 – PAYLOAD ANTENNA ANOMALY

This use-case is selected since the classical FDIR was not able to detect a unit malfunction (caused by a design error of the unit) and which ultimately lead to a permanent failure of the payload antenna with severe mission impact. The main problem for the detection of this event was that, although temperatures were indeed abruptly rising across the payload subsystem, these were still within the temperature limit ranges of the classical FDIR of all affected thermal lines. During nominal operations, the payload antenna is experiencing significant temperature variations depending on the sun illumination profile along the satellite's orbital position. However, the rising temperatures across payload components were very different from what is expected during nominal operations due to a design failure of the unit which was not detected on-ground. Especially the combinations of affected thermal lines, as well as the very steep increase of the temperatures were very clear signs of a non-nominal payload condition. These ultimately culminated with a corona discharge event causing serious damage of the antenna.

2.2.1 Payload antenna thermal control management & FDIR

The classical FDIR implementation for the thermal control of the antenna unit is fairly simple. This is based in defining an upper and a lower temperature threshold, so defined:

- The lower threshold is selected above the minimum qualification temperature of the unit. If this limit is passed the unit is switched off in order to avoid its damage from operations outside its qualification's lower limit

	ADAP Project Executive Summary Report	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 10 of 12
---	---	--

- The upper threshold is selected below the maximum qualification temperature of the unit. If this limit is passed the unit is switched off in order to avoid its damage from operations outside its qualification's upper limit

The unit's temperature is monitored by specific thermal lines which are used for its temperature regulation (accomplished by heaters in a closed-loop control). For the payload antenna, since the temperature excursion is relatively high and there is no need of a strict thermal regulation to operate it, a wide range for the upper and lower limit has been selected in order to avoid false FDIR triggering in case of nominal operations of the antenna.

In this particular use-case, the gradient monitor was not foreseen since the anomaly is originated by a design failure of the unit which had not been detected on-ground. Without the gradient observability, it is impossible for a classical FDIR design to identify such an anomaly and thus react. As for the ground, since non-contact windows can last up to 24h for Earth Observation missions and up to some days for deep space missions, it is impossible to react on time to such an anomaly due to the latency between anomaly occurrence and anomaly observation (in this specific use-case, the latency was 4h which lead to the loss of the unit).

2.2.2 ADAP system results for UC2

Anomaly Detection

This ADAP system has been deployed to monitor 49 thermal channels. Similar to UC1, for most of the channels, the model flags anomalies some time steps before the operators had flagged the start of the anomaly in the telemetry, and well ahead of the triggering of the classical FDIR. Expert evaluation of the model results confirmed that indeed the anomaly could actually have started before the start point which was identified as part of the in-orbit anomaly investigation. Furthermore, by detecting the trend change in the telemetry well in advance compared to the classical FDIR, the ADAP system could have prevented the corona mass discharge.

Anomaly Prognosis

As this was a one-off failure, the anomaly prognosis part of the ADAP system was not trained nor deployed for this use-case.

Generalization

There generalization of the model was tested using nominal data (all timesteps are marked as HEALTHY) since no anomalies on the thermal sub-system had occurred on the other satellites of the constellation. Since the model accurately identifies the data as fully nominal, there are no false positives. This is reflected in the derived metrics, and the model has an accuracy and a specificity of 100%. Note that in the absence of anomalous data, it is only possible to conclude that the model is capable of correctly identifying the nominal behaviour of the thermal sub-system, but it is not possible to confirm that the model would also perform well in the presence of an anomaly. Hence, the relevance of the results is therefore limited to the nominal behaviours or the sub-system (ground users would not be alerted by false positives).

Hardware Porting & Performance

The corona discharge use-case was selected to be implemented with the Matlab processor of the Matlab Deep Learning HDL toolbox. Since the toolbox offers the HDL block of the processor by itself, it was integrated into the overall ADAP system and interaction was achieved by means of a C application running in FreeRTOS. The network itself is trained in Python, before being transferred to Matlab in order to generate the processor. Despite the Matlab model is not a 1:1 copy of the original Python model due to the internal Matlab software adaptations, it retains a similar accuracy as the original model (performance losses are negligible). The processor itself also retains the same accuracy when deployed in the default Matlab HW image, although there is a drop-off for deployment in the ADAP system. This is an indication for potential issues regarding interference with the memory of the Matlab DLP by other HW blocks. The correct integration of the Matlab DLP into a custom HW image therefore remains an open question.

	ADAP Project Executive Summary Report	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 11 of 12
---	---	--

2.3 USE-CASE 3 - SOLAR ARRAY DEGRADATION

Although this use-case is also based on real satellite telemetry, the anomaly is manually injected to simulate the impact of a micro meteoroid on the solar array, as this was not experienced in orbit so far from the satellite constellation used in this study. However, as known from various other missions, this common event that usually creates irreversible damage to the solar array, causing a full or partial outage of one of its sections.

During the satellite lifetime, the solar array currents show large variations due to many factors, such as the mission phase (e.g. the satellite orientation to the sun), the presence of earth or moon eclipses, the varying solar activities over the year and last but not least the degradation induced by the radiation. All these factors complicate the handling of a monitoring with the classical FDIR, especially as a power drop is just another variation which under some circumstances is actually expected.

2.3.1 Electrical power system equipment management & FDIR

The electrical power system (EPS) of the satellite is one of the most critical on-board and except for failures which are linked to power distribution to specific units (which can be isolated by LCL commanding to open, via a level 1 FDIR) failures in the EPS often imply (for most missions) a Disable non-essential loads (DNEL) alarm with entry into Safe Mode. The DNEL triggering is mostly done not by directly measuring the current delivered by the solar array since this would be a complex and unreliable measurement (due to several effects such as age of the panels, eclipse entry, seasonal effects etc...) but rather through the measurements of the bus voltage or the batteries' depth of discharge. Indeed, these are more reliable estimates of the energy to supply of the satellite for its nominal operations.

Thus for the specific case presented here, using the current state of the art FDIR implementations, there is no possibility to detect in a reliable way such anomaly occurrence.

2.3.2 ADAP system results for UC3

Anomaly Detection

In this particular use-case, the anomaly has been artificially injected thus the labelling of nominal and anomalous telemetry has no ambiguities (subject to operator interpretation as in the previous use-cases). For this particular use-case, the ADAP system exploited the best results in terms of performance metrics and was always capable of correctly flagging the anomaly as soon as the drop of panel current is present. Hence, for this use-case, the ADAP system could be a valuable addition to the space segment as it could inform ground of a degradation of the power system of the satellite as soon as this occurs.

Anomaly Prognosis

As this is considered to be a one-off failure, the anomaly prognosis part of the ADAP system was not trained nor deployed for this use-case.

Generalization

The model generalization was tested using the data of another satellite of the constellation which contained telemetry of a real life anomaly. In particular, due to a wrong operator setting of the spacecraft on-board time (error in the Data Handling System and not exactly in the EPS) the orbit propagator considered the spacecraft in a different orbit position and rotated the solar panels towards the expected position of the sun. In reality the solar array was turned away from the sun and the current dropped to 0 A. The FDIR triggered only after the battery voltage dropped below the FDIR threshold leading to a DNEL alarm which brought the satellite to enter Safe mode.

The ADAP system has been capable of immediately detecting the unexpected drop of the panel current and thus would have warned the ground operator of the anomaly or could have triggered an autonomous recovery action on-board well ahead of the DNEL alarm.

Hardware Porting & Performance

For this use-case, the model is also based on the convolutional autoencoder network, therefore also Vitis AI has been used to deploy the network on the edge. The main difference with respect to the model of use case 1 is the dimensions of its layers. Indeed, with a total of 62145 parameters, this model is more than five times

	ADAP Project Executive Summary Report	ADAP-ADSF-RP-23000051149 Issue: 1 Date: 23/06/2023 Page: 12 of 12
---	---	--

larger as the one of use-case 1 (which has 11293 parameters). Also in this case, the quantization causes a drop in the model performance of about 50% compared to the original model. Nevertheless, for this use-case the quantization aware training was capable of recovering the model performance to matrices comparable to the ones of the original model (for some metrics like specificity and sensitivity, the performances are even better).