Jaliko Ltd.
Invent DCU
Glanevin
Dublin D9
Ireland

Ubotica Technologies Ltd.
Innovation House
DCU Alpha
Old Finglas Road
Glasnevin
Dublin
D11 KXN4
Ireland

Subject: **ESA CCN1 4000124100/18/NL/BJ/gp**

Title: **“CCN1 Space Qualification and Reference Design for Myriad 2 Video Processor"**

Document: **Executive Summary Report**

Prepared by	Aubrey Dunne
Reference	ESR - Executive Summary Report
Issue/Revision	1.1
Date of Issue	10/10/2022
Status	Release

Approval

Document Reference ESR - Executive Summary Report	
Issue Number 1	**Revision Number** 1
Author Aubrey Dunne	**Date** 10/10/2022
Approved by Aubrey Dunne	**Date of Approval** 10/10/2022

Change Log

Reason for change	Issue #	Revision #	Date
Initial version	1	0	21/03/2022
Update with proton results	1	1	10/10/2022

Distribution

Name/Organisation
Dr. Gianluca Furano/ESA
Dr. Antonios Tavoularis/ESA

 ©Ubotica Technologies Ltd. - CONFIDENTIAL

Contents

 ©Ubotica Technologies Ltd. - CONFIDENTIAL

1 Relevant Background and Reference Documents

2 Acronyms

AI Artificial Intelligence
API Application Programming Interface

BOM Bill Of Materials

C&C Command & Control
CAD Computer Aided Design
CAN Controller Area Network
CNN Convolutional Neural Network
COTS Commercial Off The Shelf
CSP Cubesat Space Protocol
CV Computer Vision
CVAI Computer Vision and Artificial Intelligence

DDR Double Data Rate
DR diabetic retinopathy

ECC Error Correction Code
EM Engineering Model
EO Earth Observation

GbE Gigabit Ethernet
GIS Geographic Information System
GSD Ground Sample Distance
GUI Graphical User Interface

HPC High Performance Compute

IR Infra-Red

JSON JavaScript Object Notation

LEO Low Earth Orbit
LTS Long-Term Support

MDK Myriad Development Kit
ML Machine Learning
MWIR Medium Wave Infra-Red

NIR Near Infra-Red
NMF Nanosat Mission operations Framework
NN Neural Network
NNDK Neural Network Dependability Kit

PC Personal Computer

QGIS Quantum Geographic Information System

RGB Red Green Blue

SEE Single Event Effects
SEFI Single Event Functional Interrupt
SEU Single Event Upset
SPI Serial Peripheral Interface
SSO Space Station Orbit
SWIR Short Wave Infra-Red

TID Total Ionising Dose

USB Universal Serial Bus

VIIRS Visible Infrared Imaging Radiometer Suite
VPU Vision Processing Unit

 ©Ubotica Technologies Ltd. - CONFIDENTIAL

3 Introduction

The activity "CCN 1 to Contract No. 4000124100/19/NL/BJ/gp - Space Qualification and Reference Design for Myriad2 Video Processor" has several technical objectives whose goals are to further characterise the Myriad 2 device for space flight, to develop tools and workflows for in-flight Neural Network (NN) deployment, and to facilitate Myriad 2 integration on the MANTIS cubesat mission.

Specifically the technical objectives are:

- The SEE radiation characterisation of the Myriad 2 Vision Processing Unit (VPU) over a range of energies and particles. This objective builds on the de-risking results already obtained, with the aim of conducting further test campaigns at higher energies than those used in de-risk testing.

- The design modification and integration of a Myriad 2 hardware reference board to act as the Artificial Intelligence (AI) and Computer Vision (CV) engine on the MANTIS cubesat mission from Open Cosmos. This objective aims to address the MANTIS usage requirements around image and control interfaces by facilitating their use with a new Myriad 2 UB0100 platform. In particular, the on-boarding of image data via Ethernet, as opposed to Universal Serial Bus (USB), is required.

- The design and development of a workflow engineering platform to enable inference and pre-processing on satellites and to support analysis and response on ground. This objective aims to build and demonstrate a system for end-users to perform workflows enabling on-satellite AI inference, and the optimisation of updated models for uplinking to satellite.

- The development and demonstration of a flare detection algorithm for deployment on Myriad 2. A use case application, this objective aims to develop an in-orbit gas flare detection application.

The UB0100 CubeSat-grade board from Ubotica is built around the Myriad 2 from Intel Movidius, which performs all AI inference and CV compute on the board. Ethernet (new) and USB board interfaces enable board control, transmission of images to the board, and reception of processed results from the board. The board is designed for modular Low Earth Orbit (LEO) cubesat integrations, with on-board latch-up protection and a range of sensing peripherals. Since the UB0100 is a Commercial Off The Shelf (COTS) board that uses commercial/automotive parts exclusively, the radiation characterisation of the Myriad 2 processor is of significant interest to mission designers considering the UB0100, including Open Cosmos.

The Open Cosmos MANTIS mission is a cubesat mission that, among other goals, aims to integrate AI-based cloud detection into its imaging pipeline. The applicability of AI-based flare detection for on-board batch-processing is to be determined, in order to assess flight opportunity. The sensing payload is a Satlantis Red Green Blue (RGB)+Near Infra-Red (NIR) multi-spectral sensor.

During the course of the current project, Open Cosmos assessed the suitability of the Nanosat Mission operations Framework (NMF) for mission operations on-board MANTIS. Hence the

 ©Ubotica Technologies Ltd. - CONFIDENTIAL

workflow engineering platform utilised aspects of NMF in its development.

4 Engineering Workflow

Task 1 of the Space Networks demonstrated an end-to-end workflow for the efficient update of the Myriad Processor Machine Learning (ML) model on board a nanosat, with the update training being done on a cloud High Performance Compute (HPC) system, controlled from within a standard Geographic Information System (GIS) application. In particular we describe the task of Efficient Network Updates: updating the neural model on the Myriad system using a small patch update sent over the air rather than a full network model.

In agreement with ESA, this was done with the assumption that the satellite would be using the Nanosat Management Framework, an Android-like Operating system for Nanosats. A framework was built that allows webservices on a cloud or HPC server to execute training jobs, with images on the server being downloaded and manipulated on a desktop. The Machine Learning models for the Myriad processor on the satellite are then trained on the server and pushed for installation on the satellite.

The ML models are packaged in the form of IPK packages which are apps for the satellite. A key part of the work is that the models can be shipped as patches to an existing model, shrinking their size. For example, we demonstrate the compression of a sample model, 47 MB as an Infra-Red (IR) network, to a 3.6 MB package for transmission over the air to update the satellite. On installation of the package containing the app, the model is patched and compiled on the satellite.

The user control is done by means of a plugin for Quantum Geographic Information System (QGIS), a popular open-source Graphical User Interface (GUI) tool. Within QGIS it is possible to view and manipulate Earth Observation (EO) imagery. The plugin allows the user to interact with a HPC system, uploading and downloading imagery, to set the parameters for a job to be run on a large supercomputer, executing the job remotely to produce an updated model, creating an updated package from the updated model, and installing that network on a system running NMF, the Nanosat Framework, on which a Myriad device is installed.

5 Efficient Network Updates

Since with a new EO sensor proxy data has to be used because there is no training data available, being able to create a more accurate model, using the actual data provided by the sensor after the satellite is in orbit, is highly important.

Thus, the main purpose of this work was to investigate and develop how to update remotely a model placed in a satellite in orbit and try to control or being able to specify the update size using the training parameters. Then, reduce the amount of information required to update a Convolutional Neural Network (CNN) model (in the satellite in orbit) with the result of further training it with additional data (decreasing the file size of the update).

Having a scenario where a model is running in a satellite in orbit and some bandwidth constraints exists, a difficulty occurs when a new version of that model trained with new data has to be updated due to the file size. The tests carried out in this study have shown that it is possible to create a smaller file, using different freezing techniques during training, which can be sent in less time and patched on the satellite.

As summary, in order to update a model in a satellite in orbit remotely, two methods to freeze the layers have been developed and compared in this work: a) the classic freezing of the first layers of the network, and b) a selective method in with the layers chosen to be frozen are the ones in which the weights change less significantly during a complete training phase with no frozen layers. In terms of obtaining a trade-off between accuracy improvement and patch size, the selective freezing was shown to be superior.

The developed methods can be used in conjunction with OpenVINO for the generation of model matches for Myriad devices. The methods demonstrated an improvement on the data bandwidth required to update the model while gaining accuracy on the original trained network: from an original accuracy of 78.4% to an accuracy of 78.9% after an additional training using new data, but obtaining just an update size of 18MB instead of the original model update size of 44.5MB.

6 Flare Detection

Task 2 of the Space Networks project explored both algorithmic and learning based approaches for are detection from satellite imagery. The goal was to produce models capable of on-board inference on the MANTIS platform and to assess the factors affecting the performance and accuracy of such models.

A significant portion of the project involved the collation of a dataset of images which could be used for training and testing the developed approaches. The dataset was compiled using data from the Sentinel-2 satellite, which has bands with wave-lengths matching those of the MANTIS sensor, but with lower spatial resolution. To account for this difference an up-sampled version of the dataset was prepared to give some indication of how well the developed models might generalise to higher resolution images. The Visible Infrared Imaging Radiometer Suite (VIIRS) Nightfire dataset was used to provide approximate locations of ares. However due to the much lower spatial resolution of the data used to compile the VIIRS nightfire dataset this data alone was not sufficient to fully localise the ares in the Sentinel-2 data. A combination of manual annotation augmented by automatically generated annotations followed by manual validation was used for the final dataset. Despite the significant effort invested in compiling this dataset, it was not possible to unambiguously distinguish aring events in all cases. Given more resources, access to domain specific expertise and knowledge of infrastructure locations associated with potential ares events the quality of these annotations could be greatly improved.

To provide a baseline for the learning based approaches, traditional algorithmic approaches were explored. This work consisted of a literature review of the current state-of-the-art in algorithmic are detection. This uncovered that all the existing algorithmic are detection approaches rely heavily on Short Wave Infra-Red (SWIR) and Medium Wave Infra-Red (MWIR) bands, use nighttime imagery and multiple captures taken at different times. None of these are available on the MANTIS satellite as the sensor is limited to RGB and NIR bands, only operates during daylight hours and the approaches must infer based on a single snapshot in time. Despite these challenges a number of different levelbased approaches were trialed. While these managed to detect some of the ares in the test dataset, there were significant numbers of false positives and false negatives.

Two learning based methods were explored; object detection and image segmentation. For object detection the Tiny-YOLOv3 models was used. Object detection models output a set of bounding boxes and associated probabilities for the presence of flares in each. A U-Net model with a MobileNetV2 backbone was used for the image segmentation approach. In order to compare the predictions from the image segmentation models with those of the object detection model it was necessary to adopt a common metric. Using this common metric it was found that both approaches behaved approximately equivalently, with the object detection approach producing slightly higher results on the overall metric, but the image segmentation approach producing a slightly higher recall value, of particular importance in cases that are more sensitive to false negatives than false positives.

The best performing learning based model manages an F1 score of 0:46 on the holdout test set. This means that on average approximately seven out of ten are predictions correspond to actual ares (according to the ground truth annotations) and for each set of ten ground truth are annotations the model successfully predicts three. At first glance these results do not appear very impressive. However when one looks at the predictions, takes into account the fact that many ares do not have a clear signal in the available bands and the difficulties and ambiguities involved in collating the dataset it is clear that the models are performing well in many cases and have a lot of potential in are detection applications if the issues highlighted in this report are addressed. The experiments and examples presented throughout this report provide more details and context and outline steps that can be taken to improve the models.

Timings were taken for the models on the Myriad 2 and show that the object detection Tiny-YOLOv3 models perform approximately three times faster than the U-Net model with MobileVNetV2 backbone. Taking these timings into account and the comparable performance, the object detection approach appears to be the more promising model for deployment onboard MANTIS.

7 Myriad 2 Proton Characterisation

Myriad 2 has already gained space heritage through its flight on PhiSat-1, having previously undergone Single Event Effects (SEE) testing at CERN and GSI, and Total Ionising Dose (TID) testing at ESA ESTEC. In order to realise further in-flight space applications using Myriad 2, its susceptibility to proton radiation effects, as experienced in space, was assessed. The purpose of this task was to continue the de-risk radiation characterisation of Myriad 2 by developing a broad-spectrum test codebase and performing a proton test campaign to gain further insight into Myriad 2's radiation sensitivity in LEO.

Four categories of tests were conducted – memory, functional, SIPP and inference – in order to achieve a comprehensive device characterisation. Nevertheless, complete coverage of any of these categories would require further testing.

Across all tests, the Single Event Functional Interrupt (SEFI) cross section is consistently within approximately an order of magnitude centered around 1e-9 cm2/device (with the exception of the tests for which no SEFIs occurred). The Single Event Upset (SEU) cross sections per bit for the memories are consistent with results seen from previous campaigns, where the Double Data Rate (DDR) is several orders of magnitude more resilient than all other memories and caches in the system due to the included Error Correction Code (ECC).

The inference cross sections indicate that the StarTracker network is more susceptible to bit errors than the larger networks. This is possibly due to the more extensive use of caching for the smaller StarTracker network. The StarTracker network (i.e., the weights) is only 308KB in size, whereas the DR and MobileNet networks have weights footprints at least 30x larger in size. An effect of this is that the weights for the StarTracker network can substantially fit within the caches, whereas for the other two networks the caches would effectively be flushed between inferences, since the weights are orders of magnitude larger than the caches. The cross sections of the caches (excepting the SHAVE L2 cache) are one to two orders of magnitude greater than that of the DDR and CMX, implying a higher inference cross section for those networks whose weights are persistent in the caches.

The cross section variation with energy does not, in general, reveal the knee point that would enable Weibull fitting and cross section saturation determination. This is most likely as low enough energies were not used, and therefore the low energy dropoff is not seen. As indicated by the error bars in the cross section plots, the plots always show the saturated cross-sections. Post-campaign testing of the DUTs showed persistent bit errors in the DDR. Although there are relatively few bits in error (as a percentage of the entire DDR), their presence does indicate permanent damage in the memory cells themselves, or perhaps in the controller. The errors are repeatable in that they occur for the exact same bits for a given board, but do not necessarily occur in a given order or given timeframe after each restart. The error rate was also shown to be proportional to the temperature. Leakage current in the DRAM cells increases with increasing temperature, and thus the retention time of the DRAM cells is shortened. An experiment to increase the DRAM refresh rate was conducted, but as there was no method of confirming that the refresh rate change was effected, a conclusive result could not be drawn.

The Neural Network Dependability Kit (NNDK) analysis applied to the diabetic retinopathy (DR) inference tests showed an excellent ability to detect undependable inference classifications resulting from SEUs in the network (specificity of 1.0), although the sensitivity is considerably lower. However, this sensitivity assumes that the network was operating correctly for all 'correct' inference classifications, which may not necessarily be the case.

Upset rates for an Space Station Orbit (SSO) orbit for the primary memories are presented, indicating a predicted upset rate of less than 4 bits per day. The predicted upset rate for the DR (InceptionV3) network in the same orbit is approximately 1x10-5 upsets per device per inference.

A separate Proton test campaign, conducted on Myriad X at PSI, was also performed.

 ©Ubotica Technologies Ltd. - CONFIDENTIAL

8 MANTIS Integration

8.1 Hardware

The UB0100 integration with MANTIS requires an element of mechanical integration, and a mechanical interface specification. Furthermore, the MANTIS UB0100 is a custom build, as is necessary to accommodate the connectors desired by Open Cosmos. The PC104 header on the UB0100 enables PC104 stacking, but in the case of MANTIS this stack functionality is not used. Instead, certain signals are interfaces directly to the board edge header pads, and other interfaces connect via the function-specific dedicated headers on the UB0100. In collaboration with Open Cosmos, the specific header components and their signal coverage were determined, and support for the development of a custom mechanical interface (tray) was provided. The custom tray encloses the UB0100 such that only the necessary connectors are exposed while minimising the total volume of the tray, and was facilitated by the development of a metrically accurate UB0100 Computer Aided Design (CAD) model.

In order to facilitate Controller Area Network (CAN) on the MANTIS mission, hardware updates to the UB0100 CubeSat-grade board were made to incorporate a CAN interface. The original UB0100 is an AI and CV processing engine in which all inference and vision processing requests are made over one of two primary data interfaces: Ethernet or USB. However, as part of the integration into the Open Cosmos MANTIS mission there was a request for a CAN interface for Command & Control (C&C). CAN functionality was successfully designed into the UB0100 board, using the MCP25625 CAN transceiver device. Power gating for CAN was integrated to enable overall board power consumption reduction when CAN is not in use. All bringup tests, both electrical and functional, passed on the prototype UB0100 test board, and the transfer of Cubesat Space Protocol (CSP) packets between Personal Computer (PC) and UB0100 was successfully demonstrated.

8.2 Software

The primary inference and boot interface to the UB0100 on the MANTIS platform is Ethernet. Therefore an Ethernet protocol was added to the CVAI Toolkit™ for execution on the UB0100. The CVAI Toolkit™ is a collection of interface functions that enables host applications to perform inference and image transformation operations using a connected Myriad 2 to accelerate these operations. The library is lightweight, written in C, and supports Myriad 2 devices connected over different communication protocols. The Application Programming Interface (API) software stack structure is composed of two main layers, MVNC and XLink. The XLink component handles all the communications in the host through a set of generic API functions. The MVNC exposes the inference, the management of the Myriad 2 and the Computer Vision and Artificial Intelligence (CVAI) API to applications. Adding Ethernet to the CVAI Toolkit™ involved the completion of two main tasks, which had to be carried out at each of the two layers: support for a new protocol was added, and then the Ethernet protocol itself was added.

The Ethernet interface build that was added to the CVAI Toolkit™ was updated to the OpenVINO 2020.3 Long-Term Support (LTS) version, which provides longer-term maintenance and support with a focus on stability and compatibility. The OpenVINO 2020.3 LTS version is aligned with the API and the firmware from the Myriad Development Kit (MDK) R11 release.

This MDK is only available for the Myriad X, so it was required to adapt these two components to work in the MDK 19.07, which is the last MDK for the Myriad 2. In order to accomplish this objective, three tasks were completed. Firstly, the API from OpenVINO 2020.3 (MDK R11) was adapted to the MDK 19.07. Secondly, the new firmware and capabilities from OpenVINO 2020.3 (MDK R11) were adapted to the MDK 19.07. Then, the new firmware and capabilities were adapted for their use in the Myriad 2, enabling the support of new layers, networks and models. Thirdly, the Ethernet protocol was added to the adapted API and firmware from OpenVINO 2020.3 (MDK R11). Finally, the Ethernet protocol was added to the adapted API and firmware for the Myriad 2 and the host.

Several methodologies and approaches were used to optimise the Ethernet protocol performance, required due to the strict round-trip times for an inference call from the sensor payload to the UB0100 on MANTIS. The main objective was to improve the image transmission stage, when the host sends the input tensor to the Myriad 2. After doing a thorough analysis and experiments, the main techniques used were the use of Gigabit Ethernet (GbE), disabling the Nagle's algorithm, and the modification of the maximum Myriad 2 window size. Together these modifications led to a significant increase in the efficiency of the Ethernet implementation for inference on the UB0100.

A custom SPI-CAN driver was developed for the Myriad 2, so it can interface to other devices via CAN, using the MCP25625. It is a stand-alone CAN 2.0B Controller with Integrated CAN Transceiver and Serial Peripheral Interface (SPI) interface, which is the one built into the UB0100 board. The first step was to provide the Myriad 2 with a SPI master driver allowing it to control the device. Afterwards, the logic needed to configure and effectively use the MCP25625 (CAN driver) was added. This work resulted in a fully functional Myriad 2 driver to control and use the MCP25625 to expand it communications capabilities, serving as the basis for the new CSP CAN driver.

The CSP CAN protocol is a requirement for MANTIS, and therefore its integration in CVAI Toolkit™ was necessary. It is a small network-layer delivery protocol designed for Cubesats, that allows distributed embedded systems to have a service oriented network topology. Although CSP already has a CAN driver, it was based on the SocketCAN library so it was not compatible with Myriad 2. Hence, work was conducted to add CAN to CSP on the platform - the MCP25625 driver was integrated into the CSP stack and added to the Myriad 2 MDK, providing the CVAI Toolkit™ with a powerful communication system for use on MANTIS. It enables any device to communicate to the Myriad 2 on the CAN bus using CSP packets. The integration in CVAI Toolkit™ is threaded, allowing simultaneous C&C and inference requests to the UB0100.

The University of Pisa is developing a cloud detection NN model to be deployed on the UB0100 board on MANTIS. Regular meetings were held to discuss the development of the model, with input provided by Ubotica where necessary. Development of the cloud detection network by the University of Pisa is ongoing and is behind the original project schedule, hence further input into supported layers will be provided by Ubotica as requested (beyond the scope of the current project). When the final model is available it will be benchmarked on the UB0100, in conjunction with CVAI Toolkit™, for both throughput and power. The CVAI Toolkit™ directly provides execution time - broken into tensor transfer time and execution time. Power measurement is via the onboard INA226 current sensors on the UB0100. Drivers for these sensors have been developed to enable this power assessment. The general task of supporting pre-processing for NN inference was addressed within this project, being possible independently of the cloud detection model. The CogniSatApp was developed to allow for the deployment of applications involving image processing and AI. The use of JavaScript Object Notation (JSON) configuration files allows for easy deployment of new applications without the need to write/alter any code. The design flow of the CogniSatApp addresses pre-processing in a generic manner, allowing for no or multiple pre-processing operations to be deployed. This solution was developed in order to facilitate the easy development of NN pre-processing for the MANTIS applications in advance of the final pre-processing steps being defined (these are still pending).

Finally, validation and testing was performed of the software running on the Ubotica UB0100 CubeSat grade board which is part of the MANTIS platform. All inference and vision processing requests are made over one of two primary data interfaces: Ethernet or USB. However, it also includes a CAN interface for C&C. Both interfaces, Ethernet and CAN, were verified to be fully functional on the MANTIS Engineering Model (EM). In addition, the firmware used for inference and traditional CV operations, as well as the bootloader to be integrated into the MANTIS platform, have been fully verified.

9 Conclusion

This project addressed several bodies of work related to on-orbit AI for cubesats.

A flare detection AI solution was developed, in the form of both flare detection and flare segmentation models. Results indicate a reasonable level of performance given the lack of available labelled training data. The fact that labelled data was only available at 10m Ground Sample Distance (GSD), while the MANTIS sensor has a GSD of 3.5m or less, was identified as a key inhibitor to good performance. Experiments were carried out to estimate the model performance when trained at one GSD and applied at a different GSD, in order to provide information to support the decision around flying the flare detection model on MANTIS.

An engineering workflow was developed to facilitate efficient NN model updates in flight, in a generic sense. This incorporated the development of scripts running on HPC resources that enable the training, and re-training, of generic AI models (in this case based on Keras). Training is managed via a unique method for generating model updates that is cognisant of the total update size. A patch size limit is specified by the user, and depending on whether the IR or the final model blob is patched, the training proceeds to generate optimum model updates given these constraints. The engineering workflow leverages the NMF package to support uplink and on-satellite deployment. The entire flow was validated for a sample NN model on a Myriad 2.

A proton test campaign was conducted for Myriad 2 in order to asses its susceptibility to SEEs. The test campaign was delayed due to COVID-19, and was ultimately performed remotely at PSI via remote control from Ireland. The campaign focused on Myriad 2 memories, peripherals, and function tests (with an emphasis on inference). The results complement previous heavy ion test results for the Myriad 2. SEFI and bit cross sections for each test were determined, showing cross sections approximately four orders of magnitude lower that those calculated for the heavy ion tests. Cross sections for the InceptionV3 model were used to estimate upset rates in SSO, with rates of 2.76 SEEs/device/day, or $1x10^{-5}$SEEs/device/inference.

The MANTIS cubesat mission from Open Cosmos is due for launch in Q3 2022, and the work conducted in this project involved advancements made to facilitate the integration of the UB0100 board and NN applications onto MANTIS. There is a focus on mechanical and electrical interfaces, and on pre-processing, test and validation. For MANTIS, two types of hardware customisations have been accommodated: electrical interfacing, consisting of a custom Bill Of Materials (BOM), and mechanical interfacing, supported by a custom CAD model of the UB0100. The final BOM represents a minimal configuration that satisfies the functional coverage requires on MANTIS (Ethernet for inference C&C, CAN for general C&C, power, and low-level signals for status and health indication). On the software side, Ethernet support, CAN support, CSP support, and JSON-based generic pre-processing support (via the CogniSatApp) were developed to aid with software and application integration. Test and verification of the MANTIS EM was performed for the UB0100 prior to, and subsequent to, shipment to Open Cosmos. Verification covered the USB, Ethernet and CAN interfaces, and included the execution of standard test applications that implement both image pre-processing and NN inference on the UB0100. Simultaneous inference and CAN C&C requests was also verified to operate correctly.

 ©Ubotica Technologies Ltd. - CONFIDENTIAL