REUSABLE SAMPLE TUBE ASSEMBLY (RSTA) ACQUISITION SYSTEM: PICKUP AND STOWAGE SYSTEM DEVELOPMENTS IN SFR MISSION CONTEXT

Chris Hackett*[(1)], Raul Arribas[(1)], Gianluca Borgo[(1)], Todd Cooper[(1)], Michael Dinsdale[(1)], Yvonne Pickering[(1)], Timothy Smith[(1)]

(1) Airbus Defence and Space Ltd., Stevenage, SG1 2AS, United Kingdom

*Corresponding Author - *chris.hackett@airbus.com*

ABSTRACT

Within the Mars Sample Return Campaign (MSR) the role of SFR is to retrieve samples deposited by the NASA-JPL Mars 2020 Perseverance Rover as quickly as possible. The vast distance between Earth and Mars creates a natural lag in communication and introduces the need for the system to be autonomous in its operation. The Returnable Sample Tube Assembly (RSTA) Acquisition System (RAS) consists of hardware and software necessary to perform the autonomous detection and acquisition of the sample tubes once the rover platform has been migrated to the sample depot.

1. INTRODUCTION

The Sample Fetch Rover (SFR) was studied by Airbus Defence and Space Ltd (ADS) in Stevenage, UK, for the European Space Agency (ESA) in the context of the Mars Sample Return (MSR) Campaign, a collaboration between the National Aeronautics and Space Administration (NASA) and ESA, with the objective to return samples from the surface of Mars for scientific study on Earth.

The campaign is based on three missions:

1. The caching mission, currently being carried out by NASA's Perseverance rover in Jezero Crater, to collect and seal samples inside RSTAs that are deposited in a depot for future recovery.

2. The retrieval mission, expected to launch in the late 2020s, to recover the RSTAs and launch them into Mars orbit.

3. The return mission, also expected to launch in the late 2020s, to capture the orbiting RSTAs and deliver them to Earth.

In the 2018 international MSR conceptual architecture, the retrieval mission consisted of a lander carrying SFR and the Mars Ascent Vehicle (MAV). SFR, part of the ESA contribution to MSR, would fetch up to 30 RSTAs from the depot and deliver them to MAV. The combination of orbital dynamics severely limiting the time window for fetching the RSTAs and the large distance between Mars and Earth (up to 22 light minutes) necessitated that SFR be highly autonomous.

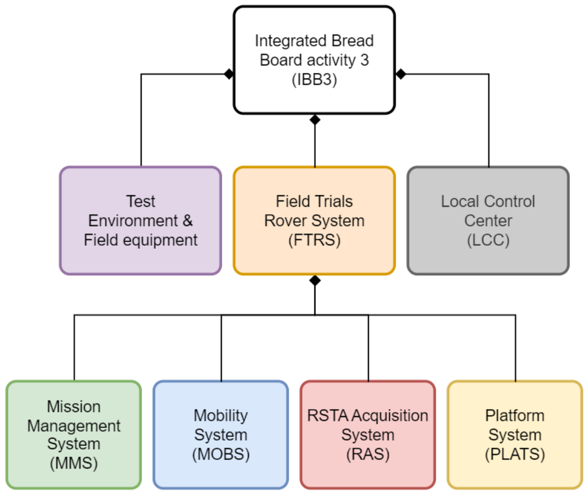

Figure 1. General chart of the IBB3 system

In 2022, the retrieval part of the campaign underwent major revision, leading to the removal of SFR and termination of its development. This was to reduce complexity and cost, driven by a combination of the space agencies reconsidering their plans for international collaboration and a considerable increase in the predicted probability of Perseverance being able to deliver the RSTAs directly to the MAV.

Despite the SFR mission cancellation ESA elected to continue development of the key SFR fetching technologies breadboards with the aim of demonstrating an end to end autonomous traverse and fetch capability in field trials conducted in a suitable Mars analogue environment, such as a quarry. This is the Integrated Breadboard 3 (IBB3) [1] project, executed by ADS for ESA. The development was kicked off in January 2023 to be concluded before the end of 2023. The field trials in September 2023 are the primary means of system validation, demonstrating the necessary capabilities of fetching RSTAs within a martian environment analogue.

Fig. 1 shows the constituent parts of the IBB3 project, with the RAS specifically being the topic of this paper. The RAS, along with the Mobility System, Mission Management System (MMS), and Platform System, form the Field Trial Rover System (FTRS) as the platform for the field trials.

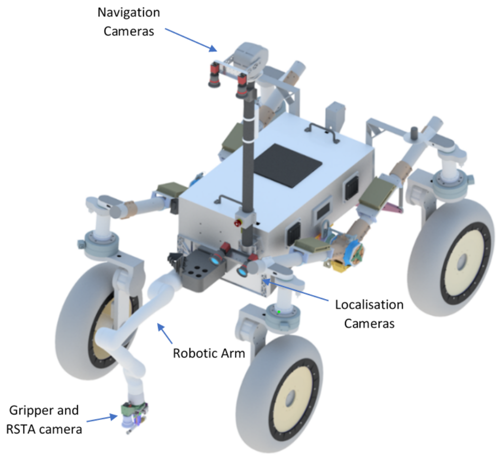

Figure 2. Overview of FRTS systems

The sequence for a typical RSTA pickup is as follows:

1. Operators, in advance, plan paths for rover to drive, from depot entry point, to each RSTA
2. Plan is sent to rover where the Mobility system autonomously guides the platform to each RSTA
3. When the rover has arrived at an RSTA the RAS is commanded to detect the RSTA on the surface
4. If not detected, the Mobility system guides the rover to reverse and approach the RSTA from another direction and for RAS to repeat detection
5. If detected, the RAS controls the robotic arm to collect the RSTA and stow on the platform
6. Once the RSTA is stowed, the mobility guides the platform to the next RSTA until no more RSTAs remain in the depot

2. FIELD TRIAL ROVER SYSTEM (FTRS)

The FTRS, depicted in Fig. 2, is a four-wheeled mobile platform equipped with multiple sets of cameras for visual localisation, navigation and RSTA detection tasks, with a Robotic arm for pickup and handling of RSTAs. All its wheels are steerable, making FTRS capable of performing Ackermann and point turn manoeuvres.

The RAS is mounted onto the front of the FTRS with the wheels driven to a 45 degrees position 'duck pose' to provide the largest workspace when operating the arm. Fig. 3 shows the wheels in RSTA pickup ready configuration.

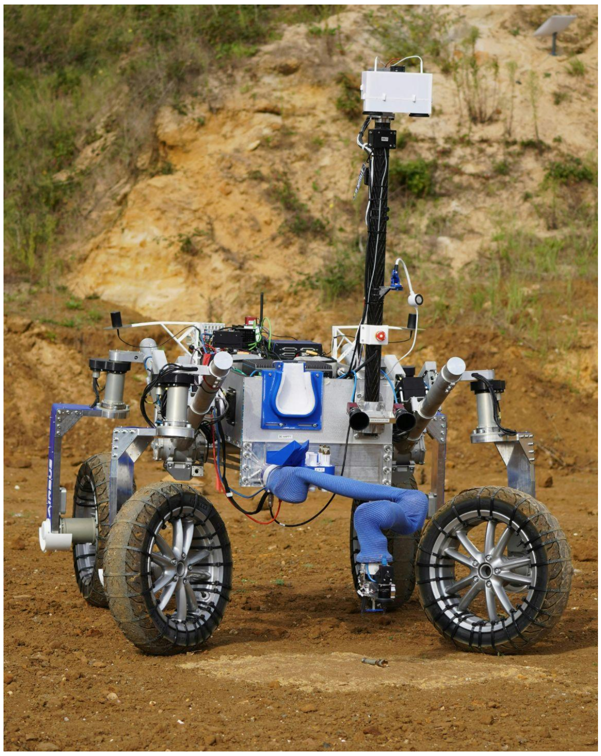

Figure 3. FTRS during RSTA pickup

The RAS and other FTRS subsystems are commanded by the FTRS Mission Management System (MMS) which queues up commands and handles failures raised by its subsystems.

3. RAS EQUIPMENT AND SUBSYSTEMS

The RAS comprises all the functions and equipment necessary to perform the end-to-end sequence from the identification of the RSTA on the ground up to the safe storage onboard the rover. Its operations are initiated once the rover has arrived at the expected RSTA location and include:

- the identification of the RSTA in the scene in front of the rover
- the planning of gripper poses to grasp the RSTA
- the deployment of the manipulator and its gripper to the RSTA and grasp it
- the re-grip of the RSTA at the RSTA Re-grip Bracket (RRB) to facilitate insertion into the storage location
- the insertion of the RSTA into the RSTA Storage Assembly (RSA)

The RAS consists of a suite of equipment, software applications and hardware including

- Acquisition Management Software (AMS)
- Vision Based Detection (VBD)
- Arm and Gripper Sub-System (AGS)

3.1 Acquisition Management Software (AMS)

The AMS function is the 'glue' that connects all the various hardware and software components of the RAS system.

The AMS coordinates the RAS subsystems through the acquisition sequence by breaking it down into a series of functional blocks. Transition to the next step is nominally triggered automatically upon completion of the previous step.

These functional blocks are the atomic steps of the AMS, and not the TCs exposed to external controllers through the remote interface (see fig. 4). TCs will contain a number of atomic steps, and steps within a TC are progressed though autonomously so long as the exit conditions are met. Failure to meet the exit conditions of an atomic step will result in a no-success being returned to the sender of the TC.

The extent of functionality encapsulated into each TC has been designed to allow flexibility in the way the pickup is run. With the software managing RAS able to handle failures and continue autonomously from different points in the sequence.

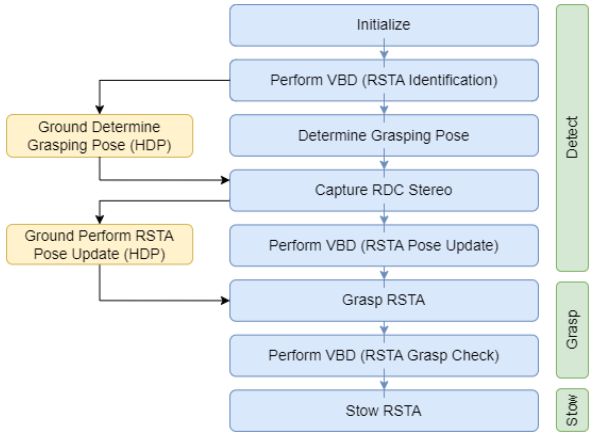

Figure 4. Simplified AMS State Machine. Blue indicates the nominal TCs available with yellow signifying the HDP alternative TCs. Green indicates which phase of the pickup each TC resides

Human Directed Pickup (HDP) TCs are available for ground operators to partially remove some of the autonomy.

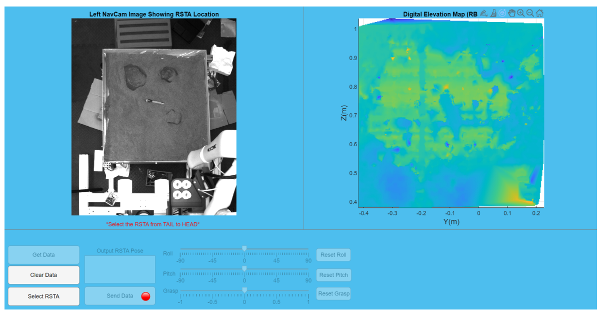

Figure 5: Screen capture of Navcam HDP Tool. Left window shows the Navcam output, and the right window shows the generated point cloud

RAS data products can be interpreted by a series of tools (see fig. 5) have been developed to read data from the interface and aid interpretation. The human operator has access to the data products necessary to assess how they'd like to command the RAS in challenging situations.

3.2 ROS2 Interface

RAS commanding and telemetry passes over a ROS2 interface [2] to the AMS (see ROS node in fig. 6). This type of software interface is used throughout rover design in the IBB3 campaign [1] in order to combine the systems onto the rover. Making communication simple for both the autonomous Rover level commands as well human directed pickup (HDP) scenarios.

The modular architecture combined with the level of autonomy allows the RAS to be compatible with many future platform demonstrations and missions.

3.3 Arm and Gripper Sub-System (AGS)

The arm and gripper subsystem consists of the manipulator, gripper, RSTA Regrip Bracket (RRB) and RSTA Storage Assembly (RSA).

The robotic arm being used for IBB3 is a Kinova Gen3 6Dof variant without the camera. The arm is connected to the FTRS by a machined aluminium piece, which also attaches the RSA and RRB. Attached to the end-effector of the arm is the gripper and the RSTA Detection Camera (RDC).

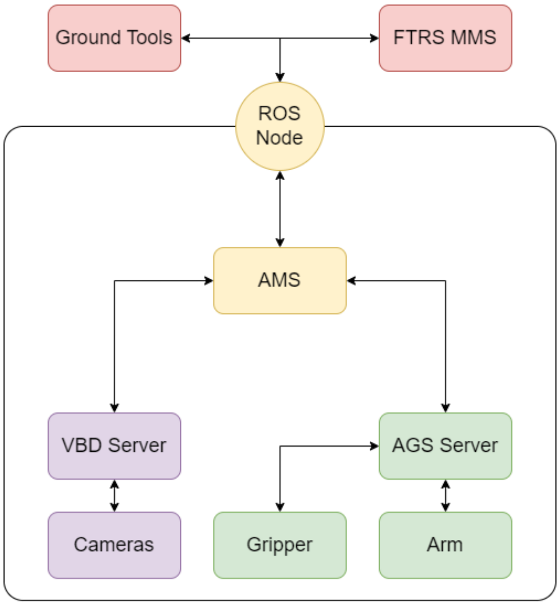

Figure 6. AMS interfacing to system with ROS node and RAS subsystems

The selection of manipulators was made primarily due to the mass constraints of the platform. Many industrial arms long listed were far too heavy for the platform and would have impacted its drivability. For much of the development a UR5 robotic arm was used so a 6 DoF arm was an easier transition, over the 7 DoF variant.

The Kinova arm provides an API that allows for high and low level control of the arm. The RAS uses velocity commands to control the arm to follow the desired trajectory over time. At the time of development there was only a ROS1 driver available for Kinova, so a custom driver was coded to allow the arm link with the ROS2 infrastructure. The planner uses MoveIt to plan safe motions around the bounding boxes that represent the rover, arm and tool.

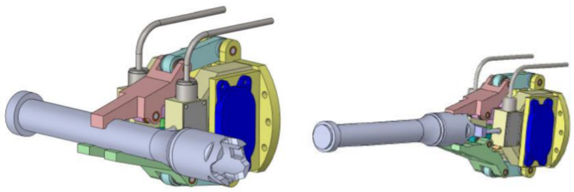

Figure 7. Gripper grasping configurations Left. Body Right. Head

The manipulator's end effector accommodates a gripper-type end effector (developed by AVS under contract by Leonardo), whose jaws are designed to cope with the external geometry of the RSTA. The gripper was designed mechanically and kinematically to grasp the RSTA both transversally (body grip) and axially (head grip), see fig. 7.

The purpose of the RRB is to provide a support where the robotic arm can deposit the RSTA and perform a regripping, i.e., move from Body grip to Head grip. The bracket has been designed to provide passive compensation of positional errors. Dropping the RSTA in from height allows the shape of the bracket to guide the RSTA into a known location which can be extracted using a fixed configuration of the arm.

Figure 8. Top view of 3D printed RSA and RRB with one RSTA slot filled

The RSA allows for up to four RSTA to be stored on the platform. The RSTA sheaths are slightly flexible at the top to allow for positional inaccuracy of the tail. The top of the sheath will deform and once the RSTA has been released from the gripper, it will centre itself.

3.4 Vision Based Detection (VBD)

The VBD system, developed by GMV, is a key enabling technology for autonomous pickup, providing three functions used during pickup:

1. RSTA Identification: Uses the navigation camera images to identify the location of the RSTA within the defined pickup area, and to provide a 5DOF estimation of the pose of the RSTA

2. RSTA Pose Update: Uses the RSTA Detection Camera (RDC), installed on the robotic arm, to take two images of the RSTA and from this produce an estimation of the 5DOF pose of the RSTA to refine the placement of the RSTA with respect to the arm

3. RSTA Grasp Check: Uses the navigation camera images to detect the presence of the RSTA in the gripper and to determine if it is held correctly using orientation and linear (1 DoF) position

The SFR data handling solution imposed strict memory usage and CPU load limitations on the VBD while the overall MSR fetching timeline imposed VBD execution time constraints. These strongly influenced the VBD algorithms selection and this is reflected in the BB VBD used for IBB3.

RSTA identification is achieved through semantic segmentation by means of Convolutional Neural Networks (CNNs). Image segmentation consists in pixel-level labelling with a set of object categories for all image pixels, grouping together similar parts of the image that belong to the same class. In the specific case of RSTA identification the categories are only two: RSTA and background, as shown in Fig. 9.

Figure 9. Image segmentation of an RSTA

After the segmentation has been performed, the next step of the algorithm is the identification of the longitudinal axis of the RSTA in the image plane. This is achieved by Principal Component Analysis (PCA), which identifies the longitudinal axis and centroid of the RSTA. Following this, the head direction is estimated by comparing the segmented RSTA with a binary template of the RSTA as seen in Fig. 10. From main, minor axes, centroid and direction, the 2D pose is obtained.

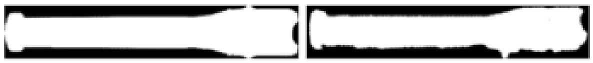

Figure 10. Binary template of the RSTA vs the segmented RSTA

From the 2D pose, the 5D pose is estimated by making use of the point cloud which is then output to the AMS together with an identification success flag and software status flag.

The RSTA pose update function takes as input two RDC images taken at two different poses, see fig. 12 and outputs an updated 5D pose estimate. For each RDC camera image, the same CNN segmentation and pose estimation process used for the NavCam identification is performed and the 2D pose in each of the images is derived.

The grasp check algorithm has the following objectives:

- Confirm that an RSTA is in the gripper
- Confirm whether the RSTA is flipped in the gripper
- Estimate the position of the RSTA relative to the gripper along the RSTA axis

The first step is the identification of the position of the gripper finger fiducial markers in the image (fig. 11) which is performed by means of template matching.

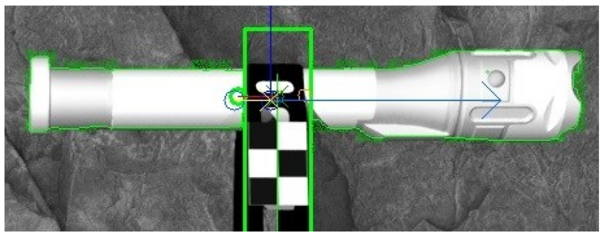

Figure 11. Estimation of the RSTA grasp offset. Fiducial markers on found on the gripper fingers

Then the CNN segmentation and pose estimation algorithm return whether an RSTA is in the image and return the centroid of the RSTA. From the knowledge of the RSTA centroid and marker position, together with the known distance from the NavCam to the

RSTA, the relative position of the RSTA with the marker is known.

4. RSTA PICK-UP SEQUENCE

The purpose of the RAS is to pick-up the RSTA. The mobility system delivers the platform to the approximate location of the RSTA, and then the RAS is commanded to perform the pick-up sequence.

This sequence can be broken down into three phases,

- Detect - Use the cameras to provide a accuracy estimation of the RSTA
- Grasp - Use the manipulator to grasp the RSTA and determine its shift in the gripper
- Stow - Change from the body grip to head grip and deposit the RSTA in the RSA

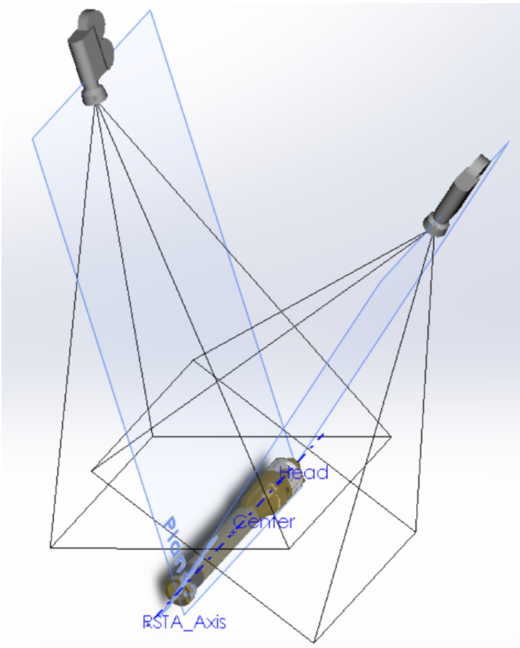

Figure 12. The ideal position for both RDC pictures

4.1. Detect Phase

At the start of the Detect phase the navigation cameras are set to an angle pointing directly downwards. A stereo image is then captured by the navigation cameras which can be used by the Visual Based Detection (VBD) software to identify the RSTA in the terrain in front of the rover. This provides a position of the RSTA with respect to the camera's coordinate frame. This is then transformed to give an initial position wrt the base of the robotic arm.

To increase the accuracy of the pose estimation, two subsequent pictures are taken using the RDC at closer proximity. A stereo pair is reconstructed by taking images 30cm above the initial position of the RSTA and at an angle of 22.5 degrees either side of the nadir, as shown in Fig. 12. These two images are then used by the VBD to produce an updated position of the RSTA wrt the centre of the two images. This estimate can again be transformed to the frame robotic arm base using the forward kinematics of the arm to provide the pose where the images were taken.

4.2. Grasp Phase

The calculated position of the RSTA is used to generate a linear path of waypoints, for the arm to follow, to intercept the RSTA on the ground. Due to depth knowledge error, the final path waypoint is beneath the estimated ground plane. Fig. 13 illustrates the path taken by the arm.

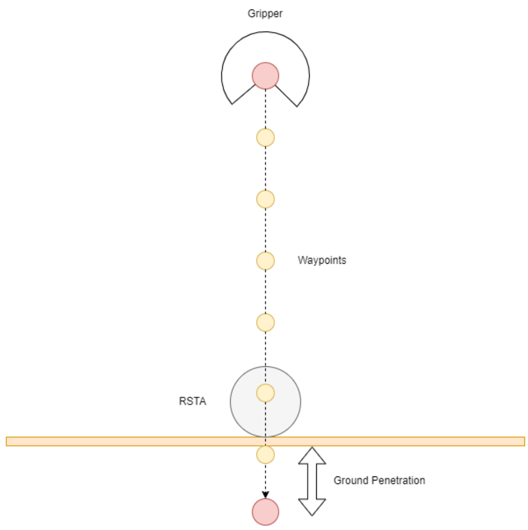

Figure 13: RSTA grasp approach waypoints. Multiple points enforce a linear motion.

Prior to following the grasping path the arm is moved to an intermediate position above the presumed location of the sample and opens the gripper.

Once the gripper has fully opened the robotic arm is commanded to follow the grasping path. Before the final waypoint is reached, collision with the RSTA should occur. This is detected by the pressure plate and triggers the arm to stop and the gripper to close. Therefore, reaching the final waypoint of the grasping path constitutes a failure as it indicates an error with the VBD.

Once the gripper has securely closed around the sample, the arm will move to a position directly below the navigation cameras and will be oriented as shown in fig. 11. The VBD will then trigger and determine the relative position with respect to the gripper. If no sample is detected, then an FDIR will be triggered.

4.3. Stow Phase

The stow phase begins after a successful grasp check. The RRB is used to change the grip from body to head. Depending on the grasp offset calculated at the grasp check, 1 of 3 predetermined poses are used to drop the RSTA in the RRB. If the RSTA is grasped further down the shaft, then the arm is positioned further back in the RRB before opening the gripper.

The RRB design is such that once dropped in the RRB the RSTA will settle such that the RSTA head is in a single position and attitude. Hence once the RSTA is dropped into the RRB, all subsequent arm motions are predetermined. The RSTA is extracted from the RRB, grasping the RSTA's head, then inserted into 1 of the 4 RSA sheaths. After insertion, the arm is parked into the cradle and an image is taken by the Navcams to be used for ground confirmation of the successful pickup.

4.4. Failure Detection and Handling

The phases described previously detail the nominal behaviour for the RSTA pickup however during operation of the RAS there are multiple reasons why a pickup can fail. It is desirable to be able to detect and recover from failures. Such failures can be grouped into either nominal or unexpected.

Nominal failures can be anticipated with the autonomous sequence handled by the higher level logic outside of the RAS. Examples of Nominal failures include:

- RSTA Not detected (Navcam)
- RSTA out of reach
- No RSTA detected at Navcam Grasp check.

In these cases the RAS will flag a FDIR and in combination with the higher level FTM, action the mode-management to attempt a recovery.

Unexpected failures are created by unexpected behaviour of the RAS. All TCs are checked for completion conditions that must be met for the TC to return success, and any condition not met will produce a unique FDIR flag that can be used by the external controller to debug and resolve. Examples of unexpected failures include:

- Unexpected contact of robotic arm
- RSTA Not detected (RDC)
- No collision free trajectory available

5. WORKSPACE ANALYSIS

A successful pickup requires that the RSTA is situated within the working area of the RAS. Since the FTRS must traverse to the site of the RSTA, its location with respect to the workspace is variable, and dependent upon the accuracy of the navigation. If navigational errors are larger than the size of the workspace, there is a risk that the RAS may be unable to access the RSTA for one or more stages of the acquisition sequence.

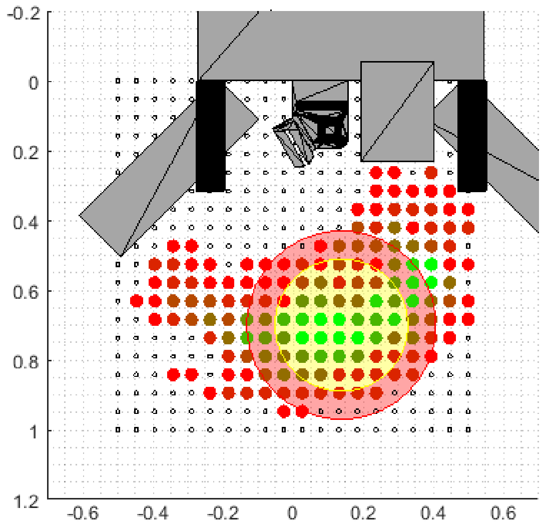

Figure 14: Top-down view showing highly constrained workspace in front of the FTRS. (Green indicates good accessibility, red indicates poor accessibility)

The workspace is limited to the work envelope of the arm, which is dependent upon multiple factors: The maximum reach of the arm, collisions between the arm and its environment (including the terrain and FTRS), and self-collisions of the arm and its attachments. All of these constraints must be satisfied while the arm performs detection using the RDC, and throughout the grasp manoeuvre to ensure the RSTA is accessible.

While the grasp manoeuvre can accommodate different RSTA orientations by rotating the end effector, the RDC detection poses must always be taken at an offset perpendicular to the RSTA. As a result, the accessible workspace will vary according to the RSTA orientation, and the workspace assessment considers 8 RSTA orientations at equal intervals of 45 degrees.

The worst case RAS workspace, shown in fig. 14, illustrates the small size of the accessible working area compared to the expected (yellow circle) and worst case (red circle) mobility errors. RSTA locations were considered in a 20x20 equidistant grid spanning a 1mx1m area directly in-front of the arm. All 8 RSTA orientations were found to only be accessible in the green region towards the centre of the workspace, while regions trending towards red become inaccessible unless the RSTA is in specific orientations.

Note that the flight SFR would have had arm limb lengths, mounting and joint positions optimised to match the arm work envelope to the RAS workspace but such optimisation was not possible with the use of the off-the-shelf Kinova arm on the FTRS.

Further investigation into increasing the accessible workspace could assess the performance of the VBD with modified RDC image poses, to determine if the position and orientation of these poses can be relaxed, while still achieving adequate accuracy to perform the RSTA pickup-sequence.

6. FIELD TRIALS EXPERIENCE

At the time of writing this paper, the IBB3 field trials have just been completed at a quarry in Leighton Buzzard, Bedfordshire, UK.

Although the detailed test results are not yet analysed, the end to end capability of the FTRS to approach, detect, grasp and store RSTAs has been demonstrated under ‘real world’, varying, Martian analogue conditions.

Some problems have been encountered, likely with the tuning of the VBD function, which was found to have limited performance under certain lighting conditions, particularly when outside with shadows. Lighter patches of sand contrast with darker sand also created conditions where the VBD failed to detect the RSTA. However, there were no false positives.

The gripper also encountered some problems. The gripper BB has been extensively used and the consequent wear and tear of the gripper gearing has reduced its ability to provide a tight grip on the RSTA. The trials have demonstrated how important a tight grip is for a robust pickup. Without a tight grip, movement of the RSTA in the gripper when manoeuvring between workspaces is possible.

Also the design of the gripper fingers can lead to creating a plough effect of sand, as it closes around the RSTA body, triggering the current threshold for the gripper controller early as it experiences resistance. With enough sand, after lifting the RSTA, a void is created which can lead to the RSTA falling from the gripper fingers.

8. CONCLUSION AND FUTURE WORK

At the time of writing this paper, the field trials have been performed and the test results are being evaluated to validate the design against the requirements and system use cases as part of IBB3 [1]. The VBD and AGS systems have been shown to be flexible in picking up RSTAs in a variety of conditions.

The flexibility of the TCs have also proven to be positive for the operation of the FTRS MMS. The modular approach and standardised interfaces meant the ground operator could attempt to recover from non-nominal failures while sticking to a ‘real life’ scenario.

Airbus has successfully breadboarded and demonstrated in the field the technologies that would have been necessary for the SFR to perform the autonomous in depot RSTA pickup and stowage operations needed in support of the MSR fetch mission.

REFERENCES

[1] V. Schaeffer. et al. (2023), Integrated Breadboard 3: rover capability evolution within SFR mission context and future planetary technology testing platform, as a service, *Proceedings of the 16th Advanced Space Technologies in Robotics and Automation (ASTRA)*.

[2] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, W. Woodall (2022), “Robot Operating System 2: Design, architecture, and uses in the wild,” *Science Robotics vol. 7, May 2022*.