Changes and Improvements
Based on Phase 2 Evaluation

ARSI Deliverable D26

March 2018

Contents

1 Introduction

The ECHORD++ Sewer Inspection PDTI follows the timeline shown in figure 1. An initial Micro Air Vehicle (MAV) prototype for sewer inspection was designed in Phase 1 (2016) and implemented in Phase 2 (2017). The prototype was demonstrated during an evaluation in Barcelona in October 2017.

While the ARSI prototype successfully passed the Phase 2 evaluation, the report (1) produced by the external ECHORD evaluators and BCASA identified a number of shortcomings which they consider severely limit the usability of the ARSI system for real-world applications.

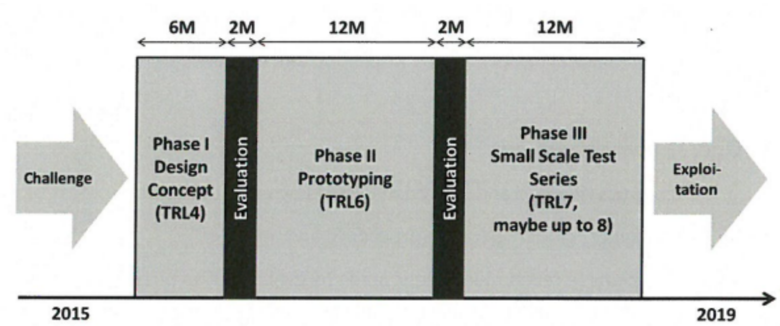

Figure 1: ECHORD PDTI Sewer Inspection time-line

1.1 Evaluation report

While the evaluators were satisfied with the semi-autonomous flight capabilities demonstrated, they considered the flight autonomy insufficient. Despite a nominal flight time of 7-9 minutes, the Phase 2 prototype had an effective autonomy of 5-7 minutes due to the fact that flying with less than 25% battery capacity is not recommended, both for safety and stability reasons.

The evaluators also pointed out that the severe weight limitations on the ARSI platform meant that it could only carry a very limited sensor payload (1 VGA camera), and the resulting data product was deemed inadequate for inspection purposes. In terms of the MAV platform itself, the report noted that no protection against water and dust had yet been implemented to protect the electronics and wiring. Furthermore, the evaluation recommended a redesign of the motor protection implemented in Phase 2, since it actually broke during the evaluation.

Wireless communications were also flagged as an area were improvement is required. The report suggests revisiting the use of repeaters in order to improve coverage and bandwidth, and to allow for more information to be fed back to the operator as inspection flights are being executed.

Finally, the inspection capabilities of the ARSI system were flagged as the area requiring most improvement. The visual data produced was considered only marginally acceptable, and the 3D reconstruction presented during the evaluation was deemed inadequate for structural inspection purposes.

1.2 Structure of the document

In the next section we first present our new MAV platform design for Phase 3, which aims at resolving most of the shortcomings flagged in by BCASA and the ECHORD evaluators. We will also detail our upgraded sensor payload, which will dramatically improve the quality of the inspection data products.

We then describe our current work to upgrade the MAV onboard software, in order to add new functionality such as backwards flight, automated obstacle avoidance, and geo-localization against GIS maps.

Finally, we detail our work on the ARSI offboard software for data visualisation and review, as well as for the generation of new data products, in particular 3D reconstruction and automated defect detection.

2 Platform improvements

2.1 MAV redesign

In the first months of Phase 3 we carried out an entire redesign of the ARSI platform, based on the experience of Phase 2 and the comments from the evaluation.

After evaluating various commercial options, we came to the conclusion that none was fully adapted to operations in the sewers, and we decided to work with a professional drone manufacturer, DroneTools in Seville. Our partnership allowed us to develop an innovative design using larger propellers (13 inches compared to 10 in Phase 2) as well as more powerful motors to increase both flight autonomy and payload capacity. The key innovation was to design a frame were propellers overlap at 50%, in order to maintain the very compact form factor required to navigate in all visitable sewer tunnels (80cm wide or more) as required by the PDTI Challenge Brief (2). The new ARSI MAV has a flight autonomy of over 15 minutes and a payload capacity of 1kg, whilst retaining a very compact form factor (62x81cm). The comparison of the Phase 3 MAV characteristics against the Phase 2 prototype is summarised in table 2.

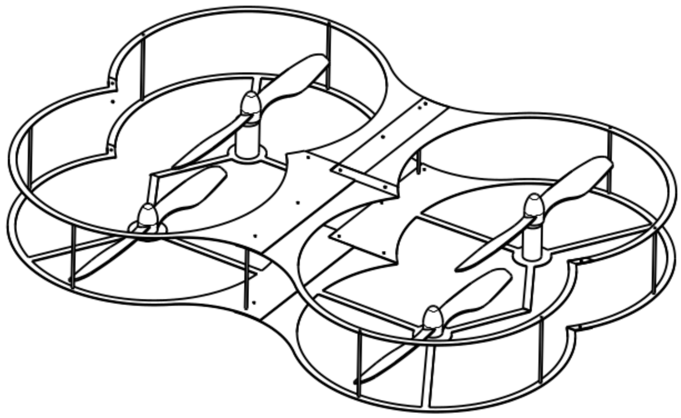

Figure 2: ARSI platform Phase 3 design

The Phase 3 design, shown in figure 2, is built around a lightweight carbon fibre frame which entirely protects the motors. The space at the centre of the platform is divided into 2 levels to mount the onboard PC, the Pixhawk Autopilot and all other internal electronic components (converters, wiring, etc.), so that they are protected from dust and water. We are currently working on mounts for the external sensors (cameras and laser), to maximally protect them from the environment. Finally, the battery will have its own adapted bracket underneath the MAV.

2.2 Sensor payload upgrade

Figure 3: Intel Realsense D435 RGBD camera

Due to critical weight limitations, the Phase 2 prototype used in the October 2017 evaluation could only carry a single Orbbec Astra camera, which produced VGA-quality imagery and weighed around 140gr. At the time, the Orbbec was the best lightweight RGBD camera available on the market. Since then, Intel released the Realsense D400 series RGBD cameras (figure 3) which showcase HD imagery, improved range and field of view, for almost half the weight (see datasheet). Table

1 clearly shows that the specifications of the Realsense D435 are vastly superior to those of the Orbbec Astra.

	Orbbec Astra	Intel Realsense D435
Weight (gr)	140	72
Range (m)	0.6-8.0	0.15-10.0
Visual image size (px)	640 x 480	1920 x 1080
Depth image size (px)	640 x 480	1280 x 720
Field of view (degrees)	60 horiz x 49.5 vert	85.2 horiz x 58 vert

Table 1: Orbbec Astra (Phase 2) vs Intel Realsense (Phase 3)

The Phase 3 ARSI platform will carry 2 Realsense D435 RGBD cameras as well as 2 Basler daA1600-60uc visual HD cameras with fish-eye lenses. The Realsense cameras are mounted at the front and back of the ARSI MAV, one looking slightly towards the ceiling and the other towards the ground, in order to provide full visual coverage of the tunnels. These cameras are extremely versatile and will be used for visual and structural inspection, for odometry and navigation, and for obstacle avoidance. The 2 Basler cameras, with a resolution of 1600x1200 and weighing no more than 40 grams each, are mounted on each side of the platform for close-range inspection of defects on the sewer walls.

Figure 4: RPLIDAR A2 laser

Another comment from the Phase 2 evaluation was that our prototype should be able to move backwards. This would allow handling situations where an obstacle is encountered and the MAV should return to the nearest entry point (manhole) to be recovered. As noted in the Phase 2 evaluation, the ability to operate backwards would also simplify mission execution, by reducing the number of manholes that need to be opened to recover the MAV or change its batteries.

To this end, we replaced the 2D laser used in Phase 2 (Hokuyo UST-10LX) with a RPLIDAR A2 (see figure 4), which has a field of view of 360 degrees. Along with the 2 Realsense, the RPLIDAR renders our MAV fully symmetrical so that it can seamlessly navigate forwards and backwards in the sewer tunnels.

Figure 5: TeraRanger Evo infrared ranger

Finally, the infrared altitude sensor mounted underneath the MAV will be upgraded to the new TeraRanger Evo (figure 5). While similar to the model used in Phase 2, the new Evo features improved close-range measurements (from 0.1m instead of 0.2m for the TeraRanger One) and simplified electronics whilst remaining extremely lightweight (12g).

3 Communications

In Phase 2, we used a single WiFi relay between the MAV and the operators on the surface. In comparison, the SIAR team deployed a network of WiFi repeaters at key points in the evaluation area, to achieve full coverage of the area from their static point of operations.

Our decision not to use repeaters in Phase 2 was motivated by the observation that MAVs have a short battery life, in our case around 7 minutes, while a ground robot like the one used by SIAR can run for up to 5 hours on a single charge (3). The MAV batteries must therefore be replaced after each flight, which already complicates logistics: operators must open a manhole, deploy a safety tripod, go into sewer to access the MAV, switch batteries, then come back up, close the manhole and remove the tripod. Since this process was unavoidable due to the battery life, we decided to move our WiFi relay as batteries were replaced, to provide coverage for the next flight without having to deploy repeaters. Our goal was to optimise an already quite complex operation by limiting the number of manholes that needed to be opened. However, the limit of this approach lies at corners and intersections in the sewers, where WiFi communication is lost after only a few meters due to the physics of signal propagation in tunnels (4).

Our strategy for MAV communications in Phase 3 is twofold. First, we will rely on the vastly improved flight autonomy (15+ minutes vs 7 in Phase 2) to simplify operations. Inspections will be planned to take advantage of the new ability to reverse flight direction, so that the MAV can be recovered at the point of deployment.

Second, we will replicate the repeater strategy used by SIAR in Phase 2, deploying devices at key locations to generate maximum WiFi coverage of the area to inspect. We will evaluate the logistics of such approach, in particular in parts of Barcelona where opening manholes to deploy repeaters is disruptive to road traffic.

Once these two strategies have been evaluated, we will bring the technical and operational experience of our consortium together to design the best possible operational procedures for ARSI, taking into account both the advantages and limitations of a MAV system.

4 Onboard software

4.1 Backwards flight

As previously discussed in section 2.2, our Phase 3 MAV will carry 2 Realsense cameras, one looking forwards and the other backwards, as well as a RPLIDAR laser with 360 degrees of field of view. With these sensor upgrades the new platform becomes fully symmetrical, and thus has the ability to fly forwards as well as backwards.

Significant work is still required to integrate this functionality into our onboard software, in particular when it comes to local path planning and obstacle avoidance. Once implemented however, operators will be able to plan missions where the MAV inspects an entire sewer tunnel from a single point of entry, significantly simplifying the logistics of inspections with ARSI.

4.2 Geo-localisation

In the context of sewer inspection, geo-localisation aims at providing ARSI operators with a real-time estimate of the MAV location on a map or GIS (Geographic Information System) as inspections are being executed. The geo-localisation solution is then refined in post-processing, to reference information such as video and structural defects against maps and GIS, and facilitate the task of the experts responsible for reviewing the inspection data.

In Phase 2, the MAV location was estimated by fusing information from the inertial sensors (accelerometers and gyroscopes), altitude measurements from the TeraRanger infrared sensor, and a Visual Odometry (VO) solution calculated from visual and depth data produced by the frontal RGBD camera.

However, the evaluators noted that this estimated MAV location quickly diverged over time, reaching errors as large as 20% of the distance travelled. This error is due to the accumulation of small sensor measurements errors over time, in particular from the Visual Odometry since the algorithm relies on the detection of standout features in the imagery. The quality and reliability of these features is affected by many factors such as the resolution of the images, the illumination of the scene, dust in the air caused by the airflow of the MAV motors, image blur caused by rapid movements, and insufficient identifiable features present on the sewer walls.

In Phase 3 we are working on reducing these various errors by using higher-end sensors (see section 2.2) and by making our odometry solution more robust. However measurement error is inevitable regardless of these improvements, and a localisation solution based on local sensors alone will always exhibit some degree of drift over time. In order to bound this drift, we need to use external features such as structural elements of the sewers whose position is known from the GIS in an absolute frame of reference (e.g. GPS or UTM coordinates). Figure 6 shows how structural features such as manholes[1] can be detected from RGBD data.

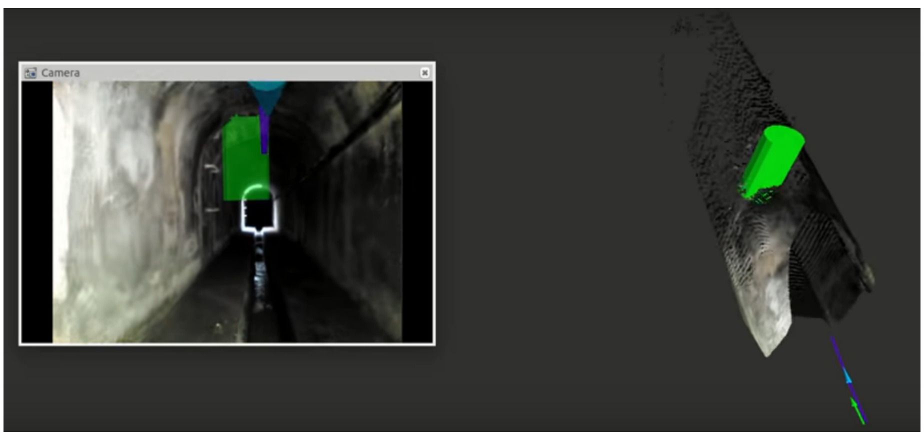

Figure 6: Automated manhole detection from RGBD data

Our approach for localisation is similar to that used by the SIAR team (5). We use a statistical framework called particle filter to represent the set of hypotheses for the real location for the MAV. These "particles" are continuously evaluated and assigned a weight based on how well they match real sewer features observed in the data. The MAV location is then estimated as a weighted sum of all particles in the sample set.

Early tests using Phase 2 data have shown promising results. Figures 7 to 10 illustrate the process of probabilistic localisation with our particle filter implementation, using data from Passatge Mercantil in Mercado del Born, Barcelona. A video[2] describing the localisation is also available.

[1]https://youtu.be/uT0S3UNuapM
[2]https://youtu.be/Yv5_HuTJIVE

The particles (in red) are overlayed on the GIS (sewers in green, manholes in grey). The spread of particles models the uncertainty of the MAV localisation solution. Figure 7 shows how the particles spread out as the uncertainty increases when navigating in a featureless tunnel.

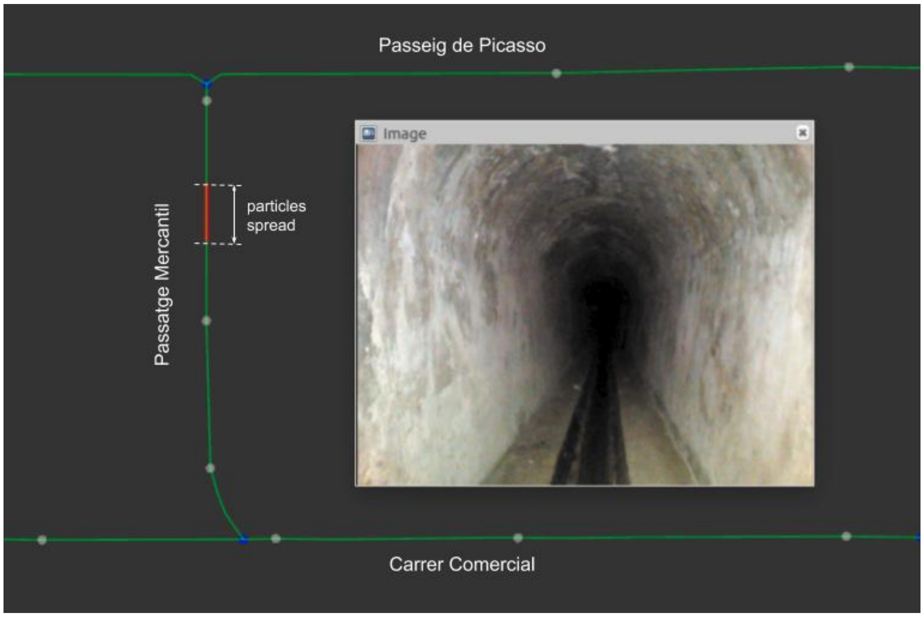

Figure 7: As the MAV progresses in a straight, featureless tunnel
the particles (in red) spread out and the localisation error increases

In figure 8, a manhole is detected in RGBD data. The detection is converted into the frame of reference of each particle, and matched against the GIS. Particles are resampled based on the detection error, so that only particles that are consistent with the GIS are retained. We can see that the spread of particles is greatly reduced, and our estimate of the MAV position is much more accurate.

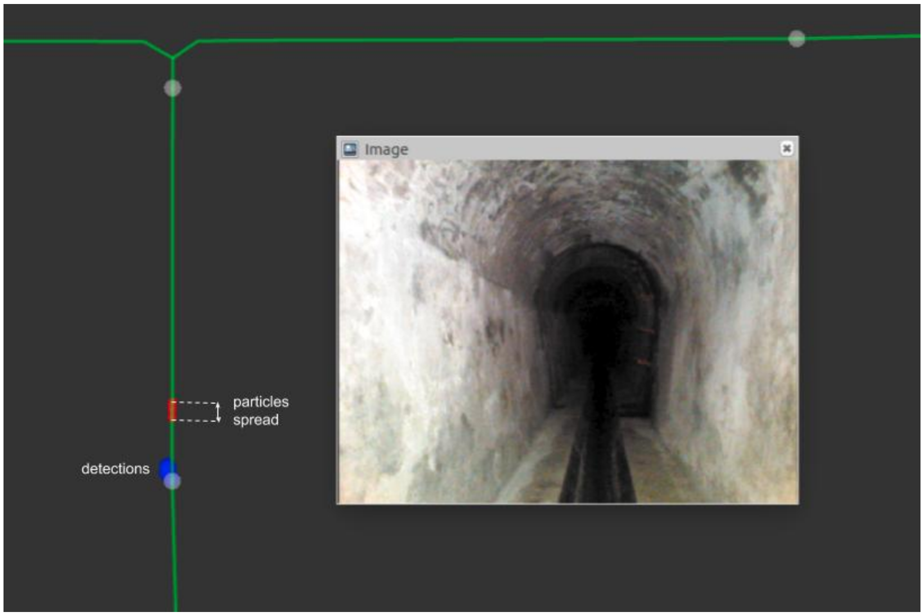

Figure 8: A manhole is detected. Detections (in blue) are matched against the GIS and the particles (in red) are resampled around the real MAV location

This process is repeated in figures 9 and 10: the particles start spreading again as the MAV progresses in the next sewer tunnel, until another manhole is detected and the particle distribution converges around the real MAV location.

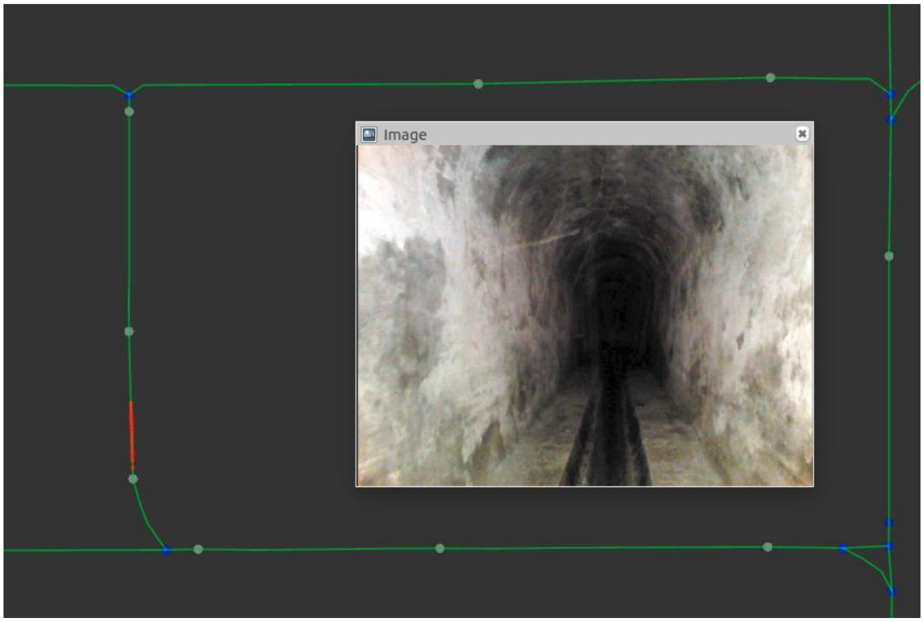

Figure 9: The particles spread out again as the MAV
navigates the next featureless sewer section

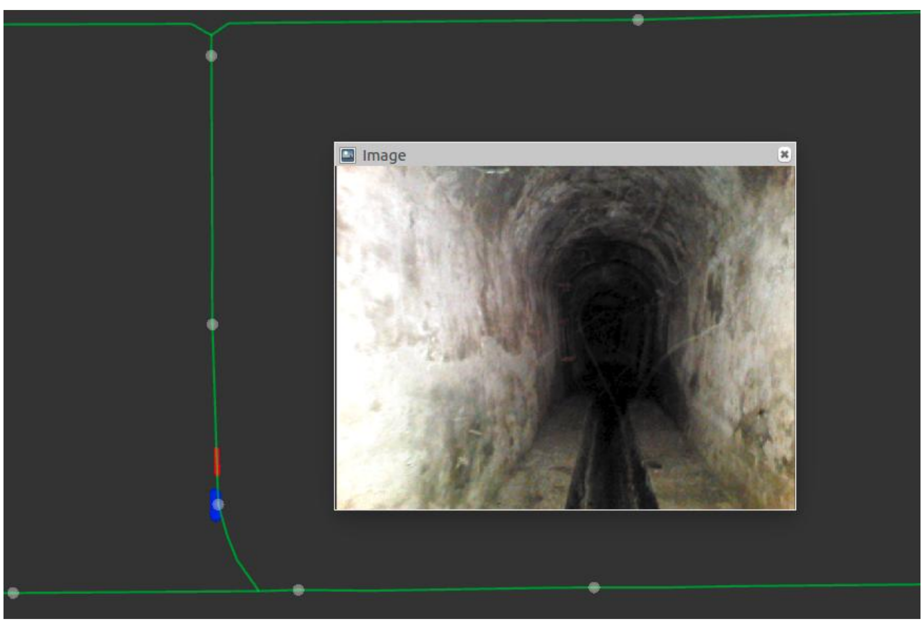

Figure 10: Another manhole is detected (in blue)
and the particles converge on the real MAV location

We are working on improving our localisation by detecting other sewer features referenced in the GIS, in particular intersections and changes of sections. This work will be helped by the upgrades to our sensor payload, in particular the improved resolution and field of view of the Realsense RGBD cameras (see section 2.2).

4.3 Obstacle avoidance

The evaluation area at Mercado del Born in Phase 2 featured several obstacles, in the shape of service pipes crossing the sewer tunnels. While our MAV was compact enough to fly underneath these obstacles, avoidance was carried out manually by the operator using the control interface to switch between standard and low flight altitudes.

Figure 11: ARSI MAV flying at low altitude
to avoid a service pipe in Mercado del Born

In Phase 3 we plan on implementing an automated 3D obstacle avoidance strategy using potential fields, as described in (6). Potential fields simulate the repulsive forces (shown in red in figure 12) generated by all obstacles in the field, be they known sewer walls or unforeseen obstacles. Likewise our target goal (specified in the inspection plan) generates an artificial attractive force, and the optimal MAV trajectory is given by the weighted sum of all attractive and repulsive forces.

As proposed in (6) we model our MAV platform as a discrete volume, in order to simulate how the potential forces affect the different parts of the structure. Following this framework, we calculate local paths that automatically deviate from obstacles, in order to keep the MAV safe at all times whilst following the overall mission path defined in advance.

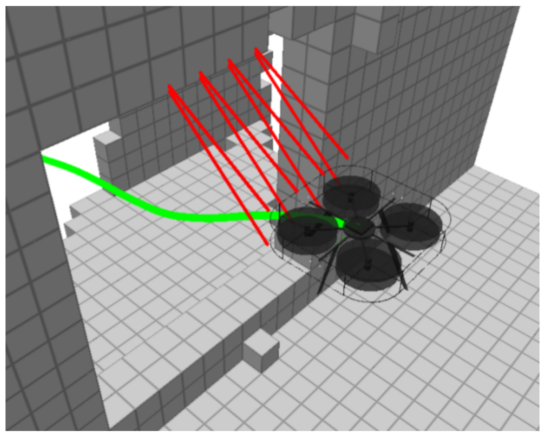

Figure 12: Obstacle avoidance and path planning using potential fields

5 Offboard software

This section describes the software that will be run after the missions (offline). This software will produce two main outcomes: a 3D map and an assessment on the serviceability and structural status of the stretch covered. Our visualisation app then allow operators to perform a detailed and efficient review of all data.

5.1 3D reconstruction

As mentioned in section 1.1 the 3D models generated by ARSI software in Phase 2 did not provide the precision and detail expected to perform inspection tasks by sewer operators. The VGA resolution of the Orbbec camera (see table 1) was the main obstacle for obtaining refined 3D reconstructions but there were other factors, such as the size limit established by the visualiser app for the 3D model, which effectively bounded the number of 3D points that generated the 3D models. In this phase, we are working to optimise every step of the 3D pipeline in order to attain high quality 3D reconstructions of the sewers.

3D model size limit on visualisation

For visualisation purposes in Phase 2 the offboard ARSI software created a textured surface mesh of the model out of the estimated 3D point cloud. However, the rendering tool used for visualising 3D models limited the number of faces of the mesh to 2 million, which forced us to bound the size of the available meshes for each mission. This ultimately affected the resolution of the 3D models displayed on the visualisation app.

In order to overcome this shortcoming we have developed a tool for splitting the 3D models into several 3D submodels, in such a way that the visualiser can work seamlessly with them. This intermediate process removes the size limits of the 3D models and allows us to focus on generating more detailed 3D structures. Figure 13 shows how this technique already allows us to render more precise 3D models.

Figure 13: Comparison between the rendering in Phase 2 (left) and Phase 3 (right), with results from the VGA Orbbec sensor

Further improvements on 3D Reconstruction

Various steps can be optimised throughout the 3D reconstruction workflow there. Crucially, the refinement of the estimated odometry can be further improved and the texture is linked to the geometry of the mesh. In this phase the texture will be built independently from the mesh, which will increase the colour resolution of the 3D reconstructions. Most importantly, the pipeline should accommodate the image quality and depth information provided by the sensor Intel Realsense D435, which will be used in Phase 3, as noted in section 2.2.

5.2 Defect detection

The main functionality expected from the offboard software in Phase 3 is defect inspection. This is a difficult task by itself, and it is even more challenging when we consider that the data is collected from an ever moving platform in a mostly textureless, badly lit environment, which in addition has been little explored by inspection techniques or research on this area.

In order to optimise the research work on this package, the ARSI consortium has proposed to the evaluators (BCASA and UPC) a list of goals to achieve in this phase. These goals are split in two areas: serviceability and defect inspection.

Serviceability

- Identification of section type of the sewer. From the information provided by the GIS and the odometry, the ARSI software will be able to match the section profile of the traversed sewer.
- Serviceability evaluation. The ARSI software will determine whether the sewer section is clean or there is any obstacle that hinders the flow of the water and where.
- Water level on sewer bucket. The ARSI software should estimate the water level on the bucket.
- Sewer assets localisation. The ARSI software will notify where different assets are located: manholes, sewer drains and bifurcations.

Defect inspection

- Structural defect identification. Using volumetric analysis of the 3D map, the ARSI software will detect when an anomaly occurs in the sewer section and will throw an incidence tag. The operator will decide whether the anomaly found is actually a defect and will establish the type of defect.
- Texture defect identification The ARSI team will develop a specific method for the classification of a typical texture defect present in the sewer, by means of deep learning techniques.

In order to detect structural defects and check the identified section type, the ARSI software will fit the section profile provided by the GIS to the reconstructed model. This fitting will be performed over the course of the sewer and will allow the detection of divergences due to a wrong type and/or location of a section according the GIS as well as the detection of any deviation from the original shape in the form of defects or obstacles.

Our first tests using a 2D shape matching approach, show that it is possible to correctly fit the section profile to different parts of a reconstructed model, which will serve as a basis for detecting any deviation from it, as shown in figure 14.

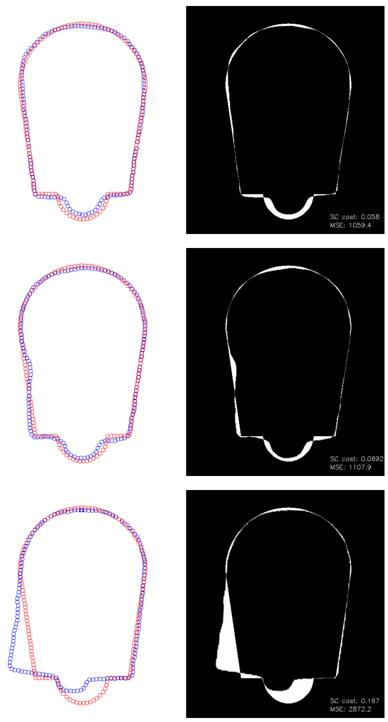

Figure 14: Fitting of sections sampled from a reconstructed model (in blue) to a section obtained from the GIS (in red). The left column shows the obtained fit; the right column shows the detected differences. The last row illustrates a wrong fit due to a sewer having a different profile.

Some defects might not be detected by means of the described volumetric approach. In particular, they may appear as irregular textures over the wall of the sewer without enough deviation from its expected profile. The ARSI team will apply the latest techniques in machine learning for image categorisation in order to detect this kind of defects. In particular, deep neural networks will be trained using a large set of representative defect samples and deployed so they can produce a heat map encoding a probability for each recorded pixel to be part of a defect.

The required steps to reach the expected goal will start by recording data at the sewer with the drone platform. After that, the team will label the images in order for the neural network to know the appearance of the defects it will be trained to recognise. For all recorded images, a manual labelling process will be produced as it is shown in figure 15. Then, training algorithms for neural networks will be used to fit the best possible model given the data and validated in an independent test set of defect images.

Figure 15: Left: sample of a textural defect that will be used to train a deep neural network for detection of its own kind. Right: manual labelling of the regions where the defect is present in the sewer's walls.

5.3 User interfaces

In Phase 3 the data analysis user interface will undergo a work of development in order to meet the standards of a user-friendly, off-the-shelf product. Specifically, our main tasks at this stage are the refinement of Phase 2 functionality and the development of new features:

- Navigation Interface
  - providing new icons and visual assets for interaction.
- Offline Mission Inspector
  - Integration of GIS data (drone positioning, cartography database, etc)
  - More tools for categorisation and manual tag processing/edition of incidences.
  - Improvements to the real time engine and specifically to the 3D datasets management.

Other tasks relate to the user experience of the Offline Mission Inspector, and the reformulation of the viewports and the information to be considered in when reviewing datasets, therefore this is currently an open task with continuous improvements during Phase 3.

6 Conclusions

Table 2 summarises the key improvements being made to the ARSI MAV and offline software since Phase 2. While most of this work is still in progress, and flight stability as well as autonomy in particular are still to be demonstrated, we believe that our new MAV platform and offline inspection functionality will be a leap forward in making ARSI a valuable tool for sewer inspection professionals.

	Phase 2	Phase 3
Flight autonomy (mins)	5-7	15+
Payload capacity (gr)	700	1000
Dimensions (cm)	62 x 72	62 x 81
HD visual cameras	0 (1 VGA)	4
HD depth cameras	0 (1 VGA)	2
Lateral cameras	No	Yes
Dust and splash resistant	No	Yes
Integral protection	No	Yes
Backwards flight	No	Yes
Geo-localisation	No	Yes
HD 3D reconstruction	No	Yes
Defect detection	No	Yes

Table 2: Summary of ARSI system improvements

References

[1] ECHORD, *ARSI Phase 2 Evaluation Report*, 2017.

[2] ECHORD, *PDTI Urban Robotics - Sewer Inspection Challenge Brief*, 2015.

[3] D. Alejo, C. Marques, F. Caballero, P. Alvito, and L. Merino, "SIAR: An Autonomous Ground Robot for Sewer Inspection," 2016.

[4] C. Rizzo, P. Cavestany, F. Chataigner, M. Soler, G. Moreno, D. Serrano, F. Lera, and J. L. Villarroel, "Wireless Propagation Characterization of Underground Sewers Towards Autonomous Inspections with Drones," 2017.

[5] D. Alejo, F. Caballero, and L. Merino, "RGBD-based Robot Localization in Sewer Networks," 2017.

[6] M. Nieuwenhuisen, D. Droeschel, D. Holz, and S. Behnke, "Omnidirectional Obstacle Perception and Collision Avoidance for Micro Aerial Vehicles," 2013.