

Phase III Tests and tests results

ARSI Deliverable D26.10

PDTI Urban Robotics - Sewer Inspection

July 27, 2018

Contents

1 Introduction	3
1.1 Structure of the document	4
2 MAV Platform	5
2.1 Platform redesign	5
2.2 Onboard software	6
2.3 Sensor payload	7
2.4 Tests and test results	9
3 Data processing	11
3.1 3D reconstruction	12
3.2 Structural assessment	13
3.3 Distance map inspection	14
3.4 Obstacles detection	15
3.5 Tests and test results	16
4 User interfaces	21
4.1 Mission Planner	21
4.2 Operator Console	24
4.3 Data Analysis Interface	25
5 Conclusions and future work	26

1 Introduction

The ECHORD++ Sewer Inspection PDTI follows the timeline shown in figure 1. A Micro Air Vehicle (MAV) prototype for sewer inspection was designed in Phase I of the project (2016) and implemented in Phase II (2017). The prototype was demonstrated to ECHORD evaluators in Barcelona in October 2017.

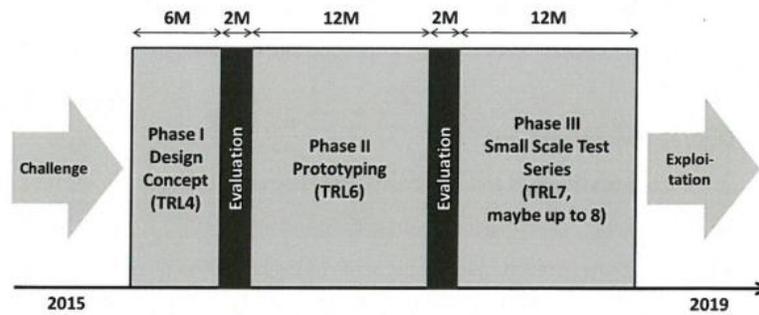


Figure 1: ECHORD PDTI Sewer Inspection time-line

Following the evaluation, the ECHORD++ evaluators issued a series of recommendations and required changes to the ARSI MAV and software suite for phase III of the project. Over the last 6 months, the ARSI consortium carried out these changes, and the updated inspection system was demonstrated during a new field test on July 3rd 2018 in the Virrei Amat area of Barcelona.

In this document we present the tests and test results for the work carried out so far during phase III of the project.



Figure 2: ARSI MAV executing a sewer inspection in Barcelona

1.1 Structure of the document

One of the main comments from the phase II evaluation was that our MAV prototype did not exhibit satisfactory payload capacity and flight autonomy for sewer inspection. Accordingly, we decided to redesign our platform to improve performance, with the help of professional drone manufacturers. Section 2 describes the work carried out in this regard.

Another key objective for phase III was to refine the 3D models of the sewers that are generated from inspection data collected by the ARSI MAV, and to use these models to carry out automated structural assessments of the sewer sections. This work is described in section 3, along with initial results of structural assessments and defect detection.

While some work had already been implemented during phase II, the 2017 evaluation highlighted that a significant effort was required to make the ARSI system more intuitive and easy to use by inspection brigades on the ground. In section 4 we present the user interfaces developed during phase III to facilitate the entire ARSI workflow, from mission planning to mission execution, data analysis, and review.

2 MAV Platform

In this section we present the changes made during this phase to the ARSI MAV platform, sensors and onboard software.

2.1 Platform redesign

The MAV quadrotor developed in phase II is shown in figure 3. It used 4 motors with 11in (28cm) propellers, powered by a single 6000mAh LiPo battery. It had a nominal flight time of 9 minutes, but field tests showed that the real flight time was closer to 5-6 minutes, after which the low battery power resulted in poor control and unstable flights. The phase II MAV had a payload limit of $\sim 700\text{g}$.



Figure 3: Phase II MAV prototype

In phase III we partnered with DroneTools, a drone manufacturing company based in Sevilla, Spain. Given our requirements, they proposed an innovative design with overlapping propellers (see figure 4), allowing us to use longer blades (14in or 36cm) whilst keeping the MAV narrow enough (61cm) to navigate in the sewers. These larger propellers, coupled with more powerful motors, deliver significantly more thrust, allowing for a larger payload capacity (1kg) and a longer flight autonomy (15 minutes).

The new platform also includes a fully rigid protection for the drone and the motors, and an enclosed space at the center where the onboard PC, the autopilot and other components can be mounted and protected from sewage water or dust.



Figure 4: Phase III MAV platform design

2.2 Onboard software

The onboard software for the ARSI MAV remained largely unchanged since phase II, as it had proved satisfactory in the 2017 evaluation. Some notable features were nonetheless implemented:

- Backward flight was added to the path planner, so that the MAV can now be commanded to return to its starting point after an inspection.
- Automated flight checks are performed onboard before each flight, to ensure that all sensors and software components are in a valid state before an inspection is initiated. Operators are notified of any issues via the Operator Console (see section 4.2).
- Missions are now created using the ARSI Mission Planner (see section 4.1) and sent to the MAV for execution from the Operator Console. The onboard software was updated to handle this new mission workflow.

2.3 Sensor payload

RGBD-cameras

The main change to the sensor payload in phase III was the integration of the new Intel Realsense D435 RGBD (color and depth) cameras, to replace the Orbbec Astra used in phase II (see figure 5). The Realsense D435 cameras are much lighter than the Orbbec (70g vs 140g) and produce HD imagery, as required in the Challenge Brief, up to a resolution of 2MP (1920×1080).



Figure 5: Intel Realsense D435 (left) and Orbbec Astra (right) RGBD cameras

However our laboratory tests showed that the RGBD data produced by the Realsense was of poor quality compared to that of the Orbbec. Figure 6 shows RGBD point-clouds for the same object (a wall of cardboard boxes) generated by the Realsense (in the back) and the Orbbec Astra (at the front) in a well-lit environment and at a range of around 1m. We can clearly see that while the front point-cloud accurately represents the geometry of the scene, the other is heavily distorted. The Realsense data was even poorer at larger ranges or shallower incidence angles.

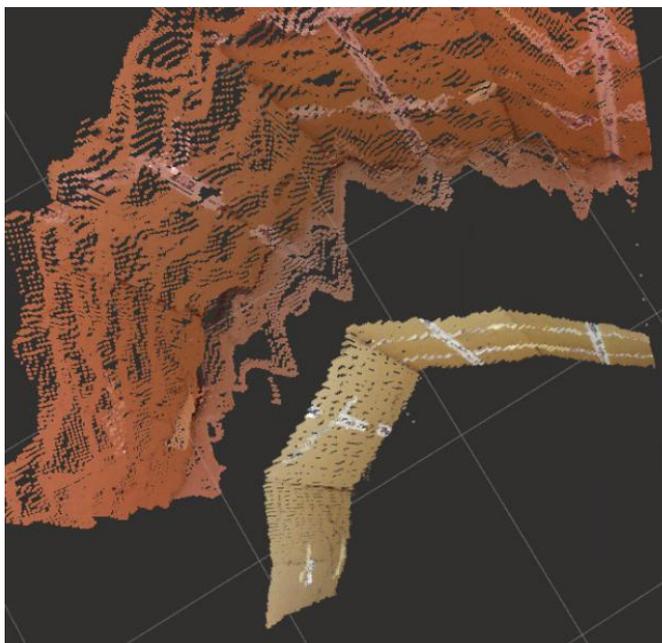


Figure 6: Comparison of point-clouds generated using the Realsense D435 (back) and the Orbbec Astra (front)

The ARSI system relies heavily on this RGBD data for visual odometry, 3D reconstruction and structural analysis. RGBD visual odometry in particular is used to estimate the MAV trajectory and real-time velocity, which are used in the low-level control loops of the autopilot. The consequences of introducing poor RGBD data into the system are therefore quite serious. Another drawback of the Realsense is that it requires post-processing of its optical and infrared stereo data to generate RGBD information, while the Orbbec performs these calculations inside the camera. During our tests we could see that the Realsense drivers were CPU-heavy, slowing down other onboard modules such as the visual odometry.

Given these shortcomings, we decided to postpone the use of the Realsense camera, pending investigation into its limitations. The July 3rd evaluation was therefore carried out using an Orbbec Astra, generating 0.3MP VGA imagery (640×480).

2D laser

In phase III we also replaced the Hokuyo UST-10LX 2D laser used in phase II with the RPLIDAR A2 (see figure 7). While the RPLIDAR is

heavier than the Hokuyo (190g vs 130g) it provides a 360 degrees field-of-view required to implement backwards flight. The RPLIDAR proved very reliable during our numerous tests in the sewers, producing accurate laser data measurements at 10-12Hz, and operating flawlessly despite the harsh sewer environment and vibrations from the motors.



Figure 7: RPLIDAR A2 laser with 360 degrees field-of-view

2.4 Tests and test results

Laboratory tests

Our first task upon reception of the new MAV prototype was to integrate all the sensors and electronic components including the cameras, laser, onboard PC, power regulators, and LEDs. We then performed several autonomy tests in an open laboratory environment (without turbulences), where we achieved flight times of 15-16 minutes (in near-hover) using 6000mAh and 7000mAh LiPo batteries.

We then tuned the gains of the various PID control loops of the Pixhawk autopilot (position, velocity and attitude loops). The objective was to achieve stable control and execute simple trajectories in our "sewer" environment, where we simulated narrow tunnels, turns and intersections using cardboard boxes, as we had done in phase II. Our goal was not to fine-tune the controllers, since we knew from previous phases that MAV control would be very different in real sewer environments, due to the turbulences generated in such confined spaces by the airflow from the motors.

Early tests in real sewers

After the laboratory tests, we planned a series of visits to real sewers in Barcelona, first at Mercado del Born, where we had worked in phase

II, then at the two designated sites for the July evaluation (Plaça Virrei Amat and Plaça Sol de Baix).

Early in our flight tests we observed that the new MAV was less stable than the prototype used in phase II. While this issue still needs to be investigated, our analysis is that it is due in part to the form factor of the new MAV, and to the increased power of the motors:

- The new MAV is 82cm long and 61cm wide, with overlapping propellers so that the lateral distance between motors is relatively short: The torque for this axis is therefore weaker, meaning that roll control is likely to be less stable.
- The larger propellers and more powerful motors generate significantly more airflow than the phase II platform. In narrow, confined sewer environments this increased airflow results in stronger turbulences, and less stable flight control.

While these issues had been foreseen, increasing the MAV thrust was the only way to improve payload capacity and flight autonomy, as required from the phase II evaluation. Despite these additional challenges, we were able to perform relatively stable flights in sewer tunnels 120cm wide, both in Mercado del Born and Virrei Amat.

Tests in narrow sewer tunnels

Operations in narrower tunnels (100cm down to 80cm wide, as required by the Challenge Brief) proved more problematic, as additional factors come into play:

- Our entire navigation and control system relies on a visual odometry algorithm to estimate the trajectory and real-time velocity of the MAV from RGBD data. While the visual odometry had performed well in all tunnels at Mercado del Born, this was not always the case in the new areas (see below)
- The less stable control and the very narrow margins (see figure 8) meant that the MAV was more likely to make contact with the walls, corners or manhole ladders, causing more instability and even occasional crashes.
- The visual odometry approach introduces a feedback loop, where unstable control caused by poor odometry leads to rapid movements and blurry images, which themselves cause poor feature detection and tracking.



Figure 8: MAV before takeoff in a T111 section (see figure 20) at Plaça Virrei Amat. The section is 60cm wide at ground level (same as the MAV) and ~ 80 cm wide at flight altitude

Both in Plaça Sol de Baix and in Virrei Amat, several sewer sections exhibited smoother walls and cleaner surfaces than we had seen in Mercado del Born, offering very few visual features for the visual odometry algorithm to detect and track. This meant that the real-time velocity estimate was generally quite poor, and the MAV would struggle to control its speed. We frequently saw it lose odometry and fly at very high speed (as observed in the onboard camera), becoming very difficult to control.

Flight autonomy

Our tests in narrow sewers showed that the flight autonomy in real operations was significantly less than the 15 minutes benchmark obtained in a laboratory environment. We were consistently able to fly ~ 8 minutes; beyond that time the reduced battery power would result in very unstable flights.

3 Data processing

In this section we present the algorithms developed in phase III to post-process visual and depth data collected by the MAV sensors and gener-

ate 3D models of the sewers, as well as automated structural inspection and defect detection.

3.1 3D reconstruction

Using the data acquired during the flight, we proceeded to reconstruct the 3D model of the sewer for its use in the data analysis interface as well as for the detection of volumetric changes during the structural assessment, as required in phase III. The reconstruction pipeline used in phase II had several shortcomings, since it required manual intervention during the reconstruction process. Texturing also presented some flaws in terms of aliasing and storage needs that were solved at this stage.

In phase III we automated the reconstruction process so that manual intervention is mostly unnecessary. This improves the reconstruction time and makes this step more transparent to the user. Concerning the texturing step, a new parametrization has been introduced that allows a continuous packing of color along the texture. This new parametrization is based on a cylindrical mapping that follows the path of the inspected sewer, allowing a much more natural parametrization (see figure 9). The obtained texture is more suitable for the analysis interface since it requires less storage and avoids aliasing artifacts present with the old texturing approach. In addition, it simplifies the structural assessment since elements and obstacles can be easily analyzed on 2D texture space, as described in section 3.3. Filling of this texture is performed taking advantage of the graphics hardware, which also makes its computation faster than before.



Figure 9: Comparison between old texture parameterization (left) and new one (right), the latter based on a continuous mapping along the sewer.

In terms of precision, the reconstructed geometry is still similar to that obtained in phase II. Due to the shortcomings present with the new RGBD sensor that was planned as a replacement for the Orbbec Astra (see section 2.3), we had to rely on the same type of data as in the

previous phase. This implied working with the same resolutions for both RGB and depth images, which are used as input for the reconstruction process. In subsequent stages, we plan to improve on this by either solving the shortcomings of the RGBD sensor or mounting a better RGB camera on the MAV platform.

3.2 Structural assessment

In order to compute the structural assessment of the inspected sewer, we need to compare the reconstructed model with the expected (theoretical) structure of the sewer. Such structure is given by the information available in the GIS database, which includes information about the type and geometry of the section profiles, and the presence of other elements such as manholes or connections with other tunnels.

By comparing the current model with the expected geometry of the sections we can detect all kinds of elements present in the current sewer, including obstacles, manholes, connections, or even structural defects. The list of sections for the inspected zone is retrieved from the GIS by the mission planner, which passes this information to the MAV platform and is subsequently retrieved during the data analysis step.

Given the geometry of a specific section (see for instance figure 20), we need to register its shape along the reconstructed sewer to detect any variations from this. At the same time, we need to detect locations where the section might change to the subsequent one, in order to always find differences with the appropriate section and check whether these locations coincide with the GIS information.

The registration process is performed at different steps along the model, computing the best transformation that aligns the reconstructed section with the theoretical one. During this process we check whether this alignment is good and hence the section is correctly identified, or whether there might be any change of section or just structural elements that make this matching differ. In the latter, bad registrations are simply ignored and their neighboring registrations are interpolated in-between.

Once the reconstructed sections have been properly aligned with the theoretical sections, we proceed to compute the distance of each point in the model with the corresponding section. This gives us a (signed) distance map that encompasses any variation from the theoretical profile along the model, which is stored as an additional (floating-point) texture. Such a texture is subsequently evaluated to detect which types of variations are present in the model as well as their location and shape, as explained in the next subsection.

Figure 10 shows an overview of these steps.

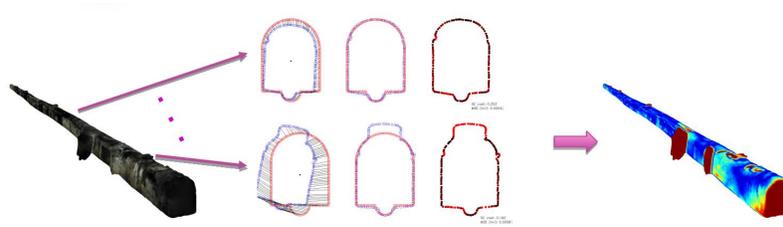


Figure 10: Structural assessment starts by registering the sewer’s theoretical profiles with the reconstructed model (left & middle). A distance map is then computed to find variations along the course of the sewer (right), here shown as a temperature map.

3.3 Distance map inspection



Figure 11: A texture image specifying the distance of each reconstructed point to the theoretical profile of the sewer.

Given the texture image specifying the positive and negative distances to the theoretical profile of the sewer (see, for instance, figure 11), detection of structural elements in the sewer is performed through image processing techniques. The texture image is an “unfolded” version of the sewer tunnel reconstruction, as if it were “cut open” along the ceiling. In this image, we detect the following elements:

- **Reconstruction errors:** first, we look for bad reconstructed regions of the sewer to avoid misclassifying them. Due to a bad collection of features, errors in the reconstruction — and, thus, in the computation of distances to the theoretical profile — might appear. It is important to detect such parts as their distance values live outside of the regular range. In practice, we look for outliers in the set of distance values and then for large compact regions of such values.
- **Manholes:** in the texture image (the 2D projection of the “unfolded” 3D reconstruction), manholes are located halfway in between the top and the bottom of the images. To locate manholes, the image is re-shaped as to have the manholes in the middle of the image, that is, as if the 3D reconstruction was mapped onto

the 2D image cutting by the *cubeta*. On this new image, thresholding techniques are used to separate manhole candidates whose depth with respect to the profile is larger than a certain value (as manholes are, in fact, holes). Then different heuristics are applied to distinguish them: in particular, size, shape and position are evaluated to decide whether a candidate is in fact a manhole.

The rest of the structural elements detected on the images are located using a similar methodology. First, the *cubeta* is located in the image so it can be deduced whether an element is close to it. This helps in the classification. Secondly, a thresholding technique is applied to locate candidates and, finally, different heuristics are used to separate among the different types of elements, in particular:

- **Junctions** (*entronques*): these are large holes in the image. The size of the hole is measured with respect to the size of the sewer, since a junction is actually another sewer connecting to the main one being analysed.
- **Gutters** (*imbornal*): a connection of this type has a particular shape, typically elongated (higher than wider). Filtering by this parameters allows locating them.
- **Connections** (*acometida*): these are smaller holes — less elongated than gutters — without the need to be close to the *cubeta*.

In all cases, whenever an element has been detected, that region is not analyzed again. That is, we do not report overlapping elements to avoid over-informing the user. We also report *profile changes* where the actual profile of the sewer has changed as detected in the registration step. Elements attached to a profile change are not reported since they might produce several false positives as the reconstruction close to a profile change is typically noisy.

3.4 Obstacles detection

To evaluate the serviceability of a particular section we look for obstacles. Obstacles are treated in two separated ways, although one prevails over the other. In the first place, we look for lacks of capacity of the sewer and, only when the sewer has a certain amount of capacity, we look for large portions of the sewer where the texture image has a piece of sediment.

- **Capacity lack**: in order to detect a section of the sewer where it has diminished its hydraulic capacity we perform a column-based

analysis of the texture image. At each reconstructed profile, we integrate the negative differences (negative values mean protrusions with respect to the theoretical shape) and compare this with the actual area determined by the theoretical capacity of the section. A ratio between the computed capacity and its theoretical one determines the % of the sewer that is at service. If this situation occurs for a certain amount of profiles, this is, along a certain distance controlled by a parameter, we consider this situation a capacity lack of the sewer.

- **Sediments:** to look for sediments, we proceed in a similar way as for the structural elements. We threshold the image (here we threshold the negative values) to have a set of candidates. The appearing blobs are selected based on their dimensions since small particles do not really constitute an obstacle and very big ones are most probably reconstruction errors, which are detected in a previous stage.

Again here, overlapping defects are only reported once to avoid producing too many alerts in very near places of the path.

3.5 Tests and test results

We evaluated our structure assessment pipeline over different reconstructed models. These models belong to the Mercado del Born area (phase II) and to Plaça Virrei Amat (phase III). For each reconstructed model we computed the corresponding signed distance map after registering it with the theoretical sections. Then, each map was converted into a texture and analyzed in order to detect structural elements, obstacles and obstructions.

Figure 12 shows different sewer models reconstructed and processed with our approach. Distance maps are displayed as heatmaps, where blue means low difference/distance against the theoretical profile and red means high difference. Structural elements such as manholes and junctions are clearly distinguishable in red.

Figure 13 shows close-up views from the inside of the sewers. Top row depicts structural elements such as manholes (left and middle-left), gutters (middle-left), junctions and connections (left and middle-right) and profile changes (right). Bottom row shows obstacles in the form of stones (left), pipes (middle-left), and deposited material (middle-right and right).

Figures 14 to 17 show different texture images for different reconstructions and their corresponding detection of structural elements and obstacles. The color legend is as in Table 1.

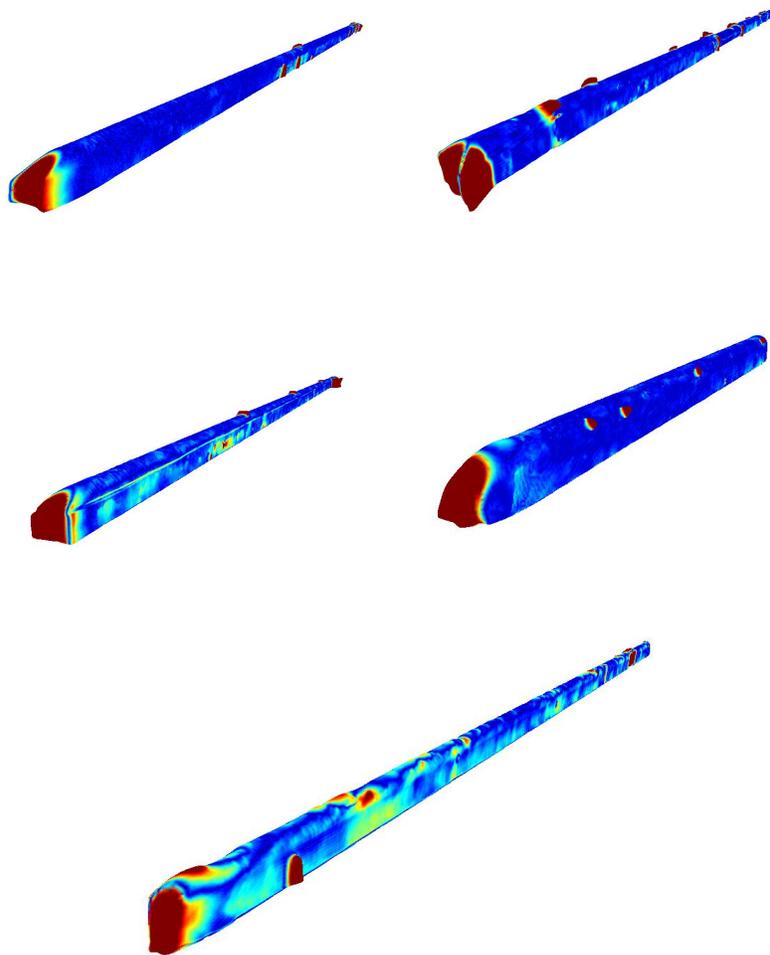


Figure 12: Distance maps computed on different sewers from Mercado del Born and Plaça Virrei Amat.

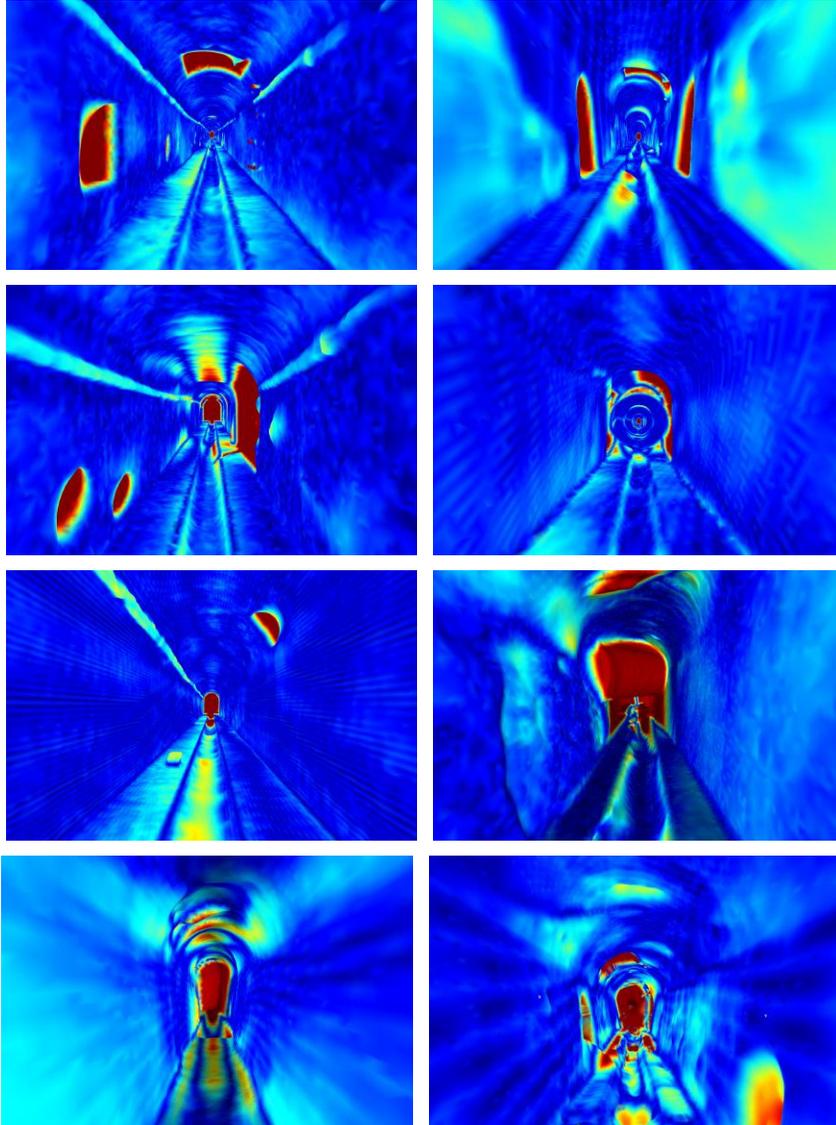


Figure 13: Several examples of structural elements (top) and obstacles (bottom) present in the analyzed sewers.

Error type	Color
Large reconstruction error	Red
Small reconstruction error	Orange
Profile change	Yellow
Manholes	Dark blue
Junctions	Green
Gutters	Purple
Connections	Light blue
Capacity lack	Pink
Obstacle or sediment	White

Table 1: Color legend for defects

The detections are mostly correct. There are very few false positives. When there is one, it is mainly because it is close to a reconstruction error. False negatives are always because of overlapping regions.

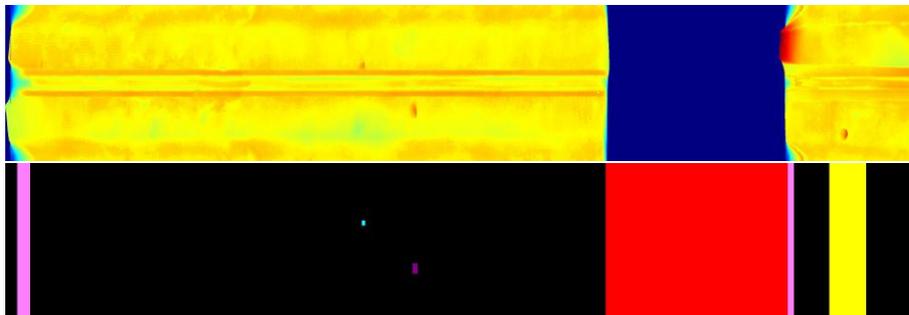


Figure 14: A texture image and the corresponding detection of structural elements and obstacles.

For instance, in Figure 16, there appear junctions and gutters close to profile changes. This is due to a bad reconstruction around a profile change. The rest of the elements are properly detected. Figure 17 is a reconstruction of a sewer section where obstacles were put on purpose by the ARSI team (see bottom-left image in Figure 13). We can see the obstacles very well reconstructed in the texture image and detected on the detection map. In particular, the largest object that was put there is as big as something that can diminish the sewer capacity by less than 95%, which is the threshold used for detection, and thus it is reported as a capacity lack defect and not as an obstacle or sediment.

The detected structural elements and obstacles are passed to the user interface through *xml* files specifying the exact 3D location and several properties of the elements, such as its size or depth.

For the next evaluation of the project, not only structural elements

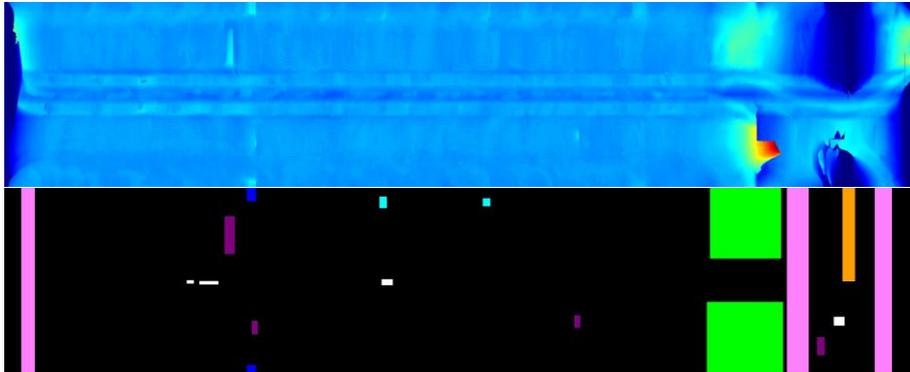


Figure 15: A texture image and the corresponding detection of structural elements and obstacles.

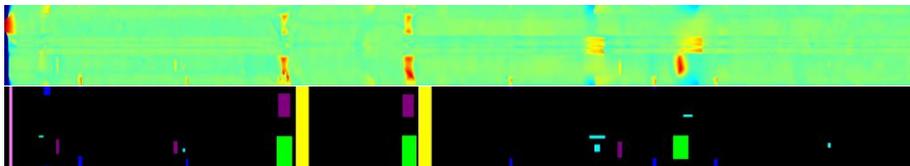


Figure 16: A texture image and the corresponding detection of structural elements and obstacles.

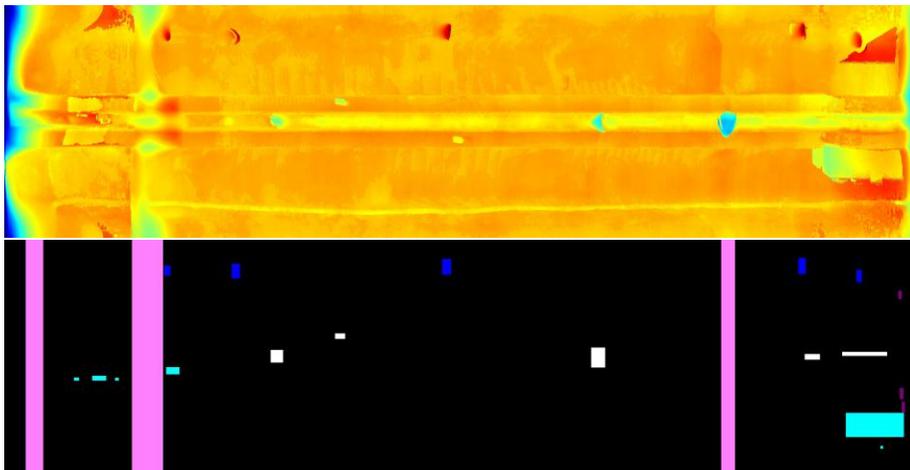


Figure 17: A texture image and the corresponding detection of structural elements and obstacles.

and obstacles will be reported but also defects on the walls, such as gravel (*áridos*) or wall erosions.

4 User interfaces

In the following sections we present the results of our effort to develop an intuitive and practical workflow allowing inspections brigades to plan and execute sewer inspections using the ARSI MAV, and to analyze the inspection data to generate informative reports for their clients.

4.1 Mission Planner

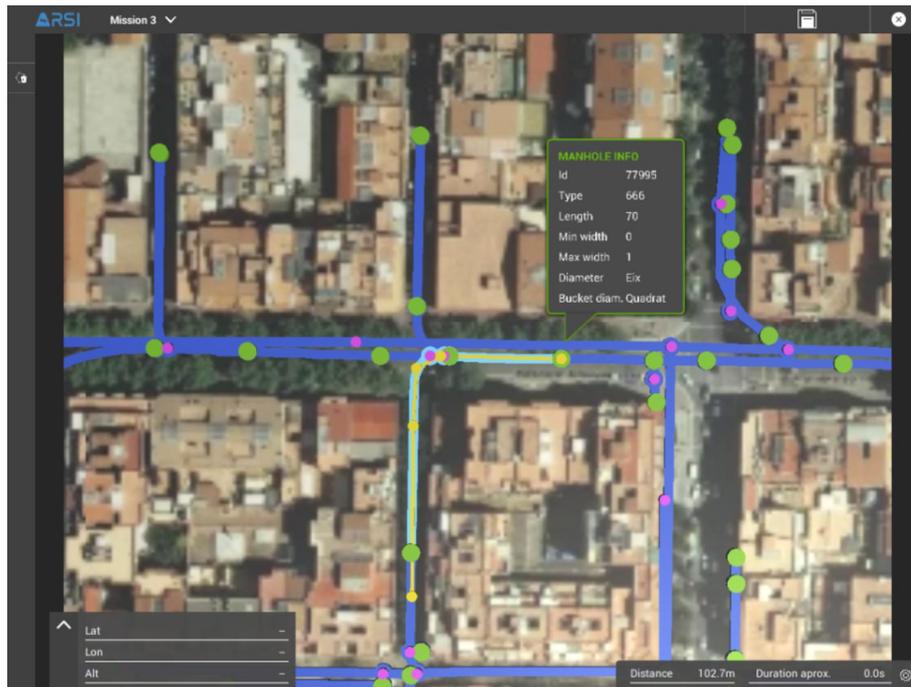


Figure 18: ARSI Mission Planner interface

A Mission Planner interface (see figure 18) was developed during phase III, allowing ARSI users to load GIS data and satellite imagery for the area selected by the client, and to plan a series of flights with the ARSI MAV to collect inspection data for the relevant sewer sections.

Each mission consists of a series of waypoints in GPS coordinates

(latitude and longitude) defining the desired flight trajectory. The first waypoint represents the MAV takeoff location, and the last waypoint represents the landing point. Each flight must start and end at an entry point into the sewers (typically a manhole) so that operators can access the MAV to deploy it, retrieve it, or replace batteries. The MAV has the ability to fly backwards and return to the same entry point after an inspection. Internally, mission files also include MAV parameters such as tolerances and PID control loop gains. These parameters are MAV-specific and are not modified by operators.

Tests and test results



Figure 19: GIS data for the Virrei Amat area (Barcelona) with sewer sections (blue polylines), manholes (red points) and singularities (blue points). Red polylines represent sewer sections that are too narrow (<70cm) to be inspected with the MAV

At the July evaluation, we demonstrated the Mission Planner by loading GIS data of the Virrei Amat area provided by BCASA (see figure 19) and by planning a series of missions with the ECHORD evaluators. Once reviewed, the mission files were copied onto a pen drive and passed to

the MAV operator for execution. In the future, mission files could be uploaded to a cloud storage, and later downloaded onto the Operator Console.

The inspection of the Virrei Amat area was carried out as follows:

1. Flight from manholes A to B
 - Open manhole A (depth: 7.5m, diameter: 70cm)
 - Deploy ARSI MAV and WiFi router
 - Section type: T141B
 - Total length: 90m approx.
 - Replace battery after flight
2. Flight from manhole A to intersection C, then backwards to A
 - Section type: T141B
 - Total length: 60m approx.
 - Retrieve MAV after flight
 - Close manhole A
3. Flight from manhole D to manhole E
 - Open manhole D (depth: 6.9m, diameter: 70cm)
 - Deploy ARSI MAV and WiFi router
 - Section type: T111
 - Total length: 70m approx
 - Recovery behavior¹
 - Open E extract the drone.
4. Flight from manholes D to F
 - Replace battery before flight
 - Section type: T111
 - Total length: 100m approx.
 - Replace battery after flight
 - Open manhole F to extract the drone

¹During this flight the MAV got trapped between the narrow walls ~25 meters before manhole E. The pilot triggered an emergency landing and the drone landed in the sewer bucket undamaged. A recovery behavior was then executed: the MAV took off safely without any help from the inspection brigade, resumed the inspection, and landed at manhole E.

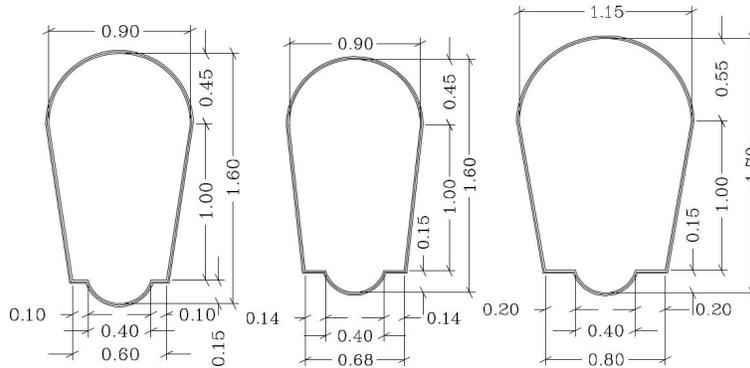


Figure 20: Sewer section types encountered in the Virrei Amat area. From left to right: T111, T114A and T141B

4.2 Operator Console

The Operator Console is the user interface used by inspection brigades on the field to carry out sewer inspections. It was significantly upgraded in phase III: the interface used for the July evaluation featured GIS data display of the sewer networks along with satellite imagery (Google Maps or other map sources such as Mapbox or OpenStreetMaps). Operators could load ARSI missions into the Console, display them on the map and query the GIS data (eg. manhole depth, type of sections, etc.) to plan operations before execution.

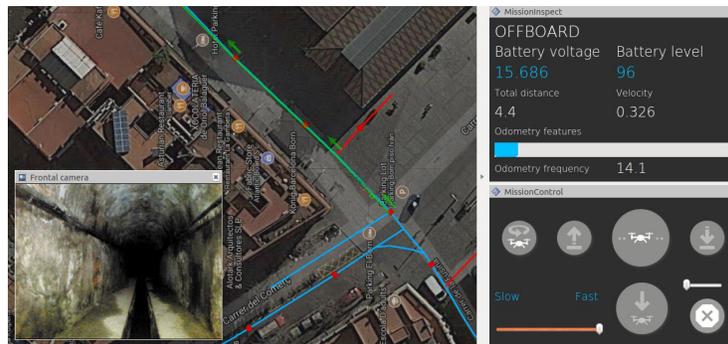


Figure 21: ARSI Operator Console with GIS and map display, live video feedback, and control panels

Tests and test results

Before each inspection, the Operator Console connects to the MAV over WiFi to perform flight checks, notifying users of any sensor or system failure. The mission is then sent to the MAV for validation, and executed from our control panel by issuing simple high-level commands such as "start mission", "pause", "land" etc. The Console provides live video feedback from the MAV and displays the mission trajectory on the map as well as sensor data. Operators without any MAV piloting experience can use the ARSI system.

4.3 Data Analysis Interface

The ARSI Data Analysis tool is the user interface that provides an integrated visual representation for completed missions, combining the data collected by the drones and the post-processed information generated by the data processing components.

In phase III, with the aim of improving the compatibility between all the components involved in the inspection process and to facilitate the automation of the entire ARSI workflow, a new generic mission data model has been defined and integrated in the analysis interface.

Also the 3D reconstruction loading process has been upgraded in phase III, in order to grant scalability when working with heavy 3d generated models. The 3d reconstruction is now segmented into blocks of constant size that can be progressively imported by the ARSI Data Analysis tool.

The map view of the mission has also been updated to show, in addition to the route and the position of the drone, the sewer infrastructure entities associated to the mission (sections, manholes and singularities), allowing access to valuable information for the operator.

A new heatmap based visualization mode has been included, which offers to the operator a faster and intuitive way to recognize defects and potential conflictive areas.

The representation of defects and structural elements has been also improved in this new version. Whereas In phase II detected incidences were almost exclusively managed through the UI timeline, in phase III the integration of incidences in the 3d reconstruction has been reinforced in terms of visualization and interaction. Defects are highlighted in a more precise way, at the exact position of the model where they were detected. Basic contextual information about incidents can be shown to help fast identification. It is also possible to interact with the 3d representation of incidences in order to obtain detailed information. By other hand new filtering capabilities have been included to facilitate

defect identification and serviceability evaluation tasks by the operator. Defects detected in the data processing stage can be filtered by typology, evaluation status, and other type specific properties like depth, area, etc.

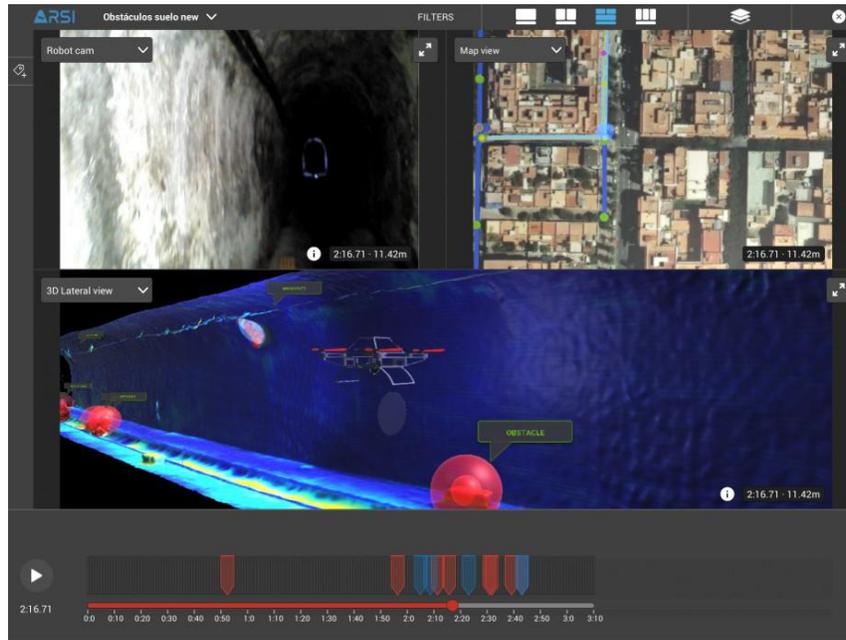


Figure 22: Data Analysis Interface. Heatmap visualization mode

Tests and test results

When a sewer inspection is completed and all the post-processed data has been generated, the resulting mission data structure can be imported into the ARSI Data Analysis Tool. After the new mission inclusion, the user interface shows it in the list of available missions. Then the operator can initiate the loading process and inspect the mission using the different viewers included in the user interface. At the July evaluation, we demonstrated the ARSI Data Analysis functionality by loading the missions carried out at Virrei Amat area.

5 Conclusions and future work

The visual odometry issues encountered in the Virrei Amat and Plaça Sol de Baix areas were unexpected, given that our approach had per-

formed well in Mercado del Born and in other areas visited during phase I and II of the project. Since the odometry estimation (velocity in particular) is central to our navigation system, it must be the focus of our work in the coming months.

We will use the data collected in these areas to evaluate if other types of visual features and feature descriptors could perform better, and investigate other visual odometry algorithms as well as alternative sensors such as laser or optical flow cameras for velocity estimation (see figure 23). We will also iterate with DroneTools to try and further reduce the platform dimensions, in order to achieve more forgiving control margins when operating in the narrowest sewer tunnels.

The odometry issues not only prevent the MAV from flying reliably in narrow sewers, they also result in poor inspection data quality due to the unstable flights. This is understandably a major concern for our clients, since the ARSI system is above all a sewer inspection system, and therefore must deliver high-quality data for 3D reconstruction, structural analysis and defect detection. Improving the quality of the inspection data will be our main objective for the remainder of this project.

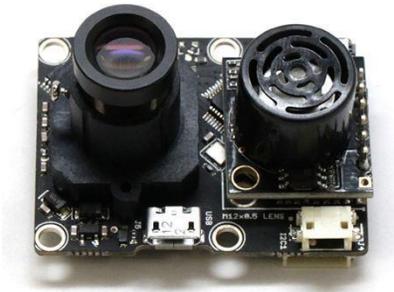


Figure 23: PX4Flow optical flow sensor for the Pixhawk autopilot