Vortex NTNU 2022 - Technical Design Report

Jonas Fillan *(Hardware Technical Lead)*, Christopher Strøm *(Software Technical Lead)*,
Tarek El-Agroudi *(Autonomous Lead)*, Nicholas Bresina *(Embedded Lead)*, Ivan Gushkov *(Perception)*

***Abstract*—For the academic year of 2022, Vortex NTNU has focused on improving upon the modular design of the 2021 Beluga AUV. The Beluga AUV has received a complete overhaul of its mechanical design, and the electrical housing and PCBs have all been reenvisioned. A computer vision software pipeline has been built from the ground up, with novel extensions of well-tested algorithms. The existing control system for dynamic positioning and 2D line following has been completed by the addition of a virtual target path following controller that allows the AUV to optimally follow arbitrary paths in 3D space. Lastly, a framework for developing solutions to individual tasks has been designed and follows a novel search-converge-execute scheme.**

***Index Terms*—Vortex NTNU, AUV, NTNU, RoboSub, Marine Robotics, Autonomy, Mission Execution, Product Development, Computer Vision, Underwater Imaging, Constrained Control, Path Following, Acoustic Localization**

I. Introduction

VORTEX NTNU is a non-profit student organization from the Norwegian University of Science and Technology in Trondheim, Norway. The focus of the organization is to introduce students to the field of marine robotics, a field which is of great importance to Trondheim, and Norway as a whole. Trondheim is home to many marine robotics companies and the fjord is important for Norway's technological development as the first ”technological playground” for all kinds of unmanned and autonomous vessels [1].

Vortex NTNU was founded in 2015 developing ROVs, but shifted our focus towards AUVs and RoboSub in 2018. In 2020, work started on a brand new AUV for the 2021 RoboSub competition and the work of this AUV continued for the 2022 competition. Two years of work and thirty to forty thousand work hours have resulted in the Beluga AUV.

II. Competition Strategy

A. Strategy: Task complexity

One of the first steps taken during the concept phase of the development was establishing a group that would evaluate competition task complexity. This group consisted of members from different parts of the organization and each competition task was evaluated with a complexity score from 1 to 10 for the different technical groups of the organization. Taking the sum of these scores and comparing them, factoring in the maximum scorable points for each task, provided a tool for us to decide what tasks to prioritize to be able to score as many points as possible.

Another important aspect of the competition is the task dependencies created by the course design and competition rules. For example, if you are not be able to complete the gate task you will not be able to continue doing the other tasks. For this reason, the task complexity evaluation alone was not enough to decide what tasks to focus on. Early in the process we created a flowchart to visualize the dependencies we had to take into account. This resulted in the chart seen in Figure 1.

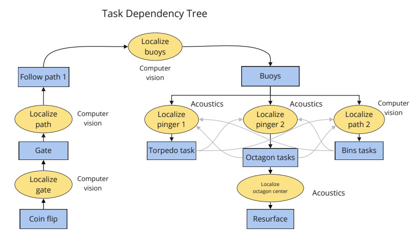

Fig. 1. The competition task flowchart used to visualize dependencies

B. Strategy: Generalization vs. specialization

While last year's keywords were modularity and adaptability, this year has seen both modularity as well as specialized solutions to perform specific tasks. This applies both for hardware and software solutions. For controls and navigation, a general 3D virtual target path following algorithm rooted in classical model-based control theory has been implemented on top of the existing dynamic positioning and straight line path following control stack. This general motion control framework was then contextualized with the 2022 competition task details, specializing the control modes and interface to allow the AUV to autonomously and accurately position itself along the 2022 course.

A state-machine approach has been chosen to implement the competition tasks. Creating stand-alone solutions for each individual task without some kind of framework was deemed to be a poor approach, since this requires exceptional cross-team cooperation and does not exploit the overlapping aspects of certain tasks. The solution the team came up with as a general mission execution framework is the so called *Search-Converge-Execute* framework. It is a framework for finite state machine structure that facilitates simpler communication and design between the different software and hardware groups, by abstracting away from task specifics. However, the framework is flexible enough to allow for specialization to each specific task, while maintaining the search-converge-execute pattern.

C. Strategy: Task execution

1) Gate: The gate state machine will spring into action once the gate is detected after the coin flip, and will initially trigger the AUV to move directly in front of the gate, while the computer vision algorithms refine the gate pose estimate. Once converged, a path that passes through the gate will be generated and followed. To maximize points, the AUV will also roll after passing through the gate. Performing the acrobatics using roll was decided since the AUV is very maneuverable in roll, and since roll is observable using the IMU, a good pose estimate can be retained by ignoring the DVL measurements that are generated while rolling.

2) Buoys: Searching and converging on pose estimates for the buoys is the same as the Gate task. However, for this task we are also required to deploy a neural net in order to touch the correct buoy.

3) Torpedoes: The Torpedo task will be completed in a similar manner to the buoys, with the differences being that the execute stage of this task will be to fire a torpedo, and that the AUV will dynamically position itself in such a manner that the torpedo mechanism aligns with the estimated poses of the cutouts in the buoys.

4) Bins and octagon: Both of these tasks involve the gripper, and were thus deemed the most complex tasks. For this reason we have not prioritized these for the 2022 competition. However, the team will attempt the marker dropper and acoustically localizing the octagon.

D. Competition execution

Based on the task complexity and task executions covered by sections II-A and II-C respectively, a best-case achievable competition run was crafted. This is a run that follows the enforced order of Figure 1. After the buoys, the acoustic localization system will be used to find the position of pinger 1. There, the AUV will execute the torpedo task. After this, the second path will be searched for and followed. The AUV will attempt to drop the marker in a bin, but will not interact with the lid. Then, the second pinger will be localized. For the octagon task, the AUV will simply position itself above the center of the octagon and resurface.

III. Design Creativity

A. Mechanical design

This year we have focused on going from a modular to a more specialization AUV design. A big design change was therefore to move away from the aluminum profile frame used last year. This design gave great flexibility and modularity but lacked rigidity and corrosion resistance. After researching different materials and methods it was decided to use 10mm water jet cut Polyoxymethylene (POM) sheets, assembled using angle brackets and machine screws. The choice of POM was due to its rigidity, workability, corrosion resistance, and prior experience with the material. Water cutting the sheets meant that holes for components had to be modeled before cutting, making it rigid but fixed with limited room for changes afterwards.

1) Pneumatic actuators: To be able to perform the manipulator tasks, various actuators are needed. This year it was decided to design new manipulators, based on the pneumatic actuation system designed together with our sponsors SMC Pneumatics last year.

To reduce the number of actuators, both the marker dropper and torpedo system were designed with elegant trigger solutions to be able to actuate two objects at different times with only one pneumatic cylinder. The torpedo system was subject to rigorous testing using AI and image recognition to find the torpedo design with the longest and straightest trajectory. The old gripper did no longer satisfy the needs for the competition tasks and had to be replaced. As the objects that were to be manipulated in the competition were not known at the time, a general purpose solution with interchangeable parts was chosen.

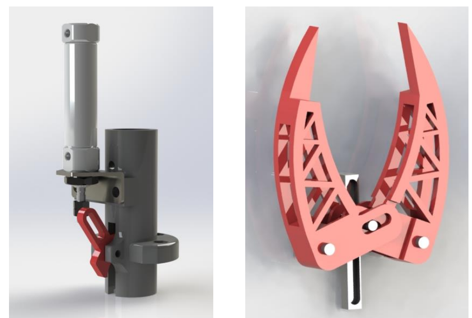

Fig. 2. Render of the new pneumatic marker dropper and gripper.

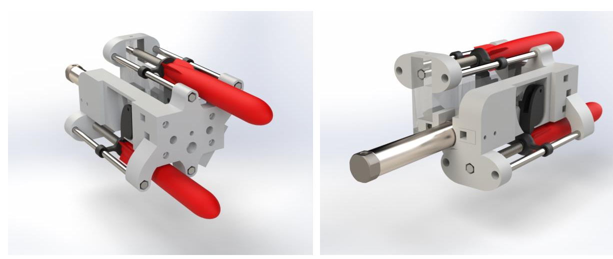

Fig. 3. Render of the new torpedo system.

2) Drone shell: The white outer shell of the drone serves several purposes and was 3D printed using white PET-G filament. Most importantly, the shell improves hydrodynamics by covering components that disrupt the flow of water through and around the drone. The drone will therefore be more efficient, stable, and easier to model for simulation. The shell is also an important piece in hiding jagged edges and making it look like a more finished design. The design was inspired by the Beluga whale and its hump, hence why it's named Beluga AUV. The shell also provides a large area for sponsor logos.

B. Electronics design

As most of the electronics design was finished last year, the focus this year was shifted towards reducing the size

Fig. 4. Render of the Beluga AUV with its shell.

of components and making the system more robust and reliable.

The first step was to create a better internal structure for the electronics housing. Last year there were several reliability issues due to lack of space in the housing, causing cable snags and shorts. This year the designer of the structure approached the problem with what was dubbed *the pizza tray principle*. This meant a board with all the components fixed to it which could slide on tracks in the housing, kind of like a pizza tray in the oven. This design, together with reduced size of many components, made it a lot easier to work on the electronics.

1) Power Distribution: One of the most significant size reductions was the new Power Distribution Board (PDB). This board is able to transform 14.8V battery voltage to 3.3V, 5V, 12V, and 24V to accommodate the different components of the drone. As the focus last year was on modularity the PDB was equipped with more connectors and transformers than needed. This made it possible to significantly reduce its size by removing the unused components. In the end, the footprint of the circuit board was reduced by more than half.

2) Acoustics: Reducing the size of the PDB also made way for a new addition to the electronic system, the acoustics board. This board is purpose-built to process the acoustic signals captured by the five hydrophones on the drone. The board is designed to do pre-filtering before converting it from analog to a digital signals using a 12bit ADC. The data is then available to be read over a parallel communication interface by the Teensy 4.1 microcontroller attached to the circuit board, which the software team uses for further processing.

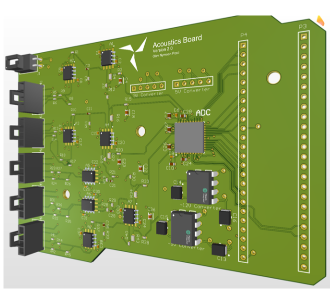

Fig. 5. The acoustics board, tasked with processing the hydrophone signals.

3) Connector lid: The last step in making the system more robust and reliable was to replace all cable penetrators with subsea connectors. The purpose of this change was so that the lid and the internal electronics structure could be completely removed from the drone and worked on without having to have the drone nearby. To make this possible MacArtney SubConn®connectors and Cobolt connectors from Blue Trail Engineering were used. The lid was also laser engraved by a local mechanical workshop.

C. High level task execution

The hardware design of the AUV has given the software team a platform that - on paper - is capable of solving a large majority of the tasks in the 2022 competition. To realize parts of this, a high level task executor was developed. This is a software module that determines which task to perform and which objects to look for at any given point in a run. The state machines for all tasks follow the same novel *search-converge-execute* framework, as introduced in section II-B. The interactions between the three states are covered in detail in appendix C-A.

A novel approach to bridge control and perception for task execution was developed in the form of the *landmark server*. The landmark server holds data and semantic information from each item that is expected to be in the pool and is updated whenever any object is detected, regardless of the current task, in order to fully utilize information that the AUV picks up during a run.

D. Computer Vision

Computer vision development is still in its early stage for Vortex NTNU, which has allowed the team to build an in-house pipeline from the bottom up. Emphasis has been placed on fully utilizing the prior information given through task descriptions, while focusing on simple and sound mathematical formulations rooted in theory from both target tracking and classical estimation theory, and of course creative applications of this. That is, most solutions rely on the age old engineering mantra that simple is good, and black-box neural networks have thus only been applied where semantic information beyond color and shape is required.

The novel developments of the computer vision team of 2022 are Principle Component Analysis (PCA) orientation estimation and filtering based on Gaussian Mixture Models (GMM).

1) PCA: Objects like the gate and buoys are planar in the sense that the variance in distribution of detectable parts of the object from its geometrical center will vary much less in one spacial dimension relative to the other two. This is exploited through PCA and singular value decomposition on features extracted using depth and color data from the front facing camera in order to estimate a two-dimensional plane in which the object lies. The third axis of this coordinate frame can be constructed from these, yielding a 3DOF orientation estimate of all planar objects detected.

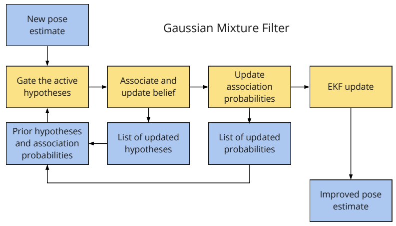

Fig. 6. The workflow of the Gaussian Mixture Filter.

2) GMM: Filtering the PCA based pose estimates to increase robustness was approached as a joint target tracking and existence detection problem. An innovative algorithm was proposed and developed based on a Gaussian Mixture Model, which admits creation, termination and probability evaluation of multiple tracks in parallel, based on association of incoming pose measurements with active track-hypotheses. Based on the probabilities the algorithm calculates, a winning track can be declared as an existing object. This solution has the goal of unifying a sensor fusion solution for robust search and final pose estimation, the workflow of which can be seen in Figure 6. Notice how GMM acts as a bayesian extension of the EKF.

E. Virtual target path following

In order to follow more complicated paths like those found in the *search* state, a novel 3D path following controller based on [2] by a former Vortex member was implemented. This is known as a virtual target controller and is made by defining a control law from the reference model of a virtual target constrained to move along the given path. The reference model is derived from a simplification of the Fossen motion model in [3] that is valid for low-velocity maneuvers. The control law includes actuator constraints which allows the algorithm to control the along-track speed of the virtual target, by formulating the inverse mapping of the control allocation as a quadratic program and solving using a trust-region based optimizer.

F. Acoustic localization

While localizing the tasks marked by pingers can be done using computer vision with random search, it is more favorable both in terms of points and robustness to localize them by the pings as intended. To solve this an in-house solution for hyperbolic positioning has been developed. Section III-B2 covers the hardware implementation that enables this. Once a pulse is detected by the microcontroller, the time differences of the received pulses on each hydrophone are determined by cross-correlation. The pinger position relative to the AUV can then be calculated from the time differences using multilateration. These position estimates will be accumulated over time which allows for refinement through smoothing and outlier detection. The entire flow of this design can be seen in Figure 7.

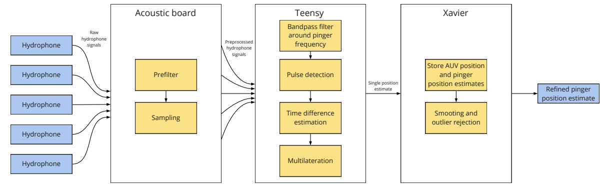

Fig. 7. The data flow and processing steps for the acoustic localization stack.

IV. Experimental Creativity

A. Hardware Corrosion test

Due to issues with corrosion on many components on the old drone, a corrosion test on the metallic components of the drone was carried out in a corrosion chamber. A cyclic program with spraying and drying was used. The componets tested included the 316L screws and nuts, 6061 Aluminium profiles and brass inserts. The ASTM G85-19 standard was followed, using synthetic seawater instead of acidic solution, and over a period of 8 days.

Test results showed some surface corrosion on the 316L screws, both when connected to the Al profiles and without galvanic contact. It was determined that this would not compromise the mechanical properties of the screws and bolts. Brass inserts corroded quite a bit when in galvanic contact with the 316L screws, and would in time make it difficult to remove the screws.

B. Pooltests schedule and strategy

Small scale pooltests were initially conducted on the weekends of October and November 2021 to verify that existing code worked. Large scale pooltests were conducted during the spring, and culminated in a two-week long easter pooltest where all parts of the AUV design was tested. Appendix B covers these plans in detail.

C. Control and localization

The virtual target path following was firstly tested in simulations. Results from this can be seen in Figure 8, which shows the reference model smoothly following the path of the virtual target in all three linear degrees of freedom, as well as for heading. The overshoots when turning are explained by the actuator constraints in the objective function for the control law. After tuning the controller for the simulator, it was used to maneuver the AUV during the easter pooltest.

D. Computer vision

Many of the novel additions to the computer vision pipeline were vigorously tested throughout the spring semester. Figures 15 and 16 in appendix C-B show two approaches to preprocessing. After preprocessing, the next step in the pipeline is feature detection. The performance of the full-scale feature detector is seen in Figure 9, which

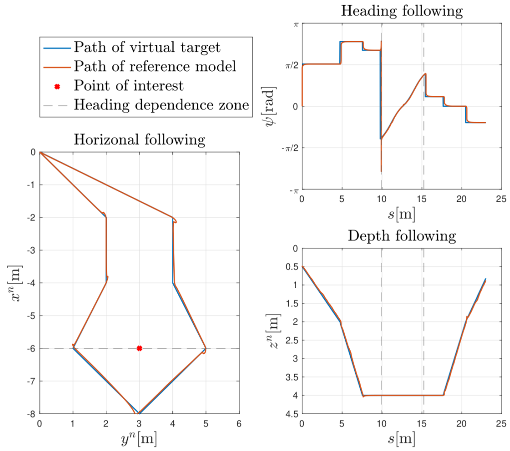

Fig. 8. Performance of the virtual target 3D path following algorithm in the simulator.

shows all parts of the gate being correctly identified, while the pool floor and surface reflections are kept out. The individual steps of this stage is also covered in appendix C-B.

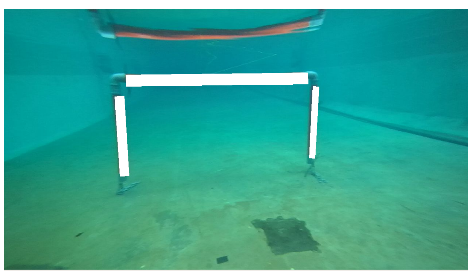

Fig. 9. Full feature detection stack applied to a recording of the gate. Note that compared to Figure 17, the outliers have been filtered out, and only the gate is detected.

The PCA approach to pose estimation as discussed in section III D was tested on the detected features directly given by the stereo camera. As seen in Figure 10, the object pose is well estimated, given the amount of outliers and zero processing done to the features.

E. Acoustics

The estimation scheme defined in section III E was implemented and tested using an in-house simulator which allows for arbitrary placement of hydrophones and sources in 3D space. Performance of the first two stages of the algorithm can be seen in figure 11. A notable observation is that the unfiltered position estimates follow a hyperbola with the center in the acoustic source. Smoothing and outlier rejection will then allow for partial resolution of the ambiguity in the position of the source. The mean of all distance gated measurements is shown as a red dot in Figure 11, and highlights how simple post-processing can yield meter-level accuracy.

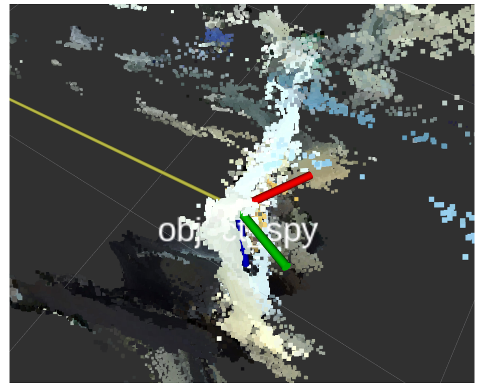

Fig. 10. PCA based on SVD to estimate the pose of the spy. A naive approach using all detected features and no outlier rejection was used to show that the algorithm still manages to converge fairly well with poor feature detection.

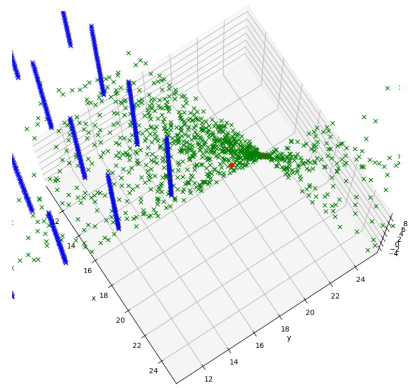

Fig. 11. Estimates of the pinger position in 3D for sampling rate of 500kHz. The red dot shows the mean of distance gated estimates. The ground truth of the source is marked with a red cross.

Acknowledgment

Vortex NTNU would like to show gratitude towards all members and alumni that has been part of this two-year project. Our members have sacrificed a lot of time and energy from their studies for this project and the results would not be possible without this dedication. The results of the work done by our members will be used for years to come and we hope the experiences and interest for marine robotics will be of use in future jobs and hobbies.

Without our sponsors, nothing of this would be possible. This year, Kongsberg Maritime decided to join our project as main sponsor and we could not be happier about it. Their support, both financial and technical, has been invaluable. Big thanks to our other sponsors [4] for the support and to the Department of Engineering Cybernetics at NTNU, for providing our office and workshop facilities.

Much of our motivation and for what we do has roots in the immense interest from the Norwegian marine community and academia. A huge thanks to all our supervisors, friends, associates, competition organizers and competitors for their willingness to help, discuss and challenge our solutions.

References

[1] G. Wolden, "Test site opens for unmanned vessels," Available at https://www.sintef.no/en/latest-news/2016/test-site-opens-for-unmanned-vessels/ (2022/06/09).

[2] A. Kristoffersen, "Guidance and control system for dynamic positioning and path following of an auv exposed to ocean currents," 2021.

[3] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*, 2nd ed. Wiley-Blackwell, 2021.

[4] "Vortex NTNU's sponsor page." Available at https://www.vortexntnu.no/sponsor (2022/06/09).

[5] D. Berman and S. Avidan, "Diving into haze-lines : Color restoration of underwater images," 2017.

APPENDIX A
COMPONENT SPECIFICATIONS

Component	**Vendor**	**Model/Type**	**Specs**	**Custom/ Purchased**	**Cost**	**Year of Purchase**
Buoyancy Control	Diab AS	Divinycell H60	Buoyancy foam	Sponsored		
Frame	ePlast AS	Polymer plates	Polyoxymethylene (POM)	Purchased	NOK 5778 ≈ \$660.00	2022
Waterproof Housing	ePlastics JM Robotics	8" Acrylic Tubing 8" flange and lid	15.75" length	Purchased	\$660.00 \$290.00	2021 2021
Waterproof Connectors	Blue Trail MacArtney	Cobalt Series SubConn®Ethernet	Connectors and Bulkhead	Purchased	\$1,685.00 \$1,762.00	2020/21 2020
Thrusters	JM Robotics	Blue Robotics T200	5.25/4.1 kgf FWD/REV Thrust@16V	Purchased	8 x NOK 1790 ≈ \$1,700.00	2021
Motor Control	Hobbywing	XRotor Micro ESC	60A 4in1 BLHeli32 6S	Purchased	NOK 1439 ≈ \$164.00	2022
High Level Control	Adafruit	PCA9685	PWM/Servo Driver 16-Channel, 12-bit, I^2C	Purchased	\$14.95	2021
Actuators	SMC	Pneumatics	4x solenoide valves, pistons, regulators and tubes	Sponsored		2021
Propellers	JM Robotics	T200 propellers	Clockwise and counter-clockwise	Included with thrusters		2021
Battery	Gens Ace	Tattu	4S - 6750mAh 25C, XT90 connector	Purchased	4x NOK 1125 ≈ \$480.00	2021
Converter						
Regulator	In-house	Regulator PCB	$V_{battery}$ to 3.3V, 5V, 12V and 24V	Custom	PCB: \$22.50 Comp.: \$93.00	2022
CPU	Nvidia	Jetson AGX Xavier Developer Kit	8-core 64-bit CPU 512-core GPU, 32GB RAM	Purchased	\$700.00	2021
Internal Comm Network		Ethernet, Serial, i2c				
External Comm Interface						
Compass						
Inertial Measurement Unit (IMU)	Sensonor	STIM300	Gyro: $\pm 400°/s$ Accelerometer: $\pm 10g$	Purchased	NOK 10 000 ≈ \$1,100.00	2021
Doppler Velocity Log (DVL)	Nortek	Nucleus1000 *Beta testing for Nortek*	**0.2-75m range, 300m depth** **100Mbit Ethernet**	Borrowed		2022
Manipulator	In-house	Custom	Pneumatic gripper, marker and torpedo	Custom		2021/22
Algorithms						
Vision	Scikit	RANSAC				
Acoustics						
Localization and mapping	Charles River Analytics	EKF				
Autonomy						
Open source software		ROS, Linux				

Appendix B
Pooltests

Small scale pooltests were initially conducted on the weekends of October and November 2021, and were mainly used to verify that existing code worked as intended, and that every member got familiar with their area of responsibility for the drone to be operational. Since the hardware teams were busy creating the final design of the AUV, the software team used the prototype developed by the previous team. Thus, the problem of balancing demands of design and testing was entirely eliminated. The first of the spring pooltests at a larger facility was conducted at the end of January, after a week-long workshop after Christmas - per Vortex tradition. The goal of this pooltest was to perform the prequalification tasks. However, the team quickly realized that the methods currently developed were not adequate to reach this goal, and so the pooltest was instead used to gather data. This data allowed the software team to develop more efficiently, by allowing us to verify correctness and performance of new algorithms. The final major pooltest happened during Easter, where two full weeks were spent testing new algorithms and rapidly prototyping solutions for the competition.

Fig. 12. Left: A view of the larger-scale testing facility - the MC Lab at Tyholt - as seen from the control room. Right: The smaller pool in the acoustic lab at our institute.

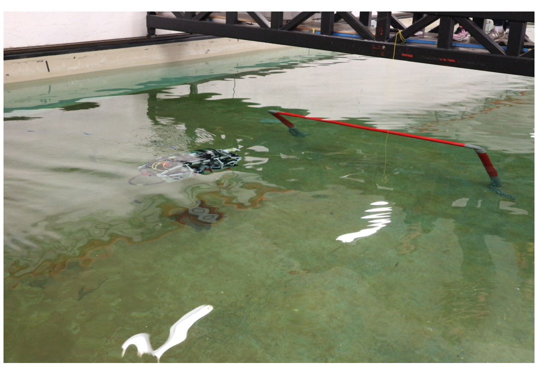

Fig. 13. The Beluga AUV attempting to estimate the gate position at a non-standard angle at the MC Lab.

Appendix C
Software development

A. Search-Converge-Execute

In the search state, the AUV moves in a general space filling search pattern while searching for objects related to

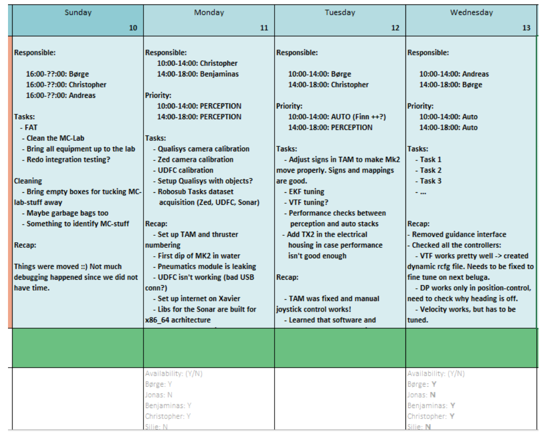

Sunday 10	Monday 11	Tuesday 12	Wednesday 13
Responsible: 16:00-??:00: Børge 16:00-??:00: Christopher 16:00-??:00: Andreas Tasks: - FAT - Clean the MC-Lab - Bring all equipment up to the lab - Redo integration testing? Cleaning - Bring empty boxes for tucking MC-lab-stuff away - Maybe garbage bags too - Something to identify MC-stuff Recap: Things were moved ::) Not much debugging happened since we did not have time.	Responsible: 10:00-14:00: Christopher 14:00-18:00: Benjaminas Priority: 10:00-14:00: PERCEPTION 14:00-18:00: PERCEPTION Tasks: - Qualisys camera calibration - Zed camera calibration - UDFC calibration - Setup Qualisys with objects? - Robosub Tasks dataset acquisition (Zed, UDFC, Sonar) Recap: - Set up TAM and thruster numbering - First dip of MK2 in water - Pneumatics module is leaking - UDFC isn't working (bad USB conn?) - Set up internet on Xavier - Libs for the Sonar are built for x86_64 acrhitecture	Responsible: 10:00-14:00: Børge 14:00-18:00: Christopher Priority: 10:00-14:00: AUTO (Finn ++?) 14:00-18:00: PERCEPTION Tasks: - Adjust signs in TAM to make Mk2 move properly. Signs and mappings are good. - EKF tuning - VTF tuning? - Performance checks between perception and auto stacks - Add TX2 in the electrical housing in case performance isn't good enough Recap: - TAM was fixed and manual joystick control works! - Learned that software and	Responsible: 10:00-14:00: Andreas 14:00-18:00: Børge Priority: 10:00-14:00: Auto 14:00-18:00: Auto Tasks: - Task 1 - Task 2 - Task 3 - ... Recap: - Removed guidance interface - Checked all the controllers: - VTF works pretty well -> created dynamic rcfg file. Needs to be fixed to fine tune on next beluga. - DP works only in position-control, need to check why heading is off. - Velocity works, but has to be tuned.
	Availability: (Y/N) Børge: Y Jonas: N Benjaminas: Y Christopher: Y Silie: N		Availability: (Y/N) Børge: Y Jonas: N Benjaminas: Y Christopher: Y Silie: N

Fig. 14. The structure that was used to plan pooltests day by day. This particular one is from parts of the easter pooltest.

the next task. Robustness is achieved by storing the starting position of the search and having a recovery state in case track of the object is lost or no object is found. The *Converge* states prepares the AUV for execution by stopping in its track and allowing computer vision algorithms to converge on the detected object pose. If the object pose does not converge, the AUV falls back to search, otherwise it enters the *Execute* state. The nature of execution will be specialized to the task at hand.

B. Other computer vision experiments

Figures 15 and 16 show two approaches to preprocessing. The former is based on contrast limited adaptive histogram equalization (CLAHE) and has the effect of making brightly colored objects like the gate - which would otherwise look washed out underwater - stand out more relative to the background. However, this stage also introduces textures in the background that may give rise to other problems further down the pipeline, especially for feature detection. The latter approach is based on the gray-world assumption, and has a similar effect of making the gate stand out more relative to the background while maintaining the textures of the original image. The downside of this approach is that there is no guarantee that the gray-world assumption holds underwater, see [5], which also may cause problems for feature detection and outlier rejection.

Figure 17 shows a feature detector based on the HSV color model and the gray-world-processed image data. While the gate is identified and masked, major parts of the pool floor is not filtered out. To reject these outliers, a Hough-transform based line detector was added to the feature detection stage. The confidence map of the stereo camera was also used as a mask in this step.

Another use of HSV and the Hough transform is detecting the path and determining where it is pointing. This can be seen in figure 18. The workflow of this algorithm is as follows: For every frame, the contour of the path is estimated using HSV feature detection. From this, a line

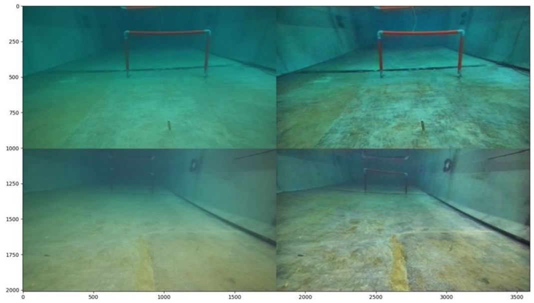

Fig. 15. Raw RGB image (left) compared to CLAHE preprocessing (right). CLAHE makes the gate stand out more compared to the background at both ranges, but also enhances textures in the background.

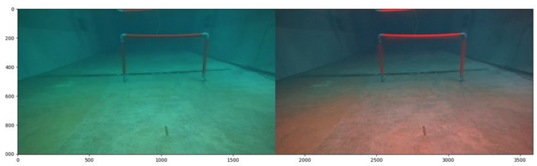

Fig. 16. Raw RGB image (left) compared to gray world preprocessing (right). Note that both the gate and parts of the pool floor is affected greatly by the preprocessing.

that points in the direction the contour is pointed to is determined. After collecting a set of lines, random sample consensus (RANSAC) is used to determine a maximal-inlier line, which is the final estimate for the path direction. In Figure 18 this is the green line.

C. Acoustics simulator

Figure 19 shows a generated view of the default configuration of the in-house acoustic simulator. 5 receivers are placed around a single source, operating at multiple frequencies.

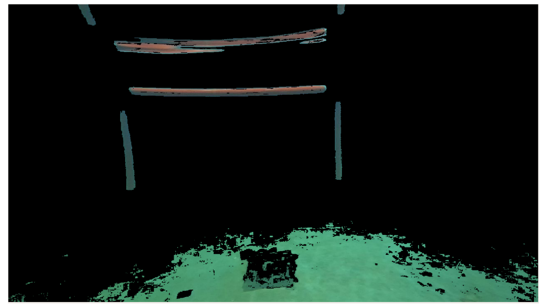

Fig. 17. HSV feature detection for the red colored gate. A naive approach using all detected features and no outlier rejection was used to show that the algorithm still picks up on the colors from the gate.

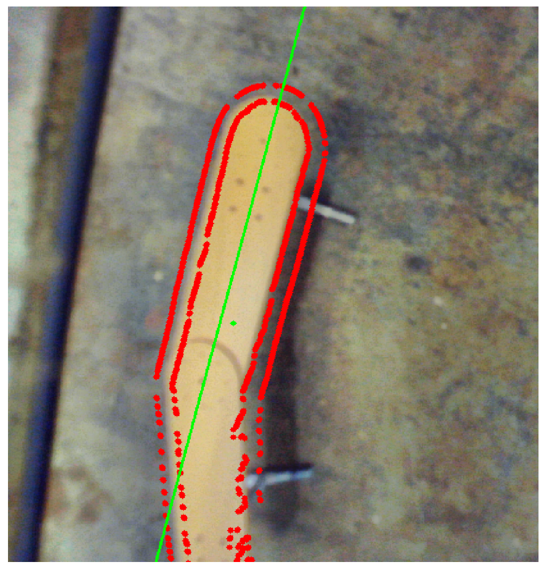

Fig. 18. Detections of the path and the estimated direction in which it points

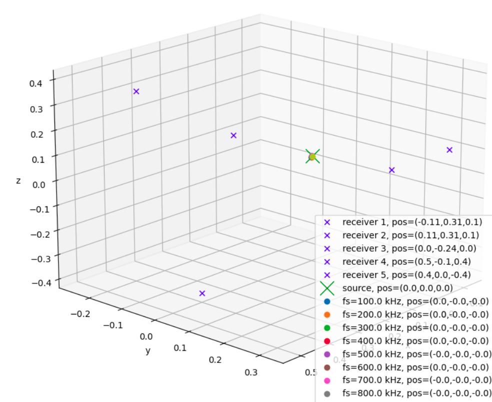

Fig. 19. A view of the simulator with a single source at multiple frequencies, and five hydrophones placed around it