Autonomous Docking for Intervention-AUVs using Sonar and Video-based Real-time 3D Pose Estimation

Jonathan Evans, Paul Redmond,
Costas Plakas, Kelvin Hamilton and David Lane

Ocean Systems Laboratory, Heriot-Watt University,
Edinburgh, EH14 4AS, SCOTLAND, UK

{j.evans ; p.redmond ; c.plakas ; k.hamilton; d.m.lane} @ hw.ac.uk

ABSTRACT

An Intervention-AUV (or I-AUV), is a hover capable AUV whose primary role is direct contact with subsea structures for measurement or physical manipulation of components. The aim of the ALIVE project is to develop an Intervention-AUV capable of docking to a subsea structure which has not been specifically modified for AUV use.

This paper describes the modular structure of the ALIVE AUV, including its distributed software architecture and in particular the **ADS (Autonomous Docking System)**. It then outlines the sonar and video sensor processing techniques used for real-time control of the AUV to perform tracking and 3D pose reconstruction. In addition, details of the system tests and practical trials used in the development process are described.

I. INTRODUCTION

Intervention-AUVs (or I-AUVs) are intended to be deployed in missions currently performed by Intermediate and Workclass-ROVs, in challenging offshore environments such as those faced by the oil and gas industry.

A relatively new concept, I-AUVs are still at a modest stage of development – but there are already a few vehicles which exist as “demonstration prototypes”. The ALIVE (Autonomous Light Intervention Vehicle) AUV project (and its earlier sister project – SWIMMER [1]) are in the forefront of the latest developments to create vehicles capable of autonomously docking and intervening on unmodified subsea structures. The commercial interest in these types of vehicle configurations is their reduced dependency on (or perhaps elimination) of expensive surface vessels, plus the highly time/cost critical nature of the work tasks to be undertaken by such AUVs.

An essential step along the road to autonomous intervention is automatic docking - the ability to automatically identify a subsea structure, transit the vehicle to within ~1m, and reliably dock (mechanically attach) the AUV. From there a range of offshore intervention tasks such as inspection, valve manipulation and subsea connector re-plugging (“hot-stabs”) are possible.

ALIVE is a world leading research project to develop a prototype Intervention-AUV. Tetherless, large (approx 4m x 3m x 2m @ 3000 Kg), and equipped with 6-DOF subsea manipulator, the vehicle is capable of autonomous docking to an unmodified subsea structure using sensor-based techniques (see Fig. 1).

The project consortium brings together expertise from the industrial and academic sectors:

- Cybernetix
- Hitec-Framnes
- Heriot-Watt University
- Ifremer
- Joint Research Centre (European Commission)

II. ALIVE: AUTONOMOUS LIGHT INTERVENTION VEHICLE

To perform automatic, precision control of a large AUV, the ALIVE project combines the following technologies:

- **Embedded, real-time SONAR-based feature tracking** to transit the AUV to within a few metres of the subsea Docking Panel.
- **High precision video-based position control loops** to manoeuvre the AUV to the required location for mechanical docking to begin.
- **INS and Video-based Dynamic Positioning (DP)** to hold the AUV at the right location while the mechanical docking proceeds autonomously.
- **High-bandwidth acoustic video transmission** to allow a pilot at the surface to monitor the progress of the autonomous docking, and intervene as necessary.

02 Jointly funded by the EU and partner companies, the ALIVE AUV is designed to dock to unmodified ROV docking panels. It does this by using robust sonar and video based control techniques to manoeuvre the vehicle from the surface deployment vessel (or platform) to within 10cm of the correct docking position. Mechanical docking is achieved using two special docking manipulators that extend up to 1m from the vehicle's body and can grab standard ROV docking bars (Fig. 2).

The docking system uses a three-stage process:

- **TRANSIT:** Information for the onboard Inertial Navigation System (INS) is used to transit from the surface vessel to a Safe Landing Zone (SLZ) up to 100m from the known location of the docking panel. The position and status of the AUV during the autonomous transit can be monitored at the surface station by periodic data transmitted using the acoustic modems (from Tritech, Aberdeen).

- **APPROACH:** In the approach phase, the AUV's Docking System uses repeated sonar scans to continuously track and calculate the vehicle's position relative to the docking panel. The vehicle's control system uses this information to manoeuvre within the SLZ to approximately 2m - 3m from the docking panel. Once the vehicle is close enough to observe the panel using the onboard video systems, the moves into the docking phase.

- **DOCKING:** During the docking phase, the combined sonar and video control systems stabilise the vehicle to within 10cm of the docking panel. Docking manipulators then extend, and grab the standard ROV docking bars.

Once docked, the intervention manipulator can be un-parked, and used to autonomously perform the required work tasks under pilot supervision using low frame rate video and sensor data transmitted over an acoustic communication link.

The AUV uses a collection of distributed computers and processes to control all aspects of its operation. This paper concentrates on the operation of the ADS (Autonomous Docking System) which is responsible for guidance of the AUV within the Safe Landing Zone, right up to the point when the docking grabbers are engaged.

The ADS (Autonomous Docking System) is a single canister "pod", with a high-performance dual-PC104 based computing platform. A modular software structure supports easy reconfiguration for different types of structures and docking panels. For primary sensors, the ADS used a high-frequency imaging sonar (Tritech, Aberdeen) and a ultra low-light CCD video camera (ROS, San Diego).

The outputs of these sensors are processed by real-time implementations of particle filter-based feature trackers and 3D pose reconstruction algorithms. The poses (relative position estimates) are then passed to the Autonomous Docking Planner (ADP). The ADP is responsible for overall planning and coordination (including fault monitoring) of the automatic docking process.

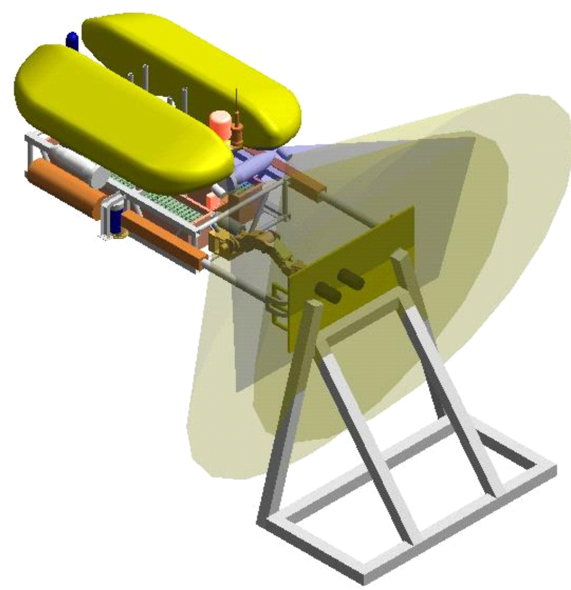

Fig. 1. Illustration of the ALIVE AUV intervening with a standard ROV Docking Panel.

Fig. 2. ALIVE AUV during initial tank tests.

III. ADS: AUTONOMOUS DOCKING SYSTEM

Heriot-Watt University (HWU) has been responsible for the design and implementation of the ADS for the ALIVE AUV. The complexity of the ADS has necessitated a distributed design involved 5 modules performing various functions (described below). The entire "Docking System" is known as the "ADS". The interface to the other AUV systems (for example, Vehicle Control, Navigation, Dynamic Position Stabilisation) is from a single module within the ADS. This is to reduce the number of dependencies between the internal and external modules. There are also some communication differences between the internal ADS modules, and those used between the other ALIVE modules.

A. ADS Internal Architecture

The Autonomous Docking System (ADS) comprises of five software processes (see Fig. 5):

- Sonar Docking System (SDS)
- Video Docking System (VDS)
- Docking Sensor Fusion (DSF)
- Autonomous Docking Planner (ADP)
- Docking System Controller (DSC)

The first four are built using a common module framework created by HWU. The DSC governs the behaviour of the other four modules while also presenting a single interface, or central point of contact, between the ADS and all external systems.

The common module framework is a suite of C++ Programming Language Classes which can be viewed as a ready to use application shell. It is simply waiting to be filled with an algorithm at a known point. Inter-Process Communication (IPC) is inherited by the Module in the form of another HWU designed component, the shared memory. The shared memory manages data-structures, encapsulating them in a message, sending them out on a communications link when they are updated, and updating the relevant data-structures when new messages are received. UDP (User Datagram Protocol) over IP (Internet Protocol) is used for communication, and distributed in a broadcast fashion over a small Ethernet LAN. Common behaviour in response to a set of generic Commands e.g. START, STOP, and the reporting of the module's state in a Generic Status message, is also inherited by the Framework Modules,

The DSC has a number of responsibilities. It controls the other four ADS modules using Generic Module Commands, such as Start and Stop. It also monitors their generic Status messages to ensure correct behaviour. The DSC also acts as an interface to the whole ADS, translating messages from external systems into Shared Memory messages which are sent to the Modules as appropriate, and translating the output message of the Modules to the correct format for the appropriate external system.

While the DSC is not directly derived from the Common Module Framework, it is constructed in a similar fashion using many common components. For example it uses one shared memory component to communicate with, and control, the four Framework Modules, and another adapted shared memory component to communicate with external systems. The two shared memories are extremely similar, due to the fact that the interface to the external ALIVE system is also based on UDP broadcast messages over Ethernet. However it is important to note that the entire ADS can be interfaced to different external architectures solely by changes within in the DSC.

B. Benefits

Within the ALIVE project, the single-point location of the interface between the ADS and the rest of the on-board AUV modules provides a number of benefits.

One of the foremost of these is the ability to easily interface the entire ADP to other external environments by simply changing the "outer" interface in the DSC.

For example, a variant of the DSC interfaces the ADS to a Cartesian robot in the laboratory's test tank facility (which happens to host a simulated vehicle); this along with the mounting of sonar and camera allows in-depth testing of the ADS modules solely within HWU premises.

In another variant the external interface is initialised in a point-to-point mode. This has allowed remote testing of the ADS with software running at the other ALIVE project partners over the Internet: Hitec (Norway) and Ifremer (France). The subsequent onsite integration work conducted in June 2003 was extremely successful.

More general benefits of the Module Framework and its inherited IPC include: greatly simplifying the software engineering task; and allowing Engineers/Researchers to focus on the core implementation issues of sonar and video processing, data fusion and path planning.

The Module Framework is highly reusable, as demonstrated by its use by HWU in the EU project AUTOTRACKER, and its adoption as the framework for HWU's own internally developed AUV/ROV "RAUVER".

C. SDS: Sonar Docking System

The Sonar Docking System (SDS) is tasked with performing the initial target acquisition and tracking while the AUV approaches the Docking Panel – possibly over ranges up to 100m. Once the AUV is within around 4m, it is used to provide initial estimates to the VDS (Video Docking System).

The module uses a particle filter (a form of Bayesian filter) to identify and then track the linear 2m structure of the Docking Panel in the real-time sequence of sonar images. Particle filters are known for their excellent model-based tracking capabilities, particularly when the target is noisy or obscured by other objects, or visual clutter.

The SDS outputs the 2D pose (relative position in the horizontal plane), plus the relative angle (yaw) between the AUV and Docking Panel at approx 0.3Hz (limited by the sweep rate of the mechanical scan). The SDS uses a Tritech Dual-Frequency Imaging Sonar, with 675Khz and 1100KHz channels. Typical sonar images obtained by the SDS are shown in Fig. 3 & Fig. 4.

The SDS pose is used by the ADP (Autonomous Docking Planner) during the approach phase to guide the AUV to within a few metres of the Docking Panel. The accuracy of the SDS is sufficient to calculate the pose of the Docking Panel to within 20cm, when 5m from the panel.

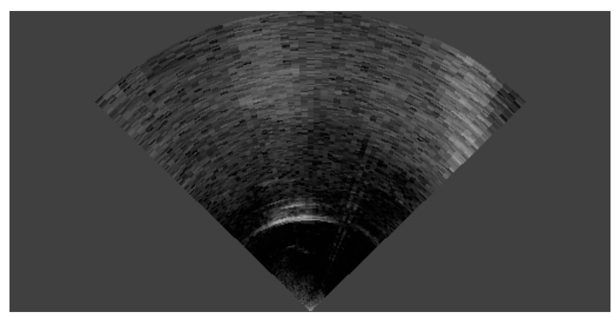

Fig. 3. Typical SDS sonar image (panel @ 30m).

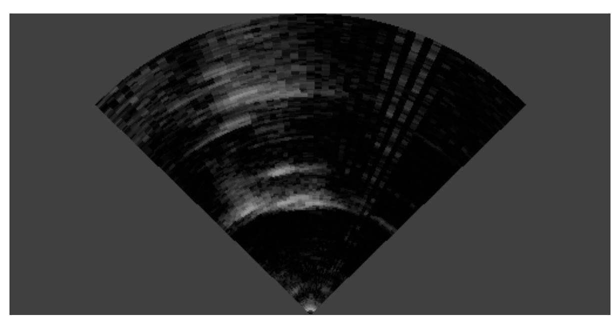

Fig. 4. Typical SDS sonar image (panel @ 5m).

.

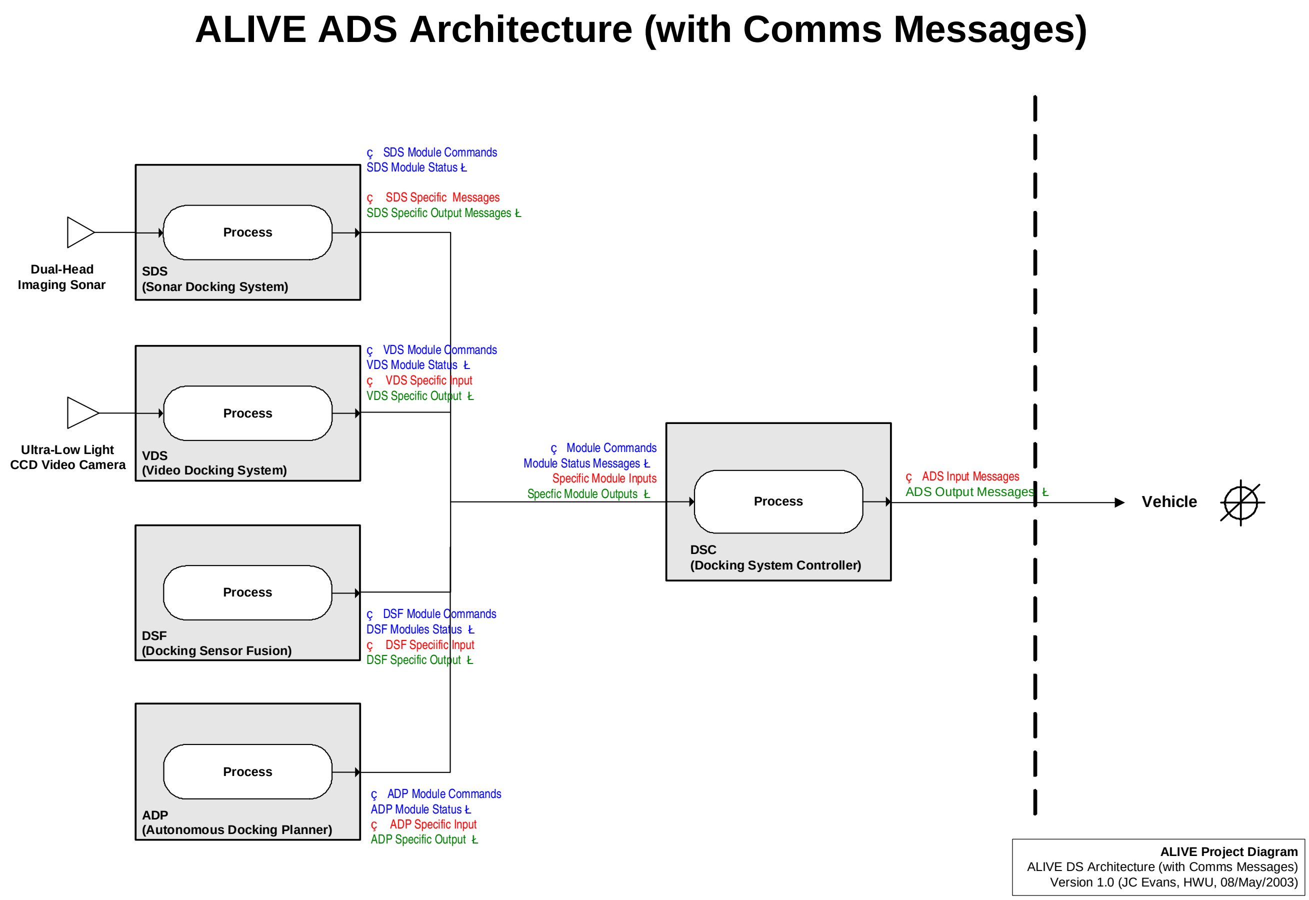

Fig. 5. Overview of Autonomous Docking System (ADS) architecture for the ALIVE Intervention-AUV.

D. VDS: Video Docking System

The Video Docking System (VDS) is tasked with performing accurate visual vehicle pose estimation during the final docking phase. The pose is calculated with respect to the docking panel (i.e. with a co-ordinate system origin centred on the panel). A typical sequence of images showing the AUV's approach to the Docking Panel can be seen in Fig. 7.

The objective of the VDS is to calculate the AUV's 3D-pose at video rates (~10Hz). This can then be used by the ADP module to control the AUV's final approach to the predefined docking position.

The VDS architecture is shown in Fig. 8. The VDS is constructed around a modular concept making it easy to change and extend. An ultra-low-light (~1x10^{-3} lux) camera is used for the image acquisition, in particular the ROS 20/20 Navigator with a focal length of 3.8mm. The system acquires greyscale images at full frame rate, with a resolution of 320x240 pixels.

The feature extraction subsystem is responsible for providing the data from which the pose will be estimated. Currently edge detection is performed utilising the Canny algorithm. Next, the detected edges are segmented into contours, and from these contours the three circles representing the outline of the circular valves on the Docking Panel are detected. It is the centres of this circles that are used as features in the subsequent pose estimation phase.

The feature selection module can easily be extended to different feature types, and part of the ongoing work is to use other feature types such as ellipses, corners, etc.

The VDS utilises a model-based procedure for pose estimation. This is made possible by the fact that a 3D CAD (Computer Aid Design) model for the target Docking Panel exists. The extracted features are paired with features in the CAD model. Then, the model features are reprojected [2, 3] onto the image given an initial pose estimate (originally coming from the SDS system), and an error is calculated between the extracted and projected features. An iterative procedure then calculates new pose estimates trying to minimise this error.

Before integration on the AUV, performance tests were run in simulation using MATLAB. The results (see Fig. 6) have shown that, in the last few metres of approach (i.e. just before docking), the performance of the VDS algorithm was not affected by the absolute range from the target.

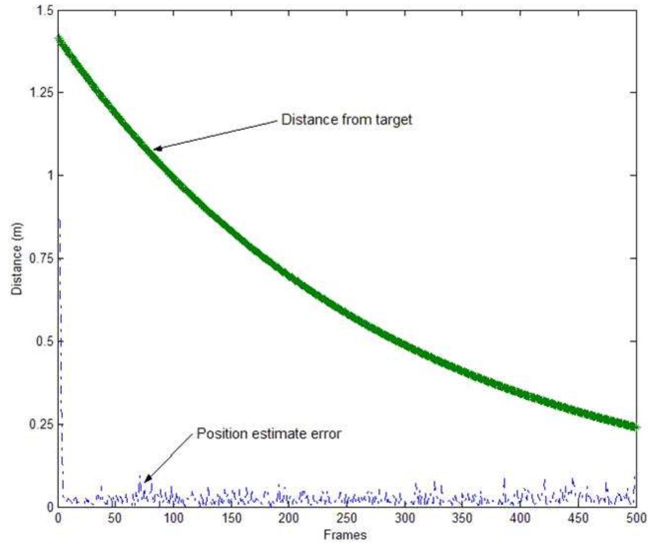

Fig. 6. Typical VDS pose error at decreasing distance from target .

E. ADP: Autonomous Docking Planner

Autonomous planners [4] have seldom been used on real-world robots due to their high computational burden, exotic technology and difficult development cycle. The ADP module is a constrained geometrical planning solution focussed solely on ALIVE's requirements.

The Autonomous Docking Planner (ADP) module (including the DSF, Docking Sensor Fusion module) is responsible for managing ALIVE's movement within the Safe Landing Zone (SLZ) - the area surrounding the Docking Panel (up to 100m radius). It decides whether to use position estimates from the SDS, VDS, or a fused combination. It receives waypoint commands from the surface, then creates and manages optimised movement plans allowing for sensor coverage, the vehicle's constraints, and the requirements of the docking systems.

This level of functionality is needed because there are several constraints active within the SLZ, and it is extremely difficult and time consuming for the human operator to allow for them. This is particularly true when working under pressure, and with the large latencies introduced by acoustic communications.

The constraints identified so far are:

- Due to the large mass of the vehicle, and the relatively low power-weight-ratio, the vehicle's position controller is better at hovering than short dynamic motions.
- The vehicle takes a significant time (up to 1 min) to completely settle into a new set-point using its AutoPOS controller – during which the damping causes the oscillatory overshoot to die down.
- Even when hovering and settled, when fed with only inertial data (INS) the vehicle still oscillates by several degrees in yaw, and tens of centimetres in translation. Although tuning work is ongoing, it is possible that these variations will remain larger than the accuracy required for docking
- The sonar and video sensors used for positioning within the SLZ, have limited fields of view (about 100 degrees). For continuous pose (position) measurements, the docking panel must remain within this field of view.
- The sonar and video docking algorithms have individual constraints of their own.
- There are illegal movements, such as trying to drive through the panel, or outside the SLZ
- Movement when close to the docking panel must be severely constrained, such as asking for (at most) 0.25m position changes, whereas when further away larger movements are acceptable to save time.
- When close to the docking panel, the vehicle must remain aligned with the panels axis, as any yaw is unacceptable at close distances, particularly with the docking arms extended

It is highly likely that more constraints will be discovered during further dock tests, and full sea trials.

The architecture of the ADP module is shown in Fig. 9. Pose estimates are continually received from the sonar & video docking systems, and fused by the DSF (Docking System Fusion) module. The fused pose estimate is then provided to the rest of the Planner module. Occasional waypoint requests are received from the surface, at which point the Plan Generator checks the requested position for legality (i.e. is it in an exclusion zone?) using the Constraint Information. If the request is legal then the Plan Generator creates a plan, consisting of a series of single-axis steps of x, y, z translations and yaw rotation. The use of single-axis motions is to minimise any cross-axis motion on the AUV's controllers that (at best) would take time to settle, or (at worse) force the sensor to lose "lock" on the Docking Panel.

The steps contained in this plan, and the coordinate frame used, depends on ALIVE's position within the Safe Landing Zone, as well as other constraints mentioned above, (Fig. 10).

From 100m range until 15m range the sonar docking system is unable to resolve the orientation of the docking panel, and so in this zone ALIVE drives straight towards the docking panel in Range Mode, whilst adjusting its heading to keep the docking panel centred.

This part of the plan is produced in ALIVE frame (see Fig. 11) with transform steps of 10m and rotation steps of 10 degrees.

At 15m range the panel can be resolved and so ALIVE moves into Cartesian Point Mode, where it always points at the centre of the docking panel to keep it in the sonar's view, and always moves so that it is aligned with the waypoint's y position. Once this is achieved it advances along the x axis until it has reached the waypoint. This part of the plan is produced in Dock Frame with transform steps of 1m and rotation steps of 5 degrees.

At 5m range ALIVE is dangerously close to the panel, and so moves into Creep Parallel Mode, where it always stays aligned with the panel's x-axis to prevent the docking arms colliding. This part of the plan is produced in Docking Panel Frame, with transform steps of 0.25m and rotation steps of 2 degrees.

IV. AUTONOMOUS DOCKING SYSTEM TRIALS

A. Shallow Water Trials

During the July 2003 ALIVE trials in IFREMER pool and dock/harbour the ADS v1 (single-step waypoint) module was tested in conjunction with all surrounding modules and docking systems, proving the underlying functionality and mathematics. It was tested in two ways, by using it as a 'tight' position control loop, and as a high-level waypoint controller. The first method proved unsuitable due to the vehicle's dynamics and the 'noise' inherent in the early-stage docking systems. The second, high level method proved to work extremely well – and under ADP control, the vehicle was brought to within docking range of the panel.

During July, August and September 2003 the ADS will be moved to version 2, with full planner functionality as described above. This will be continuously tested using the Ocean Systems Laboratory's own hover capable AUV/ROV, RAUVER, which has been upgraded to mirror ALIVE's vehicle controller software and mechanical docking arms. Weekly trials are undertaken in the Wave Tank and local lakes and reservoirs. Our distributed architecture also allows the possibility of remote trials, controlling the ALIVE vehicle in IFREMER from the Ocean Systems Lab in Edinburgh.

B. Deep Water Trials

Final deep water project trials of the ALIVE AUV prototype are planned for October 2003 in the Mediterranean Sea. The Docking Panel with be installed at flat area of seabed at roughly 200-300m water depth. The AUV will be launched from a surface ship and autonomously transit to the seabed approximately 50m from the Panel. Using a combination of the sonar and video docking systems, the AUV will approach the Panel. When the vehicle is correctly aligned, and stable the docking arm will be extended, and the autonomous dock completed. The subsea robotic manipulator can then perform the intervention (valve turn). This will be monitored using periodic images transmitted to the pilot using the high-bandwidth acoustic modem.

V. CONCLUSIONS

The ALIVE project is one of “next generation” development projects underway around the AUV community. As AUV vehicle technology (for example, thrusters, guidance, batteries etc.) matures, there is a growing demand for “intelligent” payloads which can emulate functions previously done by human pilots - or perform functions previously impossible without embedded control systems.

An essential step along the road to autonomous intervention is automatic docking - the ability to automatically identify a subsea structure, transit the vehicle to a point very close to the structure, and reliably dock. ALIVE is an important exemplar of real AUV designed to operate in this complex environment.

The ALIVE AUV is currently undergoing harbour and shallow water trials, with final deep water sea trial due in October 2003.

ACKNOWLEDGEMENTS

We would like to thank and acknowledge the important contribution all our partners have made to the project. The ALIVE research programme is partially supported by the European Union (under Framework-V GRD1-2000-25741).

Finally we would like to thank all the people, too numerous to mention, at the Ocean Systems Laboratory who have contributed in any way to the research and engineering developments described here.

REFERENCES

[1] J.C. Evans, K.M. Keller, J.S. Smith, P. Marty, V. Rigaud, “Docking techniques and evaluation trials of the SWIMMER AUV: an autonomous deployment AUV for work-class ROVs”, OCEANS 2001, MTS/IEEE Conference , Volume: 1 , 5-8 Nov. 2001, pp 520 -528.

[2] D.G. Lowe, “Fitting parameterized 3-d models to images.”, IEEE Trans. on Pattern Analysis and Machine Intelligence, 13(5):441–450, 5 1991.

[3] H. Ara´ujo, R.L. Carceroni, .M. Brown, “A fully projective formulation to improve the accuracy of Lowe’s pose-estimation algorithm. Computer Vision and Image Understanding”, CVIU, 70(2):227–238, 1998.

[4] S. Russel, P. Norvig, “Artificial Intelligence: A Modern Approach”, Prentice Hall Inc., 1995.

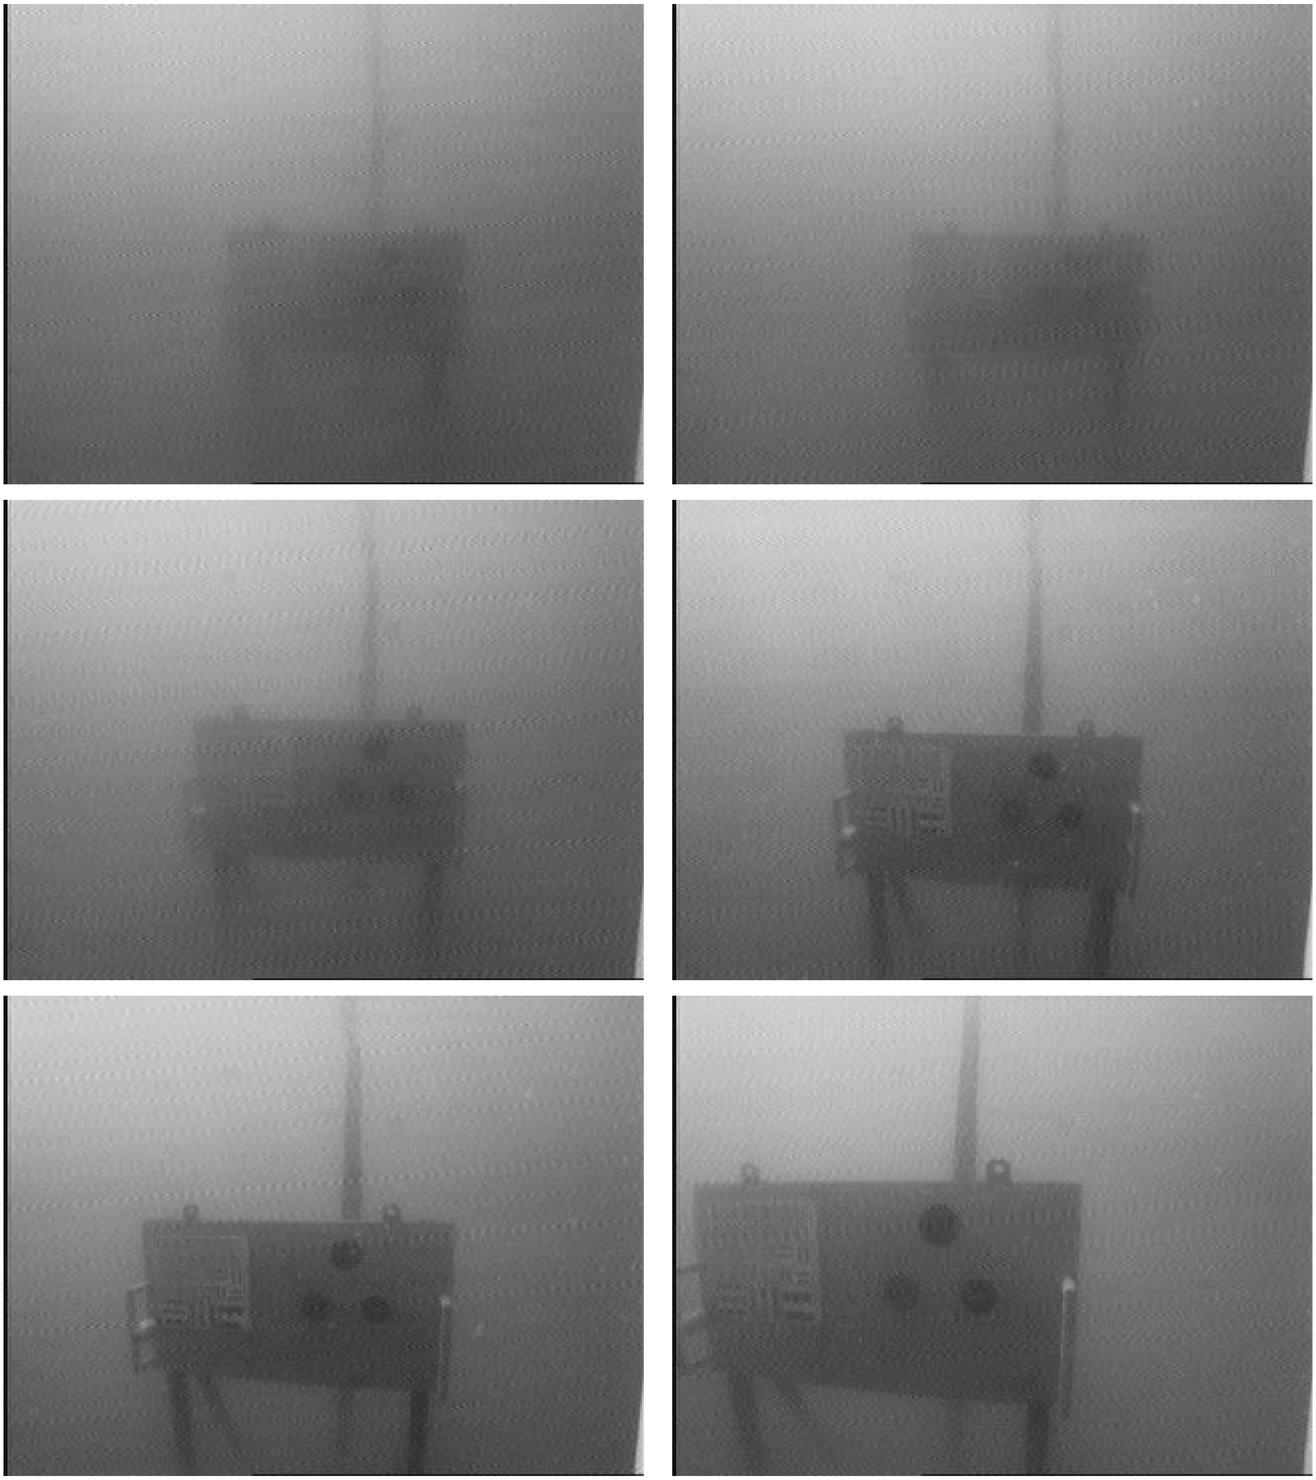

Fig. 7. Sequence of underwater video images capture during autonomous docking.

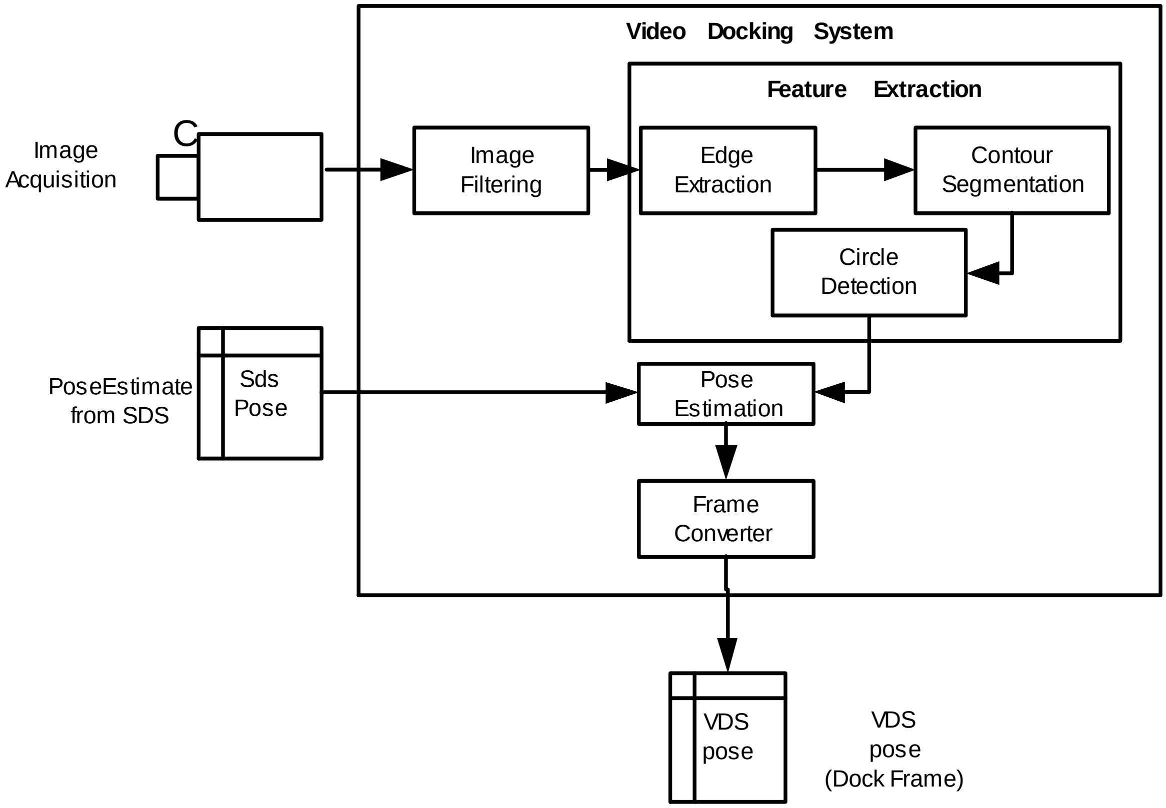

Fig. 8. VDS (Video Docking System) Internal Architecture.

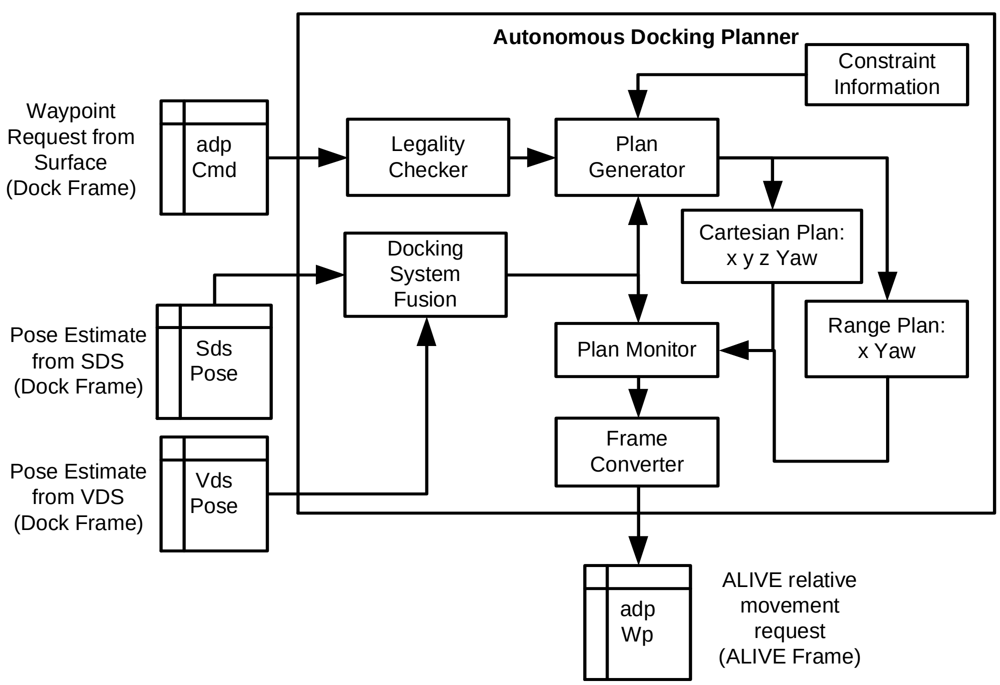

Fig. 9. ADP (Autonomous Docking Planner) Internal Architecture.

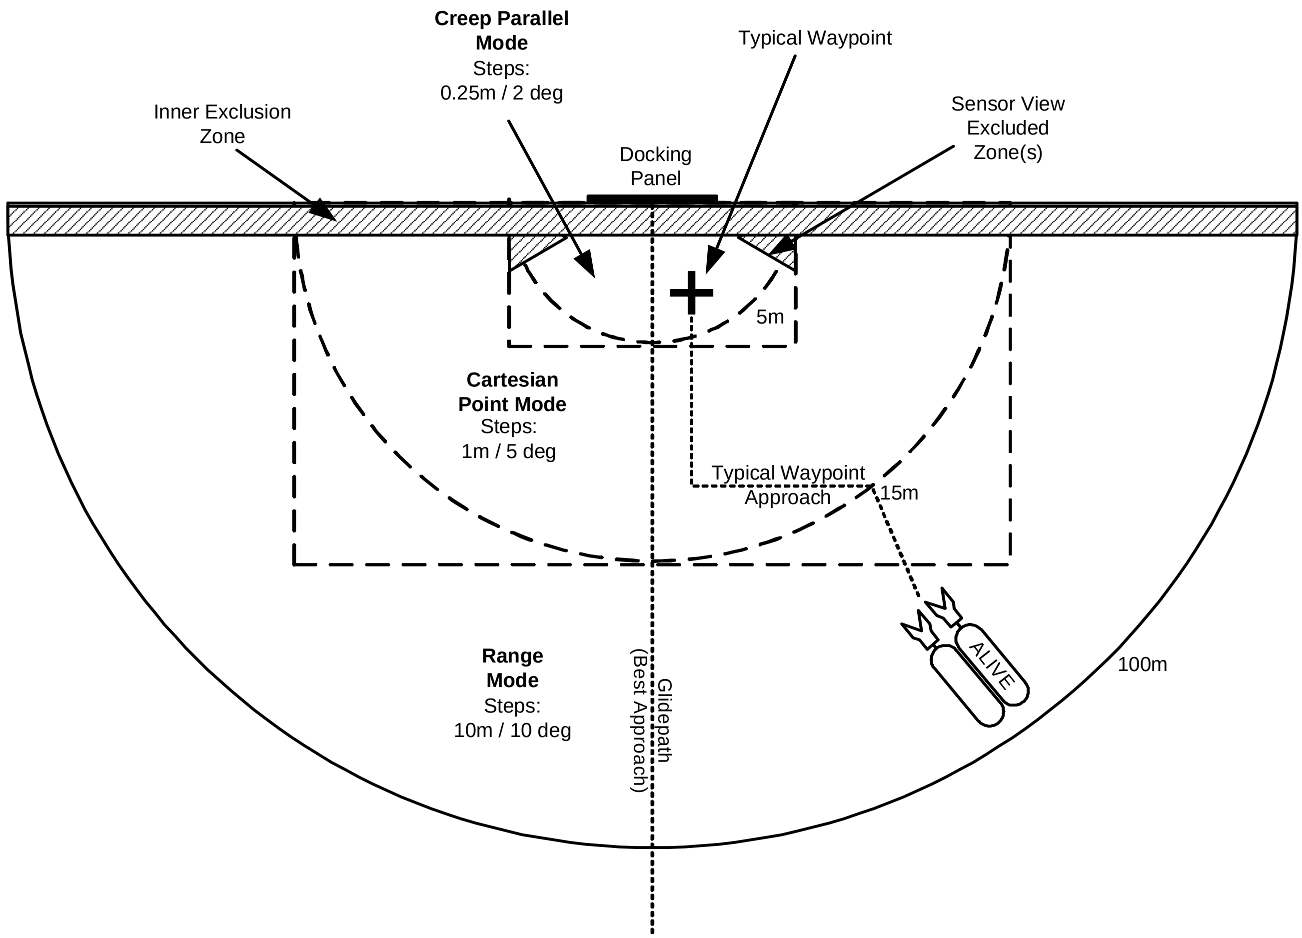

Fig. 10. Operating modes of the ADP within the SLZ (Safe Landing Zone).

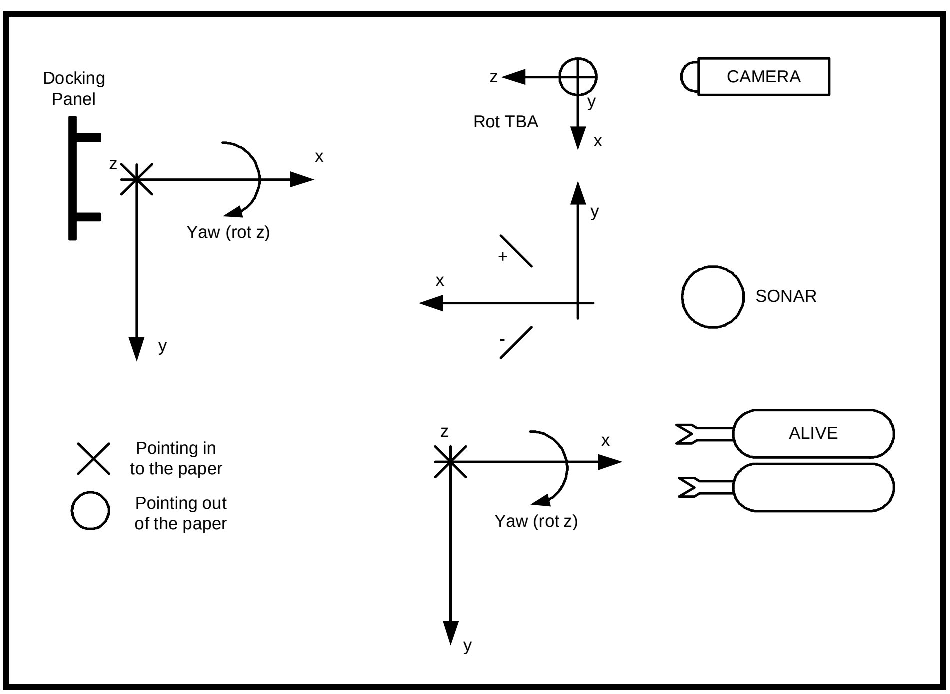

Fig. 11. Various co-ordinates frames used within the ADP.