The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

A Robotic Concept for Remote Maintenance Operations: A Robust 3D Object Detection and Pose Estimation Method and a Novel Robot Tool

Aksel A. Transeth, Øystein Skotheim, Henrik Schumann-Olsen,
Gorm Johansen, Jens Thielemann and Erik Kyrkjebø

Abstract— **Future normally-unmanned oil platforms offer potentially significantly lower commissioning and operation costs than their current manned counterparts. The ability to initiate and perform remote inspection and maintenance (I&M) operations is crucial for maintaining such platforms. This paper presents a system solution, including key components such as a 3D robot vision system, a robot tool and a control architecture for remote I&M operations on processes similar to those on topside oil platforms. In particular, a case study on how to automatically replace a battery in a wireless process sensor is investigated. A novel robot tool for removing and re-attaching the sensor lid has been designed. Moreover, a robot control architecture for remote control of industrial-type robot manipulators is presented. A 3D robot vision system for localizing the sensor lid and the battery has been developed. The system utilizes structured light, using an off-the-shelf projector and a standard machine vision camera. A novel, robust and fast vision algorithm called 3D-MaMa has been adapted to work for object localization and pose estimation in complex scenes, in our case the process equipment in our lab facility. Experimental results from our lab facility are presented which describe a series of battery replacement operations for various unknown positions of the wireless sensor, and we report on accuracies and success ratios. The experiments demonstrate that the described vision system is able to recover the full pose and orientation of an object, and that the results are directly applicable for controlling advanced robot contact operations. Moreover, the custom-built lid operation tool demonstrates successful results.**

I. INTRODUCTION

Offshore oil and gas platforms are remote and isolated places, and pose a challenging environment for their human operators due to the unsheltered maritime environment, heavy weather and unfriendly, often explosive, toxic and corrosive atmosphere [1]. Normally-unmanned automated topside platforms may be an alternative to subsea installations through increased accessibility for large maintenance operations. In addition, topside platforms may statistically recover up to 22 percent more of the oil or gas in a reservoir than a subsea alternative [2], [3]. Costs related to sound insulation, footbridges and hand-rails may be significantly reduced since there is less need for personnel.

Robot solutions for offshore inspection and maintenance (I&M) tasks have traditionally been designed for subsea tasks and performed by Remotely Operated Vehicles (ROVs) or Autonomous Underwater Vehicles (AUVs). More recently, the idea of using mobile service robots in topside offshore applications have been introduced [4], and results on autonomous navigation in such unstructured environments have been presented [1], [5], [6].

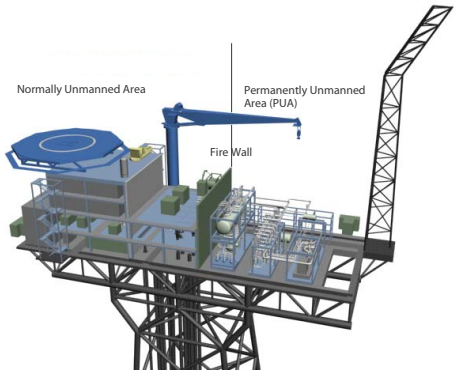

Fig. 1. The Mesa Verde Platform concept.

Similarities to remote operated I&M operations with robots offshore can be found in remote control of operations for, e.g., teleassisted surgery [7], subsea [8], in space [9], and in remote intervention tasks in nuclear facilities [10], [11]. However, there are also noticeable differences. I&M operations on oil platforms often involve manipulation of heavier objects than in other teleoperation scenarios and this requires, e.g., industrial robot manipulators. Such robots traditionally operate as preprogrammed machines and allow for very little online control. Thus, new control modes and communication infrastructure must be developed in order to facilitate remote I&M on offshore oil platforms.

Object and pose estimation of known objects in an unstructured scene is important for robot interaction applications. Existing methods in this field can be grouped into coarse and fine registration methods [12]. The goal of coarse registration methods is to compute an initial estimate of the rigid motion of the object, which eventually can be refined by more precise methods later. These later methods – like ICP [13] – are typically slow and show bad convergence properties, requiring an initial step for object detection. Methods for coarse registration include PCA, local feature based approaches like harmonic shape contexts [14], spin images [15], RANSAC-based approaches like DARCES [16] and genetic algorithms [17]. A drawback with many of the popular algorithms (like spin images) is that they require the presence of local features, i.e. sufficient local variations in geometry, in order to generate robust, local shape de-

This work was supported by Statoil and the Next Generation Robotics project (Norwegian Research Council project no. 193265/i40).

Ø. Skotheim, H. S.-Olsen and J. Thielemann work at the department of Optical measurement systems and data analysis, SINTEF ICT, N-0314 Oslo, Norway {osk,hso,jtt}@sintef.no.

A. Transeth, E. Kyrkjebø and G. Johansen work at the department of Applied Cybernetics, SINTEF ICT, N-7465 Trondheim, Norway {Aksel.A.Transeth,Gorm.Johansen,Erik.Kyrkjebo}@sintef.no

scriptors. Our approach is motivated by the fact that many industrial parts typically consist of large planar or evenly-curved surfaces, resulting in point signatures that can not discriminate well among different object poses. Instead of using local point signatures, our method is based on the search for two points a given distance apart, and with a given correspondence between their surface normals.

Salvi et al [12] found the DARCES method [16] to be the most robust method for range image registration, and DARCES is also the closest published method to the one we are using in this paper. The most important difference between DARCES and the presented method is that DARCES is based on the presence of three points, while our method only requires two points. This makes our method less computationally expensive and less vulnerable to occlusions. However, our method requires the presence of precalculated normals. See [18] for a more comprehensive review of 3D-MaMa and differences and similarities with DARCES.

In this paper we present a 3D vision system, a robot tool, and a robot control architecture employed to automatically perform robot contact operations on objects with unknown positions in unstructured environments. We present and demonstrate a *system* solution for an important challenge within remote maintenance on offshore oil platforms. This system including its key components (i.e., a 3D vision system, a robot tool, and a robot control architecture) constitute the major contribution of this paper. A case study maintenance operation involving replacing a battery in a wireless sensor is chosen and presented in this paper. Both the battery lid and the battery are detected and localized using our robust 3D object and pose detection algorithm which we call 3D-MaMa [18]. The algorithm has been adapted and the accuracy has been improved by adding a fine alignment step using the ICP algorithm. A custom designed novel tool for performing the challenging task of automatically removing and re-attaching the sensor lid has been developed. This tool is the first of its kind and can be employed for a large range of sensors. A robot control and supervision architecture for remote control of I&M operations is also presented. The proposed architecture enables an onshore operator to perform and monitor I&M operations offshore. The system for remote maintenance is demonstrated through battery replacement operations in a lab facility which resembles an unstructured industrial environment and experimental results are presented.

There is an ongoing research activity on taking the presented operations and tools offshore and this is thus beyond the scope of this paper. Still, the general approaches (e.g., the vision methods) are also applicable for other physical instruments than the ones used in this paper and can therefore also be applicable for outdoor conditions, but special care has to be taken regarding environmental disturbances, which has not been addressed in this paper.

This paper is organized as follows. In Section II we give an introduction to the concept of normally-unmanned oil platforms and an overview of a lab facility used to demonstrate this concept. Moreover, a vision system used in the lab facility is described. In Section III a description of the vision methods and robot tool presented in this paper is given. Experimental results are presented and discussed in Section IV and conclusions are stated in Section V.

II. SYSTEM OVERVIEW

In this section a short introduction to the concept of remote I&M on future normally-unmanned offshore oil platforms is given. Moreover, a lab facility is presented which is used to implement and demonstrate relevant I&M operations for such oil platforms. In particular, we detail a structured light-based vision system used in the lab to, e.g., localize objects in order to perform robot contact operations on them.

A. Background and Challenges of Remote Operations

A novel remote I&M concept for offshore oil and gas platforms was presented in [19] as an alternative to traditional offshore platforms. The platform concept is based on separating the work area accessible by human operators, and a closed permanently unmanned area (PUA) that is only serviced by robots as in Fig. 1.

The production process in the PUA is not intended to be operated by robots, but will be built on concepts developed for subsea production platforms. Its topside location will allow for easier access during I&M operations. The remotely operated platform concept is designed on the premise that robots may replace humans for the most important scheduled I&M operations inside the PUA such as gauge readings, valve and lever operations and monitoring gas levels, leakages, acoustic anomalies and surface conditions [1].

Remote offshore I&M operations pose many significant challenges, e.g.: 1) Operations must be performed on a structure subject to wind, dirt, ice and structural changes. 2) No-one is offshore to intervene if something fails. Hence, a *robust* and *adaptable* I&M system is necessary. 3) Onshore operators must be able to monitor and control I&M operations with a large range of level of detail. This requires new and versatile operator interfaces.

B. Lab Facility

A lab facility has been built in Trondheim, Norway, in order to develop, test, and demonstrate solutions for next-generation I&M operations for normally-unmanned oil platforms. An overview of the facility is given in the following.

The lab facility consists of a process structure simulating parts of a production process on a real oil-platform and

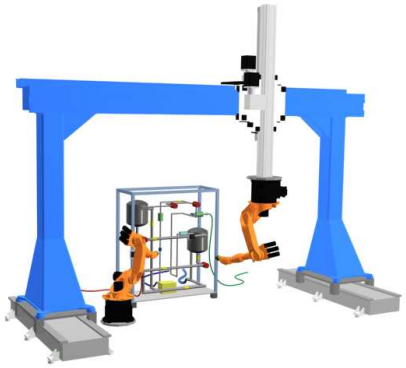

Fig. 2. Robots and process equipment in lab facility.

two robot manipulators used for I&M tasks on the process structure using available tools and sensors. See Fig. 2 for a 3D model of parts of the lab facility. All necessary information (video feeds, audio, sensor data, etc.) is relayed to human operators via the Internet.

Both robots are standard 6-axes manipulators (Kuka KR-16). One is mounted on a 3-axes Güdel gantry. The main tasks of the gantry-mounted robot (GR) is to perform I&M operations on the process equipment. This robot can connect automatically to custom-built tools and sensors such as vibration-measurement sensors, a valve-operating tool, and a lid operation tool (see Section III-B). A structured light system is installed at the top of the (black) base of the GR (see Section II-C). The floor-mounted robot (FR) is used for monitoring and assisting the GR.

The process equipment is used to simulate an offshore production process. It consists of water tanks, valves to control the water flow, and various sensors to validate the measurement readings performed by the robots. In particular, Rosemount wireless pressure and temperature sensors from Emerson are installed. A battery replacement operation on such sensors is performed automatically by the GR and the FR and this is detailed throughout this paper.

The lab facility can be remote controlled from any location via the Internet. Live video streams and continuously updated 3D models of the facility provide a remote operator with awareness of the lab operations. The remote operator can initiate high-level commands for automatic I&M routines from a graphical representation of the process equipments, or control the robots with off-the-shelf joysticks either directly or via 3D models. A model-based collisions detection system is used to ensure safe operations. See [19] for further details on the lab facility.

C. Vision System

Our structured light based 3D vision system is used for creating a 3D model of the process equipment in the lab. This is advantageous for close contact operation in a changing environment. Structured light is a relatively fast, accurate and flexible optical method for measuring 3D shapes of objects. The result is a dense cloud of points that accurately describes the shape of the illuminated surface. The custom made structured light software in our lab works with non-expensive, off-the-shelf components, in this case consisting of a BenQ multimedia projector and a Basler Scout Gbit Ethernet camera with 1280×960 resolution (see Fig. 3).

The structured light algorithms are based on a combination of Gray code and phase stepping fringe projection [20]. A measurement is performed in about 5 seconds, and the result is a dense point cloud consisting of approximately 1.2 million 3D points with sub-millimeter accuracy. Due to limited depth of focus in the camera and the projector, the measurement range for our set-up was limited to around 1 m in the z-direction.

Structured light was selected for convenience and works well in indoor testing environments. For offshore environments, however, structured light may not be the optimal technology due to strong ambient lighting and limited robustness and lifetime of the projector system. However, our data analysis algorithms (see Section III-C) work with generic point clouds, so other commercially available 3D sensors (e.g. laser scanners) may be used as well. Addressing additional outdoor challenges specific for an offshore environment will then of course also be important, but this is beyond the scope of this paper.

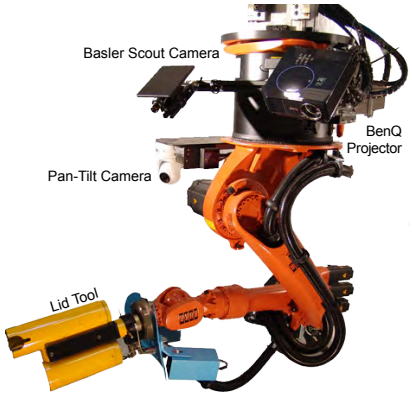

Fig. 3. KUKA robot manipulator hanging upside down from the Güdel gantry. The structured light vision system is mounted on the base and the sensor lid operation tool is mounted on the end effector.

III. METHODS AND TOOLS

In this section a control and supervision architecture for remote control of I&M operations is presented, in addition to a specialized tool – called a lid operation tool – to be used for battery replacement operations on wireless sensors. Moreover, we present a novel, robust and fast 3D object detection and pose estimation algorithm used for detecting almost arbitrary objects in complex scenes.

A. Robot Control Architecture and Sensor Integration

In the following we outline the control architecture and a sensor integration system as a step toward a complete system for remote control and monitoring of I&M operations on normally-unmanned offshore oil platforms. Both the control architecture and the sensor integration system are implemented in the lab facility described in Section II-B. Fig. 4 shows an overview of which modules that are placed "offshore" (i.e., in our lab facility) and "onshore" (i.e., at a remote location with an Internet connection).

An onshore operator may initiate and monitor I&M operations from a user interface called *Process Viewer*. The user interface gathers available high-level robot commands (e.g., "Measure vibration" or "Change sensor battery") from a Structured Query Language (SQL) database on start-up. Each command has a corresponding robot motion routine which is stored and carried out offshore. The operator may point-and-click on an illustration of the offshore process equipment displayed in Process Viewer in order to initiate I&M operations. Process Viewer also displays sensor data readings obtained during I&M tasks.

All robot commands and sensor data are relayed through a central server called *Robot Server*. This server transfers I&M commands from an operator to the robots via TCP/IP and an open interface standard called OPC. Moreover, Robot Server

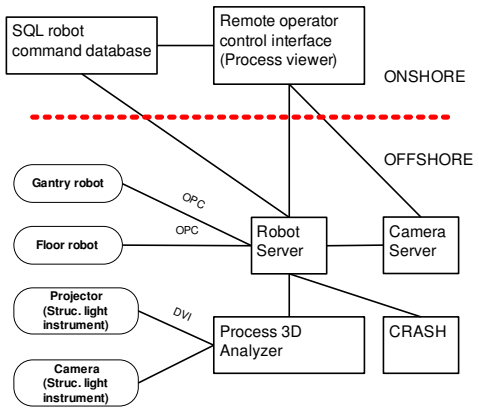

Fig. 4. Control architecture for remote inspection and maintenance. All transmission lines that are not marked use TCP/IP.

also acts as a sensor integration system in that all robot data such as tool states, robot-mounted sensor data and process-mounted sensor data are transferred to Robot Server and relayed to, e.g., Process Viewer. Robot Server communicates with a custom-designed program called *CRASH* used for model-based collision checking. The CRASH system will be important in order to avoid damage to the robots and is located offshore to increase speed and reliability.

Object detection, pose estimation and control of the structured light instrument are carried out by the *Process 3D Analyzer*. This module is located offshore to decrease the necessary amount of data traffic (e.g., high-resolution images) between offshore and onshore. The Process 3D Analyzer transfers object coordinates via Robot Server to the robots in order for them to perform, e.g., a lid operation or battery-replacement task.

Camera Server distributes live images from the offshore installation to an onshore control room. The frame rate of the live video feed transferred from Camera Server is automatically adjusted based on the available bandwidth.

B. Robot Tool Design and Operation

Wireless sensors offer a cost-effective, simple and versatile approach to process monitoring. In particular, there is a growing interest for such sensors in the offshore industry due to, e.g., reduced system weight (i.e., less cables, connection boxes, etc). Our lab facility is equipped with Rosemount wireless sensors from Emerson which are typical wireless sensor in use by the offshore industry. In this section we present a novel robot tool – called a lid operation tool – for removing and re-attaching the battery lid on Rosemount wireless sensors. In addition, parts of the operation required to replace a sensor battery which is located underneath the sensor lid will be described. The complete battery replacement operation is described in Section III-B.

Rosemount sensors are intrinsically safe and have a battery lid with a diameter of 85 mm. The lid is externally threaded and has a rubber gasket. Moreover, the top of the lid has a circular set of concavities. A battery is located behind the lid and can be removed by pulling it straight out. A new battery can then pushed straight in.

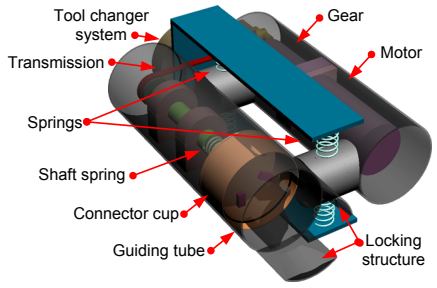

Fig. 5. CAD model of the lid operation tool.

Three main requirements regarding the battery replacement operation were defined before designing the lid operation tool: 1) The Rosemount sensors shall not be modified in any way. 2) Profibus shall be used for communication with the lid operation tool. 3) Two robots can be used for the task. The main tool challenge for this operation was how to remove and re-attach the sensor lid. Several different solutions were considered until a final design was ready.

The lid operation tool is used by the gantry-mounted robot in our lab facility in order to unscrew, keep and re-attach the lid. The floor-mounted robot replaces the battery using a standard gripper tool.

The main concept for the lid operation tool (see Fig. 5 for a CAD-model) is based on the use of spring-loaded connections together with a locking structure, a guiding tube, and a connector cup with a circular arrangement of convexities. The cup and its convexities are fitted to the shape of the sensor lid which has a corresponding circular set of cavities on the top. A lid removal procedure is initiated by the gantry-mounted robot pressing the lid operation tool in a longitudinal movement against the sensor. Then the springs and the guiding tube ensure that the lid operation tool is situated at a suitable angle relative to the sensor lid. The shaft spring in the guiding tube ensures that the connector cup and its convexities grip into cavities on the sensor lid once the connector cup is rotated. The shaft spring allows the lid to be screwed in and out without having to move the robot during the operation.

The tool is placed in a correct position and used for removing and re-attaching the sensor lid by control of the speed and torque of the motor via Profibus. While maximum torque is used to unscrew the lid, a reduced torque is used to re-attach the lid. The position feedback from the motor is used to count the number of rotations. The number of rotations and the torque are used in an error detection and retrial algorithm.

The main challenge is to re-attach the lid. The challenge partly arises from that the pitch (the distance from the crest of one thread to the next) is very low (i.e., fine threads). When the tool is in position, then the operation starts with turning the lid counter clockwise (off) in order to get the lid correctly seated on the threads of the sensor. Then it is turned clockwise (on) with a suitable torque. If the lid gets stuck before a minimum number of clockwise rotations have been completed, then the re-attaching operation has temporarily failed. However, in order to achieve a successful

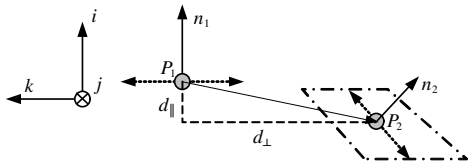

Fig. 6. Principle sketch for search for candidate planes.

re-attachment, the lid is automatically turned counter clockwise (off) again and then turned clockwise (on) once more. The springs in the tool design are crucial for fine-positioning the tool relative to the sensor when re-attaching the lid.

C. 3D Object Localization and Pose Estimation

The 3D-MaMa algorithm introduced in [18] is a method for estimating the position and orientation of one or more known 3D objects in a scene. The algorithm was originally developed by the authors for random bin picking, where the goal was to pick multiple, identical objects, each with an arbitrary position and orientation, out of a container. In this paper, we describe how the same algorithm can be applied in order to search for specific objects (in our case a sensor lid and a battery) in a point cloud of a complex scene containing process equipment for an offshore oil platform. The accuracy of the position and pose estimation is improved by subsequent application of the Iterative Closest Point [13] (ICP) algorithm for fine alignment.

The main design goal for the 3D-MaMa algorithm was that it should be suitable for industrial parts with little local texture and be robust towards missing data points. Both of these criteria make current algorithms like spin images [15], where local surface patches are used to recognize subparts of the object, less suitable for the objects we have been working on.

The input for the 3D-MaMa algorithm is a template in the form of a CAD model or a 3D mesh of the object to search for, along with a set of two points with two corresponding surface normals (hereafter referred to as the search parameters) on the surface of the template. These points can be selected either automatically or by the user, and this needs only to be done when the system is trained to handle a new object. We refer to the first selected point-normal pair, $(p_r, \hat{n}_r)$, as the *reference plane*, while the second point-normal pair, $(p_s, \hat{n}_s)$, is termed the *search plane*. The reference and the search plane should preferably be two distinct pairs of points and surface normals that are somehow characteristic for the object. Note that we use the term plane to refer to a local plane or manifold defined on a small neighborhood around a point on the 3D model, and that the reference plane and the search plane may be located on e.g. a convex or a concave surface, not necessarily on a larger, planar surface.

The first step of the algorithm is to estimate surface normals for all the points in the input point cloud. This is done by performing local plane fits to small neighborhoods around each point in the input 3D image (typically by considering a neighborhood of 5×5 points).

The second step of the algorithm is to search for candidate planes, i.e. pairs of points with a similar distance, and with corresponding surface normals having a similar relative orientation to each other as the specified reference plane and search plane.

Each plane, $(P_i, \hat{n}_i)$, in the input 3D image is checked against every other plane, $(P_j, \hat{n}_j)$, and a score is calculated by comparing their center-to-center distances and the scalar product of their surface normals. This initial matching process is illustrated in Fig. 6, where two local planes $(P_1, \hat{n_1})$ and $(P_2, \hat{n_2})$ are being evaluated. In order to tolerate a certain amount of displacement, the vector $\vec{d} = P_1P_2$ is checked whether it points into the parallelogram drawn in Fig. 6. The scalar product of the surface normals, $\hat{n}_1 \cdot \hat{n}_2$, is also checked whether it is close to the scalar product $\hat{n}_r \cdot \hat{n}_s$ of the search parameters.

If the preceding tests are satisfied, it is possible to calculate a transform, T, that represents the best possible alignment of the reference and the search plane of the template with the candidate planes $(P_1, \hat{n_1})$ and $(P_2, \hat{n_2})$. This transform is marked as a candidate for further evaluation.

In the final step of the algorithm, a cost that represents how well the transformed template points match with the scene is calculated (see [18] for details). Along with the cost for a particular position and pose, we also calculate the number of scene points, N_s, that fall within the transformed template, and a coverage ratio, $R = N_s/N_{max}$, where N_{max} is the maximum number of scene points that could be sufficiently close to the transformed template given a perfect match.

In the case of random bin picking [18], a threshold was set based on the cost and the coverage ratio. In our case, we are only interested in the one transform that provides us with the best match between the template and the scene points, hence we only keep the transform T that provides the minimum cost and with a sufficiently high coverage ratio.

The alignment using 3D-MaMa is usually accurate down to a few millimeters, which was not sufficient for the robotic operation of the battery. In order to obtain the best possible alignment of the template with the scene, we added a fine positioning step to the algorithm using the ICP algorithm, which is a standard algorithm for fine alignment of point clouds. We use a bounding box based on the result of registration with 3D-MaMa to select only those points from the scene that are believed to have a corresponding point on the CAD model. Note that the ICP algorithm needs a very good starting point in order to end up in the global minimum, providing us with the position and pose of the template that minimizes the mean-square error of distances between the points in the scene and their closest point on the template. If the starting point is too far off, the ICP algorithm might get stuck in a local minimum.

As described in [18], the registration algorithms have been highly optimized and parallelized, and position and pose estimation is usually performed within a few hundred milliseconds.

IV. EXPERIMENTS AND RESULTS

In this section the steps of action in the selected maintenance case are detailed. In addition, three sets of experiments

are presented to validate the operation and the results are discussed.

A. Case Study and Process Flow Description

The lab facility consists of, as described in Section II-B, two robots; a floor-mounted robot (FR) and a gantry-mounted robot (GR). These robots co-operate using information obtained from the structured light instrument in order to replace a battery in a Rosemount wireless sensor.

The process flow in the case study is as follows: 1) The battery-replacement operation is initiated by an operator who selects this operation in Process Viewer. This is the only necessary user interaction. 2) The GR moves the structured light instrument (SLI) in front of the wireless sensor and signals the Process 3D Analyzer (P3DA) that that the SLI is in position. 3) The P3DA uses the SLI for finding the position and orientation of the sensor lid. 4) The position of the sensor is transmitted from the P3DA to the GR. 5) The GR connects to the custom-built lid operation tool and unscrews the sensor lid if the sensor lid has been successfully localized (if not, then the operation is aborted). 6) The GR moves the SLI back in front of the sensor and signals the P3DA that it is in position. 7) The P3DA uses the SLI in order to find the position and orientation of the sensor battery which is no longer hidden by the battery lid. 8) The coordinates of the lid are transmitted to the FR. 9) The FR connects to a two-finger Schunk gripper and removes the battery based on the received coordinates if the battery is successfully localized (if not, then the lid is re-attached and the operation is aborted). The FR then picks up a new battery and inserts it into the sensor. 10) The GR uses the lid operation tool in order to re-attach the sensor lid. A short video of parts of the complete operation is available together with this paper.

All communication between the robots and the P3DA goes via Robot Server (see Fig. 4). All robot motions are pre-programmed with a teach pendant based on an initial wireless sensor position. Feedback from the P3DA is employed in order to update the robot paths when the wireless sensor has moved from its original position.

In order to detect both the sensor lid and the battery two 3D models were used as input to the vision algorithm (see Fig. 7). Since we did not have access to CAD models of the sensor lid and the battery, we obtained their 3D models by 3D digitization with a commercial ATOS structured light sensor by GOM GmbH.

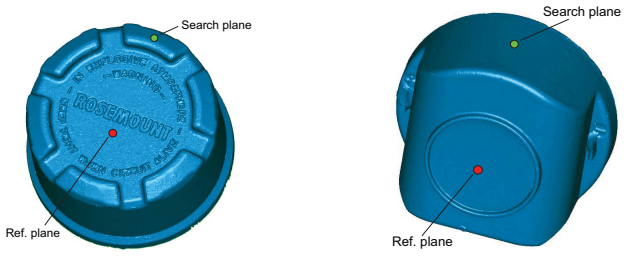

Fig. 7. 3D models used to detect position and pose for sensor lid (left) and battery (right) with indicated reference plane (red) and search plane (green) for the 3D-MaMa algorithm.

B. Experimental Results

In order to test our vision algorithm and the lid operation tool's suitability for robotic maintenance operations, three experiments have been conducted.

The first experiment was performed in order to verify the repeatability of the vision algorithms. A series of ten 3D images was captured of the sensor lid, and the result of the position and pose estimation along with mean values and standard deviations are given in Table I. The columns T_x, T_y and T_z display the detected center point of the sensor lid (in millimeters), while the columns γ, β and α display the Euler angles (in degrees) corresponding to rotations around the z, y and x axes (ZYX convention). The RMS column shows the RMS distances in millimeters after ICP alignment between the points that were considered in the scene and their closest corresponding point on the 3D template. The point density for the input 3D templates was around 0.3 mm. This implies that the RMS error is not expected to drop significantly below this value, even when we have a perfect alignment between the template and the measured points. The final column displays the coverage ratio R, which was calculated as explained in Section III-C.

TABLE I
REPEATABILITY EXPERIMENT FOR LID

#	T_x	T_y	T_z	α	β	γ	RMS	R
1	90.9	-32.8	-161.3	0.0	-0.2	48.1	0.41	0.56
2	90.9	-32.7	-161.3	0.3	0.2	167.9	0.39	0.54
3	90.9	-32.7	-161.3	0.2	0.3	167.4	0.40	0.56
4	90.8	-32.8	-161.4	0.2	0.2	167.2	0.39	0.53
5	90.8	-32.8	-161.4	0.5	0.1	-132.5	0.38	0.54
6	91.0	-32.8	-161.4	0.2	0.2	167.3	0.42	0.54
7	90.9	-32.8	-161.4	0.2	0.1	168.5	0.40	0.55
8	90.7	-32.8	-161.4	0.5	0.1	-132.6	0.38	0.53
9	90.7	-32.8	-161.4	0.6	0.2	-132.3	0.39	0.50
10	90.8	-32.8	-161.4	0.2	0.2	168.2	0.41	0.53
μ	90.8	-32.8	-161.4	0.3	0.1	65.7	0.40	0.54
σ	0.10	0.04	0.05	0.19	0.13	141.66	0.01	0.02

The results in Table I suggest that the vision algorithm is very repeatable. The detected position of the center of the sensor lid has a standard deviation of less than a 0.1 mm for the ten trials. The standard deviation in the Euler angles is less than 0.2 degrees for rotations around the x and y axes. The standard deviation in the detected rotation around the z axis (γ) is very high for the battery lid. This is due to the rotational symmetry of the lid, something that makes it impossible to determine the exact rotation around its symmetry axis. This is, however, not necessary in order for the lid operation tool to unscrew the sensor lid.

In the second experiment, the complete battery replacement operation as described in Section IV-A was evaluated. A movable sensor holder was used in order to steadily support the sensor in different positions (by sliding it up and down along a metal bar on the process equipment). In all the ten trials of this experiment, the robot was able to correctly detect and unscrew the sensor lid, detect and replace the battery and re-attach the sensor lid.

As can be seen in Table III, the angle γ for the battery also varies significantly, even though the battery is only partly symmetric. We have seen that this has been due to a lack of measured points on the bottom and on the sides of the battery (due to our rigid sensor setup on the robot base). The

TABLE II

FULL OPERATION – RESULTS FOR SENSOR LID

#	T_x	T_y	T_z	α	β	γ	RMS	R
1	91.3	-32.5	-161.6	0.1	0.2	160.0	0.40	0.53
2	90.9	-19.8	-161.7	0.2	0.1	-147.9	0.35	0.54
3	90.4	82.7	-162.0	0.2	-0.1	-180.0	0.43	0.58
4	90.1	69.1	-161.9	0.5	0.1	-114.4	0.37	0.57
5	90.3	56.0	-162.1	0.3	-0.1	-150.2	0.41	0.55
6	90.5	40.6	-161.6	-0.0	-0.1	94.9	0.40	0.55
7	90.7	26.7	-161.8	0.3	0.1	158.0	0.38	0.52
8	90.6	16.0	-161.8	-0.0	0.2	117.2	0.35	0.51
9	90.9	3.1	-161.7	0.6	-0.1	-136.4	0.45	0.56
10	90.9	-11.4	-161.6	0.3	-0.1	39.9	0.44	0.52

TABLE III

FULL OPERATION – RESULTS FOR BATTERY

#	T_x	T_y	T_z	α	β	γ	RMS	R
1	90.8	-32.1	-225.6	1.7	0.2	-2.0	0.57	0.49
2	90.5	-19.1	-225.7	2.2	0.6	5.1	0.53	0.52
3	90.5	83.3	-226.1	1.8	-0.0	2.4	0.46	0.52
4	90.2	69.6	-226.1	2.2	0.7	9.3	0.76	0.45
5	90.5	56.6	-226.0	2.2	-0.5	4.2	0.54	0.49
6	90.0	41.3	-226.0	2.0	0.7	-5.3	0.75	0.54
7	90.6	27.4	-225.9	1.9	-0.1	1.2	0.60	0.50
8	90.2	16.6	-225.9	2.0	0.9	-2.6	0.55	0.55
9	90.7	3.5	-225.8	2.0	-0.4	8.0	0.71	0.53
10	90.4	-10.7	-225.8	2.3	0.4	6.2	0.66	0.52

remaining points (on the front and on the top of battery) fit well with the template even when it is rotated around the z axis. The detected center position, however, seems to be less affected by this rotational symmetry, and the estimated position of the battery was good enough for the robot to successfully replace it for every one of the 10 trials.

In addition to Table III, we also generated 3D models for presentation on the screen for each of the experiments. A screenshot of such 3D models are shown in Fig. 8 for the sensor lid and in Fig. 9 for the battery. In these figures, the measured points from the scene are shown in gray scale, the points taken into account when calculating the RMS value are shown in green and the input 3D models (used as templates for the 3D-MaMa algorithm) are annotated on the point clouds with their detected position and pose.

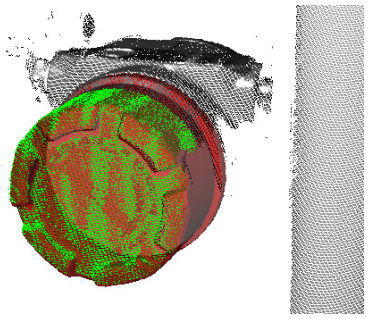

Fig. 8. Detected sensor lid on Rosemount sensor in full operation trial number 1. The 3D model of the sensor lid is annotated in red and the corresponding scene 3D points in green.

Due to lack of degrees of freedom in the sensor holder, we were not able to change the orientation of the sensor in the second experiment. Thus a third experiment was conducted in order to evaluate the robustness to different orientations. The Rosemount sensor was then attached to the process equipment in 10 different random positions and orientations

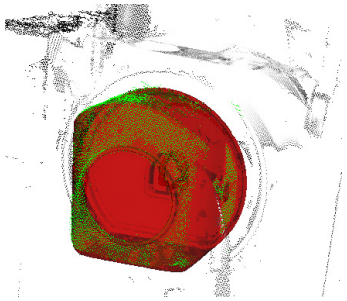

Fig. 9. Detected battery in full operation trial number 2. The 3D model of the battery is annotated in red and the corresponding scene 3D points in green.

by using plastic tie wraps, and 3D images were captured by the structured light sensor on the robot. Table IV summarizes the results of the random orientations experiments. We see that most of the experiments give an RMS error of less than 0.6 mm after alignment. That is in the same range as the first experiment and thus *probably* sufficient for the lid operation.

In trial number 5 and 10 the RMS error is a bit higher. It turns out that the ICP had converged to a local minimum giving a slightly wrong rotation around the axis passing perpendicularly through the center of the lid. In these two cases, the circular convexities of the measured sensor lid were not properly aligned with the circular convexities of the 3D model, giving rise to a bias on the RMS error. The estimated position of the center point and the remaining two degrees of freedom seem, however, to be quite good, and we believe that the lid operation tool would be successful even in this case (this has to be tested though after rebuilding the sensor holder and doing more refined experiments).

For all the 10 trials with the robot, convergence was reached by the ICP algorithm. By visual inspection of the 3D model shown on the screen, and by looking at the RMS values in Table II and III, the position and pose of the template seems to be correctly detected with average point distances of just a few tenths of a millimeter. In order to successfully replace the battery, a plastic tap on the backside of the battery has to fit in a connector socket which is less than a millimeter wider than the diameter of the tap. The fact that we successfully performed a series of 10 full operations with the robot also clearly indicates that we have achieved sub-millimeter absolute accuracy in the position and pose estimation.

Note that there is an ongoing activity on taking the presented operations and tools offshore. Since the accuracy in the position and pose estimation algorithms is expected to drop under more realistic and hazardous conditions, we are currently developing tools that are less vulnerable to small errors in the absolute positioning.

V. CONCLUSIONS AND FURTHER WORK

A system solution for remote inspection and maintenance (I&M) operations on normally-unmanned oil platforms has been presented and validated. In particular, a robot tool, a 3D vision system, and a robot control architecture have been detailed. We have demonstrated the performance and

TABLE IV

ORIENTATION AND POSE EXPERIMENT

#	T_x	T_y	T_z	α	β	γ	RMS	R
1	118.4	152.9	-172.0	-4.9	20.9	-87.0	0.36	0.59
2	93.1	138.8	-208.9	-30.7	7.7	158.7	0.40	0.51
3	112.4	48.5	-170.1	-1.2	27.0	-95.6	0.34	0.51
4	143.6	54.4	-147.3	-30.1	-9.5	-160.2	0.37	0.56
5	110.0	86.8	-217.9	-5.9	18.4	77.0	1.06	0.47
6	110.6	87.7	-238.8	-21.3	11.2	68.0	0.59	0.45
7	109.0	95.5	-261.7	-30.6	40.0	77.0	0.40	0.64
8	103.5	124.4	-250.4	-28.7	17.0	82.8	0.59	0.47
9	101.1	113.6	-238.1	-21.4	8.9	85.9	0.42	0.57
10	115.7	70.6	-207.7	6.3	-2.4	-152.0	1.63	0.58

accuracy of this system by performing an automatic maintenance operation on an industry standard Rosemount wireless sensor. The proposed system is a step toward enabling onshore operators to perform I&M operations offshore.

The maintenance operation described in this paper involves replacing a battery in the sensor and illustrates one of several maintenance operations necessary on an unmanned offshore oil platform. A novel lid operation tool for removing and re-attaching the sensor lid of the sensor has been presented and successfully tested. The casing of the sensor is employed for a large range of different sensors which suggests that this tool has a large range of applications.

To meet the stringent demands for accurate positioning during robot contact operations, a 3D vision system based on structured light is employed to acquire detailed 3D models in the form of dense point clouds of the operating environment. By basing our image analysis algorithms on 3D shape instead of 2D images, we ensure robustness to e.g. perspective, object distance, changes in the ambient lighting or in the color or appearance of surfaces, which otherwise cause problems for standard machine vision tools. A generic object localization and pose estimation algorithm called 3D-MaMa has been adapted and improved to meet the demands for accuracy in this case scenario. Experimental data show that the vision system is capable of localizing both the sensor lid and battery with 6 degrees of freedom with sub-millimeter accuracy, and that the output of the vision algorithm is directly applicable for controlling a robot.

In our experiments the point clouds are obtained by 3D imaging with structured light, but the algorithm is applicable for other 3D acquisition methods that provide dense 3D data as well, such as laser triangulation. Since 3D models (a CAD model or a 3D scan) is used as input for the object localization and pose estimation algorithms, the vision system becomes very versatile, and it is easily trained to search for new objects. For some objects, however, such as completely flat objects or objects with very complex texture, our search strategy based on two points with corresponding surface normals may not be ideal.

There are still many challenges to face for remote I&M operations offshore, such as how to enable robots and tools to operate in the harsh weather conditions of an offshore oil platforms. These challenges are topics for further research. In addition, we continue with 3D robot vision and further work will focus on increasing the accuracy and robustness of our object detection and pose estimation algorithms and to make them applicable for even more complex and noisy outdoor scenes. We are also working on algorithms for determining deviations between a CAD model and a measured point cloud of a scene, and to take advantage of this information for navigation and robot guidance.

Reduced commissioning and operation costs, together with improved Environmental, Health and Safety (EHS) are some of the potential benefits of having normally-unmanned topside oil platforms. However, such oil platforms require advanced methods and tools for remote control and monitoring of I&M operations. In this paper, we have presented a first step toward a complete system for such operations.

VI. ACKNOWLEDGMENTS

The authors thank Dr. Arne Ulrik Bindingsbø and Anders Røyrøy at Statoil for contributing with ideas and insight, Terje Mugaas, Sigurd A. Fjerdingen, and Jens Olav Nygaard at SINTEF ICT for developing many of the control software components in the system, and Dr. Geir Mathisen at SINTEF ICT for his contribution to the software design.

REFERENCES

[1] B. Graf and K. Pfeiffer, "Mobile robotics for offshore automation," in *Proc. IARP/EURON Workshop on Robotics for Risky Interventions and Environmental Surveillance*, Benicssim, Spain, 7-8 January 2008.

[2] SINTEF, "Robots taking over the job on offshore oil drilling platforms." ScienceDaily, January 2008, retrieved December 6, 2008.

[3] D. Rogers, "At arm's length," Engineering Magazine, February 2008.

[4] P. Liljebäck, T. Kavli, and H. Schumann-Olsen, "Robotic technologies for an unmanned platform," SINTEF, Report STF90 F05405, 2005.

[5] B. Graf, K. Pfeiffer, and H. Staab, "Mobile robots for offshore inspection and manipulation," in *Proc. Int. Petroleum Technology Conference*. Dubai, U.A.E.: SPE, December 4 - 6 2007.

[6] M. Bengel, K. Pfeiffer, B. Graf, A. Bubeck, and A. Verl, "Mobile robots for offshore inspection and manipulation," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2009, pp. 3317–3322.

[7] J. Hills and J. Jensen, "Telepresence technology in medicine: Principles and applications," *Proc. IEEE*, vol. 86, pp. 569 – 580, 1998.

[8] M. Hinchey and K. Muggeridge, "Potential for subsea robot control," *Ocean Engineering*, vol. 22, no. 2, pp. 223–234, February 1995.

[9] Y. Xu and T. Kanade, Eds., *Space Robotics: Dynamics and Control*. Springer, 1992.

[10] H. Roman, "Robots cut risks and costs in nuclear power plants," *IEEE Computer Applications in Power*, vol. 4, no. 3, pp. 11–15, July 1991.

[11] OC Robotics, "Snake-arm robots access the inaccessible." Nuclear Technology International, pp. 92–94, 2008.

[12] J. Salvi, C. Matabosch, D. Fofi, and J. Forest, "A review of recent range image registration methods with accuracy evaluation," *Image and Vision Computing*, vol. 25, pp. 578–596, May 2007.

[13] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proc. 3-D Dig. Imaging and Modeling*, 2001, pp. 145–152.

[14] T. Moeslund and J. Kirkegaard, "Pose estimation of randomly organized stator housings," in *Proc. Scand. Conf. on Image Analysis*, 2005, pp. 679–688.

[15] A. E. Johnson and M. Hebert, "Using spin images for efficient object recognition in cluttered 3d scenes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, pp. 433–449, 1999.

[16] C.-S. Chen, Y.-P. Hung, and J.-B. Cheng, "RANSAC-based DARCES: A new approach to fast automatic registration of partially overlapping range images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, pp. 1229–1234, 1999.

[17] K. Brunnstrom and A. Stoddart, "Genetic algorithms for free-form surface matching," in *Proc. Int. Conf. on Pattern Recognition*, vol. 4, Aug 1996, pp. 689–693.

[18] Ø. Skotheim, J. Thielemann, A. Berge, and A. Sommerfelt, "3D-MaMa: 3D pose estimation for random bin picking by pairwise manifold matching," in *Proc. 3D Image Processing (3DIP) and Applications 2010*. San Jose, CA, USA: SPIE, January 2010.

[19] E. Kyrkjebø, P. Liljebäck, and A. A. Transeth, "A robotic concept for remote inspection and maintenance on oil platforms," in *Proc. ASME 28th Int. Conf. on Ocean, Offshore and Arctic Engineering (OMAE 2009)*, Hawaii, USA, May 31 - June 5 2009.

[20] Ø. Skotheim and F. Couweleers, "Structured light projection for accurate 3D shape determination," in *Advances in Experimental Mechanics*, ser. Proc. Int. Conf. on Experimental Mechanics, C. Pappalettere, Ed., Bari, Italy, 2004, pp. 536–541.