Control Engineering Practice ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Motion planning and control of robotic manipulators on seaborne platforms

Pål J. From [a,*], Jan T. Gravdahl [a], Tommy Lillehagen [b], Pieter Abbeel [c]

[a] Department of Engineering Cybernetics, Norwegian University of Science and Technology, Trondheim, Norway
[b] Department of Strategic R&D for Oil, Gas and Petrochemicals, Division of Process Automation, ABB, Oslo, Norway
[c] Department of EECS, University of California, 253 Cory Hall, Berkeley, CA 94720-1770, USA

ARTICLE INFO

Article history:
Received 28 December 2009
Accepted 10 April 2011

Keywords:
Marine systems
Robotics
Stochastic systems
Predictive control

ABSTRACT

Robots on ships have to endure large inertial forces due to the non-inertial motion of the ship. The ship motion affects both the motion planning and control of the manipulator, and accurate predictions can improve performance substantially. It is thus important to investigate to what extent it is possible to predict the future motion of a ship. Based on these predictions, this paper presents a new approach to motion planning and control of such manipulators. It is shown that the effects of the non-inertial forces can be eliminated—in fact, the robot can even leverage the inertial forces to improve performance compared to robots on a fixed base. In particular it is shown that by including the inertial forces in the motion planning the wear and tear on the robot due to these forces can be reduced substantially. To perform realistic experiments a 9-DoF robot is used. The first five joints are used to generate the real ship motion, and the last four joints are used for motion planning. The dynamic coupling between the first five and the last four joints is thus exactly the same as the dynamic coupling between a ship and a manipulator, which allows for very realistic experiments.

1. Introduction

Robotic manipulators on non-inertial platforms such as ships have to endure large inertial forces due to the non-inertial motion of the platform. When the non-inertial platform's motion is known, motion planning and control algorithms can try to eliminate these perturbations. In some situations the motion planning algorithms can even leverage the inertial forces to more economically move to a target point. However, for many non-inertial platforms, the motion is unknown.

This paper first investigates to what extent it is possible to predict the future motion of a ship. Using real ship motion measurements one can study how the uncertainty of the prediction algorithms changes with the prediction horizon. Then a new motion planning approach that finds the optimal trajectory from an initial to a target configuration based on the predicted future motion of the ship is presented. It is also shown that by including the uncertainty in the cost function the maximum torques needed to reach the target configuration can be reduced.

Due to the stochastic nature of the ship motion and the dynamic coupling between the ship and the manipulator, empirical studies are extremely important to validate both the ship motion predictions and the motion planning algorithms. To perform realistic experiments a 9-DoF robot is used. The first five joints are used to generate the real ship motion, and the last four joints are used for motion planning. The dynamic coupling between the first five and the last four joints is thus exactly the same as the dynamic coupling between a ship and a manipulator. It is thus possible to perform very realistic experiments as ship motions measured from a real ship are used to generate the actual motion and at the same time realistic ship motion predictions are used as inputs to the motion planner. To get statistically meaningful results several simulations are performed to confirm the experimental results.

Ships and other seaborne platforms are expected to become increasingly unmanned in the future and the need for autonomously operating robots for surveillance, maintenance, and operation will continue to increase over time (Kitarovic, Tomas, & Cisic, 2005; Love, Jansen, & Pin, 2004). The demand for unmanned operation becomes even higher in harsh environments such as high sea state (Fig. 1), when it can be dangerous for human operators to be exposed. High sea environments are not only dangerous to human operators, they also pose significant challenges for robotic control. Large inertial forces will influence the manipulator and, when not anticipated and accounted for, can make the operation inaccurate, extremely energy demanding, and sometimes even impossible due to torque limits. The inertial forces thus need to be taken into account in the motion planning of the robot.

* Corresponding author at: Institutt for teknisk kybernetikk, 7491, Trondheim, Norway.
E-mail address: from@itk.ntnu.no (P.J. From).

0967-0661/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.conengprac.2011.04.007

Fig. 1. A ship in high sea. The wave forces can result in very high accelerations in the ship motion. (©2010 IEEE)

In From, Duindam, Gravdahl, and Sastry (2009) the authors solve the problem of optimal motion planning for a robot mounted on a ship *under the assumption the base motion is known for all times.* The approach includes the ship motion in the trajectory planning problem and an optimal trajectory in terms of actuator torques is found. However, in most practical situations the forces acting on the ship due to the interaction with waves and wind are very irregular and one cannot expect to know the base motion for all times (From, Gravdahl, & Abbeel, 2010).

In this paper a model predictive control scheme is adopted in the sense that a ship model is used to predict the future motion of the ship, and that these predictions are updated at every time step using feed-back from the ship measurements. The extent to which it is possible to obtain accurate ship motion predictions can thus directly affect how well the inertial forces can be compensated for or even taken advantage of. However, the accuracy of the ship motion prediction not only directly determines how optimal a solution one can achieve, it also affects the computational requirements. In a receding horizon setting, where the optimal control input sequence is re-computed at regular intervals, the computational burden will increase for an inaccurate model: for an inaccurate model the initialization point taken from the previous solution is further away from the optimal solution. In addition to affecting the choice of horizon the modeling error thus directly determines the frequency for which the optimal control or optimal trajectory can be recalculated.

An important contribution in this paper is the use of real empirical data for both the experiments and the simulations. Much of the literature on ship motion prediction uses computer generated data to verify the accuracy of the prediction algorithms, which leads to unrealistically small errors in the predictions. An important difference between the work presented here and previous work is thus the use of real full-scale ship motion data to test the performance of the ship motion prediction algorithms.

Another important contribution is the realistic experiments. For the experiments a 9-DoF robot was used. The real ship motion was fed into the first five joints, generating a very realistic ship motion, and the last four joints are then used for optimal motion planning for a 4-DoF robot on a moving base. Because the yaw motion is practically zero, the five degrees of freedom representing the ship motion and the four degrees of freedom representing the manipulator allows for experiments performed on a 4-DoF robot mounted on a base with exactly the same motion as if the robot had been mounted on a ship. The predictions used are based on the real ship motion and the experiments thus give important information on to what extent it is possible to compensate for the inertial forces.

Experiments are important to study how the motion of the base affects the manipulator dynamics. The experiments also allow for direct measurements of the torques that act on the manipulator due to the inertial forces. This paper thus presents, for the first time, a detailed study of how the inertial motion of the ship maps to the joint torques. This is used to show that the forces that act on the manipulator due to the waves are significant and cannot be ignored in the motion planning and control of the manipulator.

Stochastic uncertainty is present in a wide variety of systems, ranging from mechanical systems and process control to finance. In general, receding horizon control is a well-suited control scheme to deal with uncertainties, but most approaches do not use information about the probability distribution governing the uncertainty and they only assume that the uncertainty is bounded. Thus, the information about the probabilistic distribution is ignored and the worst-case representation of the disturbances or constraints often leads to a conservative solution.

In this paper a new motion planning algorithm that also minimizes the variance of the controlled state is presented. First, real ship motion measurements are used to calculate the variance of the predictions of the forced state. It is found that the variance is different for roll, pitch and yaw. Second, it is shown how to use the Extended Kalman filter to find how the variance in the forced state maps to the variance of the controlled state, i.e., how the uncertainties in the ship motion predictions map to uncertainties in the robot state. The general idea of this approach is to exploit the fact that there are some components of the ship motion that are more difficult to predict than others. Also, for different configurations of the robot, the inertial forces will affect the robot differently. Thus, by including the variance in the cost function one can force the motion planner to choose a trajectory that is less affected by the largest and most uncertain components of the base motion. When a receding horizon approach is applied, it is found that by augmenting the cost function to also include the variance it is possible to choose a longer horizon than when the variance is not included.

The paper is organized as follows: Section 2 gives a short introduction to ship-manipulator modeling, presented in more detail in From et al. (2009). The ship motion prediction algorithms used are presented in Section 3, and in Section 4 it is shown how to use these predictions to improve the motion planning and control of a robotic manipulator on a moving base. The simulations and empirical studies are presented in Section 5. Related research and references are discussed in Section 6. A preliminary version of the work presented here appeared in From et al. (2010).

2. Ship-manipulator modeling

In From et al. (2009) the classical dynamic equations for a serial manipulator arm with 1-DoF joints were extended to include the forced 6-DoF motion of the base. For more details on how to derive the dynamics see From et al. (2009) and Duindam and Stramigioli (2007, 2008). Consider the setup of Fig. 2 describing a general n-link robot manipulator arm attached to a moving base and choose an inertial coordinate frame Ψ_0, a frame Ψ_b rigidly attached to the moving base, and n frames Ψ_i (not shown) attached to each link i at its center of mass. Finally, choose a vector $q \in \mathbb{R}^n$ that describes the configuration of the n joints. Using standard notation (Murray, Li, & Sastry, 1994), the pose of each frame Ψ_i relative to Ψ_0 can be described as a homogeneous transformation matrix $g_{0i} \in SE(3)$. This pose can

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

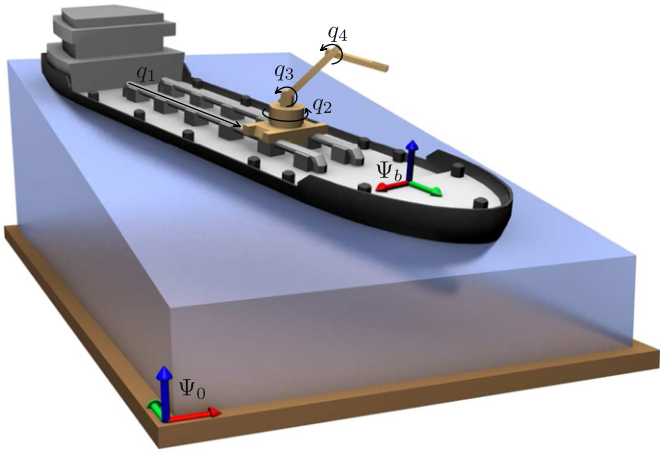

Fig. 2. Model setup for a four-link robot attached to a non-inertial base with coordinate frame Ψ_b. Frame Ψ_0 denotes the inertial reference frame. (©2010 IEEE)

also be described using the vector of joint coordinates q as

$$g_{0i} = g_{0b} g_{bi}(q). \tag{1}$$

The base pose g_{0b} and the joint positions q thus fully determine the configuration of the robot.

In a similar way, the spatial velocity of each link can be expressed using twists (Murray et al., 1994):

$$V^0_{0i} = \begin{bmatrix} v^0_{0i} \\ \omega^0_{0i} \end{bmatrix} = V^0_{0b} + V^0_{bi} = \mathrm{Ad}_{g_{0b}}(V^b_{0b} + J_i(q)\dot{q}) \tag{2}$$

where v^0_{0i} and ω^0_{0i} are the linear and angular velocities, respectively, of link i relative to the inertial frame expressed in the inertial frame, $J_i(q) \in \mathbb{R}^{6\times n}$ is the geometric Jacobian of link i relative to Ψ_b, the adjoint is defined as $\mathrm{Ad}_g := \left[\begin{smallmatrix} R & \hat{p}R \\ 0 & R \end{smallmatrix}\right] \in \mathbb{R}^{6\times 6}$, and $\hat{p} \in \mathbb{R}^{3\times 3}$ is the skew-symmetric matrix such that $\hat{p}x = p \times x$ for all $p,x \in \mathbb{R}^3$. The velocity state is thus fully determined given the twist V^b_{0b} of the base and the joint velocities $\dot{q}$.

The dynamic equations can be written in block-form as

$$\begin{bmatrix} M_{vv} & M_{qv}^T \\ M_{qv} & M_{qq} \end{bmatrix} \begin{bmatrix} \dot{V}^b_{0b} \\ \ddot{q} \end{bmatrix} + \begin{bmatrix} C_{vv} & C_{vq} \\ C_{qv} & C_{qq} \end{bmatrix} \begin{bmatrix} V^b_{0b} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} F^b_b \\ \tau \end{bmatrix} \tag{3}$$

with F^b_b the external wrench on the base link, expressed in coordinates Ψ_b (such that it is collocated with the twist V^b_{0b}).

This paper is concerned with the effects that the base pose g_{0b}, the base velocity V^b_{0b}, and the base acceleration $\dot{V}^b_{0b}$ have on the manipulator dynamics. This coupling is clearly visible when rewriting the dynamics as

$$M_{qq}\ddot{q} + C_{qq}\dot{q} + \underbrace{M_{qv}\dot{V}^b_{0b} + C_{qv}V^b_{0b}}_{\text{inertial forces}} = \tau. \tag{4}$$

Finally, the way the gravitational forces map to joint torques depends on the configuration of the base and is added to the right hand side of (4). The torque associated with link i is given by

$$\tau^i_g = J_i(q)\mathrm{Ad}^T_{g_{bi}}(q)F^i_g(g_{0i},q). \tag{5}$$

F^i_g is given by

$$F^i_g = \begin{bmatrix} f_g \\ \hat{r}^i_g f_g \end{bmatrix} = -m_i g \begin{bmatrix} R_{0i}e_z \\ \hat{r}^i_g R_{0i} e_z \end{bmatrix} \tag{6}$$

where $e_z = [0\ 0\ 1]^T$, r^i_g is the center of mass of link i expressed in frame Ψ_i, and $\hat{r}$ is the skew-symmetric matrix representation of r. Note that R_{0i} depends on the configuration of the base with respect to the inertial frame.

3. Ship motion prediction

Due to the stochastic nature of the forces that act on ships, ship motion prediction is a very difficult problem. This section presents two simple and computationally efficient methods for predicting the future motion of the ship: the auto-regressive (AR) predictor and a predictor using a superposition of sinusoidal waves representation. In this paper the focus is on these simple and computationally efficient methods as they do not require a ship or wave model, nor external sensors such as wave cameras or force sensors. Adding sensors or using more advanced methods will allow for more accurate predictions and the results presented in the following sections can thus probably be improved if external sensors are added or ship and wave models are available.

The auto-regressive (AR) predictor is an all-pole model (i.e., no inputs) and gives an estimate of the output directly without the need for information about the forces that cause the motion. Write

$$y(t) = -a_1 y(t-1) - a_2 y(t-2) - \cdots - a_n y(t-n) \tag{7}$$

and define

$$\phi(t) = [-y(t)\ -y(t-1)\ -y(t-2)\ \cdots\ -y(t-n+1)]^T,$$

$$\theta = [a_1\ a_2\ \cdots a_n]^T. \tag{8}$$

Collecting N samples and stacking ϕ in Φ and y in Y one can find the optimal parameters θ in the least squares sense by

$$\theta = (\Phi^T\Phi)^{-1}\Phi^T Y. \tag{9}$$

The prediction problem is then solved by

$$y(t+1) = \phi(t)^T\theta. \tag{10}$$

Alternatively one can fit the superposition of N sine waves to the measurements in the least squares sense. Following the approach in Chung, Bien, and Kim (1990) write

$$\xi(t) = \sum_{i=1}^{N} A_i \sin(\omega_i t + b_i) \tag{11}$$

where A_i is the amplitude of the sines, ω_i is the frequency and b_i is the phase. Assuming the frequencies are found from the peaks in the frequency spectrum, the problem amounts to finding A_i and b_i. Note that

$$\xi(t) = \sum_{i=1}^{N} a_{2i-1}\sin(\omega_i t) + a_{2i}\cos(\omega_i t) \tag{12}$$

where $a_{2i-1} = A_i\cos(b_i)$ and $a_{2i} = A_i\sin(b_i)$ are used to handle phase shifts. The parameters

$$\theta = [a_1\ a_2\ \cdots\ a_{2N}]^T \tag{13}$$

representing the best fit in the least squares sense are then found from (9) with

$$\phi(t) = [\sin(\omega_1 t)\ \cos(\omega_1 t)\ \cdots\ \sin(\omega_N t)\ \cos(\omega_N t)]^T.$$

4. Motion planning and control in a stochastic environment

This section discusses the motion planning problem, i.e., to take the manipulator from an initial configuration to a target configuration using as little torque as possible. By planning the motion so that the inertial forces contribute to the motion instead of working against them, it is possible, in addition to save energy, to achieve more accurate trajectory tracking and reduce the strain and tension on the manipulator.

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

4.1. Motion planning

Consider the control law

$$\tau = \tau_{ff} + \tau_{PD} \tag{14}$$

where

$$\begin{aligned}\tau_{ff} = &\underbrace{M_{qq}\ddot{q}_d + C_{qq}\dot{q}_d}_{\text{tracking terms}} \\ &+ \underbrace{M_{qV}\dot{V}_{0b}^b + C_{qV}V_{0b}^b}_{\text{compensation for inertial forces}} \\ &- \underbrace{\sum_{n=1}^{n}(J_i \mathrm{Ad}_{g_{bi}}^{\mathrm{T}} F_g^i)}_{\text{gravity compensation}}\end{aligned} \tag{15}$$

$$\tau_{PD} = \underbrace{K_P(q_d - q) + K_D(\dot{q}_d - \dot{q})}_{\text{PD-controller}}. \tag{16}$$

This is the standard augmented PD control law which in this case also compensates for the inertial forces. Based on the predictions of g_{0b}, V_{0b}^b, and $\dot{V}_{0b}^b$ for a given horizon this control law tries to cancel these disturbances regardless of whether they contribute to the desired motion or not.

For trajectory tracking this is in general a very energy demanding solution. When large inertial forces are present, simply canceling these terms as in τ_{ff} in (15) may require excessive joint torques. Note that due to the high gain individual joint controllers normally found in industrial manipulators, trajectory tracking will in general be very good. In the setting of this paper, however, this solution is very energy demanding and the inertial forces may cause stress and tension on the robot joints. Thus, instead of regarding these terms as disturbances, the prediction of the ship motion can be included in the motion planning of the robot. The planner can then use this information to calculate the trajectory that requires the least actuator torque for the given base motion. As an example, consider a manipulator that is to move from the left to the right on the ship. If it chooses to start the trajectory at a time when the inertial forces contribute to the desired motion it can potentially get an almost free ride from one side to the other. If it simply chooses to cancel these disturbances, for example by (15), it might end up following a trajectory for which the inertial forces are working against the desired motion for the entire interval. One intuitive situation where this can occur is when the manipulator moves uphill instead of downhill for the entire motion and will thus not take advantage of the gravitational forces.

This paper follows the approach presented in From et al. (2009) and solves the motion planning problem by numerically minimizing an objective function representing the joint torques with respect to the start and end time. The motion is restricted to lie in the interval (T_0,T_1) (constraint of the optimization problem) while the interval for which the actual motion occurs (the solution to the optimization problem) is denoted (t_0,t_1). The cost function is found from the whole interval (T_0,T_1) and not only the interval when the motion occurs, i.e.,

$$P_{tor} = \min_{t_0,t_1} \int_{t=T_0}^{T_1} \tau^{\mathrm{T}} D \tau \, dt \tag{17}$$

where $T_0 \le t_0 < t_1 \le T_1$ and D is a positive definite matrix that defines a metric in τ-space. In From et al. (2009) this was solved assuming the base motion was known. This paper uses realistic predictions of the base motion in the cost function. These predictions will become less accurate as the horizon increases, and it is investigated how the choice of motion planning algorithm and prediction horizon affects the performance of the motion planner.

4.2. Stochastic MPC

As will be clear in Section 5.1 the accuracy of the predictions is different for the different axes. For example for the AR predictor the predictions of the angular acceleration about the y-axis (pitch) is less accurate than the x- and z-axes (roll and yaw). Also, the linear acceleration in the direction of the x-axis (surge) is far more accurate than the y- and z-axes (sway and heave). This section presents a modified cost function that minimizes also the expected variance on the output (the robot state) assuming information about how the uncertainty evolves with time is available for the different components of the ship motion.

In Cannon, Couchman, and Kouvaritakis (2007) the control objective of the stochastic model predictive control (stochastic MPC) law is to regulate the expected value and variance of the output state. In this section the same ideas are applied and by including the covariance matrix in the cost function a trajectory that also minimizes the variance is chosen. The cost function as defined in Cannon et al. (2007) is given by

$$P = \sum_{j=0}^{N-1} l(k+j|k) + L(k+N|k) \tag{18}$$

where L is the cost-to-go function and

$$l(k+j|k) = \bar{x}^2(k+j|k) + \kappa^2 \sigma_x^2(k+j|k) \tag{19}$$

with

$$\bar{x}(k+j|k) = \mathbb{E}_k(x(k+j|k)) \tag{20}$$

$$\sigma_x^2(k+j|k) = \mathbb{E}_k[x(k+j|k) - \bar{x}(k+j|k)]^2 \tag{21}$$

denoting the expected value and variance of the state x at time $k+j$ given the information available at time k, denoted $x(k+j|k)$. x is the state that is to be controlled, which in the setting of this paper is the state of the robot, i.e., $x=q$. The relative weighing of the expected value and the variance can thus be controlled directly through the parameter κ.

In the following the same ideas are applied and the covariance matrix is included in the cost function so that a trajectory that also minimizes the covariance, in addition to the joint torques, is chosen. A cost function similar to the one found in Cannon et al. (2007) is given by

$$P = \sum_{j=0}^{N-1} l(k+j|k) + L(k+N|k) \tag{22}$$

where $L(k+N|k)$ is the cost-to-go and

$$l(k+j|k) = \bar{q}^2(k+j|k) + \kappa \| \Sigma_q(k+j|k) \|_r \tag{23}$$

with $\bar{q}(k+j|k) = \mathbb{E}[q(k+j|k)]$ denoting the expected state of the robot and

$$\Sigma_q = \begin{bmatrix} \mathbb{E}[(q_1 - \bar{q}_1)(q_1 - \bar{q}_1)] & \cdots & \mathbb{E}[(q_1 - \bar{q}_1)(q_n - \bar{q}_n)] \\ \vdots & \ddots & \vdots \\ \mathbb{E}[(q_n - \bar{q}_n)(q_1 - \bar{q}_1)] & \cdots & \mathbb{E}[(q_n - \bar{q}_n)(q_n - \bar{q}_n)] \end{bmatrix} \tag{24}$$

so that $\Sigma_q(k+j|k)$ is the covariance matrix of $q(k+j|k)$ given the measurements of the robot state $q(i)$ and the ship state $g_{0b}(i)$ for $i = k_0 \cdots k$, and $\| \cdot \|_r$ denotes the Euclidean norm of each row. Similarly to the approach in Cannon et al. (2007), also shown in (19), the uncertainty can be added to the cost in (17) to get a cost function in the form

$$P_{var} = \min_{t_0,t_1} \int_{t=T_0}^{T_1} (\tau^{\mathrm{T}}(t) D \tau(t) + \kappa \| \Sigma_q(t|T_0) \|) \, dt. \tag{25}$$

where $\| \Sigma_q \|$ denotes the Euclidean norm of the covariance matrix. The problem is thus to find the start time $t_{i,0}$ and the end time $t_{i,1}$ for the motion of each joint i subject to the restriction

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

$T_0 \le t_{i,0} < t_{i,1} \le T_1$. The cost, however, sums over the entire predefined interval (T_0, T_1), i.e., also when the joints do not move and $\dot{q}_i(t) = 0$. Note that in this case the joint torques are not necessarily zero because they will need to compensate for the inertial forces.

The optimization problem in (25) is thus solved for each time step. A receding horizon control scheme is obtained by introducing a feed-back by updating the ship and robot position and recalculating the predicted ship motions before the optimization problem is solved at each time step.

Assume that each degree of freedom of the ship motion has a distribution in the form

$$\begin{bmatrix} \dot{u} \\ \dot{v} \\ \vdots \\ \dot{r} \end{bmatrix} \sim \begin{bmatrix} \mathcal{N}(\bar{\dot{u}}, \sigma^2_{\dot{u}}, t) \\ \mathcal{N}(\bar{\dot{v}}, \sigma^2_{\dot{v}}, t) \\ \vdots \\ \mathcal{N}(\bar{\dot{r}}, \sigma^2_{\dot{r}}, t) \end{bmatrix}. \tag{26}$$

The time dependency states that the uncertainty changes with the prediction horizon. As not only the expected value, but also the variance is known for each degree of freedom, this represents valuable information that can be included in the motion planning and control to improve performance. Examples of the expected value and the variance for the acceleration of the ship are shown in Figs. 4 and 5, respectively. Similar relations can also be found for the position and velocity of the ship.

For linear systems the Kalman filter can be used to find the expected state and the error covariance. For non-linear systems the extended Kalman filter is implemented, i.e., linearizing around the mean value and finding the expected state of the robot and the covariance matrix used in (25). It is then possible to include both the expected state and the covariance at time $(k+j)$ given the measurements available at time k also for non-linear systems. This is then included in the cost function as in (25) and the optimal solution is found by minimizing the weighed cost of the expected value and the covariance.

5. Simulations and experimental studies

This section presents the simulation results and the results from experiments performed in the lab. Due to the stochastic nature of the problem, both experimental and simulation studies are important. The experimental studies are carried out in order to gain insight into how the moving base and errors in the prediction of the base motion affect the motion planning and control of the manipulator. The experiments make it possible to measure the torques directly and get valuable insight into the complex coupling between the ship and manipulator and how this coupling affects the motion planning and control of the manipulator. Also, due to the stochastic nature of the disturbances, simulations are important to be able to perform a sufficient number of runs and hence get statistically meaningful results.

5.1. Empirical data of the ship motion

The need for empirical data is of utmost importance when verifying the performance of prediction algorithms. Most publications on the topic of ship motion prediction use computer generated data such as a combination of sines, a wave model, or a sine with added noise. This will not give a good performance indicator because of the stochastic nature of the waves. In this work measurements from a real full-scale ship are used. This makes it possible to compare the different prediction algorithms on real data and, most importantly, it gives valuable information about the accuracy of the predictions for different prediction horizons. The ship used to collect the measurements was the RS 113 "Erik Bye", shown in Fig. 3, which is a 20.4 m long Emmy Dyvi class ship and weighs 96 tons. The wave height at the time of the measurements was about 1 m. The ship is owned by Redningsselskapet AS, Norway.

The most important information when including the future motion of the ship in a model predictive control approach is to have as accurate predictions as possible of the velocity and the acceleration entries of the state. If gravity plays an important role, the attitude of the ship should also be included. Fig. 4 shows an example of the true and estimated angular acceleration (roll) of the ship. Note that the predicted acceleration needs to be estimated at short time intervals to maintain a low prediction error. In general one can obtain very good results when the predictions are computed every 0.5 or 1.0 s. For predictions up to 3 s the predictions are also reasonable and no large errors occur. The prediction accuracy depends on the frequency of the waves. In general, predictions with horizons longer than one wave length, which in this case was 3–5 s, are not very reliable. This

Fig. 3. The Emmy Dyvi class "Erik Bye" used to measure the ship motions. "Erik Bye" is owned by Redningsselskapet, Norway. (©2010 IEEE)

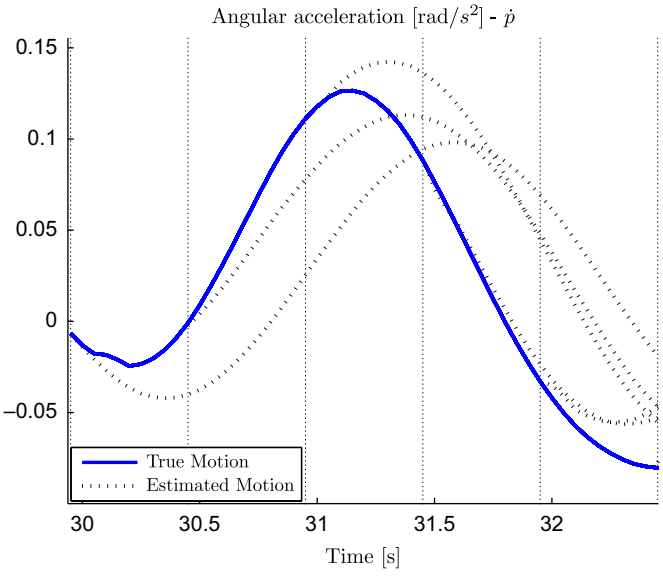

Fig. 4. A typical example of true and predicted motion (AR) where a new prediction is calculated every 0.5 s. This clearly shows the need for re-computation of the predicted motion at short time intervals. The angular acceleration around the x-axis (roll) is shown. (©2010 IEEE)

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

is mainly due to the fact that these may be out of phase which may lead to very large errors.

To get a more structured formulation of how the error changes with the prediction horizon one can look at how the standard deviation evolves over time. This is important as it allows the path planner and controller to include uncertainty in the cost function and minimize this. Fig. 5 shows the standard deviation for the six degrees of freedom of the velocity state denoted $v = [u\ v\ w\ p\ q\ r]^{\mathrm{T}}$. Note that the AR predictions are more accurate than the superposition of sines, except for horizons of 0.2 s or shorter. The AR method is thus chosen to predict the ship motions to be used in the motion planning and control presented in the next sections. The main problem in ship motion prediction is the irregularity of the waves. Due to these irregularities it is very difficult to predict the ship motion as most predictors try to fit sinusoidal waves to the measured ship motion. When these irregularities occur the prediction errors can become very large and the predictors used in this paper will not perform very well. To get good performance also in these cases, external sensors or wave cameras need to be added.

5.2. Experimental setup

The experiments are performed using true motion data from the full-scale ship "Erik Bye" to generate the ship motion, and predictions from the AR model as presented in Section 3 to compute the optimal trajectories. All the data is obtained using inertial sensing in six degrees of freedom. One example of the measured ship motions in six degrees of freedom and the predicted ship motions using the AR model can be seen in Fig. 7. For the experimental setup a 9-DoF robot (3-DoF gantry crane and 6-DoF industrial manipulator, see Fig. 6) is used.

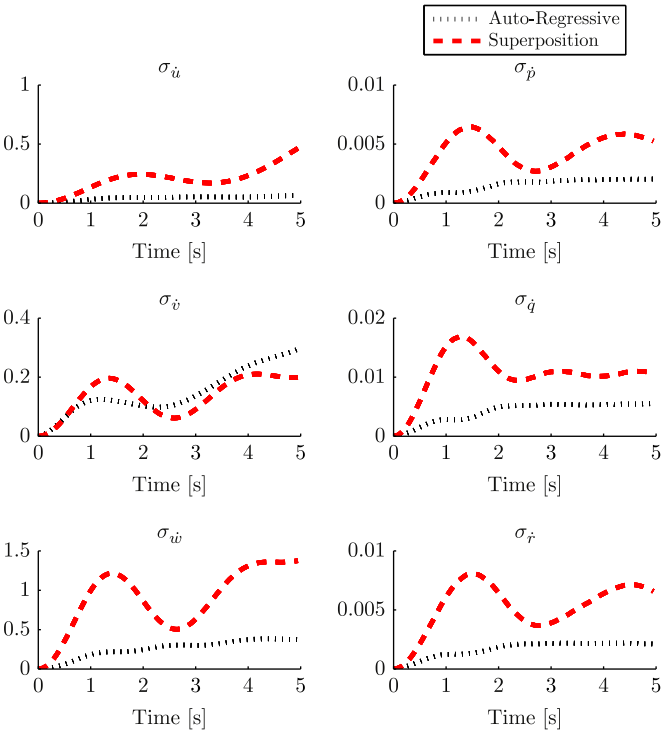

Fig. 5. The standard deviation as a function of the prediction horizon. Note that the standard deviation is smaller with the AR model than for the superposition of sines, except for predictions of 0.2 s or less (not visible on the figure). The standard deviation is calculated from 200 samples, for "Erik Bye" moving at 15 knots at 1 m wave height. (©2010 IEEE)

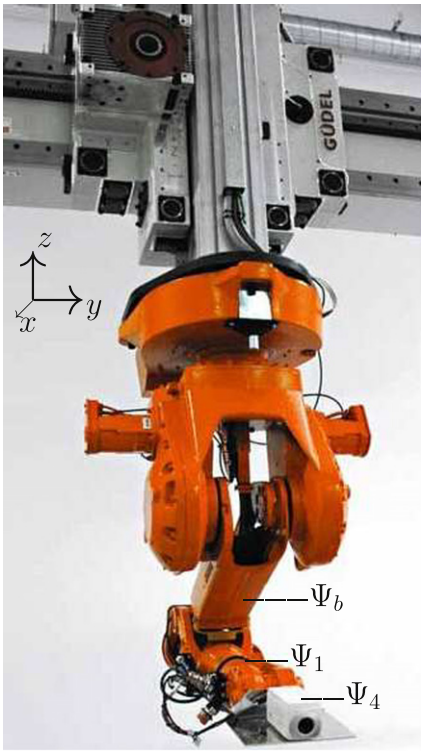

Fig. 6. The ABB IRB-2400 robot mounted on a Güdel gantry crane. The first 5 joints (the 3 of the gantry crane and the first 2 of the robot) are used to generate the ship motion, represented by Ψ_b, and the last four joints of the robot, represented by $\Psi_{1,2,3,4}$, are used for optimal motion planning of a 4-DoF robot. Links 2 and 3 of the robot are hidden in the wrist and cannot be seen in the figure. Courtesy ABB Strategic R&D for Oil, Gas and Petrochemicals. (©2010 IEEE)

A Güdel gantry crane with three translational degrees of freedom is used to hold an ABB IRB-2400 with six rotational degrees of freedom which is mounted upside down as in Fig. 6. The end effector of the robot is a camera which weighs about 1 kg. Due to limitations in the workspace of the gantry crane some of the components of the base motion had to be scaled down. The sway motion is scaled down with about 10% and the heave motion is scaled down with a factor of 5 to avoid collision between the manipulator and the floor.

The first 5-DoF were used to generate the ship motion. The robot thus generates the surge, sway and heave motion with the linear actuators of the gantry crane and the roll and pitch motion with the first two rotational joints of the manipulator. The yaw motion is very small and can be neglected. The last four links of the manipulator are then considered a standard manipulator on which the control and motion planning algorithms are tested. The motion of the "base", i.e., link 2 of the manipulator (see Fig. 6), is set to the same as the measurements taken from the full-scale ship. This setup thus allows for very realistic experiments as the base link has exactly the same motion as the real ship. Due to the difference in the inertia between the ship and the manipulator the motion of the manipulator does not affect the motion of the ship. The dynamic coupling between the base link and the last four links of the manipulator is thus the same as for a 4-DoF manipulator mounted on a moving base.

5.3. Experimental results based on predicted ship motions

This section shows how to exploit the inertial forces and choose a trajectory that minimizes the torques and the strain

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

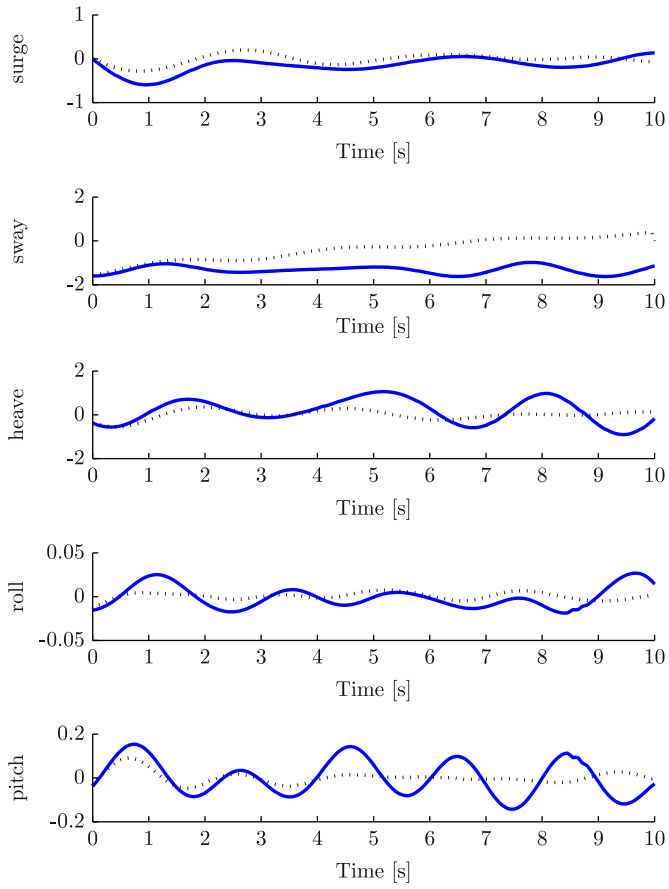

Fig. 7. The motion collected from the ship measurements and given as input to the first five joints of the manipulator. The full line shows the measured accelerations and the dotted line shows the predicted ship motion used in the motion planners. Note that the predictions have a damping effect as the predictions become less accurate.

and tension on the manipulator. The approach is based on From et al. (2009) and the motion planning problem is solved by numerically minimizing an objective function. Two different objective functions are used: (i) the objective function in (17) which minimizes torques only, and (ii) the objective function in (25) which minimizes torques and the variance (Fig. 8).

For a given interval $(T_0,T_1)=(0,10)$, the optimization problem is then to find the start and the end time for each joint with the restriction that $T_0 \le t_{i,0} < t_{i,1} \le T_1$ for all i. The shape, but not the starting time or length, of the trajectory is thus assumed fixed. The motion is given by shaping the acceleration as one period of a sinusoidal motion (From et al., 2009) and is illustrated in Fig. 8.

By fixing the shape, a sub-optimal solution that can be solved efficiently and in real time is found. Finding the optimal solution over all trajectories is a huge optimization problem and computationally too demanding to be solved in real time. The start and target configurations are chosen as

$$q_d(0) = [0\ 0\ 0\ 0]^{\mathrm{T}},$$

$$q_d(T) = \left[\frac{\pi}{2}\ \frac{\pi}{2}\ \frac{\pi}{2}\ \frac{\pi}{2}\right]^{\mathrm{T}}.$$

The motion planning problem thus amounts to finding the eight parameters (one start and end time for each joint) that minimize the cost integrated over a fixed time interval while the robot start and target configurations are satisfied. The optimization problem is solved for all the joints at the same time to be able to control the resulting end-effector trajectory. It is thus possible to include aspects such as collision avoidance and to more actively control the trajectory of the end effector. In this paper, however, collision avoidance is not addressed and it is assumed that the robot moves in the free space.

For the first cost function (17) let D be a positive definite matrix that defines a metric in τ-space. Choose $D = 10^{-6} \cdot \mathrm{diag}[10\ 2\ 2\ 5]$ reflecting the masses of the links. For the second cost function (25) the weight is chosen as $\kappa = 10$ to enhance the effect of adding the variance, but otherwise the same weighing as for (17) is used.

First the experiments are performed with a baseline trajectory, i.e., a starting time $t_{i,0} = T_0$ and end time $t_{i,1} = T_1$ for all i. This is how one would choose the trajectory when no information about the motion of the ship is available. Second, the optimization problem is solved assuming complete knowledge about the ship motion, including future motion. This gives information about how well one can expect to perform if accurate motion predictions are available. Finally, two experiments are performed where the trajectory is calculated based on the predicted base motion. The base motion used in the experiments is always exactly the motion of the ship while the predictions of the base motion are used in the motion planner. This is done with the cost function given in (17) and in (25). Thus, the four algorithms in Table 1 were tested.

At the beginning of the time interval the ship motion is predicted based on the previous 5 s of known ship motion. As the predictions become increasingly inaccurate, only the ship motion predictions in the interval where these are sufficiently accurate are used. The optimal horizon for the different cost functions are found in Section 5.5. In the remaining of the interval, where no accurate predictions are available, the algorithm assumes that the base does not move. Then, every 1 s the current position of the ship is updated and the ship motion predictions are re-computed for the optimal horizon. Finally the optimal trajectory is re-computed.

To compare the energy used for the different algorithms the following cost function (power) is used:

$$P_i = \int_{t=T_0}^{T_1} \tau_i \dot{q}_i \, dt \tag{27}$$

for each joint $i = 1 \cdots 4$. This gives a good indication of how much energy is used to take the manipulator from the start to the target configuration. It also gives a good indication of the strain and tension on the manipulator due to the moving base.

The positions and the velocities of the four joints for the four different trajectories are found in Fig. 8. Note that the trajectories found for the four different approaches are quite different. First, the trajectories based on the true and predicted base motions are quite different, especially for joints 3 and 4. The trajectory based on the true ship motion (Alg. 2) takes advantage of a favorable motion at the end of the interval while the trajectories based on the predicted motion (Alg. 3 and 4) take advantage of a favorable motion in the beginning of the interval. This difference arises because information about the ship motion towards the end of the interval is not available to the algorithms based on the predicted ship motion, due to the long horizon. Alg. 2 chooses to wait to start the motion until the very last minute. This probably has a simple explanation. All the joints start out in the initial condition, as shown in Fig. 6, and end up with a rotation of $\pi/2$ radians for all the joints. The largest components of the ship motion lie in the xz-plane, i.e., the pitch is far bigger than the roll, and the sway is fairly constant. It can be seen from Fig. 6 that the inertial forces affect the manipulator less for the initial configuration than for the target configuration and thus less energy is used if the motion is delayed until the end of the interval.

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

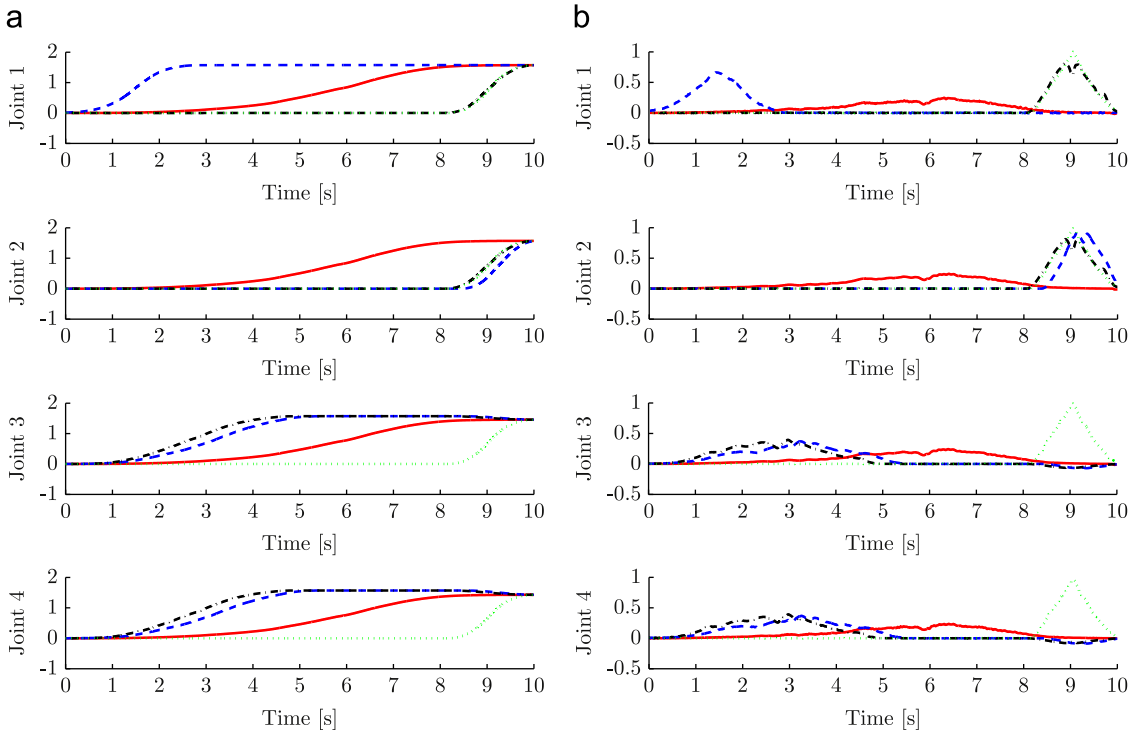

Fig. 8. Four different trajectories are shown: Alg. 1—a baseline trajectory with maximum motion duration (full, red lines), Alg. 2—an optimal trajectory taking the correct base motion into account (dotted, green lines), Alg. 3—an optimized trajectory taking the predicted base motion into account (dashed, blue lines), and Alg. 4—an optimized trajectory taking the predicted base motion into account where the variance is included in the cost function (dot-dashed, black lines). (a) Joint Positions, (b) Joint Velocities. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
The four algorithms studied. Note the color coding used in the figures in this section.

Alg. 1 (——)	A baseline trajectory with maximum motion duration for all the joints
Alg. 2 (· · · · ·)	An optimized trajectory taking the correct base motion into account
Alg. 3 (- - - -)	An optimized trajectory taking the predicted base motion into account with cost function (17)
Alg. 4 (— · —·)	An optimized trajectory taking the predicted base motion into account with cost function (25)

There is also an interesting difference between the torque-based cost function (Alg. 3) and the cost function that is based on both torque and variance (Alg. 4). Note that Alg. 4. calculates an optimal trajectory that is closer to Alg. 2, which is considered the optimal trajectory. This can be seen in joint 1, which is the most important joint due to its inertia and also the joint with the highest weight in the cost function. As the uncertainty is biggest in surge, heave, and pitch, and the variance is included in the cost function, this algorithm will try to keep a configuration for which these components of the ship motion do not affect the manipulator. In other words, the algorithm will choose a trajectory where the mapping from the components of the ship motion with the largest uncertainties (surge, heave, and pitch) to the internal forces of the manipulator is as small as possible, which is the reason for delaying the motion for joint 1.

Fig. 9 shows the value of the cost (27) for the four algorithms, i.e., torque times velocity for each joint. The most apparent observation here is the difference between joint 1, which carries a lot of weight, and joints 2, 3 and 4, whose inertia around the axis of rotation are very low, i.e., about 10% of that of joint 1. Note that even though both joint 1 and 2 carry a lot of weight, the axis of rotation of joint 2 is along the axis with the smallest inertia but for joint 1 it is along the axis with the largest inertia. Note that due to the motion of the base, joint 1 needs to use a substantial amount of torque just to keep the arm fixed. In fact, the largest amount of torque is used to compensate for the base motion, and not to move the joint from the start to the target configuration. This is a surprising result as the base motion simulates only 1 m wave height and because some components of the motion are scaled down. One can thus conclude that for a manipulator mounted on a ship in high sea, the amount of torque needed to compensate for the inertial forces will be substantial and needs to be included in the motion planning and control of the manipulator in order to obtain energy efficient control and to minimize the wear and tear on the joints.

Table 2 shows the values of the cost function (27) for each joint for the four algorithms studied. The average values of the cost function (27) for all joints can be found in Table 3. All the entries are scaled so that the maximum value equals 1, so the values of the different joints cannot be compared directly. Note that all the optimization algorithms perform better than the baseline trajectory, in terms of power. Also note that the optimization based on the predicted base motion performs better than the optimization based on the real base motion when evaluating the cost given in (27), which may come as a surprise.

On the other hand, as can be seen in Table 4, this is not the case when the square of the torque is used. In this case the optimization based on the real base motion performs best, as should be expected as this is the cost function used in the optimization algorithm. Also note that the cost with torque only performs better than when the variance is added. This is also as should be expected as the cost function used in Alg. 3 is the one found in (17).

Tables 6 and 7 show the maximum torques for the four algorithms. In this case, Alg. 4 performs better than Alg. 3.

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

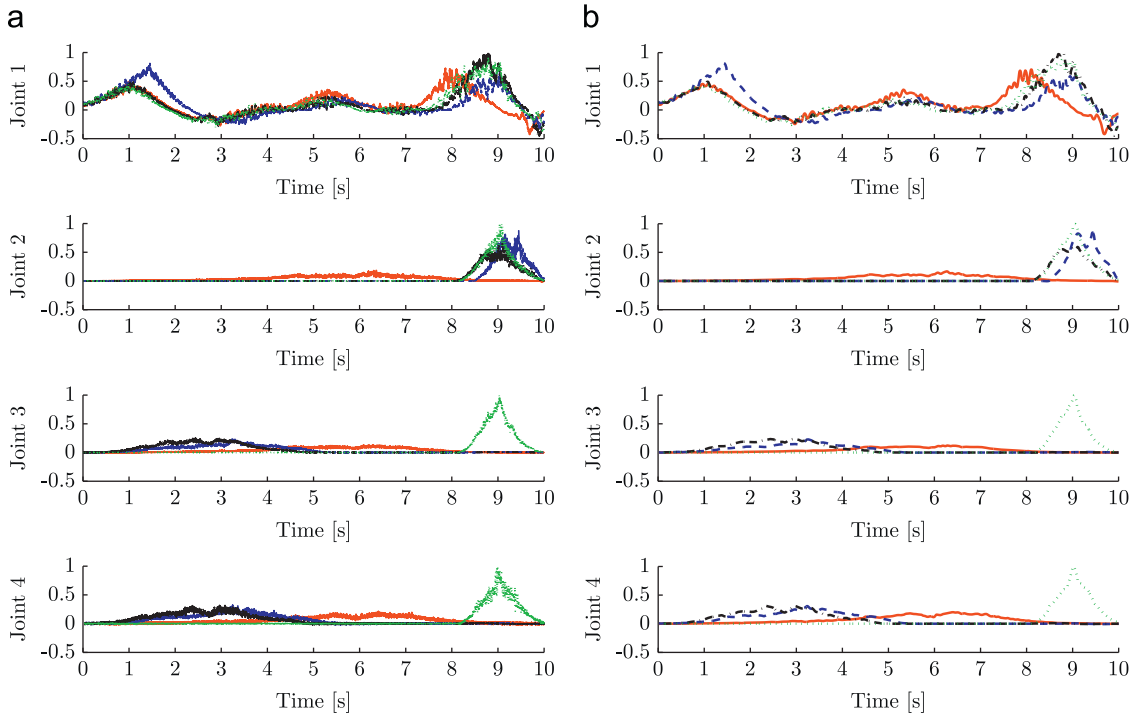

Fig. 9. The torques times velocity over the trajectory. Four different trajectories are shown: Alg. 1—a baseline trajectory with maximum motion duration (full, red lines), Alg. 2—an optimal trajectory taking the correct base motion into account (dotted, green lines), Alg. 3—an optimized trajectory taking the predicted base motion into account (dashed, blue lines), and Alg. 4—an optimized trajectory taking the predicted base motion into account where the variance is included in the cost function (dot-dashed, black lines). (a) Measured joint torques times velocities, (b) Low pass filtered torques times velocities. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
Square of the torque times velocity over the trajectory for each joint.

Joint	Alg. 1	Alg. 2	Alg. 3	Alg. 4
1	1.000	0.787	0.627	0.856
2	1.000	0.061	0.036	0.817
3	1.000	0.652	0.331	0.960
4	1.000	0.709	0.657	0.669

Table 3
Average of the square of the torque times velocity over the trajectory.

	Alg. 1	Alg. 2	Alg. 3	Alg. 4
Average	1.0000	0.552	0.413	0.825

Table 4
Square of the torques over the trajectory for each joint.

Joint	Alg. 1	Alg. 2	Alg. 3	Alg. 4
1	1.000	0.807	0.844	0.948
2	1.000	0.292	0.417	0.292
3	1.000	0.756	0.511	0.689
4	1.000	0.571	0.327	0.612

Table 5
Average values of the square of the torques over the trajectory.

	Alg. 1	Alg. 2	Alg. 3	Alg. 4
Average	1.000	0.789	0.799	0.910

Table 6
Maximum of the torques for each joint.

Joint	Alg. 1	Alg. 2	Alg. 3	Alg. 4
1	0.945	1.000	0.968	0.935
2	0.691	0.876	0.931	1.000
3	0.706	0.768	1.000	0.666
4	0.783	0.762	1.000	0.824

Table 7
The average of the maximum of the torques.

	Alg. 1	Alg. 2	Alg. 3	Alg. 4
Average	0.873	0.935	1.000	0.903

This was the main intention of adding the variance to the cost function. When the variance is added to the cost function, the algorithm (Alg. 4) chooses a safer path in the sense that the mapping from the uncertainties of the base motion to the uncertainty in the robot state is chosen so that the uncertainty in the robot state is minimized. Also note, however, that the safest path in terms of maximum torques seems to be the one with the longest duration and thus more evenly distributing the torques over the entire interval.

5.4. Simulations based on predicted ship motions

Simulations are used to verify the experimental studies by running the motion planning algorithm several times based on

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

different motion prediction data sets. The 200 data sets that all have the same characteristics were chosen in order to be able to draw some general conclusions and compare the results to the ones found in experiments. All the sets are picked from one long sampling and are thus collected during a short period of time and in a sea state for which the dominant components of the motion were in the *xz*-plane, i.e., the pitch is far bigger than the roll, and the sway is almost zero. In other words, all the data sets are measurements of the ship moving with the same velocity, in the same sea state, and with the same attack angle on the waves.

Table 8 confirms the tendency from Table 5: all the algorithms perform better than the benchmark solution. In fact, all approaches obtain better results in the simulations than for the experiments. This is a promising result as the large number of simulations make these results meaningful also for stochastic data. Note that the approach based on the real base motion performs better in this case. This is reasonable as the amount of accurate data available should allow for a more optimal solution to be found. Thus, the small difference between Alg. 2 and 3 in Table 5 cannot be taken as a general result.

Also, the maximum values found in Table 7 can be confirmed. However, Table 9 shows that Alg. 2 performs better than Alg. 1, which is not the case for the experiments. Once again the cost function based on both the torque and the variance decreases the maximum values of the torque, but increases the overall torque used during the trajectory, which again is as expected.

5.5. Horizon length

This section discusses how to choose the horizon in a receding horizon setting when only predictions of the base motions are known. The previous sections found that the predictions are very accurate for about 0.5 s and relatively accurate for a horizon of about 3 s. When it comes to motion planning, however, it turns out that a longer horizon should be chosen. The reason for this is that the algorithm searches for the time interval in which the inertial forces contribute to the motion as much as possible. It is thus desirable to choose the horizon as long as possible, but at the same time avoid using predictions that are not accurate and may lead to large errors. Table 10 shows the optimal horizon, the total cost, and the maximum torques for (i) a cost function minimizing the torques only (Alg. 3), and (ii) a cost function minimizing the torques and variance (Alg. 4). Table 10 shows the value of the cost function based on only the torque for both algorithms in order to be able to compare the two values. The optimal horizons for the two algorithms are found to be 4 and 6 s, respectively. As accurate predictions are available for about 3 s, which is shorter than the optimal horizon, the algorithms must handle the fact that for part of the interval, the predictions are not very reliable.

When the variance is included in the cost function one can allow a longer horizon. This allows the path planner to use more of the information available, albeit inaccurate, and results in a slightly lower cost than when only the torque is used. The advantage of adding the variance is rather small when looking at the squared torque, so the main advantage of including the variance in the cost function is that the maximum values of the torque decreases, as seen in Table 9.

Table 8
The square of the torques over the trajectory for all the joints. Average values of all the joints for 200 samples.

	Alg. 1	Alg. 2	Alg. 3	Alg. 4
Average	1.000	0.670	0.788	0.829

Table 9
The maximum of the torques over the trajectory for all the joints. Average values of all the joints for 200 samples.

	Alg. 1	Alg. 2	Alg. 3	Alg. 4
Average	0.890	0.823	1.000	0.888

Table 10
The cost for different horizons for a cost function without and with variance, based on 200 simulations.

Horizon	Torque	
(s)	Alg. 3	Alg. 4
1	1.000	0.897
3	0.812	0.688
4	0.630	0.688
5	0.667	0.670
6	0.778	0.611
7	1.287	0.687

6. Related research

Research on several related topics can be found in the literature. Love et al. (2004) address the impact of wave generated disturbances on the tracking control of a manipulator mounted on a ship based on the classical Lagrangian approach. They use repetitive learning control and this results in performance improvement for purely periodic motions, but they do not present a formal derivation of the dynamics equations. Kitarovic et al. (2005) and Oh, Mankala, Agrawal, and Albus (2005) address the use of cable robots for loading and unloading cargo between two ships. In the Ampelmann project (Salzmann, 2007), a Stewart platform is mounted on a ship and is used to compensate for the motion of the ship by keeping the platform still with respect to the world frame. Lebans, Wilkie, Dubay, Crabtree, and Edmonds (1997) give a cursory description of a telerobotic shipboard handling system, and Kosuge, Okuda, and Fukuda (1992) and Kajita and Kosuge (1997) address the control of robots floating on the water utilizing vehicle restoring forces. Other related research areas are macro-micro manipulators (Yoshikawa, Harada, & Matsumoto, 1996; Bowling & Khatib, 1997), underwater vehicle-manipulator systems (McMillan, Orin, & McGhee, 1995) and spacecraft-manipulator systems (Egeland & Sagli, 1993).

Most previous work deals with robots mounted on a free-floating base. There is, however, an important difference between modeling a robot on a forced and a free-floating base. A forced base motion will add inertial forces to the dynamic equations that do not arise in the free-floating case, such as spacecraft-manipulator systems and manipulators on small AUVs.

There are some papers in the literature considering the prediction of ship motion. Yang, Pota, Garratt, and Ugrinovskii (2008a,b) discuss the problem of landing a helicopter on a ship in high sea and predict the ship motion by fitting the ship model to the measured data using recursive least squares. Khan, Bil, and Marion (2005) use artificial neural networks to solve the same problem. In Chung et al. (1990) the sea excitation is extrapolated using the superposition of sines approach and the ship motion is predicted using the ship model driven by the extrapolated forces.

Stochastic model predictive control (MPC) is discussed in detail in Cannon et al. (2007) and Couchman, Cannon, and

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007

Kouvaritakis (2006) where the output variance is included in the cost function. The relation between the input and the output variance is also important in performance assessment, which is discussed in an MPC setting in Zhang and Li (2007), and the minimum variance performance map is discussed in Harrison and Qin (2009). The output variance is also discussed in the minimum variance control of stochastic processes (Åstrøm, 1967).

7. Conclusions

This paper presents the first detailed discussion regarding several important aspects of ship-manipulator systems. First, the extent to which the ship motion can be predicted and for what horizon this can be used in the motion planning and control of the manipulator is investigated. Then several different approaches to the motion planning problem for robots mounted on ships are discussed and it is shown that the amount of torque needed to reach a desired configuration can be reduced by including the predicted base motion in the motion planner. The torques needed to reach the target can be reduced even for relatively moderate ship motions. Thus, one may conclude that for a ship in high sea it is possible to substantially improve performance and allow for efficient motion planning and accurate control in settings where this would otherwise not be possible due to large inertial forces. It is also shown that by including the variance of the predicted ship motion in the motion planner the maximum torques needed and thus also the strain and tension on the robot are reduced even further.

Several simulations and experiments are performed to validate the approaches presented. For the first time, detailed experimental results of ship-manipulator systems under the influence of inertial forces are presented. Several simulations with a large number of data sets are also performed and it is shown that the simulation results are consistent with the experimental result. Also for the first time, the effects of including the variance in the cost function in a receding horizon control law are investigated through experiments based on real ship motion data.

It is found that the inertial forces that act on a manipulator mounted on a ship in only 1 m wave height pose significant challenges in robot control and motion planning, especially for the joints that move large inertia. Also, for a ship in high sea state, the inertial forces are significantly higher and must be included in both the motion planner and the control to guarantee safe, cost efficient, and accurate manipulation. An interesting topic for future work is thus to repeat the same experiments for a ship in high sea.

Acknowledgments

P. From and T. Gravdahl wish to acknowledge the support of the Norwegian Research Council and the TAIL IO project for their continued funding and support for this research. The TAIL IO project is an international cooperative research project led by Statoil and an R&D consortium consisting of ABB, IBM, Aker Solutions and SKF. During the work with this paper the first author was with the University of California at Berkeley. The authors would also like to thank Redningsselskapet AS for allowing the use of their ship to collect the motion data.

Appendix A. Robot parameters

The lengths and the masses of the last four links of the ABB IRB-2400 industrial robot are found in Table A.1.

Table A.1
The lengths and the masses of the last four links of the ABB IRB-2400.

Link	Length (mm)	Mass (kg)
3	130	35
4	755	20
5	85	2
6	300	1

References

Åstrøm, K. J. (1967). Computer control of a paper machine—an application of linear stochastic control theory. *IBM Journal of Research and Development*, *7*, 389–396.

Bowling, A., & Khatib, O. (1997). Design of macro/mini manipulators for optimal dynamic performance. In: *IEEE International Conference on Robotics and Automation* (pp. 449–454), Albuquerque, New Mexico, USA.

Cannon, M., Couchman, P., & Kouvaritakis, B. (2007). MPC for stochastic systems. *Assessment and Future Directions of Nonlinear Model Predictive Control*, *358*, 255–268.

Chung, J. C., Bien, Z., & Kim, Y. S. (1990). A note on ship-motion prediction based on wave-excitation input estimation. *IEEE Journal of Oceanic Engineering*, *15*(3), 244–250.

Couchman, P., Cannon, M., & Kouvaritakis, B. (2006). Stochastic MPC with inequality stability constraints. *Automatica*, *42*(12), 2169–2174.

Duindam, V., & Stramigioli, S. (2007). Lagrangian dynamics of open multibody systems with generalized holonomic and nonholonomic joints. In: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3342–3347). San Diego, CA, USA.

Duindam, V., & Stramigioli, S. (2008). Singularity-free dynamic equations of open-chain mechanisms with general holonomic and nonholonomic joints. *IEEE Transactions on Robotics*, *24*(3), 517–526.

Egeland, O., & Sagli, J. R. (1993). Coordination of motion in a spacecraft/manipulator system. *International Journal of Robotics Research*, *12*(4), 366–379.

From, P.J., Duindam, V., Gravdahl, J.T., & Sastry, S. (2009). Modeling and motion planning for mechanisms on a non-inertial base. In: *International Conference of Robotics and Automation* (pp. 3320–3326), Kobe, Japan.

From, P.J., Gravdahl, J.T., & Abbeel, P. (2010). On the influence of ship motion prediction accuracy on motion planning and control of robotic manipulators on seaborne platforms. In: *International Conference of Robotics and Automation*. Anchorage, Alaska, USA.

Harrison, C. A., & Qin, S. J. (2009). Minimum variance performance map for constrained model predictive control. *Journal of Process Control*, *19*, 1199–1204.

Kajita, H., & Kosuge, K. (1997). Force control of robot floating on the water utilizing vehicle restoring force. In: *IEEE/RSJ International Conference on Intelligent Robot and Systems* (Vol. 1, pp. 162–167). Grenoble, France.

Khan, A., Bil, C., & Marion, K.E. (2005). Ship motion prediction for launch and recovery of air vehicles. In: *MTS/IEEE OCEANS* (pp. 2795–2801). Washington CD, USA.

Kitarovic, J., Tomas, V., & Cisic, D. (2005). The electronic and informatics age—a new stage in developing highly effective ships. In: *International ELMAR Symposium* (pp. 385–388). Zadar, Croatia.

Kosuge, K., Okuda, M., & Fukuda, T. (1992). Motion control of manipulator/vehicle system floating on water. In: *IEEE International Workshop on Advanced Motion Control* (pp. 560–511). Nagoya, Japan.

Lebans, G., Wilkie, K., Dubay, R., Crabtree, D., & Edmonds, T. (1997). Telerobotic shipboard handling system. In: *MTS/IEEE OCEANS* (pp. 1237–1241). Halifax, Nova Scotia, Canada.

Love, L.J., Jansen, J.F., & Pin, F.G. (2004). On the modeling of robots operating on ships. In: *IEEE International Conference on Robotics and Automation* (pp. 2436–2443). New Orleans, LA, USA.

McMillan, S., Orin, D. E., & McGhee, R. B. (1995). Efficient dynamic simulation of an underwater vehicle with a robotic manipulator. *IEEE Transactions on Systems, Man and Cybernetics*, *25*(8), 1194–1206.

Murray, R. M., Li, Z., & Sastry, S. S. (1994). *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL, USA: CRC Press.

Oh, S.-R., Mankala, K., Agrawal, S., & Albus, J. (2005). Dynamic modeling and robust controller design of a two-stage parallel cable robot. *Multibody System Dynamics*, *13*(4), 385–399.

Salzmann, D.C. (2007). Ampelmann prototype—developing a motion compensating platform for offshore access. In: *European Wind Energy Conference*. Milano, Italy.

Yang, X., Pota, H., Garratt, M., & Ugrinovskii, V. (2008). Prediction of vertical motions for landing operations of uavs. In: *IEEE Conference on Decision and Control* (pp. 5058–5063). Cancun, Mexico.

Yang, X., Pota, H., Garratt, M., & Ugrinovskii, V. (2008). Ship motion prediction for maritime flight operations. In: *IFAC World Congress* (pp. 12407–12412) Seoul, Korea.

Yoshikawa, T., Harada, K., & Matsumoto, A. (1996). Hybrid position/force control of flexible-macro/rigid-micro manipulator systems. *IEEE Transactions on Robotics and Automation*, *12*(4), 633–640.

Zhang, Q., & Li, S. (2007). Enhanced performance assessment of subspace model-based predictive controller with parameters tuning. *The Canadian Journal of Chemical Engineering*, *85*, 537–548.

Please cite this article as: From, P. J., et al. Motion planning and control of robotic manipulators on seaborne platforms. *Control Engineering Practice* (2011), doi:10.1016/j.conengprac.2011.04.007