Appendix 2 – Scientific Publications & Technical Notes

ATTACHMENTS

Attachment 1: S. Chaudhuri, H. Ramezani and J. Jouffroy, "*Remarks on augmenting stability and safety in nominal controllers: a practical perspective*," 2025 European Control Conference (ECC), Thessaloniki, Greece, 2025, pp. 2121-2128, doi: 10.23919/ECC65951.2025.11187095.

Attachment 2: S. Chaudhuri, H. Ramezani and J. Jouffroy, " *Ensuring safe roll dynamics in marine vessels based on a safety-enhanced control framework*," 2026 International Conference on Electrical, Computer and Energy Technologies (ICECET), 6-9 July 2026, Rome, Italy (Presented).

Attachment 3: S. Chaudhuri, H. Ramezani and J. Jouffroy, " *Learning Marine Model Uncertainties for Real-Time Safety-Enhanced Roll Stabilization*," Prepared for 2027 European Control Conference (ECC), 13-16 July 2027, Brussels, Belgium.

Attachment 4: B.R. Andersen, M.R Jørgensen, M.T. Frandsen (SDU Physics), "*Reinforcement learning for roll stabilization of vessels using a digital twin",* Technical Note, 2025 (Part of the submitted MSc thesis and being prepared as a conference article).

Attachment 5: Pablo Alvarez Benito, "*Catamaran Stabilization Control System Testing",* In-Company Project with Dacoma and SDU Mechatronics, 2024 - 25.

2025 23rd European Control Conference (ECC)
June 24-27, 2025. Thessaloniki, Greece

A2_Attachment 1

Remarks on augmenting stability and safety in nominal controllers: a practical perspective

Shouvik Chaudhuri[1], Hossein Ramezani[1], and Jerome Jouffroy[1]

Abstract— **This paper introduces an innovative safety-enhanced control framework that incorporates safety and stability constraints into nominal controllers through the quadratic programming approach. These constraints are defined using properly defined control Lyapunov functions and control barrier functions. In contrast to the traditional method, this framework enhances the nominal controller with stability and safety elements, eliminating the need for a controller redesign. Consequently, this increases the robustness and stability of the system while maintaining its performance. The approach is validated through a case study involving an overhead crane, which illustrates the real-world efficacy of this method.**

I. Introduction

Safety in industrial systems is paramount, particularly when dealing with scenarios where the margin for error is minimal, and the consequences of failure can be catastrophic. The emergence of safety-critical control in the last decade is a direct consequence of addressing this problem through the use of a *safety barrier* around control actions to prevent the system from breaching defined safety limits during operation [1], [2]. This technique has gained much traction in sectors such as automotive systems and mobile robotics [2], [3], where systems need to operate independently under uncertain and frequently challenging conditions. The core of safety-critical control is its capability to permit controllers to function efficiently under typical circumstances while intervening to limit the control actions as the system approaches potentially unsafe states [3]–[5]. The conventional approach to safety-critical control integrates control Lyapunov functions (CLFs) to ensure performance with control barrier functions (CBFs) to maintain safety in real time by employing a state-dependent quadratic program (QP) to determine a control solution at each instant [3], [4], [6].

In order to follow this control protocol for any non-linear system, designers must address several major challenges to ensure safety and synthesize a safe controller. A significant challenge in this context is to select an appropriate CLF and CBF, taking into account the system dynamics, which is not a trivial task. Furthermore, CLF and CBF-based methodologies depend heavily on a precise understanding of the system model, which is not always possible, thus introducing modeling/parametric uncertainties [2]. Over the last few years, numerous papers have appeared in the literature employing various methodologies to address these design challenges [5], [7]–[10]. Recent studies have also introduced approaches based on sum-of-squares (SOS) optimization-based filters [11], synthesis of compatible CLF and CBF functions without relying on nominal controllers [12], and relaxed filter-based compatibility conditions [13]. In contrast, most safety-critical control methods using the CLF-CBF-QP approach design a framework that handles both performance and safety from scratch. Many industrial systems already have nominal controllers that ensure performance within certain conditions. In such cases, a method is needed that retains the performance of the nominal controller while adding global stability and safety using the CLF-CBF-QP approach. Even though a similar concept has been proposed in CBF-QP approach to add safety to a nominal controllers [14], when it comes to CLF-QP and CLF-CBF-QP approaches where the stability constraint is included, no nominal controller has been considered in the QP problem.

In this paper, we propose a safety-enhanced control framework by redefining the CLF-CBF-QP framework to involve nominal controllers, followed by systematic augmentation of stability and safety constraints. In the conventional CLF-CBF-QP approach, the nominal controller is not included, and the cost function is focused on energy optimization which does not necessarily contribute to enhance the performance of the system. Instead, the CLF constraint is the only part of the QP which aims at improving the performance by ensuring exponential stability of the system with a tunable convergence rate. In contrast, the proposed approach allows the cost function to provide a performance close to the nominal controller, while the CLF focuses solely on ensuring stability. In practice, selecting a Lyapunov function that captures the complexities of the system is challenging, and poor selection can hinder performance improvements by making the control input infeasible, even with tunable parameters. In order to ease this problem, we propose a systematic method for determining this CLF based on the closed-loop response of the nominal controller. The only condition for this approach to work is that the nominal controller should be expressed in a state-feedback form, which can be included in the cost function formulation. Since the formulation remains in QP form, the safety condition can be incorporated as a constraint, similar to the conventional methodology.

The remainder of the paper is structured as follows. Section II provides the background for the safety-critical control and barrier functions. Section III presents the proposed safety-enhanced control framework with the aid of an illustrative example. Section IV demonstrates the effectiveness of the method through a case study involving an overhead crane

[1] Shouvik Chaudhuri, Hossein Ramezani, and Jerome Jouffroy are with Centre for Industrial Mechanics (CIM), Department of Mechanical and Electrical Engineering (DME), University of Southern Denmark (SDU), Sønderborg, Denmark chaudhuri@sdu.dk; ramezani@sdu.dk; jerome@sdu.dk

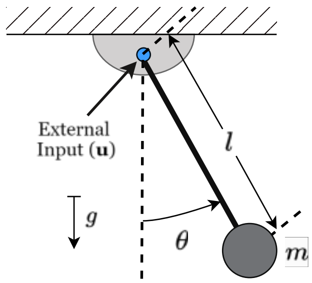

Fig. 1: Schematic of the actuated pendulum system

(MIMO system) which needs to move between different locations with limited swing. Concluding remarks are discussed in Section V.

II. Background on safety-critical control

Consider that the dynamics of a system is represented in the following non-linear control-affine form

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}(t)) + \mathbf{g}(\mathbf{x}(t))\mathbf{u} \\ y &= \mathbf{h}(\mathbf{x}(t)) \end{aligned} \tag{1}$$

with compact set $\mathcal{X} \subset \mathbb{R}^n$, state vector $\mathbf{x} \in \mathcal{X}$, $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^n$ and $\mathbf{g} : \mathbb{R}^n \to \mathbb{R}^{n \times m}$ represent the vector fields of the system which are locally Lipschitz continuous. We consider that the system is subject to input constraints such that

$$\mathbf{u} \in \mathcal{U}(\mathbf{x}(t)), \forall t \geq 0 \tag{2}$$

where, $\mathcal{U}(\mathbf{x}(t)) \subset \mathbb{R}^m$ is the state dependent set of admissible control inputs. Throughout this section, we will discuss key aspects of safety-critical control with the help of an illustrative example involving an actuated pendulum as described below.

Illustrative Example: Consider a pendulum that is actuated by an external control input (typically a torque), as shown in Fig. 1. The dynamics of this pendulum with damping and an external control input can be expressed by the model,

$$\ddot{\theta}(t) = -\frac{b}{ml^2}\dot{\theta}(t) - \frac{g}{l}\sin\theta(t) + \frac{1}{ml^2}u(t) \tag{3}$$

where, θ represents the angular displacement of the pendulum bob, m denotes the mass of the pendulum bob, l indicates the length of the pendulum rod, b is the damping coefficient, g stands for gravitational acceleration and u is the external control input. Our control objective is to maneuver the states of the system, $\mathbf{x} = [\theta, \ \dot{\theta}]$, from an initial condition, $\mathbf{x_0}$, to the equilibrium point, $\mathbf{x}_{eq} = [0, 0]$. A nominal controller of the form, $\mathbf{u} = \mathbf{k}(\mathbf{x})$, can be designed based on the linearized model of this system as given below.

$$\dot{\mathbf{x}} = \begin{bmatrix} 0 & 1 \\ -\frac{g}{l} & -\frac{b}{ml^2} \end{bmatrix} \mathbf{x} + \begin{bmatrix} 0 \\ ; \frac{1}{ml^2} \end{bmatrix} u \tag{4}$$

In this example, we select the following parameter values: $m = 1$ [kg], $b = 1$ [kgm^2/s], $1 = 1$ [m], and $g = 9.81$ [m/s^2]. The nominal controller is based on the LQR methodology considering $Q_{lqr} = 10I_2$ and $R_{lqr} = 1$.

A. Preliminaries

The closed-loop stability of the system (1) can be guaranteed if, for any small perturbation from the equilibrium point, the states of the system $\mathbf{x}(t)$, remain close to this equilibrium $\forall t \geq 0$. This can be assessed based on the notion of a control Lyapunov function (CLF), as defined below.

Definition 1 *(CLF [1], [4], [6], [15]): For the dynamical system in (1) and a Lipschitz continuous feedback law, $\mathbf{u} = \mathbf{k}(\mathbf{x})$, the exponential stability of the closed-loop system can be guaranteed by a continuously differentiable positive definite function $\mathcal{V} : \mathcal{X} \to \mathbb{R}$, called the Control Lyapunov Function (CLF), if there exists a class $\mathcal{K}$ function λ such that,*

$$\dot{\mathcal{V}}(\mathbf{x}, \mathbf{u}) = L_\mathbf{f}\mathcal{V}(\mathbf{x}) + L_\mathbf{g}\mathcal{V}(\mathbf{x})\mathbf{u} \leq -\lambda(\mathcal{V}(\mathbf{x})), \ \forall \ \mathbf{x} \in \mathcal{X} \tag{5}$$

where $L_\mathbf{f}\mathcal{V}(\mathbf{x})$ and $L_\mathbf{g}\mathcal{V}(\mathbf{x})$ are the Lie derivatives of $\mathcal{V}(\mathbf{x})$ along $\mathbf{f}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$, respectively.

The safety of the dynamical system in (1) can be ensured by maintaining the states of the system $\mathbf{x}(t)$ within a pre-defined *safe set*, $\mathcal{C}_s$ that is *forward invariant*. So, for any initial condition $\mathbf{x}_0 \triangleq \mathbf{x}(0) \in \mathbb{R}$, the states of the system in (1) strictly remain within $\mathcal{C}_s$, i.e. $\mathbf{x}(t) \in \mathcal{C}_s, \ \forall t \geq 0$. The safe set $\mathcal{C}_s$ is defined as follows.

Definition 2 *(Safe Set [1], [6], [15]): A safe set $\mathcal{C}_s$ is defined as the zero-superlevel set of a continuously differentiable function $h : \mathbb{R}^n \to \mathbb{R}$, such that,*

$$\begin{aligned} \mathcal{C}_s &= \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) \geq 0\} \\ \partial\mathcal{C}_s &= \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) = 0\} \\ Int(\mathcal{C}_s) &= \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}) > 0\} \end{aligned} \tag{6}$$

where, $\partial\mathcal{C}_s$ and $Int(\mathcal{C}_s)$ represent the boundary and the interior, respectively, of the safe set $\mathcal{C}_s$.

Definition 3 *(CBF [1], [6], [15]): Considering the dynamical system in (1) and the safe set $\mathcal{C}_s$ defined in (6), a function $\mathcal{B} : \mathcal{C}_s \to \mathbb{R}$, can be a Control Barrier Function (CBF), if there exists an extended class $\mathcal{K}_\infty$ function γ such that,*

$$\begin{gathered} \frac{\partial\mathcal{B}}{\partial\mathbf{x}} \neq 0, \ \forall \ \mathbf{x} \in \partial\mathcal{C}_s \\ \dot{\mathcal{B}}(\mathbf{x}, \mathbf{u}) = L_\mathbf{f}\mathcal{B}(\mathbf{x}) + L_\mathbf{g}\mathcal{B}(\mathbf{x})\mathbf{u} \geq -\gamma(\mathcal{B}(\mathbf{x})), \ \forall \ \mathbf{x} \in \mathcal{X} \end{gathered} \tag{7}$$

where, $L_\mathbf{f}\mathcal{B}(\mathbf{x})$ and $L_\mathbf{g}\mathcal{B}(\mathbf{x})$ are the Lie derivatives of $\mathcal{B}(\mathbf{x})$ along $\mathbf{f}(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$, respectively.

B. Quadratic Programs with CLF and CBF

In a QP formulation, the CLF and CBF conditions in (5) and (7) along with the input conditions in (2) can be added as constraints of the problem. The solution of the problem produces a feedback controller $\mathbf{u}^*$, which ensures that control objectives are met while keeping the state trajectories within the desired safe set [1], [6]. The QP formulation may be presented in multiple forms, depending on the context of the application, as detailed below.

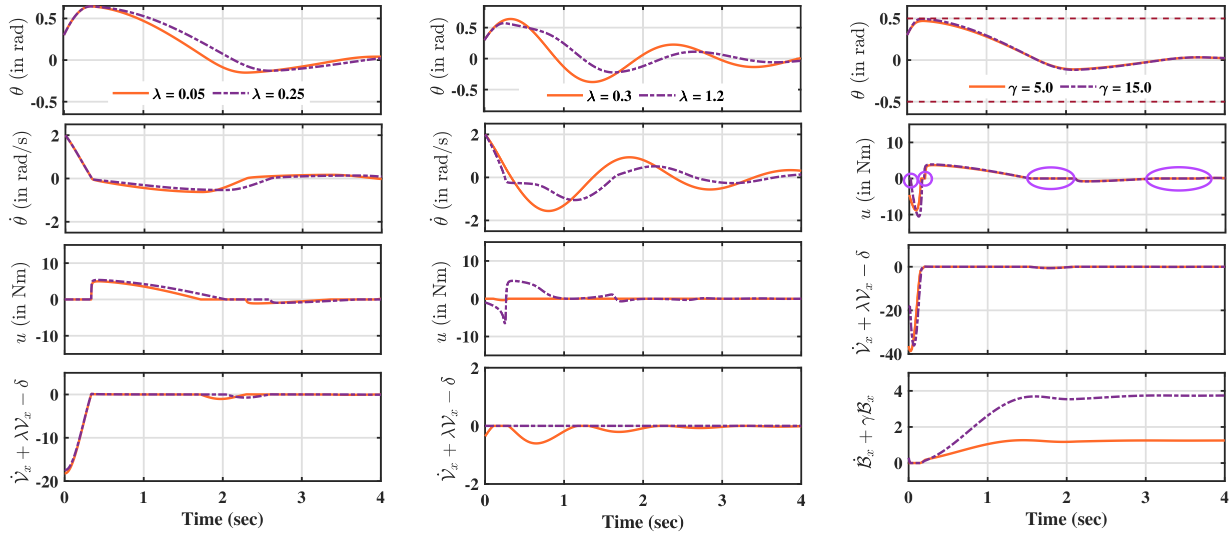

(a) **Example 1 - Case1** (b) **Example 1 - Case2** (c) **Example 2**

Fig. 2: Comparison of state variables (θ and $\dot{\theta}$), control input (u) and the CLF-CBF constraints for the actuated pendulum in (a) Example 1 - Case 1 (b) Example 1 - Case 2 (c) Example 2

1) CLF-QP: In this QP formulation, the system must adhere to the input constraints described in (2), and the CLF, which ensures the exponential stability of the system as defined in (5), is also imposed as a constraint [4]. Moreover, a slack variable δ is introduced, which relaxes the CLF constraint. This in turn establishes a sufficiently large feasible region for the intersecting constraints, thus ensuring the solution's feasibility [1], [3], [5]. The QP is structured as,

CLF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\,\delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u}-\mathbf{u}_{eq})^T\mathbf{H}(\mathbf{u}-\mathbf{u}_{eq}) + p\delta^2 \tag{8}$$

$$\text{s.t.} \quad \mathbf{u} \in \mathcal{U}(\mathbf{x}) \qquad \text{(Input Constraint)}$$

$$L_{\mathbf{f}}\mathcal{V}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{V}(\mathbf{x})\mathbf{u} \leq -\lambda\mathcal{V}(\mathbf{x}) + \delta \qquad \text{(CLF Constraint)}$$

To attain optimal performance, the parameter λ within the CLF constraint can be tuned. The cost function is modified by incorporating a quadratic cost ($p\delta^2$), where p is large positive value which may be varied to achieve an optimal solution. The aim of the cost function is to minimize the control energy while ensuring that the solution moves towards the equilibrium point, where $\mathbf{u}_{eq}$ denotes the input at the equilibrium point and $\mathbf{H} \in \mathbb{R}^m \times \mathbb{R}^m$ is a positive definite matrix.

Example 1: The CLF-QP formulation is applied to the actuated pendulum example represented by (3) with the initial value $\mathbf{x_0}$ as $(0.3\ [\text{rad}], 2\ [\text{rad/s}])$. As the focus is on the CLF constraint, we select the input constraint to be large, i.e. $\pm 20[\text{Nm}]$. We consider two cases for selecting the CLF, assuming the form $\mathcal{V} = \mathbf{x}^T\mathbf{P}\mathbf{x}$. For Case 1, we choose a canonical form with $\mathbf{P} = I_2$, whereas for Case 2, a proper selection of $\mathbf{P}$ is made as proposed later in Section III-A. The idea is to show the efficacy of the parameter λ in improving the performance of the solution. Figures 2(a) and 2(b) show the results for Cases 1 and 2, respectively. Since, a large value for λ may result in an infeasible solution, in each figure, the results are reported for two values of λ, where the largest value is close to the infeasibility limit. As expected, if $\mathbf{u}^* = 0$ satisfies the CLF constraint, the QP will select this as the optimal solution for both cases since this forms a min-norm controller. For example, in Case 2 where $\lambda = 0.3$, $\mathbf{u}^* = 0$ at all times. In summary, this approach focuses primarily on reducing control energy rather than improving system performance, as seen in both scenarios. Although λ can contribute to performance improvement, its impact is restricted.

2) CLF-CBF-QP: In this QP formulation, a CBF constraint is incorporated into the CLF-QP problem to guarantee safety of the system [3], [5]. Similar to the CLF-QP, the slack variable δ and the quadratic cost ($p\delta^2$) is applied to this problem as well. A valid solution to this QP indicates that all three constraints intersect. This QP is structured as,

CLF-CBF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\,\delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u}-\mathbf{u}_{eq})^T\mathbf{H}(\mathbf{u}-\mathbf{u}_{eq}) + p\,\delta^2 \tag{9}$$

$$\text{s.t.} \quad \mathbf{u} \in \mathcal{U}(\mathbf{x}) \qquad \text{(Input Constraint)}$$

$$L_{\mathbf{f}}\mathcal{V}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{V}(\mathbf{x})\mathbf{u} \leq -\lambda\mathcal{V}(\mathbf{x}) + \delta \qquad \text{(CLF Constraint)}$$

$$L_{\mathbf{f}}\mathcal{B}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{B}(\mathbf{x})\mathbf{u} \geq -\gamma\mathcal{B}(\mathbf{x}) \qquad \text{(CBF Constraint)}$$

Example 2: The CLF-CBF-QP formulation is applied to the same pendulum as in *Example 1*, assuming the same Case 1 CLF with $\lambda = 0.05$, and the results are displayed in 2(c) for two cases of γ. The CBF for this example is chosen as : $\mathcal{B}(\mathbf{x}) = (\theta_{safe}^2 - \theta^2) - (k_1\, \dot{\theta}^2/2\dot{\theta}_{max})$, with the safe limit $\theta_{safe} = \pm 0.5\ [rad]$ and $k_1 = 0.015$. It is noticeable that the CBF constraint manages to maintain the angular displacement within the safe limit, compared to Case 1 of *Example 1* shown in Fig. 2(a). However, as with the CLF-QP formulation, the emphasis is on minimizing energy. Consequently, the solution $\mathbf{u}^* = 0$ is often chosen as the optimal solution (shown with circled regions in the u plot), when it meets the set constraints.

III. Safety-Enhanced Control Framework

For many practical applications, the nominal controller provides an acceptable performance with local stability around the equilibrium point for closed-loop nonlinear system. Therefore, it is likely that the performance of the controller will drop when the working conditions move away from the equilibrium point, if it does not become unstable. This could be due to the fact that the controller belongs to the PID family, heuristically tuned by the operators, or it is a model-based controller designed for the linearized model of the system. On the other hand, a conventional CLF-QP algorithm, as represented in (8), while effective in ensuring global stability, might exhibit poor performance, even in proximity to the equilibrium point. This could be attributed to the selection of the CLF, leading to a scenario where even the tuning parameter λ does not enhance performance because it might result in an infeasible solution space. Moreover, the cost function in (8) does not aim to improve performance, as the focus is on energy optimization.

To address these issues, we propose to keep the nominal controller as the baseline controller in a state feedback form and add the guarantee of stability to that. This approach can be applied to both CLF-QP and CLF-CBF-QP. In the conventional formulation of both these approaches, the nominal controller is neglected and the performance is provided by the CLF itself with the aid of the parameter λ. The idea is to define a cost function that takes into account the deviation of the present control solution from the solution offered by the nominal controller, i.e. $(\mathbf{u} - \mathbf{k}(\mathbf{x}))$, where $\mathbf{k}(\mathbf{x})$ represents the nonlinear state feedback form of the nominal controller. Since the cost function aims to provide a performance close to the nominal controller, the CLF is not required to focus on the performance. This indicates that the CLF is sufficient for achieving global stability, which may or may not be exponential in nature. So, we can eliminate the term $\lambda\mathcal{V}(\mathbf{x})$ from the CLF constraint. Furthermore, we aim to identify Lyapunov functions with constraints in the QP problem that preserve the nominal controller's solution while ensuring stability.

A. Selection of the CLF

We seek a proper CLF which fulfills the following conditions:

1. *Systematic Solution*: It can be systematically selected by leveraging the solution provided by the nominal controller, offering one possible solution to the problem.
2. *Minimum Interference:* It should interfere minimally with the solution, $\mathbf{k}(\mathbf{x})$, offered by the nominal controller.

In order to achieve these, our approach leverages the structure of the closed-loop system formed by the nominal controller. Specifically, we linearize the nonlinear system around the equilibrium point and compute the Jacobians of the system dynamics $\mathbf{f}(\mathbf{x})$, input matrix $\mathbf{g}(\mathbf{x})$, and nominal controller $\mathbf{k}(\mathbf{x})$ with respect to the system vector $\mathbf{x}$, yielding $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{K}$, respectively. The closed-loop system matrix is then defined as, $\mathbf{A}_{cl} = \mathbf{A} + \mathbf{B}\,\mathbf{K}$. Then, the Lyapunov function can be selected as [16], [17],

$$\mathcal{V} = \mathbf{x}^T\, \mathbf{P}_l\, \mathbf{x} \tag{10}$$

Here, the matrix P_l can be found by solving the following Lyapunov equation

$$\mathbf{P}_l\, \mathbf{A}_{cl} + \mathbf{A}_{cl}^T\, \mathbf{P}_l = -\mathbf{Q} \tag{11}$$

where, $\mathbf{Q} \in \mathbb{R}^{n\times n}$ is a properly chosen positive definite matrix. For reasons of simplicity, we can choose $\mathbf{Q} = \alpha I$, where $\alpha \in \mathbb{R}$. It can be shown that the choice of α has no impact on the performance of the controller in the QP formulation, since the nature of the CLF constraint stays negative definite/ semi-definite with α behaving only as a multiplying factor.

B. Proposed CLF-QP Formulation

Considering the insights provided regarding the nominal controller and the choice of the CLF, we can formulate a CLF-QP problem as given below.

Proposed CLF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\, \delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u} - \mathbf{k}(\mathbf{x}))^T\mathbf{H}(\mathbf{u} - \mathbf{k}(\mathbf{x})) + p\,\delta^2 \tag{12}$$

$$\begin{aligned} \text{s.t.} \quad & \mathbf{u} \in \mathcal{U}(\mathbf{x}) && \text{(Input Constraint)} \\ & L_{\mathbf{f}}\mathcal{V}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{V}(\mathbf{x})\mathbf{u} \leq \delta && \text{(CLF Constraint)} \end{aligned}$$

where, CLF is introduced as a constraint similar to (8), incorporating the slack variable δ. Here, δ creates a large enough feasible region of the intersecting constraints by allowing the $\Pi_V(\mathbf{x}, \mathbf{u})$ plane to shift. In Fig. 3(a), the blue region shows the feasible input space, $\mathcal{U}(\mathbf{x})$, and the green region shows the stability region of the chosen $\mathcal{V}(\mathbf{x})$. We will refer to the intersection of these two regions as the *solution space* $(\mathcal{U}(\mathbf{x}) \cap \Pi_V(\mathbf{x}, \mathbf{u}))$. The trajectory of the control solution offered by the nominal controller is also shown in Fig. 3(a) in red, where $\mathbf{u}(0)$ is the control input at the initial time instant $t(0)$ and $\mathbf{u}_{eq}$ is the control input at the equilibrium point. It can be seen that part of the trajectory of $\mathbf{u}$ may remain outside the solution space , which illustrates a situation where the trajectory of the nominal controller will be appropriately modified by the CLF constarints. This in turn guarantees the stability of the system without compromising performance, so that it remains close to the solution provided by the nominal controller, $\mathbf{k}(\mathbf{x})$.

Example 3: We demonstrate the impact of the proposed CLF-QP using the pendulum example, with the results shown in Fig. 4(a). To design the nominal controller, the LQR technique is used considering $Q_{lqr} = 10I_2$ and $R_{lqr} = 1$. To show the importance of the choice of CLF, again two cases are considered. In Case 1, $\mathbf{P}_l = I_2$ and in Case 2, $\mathbf{P}_l$ is calculated based on the protocol outlined in (10) and (11), assuming $\mathbf{Q} = I_2$. In this example, if we use the same initial state values as in Examples 1 and 2, the CLF constraint in Case 2 does not become activated. This means that the nominal controller already fulfills the stability requirements. This occurs because the CLF in Case 2 is based on the linearized closed-loop system and when initial states are near the equilibrium point, $\mathbf{x}_{eq}$, the CLF constraint is met. To demonstrate the impact of the CLF constraint, we select an initial condition further from equilibrium point, i.e. $\mathbf{x_0} = (1\ [\text{rad}], 9\ [\text{rad/s}])$. Since the initial conditions are far from the equilibrium point, the nominal controller does not satisfy the stability constraint at all times for either of the selected CLFs. However, modification of the nominal controller by the stability constraint is only needed for a short time if the CLF is selected properly (as in Case 2), as shown in the zoomed-in section of the initial moments in the inset of the CLF plot. This modification is needed for longer time in Case 1, resulting in more deviation from the nominal controller. However, in Case 2 this modification takes place for a fraction of time, resulting in a trajectory that closely aligns with the nominal controller.

C. Proposed CLF-CBF-QP Formulation

The CLF-QP formulation can be further modified to include the CBF constraint similar to (9), where the CBF acts as a hard constraint which must be strictly satisfied $\forall\ t \geq t(0)$, as shown in (13). In Fig. 3(b), the region corresponding to the CBF constraint $(\Pi_B(\mathbf{x}, \mathbf{u}))$ is shown in yellow, which represents the *safe region* and the intersection of this safe region with the solution space forms the *safe solution space*, that is, $(\mathcal{U}(\mathbf{x}) \cap \Pi_V(\mathbf{x}, \mathbf{u}) \cap \Pi_B(\mathbf{x}, \mathbf{u}))$. This illustration shows that two segments of the trajectory of $\mathbf{u}$ lie outside the safe solution space. One segment breaches the $(\mathcal{U}(\mathbf{x}) \cap \Pi_V(\mathbf{x}, \mathbf{u}))$ intersection but adheres to the CBF constraint. In contrast, the other segment of the trajectory violates the intersection $(\mathcal{U}(\mathbf{x}) \cap \Pi_B(\mathbf{x}, \mathbf{u}))$, but remains well within the CLF constraint. As mentioned in the previous section, the breach of the solution space can be managed by the CLF constraint. In order to address the breach of the safe solution space, the CBF constraint must adjust the trajectory accordingly. This would maintain the safety of the system while slightly compromising performance.

Proposed CLF-CBF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u} \in \mathbb{R}^m,\, \delta \in \mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u} - \mathbf{k}(\mathbf{x}))^T \mathbf{H} (\mathbf{u} - \mathbf{k}(\mathbf{x})) + p\,\delta^2 \tag{13}$$

$$\begin{aligned} \text{s.t.} \quad & \mathbf{u} \in \mathcal{U}(\mathbf{x}) && \text{(Input Constraint)} \\ & L_{\mathbf{f}}\mathcal{V}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{V}(\mathbf{x})\mathbf{u} \leq \delta && \text{(CLF Constraint)} \\ & L_{\mathbf{f}}\mathcal{B}(\mathbf{x}) + L_{\mathbf{g}}\mathcal{B}(\mathbf{x})\mathbf{u} \geq -\gamma \mathcal{B}(\mathbf{x}) && \\ & && \text{(CBF Constraint)} \end{aligned}$$

Example 4: The efficacy of the proposed CLF-CBF-QP formulation is illustrated by modifying the pendulum example in *Example 2* with the same initial condition, safety limits and CBF form with $\gamma = 5$. As for the CLF, we consider the one selected in Case 2 of *Example 3*. The results for this example are shown in Fig. 4(b). The figure shows that the CBF constraint becomes activated as θ begins to approach the safe limits within 0.1 seconds. This is reflected in the change in the control signal (u) at the beginning, which diverges from the solution of the nominal controller. Despite this

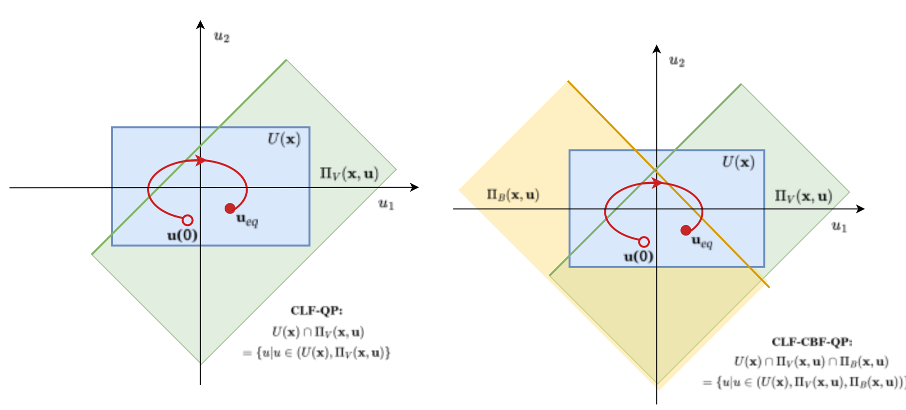

(a) **Proposed CLF-QP Formulation** (b) **Proposed CLF-CBF-QP Formulation**

Fig. 3: Schematic illustrating the intersection of the constraints for the proposed QP formulations and the trajectory of the control solution offered by the nominal controller

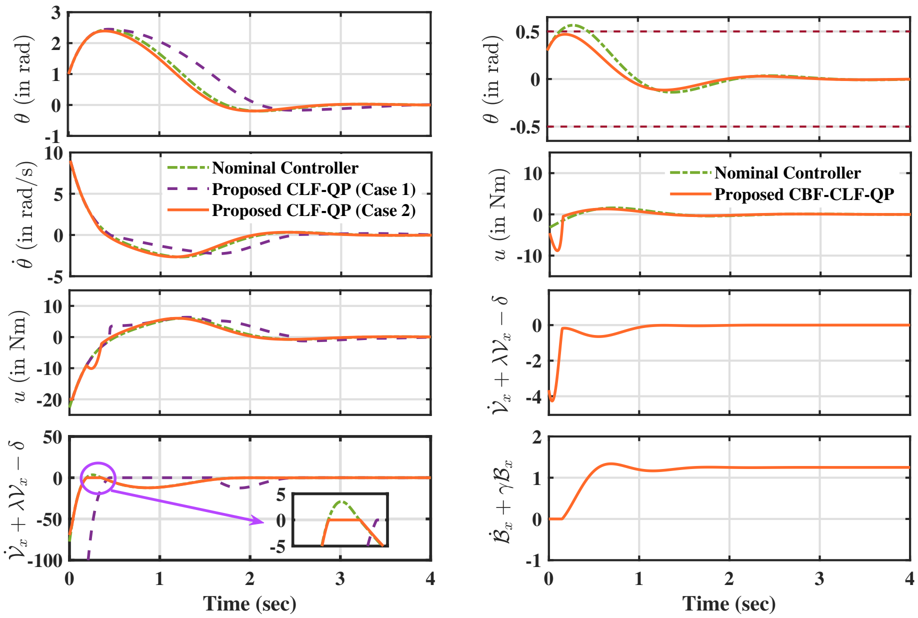

(a) **Proposed CLF-QP approach (Example 3)** (b) **Proposed CLF-CBF-QP approach (Example 4)**

Fig. 4: Comparison of state variables (θ and $\dot{\theta}$), control input (u) and the CLF-CBF constraints for the actuated pendulum for (a) Proposed CLF-QP approach (Example 3) (b) Proposed CLF-CBF-QP approach (Example 4)

necessary divergence of the control signal, the trajectory of θ remains closely aligned with that of the nominal controller, while meeting the safety limits. This shows that the proposed CLF-CBF-QP approach is capable of closely maintaining the system performance offered by the nominal controller without compromising safety.

IV. Simulation Results and Discussion

In this section, we numerically validate our proposed approach through a tracking problem involving an overhead crane (MIMO system) that demonstrates the capabilities of the proposed CLF-CBF-QP in (13), with a well-defined CBF, $\mathcal{B}(\mathbf{x})$, keeping the system within the safety limits.

A. Overhead Crane (System Description)

To study the augmented formulation with the CBF, we consider a 2D overhead crane (MIMO system) as illustrated in Fig. 5. The dynamics of OC can be expressed by the following nonlinear equations [18]–[20]:

$$\ddot{x_T} = \frac{J\sin(\theta)\left(g\cos\theta + L\dot{\theta}^2\right)}{d_\theta} + \frac{J_0}{Md_\theta}u_1 + \frac{R\sin\theta}{d_\theta}u_2 \tag{14a}$$

$$\ddot{L} = \frac{MR^2\left(g\cos\theta + L\dot{\theta}^2\right)}{d_\theta} - \frac{R^2\sin\theta}{d_\theta}u_1 - \frac{R\left(\frac{M}{m} + \sin^2\theta\right)}{d_\theta}u_2 \tag{14b}$$

$$\ddot{\theta} = \frac{-2d_\theta\dot{\theta}\dot{L} - g\sin\theta(J+J_0) - JL\sin\theta\cos\theta\dot{\theta}^2}{Ld_\theta} - \frac{J_0\cos\theta}{MLd_\theta}u_1 - \frac{R\sin\theta\cos\theta}{d_\theta}u_2 \tag{14c}$$

with, $J_0 = M\left(\frac{J}{m} + R^2\right)$ and $d_\theta = J_0 + J\sin^2(\theta)$. Here,

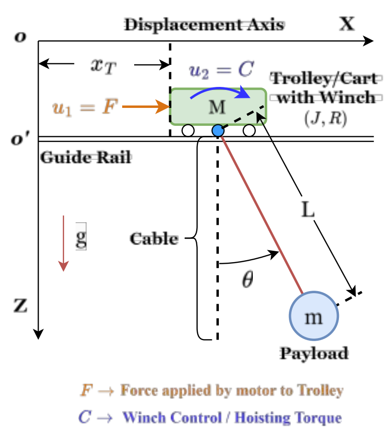

Fig. 5: Schematic representation of the overhead crane (OC) system

M and m represent the masses of the trolley/cart and the payload (in kg), respectively, L represents the length of the cable (in m), x_T is the position of the cart (in m), θ is the swing angle of the cable w.r.t. the vertical axis, g is the acceleration due to gravity, J is the moment of inertia of the winch mechanism (in kgm^2), R is the radius of the winch mechanism (in m), F is the horizontal force applied to the cart actuated by the first motor, which acts as the first control input (u_1) and C denotes the hoisting torque exerted on the winch by a second motor which acts as the second control input (u_2). For this case, the following values are chosen for the parameters: $m = 1$ kg, $M = 3$ kg, $L = 1.5$ m(initial), $R = 0.1$ m, $J = 1$ kgm^2 and $g =$

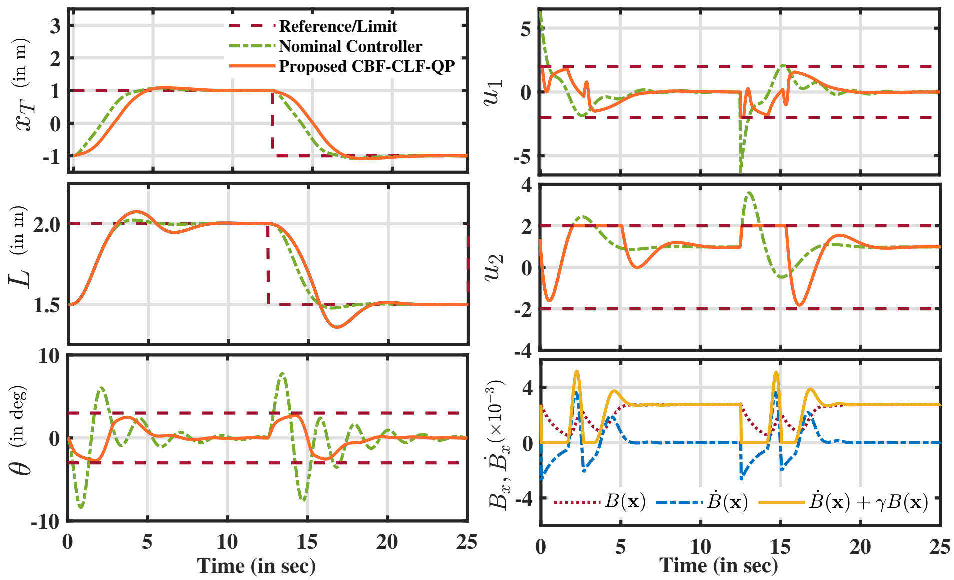

Fig. 6: Variation of state variables, control input, and the rate of change of the CLF for a 3D crane system with input constraints ($u_{1,lim} = \pm 2$N; $u_{2,lim} = \pm 2$Nm) and the swing angle safety limit ($\theta_{safe} = \pm 3°$).

9.81 ms^{-2}. The control objective is to move the suspended mass between two points, $P_1(x = -1m, L = 1.5m, \theta = 0)$ and $P_2(x = 1m, L = 2m, \theta = 0)$, with a focus on fast movement and quick settling. As the nominal controller will be designed based on the linearized model, the nonlinear model of the system is linearized as follows.

$$\dot{\mathbf{x}} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & \frac{Jg}{J_0} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{gJ_a}{L} & 0 & 0 & 0 \end{bmatrix} \mathbf{x} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{1}{M} & 0 \\ 0 & -\frac{RM}{mJ_0} \\ -\frac{1}{ML} & 0 \end{bmatrix} \mathbf{u} \tag{15}$$

where, the state vector $\mathbf{x} = [x_T \ L \ \theta \ \dot{x_T} \ \dot{L} \ \dot{\theta}]^T$ and $J_a = (J/J_0) + 1$. Since the Jacobian matrices are independent of both x_T and L, the linear model of (15) is valid at both P_1 and P_2.

B. Results and Discussion

This is a tracking control problem in which the reference trajectories for the position of the cart and the length of the cable are denoted by x_T^d and L^d, respectively. By defining two new state variables as $x_7 = \int(x_T - x_T^d)dt$ and $x_8 = \int(L - L^d)dt$, an LQI controller is designed as the nominal controller. The performance of this controller is shown in dotted green in Fig. 6 for the system outputs, (x_T, L, θ) and the control inputs (u_1, u_2). The red dotted lines show the reference trajectories x_T^d and L^d in their respective plots.

To choose the Lyapunov function, given that this is a tracking problem, the CLF presented in (10) needs to be modified to incorporate the state error, $\mathbf{x}_e$, as follows.

$$\mathcal{V}(\mathbf{x}_e) = \mathbf{x}_e^T \, \mathbf{P}_l \, \mathbf{x}_e \tag{16}$$

where $x_e = x_a - x_a^d$ in which x_a is the augmented state defined as $x_a = [x^T \ x_7 \ x_8]^T$. The value of $\mathbf{P}_l$ is found using the Lyapunov function described in (11), where, $\mathbf{A}_{cl}$ is the closed loop system matrix considering the nominal LQI controller designed before.

To select the CBF for this system, we consider that the safety objective is to limit the swing of the payload to $\pm 3°$, denoted as θ_{safe}. To satisfy this safety requirement, we must select a CBF that enforces constraints on both actuators, ensuring the system remains within the safe solution space. Consequently, we implement a CBF composed of two components inspired by the works of [4], [5]. The first term, $(\theta_{safe}^2 - \theta^2)$ ensures that $|\theta| \leq \theta_{safe}$ is not violated. Although in principle, the first term should be enough, however, when the boundary of CBF is reached, the required u to avoid entering the unsafe solution space, could be more than what the actuators allow. In order to avoid this situation, as suggested in [5], $\mathcal{B}(\mathbf{x})$ should include a second term which reduces the $\mathcal{B}(\mathbf{x})$ based on the actuator constraints and the angular velocity of the pendulum. Based on this, we represent the CBF for this case as,

$$\mathcal{B}(\mathbf{x}) = (\theta_{safe}^2 - \theta^2) - \frac{1}{2}\left(\frac{k_{b1}}{u_1^{max}} + \frac{k_{b2}}{u_2^{max}}\right)\dot{\theta}^2 \tag{17}$$

where, u_1^{max} and u_2^{max} are the maximum input limitation of the two actuators, which were chosen to be 2N and 2Nm, respectively. Also, the parameters k_{b1} and k_{b2} can be chosen depending on the system specifications.

Fig. 6 shows a performance comparison of the nominal LQI controller with the proposed CLF-CBF-QP approach considering the CLF in (16) and the CBF in (17). The plot in the top-left illustrates the system's position response (x_T) in comparison to the reference trajectory (x_T^d), shown in red dotted lines. Although the response generated by the proposed CLF-CBF-QP method aligns closely with that of the nominal LQI controller, it compromises the performance to maintain the safety and the input constraints, leading to a slightly slower response. This occurs as a direct consequence of the restrictions imposed on the control input u_1, as seen beside the position response. It can be observed from this figure that the nominal controller demands more control effort to reach an optimal performance. However, the proposed approach restricts the available control effort, leading to a slightly sluggish response. A pattern can be observed in the cable length (L), where its behavior exhibits increased oscillations around the desired trajectory (L^d) when utilizing the proposed method. This phenomenon can be linked to the constraints placed on u_2, as illustrated in the accompanying figure. Notably, u_2 behaves similarly for both controllers until it reaches the limit, after which a transient change appears in the proposed method's response.

However, as expected, the swing angle of the payload (θ) remains within the safe limit (θ_{safe}) of $\pm 3°$ and displays far fewer oscillations with the proposed approach compared to the nominal controller. This is attributed to the variations of the CBF, $\mathcal{B}(\mathbf{x})$ and the constraint, $\dot{\mathcal{B}}(\mathbf{x}) + \gamma\mathcal{B}(\mathbf{x})$, which remains positive at all times, as seen in the adjacent figure. As the swing angle approaches the boundary, it is important to observe that $\dot{\mathcal{B}}(\mathbf{x})+\gamma\mathcal{B}(\mathbf{x})$, approaches zero. This prevents any further increase in the swing angle, thereby ensuring safety limits are upheld. In conclusion, it can be stated that the proposed CLF-CBF-QP controller not only stabilizes the system, but also ensures that it remains safe.

V. Conclusion

This paper presents a novel safety-enhanced control framework that systematically augments nominal controllers with stability and safety constraints through the CLF-CBF-QP approach, without requiring the redesign of controllers as typically needed in the conventional CLF-CBF-QP approach. Using a systematic approach for selecting the CLF for a system with a nominal control, along with a properly chosen CBF in the QP framework, the proposed method offers a practical solution to enhance both safety and stability in non-linear systems. The applicability of the framework is demonstrated through a case study involving an overhead crane, showing its ability to ensure stability while addressing practical safety concerns.

In future endeavors, the framework offers potential for enhancing the safety and robustness of AI-based controllers [21], [22], which often face challenges of over-fitting to particular nominal scenarios, thus limiting their appeal for critical applications. By representing these controllers in a nonlinear state feedback form, the proposed framework can effectively manage unforeseen situations, preserving the integrity of the system. Therefore, this approach bridges a significant gap in control system design by providing a scalable, practical method for enhancing safety in both traditional and AI-driven controllers.

References

[1] A. D. Ames and M. Powell, *Towards the Unification of Locomotion and Manipulation through Control Lyapunov Functions and Quadratic Programs*. Heidelberg: Springer International Publishing, 2013, pp. 219–240. [Online]. Available: https://doi.org/10.1007/978-3-319-01159-2_12

[2] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[3] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[4] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in *2016 American Control Conference (ACC)*, 2016, pp. 322–328.

[5] J. Choi, F. Castañeda, C. J. Tomlin, and K. Sreenath, "Reinforcement learning for safety-critical control under model uncertainty, using control lyapunov functions and control barrier functions," 2020. [Online]. Available: https://arxiv.org/abs/2004.07584

[6] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 6271–6278.

[7] S.-C. Hsu, X. Xu, and A. D. Ames, "Control barrier function based quadratic programs with application to bipedal robotic walking," in *2015 American Control Conference (ACC)*, 2015, pp. 4542–4548.

[8] C. Wang, Y. Meng, Y. Li, S. L. Smith, and J. Liu, "Learning control barrier functions with high relative degree for safety-critical control," in *2021 European Control Conference (ECC)*, 2021, pp. 1459–1464.

[9] A. J. Taylor and A. D. Ames, "Adaptive safety with control barrier functions," in *2020 American Control Conference (ACC)*, 2020, pp. 1399–1405.

[10] S. Byeon and G. Park, "Safe and robust stabilization of uncertain nonlinear systems via control lyapunov-barrier function and disturbance observer: A preliminary study," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, 2023, pp. 6525–6532.

[11] M. Schneeberger, S. Mastellone, and F. Dörfler, "Advanced safety filter based on sos control barrier and lyapunov functions," 2025. [Online]. Available: https://arxiv.org/abs/2401.06901

[12] H. Dai, C. Jiang, H. Zhang, and A. Clark, "Verification and synthesis of compatible control lyapunov and control barrier functions," 2024. [Online]. Available: https://arxiv.org/abs/2406.18914

[13] H. Wang, K. Margellos, and A. Papachristodoulou, "Safe and stable filter design using a relaxed compatibitlity control barrier – lyapunov condition," 2024. [Online]. Available: https://arxiv.org/abs/2407.00414

[14] A. Alan, A. J. Taylor, C. R. He, A. D. Ames, and G. Orosz, "Control barrier functions and input-to-state safety with application to automated vehicles," *IEEE Transactions on Control Systems Technology*, vol. 31, no. 6, pp. 2744–2759, 2023.

[15] J. Zeng, B. Zhang, Z. Li, and K. Sreenath, "Safety-critical control using optimal-decay control barrier functions with guaranteed pointwise feasibility," in *2021 American Control Conference (ACC)*, 2021.

[16] J.-J. E. Slotine, W. Li *et al.*, *Applied nonlinear control*. Prentice hall Englewood Cliffs, NJ, 1991, vol. 199, no. 1.

[17] H. K. Khalil, *Nonlinear control*. Pearson New York, 2015.

[18] B. d'Andréa Novel and J. Lévine, *Modelling and Nonlinear Control of an Overhead Crane*. Boston, MA: Birkhäuser Boston, 1990, pp. 523–529. [Online]. Available: https://doi.org/10.1007/978-1-4612-4484-4_52

[19] J. Levine, *Analysis and control of nonlinear systems: A flatness-based approach*. Springer Science & Business Media, 2009.

[20] M. R. Mojallizadeh, B. Brogliato, and C. Prieur, "Modeling and control of overhead cranes: A tutorial overview and perspectives," *Annual Reviews in Control*, vol. 56, p. 100877, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1367578823000056

[21] W. Xiao, T.-H. Wang, R. Hasani, M. Chahine, A. Amini, X. Li, and D. Rus, "Barriernet: Differentiable control barrier functions for learning of safe robot control," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 2289–2307, 2023.

[22] Y. Yang, Y. Jiang, Y. Liu, J. Chen, and S. E. Li, "Model-free safe reinforcement learning through neural barrier certificate," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1295–1302, 2023.

VI. International Conference on Electrical, Computer and Energy Technologies (ICECET 2026)
6-9 July 2026, Rome-Italy

Ensuring safe roll dynamics in marine vessels based on a safety-enhanced control framework

Shouvik Chaudhuri
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0001-8957-5086

Hossein Ramezani
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0002-2273-1520

Jerome Jouffroy
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0002-3097-0679

***Abstract*—This paper presents a safety-enhanced control framework for ensuring safe roll dynamics in monohull marine vessels equipped with canting keels. Motivated by the regulatory shift introduced through the IMO's Second Generation Intact Stability Criteria (SGISC), dynamic stability limits associated with parametric rolling and excessive lateral acceleration are recast as enforceable constraints within a unified quadratic program (QP). The canting keel mechanism introduces additional challenges due to actuator saturation, electromechanical limits, and significant rotational inertia. A control architecture combining Control Lyapunov Functions (CLFs) and Control Barrier Functions (CBFs) is formulated to guarantee stability and forward invariance of the roll-angle safe set associated with parametric rolling. The excessive-acceleration criterion, being algebraically dependent on the control input, is incorporated into the set of admissible inputs as a state-dependent affine constraint. Numerical simulations under stochastic wave excitation demonstrate that the framework truncates resonance cycles during parametric rolling and actively bounds lateral accelerations under synchronous rolling conditions, while preserving baseline roll stabilization performance.**

***Keywords*—Dynamic stability, Marine vessels, Quadratic programming, Roll stabilization, Safety-critical control**

I. Introduction

In the maritime sector, ship stability and safety constitute fundamental prerequisites for the protection of life, property, and the marine environment, traditionally ensured through prescriptive design and operational criteria established by the International Maritime Organization (IMO) [1]. In recent years, the advancement of maritime autonomy and the increasing complexity of modern hull forms have necessitated a paradigm shift in the assessment of ship stability and safety. Historically, safety against roll instability, one of the most critical dynamic hazards, was evaluated primarily through static means by analyzing the righting arm curve (GZ) in calm water to ensure a vessel possessed sufficient energy to resist heeling moments. However, the dynamic realities of the seaway, stochastic wave patterns, nonlinear fluid-structure interactions, and coupling between degrees of freedom, render static analyses insufficient for predicting dynamic failure modes such as parametric roll, pure loss of stability, excessive accelerations in waves, or catastrophic failures such as capsizing or broaching. This inadequacy has led to the development of the *Second Generation Intact Stability Criteria* (SGISC) by the International Maritime Organization (IMO) [2], [3] and supported by experimental procedures from the International Towing Tank Conference (ITTC) [4], focusing on physics-based vulnerability assessment rather than static compliance.

For maritime researchers, these regulatory criteria represent more than mere compliance targets; they delineate the fundamental mathematical boundaries of the vessel's safe operational envelope. Within this context, the SGISC dynamic stability criteria, particularly those associated with Parametric Rolling and Excessive Lateral Acceleration, can be employed as a rigorous mathematical foundation for real-time control constraints. Rather than serving exclusively as post hoc compliance verifications, these regulatory limits explicitly define the vessel's safe operational domain under time-varying and dynamically coupled operating conditions. Embedding such limits as inequality constraints within advanced control schemes enables the controller to actively enforce these safety margins; in other words, control actions are continuously optimized to avoid violating dynamic stability limits, thereby upholding IMO-defined criteria in real time.

While these mathematical boundaries provide a rigorous theoretical foundation for defining safety, the control architectures required to actively enforce them in real-time largely remain in their infancy in the maritime domain, particularly when juxtaposed with the automotive or aerospace sectors, where active safety systems are ubiquitous and governed by rigorous functional safety standards. In contrast, marine roll stabilization strategies, ranging from passive tanks to active fins and gyroscopes [5]–[7], have traditionally prioritized performance (i.e., comfort and seasickness reduction). This performance-centric paradigm presents a hidden risk. Conventional stabilizers effectively reduce roll variance under nominal conditions but lack the mathematical machinery to

This paper is part of Project SAFEMARVEL which is funded by the Den Danske Maritime Fond (DDMF) (Project nr. 95-45002-3410157)

strictly enforce safety constraints. Under extreme sea states or aggressive maneuvering, a performance-optimizing controller might inadvertently saturate actuators or induce destabilizing moments, effectively driving the vessel toward the very failure modes the SGISC seeks to prevent. The result is a *critical gap*: vessels may exhibit adequate performance on average yet remain vulnerable to dynamic failure modes such as parametric rolling or excessive acceleration during transients or severe disturbances. This gap is particularly evident in vessels utilizing actuation mechanisms with significant inertia, such as the canting keel considered in this study. Unlike active fins, which rely on lift proportional to speed, a canting keel generates a restoring moment by shifting the vessel's center of gravity [8]. While effective at low speeds, the keel's large inertia and electromechanical limits introduce complex constraints. Violating these constraints can lead to mechanical failure or, paradoxically, destabilize the hull.

In this context, Safety-critical control (SCC) architectures offer a promising pathway to bridge this gap by embedding these regulatory constraints directly into the control loop. Such frameworks enable real-time enforcement of stability criteria, transforming safety margins from passive guidelines into active invariants of the control policy. This paper presents a safety-enhanced control framework for a monohull vessel equipped with a canting keel, embedding the SGISC regulatory criteria as enforceable constraints within the control architecture. Specifically, we recast the stabilization problem as a safety-critical control task, using a nonlinear dynamic model of a canting-keel vessel [8], [9]. Within this framework, a Control Barrier Function (CBF) enforces forward invariance of the roll-angle safe set corresponding to the parametric rolling limit, while the excessive-acceleration criterion is embedded directly into the admissible input set as a state-dependent constraint. Together, these mechanisms ensure that the vessel respects both dynamic stability envelopes under admissible conditions. The resulting controller acts as an intelligent safeguard, intervening only when performance optimizations risk breaching safety boundaries, unifying compliance, safety, and performance in a single structure.

II. System Description

This study focuses on the roll stabilization of a monohull marine vessel equipped with a canting keel mechanism. The system configuration comprises a rigid hull and a keel bulb actuated by a DC servo motor via a hinge point at the hull's bottom as shown in Fig. 1. We adopt the coordinate systems and lumped parameter modeling approach detailed in [8].

A. Vessel Model

Assuming negligible coupling between roll and other degrees of freedom (1-DoF), as is common in roll stabilization studies near resonance where roll dominates, the nonlinear roll dynamics are governed by hydrodynamic forces, hydrostatic restoring moments, and external torques. The equation of motion, as defined in [9], [10] is given as:

$$(I_x + \Delta I_x)\ddot{\phi} + B_\phi(\dot{\phi}) + C_\phi(\phi) = \tau_k + \tau_d + \tau_{wi} \tag{1}$$

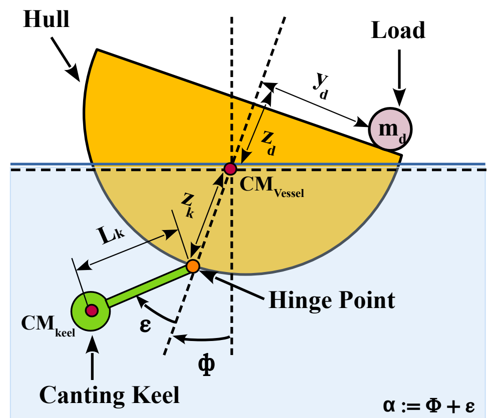

Fig. 1. Schematic representation of a monohull marine vessel equipped with a canting keel mechanism

where ϕ and $\dot{\phi}$ denote the roll angle and roll rate, respectively. The terms I_x and ΔI_x represent the moment of inertia and the hydrodynamic added mass. The damping term $B_\phi(\dot{\phi}) = K_p\dot{\phi} + K_{p|p|}\dot{\phi}|\dot{\phi}|$ captures linear and quadratic drag, while $C_\phi(\phi) = \rho g \nabla \overline{GM}_m \phi + K_{\phi^3}\phi^3$ represents the linear and nonlinear hydrostatic restoring moments. The right-hand side of (1) aggregates the moments acting on the hull, as discussed below:

1. **Keel Torque** (τ_k): The control torque generated by the canting keel, which includes the gravitational moment of the keel bulb (τ_w), the paddling hydrodynamics (τ_p), and the reaction torque from the servo motor (τ_m), as defined in [8], is given below:

$$\begin{cases} \tau_k = \tau_w - \tau_p - \tau_m \\ \tau_w = -\left(m_k - \rho\nabla_k\right) g\left(z_k \sin\phi + L_k \sin(\varepsilon + \phi)\right) \\ \tau_p = K_k(\dot{\varepsilon} + \dot{\phi}) + K_{k|k|}(\dot{\varepsilon} + \dot{\phi})|\dot{\varepsilon} + \dot{\phi}| + \Delta m_k(\ddot{\varepsilon} + \ddot{\phi}) \\ \tau_m = J_{km}\ddot{\varepsilon} + B_{km}\dot{\varepsilon} = k_t I_a \end{cases} \tag{2}$$

The constants, m_k and ∇_k, denote the mass and displaced volume of the keel system, respectively. The geometric parameters z_k and L_k represent the distances from the vessel's center of mass to the keel hinge and from the hinge to the keel's mass–buoyancy center, respectively (see Fig. 1). Also, $\dot{\varepsilon}$ denotes the keel angular rate, while K_k and $K_{k|k|}$ are the linear and nonlinear hydrodynamic drag coefficients. The term Δm_k accounts for the added mass associated with keel motion.

2. **Load Torque** (τ_d): A time-varying disturbance torque arising from a shifting deck load, modeled as a mass m_d with a time-varying lateral position $y_d(t)$ and constant vertical position from the center of mass of the vessel $z_d(t)$, following the formulation in [9], shown below:

$$\tau_d = m_d g\left(z_d \sin\phi + y_d \cos\phi\right) \tag{3}$$

3. **Wave-induced Torque** (τ_{wi}): The wave-induced disturbance is treated as a stochastic process defined by

the JONSWAP spectrum ($S(\omega)$) [10]. Under a lumped-parameter approximation and assuming a block-shaped vessel [11], the wave-induced torque is modeled as a superposition of N regular wave components, with amplitudes obtained from the wave spectral density defined in [8], as given below:

$$\tau_{wi}(t) = \alpha_\tau \sum_{j=1}^{N_f}\sum_{i=1}^{N_d} k_j^2\ sin(2\psi_i)\ A_j\ \cos(\omega_j t + \varepsilon_j \\ -k_j\ x\cos(\psi_i) - k_j y\sin(\psi_i)) \tag{4}$$

B. Actuator Dynamics and Control Form

The keel angle (ε) is actuated by a DC servo motor. Neglecting armature inductance and assuming a rigid gearbox, the actuator dynamics can be described as:

$$\ddot{\varepsilon} = \alpha_1\varepsilon + \alpha_2\dot{\varepsilon} + \beta_\varepsilon u \tag{5a}$$

where u is the commanded keel angle ε_d (control input), and $\alpha_{1,2}, \beta_\varepsilon$ are lumped parameters derived from the motor's electromechanical constants (inertia J_{km}, damping B_{km}, back-emf k_b, torque constant k_t, armature resistance R_a and current I_a) and control gains (k_{pm} and k_{dm}), as shown below:

$$\alpha_1 = \frac{-k_t k_{pm}}{J_{km}R_a}, \quad \beta_\varepsilon = \frac{k_t k_{pm}}{J_{km}R_a} \tag{5b}$$

$$\alpha_2 = \frac{-k_t(k_b + k_{dm}) - R_a B_{km}}{J_{km}R_a} \tag{5c}$$

By substituting the actuator dynamics (5a) into the derivative of (1), the coupled system can be expressed as a fourth-order control-affine nonlinear system composed of two interconnected subsystems: a second-order nonlinear vessel-roll subsystem with states $(\phi, \dot{\phi})$ driven by the net applied torque, and a second-order linear actuator–keel subsystem with states $(\varepsilon, \dot{\varepsilon})$ driven by the commanded keel angle (ε_d), which also serves as the overall system input. Accordingly, the dynamics can be expressed in the control-affine form suitable for safety-critical controller synthesis,

$$\ddot{\phi} = f(\mathbf{x}) + g(\mathbf{x})u + d(t), \tag{6}$$

where the state vector is $\mathbf{x} = [\phi, \dot{\phi}, \varepsilon, \dot{\varepsilon}]^T$, $f(\mathbf{x})$ captures the unforced hydrodynamics and keel interaction terms, $g(\mathbf{x})$ maps the control input through the vessel inertia, and $d(t)$ represents external wave and load disturbances.

III. Safety-enhanced Control Framework

The objective of the proposed control framework is to reconcile the competing requirements of operational performance (roll reduction for comfort) and dynamic survivability (adherence to safety limits). Traditional stabilizers prioritize performance under nominal conditions but often lack the mathematical machinery to strictly enforce safety boundaries during extreme events. To bridge this gap, we adopt a safety-critical control framework that acts as a supervisory filter, augmenting a performance-tuned nominal controller with formal stability and safety guarantees. Here, the nominal controller refers to an existing control law, typically designed and tuned using conventional methods that achieve the desired performance objectives within a predefined operating envelope, without explicitly accounting for hard safety constraints. This section first transforms the relevant SGISC failure modes into explicit mathematical stability and safety constraints. Subsequently, it defines the augmentation framework, which leverages a Quadratic Program (QP) to minimally alter a nominal performance controller only when the vessel approaches the boundaries of the safe operational envelope.

A. Transformation of IMO Dynamic Stability Criteria into Safety Constraints

To define the vessel's safe set, the vulnerability criteria outlined in the IMO SGISC [2] must be recast as state-dependent inequality constraints. Within the framework of Control Barrier Functions (CBF), the system is considered safe if its state ($\mathbf{x}$) remains within a safe set $\mathcal{C}$, defined as the superlevel set of a continuously differentiable function $\mathcal{B}(\mathbf{x})$:

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n : \mathcal{B}(\mathbf{x}) \geq 0\} \tag{7}$$

Based on the roll dynamics identified in [9], we address the critical dynamic stability failure modes relevant to roll dynamics: Parametric Rolling and Excessive Lateral Acceleration. As detailed below, parametric rolling admits a standard CBF formulation, whereas the lateral acceleration constraint is incorporated into the admissible input set.

1) Parametric Rolling Constraint ($\mathcal{B}_{PR}$): Parametric rolling is a dynamic instability characterized by the amplification of roll motion due to the periodic variation of the ship's transverse stability (ΔGM) as it encounters longitudinal waves [2]. This phenomenon is particularly critical for vessels with fine bow flares and wide transom sterns. While the IMO Level 1 vulnerability criterion requires that stability variations remain below a damping-defined threshold, real-time control requires enforcing constraints on the maximum roll angle $\phi_{PR,limit}$ to prevent large-amplitude resonance cycles predicted by Mathieu instability regions [4]. This leads us to the simplified barrier function, $\mathcal{B}_{PR}$, of the form:

$$\mathcal{B}_{PR}(\mathbf{x}) = \phi_{PR,limit} - |\phi| \geq 0 \tag{8}$$

where $\phi_{PR,limit}$ is the maximum safe roll angle (typically 25° or the angle of vanishing stability) as defined in the SGISC Level 2 assessment checks [2].

While $\mathcal{B}_{PR}$ defines the safe set in the configuration space, a static barrier function does not account for the rotational momentum of the vessel or the finite control authority of the heavy 4-ton canting keel. To ensure forward invariance in high-sea states, we implement a *Braking Barrier Function* ($\mathcal{B}_{PRb}$) that incorporates a velocity-dependent stopping distance derived from the physical limits of the restorative torque. By utilizing a squared formulation to ensure a smooth safety manifold, the augmented barrier function is defined as:

$$\mathcal{B}_{PRb}(\phi, \dot{\phi}) = (\phi_{PR,limit}^2 - \phi^2) - k_r\left(T_h\dot{\phi}\right)^2 - \\ k_{bd}\left(\frac{\dot{\phi}^2}{2\ddot{\phi}_{max}}\right)^2 \geq 0 \tag{9}$$

where k_r and k_{bd} are conservative braking gains, and $\ddot{\phi}_{max}$ is the maximum effective roll acceleration that the canting keel can generate against the buoyancy of the bulb. This provides a robust safety buffer as the system state nears the boundary of the admissible control set.

2) Excessive Lateral Acceleration Constraint: Excessive lateral acceleration is a dynamic failure mode primarily associated with synchronous rolling in beam sea conditions ($\mu \approx 90°$). Unlike parametric rolling, which arises from periodic variations in transverse stability (ΔGM), synchronous rolling is a forced-resonance phenomenon occurring when the wave encounter frequency approaches the vessel's natural roll frequency. Vessels with relatively high metacentric height (GM) exhibit stiff restoring characteristics that, combined with roll angular acceleration, generate significant lateral accelerations (a_y) at elevated locations such as the bridge or cargo deck. The lateral acceleration at a height h_r above the roll axis is expressed as,

$$a_y = k_L \left(g \sin\phi + h_r \ddot{\phi} \right), \tag{10}$$

where $g \sin\phi$ is the gravity projection component, $h_r\ddot{\phi}$ captures the inertial contribution due to roll angular acceleration, and k_L is a dimensionless coupling factor. Since $\ddot{\phi}$ depends on u through the roll dynamics (6), the lateral acceleration is affine in the control input,

$$a_y(\mathbf{x}, u) = k_L \big(g \sin\phi + h_r \left[f(\mathbf{x}) + g(\mathbf{x})\, u + d(t) \right] \big). \tag{11}$$

According to the IMO SGISC criterion [2], a_y at critical locations must not exceed $R_{EA,limit} = 4.64$ m/s^2. This regulatory threshold is therefore enforced as a pair of linear inequality constraints on the control input,

$$-R_{EA,limit} \;\leq\; a_y(\mathbf{x}, u) \;\leq\; R_{EA,limit}, \tag{12}$$

which are absorbed into the state-dependent admissible input set $\mathcal{U}(\mathbf{x})$ of the QP formulated in the next section.

B. Nominal Controller Augmentation via CLF-CBF-QP

For many practical applications, nominal controllers tuned for performance (e.g., PID or LQR) provide local stability but may exhibit degraded performance or instability when working conditions move away from the equilibrium point. Conversely, conventional safety-critical algorithms often prioritize energy optimization over performance. To address these issues, the nominal controller represented in nonlinear state feedback form is retained as $\mathbf{k}(\mathbf{x})$ as the performance baseline. Stability and safety guarantees are systematically added thereafter.

1) Systematic Selection of the CLF for stability: We seek a Control Lyapunov Function (CLF) that fulfills the conditions of systematic selection and minimum interference with the nominal solution $\mathbf{k}(\mathbf{x})$. To achieve this, we leverage the structure of the closed-loop system formed by the nominal controller. By linearizing the nonlinear system around the equilibrium point, we compute the Jacobian of the system dynamics $\mathbf{f}(\mathbf{x})$, input matrix $\mathbf{g}(\mathbf{x})$, and nominal controller $\mathbf{k}(\mathbf{x})$, yielding $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{K}$, respectively. The closed-loop system matrix is defined as $\mathbf{A}_{cl} = \mathbf{A} + \mathbf{BK}$. The Lyapunov function is then selected as [12], [13]:

$$\mathcal{V}(\mathbf{x}) = \mathbf{x}^T \mathbf{P}_l \mathbf{x} \tag{13}$$

where $\mathbf{P}_l$ is found by solving the Lyapunov equation,

$$\mathbf{P}_l \mathbf{A}_{cl} + \mathbf{A}_{cl}^T \mathbf{P}_l = -\mathbf{Q} \tag{14}$$

with $\mathbf{Q} = \alpha \mathbf{I}$. This choice ensures that the CLF constraint preserves the nominal controller's solution in the vicinity of the equilibrium while providing global stability guarantees.

2) Safety-Enhanced QP Formulation: The safety-enhanced control problem is formulated as a state-dependent QP that minimizes the deviation of the actual control input u from the nominal control input $\mathbf{k}(\mathbf{x})$. Since the objective function aims to maintain performance close to the nominal controller, the CLF is no longer required to focus on exponential convergence rates. Thus, the $\lambda\mathcal{V}(\mathbf{x})$ term is eliminated from the stability constraint [14].

To address safety, the CBF constraint derived from the parametric-rolling barrier $\mathcal{B}_{PRb}$ in (9) guarantees forward invariance of the roll-angle safe set. The excessive-acceleration criterion from (12) is absorbed into the admissible input set, which is now state-dependent:

$$\mathcal{U}(\mathbf{x}) = \Big\{ u \in \mathbb{R}^m : u_{\min} \leq u \leq u_{\max}, \; |a_y(\mathbf{x}, u)| \leq R_{EA,\text{limit}} \Big\} \tag{15}$$

where $u_{\min}$ and $u_{\max}$ are the actuator saturation bounds and the acceleration constraint follows from (11). Building on the nominal controller and the selection of the CLF and CBF, the safety-enhanced QP is formulated as given in [14].

Proposed Safety-Enhanced QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\, \delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u} - \mathbf{k}(\mathbf{x}))^T \mathbf{H} (\mathbf{u} - \mathbf{k}(\mathbf{x})) + p\,\delta^2 \tag{16}$$

$$\begin{aligned} \text{s.t.} \quad & \mathbf{u} \in \mathcal{U}(\mathbf{x}) && \text{(Admissible Input Set)} \\ & L_\mathbf{f}\mathcal{V}(\mathbf{x}) + L_\mathbf{g}\mathcal{V}(\mathbf{x})\mathbf{u} \leq \delta && \text{(CLF Constraint)} \\ & L_\mathbf{f}\mathcal{B}_{PRb}(\mathbf{x}) + L_\mathbf{g}\mathcal{B}_{PRb}(\mathbf{x})\mathbf{u} \geq -\gamma\mathcal{B}_{PRb}(\mathbf{x}) && \text{(CBF Constraint)} \end{aligned}$$

---

where $\mathcal{U}(\mathbf{x})$ is defined in (15) and encompasses both actuator saturation limits and the lateral acceleration bound; $\mathbf{H}$ is a weighting matrix; δ is a slack variable with penalty weight p that relaxes the stability constraint; and $\gamma > 0$ is a class-$\mathcal{K}$ gain for the CBF [14].

IV. Simulation Results and Discussion

The proposed safety-enhanced control framework is validated through two deterministic scenarios: (i) parametric rolling induced by periodic metacentric-height modulation, and (ii) excessive lateral acceleration during aggressive recovery from a large initial roll angle. Each scenario compares three controllers: open loop (u=0), a nominal LQR state-feedback controller, and the safety-enhanced QP formulation of Section III.

Simulations are conducted in MATLAB using the nonlinear roll-keel model described in Section II with fixed-step RK4 integration at $\Delta t = 5$ ms. Since the focus of this study is on the safety filter's interaction with the dynamic stability constraints, wave excitation is omitted (τ_{wi}=0) and the instability mechanisms are activated deterministically: parametric rolling is triggered by direct modulation of the restoring coefficient $C_1(t)$, while the excessive-acceleration scenario uses a large initial displacement. The physical parameters are summarized in Table I.

The nominal LQR controller is deliberately tuned with moderate gains ($\mathbf{Q} = \mathrm{diag}(50, 10, 1, 1)$, R=20) rather than maximizing roll-reduction performance. This design choice is motivated by three considerations: (i) energy efficiency, since aggressive keel actuation increases power consumption and mechanical wear; (ii) robustness, because high-gain controllers amplify modeling errors and sensor noise; and (iii) the fact that parametric rolling is a rare event triggered only by specific encounter-frequency and sea-state combinations. An aggressive nominal controller would suppress parametric rolling through increased closed-loop damping alone, thereby masking the very phenomenon the CBF is designed to handle.

Scenarios 1 and 2 address physically distinct failure modes that are unlikely to occur simultaneously, and accordingly employ different parameter sets to activate each phenomenon within the same vessel family, as summarized in Table II.

TABLE I
COMMON VESSEL MODEL PARAMETERS (SCENARIOS 1–2)

Parameter	Value
Monohull Vessel	
Roll moment of inertia (I_x)	125,000 kgm^2
Hydrodynamic added mass (ΔI_x)	31,250 kgm^2
Displacement mass (∇)	20,000 kg
Mean metacentric height ($\overline{GM}_m$)	0.75 m
Quadratic drag in roll ($K_{\dot{\phi}\|\dot{\phi}\|}$)	5×10^3 kgm^2
Keel Mechanism	
Keel system mass (m_k)	6,000 kg
Keel displaced volume (∇_k)	2.0 m^3
CoM-to-hinge distance (z_k)	1.4 m
Drive Motor	
Torque constant (k_t)	100 Nm/A
Back-emf constant (k_b)	375 Vs/rad
Armature resistance (R_a)	1 Ω
Motor inertia (J_{km})	20×10^3 kgm^2
Motor damping (B_{km})	100 Nms/rad
Safety & Control	
Commanded keel angle bounds (u)	±45°
IMO roll-angle limit ($\phi_{PR,limit}$)	25°
IMO lateral accel. limit (R_{EA})	4.64 m/s^2
CBF class-$\mathcal{K}$ gain (γ)	1.0

TABLE II
SCENARIO-SPECIFIC PARAMETERS
(VALUES THAT DIFFER ACROSS SCENARIOS)

Parameter	Sc. 1	Sc. 2	Unit
Linear roll damping (K_p)	4×10^4	1×10^5	kgm^2/s
Keel arm length (L_k)	2.5	3.5	m
Bridge height (h_r)	2.5	7.5	m
LQR $\mathbf{Q}$ diagonal	(50,10,1,1)	(500,10,1,1)	–
LQR R	20	2	–
Scenario 1 only			
GM variation (ΔGM)	0.70		m
Encounter freq. (ω_e)	$2\omega_{n,OL}$		rad/s
Initial roll (ϕ_0)	5		deg
Sim. duration	120		s
Scenario 2 only			
Initial roll (ϕ_0)	30		deg
Sim. duration	10		s

A. Parametric Rolling Suppression (Scenario 1)

Parametric rolling is triggered by modulating the hull restoring coefficient as $C_1(t) = C_1[1 + (\Delta GM/\overline{GM}_m)\cos(\omega_e t)]$ with ΔGM=0.70 m and ω_e=$2\omega_{n,OL}$, starting from $\phi(0)$=5°. Floquet analysis confirms that both the open-loop and closed-loop (LQR) plants are Mathieu-unstable at this excitation depth, with Floquet multipliers of 1.175 and 1.052, respectively. The moderate LQR adds damping but does not eliminate the instability; it merely slows its growth rate. This is by design: an aggressive LQR would suppress parametric rolling through high closed-loop damping, rendering the CBF superfluous for this scenario.

Fig. 2 shows the results over a 120 s simulation. In the upper panel, both the open-loop and nominal responses exhibit growing oscillations that eventually exceed the ±25° IMO limit, with the open-loop diverging faster. The safety-enhanced controller, by contrast, enforces the barrier constraint $\mathcal{B}_{PRb} \geq 0$ and keeps the roll angle within the prescribed limits throughout. Importantly, the safety filter does not fully damp the oscillations; rather, it bounds their amplitude at the safety boundary, intervening only when needed. The middle panel shows the keel commands; the zoomed inset reveals that the nominal and safety-enhanced inputs diverge as the roll approaches the limit, with the safety filter applying larger corrective commands to prevent boundary violations. The bottom panel confirms that the barrier function $\mathcal{B}_{PRb}$ remains non-negative under the safety-enhanced controller, while it goes negative under both the open-loop and nominal controllers, indicating breach of the safe set.

B. Excessive Lateral Acceleration Mitigation (Scenario 2)

This scenario evaluates the acceleration constraint on a larger vessel configuration (Table II) with the bridge at h_r=7.5 m, a longer keel arm (L_k=3.5 m), and higher hydrodynamic damping (K_p=10^5) representative of a 30–40 m hull. After a large wave impact leaves the vessel at $\phi(0)$=30°, the LQR drives the keel to return to upright.

Fig. 3 presents the results. The open-loop response, governed solely by passive restoring and damping forces, recovers slowly with a peak $|a_y|$ that remains below R_{EA}. The nominal LQR, however, commands the keel to accelerate recovery,

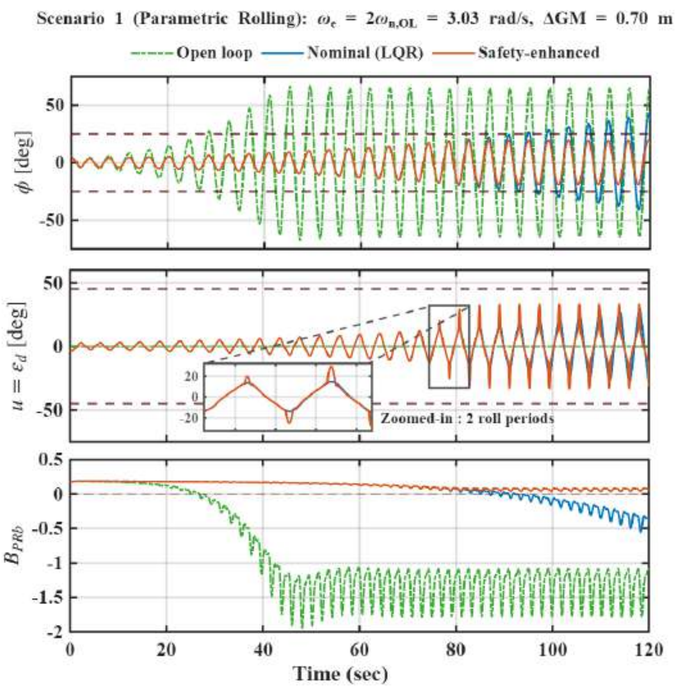

Fig. 2. Scenario 1: Parametric rolling suppression. Top: roll angle with $\pm 25°$ IMO limit. Middle: commanded keel angle with zoomed inset showing the divergence between nominal and safety-enhanced inputs over two roll periods. Bottom: Barrier function $\mathcal{B}_{PRb}$ (negative indicates safety-set violation).

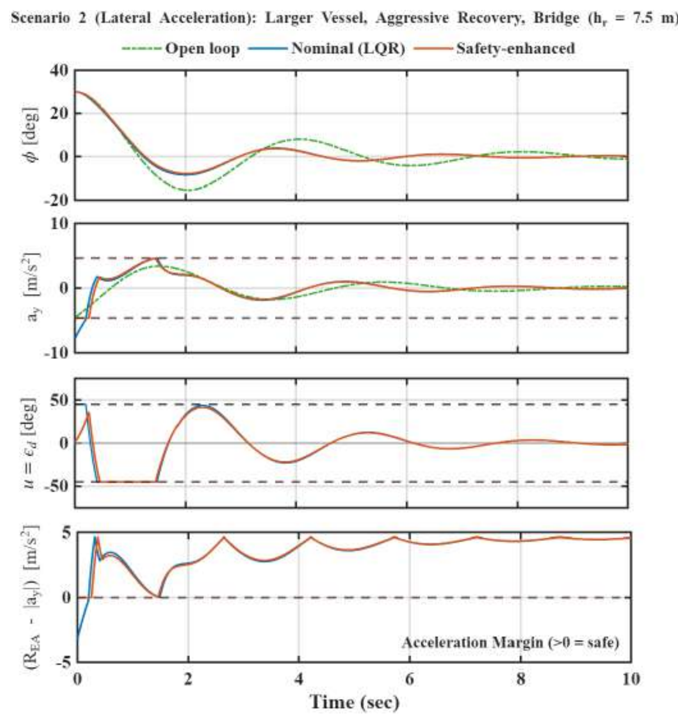

Fig. 3. Scenario 2: Lateral acceleration mitigation (larger vessel, h_r=7.5 m). Top: roll angle. Second: lateral acceleration with $\pm R_{EA}$ limit. Third: keel angle with actuator bounds ($\pm 45°$). Bottom: acceleration margin (> 0 = safe).

causing the lateral acceleration to significantly exceed the IMO limit during the initial transient. The safety-enhanced controller recognizes the imminent violation and backs off the keel command, keeping $|a_y| \leq R_{EA}$ at every control update. When the acceleration constraint is tighter than the actuator limits, the safety filter reduces the keel command below what the nominal LQR requests, trading slightly slower roll recovery for guaranteed compliance with the IMO criterion.

V. Conclusions

This paper presents a safety-enhanced control framework for monohull vessels equipped with canting keels, integrating the IMO's Second Generation Intact Stability Criteria (SGISC) directly into a real-time control architecture. Parametric rolling is enforced via a Control Barrier Function (CBF), which guarantees forward invariance of the roll-angle safe set, ensuring that the vessel trajectory never exits the IMO-prescribed envelope. The excessive lateral acceleration criterion is incorporated into the state-dependent admissible input set $\mathcal{U}(\mathbf{x})$, which jointly enforces actuator saturation limits and the IMO acceleration threshold at every control update. Numerical simulations demonstrate that the framework effectively modulates control inputs to prevent safety-critical breaches that nominal controllers fail to address. The results also highlight a fundamental challenge when applying CBFs to heavy maritime machinery: actuator latency and significant rotational inertia can cause the barrier function to approach its boundary during aggressive transients, even though the safety guarantee is mathematically satisfied.

References

[1] International Maritime Organization, "Revised guidelines for formal safety assessment (FSA) for use in the IMO rule-making process," International Maritime Organization, London, UK, Circular MSC-MEPC.2/Circ.12/Rev.2, 2018.

[2] International Maritime Organization , "Interim guidelines on second generation intact stability criteria," International Maritime Organization, London, UK, Circular MSC.1/Circ.1627, 2020.

[3] M. A. S. Neves, V. L. Belenky, J. O. De Kat, K. Spyrou, and N. Umeda, "Contemporary ideas on ship stability and capsizing in waves," 2011.

[4] International Towing Tank Conference, "Stability," ITTC Recommended Procedures and Guidelines, Tech. Rep. 7.5-02-07-04, 2017, rev. 02.

[5] R. Moaleji and A. R. Greig, "On the development of ship anti-roll tanks," *Ocean Engineering*, vol. 34, no. 1, pp. 103–121, 2007.

[6] N. C. Townsend and R. A. Shenoi, "Control strategies for marine gyrostabilizers," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 2, pp. 243–255, 2013.

[7] R. Li, T. Li, W. Bai, and X. Du, "An adaptive neural network approach for ship roll stabilization via fin control," *Neurocomputing*, vol. 173, pp. 953–957, 2016.

[8] H. Ramezani, S. Chaudhuri, J. Jouffroy, A. Baurichter, and S. Hansen, "On roll stabilization using a canting keel," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 2, pp. 539–553, 2025.

[9] S. Chaudhuri, H. Ramezani, and J. Jouffroy, "Modelling and control of a canting keel-based ship roll stabilization system for crane operations," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2022, pp. 1904–1909.

[10] T. I. Fossen, *Handbook of marine craft hydrodynamics and motion control*. John Wiley & Sons, 2011.

[11] C. G. Kallstrøm, "Identification and adaptive control applied to ship steering," Lund University, Sweden, Tech. Rep., 1982.

[12] J.-J. E. Slotine, W. Li *et al.*, *Applied nonlinear control*. Prentice hall Englewood Cliffs, NJ, 1991, vol. 199, no. 1.

[13] H. K. Khalil, *Nonlinear control*. Pearson New York, 2015.

[14] S. Chaudhuri, H. Ramezani, and J. Jouffroy, "Remarks on augmenting stability and safety in nominal controllers: a practical perspective," in *2025 European Control Conference (ECC)*, 2025, pp. 2121–2128.

Learning Marine Model Uncertainties for Real-Time Safety-Enhanced Roll Stabilization

Shouvik Chaudhuri
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0001-8957-5086

Hossein Ramezani
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0002-2273-1520

Jerome Jouffroy
Centre for Industrial Mechanics (CIM)
University of Southern Denmark (SDU)
Sønderborg, Denmark
ORCID ID: 0000-0002-3097-0679

***Abstract*—This paper presents a reinforcement learning (RL)-based safety-enhanced control framework for the roll dynamics of monohull marine vessels equipped with high-inertia canting keels. Control Barrier Function (CBF) and Control Lyapunov Function (CLF) architectures can enforce the IMO Second Generation Intact Stability Criteria (SGISC) as real-time constraints, but their forward-invariance guarantee holds only for the model they are synthesized on, and marine roll damping is trustworthy to no better than ± 20–30%. We quantify the consequence: under parametric excitation, a CLF-CBF-QP safety filter built on an Ikeda-consistent nominal model loses the invariance of the safe set. To close this gap, a Deep Deterministic Policy Gradient (DDPG) agent learns the mismatch terms (Δ_1, Δ_2) between the nominal and true Lie derivatives of the CLF and CBF constraints, together with the analogous correction to the lateral-acceleration input constraint. The learned terms enter the quadratic program as state-dependent affine corrections, preserving convexity and real-time solvability. In numerical case studies on a canting-keel research vessel, the RL-augmented filter restores barrier invariance under parametric rolling and re-establishes compliance with the SGISC excessive-acceleration threshold, matching an oracle filter that knows the true plant.**

***Keywords*—Dynamic stability, Marine vessels, Quadratic programming, Reinforcement learning, Safety-enhanced framework.**

I. Introduction

In the maritime sector, ship stability and safety are fundamental prerequisites for the protection of life, property, and the marine environment, historically ensured through prescriptive design rules established by the International Maritime Organization (IMO) [1]. Safety against roll instability, one of the most critical dynamic hazards, was traditionally assessed through static means, primarily by analyzing the righting-arm (GZ) curve in calm water. However, the dynamic realities of the seaway (stochastic wave excitation, nonlinear fluid–structure interaction, and coupling between degrees of freedom) render static analyses insufficient for predicting dynamic failure modes such as parametric roll, pure loss of stability, excessive lateral accelerations, or catastrophic capsizing and broaching [2]. This inadequacy motivated the IMO's *Second Generation Intact Stability Criteria* (SGISC) [1], [3], a multi-tiered, physics-based vulnerability framework supported by experimental procedures of the International Towing Tank Conference (ITTC) [4].

For control-oriented researchers, the SGISC criteria are more than compliance targets: they delineate the mathematical boundaries of a vessel's safe operational envelope. In particular, the criteria for *parametric rolling* and *excessive lateral acceleration* can be recast as real-time inequality constraints on the control input [5], [6]. Embedding such limits directly within the control loop transforms safety margins from passive, post-hoc guidelines into active invariants of the control policy, so that control actions are continuously optimized to avoid violating the dynamic stability envelope.

Enforcing these limits in real time is, however, at odds with how marine roll stabilization has historically evolved: stabilization devices and the controllers designed around them have overwhelmingly prioritized *performance* (comfort and seasickness reduction) rather than formally guaranteed safety (Section II). Safety-critical control offers a principled bridge between the two objectives. Control Barrier Functions (CBFs), unified with Control Lyapunov Functions (CLFs) inside a quadratic program (CLF-CBF-QP), render a safe set forward invariant while allowing a performance-tuned nominal controller to operate freely inside the safe region, intervening only near the stability boundary [7]. In our prior work [8], [9], we mapped the SGISC parametric-roll and excessive-acceleration criteria onto exactly such a safety filter for a vessel actuated by a *canting keel*. This mechanism generates its restoring moment by shifting the vessel's center of gravity and, unlike active fins, remains effective at zero forward speed [10], at the price of large rotational inertia and strict electromechanical actuator limits.

Despite the theoretical rigor of the CLF-CBF-QP paradigm, a critical gap remains in practical maritime implementation: the forward-invariance guarantee of a CBF is only as trustworthy as the model of the Lie derivatives it is built on, and marine roll is precisely the motion for which trustworthy models are hardest to obtain. Section III-A traces this difficulty to the anatomy of standard roll-damping practice, namely Ikeda's component method and its documented ± 20–30% error band, and to effects such methods omit. Compounding the modeling problem, high-inertia actuation such as the canting keel adds substantial latency: the servo-driven keel cannot change its

This paper is part of Project SAFEMARVEL which is funded by the Den Danske Maritime Fond (DDMF) (Project nr. 95-45002-3410157)

restoring moment instantaneously, so a safety filter designed on nominal dynamics may command an input that is physically unrealizable within the available reaction time, allowing the barrier to be breached during aggressive transients. Under extreme sea states, such model–plant mismatch can drive a vessel toward the very failure modes the SGISC seeks to prevent.

Reinforcement learning (RL) offers a data-driven route around this modeling bottleneck: rather than deriving a higher-fidelity physical model for every hull and loading condition, the missing dynamics can be learned from closed-loop data [11]. Yet marine applications of RL have so far optimized performance without formal safety guarantees, while safe-RL architectures that couple learning with control-theoretic certificates have been demonstrated almost exclusively on robotic and automotive benchmarks; their transfer to marine dynamic-stability enforcement, and specifically to a SGISC-constrained, high-inertia canting-keel vessel, remains open (Section II).

Contributions. To bridge this gap, this paper integrates RL into the SGISC-constrained safety-enhanced framework of [8]. A nominal controller provides the performance baseline, a CLF-CBF-QP enforces the parametric-roll and excessive-acceleration criteria, and a Deep Deterministic Policy Gradient (DDPG) [11] agent learns the model-uncertainty terms (Δ_1, Δ_2), the mismatch in the Lie derivatives of the CLF and CBF constraints, together with the analogous correction to the acceleration input constraint, so that the safety filter reasons about the *true* rather than the nominal plant. Building on the RL-CBF-CLF-QP concept of [12] and our canting-keel modeling [8], [10], the paper makes three contributions:

- **Quantifying the cost of model mismatch**: we show that an Ikeda-consistent set of modeling errors (damping within $\pm 30\%$, neglected cubic restoring, neglected keel added mass, unknown metacentric modulation) is sufficient to void the forward-invariance guarantee of the published safety filter [8] and, if the filter is not backed by an emergency fallback, to let parametric rolling grow unchecked.
- **Learning the constraint-level uncertainty**: the DDPG agent estimates $(\hat{\Delta}_1, \hat{\Delta}_2)$ online from measured constraint derivatives, entering the QP affinely so the problem remains a convex program solvable at the embedded control rate.
- **Restoring SGISC compliance**: across parametric-rolling and excessive-acceleration case studies, the RL-augmented filter recovers oracle-level safety (zero barrier violations; lateral acceleration within the IMO threshold) while retaining the nominal controller's stabilization performance.

The remainder of this paper is organized as follows. Section II reviews related work; Section III describes the nonlinear vessel dynamics and canting-keel constraints; Section IV details the RL-augmented safety-enhanced framework and the uncertainty-learning process; Section V validates the approach through numerical simulations; and the final section provides concluding remarks.

II. Related Work

Dynamic stability and the SGISC. The SGISC establish a multi-tier vulnerability framework covering five dynamic failure modes [1], [3]. Simplified operational guidance exists for parametric roll [5] and for the excessive-acceleration criterion across arbitrary speed and heading [6], and such guidance is entering operational layers such as weather routing [13]; model-based controllers suppress parametric resonance directly [14], and data-driven surrogates forecast roll and parametric-roll episodes for decision support [15]. These works define, verify, or forecast the safe envelope but stop short of embedding it as a hard real-time control constraint.

Roll stabilization actuation. A broad spectrum of roll-mitigation devices has been developed, ranging from passive bilge keels and anti-roll tanks to active fin stabilizers and gyrostabilizers, and comparative studies reveal a persistent performance-centric bias [16]–[19]. Advanced roll controllers, including disturbance-observer-enhanced model predictive control [20] and adaptive robust fin control during turns [21], improve roll reduction and robustness, yet likewise stop short of enforcing regulatory limits as hard constraints. The canting keel provides a speed-independent restoring moment [10], [22], [23]; this advantage is offset by large rotational inertia, electromechanical actuator limits, and nonminimum-phase zero dynamics in the positively buoyant "airkeel" configuration. These effects make constraint enforcement delicate, and their violation can cause mechanical failure or paradoxically destabilize the hull [10].

Safety-critical control under uncertainty. CBF-QP filters give formal invariance guarantees for control-affine systems [7], [24], but the guarantees degrade under model uncertainty. Robust, adaptive, and data-driven formulations respond by estimating structured parametric uncertainty [25], [26], learning residual dynamics via Gaussian processes [27], [28], bounding model error distribution-free via conformal prediction [29], or synthesizing verified neural barrier certificates [30].

Learning-based control. In the marine domain, deep RL has targeted anti-roll actuation [31]–[33], path following, formation control, and course keeping of unmanned surface vessels [34], [35], COLREGs-compliant collision avoidance [36], [37], and offshore crane and active heave compensation [38]; adaptive-neural and data-driven schemes have likewise been proposed for roll stabilization and roll-motion prediction [15], [39]. These approaches optimize performance and offer no dynamic-stability guarantee. Conversely, the safe-RL community couples learning with safety certificates: constrained Markov decision processes [40], CBF-filtered policy learning [41]–[44], model-predictive shielding with a robust backup controller [45], learning the model-uncertainty terms that appear in the CBF/CLF constraints [12], and observer-based disturbance compensation inside the filter [46]. These have

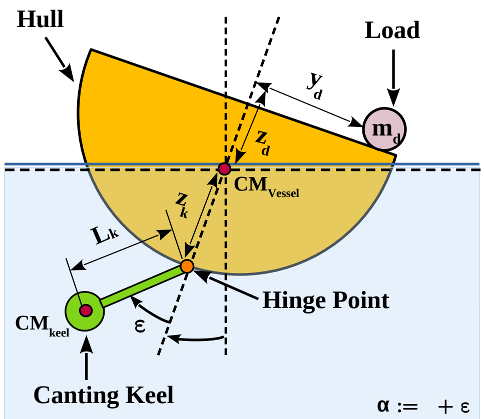

Fig. 1. Schematic representation of a monohull marine vessel equipped with a canting keel mechanism

TABLE I
ROLL-DAMPING MODELING UNCERTAINTY BY HULL FORM AND ITS ROLE AS THE LEARNED RESIDUAL (Δ_1, Δ_2).

Hull form	**Dominant unmodeled effect**	**Ikeda validity**
Monohull	Large-angle eddy/bilge-keel nonlinearity; high-KG error; fluid memory	Within envelope (±20–30%) [52], [54]
Catamaran	Demi-hull hydrodynamic interference on damping/lift	Outside scope [56], [57]
Trimaran	Side-hull emergence; strongly nonlinear interference damping	Outside scope [58], [59]

been demonstrated almost exclusively on robotic and automotive benchmarks. Closest to the present setting, safety-critical marine control under uncertainty has recently been addressed with adaptive CBFs for automated surface vessels [47] and with extended-state-observer-based CBF control for capsizing prevention of an unmanned sailboat [48]; both estimate a lumped disturbance with observer or adaptation machinery rather than learning the constraint-level Lie-derivative mismatch, and neither addresses the SGISC failure modes or high-inertia canting-keel actuation. To our knowledge, the present work is the first to close this loop for a SGISC-constrained canting-keel vessel by learning the Lie-derivative uncertainty inside the safety filter.

III. System Description

This study focuses on the roll stabilization of a monohull marine vessel equipped with a canting keel mechanism. The system configuration comprises a rigid hull and a keel bulb actuated by a DC servo motor via a hinge point at the hull's bottom as shown in Fig. 1. We adopt the coordinate systems and lumped parameter modeling approach detailed in [10].

A. Sources of Model Uncertainty in Roll

The forward-invariance guarantee of a CBF is only as trustworthy as the model of the Lie derivatives it is built on. Roll is the least-damped and most safety-critical rigid-body motion, without the large hydrostatic or wave-radiation damping available in heave and pitch [49], [50]. Its dominant term is the roll damping moment B_{44}, for which the standard engineering model is Ikeda's component method: the total damping is decomposed into physically separable friction, eddy-making, lift, wave, and bilge-keel contributions, several of which are strongly nonlinear and amplitude- and frequency-dependent [51]. Because the quadratic component enters the roll equation as $B_2\dot{\phi}|\dot{\phi}|$, it is commonly replaced by an energy-equivalent linear coefficient, yielding an amplitude- and frequency-dependent effective damping that is naturally read as *parametric uncertainty* in a control model [50]. The regressed Simplified Ikeda Method is accurate only to about ± 20–30% within a bounded envelope of conventional monohulls and loses fidelity at large roll angles and high centers of gravity [52]–[54], i.e., exactly in the parametric-roll and dead-ship regimes the SGISC addresses [55].

This uncertainty is hull-form dependent (Table I). For a **monohull** the Ikeda model provides a usable nominal prior within its envelope. For a **catamaran**, hydrodynamic interference between the demi-hulls significantly alters the roll-damping and lift components, an effect outside the scope of a single-hull calibration [56], [57]. For a **trimaran**, the emergence of the side hulls during roll produces strongly nonlinear flow-interference damping that requires bespoke treatment [58], [59]. Across all three classes the method is a steady-state, forced-roll description that neglects fluid-memory (radiation-convolution) effects governing transient and broadband response [60]. Consequently, deriving a high-fidelity nominal model per hull form is impractical, motivating a data-driven residual: rather than re-identify each hull, the agent of Section IV learns the mismatch (Δ_1, Δ_2) between an Ikeda-seeded nominal model and the true plant, so that the same safety-enhanced architecture can in principle transfer across monohull, catamaran, and trimaran configurations [12]; the present paper demonstrates the monohull case.

B. Vessel Model and Control Form

Assuming negligible coupling between roll and the remaining degrees of freedom (1-DoF), as is common in roll-stabilization studies near resonance where roll dominates, the roll dynamics follow [23], [49]:

$$(I_x + \Delta I_x)\ddot{\phi} + B_\phi(\dot{\phi}) + C_\phi(\phi) = \tau_k + \tau_d + \tau_{wi} \quad (1)$$

where ϕ is the roll angle, I_x and ΔI_x are the rigid-body inertia and hydrodynamic added mass, $B_\phi(\dot{\phi}) = K_p\dot{\phi} + K_{p|p|}\dot{\phi}|\dot{\phi}|$ collects linear and quadratic damping, and $C_\phi(\phi) = \rho g \nabla \overline{GM}_m \phi + K_{\phi^3}\phi^3$ the linear and cubic restoring moments. The right-hand side aggregates the keel torque τ_k (gravitational moment of the bulb, paddling hydrodynamics with added mass Δm_k, and servo reaction torque), the shifting-load disturbance τ_d, and the wave-induced torque τ_{wi}; the complete expressions, including the keel kinematics of Fig. 1, are given in [8], [10].

The keel angle ε is driven by a DC servo motor with proportional-derivative control of the commanded angle ε_d.

TABLE II
NOMINAL-MODEL ERRORS RELATIVE TO THE TRUE PLANT.

Quantity	**True plant**	**Nominal model**
Linear roll damping K_p	4×10^4	5.2×10^4 (+30%)
Quadratic roll damping $K_{p\|p\|}$	5×10^3	2.5×10^3 (−50%)
Cubic restoring K_{ϕ^3}	-5×10^4	0 (neglected)
Keel added mass Δm_k	2×10^3	0 (neglected)
Quadratic keel drag $K_{k\|k\|}$	50	0 (neglected)
Restoring coefficient	$C_1(t)$ modulated	constant C_1

Neglecting armature inductance and assuming a rigid gearbox, the actuator dynamics reduce to the linear form $\ddot{\varepsilon} = \alpha_1\varepsilon + \alpha_2\dot{\varepsilon} + \beta_\varepsilon u$ with $u = \varepsilon_d$ and lumped electromechanical parameters $\alpha_{1,2}, \beta_\varepsilon$ [10]. Substituting the actuator dynamics into (1) yields a fourth-order control-affine system,

$$\ddot{\phi} = f(\mathbf{x}) + g(\mathbf{x})u + d(t), \tag{2}$$

with state $\mathbf{x} = [\phi, \dot{\phi}, \varepsilon, \dot{\varepsilon}]^T$, drift $f(\mathbf{x})$ collecting the unforced hydrodynamics and keel interaction, input map $g(\mathbf{x})$, and external disturbances $d(t)$.

C. Nominal Model versus True Plant

The safety filter of Section IV is synthesized on a *nominal* instance of (2), denoted $(\tilde{f}, \tilde{g})$, while the vessel evolves according to the *true* dynamics (f, g). The nominal model deviates from the true plant through a set of errors chosen to be representative of standard marine practice rather than adversarial: the linear roll damping is overestimated by 30% and the quadratic (eddy-making) component underestimated by 50%, consistent with the documented accuracy band of the Simplified Ikeda Method [52], [53]; the cubic softening of the restoring curve is neglected ($K_{\phi^3} = 0$); the keel added mass Δm_k and the quadratic keel drag $K_{k|k|}$ are neglected, which misrepresents both the effective roll inertia and the actuator-side reaction torque through which the keel command acts; and the periodic metacentric modulation responsible for parametric rolling is unknown to the model, which assumes a constant restoring coefficient. Table II summarizes the mismatch set.

IV. REINFORCEMENT LEARNING BASED SAFETY-ENHANCED CONTROL FRAMEWORK

A. Nominal Controller Augmentation via CLF-CBF-QP

Nominal controllers tuned for performance (e.g., PID or LQR) provide local stability but can degrade away from the equilibrium point, while conventional safety-critical algorithms often prioritize energy optimization over performance. The nominal controller, in nonlinear state-feedback form $\mathbf{k}(\mathbf{x})$, is therefore retained as the performance baseline, and stability and safety guarantees are added around it.

1) Systematic Selection of the CLF for Stability: We seek a Control Lyapunov Function (CLF) that is systematically selectable and interferes minimally with the nominal solution $\mathbf{k}(\mathbf{x})$. Linearizing the system and the nominal controller at the equilibrium yields $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{K}$, and the closed-loop matrix $\mathbf{A}_{cl} = \mathbf{A} + \mathbf{B}\mathbf{K}$. The Lyapunov function is then selected as [61], [62]:

$$\mathcal{V}(\mathbf{x}) = \mathbf{x}^T \mathbf{P}_l \mathbf{x} \tag{3}$$

where $\mathbf{P}_l$ is found by solving the Lyapunov equation,

$$\mathbf{P}_l \mathbf{A}_{cl} + \mathbf{A}_{cl}^T \mathbf{P}_l = -\mathbf{Q} \tag{4}$$

with $\mathbf{Q} = \alpha\mathbf{I}$, so that the CLF constraint preserves the nominal solution near the equilibrium.

2) Safety-Enhanced QP Formulation: The safety-enhanced control problem is formulated as a state-dependent QP that minimizes the deviation of the actual control input u from the nominal control input $\mathbf{k}(\mathbf{x})$. Since the objective function aims to maintain performance close to the nominal controller, the CLF is no longer required to focus on exponential convergence rates. Thus, the $\lambda\mathcal{V}(\mathbf{x})$ term is eliminated from the stability constraint [9].

To address safety, the formulation explicitly incorporates the CBF constraint, which acts as a rigorous safety filter, ensuring that the system trajectories remain within a forward-invariant safe set $\mathcal{C}$ at all times. This constraint allows the performance controller to operate freely within the safe region, only intervening to adjust the control signal when the vessel approaches the predefined stability boundaries defined by the IMO criteria. For parametric rolling, the barrier is the braking barrier function $\mathcal{B}_{PRb}$ of [8], which augments the roll-angle safe set with velocity- and braking-distance terms reflecting the finite control authority of the heavy keel; the excessive-acceleration criterion enters the admissible input set $\mathcal{U}(\mathbf{x})$ as the state-dependent affine constraint $|a_y(\mathbf{x}, u)| \leq R_{EA,limit}$. Building on the insights gained about the nominal controller and the selection of the CLF and CBF, we can set up a CLF-CBF-QP problem as given in [8], [9].

Proposed CLF-CBF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\, \delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u} - \mathbf{k}(\mathbf{x}))^T \mathbf{H} (\mathbf{u} - \mathbf{k}(\mathbf{x})) + p\,\delta^2 \tag{5}$$

$$\begin{aligned} \text{s.t.} \quad & \mathbf{u} \in \mathcal{U}(\mathbf{x}) && \text{(Input Constraint)} \\ & L_\mathbf{f}\mathcal{V}(\mathbf{x}) + L_\mathbf{g}\mathcal{V}(\mathbf{x})\mathbf{u} \leq \delta && \text{(CLF Constraint)} \\ & L_\mathbf{f}\mathcal{B}(\mathbf{x}) + L_\mathbf{g}\mathcal{B}(\mathbf{x})\mathbf{u} \geq -\gamma\mathcal{B}(\mathbf{x}) && \text{(CBF Constraint)} \end{aligned}$$

where $\mathcal{U}(\mathbf{x}(t)) \subset \mathbb{R}^m$ denotes the set of admissible control inputs at time t, dependent on the current state $\mathbf{x}(t)$; $\mathbf{H}$ is a weighting matrix and δ is a slack variable with penalty weight p that relaxes the stability constraint to ensure the feasibility of the safe solution space $(\mathcal{U}(\mathbf{x}) \cap \Pi_\mathcal{V} \cap \Pi_\mathcal{B})$ [9]. In implementation, the CBF constraint is imposed while the state is inside the safe set ($\mathcal{B} > 0$); if the barrier is nevertheless breached, the filter switches to a best-effort recovery input that maximizes the (corrected) barrier derivative subject to the remaining input constraints, and hands back to the QP once $\mathcal{B} \geq 0$ is restored. This fallback is shared by every controller variant compared in Section V, so differences between them isolate the effect of the model knowledge inside the QP.

B. Learning Model Uncertainty with a DDPG Agent

The QP in (5) is synthesized on the *nominal* model $(\tilde{f}, \tilde{g})$ of Section III-C. Consequently, the Lie derivatives that define the CLF and CBF constraints are uncertain. Following the RL-CBF-CLF-QP concept of [12], we write the *true* constraint derivatives as their nominal values augmented by additive uncertainty terms:

$$\dot{\mathcal{V}}(\mathbf{x},\mathbf{u}) = L_{\tilde{\mathbf{f}}}\mathcal{V}(\mathbf{x}) + L_{\tilde{\mathbf{g}}}\mathcal{V}(\mathbf{x})\,\mathbf{u} + \Delta_1(\mathbf{x}), \tag{6a}$$

$$\dot{\mathcal{B}}(\mathbf{x},\mathbf{u}) = L_{\tilde{\mathbf{f}}}\mathcal{B}(\mathbf{x}) + L_{\tilde{\mathbf{g}}}\mathcal{B}(\mathbf{x})\,\mathbf{u} + \Delta_2(\mathbf{x}), \tag{6b}$$

where Δ_1, Δ_2 lump the effects of the residual dynamics on the stability and safety constraints, respectively; the analogous additive term $\Delta_a(\mathbf{x})$ corrects the model-based part of the lateral-acceleration constraint in $\mathcal{U}(\mathbf{x})$. The additive, state-dependent form absorbs the drift mismatch exactly and approximates the (small) input-gain mismatch; in the present vessel the error in $L_{\mathbf{g}}\mathcal{B}$ is below 10%, while the drift error dominates.

A Deep Deterministic Policy Gradient (DDPG) agent [11], with actor μ_θ and critic Q_ψ, observes the state $\mathbf{x} = [\phi, \dot{\phi}, \varepsilon, \dot{\varepsilon}]^T$, augmented with the measured wave-encounter phase $(\sin\omega_e t, \cos\omega_e t)$ so that excitation-synchronous mismatch is representable, and outputs the estimates $(\hat{\Delta}_1, \hat{\Delta}_2)$. These estimates are injected into the QP as state-dependent affine corrections, yielding the RL-augmented safety filter:

RL-Augmented CLF-CBF-QP

$$\mathbf{u}^*(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m,\,\delta\in\mathbb{R}}{\operatorname{argmin}} \quad \frac{1}{2}(\mathbf{u}-\mathbf{k}(\mathbf{x}))^T\mathbf{H}(\mathbf{u}-\mathbf{k}(\mathbf{x})) + p\,\delta^2 \tag{7}$$

$$\begin{aligned} \text{s.t.}\quad & \mathbf{u}\in\mathcal{U}(\mathbf{x};\hat{\Delta}_a) \\ & L_{\tilde{\mathbf{f}}}\mathcal{V} + L_{\tilde{\mathbf{g}}}\mathcal{V}\,\mathbf{u} + \hat{\Delta}_1(\mathbf{x}) \le \delta \\ & L_{\tilde{\mathbf{f}}}\mathcal{B} + L_{\tilde{\mathbf{g}}}\mathcal{B}\,\mathbf{u} + \hat{\Delta}_2(\mathbf{x}) \ge -\gamma\mathcal{B}(\mathbf{x}) \end{aligned}$$

Because $\hat{\Delta}_1, \hat{\Delta}_2, \hat{\Delta}_a$ enter (7) affinely, the problem remains a convex QP solvable at the embedded control sample rate; for the scalar keel command it reduces to a piecewise-quadratic minimization over an interval and is solved in closed form.

Reward design. The agent is trained on episodes of the closed-loop system with the RL-augmented filter in the loop. The reward combines three terms: (i) the prediction error of the corrected constraint derivatives, $|\dot{\mathcal{B}}_{meas} - (\dot{\mathcal{B}}_{nom} + \hat{\Delta}_2)|^2$ and its CLF counterpart, where $\dot{\mathcal{B}}_{meas}$ is obtained by finite differences of the measured barrier value over one control interval; (ii) a hinged safety penalty $\max(0, m - \mathcal{B})$ with margin $m > 0$, so that shallow grazing violations of the barrier remain visible to the learning signal rather than vanishing quadratically; and (iii) a penalty on the deviation of the realized input from the nominal command, which protects baseline stabilization performance. The prediction-error term anchors $(\hat{\Delta}_1, \hat{\Delta}_2)$ to the physically measurable mismatch, in the spirit of [12]; the margin term shapes the estimates conservatively near the boundary of the safe set.

V. Simulation Results and Discussion

The framework is validated on the Klara 2 research vessel model of [8] (length overall 14 m, displacement 20 t; physical parameters as in Table I of [8]). The true plant integrates the full nonlinear dynamics of Section III with a fixed-step fourth-order Runge–Kutta scheme at 5 ms; the safety filter and the agent run at a 20 ms control period, consistent with embedded marine implementations. The commanded keel angle is bounded to $\pm 45°$. As in [8], wave forcing is omitted and the two SGISC failure mechanisms are activated deterministically, so that the interaction between the safety filter and the stability boundary can be isolated; the nominal LQR is deliberately tuned with moderate gains so that parametric rolling is not masked by closed-loop damping alone. Four controller variants are compared: (a) the nominal LQR without safety filter; (b) the CLF-CBF-QP of [8] synthesized on the nominal (mismatched) model of Table II; (c) the proposed RL-augmented CLF-CBF-QP, i.e., variant (b) plus the learned corrections; and (d) an *oracle* filter synthesized on the true plant, which upper-bounds the achievable performance. The DDPG actor and critic are two-hidden-layer perceptrons (64 neurons per layer); training uses 60 episodes of 60 s for the parametric-roll scenario and 120 episodes of 10 s for the acceleration scenario, with exploration noise decayed across episodes.

A. Parametric Rolling under Model Mismatch (Scenario 1)

Parametric rolling is triggered by modulating the hull restoring coefficient as $C_1(t) = C_1[1 + (\Delta GM/\overline{GM}_m)\cos(\omega_e t)]$ with $\Delta GM = 0.70$ m and $\omega_e = 2\omega_{n,OL} = 3.03$ rad/s, starting from $\phi(0) = 5°$; Floquet analysis in [8] confirms that both the open-loop and LQR-controlled plants are Mathieu-unstable at this excitation depth. The filter variants (b)–(c) do *not* know $C_1(t)$: their constraint derivatives are evaluated on the constant-C_1 nominal model.

B. Excessive Lateral Acceleration under Model Mismatch (Scenario 2)

The second scenario evaluates the acceleration constraint on the larger vessel configuration of [8] (bridge height $h_r = 7.5$ m, keel arm $L_k = 3.5$ m, higher hull damping), with an aggressive recovery LQR returning the vessel from $\phi(0) = 30°$ after a large wave impact. The lateral acceleration at the bridge, $a_y = k_L(g\sin\phi + h_r\ddot{\phi})$, is affine in the control input and must satisfy $|a_y| \le R_{EA,limit} = 4.64$ m/s^2 [1].

C. Discussion

VI. Conclusions

This paper extended the SGISC-constrained safety-enhanced control framework for canting-keel vessels with a reinforcement-learning layer that learns the model-uncertainty terms appearing in the CLF, CBF, and acceleration constraints of the governing quadratic program.

References

[1] International Maritime Organization, "Interim guidelines on second generation intact stability criteria," International Maritime Organization, London, UK, Tech. Rep. Circular MSC.1/Circ.1627, 2020.

[2] M. A. S. Neves, V. L. Belenky, J. O. De Kat, K. Spyrou, and N. Umeda, *Contemporary Ideas on Ship Stability and Capsizing in Waves*. Springer, 2011.

[3] G. Bulian and A. Francescutto, "The development of second generation intact stability criteria," *Sustainable Development and Innovations in Marine Technologies*, 2019.

[4] International Towing Tank Conference, "Stability," ITTC Recommended Procedures and Guidelines, Tech. Rep. 7.5-02-07-04, 2017, rev. 02.

[5] E. Begovic *et al.*, "Simplified operational guidance for second generation intact stability criteria," *Ocean Engineering*, 2023.

[6] ——, "Excessive acceleration simplified operational guidance," *International Journal of Naval Architecture and Ocean Engineering*, 2022.

[7] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[8] S. Chaudhuri, H. Ramezani, and J. Jouffroy, "Ensuring safe roll dynamics in marine vessels based on a safety-enhanced control framework," in *VI. International Conference on Electrical, Computer and Energy Technologies (ICECET)*, Rome, Italy, 2026.

[9] ——, "Remarks on augmenting stability and safety in nominal controllers: A practical perspective," in *2025 European Control Conference (ECC)*, 2025, pp. 2121–2128.

[10] H. Ramezani, S. Chaudhuri, J. Jouffroy, A. Bauricher, and S. Hansen, "On roll stabilization using a canting keel," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 2, pp. 539–553, 2025.

[11] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[12] J. J. Choi, F. Castañeda, C. J. Tomlin, and K. Sreenath, "Reinforcement learning for safety-critical control under model uncertainty, using control lyapunov functions and control barrier functions," *Robotics: Science and Systems (RSS)*, 2020.

[13] M. Życzkowski *et al.*, "Implementation of the operational guidance as a part of the second generation intact stability criteria to the deterministic ship weather routing," *Ocean Engineering*, 2025.

[14] S.-D. Lee *et al.*, "Active control synthesis for parametric instability of container ship model," *Acta Mechanica*, 2024.

[15] D. Xu *et al.*, "An enhanced hybrid scheme for ship roll prediction using support vector regression and TVF-EMD," *Ocean Engineering*, 2024.

[16] R. Moaleji and A. R. Greig, "On the development of ship anti-roll tanks," *Ocean Engineering*, vol. 34, no. 1, pp. 103–121, 2007.

[17] N. C. Townsend and R. A. Shenoi, "Control strategies for marine gyrostabilizers," *IEEE Journal of Oceanic Engineering*, vol. 39, no. 2, pp. 243–255, 2013.

[18] H. A. Ahmad Naziri *et al.*, "Multi-stabilizer devices for marine vessel, design and control – a review," *IEEE Int. Conf. on Underwater System Technology (USYS)*, 2022.

[19] A. Giallanza and G. Marannano, "Interactive roll stabilization comparative analysis for large yacht: Gyroscope versus active fins," *International Journal on Interactive Design and Manufacturing*, vol. 14, pp. ON–LINE, 2020.

[20] I. A. Jimoh, S. Küchler *et al.*, "Fin control for ship roll motion stabilisation based on observer enhanced MPC with disturbance rate compensation," *Ocean Engineering*, vol. 224, 2021.

[21] S. Zhang *et al.*, "L2-gain based adaptive robust heel/roll reduction control using fin stabilizer during ship turns," *Journal of Marine Science and Engineering*, vol. 9, no. 8, 2021.

[22] H. Ramezani and J. Jouffroy, "On roll stabilisation using a canting keel," in *62nd IEEE Conference on Decision and Control (CDC)*, 2023.

[23] S. Chaudhuri, H. Ramezani, and J. Jouffroy, "Modelling and control of a canting keel-based ship roll stabilization system for crane operations," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2022, pp. 1904–1909.

[24] W. Xiao and C. Belta, "High-order control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no. 7, pp. 3655–3662, 2022.

[25] B. T. Lopez, J.-J. E. Slotine, and J. P. How, "Robust adaptive control barrier functions: An adaptive and data-driven approach to safety," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1031–1036, 2021.

[26] T. Kim *et al.*, "Learning to refine input constrained control barrier functions via uncertainty-aware online parameter adaptation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.

[27] Y. Hu *et al.*, "Decentralized robust collision-avoidance for cooperative multirobot systems: A gaussian process-based control barrier function approach," *IEEE Transactions on Control of Network Systems*, 2023.

[28] J. Li *et al.*, "Robust safe learning and control in an unknown environment: An uncertainty-separated control barrier function approach," *IEEE Robotics and Automation Letters*, 2023.

[29] K. Y. Chee, M. A. Hsieh, and N. Matni, "Uncertainty quantification and robustification of model-based controllers using conformal prediction," *arXiv / Learning for Dynamics and Control (L4DC)*, 2024.

[30] H. Zhao, X. Zeng, T. Chen, and Z. Liu, "Learning safe neural network controllers with barrier certificates," *Formal Aspects of Computing*, vol. 33, pp. 437–455, 2020.

[31] J. Lin *et al.*, "Intelligent ship anti-rolling control system based on a deep deterministic policy gradient algorithm and the magnus effect," *Physics of Fluids*, vol. 34, no. 5, 2022.

[32] W. Xu *et al.*, "Study on the intelligent anti-roll control of a ship at zero speed based on flapping fins," *Ocean Engineering*, 2025.

[33] Y. Zou *et al.*, "Sea trial of anti-rolling performance of a zero-speed ship model in coastal waves using PPO-based reinforcement learning control of flapping fins," *Ocean Engineering*, 2026.

[34] Y. Zhao *et al.*, "Path following optimization for an underactuated USV using smoothly-convergent deep reinforcement learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 10, pp. 6208–6220, 2021.

[35] ——, "USV formation and path-following control via deep reinforcement learning with random braking," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 5468–5478, 2021.

[36] D.-H. Chun, M.-I. Roh *et al.*, "Deep reinforcement learning-based collision avoidance for an autonomous ship," *Ocean Engineering*, vol. 234, 2021.

[37] Y. Wang *et al.*, "Deep reinforcement learning based collision avoidance system for autonomous ships," *Ocean Engineering*, 2024.

[38] J.-H. Bae *et al.*, "Heave reduction of payload through crane control based on deep reinforcement learning using dual offshore cranes," *Journal of Computational Design and Engineering*, 2022.

[39] V. S. Nguyen *et al.*, "Fixed-time adaptive RBF neural network controller via minimum learning parameter for ship roll stabilization," *Applied Ocean Research*, 2025.

[40] A. Wachi and Y. Sui, "Safe reinforcement learning in constrained markov decision processes," *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.

[41] Z. Marvi and B. Kiumarsi, "Safe reinforcement learning: A control barrier function optimization approach," *International Journal of Robust and Nonlinear Control*, vol. 31, no. 6, pp. 1923–1940, 2021.

[42] L. Yang *et al.*, "CBF-RL: Safety filtering reinforcement learning in training with control barrier functions," *arXiv preprint arXiv:2510.14959*, 2025.

[43] X. Kong *et al.*, "Differential high order control barrier function-based safe reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 10, pp. 7524–7531, 2025.

[44] B. Gangopadhyay, P. Dasgupta, and S. Dey, "Safe and stable RL (S2RL) driving policies using control barrier and control lyapunov functions," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 2, pp. 1889–1899, 2023.

[45] O. Bastani, "Safe reinforcement learning with nonlinear dynamics via model predictive shielding," in *American Control Conference (ACC)*, 2021.

[46] Z. Hou *et al.*, "Safe reinforcement learning for autonomous driving by using disturbance-observer-based control barrier functions," *IEEE Transactions on Intelligent Vehicles*, 2025.

[47] Y. Zhang *et al.*, "Safety-critical automated surface vessels MIMO control with adaptive control barrier functions under model uncertainties," *IEEE Transactions on Automation Science and Engineering*, 2025.
[48] B. Peng *et al.*, "Safety-critical disturbance rejection control for capsizing prevention of an unmanned sailboat," *IEEE Transactions on Automation Science and Engineering*, 2026.
[49] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. John Wiley & Sons, 2011.
[50] T. Pérez, *Ship Motion Control: Course Keeping and Roll Stabilisation Using Rudder and Fins*. Springer, London, 2005.
[51] Y. Ikeda, Y. Himeno, and N. Tanaka, "A prediction method for ship roll damping," Department of Naval Architecture, University of Osaka Prefecture, Tech. Rep. 00405, 1978.
[52] Y. Kawahara, K. Maekawa, and Y. Ikeda, "A simple prediction formula of roll damping of conventional cargo ships on the basis of Ikeda's method and its limitation," *Journal of Shipping and Ocean Engineering*, vol. 2, no. 4, pp. 201–210, 2012.
[53] M. A. R. Irkal, S. Nallayarasu, and S. K. Bhattacharyya, "Numerical prediction of roll damping of ships with and without bilge keel," *Ocean Engineering*, vol. 179, pp. 114–134, 2019.
[54] S. S. Kianejad, H. Enshaei, J. Duffy, and N. Ansarifard, "Ship roll damping coefficient prediction using CFD," *Journal of Ship Research*, vol. 63, no. 4, pp. 237–250, 2019.
[55] N. Petacco, D. Villa, S. Brizzolara, and P. Gualeni, "An insight about roll damping within the second generation intact stability criteria," *Journal of Ocean Engineering and Marine Energy*, 2024.
[56] Y. Zhang *et al.*, "Nonlinear roll damping and roll motion study of asymmetric catamaran with variable layouts," *Proc. ASME Int. Conf. on Ocean, Offshore and Arctic Engineering (OMAE)*, 2020.
[57] M. Araki *et al.*, "A study on the hydrodynamic interactions between demi-hulls on roll damping of a catamaran-type crew transfer vessel using RANS simulations," *Proc. ASME Int. Conf. on Ocean, Offshore and Arctic Engineering (OMAE), Vol. 4: CFD, FSI and AI*, 2025.
[58] Y. Zhang *et al.*, "Numerical and experimental study on nonlinear roll damping characteristics of trimaran vessel," *Ocean Engineering*, vol. 261, 2022.
[59] J. Gong, Y. Li *et al.*, "The effect of side-hull position on the seakeeping performance of a trimaran at various headings," *Ocean Engineering*, vol. 239, 2021.
[60] T. Pérez and T. I. Fossen, "Time- vs. frequency-domain identification of parametric radiation force models for marine structures at zero speed," *Modeling, Identification and Control*, vol. 29, no. 1, pp. 1–19, 2008.
[61] J.-J. E. Slotine and W. Li, *Applied Nonlinear Control*. Englewood Cliffs, NJ: Prentice Hall, 1991, vol. 199, no. 1.
[62] H. K. Khalil, *Nonlinear Control*. New York: Pearson, 2015.

Reinforcement learning for roll stabilization of vessels using a digital twin

Lucas Beck Røn Andersen, Martin Rosenlyst Jørgensen, Mads Toudal Frandsen

Abstract

In this note, the use of reinforcement learning to control an airkeel system in a simulated ship model built in Simulink is explored. A Deep Q-Network (DQN) agent was trained to directly control the keel, showing improved performance over a DDPG agent. With a carefully designed reward function and added penalties for abrupt actions, the final DQN agent demonstrated effective roll stabilization and reduced output jitter in simulation.

I. Introduction

A *pure* reinforcement learning method for controlling the airkeel has been tested. 'Pure' because the RL agent directly controls the airkeel, whereas other methods have also been discussed, such as having the RL agent tune the parameters of a control algorithm such as a PID.

The reinforcement learning has been done in a simulated model of the ship, made in Simulink, where the training of the agent I also done. This is done to keep the system within one piece of software, which is also practical when it comes to deploying the agent on a real-life model boat.

II. Simulation Environment Setup for Reinforcement Learning

The simulated boat model provided by SDU Mechatronics has been modified to include the necessary Simulink blocks to support reinforcement learning. The entire system can be seen in Fig. 1. Here, the environment (the model of the physical system) can be seen on the right. In the middle is the RL agent, observation block, reward function, and stop function. On the left the observables are being prepared/collected.

The authors are with SDU Physics, DDepartment of Physics, Chemistry and Pharmacy, Odense, Denmark (e-mail: luand20@student.sdu.dk; mrosenlyst@sdu.dk; frandsen@cp3.sdu.dk).

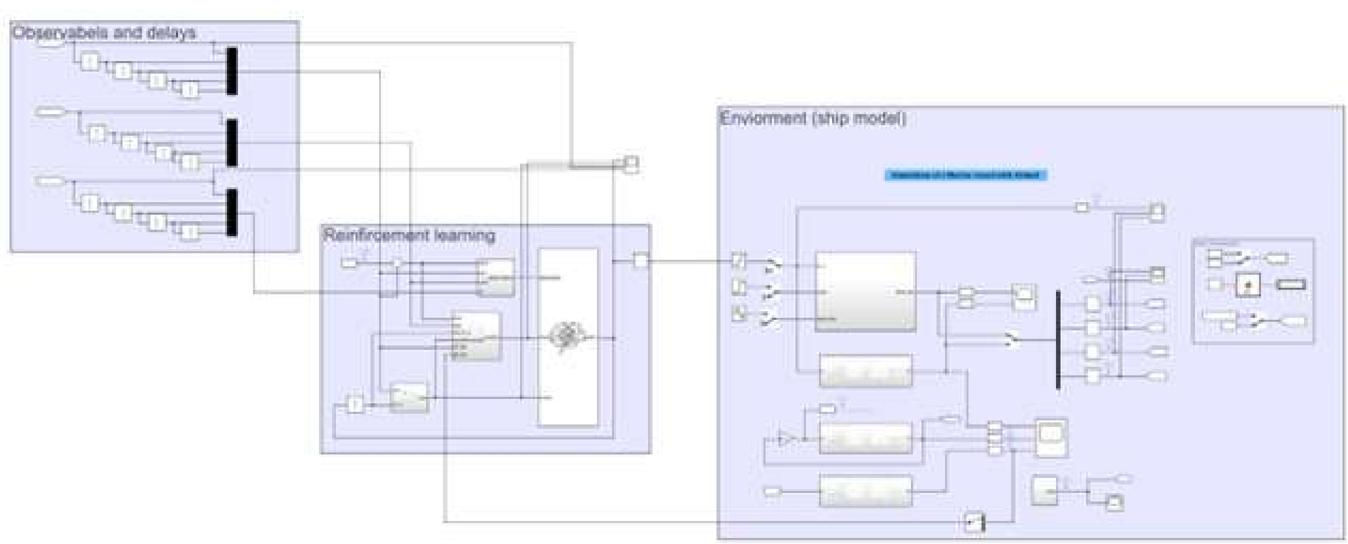

Fig. 1. Overview of the Simulink environment setup for reinforcement learning.

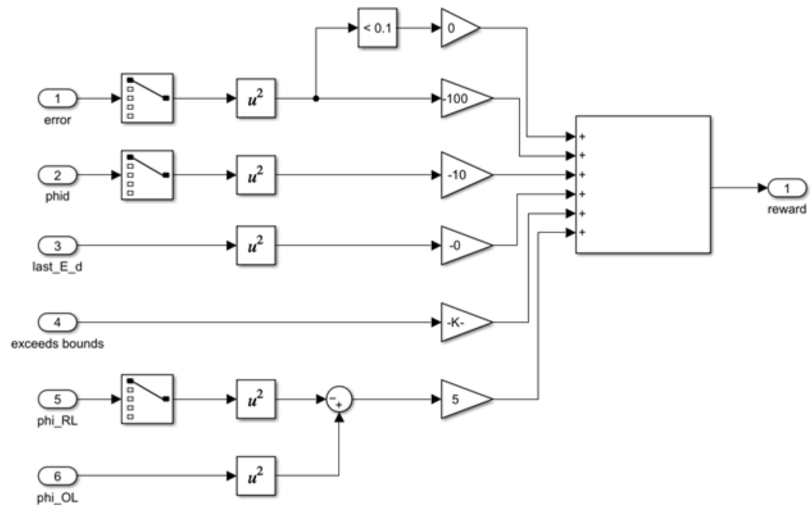

Fig. 2. Simulink reward function block.

As for the observable used, the current roll angle, roll velocity, and the keel angle are used, as well as the 4 previous ones, to help the model better understand the system dynamics. Furthermore, the steering error and the integrated error are also used as observables. Here, the error is the difference between the roll angle and the desired angle, which is 0.

A. *The reward function*

The reward function block can be seen in Fig. 2. In this function, the reward given to the agent pr. step is calculated. Inputs include: the error (as described earlier), roll velocity, previous steering input, if the model exceeds some steering bounds, the roll angle of the RL controlled system and the roll angle of an open loop system.

All the inputs are squared, to account for the possibility of plus and minus signs, and to punish large errors. All of the inputs are passed though a gain, that has been tuned from looking at other examples, as well as from experience. Some are set to zero, as it was found these had no or negative effect on the training. The uppermost gain in the figure was used to reward the agent extra, if the roll was kept close to zero, but was ultimately dropped to smoothen the reward function. The ‘exceeds bounds’ input is used to punish the agent if the model gives an output of more than ± 60 degrees to the keel, or if the boat roll become too large. If the model exceeds the bounds, the training episode is stopped, and the model is given a negative reward of 10^5, ensuring that the model learns not to give out an output larger than the keel limit.

The last two inputs for the reward function are the roll of the RL control system as well as the roll of an open loop system. This is to reward the agent, when it performs better than the open loop system, which corresponds to not doing anything, and punish the agent for performing worse. The agent might perform an action resulting in less reward, but if this action results in a steady boat in the long run, the short while punishment might be worth it in the long run.

B. *Smoothness Penalty*

An additional penalty term was later included to discourage abrupt changes in output, as shown in Fig. 3. If the agent’s action changes by more than $10°$ from the previous step, a proportional penalty is applied.

C. *Stop Condition*

A dedicated block halts training if bounds are exceeded, reinforcing the severe penalty described earlier.

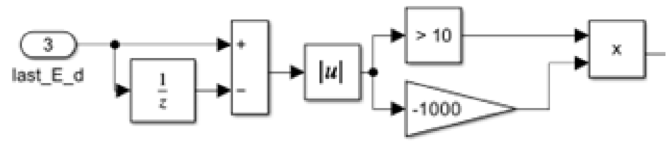

Fig. 3. Smoothness Penalty in the reward function.

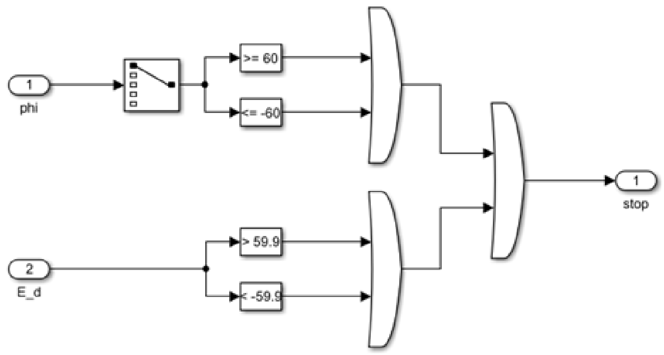

Fig. 4. Stop condition for training

III. Agent Setup

A MATLAB script is used to configure the agent, including environment setup and neural network design. Two agent types were tested:

- **Deterministic Deep Policy Gradient (DDPG)** – a continuous action-space agent, meaning that the model in this case only have a single output [1].
- **Deep Q-Network (DQN)** – a discrete action-space agent, meaning that the possible output of the network needs to be defined [2].

For DQN, the output space was discretized into:

- $-20°$ to $20°$: steps of $1°$
- $\pm 20°$ to $\pm 60°$: steps of $2°$

Finer granularity increases training time substantially.

IV. Training the agents

During the training of the network, 3 metrics are monitored to determine how well the process is going. The episode reward and the average episode reward shows how well the agents performs in the environment. The reward is expected to fluctuate, both since each episode is initiated differently, but also due to the agent exploring different policies to find the optimum such the reward can be maximized. The last metric is the Episode Q0, which is the reward that the agent expects to get from the episode. The Episode Q0 is expected to converge towards the episode reward, to have a properly trained network.
The training of the DDPG can be seen below in Fig. 5 in the figure showing the training monitor. Here it is seen that the episode Q0 diverges, while the reward keeps at a much smaller number. This suggest that the agent does not train properly, which is proven later when testing out the model. There can be many reasons for this including the agent used, the complexity of the network, hyperparameters etc. Ultimately no tuning of the network or hyperparameters showed improvement to this model, therefor the DQN agent was tested.
The DQN model trained much better compared to the DDPG, since the episode Q0 matches much better than the episode reward as seen in the Fig. 6. Ultimately, the only way to know if the agent performs as expected is to test it in the environment.

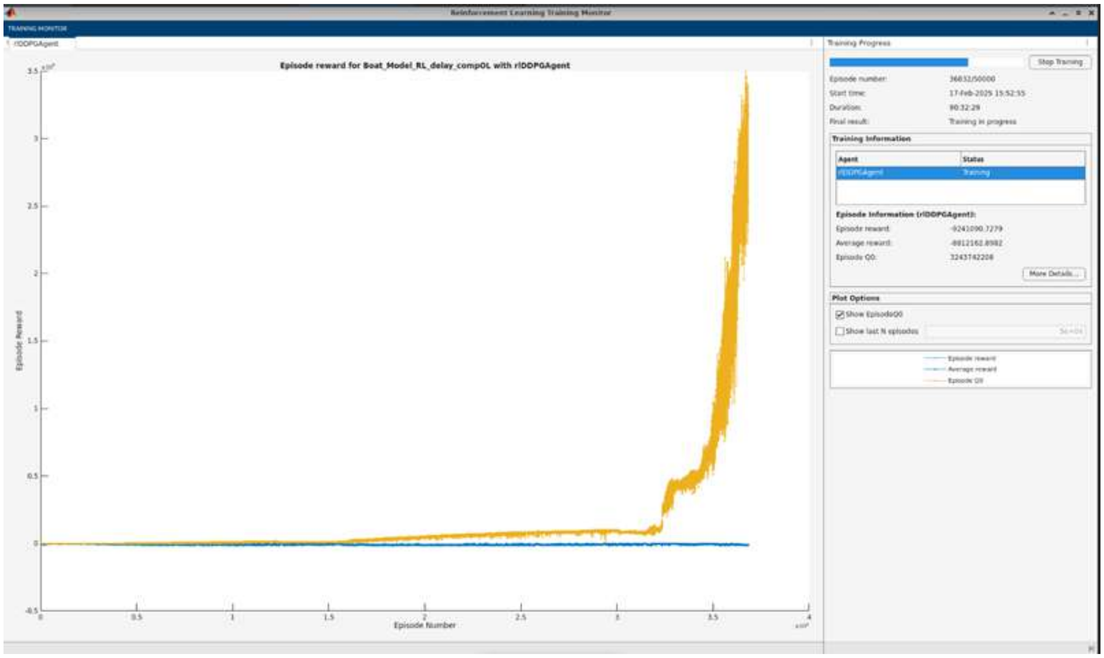

Fig. 5. Training monitor for DDPG.

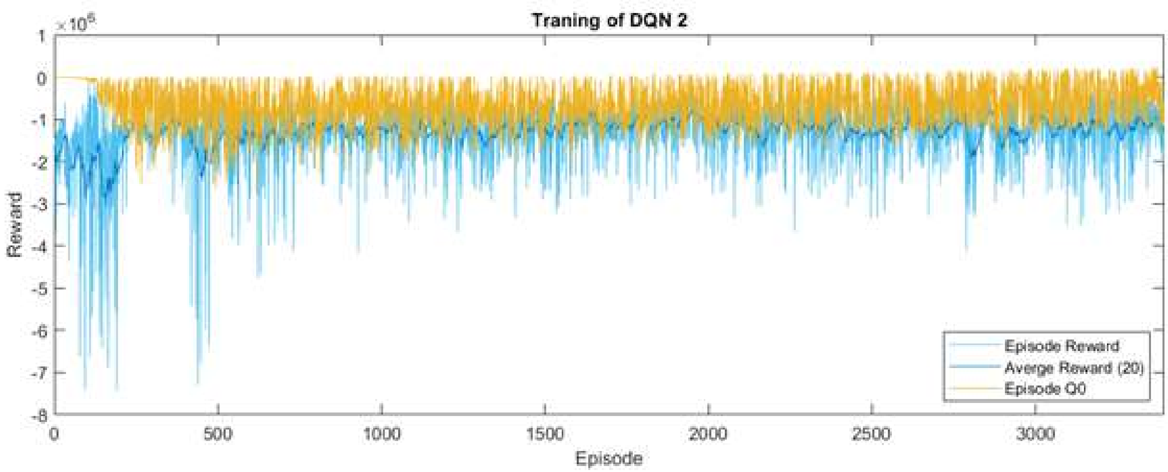

Fig. 6. Training monitor for DQN.

V. Training Results

Testing the DQN network in a single simulation, it is seen that there is a substantial 'jitter' to the output resulting in oscillations, especially after 7 seconds. However, the initial tendencies of the simulation suggest that the agent tries to stabilize the ship at least partially successfully, as the boat never rolls beyond 0° s, and hence the agent stops the roll motion. Though much can be improved, the agent shows potential to stabilize the ship in the simulation as shown in Fig. 7

When running more simulations to test the DQN agents' performance, there are examples of similar behavior, however also examples where the oscillations make the boat continue to roll.

VI. Improving the Agent

To punish the model for these unwanted oscillations in the output, a DQN network with the addition of the smoothing term in the reward described earlier is tried out.

The training can be seen in Fig. 8. Here it is again seen that the episode Q0 after some training time

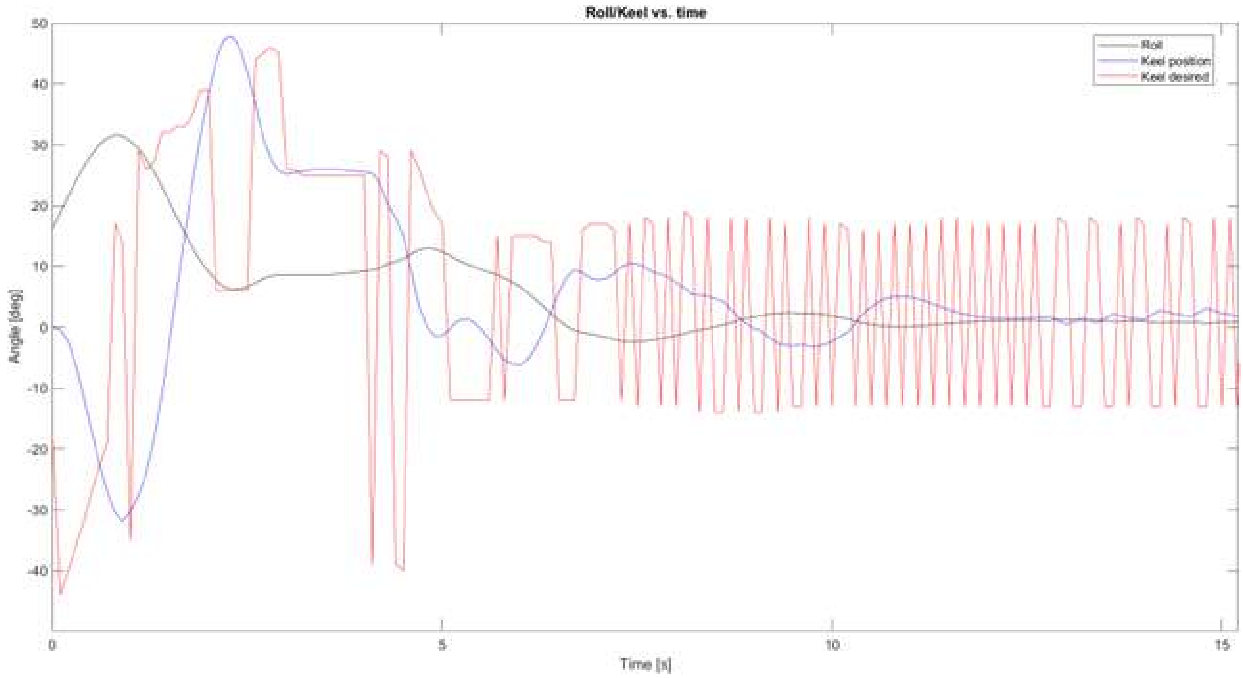

Fig. 7. Roll/Keel angle after training the agent

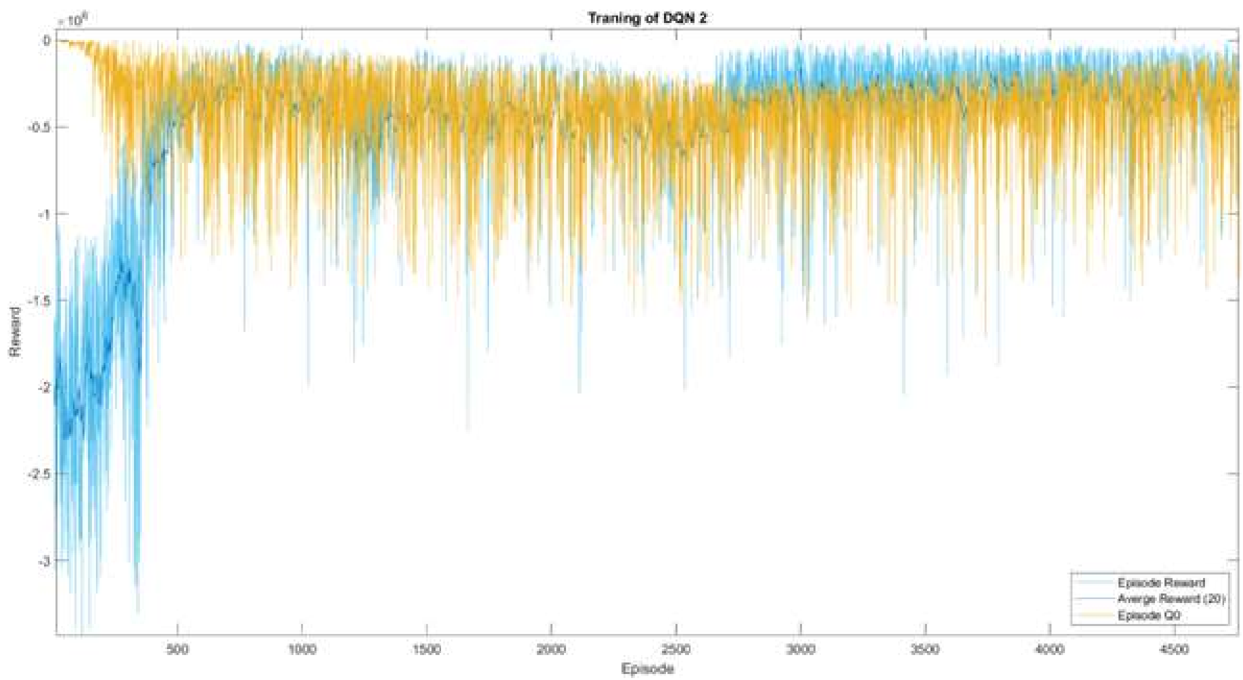

Fig. 8. Training of improved DQN agent

converges towards the episode reward. For the first roughly 700 episode the reward clearly increases, as the training progresses, however from episode 700 to 2600 no clear improvements are seen in the average reward or episode reward. At episode 2600 some episode rewards show clear improvement, with the average reward also slowly following. When the ‘step’ in episode reward happens at episode 2600 it is seen that the episode Q0 underestimates, the reward, which is likely since the agent’s estimation of the reward needs to adjust for the improvement of the network. In the end of the training run, the episode Q0 estimates the expected reward much better.

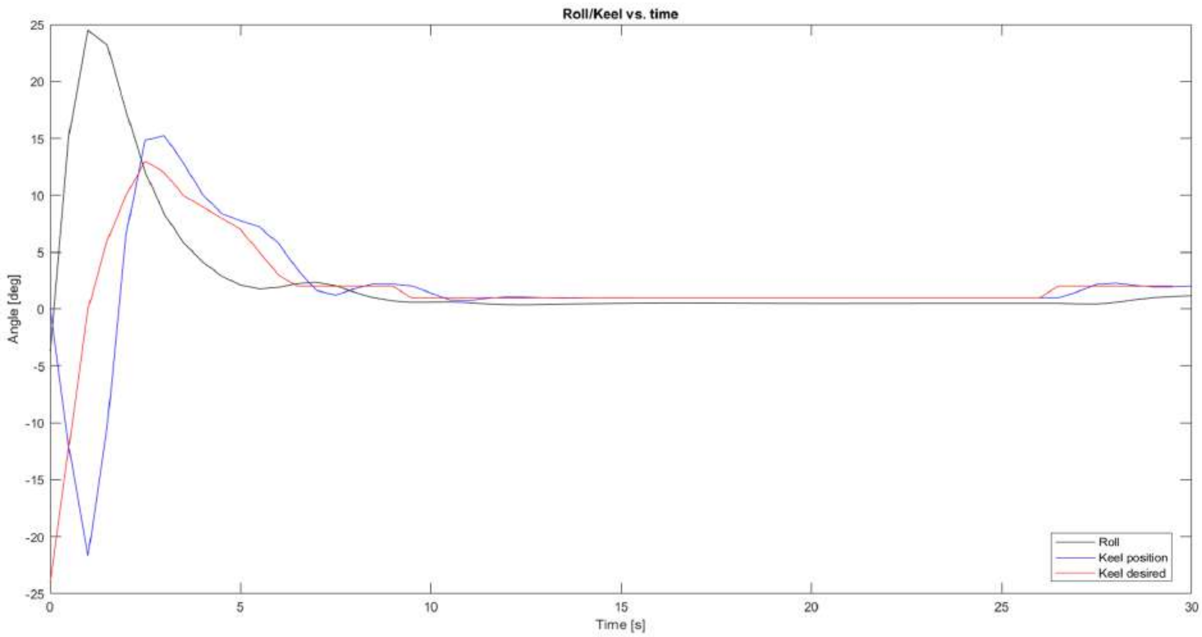

Fig. 9. Improvement in roll/keel angle after training improved DQN agent

VII. Final Results

Inspecting the performance of the new agent in a simulation run show that agent performs much better, and that the jitter has been greatly reduced. The output of the agent is much smoother. The roll motion of the boat is arrested much quicker, and 0 degrees is reached much faster.
Inspecting other simulations also suggest much less jitter and oscillations, however some oscillations still occur in some of the simulations. Some of the simulations also suggest that the agent does not stabilise the ship as well as in the shown simulation in Fig. 9.

References

[1] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” 2019. [Online]. Available: https://arxiv.org/abs/1509.02971
[2] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, “Playing atari with deep reinforcement learning,” 2013. [Online]. Available: https://arxiv.org/abs/1312.5602

Catamaran Stabilization Control System Testing

In-Company Project

Pablo Álvarez Benito
paulv23@student.sdu.dk

Company supervisor: Arnd Baurichter
SDU supervisor: Hossein Ramezani

1 Introduction

In maritime operations, one of the most persistent challenges is maintaining vessel stability in the face of rolling motions caused by waves. This natural movement not only affects the comfort and safety of passengers and crew but also limits operational efficiency, particularly in rough seas.

Dacoma ApS addresses this problem with its revolutionary Airkeel stabilization technology. The Airkeel is an air-filled appendage mounted at the bottom of the vessel's hull. It swings in an inverse-pendulum motion from side to side, actively counteracting roll motions. To optimize the swinging motion, a control system is implemented, ensuring a fast and effective damping of the rolling motion of the boat.

This patented system not only enhances safety and comfort, but also reduces fuel consumption and CO_2 emissions by decreasing ship resistance. The Airkeel is an ideal solution for offshore professionals, as it helps with motion sickness, which often leads to loss of industrial staff productivity, resulting in further costs and eventually project delays.

2 CTV Model Boat

The main function of the Crew Transfer Vessel (CTV) Model Boat is to be able to test different keel solutions, wave impact reduction techniques and control algorithms with a boat similar to a CTV in down scale.

The CTV Model Boat allows for the opportunity to fit and test these different solutions which can be later upscaled to a real size CTV. The different components of the CTV Model Boat will be explained in the following sections.

Figure 1: CTV Model Boat

Figure 2: CTV Model Boat side view

2.1 The Boat

The boat is a catamaran, which is a type of boat characterized by its two parallel hulls of equal size connected by a deck. This design distinguishes from monohull boats, which have a single hull. Catamarans are widely used for both recreational and commercial purposes, and they offer several advantages due to their unique structure.

The dual-hull configuration reduces the rolling motion caused by waves and provides good stability, what makes catamarans ideal to prevent seasickness.

Another advantage is that catamarans usually have a shallow draft, so they can navigate in shallow waters that monohulls may not access.

The main structural parameters of the boat can be found in Table 1.

Parameter	**Value**
Length Overall (LOA)	1.55 m
Length Between Perpendiculars (LPP)	1.50 m
Breadth	0.66 m
Breadth at Design Waterline (BDWL)	0.60 m
Draft at design waterline load (45 kg)	0.11 m
Displacement (light ship)	ca. 20 kg
Approximate placement of weights (2 kg) from centerline to port	0.275 m
Transverse Metacentric Height (GMt) at waterline load (45 kg)	0.394 m

Table 1: CTV Model Boat Parameters

2.2 The Airkeel

The Airkeel is an air-filled appendage mounted at the bottom of the vessel hull. It functions as a roll damping buoyancy positive object under the water, moving from side to side to counteract the roll motions by changing the lateral buoyancy of the boat.

The Airkeel of the model boat is put together with eight different 3D printed parts. The parts are printed in PETG and consist of five parts for the keel and two parts for the fin. These parts have been treated with epoxy and further painted to ensure impermeability to water.

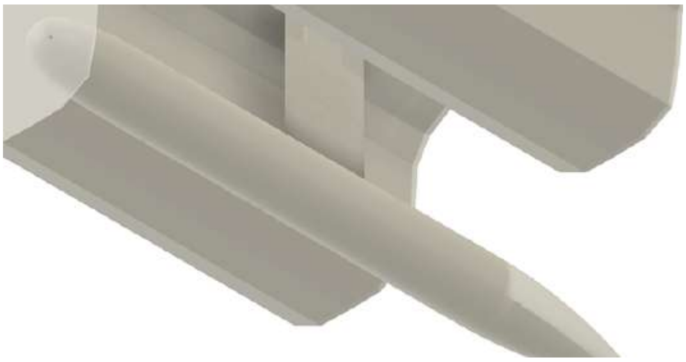

Figure 3: Airkeel 3D model

The top of the fin has another printed part which holds the whole body in place on the model boat with two bearings. These two bearings are connected to the motor.

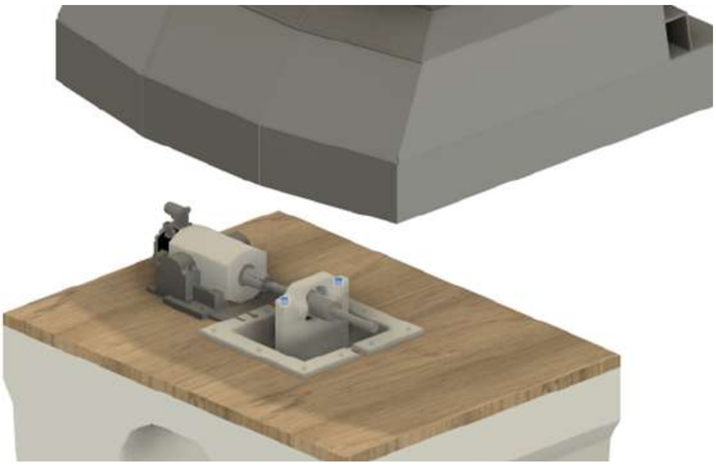

Figure 4: Airkeel connection to the motor

2.3 Motor

The Dunkermotoren BG66x50 motor is used to move the Airkeel from side to side. It moves the fin through the linear ball bearing screw. It has an integrated motor controller and an encoder used to read the position, velocity and torque of the motor through CAN bus.

Two limit switches are mounted with the Airkeel and the motor, one on each side of the Airkeel. The purpose of the limit switches is to provide safety for the motor and the Airkeel. If the motor moves the keel in a too wide angle, the limit switches will be activated and the motor will stop.

2.4 IMU

The Inertial Measurement Unit (IMU) is used to monitor the angular position of the boat. It is a TransducerM TM362 from SYD Dynamics.

Like the motor controller, it sends data to the microcontroller through CAN bus.

Figure 5: Inertial Measurement Unit (IMU)

2.5 Control Box

The control box is located on the deck of the CTV Model Boat. This box contains a microcontroller, a Teensy 4.1 and a Raspberry Pi 4, among other electronic components. On the side of the control box, a power button is placed, used to activate/deactivate the electrical components of the system by providing 24 V DC supply through the power line.

2.5.1 Teensy 4.1

The Teensy 4.1 microcontroller contains a C++ code that controls the motor. It reads data from the motor and sends commands through CAN bus. The Teensy 4.1 also communicates with the IMU via CAN bus to get the roll angle of the boat, and with the Raspberry Pi, where the control system to correct the roll angle of the boat is running, via Inter-Integrated Circuit (I2C).

Different actions can be performed by connecting the Teensy 4.1 to a laptop and using a serial monitor for user input. The most used ones are the following:

- Enable/disable motor
- Move the Airkeel to the zero position
- Move the motor +/- 10000 steps

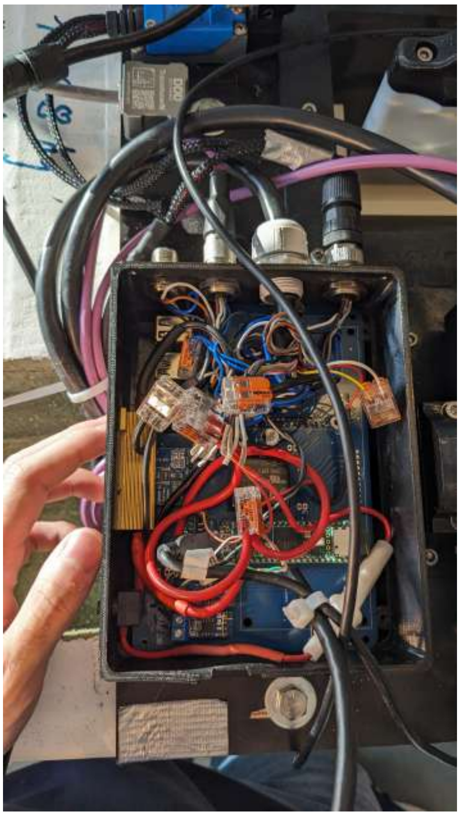

Figure 6: Control box

- Set current position as the new zero position
- Start/stop logging to save data in a microSD card
- Enable/disable debug to read data from the serial monitor
- Start/stop the control system to control the motor with the Simulink program running on the Raspberry Pi

2.5.2 Raspberry Pi 4

The Raspberry Pi 4 is used to run the control system. When the control system is running, the motor moves the Airkeel automatically to maintain the roll angle of the boat at 0º. To achieve this, the Airkeel is constantly moving from side to side to counteract the effect of the waves.

The control system is implemented in Simulink, which can be run on the Raspberry Pi from a laptop via Wi-Fi. To run the control system, both the laptop and the Raspberry Pi are connected to a hotspot with SSID "DACOMA".

2.6 Battery

The battery on the boat is a 6S Li-Ion battery which has nominal voltage of 25,2 V DC. It provides supply voltage to power every electrical component and has a built-in Battery Management System (BMS) which monitors the temperature across the pack and makes sure that the cell voltage is evenly distributed to provide the optimal battery performance.

There is also a DC-DC converter to convert the supply voltage from 25,2 V DC to 5 V DC, in order to provide the correct voltage for the Raspberry Pi, the Teensy, and the IMU.

A diagram with the system architecture is shown in Figure 7.

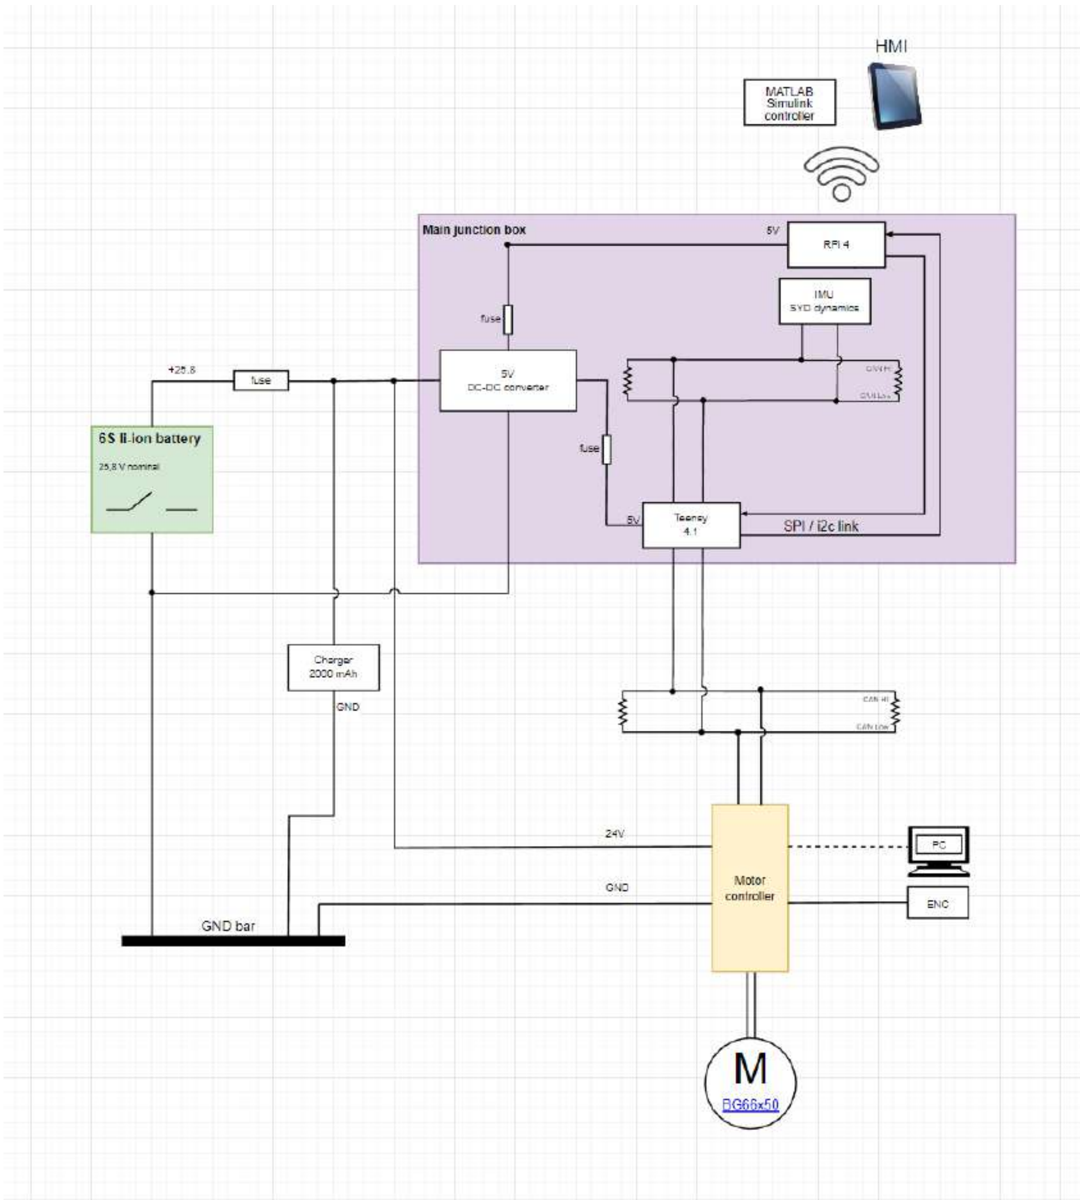

Figure 7: System architecture

3 Control System

The control system is implemented in MATLAB & Simulink. It was previously designed based on the model of the boat.

The designed controller is a Linear-Quadratic-Integral (LQI) controller, which is a type of advanced control system used to ensure the desired performance and stability of the system. LQI control extends the principles of the Linear-Quadratic Regulator (LQR) by incorporating integral action.

The LQI controller optimizes a quadratic cost function of the following form that balances performance and energy usage:

$$J = \int_0^\infty \left(\mathbf{x}^T \mathbf{Q} \mathbf{x} + \mathbf{u}^T \mathbf{R} \mathbf{u} \right) \, dt \tag{1}$$

where,

- J: Cost function, representing the trade-off between state deviations and control effort
- $\mathbf{x}$: State vector
- $\mathbf{u}$: Control input
- $\mathbf{Q}$: State weight matrix (penalizes state deviations)
- $\mathbf{R}$: Control input weight matrix (penalizes large control efforts)

The control input has the following form:

$$\mathbf{u}(t) = -\mathbf{K_x} \mathbf{x}(t) - K_I \int e(t) \, dt \tag{2}$$

where $\mathbf{K_x}$ and K_I are the feedback gains that minimize the cost function, and $e(t) = r(t) - y(t)$ is the tracking error (where $r(t)$ is the desired roll angle, in this case 0^{o}, and $y(t)$ is the actual roll angle of the boat).

The controller was initially pre-tuned based on previous observations, model calculations and simulations. There are a total of 6 parameters, which are elements of the weight matrices for the system states and the control input. The initial tuning of 5 of these parameters was based on Bryson's rule, which suggests that the weight for each state or control input should be proportional to the inverse square of the maximum expected value of that variable. Starting from this initial tuning, they were tuned further based on observations of the system.

4 Tests and Results

After the simulation tests, the goal was to test the performance of the real boat on the water and optimize the control system parameters under realistic model test conditions.

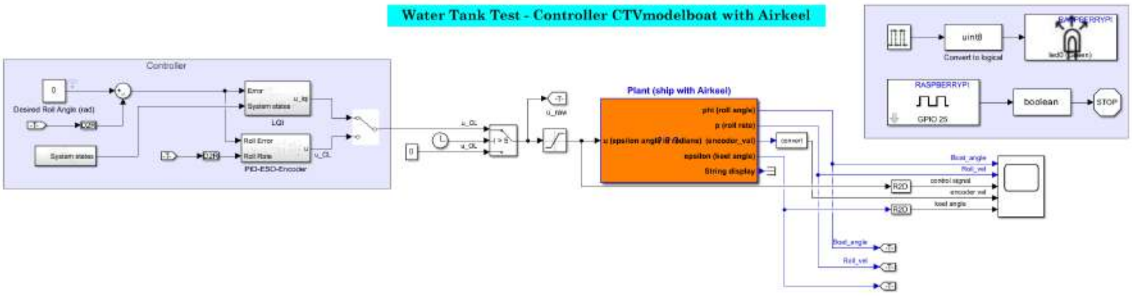

Figure 8: Simulink control system

During most of the testing process there was an issue with the motor. It was constantly stopping without any apparent reason, what made the tasks tedious and difficult. This problem is further explained in Section 4.3.

4.1 Harbor Test

The boat was first taken to a harbor area in Svendborg to test the control system with natural waves. The boat was first placed on calm water to observe the behavior of the boat with and without the Airkeel, and it was later taken to a place with waves to test the control system.

Figure 9: Boat in the water during the harbor test

Unfortunately, the control system did not damp the roll angle at all. The Airkeel moved all the way to one of the sides of the boat and it kept doing small movements in that side.

Furthermore, the data logging functionality implemented in the microcontroller was not very well programmed. As a result, the collected data were too chaotic and difficult to understand, so there were not clear conclusions.

After these tests, the code was improved so data collection is performed in a clearer way.

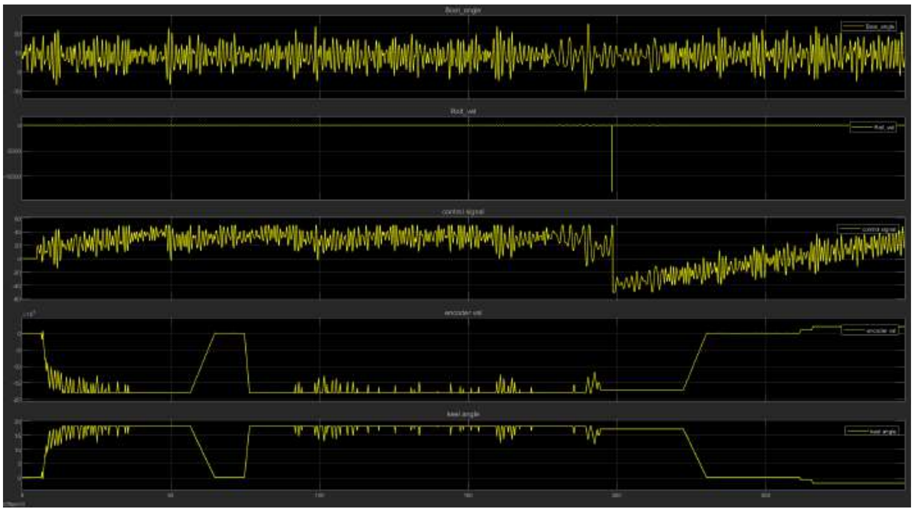

Figure 10: Control system harbor test. From top to bottom: Roll angle, roll velocity, control signal, encoder value and keel angle

4.2 Water Tank Tests

4.2.1 Control System Tests

After the unsuccessful experiments in the harbor, the following tests were carried out in a water tank with calm water, which made the testing process much easier and more convenient.

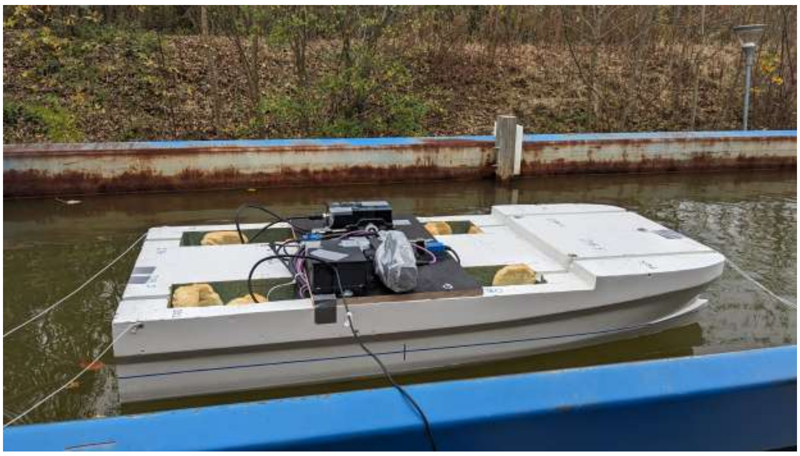

Figure 11: Boat in the water tank for testing

After a couple of days of testing in the water tank, the following results were obtained from different tests.

2kg of weight was placed at the port side, about 30 cm from the center line

Before starting the tests, 2 kg of weight was placed on the port side to level the boat, about 30 cm from the centerline, and the reference position for the Airkeel was set so that the boat's roll angle is 0º when the Airkeel is in the zero position.

Test 1: Airkeel fixed at zero position

For the first test, the airkeel was maintained at zero position to test whether the boat was properly leveled. The result was very good, with the roll angle quite stable at 0º.

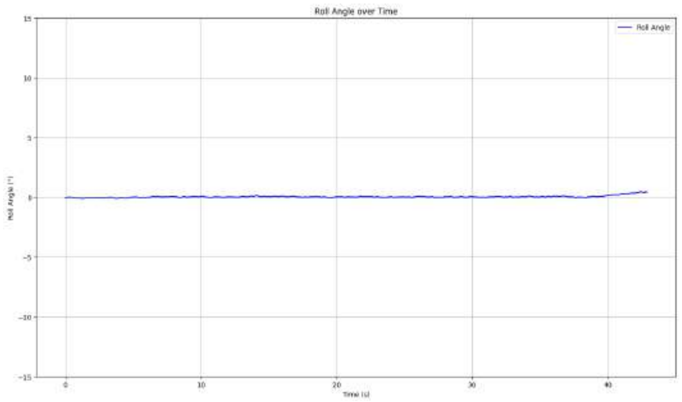

Figure 12: Roll angle with the Airkeel fixed at zero position

Test 2: Airkeel from zero position to starboard side

For this test, the Airkeel was moved from the zero position to the starboard side to see how it affects the roll angle of the boat. The roll angle increased up to around 1.5º for -18º of the Airkeel.

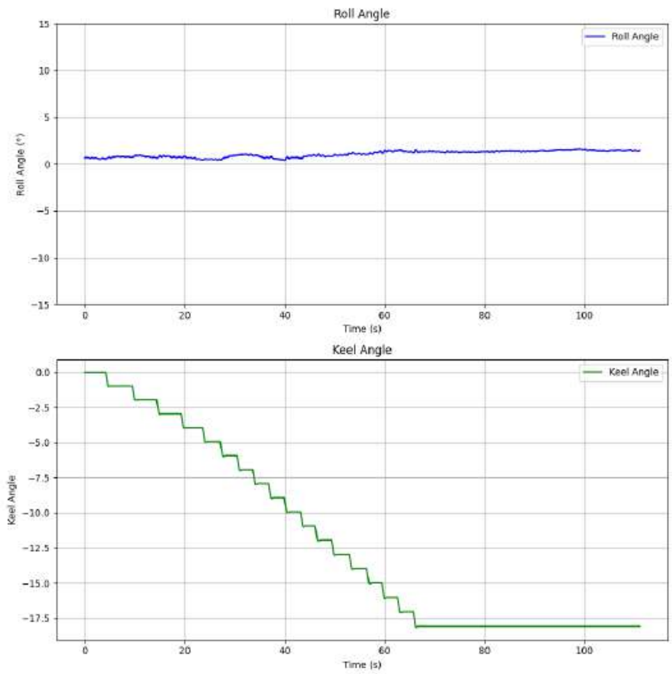

Figure 13: Roll and keel angles with the Airkeel from zero position to starboard side

Test 3: Airkeel from zero position to port side

This time, the Airkeel was moved from the zero position to the port side, getting similar results. The roll angle decreased to around -1.5º for 18º of the Airkeel.

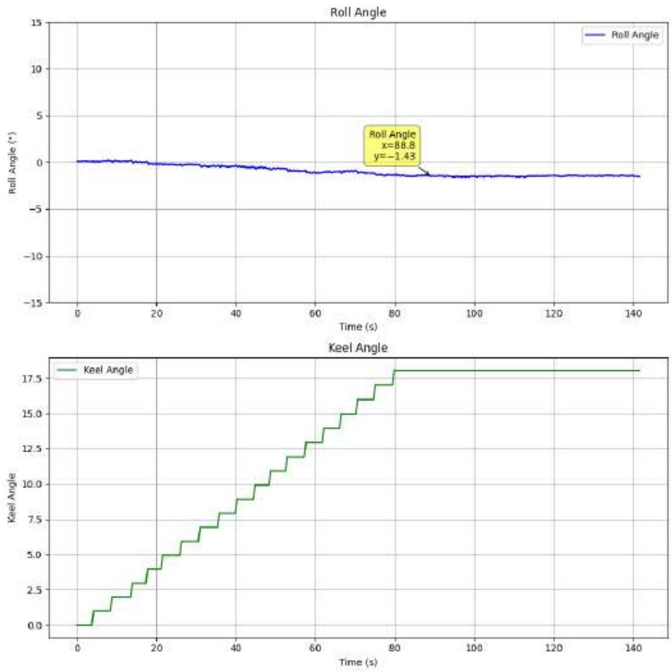

Figure 14: Roll and keel angles with the Airkeel from zero position to port side

Test 4: Waves without control system

For this test, the boat was pushed down from the starboard side and the roll angle was registered with the Airkeel fixed at zero position, in order to observe the effect of the waves on the boat without any damping.

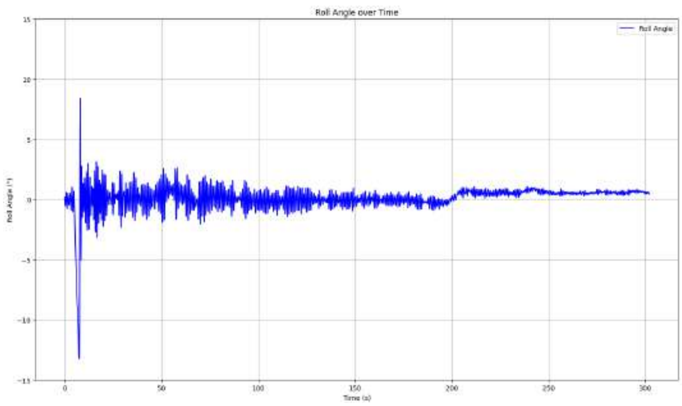

Figure 15: Roll and keel angles with waves and the Airkeel fixed at zero position

Test 5: LQI controller

For this test, the boat was pushed down from the starboard side and the control system was started shortly after. It can be observed in the graph that the control system did not damp the oscillations but excited roll motions instead, which is the same behavior that it showed during the harbor tests.

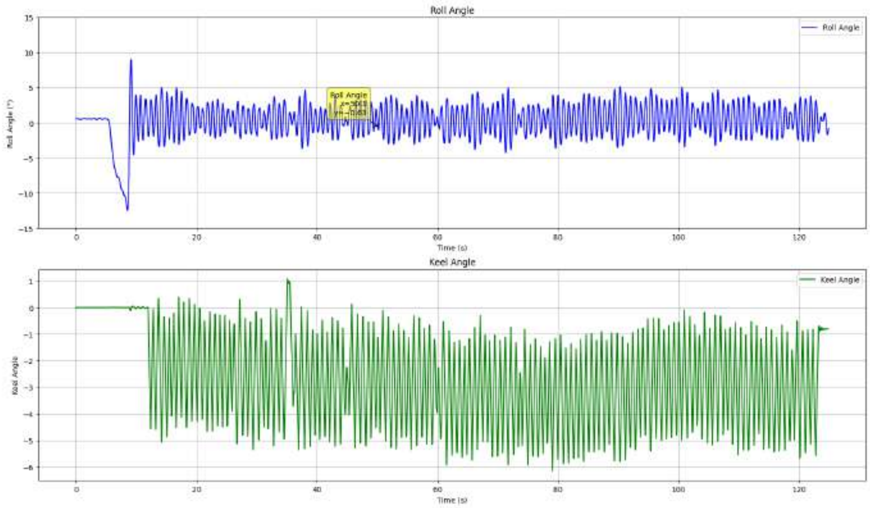

Figure 16: Roll and keel angles with the LQI controller running

Test 6: LQI controller with inverted sign

For the next test, the sign of the controller was inverted to change the phase and the same experiment was carried out. The boat was released at second 11 and the control system was started 2 seconds later. This time, it worked much better, but the controller was a bit slow. After 55 seconds, the motor stopped working.

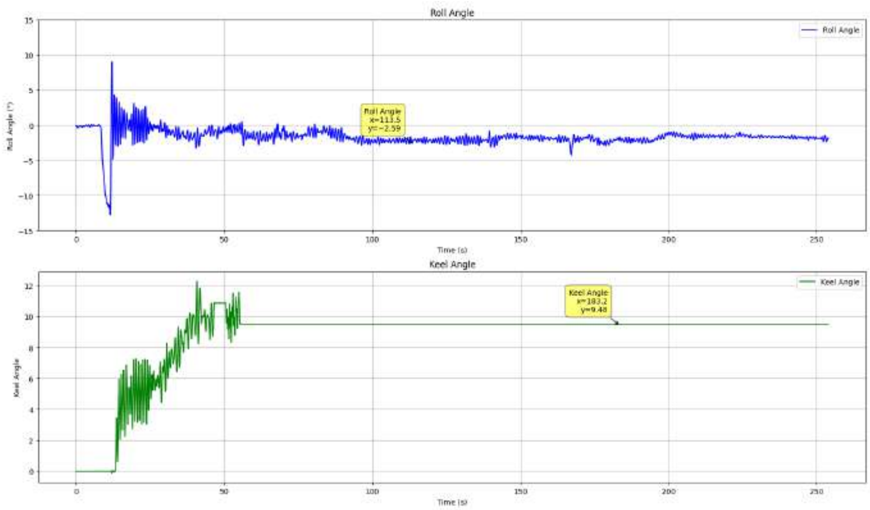

Figure 17: Roll and keel angles with the LQI controller with inverted sign running

Test 7: LQI controller with inverted sign (2nd time)

This test was performed again. This time, the boat was released at second 8 and the control system was started 2 seconds later. The boat was pushed again at second 23, when the boat was already more or less stable. The motor stopped working at second 30.

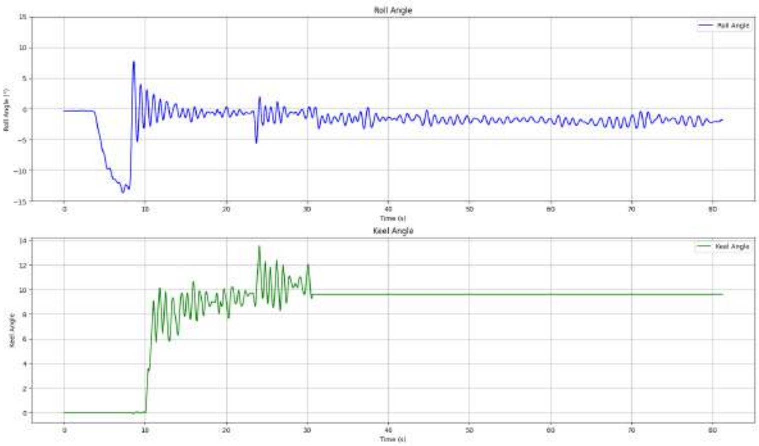

Figure 18: Roll and keel angles with the LQI controller with inverted sign running (2nd time

Test 8: LQI controller with inverted sign (3rd time)

The test was repeated one last time. The boat was released at second 5.5 and the control system was started 2 seconds later. This time, the boat was pushed again several times after the boat stabilized, at seconds 22.5, 30 and 44.5. The motor stopped working at second 55.

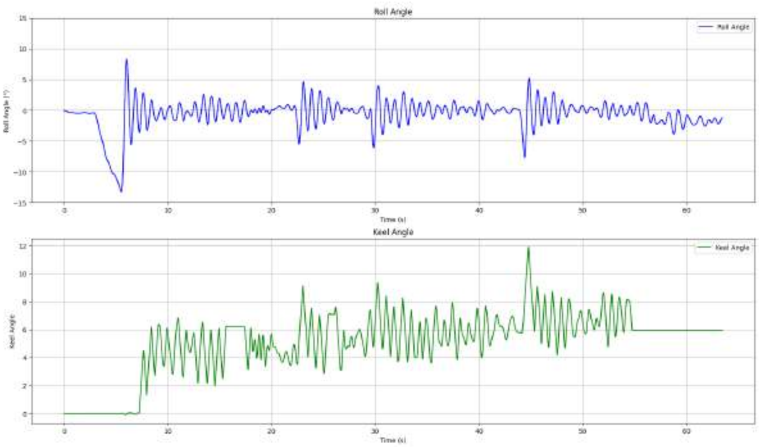

Figure 19: Roll and keel angles with the LQI controller with inverted sign running (3rd time

Discussion

The Airkeel-boat system can be described as a damped harmonic oscillator. Hence, its dynamics depend on the phase between the roll motion and the Airkeel motion. If the phase is wrong, the system excites roll motions, like in a driven harmonic oscillator.

In a driven harmonic oscillator, the phase of the driver and the phase of the oscillator are not always the same. At low frequencies, the oscillation is in phase with the driver, but at high frequencies, the oscillation and the driver are out of phase.

The results of the tests, that correspond to higher frequencies, suggest that the controller is out of phase for those frequencies, as it shows a good behavior when changing its phase by inverting the sign.

The system behaves differently for different frequencies, and the controller should work properly for both low and high frequencies, and it is not doing so.

The transfer function of the system was obtained for a different load scenario, which obviously changes the performance. Most likely, this is due to a shift of the natural roll frequency, where the phase change between roll motion and Airkeel motion is most sensitive. A slight shift of the natural roll frequency can hence turn the system from the damping motion to ”anti-damping”, which means exciting roll motions.

Therefore, the next step is to perform a sine wave test and get a new transfer function to ensure that it corresponds to the system as it is in the moment.

4.2.2 Sine Wave Test

The sine wave test is an experiment in which the system is excited with an input sine signal with different amplitudes and frequencies, and the response is registered to analyze the behavior of the system.

With this information, a transfer function that accurately represents the system can be obtained.

The sine wave test was carried out on the water tank and the result can be seen in Figure 20. There was some wind and there were some problems with the motor because it stopped for a few seconds at some point, so it may be necessary to repeat the test.

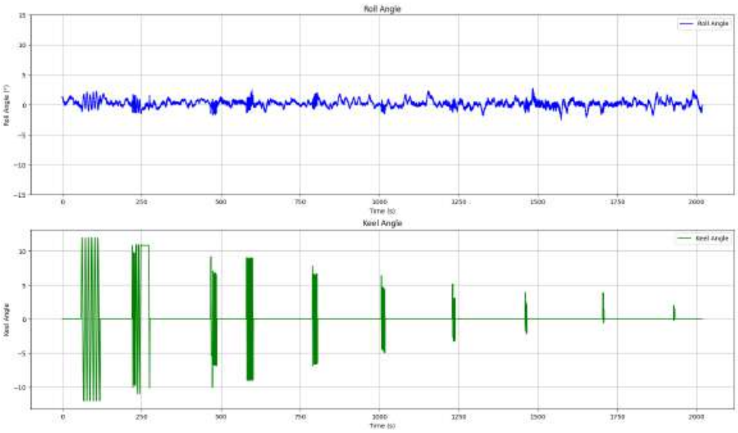

Figure 20: Roll and keel angles for the sine wave test

4.3 Motor issue

There was a persistent issue with the motor that was very problematic, as it was making the data collection almost impossible sometimes, because it was unexpectedly stopping in the middle of the tests.

It was possible to find a hint using the software Drive Assistant 5 from Dunkermotoren. With this software, it is possible to connect to the motor using a serial CAN bus dongle. There was some issue with the device firmware, but it was still quite difficult to solve, as there was also problems to update it.

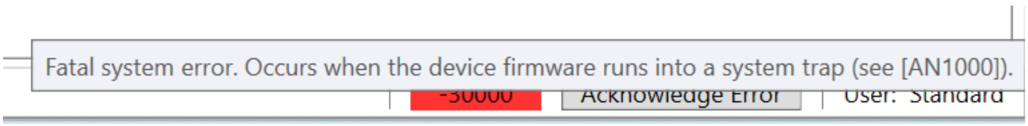

Figure 21: Motor error

After a few emails and attempts to contact the manufacturer, it was possible to arrange a meeting with them and they could finally solve the issue via TeamViewer.

5 Conclusion

After a long time without usage, the CTV Model Boat was set up again as a stable platform to perform tests that can be useful to improve the control system. One of the main problems at the moment is that it is not easy to move the boat, so it takes a lot of time to prepare the experiments. For this reason, it would be very advantageous to have a more accessible water tank, so more time can be used for the actual tests instead of preparing for them.

Another problem that caused a significant loss of time was the issue with the motor. This and other smaller problems that emerged during the tests were solved, so it is now a reliable system that makes testing easier.

The previously designed control system was tested, and it was found that it is not working properly, so future tests and improvements need to be done. To continue with the work, the first step is performing more sine wave tests with different load scenarios, as the behavior of the system changes with different loads. Once some results are obtained, the idea is to design an adaptive controller that can adjust to different load scenarios, using a machine learning algorithm that would learn from the obtained data and adjust the controller to any possible scenario.