IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Limiting Kinetic Energy through Control Barrier Functions: Analysis and Experimental Validation

Federico Califano[1], Daniël Logmans, and Wesley Roozing[1]

***Abstract*— In the context of safety-critical control, we propose and analyse the use of Control Barrier Functions (CBFs) to limit the kinetic energy of torque-controlled robots. The proposed scheme is able to modify a nominal control action in a minimally invasive manner to achieve the desired kinetic energy limit. We show how this safety condition is achieved by appropriately injecting damping in the underlying robot dynamics independently of the nominal controller structure. We present an extensive experimental validation of the approach on a 7-Degree of Freedom (DoF) Franka Emika Panda robot. The results demonstrate that this approach provides an effective, minimally invasive safety layer that is straightforward to implement and is robust in real experiments.**

I. Introduction

Collaborative robots, or *cobots*, are gaining traction across a wide range of applications. Safety is a critical control objective when these robots share space with humans [1], [2]. The recent rise of learning-based controllers, which typically only provide probabilistic safety guarantees, underscores the need for safety-critical approaches [3]. ISO standards [4] attempt to formalise safety hazards in this setting, and their mitigation is an active research area. Some works prevent interaction, by enforcing a speed-dependent separation distance between robot and operator, assuming reliable detection [5], [6]. Other works limit interaction power and force by implementing, e.g., impedance control [7], [8]. Yet other approaches explore dynamic human-robot impact scenarios and relate the impact velocity to the risk of injury [9].

In this work safety is addressed by bounding the kinetic energy that could potentially be transferred to a human operator, in order to prevent injury in collision scenarios. The importance of this choice is backed by numerous publication which relate directly relevant safety metrics to the energy flow generated from the interaction [2], [10]. Furthermore, the *power and force limiting* (PFL) conditions in the ISO/TS 15066 [4], which are the only collaborative conditions in which contact between humans and robots are considered, are addressed through energetic constraints.

We propose a method that takes the form of a *safety filter*, enforcing a bound on maximum kinetic energy while minimally altering a desired control input. We make use of *Control Barrier Functions* (CBFs), a safety-critical control algorithm able to constrain the robot to a region of its state space representing safe operating conditions [11]. Most CBF implementations in robotics apply to safety-critical kinematic control (i.e., tasks in which the safety constraint represents obstacle avoidance conditions) and rely on lower-level controllers to handle system dynamics [12]–[14]. Instead, we investigate the use of energy-based CBFs and, different from previous works such as [15]–[17], utilize them to directly limit the kinetic energy of a torque-controlled robot.

We recognise relevant related works proposing schemes to limit the kinetic or total energy of torque-controlled manipulators for safety objectives. These works are motivated by *energy-aware* and *passivity* arguments [18]–[20], stressing the fact that safety measures are closely related to energy- and power-based metrics. The recent work [21] presents a control algorithm that is able to limit the kinetic energy, achieved by using higher-order CBFs in a system augmented with energy tanks [22], used to enforce passivity of the overall scheme. Other approaches attempt to limit kinetic energy [23] and total energy [24] of controlled robots, also using energy-tank based arguments to recover passivity.

In this work, we present a novel approach that avoids considering passivity as a strict constraint to be achieved at design phase. Instead, we achieve the kinetic energy bound directly through the proposed CBF-based algorithm. We analytically and experimentally show that the proposed CBF operates solely by injecting damping into the system, ensuring that the safety-critical control action inherently preserves the passivity of any nominal passive closed-loop system. This eliminates the need for supplementary tools such as energy tanks, making the proposed scheme significantly simpler than most of the state-of-the-art solutions. The energetic analysis in this work constitutes a contribution on its own, even if from a CBF-based methodological perspective, the used CBF is analytically related (and shares technical advantages) to those used in [16], [17]. Finally, although with a different controller, [17], [25] demonstrate passivity in a kinematic safety-critical setting.

The main contributions of this paper are:

1) A kinetic energy-limiting CBF-based safety filter and analysis of its energetic properties.
2) Extensive experimental validation on a 7-DoF robot manipulator of the proposed safety-critical control system.

The remainder of this paper is outlined as follows. Sec. II contains the mathematical background and analysis involving the specific CBF used in this work. We present extensive experimental results in four scenarios in Sec. III. Finally, Sec. IV and V discuss the results and conclude the paper.

II. Methods

In this section, we introduce the CBF control approach (we refer to [11], [26], [27] and references therein for a detailed

[1]Robotics & Mechatronics (RaM), Robotics Centre, University of Twente, The Netherlands. E-mail: {w.roozing,f.califano} at utwente.nl.

overview), followed by the specific CBF used in this work.

We denote vectors, matrices, vector-valued maps and matrix-valued maps with bold font letters; scalars and real-valued maps with regular letters; and sets with calligraphic letters. The norm of a vector is denoted by $||\cdot||$.

A. Background

1) Control Barrier Functions: Consider a nonlinear control-affine system in standard form

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{g}(\boldsymbol{x})\boldsymbol{u} \tag{1}$$

with system state $\boldsymbol{x} \in \mathcal{D} \subset \mathbb{R}^n$ and control input $\boldsymbol{u} \in \mathcal{U} \subset \mathbb{R}^m$. All variables are assumed to have a degree of continuity such that the right-hand side of (1) is locally Lipschitz, to guarantee the existence and uniqueness of the solutions.

Control barrier functions (CBFs) serve to achieve *forward invariance* of a set $\mathcal{S}$, referred to as *safe set*, i.e.,

$$\forall \boldsymbol{x}(0) \in \mathcal{S} \implies \boldsymbol{x}(t) \in \mathcal{S} \ \ \forall t > 0. \tag{2}$$

The safe set $\mathcal{S}$ is built as the superlevel set of a continuously differentiable function $h : \mathcal{D} \to \mathbb{R}$, i.e.,

$$\mathcal{S} = \{\boldsymbol{x} \in \mathcal{D} : h(\boldsymbol{x}) \geq 0\}.$$

The function $h(\boldsymbol{x})$ is then defined as a CBF on $\mathcal{D}$ if $\partial_{\boldsymbol{x}} h(\boldsymbol{x}) \neq 0, \forall \boldsymbol{x} \in \partial\mathcal{S}$ and

$$\sup_{\boldsymbol{u}\in\mathcal{U}} \underbrace{\left[\frac{\partial h}{\partial \boldsymbol{x}}\boldsymbol{f}(\boldsymbol{x}) + \frac{\partial h}{\partial \boldsymbol{x}}\boldsymbol{g}(\boldsymbol{x})\boldsymbol{u}\right]}_{\dot{h}(\boldsymbol{x},\boldsymbol{u})} \geq -\alpha(h(\boldsymbol{x})) \tag{3}$$

for all $\boldsymbol{x} \in \mathcal{D}$ and some *extended class $\mathcal{K}$ function*[1] α. Here we denote the term in the square bracket, i.e., the variation of h along the solution of (1), by $\dot{h}(\boldsymbol{x}, \boldsymbol{u})$.

The link between the existence of a CBF and the forward invariance of the related safe set is established by the following key result.

Theorem 1 ([26])**.** *Let $h(\boldsymbol{x})$ be a CBF on $\mathcal{D}$ for* (1)*. Any locally Lipschitz controller $\boldsymbol{u} = \boldsymbol{k}(\boldsymbol{x})$ such that $\dot{h}(\boldsymbol{x}, \boldsymbol{u}) \geq -\alpha(h(\boldsymbol{x}))$ provides forward invariance of the safe set $\mathcal{S}$. Additionally the set $\mathcal{S}$ is asymptotically stable on $\mathcal{D}$.*

The way controller synthesis induced by CBFs are implemented in practice is to use them as *safety filters* (Fig. 1), transforming a nominal state-feedback control input $\boldsymbol{u}_{\text{nom}}(\boldsymbol{x})$ into a new state-feedback control input $\boldsymbol{u}(\boldsymbol{x})$ in a minimally invasive fashion in order to guarantee forward invariance of $\mathcal{C}$. In practice, the following Quadratic Program (QP) is solved:

$$\begin{aligned} \boldsymbol{u}(\boldsymbol{x}) = &\underset{\boldsymbol{u}\in\mathcal{U}}{\text{argmin}} \quad ||\boldsymbol{u} - \boldsymbol{u}_{\text{nom}}(\boldsymbol{x})||^2 \\ &\text{s.t.} \quad \dot{h}(\boldsymbol{x}, \boldsymbol{u}) \geq -\alpha(h(\boldsymbol{x})) \end{aligned} \tag{4}$$

The transformation of the nominal control input $\boldsymbol{u}_{\text{nom}}(\boldsymbol{x})$ into the new state-dependent control $\boldsymbol{u}(\boldsymbol{x})$ by solving (4) is referred to as *safety-critical control*, and the fact that the constraint is linear in the input allows for efficient real-time implementations of such a controller. In this work we will take advantage of the additive decomposition of the safety-critical control:

$$\boldsymbol{u}(\boldsymbol{x}) = \boldsymbol{u}_{\text{nom}}(\boldsymbol{x}) + \boldsymbol{u}_{\text{safe}}(\boldsymbol{x}), \tag{5}$$

resulting from the solution of (4), where clearly $\boldsymbol{u}_{\text{safe}}(\boldsymbol{x}) = 0$ when $\dot{h}(\boldsymbol{x}, \boldsymbol{u}) \geq -\alpha(h(\boldsymbol{x}))$.

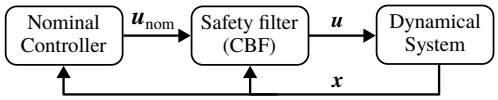

Fig. 1: CBF-based safety filter.

2) Robots and their energetic analysis: In this work we will consider fully actuated n-Degree-of-Freedom torque-controlled robots, represented as a $2n$-dimensional system whose state belongs to the tangent bundle of the robot's configuration manifold, expressed in the usual Lagrangian canonical coordinates $\boldsymbol{q} \in \mathbb{R}^n$ and $\dot{\boldsymbol{q}} \in \mathbb{R}^n$. The system dynamics are expressed by the Euler-Lagrange equations:

$$\boldsymbol{D}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}} + \boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{B}\boldsymbol{u}, \tag{6}$$

with $\boldsymbol{D}(\boldsymbol{q})$ the inertia matrix, $\boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})$ the Coriolis matrix, $\boldsymbol{g}(\boldsymbol{q})$ the gravity vector, and $\boldsymbol{B}$ the full rank actuation matrix. System (6) admits a representation as a control affine system (1) when choosing $\boldsymbol{x} = (\boldsymbol{q}^\top, \dot{\boldsymbol{q}}^\top)^\top \in \mathbb{R}^{2n}$, and as such CBF-based algorithms can be applied.

In the sequel we will perform analysis involving the kinetic energy:

$$K_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = \frac{1}{2}\dot{\boldsymbol{q}}^\top \boldsymbol{D}(\boldsymbol{q})\dot{\boldsymbol{q}}. \tag{7}$$

Using (6) and the skew symmetry of the matrix $\dot{\boldsymbol{D}} - 2\boldsymbol{C}$ (see e.g., [16]), it is straightforward to verify that the rate of change of the kinetic energy along solutions of (6) verifies:

$$\dot{K}_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = -\dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) + \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}, \tag{8}$$

where the last term $\dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}$ represents the instantaneous mechanical power that the controller injects into the robot.

B. Bounding kinetic energy with CBFs

Since we aim to set an upper bound $K_{\max}$ on the robot's kinetic energy, we propose to encode the safe set $\mathcal{S} = \{(\boldsymbol{q}^\top, \dot{\boldsymbol{q}}^\top)^\top \in \mathbb{R}^{2n} : K_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) \leq K_{\max}\}$ through the CBF

$$h(\boldsymbol{q}, \dot{\boldsymbol{q}}) = K_{\max} - K_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = K_{\max} - \frac{1}{2}\dot{\boldsymbol{q}}^\top \boldsymbol{D}(\boldsymbol{q})\dot{\boldsymbol{q}}. \tag{9}$$

Using (8), it immediately follows that

$$\dot{h}(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}) = \dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) - \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}. \tag{10}$$

If (9) is a valid CBF for the controlled robot, then Theorem 1 guarantees that, if we are able to solve (4), the kinetic energy limit is always respected.

For the following analysis, let us define the quantity

$$\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}) = \dot{h}(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}) + \alpha(h(\boldsymbol{q}, \dot{\boldsymbol{q}})), \tag{11}$$

[1] A function $\alpha : (-b, a) \to (-\infty, \infty)$ with $a, b > 0$, which is continuous, strictly increasing, and $\alpha(0) = 0$.

so that the constraint in (4) reads $\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}) \geq 0$, or, when expanded using (9) and (10):

$$-\dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u} + \dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) \geq -\alpha\left(K_{\max} - \frac{1}{2}\dot{\boldsymbol{q}}^\top \boldsymbol{D}(\boldsymbol{q})\dot{\boldsymbol{q}}\right). \quad (12)$$

Notice that the left hand side of this expression is the negative of the net power flowing into the robot, due to gravitational effects and due to the control $\boldsymbol{u}$.

The following result states an interesting fact: when using a CBF in the form (9), the safety-critical component $\boldsymbol{u}_{\text{safe}}$ in the decomposition (5) only injects negative power into the underlying closed-loop system (6) controlled with $\boldsymbol{u}_{\text{nom}}$.

Theorem 2. *Let* (9) *be the CBF acting on a system* (6)*, controlled with nominal input* $\boldsymbol{u}_{\text{nom}}(\boldsymbol{q}, \dot{\boldsymbol{q}})$ *and resulting in the safety-critical decomposition* (5)*. The total power injected by the safety filter is always non-positive, that is,*

$$P_{\text{safe}} := \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}_{\text{safe}} \leq 0. \quad (13)$$

Proof. We distinguish two cases. First, if $\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) \geq 0$, the safety constraint in (4) is satisfied with the trivial solution $\boldsymbol{u} = \boldsymbol{u}_{\text{nom}}$, and as such $\boldsymbol{u}_{\text{safe}} = 0$ and $P_{\text{safe}} = 0$.

Secondly, we consider the case $\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) < 0$. We need to show that (13) holds, which can be rewritten using (5) as

$$\dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u} \leq \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}_{\text{nom}}. \quad (14)$$

The CBF enforces the inequality (12), which is rewritten as:

$$\dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u} \leq \alpha(h(\boldsymbol{q}, \dot{\boldsymbol{q}})) + \dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}). \quad (15)$$

The task is now to show that (15) implies (14). A sufficient condition for the latter proposition to be true is that:

$$\alpha(h(\boldsymbol{q}, \dot{\boldsymbol{q}})) + \dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) \leq \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}_{\text{nom}}, \quad (16)$$

which can be rewritten as

$$\dot{h}(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) + \alpha(h(\boldsymbol{q}, \dot{\boldsymbol{q}})) = \Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) \leq 0, \quad (17)$$

which is true by assumption. □

It is worth noticing that the previous result holds independently on the specific control law $\boldsymbol{u}_{\text{nom}}(\boldsymbol{q}, \dot{\boldsymbol{q}})$: whatever design is chosen for the nominal controller, if a CBF in the form (9) is used, the safety-critical component of the closed-loop system $\boldsymbol{u}_{\text{safe}}$ will always act in a way to extract mechanical energy from the system.

Remark 1 (Relation to energy-based CBFs in [16])**.** *The CBF* (9) *shares some properties with so called energy-based CBFs as introduced in [16] for safety-critical kinematic control, defined as* $h(\boldsymbol{q}, \dot{\boldsymbol{q}}) = -K_e(\boldsymbol{q}, \dot{\boldsymbol{q}}) + \beta\bar{h}(\boldsymbol{q})$ *with* $\beta > 0$*. In particular the proposed CBF (9) shares a technical advantage with the one in [16]: the safety constraint* (12) *is independent of the Coriolis matrix, reducing model dependence and computational complexity while solving* (4)*.*

Remark 2 (Passivity/Stability preservation property of (9))**.** *A convenient consequence of a negative power injection by the safety-critical control component* $\dot{\boldsymbol{q}}^\top \boldsymbol{B}\,\boldsymbol{u}_{\text{safe}} \leq 0$ *is the following: if the controlled system with* $\boldsymbol{u}_{\text{nom}}(\boldsymbol{q}, \dot{\boldsymbol{q}})$ *is passive (or stable), then the safety-critical control preserves the closed-loop passivity (or stability) properties of the nominal controller. This fact allows for assessing passivity of the critically controlled closed-loop system without the use of extra passivising framework such as energy tanks as done in e.g., [21], [24]. In [28] a variation of this result (with different generality and different proof) was indeed given in the framework of passivity-based control.*

Another important result in the CBF framework is that, under the conditions of Theorem 2, together with the extra assumption that $\mathcal{U} = \mathbb{R}^m$, the decomposition (5) admits the closed-form solution [16], [27]:

$$\boldsymbol{u} = \boldsymbol{u}_{\text{nom}} + \begin{cases} \frac{\boldsymbol{B}^\top \dot{\boldsymbol{q}}}{||\boldsymbol{B}^\top \dot{\boldsymbol{q}}||^2}\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) & \text{if } \Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) < 0\,, \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

This expression is useful since it induces an analytic expression for the power injected in the system by the controller, clearly displaying the role of the function $\alpha(\cdot)$ in the CBF algorithm. For example, it is simple to see that using (18) the expression (13) becomes $P_{\text{safe}} = \Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}})$ (when $\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) < 0$), providing a quantitative measure on how much damping the safety-critical control injects into the system and how it can be modulated with $\alpha(\cdot)$. In the next section we will use (18) to perform a thorough power analysis involving unmodelled external interactions.

C. External interaction forces

The dynamic system (6) does not include external torques caused by disturbances or interaction. For an external generalised force vector $\boldsymbol{\tau}_{\text{ext}} \in \mathbb{R}^n$, the system becomes:

$$\boldsymbol{D}(\boldsymbol{q})\ddot{\boldsymbol{q}} + \boldsymbol{C}(\boldsymbol{q}, \dot{\boldsymbol{q}})\dot{\boldsymbol{q}} + \boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{B}\boldsymbol{u} + \boldsymbol{\tau}_{\text{ext}} \quad (19)$$

where typically the external torques are expressed as $\boldsymbol{\tau}_{\text{ext}} = \boldsymbol{J}(\boldsymbol{q})^\top \mathbf{f}$, where $\mathbf{f}$ are interaction forces applied at the end-effector of the robot and $\boldsymbol{J}(\boldsymbol{q})$ denotes the end-effector Jacobian matrix.

Obtaining these external interaction forces and including them in the model leads to the substitution of (12) with the new constraint:

$$\underbrace{\dot{\boldsymbol{q}}^\top(-\boldsymbol{B}\boldsymbol{u} - \boldsymbol{\tau}_{\text{ext}} + \boldsymbol{g}(\boldsymbol{q}))}_{\dot{h}(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u})} \geq -\alpha\left(K_{\max} - \frac{1}{2}\dot{\boldsymbol{q}}^\top \boldsymbol{D}(\boldsymbol{q})\dot{\boldsymbol{q}}\right). \quad (20)$$

The challenge lies in measuring $\boldsymbol{\tau}_{\text{ext}}$, or, in practice, finding a good estimate $\hat{\boldsymbol{\tau}}_{\text{ext}}$. If accurately estimated, the previously presented control scheme using (9) and (20) will keep the kinetic energy below the desired limit. We refer to this case in the sequel as *interaction-aware*. If these torques are *not* taken into account in the model (i.e., the the CBF algorithm is implemented with the constraint (12)), system invariance cannot in general be guaranteed. We refer to this case in the sequel as *interaction-agnostic*. In the following we state a result that gives insight on the behaviour of the controlled system in interaction-agnostic case.

Theorem 3. *Consider system* (19) *with unknown external torques* $\boldsymbol{\tau}_{\text{ext}}$ *producing a positive power inflow* $P_{\text{ext}} = \dot{\boldsymbol{q}}^\top \boldsymbol{\tau}_{\text{ext}} \geq 0$. *Let the system* (19) *be controlled with the CBF* (9) *and assume* $\mathcal{U} = \mathbb{R}^m$. *If the system converges to a positive constant kinetic energy value with* $\boldsymbol{u}_{\text{safe}} \neq 0$, *the kinetic energy error* $K_{\text{e}} - K_{\max}$ *converges to the relation:*

$$K_{\text{e}} - K_{\max} = \alpha^{-1}(P_{\text{ext}}). \tag{21}$$

Proof. The variation of kinetic energy of (19) results in:

$$\dot{K}_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = -\dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) + \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u} + P_{\text{ext}}, \tag{22}$$

where $P_{\text{ext}} > 0$ contributes to a transient increase of kinetic energy. Substituting the closed-form solution (18) (for the case $\Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) < 0$ because $\boldsymbol{u}_{\text{safe}} \neq 0$) and assuming non-zero kinetic energy (i.e., $\dot{\boldsymbol{q}} \neq 0$) one obtains:

$$\dot{K}_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}}) = -\dot{\boldsymbol{q}}^\top \boldsymbol{g}(\boldsymbol{q}) + \dot{\boldsymbol{q}}^\top \boldsymbol{B}\boldsymbol{u}_{\text{nom}} + \Psi(\boldsymbol{q}, \dot{\boldsymbol{q}}, \boldsymbol{u}_{\text{nom}}) + P_{\text{ext}}. \tag{23}$$

Now, using (10) and (11), the power balance simplifies to[2]

$$\dot{K}_{\text{e}} = \alpha(K_{\max} - K_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}})) + P_{\text{ext}}\,, \tag{24}$$

which in steady-state condition ($\dot{K}_{\text{e}} = 0$) reduces to

$$\alpha(K_{\max} - K_{\text{e}}(\boldsymbol{q}, \dot{\boldsymbol{q}})) = -P_{\text{ext}}, \tag{25}$$

from which (21) follows, due to the invertibility of the class $\mathcal{K}$ function $\alpha(\cdot)$. Notice that with $\alpha(h) = \gamma\, h$ this describes first-order dynamics for the kinetic energy error. □

Remark 3. *The last result gives insight on the role of the class* $\mathcal{K}$ *function* α. *For example, as easily demonstrated using a linear function* $\alpha(h) = \gamma h$ *with* $\gamma > 0$, *an increase in* γ *reduces the kinetic energy error that is incurred due to unmodelled power flows, as* $K_{\text{e}} - K_{\max} \leq P_{\text{ext}}/\gamma$. *At a design stage, this fact needs to be traded with the advantages of choosing a lower value of* γ. *Intuitively, lowering* γ *induces a more conservative behaviour to achieve invariance in nominal conditions, tending to push the system state towards the safe set before reaching its boundary (see e.g., [11]). As a consequence, lowering* γ *corresponds to smoother closed-loop behaviour in nominal conditions, but also to poorer rejection in case of unmodelled external disturbances.*

III. Experimental results

To assess the practical applicability of the approach presented in this paper in the context of imperfect torque tracking, sensor errors, discrete-time control, and communication delays, we present extensive experimental results. This comprises four different experiments:

1) A Cartesian step response;
2) Contact loss with the environment;
3) External power input by human-robot interaction;
4) Kinetic energy error validation (Theorem 3).

The nominal control action for the first two experiments is generated by an underdamped Cartesian impedance controller with a stiffness of 200 N/m and a damping of 6 N s m^{-1}. For the latter two experiments, the nominal controller is deactivated[3]. For all experiments, we choose $\alpha(h) = \gamma h$, and we will investigate its influence: when decreasing γ, we expect increasingly conservative behaviour (see Remark 3).

A. Experimental Setup

Experiments are performed on a Franka Emika Panda 7-DoF robotic arm. The Franka Control Interface (FCI) provides a ROS-interface for joint torque commands at 1000 Hz, with built-in gravity and friction compensation active by default. As a result, we set $\boldsymbol{g}(\boldsymbol{q}) = \boldsymbol{0}$ in Eqs. (6) and (19). The interface provides the inertia tensor $\boldsymbol{D}(\boldsymbol{q})$, external torque estimate $\hat{\boldsymbol{\tau}}_{\text{ext}}$ and state information $\boldsymbol{q}$, $\dot{\boldsymbol{q}}$. A schematic representation of the architecture is shown in Fig. 2.

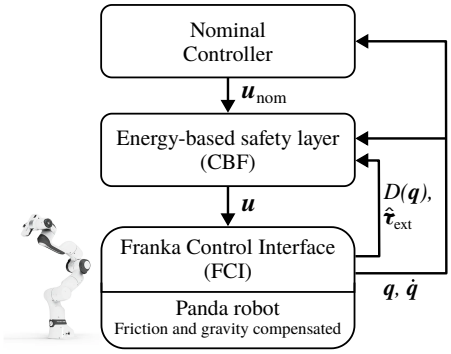

Fig. 2: Control architecture of the experimental setup.

To reduce the effect of sensor noise on the velocity estimates, we use a discrete joint acceleration rate limiter:

$$\dot{\boldsymbol{q}}_k = \dot{\boldsymbol{q}}_{k-1} + \min(\max(\dot{\boldsymbol{q}}_k - \dot{\boldsymbol{q}}_{k-1}, -\Delta_{\text{t}}\ddot{\boldsymbol{q}}_{\max}), \Delta_{\text{t}}\ddot{\boldsymbol{q}}_{\max}),$$

where the maximum joint acceleration $\ddot{\boldsymbol{q}}_{\max}$ is set to the robot's documented limits plus a 20% margin. Δ_{t} denotes the time interval between consecutive measurements $\dot{\boldsymbol{q}}_{k-1}$ and $\dot{\boldsymbol{q}}_k$. The benefit of an acceleration saturation filter is that it does not introduce delay and enforces an upper bound on the noise amplitude without attenuating the signal itself.

Despite the existence of an analytical solution (18) to the QP (4), our implementation leverages the OSQP quadratic program solver [29]. The reason is that the analytical solution does not allow including input saturation limits or stacking of additional (CBF) constraints. This allows straightforward extension to a more elaborate safety filter in the future. The solver reaches sufficient convergence to the analytical solution well within the 1 ms sample time.

B. Experiment 1: Step response

In this experiment, the equilibrium setpoint of the nominal impedance controller is moved by 40 cm in the (horizontal) y-direction by a square wave signal. As a result, it will attempt to inject a significant amount of virtual potential energy into the physical robot. When the safety filter is active, the kinetic energy limit $K_{\max} = 1$ J. We repeat the experiment for $\gamma \in \{1, 2, 10, 50\}$, and with the CBF disabled.

[2]Notice that (24) and (25) do not depend on $\boldsymbol{u}_{\text{nom}}$, while (18) does.

[3]Note that gravity and friction compensation are still present as part of the lower-level Franka Control Interface (FCI).

The end-effector trajectories (y-position) are shown in Fig. 3, with the corresponding kinetic energy in Fig. 4. The latter shows that the CBF effectively limits the kinetic energy, becoming more conservative with lower values of γ. In contrast, for the case without safety filter, the kinetic energy reaches up to 2.3 J.

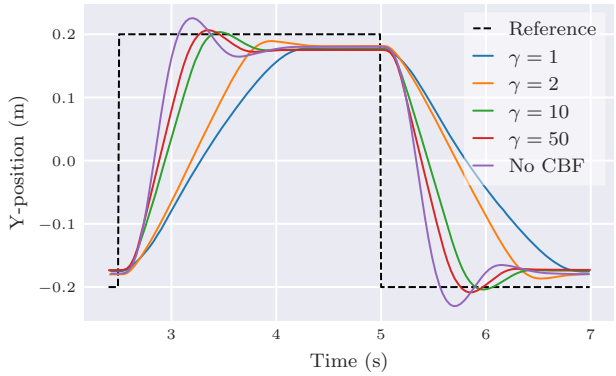

Fig. 3: Experiment 1 (Step response): End-effector y-position.

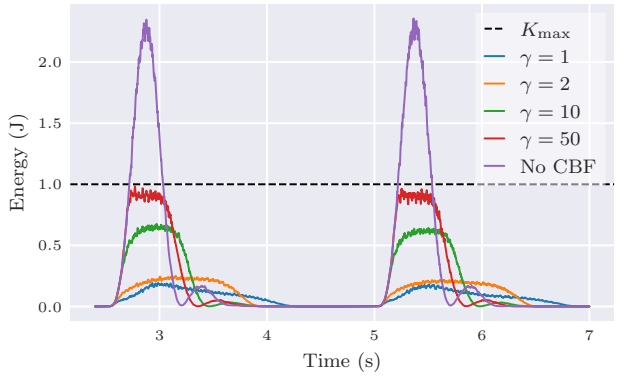

Fig. 4: Experiment 1 (Step response): Total kinetic energy.

Fig. 5 shows the power input of the safety filter, given by $\dot{\boldsymbol{q}}^\top(\boldsymbol{u}_{\text{nom}} - \boldsymbol{u})$. For all experiments, total safety filter power can be observed to be non-positive, demonstrating that the safety filter only applies damping to the system as predicted by Theorem 2. Notice that for an individual joint the injected power may indeed be positive, however the total power input is always non-positive. Fig. 6 shows the associated joint control torques for the experiment with $\gamma = 50$. The commanded input $\boldsymbol{u}$ is identical to the desired input $\boldsymbol{u}_{\text{nom}}$ until intervention is necessary, demonstrating that the safety filter is minimally invasive.

C. Experiment 2: Contact loss

In this experiment, a string is attached to the end-effector, fixed to the base on the other end, and then brought under 50 N of tension by lifting the equilibrium setpoint of the Cartesian impedance controller up by 25 cm, resulting in approximately 6.25 J of stored energy in the virtual spring. The fixed end of the string is then suddenly released, allowing the robot to accelerate upwards. This experiment is again repeated for $\gamma \in \{1, 2, 10, 50\}$ and without the CBF.

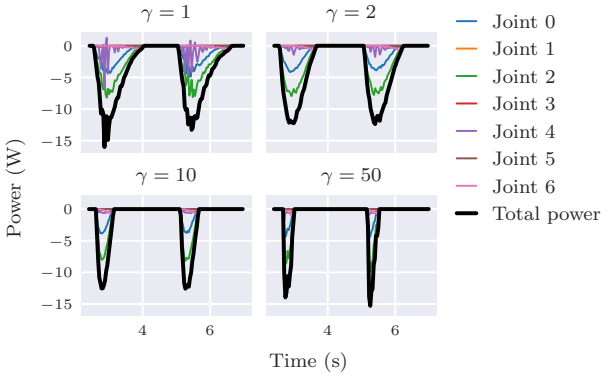

Fig. 5: Experiment 1 (Step response): Safety filter power injection.

The resulting end-effector motion is shown in Fig. 9, as snapshots of the experiment without and with the safety filter using $\gamma = 50$. Upon release of the string, the stored control energy is released and the end-effector rapidly moves up towards the equilibrium. This is similar to the robot slipping off a surface it is pushing against in a sudden loss-of-contact scenario. The total kinetic energy is shown in Fig. 7 and the CBF-induced power injection is shown in Fig. 8, both for all values of γ respectively. The former shows that approx. 1.7 J of kinetic energy is injected by the nominal controller without safety filter, and that the excess is effectively dissipated when the safety filter is activated. For $\gamma = 50$ the energy limit is momentarily exceeded, which we attribute to limited torque tracking capability of the robot's actuators. However, this breach is small, and cases with more conservative values of γ remain far from the boundary, suggesting that $\gamma = 50$ might be slightly too high for the capabilities of this system. Fig. 8 shows the total safety filter power input, which is negative for all experiments as expected due to its damping nature.

D. Experiment 3: External interaction

In the external interaction experiment, we disable the nominal controller (i.e., $\boldsymbol{u}_{\text{nom}} = \boldsymbol{0}$) and subject the robot to an unmodelled external power input, by physically pushing the end-effector by hand. We compare three different cases:

1) Without safety filter;
2) With interaction-agnostic safety filter (Eq. (12));
3) With interaction-aware safety filter (Eq. (20)).

In all cases the kinetic energy limit is set to $K_{\max} = 0.3$ J and $\gamma = 50$ for the safety filter. The external interaction torque $\boldsymbol{\tau}_{\text{ext}}$ is estimated internally and provided by the Franka API.

The kinetic energy for all three cases is shown in Fig. 10. Fig. 11 shows the relevant corresponding power flows: 1) power injected by the operator, 2) power injected by the safety filter, and 3) their sum, which is the net power input into the system. Although the three experiments are not identical, the maximum operator power input is of comparable magnitude and duration. The red zones in Fig. 11 indicate when the kinetic energy limit is exceeded (as per Fig. 10).

Considering first Fig. 10, we observe that both the interaction-agnostic and interaction-aware safety filters de-

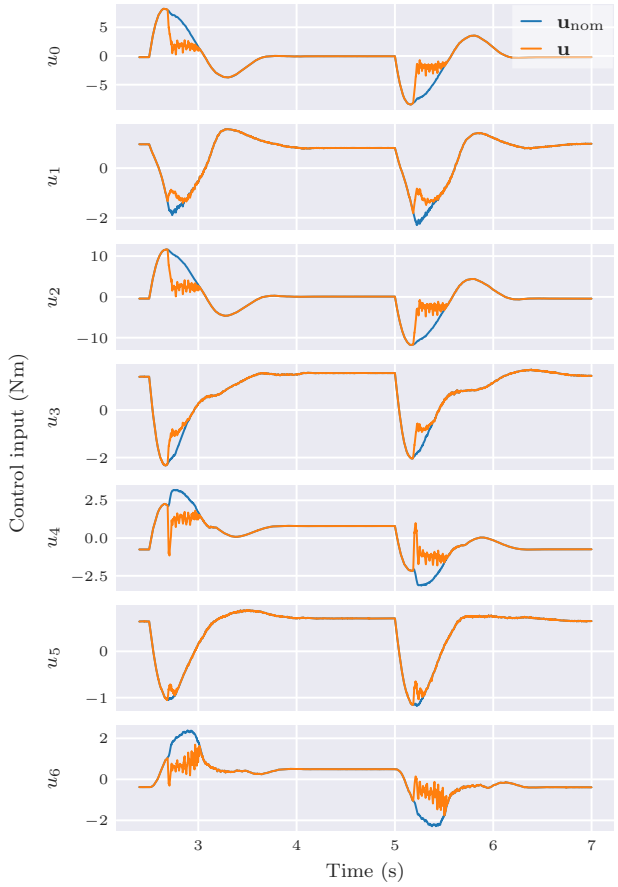

Fig. 6: Experiment 1 (Step response): Nominal desired control action $\boldsymbol{u}_{\text{nom}}$ compared to the filtered control action $\boldsymbol{u}$ ($\gamma = 50$).

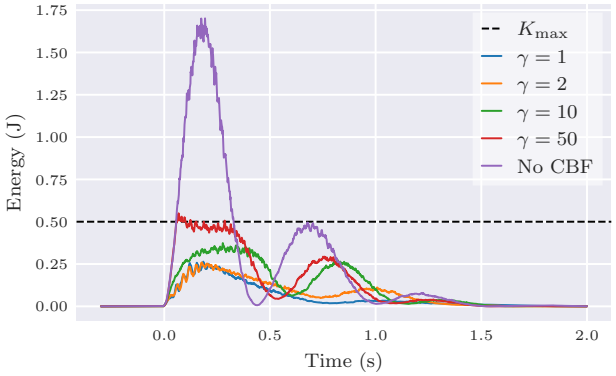

Fig. 7: Experiment 2 (Contact loss): Total kinetic energy.

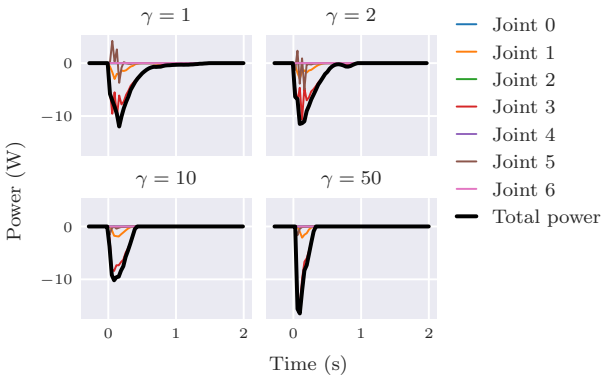

Fig. 8: Experiment 2 (Contact loss): Safety filter power.

crease the kinetic energy error compared to the experiment without safety filter. Critically, Fig. 10 shows that incorporating the estimate of the external power input reduces the error to near the kinetic energy limit, even if it is still momentarily exceeded. We attribute the latter to the relatively poor quality of the external torque estimate $\hat{\boldsymbol{\tau}}_{\text{ext}}$.

Now considering Fig. 11, we observe that the safety filter produces a larger negative power when provided with the estimate of external power input. This is consistent with the reduced kinetic energy error observed in Fig. 10.

E. Experiment 4: Kinetic energy error validation

With this final experiment we aim to validate the kinetic energy error ($K_{\text{e}} - K_{\max}$) as predicted by Theorem 3. This requires precise knowledge of the unmodelled power input, beyond what can be achieved through the robot's external torque estimation $\hat{\boldsymbol{\tau}}_{\text{ext}}$ which can be of poor quality especially during dynamic motions. Hence, we achieve this by adding an additional term $\boldsymbol{u}_{\text{err}}$ to the torque inputs generated by the CBF-based safety filter, as shown in Fig. 12.

The unmodelled power is injected by applying a virtual force horizontally in the y-direction at the end-effector. After an initial push to initiate motion, this force is regulated with velocity to provide a constant power input P_{ext}:

$$\boldsymbol{f}_{\text{ee}} = \begin{pmatrix} 0 & P_{\text{ext}}/\dot{x}_{\text{ee,y}} & 0 \end{pmatrix}, \tag{26}$$

$$\boldsymbol{u}_{\text{err}} = \boldsymbol{J}^{\top}(\boldsymbol{q})\boldsymbol{f}_{\text{ee}} \tag{27}$$

where $\dot{x}_{\text{ee,y}}$ denotes the y-direction component of the end-effector velocity $\dot{\boldsymbol{x}}_{\text{ee}} = \boldsymbol{J}(\boldsymbol{q})\dot{\boldsymbol{q}}$, and $\mathbf{f}_{\text{ee}}$ denotes the applied virtual force. The kinetic energy is then measured at steady-state, which occurs when the external input power and CBF-induced (damping) power are at equilibrium. We set $K_{\max} = 0\,\text{J}$, and perform the experiment for various power input values and $\gamma \in \{5, 10, 20, 30, 40, 50\}$.

Fig. 13 shows the resulting steady-state kinetic energy (error), as function of the input power and for different values of γ. The coloured circles indicate the data, and the dashed lines show a linear least-squares fit per value of γ. The magnitude of the kinetic energy (and, as $K_{\max} = 0\,\text{J}$, its error) is lower for higher values of γ. In addition, the linear fits closely match the the linear relation predicted by Eq. (21), as the slope of each curve is approx. γ^{-1}. This confirms the prediction of Theorem 3, that although a lower value of γ produces more conservative behaviour regarding forward invariance of the safe set, rejection of unmodelled (external) disturbances is indeed reduced.

IV. Discussion

The extensive experimental results presented in the previous Section demonstrate that a CBF-based safety filter is

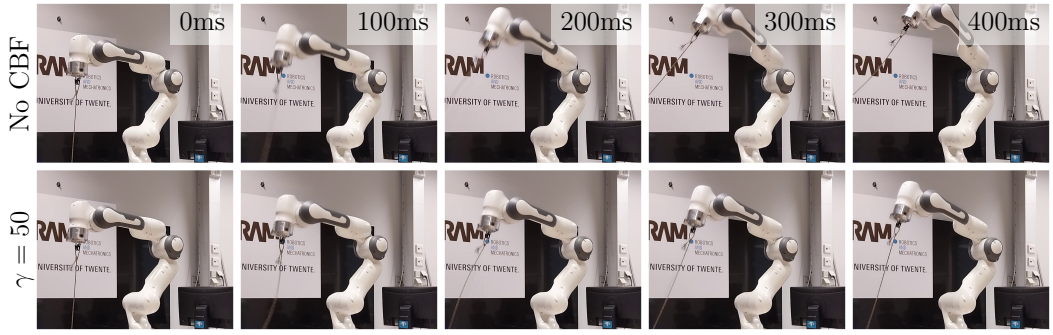

Fig. 9: Experiment 2 (Contact loss): Snapshots of the experiment. Top: No safety filter. Bottom: Safety filter with $\gamma = 50$.

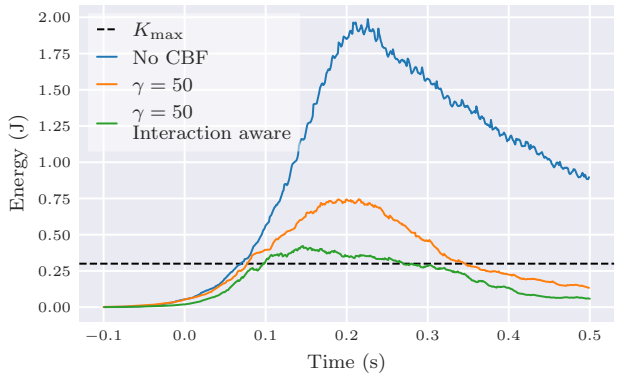

Fig. 10: Experiment 3 (External interaction): Total kinetic energy.

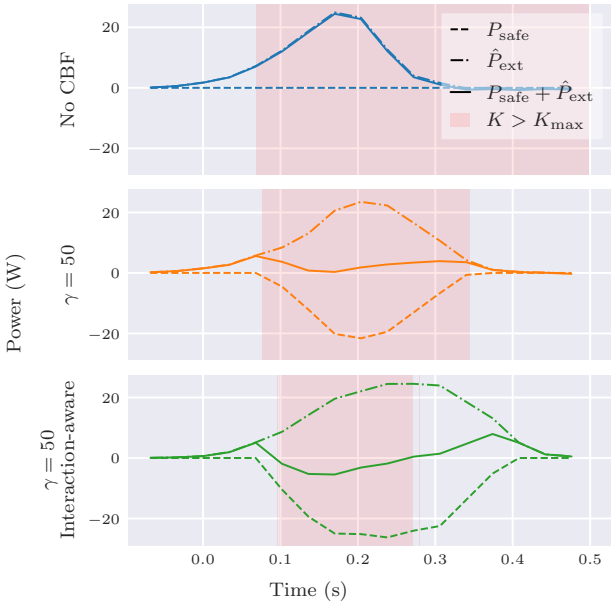

Fig. 11: Experiment 3 (External interaction): Internal and external power injection for $\gamma = 50$. The red zone indicates where the robot exceeds its kinetic energy limit.

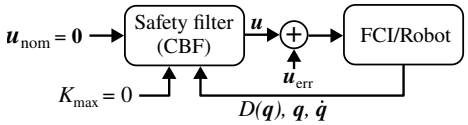

Fig. 12: Experiment 4: Unmodelled power input through u_{err}.

an effective approach to limiting the kinetic energy of a robot. We found in our experiments that it is easy to tune, requiring only a single parameter γ and no knowledge of the nominal controller, and that it provides robust performance in a variety of situations. The parameter γ can be conveniently interpreted as a measure of conservatism, with lower values being increasingly conservative. This was clearly demonstrated by Experiments 1 and 2, in which reducing γ kept the robot further away from the kinetic energy limit. In the case of a torque-controlled manipulator such as here, its value is practically upper bounded by the capability of the robot to achieve rapid changes in desired joint torques as the robot reaches the kinetic energy limit, as observed for the case with $\gamma = 50$ in Experiment 2.

Interestingly, as shown by Theorem 3 and Experiment 4, decreasing γ also reduces robustness against unmodelled (external) disturbances, in the sense of increased kinetic energy errors beyond the chosen limit. Experiment 3 demonstrated that incorporating an estimate of external interaction torques $\hat{\boldsymbol{\tau}}_{\text{ext}}$ into the computation of the CBF (which we called *interaction-aware*) can reduce or remove such errors, however we reiterate that it can be difficult to obtain accurate estimates of such external disturbances.

The choice of the kinetic energy limit $K_{\max}$ itself is almost entirely task- and situation-dependent, and outside the scope of this work. Considering safety, we envision that its value would be determined by external systems, e.g. planning and/or vision systems that assess the level of danger in a given situation, such as human proximity.

V. Conclusions and Future Work

We have presented, analysed, and experimentally validated a CBF-based approach to limit the kinetic energy of torque-controlled robots. Its energetic and disturbance rejection properties were thoroughly analysed. Taking the form of a *safety filter*, it requires zero knowledge of the nominal controller, which enables its use with black box (e.g., learning) controllers, providing them with strong guarantees on closed-loop energetic behaviour. Furthermore, the approach is minimally invasive, that is, the behaviour of the nominal

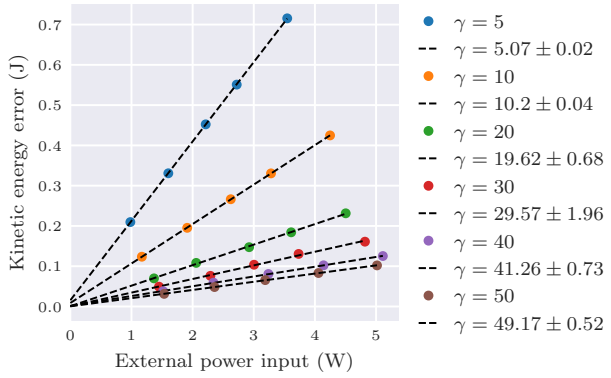

Fig. 13: Experiment 4: Kinetic energy (error) versus external unmodelled power input, with linear least-squares fits.

controller is unaltered until intervention is necessary to keep the system within the safe set. These properties make such a safety filter attractive and straightforward to implement.

We are working towards extending the proposed schemes to consider input constraints and limit not only the manipulator's total kinetic energy but also the kinetic energy transferable in specific task-space directions. In this way it will be possible to address protocols for safety hazards by restricting energy transfer in directions where human operators are present, reducing conservatism.

References

[1] S. Robla-Gomez *et al.*, "Working Together: A Review on Safe Human-Robot Collaboration in Industrial Environments," *IEEE Access*, vol. 5, pp. 26 754–26 773, 2017.

[2] M. Hamad *et al.*, "A Concise Overview of Safety Aspects in Human-Robot Interaction," no. 101017274, pp. 1–15, 2023.

[3] L. Brunke *et al.*, "Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 411–444, May 2022.

[4] International Organization for Standardization, *Robots and robotic devices - Collaborative robots (ISO Standard No. 15066:2016)*, Feb. 2016.

[5] T. Malm *et al.*, "Dynamic safety system for collaboration of operators and industrial robots," *Open Engineering*, vol. 9, no. 1, pp. 61–71, Mar. 2019.

[6] N. Mansfeld *et al.*, "Safety map: A unified representation for biomechanics impact data and robot instantaneous dynamic properties," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1880–1887, Jul. 2018.

[7] D. Tsetserukou *et al.*, "Towards Safe Human-Robot Interaction: Joint Impedance Control of a New Teleoperated Robot Arm," in *RO-MAN 2007 - The 16th IEEE International Symposium on Robot and Human Interactive Communication*, IEEE, 2007, pp. 860–865.

[8] W. Roozing, S. S. Groothuis, and S. Stramigioli, "Energy-based safety in series elastic actuation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, Paris, France: IEEE, May 2020, pp. 914–920.

[9] S. Haddadin and E. Croft, "Physical Human–Robot Interaction," in *Springer Handbook of Robotics*, 2016, pp. 1839–1840.

[10] T. S. Tadele, T. de Vries, and S. Stramigioli, "The Safety of Domestic Robotics: A Survey of Various Safety-Related Publications," *IEEE Robotics & Automation Magazine*, vol. 21, no. 3, pp. 134–142, Sep. 2014.

[11] A. D. Ames *et al.*, "Control barrier functions: Theory and applications," *2019 18th European Control Conference, ECC 2019*, pp. 3420–3431, 2019.

[12] W. S. Cortez and D. V. Dimarogonas, "Correct-by-Design Control Barrier Functions for Euler-Lagrange Systems with Input Constraints," in *2020 American Control Conference (ACC)*, IEEE, Jul. 2020, pp. 950–955.

[13] C. T. Landi *et al.*, "Safety barrier functions for human-robot interaction with industrial manipulators," in *2019 18th European Control Conference, ECC 2019*, 2019.

[14] M. Rauscher, M. Kimmel, and S. Hirche, "Constrained robot control using control barrier functions," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, Oct. 2016, pp. 279–285.

[15] F. Ferraguti *et al.*, "Safety and Efficiency in Robotics: The Control Barrier Functions Approach," in 2022, pp. 15–30.

[16] A. Singletary, S. Kolathaya, and A. D. Ames, "Safety-Critical Kinematic Control of Robotic Systems," *Proceedings of the American Control Conference*, vol. 2021-May, pp. 14–19, 2021.

[17] W. S. Cortez, C. K. Verginis, and D. V. Dimarogonas, "Safe, passive control for mechanical systems with application to physical human-robot interactions," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3836–3842.

[18] B. Capelli, C. Secchi, and L. Sabattini, "Passivity and Control Barrier Functions: Optimizing the Use of Energy," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1356–1363, Apr. 2022.

[19] R. Ortega *et al.*, "Putting energy back in control," *IEEE Control Systems Magazine*, vol. 21, no. 2, pp. 18–33, 2001.

[20] S. Stramigioli, "Energy-aware robotics," in *Mathematical Control Theory I*, M. K. Camlibel *et al.*, Eds., Cham: Springer International Publishing, 2015, pp. 37–50.

[21] Y. Michel, M. Saveriano, and D. Lee, "A Novel Safety-Aware Energy Tank Formulation Based on Control Barrier Functions," *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5206–5213, Jun. 2024.

[22] F. Califano *et al.*, *On the Use of Energy Tanks for Robotic Systems*, P. Borja *et al.*, Eds. Cham: Springer International Publishing, 2023, pp. 174–188.

[23] F. Benzi, F. Ferraguti, and C. Secchi, "Energy tank-based control framework for satisfying the iso/ts 15066 constraint," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 1288–1293, 2023, 22nd IFAC World Congress.

[24] G. Raiola *et al.*, "Development of a Safety- and Energy-Aware Impedance Controller for Collaborative Robots," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 1237–1244, Apr. 2018.

[25] W. Shaw Cortez, X. Tan, and D. V. Dimarogonas, "A robust, multiple control barrier function framework for input constrained systems," *IEEE Control Systems Letters*, vol. 6, pp. 1742–1747, 2022.

[26] A. D. Ames *et al.*, "Control Barrier Function Based Quadratic Programs for Safety Critical Systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[27] X. Xu *et al.*, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015, Analysis and Design of Hybrid Systems ADHS.

[28] F. Califano, "Passivity-Preserving Safety-Critical Control Using Control Barrier Functions," *IEEE Control Systems Letters*, vol. 7, pp. 1742–1747, 2023.

[29] B. Stellato *et al.*, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, Dec. 2020.