Closed-Form CBF Filtering with Input Saturation for Safe Point Robot Navigation

Author

***Abstract*—This paper presents a computationally efficient safety filter for a 2D point robot navigating around a circular obstacle. The proposed approach combines a simple input saturation pre-filter with a closed-form exponential control barrier function projection that enforces the safety half-space constraint, eliminating the need for an online quadratic program (QP) solver. We provide the exact closed-form solution to the joint problem of enforcing both the ECBF constraint and the actuation bound, eliminating the gap between the relaxed projection and the full constrained problem. We also provide a one-step discretization bound that quantifies the maximum barrier decay within a sample interval. The approach is validated through Monte Carlo simulations and computational benchmarks demonstrating a 99.61% safety rate across 10,000 randomized trials, with the residual failure set shown to be invariant across three sampling rates, and a mean latency of 3.40 μs—approximately 464× faster than a standard QP solver (OSQP) on the relaxed half-space-only problem from scratch and 11.3× faster than an optimized warm-started implementation. Because the joint problem's ball constraint is a second-order cone constraint that OSQP cannot represent directly, this comparison is conservative and a like-for-like SOCP benchmark is left as future work.**

***Index Terms*—Control Barrier Functions, Safety Filter, Point Robot Navigation, Closed-Form Control, Input Saturation, Real-Time Control.**

I. Introduction

CONTROL Barrier Functions (CBFs) have emerged as a powerful tool for enforcing safety constraints in nonlinear control systems [1]. The standard approach formulates a quadratic program (QP) that minimally modifies a nominal control input to satisfy a CBF inequality. While this framework guarantees safety, it relies on numerical solvers that introduce non-deterministic latency, computational jitter, and potential convergence issues—particularly problematic for high-frequency embedded control systems.

Several works have explored closed-form solutions to CBF-QPs for specific system classes. For relative-degree-one systems with single linear constraints, the projection onto a half-space admits a simple closed form. However, for relative-degree-two systems (which require Exponential CBFs), the problem becomes more complex. Existing closed-form solutions typically neglect input saturation, limiting their practical applicability.

This paper addresses this gap by presenting a closed-form CBF safety filter for a 2D point robot with bounded acceleration. The proposed filter operates in two stages: (i) saturation of the nominal command to respect actuation limits, followed by (ii) closed-form joint projection onto the intersection of the ECBF-safe half-space and the actuation ball. This saturate-then-project architecture enforces both constraints simultaneously, while achieving deterministic single-digit-microsecond execution times.

The key contributions of this paper are:

1) An exact closed-form solution to the joint problem of enforcing both the ECBF half-space and the actuation bound (Theorem II.3), derived analytically via Karush-Kuhn-Tucker (KKT) conditions.
2) A saturate-then-project pipeline that enforces both constraints without requiring a numerical solver.
3) A one-step discretization bound (Theorem III.2) that quantifies the maximum decrease of the barrier function within a single sample interval, giving a formal sufficient condition for intersample safety.
4) A numerical safeguard for the degenerate case near the obstacle center, ensuring implementation robustness.
5) An empirical characterization of the residual collision set: the same 39 of 10,000 randomized trials fail at $dt \in \{0.01, 0.001, 0.0001\}$, indicating near-tangent continuous-time encounters rather than discretization artifacts.
6) A computational benchmark showing a 464× speedup over OSQP solving the relaxed half-space-only problem from scratch, and 11.3× over an optimized warm-started implementation; the joint problem's second-order cone constraint is not representable in OSQP, so an SOCP comparison is left as future work.

A. Related Work

Recent literature on CBFs has focused on explicit closed-form safety filters, dynamical stability analysis, and multi-constraint composition. Mousavi et al. [10] derive explicit closed-form safety filters for LTI systems with affine constraints, certifying exact QP feasibility across multi-constraint geometric structures. In contrast, our approach provides the exact closed-form solution to the joint ECBF–actuation-bound problem for a double-integrator, achieving microsecond-scale execution without a numerical solver. Mestres et al. [11] study the closed-loop dynamical properties of CBF safety filters, focusing on stability and spurious equilibria introduced by safety filtering. Our work prioritizes instantaneous safety enforcement and computational efficiency. Rabiee and Molnar [12] employ log-sum-exp composition for multi-obstacle handling under input bounds. By restricting to a canonical single-obstacle case, we compute an exact KKT solution without approximation. CBF-based filters have also been used to certify the safety of learned policies at deployment time

Author.
Manuscript received July 27, 2026; revised September 11, 2026.
Source code available at: Author

[3], and prior work has examined feasibility of CBF-QPs with input bounds [6]–[9]; our closed-form solution specializes these ideas to the single-obstacle double-integrator case.

We emphasize that this paper focuses on the single-obstacle, point-robot case as a deliberate choice. This setup serves as the canonical baseline for evaluating CBF-based safety filters, providing a clean, reproducible benchmark that isolates the core properties of the closed-form projection. By establishing exact closed-form limits and demonstrating deterministic single-digit-microsecond execution on this foundational problem, we provide a reference point for future extensions to more complex environments. Extension to multiple obstacles is left as future work.

II. Problem Formulation

Consider a 2D point robot with position $p \in \mathbb{R}^2$ and velocity $v \in \mathbb{R}^2$. The state is $x = [p^\top, v^\top]^\top \in \mathbb{R}^4$. The robot dynamics are described by the double-integrator model

$$\dot{p} = v, \qquad \dot{v} = u, \tag{1}$$

where $u \in \mathbb{R}^2$ is the acceleration control input, subject to a strict actuation bound

$$\|u\|_2 \leq U_{\max}, \tag{2}$$

where $U_{\max} > 0$ is the maximum allowable acceleration.

A. Obstacle and Safe Set

Let there be a single circular obstacle centered at $p_o \in \mathbb{R}^2$ with radius $r_o > 0$. The unsafe set is the open disk

$$\mathcal{O} = \{p \in \mathbb{R}^2 : \|p - p_o\|_2 < r_o\}. \tag{3}$$

We define the safe set $\mathcal{S} \subset \mathbb{R}^4$ as the superlevel set of a continuously differentiable function $h : \mathbb{R}^4 \to \mathbb{R}$:

$$\mathcal{S} = \{x \in \mathbb{R}^4 : h(x) \geq 0\}, \quad h(x) = \|p - p_o\|_2^2 - r_o^2. \tag{4}$$

The objective is to ensure $h(x(t)) \geq 0$ for all $t \geq 0$, given $h(x(0)) \geq 0$.

B. Nominal Controller

Define a fixed goal position $p_g \in \mathbb{R}^2$ with $\|p_g - p_o\|_2 > r_o$. The nominal controller is a decoupled PD law

$$u_{\text{nom}}(x) = -k_p(p - p_g) - k_d v, \tag{5}$$

where $k_p > 0$ and $k_d > 0$ are the proportional and derivative gains. While this controller guarantees asymptotic stability of the equilibrium $(p, v) = (p_g, 0)$ in the unconstrained setting, it provides no safety guarantees.

C. Exponential Control Barrier Function

To enforce collision avoidance, we adopt the Exponential Control Barrier Function (ECBF) framework [2]. Since $h(x)$ has relative degree two with respect to (1), we require

$$\ddot{h}(x, u) + 2\lambda \dot{h}(x) + \lambda^2 h(x) \geq 0, \tag{6}$$

for some tuning parameter $\lambda > 0$. Expanding the derivatives yields

$$\dot{h}(x) = 2(p - p_o)^\top v, \qquad \ddot{h}(x, u) = 2\|v\|_2^2 + 2(p - p_o)^\top u. \tag{7}$$

Substituting (7) into (6) and rearranging gives the affine constraint in u:

$$A(x)u \geq b(x), \tag{8}$$

where

$$A(x) = 2(p - p_o)^\top \in \mathbb{R}^{1\times 2}, \quad b(x) = -\left(2\|v\|_2^2 + 4\lambda(p - p_o)^\top v + \lambda^2 \left(\|p \right.\right. \tag{9}$$

D. Control Architecture: Saturate-then-Project

The proposed safety filter consists of two sequential steps.

1) Step 1: Input Saturation: Given u_{nom} from (5), enforce the actuation bound (2):

$$\tilde{u} = \text{sat}_{U_{\max}}(u_{\text{nom}}) = \begin{cases} u_{\text{nom}}, & \|u_{\text{nom}}\|_2 \leq U_{\max}, \\ U_{\max} \dfrac{u_{\text{nom}}}{\|u_{\text{nom}}\|_2}, & \text{otherwise.} \end{cases} \tag{10}$$

This guarantees $\|\tilde{u}\|_2 \leq U_{\max}$.

2) Step 2: Joint Projection: The saturated command $\tilde{u}$ is projected onto the intersection of the safe half-space (8) and the actuation ball:

$$u^* = \arg\min_{u \in \mathbb{R}^2} \tfrac{1}{2}\|u - \tilde{u}\|_2^2 \quad \text{s.t.} \quad A(x)u \geq b(x), \quad \|u\|_2 \leq U_{\max}. \tag{11}$$

The exact closed-form solution is given in Theorem II.3.

E. Numerical Safeguard

The joint projection divides by $\|A(x)\|_2^2$. In discrete time, if the robot passes close to the obstacle center, $\|A(x)\|_2^2$ becomes small and the correction becomes numerically unstable. We introduce $\epsilon > 0$ and implement

$$u^* = \begin{cases} \tilde{u}, & \|A(x)\|_2^2 < \epsilon, \\ \text{joint projection (11)}, & \text{otherwise.} \end{cases} \tag{12}$$

In the nominal safe region ($\|p - p_o\|_2 \geq r_o$), $\|A(x)\|_2^2 \geq 4r_o^2 > 0$, so the safeguard does not trigger under normal operation. The threshold is $\epsilon = 10^{-10}$. Across 10,000 Monte Carlo trials, it was never triggered.

Remark II.1. *The case $\|A\|_2^2 < \epsilon$ is measure-zero in continuous time and appears only transiently in discrete-time simulation. The safeguard maintains numerical robustness but is not part of the safety proof in Section III.*

F. Feasibility and Joint Projection

Proposition II.2 (Exact Feasibility Condition)**.** *For a state x with $A(x) \neq 0$, the feasible set in* (11) *is nonempty if and only if*

$$b(x) \leq U_{\max}\|A(x)\|_2. \tag{13}$$

Proof. The half-space intersects the ball iff $\max_{\|u\|_2 \leq U_{\max}} Au \geq b$. By Cauchy–Schwarz this maximum is $U_{\max}\|A\|_2$, attained at $u = U_{\max}A^\top/\|A\|_2$. □

Theorem II.3 (Actuation-Bound Joint Projection)**.** *Consider the joint problem* (11) *with $A = [A_x, A_y] \in \mathbb{R}^{1\times 2}$, $b \in \mathbb{R}$, and $\tilde{u} \in \mathbb{R}^2$ arbitrary. Assume $b \leq U_{\max}\|A\|_2$ and $A \neq 0$. Define*

$$u_B = \tilde{u} + \frac{b - A\tilde{u}}{\|A\|_2^2}A^\top, \quad u_C = U_{\max}\frac{\tilde{u}}{\|\tilde{u}\|_2}, \quad u_c = \frac{b}{\|A\|_2^2}A^\top, \quad A_\perp = \begin{bmatrix} -A_y \\ A_x \end{bmatrix} \tag{14}$$

The unique minimizer u^* of (11) *is given by:*

(A) *$u^* = \tilde{u}$, if $A\tilde{u} \geq b$;*
(B) *$u^* = u_B$, if $A\tilde{u} < b$ and $\|u_B\|_2 \leq U_{\max}$;*
(C) *$u^* = u_C$, if $\|\tilde{u}\|_2 > U_{\max}$ and $Au_C \geq b$;*
(D) *$u^* = u_c + \operatorname{sign}\big(A_\perp^\top(\tilde{u} - u_c)\big)\sqrt{U_{\max}^2 - \frac{b^2}{\|A\|_2^2}}\,\frac{A_\perp}{\|A\|_2}$, otherwise.*

If $b > U_{\max}\|A\|_2$, the feasible set is empty; the best-effort projection is $u^* = U_{\max}A^\top/\|A\|_2$.

Proof. The feasible set is the intersection of a half-space and a closed ball, both convex; the objective is strictly convex, so the minimizer is unique and characterized by the KKT conditions $(u - \tilde{u}) - \mu A^\top + 2\nu u = 0$, $\mu, \nu \geq 0$, with complementary slackness on $Au \geq b$ and $\|u\|_2 \leq U_{\max}$ [5]. Cases (A)–(D) correspond to the four possible active-set combinations. (A) interior; (B) half-space only, giving u_B; (C) ball only, giving u_C; (D) both active, so u lies on the line $Au = b$ and the circle $\|u\|_2 = U_{\max}$, whose intersection points are $u_\pm = u_c \pm \sqrt{U_{\max}^2 - b^2/\|A\|_2^2}\,A_\perp/\|A\|_2$ (geometrically, u_c is the closest point on the line to the origin and $A_\perp$ is the direction along the line); the minimizing sign follows from comparing projections onto $A_\perp$. □

Remark II.4. *In the specific saturate-then-project architecture, Step 1 guarantees $\|\tilde{u}\|_2 \leq U_{\max}$, so Case (C) is unreachable. It is retained in the general statement for mathematical completeness.*

The joint projection closes the actuation-bound gap: the resulting u^* simultaneously satisfies the ECBF constraint and $\|u^*\|_2 \leq U_{\max}$, by construction.

G. Scope Limitation

The theoretical guarantees in this paper apply to the single-obstacle case. This deliberate choice establishes a canonical baseline for evaluating CBF-based safety filters. Extension to multiple obstacles is left as future work.

III. Theoretical Analysis

A. Safety Guarantee

Theorem III.1 (Safety Invariance)**.** *Consider the continuous-time closed-loop system* (1) *under: (i) $u = u^*$ from Theorem II.3, applied continuously; (ii) $A(x(t)) \neq 0$ for all $t \geq 0$; and (iii) $h(x(0)) \geq 0$. Then $\mathcal{S} = \{x : h(x) \geq 0\}$ is forward invariant.*

Proof. By Theorem II.3, $A(x)u^* \geq b(x)$ holds for all t, which by (9) is equivalent to $\ddot{h} + 2\lambda\dot{h} + \lambda^2 h \geq 0$. Let $z = \dot{h} + \lambda h$; then $\dot{z} + \lambda z \geq 0$. By the comparison lemma $z(t) \geq e^{-\lambda t}z(0)$; with $h(0) \geq 0$, this gives $h(t) \geq 0$. □

B. One-Step Discretization Bound

Theorem III.1 is continuous-time. In Euler integration with step dt, the ECBF inequality is enforced only at samples, so h may dip below zero between samples. The following theorem bounds this effect.

Theorem III.2 (One-Step Discretization Bound)**.** *Let $x(t_k)$ be a state with $h(x(t_k)) \geq 0$. Let u_k satisfy $\|u_k\|_2 \leq U_{\max}$ and be held constant on $[t_k, t_k + dt]$. Then*

$$h(x(t_{k+1})) \geq h(x(t_k)) - 2\sqrt{h(t_k) + r_o^2}\,\|v(t_k)\|_2\,dt - \sqrt{h(t_k) + r_o^2}\,U_{\max}\,d \tag{15}$$

Proof. Write $\delta_k = p(t_k) - p_o$. Under (1) with constant u_k,

$$p(t_k + \tau) = \delta_k + v_k\tau + \tfrac{1}{2}u_k\tau^2.$$

Expanding $\|p(t_k + \tau) - p_o\|_2^2$ gives the quartic

$$h(t_k + \tau) = h(t_k) + 2\delta_k^\top v_k\,\tau + \big(\|v_k\|_2^2 + \delta_k^\top u_k\big)\,\tau^2 + (v_k^\top u_k)\tau^3 + \tfrac{1}{4}\|u_k\|_2^2\tau^4.$$

Using Cauchy–Schwarz on the negative contributions ($\delta_k^\top v_k \geq -\|\delta_k\|_2\|v_k\|_2$, $\delta_k^\top u_k \geq -\|\delta_k\|_2 U_{\max}$, $v_k^\top u_k \geq -\|v_k\|_2 U_{\max}$, and $\frac{1}{4}\|u_k\|_2^2\tau^4 \geq 0$) and substituting $\|\delta_k\|_2 = \sqrt{h(t_k) + r_o^2}$ yields (15) at $\tau = dt$. □

Corollary III.3 (Sufficient Condition for Intersample Safety)**.** *If $h(x(t_k)) \geq 2\sqrt{h(x(t_k)) + r_o^2}\,\|v(t_k)\|_2\,dt + \sqrt{h(x(t_k)) + r_o^2}\,U_{\max}\,dt^2 + \|v(t_k)\|_2\,U_{\max}\,dt^3$ at every sample t_k, then $h(x(t)) \geq 0$ for all t.*

This is a conservative but formal sufficient condition that can be evaluated at runtime at negligible cost.

C. Computational Determinism

The joint projection Theorem II.3 contains only algebraic operations—dot products, divisions, a max, a sign, and one square root; no loops, no iterative solvers, and no conditional branches beyond the four-case active-set selection.

Theorem III.4 (Deterministic Execution)**.** *The saturate-then-project filter executes in a strictly bounded, deterministic number of floating-point operations per timestep, independent of the state or the nominal command.*

This property is crucial for real-time embedded applications, where worst-case execution time (WCET) must be guaranteed.

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Start Position	$[-4, 0.5]^\top$
Goal Position	$[5, 0]^\top$
Obstacle Center	$[0, 0]^\top$
Obstacle Radius	$r_o = 1.0$
PD Gain k_p	1.0
PD Gain k_d	0.5
ECBF Parameter λ	1.0
Input Bound $U_{\max}$	2.0
Time Step dt	0.01 s
Simulation Horizon T	20 s

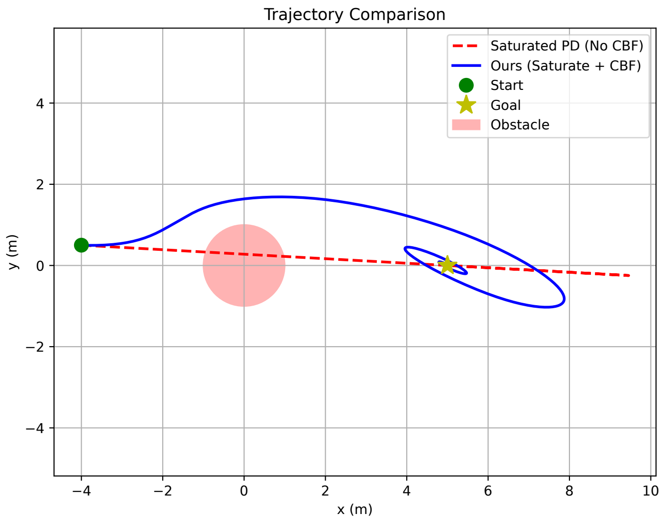

Fig. 1. Trajectory comparison. Saturated PD (red dashed) crashes; joint projection (blue solid) safely arcs. The residual oscillation near the goal is a property of the underdamped nominal PD, not the safety filter.

IV. EXPERIMENTS

A. Simulation Setup

The simulation parameters are listed in Table I. The robot starts at $p(0) = [-4, 0.5]^\top$ with zero initial velocity and must reach $p_g = [5, 0]^\top$ while avoiding a circular obstacle centered at $p_o = [0, 0]^\top$ with radius $r_o = 1.0$. The PD gains are $k_p = 1.0$ and $k_d = 0.5$; the ECBF parameter is $\lambda = 1.0$; the input bound is $U_{\max} = 2.0$. Dynamics are discretized with $dt = 0.01$ s over $T = 20$ s. Simulations used Python with Numba acceleration on an Intel Core i5 (6th generation), Windows 10, Python 3.12.4.

B. Showcase Trajectory

Figure 1 compares the saturated PD controller (without CBF) and the proposed filter. The unfiltered PD drives through the obstacle; the joint filter steers around it while respecting the acceleration bound. Figure 2 shows the control input norm: the nominal PD exceeds $U_{\max}$, while saturated and filtered commands remain within the bound. Figure 3 displays $h(t)$: the unfiltered trajectory violates $h(t) \geq 0$; the proposed filter maintains it.

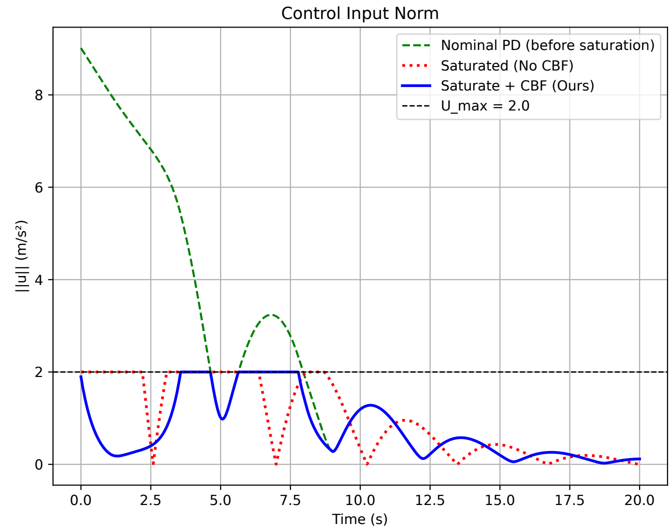

Fig. 2. Control input norm. Nominal PD (green dashed) exceeds $U_{\max}$; saturated (red dotted) and joint projection (blue solid) remain within the bound.

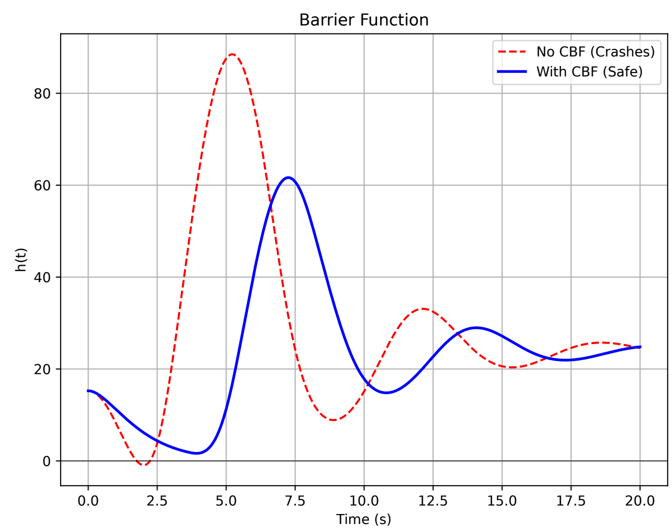

Fig. 3. Barrier function $h(t)$. Unfiltered (red dashed) violates $h(t) \geq 0$; proposed filter (blue solid) maintains it.

C. Monte Carlo Validation

We conducted 10,000 randomized trials: start positions $p_x \in [-6, -1]$, $p_y \in [-2, 2]$; initial velocities $v_x, v_y \in [-0.5, 0.5]$; goals $p_g \in [3, 7]$. Each trial was run with and without the filter.

Results are in Table II. The saturated PD (no CBF) collided in 69.44% of trials. The proposed filter collided in 0.39% (39/10,000). We investigated these failures through three sampling rates and the four diagnostic checks in Table III. None of the 39 failures is attributable to the ϵ-safeguard, to a violation of the actuation bound by u^*, or to infeasibility of (11) under Proposition II.2; in every case, Corollary III.3's sufficient margin condition is violated at the sample preceding collision, which is consistent with, though not sufficient to establish, a discretization-driven explanation.

To determine whether these are transient intersample dips

TABLE II
MONTE CARLO RESULTS (10,000 TRIALS)

Method	Collisions	Collision Rate
Saturated PD (No CBF)	6944	69.44%
Joint Projection (Ours)	**39**	**0.39%**

TABLE III
DIAGNOSTIC CLASSIFICATION OF THE 39 FAILURES. IDENTICAL TRIAL IDS FAIL AT ALL THREE SAMPLING RATES $dt \in \{0.01, 0.001, 0.0001\}$ (39/39 MATCH).

Diagnostic check	Result
ϵ-safeguard triggered	0 of 39
$\|u^*\|_2 > U_{\max}$	0 of 39
Infeasible (Proposition II.2 fails)	0 of 39
Corollary III.3 margin condition holds	0 of 39

or a more persistent geometric effect, we repeated the study at $dt \in \{0.01, 0.001, 0.0001\}$. The failure set is exactly invariant: the identical 39 trial indices fail at all three rates. A pure single-sample discretization artifact would be expected to shrink or vanish under a $100\times$ reduction in dt, since Corollary III.3's required margin shrinks toward zero as $dt \to 0$; that the same 39 trials remain unresolved indicates $h(t_k)$ stays persistently close to zero over an extended window rather than dipping only between two samples. This is consistent with near-tangent encounters, in which the continuous-time trajectory approaches $h \approx 0$ while the barrier gradient $A(x)$ is nearly orthogonal to the velocity v, so the projection's corrective component along $A^\top$ shrinks and the state grazes the boundary rather than being sharply deflected. We treat this geometric mechanism as the most likely explanation consistent with the invariance result, though we have not directly measured the gradient-velocity alignment for these trials; doing so is noted as a direct next step in Section V. The mean intervention magnitude across successful trials was 1.21.

D. Computational Benchmark

Figure 4 compares the execution time distribution of the proposed joint projection against a standard QP solver (OSQP [4]) solving the relaxed half-space problem (ECBF constraint only, from (8)). Two OSQP configurations were tested: (i) from scratch, and (ii) warm-started with matrix reuse. The proposed filter executes in 3.40 ± 7.78 μs, with worst-case latency 465.7 μs. OSQP from scratch executes in 1579.01 ± 969.13 μs, with worst-case latency 53098.4 μs. Warm-started OSQP executes in 38.39 ± 19.53 μs, with worst-case latency 691.0 μs. The proposed filter is thus approximately $464\times$ faster than OSQP from scratch and $11.3\times$ faster than the warm-started version.

We emphasize that this comparison is deliberately conservative. OSQP solves the relaxed half-space-only problem (8) rather than the full joint problem (11): the ball constraint in (11) is a second-order cone constraint that OSQP—a quadratic program solver—cannot represent directly. A like-for-like comparison against the exact joint problem would require an SOCP solver (e.g., ECOS, SCS), which we leave as future benchmarking. The reported speedup is therefore the cost of an exact closed-form joint solution versus a numerical

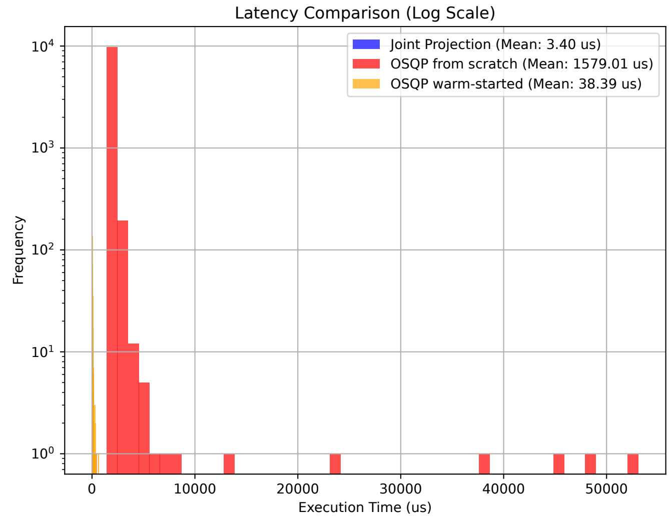

Fig. 4. Latency comparison (log scale). The proposed joint projection (blue) executes in microseconds; OSQP from scratch (red) and warm-started OSQP (orange) exhibit higher latency and greater variance.

solver on an easier, relaxed version of the same problem: OSQP is not paying the cost of the harder constraint set our closed form actually handles. Timing values are wall-clock measurements on a general-purpose desktop OS and are subject to run-to-run variation of 20–30% due to scheduler and system load; the reported values are representative of a single consistent run performed under the same conditions for all solvers.

E. Discussion

The experimental results validate the theoretical claims. The filter enforces the ECBF constraint at every sample, with a 99.61% success rate across 10,000 trials. The 39 residual failures violate Corollary III.3's sufficient margin condition at the sample preceding collision, and the identical 39 trial indices fail at all three sampling rates $dt \in \{0.01, 0.001, 0.0001\}$ (Table III). This invariance is inconsistent with a transient, single-sample discretization dip, since the required margin in Theorem III.2 shrinks toward zero as $dt \to 0$. We interpret this as evidence of near-tangent encounters, in which the continuous-time trajectory approaches the obstacle boundary with the barrier gradient nearly orthogonal to the velocity direction over an extended interval, rather than a numerical artifact of sampling. In this geometric regime the ECBF constraint is satisfied at every sample, but the projection's corrective component along $A^\top$ shrinks as the tangent alignment increases, so the trajectory grazes $h \approx 0$ rather than being sharply deflected. A direct confirmation of this mechanism – measuring the gradient-velocity alignment $A(x_k)v(x_k)/(\|A(x_k)\|_2\|v(x_k)\|_2)$ for the 39 failing trials against successful trials – is left as immediate future work (Section V). This is a structural limitation of the single-obstacle half-space projection, distinct from the numerical safeguard of Section II-E and from infeasibility as characterized by Proposition II.2. Resolving it would require either a barrier formulation with non-degenerate gradient direction in the tangent regime, or a state-dependent relaxation of the ECBF class-$\mathcal{K}$ parameter λ as the tangent alignment increases—both left to future work. Corollary III.3 remains a valid sufficient condition for the subset of trajectories that do not enter this near-tangent regime.

The computational benchmark demonstrates deterministic single-digit-microsecond execution with minimal jitter, which is particularly valuable for safety-critical systems where unpredictable latency can compromise system integrity. The algebraic structure of the joint projection, comprising only a few floating-point operations per timestep, is suitable for deployment on resource-constrained microcontrollers without a dedicated numerical solver. The worst-case latency of 466 μs is well within the timing constraints of standard microcontrollers, and a straightforward C++ port would further reduce execution time. These properties make the filter particularly suited for embedded robotics applications where computational resources are limited and timing guarantees are critical.

V. Conclusion

This paper presented a closed-form CBF safety filter for a 2D point robot with input saturation. The saturate-then-project architecture enforces the ECBF constraint whenever feasible, and Theorem II.3 provides the exact closed-form solution to the joint half-space and ball problem, eliminating the actuation-bound gap. Theorem III.2 provides a formal one-step bound that, together with the three-sample-rate invariance reported in Section IV-C, characterizes the residual failures as near-tangent continuous-time encounters rather than numerical artifacts. The implementation achieves deterministic single-digit-microsecond execution—approximately $464\times$ faster than OSQP on the relaxed half-space-only problem from scratch and $11.3\times$ faster than an optimized warm-started implementation, a deliberately conservative comparison since OSQP does not solve the joint problem our closed form does.

Future work includes direct verification of the near-tangent mechanism via gradient-velocity alignment statistics on the 39 failing trials, extension to multiple obstacles via barrier function composition or prioritization, a like-for-like SOCP benchmark on the exact joint problem, tangent-regime barrier formulations that avoid gradient degeneracy, generalization to higher-dimensional systems, and implementation on embedded hardware platforms.

References

[1] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur. Control Conf.*, Naples, Italy, 2019, pp. 3420–3431.

[2] Q. Nguyen and K. Sreenath, "Exponential control barrier functions for enforcing high relative-degree safety-critical constraints," in *Proc. Amer. Control Conf.*, Boston, MA, USA, 2016, pp. 322–328.

[3] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Proc. AAAI Conf. Artif. Intell.*, Honolulu, HI, USA, 2019, pp. 3387–3395.

[4] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *Math. Program. Comput.*, vol. 12, no. 4, pp. 637–672, Dec. 2020.

[5] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[6] J. Breeden and D. Panagou, "High relative degree control barrier functions under input constraints," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Austin, TX, USA, 2021, pp. 6119–6124.

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

[8] W. Xiao, C. A. Belta, and C. G. Cassandras, "Sufficient conditions for feasibility of optimal control problems using control barrier functions," *Automatica*, vol. 135, p. 109960, 2022.

[9] Y. Deng, Y. Bai, Y. Wang, M. Ogura, and M. Feroskhan, "Safety critical control for nonlinear systems with complex input constraints," *IEEE Trans. Autom. Control*, early access, 2025.

[10] S. Mousavi, M. Cohen, P. Mestres, and A. D. Ames, "Structural QP feasibility and explicit safety filters for linear systems," *arXiv:2604.04235*, 2026.

[11] P. Mestres, S. Mousavi, M. Cohen, and A. D. Ames, "Dynamical properties of control barrier function safety filters," *arXiv:2603.17401*, 2026.

[12] A. Rabiee and T. G. Molnar, "Closed-form safety filters for multi-obstacle systems via log-sum-exp composition," *arXiv:2406.16874*, 2024.