Anytime Solution of Constrained Nonlinear Programs via Control Barrier Functions

Ahmed Allibhoy Jorge Cortés

Abstract— **This paper considers the problem of designing a dynamical system to solve constrained nonlinear optimization problems such that the feasible set is forward invariant and asymptotically stable. The invariance of the feasible set makes the dynamics anytime, when viewed as an algorithm, meaning that it is guaranteed to return a feasible solution regardless of when it is terminated. Such property is of critical importance in feedback control since controllers are often implemented as solutions to constrained programs that must be solved in real time. The proposed design builds on the basic insight of following the gradient flow of the objective function while keeping the state evolution within the feasible set using techniques from the theory of control barrier functions. We show that the resulting closed-loop system can be interpreted as a continuous approximation of the projected gradient flow, establish the monotonic decrease of the objective function along the feasible set, and characterize the asymptotic convergence properties to the set of critical points. Various examples illustrate our results.**

I. Introduction

Optimization problems arise naturally in many engineering applications and much research effort in applied mathematics and engineering is devoted to finding efficient methods that scale well with the problem dimension. This is of particular importance in control applications since feedback controllers are often implemented as the solution to an optimization that must solved in real time. Real-time implementations create additional challenges in the design of optimization methods when the program involves constraints on the decision variables. This is because an algorithm solving the problem may be terminated at any time, and hence feasibility must be maintained at all times. This paper addresses these challenges by adopting a novel control-theoretic approach that combines continuous-time gradient flows to optimize the objective function with techniques from control barrier functions to keep the evolution within the feasible set.

Literature Review: Dynamical systems and optimization often go hand in hand [1], [2], [3], going back all the way to the use of gradient descent techniques to solve unconstrained optimization problems. Of particular relevance to the present paper are works that have explored the use of dynamical systems to solve constrained optimization problems while ensuring that the evolution remains feasible. For problems involving only equality constraints, [4] employs differential geometric techniques to design a vector field that maintains feasibility along the flow, makes the constraint set asymptotically stable, and whose solutions converge to critical points of the objective function. The work [5] introduces a generalized form of this vector field to deal with inequality constraints in the form of a differential algebraic equation (DAE), and explores links between the DAE and sequential quadratic programming (SQP). Recently, this work has been extended in [6] to study the region of attraction of local minima, showing that the introduction of a stochastic perturbation allows solutions to escape sharp local minimizers.

An alternative approach for solving optimization problems in continuous time makes use of projected dynamical systems [7] by projecting the gradient of the objective function onto the cone of feasible descent directions, see e.g., [8], [9]. Such projection ensures feasibility at all times, which is particularly useful in applications where the optimization problem is in a feedback loop with a plant (e.g, providing setpoints, specifying optimization-based controller). Examples of this setting are numerous in power systems [10], [11], network congestion control [12], and transportation [13]. However, projected dynamical systems are, in general, discontinuous, which from an analysis viewpoint requires properly dealing with notions and existence of solutions, cf. [14], and from a practical viewpoint raises challenges for implementation. Our approach here uses techniques from safety-critical control, namely the concept of control barrier function [15], [16]. The typical formulation of a safety-critical control problem involves a subset of the state-space that represents the region where the system can operate safety. The concept of control barrier function allows to identify the range of inputs that keep the state safe without overconstraining the system evolution, leading to the synthesis of feedback controller that enforce forward invariance and asymptotic stability of the safe set.

Statement of Contributions: In this paper we introduce a continuous-time dynamical system to solve constrained optimization problems such that feasible set is forward invariant and asymptotically stable. Our technical approach demonstrates that the framework of safety-critical control naturally carries over to the setting of constrained optimization. The basic intuition is to combine the standard gradient flow to optimize the objective function with the idea of keeping the collection of feasible points safe. To maintain safety, we augment the gradient flow with an input whose role is to make sure the feasible set remains safe. We do this by using a control barrier function approach to design an optimization-based feedback controller that ensures forward invariance and asymptotic stability of the feasible set. The proposed approach can handle both equality and inequality constraints. We show that the resulting closed-loop system is locally Lipschitz, is well defined on an open set containing the feasibility region, and is a continuous approximation of the projected

This work was supported by NSF Award IIS-2007141.

A. Allibhoy and J. Cortés are with the Department of Mechanical and Aerospace Engineering, UC San Diego, {aallibho,cortes}@ucsd.edu

gradient flow. Further, we establish that its equilibria exactly correspond to the critical points of the optimization problem, the objective function is monotonically decreasing along the feasible set, and identify conditions for convergence to and stability of the minimizers. Finally, we illustrate the proposed approach on both convex and nonconvex example problems. For reasons of space, the proofs are omitted and will appear elsewhere.

II. Preliminaries

We introduce here the notation and stability definitions, and recall basic facts from control barrier functions.

Notation: For $v, w \in \mathbb{R}^n$, $v \leq w$ (resp. $v < w$) denotes $v_i \leq w_i$ (resp. $v_i < w_i$) for all $i = 1, 2, \dots, n$. Given $g : \mathbb{R}^n \to \mathbb{R}^m$, we denote its Jacobian by $\frac{\partial g}{\partial x}$. In case where $m = 1$, we denote the gradient by ∇g. We let $I_0(x) = \{1 \leq i \leq m \mid g_i(x) = 0\}$ denote the active constraint set. The matrix whose rows are $\nabla g_i(x)^\top$ for $i \in I_0(x)$ is denoted $\frac{\partial g_I(x)}{\partial x}$. The Lie derivative of $g : \mathbb{R}^n \to \mathbb{R}$ along a vector field $F : \mathbb{R}^n \to \mathbb{R}^n$ is $\mathcal{L}_F g : \mathbb{R}^n \to \mathbb{R}$ defined by $\mathcal{L}_F g(x) = \nabla g(x)^\top F(x)$. Given $A \in \mathbb{R}^{n \times m}$, the Moore-Penrose pseudoinverse of A is denoted $A^\dagger$.

Invariance and Stability Notions: We recall basic definitions from the theory of ordinary differential equations [17]. Let $F : \mathbb{R}^n \to \mathbb{R}^n$ be a locally Lipschitz continuous vector field and consider the dynamical system $\dot{x} = F(x)$. Local Lipschitz continuity ensures that for every initial condition $x_0 \in \mathbb{R}^n$, there exists a $T > 0$ and a unique function $x : [0, T] \to \mathbb{R}^n$ such that $x(0) = x_0$ and $\dot{x}(t) = F(x(t))$. The system is *forward complete* if it admits a solution on the interval $[0, \infty)$ for all initial conditions. The *flow map* is the function $\Phi_t : \mathbb{R}^n \to \mathbb{R}^n$ such that $\Phi_t(x) = x(t)$ where $x(t)$ is the unique solution with $x(0) = x$.

A point $x^* \in \mathbb{R}^n$ such that $F(x^*) = 0$ is an *equilibrium*. A set $\mathcal{M} \subset \mathbb{R}^n$ is *forward invariant* if $x \in \mathcal{M}$ implies that $\Phi_t(x) \in \mathcal{M}$ for all $t \geq 0$. If $\mathcal{M}$ is forward invariant and $x^* \in \mathcal{M}$ is an equilibrium, we say that x^* is *Lyapunov stable relative to* $\mathcal{M}$ if for every open set U containing x, there exists an open set $\tilde{U}$ also containing x such that for all $y \in \tilde{U} \cap \mathcal{M}$, $\Phi_t(y) \in U \cap \mathcal{M}$ for all $t > 0$. We say that x^* is *asymptotically stable relative to* $\mathcal{M}$ if it is Lyapunov stable relative to $\mathcal{M}$ and $\Phi_t(y) \to x^*$ as $t \to \infty$ for all $y \in \tilde{U} \cap \mathcal{M}$. Analogous definitions of relative stability can be made for subsets.

Control Barrier Functions: We recall here basic notions on control barrier functions, and refer the reader to [16] for further details. Consider the control affine system

$$\dot{x} = F_0(x) + \sum_{i=1}^{m} u_i F_i(x) \tag{1}$$

where $F_i : \mathbb{R}^n \to \mathbb{R}^n$, $0 \leq i \leq m$ are locally Lipschitz. Let $\mathcal{S} \subset \mathbb{R}^n$ be the *safe set*, representing the set of states where the system can operate safely. Let U be an open set containing $\mathcal{S}$ and $u : U \to \mathbb{R}^m$ a locally Lipschitz feedback controller. The closed loop system is called *safe* with respect to $\mathcal{S}$ if $\mathcal{S}$ is forward invariant.

Let $g : \mathbb{R}^n \to \mathbb{R}^m$ and define $\mathcal{S} = \{x \in \mathbb{R}^n \mid g(x) \leq 0\}$. For $\alpha > 0$, define the admissible control set at $x \in \mathbb{R}^n$ as,

$$K(x) = \Big\{ u \in \mathbb{R}^m \mid \mathcal{L}_{F_0} g(x) + \sum_{i=1}^{m} u_i \mathcal{L}_{F_i} g(x) + \alpha g(x) \leq 0 \Big\}.$$

If there exists an $\alpha > 0$ and an open set U containing $\mathcal{S}$ such that $K(x)$ is nonempty for all $x \in \mathcal{S}$, we say that g is a *zeroing control barrier function (ZCBF)* of $\mathcal{S}$ on U. With a feedback controller $u : U \to \mathbb{R}^m$ satisfying $u(x) \in K(x)$ for all $x \in U$, it follows that $\dot{g}(x(t)) \leq -\alpha g(x(t))$ along the trajectories of the closed-loop system and the set $\mathcal{S}$ is forward invariant and asymptotically stable.

Remark 2.1: (Standard notion of control barrier function): The original definition [15] of a ZCBF considers safe sets parameterized by real-valued functions but can be readily extended to safe sets parameterized by vector-valued ones. Additionally, the use of a positive constant α here is a specific instance of the more general notion of class $\mathcal{K}$ function employed in [15]. •

A similar strategy can be employed to ensure safety when $\mathcal{S}$ is parameterized by both equality and inequality constraints. Let $g : \mathbb{R}^n \to \mathbb{R}^m$, $h : \mathbb{R}^n \to \mathbb{R}^k$ and define

$$\mathcal{S} = \{x \in \mathbb{R}^n \mid g(x) \leq 0, h(x) = 0\}. \tag{2}$$

For $\alpha > 0$, the admissible control set at $x \in \mathbb{R}^n$ is

$$K(x) = \Big\{ u \mid \mathcal{L}_{F_0} g(x) + \sum_{i=1}^{m} u_i \mathcal{L}_{F_i} g(x) + \alpha g(x) \leq 0, \quad \mathcal{L}_{F_0} h(x) + \sum_{i=1}^{m} u_i \mathcal{L}_{F_i} h(x) + \alpha h(x) = 0 \Big\}.$$

For a feedback $u : U \to \mathbb{R}^m$ satisfying $u(x) \in K(x)$, it follows that $\dot{h}(x(t)) = -\alpha h(x(t))$, and $\dot{g}(x(t)) \leq -\alpha g(x(t))$ along the solutions of the closed-loop system. Note that when $x(0) \in \mathcal{S}$ then $h(x(t)) = 0$ for solutions of (1) and when $x(0) \in U \setminus \mathcal{S}$, solutions converge exponentially to $\mathcal{S}$, so once again $\mathcal{S}$ is forward invariant and asymptotically stable. This is summarized in the following result, which is a slight adaptation of [16, Theorem 2].

Theorem 2.2 (Safe Feedback Control): Consider the system (1) with safety set $\mathcal{S}$. Suppose there exist $\alpha > 0$ and an open set U containing $\mathcal{S}$ such that $K(x)$ is nonempty for all $x \in U$. Then a Lipschitz continuous feedback controller $u : U \to \mathbb{R}^m$ satisfying $u(x) \in K(x)$ renders $\mathcal{S}$ forward invariant and asymptotically stable.

While Theorem 2.2 gives sufficient conditions for the existence of a safe feedback controller, it does not specify how to synthesize it. A common technique [15] is to define, for each $x \in U$, $u(x)$ as the minimum-norm element of $K(x)$. Alternatively, one may consider

$$u(x) \in \underset{u \in K(x)}{\operatorname{argmin}} \left\{ \left\| \sum_{i=1}^{m} u_i F_i(x) \right\|^2 \right\}. \tag{3}$$

This has the interpretation of finding a controller which guarantees safety while modifying the drift term in (1) as little as possible. Because the equations parameterizing K are affine as a function of u, (3) is a quadratic program in u.

III. Problem Formulation

Given $f : \mathbb{R}^n \to \mathbb{R}$, $g : \mathbb{R}^n \to \mathbb{R}^m$, $h : \mathbb{R}^n \to \mathbb{R}^k$, consider the constrained nonlinear programming problem

$$\begin{aligned} \underset{x\in\mathbb{R}^n}{\text{minimize}} \quad & f(x) \\ \text{subject to} \quad & g(x) \leq 0 \\ & h(x) = 0. \end{aligned} \tag{4}$$

We let $\mathcal{M} = \{x \in \mathbb{R}^n \mid g(x) \leq 0, h(x) = 0\}$ denote the feasible set. We make the following assumption.

Assumption 1: (Linear Independence Constraint Qualification Condition): For all $x \in \mathbb{R}^n$, f, g, and h are continuously differentiable and $\{\nabla g_i(x)\}_{i\in I_0(x)} \cup \{\nabla h_i(x)\}_{1\leq i\leq k}$ is linearly independent.

We call $x^* \in \mathbb{R}^n$ a *Karash-Kuhn-Tucker (KKT) point* if there exist $u^* \in \mathbb{R}^m$ and $v^* \in \mathbb{R}^p$ such that

$$\nabla f(x^*) + \frac{\partial g(x^*)}{\partial x}^\top u^* + \frac{\partial h(x^*)}{\partial x}^\top v^* = 0 \tag{5a}$$

$$g(x^*) \leq 0 \tag{5b}$$

$$h(x^*) = 0 \tag{5c}$$

$$u^* \geq 0 \tag{5d}$$

$$(u^*)^\top g(x^*) = 0 \tag{5e}$$

We denote by X_{KKT} the collection of all KKT points. Under the LICQ condition, the KKT conditions are necessary conditions for optimality of (4).

Our goal is to design a locally Lipschitz-continuous vector field $F : \mathbb{R}^n \to \mathbb{R}^n$ such that the feasible set $\mathcal{M}$ is forward invariant and asymptotically stable with respect to $\dot{x} = F(x)$, and all trajectories converge to X_{KKT}. This dynamical system is then an *anytime* algorithmic solution to (4), meaning that it is guaranteed to return a feasible solution regardless of when it is terminated. The anytime property is particularly important for real-time applications when a feasible solution must be obtained within a fixed time horizon. Further, the anytime property is desirable for settings where the optimization algorithm is in a feedback loop with a dynamical process, such as in model predictive control.

IV. Design of Safe Gradient Descent

In this section we propose a solution to the synthesis problem described in Section III. We refer to it as the *safe gradient flow* for reasons that will become clear in the course of its design. Our strategy is to identify a control barrier function to enforce forward invariance and asymptotic stability of the feasible set $\mathcal{M}$ and then synthesize a feedback controller using a quadratic program. In the next section, we analyze the continuity properties of the proposed design. For clarity of exposition, we first consider the case with only equality constraints and then we move on to the general case.

A. Design with Only Equality Constraints

Here, we consider the special case where there are only equality constraints in (4). Consider the control-affine system

$$\dot{x} = -\nabla f(x) - \sum_{i=1}^{k} u_i \nabla h_i(x) = -\nabla f(x) - \frac{\partial h(x)}{\partial x}^\top u. \tag{6}$$

One can view this dynamics as the standard gradient flow of f modified by an "input". The intuition is that the gradient direction takes care of optimizing f towards a minimizer inside the feasible set, and this direction only needs to be modified if the state gets close to the boundary of the feasible set. Following Section II, the admissible control set is

$$\begin{aligned} K(x) &= \Big\{u \in \mathbb{R}^k \mid -\mathcal{L}_{\nabla f} h(x) - \sum_{i=1}^{k} \mathcal{L}_{\nabla h_i} h(x) = -\alpha h(x)\Big\} \\ &= \Big\{u \in \mathbb{R}^k \mid -\frac{\partial h}{\partial x}\nabla f(x) - \frac{\partial h}{\partial x}\frac{\partial h}{\partial x}^\top u = -\alpha h(x)\Big\}. \end{aligned}$$

Note that in this case, the admissible control set is a singleton, and the control can be expressed in closed form as

$$u(x) = -\Big(\frac{\partial h(x)}{\partial x}\frac{\partial h(x)}{\partial x}^\top\Big)^{-1}\Big(\frac{\partial h(x)}{\partial x}\nabla f(x) - \alpha h(x)\Big).$$

Remark 4.1: (Connection with the Literature): The problem with equality constrains has also been considered in [4] from a differential geometric perspective, and extended in [5], [6], and in fact, the proposed solution exactly corresponds to the one here, as we explain next. Under the assumption that $h \in \mathcal{C}^r$ and LICQ holds, the feasible set $\mathcal{M} = \{x \in \mathbb{R}^n \mid h(x) = 0\}$ is an embedded C^r submanifold of $\mathbb{R}^n$ of codimension k. For each $x \in \mathcal{M}$, let $\mathcal{T}_x\mathcal{M} = \ker \frac{\partial h}{\partial x}$ be the tangent space to $\mathcal{M}$. The approach in [4] proceeds by identifying a vector field $F : \mathbb{R}^n \to \mathbb{R}^n$ satisfying: (i) $F \in \mathcal{C}^r$ and $F(x) \in \mathcal{T}_x\mathcal{M}$ for all $x \in \mathcal{M}$; and (ii) $\dot{h}(x) = -\alpha h(x)$ along the trajectories of $\dot{x} = F(x)$. The vector field satisfying both properties is

$$F(x) = -\Big(I - \frac{\partial h(x)}{\partial x}^\dagger\frac{\partial h(x)}{\partial x}\Big)\nabla f(x) - \alpha\frac{\partial h(x)}{\partial x}^\dagger h(x). \tag{7}$$

Under Assumption 1, all bounded trajectories of $\dot{x} = F(x)$ converge to KKT points, and $\mathcal{M}$ is forward invariant and asymptotically stable. Note that (6) with the controller u^* is exactly (7). This provides an alternative interpretation of the design in [4] from a control-theoretic perspective. •

B. Design with Equality and Inequality Constraints

Here we consider the general case with both inequality and equality constraints in (4). Consider

$$\dot{x} = -\nabla f(x) - \frac{\partial g(x)}{\partial x}^\top u - \frac{\partial h(x)}{\partial x}^\top v. \tag{8}$$

This has the same interpretation as before: the gradient of f to minimize the function along with additional inputs to ensure the feasible set $\mathcal{M}$ remains safe. In this case, the admissible control set is

$$\begin{aligned} K(x) = \Big\{(u, v) \in \mathbb{R}^m \times \mathbb{R}^k \Big| \\ -\frac{\partial g}{\partial x}\frac{\partial g}{\partial x}^\top u - \frac{\partial g}{\partial x}\frac{\partial h}{\partial x}^\top v \leq \frac{\partial g}{\partial x}\nabla f(x) - \alpha g(x) \\ -\frac{\partial h}{\partial x}\frac{\partial g}{\partial x}^\top u - \frac{\partial h}{\partial x}\frac{\partial h}{\partial x}^\top v = \frac{\partial h}{\partial x}\nabla f(x) - \alpha h(x)\Big\}. \end{aligned}$$

In the presence of both inequality and equality constraints, the set of admissible controls is no longer a singleton. To synthesize a feedback control, we use a controller of the form (3), resulting in the following quadratic program:

$$\begin{bmatrix} u(x) \\ v(x) \end{bmatrix} = \underset{u,v \in K(x)}{\operatorname{argmin}} \left\{ \left\| \frac{\partial g(x)}{\partial x}^\top u + \frac{\partial h(x)}{\partial x}^\top v \right\|^2 \right\} \quad (9)$$

Unlike the case with only equality constraints, the feedback controller (9) no longer has a convenient closed-form expression. However, we remark that the computation of u^* and projected gradient methods are similar in terms of computational complexity. Furthermore, we show in our ensuing discussion that, under mild assumptions on f, g, h, the controller (9) is Lipschitz continuous and implemented over the dynamics (8) makes all trajectories converge to KKT points while rendering $\mathcal{M}$ forward invariant and asymptotically stable.

Remark 4.2: (Inequality Constraints via Quadratic Slack Variables): The parallelism outlined in Remark 4.1 between the differential geometric and control-theoretic approaches justifies interpreting (8) with the controller (9) as the natural extension of the treatment in [4] to the case with inequality and equality constraints. The work [5] pursues a different approach and instead of dealing with inequality constraints directly, reduces them to equality constraints by introducing quadratic slack variables. Formally, one replaces the constraints $g_i(x) \geq 0$ with equality constraints $g_i(x) = -y_i^2$, and solves the equality-constrained optimization problem in the variables $(x, y) \in \mathbb{R}^{n+m}$ with a flow of the form (7). While this method can be expressed in closed-form, there are several drawbacks with it. First, this increases the dimensionality of the problem. Second, adding quadratic slack variables also introduces equilibrium points to the resulting flow which do not correspond to KKT points of the original problem. •

V. Analysis of Safe Gradient Descent

In this section we analyze the properties of the safe gradient flow introduced in Section IV. We first show that the system (8) with the controller (9) can be interpreted as a continuous approximation of the projected gradient flow. This alternative interpretation is then used to analyze the stability and convergence properties of the closed-loop system.

A. Relationship with Projected Gradient Flow

Here we show how the safe gradient descent can be thought of as a continuous approximation of the projected gradient flow. The latter is a discontinuous dynamical system obtained by projecting the gradient of the objective function onto the set of feasible descent directions.

When Assumption 1 is satisfied and $f, g \in \mathcal{C}^r$, the feasible set $\mathcal{M}$ is an embedded $\mathcal{C}^r$ submanifold with boundary of $\mathbb{R}^n$, cf. [18], having codimension k [19, Lemma 3.1.12]. Given $x \in \mathcal{M}$, the tangent cone at x is defined as

$$\mathcal{T}_x\mathcal{M} = \left\{ \xi \in \mathbb{R}^n \;\middle|\; \frac{\partial h}{\partial x}\xi = 0, \frac{\partial g_I}{\partial x}\xi \leq 0 \right\}.$$

For $x \in \mathcal{M}$, the projection onto $\mathcal{T}_x\mathcal{M}$ is defined by $\Pi_x(\xi) = \arg\min_{\xi' \in \mathcal{T}_x\mathcal{M}}\{\|\xi' - \xi\|\}$. For a general set, the projection operator is a set-valued map, but the fact that $\mathcal{T}_x\mathcal{M}$ is closed and convex makes the projection onto $\mathcal{T}_x\mathcal{M}$ unique in this case. The projected gradient flow is the dynamical system

$$\begin{aligned} \dot{x} = \; & \Pi_x(-\nabla f(x)) \\ = \; & \underset{\xi \in \mathbb{R}^n}{\operatorname{argmin}} \frac{1}{2} \|\xi + \nabla f(x)\|^2 \\ & \text{subject to} \frac{\partial g_I(x)}{\partial x}\xi \leq 0, \frac{\partial h(x)}{\partial x}\xi = 0. \end{aligned} \quad (10)$$

In general, this system is discontinuous, so one must resort to appropriate notions of solution trajectories and establish their existence, see e.g., [14]. Here, we consider Carathéodory solutions, which are absolutely continuous functions that satisfy (10) almost everywhere. When Carathéodory solutions exist in $\mathcal{M}$, then the local minimizers of (4) are equilibria of (10). Additionally, the asymptotically stable equilibria are strict local minimizers, and strict local minimizers are always stable (though not necessarily asymptotically stable, cf. [20]).

To derive a continuous approximation of (10), let $\alpha > 0$ and define $F_\alpha : \mathbb{R}^n \to \mathbb{R}^n$ by

$$\begin{aligned} F_\alpha(x) = \underset{\xi \in \mathbb{R}^n}{\operatorname{argmin}} \quad & \frac{1}{2} \|\xi + \nabla f(x)\|^2 \\ \text{subject to} \quad & \frac{\partial g(x)}{\partial x}\xi \leq -\alpha g(x) \\ & \frac{\partial h(x)}{\partial x}\xi = -\alpha h(x) \end{aligned} \quad (11)$$

Note that (11) has a similar form to (10), with the key difference that $\Pi_x(-\nabla f)$ is defined only on $\mathcal{M}$ whereas, as we show below, F_α is well defined on an open set U containing $\mathcal{M}$, allowing for infeasible initial conditions. Furthermore, under mild regularity assumptions on f, g, and h, the vector field F_α is Lipschitz continuous.

Proposition 5.1: (Feasibility and regularity of F_α): Suppose f, g and h are continuously differentiable, and their derivatives are locally Lipschitz. Then

(i) there exists an open neighborhood $\mathcal{U}$ containing $\mathcal{M}$ such that (11) is well defined;
(ii) F_α is locally Lipschitz on $\mathcal{U}$.

Remark 5.2: (Global Feasibility of Continuous Approximation): The open neighborhood U where (11) is well defined depends on the choice of α. In many examples, we observe that these parameters can be chosen so that the domain of the controller is $U = \mathbb{R}^n$. •

Next, we show that F_α approximates the projected gradient in the sense that, for $x \in \mathcal{M}$, $F_\alpha(x) \to \Pi_x(-\nabla f(x))$ as $\alpha \to \infty$. Intuitively, this is because for inactive constraints $j \notin I_0(x)$, one has $g_j(x) < 0$ and hence the jth inequality constraint in (11), $\nabla g_j(x)^\top \xi \leq -\alpha g_j(x)$, becomes $\nabla g_j(x)^\top \xi \leq \infty$ as $\alpha \to \infty$ and the constraint is effectively removed, reducing the problem to (10).

Proposition 5.3: (F_α approximates the projected gradient): For all $x \in \mathcal{M}$, $F_\alpha(x) \in \mathcal{T}_x\mathcal{M}$ and $\lim_{\alpha\to\infty} F_\alpha(x) = \Pi_x(-\nabla f(x))$.

Next, we show that the vector field F_α is equivalent to the closed-loop system resulting from implementing the controller (9) over the system (8).

Proposition 5.4: (Equivalence between Control Barrier QP-Controller and Continuous Approximation to Projected Gradient): Let $\mathcal{U}$ be an open set containing $\mathcal{M}$ on which (11) is well-defined. We have the following for all $x \in \mathcal{U}$:

(i) If $u \in \mathbb{R}^m$ and $v \in \mathbb{R}^k$ are Lagrange multipliers of (11) corresponding to x, then (u, v) solves (9)

(ii) For any (u, v) that solves (9) we have

$$F_\alpha(x) = -\nabla f(x) - \frac{\partial g(x)}{\partial x}^\top u - \frac{\partial h(x)}{\partial x}^\top v.$$

We note that because both (10) and (11) are least-squares problems of the same dimension subject to affine constraints, the computational complexity of computing either is equivalent. Furthermore, there are several distinct advantages of F_α as compared to the projected gradient flow. First, by Proposition 5.1, F_α is locally Lipschitz, so classical solutions to the dynamics $\dot{x} = F_\alpha(x)$ are guaranteed to exist, and the continuous-time flow can be numerically solved using standard ODE discretization schemes. Secondly, $F_\alpha(x)$ is defined for initial conditions outside $\mathcal{M}$, allowing us to guarantee convergence to a local minimizer starting from infeasible initial conditions.

B. Stability Properties of Safe Gradient Descent

Here we establish the stability properties of the safe gradient flow. We first show that the that set of equilibria of F_α is exactly X_{KKT}. This can be seen by comparing the KKT conditions of (11) with (5)

Proposition 5.5: (Equilibria of Safe Gradient Flow Correspond to KKT Points): $F_\alpha(x^*) = 0$ if and only if $x^* \in X_{\text{KKT}}$.

Next, by Proposition 5.4 the safe gradient flow satisfies the conditions in Theorem 2.2, from which it follows that $\mathcal{M}$ is safe.

Proposition 5.6 (Asymptotic Stability of $\mathcal{M}$): The feasible set $\mathcal{M}$ is forward invariant, and $\mathcal{M}$ is asymptotically stable on $\mathcal{U}$.

To show stability, we first note that f is monotonically decreasing along the safe gradient flow on $\mathcal{M}$.

Lemma 5.7: (Objective function decreases on $\mathcal{M}$): For $x \in \mathcal{M}$, $\mathcal{L}_{F_\alpha} f(x) \leq 0$ with equality if and only if $x \in X_{\text{KKT}}$.

Finally, by combining the monotonic decrease of the objective function on the feasible set along with Proposition 5.5, we can show the stability of isolated local minimizers using f as a Lyapunov function.

Proposition 5.8: (Stability of isolated KKT points):

(i) If $x^* \in X_{\text{KKT}}$ is a strict local minimizer of f and an isolated equilibrium then x^* is asymptotically stable relative to $\mathcal{M}$.

(ii) If x^* is asymptotically stable relative to $\mathcal{M}$, then x^* is a local minimum.

VI. Examples

In this section we demonstrate the safe gradient flow on several example problems. To illustrate the versatility of the approach outlined here, we consider both convex and nonconvex problems.

A. Quadratic Program

Consider the problem of minimizing a quadratic function subject to affine constraints:

$$\begin{aligned}
\underset{x\in\mathbb{R}^2}{\text{minimize}} \quad & \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}^\top \begin{bmatrix} 0.148 & -0.033 \\ -0.033 & 0.101 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \\
\text{subject to} \quad & \begin{bmatrix} -0.756 & 0.648 \\ 0.320 & -0.895 \\ 0.382 & 0.263 \end{bmatrix} x \leq \begin{bmatrix} 0.659 \\ 0.703 \\ 0.413 \end{bmatrix} \\
& \begin{bmatrix} -0.908 & 0.919 \end{bmatrix} x = 0.413
\end{aligned} \tag{12}$$

Figure 1(left) shows the phase portrait of the safe gradient descent, $\dot{x} = F_\alpha(x)$, with $\alpha = 1.0$, with the level sets of the objective function overlaid in green. Figure 1(right) shows a surface plot of the objective function. In both plots, the blue shaded region is where the inequality constraints are satisfied. All trajectories converge to the unique global minimizer.

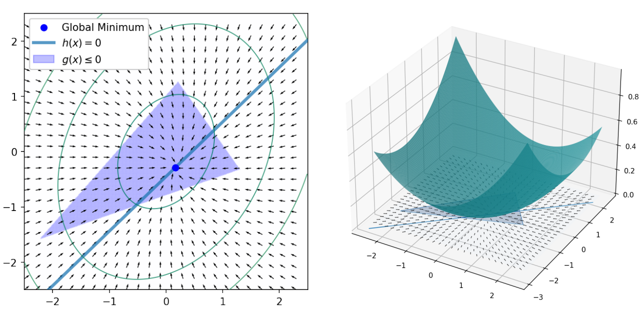

Fig. 1: (Left) Phase portrait of safe gradient descent with feasible region in blue and level sets of the objective function overlaid. (Right) Surface plot of the objective function.

B. Rosenbrock Function With Either Convex or Nonconvex Constraints

Here we consider two examples with the Rosenbrock function as objective. First, from [21], let

$$\begin{aligned}
\underset{x\in\mathbb{R}^2}{\text{minimize}} \quad & (1-x_1)^2 + 100(x_2 - x_1^2)^2 \\
\text{subject to} \quad & x_1^2 + x_2^2 \leq 2
\end{aligned} \tag{13}$$

Figure 2(left) shows the phase portrait of the safe gradient descent, $\dot{x} = F_\alpha(x)$, with $\alpha = 1.0$ (the parameter α is irrelevant since there are only inequality constraints). With this choice of parameter value, the dynamics are well-defined on $U = \mathbb{R}^2$. Figure 2(right) shows a surface plot of the objective function.

Next, consider from [22]

$$\begin{aligned}
\underset{x\in\mathbb{R}^2}{\text{minimize}} \quad & (1-x_1)^2 + 100(x_2 - x_1^2)^2 \\
\text{subject to} \quad & (x_1 - 1)^3 - x_2 + 1 \leq 0 \\
& x_1 + x_2 - 2 \leq 0
\end{aligned} \tag{14}$$

Figure 3(left) shows the phase portrait of the safe gradient descent, $\dot{x} = F_\alpha(x)$, with $\alpha = 1.0$. With this choice of parameter value, the dynamics are well-defined on $U = \mathbb{R}^2$. Figure 3(right) shows a surface plot of the objective function.

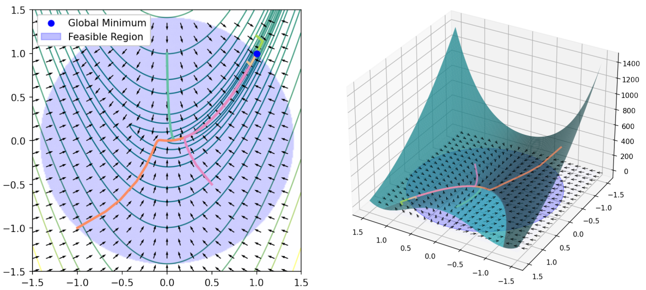

Fig. 2: (Left) Phase portrait of safe gradient descent with feasible region in blue and level sets overlaid. Solutions of the flow are plotted in pink, orange, yellow, green, and blue. (Right) Surface plot of the objective function.

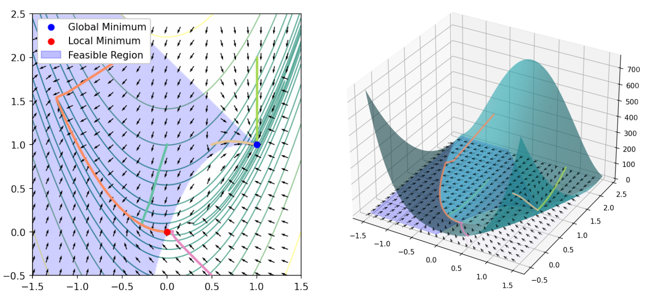

Fig. 3: (Left) Phase portrait of safe gradient descent with feasible region in blue and level sets overlaid. Solutions of the flow are plotted in pink, orange, yellow, green, and blue. (Right) Surface plot of the objective function.

Both problems (13) and (14) have a unique strict global minimizer $f(1.0, 1.0) = 0.0$. The simulations reveal that the global minimizer is asymptotically stable. For problem (13), all trajectories converge to the global minimizer. For problem (14), stability only holds locally, and its possible for trajectories with both feasible and infeasible initial conditions to converge to a local minimizer.

VII. Conclusions

We have introduced a continuous-time dynamical system to solve constrained optimization problems while making the feasible set forward invariant. We showed that the proposed dynamics is a continuous approximation of the projected gradient flow, with equilibria corresponding to the critical points of the optimization problem and monotonically decreasing the objective function while evolving in the feasible set. We also established the asymptotic stability of the feasible set (meaning that the dynamics can be initialized at infeasible points) and identified conditions that guarantee convergence to the set of minimizers. Future work will explore the relationship between the domain of the feedback controller and the design parameters, study the input-to-state stability properties of the proposed dynamics, and develop discretizations and their relationship with discrete-time iterative methods for nonlinear programming. We also hope to extend this framework to Newton-like flows for nonlinear programs which incorporate higher-order information and apply the results to real-time optimal feedback control problems.

References

[1] K. Arrow, L. Hurwitz, and H. Uzawa, *Studies in Linear and Non-Linear Programming*. Stanford, CA: Stanford University Press, 1958.

[2] R. W. Brockett, "Dynamical systems that sort lists, diagonalize matrices, and solve linear programming problems," *Linear Algebra and its Applications*, vol. 146, pp. 79–91, 1991.

[3] U. Helmke and J. B. Moore, *Optimization and Dynamical Systems*. Springer, 1994.

[4] K. Tanabe, "A geometric method in nonlinear programming," *Journal of Optimization Theory and Applications*, vol. 30, no. 2, pp. 181–210, 1980.

[5] J. Schropp and I. Singer, "A dynamical systems approach to constrained minimization," *Numerical functional analysis and optimization*, vol. 21, no. 3-4, pp. 537–551, 2000.

[6] H. Feng, H. Zhang, and J. Lavaei, "A dynamical system perspective for escaping sharp local minima in equality constrained optimization problems," in *IEEE Conf. on Decision and Control*, (Jeju Island, Republic of Korea), pp. 4255–4261, IEEE, 2020.

[7] A. Nagurney and D. Zhang, *Projected Dynamical Systems and Variational Inequalities with Applications*, vol. 2 of *International Series in Operations Research and Management Science*. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1996.

[8] A. Hauswirth, S. Bolognani, and F. Dörfler, "Projected dynamical systems on irregular, non-Euclidean domains for nonlinear optimization," *SIAM Journal on Control and Optimization*, vol. 59, no. 1, pp. 635–668, 2021.

[9] V. Shikhman and O. Stein, "Constrained optimization: projected gradient flows," *Journal of Optimization Theory & Applications*, vol. 140, no. 1, pp. 117–130, 2009.

[10] A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler, "Optimization algorithms as robust feedback controllers," *preprint*, 2021.

[11] M. Colombino, E. Dall'Anese, and A. Bernstein, "Online optimization as a feedback controller: Stability and tracking," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 1, pp. 422–432, 2020.

[12] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems*, vol. 22, no. 1, pp. 28–43, 2002.

[13] G. Bianchin, J. Cortés, J. I. Poveda, and E. Dall'Anese, "Time-varying optimization of LTI systems via projected primal-dual gradient flows," *IEEE Transactions on Control of Network Systems*, 2021. To appear.

[14] J. Cortés, "Discontinuous dynamical systems - a tutorial on solutions, nonsmooth analysis, and stability," *IEEE Control Systems*, vol. 28, no. 3, pp. 36–73, 2008.

[15] 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.

[16] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: theory and applications," in *European Control Conference*, (Naples, Italy), pp. 3420–3431, June 2019.

[17] G. Teschl, *Ordinary differential equations and dynamical systems*, vol. 140. American Mathematical Soc., 2012.

[18] J. M. Lee, "Smooth manifolds," in *Introduction to Smooth Manifolds*, pp. 1–31, Springer, 2013.

[19] H. T. Jongen, P. Jonker, and F. Twilt, *Nonlinear optimization in finite dimensions: Morse theory, Chebyshev approximation, transversality, flows, parametric aspects*, vol. 47. Springer Science & Business Media, 2013.

[20] A. Hauswirth, S. Bolognani, G. Hug, and F. Dörfler, "Projected gradient descent on Riemannian manifolds with applications to online power system optimization," in *Allerton Conf. on Communications, Control and Computing*, (Monticello, IL), pp. 225–232, IEEE, 2016.

[21] Mathworks Inc., "Solve a constrained nonlinear problem, solver-based," *MATLAB Optimization Toolbox Documentation*, 2021.

[22] P.-A. Simionescu and D. G. Beale, "New concepts in graphic visualization of objective functions," in *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, vol. 36223, pp. 891–897, 2002.