

OpEn: Code Generation for Embedded Nonconvex Optimization^{*}

Pantelis Sopasakis^{*}, Emil Fresk^{**} and Panagiotis Patrinos^{***}

^{*} *Queen's University Belfast, School of Electronics, Electrical Engineering and Computer Science (EEECs) and Centre for Intelligent and Autonomous Manufacturing Systems (i-AMS), Ashby Building, Stranmillis Road, BT9 5AG Belfast, Northern Ireland*

^{**} *WideFind AB, Aurorum 1C, 977 75 Luleå, Sweden*

^{***} *KU Leuven, Department of Electrical Engineering (ESAT), STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics & Optimization in Engineering (OPTEC), Kasteelpark Arenberg 10, 3001 Leuven, Belgium.*

Abstract: We present Optimization Engine (OpEn): an open-source code generation framework for real-time embedded nonconvex optimization, which implements a novel numerical method. OpEn combines the proximal averaged Newton-type method for optimal control (PANOC) with the penalty and augmented Lagrangian methods to compute approximate stationary points of nonconvex problems. The proposed method involves very simple algebraic operations such as vector products, has a low memory footprint and exhibits very good convergence properties that allow the solution of nonconvex problems on embedded devices. OpEn's core solver is written in Rust — a modern, high-performance, memory-safe and thread-safe systems programming language — while users can call it from Python, MATLAB, C, C++, ROS or over a TCP socket.

Keywords: Embedded numerical optimization; nonconvex optimization problems; code generation; model predictive control; moving horizon estimation; Rust; Robot Operating System

1. INTRODUCTION

In embedded applications it is often necessary to solve optimization problems in real time. Typical examples involve moving horizon estimation, model predictive control and online learning (Ferreau et al., 2017). The distinctive features of embedded real time computing are the limited computing and storage capabilities of hardware devices, the presence of stringent runtime requirements, the need for memory safety and reliability, the need for simple numerical methods which can be verified easily, and the need for design paradigms that involve code generation.

A significant research effort has been dedicated to numerical optimization software with code generation capabilities for convex problems such as CVXGen (Mattingley and Boyd, 2012), SPLIT (Shukla et al., 2017), OSQP (Banjac et al., 2017), and μ AO-MPC (Zometa et al., 2013). Kouzoupis et al. (2015) provide benchmarks of numerical methods for embedded quadratic programming.

Nonconvex optimization problems are typically solved with sequential convex programming (SCP) (Zillober et al., 2004) and interior point (IP) methods (Zanelli et al., 2017). Several software are available for nonconvex opti-

mization such as SNOPT (Gill et al., 2005), Acado (Verschueren et al., 2018) and IPOPT (Wächter and Biegler, 2006).

Nonetheless, SCP and IP methods have a high per-iteration computation cost and involve expensive operations such as the solution of quadratic programming problems or the solution of linear systems. Was it not for their slow convergence, first-order methods such as the projected gradient method, would be ideal for embedded applications. Stella et al. (2017) proposed the proximal averaged Newton-type method for optimal control (PANOC), which uses exactly the same oracle as the projected gradient method but exhibits superior convergence characteristics and has been found to outperform SCP and IP methods (Sathya et al., 2018; Small et al., 2019).

PANOC can solve problems with a smooth cost function and simple constraints on which one can compute projections. More general constraints can be relaxed and replaced by penalty functions (soft constraints). The penalty method can then be applied to obtain near-feasible solutions (Hermans et al., 2018). However, as the penalty parameter needs to grow unbounded, it is likely that it will overshadow the cost function and lead to highly ill conditioned problems. In this paper we propose a numerical method to determine stationary points of constrained nonconvex optimization problems by combining the penalty and the augmented Lagrangian methods. Moreover, we present OpEn: an open-source code generation tool for real-time embedded nonconvex optimization with guaranteed

^{*} The work of P. Patrinos was supported by: FWO projects: G086318N, G086518N, G0A0920N; Fonds de la Recherche Scientifique – FNRS, the Fonds Wetenschappelijk Onderzoek – Vlaanderen under EOS Proj. 30468160 (SeLMA), Research Council KU Leuven C1 proj. No. C14/18/068 and the Ford–KU Leuven Research Alliance project. Corresp. author email: p.sopasakis@qub.ac.uk.

safe memory management. In Section 5 we present two simulation examples – a model predictive controller for obstacle avoidance and a constrained nonlinear estimator – and provide comparisons with IPOPT and `scipy`'s SQP solver.

Preliminaries and Notation

We define $\mathbb{N}_N = \{0, 1, \dots, N\}$. We denote the set of extended-real numbers by $\mathbb{R} = \mathbb{R} \cup \{\infty\}$. The Jacobian of a mapping $F : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is denoted by JF . For two vectors, $x, x' \in \mathbb{R}^n$, the relation $x \leq x'$ is meant in the pointwise sense. For $x \in \mathbb{R}$, we define $[x]_+ = \max\{0, x\}$ and for $x \in \mathbb{R}^n$, $[x]_+$ is defined element wise. The projection of a vector $x \in \mathbb{R}^n$ on a nonempty, closed, convex set $C \subseteq \mathbb{R}^n$ is defined as $\Pi_C(x) = \operatorname{argmin}_{y \in C} \|y - x\|$ and the distance of $x \in \mathbb{R}^n$ from C is defined as $\operatorname{dist}_C(x) = \min_{y \in C} \|y - x\|$. The support function of a nonempty, closed, convex set $C \subseteq \mathbb{R}^n$ is a function $\delta_C^* : \mathbb{R}^n \rightarrow \overline{\mathbb{R}}$ defined as $\delta_C^*(y) = \sup_{x \in C} y^\top x$. For a cone $K \subseteq \mathbb{R}^n$, we define its polar to be the cone $K^\circ = \{y \in \mathbb{R}^n : y^\top x \leq 0\}$. The normal cone of a set $C \subseteq \mathbb{R}^n$ is defined as that subdifferential of δ_C , that is, $N_C(x) = \partial\delta_C(x)$, for $x \in \mathbb{R}^n$.

2. PROBLEM STATEMENT

Consider the following parametric optimization problem

$$\mathbb{P}(p) : \underset{u \in U}{\text{minimize}} f(u, p), \quad (1a)$$

$$\text{subj. to: } F_1(u, p) \in C, \quad (1b)$$

$$F_2(u, p) = 0, \quad (1c)$$

where for each $p \in \mathbb{R}^{n_p}$, $f(\cdot, p) : \mathbb{R}^n \rightarrow \mathbb{R}$ is a (possibly nonconvex) continuously differentiable function with L_f -Lipschitz gradient and $U \subseteq \mathbb{R}^n$ is a nonempty, closed — but not necessarily convex — set such that we can compute projections on it. Moreover, $F_1(\cdot, p) : \mathbb{R}^n \rightarrow \mathbb{R}^{n_1}$ is a smooth mapping with Lipschitz-continuous Jacobian which is bounded on U . Set $C \subseteq \mathbb{R}^{n_1}$ is a convex set from which we can compute distances. Mapping $F_2(\cdot, p) : \mathbb{R}^n \rightarrow \mathbb{R}^{n_2}$ is such that $\|F_2(\cdot, p)\|^2$ is a continuously differentiable function with Lipschitz-continuous gradient. In particular, constraints (1b) will be treated using the proposed augmented Lagrangian method, while constraints (1c) will be accounted for using the penalty method as we shall discuss in Section 3. Constraints of both types can be present at the same time. Constraints (1c) can accommodate multiple vertical complementarity constraints (Izmailov et al., 2012; Scheel and Scholtes, 2000) of the form

$$\prod_{i=1}^m [h_i(u, p)]_+ = 0, \quad (2)$$

a special case of which arises in collision avoidance problems (Hermans et al., 2018; Sathya et al., 2018; Small et al., 2019). Such constraints do not satisfy the smoothness conditions, so they cannot be described in terms of the ALM-type constraints in (1b).

Set U can be a Euclidean ball of radius r centered at a point u_c , $\mathcal{B}_2(u_c, r) = \{u \in \mathbb{R}^n : \|u - u_c\|_2 \leq r\}$, an infinity ball, $\mathcal{B}_\infty(u_c, r) = \{u \in \mathbb{R}^n : \|u - u_c\|_\infty \leq r\}$, a rectangle, $\{u \in \mathbb{R}^n : u_{\min} \leq u \leq u_{\max}\}$ (where some coordinates of u_{\min} and u_{\max} are allowed to be equal to

$-\infty$ and ∞ respectively), a finite set $\{u^{(1)}, u^{(2)}, \dots, u^{(a)}\}$, or a Cartesian product of any such sets.

The above formulation can accommodate equality constraints, either through constraint (1c), or by using $C = \{0\}$. It can also accommodate inequality constraints of the general form

$$H(u, p) \leq 0, \quad (3)$$

by using either $F_1(u, p) = H(u, p)$ and $C = \{v \in \mathbb{R}^{n_1} : v \leq 0\}$, or $F_2(u, p) = [H(u, p)]_+$. Furthermore, constraint (1b) allow the designer to encode norm bounds using $C = \mathcal{B}_2(u_c, r)$ or $C = \mathcal{B}_\infty(u_c, r)$, or constraints involving second-order cones, $C = \operatorname{SOC}_\alpha^{n_1} := \{u = (x, t), x \in \mathbb{R}^{n_1-1}, t \in \mathbb{R} : \|x\| \leq \alpha t\}$, for some $\alpha > 0$.

Optimal control and estimation problems can be written in the form of Problem (1) in a number of different ways. For instance, consider the following discrete-time optimal control problem

$$\mathbb{P}(x_0) : \underset{u_t}{\text{minimize}} \ell_N(x_N) + \sum_{t=0}^{N-1} \ell_t(x_t, u_t) \quad (4a)$$

$$\text{subj. to: } u_t \in U_t, \quad (4b)$$

$$x_{t+1} = \Phi_t(x_t, u_t), \quad (4c)$$

$$H_t(x_t, u_t) \leq 0, \quad (4d)$$

$$H_N(x_N) \leq 0, \quad (4e)$$

for $t \in \mathbb{N}_{N-1}$, where U_t are sets on which one can easily compute projections. One may follow the *single shooting* approach, in which the minimization is carried out over the sequence of control actions, $u = (u_0, \dots, u_{N-1})$, and the sequence of states is eliminated by defining a sequence of function $F_0(u, x_0) = x_0$ and $F_{t+1}(u, x_0) = \Phi_t(F_t(u, x_0), u_t)$ for $t \in \mathbb{N}_{N-1}$. We may then define the cost function

$$f(u, x_0) = \ell_N(F_N(u, x_0)) + \sum_{t=0}^{N-1} \ell_t(F_t(u, x_0), u_t). \quad (5)$$

A second alternative is to follow the *multiple shooting* approach where the minimization is carried out over $u = (u_0, \dots, u_{N-1}, x_1, \dots, x_N)$ and the system dynamics in Equation (4c) are treated as constraints. Joint state-input constraints of the form (4d) and terminal constraints of the form (4e) can be cast using either F_1 or F_2 as discussed above.

3. NUMERICAL ALGORITHM

3.1 Augmented Lagrangian and Penalty Methods

In order to solve the original optimization problem, we shall use the augmented Lagrangian method (Birgin and Martínez, 2014). Hereafter, we shall drop the parameter p for the sake of simplicity. We introduce the *augmented Lagrangian function*

$$L_c(u, v, y) := f(u) + \langle y, F_1(u) - v \rangle + \frac{c}{2} \|F_1(u) - v\|^2 + \frac{c}{2} \|F_2(u)\|^2, \quad (6)$$

defined for $u \in U$ and $v \in C$, where $c > 0$ is a constant penalty parameter and $y \in \mathbb{R}^{n_1}$ is the vector of Lagrange multipliers for constraints (1b). We can then show that

Proposition 1. The following holds for all $y \in \mathbb{R}^{n_1}$ and $c > 0$

$$\min_{u \in U, v \in C} L_c(u, v, y) = -\frac{1}{2c} \|y\|^2 + \min_{u \in U} \psi(u; c, y), \quad (7)$$

where

$$\psi(u; c, y) := f(u) + \frac{c}{2} \left[\text{dist}_C^2(F_1(u) + \frac{1}{c}y) + \|F_2(u)\|^2 \right]. \quad (8)$$

Proof. For $u \in U$ and $v \in C$, we have that

$$\begin{aligned} & \langle y, F_1(u) - v \rangle + \frac{c}{2} \|F_1(u) - v\|^2 \\ &= c \langle \frac{1}{c}y, F_1(u) - v \rangle + \frac{c}{2} \|F_1(u) - v\|^2 + \frac{1}{2c} \|y\|^2 - \frac{1}{2c} \|y\|^2 \\ &= \frac{c}{2} \|F_1(u) - v + \frac{1}{c}y\|^2 - \frac{1}{2c} \|y\|^2, \end{aligned}$$

therefore,

$$\begin{aligned} \min_{u \in U, v \in C} L_c(u, v, y) &= -\frac{1}{2c} \|y\|^2 + \min_{u \in U} \left\{ f(u) \right. \\ &\quad \left. + \frac{c}{2} \|F_2(u)\|^2 + \underbrace{\min_{v \in C} \frac{c}{2} \|F_1(u) - v + \frac{1}{c}y\|^2}_{\text{dist}_C^2(F_1(u) + c^{-1}y)} \right\}, \quad (9) \end{aligned}$$

which proves the assertion. \square

The above choice of an augmented Lagrangian function leads to Algorithm 1 where the *inner problem* takes the form

$$\mathbb{P}_{\text{in}}(c, y) : \underset{u \in U}{\text{minimize}} \psi(u; c, y). \quad (10)$$

This problem has a smooth cost function and simple constraints, therefore, it can be solved using PANOC as we discuss in Section 3.2. The algorithm updates both a penalty parameter, c , as well as a vector of Lagrange multipliers, $y \in \mathbb{R}^{n_1}$, corresponding to constraint (1b) according to

$$y^{\nu+1} = \bar{y}^\nu + c_\nu (F_1(u^{\nu+1}) - \Pi_C(F_1(u^{\nu+1}) + c_\nu^{-1} \bar{y}^\nu)). \quad (11)$$

Algorithm 1 Augmented Lagrangian and penalty method

Input: $u^0 \in \mathbb{R}^{n_u}$ (initial guess), $p \in \mathbb{R}^{n_p}$ (parameter), $y^0 \in \mathbb{R}^{n_1}$ (initial guess for the Lagrange multipliers), $\epsilon, \delta > 0$ (tolerances), β (tolerance decrease coefficient), ρ (penalty update coefficient), θ (sufficient decrease coefficient), $Y \subseteq \text{dom } \delta_C^*$ (compact set)

Output: (ϵ, δ) -approximate solution (u^*, y^*)

- 1: $\bar{\epsilon}_0 = \epsilon_0$
 - 2: **for** $\nu = 0, \dots, \nu_{\max}$ **do**
 - 3: $\bar{y}^\nu = \Pi_Y(y^\nu)$
 - 4: $u^{\nu+1}$ is a solution of $\mathbb{P}_{\text{in}}(c_\nu, \bar{y}^\nu)$ with tolerance $\bar{\epsilon}$ and initial guess u^ν , using Algorithm 2
 - 5: $y^{\nu+1} = \bar{y}^\nu + c_\nu (F_1(u^{\nu+1}) - \Pi_C(F_1(u^{\nu+1}) + c_\nu^{-1} \bar{y}^\nu))$
 - 6: $z_{\nu+1} = \|y^{\nu+1} - \bar{y}^\nu\|_\infty$, $t_{\nu+1} = \|F_2(u^{\nu+1})\|_\infty$
 - 7: **if** $z_{\nu+1} \leq c_\nu \delta$, $t_{\nu+1} \leq \delta$ **and** $\bar{\epsilon}_\nu \leq \epsilon$ **then**
 - 8: **return** $(u^*, y^*) = (u^{\nu+1}, y^{\nu+1})$
 - 9: **else if** $\nu > 0$, $z_{\nu+1} > \theta z_\nu$, $t_{\nu+1} > \theta t_\nu$ **then**
 - 10: $c_{\nu+1} = \rho c_\nu$
 - 11: $\bar{\epsilon}_{\nu+1} = \beta \bar{\epsilon}_\nu$
-

The most critical tuning parameter of the algorithm is the *penalty update factor*, $\rho > 1$. Typical values are between 1.1 and 20. Low values lead to sequences of inner problems which are “similar” enough so that u^ν is a good initial guess, however, at the expense of slow convergence of the infeasibility as quantified by z_ν and t_ν . On the other hand, larger values of ρ create more dissimilar inner problems thus deeming u^ν potentially not too good initial guesses, but lead to fewer outer iterations.

Note that if $u^{\nu+1}$ and $y^{\nu+1}$ lead to a significant improvement in infeasibility as measured by $z_{\nu+1}$ and $t_{\nu+1}$, then the penalty parameter is not increased. Moreover, the tolerance ϵ of the inner problems can be relaxed by starting from a given initial tolerance ϵ_0 and decreasing it gradually until it reaches the target inner tolerance, ϵ .

Set Y is taken to be a compact subset of $\text{dom } \delta_C^*$. If C is the Minkowski sum of a compact set and a cone $K \subseteq \mathbb{R}^n$, then $\text{dom } \delta_C^* = K^\circ$. For example, if C is compact, we may select $Y = [-M, M]^{n_1}$, whereas if $C = \mathbb{R}_+^{n_1}$, we select $Y = [-M, 0]^{n_1}$.

3.2 Inner Problems

Problem $\mathbb{P}_{\text{in}}(c, y)$ has a smooth cost function with gradient

$$\begin{aligned} \nabla_u \psi(u; c, y) &= \nabla f(u) + c J F_1(u)^\top \left[F_1(u) + c^{-1}y \right. \\ &\quad \left. - \Pi_C(F_1(u) + c^{-1}y) \right] + \frac{c}{2} \nabla \|F_2(u)\|^2. \quad (12) \end{aligned}$$

Under the assumptions on F_1 and F_2 , $\nabla_u \psi$ is Lipschitz in u with some Lipschitz modulus $L_\psi > 0$, therefore, \mathbb{P}_{in} falls into the framework of PANOC (Stella et al., 2017). This gradient can be computed by automatic differentiation software such as CasADi (Andersson et al., 2019).

Our aim is to determine a u^* that satisfies the first-order optimality conditions of $\mathbb{P}_{\text{in}}(c, y)$

$$u^* = T_\gamma(u^*; c, y), \quad (13)$$

where $T_\gamma : \mathbb{R}^n \rightarrow U$, with $\gamma > 0$, is the projected gradient operator, $T_\gamma(u; c, y) = \Pi_U(u - \gamma \nabla_u \psi(u; c, y))$. To that end, we may use the proximal gradient method,

$$u^{\nu+1} = T_\gamma(u^\nu; c, y), \quad (14)$$

the accumulation points of which satisfy the fixed point condition (13) whenever $\gamma < 2/L_\psi$ (Nocedal and Wright, 2006, Prop. 2.3.2). The iterations are very simple, however, convergence can be particularly slow, especially for ill-conditioned problems.

Instead, we follow the approach of PANOC (Stella et al., 2017): we define the *fixed point residual* operator $R_\gamma : \mathbb{R}^n \rightarrow \mathbb{R}^n$ as

$$R_\gamma(u; c, y) = \gamma^{-1} (u - T_\gamma(u; c, y)). \quad (15)$$

Then a $u^* \in \mathbb{R}^n$ satisfies (13) if and only if it is a zero of R_γ , that is, if it solves the nonlinear equation $R_\gamma(u^*; c, y) = 0$. The main idea is that thereon we can apply a Newton-type method of the form $u^{\nu+1} = u^\nu - H_\nu R_\gamma(u^\nu; c, y)$, where $H_\nu : \mathbb{R}^n \rightarrow \mathbb{R}^n$ are invertible linear operators that satisfy the secant condition $u^{\nu+1} - u^\nu = H_\nu (R_\gamma(u^{\nu+1}; c, y) - R_\gamma(u^\nu; c, y))$. This motivates the use of quasi-Newtonian directions such as L-BFGS, which is known to yield a good convergence rate at a low memory footprint (Nocedal and Wright, 2006, Sec. 7.2). However, convergence can only be guaranteed when u^0 is in a neighborhood of a critical point u^* . To overcome this limitation, PANOC employs the *forward-backward envelope*: a real-valued, continuous merit function $\varphi_\gamma : \mathbb{R}^n \rightarrow \mathbb{R}$ for Problem $\mathbb{P}_{\text{in}}(c, y)$ (Themelis et al., 2018). For $0 < \gamma < 1/L_\psi$, φ_γ shares the same (local/strong) minima with \mathbb{P}_{in} (Stella et al., 2017). Function φ_γ is given by

$$\begin{aligned} \varphi_\gamma(u; c, y) &= \psi(u; c, y) - \frac{\gamma}{2} \|\nabla \psi(u; c, y)\|^2 \\ &\quad + \frac{1}{2\gamma} \text{dist}_U^2(u - \gamma \nabla \psi(u; c, y)). \end{aligned}$$

Note that $\varphi_\gamma(u; c, y)$ is computed at the cost of one projected gradient step on ψ .

PANOC is shown in Algorithm 2. PANOC combines *safe* projected gradient updates with *fast* quasi-Newtonian directions, which are computed with L-BFGS (see line 11 in Algorithm 2) based on the modification of L-BFGS proposed by Li and Fukushima (2001). With this choice of averaged directions and under mild conditions, eventually only fast quasi-Newtonian directions will be activated leading to very fast convergence (Stella et al., 2017).

PANOC uses the same oracle as the projected gradient method — it only requires ψ , $\nabla\psi$ and Π_U — and involves a very simple decrease criterion on φ_γ . The computation of the L-BFGS directions, d^ν , requires merely $4\mu n$, where μ is the L-BFGS memory.

Algorithm 2 PANOC method for solving the inner problem, $\mathbb{P}_{\text{in}}(c, y)$

Input: $u^0 \in \mathbb{R}^n$ (initial guess), $L_{\psi,0} > 0$ (estimate of the Lipschitz constant of $\nabla\psi(\cdot; c, y)$), μ (memory length of L-BFGS), $\epsilon > 0$ (tolerance), ν_{max} (max. iterations)

Output: Approximate solution u^*

- 1: $L \leftarrow L_{\psi,0}$, choose $\gamma \in (0, 1/L)$, $\sigma \in (0, \frac{\gamma}{2}(1 - \gamma L))$
 - 2: **for** $\nu = 0, 1, \dots, \nu_{\text{max}}$ **do**
 - 3: Compute $\nabla\psi(u^\nu; c, y)$ with automatic differentiation
 - 4: $u^{\nu+1/2} \leftarrow T_\gamma(u^\nu; c, y)$, $r^\nu \leftarrow u^\nu - u^{\nu+1/2}$
 - 5: **if** $\|\gamma^{-1}r^\nu + \nabla\psi(u^{\nu+1/2}; c, y) - \nabla\psi(u^\nu; c, y)\|_\infty < \epsilon$
 - 6: **exit** and return $u^* = u^{\nu+1/2}$
 - 7: **while** $\psi(u^{\nu+1/2}; c, y) > \psi(u^\nu; c, y) - \nabla\psi(u^\nu; c, y)^\top r^\nu + \frac{L}{2}\|r^\nu\|^2$ **do** (Update Lipschitz constant estimate)
 - 8: Empty the L-BFGS buffers
 - 9: $L \leftarrow 2L$, $\sigma \leftarrow \sigma/2$, $\gamma \leftarrow \gamma/2$
 - 10: $u^{\nu+1/2} \leftarrow T_\gamma(u^\nu; c, y)$
 - 11: $d^\nu \leftarrow -H_\nu r^\nu$ using L-BFGS
 - 12: $u^{\nu+1} \leftarrow u^\nu - (1 - \tau_\nu)r^\nu + \tau_\nu d^\nu$, where τ_ν is the largest number in $\{1/2^i : i \in \mathbb{N}\}$ such that
- $$\varphi_\gamma(u^{\nu+1}; c, y) \leq \varphi_\gamma(u^\nu; c, y) - \sigma\|\gamma^{-1}r^\nu\|^2$$

Note that Problem (1) can be written equivalently as

$$\underset{u \in U}{\text{minimize}} \quad f(u) + h(F(u)), \quad (16)$$

where $h(z_1, z_2) = \delta_C(z_1) + \delta_{\{0\}}(z_2)$, for $z_1 \in \mathbb{R}^{n_1}$ and $z_2 \in \mathbb{R}^{n_2}$, $F(u) = [F_1(u) \ F_2(u)]^\top$, and the associated Lagrangian is

$$L(u, y) = f(u) + \langle y, F(u) \rangle - h^*(y), \quad (17)$$

and the first order necessary conditions for optimality can be written as

$$0 \in \partial_x L(u, y), \quad 0 \in \partial_y [-L(u, y)]. \quad (18)$$

Provided that F_2 is differentiable, the optimality conditions read

$$0 \in \nabla f(u) + N_U(u) + JF(u)^\top y, \quad y \in \partial h(F(u)), \quad (19)$$

or

$$0 \in \nabla f(u) + N_U(u) + JF_1(u)^\top y + JF_2(u)^\top y_2, \quad (20a)$$

$$y \in \partial N_C(F(u)), \quad y_2 \in \mathbb{R}^{n_2}. \quad (20b)$$

It can be shown that if a pair (\bar{u}, \bar{y}) satisfies the termination conditions of Algorithms 1 and 2, it is an (ϵ, δ) -approximate KKT point of Problem (1) in the sense that

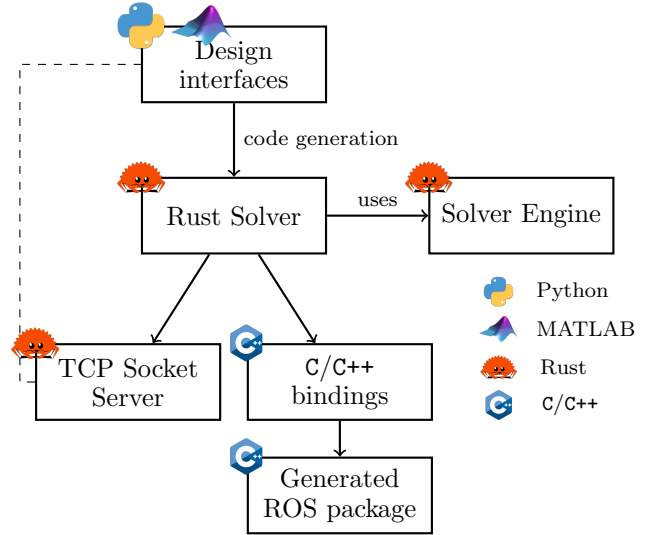


Fig. 1. Structure of the prototyping and embedded code generation framework of OpEn.

there exist $v \in \partial_u L(\bar{u}, \bar{y})$ and $w \in \partial_y [-L(\bar{u}, \bar{y})]$ with $\|v\| \leq \epsilon$ and $\|w\| \leq \delta$. In particular, note that δ controls the infeasibility of the approximate solution. The convergence properties of formulations without F_1 where F_2 is not differentiable are studied in (Hermans et al., 2019).

4. EMBEDDED CODE GENERATION

OpEn is an open-source software for embedded nonconvex optimization which involves several independent modules. The core numerical solver, which implements Algorithm 1, is implemented in Rust. Rust is a modern systems programming language with memory safety guarantees — an important feature in embedded applications.

OpEn provides a Python library (`opengen`) and a MATLAB toolbox for fast prototyping of embedded optimization applications; these can be used to generate Rust code for user-defined problems. OpEn can cross-compile the generated optimizer for several target systems (e.g., ARM Cortex-A processors). The generated optimizer can either be used in Rust or consumed through an auto-generated C or C++ interface. This allows the incorporation of the generated code in robot operating system (ROS) projects. OpEn can produce a very-low-latency TCP socket server, written in Rust, that uses a simple JSON data format. The TCP server can be invoked from Python, MATLAB and virtually any programming language, while it facilitates the deployment of edge intelligence applications and the solution of problems by a distributed network of agents.

In order to use OpEn in Python, the user only needs to install the Rust compiler, `clang` and `opengen` using `pip install opengen`. To use `opengen` in Python, one first needs to import it alongside CasADi (which is co-installed with `opengen`)

```

import opengen as og
import opengen.constraints as ogc
import casadi.casadi as cs
    
```

Let us give a simple example of using OpEn to generate a parametric optimizer. Consider the following problem of

constrained minimization of the Rosenbrock function:

$$\mathbb{P}(p) : \underset{u \in \mathbb{R}^5}{\text{minimize}} \sum_{i=1}^4 p_2 (u_{i+1} - u_i^2)^2 + (p_1 - u_i)^2 \quad (21a)$$

$$\text{subj. to: } \|u\| \leq 0.73 \quad (21b)$$

$$p_3 \sin(u_1) = \cos(u_2 + u_3) \quad (21c)$$

$$u_3 + u_4 \leq 0.2 \quad (21d)$$

with $p \in \mathbb{R}^3$. This is a nonconvex parametric optimization problem which can be written in multiple ways in the form of Problem (1). One possible choice is $U = \{u \in \mathbb{R}^5 : \|u\| \leq 0.73\}$ and $F_2 : \mathbb{R}^5 \times \mathbb{R}^3 \rightarrow \mathbb{R}^2$ is given by

$$F_2(u, p) = \begin{bmatrix} p_3 \sin(u_1) - \cos(u_2 + u_3) \\ [u_3 + u_4 - 0.2]_+ \end{bmatrix}. \quad (22)$$

In this case, both constraints will be handled with the penalty method.

```

u = cs.SX.sym("u", 5) # decision variable
p = cs.SX.sym("p", 3) # parameter
# Define cost function...
f = sum([p[1] * (u[i+1] - u[i]**2)**2
        + (p[0] - u[i])**2 for i in range(4)])
# Define constraints...
c1 = p[2]*cs.sin(u[0]) - cs.cos(u[1] + u[2])
c2 = cs.fmax(u[2] + u[3] - 0.2, 0)
c = cs.vertcat(c1, c2)
bounds = ogc.Ball2(radius=0.73)
# Formulate problem...
problem = og.builder.Problem(u, p, f) \
    .with_penalty_constraints(c) \
    .with_constraints(bounds)
# Configure the build and the solver...
build_cfg = og.config.BuildConfiguration() \
    .with_build_directory("build") \
    .with_tcp_interface_config() \
    .with_build_c_bindings()
meta = og.config.OptimizerMeta() \
    .with_optimizer_name("my_optimizer")
solver_cfg = og.config.SolverConfiguration() \
    .with_tolerance(1e-5) \
    .with_delta_tolerance(1e-4) \
    .with_initial_penalty(1e3)
builder = og.builder.OpEnOptimizerBuilder(
    problem, meta, build_cfg, solver_cfg)
builder.build() # Generate code
    
```

A second option is to handle both constraints with the augmented Lagrangian method by defining

$$F_1(u, p) = \begin{bmatrix} p_3 \sin(u_1) - \cos(u_2 + u_3) \\ u_3 + u_4 - 0.2 \end{bmatrix} \quad (23)$$

and $C = \{0\} \times (-\infty, 0]$. In that case, set Y is computed internally and is of the form $Y = [-M, M] \times [0, M]$, where $M = 10^{12}$ (however, the user can override this and provide her own set compact Y).

```

# Set C = {0} x (-inf, 0]
n1 = 2
Ca = ogc.Zero() # set {0}
Cb = ogc.Rectangle(None, [0]) # set (-inf, 0]
C = ogc.CartesianProduct(n1, [0, 1], [Ca, Cb])
# Custom set Y (Optional)...
M = 1e10
Y = ogc.Rectangle([-M, 0], [M, M])
# Problem specification...
problem = og.builder.Problem(u, p, f) \
    .with_aug_lagrangian_constraints(c, C, Y) \
    .with_constraints(bounds)
    
```

The solver will be generated and all relevant files will be stored in `build/my_optimizer`. The solver can be interfaced through its C/C++ interface (enabled with `.with_build_c_bindings`) using the auto-generated bindings `build/my_optimizer/my_optimizer_bindings.h`. It can also be invoked over a generated TCP socket interface which is enabled using `with_tcp_interface_config` (the default IP and port are 127.0.0.1 and 8333, but this can be configured).

For $p = (1, 50, 1.5)$, the formulation with the penalty method led to 7 outer iterations, 647 inner iterations in total and ran in 3.5ms. The formulation based on the augmented Lagrangian method ran in 1.4 ms after 5 outer and 175 total inner ones. IPOPT solves the same problem in 8.2ms and `scipy`'s SQP solver, SLSQP, which is based on (Kraft, 1988), in 15.3ms.

While the solver itself is written in Rust, the generation of code for the cost function, its gradient, and functions F_1 and F_2 is delegated to CasADi (Andersson et al., 2019) which generates C code. This optimized C code is compiled into the Rust binary using the Rust tools `bindgen` and `cc`, which gives a Rust interface to the generated code. The generation of both the Rust solver and the TCP socket server is facilitated by the *Jinja2 template language*. C/C++, as well as any language that supports the C/C++ application binary interface (ABI), can consume the generated optimizers. To that end, the Rust tool `cbindgen` is used to generate appropriate bindings. This allows the solver to be readily integrated with the robot operating system. These design choices allow the entire solver and library to run in Rust thus retaining its strong memory safety guarantees.

Moreover, the solver is thoroughly and continuously being tested in *Travis CI* (108 Rust tests, 35 Python tests) with 96.3% code coverage, alongside extensive documentation of both `OpEn`'s Rust code and `openen`'s Python code, and a large collection of examples available on the website at <http://doc.optimization-engine.xyz>.

5. SIMULATIONS

5.1 Obstacle Avoidance

The dynamics of an autonomous ground vehicle can be described by the following nonlinear bicycle model with four state and two input variables

$$\dot{p}_x = v \cos \psi, \dot{p}_y = v \sin \psi, \dot{\psi} = v/L \tan \delta, \dot{v} = \alpha(a - v),$$

where (p_x, p_y) are the position coordinates of the vehicle with respect to an Earth-fixed frame of reference, ψ is the vehicle's orientation and v is the longitudinal velocity, which follows a first-order dynamics with parameter $\alpha = 0.25$. The vehicle is controlled by its longitudinal acceleration, a , and the steering angle δ . We define the state vector $x = (p_x, p_y, \psi, v)$ and the input vector $u = (a, \delta)$. The length of the vehicle is $L = 0.5$ m. The above dynamical system is discretized with the Euler method with sampling time $T_s = 50$ ms. The vehicle must navigate from an initial position, pose and velocity, x , to a target position and pose, $x^{\text{ref}} = (p_x^{\text{ref}}, p_y^{\text{ref}}, \psi^{\text{ref}}, 0)$, while avoiding a cylindrical obstacle of radius $r = 0.65$ positioned at $(-3, 0.2)$.

We formulate an MPC problem with prediction horizon $N = 100$ and stage cost function $\ell(x, u) = 18(p_x^2 + p_y^2) + 2\psi^2 + 5v^2$ and the terminal cost function $\ell_N(x) = 1500(p_x^2 + p_y^2) + 500\psi^2 + 10v^2$. In order to avoid aggressive steering or jerks we introduce the additional stage cost function $\ell_\Delta(u_t, u_{t-1}) = 100(a_t - a_{t-1})^2 + 30(\delta_t - \delta_{t-1})^2$. Furthermore, we impose the actuation constraints $-1 \leq a \leq 2 \text{ m/s}^2$ and $|\delta| \leq 0.25 \text{ rad} \approx 14.32^\circ$. Then, the state sequence is eliminated as in Equation (5). The obstacle avoidance constraints can be formulated as follows

$$[r^2 - p_x^2 - p_y^2]_+ = 0, \quad (24)$$

which can be accommodated using the penalty method. Alternatively, in this case we can treat the obstacle constraints as inequality constraints of the form $r^2 - p_x^2 - p_y^2 \leq 0$ and describe them using F_1 . The two approaches lead to comparable results as shown in Figure 4.




Fig. 2. (Solid line) Path of NMPC-controlled ground vehicle. (Dashed line) Circular obstacle at $(x_c, y_c) = (-3, 0.2)$ of radius $r = 0.65 \text{ m}$.

The NMPC-controlled autonomous vehicle manages to avoid the obstacle, arrive at the target position and assume the desired orientation as shown in Figures 2 and 3.

OpEn outperforms significantly the interior point solver IPOPT and the SQP solver of Python's package `scipy`, SLSQP, as illustrated in Figure 4. All software make use of the same auto-generated code to evaluate the cost function, f , its gradient, ∇f , and function F_2 and the tolerance is set to $\delta = 10^{-3}$ and $\epsilon = 10^{-4}$. The selected solver parameters are $\epsilon_0 = 10^{-4}$, $\rho = 5$, $c_0 = 500$ and L-BFGS memory $\mu = 20$.

5.2 Constrained Nonlinear Estimation

Consider Lorenz's chaotic system




Fig. 3. Closed-loop trajectories for the NMPC-controlled vehicle.




Fig. 4. Computation time for the NMPC controlling the autonomous ground vehicle. (Solid lines) Execution times for OpEn, IPOPT and SLSQP, (Dashed line) sampling time.

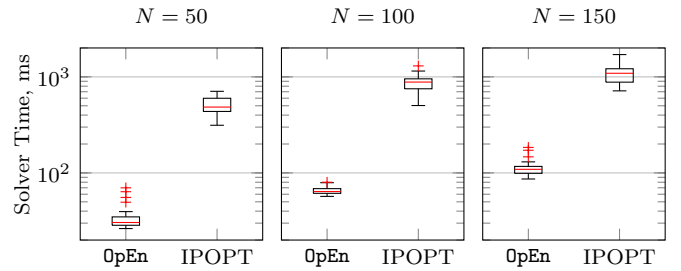


Fig. 5. Runtime of OpEn and IPOPT; box plots for 30 random trials.

$$\dot{x} = \begin{bmatrix} a_1(x_2 - x_1) \\ x_1(a_2 - x_3) - x_2 \\ x_1x_2 - a_3x_3 \end{bmatrix} + w(t), \quad (25)$$

where w is a disturbance signal, and parameters $a_1 = 10$, $a_2 = 14$, $a_3 = 8/3$, which we discretize by means of the fourth-order Runge-Kutta method with integration step $h = 0.1$ leading to a discrete-time system $x_{t+1} = \Phi(x_t) + w_t$. It is known that $-1 \leq w \leq 1$. The system output is given by

$$y = \begin{bmatrix} 2x_1 \\ x_2 + x_3 \end{bmatrix} + v, \quad (26)$$

$G(x)$

where v is a noise signal for which it is known that $-1.5 \leq v \leq 1.5$. We formulate the following constrained nonconvex estimation problem:

$$\mathbb{P}(\mathbf{y}) : \underset{\hat{x}, \hat{w}, \hat{v}}{\text{minimize}} \sum_{t=0}^{N-1} \|\hat{w}_t\|_Q^2 + \|\hat{v}_t\|_R^2 \quad (27a)$$

$$\text{subj. to: } w_{\min} \leq \hat{w}_t \leq w_{\max}, t \in \mathbb{N}_N \quad (27b)$$

$$v_{\min} \leq \hat{v}_t \leq v_{\max}, t \in \mathbb{N}_N \quad (27c)$$

$$\hat{x}_{t+1} = \Phi(\hat{x}_t) + \hat{w}_t, t \in \mathbb{N}_{N-1} \quad (27d)$$

$$y_t = G(\hat{x}_t) + \hat{v}_t, t \in \mathbb{N}_N \quad (27e)$$

where Q and R are symmetric positive definite matrices. This is a problem with decision variables $\hat{x} = (\hat{x}_t)_{t=0}^N$, $\hat{w} = (\hat{w}_t)_{t=0}^N$ and $\hat{v} = (\hat{v}_t)_{t=0}^N$, that is, $u = (\hat{x}, \hat{w}, \hat{v})$, and parameter $\mathbf{y} = (y_t)_{t=0}^N$. Constraints (27b) and (27c) define a rectangle, on which one can easily compute projections. The equality constraints (27d) and (27e) are treated with the augmented Lagrangian method by defining

$$F_1(u, \mathbf{y}) = \begin{bmatrix} (\hat{x}_{t+1} - \Phi(\hat{x}_t) - \hat{w}_t)_{t=0}^{N-1} \\ (y_t - G(\hat{x}_t) - \hat{v}_t)_{t=0}^N \end{bmatrix}, \quad (28)$$

and $C = \{0\}$.




Fig. 6. Constrained state estimation for Lorenz's system with $N = 100$

In Figure 5 we compare the runtime of **OpEn** with IPOPT for $N = 50$, $N = 100$ and $N = 150$. The SQP solver of **scipy** was significantly slower compared to both **OpEn** and IPOPT and was therefore omitted. Estimated trajectories for $N = 100$ are shown in Figure 6. The tuning parameters for **OpEn** were chosen to be $c_0 = 200$, $\rho = 1.8$, $\epsilon_0 = 0.1$, $\mu = 15$ and tolerances $\delta = 10^{-5}$ and $\epsilon = 10^{-4}$. **OpEn** performed no more than 7 outer iterations (in the majority of cases, fewer than 5) keeping the penalty parameter, c_ν , below 39672.

6. CONCLUSIONS AND FUTURE WORK

In this paper we presented a code generation software that can be easily used in Python and MATLAB and will empower engineers to build embedded optimization modules that will underpin complex control, estimation and signal processing developments. The proposed algorithm can handle nonconvex smooth cost functions and nonlinear constraints by means of the augmented Lagrangian and penalty methods, the latter being particularly suitable for obstacle/collision avoidance problems. Future work will focus on the study of the convergence properties of problems that involve constraints of the form (1b) and (1c) under weak assumptions.

REFERENCES

- J.A.E. Andersson, J. Gillis, G. Horn, J.B. Rawlings, and M. Diehl. CasADi: a software framework for nonlinear optimization and optimal control. *Math Prog Comp*, 11(1):1–36, 2019.
- G. Banjac, B. Stellato, N. Moehle, P. Goulart, A. Bemporad, and S. Boyd. Embedded code generation using the OSQP solver. In *IEEE CDC*, 2017.
- E.G. Birgin and J.M. Martínez. *Practical Augmented Lagrangian Methods for Constrained Optimization (Fundamentals of Algorithms)*. SIAM, 2014.
- H.J. Ferreau, S. Almér, R. Verschueren, M. Diehl, D. Frick, A. Domahidi, J.L. Jerez, G. Stathopoulos, and C. Jones. Embedded optimization methods for industrial automatic control. *IFAC WC*, 50(1):13194 – 13209, 2017.
- P.E. Gill, W. Murray, and M.A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, January 2005.
- B. Hermans, P. Patrinos, and G. Pipeleers. A penalty method based approach for autonomous navigation using nonlinear model predictive control. *IFAC WC*, 51(20):234 – 240, 2018.
- B. Hermans, G. Pipeleers, and P. Patrinos. A penalty method for nonlinear programs with set exclusion constraints. Submitted to *Automatica*, 2019.
- A.F. Izmailov, M.V. Solodov, and E.I. Uskov. Global convergence of augmented lagrangian methods applied to optimization problems with degenerate constraints, including problems with complementarity constraints. *SIAM J Optim*, 22(4):1579–1606, 2012.
- D. Kouzoupis, A. Zanelli, H. Peyrl, and H. J. Ferreau. Towards proper assessment of QP algorithms for embedded model predictive control. In *ECC*, pages 2609–16, 2015.
- D. Kraft. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, DLR German Aerospace Center – Institute for Flight Mechanics, Köln, Germany, 1988.
- D.-H. Li and M. Fukushima. On the global convergence of the BFGS method for nonconvex unconstrained optimization problems. *SIAM J Optim*, 11(4):1054–64, 2001.
- J. Mattingley and S. Boyd. CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, Mar 2012.
- J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer New York, 2006.
- A. Sathya, P. Sotasakis, R. Van Parys, A. Themelis, G. Pipeleers, and P. Patrinos. Embedded nonlinear model predictive control for obstacle avoidance using PANOC. In *ECC*, pages 1523–1528, 2018.
- H. Scheel and S. Scholtes. Mathematical programs with complementarity constraints: Stationarity, optimality, and sensitivity. *Math. Op. Res.*, 25(1):1–22, 2000.
- H.A. Shukla, B. Khusainov, E.C. Kerrigan, and C.N Jones. Software and hardware code generation for predictive control using splitting methods. *IFAC WC*, 50(1):14386 – 14391, 2017.
- E. Small, P. Sotasakis, E. Fresk, P. Patrinos, and G. Nikolakopoulos. Aerial navigation in obstructed environments with embedded nonlinear model predictive control. In *ECC*, pages 3556–3563, 2019.
- L. Stella, A. Themelis, P. Sotasakis, and P. Patrinos. A simple and efficient algorithm for nonlinear model predictive control. In *IEEE CDC*, pages 1939–44, 2017.
- A. Themelis, L. Stella, and P. Patrinos. Forward-backward envelope for the sum of two nonconvex functions: Further properties and nonmonotone linesearch algorithms. *SIAM J Optim*, 28(3):2274–2303, 2018.
- R. Verschueren, G. Frison, D. Kouzoupis, N. van Duijckeren, A. Zanelli, R. Quirynen, and M. Diehl. Towards a modular software package for embedded optimization. *IFAC WC*, 51(20):374 – 380, 2018.
- A. Wächter and L.T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math Prog*, 106(1):25–57, 2006.
- A. Zanelli, R. Quirynen, J. Jerez, and M. Diehl. A homotopy-based nonlinear interior-point method for NMPC. *IFAC WC*, 50(1):13188 – 13193, 2017.
- Ch. Zillober, K. Schittkowski, and K. Moritzen. Very large scale optimization by sequential convex programming. *Optim Meth Soft*, 19(1):103–120, February 2004.
- P. Zometa, M. Kögel, and R. Findeisen. μ AO-MPC: A free code generation tool for embedded real-time linear model predictive control. In *2013 IEEE ACC*, pages 5320–5325, June 2013.