



**QUEEN'S
UNIVERSITY
BELFAST**

Parallelisable computation of the gradient in nonlinear stochastic optimal control problems

Lin, J., Moran, R., & Sopasakis, P. (2023). Parallelisable computation of the gradient in nonlinear stochastic optimal control problems. In *2023 34th Irish Signals and Systems Conference (ISSC): Proceedings* (Irish Signals and Systems Conference (ISSC): Proceedings). Institute of Electrical and Electronics Engineers Inc.. <https://doi.org/10.1109/ISSC59246.2023.10162071>

Published in:

2023 34th Irish Signals and Systems Conference (ISSC): Proceedings

Document Version:

Peer reviewed version

Queen's University Belfast - Research Portal:

[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

Copyright 2023 IEEE.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access

This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

Parallelisable computation of the gradient in nonlinear stochastic optimal control problems

Jie Lin

School of EECS
Queen's University Belfast
Belfast, Northern Ireland, UK
jlin20@qub.ac.uk

Ruairi Moran

School of EECS & i-AMS
Queen's University Belfast
Belfast, Northern Ireland, UK
rmoran05@qub.ac.uk

Pantelis Sopasakis

School of EECS & i-AMS
Queen's University Belfast
Belfast, Northern Ireland, UK
p.sopasakis@qub.ac.uk

Abstract—Nonlinear (deterministic and stochastic) optimal control problems are often solved on embedded devices using first-order numerical optimisation methods. The gradient computation accounts for a significant part of the computation cost per iteration; this is often performed with reverse-mode automatic differentiation and software libraries such as CasADi. CasADi can be used to generate C code for this computation. In this paper, we propose a simple *ad hoc* and highly parallelisable algorithm for the computation of the gradient of the total cost for deterministic and stochastic scenario-based optimal control problems. We also present `gradgen`: an open-source Python package that generates Rust code for the gradient computation. The proposed method leads to a faster performance compared to CasADi and a significant reduction in generated code.

Index Terms—Automatic differentiation, optimal control, scenario-based stochastic optimal control

I. INTRODUCTION

A. Background and motivation

Gradient-based methods — such as line search methods and various quasi-Newtonian methods [1] — are popular for the solution of nonlinear optimal control problems (OCPs), especially on embedded devices since the involved steps are simple (typically no linear systems need to be solved), and the memory requirements are low. Then, the computation of the gradient of the cost function accounts for a major part of the computational cost at each iteration of the method.

The gradient can be computed numerically, but that would introduce errors that could compromise the convergence of the numerical optimisation method. Gradients can also be computed using symbolic algebra methods, but symbolic differentiation often leads to overly lengthy expressions, which are inefficient for applications of numerical optimisation.

Automatic differentiation (AD) is a class of methods for the algorithmic computation of the gradient [2]. There are two main AD approaches: forward-mode — which is, in essence, equivalent to symbolic differentiation [3] — and reverse-mode AD. In forward-mode AD, the computational graph of a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is traversed forwards, from the inputs x_1, \dots, x_n to the output $f(x)$, while the partial gradients $\partial f / \partial x_i, i = 1, \dots, n$, are computed. In reverse-mode AD, the

computation that defines the function is traversed backwards, from $f(x)$ to x , while the chain rule is applied. Reverse-mode AD is more efficient compared to forward-mode [4].

Several open-source AD implementations are available, such as JAX [5] and within TensorFlow [6]. CasADi [7] is a Python package that can generate library-free C code that has been widely used in embedded applications such as autonomous ground vehicles [8], [9], aerial vehicles [10], [11], and quadruped robots [12] to name a few. CasADi is also used to compute gradients in numerical optimisation software such as OpEn [13] and acados [14]. However, the C code auto-generated by CasADi can often be tens to hundreds of thousands of lines of code long.

The efficient computation of the cost gradient becomes even more important in problems of very large dimension such as stochastic scenario-based optimal control where the number of control actions increases exponentially with the prediction horizon [15]. The structure of scenario-based problems has been exploited to devise parallelisable numerical algorithms [15], [16] for convex formulations, but to the best of the authors' knowledge such developments have not been generalised to nonlinear formulations.

In this paper, we propose reverse-mode-type AD algorithms in the deterministic and stochastic cases, that allow fast and low-memory computation of the gradient of an OCP's cost function with respect to the control actions. We show that the complexity of the proposed algorithms scales gracefully with the problem size. In the stochastic case, the proposed algorithm can be massively parallelised. We also present an open-source code generator that generates fast memory-safe Rust code [17] that is suitable for embedded applications.

B. Notation

Let $\mathbb{N}_{[k_1, k_2]}$ denote the integers in $[k_1, k_2]$. We denote the transpose of a matrix A by A^\top . Let $g : X \rightarrow Y$ and $f : Y \rightarrow Z$. Then we define the composition $f \circ g : X \rightarrow Z$ with $(f \circ g)(x) = f(g(x))$. For $x \in \mathbb{R}^n$, note that $\|x\|$ denotes the Euclidean norm, that is, $\|x\| = \sqrt{x^\top x}$.

C. The Jacobian matrix and the chain rule

Consider a function $f : \mathbb{R}^n \ni x \mapsto (f_1(x), \dots, f_m(x)) \in \mathbb{R}^m$. The Jacobian matrix of a function f is defined as the

function $Jf : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times n}$ with

$$Jf(x) = \begin{bmatrix} \frac{\partial f_1(x)}{\partial x_1} & \dots & \frac{\partial f_1(x)}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m(x)}{\partial x_1} & \dots & \frac{\partial f_m(x)}{\partial x_n} \end{bmatrix}, \quad (1)$$

provided that all partial derivatives exist. For functions $f : \mathbb{R}^n \rightarrow \mathbb{R}$ the gradient of f is defined as the function $\nabla f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined by $\nabla f(x) = Jf(x)^\top$.

Suppose that $f : \mathbb{R}^p \rightarrow \mathbb{R}^m$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^p$ are differentiable. Then, the Jacobian matrix of the composition $f \circ g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is given by

$$J(f \circ g)(x) = Jf(g(x))Jg(x). \quad (2)$$

This is the chain rule for Jacobian matrices. The notation $Jf(g(x))$ means the Jacobian of f evaluated at $g(x)$.

Let $h_1 : \mathbb{R}^n \ni x \mapsto s_1 \in \mathbb{R}^{n_1}$, and $h_2 : \mathbb{R}^n \ni x \mapsto s_2 \in \mathbb{R}^{n_2}$. Define $g : \mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \ni (s_1, s_2) \mapsto g(s_1, s_2) \in \mathbb{R}^m$. Then, g is a function composed of two functions, and so we define $f(x) = g \circ (h_1(x), h_2(x))$. The Jacobian of f at x is

$$Jf(x) = J_{s_1}g(h_1(x), h_2(x))Jh_1(x) + J_{s_2}g(h_1(x), h_2(x))Jh_2(x). \quad (3)$$

II. DETERMINISTIC OCPs

Consider a — generally nonlinear — OCP of the general form

$$\mathbb{P}_N(x) : \underset{\substack{u_0, \dots, u_{N-1} \\ x_0, \dots, x_N}}{\text{Minimise}} \sum_{t=0}^{N-1} \ell(x_t, u_t) + V_f(x_N), \quad (4a)$$

$$\text{subject to: } x_{t+1} = f(x_t, u_t), t \in \mathbb{N}_{[0, N-1]}, \quad (4b)$$

$$x_0 = x, \quad (4c)$$

where $x_t \in \mathbb{R}^{n_x}$ is the system state, $u_t \in \mathbb{R}^{n_u}$ is the input, and the system dynamics $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, the stage cost function $\ell : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \rightarrow \mathbb{R}^{n_x}$, and the terminal cost function $V_f : \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$ are differentiable.

We can cast this problem as an unconstrained optimisation problem by eliminating the sequence of states. This is referred to as the single shooting formulation [7]. We define $u = (u_0, u_1, \dots, u_{N-1}) \in \mathbb{R}^{Nn_u}$, $F_0(x, u) = x$, and recursively

$$F_{t+1}(x, u) = f(F_t(x, u), u_t). \quad (5)$$

We see that $x_t = F_t(x, u)$ is the state at time t starting from the initial state x upon the action of the sequence u . This allows us to define the total cost function $V_N : \mathbb{R}^{n_x} \times \mathbb{R}^{Nn_u} \rightarrow \mathbb{R}$ as

$$V_N(x, u) = \sum_{t=0}^{N-1} \ell(F_t(x, u), u_t) + V_f(F_N(x, u)), \quad (6)$$

and \mathbb{P}_N can be written as

$$\mathbb{P}_N(x) : \underset{u \in \mathbb{R}^{Nn_u}}{\text{Minimise}} V_N(x, u). \quad (7)$$

In embedded optimisation we typically aim at determining a stationary point of this problem, that is, a sequence of

control actions $u^*(x)$ such that $\nabla V_N(x, u^*(x)) = 0$. To that end, any first-order optimisation method can be used; e.g., a line search gradient method, a quasi-Newtonian method [1], or the bespoke PANOC method for OCPs [18]. In all these methods, the most computationally demanding part is typically the gradient of the total cost function.

A straightforward application of the chain rule for the computation of the gradient of V_N is possible, but would necessitate the computation and storage of several matrices as well as several matrix-matrix/vector products. Instead, we will exploit the structure of the problem and will work in a fashion similar to the reverse-mode AD method [4].

First, let us introduce some convenient notation. Given a vector $d \in \mathbb{R}^{n_x}$ we define $f_t^x(d) = J_x f(x_t, u_t)^\top d$ and $f_t^u(d) = J_u f(x_t, u_t)^\top d$, for $t \in \mathbb{N}_{[0, N-1]}$, where $x_t = F_t(x, u)$. We also define $\ell_t^x = \nabla_x \ell(x_t, u_t)$ and $\ell_t^u = \nabla_u \ell(x_t, u_t)$, for $t \in \mathbb{N}_{[0, N-1]}$, and $V_f' = \nabla V_f(x_N)$.

We start by determining the gradient of V_N with respect to u_{N-1} , which, by (3), is

$$\begin{aligned} \nabla_{u_{N-1}} V_N(x, u) &= \nabla_{u_{N-1}} [\ell(x_{N-1}, u_{N-1}) + V_f(f(x_{N-1}, u_{N-1}))] \\ &= \ell_{N-1}^u + f_{N-1}^u(V_f'). \end{aligned} \quad (8)$$

Likewise, we can determine $\nabla_{u_{N-2}} V_N$. It is

$$\nabla_{u_{N-2}} V_N(x, u) = \ell_{N-2}^u + f_{N-2}^u(\ell_{N-1}^x + f_{N-1}^x(V_f')). \quad (9)$$

Note, however, that V_f' has already been computed. Moving on to stage $N-3$, we have

$$\begin{aligned} \nabla_{u_{N-3}} V_N(x, u) &= \ell_{N-3}^u + f_{N-3}^u(\ell_{N-2}^x + f_{N-2}^x(\ell_{N-1}^x + f_{N-1}^x(V_f'))). \end{aligned} \quad (10)$$

Note that the quantity $\ell_{N-1}^x + f_{N-1}^x(V_f')$ has already been computed in (9). A pattern is emerging, which we state in the following proposition.

Proposition 1: Let $a_{N-1} = V_f'$. Then,

$$\nabla_{u_t} V_N(x, u) = \ell_t^u + f_t^u(a_t), \quad (11a)$$

$$a_{t-1} = \ell_t^x + f_t^x(a_t), \quad (11b)$$

for $t = N-1, N-2, \dots, 0$.

This leads to Alg. 1 for the computation of the gradient of the total cost function, which is a special case of [18, Alg. 1].

The variables a_t in Alg. 1 correspond to the adjoint variables in reverse-mode AD. Note that Alg. 1 requires exactly N invocations of f and the Jacobians of f and ℓ with respect to x and u , that is, it scales linearly with N .

III. STOCHASTIC SCENARIO-BASED OCPs

In this section, we turn our attention to stochastic scenario-based OCPs [19]. We consider a dynamical system of the form

$$x_{t+1} = f(x_t, u_t, w_t), \quad (12)$$

for $t \in \mathbb{N}_{[0, N-1]}$, where w_t is a finitely-supported random process acting on the system as a disturbance. At every time t , having measured x_t , we apply a control action $u_t = u_t(x_t)$.

Algorithm 1 Computation of the gradient of V_N with respect to u in the deterministic case

Input: Initial state, x , and sequence of control actions, u

Output: $\nabla V_N(x, u)$

- 1: **for** $t = 0$ to $N - 1$ **do**
- 2: $x_{t+1} \leftarrow f(x_t, u_t)$
- 3: $a_{N-1} \leftarrow V'_f$
- 4: **for** $t = N - 1$ to 1 **do**
- 5: $\nabla_{u_t} V_N(x, u) \leftarrow \ell_t^u + f_t^u(a_t)$
- 6: $a_{t-1} \leftarrow \ell_t^x + f_t^x(a_t)$
- 7: $\nabla_{u_0} V_N(x, u) \leftarrow \ell_0^u + f_0^u(a_0)$
- 8: **return** $\nabla V_N(x, u) = (\nabla_{u_t} V_N(x, u))_{t=0, \dots, N-1}$

This entails a cost $\ell(x_t, u_t, w_t)$, which is a random variable. At the end of the horizon — at $t = N$ — a terminal cost, $V_f(x_N)$, is incurred.

The evolution of the system states and costs can be captured by the structure of a scenario tree as shown in Fig. 1. The possible realisations of the state of the system correspond to the nodes of the scenario tree, while the edges correspond to the realisations of the disturbance, w_t .




Fig. 1. Scenario tree with $N = 2$. Regarding the system dynamics, note that $x^1 = f(x^0, u^0, w^1)$ and $x^2 = f(x^0, u^0, w^2)$. Likewise, for the stage costs note that $\ell^2 = \ell(x^0, u^0, w^2)$. For the terminal costs, note that $V_f^3 = V_f(x^3)$.

The scenarios are enumerated with an index i , where $i = 0$ corresponds to the unique node at stage $t = 0$ which is referred to as the *root* node. The set of nodes at stage $t \in \mathbb{N}_{[0, N]}$ are denoted by $\mathbf{nodes}(t)$, where $\mathbf{nodes}(0) = \{0\}$. The nodes at stage $t = N$ are referred to as the *leaf* nodes of the tree, and the nodes that are not at this stage are referred to as the *nonleaf* nodes. We also denote $\mathbf{nodes}(t_1, t_2) = \bigcup_{\tau=t_1}^{t_2} \mathbf{nodes}(\tau)$. Each node $i \in \mathbf{nodes}(1, N)$ is connected to a single node at stage $t - 1$ known as the ancestor of i , denoted by $\mathbf{anc}(i)$. Conversely, every node $i \in \mathbf{nodes}(0, N - 1)$ at a stage t leads up to a set of nodes $\mathbf{ch}(i) \subseteq \mathbf{nodes}(t + 1)$. The probability

of occurrence of node i is denoted by $\pi^i > 0$. Lastly, every node of the tree is associated with a state x^i , every nonleaf node with an input u^i , and every edge — so, all but the root node — with a disturbance w^i . Note that w^i is the event that happens at node $\mathbf{anc}(i)$ that causes node i to happen.

The dynamics (12) on the scenario tree can be written

$$x^{i+} = f(x^i, u^i, w^{i+}), \quad (13)$$

for $i \in \mathbf{nodes}(0, N - 1)$ and $i_+ \in \mathbf{ch}(i)$. Equivalently,

$$x^i = f(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i), \quad (14)$$

for $i \in \mathbf{nodes}(1, N)$.

We can now formulate the stochastic scenario-based OCP

$$\begin{aligned} \text{Minimise} \quad & \sum_{\substack{(u^i)_{i \in \mathbf{nodes}(0, N-1)}, \\ (x^i)_{i \in \mathbf{nodes}(0, N)}}} \sum_{i \in \mathbf{nodes}(1, N-1)} \pi^i \ell(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i) \\ & + \sum_{j \in \mathbf{nodes}(N)} \pi^j V_f(x^j), \end{aligned} \quad (15a)$$

$$\text{subject to: } x^i = f(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i), \quad (15b)$$

$$x^0 = x. \quad (15c)$$

Hereafter we assume that f and ℓ are differentiable in (x, u) and V_f is differentiable.

Let $u = (u^i)_{i \in \mathbf{nodes}(0, N-1)}$. We define $V_N(x, u)$ to be the total cost function in the above problem. We will determine the gradient of V_N with respect to all u^i . At this point, let us introduce the notation $\ell_u^i = \nabla_u \ell(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i)$ and $\ell_x^i = \nabla_x \ell(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i)$, for $i \in \mathbf{nodes}(1, N)$. Moreover, for $d \in \mathbb{R}^{n_x}$ we define $f_x^i(d) = J_x f(x^{\mathbf{anc}(i)}, u^{\mathbf{anc}(i)}, w^i) \top d$ and similarly, $f_u^i(d)$, for $i \in \mathbf{nodes}(1, N)$. Lastly, we define $V_f^i = \nabla V_f(x^i)$, for $i \in \mathbf{nodes}(N)$.

Firstly, for all $k_0 \in \mathbf{nodes}(N)$ we compute $a^{k_0} = V_f^{k_0}$. Then for $k_1 \in \mathbf{nodes}(N - 1)$, we compute

$$\nabla_{u^{k_1}} V_N = \sum_{k_0 \in \mathbf{ch}(k_1)} \pi^{k_0} (\ell_u^{k_0} + f_u^{k_0}(\underbrace{V_f^{k_0}}_{a^{k_0}})), \quad (16a)$$

and

$$a^{k_1} = \sum_{k_0 \in \mathbf{ch}(k_1)} \pi^{k_0} (\ell_x^{k_0} + f_x^{k_0}(a^{k_0})). \quad (16b)$$

Note that a^{k_1} is similar to the gradient in (16a), except with respect to x instead of u . Next, let us take $k_2 \in \mathbf{nodes}(N - 2)$.

We compute

$$\begin{aligned}
\nabla_{u^{k_2}} V_N &= \sum_{k_1 \in \text{ch}(k_2)} \pi^{k_1} \ell_u^{k_1} \\
&\quad + \sum_{\substack{k_1 \in \text{ch}(k_2) \\ k_0 \in \text{ch}(k_1)}} \pi^{k_0} (f_u^{k_1}(\ell_x^{k_0}) + f_u^{k_1}(f_x^{k_0}(V_f^{k_0}))) \\
&= \sum_{k_1 \in \text{ch}(k_2)} \pi^{k_1} \ell_u^{k_1} \\
&\quad + \sum_{k_1 \in \text{ch}(k_2)} \underbrace{f_u^{k_1} \left(\sum_{k_0 \in \text{ch}(k_1)} \pi^{k_0} (\ell_x^{k_0} + f_x^{k_0}(V_f^{k_0})) \right)}_{a^{k_1}} \\
&= \sum_{k_1 \in \text{ch}(k_2)} \pi^{k_1} \ell_u^{k_1} + f_u^{k_1}(a^{k_1}), \tag{17}
\end{aligned}$$

and

$$a^{k_2} = \sum_{k_1 \in \text{ch}(k_2)} \pi^{k_1} \ell_x^{k_1} + f_x^{k_1}(a^{k_1}). \tag{18}$$

Again, note that a^{k_2} is similar to the gradient in (17), except with respect to x instead of u . For $\nabla_{u^i} V_N$ with $i \in \text{nodes}(N-1)$, we use (16). Then, moving backwards, we can show recursively that for $i \in \text{nodes}(0, t)$ with $t \in \mathbb{N}_{[0, N-2]}$

$$\nabla_{u^i} V_N = \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} \ell_{u^+}^{i_+} + f_{u^+}^{i_+}(a^{i_+}), \tag{19a}$$

where the adjoint vectors a^i , for $i \in \text{nodes}(1, N-2)$ are given by

$$a^i = \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} \ell_x^{i_+} + f_x^{i_+}(a^{i_+}). \tag{19b}$$

Many of the involved computations can be parallelised. This is evident in Alg. 2 where the parallelisation is represented using the fork-join model [20, Chapter 2]. In fact, the computation of the gradient of the total cost lends itself to a lockstep implementation on a GPU or FPGA.

Note that the complexity of Alg. 2 is linear with respect to the total number of nodes of the tree.

IV. GRADGEN: CODE GENERATION FOR EMBEDDED COMPUTATIONS

The open-source Python package `gradgen`¹ is distributed under the MIT License and can be easily installed with `pip`.

The user can then import `gradgen` to their Python project, where they specify the system dynamics and cost function in symbolic form. The Jacobians ℓ^x , ℓ^u , f^x , f^u , and V_f' are determined using `CasADi`. Then using `Jinja2` templates, `gradgen` generates Rust code that implements Alg. 1 or 2. The modular structure of the generated code is presented in Fig. 2. The top layer is the generated `gradgen` API, comprising Rust implementations of Alg. 1 and 2.

¹The source code is available at <https://github.com/QUB-ASL/gradgen>.

Algorithm 2 Parallelisable computation of the gradient of V_N with respect to u in the stochastic case

Input: Initial state, x , and sequence of control actions, u

Output: $\nabla V_N(x, u)$

```

1: for  $i \in \text{nodes}(1, N)$  do
2:    $x^i \leftarrow f(x^{\text{anc}(i)}, u^{\text{anc}(i)}, w^i)$ 
3: for  $i \in \text{nodes}(N)$  (in parallel) do
4:   fork  $a^i \leftarrow V_f^{i_i}$ 
5: join
6: for  $i \in \text{nodes}(N-1)$  (in parallel) do
7:   fork:
8:      $\nabla_{u^i} V_N \leftarrow \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} (\ell_{u^+}^{i_+} + f_{u^+}^{i_+} a^{i_+})$ 
9:      $a^i \leftarrow \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} (\ell_x^{i_+} + f_x^{i_+} a^{i_+})$ 
10:  join
11: for  $t = N-2$  to 1 do
12:   for  $i \in \text{nodes}(t)$  (in parallel) do
13:     fork:
14:        $\nabla_{u^i} V_N \leftarrow \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} \ell_{u^+}^{i_+} + f_{u^+}^{i_+} a^{i_+}$ 
15:        $a^i \leftarrow \sum_{i_+ \in \text{ch}(i)} \pi^{i_+} \ell_x^{i_+} + f_x^{i_+} a^{i_+}$ 
16:     join
17:  $\nabla_{u^0} V_N \leftarrow \sum_{i_+ \in \text{ch}(0)} \pi^{i_+} \ell_{u^+}^{i_+} + f_{u^+}^{i_+} a^{i_+}$ 
18: return  $\nabla V_N(x, u) = (\nabla_{u^i} V_N(x, u))_{i \in \text{nodes}(0, N-1)}$ 

```

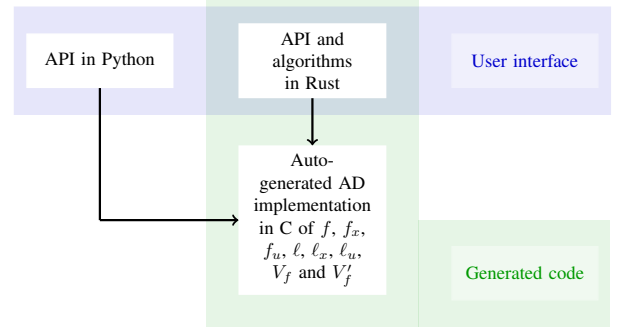


Fig. 2. Multi-layered structure of the generated code and Python API.

V. RESULTS

Here we demonstrate how `gradgen` performs on two systems: an inverted pendulum (two states, one input) and a ball-and-beam (four states, one input), in both deterministic and stochastic cases. All simulations are run on an Apple Macbook Pro with a 2 GHz quad-core Intel core i5 processor and 16 GB of RAM. Below, the system state is denoted x and the control action u . For both examples, the stage cost is $\ell(x, u) = \|x\|^2 + \|u\|^2$, and the terminal cost is $V_f(x) = 10\|x\|^2$.

A. System I: inverted pendulum

Consider the inverted pendulum system in Fig. 3. The car can move along the x axis, while the rod is free to rotate about the fulcrum, A , which is fixed to the car. The deviation angle, θ , is the angle the rod makes with the y axis. The car moves under the horizontal force, F , in an attempt to keep the rod upright, that is, $\theta = 0$.

The car has a mass M . The rod has a mass m and half-length L and the rod is isotropic. The friction of rotation of the rod is negligible.




Fig. 3. System I: inverted pendulum ($n_x = 2, n_u = 1$). The mid-point of the rod is denoted by B . The force applied to the car is denoted by F . The deviation angle, θ , is the angle the rod makes with the y axis.

The dynamics of the inverted pendulum is [21]

$$\dot{\theta} = \omega, \quad (20a)$$

$$\dot{\omega} = -3 \frac{\frac{1}{2}mL\omega^2 \sin(2\theta) + F \cos \theta - M_{\text{tot}}g \sin \theta}{(4M_{\text{tot}} - 3m \cos^2 \theta)L}. \quad (20b)$$

where $M_{\text{tot}} = M + m$.

We discretise this system using the Euler method with sampling time T_s to obtain a discrete-time system in the form of (4b). We use the values $m = 1$ kg, $M = 3$ kg, $g = 9.81$ m/s², and $T_s = 0.01$ s.

B. System II: ball-and-beam

Consider a ball of mass m is placed on a beam which is poised on a fulcrum at its middle as shown in Fig. 4. We can control the system by applying a torque u with respect to the fulcrum point. The moment of inertia of the beam is denoted by I . The displacement x of the ball from the midpoint can be measured with an optical sensor. The dynamical system is described by the following nonlinear differential equations

$$\frac{7}{5}\ddot{x} + g \sin \theta - x\dot{\theta}^2 = 0, \quad (21a)$$

$$(mx^2 + I)\ddot{\theta} + 2mx\dot{x}\dot{\theta} + mgx \cos \theta = u. \quad (21b)$$

This is a continuous-time system with state vector $(x, \dot{x}, \theta, \dot{\theta})$ and input u , which can be discretised with the Euler method and sampling period $T_s = 0.01$ s.

C. Stochastic systems

Here we consider stochastic variants of the above systems by considering the case of random communication delays that render the sampling time, T_s , a random variable [22]. We assume that T_s can take the values 10 ms, 20 ms, and 100 ms, which are governed by the Markov chain shown in Fig. 5 with uniform initial distribution.




Fig. 4. System II: ball and beam ($n_x = 4, n_u = 1$). The position of the ball is denoted by x , the angle of the beam is θ . The manipulated input is the torque u .




Fig. 5. Markov chain of random sampling time. The sampling time $T_s = 10$ ms corresponds to normal operation, whereas 20 ms and 100 ms correspond to random time delays.

D. Benchmark results

In Figures 6 and 7 we present the average execution time for the computation of the gradient using `gradgen`'s implementation of Alg. 1 and `CasADi` (in C) for the inverted pendulum and the ball-and-beam systems. In both figures we can observe that the execution time of `gradgen` increases linearly with the prediction horizon. Indeed, the complexity of Alg. 1 is $\mathcal{O}(N)$. Instead, the execution time of `CasADi` seems to be growing superlinearly with the prediction horizon.




Fig. 6. Inverted pendulum example: average computation time for the gradient and comparison with `CasADi` (based on 10^6 runs).

It should be noted that `CasADi` generates large C code files. For instance, for the ball-and-beam system with $N = 60$, the `CasADi` auto-generated C file counts 46k single lines of




Fig. 7. Ball-and-beam example: average computation time for the gradient and comparison with CasADi (based on 10^6 runs).

code (SLOCs). At $N = 100$, it generates 125k SLOCs and the compilation time exceeds 20mins. On the other hand, the total SLOCs of gradgen (including header files) is 1188, which is independent of the prediction horizon. For the inverted pendulum model, the auto-generated SLOCs by CasADi at $N = 50$ and $N = 100$ are 12.5k and 44.9k, respectively.

In the stochastic case (Sec. V-C), we compute the gradient of the total cost function, ∇V_N , for $N = 3, \dots, 7$, using gradgen’s implementation of Alg. 2. The computation times are shown in Fig. 8 where we see a linear increase with the number of nodes. Fig. 9 shows the number of SLOCs of CasADi’s auto-generated C file.




Fig. 8. Execution time for the computation of the gradient of the stochastic optimal control problem against the number of nodes.




Fig. 9. Ball-and-beam stochastic system: lines of code for the computation of the gradient of the stochastic optimal control problem against the number of nodes and comparison with CasADi.

VI. CONCLUSIONS AND FUTURE WORK

In this paper we derived an algorithm for computing the gradient of a cost function with respect to a sequence of

control actions in the deterministic case. We also derived a parallelisable algorithm for computing the gradient in the stochastic case. We simulated two systems that demonstrate the implementations of the deterministic and stochastic algorithms in the new open-source package gradgen, which is available at <https://github.com/QUB-ASL/gradgen>. The deterministic algorithm is benchmarked against CasADi. The simulation results indicate that gradgen can outperform CasADi, especially at large prediction horizons, gradgen generates much smaller files than CasADi, and that the run time of gradgen scales well with the problem size. Future work will focus on implementing the stochastic algorithm in parallel on a GPU.

REFERENCES

- [1] J. Nocedal and S. Wright, *Numerical Optimization*. Springer, 2006.
- [2] C. C. Margossian, “A review of automatic differentiation and its efficient implementation,” *WIREs Data Min Knowl Discov*, vol. 9, Mar. 2019.
- [3] S. Laue, “On the equivalence of automatic and symbolic differentiation,” *arXiv*, no. 1904.02990, 2019.
- [4] M. B. Giles, “Collected matrix derivative results for forward and reverse mode algorithmic differentiation,” in *Advances in Automatic Differentiation*, pp. 35–44, Springer Berlin Heidelberg, 2008.
- [5] J. Bradbury, R. Frostig *et al.*, “JAX: composable transformations of Python+NumPy programs,” 2018.
- [6] M. Abadi, A. Agarwal *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015. Available from tensorflow.org.
- [7] 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. Program. Comput.*, vol. 11, no. 1, pp. 1–36, 2019.
- [8] 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*, pp. 1523–1528, June 2018.
- [9] V. Cataffo, G. Silano, L. Iannelli, V. Puig, and L. Glielmo, “A nonlinear model predictive control strategy for autonomous racing of scale vehicles,” in *SMC*, pp. 100–105, 2022.
- [10] S. S. Mansouri, C. Kanellakis, E. Fresk, B. Lindqvist, D. Kominiak, A. Koval, P. Sotasakis, and G. Nikolakopoulos, “Subterranean MAV navigation based on nonlinear MPC with collision avoidance constraints,” in *IFAC WC*, (Germany), 2020.
- [11] E. Small, P. Sotasakis, E. Fresk, P. Patrinos, and G. Nikolakopoulos, “Aerial navigation in obstructed environments with embedded nonlinear model predictive control,” in *ECC*, pp. 3556–3563, June 2019.
- [12] Q. Nguyen, M. J. Powell, B. Katz, J. D. Carlo, and S. Kim, “Optimized jumping on the MIT Cheetah 3 robot,” in *ICRA*, 2019.
- [13] P. Sotasakis, E. Fresk, and P. Patrinos, “OpEn: Code generation for embedded nonconvex optimization,” in *IFAC WC*, (Germany), 2020.
- [14] R. Verschuere, G. Frison *et al.*, “acados—a modular open-source framework for fast embedded optimal control,” *Math. Program. Comput.*, vol. 14, pp. 147–183, Oct. 2021.
- [15] A. K. Sampathirao, P. Patrinos, A. Bemporad, and P. Sotasakis, “Massively parallelizable proximal algorithms for large-scale stochastic optimal control problems,” *arXiv*, no. 2107.01745, 2021.
- [16] A. K. Sampathirao, P. Sotasakis, A. Bemporad, and P. Patrinos, “GPU-accelerated stochastic predictive control of drinking water networks,” *IEEE Trans Control Syst Technol*, vol. 26, no. 2, pp. 551–562, 2018.
- [17] N. D. Matsakis and F. S. Klock II, “The rust language,” in *ACM SIGAda Ada Letters*, vol. 34, pp. 103–104, ACM, 2014.
- [18] L. Stella, A. Themelis, P. Sotasakis, and P. Patrinos, “A simple and efficient algorithm for nonlinear model predictive control,” in *CDC*, pp. 1939–1944, Dec 2017.
- [19] A. Shapiro, D. Dentcheva, and Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, second ed., 2014.
- [20] P. Pacheco and M. Malensek, *Introduction to Parallel Programming*. Morgan Kaufman, 2011.
- [21] P. Sotasakis, *Control systems: an introduction*. Applied Mathematics Press, 2023.
- [22] P. Patrinos, P. Sotasakis, and H. Sarimveis, “Stochastic model predictive control for constrained networked control systems with random time delay,” *IFAC Proc. Volumes*, vol. 44, no. 1, pp. 12626–12631, 2011.