

# Ungar – A C++ Framework for Real-Time Optimal Control Using Template Metaprogramming

Flavio De Vincenti and Stelian Coros

**Abstract**—We present Ungar, an open-source library to aid the implementation of high-dimensional optimal control problems (OCPs). We adopt modern template metaprogramming techniques to enable the compile-time modeling of complex systems while retaining maximum runtime efficiency. Our framework provides syntactic sugar to allow for expressive formulations of a rich set of structured dynamical systems. While the core modules depend only on the header-only Eigen and Boost.Hana libraries, we bundle our codebase with optional packages and custom wrappers for automatic differentiation, code generation, and nonlinear programming. Finally, we demonstrate the versatility of Ungar in various model predictive control applications, namely, four-legged locomotion and collaborative loco-manipulation with multiple one-armed quadruped robots. Ungar is available under the Apache License 2.0 at <https://github.com/fdevinc/ungar>.

## I. INTRODUCTION

The advancements in model predictive control (MPC) methods have endowed robots with exceptional athletic skills. Recent displays of humanoid [8] and quadruped robots [4, 11] have shown feats that were once prerogative to science fiction. However, significant engineering efforts are still necessary to make such practical MPC implementations possible. The solution of big nonlinear programming (NLP) problems at real-time rates clashes with the inherent high dimensionality and fast dynamics of mechanical systems. These conflicting aspects translate to onerous computational costs to be met in fractions of seconds, thus calling for complex data structures and ingenious software designs.

We seek to facilitate the manual endeavors flowing into the development of MPC controllers. In our vision, ease of use and relevance to a broad range of applications are of paramount importance. These objectives are only achievable by carefully averting any runtime computational overheads. At the same time, user interfaces must provide an intuitive syntax that mirrors standardized, mathematical formulations of optimal control problems (OCPs).

Achieving robust MPC performance is challenging in many ways. Efficient NLP solvers must be coupled with fast derivative computations. Given the large numbers and interrelation-

This research was supported by the Swiss National Science Foundation through the National Centre of Competence in Digital Fabrication (NCCR DFAB) and through Grant No. 200021 200644, and by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant No. 866480).

We express our gratitude to Miguel Angel Zamora Mora and Zijun Hui for their precious contributions to the development of Ungar.

The authors are with the Computational Robotics Lab, ETH Zurich, Switzerland. [flavio.devincenzi@inf.ethz.ch](mailto:flavio.devincenzi@inf.ethz.ch)

ships of the state and control variables, manual implementation of first and second-order derivatives would result in a tedious, error-prone process. Also, although mature automatic differentiation (AD) and NLP libraries exist, most implementations require all the variables stacked in a single vector, which necessitates some index-keeping logic. This fact begs the question of what data structures could store them while guaranteeing zero-cost access operations and adaptability to different system designs.

With Ungar, we provide a metalanguage that addresses these modeling challenges. Our solution introduces constructs that significantly simplify the definition of the NLP problems typically arising in optimal control. We use template metaprogramming (TMP) to delegate the generation of the necessary implementations to the compiler, while users can focus exclusively on the architectural details of a desired MPC application. Our approach makes the transcription into code of structured variable sets seamless while encoding all hierarchical and indexing information at compile time. Consequently, all read/write operations acting on correspondingly created objects incur no additional costs, just like ad hoc programming solutions. Since the core of Ungar is header-only, its integration in C++ projects is effortless. However, we also include an optional interface to CppADCodeGen [3, 15] for automatically generating derivatives and an optional sequential quadratic programming (SQP) solver using OSQP [20] as a backend; if enabled, all external dependencies are automatically downloaded through CMake. Finally, we illustrate the capabilities of Ungar by implementing MPC controllers for increasingly complex systems, including quadruped robots and teams of four-legged manipulators cooperatively carrying an object.

## Related Work

There exist many open-source packages to assist in the creation of MPC controllers. A noncomprehensive list includes frameworks for modeling, simulation, and optimization-based control of mechanical systems, such as the Control Toolbox [10] and its successor OCS2 [9], Drake [21], Crocodyl [18], TOWR [25], FROST [13], Quad-SDK [19], SymForce [17], etc. These libraries are geared toward robotic applications and the OCPs they solve require specific structures. In contrast, NLP-oriented frameworks address more general classes of nonlinear optimization problems; notable examples are IPOPT [24], ACADOS [14, 22], PSOPT [2], CasADi [1], etc. Ungar

```

1 // Define integral constants.
2 constexpr auto N = 30_c;
3 constexpr auto NUM_ROTORS = 4_c;
4
5 // Define "leaf" variables.
6 constexpr auto position = var_c<"position", 3>;
7 constexpr auto orientation = var_c<"orientation", Q>;
8 constexpr auto linear_velocity = var_c<"linear_velocity", 3>;
9 constexpr auto angular_velocity = var_c<"angular_velocity", 3>;
10 constexpr auto rotor_speed = var_c<"rotor_speed", 1>;
11
12 // Define "branch" variables.
13 constexpr auto x = var_c<"x"> <<= (position, orientation, linear_velocity, angular_velocity);
14 constexpr auto X = var_c<"X"> <<= (N + 1_c) * x;
15 constexpr auto u = var_c<"u"> <<= NUM_ROTORS * rotor_speed;
16 constexpr auto U = var_c<"U"> <<= N * u;
17 constexpr auto decision_variables = var_c<"decision_variables"> <<= (X, U);

```

Listing 1. Decision variables of an OCP for controlling a quadrotor with Ungar.




Fig. 1. Hierarchical relationships among state and input variables of an OCP for controlling a quadrotor and underlying memory representation. At every time step  $k$ , we define the robot’s state  $\mathbf{x}_k$  as a stacked vector containing its position  $\mathbf{p}_k \in \mathbb{R}^3$ , orientation  $\mathbf{q}_k \in \mathbb{S}^3 \subseteq \mathbb{R}^4$ , linear velocity  $\dot{\mathbf{p}}_k \in \mathbb{R}^3$ , and angular velocity  $\boldsymbol{\omega}_k \in \mathbb{R}^3$ . The inputs  $\mathbf{u}_k$  consist of the four rotor speeds  $\omega_k^i \in \mathbb{R}, \forall i \in \{1, 2, 3, 4\}$ . Finally,  $\mathbf{X} \in \mathbb{R}^{13(N+1)}$  and  $\mathbf{U}^{4N}$  denote the stacked state and input vectors, where  $N \in \mathbb{N}$  is the discrete time horizon—in the diagram,  $N = 2$ . Using TMP techniques, Ungar supports the generation of data structures for efficiently manipulating raw data arrays.

complements the above libraries by providing novel system modeling features. In particular, it allows for quickly setting up the data structures required by widely adopted AD implementations and OCP solvers. Eventually, while MPC is its primary focus, the design of Ungar makes it suitable for any application requiring the solution of finite dimensional NLP problems.

## II. DATA STRUCTURES

In this section, we introduce the two main data structures at the heart of Ungar: variables and maps. We accompany their descriptions with motivating examples in the robotics domain. We remark that the classes we discuss build on only two external dependencies, namely, Eigen [12] and Boost.Hana [7]. The former is a linear algebra template library ubiquitous in robotics codebases due to its fast performance and versatility; the latter is a collection of utilities and containers that greatly simplify the implementation of TMP algorithms. Since both

are header-only, we wrap them within Ungar to make its integration in C++ projects as straightforward as possible.

### A. Variables

At the core of Ungar lies the Variable template class. Variables describe the structure of quantities of interest such as states, inputs, or parameters. Each variable has a name and a kind<sup>1</sup>, and it is related to all other variables through hierarchical relationships. When a Variable object is instantiated, all this information is encoded in its type using three template parameters: a compile-time string for the name, an integer number identifying its kind, and a compile-time map representing the variable hierarchy. We employ `boost::hana::string` and `boost::hana::map` types to designate the name and the variable hierarchy, respectively. In particular, the latter maps variable names to sub-Variable objects or arrays thereof.

To clarify and explain the above design choices, let us consider the decision variables of a finite horizon MPC for controlling a quadrotor. The state  $\mathbf{x}_k$  of the system at time step  $k$  consists of the robot’s pose and velocity; the control inputs  $\mathbf{u}_k$  comprise four rotor speeds. Given a discrete time horizon  $N \in \mathbb{N}$  and assuming a direct multiple shooting formulation [6], the optimization variables amount to the stacked state and input vectors  $\mathbf{X} = [\mathbf{x}_0^\top \mathbf{x}_1^\top \dots \mathbf{x}_N^\top]^\top$  and  $\mathbf{U} = [\mathbf{u}_0^\top \mathbf{u}_1^\top \dots \mathbf{u}_{N-1}^\top]^\top$ , respectively. Then, we can use the Ungar template object `var_c` to instantiate the relevant variables as shown in Listing 1.

The `var_c` construct takes a fixed string denoting the name of the variable and an optional integer parameter representing its kind: if present, it instantiates a “leaf” variable, i.e., a variable that has no additional substructures; otherwise, it creates a “branch” variable, which is only defined in relation to its subvariables. We identify the kind of a leaf variable with 1 for scalars, an implementation-defined constant `Q` for unit quaternions, and any positive integer for a correspondingly

<sup>1</sup>In this paper, we assign distinct meanings to the terms *kind* and *type*. To clarify, *kind* denotes the mathematical group to which a variable belongs, while *type* exclusively refers to C++ data types.

sized vector. As schematically represented in Fig. 1, Ungar allows mapping the structures encoded by variables to contiguous memory buffers (see Section II-B).

In our example, the lines 1-3 define a discrete time horizon  $N$  of 30 time steps and the number of rotors as integral constants [7] through the user-defined literal `_c`. The lines 5-10 define the leaf variables for the pose and velocity of a quadrotor, as well as its rotor speeds. Instead, the lines 12-17 introduce the branch variables for  $\mathbf{x}_k$ ,  $\mathbf{X}$ ,  $\mathbf{u}_k$ , and  $\mathbf{U}$ , respectively. We observe that  $\mathbf{x}_k$  consists of the stacked position, orientation, linear velocity, and angular velocity of the robot, while  $\mathbf{X}$  contains  $N+1$  stacked states. Then, we can immediately express the above structural information using the overloaded functions `operator<<=`, `operator,`, and `operator*` as shown at lines 13-14. Similarly, we define the branch variables for  $\mathbf{u}_k$  and  $\mathbf{U}$ , and finally we stack  $\mathbf{X}$  and  $\mathbf{U}$  inside the object `decision_variables` at line 17.

All the variables defined in Listing 1 are `constexpr` and can be queried at compile-time for two main pieces of information: their sizes and their indices within the hierarchy. For instance, with reference to the diagram in Fig. 1 and for  $N = 30$ , we can write

```
1 static_assert(
2     x.Size() == 13 &&
3     X.Size() == 403 &&
4     u.Size() == 4 &&
5     U.Size() == 120 &&
6     decision_variables.Size() == 523
7 );
```

and

```
1 static_assert(
2     X(x, 0).Index() == 0 &&
3     X(x, 1).Index() == x.Size() &&
4     X(x, 1, linear_velocity).Index() == 20 &&
5     decision_variables(U).Index() == X.Size() &&
6     U(u, 0).Index() == 0 &&
7     U(u, 1).Index() == 4 &&
8     U(u, 1, rotor_speed, 0).Index() == 4 &&
9     U(u, 1, rotor_speed, 1).Index() == 5
10 );
```

As shown in the above listings, we can access all subvariables of any branch variable using the function call `operator()`. If multiple copies of a subvariable exist, we must write the zero-based index of the subvariable we are interested in. Most importantly, if there is no ambiguity in the path from a branch variable to any of its subvariables, we can bypass all intermediate variables as shown in Listing 2: this feature is very convenient when the variables are organized in deep hierarchies. Finally, Ungar provides macros for defining variables according to the convention that their name is identical to the corresponding object name. Thus, Listing 1 can be rewritten in a more succinct way as shown in Listing 3.

## B. Variable Maps

The `Variable` framework provides the means to describe complex systems with a minimal yet expressive syntax without incurring runtime costs. To turn system descriptions into useful data structures, Ungar offers the template class

`VariableMap`. A variable map associates a variable with an array of scalars and each subvariable with a subarray. To perform the various mappings, we adopt the `Eigen::Map` class [12], which allows interfacing raw buffers with dense matrix expressions seamlessly. All necessary Eigen maps are created during the execution of the `VariableMap` constructor; therefore, accessing any subvariable data has no runtime cost. This is only possible due to our adoption of TMP techniques and makes Ungar akin to a metalanguage.

Given some user-defined scalar type `scalar_t`, we can create a variable map for the quadrotor decision variables introduced in Section II-A as:

```
1 auto vars = MakeVariableMap<scalar_t>(
2     decision_variables);
```

Then, we can access all submaps by passing corresponding subvariables to the `Get` method. For example, we can initialize all unit quaternions to identity rotations and all remaining decision variables to zero [12] as:

```
1 vars.Get(X).setZero();
2 for (auto k = 0; k < N + 1; ++k) {
3     vars.Get(orientation, k).setIdentity();
4 }
5 vars.Get(U).setZero();
6 static_assert(
7     std::same_as<
8     decltype(vars.Get(U)),
9     Eigen::Map<Eigen::VectorX<scalar_t>>&
10 );
```

We remark that all objects returned by the `Get` method have reference types, hence they do not perform copies and directly manipulate the underlying data. Also, the returned type depends on the kind of the corresponding variable, so it can be a reference to `scalar_t`, an Eigen map to a unit quaternion, or an Eigen map to a vector. Branch variables are always mapped to vectors spanning all the corresponding subvariables (line 6).

We enable this flexibility by internally adopting compile-time maps associating subvariables to corresponding data subarrays. While this solution ensures the best possible runtime performance, it can be demanding in terms of compile time and memory consumption. This can be undesirable if, for instance, such variable maps are employed only for intermediate code generation steps; indeed, most code generators optimize the input code, thus making any performance optimizations unnecessary at this stage. We address this need with a lazy version of `VariableMap` named `VariableLazyMap`. Lazy maps instantiate Eigen maps on demand, which is a cheap operation involving the copies of two integer numbers. Their constructors require a data buffer with the correct size as:

```
1 VectorXr underlying{decision_variables.Size()};
2 auto lvars = MakeVariableLazyMap(underlying,
3     decision_variables);
```

In particular, we can rewrite our initialization example as:

```
1 lvars.Get(X).setZero();
2 for (auto k = 0; k < N + 1; ++k) {
3     lvars.Get(orientation, k).setIdentity();
```

```

1 static_assert(X(x, 1, linear_velocity) == X(linear_velocity, 1));
2 static_assert(U(u, 1, rotor_speed, 0) == U(rotor_speed, 1, 0));
3 static_assert(U(u, 1, rotor_speed, 1) == U(rotor_speed, 1, 1));
4 static_assert(decision_variables(X, x, 1, linear_velocity) == decision_variables(linear_velocity, 1));
5 static_assert(
6     decision_variables(U, u, 2, rotor_speed, 3) == decision_variables(u, 2, rotor_speed, 3) &&
7     decision_variables(U, u, 2, rotor_speed, 3) == decision_variables(rotor_speed, 2, 3)
8 );

```

Listing 2. Equivalent expressions for unambiguous variable hierarchies.

```

1 // Define integral constants.
2 constexpr auto N = 30_c;
3 constexpr auto NUM_ROTORS = 4_c;
4
5 // Define "leaf" variables.
6 UNGAR_VARIABLE(position, 3);
7 UNGAR_VARIABLE(orientation, Q);
8 UNGAR_VARIABLE(linear_velocity, 3);
9 UNGAR_VARIABLE(b_angular_velocity, 3);
10 UNGAR_VARIABLE(rotor_speed, 1);
11
12 // Define "branch" variables.
13 UNGAR_VARIABLE(x) <<= (position, orientation, linear_velocity, b_angular_velocity);
14 UNGAR_VARIABLE(X) <<= (N + 1_c) * x;
15 UNGAR_VARIABLE(u) <<= NUM_ROTORS * rotor_speed;
16 UNGAR_VARIABLE(U) <<= N * u;
17 UNGAR_VARIABLE(decision_variables) <<= (X, U);

```

Listing 3. Equivalent transcription of Listing 1 with improved readability using macros.

```

4 }
5 lvars.Get(U).setZero();
6 static_assert(
7     std::same_as<
8         decltype(lvars.Get(U)),
9         Eigen::Map<Eigen::VectorX<scalar_t>>
10     >
11 );

```

We highlight that the only difference between `VariableMap` and `VariableLazyMap` objects lies in the return types of `Get`. As shown in the static assertion of line 6, lazy maps return Eigen maps by value instead of reference.

### III. EXPERIMENTS

We validate Ungar by implementing two MPC schemes for different but related systems. The former is a quadrupedal locomotion controller based on the single rigid body dynamics (SRBD) model [5]; the latter is a centralized MPC controller for collaborative locomotion and manipulation using one-armed quadruped robots [23]. In particular, the cooperative loco-manipulation example extends the simple locomotion controller to multiple arm-endowed robots carrying a shared payload. Our accompanying video presents recordings of simulation experiments using both controllers. These applications allow us to showcase the versatility of Ungar in defining completely different variable hierarchies by changing only a few lines of code. Although we do not provide the source code for the controllers we discuss, we bundle the library with multiple examples, including nonlinear MPC implementations for a quadrotor and a miniature radio-controlled car [16].

#### A. Implementation Details

We implement our controllers relying exclusively on the functionalities provided by Ungar. For our tests, we employ the core data structures alongside the optional `CppADCodeGen` wrapper [15] and SQP solver. We base our solver on the recent work by Grandia et al. [11] and refer the interested reader to [23] for more details.

To generate the necessary derivatives, `CppADCodeGen` requires all functions to depend on a single array of data including both independent variables and parameters. We comply with this interface through Ungar maps as described in Section II-B. For each controller, we define the variable hierarchies `decision_variables` and `parameters`. The decision variables include states and control inputs, while the parameters contain inertial properties, reference trajectories, physical constants, etc. In the following sections, we will show the definitions of only the decision variables of the various MPC controllers to keep the exposition succinct. Nevertheless, we consistently adopt this subdivision in all our implementations, and we invite the reader to explore the examples in the library for a detailed walk-through.

Ungar requires compilers with C++20 support<sup>2</sup>. We implemented and thoroughly tested our controllers on Ubuntu 22.04.2 LTS with GCC 11, but Ungar’s core modules do not depend on any OS-specific instructions. However, we note that our optional `CppADCodeGen` wrapper uses runtime

<sup>2</sup>For convenience, we provide a library version compliant with C++17 on a separate Git branch. Although this adaptation lacks some features, it implements all the most important functionalities of Ungar. We refer the reader to the documentation on the GitHub webpage for further information.

compilation and dynamic linking features offered by Linux.

*Compile Times:* Ungar map objects have no runtime overheads by construction. Thus, we only provide a benchmark of the compile times for systems with different sizes. For this purpose, we measure the time required to build a `VariableMap` and a `VariableLazyMap` object for the decision variables defined in Listing 3. We take this measurement for different values of  $N$  and `NUM_ROTORS` on a laptop computer with an i7-11800H, 2.30 GHz, 16-core CPU, and we plot corresponding heatmaps in Fig. 2. We can see that the lazy map has more favorable compile times, requiring only 14s to build the data structures for an octocopter with  $N = 390$ . In contrast, for the same setup, the compile time of `VariableMap` is 39s long. Although `VariableMap` scales less favorably compared to `VariableLazyMap`, we can notice that they have similar compile time performance for time horizons smaller than 100 time steps. Nevertheless, we recommend using lazy maps for prototyping and switching to `VariableMap` for production code to get the fastest runtime performance.

### B. Quadrupedal Locomotion

We base our MPC formulation for quadrupedal locomotion on the controller of Bledt and Kim [5], but with three notable differences. Firstly, we use the nonlinear SRBD equations without linearizations or simplifications. Secondly, we represent orientations with unit quaternions instead of Euler angles to prevent singularity issues. Lastly, we employ a Lie group time-stepping method to integrate the dynamics conserving quaternion unit-norm constraints [23].

We define our variable hierarchy as shown in Listing 4. In particular, we can see that less than 19 lines of code suffice for generating the data structures required to manipulate the robot’s states and inputs. While an MPC locomotion controller requires additional components to be of practical use, such as gait planners, inverse kinematics solvers, and whole-body controllers, we can already appreciate the potential of Ungar in simplifying the formulation of NLP problems with elaborate structures.

### C. Collaborative Loco-Manipulation

We use Ungar to implement an MPC controller that simultaneously optimizes ground reaction forces, manipulation wrenches, stepping locations, and body trajectories for a team of one-armed quadruped robots collectively manipulating an object. The resulting optimal control problem is very high-dimensional and presents coupled dynamics and deep hierarchies among states and inputs. For instance, each robot has 4 legs, and each leg is associated with a ground reaction force and a stepping location at every time step; also, each robot has an arm that, in turn, corresponds to a manipulation force and torque.

We formulate our MPC for collaborative loco-manipulation (CLM) as an extension of the SRBD to multi-agent systems and refer the interested reader to [23] for a detailed description of our model. As shown in Listing 5, the creation of Ungar variables for the seemingly involved CLM setting requires only minor changes to the locomotion control problem of Listing 4.

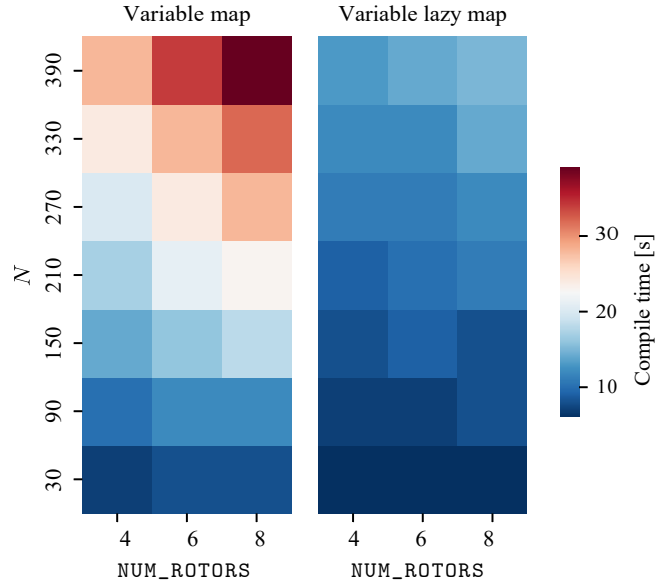


Fig. 2. Compile times to generate the implementations of a `VariableMap` (left) and a `VariableLazyMap` (right) for the multirotor example of Listing 3. We benchmark Ungar against different time horizons and numbers of rotors by varying the lines 2 and 3. The heatmaps manifest the more desirable compile times of the lazy map compared to its non-lazy alternative. For this reason, the `VariableMap` type should only be employed when seeking the best MPC performance possible.

### D. Limitations

The optimal runtime efficiency of Ungar is due to a large number of compile-time computations. However, if the variable hierarchies become too deep or nested, then compile times may become significantly high. Also, we observed compiler crashes when instantiating `VariableMap` objects for very large OCPs, once again caused by the excessive amount of compile-time computations. In these rare cases, it is sufficient to adopt the `VariableLazyMap` type, which is considerably less computation-intensive and provides almost the same performance as its non-lazy version. For future work, we will optimize the design of Ungar to improve its compile times. We additionally plan to expand the library with more tools to facilitate the implementation of high-performance MPC controllers.

## IV. CONCLUSION

In this paper, we introduced Ungar, a C++ template library for real-time MPC applications. Our framework uses TMP techniques to address modeling needs overlooked in existing NLP and optimal control software packages. In particular, it provides a metalanguage to describe complex systems in terms of variable hierarchies. Then, it lets compilers produce highly efficient code for manipulating raw data buffers based on these hierarchies. As shown in our quadruped locomotion and collaborative loco-manipulation experiments, these features enable great flexibility in formulating NLP problems and simplify AD-compliant implementations.

```

1 // Define integral constants.
2 constexpr auto N = 30_c;
3 constexpr auto NUM_LEGS = 4_c;
4
5 // Define "leaf" variables.
6 UNGAR_VARIABLE(position, 3);
7 UNGAR_VARIABLE(orientation, Q);
8 UNGAR_VARIABLE(linear_velocity, 3);
9 UNGAR_VARIABLE(angular_velocity, 3);
10 UNGAR_VARIABLE(force, 3);
11 UNGAR_VARIABLE(relative_position, 3);
12
13 // Define "branch" variables.
14 UNGAR_VARIABLE(leg_input) <<= (force, relative_position);
15 UNGAR_VARIABLE(x) <<= (position, orientation, linear_velocity, angular_velocity);
16 UNGAR_VARIABLE(X) <<= (N + 1_c) * x;
17 UNGAR_VARIABLE(u) <<= NUM_LEGS * leg_input;
18 UNGAR_VARIABLE(U) <<= N * u;
19 UNGAR_VARIABLE(decision_variables) <<= (X, U);

```

Listing 4. Decision variables of an OCP for quadrupedal locomotion using the single rigid body model.

```

1 // Define integral constants.
2 constexpr auto N = 10_c;
3 constexpr auto NUM_ROBOTS = 2_c;
4 constexpr auto NUM_LEGS = 4_c;
5
6 // Define "leaf" variables.
7 UNGAR_VARIABLE(position, 3);
8 UNGAR_VARIABLE(orientation, Q);
9 UNGAR_VARIABLE(linear_velocity, 3);
10 UNGAR_VARIABLE(angular_velocity, 3);
11 UNGAR_VARIABLE(force, 3);
12 UNGAR_VARIABLE(relative_position, 3);
13 UNGAR_VARIABLE(torque, 3);
14
15 // Define "branch" variables.
16 UNGAR_VARIABLE(leg_input) <<= (force, relative_position);
17 UNGAR_VARIABLE(arm_input) <<= (force, torque);
18 UNGAR_VARIABLE(robot_input) <<= (NUM_LEGS * leg_input, arm_input);
19 UNGAR_VARIABLE(payload_state) <<= (position, orientation, linear_velocity, angular_velocity);
20 UNGAR_VARIABLE(robot_state) <<= (position, orientation, linear_velocity, angular_velocity);
21 UNGAR_VARIABLE(x) <<= (payload_state, NUM_ROBOTS * robot_state);
22 UNGAR_VARIABLE(X) <<= (N + 1_c) * x;
23 UNGAR_VARIABLE(u) <<= NUM_ROBOTS * robot_input;
24 UNGAR_VARIABLE(U) <<= N * u;
25 UNGAR_VARIABLE(decision_variables) <<= (X, U);

```

Listing 5. Decision variables of an OCP for collaborative loco-manipulation with two robots modeled as single rigid bodies. We highlight the differences from the locomotion controller formulated in Listing 4. In particular, we mark newly added variables in yellow and modified variables in light blue.

## REFERENCES

- [1] Joel A E Andersson, Joris Gillis, Greg Horn, James B Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019. doi: 10.1007/s12532-018-0139-4.
- [2] Victor M. Becerra. Solving complex optimal control problems at no cost with psopt. *2010 IEEE International Symposium on Computer-Aided Control System Design*, pages 1391–1396, 2010.
- [3] Bradley M. Bell. CppAD: a package for c++ algorithmic differentiation. [Online]. Available: <https://github.com/coin-or/CppAD>.
- [4] Marko Bjelonic, Ruben Grandia, Moritz Geilinger, Oliver Harley, Vivian Suzano Medeiros, Vuk Pajovic, Edo Jelavic, Stelian Coros, and Marco Hutter. Offline motion libraries and online mpc for advanced mobility skills. *The International Journal of Robotics Research*, 41:903–924, 2022.
- [5] G. Bledt and Sangbae Kim. Implementing regularized predictive control for simultaneous real-time footstep and ground reaction force optimization. *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6316–6323, 2019.
- [6] M. Diehl, H. Bock, H. Diedam, and Pierre-Brice Wieber. Fast direct multiple shooting algorithms for optimal robot control. 2005.
- [7] Louis Dionne et al. Boost.Hana: Your standard library for metaprogramming. [Online]. Available: <https://github.com/boostorg/hana>.
- [8] Boston Dynamics. Atlas Gets a Grip — Boston Dynamics. [https://www.youtube.com/watch?v=e1\\_QhJ1EhQ](https://www.youtube.com/watch?v=e1_QhJ1EhQ), January 2023.
- [9] Farbod Farshidian et al. OCS2: An open source library for optimal control of switched systems. [Online]. Available: <https://github.com/leggedrobotics/ocs2>.
- [10] Markus Gifftaler, Michael Neunert, M. Stäuble, and Jonas Buchli. The control toolbox — an open-source c++ library for robotics, optimal and model predictive control. *2018 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAN)*, pages 123–129, 2018.
- [11] Ruben Grandia, Fabian Jenelten, Shao-Hua Yang, Farbod Farshidian, and Marco Hutter. Perceptive locomotion through nonlinear model predictive control. *ArXiv*, abs/2208.08373, 2022.
- [12] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [13] Ayonga Hereid and Aaron D. Ames. Frost: Fast robot optimization and

- simulation toolkit. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, BC, Canada, September 2017. IEEE/RSJ.
- [14] Boris Houska, Hans Joachim Ferreau, and Moritz Diehl. Acado toolkit—an open-source framework for automatic control and dynamic optimization. *Optimal Control Applications and Methods*, 32, 2011.
- [15] Joao Rui Leal et al. CppADCodeGen. [Online]. Available: <https://github.com/joaoleal/CppADCodeGen>.
- [16] Alexander Liniger, Alexander Domahidi, and Manfred Morari. Optimization-based autonomous racing of 1:43 scale rc cars. *Optimal Control Applications and Methods*, 36:628 – 647, 2015. URL <https://api.semanticscholar.org/CorpusID:11242645>.
- [17] Hayk Martiros, Aaron Miller, Nathan Bucki, Bradley Solliday, Ryan Kennedy, Jack Zhu, Tung Dang, Dominic Pattison, Harrison Zheng, Teo Tomic, Peter Henry, Gareth Cross, Josiah VanderMey, Alvin Sun, Samuel Wang, and Kristen Holtz. SymForce: Symbolic Computation and Code Generation for Robotics. In *Proceedings of Robotics: Science and Systems*, 2022. doi: 10.15607/RSS.2022.XVIII.041.
- [18] Carlos Mastalli, Rohan Budhiraja, Wolfgang Merkt, Guilhem Saurel, Bilal Hammoud, Maximilien Naveau, Justin Carpentier, Ludovic Righetti, Sethu Vijayakumar, and Nicolas Mansard. Crocodyl: An Efficient and Versatile Framework for Multi-Contact Optimal Control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [19] Joseph Norby, Yanhao Yang, Ardalan Tajbakhsh, Jiming Ren, Justin K. Yim, Alexandra Stutt, Qishun Yu, Nikolai Flowers, and Aaron M. Johnson. Quad-SDK: Full stack software framework for agile quadrupedal locomotion. In *ICRA Workshop on Legged Robots*, May 2022.
- [20] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: an operator splitting solver for quadratic programs. *Mathematical Programming Computation*, 12(4): 637–672, 2020. doi: 10.1007/s12532-020-00179-2. URL <https://doi.org/10.1007/s12532-020-00179-2>.
- [21] Russ Tedrake and the Drake Development Team. Drake: A planning, control, and analysis toolbox for nonlinear dynamical systems, 2016. URL <http://drake.mit.edu>.
- [22] Robin Verschueren, Gianluca Frison, Dimitris Kouzoupis, Jonathan Frey, Niels van Duijkeren, Andrea Zanelli, Branimir Novoselnik, Thivaharan Albin, Rien Quirynen, and Moritz Diehl. acados – a modular open-source framework for fast embedded optimal control. *Mathematical Programming Computation*, Oct 2021. ISSN 1867-2957. doi: 10.1007/s12532-021-00208-8. URL <https://doi.org/10.1007/s12532-021-00208-8>.
- [23] Flavio De Vincenti and Stelian Coros. Centralized model predictive control for collaborative loco-manipulation. In *Robotics: Science and Systems*, 2023. URL <https://api.semanticscholar.org/CorpusID:259319056>.
- [24] Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006.
- [25] Alexander W Winkler, Dario C Bellicoso, Marco Hutter, and Jonas Buchli. Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robotics and Automation Letters (RA-L)*, 3:1560–1567, July 2018. doi: 10.1109/LRA.2018.2798285.