



HAL
open science

The Open Stack of Tasks Library: OpenSoT

Enrico Mingo Hoffman, Arturo Laurenzi, Nikos G. Tsagarakis

► **To cite this version:**

Enrico Mingo Hoffman, Arturo Laurenzi, Nikos G. Tsagarakis. The Open Stack of Tasks Library: OpenSoT. 2024. hal-04621130v1

HAL Id: hal-04621130

<https://hal.science/hal-04621130v1>

Preprint submitted on 24 Jun 2024 (v1), last revised 12 Dec 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Open Stack of Tasks Library: OpenSoT

Enrico Mingo Hoffman^{1,2}, Arturo Laurenzi², and Nikos G. Tsagarakis²

Abstract—The OpenSoT library is a state-of-the-art framework for instantaneous whole-body motion planning and control based on Quadratic Programming optimization. The library is designed to enable users to easily write and solve a variety of complex instantaneous whole-body control problems with minimal input, facilitating the addition of new tasks, constraints, and solvers. OpenSoT has been designed to be real-time safe and can be conveniently interfaced with other software components such as ROS or other robotic-oriented frameworks. This paper aims to present the usage of the OpenSoT library to a large audience of researchers, engineers, and practitioners, as well as to provide insights into its software design, which has matured over nearly 10 years of development.

Index Terms—Constrained reactive control, whole-body control, model-based control, software frameworks

I. INTRODUCTION

In recent years, we have witnessed the emergence of sophisticated robotic systems characterized by multiple degrees of freedom (DOFs) and limbs, enabling them to navigate through and interact with their environment. For instance, mobile platforms equipped with manipulators, quadrupeds, and more recently, full humanoids, are transitioning from laboratory settings to practical applications [1]. These systems surpass classical industrial manipulators in terms of capabilities, including a broader workspace, extended range of motions, and enhanced control of interaction forces through their whole body. However, it is widely acknowledged that these robots demand sophisticated motion and interaction controllers. These controllers must consider hardware limitations, and manage a large number of DOFs and multiple tasks simultaneously.

One of the most well-established control paradigms in robotics is based on computing the next control action by solving an inverse problem linearized at the robot's current state. Following this paradigm, various control schemes have been implemented in real robotic systems, including resolved-rate control, closed-loop inverse kinematics, operational space control, and Cartesian impedance control, among others. These methods also allow for the management of actuation and force redundancies and have been adapted for use with both fixed- and floating-base robots. Optimization techniques, particularly Quadratic Programming (QP), have become the standard approach for efficiently solving these types of control problems with the possibility to include linear constraints, e.g. joint limits or collision avoidance. The development of

these methods has been particularly active in the context of whole-body control where multiple tasks and constraints need to be resolved simultaneously, for example when controlling humanoids and quadrupeds.




Fig. 1: Some of the robots that utilized OpenSoT as a whole-body motion engine include MARM in the top-left, COMAN+ in the bottom-left, CENTAURO in the center, and WALK-MAN on the right.

This article describes an open-source software library written in C++ to write instantaneous whole-body control problems and resolve them using optimization, designed for research, educational, and industrial applications, named *Open Stack of Tasks: OpenSoT*. OpenSoT is aimed at different audiences: motion planning and control researchers, robotics educators, practitioners, and end users in the robotics industry. Within the robotics community, showcasing the superiority of a new task, constraint, solver, or formulation, over existing ones, can be a daunting task. It requires a considerable investment of time and effort for a researcher to implement these components and conduct a thorough comparison with the existing state-of-the-art. Furthermore, the whole-body algorithms should effectively and efficiently address motion control problems for systems with numerous DOFs in real-time. To address this challenge, we developed OpenSoT to simplify this process and facilitate the exploration of innovative ideas. The division between abstract base classes, which solely define the interface, and derived classes responsible for implementing the specified functionality, contributes to a clearer comprehension of general concepts in whole-body planning and control. From its inception, OpenSoT has been designed to be practical and applicable in real-world scenarios, and it has been deployed in many real robotic systems, e.g. see Figure 1. Another crucial requirement is the seamless integration of OpenSoT with other software components on a robot, including high-level planning, perception, kinematics and dynamics modeling, and low-level control. This paper aims to introduce the software's evolution, illuminate key lessons learned throughout its development, and underscore its enduring significance in

¹ Université de Lorraine, CNRS, Inria, LORIA, France, 615 rue du Jardin Botanique, 54600 Villers les Nancy, France. E-mail: enrico.mingo-hoffman@inria.fr

² Humanoid and Human Centered Mechatronics (HHCM) Lab, Istituto Italiano di Tecnologia (IIT), Via Morego 30, 16163 Genova, Italy. E-mail: {arturo.laurenzi, nikos.tsagarakis}@iit.it

current robotic applications.

II. BACKGROUND

We address control problems involving the *linear mapping* of quantities from operational space, such as Cartesian space, to generalized coordinates. These problems comprise various atomic elements that can be categorized into two main groups: *tasks* and *constraints*. Generally, a task can be accurately mapped into generalized coordinates when constraints are absent. Tasks can be mathematically formalized as linear least-squares pieces of a cost function to be minimized. Each task can be characterized by a task matrix $\mathbf{A}_i \in \mathbb{R}^{m \times n}$, a task vector $\mathbf{b}_i \in \mathbb{R}^m$, and a task weight matrix $\mathbf{W}_i \in \mathbb{R}^{m \times m}$. Constraints are in the form of linear equalities and/or inequalities characterized by a constraint matrix $\mathbf{C} \in \mathbb{R}^{l \times n}$ and lower and upper constraint vectors, respectively $\mathbf{l}, \mathbf{u} \in \mathbb{R}^l$:

$$\begin{aligned} \min_{\mathbf{x}} \sum_{i=0}^N \|\mathbf{A}_i \mathbf{x} - \mathbf{b}_i\|_{\mathbf{W}_i} + \epsilon \|\mathbf{x}\| \\ \text{s.t. } \mathbf{l} \leq \mathbf{C}\mathbf{x} \leq \mathbf{u}, \end{aligned} \quad (1)$$

with $\mathbf{x} \in \mathbb{R}^n$ the vector of unknowns, e.g. generalized coordinates, and m the task size. The second term in the cost function in (1) is referred to as Tikhonov regularization. Its purpose is to regularize ill-posed problems that may arise during the computation of the Lagrangian function associated with the optimization problem. Without this regularization, the Hessian matrix $\mathbf{H} = \mathbf{A}^T \mathbf{W} \mathbf{A} \in \mathbb{R}^{n \times n}$ would lack invertibility to solve the Karush-Kuhn-Tucker (KKT) conditions, as $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{H}) = m$, when $m < n$. This particular condition is quite common when mapping a Cartesian task in a redundant robot.

When multiple tasks need to be mapped together, or if the generalized coordinates do not provide enough DOFs, *priorities* need to be established between them. Priorities can be divided into two types: *Soft* and *Hard*. Soft priorities can be mathematically formalized as weighted sums in the cost function, where high weights are assigned to the most important tasks, as in (1). Hard priorities can be mathematically formalized as low-priority tasks resolved in the null-space of high-priority ones. This can be implemented in multiple ways; we will briefly address common methods to implement hard priorities between tasks throughout the paper. The distinction between hard and soft priorities lies in the fact that, in soft priorities, all tasks are concurrently minimized (in the least-square sense), and their optimality is contingent upon the relative weights assigned to each task, leading to average-like solutions when tasks conflict. In contrast, hard priorities involve minimizing secondary tasks *without* compromising the optimality of the primary task. In robotic systems with many DOFs, normally hard and soft priorities are mixed. Constraints are usually hard unless slack variables are used. Therefore, when two constraints conflict, the QP problem defined in (1) becomes *infeasible*. Two major families of mapping problems in robotics use the QP problem in (1): *differential inverse kinematics* and *inverse dynamics*.

A. Differential Inverse Kinematics

In Differential Inverse Kinematics (DIK) *instantaneous* Operational Space velocities are mapped into joint space velocities. The QP problem in the DIK case with multiple soft tasks is formulated as follows:

$$\begin{aligned} \dot{\mathbf{q}}^* = \operatorname{argmin}_{\dot{\mathbf{q}}} \sum_{i=0}^N \|\mathbf{J}(\mathbf{q})_i \dot{\mathbf{q}} - \mathbf{v}_i\|_{\mathbf{W}_i} + \epsilon \|\dot{\mathbf{q}}\| \\ \text{s.t. } \mathbf{l} \leq \mathbf{C}(\mathbf{q})\dot{\mathbf{q}} \leq \mathbf{u}. \end{aligned} \quad (2)$$

First-order integration of $\dot{\mathbf{q}}^*$ permits to solve the non-linear Inverse Kinematics (IK) problem for the *ith* task given by:

$$\mathbf{T}_{i,d} = \mathbf{f}_i(\mathbf{q}), \quad (3)$$

where $\mathbf{T}_{i,d}$ is a desired Cartesian pose and $\mathbf{f}_i(\cdot)$ representing the Forward Kinematic (FK) function for the *ith* task. The actual and desired poses are often utilized to compute an appropriate error, which is then employed as a Cartesian velocity reference in the task. This error exponentially tends to zero based on a positive definite proportional gain.

B. Inverse Dynamics

Similarly, it is possible to formulate the QP problem in terms of joint acceleration variables and subsequently use the optimized joint accelerations to solve the IK problem through second-order integration. Another application of the optimized joint accelerations is the computation of joint torques associated with the motion, i.e., the Inverse Dynamics (ID).

An interesting case is the floating-base ID where other than the joint accelerations, also the contact forces are computed within the QP optimization:

$$\begin{aligned} \min_{\dot{\nu}, \mathbf{f}} \sum_{i=0}^N \|\mathbf{J}_i(\mathbf{q})\dot{\nu} + \dot{\mathbf{J}}_i(\mathbf{q}, \nu)\nu - \mathbf{a}_i\|_{\mathbf{W}_i} + \epsilon_\nu \|\dot{\nu}\| + \epsilon_f \|\mathbf{f}\| \\ \text{s.t. } \mathbf{M}_b(\mathbf{q})\dot{\nu} + \mathbf{h}_b(\mathbf{q}, \nu)\nu = \mathbf{J}_{c,b}^T(\mathbf{q})\mathbf{f} \\ \mathbf{J}_c(\mathbf{q})\dot{\nu} + \dot{\mathbf{J}}_c(\mathbf{q}, \nu)\nu = \mathbf{0} \\ \mathbf{l} \leq \mathbf{C}(\mathbf{q}, \nu)\dot{\nu} \leq \mathbf{u}, \end{aligned} \quad (4)$$

with $\mathbf{q} \in \mathbb{R}^n + SO(3)$, $\nu \in \mathbb{R}^{n+6}$, $\mathbf{J}_c(\mathbf{q}) \in \mathbb{R}^{3c \times n+6}$ the stack of the Jacobians of the contacts and $\mathbf{J}_{c,b}(\mathbf{q}) \in \mathbb{R}^{3c \times 6}$ its floating-base part, $\mathbf{M}_b(\mathbf{q}) \in \mathbb{R}^{6 \times n+6}$, $\mathbf{h}_b(\mathbf{q}, \nu)\nu \in \mathbb{R}^6$ the floating-base dynamics and non-linear terms, respectively, and $\mathbf{f} \in \mathbb{R}^{3c}$ where c are the number of contacts. Also in this case, the optimized joint accelerations and contact forces can be used to compute the joint torques:

$$\mathbf{M}(\mathbf{q})\dot{\nu} + \mathbf{h}(\mathbf{q}, \nu) = \boldsymbol{\tau} + \mathbf{J}_{c,j}^T(\mathbf{q})\mathbf{f}, \quad (5)$$

with $\mathbf{J}_{c,j}(\mathbf{q}) \in \mathbb{R}^{3c \times n}$ the joint space part of the stacked contact Jacobians.

III. STATE OF THE ART

Considerable effort has been dedicated to both algorithm and software development for instantaneous whole-body motion planning and control. The primary objectives of these software frameworks are twofold: first, to simplify the formulation of optimization problems like those in (2) or (4); second,

to enable efficient solving of such problems. The first goal is typically addressed by segregating each task and constraint into atomic entities that can be flexibly combined to construct the optimization problem. This optimization problem is then tackled by a solver, which is often partitioned into a *front-end*, responsible for generating elements (matrices and vectors), and a *back-end*, that is, an off-the-shelf optimization library that uses these elements to solve the problem.

A pioneering work in this direction has been the *Stack of Tasks* [2] (SoT), which allowed the formulation and resolution of complex Inverse Kinematics (IK) problems with hard priorities in humanoid robots, eventually extended to Inverse Dynamics (ID) [3]. In the SoT, each task is represented as an entity class loaded through a plugin within a scripting language, enabling the construction of a *stack* defining the problem to be solved, where each level of the stack corresponds to a priority level. As the reader may infer, OpenSoT draws significant inspiration from this foundational work. Similar software frameworks to the SoT have been developed within several laboratories. For instance, the Optimization-Based Controllers for Robotics Applications (OCRA) [4] was developed within the YARP community for the humanoid robot iCub. Another example is the ihmc-whole-body-controller [5], developed for the DARPA Robotics Challenge (DRC) and based on Java. Other notable frameworks include Drake [6], ControllIt! [7], TSID [8], and iTask [9]. While many of these software frameworks are intended to function as comprehensive packages for robotics, OpenSoT was specifically crafted to focus solely on generic formulations of whole-body controllers. OpenSoT provides an extensive library of implemented tasks, constraints, and solvers, along with a clearly defined developer API for adding new ones. Its abstraction level and generic components make it straightforward to integrate OpenSoT with various other libraries for creating new tasks and constraints, as well as with other software frameworks.

IV. OPENSOT: OVERVIEW AND MAIN COMPONENTS

In contrast to other existing whole-body control software libraries, OpenSoT does not incorporate a representation of the robot model and avoids assuming a specific controller formulation or priority resolution strategy. The advantage of this minimalist approach is that it enables the design of a library suitable for the generic creation and resolution of hierarchical QP and Linear Programming (LP) problems subject to linear constraints. For instance, beyond the conventional formulations of whole-body controllers discussed in Section II, OpenSoT can be utilized to address problems like contact force distribution or floating-base estimation.

The OpenSoT framework is built upon the *hierarchical paradigm*: initially, an external model of the robot is updated, allowing the computation of kinematics and dynamics quantities necessary for various *tasks* and *constraints*. Subsequently, tasks and constraints are updated to compute all the matrices and vectors required to formulate and solve the optimization problem using a dedicated *Solver*. While the model can be updated using measured or derived quantities, each task and constraint contains a reference to the model to retrieve all

the necessary matrices and vectors. Additionally, each task and constraint independently manages various settings, such as gains, parameters, and high-level references.

Tasks and constraints are organized into a *Stack*, providing a convenient means to specify types and levels of priorities. Now, we will examine the fundamental concepts associated with OpenSoT's component (see Fig. 2). Throughout the following sections, references to the robot model will be made. However, it's crucial to note that, as introduced earlier, the OpenSoT API doesn't encompass classes for representing robots. This design choice offers the advantage of creating a library that can function independently and seamlessly integrate with various generic model libraries, such as RBDL [10], or Pinocchio [11]. For tasks and constraints implemented in the library that utilize the robot's model, we depend on a generic `ModelInterface` class, developed within the XBotCore [12] software suite, which can be easily specialized using the aforementioned kinematic and dynamic modeling libraries. The development of the OpenSoT library commenced in 2015 as the motion engine for the COMAN and WALKMAN robot, participating in the DARPA Robotics Challenge (DRC) Finals. The initial version of OpenSoT was closely tied to the YARP robotics framework [13], utilizing `yarp::sig` matrices and vectors, and the `iDynTree` model library. Its functionality was limited to resolving Inverse Kinematics and Differential Inverse Kinematics problems. Following the DRC, the library underwent significant restructuring, transitioning to the `Eigen` library for all linear algebra computations, and adding further capabilities.

A. Tasks

OpenSoT employs a different expression for the cost function compared to similar software frameworks:

$$\|\mathbf{A}_i \mathbf{x} - \mathbf{b}_i\|_{\mathbf{W}_i} + \mathbf{c}_i^T \mathbf{x}, \quad (6)$$

where $\mathbf{c}_i \in \mathbb{R}^n$, allowing the specification of both least-squares ($\mathbf{c}_i = \mathbf{0}$) and linear problems ($\mathbf{A}_i = \mathbf{0}$), e.g., for Lasso regression. A task object is inherited from the base class `OpenSoT::Task`, where the method `update()` must be implemented to assign the matrix \mathbf{A} and vectors \mathbf{b} and \mathbf{c} , at every control loop. As mentioned earlier, this assignment is typically done by requesting kinematics and dynamics quantities from a reference to the model, which is updated outside the task. The weight matrix \mathbf{W} can be set using a dedicated method exposed to users. The `GenericTask` class provides users with methods to assign all the internal matrices and vectors, eliminating the need to derive from the base class. For this reason, the `GenericTask` is a convenient way to easily interface with libraries for prototyping reasons. To exclude joints that belong to a kinematic chain inside a task, it is possible to use the `applyActiveJointsMask()` method, which selectively sets columns in the \mathbf{A} matrix of the task to $\mathbf{0}$. Finally, a task can be activated or deactivated using the `setActive()` method. When a task is not active, its \mathbf{A} matrix is set to $\mathbf{0}$.

A *SubTask* comprises a specific number of rows from a task and is constructed from a task object along with the rows

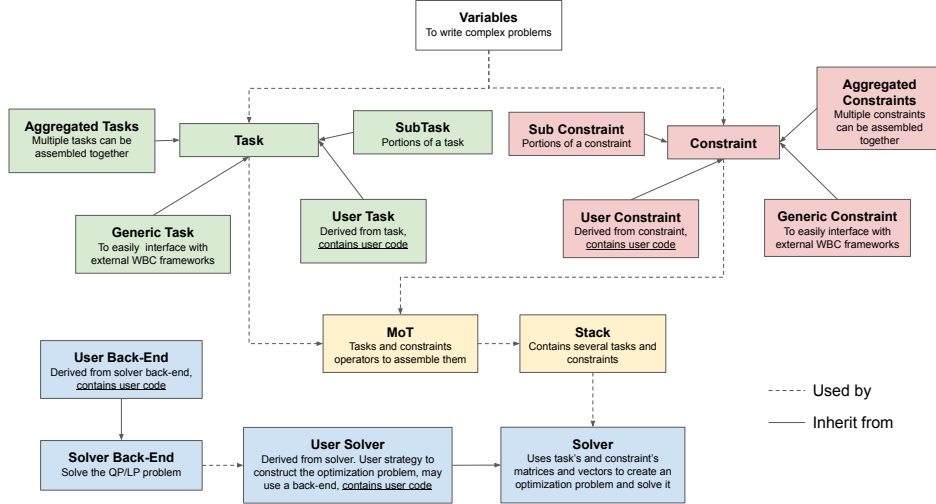


Fig. 2: Main components in OpenSoT.

of interest. The `SubTask` class facilitates the selection of adjacent and non-adjacent rows from a task by extracting submatrices from \mathbf{A} and \mathbf{W} , along with a sub-vector from \mathbf{b}^1 . A `SubTask` holds a reference to the task that was used to instantiate it; therefore, the original task is still used to handle settings and references. As a result, any modifications applied to the original task are reflected in the `SubTask`, and vice versa. The referenced task is also updated when the `update()` method is invoked on the `SubTask`. A common application of a `SubTask` is to focus on a specific portion of a task. For instance, in a Cartesian task, a `SubTask` can be employed to isolate the positional component while disregarding the orientation. A library of ready-to-use tasks is available in OpenSoT. The most important available tasks are reported in Table I with the implemented formulation. Tasks with “*” are mostly used as equality constraints.

TABLE I: Implemented tasks in OpenSoT

Task	Formulation
Cartesian	Velocity/Acceleration
Postural	Velocity/Acceleration
CoM	Velocity/Acceleration
Angular Momentum	Velocity/Acceleration
Gaze	Velocity
Pure Rolling*	Velocity
Floating-Base Dynamics*	Acceleration + Force
Min Effort	Velocity
Manipulability	Velocity

B. Constraints

A constraint object is inherited from the base class `OpenSoT::Constraint`, where the method `update()` must be implemented to assign the matrix \mathbf{C} , and the vectors \mathbf{c}_l and \mathbf{c}_u , for the lower and upper bounds respectively. As for the Subtask, a `SubConstraint` comprises a specific number of rows from a constraint, and the `GenericConstraint` permits assigning internal matrices and vectors without the need to derive from the base `Constraint` class. A library

of ready-to-use constraints is also available in OpenSoT. In Table II are reported the most significant constraints with the implemented formulation and typology. In particular, joint po-

TABLE II: Implemented constraints in OpenSoT

Constraint	Formulation	Type
Joint Pos./Vel./Acc. Limits	Velocity/Acceleration	Inequality
Torque Limits	Acceleration + Force	Inequality
CoP Limits	Wrench	Inequality
Friction Cone Limits	Force	Inequality
Normal Torque Limits	Wrench	Inequality
Collision Avoidance	Velocity	Inequality
OmniWheels	Velocity	Equality

sition, velocity, and acceleration limits are implemented using exponential control barrier functions and invariance. Friction cones use pyramidal approximation, and normal torque limits implement bounds on the contact normal torque as in [14].

C. Variables

In OpenSoT, the matrix \mathbf{A} and the vectors \mathbf{b} and \mathbf{c} completely define a task, and similarly, the matrix \mathbf{C} and the vectors \mathbf{c}_l and \mathbf{c}_u completely define a constraint. When formulating the optimization problem using these matrices, the *optimization variables are implicit* and depend on how these matrices and vectors are written. In other words, we do not explicitly write the variables of the problem; instead, we infer them from the structure of these matrices and vectors.

For example, in the DIK problem (2) the generic Cartesian task is defined by the task matrix $\mathbf{A} = \mathbf{J}$, and the variables are the joint velocities $\mathbf{x} = \dot{\mathbf{q}}$. In the ID problem (4), the generic Cartesian task is defined by the task matrix $\mathbf{A} = [\mathbf{J} \ \mathbf{0}]$, whose size depends on the number of contacts, in fact the variables of the problem are $\mathbf{x} = [\dot{\mathbf{v}}^T \ \mathbf{f}^T]^T$. From a developer’s point of view, this is not convenient because, depending on the number of contact forces, the size of the task matrix will vary, requiring more code to handle a variable that is not used in that particular task. This also complicates the process of reusing the task, particularly in scenarios like differential

¹SubTasks do not make use of the \mathbf{c} vector

inverse kinematics at the acceleration level, where contact forces and dynamics are not present. Ideally, a developer should be able to implement the task irrespective of other variables, focusing solely on joint accelerations.

To address this issue, OpenSoT introduces the concept of *explicit variables* through the `OpenSoT::AffineHelper` class. An *affine* variable is defined as:

$$\mathbf{y} = \mathbf{M}\mathbf{x} + \mathbf{p} \quad (7)$$

and can be used to define (or extract) explicit variables from \mathbf{x} . For example, considering the variables in (4):

$$\begin{aligned} \dot{\mathbf{v}} &= \mathbf{M}_{\dot{\mathbf{v}}}\mathbf{x} \\ \mathbf{f} &= \mathbf{M}_{\mathbf{f}}\mathbf{x} \end{aligned} \quad (8)$$

with $\mathbf{M}_{\dot{\mathbf{v}}} = [\mathbf{I}_{n+6 \times n+6} \quad \mathbf{0}_{n+6 \times 3c}]$ and $\mathbf{M}_{\mathbf{f}} = [\mathbf{0}_{3c \times n+6} \quad \mathbf{I}_{3c \times 3c}]$. In this way is possible to write tasks and constraints that are independent of the definition of the problem optimization variables. If we consider again the Cartesian task in acceleration, the task matrix will be $\mathbf{A} = \mathbf{J}\mathbf{M}_{\dot{\mathbf{v}}}$.

The `AffineHelper` class exposes an easy-to-use API to create the vector of variables \mathbf{x} and compute the matrices to extract these variables. The `VariableVector` type allows stacking multiple variables together using the `emplace_back(var_name, var_size)` method. It is then possible to retrieve the variable by name using the `getVariable(var_name)` method along with the associated \mathbf{M} matrix and \mathbf{p} vector.

Some tasks in OpenSoT can be constructed using a variable, for example, all the tasks formulated in joint acceleration or contact forces, to be easy to combine. Utilities to transform tasks and constraints that do not make use of variables are provided, allowing the inclusion of implicit tasks and constraints in explicit stacks.

D. Stack

In OpenSoT a Stack consists of one or more tasks and constraints and their relations. Tasks and constraints can be *aggregated* together through opportune stacking of the associated matrices and vectors. Notice that such operation can be seen from the point of view of the cost function as a *sum* of tasks. OpenSoT uses a Domain Specific Language (DSL) called *Math of Tasks* [15] which permits to easy specify various types of relations between tasks and constraints objects. For example, the sum of multiple tasks is done by the “+” (sum) operator, while relative weights can be set by using the “*” (product) operator. A `SubTask` can be specified inside a stack using the “%” (modulo) operator. Finally, constraints can be inserted into a stack using the “<<” (left shift) operator².

Applying these operators creates a stack object, implemented through the `OpenSoT::AutoStack` class. An `autostack` carries pointers to the associated tasks and constraints, therefore when the `autostack` is updated through the `update()` method, also the internal tasks and constraints are updated.

²This operator can be used also with tasks to create equality constraints

E. Solvers

Solvers are responsible for resolving the optimization problem presented in Equation (1). A user solver is derived from the base class `OpenSoT::Solver`, where the user is required to implement the method `solve()`. Solvers typically consist of two components: a *Front-end* and a *Back-end*: the front-end utilizes the stack to construct the matrices and vectors that will be utilized by the back-end to solve the problem. For instance, the inequality Hierarchical QP (iHQP) solver [16] adheres to this framework. The front-end calculates the Hessian and constraints and sets up a series of QP problems corresponding to the number of hard priorities specified in the stack, which are sequentially addressed. At each priority level, the front-end incorporates an additional constraint considering the solution from the preceding solved QP. Each QP is tackled by a dedicated back-end, implemented by an out-of-the-shelf QP solving library. This framework can be applied to various techniques implemented in the front-end, while the back-ends can be reused, thereby promoting reusability.

Back-ends can be developed by utilizing the base class `OpenSoT::BackEnd`. When creating a back-end, the methods `initProblem()` and `solve()` need to be implemented. Additional methods are mandatory: `setOptions()` and `getOptions()` to transmit and retrieve options to and from the solvers, respectively, and the method `getObjective()` to retrieve the residual from the optimization. Solvers that deviate from this pattern can be directly implemented using the `Solver` base class. Table III lists the solvers already integrated into OpenSoT. Specifically, we interface with the HCOD implementation available in the `soth` package. While Table IV lists the available back-ends in

TABLE III: Implemented solvers in OpenSoT

Solver	Constraints Type	Use Back-end
eHQP [17]	equality only	×
iHQP [16]	equality/inequality	✓
nHQP [18]	equality/inequality	✓
HCOD [19]	equality/inequality	×

OpenSoT. It’s worth noting that unlike many similar software packages, we also provide LP and MIP solvers, expanding the range of applications for these formulations in whole-body control problems, as well as facilitating their integration with QP formulations (as done by the authors in [20]). We refer

TABLE IV: Implemented back-ends in OpenSoT

Solver	Type
qpOASES [21]	QP/LP
OSQP [22]	QP/LP
proxQP [23]	QP/LP
qpSWIFT [24]	QP
eiQuadProg [25]	QP
GLPK [26]	LP/MIP

to the cited papers for detailed descriptions of the solvers and back-ends. Finally in Figure 3 are reported benchmark results on a IK problem casted on a 29 DOFs floating-base humanoid robot. We considered 4 possible tasks varying the number of hard priorities (from single layer to 4 priority levels) including joint position and velocity limits and contact constraints. The




Fig. 3: 29 DOFs humanoid robot benchmark.

tasks include the positioning of the Center of Mass (CoM), left and right hand poses, and a joint level posture. We randomly selected 30 IK problems, each involving a transition from an initial configuration to a goal configuration defined by a specific Cartesian pose. Each IK problem was solved using the iHQP and nHQP front-end solvers, using the back-ends: OSQP, qpOASES, eigQuadProg, proxQP, and qpSWIFT. We compared the average time taken to solve a single instance of a hierarchical problem³, including the time needed to compute the Hessian and gradient, as well as to fill the solver matrices. Additionally, we calculated the success rate (SR), defined as the number of times the back-end successfully found a solution within certain conditions, divided by the total number of runs. To confirm that the back-end successfully performs the task, we consider the task error norm to be less than or equal to $1e^{-3}$ within the specified maximum number of iterations, set to 1000. Local minima as well as near singular configurations may reduce the SR. For robots with multiple DOFs, simple constrained IK problems can be solved fast, within 1 *ms* and 10 *ms*. The success rate is for most of the cases between 80% and 100%. Notably, in the case of the iHQP, the time taken to resolve multiple layers does not change significantly when increasing the number of layers to values typically used in complex robotics systems (4-5 layers). Notice that these times do not take into account the time needed to update the model and stack. This duration can vary depending on the model library utilized and the implementation of various tasks and constraints. The tests were conducted without fully leveraging the extensive capabilities of the examined QP solvers, both in terms of tunable options and optimal implementation⁴. For example, some of the QP solvers only offer a sparse interface, necessitating the conversion of dense matrices, which are typical of the problems in the whole-body framework, into sparse format. Consequently, the obtained results may be affected by suboptimal option tuning and less-than-optimal implementation. For these reasons, these results are provided

solely to offer a preliminary understanding of the performance in solving specific control problems using OpenSoT.

V. MAIN EXTENSIONS AND INTEGRATION WITH OTHER FRAMEWORKS

The OpenSoT control library has been integrated with other frameworks to enhance its usability and capabilities, as well as facilitate the development of new tasks or constraints. Specifically, we will here discuss four of the major extensions of OpenSoT.

A. *Cartesi/O* & *ROS*

OpenSoT doesn't prescribe a specific structure for controllers but offers components for crafting controllers with an intuitive programmatic interface. However, there are specific scenarios where the controller's structure remains consistent, and the only variables are the stack and/or solver utilized. For instance, in IK, the controller's framework involves initializing the reference to the current robot state, solving the optimization problem, and integrating to calculate the subsequent reference. *Cartesi/O* [28] is a high-level framework designed for two primary objectives: firstly, to provide user-friendly high-level interfaces for tasks, constraints, and solvers, and secondly, to offer readily configurable controllers. The core concept is to furnish base classes for tasks and constraints, necessitating the implementation of an *interface* class and a *computation* class. In particular, interface classes are ROS-based, while computation classes utilize OpenSoT. Consequently, ROS topics and services are exposed for each task and constraint, facilitating easy parameter tuning and reference setting through interactive markers, among other functionalities. Each component functions as a plugin that can be loaded with a YAML file, wherein initial parameter values, stack structure, solver types, and other pertinent information are specified. This YAML file, also known as a *.stack* file, forms the basis of a controller template, automatically configured to the user's requirements. This approach enables users to construct,

³On an AMD® Ryzen 9 4900HS with 32 GiB of RAM

⁴An accurate comparison of QP solvers can be found in [27].

configure, and interact with controllers without the necessity of low-level coding. Additionally, CartesI/O provides Python bindings and integrates the Reflexxes motion library to smooth commands for Cartesian tasks in OpenSoT.

B. Visp

The Visual Servoing Platform (Visp) is a well-known framework in robotics dedicated to visual servoing [29]. Within OpenSoT, Visp is utilized for computing interaction matrices, which map various type of visual features into end-effector velocities, and for managing data structures to handle these features. The `opensot_visual_servoing` package implements a velocity task in OpenSoT, enabling the tracking of desired visual features using image feedback. The `tasks::velocity::VisualServoing` task contains a `tasks::velocity::Cartesian` task object responsible for computing Jacobians related to the specified camera frame. The camera frame Jacobian is pre-multiplied by the interaction matrix computed by Visp based on camera parameters. Measured visual features are used as feedback and can be provided to the task via the `setFeatures()` method at every control loop. Visual features are computed from an RGB image using a Python script, which then transmits the results through a ROS topic to a CartesI/O interface designed for the `VisualServoing` task.

C. HPP-FCL

HPP-FCL is a well-known library utilized for efficiently computing distances between shapes [30]. This library has been incorporated into the model interface library to enable the utilization of self-collision and object-collision avoidance inequality constraints in OpenSoT. Self-collision avoidance is supported for any type of collision meshes. Through CartesI/O, the `PlanningScene` from ROS is leveraged to seamlessly insert or remove shapes and point clouds into the constraint.

D. OMPL

The Open Motion Planning Library (OMPL) is a comprehensive collection of state-of-the-art sampling-based motion planning algorithms [31]. Within OMPL, OpenSoT is utilized to project random state configurations for floating-base systems onto the manifold defined by the contacts through IK, implemented inside `ompl::base::Constraint`. This projection is performed while simultaneously considering secondary tasks and constraints in various forms, such as stability checks, collisions, and generalized preferred postures.

VI. AN EXAMPLE

In this section, we now describe an example of the usage of the OpenSoT library to create and solve an ID stack for the humanoid robot TALOS. Despite OpenSoT provides Python bindings for its main components, our example will be written using the C++ API. First, we create the model of the robot starting from the URDF and SRDF descriptions:




Fig. 4: Open loop whole-body inverse dynamics performed on TALOS model.

```
using namespace XBot;
auto model = ModelInterface::getModel(
    ReadFile("robot.urdf"),
    ReadFile("robot.srdf"), "pin");
```

The SRDF model is used to specify further parameters than the URDF such as initial configurations, kinematics chains, self-collision maps, etc... The third parameter permits to specify the kinematics/dynamics library to be used, i.e. Pinocchio.

Second, we declare the variables we will use, in particular, joint accelerations and contact forces for each contact point:

```
using namespace OpenSoT;
using namespace std;

OptvarHelper::VariableVector variable_vec;
variable_vec.emplace_back("qddot", model.getNv());
for(auto contact : contact_points)
    variable_vec.emplace_back(contact, 3);

auto variables =
    make_shared<OptvarHelper>(variable_vec);

auto qddot = variables->getVariable("qddot");
vector<AffineHelper> contact_forces;
for(auto contact : contact_points)
    contact_forces.push_back(
        variables->getVariable(contact));
```

Third, we declare the tasks and constraints that will compose the stack:

```
#define MS(x) std::make_shared<typename x>

using namespace OpenSoT::tasks::acceleration;
using namespace OpenSoT::constraints::acceleration;
using namespace OpenSoT::constraints::force;

// Cartesian tasks
auto l_sole =
    MS(Cartesian)("LS", model, "l_sole", qddot);
...
// CoM task
auto com = MS(CoM)(model, qddot);

// Postural task
auto postural = MS(Postural)(model, qddot);
...

// Underactuation constraint
auto floating_base = MS(DynamicFeasibility)(
    "floating_base", model, qddot,
    contact_forces, contact_points);

// Friction cones
double mu = 0.8;
```



```

auto LFULFc = MS(FrictionCone)(
    contact_points[0], contact_force[0], model, mu);
...

```

Once the tasks and the constraints are declared, we can compose them in a stack, for example, considering a single-level stack:

```

// Concatenate tasks
auto stack =
    LS + RS + 1e-3*postural + base%{3,4,5} + ...

// Add constraints
stack<<floating_base<<v_lims<<j_lims<<LFULFc<<...

```

notice that we are setting a relative weight of $1e^{-3}$ for the `postural` task and we are considering the orientation part for the `base` task using the operator “%” that creates a subtask from the original task using the provided list of indices.




Fig. 5: Comparison between different back-ends running the whole-body inverse dynamics on the TALOS robot using the iHQP front-end. The run time considers the solver time, the update of the model and the stack, and a small overhead due to the communication to ROS provided by the CartesI/O framework.

Finally, we instantiate the solver that will compute the optimal joint accelerations and contact forces:

```

auto solver = MS(solvers:iHQP)(
    autostack, solvers::solver_back_end::proxQP);

```

The control loop consists of the following:

```

while (control_on)
{
    // Update model and stack
    model->setJointPosition(q);
    model->setJointVelocity(qdot);
    model->update();

    com->setReference(com_p_ref, com_v_ref);
    ...
    autostack->update();

    // Solve, retrieve solution and compute ID
    if (!solver->solve(x))
        // do something...

    qddot->getValue(x, qddot_val);
    for (int i = 0; i < contact_forces.size(); ++i)
        contact_forces[i]->getValue(x, f_val[i]);
}

```

```

auto tau = ID(qddot_val, f_val, contact_points);
...
}

```

notice that any reference or parameter variation in tasks or constraints has to be done *before* the autostack update to be actualized. The `getValue()` method is used to retrieve the variable values from the solution vector x .

The final stack is composed of Cartesian tasks for the arm’s and leg’s end-effectors, CoM XY positions, centroidal angular momentum regulation to zero, orientation of the base, a postural for the upper body, and regularization of contact forces and joint accelerations. The constraints are the floating-base dynamics, joint torques, velocity and position limits, force limits, pyramidal friction cones, and the unilaterality of contact forces. The problem has a total of 60 variables, 24 contact forces and 36 joint accelerations, and 172 constraints. The stack has been implemented through a CartesI/O stack file permitting the exploitation of the Python CartesI/O API and ROS for setting references to tasks, and RVIZ for visualization. For the sake of simplicity, we implemented an open-loop scheme where optimal joint accelerations are integrated twice and fed back to the model. To solve the optimization problem we use the iHQP front-end. We command the right arm to follow a linear path involving also a movement of the entire body, as shown in Figure 4. We compared different back-end solvers, as shown in Figure 5. In Figure 6 are reported the computed torques and in Figure 7 the related contact forces.




Fig. 6: Computed torques.

VII. CONCLUSION AND FUTURE WORKS

This paper describes OpenSoT, an open-source library for reactive whole-body control on generic robotic systems. The main components of the library are discussed, including the OpenSoT API, and its integration with other software libraries and robotics frameworks. OpenSoT has been adopted in various robotics platforms, including manipulators, humanoids, and quadrupedal robots. Its integration with CartesI/O provides an introduction to the complexity of whole-body control: without writing any code, users can construct whole-body problems by combining ready-to-use tasks and constraints written for different formulations, and solve such problems using different solvers, vary parameters, and perform extensive

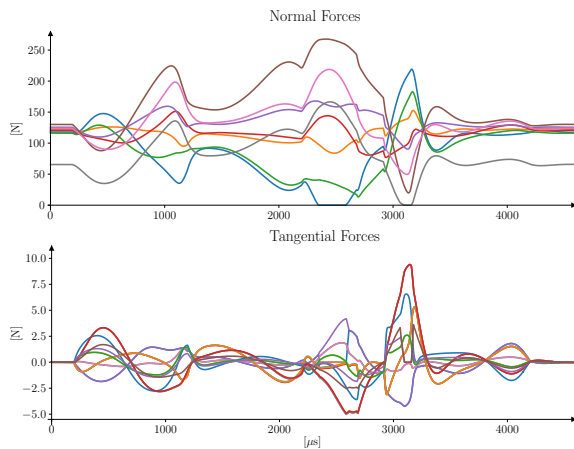


Fig. 7: Contact forces.

benchmarking experiments. Through ongoing development in our group and contributions from others, we expect OpenSoT to become a valuable tool for researchers, robotics industry users, and students. We envision utilizing OpenSoT in various scenarios except from traditional whole-body control. It can be leveraged for modeling and analyzing the design of complex robotics systems, conducting preliminary studies, and offering valuable insights. We have already discussed its integration with OMPL for manifold projection, considering priorities. Another potential application of OpenSoT could involve state estimation, such as floating-base estimation. Lastly, we believe that it could play a role in Reinforcement Learning by reducing the size of the input of the policy through Cartesian mapping into generalized coordinates.

REFERENCES

- [1] E. M. Hoffman, S. Caron, F. Ferro, L. Sentis, and N. G. Tsagarakis, “Developing humanoid robots for applications in real-world scenarios [from the guest editors],” *IEEE Robotics & Automation Magazine*, vol. 26, no. 4, pp. 17–19, 2019.
- [2] N. Mansard, O. Stasse, P. Evrard, and A. Kheddar, “A versatile generalized inverted kinematics implementation for collaborative working humanoid robots: The stack of tasks,” in *International Conference on Advanced Robotics*, 2009, pp. 1–6.
- [3] O. E. Ramos, N. Mansard, O. Stasse, C. Benazeth, S. Hak, and L. Saab, “Dancing humanoid robots: Systematic use of osid to compute dynamically consistent movements following a motion capture pattern,” *IEEE Robotics & Automation Magazine*, vol. 22, no. 4, pp. 16–26, 2015.
- [4] J. G. Eljaik, R. Lober, A. Hoarau, and V. Padois, “Optimization-based controllers for robotics applications (ocra): The case of icub’s whole-body control,” *Frontiers in Robotics and AI*, vol. 5, p. 24, 2018.
- [5] T. Koolen, S. Bertrand, G. Thomas, T. De Boer, T. Wu, J. Smith, J. Engelsberger, and J. Pratt, “Design of a momentum-based control framework and application to the humanoid robot atlas,” *International Journal of Humanoid Robotics*, vol. 13, 03 2016.
- [6] R. Tedrake, “the drake development team,” *Drake: A planning, control, and analysis toolbox for nonlinear dynamical systems*, vol. 5, p. 8, 2016.
- [7] C.-L. Fok, G. Johnson, J. D. Yamokoski, A. Mok, and L. Sentis, “ControlIt!—a software framework for whole-body operational space control,” *International Journal of Humanoid Robotics*, vol. 13, no. 01, 2016.
- [8] A. Del Prete, N. Mansard, O. E. Ramos, O. Stasse, and F. Nori, “Implementing torque control with high-ratio gear boxes and without joint-torque sensors,” *International Journal of Humanoid Robotics*, vol. 13, no. 01, 2016.
- [9] J. De Schutter, T. De Laet, J. Rutgeerts, W. Decré, R. Smits, E. Aertbeliën, K. Claes, and H. Bruyninckx, “Constraint-based task specification and estimation for sensor-based robot systems in the presence of geometric uncertainty,” *The International Journal of Robotics Research*, vol. 26, no. 5, pp. 433–455, 2007.
- [10] M. L. Felis, “Rbd: an efficient rigid-body dynamics library using recursive algorithms,” *Autonomous Robots*, vol. 41, no. 2, pp. 495–511, 2017.
- [11] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, “The pinocchio c++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives,” in *IEEE/SICE International Symposium on System Integration (SII)*, 2019, pp. 614–619.
- [12] L. Muratore, A. Laurenzi, E. Mingo Hoffman, and N. G. Tsagarakis, “The xbot real-time software framework for robotics: From the developer to the user perspective,” *IEEE Robotics & Automation Magazine*, vol. 27, no. 3, pp. 133–143, 2020.
- [13] G. Metta, P. Fitzpatrick, and L. Natale, “Yarp: yet another robot platform,” *International Journal of Advanced Robotic Systems*, vol. 3, no. 1, p. 8, 2006.
- [14] S. Caron, Q.-C. Pham, and Y. Nakamura, “Stability of surface contacts for humanoid robots: Closed-form formulae of the contact wrench cone for rectangular support areas,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 5107–5112.
- [15] E. M. Hoffman and N. G. Tsagarakis, “The math of tasks: A domain specific language for constraint-based task specification,” *International Journal of Humanoid Robotics*, vol. 18, no. 03, 2021.
- [16] O. Kanoun, F. Lamiroux, and P.-B. Wieber, “Kinematic control of redundant manipulators: Generalizing the task-priority framework to inequality task,” *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 785–792, 2011.
- [17] F. Flacco and A. De Luca, “Discrete-time redundancy resolution at the velocity level with acceleration/torque optimization properties,” *Robotics and Autonomous Systems*, vol. 70, pp. 191–201, 2015.
- [18] M. de Lasa and A. Hertzmann, “Prioritized optimization for task-space control,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 5755–5762.
- [19] A. Escande, N. Mansard, and P.-B. Wieber, “Hierarchical quadratic programming: Fast online humanoid-robot motion generation,” *The International Journal of Robotics Research*, vol. 33, no. 7, pp. 1006–1028, 2014.
- [20] E. M. Hoffman, M. P. Polverini, A. Laurenzi, and N. G. Tsagarakis, “A study on sparse hierarchical inverse kinematics algorithms for humanoid robots,” *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 235–242, 2020.
- [21] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, “qpOASES: a parametric active-set algorithm for quadratic programming,” *Math. Program. Comput.*, vol. 6, no. 4, pp. 327–363, 2014.
- [22] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, “OSQP: an operator splitting solver for quadratic programs,” *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020.
- [23] A. Bambade, S. El-Kazdadi, A. Taylor, and J. Carpentier, “Prox-qp: Yet another quadratic programming solver for robotics and beyond,” in *Robotics: Science and Systems*, 2022.
- [24] A. G. Pandala, Y. Ding, and H.-W. Park, “qpswift: A real-time sparse quadratic program solver for robotic applications,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3355–3362, 2019.
- [25] D. Goldfarb and A. Idnani, “A numerically stable dual method for solving strictly convex quadratic programs,” *Mathematical programming*, vol. 27, no. 1, pp. 1–33, 1983.
- [26] “Gnu linear programming kit.” [Online]. Available: <http://www.gnu.org/software/glpk/glpk.html>
- [27] S. Caron, A. Zaki, P. Otta, D. Arnström, J. Carpentier, and F. Yang, “qpbenchmark: Benchmark for quadratic programming solvers available in Python.” 2024. [Online]. Available: <https://github.com/qpsolvers/qpbenchmark>
- [28] A. Laurenzi, E. Mingo Hoffman, L. Muratore, and N. G. Tsagarakis, “Cartesi/o: A ros based real-time capable cartesian control framework,” in *International Conference on Robotics and Automation (ICRA)*, 2019, pp. 591–596.
- [29] E. Marchand, F. Spindler, and F. Chaumette, “Visp for visual servoing: a generic software platform with a wide class of robot control skills,” *IEEE Robotics and Automation Magazine*, vol. 12, no. 4, pp. 40–52, 2005.
- [30] L. Montaut, Q. L. Lidec, V. Petrik, J. Sivic, and J. Carpentier, “Collision detection accelerated: An optimization perspective,” *arXiv preprint*, 2022.
- [31] I. A. Şucan, M. Moll, and L. E. Kavraki, “The Open Motion Planning Library,” *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, December 2012.