



Safe and compliant control of redundant robots using superimposition of passive task-space controllers

Carlo Tiseo · Wolfgang Merkt ·
Wouter Wolfslag · Sethu Vijayakumar ·
Michael Mistry

Received: 4 April 2023 / Accepted: 13 October 2023 / Published online: 27 November 2023
© The Author(s) 2023

Abstract Animals are capable of robust and reliable control in unstructured environments, where they effortlessly overcome the uncertainty of interaction and are capable of exploiting singularities. These conditions are a well-known challenge for robots due to the limitations of projected dynamics, which requires

accurate modelling and is susceptible to singularities. This work proposes a compliant passive control method for redundant manipulators based on a superimposition of multiple passive task-space controllers in a hierarchy without requiring any knowledge of the robot dynamics. The proposed control framework of passive controllers is inherently stable, numerically well-conditioned (as no matrix inversions are required), and computationally inexpensive (as no optimisation is used). This method leverages and introduces a novel stiffness profile for a recently proposed passive controller with smooth transitions between the divergence and convergence phases making it particularly suitable when multiple passive controllers are combined through superimposition. The experimental results demonstrate that the proposed method achieves sub-centimetre tracking performance during demanding dynamic tasks with fast-changing references, while remaining safe to interact with and robust to singularities. The data further show that the robot can fully take advantage of the redundancy to maintain the primary task accuracy while compensating for unknown environmental interactions, which is not possible from current frameworks that require accurate contact information.

Carlo Tiseo and Wolfgang Merkt have contributed equally to this work.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11071-023-09045-x>.

C. Tiseo
School of Engineering and Informatics, University of Sussex,
Brighton BN1 9RH, UK

C. Tiseo (✉) · W. Merkt · W. Wolfslag · S. Vijayakumar ·
M. Mistry
School of Informatics, University of Edinburgh, 10 Crichton St
Edinburgh, Edinburgh EH8 9AB, UK
e-mail: c.tiseo@sussex.ac.uk

W. Merkt
e-mail: wolfgang@robots.ox.ac.uk

W. Wolfslag
e-mail: wjwolfslag@gmail.com

S. Vijayakumar
e-mail: sethu.vijayakumar@ed.ac.uk

M. Mistry
e-mail: michael.mistry@ed.ac.uk

W. Merkt
Oxford Robotics Institute, University of Oxford, 23 Banbury Rd,
Oxford OX2 6NN, UK

Keywords Impedance control · Interaction control ·
Model-free control

1 Introduction

Redundant robots—which have more degrees of freedom (DoF) than required for a task—have been widely studied and deployed due to their intrinsic versatility. The higher dimensionality of the joint configuration space with respect to (w.r.t) the task-space makes these systems more adaptable as multiple solutions can be found.

However, this flexibility introduces a higher complexity for both planning and control that rapidly increases with the system and task dimensionality. For example, computing the joint configuration from a task-space pose, i.e. inverse kinematics (IK), becomes increasingly more challenging with the increase of the redundancy dimension [1]. This problem is also encountered when dealing with the inverse dynamics problem, which is used to derive the control laws used in interaction control [1–3].

On the other hand, biological systems reliably interact with uncertain dynamics despite their highly redundant structures, which are also characterised by the highly nonlinear dynamics introduced by soft tissues, such as muscles and connective tissues. To address these inverse problems, multiple optimisation frameworks and approaches to deal with challenges arising from numerical conditioning have been developed. Currently, inverse problems for soft robots¹ and optimisation of the task-space dynamics are still open problems [4]. A recent approach to robustness for achieving task-space compliance behaviours includes systems for increasing the robustness of projections through software by modulating/adapting the references [2]. This robustness can also be achieved by exploiting more complex hardware design that embeds variable mechanical compliance directly into the robot structure and its actuation [3]. As an example, *Keppler et al.* have proposed a control architecture that enables to retain both robustness and accurate tracking [3], while related work has algorithmically optimised the spatiotemporal modulation of impedance to achieve tasks more efficiently [5]. These approaches take advantage of hardware equipped with variable stiffness actuators (VSAs)—these systems allow better exploitation of the hardware to achieve more dynamic movements; how-

ever, this comes with increased complexity and cost of motion planning in addition to the need for very accurate modelling of the VSA structures.

Inverse kinematics solutions and task-space dynamics projections are required for controlling redundant robots and they share similar challenges, as analysed in depth in [6,7]. In summary, both problems rely on the inversion of the Jacobian matrix, which is non-square in redundant manipulator due to the different task and joint-space dimensions [1]. The pseudo-inverse is a transformation that solves such a problem. It separates the information regarding robot states into two orthogonal sub-spaces (task-space and null-space), which are not expected to exchange information (i.e. energy). Therefore, retaining the orthogonality between these two sub-spaces is paramount for the algorithms' stability [6,7]. Maintaining this orthogonality depends on both the robot kinematics and the task—which can be quite difficult to achieve and maintain, especially during highly variable situations, such as sudden changes in contacts and dynamic interactions (Fig. 1). As one example, the dynamically consistent inverse obtains orthogonality via the minimisation of the kinetic energy projected by the null-space into the task-space [7]. These methods have found application in many architectures exploiting null-space projection to control redundant manipulators; however, they all face the same problem of the degeneration of the projection close to singularities [8–10].

Passive controllers have been proposed to theoretically guarantee interaction stability under uncertain interaction conditions, using for instance virtual tanks as energy storage (i.e. path integral) for the non-conservative energy of the controller. However, their passive behaviour trades off tracking performances to retain the safety of interaction, making this framework difficult to deploy in highly variable environments [8]. The result presented in their manuscript focuses on verifying passivity and safety of interaction and does not provide a clear quantitative analysis of task-space tracking performance. However, based on the plots of the Cartesian tracking error (Fig. 5 in [8]) reported for the simulation results for a 4-DoF planar manipulator, it seems that it can be expected about 1 cm residual pose error in the best case scenario (i.e. when there is energy available in the tank). Moreover, virtual tank impedance controllers are only passive if there is energy left in their virtual tanks. Due to passivity constraints, the tanks' energy can only be charged from external

¹ Here, soft robots are used to refer to both systems made from non-rigid materials and those with compliant control, e.g. collaborative robots.




Fig. 1 Stack of passive controllers executing a reference motion to follow a line trajectory while an unknown external disturbance is introduced to the elbow joint. The controller adapts safely and compliantly to the disturbance and degrades tracking performance gracefully. The end-effector controller $\mathbf{h}_e^e = K(\bar{\mathbf{x}}^e)\bar{\mathbf{x}}^e = K(\bar{\mathbf{x}}^e)(\bar{\mathbf{x}}_d^e - \bar{\mathbf{x}}^e)$ ensures that the end-effector accurately follows the trajectory. The secondary controller on the elbow $\mathbf{h}_e^{el} = K(\bar{\mathbf{x}}^{el})\bar{\mathbf{x}}^{el} = K(\bar{\mathbf{x}}^{el})(\bar{\mathbf{x}}_d^{el} - \bar{\mathbf{x}}^{el})$, being less stiff,

ensures the robot can deform and recover in the null-space in response to external perturbations. In this case, the external perturbation ($\vec{\mathbf{f}}_{Ext}$) is applied by the human pushing a box against the link connecting the shoulder joints to the elbow joint. The control signal is the joint torque vector computed by the superimposition of the two controllers $\vec{\mathbf{\tau}}_{tot} = J_e^T \mathbf{h}_e^e + J_{el}^T \mathbf{h}_e^{el}$, where J_e^T and J_{el}^T are the end-effector and the elbow Jacobians, respectively

energy sources [8]. Realising passive control is made even more difficult when dealing with null-space and task-space controllers. In fact, as they are orthogonal to each other, tracking the total energy exchanged by the manipulator is challenging [11]. Another challenge to the stability of virtual tank controllers is to maintain the orthogonality between null-space and task-spaces during highly variable tasks. Higher nonlinearity in the dynamics reduces the accuracy in the computation of the orthogonal projection that, consequentially, generates unaccounted energy transfer between the two sub-spaces [7, 11].

This work investigates the possibility of using the superimposition of passive task-space controllers to drive redundant manipulators rather than relying on null-space controllers, cf. Fig. 1. Conceptually, this follows the idea to generate task-space wrenches at multiple links and map them back to joint-level torques. To do so, the proposed solution will not need to rely on any mathematical projections (and implicitly, matrix inversions) required by the null-space projections. Virtual mechanical constraints are instead generated using

the superimposition of task-space controllers to control task-space and the redundant degrees of freedom of the manipulators. However, implementing such a solution will require a controller framework that is intrinsically stable. The recently proposed fractal impedance controller (FIC) [11] is a passive controller meeting this requirement. It relies upon a nonlinear stiffness behaviour in the task-space to track the energy exchanged between the robot and the environment and treats the unexpected energy flow from the null-space as an external perturbation. The controller uses the concept of *fractal impedance* for the implementation of a passive controller that can provide good performances in both trajectory and force tracking. Thereby, it detaches the robot stability from the postural optimisation, which are currently bounded for interaction controllers relying on quadratic programming (QP) optimisation [2]. Furthermore, the proposed method is independent of any specific type of actuation and thus can be used in any torque/force-controlled robot.

In summary, the contributions are:

1. *Superimposition of Passive Task-Space Controllers* to preserve the primary task by sacrificing redundancy tasks through the exploitation of the mechanical redundancy. The priority of the controller is determined by the maximum force exertable by the controller, as will be explained in Sect. 2. As the framework only relies on the forward computation of kinematics and geometric Jacobian, it is numerically stable and computationally inexpensive. Further, it can be used with uncertain and imprecise dynamics models as it only relies on the kinematics model (2.5).
2. Proposal of a new force profile for a fractal attractor which enables a smooth transition between convergence and divergence phases 2.4.
3. Validation of the proposed approach both in simulation to test contact interaction with unknown obstacles and in real hardware experiments 3 to evaluate reference tracking performance for fast reference motion 4. As all open parameters have a physically tractable meaning and the controller is intrinsically stable, online tuning can safely be performed.

The intention is to release the implementation for simulation and hardware experiments open source with the publication of this manuscript.

2 Method

The null-space of a redundant manipulator is the set of solutions that do not lead to a variation in the task output, e.g. the set of joint-space configurations that achieve the same end-effector pose. On highly redundant robots, null-space optimisation frameworks are used to identify the optimal joint-space configuration for a given task. The null-space can then be used to fulfil redundancy tasks, e.g. to minimise energy or difference to a nominal configuration, or to adapt to changes imposed by the environment. A popular approach is stack-of-tasks (SoT) optimisation methods [12]. SoT approaches are iterative algorithms which apply null-space optimisation to a hierarchy of tasks: Subsequent tasks are projected onto and solved in the null-space of the higher-order task. As thus, this hierarchical approach guarantees that lower-order tasks do not lead to a deterioration in performance (or violation of) a higher-order task. Their main limitation, however, is

that null-space projections are inserted in the control loop, rendering the controllers susceptible to numerical instability connected with the null-space projections.

The proposed method, shown in Fig. 1, aims to remove null-space projections from the control loop. Null-space projections are used to account for the external interaction in the whole-body control optimisation problem [2]. This type of formulation requires not only to make a priori assumptions on the environmental interaction, but it also renders the controller stability dependent on their accuracy. Thus, the controller stability is highly susceptible to erroneous assumptions [2].

The proposed method unravels the co-dependency between stability and assumptions made on the external environment by using the superimposition of task-space controllers that generate virtual force fields (i.e. soft mechanical constraints) that pull the robot towards the desired configuration. This is a different approach to handling redundancy compared to the null-space approach. The superimposition of the controller will bias the robot to move towards a certain preferred posture, without guaranteeing that this particular configuration will be reached. In fact, the controller will continuously maintain a mechanical equilibrium between the virtual forces and the environmental interaction without requiring any assumption on the environmental interaction. This implies that the controller is robust to unknown environmental interactions, but is not guaranteed to be in a global optimum. However, it is likely to settle in the closest minimum in the system's energetic manifold (i.e. the closest state with mechanical equilibrium). In contrast to weighted optimisation approaches, the particular structure of the fractal attractor (cf. Sect. 2.4) in this passive control approach ensures that the hierarchy of tasks is maintained, and a large violation of a lower-order task does not lead to performance deterioration in a higher-order task.

In synthesis, the relative strength of these virtual constraints will determine the trade-off between the tasks assigned to the controllers and, consequentially, the order in which the task accuracy will be sacrificed. While this method can be applied with any type of task-space controller, using passive controllers guarantees stability by independently verifying that all the superimposed controllers are stable. Among the different passive controllers, the fractal impedance controller thanks to its explicit formulation of the tasks in terms of virtual mechanical constraints (i.e. desired

force/displacement behaviour) enables direct control of the controller trade-off policies.

The proposed approach makes use of the kineto-static duality to address the inverse problem, which is recapped in Sect. 2.1. The fractal impedance controller is described in Sect. 2.3, for which a novel force profile is introduced that has a more robust formulation, which is essential when using multiple superimposed tasks in Sect. 2.4. Finally, the concept of superimposition for the control of redundant robots is introduced in Sect. 2.5. The proposed framework is able to track task-space references for each of the tasks individually. For the experiments, these were obtained from an inverse-kinematics-based postural optimisation (Sect. 2.6), however, highlighting that the task-space targets for the individual tasks can be obtained from any planning method and do not affect the proposed control approach.

2.1 Inverse problem and kineto-static duality

The inverse problem is not possible for the redundant robots due to the asymmetry of the problem dimensionality [1, 7, 13, 14]. However, an approximate solution can be identified by introducing constraints to the equation using the Moore–Penrose inverse or pseudoinverse [1, 7, 13]. Thus, it became an essential tool for solving inverse kinematic and projected inverse dynamics in redundant robots [1, 7, 14]. These constraints of the pseudoinverse define a set, called null-space, made of all the possible states in generalised coordinates (i.e. joint-space) for which the robot’s end-effector remains static. Thus, the constraint introduced is the minimisation of the end-effector twist (i.e. linear and angular velocities) [1, 7, 13] to deal with the redundancy. Kathib has proposed to scale this approach with the robot’s inertia that changes the optimisation criterion to the minimisation of the kinetic energy projected in the task-space [1, 7, 13, 14]. These characteristics make the control of the null-space critical in torque-controlled robots allowing the modulation of the robot manipulability and interaction response to interactions that bypass the end-effector.

The generalised inverse of $\mathbf{A} \in \mathbb{R}^{n \times m}$ is defined as any matrix $\mathbf{G} \in \mathbb{R}^{m \times n}$ that satisfies the following equations:

$$\begin{cases} \vec{\mathbf{a}} = \mathbf{G}\vec{\mathbf{b}} + (\mathbf{I}_n - \mathbf{G}\mathbf{A})\vec{\mathbf{a}}_\epsilon = \mathbf{G}\vec{\mathbf{b}} + \mathbf{P}\vec{\mathbf{a}}_\epsilon \\ \mathbf{A}\mathbf{G}\mathbf{A} - \mathbf{A} = \mathbf{0} \end{cases} \quad (1)$$

where $\vec{\mathbf{a}} \in \mathbb{R}^n$, $\vec{\mathbf{b}} \in \mathbb{R}^m$, $\vec{\mathbf{a}}_\epsilon \in \mathbb{R}^n$ and $\mathbf{I}_n \in \mathbb{R}^{n \times n}$ is the identity matrix. \mathbf{P} is a projection matrix that projects a generic vector \mathbf{a}_ϵ into the null-space of \mathbf{A} , $\mathcal{N}(\mathbf{A})$. It is also worth noting that Equation 1 is a solution for $\mathbf{A}\vec{\mathbf{a}} = \vec{\mathbf{b}}$ [1, 2]. Redundant robots are more versatile than non-redundant systems; however, they do not have a bijective transformation between generalised coordinates and task-space. Thus, control algorithms rely on numerical optimisation to solve the inverse problem and identify viable strategies. This is task-dependent and degenerates when \mathbf{A} drops rank (i.e. $\det(\mathbf{A}) = 0$) [2, 7, 15]. Specifically, the rank of the inverse projection matrix drops if the robot is in a singular configuration or the task constraints are violated (e.g. unexpected sudden loss of contact) [7].

The idea of taking advantage of the kineto-static duality to address the inverse problem has been introduced with the concept of Port-Hamiltonian control in [16]. In fact, the kinematic joint-space information can be used to derive task-space behaviour and task-space force interaction can be used to relate back to joint-space torques:

$$\begin{cases} \vec{\mathbf{v}} = \mathbf{J}\dot{\vec{\mathbf{q}}} \\ \vec{\boldsymbol{\tau}} = \mathbf{J}^T\vec{\mathbf{h}} \end{cases} \quad (2)$$

where $\mathbf{J} \in \mathbb{R}^{n \times m}$ is the geometric Jacobian matrix, $\vec{\mathbf{v}} \in \mathbb{R}^n$ is the end-effector twist, $\dot{\vec{\mathbf{q}}} \in \mathbb{R}^m$ is the joint velocities’ vector, $\vec{\mathbf{h}} \in \mathbb{R}^n$ is the end-effector wrench and $\vec{\boldsymbol{\tau}} \in \mathbb{R}^m$ is the joint torques’ vector [1].

2.2 Passive and active systems

A passive system is a system that either conserves or dissipates its energy [11, 17]. An active system is a system that is capable of generating energy [11, 17]. Therefore, a controlled system can be defined passive as long as all the energy the controllers generate is either conservative or dissipated via non-conservative terms (i.e. damping). It is worth noting that what are commonly referred to as energy generators are systems transforming one form of energy into another. For example, a petrol engine transforms the chemical energy in the fuel into mechanical energy. As a matter of fact, all known physical systems are passive, but this does not apply to all controllers [11, 17]. A benefit of passive systems is that the superimposition of multiple stable passive systems retains both properties [11, 17, 18], which decouples the stability problem, making the problem com-

plexity scale linearly rather than exponentially with the system dimension.

A controller is passive if it does not track the system velocity term because acceleration and position terms are associated with conservative energy forms, kinetic and potential energy, respectively [11].

A mono-dimensional system, in its most general formulation, is described by the following equation:

$$M(\tilde{x})\ddot{\tilde{x}} - D(\tilde{x}, \dot{\tilde{x}})\dot{\tilde{x}} + K(\tilde{x})\tilde{x} = \mathbf{f}_{Ext} \quad (3)$$

where $M(\tilde{x})$ is the inertia, $D(\tilde{x}, \dot{\tilde{x}})$ is a positively defined damping, $K(\tilde{x})$ is a nonlinear stiffness, x is the system measured trajectory, and $\tilde{x} = x_d - x$ is the error between the desired trajectory (x_d) and x .

The equation of a mono-dimensional active controller can be in its most general form is

$$M(\tilde{x})\ddot{\tilde{x}} + D(\tilde{x}, \dot{\tilde{x}})\dot{\tilde{x}} + K(\tilde{x})\tilde{x} = \mathbf{f}_{Ext} \quad (4)$$

2.2.1 Passivisation of an active system

An active system can be passivised using a conservative observer, known as a virtual tank, which is a path integral that acts as a virtual spring tracking the non-conservative energy exchanged by the robot [11, 17]. The energy in the virtual tank defines the maximum amount of non-conservative energy that the controller can use to track velocities without violating passivity [11, 17]. However, being path integrals, virtual tanks retain the issues associated with model errors and integration drift of any other numerical integration in a discrete domain [11]. A redundant system's passivisation is more challenging than in a fully defined system due to the need to rely on Eq. 1 to transform the information between task-space and joint-space, being the null-space not observable in the task-space [1, 11, 17]. Tank-based controllers increase interaction robustness, compared to traditional interaction controllers, by sacrificing tracking accuracy when the virtual tank runs out of energy. Nevertheless, they still introduce a path integral that requires tracking all the energy exchanged that could be neglected if using a passive controller. For example, the robot centrifugal and centripetal forces (i.e. Coriolis Matrix) are passive forces introduced by the relative movements between the robot links reference frame [1]. These components of the robot dynamic equation are derived from the kinetic energy that can be discarded when computing the system's energy without controllers (i.e. Lagrangian) [1] and

when the robot is controlled using passive controllers [11]. Finally, another limit that passivised and active controllers shared is they need to measure all the energy exchanged with the environment to guarantee stability; thus, they rely on the accuracy of the contact models and sensing, which is one of the major concerns in their robustness [2, 11, 12, 15, 17, 19]. Figure 2 illustrates the difference between active, passive and passivised systems based on the tracking of the velocity error ($\dot{\tilde{x}}$).

2.3 Fractal impedance controller

The FIC controls the robot as a nonlinear mass–spring system and generates the attractor in Fig. 6a around the desired state [11]. The fractal attractor allows to nonlinear impedance profile that cannot run with traditional impedance controllers. These nonlinear profiles enable high tracking accuracy without requiring velocity tracking, thus allowing the controller passivity to be retained without compromising accuracy [11]. The equivalent mechanical system equation for the controlled robot is:

$$\Lambda_c(\tilde{\mathbf{q}})\ddot{\tilde{\mathbf{x}}} + \mathbf{n}(\tilde{\mathbf{q}}, \dot{\tilde{\mathbf{q}}}) + \mathbf{K}(\tilde{\mathbf{x}})\tilde{\mathbf{x}} = \tilde{\mathbf{f}}_{Ext} \quad (5)$$

where $\Lambda_c(\tilde{\mathbf{q}})$ is the projection of the task-space inertia matrix at the end-effector, $\mathbf{n}(\tilde{\mathbf{q}}, \dot{\tilde{\mathbf{q}}})$ is the nonlinear robot dynamics, and $\tilde{\mathbf{f}}_{Ext}$ is the external force. $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_d - \tilde{\mathbf{x}}$ is the pose error, where $\tilde{\mathbf{x}}_d$ is the desired state and $\tilde{\mathbf{x}}$ is the current state. The FIC command in Eq. 5 is $\tilde{\mathbf{h}}_e = \mathbf{K}(\tilde{\mathbf{x}})\tilde{\mathbf{x}}$. The state-dependent stiffness gain $\mathbf{K}(\tilde{\mathbf{x}})$ is derived from the desired end-effector interaction properties (i.e. force/displacement), which can be regulated online without affecting stability [11]. The nonlinear stiffness of the controller is evaluated independently for each of the controller dimensions, and it is anisotropic between the divergence (i.e. tracking error is increasing) and the convergence (i.e. tracking error is decreasing) phases [11]. Intuitively, during the divergence phase, the energy is accumulated in the spring and, subsequently, released in convergence. The attractor is implemented using a switching behaviour that introduces an additional nonlinear spring which triggers when the system starts converging (i.e. zero crossing of $\tilde{\mathbf{x}}$). The updated impedance conserves the energy accumulated in the controller while diverging, and it redistributes the energy altering the trajectory during the convergence, as shown in Fig. 3. Therefore, the stability of the controller is guaranteed by the fractal attractor (Fig. 3). This determines the passivity of

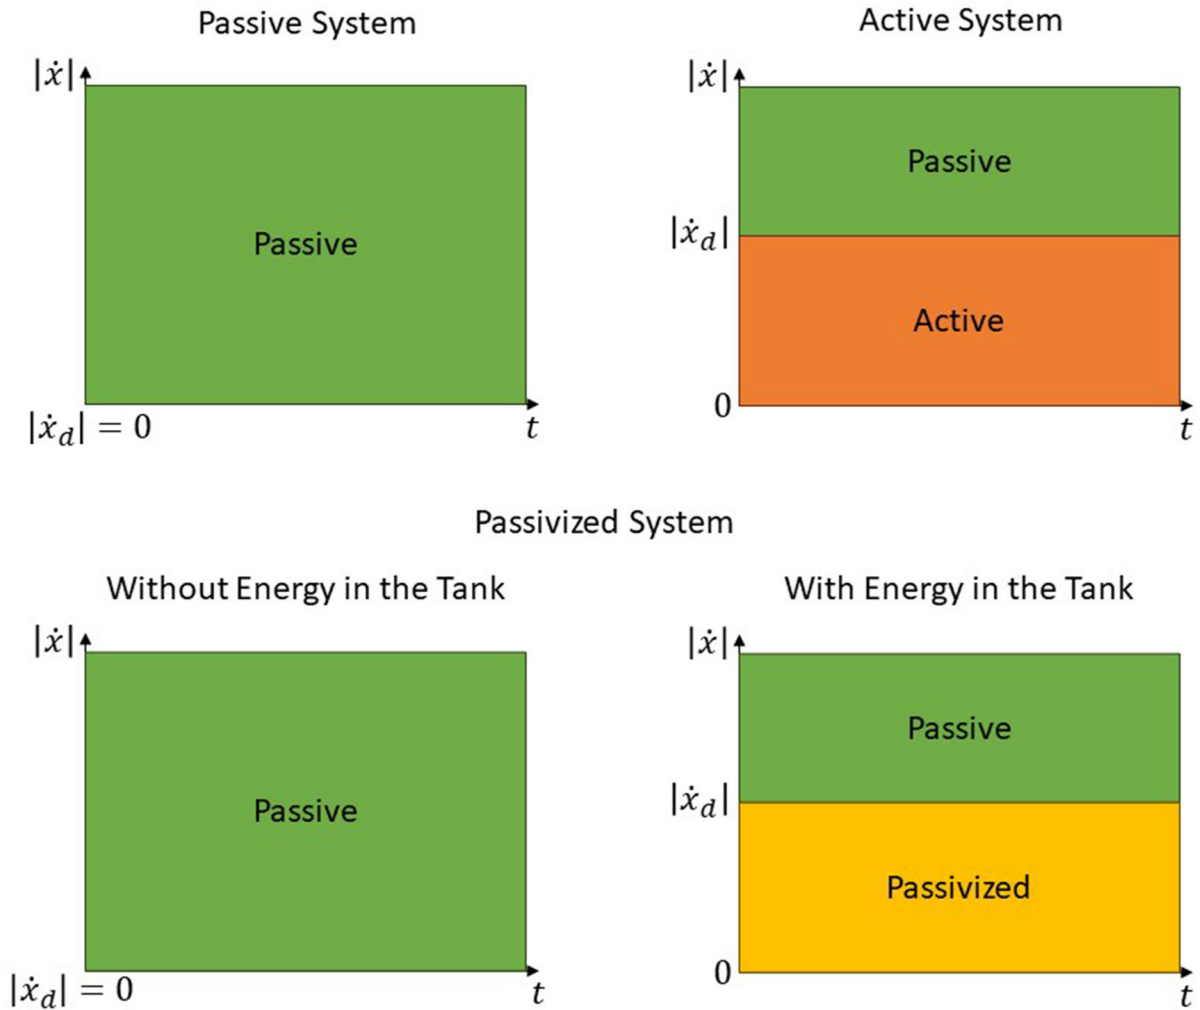


Fig. 2 Difference between the passive, active, and passivized system based on the evolution of the absolute value of the velocity error (\dot{x}). A passive system is always passive because the magnitude of the system velocity is always greater than the desired velocity $\dot{x}_d = 0$. An active system is active when $|\dot{x}| < |\dot{x}_d| \neq 0$,

and it is passive when $|\dot{x}| > |\dot{x}_d| \neq 0$. A passivized system behaves as an active system when there is energy in the tank and a passive one when it runs out of stored energy. Thus, when the tank is out of energy, a passivized controller compromises tracking accuracy to retain stability

the controller and the online adaptability; it is independent of the chosen impedance. For further details about the properties and stability of the fractal attractor, the reader can refer to [11].

For each DoF in the task-space, the FIC is given in Algorithm 1. The control torques ($\vec{\tau}_{ctr}$) can be calculated from $\vec{h}_e \in \mathbb{R}^6$ using (2). Differently from the FIC control scheme introduced in [11] on a sharp force/torque saturation, this manuscript introduces a more versatile force profile. The new force profile allows to independently tune the linear, nonlinear and

saturation behaviours of the controller wrench, making it easier to tune the controller for different tasks.

2.4 Sigmoidal force profile for fractal impedance

The FIC relies on a stiffness profile. The profile proposed in [11] results in fast changes in stiffness, and only allows limited task-dependent tuning of the profile. The sole requirement for guaranteeing the FIC stability is that the force profile is a continuous upper-




Fig. 3 **a** Attractor’s phase portrait—fractal impedance controller using the proposed spring (2.4), which, differently from the linear spring, has a compliant region near the fixed point in the desired state. The trajectories beyond $\tilde{x} > 0.3$ have been omitted **b** Proposed force profile. **c** Energy profile associated with the proposed force profile

Algorithm 1 Mono-dimensional FIC

input : Convergence/Divergence, \tilde{x} , \tilde{x}_{\max}

output: h_c

1 **if** diverging from x_d **then**

2 | $h_c = F_K(\tilde{x}) = K_d(\tilde{x})\tilde{x}$

3 **else**

4 | $h_c = K_c(\tilde{x}_{\max})(0.5\tilde{x}_{\max} - \tilde{x})$

where:

$K_d(\tilde{x}) = \frac{F_K}{\tilde{x}}$ is the stiffness during divergence

$K_c(\tilde{x}) = \frac{4E_K(\tilde{x}_{\max})}{\tilde{x}_{\max}^2}$ is the stiffness during convergence

\tilde{x}_{\max} is the maximum error recorded at the end of the divergence phase

E_K is the energy associated with \tilde{x}_{\max}

bounded function [11]. Therefore, a sigmoidal force profile is proposed for an easier definition of the stiffness profile, allowing users to better adapt the robot impedance behaviour to the different tasks. Similarly to the profile proposed by [11], the sigmoidal profile is fully determined based on the maximum force (F_{Max}) to be exerted at a chosen position error ($|\tilde{x}| = \tilde{x}_b$). Here, the position error is defined as the difference between the desired end-effector pose and the current pose ($\tilde{x} = x_d - x$). An additional displacement parameter ($|\tilde{x}| = \tilde{x}_0$) is added. It describes the minimum displacement to activate the nonlinear impedance, as shown in Fig. 3b. The proposed force profile thus becomes:

$$F_K = \begin{cases} K_0\tilde{x}, & |\tilde{x}| < \tilde{x}_0 \\ \text{sgn}(\tilde{x})(\Delta F(1 - e^{-(|\tilde{x}|-\tilde{x}_0)/b}) + K_0\tilde{x}_0), & \tilde{x}_0 \leq |\tilde{x}| < \tilde{x}_b \\ \text{sgn}(\tilde{x})F_{\text{Max}}, & \text{Otherwise} \end{cases} \quad (6)$$

where K_0 is the constant stiffness, $b = (\tilde{x}_b - \tilde{x}_0)/S$ is the characteristic length, S determines the shape of the sigmoid curve and $\Delta F = (F_{\text{Max}} - K_0\tilde{x}_0)$. In this work, $S = 20$ ensures force saturation before \tilde{x}_b .

The proposed force profile can further be associated with an energy (Fig. 4b) that is an unbounded Lipschitz function. It, therefore, respects the requirement for Lyapunov’s stability by the fractal attractor controller [11]. For the proposed force profile, this becomes:

$$E_K = \begin{cases} 0.5K_0\tilde{x}^2, & |\tilde{x}| < \tilde{x}_0 \\ F_{\text{Max}}|\tilde{x}| - F_{\text{Max}}\tilde{x}_0 + (K_0\tilde{x}_0^2)/2 - (1 - e^{-(|\tilde{x}|-\tilde{x}_0)/b})b\Delta F, & \tilde{x}_0 \leq |\tilde{x}| < \tilde{x}_b \\ F_{\text{Max}}|\tilde{x}| - F_{\text{Max}}\tilde{x}_0 + (K_0\tilde{x}_0^2)/2 - (1 - e^{-(\tilde{x}_b-\tilde{x}_0)/b})b\Delta F, & \text{Otherwise} \end{cases} \quad (7)$$

2.5 Controller superimposition for the control of redundant robots

The proposed method aims to achieve a hierarchy of tasks by using virtual soft mechanical constraints generated by the superimposition of task-space controllers that drive the robot to assume a commanded reference

posture. The benefit of using impedance controllers based on fractal impedance is that their passivity allows for superimposition without compromising overall system stability. Therefore, the total torque vector ($\vec{\tau}_{\text{tot}}$) can be computed by the superimposition of controllers as:

$$\vec{\tau}_{\text{tot}} = \sum_{i=1}^n \mathbf{J}_i^T \vec{\mathbf{h}}_e^i \quad (8)$$

where \mathbf{J}_i and $\vec{\mathbf{h}}_e^i$ are the Jacobian and the wrench generated by the impedance controller of the i^{th} -link, as depicted in Fig. 1 for our experiments that has a controller on the end-effector and another in the elbow joint. As a consequence of Eq. 8, the controller command does not contain any inversion of the Jacobian matrix, which renders the proposed method robust to drops in Jacobian rank (i.e. singularities). This is the main difference between a superimposition of FIC controllers and the artificial potential field, which would require performing a nonlinear inverse optimisation to verify stability.

2.6 Postural optimisation

For the scope of this paper, an optimisation-based inverse kinematics algorithm is used to obtain the reference configuration (*postural optimisation*). In particular, the optimisation identifies the joint configuration \mathbf{q} which minimises the cost to achieve the desired target task-space 6-DoF pose as obtained from forward kinematics using EXOTica [20]. Redundancy resolution either uses regularisation to a nominal target configuration or an energy minimisation proxy. First-order methods are leveraged to solve the unconstrained optimisation problem and use forward kinematics to subsequently extract the task-space references for the individual superimposed controllers. In place of this postural optimisation, more comprehensive planners and frameworks could be used to provide and update the reference configurations.

3 Evaluation

The proposed method is evaluated using a 7-DoF torque-controlled Kuka LWR3+² manipulator in both

² A description of the robot can be found at web.inf.ed.ac.uk/slmc/lab-resources/hardware.

simulation and hardware experiments. The same superimposition of two task-space controllers has been applied in both cases: A 6-DoF FIC controller at the end-effector (7th link) and a 3-DoF FIC controller at the elbow (4th link of the Kuka URDF) for postural control.

To generate pose references for each of the controllers, an optimisation to obtain a configuration satisfying the end-effector reference is used. Specifically, one-step variant of approximate inference control (AICO) is used [21]. Note while the end-effector pose reference can be passed in directly to the end-effector controller, a postural optimisation is used in this case to obtain a pose reference for the null-space or additional superimposed controllers. The reference pose for each of the controllers is extracted using forward kinematics.

3.1 Reference trajectories

The figure-of-8 (i.e. lemniscate) trajectory has been selected to show the dynamic behaviour of the robot. The trajectory is composed of two orthogonal sinusoidal trajectories. The vertical trajectory has an amplitude of 0.2m and the transverse trajectory amplitude is 0.1 m. The figure-of-8 trajectory is particularly demanding due to its multiple velocity inversions and wide joint movement range, thus introducing high variability of both the Jacobian and the inertial behaviour of the robot. The figure-of-8 reference motion is tested in both simulation and hardware experiments.

The hardware experiments also test a sinusoidal trajectory with an amplitude of 0.5 m and velocities up to about 0.7 m s⁻¹. The straight-line experiment enabled us to test interaction and robustness at higher speeds.

3.2 Simulation experiments

The robot simulations are in the Gazebo physics simulator and apply the superimposition of passive task-space controllers control scheme directly without compensating for gravity, Coriolis, or other dynamic effects (in contrast to [11]), i.e. as a model-free compliant controller. The simulation experiments compare nominal tracking performance with an interaction scenario where an unsensed environment obstacle has been introduced, cf. Fig. 5.




Fig. 4 Diagram of the proposed method. \vec{x}_d^e is the desired pose of the end-effector. \vec{x}_d^{el} is the desired pose of the elbow. \vec{x}^e is the current pose of the end-effector. \vec{x}^{el} is the pose of the elbow. \vec{h}_e^e is the desired wrench of the end-effector. \vec{h}_e^{el} is the desired wrench

of the elbow. $\vec{\tau}^e$ are the joint torques generated by the controller at the end-effector for the primary task. $\vec{\tau}^{el}$ are the joint torques generated by the controller at the elbow for the secondary task. $\vec{\tau}_{tot}$ is the torque command sent to the robot. \vec{q} are the joint positions




Fig. 5 Interaction with an unsensed object in simulation. The planned reference motion is shown with an alpha value while the robot state in simulation is shown in solid, with the trace of

the end-effector for both overlaid. Note the controller remains in contact and follows the surface of the obstacle

3.3 Hardware experiments

The hardware experiments are conducted with a Kuka LWR3+ robot. The manipulator is controlled using the Fast Research Interface (FRI) at 333.3 Hz in *joint impedance* mode with all gains set to zero to enable feed-forward torque control. Note unlike the simulation experiments the Kuka's built-in controller compensates for dynamic effects and gravity. On the real robot, the tracking of the figure-of-8 trajectory has also been tested with and without a human operator applying random perturbations. The values used in the controller for the simulation and the experiments are reported in Table 1. It shall also be noted that during the experiment, it has been kept the minimum set of controlled DoF required to fully control the 3-DoF of redundancy for the assigned tasks, being the task invariant to the configuration of the 7th DoF due to the symmetric geometry of the end-effector in the manipulator (Fig. 1).

4 Results

To complement the plots in this section, the reader is recommended to watch the supplementary video demonstrating the tracking and interaction both in simulation and hardware experiments. It has also been included a sequence demonstrating the safe behaviour of the controller during calibration of the FIC parameters given in Table 1.

The simulation results are shown in Fig. 6 for the free motion, and in Fig. 7 for the interaction behaviour with the obstacle. They show that the robot can be successfully controlled without dynamic compensation and that it can achieve dexterous dynamic behaviours. The tracking root mean square error (RMSE) at the end-effector is recorded without interaction as $RMSE_x = 5.6$ mm, $RMSE_y = 4.6$ mm, and $RMSE_z = 6.1$ mm. For simulation with interaction with an obstacle: $RMSE_x = 13.2$ mm, $RMSE_y = 4.2$ mm, and $RMSE_z = 5.5$ mm. That is, the tracking performance

Table 1 FIC parameters used for simulation and hardware experiments

	Elbow controller			End-effector controller					
	x	y	z	x	y	z	θ_x	θ_y	θ_z
<i>Simulations</i>									
x_0 (m or rad)	.01	.01	.01	.005	.005	.005	.057	.057	.057
x_b (m or rad)	.011	.011	.011	.006	.006	.006	.63	.63	.63
F_{Max} (N or N m)	50	50	50	150	150	150	100	100	100
K_0 (N m ⁻¹ or N m rad ⁻¹)	5000	5000	5000	5000	5000	5000	100	100	100
<i>Robot</i>									
x_0 (m or rad)	.01	.01	.01	.06	.06	.06	NA	NA	NA
x_b (m or rad)	.011	.011	.011	.08	.08	.08	NA	NA	NA
F_{Max} (N or N m)	10	10	10	100	100	100	NA	NA	NA
K_0 (N m ⁻¹ or N m rad ⁻¹)	630	630	630	800	800	800	NA	NA	NA

XYZ Euler angles are used for describing the system orientation. The NA indicates controllers that have not been used in the robot experiment because the orientation controllers are redundant in the considered robot due to the end-effector symmetry




Fig. 6 Simulation: Fast figure-of-8, no interaction. In **a** and **b**, reference and executed trajectories are blue and orange, respectively. Note as the simulation experiments are fully model-free (i.e. without gravity or nonlinear effects compensation), the plotted torques are higher compared with Fig. 8 as the FIC automati-

cally compensates for gravity and dynamic effects. This demonstrates that the proposed controller can handle gravity/nonlinear effects compensation as external disturbances F_{ext} without the need for a dynamics model




Fig. 7 Simulation: Fast figure-of-8, interaction with unsensed obstacle (cf. Fig. 5). In **a** and **b**, reference and executed trajectories are blue and orange, respectively. The stack of passive task-space controllers compensates for the obstacle using the

redundancy task by smoothly sliding across its surface, cf. (b). Note that despite the interaction there are no torque spikes beyond the system limits or instability, cf. (c)




Fig. 8 Hardware experiments: Fast figure-of-8, no interaction. In **a** and **b**, reference and executed trajectories are blue and orange, respectively. The controller tracks the reference closely for both the primary and redundancy tasks

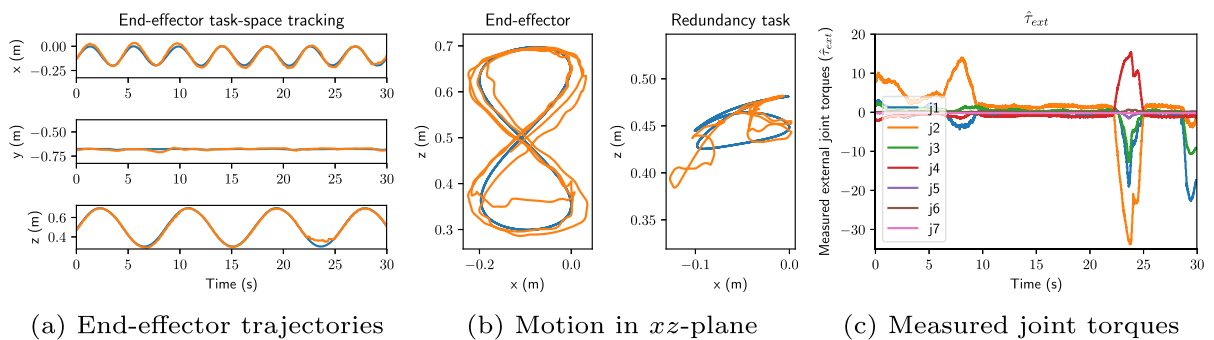


Fig. 9 Hardware experiments: Fast figure-of-8, with random interaction. In **a** and **b**, reference and executed trajectories are blue and orange, respectively. Note the significant external joint torque forces compared with Fig. 8 due to the interaction/perturbations

degrades in one dimension impacted by the obstacle (x), while being virtually unaffected in y and z .

The experimental data for the hardware experiments of the figure-of-8 trajectory are shown in Figs. 8 and 9. The errors recorded during free motion are: $\text{RMSE}_x = 7.6$ mm, $\text{RMSE}_y = 1.5$ mm, and $\text{RMSE}_z = 8.6$ mm. The perturbations do not affect the tracking performance at the end-effector task, but they are fully compensated by the deflection from the redundancy task target at the elbow joint, as shown in Fig. 9b. The RMSE during interaction is $\text{RMSE}_x = 20.3$ mm, $\text{RMSE}_y = 7.9$ mm, and $\text{RMSE}_z = 9.1$ mm.

The results for the straight-line trajectory experiment (Figs. 10 and 11) show the ability of the controller to complete the task and reject perturbations by reducing tracking on the redundancy task. The errors are $\text{RMSE}_x = 6.1$ mm, $\text{RMSE}_y = 4.6$ mm, and $\text{RMSE}_z = 6.7$ mm. The RMSE for interaction is $\text{RMSE}_x = 9.6$ mm, $\text{RMSE}_y = 7.7$ mm, and $\text{RMSE}_z = 7.3$ mm.

It shall also be remarked how the robot remained safe to interact with despite the high joint feed-forward

torques involved in the motions, which reached ≈ 30 N m for both the figure-of-8 trajectory and the linear trajectory.

5 Discussion

The results show the proposed method enables an intrinsically stable control framework for redundant robots which does not rely on inverse dynamics and projection matrices. The proposed method is robust to unknown environmental interactions and singularities, where safe means that the robot does not show erratic behaviours even while perturbed or when there is a sudden change in the desired task (e.g. sudden acceleration/deceleration). The RMSE data show how the robot keeps the minimum tracking accuracy (i.e. maximum error) contained under 1 cm for the unperturbed experiments, which is in line with the task requirement set in the controller parameters $\tilde{x}_b = 1.1$ cm. These results are in line with the results obtained in [11], and they are lower than other impedance controller frame-

works that usually have an error of a few centimetres, as can be seen in Fig. 7 in [6] and Fig. 5 in [8].

The introduction of significant perturbation degrades the minimum tracking accuracy to 2 cm, but the controller remains stable, and it is able to recover once the perturbation ends. It shall be remembered that the trade-off between accuracy and robustness is a known trade-off in interaction control frameworks [22]. While admittance control may provide better accuracy, it requires accurate knowledge of interaction force intensity and direction in all the points of contact with the environment. On the other hand, impedance controllers provide better safety of interaction but sacrifice tracking accuracy in favour of compliance [2,22]. Variable impedance controllers have been proposed as a solution to this dilemma, but the stability requirements on the impedance updates are often very stringent and difficult to retain under highly variable environmental conditions [19,23]. The proposed framework provides the best of both worlds providing good tracking accuracy while retaining the robustness typical of impedance controllers; furthermore, it enables online adjustment of the impedance profiles [11].

The data also confirm the hypothesis that redundant robot interaction behaviour can be accurately defined without any a priori knowledge of the system dynamics model, being Eq. 8 the control command. The simulation experiments show that the tracking performance in this work is similar to the results reported in [11] that relied on a compensation of the robot dynamics and the use of a null-space controller. This latest result is particularly important because robots' mechanical properties such as inertia and joints friction matrices are often difficult to retrieve and highly unreliable [4]. The knowledge of the actuation characteristics and the kinematic structure are still necessary for the implementation of the proposed method. However, they are both normally accurate and easier to obtain if not available.

The FIC generates an asymptotically stable potential field around the target state that enables the direct superimposition of multiple controllers without compromising the system stability. The controller superimposition generates a force field that acts as a trade-off cost function determining the preferred path of motion in the robot's configuration space. The force upper bound of the controllers guarantees that the loss of accuracy in the main task is contained \tilde{x}_b until the condition is compatible with the mechanical characteristics of

the system. Especially, if it is considered that the proposed method is a compliant postural controller, where accurate tracking is subordinate to the robustness of interaction. In other words, the controller stabilises the robot around the desired posture relying on the nonlinear stiffness profile to compensate for its nonlinear and environmental interaction, sacrificing the redundancy task before degrading the end-effector task beyond the selected accuracy. This is confirmed by both the simulation and the experimental data, showing how the FIC tries and successfully keeps the accuracy under 0.11 cm in unperturbed conditions. The data also describe that the controller fully sacrifices the redundancy task in the attempt to retain the same accuracy while experiencing external perturbations that exceed its mechanical limits, as shown in Figs. 7, 9, and 11.

The data show that the proposed method can achieve a highly dynamic interaction using variable impedance at the controller level. The FIC also enables online tuning of the impedance behaviour and is robust to reduce bandwidth in the feedback signals [11], allowing to switch from rigid to soft behaviours seemingly. Nevertheless, the performances are strictly related to the physical hardware capabilities, and a higher band-pass in the mechanics of the robot implies a higher stiffness to mass ratio. Therefore, it will be interesting to study these capabilities by deploying in hardware equipped with VSA to conduct a systematic experiment on these properties. Furthermore, it may also enable the switching from a model-based (e.g. the one proposed in [3]) to a data-driven control of their nonlinear actuators' dynamics. It is worth also noting that, at the current stage, it is impossible to compare the proposed method with VSA control architectures due to their different hardware requirements. Nevertheless, it can be said that both of them achieve nonlinear impedance behaviour and robustness to highly dynamic interactions. The fractal impedance controller hardware requirements are less stringent, and it has a simpler formulation. The results presented in [3] indicate that the DLR robot can achieve a stiffer behaviour. However, it is impossible to discriminate if they are connected solely to hardware superiority or if there is also a controller component in play.

The superimposition of task-space controllers also opens new possibilities for improving controllability and dexterity for compliant robots, developing human motor control theory, and robust control architecture for learning algorithms. In fact, the dynamics of soft




Fig. 10 Hardware experiments: Straight line reference, no interaction: The robot tracks the task-space references for both the primary and redundancy task closely joint configurations are sac-

rified. In **a** and **b**, reference and executed trajectories are blue and orange, respectively

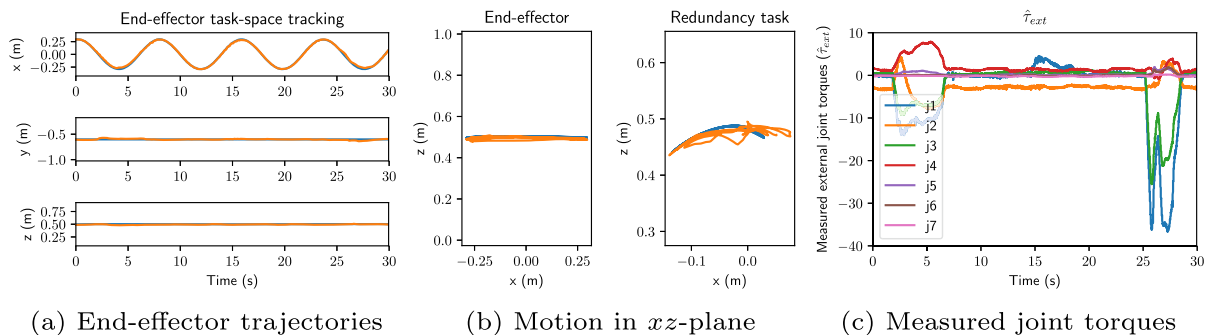


Fig. 11 Hardware experiments: Straight line reference (controlled via a slider in the graphic interface in the interval 0.4 m s^{-1} to 0.7 m s^{-1}), with interaction: Robot joint trajectories and effort during interaction experiments. The robot joints diverged significantly from the planned configuration, and reached posi-

tion across singularities without any evident impact on the robot performance in terms of robustness and stability. In **a** and **b**, reference and executed trajectories are blue and orange, respectively

robots are even more challenging to model than rigid dynamics [4], as the dynamic modelling of robots is founded on the assumption of rigidity [1]. Regarding human motor control having a framework that enables robustness and dexterity of interaction will overcome the current limitation of the passive motion paradigm (PMP) model, which still relies on inverse matrices for trajectory optimisation [24]. Finally, the learning algorithms are currently facing the challenges of performing a system identification to guarantee the stability of the learned behaviour. The proposed method removes this challenge, and the learning component can focus on learning how to synchronise the task-space controller to maximise the efficiency and dexterity of the robot.

6 Conclusions

The experimental results confirmed the proposed hypothesis that it is possible to control a redundant robot with a superposition of task-space controllers. This

approach renders the architecture intrinsically robust to singularity and fully passive, which guarantees stability. It is important to properly balance the strength of the controllers to guarantee that the redundancy tasks do not interfere with the end-effector controller, which may result challenging under certain conditions. Nevertheless, unbalanced controllers may interfere with the action efficacy but not with the robustness and stability of interaction.

The proposed framework does not require any a priori knowledge of the system dynamics parameters (i.e. inertia, friction, and gravity). It is suited for applications where the stability of interaction with unpredictable environments is more critical than the tracking accuracy. Future work will focus on improving the coordination among the redundancy task-space controllers to improve the tracking accuracy of the end-effector task when coupling effects are introduced.

Author contributions CT theorised the approach. All the authors were involved in the formulation of the control architecture and experimental design. WM has developed code for the control architecture, and CT has helped troubleshoot the implementation. CT and WM have run the simulations and performed the data analysis. CT and WM have written the manuscript. WW, SV, and MM have revised it SV and MM provided the funding.

Funding This work has been supported by EPSRC UK RAI Hub ORCA (EP/R026173/1), National Centre for Nuclear Robotics (EPR02572X/1), THING project (ICT-2017-1), and MEMMO (780684) in the EU Horizon 2020.

Data Availability Statement The code used for the experiments is available at: https://github.com/TRL-Sussex/superimposition_of_passive_taskspace_controllers.git. Further inquiries about the data can be directed to the corresponding authors.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Siciliano, B., Sciavicco, L., Villani, L., Oriolo, G.: Robotics: Modelling, Planning and Control. Springer, London (2010)
- Xin, G., Wolfslag, W., Lin, H.-C., Tiseo, C., Mistry, M.: An optimization-based locomotion controller for quadruped robots leveraging cartesian impedance control. *Front. Robot. AI* **7**, 48 (2020)
- Keppler, M., Lakatos, D., Ott, C., Albu-Schäffer, A.: Elastic structure preserving (ESP) control for compliantly actuated robots. *IEEE Trans. Rob.* **34**(2), 317–335 (2018)
- Bruder, D., Remy, C.D., Vasudevan, R.: Nonlinear system identification of soft robot dynamics using koopman operator theory. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 6244–6250 (2019). <https://doi.org/10.1109/ICRA.2019.8793766>
- Nakanishi, J., Rawlik, K., Vijayakumar, S.: Stiffness and temporal optimization in periodic movements: an optimal control approach. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 718–724 (2011). IEEE
- Dietrich, A., Ott, C., Albu-Schäffer, A.: An overview of null space projections for redundant, torque-controlled robots. *Int. J. Robot. Res.* **34**(11), 1385–1400 (2015)
- Moura, J., Ivan, V., Erden, M.S., Vijayakumar, S.: Equivalence of the projected forward dynamics and the dynamically consistent inverse solution. In: Proceedings of Robotics: Science and Systems, Freiburg im Breisgau, Germany (2019). <https://doi.org/10.15607/RSS.2019.XV.036>
- Dietrich, A., Ott, C., Stramigioli, S.: Passivation of projection-based null space compliance control via energy tanks. *IEEE Robot. Automat. Lett.* **1**(1), 184–191 (2016). <https://doi.org/10.1109/LRA.2015.2512937>
- Vigoriti, F., Ruggiero, F., Lippiello, V., Villani, L.: Control of redundant robot arms with null-space compliance and singularity-free orientation representation. *Robot. Auton. Syst.* **100**, 186–193 (2018)
- Sadeghian, H., Villani, L., Keshmiri, M., Siciliano, B.: Task-space control of robot manipulators with null-space compliance. *IEEE Trans. Rob.* **30**(2), 493–506 (2013)
- Babarahmati, K.K., Tiseo, C., Smith, J., Lin, H.C., Erden, M.S., Mistry, M.: Fractal impedance for passive controllers: a framework for interaction robotics. *Nonlinear Dyn.* (2022)
- Mansard, N., Khatib, O., Kheddar, A.: A unified approach to integrate unilateral constraints in the stack of tasks. *IEEE Trans. Rob.* **25**(3), 670–685 (2009). <https://doi.org/10.1109/TRO.2009.2020345>
- Khatib, O.: The operational space framework. *JSME Int. J. C Dyn. Control Robot. Des. Manuf.* **36**(3), 277–287 (1993)
- Tchoń, K., Ratajczak, J.: Dynamically consistent Jacobian inverse for non-holonomic robotic systems. *Nonlinear Dyn.* **85**, 107–122 (2016)
- Escande, A., Mansard, N., Wieber, P.-B.: Hierarchical quadratic programming: Fast online humanoid-robot motion generation. *Int. J. Robot. Res.* **33**(7), 1006–1028 (2014)
- Hogan, N.: Impedance control: an approach to manipulation: part ii - implementation. *J. Dyn. Syst. Meas. Contr.* **107**(1), 8–16 (1985)
- Secchi, C., Stramigioli, S., Fantuzzi, C.: Control of Interactive Robotic Interfaces: a port-Hamiltonian approach vol. 29. Springer (2007)
- Strogatz, S.H.: Nonlinear dynamics and chaos with student solutions manual: with applications to physics, biology, chemistry, and engineering. CRC press (2018)
- Angelini, F., Xin, G., Wolfslag, W.J., Tiseo, C., Mistry, M., Garabini, M., Bicchi, A., Vijayakumar, S.: Online optimal impedance planning for legged robots. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 6028–6035 (2019). IEEE
- Ivan, V., Yang, Y., Merkt, W., Camilleri, M.P., Vijayakumar, S.: In: Koubaa, A. (ed.) EXOTica: An Extensible Optimization Toolset for Prototyping and Benchmarking Motion Planning and Control, pp. 211–240. Springer, Cham (2019). https://doi.org/10.1007/978-3-319-91590-6_7
- Toussaint, M.: Robot trajectory optimization using approximate inference. In: International Conference on Machine Learning (ICML), pp. 1049–1056 (2009). ACM
- Ott, C., Mukherjee, R., Nakamura, Y.: Unified impedance and admittance control. In: 2010 IEEE International Conference on Robotics and Automation, pp. 554–561 (2010). IEEE

23. Li, Y., Ganesh, G., Jarrassé, N., Haddadin, S., Albu-Schaeffer, A., Burdet, E.: Force, impedance, and trajectory learning for contact tooling and haptic identification. *IEEE Trans. Rob.* **34**(5), 1170–1182 (2018)
24. Tiseo, C., Veluvolu, K., Ang, W.: The bipedal saddle space: modelling and validation . *Bioinspirat. Biomimet.* **14**(1), 015001 (2018)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.