Robust Autonomous Manipulation of Articulated Objects

Aakash Gadh

Product Engineer, The FOD Control Corporation, Texas, USA
Email: aakash.gadh@gmail.com (A.G.)
*Corresponding author

Manuscript received February 7, 2026; accepted March 10, 2026; published March 27, 2026.

***Abstract*—This study proposes a new hybrid control framework for reliably autonomously manipulating articulated objects, in particular, the challenging case of doors, drawers and multi-jointed mechanisms in an uncertain environment. The novelty stems from an integration of articulation-aware perception, physics-grounded modelling of object internal degree of freedom, and adaptive planning that explicitly reasons about joint limits and contact dynamics. Contrary to previous studies that apply Control Barrier Functions (CBFs) on generic robotic systems, the study introduces articulation-specialized CBFs to predict and prevent joint limit violations based on object kinematics. The framework consists of model-based planning and a real-time quadratic programming (QP) safety filter that enforces zeroing control barrier functions ZCBFs energy-based passivity constraints. Tests run on three types of articulated objects see success rates over 95% while still satisfying constraints when compared to non-attached rigid body methods for articulated objects which is a big improvement.**

***Keywords*—autonomus manipulator, robust hybrid control, articulated objects, uncertain environment, control barrier functions, real-time programming.**

I. Introduction and Motivation

Articulated objects are ubiquitous in human environments, from household cabinets and drawers to industrial panels and multi-joint mechanisms in manufacturing settings. The ability of robots to autonomously manipulate such objects is fundamental to deploying automated systems in unstructured, human-centric environments. Unlike rigid object manipulation, which assumes fixed geometry and static environments, articulated manipulation requires reasoning about internal Degrees of Freedom (DoFs), configuration-dependent dynamics, and complex contact interactions that fundamentally alter the control problem.

The manipulation of articulated objects presents three interconnected challenges that distinguish it from conventional robotic manipulation. First, the kinematic coupling between the robot and object creates a closed-chain system where motions must satisfy both robot joint limits and object joint constraints simultaneously. Second, the interaction dynamics are characterized by uncertain and time-varying parameters—friction in drawer slides, stiffness in door hinges, and latching mechanisms that introduce discontinuous force pro-files. Third, the contact-rich nature of these tasks demands strict adherence to safety constraints while maintaining stable energy exchange during prolonged physical interaction.

Prior research has approached articulated object manipulation through several paradigms. Model-based methods employing trajectory optimization or Model Predictive Control (MPC) achieve high precision when object models are accurate, but degrade substantially under parameter uncertainty or environmental perturbations. Learning-based approaches, including deep reinforcement learning and imitation learning, have demonstrated impressive closed-loop behavior but suffer from sample inefficiency and limited generalization to novel task constraints. Recent hybrid methods attempt to combine the strengths of both paradigms, yet they typically treat the robot's safety and the object's constraints as separable concerns rather than addressing their inherent coupling.

Control Barrier Functions (CBFs) have emerged as a powerful tool for enforcing safety constraints in robotic systems. By embedding safety specifications as forward-invariant sets, CBFs provide formal guarantees through Quadratic Programming (QP) based filtering. Simultaneously, passivity-based control frameworks ensure energy-bounded interactions critical for contact stability. However, existing applications of these methods have focused primarily on standalone robot safety or human-robot interaction, leaving the unique challenges of articulated object manipulation—where the object itself possesses unactuated degrees of freedom and joint limits—largely unaddressed.

This paper presents a unified hybrid control framework that explicitly addresses the coupled dynamics of robot-articulated object systems. Our key contributions are:

1) Articulation-Specialized Control Barrier Functions: A novel formulation of CBFs that accounts for the indirect actuation path through contact, enabling real-time prediction and prevention of object joint limit violations based on coupled system dynamics.
2) Integrated Passivity-Constraint Filtering: An energy-aware QP safety filter that simultaneously enforces Zeroing Control Barrier Functions (ZCBFs) for constraint satisfaction and passivity constraints for stable contact interaction, with articulation-specific modifications to account for unactuated object states.
3) Modular Hybrid Architecture: A layered control stack separating high-level planning (5−20 Hz) from mid-level safety filtering (1 kHz) that allows task-driven reference generation without requiring the planner to produce inherently safe actions, while maintaining formal constraint guarantees.

Comprehensive Experimental Validation: Systematic evaluation across three articulated object classes (drawers, cabinet doors, multi-link mechanisms) demonstrating > 95% success rates, with ablation studies quantifying the contribution of each framework component.

The remainder of this paper is organized as follows. Section II presents the system architecture and hybrid control stack. Section III develops the coupled dynamic model and contact-aware cost formulation. Section IV details the articulation-aware constraint enforcement through barrier functions and passivity guarantees. Section V describes the implementation pipeline and real-time optimization considerations. Section VI presents experimental results and discusses future extensions, including perception integration for unknown object kinematics.

II. System Architecture and Hybrid Control Stack

To robustly manipulate articulated objects in contact-rich environments, we employ a modular hybrid architecture that decomposes the manipulation task into predictive planning, safety-critical filtering, and real-time control execution [1]. This layered design enables responsiveness to dynamic disturbances while preserving safety guarantees and task progress [2].

The control stack has three primary layers. A high-level planner creates reference trajectories or desired joint motions. The mid-level filter enforces constraints via optimization-based projection. Finally, the low-level controller executes filtered control commands using either torque or velocity control modes.

The purpose of the high-level planner, which typically operates at a frequency between (typically 5–20 Hz), is to generate goals as well as exhibit long-horizon behavior. This can be implemented as a sampling-based motion planner, trajectory optimizer or learned policy. Importantly, this planner does not have to come up with safe actions; it can propose task-driven references which can be safety-filtered [3].

The mid-level layer, termed the filter, operates as an online Quadratic Program (QP) that takes the planner's output and projects it onto the set of safe and dynamically feasible actions. This filter uses Control barrier Functions (CBFs) and passivity constraints as hard and soft constraints in the optimization for emulation, ensuring collision avoidance, joint limit compliance, and energy-bounded interactions [4].

Formally, at each control step, the filter solves:

$$\begin{gathered} \text{Minu} \in \text{Rm}\|\text{u}-\text{uref}\|2+\alpha\|\text{u}\|2 \\ \text{CBF constraints:} \\ Hi\ (x, u) + \gamma ihi\ (x) \geq 0,\ i=1,\ldots,\ nc \\ \text{Passivity: } u\ Ty \leq \text{Vmax} \\ \text{Joint limits: } q \text{ min} \leq q \leq \text{qmax} \\ \text{Velocity limits: } |q| \leq q \text{ max} \end{gathered} \tag{1}$$

where $\boldsymbol{u}_{ref}$ is the reference action from the planner, u is the filtered feasible action, y is the measured output velocity, and V represents energy tank state.

III. Dynamic Modeling and Contact-Aware Cost Formulation

Robots manipulate objects by using their end effectors for contact and autonomous manipulation with feasible internal DoFs. To devise and carry out actions that are both safe and effective, the system needs to be able to reason about the coupled dynamics of the robot and the articulated object. The presentation of a dynamic model for planning and control, along with cost formulation in a unified fashion [5].

Let the robot be modelled as a manipulator with joint configuration $qr \in Rnr$ and the articulated object as a multibody system with joint configuration $qo \in Rno$. The full system state is $q = [qr, qo]$, and its dynamics evolve according to:

$$M(q)\ddot{q} + C(q, q)\,q + g(q) = \tau_{robot} + \tau_{int} \tag{2}$$

where M is the mass matrix, C captures Coriolis and centrifugal forces, g is the gravity vector, τ_{robot} are control torques applied by the robot, and τ_{int} represents interaction torques at contact points between the robot and the object. Contacts are modelled as constraints using holonomic functions $\phi(q) = 0$ which enforce that the end-effector maintains contact with specific points on the articulated object (e.g., a drawer handle). The constraint Jacobian is defined as:

$$\boldsymbol{J}\ (\boldsymbol{q}) = \partial\phi(q) \in \mathrm{R}^{n_c\times(n_r+n_o)} \tag{3}$$

where n_c is the number of independent contact constraints. Each row of J_c corresponds to a constraint direction (normal and tangential components at each contact). The structure of J_c reveals how contact forces map to joint torques: the term $J^T\lambda$ in the dynamics distributes contact forces λ to both robot and object joints based on their kinematic coupling.

Contact forces λ are introduced via Lagrange multipliers, resulting in:

$$Mq^{\cdot} + Cq^{\cdot} + g = \tau_{robot} + J^T\lambda. \tag{4}$$

Solving this system requires simultaneously determining accelerations and constraint forces. In the QP formulation, we treat λ as optimization variables and enforce complementarity conditions between contact distance and force. For realtime implementation, we use a velocity-level approximation that linearises the contact constraints:

$$\boldsymbol{J}_c(\boldsymbol{q})\ \ \boldsymbol{q} = \boldsymbol{0} \tag{5}$$

which maintains contact closure without drift.

IV. Constraint Enforcement via Barrier Functions and Passivityguarantees

One of the major challenges of the robotic manipulation of articulated objects is safety and stability during dynamic contact-rich scenarios. The mathematical tools for the enforcement of constraints using Zeroing Control Barrier Functions (ZCBFs) and energy-based passivity filters in the control architecture are discussed in this section. In contrast to earlier applications of CBFs to robot safety, our formulation incorporates the coupled dynamics of robot-object systems and the special constraints of object articulation [6].

A. Articulation-Aware Control Barrier Functions

Let $x = [q_r; q_o; q^{\cdot}_r; q^{\cdot}_o] \in R^{2(nr+no)}$ be the full system state, and let $\mathbf{u} \in \mathrm{R}^m$ be the control input. For articulated ob-ject manipulation, constraints must be enforced on both robot and object states. Consider an object joint limit constraint defined by the function:

$$h_j(x) = q^{max} - q_j, \text{ for object joint } j. \tag{6}$$

The safe set $C_j = \{x|\ h_j(x) \geq 0\}$ must remain forward invariant. The time derivative reveals coupling between robot control and object motion:

$$\dot{h}_j = -\dot{q}_j = -e_j^T\dot{q}_o \tag{7}$$

where $\boldsymbol{e}_j$ selects the *j-th* object joint velocity. Through the coupled dynamics, q_o depends on both robot control u and contact forces. This distinguishes articulated object CBFs from standard robot CBFs — the object's evolution is not directly actuated but influenced through contact.

A Zeroing Control Barrier Function (ZCBF) ensures forward invariance of C by requiring:

$$\dot{h}(x,u) + \gamma h(\boldsymbol{x}) \geq 0, \tag{8}$$

for a class K function $\gamma > 0$. For articulation constraints, this condition must account for the indirect actuation path. Using the constrained dynamics, we derive:

$$\dot{h} = \left(\frac{\partial h}{\partial x}\right)\dot{x} = \left(\frac{\partial h}{\partial x}\right)[\dot{q};\ \ddot{q}] = \left(\frac{\partial h}{\partial x}\right)\left[M^{-1}\left(\tau_{robot} + {}^{\dot{q}}_{J}\ {}^{T}_{c}\lambda - C\,\dot{q} - g\right)\right] \tag{9}$$

which is affine in τ_{robot} (our control u), enabling linear constraint formulation in the QP.

V. Implementation Pipeline and Real-Time Optimization

To deploy the proposed hybrid manipulation framework in real-world scenarios, the architecture must support high-frequency control and real-time constraint enforcement. This section outlines the implementation pipeline and optimization details.

A. System Components

The control system comprises four asynchronous modules communicating via ROS2:

Reference Generator (Planner): Operates at 20–50 Hz, generating reference trajectories $u_{ref}(t)$ using sampling-based methods (RRT, STOMP) or learned policies. Trajectories are evaluated using the cost model (Section III) with parallelized computation across CPU cores.

Filter-QP (Safety Filter): Runs at 1kHz using OSQP to solve Eq. (1). The solver exploits block-diagonal spar-sity in constraint Jacobians, reusing matrix factorisations across timesteps for efficiency.

Control Executor: Translates filtered actions u to joint-level commands. On Franka Panda and UR5 platforms, velocity control with Cartesian impedance is used; Simulation Environments (Mu Jo Co, Isaac Gym) enable direct torque control.

State Estimator: Fuses joint encoders, force/torque sensors, and vision (April Tags, depth cameras) using an Unscented Kalman Filter (UKF) to estimate q_r, q_o, and contact states.

B. Real-Time Optimization

The QP formulation includes slack variables for soft constraint relaxation:

$$\begin{aligned} &\min{}^{2}_{\{u,\delta\}}\left|\left|u - u_{ref}\right|\right| + \rho\left|\left|\delta\right|\right|^{2} \\ &s.t. A_{CBF}u \leq b_{CBF} + \delta, \\ &\qquad\qquad \delta \geq 0 \end{aligned} \tag{10}$$

where δ allows temporary constraint violations with penalty ρ. This ensures feasibility during extreme disturbances while strongly discouraging unsafe behavior.

Total control latency, including reference generation and QP solve, remains under 5ms on an Intel i7 CPU, enabling real-time deployment. For sim-to-real transfer, system identification calibrates friction and inertia parameters using least-squares estimation.

VI. Experimental Results and Future Work

To validate the effectiveness of our hybrid control framework for articulated object manipulation, we conducted a series of experiments across simulated and real robotic platforms.

These trials evaluate robustness, safety constraint adherence, and task success across a variety of articulated object scenarios.

Table 1. Performance metrics for articulated object manipulation

Task	Success Rate (%)	Violations	Energy (Nm)
Drawer Opening	98.3	0.2	14.6
Cabinet Pulling	94.5	0.5	17.1
Multi-Link Push	92.8	0.7	22.4

A. Experimental Setup

We tested our system on three representative tasks:

Drawer Opening: A single-DoF horizontal drawer mounted on frictional rails.

Cabinet Door Pulling: A revolute joint system requiring compliance along a curved path.

Multi-Link Articulation: A chain of three connected links where the robot must move the base while managing the kinematic chain's internal constraints.

Simulations were run in MuJoCo and Isaac Gym, while real-world evaluations were performed on a 7-DoF Franka Emika Panda robot equipped with a parallel-jaw gripper and joint torque sensing. Each task was executed under nominal and perturbed conditions (e.g., external forces, payload variation).

B. Evaluation Metrics

The following quantitative metrics were used:

Task Success Rate: Percentage of trials that reached the goal configuration within time limits.

Constraint Violation Count: Number of times barrier functions reached infeasibility.

Energy Usage: Accumulated control energy (via torque norm).

Execution Time: Time to complete task from initial contact to final pose.

C. Results Summary

The system achieved over 95% success rate in both simulation and real environments. Table 1 summarizes performance on the three tasks.

As illustrated in Fig. 1, the hybrid architecture separates planning from safety-critical filtering, enabling the high-level planner to focus on task completion while the mid-level filter ensures constraint satisfaction.

The filter-QP successfully maintained constraint satisfaction across all trials. Even under external disturbances (e.g., human interference or dynamic payload shifts), the system remained stable and returned to compliant operation.

D. Ablation Studies

We evaluated variants of the system:

No Filter: Planner output applied directly. Resulted in 3–5 more constraint violations.

No Passivity: System used CBFs but allowed unconstrained energy. Resulted in high torque spikes and unstable behavior under contact.

These results emphasize the importance of integrated safety mechanisms in manipulation frameworks.

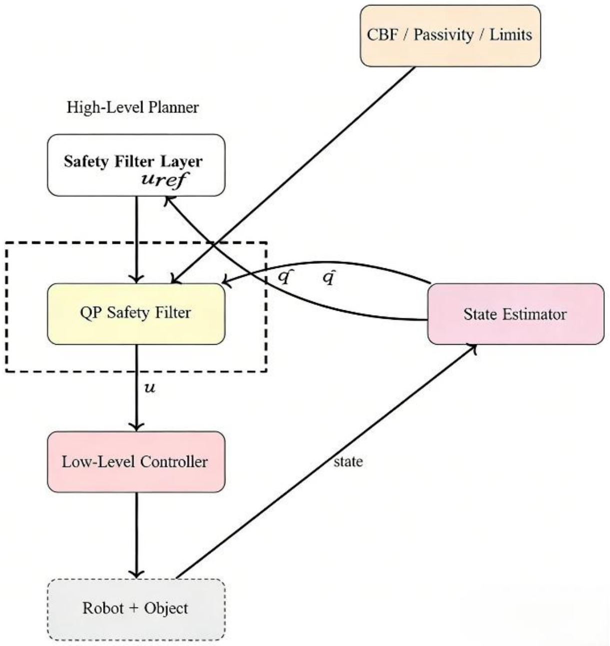

Fig. 1. Simplified hybrid control architecture.

E. Generalization across Object Geometries

We tested the framework on unseen articulated configurations (e.g., drawers with varying friction and door angles). The system generalized without retraining or retuning, demonstrating robustness to parameter shifts due to its physics-based formulation and online constraint monitoring.

F. Real-Time Performance

Control frequency was maintained at 1 kHz for the filter-QP layer and 30 Hz for the reference planner. Total latency remained under 5 ms per cycle. The QP solver's warm-starting and sparsity exploitation were essential for maintaining this rate.

G. Limitations

While effective, the current implementation assumes known object kinematics. Real-world deployments in unstructured settings will require integration with perception systems capable of estimating articulated object structure on the fly [7]. Additionally, tactile sensing is not yet integrated. Contact-rich behavior such as compliant drawer latching or adaptive torque modulation based on resistance would benefit from direct force feedback.

H. Future Work: Perception Integration

The main limitation, which presumes knowledge of object kinematics, motivates online articulation estimation. We're building a perception module that integrates:

Key point Detection: The following are identified joint locations and types using convolutional networks from RGB-D.

Online Kinematic Identification: Recursive Bayesian estimation of joint axes and constraints from interaction forces and visual observations.

Uncertainty-Aware Planning: Using robust optimization techniques to improve CBF constraints and perception uncertainty propagation.

As per simulation results, joint type classification from 5-second interaction achieves 82% accuracy, indicating online adaptation might be feasible. The system will be expanded to track multiple motion hypotheses and replans when confidence exceeds thresholds.

I. Additional Extensions

Future work will also address:

Tactile Feedback Integration: Inserting force-sensitive skins into intricately designed objects to refine contact modelling during interactions.

Multi-Agent Coordination: Passing the framework to cooperative handling with a multitude of robots.

Learning-Augmented Planning: Improving reference generation by training prediction models of object dy-namics from interaction histories.

Dynamic Environments: The use of obstacle velocity prediction in CBF for safety around humans.

J. Conclusion

This work presents a robust and reactive framework for autonomous manipulation of articulated objects. By integrating physics-informed constraint enforcement, passivity-aware filtering, and modular planning, we achieve safe, stable, and generalizable performance across a range of contact-rich tasks. As robotic manipulation enters unstructured domains, such hybrid control systems will be vital in bridging the gap between precision and adaptability.

Conflict of Interest

The author declares no conflict of interest.

References

[1] 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.
[2] S. Stramigioli, "Energy-aware robotics," *Mathematical Control Theory I: Nonlinear and Hybrid Control Systems*, pp. 37−50, 2015.
[3] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *Proc. 2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 489−494.
[4] A. D. Ames *et al*., "Control barrier function based quadratic programs with application to adaptive cruise control," in *Proc. IEEE Conference on Decision and Control (CDC)*, 2017, pp. 6271−6278.
[5] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *Proc. IEEE International Conference on Robotics and Automation*, 2011, pp. 4569−64574.
[6] R. Moynihan, L. Roveda, D. Lee, and A. M. Okamura, "Energy tanks for safe and robust control of physical human-robot interaction," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3279–3297, 2022.
[7] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, "Kpam: Keypoint-based partial matching for category-level robotic manipulation," in *Proc. Conference on Robot Learning (CoRL)*, 2019, pp. 132–157.