

Multi-Quadruped Cooperative Object Transport: Learning Decentralized Pinch-Lift-Move

Bikram Pandit, Aayam Kumar Shrestha, and Alan Fern

Abstract—We study decentralized cooperative transport using teams of N -quadruped robots with arm that must pinch, lift, and move ungraspable objects through physical contact alone. Unlike prior work that relies on rigid mechanical coupling between robots and objects, we address the more challenging setting where mechanically independent robots must coordinate through contact forces alone without any communication or centralized control. To this end, we employ a hierarchical policy architecture that separates base locomotion from arm control, and propose a constellation reward formulation that unifies position and orientation tracking to enforce rigid contact behavior. The key insight is encouraging robots to behave as if rigidly connected to the object through careful reward design and training curriculum rather than explicit mechanical constraints. Our approach enables coordination through shared policy parameters and implicit synchronization cues – scaling to arbitrary team sizes without retraining. We show extensive simulation experiments to demonstrate robust transport across 2-10 robots on diverse object geometries and masses, along with sim2real transfer results on lightweight objects.

I. INTRODUCTION

Object transport in local environments is central to construction, logistics, and emergency response. Legged robots are suited for such domains, as they can traverse terrain where wheeled platforms fail [1], [2]. Yet many real-world objects—boxes, furniture, construction materials, logs, irregular containers—cannot be easily grasped or rigidly attached. Recent work in learning-based loco-manipulation shows that a single humanoid or quadruped can lift and transport boxes through whole-body coordination [3]–[5], but only for payloads within a single robot’s capacity. Larger/heavier objects require coordinated teams manipulating through sustained contact, raising a key question: how can multiple robots maintain stable contact, distribute forces, and coordinate motion without communication or centralized control?

Existing decentralized approaches succeed mainly when robots are rigidly coupled to objects, where mechanical connections passively distribute forces and enforce coordination [6]–[8]. In contrast, manipulating ungraspable objects through contact alone is fundamentally harder: robots must establish and maintain contact, synchronize lifting, and stabilize payloads during motion—without mechanical guarantees or explicit communication. Our key insight is that coordination can still emerge if reward shaping encourages

robots to behave *as if rigidly attached* to the payload, even though they are only connected through pinching forces.

We propose **decPLM (Decentralized Pinch-Lift-Move)**, a hierarchical control framework with a constellation-based reward that enables quadruped–arm teams to learn robust pinch–lift–move behaviors without communication/centralization. By aligning base and payload motion while maintaining relative poses, N -robot teams achieve the benefits of rigid coupling while retaining the flexibility of contact-based manipulation. *Our main contributions are:* (1) a constellation reward unifying position and orientation alignment to enforce rigid-like contact; (2) decentralized training that induces implicit synchronization without communication; (3) evidence that policies trained with small teams ($N=2$) generalize to larger teams ($N=10$) and diverse payloads. We further analyze when continuous payload tracking is necessary, and demonstrate sim-to-real transfer on physical quadruped–arm robots.

II. RELATED WORK

We review 3 types of prior work: (1) rigid mechanical coupling between robots and objects, (2) grasping-based manipulation with prehensile contacts, and (3) contact-only manipulation that typically involves planar pushing. In contrast, our approach combines decentralized coordination with contact-only lifting and transport of ungraspable objects.

Rigidly Coupled Loco-Manipulation. Many successful systems assume mechanical constraints between robots and payloads—through closed chains, fixtures, or passive arms—which greatly simplifies force distribution and stability control [9]–[13]. Centralized approaches use model predictive control and hierarchical planning for mechanically coupled systems [6], [7], while passive-arm methods provide stabilization through mechanical compliance [8]. Recent work has shown that decentralized control can succeed when robots remain rigidly attached to payloads [14]. In contrast, our approach eliminates mechanical coupling entirely: robots must establish and maintain contact purely through learned policies, changing the coordination problem from force distribution to contact formation and maintenance.

Grasping-Based and Tethered Systems. A second category relies on prehensile grasping or non-rigid connections like cables. Multi-robot grasp planning enables coordinated assembly and regrasping but requires reliable graspable features [15]. Cable-towed systems coordinate teams through explicit link dynamics, with taut/slack mode switching that simplifies connection maintenance during transport [16], [17]. While some fully decentralized approaches exist [18],

Project page: <https://decplm.github.io>

*This work is supported by DARPA contract HR0011-24-9-0423 and an NVIDIA Academic Grant. We thank Mohit Gadde and Ashutosh Gupta for their important technical support.

All authors are with Collaborative Robotics and Intelligent Systems Institute (CoRIS), Oregon State University, Corvallis, Oregon, 97331, USA. Email: {panditb, aayam.shrestha, afern}@oregonstate.edu.

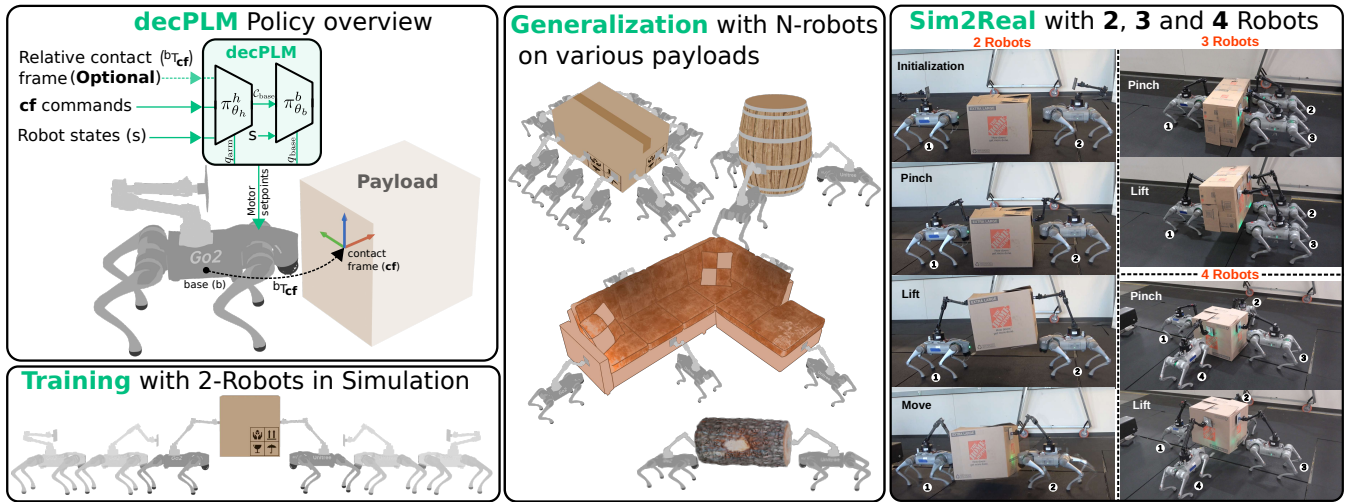


Fig. 1: Overview of **decPLM** (Decentralized Pinch-Lift-Move). **Top-Left:** Policy structure. Each robot runs the same decentralized policy that receives local proprioception (s), contact-frame command c_{cf} , and (optionally) contact frame pose in base frame ${}^bT_{cf}$, and outputs arm joint targets and base velocity commands. **Bottom-Left:** Training setup with two robots in simulation on a box. **Middle:** Generalization to larger teams and diverse payloads, including a box, a log, a barrel, and a couch. **Right:** Sim2Real Demonstration with 2, 3 and 4 Robots, each running **decPLM(const⁺, cf^{init})** policy independently, without any inter-robot communication.

they still assume reliable attachment mechanisms. These approaches sidestep the core challenge we address: maintaining coordinated contact without mechanical fixtures or tethers.

Contact-Only, Decentralized Manipulation. Non-prehensile manipulation through contact alone has primarily focused on planar pushing of supported objects, avoiding the challenges of lifting and three-dimensional transport [19]–[21]. While Single-robot whole-body manipulation has achieved impressive force and position control [22]–[25], scaling to cooperative lifting remains unexplored. Moreover multi-agent reinforcement learning has developed relevant coordination mechanisms - hierarchical policies, shared representations, and scalable architectures [26]–[30] - but has not addressed the specific challenge of maintaining coordinated contact during cooperative lifting. Our approach fills this gap by enabling multiple robots to perform coordinated pinch-lift-move behaviors through contact alone, without mechanical coupling or communication.

III. PROBLEM FORMULATION

We consider decentralized cooperative transport where a team of N quadrupedal robots must transport a single rigid-body object, which we refer to as the *payload*. The payload is assumed to be ungraspable, that is, it cannot be rigidly attached to any robot and must instead be manipulated through physical contact at its surfaces using contact pads attached to each robot’s arm. The robots must coordinate to *pinch*, *lift*, and *move* the payload so that it tracks a commanded motion trajectory of the payload.

In this paper, we focus on the decentralized pinch-lift-move (PLM) control problem and assume a higher-level planner that provides: (1) contact frames on the payload, one per robot, that allows force closure; (2) a synchronized start signal to initiate the collaborative pinch and lift maneuver; and (3) velocity commands for the payload transformed into contact-frame coordinates for each robot. The planner

specifies only geometrically feasible contact locations and transport commands. Contact maintenance, force regulation, synchronization, and dynamic stability during transport are not enforced by the planner and must be achieved through decentralized control. We assume that each pinch-lift-move episode begins with robots positioned within a specified tolerance of their contact frames.

A. Technical Formulation

We assume a team of N quadruped robots, each equipped with a single arm whose end-effector is a flat contact pad. The N -robot PLM problem is specified through N contact frames $\{{}^bT_{cf}^{(r)}\}_{r=1}^N$, where each contact frame ${}^bT_{cf}^{(r)} \in SE(3)$ ¹ specifies a surface location on the payload (relative to the base) where robot r should engage its contact pad.

Once the payload has been lifted, its motion is driven by an instantaneous team-level payload transport command $\mathcal{C}_{pl} = (v_{pl}, \omega_{pl}, h_{pl})$, which specifies joystick-like inputs for planar velocity, yaw rate, and height of the payload’s root frame. Because the team-level command \mathcal{C}_{pl} depends on the payload’s geometry, which is not known to the robots, it is converted into per-robot contact-frame commands $\mathcal{C}_{cf}^{(r)} = (v_{cf}^{(r)}, \omega_{cf}^{(r)}, h_{cf}^{(r)})$. Specifically, for robot r with contact frame at fixed offset $p_{offset}^{(r)}$ from the payload root (where \mathcal{C}_{pl} is defined), we compute:

$$v_{cf}^{(r)} = v_{pl} + \omega_{pl} \times p_{offset}^{(r)}; \quad \omega_{cf}^{(r)} = \omega_{pl}; \quad h_{cf}^{(r)} = h_{pl} + p_{offset,z}^{(r)}$$

Importantly, the offsets $p_{offset}^{(r)}$ do not change with motion, so the transformation from \mathcal{C}_{pl} to $\mathcal{C}_{cf}^{(r)}$ is state-free, inexpensive to compute, and identical across time.

IV. SYSTEM OVERVIEW

Cooperative transport requires a team of N quadruped-arm robots to perform coordinated *pinch*, *lift*, and *move*

¹Represented in the observation as a 7D vector (position and quaternion).

operations. Fundamentally, this is a high-dimensional control problem where each robot must coordinate all leg and arm joints to maintain locomotion while regulating contact forces on the payload. To simplify this challenge, we adopt a hierarchical policy architecture for each robot that builds on well-established low-level locomotion control for quadrupeds.

Hierarchical Policy Architecture: We employ a two-level hierarchy that separates locomotion from manipulation.

(1) *Low-Level Policy.* The base locomotion policy $\pi_{\theta_b}^b$ provides a velocity-control abstraction to the high-level policy, where θ_b are the learned policy parameters. The input to $\pi_{\theta_b}^b$ is the proprioceptive state $s^{(r)}$ and desired base velocity commands $\mathcal{C}_{\text{base}}^{(r)} = (v_{\text{base}}^{(r)}, \omega_{\text{base}}^{(r)})$ of robot r , specifying the target linear and angular velocities. The output is the target joint positions $q_{\text{base}}^{(r)}$ for the quadruped base, which are then tracked by PD controllers to produce motor torques.

(2) *High-Level Policy.* As illustrated in Figure 1, the high-level policy $\pi_{\theta_h}^h$ coordinates task-level behavior. Its inputs when used for robot r include: proprioceptive state $s^{(r)}$, contact-frame pose ${}^bT_{\text{cf}}^{(r)}$ (optionally available) in the robot’s base frame, contact-frame command $\mathcal{C}_{\text{cf}}^{(r)}$, and a temporal synchronization signal $t^{\text{sync}} \in [0, 5]$ that provides weak coordination cues during initial contact formation. In particular, when the policy is started before pinching $t^{\text{sync}} = 0$ and then increments for 5 seconds (see Section V for details).

Due to the challenges of accurate tracking of the contact-frame pose ${}^bT_{\text{cf}}^{(r)}$, we consider two modes of operation: (1) the contact frame pose is continuously available, and (2) the contact frame pose is only available at the initial timestep and then masked thereafter. In the latter case, the robot must rely on proprioception to infer information about the contact frame for effective force regulation.

Decentralized Policy Execution: At execution time, each robot independently runs $\pi_{\theta_b}^b$ and $\pi_{\theta_h}^h$ at 50Hz using only local information. First, the high-level policy produces arm joint targets and a base command:

$$(q_{\text{arm}}^{(r)}, \mathcal{C}_{\text{base}}^{(r)}) = \pi_{\theta_h}^h \left(s^{(r)}, {}^bT_{\text{cf}}^{(r)}, \mathcal{C}_{\text{cf}}^{(r)}, t^{\text{sync}} \right).$$

Next, the low-level policy maps the base command into joint targets: $q_{\text{base}}^{(r)} = \pi_{\theta_b}^b \left(s^{(r)}, \mathcal{C}_{\text{base}}^{(r)} \right)$. Each robot then sends $(q_{\text{arm}}^{(r)}, q_{\text{base}}^{(r)})$ to the joint-level PD controllers.

Robots are initialized near their respective payload contact frames. A synchronized start signal sets $t^{\text{sync}} = 0$, initiating the pinch-lift phase. After $t^{\text{sync}} = 5$, this phase ends and $\pi_{\theta_h}^h$ begins receiving transport (move) commands.

This architecture learns complex coordination behaviors - contact placement, force regulation, and lift synchronization - in the high-level layer while ensuring reliable locomotion through the low-level controller. Coordination emerges through shared parameters, common objectives, and physical interaction, without requiring explicit communication.

V. TRAINING APPROACH

The base velocity controller $\pi_{\theta_b}^b$ is pre-trained using standard locomotion rewards and PPO [31], following established approaches for legged locomotion [32]. We use a

3-layer 128x128x128 MLP with ELU units as the policy architecture. Details can be found in the extended paper. The pre-trained $\pi_{\theta_b}^b$ is frozen during high-level policy training.

For the high-level policy $\pi_{\theta_h}^h$, we use Multi-Agent Proximal Policy Optimization (MAPPO) [33] within the Centralized Training and Decentralized Execution (CTDE) [34] paradigm. We also explored IPPO [35] training without centralized critics. In this setting, learning was unstable due to non-stationarity introduced by multiple simultaneously adapting agents, and policies failed to achieve synchronized lifting or stable transport. We therefore adopt CTDE to stabilize coordination learning. Training can be conducted with any number of robots, with update time scaling roughly linearly in N . Surprisingly, as our results show, training with only $N = 2$ robots is sufficient to produce policies that generalize well to much larger teams.

During training, we use asymmetric critic networks [36], which include privileged information that is unavailable to the actor. In particular, each robot maintains its own critic network that observes both local inputs available to the actor and the local inputs of all other robots to provide richer learning signals. At execution time, robots use only the shared policy $\pi_{\theta_h}^h$ with local observations, ensuring fully decentralized operation for an arbitrary number of robots.

The policy $\pi_{\theta_h}^h$ is implemented as a 2-layer 128x128 MLP that concatenates the local observation vector as indicated in IV with the previous action (policy output). This action history provides important temporal context for maintaining contact stability, particularly in the execution mode where contact frame pose information is masked out. All robots share policy parameters θ_h while maintaining separate critic networks during training. We assume homogeneous robots and shared policy parameters; heterogeneity in contact geometry arises through randomized contact frame assignments rather than distinct policy roles.

All policies are trained in the IsaacLab [37] simulator with 2048 parallel environments, each containing a team of N Unitree Go2 quadrupeds equipped with Unitree D1 arms that jointly manipulate a box. Training builds on the RSL-RL framework [38]. The following subsections detail two key design choices that critically affect learning outcomes: the reward function and the training curriculum.

A. Reward Design

Our reward function addresses the challenge of learning rigid contact behavior and coordinated motion without mechanical constraints. A novelty of our design is a *constellation reward* which provides a unified supervision for contact formation and velocity tracking. This is supplemented by more traditional task-specific tracking rewards and auxiliary penalties that stabilize training. Fig 2 shows the reward components and when they are activated across training phases and episode time (phase/time details in Sec. V-B).

1) *Constellation Reward:* The *constellation reward* provides a unified mechanism for enforcing both positional and orientational alignment between each robot and its designated contact frame. We define two constellations of

Commands Modality			Reward Terms	Episode Length (s)				
Pinch	Pinch + Lift	Pinch + Lift + Move		1	2	3	4	5-14
✓	✓	✓	Action Smoothness, Joint Vel/Acc/Torque, Box Acc.	■	■	■	■	■
✓	✓	✓	Arm Action Penalty	■	■	■	■	■
✓	✓	✓	Leg Action Penalty	■	■	■	■	■
✓	✓	✓	Outside Box Range Penalty	■	■	■	■	■
✓	✓	✓	Base Pose Tracking Reward	■	■	■	■	■
✓	✓	✓	Arm Constellation Reward	■	■	■	■	■
✓	✓	✓	Contact Reward	■	■	■	■	■
×	✓	✓	Box Height Tracking Reward	■	■	■	■	■
×	✓	✓	Box Levelness Penalty	■	■	■	■	■
×	×	✓	Box Velocity Tracking Reward	■	■	■	■	■
×	×	✓	Base Constellation Reward	■	■	■	■	■

Fig. 2: Reward activation schedule: Each row lists a reward term along with the command phase (pinch, lift, move) in which it is used. The blue bars show when, within an episode, each reward is active. This illustrates how different rewards are introduced and shaped across both training phases and episode time.




Fig. 3: Constellation reward illustration. Source points (blue) on the pad and base are matched to their corresponding target points (green). Dotted orange lines show the errors that the policy must minimize, with the two groups representing End-effector Contact Constellation and Base Tracking Constellation.

points, one anchored to the robot’s base and one to arm end-effector, each paired with a corresponding target constellation anchored to the payload. This constellation-based approach is inspired by point-set registration methods in computer vision [39]–[41], where both translation and rotation are implicitly encoded through the alignment of two point sets. The intuition is that by maintaining these, the team behaves as if the robots were rigidly attached to the payload.

Formally, we represent a constellation as a set of landmark points $P = \{p_1, \dots, p_K\}$ rigidly anchored to a reference frame, and define a corresponding target constellation $P^* = \{p_1^*, \dots, p_K^*\}$ anchored to the payload. The constellation distance is then defined as the mean squared error between matched landmarks:

$$d_{\text{con}}(P, P^*) = \frac{1}{K} \sum_{i=1}^K \|p_i - p_i^*\|_2^2,$$

which provides a smooth signal that couples position and orientation through the geometry of the point set.

This formulation allows us to construct constellations that enforce different aspects of the cooperative transport task. In the following, we define two specific choices: a constellation anchored to the robot base to encourage global alignment with the payload, and a constellation anchored to the end-effector to regulate local contact behavior.

End-Effector Contact Constellation. For each arm con-

tact pad and its designated contact frame ${}^bT_{\text{cf}}^{(r)}$, we define corresponding constellations. As illustrated in Fig 3, for robot r we spawn five *colinear* points $P_{\text{pad}}^{(r)}$ spanning 50cm uniformly along the arm’s contact pad normal and corresponding colinear points $P_{\text{cf}}^{(r)}$ spanning the contact frame normal. Matching these constellations enforces both position alignment and surface normal consistency, but allows flexibility in tangential orientations. The associated reward term is

$$R_{\text{contact}} = -d_{\text{con}}(P_{\text{pad}}^{(r)}, P_{\text{cf}}^{(r)})$$

which drives the contact pad and contact frame constellations to match, effectively emulating rigid attachment.

Base Tracking Constellation. To enforce the robot body to maintain consistent alignment with the payload during transport we define a constellation of three *non-collinear* points on the robot base $P_{\text{base}}^{(r)}$, chosen to resemble adjacent corners of a cube of size (10 × 10 × 10) cm (see Fig 3). The target constellation of three corresponding points $P_{\text{rigid}}^{(r)}$ are defined to represent the target positions of $P_{\text{base}}^{(r)}$ under the assumption that the robot is attached through a rigid kinematic chain through the arm. In particular, the points in $P_{\text{rigid}}^{(r)}$ are computed at each timestep to be the positions that the points in $P_{\text{base}}^{(r)}$ should be given the current payload target command. The base tracking reward is then defined as:

$$R_{\text{track}}^{(r)} = -d_{\text{con}}(P_{\text{base}}^{(r)}, P_{\text{rigid}}^{(r)})$$

This encourages base motion that is kinematically consistent with the commanded payload motion.

Together, these terms define the overall constellation reward for robot r (Fig. 3),

$$R_{\text{constellation}}^{(r)} = R_{\text{contact}}^{(r)} + R_{\text{track}}^{(r)}.$$

2) Task-Specific Tracking Rewards and Regularization:

Beyond the constellation formulation, which unifies position and orientation alignment into a single smooth objective, we include targeted rewards for specific aspects of the transport task. During contact establishment, a binary **contact reward** encourages initial surface engagement. For payload control, we track **height accuracy** during lift/move phases and **velocity tracking** during transport. These rewards are implemented via exponential kernels over the tracking errors.

To stabilize training and encourage smooth policies, we include regularization terms. As shown in Figure 2, these penalties are selectively activated based on training phase and episode timing. These include **torque and joint motion penalties** to prevent excessive actuator effort and high velocities, **action smoothness penalties** to discourage abrupt control changes, **phase-specific constraints** such as leg motion penalties during contact formation (seconds 0-4). Additional safety terms include **levelness penalties** to prevent payload tilting, **box acceleration penalties** during transport phases and outside-range penalties to constrain robots within assigned contact regions. These auxiliary terms provide essential regularization while the constellation and tracking rewards drive the primary learning objectives of




Fig. 4: Robot arrangements with different team sizes. Robots are evenly distributed around the box ($1.0 \times 1.5 \times 0.7$ m, mass 2 kg), as used in the experiments of Fig. 5. For the two-robot case, a smaller box ($0.5 \times 0.4 \times 0.7$ m) is used. Robot indexing, when referenced, starts from the *top-right corner* and increases clockwise around the box.

rigid contact and coordinated transport. Full details of all reward terms are in the extended paper.

B. Training Curriculum

The complexity of coordinated manipulation necessitates a carefully designed curriculum that progressively increases the difficulty of the learning problem. We structure this by carefully defining the generation of training episodes and incrementally adding command modalities to the training phases, as shown in Figure 2.

Training episode generation. Training uses a box payload with contact frames $\{{}^bT_{cf}^{(r)}\}$ randomly sampled on its surface subject to force-closure feasibility. Robots start near their assigned frames so that arms can reach without base motion, enabling early episodes to focus on contact formation.

To support robustness and sim-to-real transfer, we apply domain randomization over robot dynamics, contact parameters, payload mass (0.1-2 kg), external forces, initial poses, and sensor/action noise (details in the extended paper).

We train both controllers with persistent contact-frame observations and controllers where pose information is only given for the first 2 s. Directly training the latter proved difficult, so we begin with full 50 Hz updates and progressively anneal the update rate (25 Hz 5 Hz 0.25 Hz 0 Hz), which yields faster and more stable learning.

Training Phase 1: Contact Formation (Pinch). In this phase, we limit the command pool for the velocity and the height of the box to be 0, i.e., $\mathcal{C}_{obj} = (0, 0, 0, 0)$, hence the robots learn only to establish and maintain rigid contact at assigned contact frames. Here, the reward is designed in such a way that the robot is incentivized to reach the box contact surface within 4 seconds - hence, we limit the episode length to 7 seconds for this training phase. Moreover, for this phase, the box mass is set to a very large value (100 kg) to prevent premature motion, and leg movement is heavily penalized since locomotion is unnecessary. This phase focuses purely on arm control and contact regulation.

Training Phase 2: Coordinated Lifting (Pinch and Lift).

In this phase we extend the command pool to include payload height commands in addition to the pinch setting. Thus $\mathcal{C}_{obj} = (0, 0, 0, h_{obj})$ where $h_{obj} \geq 0$, so the policy must handle both cases: continuing contact maintenance with $h_{obj} = 0$ (as in Phase 1) and executing vertical lifting when $h_{obj} > 0$. Planar velocities remain zero. The box mass is randomized from 0.1 to 2 kg and the episode length is extended to 14 seconds. Robots receive vertical lifting commands and must coordinate to lift the payload while maintaining stable contact, with emphasis on smooth, synchronized motion without oscillation.

Training Phase 3: Full Transport (Pinch, Lift and Move). In this phase, we employ the complete command pool, which includes nonzero planar velocity as well as height commands, i.e., $\mathcal{C}_{obj} = (v_{obj}^x, v_{obj}^y, \omega_{obj}^z, h_{obj})$ with each component permitted to be either zero (as in earlier phases) or nonzero. This command modality requires full coordination of base locomotion and arm control. The episode length remains 14 seconds, but leg movement penalties are removed since robots must demonstrate pinch-lift-move capabilities.

VI. EXPERIMENTS AND RESULTS

We evaluate our approach through experiments in sim and demonstrations of sim-to-real transfer. Specifically, we examine: (1) the importance of the constellation reward, (2) robustness to limited contact frame information, (3) generalization across different team sizes and payload masses, and (4) transfer to out-of-distribution payloads.

A. Experimental Setup

We conduct experiments in simulation with teams of $N \in \{2, \dots, 10\}$ quadruped-arm robots cooperatively transporting payloads. Our primary experiments use a box as the payload with dimensions $1.0 \times 1.5 \times 0.7$ m and mass 2 kg with the exception of using a slightly smaller box ($0.5 \times 0.4 \times 0.7$ m) for $N = 2$. Figure 4 shows how different numbers of robots are positioned around the box.

Each evaluation episode is run for 14 seconds with each experimental configuration being run over 500 episodes. Episodes are initialized with robots near their assigned contact frames with small random offsets to test robustness under positioning uncertainty. Episodes then proceed with randomized commands: box velocities $v_x, v_y \sim \mathcal{U}[-0.4, 0.4]$ m/s, angular velocity $\omega_z \sim \mathcal{U}[-0.4, 0.4]$ rad/s, and height $h \sim \mathcal{U}[0.1, 0.3]$ m above ground.

To ensure robustness under realistic deployment conditions, all evaluations are performed with dynamics randomization and observation noise enabled. Gaussian noise is injected into all observation channels used by the high-level policy π_θ^h and low-level policy π_θ^b , including proprioceptive state $s^{(r)}$, contact-frame pose ${}^bT_{cf}^{(r)}$, and the temporal synchronization signal t^{sync} . These perturbations simulate sensing imperfections and environmental variability.

We evaluate performance using five metrics: (i) linear velocity tracking error (RMS error in x and y directions), (ii) angular velocity tracking error about z -axis, (iii) height tracking error, (iv) drop rate (percentage of episodes where the payload is dropped or fails to lift), and (v) robot failure rate (percentage of episodes with robot collisions or falls).




Fig. 5: We compare performance across different team sizes for four model variants: $\text{decPLM}(\text{const}^-, \text{cf}^{\text{init}})$, $\text{decPLM}(\text{const}^-, \text{cf}^+)$, $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$, and $\text{decPLM}(\text{const}^+, \text{cf}^+)$. Line style indicates access to contact-frame pose (solid = cf^+ vs. dotted = cf^{init}), marker shape indicates use of the constellation reward (circle = const^+ vs. cross = const^- , and colors correspond to the methods as listed. We also include a **1-robot baseline**, which does not carry any load but is provided as an idealized reference value. Finally, decPLM_{3r} is a variant trained with three robots instead of two. For details, see Subsection VI-C and VI-E.




Fig. 6: The figure shows the average normal force exerted on the box by each robot after the lifting phase for $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$. For details see second paragraph on Subsection VI-D.

B. Model Variants

To evaluate our key design choices, we focus on two questions: (1) does the constellation reward improve collaborative PLM performance, and (2) how important is continual access to contact frame information versus only at initialization? To answer these, we train four model variants that cross reward formulation (const^+ vs. const^-) with contact frame observability (cf^+ vs. cf^{init}), denoted as

$$\text{decPLM}(X, Y), \quad X \in \{\text{const}^+, \text{const}^-\}, \quad Y \in \{\text{cf}^+, \text{cf}^{\text{init}}\}.$$

By default all decPLM controllers are trained with 2 robots.

Since no prior methods address the decentralized PLM problem, we compare against three baselines. First, a 1-robot system with no payload provides an idealized reference for linear velocity tracking, angular velocity tracking, and fall rate. Second, we train with three robots, which we denote by decPLM_{3r} . In later experiments, we evaluate a centralized architecture to assess the impact of decentralization.

C. Constellation Reward Effectiveness

Figure 5 compares performance across the model variants for teams of 2-10 robots. The constellation-based methods substantially outperform traditional approaches. Both $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ and $\text{decPLM}(\text{const}^+, \text{cf}^+)$ achieve lower tracking errors and drop rates compared to their non-constellation counterparts across all team sizes.

Surprisingly, by comparing $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ to $\text{decPLM}(\text{const}^-, \text{cf}^+)$ we see that the constellation reward is more critical to performance than continuous

contact frame information. For example, with 5 robots, $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ is approximately 50% better at tracking velocity despite the contact pose information being masked. This indicates that a constellation reward structure enables robust contact maintenance even when explicit contact frame information is unavailable after the pinch phase.

D. Team Size Generalization and Force Distribution

Our policies, trained exclusively with 2-robot, generalize effectively to larger configurations without retraining. Figure 5 shows that performance improves consistently as team size increases: linear velocity errors decrease from 0.1 to 0.02 m/s when scaling from 2 to 10 robots, which is an 80% drop in error, while box drop rates fall from 5% to less than 1%. This improvement stems from better force distribution and increased redundancy provided by additional robots, as evidenced by our force below.

Comparing to $\text{decPLM}_{3r}(\text{const}^+, \text{cf}^+)$ (i.e. 3-robot training), we find 2-robot training performs as well or better across team sizes, with only a slight improvement in collision/fall rate for 3-robot training. One anomaly occurs where the 3-robot model performs worse on 3-robot teams, warranting further study, but overall there is little evidence that 3-robot training justifies its added cost. We also observe a small increase in linear velocity error (and, to a lesser extent, other metrics) when moving from even to odd team sizes. This is due to geometric asymmetry in contact placement: even teams allow approximately opposing force pairs, whereas odd teams require uneven moment balancing. The learned policy compensates by redistributing normal forces (Figure 6), but this introduces slightly higher tracking error.

We also analyze how contact forces are distributed across team members using $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$. As shown in Figure 6, in asymmetric team configurations - e.g., 3-robots with one robot on one side and 2 on the other (see Fig 4 for team arrangement) - the co-located robots together apply approximately the same total force as the single robot on the opposite side, indicating that the policy generalizes well by balancing forces despite geometric asymmetry. In general, we observe this asymmetric balancing behavior most clearly in smaller teams, but as the number of robots increases, individual force contributions become more uniform, indicating that larger teams distribute loads more evenly and rely less




Fig. 7: Performance of $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ and $\text{decPLM}(\text{const}^+, \text{cf}^+)$ on different team sizes under varying payload masses from 2 to 20 kg, with curve shading indicating mass progression from lighter (2 kg) to heavier (20 kg) payloads. For details see Subsection VI-E.



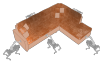
Scene	Robots	Methods	Linear vel. error	Ang vel. error	Height error	Payload Dropped %	Robot falling %
	2	$\text{decPLM}(\text{const}^+, \text{cf}^+)$	0.103	0.290	0.017	0.3	0.0
		$\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$	0.117	0.278	0.049	2.6	0.0
		$\text{decPLM}_{3r}(\text{const}^+, \text{cf}^+)$	0.134	0.279	0.071	12.2	0.0
		specialized-wooden-log	0.099	0.288	0.052	0.0	0.0
		centralized-2-robots	0.119	0.278	0.056	6.7	0.4
	3	$\text{decPLM}(\text{const}^+, \text{cf}^+)$	0.184	0.293	0.031	0.1	0.0
		$\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$	0.135	0.310	0.080	6.0	0.2
		$\text{decPLM}_{3r}(\text{const}^+, \text{cf}^+)$	0.116	0.278	0.046	1.0	0.0
		specialized-barrel	0.097	0.271	0.021	0.1	0.0
	5	$\text{decPLM}(\text{const}^+, \text{cf}^+)$	0.145	0.229	0.098	15.0	0.4
		$\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$	0.264	0.268	0.187	51.5	8.8
		$\text{decPLM}_{3r}(\text{const}^+, \text{cf}^+)$	0.187	0.225	0.171	43.6	1.8
		specialized-couch	0.178	0.256	0.091	23.0	2.4

TABLE I: Comparing performance on out-of-distribution payloads. See Subsection VI-F for details.

on compensating for geometric asymmetry.

E. Payload Mass Scaling

Figure 7 evaluates performance under varying payload masses from 2 to 20 kg, with curve shading indicating mass progression. For the $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ variant - without continuous pose information, additional robots generally improve performance by providing better force distribution for heavier payloads. However, a critical threshold emerges around 15 kg: beyond this point, even larger teams struggle with the substantial normal forces required for lifting, leading to elevated drop rates and robot failures.

Including pose information significantly improves performance: $\text{decPLM}(\text{const}^+, \text{cf}^+)$ achieves consistently lower velocity errors across all masses compared to $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$. This gap becomes dramatic at 15 kg, where $\text{decPLM}(\text{const}^+, \text{cf}^+)$ maintains an 18% drop rate and 5% failure rate, while $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ degrades sharply to 60% drop rate and 40% failure rate. These results indicate that for heavy payloads requiring precise force coordination, continuous knowledge of contact frame pose becomes crucial for our current approach to maintaining system stability and performance.

F. Out-of-Distribution Payloads and Alternative Baselines

We next evaluate generalization to out-of-distribution payloads and compare against alternative baselines. Test payloads include a wooden log (2 kg, 2 robots), a barrel (3 kg, 3 robots), and a couch (10 kg, 5 robots) (Table I).

We compare against two additional baselines trained with persistent observability (cf^+) and constellation reward

(const^+): (1) *centralized control*, with joint actions from global state (practical only up to 2 robots), and (2) task-specialized decentralized controllers trained directly on each payload rather than the default box. Results in Table I.

Geometric complexity increases information requirements: As payloads become more irregular, continuous contact-frame pose information ($\text{decPLM}(\text{const}^+, \text{cf}^+)$) becomes increasingly important, with the drop-rate gap relative to $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ growing from 2.3% (log) to 5.9% (barrel) and 33.5% (couch).

Constellation reward enables strong generalization: Our general approach matches or exceeds payload-specific controllers, demonstrating effective transfer from box training to diverse geometries at much lower computational cost.

Centralized control offers limited benefit: In the two-robot setting, centralized execution does not outperform decentralized control, likely because the added dimensionality complicates optimization while CTDE already supplies strong coordination signals.

Training team size does not determine generalization: Policies trained with 2 robots generally match or outperform those trained with 3, suggesting that domain randomization with 2 robots already provides sufficient diversity for broad generalization.

VII. REAL-WORLD DEMONSTRATION

We evaluated $\text{decPLM}(\text{const}^+, \text{cf}^{\text{init}})$ on teams of two, three and four Unitree Go2 robots with D1 arms, replacing the stock end-effectors with custom rubber-covered contact pads. For safety, experiments used lightweight boxes (1-2kg).

Real-world deployment revealed key challenges: unknown arm dynamics requiring manual gain tuning, joint encoder offsets needing external calibration, limited arm torque restricting payload capacity, and box deformability disrupting contact cues. Despite these issues, the policy executed coordinated pinch-lift-move sequences (Fig. 1), demonstrating that constellation-based coordination transfers to hardware, though with a substantial sim-to-real gap due to current hardware and calibration limitations.

VIII. SUMMARY AND FUTURE WORK

We introduced the cooperative pinch-lift-move (PLM) problem for ungraspable objects and presented **decPLM**, a decentralized learning framework that requires neither communication nor rigid coupling. By combining a locomotion-manipulation hierarchy with a constellation reward that

promotes rigid-like contact behavior, the method enables a shared policy to generalize across team sizes and payloads. Simulations show that this reward is critical for coordination and that policies trained with two robots transfer to larger teams and varied payloads. Real-world experiments with two to four robot teams validate the approach while exposing practical limitations, including torque limits, calibration, and object deformability. Future work will explore high-level planning, autonomous contact assignment, richer sensing, more varied terrain, and extension to other robot platforms.

REFERENCES

- [1] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020. 1
- [2] H. Duan, B. Pandit, M. S. Gadde, B. Van Marum, J. Dao, C. Kim, and A. Fern, "Learning vision-based bipedal locomotion for challenging terrain," in *International Conference on Robotics and Automation*, 2024. 1
- [3] Y. Zhang, Y. Yuan, P. Gurunath, I. Gupta, S. Omidshafiei, A.-a. Aghamohammadi, M. Vazquez-Chanlatte, L. Pedersen, T. He, and G. Shi, "Falcon: Learning force-adaptive humanoid loco-manipulation," *arXiv preprint arXiv:2505.06776*, 2025. 1
- [4] J. Dao, H. Duan, and A. Fern, "Sim-to-real learning for humanoid box loco-manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 16930–16936, IEEE, 2024. 1
- [5] M. Arnold, L. Hildebrandt, K. Janssen, E. Ongan, P. Bürge, A. G. Gábrriel, J. Kennedy, R. Lolla, Q. Oppliger, M. Schaaf, *et al.*, "Leva: A high-mobility logistic vehicle with legged suspension," in *International Conference on Robotics and Automation (ICRA)*, 2025. 1
- [6] X. Zhang, L. Yan, T. L. Lam, and S. Vijayakumar, "Task-space decomposed motion planning framework for multi-robot loco-manipulation," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8158–8164, IEEE, 2021. 1
- [7] F. De Vincenti and S. Coros, "Centralized model predictive control for collaborative loco-manipulation," in *RSS*, 2023. 1
- [8] G. Turrisi, L. Schulze, V. S. Medeiros, C. Semini, and V. Barasuol, "Pacc: A passive-arm approach for high-payload collaborative carrying with quadruped robots using model predictive control," in *International Conference on Intelligent Robots and Systems*, 2024. 1
- [9] F. Kennel-Maushart and S. Coros, "Payload-aware trajectory optimization for non-holonomic mobile multi-robot manipulation with tip-over avoidance," *IEEE Robotics and Automation Letters*, vol. 9, no. 9, pp. 7669–7676, 2024. 1
- [10] M. Sombolostan and Q. Nguyen, "Hierarchical adaptive control for collaborative manipulation of a rigid object by quadrupedal robots," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2752–2759, IEEE, 2023. 1
- [11] M. Sombolostan and Q. Nguyen, "Hierarchical adaptive motion planning with nonlinear model predictive control for safety-critical collaborative loco-manipulation," *arXiv preprint arXiv:2411.10699*, 2024. 1
- [12] R. Tallamraju, D. H. Salunkhe, S. Rajappa, A. Ahmad, K. Karlapalem, and S. V. Shah, "Motion planning for multi-mobile-manipulator payload transport systems," in *International Conference on Automation Science and Engineering*, 2019. 1
- [13] T. G. Sugar and V. Kumar, "Control of cooperating mobile manipulators," *IEEE Trans. on Robotics and Automation*, vol. 18, 2002. 1
- [14] B. Pandit, A. Gupta, M. S. Gadde, A. Johnson, A. Shrestha, H. Duan, J. Dao, and A. Fern, "Learning decentralized multi-biped control for payload transport," in *Conference on Robot Learning*, 2025. 1
- [15] M. Dogar, A. Spielberg, S. Baker, and D. Rus, "Multi-robot grasp planning for sequential assembly operations," *Autonomous Robots*, vol. 43, no. 3, pp. 649–664, 2019. 1
- [16] C. Yang, G. N. Sue, Z. Li, L. Yang, H. Shen, Y. Chi, A. Rai, J. Zeng, and K. Sreenath, "Collaborative navigation and manipulation of a cable-towed load by multiple quadrupedal robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10041–10048, 2022. 1
- [17] W.-T. Chen, M. Nguyen, Z. Li, G. N. Sue, and K. Sreenath, "Decentralized navigation of a cable-towed load using quadrupedal robot team via marl," *arXiv preprint arXiv:2503.18221*, 2025. 1
- [18] H. Farivarnejad, A. S. Lafmejani, and S. Berman, "Fully decentralized controller for multi-robot collective transport in space applications," in *2021 IEEE Aerospace conference (50100)*, pp. 1–9, IEEE, 2021. 1
- [19] C. Fan, S. Shirafuji, and J. Ota, "Modal planning for cooperative non-prehensile manipulation by mobile robots," *Applied Sciences*, vol. 9:3, 2019. 2
- [20] Y. Feng, C. Hong, Y. Niu, S. Liu, Y. Yang, and D. Zhao, "Learning multi-agent loco-manipulation for long-horizon quadrupedal pushing," in *International Conference on Robotics and Automation*, 2025. 2
- [21] I. Dadiotis, M. Mittal, N. Tsagarakis, and M. Hutter, "Dynamic object goal pushing with mobile manipulators through model-free constrained reinforcement learning," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 13363–13369, IEEE, 2025. 2
- [22] S. Jeon, M. Jung, S. Choi, B. Kim, and J. Hwangbo, "Learning whole-body manipulation for quadrupedal robot," *IEEE Robotics and Automation Letters*, vol. 9, no. 1, pp. 699–706, 2023. 2
- [23] Z. He, K. Lei, Y. Ze, K. Sreenath, Z. Li, and H. Xu, "Learning visual quadrupedal loco-manipulation from demonstrations," in *2024 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 9102–9109, IEEE, 2024. 2
- [24] M. Liu, Z. Chen, X. Cheng, Y. Ji, R.-Z. Qiu, R. Yang, and X. Wang, "Visual whole-body control for legged loco-manipulation," *arXiv preprint arXiv:2403.16967*, 2024. 2
- [25] T. Portela, G. B. Margolis, Y. Ji, and P. Agrawal, "Learning force control for legged manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2024. 2
- [26] O. Nachum, M. Ahn, H. Ponte, S. Gu, and V. Kumar, "Multi-agent manipulation via locomotion using hierarchical sim2real," *arXiv preprint arXiv:1908.05224*, 2019. 2
- [27] J. Wu, X. Sun, A. Zeng, S. Song, S. Rusinkiewicz, and T. Funkhouser, "Spatial intention maps for multi-agent mobile manipulation," in *International Conference on Robotics and Automation*, 2021. 2
- [28] T. An, J. Lee, M. Bjelonic, F. De Vincenti, and M. Hutter, "Solving multi-entity robotic problems using permutation invariant neural networks," *arXiv preprint arXiv:2402.18345*, 2024. 2
- [29] Z. Xiong, B. Chen, S. Huang, W.-W. Tu, Z. He, and Y. Gao, "Mqe: Unleashing the power of interaction with multi-agent quadruped environment," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5918–5924, IEEE, 2024. 2
- [30] T. An, J. Lee, M. Bjelonic, F. De Vincenti, and M. Hutter, "Scalable multi-robot cooperation for multi-goal tasks using reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 10:2, 2024. 2
- [31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017. 3
- [32] "Closing the Sim-to-Real Gap: Training Spot Quadruped Locomotion with NVIDIA Isaac Lab," June 2024. 3
- [33] C. Yu, A. Velu, E. Vinitzky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, "The surprising effectiveness of ppo in cooperative multi-agent games," *Advances in neural information processing systems*, 2022. 3
- [34] C. Amato, "An introduction to centralized training for decentralized execution in cooperative multi-agent reinforcement learning," *arXiv preprint arXiv:2409.03052*, 2024. 3
- [35] C. S. De Witt, T. Gupta, D. Makoviichuk, V. Makoviychuk, P. H. Torr, M. Sun, and S. Whiteson, "Is independent learning all you need in the starcraft multi-agent challenge?," *arXiv:2011.09533*, 2020. 3
- [36] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, "Asymmetric actor critic for image-based robot learning," *arXiv preprint arXiv:1710.06542*, 2017. 3
- [37] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, *et al.*, "Orbit: A unified simulation framework for interactive robot learning environments," *IEEE Robotics and Automation Letters*, vol. 8, no. 6, pp. 3740–3747, 2023. 3
- [38] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, "Learning to walk in minutes using massively parallel deep reinforcement learning," in *Conference on robot learning*, pp. 91–100, PMLR, 2022. 3
- [39] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor fusion IV: control paradigms and data structures*, vol. 1611, pp. 586–606, Spie, 1992. 4
- [40] Y. Chen and G. Medioni, "Object modelling by registration of multiple range images," *Image and vision computing*, vol. 10, no. 3, 1992. 4
- [41] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-d point sets," *IEEE Transactions on pattern analysis and machine intelligence*, no. 5, pp. 698–700, 1987. 4