2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Consensus Driven Dynamical Systems Control for Dual-Arm Handover

Debojit Das[1,*], Siddhi Jain[2], Rajesh Kumar[2] and Harish J. Palanthandalam-Madapusi[1]

***Abstract*— Robot bimanual handovers (transferring an object between two arms) require careful coordination of timing, motion, and obstacle avoidance. Efficient, human-like object transfer between cooperating robots demands both spatial and tight temporal coordination. Existing approaches treat these requirements in isolation or rely on pre-computed trajectories that fail when obstacles/disturbances appear, degrading performance, segmented behavior, and introducing desynchronization. This paper introduces a dynamical systems framework that transitions each arm from independent asynchronous motion to coupled synchronous coordination. In this context, coupling denotes both the spatial coordination of the arms and their temporal synchronization. The framework's coordination and synchrony are robust to obstacles/disturbances along its path. Experiments on an upper torso dual-arm platform and on traditional manipulators show seamless handovers that remain stable despite obstructions, always preserving spatial coordination and temporal synchrony. Supplementary videos and source code are available at: `https://debojit-d.github.io/consensus-handover/`**

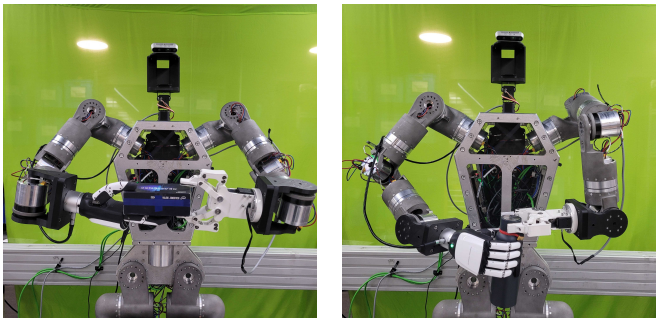

(a) Handing over a box. **(b)** Handing over a bottle.

Fig. 1: Real-world demonstrations of the proposed coordination framework on an upper torso dual-arm platform. The system successfully performs handovers with the proposed framework for two example cases.

I. INTRODUCTION

Robot handovers, the act of passing an object from one manipulator to another, are fundamental for efficient multi-robot cooperation [1]–[3]. Such tasks are ubiquitous, occurring in scenarios like transferring parts between manipulators in industrial assembly lines, repositioning items during warehouse picking (where one arm picks up an item and quickly transfers it to the other arm), or enabling humanoid robots to seamlessly shift objects between their hands without repositioning their entire body [4], [5]. Similarly, in bimanual sorting tasks, objects located outside the reachable workspace of one manipulator can be efficiently passed to the other arm, ensuring that productivity remains high [6]. Traditional robotic handover strategies typically segment the process into discrete stages, such as approach, grasp, transfer, and retreat, which inherently lack continuous temporal and spatial synchronization and coordination. [1], [7]. Although off-line-computed trajectories can enforce such behaviors, they suffer from rigidity due to pre-planning, failing to adapt in real-time to dynamic disturbances or obstacles in cluttered environments [2].

Building on this observation, we seek a unified control strategy that transforms segmented handover into a single, closed-loop behavior. Specifically, the problem statement we address is the following: How can two robotic manipulators perform an object handover in real-time such that they (i) converge to an identical handover pose at precisely the same instant, and (ii) maintain this temporal–spatial synchrony and coordination even when the trajectory of one or both the manipulators is locally deformed due to online obstacle-avoidance modulation or unexpected external disturbances?

Human hand-to-hand transfers naturally involve precise coordination and synchronization in both time and space between the giver and receiver [8]. Studies reveal that humans synchronize "what, when, and where" during the exchange, demonstrating exceptional proficiency in object transfers by continuously closing perceptual feedback loops at multiple levels [9], [10]. Gaze anticipates partner movements, grip force dynamically adjusts to object weight, and arm trajectories co-adapt in real-time [11], [12]. Inspired by this natural fluidity, our approach adopts real-time co-adaptation based not on visual feedback but on task-progress estimation, easily computed for robotic manipulators through accurate tracking progress of their end-effectors from initial to final states.

Emulating human-like fluidity in robot handovers, temporal synchronization and spatial coordination significantly improves upon traditional segmented approaches. Synchronizing the timing of the transfer eliminates idle hovering of one manipulator, thereby reducing energy consumption and unnecessary grip force adjustments. Furthermore, this temporal synchronization and spatial coordination creates a predictable transfer window, precisely defining when and where grippers should open and close, thus ensuring instantaneous exchanges.

Unlike traditional hand-over pipelines that let the receiver

[1]IITGN Robotics Laboratory, Department of Mechanical Engineering, Indian Institute of Technology Gandhinagar, Gujarat 382355, India.
[2]Addverb Technologies Private Limited, Noida, Uttar Pradesh 201305, India.
Emails: debojit.das@iitgn.ac.in, siddhi.jain@addverb.com, rajesh.kumar01@addverb.com, harish@iitgn.ac.in,
*Corresponding author.
This research was supported by Addverb Technologies Limited (Project No: RES/ATL/ME/P0079/2425/0045).

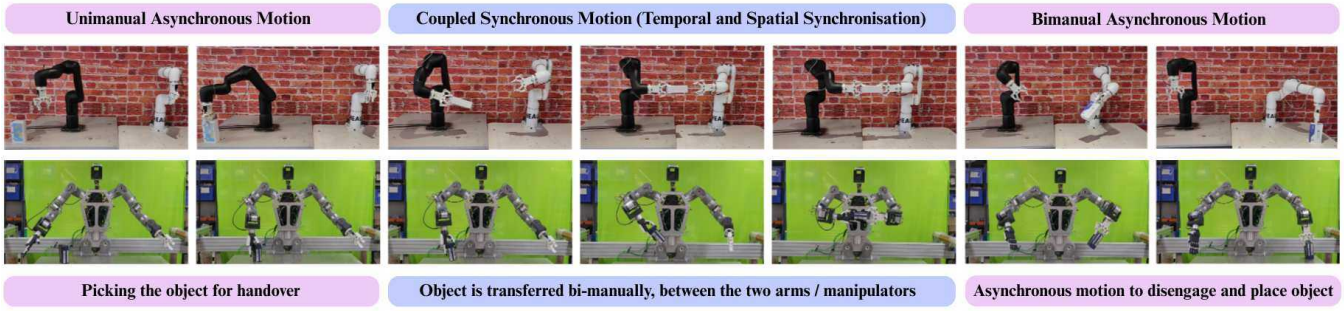

Fig. 2: Snapshots of the proposed handover and transport pipeline. The top row shows a rigid box being handed over between two single-arm manipulators in three phases: (i) unimanual pick, (ii) bimanual transfer, and (iii) asynchronous placement and disengagement. The bottom row demonstrates equivalent motions on an upper torso dual-arm platform.

"wait" at the transfer pose, our controller eliminates idle phases altogether. Each arm closes a local Dynamical System (DS) that is phase-locked through a single scalar clock. Disturbances or dynamic obstacle detours perturb this clock equally for both robots, maintaining a shared consensus of progress so that the meeting point in space–time is preserved without replanning. The tight coupling produces a predictable transfer window that reduces idle hovering and aligns with qualitative observations of coordinated and synchronized motion in human handovers.

II. RELATED WORK

Using coupled dynamical systems (DS) to synchronize multiple manipulators is well-established. Early results showed that adding diffusive coupling terms between independent Lagrangian systems guarantees convergence of the synchronization error [13].

A unified linear–parameter-varying DS framework capable of switching between asynchronous point-to-point reaches and perfectly synchronous interception of a moving object-while embedding a quadratic-program inverse-kinematics layer for real-time self-collision avoidance, was presented in [14]. While laying important groundwork for hybrid sync/async coordination, the formulation concentrates on positional synchrony for interception.

A complementary approach is presented in [15], where a humanoid robot employs state-dependent DS primitives to coordinate whole-body reaching, grasping, navigation, and cooperative manipulation. For bimanual tasks such as lifting boxes, the grasp controller models the two hands as rigidly connected via a virtual object, an effective abstraction for cooperative lifting, but the work targets a different objective from the temporally phase-locked handover addressed here and does not explicitly focus on timing synchrony between unconstrained arms.

An intermediate-DS coupling method for in-hand manipulation was introduced in [16], combining multi-finger synchronization with an adaptive torque controller to compensate for dynamic uncertainties. The method is convincingly validated at the finger scale; however, it primarily addresses grasping and finger-level manipulation and does not explicitly study phase-locked coordination of full 6-DoF end-effector motions across spatially separated manipulators.

Building on prior DS-based approaches, this work presents the following contributions:

1) A dynamical-systems control framework for robotic handover that replaces segmented pipeline stages with a single formulation, enabling seamless, human-like transfers.
2) Extension of the intermediate DS coupling scheme, previously applied at the finger scale [16], to two spatially separated manipulators. The formulation extends coupling to full 6-DoF end-effector states by introducing a quaternion-based orientation DS, ensuring both translational and rotational synchrony at the handover pose.
3) A single consensus clock that couples translational and rotational progress variables, guaranteeing simultaneous, phase-locked convergence of position and orientation even under disturbances or obstacle-induced path deviations.
4) Experimental validation on two distinct platforms: two independent 6-DoF manipulator setups and an upper torso dual-arm system.

III. DYNAMICAL SYSTEMS FRAMEWORK

We define first-order autonomous dynamical systems for each manipulator $i \in \{1, \dots, N\}$, governing the evolution of its end-effector state

$$x_{\text{pose}}^{(i)} = \big(x^{(i)},\, q^{(i)}\big),$$

where $x^{(i)} \in \mathbb{R}^3$ denotes the position and $q^{(i)} \in \mathbb{H}$ is the unit quaternion representing orientation.

A. Asynchronous Motion

In the asynchronous (async) setting, each end-effector evolves independently toward its goal configuration. Let $x_s^{(i)}, x_f^{(i)} \in \mathbb{R}^3$ denote the start and goal positions, and $q_s^{(i)}, q_f^{(i)} \in \mathbb{H}$ denote the corresponding start and goal orientations.

Because no inter-arm synchronisation is enforced, the flow at any instant depends only on the current state. We write the generic first-order law

$$\dot{x}_{pose}^{(i)} = f^{(i)}\big(x_{pose}^{(i)}\big), \tag{1}$$

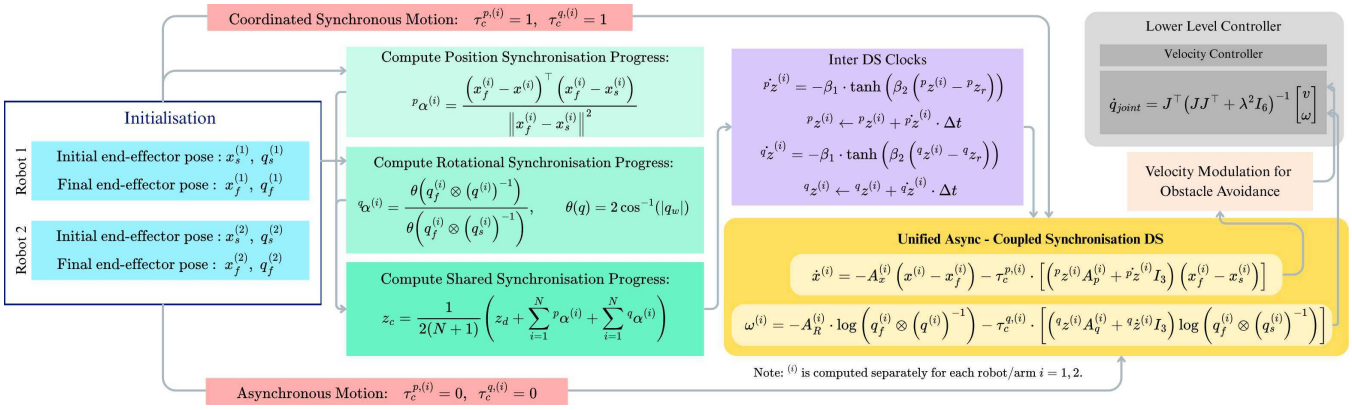

Fig. 3: Overview of the proposed consensus driven dynamical systems control. Given the start and goal poses for each robot, synchronization progress variables are computed independently for position and orientation. These feed into a shared consensus clock z_c, which drives the inter-DS clocks ${}^{p}z^{(i)}$, ${}^{q}z^{(i)}$, and ensures coordinated convergence of both motion modalities. The final task-space velocities $\dot{x}^{(i)}$ and angular velocities $\omega^{(i)}$ are passed to a lower-level controller, here implemented as a damped least-squares velocity controller. The architecture supports both Coordinated Synchronous Motion ($\tau_c^{p,(i)}, \tau_c^{q,(i)} = 1$) and Asynchronous Motion ($\tau_c^{p,(i)}, \tau_c^{q,(i)} = 0$) by simply updating coupling gates.

where $f^{(i)}$ is chosen so that $x_f^{(i)}$ is a (globally) asymptotically stable equilibrium. We now instantiate this law separately for translation and rotation.

1) Position Async DS: We define the error and impose a stable linear DS:

$$e_p^{(i)} = x^{(i)} - x_f^{(i)}, \tag{2}$$
$$\dot{e}_p^{(i)} = -A_x^{(i)}\, e_p^{(i)}, \tag{3}$$
$$\dot{x}^{(i)} = -A_x^{(i)}\left(x^{(i)} - x_f^{(i)}\right), \tag{4}$$

where $A_x^{(i)} \in \mathbb{R}^{3\times3}$ is a symmetric positive-definite (SPD) gain. In our implementation, we use the isotropic choice $A_x^{(i)} = k^{(i)} I_3$, with $k^{(i)} > 0$.

2) Orientation Async DS: We define the orientation error via the quaternion logarithmic map:

$$e_q^{(i)} = \log\left(q_f^{(i)} \otimes \left(q^{(i)}\right)^{-1}\right), \tag{5}$$
$$\dot{e}_q^{(i)} = -A_R^{(i)}\, e_q^{(i)}, \tag{6}$$
$$\omega^{(i)} = -A_R^{(i)}\, e_q^{(i)}, \tag{7}$$

where $\omega^{(i)} \in \mathbb{R}^3$ denotes the desired body-frame angular velocity and $A_R^{(i)} \in \mathbb{R}^{3\times3}$ is a symmetric positive-definite (SPD) gain. In our implementation, we use the isotropic choice $A_R^{(i)} = a^{(i)} I_3$ with $a^{(i)} > 0$.

Equations (4) and (7) drive each manipulator's position and orientation independently to their respective targets.

B. Coupled Synchronous Motion

Coupled synchronous (sync) actions require all manipulators to reach their goals with both temporal and spatial agreement. For each arm i we maintain two local progress variables,

$$^{p}z^{(i)} \in [0,1], \qquad ^{q}z^{(i)} \in [0,1],$$

that evolve from 1 (at the start) to 0 (at the goal) for both translation and orientation.

The dynamics of either modality can be written in the generic task–space form

$$\dot{x}_{pose}^{(i)} = f^{(i)}\left(x_{pose}^{(i)},\ ^{p,q}z^{(i)},\ ^{p,q}\dot{z}^{(i)}\right), \tag{8}$$

1) Position Coupled Sync DS: For the position-coupled DS, we define a virtual goal that slides along the start-goal line,

$$x^{*(i)} = x_f^{(i)} - {}^{p}z^{(i)}\left(x_f^{(i)} - x_s^{(i)}\right). \tag{9}$$

The geometric projection, which is later used for coupling

$$^{p}\alpha^{(i)} = \frac{\left(x_f^{(i)} - x^{(i)}\right)^\top \left(x_f^{(i)} - x_s^{(i)}\right)}{\|x_f^{(i)} - x_s^{(i)}\|^2}, \tag{10}$$

reacts immediately to disturbances. This quantity captures the normalized scalar projection of the instantaneous displacement $x_f^{(i)} - x^{(i)}$ onto the start–goal direction $x_f^{(i)} - x_s^{(i)}$, that updates immediately under perturbations/disturbances.

Each arm steers its position clock toward a reference $z_r^{(i)}$ (defined in Sec. III-B.3) through the sigmoid law

$$\dot{}^{p}z^{(i)} = -\beta_1 \tanh\left(\beta_2\left({}^{p}z^{(i)} - {}^{p}z_r^{(i)}\right)\right), \qquad \beta_1, \beta_2 > 0, \tag{11}$$

which provides fast correction for large phase deviations and smooth convergence near synchrony.

Enforcing the error dynamics $\dot{e}_p^{(i)} = -A_x^{(i)}\, e_p^{(i)}$, $e_p^{(i)} = x^{(i)} - x^{*(i)}$. yields

$$\dot{x}_d^{(i)} = -A_x^{(i)}\left(x^{(i)} - x_f^{(i)}\right) - \left[{}^{p}z^{(i)} A_p^{(i)} + \dot{}^{p}z^{(i)} I_3\right]\left(x_f^{(i)} - x_s^{(i)}\right), \tag{12}$$

where $A_p^{(i)} \in \mathbb{R}^{3\times3}$ is a symmetric positive-definite (SPD) gain.

A step-by-step derivation of (12) is provided in the Appendix.

2) Orientation Coupled Sync DS: Let $q_s^{(i)}, q_f^{(i)} \in \mathbb{H}$ be start and goal quaternions. The fixed displacement

$$ {}^{q}\Delta^{(i)} = \log\Big(q_f^{(i)} \otimes (q_s^{(i)})^{-1}\Big) \in \mathbb{R}^3 \tag{13}$$

is constant for the entire motion. A sliding target quaternion is defined as

$$ q^{*(i)} = \exp\big(-\,{}^{q}z^{(i)}\,{}^{q}\Delta^{(i)}\big) \otimes q_f^{(i)}. \tag{14}$$

The body-frame orientation error is $e_q^{(i)} = \log\big(q_f^{(i)} \otimes (q^{(i)})^{-1}\big)$, and the desired angular velocity follows from a first-order log-space DS:

$$ \omega_d^{(i)} = -A_R^{(i)}\, e_q^{(i)} - \Big[{}^{q}z^{(i)} A_q^{(i)} + {}^{q}\dot{z}^{(i)} I_3\Big]\,{}^{q}\Delta^{(i)}, \tag{15}$$

where $A_q^{(i)} \in \mathbb{R}^{3\times3}$ is a symmetric positive-definite (SPD) gain. A step-by-step derivation of (15) is provided in the Appendix.

The instantaneous rotational progress used later for coupling is

$$ {}^{q}\alpha^{(i)} = \frac{\theta\big(q_f^{(i)} \otimes (q^{(i)})^{-1}\big)}{\theta\big(q_f^{(i)} \otimes (q_s^{(i)})^{-1}\big)}, \tag{16}$$

where $\theta(q) = 2\arccos\big(|q_w|\big)$, with q_w denoting the scalar component of the unit quaternion q. All quaternion products are hemisphere-aligned prior to evaluation to ensure shortest-arc interpolation.

The rotational clock evolves analogously to the translational clock as

$$ {}^{q}\dot{z}^{(i)} = -\beta_1 \tanh\big(\beta_2\big({}^{q}z^{(i)} - {}^{q}z_r^{(i)}\big)\big), \tag{17}$$

with the same positive gains $\beta_1, \beta_2 > 0$.

3) Shared Synchronisation Clock: After evaluating the instantaneous translational and rotational progress measures ${}^{p}\alpha^{(i)}$ from (10) and ${}^{q}\alpha^{(i)}$ from (16), we define a common reference clock shared across all manipulators:

$$ z_r^{(i)} = z_c = \frac{1}{2\,(N+1)}\Big(z_d + \sum_{j=1}^{N} {}^{p}\alpha^{(j)} + \sum_{j=1}^{N} {}^{q}\alpha^{(j)}\Big), \tag{18}$$

which couples translational and rotational progress across agents.

The reference $z_r^{(i)}$ is then substituted into the sigmoid laws (11) and (17) for translation and rotation, respectively ($z^{(i)} \mapsto {}^{p}z^{(i)}$ and $z^{(i)} \mapsto {}^{q}z^{(i)}$).

- If all arms progress identically, $z_c = z_d$ and the group advances in lock-step.
- Any lag in position or orientation reduces z_c; faster arms decelerate while slower arms accelerate until synchrony is restored.

By construction, the velocity laws (12) and (15), together with the shared clock (18), yield phase-locked translational and rotational motion, including when trajectories are locally deformed by obstacle-avoidance modulation.

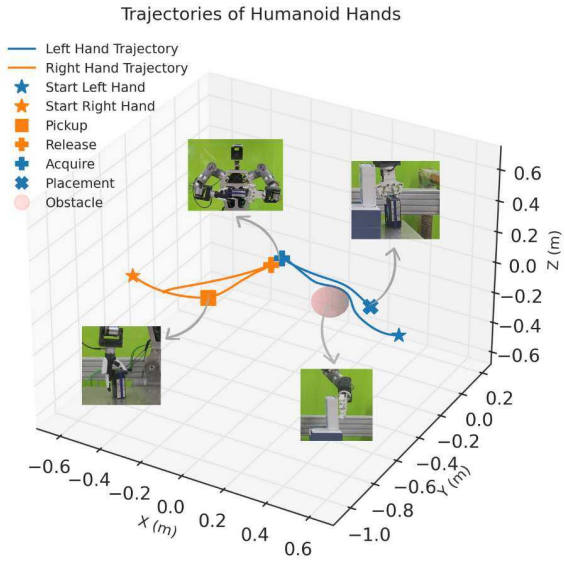

(a) Box hand-over

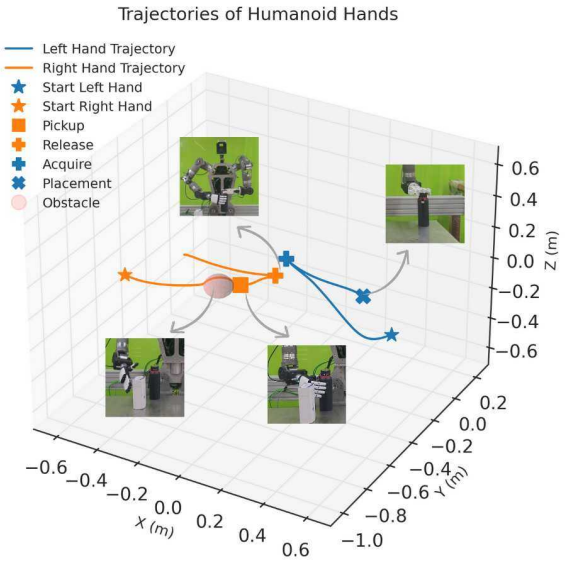

(b) Bottle hand-over

Fig. 4: End-effector trajectories for the two hand-over tasks. Each trace highlights the pick-up, transfer, release, and acquisition points for both manipulators separately.

C. Modulation for Obstacle Avoidance

We wrap the translational velocity with the dynamic-modulation principle of [17]. A static obstacle is approximated by an n-sphere centered at $c \in \mathbb{R}^3$ with radius $r > 0$. For manipulator i we write the relative position $\tilde{p}^{(i)} = x^{(i)} - c$ and define the analytic map

$$ \varphi_s\big(\tilde{p}^{(i)};\, r\big) = \Big(1 + \frac{r^2}{\tilde{p}^{(i)\top}\tilde{p}^{(i)}}\Big)\tilde{p}^{(i)}. \tag{19}$$

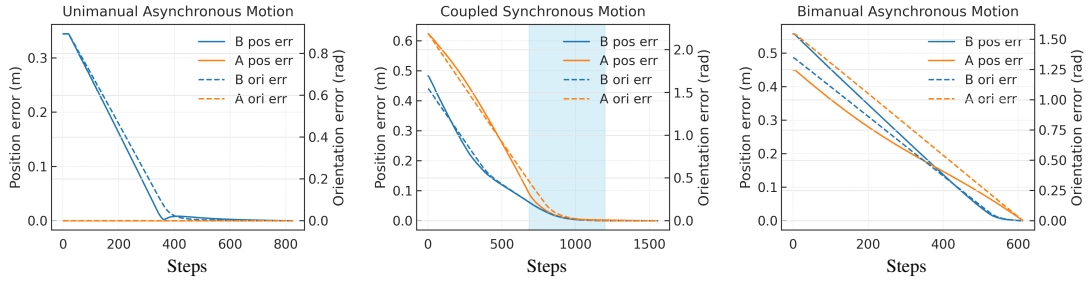

(a) Handover of box without any obstacles

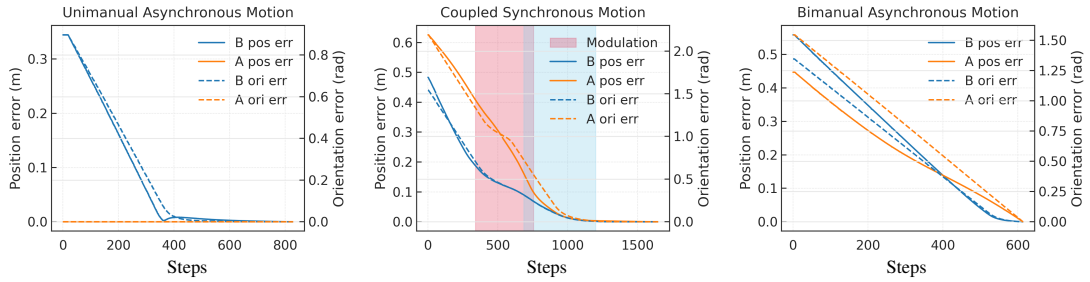

(b) Handover of box with obstacles

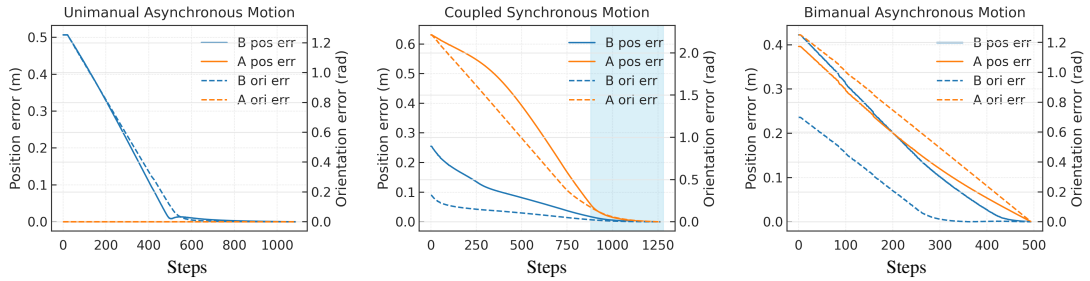

(c) Handover of bottle without any obstacles

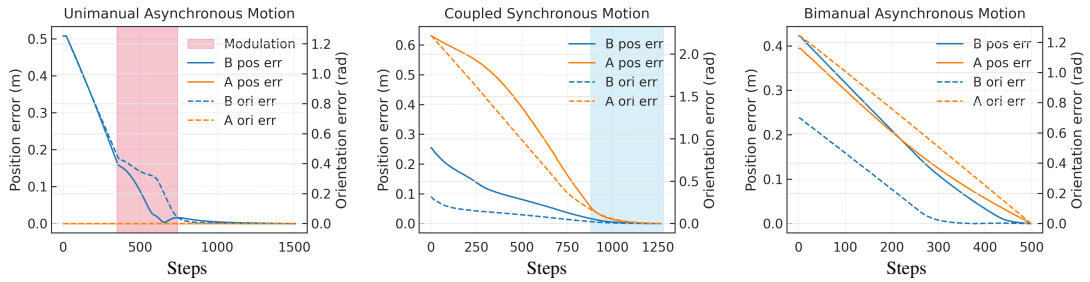

(d) Handover of bottle with obstacles

Fig. 5: Each row shows three contiguous phases of a single continuous trial: async pick-up → coupled synchronous handover → async retreat. Blue bands mark the synchronous handover window; red highlights indicate obstacle-induced velocity modulation. Phase transitions are effected via piecewise-constant coupling gates $\tau_c^p, \tau_c^q \in \{0, 1\}$ applied in the position and orientation laws (Eqs. (22), (23)).

Its Jacobian gives the dynamic-modulation matrix

$$M_s(\tilde{p}^{(i)};\, r) = I_3 + \left(\frac{r}{\tilde{p}^{(i)\top}\tilde{p}^{(i)}}\right)^2 \left[(\tilde{p}^{(i)\top}\tilde{p}^{(i)})I_3 - 2\,\tilde{p}^{(i)}\tilde{p}^{(i)\top}\right]. \tag{20}$$

By left-multiplying the desired velocity,

$$\dot{x}^{(i)} = M_s(\tilde{p}^{(i)};\, r)\,\dot{x}_d^{(i)}, \tag{21}$$

the flow bends smoothly around the obstacle while the global attractor of the dynamical system is preserved (see stability proof in [17]).

The orientation dynamics (Sec. III-B) remain unaltered; thus the quaternion-based synchronization continues to converge even when the translational path is locally deformed.

D. Blending Async and Sync Dynamics

To switch smoothly from the independent (async) DS to the fully coupled (sync) DS we introduce time-varying coupling gains $\tau_c^{p,(i)}(t),\, \tau_c^{q,(i)}(t) \in [0, 1]$ for each manipulator. Setting $\tau_c^{p,(i)} = \tau_c^{q,(i)} = 0$ recovers the pure asynchronous flow, while $\tau_c^{p,(i)} = \tau_c^{q,(i)} = 1$ yields the fully coupled equations of Sec. III-B.

Position dynamics: Scaling only the coupling term in (12) gives

$$\begin{aligned}\dot{x}^{(i)} = &-A_x^{(i)}\left(x^{(i)} - x_f^{(i)}\right) \\ &- \tau_c^{p,(i)}\left[{}^{p}z^{(i)}A_p^{(i)} + {}^{p}\dot{z}^{(i)}I_3\right]\left(x_f^{(i)} - x_s^{(i)}\right).\end{aligned} \tag{22}$$

Orientation dynamics: Applying the same scaling gives

$$\boldsymbol{\omega}^{(i)} = -A_R^{(i)} \log\left(q_f^{(i)} \otimes \left(q^{(i)}\right)^{-1}\right) - \tau_c^{q,(i)} \left[{}^{q}\dot{z}^{(i)} A_q^{(i)} + {}^{q}\dot{z}^{(i)} I_3\right] \log\left(q_f^{(i)} \otimes \left(q_s^{(i)}\right)^{-1}\right). \tag{23}$$

Choice of $\tau_c^{p,(i)}$, $\tau_c^{q,(i)}$: Any C^1 function that increases monotonically from 0 to 1 can be used. Typical choices are $\tau_c^{p,(i)}(t) = \sigma\big(\kappa(t - t_0)\big)$, with $\sigma(s) = \frac{1}{1+e^{-s}}$, or a progress-based gate $\tau_c^{p,(i)} = \sigma\big(\kappa({}^{p}\alpha^{(i)} - \alpha_{\text{th}})\big)$, where $\kappa > 0$ sets transition sharpness and α_{th} is the point at which the system starts coupling. The same constructions apply to $\tau_c^{q,(i)}$. Because the scaling is confined to the coupling term, the underlying async dynamics remain globally stable and the overall blended field inherits stability for all $\tau_c^{p,(i)}, \tau_c^{q,(i)} \in [0, 1]$.

IV. Experimental Validation

The proposed framework is evaluated on two hardware platforms:

1) Two independent 6-DoF manipulators, each equipped with a parallel-jaw gripper.
2) An upper torso dual arm platform; the right arm carries an Inspire-Robots RH56E2 6-DoF dexterous hand and the left arm carries a parallel-jaw gripper.

For every trial the grasp position and orientation of the object are specified in advance. Although these parameters could be generated automatically in future work by integrating vision-based perception and grasp-planning modules, manual specification does not affect the current framework. Obstacle geometries are likewise predefined, but the controller already provides an interface for subscribing to real-time obstacle information from external vision pipelines. The framework outputs end-effector velocities, which are subsequently converted to joint-level velocity commands and executed on both platforms by the underlying joint-velocity controller.

A. Coupled Synchronous Motion

We evaluate the coupled controller across the three hand-over phases, pick-up, transfer, and release, using two objects (box and bottle) in environments with and without obstacle modulation. With both translational and rotational coupling active, Fig. 5 shows position and orientation errors for both manipulators. In all four conditions the error traces decay concurrently and reach zero within the same temporal window, evidencing phase-locked convergence of both modalities across both arms. The highlighted bands emphasize that, even when the translational velocity is locally deformed to skirt obstacles, the rotational DS and the shared clock keep the group synchronized; faster arms automatically decelerate while lagging arms accelerate until synchrony is restored. The corresponding end-effector paths in Fig. 4 remain smooth and collision-free, exhibiting local detours yet identical arrival times at the transfer pose without idle hovering.

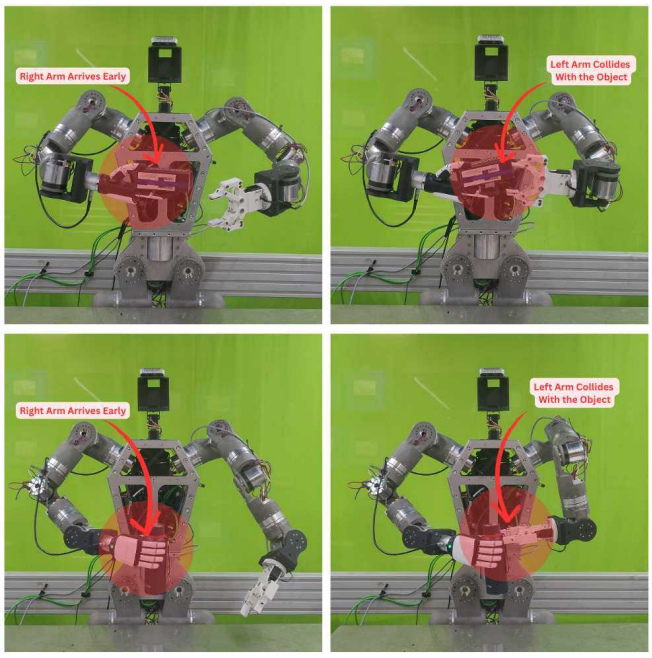

Fig. 6: Ablation of the no-coupling baseline ($\tau_c^{p,(i)} = \tau_c^{q,(i)} = 0$; purely asynchronous). The 2×2 grid shows two tasks (top: box, bottom: bottle). In each row, the left panel illustrates one arm arriving early at the transfer pose, and the right panel shows the resulting collision of the other arm with the object. Red overlays highlight the contact region and the failure mode induced by loss of synchronisation.

B. Ablation Study

To isolate the contribution of the proposed coupling, we compare two controllers on both objects (box, bottle): (i) with coupling ($\tau_c^p = \tau_c^q = 1$), and (ii) no coupling ($\tau_c^p = \tau_c^q = 0$). Each condition was tested in five runs per object under varied initial placements; a run was counted as a success if the object was transferred at the designated hand-over pose without collisions. Table I reports these representative outcomes, while additional runs under other initial conditions confirmed the same overall trend.

TABLE I: Ablation study: successful transfers out of five runs per object. Additional trials with varied initial conditions showed consistent results.

Variant	Box (5 runs)	Bottle (5 runs)
With coupling	5	5
No coupling	0	1

Across both objects, the coupled controller consistently succeeds $(10/10)$, whereas the no-coupling baseline rarely succeeds $(1/10)$. The dominant failure mode without coupling is temporal misalignment: one arm reaches the transfer pose early and idles or intrudes into the partner's workspace, after which the other arm collides with the object or the first arm. This behaviour is visible in Fig. 6 and in the error traces of Fig. 7 (b,d), where position/orientation errors decay at different rates for the two manipulators.

In contrast, with coupling enabled the progress clocks are driven by the shared consensus value, forcing faster arms to

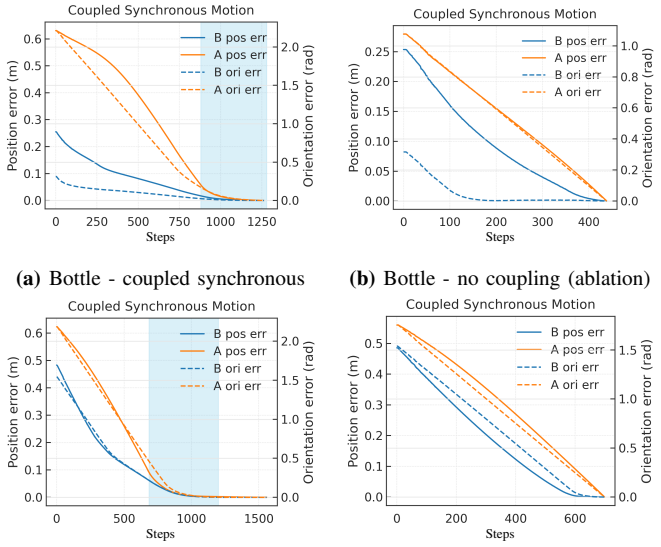

(a) Bottle - coupled synchronous **(b)** Bottle - no coupling (ablation)

(c) Box - coupled synchronous **(d)** Box - no coupling (ablation)

Fig. 7: Coupled synchronisation vs. no-coupling (ablation). Top row: bottle; bottom row: box. Left column uses the proposed coupled controller; the right column disables coupling. In the coupled case, the blue shaded region indicates the handover window where both modalities converge simultaneously; without coupling, errors decay at different rates, leading to loss of temporal synchronisation.

decelerate and lagging arms to accelerate until both position and orientation errors collapse within the same temporal window. This phase-locking is evident in Fig. 7 (a,c): the shaded region marks simultaneous convergence of all error channels and coincides with a collision-free, zero-idle exchange. Since this ablation was conducted at the velocity-control layer, the observed behaviours are primarily kinematic in nature, but the consistency across varied initialisations supports the conclusion that clock-level coupling is beneficial to maintain temporal synchronisation and spatial coordination during handover.

V. DISCUSSIONS AND CONCLUSIONS

This work introduced a control-centric framework for dual-arm handover that replaces the customary segmented pipeline (approach/hover/transfer/retreat) with a single-loop dynamical system. Two key design choices underpin the observed behavior.

First, translation and rotation are synchronized through a shared consensus clock z_c, coupling the local progress variables ${}^{p}z^{(i)}$ and ${}^{q}z^{(i)}$ across arms and modalities. At the velocity level, explicit dependence on ${}^{p}z^{(i)}, {}^{p}\dot{z}^{(i)}$ and ${}^{q}z^{(i)}, {}^{q}\dot{z}^{(i)}$ (Eqs. (12), (15)) causes faster agents to decelerate and slower ones to accelerate, restoring phase-locking. This produces a predictable, short transfer window in which both grippers reach the handover pose simultaneously.

Second, we preserve the simplicity of first-order DS control by shaping translation with a moving attractor along the start–goal segment and rotation with a quaternion log-space DS that tracks a sliding target along the shortest arc. Coupling gates influence only the clock-dependent terms, while the asynchronous DS remains globally stable. As a result, the vector field stays well behaved for all gate values and allows smooth transitions between asynchronous and fully coupled regimes without replanning. The controller therefore preserves goal-attractor stability for any $\tau_c^{p,(i)}, \tau_c^{q,(i)} \in [0,1]$ while enforcing phase-locked convergence of translational and rotational errors during handover.

The benefits are evident under obstacle-induced path deformation: position flows bend around obstacles while orientation tracking and synchrony remain intact. Disabling coupling reproduces common segmented handover failures—one arm arrives early and hovers while the other collides with the object or partner. Our ablation (Fig. 6, right column of Fig. 7) and the success rates in Table I (5/5 vs. 0/5 for box, 5/5 vs. 1/5 for bottle, with vs. without coupling) quantify this effect. Although conducted at the velocity-control layer and primarily kinematic, these experiments consistently demonstrate that clock-level coupling is essential for maintaining temporal synchronisation and spatial coordination.

Unlike trajectory-based pipelines, the proposed controller does not rely on pre-timed plans; timing emerges from the coupled clocks. Compared to earlier DS approaches focusing on positional synchrony or rigid templates, our formulation couples full 6-DoF end-effector states and introduces a quaternion DS for rotational synchrony. Across all coupled trials, clock-level coupling significantly improves temporal synchronisation and spatial coordination.

VI. FUTURE WORK

Future extensions will generalise the controller from simultaneous coupling to task-dependent sequencing. Many bimanual skills (e.g., tilt–pour) require position to converge before rotation, keeping a container upright until it is stably over the receptacle, then releasing the orientation clock. Phase-conditioned gates will trigger rotation once a positional threshold and contact confidence are met, within the same DS structure. Beyond same-side grasps, we will test cross-body exchanges and reach-through motions that stress clearance, self-collision avoidance, and heterogeneous kinematics. The broader aim is to span a family of bimanual behaviours, handover-in-motion, cooperative transport, tray carrying, insertion, tilt–pour, by selecting clock schedules and coupling patterns instead of replanning trajectories.

A complementary line targets autonomy and robustness, in which the coupling gate becomes state dependent. Specifically, we will design $\tau_c^{p,(i)} = \tau_c^{p,(i)}\big(\text{distance, relative phase, contact confidence}\big)$ and $\tau_c^{q,(i)} = \tau_c^{q,(i)}\big(\text{distance, relative phase, contact confidence}\big)$ so the controller exhibits reflex-like behavior: when a drop or large separation is detected, coupling is reduced, the system reverts to asynchronous re-acquisition, and synchrony is re-engaged once conditions recover. In parallel, online perception and lightweight reasoning (VLM/LLM) will infer object poses and select DS templates (e.g., desired handover pose, which modalities couple, and when), thereby instantiating goals and coupling schedules on the fly.

APPENDIX

This appendix provides cohesive derivations of the closed-form velocity fields for translation and rotation used by the synchronous controller. Unless required for clarity, superscripts (i) are omitted; x_s, x_f and q_s, q_f denote the fixed start and goal configurations of a single manipulator.

A. Translational dynamics (deriving Eq. (12)*)*

We begin by defining a sliding attractor that moves on the start–goal segment as a function of the clock ${}^{p}z$:

$$x^*({}^{p}z) \;=\; x_f - {}^{p}z\,\big(x_f - x_s\big), \qquad \dot{x}^* \;=\; -\,{}^{p}\dot{z}\,\big(x_f - x_s\big).$$

This choice enforces $x^*(1) = x_s$ and $x^*(0) = x_f$, encoding phase evolution along the line connecting start and goal. Let the error be $e_p := x - x^*$. We impose first-order error decay $\dot{e}_p = -A_x\, e_p$ with $A_x \succ 0$. Since x_s and x_f are constant, the attractor velocity is as above. Substituting $\dot{e}_p = \dot{x} - \dot{x}^*$ and expanding $e_p = x - x_f + {}^{p}z\,(x_f - x_s)$ yields

$$\dot{x} \;=\; -A_x\big(x - x_f\big) \;-\; \big[\,{}^{p}z\,A_p + {}^{p}\dot{z}\,I_3\,\big]\,(x_f - x_s),$$

which recovers Eq. (12). The first term drives direct convergence toward x_f, while the second term couples motion to the shared clock through ${}^{p}z$ and ${}^{p}\dot{z}$, regulating phase so that all agents arrive simultaneously even when trajectories are locally deformed (e.g., via obstacle modulation).

B. Rotational dynamics (deriving Eq. (15)*)*

For orientation, let $q_s, q_f \in \mathbb{H}$ be the start and goal quaternions, hemisphere-aligned. The fixed axis–angle displacement from start to goal is

$${}^{q}\Delta \;=\; \log\!\Big(q_f \otimes q_s^{-1}\Big) \;\in\; \mathbb{R}^3,$$

which remains constant throughout the motion. Mirroring the translational construction, we define a sliding target quaternion that regresses along ${}^{q}\Delta$ according to the rotational clock ${}^{q}z$:

$$q^*({}^{q}z) \;=\; \exp\!\big(-\,{}^{q}z\,{}^{q}\Delta\big) \otimes q_f,$$

so that $q^*(1) = q_s$ and $q^*(0) = q_f$. We track this moving target in the tangent space via the log-space error

$${}^{q}\phi := \log\!\Big(q^* \otimes q^{-1}\Big).$$

A first-order tracking DS in the tangent space, ${}^{q}\dot{\phi} = -A_R\,{}^{q}\phi$ with $A_R \succ 0$, yields a stabilizing law in a Euclidean vector space. Using first-order properties of the quaternion log/exp maps around the identity, we write the local approximation

$${}^{q}\phi \approx -\,{}^{q}z\,{}^{q}\Delta + e_q, \qquad e_q := \log\!\big(q_f \otimes q^{-1}\big),$$

and differentiate to obtain ${}^{q}\dot{\phi} \approx -\,{}^{q}\dot{z}\,{}^{q}\Delta + \dot{e}_q$. Equating with ${}^{q}\dot{\phi} = -A_R\,{}^{q}\phi$ and identifying $\dot{e}_q$ with the desired body-frame angular velocity ω_d (small-angle regime) gives the closed-form control:

$$\omega_d \;=\; -A_R\, e_q - \big[\,{}^{q}z\,A_q + {}^{q}\dot{z}\,I_3\,\big]\,{}^{q}\Delta,$$

which recovers Eq. (15). As in translation, the first term stabilizes residual error relative to q_f, while the clock-dependent terms regulate phase progression along the start–goal rotation.

Translation and rotation share the same recipe: a sliding start–goal (geodesic) attractor, a first-order DS on the log-space error, and explicit clock terms ${}^{p}z, {}^{p}\dot{z}$ and ${}^{q}z, {}^{q}\dot{z}$ so a shared consensus clock phase-locks all manipulators. Since only the clock terms are scaled, the system blends smoothly between asynchronous and fully synchronous behavior while preserving error-to-goal stability.

REFERENCES

[1] M. Costanzo, G. De Maria, and C. Natale, "Handover control for human-robot and robot-robot collaboration," Frontiers in Robotics and AI, vol. 8, p. 672995, 2021.

[2] V. Ortenzi, A. Cosgun, T. Pardi, W. P. Chan, E. Croft, and D. Kulić, "Object handovers: a review for robotics," IEEE Transactions on Robotics, vol. 37, no. 6, pp. 1855–1873, 2021.

[3] H. Zhang, S.-H. Chan, J. Zhong, J. Li, P. Kolapo, S. Koenig, Z. Agioutantis, S. Schafrik, and S. Nikolaidis, "Multi-robot geometric task-and-motion planning for collaborative manipulation tasks," Autonomous Robots, vol. 47, no. 8, pp. 1537–1558, 2023.

[4] J. Kurosu, A. Yorozu, and M. Takahashi, "Simultaneous dual-arm motion planning for minimizing operation time," Applied Sciences, vol. 7, no. 12, p. 1210, 2017.

[5] Z. Geng, Z. Yang, W. Xu, W. Guo, and X. Sheng, "A globally guided dual-arm reactive motion controller for coordinated self-handover in a confined domestic environment," Biomimetics, vol. 9, no. 10, p. 629, 2024.

[6] Y. Li, C. Pan, H. Xu, X. Wang, and Y. Wu, "Efficient bimanual handover and rearrangement via symmetry-aware actor-critic learning," in 2023 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2023, pp. 3867–3874.

[7] G. Vezzani, M. Regoli, U. Pattacini, and L. Natale, "A novel pipeline for bi-manual handover task," Advanced Robotics, vol. 31, no. 23-24, pp. 1267–1280, 2017.

[8] L. Zhu, C. Y. Wong, and A. Moon, "Did you notice? unveiling robot-induced synchronization in human-robot interaction," in 2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids), 2025, pp. 1–8.

[9] P. Basili, M. Huber, T. Brandt, S. Hirche, and S. Glasauer, Investigating Human-Human Approach and Hand-Over. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 151–160.

[10] K. Strabala, M. K. Lee, A. Dragan, J. Forlizzi, S. S. Srinivasa, M. Cakmak, and V. Micelli, "Toward seamless human-robot handovers," Journal of Human-Robot Interaction, vol. 2, no. 1, pp. 112–132, 2013.

[11] A. Moon, D. M. Troniak, B. Gleeson, M. K. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft, "Meet me where i'm gazing: how shared attention gaze affects human-robot handover timing," in Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction, 2014, pp. 334–341.

[12] A. H. Mason and C. L. MacKenzie, "Grip forces when passing an object to a partner," Experimental brain research, vol. 163, pp. 173–187, 2005.

[13] S.-J. Chung and J.-J. E. Slotine, "Cooperative robot control and concurrent synchronization of lagrangian systems," IEEE transactions on Robotics, vol. 25, no. 3, pp. 686–700, 2009.

[14] S. S. Mirrazavi Salehian, N. Figueroa, and A. Billard, "A unified framework for coordinated multi-arm motion planning," The International Journal of Robotics Research, vol. 37, no. 10, pp. 1205–1232, 2018.

[15] N. Figueroa, S. Faraji, M. Koptev, and A. Billard, "A dynamical system approach for adaptive grasping, navigation and co-manipulation with humanoid robots," in 2020 IEEE International conference on robotics and automation (ICRA). IEEE, 2020, pp. 7676–7682.

[16] F. Khadivar and A. Billard, "Adaptive fingers coordination for robust grasp and in-hand manipulation under disturbances and unknown dynamics," IEEE Transactions on Robotics, vol. 39, no. 5, pp. 3350–3367, 2023.

[17] S. M. Khansari-Zadeh and A. Billard, "A dynamical system approach to realtime obstacle avoidance," Autonomous Robots, vol. 32, pp. 433–454, 2012.