

Robot Teams Sharing a Load: Compositional Decay Certificates

Mehdi Heydari Shahna

mehdi.heydarishahna@tuni.fi

Tampere University

Soyi Jung

Ajou University

Soohyun Park

Sookmyung Women's University

Joongheon Kim

Korea University

Jouni Mattila

Tampere University

Research Article

Keywords: Compositional stability, dissipativity, chordal decomposition, constrained mechanical systems, cooperative manipulation, virtual decomposition control

Posted Date: October 19th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11401773/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Robot Teams Sharing a Load: Compositional Decay Certificates[★]

Mehdi Heydari Shahna^a, Soyi Jung^b, Soohyun Park^c, Joongheon Kim^d,
Jouni Mattila^a

^a*Faculty of Engineering and Natural Sciences, Tampere University, Tampere, Finland*

^b*Department of Electrical and Computer Engineering, Ajou University, Suwon, Republic of Korea*

^c*Division of Computer Science, Sookmyung Women's University, Seoul, Republic of Korea*

^d*Department of Electrical and Computer Engineering, Korea University, Seoul, Republic of Korea*

Abstract

Robots carrying a shared payload form a constrained mechanical system whose velocity-error decay rate depends on all robots and constraints. We study when this rate can be certified, and gains achieving it designed, from inequalities on each robot's own constraint space. For cell architectures with a join tree, local certificates exist for every rate below the centralized one, and attain it when the assembly realizes every combined velocity of the modules shared by adjacent cells, but not always otherwise. Every architecture without a join tree admits an assembly, possibly with constraints joining several modules, that has a positive centralized rate and no local certificate of positive rate. Because dissipation is affine in the gains, minimum-gain and maximum-rate designs decompose into robot-level linear matrix inequalities coupled through payload-sized supplies, without loss under the endpoint condition. An error-power identity admitting redundant constraints and a gain-weighted fit of independent joint and wheel requests make the certified rates decay rates of joint-level loops; for skid-steered bases, wheel friction supplies the missing damping. On four compiled mobile-manipulator teams, the local test reproduces the centralized rates, robot-level synthesis certifies the mission rate with 40–44% less damping gain, and complete missions show the loops in operation.

Key words: Compositional stability; dissipativity; chordal decomposition; constrained mechanical systems; cooperative manipulation; virtual decomposition control.

1 Introduction

Cooperative manipulation connects several robots through a shared object, producing a constrained mechanical system. Grasp closure and ground contacts restrict the admissible velocities, while the object contributes inertia that the robots accelerate jointly. Quantitative tracking analysis therefore depends on both the actuator controllers and the mechanical interconnection.

Constrained robot control uses reduced coordinates and projection operators to separate admissible motion from constraint reactions [1, 14, 21, 35]. Cooperative controllers address uncertain payload dynamics, decentralized adaptation, mobile platforms, and the distribution of internal forces [11, 13, 30, 34]. For redundant actuation, antagonism filters separate motion-producing efforts from internal loading, including weighted projections determined by the controller gains [16, 24]. Composite-error control provides a kinetic-energy tracking certificate [29], and projected adaptive control of closed-chain co-manipulation constructs a reference admissible at the measured configuration [12]. Virtual decomposition control (VDC) assigns subsystem error-storage balances whose interface terms are virtual power flows [40, 42], with cooperative, subsystem-based, observer, and geometric formulations [15, 17, 19, 41].

[★] This paper was not presented at any IFAC meeting. Corresponding author M. H. Shahna. Supplementary material for this paper is available at <https://github.com/Mehdi-Heydari-Shahna/CompositionalDecayCertificates>.

Email address: mehdi.heydarishahna@tuni.fi (Mehdi Heydari Shahna).

Compositional certification replaces a network-level dissipation inequality by subsystem inequalities coupled through supply rates [5, 23, 36]; the supplies can be searched by distributed optimization [22], and port-based and passive decompositions give mechanical counterparts [9, 20]. When the network certificate is a sparse semidefinite constraint, chordal decomposition replaces it without loss by clique-sized constraints [2, 33], which has been used for robust stability analysis and structured controller design of sparsely interconnected systems [4, 39]; for non-chordal patterns, factor-width decompositions are restrictions [7, 38].

In a constrained assembly, however, the certificate is positivity of $v^\top (\mathcal{D} - \kappa M)v$ on the admissible space $Av = 0$, not on the coordinate space. The unactuated shared inertia makes the full-space form indefinite, and each robot can verify an inequality only on its own constraint space, which contains motions that the complete assembly forbids. For rates strictly below the centralized rate, a constraint-penalty (Finsler) lift turns positivity on the admissible space into a sparse positive-definite decomposition, to which the chordal results apply. The lift leaves three questions open, and they are the subject of this paper: whether local certificates attain the centralized rate itself, which prescribed cell architectures are lossless for every assembly, and whether the local inequalities can design gains rather than only verify them.

The contributions are:

- (1) *Endpoint condition and architecture characterization.* For cells with a join tree, local certificates exist for every rate below the centralized rate $\underline{\kappa}$ (Theorem 11(a)) and, when the assembly realizes every velocity of each separator, at $\underline{\kappa}$ itself (Theorem 11(b)); without this realizability the endpoint can fail even for a two-robot payload star (Proposition 12). For payload stars the construction is a payload-space test with $n_0(n_0 + 1)/2$ -entry summaries per robot (Proposition 7, Theorem 8). Conversely, every architecture without a join tree admits an assembly with $\underline{\kappa} > 0$ and no certificate of positive rate (Theorem 13, Remark 15).
- (2) *Compositional synthesis.* Since dissipation is affine in the actuator gains, the local inequalities are linear matrix inequalities (LMIs) in each robot's gains and supplies, and minimum-gain and maximum-rate designs over finite configuration sets decompose into robot-level LMIs coupled through payload-sized supplies, without loss under the hypotheses of Theorem 11(b) (Theorem 16).
- (3) *Dynamic meaning and realization.* An error-power identity that admits redundant constraints (Propositions 3–5), together with a gain-weighted fit of independently generated joint and wheel requests (Theorem 19), makes the certified rates decay rates of joint-level loops with the fitted feedforward, which builds on [12, 16, 29]; a reset law bounds

the storage across grasp and release events and under sampled renewal (Proposition 22), and for skid-steered bases a friction law with a convex dissipation potential adds a secant damping that certifies a rate where the actuation rank condition fails (Proposition 23).

Sections 2–6 present the model, the identity and rates, the certificates, the synthesis, and the realization; Section 7 evaluates them on six compiled teams, from single configurations to complete missions.

In the setting of Fig. 1, RoboCompiler [28] attaches the component mechanism graphs of the robots, the payload, and the ground into one mode-dependent team graph and supplies the constraint Jacobian, admissible-velocity lift, actuator map, and inverse dynamics, in which the results are stated. Each robot is a module with its own storage W_i , the grasp and rolling edges are the power ports through which the modules interact, and the certificates are inequalities on these modules and ports. Certification of contact plants and of instantaneous load limits is outside the scope of this paper.

Analysis domain. Within each mode, the mechanical coefficients are smooth and the ideal bilateral constraints have constant rank; holonomic and nonholonomic rolling constraints are admitted, and the planar wheel forces of skid-steered bases follow a friction law at the wheel loads of a load model (Section 6.4). The nominal controller uses configuration and velocity (for skid-steered bases also the wheel loads), exact dynamics, and unrestricted actuator efforts; model and implementation perturbations enter through explicitly defined residuals. The gains are constant within the mode. All certificates are pointwise in the configuration (for skid-steered bases, in the state and the loads); trajectory statements integrate the pointwise rates, and designs are certified on stated finite configuration sets.

Notation. $|\cdot|$ is the Euclidean norm and the matrix norm it induces, $\langle X, Y \rangle = \text{tr}(X^\top Y)$, $|x|_Q = \sqrt{x^\top Q x}$ for $Q = Q^\top \succeq 0$, $\text{sym } X = \frac{1}{2}(X + X^\top)$, and $X^\dagger$ the Moore–Penrose inverse. $\lambda_{\min}(X, Y)$ and $\lambda_{\max}(X, Y)$ are the extreme generalized eigenvalues of a symmetric X relative to $Y = Y^\top \succ 0$, and blkdiag denotes a block-diagonal matrix. The subscript r marks required quantities, and a bar reduced matrices ($\bar{M}$, $\bar{D}$, $\bar{R}$).

2 Constrained mechanical model

Robot i , the payload ($i = 0$), and the ground are modules of one mechanism graph [28]. A mode ν specifies the active grasp, support, and rolling interfaces; each is an ideal edge and a VDC power port. Module i is a tree of rigid bodies with coordinates $q_i \in \mathbb{R}^{n_i}$; $q = (q_0, \dots, q_N) \in \mathcal{Q} \subset \mathbb{R}^n$ and $v = \dot{q}$, where $n = \sum_{i=0}^N n_i$; all coordinates




Fig. 1. Unified graph assembly following RoboCompiler. The component mechanism graphs retain their bodies, joints, and ports; grasp, ground, and support edges attach them into one mode-dependent graph. Its module view exposes the VDC power interfaces; the compiled matrices serve both the fitted controller and the certificates. The sketch generalizes to N robots.

are taken in regular local charts. The kinetic energy is $\frac{1}{2}v^\top M(q)v$ with $M = \text{blkdiag}(M_0, \dots, M_N) \succ 0$, M_i depending on q_i alone; $C(q, v)$ is the Coriolis matrix in Christoffel form, so that

$$\frac{d}{dt}M(q) = C(q, v) + C(q, v)^\top, \quad (1)$$

$g(q)$ collects the potential forces and $R = \text{blkdiag}(R_i) \succeq 0$ the viscous damping of the joints. The constant matrix $B = \text{blkdiag}(B_i) \in \mathbb{R}^{n \times p}$ maps the actuator efforts $\tau \in \mathbb{R}^p$, the payload rows of B vanishing, so that $B^\top q$ are the actuator coordinates. An ideal edge e contributes rows $A_e(q)v = 0$, holonomic ($A_e = \partial\Phi_e/\partial q$, $\Phi_e(q) = 0$) or not, which act only on the modules it joins (usually two; a cable of fixed length routed over several modules joins all of them); $A_\nu(q)$ is the stack of the m rows of mode ν . In a mode,

$$M\dot{v} + C(q, v)v + g + Rv = B\tau + A^\top\lambda, \quad A(q)v = 0, \quad (2)$$

with $A \in \mathbb{R}^{m \times n}$ and $\lambda \in \mathbb{R}^m$ the stacked constraint Jacobian and its multipliers; dependent rows are permitted.

Assumption 1 (Regular configurations). *The holonomic rows of mode ν vanish on an embedded submanifold $\mathcal{Z}_\nu$ whose tangent space is the kernel of their Jacobian, and on a set $\mathcal{Q}_\nu$, open in $\mathcal{Z}_\nu$, that contains the motions considered in the mode, A_ν is smooth and has constant rank. The configurations in $\mathcal{Q}_\nu$ are called regular.*

No independence of the rows is required: dependent rows occur when connections are duplicated, as with one contact row for each of four wheels under a rigid chassis. The admissible velocities are $\mathcal{T}(q) = \ker A(q)$, of dimension $d = n - \text{rank } A$.

The compiler returns, at a regular configuration, $E(q) \in \mathbb{R}^{n \times d}$ with $\text{range } E = \mathcal{T}(q)$, smooth in q , and from it

$$\begin{aligned} L &= B^\top E, & \bar{M} &= E^\top M E, & \bar{R} &= E^\top R E, \\ P &= E \bar{M}^{-1} E^\top M. \end{aligned} \quad (3)$$

P is the projector onto $\mathcal{T}(q)$ that is orthogonal in the kinetic-energy metric; equivalently $P = I - M^{-1}A^\top(AM^{-1}A^\top)^\dagger A$, a form that admits dependent rows [31]. The fitted reference, W , and the rates below do not depend on the choice of E (they are unchanged under $E \mapsto ET(q)$ with invertible T). With $v = Eu$, (2) reads $\bar{M}\dot{u} + E^\top((CE + M\dot{E})u + g + REu) = L^\top\tau$.

Lemma 2 (Redundant rows). *Under Assumption 1, for every state with $A(q)v = 0$ and every τ , the acceleration $\dot{v}$ and the total constraint force $A^\top\lambda$ are unique; λ is unique up to $\ker A^\top$. A component λ_j is the same for all solutions if and only if row j of A is not a combination of the other rows.*

Proof. The reduced equation determines $\dot{u}$, hence $\dot{v} = E\dot{u} + \dot{E}u$. Then $f = M\dot{v} + Cv + g + Rv - B\tau$ satisfies $E^\top f = 0$, that is $f \in (\text{range } E)^\perp = \text{range } A^\top$, which gives existence of λ ; two solutions differ by an element of $\ker A^\top$. The j -th component vanishes on $\ker A^\top$ if and only if $\hat{e}_j \in (\ker A^\top)^\perp = \text{range } A$, and $Ax = \hat{e}_j$ is solvable exactly when row j is independent of the others. $\square$

The second statement is the criterion of [37] for reactions under redundant constraints.

3 VDC error-power balance and decay rates

Let $v_r(t) \in \mathbb{R}^n$ be a *required velocity*, absolutely continuous along the motion (it may depend on time and on the measured state). A triple (v_r, τ_r, λ_r) satisfies the *required balance* if

$$M(q)\dot{v}_r + C(q, v)v_r + g(q) + Rv_r = B\tau_r + A(q)^\top\lambda_r, \quad (4)$$

with the actual q and v in M, C, g and A . The law is

$$\tau = \tau_r - KB^\top(v - v_r), \quad K = \text{blkdiag}(K_i) \succ 0, \quad (5)$$

which every robot evaluates with its own rates. Write $s = v - v_r$, $W = \frac{1}{2}|s|_M^2$ and $\mathcal{D} = BKB^\top + R \succeq 0$: the VDC

velocity error, aggregate error storage, and dissipation matrix. The norm at time t uses $M(q(t))$.

Proposition 3 (Error-power identity). *Let Assumption 1 hold and consider a motion of (2) under (5) with (4); v_r need not be admissible. Then, for almost all t :*

- (i) $\dot{W} = -s^\top \mathcal{D}s - \langle A(q)v_r, \lambda - \lambda_r \rangle$, which depends on the reactions only through the unique $A^\top(\lambda - \lambda_r)$.
- (ii) If $A(q(t))v_r(t) = 0$ for all t , then

$$\dot{W} = -s^\top \mathcal{D}s \quad (6)$$

whatever the topology of the graph, the number and the types of its ideal edges and the rank of A . If $A(q)v_r \neq 0$ at a state, then $\dot{W}$ takes every real value there as $\dot{v}_r$ varies, the instantaneous actuator effort and actual acceleration being fixed.

- (iii) $|s|_M \geq |A(q)v_r|_{(AM^{-1}A^\top)^\dagger} = |(I - P)v_r|_M$, the distance of v_r to $\mathcal{T}(q)$ in the kinetic-energy metric.
- (iv) $W = \sum_{i=0}^N W_i$ with $W_i = \frac{1}{2}|s_i|_{M_i}^2$, and $\dot{W}_i = -s_i^\top \mathcal{D}_i s_i + \sum_{e \ni i} \pi_{e,i}$, where $\mathcal{D}_i = B_i K_i B_i^\top + R_i$ and $\pi_{e,i} = \langle A_e^{(i)} s_i, \lambda_e - \lambda_{r,e} \rangle$ is, for the chosen multipliers, the error power that edge e delivers to module i ($A_e^{(i)}$: the columns of A_e that belong to module i). For every edge, $\sum_i \pi_{e,i} = -\langle A_e v_r, \lambda_e - \lambda_{r,e} \rangle$, which is zero when the required velocities of the modules that e joins are compatible at the actual configuration.

Proof. Subtracting (4) from (2) and using (5),

$$M\dot{s} + C(q, v)s + \mathcal{D}s = A^\top(\lambda - \lambda_r). \quad (7)$$

By (1), $\dot{W} = s^\top (M\dot{s} + Cs) = -s^\top \mathcal{D}s + (As)^\top(\lambda - \lambda_r)$, and $As = -Av_r$: (i). (ii) Replacing $\dot{v}_r$ by $\dot{v}_r + M^{-1}A^\top \gamma$ changes the left side of (4) by $A^\top \gamma$; with the same τ_r , λ_r changes by γ , and the term of (i) by $\langle Av_r, \gamma \rangle$. (iii) s satisfies $As = -Av_r$, and the least value of $|x|_M$ subject to $Ax = b$, $b \in \text{range } A$, is $|b|_{(AM^{-1}A^\top)^\dagger}$, attained at $x = M^{-1}A^\top(AM^{-1}A^\top)^\dagger b = -(I - P)v_r$ for $b = -Av_r$. (iv) is (i) written module by module, M, R, B, K being block diagonal and $C = \text{blkdiag}(C_i(q_i, v_i))$. $\square$

The identity combines the energy computation of the composite-error law [29] with explicit constraint forces and the VDC subsystem construction: (iv) is the cancellation of virtual power flows [42]. Its decisive hypothesis is reference admissibility at the measured configuration, $v_r = E(q)u_r$; a reference admissible at a desired configuration $q_d \neq q$ in general violates it, and (iii) gives the unavoidable error.

Except for the fully constrained case in Theorem 11, assume $d \geq 1$. With an admissible reference $v_r = E(q)u_r$,

$s = E\sigma$ with $\sigma = u - u_r$, $W = \frac{1}{2}\sigma^\top \bar{M}\sigma$ and, by (6), $\dot{W} = -\sigma^\top \bar{\mathcal{D}}\sigma$, $\bar{\mathcal{D}} = E^\top \mathcal{D}E = L^\top KL + \bar{R}$. Let

$$\underline{\kappa}(q) = \lambda_{\min}(\bar{\mathcal{D}}, \bar{M}), \quad \bar{\kappa}(q) = \lambda_{\max}(\bar{\mathcal{D}}, \bar{M}). \quad (8)$$

Equivalently, $\underline{\kappa}(q) = \max\{\kappa : v^\top (\mathcal{D} - \kappa M)v \geq 0 \text{ for all } v \in \mathcal{T}(q)\}$, the form whose decomposition is studied in Section 4.

Proposition 4 (Rank condition and rates). *Let Assumption 1 hold.*

(a) Characterization. *At a regular configuration the following are equivalent: (a1) $\ker A(q) \cap \ker B^\top = \{0\}$, i.e., with all actuators locked the mechanism has no admissible motion; (a2) $\text{rank } L(q) = d$; (a3) the required balance (4) has a solution (τ_r, λ_r) for every admissible v , $v_r = Eu_r$ and every $\dot{u}_r \in \mathbb{R}^d$; (a4) $L^\top KL \succ 0$ for one, equivalently for every, $K \succ 0$. They imply $\underline{\kappa}(q) \geq \lambda_{\min}(L^\top KL, \bar{M}) > 0$. In general $\underline{\kappa}(q) > 0$ if and only if $\ker A(q) \cap \ker B^\top \cap \ker R = \{0\}$.*

(b) Two-sided rate. *Along every motion with an admissible reference, $-2\bar{\kappa}(q)W \leq \dot{W} \leq -2\underline{\kappa}(q)W$, both inequalities being attained at every configuration. Hence for $t \geq t_0$*

$$\begin{aligned} e^{-\int_{t_0}^t \bar{\kappa}(q(r)) dr} |s(t_0)|_M &\leq |s(t)|_M \\ &\leq e^{-\int_{t_0}^t \underline{\kappa}(q(r)) dr} |s(t_0)|_M. \end{aligned} \quad (9)$$

(c) Necessity. *If (a1) fails at q , the accelerations $\dot{u}$ that efforts can produce at (q, u) form an affine subspace of dimension $\text{rank } L < d$; there are admissible references that no control follows exactly, and errors $\sigma \neq 0$ at which the actuators dissipate nothing.*

(d) Internal-effort invariance. *The reduced error obeys $\bar{M}\dot{\sigma} + E^\top(CE + M\dot{E})\sigma + \bar{\mathcal{D}}\sigma = 0$. For the same admissible reference and initial state, different solutions of (4) generate the same motion; their efforts form an affine family with direction $\ker L^\top$, of dimension $p - \text{rank } L$.*

Proof. (a) $L\sigma = 0$ with $\sigma \neq 0$ means $0 \neq E\sigma \in \ker A \cap \ker B^\top$: (a1) $\Leftrightarrow$ (a2). A vector x lies in $B\mathbb{R}^p + \text{range } A^\top$ if and only if $E^\top x \in \text{range } L^\top$. Premultiplied by $E^\top$, the left side of (4) is $\bar{M}\dot{u}_r$ plus terms independent of $\dot{u}_r$; it sweeps $\mathbb{R}^d$, so (4) is always solvable exactly when $\text{range } L^\top = \mathbb{R}^d$: (a3). (a4) is (a2). If $L\sigma = 0$, $\sigma \neq 0$, then $\sigma^\top \bar{\mathcal{D}}\sigma = \sigma^\top \bar{R}\sigma$, which gives the condition with R . (b) follows from $\underline{\kappa}\sigma^\top \bar{M}\sigma \leq \sigma^\top \bar{\mathcal{D}}\sigma \leq \bar{\kappa}\sigma^\top \bar{M}\sigma$, with equality for the extreme eigenvectors. (c) The reduced plant gives $\bar{M}\dot{u} = L^\top \tau + (\text{terms without } \tau)$. (d) Premultiply (7) by $E^\top$ and use $\dot{s} = E\dot{\sigma} + \dot{E}\sigma$. $\square$

Condition (a1) is a statement about the assembled graph and is checked after assembly. The number $p - d$ of superfluous actuators does not enter the certificate: it is the dimension of the internal efforts, available for load distribution at a fixed requested motion [26].

Proposition 5 (Residuals). *Suppose that (4) holds up to a measurable, locally bounded residual $\delta(t) \in \mathbb{R}^n$, its right-hand side being $B\tau_r + A^\top \lambda_r + \delta(t)$ (model error, omitted forces, sampling). Then $\dot{W} = -\sigma^\top \bar{D}\sigma - \sigma^\top E^\top \delta$. If $E^\top \delta = \Psi(t)\sigma + \eta(t)$, with Ψ and η measurable and locally bounded, and $\lambda_{\min}(\bar{D} + \text{sym } \Psi, M) \geq \kappa_0 > 0$, then*

$$|s(t)|_M \leq e^{-\kappa_0(t-t_0)} |s(t_0)|_M + \frac{1}{\kappa_0} \sup_{[t_0, t]} |\eta|_{\bar{M}^{-1}}. \quad (10)$$

Moreover $|E^\top x|_{\bar{M}^{-1}} = |P^\top x|_{M^{-1}} \leq |x|_{M^{-1}}$: a residual in range $A^\top$ (an error of the internal forces or of the load distribution) has no effect on the motion.

Proof. The residual adds $-\delta$ to the right of (7), and $\dot{W} = -\sigma^\top (\bar{D} + \text{sym } \Psi)\sigma - \sigma^\top \eta \leq -2\kappa_0 W + |s|_M |\eta|_{\bar{M}^{-1}}$ by the Cauchy–Schwarz inequality in the metric M . With D^+ the upper right Dini derivative, $D^+ |s|_M \leq -\kappa_0 |s|_M + |\eta|_{\bar{M}^{-1}}$ for almost all t ; integrating gives (10). The norm identity follows from $P^\top = M E M^{-1} E^\top$. $\square$

4 Compositional certificates

4.1 Payload stars

Consider a payload held by N robots, every edge joining one robot with the payload or with the ground. Then $\mathcal{T} = \{v : (v_0, v_i) \in \mathcal{T}_i, i = 1, \dots, N\}$, where $\mathcal{T}_i \subset \mathbb{R}^{n_0+n_i}$ is the admissible space of robot i holding the payload alone, and $\mathcal{D} = \text{blkdiag}(0, \mathcal{D}_1, \dots, \mathcal{D}_N)$, the payload having neither actuators nor damping.

Definition 6. For $\kappa \geq 0$ and a symmetric $\Theta \in \mathbb{R}^{n_0 \times n_0}$, module i is (κ, Θ) -dissipative at q if

$$v_i^\top (\mathcal{D}_i - \kappa M_i) v_i \geq v_0^\top \Theta v_0 \quad \text{for all } (v_0, v_i) \in \mathcal{T}_i.$$

For $J = J^\top \succeq 0$ let $\kappa_i(J)$ be the supremum of the κ for which module i is $(\kappa, \kappa J)$ -dissipative: the certified standalone rate for inertia share J .

Θ measures the dissipation that the robot, at rate κ , has left for the payload; by Theorem 8(e), $\pi_i - s_0^\top \Theta_i s_0$ is a supply rate at the grasp port in the sense of [36], with storage W_i and dissipation rate $2\kappa W_i$.

Proposition 7 (Payload-space test). *Fix a regular configuration, its mode and gains, and $\kappa \geq 0$. Suppose every robot can realize every payload velocity. Parametrize its*

local space by $v_i = X_i v_0 + Z_i x_i$, where Z_i spans the local self-motions with $v_0 = 0$, and set

$$Q_i = \mathcal{D}_i - \kappa M_i, \quad F_i = Z_i^\top Q_i Z_i, \quad \Lambda_i = Z_i^\top Q_i X_i, \quad (11)$$

$$\Xi_i = X_i^\top Q_i X_i - \Lambda_i^\top F_i^\dagger \Lambda_i, \quad S(\kappa) = \sum_i \Xi_i - \kappa M_0. \quad (12)$$

Then $\underline{\kappa} \geq \kappa$ if and only if

$$F_i \succeq 0, \quad \text{range } \Lambda_i \subseteq \text{range } F_i \quad (\forall i), \quad S(\kappa) \succeq 0. \quad (13)$$

When (13) holds, $\Theta_i = \Xi_i$ satisfies Definition 6 and $\sum_i \Theta_i \succeq \kappa M_0$; for agreed $\alpha_i \geq 0$ with $\sum_i \alpha_i = 1$, $\Theta_i = \Xi_i - \alpha_i S$ has sum κM_0 .

Proof. At fixed v_0 the local form is $x_i^\top F_i x_i + 2x_i^\top \Lambda_i v_0 + v_0^\top X_i^\top Q_i X_i v_0$. Its infimum is finite for every v_0 exactly under the two local conditions in (13), and it equals $v_0^\top \Xi_i v_0$ [3]. Payload realizability completes any local choice to a team velocity, so global nonnegativity requires every local infimum to be finite; the self-motions of different robots are independent, and their separate minimization gives the team infimum $v_0^\top S(\kappa) v_0$. $\square$

Theorem 8 (Rates compose). *At a regular configuration:*

(a) Composition. *If module i is (κ, Θ_i) -dissipative for $i = 1, \dots, N$ and $\sum_i \Theta_i \succeq \kappa M_0$, then $\underline{\kappa} \geq \kappa$. In particular, for inertia shares $J_i \succeq 0$ with $\sum_i J_i = M_0$,*

$$\underline{\kappa} \geq \min_i \kappa_i(J_i), \quad \underline{\kappa} \geq \min_i \kappa_i(M_0/N) \geq \min_i \kappa_i(M_0). \quad (14)$$

(b) Exactness. *If every payload velocity can be followed by every robot and $\underline{\kappa} \geq \kappa$, there are symmetric $\Theta_1, \dots, \Theta_N$ with $\sum_i \Theta_i = \kappa M_0$ such that every module is (κ, Θ_i) -dissipative: the modular certificate loses nothing, including at $\kappa = \underline{\kappa}$.*

(c) Plug-and-play. *If a robot is attached to the payload of a team with rate κ_N , the gains of the others unchanged, then $\underline{\kappa}_{N+1} \geq \min\{\underline{\kappa}_N, \kappa_{N+1}(0)\}$.*

(d) Constraints. *If edges are added, $\mathcal{T}' \subseteq \mathcal{T}$ with the same M and $\mathcal{D}$, then $\underline{\kappa}' \geq \underline{\kappa}$ and $\bar{\kappa}' \leq \bar{\kappa}$.*

(e) Dissipation along motions. *Let the reference be admissible and module i be (κ, Θ_i) -dissipative at the configurations of a motion. Then $\dot{W}_i \leq -2\kappa W_i - s_0^\top \Theta_i s_0 + \pi_i$, where π_i is the error power that the grasp of robot i delivers to it (Proposition 3(iv)), and $\dot{W}_0 = -\sum_i \pi_i$. If $\sum_i \Theta_i \succeq \kappa M_0$, the sum gives $\dot{W} \leq -2\kappa W$.*

Proof. (a) For $s \in \mathcal{T}$, $s^\top(\mathcal{D} - \kappa M)s = \sum_{i \geq 1} s_i^\top(\mathcal{D}_i - \kappa M_i)s_i - \kappa s_0^\top M_0 s_0 \geq s_0^\top(\sum_i \Theta_i - \kappa M_0)s_0 \geq 0$. With $\kappa = \min_i \kappa_i(J_i)$ module i is $(\kappa, \kappa J_i)$ -dissipative, and $\kappa_i(J)$ does not increase with J . (b) Proposition 7 gives $S(\kappa) \succeq 0$; take $\Theta_i = \Xi_i - S/N$. (c) The restriction of $s \in \mathcal{T}_{(N+1)}$ to the first N robots and the payload lies in $\mathcal{T}_{(N)}$; add the two inequalities. (d) The extreme values of a Rayleigh quotient over a subspace lie between those over the space. (e) By Proposition 3(iv), $\dot{W}_i = -s_i^\top \mathcal{D}_i s_i + \pi_i$, the error powers of the ground edges vanishing and those of a grasp adding up to zero; Definition 6 gives $-s_i^\top \mathcal{D}_i s_i \leq -\kappa s_i^\top M_i s_i - s_0^\top \Theta_i s_0$. $\square$

Remark 9 (Scalar and matrix shares). Once M_0 and scalar shares α_i have been agreed, each robot tests its own share without the other robots' matrices. Scalar shares can be very conservative: a robot alone must support its share of every payload rotation, whereas the team also uses coordinated grasp translations. Matrix supplies retain this directional coupling and may be indefinite. The augmented-object inertia [18] is a related mechanical aggregation; the present inequalities quantify closed-loop error dissipation.

Information and membership. At a trial rate, robot i forms (11) from its own $M_i, \mathcal{D}_i$, grasp/ground rows, and the shared payload configuration, and sends its feasibility status and $\Xi_i, n_0(n_0 + 1)/2$ numbers in an agreed payload-velocity chart. A coordinator tests $S \succeq 0$ with M_0 ; neither the robot matrices nor the assembled lift E are transmitted. At an unchanged state, cached summaries give

$$S_{\text{add}} = S + \Xi_{\text{new}}, \quad S_{\text{remove}} = S - \Xi_j, \quad (15)$$

while moving cells refresh their summaries. The benefit of the summaries is locality of information rather than arithmetic, and Section 5 uses it for design.

4.2 General cell architectures

For a general assembly let the modules be grouped into *cells* $V_1, \dots, V_k$ that cover them, every edge being assigned to one cell that contains the modules it joins. Write $v|_{V_c}$ for the velocities of the modules of cell c , $\mathbb{R}^{V_c}$ for their space, and $\mathcal{T}_c \subseteq \mathbb{R}^{V_c}$ for the admissible velocities of the edges of the cell, so that $\mathcal{T} = \{v : v|_{V_c} \in \mathcal{T}_c \text{ for all } c\}$. For a set S of modules, π_S restricts velocities to $\mathbb{R}^S$.

Definition 10. A *modular certificate of rate κ* is a family of quadratic forms θ_c on $\mathbb{R}^{V_c}$ with $\theta_c \geq 0$ on $\mathcal{T}_c$ and $\sum_c \theta_c(v|_{V_c}) = v^\top(\mathcal{D} - \kappa M)v$ for all $v \in \mathbb{R}^n$.

Each cell checks its own inequality, and a modular certificate of rate κ gives $\underline{\kappa} \geq \kappa$. For the cells $\{0, i\}$ of Theorem 8 the certificates are the supplies of Definition 6

with $\sum_i \Theta_i = \kappa M_0$. A *join tree* of the cells is a tree whose vertices are the cells and in which, for every module, the cells that contain it form a connected subtree; for a tree edge $\{c, c'\}$, $S_{cc'} = V_c \cap V_{c'}$ is its *separator*.

Theorem 11 (Join trees). *Let the cells admit a join tree.*

- (a) *For every $\kappa < \underline{\kappa}$ there is a modular certificate of rate κ (for every κ if $d = 0$).*
- (b) *Suppose in addition that the join tree can be chosen so that every separator velocity is realizable by the assembly: $\pi_{S_{cc'}}(\mathcal{T}) = \mathbb{R}^{S_{cc'}}$ for every tree edge $\{c, c'\}$. Then a modular certificate of rate κ exists for every $\kappa \leq \underline{\kappa}$, in particular at the endpoint $\kappa = \underline{\kappa}$. One is obtained by eliminating the leaves of the join tree with generalized Schur complements; for a payload star it gives the supplies of Proposition 7 with the slack $S(\kappa)$ assigned to the root cell.*

Proof. (a) Stack each cell's constraint rows in A_c . Strict positivity of $\mathcal{D} - \kappa M$ on $\mathcal{T} \setminus \{0\}$ gives, by Finsler's lemma, a finite $\beta > 0$ with

$$Q_\beta = \mathcal{D} - \kappa M + \beta \sum_c A_c^\top A_c \succ 0,$$

each A_c embedded in the assembled coordinates (otherwise unit vectors violating the inequality for $\beta \rightarrow \infty$ would have a limit in $\mathcal{T}$ contradicting strict positivity). The join-tree property makes the graph obtained by completing each cell a clique chordal, with every maximal clique contained in a cell, and classical clique-supported decomposition [33, Thm. 9.2] gives $Q_\beta = \sum_c U_c$, $U_c \succeq 0$ supported on V_c . Then $\theta_c = U_c - \beta A_c^\top A_c$ is nonnegative on $\ker A_c$, and the forms sum to $\mathcal{D} - \kappa M$.

(b) Let $\kappa \leq \underline{\kappa}$, so that $q(v) = v^\top(\mathcal{D} - \kappa M)v \geq 0$ on $\mathcal{T}$, and write $q = \sum_c p_c$ with p_c a form supported on V_c (assign the form of each module to one cell that contains it). We prove, by induction on the number of cells, that every family of cell-supported forms whose sum is nonnegative on $\mathcal{T}$ admits a modular certificate when separator realizability holds. One cell: $\theta = p$. Otherwise let c be a leaf of the join tree with neighbor c' , $S = S_{cc'}$, and $P = V_c \setminus S$. By the join-tree property no module of P belongs to another cell. Let $\mathcal{T}'$ be the admissible space of the other cells, on the velocities of $V \setminus P$; then $\mathcal{T} = \{(w, v_P) : w \in \mathcal{T}', (w_S, v_P) \in \mathcal{T}_c\}$. Since $\pi_S(\mathcal{T}) = \mathbb{R}^S$, also $\pi_S(\mathcal{T}_c) = \mathbb{R}^S$ and $\pi_S(\mathcal{T}') = \mathbb{R}^S$, and $\mathcal{T}_c = \{(\varsigma, X\varsigma + Zx)\}$ for a particular-solution matrix X and a basis Z of the motions of P with $\varsigma = 0$. Write $p_c(\varsigma, X\varsigma + Zx) = x^\top Fx + 2x^\top \Lambda\varsigma + \varsigma^\top G\varsigma$. Nonnegativity of q on $\mathcal{T}$ requires, for each $w \in \mathcal{T}'$, a finite infimum over x ; as w_S ranges over $\mathbb{R}^S$, this means $F \succeq 0$ and range $\Lambda \subseteq \text{range } F$, and the infimum is $\varsigma^\top \Xi \varsigma$ with $\Xi = G - \Lambda^\top F^\dagger \Lambda$ [3]. Put $\theta_c(\varsigma, v_P) = p_c(\varsigma, v_P) - \varsigma^\top \Xi \varsigma$, which is nonnegative on $\mathcal{T}_c$, and replace $p_{c'}$ by $p_{c'} + \varsigma^\top \Xi \varsigma$, admissible because $S \subseteq V_{c'}$. Every $w \in \mathcal{T}'$ extends to $(w, Xw_S) \in \mathcal{T}$, so $\mathcal{T}' = \pi_{V \setminus P}(\mathcal{T})$, and the remaining

forms are nonnegative on $\mathcal{T}'$: for $w \in \mathcal{T}'$ they sum to $\inf_x q(w, Xw_S + Zx) \geq 0$; the remaining cells keep a join tree, and their separators S' satisfy $\pi_{S'}(\mathcal{T}') = \pi_{S'}(\mathcal{T}) = \mathbb{R}^{S'}$. Induction gives forms for the remaining cells; each Ξ is subtracted once and added once, so the sum is q . For a payload star, S is the payload, separator realizability is payload realizability, and $\Xi = \Xi_i$. $\square$

Theorem 11(a) is where classical chordal theory enters. At $\kappa = \underline{\kappa}$ the form is only semidefinite on $\mathcal{T}$ and a finite penalty need not exist, so (b) needs additional structure. The realizability hypothesis cannot be dropped:

Proposition 12 (Endpoint failure without realizability). *Even for a two-robot payload star with positive-definite gains and full team actuation, modular certificates need not exist at $\underline{\kappa}$. Let the payload velocity be y , $v_1 = (x, u)$, $v_2 = (z, w)$, and*

$$\begin{aligned} M_0 &= 1, \quad M_1 = \begin{bmatrix} 1 & -1 \\ -1 & 2 \end{bmatrix}, \quad M_2 = I_2, \\ \mathcal{D}_1 &= \text{diag}(1, 2), \quad \mathcal{D}_2 = 2I_2, \\ \mathcal{T}_1 &= \{(y, x, u) : u = y\}, \\ \mathcal{T}_2 &= \{(y, z, w) : w = y, w = 0\}. \end{aligned} \quad (16)$$

Then $d = 2$ and $\underline{\kappa} = 1$. No finite supplies attain $\kappa = 1$, supplies exist for every $0 \leq \kappa < 1$, and every family approaching $\kappa = 1$ has unbounded supply coefficients.

Proof. With $B = [0_{1 \times 4}; I_4]$, $K = \text{blkdiag}(\mathcal{D}_1, \mathcal{D}_2)$ and $R = 0$, the payload is unactuated and rank $L = d = 2$. On $\mathcal{T}$, $y = u = w = 0$, so the kinetic and dissipation forms are $x^2 + z^2$ and $x^2 + 2z^2$; hence $\underline{\kappa} = 1$. At this rate robot 1's local form is $2xy$, unbounded below for fixed $y \neq 0$, so no finite Θ_1 satisfies Definition 6. For $\delta = 1 - \kappa > 0$ its local form is $\delta(x + \kappa y / \delta)^2 + (2\delta - \kappa^2 / \delta)y^2$, so $\Theta_1 = 2\delta - \kappa^2 / \delta$ and $\Theta_2 = \kappa - \Theta_1$ are feasible (robot 2 has $y = 0$ and local form $(2 - \kappa)z^2 \geq 0$). Every feasible Θ_1 is at most $2\delta - \kappa^2 / \delta \rightarrow -\infty$, and $\Theta_1 + \Theta_2 \geq \kappa$ forces $\Theta_2 \rightarrow +\infty$. $\square$

Here robot 2 pins the payload, which robot 1's own constraint space does not see: the endpoint fails because a cell certifies motions that the assembly forbids.

The next result shows that the join tree is also necessary. Let G be the graph on the modules in which two modules are adjacent when some cell contains both.

Theorem 13 (Cell architectures). *For a prescribed family of cells, the following are equivalent:*

(i) *the cells admit a join tree;*

(ii) *for every choice of module coordinates, inertias, dissipations, and edges assigned to the cells (an edge may join any modules of its cell), a modular certificate exists for every $\kappa < \underline{\kappa}$.*

If (i) fails, there are assemblies with these cells and $\underline{\kappa} > 0$ for which modular certificates exist only for $\kappa \leq 0$.

Proof. (i) $\Rightarrow$ (ii) is Theorem 11(a). For the converse, suppose that G is chordal and that every clique of G lies in a cell. A maximal clique of G lies in a cell, which is itself a clique of G ; by maximality the two coincide, so every maximal clique is a cell. A chordal graph has a clique tree on its maximal cliques [33]; attaching every other cell to a maximal clique containing it gives a join tree (this is the classical characterization of acyclic hypergraphs [6]). Therefore, if (i) fails, either (A) G has a chordless cycle $j_1, \dots, j_k$, $k \geq 4$, or (B) some clique of G lies in no cell; in case (B) let K be such a clique of least cardinality m , so that $m \geq 3$ and every proper subset of K lies in a cell. Let J be the modules of the cycle or of K .

Let every module have velocities $(x_j, z_j) \in \mathbb{R}^2$ and $M_j = I$, with $\mathcal{D}_j = \text{diag}(0, \kappa_z)$ for $j \in J$ and $\mathcal{D}_j = \kappa_z I$ otherwise, $\kappa_z > 0$. Edges act only on the x_j , $j \in J$: in case (A), $x_{j_i} = x_{j_{i+1}}$ for $i < k$ and $x_{j_k} = -x_{j_1}$, each assigned to a cell that contains its two modules; in case (B), for every $j \in K$ the edge $\sum_{i \in K \setminus \{j\}} x_i = 0$, assigned to a cell that contains $K \setminus \{j\}$ (for $m = 3$ it joins two modules; for $m \geq 4$, e.g., a fixed-length cable over the modules of $K \setminus \{j\}$). In both cases the edges force $x_J = 0$ (in (B), subtracting two edges gives $x_j = x_{j'}$), so $\underline{\kappa} = \kappa_z$.

Define Y , symmetric and supported on the x_J coordinates: in case (A), $Y_{jj} = 1$, $Y_{j_i j_{i+1}} = Y_{j_{i+1} j_i} = 1$ for $i < k$, $Y_{j_k j_1} = Y_{j_1 j_k} = -1$, and all other entries zero; in case (B), $Y = (1+t)I - t\mathbf{1}\mathbf{1}^\top$ on x_K with $t = 1/(m-2)$. For every cell c , the principal submatrix Y_c of Y on the coordinates of V_c is positive semidefinite with range in $\mathcal{T}_c$. In case (A), a cell contains at most two modules of the chordless cycle, consecutive if two; Y_c is then 1, or $\begin{bmatrix} 1 & \pm 1 \\ \pm 1 & 1 \end{bmatrix}$ with range spanned by $(1, \pm 1)$, which satisfies the edge between the two modules whether or not c carries it. In case (B), a cell contains at most $m-1$ modules of K , and an $r \times r$ principal submatrix of Y has eigenvalues $1+t-tr \geq 0$ and $1+t$. A cell that carries the edge of j contains $K \setminus \{j\}$ and no other edge (two would place K in the cell), and its $Y_c = (1+t)I_{m-1} - t\mathbf{1}\mathbf{1}^\top$ has kernel spanned by $\mathbf{1}$, so its range is that edge's constraint set.

Each Y_c is a nonnegative combination of $yy^\top$ with $y \in \mathcal{T}_c$. Hence, for a modular certificate of rate κ ,

$$0 \leq \sum_c \langle \theta_c, Y_c \rangle = \langle \mathcal{D} - \kappa M, Y \rangle = -\kappa \text{tr} Y = -\kappa |J|,$$

because $\mathcal{D} - \kappa M$ is diagonal with entries $-\kappa$ on the x_J coordinates. Thus $\kappa \leq 0$. $\square$

Remark 14 (Duality and factor width). The matrix Y certifies the failed decomposition, as in the duality between clique decomposition and positive semidefinite completion [2, 33]. For pairwise cells on a ring, the lift of Theorem 11(a) is $\beta L_s - \kappa I$ on the x -block, with L_s the signed Laplacian of an unbalanced cycle; its comparison matrix has the eigenvalue $-\kappa$, so it has no factor-width-two decomposition [7]; Theorem 13 extends this obstruction to arbitrary cell architectures.

Remark 15 (Edge arity). Construction (A), and (B) with $m = 3$, use edges that join two modules; for $m \geq 4$, (B) uses edges that join $m - 1$ modules, and this cannot be avoided in general. The matrices Y supported on the cell blocks, with every block Y_c positive semidefinite with range in $\mathcal{T}_c$, form the dual cone of the sums of cell forms, so a modular certificate of rate κ exists if $\langle \mathcal{D} - \kappa M, Y \rangle > 0$ for every such $Y \neq 0$. Let the cells be the $(m - 1)$ -subsets of $m \geq 4$ modules, and let every edge join two modules. A block $Y_c \succeq 0$ annihilated by the rows of an edge forces the same on every block that contains both of its modules, and every module shares a cell with any two others; hence range $Y \subseteq \mathcal{T}$. In a basis of $\mathcal{T}$ adapted to the subspaces $\mathcal{U}_j$ of motions of module j alone, $Y = NZN^\top$ with Z positive semidefinite on the coordinates outside each $\mathcal{U}_j$, whereas $N^\top(\mathcal{D} - \kappa M)N \succ 0$ ($\kappa < \underline{\kappa}$) couples no two $\mathcal{U}_j$ and splits into positive definite parts, each supported outside some $\mathcal{U}_{j'}$. Thus $\langle \mathcal{D} - \kappa M, Y \rangle \geq 0$, with equality only if $Z = 0$, since any two $\mathcal{U}_i, \mathcal{U}_k$ lie outside a common $\mathcal{U}_{j'}$. Certificates therefore exist for every $\kappa < \underline{\kappa}$ although the cells admit no join tree: for two-module edges, property (ii) of Theorem 13 holds for a strictly larger class of architectures, which we do not characterize here.

5 Compositional gain synthesis

The cell inequalities are linear in the dissipation, and the dissipation is affine in the actuator gains: $\mathcal{D}_i(K_i) = B_i K_i B_i^\top + R_i$. The certificate therefore extends from verification to design. Let $\mathcal{K}_i$ be a closed convex set of admissible gains of robot i (for example, diagonal positive gains within actuator limits), $\mathcal{K} = \{\text{blkdiag}(K_i) : K_i \in \mathcal{K}_i\}$, and let $\mathcal{C}$ be a finite set of regular configurations of a mode, sampled from the motions to be certified. Write $\underline{\kappa}(K; q)$ for the rate (8) with gains K at q , and consider

$$\mathcal{G}(\kappa) = \{K \in \mathcal{K} : \underline{\kappa}(K; q) \geq \kappa \text{ for all } q \in \mathcal{C}\}, \quad (17)$$

$$\mathcal{G}_m(\kappa) = \{K \in \mathcal{K} : \text{for every } q \in \mathcal{C} \text{ there is a modular certificate of rate } \kappa \text{ for } \mathcal{D}(K) \text{ at } q\}. \quad (18)$$

Theorem 16 (Compositional gain synthesis). *Let the cells admit a join tree at every $q \in \mathcal{C}$.*

- (a) $\mathcal{G}_m(\kappa) \subseteq \mathcal{G}(\kappa)$. Both sets are described by LMIs that are jointly linear in the gains, the cell forms, and κ : $E^\top(\mathcal{D}(K) - \kappa M)E \succeq 0$ for $\mathcal{G}$, and $N_c^\top \theta_c N_c \succeq 0$ with

range $N_c = \mathcal{T}_c(q)$ together with the linear identity of Definition 10 for $\mathcal{G}_m$.

- (b) If separator realizability holds at every $q \in \mathcal{C}$, then $\mathcal{G}_m(\kappa) = \mathcal{G}(\kappa)$. Without it, $\mathcal{G}_m(\kappa) \supseteq \{K \in \mathcal{K} : \underline{\kappa}(K; q) > \kappa \forall q \in \mathcal{C}\}$.
(c) Let $J : \mathcal{K} \rightarrow \mathbb{R}$ be convex. The infima of J over $\mathcal{G}(\kappa)$ and over $\mathcal{G}_m(\kappa)$ are equal under the hypothesis of (b), and without it whenever some $K \in \mathcal{K}$ has $\underline{\kappa}(K; q) > \kappa$ for all $q \in \mathcal{C}$. For every budget b , the suprema of the rates κ for which $\mathcal{G}(\kappa)$, respectively $\mathcal{G}_m(\kappa)$, contains a gain with $J(K) \leq b$ are equal.
(d) For a payload star, $\mathcal{G}_m(\kappa)$ is the set of gains for which there are supplies Θ_i^q with, for every $q \in \mathcal{C}$ and all $(v_0, v_i) \in \mathcal{T}_i(q)$,

$$v_i^\top (B_i K_i B_i^\top + R_i - \kappa M_i(q)) v_i \geq v_0^\top \Theta_i^q v_0, \quad (19)$$

$$\sum_{i=1}^N \Theta_i^q \succeq \kappa M_0(q). \quad (20)$$

Each (19) is an LMI of order $\dim \mathcal{T}_i(q)$ in $(K_i, \Theta_i^q, \kappa)$ that involves only robot i 's data; the robots are coupled only through (20), of order n_0 .

Proof. (a) A modular certificate of rate κ gives $\underline{\kappa} \geq \kappa$ at its configuration. Linearity: $\mathcal{D}(K)$ is affine in K , and positivity of θ_c on $\mathcal{T}_c$ is equivalent to $N_c^\top \theta_c N_c \succeq 0$. (b) For $K \in \mathcal{G}(\kappa)$, Theorem 11(b) at each q gives a certificate of rate $\kappa \leq \underline{\kappa}(K; q)$; without realizability Theorem 11(a) applies when $\kappa < \underline{\kappa}(K; q)$. (c) Under the hypothesis of (b) the sets coincide. Otherwise let $K_s \in \mathcal{K}$ have $\underline{\kappa}(K_s; q) > \kappa$ on $\mathcal{C}$, and let $K \in \mathcal{G}(\kappa)$. For $t \in (0, 1]$, $K_t = (1 - t)K + tK_s \in \mathcal{K}$, and $E^\top(\mathcal{D}(K_t) - \kappa M)E \succ 0$ at every $q \in \mathcal{C}$ because $\mathcal{D}$ is affine; so $K_t \in \mathcal{G}_m(\kappa)$ by (b), and $J(K_t) \leq (1 - t)J(K) + tJ(K_s) \rightarrow J(K)$ as $t \rightarrow 0$. With (a), the infima are equal. If κ is below the supremum for $\mathcal{G}$, some K with $J(K) \leq b$ has $\underline{\kappa}(K; q) > \kappa$ on $\mathcal{C}$ and lies in $\mathcal{G}_m(\kappa)$ by (b). (d) For the cells $\{0, i\}$, a modular certificate is a family of supplies of Definition 6, with $\mathcal{D}_i = B_i K_i B_i^\top + R_i$, that sum to κM_0 (Section 4.2); the inequality in (20) is equivalent, since lowering one Θ_i^q preserves (19). $\square$

Remark 17 (Distributed solution). In (d), a minimum-gain design $\min \sum_i J_i(K_i)$ has the sharing form of [8, Sec. 7.3]: each robot solves its own LMIs (19) for (K_i, Θ_i^q) and exchanges only the supplies, $n_0(n_0 + 1)/2$ numbers per configuration, with a coordinator that projects their sum onto (20), as in the supply-rate search of [22]. A robot that joins a team with spare supply needs only its own inequality (Theorem 8(c)). The guarantee is for the configurations in $\mathcal{C}$; certification over continuous regions of configurations, for instance by polynomial certificates in q , is not addressed here.

6 Realization by joint-level loops

6.1 Fitted reference

The certificates require an admissible reference and its required balance. Joint and wheel controllers, however, are given requested actuator rates. A position loop on a joint with coordinate a , gains k_p , k_d and target $a_d(t)$ applies $-k_d(\dot{a} - w)$ with $w = \dot{a}_d + (k_p/k_d)(a_d - a)$, and a speed loop on a wheel applies $-k_w(\dot{a} - w)$. Such loops are

$$\tau = \tau_f - K(B^\top v - w), \quad K = \text{blkdiag}(K_i) \succ 0, \quad (21)$$

with requested rates $w \in \mathbb{R}^p$ and a feedforward τ_f . When $p > d$, w is in general not of the form $L(q)u$: the loops ask for a motion that the constraints do not admit. In a team, each robot generates its requests from its own targets, so the requests of different robots are generally inconsistent with the grasp and rolling constraints that couple them.

Assumption 18 (Requests). $w = \omega(t, q, \xi)$ and $\dot{\xi} = \chi(t, q, v, \xi, y_m(t))$, with ω and χ continuously differentiable, controller states ξ , and measured signals y_m . Signals of y_m that the present effort influences, such as a wrist wrench, enter only the rates χ_j of states ξ_j on which ω does not depend, so that $\partial_{\xi} \omega \chi$ does not depend on them.

Then $\dot{w} = \partial_t \omega + \partial_q \omega v + \partial_{\xi} \omega \chi$ is computable from time and measurements before the present effort is applied, so the computation below contains no algebraic loop. The fitted reference of the loops is

$$u_r = H^{-1} L^\top K w, \quad v_r = E(q) u_r, \quad w_\perp = w - L u_r, \quad (22)$$

with $H = L^\top K L$, under the full-rank condition of Proposition 4(a): $L u_r$ uniquely minimizes $|L u - w|_K^2$. The fit is the admissible motion closest to the requests in the metric of the gains: a stiff loop pulls the reference toward its own request more than a soft one.

Theorem 19 (Joint-level loops). *Let Assumptions 1 and 18 and (a1) hold along a motion of (2) under (21), with u_r , v_r , $w_\perp$ given by (22).*

- (a) Decomposition. $L^\top K w_\perp = 0$ and $\tau = (\tau_f + K w_\perp) - K B^\top (v - v_r)$: the loops are the law (5) for the admissible reference v_r , with the internal effort $K w_\perp \in \ker L^\top$ added to the feedforward, and u_r is the only u for which $K(w - L u)$ is an internal effort. Relative to loops asked for $L u_r$, the original loops add $K w_\perp$ to the effort and $-B K w_\perp$ to the total constraint force, at the same reduced acceleration.
- (b) Fitted feedforward. Along the motion $\dot{v}_r = E \dot{u}_r + \dot{E} u_r$ with

$$\dot{u}_r = H^{-1} (\dot{L}^\top K w_\perp + L^\top K (\dot{w} - \dot{L} u_r)), \quad \dot{L} = B^\top \dot{E}, \quad (23)$$

a function of time and measured quantities. If τ_f solves (4) for this v_r (one exists by (a3)), then $\dot{W} = -s^\top \mathcal{D} s$ with $s = v - v_r$, and Proposition 4(b),(d) and the certificates of Sections 4–5 apply, for every $K \succ 0$ and all requests of Assumption 18.

- (c) Other feedforwards. For any τ_f , $\dot{W} = -\sigma^\top \bar{\mathcal{D}} \sigma - \sigma^\top E^\top \delta$ with the defect $\delta = M \dot{v}_r + C(q, v) v_r + g + R v_r - B \tau_f$, to which Proposition 5 applies.
- (d) Desired-motion feedforward at rest. Let the loops be position loops with targets at rest, $w = K^{-1} K_p (a_d - B^\top q)$ with $K_p = K_p^\top \succeq 0$, and let τ_f solve (4) for the fitted reference with a symmetric $\hat{M}(q)$ in place of M , all other terms exact. Then, at every state with $K_p (a_d - B^\top q) = 0$,

$$\dot{W} = -\sigma^\top (\bar{\mathcal{D}} + E^\top (\hat{M} - M) E G_p) \sigma, \quad (24)$$

with $G_p = H^{-1} L^\top K_p L$. The feedforward of the desired motion, which supplies at rest only the effort that rest requires $(B \tau_f + A^\top \lambda_f = g)$, is the case $\hat{M} = 0$, with the form $\bar{\mathcal{D}} - \bar{M} G_p$, which can be indefinite.

Proof. (a) $L^\top K(w - L u) = L^\top K w - H u$ vanishes for $u = u_r$ only, and on admissible velocities $B^\top v - w = L(u - u_r) - w_\perp$. Since $E^\top B K w_\perp = 0$, the generalized force $B K w_\perp$ lies in range $A^\top$; by Lemma 2, $A^\top \lambda = M \dot{v} + C v + g + R v - B \tau$, which changes by $-B K w_\perp$ when τ changes by $K w_\perp$ at the same motion. (b) Differentiating $H u_r = L^\top K w$ with $\dot{H} = \dot{L}^\top K L + L^\top K \dot{L}$ gives (23). Let (τ_f, λ_f) solve (4) for v_r and $A^\top \lambda_\perp = -B K w_\perp$, solvable by (a). Then $(v_r, \tau_f + K w_\perp, \lambda_f + \lambda_\perp)$ satisfies (4), the loops are (5) with this triple by (a), and v_r is admissible: Proposition 3(ii). (c) With $\tau_r = \tau_f + K w_\perp$ and $\lambda_r = \lambda_\perp$ the balance (4) holds up to δ . (d) Here $u_r = H^{-1} L^\top K_p e_a$ with $e_a = a_d - B^\top q$, and $\dot{e}_a = -L u$. Where $K_p e_a = 0$: $u_r = 0$, $v_r = 0$, $\sigma = u$, $\dot{u}_r = H^{-1} L^\top K_p \dot{e}_a = -G_p u$ and $\dot{v}_r = E \dot{u}_r$, so that $\delta = (M - \hat{M}) E \dot{u}_r$ up to an element of range $A^\top$, and $E^\top \delta = -E^\top (M - \hat{M}) E G_p \sigma$; insert in (c). $\square$

Remark 20 (Connections and gain-dependent damping). The effort split in (a) is the weighted projection of [16]: its force projector $K L (L^\top K L)^{-1} L^\top$ maps $K(w - L u)$ to $K(L u_r - L u)$ and leaves $K w_\perp$ as internal loading. Theorem 19(b) differentiates the fitted request and evaluates its required balance, extending composite-reference inverse dynamics [12, 29] to requests that the constraints do not admit. For one joint of inertia m with $R = 0$ and rate gain K , the fitted effort is $m(\ddot{a}_d + (k_p/K)(\dot{a}_d - \dot{a}))$, so the total rate feedback is $K + m k_p/K$, a gain-dependent damping that the certificate accounts for. With $k_d = 2\sqrt{k_p m}$, the fitted loop with $K = k_d$ has damping ratio 1.25, and the fitted loop with $K = k_d/2$ has the ratio 1 of the




Fig. 2. Fitted-reference realization of the joint and wheel loops. The requests are projected by (22); (23) and the required balance (4) give τ_f ; each actuator adds damping of its measured rate toward its request. The fit uses the assembled lift and the gains, the feedforward also the assembled dynamics; $Kw_\perp$ changes only the internal loading.

desired-motion loop with $K = k_d$. The certificate and the synthesis are compositional, whereas evaluating τ_f uses the team's E , $\dot{E}$, H^{-1} , and inverse dynamics.

6.2 Controller realization

In the missions of Section 7, arm and lift position loops and wheel-speed loops generate the requests in (21). An arm target solves its grasp closure for the desired payload pose at the measured base and lift pose. A force admittance driven by the filtered internal wrist wrench adjusts this target; its offset and filter belong to ξ , as Assumption 18 requires. Wheel requests follow the compiled rolling rows for a base pose controller with integral action. The fit reconciles these independent requests at the actual configuration (Fig. 2).

Every 10 ms the controller evaluates E , the fitted reference, $\dot{E}$ by centered differences along the measured v , $\dot{w}$, and the required effort; tree inverse dynamics evaluates $C(q, v)v_r$ by polarization and distributes the payload wrench by the declared load-sharing rule. In every plant step the loops use advanced requests and the effort $\tau_{f,k} + \Gamma_k(\dot{w}(t) - \dot{w}_k)$, $\Gamma = K L H^{-1} \bar{M} H^{-1} L^T K$, with Γ_k frozen at renewal k ; since $L^T \Gamma = \bar{M} H^{-1} L^T K$, this advances the reduced effort required by $\dot{w}$. Renewal every 10 ms makes the implementation a sampled version of Theorem 19(b); (26) bounds the cost of sampling, and Section 7.5 compares it with per-step renewal.

6.3 Mode changes and sampled renewal

At given event times t_k the mode changes from ν^- to ν^+ ; the configuration is continuous, and the holonomic rows of ν^+ hold at $q(t_k)$. Edges that open cause no impulse. Edges that close do so plastically: $M(v^+ - v^-) = A^{+\top} \gamma$, $A^+ v^+ = 0$.

Lemma 21 (Carnot, projector form). $v^+ = P^+ v^-$, where P^+ is the projector (3) of the mode ν^+ at $q(t_k)$; v^+

is unique although γ need not be, and the kinetic energy drops by $\frac{1}{2}|v^- - v^+|_M^2$. If $\mathcal{T}^- \subseteq \mathcal{T}^+$ (edges only open), $v^+ = v^-$.

Proof. $v^+ - v^- \in M^{-1} \text{range } A^{+\top}$, the M -orthogonal complement of $\mathcal{T}^+$, and $v^+ \in \mathcal{T}^+$: v^+ is the M -orthogonal projection of v^- ; the rest is Pythagoras' theorem. $\square$

This is the classical rule for suddenly imposed constraints [10], in a form that admits dependent rows. A holonomic edge that closes by *latching at the present relative pose*, its constraint function being defined by $q(t_k)$, holds in position by construction; this is how a gripper closes.

Proposition 22 (Events). Let $v_r^\pm \in \mathcal{T}^\pm(q(t_k))$ be the required velocities before and after an event and $\Delta_r = v_r^- - v_r^+$.

- (a) Reset law. $s^+ = P^+(s^- + \Delta_r)$, hence $W^+ = \frac{1}{2}|P^+(s^- + \Delta_r)|_M^2$ and $|s^+|_M \leq |s^-|_M + |P^+ \Delta_r|_M$.
- (b) Reference reset with the plant. If $v_r^+ = P^+ v_r^-$, then $s^+ = P^+ s^-$ and $W^+ = W^- - \frac{1}{2}|(I - P^+)s^-|_M^2 \leq W^-$. In particular, if the reference is continuous and admissible in both modes ($v_r^+ = v_r^- \in \mathcal{T}^+$), then $W^+ = W^- - \frac{1}{2}|v^- - v^+|_M^2$: the error energy drops by the Carnot loss. At an opening with a continuous reference it is unchanged.
- (c) Event sequences. Let the events $t_1 < t_2 < \dots$ have no finite accumulation point, let $\lambda_{\min}(\bar{D} + \text{sym } \Psi, M) \geq \kappa_0 > 0$ in all modes visited, the residual between events being $E^\top \delta = \Psi \sigma + \eta$ as in Proposition 5. Then for $t \geq t_0$

$$|s(t)|_M \leq e^{-\kappa_0(t-t_0)}|s(t_0)|_M + \frac{1}{\kappa_0} \sup_{[t_0, t]} |\eta|_{\bar{M}^{-1}} + \sum_{t_0 < t_k \leq t} e^{-\kappa_0(t-t_k)} |P_k^+ \Delta_{r,k}|_M. \quad (25)$$

With an exact model and a reference that is continuous at every event, $|s(t)|_M \leq e^{-\kappa_0(t-t_0)}|s(t_0)|_M$ for every switching sequence: $\bar{W}$ is a common certificate and no dwell time is needed.

Proof. (a) $s^+ = P^+ v^- - v_r^+ = P^+(v^- - v_r^+) = P^+(s^- + v_r^- - v_r^+)$ because $P^+ v_r^+ = v_r^+$; an M -orthogonal projector does not increase the M -norm. (b) With $v_r^+ = P^+ v_r^-$, $\Delta_r = (I - P^+)v_r^-$ and $P^+ \Delta_r = 0$, so $s^+ = P^+ s^-$; Pythagoras. If $v_r^- \in \mathcal{T}^+$, then $(I - P^+)s^- = (I - P^+)v^- = v^- - v^+$. At an opening $P^+ s^- = s^-$. (c) Between events, $|s(t)|_M \leq e^{-\kappa_0(t-\tau)}|s(\tau)|_M + \int_\tau^t e^{-\kappa_0(t-r)} |\eta(r)|_{\bar{M}^{-1}} dr$ from the proof of Proposition 5, with τ the last event before t (or t_0) and $s(\tau)$ the value after it; at events, (a); induction over the events, and $\int_{t_0}^t e^{-\kappa_0(t-r)} dr \leq 1/\kappa_0$. $\square$




Fig. 3. Error storage of **team3** across a release ($t = 1$ s) and a plastic latch at the present relative pose ($t = 2$ s) on the compiled ideal-constraint model, with a reference continuous at the events and with a retreat of the released arm that starts with a reference jump. Circles: predictions of Proposition 22.

Part (b) is the rule for closings: reset the reference by the projection that resets the plant, in the spirit of reference spreading [27, 32]; then W cannot increase, and by Theorem 8(d) a closing cannot lower the certified rate when M and $\mathcal{D}$ are unchanged. The fitted reference is in general not continuous at events; (25) bounds what its jumps cost.

Sampled renewal. A controller that renews its references at the instants $t_k = t_0 + kh$ and advances them in between has an event of the same kind at every renewal, without a change of mode ($P^+ = I$ on $\mathcal{T}$): the fitted reference jumps by what the advance has missed, and the held part of the feedforward leaves a residual. If the jumps satisfy $|\Delta_{r,k}|_M \leq \bar{\Delta}$ and the residual satisfies (c) with $|\eta|_{\bar{M}-1} \leq \bar{\eta}$, then (25) and $\sum_k e^{-\kappa_0 kh} = 1/(1 - e^{-\kappa_0 h})$ give

$$\limsup_{t \rightarrow \infty} |s(t)|_M \leq \frac{\bar{\Delta}}{1 - e^{-\kappa_0 h}} + \frac{\bar{\eta}}{\kappa_0}. \quad (26)$$

If the reference advance and held feedforward satisfy $\bar{\Delta} \leq c_\Delta h^2$ and $\bar{\eta} \leq c_\eta h$ on the considered compact set of regular states, with constants independent of h , then (26) gives an ultimate error of $O(h)$ and ultimate storage of $O(h^2)$; Section 7.7 shows this scaling emerging at the longer intervals. The constants are not explicit, and stepped actuator feedback is not covered; sample-period design for fully held feedback has a separate emulation literature [25].

Fig. 3 illustrates Proposition 22 on **team3** under the law (5), simulated on the compiled ideal-constraint model (Runge–Kutta, 2 ms steps) with $K = 10 \text{ s}^{-1}$ times the inertia seen at each actuator port and the admissible reference $u_r = \dot{y}_d - 4(y - y_d)$ on the independent coordinates y . The last robot releases the payload at $t = 1$ s and latches again at its present relative pose at $t = 2$ s. With a continuous reference, W is unchanged at the release and drops by the Carnot loss at the latch ($4.56 \times 10^{-7} \rightarrow 4.28 \times 10^{-7}$ J). When the retreat of the released arm starts with a reference

jump of mass norm $0.40 \text{ J}^{1/2}$, W jumps at the release ($4.06 \times 10^{-5} \rightarrow 8.27 \times 10^{-2}$ J), as (a) states, and drops by the Carnot loss at the latch. In all four events, W^+ equals the prediction to a relative 10^{-12} .

6.4 Friction-driven bases

A skid-steered base turns by sliding its wheels, so its planar wheel interfaces are not rolling edges: only its vertical support is an ideal edge, and its wheels transmit planar forces through friction, which enters (2) as the generalized force

$$Q_\mu(q, v; \varrho) = - \sum_k \mu_k \varrho_k S_k(q)^\top \phi(S_k(q)v), \quad (27)$$

where S_k maps v to the planar slip velocity of wheel k and acts only on the coordinates of its robot, $\varrho_k \geq 0$ is the normal load and $\mu_k > 0$ the friction coefficient of the wheel, and $\phi = \nabla \psi$ for a convex C^2 dissipation potential ψ ; the smoothed Coulomb law has $\psi(x) = (|x|^2 + \epsilon^2)^{1/2}$ with a regularization speed $\epsilon > 0$. The loads are not determined by (2), since a base has more wheel loads than independent support rows (Lemma 2), and they depend on the efforts and on the friction itself; a load model fixes them, and it is assumed to have a measurable, locally bounded solution $\varrho(t)$ along the motion. With the actuators locked, such a team can still slide on the plane, so $\text{rank } L < d$ (Proposition 4(a)).

Proposition 23 (Friction-driven bases). *Let Assumption 1 hold, and consider a motion of (2) with $Q_\mu(q, v; \varrho)$ on its right side under (5), with (4) containing $Q_\mu(q, v_r; \varrho)$ at the same loads and v_r admissible.*

(a) $\dot{W} = -s^\top (\mathcal{D} + \mathcal{D}_\mu)s$, with the positive semidefinite secant damping

$$\mathcal{D}_\mu = \sum_k \mu_k \varrho_k S_k^\top \left(\int_0^1 \nabla^2 \psi(S_k(v_r + \vartheta s)) d\vartheta \right) S_k.$$

(b) $\dot{W} \leq -2\kappa_\mu W$ with $\kappa_\mu = \lambda_{\min}(E^\top (\mathcal{D} + \mathcal{D}_\mu) E, \bar{M})$, and $\kappa_\mu > 0$ if and only if $\ker A \cap \ker B^\top \cap \ker R \cap \ker \mathcal{D}_\mu = \{0\}$, which can hold when $\text{rank } L < d$.

(c) If (4) uses loads $\tilde{\varrho}$ instead of ϱ , it holds up to the residual $\sum_k \mu_k (\varrho_k - \tilde{\varrho}_k) S_k^\top \phi(S_k v_r)$ of Proposition 5, with $\mathcal{D} + \mathcal{D}_\mu$ in place of $\mathcal{D}$; its contribution to $\dot{W}$ has either sign.

Proof. (a) The friction adds $Q_\mu(q, v; \varrho) - Q_\mu(q, v_r; \varrho)$ to the right of (7), and $\phi(S_k v) - \phi(S_k v_r) = \int_0^1 \nabla^2 \psi(S_k(v_r + \vartheta s)) d\vartheta S_k s$ with $\nabla^2 \psi \succeq 0$; the rest is the proof of Proposition 3. (b) As Proposition 4(b); for positive semidefinite matrices, $\ker(\mathcal{D} + \mathcal{D}_\mu) = \ker B^\top \cap \ker R \cap \ker \mathcal{D}_\mu$. (c) The balance holds up to that term. $\square$

For the smoothed Coulomb law, $\nabla^2 \psi(x)$ has the eigenvalue $(|x|^2 + \epsilon^2)^{-1/2}$ across x and $\epsilon^2(|x|^2 + \epsilon^2)^{-3/2}$ along it, so the secant damping of a wheel is of order $\mu_k \varrho_k / \epsilon$ where it grips and, where it slides, of order $\mu_k \varrho_k / |S_k v|$ across the slip but only $\mu_k \varrho_k \epsilon^2 / |S_k v|^3$ along it if the reference slides with it: the rate is high while the wheels grip and low while they slide. It depends on s , v_r , and ϱ , so a decay bound along a motion needs a lower bound on $\mathcal{D}_\mu$ there. Since each S_k acts on one robot, $\mathcal{D}_\mu = \text{blkdiag}(0, \mathcal{D}_{\mu,1}, \dots, \mathcal{D}_{\mu,N})$, and Section 4 applies at each state with $\mathcal{D}_i + \mathcal{D}_{\mu,i}$ in place of $\mathcal{D}_i$. As rank $L < d$, (a3) of Proposition 4 fails: the projection of $E^\top$ times (4) on $\ker L$ contains no effort, so along $\ker L$ the reference must accelerate as the unactuated plant would at the reference velocity. A controller integrates this projection, with the loads solved together with the efforts; the fitted reference of Theorem 19 needs rank $L = d$.

7 Numerical evaluation

Sections 7.2–7.4 evaluate the certificates and the synthesis at recorded configurations and on synthetic examples, Sections 7.5–7.7 the loops in time, and Section 7.8 the friction-driven teams; the supplementary material is available in the GitHub repository linked in the title footnote.

7.1 Teams, missions, and controllers

The compiled models contain three mobile-manipulator families in two- and three-robot assemblies carrying a beam (Fig. 4). The RB-ROBOUT+/UR30 and Ridgeback/UR10e assemblies satisfy rank $L = d$. The Husky A300/UR10e assemblies have rank $L = d - 3$: three admissible motions remain with the actuators locked (Proposition 4(a)), and their skid-steered bases are driven through wheel friction (Section 6.4). The full-rank teams are `team2`, `team3`, `ridgeback2`, and `ridgeback3`, the friction-driven ones `husky2` and `husky3`. Sections 7.2–7.7 use the former, and only the event checks of Section 7.2 also cover the latter.

The complete missions extract a beam (5 and 5.5 m long for two and three robots) from a rack, carry it along a route with an aisle turn, a 1.7 m doorway, a crossing cart, and a 3° ramp, place it on two supports through guides with 6 mm clearance, and release it (Fig. 5). Grasp acquisition, lift-off, touchdown, and release change the active constraints. Two plants use the same physical models: a Pinocchio-based plant that enforces ideal grasp and rolling constraints by projecting its state onto them, and MuJoCo 3.3.7, with contact dynamics and approximate constraint enforcement. Unless stated otherwise, both use 2 ms steps, references and feedforward are renewed every 10 ms, and feedback is applied in every step.

The arm and lift rate gains are $k_{d,j} = 2\sqrt{k_{p,j}m_j}$ (Table 1), with m_j the joint-diagonal entry of the robot mass

Table 1

Mission parameters. Beam masses are for two/three robots.

	RB-ROBOUT+/UR30	Ridgeback/UR10e
Beam mass (kg)	43.2/59.4	11.7/14.6
Arm k_p (10^3 N m/rad)	14, 14, 10, 3, 3, 3	9, 9, 6, 1.5, 1.5, 1.5
Lift k_p (N/m)	2.5×10^5	–
Wheel gain (N ms/rad)	60	12
Base K_b (s^{-1}), K_I (s^{-2})	1.6, 0.4	3.0, 1.0
Grasp force/moment bounds	450 N/45 N m	85 N/10 N m

matrix at the nominal closed configuration, augmented by $(m_0/N)(0.5 \text{ m}^2)$ on the first three arm joints and by m_0/N on the lift (m_0 : beam mass); wrist wrenches enter the admittance through a 6 Hz filter. Four controllers are compared: D , the desired-motion feedforward of the supplied mission controller with k_d and unfiltered rates; F , the fitted loops of Theorem 19 with k_d ; $F/2$, the fitted loops with $k_d/2$; and D_R , the desired-motion feedforward with arm damping retuned for the closed-chain inertia and the supplied 15 Hz rate filter, which also gives the filtered D , F , and $F/2$ below. By Remark 20, D and $F/2$ have the single-joint damping ratio 1 and F has 1.25: D – $F/2$ compares matched single-joint damping, D – F equal gains.

The certificates are evaluated at the nominal carrying configuration at route station 2 m, with all grasps active and no payload support, and at 24 snapshots of the navigation phase of a recorded ideal-plant mission with F , anchored by the recorded poses and coordinates; grasp closure changes the arm coordinates by at most 1.46×10^{-4} rad.

7.2 Exact composition and refresh

At random states, Proposition 3 holds to a relative 9×10^{-15} and Theorem 19(a) to 2×10^{-13} on the full-rank teams, the Carnot rule of Lemma 21 and the reset law of Proposition 22 hold to 10^{-12} on all six teams, and `team2` with one contact row per wheel (28 rows of rank 26) satisfies Lemma 2 to 2×10^{-13} .

At the nominal configurations, every robot’s local rows admit every six-dimensional payload velocity, so Proposition 7 and Theorem 8(b) apply. For the mission gains, the payload-space test reproduces the centralized rates, 15.6171, 16.3750, 28.8806, and 29.1350 s^{-1} in team order. At $\kappa = \underline{\kappa}$, the supplies $\Theta_i = \Xi_i - S/N$ sum to $\underline{\kappa}M_0$ and satisfy every robot’s inequality, both to a relative 10^{-15} , and they are indefinite, as directional supplies may be.




Fig. 4. Compiled platform families. Tables give the assembled coordinate dimension n , admissible dimension d , actuator count p , and rank L for two- and three-robot carrying modes. The RB-ROBOUT+ and Ridgeback teams are fully actuated on the admissible space; Husky has rank deficit three.




Fig. 5. Recorded snapshots of one three-robot RB-ROBOUT+/UR30 MuJoCo mission with the fitted-reference controller: (a) extraction from the rack, (b) transport with three active grasps, and (c) delivery and load transfer to the mounting supports. Labels indicate approximate simulation times.




Fig. 6. Rates certified at the nominal carrying configurations by the payload-space test (equal to the centralized rate) and by equal and optimized scalar shares, (a) for gains designed robot by robot for a target of 10 s^{-1} with equal payload shares and (b) for the mission gains.

How much scalar shares lose depends on the gains (Fig. 6). For gains designed robot by robot—each robot sets K_i to 10 s^{-1} times the diagonal of the inertia seen at its own ports, with an equal share M_0/N of the payload inertia—optimized scalar shares lose a factor of 1.9–2.3. For the mission gains, designed joint by joint for design inertias, they lose a factor of 26–38. Matrix supplies lose nothing in either case (Remark 9).

Along the 24 snapshots, where the centralized rates range over $[15.6162, 15.9239]$, $[16.3726, 16.7445]$, $[28.4088, 28.8806]$, and $[28.5604, 29.1350] \text{ s}^{-1}$ in team

order, every robot refreshes its summary, and a target of 90% of each team’s rate at the first snapshot passes the payload-space test at every snapshot; reusing the previous snapshot’s summaries would change the aggregate by up to 16.2% in relative Frobenius norm, so moving cells must refresh. For membership, at 95% of the nominal team rate, removing robot 3 from **team3** by (15) lowers the largest certified rate to 12.0846 s^{-1} , below the 15.5562 s^{-1} target, whereas for **ridgeback3** it remains 28.3268 s^{-1} , above the 27.6782 s^{-1} target. These are frozen-configuration checks.

7.3 Endpoint and architecture examples

For the data (16), supplies bounded by $|\Theta_i| \leq 10, 100$, and 1000 certify at best 0.911, 0.990, and 0.999, approaching $\underline{\kappa} = 1$ only as the bound grows (Proposition 12). On 136 random join-tree assemblies with 3–7 cells, separator realizability, and module damping of random rank, the leaf-to-root construction of Theorem 11(b) at $\kappa = \underline{\kappa}$ gives cell forms that sum to $\mathcal{D} - \underline{\kappa}M$ and are nonnegative on their cells, both to a relative 4×10^{-13} . For the constructions of Theorem 13 with $\kappa_z = 2$ (chordless rings, $k = 4, 6$; uncovered cliques, $m = 3, 4, 5$), a semidefinite program over all modular certificates returns best rates below 10^{-8} in magnitude, whereas one cell containing the clique recovers 2; with only two-module edges on the cells of the $m = 4, 5$ cliques (all $(m - 1)$ -subsets), it returns $\underline{\kappa}$ to a relative 2×10^{-6} on 218 random assemblies, as Remark 15 states.

Table 2

Synthesis over the 24 snapshots. Normalized gain for P1 (mission: n_J) and certified rate for P2 (s^{-1}); the centralized and robot-level optima agree to within 2×10^{-7} relative.

Team	n_J	κ_{mis}	P1		P2	
			gain (cen. = rob.)	cen.	rob.	
team2	14	15.6162	7.9141	22.3761	22.3761	
team3	21	16.3726	11.8626	22.8662	22.8662	
ridgeback2	12	28.4088	7.1933	29.5849	29.5849	
ridgeback3	18	28.5604	10.4173	29.7455	29.7455	

7.4 Compositional synthesis

Theorem 16 is applied with $\mathcal{C}$ the 24 snapshots of each team. The decision variables are the arm and lift rate gains, within $[0.1, 2]$ times their mission values; the wheel gains stay at their mission values. Let κ_{mis} be the smallest rate certified by the mission gains over $\mathcal{C}$. Problem P1 minimizes the normalized gain $\sum_j k_{d,j}/k_{d,j}^{\text{mis}}$ subject to a certified rate of at least κ_{mis} at every snapshot; Problem P2 maximizes the certified rate at the mission's normalized budget, $\sum_j k_{d,j}/k_{d,j}^{\text{mis}} \leq n_J$, where n_J is the number of arm and lift joints. Both are solved centrally with $\mathcal{G}$ and in the robot-level form of Theorem 16(d). At every snapshot each robot realizes every payload velocity, so Theorem 16(b) applies.

Table 2 shows that the robot-level and centralized optima agree to the solver tolerance, as Theorem 16(c) states, and every returned gain set, rechecked by generalized eigenvalues at all snapshots, certifies at least its rate, up to a relative shortfall below 10^{-9} . P1 certifies the mission rate with 40–44% less normalized gain: the mission gains, chosen joint by joint for critical damping of a design inertia, are not distributed where the certificate needs them. P2 raises the certified rate by 43% and 40% for the RB-ROBOUT+ teams but by only 4% for the Ridgeback teams, whose rates are limited by the fixed wheel loops: with arm gains multiplied by 100 the **ridgeback2** rate rises only from 28.88 to 30.36 s^{-1} , whereas doubling its wheel gains gives 40.52 s^{-1} . A sharing ADMM [8] for P1 on **team2**, in which each robot solves only its own LMIs and exchanges 24×21 supply entries per iteration, comes within 0.24% of the optimum after 200 iterations with gains that certify $15.639 s^{-1} \geq \kappa_{\text{mis}}$.

7.5 Decay after a velocity jump

Fig. 7 applies the fitted loops of Theorem 19(b) with the mission, P1, and P2 gains on the ideal plant at 0.5 ms steps with renewal in every step, approximating the continuous-time loops. After settling at route station 2 m, each run receives an admissible velocity jump of mass norm 0.1 $J^{1/2}$ along the slowest certified direction




Fig. 7. Normalized error storage after an admissible velocity jump along the slowest certified direction, on the ideal-constraint plant (0.5 ms steps, renewal in every step), for the mission gains and the minimum-gain (P1) and maximum-rate (P2) designs of Table 2. Dashed: the certified bounds $e^{-2 \int \kappa}$ of Proposition 4(b).

of its gains. In all twelve runs, $W/W(0)$ stays at or below its certified bound $e^{-2 \int \kappa}$, evaluated with the measured configuration, until $W/W(0) = 10^{-3}$. The margin is an integration effect: at the end of the window, $W/W(0)$ lies 4–7% below the bound at 0.5 ms and 8–13% below it at 1 ms, and a first-order extrapolation to zero step places it within 0.2% of the bound (up to 0.19% above it, which two steps cannot resolve). The bound is thus nearly attained along this direction, consistent with Proposition 4(b). The P1 gains reach 10^{-3} within 0.003 s of the mission gains, and the P2 gains reach it in 0.1535 and 0.1500 s instead of 0.2200 and 0.2100 s for the RB-ROBOUT+ teams. With the 10 ms renewal of the missions, the bound is not guaranteed (Theorem 19(c)): at a 0.5 ms step, $W/W(0)$ exceeds it by up to 11.1% and 8.0% for the P1 gains of **team2** and **team3** and stays below it in the other ten runs; at a 2 ms step the coarser integration hides this excess by adding damping (fitted rates up to 3.1% higher).

Fig. 8 repeats the jump test with D , F , and $F/2$ at 2 ms steps and 10 ms renewal, with jumps of mass norm 0.1 and 0.6 $J^{1/2}$ on the ideal plant and on MuJoCo. The adverse direction minimizes the generalized eigenvalue of $\text{sym}(\bar{D} - \bar{M}G_p)$ relative to $\bar{M}$, the form of Theorem 19(d) for a desired-motion feedforward at rest; three random directions give further tests. This eigenvalue is -9.77 and $-8.69 s^{-1}$ for the RB-ROBOUT+ teams and 11.31 and 15.26 s^{-1} for the Ridgeback teams. Accordingly, with D the storage of the fitted reference first rises, by 10.6% and 8.0% on the RB-ROBOUT+ teams on the ideal plant and by 9.7% and 7.2% on MuJoCo, whereas on the Ridgeback teams it stays below $W(0)$. In all 32 runs of F and $F/2$ on each plant, W stays below $W(0)$. On the ideal plant, both stay within their own certified bands until $W/W(0) = 10^{-3}$. On MuJoCo, F stays within its band and $F/2$ exceeds its own bound by at most 0.49% of the bound before $W/W(0) = 0.05$; af-




Fig. 8. Normalized error storage of the fitted reference after an adverse velocity jump with the mission controllers D , F , and $F/2$ (2 ms steps, 10 ms renewal) on (a) the ideal-constraint plant and (b) MuJoCo. Shaded: the band certified for F ($F/2$ has other gains; D is not certified). Cutoffs: 10^{-3} (a), 0.05 (b).




Fig. 9. Error storage W of the fitted reference along the complete missions on the ideal-constraint plant (top) and MuJoCo (bottom), for the desired-motion feedforward with the 15 Hz rate filter and the fitted loops with k_d and $k_d/2$. Shaded: carrying phase.

terward contact compliance and approximate constraint enforcement leave a residual of 0.6–2.3% of $W(0)$, and the band is only a reference curve. In one MuJoCo run ($F/2$, **team2**, adverse direction), an effort is clipped at its bound in the first step after the jump.

7.6 Complete missions

All 40 missions— D with the 15 Hz filter, F , and $F/2$ on the ideal plant, and on MuJoCo D , F , and $F/2$ with and without that filter and D_R —complete in 246.53–265.98 s of simulated time, with placement errors below 0.25 mm. Along the missions (Fig. 9), the median storage over the reference updates on the ideal plant is 1.1×10^{-8} to 2.2×10^{-8} J with F , 2.7×10^{-8} to 1.1×10^{-7} J with $F/2$, and 0.9×10^{-6} to 3.1×10^{-6} J with the filtered D , with isolated peaks below 4.2×10^{-2} J during navigation and lowering. On MuJoCo the medians are 4×10^{-3}

to 1.7×10^{-2} J for all seven controllers; there the storage is dominated by contact and constraint-enforcement effects, as in Fig. 8(b), and serves only as a diagnostic.

Fig. 10 compares the controllers on MuJoCo. With matched single-joint damping, $F/2$ lowers the peak carried-payload position error of D by 37–52% and the rms error by 17–34%; with equal gains, F lowers them by 8.5–28% and 5.7–15%. D_R changes the peak of D by at most 0.043 mm. Peaks of single runs are sensitive, however: repeated with the same code on another computer, the $F/2$ missions change their peak errors by up to 27% (**team2**) and their rms errors by at most 2.3%, so the rms comparison is the more robust one.

To check the resolution of these 2 ms realizations, a carrying segment is repeated at plant steps of 2, 1, 0.5, and 0.25 ms on both plants: from rest at route station




Fig. 10. MuJoCo missions of the four full-rank teams: (a) peak and (b) rms carried-payload position error, (c) largest grasp moment and (d) largest actuator effort relative to their declared bounds; controllers as in Section 7.1. Grasp moments are maxima of 20 ms interval means of the wrist measurements after tool-inertia removal and cannot certify instantaneous limits.

Table 3
Carrying segment: payload position error over the 4 s of motion, peak/rms (mm), at a plant step of 0.25 ms.

Plant	Team	D	F	$F/2$
MuJoCo	team2	.470/.231	.341/.163	.191/.090
	team3	.474/.226	.338/.162	.191/.089
	ridgeback2	.521/.243	.404/.194	.232/.103
	ridgeback3	.513/.239	.400/.192	.231/.103
Ideal plant	team2	.154/.082	.084/.041	.053/.022
	team3	.151/.081	.083/.040	.051/.022
	ridgeback2	.093/.040	.048/.018	.029/.010
	ridgeback3	.091/.039	.047/.018	.028/.010

1 m, the team holds the beam for 2 s and accelerates along the route for 4 s to 0.45 m/s, with renewal every 10 ms and feedback in every step (Table 3: errors over the motion at 0.25 ms). On the ideal plant, the errors grow by 10–24% from 2 to 0.25 ms, with first-order convergence of the peaks (the last halving adds 2.4–3.2%), while their ratios to D change by at most 2.4% in peak and 11% in rms. On MuJoCo, the last halving changes the errors by at most 1.4%, but at 2 ms the peak errors of the Ridgeback teams are up to 27% too low. In every run, $F/2 < F < D$ in both peak and rms; at 0.25 ms, $F/2$ lowers the peak error of D by 66–69% on the ideal plant and 55–60% on MuJoCo, and F by 45–48% and 22–29%. The ordering of the missions thus persists under step refinement, whereas their percentages at 2 ms carry resolution errors of this size.

In the recorded runs, the grasp-moment maxima are

9.25, 11.92, 6.81, and 8.37 N m with $F/2$ in team order, against 11.06, 17.59, 8.79, and 10.60 N m with F and 11.29, 18.80, 7.91, and 13.00 N m with D (Fig. 10(c)); ridgeback3 with F and with D exceeds its 10 N m bound. These 20 ms interval means change by up to 21% in the repeated $F/2$ missions, where ridgeback3 also exceeds the bound (10.09 N m). Inferred at every 2 ms step (within 2×10^{-3} N m of the solver’s grasp wrench on the carrying segment), the moments of the repetitions peak at 17.7, 25.3, 31.1, and 32.4 N m during navigation: below the 45 N m bound of the RB-ROBOUT+ teams, but above the 10 N m bound of the Ridgeback teams in 0.8% and 1.0% of the latched grasp steps. Differences between controllers below about 20% in Fig. 10(c) are within the variation between repetitions.

7.7 Sampling, model errors, noise, and delay

Fig. 11 runs a carry stage on the ideal plant under conditions that the statements leave out: from rest at route station 1 m, the team speeds up to 0.45 m/s, and the metrics are taken over the cruise, $t > 3$ s. (a) As the renewal interval grows from 2 to 20 ms, the median storage grows from 3.6×10^{-10} and 4.3×10^{-10} J to 1.2×10^{-8} and 1.5×10^{-8} J (team2, ridgeback3), with a local slope of $\log W$ against $\log h$ that rises from 1.2 (2–4 ms) to 1.8 and 1.9 (10–20 ms), approaching the h^2 scaling of (26). (b) A beam model with 0.7–1.3 times the mass causes a steady sag of up to 1.6 mm (5.0 mm for team2), a residual in the sense of Proposition 5, while the varying part of the error stays at 2.7–8.9 μ m with the fitted loops and 13.5–16.6 μ m with the filtered D . (c) White noise of 10^{-3} and 5×10^{-3} (rad/s, m/s) on the measured rates raises the rms grasp moment from about 0.014 N m to 0.48 and 2.38 N m with F , 0.33 and 1.63 N m with $F/2$,




Fig. 11. Carry stage on the ideal-constraint plant under conditions outside the statements: (a) median storage over the cruise against the renewal interval h (**team2**, **ridgeback3**); for **ridgeback3**, (b) varying part of the payload error when the controller’s model of the beam has 0.7–1.3 times its mass, (c) rms grasp moment under white noise on the measured rates, and (d) median storage with measurement delays of 2 and 4 ms. Legends: “fitted, $K = k_d$ ” and “ $K = k_d/2$ ” are F and $F/2$, “desired motion” is the filtered D unless marked “no filter”; “filter”: the 15 Hz rate filter.




Fig. 12. Friction-driven teams. (a)–(c) MuJoCo mission of **husky3** with the supplied controller; the bases turn by sliding their wheels. (d) Error storage of **husky2** on the compiled model under (5) with the friction in the required balance, the reference integrated along $\ker L$, and the loads solved at every state, for three wheel conditions; dashed: $e^{-2\kappa_\mu t}$ with $\kappa_\mu = 1.455 \text{ s}^{-1}$ (Proposition 23(b), zero slip, standing loads), shifted for visibility.

and 0.12 and 0.58 N m with the filtered D ; filtering the damping rates of F halves its values. (d) Under a 4 ms measurement delay, both feedforwards with the 15 Hz filter fail on both teams (efforts at their bounds); without it, D , F , and $F/2$ keep the median below 1.2×10^{-8} J for **ridgeback3**, but for **team2** only $F/2$ keeps it small (2.2×10^{-8} J, against 0.22 J for F and 0.32 J for D). The fitted loops are thus more sensitive to rate noise, and the filter combined with delay limits both feedforwards.

7.8 Friction-driven teams

For the Husky teams, $\ker L$ is spanned, to 2×10^{-15} , by the planar rigid motions of the formation, whose rate is zero without friction. With the smoothed Coulomb law of the compiled model ($\epsilon = 1 \text{ mm/s}$, $\mu = 0.8$) and the gain rule of Section 6.3, Proposition 23(b) certifies 1.455 and 1.472 s^{-1} (**husky2**, **husky3**) at zero slip and standing loads, limited by the team rolling on gripping wheels, which the wheel gains damp (95% of the dissipation). At 20 random states per team with sliding wheels and loads of 0.3–1.7 times standing, it certifies 0.12–1.40 and 0.60–1.39 s^{-1} , Proposition 23(a) holds to a relative 5×10^{-15} , and the residual of (c) takes both signs.

Fig. 12(d) shows three 1 s runs of **husky2**, asked to drive straight at 0.15 m/s, that start with the bases sliding at 23 and 56 mm/s, yawing at 0.02 rad/s, and every wheel

0.1 rad/s off its rate: with $\mu = 0.25$; with $\mu = 0.25$ and the front wheels driven faster and the rear wheels slower by up to 40 mm/s (scrub); and with $\mu = 0.8$, where two wheel loads vanish at the start. W never increases; friction dissipates 68–70% of $W(0)$, the actuators and joint damping 27–31%. Late in the runs where the wheels grip, $|s|_M$ decays at 1.455 s^{-1} , the rate at zero slip; in the scrub run, the front wheels slide with the reference at 67–87 mm/s, and the decay slows to 0.79 s^{-1} .

The fitted loops are not available for these teams. With the uncertified supplied controller (the filtered D with the beams, arm gains, and grasp bounds of the Ridgeback missions, and integral wheel loops of gain 60 N ms/rad that track wheel rates from a skid model at the present loads), both teams complete the missions (Fig. 12(a)–(c)) on MuJoCo and on the ideal plant, where these bases stand on tire compliance, in 268.6–298.3 s, with peak payload errors of 0.70–0.80 mm (rms 0.064–0.079 mm), placement errors below 0.06 mm, and grasp-moment maxima (20 ms means) of 7.20–9.32 N m, within the 10 N m bound. Inferred at every 2 ms step on MuJoCo as in Section 7.6 (accurate to 5×10^{-4} N m), the moments peak at 11.8 and 14.2 N m where a rear axle meets the foot of the ramp and exceed the bound in 7 of 1.8×10^5 and 12 of 2.9×10^5 latched grasp steps. The bases turn by sliding their wheels (up to 95.3 mm/s, 2.03–2.91 kJ per base). On rigid ground instead of tires,

with the loads of the previous step, both teams exceed their grasp-load bounds at the foot of the ramp, and the **husky3** bases rock between wheel diagonals in a skid turn, where friction shifts load through the height of the center of mass with a loop gain of 0.83–0.85, amplifying load changes six- to sevenfold: the loads, a residual in Proposition 23(c), matter.

Isaac Sim. The supplied controller (the filtered D) also completes all six missions in Isaac Sim 6.1.0 (PhysX, compliant grasps), with peak carried-payload position errors of 0.88–1.71 mm and every declared bound met. For brevity, these runs are only summarized here; their code and records are in the repository.

8 Conclusion

For constrained robot teams, the velocity-error decay rate composes from local inequalities, for every assembly with constraints of arbitrary arity, exactly when the cells admit a join tree; with two-module constraints, more architectures are lossless. The centralized rate is attained when the assembly realizes every velocity of each separator, a hypothesis that cannot be dropped. Because the cell inequalities are linear in the actuator gains, minimum-gain and maximum-rate syntheses decompose into robot-level problems coupled through payload-sized supplies, without loss under the endpoint condition. The error-power identity and the fitted reference make the certified rates decay rates of joint-level loops; a reset law carries the bound across grasp and release events and over sampled renewal, with an $O(h)$ ultimate error under the stated uniform local error bounds; for skid-steered bases, friction with a convex dissipation potential supplies the damping that the actuators cannot. On the four full-rank teams, the local tests reproduce the centralized rates, robot-level synthesis certifies the mission rate with markedly less gain, and complete missions run with the fitted loops on two plants, although the Ridgeback grasp moments exceed their bound in single steps, and those of **ridgeback3** also in 20 ms means; two friction-driven teams complete them with an uncertified controller, exceeding the grasp-moment bound in a few single steps. Certificates for contact plants, instantaneous loads, and continuous configuration regions, certified loops for friction-driven teams, explicit sampling constants, and a decentralized evaluation of the required effort remain open.

Acknowledgements

This work was supported by the Korean Ministry of Trade, Industry and Energy (MOTIE) under Grant RS-2024-00442168, through the Mechanical Equipment Industry Technology Development Program administered by KEIT.

References

- [1] F. Aghili. A unified approach for inverse and direct dynamics of constrained multibody systems based on linear projection operator: Applications to control and simulation. *IEEE Trans. Robot.*, 21(5):834–849, 2005.
- [2] J. Agler, J. W. Helton, S. McCullough, and L. Rodman. Positive semidefinite matrices with a given sparsity pattern. *Linear Algebra Appl.*, 107:101–149, 1988.
- [3] A. Albert. Conditions for positive and nonnegative definiteness in terms of pseudoinverses. *SIAM J. Appl. Math.*, 17(2):434–440, 1969.
- [4] M. S. Andersen, S. K. Pakazad, A. Hansson, and A. Rantzer. Robust stability analysis of sparsely interconnected uncertain systems. *IEEE Trans. Autom. Control*, 59(8):2151–2156, 2014.
- [5] M. Arcak, C. Meissen, and A. Packard. *Networks of Dissipative Systems: Compositional Certification of Stability, Performance, and Safety*. SpringerBriefs in Electrical and Computer Engineering. Springer, 2016.
- [6] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *J. ACM*, 30(3):479–513, 1983.
- [7] E. G. Boman, D. Chen, O. Parekh, and S. Toledo. On factor width and symmetric H-matrices. *Linear Algebra Appl.*, 405:239–248, 2005.
- [8] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2011.
- [9] J. Cervera, A. J. van der Schaft, and A. Baños. Interconnection of port-Hamiltonian systems and composition of Dirac structures. *Automatica*, 43(2):212–225, 2007.
- [10] J. Cortés, M. de León, D. Martín de Diego, and S. Martínez. Mechanical systems subjected to generalized non-holonomic constraints. *Proc. R. Soc. Lond. A*, 457(2007):651–670, 2001.
- [11] P. Culbertson, J.-J. Slotine, and M. Schwager. Decentralized adaptive control for collaborative manipulation of rigid bodies. *IEEE Trans. Robot.*, 37(6):1906–1920, 2021.
- [12] R. Danesh, F. Janabi-Sharifi, and F. Aghili. Real-time projected adaptive control for closed-chain co-manipulative continuum robots. arXiv:2604.04286v1, 2026.
- [13] S. Erhart and S. Hirche. Internal force analysis and load distribution for cooperative multi-robot manipulation. *IEEE Trans. Robot.*, 31(5):1238–1243, 2015.
- [14] S. Erhart and S. Hirche. Model and analysis of the interaction dynamics in cooperative manipulation tasks. *IEEE Trans. Robot.*, 32(3):672–683, 2016.
- [15] M. Hejrati, A. H. Barjini, G. Alcan, and J. Mattila. Adaptive modular geometric control of robotic ma-

- nipulators. arXiv:2603.03965v3, 2026.
- [16] T. Hufnagel and A. Müller. A projection method for the elimination of contradicting decentralized control forces in redundantly actuated PKM. *IEEE Trans. Robot.*, 28(3):723–728, 2012.
 - [17] J.-P. Humaloja, J. Koivumäki, L. Paunonen, and J. Mattila. Decentralized observer design for virtual decomposition control. *IEEE Trans. Autom. Control*, 67(5):2529–2536, 2022.
 - [18] O. Khatib. Inertial properties in robotic manipulation: An object-level framework. *Int. J. Robot. Res.*, 14(1):19–36, 1995.
 - [19] J. Koivumäki, J.-P. Humaloja, L. Paunonen, W.-H. Zhu, and J. Mattila. Subsystem-based control with modularity for strict-feedback form nonlinear systems. *IEEE Trans. Autom. Control*, 68(7):4336–4343, 2023.
 - [20] D. Lee and P. Y. Li. Passive decomposition of mechanical systems with coordination requirement. *IEEE Trans. Autom. Control*, 58(1):230–235, 2013.
 - [21] N. H. McClamroch and D. Wang. Feedback stabilization and tracking of constrained robots. *IEEE Trans. Autom. Control*, 33(5):419–426, 1988.
 - [22] C. Meissen, L. Lessard, M. Arcak, and A. Packard. Compositional performance certification of interconnected systems using ADMM. arXiv:1410.2006, 2014.
 - [23] P. J. Moylan and D. J. Hill. Stability criteria for large-scale systems. *IEEE Trans. Autom. Control*, 23(2):143–149, 1978.
 - [24] A. Müller. Consequences of geometric imperfections for the control of redundantly actuated parallel manipulators. *IEEE Trans. Robot.*, 26(1):21–31, 2010.
 - [25] D. Nešić, A. R. Teel, and D. Carnevale. Explicit computation of the sampling period in emulation of controllers for nonlinear sampled-data systems. *IEEE Trans. Autom. Control*, 54(3):619–624, 2009.
 - [26] L. Righetti, J. Buchli, M. Mistry, M. Kalakrishnan, and S. Schaal. Optimal distribution of contact forces with inverse-dynamics control. *Int. J. Robot. Res.*, 32(3):280–298, 2013.
 - [27] A. Saccon, N. van de Wouw, and H. Nijmeijer. Sensitivity analysis of hybrid systems with state jumps with application to trajectory tracking. In *Proc. 53rd IEEE Conf. Decis. Control*, pages 3065–3070, 2014.
 - [28] M. H. Shahna, J. Kim, and J. Mattila. RoboCompiler: Graph-native compilation of closed-chain robots for consistent modeling, control, and simulation. arXiv:2609.35717v2, 2026.
 - [29] J.-J. E. Slotine and W. Li. On the adaptive control of robot manipulators. *Int. J. Robot. Res.*, 6(3):49–59, 1987.
 - [30] H. G. Tanner, S. G. Loizou, and K. J. Kyriakopoulos. Nonholonomic navigation and control of cooperating mobile manipulators. *IEEE Trans. Robot. Autom.*, 19(1):53–64, 2003.
 - [31] F. E. Udwardia and R. E. Kalaba. A new perspective on constrained motion. *Proc. R. Soc. Lond. A*, 439:407–410, 1992.
 - [32] J. J. van Steen, N. van de Wouw, and A. Saccon. Robot control for simultaneous impact tasks via quadratic programming-based reference spreading. In *Proc. Amer. Control Conf.*, pages 3865–3872, 2022.
 - [33] L. Vandenberghe and M. S. Andersen. Chordal graphs and semidefinite optimization. *Found. Trends Optim.*, 1(4):241–433, 2015.
 - [34] C. K. Verginis, M. Mastellaro, and D. V. Dimarogonas. Robust cooperative manipulation without force/torque measurements: Control design and experiments. *IEEE Trans. Control Syst. Technol.*, 28(3):713–729, 2020.
 - [35] J. T. Wen and K. Kreutz-Delgado. Motion and force control of multiple robotic manipulators. *Automatica*, 28(4):729–743, 1992.
 - [36] J. C. Willems. Dissipative dynamical systems part I: General theory. *Arch. Ration. Mech. Anal.*, 45(5):321–351, 1972.
 - [37] M. Wojtyra. Joint reaction forces in multibody systems with redundant constraints. *Multibody Syst. Dyn.*, 14:23–46, 2005.
 - [38] Y. Zheng, G. Fantuzzi, and A. Papachristodoulou. Chordal and factor-width decompositions for scalable semidefinite and polynomial optimization. *Annu. Rev. Control*, 52:243–279, 2021.
 - [39] Y. Zheng, R. P. Mason, and A. Papachristodoulou. Scalable design of structured controllers using chordal decomposition. *IEEE Trans. Autom. Control*, 63(3):752–767, 2018.
 - [40] W.-H. Zhu. *Virtual Decomposition Control: Toward Hyper Degrees of Freedom Robots*, volume 60 of *Springer Tracts in Advanced Robotics*. Springer, Berlin, Heidelberg, 2010.
 - [41] W.-H. Zhu, Z. Bien, and J. De Schutter. Adaptive motion/force control of multiple manipulators with joint flexibility based on virtual decomposition. *IEEE Trans. Autom. Control*, 43(1):46–60, 1998.
 - [42] W.-H. Zhu, Y.-G. Xi, Z.-J. Zhang, Z. Bien, and J. De Schutter. Virtual decomposition based control for generalized high dimensional robotic systems with complicated structure. *IEEE Trans. Robot. Autom.*, 13(3):411–436, 1997.



Mehdi Heydari Shahna received the D.Sc. (Tech.) degree with distinction in Automation Science and Engineering from Tampere University, Finland, in 2025. He is currently a postdoctoral researcher at Tampere University. His research interests include physical AI, robot learning, safe control, and robust control.



Soyi Jung (Senior Member, IEEE) received the B.S., M.S., and Ph.D. degrees in electrical and computer engineering from Ajou University, Suwon, South Korea, in 2013, 2015, and 2021, respectively. She has been an Associate Professor with the Department of Electrical and Computer Engineering, Ajou University, since March 2025.



Soohyun Park (Member, IEEE) received the B.S. degree in computer science and engineering from Chung-Ang University, Seoul, South Korea, in February 2019, and the Ph.D. degree in electrical and computer engineering from Korea University, Seoul, in August 2023. She has been an Assistant Professor with the Division of Computer Science, Sookmyung Women's University, Seoul, since March 2025.



Joongheon Kim (Senior Member, IEEE) received the B.S. and M.S. degrees in computer science and engineering from Korea University, Seoul, South Korea, in 2004 and 2006, respectively, and the Ph.D. degree in computer science from the University of Southern California, Los Angeles, CA, USA, in 2014. He has been a Professor with the Department of Electrical and Computer Engineering, Korea University, since September 2019.



Jouni Mattila received a Ph.D. in automation engineering from Tampere University of Technology, Tampere, Finland, in 2000. He is currently a professor of machine automation with the Unit of Automation Technology and Mechanical Engineering at Tampere University. His research interests include machine automation, nonlinear-model-based control of robotic manipulators, and energy-efficient control of heavy-duty mobile manipulators.