

Multi-Objective Policy Generation for Multi-Robot Systems Using Riemannian Motion Policies

Anqi Li, Mustafa Mukadam, Magnus Egerstedt, and Byron Boots

Georgia Institute of Technology, Atlanta, Georgia 30332, USA
{anqi.li, mhmukadam, magnus}@gatech.edu, bboots@cc.gatech.edu

Abstract. In many applications, multi-robot systems are required to achieve multiple objectives. For these multi-objective tasks, it is often-times hard to design a single control policy that fulfills all the objectives simultaneously. In this paper, we focus on multi-objective tasks that can be decomposed into a set of simple subtasks. Controllers for these subtasks are individually-designed and then combined into a control policy for the entire team. One significant feature of our work is that the subtask controllers are designed along with their underlying manifolds. When a controller is combined with other controllers, their associated manifolds are also taken into account. This formulation yields a policy generation framework for multi-robot systems that can combine controllers for a variety of objectives while implicitly handling the interaction among robots and subtasks. To describe controllers on manifolds, we adopt Riemannian Motion Policies (RMPs), and propose a collection of RMPs for common multi-robot subtasks. Centralized and decentralized algorithms are designed to combine these RMPs into a final control policy. Theoretical analysis shows that the system under the control policy is stable. Moreover, we prove that many existing multi-robot controllers can be closely approximated by the framework. The proposed algorithms are validated through both simulated tasks and robotic implementations.

Keywords: Multi-Robot Systems, Motion Planning and Control

1 Introduction

Multi-robot control policies are often designed through performing gradient descent on a potential function that encodes a *single* team-level objective, e.g. forming a certain shape, covering an area of interest, or meeting at a common location [1–3]. However, many problems involve a diverse set of objectives that the robotic team needs to fulfill simultaneously. For example, collision avoidance and connectivity maintenance are often required in addition to any primary tasks [4]. One possible solution is to encode the multi-objective problem as a single motion planning problem with various constraints [5–8]. However, as more objectives and robots are considered, it can be difficult to directly search for a solution that can achieve *all* of the objectives simultaneously.

An alternative strategy is to design a controller for each individual objective and then *combine* these controllers into a single control policy. Different schemes

for combining controllers have been investigated in the multi-robot systems literature. For example, one standard treatment for inter-robot collision avoidance is to let the collision avoidance controller take over the operation if there is a risk of collision [9]. A fundamental challenge for such construction is that unexpected interaction between individual controllers can yield the overall system unstable [10, 11]. Null-space-based behavioral control [12, 13] forces low priority controllers to not interfere with high priority controllers. However, when there are a large number of objectives, the system may not have the sufficient degrees of freedom to consider all the objectives simultaneously. Another example is the potential field method, which formulates the overall controller as a weighted sum of controllers for each objective [9, 14]. While the system is guaranteed to be stable when the weights are constant, careful tuning of these constant weights are required to produce desirable behaviors.

Methods based on Control Lyapunov Functions (CLFs) and Control Barrier Functions (CBFs) [4, 10, 11] seek to optimize primary task-level objectives while formulating secondary objectives, such as collision avoidance and connectivity maintenance, as CLF or CBF constraints and solve via quadratic programming (QP). While this provides a computational framework for general multi-objective multi-robot tasks, solving the QP often requires centralized computation and can be computationally demanding if the number of robots or constraints is large. Although the decentralized safety barrier certificate [11] is a notable exception, it only considers inter-robot collision avoidance and it has not been demonstrated how the same decentralized construction can be applicable to other objectives.

In this paper, we return to the idea of combining controllers and rethink how an objective and its corresponding controller are defined: instead of defining objectives directly on the configuration space, we define them on non-Euclidean *manifolds*, which can be lower-dimensional than the configuration space. When combining individually-designed controllers, we consider the outputs of the controllers and their underlying manifolds. In particular, we adopt Riemannian Motion Policies (RMPs) [15], a class of manifold-oriented control policies that has been successfully applied to robot manipulators, and RMPflow [16], the computational framework for combining RMPs. This framework, where each controller is associated with a matrix-value and state-dependent weight, can be considered as an extension to the potential field method. This extension leads to new geometric insight on designing controllers and more freedom to combine them. While the RMPflow algorithm is centralized, we provide a decentralized version and establish the stability analysis for the decentralized framework.

There are several major advantages to defining objectives and controllers on manifolds for multi-robot systems. First, this formulation provides a general formula for the construction of controllers: the key step is to design the manifold for each subtask, as controllers/desired behaviors can be viewed as a natural outcome of their associated manifolds. For example, obstacle avoidance behavior is closely related to the geodesic flow in a manifold where the obstacles manifest as holes in the space. Second, since we design controllers in the manifolds most relevant to their objectives, these manifolds are usually of lower dimension than

the configuration space. When properly combined with other controllers, this can provide additional degrees of freedom that help controllers avoid unnecessary conflicts. This is particularly important for multi-robot systems where a large number of controllers interact with one another in a complicated way. Third, it is shown in [16] that Riemannian metrics on manifolds naturally provide a notion of importance that enables the stable combination of controllers. Finally, RMPflow is coordinate-free [16], which allows the proposed framework to be directly generalized to heterogeneous multi-robot teams.

We present four contributions in this paper. First, we present a centralized solution to combine controllers for solving multi-robot tasks based on RMPflow. Second, we design a collection of RMPs for simple and common multi-robot subtasks that can be combined to achieve more complicated tasks. Third, we draw a connection between some of the proposed RMPs and a large group of existing multi-robot distributed controllers. Finally, we introduce a decentralized extension to RMPflow, with its application to multi-robot systems, and establish the stability analysis for this decentralized framework.

2 Riemannian Motion Policies (RMPs)

We briefly review Riemannian Motion Policies (RMPs) [15], a mathematical representation of policies on manifolds, and RMPflow [16], a recursive algorithm to combine RMPs. In this section, we start with RMPs and RMPflow for single robots, for which these concepts are initially defined [15, 16]. We will later on consider them in the context of multi-robot systems in subsequent sections.

Consider a robot (or a group of robots in later sections) with its configuration space \mathcal{C} being a smooth d -dimensional manifold. For the sake of simplicity, we assume that \mathcal{C} admits a global¹ *generalized coordinate* $\mathbf{q} : \mathcal{C} \rightarrow \mathbb{R}^d$. As is the case in [16], we assume that the system can be feedback linearized in such a way that it is controlled directly through the generalized acceleration, $\ddot{\mathbf{q}} = \mathbf{u}(\mathbf{q}, \dot{\mathbf{q}})$. We call \mathbf{u} a *policy* or a *controller*, and $(\mathbf{q}, \dot{\mathbf{q}})$ the *state*.

RMPflow [16] assumes that the task is composed of a set of *subtasks*, for example, avoiding collision with an obstacle, reaching a goal, tracking a trajectory, etc. In this case, the task space, denoted \mathcal{T} , becomes a collection of multiple *subtask spaces*, each corresponding to a subtask. We assume that the task space \mathcal{T} is related to the configuration space \mathcal{C} through a smooth *task map* $\psi : \mathcal{C} \rightarrow \mathcal{T}$. The goal of RMPflow [16] is to generate policy \mathbf{u} in the configuration space \mathcal{C} so that the trajectory exhibits desired behaviors on the task space \mathcal{T} .

2.1 Riemannian Motion Policies

Riemannian Motion Policies (RMPs) [15] represent policies on manifolds. Consider an m -dimensional manifold \mathcal{M} with generalized coordinate $\mathbf{x} \in \mathbb{R}^m$. An RMP on \mathcal{M} can be represented by two forms, its *canonical form* and its *natural form*. The *canonical form* of an RMP is a pair $(\mathbf{a}, \mathbf{M})^{\mathcal{M}}$, where $\mathbf{a} : (\mathbf{x}, \dot{\mathbf{x}}) \mapsto$

¹ In the case when \mathcal{C} does not admit a global coordinate, a similar construction can be done locally on a subset of the configuration space \mathcal{C} .

$\mathbf{a}(\mathbf{x}, \dot{\mathbf{x}}) \in \mathbb{R}^m$ is the desired acceleration, i.e. control input, and $\mathbf{M} : (\mathbf{x}, \dot{\mathbf{x}}) \mapsto \mathbf{M}(\mathbf{x}, \dot{\mathbf{x}}) \in \mathbb{R}_+^{m \times m}$ is the inertial matrix which defines the importance of the RMP when combined with other RMPs. Given its canonical form, the *natural form* of an RMP is the pair $[\mathbf{f}, \mathbf{M}]^{\mathcal{M}}$, where $\mathbf{f} = \mathbf{M}\mathbf{a}$ is the desired force. The natural forms of RMPs are introduced mainly for computational convenience.

RMPs on a manifold \mathcal{M} can be naturally (but not necessarily) generated from a class of systems called *Geometric Dynamical Systems* (GDSs) [16]. GDSs are a generalization of the widely studied classical *Simple Mechanical Systems* (SMSs) [17]. In GDSs, the kinetic energy metric, \mathbf{G} , is a function of both the configuration and velocity, i.e. $\mathbf{G}(\mathbf{x}, \dot{\mathbf{x}}) \in \mathbb{R}_+^{m \times m}$. This allows kinetic energy to be dependent on the direction of motion, which can be useful in applications such as obstacle avoidance [15, 16]. The dynamics of GDSs are in the form of

$$(\mathbf{G}(\mathbf{x}, \dot{\mathbf{x}}) + \boldsymbol{\Xi}_{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}})) \ddot{\mathbf{x}} + \boldsymbol{\xi}_{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}}) = -\nabla_{\mathbf{x}}\Phi(\mathbf{x}) - \mathbf{B}(\mathbf{x}, \dot{\mathbf{x}})\dot{\mathbf{x}}, \quad (1)$$

where we call $\mathbf{B} : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}_+^{m \times m}$ the *damping matrix* and $\Phi : \mathbb{R}^m \rightarrow \mathbb{R}$ the *potential function*. The curvature terms $\boldsymbol{\Xi}_{\mathbf{G}}$ and $\boldsymbol{\xi}_{\mathbf{G}}$ are induced by metric \mathbf{G} ,

$$\begin{aligned} \boldsymbol{\Xi}_{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}}) &:= \frac{1}{2} \sum_{i=1}^m \dot{x}_i \partial_{\dot{\mathbf{x}}} \mathbf{g}_i(\mathbf{x}, \dot{\mathbf{x}}), \\ \boldsymbol{\xi}_{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}}) &:= \dot{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}})\dot{\mathbf{x}} - \frac{1}{2} \nabla_{\mathbf{x}}(\dot{\mathbf{x}}^\top \mathbf{G}(\mathbf{x}, \dot{\mathbf{x}})\dot{\mathbf{x}}), \end{aligned} \quad (2)$$

with $\dot{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}}) := [\partial_{\dot{\mathbf{x}}} \mathbf{g}_i(\mathbf{x}, \dot{\mathbf{x}})\dot{\mathbf{x}}]_{i=1}^m$, \mathbf{g}_i denoting the i th column of \mathbf{G} , x_i denoting the i th component of \mathbf{x} , and $[\cdot]$ denoting matrix composition through horizontal concatenation of vectors. Given a GDS (1), there is an RMP $(\mathbf{a}, \mathbf{M})^{\mathcal{M}}$ naturally associated with it given by $\mathbf{a} = \ddot{\mathbf{x}}$ and $\mathbf{M}(\mathbf{x}, \dot{\mathbf{x}}) = (\mathbf{G}(\mathbf{x}, \dot{\mathbf{x}}) + \boldsymbol{\Xi}_{\mathbf{G}}(\mathbf{x}, \dot{\mathbf{x}}))$. Therefore, the velocity dependent metric \mathbf{G} provides velocity dependent importance weight \mathbf{M} when combined with other RMPs.

2.2 RMPflow

RMPflow [16] is an algorithm to generate control policies on the configuration space given the RMPs for all subtasks, for example, collision avoidance with a particular obstacle, reaching a goal, etc. Given the state information of the robot in the configuration space and a set of individually-designed controllers (RMPs) for the subtasks, RMPflow produces the control input on the configuration space through combining these controllers.

RMPflow introduces: i) a data structure, the *RMP-tree*, to describe the structure of the task map ψ , and ii) a set of operators, the *RMP-algebra*, to propagate information across the RMP-tree. An RMP-tree is a directed tree. Each node \mathbf{u} in the RMP-tree is associated with a state $(\mathbf{x}, \dot{\mathbf{x}})$ defined over a manifold \mathcal{M} together with an RMP $(\mathbf{f}_{\mathbf{u}}, \mathbf{M}_{\mathbf{u}})^{\mathcal{M}}$. Each edge \mathbf{e} in the RMP-tree is augmented with a smooth map from the parent node to the child node, denoted as $\psi_{\mathbf{e}}$. An example RMP-tree is shown in Fig. 1. The root node of the RMP-tree \mathbf{r} is associated with the state of the robot $(\mathbf{q}, \dot{\mathbf{q}})$ and its control policy on the configuration space $(\mathbf{f}_{\mathbf{r}}, \mathbf{M}_{\mathbf{r}})^{\mathcal{C}}$. Each leaf node \mathbf{l}_k corresponds to a subtask with its control policy given by an RMP $(\mathbf{f}_{\mathbf{l}_k}, \mathbf{M}_{\mathbf{l}_k})^{\mathcal{T}_k}$, where \mathcal{T}_k is a subtask space.

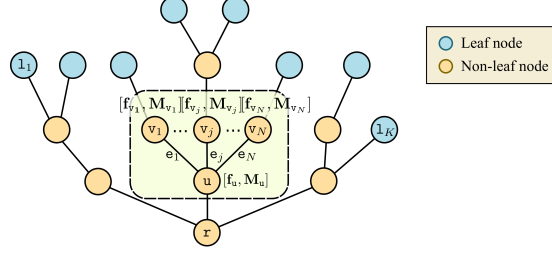


Fig. 1: An example of an RMP-tree. See text for details.

The RMP-algebra consists of three operators: **pushforward**, **pullback** and **resolve**. To illustrate how they operate, consider a node u with N child nodes, denoted as $\{v_j\}_{j=1}^N$. Let $\{e_j\}_{j=1}^N$ be the edges from u to the child nodes (Fig. 1). Suppose that u is associated with the manifold \mathcal{M} , while each child node v_j is associated with the manifold \mathcal{N}_j . The RMP-algebra works as follows:

1. The **pushforward** operator forward propagates the *state* from the parent node u to its child nodes $\{v_j\}_{j=1}^N$. Given the state $(\mathbf{x}, \dot{\mathbf{x}})$ associated with u , the **pushforward** operator at v_j computes its associated state as $(\mathbf{y}_j, \dot{\mathbf{y}}_j) = (\psi_{e_j}(\mathbf{x}), \mathbf{J}_{e_j}(\mathbf{x}) \dot{\mathbf{x}})$, where $\mathbf{J}_{e_j} = \partial_{\mathbf{x}} \psi_{e_j}$ is the Jacobian matrix of the map ψ_{e_j} .
2. The **pullback** operator combines the RMPs from the child nodes $\{v_j\}_{j=1}^N$ to obtain the RMP associated with the parent node u . Given the RMPs from the child nodes, $\{[\mathbf{f}_{v_j}, \mathbf{M}_{v_j}]^{\mathcal{N}_j}\}_{j=1}^N$, the RMP associated with node u , $[\mathbf{f}_u, \mathbf{M}_u]^{\mathcal{M}}$, is computed by the **pullback** operator as,

$$\mathbf{f}_u = \sum_{j=1}^N \mathbf{J}_{e_j}^\top (\mathbf{f}_{v_j} - \mathbf{M}_{v_j} \mathbf{J}_{e_j} \dot{\mathbf{x}}), \quad \mathbf{M}_u = \sum_{j=1}^N \mathbf{J}_{e_j}^\top \mathbf{M}_{v_j} \mathbf{J}_{e_j}.$$

3. The **resolve** operator maps a natural-formed RMP to its canonical form. Given the natural-formed RMP $[\mathbf{f}_u, \mathbf{M}_u]^{\mathcal{M}}$, the operator produces $(\mathbf{a}_u, \mathbf{M}_u)^{\mathcal{M}}$ with $\mathbf{a}_u = \mathbf{M}_u^\dagger \mathbf{f}_u$, where \dagger denotes Moore-Penrose inverse.

With the RMP-tree specified, RMPflow can perform control policy generation through the following process. First, RMPflow performs a forward pass: it recursively calls **pushforward** from the root node to the leaf nodes to update the state information associated with each node in the RMP-tree. Second, every leaf node l_k *evaluates* its corresponding natural-formed RMP $\{[\mathbf{f}_{l_k}, \mathbf{M}_{l_k}]^{\mathcal{T}_k}\}$, possibly given by a GDS. Next, RMPflow performs a backward pass: it recursively calls **pullback** from the leaf nodes to the root node to back propagate the RMPs in the natural form. After that, RMPflow calls **resolve** at the root node to transform the RMP $[\mathbf{f}_r, \mathbf{M}_r]^{\mathcal{C}}$ into its canonical form $(\mathbf{a}_r, \mathbf{M}_r)^{\mathcal{C}}$. Finally, the robot executes the control policy by setting $\ddot{\mathbf{q}} = \mathbf{u} = \mathbf{a}_r$.

2.3 Stability Properties of RMPflow

To establish the stability results of RMPflow, we assume that every leaf node is associated with a GDS. Before stating the stability theorem, we need to define the metric, damping matrix, and potential function for a *node* in the RMP-tree.

Definition 1. If a node is a leaf, its metric, damping matrix and potential function are defined as in its associated GDS (1). Otherwise, let $\{\mathbf{v}_j\}_{j=1}^N$ and $\{\mathbf{e}_j\}_{j=1}^N$ denote the set of all child nodes of \mathbf{u} and associated edges, respectively. Suppose that $G_{\mathbf{v}_j}$, $B_{\mathbf{v}_j}$ and $\Phi_{\mathbf{v}_j}$ are the metric, damping matrix, and potential function for the child node \mathbf{v}_j . Then, the metric $\mathbf{G}_{\mathbf{u}}$, damping matrix $\mathbf{B}_{\mathbf{u}}$ and potential function $\Phi_{\mathbf{u}}$ for the node \mathbf{u} are defined as,

$$\mathbf{G}_{\mathbf{u}} = \sum_{j=1}^N \mathbf{J}_{\mathbf{e}_j}^\top \mathbf{G}_{\mathbf{v}_j} \mathbf{J}_{\mathbf{e}_j}, \quad \mathbf{B}_{\mathbf{u}} = \sum_{j=1}^N \mathbf{J}_{\mathbf{e}_j}^\top \mathbf{B}_{\mathbf{v}_j} \mathbf{J}_{\mathbf{e}_j}, \quad \Phi_{\mathbf{u}} = \sum_{j=1}^N \Phi_{\mathbf{v}_j} \circ \psi_{\mathbf{e}_j}, \quad (3)$$

where \circ denotes function composition.

The stability results of RMPflow are stated in the following theorem.

Theorem 1 (Cheng et al. [16]). Let $\mathbf{G}_{\mathbf{r}}$, $\mathbf{B}_{\mathbf{r}}$, and $\Phi_{\mathbf{r}}$ be the metric, damping matrix, and potential function of the root node defined in (3). If $\mathbf{G}_{\mathbf{r}}, \mathbf{B}_{\mathbf{r}} \succ 0$, and $\mathbf{M}_{\mathbf{r}} = (\mathbf{G}_{\mathbf{r}} + \boldsymbol{\Xi}_{\mathbf{G}_{\mathbf{r}}})$ is nonsingular, the system converges to a forward invariant set $\mathcal{C}_{\infty} := \{(\mathbf{q}, \dot{\mathbf{q}}) : \nabla_{\mathbf{q}} \Phi_{\mathbf{r}} = 0, \dot{\mathbf{q}} = 0\}$.

3 Centralized Control Policy Generation

We begin by formulating a control policy generation algorithm for multi-robot systems directly based on RMPflow. This algorithm is *centralized* because it requires a centralized processor to collect the states of *all* robots and solve for the control input for *all* robots jointly given *all* the subtasks. In Section 4, we introduce a decentralized algorithm and analyze its stability properties.

Consider a potentially heterogeneous² team of N robots indexed by $\mathcal{I} = \{1, \dots, N\}$. Let \mathcal{C}_i be the configuration space of robot i with \mathbf{q}_i being a generalized coordinate on \mathcal{C}_i . The configuration space is then the product manifold $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_N$. As in Section 2, we assume that each robot is feedback linearized and we model the control policy for each robot as a second-order differential equation $\ddot{\mathbf{q}}_i = \mathbf{u}_i(\mathbf{q}_i, \dot{\mathbf{q}}_i)$. An obvious example is a team of mobile robots with double integrator dynamics on \mathbb{R}^2 . Note, however, that the approaches proposed in this paper is not restricted to mobile robots in Euclidean spaces.

Let $\mathcal{K} = \{1, \dots, K\}$ denote the index set of all subtasks. For each subtask $k \in \mathcal{K}$, a controller is individually designed to generate RMPs on the subtask manifold \mathcal{T}_k . Here we assume that the subtasks are *pre-allocated* in the sense that each subtask $\mathbf{1}_k$ is defined for a specified subset of robots \mathcal{I}_k . Examples of subtasks include collision avoidance between a pair of robots (a binary subtask), trajectory following for a robot (a unitary subtask), etc.

The above formulation gives us an alternative view of multi-robot systems with emphasis on their multi-task nature. Rather than encoding the team-level task as a global potential function (as is commonly done in the multi-robot literature), we decompose the task as *local* subtasks defined for subsets of robots, and design policies for individual subtasks. The main advantage is that as the task

² Please see Appendix D for a discussion about heterogeneous teams.

becomes more complex, it becomes increasingly difficult to design a single potential function that renders the desired global behavior. However, it is often natural to decompose global tasks into local subtasks, even for complex tasks, since multi-robot tasks can often come from local specifications [3, 4]. Therefore, this formulation provides a straightforward generalization to multi-objective tasks. Moreover, this subtask formulation allows us to borrow existing controllers designed for single-robot tasks, such as collision avoidance, goal reaching, etc.

Recall from Section 2 that RMPflow operates on an RMP-tree, a tree structure describing the task space. The main objective of this section is thus to construct an RMP-tree for general multi-robot problems. Note that given a set of subtasks, the construction of the RMP-tree is not unique. One way to construct an RMP-tree is to use non-leaf nodes to represent subset of the team:

- The root node corresponds to the joint configuration space $\mathcal{C} = \mathcal{C}_1 \times \dots \times \mathcal{C}_N$ and its corresponding control policy.
- Any leaf node 1_k is augmented with a user-specified policy represented as an RMP on the subtask manifold \mathcal{T}_k .
- Every non-leaf node is associated with a product space of the configuration spaces for a subset of the team.
- The parent of any leaf RMP 1_k is associated with the joint configuration space $\prod_{i \in \mathcal{I}_k} \mathcal{C}_i$, where \mathcal{I}_k are the robots that subtask 1_k is defined on.
- Consider two non-leaf nodes u and v such that v is a decedent of u in the RMP-tree. Let \mathcal{I}_u and \mathcal{I}_v be the subset of robots corresponds to node u and v , respectively. Then $\mathcal{I}_v \subseteq \mathcal{I}_u$.

Fig. 2 shows an example RMP-tree for a team of three robots. The robots are tasked with forming a certain shape and reaching a goal while avoiding inter-robot collisions. The root of the RMP-tree is associated with the configuration space for the team, which is the product of the configuration spaces for all three robots. On the second level, the nodes represent subsets of robots which, in this case, are pairs of robots. Several leaf nodes, such as the ones corresponding to collision avoidance and distance preservation, are children of these nodes as they are defined on pairs of robots. One level deeper is the node corresponding to the configuration space for robot 1. The goal attractor leaf node is a child of it since the goal reaching subtask is assigned only to robot 1.

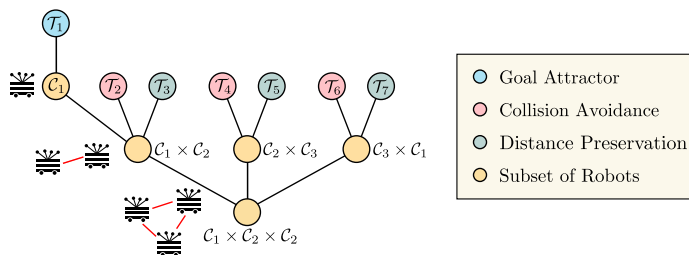


Fig. 2: An example of an RMP-tree for a group of three robots performing a formation preservation task. See text for details.

Note that branching in the RMP-tree does not necessarily define a partition over robots. Let \mathbf{v}_i and \mathbf{v}_j be children of the same node \mathbf{u} and let $\mathcal{I}_{\mathbf{v}_i}$ and $\mathcal{I}_{\mathbf{v}_j}$ be the subset of robots for node \mathbf{v}_i and \mathbf{v}_j , respectively. Then it is *not* necessary that $\mathcal{I}_{\mathbf{v}_i} \cap \mathcal{I}_{\mathbf{v}_j} = \emptyset$. For example, in Fig. 2, the three nodes on the second level are defined for subsets $\{1, 2\}$, $\{2, 3\}$, and $\{3, 1\}$, respectively. The intersection of any two of them is not empty. In fact, if a branching is indeed a partition, then the problem can be split into *independent* sub-problems. For multi-robot systems, this means that the team consists of independent sub-teams with completely independent tasks. This rarely occurs in practice.

According to Theorem 1, if all the leaf node controllers are designed through GDSs, it is guaranteed that the controller generated by RMPflow drives the system to a forward invariant set $\mathcal{C}_\infty := \{(\mathbf{q}, \dot{\mathbf{q}}) : \nabla_{\mathbf{q}} \Phi_{\mathbf{r}} = 0, \dot{\mathbf{q}} = 0\}$ if $\mathbf{G}_{\mathbf{r}}, \mathbf{B}_{\mathbf{r}} \succ 0$ and $\mathbf{M}_{\mathbf{r}}$ is non-singular. In other words, this guarantees that the resulting system is stable, which is important: unstable behaviors such as high-frequency oscillation are avoided and, more importantly, stability provides formal guarantees on the performance of certain types of subtasks such as collision avoidance, which is discussed in Section 3.1.

To elucidate the process of designing RMPs and to connect to relevant multi-robot tasks, we provide examples of RMPs for multi-robot systems that can produce complex behaviors when combined. In the following examples, we use \mathbf{x}_i to denote the coordinate of robot i in \mathbb{R}^2 . An additional map can be composed with the given task maps if robots possess different kinematic structures.

3.1 Pairwise Collision Avoidance

To ensure safety operation of the robotic team, inter-robot collisions should be avoided. We formulate collision avoidance as ensuring a minimum safety distance d_S for *every* pair of robots. To generate collision-free motions, for any two robots $i, j \in \mathcal{I}$, we construct a collision avoidance leaf node for the pair. The subtask space is the 1-d distance space, i.e. $z = \psi(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|/d_S - 1$. Here, we use z (italic) to denote that it is a scalar on the 1-d space.

To ensure a safety distance between the pair, we use a construction similar to the collision avoidance RMP for static obstacles in [16]. The metric for the pairwise collision avoidance RMP is defined as $\mathbf{G}(z, \dot{z}) = w(z)u(\dot{z})$, where $w(z) = \frac{1}{z^4}$, $u(\dot{z}) = \epsilon + \min(0, \dot{z})\dot{z}$ with a small positive scalar $\epsilon > 0$. The metric retains a large value when the robots are close to each other (z is small), and when the robots are moving fast towards each other ($\dot{z} < 0$ and $|\dot{z}|$ is large). Conversely, the metric decreases rapidly as z increases. Recall that the metric is closely related to the inertial matrix, which determines the importance of the RMP when combined with other policies. This means that the collision avoidance RMP dominates when robots are close to each other or moving fast towards each other, while it has almost no effect when the robots are far from each other.

We next design the GDS that generates the collision avoidance RMP. The potential function is defined as $\Phi(z) = \frac{1}{2}\alpha w(z)^2$ and the damping matrix is defined as $\mathbf{B}(z, \dot{z}) = \eta \mathbf{G}(z, \dot{z})$, where α, η are positive scalars. As the robots approach the safety distance, the potential function $\Phi(z)$ approaches infinity.

Due to the stability guarantee of RMPflow, this barrier-type potential will always ensure that the distance between robots is greater than d_S . This means that the resulting control policy from RMPflow is *always* collision-free.

3.2 Pairwise Distance Preservation

Another common task for multi-robot systems is to form a specified shape or formation. This can be accomplished by maintaining the inter-robot distances between certain pairs of robots. Therefore, formation control can be induced by a set of leaf nodes that maintain distances. Such an RMP can be defined on the 1-d distance space, $z = \psi(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\| - d_{ij}$, where d_{ij} is the desired distance between robot i and robot j . For the GDS, we use a constant metric $\mathbf{G} \equiv c \in \mathbb{R}_{++}$. The potential function is defined as $\Phi(z) = \frac{1}{2} \alpha z^2$ and the damping is $\mathbf{B}(\mathbf{x}, \dot{\mathbf{x}}) \equiv \eta$, with $\alpha, \eta > 0$. We will refer to this RMP as *Distance Preservation RMPa* in later sections.

Note that the above RMP is not equivalent to the potential-based formation controller in, e.g. [2, 3]. However, there does exist an RMP that has very similar behavior. It is defined on the product space, $\mathbf{z} = (\mathbf{x}_i, \mathbf{x}_j)$. The metric for the RMP is also constant, $\mathbf{G} \equiv c\mathbf{I}$, where $c \in \mathbb{R}_{++}$ and \mathbf{I} denotes the identity matrix. The potential function is defined as $\Phi(\mathbf{z}) = \frac{1}{2} \mathcal{E}_{ij}(\|\mathbf{x}_i - \mathbf{x}_j\|)$, where $\mathcal{E}_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable and achieves its minimum at d_{ij} . Common choices include $\mathcal{E}_{ij}(s) = (s - d_{ij})^2$ and $\mathcal{E}_{ij}(s) = (s^2 - d_{ij}^2)^2$ [2]. The damping matrix is defined as $\mathbf{B} \equiv \eta\mathbf{I}$, with $\eta > 0$. This RMP will be referred to as *Distance Preservation RMPb* in later sections.

When there are only distance preserving RMPs in the RMP-tree, the resulting individual-level dynamics are given by

$$\ddot{\mathbf{x}}_i = -\frac{\alpha}{cD_i} \sum_{j:(i,j) \in E} \nabla_{\mathbf{x}_i} \mathcal{E}_{ij}(\|\mathbf{x}_i - \mathbf{x}_j\|) - \frac{\eta}{c} \dot{\mathbf{x}}_i, \quad (4)$$

where E represents the set of edges in the formation graph, and $D_i = |\{j : (i, j) \in E\}|$ is the degree of robot i . This is closely related to the gradient descent update rule over the potential function $\mathcal{E}(x) = \frac{1}{2} \sum_{(i,j) \in E} \mathcal{E}_{ij}(\|\mathbf{x}_i - \mathbf{x}_j\|)$ with an additional damping term, and normalized by the degree of the robot. We will later prove in Section 3.3 that the degree-normalized potential-based controller and the original potential-based controller have similar behaviors in the sense that the resulting systems converge to the same invariant set.

The main difference between the two distance preserving RMPs is the space on which they are defined. The first RMP is defined on a 1-d distance space while the second RMP is defined on a higher dimensional space. Therefore, the first RMP is more permissive in the sense that it only specifies desired behaviors in a one dimensional submanifold of the configuration space. This is illustrated through a simulated formation preservation task in Section 5.

3.3 Potential-based Controllers from RMPs

Designing controllers based on the gradient descent rule of a potential function is very common in the multi-robot systems literature, e.g. [2, 3, 18]. Usually, the

overall potential function \mathcal{E} is the sum of a set of symmetric, pairwise potential functions $\mathcal{E}_{ij}(\|\mathbf{x}_i - \mathbf{x}_j\|)$ between robot i and robot j that are adjacent in an underlying graph structure. When the robots follow double-integrator dynamics, a damping term is typically introduced to guarantee convergence to an invariant set. Let \mathbf{x} be the ensemble-level state of the team. The controller is given by, $\ddot{\mathbf{x}} = \mathbf{u} = -\nabla\mathcal{E} - \eta\dot{\mathbf{x}}$, where η is a positive scalar. We define a *degree-normalized* potential-based controller as, $\mathbf{u} = -\mathbf{\Gamma}(\nabla\mathcal{E} + \eta\dot{\mathbf{x}})$, where $\mathbf{\Gamma}$ is a diagonal matrix with $\mathbf{\Gamma}_{ii} = 1/D_i$ and D_i is the degree of robot i in the graph.

Theorem 2. *Both the degree-normalized controller and the original potential-based controller converge to the invariant set $\{(\mathbf{x}, \dot{\mathbf{x}}) : \nabla\mathcal{E} = 0, \dot{\mathbf{x}} = 0\}$.*

Proof. For the original controller, consider the Lyapunov function candidate $V(\mathbf{x}, \dot{\mathbf{x}}) = \frac{1}{2}\|\dot{\mathbf{x}}\|^2 + \mathcal{E}(\mathbf{x})$. Then $\dot{V} = \dot{\mathbf{x}}^\top(\ddot{\mathbf{x}} + \nabla\mathcal{E}) = -\eta\|\dot{\mathbf{x}}\|^2$. By LaSalle's invariance principle [19], the system converges to the set $\{(\mathbf{x}, \dot{\mathbf{x}}) : \nabla\mathcal{E} = 0, \dot{\mathbf{x}} = 0\}$. For the degree-normalized controller, consider the Lyapunov function candidate $V(\mathbf{x}, \dot{\mathbf{x}}) = \frac{1}{2}\dot{\mathbf{x}}^\top\mathbf{\Gamma}^{-1}\dot{\mathbf{x}} + \mathcal{E}(\mathbf{x})$. Then $\dot{V} = \dot{\mathbf{x}}^\top(\mathbf{\Gamma}^{-1}\ddot{\mathbf{x}} + \nabla\mathcal{E}) = -\eta\|\dot{\mathbf{x}}\|^2$. The system also converges to the same set by LaSalle's invariance principle [19]. ■

Therefore, similar to potential-based formation control, one can directly implement the degree-normalized version of these potential-based controllers by RMPs defined on the product space, $\mathbf{z} = (\mathbf{x}_i, \mathbf{x}_j)$. The potential function for the RMP is defined as $\Phi(\mathbf{z}) = \mathcal{E}_{ij}(\|\mathbf{x}_i - \mathbf{x}_j\|)$. Constant metric and damping can be used, e.g. $\mathbf{G} \equiv c\mathbf{I}$, and $\mathbf{B} \equiv \eta\mathbf{I}$, where c and η are positive scalars. Moreover, similar to formation control, one can also define RMPs on the distance space $z = \psi(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|$ with potential function $\Phi = \mathcal{E}_{ij}$. Since RMPs are defined on a lower-dimensional manifold, this approach may provide additional degrees of freedom when these RMPs are combined with other policies.

4 Decentralized Control Policy Generation

Although the centralized RMPflow algorithm can be used to generate control policies for multi-robot systems, it can be demanding in both communication and computation. Therefore, we develop a decentralized approximation of RMPflow that only relies on local communication and computation.

Before discussing the algorithm, a few definitions and assumptions are needed. Given the set of all subtasks \mathcal{K} , we say two robots i and j are neighbors if and only if there exists a subtask such that both robots are involved in. We then say that the algorithm is decentralized if only the *state* information of the robot's direct neighbors is required to solve for its control input. Note that here we implicitly assume that the robots are equipped with the sensing modality or communication modality to access the *state* of the neighbors. We also assume that the map and the Jacobian matrix for a subtask are known to the robot if the robot is involved in the subtask. For example, for the formation control task, the robot should know how to calculate distance between two robots given their states, and also know the partial derivatives of the distance function.

The major difference between the decentralized algorithm and the centralized RMPflow algorithm is that, in the decentralized algorithm, there is no longer

a centralized root node that can generate control policies for *all* robots. Instead, each robot should have its *own* RMP-tree that generates policies based on the information available locally. Therefore, the decentralized algorithm actually operates on a *forest* with N RMP-trees, called the RMP-forest. An example RMP-forest is shown in Fig. 3. There are three robots performing the same formation preservation task as in Fig. 2. Hence, there are three RMP-trees in the RMP-forest. For each RMP-tree, there are leaf nodes for every subtask relevant to the robot associated with the RMP-tree. As a result, there are multiple copies of certain subtasks in the RMP-forest, for example, the collision avoidance node for robot 1 and 2 appears twice: once in the RMP-tree of robot 1, and once in the RMP-tree of robot 2. However, these copies do not share information.

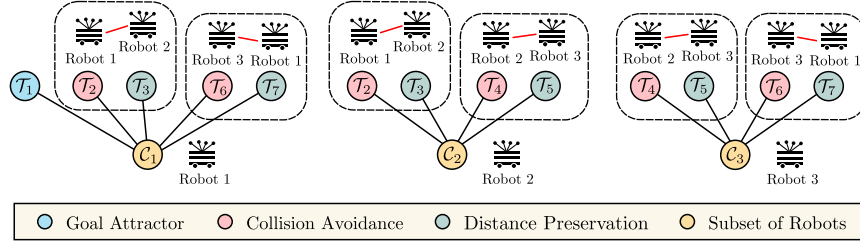


Fig. 3: A decentralized RMP-forest. The three robots are tasked with the same formation preservation task as in Fig. 2. For the decentralized algorithm, each robot has an individual RMP-tree to solve for its control input. All the leaf RMP nodes that are relevant to the robot are in its RMP-tree. For example, for robot 1, collision avoidance and distance preservation RMPs for both the pair $\{1, 2\}$ and $\{3, 1\}$ are introduced. There is a goal attractor for robot 1 since it is the leader. Note that there are several copies of the same subtasks in the forest, however, these copies do not share information with each other.

We call the decentralized approximation *partial RMPflow*. In partial RMPflow, every subtask is viewed as a time-varying unitary task. Therefore, following the RMP-tree construction in the previous section, it is natural to consider one-level RMP-trees, where the leaf nodes are direct children of the root nodes.

Notationwise, let \mathcal{K}_i be the set of subtasks that robot i participates in. Since there are multiple copies of the same subtasks in the RMP-forest, we use 1_k^i to denote the node corresponds to the copy of subtask k in the tree of robot i while let e_k^i denote the edge from the root of tree i to the leaf node 1_k^i . We let ψ_{e_k} denote the smooth map from the joint configuration space $\prod_{j \in \mathcal{I}_k} \mathcal{C}_j$ to the subtask space \mathcal{T}_k (which is the same across trees in the RMP-forest) and let $\mathbf{J}_{e_k}^i$ be the Jacobian matrix of ψ_{e_k} with respect to \mathbf{q}_i only, i.e. $\mathbf{J}_{e_k}^i = \partial_{\mathbf{q}_i} \psi_{e_k}$.

To compute the control input for robot i , an algorithm similar to RMPflow is applied in RMP-tree i :

- **pushforward**: Let $\{\mathbf{z}_k^i\}_{k \in \mathcal{K}_i}$ be the coordinates of the leaf nodes of RMP-tree i . Given the state of the root $(\mathbf{q}_i, \dot{\mathbf{q}}_i)$, its state is computed as, $\mathbf{z}_k^i = \psi_{e_k}(\mathbf{q}_{\mathcal{I}_k})$, $\dot{\mathbf{z}}_k^i = \mathbf{J}_{e_k}^i \dot{\mathbf{q}}_i$, where $\mathbf{q}_{\mathcal{I}_k} = \{\mathbf{q}_j : j \in \mathcal{I}_k\}$. It is worth noting that $\dot{\mathbf{z}}_k^i \neq \frac{d}{dt} \mathbf{z}_k^i$, since the other robots are considered static when computing $\dot{\mathbf{z}}_k^i$.

- **Evaluate:** Let \mathbf{M}_{1_k} , \mathbf{B}_{1_k} , and Φ_{1_k} be the user-designed metric, damping matrix, and potential function for subtask 1_k . For notational simplicity, we denote $\mathbf{G}_{1_k}^i = \mathbf{G}_{1_k}(\mathbf{z}_{1_k}^i, \dot{\mathbf{z}}_{1_k}^i)$, $\mathbf{B}_{1_k}^i = \mathbf{B}_{1_k}(\mathbf{z}_{1_k}^i, \dot{\mathbf{z}}_{1_k}^i)$, and $\Phi_{1_k}^i = \Phi_{1_k}(\mathbf{z}_{1_k}^i)$. At leaf node 1_k^i , the RMP is given by the following system (similar to GDS),

$$\mathbf{f}_{1_k}^i = -\nabla_{\mathbf{z}_{1_k}^i} \Phi_{1_k}^i - \mathbf{B}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i - \frac{1}{2} \mathbf{G}_{1_k}^i(\mathbf{z}_{1_k}^i, \dot{\mathbf{z}}_{1_k}^i) \dot{\mathbf{z}}_{1_k}^i, \quad \mathbf{M}_{1_k}^i = \mathbf{M}_{1_k}(\mathbf{y}_{1_k}^i, \dot{\mathbf{y}}_{1_k}^i) \quad (5)$$

Note that, to provide stability, the RMP is no longer generated by a GDS. In particular, the curvature term compensates for the motion of other robots.

- **pullback:** Given the RMPs from the leaf nodes of tree i , $\{[\mathbf{f}_{1_k}^i, \mathbf{M}_{1_k}^i]^{\top}\}_{k \in \mathcal{K}_i}$, the **pullback** operator calculates the RMP for the root node of tree i ,

$$\mathbf{f}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^{\top} (\mathbf{f}_{1_k}^i - \mathbf{M}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i), \quad \mathbf{M}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^{\top} \mathbf{M}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i. \quad (6)$$

- **resolve:** The control input is given by $\mathbf{u}_i = \mathbf{a}_r^i = (\mathbf{M}_r^i)^{\dagger} \mathbf{f}_r^i$.

Note that when all the metrics are constant diagonal matrices and all the Jacobian matrices are identity matrices, the decentralized partial RMPflow framework has exactly the same behavior as RMPflow. This, in particular, holds for the degree-normalized potential-based controllers discussed in Section 3.3. Therefore, the decentralized partial RMPflow framework can also reconstruct a large number of multi-robot controllers up to degree normalization.

Partial RMPflow has a stability result similar to RMPflow, which is stated in the following theorem.

Theorem 3. Let $\mathbf{G}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^{\top} \mathbf{G}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i$, $\mathbf{B}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^{\top} \mathbf{B}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i$, and $\Phi_r^i = \sum_{k \in \mathcal{K}_i} \Phi_{1_k}^i \circ \psi_{\mathbf{e}_k}$ be the metric, damping matrix, and potential of the tree i 's root node. If $\mathbf{G}_r^i, \mathbf{B}_r^i \succ 0$ and \mathbf{M}_r^i is nonsingular for all $i \in \mathcal{I}$, the system converges to a forward invariant set $\mathcal{C}_{\infty} := \{(\mathbf{q}, \dot{\mathbf{q}}) : \nabla_{\mathbf{q}_i} \Phi_r^i = 0, \dot{\mathbf{q}}_i = 0, \forall i \in \mathcal{I}\}$.

Proof. See Appendix A. ■

5 Experimental Results

We evaluate the multi-robot RMP framework through both simulation and robotic implementation. The detailed choice parameters in the experiments and additional simulation results can be found in Appendix B and Appendix C.

5.1 Simulation Results

Formation preservation tasks [20], where robots must maintain a certain formation while the leader is driven by some external force, are considered harder than formation control tasks since one needs to carefully balance the external force and the formation controller. However, since translations and rotations can still preserve shape, the team *should* have the capability of maintaining the formation regardless of the motion of the leader.

We consider a formation preservation task in simulation where a team of five robots are tasked with forming a regular pentagon while the leader has the additional task of reaching a goal. The two distance preservation RMPs

introduced in Section 3.2 are compared. Distance preservation RMPs are defined for all edges in the formation graph. To move the formation, an additional goal attractor RMP is defined for the leader robot, where the construction of the goal attractor RMP can be found in Appendix B.1 (referred to as Goal Attractor RMPa) or [16]. We use a damper RMP defined by a GDS on the configuration space of every single robot with only damping so that the robots can reach a full stop at the goal. Fig. 4a shows the resulting behavior for the distance preserving RMPa. The robots manage to preserve shape while the leader robot is reaching the goal since the subtasks are defined on lower-dimensional manifolds. By contrast, the behavior for distance preservation RMPb (which is equivalent to the degree-normalized potential-based controller) is shown in Fig. 4b. This distance preservation RMP fails to maintain the formation when the leader robot is attracted to the goal.

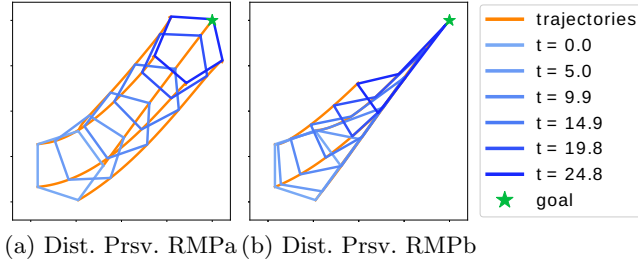


Fig. 4: (a) The behavior of distance preservation RMPa when combined with a goal attractor RMP. The blue pentagons from light to dark denote the shape from $t = 0.0s$ to $t = 24.8s$. The orange curves show the trajectories of the robots. The robots manage to reach the goal while maintaining the shape. (b) The same task with distance preservation RMPb. The robots fail to maintain the shape.

5.2 Robotic Implementations

We present several experiments (video: <https://youtu.be/VZHR5SN9wXk>) conducted on the Robotarium [21], a remotely accessible swarm robotics platform. Since the centralized RMPflow framework and the decentralized partial RMPflow frameworks have their own features, we design a separate experiment for each framework to show their full capability.

Centralized RMPflow Framework The main advantage of the centralized RMPflow framework is that the subtask spaces are jointly considered and hence the behavior of each controller is combined consistently. To fully exploit this feature, we consider formation preservation with two sub-teams of robots. The two sub-teams are tasked with maintaining their formation while moving back and forth between two goal points A and B . The five robots in the first sub-team are assigned a regular pentagon formation and the four robots in the second sub-team must form a square. At the beginning of the task, goal A is assigned to the first sub-team and goal B to the second sub-team. The robots negotiate their path so that their trajectories are collision free.

A combination of distance preservation RMPs, collision avoidance RMPs, goal attractor RMPs, and damper RMPs are used to achieve this behavior. The construction of the RMP-tree is similar to Fig. 2. A distance preservation RMPa is assigned to every pair of robots that corresponds to an edge in the formation graph, while collision avoidance RMPs are defined for every pair of robots. For each sub-team, we define a goal attractor RMP for the leader, where the construction of the goal attractor RMP is explained in Appendix B.1. We also use a damper RMP defined by a GDS on the configuration space of every single robot so that the robots can reach a full stop at the goal. Fig. 5 shows several snapshots from the experiment. We see that the robots are able to maintain their corresponding formations while avoiding collision. The two sub-teams of robots rotates around each other to avoid potential collision, which shows that the full degrees of freedom of the task is exploited.

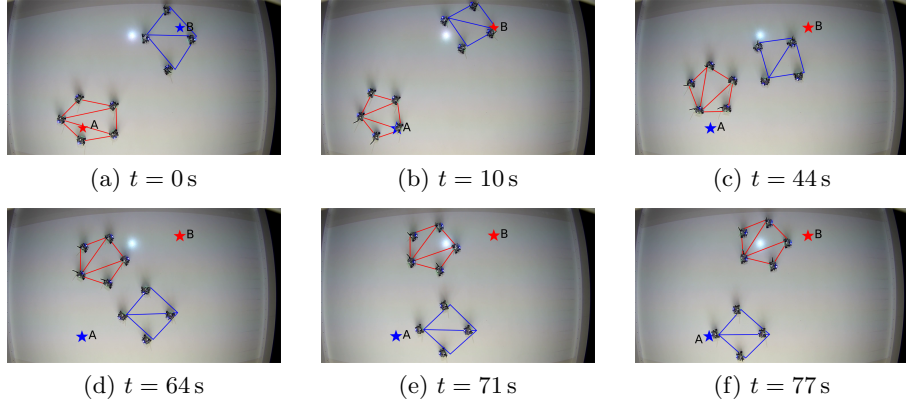


Fig. 5: The snapshots from the formation preservation experiment with the centralized RMPflow framework. Goal positions and the formation graphs are projected onto the arena by an overhead projector. The colors of the graphics are augmented in the figures for the purpose of visualization. The two sub-teams of robots are tasked with maintaining the formation while moving back and forth between two goal points in arena. The red and blue lines in the figure denote the formation graphs. The red and blue stars are the current goal positions for sub-team 1 and sub-team 2, respectively.

Decentralized Partial RMPflow Framework For the decentralized partial RMPflow framework, we consider a team of eight robots. The robots are divided into two sub-teams. The task of the first sub-team is to achieve cyclic pursuit behavior for a circle of radius 1 m centered at the origin. The other sub-team is designed to go through the circle surveilled by the other sub-team. To achieve the cyclic pursuit behavior, each robot in the first sub-team follows a point moving along the circle through a goal attractor RMP (defined in Appendix B.1). The RMP-forest for the second sub-team follows a similar structure as Fig. 3. For each single robot, there are collision avoidance RMPs for all other robots. Snapshots from the experiment are shown in Fig. 6. The robots from the second sub-team manage to pass through the circle under the decentralized framework.

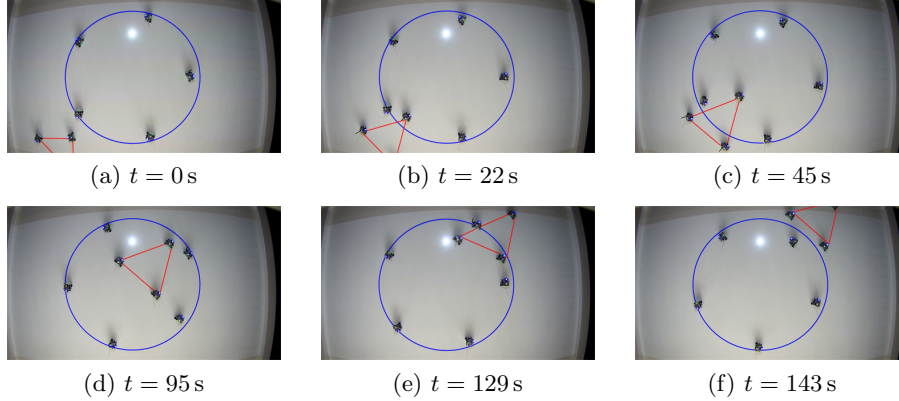


Fig. 6: The snapshots from the experiment for the decentralized multi-robot RMP framework. The circle surveilled by the first sub-team and the formation graph for the second sub-team are projected onto the environment. Robots were divided into two sub-teams. The first sub-team of five robots performed a cyclic pursuit behavior for a circle of radius 1 centered at the origin. The other sub-team passes through the circle surveilled by the other sub-team.

6 Conclusions

In this paper, we consider multi-objective tasks for multi-robot systems. We argue that it is advantageous to define controllers for single subtasks on their corresponding manifolds. We propose centralized and decentralized algorithms to generate control policies for multi-robot systems by combining control policies defined for individual subtasks. The multi-robot system is proved to be stable under the generated control policies. We show that many existing potential-based multi-robot controllers can also be approximated by the proposed algorithms. Several subtask policies are proposed for multi-robot systems. The proposed algorithms are tested through simulation and deployment on real robots.

Acknowledgments This work was supported in part by the grant ARL DCIST CRA W911NF-17-2-0181.

References

1. Francesco Bullo, Jorge Cortes, and Sonia Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*, volume 27. Princeton University Press, 2009.
2. Mehran Mesbahi and Magnus Egerstedt. *Graph theoretic methods in multiagent networks*, volume 33. Princeton University Press, 2010.
3. Jorge Cortés and Magnus Egerstedt. Coordinated control of multi-robot systems: A survey. *SICE Journal of Control, Measurement, and System Integration*, 10(6):495–503, 2017.
4. Li Wang, Aaron D Ames, and Magnus Egerstedt. Multi-objective compositions for collision-free connectivity maintenance in teams of mobile robots. In *IEEE 55th Conference on Decision and Control*, pages 2659–2664. IEEE, 2016.

5. Vishnu R Desaraju and Jonathan P How. Decentralized path planning for multi-agent teams with complex constraints. *Autonomous Robots*, 32(4):385–403, 2012.
6. Glenn Wagner and Howie Choset. Subdimensional expansion for multirobot path planning. *Artificial Intelligence*, 219:1–24, 2015.
7. Wenhao Luo, Nilanjan Chakraborty, and Katia Sycara. Distributed dynamic priority assignment and motion planning for multiple mobile robots with kinodynamic constraints. In *American Control Conference*, pages 148–154. IEEE, 2016.
8. Siddharth Swaminathan, Mike Phillips, and Maxim Likhachev. Planning for multi-agent teams with leader switching. In *IEEE International Conference on Robotics and Automation*, pages 5403–5410. IEEE, 2015.
9. Ronald C Arkin, Ronald C Arkin, et al. *Behavior-based robotics*. MIT press, 1998.
10. Aaron D Ames, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *IEEE 53rd Annual Conference on Decision and Control*, pages 6271–6278. IEEE, 2014.
11. Li Wang, Aaron D Ames, and Magnus Egerstedt. Safety barrier certificates for collisions-free multirobot systems. *IEEE Transactions on Robotics*, 33(3):661–674, 2017.
12. Bradley E Bishop. On the use of redundant manipulator techniques for control of platoons of cooperating robotic vehicles. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 33(5):608–615, 2003.
13. Gianluca Antonelli, Filippo Arrichiello, and Stefano Chialverini. Flocking for multi-robot systems via the null-space-based behavioral control. *Swarm Intelligence*, 4(1):37, 2010.
14. O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. In *IEEE International Conference on Robotics and Automation*, volume 2, pages 500–505, Mar 1985.
15. Nathan D Ratliff, Jan Issac, Daniel Kappler, Stan Birchfield, and Dieter Fox. Riemannian motion policies. *arXiv preprint arXiv:1801.02854*, 2018.
16. Ching-An Cheng, Mustafa Mukadam, Jan Issac, Stan Birchfield, Dieter Fox, Byron Boots, and Nathan Ratliff. RMPflow: A computational graph for automatic motion policy generation. In *The 13th International Workshop on the Algorithmic Foundations of Robotics*, 2018.
17. Francesco Bullo and Andrew D Lewis. *Geometric control of mechanical systems: modeling, analysis, and design for simple mechanical control systems*, volume 49. Springer Science & Business Media, 2004.
18. Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.
19. Hassan K Khalil. Nonlinear systems. *Prentice-Hall, New Jersey*, 2(5):5–1, 1996.
20. Brian DO Anderson, Changbin Yu, Baris Fidan, and Julien M Hendrickx. Rigid graph control architectures for autonomous formations. *IEEE Control Systems Magazine*, 28(6):48–63, 2008.
21. Daniel Pickem, Paul Glotfelter, Li Wang, Mark Mote, Aaron Ames, Eric Feron, and Magnus Egerstedt. The Robotarium: A remotely accessible swarm robotics research testbed. In *IEEE International Conference on Robotics and Automation*, pages 1699–1706. IEEE, 2017.

Appendices

A Proof of Theorem 3

Theorem 3. Let $\mathbf{G}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{G}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i$, $\mathbf{B}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{B}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i$, and $\Phi_r^i = \sum_{k \in \mathcal{K}_i} \Phi_{1_k}^i \circ \psi_{\mathbf{e}_k}$ be the metric, damping matrix, and potential of the tree i 's root node. If $\mathbf{G}_r^i, \mathbf{B}_r^i \succ 0$ and \mathbf{M}_r^i is nonsingular for all $i \in \mathcal{I}$, the system converges to a forward invariant set $\mathcal{C}_\infty := \{(\mathbf{q}, \dot{\mathbf{q}}) : \nabla_{\mathbf{q}_i} \Phi_r^i = 0, \dot{\mathbf{q}}_i = 0, \forall i \in \mathcal{I}\}$.

Proof. Let Φ be the total potential function for all subtasks, i.e., $\Phi = \sum_{k \in \mathcal{K}} \Phi_{1_k} \circ \psi_{\mathbf{e}_k}$. Consider the Lyapunov function candidate $V = \left(\sum_{i=1}^N K_i\right) + \Phi$, where $K_i = \frac{1}{2} \dot{\mathbf{q}}_i^\top \mathbf{G}_r^i \dot{\mathbf{q}}_i$. Then, following a derivation similar to [16], we have,

$$\begin{aligned} \frac{d}{dt} K_i &= \frac{d}{dt} \left(\frac{1}{2} \dot{\mathbf{q}}_i^\top \left(\sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{G}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i \right) \dot{\mathbf{q}}_i \right) = \frac{d}{dt} \left(\frac{1}{2} \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{G}_{1_k}^i \ddot{\mathbf{z}}_{1_k}^i + \frac{1}{2} \left(\frac{d}{dt} \mathbf{G}_{1_k}^i \right) \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{M}_{1_k}^i \ddot{\mathbf{z}}_{1_k}^i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{M}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i \ddot{\mathbf{q}}_i + \mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \dot{\mathbf{q}}_i^\top \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{M}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i \ddot{\mathbf{q}}_i + \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right). \end{aligned} \quad (7)$$

By definition of the pullback operator, we have $\mathbf{M}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{M}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i$,

$$\begin{aligned} \frac{d}{dt} K_i &= \dot{\mathbf{q}}_i^\top \mathbf{M}_r^i \ddot{\mathbf{q}}_i + \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \dot{\mathbf{q}}_i^\top \mathbf{f}_r^i + \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right). \end{aligned} \quad (8)$$

Also by definition of pullback, $\mathbf{f}_r^i = \sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top (\mathbf{f}_{1_k}^i - \mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i)$, hence,

$$\begin{aligned} \frac{d}{dt} K_i &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{f}_{1_k}^i - \mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \mathbf{M}_{1_k}^i \dot{\mathbf{J}}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(\mathbf{f}_{1_k}^i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \\ &= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(-\nabla_{\mathbf{z}_k^i} \Phi_{1_k}^i - \mathbf{B}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i - \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i + \frac{1}{2} \mathbf{G}_{1_k}^i \dot{\mathbf{z}}_{1_k}^i \right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \left(-\nabla_{\mathbf{z}_k^i} \Phi_{1_k}^i - \mathbf{B}_{1_k}^i \dot{\mathbf{z}}_k^i \right) \\
&= - \left(\sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \nabla_{\mathbf{z}_k^i} \Phi_{1_k}^i \right) - \dot{\mathbf{q}}_i^\top \left(\sum_{k \in \mathcal{K}_i} (\mathbf{J}_{\mathbf{e}_k}^i)^\top \mathbf{B}_{1_k}^i \mathbf{J}_{\mathbf{e}_k}^i \right) \dot{\mathbf{q}}_i \\
&= - \left(\sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \nabla_{\mathbf{z}_k^i} \Phi_{1_k}^i \right) - \dot{\mathbf{q}}_i^\top \mathbf{B}_r^i \dot{\mathbf{q}}_i \\
&= - \left(\sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_{1_k}^i)^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} \right) - \dot{\mathbf{q}}_i^\top \mathbf{B}_r^i \dot{\mathbf{q}}_i,
\end{aligned} \tag{9}$$

where we denote $\mathbf{z}_k = \psi_{1_k}(\mathbf{q})$ and the last equation follows from the fact that $\mathbf{z}_k^i = \psi_{1_k} = \mathbf{z}_k$ for all i .

Therefore, for the Lyapunov function candidate V , we have,

$$\begin{aligned}
\frac{d}{dt} V &= \sum_{i=1}^N \frac{d}{dt} K_i + \sum_{k \in \mathcal{K}} \dot{\mathbf{z}}_k^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} \\
&= - \sum_{i=1}^N \left(\sum_{k \in \mathcal{K}_i} (\dot{\mathbf{z}}_k^i)^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} \right) + \sum_{k \in \mathcal{K}} \dot{\mathbf{z}}_k^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} - \sum_{i=1}^N \dot{\mathbf{q}}_i^\top \mathbf{B}_r^i \dot{\mathbf{q}}_i \\
&= - \sum_{k \in \mathcal{K}} \left(\sum_{i \in \mathcal{I}_k} (\dot{\mathbf{z}}_k^i)^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} \right) + \sum_{k \in \mathcal{K}} \dot{\mathbf{z}}_k^\top \nabla_{\mathbf{z}_k} \Phi_{1_k} - \sum_{i=1}^N \dot{\mathbf{q}}_i^\top \mathbf{B}_r^i \dot{\mathbf{q}}_i \\
&= - \sum_{i=1}^N \dot{\mathbf{q}}_i^\top \mathbf{B}_r^i \dot{\mathbf{q}}_i,
\end{aligned} \tag{10}$$

where the last equation follows from $\dot{\mathbf{z}}_k = \sum_{i \in \mathcal{I}_k} \mathbf{J}_{\mathbf{e}_k}^i \dot{\mathbf{q}}_i = \sum_{i \in \mathcal{I}_k} \dot{\mathbf{z}}_k^i$. Then by LaSalle's invariance principle [19], the system converges to a forward invariant set $\mathcal{C}_\infty := \{(\mathbf{q}, \dot{\mathbf{q}}) : \nabla_{\mathbf{q}_i} \Phi_{1_k}^i = 0, \dot{\mathbf{q}}_i = 0, \forall i \in \mathcal{I}\}$. ■

B Details of the Experiments

In this appendix, we introduce the construction of unitary goal attractor RMP, which is used in many of the experiments, and provide the choice of parameters for the simulation and experiments.

B.1 Unitary Goal Attractor RMP

In multi-robot scenarios, instead of planning paths for *every* robot, it is common to plan a path or assign a goal to one robot, called the *leader*. The other robots may simply follow the leader or maintain a given formation depending on other subtasks assigned to the team. In this case, a goal attractor RMP may be assigned to the leader. A number of controllers for multi-robot systems are also based on going to a goal position, such as the cyclic pursuit behavior [3] and Voronoi-based coverage controls [3, 18].

There are several design options for goal attractor RMPs. We will discuss two examples. The first goal attractor RMP is introduced in [16]. The attractor RMP for robot i is defined on the subtask space $\mathbf{z} = \mathbf{x}_i - \mathbf{g}_i$, where \mathbf{g}_i is the desired configuration for the robot. The metric is designed as $\mathbf{G}(\mathbf{z}) = w(\mathbf{z})\mathbf{I}$. The weight function $w(\mathbf{z})$ is defined as $w(\mathbf{z}) = \gamma(\mathbf{z})w_u + (1 - \gamma(\mathbf{z}))w_l$, with $0 \leq w_l \leq w_u < \infty$ and $\gamma(\mathbf{z}) = \exp(-\frac{\|\mathbf{z}\|^2}{2\sigma^2})$ for some $\sigma > 0$. The weights w_l and w_u control the importance of the RMP when the robots are far from the goal and close to the goal, respectively. As the robot approaches the goal, the weight $w(\mathbf{z})$ will smoothly increase from w_l to w_u . The parameter σ determines the characteristic length of the metric. The main intuition for the metric is that when the robot is far from the goal, the attractor should be permissive enough for other subtasks such as collision avoidance, distance preservation, etc. However, when the robot is close to the goal, the attractor should have high importance so that the robot can reach the goal. The potential function is designed such that,

$$\nabla_{\mathbf{z}}\Phi(\mathbf{z}) = \beta w(\mathbf{z}) \left(\frac{1 - e^{-2\alpha\|\mathbf{z}\|}}{1 + e^{-2\alpha\|\mathbf{z}\|}} \right) \hat{\mathbf{z}} = \beta w(\mathbf{z}) s_{\alpha}(\|\mathbf{z}\|) \hat{\mathbf{z}}, \quad (11)$$

where $\beta > 0$, $s_{\alpha}(0) = 0$ and $s_{\alpha}(r) \rightarrow 1$ as $r \rightarrow \infty$. The parameter α determines the characteristic length of the potential. The potential function defined in (11) provides a *soft-normalization* for x so that the transition near the origin is smooth. The damping matrix is $\mathbf{B}(\mathbf{z}) = \eta w(\mathbf{z})\mathbf{I}$, where $\eta > 0$ is a positive scalar. We will refer to this goal attractor RMP as *Goal Attractor RMPa* in subsequent sections. Although more complicated, it produces better results when combined with other RMPs, especially collision avoidance RMPs (see Appendix C).

Another possible goal attractor RMP is based on a PD controller. This RMP is also defined on the subtask space $\mathbf{z} = \psi(\mathbf{x}_i) = \mathbf{x}_i - \mathbf{g}_i$. The metric is a constant times identity matrix, $\mathbf{G} = c\mathbf{I}$ with some $c > 0$. The potential function is defined as $\Phi(\mathbf{z}) = \frac{1}{2}\alpha\|\mathbf{z}\|^2$ and the damping is $\mathbf{B} \equiv \eta\mathbf{I}$, with $\alpha, \eta > 0$. This RMP is equivalent to a PD controller with $k_p = \alpha/c$ and $k_d = \eta/c$. This goal attractor will be referred to as *Goal Attractor RMPb* in subsequent sections.

B.2 Choice of Parameters

Simulated Formation Preservation Task For each robot, we define a goal attractor RMPa, with parameters $w_u = 10$, $w_l = 1$, $\sigma = 0.1$, $\beta = 0.1$, $\alpha = 10$, $\eta = 1$. We use a damper RMP with $\mathbf{G} \equiv 0.01\mathbf{I}$, $\mathbf{B} \equiv \mathbf{I}$, $\Phi \equiv 0$. For the distance preserving RMPa we set parameters $\mathbf{G} = c = 1$ and $\eta = 2$. For distance preservation RMPb (which is shown to be equivalent to the potential-based controller in the previous simulation), we choose parameters $\mathbf{G} = \mathbf{I}$ and $\eta = 2$.

Centralized RMPflow Framework We use distance preservation RMPa's with $\mathbf{G} = c = 10$, and $\eta = 5$. The safety distance between robots is $d_S = 0.18$. The parameters for the collision avoidance RMPs are set as $\alpha = 1e - 5$, $\epsilon = 1e - 8$, and $\eta = 0.5$. For goal attractors, we use goal attractor RMPa's with $w_u = 10$, $w_l = 0.01$, $\sigma = 0.1$, $\beta = 1$, $\alpha = 1$, and $\eta = 1$. The damping RMPs have parameters $\mathbf{G} \equiv 0.01\mathbf{I}$, $\mathbf{B} \equiv \mathbf{I}$, $\Phi \equiv 0$.

Decentralized Partial RMPflow Framework For the cyclic pursuit tasks, the robots are attracted by points moving along the circle with angular velocity 0.06 rad/s. The parameters for the associated goal attractor RMPa's are $w_u = 10$, $w_l = 0.01$, $\sigma = 0.1$, $\beta = 1$, $\alpha = 1$, $\eta = 1$. For robots from sub-team 2, the distance preservation RMPa's have parameters $\mathbf{G} = c = 10$, $\eta = 2$. The goal attractor are goal attractor RMPa's with parameters $w_u = 10$, $w_l = 1$, $\sigma = 0.1$, $\beta = 1$, $\alpha = 10$, $\eta = 2$. The parameters for the collision avoidance RMPs are $\alpha = 1e-5$, $\epsilon = 1e-8$, and $\eta = 1$, with safety distance $d_S = 0.18$.

C Additional Simulation Results

RMPs & Potential-based Controllers As is discussed in Section 3, many potential-based multi-robot controllers can be reconstructed by the RMP framework up to degree normalization. In this example, we consider a formation control task with five robots. The robots are tasked with forming a regular pentagon with circumcircle radius 0.4. The robots are initialized with a regular pentagon formation, but with a larger circumcircle radius of 1.

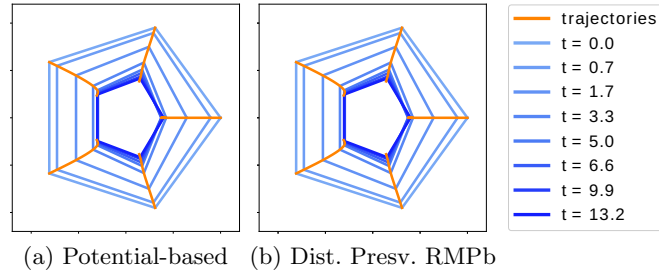


Fig. 7: (a) The potential-based controller introduced in [2]. The blue pentagons from light to dark denotes the shape from $t = 0.0s$ to $t = 13.2s$. The orange curves represent the trajectories of the robots. (b) The controller generated from the centralized RMP framework with Distance Preservation RMPa. The fact that (a) and (b) are identical shows that the two controllers have the same behavior.

We consider a degree-normalized potential field controller from [2],

$$\begin{aligned} \mathbf{u}_i &= -\frac{1}{D_i} \sum_{j:(i,j) \in E} \left(\nabla_{\mathbf{x}_i} \left\{ \frac{1}{2} (\|\mathbf{x}_i - \mathbf{x}_j\| - d_{ij})^2 \right\} - \eta \dot{\mathbf{x}}_i \right) \\ &= -\frac{1}{D_i} \sum_{j:(i,j) \in E} \left(\frac{\|\mathbf{x}_i - \mathbf{x}_j\| - d_{ij}}{\|\mathbf{x}_i - \mathbf{x}_j\|} (\mathbf{x}_i - \mathbf{x}_j) - \eta \dot{\mathbf{x}}_i \right), \end{aligned} \quad (12)$$

where d_{ij} is the desired distance between robot i and robot j , E is the set of edges in the formation graph, and D_i is the degree of robot i in the formation graph. For the RMP implementation, we use the controller given by (4). The potential-based controller (12) is equivalent to the controller generated by the distance preservation RMP given by (4) when choosing $c = \alpha = 1$. In simulation,

we choose $\eta = 2$ for both the RMP controller and the potential-based controller. The trajectories of the robots under the two controllers are displayed in Fig. 7a and Fig. 7b, respectively. The results are *identical* implying the controllers have exactly the same behavior.

Goal Attractor RMPs & Collision Avoidance RMPs An advantage of the multi-robot RMP framework is that it can leverage existing single-robot controllers, which may have desirable properties, especially when combined with other controllers. In this simulation task, the performance of the two goal attractor RMPs are compared when combined with pairwise collision avoidance RMPs. In the simulation, three robots are tasked with reaching a goal on the other side of the field while avoiding collisions with each of the others. The parameters for the collision avoidance RMPs are set as $\alpha = \epsilon = 1e - 5$, and $\eta = 0.2$. Fig. 8a and Fig. 8b show the behavior of the resulting controllers with the two choices goal attractor RMPs discussed in Section B.1, respectively. For goal attractor a, we use parameters $w_u = 10, w_l = 0.01, \sigma = 0.1, \alpha = 1$, and $\eta = 1$. we use $c = 1, \alpha = 1$, and $\eta = 2$ for goal attractor RMPb. We notice that goal attractor RMPa generates smoother trajectories compared to goal attractor RMPb.

The Centralized & Decentralized RMP Algorithms The centralized and the decentralized RMP algorithms are also compared through the same simulation of three robots reaching goals. Goal attractor RMPa's with the same parameters were used. For the collision avoidance RMP, we set $\alpha = \epsilon = 1e - 5$, and $\eta = 0.2$. The trajectories of the robots under the decentralized algorithm are illustrated in Fig. 8c. Compared to trajectories generated from centralized RMPs, the robots oscillate slightly when approaching other robots, and made aggressive turns to avoid collisions.

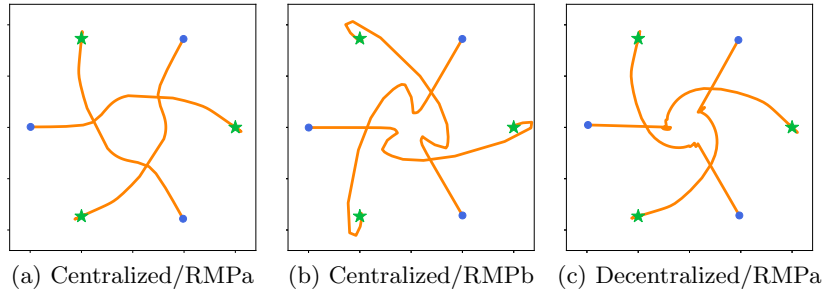


Fig. 8: The performance of goal attractor RMPs combined with pairwise collision avoidance RMPs. The blue dots and green stars denote the initial and goal positions, respectively. The trajectories are represented by orange curves. (a) The more sophisticated goal attractor RMPa generates smooth trajectories when combined with collision avoidance RMP in the centralized framework. (b) Goal attractor RMPb, which is equivalent to a PD controller, generates more distorted trajectories. (c) Under the decentralized RMP framework, the robots oscillate slightly near the origin and turn abruptly against each other.

D On Heterogeneous Robotic Teams

A significant feature of RMPs is that they are intrinsically coordinate-free [16]. Consider two robots i and j with configuration space \mathcal{C}_i and \mathcal{C}_j , respectively. Assume that there exists a smooth map ψ from \mathcal{C}_j to \mathcal{C}_i . Then the RMP-tree designed for one robot i can be directly transferred to robot j by connecting the tree to the root node of robot j through the map ψ . Therefore, RMPs provides a level of abstraction for heterogeneous robotic teams so that the user only needs to design desired behaviors for a homogeneous team with simple dynamics models, for example, double integrator dynamics, and seamlessly transfer it to the heterogeneous team. This insight could bridge the gap between theoretical results, which are usually derived for homogeneous robotic teams with simple dynamics models, and real robotics applications.