A Distributed Pipeline for Scalable, Deconflicted Formation Flying

Parker C. Lusk, Xiaoyi Cai, Samir Wadhwania, Aleix Paris, Kaveh Fathian, Jonathan P. How

Abstract—Reliance on external localization infrastructure and centralized coordination are main limiting factors for formation flying of vehicles in large numbers and in unprepared environments. While solutions using onboard localization address the dependency on external infrastructure, the associated coordination strategies typically lack collision avoidance and scalability. To address these shortcomings, we present a unified pipeline with onboard localization and a distributed, collision-free formation control strategy that scales to a large number of vehicles. Since distributed collision avoidance strategies are known to result in gridlock, we also present a distributed task assignment solution to deconflict vehicles. We experimentally validate our pipeline in simulation and hardware. The results show that our approach for solving the optimization problem associated with formation control gives solutions within seconds in cases where general purpose solvers fail due to high complexity. In addition, our lightweight assignment strategy leads to successful and quicker formation convergence in 96–100% of all trials, whereas indefinite gridlocks occur without it for 33–50% of trials. By enabling large-scale, deconflicted coordination, this pipeline should help pave the way for anytime, anywhere deployment of aerial swarms.

Index Terms—Swarms; Distributed Robot Systems; Multi-Robot Systems

Supplementary Material

Video and open-source implementation available at `https://github.com/mit-acl/aclswarm`.

I. Introduction

TWO main challenges in the deployment of large-scale swarms are the localization and coordination of vehicles. Localization methods that rely on external infrastructure (e.g., GPS) are prone to systematic errors (e.g., multipath effect) and may not always be available. Coordination strategies that are centralized can deconflict motion plans to prevent collisions and gridlock, but introduce a single point of failure and are difficult to scale in swarm size due to communication bandwidth limitations.

This paper presents a unified formation flying pipeline for unmanned aerial vehicles (UAVs). Our pipeline uses *onboard* sensors for localization, which eliminate the need for

Fig. 1. Six multirotors in a slanted plane formation. Vehicles communicate with each other, make distributed decisions onboard, and use VIO for localization.

external positioning systems, and *distributed* techniques for coordination, which enable each vehicle to make decisions independently while communicating their state to a subset of the team. For *localization*, we use an off-the-shelf commercial visual inertial odometry (VIO) package [1] that fuses inertial measurement unit (IMU) and downward-facing monocular camera measurements to estimate changes in the vehicle pose. For *coordination*, we present distributed formation control and task assignment strategies that run onboard the vehicles, do not rely on a common reference frame, and use vehicle-to-vehicle communication. Key features of our formation control strategy include scalability to a large number of vehicles and robustness to disturbances. The latter is crucial for reaching the desired formations with sensing imperfections. Our task assignment strategy uses an auction-based algorithm to guarantee conflict-free assignments. This algorithm can deconflict vehicle gridlocks resulting from distributed collision avoidance (type 3 deadlock [2]) and is well-suited for vehicles with limited computational capability and low-bandwidth communication.

A. Contributions

This research extends our previous work on UAV formations [3] and presents a unified pipeline consisting of *onboard localization* and *distributed coordination*. The three main contributions of this work are:

1) scalable formulation of control design suitable for onboard sensing without a common reference frame;
2) algorithms for deconfliction via distributed task assignment of vehicles to desired formation points;
3) simulation- and hardware-ready open-source pipeline.

Our pipeline is tested in hardware with six multirotors (see Fig. 1), and to our knowledge is the first demonstration of formation flying that does not rely on external sensing, fiducial markers for localization, a common reference frame, or a centralized base station for coordination. The only requirements for the presented pipeline are that the vehicles can

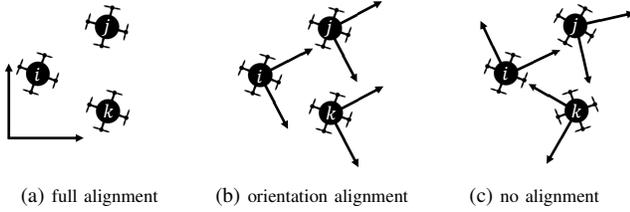(a) full alignment (b) orientation alignment (c) no alignment

Fig. 2. Required alignment of UAV frames in existing swarm strategies: (a) the most restrictive case requiring a common reference frame, i.e., orientation and origin of the frames must be aligned; (b) only the orientation of the frames must be aligned; (c) no alignment restrictions (this work).
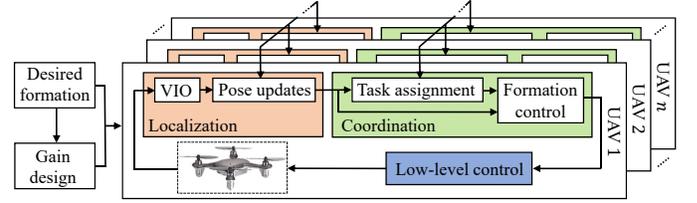

Fig. 3. Modules of our formation flying pipeline. The desired formation is used to design required gains for formation control. The localization framework provides self and relative pose measurements. The coordination framework assigns each UAV to a formation point and plans its motion to attain formation flying.

communicate, can find the transformation between their VIO start frames, and the environment is sufficiently textured—a standard assumption for VIO systems. As such, this framework paves the way for future, real-world deployments of aerial vehicle swarms in large numbers and without requiring external localization infrastructure.

B. Related Work

Existing aerial swarms can be grouped based on the coordination (centralized vs. distributed) and localization (external vs. onboard) methods used. It is further crucial to distinguish these methods based on the level of alignment required for the vehicle coordinate frames; see Fig. 2.

Works with *centralized* coordination and *external* localization include [4]–[6], which are based on lightweight UAVs with limited onboard computational capability and therefore rely on an external motion capture system and a base station. Works with *distributed* coordination and *external* localization include [7], [8], where robots execute distributed controls based on external localization by motion capture and ultrasonic beacons, respectively. Works with *centralized* coordination and *onboard* localization include [9], [10], which use a ground station for task assignment among vehicles. In [11], formation flying based on VIO is demonstrated, where motion planning and assignment are run on a base station to ensure collision-free trajectories. The coordination strategies used in aforementioned works require a *common reference frame* (Fig. 2a).

Despite the large body of work on formation control [12], and the variety of onboard sensing solutions for localization (e.g., VIO [13]), few frameworks demonstrated formation flying with *distributed* coordination and *onboard* localization. A key reason is reliance of many distributed control and assignment algorithms on aligned frames (Fig. 2a, 2b), which require computation-expensive and/or communication-intensive synchronization/consensus steps for frame alignment. Equally important, dependence on alignment in existing methods [2], [14]–[16] diminishes robustness to inherent noise and unobservable errors that cannot be corrected (e.g., disparities between the actual and estimated body frame *orientation* caused by VIO drift). Leveraging coordination methods that are *robust to misaligned frames* is hence crucial and a focus of this work.

Examples of other pipelines with distributed coordination and onboard localization include [17], [18]. Both works demonstrated formation flying on three UAVs, required information from an external motion capture system due to hardware limitations, did not incorporate collision avoidance,

and required frame alignment. Note that while [17], [18] can achieve formations with arbitrary headings as illustrated in Fig. 2c, knowledge of relative orientations is still required; therefore, they belong to the category of Fig. 2b.

II. SYSTEM OVERVIEW

A schematic representation of our pipeline is depicted in Fig. 3 for a swarm of n multirotor UAVs. The key components of this pipeline include modules for localization and coordination of the vehicles, which require exchanging information between a subset of UAV peers referred to as *neighbors*. The main goal of the pipeline is formation flying. We assume a desired formation shape is specified by an operator. This desired formation is used to design the required gains for formation control and both are given as input to the vehicles.

With onboard localization, the pose of a UAV with respect to its own start frame, which is fixed at its initial pose, is estimated using VIO. These self pose estimates provide feedback to the low-level controller, which stabilizes the UAV and tracks a reference velocity specified by the high-level formation control strategy. Through inter-vehicle pose updates, a UAV acquires relative pose estimates of its neighbors by transforming their poses into its own start frame. This process requires knowledge of the transformations that relate the UAVs' start frames. Several methods can be used to obtain these transformations. For instance, if two vehicles have a common field of view, once the correspondence among the 3D landmarks reconstructed by VIO is determined, the relative pose between the UAVs' start frames can be found using Arun's method [19]. For simplicity, the transformations are obtained in our experiments by initializing the UAVs at pre-specified locations. Note that the UAVs do *not* require the transformations to *non-neighboring* vehicles.

The coordination framework handles formation flying of the UAV swarm. This framework consists of the task assignment and formation control modules, as depicted in Fig. 3. Formation control is concerned with finding collision-free trajectories that bring the UAVs to a desired formation. A *desired formation* is defined by a graph $\mathcal{G}$ with vertices located at 3D points $p_1, \ldots, p_n$ and edges connecting the vertices to indicate neighbors. Graph $\mathcal{G}$ is used in designing the formation control and is also broadcasted to the UAVs from the base station. The vehicles aim to achieve the overall geometric shape specified by the points p_i (rather than the exact location and orientation of this point configuration in

the space). Throughout this paper, we assume that $\mathcal{G}$ is undirected, connected, and *universally rigid* [20]. Informally, rigidity implies that $\mathcal{G}$ cannot be deformed without violating the desired distances between formation points.

The goal of task assignment is to uniquely allocate each UAV to a point in the desired formation. Each UAV i is assigned to a formation point p_j using a one-to-one assignment map σ with $\sigma(i) = j$ (see Section IV). The set of neighbors of UAV i, denoted by $\mathcal{N}_i$, is defined as the set of UAVs j such that $p_{\sigma(j)}$ is connected to $p_{\sigma(i)}$ by an edge in $\mathcal{G}$. UAV i and its neighbors communicate to attain relative pose measurements using the localization framework. We emphasize that $\mathcal{N}_i$ is defined according to the current assignment map σ and formation graph $\mathcal{G}$, and that these neighbors are used for both formation control and communication. As tasks are reassigned, the set of neighbors, and therefore communication links, may change.

III. DISTRIBUTED FORMATION CONTROL

To achieve a given desired formation, we require a distributed strategy in which the UAVs execute their motions independently and are robust to misalignments of UAV frames. To make the paper self-contained, we first review a candidate strategy and then present a solution to address the scalability issue that arises for large-scale formations.

A. Overview of Formation Control

The framework for achieving a formation using only relative and local position measurements is based on [21], [22] and our previous work [3], [23]. For each UAV, the key steps in this strategy can be summarized as follows.

1) The UAV calculates the position vectors from itself to each of its neighbors in its own body frame.
2) Each vector is scaled and rotated about the z-axis of the UAV's body frame. The amount of scaling and rotation is pre-specified and depends on the desired formation.
3) These scaled and rotated vectors are then summed to obtain a resultant velocity vector command.

We emphasize that the above strategy does *not* rely on a common reference frame (Fig. 2a) and the scaling and rotation are performed in each respective UAV body frame. To formulate and analyze this framework mathematically, however, we consider a common reference frame, in which we express the position of UAV i by $q_i \in \mathbb{R}^3$ and the vector connecting UAV i to its neighbor $j \in \mathcal{N}_i$ by $q_j - q_i$. The scaling and rotation of this vector is expressed as $A_{ij}(q_j - q_i)$, where $A_{ij} \in \mathbb{A}(3)$ is called a *gain matrix* and belongs to the set of scaled rotation matrices along the z-axis denoted by

$$\mathbb{A}(3) \stackrel{\text{def}}{=} \left\{ \begin{bmatrix} a & -b & 0 \\ b & a & 0 \\ 0 & 0 & c \end{bmatrix} : a, b, c \in \mathbb{R} \right\}. \quad (1)$$

Consequently, the motion of UAV i can be expressed as

$$\dot{q}_i = \sum_{j \in \mathcal{N}_i} A_{ij} (q_j - q_i), \quad (2)$$

where $\dot{q}_i$ is the velocity vector that encapsulates the desired speed and direction of motion for the vehicle. While we consider single-integrator dynamics for simplicity of motion planning, higher-order dynamics can be utilized [23].

In (2), we assume that the z-axes of UAVs' body frames (and the reference frame used for the analysis) are aligned. In practice, the direction of gravity can be used to align these axes, and, as we will discuss in Section III-C, small misalignments caused by measurement errors or acceleration effects do not affect the convergence. Note that we do not require that the x-y axes be aligned; the UAVs can have arbitrary yaw orientations (Fig. 2c). This point distinguishes (2) from the consensus-based [24], bearing-based [25], or similar distributed control [26], in which convergence guarantees rely on orientation alignment or consensus of the UAV body frames (Fig. 2b).

To analyze the trajectory of the swarm, we define

$$q \stackrel{\text{def}}{=} \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_n \end{bmatrix}, \quad A \stackrel{\text{def}}{=} \begin{bmatrix} -\sum_j A_{1j} & A_{12} & \cdots & A_{1n} \\ A_{21} & -\sum_j A_{2j} & \cdots & A_{2n} \\ \vdots & & \ddots & \vdots \\ A_{n1} & A_{n2} & \cdots & -\sum_j A_{nj} \end{bmatrix}, \quad (3)$$

where q is the aggregate vector of UAV positions and A consists of gain matrices. Here, if UAVs i and j are not neighbors, the corresponding A_{ij} is defined as a zero matrix. Based on (2) and by using the notation in (3), swarm motion can be expressed by $\dot{q} = A q$, which determines the trajectories that the UAVs traverse.

Given a desired formation expressed via the set of points $p_i = [x_i, y_i, z_i]^\top \in \mathbb{R}^3$, we define

$$N \stackrel{\text{def}}{=} \begin{bmatrix} p_1^x & p_1^y & p_1^z & e^x & e^y & e^z \\ p_2^x & p_2^y & p_2^z & e^x & e^y & e^z \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ p_n^x & p_n^y & p_n^z & e^x & e^y & e_z \end{bmatrix} \in \mathbb{R}^{3n \times 6}, \quad (4)$$

where $p_i^x \stackrel{\text{def}}{=} [x_i, y_i, 0]^\top$, $p_i^y \stackrel{\text{def}}{=} [-y_i, x_i, 0]^\top$, $p_i^z \stackrel{\text{def}}{=} [0, 0, z_i]^\top$, $e^x \stackrel{\text{def}}{=} [1, 0, 0]^\top$, $e^y \stackrel{\text{def}}{=} [0, 1, 0]^\top$, and $e^z \stackrel{\text{def}}{=} [0, 0, 1]^\top$. Convergence to the desired shape is guaranteed if $A N = 0$ (i.e., columns of N are null vectors of A) and if all remaining eigenvalues of A not associated with N are strictly negative. For such an A to exist, each UAV should have a sufficient number of neighbors. Specifically, if vertices and edges represent the UAVs and their neighboring relations in the formation graph $\mathcal{G}$, A exists if $\mathcal{G}$ is universally rigid (see [22, Thm. 3.2]).

Our approach consists of *design* and *execution* phases. In the *design* phase, a gain matrix A as in (3) is computed offline based on the specification of a formation graph $\mathcal{G}$ and serves as an input to the distributed algorithms that run onboard the UAVs. The *execution* phase is entirely distributed, where the UAVs plan their trajectories independently using relative translation measurements to their neighboring UAVs.

B. Scalable Gain Design

Given a desired formation, the gain matrix A that meets the aforementioned constraints can be computed from

$$
\begin{aligned}
\underset{A \in \mathbb{S}_{3n}^-}{\text{minimize}} \quad & \lambda_{\max}\left(Q^\top A Q\right) \\
\text{subject to} \quad & A N = 0 \\
& A_{ij} \in \mathbb{A}(3) & \forall_{i,j} \\
& A_{ij} = 0 & \forall_i \; \forall_{j \notin \mathcal{N}_i} \\
& \text{tr}(A) = \text{constant}
\end{aligned}
\quad (5)
$$

where $\lambda_{\max}$ denotes the largest eigenvalue of a matrix, $Q \in \mathbb{R}^{3n \times (3n-6)}$ is the orthogonal complement of N (i.e., $N^\top Q = 0$), which is found from the singular value decomposition of N, and $\mathbb{S}_{3n}^-$ is the space of symmetric negative semidefinite matrices of dimension $3n$. The objective of (5) is to make the nonzero eigenvalues of A as negative as possible ($Q^\top A Q$ is the restriction of A on the subspace Q and eliminates the zero eigenvalues associated with N). By doing so, stability and robustness of the formation to noise, measurement errors, and disturbances is increased. We note that the last constraint in (5) sets the trace of A to a constant value to ensure that the problem is bounded (without this constraint, if $A \in \mathbb{S}_{3n}^-$ is a solution, so is cA for any $c > 0$ with a better objective value). The universal rigidity assumption on the formation graph [22, Thm. 3.2] is sufficient to ensure that (5) is feasible and that all remaining eigenvalues of A not associated with N are strictly negative.

The formulation (5) was presented in our earlier work [3] and can be solved relatively quickly using existing SDP solvers for small number of vehicles. However, for large-scale problems (e.g., more than 50 UAVs) it becomes challenging, or even impossible, to solve. We address this issue by exploiting the problem structure to derive a solution based on the alternating direction method of multipliers (ADMM).

We observe from (1) that $A_{ij} \in \mathbb{A}(3)$ has a block diagonal structure, which can be expressed by $A_{ij} = \text{blkdiag}(D_{ij}, c_{ij})$, where the 2×2 matrix D_{ij} consists of the first two rows and columns, and the scalar c_{ij} is the entry in the last row and column of A_{ij}. Due to this structure, we conclude from (2) that vehicle trajectories along the x-y and z components are decoupled and depend only on D_{ij} and c_{ij}, respectively. This observation allows us to split (5) into two subproblems with lower dimensions, leading to reduced computational effort. By defining

$$B \overset{\text{def}}{=} \begin{bmatrix} -\sum_j c_{1j} & c_{12} & \cdots & c_{1n} \\ c_{21} & -\sum_j c_{2j} & \cdots & c_{2n} \\ \vdots & & \ddots & \vdots \\ c_{n1} & c_{n2} & \cdots & -\sum_j c_{nj} \end{bmatrix}, \quad M \overset{\text{def}}{=} \begin{bmatrix} z_1 & 1 \\ z_2 & 1 \\ \vdots & \vdots \\ z_n & 1 \end{bmatrix}, \quad (6)$$

which correspond to the z components of A in (3) and N in (4), the problem of finding c_{ij} is formulated as

$$\begin{aligned} \underset{B \in \mathbb{S}_n^-}{\text{minimize}} \quad & \lambda_{\max}\left(R^\top B R\right) \\ \text{subject to} \quad & B M = 0 \\ & c_{ij} = 0 \qquad \forall_i \ \forall_{j \notin \mathcal{N}_i} \\ & \text{tr}(B) = \text{constant} \end{aligned} \qquad (7)$$

where $R \in \mathbb{R}^{n \times (n-2)}$ is the orthogonal complement of $M \in \mathbb{R}^{n \times 2}$. The optimization problem for finding D_{ij} is formulated similarly to (7), with an additional constraint that the diagonal entries of D_{ij} must be equal and the off-diagonal entries must have the same absolute value with different signs. With this point in mind, we henceforth focus our attention on (7). The following proposition brings (7) into the standard form suitable for applying ADMM.

Proposition 1. *Problem* (7) *can be formulated as*

$$\begin{aligned} \underset{X \in \mathbb{S}_{2n-4}^+}{\text{minimize}} \quad & \langle C, X \rangle \\ \text{subject to} \quad & \mathcal{A}(X) = b \end{aligned} \qquad (8)$$

where $\langle C, X \rangle \overset{\text{def}}{=} \text{tr}(C^\top X)$, $C \overset{\text{def}}{=} \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix}$ *with* I *as the identity matrix of size* $n - 2$, *and* $b \in \mathbb{R}^m$. *The operator* $\mathcal{A}(X)$ *represents a set of linear constraints on* X *and enforces it to have the block diagonal structure* $X \overset{\text{def}}{=} \begin{bmatrix} \gamma I & I \\ I & Z \end{bmatrix}$, *where* $\gamma \geq 0$ *and* $Z \in \mathbb{S}_{n-2}^+$. *The solution of* (7) *is obtained from* X *as* $B = -M Z M^\top$.

Proposition 1 is proved in the appendix of [27]. We now leverage the ADMM technique in [28] to solve (8). From [28], the augmented Lagrangian associated with (8) is

$$\mathcal{L} \overset{\text{def}}{=} -\langle y, b \rangle + \langle \mathcal{A}^*(y) + S - C, X \rangle + \frac{1}{2\mu} \| \mathcal{A}^*(y) + S - C \|, \quad (9)$$

where $S \in \mathbb{S}_{2n-4}^+$ and y are dual variables associated with constraints $X \in \mathbb{S}_{2n-4}^+$ and $\mathcal{A}(X) = b$, respectively, $\mathcal{A}^*$ is the adjoint of $\mathcal{A}$, and $\mu > 0$ is a penalty parameter that balances the standard Lagrangian and the augmented term. ADMM then proceeds by alternatively optimizing each primal and dual variable with others fixed, which results in a closed-form solution for each subproblem. Denoting by the superscript k the iteration number, the ADMM iterative update procedure is given as

$$\begin{aligned} y^{k+1} &= (\mathcal{A}\mathcal{A}^*)^{-1}\left(\mathcal{A}(C - S^k - \mu X^k) + \mu b\right), \\ W^{k+1} &= C - \mathcal{A}^*(y^{k+1}) - \mu X^k, \\ S^{k+1} &= \mathcal{P}_{\text{psd}}\left(W^{k+1}\right), \\ X^{k+1} &= \frac{1}{\mu}\left(S^{k+1} - W^{k+1}\right). \end{aligned} \qquad (10)$$

In (10), the operator $\mathcal{P}_{\text{psd}}$ denotes the projection onto the positive semidefinite cone $\mathbb{S}^+$, and is computed via eigendecomposition (see [28] for details). ADMM typically converges in a reasonable time to a solution with acceptable accuracy. The number of ADMM iterations required for convergence depends on the desired accuracy as well as the formation graph (e.g., when the formation graph is complete, it is straightforward to show that ADMM converges to the optimal solution in a single iteration). The time comparisons between an existing SDP solver for (5) and the presented ADMM method (10) are given in Section V.

C. Robustness, Collision Avoidance, and Formation Size

Gain matrices are recomputed only when a new desired formation is specified. During execution, vehicles use the gains and the relative position of their neighbors to compute the velocity vector u_i in (2) at each time instance. Having u_i computed, the vehicle's low-level controller is tasked with tracking the direction and speed specified by this vector.

One can show that 1) *any positive* scaling; and 2) *any rotation less than 90 degrees* of the velocity vector u_i does not void the convergence guarantees of the formation control strategy (see [3, Thm. 2]). These key properties indicate extreme robustness to errors and disturbances. For example, discrepancies between the actual and desired velocity of a vehicle caused by imperfect tracking, unmodeled dynamics, or small misalignments in z-axes of UAV body frames can be modeled as a positive scaling and small rotation of the nominal u_i, for which convergence to the desired formation is unaffected. The aforementioned properties can be further

used for collision avoidance, where velocity vectors that lead vehicles to close proximity are modified to prevent collisions. More specifically, (2) can be modified as

$$u_i \stackrel{\text{def}}{=} c_i R_i \sum_{j \in \mathcal{N}_i} A_{ij} (q_j - q_i), \tag{11}$$

where the rotation matrix R_i, which is limited to 90 degrees, is chosen to rotate any velocity vector that brings two vehicles closer than a specified distance. If there is no feasible direction of motion within this range, the scalar c_i, which is normally set to one, is set to zero to stop the vehicle.

The collision avoidance strategy (11) runs onboard, but comes at the cost of losing convergence guarantees since vehicles can become gridlocked due to the unavailability of motion directions (allowing $c_i = 0$ in (11) violates the aforementioned property in which $c_i > 0$ is required to ensure convergence). Optimal assignment of vehicles to target formation points guarantees non-intersecting lines from their current positions to their assigned points (see [14, Thm. 3.1]). Since the vehicle body frames can be nonaligned and the control is distributed, vehicles are not expected to move perfectly in a straight line under our control strategy. However, assignment can still help deconflict the swarm by increasing the availability of motion directions. The impact of assignment on resolving gridlocks is shown in Section V.

Finally, note that (11) leads to achieving the desired formation *shape*, but the formation *size* is not regulated and depends on the initial position of the vehicles. To control the size, (11) can be augmented to contract (or expand) the formation when the vehicles are farther (or closer) than the desired distance. This augmented strategy, and its theoretical convergence guarantees, is discussed in our earlier work for 2D formations (see (74) in [23]). Since the extension to 3D formations considered in this work is straightforward, we omit this discussion for brevity.

IV. DISTRIBUTED TASK ASSIGNMENT

The goal of task assignment is to uniquely allocate each UAV to a point in the desired formation. A natural objective for this task is to minimize the overall distance from the UAVs to their assignments in the desired formation. We are interested in the final 3D geometric shape rather than the exact location and yaw of the end formation. To this effect, we allow a rotation R around the z-axis, and translation t of the desired formation coordinates that minimize the overall distance from the UAVs to the rotated and translated desired formation. More precisely, our objective is to find the assignment map σ that solves

$$\underset{\substack{R \in \mathcal{R}_z, \, t \in \mathbb{R}^3 \\ \sigma \in S_n}}{\text{minimize}} \quad \sum_{i=1}^{n} \|q_i - (R \, p_{\sigma(i)} + t)\|^2, \tag{12}$$

where q_i denotes the UAV positions, S_n is the symmetric group of all permutations from the set $\{1, \ldots, n\}$ to itself, and $\mathcal{R}_z$ is the set of rotation matrices around the z-axis. Recall that the z-axes of UAV body frames are assumed to be aligned, as per Section III.

Three challenges arise in finding a distributed solution for (12). First, the objective of (12) includes the positions
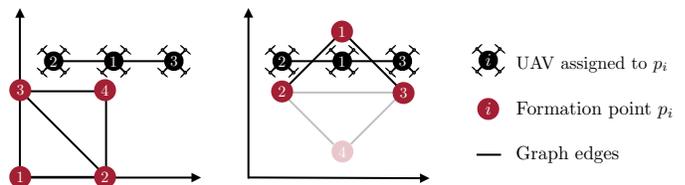

Fig. 4. Illustrative 2D alignment example with four vehicles from UAV 1's perspective. UAV 4 is not shown because it does not communicate with UAV 1. **(left)** New formation graph, UAV 1 and its neighbors before the alignment. **(right)** Aligned formation based on UAV 1, its neighbors and their corresponding formation points. The formation point associated to UAV 4 is faded to indicate that UAV 1 does not have information about it.

of all UAVs, whereas each UAV only obtains the positions of its neighbors. Second, the UAV body frames are not aligned and the UAVs only know the transformations between their body frames and their neighbors. Lastly, (12) is a nonconvex mixed integer program, for which finding the global optimizer becomes intractable for large n. As computational efficiency and scalability are of utmost concern for UAV platforms, we settle with obtaining a suboptimal answer via a coordinate descent approach and an auction assignment strategy. This approach is inspired by [29], which in contrast uses a centralized Hungarian algorithm for assignment. Every iteration of our algorithm consists of an *alignment* stage and an *assignment* stage, where the assignment is fixed as we solve for an alignment, and vice versa.

A. Alignment

Given an assignment σ^* (e.g., identity assignment $\sigma^*(i) = i$ for every new formation, or prior assignment computed for the same formation), UAV i solves a distributed formulation of (12) given by

$$\underset{R_i \in \mathcal{R}_z, \, t_i \in \mathbb{R}^3}{\text{minimize}} \quad \sum_{j \in \mathcal{N}'_i} \|q_j - (R_i \, p_{\sigma^*(j)} + t_i)\|^2, \tag{13}$$

where $\mathcal{N}'_i \stackrel{\text{def}}{=} \mathcal{N}_i \cup \{i\}$, and positions q_j are in UAV i's start frame. In (13), UAV i aims to align the desired formation to minimize the distance to its own and neighbors' positions based on the given assignment σ^*. Fig. 4 gives an illustrative example of this stage. Problem (13) is the well-known point cloud alignment problem, for which the optimal solution (R_i^*, t_i^*) is obtained from Arun's method [19] using the projection of q_j and $p_{\sigma^*(j)}$ on the x-y plane for the rotation.

B. Assignment

In this stage, the UAVs aim to collaboratively find an assignment based on the results obtained from their alignment stage. The assignment problem is formulated as

$$\underset{\sigma \in S_n}{\text{minimize}} \quad \sum_{i=1}^{n} \|q_i - (R_i^* \, p_{\sigma(i)} + t_i^*)\|^2. \tag{14}$$

Problem (14) is a linear sum assignment problem that can be solved optimally by methods such as the distributed Hungarian algorithm [30], [31], which converges in $O(n^3)$ iterations. Due to onboard resource constraints, we trade optimality for computational efficiency by using our prior work, the

consensus-based auction algorithm (CBAA) [32]. CBAA is guaranteed to converge in at most nd iterations, where d is the diameter of the formation graph, $\mathcal{G}$.

To bring (14) into the standard form for applying CBAA, let binary variables x_{ij} represent the assignment σ, where $x_{ij} = 1$ if $\sigma(i) = j$ and 0 otherwise. Further, let $c_{ij} \overset{\text{def}}{=} 1 / \|q_i - (R_i^* p_j + t_i^*)\|^2$ denote the positive score for assigning UAV i to formation point j. In practice, a small positive number can be added to the denominator of c_{ij} to avoid division by zero. It is straightforward to show that (14) can be expressed as the integer program

$$
\begin{aligned}
\underset{x_{ij} \in \{0,1\}}{\text{maximize}} \quad & \sum_{i,j=1}^{n} c_{ij}\, x_{ij}, \\
\text{subject to} \quad & \sum_{i=1}^{n} x_{ij} = 1,\ \forall_j \\
& \sum_{j=1}^{n} x_{ij} = 1,\ \forall_i
\end{aligned}
\tag{15}
$$

where the constraints on x_{ij} enforce conflict-free and one-to-one assignment captured by $\sigma \in S_n$ in (14), and maximizing (15) is equivalent to minimizing the overall distance from the UAVs to the rotated and translated formation in (14).

In executing CBAA, UAV i stores and updates its own assignment and a list of winning bids (initialized as zeros) for all formation points. Each iteration of CBAA consists of an auction phase and a consensus phase. In the auction phase, UAV i determines which formation point it would like to be assigned to in three steps: (1) check if any formation point p_j produces a score c_{ij} higher than its current winning bid; (2) of those formation points, set $x_{ij} = 1$ for the p_j that produces the highest score; (3) update the bid for the winning p_j with new score c_{ij}. In the consensus phase, vehicles converge on a common winning bid list. UAV i exchanges its winning bid list with its neighbors and updates its list with the highest values from its own and all received lists. It sets $x_{ij} = 0$ if the new winning bid for p_j is higher than c_{ij}, implying that a different vehicle has been assigned to p_j.

Since CBAA is distributed, no central authority exists to affirm convergence. Therefore, we enforce a synchronous execution to terminate the algorithm in nd iterations, which is the maximum number of iterations required to guarantee convergence. The final assignment is recovered by letting $\sigma^*(i) = j$ for each $x_{ij} = 1$. CBAA guarantees a conflict-free assignment even though UAVs do not have a common reference frame and may have inconsistent position estimates or different R_i^* and t_i^* for alignment. We emphasize that although c_{ij} is calculated by each UAV using only local knowledge, CBAA assigns UAVs to formation points without conflict. Further, it retains at least 50% of the optimal performance; that is, given the optimal overall score C^* of (15) and the C resulted from CBAA, $C/C^* \geq 0.5$.

V. EXPERIMENTAL RESULTS

This section shows that our distributed formation control and distributed task assignment solutions scale with the number of UAVs, resolve gridlocks resulting from collision avoidance, and reduce the total distance traveled.

First, we investigate scalability by comparing the runtime of our ADMM-based solver (10) with the interior-point method

TABLE I. Execution time of the CVX solver used for (5) vs. our ADMM solver (10) for obtaining formation gains for different number of vehicles. Reported times are in seconds and rounded to two decimals.

Algorithm	Number of Vehicles				
	5	20	50	100	200
CVX-SDP time	0.54	32.48	8684.24	OOM	OOM
ADMM time (ours)	0.01	0.03	1.31	12.26	134.67

OOM: Out of memory

TABLE II. Simulation results for 30 vehicles over 100 Monte Carlo trials. Using our distributed assignment algorithm, we obtain results closer to the optimal, but centralized, Hungarian approach.

		Distance Traveled (m)		Convergence Time (s)		Success
		mean	std	mean	std	
NA	nc	28.2	4.1	131.0	30.0	58%
	c	28.6	3.6	134.0	25.2	66%
A	nc	10.9	2.2	64.7	38.6	98%
	c	9.9	2.0	68.1	43.7	96%
H	c	5.1	1.0	40.6	53.5	100%

NA: no assignment   A: distributed assignment (ours)   H: centralized Hungarian
c: complete graph   nc: non-complete graph

used in CVX (http://cvxr.com/cvx) to solve the SDP formulation (5). These results are shown in Table I, and are generated in MATLAB using an Intel Core i7-7700K with 32 GB RAM. While the interior-point method becomes intractable for formations with more than 50 vehicles, our ADMM approach can solve for the control gains in seconds.

Second, we use software-in-the-loop simulations and hardware demonstrations to highlight how task assignment leads to quicker formation convergence with nearly 100% success. Our pipeline is implemented in C++ using Robot Operating System (ROS) [33]. Hardware demonstrations use a team of custom-built hexarotors, each with a diameter of 0.5 m and an all-up-weight of 1.1 kg. Code runs onboard the Qualcomm Snapdragon Flight board that includes a platform-optimized VIO package that outputs odometry at 30 Hz [1]. For simplicity of the implementation and the safety of the vehicles, we use our localization module (see Fig. 3) to inform each vehicle of every other vehicle's position. However, the information about non-neighbors is only used for collision avoidance and could alternatively be found using, for example, onboard cameras.

A. Simulations

We perform Monte Carlo trials to measure the impact of distributed task assignment on a large team of vehicles. A trial consists of randomly initializing 30 UAVs in a 20×20 m area, where the minimum distance between initial positions is 1.5 m. A random formation is generated for each trial within a $15 \times 15 \times 2$ m volume, with a minimum distance between formation points of 2 m. A trial is completed once the swarm has successfully reached the formation from the random initialization. If the swarm is trapped in a gridlock for more than 90 s, the trial is considered indefinitely gridlocked and is aborted.

For each trial, our pipeline is tested in three main configurations: with centralized assignment, with distributed assignment, and without assignment. Except for centralized
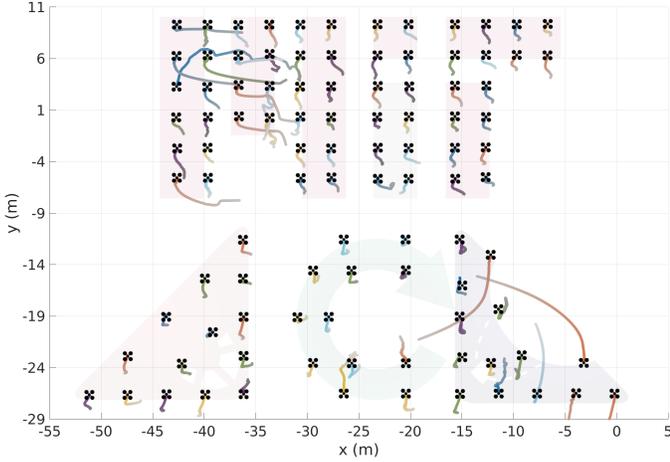
Fig. 5. Large-scale simulation with 100 UAVs. The last $40\,\text{s}$ of motion are shown and UAVs are depicted at 2x scale for better visibility.

(a) Pentagonal pyramid   (b) Triangular prism   (c) Slanted plane

Fig. 6. Non-complete formation graphs used in the hardware experiments.

— Trajectory
➚ Desired velocity
◯ Keep-out region

Fig. 7. Without assignment, UAVs attempting to achieve the pyramid formation are gridlocked due to collision avoidance.

assignment, each configuration is further tested with both a complete and randomly generated non-complete formation graph. Centralized assignment provides an optimal baseline for comparison and is performed using the Hungarian algorithm with a complete graph. The assignment algorithms are executed at a period of $2\,\text{s}$, allowing the swarm to resolve gridlocks by enabling new collision-free motion directions.

Table II shows the comparison results, where for successful trials the average distance traveled and average flying time to converge to the desired formation are reported. As expected, the centralized Hungarian approach obtains $100\,\%$ success rate with the shortest distance traveled and only an average of 2.0 reassignments to converge to the formation. However, this approach has a computational cost of $O(n^3)$ in the number of vehicles and relies on a centralized coordinator in a common reference frame with complete knowledge of the swarm (Fig. 2a). On the other hand, our CBAA-based assignment algorithm is a more scalable deconfliction strategy that is executed in the non-aligned frames of the UAVs (Fig. 2c) and is nearly optimal in practice as confirmed by the $97\,\%$ average convergence rate in Table II, with an average of 11.6 reassignments. Compared to formation control without assignment, our algorithm allows the swarm to achieve formation convergence in nearly every case and in half the amount of time, on average. The results also show that, on average, there is no significant performance decrease between complete and non-complete formation graphs. Thus, non-complete formation graphs can be used to reduce communication overhead without sacrificing the convergence rate or ability to reach the desired formation.

We remark that symmetric formations may lead our task assignment strategy (12) to exhibit momentary swapping behavior. However, noise in each UAV's sensing is included in the simulation and we have not observed convergence failure of (12) in simulation or hardware. We believe this swapping behavior is caused by ignoring the vehicle dynamics in (12) and consider this in future work.

To demonstrate scalability, we performed a large-scale simulation with 100 UAVs randomly initialized in a $60 \times 30\,\text{m}$ area. This simulation was performed using Amazon Web Services. Vehicles achieve the `MIT ACL` formation using a sparse fo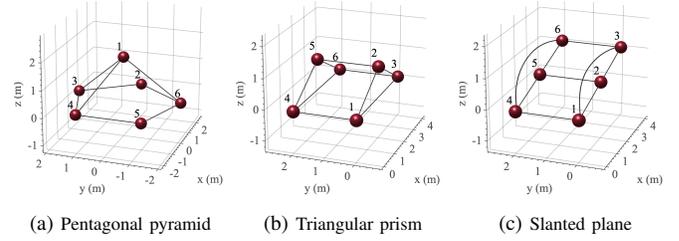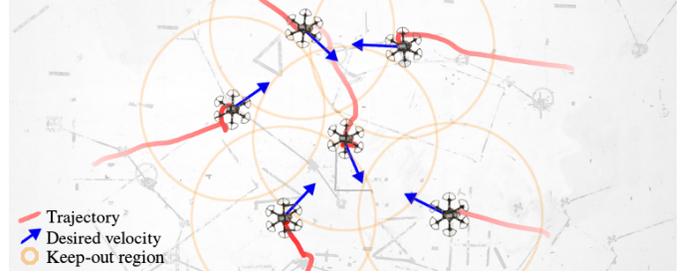rmation graph with only $24\,\%$ of the edges in a complete graph, which is beneficial for bandwidth-limited communication. The last $40\,\text{s}$ are shown in Fig. 5, where the motion traces indicate deconfliction due to reassignment.

B. Hardware Demonstrations

We demonstrate formation flight with six UAVs by cycling through the three formations illustrated in Fig. 6. The minimum distance between desired formation points is $2\,\text{m}$ for each formation. Because the time required to calculate the formation gains from (10) is small, in our experiments each UAV independently calculates the formation gains onboard in $20\,\text{ms}$. In this case, the base station is used only to dispatch the desired formation graph to the UAVs.

The UAVs are initialized at pre-specified locations so that the transforms between vehicles' VIO start frames are known. After taking off and hovering, an operator dispatches each desired formation to the swarm. Four configurations are tested by cycling through the formations twice. For each configuration, a total of six trials are recorded and averaged over, where a trial is the transition from the current swarm state to the next desired formation. When assignment is used, the period of reassignment is $1.2\,\text{s}$.

Consistent with the simulations, the results in Table III indicate that without our assignment strategy, the vehicles fail to achieve the desired formation in up to $50\,\%$ of the trials, while every trial using assignment was successful. An example of convergence failure is shown in Fig. 7.

The supplementary video provides insights into the qualitative behavior of our system. Note that the achieved formations in the video are occasionally inverted from the desired formations shown in Fig. 6. Recall that our formation control aims to achieve the desired *shape*. The inverted formations seen in experiments are due to negative scaling in the z-axis. We also point out that the formations shown in Fig. 6 are not universally rigid. Universal rigidity is a sufficient condition for our gain design, but not necessary. In practice, formations with sparser graphs can be used, so long as the recovered gain

TABLE III.   Hardware results. Our distributed assignment algorithm successfully breaks gridlock and converges to every desired formation.

		Distance Traveled (m)		Convergence Time (s)		Success
		mean	std	mean	std	
NA	nc	0.8	0.5	15.5	9.2	50 %
	c	0.9	1.0	11.3	3.0	67 %
A	nc	1.4	0.8	14.3	8.7	100 %
	c	0.8	0.6	10.1	4.8	100 %

NA: no assignment   A: distributed assignment (ours)
c: complete graph   nc: non-complete graph

matrix leads to a negative objective (8). This helps to alleviate communication load across the swarm.

The transmission requirements for localization and assignment are 5.2 kbps per neighbor and $0.064nd(n + 1)$ kb per neighbor at the reassignment period, respectively. For example, in our experiments with non-complete graphs, the theoretical bandwidth between each vehicle is approximately 9 kbps. Using a sparse graph, a reassignment period of $30\,\mathrm{s}$, and midgrade WiFi connectivity, the expected upper bound before channel saturation is 800 UAVs. These numbers are supported by our simulation of 30 UAVs in a non-complete formation graph, where we measured 2161 kbps of communication between a UAV and its neighbors.

VI. CONCLUSION AND FUTURE WORK

We presented a unified formation flying pipeline with distributed formation control and task assignment solutions that run onboard the vehicles and uses VIO for localization. Our ADMM solver addressed the scalability issue of general solvers for obtaining formation gains and the auction-based algorithm generated non-conflicting assignment solutions in a computationally efficient manner. Simulation and hardware tests demonstrated formation convergence in $96\text{--}100\,\%$ of cases that gridlocked when assignment was not used. Noteworthy future extensions include incorporating an assignment strategy that considers vehicle dynamics to minimize the total *predicted* distance traveled, and addition of distributed pose graph optimization to obtain consistent VIO pose estimates.

REFERENCES

[1] https://developer.qualcomm.com/software/machine-vision-sdk.
[2] L. Wang, A. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE TRO*, vol. 33, no. 3, pp. 661–674, 2017.
[3] K. Fathian, S. Safaoui, T. H. Summers, and N. R. Gans, "Robust 3D distributed formation control with collision avoidance and application to multirotor aerial vehicles," *IEEE ICRA*, pp. 9209–9215, 2019.
[4] J. Preiss, W. Honig, G. Sukhatme, and N. Ayanian, "Crazyswarm: A large nano-quadcopter swarm," in *IEEE ICRA*, 2017, pp. 3299–3304.
[5] W. Hönig, J. A. Preiss, T. S. Kumar, G. S. Sukhatme, and N. Ayanian, "Trajectory planning for quadrotor swarms," *IEEE TRO*, vol. 34, no. 4, pp. 856–869, 2018.
[6] X. Du, C. Luis, M. Vukosavljev, and A. Schoellig, "Fast and in sync: Periodic swarm patterns for quadrotors," in *IEEE ICRA*, 2019, pp. 9143–9149.
[7] S. Wilson, P. Glotfelter, L. Wang, S. Mayya, G. Notomista, M. Mote, and M. Egerstedt, "The Robotarium: Globally Impactful Opportunities, Challenges, and Lessons Learned in Remote-Access, Distributed Control of Multirobot Systems," *IEEE CSM*, vol. 40(1), pp. 26–44, 2020.
[8] J. Enright, M. Hilstad, A. Saenz-Otero, and D. Miller, "The SPHERES guest scientist program: Collaborative science on the ISS," in *IEEE Aerospace Conference Proceedings*, vol. 1, 2004.
[9] C. Forster, S. Lynen, L. Kneip, and D. Scaramuzza, "Collaborative monocular SLAM with multiple micro aerial vehicles," in *IEEE/RSJ IROS*, 2013, pp. 3962–3970.
[10] G. Loianno, Y. Mulgaonkar, C. Brunner, D. Ahuja, A. Ramanandan, M. Chari, S. Diaz, and V. Kumar, "A swarm of flying smartphones," in *IEEE/RSJ IROS*, 2016, pp. 1681–1688.
[11] A. Weinstein, A. Cho, G. Loianno, and V. Kumar, "Visual inertial odometry swarm: An autonomous swarm of vision-based quadrotors," *IEEE RA-L*, vol. 3, no. 3, pp. 1801–1807, July 2018.
[12] K.-K. Oh, M.-C. Park, and H.-S. Ahn, "A survey of multi-agent formation control," *Automatica*, vol. 53, pp. 424–440, 2015.
[13] J. Delmerico and D. Scaramuzza, "A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots," in *IEEE ICRA*, 2018, pp. 2502–2509.
[14] M. Turpin, N. Michael, and V. Kumar, "Capt: Concurrent assignment and planning of trajectories for multiple robots," *IJRR*, vol. 33, no. 1, pp. 98–112, 2014.
[15] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
[16] D. Morgan, G. P. Subramanian, S.-J. Chung, and F. Y. Hadaegh, "Swarm assignment and trajectory optimization using variable-swarm, distributed auction assignment and sequential convex programming," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1261–1285, 2016.
[17] E. Montijano, E. Cristofalo, D. Zhou, M. Schwager, and C. Saguees, "Vision-based distributed formation control without an external positioning system," *IEEE TRO*, vol. 32, no. 2, pp. 339–351, 2016.
[18] R. Tron, J. Thomas, G. Loianno, K. Daniilidis, and V. Kumar, "A distributed optimization framework for localization and formation control: Applications to vision-based measurements," *IEEE CSM*, vol. 36, no. 4, pp. 22–44, Aug 2016.
[19] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE TPAMI*, no. 5, pp. 698–700, 1987.
[20] S. J. Gortler and D. P. Thurston, "Characterizing the universal rigidity of generic frameworks," *Disc. & Comp. Geom.*, vol. 51, no. 4, pp. 1017–1036, 2014.
[21] Z. Lin, L. Wang, Z. Han, and v. Minyue Fu, "A Graph Laplacian Approach to Coordinate-Free Formation Stabilization for Directed Networks," *IEEE TAC*, vol. 61, no. 5, pp. 1269–1280, May 2016.
[22] Z. Lin, L. Wang, Z. Chen, M. Fu, and Z. Han, "Necessary and sufficient graphical conditions for affine formation control," *IEEE TAC*, vol. 61, no. 10, pp. 2877–2891, 2016.
[23] K. Fathian, S. Safaoui, T. H. Summers, and N. R. Gans, "Robust distributed planar formation control for higher-order holonomic and nonholonomic agents," *arXiv preprint, arXiv:1807.11058*, 2018.
[24] W. Ren, "Consensus strategies for cooperative control of vehicle formations," *IET CTA*, vol. 1, no. 2, pp. 505–512, 2007.
[25] S. Zhao and D. Zelazo, "Bearing rigidity theory and its applications for control and estimation of network systems: Life beyond distance rigidity," *IEEE CSM*, vol. 39, no. 2, pp. 66–83, 2019.
[26] E. Montijano, D. Zhou, M. Schwager, and C. Sagues, "Distributed formation control without a global reference frame," in *IEEE ACC*, 2014, pp. 3862–3867.
[27] P. C. Lusk, X. Cai, S. Wadhwania, A. Paris, K. Fathian, and J. P. How, "A distributed pipeline for scalable, deconflicted formation flying," 2020, https://arxiv.org/abs/2003.01851.
[28] Z. Wen, D. Goldfarb, and W. Yin, "Alternating direction augmented lagrangian methods for semidefinite programming," *Mathematical Programming Computation*, vol. 2, no. 3-4, pp. 203–230, 2010.
[29] E. A. Macdonald, "Multi-robot assignment and formation control," Master's thesis, Georgia Institute of Technology, 2011.
[30] S. Giordani, M. Lujak, and F. Martinelli, "A distributed algorithm for the multi-robot task allocation problem," in *IEA/AIE*. Springer, 2010, pp. 721–730.
[31] S. Chopra, G. Notarstefano, M. Rice, and M. Egerstedt, "A distributed version of the hungarian method for multirobot assignment," *IEEE TRO*, vol. 33, no. 4, pp. 932–947, 2017.
[32] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE TRO*, vol. 25, no. 4, pp. 912–926, 2009.
[33] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source Robot Operating System," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.

APPENDIX

Proof of Proposition 1. Consider problem (7). The facts that $BM = 0$ and $R \in \mathbb{R}^{n \times (n-2)}$ is the orthogonal complement of M imply that B can be factored as $B = -R Z R^\top$, where $Z \in \mathbb{S}_{n-2}^+$. Substituting B with $-R Z R^\top$ in (7) and simplifying yields

$$\begin{aligned} \underset{Z \in \mathbb{S}_{n-2}^+}{\text{maximize}} \quad & \lambda_{\min}(Z) \\ \text{subject to} \quad & \left[R Z R^\top \right]_{ij} = 0 \quad \forall_i \ \forall_{j \notin \mathcal{N}_i} \\ & \mathrm{tr}(Z) = \text{constant} \end{aligned} \quad (16)$$

which reduces the dimension of the optimization variable from n to $n - 2$. Note that the constraint $BM = 0$ in (7) is automatically satisfied in (16) as $R^\top M = 0$ by orthogonality. The objective of (16), i.e., maximizing the smallest eigenvalue of the positive semidefinite matrix Z, can be expressed equivalently as finding Z and the smallest $\gamma \geq 0$ such that $Z - \gamma^{-1} I$ remains positive semidefinite (this statement can be proved by diagonalizing Z). Hence, (16) can be expressed as

$$\begin{aligned} \underset{Z \in \mathbb{S}_{n-2}^+}{\text{minimize}} \quad & \gamma \\ \text{subject to} \quad & \gamma \geq 0, \quad Z - \gamma^{-1} I \succeq 0 \\ & \left[R Z R^\top \right]_{ij} = 0 \quad \forall_i \ \forall_{j \notin \mathcal{N}_i} \\ & \mathrm{tr}(Z) = \text{constant} \end{aligned} \quad (17)$$

Let $C \overset{\text{def}}{=} \left[\begin{smallmatrix} I & 0 \\ 0 & 0 \end{smallmatrix} \right]$ and $X \overset{\text{def}}{=} \left[\begin{smallmatrix} \gamma I & I \\ I & Z \end{smallmatrix} \right]$, where the size of the identity matrix I is the same as Z. The Schur complement condition for positive semidefinite matrices states that $X \succeq 0$ if and only if $\gamma I \succeq 0$ and $Z - I (\gamma I)^{-1} I \succeq 0$. The latter implies that $\gamma \geq 0$ and $Z - \gamma^{-1} I \succeq 0$, which are the constraints in (17). Consequently, (17) can be written concisely as

$$\begin{aligned} \underset{X \in \mathbb{S}_{2n-4}^+}{\text{minimize}} \quad & \langle C, X \rangle \\ \text{subject to} \quad & \mathcal{A}(X) = b \end{aligned} \quad (18)$$

where $\langle C, X \rangle$ is the Frobenius inner product, and $\mathcal{A}(X) = b$ captures the linear constraints on both the structure of X, i.e., the identity blocks and the last two constraints in (17). $\square$