


Persistent surveillance for unmanned aerial vehicles subject to charging and temporal logic constraints

Kevin Leahy¹  · Dingjiang Zhou¹ · Cristian-Ioan Vasile¹ · Konstantinos Oikonomopoulos¹ · Mac Schwager¹ · Calin Belta¹

Received: 11 April 2015 / Accepted: 22 October 2015
© Springer Science+Business Media New York 2015

Abstract In this work, we present a novel method for automating persistent surveillance missions involving multiple vehicles. Automata-based techniques are used to generate collision-free motion plans for a team of vehicles to satisfy a temporal logic specification. Vector fields are created for use with a differential flatness-based controller, allowing vehicle flight and deployment to be fully automated according to the motion plans. The use of charging platforms with the vehicles allows for truly persistent missions. Experiments were performed with two quadrotors for two different missions over 50 runs each to validate the theoretical results.

Keywords Persistent monitoring · Multi-robot systems · Aerial robotics · Formal methods

1 Introduction

In this paper, we investigate the automatic deployment of multiple quadrotors under resource constraints. The short battery life in many unmanned aerial vehicles (UAVs) presents a significant barrier to their use in complex, long term surveillance missions. Moreover, the use of multiple vehicles allows for more complex behavior and longer mission horizons, but further complicates the task of deploying those vehicles given limited flight time. We present an algorithm that generates a feedback controller for multiple quadrotors with charging constraints to meet a complex temporal logic specification. The algorithm comprises a three-part tool chain that first plans a high level routing sched-

ule for the quadrotors, then generates a vector field control input for the quadrotors to accomplish the schedule, and finally controls the quadrotors' nonlinear dynamics to follow the vector field with a feedback controller. The performance of the complete system, with its three interacting parts, is investigated in 50 experimental runs using two quadrotors and three charging stations in a motion capture environment as well as in several longer horizon experiments to test the efficacy of the system.

We consider the following problem: given an environment and a temporal logic mission specification with time deadlines that needs to be satisfied infinitely often, generate control policies for a team of quadrotors to complete the mission, while ensuring vehicles remain charged and collisions are avoided. The solution to this problem requires the use of several sophisticated systems, whose interaction both at a theoretical level and an experimental level produces many unique challenges.

The environment shown in Fig. 1 is presented as a motivating example, consisting of three charging stations, three regions of interest, and two aerial vehicles. Vehicle battery life is 40 time units, and charging takes 120 time units, where time units are a generic unit that can be instantiated based on a particular implementation. Given this environment and these battery and charging constraints, the vehicles must perform a persistent surveillance mission defined by a rich linear temporal logic formula which imposes time bounds on each loop of the vehicles' (infinite) runs. Thus, the specification is given as a bounded time formula which needs to be satisfied infinitely often. An example of such a mission specification to be satisfied *infinitely often* by the multi-robot system is: "within 16 time units observe Region R3 for at least 3 time units; within 28 time units, observe Region R1 for at least 2 time units; and within 46 time units, observe Region R2 for at least 2 time units then within 8 time units observe Region R1

✉ Kevin Leahy
kjleahy@bu.edu

¹ Boston University, Boston, MA 02215, USA

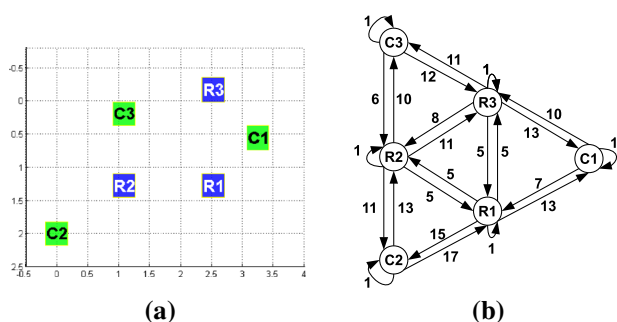


Fig. 1 **a** Partitioned environment viewed from above and **b** transition system. Green squares are charging stations, while blue squares are regions of interest. States in the transition system are charging stations and regions of interest. Weights on transitions are based on analytically calculated time bounds (Color figure online)

or Region R3 for at least 2 time units.” We seek a method to generate a control policy ensuring that vehicles can be automatically deployed to successfully complete this mission in the specified environment. Our solution is a general method for solving problems of this type, with complex missions to be automatically satisfied by a team of robots subject to charging constraints.

Our approach is related to the vehicle routing problem (VRP) (Dantzig and Ramser 1959), which can be summarized as: given a number of identical vehicles at a depot and the distances among all sites and the depot, find a minimum distance tour for each vehicle such that it begins and ends at the depot and visits each site at least once. With time bounds on when each site must be visited, the VRP becomes a problem known as the Time Window VRP (VRPTW) (Toth and Vigo 2001). Multi-agent control for the VRPTW has also been considered without temporal logic constraints in Michael et al. (2011), Stump and Michael (2011). Our work uses temporal logic constraints for the VRPTW with richer specifications, providing a framework for automatic satisfaction of complex, persistent, multi-agent routing problems.

The most closely related recent work includes Karaman and Frazzoli (2008) in which the authors propose a fragment of metric temporal logic, which restricts temporal operators to atomic propositions and their negation. In that work, each site may be visited only once, and bounds on transition duration are not allowed. Additionally, their work does not take into account resource constraints, and optimizes a weighted sum of distance traveled over a finite horizon. Our approach allows for a vehicle to visit a site multiple times during a tour if it is required, capturing resource constraints, and allowing bounds on transition durations.

Temporal logic and formal methods (Baier and Katoen 2008) have been used for robot motion planning and control in persistent surveillance in Smith et al. (2011), Ulusoy et al. (2013). These works, while considering optimal persistent surveillance with temporal logic constraints, do not consider

battery constraints. These works also do not consider time windows, which we use in this paper. Temporal logic has been used to consider resource constraints in Ozay et al. (2011), in which the authors consider constraints on peak power consumption. Our work does not take into account peak power consumption, but instead considers resource constraints in the form of total energy available for flight.

Resource constraints have been modeled in the routing problem for one vehicle without temporal logic constraints in Sundar and Rathinam (2014). Resource constraints have also been modeled for persistent monitoring in Mulgaonkar and Kumar (2014), in which the authors present a platform for autonomous charging of UAVs, including an algorithm for persistent surveillance for multiple vehicles without temporal logic constraints. Our work allows for richer mission specifications while still modeling resource constraints.

Related methods for creating routing plans appear in Vasile and Belta (2014), where a specialized logic, called Time Window Temporal Logic (TWTL) was used as a specification language. In contrast, in this work, we use a fragment of an off-the-shelf temporal logic, called bounded linear temporal logic (BLTL). BLTL and TWTL are equally expressive languages (in the sense of language equivalence), although specifications given in one language may be more concise or natural in the other. For example, in BLTL it is difficult to express that a task must be performed within a time window (i.e., after some time t_1 but before some time t_2), whereas such a specification can be stated easily in TWTL. The fragment of BLTL we consider does not permit the expression of such time window tasks, instead allowing only the expression of deadlines. In addition, in this paper we consider the continuous dynamics of the vehicles, while in Vasile and Belta (2014) the vehicles were assumed to move on a finite graph-like environment. Details on the differential flatness approach to vehicle control appear in Zhou and Schwager (2014). A preliminary version of this work appears in Leahy et al. (2014), which included fewer experimental results and no technical proofs of the differential flatness controller, vector field time bounds, or vector field derivatives.

2 Problem formulation and approach

2.1 Environment and vehicle models

Generating a control policy for our persistent surveillance problem first requires creating an abstraction of the environment and quadrotor behavior, including a model of the quadrotor battery charging and discharging. By specifying the mission using a temporal logic formula (see Sect. 2.3), we are able to use automata theoretic techniques in conjunction with these abstractions to synthesize a control policy.

We consider a team made of N identical quadrotors. A finite abstraction of the environment is given as a graph $G = (V = \mathcal{S} \cup \mathcal{C}, E, w)$, where \mathcal{S} is the set of sites and \mathcal{C} is the set of charging stations or depots. An edge $e \in E \subseteq V \times V$ denotes that travel is possible between the source and destination of the edge. Edges represent the fact that a vector field can be constructed to fly a quadrotor between the regions labeled by those two nodes (see Sect. 4). Quadrotors can deterministically choose to traverse the edges of G , stay at a site for service, or stay docked in a charging station. A duration is associated with each edge, which represents the flight time and includes docking or undocking, if applicable, and is given by $w : E \rightarrow \mathbb{Z}_{\geq 1}$. The construction of the environment graph G is described in Sect. 4.

In this paper we assume that the team has a mutually exclusive operation mode, i.e. at any moment in time at most one quadrotor is flying. Thus, collision avoidance is conservatively guaranteed. Mutually exclusive operation is useful for experimentally demonstrating our method, including the ability to charge vehicles to prolong mission horizon. Because fully concurrent operation limits mission horizon in the absence of more vehicles and charging stations, we only consider mutually exclusive operation. However it should be noted that our method may be extended to fully concurrent operation, as presented in Vasile and Belta (2014), and we plan to extend our experiments to include fully concurrent operation in the future.

Each vehicle has a limited amount of battery life, specified as an integer value, and must regularly return to a charging station. The maximum operation time starting with a fully charged battery is denoted by t_{op} , while the maximum charging time starting with an empty battery is denoted by t_{ch} . The charge-discharge ratio, which denotes the amount of time required to charge the battery vs. how long the vehicle may fly on a fully-charged battery, is $\gamma = \lceil \frac{t_{ch}}{t_{op}} \rceil \geq 1$. Using the ceiling operator provides a conservative ratio that only takes integer values. For simplicity, we assume that time is discretized, and all durations (e.g., $w(e)$, t_{op} , t_{ch}) are expressed as an integer multiple of a time interval Δt .

A battery is abstracted by a discrete *battery state* $b_t(i) \in \{0, \dots, t_{ch}\}$, corresponding to quadrotor i at time $t \in \mathbb{Z}_{\geq 0}$, and an update rule, which specifies the change of charge after d time units:

$$b_{t+d}(i) = \begin{cases} \min\{b_t(i) + d, t_{ch}\} & \text{vehicle } i \text{ is docked} \\ b_t(i) - \gamma d & \text{otherwise} \end{cases} \quad (1)$$

It is assumed that the quadrotors are equipped with identical batteries. The batteries may be charged at any of the unoccupied charging stations \mathcal{C} . Charging may start and stop at any battery state. Once a quadrotor is fully charged, it will remain fully charged until it leaves the charging station. We

assume that at the start of the mission all quadrotors are fully charged and docked at charging stations.

We will say that a quadrotor is *active* if it is flying, i.e. moving between sites and charging stations or servicing a request. A request at a site is said to be serviced if a quadrotor hovers above it. The time bounds in (2) represent the duration for which each site is to be serviced. A time interval in which all vehicles are docked and none are charging is called *idle time*. Note that idle time is therefore a property of the multi-robot system and not a property of any particular vehicle.

2.2 Routing policy

For $q \in V$, we use \vec{q} to denote that a quadrotor is flying towards q . Let $\vec{V} = \{\vec{q} \mid q \in V\}$. A *control policy* for the team of quadrotors is a sequence $\mathbf{v} = v_1 v_2 \dots$ where $v_t \in (V \cup \vec{V})^N$ specifies at each time $t \in \mathbb{Z}_{\geq 0}$ and for each quadrotor $i \in \{1, \dots, N\}$ if quadrotor i is at a site or charging station or if it is moving. Let $v_t(i)$ and $v(i)$, $i \in \{1, \dots, N\}$, denote the control value for quadrotor i at time t and the control policy for quadrotor i (i.e., the sequence of control values), respectively. Then a transition $(q_1, q_2) \in E$ performed by quadrotor i starting at time t will correspond to $v_t(i) = q_1$, $v_{t+d}(i) = q_2$ and $v_{t+k}(i) = \vec{q}_2$, $k \in \{1, \dots, d-1\}$, where $d = w((q_1, q_2))$ is the duration of the transition. Servicing or charging for one time interval (Δt time) by quadrotor i at time t corresponds to $v_t(i) = v_{t+1}(i) \in V$. A control policy $\mathbf{v} = v_1 v_2 \dots$ determines an *output word* $\mathbf{o} = o_1 o_2 \dots$ such that $o_t = \{v_t(i) \mid v_t(i) \in \mathcal{S}, i \in \{1, \dots, N\}\}$ is the set of all sites occupied by the N quadrotors at time $t \in \mathbb{Z}_{\geq 0}$. We use ϵ to denote that no site is occupied. Note o_t is either ϵ or a singleton set, because of the mutually exclusive operation mode assumption. Let $q^{[d]}$ and q^ω denote d and infinitely many repetitions of q , respectively.

Let \mathbf{v} be a control policy. We say that \mathbf{v} is *feasible* if at each moment in time all N quadrotors have non-negative battery states, i.e., $b_t(i) \geq 0$ for all $i \in \{1, \dots, N\}$ and $t \in \mathbb{Z}_{\geq 0}$.

2.3 Bounded linear temporal logic

To capture the richness of the specifications we consider, we use a fragment of bounded linear temporal logic (BLTL) (Jha et al. 2009), a temporal logic with time bounds on each of its temporal operators. The mission specification presented in Sect. 1 can be expressed as $\mathbb{G}\phi_1$, where ϕ_1 is given in (2) as a BLTL formula and the \mathbb{G} operator indicates that ϕ_1 should be satisfied infinitely often.

$$\begin{aligned} \phi_1 = & \mathbf{F}^{16} \mathbf{G}^3 R3 \wedge \mathbf{F}^{28} \mathbf{G}^2 R1 \\ & \wedge \mathbf{F}^{46} (\mathbf{G}^2 R2 \wedge \mathbf{F}^{10} \mathbf{G}^2 (R1 \vee R3)) \end{aligned} \quad (2)$$

In (2), \wedge and \vee are the usual Boolean operators indicating conjunction and disjunction, while \mathbf{F} and \mathbf{G} are the temporal

operators “eventually” and “always”, respectively. Superscripts on the temporal operators are time bounds on those operators. The fragment we consider consists of BLTL formulas in positive normal form (Baier and Katoen 2008) and using only the **F** and **G** temporal operators. This fragment allows expression of missions with deadlines on performing tasks as well as lower bounds on dwell time. Each R_i is a request associated with the region. A control policy is said to *satisfy* a persistent surveillance specification $\mathbb{G}\phi$, where ϕ is a BLTL formula, if the generated output word satisfies the BLTL formula ϕ infinitely often and there is no idle time—as defined in Sect. 2.1—between any two consecutive satisfactions of ϕ . Note that, between successive satisfactions of ϕ , the quadrotors may recharge their batteries, i.e. at least one may not be idle, because it charges its battery.

2.4 Problem formulation

The problem as informally stated in Sect. 1 is formulated in Problem 1:

Problem 1 Given an environment $G = (V = \mathcal{S} \cup \mathcal{C}, E, w)$, N quadrotors with operation time t_{op} and charging time t_{ch} , and a BLTL formula ϕ over \mathcal{S} , find a feasible control policy that satisfies $\mathbb{G}\phi$ if one exists, and design a controller to automatically deploy the quadrotors to carry out the control policy. If such a control policy does not exist, report failure.

2.5 Technical approach

There are three main components in our system: control policy generation, vector field construction, and differential flatness-based flight control, each corresponding to a different subsystem. These interacting subsystems are shown in Fig. 2. There are two steps in the solution, an offline planning stage and the online execution of the system, shown in the figure in the red and blue boxes. The solution is outlined as follows: first, a vector field is constructed offline for navigating the quadrotors, from which a finite representation in the form of a transition system is abstracted as explained in Sect. 4. Next, motion plans are generated offline to satisfy the mission specification in Sect. 3 using timing information and the transition system from the vector field subsystem. Finally, during execution, a differential flatness-based approach is used to control the vehicles through the previously constructed vector field, as presented in Sect. 5.

3 Control policy generation

The proposed approach to Problem 1 is based on automata techniques (Baier and Katoen 2008). The motion model of

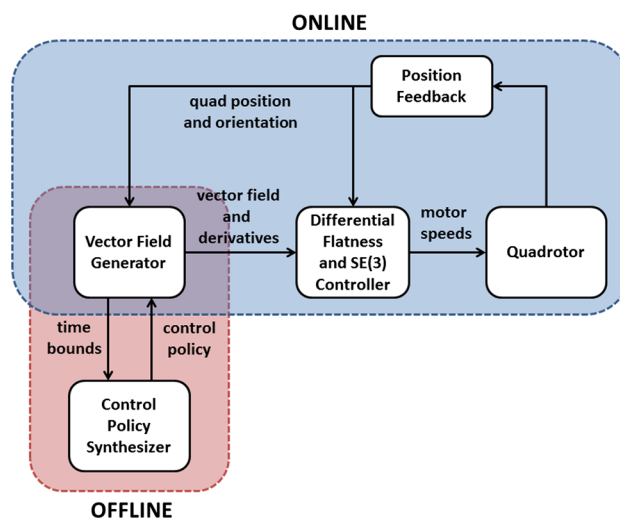



Fig. 2 A diagram of the online and offline components of the system. The red rectangle indicates the components involved in the offline planning stage, and the blue rectangle indicates those which are used during execution of the flight mission (Color figure online)

the quadrotor team is represented as a product transition system between N copies of G which is pruned of any states and transitions which violate the mutually exclusive operation mode. The product transition system is then composed with a finite state automaton which captures the charging constraints. The resulting product model is then composed with another finite state automaton which accepts the satisfying language corresponding to the given BLTL formula ϕ . The finite state automaton encoding ϕ is obtained by first translating it (Tkachev and Abate 2013) to a syntactically co-safe Linear Temporal Logic formula (Kupferman and Vardi 2001) and then to an automaton using the *scheck* tool (Latvala 2003).

Let \mathbf{v} be a feasible control policy satisfying $\mathbb{G}\phi$. We define a *loop* as a finite subsequence of \mathbf{v} starting with the satisfaction of the formula ϕ and ending before the next satisfaction of ϕ . The satisfiability problem (Problem 1) is solved on the resulting product automaton by considering all possible states of the team at the start of a loop and paths between these states obtained with Dijkstra’s algorithm. For more details about the procedure, including fully concurrent flight, see Vasile and Belta (2014), where the authors prove the completeness of the proposed approach for TWTL. Although this work considers BLTL instead of TWTL, the expressiveness of the two logics is identical, and therefore the results of Vasile and Belta (2014) apply to this work as well.

The worst case computational complexity for generating the control policy is $\mathcal{O}(2^N + N^{k+1}2^N)$, where N is the number of vehicles and k is the difference between the number of vehicles and depots (charging stations). This process is performed offline, before deploying the vehi-

Fig. 3 Vector field detail and quadrotor flight data. The cube at the top left shows a control-to-facet vector field, and the cube at the bottom left shows a stay-in-cell vector field. One of these two kinds of fields is given to the quadrotor in each cell along its path to guide it through the desired trajectory



cles. Implementing the control policy using a vector field is computationally efficient and can be performed in real-time.

4 Vector field and transition system weights

We use a vector field for the implementation of the control policies synthesized as explained in Sect. 3, because it allows for the discrete environment model to be combined with the continuous dynamics necessary for vehicle navigation. Additionally, once the vector field has been created, upper limits on travel times through the vector field provide the weights w for the environment graph G such that a control policy can be synthesized.

4.1 Partition

To generate the vector field, we first partition the environment into cubes. Each cube is defined by two vectors, $a = (a_1, a_2, a_3)$ and $b = (b_1, b_2, b_3)$ where $a_i < b_i$ for all $i = 1, 2, 3$. These vectors represent the corners of the cube closest and farthest from the origin, respectively. Thus, each cube may be written as

$$C(a, b) = \{x \in \mathbb{R}^3 | \forall i \in \{1, 2, 3\} : a_i \leq x_i \leq b_i\}. \quad (3)$$

Paths made by edges in the environment are found as sequences of these cubes. The paths are constrained such that quadrotors fly to a fixed height from the charging stations and perform all observations from that fixed altitude. From these paths, we generate vector fields to ensure each sequence of cubes is followed.

4.2 Vector field construction

A vector field everywhere inside a given cube can be created as a convex combination of a set of vectors at its vertices (Belta and Habets 2006), expressed as

$$h(x_1, x_2, x_3) = \sum_{v \in \mathbf{V}(a,b)} \prod_{i=1}^3 \left(\frac{x_i - a_i}{b_i - a_i} \right)^{\xi_i(v_i)} \times \left(\frac{b_i - x_i}{b_i - a_i} \right)^{1 - \xi_i(v_i)} h(v), \quad (4)$$

where x_i is the coordinate in the i th dimension of a point in the cube, $\mathbf{V}(a, b)$ are the vertices of cube $C(a, b)$, $h(v)$ are the vectors at each vertex $v \in \mathbf{V}(a, b)$, and $\xi_i(v_i)$ is an indicator function such that $\xi_i(a_i) = 0$ and $\xi_i(b_i) = 1$. Such a vector field can be used to keep the vehicle from leaving the cube (stay-in-cell) or to force it to leave through a given facet (control-to-facet), as displayed in Fig. 3. Constructing the vector field of the form (4) serves several purposes. First, it allows for calculation of upper bounds on exit time, as explained in Sect. 4.3, which allows for control policy generation. Second, it permits analytical calculation of the vector field and its derivatives at any point in the vector field, as explained in Sect. 5.2. Finally, such a vector field can be designed to be continuous, allowing for smooth flight during experiments.

For each cube in any given path, we create a control-to-facet vector field to lead to the next cube in the path. Because discontinuities in the vector field could result in undesirable behavior of the quadrotors, we must ensure that velocity is continuous from one cube to the next. We ensure continuity by examining vectors at the facet where cubes meet. For each corner of such a facet, the vectors from




Fig. 4 Two-dimensional example of combining vectors. **a** Control-to-facet vector field from A to B and B to C, and stay-in-cell vector field for cell C. **b** Vector conflict where A, B and C meet. **c** Final vector field, keeping only non-conflicting vector components




Fig. 5 Two-dimensional example of vector field configurations from A to B. **a** Allowable configuration results in vectors with some zero-magnitude components, while resulting in no vectors with zero-magnitude. **b** Not allowable configuration with an occurrence of zero-magnitude for all components (circled)

the two cubes are compared to each other. Only the vector components that the two vectors have in common are kept. This process is illustrated in Fig. 4. In the figure, cells A, B, and C are joined together, and B then shares a facet with A and C. The vectors for cell B and C on their shared facet are identical, and continuity is ensured. But the vectors on A's shared facet with B are different (Fig. 4b). Thus the vertical components of these vectors are discarded, but the horizontal components, which are identical, are kept (Fig. 4c). Because of this process, there are limitations to the types of arrangements of cubes that can be constructed, because they would result in a vector of zero magnitude (see Fig. 5b), but in practical examples, such arrangements are unlikely to be desirable and can be avoided by using a finer partition of the environment if necessary. It should be noted that discontinuities are inevitable when the vehicle stops and starts again in a given region. Such discontinuities are less problematic than those that arise between cubes, because the vector field can be designed with arbitrarily small initial velocity within a cube, whereas the discontinuity between two cubes may be unpredictable. Another consideration when creating the environment partition is the size of the cubes relative to the localization capabilities of the vehicles for which the vector field is being designed. As the size of the cubes approaches the localization resolution of the vehicles, the more likely the vehicle is to suffer from incorrect velocity input from the vector field.

4.3 Weights

Because satisfaction of a BLTL formula depends on the time to travel among the regions of the environment, these times must be known. We can calculate the upper bound on the travel time between any two regions, which are captured as weights on the transition system as shown in Fig. 1. This section presents the method for computing those weights. We model hovering over a region or charging as self-loop transitions of weight 1. Calculating the upper time bound for leaving a cube depends on the vectors at the vertices. If none of these vectors has a component of magnitude zero, we calculate the time bound for exiting the cube through facet F as

$$T^F = \ln \left(\frac{s_F}{s_{\bar{F}}} \right) \frac{b_i - a_i}{s_F - s_{\bar{F}}}, \quad (5)$$

where \bar{F} is the facet opposite to F , and $s_F, s_{\bar{F}}$ are the minimum vector components in the i th direction on facet F and \bar{F} , respectively. Note that \bar{F} need not be the facet through which the cube is entered. Because \bar{F} is opposite the exit facet, considering the vectors at the corners of F and \bar{F} accounts for all 8 corners of a given cube, and hence includes all of the vectors that contribute to the vector field in a given cell. A complete derivation of this bound can be found in Aydin Gol and Belta (2013). In the event that s_F approaches $s_{\bar{F}}$, T^F approaches $(b_i - a_i) / s_{\bar{F}}$.

Because of the continuity requirements on the vector field, it is possible to have a vector with a component of magnitude zero (i.e. as seen in Fig. 5a). In this case, as long as there remains a non-zero component in another direction, there is a guaranteed upper bound on the time to leave the cell. This time bound, in the case of a zero-magnitude component in the i th direction and a non-zero component in the j th direction, while exiting in the i th direction through the facet containing the zero-magnitude component, can be expressed as

$$\begin{aligned} T^F &= T_i^F + T_j^F \\ &= \left(\frac{b_i - a_i}{s_F \left(\frac{M}{2} - 1 \right)} \right) \ln \left(\frac{M}{2} \right) \\ &\quad + \left(\frac{b_j - a_j}{-2s_{\bar{F}}} \right) \ln (1 - M), \end{aligned} \quad (6)$$

where $0 < M < 1$ is a parameter that affects the tightness of the bound, due to the asymptotic nature of the solution approaching the zero-magnitude component in the i th direction. Although analytical calculation a value of M that results in the tightest bound is difficult, it can be found numerically by solving

$$2 \ln \left(\frac{M}{2} \right) AM^2 + BM^3 - (A + 4) M^2 + 3(A + 1) \times M - 2A = 0 \tag{7}$$

for M , where $A = \left(\frac{b_i - a_i}{s_{\bar{F}}} \right)$ and $B = \left(\frac{b_j - a_j}{2s_{\bar{F}}} \right)$. Proof of this time bound and the optimal value of M can be found in Appendix 1.

5 Vector field following

Motion planning often involves the use of vector fields to be followed by a robot. This is easily accomplished with most ground robots as well as slow aerial robots. In our experiments however, we use quadrotors, which cannot easily follow a vector field because of their high dimensional, nonlinear dynamics. Thus, we exploit the differential flatness of quadrotor dynamics to design a controller which will allow the quadrotor to follow the vector field constructed in Sect. 4.2, compensating for the quadrotor’s nonlinear dynamics (Zhou and Schwager 2014).

5.1 Differential flatness

Differential flatness is a property of some nonlinear systems allowing the state vector and input vector to be written in terms of a smaller number of flat outputs and their time derivatives. The function mapping the flat outputs and their derivatives to the states and inputs is known as the endogenous transformation (Wongpiromsarn et al. 2009). Quadrotor dynamics are known to be differentially flat, and we use this property to find a closed-loop controller to drive a quadrotor as if it were a simple integrator traveling through a desired velocity vector field.

Formally, a nonlinear system $\dot{\xi} = f(\xi, \mu)$ is said to be differentially flat if there exists an invertible function α such that

$$\sigma = \alpha \left(\xi, \mu, \dot{\mu}, \dots, \mu^{(d_\mu)} \right) \tag{8}$$

for a finite number of derivatives, d_μ , where σ is called the flat output. The inverse of α yields the trajectories of ξ and μ as functions of the flat outputs and d_σ of their time derivatives

$$\xi = \beta \left(\sigma, \dot{\sigma}, \dots, \sigma^{(d_\sigma)} \right) \tag{9}$$

$$\mu = \gamma \left(\sigma, \dot{\sigma}, \dots, \sigma^{(d_\sigma)} \right). \tag{10}$$

Taken together, β and γ are known as the endogenous transformation. Below, we present the endogenous transformation for a quadrotor, with the position and yaw angles as its flat output.

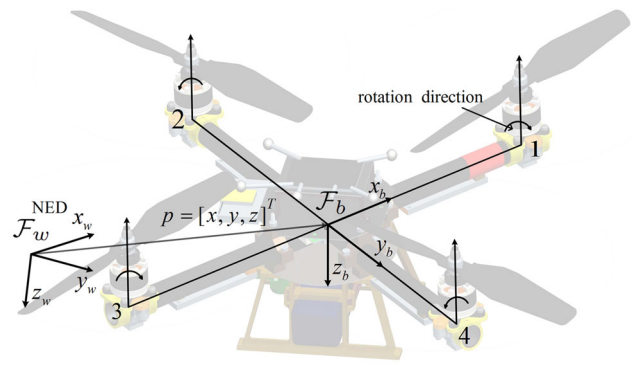


Fig. 6 Quadrotor coordinate frames with a North-East-Down world coordinate system. The world frame is denoted \mathcal{F}_w , and the aircraft body-fixed frame is \mathcal{F}_b

A quadrotor can be modeled as a rigid body with forces and torques produced by its four motors and gravity (Mellinger and Kumar 2011). The forces, moments, and coordinate frames in such a model are displayed in Fig. 6. We define the rotation matrix from the body frame to the world frame using ZYX Euler angles as

$$R = R_{(z, \psi)} R_{(y', \theta)} R_{(x'', \phi)}$$

$$= \begin{bmatrix} C\theta C\psi & S\phi S\theta C\psi - C\phi S\psi & C\phi S\theta C\psi + S\phi S\psi \\ C\theta S\psi & S\phi S\theta S\psi + C\phi C\psi & C\phi S\theta S\psi - S\phi C\psi \\ -S\theta & S\phi C\theta & C\phi C\theta \end{bmatrix}, \tag{11}$$

where ϕ is roll, θ is pitch, ψ is yaw, and $S \cdot$ and $C \cdot$ indicate $\sin(\cdot)$ and $\cos(\cdot)$, respectively. The dynamics of a quadrotor are then given by the nonlinear system of equations

$$\begin{cases} \dot{v} = g e_3 + \frac{1}{m} R f_z e_3 \end{cases} \tag{12}$$

$$\begin{cases} \dot{R} = R \Omega \end{cases} \tag{13}$$

$$\begin{cases} \dot{\omega}_b = J^{-1} \tau - J^{-1} \Omega J \omega_b \end{cases} \tag{14}$$

$$\begin{cases} \dot{p} = v, \end{cases} \tag{15}$$

where $v = [v_x, v_y, v_z]^T$ is the velocity in the world frame, g is the acceleration due to gravity, m is the mass, f_z is the total thrust force from the rotors, $e_3 = [0, 0, 1]^T$, and hence $f_z e_3$ is aligned with the negative vertical direction of the body frame, $-z_b$. R is the rotation matrix from the world frame to the body frame, defined in terms of Euler angles ψ, θ , and ϕ . The angular velocity of the quadrotor expressed in the body frame is $\omega_b = [\omega_x, \omega_y, \omega_z]^T$, and $\Omega = \omega_b^\wedge$ is the tensor form of ω_b . The torque on the quadrotor is given by τ in the body frame \mathcal{F}_b . J is the inertia matrix of the quadrotor, and $p = [x, y, z]^T$ is the position of the quadrotor in the world frame.

The system as defined in (12)–(15) has a 12-dimensional state, $\xi = [x, y, z, v_x, v_y, v_z, \psi, \theta, \phi, \omega_x, \omega_y, \omega_z]^T$, and

4-dimensional input, $\mu = [f_z, \tau_x, \tau_y, \tau_z]^T$, which is the total thrust and three torques. The state and input are differentially flat. Their flat outputs

$$\sigma = [\sigma_1, \sigma_2, \sigma_3, \sigma_4]^T := [x, y, z, \psi]^T, \quad (16)$$

consisting of position and yaw, are such that the state, ξ is a function of these outputs and their derivatives. More precisely, $\xi = \beta(\sigma, \dot{\sigma}, \ddot{\sigma}, \ddot{\ddot{\sigma}})$, with

$$\begin{cases} [x, y, z, v_x, v_y, v_z, \psi]^T \\ = \beta_{1:7}(\sigma, \dot{\sigma}) = [\sigma_1, \sigma_2, \sigma_3, \dot{\sigma}_1, \dot{\sigma}_2, \dot{\sigma}_3, \sigma_4]^T \\ \theta = \beta_8(\sigma, \dot{\sigma}, \ddot{\sigma}) = \text{atan2}(\beta_a, \beta_b) \\ \phi = \beta_9(\sigma, \dot{\sigma}, \ddot{\sigma}) = \text{atan2}(\beta_c, \sqrt{\beta_a^2 + \beta_b^2}) \\ [\omega_x, \omega_y, \omega_z]^T = \beta_{10:12}(\sigma, \dot{\sigma}, \ddot{\sigma}, \ddot{\ddot{\sigma}}) = (R^T \dot{R})^\vee, \end{cases} \quad (17)$$

where

$$\begin{cases} \beta_a = -\cos \sigma_4 \ddot{\sigma}_1 - \sin \sigma_4 \ddot{\sigma}_2 \\ \beta_b = -\ddot{\sigma}_3 + g \\ \beta_c = -\sin \sigma_4 \ddot{\sigma}_1 + \cos \sigma_4 \ddot{\sigma}_2, \end{cases} \quad (18)$$

and R is the rotation matrix with the Euler angles (ϕ, θ) defined in (17). Furthermore, the input, μ , is also a function of the flat outputs, expressed as $\mu = \gamma(\sigma, \dot{\sigma}, \ddot{\sigma}, \ddot{\ddot{\sigma}}, \ddot{\ddot{\ddot{\sigma}}})$, with

$$\begin{cases} f_z = \gamma_1(\sigma, \dot{\sigma}, \ddot{\sigma}) = -m \|\ddot{\sigma}_{1:3} - g e_3\| \\ [\tau_x, \tau_y, \tau_z]^T = \gamma_{2:4}(\sigma, \dot{\sigma}, \ddot{\sigma}, \ddot{\ddot{\sigma}}, \ddot{\ddot{\ddot{\sigma}}}) \\ = J(\dot{R}^T \dot{R} + R^T \ddot{R})^\vee + R^T \dot{R} J(R^T \dot{R})^\vee, \end{cases} \quad (19)$$

where $\ddot{\sigma}_{1:3} = [\ddot{\sigma}_1, \ddot{\sigma}_2, \ddot{\sigma}_3]^T$ for short and the $^\vee$ map is the inverse operation of $^\wedge$. For details and a proof, please refer to Zhou and Schwager (2014).

With the flat outputs and their derivatives obtained as described below, the above equations can generate all the states and inputs. A standard $SE(3)$ controller (Lee et al. 2010) can be implemented to control the quadrotor flight along the vector field using the states and inputs as a control reference. The control architecture incorporates the open-loop inputs from the differential flatness procedure as a feed-forward element, while the reference states from the differential flatness procedure are combined with the measured states to produce an error signal for the $SE(3)$ feedback controller. This feed-forward, feedback architecture is shown in Fig. 7.




Fig. 7 Block diagram for quadrotor control, with differential flatness feed-forward element and $SE(3)$ feedback controller

5.2 Vector field derivatives

The inputs described in (19) require knowledge of velocity, acceleration, jerk, and snap. Hence it is necessary to find the time derivatives ($\dot{\sigma}$, $\ddot{\sigma}$, $\ddot{\ddot{\sigma}}$, $\ddot{\ddot{\ddot{\sigma}}}$) by taking spatial derivatives of the vector field. We only consider vector fields which do not specify rotation, hence the yaw angle σ_4 is irrelevant. We arbitrarily set $\sigma_4(t) \equiv 0$. In general, the flat output derivatives $\dot{\sigma}_{1:3}$, $\ddot{\sigma}_{1:3}$, $\ddot{\ddot{\sigma}}_{1:3}$, $\ddot{\ddot{\ddot{\sigma}}}_{1:3}$ at any point p in a vector field $h(p)$ can be recursively calculated by

$$\begin{cases} \dot{\sigma}_{1:3}(p) = h(p) \\ \ddot{\sigma}_{1:3}(p) = \mathcal{J}(\dot{\sigma}_{1:3}(p), p) \dot{\sigma}_{1:3}(p) \\ \ddot{\ddot{\sigma}}_{1:3}(p) = \mathcal{J}(\ddot{\sigma}_{1:3}(p), p) \ddot{\sigma}_{1:3}(p) \\ \ddot{\ddot{\ddot{\sigma}}}_{1:3}(p) = \mathcal{J}(\ddot{\ddot{\sigma}}_{1:3}(p), p) \ddot{\ddot{\sigma}}_{1:3}(p), \end{cases} \quad (20)$$

where $\mathcal{J}(f(p), p)$ denotes the Jacobian matrix of the function $f(p)$.

The velocity is obtained directly from the vector field described by (4), from which the derivatives required for the differential flatness controller given in (20) can be derived analytically. First (4) is rewritten in matrix form as

$$h(p_1, \dots, p_3) = [c_1, \dots, c_8] \begin{bmatrix} h_{1p_1} & h_{1p_2} & h_{1p_3} \\ \vdots & \vdots & \vdots \\ h_{8p_1} & h_{8p_2} & h_{8p_3} \end{bmatrix}. \quad (21)$$

In this form, the coefficients c are functions of position, but the values of h are fixed for any given cube. This form is therefore convenient for computation of the acceleration and other vector field derivatives.

In general, the acceleration at p is given by

$$a(p) = \mathcal{J}(v(p), p)v(p), \quad (22)$$

where $\mathcal{J}(f(p), p)$ denotes the Jacobian matrix of the function $f(p)$, which is a 3×3 matrix with entries

$$\mathcal{J}_{ij} = \frac{\partial v_i}{\partial p_j} = h_{1p_i} \frac{\partial c_1}{\partial p_j} + \dots + h_{8p_i} \frac{\partial c_8}{\partial p_j}. \quad (23)$$

Through straightforward calculation, acceleration is therefore given by

$$a_i = \sum_{j=1}^3 \left(\sum_{k=1}^8 h_{kp_i} \frac{\partial c_k}{\partial p_j} \right) v_j. \quad (24)$$

It should be noted that the vector fields for acceleration, jerk, and snap are continuous everywhere within a given cube but may be discontinuous at the facets between cubes. Similar calculations can be made for jerk and snap, and are presented in Appendix 2.

6 Results and experiments

The partitioned environment (Figs. 1 and 12) consists of 385 cubes each with edge length 0.36 m. Control policies for $\mathbb{G}\phi_1$ —where ϕ_1 is given as (2)—were calculated over the transition system displayed in Fig. 1. The computation time, excluding encoding of (2), was 301.7 s on a Linux system with a 2.1 GHz processor and 32 GB memory, and the final product automaton had 579,514 nodes and 2,079,208 edges. No solutions were found for quadrotors starting on Chargers C2 and C3, but all other combinations of starting positions yielded solutions. The parameters t_{ch} and t_{op} were 120 and 40 time units, respectively.

Control policies were also computed for $\mathbb{G}\phi_2$, where ϕ_2 is given as

$$\begin{aligned} \phi_2 = & F^7(G^2R2 \wedge F^7G^4R1) \\ & \wedge F^{45}(G^2R2 \wedge F^{14}G^2R3). \end{aligned} \tag{25}$$

This new specification can be understood as “within 7 seconds observe Region R2 for at least 2 seconds then within 5 seconds observe Region R1 for at least 4 seconds and within 45 seconds observe Region R2 for at least 2 seconds then within 12 seconds observe Region R3 for at least 2 seconds.” The transition system, t_{ch} , and t_{op} were the same as for $\mathbb{G}\phi_1$. The final product automaton had 284,550 nodes and 998,574 edges, and required 180 s to compute, excluding encoding of (25).

Experiments were performed in the Boston University Multi-robot Systems Lab. The lab consists of a flight space with IR cameras to track reflective markers on the quadrotors using an OptiTrack system. This system allows for real-time localization of the quadrotors during experiments. Two K500 quadrotors from KMeI robotics were used to execute the control policies described in Sect. 6.

Charging stations (Fig. 8) were designed and built at Boston University for automatic docking and charging of quadrotors. These platforms allow a vehicle to land when its battery requires charging. When using multiple such platforms, another vehicle can then take off, ensuring continuous monitoring in situations where one vehicle may not be able to satisfy a persistent monitoring mission specification on its own. A screenshot of the GUI is shown in Fig. 9.

The charging stations are made of laser cut acrylic parts connected with PLA plastic 3D printed parts. The electronics of the station consist of the Hyperion EOS0720i Net3AD charger, modified to enable control by MATLAB. To secure a robust connection with the stainless steel pads of the charging station, the quadrotors are equipped with stainless steel contacts mounted on springs with magnets. The platform is entirely controlled by MATLAB via USB connection, allowing for the detection of the presence of a quadrotor, real-




Fig. 8 Quadrotor resting on charging station

time monitoring of battery and charging status, and control of the charging parameters including battery type, capacity, and charging rate. The maximum charging rate that can be achieved is 8 Amperes.

Two sets of experiments were performed. In the first, a shorter version of the persistent surveillance mission was run 50 times each for $\mathbb{G}\phi_1$ and $\mathbb{G}\phi_2$ to validate the satisfaction of the mission specifications, specifically with respect to time bounds. The second set of experiments consisted of running the system for $\mathbb{G}\phi_1$ until loss of battery power in order to demonstrate the persistent abilities afforded the team by the charging stations. The two sets of experiments are presented in Sects. 6.1 and 6.2.

6.1 Short horizon experiments

Figures 10 and 11 show the results of a flight by two quadrotors for $\mathbb{G}\phi_1$ and $\mathbb{G}\phi_2$, respectively. Seconds were used as the time units for these experiments so flights could be rapidly performed and analyzed.

For $\mathbb{G}\phi_1$, the quadrotors, shown in red (Quad 1) and blue (Quad 2) in Fig. 12, start fully charged from the charging stations C1 and C2, respectively. The control policy \mathbf{v} for the two quadrotors, generated as described in Sect. 3, is the following:

$$\begin{aligned} v(1) = & C1^{[1]} \vec{R1}^{[6]} R1^{[3]} \vec{R3}^{[4]} R3^{[4]} \vec{C3}^{[10]} C3^{[35]} \\ & \left(C3^{[31]} \vec{R2}^{[5]} R2^{[3]} \vec{R1}^{[4]} R1^{[3]} \vec{R2}^{[4]} R2 \vec{C3}^{[9]} \right)^\omega \\ v(2) = & C2^{[29]} \vec{R2}^{[12]} R2^{[3]} \vec{R1}^{[4]} R1^{[3]} \vec{C1}^{[12]} \\ & \left(C1^{[1]} \vec{R1}^{[6]} R1^{[3]} \vec{R3}^{[4]} R3^{[4]} \vec{C1}^{[12]} C1^{[30]} \right)^\omega. \end{aligned} \tag{26}$$

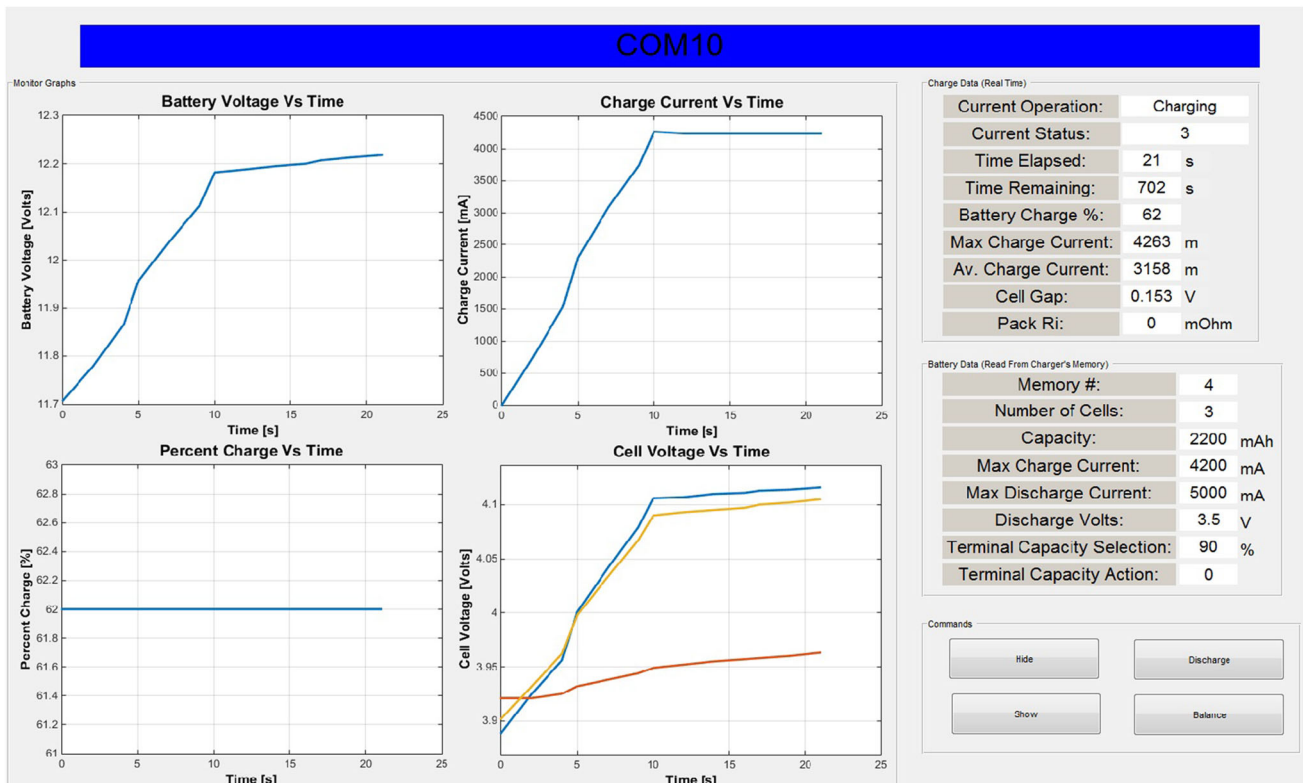


Fig. 9 Graphical user interface for charging stations. Interface displays graphs of battery voltage, battery current, percent of full charge, and individual cell voltages versus time. It also displays other battery information on the right hand side




Fig. 10 Timeline of quadrotor flights for two loops for ϕ_1 . The first two rows display the first loop, with Quadrotor 1 flying before Quadrotor 2. The next two rows show the second loop, with Quadrotor 2 flying first

Under control strategy (26), in the first loop Quadrotor 1 (red) takes off first and services sites $R1$ and $R3$ and Quadrotor 2 (blue) completes the loop by servicing sites $R2$ and $R1$. In all subsequent loops, Quadrotor 2 (blue) takes-off first and services sites $R1$ and $R3$ and Quadrotor 1 completes the loop by servicing sites $R2$ and $R1$. After the first loop, Quadrotors 1 and 2 always return to $C3$ and $C1$, respectively. The corresponding output word is

$$o = \epsilon^{[7]} R1^{[3]} \epsilon^{[4]} R3^{[4]} \epsilon^{[23]} R2^{[3]} \epsilon^{[4]} R1^{[3]} \epsilon^{[4]} R2 \epsilon^{[9]} \left(\epsilon^{[7]} R1^{[3]} \epsilon^{[4]} R3^{[4]} \epsilon^{[18]} R2^{[3]} \epsilon^{[4]} R1^{[3]} \epsilon^{[4]} R2 \epsilon^{[9]} \right)^\omega$$

The flights presented in the experiments consist of the first two loops each satisfying ϕ_1 . Any subsequent loop would be identical to the second loop. Since ϕ_1 can be satisfied repeatedly, these flights can satisfy the mission specification, $\mathbb{G}\phi_1$.

Figure 10 shows that the specification was satisfied for both loops in the flight. Region $R1$ was visited in Loop 1 and 7.48 s in Loop 2, ahead of the 28 s deadline. Likewise, Region $R3$ was visited in 12.44 and 12.64 s ahead of the 16 s deadline. In the second portion of each loop, Region $R2$ was visited in 34.00 and 30.27 s with a deadline


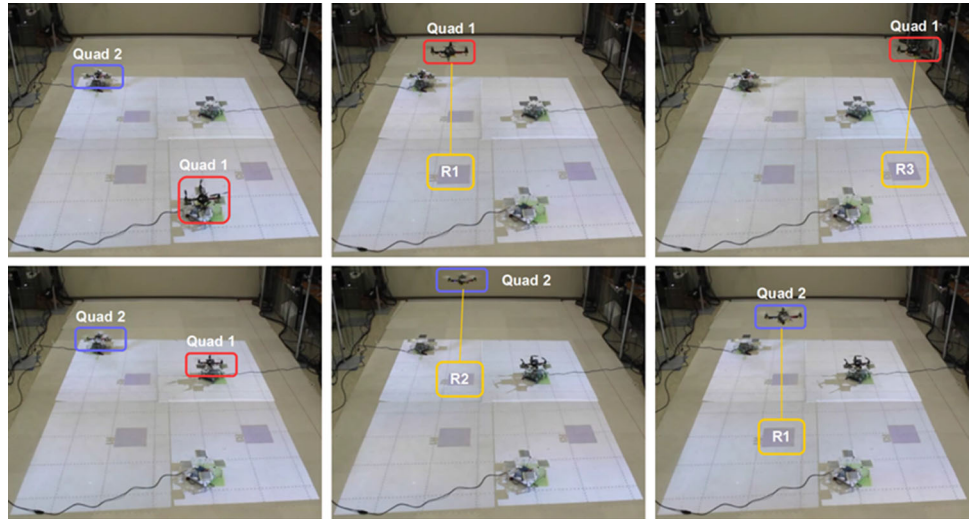


Fig. 11 Timeline of quadrotor flights for two loops for ϕ_2 . The first two rows display the first loop, with Quadrotor 1 flying before Quadrotor 2. The next two rows show the second loop, with Quadrotor 2 flying first

Fig. 12 Screenshots of the first flight loop for $\mathbb{G}\phi_1$



of 46 s, and Region $R1$ was visited within the 8 s deadline after each visit to Region $R2$.

Likewise, for specification $\mathbb{G}\phi_2$, the control policy for the two quadrotors was:

$$\begin{aligned}
 v(1) &= \left(C3^{[1]} \vec{R}2^{[5]} R2^{[3]} \vec{R}1^{[4]} R1^{[5]} \vec{C}1^{[12]} C1^{[71]} \right. \\
 &\quad \left. \vec{R}1^{[6]} R1 \vec{R}2^{[4]} R2^{[3]} \vec{R}1^{[4]} R1 \vec{R}3^{[4]} R3^{[3]} \vec{C}3^{[10]} \right)^\omega \\
 v(2) &= \left(C2^{[31]} \vec{R}2^{[12]} R2^{[3]} \vec{R}1^{[4]} R1 \vec{R}3^{[4]} R3^{[3]} \vec{C}3^{[10]} \right. \\
 &\quad \left. C3^{[1]} \vec{R}2^{[5]} R2^{[3]} \vec{R}1^{[4]} R1^{[5]} \vec{C}2^{[14]} C2^{[37]} \right)^\omega. \quad (27)
 \end{aligned}$$

Under this control policy, Quadrotor 1 starts on Charger $C3$ and takes off first, servicing Regions $R2$ and $R1$, and then landing on Charger $C1$. Next, Quadrotor 2, starting on Charger $C2$, takes off and services Regions $R2$ and $R3$, before landing on Charger $C3$. For the second loop, Quadrotor 2 begins, servicing Regions $R2$ and $R1$, followed by Quadrotor 1 servicing regions $R2$ and $R3$. After the second loop, the quadrotors are in their initial configuration. Any subsequent two-loop flight would be identical to this

sequence of two loops. The output word associated with these control policies is:

$$\begin{aligned}
 o = & \left(\epsilon^{[6]} R2^{[3]} \epsilon^{[4]} R1^{[5]} \epsilon^{[25]} R2^{[3]} \epsilon^{[4]} R1 \epsilon^{[4]} R3^{[3]} \epsilon^{[16]} \right. \\
 & \left. R2^{[3]} \epsilon^{[4]} R1^{[5]} \epsilon^{[21]} R1 \epsilon^{[4]} R2^{[3]} \epsilon^{[4]} R1 \epsilon^{[4]} R3^{[3]} \epsilon^{[10]} \right)^\omega.
 \end{aligned}$$

Figure 11 shows the satisfaction of the specification for both flight loops. Region $R2$ was serviced in 6.05 and 4.59 s in Loop 1 and Loop 2, respectively, ahead of the 7 s deadline. Region $R1$ was visited in 11.75 and 10.85 s for Loops 1 and 2, with a deadline of 14 s. For the second portion of Loops 1 and 2, Region $R2$ was visited in 36.24 and 39.84 s, ahead of the deadline of 45 s. Finally, Region $R3$ was visited in 48.90 and 48.76 s, with a deadline of 59 s.

The two-loop flights described above was performed 50 times for each specification, and both quadrotors were consistent in their flight times. The standard deviation in the length of each portion of the flight time was on the order of 0.1 s for both specifications. Despite this consistency, the time bound on flying from Charger $C1$ to Region $R1$ was violated by

Table 1 Results of long horizon experiments

Batteries	Init. voltage (V)	Flight time w/o charging (min:s)	Flight time w/ charging (min:s)	Percent increase
Old	12.5, 12.6	19:09	24:22	27
New	12.5, 12.5	22:53	34:36	51

the second quadrotor in each flight, while not being violated by the first quadrotor for specification ϕ_1 . While the vehicles were nominally identical, small physical differences between them required the controllers to be tuned using different values. Because both quadrotors followed the same vector field using the same controller, this time bound violation suggests some potential for better tuning of the controllers. No such inconsistency occurred for specification ϕ_2 .

6.2 Long horizon experiments

The missions we consider in this work require that a BLTL formula ϕ be satisfied infinitely often. For two vehicles to satisfy this specification perpetually, there must be a period in which both vehicles are charging and neither is flying. This requirement follows from the fact that the time required to fully charge a battery is in general about three times longer than the flight time for a fully charged battery. Therefore, if we wish to have at least one vehicle airborne at any given time (i.e., constant surveillance), two vehicles are insufficient to perpetually satisfy ϕ . The charging stations should nonetheless extend the feasible mission horizon when at least one vehicle is airborne at all times, despite the fact that loss of battery power is inevitable with constant flight for only two vehicles.

Two experiments were performed to test the extra endurance afforded by the use of charging stations: one with new batteries, and one with batteries that have been used on the quadrotors previously, both for specification $\mathbb{G}\phi_1$. In both experiments, the batteries started fully charged. Performing experiments with two sets of batteries allows us to control for effects due to the age of the batteries. With each set of batteries, the system was tested until failure occurred—that is, until a quadrotor ran out of charge—using the charging stations to recharge the batteries during mission execution and without using the charging stations. Results from these experiments are displayed in Table 1. With both the new and old batteries, charging increased mission horizon substantially, with greater increase in flight time with new batteries (51 vs. 27%).

7 Conclusion

We presented a method for automatic control policy synthesis and vehicle deployment for a persistent surveillance mission

for agents with charging constraints. The implementation of the persistent surveillance framework required three systems to be integrated together: a BLTL control synthesis algorithm, a vector field generation algorithm, and a quadrotor differential flatness controller. Because a conservative approach was used, such as using upper bounds on travel time rather than expected travel time, the system met the specifications reliably and predictably.

Our method easily and effectively accommodates rapid experimentation for different mission specifications, environments, or numbers of vehicles. By using the environment partition and transition system generation with time bounds, minimal human input is required to execute such missions. That is, if the user specifies a surveillance mission as well as the locations of regions of interest, charging stations, and vehicles, execution of the mission requires no further human intervention. Further, the inclusion of charging stations, whose performance can be modeled using automata, allows us to extend the feasible horizon of such missions. With an appropriate number of vehicles, the charging stations should also accommodate perpetual surveillance missions.

These experiments establish a framework that can be extended to a variety of future work. Our deterministic model of battery life is limited and one future research direction involves robust mission planning with a stochastic battery model. We are particularly interested in conducting experiments involving missions that require multiple vehicles to be airborne simultaneously. Such missions would involve more complex distributed tasks, such as simultaneously servicing several sites, or distributing tasks among subgroups of agents. Along those lines, we are also interested in extending this work to longer mission horizons with the use more vehicles, especially perpetual flight with at least one agent airborne at all times.

Acknowledgments This work was supported in part by NSF Grant Numbers CNS-1035588, NRI-1426907 and CMMI-1400167 and ONR Grant Numbers N00014-12-1-1000, MURI N00014-10-10952 and MURI N00014-09-1051.

Appendix 1: Derivation of time bounds

The derivation for (6) follows the same structure as that of (5), which can be found in Aydin Gol and Belta (2013). That derivation involves finding the minimum velocity vector towards the exit facet, and solving a linear system to find




Fig. 13 Vector field with zero-magnitude component for deriving time bounds

the time taken to exit at that velocity. In our work, however, the minimum velocity towards the exit facet may be zero, and so an alternate method must be used to compute the time bound. For this derivation, we assume that positive x is the direction of the desired exit facet, as displayed in Fig. 13. In the event that the minimum magnitude of velocity towards the exit facet is zero, we restrict velocity in one of the other coordinates to be non-zero away from the other facets, which in this figure is the y direction, but holds also for the z direction. Following from (4),

$$\dot{x} = \frac{b_i - x}{b_i - a_i} s_{\bar{F}} + \left(1 - \frac{b_i - x}{b_i - a_i}\right) s_F, \tag{28}$$

where s_F and $s_{\bar{F}}$ are the vectors in the x direction away from the exit facet F and the opposite facet \bar{F} . But the magnitude of \dot{x} depends on the y position through s_F and $s_{\bar{F}}$.

We separate the x and y directions in order to bound the time to exit the cube without needing to solve the coupled nonlinear equations of the vector field. First, we note that in (28),

$$s_F = \left(1 - \frac{b_j - y}{b_j - a_j}\right) h, \tag{29}$$

where h is the magnitude of the vector at the corner of the cube in the y direction. We write the dynamics for the y direction as

$$\dot{y} = \frac{b_j - y}{b_j - a_j} h + \left(1 - \frac{b_i - y}{b_j - a_j}\right) (-h), \tag{30}$$

which rearranges to

$$\dot{y} = -\frac{2h}{b_j - a_j} y + h. \tag{31}$$

This equation asymptotically approaches equilibrium at

$$y = \frac{b_j - a_j}{2} \tag{32}$$

which means that getting a finite solution for time to equilibrium is not possible. However, we can solve for the time to some fraction of its equilibrium, $y^* = M \frac{b_j - a_j}{2}$, where $0 < M < 1$. The linear system in (31) can be solved explicitly for the time to reach y^* as

$$t_y = \frac{b_j - a_j}{-2h} \ln(1 - M). \tag{33}$$

Then, we can substitute $M \frac{b_j - a_j}{2}$ for y in (29) to get

$$\dot{x} = \frac{\frac{M}{2} - 1}{b_i - a_i} hx + h, \tag{34}$$

which can be solved explicitly for the time to reach $x = b_i$, yielding

$$t_x = \frac{b_i - a_i}{h \left(\frac{M}{2} - 1\right)} \ln\left(\frac{M}{2}\right). \tag{35}$$

Adding (33) and (35) yields the time bound in (6).

To solve for the value of M that gives the tightest bound, we must take the derivative of (33) and (35). Starting with (33), we find

$$\frac{dt_y}{dM} = \left(\frac{b_j - a_j}{-2s_{\bar{F}}}\right) \left(\frac{1}{M - 1}\right). \tag{36}$$

Similarly, taking the derivative of (35) yields

$$\begin{aligned} \frac{dt_x}{dM} &= \frac{1}{M} \left(\frac{b_i - a_i}{s_F \left(\frac{M}{2} - 1\right)}\right) \\ &\quad - 2 \ln\left(\frac{M}{2}\right) \left(\frac{b_i - a_i}{s_F}\right) \frac{1}{(M - 2)^2}. \end{aligned} \tag{37}$$

The quantities $b_i - a_i$, $b_j - a_j$, s_F , and $s_{\bar{F}}$ are all non-negative, and hence we can replace $\left(\frac{b_i - a_i}{s_F}\right)$ with A and $\left(\frac{b_j - a_j}{2s_{\bar{F}}}\right)$ with B and rearrange to get (7). Since (6) is convex, the solution to (7) corresponds to a value of M such that the time bound given by (6) is minimized.

Appendix 2: Analytical calculation of vector field derivatives

As with calculation of acceleration in (22), jerk j can be computed by applying the Jacobian to the acceleration as

$$\begin{aligned} j &= \dot{a}(v(p(t)), p(t)) = \frac{\partial a}{\partial p} \frac{dp}{dt} \\ &= \begin{bmatrix} \frac{\partial a_1}{\partial p_1} & \frac{\partial a_1}{\partial p_2} & \frac{\partial a_1}{\partial p_3} \\ \vdots & \vdots & \vdots \\ \frac{\partial a_3}{\partial p_1} & \frac{\partial a_3}{\partial p_2} & \frac{\partial a_3}{\partial p_3} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}. \end{aligned} \tag{38}$$

The partial derivatives of acceleration can be solved by differentiating the terms for acceleration to get

$$\frac{\partial a_i}{\partial p_j} = \begin{bmatrix} \frac{\partial J_{i1}}{\partial p_j} & \frac{\partial J_{i2}}{\partial p_j} & \frac{\partial J_{i3}}{\partial p_j} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} + \begin{bmatrix} J_{i1} & J_{i2} & J_{i3} \end{bmatrix} \begin{bmatrix} J_{1j} \\ J_{2j} \\ J_{3j} \end{bmatrix}, \quad (39)$$

where J_{ij} is the Jacobian as defined in (23). These partial derivatives can be solved as

$$\frac{\partial J_{ij}}{\partial p_k} = \begin{bmatrix} \frac{\partial^2 c_1}{\partial p_j \partial p_k} & \cdots & \frac{\partial^2 c_8}{\partial p_j \partial p_k} \end{bmatrix} \begin{bmatrix} h_{1i} \\ \vdots \\ h_{8i} \end{bmatrix}. \quad (40)$$

In this equation, h_{ij} is the p_j th component of the vector at the i th vertex, and the c_i 's are the coefficients calculated in (21).

The same process is used to calculate snap:

$$\frac{\partial j}{\partial p} \frac{dp}{dt} = \begin{bmatrix} \frac{\partial j_1}{\partial p_1} & \frac{\partial j_1}{\partial p_2} & \frac{\partial j_1}{\partial p_3} \\ \vdots & \vdots & \vdots \\ \frac{\partial j_3}{\partial p_1} & \frac{\partial j_3}{\partial p_2} & \frac{\partial j_3}{\partial p_3} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \quad (41)$$

$$\begin{aligned} \frac{\partial j_i}{\partial p_j} &= \begin{bmatrix} \frac{\partial^2 a_i}{\partial p_1 \partial p_j} & \frac{\partial^2 a_i}{\partial p_2 \partial p_j} & \frac{\partial^2 a_i}{\partial p_3 \partial p_j} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} + \begin{bmatrix} \frac{\partial a_i}{\partial p_1} & \frac{\partial a_i}{\partial p_2} & \frac{\partial a_i}{\partial p_3} \end{bmatrix} \\ &\times \begin{bmatrix} J_{1j} \\ J_{2j} \\ J_{3j} \end{bmatrix} \end{aligned} \quad (42)$$

All of the terms in (42) have been calculated previously in (4), (23), and (39), except the second partial derivatives of acceleration, which can be expressed as

$$\begin{aligned} \frac{\partial^2 a_i}{\partial p_j \partial p_k} &= \begin{bmatrix} \frac{\partial^2 J_{i1}}{\partial p_j \partial p_k} & \frac{\partial^2 J_{i2}}{\partial p_j \partial p_k} & \frac{\partial^2 J_{i3}}{\partial p_j \partial p_k} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \\ &+ \begin{bmatrix} \frac{\partial J_{i1}}{\partial p_k} & \frac{\partial J_{i2}}{\partial p_k} & \frac{\partial J_{i3}}{\partial p_k} \end{bmatrix} \begin{bmatrix} J_{1j} \\ J_{2j} \\ J_{3j} \end{bmatrix} \\ &+ \begin{bmatrix} \frac{\partial J_{i1}}{\partial p_k} & \frac{\partial J_{i2}}{\partial p_k} & \frac{\partial J_{i3}}{\partial p_k} \end{bmatrix} \begin{bmatrix} J_{1j} \\ J_{2j} \\ J_{3j} \end{bmatrix} \\ &+ \begin{bmatrix} J_{i1} & J_{i2} & J_{i3} \end{bmatrix} \begin{bmatrix} \frac{\partial J_{1j}}{\partial p_k} \\ \frac{\partial J_{2j}}{\partial p_k} \\ \frac{\partial J_{3j}}{\partial p_k} \end{bmatrix}. \end{aligned} \quad (43)$$

Again, each of these terms is known except the second partial derivatives of the elements of the Jacobian matrix, which are

written as

$$\frac{\partial^2 J_{ij}}{\partial p_k \partial p_l} = \begin{bmatrix} \frac{\partial^2 c_1}{\partial p_j \partial p_k \partial p_l} & \cdots & \frac{\partial^2 c_8}{\partial p_j \partial p_k \partial p_l} \end{bmatrix} \begin{bmatrix} h_{1i} \\ \vdots \\ h_{8i} \end{bmatrix}. \quad (44)$$

Thus all elements are known, and acceleration, jerk and snap can be expressed as functions of position, velocity, and partial derivatives of the coefficients calculated in (21). Analytical computation in these forms allows for efficient online computation of the parameters needed for the vector field based controller used in the experiments.


References

- Aydin Gol, E., & Belta, C. (2013). Time-constrained temporal logic control of multi-affine systems. *Nonlinear Analysis: Hybrid Systems*, 10, 21–33.
- Baier, C., & Katoen, J. (2008). *Principles of model checking*. Cambridge: MIT Press.
- Belta, C., & Habelts, L. C. G. J. M. (2006). Controlling a class of nonlinear systems on rectangles. *IEEE Transactions on Automatic Control*, 51(11), 1749–1759.
- Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6(1), 80–91.
- Jha, S., Clarke, E., Langmead, C., Legay, A., Platzer, A., & Zuliani, P. (2009). A Bayesian approach to model checking biological systems. In *Proceedings of the 7th International Conference on Computational Methods in Systems Biology, CMSB '09* (pp. 218–234). Berlin: Springer.
- Karaman, S., & Frazzoli, E. (2008). Vehicle routing problem with metric temporal logic specifications. In *IEEE conference on decision and control* (pp. 3953–3958).
- Kupferman, O., & Vardi, M. (2001). Model checking of safety properties. *Formal Methods in System Design*, 19(3), 291–314.
- Latvala, T. (2003). Efficient model checking of safety properties. In *10th international SPIN workshop, model checking software* (pp. 74–88). Springer.
- Leahy, K., Zhou, D., Vasile, C., Oikonomopoulos, K., Schwager, M., & Belta, C. (2014). Provably correct persistent surveillance for unmanned aerial vehicles subject to charging constraints. In *Proceedings of the international symposium on experimental robotics (ISER)*.
- Lee, T., Leoky, M., & McClamroch, N. (2010). Geometric tracking control of a quadrotor UAV on SE(3). In *49th IEEE conference on decision and control (CDC), 2010* (pp. 5420–5425). IEEE.
- Mellinger, D. & Kumar, V. (2011). Minimum snap trajectory generation and control for quadrotors. In *IEEE international conference on robotics and automation (ICRA), 2011* (pp. 2520–2525). IEEE.
- Michael, N., Stump, E., & Mohta, K. (2011). Persistent surveillance with a team of MAVs. In *Proceedings of the international conference on intelligent robots and systems (IROS 11)* (pp. 2708–2714). IEEE.
- Mulgaonkar, Y. & Kumar, V. (2014). Autonomous charging to enable long-endurance missions for small aerial robots. In *Proceedings of SPIE-DSS* (p. 90831S).
- Ozay, N., Topcu, U., & Murray, R. M. (2011). Distributed power allocation for vehicle management systems. In *50th IEEE conference on decision and control and European control conference (CDC-ECC), 2011* (pp. 4841–4848). IEEE.

- Smith, S., Tumova, J., Belta, C., & Rus, D. (2011). Optimal path planning for surveillance with temporal logic constraints. *International Journal of Robotics Research*, 30(14), 1695–1708.
- Stump, E. & Michael, N. (2011). Multi-robot persistent surveillance planning as a vehicle routing problem. In *Proceedings of the IEEE conference on automation science and engineering (CASE)* (pp. 569–575). IEEE.
- Sundar, K., & Rathinam, S. (2014). Algorithms for routing an unmanned aerial vehicle in the presence of refueling depots. *IEEE Transactions on Automation Science and Engineering*, 11(1), 287–294.
- Tkachev, I. & Abate, A. (2013). Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems. In *Proceedings of the 16th international conference on hybrid systems: Computation and control* (pp. 283–292). Philadelphia, PA.
- Toth, P., & Vigo, D. (2001). *The vehicle routing problem*. Philadelphia: SIAM.
- Ulusoy, A., Smith, S. L., Ding, X. C., Belta, C., & Rus, D. (2013). Optimality and robustness in multi-robot path planning with temporal logic constraints. *International Journal of Robotics Research*, 32(8), 889–911.
- Vasile, C. & Belta, C. (2014). An automata-theoretic approach to the vehicle routing problem. In *Robotics: science and systems conference (RSS)*. Berkeley, CA, USA.
- Wongpiromsarn, T., Topcu, U., & Murray, R. M. (2009). Receding horizon temporal logic planning for dynamical systems. *Conference on Decision and Control (CDC), 2009*, 5997–6004.
- Zhou, D. & Schwager, M. (2014). Vector field following for quadrotors using differential flatness. In *Proceedings of the international conference on robotics and automation (ICRA)* (pp. 6567–6572).



Kevin Leahy is a PhD student in the Department of Mechanical Engineering at Boston University in the Hybrid and Networked Systems (HyNeSs) and Multi-Robot Systems Laboratories. He received his B.A. degree in Economics from Boston University in 2009. His research interests include formal methods, persistent surveillance, and information based planning for teams of robots.



Dingjiang Zhou received the BS and MS degrees from the department of mechanical engineering, Harbin Institute of Technology, Harbin, China, in 2009, and 2011, respectively. He is currently working toward his Ph.D. degree at the Multi-robot Systems Laboratory, the department of mechanical engineering, Boston University, Boston. His research interests include quadrotor dynamics, control, agile maneuvering, trajectory planning, formation flight, on-board autonomous flight systems and embedded systems.



Cristian-Ioan Vasile is a PhD student in the Division of Systems Engineering at Boston University and the Hybrid and Networked Systems (HyNeSs) laboratory. He obtained his BS degree in Computer Science in 2009 and his MEng in Intelligent Control Systems in 2011 both from the Faculty of Automatic Control and Computers, Politehnica University of Bucharest. His research interests include formal methods, motion and path planning, distributed and decentralized control with applications to robotics, networked systems and systems biology. He is a student member of the IEEE.



Konstantinos Oikonomopoulos obtained his Bachelor's Degree in Mechanical Engineering from Boston University in 2014. While an undergraduate, Konstantinos participated in BU's SURF (2012) and UROP (2012) undergraduate research programs, where he developed Micro Aerial Vehicles as well as relevant infrastructure for the HyNeSs and IML labs. In 2013, he was awarded the Dean Lutchen Fellowship to participate in research on autonomously charging MAVs under Professor Calin Belta. Konstantinos is currently a Mechanical Engineer working for Formlabs Inc.



Mac Schwager is an assistant professor in the Department of Mechanical Engineering and the Division of Systems Engineering at Boston University. He obtained his BS degree in 2000 from Stanford University, his MS degree from MIT in 2005, and his PhD degree from MIT in 2009. He was a postdoctoral researcher working jointly in the GRASP lab at the University of Pennsylvania and CSAIL at MIT from 2010 to 2012. His research interests are in distributed algorithms for control, perception, and learning in groups of robots and animals. He received the NSF CAREER award in 2014.



Calin Belta is a Professor in the Department of Mechanical Engineering, Department of Electrical and Computer Engineering, and the Division of Systems Engineering at Boston University, where he is also affiliated with the Center for Information and Systems Engineering (CISE) and the Bioinformatics Program. His research focuses on dynamics and control theory, with particular emphasis on hybrid and cyber-physical systems, formal synthesis and verification, and

applications in robotics and systems biology. Calin Belta is a Senior Member of the IEEE and an Associate Editor for the SIAM Journal on Control and Optimization (SICON) and the IEEE Transactions on Automatic Control. He received the Air Force Office of Scientific Research Young Investigator Award and the National Science Foundation CAREER Award.