

Network integrity via coordinated motion of stratospheric vehicles

Michael Ouimet Jorge Cortés Sonia Martínez

Abstract—This paper considers the task of deploying mobile wireless repeaters on aerial vehicles in the upper air of Earth to increase the Internet connectivity of users with little to no network infrastructure. By routing data between a user and a strong connection to the Internet through the vehicle network, this ad hoc infrastructure can provide Internet access to those currently without it. We consider two different types of vehicles, altitude-actuating balloons and fixed-altitude gliders, and examine the task of finding optimal vehicle trajectories under lateral wind dynamics employing a throughput-based performance function. Given the complexity of computing the optimal trajectories of all vehicles simultaneously, we introduce an approximation that allows us to determine the optimal trajectory of a single vehicle. We use this optimal trajectory as a reference for the network and develop a coordination algorithm that makes all vehicles follow it while being spaced out equally in time, providing good overall coverage to users. For gliders, we show that, if cost of control is small enough, these vehicles converge to a stationary formation rather than flow with the wind. Simulations validate the throughput performance of the proposed periodic trajectories.

I. INTRODUCTION

As of December 31, 2013 only 40% of the world’s population has access to the Internet. Many others are connected but with very slow speeds. To combat this, Google unveiled its Loon project in 2013, which aims to bring the Internet to areas of the world currently without it through an ad hoc network of high altitude balloons equipped with wireless antennae and receivers. The basic idea is for a user in a remote location to wirelessly send/receive data to/from a balloon above them using a specialized antenna. Then, the network of balloons route the information between themselves until the data can be sent to/received from a fast connection to the Internet; Figure 1 depicts this scenario. High-altitude weather balloons are being proposed for this project because they are cheap and reliable. However, their simplicity brings about interesting control challenges: the balloons can change their altitude by controlling their internal pressure using a pump, but their latitudinal and longitudinal velocity is governed by the airspeed at their current altitude. Since these balloons are underactuated, one cannot just determine the set of optimal static locations that provide the best service to the users. Instead one must consider the performance of trajectories.

In April 2014, Google further acquired Titan Aerospace, a company that produces high-altitude, solar-powered drones, to aid the balloons in providing Internet coverage. In contrast to the balloons, these drones are capable of counteracting the latitudinal and longitudinal winds at the cost of actuation. Motivated by this, we also consider finding the trajectories

that optimize the average coordinated performance of balloons and drones over one revolution around the Earth to provide better network services to users.




Fig. 1. Central dark circle represents a latitude of Earth, looking down the axis of rotation. There are users and infrastructure with large distances in between. The balloons/gliders, depicted as gray circles, act as wireless repeaters to bring the signal between users and infrastructure. The balloons are able to change their altitude within an upper and lower limit. The wind’s direction changes with altitude and location, so balloons actively control their altitude to move in good directions. Similarly, the drones may counteract the winds at the expense of actuation. In both cases, vehicles seek to move in the direction which provides the best user service.

Literature review: Current work is being devoted to understanding how a repeater’s motion can enhance a wireless network’s performance. In [1], the authors consider the task of mobile robots moving through an area to complete a generic task while respecting each agent’s desire to have a minimum communication rate out of their network to fixed access points. Recent experiments have quantified the effect of actively moving wireless infrastructure (WiFi routers) to increase the network throughput for home and office environments. [2]. Furthermore, the idea of high-altitude platforms (HAPs) to improve wireless communications has been explored for at least 20 years, see e.g., [3], [4], [5], [6], as their use has the benefit of less lag and more bandwidth over satellites, and a larger range of coverage over a terrestrial antenna. These works use the Shannon-Hartley equation [7] to model the wireless capacity of a link where each type of connection (HAP-to-HAP, HAP-to-user, and HAP-to-ground station) operates at different bandwidths. In contrast with balloons, a HAP is typically an aircraft such as a blimp or plane that controls itself to remain in an approximately fixed position. This allows for simpler guarantees on network performance than when using balloons. While balloons share the benefits of traditional HAPs, their cost is much lower. The work [8] solves the problem of trajectory planning for a single hot-air balloon in a linear wind field using optimal control-based techniques.

Research supported by NSF award CMMI-1300272 and L3 Communications through the UCSD Cymer Center for Control Systems and Dynamics. The authors are with the Department of Mechanical and Aerospace Engineering, University of California, San Diego, CA 92093, USA, {mouimet, cortes, soniamd}@ucsd.edu

Finally, for balloons [9] and UAVs [10], work has been done finding feasible trajectories for motion planning in the presence of uncertain wind fields. The control techniques employed in this paper for the coordination of vehicles employ Laplacian agreement [11] and backstepping [12] to make the vehicles follow the desired optimal trajectories. Our coordinated trajectories are related to cyclic pursuit and present an alternative to [13], which result into appropriate vehicle spacing for sampling spatial phenomena. The latter work consider planar vehicles with steering control and fixed velocity and proposes a solution to reducing the high-dimensional functional space based on considering the optimization over a family of parameterized trajectories.

Statement of contributions: We propose a solution to a problem in global network integrity by coordinating the motions of high-altitude platforms acting as an ad hoc network of wireless communication repeaters. To measure the network quality provided by the vehicles, we propose a simplified throughput-based performance metric that allows us to find the routing policy that maximizes the flow between users and infrastructure for any static configuration of vehicles. Using the assumptions that communication between the wired infrastructure of the Internet is fast and less costly and that all bandwidth allocations are constant, the problem is efficiently routing the information from a user to any infrastructure node (and back). This allows us to formulate the capacity of the network as the optimizer of a linear program. Because the vehicles are constantly moving, we seek to optimize the average capacity over one revolution around the Earth. Since the associated optimal control problem is intractable for large numbers of vehicles, we derive an approximate performance metric and determine an optimal periodic trajectory for a single vehicle. Our design strategy consists of synthesizing a distributed algorithm for all vehicles to converge to this same optimal trajectory while being spaced out equally in time. Simulations show that this periodic formation leads to good network throughput. For the glider case, we provide a bound on the critical actuation cost for which optimal trajectories result into static deployments.

II. PRELIMINARIES

This section contains notions on graph theory, a wind model, and the vehicle dynamic models. \mathbb{R} is the set of real numbers, $\mathbb{Z}_{\geq 1}$ the positive integers, and \mathbb{S} the space of angles parameterizing a circle. For any $d \in \mathbb{Z}_{\geq 1}$, $\mathbf{1}_d$ and $\mathbf{0}_d$ are the d -dimensional vectors of all ones and zeros, respectively.

A. Graph theory

We review basic notions on graph theory following [14]. A weighted digraph of order n is a triplet $G = (V, E, A)$, where V is a set of n elements called the vertices (or nodes) and E is a set of ordered pairs of vertices called edges, where $E \subseteq V \times V$. For $u, v \in V$, the ordered pair (u, v) denotes an edge from u to v . The nonnegative matrix $A \in \mathbb{R}_{>0}^{n \times n}$ has the following property: for $i, j \in \{1, \dots, n\}$, the entry $a_{ij} > 0$ if (v_i, v_j) is an edge of G , and $a_{ij} = 0$ otherwise. The sources of a weighted digraph G are the nodes with no incoming edges. Similarly, the sinks of G are the nodes that contain no outgoing edges. An undirected graph is one where

an edge from node u to v implies an edge from v to u and an unweighted graph has all positive $a_{ij} = 1$. We refer to an unweighted, undirected graph as simply a graph. A path in weighted digraph G is an ordered sequence of vertices such that any pair of consecutive vertices in the sequence is an edge of the graph. A graph is connected if there exists a path between any two vertices. The Laplacian matrix $L(G) = (\ell_{i,j})_{n \times n}$ associated with an undirected unweighted graph G has off-diagonal elements defined by

$$\ell_{i,j} = \begin{cases} -1, & \text{if } i \neq j \text{ and } (i, j) \in E, \\ 0, & \text{otherwise,} \end{cases}$$

and positive diagonal elements defined so that $L(G)\mathbf{1}_n = \mathbf{0}_n$. This matrix is symmetric and positive semidefinite, with the multiplicity of 0 corresponding to the number of connected components of G when G is undirected. In the ring graph, each vertex is connected to the vertices immediately on either side of it. The associated Laplacian matrix is

$$L = M^T M, \quad M = \begin{bmatrix} 1 & 0 & 0 & \dots & -1 \\ -1 & 1 & 0 & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & -1 & 1 & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}.$$

B. Global wind patterns

The Earth's winds are a chaotic system that is difficult to model over long timescales. However, historical and current weather balloon data [15] allows one to approximate the state of the wind fields $w(t, \theta, \phi, h)$ as a function of time, latitude (θ), longitude (ϕ), and altitude (h). We denote the components of the windspeed along the corresponding directions by w_θ , w_ϕ , and w_h , respectively, and assume they are continuously differentiable. At latitudes close to the equator, winds can be approximated as moving completely east-to-west or west-to-east depending on the time of year. Additionally, in general, the airspeed increases with altitude and is constant over time [16]. Thus, for the rest of the paper, we model $w_\theta = Kh + C_\theta$, with $K > 0$, and $w_\phi = 0$.

C. Balloon dynamics

Consider an altitude-actuating balloon in the stratosphere. Balloons float at an altitude that they are neutrally buoyant. In the stratosphere, the air pressure and temperature monotonically decreases with height. Thus, balloons are capable of changing their altitude by changing their internal pressure via a pump. Considering this, we assume that each balloon may directly control its rate of altitude change. However, having no actuation in the lateral directions, a balloon's motion in these directions is governed by the winds in the stratosphere. Because the balloon's mass is small, we assume that it is Lagrangian, i.e., its lateral velocity is equal to the wind velocity at that location. An accurate modeling of the dynamics of high-altitude balloons is involved [17], [18]. In this paper, we abstract away many of the physical details for the sake of tractability. The position of a balloon i is $x_i^B = (\theta_i^B, h_i) \in \mathbb{S} \times \mathbb{R}$ and its dynamics given by

$$\dot{\theta}_i^B = w_\theta(\theta_i^B, h_i), \quad \dot{h}_i = u_i^B. \quad (1)$$

We assume that the vertical wind and solar radiation-induced forcing are already canceled by the control authority and the remainder is u_i^B . To ensure the safety of air traffic, the vehicles have a lower bound on the altitude $h_{\min} > 0$ that it can drift at as well as an upper bound $h_{\max} > 0$ imposed by physics. We also assume that the balloon's control u_i^B is bounded by a value $u_{\max} > 0$ for all time.

D. Glider dynamics

Here, we consider a simple model for a fixed-wing aircraft i with limited propulsion control:

$$\dot{\theta}_i^G = w_\theta(\theta_i^G, h_i) + u_i^G, \quad \dot{h}_i = 0. \quad (2)$$

The glider maintains a fixed altitude and is able to change its speed relative to the wind by exerting control effort. We also assume that the magnitude of the glider's control u_i is uniformly bounded by a value $u_{\max} > 0$.

III. PROBLEM STATEMENT

We assume there are N_I infrastructure locations, N_U users, and N_V vehicles (balloons or gliders). Infrastructure locations have antennae with large capacity wireless channels that can communicate with the vehicles and route the users' Internet data to and from them. The static location of the infrastructure $i \in \mathcal{N}_I = \{1, \dots, N_I\}$ is $x_i^I = \theta_i^I \in \mathbb{S}$. Here, we model groups of individual consumers together as one user. The position of user $i \in \mathcal{N}_U = \{1, \dots, N_U\}$ is $x_i^U = \theta_i^U \in \mathbb{S}$. For simplicity, we neglect the elevation of users and infrastructure because their variation is assumed small compared to the altitude of the vehicles. Dependent on the number of people and their preferences, each user i has a desired capacity preference κ_i that it would like to receive. Finally, the position of vehicle $i \in \mathcal{N}_V = \{1, \dots, N_V\}$ is $x_i^V = (\theta_i^V, h_i) \in \mathbb{S} \times \mathbb{R}$. Then, the network dynamics is

$$\dot{x}^V = g(x^V, u), \quad (3)$$

where g_i is (1) or (2) depending on i 's vehicle type.

We consider the objective of providing optimal network capacity to the users using the vehicles to connect them to the infrastructure. Because the vehicles are constantly moving and this motion affects the network's throughput, the goal is to determine a control strategy maximizes the achievable average network capacity over one revolution around the Earth while minimizing control effort. Given an initial vehicle configuration $x^V(0)$, let $\mathcal{U}(x^V(0))$ be the set of admissible control trajectories that result in trajectories of (3) that respect the height constraints, i.e.,

$$\mathcal{U}(x^V(0)) = \{u : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^{N_V} \mid u(t) \in [-u_{\max}, u_{\max}]^{N_V}, \\ x^V(t) \in [h_{\min}, h_{\max}]^{N_V}, \forall t \geq 0\}.$$

Given a function $R : (\mathbb{S} \times \mathbb{R})^{N_V} \times (\mathbb{S})^{N_I} \times (\mathbb{S})^{N_U} \rightarrow \mathbb{R}$ that captures the quality of service for a static set of vehicle positions and a function $G : \mathbb{R}^{N_V} \rightarrow \mathbb{R}$ that penalizes energy consumption, we wish to find the optimal trajectories generated by the control inputs that solve

$$J(x^V(0); T; x^I; x^U) = \\ \max_{u \in \mathcal{U}(x^V(0))} \frac{1}{T} \int_0^T (R(x^V(t); x^I; x^U) + G(u(t))) dt. \quad (4)$$

T is a design parameter which can correspond to, for instance, the maximum time that it would take a vehicle to

circle the Earth. Later, we will take it to be a variable that we optimize over, so that T is exactly the time it takes the network of balloons to make one complete revolution with the given control strategy.

IV. NETWORK PERFORMANCE MODEL

Now we define a quality-of-service function R for (4).

A. Wireless communication model for channel capacity

We introduce here the wireless communication model for the channel capacity between two agents (users, vehicles, and infrastructure nodes) i and j located at positions x_i and x_j . The Shannon-Hartley theorem [7] provides an upper bound on the channel capacity C_{ij} (bits per second) for point-to-point line-of-sight radio communication from i to j ,

$$C_{ij} = B_{ij} \log \left(1 + \frac{P_{ij}}{N_0 B_{ij}} \right), \quad (5)$$

where B_{ij} is the allowable bandwidth, P_{ij} is the received power, and N_0 is the noise power per unit bandwidth. Here, we assume a fixed transmitted power P and free-space path loss [19] leading to the power of the received signal $P_{ij} \propto \frac{P}{\|x_i - x_j\|^2}$. In principle, given a network of agents communicating wirelessly, bandwidth allocation policies may allow for time-varying bandwidth allocations. However, we make the simplifying assumption that the bandwidth allocations are fixed and equal for all channels of the same type (vehicle-to-vehicle, user-to-vehicle, and vehicle-to-infrastructure).

B. Linear program formulation of network capacity

We now define the model to measure the capacity provided by a static vehicle configuration to users. Because routing information is cheaper and faster amongst the wired infrastructure, we simplify the capacity problem. We consider the goal of the vehicle network to route the most capacity from each user to any infrastructure node. We view the network as a graph with the vehicles, users, and infrastructure as nodes. Two nodes are connected with an edge if they can communicate wirelessly with each other and the edge's weight equal is to their link capacity according to (5). Then, one natural measure for the capacity of given network configuration is a (modified) max flow through the network when viewing the users as sources and the infrastructure as sinks. To do this, we artificially set the allowable bandwidth from vehicles to users and from infrastructure nodes to vehicles to 0. Then, the maximum flow between the users and the infrastructure nodes can be formally written as the optimal solution of a linear program, as described next.

Let $F^{VV} \in \mathbb{R}_{\geq 0}^{N_V \times N_V}$, $F^{VI} \in \mathbb{R}_{\geq 0}^{N_V \times N_I}$, and $F^{UV} \in \mathbb{R}_{\geq 0}^{N_U \times N_V}$ represent the information flow between vehicles, between vehicles and infrastructure nodes, and users and vehicles, respectively. The ij th entry is the amount of flow from i to j . Similarly, define $C^{VV} \in \mathbb{R}_{\geq 0}^{N_V \times N_V}$, $C^{VI} \in \mathbb{R}_{\geq 0}^{N_V \times N_I}$, and $C^{UV} \in \mathbb{R}_{\geq 0}^{N_U \times N_V}$ as the channel capacities of the links as in (5). Then, we define the following constraints

$$F_{ij}^{VV} + F_{ji}^{VV} \leq C_{ij}^{VV}, \quad \forall i, j \in \mathcal{N}_V \quad (6a)$$

$$F_{ij}^{VI} \leq C_{ij}^{VI}, \quad \forall i \in \mathcal{N}_V, j \in \mathcal{N}_I \quad (6b)$$

$$F_{ij}^{UV} \leq C_{ij}^{UV}, \quad \forall i \in \mathcal{N}_U, j \in \mathcal{N}_V \quad (6c)$$

$$F^{VV} \geq 0, \quad F^{UV} \geq 0, \quad F^{VI} \geq 0, \quad (6d)$$

stating that the information flow across each link must be non-negative and bounded by the link capacity. and vehicle-to-vehicle channels are shared between each direction of communication. No information can be saved at a vehicle:

$$((F^{VV} - (F^{VV})^T) \mathbf{1}_{N_V} + F^{VI} \mathbf{1}_{N_I} - (F^{UV})^T \mathbf{1}_{N_U}) = \mathbf{0}_{N_V}. \quad (7)$$

The maximum flow from users to infrastructure nodes is

$$R_{\text{wmf}}(x^V, x^I, x^U) = \min_{F^{VV}, F^{VI}, F^{UV}} -d^T F^{UV} \mathbf{1}_{N_V}, \quad (8)$$

subject to (6) and (7). Note that this function sums the outgoing information from the users to all the vehicles and weights them by $d \in \mathbb{R}_{\geq 1}^{N_U}$. For a feasible set of flows F^{VV}, F^{UV}, F^{VI} , all information leaving every source reaches an infrastructure node, so summing all of the flows out of the sources is the network's maximum flow. Choosing some elements of d strictly greater than 1 gives them more importance, incentivizing the network to provide more capacity to those users. One should choose d as a function of the capacity preference κ . Alternatively, one could handle the capacity preferences directly with the constraint $F^{UV} \mathbf{1}_{N_V} \geq \kappa$. However, this may make the linear program infeasible.

C. Approximation of network capacity

We assume the limiting link between the users and the infrastructure positions are the uplinks and downlinks between the vehicles and the surface (users and infrastructure positions). For instance, the vehicles could efficiently communicate using free space optics, which has a much larger range and throughput than air-to-surface communications via radio waves. This approximation becomes better the more vehicles there are. Hence, we only consider the links between vehicles and ground features (users and infrastructure locations) and approximate the capacity in (8) by

$$\hat{R}_{\text{wmf}}(x^V, x^I, x^U) = \frac{1}{2} (\mathbf{1}_{N_V}^T C^{VI} \mathbf{1}_{N_I} + \mathbf{1}_{N_U}^T C^{UV} \mathbf{1}_{N_V}). \quad (9)$$

V. DYNAMIC DEPLOYMENT

In this section, we describe our solution to the problem stated in Section III. Our discussion consists of the following parts: we find an optimal periodic trajectory for a single vehicle, provide justification for a solution that spaces out all vehicles homogeneously along it, and finally introduce a coordination algorithm that achieves this objective.

A. Dynamic deployment on optimal single-vehicle trajectory

We employ the approximation (9) as the quality-of-service function for the single vehicle optimization problem:

$$\max_{x_i^V(0), u \in \mathcal{U}(x_i^V(0)), T} \frac{1}{T} \int_0^T \frac{1}{2} (C^{VI} \mathbf{1}_{N_I} + (\mathbf{1}_{N_U}^T C^{UV})^T)_i - G_i(u(t)) dt, \quad (10)$$

subject to the dynamics ((1) or (2)) and enforcing $x_i^V(0) = x_i^V(T)$ to generate a periodic reference trajectory. Since the dynamics are autonomous, maximizing (10) yields a

trajectory $g : [0, T] \rightarrow \mathbb{S} \times \mathbb{R}$. The vehicle's complete trajectory is then the periodic function where $g(t+T) = g(t)$. This trajectory, in turn, can be expressed as two coupled differential equations. Although θ_i and h_i are explicitly functions of time, they can be parameterized in terms of θ_i , yielding a family of goal trajectories satisfying

$$\dot{\theta}_i = f_\theta(\theta_i), \quad \dot{h}_i = f_h(\theta_i), \quad (11)$$

for i in $\{1, \dots, N\}$, where $(\theta_i(0), h_i(0))$ can be $g(\tau)$, for any $\tau \in [0, T]$. Choosing τ corresponds to the vehicle following the shifted trajectory where $(\theta_i(t), h_i(t)) = g(t + \tau)$.

The validity of (9) is based on reliable connections between the vehicles. Thus, the vehicles should not only follow the prescribed trajectory, but also be well spaced to ensure that inter-vehicle communication is not the limiting factor and justify that links through the vehicle network exist between a user and infrastructure, even though they are not explicitly modeled. To aid in this goal, we define the total travel time \mathcal{T}_{tot} to make one revolution while following the optimal reference trajectory as $\mathcal{T}_{\text{tot}} = \int_0^{2\pi} \frac{1}{f_\theta(\zeta)} d\zeta$. Now, we define the travel time $\mathcal{T} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ between vehicle i at longitude θ_i and vehicle $i+1$ at θ_{i+1} for a vehicle following the desired trajectory $\dot{\theta} = f_\theta(\theta)$ and starting at θ_i as

$$\mathcal{T}_i(\theta) = \int_{\theta_i}^{\theta_{i+1}} \frac{1}{f_\theta(\zeta)} d\zeta, \quad \forall i \in \{1, \dots, N-1\},$$

$$\mathcal{T}_N(\theta) = \mathcal{T}_{\text{tot}} - \int_{\theta_1}^{\theta_N} \frac{1}{f_\theta(\zeta)} d\zeta.$$

Note that this function only corresponds to the true travel time if $\theta_i < \theta_j$ and $\theta_1 < \theta_N$. Thus, the goal is for all of the vehicles to follow an optimal trajectory, i.e., follow the dynamics (11) while also equally spaced out: $\mathcal{T}_i(\theta) = \mathcal{T}_j(\theta)$, for all $i, j \in \{1, \dots, N\}$.

B. Distributed coordination under wind dynamics

Here we describe our distributed strategy to make the vehicles achieve the objective described in the previous section. We consider first the case of balloon dynamics. Our design first deals with a "virtual" control law that assumes angular velocity can be directly actuated (even though the balloons are actually controlled by changing their vertical velocity). Once this is in place, we design the altitude control law to converge to the virtual control law using backstepping.

For each $i \in \{1, \dots, N\}$, let the "virtual" dynamics be

$$\dot{\theta}_i = f_\theta(\theta_i) + v_i. \quad (12)$$

When $v_i = 0$, balloon i follows the goal trajectory in the θ -direction. Because the trajectory is periodic, the balloon can be following any trajectory in the family of trajectories that satisfy $\dot{\theta}_i = f_\theta(\theta_i)$. By choosing v_i non-zero for some time, the balloon changes the specific trajectory in this family for which it is following. Since the goal is for all of the balloons to follow a trajectory in this family while being equally spaced out in time, we construct a control law v_i that speeds up/slows down a balloon depending on the travel time between the balloons, $v_i = f_\theta(\theta_i)(\alpha(\mathcal{T}_i(\theta) - \mathcal{T}_{i-1}(\theta)))$, or

$$v = \alpha D(\theta) M \mathcal{T}, \quad (13)$$

more compactly, where $D(\theta)_{ii} = f_\theta(\theta_i)$ and is 0 otherwise and $\alpha \in \mathbb{R}_{>0}$ is a design parameter. From the structure of v ,

note that if all of the travel times are equal, $v = 0$, and so, the balloons simply follow the optimal reference trajectory.

Theorem 5.1: For a network of N agents with the dynamics of (12) and starting from any $\theta(0) \in [0, 2\pi)^N$ such that $\theta_i(0) < \theta_{i+1}(0)$ for all $i \in \{1, \dots, N-1\}$, the control law (13) causes the travel times of all agents to asymptotically converge to $\frac{T_{\text{tot}}}{N}$. Furthermore, each balloon asymptotically follows the desired trajectory.

We now consider rewriting (3) to capture the error between the actual θ -dynamics and the desired one:

$$\begin{aligned}\dot{\theta} &= D(\theta)(\mathbf{1}_n + \alpha MT(\theta)) + e, \\ \dot{e} &= Ku - \frac{d}{dt}[D(\theta)(\mathbf{1}_n + \alpha MT(\theta))],\end{aligned}$$

where $e = C_\theta \mathbf{1}_n + Kh - D(\theta)(\mathbf{1}_n + \alpha MT(\theta))$. We design the control law

$$u = \frac{1}{K} \frac{d}{dt}[D(\theta)(\mathbf{1}_n + \alpha MT(\theta))] + D^{-1}(\theta) M^T L T(\theta) - e \quad (15)$$

leading to closed-loop dynamics of

$$\dot{\theta} = D(\theta)(\mathbf{1}_n + \alpha MT(\theta)) + e, \quad \dot{e} = D^{-1}(\theta) M^T L T(\theta) - e. \quad (16)$$

Expressing the \mathcal{T} dynamics as in Theorem 5.1 leads to

$$\dot{\mathcal{T}} = -\alpha L T(\theta) - M^T D^{-1}(\theta) e, \quad \dot{e} = D^{-1}(\theta) M L T - e,$$

Note that the control law (15) is distributed. Each agent i requires only the state of its 2-hop neighbors, making the complexity of each agent's control law independent of the size of the network. The next result establishes that this controller makes the overall network achieve a uniform deployment in time over the desired trajectory.

Theorem 5.2: For any $\theta(0) \in [0, 2\pi)^N$ such that $\theta_i(0) < \theta_{i+1}(0)$ for all $i \in \{1, \dots, N-1\}$ and $h(0) \in \mathbb{R}_{>0}^N$, trajectories of the dynamics (16) asymptotically satisfy $\lim_{t \rightarrow \infty} \mathcal{T}(\theta_i(t), \theta_{i+1}(t)) = \frac{T_{\text{tot}}}{N}$, for all $i \in \{1, \dots, N\}$ and all balloons asymptotically follow the desired trajectory.

Given that the vehicles all converge to the same periodic trajectory and are equally spaced in time, this allows us to guarantee that a vehicle will be directly overhead every $\frac{T_{\text{tot}}}{N}$ seconds. By adding more vehicles, the time between vehicles decreases causing the average throughput to increase

Remark 5.3: (Glider dynamics analysis) The gliders may actuate in the θ -direction, see (2), and so, for a glider trajectory optimizing (10), Theorem 5.1 may directly be used to cause the gliders to spread out equally in time. •

VI. STATIC DEPLOYMENT FOR GLIDERS

The trajectory optimizing (10) depends on the relative importance of the benefit of providing service to the cost of motion. The next result states that for sufficiently small cost of actuation, the optimal choice for a group of gliders is to remain at a fixed location rather than maintain a periodic trajectory, regardless of the network performance function.

Lemma 6.1: Given a team of N gliders each following the dynamics of (2) with a network performance function $R : \mathbb{S}^N \rightarrow \mathbb{R}_{>0}$ and cost of actuation $G : \mathbb{R}^N \rightarrow \mathbb{R}_{>0}$, define an optimal trajectory by its initial location and the control strategy, i.e., $(x^V(0)^*, u^*) \in$

$$\operatorname{argmax}_{x^V(0), u \in \mathcal{U}(x^V(0))} \frac{1}{T} \int_0^T (R(x^V(t)) - \alpha G(u(t))) dt.$$

For α sufficiently small, an optimal initial location x^{V*} is a set of vehicle locations that optimize R , i.e., $x^{V*} \in$

$\operatorname{argmax}_{x^V \in [0, 2\pi)^{N_V}} R(x^V)$ and its corresponding optimal control strategy is $u(t) = -w_\theta(x^{V*})$, for all $t \in [0, T]$.

This result motivates the desire to find local optimizers of R_{wmf} . To this end, we examine gradient-based optimization of R_{wmf} . However, in general, R_{wmf} is not always differentiable. We note that if it is differentiable at x^V , then

$$\begin{aligned}\frac{\partial R_{\text{wmf}}}{\partial x_k^V}(x^V) &= - \sum_{ij} \lambda_{ij}^{VV}(x^V) \frac{\partial C_{ij}^{VV}}{\partial x_k^V}(x^V) \\ &\quad - \sum_{ij} \lambda_{ij}^{VI}(x^V) \frac{\partial C_{ij}^{VI}}{\partial x_k^V}(x^V) - \sum_{ij} \lambda_{ij}^{UV}(x^V) \frac{\partial C_{ij}^{UV}}{\partial x_k^V}(x^V),\end{aligned}$$

where λ_{ij}^{VV} is, for example, the optimal dual variable corresponding to the constraint $F_{ij}^{VV} + F_{ji}^{VV} \leq C_{ij}^{VV}$. This is a result of the fact that the optimal dual variables directly describe the sensitivity of the optimal value with respect to perturbations of the constraints, as seen in [20, (5.58)]. Thus, a network of gliders could implement $\dot{x}^V = \frac{dR_{\text{wmf}}}{dx^V}(x^V)$ to converge to a local optimizer of R_{wmf} .

VII. SIMULATIONS

We consider the simple but illustrative example of vehicles deploying on a ring, i.e., fixed latitude, where the latitudinal velocity increases with altitude. In this case, there is one user at 0 and one infrastructure located at π and the network of 20 balloons are constrained to remain between 25 and 30 km above the Earth's surface. The wind flowfield (in $\frac{\text{rad}}{\text{hr}}$) is given by $w_\theta = .0017h - .0388$. The bandwidth B between a vehicle and a ground feature is 100 and $\frac{P}{N_0 B} = 10$. Because vehicles can communicate more efficiently by laser optics, the bandwidth between vehicles is 1000. The cost of control is $G(u) = .1\|u\|^2$. Figure 2(a) shows the capacity that a balloon can deliver to/from users as a function of its latitude and altitude as given by (9). Note that we have also introduced a term that allows for no capacity when the balloon is more than $\frac{\pi}{10}$ radians away from a surface feature to account for its occlusion beyond the horizon. Similarly, vehicles can only send data between each other if they are less than $\frac{\pi}{5}$ radians apart. The plot highlights the intuitive fact that higher altitudes contribute to less capacity due to the increased distance for the approximate capacity function (9). Figure 2(b) depicts a locally optimal solution of (10) using (9) for one balloon. Note that the balloon is at a lower altitude (with better capacity and slower wind) while over the user or infrastructure and then moves to the higher altitude while in between. This results in a higher average capacity, as seen in Figure 2(c), which compares the average capacity for a balloon remaining at a fixed altitude (the maximum is around .2) to the average capacity of a balloon following the locally optimal periodic trajectory in Figure 2(b), which is .3.

Figure 3(b) shows the convergence of the travel time between each balloon and the one ahead of it to the same value. Figure 3(c) shows a balloon configuration after the dynamic deployment algorithm as in Theorem 5.2 has been run. The smallest circle is the Earth with a user at 0 and an infrastructure at π . Notice that more balloons are located over the user and infrastructure than above the open space in between them. This allows more balloons to be passing data to/from the user/infrastructure but with enough balloons




Fig. 2. (a) depicts the approximate capacity (9) that one balloon can provide to a user/infrastructure location while at fixed altitudes of 25 km and 30 km. Note that in this simplified scenario there is one user at 0 and one infrastructure located at π . (b) shows a locally optimal solution to (10). The balloon maintains a low, slow altitude while over surface features (near 0 and π as in (a)) and moves to a higher, fast altitude to quickly move to the next feature. (c) compares the average capacity of a balloon to the ground as a function of the fixed altitude that it maintains to the average capacity of the trajectory in (b). Note that the optimal trajectory outperforms all fixed altitude trajectories.

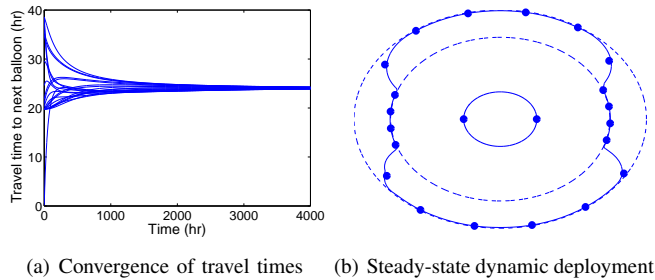


Fig. 3. (a) shows the convergence of the travel times to the next balloon to $\frac{T_{tot}}{N}$, which makes all balloons equally spaced in time. (b) depicts one instant in the steady-state dynamic deployment of the balloons after they have converged to being equally spaced in time as in Theorem 5.2. Balloons are more densely spaced above the user and infrastructure. Since the Earth's radius is larger than the balloons' range of motion, the figure is not to scale.

to also bridge the gap between the user and infrastructure. Because the Earth's radius is much larger than the balloons' range of altitude, the figure is not to scale.

VIII. CONCLUSIONS

We have proposed an algorithmic solution to a global network integrity problem that relies on coordinated motion control of high-altitude platforms acting as an ad hoc network of wireless communication repeaters. We proposed a measure of network quality based on a simplified throughput model. Because the vehicles are in constant motion, we considered the average performance over one period around the Earth. Since the associated optimal control was intractable for large numbers of vehicles, we derived an approximate metric that allowed us to determine an optimal periodic trajectory for one vehicle. Given this trajectory, we designed a distributed method for all vehicles to converge to it while being spaced out equally in time. Simulations show that this formation leads to good network throughput. For the case of gliders, we also provided a bound on the critical actuation cost for which the optimal trajectories result in static deployments. Future work will include investigating better numerical methods for determining the optimal periodic solutions of these high-dimensional, nonconvex optimal control problems, considering more general wind dynamics, and determining strategies for heterogeneous teams of both gliders and balloons.

REFERENCES

- [1] M. Zavlanos, A. Ribeiro, and G. Pappas, "Network integrity in mobile robotic networks," *IEEE Transactions on Automatic Control*, vol. 58, pp. 3–18, Jan. 2013.
- [2] M. Gowda, N. Roy, and R. R. Choudhury, "Infrastructure mobility: a what-if analysis," in *ACM HOTNETS*, Oct. 2014.
- [3] A. Mohammed, A. Mehmood, F. N. Pavlidou, and M. Mohorcic, "The role of high-altitude platforms in the global wireless connectivity," *Proceedings of the IEEE*, vol. 99, pp. 1939–1953, Nov. 2011.
- [4] G. M. Djuknic, J. Freidenfelds, and Y. Okunev, "Establishing wireless communications services via high-altitude aeronautical platforms: a concept whose time has come?," *IEEE Communications Magazine*, vol. 35, pp. 128–135, Sept. 1997.
- [5] T. C. Tozer and D. Grace, "High-altitude platforms for wireless communications," *Electronics & Communication Engineering Journal*, vol. 13, pp. 127–137, Mar. 2001.
- [6] S. Karapantazis and F. Pavlidou, "Broadband communications via high-altitude platforms: a survey," *IEEE Communications Surveys & Tutorials*, vol. 7, no. 1, pp. 2–31, 2005.
- [7] C. Shannon, "Communication in the presence of noise," *Proceedings of the IRE*, vol. 37, pp. 10–21, Jan. 1949.
- [8] T. Das, R. Mukherjee, and J. Cameron, "Optimal trajectory planning for hot-air balloons in linear wind field," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 26, pp. 416–424, June 2003.
- [9] M. T. Wolf, L. Blackmore, Y. Kuwata, N. Fathpour, A. Elfes, and C. Newman, "Probabilistic motion planning of balloons in strong, uncertain wind fields," in *IEEE Int. Conf. on Robotics and Automation*, (Anchorage, AK), pp. 1123–1229, May 2010.
- [10] W. Al-Sabban, L. F. Gonzalez, and R. N. Smith, "Wind-energy based path planning for electric unmanned aerial vehicles using Markov Decision Processes," in *IEEE Int. Conf. on Robotics and Automation*, (Karlsruhe, Germany), pp. 784–789, May 2012.
- [11] W. Ren and R. W. Beard, *Distributed Consensus in Multi-Vehicle Cooperative Control*. Communications and Control Engineering, Springer, 2008.
- [12] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 3 ed., 2002.
- [13] N. E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. Davis, "Collective motion, sensor networks and ocean sampling," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, 2007.
- [14] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*. Applied Mathematics Series, Princeton University Press, 2009. Electronically available at <http://coordinationbook.info>.
- [15] National Weather Service, "Environmental modeling center." <http://www.emc.ncep.noaa.gov/>, 2015.
- [16] C. Beccario, "Global wind visualization." <http://earth.nullschool.net>.
- [17] Q. Dai, X. Fang, X. Li, and L. Tian, "Performance simulation of high altitude scientific balloons," *Advances in Space Research*, vol. 49, pp. 1045–1052, Mar. 2012.
- [18] L. A. Carlson and W. J. Horn, "New thermal and trajectory model for high-altitude balloons," *Journal of Aircraft*, vol. 20, no. 6, pp. 500–507, 1983.
- [19] P. Santi, "Topology control in wireless ad hoc and sensor networks," *ACM Computing Surveys*, vol. 37, pp. 164–194, 2005.
- [20] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2009.