Safe Open-Loop Strategies for Handling Intermittent Communications in Multi-Robot Systems*

Siddharth Mayya[1] and Magnus Egerstedt[2]

***Abstract*— In multi-robot systems where a central decision maker is specifying the movement of each individual robot, a communication failure can severely impair the performance of the system. This paper develops a motion strategy that allows robots to safely handle critical communication failures for such multi-robot architectures. For each robot, the proposed algorithm computes a time horizon over which collisions with other robots are guaranteed not to occur. These safe time horizons are included in the commands being transmitted to the individual robots. In the event of a communication failure, the robots execute the last received velocity commands for the corresponding safe time horizons leading to a provably safe open-loop motion strategy. The resulting algorithm is computationally effective and is agnostic to the task that the robots are performing. The efficacy of the strategy is verified in simulation as well as on a team of differential-drive mobile robots.**

I. Introduction

Communication in multi-robot teams is not only essential for sharing sensor measurements and performing diagnostics, it is often an integral part of the closed loop control mechanism [1], [15]. In fact, in many applications and scenarios such as extra-terrestrial exploration [18], high precision manufacturing [2], and multi-robot testbeds [12], the robots frequently rely on communicating with a centralized decision maker for their velocity or position commands. In such situations, a failure in the communication network can severely hinder the motion of the robots and performance of the algorithm. This is the premise behind the work performed in this paper, whereby the adverse effects caused by intermittently failing communication networks are mitigated.

The inevitability of occasional failures in wireless communication channels [13] raises the following question: *What should a robot do in case a communication failure prevents it from receiving critical motion commands from a central decision maker?* Many different techniques have been explored to handle intermittent communications in multi-robot teams (see for e.g., [4], [11], [16], [17]). In many cases, the robots are assumed to have significant decision making capabilities, have sensors to maneuver around obstacles, or have knowledge about the positions of other robots. Furthermore, some developed communication recovery techniques do not provide formal collision-avoidance guarantees in case of unforeseen communication failures.

An existing technique used to handle communication failures, mentioned in [16], is to stop the robot when critical data is not received. While this behavior preserves safety, it could cause the robot to behave erratically. For example, if only intermittent velocity commands are received, the robot could move in a jerky "start-stop" fashion. This problem was observed on the Robotarium: a remote-access multi-robot testbed being developed at Georgia Tech [12]. During experiments, it was observed that temporary failures in the communication channels prevented robots from receiving velocity or position commands which caused them to abruptly stop moving, leading to a disruption in the coordination algorithm being executed.

Motivated by the need to alleviate such problems in general, and resolve issues with the Robotarium in particular, this paper proposes a strategy that allows differential-drive robots without sensory or decision-making capabilities, to continue moving safely for a specific amount of time even when velocity commands from a central decision maker are not received. For each robot, the central decision maker computes a time horizon over which collisions with other robots are guaranteed not to occur. This is called the safe time horizon. During normal operations, the desired velocity and the safe time horizon are transmitted to the robots periodically. If a robot stops receiving data due to a communication failure, it executes the last received velocity command for the duration of the last received safe time horizon. This allows the robot to continue moving in an open-loop yet provably collision-free manner despite having no updated information about the environment. Beyond the safe time horizon, it stops moving. By allowing the robots to move safely even during a communication failure, the algorithm is able to avoid jerky motions in robots during short-duration communication failures.

In order to calculate the safe time horizon, we first compute the set of all possible locations that can be reached by a robot within a given time (i.e., the reachable set [3], [14]). While reachability based techniques have been used to obtain connectivity guarantees in multi-robot teams, e.g., [5], in this paper, reachability analysis is used to ensure safety while allowing the robots to move in an open-loop manner. For each robot, we compute a time horizon for which it lies outside the reachable sets of other robots. Owing to the computational complexity of performing set-membership tests on the non-convex reachable sets of differential-drive robots, we over-approximate the reachable set by enclosing

*This research was sponsored by Grant No. 1544332 from the U.S. National Science Foundation

[1]Siddharth Mayya is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA siddharth.mayya@gatech.edu

[2]Magnus Egerstedt is a professor at the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA magnus@gatech.edu

it within an ellipse whose convex and simp[illegible] allows for simpler set-membership tests and finite representation [9]. By minimizing the area of the ellipse enclosing the convex hull of the reachable set, we obtain the best ellipsoidal over-approximation of the reachable set in terms of the accuracy and effectiveness of set-membership tests.

The outline of this paper is as follows: Section II derives an ellipsoidal approximation of the reachable set of a differential-drive robot. Section III formally defines the safe time horizon, outlines the algorithm used by the robots, and proves the safety guarantees that it provides. In Section IV, the dev[illegible]mplemented in simulation, followi[illegible]rification is performed on a team [illegible] robots. Finally, Section V conclud[illegible]

this figure is a bit ugly.

Here point out that a version involving all the proofs is available on the ArXiv in [10]

[illegible]Y ANALYSIS

This section introduces the notion of a reachable set of a robot, and derives an ellipsoidal approximation for it. Since the Robotarium [12], in its current form, is populated with differential-drive mobile robots, this paper investigates the reachability of two-wheeled differential-drive robots with non-holonomic dynamics.

For such systems, let $z = (x, y) \in \mathbb{R}^2$ denote the position of the robot in the 2D plane, and let $\phi \in [-\pi, \pi]$ denote its orientation with respect to the horizontal axis. As such, the robot is described as a point (z, ϕ) in the configuration space $\mathbb{R}^2 \times S^1$. The motion of the robots can be captured using the unicycle dynamics model:

$$\dot{x} = v\cos(\phi),\ \dot{y} = v\sin(\phi),\ \dot{\phi} = \omega, \qquad (1)$$
$$|v| \leq 1,\ |\omega| \leq 1.$$

The bounds on the linear velocity v and the angular velocity ω represent the physical limitations of the robots, and are normalized to 1 without loss of generality.

Since collisions are ultimately defined by posi[illegible] than orientations, we consider the output of the [illegible] be z. Relative to this output, the reachable set, denoted as $\mathcal{R}(t; z_0, \phi_0)$ is the set of all positions in the 2D plane that can be reached at time t by a robot starting in the configuration (z_0, ϕ_0) at $t = 0$. For $z_0 = (0, 0), \phi_0 = 0$, denote the reachable set as $\mathcal{R}(t)$. Since the dynamics in (1) is drift-free, the structure of the reachable set does not depend on z_0 and ϕ_0, i.e.,

$$\mathcal{R}(t; z_0, \phi_0) = z_0 + \Pi_{\phi_0}\mathcal{R}(t),$$

where Π_{ϕ_0} is a rotation by ϕ_0. Therefore, the study of reachable sets can be restricted to the case when $z_0 = (0, 0)$ and $\phi_0 = 0$. Fig. 1 portrays the reachable set for a robot with dynamics given in (1).

As discussed in the introduction, the complexity of performing set-membership tests with respect to the non-convex set $\mathcal{R}(t)$ increases the computational burden associated with the safe time horizon algorithm. Consequently, we compute an ellipsoidal approximation of the reachable set, which is not only convex, but allows for easy set-membership tests and finite representation.

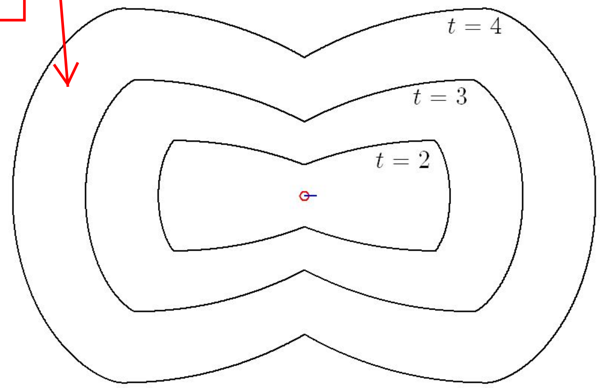

Fig. 1: The reachable set $\mathcal{R}(t)$ of a robot with dynamics given in (1), is depicted for varying time horizons. The robot is represented as a circle at the center.

Given any convex set $K \subset \mathbb{R}^n$, there exists a unique ellipsoid of minimum volume circumscribing it [8]. This ellipsoid is denoted as $\xi(K)$. One way to derive an ellipsoidal approximation of the reachable set would be to compute the minimum area ellipse $\xi(conv(\mathcal{R}(t)))$, where $conv(\mathcal{R}(t))$ denotes the convex hull of $\mathcal{R}(t)$. Additionally, the derivation of analytical expressions for the ellipse $\xi(conv(\mathcal{R}(t)))$, if at all possible, will enable its efficient and fast computation in the safe time horizon algorithm.

In order to allow for an analytical expression for the ellipse [6], we introduce a new set $\mathcal{K}(t)$ enclosing $conv(\mathcal{R}(t))$, which allows for an easier computation of $\xi(\mathcal{K}(t))$. Essentially, this enables us to swap the problem of computing $\xi(conv(\mathcal{R}(t)))$ with the simpler problem of computing $\xi(\mathcal{K}(t))$.

Furthermore, this approximation is justified by showing that the dissimilarity between $conv(\mathcal{R}(t))$ and $\mathcal{K}(t)$, as [illegible]y the Jaccard distance metric [7], asymptotically [illegible]o over time. For the sets $X, Y \subset \mathbb{R}^2$, the Jaccard [illegible]sed on the area measure is given by,

(For the proof, see [10].)

$$d_J(X, Y) = 1 - \frac{\mathbf{A}(X \cap Y)}{\mathbf{A}(X \cup Y)},$$

where $\mathbf{A}(\cdot)$ denotes the area of the set. The following proposition formally introduces the set $\mathcal{K}(t)$.

Proposition 1. *Let $\mathcal{K}(t)$ be given by,*

$$\mathcal{K}(t) = \Bigg\{ p = \begin{bmatrix} p_x \\ p_y \end{bmatrix} \in \mathbb{R}^2 : \|p\|_2 \leq t,$$
$$|p_y| \leq \begin{cases} 1 - \cos(t), & \text{if } \ 0 < t \leq \pi/2 \\ t - \pi/2 + 1, & \text{if } \ t > \pi/2 \end{cases} \Bigg\}.$$

Then, $conv(\mathcal{R}(t)) \subset \mathcal{K}(t)$ and

$$\lim_{t \to \infty} d_J(conv(\mathcal{R}(t)), \mathcal{K}(t)) = 0,$$

where d_J is the Jaccard distance.

~~*Proof.* See [10], *Proposition 1.*~~ ∎

The set $\mathcal{K}(t)$ not only allows us to derive analytical expressions for the minimum area ellipse enclosing it, the asymptotic reduction in the dissimilarity between $\mathcal{K}(t)$ and $conv(\mathcal{R}(t))$ implies that, the impact of using $\xi(\mathcal{K}(t))$ instead of $\xi(conv(\mathcal{R}(t)))$ on the accuracy of set-membership tests in the safe time horizon algorithm, is minimal.

In $\mathbb{R}^n$, an ellipsoid can be represented uniquely by its center c and a positive-definite matrix H: $E(c, H) = \{x \in \mathbb{R}^n : (x-c)^T H(x-c) \leq 1\}$. According to results presented in [8], at any given time t, the minimum area ellipse circumscribing $\mathcal{K}(t)$ can be obtained by solving the following semi-infinite programming problem:

$$\min_{c,H} -\log\ det(H) \qquad (2)$$
$$s.t.\ \ (z-c)^T H(z-c) \leq 1, \forall z \in \mathcal{K}(t).$$

The rest of this section formulates an analytical solution to this semi-infinite programming problem.

In Lemma 1, we utilize the symmetry properties of $\mathcal{K}(t)$ to determine the center and orientation of the ellipse $\xi(\mathcal{K}(t))$.

Lemma 1. *The ellipse $\xi(\mathcal{K}(t))$ has the form $E(c, H(t))$, where $c = (0, 0)$ and $H(t) = diag(A(t), B(t))$ for some $A(t), B(t) \in \mathbb{R}$, such that $A(t) > 0, B(t) > 0, \forall t > 0$.*

Proof. For a convex set $\mathbf{K} \subset \mathbb{R}^n$, denote $\mathcal{O}(\mathbf{K})$ as a set of affine transformations which leave the set $\mathbf{K}$ invariant, i.e., the automorphism group of $\mathbf{K}$. Applying results from [6], we know that, $\mathcal{O}(\mathbf{K}) \subseteq \mathcal{O}(\xi(\mathbf{K}))$. Since $\mathcal{K}(t)$ is symmetric about both the x and y axes, the transformation $\mathbf{T}(x) = -I_2 x$, where I_2 is the identity matrix, lies in $\mathcal{O}(\mathcal{K}(t))$, and hence in $\mathcal{O}(\xi(\mathcal{K}(t)))$. Furthermore, if $\mathbf{T} \in \mathcal{O}(\xi(\mathcal{K}(t)))$ then $\mathbf{T}(c) = c$. Applying this result with the transformation $\mathbf{T}(x) = -I_2 x$, we get $c = (0, 0)$.
We know from [6], that $\mathbf{T} \in \mathcal{O}(\xi(\mathcal{K}(t))) \implies P^T H(t) P = H(t)$. The structure of $H(t)$ appears by applying

$$P = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

and the fact that $H(t)$ is always positive definite. □

Applying results from [8] for the case of $\mathcal{K}(t) \subset \mathbb{R}^2$, we know that, there exist contact points $\{q_i\}_1^h$, $0 < h \leq 5$ satisfying $q_i \in \partial\mathcal{K}(t) \cap \partial\xi(\mathcal{K}(t)), i = 1, \ldots, h$. Furthermore, these contact points cannot all lie in any closed halfspace whose bounding hyperplane passes through the center of $\xi(\mathcal{K}(t))$. Using these facts, the following lemma can be stated:

Lemma 2. *There are 4 contact points between $\mathcal{K}(t)$ and $\xi(\mathcal{K}(t))$.*

Proof. We seek to find the number of contact points h. Since $\mathcal{K}(t)$ and $\xi(\mathcal{K}(t))$ are symmetric about the x and y axes (see Lemma 1), and $0 < h \leq 5$, h can only take values 2 or 4. If $h = 2$, both contact points must lie on the x or y axes (otherwise symmetry in all the 4 quadrants is not possible). But, the contact points cannot lie in any closed halfspace whose bounding hyperplane passes through the center of $\xi(\mathcal{K}(t))$. Hence $h = 4$. □

Given the structure of $H(t)$ from Lemma 1, and the number of contact points from Lemma 2, the semi-infinite programming problem (2) will now be re-formulated as a convex optimization problem.

Lemma 3. *The matrix $H(t) = diag(A(t), B(t))$ can be expressed as the solution of a convex optimization problem. At any given time $t > 0$, $A(t)$ is given as,*

$$A(t) = \underset{\mathbf{X}_t}{\text{argmin}}\ -\log\ \mathbf{X}_t - \log\ \frac{1 - \mathbf{X}_t(t^2 - \alpha(t)^2)}{\alpha(t)^2}$$
$$s.t.\ \ 0 < \mathbf{X}_t \leq \frac{1}{t^2}.$$

Furthermore,

$$B(t) = \frac{1 - A(t)(t^2 - \alpha(t)^2)}{\alpha(t)^2}$$

where $\alpha(t) = \begin{cases} 1 - \cos(t), & \text{if}\ \ 0 < t \leq \pi/2 \\ t - \pi/2 + 1, & \text{if}\ \ \ t > \pi/2 \end{cases}$.

Proof. See [10], *Lemma 3.* □

The next theorem solves the convex optimization problem outlined above to obtain analytical expressions for the minimum area ellipse enclosing $\mathcal{K}(t)$ (Fig. 2).

Theorem 1. *The minimum area ellipse $\xi(\mathcal{K}(t))$ has the form $E(c, H(t))$, where $c = (0, 0)$ and $H(t) = diag(A(t), B(t))$. $A(t)$ and $B(t)$ are given by the following expressions:*

1) *If $0 < t \leq \pi/2$, then*
$$A(t) = \frac{1}{2(t^2 - \alpha(t)^2)}\ \text{and}\ B(t) = \frac{1}{2\alpha(t)^2},$$
where $\alpha(t) = 1 - \cos(t)$.
2) *If $\pi/2 < t \leq (1 + \frac{1}{\sqrt{2}})(\pi - 2)$, then*
$$A(t) = \frac{1}{2(t^2 - \alpha(t)^2)}\ \text{and}\ B(t) = \frac{1}{2\alpha(t)^2},$$
where $\alpha(t) = t - \pi/2 + 1$.
3) *If $t > (1 + \frac{1}{\sqrt{2}})(\pi - 2)$,*
$$A(t) = \frac{1}{t^2}\ \text{and}\ B(t) = \frac{1}{t^2}$$

Proof. See [10], *Theorem 1.* □

The following theorem justifies the ellipsoidal approximation by showing that the dissimilarity between $conv(\mathcal{R}(t))$ and $\xi(\mathcal{K}(t))$ asymptotically goes to zero as time grows larger.

Theorem 2. *Let $\mathcal{R}(t)$ denote the reachable set of a differential-drive robot with dynamics given in* (1), *$conv(\mathcal{R}(t))$ denote its convex hull, $\mathcal{K}(t)$ denote an approximation of the convex hull as defined in Proposition 1, and let $\xi(t)$ be the minimum area ellipse circumscribing $\mathcal{K}(t)$. Then,*

$$\lim_{t \to \infty} d_J(conv(\mathcal{R}(t)), \xi(t)) = 0,$$

where d_J is the Jaccard distance.

low res?

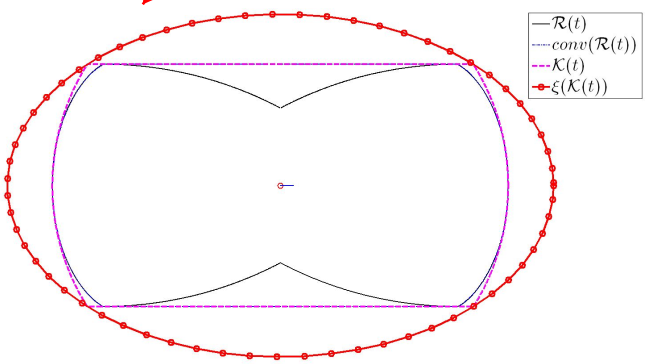

inimum area ellipse enclosing $\mathcal{K}(t)$

Proof. See [10], *Theorem 2.* □

Fig. 3 shows the evolution of $\mathcal{R}(t), conv(\mathcal{R}(t)), \mathcal{K}(t)$ and $\xi(\mathcal{K}(t))$ for different values of time. As predicted by the result, the dissimilarity between $conv(\mathcal{R}(t))$, $\mathcal{K}(t)$ and $\xi(\mathcal{K}(t))$ asymptotically goes to zero.

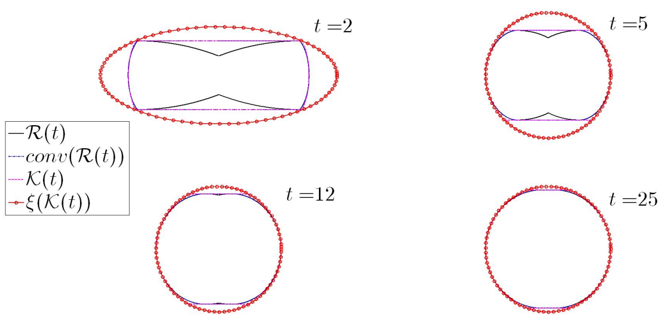

Fig. 3: Evolution of $\mathcal{R}(t), conv(\mathcal{R}(t)), \mathcal{K}(t)$ and $\xi(\mathcal{K}(t))$ for $t = 2, 5, 12, 25$. As predicted by Theorem 2, the dissimilarity between $conv(\mathcal{R}(t)$ and $\xi(\mathcal{K}(t))$ asymptotically goes to zero. The scale for each figure is different.

Using the ellipsoidal approximation of the reachable set developed in this section, Section III outlines a safe time horizon based open-loop motion strategy for the robots.

III. A Safe Open-Loop Motion Strategy

Let $\mathcal{M} = \{1, \ldots, N\}$ be a set of N differential-drive robots, where each robot moves according to the dynamics specified in (1). Let $(z_i(t), \phi_i(t))$ denote the configuration of robot $i \in \mathcal{M}$ at time t. At regular intervals of time $t_k = k\delta, k \in \mathbb{N}$, the central decision maker transmits the desired velocities $u_i(t_k) = (v_i(t_k), \omega_i(t_k))$ and the corresponding safe time horizon $s_i(t_k)$ to each robot $i \in \mathcal{M}$ via a wireless communication channel. $1/\delta$ is called the update frequency.

When a robot experiences communication failure, it executes the last received velocity command repeatedly for the duration of the corresponding safe time horizon. This causes the robot to follow a circular trajectory. Let $Z_i(\mu, t_k)$ denote the position of robot i along this circular trajectory, where t_k is the time of the last received command and μ is the time elapsed since the communication failure. The expressions for $Z_i(\mu, t_k)$ can be obtained by integrating (1) for constant velocity inputs.

In order to ensure the scalability and computational tractability of the safe time horizon algorithm, we introduce the notion of a neighborhood set for each robot. To do this, the safe time horizon for each robot is upper-bounded by a pre-specified value L. This allows us to introduce the neighborhood set of robot i at time t as:

$$N_i(t) = \{j \in \mathcal{M}, j \neq i : \|z_i(t) - z_j(t)\| < 2L\},$$

where $\|.\|$ denotes the l_2 norm. If robot i and robot j are not neighbors, they cannot collide within the maximum safe time horizon L.

The safe time horizon $s_i(t_k)$ can be defined as,

$$s_i(t_k) = \min_{j \in N_i} s_{ij}(t_k),$$

where $s_{ij}(t_k)$ is called the pair-wise safe time and is defined as,

$$s_{ij}(t_k) = \max_{\lambda} \int_0^{\lambda} 1 \, d\lambda \tag{3}$$
$$\text{s.t. } Z_i(\mu, t_k) \notin \mathcal{R}(\mu; z_j(t_k), \phi_j(t_k)), \forall \mu \in [0, \lambda]$$
$$\text{and } \lambda \leq L.$$

Thus, the safe time horizon is the longest amount of time for which the trajectory of the robot after communication failure, does not intersect the reachable sets of its neighbors. But, as motivated in Section II, the definition of $s_{ij}(t_k)$ can be modified by replacing $\mathcal{R}(\mu; z_j(t_k), \phi_j(t_k))$ with $\xi_j(\mu, t_k)$ in (3), where $\xi_j(\mu, t_k)$ denotes the ellipsoidal approximation corresponding to $\mathcal{R}(\mu; z_j(t_k), \phi_j(t_k))$ as derived in Section II. Next, we discuss how the safe time horizon is incorporated into the motion strategy of the robots.

As discussed earlier, the central decision maker transmits $u_i(t_k)$ and $s_i(t_k)$ to all the robots $i \in \mathcal{M}$ at regular time intervals t_k, $k \in \mathbb{N}$. Let c_i represent the status of the communication link of robot i:

$$c_i(t_k) = \begin{cases} 1, & \text{if } (u_i(t_k), s_i(t_k)) \text{ was received} \\ 0, & \text{if } (u_i(t_k), s_i(t_k)) \text{ was not received.} \end{cases}$$

Algorithm 1 outlines the motion strategy that robot i employs.

Fig. 4 illustrates the rationale behind the safe time horizon algorithm. As long as the robot experiencing communication failure is outside the ellipsoidal reachable sets of its neighbors, it can safely move. The end of the safe time horizon corresponds to the time when the robot reaches the boundary of one of the ellipses. At this point, the robot stops moving.

In order to state formal safety guarantees regarding Algorithm 1, we make mild assumptions on the capability of the control algorithm executing on the central decision maker. We assume that, the control algorithm ensures collision avoidance between communicating robots as well as between communicating robots and stationary obstacles. In particular,

Algorithm 1 Safe Time Horizon based Open-Loop Motion Strategy

$k = 1$, $l = 1$; $u_i(0) = 0, s_i(0) = 0$
while `true` **do**
  if $c_i(t_k) = 1$ **then**
    `Execute` $u_i(t_k)$
    $l = k$
  else if $t_k - t_l < s_i(t_l)$ **then**
    `Execute` $u_i(t_l)$
  else
    `Stop Moving`
  end if
  $k = k + 1$
end while

if $\exists i, j \in \mathcal{M}$ such that $c_i(t_k) = 1$ and $c_j(t_k) = 1$, then $u_i(t_k)$ and $u_j(t_k)$ guarantee that,

$$\|z_i(t_k) - z_j(t_k)\| > 0 \implies \|z_i(t_{k+1}) - z_j(t_{k+1})\| > 0.$$

Let z_O denote the position of a stationary obstacle. If $c_i(t_k) = 1$ for any $i \in \mathcal{M}$,

$$\|z_i(t_k) - z_O\| > 0 \implies \|z_i(t_{k+1}) - z_O\| > 0. \quad (4)$$

Utilizing these assumptions, the following theorem outlines the safety guarantees provided by Algorithm 1.

Theorem 3. *If robot i does not receive any commands from the central decision maker after time t_k, i.e., $c_i(t_k) = 1$ and $c_i(t_m) = 0\ \forall m > k$, then Algorithm 1 ensures that,*

$$\|z_i(t_k) - z_j(t_k)\| > 0 \implies \|z_i(t_k + \mu) - z_j(t_k + \mu)\| > 0,$$
$$\forall \mu \in [0, s_i(t_k)],\ \forall j \in \mathcal{M}, j \neq i.$$

Proof. See [10], *Theorem 3.* □

Beyond the safe time horizon, the robot stops moving, and (4) ensures that no collisions occur with the stationary robot. Thus, the original safety guarantee of the control algorithm is extended to situations where the robot is moving without commands from the central decision maker within the safe time horizon.

IV. Results

This section provides validation of the safe time horizon algorithm using simulations as well as experiments conducted on a multi-robot testbed.

Fig. 5 simulates the motion of 6 robots during a communication failure with and without the safe time horizon algorithm. In the case where safe time horizons are not utilized, shown by Fig. 5a and Fig. 5b, the robots experiencing communication failure abruptly stop moving, thus exhibiting a jerky motion pattern. When safe time horizons are utilized, the robots experiencing communication failure execute their last received velocity command for the duration of the safe time horizon (Fig. 5c and Fig. 5d), thus avoiding jerky "start-stop" motion behaviors.

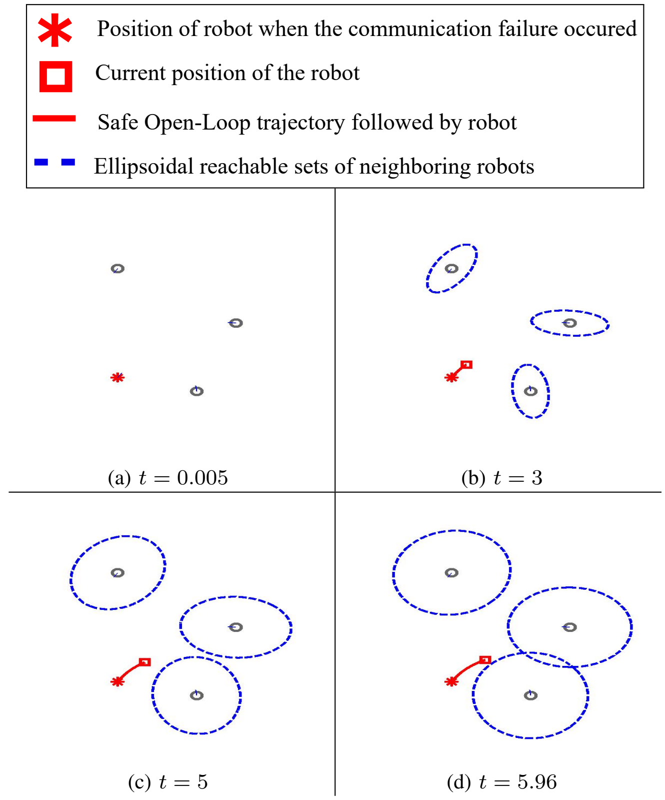

(a) $t = 0.005$ (b) $t = 3$ (c) $t = 5$ (d) $t = 5.96$

Fig. 4: The safe time horizon represents the longest time duration for which the robot lies outside the ellipsoidal reachable sets of other robots. Thus, the robot experiencing communication failure can execute its last received velocity command for the corresponding safe time horizon and remain safe. Beyond this, the robot stops moving.

The safe time horizon algorithm is also implemented on a multi-robot testbed with 4 Khepera III robots and an Optitrack motion capture system. In the first scenario (Fig. 6a and Fig. 6b), the robot experiencing communication failure executes its last received velocity command for the duration of the safe time horizon, after which it becomes stationary. In the second scenario, (Fig. 6c and Fig. 6d), the robot experiences a short duration communication failure and keeps moving safely through it. In particular, the safe time horizon algorithm successfully combats issues pertaining to intermittent communications on the Robotarium, and enables the seamless execution of coordination algorithms in the face of such failures.

V. Conclusions

The safe time horizon algorithm not only provides a technique for multi-robot systems to safely handle intermittent communication failures, it demonstrates the feasibility of reachability analysis as a powerful tool for multi-robot algorithms. The minimum area ellipse derived in Section II provides a compact and efficient way to represent the reachable set of a differential-drive robot and can be used in other robotics algorithms as an abstraction of the reachable set itself.

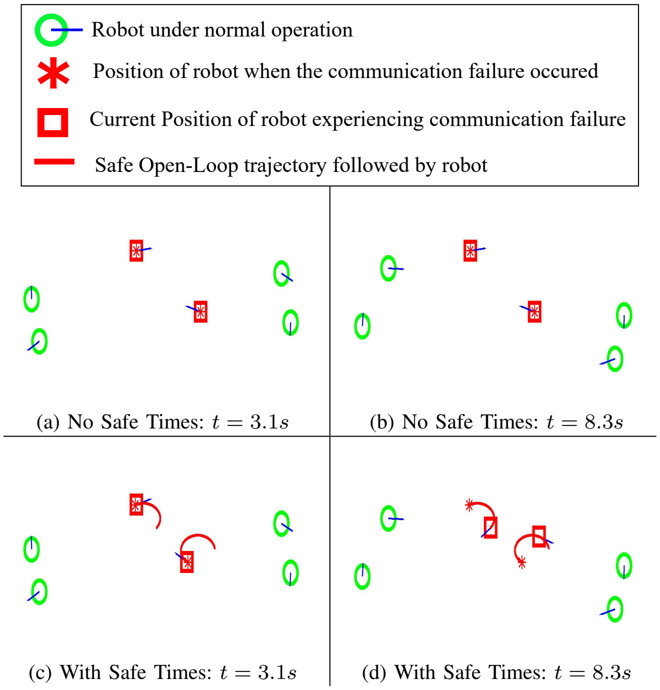

(a) No Safe Times: $t = 3.1s$ (b) No Safe Times: $t = 8.3s$

(c) With Safe Times: $t = 3.1s$ (d) With Safe Times: $t = 8.3s$

Fig. 5: Simulation results comparing of the motion of robots with and without safe time horizons. Two robots experience communication failure from $t = 3.1s$ to $t = 8.3s$. In the case when safe time horizons are not used (Fig. 5a and Fig. 5b), the robots exhibit jerky motion behavior, since they abruptly stop during the communication failure. When safe time horizons are used (Fig. 5c and Fig. 5d), the robots continue moving by executing their last received velocity command for the corresponding safe time horizon, thus demonstrating the ability of the safe time horizon algorithm to effectively handle communication failures.

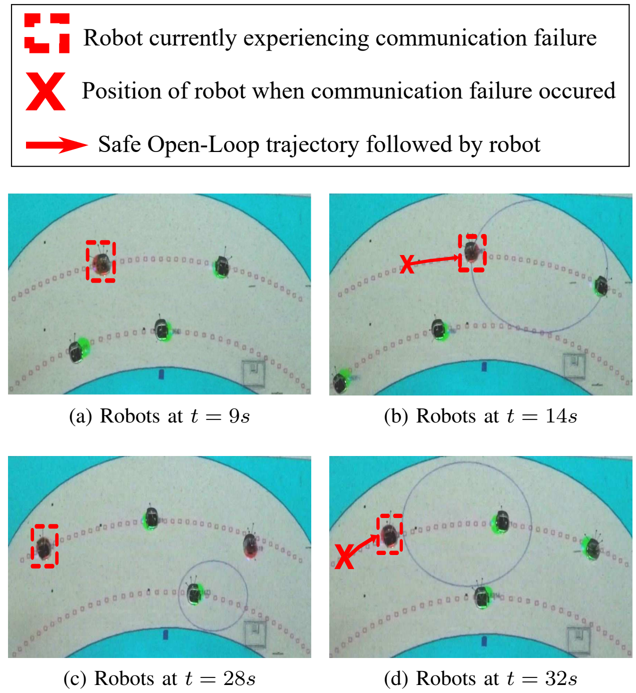

(a) Robots at $t = 9s$ (b) Robots at $t = 14s$

(c) Robots at $t = 28s$ (d) Robots at $t = 32s$

Fig. 6: Four Khepera III robots are shown patrolling a U-shaped corridor. At $t = 9s$, a robot experiences communication failure and executes its last received velocity command until $t = 14s$. Similarly, a robot loses communication at $t = 28s$ (Fig. 6c), but the developed algorithm allows it to continue moving until communication is restored (Fig. 6d). This experiment demonstrates how the safe time horizon algorithm can prevent disruptions in robot motion caused due to intermittent communications on the Robotarium. A video of this experiment can be found at `www.youtube.com/watch?v=Gyz861xwaHY`

REFERENCES

[1] Tucker Balch and Ronald C Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

[2] M Dogar, R A Knepper, A Spielberg, C Choi, H I Christensen, and D Rus. Multi-scale assembly with robot teams. *The International Journal of Robotics Research*, 34(13):1645–1659, 2015.

[3] A Fedotov, V Patsko, and V Turova. *Reachable Sets for Simple Models of Car Motion*. INTECH Open Access Publisher, 2011.

[4] Jonathan Fink, Alejandro Ribeiro, and Vijay Kumar. Robust control for mobility and wireless communication in cyber–physical systems with application to robot teams. *Proceedings of the IEEE*, 100(1):164–178, 2012.

[5] Stephanie Gil, Dan Feldman, and Daniela Rus. Communication coverage for independently moving robots. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4865–4872. IEEE, 2012.

[6] O Güler and F Gürtuna. Symmetry of convex sets and its applications to the extremal ellipsoids of convex bodies. *Optimization Methods and Software*, 27(4-5):735–759, 2012.

[7] Paul Jaccard. *Distribution de la Flore Alpine: dans le Bassin des dranses et dans quelques régions voisines*. Rouge, 1901.

[8] F John. Extremum problems with inequalities as subsidiary conditions. In *Traces and Emergence of Nonlinear Programming*, pages 197–215. Springer, 2014.

[9] Oded Maler. Computing reachable sets: An introduction. Technical report, 2008.

[10] Siddharth Mayya and Magnus Egerstedt. Safe open-loop strategies for handling intermittent communications in multi-robot systems. *arXiv pre-print, arXiv:1702.03466*, 2017.

[11] Suruz Miah, Bao Nguyen, Alex Bourque, and Davide Spinello. Nonuniform coverage control with stochastic intermittent communication. *IEEE Transactions on Automatic Control*, 60(7):1981–1986, 2015.

[12] D Pickem, L Wang, P Glotfelter, Y Diaz-Mercado, M Mote, A Ames, E Feron, and M Egerstedt. Safe, remote-access swarm robotics research on the robotarium. *arXiv preprint arXiv:1604.00640*, 2016.

[13] R Ramanathan and J Redi. A brief overview of ad hoc networks: challenges and directions. *IEEE communications Magazine*, 40(5):20–22, 2002.

[14] P Souères, J Y Fourquet, and J Pl Laumond. Set of reachable positions for a car. *Automatic Control, IEEE Transactions on*, 39(8):1626–1630, 1994.

[15] Ethan Stump, Ali Jadbabaie, and Vijay Kumar. Connectivity management in mobile robot teams. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, pages 1525–1530. IEEE, 2008.

[16] Patrick Ulam and Ronald C Arkin. When good communication go bad: communications recovery for multi-robot teams. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 4, pages 3727–3734. IEEE, 2004.

[17] Guanghui Wen, Zhisheng Duan, Wei Ren, and Guanrong Chen. Distributed consensus of multi-agent systems with general linear node dynamics and intermittent communications. *International Journal of Robust and Nonlinear Control*, 24(16):2438–2457, 2014.

[18] Logan Yliniemi, Adrian K Agogino, and Kagan Tumer. Multirobot coordination for space exploration. *AI Magazine*, 35(4):61–74, 2014.