

Efficient Safety Control Synthesis with Imperfect State Information

Liren Yang

Necmiye Ozay

Abstract—We consider synthesizing safety controllers for discrete-time dynamical systems with imperfect (i.e., noisy) state measurements. In order to find the actual winning set of a safety game for such systems, one needs to solve a partial information game via power set construction, which, in general, is computationally intractable. In this paper, we propose two conservative but computationally more efficient approaches by computing sets that can be rendered invariant with the noisy measurement. This is achieved by considering a perfect information safety game for the dynamics of an estimated state. The invariant set for this alternative game is shown to be equivalent to a noise-adapted contractive set for the original system. The controllers associated with the two proposed approaches require different knowledge of the initial states: one requires only the initial measurement and the other also requires knowing the initial state exactly. In general, the resulting controlled invariant sets by these two approaches are not comparable, and depending on the problem in hand either one can be preferable. The efficacy of the proposed approach is illustrated with an aircraft taxiing example, where the state estimation task is performed by a perception module.

I. INTRODUCTION

In this paper we consider synthesizing safety controllers for dynamical systems to ensure that their states remain in a safe set for all time. To find where such a safety controller can be initiated and operated, one needs to compute the maximal robust controlled invariant set contained by the safe set. The properties and computation of robust controlled invariant sets are well studied in the literature of control theory, both for discrete-state (see e.g., [12]) and continuous-state systems (see e.g., [2], [3] and the references therein).

At runtime, the safety controller defined on the controlled invariant set takes the true state x as input and maps x to a set of admissible control inputs u , which guarantee that the next state will stay in the controlled invariant set. In practice, however, the true value of the state x may not be known exactly. Instead, we observe an estimate \hat{x} of the true state x with certain measurement noise. Such noisy estimate \hat{x} may come from an interval observer, or even a perception module that estimates the system's state from camera images. The safety controller must make decisions based on the noisy measurement \hat{x} instead of x but still guarantee safety. This hence requires computing sets that can be made controlled invariant with noisy state measurement.

Safety control synthesis with noisy measurements is a partial information game as each observation \hat{x} may be valid for multiple values of the true state x . Complete (i.e., non-conservative) solutions of such partial information games

usually require lifting the system into the belief space via power set construction. This allows the controller to estimate the true state as accurately as possible using the entire history of the collected data instead of the latest measurement, which is enough in full information settings. Although this is computationally expensive, it is still theoretically possible for discrete-state systems, see e.g., [6], [13], [16]. Similar ideas are explored for continuous-state systems in [1], [10], but the heavy computational requirements make it hard to employ these techniques for real applications.

A computationally more efficient approach tackling this problem is to design conservative controllers by enforcing some form of contractivity to overcome the impact of the measurement noise. For example, set invariance for noisy continuous-time continuous-state systems is studied in [8] with this idea, using control barrier functions as a means to enforce a Nagumo type condition strengthened with extra contraction. Similar contractivity-based ideas are also used in abstraction-based synthesis approaches for more general linear temporal logic (LTL) specifications [9], and opacity verification [17]. A slightly different yet highly related approach is to use contraction to design bounded-error observers and achieve the control objective robust against any noise within the bounds. For example, this is used for two-player games [11] and path planning [15] with LTL specifications with imperfect measurement.

In this paper, we consider synthesizing safety controllers for discrete-time systems with imperfect state measurement, which may have discrete or continuous state space. In particular, we follow the second line of research that uses contraction to avoid the expensive power set construction. We propose two efficient approaches to compute sets that can be rendered invariant with the imperfect measurement. This is done by solving a perfect information safety game for the dynamics of an estimated state, which is described using a generalization of the Minkowski set arithmetic. We prove that the obtained invariant sets are equivalent to some noise-adapted contractive sets for the original system. Such contraction ensures robustness against noise with a simple controller that only uses the latest measurement, and hence avoids complex state estimation algorithms involving the power set construction. In particular, our generalization of the Minkowski arithmetic enables a uniform approach to handle systems with both continuous and discrete states, and can adapt to systems with non-uniform state-dependent measurement noise as well.

To illustrate the efficacy of our approach, we present a case study on an aircraft taxiing system, whose state is estimated from camera images at runtime by a perception module.

The authors are with the Department of Electrical Engineering and Computer Science, Univ. of Michigan, Ann Arbor, MI 48109, USA. Emails: yliren@umich.edu, necmiye@umich.edu.

The state estimation is imperfect due to the noise in image formation and the error introduced by the vision algorithms. Using a realistic flight simulator, X-plane, we show that the obtained controller is able to achieve safety despite the aforementioned sensing and perception inaccuracies.

Notation: Let S be a set, 2^S denotes the power set of S , and S^* (S^ω , respectively) denotes the set of all finite (infinite, respectively) sequences of elements in S . Throughout this paper, we will use bold font letters, e.g., $\mathbf{s} = s_0s_1s_2\dots$, to denote an infinite sequence, and use $s_t = s_0s_1s_2\dots s_t$ to denote the finite prefix of \mathbf{s} until t .

II. SYSTEM WITH NOISY STATE MEASUREMENT

In this paper, we consider systems in the following form:

$$\Sigma : x_{t+1} \in \tau(x_t, u_t), \quad (1)$$

$$\hat{x}_t \in \mu(x_t), \quad (2)$$

where $x \in X$ is the state, $u \in U$ is the control input, $\hat{x} \in \hat{X} = X$ is the noisy state measurement, $\tau : X \times U \rightarrow 2^X$ is the state transition mapping, and $\mu : X \rightarrow 2^{\hat{X}}$ is the measurement mapping. The uncertainty in the evolution and state measurement is captured by the fact that $\tau(x, u)$ and $\mu(x)$ are sets. We will overload the notation of τ and μ for set inputs as well, i.e., for sets $X_0 \subseteq X$, $U_0 \subseteq U$, $\tau(X_0, U_0) := \bigcup_{x \in X_0, u \in U_0} \tau(x, u)$ and $\mu(X_0) := \bigcup_{x \in X_0} \mu(x)$. We also make the following assumption on the measurement map μ .

Assumption 1: At least one measurement is available at every state, i.e., $\mu(x) \neq \emptyset$ for all $x \in X$.

Let $\pi : \hat{X}^* \times U^* \rightarrow U$ be the control policy and $x \in X$ be a state, an infinite sequence $\mathbf{x} = x_0x_1x_2\dots \in X^\omega$ is a trajectory generated by the closed-loop system Σ_π starting from state x if there exists $\hat{\mathbf{x}} = \hat{x}_0\hat{x}_1\hat{x}_2\dots \in \hat{X}^\omega$, $\mathbf{u} = u_0u_1u_2\dots \in U^\omega$ such that

$$\text{i) } x_0 = x,$$

$$\text{ii) } \forall t : x_{t+1} \in \tau(x_t, u_t), \hat{x}_t \in \mu(x_t), u_t = \pi(\hat{\mathbf{x}}_t, \mathbf{u}_{t-1}).$$

We will define $\mathfrak{B}(\Sigma_\pi, x) := \{(\mathbf{x}, \hat{\mathbf{x}}) \mid \exists \mathbf{u} : \text{i) and ii) hold}\}$ to be the set of behaviors¹ of the closed-loop system from initial state x , and define $\mathfrak{B}(\Sigma_\pi, X_0) := \bigcup_{x \in X_0} \mathfrak{B}(\Sigma_\pi, x)$.

III. PROBLEM STATEMENT

In this section, we define the safety control synthesis problem for systems with noisy state measurements. The problem can be viewed as a game between the controller and the environment, which ‘‘chooses’’ the uncertainties in the system and aims at violating safety requirement of the system. A set of states from where safety can be enforced by the controller is hence called a winning set and is formally defined below.

Definition 1: (Winning Condition of Safety Game) Given a set X_{safe} of safe states, a set $W \subseteq X_{\text{safe}}$ is *winning* w.r.t. the safety specification under the dynamics in (1), (2) if there exists a control policy $\pi : \hat{X}^* \times U^* \rightarrow U$, such that

$$\forall (\mathbf{x}, \hat{\mathbf{x}}) \in \mathfrak{B}(\Sigma_\pi, W), (\mathbf{x}', \hat{\mathbf{x}}') \in \mathfrak{B}(\Sigma_\pi, X_{\text{safe}}) \text{ s.t. } \hat{x}'_0 = \hat{x}_0 : \mathbf{x}' \in X_{\text{safe}}^\omega \quad (3)$$

¹Since the control sequence \mathbf{u} is fully determined by the observation $\hat{\mathbf{x}}$ given controller π , there is no need to include \mathbf{u} in the behavior.

and a set $W \subseteq X_{\text{safe}}$ is *weakly winning* if there exists a control policy $\pi : \hat{X}^* \times U^* \rightarrow U$ such that

$$\forall (\mathbf{x}, \hat{\mathbf{x}}) \in \mathfrak{B}(\Sigma_\pi, W) : \mathbf{x} \in X_{\text{safe}}^\omega, \quad (4)$$

Remark 1: Clearly Eq. (3) implies Eq. (4). There exists a maximal (in the set inclusion sense) winning set W_{max} , which is the union of all set satisfying Eq. (3). However, a maximal weakly winning set does not necessarily exist. In fact, the union of two sets W_1, W_2 that satisfy Eq. (4) under controller π_1, π_2 respectively does not necessarily satisfy Eq. (4). This is because in general it is impossible to find a controller π under which $W_1 \cup W_2$ is weakly winning, unless the initial state is known exactly, in which case we can choose to use π_1 or π_2 accordingly.

The safety control synthesis problem is stated as follows.

Problem 1: Suppose Assumption 1 holds, given a system Σ in Eq. (1), (2), and a set X_{safe} of safe states, find a set $W \subseteq X_{\text{safe}}$, a feedback control mapping $\pi : \hat{X}^* \times U^* \rightarrow U$ such that W is winning (or weakly winning) under π .

A concept closely related to the winning set of safety games is the controlled invariant set.

Definition 2: A set $C \subseteq X$ is *controlled invariant* w.r.t. system Σ if there exists π such that $\mathfrak{B}(\Sigma_\pi, C) \subseteq C^\omega$.

Remark 2: For systems with perfect state measurement (i.e., $\mu(x) = \{x\}$), a maximal controlled invariant set C_{max} , contained by X_{safe} , exists and $W_{\text{max}} = C_{\text{max}}$. This is not the case when the measurement is imperfect. In fact, trajectory starting from W_{max} may leave W_{max} under the winning strategy [1] in this setting. Under imperfect information, a controlled invariant set contained by X_{safe} is weakly winning by definition but not necessarily winning.

IV. SOLUTION APPROACH

Problem 1 can be viewed as a partial information game and can be solved in a sound and complete way. That is, the maximal winning set (or the maximal weakly winning set whenever the initial state is known exactly) can be found by lifting the system into the belief space via power set construction. This approach, however, is computationally expensive because the lifted system has a state space whose size is exponential in that of the original system. Moreover, for continuous state systems, synthesis in the belief space is generally intractable. In this section, we propose two conservative yet efficient ways to solve Problem 1, one computes a winning set and the other computes a weakly winning set. Both sets are not necessarily maximal.

A. Measurement Mapping and Inverse Mappings

We first introduce several mappings that relate true states with their measurements.

- the set of possible true states given measurement \hat{x} :

$$\mu^{\oplus-1}(\hat{x}) := \{x \mid \hat{x} \in \mu(x)\}, \quad (5)$$

$$\text{and } \mu^{\oplus-1}(\hat{S}) := \bigcup_{\hat{x} \in \hat{S}} \mu^{\oplus-1}(\hat{x}).$$

- the set of true states whenever we know the measurement cannot go beyond \hat{S} regardless of the noise:

$$\mu^{\ominus 1}(\hat{S}) := \{x \mid \mu(x) \subseteq \hat{S}\}. \quad (6)$$

- the set of measurement whose associated true states must be in S :

$$\mu^{\ominus-1}(S) := \{\hat{x} \mid \mu^{\oplus-1}(\hat{x}) \subseteq S\}. \quad (7)$$

We now provide several remarks on the above notations. First, map μ and $\mu^{\ominus-1}$ are similar in the sense that their superscripts both result in “positive one” and both μ and $\mu^{\ominus-1}$ map a set of true states to a set of measurements; whereas $\mu^{\oplus-1}$ and $\mu^{\ominus1}$, whose superscripts result in “negative one”, are certain inverse maps that bring a set of measurements to a set of true states.

Second, the difference between $\mu^{\ominus1}(\hat{S})$ and $\mu^{\oplus-1}(\hat{S})$ is:

- 1° if we only know that the measurement $\hat{x} \in \hat{S}$, we can only say that the true state $x \in \mu^{\oplus-1}(\hat{S})$;
- 2° if we know that a) the measurement $\hat{x} \in \hat{S}$, and b) the measurement \hat{x} must be in \hat{S} no matter what noise the environment picks, we can conclude that the true state is in $\mu^{\ominus1}(\hat{S})$.

Finally, for the readers who are familiar with Minkowski sum and difference [14], it might be helpful to think of the above notations as such.

Example 1: Let $X \subseteq \mathbb{R}^n$, and $\mu(x) = x \oplus V = \{x + v \mid v \in V\}$ be the Minkowski sum of x and a set $V \subseteq \mathbb{R}^n$ of admissible noise. In this case,

- $\mu^{\oplus-1}(\hat{S}) = \hat{S} \oplus -V$;
- $\mu^{\ominus1}(\hat{S}) = \hat{S} \ominus V := \{x \mid x \oplus V \subseteq \hat{S}\}$, which is the Minkowski difference between \hat{S} and V ;
- $\mu^{\ominus-1}(S) = S \ominus -V$.

The following Lemmas can be easily proven by definitions and will be useful in later proofs.

Lemma 1: $\mu(A) \subseteq \hat{B}$ is equivalent to $A \subseteq \mu^{\ominus1}(\hat{B})$.

Lemma 2: Under Assumption 1, i.e., $\mu(x) \neq \emptyset$ for all $x \in X$, $\mu(A) \subseteq \hat{B}$ implies that $A \subseteq \mu^{\oplus-1}(\hat{B})$.

Note that Lemma 2 does not necessarily hold without Assumption 1. To see this, let $A = \{x\}$ such that $\mu(x) = \emptyset$ and let $\hat{B} = \emptyset$. Clearly, $\mu(A) = \emptyset \subseteq \hat{B}$ but $A \not\subseteq \mu^{\oplus-1}(\hat{B})$.

Lemma 3: Under Assumption 1, $\mu^{\ominus1}(\hat{B}) \subseteq \mu^{\oplus-1}(\hat{B})$.

Lemma 4: $\mu^{\oplus-1}(\mu^{\ominus-1}(A)) \subseteq A \subseteq \mu^{\ominus1}(\mu(A))$.

B. Synthesis without Power Set Construction

In the rest of this section, we propose two safety control synthesis algorithms for noisy systems without power set construction. The key idea is, instead of considering system Σ , we will consider an auxiliary system $\hat{\Sigma}$ that captures the dynamics of measurement \hat{x} :

$$\hat{\Sigma} : \hat{x}_{t+1} \in \mu\left(\tau(\mu^{\oplus-1}(\hat{x}_t), u_t)\right). \quad (8)$$

Then we can use off-the-shelf tools² to efficiently compute a (approximately) maximal controlled invariant set of $\hat{\Sigma}$ (the set is called “ \hat{x} -invariant”) and recover a controlled invariant set of Σ because x and \hat{x} are closely related by the measurement mapping μ . The above intuition leads to two slightly different algorithms, one gives a winning set and a

²For example, the implementation used in this paper is available in <https://github.com/pettini/pcis>.

controller that only uses the latest noisy measurement for decision making, while the other leads to a weakly winning set and a controller that also requires the exact knowledge of the initial state. We will present the two algorithms separately and compare the obtained sets in what follows.

1) *Computing Winning Set:* We first consider a set of states that can be rendered invariant with a controller that only uses the last noisy measurement as input. The set and the associated controller can be found by Algorithm 1.

Algorithm 1 $[C_1, \pi] = \text{Win}_1(X_{\text{safe}}, \Sigma)$

- 1: Find the largest set $\hat{C} \subseteq \mu(X_{\text{safe}})$ and a controller $\hat{\pi} : \hat{C} \rightarrow U$ such that

$$\forall \hat{x} \in \hat{C} : \mu\left(\tau(\mu^{\oplus-1}(\hat{x}), \hat{\pi}(\hat{x}))\right) \subseteq \hat{C}. \quad (9)$$
 - 2: $C_1 = \mu^{\ominus1}(\hat{C})$
 - 3: define π to be s.t. $\pi(\hat{\mathbf{x}}_t, \mathbf{u}_{t-1}) = \hat{\pi}(\hat{x}_t)$
 - 4: **return** C_1, π
-

With a slight abuse of notation, we will write $\pi(\hat{\mathbf{x}}_t, \mathbf{u}_{t-1})$ as $\pi(\hat{x}_t)$ (or just $\pi(\hat{x})$ when the time information is not important) in the sequel.

We have the following results regarding Algorithm 1.

Proposition 1: Let C_1 and π be returned by Algorithm 1,

$$\forall \hat{x} \in \hat{C}, x \in \mu^{\oplus-1}(\hat{x}) : \tau(x, \pi(\hat{x})) \subseteq C_1. \quad (10)$$

Proof: First, since $x \in \mu^{\oplus-1}(\hat{x})$, we have

$$\tau(x, \pi(\hat{x})) \subseteq \tau(\mu^{\oplus-1}(\hat{x}), \pi(\hat{x})) \quad (11)$$

Secondly, by definition of \hat{C} :

$$\begin{aligned} \hat{x} \in \hat{C} &\Rightarrow \mu\left(\tau(\mu^{\oplus-1}(\hat{x}), \pi(\hat{x}))\right) \subseteq \hat{C} \\ &\Leftrightarrow \tau(\mu^{\oplus-1}(\hat{x}), \pi(\hat{x})) \subseteq \mu^{\ominus1}(\hat{C}) \quad (\text{Lemma 1}) \\ &\Leftrightarrow \tau(\mu^{\oplus-1}(\hat{x}), \pi(\hat{x})) \subseteq C_1 \end{aligned} \quad (12)$$

Combining Eq. (11), (12) yields

$$\tau(x, \pi(\hat{x})) \subseteq \tau(\mu^{\oplus-1}(\hat{x}), \pi(\hat{x})) \subseteq C_1, \quad (13)$$

which is what we want to prove. \blacksquare

Remark 3: Proposition 1 requires the set C_1 to have a property stronger than invariance. In fact, set C_1 is not only controlled invariant, but also contracting. That is, as Proposition 1 suggests, even for a state out of C_1 , as long as it may generate a measurement $\hat{x} \in \hat{C}$, there is a control action $\pi(\hat{x})$ that brings the true state x into C_1 in one step. This means that, unless a strong enough contraction can be enforced for any set of states, algorithm will return $C_1 = \emptyset$. We illustrate this with the following example.

Example 2: Consider the finite transition system in Fig. 1. Let $X = \{x_1, x_2, \dots, x_6\}$, $U = \{u_1\}$ and let $X_{\text{safe}} = \{x_3, x_4, x_5, x_6\}$. The transition mapping τ is such that $\tau(x_i) = \{x_i\}$ for $i = 1, 2, \dots, 6$, and the measurement mapping μ is such that

$$\begin{aligned} \mu(x_1) &= \{x_1, x_2\}, \\ \mu(x_i) &= \{x_{i-1}, x_i, x_{i+1}\}, \text{ for } i = 2, 3, 4, 5, \end{aligned}$$




Fig. 1: Illustration of Example 2. Blue arrows mark the transition under control input u_1 .

$$\mu(x_6) = \{x_5, x_6\}.$$

In this example, it can be seen that, the largest winning set found with power set construction is $\{x_4, x_5, x_6\}$. However, Algorithm 1 returns $C_1 = \emptyset$ because the system lacks of necessary contraction suggested by Proposition 1.

Remark 4: Note that C_1 is not necessarily contained by X_{safe} in general although $\hat{C} \subseteq \mu(X_{\text{safe}})$ by construction. We illustrate this situation by the following example.

Example 3: Consider the finite transition system in Fig. 2. Let $X = \{x_1, x_2, x_3\}$, $U = \{u_1, u_2\}$, and let the safe set $X_{\text{safe}} = \{x_1, x_3\}$. Suppose τ is such that

$$\begin{aligned} \tau(x_1, u_1) &= \{x_3\}, \tau(x_1, u_2) = \{x_2\}, \\ \tau(x_2, u_1) &= \{x_2\}, \tau(x_2, u_2) = \{x_2\}, \\ \tau(x_3, u_1) &= \{x_2\}, \tau(x_3, u_2) = \{x_1\}; \end{aligned}$$

and μ is such that

$$\mu(x_1) = \{x_1, x_2\}, \mu(x_2) = \{x_2\}, \mu(x_3) = \{x_2, x_3\}.$$

In this case, it can be seen that $\mu(X_{\text{safe}}) = \{x_1, x_2, x_3\}$, and $\hat{C} = \{x_1, x_2, x_3\}$ is the largest controlled invariant set under $\hat{\Sigma}$ contained by $\mu(X_{\text{safe}})$. However, $C_1 = \mu^{\ominus 1}(\hat{C}) = \{x_1, x_2, x_3\} \not\subseteq X_{\text{safe}}$. This is true because if the measurement $\hat{x} = x_2$, there is no chance that we can tell if the true state is unsafe. As a result, Algorithm 1, which only yields a controller that makes its decision based on the latest measurement \hat{x}_t , will not remove x_2 from \hat{C} .




Fig. 2: Illustration of Example 3. Blue arrows mark the transition under control input u_1 , red arrows mark the transitions under u_2 .

To remove the above situation, we make the following extra assumption on the safe set X_{safe} and the measurement mapping μ .

Assumption 2: The safe set X_{safe} and μ are ‘‘completely consistent’’, i.e., $\mu^{\ominus 1}(\mu(X_{\text{safe}})) = X_{\text{safe}}$.

Note that, by Lemma 4, $\mu^{\ominus 1}(\mu(X_{\text{safe}})) \supseteq X_{\text{safe}}$ is always true. Hence the nontrivial part of Assumption 2 is that $\mu^{\ominus 1}(\mu(X_{\text{safe}})) \subseteq X_{\text{safe}}$ also holds.

Assumption 2 may look a bit strong at first sight. However, without this assumption, there will be unsafe states that are indistinguishable from the safe states in any case. Moreover, Assumption 2 holds for the following setting that is widely considered in the literature: X_{safe} is a convex and bounded set and $\mu(x) = x \oplus V$ where V is a convex bounded set of admissible noises [14].

Under Assumption 2, Algorithm 1 returns a winning set. We summarize this result with the following theorem.

Theorem 1: If Assumption 2 holds, then set C_1 and controller π be generated by Algorithm 1 is contained by X_{safe} and is a winning set under π .

Proof: First, we have $C_1 \subseteq X_{\text{safe}}$ because

$$\begin{aligned} C_1 &= \mu^{\ominus 1}(\hat{C}) && \text{(definition of } C_1) \\ &\subseteq \mu^{\ominus 1}(\mu(X_{\text{safe}})) && (\hat{C} \subseteq \mu(X_{\text{safe}})) \\ &= X_{\text{safe}} && \text{(Assumption 2)} \end{aligned} \quad (14)$$

Next, we show that the winning condition (3) also holds for set C_1 . Let $(\mathbf{x}, \hat{\mathbf{x}}) \in \mathfrak{B}(\Sigma_\pi, C_1)$ and $(\mathbf{x}', \hat{\mathbf{x}}') \in \mathfrak{B}(\Sigma_\pi, X_{\text{safe}})$ be such that $\hat{x}_0 = \hat{x}'_0$. First, since $x_0 \in C_1$, we have

$$\hat{x}'_0 \in \hat{C} \quad (15)$$

because $\hat{x}'_0 = \hat{x}_0 \in \mu(x_0) \subseteq \hat{C}$ by definition of C_1 . Clearly

$$x'_0 \in \mu^{\oplus -1}(\hat{x}'_0) \quad (16)$$

also holds by definition of \mathfrak{B} . Combining Eq. (15),(16) and applying Proposition 1 shows that $x'_1 \in \tau(x'_0, \pi(\hat{x}'_0)) \subseteq C_1 \subseteq X_{\text{safe}}$, and this implies that $\hat{x}'_t \in \hat{C} \subseteq X_{\text{safe}}$ for all $t \geq 1$ by an inductive argument. Finally, note that $x'_0 \in X_{\text{safe}}$, hence $\mathbf{x}' \in X_{\text{safe}}^\omega$ and this verifies the winning condition. ■

The winning set C_1 obtained by the presented approach may not be maximal in general. In fact, as suggested by Example 2, sometimes C_1 can be empty while the maximal winning set W_{max} found via power set construction is nonempty. This is the case whenever the necessary contraction suggested by Proposition 1 cannot be fulfilled for any nonempty subset of X_{safe} . Hence, in general, it is impossible to quantify the level of conservatism of Algorithm 1 unless a contracting subset can be found. In the following theorem, we will show that finding such contracting subset that is nonempty is also sufficient for C_1 to be nonempty. With this result, quantifying the conservatism essentially amounts to asking what is the gap between W_{max} and the largest subset of X_{safe} that can be made contracting enough to overcome the effect of the noise.

Theorem 2: (Characterization of Set C_1) We say a set $\tilde{C}_1 \subseteq X_{\text{safe}}$ is ‘‘contracting enough w.r.t. τ and μ to overcome the noise’’ if

$$\forall \hat{x} \in \mu(\tilde{C}_1) : \exists u \in U : \forall x \in \mu^{\oplus -1}(\hat{x}) : \tau(x, u) \subseteq \tilde{C}_1, \quad (17)$$

If Assumption 2 holds, set C_1 returned by Algorithm 1 is the largest (in set inclusion sense) subset of X_{safe} that satisfies Eq. (17).

Proof: We first show that $\tilde{C}_1 \subseteq C_1$ holds for any \tilde{C}_1 that satisfies Eq. (17). Note that

$$\tilde{C}_1 \subseteq X_{\text{safe}} \Rightarrow \mu(\tilde{C}_1) \subseteq \mu(X_{\text{safe}}). \quad (18)$$

By Eq. (17), we also have

$$\forall \hat{x} \in \mu(\tilde{C}_1) : \exists u \in U : \tau(\mu^{\oplus -1}(\hat{x}), u) \subseteq \tilde{C}_1, \quad (19)$$

which implies

$$\forall \hat{x} \in \mu(\tilde{C}_1) : \exists u \in U : \mu(\tau(\mu^{\oplus -1}(\hat{x}), u)) \subseteq \mu(\tilde{C}_1). \quad (20)$$

that is, $\mu(\tilde{C}_1)$ satisfies Eq. (9). But note that, by construction, \hat{C} in Algorithm 1 is the largest subset of $\mu(X_{\text{safe}})$ that satisfies the condition in Eq. (9), hence

$$\mu(\tilde{C}_1) \subseteq \hat{C}. \quad (21)$$

Finally,

$$\begin{aligned} \tilde{C}_1 &\subseteq \mu^{\ominus 1}(\mu(\tilde{C}_1)) && \text{(Lemma 4)} \\ &\subseteq \mu^{\ominus 1}(\hat{C}) && \text{(apply } \mu^{\ominus 1} \text{ to Eq. (21))} \\ &= C_1. && \text{(definition of } C_1) \end{aligned} \quad (22)$$

We now show that C_1 satisfies Eq. (17). First, under Assumption 2, $C_1 \subseteq X_{\text{safe}}$ (see the proof of Theorem 1). Moreover, we have

$$C_1 = \mu^{\ominus 1}(\hat{C}) \Rightarrow \mu(C_1) \subseteq \hat{C} \quad (\text{Lemma 1}). \quad (23)$$

Eq. (23) together with Eq. (10) in Proposition 1 immediately implies the condition in Eq. (17). \blacksquare

2) *Computing Weakly Winning Set with Exactly Known Initial State:* In this part, we present Algorithm 2 that solves a problem slightly different from the one solved by Algorithm 1. In particular, we consider the case where the initial state of the system is known exactly. This is the case, for example, when one can set the initial state before starting the system. In this setting, it makes sense to consider the notion of weakly winning set.

Algorithm 2 $[C_2, \bar{\pi}] = \text{Win}_2(X_{\text{safe}}, \Sigma)$

- 1: Find a set $\check{C} \subseteq \mu^{\ominus 1}(X_{\text{safe}})$ and a controller $\check{\pi} : \check{C} \rightarrow U$ such that

$$\forall \check{x} \in \check{C} : \mu\left(\tau(\mu^{\oplus 1}(\check{x}), \check{\pi}(\check{x}))\right) \subseteq \check{C}. \quad (24)$$

- 2: $C_2 := \mu^{\ominus 1}(\check{C})$
 - 3: π is s.t. $\pi(\hat{\mathbf{x}}_t, \mathbf{u}_{t-1}) := \check{\pi}(\hat{x}_t)$
 - 4: **return** C_2, π
-

We use $\pi(\hat{x}_t)$ as a short notation for $\pi(\hat{\mathbf{x}}_t, \mathbf{u}_{t-1})$, and give the following result regarding Algorithm 2.

Proposition 2: Let C_2 and π be returned by Algorithm 2,

$$\forall x \in C_2 : \exists \check{x} \in \check{C} : \tau(x, \pi(\check{x})) \subseteq \mu^{\ominus 1}(\check{C}) \subseteq C_2. \quad (25)$$

In Eq. (25), \check{x} is called a ‘‘pseudo measurement’’ of x .

Proof: Let $x \in C_2$ be arbitrary, we have

$$\begin{aligned} x \in C_2 &\Leftrightarrow x \in \mu^{\oplus 1}(\check{C}) \\ &\Leftrightarrow x \in \bigcup_{\check{x} \in \check{C}} \mu^{\oplus 1}(\check{x}) \\ &\Leftrightarrow \exists \check{x} \in \check{C} : x \in \mu^{\oplus 1}(\check{x}) \end{aligned} \quad (26)$$

First, $x \in \mu^{\oplus 1}(\check{x})$ implies

$$\tau(x, \pi(\check{x})) \subseteq \tau(\mu^{\oplus 1}(\check{x}), \pi(\check{x})). \quad (27)$$

Secondly, since $\check{x} \in \check{C}$, we have

$$\begin{aligned} &\mu\left(\tau(\mu^{\oplus 1}(\check{x}), \pi(\check{x}))\right) \subseteq \check{C} \quad (\text{construction of } \check{C}). \\ \Rightarrow &\tau(\mu^{\oplus 1}(\check{x}), \pi(\check{x})) \subseteq \mu^{\ominus 1}(\check{C}). \end{aligned} \quad (\text{Lemma 1}) \quad (28)$$

Combining Eq. (27), (28) yields

$$\tau(x, \pi(\check{x})) \subseteq \tau(\mu^{\oplus 1}(\check{x}), \pi(\check{x})) \subseteq \mu^{\ominus 1}(\check{C}). \quad (29)$$

Finally, by Lemma 3, we have

$$\mu^{\ominus 1}(\check{C}) \subseteq \mu^{\oplus 1}(\check{C}) = C_2, \quad (30)$$

which completes the proof. \blacksquare

Remark 5: Note that C_2 is contracting under π in the sense that $\tau(x, \pi(\check{x})) \subseteq \mu^{\ominus 1}(\check{C}) \subseteq C_2$ for all $x \in C_2$. Very similar to the case in Proposition 1, we need such extra contraction to overcome the effect of imperfect state measurements. In fact, similar to Theorem 2, we can prove that set \check{C} is the largest (in set inclusion sense) subset of $\mu^{\ominus 1}(X_{\text{safe}})$ that satisfies certain contraction condition, i.e.,

$$\forall x \in \check{C} : \exists u \in U : \forall x' \in \mu^{\oplus 1}(\hat{x}) : \tau(x', u) \subseteq \mu^{\ominus 1}(\check{C}). \quad (31)$$

Remark 6: Proposition 2 does *not* say that C_2 is controlled invariant under π . Because whenever the true state x is initiated in C_2 , the actual measurement \hat{x} may not be necessarily equal to a pseudo measurement $\check{x} \in \check{C}$ that will lead to the right control input $\pi(\check{x})$ keeping the next true state in C_2 . To set $\hat{x} = \check{x}$, we need to know the true initial state x . This extra requirement is captured by the following assumption.

Assumption 3: The value of the true initial state is known.

Assumption 3 holds in the situations where we can initialize the state very accurately before starting the system.

Theorem 3: Let \check{C}, π, C_2 be from Algorithm 2, and let $x_0 \in C_2$ be the initial state, which is known by Assumption 3. Then C_2 is contained by the safe set X_{safe} and is controlled invariant (hence weakly winning) under $\bar{\pi}$ defined below

$$\begin{cases} \bar{\pi}(\hat{\mathbf{x}}_0) = \pi(\check{x}) \text{ where } \check{x} \in \mu(x_0) \cap \check{C} \\ \bar{\pi}(\hat{\mathbf{x}}_t, \mathbf{u}_t) = \pi(\hat{x}_t) \text{ when } t \geq 1 \end{cases}. \quad (32)$$

Proof: By Proposition 2 and Eq. (32), set C_2 is controlled invariant under the controller $\bar{\pi}$. Moreover,

$$\begin{aligned} C_2 &= \mu^{\oplus 1}(\check{C}) \\ &\subseteq \mu^{\oplus 1}(\mu^{\ominus 1}(X_{\text{safe}})) \quad (\check{C} \subseteq \mu^{\ominus 1}(X_{\text{safe}})) \\ &\subseteq X_{\text{safe}}. \end{aligned} \quad (\text{Lemma 4}) \quad (33)$$

Hence $C_2 \subseteq X_{\text{safe}}$ is controlled invariant. \blacksquare

Note that, to use controller $\bar{\pi}$ in set C_2 , we only need the exact initial state x_0 at the very beginning because we need to pick the pseudo measurement $\check{x}_0 \in \check{C} \cap \mu(x_0)$ to determine the right control input $u_0 \in \pi(\check{x}_0)$. However, it is enough to know the noisy measurement afterwards as the contracting condition in Eq. (31) will make sure $\hat{x}_t \in \check{C}$ for $t \geq 1$.

3) *Comparing Sets C_1 and C_2 :* We hereby compare sets C_1 and C_2 and the associated results. First, both sets require a certain amount of contraction to overcome the effect of the noise. Set C_1 is winning under a controller that only takes the latest noisy measurement as input, while rendering set C_2 safe also requires the exact knowledge of the initial state (i.e., Assumption 3). However, set C_1 is not necessarily a subset of the safe set X_{safe} unless Assumption 2 holds, whereas $C_2 \subseteq X_{\text{safe}}$ is true without any further assumptions.

Since set C_2 is rendered invariant under a controller that has more information than the controller associated with set C_1 , another natural question to ask is whether $C_1 \subseteq C_2$. It turns out that the two sets are not comparable in general. We will show this by providing two examples, where $C_1 \subsetneq C_2$ in one and $C_2 \subsetneq C_1$ in the other. In both of these examples, we consider discrete-time linear control systems

$$x_{t+1} = Ax_t + Bu_t + Ew_t, \quad (34)$$

$$\hat{x}_t = x_t + v_t. \quad (35)$$

where $w \in W$ is the disturbance and $v \in V$ is the measurement noise. Sets W , V , and X_{safe} are assumed to be polytopes (i.e., a bounded set of points from a Euclidean space that satisfies finitely many linear inequalities). This system can be written in the form in Eq. (1), (2) with $\tau(x_t, u_t) := Ax_t + Bu_t \oplus W$ and $\mu(x_t) := x_t \oplus V$. As pointed out earlier, it can be proved that Assumption 2 holds for $\mu(x) = x \oplus V$ and X_{safe} when V and X_{safe} are convex and bounded sets. Hence set $C_1 \subseteq X_{\text{safe}}$ by Theorem 1.

Example 4: Consider a 2-dimensional system in the form of Eq. (34), (35), with $E = [0 \ 0]^T$,

$$A = \begin{bmatrix} 0.9930 & 0.0358 \\ -0.2240 & 0.9930 \end{bmatrix}, \quad B = \begin{bmatrix} -0.0053 \\ 0.1205 \end{bmatrix}. \quad (36)$$

Here we assume that $X_{\text{safe}} = [-1.5, 1.5] \times [-1, 1]$, $U = [-1, 1]$, $W = \{0\}$ and $V = [-0.01, 0.01] \times \{0\}$. In this example, it can be verified that $\hat{C} = \check{C}$, and hence $C_1 = \hat{C} \ominus V \subsetneq \check{C} \oplus V = C_2$.

Example 5: Consider a 2-dimensional system, again, in the form of Eq. (34), (35), with $A = I$, $B = I$ and $E = [0 \ 0]^T$. $X_{\text{safe}} = [-5, 5] \times [-5, 5]$, $U = [-2, 2] \times [-2, 2]$, $W = \{0\}$ and $V = \{v \in \mathbb{R}^2 \mid \|d\|_1 \leq 1\}$. In this example, it can be verified that $C_1 = X_{\text{safe}}$, while $C_2 = (X_{\text{safe}} \ominus V) \oplus V \subsetneq X_{\text{safe}} = C_1$.

V. CASE STUDY

We consider an autonomous taxiing system for a Baron 58 aircraft that aims to implement a ‘‘cockpit over centerline’’ specification. That is, the controller aims to keep the aircraft on the taxiway within a pre-specified distance from the taxiway centerline. The system model is taken to be a 4-dimensional lateral dynamics model of the form $\dot{x} = Ax + B\delta_f + Er_d$ from [18]. The state $x = [y, v, \Delta\psi, v]$ consists of the lateral deviation y (m) from the centerline, the lateral velocity v (m/s), the yaw-angle deviation $\Delta\psi$ (rad) in centerline-fixed coordinates, and the yaw rate r (rad/s). The control input δ_f (rad) is the steering angle of the front gear, and r_d (rad) is the desired yaw rate computed from centerline curvature and treated as an external disturbance. The values of the A , B , E matrices are picked according to the data of a Baron 58 and can be found in [18]. To obtain a linear difference equation in the form of Eq. (34), (35), we discretize the continuous-time system with a sampling rate $\Delta t = 0.1$ s and a nominal longitudinal speed $v_0 = 5$ m/s.

We formulate this aircraft lateral control problem as a safety control problem. The ‘‘cockpit over centerline’’ specification and other comfort specifications are captured by

requiring the state $x = [y, v, \Delta\psi, r]$ to stay in a rectangular set $X_{\text{safe}} = [-1, 1] \times [-1, 1] \times [-0.2, 0.2] \times [-0.2, 0.2]$. The admissible set of control inputs and the disturbance set are taken to be $U = [-0.6981, 0.6981]$ and $W = [-0.02, 0.02]$.

In this example, a down-facing belly camera is used to get the estimates of the deviation (y) from the centerline and yaw angle ($\Delta\psi$). Some sample camera images generated by the X-plane simulator are shown in Fig. 4. The estimation is based on a line-segment detection algorithm using Hough transform [7] implemented in MATLAB (function `hough`). The boundaries of the yellow and black strips in the image are detected as line segments and are then used to approximate the centerline of the lane (Fig. 4, left). The yaw angle $\Delta\psi$ can be easily computed from this approximated centerline, and y is estimated as the distance from the approximated centerline to the aircraft’s center of mass. To make sure the black and yellow strips (or at least part of them) are in the field of view of the camera, we further restrict the state $x = [y, v, \Delta\psi, r]$ within $X_{\text{safe, vision}} = \{x \in X_{\text{safe}} \mid |0.444y + \Delta\psi| \leq 0.444\}$, whose edges are plotted with the red solid lines in Fig. 3. For most of the images generated by the X-plane simulator, the algorithm is able to achieve an estimation error ± 0.02 m in y and ± 0.007 rad (i.e., 0.4 deg) in $\Delta\psi$. We assume the other states are known exactly (e.g., from a gyroscope measurement) and define the admissible noise set $V = [-0.02, 0.02] \times \{0\} \times [-0.007, 0.007] \times \{0\}$.

We compute sets C_1 , C_2 for the aircraft taxiing system. Sets \hat{C} and \check{C} are computed by an implementation based on the algorithm proposed in [5]. The obtained sets are plotted in Fig. 3. The blue transparent polytope is the robust controlled invariant set with no measurement noise, the blue solid polytope is the set C_1 obtained by Algorithm 1. Set C_2 returned by Algorithm 2 is not plotted as it looks almost identical to C_1 . However, it can be checked numerically that $C_2 \subsetneq C_1$ in this case.

We simulate the closed-loop system with the controller induced by set C_1 and the vision-based state estimation module in the loop. At run time, the controller returns a set $\pi(\hat{x}_t) \subseteq U$ of control inputs that maintain safety, we further use a one step MPC controller to pick a control action from the set $\pi(\hat{x}_t)$ while minimizing the distance to the origin to improve tracking performance. Fig. 5 shows the simulation result. The plotted states (i.e., y and $\Delta\psi$) stay in their bounds, meanwhile the vision module is able to provide estimations close to the true values because the yellow and black strips will remain in the camera field of view by construction. It can be seen that the estimation error is larger (and sometimes larger than the allowable bounds used for controlled invariant set computation) whenever the state gets closer to the origin. This is because the yellow and black strips are partially occluded by the landing gear of the aircraft whenever the value of y and $\Delta\psi$ are small, which interfere the line detection algorithm. However, the safety is not likely to be violated in those cases because the system can actually tolerate larger measurement noise when the state is in the middle of the controlled invariant set.




Fig. 3: The safe set $X_{\text{safe,vision}}$ (red solid line skeleton) and the controlled invariant sets of the aircraft taxiing system with exact state measurement (blue transparent) and noisy state measurement (blue solid).




Fig. 4: Right: images generated in X-Plane via a down-facing camera attached to the belly of a Baron 58 aircraft right above the landing gear. Left: outcome of the line detection and centerline estimation algorithm.




Fig. 5: Simulation results.

VI. CONCLUSION

This paper presents theoretical results that enable efficient computation of controlled invariant sets in imperfect information settings. Such invariant sets are relevant for systems that only have access to inaccurate state information, like noisy measurements or outputs of perception modules, at run-time. We demonstrated how these invariant sets can be used for synthesizing safe-by-construction controllers for an autonomous taxiing system. These controllers guarantee that the closed-loop system stays safe (within a pre-specified distance from the taxiway centerline) as long as it starts in the computed set and the assumptions on the system and measurement models remain valid. For future work, we plan to investigate the use of these sets for generating corner cases, as is done in the perfect information setting in [4] and to derive conditions on perception modules for safety.

Acknowledgments: The authors would like to thank Prof. Jean-Baptiste Jeannin for useful discussions and Yuhao Zhang for his help with the example. This work is supported in part by Collins Aerospace, a unit of United Technologies

Corp. and ONR grant N00014-18-1-2501. The article solely reflects the opinions and conclusions of its authors and not Collins Aerospace.

REFERENCES

- [1] Z. Artstein and S. V. Raković. Set invariance under output feedback: a set-dynamics approach. *International Journal of Systems Science*, 42(4):539–555, 2011.
- [2] J.-P. Aubin. A survey of viability theory. *SIAM Journal on Control and Optimization*, 28(4):749–788, 1990.
- [3] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.
- [4] G. Chou, Y. E. Sahin, L. Yang, K. J. Rutledge, P. Nilsson, and N. Ozay. Using control synthesis to generate corner cases: A case study on autonomous driving. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37(11):2906–2917, 2018.
- [5] E. De Santis, M. D. Di Benedetto, and L. Berardi. Computation of maximal safe sets for switching systems. *IEEE Transactions on Automatic Control*, 49(2):184–195, 2004.
- [6] M. De Wulf, L. Doyen, and J.-F. Raskin. A lattice theory for solving games of imperfect information. In *International Workshop on Hybrid Systems: Computation and Control*, pages 153–168. Springer, 2006.
- [7] D. A. Forsyth and J. Ponce. *Computer vision: a modern approach*. Prentice Hall Professional Technical Reference, 2002.
- [8] T. Gurriet, P. Nilsson, A. Singletary, and A. D. Ames. Realizable set invariance conditions for cyber-physical systems. In *2019 American Control Conference (ACC)*, pages 3642–3649. IEEE, 2019.
- [9] J. Liu and N. Ozay. Finite abstractions with robustness margins for temporal logic-based control synthesis. *Nonlinear Analysis: Hybrid Systems*, 22:1–15, 2016.
- [10] R. Majumdar, N. Ozay, and A.-K. Schmuck. On abstraction-based controller design with output feedback. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, pages 1–11, 2020.
- [11] O. Mickelin, N. Ozay, and R. M. Murray. Synthesis of correct-by-construction control protocols for hybrid systems using partial state information. In *2014 American Control Conference*, pages 2305–2311. IEEE, 2014.
- [12] P. J. Ramadge and W. Wonham. Modular feedback logic for discrete event systems. *SIAM Journal on Control and Optimization*, 25(5):1202–1218, 1987.
- [13] S. Sadraddini and C. Belta. Formal methods for adaptive control of dynamical systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 1782–1787. IEEE, 2017.
- [14] R. Schneider. *Convex bodies: the Brunn–Minkowski theory*. Number 151. Cambridge university press, 2014.
- [15] L. Yang and N. Ozay. Fault-tolerant output-feedback path planning with temporal logic constraints. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4032–4039. IEEE, 2018.
- [16] X. Yin and S. Lafortune. Synthesis of maximally permissive supervisors for partially-observed discrete-event systems. *IEEE Transactions on Automatic Control*, 61(5):1239–1254, 2016.
- [17] K. Zhang, X. Yin, and M. Zamani. Opacity of nondeterministic transition systems: A (bi) simulation relation approach. *IEEE Transactions on Automatic Control*, 2019.
- [18] Y. Zhang, G. Poupart-Lafarge, H. Teng, J. Wilhelm, J.-B. Jeannin, N. Ozay, and E. Scholte. A software architecture for autonomous taxiing of aircraft. In *AIAA Scitech 2020 Forum*, page 0139, 2020.