Temporal Robustness of Stochastic Signals

Lars Lindemann*
larsl@seas.upenn.edu
University of Pennsylvania
Philadelphia, Pennsylvania, USA

Alëna Rodionova*
alena.rodionova@seas.upenn.edu
University of Pennsylvania
Philadelphia, Pennsylvania, USA

George J. Pappas
pappasg@seas.upenn.edu
University of Pennsylvania
Philadelphia, Pennsylvania, USA

ABSTRACT

We study the temporal robustness of stochastic signals. This topic is of particular interest in interleaving processes such as multi-agent systems where communication and individual agents induce timing uncertainty. For a deterministic signal and a given specification, we first introduce the synchronous and the asynchronous temporal robustness to quantify the signal's robustness with respect to synchronous and asynchronous time shifts in its sub-signals. We then define the temporal robustness risk by investigating the temporal robustness of the realizations of a stochastic signal. This definition can be interpreted as the risk associated with a stochastic signal to not satisfy a specification robustly in time. In this definition, general forms of specifications such as signal temporal logic specifications are permitted. We show how the temporal robustness risk is estimated from data for the value-at-risk. The usefulness of the temporal robustness risk is underlined by both theoretical and empirical evidence. In particular, we provide various numerical case studies including a T-intersection scenario in autonomous driving.

CCS CONCEPTS

• **Computer systems organization** → *Real-time system specification*; • **Mathematics of computing** → *Stochastic processes*; • **Theory of computation** → *Modal and temporal logics*; *Formal languages and automata theory*.

KEYWORDS

Risk-aware decision making; stochastic system verification; temporal robustness; signal temporal logic.

ACM Reference Format:
Lars Lindemann, Alëna Rodionova, and George J. Pappas. 2022. Temporal Robustness of Stochastic Signals. In *25th ACM International Conference on Hybrid Systems: Computation and Control (HSCC '22), May 4–6, 2022, Milan, Italy*. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3501710.3519504

1 INTRODUCTION

In this paper, we are interested in analyzing the robustness of time-critical systems, i.e., systems that need to satisfy stringent real-time constraints. Examples of time-critical systems include, but are not limited to, medical devices, autonomous driving, and air traffic control. When real-time specifications are interpreted over deterministic or stochastic signals, it is natural to study robustness with respect to timing uncertainty in the corresponding signals.

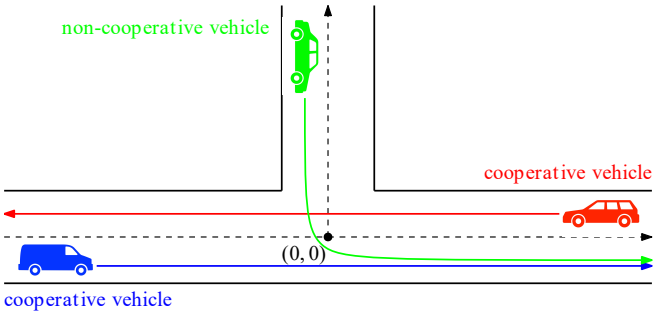


Figure 1: A T-intersection involving three cars that should pass the intersection at different times.

For instance, consider a cooperative driving scenario in which a group of autonomous vehicles communicate to pass through a blind intersection (see Figure 1 for an illustration of one of our case studies). Based on the transmitted information, a temporal order is assigned to the vehicles in which they must pass through the intersection. There may, however, be various reasons why an individual vehicle cannot comply with this predefined temporal order, e.g., due to delays or due to random changes in vehicle speed or the environment. In these cases, robustness of the predetermined temporal order to such signal timing uncertainties is crucial.

1.1 Related Work

Time-critical systems have been studied from various perspectives, and analysis tools for time-critical systems can be classified into three categories.

One category focuses on the study of real-time systems, see [26] and [30] for an overview. A particular interest in the study of real-time systems is in scheduling algorithms [41] that aim at finding an execution order for a set of tasks with corresponding deadline. Solutions consist, for instance, of periodic scheduling [40] or of event triggered scheduling [43]. The tools presented in our paper can be seen as complementary to these works as they allow to analyze the risk of scheduling algorithms.

The second category considers timed automata which allow to analyze real-time systems [2, 10]. Timed automata enable automatic system verification using model checking tools such as UPPAAL [8]. Robustness of timed automata was investigated in [21] and [9], while control of timed automata was considered in [33] and [5]. A

*Both authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
HSCC '22, May 4–6, 2022, Milan, Italy
© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9196-2/22/05…$15.00
https://doi.org/10.1145/3501710.3519504

connection between the aforementioned scheduling algorithms and timed automata was made in [18]. To the best of our knowledge, no one has yet considered the effect of stochastic timing uncertainty and the integration of axiomatic risk theory into a timed automata.

The last category considers formal specification languages such as real-time temporal logics. Metric interval temporal logic (MITL) [3] and signal temporal logic (STL) [32] enable specifying real-time constraints. It was shown in [3] and [28] that checking satisfiability of MITL and STL can be transformed into an emptiness checking problem of a timed automaton. Spatial robustness of MITL specifications over deterministic signals was considered in [17], while spatial robustness over stochastic signals was analyzed in [6] and [7]. Robust linear temporal logic was studied in [4]. Less attention has been spent on analyzing the temporal robustness of real-time systems subject to temporal logic specifications. Notions of system conformance were introduced in [1, 13, 19] to quantify closeness of systems in terms of spatial and temporal closeness of system trajectories. Conformance, however, only allows to reason about temporal robustness with respect to synchronous time shifts of a signal and not with respect to asynchronous time shifts in its sub-signals as captured by our definition of asynchronous temporal robustness. Our notion of asynchronous temporal robustness is especially important in multi-agent systems where indiviual agents may be delayed and subject to individual timing uncertainties. An orthogonal direction that is worth mentioning is followed by the authors in [36] and [24] to find temporal relaxations in control design for time window temporal logic specifications [44]. Temporal robustness for STL specifications over deterministic signals was presented in [14] and used for control design in [27] and in our recent work [38]. The notion of asynchronous temporal robustness as presented in this paper is fundamentally different and is interpretable in the sense that it quantifies the permissible timing uncertainty of a signal in terms of maximum time shifts in its sub-signals.

The interest in axiomatic risk theory is largely inspired by its longstanding successful use in finance, see e.g., [37]. More recently, risk measures were considered for decision making and control in robotics due to their systematic axiomatization [31]. The authors in [22] consider risk assessment in autonomous driving for motion prediction by deep neural networks, while [35] consider risk-aware planning for mobile robots in hospitals. Risk-constrained reinforcement and imitation learning was considered in [12] and [25], respectively. For signal temporal logic specifications, the use of risk was proposed in our recent work [29]. Risk-aware control was further considered in [42], [11], and [39] by using various risk measures, while risk-aware estimation was considered in [23].

We conclude with the observation that there is a set of influential works on the analysis and design of real-time systems, but that no one has yet studied the temporal robustness of stochastic signals in this context. Furthermore, no one has yet considered the integration of axiomatic risk theory into the analysis of real-time systems.

1.2 Contributions and Paper Outline

Our goal is to analyze the temporal robustness of stochastic signals, and to quantify the risk associated with the signal to not satisfy a specification robustly in time. We make the following contributions.

- We define the synchronous and the asynchronous temporal robustness with respect to synchronous and asynchronous time shifts in the sub-signals of a deterministic signal. This definition captures the permissible timing uncertainty.
- We define the synchronous and asynchronous temporal robustness risk of a stochastic signal. We permit the use of various risk measures such as the value-at-risk and show how the temporal robustness risk is estimated from data.
- We show that the temporal robustness risk of a stochastic signal under timing uncertainty is upper bounded by the sum of the temporal robustness risk of the nominal stochastic signal and the "maximum timing uncertainty".
- We extend the previous definitions to signal temporal logic (STL) specifications and define the STL temporal robustness risk.
- We provide empirical evidence and illustrate how the temporal robustness risk can be used in decision making and verification. We also show that, in scenarios where satisfying real-time constraints is of paramount importance, temporal robustness is advantageous over space robustness notions that are primarily considered in the literature.

Section 2 provides background on stochastic processes and risk measures. In Section 3, the synchronous and asynchronous temporal robustness is introduced. The temporal robustness risk is defined in Section 4, while the extension to STL specifications is presented in Section 5. Section 6 shows how the robustness risk can be estimated from data. Section 7 presents two case studies on cooperative driving and mobile robots. We conclude in Section 8.

2 BACKGROUND

Let $\mathbb{R}$ and $\mathbb{Z}$ be the set of real numbers and integers, respectively, and let $\mathbb{N}$ be the set of natural numbers including zero. Let $\mathbb{R}^n$ be the n-dimensional real vector space. We denote by $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$ and $\overline{\mathbb{Z}} := \mathbb{Z} \cup \{-\infty, \infty\}$ the extended real numbers and integers, respectively. Let $\mathfrak{F}(T, S)$ denote the set of all measurable functions mapping from the domain T into the domain S.

In this paper, we aim to define the temporal robustness risk of stochastic signals. Let us therefore first provide some brief background on stochastic processes and risk measures. We remark that all proofs of our technical results can be found in the appendix.

2.1 Random Variables and Stochastic Processes

Consider the *probability space* $(\Omega, \mathcal{F}, P)$ where Ω is the sample space, $\mathcal{F}$ is a σ-algebra of Ω, and $P : \mathcal{F} \to [0, 1]$ is a probability measure. Let Z denote a real-valued *random vector*, i.e., a measurable function $Z : \Omega \to \mathbb{R}^n$. When $n = 1$, we say Z is a *random variable*. We refer to $Z(\omega)$ as a realization of Z where $\omega \in \Omega$. Since Z is a measurable function, a probability space can be defined for Z so that probabilities can be assigned to events related to values of Z. Consequently, a cumulative distribution function (CDF) $F_Z(z)$ and probability density function (PDF) $f_Z(z)$ can be defined for Z.

A *stochastic process* is a function $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ where $X(t, \cdot)$ is a random vector for each fixed time $t \in \mathbb{Z}$. A stochastic process can be viewed as a collection of random vectors $\{X(t, \cdot) | t \in \mathbb{Z}\}$ defined on a common probability space $(\Omega, \mathcal{F}, P)$. For a fixed $\omega \in \Omega$, the function $X(\cdot, \omega)$ is a *realization* of the stochastic process. Another equivalent

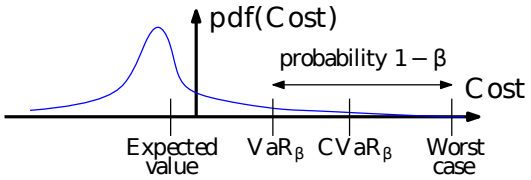


Figure 2: Illustration of various risk measures.

definition is that a stochastic process is a collection of deterministic functions of time that are indexed by Ω as $\{X(\cdot,\omega)|\omega \in \Omega\}$.

2.2 Risk Measures

A *risk measure* is a function $R : \mathfrak{F}(\Omega,\mathbb{R}) \to \mathbb{R}$ that maps from the set of real-valued random variables to the real numbers. In particular, we refer to the input of a risk measure R as the *cost random variable*. Risk measures allow for a risk assessment in terms of such cost random variables. There exists various risk measures, while Figure 2 particularly illustrates the expected value, the value-at-risk VaR_β, and the conditional value-at-risk $CVaR_\beta$ at risk level $\beta \in (0,1)$ which are commonly used risk measures. The VaR_β of a random variable $Z \in \mathfrak{F}(\Omega,\mathbb{R})$ is defined as

$$VaR_\beta(Z) := \inf\{\alpha \in \mathbb{R} | F_Z(\alpha) \geq \beta\},$$

i.e., the $1-\beta$ quantile of Z. The $CVaR_\beta$ of Z is defined as

$$CVaR_\beta(Z) := \inf_{\alpha\in\mathbb{R}} \{\alpha + (1-\beta)^{-1}E([Z-\alpha]^+)\}$$

where $[Z-\alpha]^+ := \max(Z-\alpha, 0)$. When the CDF F_Z of Z is continuous, it holds that $CVaR_\beta(Z) := E(Z|Z \geq VaR_\beta(Z))$, i.e., $CVaR_\beta(Z)$ is the expected value of Z conditioned on the events where Z is greater or equal than $VaR_\beta(Z)$.

Towards an axiomatic risk theory, various properties of risk measures are of interest, see [31] for an overview. While different properties may be desired in certain applications, we particularly use the monotonicity and translational invariance property.

- For two cost random variables $Z, Z' \in \mathfrak{F}(\Omega,\mathbb{R})$, the risk measure R is *monotone* if
$$Z(\omega) \leq Z'(\omega) \text{ for all } \omega \in \Omega \implies R(Z) \leq R(Z').$$
- For a random variable $Z \in \mathfrak{F}(\Omega,\mathbb{R})$, the risk measure R is *translationally invariant* if, for all $c \in \mathbb{R}$, it holds that
$$R(Z+c) = R(Z) + c.$$

Both $VaR_\beta(Z)$ and $CVaR_\beta(Z)$ satisfy these two properties.

3 TEMPORAL ROBUSTNESS OF DETERMINISTIC SIGNALS

We introduce two notions of temporal robustness to quantify how robustly a deterministic signal $x : \mathbb{Z} \to \mathbb{R}^n$ satisfies a real-time constraint with respect to time shifts in x.[1] These notions are referred to as the *synchronous temporal robustness* and the *asynchronous temporal robustness*. We express real-time constraints as constraints $c(x(t),t) \geq 0$ for all $t \in \mathbb{Z}$ where $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ is a measurable constraint function. Later in Section 5, we consider signal temporal logic (STL) as a more general means to express complex real-time constraints. The next example is used throughout this section.

Example 1. *Consider a signal x that consists of the two sub-signals*

$$x_{\{1\}}(t) := \sin(a\pi t)$$
$$x_{\{2\}}(t) := -b\sin(1.5a\pi t)$$

that are illustrated in red and blue in the left part of Figure 3.[2] *We require these two sub-signals to be ϵ-close within the time interval $[T_1,T_2]$ with parameters $T_1 := 145$, $T_2 := 155$, $\epsilon := 1$, $a := 0.04$ and $b = 1.05$.*[3] *This requirement is expressed by the constraint function*

$$c(x(t),t) := \begin{cases} 1 & \text{for all } t < T_1 \text{ and all } t > T_2 \\ h(x(t)) & \text{for all } t \in [T_1,T_2] \end{cases}$$

where $h : \mathbb{R}^2 \to \mathbb{R}$ encodes ϵ-closeness of $x_{\{1\}}$ and $x_{\{2\}}$ as

$$h(x(t)) := \epsilon - |x_{\{1\}}(t) - x_{\{2\}}(t)|.$$

so that $c(x(t),t) \geq 0$ for all $t \in \mathbb{Z}$ is equivalent to $h(x(t)) \geq 0$ for all $t \in [T_1,T_2]$. In the left part of Figure 3, the function $h(x(t))$ is illustrated in black and we can conclude that the requirement is satisfied, i.e., $c(x(t),t) \geq 0$ for all $t \in \mathbb{Z}$.

To indicate whether or not a signal x satisfies the constraint $c(x(t),t) \geq 0$ for all $t \in \mathbb{Z}$, let us define the satisfaction function $\beta^c : \mathfrak{F}(\mathbb{Z},\mathbb{R}^n) \to \{-1,1\}$ as

$$\beta^c(x) := \begin{cases} 1 & \text{if } \inf_{t\in\mathbb{Z}} c(x(t),t) \geq 0 \\ -1 & \text{otherwise.} \end{cases}$$

3.1 Synchronous Temporal Robustness

Let us first focus on synchronous time shifts across all sub-signals, i.e., each $x_{\{i\}}(t)$ in $x(t) = (x_{\{1\}}(t),\ldots,x_{\{n\}}(t))$ is shifted in time by the same amount. The *synchronous temporal robustness* is captured by the function $\eta^c : \mathfrak{F}(\mathbb{Z},\mathbb{R}^n) \to \overline{\mathbb{Z}}$ which we define as

$$\eta^c(x) := \beta^c(x)\sup\{\tau \in \mathbb{N} | \forall \kappa \in [-\tau,\tau], \beta^c(x_\kappa) = \beta^c(x)\}$$

where $x_\kappa(t) := x(t+\kappa)$, which is equivalent to x shifted by κ time units. Intuitively, $\eta^c(x)$ quantifies the maximum amount by which we can *synchronously* shift x in time while retaining the same behavior as x with respect to the constraint c. Note that $\eta^c(x)$ is non-negative (non-positive) if $\beta^c(x)$ is positive (negative) and that $\eta^c(x) > 0$ implies $\beta^c(x) = 1$ while $\eta^c(x) < 0$ implies $\beta^c(x) = -1$. The following result is a straightforward consequence of the definition of $\eta^c(x)$.

Corollary 1. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a constraint function. For $\xi \in \mathbb{Z}$, it holds that*

$$|\xi| \leq |\eta^c(x)| \implies \beta^c(x_\xi) = \beta^c(x).$$

Example 2. *Let us continue with Example 1 and take a look again at Figure 3. The synchronous temporal robustness is $\eta^c(x) = 12$, i.e., $x_{\{1\}}$ and $x_{\{2\}}$ can be shifted synchronously by 12 time units while still satisfying the constraint c. The middle part of Figure 3 shows the*

[1] We assume here that the time domain $\mathbb{Z}$ of the signal x is the set of integers, i.e., the time domain is unbounded in both directions. This assumption is made without loss of generality and in order to avoid technicalities.

[2] We choose the non-standard notation of $x_{\{i\}}$ to denote the ith sub-signal of x and reserve the notation of x_i for time shifted versions of x as defined in the remainder.

[3] As this paper is concerned with discrete-time signals, we implicitly assume that a time interval $[T_1, T_2]$ encodes $[T_1, T_2] \cap \mathbb{Z}$.

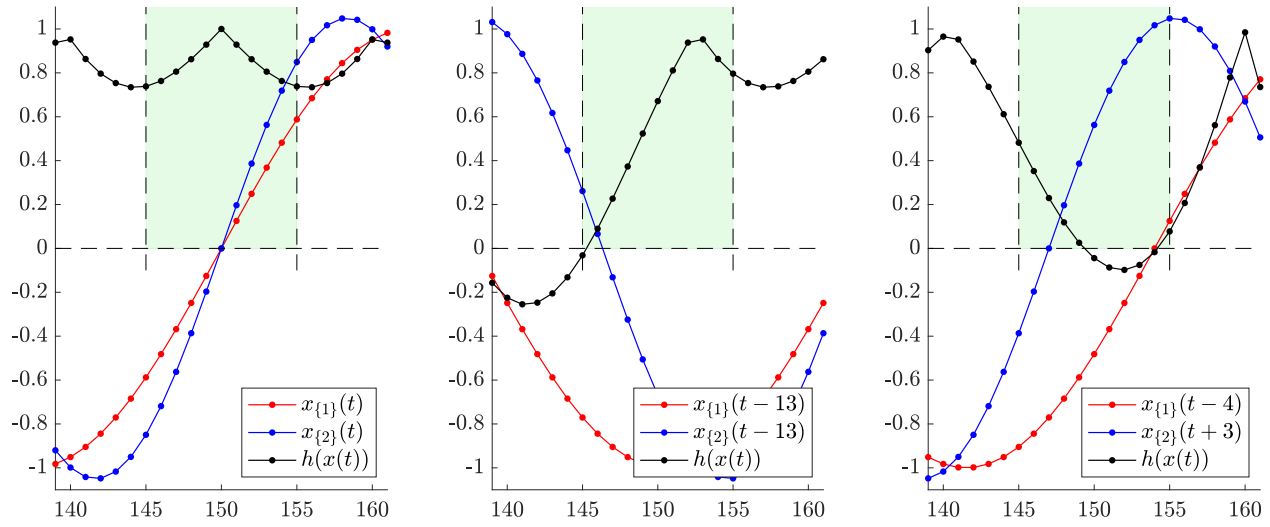


Figure 3: Example 1 in which the constraint c requires the two signals $x_{\{1\}}(t)$ and $x_{\{2\}}(t)$ to be ϵ-close within the time interval $[145, 155]$. Left: The constraint c is satisfied with synchronous temporal robustness $\eta^c(x) = 12$ and asynchronous temporal robustness $\theta^c(x) = 3$. Middle: The signals $x_{\{1\}}(t)$ and $x_{\{2\}}(t)$ are synchronously shifted by -13 so that the constraint c is violated. Right: The signals $x_{\{1\}}(t)$ and $x_{\{2\}}(t)$ are asynchronously shifted by -4 and 3, respectively, so that the constraint c is violated.

case where the signals are synchronously shifted by -13 time units, i.e., beyond this limit, so that the constraint c is violated.

We next provide a result that establishes a worst case lower bound for $\eta^c(x_{\bar{\xi}})$ based on $\eta^c(x)$ and $\bar{\xi}$. This result is needed later.

Lemma 1. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a constraint function. For all $\bar{\xi} \in \mathbb{Z}$, it holds that*

$$\eta^c(x_{\bar{\xi}}) \geq \eta^c(x) - |\bar{\xi}|.$$

3.2 Asynchronous Temporal Robustness

Besides the notion of synchronous temporal robustness, one can think of a notion with asynchronous time shifts across all subsignals, i.e., each $x_{\{i\}}(t)$ in $x(t) = (x_{\{1\}}(t), \dots, x_{\{n\}}(t))$ is allowed to be shifted in time by a different amount as opposed to being shifted by the same amount. We capture the *asynchronous temporal robustness* by the function $\theta^c : \mathfrak{F}(\mathbb{Z}, \mathbb{R}^n) \to \overline{\mathbb{Z}}$ which we define as

$$\theta^c(x) := \beta^c(x) \sup\{\tau \in \mathbb{N} | \forall \kappa_1, \dots, \kappa_n \in [-\tau, \tau], \beta^c(x_{\bar{\kappa}}) = \beta^c(x)\},$$

where $\bar{\kappa} := (\kappa_1, \dots, \kappa_n)$ and

$$x_{\bar{\kappa}}(t) := (x_{\{1\}}(t + \kappa_1), \dots, x_{\{n\}}(t + \kappa_n)).$$

The interpretation of $\theta^c(x)$ is that it quantifies the amount by which we can *asynchronously shift* x in time while retaining the same behavior as x with respect to the constraint c. The next corollary is a consequence of this definition and we particularly use the notation $\bar{\xi} := (\xi_1, \dots, \xi_n)$ for $\xi_1, \dots, \xi_n \in \mathbb{Z}$.

Corollary 2. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a constraint function. For $\xi_1, \dots, \xi_n \in \mathbb{Z}$, it holds that*

$$\max(|\xi_1|, \dots, |\xi_n|) \leq |\theta^c(x)| \implies \beta^c(x_{\bar{\xi}}) = \beta^c(x).$$

Example 3. *Recall Example 1 and look again at Figure 3. The asynchronous temporal robustness is $\theta^c(x) = 3$, i.e., $x_{\{1\}}$ and $x_{\{2\}}$ can be shifted asynchronously by 3 time units while still satisfying the constraint c. The right part of Figure 3 shows the case where the signals are shifted by -4 and 3, respectively, so that the constraint c is violated.*

We next state a result that is similar to Lemma 1 but for the asynchronous temporal robustness. It establishes a worst case lower bound for $\theta^c(x_{\bar{\xi}})$ based on $\theta^c(x)$ and $\bar{\xi}$.

Lemma 2. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a constraint function. For all $\xi_1, \dots, \xi_n \in \mathbb{Z}$, it holds that*

$$\theta^c(x_{\bar{\xi}}) \geq \theta^c(x) - \max(|\xi_1|, \dots, |\xi_n|).$$

Some remarks are in place regarding the presented notions and computational aspects.

Remark 1. *To calculate $\eta^c(x)$ and $\theta^c(x)$ in practice, we let $r \in \mathbb{N}$ be a bound on the maximum temporal robustness (in absolute value) that we are interested in. This allows us to replace the statement $\tau \in \mathbb{N}$ in the definitions of $\eta^c(x)$ and $\theta^c(x)$ by the statement $\tau \in [0, r]$. The complexity of calculating $\eta^c(x)$, measured in the number of shifted versions of x that one may need to construct, scales polynomially with r and is $O(r)$. In particular, at most $2r + 1$ shifted versions of x need to be constructed. On the other hand, the complexity of $\theta^c(x)$ scales polynomially in r and exponentially in the state dimension n and is $O(r^n)$, i.e., at most $(2r + 1)^n$ shifted versions of x need to be constructed.*

Remark 2. *There are various ways to reduce the complexity of calculating $\theta^c(x)$. First, observe that subsets of states of dynamical systems, e.g., position and velocity of a vehicle, are often subject to the same timing uncertainty. Under this assumption, we can group states in these subsets together and shift them synchronously which simplifies the calculation of $\theta^c(x)$. More explanation is provided Section 7. Second, a natural way to reduce the complexity in practice is to*

downsample the signal x. Third, depending on the coupling of states in $c(x(t), t)$ one may be able to decompose the calculation of $\theta^c(x)$ into smaller sub-problems that can be solved in parallel.

Remark 3. *These notions of temporal robustness focus on shifts of x in time, while one may also be interested in time scaling effects, e.g., $x(\tau \cdot t)$ for some $\tau \in \mathbb{N}$, which we are not considering in this paper.*

We conclude this section by comparing the two presented definitions. The asynchronous temporal robustness $\theta^c(x)$ is sensitive to asynchronous time shifts and hence more general than the synchronous temporal robustness $\eta^c(x)$. The synchronous temporal robustness, however, is easier to compute as argued in Remark 1 and provides an upper bound to $\theta^c(x)$ as shown next.

Corollary 3. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a constraint function. Then it holds that*

$$|\theta^c(x)| \leq |\eta^c(x)|.$$

4 TEMPORAL ROBUSTNESS OF STOCHASTIC SIGNALS

So far, we evaluated the constraint function $c(x(t), t)$ over signals $x : \mathbb{Z} \to \mathbb{R}^n$. We are now interested in evaluating $c(X(t, \cdot), t)$ over a stochastic process X. As c is a measurable function and $X(t, \cdot)$ is a random variable, it follows that $c(X(t, \cdot), t)$ is a random variable. We next show that $\eta^c(X)$ and $\theta^c(X)$ are also random variables.

Theorem 1. *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a measurable constraint function. Then $\eta^c(X(\cdot, \omega))$ and $\theta^c(X(\cdot, \omega))$ are measurable in ω, i.e., $\eta^c(X)$ and $\theta^c(X)$ are random variables.*

In general, it is desirable if the distributions of $\eta^c(X)$ and $\theta^c(X)$ are concentrated around larger values to be robust with respect to time shifts in X. This is why we view $-\eta^c(X)$ and $-\theta^c(X)$ as cost random variables. Following this interpretation, we define the *synchronous temporal robustness risk* and the *asynchronous temporal robustness risk* by applying risk measures R to $-\eta^c(X)$ and $-\theta^c(X)$.

Definition 1 (Synchronous and Asynchronous Temporal Robustness Risk). *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a measurable constraint function. The synchronous temporal robustness risk is defined as*

$$R(-\eta^c(X))$$

while the asynchronous temporal robustness risk is defined as

$$R(-\theta^c(X)).$$

The previous definitions quantify the risk of the stochastic process X to not satisfy the constraint c robustly with respect to synchronous and asynchronous time shifts in X. Note particularly that $R(-\eta^c(X)) < 0$ and $R(-\theta^c(X)) < 0$ can be given a sound interpretation, e.g., for the value-at-risk VaR_β at level β (the $1-\beta$ worst case quantile) the interpretation is that at most $1-\beta$ percent of the realizations of X violate the constraint c.

We next analyze the properties of the synchronous and the asynchronous temporal robustness risks. For $d \in \mathbb{N}$ and a shift random variable $\xi \in \mathfrak{F}(\Omega, [-d, d])$, let us now define the time shifted version X_ξ of X in analogy to x_ξ from the previous section as

$$X_\xi(t, \omega) := X(t + \xi(\omega), \omega).$$

In other words, X_ξ is equivalent to the stochastic process X synchronously shifted in time by the shift random variable ξ. The meaning of the synchronous temporal robustness risk with respect to the shift random variable ξ is shown in the next result.

Theorem 2. *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a measurable constraint function. Let the risk measure R be monotone and translationally invariant. For $d \in \mathbb{N}$ and a shift random variable $\xi \in \mathfrak{F}(\Omega, [-d, d])$, it holds that*

$$R(-\eta^c(X_\xi)) \leq R(-\eta^c(X)) + d.$$

The synchronous temporal robustness risk $R(-\eta^c(X_\xi))$ of X_ξ can hence not exceed the sum of the synchronous robustness risk $R(-\eta^c(X))$ of X and the bound d of the shift random variable ξ. Importantly, Theorem 2 tells us that

$$R(-\eta^c(X)) < 0 \quad \Longrightarrow \quad R(-\eta^c(X_\xi)) < 0$$

for all $\xi \in \mathfrak{F}(\Omega, [-d, d])$ where d is such that $d < |R(-\eta^c(X))|$. The next corollary specializes Theorem 2 to deterministic ξ.

Corollary 4. *Under the conditions of Theorem 2 and for $\xi \in \mathbb{Z}$, it holds that*

$$R(-\eta^c(X_\xi)) \leq R(-\eta^c(X)) + |\xi|.$$

Similar results and interpretations can be obtained for the asynchronous robustness risk. For $d \in \mathbb{N}$ and shift random variables $\xi_1, \ldots, \xi_n \in \mathfrak{F}(\Omega, [-d, d])$, define next $X_{\bar{\xi}}$ in analogy to $x_{\bar{\xi}}$ from the previous section as

$$X_{\bar{\xi}}(t, \omega) := (X_{\{1\}}(t + \xi_1(\omega), \omega), \ldots, X_{\{n\}}(t + \xi_n(\omega), \omega)).$$

We can now present the analogous result of Theorem 2.

Theorem 3. *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and $c : \mathbb{R}^n \times \mathbb{Z} \to \mathbb{R}$ be a measurable constraint function. Let the risk measure R be monotone and translationally invariant. For $d \in \mathbb{N}$ and shift random variables $\xi_1, \ldots, \xi_n \in \mathfrak{F}(\Omega, [-d, d])$, it holds that*

$$R(-\theta^c(X_{\bar{\xi}})) \leq R(-\theta^c(X)) + d.$$

The asynchronous temporal robustness risk $R(-\theta^c(X_{\bar{\xi}}))$ of $X_{\bar{\xi}}$ can hence not exceed the sum of the asynchronous robustness risk $R(-\theta^c(X))$ of X and the bound d of the shift random variables $\bar{\xi}$. Importantly, Theorem 3 now tells us that

$$R(-\theta^c(X)) < 0 \quad \Longrightarrow \quad R(-\theta^c(X_{\bar{\xi}})) < 0$$

for all $\xi_1, \ldots, \xi_n \in \mathfrak{F}(\Omega, [-d, d])$ where d is such that $d < |R(-\theta^c(X))|$. The next corollary specializes Theorem 3 to deterministic $\xi_1, \ldots, \xi_n$.

Corollary 5. *Under the conditions of Theorem 3 and for $\xi_1, \ldots, \xi_n \in \mathbb{Z}$, it holds that*

$$R(-\theta^c(X_{\bar{\xi}})) \leq R(-\theta^c(X)) + \max(|\xi_1|, \ldots, |\xi_n|).$$

5 THE STL TEMPORAL ROBUSTNESS RISK

We now extend the notion of synchronous and asynchronous temporal robustness to general specifications formulated in signal temporal logic (STL) introduced in [32]. STL specifications are constructed

from predicates $\mu : \mathbb{R}^n \to \{-1, 1\}$. Typically, these predicates are defined via predicate functions $h : \mathbb{R}^n \to \mathbb{R}$ as

$$\mu(\zeta) := \begin{cases} 1 & \text{if } h(\zeta) \geq 0 \\ -1 & \text{otherwise} \end{cases}$$

for $\zeta \in \mathbb{R}^n$. The syntax of STL is then recursively defined as

$$\phi \ ::= \ \top \mid \mu \mid \neg\phi \mid \phi' \wedge \phi'' \mid \phi' U_I \phi'' \mid \phi' \underline{U}_I \phi''$$

where $\top$ is the logical true element, ϕ' and ϕ'' are STL formulas and where U_I is the future until operator with $I \subseteq \mathbb{R}_{\geq 0}$, while $\underline{U}_I$ is the past until-operator. The operators $\neg$ and $\wedge$ encode negation and conjunction. One can further define the operators $\phi' \vee \phi'' := \neg(\neg\phi' \wedge \neg\phi'')$ (disjunction), $F_I\phi := \top U_I \phi$ (future eventually), $\underline{F}_I\phi := \top \underline{U}_I \phi$ (past eventually), $G_I\phi := \neg F_I \neg\phi$ (future always), $\underline{G}_I\phi := \neg \underline{F}_I \neg\phi$ (past always).

Semantics. We now define the satisfaction function $\beta^\phi : \mathfrak{F}(\mathbb{Z}, \mathbb{R}^n) \times \mathbb{Z} \to \{-1, 1\}$. In particular, $\beta^\phi(x, t) = 1$ indicates that the signal x satisfies the formula ϕ at time t, while $\beta^\phi(x, t) = -1$ indicates that x does not satisfy ϕ at time t. Following [32], the semantics $\beta^\phi(x, t)$ of an STL formula ϕ are inductively defined as

$$\begin{aligned}
\beta^\top(x, t) &:= 1, \\
\beta^\mu(x, t) &:= \begin{cases} 1 & \text{if } h(x(t)) \geq 0 \\ -1 & \text{otherwise,} \end{cases} \\
\beta^{\neg\phi}(x, t) &:= \neg\beta^\phi(x, t), \\
\beta^{\phi' \wedge \phi''}(x, t) &:= \min(\beta^{\phi'}(x, t), \beta^{\phi''}(x, t)), \\
\beta^{\phi' U_I \phi''}(x, t) &:= \sup_{t'' \in (t \oplus I) \cap \mathbb{Z}} \left(\min \left(\beta^{\phi''}(x, t''), \inf_{t' \in [t, t'']} \beta^{\phi'}(x, t') \right) \right), \\
\beta^{\phi' \underline{U}_I \phi''}(x, t) &:= \sup_{t'' \in (t \ominus I) \cap \mathbb{Z}} \left(\min \left(\beta^{\phi''}(x, t''), \inf_{t' \in [t'', t]} \beta^{\phi'}(x, t') \right) \right)
\end{aligned}$$

where $\oplus$ and $\ominus$ denote the Minkowski sum operator and the Minkowski difference operator, respectively.

STL Temporal Robustness. Besides calculating $\beta^\phi(x, t)$ to determine satisfaction of ϕ by x at time t, one can again calculate how robustly x satisfies ϕ at time t with respect to synchronous and asynchronous time shifts in x. The *synchronous STL temporal robustness* at time t is

$$\eta^\phi(x, t) := \beta^\phi(x, t) \sup\{\tau \in \mathbb{N} | \forall \kappa \in [-\tau, \tau], \beta^\phi(x_\kappa, t) = \beta^\phi(x, t)\}$$

while the *asynchronous STL temporal robustness* at time t is

$$\theta^\phi(x, t) := \beta^\phi(x, t) \sup\{\tau \in \mathbb{N} | \forall \kappa_1, \ldots, \kappa_n \in [-\tau, \tau], \beta^\phi(x_{\bar{\kappa}}, t) = \beta^\phi(x, t)\}.$$

For a fixed t, one can easily show that the same results presented for $\eta^c(x)$ and $\theta^c(x)$ in Corollaries 1 and 2 also hold for $\eta^\phi(x, t)$ and $\theta^\phi(x, t)$, respectively.

Corollary 6. *Let $x : \mathbb{Z} \to \mathbb{R}^n$ be a signal and ϕ be an STL formula. For $\xi \in \mathbb{Z}$, it holds that*

$$|\xi| \leq |\eta^\phi(x, t)| \implies \beta^\phi(x_\xi, t) = \beta^\phi(x, t).$$

Furthermore, for $\xi_1, \ldots, \xi_n \in \mathbb{Z}$ it holds that

$$\max(|\xi_1|, \ldots, |\xi_n|) \leq |\theta^\phi(x, t)| \implies \beta^\phi(x_{\bar{\xi}}, t) = \beta^\phi(x, t).$$

Remark 4. *The synchronous and asynchronous STL temporal robustness differ from the temporal robustness defined in [14]. We consider time shifts in x over ϕ via $\beta^\phi(x, t)$ instead of time shifts in x over each predicate μ in ϕ via $\beta^\mu(x, t)$. This enables Corollary 6 which does not hold for the definition in [14].*

STL Temporal Robustness Risk. We now interpret the STL formula ϕ over the stochastic process X instead of a deterministic signal x. In [29, Theorem 1], it was shown that $\beta^\phi(X, t)$ is a random variable. As in Theorem 1, it can hence be shown that $\eta^\phi(X, t)$ and $\theta^\phi(X, t)$ are random variables.

Definition 2 (Synchronous and Asynchronous STL Temporal Robustness Risk). *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and ϕ be an STL formula. The synchronous STL temporal robustness risk is defined as*

$$R(-\eta^\phi(X, t))$$

while the asynchronous STL temporal robustness risk is defined as

$$R(-\theta^\phi(X, t)).$$

The following result can be derived similarly to the results in Sections 4 and is stated without a proof.

Theorem 4. *Let $X : \mathbb{Z} \times \Omega \to \mathbb{R}^n$ be a discrete-time stochastic process and ϕ be an STL formula. Let the risk measure R be monotone and translationally invariant. For $d \in \mathbb{N}$ and a shift random variable $\xi \in \mathfrak{F}(\Omega, [-d, d])$, it holds that*

$$R(-\eta^\phi(X_\xi, t)) \leq R(-\eta^\phi(X, t)) + d.$$

For shift random variables $\xi_1, \ldots, \xi_n \in \mathfrak{F}(\Omega, [-d, d])$, it holds that

$$R(-\theta^\phi(X_{\bar{\xi}}, t)) \leq R(-\theta^\phi(X, t)) + d.$$

6 ESTIMATION OF THE TEMPORAL ROBUSTNESS RISK FROM DATA

Motivated by recent interest in data-driven verification, we show how the temporal robustness risk $R(-\eta^c(X))$ and $R(-\theta^c(X))$ as well as the STL temporal robustness risk $R(-\eta^\phi(X, t))$ and $R(-\theta^\phi(X, t))$ can be estimated from data. Let us, for convenience, define a random variable Z that corresponds to

$$Z \in \{-\eta^c(X), -\theta^c(X), -\eta^\phi(X, t), -\theta^\phi(X, t)\}$$

depending on the case of interest. For further convenience, let us define the tuple

$$\mathcal{Z} := (Z^1, \ldots, Z^N)$$

where again depending on the case of interest

$$Z^i \in \{-\eta^c(X^i), -\theta^c(X^i), -\eta^\phi(X^i, t), -\theta^\phi(X^i, t)\}$$

and where $X^1, \ldots, X^N$ are N observed realizations of X. Without loss of generality, assume that the tuple $\mathcal{Z}$ is sorted in increasing order, i.e., $Z^i \leq Z^{i+1}$ for each $i \in \{1, \ldots, N-1\}$.

We consider the value-at-risk (VaR) as a risk measure. For a risk level of $\beta \in (0, 1)$, recall that $VaR_\beta(Z) := \inf\{\alpha \in \mathbb{R} | F_Z(\alpha) \geq \beta\}$. Note now that the distribution function F_Z is discontinuous as we consider discrete-time stochastic processes X. To estimate $VaR_\beta(Z)$, we use [34, Lemma 3] to be able to deal with discontinuous distribution functions F_Z using order statistics. The following result follows from [34, Lemma 3] and is stated without proof.

Proposition 1. *Let $\delta \in (0,1)$ be a failure probability, $\beta \in (0,1)$ be a risk level, and $\gamma := \sqrt{\ln(2/\delta)/2N}$ be such that $\gamma \leq \min\{\beta, 1-\beta\}$. Then, with probability of at least $1-\delta$, it holds that*

$$\underline{VaR}_\beta(\mathcal{Z}) \leq VaR_\beta(Z) \leq \overline{VaR}_\beta(\mathcal{Z})$$

where the upper and lower bounds are defined as[4]

$$\overline{VaR}_\beta(\mathcal{Z}) := Z^{\lceil N(\beta+\gamma)\rceil}$$
$$\underline{VaR}_\beta(\mathcal{Z}) := Z^{\lfloor N(\beta-\gamma)\rfloor}.$$

In case that Z is a random variable with bounded support, we can derive similar estimators as in Proposition 1 for the conditional value-at-risk, see e.g., [45], which is omitted here for brevity.

7 CASE STUDIES

We present two case studies in this section. The first case study considers a T-intersection in a cooperative driving scenario with a constraint function c encoding collision avoidance at the intersection. The second case study considers two mobile robots under an STL specification ϕ encoding servicing tasks. Our code is available at https://tinyurl.com/temporalrob and animations of the two case studies can be found at https://tinyurl.com/temporalsim.

7.1 Cooperative Driving at a T-Intersection

This example is inspired by the Grand Cooperative Driving Challenge that aimed at communication-enabled cooperative autonomous driving [16]. We consider a T-intersection scenario where timing uncertainty, e.g., resulting from communication delays or communication errors caused by malfunctioning communication protocols, may lead to a violation of the specification.

Particularly, consider three cars at a T-intersection as illustrated in Figure 1. At time $t = 0$, each car is supposed to transmit their estimated positions. The green vehicle is non-cooperative, e.g., a human driver, while the blue and red vehicles are cooperative in the sense that they adjust their speeds based on the transmitted information to avoid collisions. For an intersection of size $\epsilon_1 := 10$ and a minimum safety distance of $\epsilon_2 := 15$, the safety specification is based upon

$$h_1(x) := \epsilon_1 - \|x_{\text{green}}\|$$
$$h_2(x) := \|x_{\text{green}} - x_{\text{red}}\| - \epsilon_2$$
$$h_3(x) := \|x_{\text{green}} - x_{\text{blue}}\| - \epsilon_2.$$

When $h_1(x) \geq 0$, the green car is within ϵ_1 radius of the intersection, while $h_2(x) \geq 0$ and $h_3(x) \geq 0$ indicate that a minimum safety distance of ϵ_2 is violated between cars. We consider the Boolean specification

$$(h_1(x) \geq 0) \implies [(h_2(x) \geq 0) \wedge (h_3(x) \geq 0)]$$

which is expressed by the constraint function

$$c(x,t) := \max(-h_1(x), \min(h_2(x), h_3(x))).$$

Let us assume initial positions of $x_{\text{green}}(0) := (-5, 300)$, $x_{\text{red}}(0) := (300, 5)$, and $x_{\text{blue}}(0) := (-300, -5)$. We have two scenarios S_1 and S_2. In both S_1 and S_2, the non-collaborative car is driving with velocity $v_{\text{green}} := 15$. In S_1, let $v_{\text{red}} := 12$ and $v_{\text{blue}} := 18$ so that first the blue, then the green, and then the red car pass the intersection. In S_2, let $v_{\text{red}} := 18$ and $v_{\text{blue}} := 12$ so that first the red, then the green, and the blue car pass the intersection. Note that we have chosen a simulation horizon of 60 time units and a time discretization of 0.1.[5] In both scenarios, we have $\beta^c(x) = 1$, i.e., the specification is satisfied, and $\eta^c(x) = +\infty$ for the synchronous temporal robustness. The asynchronous temporal robustness, however, is $\theta^c(x) = 10$ in S_1 and $\theta^c(x) = 8$ in S_2. This indicates that scenario S_1 is preferable from a temporal robustness point of view. Interestingly, the spatial robustness, here defined as $\inf_{t\in\mathbb{Z}} c(x(t), t)$, is 8.5 in S_1 and 16.87 in S_2 indicating that S_2 is preferable. This example highlights the difference between temporal and spatial robustness.

7.1.1 Communication Delays. Let us first model communication delays and errors. Let therefore the initial time of each car be random, i.e., let $x_{\text{green}}(t_0^{\text{g}}) := (-5, 300)$, $x_{\text{red}}(t_0^{\text{r}}) := (300, 5)$, and $x_{\text{blue}}(t_0^{\text{b}}) := (-300, -5)$ where t_0^{r}, t_0^{g}, and t_0^{b} are uniformly distributed within the interval $[-d, d] \subseteq \mathbb{Z}$. We consider the cases $d \in \{0, 8, 10, 12\}$, set $\delta := 0.01$ and $Z := -\theta^c(X)$, and collected $N := 10000$ realizations of X. We additionally calculate estimates of the conditional value-at-risk based on [37], denoted by $\widehat{CVaR}_\beta(\mathcal{Z})$.[6] The numerical results for these risk measures are shown in Table 1, while Figure 4 shows the histograms of $-\theta^c(X)$. We make the following observations:

- As lower values of $R(-\theta^c(X))$ imply less risk, scenario S_1 is preferable across all considered risk measures. This illustrates how the temporal robustness risk can be used to make risk-aware decisions.
- With S_1 and $d = 0$, we have that $R(-\theta^c(X)) = -10$. By Theorem 3, we expect $R(-\theta^c(X)) \leq 0$ for S_1 and $d = 10$. Indeed, this is shown empirically in the table below. Similarly, as $R(-\theta^c(X)) = -8$ for S_2 and $d = 0$, it holds that $R(-\theta^c(X)) \leq 0$ for S_2 and $d = 8$.
- The upper and lower bounds $\overline{VaR}_\beta$ and $\underline{VaR}_\beta$ are tight and provide a good approximation of VaR_β.
- The last column shows $\# := \sum_{i=1}^N \mathbb{I}(\beta^c(X^i) = -1)$, i.e., the number of realizations that violate c. Note that the numbers are in line with Corollary 2 and the interpretation of VaR.

Risk	$\overline{VaR}_{0.95}$	$\underline{VaR}_{0.95}$	$\overline{VaR}_{0.98}$	$\underline{VaR}_{0.98}$	$\widehat{CVaR}_{0.95}$	#
$S_1, d = 0$	-10	-10	-10	-10	-10	0
$S_1, d = 8$	-4	-5	-3	-4	-4	0
$S_1, d = 10$	-3	-4	-1	-3	-2.4	0
$S_1, d = 12$	-1	-2	0	-1	-0.8	57
$S_2, d = 0$	-8	-8	-8	-8	- 8	0
$S_2, d = 8$	-2	-3	-1	-2	-2	0
$S_2, d = 10$	-1	-2	0	-1	-0.5	78
$S_2, d = 12$	0	-1	2	0	0.5	404

Table 1: Risk estimates of $-\theta^c(X)$ where the initial time of each car is uniformly distributed within $[-d, d]$.

[4] We let $\lceil N(\beta+\gamma)\rceil$ and $\lfloor N(\beta-\gamma)\rfloor$ be the nearest integer to $N(\beta+\gamma)$ from above and the nearest integer to $N(\beta-\gamma)$ from below, respectively.

[5] To obtain signals that are defined on $\mathbb{Z}$, we extend all simulated trajectories that are defined on the interval [0, 60] to the left and right with the corresponding left and right endpoints of the trajectory.

[6] Note that $\widehat{E}$ and $\widehat{CVaR}$ are estimates and not an upper bound like $\overline{VaR}$.

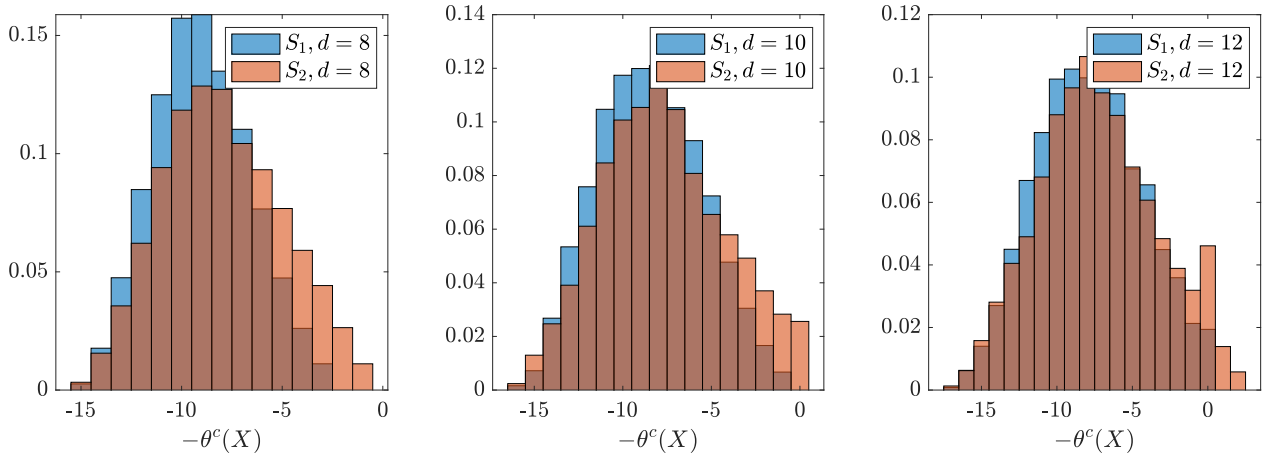


Figure 4: Histogram of $-\theta^c(X)$ in the T-intersection scenario in case of uncertainty in the communication between cars.

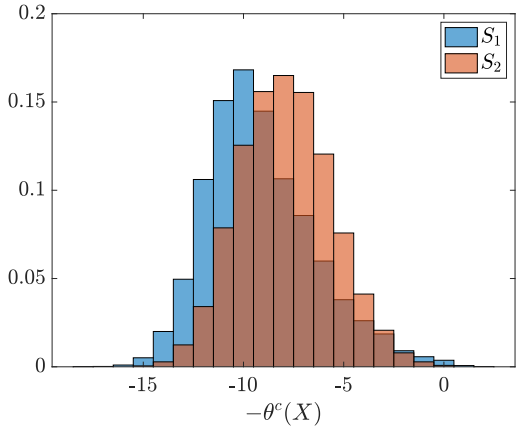


Figure 5: Histogram of $-\theta^c(X)$ in the T-intersection scenario in case of uncertainty in the cruise controllers.

7.1.2 Communication Delays and Control Disturbances. Usually, the cruise controller of a car is not exact. We model such errors by adding Gaussian distributed noise to v_{green}, v_{red}, and v_{blue}. For $Z := -\theta^c(X)$, $N := 10000$, and $\delta := 0.01$, the results for various risk measures and $-\theta^c(X)$ are shown in Table 2.

Risk	$\overline{VaR}_{0.85}$	$\overline{VaR}_{0.9}$	$\overline{VaR}_{0.95}$	$\overline{VaR}_{0.98}$	#
S_1	-6	-5	-3	0	20
S_2	-5	-5	-4	-2	1

Table 2: Risk estimates of $-\theta^c(X)$ where the velocity of each car is affected by additive Gaussian noise with variance 0.3^2.

Additionally, Figure 5 shows the histograms of $-\theta^c(X)$. Interestingly, the distribution of S_2 is Gaussian shaped, while the distribution of S_1 has a long tail of "bad" events. This is reflected in the calculated risks, i.e., for $\overline{VaR}_{0.85}$ and $\overline{VaR}_{0.9}$ the S_1 scenario is preferable. However, $\overline{VaR}_{0.95}$ and $\overline{VaR}_{0.98}$ detect the non-safe tail events in S_1 (the blue tail in Figure 5) and label S_2 as preferable. This is confirmed by the number of realizations that violate c in the rightmost column of the table and denoted by #.

So far, only the velocities were affected by noise in this subsection. To utilize the obtained information in terms of Corollary 5, we consider the starting times $t_0^{\text{r}} := -5$, $t_0^{\text{g}} := 5$, and $t_0^{\text{b}} := 2$. The results for $-\theta^c(X)$, aligning with Corollary 5, are shown in Table 3.

Risk	$\overline{VaR}_{0.85}$	$\overline{VaR}_{0.9}$	$\overline{VaR}_{0.95}$	$\overline{VaR}_{0.98}$	#
S_1	-1	0	1	4	604
S_2	-9	-8	-7	-5	1

Table 3: Risk estimates of $-\theta^c(X)$ where the velocity of each car is affected by additive Gaussian noise with variance 0.3^2 while the initial times of the red, green, and blue car are -5, 5, and 2, respectively.

Remark 5. *In this case, the dimension of x is $n = 6$ as each of the three cars consists of two states. As mentioned in Remark 2, we can reduce the computational complexity for calculating $\theta^c(x)$ under the assumption that the states of each car are subject to the same time shifts. In other words, the states of each car can be grouped and shifted synchronously so that the complexity is $O(r^3)$ instead of $O(r^6)$, which is a significant reduction. This way, the computation of $\theta^c(X^i)$ for a realization X^i only took around 0.25 s on a 1,4 GHz Intel Core i5.*

7.2 Cooperative Servicing

Consider two mobile robots employed in a cooperative servicing mission, see Figure 6 for an overview. The robots operate in a two-dimensional workspace and are described by the state $x_{\{i\}}$ consisting of its position p_i and its velocity v_i, i.e., $x_{\{i\}} := (p_i, v_i) \in \mathbb{R}^4$ for robot i. The state $x := (x_{\{1\}}, x_{\{2\}})$ is driven by discrete-time double integrator dynamics

$$x_{\{i\}}(t+1) := Ax_{\{i\}}(t) + Bu_i(t) \tag{1}$$

where $u_i : \mathbb{N} \to \mathbb{R}^2$ is the control input and where A and B are the discretized system and input matrices for two-dimensional double

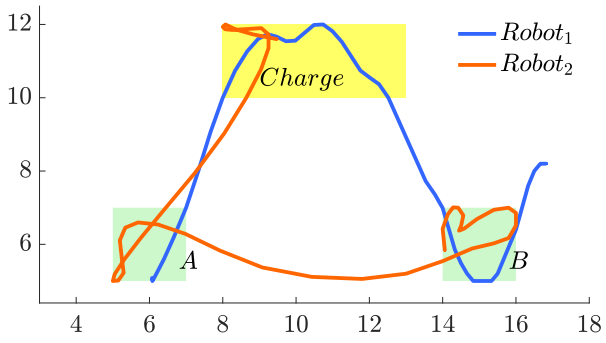


Figure 6: Nominal trajectories of the two robots under the control policies $u_i(t)$, i.e., system (1) without any time delays. The orange and blue stars denote initial conditions.

integrator dynamics under sampling time $t_s := 0.1$. Assume now that we are given the STL specification

$$\phi := F_{[0,1]}\big((p_1 \in A \vee p_2 \in A) \wedge F_{[1,5]}(p_1 \in A \vee p_2 \in A)\big) \tag{2a}$$
$$\wedge F_{[1,6]}(p_1 \in B \wedge p_2 \in B) \tag{2b}$$
$$\wedge F_{[0,2]}G_{[0,0.5]}(p_1 \in \text{Charge}) \wedge F_{[0,2]}G_{[0,0.5]}(p_2 \in \text{Charge}). \tag{2c}$$

The first line (2a) encodes a sequencing task, i.e., eventually within 1 time unit either robot 1 or robot 2 is in region A, followed by eventually within $[1, 5]$ time units either robot 1 or robot 2 is again in region A. The second line (2b) encodes that eventually within $[1, 6]$ time units both robots meet in region B. The third line (2c) encodes that both robots eventually within 2 time units visit the charging areas for at least 0.5 time units.

We first consider the deterministic system in (1) and synthesize a sequence of control inputs $u_i(t)$ for each robot i following the approach presented in [38]. For the deterministic system, we have that $\eta^\phi(X, 0) = 7$ time steps and $\theta^\phi(X, 0) = 6$ time steps. Note here the particular interplay between time units and the sampling time t_s, i.e., each time unit consists of 10 time steps.

We now delay the initial time of each robot by d_i similarly to the previous example, but here d_i follows a Poisson distribution with probability density function $\frac{\lambda_i^{d_i} \exp(-\lambda)}{d_i!}$ and parameter $\lambda_i > 0$. Our goal is now to verify the STL temporal robustness risk under the sequence of control inputs u_i. The results for $\eta^\phi(X, 0)$ are presented in Table 4 where λ_i corresponds to the Poisson parameter for d_i.

Risk	$\overline{VaR}_{0.85}$	$\overline{VaR}_{0.9}$	$\overline{VaR}_{0.95}$	$\overline{VaR}_{0.98}$	#
$\lambda_1 = \lambda_2 = 1$	-7	-7	-7	-6	0
$\lambda_1 = \lambda_2 = 3$	-5	-5	-4	-2	1
$\lambda_1 = \lambda_2 = 5$	-3	-2	-1	1	139
$\lambda_1 = \lambda_2 = 7$	0	0	1	4	1002
$\lambda_1 = 1, \lambda_2 = 5$	-7	-7	-7	-6	9
$\lambda_1 = 5, \lambda_2 = 1$	-3	-2	-1	1	137

Table 4: Risk estimates of $-\eta^\phi(X, 0)$ where the initial time of each robot is Poisson distributed with parameters λ_i.

Note here that similar general observation as in the T-intersection scenario can be made, i.e., lower risks are preferable so that naturally the cases with lower λ_i are preferable. The results for $\theta^\phi(X, 0)$ are presented in Table 5.

Risk	$\overline{VaR}_{0.85}$	$\overline{VaR}_{0.9}$	$\overline{VaR}_{0.95}$	$\overline{VaR}_{0.98}$	#
$\lambda_1 = \lambda_2 = 1$	-5	-5	-4	-4	0
$\lambda_1 = \lambda_2 = 3$	-4	-4	-3	-2	1
$\lambda_1 = \lambda_2 = 5$	-2	-2	-1	1	139
$\lambda_1 = \lambda_2 = 7$	0	0	1	4	1002
$\lambda_1 = 1, \lambda_2 = 5$	-2	-2	-1	0	9
$\lambda_1 = 5, \lambda_2 = 1$	-3	-2	-1	1	137

Table 5: Risk estimates of $-\theta^\phi(X, 0)$ where the initial time of each robot is Poisson distributed with parameters λ_i.

Tables 4 and 5 are according to Corollary 3 which can be shown to also hold for STL specifications ϕ, i.e., $|\theta^\phi(X, 0)| \leq |\eta^\phi(X, 0)|$. Notable are the cases $\lambda_1 = 1, \lambda_2 = 5$ and $\lambda_1 = 5, \lambda_2 = 1$. For the latter case, the bound $|\theta^\phi(X, 0)| \leq |\eta^\phi(X, 0)|$ proves to be conservative and shows the sensitivity of $\theta^\phi(X, 0)$ to asynchronous time shifts.

8 CONCLUSION

This paper studied temporal robustness of stochastic signals. We presented notions of temporal robustness for deterministic and stochastic signals. Our framework estimates the risk associated with a stochastic signal to not satisfy a specification robustly in time. We particularly permit general forms of risk measures to enable an axiomatic risk assessment. Case studies complemented the theoretical results highlighting the importance in applications and for risk-aware decision making and system verification.

ACKNOWLEDGMENTS

This research was supported by AFOSR grant FA9550-19-1-0265 and NSF award CPS-2038873.

REFERENCES

[1] Houssam Abbas, Hans Mittelmann, and Georgios Fainekos. 2014. Formal property verification in a conformance testing framework. In *Proceedings of the Conference on Formal Methods and Models for Codesign*. Lausanne, Switzerland, 155–164.
[2] Rajeev Alur and David L. Dill. 1994. A theory of timed automata. *Theoretical Computer Science* 126, 2 (1994), 183–235.
[3] Rajeev Alur and Thomas A Henzinger. 1996. The Benefits of Relaxing Punctuality. *Journal of the ACM* 43, 1 (1996), 116–146.
[4] Tzanis Anevlavis, Matthew Philippe, Daniel Neider, and Paulo Tabuada. 2022. Being correct is not enough: efficient verification using robust linear temporal logic. *ACM Transactions on Computational Logic* 23, 2 (2022), 1–39.
[5] Eugene Asarin, Oded Maler, Amir Pnueli, and Joseph Sifakis. 1998. Controller synthesis for timed automata. In *Proceedings of the Conference on System Structure and Control*. Nantes, France, 447–452.
[6] Ezio Bartocci, Luca Bortolussi, Laura Nenzi, and Guido Sanguinetti. 2013. On the Robustness of Temporal Properties for Stochastic Models. In *Proceedings of the Workshop on Hybrid Systems and Biology*, Vol. 125. Taormina, Italy, 3–19.
[7] Ezio Bartocci, Luca Bortolussi, Laura Nenzi, and Guido Sanguinetti. 2015. System design of stochastic models using robustness of temporal properties. *Theoretical Computer Science* 587 (2015), 3–25.
[8] Gerd Behrmann, Alexandre David, and Kim G. Larsen. 2004. A tutorial on Uppaal. *Formal methods for the design of real-time systems* (2004), 200–236.
[9] Jaroslav Bendík, Ahmet Sencan, Ebru Aydin Gol, and Ivana Černá. 2021. Timed Automata Robustness Analysis via Model Checking. *arXiv preprint arXiv:2108.08018* (2021).
[10] Johan Bengtsson and Wang Yi. 2003. Timed automata: Semantics, algorithms and tools. In *Advanced Course on Petri Nets*. Springer, 87–124.

[11] Margaret P. Chapman, Jonathan Lacotte, Aviv Tamar, Donggun Lee, Kevin M. Smith, Victoria Cheng, Jaime F. Fisac, Susmit Jha, Marco Pavone, and Claire J. Tomlin. 2019. A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems. In *Proceedings of the 2019 American Control Conference*. Philadelphia, PA, 2958–2963.

[12] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. 2017. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research* 18, 1 (2017), 6070–6120.

[13] Jyotirmoy V. Deshmukh, Rupak Majumdar, and Vinayak S. Prabhu. 2015. Quantifying conformance using the Skorokhod metric. In *Proceedings of the Conference on Computer Aided Verification*. San Francisco, CA, 234–250.

[14] Alexandre Donzé and Oded Maler. 2010. Robust Satisfaction of Temporal Logic over Real-valued Signals. In *Proceedings of the Conference on Formal Modeling and Analysis of Timed Systems*. Klosterneuburg, Austria, 92–106.

[15] Rick Durrett. 2019. *Probability: theory and examples*. Vol. 49. Cambridge university press.

[16] Cristofer Englund, Lei Chen, Jeroen Ploeg, Elham Semsar-Kazerooni, Alexey Voronov, Hoai Hoang Bengtsson, and Jonas Didoff. 2016. The grand cooperative driving challenge 2016: boosting the introduction of cooperative automated vehicles. *IEEE Wireless Communications* 23, 4 (2016), 146–152.

[17] Georgios E. Fainekos and George J. Pappas. 2009. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science* 410, 42 (2009), 4262–4291.

[18] Elena Fersman, Leonid Mokrushin, Paul Pettersson, and Wang Yi. 2006. Schedulability analysis of fixed-priority systems using timed automata. *Theoretical Computer Science* 354, 2 (2006), 301–317.

[19] Maciej Gazda and Mohammad Reza Mousavi. 2020. Logical characterisation of hybrid conformance. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*. Saarbrücken, Germany.

[20] A Hitchhiker's Guide. 2006. *Infinite dimensional analysis*. Springer.

[21] Vineet Gupta, Thomas A. Henzinger, and Radha Jagadeesan. 1997. Robust timed automata. In *Proceedings of the Workshop on Hybrid and Real-Time Systems*. Grenoble, France, 331–345.

[22] Ashkan Jasour, Xin Huang, Allen Wang, and Brian C Williams. 2021. Fast nonlinear risk assessment for autonomous vehicles using learned conditional probabilistic models of agent futures. *Autonomous Robots* (2021), 1–14.

[23] Dionysios S. Kalogerias, Luiz F. O. Chamon, George J. Pappas, and Alejandro Ribeiro. 2020. Better Safe Than Sorry: Risk-Aware Nonlinear Bayesian Estimation. In *Proceedings of the Conference on Acoustics, Speech and Signal Processing*. Barcelona, Spain, 5480–5484.

[24] Disha Kamale, Eleni Karyofylli, and Cristian-Ioan Vasile. 2021. Automata-based optimal planning with relaxed specifications. In *Proceedings of the International Conference on Intelligent Robots and Systems*. Prague, Czech Republic, 6525–6530.

[25] Jonathan Lacotte, Mohammad Ghavamzadeh, Yinlam Chow, and Marco Pavone. 2019. Risk-sensitive generative adversarial imitation learning. In *Proceedings of the Conference on Artificial Intelligence and Statistics*. Okinawa, Japan, 2154–2163.

[26] Phillip A. Laplante et al. 2004. *Real-time systems design and analysis*. Wiley New York.

[27] Zhenyu Lin and John S. Baras. 2020. Optimization-based Motion Planning and Runtime Monitoring for Robotic Agent with Space and Time Tolerances. In *Proceedings of the 21st IFAC World Congress*. Berlin, Germany, 1900–1905.

[28] Lars Lindemann and Dimos V. Dimarogonas. 2020. Efficient automata-based planning and control under spatio-temporal logic specifications. In *Proceedings of the 2020 American Control Conference*. Denver, CO, 4707–4714.

[29] Lars Lindemann, Nikolai Matni, and George J. Pappas. 2021. STL Robustness Risk over Discrete-Time Stochastic Processes. In *Proceedings of the Conference on Decision and Control*. Austin, Texas.

[30] Jane W. S. Liu. 2000. *Real-time systems design and analysis*. Prentice Hall.

[31] Anirudha Majumdar and Marco Pavone. 2020. How should a robot assess risk? Towards an axiomatic theory of risk in robotics. In *Robotics Research*. Springer, 75–84.

[32] Oded Maler and Dejan Nickovic. 2004. Monitoring temporal properties of continuous signals. In *Proceedings of the Conference on Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*. Grenoble, France, 152–166.

[33] Oded Maler, Amir Pnueli, and Joseph Sifakis. 1995. On the synthesis of discrete controllers for timed systems. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*. Munich, Germany, 229–242.

[34] Konstantinos E Nikolakakis, Dionysios S Kalogerias, Or Sheffet, and Anand D Sarwate. 2021. Quantile Multi-Armed Bandits: Optimal Best-Arm Identification and a Differentially Private Scheme. *IEEE Journal on Selected Areas in Information Theory* 2, 2 (2021), 534–548.

[35] Roya Sabbagh Novin, Amir Yazdani, Andrew Merryweather, and Tucker Hermans. 2021. Risk-Aware Decision Making for Service Robots to Minimize Risk of Patient Falls in Hospitals. In *Proceedings of the Conference on Robotics and Automation*. Xi'an, China, 3299–3305.

[36] Francisco Penedo, Cristian-Ioan Vasile, and Calin Belta. 2020. Language-guided sampling-based planning using temporal relaxation. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*. Helsinki, Finland, 128–143.

[37] R. Tyrrell Rockafellar and Stanislav Uryasev. 2000. Optimization of conditional value-at-risk. *Journal of risk* 2 (2000), 21–42.

[38] Alena Rodionova, Lars Lindemann, Manfred Morari, and George J. Pappas. 2021. Time-Robust Control for STL Specifications. In *Proceedings of the Conference on Decision and Control*. Austin, Texas, 1329–1335.

[39] Mathijs Schuurmans and Panagiotis Patrinos. 2020. Learning-based distributionally robust model predictive control of markovian switching systems with guaranteed stability and recursive feasibility. In *Proceedings of the Conference on Decision and Control*. Jeju Island, Republic of Korea, 4287–4292.

[40] Paolo Serafini and Walter Ukovich. 1989. A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics* 2, 4 (1989), 550–581.

[41] Lui Sha, Tarek Abdelzaher, Anton Cervin, Theodore Baker, Alan Burns, Giorgio Buttazzo, Marco Caccamo, John Lehoczky, and Aloysius K. Mok. 2004. Real time scheduling theory: A historical perspective. *Real-time systems* 28, 2 (2004), 101–155.

[42] Sumeet Singh, Yinlam Chow, Anirudha Majumdar, and Marco Pavone. 2018. A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms. *IEEE Transactions on Automatic Control* 64, 7 (2018), 2905–2912.

[43] Paulo Tabuada. 2007. Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Transactions on Automatic Control* 52, 9 (2007), 1680–1685.

[44] Cristian-Ioan Vasile, Derya Aksaray, and Calin Belta. 2017. Time window temporal logic. *Theoretical Computer Science* 691 (2017), 27–54.

[45] Ying Wang and Fuqing Gao. 2010. Deviation inequalities for an estimator of the conditional value-at-risk. *Operations Research Letters* 38, 3 (2010), 236–239.

A PROOF OF LEMMA 1

Assume that the synchronous temporal robustness is such that $|\eta^c(x)| = \zeta \geq 0$. By Corollary 1, it follows that

$$\forall \kappa \in [-\zeta, \zeta],\ \beta^c(x_\kappa) = \beta^c(x). \tag{3}$$

Now, take any $\xi \in \mathbb{Z}$. By definition, we have

$$\eta^c(x_\xi) := \beta^c(x_\xi) \sup\{\tau \in \mathbb{N} | \forall \kappa \in [-\tau, \tau], \beta^c(x_{\xi+\kappa}) = \beta^c(x_\xi)\}. \tag{4}$$

First, assume that $\beta^c(x_\xi) = 1$ so that $\eta^c(x_\xi) \geq 0$.

- For $|\xi| > \zeta$, it trivially follows that $\eta^c(x_\xi) \geq 0 > \zeta - |\xi| \geq \eta^c(x) - |\xi|$.
- For $|\xi| \leq \zeta$, it holds that $\beta^c(x) = \beta^c(x_\xi) = 1$ due to (3). From (3) it follows that $\forall \kappa \in [-\zeta + |\xi|, \zeta - |\xi|]$, $\beta^c(x_{\xi+\kappa}) = \beta^c(x) = \beta^c(x_\xi)$ since $\xi + \kappa \in [\xi - \zeta + |\xi|, \xi + \zeta - |\xi|] \subseteq [-\zeta, \zeta]$. Together with the definition of $\eta^c(x_\xi)$ in (4) this leads to $\eta^c(x_\xi) \geq \zeta - |\xi| \geq \eta^c(x) - |\xi|$.

Next, assume that $\beta^c(x_\xi) = -1$ so that $\eta^c(x_\xi) \leq 0$.

- For $\beta^c(x) = 1$, we have $\eta^c(x) = \zeta \geq 0$. This implies $\forall \kappa \in [-\zeta, \zeta]$, $\beta^c(x_\kappa) = 1$ by (3) so that, as we assume $\beta^c(x_\xi) = -1$, it has to hold that $|\xi| > \zeta$. Note that there exists a κ with $|\kappa| = ||\xi| - \zeta|$ so that $|\xi + \kappa| = \zeta$ and for which hence $\beta^c(x_{\xi+\kappa}) = 1$ so that $\eta^c(x_\xi) \geq -(||\xi| - \zeta| - 1) \geq -(|\xi| - \zeta) = \eta^c(x) - |\xi|$ by (4).
- For $\beta^c(x) = -1$, note that there exists $\kappa := -\xi + b \pm 1$ where $|b| = \zeta$ is such that $\beta^c(x_{\xi+\kappa}) = 1$ so that $\eta^c(x_\xi) \geq -(|\xi| + \zeta) = \eta^c(x) - |\xi|$ by (4).[7]

It hence follows that $\eta^c(x_\xi) \geq \eta^c(x) - |\xi|$.

B PROOF OF LEMMA 2

The proof follows similarly to the proof of Lemma 1. Assume that the asynchronous temporal robustness is such that $|\theta^c(x)| = \zeta \geq 0$. By Corollary 2, it follows that

$$\forall \kappa_1, \ldots, \kappa_n \in [-\zeta, \zeta],\ \beta^c(x_{\bar{\kappa}}) = \beta^c(x). \tag{5}$$

[7] The intuition of κ here is that one can shift x_ξ to x by $-\xi$ shifts first. There then has to exist a shift $b \pm 1$ of x so that $\beta^c(x_{\xi+\kappa}) = 1$.

Now, take any $\xi_1, \dots, \xi_n \in \mathbb{Z}$. By definition, we have

$$\theta^c(x_{\bar{\xi}}) := \beta^c(x_{\bar{\xi}}) \sup\{\tau \in \mathbb{N} | \forall \kappa_1, \dots, \kappa_n \in [-\tau, \tau], \beta^c(x_{\bar{\xi}+\bar{\kappa}}) = \beta^c(x_{\bar{\xi}})\} \tag{6}$$

where $\bar{\xi} + \bar{\kappa} := (\xi_1 + \kappa_1, \dots, \xi_n + \kappa_n)$. For brevity, let us also denote $\xi^* := \max(|\xi_1|, \dots, |\xi_n|)$.

First, assume that $\beta^c(x_{\bar{\xi}}) = 1$ so that $\theta^c(x_{\bar{\xi}}) \geq 0$.

- For $\xi^* > \zeta$, it trivially follows that $\theta^c(x_{\bar{\xi}}) \geq 0 > \zeta - \xi^* \geq \theta^c(x) - \xi^*$.
- For $\xi^* \leq \zeta$, it holds that $\beta^c(x) = \beta^c(x_{\bar{\xi}}) = 1$ due to (5). From (5) it follows that $\forall \kappa_i \in [-\zeta + \xi^*, \zeta - \xi^*], \beta^c(x_{\bar{\xi}+\bar{\kappa}}) = \beta^c(x) = \beta^c(x_{\bar{\xi}})$ since $\xi_i + \kappa_i \in [\xi_i - \zeta + \xi^*, \xi_i + \zeta - \xi^*] \subseteq [-\zeta, \zeta]$. Together with the definition of $\theta^c(x_{\bar{\xi}})$ in (6) this leads to $\theta^c(x_{\bar{\xi}}) \geq \zeta - \xi^* \geq \theta^c(x) - \xi^*$.

Next, assume that $\beta^c(x_{\bar{\xi}}) = -1$ so that $\theta^c(x_{\bar{\xi}}) \leq 0$.

- For $\beta^c(x) = 1$, we have $\theta^c(x) = \zeta \geq 0$. This implies that $\forall \kappa_1, \dots, \kappa_n \in [-\zeta, \zeta]$, $\beta^c(x_{\bar{\kappa}}) = 1$ by (5) so that, as we assume $\beta^c(x_{\bar{\xi}}) = -1$, it has to hold that $\xi^* > \zeta$. Note that there exists $\kappa_1, \dots, \kappa_n$ with $|\kappa_i| \leq |\xi^* - \zeta|$ so that $|\xi_i + \kappa_i| \leq \zeta$ and for which hence $\beta^c(x_{\bar{\xi}+\bar{\kappa}}) = 1$ so that $\theta^c(x_{\bar{\xi}}) \geq -(|\xi^* - \zeta| - 1) \geq -(\xi^* - \zeta) = \theta^c(x) - \xi^*$ by (6).
- For $\beta^c(x) = -1$, note that there exists $\kappa_i := a_i + b_i \pm 1$ where $|a_i| \leq \xi^*$ and $|b_i| \leq \zeta$ such that $\beta^c(x_{\bar{\xi}+\bar{\kappa}}) = 1$ so that $\theta^c(x_{\bar{\xi}}) \geq -(\xi^* + \zeta) = \theta^c(x) - \xi^*$ by (6).

It hence follows that $\theta^c(x_{\bar{\xi}}) \geq \theta^c(x) - \xi^*$.

C PROOF OF COROLLARY 3

Assume that the asynchronous temporal robustness is such that $|\theta^c(x)| = \zeta \geq 0$. By Corollary 2, we again have that (5) holds. Therefore, $\beta^c(x_{\bar{\kappa}}) = \beta^c(x)$ if $\kappa_1 = \dots = \kappa_n$ and $\kappa_1 \in [-\zeta, \zeta]$. Consequently, it holds $\beta^c(x_\kappa) = \beta^c(x)$ for all $\kappa \in [-\zeta, \zeta]$. By the definition of $\eta^c(x)$, it follows that $|\eta^c(x)| \geq \zeta = |\theta^c(x)|$.

D PROOF OF THEOREM 1

Step 1 - Measurability of $\beta^c(X(\cdot, \omega))$: Recall that

$$\beta^c(X(\cdot, \omega)) = \begin{cases} 1 & \text{if } \inf_{t \in \mathbb{Z}} c(X(t, \omega), t) \geq 0 \\ -1 & \text{otherwise} \end{cases}$$

where $c(X(t, \cdot), t)$ is a random variable as noted before, i.e., the function $c(X(t, \omega), t)$ is measurable in ω for each fixed $t \in \mathbb{Z}$. It holds that $\inf_{t \in \mathbb{Z}} c(X(t, \omega), t)$ is measurable in ω as the infimum over a countable number of measurable functions is again measurable, see e.g., [20, Theorem 4.27]. Note next that the characteristic function of a measurable set is measurable again (see e.g., [15, Chapter 1.2]) so that $\beta^c(X(\cdot, \omega))$ is measurable in ω.

Step 2 - Measurability of $\eta^c(X(\cdot, \omega))$: First, let $X_\kappa(t, \omega) := X(t + \kappa, \omega)$ in direct analogy to x_κ defined previously. It holds that $X_\kappa(t, \omega)$ is measurable in ω since $X(t + \kappa, \omega)$ is measurable in ω. Recall that

$$\eta^c(X(\cdot, \omega)) := \beta^c(X(\cdot, \omega)) \sup\{\tau \in \mathbb{N} | \forall \kappa \in [-\tau, \tau], \beta^c(X_\kappa(\cdot, \omega)) = \beta^c(X(\cdot, \omega))\}$$

Next, let us rewrite the supremum within the previous expression as $\sup_{\tau \in \mathbb{N}} g_\tau(X(\cdot, \omega))$ where each $g_\tau(X(\cdot, \omega))$ is such that

$$g_\tau(X(\cdot, \omega)) = \begin{cases} \tau & \text{if } \forall \kappa \in [-\tau, \tau], \beta^c(X_\kappa(\cdot, \omega)) = \beta^c(X(\cdot, \omega)) \\ 0 & \text{otherwise.} \end{cases}$$

In particular, the function g_τ is equivalent to

$$g_\tau(X(\cdot, \omega)) := \begin{cases} \tau & \text{if } \sum_{\kappa=-\tau}^{\tau} |\beta^c(X_\kappa(\cdot, \omega)) - \beta^c(X(\cdot, \omega))| = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Since $\sum_{\kappa=-\tau}^{\tau} |\beta^c(X_\kappa(\cdot, \omega)) - \beta^c(X(\cdot, \omega))|$ is measurable in ω and since the characteristic function of a measurable set is again measurable, it holds that $g_\tau(X(\cdot, \omega))$ is measurable in ω. It follows that $\sup_{\tau \in \mathbb{N}} g_\tau(X(\cdot, \omega))$ is measurable in ω as the supremum over a countable number of measurable functions is measurable, see e.g., [20, Theorem 4.27]. It consequently holds that $\eta^c(X(\cdot, \omega))$ is measurable in ω.

Step 3 - Measurability of $\theta^c(X(\cdot, \omega))$: Measurability of $\theta^c(X(\cdot, \omega))$ can be shown similarly to measurability of $\eta^c(X(\cdot, \omega))$ by changing the definition of g_τ to include a nested sum over $\kappa_1, \dots, \kappa_n$.

E PROOF OF THEOREM 2

Every realization $X(\cdot, \omega)$ for $\omega \in \Omega$ leads to a synchronous temporal robustness $\eta^c(X(\cdot, \omega))$. Additionally, every realization $\xi(\omega)$ for $\omega \in \Omega$ leads to a synchronous temporal robustness $\eta^c(X_{\bar{\xi}}(\cdot, \omega))$.[8] According to Lemma 1 and since $|\xi(\omega)| \leq d$, it holds that

$$\eta^c(X_{\bar{\xi}}(\cdot, \omega)) \geq \eta^c(X(\cdot, \omega)) - |\xi(\omega)| \geq \eta^c(X(\cdot, \omega)) - d$$

so that $-\eta^c(X_{\bar{\xi}}(\cdot, \omega)) \leq -\eta^c(X(\cdot, \omega)) + d$. As R is monotone and translationally invariant, it follows that

$$R(-\eta^c(X_{\bar{\xi}})) \leq R(-\eta^c(X) + d) = R(-\eta^c(X)) + d.$$

F PROOF OF THEOREM 3

The proof follows similarly to the proof of Theorem 2. According to Lemma 2 and as $\max(|\xi_1(\omega)|, \dots, |\xi_n(\omega)|) \leq d$, it holds that

$$\theta^c(X_{\bar{\xi}}(\cdot, \omega)) \geq \theta^c(X(\cdot, \omega)) - d$$

so that $-\theta^c(X_{\bar{\xi}}(\cdot, \omega)) \leq -\theta^c(X(\cdot, \omega)) + d$. As R is monotone and translationally invariant, it follows that

$$R(-\theta^c(X_{\bar{\xi}})) \leq R(-\theta^c(X) + d) = R(-\eta^c(X)) + d.$$

[8] Recall that $X_{\bar{\xi}}(t, \omega) = X(t + \xi(\omega), \omega)$. Note that it can be shown that $X_{\bar{\xi}}(\cdot, \omega)$ is measurable in $\omega \in \Omega$.