A Small Gain Theorem for Parametric Assume-Guarantee Contracts*

Eric S. Kim, Murat Arcak, Sanjit A. Seshia
{eskim, arcak, sseshia}@eecs.berkeley.edu
University of California at Berkeley, Berkeley, CA, USA
Department of Electrical Engineering and Computer Sciences

ABSTRACT

The problem of verifying properties of large, networked cyber-physical systems (CPS) is beyond the reach of most computational tools today. Two common "divide-and-conquer" techniques for CPS verification are assume-guarantee contracts from the formal methods literature and input-output properties from the control theory literature. Combining these two approaches, we first introduce the notion of a parametric assume-guarantee contract, which lets one reason about system behavior abstractly in a parameter domain. We next show how a finite gain property can be encoded in this form and provide a generalized small-gain theorem for parametric assume-guarantee contracts.

This theorem recovers the classical small gain theorem as a special case and its derivation highlights the connection between assume-guarantee reasoning and small-gain results. This new small-gain theorem applies to behaviors beyond bounded deviation from a nominal point to include a fragment of linear temporal logic with parametrized predicates that can encode safety, recurrence, and liveness properties. Our results are validated with an example which certifies that the interconnection of two freeway segments experiences intermittent congestion.

Keywords

Assume-Guarantee Contracts; Parametric Temporal Logic; Small Gain Theorem; Robustness

1. INTRODUCTION

Exploiting compositionality is a common procedure for the design and verification of systems consisting of a large number of interconnected components. Such techniques leverage higher level representations of component behavior to reason about the complete interconnected system's behavior. In the broader cyber-physical systems (CPS) literature, these higher level representations are commonly in the form of an input-output property or an assume-guarantee contract. The notion of assume-guarantee contract resembles the concepts of an input-output robustness property because both encode relationships between an environment and the behaviors exhibited by a system. Both theories are leveraged for compositional design and verification for an interconnection of systems.

We first introduce the notion of a parametric assume-guarantee contract. Contracts of this form permit us to write tighter guarantees on system behaviors as a response to the environment a system *actually* experiences once implemented and deployed; contracts without this form typically have coarse guarantees because assumptions need to account for all possible environments it *could* experience. Additionally, the parameterization of assumptions and guarantees permits us to reason about system behavior. A finite gain property can be encoded as a parametric assume-guarantee contract.

This paper's core technical contribution is to provide a small gain theorem for system behaviors satisfying a parametric assume-guarantee contract. The classical small-gain theorem, which establishes bounded input bounded output (BIBO) stability for a feedback interconnection is recovered as a special case. The new result opens a door to small gain-like results for a broader class of specifications beyond BIBO stability. This broader class of specifications requires a mild technical condition, which we show to also be satisfied by a fragment of linear temporal logic with parametrized predicates. This fragment may encode objectives such as safety, recurrence, and liveness properties.

We next discuss methods to verify that a system satisfies a parametric assume-guarantee contract. Our results are demonstrated with a freeway traffic flow example with hybrid dynamics. Two freeway segments are individually certified to have intermittent congestion. The concatenation is also certified to have intermittent congestion and a quantitative upper bound is established on its severity.

1.1 Related Work

The notions of input-output stability and robustness from the control theory literature have been extended to cyber-physical systems for both verification and controller synthesis. Definitions of robustness for systems with discrete input-output alphabets and associated verification algorithms were provided by Tarraf et al. [23] and Tabuada et al. [22]. Small-gain conditions are leveraged by Rungger et al. [21] for com-

*This work was supported in part by NSF grant CNS-1446145, the NSF Graduate Research Fellowship Program, NSF Expeditions grant CCF-1139138 and by STARnet, a Semiconductor Research Corporation program, sponsored by MARCO and DARPA.

HSCC'17, April 18 - 20, 2017, Pittsburgh, PA, USA

ISBN 978-1-4503-4590-3/17/04…$15.00

DOI: http://dx.doi.org/10.1145/3049797.3049805

positional construction of approximate discrete abstractions of continuous systems and Dallal et al. [6] to design customized compositional abstractions for a persistency specification. Majumdar et al. [16] argue that graceful degradation in performance in the presence of errors is a desirable property to enforce in systems. Bloem et al. [4] advocate for objectives with quantitative measures of "goodness" to distinguish between control strategies that both satisfy a Boolean specification. Indeed, robustness has been introduced into temporal logic to serve as a qualitative measure of satisfaction [10][9].

Compositionality in the formal methods literature has been approached through the notion of a contract which specifies a set of environments under which a component is guaranteed to exhibit a desired behavior [3]. Assume-guarantee contracts have been used for the verification of aircraft electrical systems [18], controller synthesis in traffic networks [13], and to certify stability for embedded systems in the presence of timing uncertainty [1].

To date, assume-guarantee contracts have been Boolean properties and do not incorporate notions of robustness as described above. To the authors' knowledge, there are no existing results that quantify how robust an interconnection of two systems is upon interconnection. The literature on assume-guarantee contracts with quantitative values primarly concerns verifying that a stochastic system satisfies a property with sufficiently high probability [14].

2. PRELIMINARIES

For a set $\mathcal{P}$, let $|\mathcal{P}|$, $2^{\mathcal{P}}$, $\mathcal{P} \times \mathcal{Q}$, $\mathcal{P}^*$, and $\mathcal{P}^\omega$ respectively represent $\mathcal{P}$'s cardinality, powerset (set of all subsets), Cartesian product with $\mathcal{Q}$, and sets of finite and infinite sequences of elements of $\mathcal{P}$. We let $\mapsto$ denote a functional map from a domain to codomain, and $\Longrightarrow$ represent Boolean implication. Boolean true and false are denoted by $\top$ and $\bot$. For two functions f, g with appropriate domains and codomains, $(f \circ g)(x)$ denotes the function composition $f(g(x))$. The Boolean negation of a proposition a is $\neg a$ and we have logical operations $\wedge$ (and/conjunction) and $\vee$ (or/disjunction). The implication $A \Longrightarrow B$ is equivalent to the logical statement $\neg A \vee B$; in other words, $A \Longrightarrow B$ is violated only when A is true and B is false.

Given a space $\mathcal{X}$ equipped with a distance metric $d : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}_{\geq 0}$, the closure of the set $\mathcal{L} \subseteq \mathcal{X}$ is denoted $\mathbf{cl}(\mathcal{L})$. The $\epsilon-$expansion of $A \subseteq \mathcal{X}$ for $\epsilon \geq 0$ is $\mathcal{B}_\epsilon(A) = \bigcup_{x \in A}\{y \in \mathcal{X} : d(x, y) \leq \epsilon\}$. A point x' is in the $\omega-$limit of a sequence $x[\cdot] = x[0]x[1]\ldots$ if and only if there exists a subsequence that converges to x'.

For a space $\mathcal{X}$ and an interval $I = [a, b]$ (where $a \leq b$, $a \in \mathbb{Z}_{\geq 0}, b \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$) the space of signals $\mathcal{X}[\cdot]$ is given by a Cartesian product indexed by elements of I:

$$\mathcal{X}[\cdot] = \prod_{k \in I} \mathcal{X}. \tag{1}$$

For a signal $x[\cdot] \in \mathcal{X}[\cdot]$, let $x[k]$ represent its value at time k.

In this paper, a dynamical system Σ is viewed as a relation between inputs and output signals $\Sigma \subseteq \mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$. We assume that any input $u[\cdot]$ is paired with at least one $y[\cdot]$ via the relation Σ. Such a $y[\cdot]$ is unique if Σ is deterministic and we say $y[\cdot] = \Sigma(u[\cdot])$ holds. If Σ is non-deterministic then $y[\cdot]$ is not necessarily unique and we say that $y[\cdot] \in \Sigma(u[\cdot])$.

3. SPECIFICATIONS

An input-output specification $\phi \subseteq \mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$ for a dynamical system is a logical statement describing a set of desirable input-output behaviors. An input specification is a subset of $\mathcal{U}[\cdot]$ and an output specification is a subset of $\mathcal{Y}[\cdot]$. We signify that an input (resp. output) signal $z[\cdot]$ *satisfies* an input (resp. output) specification ϕ by $z[\cdot] \models \phi$. A system Σ satisfies ϕ if $\Sigma \subseteq \phi$. A specification ϕ over $\mathcal{Z}[\cdot]$ is *satisfiable* if there exists a $z[\cdot] \in \mathcal{Z}[\cdot]$ such that $z[\cdot] \models \phi$. A specification ϕ has an evaluation time $|\phi|$, which represents the minimum amount of time to determine if a signal satisfies or violates ϕ. If specification $\phi \subseteq \mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$ is associated with the time interval $I = [a, b]$ then generally $|\phi| = b - a$. Specification ϕ has *bounded evaluation time* if $|\phi| < \infty$.

One can project from a Boolean view of the specification ϕ to a set point of view with $\phi = \{z[\cdot] \in \mathcal{Z}[\cdot] : z[\cdot] \models \phi\}$. As expected $z[\cdot] \models \phi$ if and only if $z[\cdot] \in \phi$ (using the set theoretic definition of ϕ). It is typically easier to manipulate specifications as logical objects, but some notation overloading between the set/logic points of view may occur and is pointed out when appropriate.

A parametric input-output specification $\psi : \mathcal{P} \mapsto 2^{\mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]}$ is a collection of specifications indexed by some parameter space $\mathcal{P}$. Parametric input specifications and parametric output specifications are defined analogously. For instance the input specification (in set form) $\psi(p) = \{u[\cdot] : \sqrt[p]{u[1]} > 4\}$ has parameter $p \in \mathbb{R}$. An example of a parametric output specification may be the set $\psi(p) = \{y[\cdot] : ||y[\cdot]||_2 < p\}$ of signals with Euclidean norm bounded by $p \in \mathbb{R}_{\geq 0}$.

We use the Hausdorff pseudo-metric to measure the difference between satisfiable specifications. Given a metric $d : \mathcal{X}[\cdot] \times \mathcal{X}[\cdot] \mapsto \mathbb{R}_{\geq 0} \cup \{\infty\}$ between signals, the Hausdorff between specifications ϕ_a and ϕ_b is:

$$\begin{aligned} d_H(\phi_a, \phi_b) &:= \inf\{\epsilon \geq 0 : \phi_a \subseteq \mathcal{B}_\epsilon(\phi_b) \text{ and } \phi_b \subseteq \mathcal{B}_\epsilon(\phi_a)\} \\ &:= \max\{\sup_{a \in \phi_a} \inf_{b \in \phi_b} d(a[\cdot], b[\cdot]), \sup_{b \in \phi_b} \inf_{a \in \phi_a} d(a[\cdot], b[\cdot])\} \end{aligned}$$

If $d_H(\phi_a, \phi_b) < \epsilon$, then for each signal that satisfies ϕ_a is at most ϵ for some signal that satisfies ϕ_b and vice versa. The Hausdorff distance can assume infinite values and is a pseudo-metric because $d_H(\phi_a, \phi_b) = 0$ implies $\mathbf{cl}(\phi_a) = \mathbf{cl}(\phi_b)$ rather than $\phi_a = \phi_b$. A parametric specification with a metric-equipped parameter space is Hausdorff continuous if it satisfies the standard $\epsilon - \delta$ definition. That is, any arbitrary small bound on the Hausdorff distance between specifications is satisfied for a sufficiently small parameter difference.

3.1 Assume-Guarantee Contracts

Assume-guarantee reasoning is a common way to abstract a system by encoding what behaviors can be expected under suitable assumptions [3][17]. The assumption is often viewed as an environment experienced by a system.

DEFINITION 1. *(Assume-Guarantee Contract) An assume guarantee contract $\mathcal{C}$ is a pair (ϕ_a, ϕ_g) consisting of an assumption ϕ_a and guarantee ϕ_g that encodes the requirement that the logical implication $\phi_a \Longrightarrow \phi_g$ holds.*

A system $\Sigma \subseteq \mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$ satisfies $\mathcal{C} = (\phi_a, \phi_g)$ if $\Sigma \cap \phi_a \subseteq \phi_g$ (where ϕ_a, ϕ_g are viewed as sets) and satisfaction is depicted in Figure 1. Note that an assume-guarantee specification is automatically satisfied if the assumptions are not true; it

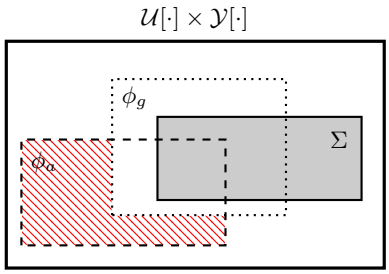

Figure 1: Illustration of system Σ (shaded box) satisfying an unsaturated contract $\mathcal{C} = (\phi_a, \phi_g)$. The space $\mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$ represets the set of all possible behaviors. The system's set of feasible behaviors does not always satisfy the guarantee (dotted box), but this is permitted because the assumption (dashed box) is not true when the violation occurs. A violation occurs only in the patterned region.

can only be violated when the assumption is true and the guarantee is false.

Any contract (ϕ_a, ϕ_g) can be transformed into its *saturated form* (ϕ'_a, ϕ'_g) where $\phi'_a := \phi_a$ and $\phi'_g := (\phi_a \implies \phi_g)$. Logically, a contract and its saturated form are equivalent. However, unlike the original guarantee ϕ_g which is permitted to be false when the assumption is false, the new guarantee ϕ'_g is false only when the system violates the contract. With respect to Figure 1, the new guarantee $\phi'_g := \neg\phi_a \vee \phi_g$ is the complement of the patterned region that signifies a contract violation.

The conjunction of saturated contracts is denoted by $\mathcal{C}^1 \wedge \mathcal{C}^2 = (\phi^1_a \vee \phi^2_a, \phi^1_g \wedge \phi^2_g)$ [3]. Thus a system that satisfies the conjunction of contracts can satisfy tighter guarantees under a wider range of environments/assumptions.

3.2 Parametric Assume Guarantee Contracts

Most assume-guarantee contracts make worst case assumptions about an environment's behavior at design time. A system's guarantee as a result is coarse in order to compensate for the uncertainty about which environment a system will experience once deployed.

In order to make guarantees more precise, we use parametric specifications to divide the assumption ϕ_a into smaller regions $\psi_a(p_a)$ where each p_a can be thought of as an "environmental scenario" parametrized over a set $\mathcal{P}_a$. Systems can provide a finer guarantee $\psi_g(p_g)$ on response to a smaller set of environment behaviors $\psi_a(p_a)$. The relationship between the assumption and guarantee is specified by the parameter map $\lambda : \mathcal{P}_a \mapsto \mathcal{P}_g$.

We define a set of contracts $\mathcal{C}(p_a)$ that correspond to each of these environment scenarios.

$$\mathcal{C}(p_a) := (\psi_a(p_a), \psi_a(p_a) \implies \psi_g(\lambda(p_a))). \tag{2}$$

Definition 2. *(Parametric Assume-Guarantee Contract) Assume-guarantee contract $\mathcal{C} = (\phi_a, \phi_g)$ is in parametric form if there exists a parametric specification $\psi_a : \mathcal{P}_a \mapsto 2^{\mathcal{U}[\cdot]}$, parametric specification $\psi_g : \mathcal{P}_g \mapsto 2^{\mathcal{Y}[\cdot]}$, and parameter map $\lambda : \mathcal{P}_a \mapsto \mathcal{P}_g$ from assumption parameter space $\mathcal{P}_a$ to guarantee parameter space $\mathcal{P}_g$ such that:*

$$\phi_a = \bigvee_{p_a \in \mathcal{P}_a} \psi_a(p_a) \tag{3}$$

$$\phi_g = \bigwedge_{p_a \in \mathcal{P}_a} \big(\psi_a(p_a) \implies \psi_g(\lambda(p_a))\big). \tag{4}$$

The parametric contract can be viewed as a conjunction of smaller contracts $\mathcal{C} = \bigwedge_{p_a \in \mathcal{P}_a} \mathcal{C}(p_a)$.

For all the "assumption scenarios" that are satisfied by the environment, a corresponding guarantee is triggered. Likewise, unsatisfied assumptions do not trigger an obligation to satisfy a guarantee. Parametric assume-guarantee contracts calibrate the guarantees in response to only those environmental scenarios that are satisfied. They are robust in the sense that they are able to provide *some* guarantee despite uncertainties at design time about which environment assumptions will be satisfied after system deployment.

System Σ satisfies the parametric assume-guarantee contract if for all $p_a \in \mathcal{P}_a$:

$$\Sigma \cap \psi_a(p_a) \subseteq \psi_g(\lambda(p_a)).$$

We now show that a finite gain property can be encoded as a parametric assume-guarantee contract. Let $\mathcal{U}[\cdot]$ and $\mathcal{Y}[\cdot]$ be vector spaces equipped with a norm $||\cdot||$.

Example 1 (Bounded Gain). *The bounded gain condition $||y[\cdot]|| \leq \gamma||u[\cdot]|| + \beta$ can be encoded as a parametric assume-guarantee contract (ϕ_a, ϕ_b) where $\mathcal{P}_a = \mathcal{P}_g = \mathbb{R}_{\geq 0} \cup \{\infty\}$, $\psi_a(p_a) := ||u[\cdot]|| \leq p_a$, $\psi_g(p_g) := ||y[\cdot]|| \leq p_g$, and $\lambda(p_a) = \gamma p_a + \beta$.*

$$\phi_a := \bigvee_{p_a \in \mathcal{P}_a} (||u[\cdot]|| \leq p_a)$$

$$\phi_g := \bigwedge_{p_a \in \mathcal{P}_a} (||u[\cdot]|| \leq p_a \implies ||y[\cdot]|| \leq \gamma p_a + \beta)$$

When $||u[\cdot]|| = \infty$ then $||y[\cdot]||$ has a trivial upper bound ∞.

Example 1 highlights how the system's guarantee now adjusts to the environment it is in. If $u[\cdot]$ has a large norm, then $y[\cdot]$ will as well. The parametric assume-guarantee contract is also tight in the sense that a stricter norm bound on $u[\cdot]$ will also incur a stricter norm bound on $y[\cdot]$ automatically.

Note that when parameter spaces $\mathcal{P}_a, \mathcal{P}_g$ are singletons then the parametric contract is a regular assume-guarantee contract as detailed in the previous section. Although parametric assume-guarantee contracts permit us to adjust guarantees in response to assumptions, establishing that a system Σ satisfies such contracts may be difficult. Given a system Σ, calibrating the parameter map λ may require domain specific knowledge from the user. Parametric contracts also have more complex encodings which could incur a computational cost during verification. Some rules of thumb for picking the type of parametric specification and techniques to verify contract satisfaction are provided in latter sections.

4. A SMALL GAIN THEOREM FOR PARAMETRIC CONTRACTS

Consider the interconnection in Figure 2, which contains an exogenous environment and a feedback loop. Suppose for

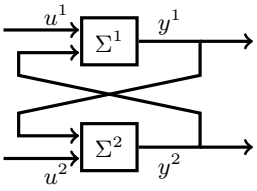

Figure 2: An interconnection with exogenous environment and feedback.

each system Σ^i where $i \in \{1,2\}$ the input space $\mathcal{U}^i = \mathcal{U}^i_e \times \mathcal{U}^i_f$ is partitioned into an exogenous environment $\mathcal{U}^i_e$ component and a feedback component $\mathcal{U}^i_f$. In order for the interconnection to be valid the spaces must match $\mathcal{U}^1_f = \mathcal{Y}^2$ and $\mathcal{U}^2_f = \mathcal{Y}^1$. The assume-guarantee contract framework [3][17] in its full generality ignores the roles of ports as inputs and outputs. With Figure 2 in mind, we place the following restriction on assumptions and guarantees.

ASSUMPTION 1. *Let $\psi_{af} : \mathcal{P}_{af} \mapsto 2^{\mathcal{U}_{af}[\cdot]}$ and $\psi_{ae} : \mathcal{P}_{ae} \mapsto 2^{\mathcal{U}_{ae}[\cdot]}$ be input parametric specifications over appropriate domains, and $\psi_g : \mathcal{P}_g \mapsto 2^{\mathcal{Y}[\cdot]}$ be an output parametric specification.*

The assumption parameter space is partitioned into two components $\mathcal{P}^i_a = \mathcal{P}^i_{ae} \times \mathcal{P}^i_{af}$ and the parametric input assumptions $\psi_a : \mathcal{P}^i_{ae} \times \mathcal{P}^i_{af} \mapsto 2^{\mathcal{U}^i[\cdot]}$ is the conjunction of two components

$$\psi^i_a(p^i_{ae}, p^i_{af}) := \psi^i_{ae}(p^i_{ae}) \wedge \psi^i_{af}(p^i_{af}). \tag{5}$$

The parameter map $\lambda^i : \mathcal{P}^i_{ae} \times \mathcal{P}^i_{af} \mapsto \mathcal{P}^i_g$ is adjusted to account for this decomposition of the input space.

4.1 Main Results

Due to the feedback loop in Figure 2, a few additional assumptions are required to derive a new assume-guarantee contract for the interconnected system. First, the exogenous environment and internal feedback assumptions for at least one system need to be satisfied. Second, the guarantees from one system need to imply that the assumptions for the other system hold.

THEOREM 1. *Consider the interconnection of two systems Σ^1, Σ^2 depicted in Figure 2. We assume the following.*

1. *Both systems satisfy their parametric assume guarantee contracts. That is for $i \in \{1,2\}$, $\Sigma_i \cap \phi^i_a \subseteq \phi^i_g$ where ϕ^i_a and ϕ^i_g are defined as*

$$\phi^i_a = \bigvee_{(p^i_{ae}, p^i_{af}) \in \mathcal{P}^i_a} \left(\psi^i_{ae}(p^i_{ae}) \wedge \psi^i_{af}(p^i_{af}) \right) \tag{6}$$

$$\phi^i_g = \bigwedge_{(p^i_{ae}, p^i_{af}) \in \mathcal{P}^i_a} \big(\psi^i_{ae}(p^i_{ae}) \wedge \psi^i_{af}(p^i_{af}) \implies \psi^i_g(\lambda^i(p^i_{ae}, p^i_{af})) \big). \tag{7}$$

2. *The guarantee parameter spaces are subsets of the feedback components of the assumption parameter spaces, i.e., $\mathcal{P}^2_g \subseteq \mathcal{P}^1_{af}$ and $\mathcal{P}^1_g \subseteq \mathcal{P}^2_{af}$. Moreover, the guarantee $\psi^2_g(\cdot)$ from one system implies the feedback assumption $\psi^1_{af}(\cdot)$ and vice versa:*

$$\forall p \in \mathcal{P}^1_g \quad \psi^2_g(p) \implies \psi^1_{af}(p) \tag{8}$$

$$\forall p \in \mathcal{P}^2_g \quad \psi^1_g(p) \implies \psi^2_{af}(p). \tag{9}$$

This condition is trivially satisfied if $\psi^2_g(\cdot) = \psi^1_{af}(\cdot)$ and vice versa.

3. *There exist environment parameters $p^1_{ae} \in \mathcal{P}^1_{ae}$ and $p^2_{ae} \in \mathcal{P}^2_{ae}$ such that $\psi^1_{ae}(p^1_{ae})$ and $\psi^2_{ae}(p^2_{ae})$ are satisfied.*

4. *For either $i = 1, 2$ there exists a feedback parameter $p^i_{af}[0] \in \mathcal{P}^i_{af}$ such that $\psi^i_{af}(p^i_{af}[0])$ is true.*

Without loss of generality let $i = 1$, define new feedback parameter maps $\hat{\lambda}^1(\cdot) = \lambda^1(p^1_{ae}, \cdot)$ and $\hat{\lambda}^2(\cdot) = \lambda^1(p^2_{ae}, \cdot)$ associated with exogenous environment assumptions p^1_{ae}, p^2_{ae}, and define guarantee parameter iterations

$$p^1_g[k+1] = (\hat{\lambda}^1 \circ \hat{\lambda}^2)(p^1_g[k]) \tag{10}$$

$$p^2_g[k+1] = (\hat{\lambda}^2 \circ \hat{\lambda}^1)(p^2_g[k]) \tag{11}$$

with initializations $p^1_g[0] = \hat{\lambda}^1(p^1_{af}[0])$, $p^2_g[0] = \hat{\lambda}^2(p^1_g[0])$. Then the guarantee simplifies to

$$\bigwedge_{k=0}^{\infty} \psi^1_g\left(p^1_g[k]\right)) \wedge \bigwedge_{k=0}^{\infty} \psi^2_g\left(p^2_g[k]\right). \tag{12}$$

For the case when $i = 2$ then a similar guarantee can be obtained by switching the indexes in (10), (11), and (12).

PROOF. Without loss of generality let $i = 1$. The existence of satisfying p^1_{ae}, p^2_{ae} and $p^1_{af}[0]$ ensure that we can bootstrap an infinite sequence of implications from (7), (8) and (9). The parameters in this implication are generated from the sequences (10) and (11).

$$\begin{aligned}
&\psi^1_{ae}(p^1_{ae}) \wedge \psi^2_{ae}(p^2_{ae}) \wedge \psi^1_{af}(p^1_{af}[0]) \\
&\wedge \big(\psi^1_{ae}(p^1_{ae}) \wedge \psi^1_{af}(p^1_{af}[0]) \implies \psi^1_g(p^1_g[0])\big) \\
&\wedge \big(\psi^1_g(p^1_g[0]) \implies \psi^2_a(p^1_g[0])\big) \\
&\wedge \big(\psi^2_{ae}(p^2_{ae}) \wedge \psi^2_{af}(p^1_g[0]) \implies \psi^2_g(p^2_g[0])\big) \\
&\wedge \big(\psi^2_g(p^2_g[0]) \implies \psi^1_a(p^2_g[0])\big) \\
&\wedge \big(\psi^1_{ae}(p^1_{ae}) \wedge \psi^1_{af}(p^2_g[0]) \implies \psi^1_g(p^1_g[1])\big) \\
&\wedge \dots
\end{aligned}$$

This infinite conjunction sequence contains within it (12) as a subsequence. □

The guarantee (12) can be simplified dramatically by investigating the contraction properties of the parameter iterations (10) and (11). To achieve this simplification we assume that the parameter sets $\mathcal{P}^i_g$ for $i = 1, 2$ are equipped with distance metrics $d^i_{\mathcal{P}} : \mathcal{P}^i_g \times \mathcal{P}^i_g \mapsto \mathbb{R}_{\geq 0}$ and the input and output spaces $\mathcal{U}[\cdot] \times \mathcal{Y}[\cdot]$ are equipped with a distance metric.

THEOREM 2. *(Small Gain Theorem for Parametric Assume-Guarantee Contracts) Let $\mathcal{P}^1_g$ and $\mathcal{P}^2_g$ be metric spaces and $\psi^1_g(\cdot), \psi^2_g(\cdot)$ be specifications on a metric space. If in addition to the assumptions of Theorem 1 the following are also true:*

1. *Sequences generated by the iterations (10), (11) have nonempty $\omega-$limit sets W^1, W^2 respectively.*

2. *The specifications ψ_g^1, ψ_g^2 vary continuously with parameters everywhere in $\mathcal{P}_g^1, \mathcal{P}_g^2$, where the Hausdorff distance d_H is used as a metric between specifications. In other words for both $i = 1, 2$ for all $\epsilon^i > 0$ and $p \in \mathcal{P}_g^i$ there exists a $\delta^i > 0$ such that*

$$d^i_{\mathcal{P}}(p, \hat{p}^i) < \delta^i \Longrightarrow d_H(\psi_g^i(p), \psi_g^i(\hat{p}^i)) < \epsilon^i$$

then the guarantee (12) is over-approximated by:

$$\bigwedge_{p^1 \in W^1} \mathbf{cl}(\psi_g^1(p^1)) \wedge \bigwedge_{p^2 \in W^2} \mathbf{cl}(\psi_g^2(p^2)) \tag{13}$$

where $\mathbf{cl}(\psi)$ is the closure of the specification set ψ.

PROOF. Without loss of generality, we seek to prove that the ψ_g^1 component of formula (12) implies $\mathbf{cl}(\psi_g^1(p^1))$ for a p^1 in the ω-limit set W^1. Suppose $\epsilon > 0$. By Hausdorff continuity of ψ_g^1, there exists a δ such that $|p - p^1| < \delta$ implies $d_H(\psi_g^1(p), \psi_g^1(p^1)) < \epsilon$. Because $p^1 \in W^1$, there is a subsequence of (10) that converges to p^1 and for arbitrary δ. It follows that (12) consists of an infinite sequence of intersections that converge to $\psi_g^1(p^1)$ with arbitrary precision. Because of Hausdorff distance of zero implies that the closure of the two sets are equivalent, this infinite intersection then implies that $\mathbf{cl}(\psi_g^1(p^1))$ holds. Similar arguments can be made for any $p^1 \in W^1$ and for ψ_g^2. □

If the iterations (10) and (11) are contractions to a single point, then we can declare a new parametric assume-guarantee contract for the interconnected system in Figure 2 that is expressed between the exogenous inputs and the outputs.

COROLLARY 1. *If in addition to the assumptions of Theorem 2, the iterations (10) and (11) globally converge to fixed points for any initial parameters $p_{af}^1[0], p_{af}^2[0]$ then the interconnected system of Figure 2 satisfies the parametric assume-guarantee contract associated with*

$$\begin{aligned}
\mathcal{U} &:= \mathcal{U}_e^1 \times \mathcal{U}_e^2 \\
\mathcal{Y} &:= \mathcal{Y}^1 \times \mathcal{Y}^2 \\
\mathcal{P}_a &:= \mathcal{P}_{ae}^1 \times \mathcal{P}_{ae}^2 \\
\mathcal{P}_g &:= \mathcal{P}_g^1 \times \mathcal{P}_g^2 \\
\psi_a(p_{ae}^1, p_{ae}^2) &:= \psi_{ae}^1(p_{ae}^1) \wedge \psi_{ae}^2(p_{ae}^2) \\
\psi_g(p_g^1, p_g^2) &:= \mathbf{cl}(\psi_g^1(\lambda^1(p_{ae}^1, p_{ae}^2))) \wedge \mathbf{cl}(\psi_g^2(\lambda^2(p_{ae}^1, p_{ae}^2))))
\end{aligned}$$

and $\lambda^1 : \mathcal{P}_{ae}^1 \times \mathcal{P}_{ae}^2 \mapsto \mathcal{P}_g^1$ and $\lambda^2 : \mathcal{P}_{ae}^1 \times \mathcal{P}_{ae}^2 \mapsto \mathcal{P}_g^1$ are the respective limit points of the iterations (10) and (11) as a function of exogenous environment assumptions p_{ae}^1 and p_{ae}^2.

4.2 Ensuring that Guarantees are Satisfiable

One technical issue with applying Theorem 2 to richer sets of behaviors is determining whether the guarantees (12) or (13) are nonempty sets and satisfiable. It is advantageous to design parametric specifications to ensure that satisfiability is maintained.

We link parameters to set containment through the notion of monotone specifications, which were previously advocated in the context of requirement mining [12]. Given a partially ordered parameter space $\mathcal{P}_g^1$ equipped with an ordering relation $\leq_{\mathcal{P}_g^1}$, the parametric output specification $\psi_g^1 : \mathcal{P}_g^1 \mapsto 2^{\mathcal{Y}[\cdot]}$ is monotone if $a \leq_{\mathcal{P}_g^1} b$ implies $\psi_g^1(a) \subseteq \psi_g^1(b)$.

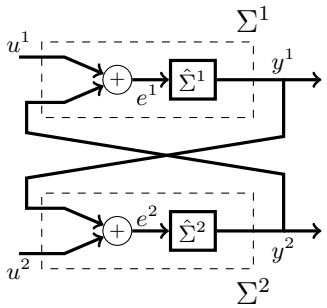

Figure 3: System interconnection for the classical small gain theorem. This interconnection is related to Figure 2 via composite systems Σ^1, Σ^2 which incorporate the addition blocks.

Proposition 1 uses these notions of monotonicity and set containment to give a sufficient condition for the guarantees to be nonempty.

PROPOSITION 1. *Suppose that*

1. *For both $i = 1, 2$ and all nonempty subsets $\mathcal{L}$ of parameter space $\mathcal{P}_g^i$, there exists a lower bound $p \in \mathcal{P}_g^i$ such that $p \leq_{\mathcal{P}_g^i} q$ for all $q \in \mathcal{L}$.*

2. *Parametric output guarantees $\psi_g^1(\cdot)$ and $\psi_g^2(\cdot)$ are monotone specifications and for all parameters $p^1 \in \mathcal{P}_g^1, p^2 \in \mathcal{P}_g^2$ the guarantees $\psi_g^1(p^1), \psi_g^2(p^2)$ are nonempty sets.*

Then the guarantees (12) and (13) are satisfiable/nonempty.

PROOF. Parameters that appear within the sequences (10) and (11) have a lower bound from the first condition and second conditions. Let these lower bounds be denoted as l^1 and l^2 respectively. The sets $\psi_g^1(l^1)$ and $\psi_g^2(l^2)$ are nonempty and $\psi_g^1(l^1) \subseteq \psi_g^1(p_g^1[k])$ and $\psi_g^2(l^2) \subseteq \psi_g^2(p_g^2[k])$ for all $k \in \mathbb{Z}_{\geq 0}$. Therefore, because $\psi_g^1(l^1) \subseteq \cap_{k \in \mathbb{Z}_{\geq 0}} \psi_g^1(p_g^1[k])$ and $\psi_g^2(l^2) \subseteq \cap_{k \in \mathbb{Z}_{\geq 0}} \psi_g^2(p_g^2[k])$, the guarantees (12) and (13) correspond to nonempty sets and are hence satisfiable. □

4.3 Classical Small Gain Theorem as a Special Case

Theorem 2 and Proposition 1 recover the small-gain theorem as stated in [7]. Given some norm, let $\mathcal{L}$ be the space of norm bounded signals. A signal $x[\cdot]$ has an associated T-truncated norm $|x[\cdot]|_T = |\mathbb{I}_T x[\cdot]|$. The signal $x[\cdot]$ is pointwise multiplied with $\mathbb{I}_T$ the indicator function on the time interval $[0, T]$ before the signal norm is taken. The $\mathcal{L}$-extended space $\mathcal{L}_e$ is defined as $\{x[\cdot] : \forall T > 0,\ |x[\cdot]|_T < \infty\}$ and it is clear that $\mathcal{L}$ is a strict subset of $\mathcal{L}_e$.

COROLLARY 2. *(Classical small gain theorem [7]) Let systems $\hat{\Sigma}^1, \hat{\Sigma}^2$ be input-output maps $\hat{\Sigma}^i : \mathcal{L}_e \mapsto \mathcal{L}_e$ and interconnected as in Figure 3. Let $e^1[\cdot], e^2[\cdot] \in \mathcal{L}_e$ and $u^1[\cdot], u^2[\cdot]$ be defined such that*

$$u^1[\cdot] = e^1[\cdot] - y^2[\cdot] \tag{14}$$

$$u^2[\cdot] = e^2[\cdot] - y^1[\cdot]. \tag{15}$$

Suppose there are four constants $\gamma^1, \gamma^2, \beta^1, \beta^2 \geq 0$ such that

$$|y^1[\cdot]|_T \leq \gamma^1 |e^1[\cdot]|_T + \beta^1 \tag{16}$$

$$|y^2[\cdot]|_T \leq \gamma^2 |e^2[\cdot]|_T + \beta^2 \tag{17}$$

for all T. If $\gamma^1\gamma^2 < 1$, then for all T:

$$|y^1[\cdot]|_T \leq \frac{1}{1-\gamma^1\gamma^2}\left(\gamma^1|u^1[\cdot]|_T + \gamma^1\gamma^2|u^2[\cdot]|_T + \gamma^1\beta^2 + \beta^1\right) \tag{18}$$

$$|y^2[\cdot]|_T \leq \frac{1}{1-\gamma^1\gamma^2}\left(\gamma^2|u^2[\cdot]|_T + \gamma^1\gamma^2|u^1[\cdot]|_T + \gamma^2\beta^1 + \beta^2\right). \tag{19}$$

PROOF. The interconnection defined by (14) and (15) is depicted in Figure 3 where the dashed boxes correspond to Σ^1, Σ^2 in Figure 2 used in Theorem 2. Via the triangle inequality, the bounds (16) and (17) are replaced with

$$|y^1[\cdot]|_T \leq \gamma^1|u^1[\cdot]|_T + \gamma^1|y^2[\cdot]|_T + \beta^1 \tag{20}$$

$$|y^2[\cdot]|_T \leq \gamma^2|u^2[\cdot]|_T + \gamma^2|y^1[\cdot]|_T + \beta^2. \tag{21}$$

The assumption parameters associated with the exogenous inputs u^1, u^2 are bounds on their truncated norms. The feedback assumptions pertain to norm bounds on y^1, y^2. The parameter spaces are $\mathcal{P}^1_{ae}, \mathcal{P}^1_{af}, \mathcal{P}^1_g = \mathbb{R}_{\geq 0} \cup \{\infty\}$. For system Σ^1, define the exogenous assumption, feedback assumption, and guarantee as

$$\psi^1_{ae}(p) = (|u^1[\cdot]|_T \leq p) \tag{22}$$

$$\psi^1_{af}(r) = (|y^2[\cdot]|_T \leq r) \tag{23}$$

$$\psi^1_g(r) = (|y^1[\cdot]|_T \leq r) \tag{24}$$

with the parameter iteration map $\lambda^1(p, r) = \gamma^1 p + \gamma^1 r + \beta^1$. With the above definitions, the bounds (20) can be replaced with a parametric assume-guarantee contract. Analogous definitions for Σ^2 lead to a similar reformulation of (21). The first condition of Theorem 1 is therefore satisfied. The second condition is satisfied because both the guarantees and feedback assumptions are of the same form. The third and fourth conditions of Theorem 1 are satisfied because the existence of $e^1[\cdot], e^2[\cdot] \in \mathcal{L}_e$ implies that their T-truncated norm is finite for some T. Via (16) and (17), $y^1[\cdot], y^2[\cdot]$ also have finite T-truncated norm for an identical T. Via the triangle inequality, $u^1[\cdot], u^2[\cdot]$ must have finite T-truncated norm and satisfy (22).

For fixed norm bounds on $u^1[\cdot], u^2[\cdot]$, the feedback iteration functions become $\hat{\lambda}^1(r) := \gamma^1|u^1[\cdot]|_T + \gamma^1 r + \beta^1$ and $\hat{\lambda}^2(r) := \gamma^2|u^2[\cdot]|_T + \gamma^2 r + \beta^2$. When $\gamma^1\gamma^2 < 1$ the parameter iterations converge to a pair of fixed points, which are given by the right hand sides of (18) and (19). Theorem 2 certifies that these bounds are in fact enforced. We know these guarantees are satisfiable via Proposition 1 because any subset of $\mathcal{P}^1_g, \mathcal{P}^2_g = \mathbb{R}_{\geq 0} \cup \{\infty\}$ has a lower bound within $\mathbb{R}_{\geq 0} \cup \{\infty\}$ and the guarantees ψ^1_g, ψ^2_g are non-empty for all parameters. □

5. HAUSDORFF CONTINUITY OF PARAMETRIC LINEAR TEMPORAL LOGIC

The results from the previous section place relatively mild conditions on guarantee specifications $\psi_g(\cdot)$ to provide a small gain result. These were satisfied when the parameteric specification corresponded to sublevel sets of a norm on signals. In this section, we consider a parametric temporal logic variant that can also be used by Theorem 2.

Temporal logic [19] is a powerful formalism to encode complex timing requirements and has been used as a specification language for controller synthesis and verification of cyber-physical systems. Linear temporal logic (LTL) is a specification language for discrete time systems. Predicates are encoded as statements that are true at a specific instant in time, and a set of temporal operators allow one to make statements that incorporate temporal constraints. In our problem formulation, LTL formulas $\phi \subseteq \mathcal{U}^\omega \times \mathcal{Y}^\omega$ can be thought of as sets of infinite length input-output sequences. We consider an LTL variant where input predicates are subsets of $\mathcal{U}$ of the form $f(x) \sim p$ where $\sim \in \{\leq, \geq\}$ and $f : \mathcal{U} \mapsto \mathbb{R}$ is a real-valued function. Output predicates are defined analogously. Definition 3 provides a syntax for constructing LTL formulas.

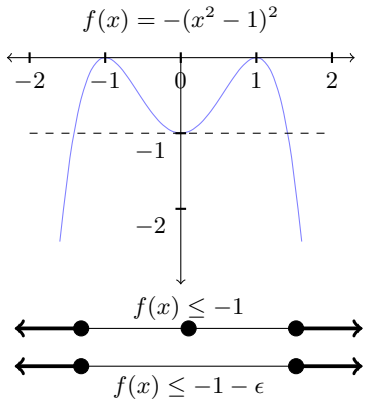

Figure 4: Parametric sublevel sets of non-convex continuous functions are not necessarily Hausdorff continuous.

DEFINITION 3. *Linear temporal logic formulas are constructed with the syntax below*

$$\phi = \top \mid f(\cdot) \sim p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2 \mid \mathbf{X}\phi \mid \mathbf{F}\phi \mid \mathbf{G}\phi \mid \phi_1 \mathbf{U} \phi_2.$$

with parametric predicates $f(\cdot) \sim p$ where $p \in \mathbb{R}$ and $\sim \in \{\leq, \geq\}$. The semantics of the temporal operators $\mathbf{X}, \mathbf{F}, \mathbf{G}$, and $\mathbf{U}$ are summarized below:

- *Specification $\mathbf{X}\phi$ with the "next" operator $\mathbf{X}$ is true if and only if ϕ is true at the next time step.*
- *Specification $\mathbf{F}\phi$ with the "eventually" operator $\mathbf{F}$ is true if and only if ϕ is true at the current time step or there exists a future time step when ϕ is true.*
- *Specification $\mathbf{G}\phi$ with the "always" operator $\mathbf{G}$ is true if and only if ϕ is true at the present and all future time steps.*
- *Specification $\phi_1 \mathbf{U} \phi_2$ with the "until" operator $\mathbf{U}$ is true if and only if ϕ_2 is eventually true and ϕ_1 is true for all future time instances before ϕ_2 becomes true.*

The simplest parametric LTL formula is a parametric predicate, which takes the form of sublevel or superlevel set of a function. The parameter p corresponds to the level. Unfortunately, even for continuous functions $f(\cdot)$, the Hausdorff distance between sublevel sets does not vary continuously for all $p \in \mathbb{R}$. Consider the example given in Figure 4. Due to the presence of a spurious local minimum, perturbing p from $p = -1$ to $p = -1 - \epsilon$ for any $\epsilon > 0$ causes the point at zero to vanish from the sublevel set. The Hausdorff distance

between the two sublevel sets is lower bounded in this example by $\sqrt{2}$ for all sufficiently small neighborhoods around $p = -1$.

To alleviate this issue of disconnected sublevel sets appearing with spurious local minima, we consider a fragment of LTL where the predicates are compact convex sets.

DEFINITION 4. *LTL formulas with convex parametric predicates are constructed with the following syntax.*

$$\phi = \top \mid f(\cdot) \leq p \mid g(\cdot) \geq q \mid \phi_1 \vee \phi_2 \mid \mathbf{X}\phi \mid \mathbf{F}\phi \mid \mathbf{G}\phi \mid \phi_1 \mathbf{U} \phi_2.$$

where for each predicate $f(\cdot) \leq p$ associated with a convex $f(\cdot)$ we restrict p to the domain $[\min_x f(x), \infty]$ and similarly $g(\cdot)$ is concave with $q \in [-\infty, \max_x g(x)]$.

We make the mild technical assumption that $f(\cdot)$ is uniformly continuous; that is, for all $\epsilon > 0$ there exists a δ where $d(x, y) < \delta$ implies $|f(x) - f(y)| < \epsilon$ for all appropriate x, y.

PROPOSITION 2. *The sublevel set of uniformly continuous convex predicates $f(\cdot) \leq p$ varies continuously with parameter $p \in [\min_x f(x), \infty]$ when the distance between predicates is given by the Hausdorff metric.*

PROOF. Consider two predicates $\psi(p) = f(x) \leq p$ and $\psi(p') = f(x) \leq p'$. Without loss of generality, we assume that $p < p'$. Suppose $\epsilon > 0$. Let $m^+ \geq 0$ be defined as

$$m^+ = \left(\sup_{x \in \mathcal{B}_\epsilon(\psi(p))} f(x) \right) - p.$$

m^+ is the least upper bound on how much $f(\cdot)$ may increase by bloating the set $\psi(p)$ by ϵ. Due to uniform continuity of $f(\cdot)$, m^+ is finite. Because $f(\cdot)$ is convex its sublevel sets cannot consist of many disjoint regions and $\{x : f(x) \leq p + am^+\} \subseteq \mathcal{B}_\epsilon(\psi(p))$ for all $0 < a < 1$. If $|p - p'| < am^+$ then $\psi(p') \subseteq \mathcal{B}_\epsilon(\psi(p))$. An identical argument can be made when $p' < p$. Thus, if the parameters $|p - p'| < am^+$ then the Hausdorff distance of the sublevel sets are bounded above by ϵ. □

Note that the syntax in Definition 4 makes the curious choice of permitting disjunctions $\vee$ and not conjunctions $\wedge$. This choice was made due to the following property of the Hausdorff distance

$$d_H(A \cup B, C \cup D) \leq \max(d_H(A, C), d_H(B, D)) \qquad (25)$$

which upper bounds the distance between sets after a union. No analogous property exists for set intersections because they may be empty and the Hausdorff distance is ill defined. The potential loss of convexity under unions and disjunctions is not an issue because convexity of predicates in Definition 4 simply serves as a sufficient condition for predicates to be Hausdorff continuous and is not necessary.

THEOREM 3. *Let $\psi(\cdot)$ be a parametric specification constructed with the convex predicate LTL grammar from Definition 4. Define $N \in \mathbb{Z}_{\geq 0}$ to be the number of times a predicate appears in $\psi(\cdot)$ and parameter space $\mathcal{P} = [\min_x f_1(x), \infty] \times \ldots \times [\min_x f_N(x), \infty]$. Specifications constructed with the grammar $\psi(\cdot)$ are Hausdorff continuous where signals distances are measured with a supremum metric, $d(x[\cdot], y[\cdot]) = \sup_{k \in \mathbb{Z}_{\geq 0}} d(x[k], y[k])$*

PROOF. Suppose $\epsilon > 0$. Let m_i be defined as it appears at the end of the proof of Proposition 2 for predicate $f_i(p_i) \leq p_i$. Let $m = \min_{i \in \{1,\ldots,N\}} m_i$. Suppose that $x[\cdot] \models \psi(p)$ but $x[\cdot] \not\models \psi(p')$ and $\max_i(|p_i - p_i'|) < m$. Each predicate p_i in formula $\psi(p)$ has an associated infinite Boolean sequence where the k-th value is $\top$ if and only if $x[k] \models p_i$. For some time k, there must be at least one predicate that is different; for it to be otherwise would contradict the assertion that $x[\cdot] \not\models \psi(p')$. Given such a time step k, Proposition 2 and (25) guarantee that for any time step k when the difference arises, $x[k]$ must be less than a distance ϵ away from a point $y[k]$ that satisfies the same set of predicates for p'. Thus, $\psi(p) \subseteq \mathcal{B}_\epsilon(\psi(p'))$ where the $\epsilon-$expansion of $\psi(p')$ is with respect to the supremum metric. A similar argument can be made for the case when $x[\cdot] \models \psi(p')$ and $x[\cdot] \not\models \psi(p)$. □

Theorem 3 augments Theorem 2 by providing a concrete instantiation of a class of Hausdorff continuous specifications with temporal logic operators.

6. CERTIFICATION OF PARAMETRIC CONTRACTS

To apply the results from previous sections we need to show that each system satisfies a parametric assume-guarantee contract. We pose a falsification problem that seeks to construct a violation of the contract. Consider a system Σ with a state space $\mathcal{X}$ and initial state set $\mathcal{X}_0$. The notation $y[\cdot] \in \Sigma(x[0], u[\cdot])$ signifies that output $y[\cdot]$ satisfies the dynamics Σ permitted by $x[0]$ and $u[\cdot]$. Let (ϕ_a, ϕ_g) be a parametric contract obtained from parametric specifications ψ_a, ψ_g and parameter map $\lambda : \mathcal{P}_a \mapsto \mathcal{P}_g$.

PROBLEM 1. *If there exist p, $x[0]$, and $u[\cdot]$ that satisfy constraints (27), (28) and (29) then Σ does not satisfy the parametric assume-guarantee contract (ϕ_a, ϕ_g).*

$$\begin{aligned} \textit{find} \quad & p \in \mathcal{P}_a, x[0] \in \mathcal{X}_0, u[\cdot] && (26) \\ \textit{subject to} \quad & u[\cdot] \models \psi_a(p) && (27) \\ & y[\cdot] \not\models \psi_g(\lambda(p)) && (28) \\ & y[\cdot] \in \Sigma(x[0], u[\cdot]). && (29) \end{aligned}$$

The proper falsification engine to solve Problem 1 is implementation specific and depends on both the system dynamics and specification representation.

For black-box systems and systems exhibiting complex, hybrid, and non-linear dynamics, simulation-based falsification is the most practical method to prove that an assume-guarantee contract is satisfied. Most existing simulation-based falsification algorithms are sound but typically not complete. However, the failure to falsify a contract is evidence suggesting that the contract in fact holds. Simulation-based falsification tools are built into toolboxes `S-TaLiRo`[2] and `Breach`[8] for metric and signal temporal logic.

If the falsification algorithm is complete and no violating p, $x[0]$, and $u[\cdot]$ exist, then Σ satisfies (ϕ_a, ϕ_g). The examples in the next section use a component of the `BluSTL` toolbox [20] to translate bounded time temporal logic specifications (27) and (28) into mixed integer constraints for the optimization toolbox `YALMIP` [15].

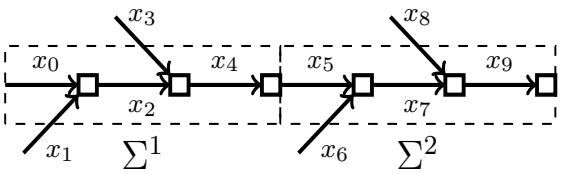

Figure 5: An example network with two on-ramps x_1, x_3. Dashed arrows are exogenous network links.

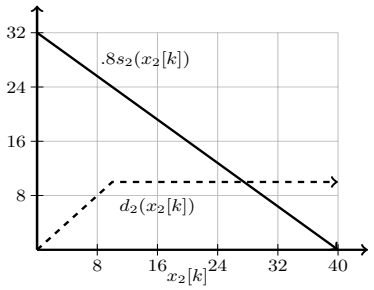

Figure 6: Supply (solid) that link x_2 provides to link x_0 and Demand (dashed) that link x_2 creates for link x_4

7. FREEWAY EXAMPLE

This section applies Theorems 1 and 2 to a freeway traffic example. Consider the two freeway segments depicted in Fig. 5, where the left segment has a main stretch of three links x_0, x_2, x_4 and two on-ramps x_1, x_3. The right segment has identical dynamics. We use the cell transmission model (CTM) [5][11], a macroscopic fluid-like model of freeway dynamics. Individual vehicles are not a component of this model. Each discrete time instant represents a five minute interval.

7.1 Freeway Dynamics

We first describe the dynamics for Σ^1, which are identical to the dynamics of Σ^2 besides a variable renaming, and subsequently describe how the interconnected networks resemble the small-gain interconnections (e.g. Figure 2) from previous sections.

Freeway segment Σ^1's state space $\mathcal{X}^1 \subset \mathbb{R}^5_{\geq 0}$ represents the average occupancy over the five minute period in each of the five links. We overload notation and refer to links and their occupancy values using the same variable. The upper bound on occupancy is encoded with a vector $x^{\max} = [40, 20, 40, 20, 40]$. The state update equations arise from conservation of mass:

$$\begin{aligned}
x_0[k+1] &= x_0[k] - f_0^{\text{out}}[k] + f_0^{\text{in}}[k] \\
x_1[k+1] &= x_1[k] - f_1^{\text{out}}[k] + f_1^{\text{in}}[k] \\
x_2[k+1] &= x_2[k] - f_2^{\text{out}}[k] + f_0^{\text{out}}[k] + f_1^{\text{out}}[k] \\
x_3[k+1] &= x_3[k] - f_3^{\text{out}}[k] + f_3^{\text{in}}[k] \\
x_4[k+1] &= x_4[k] - f_4^{\text{out}}[k] + f_2^{\text{out}}[k] + f_3^{\text{out}}[k]
\end{aligned}$$

where $f_i^{\text{out}}[k]$ and $f_i^{\text{out}}[k]$ respectively represent the flows exiting and entering link x_i at time k.

The flows into and out of a link are determined by *demand* and *supply*. A link's demand is the rate at which it would like to send vehicles to downstream links. The demand $d_i(x_i[k])$ that link x_i exhibits is a non-decreasing function

$$d_i(x_i[k]) = \min(c_i, x_i[k]) \tag{30}$$

where c_i is a saturation rate.The primary links have saturation rates $c_0 = c_2 = c_4 = 10$ and on-ramps have saturation rates $c_1 = c_3 = 5$. All links also exhibit a supply function

$$s(x_i[k]) = x_i^{\max} - x_i[k], \tag{31}$$

which is the rate of incoming vehicles that it can accept from upstream. A link's supply is partitioned among upstream links, with links x_2, x_4 allocating 80% of their supply to an upstream highway link and 20% to on-ramps. Link x_2's supply and demand functions are depicted in Figure 6. Congestion occurs when demand exceeds supply and the left term in the minimization is active. The flows out of links $0, 1, 2, 3$ are the minimum between the supply available to them and their demand:

$$\begin{aligned}
f_0^{\text{out}}[k] &= \min\left(.8(40 - x_2[k]), 10, x_0[k]\right) \\
f_1^{\text{out}}[k] &= \min\left(.2(40 - x_2[k]), 5, x_1[k]\right) \\
f_2^{\text{out}}[k] &= \min\left(.8(40 - x_4[k]), 10, x_2[k]\right) \\
f_3^{\text{out}}[k] &= \min\left(.2(40 - x_4[k]), 5, x_3[k]\right)
\end{aligned}$$

7.2 Interconnection between Networks

Figure 7 summarizes the input and output variables for each network. Network Σ^1 has a vector of demands as its exogenous input $u_{ae}^1 = (d^{\text{exog}}, d_1^{\text{on}}, d_3^{\text{on}})$ and the feedback input $u_{af}^1 = (s_5)$ is the supply from downstream. Similarly, Σ^2 has exogenous input $u_{ae}^2 = (s^{\text{exog}}, d_6^{\text{on}}, d_8^{\text{on}})$ and feedback input $u_{ae}^2 = (d_4)$. The outputs can be identified in a similar manner.

With the notions of demand and supply in mind, we can now consider how both networks are affected by their interconnection and by exogenous environments. The flow f_4^{out} between Σ^1 and Σ^2 is determined by d_4 and s_5

$$f_4^{\text{out}}[k] = \min\left(.8(40 - x_5[k]), 10, x_4[k]\right).$$

Both systems experience an exogenous environment via the on ramp demands. The upstream system Σ^1 also experiences a demand d^{exog} for link x_0 and the downstream network Σ^2 experiences an exogenous supply for link x_9.

Link x_0 allocates 80% of its supply to the exogenous environment. The flow into x_0 is therefore

$$f_0^{\text{in}}[k] = \min(.8(40 - x_0[k]), d^{\text{exog}}).$$

The onramps x_i with $i \in \{1, 3, 6, 8\}$ allocate all supply to the environment so

$$f_i^{\text{in}}[k] = \min((20 - x_i[k]), d_i^{\text{exog}}).$$

Similarly, link x_9's outflow is governed by an exogenous environment so

$$f_9^{\text{out}}[k] = \min(.8s^{\text{exog}}, 10, x_9[k]).$$

7.3 Certifying Intermittent Congestion

Congestion is shown to be intermittent after the two segments are interconnected. Intermittency is encoded via "always" and "eventually" temporal operators augmented with

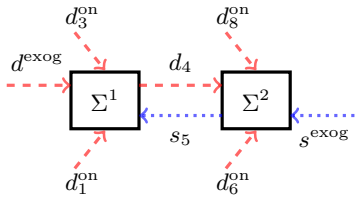

Figure 7: Although vehicular flow in Figure 5 is from left to right, the right network also affects the left network. Demand's influence (dashed lines) is directed forward while supply's influence (dotted lines) is directed backward.

intervals

$$\mathbf{G}_{[0,3]}\phi := \phi \wedge \mathbf{X}\phi \wedge \mathbf{XX}\phi \wedge \mathbf{XXX}\phi \tag{32}$$
$$\mathbf{F}_{[0,2]}\phi := \phi \vee \mathbf{X}\phi \vee \mathbf{XX}\phi. \tag{33}$$

For both systems, all onramp demands are limited to always be less than 3. That is,

$$\mathbf{G}(d_i^{\text{on}} \leq 3) \tag{34}$$

for all $i \in \{1, 3, 6, 8\}$. All links are assumed to have an initial occupancy less than 5, i.e., $\mathcal{X}_0 = \prod_{i=0,\ldots,4}[0, 5] \subset \mathcal{X}$.

From Figure 7, it's clear that the upstream network Σ^1 is subjected to an exogenous mainline demand d^{exog} and the supply availability from downstream network Σ^2. A static exogenous environment contract is imposed by assuming that the main line demand satisfies the assumption with no free parameters

$$\mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d^{\text{exog}} \leq 15). \tag{35}$$

How does the supply from Σ^2 affect the demand outputted by Σ^1? Via monotonicity of the network dynamics, a greater supply availability means that Σ^1 can expel vehicles quicker and will be able to lower the demand it outputs. This relationship is encoded in the parametric assume-guarantee contract below:

$$\phi_a^1 := \bigvee_{s\geq 0} \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(s_5(x_5) \geq 10 - s) \tag{36}$$
$$\phi_g^1 := \bigwedge_{s\geq 0} \Big(\mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(s_5(x_5) \geq 10 - s) \tag{37}$$
$$\implies \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d_4(x_4) \leq \lambda_1(s))\Big) \tag{38}$$
$$\lambda_1(s) := .9s + 4. \tag{39}$$

The falsification procedure encoded in Problem 1 failed to violate the assume-guarantee contract for any parameter $s \geq 0$ and hence Σ^1 satisfies the parametric contract (ϕ_a^1, ϕ_g^1).

Similarly the downstream network Σ^2 is subjected to the demand from x_4 and exogenous supply. The exogenous supply has a fixed assumption

$$\mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(s^{\text{exog}} \geq 5). \tag{40}$$

It influences Σ_1 by outputting supply from x_5 and the contract is

$$\phi_a^2 := \bigvee_{d\geq 0} \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d_4(x_4) \leq d) \tag{41}$$
$$\phi_g^2 := \bigwedge_{d\geq 0} \Big(\mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d_4(x_4) \leq d) \tag{42}$$
$$\implies \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(s_5(x_5) \geq 10 - \lambda_2(d))\Big) \tag{43}$$
$$\lambda_2(d) := .2d. \tag{44}$$

Again, Problem 1 failed to violate the contract and Σ^2 therefore satisfies the contract (ϕ_a^2, ϕ_g^2).

Each of the conditions for Theorem 1 have been proven to hold in this section.

1. The parametric contracts are satisfied for each network.
2. Guarantees from one network imply the feedback assumptions of the other network because they are of the same form. In other words, pairs (41), (38) and (36), (43) are identical parametric specifications.
3. The exogenous assumptions are satisfied via (34), (35), and (40).
4. Let $i = 2$. For a large enough $d \geq 0$, the feedback assumption $\psi_{af}^2(d) = \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d_4(x_4) \leq d)$ is satisfied because d_4 has a maximum value of 10.

The composition of the parameter mapping functions λ_1, λ_2 is a contraction and hence converges in the limit to a fixed point $(d, s) = (4.878, .975)$. Thus, via Theorem 2 the following statement must also hold:

$$\mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(d_4(x_4) \leq 4.878) \wedge \mathbf{G}_{[0,3]}\mathbf{F}_{[0,2]}(s_5(x_5) \geq 9.025).$$

8. CONCLUSION

This paper connects two formalisms for compositional reasoning in the CPS literature: small gain theorems and assume-guarantee contracts. We have incorporated continuous parameters into contracts and showed how to derive an analog of the small gain theorem with broader applicability. We showed that a fragment of parametrized linear temporal logic is Hausdorff continuous and can hence be levered by this small gain theorem, but richer fragments and analogous results for continuous time specifications may exist. Future work will also investigate applying the parametric assume-guarantee framework for additional variants of the small gain theorem (such as those involving input-to-state stability) and broader classes of input-output properties.

9. ACKNOWLEDGEMENTS

The authors would like to thank Shromona Ghosh for helpful discussions about assume-guarantee contracts.

10. REFERENCES

[1] M. Al Khatib, A. Girard, and T. Dang. Verification and synthesis of timing contracts for embedded controllers. In *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*, pages 115–124. ACM, 2016.

[2] Y. Annpureddy, C. Liu, G. Fainekos, and S. Sankaranarayanan. S-taliro: A tool for temporal logic falsification for hybrid systems. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 254–257. Springer, 2011.

[3] A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J.-B. Raclet, P. Reinkemeier, A. Sangiovanni-Vincentelli, W. Damm, T. Henzinger, and K. Larsen. Contracts for systems design. 2012.

[4] R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In *International Conference on Computer Aided Verification*, pages 140–156. Springer, 2009.

[5] C. F. Daganzo. The cell transmission model: A Dynamic Representation of Highway Traffic Consistent with the Hydrodynamic Theory. *Transportation Research*, 28:269–287, 1994.

[6] E. Dallal and P. Tabuada. On compositional symbolic controller synthesis inspired by small-gain theorems. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 6133–6138. IEEE, 2015.

[7] C. A. Desoer and M. Vidyasagar. *Feedback systems: input-output properties*, volume 55. Siam, 2009.

[8] A. Donzé. Breach, a toolbox for verification and parameter synthesis of hybrid systems. In *International Conference on Computer Aided Verification*, pages 167–170. Springer, 2010.

[9] A. Donzé and O. Maler. Robust satisfaction of temporal logic over real-valued signals. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 92–106. Springer, 2010.

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

[11] G. Gomes and R. Horowitz. Optimal freeway ramp metering using the asymmetric cell transmission model. *Transportation Research Part C: Emerging Technologies*, 14(4):244 – 262, 2006.

[12] X. Jin, A. Donzé, J. V. Deshmukh, and S. A. Seshia. Mining requirements from closed-loop control models. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(11):1704–1717, 2015.

[13] E. S. Kim, M. Arcak, and S. A. Seshia. Compositional controller synthesis for vehicular traffic networks. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 6165–6171. IEEE, 2015.

[14] M. Kwiatkowska, G. Norman, D. Parker, and H. Qu. Compositional probabilistic verification through multi-objective model checking. *Information and Computation*, 232:38–65, 2013.

[15] J. Lofberg. Yalmip: A toolbox for modeling and optimization in matlab. In *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*, pages 284–289. IEEE, 2004.

[16] R. Majumdar, E. Render, and P. Tabuada. A theory of robust omega-regular software synthesis. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(3):48, 2013.

[17] P. Nuzzo, A. L. Sangiovanni-Vincentelli, D. Bresolin, L. Geretti, and T. Villa. A platform-based design methodology with contracts and related tools for the design of cyber-physical systems. *Proceedings of the IEEE*, 103(11):2104–2132, 2015.

[18] P. Nuzzo, H. Xu, N. Ozay, J. B. Finn, A. L. Sangiovanni-Vincentelli, R. M. Murray, A. Donzé, and S. A. Seshia. A contract-based methodology for aircraft electric power system design. *IEEE Access*, 2:1–25, 2014.

[19] A. Pnueli. The temporal logic of programs. In *Foundations of Computer Science, 1977., 18th Annual Symposium on*, pages 46–57. IEEE, 1977.

[20] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia. Reactive synthesis from signal temporal logic specifications. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, pages 239–248. ACM, 2015.

[21] M. Rungger and M. Zamani. Compositional construction of approximate abstractions. In *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, pages 68–77. ACM, 2015.

[22] P. Tabuada, S. Y. Caliskan, M. Rungger, and R. Majumdar. Towards robustness for cyber-physical systems. *IEEE Transactions on Automatic Control*, 59(12):3151–3163, Dec 2014.

[23] D. C. Tarraf, A. Megretski, and M. A. Dahleh. A framework for robust stability of systems over finite alphabets. *IEEE Transactions on Automatic Control*, 53(5):1133–1146, June 2008.