

Hybrid System Falsification and Reinforcement Learning

Formal Method for Cyber-Physical Systems

Clovis Eberhart David Sprunger

National Institute of Technology, Japan

SOKENDAI lesson, July 1, 8, and 22



Lecture structure

Lectures

- 1st: falsification (problem, framework, logics. . .), by me
- 2nd: deep learning for falsification (learning, reinforcement learning, application to falsification. . .), by David
- 3rd: advanced techniques in falsification and reinforcement learning, by David and me

Evaluation

Easy practical assignment (in Python).

Questions?

- Ask them during the lesson.
- Find me at my desk (Palaceside building).
- `clovis.eberhart@gmail.com`

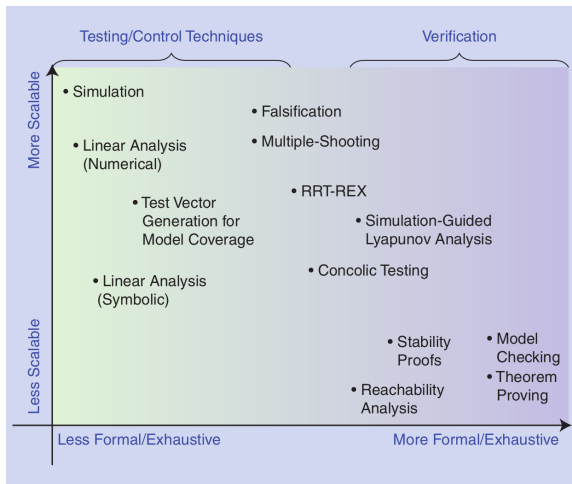
First lecture overview

- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems
- 4 Formulas
- 5 Optimisation

Table of Contents

- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems
- 4 Formulas
- 5 Optimisation

Formal method landscape



Kapinski, Deshmukh, Jin, Ito, Butts, *Simulation-Based Approaches for Verification of Embedded Control Systems*, IEEE Control Magazine, 2010

Testing

Testing

Given: a system \mathcal{S} , a property φ .

Goal: generate a test suite $\{t_i\}_{i \in I}$.

Testing

Testing

Given: a system \mathcal{S} , a property φ .

Goal: generate a test suite $\{t_i\}_{i \in I}$.

Characteristics

- **simple** (run tests)
- **black-box** (unknown systems)
- **versatile** (guarantees, explainable failures...)
- **no formal guarantee**
- **too general**

Verification

Verification

Given: a model \mathcal{M} , a property φ .

Goal: automatically prove that $\mathcal{M} \models \varphi$.

Verification

Verification

Given: a model \mathcal{M} , a property φ .

Goal: automatically prove that $\mathcal{M} \models \varphi$.

Characteristics

- **complex** (design model, use specific techniques, so typically not used by engineers)
- **white-box** (known systems only)
- **formal proof** (strong guarantee)
- **ill-suited to CPS** (continuous systems)

Verification

Verification

Given: a model \mathcal{M} , a property φ .

Goal: automatically prove that $\mathcal{M} \models \varphi$.

Characteristics

- **complex** (design model, use specific techniques, so typically not used by engineers)
- **white-box** (known systems only)
- **formal proof** (strong guarantee)
- **ill-suited to CPS** (continuous systems)

Theorem proving

Given: a model \mathcal{M} , a property φ .

Goal: prove that $\mathcal{M} \models \varphi$.

Falsification

Falsification

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Falsification

Falsification

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Characteristics

- particular case of testing
- **black-box** (unknown systems)
- **relatively simple**
- **no proof** (no formal guarantee)

Verification versus falsification

- Verification:
 - finds a proof: **system verifies property**,
 - finds nothing:
- Falsification:
 - finds a counterexample: **system violates property**,
 - finds nothing:

Verification versus falsification

- Verification:
 - finds a proof: **system verifies property**,
 - finds nothing: **nothing can be said**.
- Falsification:
 - finds a counterexample: **system violates property**,
 - finds nothing: **nothing can be said**.

Verification versus falsification

- Verification:
 - finds a proof: **system verifies property**,
 - finds nothing: **nothing can be said**.
- Falsification:
 - finds a counterexample: **system violates property**,
 - finds nothing: **nothing can be said**.

Interaction:

- verification for falsification: constraining state space by reachability analysis,
- falsification for verification: coverage-based falsification techniques.

Table of Contents

- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems
- 4 Formulas
- 5 Optimisation

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system?

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system? \rightsquigarrow hybrid system

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system? \rightsquigarrow hybrid system
- What is a property?

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system? \rightsquigarrow hybrid system
- What is a property? \rightsquigarrow logical formula

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system? \rightsquigarrow hybrid system
- What is a property? \rightsquigarrow logical formula
- What is a counterexample?

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.

Questions

- What is a system? \rightsquigarrow hybrid system
- What is a property? \rightsquigarrow logical formula
- What is a counterexample? \rightsquigarrow an input (and output) signal to the system that violates the property

General framework

Reminder

Given: a system \mathcal{S} , a property φ .

Goal: generate a counterexample to $\mathcal{S} \models \varphi$.


Questions

- What is a system? \rightsquigarrow hybrid system
- What is a property? \rightsquigarrow logical formula
- What is a counterexample? \rightsquigarrow an input (and output) signal to the system that violates the property

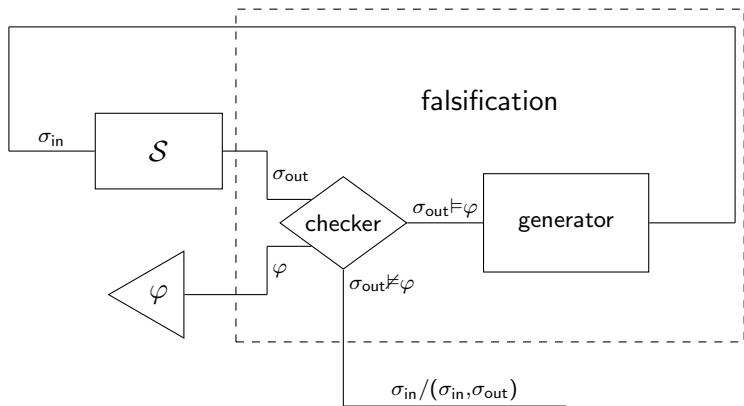
Challenges

- infinite (and high-dimensional) search space
- non-linear dynamics

The falsification loop



The falsification loop



The falsification algorithm

Input: A system \mathcal{S} , a formula φ , a satisfaction predicate \models ,
and a timeout t_{\max}

Output: A signal σ_{in} such that $\mathcal{S}(\sigma_{\text{in}}) \not\models \varphi$
found = false;

while *not*(found) and $t < t_{\max}$ **do**

| $\sigma_{\text{in}} = \text{generate}();$

| $\sigma_{\text{out}} = \mathcal{S}(\sigma_{\text{in}});$

| found = $\sigma_{\text{out}} \not\models \varphi;$

end

if found **then**

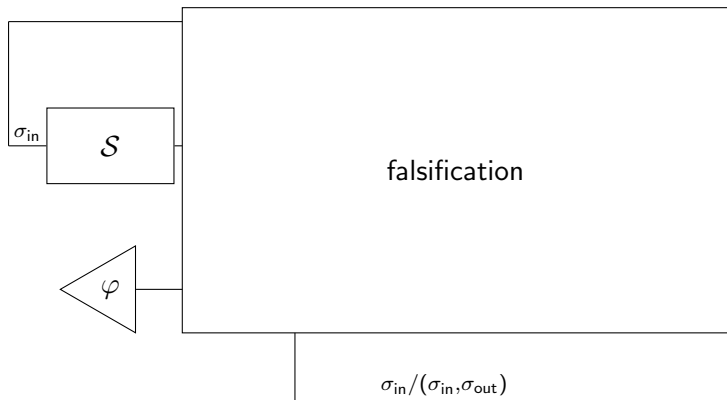
| **return** σ_{in}

else

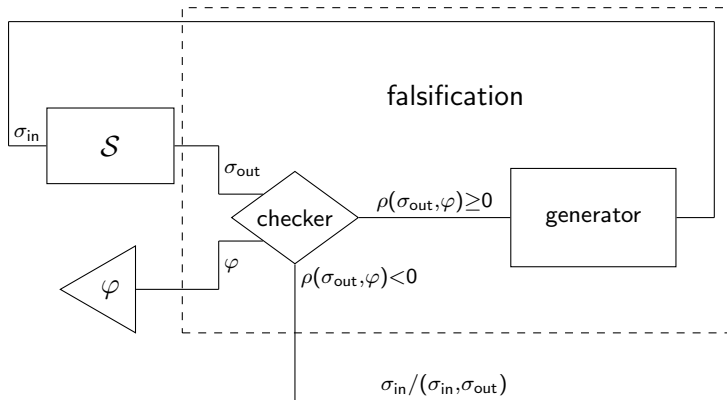
| **return** "timeout"

end

Optimisation-based falsification



Optimisation-based falsification



The optimisation-based falsification algorithm

Input: A system \mathcal{S} , a formula φ , a **robustness function** ρ , and a timeout t_{\max}

Output: A signal σ_{in} such that $\mathcal{S}(\sigma_{\text{in}}) \not\models \varphi$

found = false;

while *not*(found) *and* $t < t_{\max}$ **do**

$\sigma_{\text{in}} = \text{search_minimum}(\rho)$;

$\sigma_{\text{out}} = \mathcal{S}(\sigma_{\text{in}})$;

 found = $\rho(\sigma_{\text{out}}, \varphi) < 0$;

end

if found **then**

return σ_{in}

else

return "timeout"

end

Required: $\rho(\sigma, \varphi) \geq 0 \iff \sigma \models \varphi$

Table of Contents


- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems**
- 4 Formulas
- 5 Optimisation

Hybrid systems

Definition

A **hybrid system** is a dynamical system that exhibits both continuous and discrete dynamic behavior – a system that can both *flow* (described by a differential equation) and *jump* (described by a state machine or automaton).

Wikipedia



Fehnker, Ivančić, *Benchmarks for Hybrid Systems Verification*, Hybrid Systems: Computation and Control, pp 326–341


Hybrid system: definition

Hybrid system


A **Hybrid system** is a tuple $\mathcal{H} = (Q, X, \text{GUARD}, \text{JUMP}, U, \text{FLOW})$ of:

- a finite set of **modes** Q ,
- a family of continuous state spaces $X = \{X_q \subseteq \mathbb{R}^{n_q} \mid q \in Q\}$,
- $\text{GUARD}_{q,q'} \subseteq X_q$ is the set of states in X_q that can transition to mode q' ,
- $\text{JUMP}_{q,q'}: X_q \rightarrow X_{q'}$ describes the transition from q to q' ,
- U is the **input space**,
- FLOW_q is a set of differential equations in X_q and U , seen as a function $X_q \times U \times \mathbb{R}_{\geq 0} \rightarrow X_q$,

Example of hybrid system: thermostat




Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$
- $\text{GUARD}_{q_1, q_0} = \{M\}$

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$
- $\text{GUARD}_{q_1, q_0} = \{M\}$
- $\text{JUMP}_{q, q'}(T) = T$

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$
- $\text{GUARD}_{q_1, q_0} = \{M\}$
- $\text{JUMP}_{q, q'}(T) = T$
- $U = \emptyset$

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$
- $\text{GUARD}_{q_1, q_0} = \{M\}$
- $\text{JUMP}_{q, q'}(T) = T$
- $U = \emptyset$
- $\text{FLOW}_{q_0} = (\frac{dT}{dt} = -KT)$

Example of hybrid system: thermostat




- $Q = \{q_0, q_1\}$,
- $X_{q_0} = \{T \in \mathbb{R} \mid T \geq T_0\}$,
- $X_{q_1} = \{T \in \mathbb{R} \mid T \leq T_1\}$,
- $\text{GUARD}_{q_0, q_1} = \{m\}$
- $\text{GUARD}_{q_1, q_0} = \{M\}$
- $\text{JUMP}_{q, q'}(T) = T$
- $U = \emptyset$
- $\text{FLOW}_{q_0} = (\frac{dT}{dt} = -KT)$
- $\text{FLOW}_{q_1} = (\frac{dT}{dt} = K(h - T))$

Example of hybrid system: timed automata




Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,

Example of hybrid system: timed automata



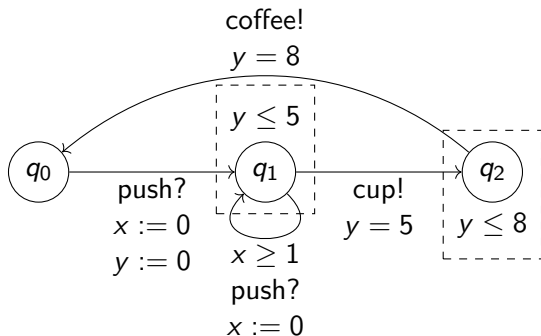
- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$
- $\text{GUARD}_{q_1, q_2} = \mathbb{R} \times \{5\}$

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$
- $\text{GUARD}_{q_1, q_2} = \mathbb{R} \times \{5\}$
- $\text{JUMP}_{q_0, q_1}(x, y) = (0, 0)$

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$
- $\text{GUARD}_{q_1, q_2} = \mathbb{R} \times \{5\}$
- $\text{JUMP}_{q_0, q_1}(x, y) = (0, 0)$
- $\text{JUMP}_{q_1, q_1}(x, y) = (0, y)$

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$
- $\text{GUARD}_{q_1, q_2} = \mathbb{R} \times \{5\}$
- $\text{JUMP}_{q_0, q_1}(x, y) = (0, 0)$
- $\text{JUMP}_{q_1, q_1}(x, y) = (0, y)$
- $U = \emptyset$

Example of hybrid system: timed automata




- $Q = \{q_0, q_1, q_2\}$,
- $X_{q_0} = \mathbb{R}^2$,
- $X_{q_1} = \{(x, y) \in \mathbb{R}^2 \mid y \leq 5\}$,
- $\text{GUARD}_{q_0, q_1} = \mathbb{R}^2$
- $\text{GUARD}_{q_1, q_2} = \mathbb{R} \times \{5\}$
- $\text{JUMP}_{q_0, q_1}(x, y) = (0, 0)$
- $\text{JUMP}_{q_1, q_1}(x, y) = (0, y)$
- $U = \emptyset$
- $\text{FLOW}_q = \left(\frac{dx}{dt} = \frac{dy}{dt} = 1\right)$

Example of hybrid system: navigation



- $Q = \{(i, j) \mid i, j \in [0..4]\}$
- X :
 - $X_{(i,j)} = [0, 1] \times [0, 1] \times \mathbb{R} \times \mathbb{R} \quad (i, j \in [1..3])$
 - $X_{(0,j)} = (-\infty, 1] \times [0, 1] \times \mathbb{R} \times \mathbb{R} \quad (j \in [1..3])$
 - $X_{(4,j)} = [0, \infty) \times [0, 1] \times \mathbb{R} \times \mathbb{R} \quad (j \in [1..3])$
 - $X_{(i,0)} = [0, 1] \times (-\infty, 1] \times \mathbb{R} \times \mathbb{R} \quad (i \in [1..3])$
 - $X_{(i,4)} = [0, 1] \times [0, \infty) \times \mathbb{R} \times \mathbb{R} \quad (i \in [1..3])$
 - $X_{(0,0)} = (-\infty, 1] \times (-\infty, 1] \times \mathbb{R} \times \mathbb{R}$
 - ...

Example of hybrid system: navigation




- GUARD:

- $\text{GUARD}_{(i,j),(i-1,j)} = \{(0, y, v_x, v_y)\}$
- $\text{GUARD}_{(i,j),(i+1,j)} = \{(1, y, v_x, v_y)\}$
- $\text{GUARD}_{(i,j),(i,j-1)} = \{(x, 0, v_x, v_y)\}$
- $\text{GUARD}_{(i,j),(i,j+1)} = \{(x, 1, v_x, v_y)\}$

- JUMP:

- $\text{JUMP}_{(i,j),(i-1,j)}(x, y, v_x, v_y) = (x + 1, y, v_x, v_y)$
- $\text{JUMP}_{(i,j),(i+1,j)}(x, y, v_x, v_y) = (x - 1, y, v_x, v_y)$
- $\text{JUMP}_{(i,j),(i,j-1)}(x, y, v_x, v_y) = (x, y + 1, v_x, v_y)$
- $\text{JUMP}_{(i,j),(i,j+1)}(x, y, v_x, v_y) = (x, y - 1, v_x, v_y)$

Example of hybrid system: navigation



- $U = ([-0.1, 0.1]^{\mathbb{R}_{\geq 0}})^2$
- $\text{FLOW}_{i,j}$:
 - $\frac{dx}{dt} = v_x + u_x(t)$,
 - $\frac{dy}{dt} = v_y + u_y(t)$,
 - $\frac{dv_x}{dt} = 0.1(v_y - v_y^{(i,j)}) - 1.2(v_x - v_x^{(i,j)})$,
 - $\frac{dv_y}{dt} = 0.1(v_x - v_x^{(i,j)}) - 1.2(v_y - v_y^{(i,j)})$,

where

- $(u_x, u_y) \in U$,
- $v_x^{(i,j)}$ and $v_y^{(i,j)}$ are constants.


Run of a hybrid automaton

Run

Finite or infinite sequence $(q_0, x_0) \rightarrow_{t_0} (q_1, x_1) \rightarrow_{t_1} \dots$ such that, for each $i \in [1..n - 1]$:

- $y_i \in \text{GUARD}(q_i, q_{i+1})$,
- $\text{JUMP}_{q_i, q_{i+1}}(y_i) = x_{i+1}$,

where $y_i = \text{FLOW}_{q_i}(x_i, u, t_i - t_{i-1})$.



- $((1, 2), x_0) \rightarrow_{t_0} ((2, 2), x_1) \rightarrow_{t_1} ((2, 1), x_2) \rightarrow_{t_2} ((2, 0), x_3)$
- $((0, 1), x'_0) \rightarrow_{t'_0} ((0, 0), x'_1) \rightarrow_{t'_1} ((1, 0), x'_2) \rightarrow_{t'_2} ((2, 0), x'_3)$

Runs as signals

It is often more practical to consider runs of a hybrid system as timed signals.

Signal

A **signal** on a set of variables $\{x_i \mid i \in [1..n]\}$, where each x_i takes values in X_i , is a function $\mathbb{R}_{\geq 0} \rightarrow \prod_{i=1}^n X_i$.

Translation

In the case where all X_q 's are subspaces of a given X , the following signal corresponds to the run $\rho = (q_0, x_0) \rightarrow_{t_0} (q_1, x_1) \rightarrow_{t_1} \dots$ of a hybrid automaton under input signal u :

$$\sigma_\rho: \begin{array}{l} [0, \sum_{i=1}^n t_i) \rightarrow Q \times X \\ \left(\sum_{i=1}^k t_i \right) + t \mapsto (q_k, \text{FLOW}_{q_k}(x_k, u, t)), \end{array}$$

where $t < t_{k+1}$.

Signals in practice

Time-boundedness

Since simulation cannot be run for an infinite amount of time, so all considered signals are time-bounded: they are not functions $\mathbb{R}_{\geq 0} \rightarrow U$ but $[0, T) \rightarrow U$ for some $T \in \mathbb{R}$.

Finite representation of input signals

We must also stick to classes of signals that can be represented by finite means, e.g., piecewise constant signals, piecewise affine signals, spline (piecewise polynomial) signals. . .

Control points (those points between which the function is interpolated) can be chosen equidistant, or according to other policies.

Discretisation of signals

Discretisation of output signals

Output signals may not be finitely representable, so they are discretised.

Attention

Discretisation of signals can lead to:

- false positives: e.g., for formula $F \varphi$, if $\sigma(t_i) \not\models \varphi$, but $\sigma(t) \models \varphi$ for some $t_i < t < t_{i+1}$, the algorithm will return σ as a (wrong) counterexample,
- false negatives: for dual reasons.

Initial position and signals

Two kinds of falsifications:

- falsification on initial position,
- falsification on signals.

Differences:

- An initial position can be seen as a constant signal, so falsification on initial positions is easier.
- Many falsification methods incrementally modify the shape of the signal (learning approaches, Monte-Carlo Tree Search. . .), so they are ill-suited to falsification on initial positions.

Table of Contents

- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems
- 4 Formulas**
- 5 Optimisation

Signal temporal logic

Syntax

$$\varphi ::= \top \mid f \sim 0 \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi \mathbf{U}_I \psi$$

where $I = [a, \infty)$ or $[a, b]$ for $a < b$, $f: S \rightarrow \mathbb{R}$, and $\sim \in \{>, =\}$ is a comparison operator.

Syntactic sugar

- $\perp = \neg\top$, $\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi)$...
- $f \geq 0 = f > 0 \vee f = 0$, $f \leq 0 = \neg f > 0$...
- $F_I \varphi = \top \mathbf{U}_I \varphi$, $G_I \varphi = \neg F_I \neg\varphi$
- $\mathbf{U} \varphi = \mathbf{U}_{[0, \infty)} \varphi$...

STL examples

To avoid awkward notations, $f \sim 0$ is never written as a function, but as a computation on states. Thus, if

$f: (q, x, y, v_x, v_y) \mapsto x^2 + y^2 - 1$, $f < 0$ would be written $x^2 + y^2 - 1 < 0$ or $x^2 + y^2 < 1$.

Examples

- $q \neq (3, 1) \cup q = (1, 3)$
- $G(\text{gear} = 2 \rightarrow \neg F_{[0, \varepsilon]}(\text{gear} \neq 2 \wedge F_{[0, \tau]} \text{gear} = 2))$
- $(q = q_i \wedge x \in [x_i^-, x_i^+]) \wedge F_{[0, T]}(q = q_f \wedge x \in [x_f^-, x_f^+])$
- $(q = q_i \wedge x \in [x_i^-, x_i^+]) \wedge G_{[0, T]} \neg(q = q_f \wedge x \in [x_f^-, x_f^+])$
- $G(\text{danger} \rightarrow F_{[0, t]} \neg \text{danger})$ with
 $\text{danger} = (q = q_d \wedge x \in [x_d^-, x_d^+])$

Boolean semantics of STL

The semantics of STL formulas is defined over timed signals:

Boolean semantics

$$\sigma, t \models \top \iff \top$$

$$\sigma, t \models f \sim 0 \iff f(\sigma(t)) \sim 0$$

$$\sigma, t \models \neg\varphi \iff \sigma, t \not\models \varphi$$

$$\sigma, t \models \varphi \wedge \psi \iff \sigma, t \models \varphi \text{ and } \sigma, t \models \psi$$

$$\sigma, t \models \varphi \mathbf{U}_{[a,b]} \psi \iff \text{there is } t' \in [a, b] \text{ such that } \sigma, t' \models \psi \\ \text{and for all } t'' \in [a, t'), \sigma, t'' \models \varphi$$

$\sigma \models \varphi$ stands for $\sigma, 0 \models \varphi$.

Robustness semantics of STL

We can define a robustness semantics, whose value is not a boolean, but a real:

Robustness semantics

$$\rho(\sigma, \top, t) = \infty$$

$$\rho(\sigma, f > 0, t) = f(\sigma(t))$$

$$\rho(\sigma, f = 0, t) = -|f(\sigma(t))|$$

$$\rho(\sigma, \neg\varphi, t) = -\rho(\sigma, \varphi, t)$$

$$\rho(\sigma, \varphi \wedge \psi, t) = \min(\rho(\sigma, \varphi, t), \rho(\sigma, \psi, t))$$

$$\rho(\sigma, \varphi \cup_{[a,b]} \psi, t) = \sup_{t' \in [a,b]} \min \left(\rho(\sigma, \psi, t'), \inf_{t'' \in [a,t']} \rho(\sigma, \varphi, t'') \right)$$

$\rho(\sigma, \varphi)$ stands for $\rho(\sigma, \varphi, 0)$.

Soundness of robustness semantics

Theorem

If $\rho(\sigma, \varphi) > 0$, then $\sigma \models \varphi$.

If $\rho(\sigma, \varphi) < 0$, then $\sigma \not\models \varphi$.

Proof.

Structural induction. □

Remark

If $\rho(\sigma, \varphi) = 0$, nothing can be said: if $\sigma(x)(0) = x_0$ for $\varphi = (x = x_0)$:

- $\rho(\sigma, \varphi) = \rho(\sigma, \neg\varphi) = 0$, but
- $\sigma \models \varphi$, $\sigma \not\models \neg\varphi$.

Advantage: if functions f are continuous or smooth, we can use optimisation techniques to find the minimal robustness.

Computing the robustness semantics

Computing ρ

Given: piecewise-constant σ (because discretised).

Goal: compute $\rho(\sigma, \varphi)$.

Naive idea: compute $\rho(\sigma, \varphi, t)$ inductively.

Problem

Complexity to compute $\rho(\sigma, \varphi \cup_{[a,b]} \psi, -)$ is

$\mathcal{O}(\text{number of control points} \times \text{number of control points in } [a, b])$

Efficiently computing $\rho(\sigma, \varphi, -)$

First step: get rid of U_I :

- $\varphi U_{[a,b]} \psi \equiv F_{[a,b]} \varphi \wedge G_{[0,a]} (\varphi U \psi)$
- $\varphi U_{[a,\infty)} \psi \equiv G_{[0,a]} (\varphi U \psi)$

Goal

The function $t \mapsto \rho(\sigma, \varphi, t)$ is a signal $\mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ denoted $\rho(\sigma, \varphi, -)$. Our goal is to compute it from $\rho(\sigma, \psi, -)$ for subformulas ψ of φ .

Need to know how to recursively compute signals for \top , $f \sim 0$, $\neg\varphi$, $\varphi \vee \psi$, $\varphi U \psi$, and $F_{[a,b]} \varphi$ in $\mathcal{O}(\text{number of control points})$.

Computing $\rho(\sigma, \varphi \cup \psi, -)$

Without loss of generality, signals y and y' representing $\rho(\sigma, \varphi, -)$ and $\rho(\sigma, \psi, -)$ have the same control points t_i (otherwise, take the union of control points).

Let z be the signal corresponding to $\rho(\sigma, \varphi \cup \psi, -)$, then:

$$\begin{aligned} z(t_i) &= \sup_{t \in [t_i, \infty)} \min\{y'(t), \inf_{[t_i, t)} y\} \\ &= \max\{\min\{y'(t_i), y(t_i)\}, \sup_{t \in [t_{i+1}, \infty)} \min\{y'(t), \inf_{[t_i, t)} y\}\} \\ &= \max\{\min\{y'(t_i), y(t_i)\}, \min\{y(t_i), \sup_{t \in [t_{i+1}, \infty)} \min\{y'(t), \inf_{[t_{i+1}, t)} y\}\}\} \\ &= \min\{y(t_i), \max\{y'(t_i), z(t_{i+1})\}\} \end{aligned}$$

Thus, there is an algorithm to compute $\rho(\sigma, \varphi \cup \psi, -)$ from $\rho(\sigma, \varphi, -)$ and $\rho(\sigma, \psi, -)$ whose complexity is linear in the number of control points.


Computing $\rho(\sigma, F_{[a,b]} \varphi, -)$

We need to compute $z(t) = \sup_{t+[a,b]} y = \max_{t_i \in t+[a,b]} \{y(t_i)\}$.

Idea

Compute M such that $i \in M$ iff $t_i \in t + [a, b]$ and for all $t_j \in t + [a, b]$, $y(t_j) < y(t_i)$.

Thus: $y(t_{\min M}) = \max_{t_i \in t+[a,b]} \{y(t_i)\}$.



Alexandre Donz , Thomas Ferr re, Oded Maler, *Efficient Robust Monitoring for STL*, International Conference on Computer Aided Verification 2013

Analysis

- M can be implemented to have all operations in $\mathcal{O}(1)$ (doubly-linked list),
- all control points are popped at most once,
- in total, the number of comparisons (done when searching for elements to pop) is at most $2n$,

therefore, computing $\rho(\sigma, F_{[a,b]} \varphi, -)$ can be done in time linear with respect to the number of control points.

Overall complexity

The number of control points is at most $d^{h(\varphi)}|\sigma|$, so the whole complexity is $\mathcal{O}(|\varphi|d^{h(\varphi)}|\sigma|)$.

Generalisation

The same argument (but more complex) applies to more general signals (say, piecewise affine).

Table of Contents

- 1 Formal methods landscape
- 2 Framework
- 3 Hybrid systems
- 4 Formulas
- 5 Optimisation**

Optimisation techniques

Falsification is based on a number of optimisation techniques:

- ant colony,
- CMA-ES,
- cross-entropy,
- gradient descent,
- hill-climbing,
- Nelder-Mead,
- simulated annealing,
- ...

Gradient descent

Goal: find a *local* minimum to a function f .

```
1  $i = 0$ ;  
2 while continue do  
3    $x_{i+1} = x_i - \gamma_i \cdot \nabla(f)(x_i)$ ;  
4    $best = x_{i+1}$ ;  
5   if  $f(x_{i+1}) > f(x_i)$  then  
6      $continue = false$ ;  
7      $best = x_i$ ;  
8   end  
9    $i ++$ ;  
end  
9 return best
```

Finds a local minimum: may be useful in certain cases.

Hill climbing

Goal: find a *local* maximum to a function f .

```
1  $i = 0$ ;  
2 while continue do  
3   for  $x' \in \text{neighbours}(x_i)$  do  
4     if  $f(x') > f(x_i)$  then  
5        $x_{i+1} = x'$ ;  
6       break;  
     end  
   end  
7   continue =  $x_{i+1} \neq x_i$ ;  
8    $i++$ ;  
end  
9 return  $x_{i-1}$ 
```

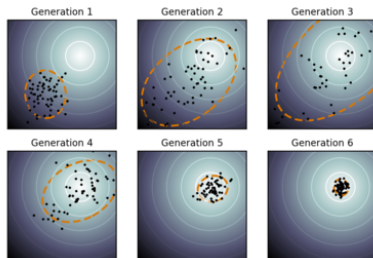
- Finds a local maximum: may be useful in certain cases.
- Simpler than gradient descent (no derivatives to compute), but less efficient.

CMA-ES

CMA-ES: Covariance Matrix Adaptation Evolution Strategy

Goal: find the maximum of a function f on a space X .

```
1 while continue do  
2    $x_1, \dots, x_n = \text{sample-multivariate-normal}(m, \sigma^2 C);$   
3    $x_1, \dots, x_n = \text{sort}(x_1, \dots, x_n, f);$   
4    $m, \sigma, C = \text{update}(x_1, \dots, x_n, m, \sigma, C);$   
   end  
5 return  $m$ 
```



Wikipedia

CMA-ES for falsification

Characteristics

- evolutionary algorithm (the parameters evolve towards a better sampler)
- sampling-based \rightsquigarrow no need to compute derivatives

Adaptation

Take X to be the space of signals σ and f to be $-\rho(\sigma, \varphi)$.

Remark

This is a way to adapt sampling-based optimisation to falsification, so it works for other such algorithms, such as the cross-entropy method.

Nelder-Mead method

Goal: find the maximum of a function f on a space X of dimension n .

```
1  $x_1, \dots, x_{n+1} = \text{simplex}()$ ;  
2 while continue do  
3    $x_i = \text{worst-vertex}(x_1, \dots, x_{n+1}, f)$ ;  
4    $x'_i = \text{reflect}(x_i, (x_j)_{j \neq i})$ ;  
   end  
5 return  $\text{best-vertex}(f)$ 
```

- 1: the structure is just a simplex.
- 3: when reflecting a vertex, if the new value is much better than the previous one, we keep stretching, otherwise, we shrink.

Adaptation to falsification

Take X a subspace of signals of finite dimension (e.g., fix a shape and control points), and $f = -\rho(\sigma, \varphi)$.

Ant colony

Goal: find a path in a graph that maximises performance.

```
1 while continue do
2   foreach "ant" i in colony do
3      $p_i = \text{construct-solution}(i)$ ;
4      $\text{local-update-pheromones}(i, f)$ ;
5   end
6    $p = \text{argmax}(p_i, f)$ 
7 end
8 return  $p$ 
```

- 3: ants walk randomly on the graph, choosing neighbours with more pheromone more often.
- 4: ants put pheromone on their chosen edge, making it more attractive; the better the ant's solution (i.e., the smaller $f(p_i)$), the more pheromone she puts.
- 5: pheromone gets put on all edges taken by the best solution.

Ant colony for falsification

Adaptation

Take G to be the graph:

- whose vertices are input signals σ ,
- there is an edge between σ and τ when they are “close enough”, say, $\|\sigma - \tau\|_\infty < C$, i.e., for all $t < t_{\max}$ and component x of U , $|\sigma(t)(x) - \tau(t)(x)| < C$,

and $f(\sigma_1 \dots \sigma_n)$ to be $-\rho(\sigma_n, \varphi)$.

The ant colony algorithm tries to find a path that maximises f , i.e., minimises $\rho(\sigma, \varphi)$.

Remark

This is a way to adapt optimisation of a function on a graph to falsification, so it works for other optimisation algorithms that work on graphs, such as simulated annealing.

Simulated annealing

Goal: find a global maximum of a function f on a space X .

$T = T_0$;

$i = 0$;

while $T(i) > 0$ **do**

$x_{i+1} = \text{neighbour}(x_i)$;

if $E(x_{i+1}) < E(x_i)$ and $\text{rand}() < P(E(x_{i+1}), E(x_i), T)$

then

$x_{i+1} = x_i$;

end

$i++$;

end

return x_{i-1}

- Avoids local minima by “cooling” the system down progressively.
- Many parameters \rightsquigarrow adaptive, but tricky.

Conclusion

- Falsification:
 - method to find counterexamples to a property,
 - useful in the world of formal methods,
 - black-box method,
 - relies on optimisation algorithms.
- Hybrid system:
 - continuous and discrete parameters,
 - non-linear behaviour,
 - very expressive.
- Formulas:
 - expressed in a temporal logic,
 - boolean and robustness semantics.