Quantitative Robustness Analysis of Sensor Attacks on Cyber-Physical Systems

Stephen Chong[1] Ruggero Lanotte[2] Massimo Merro[3]

Simone Tini[2] Jian Xiang[1]

[1] Harvard University
Boston, MA

[2] University of Insubria
Como, IT

[3] University of Verona
Verona, IT

Berinoro

June 26, 2023

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?
- Q2: How the loss of safety can be quantified in case of perturbations ?

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?
- Q2: How the loss of safety can be quantified in case of perturbations ?
- A1: We propose the notions of
  - quantitative forward safety, and
  - quantitative backward safety (not discussed in the talk).

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?
- Q2: How the loss of safety can be quantified in case of perturbations ?
- A1: We propose the notions of
  - quantitative forward safety, and
  - quantitative backward safety (not discussed in the talk).
- A2: We propose the notions of
  - forward robustness and
  - backward robustness (not discussed in the talk).

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?
- Q2: How the loss of safety can be quantified in case of perturbations ?
- A1: We propose the notions of
  - quantitative forward safety, and
  - quantitative backward safety (not discussed in the talk).
- A2: We propose the notions of
  - forward robustness and
  - backward robustness (not discussed in the talk).
- We work with:
  - Platzer's Hybrid Programs formalism,
  - Platzer's differential dynamic logic specification language

Scientific questions to be addressed and hypothesis

- Q1: How the safety of a CPS can be quantified in terms of preconditions/postconditions?
- Q2: How the loss of safety can be quantified in case of perturbations ?
- A1: We propose the notions of
  - quantitative forward safety, and
  - quantitative backward safety (not discussed in the talk).
- A2: We propose the notions of
  - forward robustness and
  - backward robustness (not discussed in the talk).
- We work with:
  - Platzer's Hybrid Programs formalism,
  - Platzer's differential dynamic logic specification language
- Perturbations will be attacks on sensors.

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

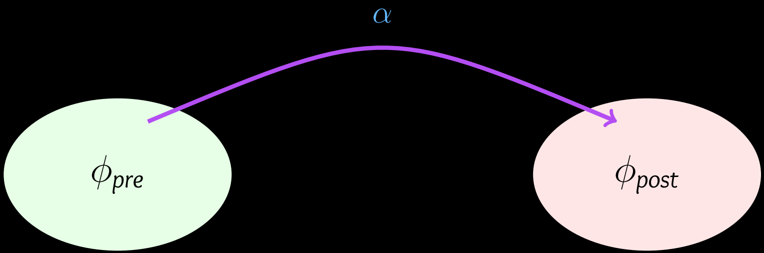

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

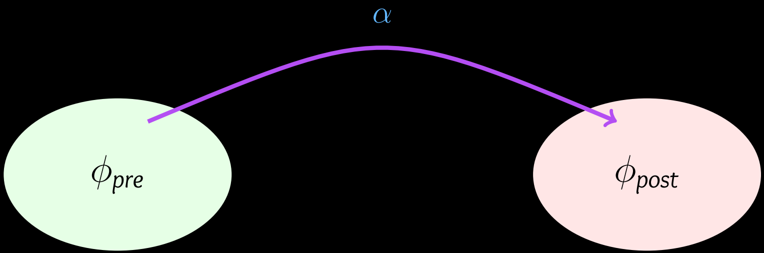

This definition does not provide any info about *how "good"* α is.

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

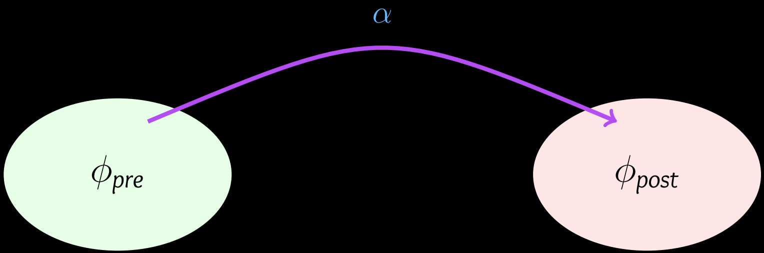

This definition does not provide any info about *how "good"* α is.
Two questions have no answer:

- Can we strenghten ϕ_{post}? How much?
- Can we weaken ϕ_{pre}? How much?

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

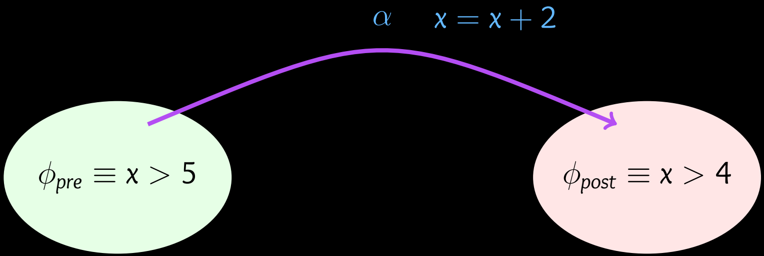

This definition does not provide any info about *how "good"* α is.
Two questions have no answer:

- Can we strenghten ϕ_{post}? How much?
- Can we weaken ϕ_{pre}? How much?

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

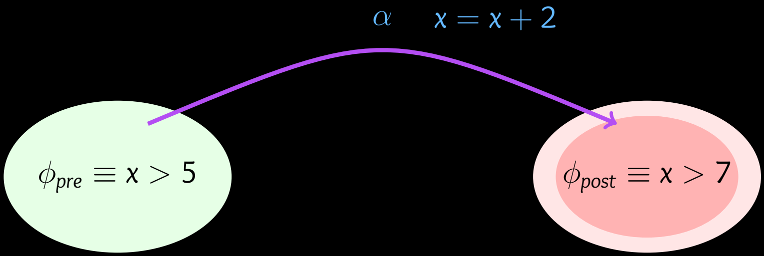

This definition does not provide any info about *how "good"* α is.
Two questions have no answer:

- Can we strenghten ϕ_{post}? How much?
- Can we weaken ϕ_{pre}? How much?

Qualitative safety

A program α is *safe* for ϕ_{post} assuming ϕ_{pre}, if $\phi_{pre} \rightarrow [\alpha]\phi_{post}$ holds.

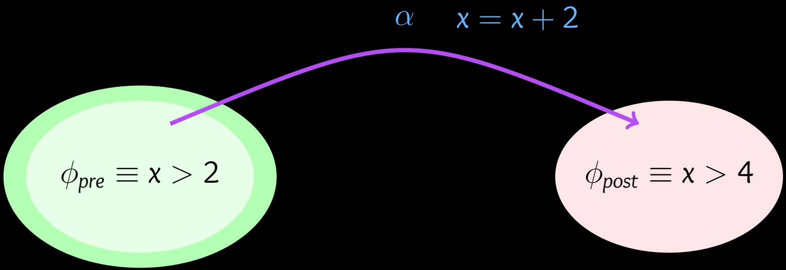

This definition does not provide any info about *how "good"* α is.
Two questions have no answer:

- Can we strengthen ϕ_{post}? How much?
- Can we weaken ϕ_{pre}? How much?

Forward quantitative safety

Assume a notion of *distance* between states.

Forward quantitative safety

Assume a notion of *distance* between states.

Let $\phi\langle\alpha\rangle$ denote the *strongest postcondition* after the execution of α in states satisfying ϕ.

Forward quantitative safety

Assume a notion of *distance* between states.

Let $\phi\langle\alpha\rangle$ denote the *strongest postcondition* after the execution of α in states satisfying ϕ.

Then, α is forward u-safe for ϕ_{pre} and ϕ_{post}, written $\text{F-SAFE}_u(\alpha, \phi_{pre}, \phi_{post})$, if $u = \inf\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle\alpha\rangle]\!]\}$.

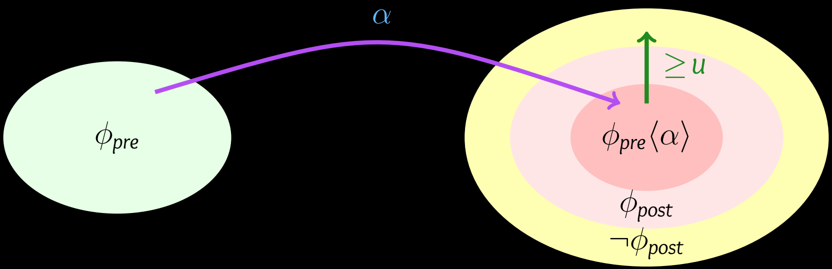

Forward quantitative safety

Assume a notion of *distance* between states.

Let $\phi\langle\alpha\rangle$ denote the *strongest postcondition* after the execution of α in states satisfying ϕ.

Then, α is forward u-safe for ϕ_{pre} and ϕ_{post}, written $\text{F-SAFE}_u(\alpha, \phi_{pre}, \phi_{post})$, if $u = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle\alpha\rangle]\!]\}$.

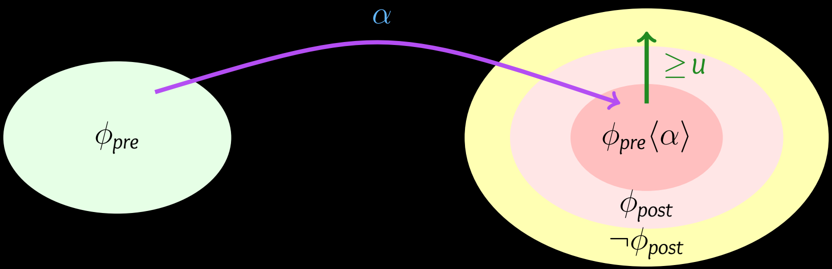

- u estimates how strong $\phi_{pre}\langle\alpha\rangle$ is with respect to ϕ_{post}.
- u estimates how much ϕ_{post} can be strengthened w.r.t. $\phi_{pre}\langle\alpha\rangle$.
- The bigger u is, the *safer* the program α is.

Forward robustness (w.r.t. some perturbation)

Assume a program α and a perturbation $\mathcal{P}$ s.t. $\phi_{pre}\langle\alpha\rangle \subseteq \phi_{pre}\langle\mathcal{P}(\alpha)\rangle$.
Then, α is forward δ-robust for ϕ_{pre} and ϕ_{post}, under $\mathcal{P}$, if

- $\text{F-SAFE}_u(\alpha, \phi_{pre}, \phi_{post})$
- $\text{F-SAFE}_{u_1}(\mathcal{P}(\alpha), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

Forward robustness (w.r.t. some perturbation)

Assume a program α and a perturbation $\mathcal{P}$ s.t. $\phi_{pre}\langle\alpha\rangle \subseteq \phi_{pre}\langle\mathcal{P}(\alpha)\rangle$.
Then, α is forward δ-robust for ϕ_{pre} and ϕ_{post}, under $\mathcal{P}$, if

- F-SAFE$_u(\alpha, \phi_{pre}, \phi_{post})$
- F-SAFE$_{u_1}(\mathcal{P}(\alpha), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

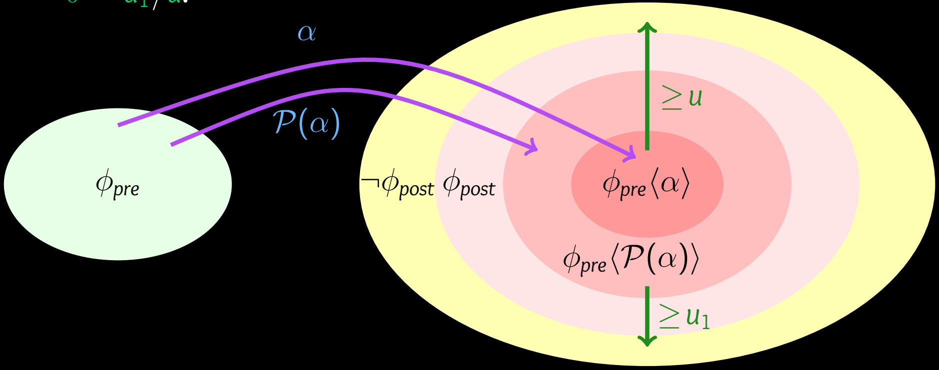

Forward robustness (w.r.t. some perturbation)

Assume a program α and a perturbation $\mathcal{P}$ s.t. $\phi_{pre}\langle\alpha\rangle \subseteq \phi_{pre}\langle\mathcal{P}(\alpha)\rangle$.
Then, α is forward δ-robust for ϕ_{pre} and ϕ_{post}, under $\mathcal{P}$, if

- $\text{F-SAFE}_u(\alpha, \phi_{pre}, \phi_{post})$
- $\text{F-SAFE}_{u_1}(\mathcal{P}(\alpha), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

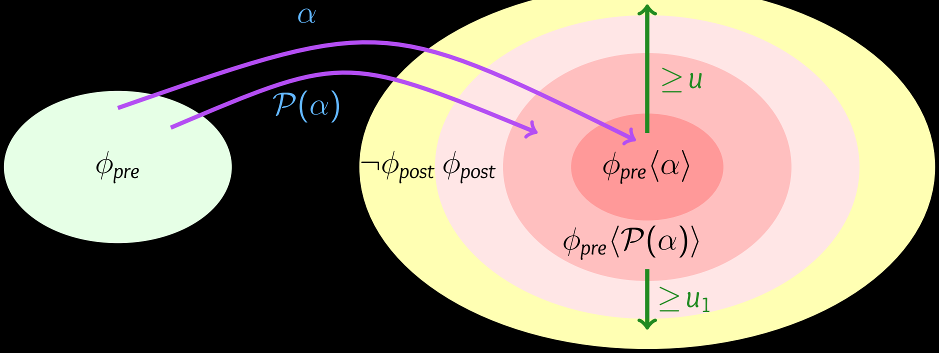

- δ is the percentage of forward safety that is maintained under $\mathcal{P}$.
- The closer δ is to 1, the more robust the system is.

An example: Platzer's autonomous vehicle

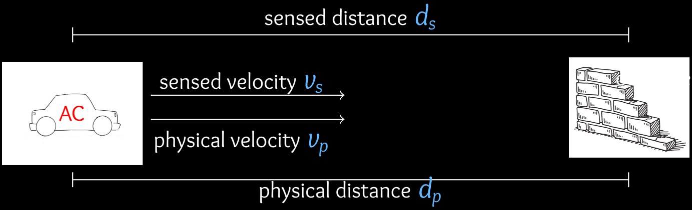

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

An example: Platzer's autonomous vehicle

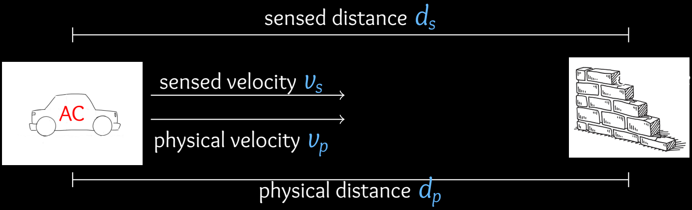

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

An example: Platzer's autonomous vehicle

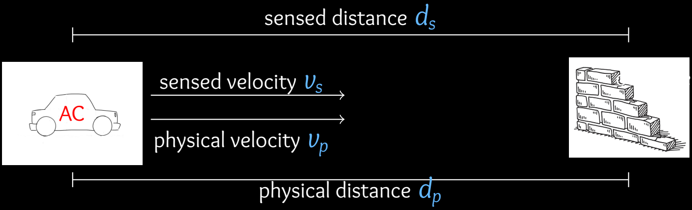

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

$$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl\,;\ plant)^*]\phi_{post}$$

An example: Platzer's autonomous vehicle

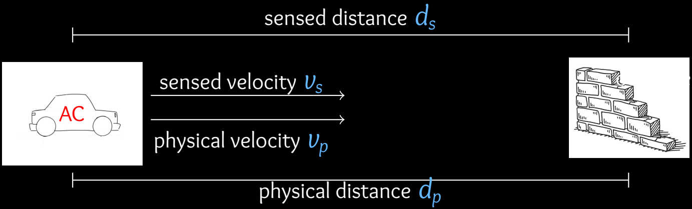

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

$$ctrl \equiv d_s := d_p \,;\, v_s := v_p \,;\, (accel \cup brake)$$

$$\phi_{safety} \equiv \phi_{pre} \to [(ctrl \,;\, plant)^*]\phi_{post}$$

An example: Platzer's autonomous vehicle

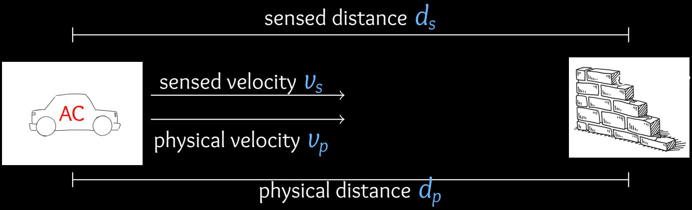

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

$$brake \equiv a := -B$$

$$ctrl \equiv d_s := d_p\,;\; v_s := v_p\,;\,(accel \cup brake)$$

$$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl\,;\; plant)^*]\phi_{post}$$

An example: Platzer's autonomous vehicle

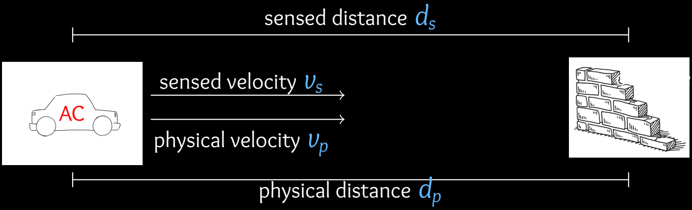

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p{+}2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

$\psi \equiv 2Bd_s > ((v_s{+}2)^2 + (A+B)(A\epsilon^2 + 2(v_s{+}2)\epsilon)$

$accel \equiv ?\psi \,;\, a := A$ // acceleration guarded by ψ

$brake \equiv a := -B$

$ctrl \equiv d_s := d_p \,;\, v_s := v_p \,;\, (accel \cup brake)$

$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl \,;\, plant)^*]\phi_{post}$

An example: Platzer's autonomous vehicle

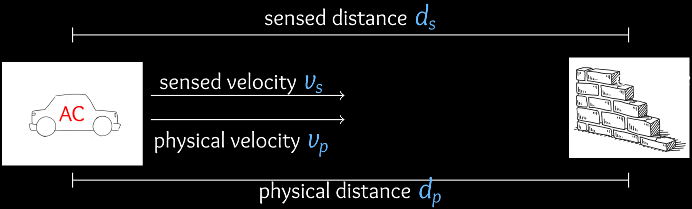

System constants:

accel. rate $A = 1$

braking rate $B = 1$

time constant $\epsilon = 1$

$\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$ // no crash if we break immediately

$\phi_{post} \equiv d_p > 0$ // there is no crash!

$\psi \equiv 2Bd_s > ((v_s+2)^2 + (A+B)(A\epsilon^2 + 2(v_s+2)\epsilon)$

$accel \equiv ?\psi \,;\, a := A$ // acceleration guarded by ψ

$brake \equiv a := -B$

$ctrl \equiv d_s := d_p \,;\, v_s := v_p \,;\, (accel \cup brake)$

$plant \equiv d_p{}' = -v_p, v_p{}' = a, t' = 1 \,\&(v_p \geq 0 \wedge t \leq \epsilon)$

$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl \,;\, plant)^*]\phi_{post}$

Vehicle's safety

Given

- postcondition $\phi_{post} \equiv d_p > 0$
- precondition $\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$

the autonomous vehicle enjoys forward 2-safety:

$$\text{F-SAFE}_2((ctrl\,;\ plant)^*, \phi_{pre}, \phi_{post})$$

Vehicle's safety

Given

- postcondition $\phi_{post} \equiv d_p > 0$
- precondition $\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$

the autonomous vehicle enjoys forward 2-safety:

$$\text{F-SAFE}_2((ctrl\,;\,plant)^*, \phi_{pre}, \phi_{post})$$

Notice that:

- safety is guaranteed since *accel* is guarded by ψ:
  $accel \equiv ?\psi\,;\,a := A$ with
  $\psi \equiv 2Bd_s > (v_s+2)^2 + (A+B)(A\epsilon^2 + 2(v_s+2)\epsilon)$
- without $+2$ there would be no room for perturbations.

Graphical intuition of vehicle's safety

Property F-SAFE$_2$($(ctrl\,;\ plant)^*$, ϕ_{pre}, ϕ_{post}) can be represented as:

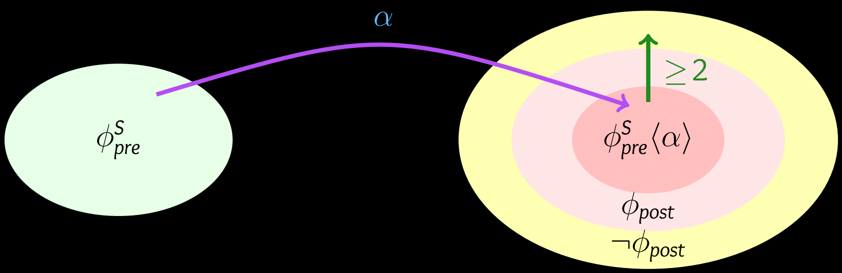

where:

- $\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$
- $\phi_{post} \equiv d_p > 0$
- $\alpha = (ctrl\,;\ plant)^*$

Bounded attack on velocity sensor

Assume an attack deviating the readings of v_s from v_p up to 1 m/s:

$$\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$$

$$\phi_{post} \equiv d_p > 0$$

$$\psi \equiv 2Bd_s > (v_s+2)^2 + (A+B)(A\epsilon^2 + 2(v_s+2)\epsilon)$$

$$accel \equiv ?\psi\,;\ a := A$$

$$brake \equiv a := -B$$

$$ctrl_A \equiv d_s := d_p\,;\ v_s := *;\ ?v_s \leq v_p + 1 \wedge v_s \geq v_p - 1\,;\ (accel \cup brake)$$

$$plant \equiv d_p{}' = -v_p, v_p{}' = a, t' = 1\,\&(v_p \geq 0 \wedge t \leq \epsilon)$$

$$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl_A\,;\ plant)^*]\phi_{post}$$

Bounded attack on velocity sensor

Assume an attack deviating the readings of v_s from v_p up to 1 m/s:

$$\phi_{pre} \equiv 2Bd_p > (v_p+2)^2 \wedge v_p \geq 0$$
$$\phi_{post} \equiv d_p > 0$$
$$\psi \equiv 2Bd_s > (v_s+2)^2 + (A+B)(A\epsilon^2 + 2(v_s+2)\epsilon)$$
$$accel \equiv ?\psi\,;\ a := A$$
$$brake \equiv a := -B$$
$$ctrl_A \equiv d_s := d_p\,;\ v_s := *;\ ?v_s \leq v_p + 1 \wedge v_s \geq v_p - 1\,;\ (accel \cup brake)$$
$$plant \equiv d_p{}' = -v_p, v_p{}' = a, t' = 1\,\&(v_p \geq 0 \wedge t \leq \epsilon)$$
$$\phi_{safety} \equiv \phi_{pre} \rightarrow [(ctrl_A\,;\ plant)^*]\phi_{post}$$

- The safety property $\text{F-SAFE}_2((ctrl_A\,;\ plant)^*, \phi_{pre}, \phi_{post})$ does not hold anymore.

Reasoning about robustness

We know that $\textsc{F-safe}_2((ctrl\,;\ plant)^*, \phi_{pre}, \phi_{post})$.
What about the robustness of $(ctrl\,;\ plant)^*$ under attack $\mathcal{P}$?

Reasoning about robustness

We know that $\text{F-SAFE}_2((\mathit{ctrl}\,;\ \mathit{plant})^*, \phi_{pre}, \phi_{post})$.
What about the robustness of $(\mathit{ctrl}\,;\ \mathit{plant})^*$ under attack $\mathcal{P}$?
According to the definition, $(\mathit{ctrl}\,;\ \mathit{plant})^*$ is forward δ-robust for properties ϕ_{pre} and ϕ_{post}, under that perturbation $\mathcal{P}$, if

- $\text{F-SAFE}_u((\mathit{ctrl}\,;\ \mathit{plant})^*, \phi_{pre}, \phi_{post})$,
- $\text{F-SAFE}_{u_1}(\mathcal{P}((\mathit{ctrl}\,;\ \mathit{plant})^*), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

Reasoning about robustness

We know that $\text{F-SAFE}_2((ctrl\,;\,plant)^*, \phi_{pre}, \phi_{post})$.
What about the robustness of $(ctrl\,;\,plant)^*$ under attack $\mathcal{P}$?
According to the definition, $(ctrl\,;\,plant)^*$ is forward δ-robust for properties ϕ_{pre} and ϕ_{post}, under that perturbation $\mathcal{P}$, if

- $\text{F-SAFE}_u((ctrl\,;\,plant)^*, \phi_{pre}, \phi_{post})$,
- $\text{F-SAFE}_{u_1}(\mathcal{P}((ctrl\,;\,plant)^*), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

Namely:

- $u = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle(ctrl\,;\,plant)^*\rangle]\!]\}$
- $u_1 = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle\mathcal{P}((ctrl\,;\,plant)^*)\rangle]\!]\}$

Reasoning about robustness

We know that $\textsf{F-SAFE}_2((\mathit{ctrl}\,;\ \mathit{plant})^*, \phi_{pre}, \phi_{post})$.
What about the robustness of $(\mathit{ctrl}\,;\ \mathit{plant})^*$ under attack $\mathcal{P}$?
According to the definition, $(\mathit{ctrl}\,;\ \mathit{plant})^*$ is forward δ-robust for properties ϕ_{pre} and ϕ_{post}, under that perturbation $\mathcal{P}$, if

- $\textsf{F-SAFE}_u((\mathit{ctrl}\,;\ \mathit{plant})^*, \phi_{pre}, \phi_{post})$,
- $\textsf{F-SAFE}_{u_1}(\mathcal{P}((\mathit{ctrl}\,;\ \mathit{plant})^*), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

Namely:

- $u = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle(\mathit{ctrl}\,;\ \mathit{plant})^*\rangle]\!]\}$
- $u_1 = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle\mathcal{P}((\mathit{ctrl}\,;\ \mathit{plant})^*)\rangle]\!]\}$

Computing these inf may be difficult, in particular for u_1, since $\mathcal{P}$ replaces a real with an element in a set of reals.

Reasoning about robustness

We know that $\text{F-SAFE}_2((ctrl\,;\, plant)^*, \phi_{pre}, \phi_{post})$.
What about the robustness of $(ctrl\,;\, plant)^*$ under attack $\mathcal{P}$?
According to the definition, $(ctrl\,;\, plant)^*$ is forward δ-robust for properties ϕ_{pre} and ϕ_{post}, under that perturbation $\mathcal{P}$, if

- $\text{F-SAFE}_u((ctrl\,;\, plant)^*, \phi_{pre}, \phi_{post})$,
- $\text{F-SAFE}_{u_1}(\mathcal{P}((ctrl\,;\, plant)^*), \phi_{pre}, \phi_{post})$
- $\delta = u_1/u$.

Namely:

- $u = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle(ctrl\,;\, plant)^*\rangle]\!]\}$
- $u_1 = \mathbf{inf}\{\mathbf{Dist}(\nu, [\![\phi_{post}]\!]) \mid \nu \in [\![\phi_{pre}\langle\mathcal{P}((ctrl\,;\, plant)^*)\rangle]\!]\}$

Computing these inf may be difficult, in particular for u_1, since $\mathcal{P}$ replaces a real with an element in a set of reals.
Possible solution: provide a notion of simulation distance between programs allowing us to give an upper bound to the loss of safety $u - u_1$.

Forward simulation distance

Assume a set of variables $\mathcal{H}$ and a distance over states $\rho_{\mathcal{H}}$.

E.g., for states ω and ν: $\rho_{\mathcal{H}}(\omega, \nu) = \sqrt{\sum_{x \in \mathcal{H}} (\omega(x) - \nu(x))^2}$.

Forward simulation distance

Assume a set of variables $\mathcal{H}$ and a distance over states $\rho_{\mathcal{H}}$.

E.g., for states ω and ν: $\rho_{\mathcal{H}}(\omega, \nu) = \sqrt{\sum_{x \in \mathcal{H}} (\omega(x) - \nu(x))^2}$.

Definition: Two programs α and β are at forward simulation distance d w.r.t. a formula ϕ_{pre} and $\mathcal{H}$, written

$$\beta \sqsubseteq^{\mathrm{F}}_{\phi_{pre}, \mathcal{H}, d} \alpha$$

if $\forall \nu_1 \in [\![\phi_{pre}\langle\beta\rangle]\!]\ \exists \nu_2 \in [\![\phi_{pre}\langle\alpha\rangle]\!]$ such that $\rho_{\mathcal{H}}(\nu_1, \nu_2) \leq d$.

Forward simulation distance

Assume a set of variables $\mathcal{H}$ and a distance over states $\rho_{\mathcal{H}}$.

E.g., for states ω and ν: $\rho_{\mathcal{H}}(\omega, \nu) = \sqrt{\sum_{x \in \mathcal{H}} (\omega(x) - \nu(x))^2}$.

Definition: Two programs α and β are at forward simulation distance d w.r.t. a formula ϕ_{pre} and $\mathcal{H}$, written

$$\beta \sqsubseteq^{\mathrm{F}}_{\phi_{pre}, \mathcal{H}, d} \alpha$$

if $\forall \nu_1 \in [\![\phi_{pre}\langle\beta\rangle]\!]\ \exists \nu_2 \in [\![\phi_{pre}\langle\alpha\rangle]\!]$ such that $\rho_{\mathcal{H}}(\nu_1, \nu_2) \leq d$.

Example, for $\beta = \mathcal{P}(\alpha)$:

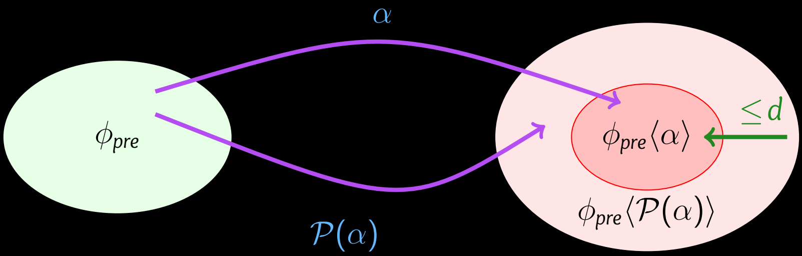

Upper bound to loss of safety

Theorem. Assume a hybrid program α and formulas ϕ_{pre} and ϕ_{post}. If

- $\text{F-SAFE}_{u}(\alpha, \phi_{pre}, \phi_{post})$ and
- $\mathcal{P}(\alpha) \sqsubseteq^{\mathrm{F}}_{\phi_{pre}, \text{VAR}(\phi_{post}), d} \alpha$

then: $\text{F-SAFE}_{\gamma}(\mathcal{P}(\alpha), \phi_{pre}, \phi_{post})$, with $\gamma \geq u - d$

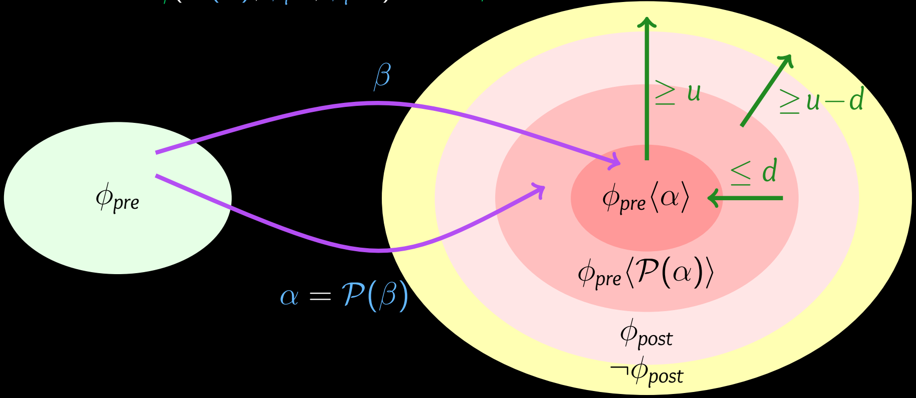

Upper bound to loss of safety

Theorem. Assume a hybrid program α and formulas ϕ_{pre} and ϕ_{post}. If

- $\text{F-SAFE}_u(\alpha, \phi_{pre}, \phi_{post})$ and
- $\mathcal{P}(\alpha) \sqsubseteq^{\mathrm{F}}_{\phi_{pre}, \mathrm{VAR}(\phi_{post}), d} \alpha$

then: $\text{F-SAFE}_\gamma(\mathcal{P}(\alpha), \phi_{pre}, \phi_{post})$, with $\gamma \geq u - d$

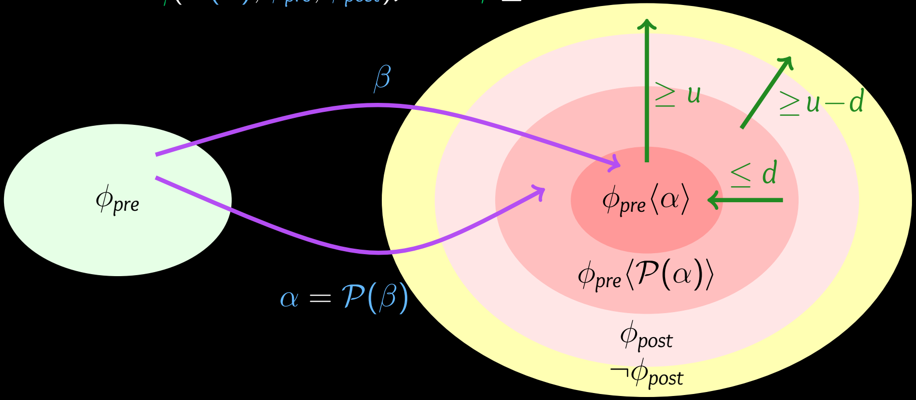

In words, d is an upper bound of the loss of forward safety.
Notice that α is γ-robust for $\gamma = {}^{(u-d)}/_{u}$.

Applying the theorem: An attempt

- By hand, we have computed

$$\mathcal{P}((ctrl\,;\,plant)^*) \sqsubseteq^{\mathrm{F}}_{\phi_{pre},\{d_p\},d} (ctrl\,;\,plant)^* \text{ with } d \leq 1.5$$

- Now, from

$$\text{F-SAFE}_2((ctrl\,;\,plant)^*, \phi_{pre}, \phi_{post})$$

and

$$\mathcal{P}((ctrl\,;\,plant)^*) \sqsubseteq^{\mathrm{F}}_{\phi_{pre},\{d_p\},d} (ctrl\,;\,plant)^* \text{ with } d \leq 1.5$$

we can conclude that:

$$\text{F-SAFE}_\gamma(\mathcal{P}((ctrl\,;\,plant)^*), \phi_{pre}, \phi_{post}) \text{ with } \gamma \geq 0.5$$

Open problem: How to compute forward simulation

Attempt 1: Encoding simulation distances with formulas

- Forward simulation distance is computed on states satisfying $\phi_{pre}\langle\alpha\rangle$ and $\phi_{pre}\langle\mathcal{P}(\alpha)\rangle$ and is encodable in a *forall exists* manner.

Open problem: How to compute forward simulation

Attempt 1: Encoding simulation distances with formulas

- Forward simulation distance is computed on states satisfying $\phi_{pre}\langle\alpha\rangle$ and $\phi_{pre}\langle\mathcal{P}(\alpha)\rangle$ and is encodable in a *forall exists* manner.
- More precisely:

$$(\phi_{pre}\langle\mathcal{P}(\alpha)\rangle \wedge (\bar{y} = \bar{x})) \rightarrow \exists\bar{x}.\ (\phi_{pre}\langle\alpha\rangle \wedge (\rho_{\mathcal{H}}(\bar{y}, \bar{x}) \leq d))$$

with $\bar{x}$ the variable in the formulae and $\bar{y}$ the fresh variables, implicely quantified universally, used to store the value of $\bar{x}$.

Open problem: How to compute forward simulation

Attempt 1: Encoding simulation distances with formulas

- Forward simulation distance is computed on states satisfying $\phi_{pre}\langle\alpha\rangle$ and $\phi_{pre}\langle\mathcal{P}(\alpha)\rangle$ and is encodable in a *forall exists* manner.
- More precisely:

$$(\phi_{pre}\langle\mathcal{P}(\alpha)\rangle \wedge (\bar{y} = \bar{x})) \rightarrow \exists\bar{x}.\ (\phi_{pre}\langle\alpha\rangle \wedge (\rho_{\mathcal{H}}(\bar{y}, \bar{x}) \leq d))$$

  with $\bar{x}$ the variable in the formulae and $\bar{y}$ the fresh variables, implicely quantified universally, used to store the value of $\bar{x}$.
- Unfortunately, this formula cannot be verified by using KeYmaera X.

Open problem: How to compute forward simulation-II

Attempt 1: Encoding simulation distances with formulas-II

- In our example, working by hand works:
- Having $\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$
  we have
  $\phi_{pre}\langle\mathcal{P}(\alpha)\rangle \equiv 2Bd_p > (v_p + 1)^2 \wedge v_p \geq 0$
  by using KeYmaera X we have proved that

$$2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0 \wedge d_p = fd_p \rightarrow$$

$$\exists d_p.(2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0 \wedge \sqrt{(d_p - fd_p)^2} \leq 1.5))$$

Open problem: How to compute forward simulation-II

Attempt 1: Encoding simulation distances with formulas-II

- In our example, working by hand works:
- Having $\phi_{pre} \equiv 2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0$
  we have
  $\phi_{pre}\langle \mathcal{P}(\alpha) \rangle \equiv 2Bd_p > (v_p + 1)^2 \wedge v_p \geq 0$
  by using KeYmaera X we have proved that

$$2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0 \wedge d_p = fd_p \rightarrow$$
$$\exists d_p.(2Bd_p > (v_p + 2)^2 \wedge v_p \geq 0 \wedge \sqrt{(d_p - fd_p)^2} \leq 1.5))$$

- From F-SAFE$_2((ctrl\,;\ plant)^*, \phi_{pre}, \phi_{post})$ and
  $\mathcal{P}((ctrl\,;\ plant)^*) \sqsubseteq^{\mathrm{F}}_{\phi_{pre},\{d_p\},d} (ctrl\,;\ plant)^*$ with $d \leq 1.5$ we get
  F-SAFE$_\gamma(\mathcal{P}((ctrl\,;\ plant)^*), \phi_{pre}, \phi_{post})$ with $\gamma \geq 0.5$.

Open problem: How to compute forward simulation-III

Attempt 2: Encoding simulation distances with modalities

- By using modalities we can directly encode program executions:

$$(\phi_{pre} \wedge \langle \mathcal{P}(\alpha) \rangle (\bar{y} = \bar{x})) \rightarrow$$

"for each state reachable from ϕ_{pre} by $\mathcal{P}(\alpha)$"

$$(\exists \bar{x}.\ \phi_{pre} \wedge \langle \alpha \rangle (\rho_{\mathcal{H}}(\bar{y}, \bar{x}) \leq d))$$

"there is an execution of α to a state at distance bounded by d."

Open problem: How to compute forward simulation-III

Attempt 2: Encoding simulation distances with modalities

- By using modalities we can directly encode program executions:

  $(\phi_{pre} \wedge \langle \mathcal{P}(\alpha) \rangle (\bar{y} = \bar{x})) \rightarrow$

  "for each state reachable from ϕ_{pre} by $\mathcal{P}(\alpha)$"

  $(\exists \bar{x}.\ \phi_{pre} \wedge \langle \alpha \rangle (\rho_{\mathcal{H}}(\bar{y}, \bar{x}) \leq d))$

  "there is an execution of α to a state at distance bounded by d."

- This is admitted by KeYmaera X syntax, but, in general, we have no answer

Open problem: How to compute forward simulation-III

Attempt 2: Encoding simulation distances with modalities

- By using modalities we can directly encode program executions:

$(\phi_{pre} \wedge \langle \mathcal{P}(\alpha) \rangle (\bar{y} = \bar{x})) \rightarrow$

"for each state reachable from ϕ_{pre} by $\mathcal{P}(\alpha)$"

$(\exists \bar{x}.\ \phi_{pre} \wedge \langle \alpha \rangle (\rho_{\mathcal{H}}(\bar{y}, \bar{x}) \leq d))$

"there is an execution of α to a state at distance bounded by d."

- This is admitted by KeYmaera X syntax, but, in general, we have no answer
- What we need is a proof system allowing us to give some upper bound to the simulation distance. We are on this but, presently, we have no solution.

Open problems - a more general view

- Developing a proof system for verifying properies encoding the simulation distance between programs.
- Dealing with more sophisticated sensor attacks, e.g. periodic attacks with several attack windows characterised by different tamperings.
- Dealing with different attacks.