

A Formally Verified Monitor for Metric First-Order Temporal Logic

Joshua Schneider, David Basin, Srđan Krstić, and Dmitriy Traytel

Department of Computer Science

ETH zürich

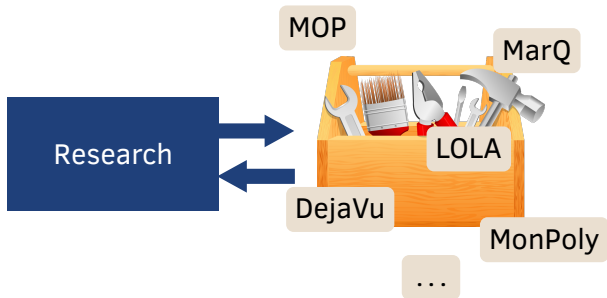
Goals of Runtime Verification

to study whether runtime application of formal methods is a viable **complement to** the traditional methods **proving** programs correct [...]


to study whether formality **improves** traditional **ad-hoc monitoring** techniques [...]

Source: www.runtime-verification.org (28/08/19, emphasis added)

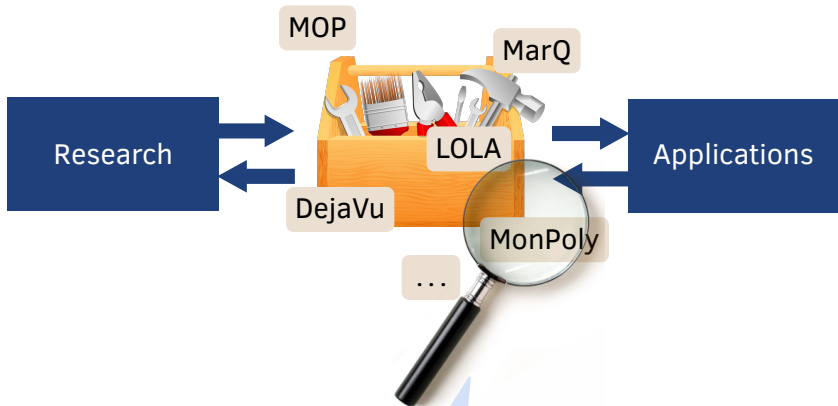
RV Tools



RV Tools



Verifying RV Tools



How can we prove that our tools are trustworthy?
Who guards the guardians?

Why Theorem Proving?

Machine-checked theorem proving is suitable for RV tools:



Criticality

Why Theorem Proving?

Machine-checked theorem proving is suitable for RV tools:



Criticality



Small size

Why Theorem Proving?


Machine-checked theorem proving is suitable for RV tools:



Criticality



Small size



Clear specification

Why Theorem Proving?


Machine-checked theorem proving is suitable for RV tools:



Criticality



Small size



Clear specification

All RV tools and should be verified formally.



Gain understanding of assumptions and guarantees!

Related Work

| | Language | Verified with | User effort |
|--------------------------|----------|---------------|--------------|
| Blech et al. (2012) | Regex | Coq | manual proof |
| Völlinger & Akili (2017) | — | Coq | manual proof |

Related Work

| | Language | Verified with | User effort |
|--------------------------|----------|---------------|--------------|
| Blech et al. (2012) | Regex | Coq | manual proof |
| Völlinger & Akili (2017) | — | Coq | manual proof |
| Laurent et al. (2015) | Copilot | SMT | semi-autom. |

Related Work

| | Language | Verified with | User effort |
|--------------------------|-----------------|------------------------------------|--------------|
| Blech et al. (2012) | Regex | Coq | manual proof |
| Völlinger & Akili (2017) | — | Coq | manual proof |
| Laurent et al. (2015) | Copilot | SMT | semi-autom. |
| Rizaldi et al. (2017) | LTL | Isabelle/HOL | none |
| Bohrer et al. (2018) | d \mathcal{L} | KeYmaera X Isabelle/HOL HOL4 | none |

Related Work

| | Language | Verified with | User effort |
|--------------------------|-----------------|------------------------------------|--------------|
| Blech et al. (2012) | Regex | Coq | manual proof |
| Völlinger & Akili (2017) | — | Coq | manual proof |
| Laurent et al. (2015) | Copilot | SMT | semi-autom. |
| Rizaldi et al. (2017) | LTL | Isabelle/HOL | none |
| Bohrer et al. (2018) | d \mathcal{L} | KeYmaera X Isabelle/HOL HOL4 | none |
| this work | MFOTL | Isabelle/HOL | none |

Our Contribution

Verimon: verified **MonPoly** (w/o optimizations)

- Formally verified monitor for metric first-order temporal logic (MFOTL)
- Expressive language with intervals and data quantification
- Proved correct for all instances of the monitor
- Explain and clarify MonPoly's algorithm



Our Contribution

Verimon: verified MonPoly (w/o optimizations)

- Formally verified monitor for metric first-order temporal logic (MFOTL)
- Expressive language with intervals and data quantification
- Proved correct for all instances of the monitor
- Explain and clarify MonPoly's algorithm



Basis for exploration:

- Monitor state manipulation
[ATVA'19]
- Foundation for future
extensions and optimizations

Our Contribution

Verimon: verified MonPoly (w/o optimizations)

- Formally verified monitor for metric first-order temporal logic (MFOTL)
- Expressive language with intervals and data quantification
- Proved correct for all instances of the monitor
- Explain and clarify MonPoly's algorithm



Basis for exploration:


- Monitor state manipulation [ATVA'19]
- Foundation for future extensions and optimizations

Differential testing case study:


- Used Verimon as oracle to test unverified implementations
- Tested MonPoly and DejaVu
- Found bugs!



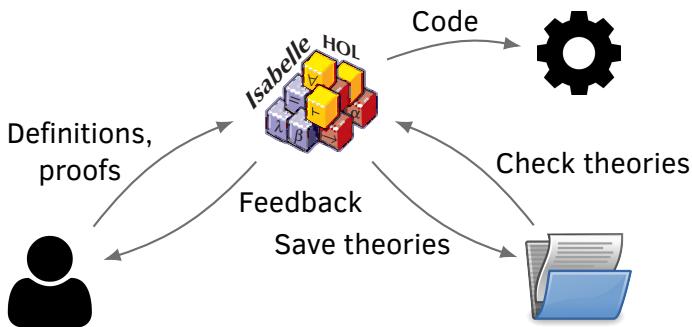
Background: Isabelle/HOL




Background: Isabelle/HOL



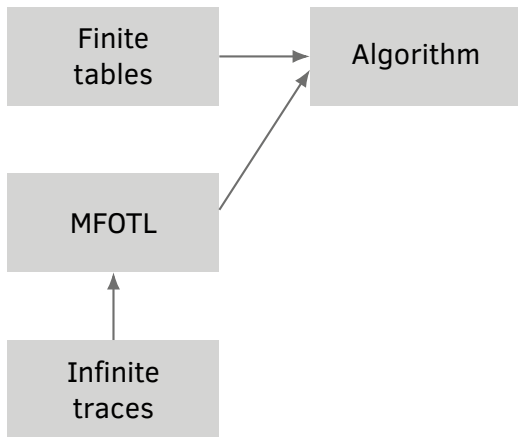
Background: Isabelle/HOL




Formalization Overview



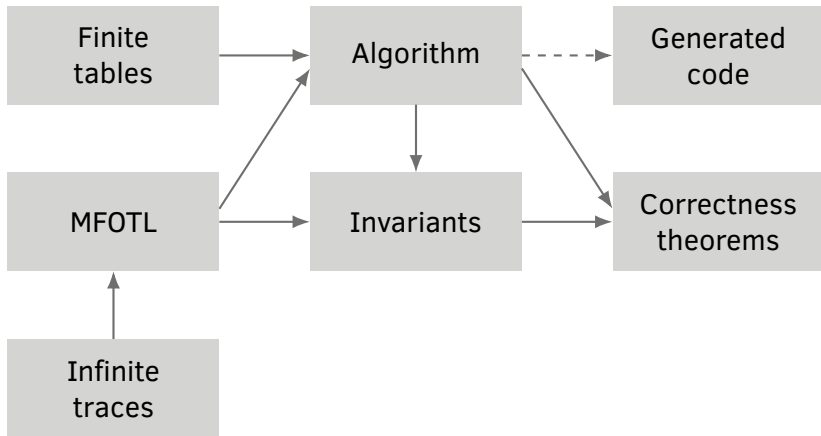
Formalization Overview



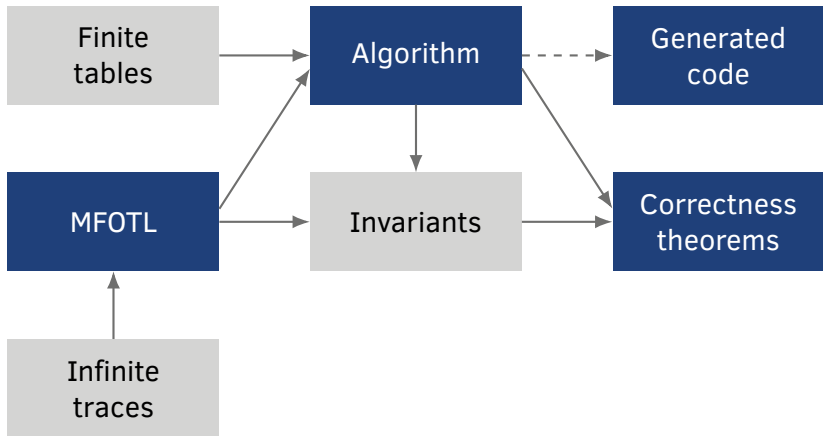
Formalization Overview



Formalization Overview



Formalization Overview



Background: MFOTL

TL

Temporal Logic: linear time ● ○ ◆ ◇ ■ □ S U

□(access → (¬release S acquire))

Background: MFOTL

FO

First-Order: data and quantification

$\square \forall x. \text{access}(x) \rightarrow (\neg \text{release}(x) \text{ S acquire}(x))$

TL

Temporal Logic: linear time ● ○ ◆ ◇ ■ □ S U

$\square (\text{access} \rightarrow (\neg \text{release} \text{ S acquire}))$

Background: MFOTL

M

Metric: time intervals

$\Box \forall x. \text{access}(x) \rightarrow (\neg \text{release}(x) S_{[0,1s]} \text{acquire}(x))$

FO

First-Order: data and quantification

$\Box \forall x. \text{access}(x) \rightarrow (\neg \text{release}(x) S \text{acquire}(x))$

TL

Temporal Logic: linear time ● ○ ◆ ◇ ■ □ S U

$\Box (\text{access} \rightarrow (\neg \text{release} S \text{acquire}))$

Background: MFOTL

M

Metric: time intervals

$\neg \forall x \text{ access}(x) \rightarrow (\neg \text{release}(x) \text{ S acquire}(x))$

Monitorable fragment:

- Safety properties ($\Box \phi$ with bounded future)

NOT: $\Box (\text{open} \rightarrow \Diamond \text{close})$

- Finitely evaluable violations

NOT: $\Box \forall x. P(x)$

$\Box (\text{access} \rightarrow (\neg \text{release S acquire}))$

Monitor Output

Checking the specification $\Box \forall \bar{x}. \varphi(\bar{x})$:

output whether $\models_{\sigma} \Box \forall \bar{x}. \varphi(\bar{x})$

Monitor Output

Checking the specification $\square \forall \bar{x}. \varphi(\bar{x})$:

output whether $\models_{\sigma} \square \forall \bar{x}. \varphi(\bar{x})$

Reporting violating points of $\forall \bar{x}. \varphi(\bar{x})$:

output all i s.t. $i \not\models_{\sigma} \forall \bar{x}. \varphi(\bar{x})$

Monitor Output

Checking the specification $\Box \forall \bar{x}. \varphi(\bar{x})$:

output whether $\models_{\sigma} \Box \forall \bar{x}. \varphi(\bar{x})$

Reporting violating points of $\forall \bar{x}. \varphi(\bar{x})$:

output all i s.t. $i \not\models_{\sigma} \forall \bar{x}. \varphi(\bar{x})$

\iff Reporting satisfying points of the negation $\exists \bar{x}. \neg \varphi(\bar{x})$:

output all i s.t. $i \models_{\sigma} \exists \bar{x}. \neg \varphi(\bar{x})$

Monitor Output

Checking the specification $\Box \forall \bar{x}. \varphi(\bar{x})$:

output whether $\models_{\sigma} \Box \forall \bar{x}. \varphi(\bar{x})$

Reporting violating points of $\forall \bar{x}. \varphi(\bar{x})$:

output all i s.t. $i \not\models_{\sigma} \forall \bar{x}. \varphi(\bar{x})$


\iff Reporting satisfying points of the negation $\exists \bar{x}. \neg \varphi(\bar{x})$:

output all i s.t. $i \models_{\sigma} \exists \bar{x}. \neg \varphi(\bar{x})$


Reporting satisfying points and assignments of $\neg \varphi(\bar{x})$:

output all (i, \bar{x}) s.t. $i, \bar{x} \models_{\sigma} \neg \varphi(\bar{x})$


Monitor Interface



Monitor Interface



Monitor Interface




type *event* = *string* × *domain list*

type *database* = *event set*

type *ts* = *nat*

type *prefix* = {*p* :: (*database* × *ts*) *list*. sorted (map snd *p*)}

Monitor Interface



type *output* = (*nat* × *tuple*) set

all pairs (i, \bar{x}) such that $i, \bar{x} \models_{\sigma} \neg\varphi(\bar{x})$

type *event* = *string* × *domain list*

type *database* = *event set*

type *ts* = *nat*

type *prefix* = {*p* :: (*database* × *ts*) list. sorted (map snd *p*)}

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec} :: \text{formula} \Rightarrow \text{prefix} \Rightarrow \text{output}$ **where**

$\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi$ *t is assignment to free variables*

$\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi$ *t is assignment to free variables*

$\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi = \{(i, t) \mid \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Specification

Define the expected output of the monitor algorithm:

definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$

$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$


$(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Specification

Define the expected output of the monitor algorithm:


definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$
 $(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Specification

Define the expected output of the monitor algorithm:


definition $\text{spec } \varphi \pi = \{(i, t). \text{wf_tuple } \varphi t \wedge$
 $(\forall \sigma. \text{prefix_of } \pi \sigma \rightarrow i < \text{progress } \sigma \varphi (\text{len } \pi) \wedge \text{sat } \sigma t i \varphi)\}$

all infinite extensions σ of π

MFOTL semantics

delay verdicts for formulas with future modalities

$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Correctness (1)

Does the `spec` function characterize a reasonable monitor?

Correctness (1)

Does the `spec` function characterize a reasonable monitor?

Fix an event stream σ and a prefix π (i.e., `prefix_of π σ` is true).

Soundness:

$(i, t) \in \text{spec } \varphi \pi$ implies `sat σ t i φ` .

Correctness (1)

Does the `spec` function characterize a reasonable monitor?

Fix an event stream σ and a prefix π (i.e., `prefix_of π σ` is true).


Soundness:

$(i, t) \in \text{spec } \varphi \pi$ implies `sat σ t i φ` .

Eventual completeness:

If $i < \text{len } \pi$ and `wf_tuple φ t` and $\forall \sigma'. \text{prefix_of } \pi \sigma' \rightarrow \text{sat } \sigma' t i \varphi$, then there exists a prefix π' of σ such that $(i, t) \in \text{spec } \varphi \pi'$.

$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$



Correctness (1)

Does the `spec` function characterize a reasonable monitor?

Fix an event stream σ and a prefix π (i.e., `prefix_of π σ` is true).


Soundness:

$(i, t) \in \text{spec } \varphi \pi$ implies `sat σ t i φ` .

Eventual completeness:

If $i < \text{len } \pi$ and `wf_tuple φ t` and $\forall \sigma'. \text{prefix_of } \pi \sigma' \rightarrow \text{sat } \sigma' t i \varphi$, then there exists a prefix π' of σ such that $(i, t) \in \text{spec } \varphi \pi'$.


$$\varphi \equiv P \wedge \Diamond_{[0,10s]} Q$$



Implementation

Online interface (unbounded stream):

definition `init` :: *formula* \Rightarrow *state*




Implementation

Online interface (unbounded stream):

definition $\text{init} :: \text{formula} \Rightarrow \text{state}$

definition $\text{step} :: \text{database} \times \text{ts} \Rightarrow \text{state} \Rightarrow \text{output} \times \text{state}$




Implementation

Online interface (unbounded stream):

definition $\text{init} :: \text{formula} \Rightarrow \text{state}$

definition $\text{step} :: \text{database} \times \text{ts} \Rightarrow \text{state} \Rightarrow \text{output} \times \text{state}$




Implementation

Online interface (unbounded stream):


definition $\text{init} :: \text{formula} \Rightarrow \text{state}$

definition $\text{step} :: \text{database} \times \text{ts} \Rightarrow \text{state} \Rightarrow \text{output} \times \text{state}$



Offline interface (finite prefix):

definition $\text{monitor} :: \text{formula} \Rightarrow \text{prefix} \Rightarrow \text{output}$



Correctness (2)

Is the implementation correct?

1. `init` establishes the invariant `wf_state`:

φ is monitorable implies `wf_state` φ [] (`init` φ).

Correctness (2)

Is the implementation correct?

1. **init** establishes the invariant `wf_state`:

φ is monitorable implies `wf_state` φ [] (**init** φ).

2. **step** preserves the invariant:

Let **step** $(db, \tau) st = (X, st')$. If `wf_state` $\varphi \pi st$ and `last_ts` $\pi \leq \tau$, then `wf_state` $\varphi (\pi @ [(db, \tau)]) st', \dots$

Correctness (2)

Is the implementation correct?

1. **init** establishes the invariant **wf_state**:

φ is monitorable implies **wf_state** φ [] (**init** φ).

2. **step** preserves the invariant:

Let **step** $(db, \tau) st = (X, st')$. If **wf_state** $\varphi \pi st$ and **last_ts** $\pi \leq \tau$, then **wf_state** $\varphi (\pi @ [(db, \tau)]) st', \dots$

3. **step**'s output corresponds to **spec**:

\dots and $X = \text{spec } \varphi (\pi @ [(db, \tau)]) - \text{spec } \varphi \pi$.

Correctness (2)

Is the implementation correct?

1. **init** establishes the invariant **wf_state**:

φ is monitorable implies **wf_state** φ [] (**init** φ).

2. **step** preserves the invariant:

Let **step** $(db, \tau) st = (X, st')$. If **wf_state** $\varphi \pi st$ and **last_ts** $\pi \leq \tau$, then **wf_state** $\varphi (\pi @ [(db, \tau)]) st', \dots$

3. **step**'s output corresponds to **spec**:

\dots and $X = \mathbf{spec} \varphi (\pi @ [(db, \tau)]) - \mathbf{spec} \varphi \pi$.

4. **monitor** $\varphi \pi = \mathbf{spec} \varphi \pi$ (if φ is monitorable)

Executable Monitor

Current approach:

- Extract OCaml code from formalization using Isabelle/HOL's code generator
- Reuse MonPoly's parser/data structures and add "glue"

Executable Monitor

Current approach:

- Extract OCaml code from formalization using Isabelle/HOL's code generator
- Reuse MonPoly's parser/data structures and add "glue"

Trust assumptions:

- Isabelle's kernel and code generator
- Parser and glue code
- OCaml compiler, runtime environment etc.

Executable Monitor

Current approach:

- Extract OCaml code from formalization using Isabelle/HOL's code generator
- Reuse MonPoly's parser/data structures and add "glue"

Trust assumptions:

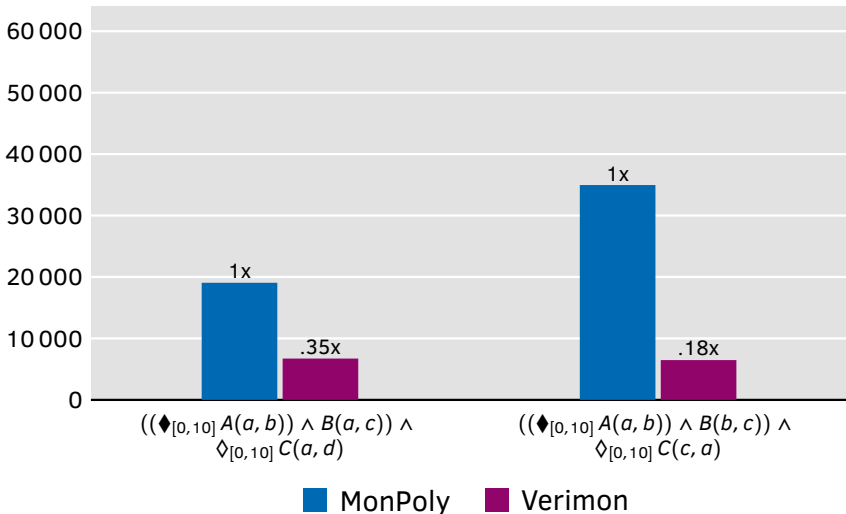
- Isabelle's kernel and code generator
- Parser and glue code
- OCaml compiler, runtime environment etc.

Satisfactory?

- The algorithm is the challenging part
- Various techniques for full-stack verification exist, for example CakeML (used in VeriPhy)

Performance

Event throughput [1/s] – higher is better



Differential Testing

Idea: Find bugs in unverified implementations by comparing their output on random inputs with Verimon.

Two targets: **MonPoly** and **DejaVu**

Differential Testing

Idea: Find bugs in unverified implementations by comparing their output on random inputs with Verimon.

Two targets: **MonPoly** and **DejaVu**

- Random formulas parameterized by size n , free variables FV
- Generated 1000 formulas each for $2 \leq n \leq 5$, $|FV| \leq 6$

Differential Testing

Idea: Find bugs in unverified implementations by comparing their output on random inputs with Verimon.

Two targets: **MonPoly** and **DejaVu**

- Random formulas parameterized by size n , free variables FV
- Generated 1000 formulas each for $2 \leq n \leq 5$, $|FV| \leq 6$
- Random prefixes with 20, 40, 60, 100 databases
- Reuse recent event parameters with probability p

Results

Two bugs found in **MonPoly**:

1. **Wrong output** for class of formulas, for example

$Q(x, y) \wedge \neg(P(x) \text{ S } Q(y, x))$ on prefix with only $Q(1, 2)$

Results

Two bugs found in **MonPoly**:

1. **Wrong output** for class of formulas, for example

$Q(x, y) \wedge \neg(P(x) S Q(y, x))$ on prefix with only $Q(1, 2)$

2. Additional violation output for finite traces

@0. (time point 0): true

@MaxTS (time point 1): true

Results

Two bugs found in **MonPoly**:

1. **Wrong output** for class of formulas, for example

$Q(x, y) \wedge \neg(P(x) \text{ S } Q(y, x))$ on prefix with only $Q(1, 2)$

2. Additional violation output for finite traces

@0. (time point 0): true

@MaxTS (time point 1): true

Documented differences in **DejaVu**'s semantics:

3. Arithmetic relations change semantics of quantifiers, e.g.,

$\neg\varphi$ vs. $\neg\exists x. \varphi \wedge x = 42$

4. Active domain does not include constants in the formula, e.g.,

$\neg\exists x. x = 42 \wedge \neg P(x)$ on $P(101)$

Ongoing and Future Work

Achieve parity with MonPoly:

- Sliding window algorithm
- Refinement to imperative data structures
- Aggregations (count, sum, max, ...)


New and verified optimizations:

- Multi-way joins (completed by Thibault Dardinier)


New features:

- State splitting and merging [ATVA'19]
- MFODL – adds regular expressions

A Formally Verified Monitor for MFOTL


| | Language | Verified with | User effort |
|----------------|----------|---|-------------|
| Verimon | MFOTL |  | none |

A Formally Verified Monitor for MFOTL

| | Language | Verified with | User effort |
|----------------|----------|---|-------------|
| Verimon | MFOTL |  | none |



A Formally Verified Monitor for MFOTL

| | Language | Verified with | User effort |
|----------------|----------|---|-------------|
| Verimon | MFOTL |  | none |



Questions?

Joshua Schneider David Basin
Srđan Krstić Dmitriy Traytel

ETH zürich