

Theory-Specific Proof Steps

Witnessing Correctness of SMT Executions

Rodrigo Otoni¹, Martin Blicha^{1,2}, Patrick Eugster^{1,3,4},
Antti E. J. Hyvärinen¹, and Natasha Sharygina¹

¹ Università della Svizzera italiana, Lugano, Switzerland

² Charles University, Prague, Czech Republic

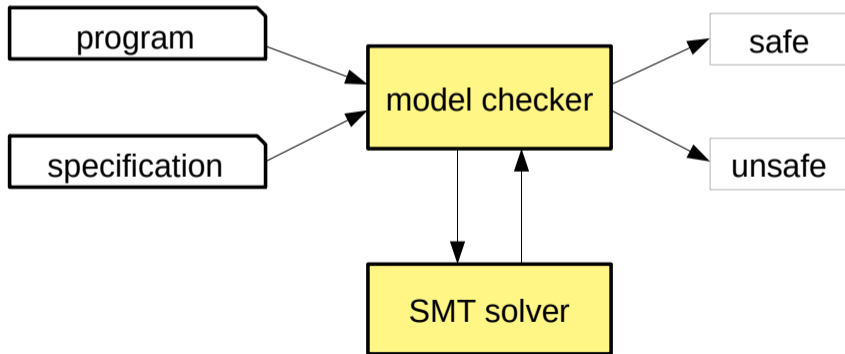
³ Technische Universität Darmstadt, Darmstadt, Germany

⁴ Purdue University, West Lafayette, USA

18th July 2021

- Software verification often relies on SMT solvers

- Software verification often relies on SMT solvers

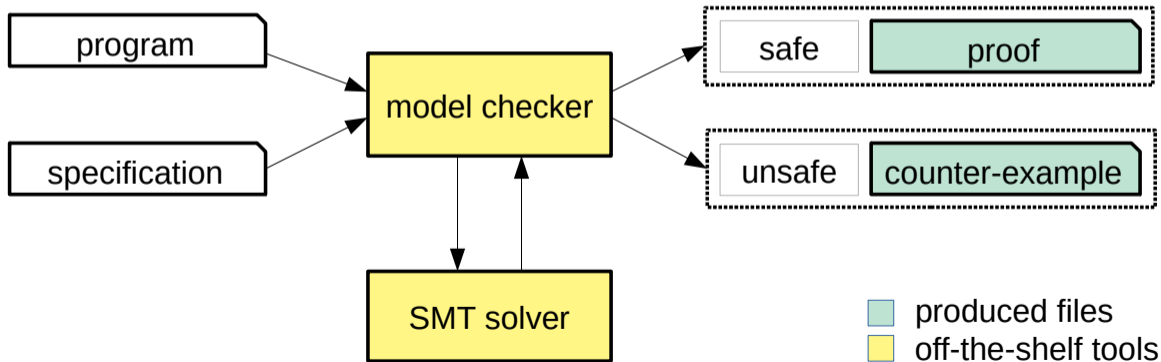


State of the Art in Verification

- Despite use in verification, SMT solvers have bugs
- At SMT-COMP'20, solvers disagreed on 149 instances with unknown status
- Guarantees are needed to trust solvers' results

State of the Art in Verification

- Despite use in verification, SMT solvers have bugs
- At SMT-COMP'20, solvers disagreed on 149 instances with unknown status
- Guarantees are needed to trust solvers' results



Proofs of unsatisfiability

- The *de facto* standard for SAT proofs is the DRAT proof format
- No standard for SMT proofs, with many competing formats available

Proofs of unsatisfiability

- The *de facto* standard for SAT proofs is the DRAT proof format
- No standard for SMT proofs, with many competing formats available

Existing solver support

- Proofs validating SMT solvers' executions vary by implementation
 - Proof-producing solvers include CVC4, veriT, and Z3
- A commonality is the focus on integration with interactive theorem provers
 - SMT proofs can be consumed by Isabelle/HOL and Coq
- Current proof formats unsuitable to automation
 - Integration to automated tools such as model checkers is not available

Lightweight witnesses validating SMT solvers' executions

- Witnesses checkable by simple automated checkers
- Checkers simple enough to support manual inspection
- A user can write a checker in a few hours/days

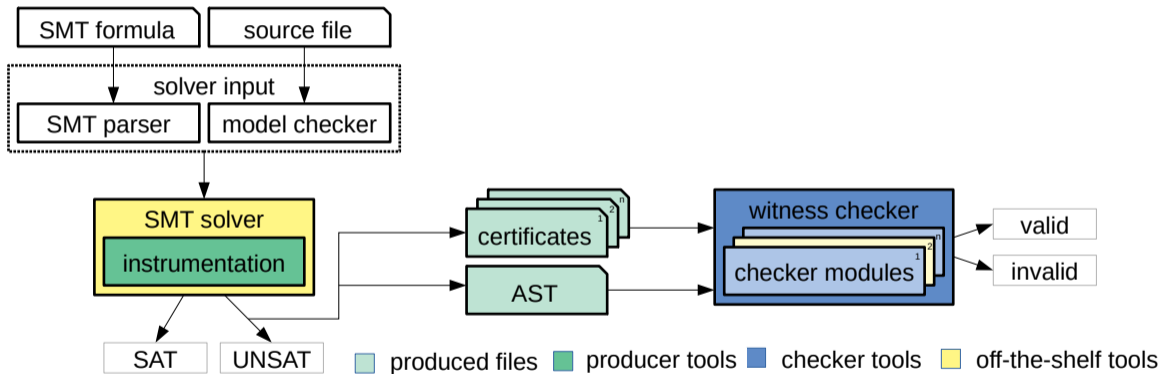
Lightweight witnesses validating SMT solvers' executions

- Witnesses checkable by simple automated checkers
- Checkers simple enough to support manual inspection
- A user can write a checker in a few hours/days

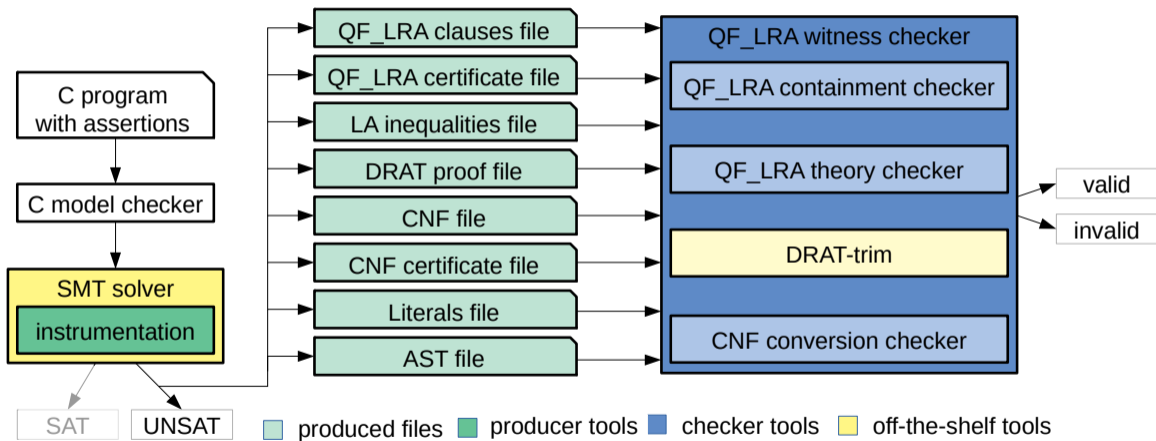
Validation based on the foundations of SMT algorithms

- DRAT proofs for propositional logic
- Theory-specific witnesses for each supported SMT theory

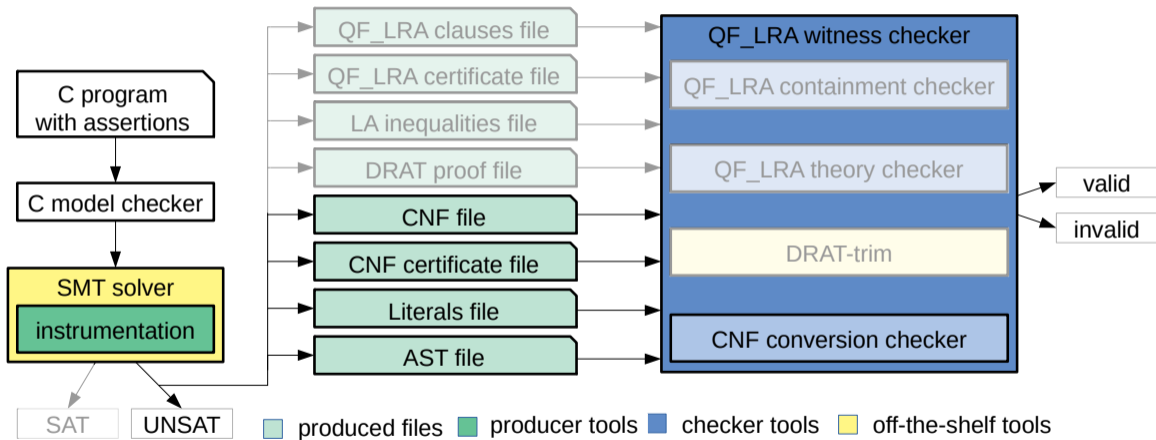
Certification Framework



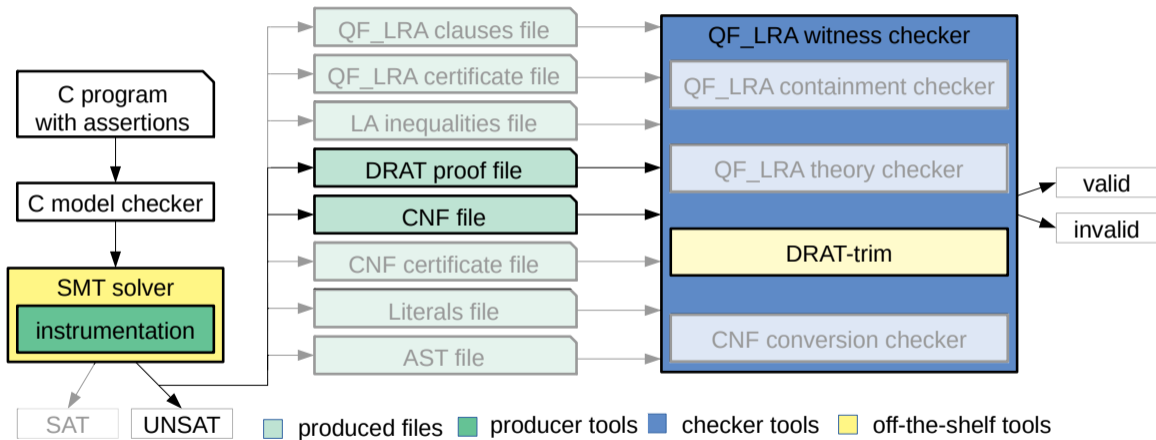
Example for QF_LRA



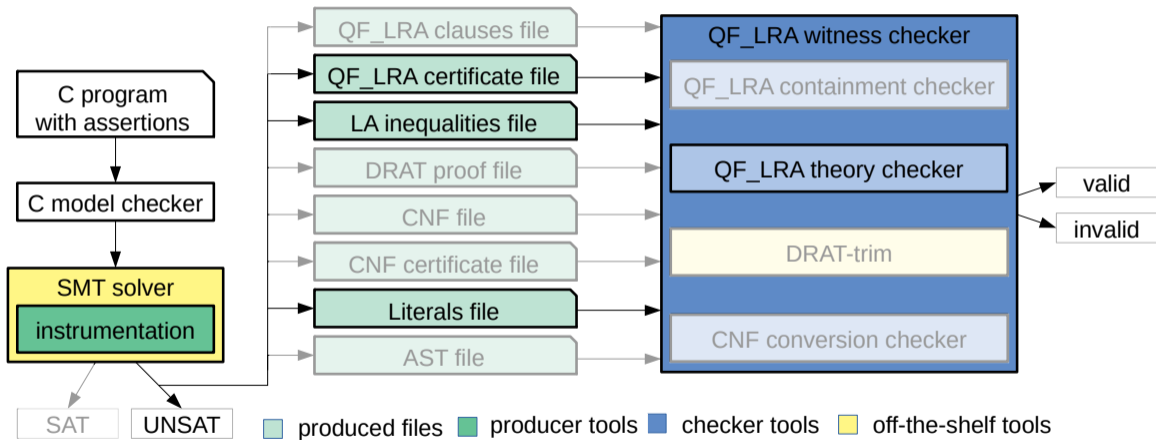
Example for QF_LRA



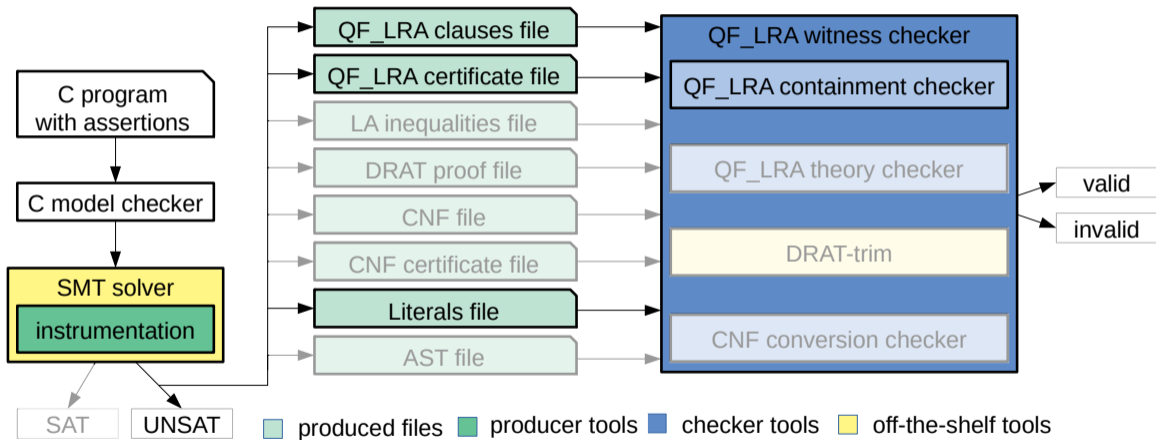
Example for QF_LRA



Example for QF_LRA



Example for QF_LRA



Example for QF_LRA - Certificate File

```
1 float x = 1;
2 if (condition) {
3     x = x - 1;
4 } else {
5     x = x + 1;
6 }
7 assert (x ≥ 0);
```

$$\begin{aligned} & x1 = 1 \wedge \\ & (condition \rightarrow x4 = x2) \wedge x2 = x1 - 1 \wedge \\ & (\neg condition \rightarrow x4 = x3) \wedge x3 = x1 + 1 \wedge \\ & \neg(x4 \geq 0) \end{aligned}$$

Example for QF_LRA - Certificate File

```
1 float x = 1;
2 if (condition) {
3     x = x - 1;
4 } else {
5     x = x + 1;
6 }
7 assert (x ≥ 0);
```

$$\begin{aligned} & x1 = 1 \wedge \\ & (condition \rightarrow x4 = x2) \wedge x2 = x1 - 1 \wedge \\ & (\neg condition \rightarrow x4 = x3) \wedge x3 = x1 + 1 \wedge \\ & \neg(x4 \geq 0) \wedge \\ & (x4 \geq 0 \vee x4 < 2) \end{aligned}$$

Learned clause

$x4 \geq 0 \vee x4 < 2$

Example for QF_LRA - Certificate File

```
1 float x = 1;
2 if (condition) {
3     x = x - 1;
4 } else {
5     x = x + 1;
6 }
7 assert (x ≥ 0);
```

$$\begin{aligned} & x1 = 1 \wedge \\ & (condition \rightarrow x4 = x2) \wedge x2 = x1 - 1 \wedge \\ & (\neg condition \rightarrow x4 = x3) \wedge x3 = x1 + 1 \wedge \\ & \neg(x4 \geq 0) \wedge \\ & (x4 \geq 0 \vee x4 < 2) \end{aligned}$$

Learned clause

$$x4 \geq 0 \vee x4 < 2$$

Witness

$$x4 < 0 \wedge x4 \geq 2$$

Example for QF_LRA - Certificate File

```
1 float x = 1;
2 if (condition) {
3     x = x - 1;
4 } else {
5     x = x + 1;
6 }
7 assert (x ≥ 0);
```

$$\begin{aligned} & x1 = 1 \wedge \\ & (condition \rightarrow x4 = x2) \wedge x2 = x1 - 1 \wedge \\ & (\neg condition \rightarrow x4 = x3) \wedge x3 = x1 + 1 \wedge \\ & \neg(x4 \geq 0) \wedge \\ & (x4 \geq 0 \vee x4 < 2) \end{aligned}$$

Learned clause

$$x4 \geq 0 \vee x4 < 2$$

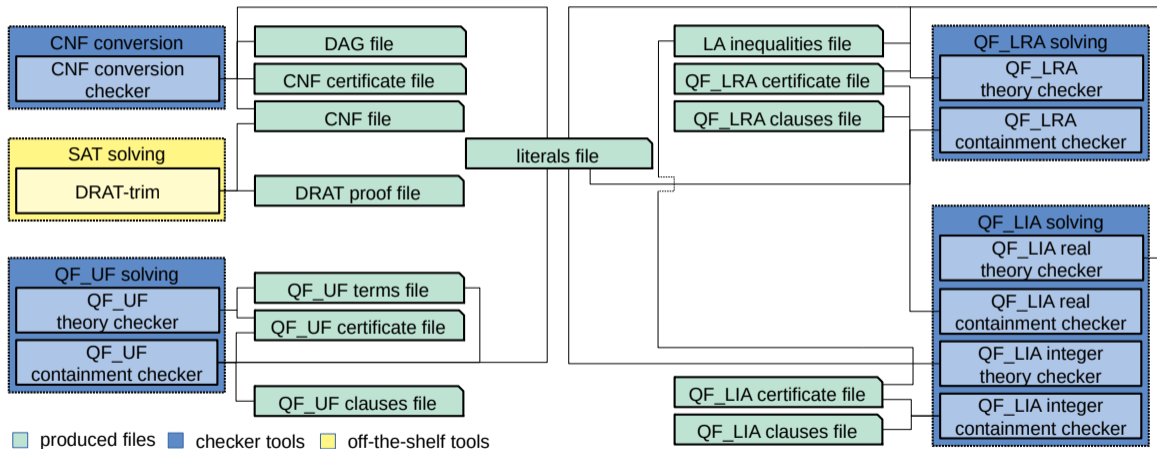
Witness

$$x4 < 0 \wedge x4 \geq 2$$

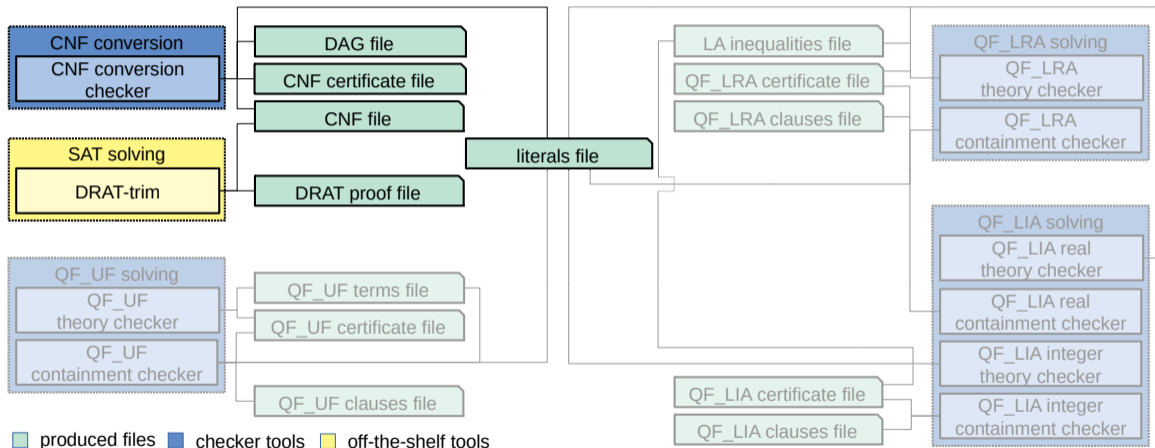
Validation

(1x)	-x4	>	0
(1x)	x4	≥	2
<hr/>			
	0	>	2

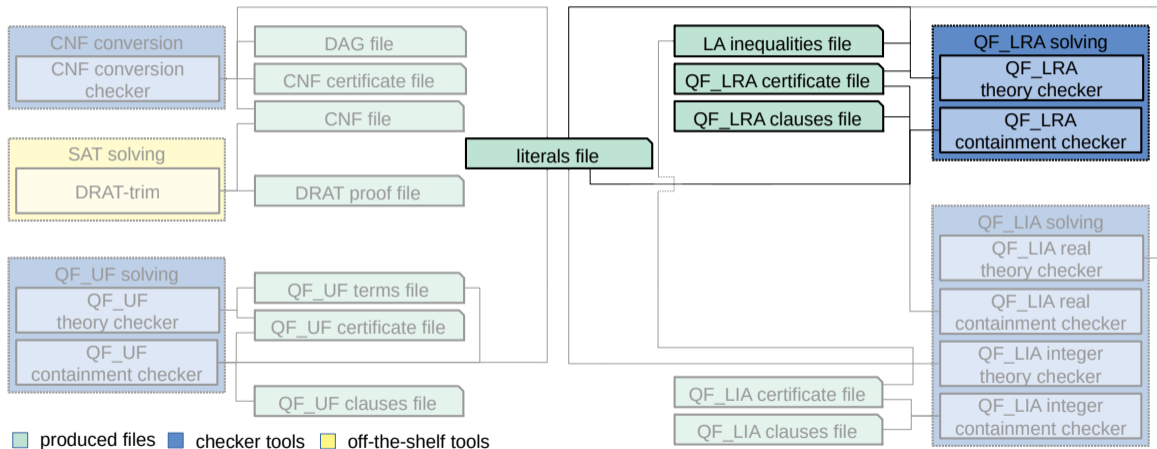
Checking Architecture



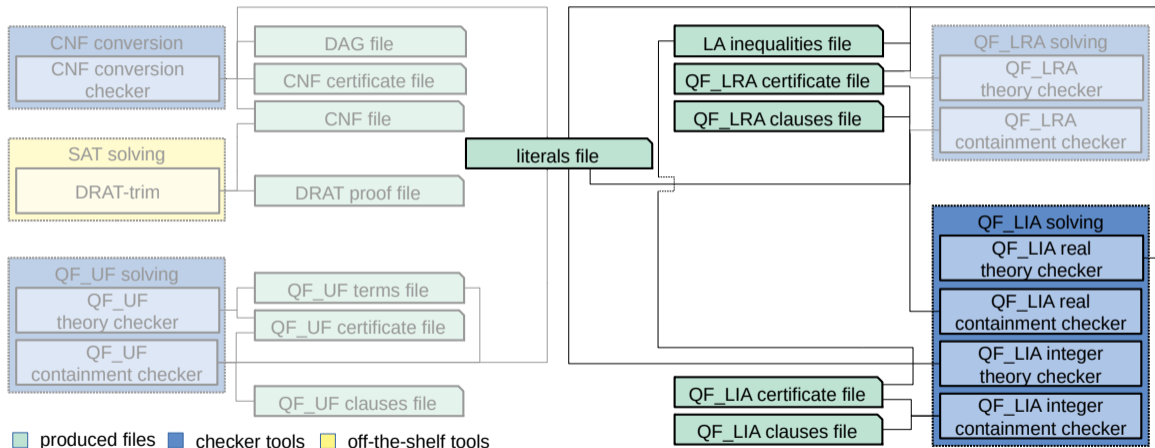
Checking Architecture



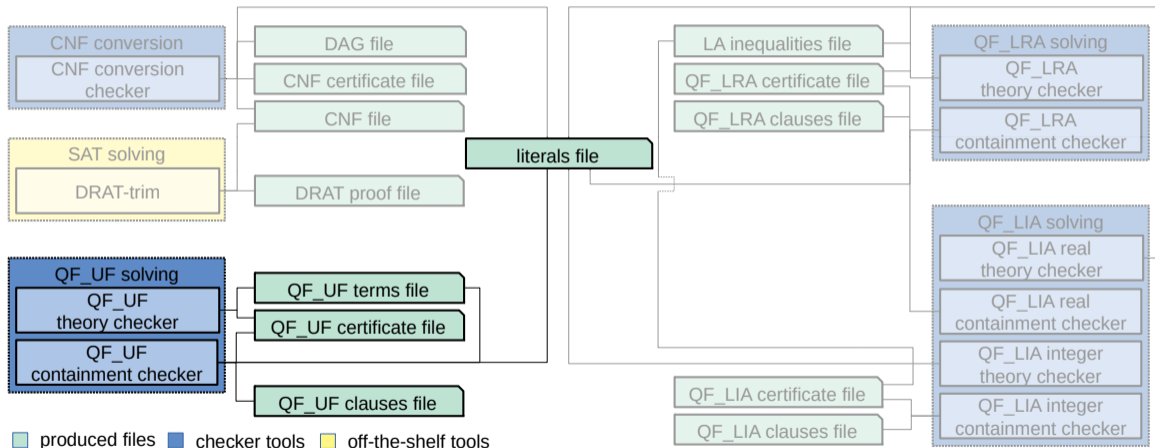
Checking Architecture



Checking Architecture



Checking Architecture



Our tools

- Witness production implemented in the OpenSMT solver
- New Theory-Specific Witness Checker (TSWC)

Our tools

- Witness production implemented in the OpenSMT solver
- New Theory-Specific Witness Checker (TSWC)

Evaluation setup

- All non-incremental SMT-LIB instances for the QF_LRA, QF_LIA, and QF_UF theories
- Witness-production comparison against CVC4, veriT, Z3
- Witness-checking comparison against CVC4's LFSC checker

Witness Production Results

		UNSAT retention (%)	Avg. runtime (s)	Avg. witness size (MB)
QF_LRA (1648 instances)	OPENSMT	99.5%	3.91/4.31	3.9
	CVC4	97.7%	5.73/6.61	13.6
	VERIT	96.7%	4.12/5.22	69.4
	Z3	97.8%	5.17/5.30	3.5
QF_LIA (6947 instances)	OPENSMT	75%	10.07/9.99	13.0
	CVC4	66.3%	6.30/ 1.30	18.8
	VERIT	95.1%	1.82 /4.05	165.9
	Z3	104.5%	5.17/6.62	24.8
QF_UF (7457 instances)	OPENSMT	99.6%	0.95/1.14	6.7
	CVC4	97.6%	0.39/1.83	6.8
	VERIT	96.0%	0.10/0.79	20.1
	Z3	97.5%	0.22/1.26	12.6

Witness Production Results

		UNSAT retention (%)	Avg. runtime (s)	Avg. witness size (MB)
QF_LRA (1648 instances)	OPENSMT	99.5%	3.91/4.31	3.9
	CVC4	97.7%	5.73/6.61	13.6
	VERIT	96.7%	4.12/5.22	69.4
	Z3	97.8%	5.17/5.30	3.5
QF_LIA (6947 instances)	OPENSMT	75%	10.07/9.99	13.0
	CVC4	66.3%	6.30/ 1.30	18.8
	VERIT	95.1%	1.82 /4.05	165.9
	Z3	104.5%	5.17/6.62	24.8
QF_UF (7457 instances)	OPENSMT	99.6%	0.95/1.14	6.7
	CVC4	97.6%	0.39/1.83	6.8
	VERIT	96.0%	0.10/0.79	20.1
	Z3	97.5%	0.22/1.26	12.6

Witness Production Results

		UNSAT retention (%)	Avg. runtime (s)	Avg. witness size (MB)
QF_LRA (1648 instances)	OPENSMT	99.5%	3.91/4.31	3.9
	CVC4	97.7%	5.73/6.61	13.6
	VERIT	96.7%	4.12/5.22	69.4
	Z3	97.8%	5.17/5.30	3.5
QF_LIA (6947 instances)	OPENSMT	75%	10.07/9.99	13.0
	CVC4	66.3%	6.30/ 1.30	18.8
	VERIT	95.1%	1.82 /4.05	165.9
	Z3	104.5%	5.17/6.62	24.8
QF_UF (7457 instances)	OPENSMT	99.6%	0.95/1.14	6.7
	CVC4	97.6%	0.39/1.83	6.8
	VERIT	96.0%	0.10/0.79	20.1
	Z3	97.5%	0.22/1.26	12.6

Witness Production Results

		UNSAT retention (%)	Avg. runtime (s)	Avg. witness size (MB)
QF_LRA (1648 instances)	OPENSMT	99.5%	3.91/4.31	3.9
	CVC4	97.7%	5.73/6.61	13.6
	VERIT	96.7%	4.12/5.22	69.4
	Z3	97.8%	5.17/5.30	3.5
QF_LIA (6947 instances)	OPENSMT	75%	10.07/9.99	13.0
	CVC4	66.3%	6.30/ 1.30	18.8
	VERIT	95.1%	1.82 /4.05	165.9
	Z3	104.5%	5.17/6.62	24.8
QF_UF (7457 instances)	OPENSMT	99.6%	0.95/1.14	6.7
	CVC4	97.6%	0.39/1.83	6.8
	VERIT	96.0%	0.10/0.79	20.1
	Z3	97.5%	0.22/1.26	12.6

Witness Checking Results

		Verified	Timeout	Error	Avg. runtime (s)
QF_LRA (567 instances)	OPENSMT + TSWC	564	3	0	3.15
	CVC4 + LFSC	471	8	88	2.85
QF_LIA (913 instances)	OPENSMT + TSWC	903	10	0	1.37
	CVC4 + LFSC	128	0	785	0.12
QF_UF (4218 instances)	OPENSMT + TSWC	4217	1	0	1.44
	CVC4 + LFSC	4157	50	11	4.01

Witness Checking Results

		Verified	Timeout	Error	Avg. runtime (s)
QF_LRA (567 instances)	OPENSMT + TSWC	564	3	0	3.15
	CVC4 + LFSC	471	8	88	2.85
QF_LIA (913 instances)	OPENSMT + TSWC	903	10	0	1.37
	CVC4 + LFSC	128	0	785	0.12
QF_UF (4218 instances)	OPENSMT + TSWC	4217	1	0	1.44
	CVC4 + LFSC	4157	50	11	4.01

Witness Checking Results

		Verified	Timeout	Error	Avg. runtime (s)
QF_LRA (567 instances)	OPENSMT + TSWC	564	3	0	3.15
	CVC4 + LFSC	471	8	88	2.85
QF_LIA (913 instances)	OPENSMT + TSWC	903	10	0	1.37
	CVC4 + LFSC	128	0	785	0.12
QF_UF (4218 instances)	OPENSMT + TSWC	4217	1	0	1.44
	CVC4 + LFSC	4157	50	11	4.01

Witness Checking Results

		Verified	Timeout	Error	Avg. runtime (s)
QF_LRA (567 instances)	OPENSMT + TSWC	564	3	0	3.15
	CVC4 + LFSC	471	8	88	2.85
QF_LIA (913 instances)	OPENSMT + TSWC	903	10	0	1.37
	CVC4 + LFSC	128	0	785	0.12
QF_UF (4218 instances)	OPENSMT + TSWC	4217	1	0	1.44
	CVC4 + LFSC	4157	50	11	4.01

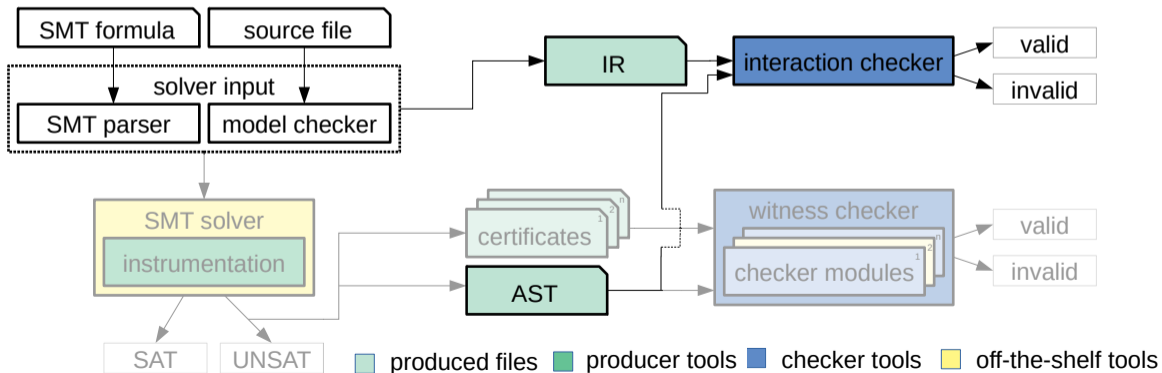
Witness Checking Results

		Verified	Timeout	Error	Avg. runtime (s)
QF_LRA (567 instances)	OPENSMT + TSWC	564	3	0	3.15
	CVC4 + LFSC	471	8	88	2.85
QF_LIA (913 instances)	OPENSMT + TSWC	903	10	0	1.37
	CVC4 + LFSC	128	0	785	0.12
QF_UF (4218 instances)	OPENSMT + TSWC	4217	1	0	1.44
	CVC4 + LFSC	4157	50	11	4.01

- Witnesses for additional theories and theory combinations
- Witnesses for parallel SMT solving
- Certify interaction between solvers and verification tools

Future Work

- Witnesses for additional theories and theory combinations
- Witnesses for parallel SMT solving
- Certify interaction between solvers and verification tools



- Lightweight witnesses that validate unsatisfiable results of SMT solvers
- Implementation of production and checking for three SMT theories
- Evaluation results indicate the compactness and easy checking of our format

Summary

- Lightweight witnesses that validate unsatisfiable results of SMT solvers
- Implementation of production and checking for three SMT theories
- Evaluation results indicate the compactness and easy checking of our format

- Details about this work can be found on our DAC'21 paper
- Check our group's website for future updates

`verify.inf.usi.ch/research/cevt`

