

# SMT-COMP 2023

## 18th International Satisfiability Modulo Theory Competition

Martin Bromberger    Jochen Hoenicke    François Bobot

CEA List, France

MPI für Informatik, Germany

Certora, Israel

July 5, 2023

# SMT-COMP

Annual competition for [SMT solvers](#)  
on (a selection of) benchmarks from [SMT-LIB](#)

Goals:

- spur development of SMT solver implementations
- promote SMT solvers and their usage
- support the SMT-LIB project
  - to promote and develop the SMT-LIB format
    - model validation
    - proof checking
  - to collect relevant benchmarks
- engage and include new members

## History

- 2005** first competition
- 2013** evaluation instead of competition
- 2014** since then hosted by [StarExec](#)

# SMT Solvers and SMT-LIB

## SMT Solver

- checks formulas in **SMT-LIB** format for **satisfiability modulo theories**

## SMT-LIB is

- ① a **language** in which benchmarks are written
- ② a community effort to **collect benchmarks**

### Non-incremental

434 212 instances (+42849)  
with 1 query each  
in 83 logics (+2).

### Incremental

43 287 instances (+2)  
with 34 036 491 queries (+37 556)  
in 38 logics.

# SMT Solvers and SMT-LIB

## SMT Solver

- checks formulas in **SMT-LIB** format for **satisfiability modulo theories**

## SMT-LIB is

- ① a **language** in which benchmarks are written
- ② a community effort to **collect benchmarks**

### Non-incremental

434 212 instances (+42849)  
with 1 query each  
in 83 logics (+2).

### Incremental

43 287 instances (+2)  
with 34 036 491 queries (+37 556)  
in 38 logics.

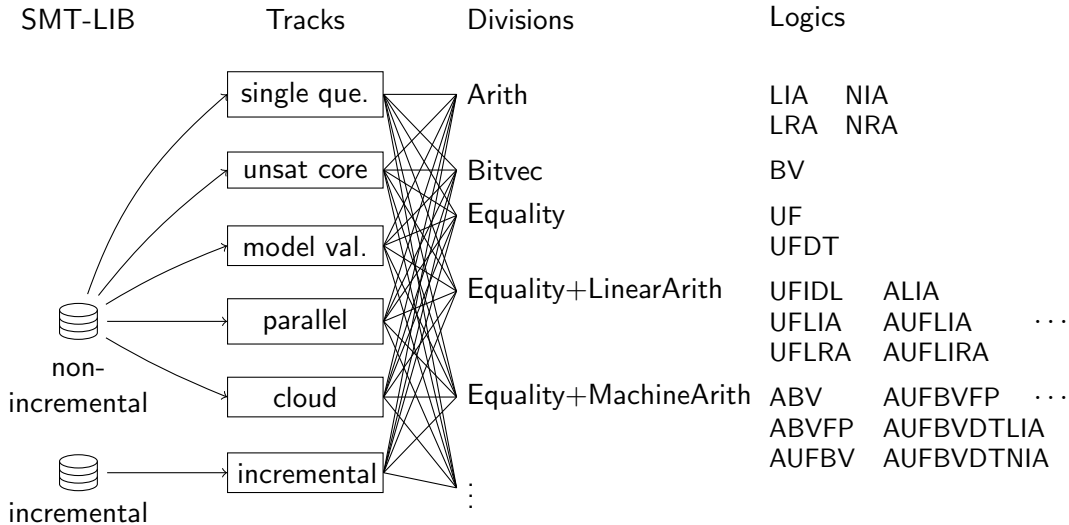
### Selected Non-incremental

227 940 instances

### Selected Incremental

22 302 instances

# Competition Overview



# SMT-COMP Tracks (traditional)

## Single Query (SQ) Track

- Determine satisfiability of one problem
- Solver answers sat/unsat/unknown

## Unsat Core Track

- Find small unsatisfiable subset of input.
- Solver answers unsat + list of formulas.

## Model Validation Track

- Find a model for a satisfiable problem.
- Solver answers sat + value for each non-logical symbol.

## Incremental Track

- Solve many small problems interactively.
- Solver acks commands and answers sat/unsat for each check.

# SMT-COMP Tracks (experimental)

## Model Validation

- Division with quantifier-free floating-point logics
- Model validation with Dolmen (thanks to Gillaume Bury and François Bobot)

## Cloud and Parallel Track (sponsored by AWS, led by Mike Whalen)

- Solve a large problem over the cloud (or a big computer)
  - 100 machines, 1600 cores, 6400 GB of memory (cloud)
  - 64 cores, 256 GB of memory (parallel)
- Solver answers sat/unsat/unknown

# SMT-COMP Tracks (experimental)

## Model Validation

- Division with quantifier-free floating-point logics
- Model validation with Dolmen (thanks to Guillaume Bury and François Bobot)

## Cloud and Parallel Track (sponsored by AWS, led by Mike Whalen)

- Solve a large problem over the cloud (or a big computer)
  - 100 machines, 1600 cores, 6400 GB of memory (cloud)
  - 64 cores, 256 GB of memory (parallel)
- Solver answers sat/unsat/unknown

## Proof Exhibition Track

- Solver submitted together with a checker for unsatisfiability proofs
- No predefined format or checker
- No ranking
- Qualitative assessment



# SMT-COMP Tracks (experimental)

## Model Validation

- Division with quantifier-free floating-point logics
- Model validation with Dolmen (thanks to Gillaume Bury and François Bobot)

## Cloud and Parallel Track (sponsored by AWS, led by Mike Whalen)

- Solve a large problem over the cloud (or a big computer)
  - 100 machines, 1600 cores, 6400 GB of memory (cloud)
  - 64 cores, 256 GB of memory (parallel)
- Solver answers sat/unsat/unknown

## Proof Exhibition Track

- Solver submitted together with a checker for unsatisfiability proofs
- No predefined format or checker
- No ranking
- Qualitative assessment

As last year the sat/unsat results from sound solvers in SQ were used to include benchmarks on the MV, UC and PE tracks.

## Tracks, Solvers, Divisions, and Benchmarks

Teams: 25 (+4)

Track	Solvers	Divisions	Benchmarks
Single Query	22(=)	19(=)	113 139
Incremental	7(-1)	17(=)	22 301
Unsat Core	6(=)	16(-1)	72 958
Model Validation	11(+3)	13(+ 6 exp.)	61 083
Proof Exhibition	4(=)	19 exp.	59 114
Parallel	3(-1)	4 exp.	400
Cloud	2(-2)	4 exp.	400

Number in parenthesis shows changes from 2022

## Participants

SMT-COMP 2022 participants rely on multiple reasoning frameworks:

- CDCL(T), Saturation, MCSAT, CP
- automata
- finite domain
- local search
- besides wrappers extending the scope of existing solvers

Seven new solvers participated:

- iProver (Korovin et al.)
- SMT-RAT-MCSAT (Jasper Nalbach et al.)
- UltimateIntBlastingWrapper+SMTInterpol (Max Barth et al.)
- Yaga (Hanák et al.)
- Z3-alpha (Lu et al.)
- Z3-Noodler (Havlena et al.)
- Z3-Owl (Ma et al.)

# Solver Presentation



**Tracks/Divisions:**  $\sim(QF\_)?(A)?(UF)?(BV|FP|FPLRA)+\$$   
in tracks **Single Query**, **Incremental**, **Unsat Core**, and **Model Validation**

## Highlights

- » **New SMT solver** for theories **A**, **BV**, **FP**, **UF** + quantifiers
- » **Rewrite from scratch**
- » previous versions were extended fork from Boolector
- » system description at **CAV 2023**

<https://bitwuzla.github.io>

## COLIBRI(2023) Bruno Marre et al

- CP solver:
  - No SAT solver
  - Combination done by explicitly building a model
- main theories: **FP**, **LIA**  $\leftrightarrow$  **BV**
  
- Some fixes
- Add some reasoning for transcendental functions

# cvc5 at the SMT Competition 2023

L. Aniva H. Barbosa C. Barrett M. Brain V. Camillo G. Kremer H. Lachnitt A. Mohamed  
M. Mohamed A. Niemetz A. Nötzli A. Ozdemir M. Preiner A. Reynolds Y. Sheng C. Tinelli A. Wilson Y. Zohar

## cvc5 1.0.5

- Support for all standardized SMT-LIB theories
- User-friendly API

## Features/Improvements

- New handwritten parser and lexer
- Bit-vector solver, integrating efficient SAT solvers, e.g., CaDiCaL, with CDCL( $\mathcal{T}$ )
- Syntax-guided and model-based quantifier instantiation
- Minor improvements to the lemma schemas used for the theory of strings
- Better performance for proof generation

## Configurations

cvc5 entered **all divisions** in **all tracks** (non-experimental).

- Single query track: Sequential portfolio
- Unsat-core track: Based on new proof module and assumptions in the SAT solver

Follow the development: <https://cvc5.github.io/>

## cvc5-NRA-LS

video: `slides-cvc5-nra-ls.mp4`



# iProver

video: slides-iProver.mp4

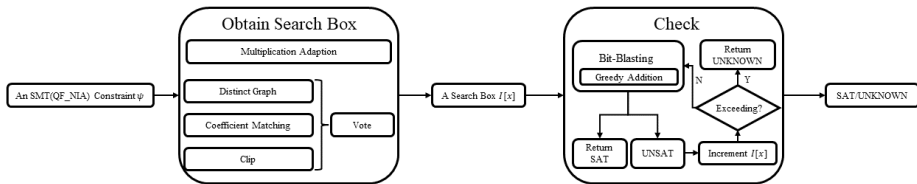
## ismt, Yices-ismt

video: slides-ismt.mp4

Fuqi Jia, Rui Han, Minghao Liu, Cunjing Ge, Pei Huang, Feifei Ma, Jian Zhang.

- ISMT is a pure bit-blasting based SMT solver, participating QF\_NIA model validation track;
- YICES-ISMT combines YICES2 for unsat reasoning, participating in QF\_NIA single query track.

#### □ ISMT:



#### □ YICES-ISMT: $YICES(\phi) \rightarrow ISMT(\phi) \rightarrow YICES(\phi \wedge \psi)$ .

- $\psi$  rule out failed space.

◆ Link: <https://github.com/MRVAPOR/BLAN>

#### Dependencies

- Yices 2.6.2
- Libpoly v0.1, 11
- CaDiCal 1.5.2

- **Interpolating** CDCL( $T$ ) SMT solver
  - Developed at University of Lugano, Switzerland
  - <https://github.com/usi-verification-and-security/opensmt>
- Support for *linear arithmetic*, uninterpreted functions, arrays
- Alternative lookahead core
- Used in Horn solver GOLEM
- New in 2023 edition 😊
  - Support for incrementality and interpolation in lookahead core
  - Theory combination with arrays
- On hold in 2023 edition 😞
  - Proof track
  - Parallel and cloud track



# Ostrich

video: slides-Ostrich.mkv

## Interpolating SMT solver

- based on CDCL(T)
- for Arrays, Uninterpreted Functions, Linear Integer and Real Arithmetic
  - plus `div` and `mod` with constants, and Data Types
- supports quantifiers
- produces models, proofs, and unsat cores
- computes sequence and tree interpolants

## SMTInterpol at SMT-COMP 2023

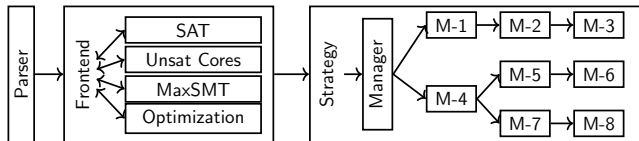
- **models** for **data type** logics
- optimize size of **proofs**
- **proof check** in single query/unsat core

Try it in your browser:



<https://tinyurl.com/smtinterpol>

# SMT-RAT 23.05



# UltimateEliminator+MathSAT

video: `slides-UltimateEliminator.mkv`





## Vampire 4.8

*Reger, Suda, Voronkov, Kovács, Bhayat, Gleiss,  
Hajdu, Hozzová, Rath, Rawson, Schoisswohl*

<https://vprover.github.io/>

- ▶ General Approach: proof search using the **Superposition Calculus** (and finite model finding in UF)
- ▶ SMT Logics: A, DT, LIA, LRA, NIA, NRA, UF (all **with quantifiers**)
- ▶ Uses a **portfolio of strategies**
  - Parallel track: parallelize strategies
  - Cloud track: randomize problem and strategies per node
- ▶ Use Z3 for ground reasoning (AVATAR)
- ▶ New for arithmetic: **ALASCA Calculus** and new simplification rules

# Yaga—MCSat-based SMT solver

- Developed at Charles University, Prague, Czech Republic
  - Drahomír Hanák, Martin Blichá, Jan Kofroň
- Student project, from December 2022
- Support for QF\_LRA
- Support for models
- Future work
  - More theories
  - Model-based interpolation
- <https://github.com/d3sformal/yaga>



Bruno Dutertre, Aman Goel, Stéphane Graham-Lengrand, Ahmed Irfan, Dejan Jovanović, Ian A. Mason

<https://yices.csl.sri.com/>

**Two solvers:** CDCL(T) & MCSAT

**Support:**

- **Quantifier-free:** non-linear arithmetic (MCSAT only), linear arithmetic, bitvectors, uninterpreted functions, and arrays.
- **With quantifiers:** uninterpreted functions only, via E-graph matching and model-based instantiation.

**Bitvectors:** Yices 2/CDCL(T) uses bitblasting.

For QF\_BV, it can optionally use third-party backend SAT solvers:

CaDiCaL, CryptoMiniSat, and Kissat (the SMT-comp version uses Kissat for single-query and model validation tracks).

**Functionalities:** incremental and push/pop modes, unsat cores, model minimization and implicants, Model-Based Over-approximations, Model-Based Under-approximations, Craig Interpolants.

**2023** 

- arrays in MCSAT
- new variable decision and clause scoring heuristics in MCSAT.

# YicesQS, an extension of Yices2 for quantifiers (SMT-comp 2023)

Stéphane Graham-Lengrand

<https://github.com/disteph/yicesQS>

Same solver as in the 2022 SMT-comp.<sup>1</sup>

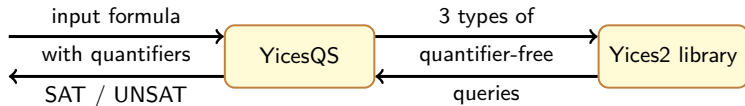
YicesQS implements a variant of the QSMA algorithm presented at CADE'2023:

<https://www.csl.sri.com/users/sgl/Work/Reports/CADE2023.pdf>

Lazy approach to quantifier elimination based on Model-Based Over-approximations (MBO) and Model-Based Under-approximations (MBU). YicesQS is written in OCaml, using Yices2 as a library via its OCaml bindings.

2023: YicesQS entered *NRA*, *NIA*, *LRA*, *LIA* and *BV* (single-track), & generally targets complete theories with procedures for answering 3 types of quantifier-free queries:

- *Satisfiability modulo assignment / modulo a model* (here relying on MCSAT)
- *MBU* (here using invertibility conditions for BV, CAD projections for arithmetic)
- *MBO* (here again relying on MCSAT, incl. CAD for arithmetic)



<sup>1</sup> YicesQS-2023 (Starexec solver 45053, too late for 2023 SMT-comp) is way better at BV, solves 805/970 instances out of the 2022 single-track BV selection.



Z3-Z3++

<https://youtu.be/fBB0Wxxf9vA>

## Other participants

- Q3B
- Q3B-pBDD
- STP
- UltimateIntBlastingWrapper+SMTInterpol
- z3-alpha
- Z3-Noodler
- Z3-Owl

# Non-Competitive Solvers

Submitted by organisers

- Best solvers, per division, from previous years (27 Solvers)

Submitted by participants

- Fixed solvers (OSTRICH, Z3-Owl, Bitwuzla, Yices2, Z3-Noodler, iProver)

# Scoring

Computing scores:

- **Single Query/Parallel/Cloud**: number of solved **instances**
- **Incremental**: number of solved **queries**
- **Unsat Core**: number of top-level assertions **removed**
- **Model Validation**: number of solved instances with correct **models**

Error scores:

- **All Tracks**: given for sat reply for unsat instance, or vice versa
- **Unsat Core**: given if returned core is satisfiable.
- **Model Validation**: given if given model evaluates formula to **false**

Error scores are **draconian**.



# Score and Ranking

In each track we collect different scores:

- **Sequential score** (SQ, UC, MV): all time limits apply to cpu time
- **Parallel score** (all): all time limits apply to wallclock time
- **SAT score** (SQ): parallel score for **satisfiable** instances
- **UNSAT score** (SQ): parallel score for **unsatisfiable** instances
- **24s** (SQ): parallel score with time limit of **24s**

Division ranking (for each score)

- For each division, one winner is declared

Two competition-wide rankings (for each score)

- **Biggest lead**: division winner with most score difference to second place
- **Largest contribution**: improvement each solver provided to a virtual best solver

## Results

T	satisfiable	$\perp$	unsatisfiable
;	sequential		parallel
24	less than 24s	inc	incremental
uc	unsat core	mv	model validation

Experimental track are added in the slides but are not present in the certificates

# Bitwuzla

## *Overall Winner*

BIGGEST LEAD(inc)

## *Winner of the Divisions*

BITVEC( $\top$ ,24), EQUALITY+MACHINEARITH(inc), FPARITH( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,uc),  
QF\_ADT+BITVEC(mv), QF\_BITVEC(inc),  
QF\_EQUALITY+BITVEC( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,mv), QF\_FPARITH( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,mv)

## *Winner of the Logics (where it did not win the corresponding division)*

ABVFP( $\top$ ,24), ABVFPLRA(24), AUFBV( $;$ , $\parallel$ , $\top$ , $\perp$ ,24),  
AUFBVFP( $;$ , $\parallel$ , $\top$ , $\perp$ ,24), QF\_ABVFP(uc), QF\_ABVFPLRA(uc),  
QF\_BVFP(uc), QF\_BVFPLRA(uc), UFBV( $;$ , $\parallel$ , $\top$ , $\perp$ ,24), UFBVFP( $;$ , $\parallel$ , $\perp$ )

# COLIBRI

*Winner of the Logics*

QF\_ABVFPLRA( $;$ ,  $\parallel$ ,  $\perp$ , 24), QF\_FP( $\perp$ , 24), QF\_FPLRA( $\perp$ , 24)

## cvc5

### *Overall Winner*

BIGGEST LEAD( $;$ , $\parallel$ , $\top$ , $\perp$ ,uc), LARGEST CONTRIBUTION( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,uc)

### *Winner of the Divisions*

ARITH( $;$ , $\parallel$ , $\perp$ ,inc,uc), BITVEC( $;$ , $\parallel$ , $\perp$ ,inc,uc), EQUALITY( $;$ , $\parallel$ , $\top$ , $\perp$ ,inc,uc,cloud),  
EQUALITY+LINEARARITH( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,uc,cloud),  
EQUALITY+MACHINEARITH( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,uc),  
EQUALITY+NONLINEARARITH( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,inc,uc), QF\_DATATYPES( $;$ , $\parallel$ , $\top$ , $\perp$ ,uc<sup>2</sup>),  
QF\_EQUALITY(inc), QF\_EQUALITY+BITVEC+ARITH(inc),  
QF\_EQUALITY+NONLINEARARITH( $;$ , $\parallel$ , $\perp$ ,24,inc,uc,mv), QF\_FPARITH(uc),  
QF\_LINEARINTARITH( $\perp$ ), QF\_LINEARREALARITH( $\perp$ ),  
QF\_NONLINEARREALARITH( $\perp$ ,24), QF\_STRINGS( $;$ , $\parallel$ , $\top$ , $\perp$ )

### *Winner of the Logics (where it did not win the corresponding division)*

LIA( $\top$ ,24), NIA( $\top$ ,24), QF\_ALIA(inc), QF\_AUFLIA(uc,mv), QF\_AX(mv),  
QF\_BVFPLRA(inc), QF\_FP( $\top$ ,mv), QF\_IDL(uc), QF\_NIRA( $;$ , $\parallel$ , $\perp$ ),  
QF\_UFBVDT( $;$ , $\parallel$ , $\top$ , $\perp$ ,24), QF\_UFDTLIRA( $;$ , $\parallel$ , $\top$ , $\perp$ ,24,uc), QF\_UFLRA(uc),

## iProver

*Winner of the Divisions*

EQUALITY(parallel), EQUALITY+LINEARARITH(parallel)

*Winner of the Logic (where it did not win the corresponding division)*

ANIA( $\perp$ )

# OpenSMT

*Winner of the Divisions*

QF\_EQUALITY+LINEARARITH( $\perp$ ), QF\_LINEARREALARITH( $\top$ ,inc,mv)

*Winner of the Logics (where it did not win the corresponding division)*

QF\_LIA( $;$ , $\parallel$ ), QF\_LRA( $;$ , $\parallel$ , $\perp$ ,24), QF\_UFIDL( $;$ , $\parallel$ , $\top$ ,mv)

# OSTRICH

*Winner of the Logic*

QF-S(;, ||, ⊥)



# SMTInterpol

## *Overall Winner*

BIGGEST LEAD(24)

## *Winner of the Divisions*

QF\_ADT+LINEARITH(mv), QF\_DATATYPES(24,uc<sup>l</sup>,mv),

QF\_EQUALITY+LINEARARITH(;,||,⊤,inc,mv), QF\_NONLINEARINTARITH(inc)

## *Winner of the Logics (where it did not win the corresponding division)*

ALIA(⊤,uc), AUFDTLIA(uc), QF\_ALIA(uc), QF\_ANIA(;,||,⊤,⊥,24,inc,uc),

QF\_AUFNIA(;,||,⊤,⊥,24,uc), QF\_LIA(⊥), QF\_UF(uc<sup>l</sup>), QF\_UFDT(uc<sup>i</sup>),

QF\_UFDTLIA(⊥,24,uc), UFIDL(⊤), UFLIA(⊤)

# STP

*Winner of the Division*

QF\_BITVEC(;;,T,24,mv)

# Vampire

## *Overall Winner*

BIGGEST LEAD(cloud,parallel)

## *Winner of the Divisions*

ARITH(cloud,parallel), EQUALITY(24),  
EQUALITY+NONLINEARARITH(cloud,parallel)

## *Winner of the Logics (where it did not win the corresponding division)*

ALIA(;;,⊥,24), AUFLIA(⊥,24,cloud,parallel), AUFLIRA(24),  
AUFNIRA(⊥,⊥,24), UF(;;,⊥), UFDTLIA(24,uc), UFDTNIA(;;,⊥,24,uc),  
UFLIA(cloud,parallel)

## Yices2

### *Winner of the Divisions*

QF\_BITVEC(uc), QF\_EQUALITY(;;,⊥,24,uc,mv),  
QF\_EQUALITY+BITVEC(uc), QF\_EQUALITY+LINEARARITH(24,uc),  
QF\_EQUALITY+NONLINEARARITH(⊥), QF\_LINEARINTARITH(24,inc,uc),  
QF\_LINEARREALARITH(;;,⊥,24,uc)

### *Winner of the Logics (where it did not win the corresponding division)*

QF\_ALIA(⊥), QF\_AUFBV(24), QF\_AUFBVLIA(inc),  
QF\_AUFLIA(;;,⊥), QF\_LIRA(;;,⊥,mv), QF\_RDL(⊥,⊥,mv),  
QF\_UFBVLIA(inc), QF\_UFLIA(⊥,inc), QF\_UFLRA(;;,⊥,inc,mv),  
QF\_UFNRA(;;,⊥,24,inc,mv)

# YicesQS

*Winner of the Division*

ARITH( $\top$ ,24)

*Winner of the Logics (where it did not win the corresponding division)*

LRA( $;$ , $\parallel$ , $\perp$ ), NRA( $;$ , $\parallel$ , $\perp$ )

## Z3-Z3++

### *Overall Winner*

BIGGEST LEAD(mv), LARGEST CONTRIBUTION(mv)

### *Winner of the Divisions*

QF\_LINEARINTARITH(;;,⊥,mv), QF\_NONLINEARINTARITH(;;,⊥,⊥,24,mv),  
QF\_NONLINEARREALARITH(;;,⊥,mv)

### *Winner of the Logic (where it did not win the corresponding division)*

QF\_IDL(⊥)

## z3-alpha

*Winner of the Division*

QF\_STRINGS(24)

*Winner of the Logics (where it did not win the corresponding division)*

QF\_S( $\top$ ), QF\_SNIA( $;$ ,  $\parallel$ ,  $\top$ )

## Checking Disagreements

- 50411 benchmarks of 227 938 have no status
- 300 benchmarks with disagreements (ABV, BV, LIA, NIA, QF\_ABV, QF\_AUFBV, QF\_BVFP, QF\_FP, QF\_NRA, QF\_SLIA, UF, UFDT, UFDTLIRA)
- We manually resolved the disagreements: authors confirmed solver unsoundness
- We had 10 solvers with soundness issues:
  - Bitwuzla (uc printing error)
  - iProver (& Fixed)
  - OSTRICH
  - Q3B
  - UltimateEliminator+MathSAT
  - UltimateIntBlastingWrapper+SMTInterpol
  - Vampire
  - Yices2 (& Fixed)
  - Z3-Noodler (& Fixed)
  - Z3-Owl (& Fixed)



## Plans for SMT-COMP 2023

- "Revelations are found in clouds" — Serge King

- MV: Algebraic number  $\implies$  non experimental

```
(root-of-with-ordering (coeffs p_0 p_1 ... p_n) i)
```

```
(root-of-with-interval (coeffs p_0 p_1 ... p_n) min max)
```

```
(root-of-with-enclosure (coeffs p_0 p_1 ... p_n) min max)
```

Thanks for the solvers that implemented one of them!

- MV: partial function more discussion needed
- MV: array theory more discussion needed

## Plans for SMT-COMP 2023?

- Proof *validation* track
- Hopefully proof exhibition this year will help
- We have to analyze the data still
  - Job only finished last week (took 18 days to run)
  - 830gb

## SMT-COMP organizing committee

Three people organize the SMT-COMP. In 2023:

- Martin Bromberger
- Jochen Hoenicke
- François Bobot

Jochen have been organizer for four-years! He can rest happy.

We need a successor for next year's competition. Contact us if you would like to volunteer!

# Acknowledgements

- Andrea Micheli: pysmt
- Guillaume Bury, FB.: Model Validator
- Clark Barrett, Pascal Fontaine, Aina Niemetz, Mathias Preiner, Hans-Jörg Schurr: SMT-LIB benchmarks
- Aaron Stump: StarExec support
- Mike Whalen and team: Cloud/Parallel Track



## Benchmark contributors

In 2023 [new benchmarks](#) were contributed by:

- Alex Coffin
- Alex Ozdemir
- Ali Uncu, James Davenport and Matthew England
- Bohan Li
- Elizabeth Polgreen
- Fuqi Jia
- Johann-Tobias Aaron and Raphael Schäg
- Matthew England and Miguel Del Rio Almajano
- Nicolas Amat
- Yannick Moy
- Yoni Zohar

Thanks

to all participants

Thanks

to all participants

and to you for listening