

Reliable and Reproducible Competition Results with BenchExec and Witnesses (Report on SV-COMP 2016)

Dirk Beyer

University of Passau, Germany

Abstract. The 5th Competition on Software Verification (SV-COMP 2016) continues the tradition of a thorough comparative evaluation of fully-automatic software verifiers. This report presents the results of the competition and includes a special section that describes how SV-COMP ensures that the experiments are reliably executed, precisely measured, and organized such that the results can be reproduced later. SV-COMP uses BENCHEXEC for controlling and measuring the verification runs, and requires violation witnesses in an exchangeable format, whenever a verifier reports that a property is violated. Each witness was validated by two independent and publicly-available witness validators. The tables report the state of the art in software verification in terms of effectiveness and efficiency. The competition used 6 661 verification tasks that each consisted of a C program and a property (reachability, memory safety, termination). SV-COMP 2016 had 35 participating verification systems (22 in 2015) from 16 countries.

1 Introduction

The annual Competition on Software Verification (SV-COMP)¹ is a continuous effort by the software-verification community. The effort consists of the following two parts: (1) The SV-COMP community defines and collects verification tasks that the researchers and developers of software verifiers find interesting and challenging; these verification problems should be used to evaluate the effectivity (soundness and completeness) and efficiency (performance) of modern verification tools. (2) The organizer of SV-COMP performs a systematic comparative evaluation of the relevant state-of-the-art tool implementations for automatic software verification with respect to effectiveness and efficiency; part of this is to define and explore standards for a reliable and reproducible execution of such a competition. This paper describes the rules, definitions, results, and other interesting facts about the execution of the competition experiments, in particular how to make the experiments reproducible. The main objectives that the community aims at by running yearly competitions are the following (taken from [5]):

1. provide an overview of the state of the art in software-verification technology and increase visibility of the most recent software verifiers,

¹ <http://sv-comp.sosy-lab.org>

2. establish a repository of software-verification tasks that is publicly available for free use as standard benchmark suite for evaluating verification software,
3. establish standards that make it possible to compare different verification tools including a property language and formats for the results, and
4. accelerate the transfer of new verification technology to industrial practice.

There is consensus that (1) and (2) are already achieved, but need continuous improvement: the community of research groups and verifiers that participate in SV-COMP is increasing, and the set of verification tasks needs even more diversity, growing, and quality assurance. The repository and the issue tracker show that there was considerable effort spent on consolidating the verification tasks, in terms of consistency and quality. Regarding (3), the simple syntax of the property language works well for SV-COMP, while it would be great to increase the supported fragment of LTL. The standard witness language as a common, exchangeable format was a big step forward in terms of standardization. The requirement in SV-COMP that bug reports are counted only if the bug is reproducible, i.e., the witness can be re-played on a different machine with a different validation tool, makes it easier to understand problems. We received positive feedback in terms of Objective (4), but we cannot evaluate this here.

Related Competitions. SV-COMP is complemented by two other competitions in the field of software verification: RERS² and VerifyThis³. While SV-COMP performs reproducible experiments in a *controlled* environment (dedicated resources, resource limits), the RERS Challenges gives more room for exploring combinations of interactive with automatic approaches without limits on the resources, and the VerifyThis Competition focuses on evaluating approaches and ideas rather than on *fully-automatic* verification. The report on SV-COMP 2014 provides a more comprehensive list of other competitions [4].

2 Procedure

The procedure for the competition organization did not change in comparison to the past SV-COMP editions [2, 3, 4, 5]. SV-COMP was again an open competition where all verification tasks were known before the submission of the participating verifiers, such that there were no surprises and developers were able to train the verifiers. In the *benchmark submission* phase, we collected and classified new verification tasks, in the *training* phase, the teams inspected verification tasks and trained their verifiers, and in the *evaluation* phase, verification runs were performed with all competition candidates and the system descriptions were reviewed by the competition jury. As in the last years, the participants received the preliminary results of their verifier per e-mail for inspection, after which the results were publicly announced.

3 Definitions, Formats, and Rules

Verification Task. The definition of verification task was not changed (taken from [4]). A verification task consists of a C program and a property. A verifica-

² <http://rers-challenge.org> ³ <http://etaps2015.verifythis.org>

tion run is a non-interactive execution of a competition candidate on a single verification task, in order to check whether the following statement is correct: “The program satisfies the property.” The result of a verification run is a triple (ANSWER, WITNESS, TIME). ANSWER is one of the following outcomes:

- TRUE: The property is satisfied (i.e., no path that violates the property exists).
- FALSE: The property is violated (i.e., there exists a path that violates the property) and a counterexample path is produced and reported as WITNESS.
- UNKNOWN: The tool cannot decide the problem, or terminates abnormally, or exhausts the computing resources time or memory (i.e., the competition candidate does not succeed in computing an answer TRUE or FALSE).

The component WITNESS [6] was this year mandatory only for FALSE answers; in the future, witnesses are also required for TRUE answers. SV-COMP was supported by the two witness validators CPACHECKER and UAUTOMIZER. TIME is measured as consumed CPU time until the verifier terminates, including the consumed CPU time of all processes that the verifier started [8].

If the wall time was larger than the CPU time, then the TIME is set to the wall time. If TIME is equal to or larger than the time limit (15 min), then the verifier is terminated and the ANSWER is set to ‘timeout’ (and interpreted as UNKNOWN).

Categories. The collection of verification tasks, which represents the current interest and abilities of tools for software verification, is arranged into categories, according to the characteristics of the programs and the properties to be verified. The assignment was proposed and implemented by the competition chair, and approved by the competition jury. For the 2016 edition of SV-COMP, a total of 10 categories were defined. The structure of categories is illustrated in Fig. 1 and described in more detail on the competition web site⁴. As a new feature of the competition, a new (meta) category *Falsification* was defined, which was meant to explore bug hunting capabilities of verifiers that are not able to construct correctness proofs. The new category consisted of all verification tasks with safety properties, and any answers TRUE were ignored. The categories, their defining category-set files, and the con-

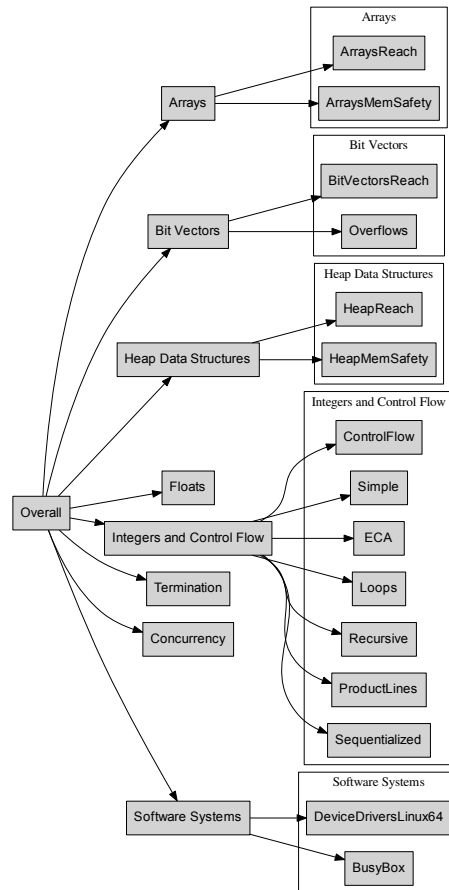


Fig. 1: Categories (generated by GRAPHVIZ)

⁴ <http://sv-comp.sosy-lab.org/2016/benchmarks.php>

Table 1: Properties used in the competition (cf. [5] for more details)

Formula	Interpretation / Syntax of property
<code>G ! call(foo())</code>	A call to function <code>foo</code> is not reachable on any finite execution. <code>CHECK(init(main()), LTL(G ! call(__VERIFIER_error())))</code>
<code>G valid-free</code>	All memory deallocations are valid (counterexample: invalid free). More precisely: There exists no finite execution of the program on which an invalid memory deallocation occurs. <code>CHECK(init(main()), LTL(G valid-free))</code>
<code>G valid-deref</code>	All pointer dereferences are valid (counterexample: invalid dereference). More precisely: There exists no finite execution of the program on which an invalid pointer dereference occurs. <code>CHECK(init(main()), LTL(G valid-deref))</code>
<code>G valid-memtrack</code>	All allocated memory is tracked, i.e., pointed to or deallocated (counterexample: memory leak). More precisely: There exists no finite execution of the program on which the program lost track of some previously allocated memory. <code>CHECK(init(main()), LTL(G valid-memtrack))</code>
<code>F end</code>	All program executions are finite and end on proposition <code>end</code> , which marks all program exits (counterexample: infinite loop). More precisely: There exists no execution of the program on which the program never terminates. <code>CHECK(init(main()), LTL(F end))</code>

Table 2: Scoring schema for SV-COMP 2016

Reported result	Points	Description
UNKNOWN	0	Failure to compute verification result
FALSE correct	+1	Violation of property in program was correctly found
FALSE incorrect	-16	Violation reported but property holds (false alarm)
TRUE correct	+2	Correct program reported to satisfy property
TRUE incorrect	-32	Incorrect program reported as correct (wrong proof)

tained programs are explained in more detail under *Verification Tasks* on the competition web site.

Properties and Their Format. For the definition of the properties and the property format we refer to the previous competition report [5]. All specifications are available as `.prp` files in the respective directories of the benchmark categories in the repository. Table 1 lists the properties and their syntax as overview.

Evaluation by Scores and Run Time. In order to reflect the steady progress towards completeness and soundness of verification tools, the scoring schema was again adjusted in order to increase the penalty for wrong results. Table 2 provides the overview. The ranking is decided based on the sum of points (normalized for meta categories) and for equal sum of points according to success run time, which is the total CPU time over all verification tasks for which the verifier reported a correct verification result. *Opt-out from Categories* and *Score Normalization for Meta Categories* was done as described previously [3] (page 597). The *Competition Jury* consists again of the chair and one member of each participating team. Team representatives of the jury are listed in Table 3.




Fig. 2: Setup: components that support reproducibility are highlighted in green

4 Reproducibility

One of the main goals of SV-COMP is to make the competition as transparent and reproducible as possible. To achieve this goal, it is necessary to control as many as possible of the variables that might influence the results. Figure 2 gives an overview over the components that contribute to the reproducible setup of SV-COMP.

BenchExec: Precise Controlling and Measurement of Resources (e). For scientifically valid experiments, we require for each verification run a reliable assignment and controlling of computing resources (cores, memory, CPU time), and a precise measurement. There are several requirements that experiments of a competition such as SV-COMP have to fulfill [8]: (i) accurate measurement and reliable enforcement of limits for CPU time and memory, (ii) reliable termination of processes (including all child processes), and (iii) correct assignment of local memory (for NUMA architectures). We use `BENCHEXEC`⁵ to perform all SV-COMP experiments, because this benchmarking framework lets us conveniently benefit from the modern resource control and measurement mechanisms that the Linux kernel offers.

Repository of Verification Tasks (a). The verification tasks are organized in a public repository⁶. The repository was moved to `GITHUB` in order to support an issue tracker and to efficiently handle contributions from the community via pull requests. The more appropriate logging of change history and issues gives credit to people that contribute. Furthermore, the continuous-integration system `TRAVIS` is used to ensure that the verification tasks are compilable by `GCC` and `CLANG`. The move to `GITHUB` also had a positive effect on the activity on the benchmark suite: more people are involved, and more fixes to verification tasks were contributed. For reproducing the results of SV-COMP, the exact versions of the verification tasks as used for SV-COMP 2016 are available via the PGP-signed tag ‘svcomp16’ in the git repository.

⁵ <https://github.com/sosy-lab/benchexec> ⁶ <https://github.com/sosy-lab/sv-benchmarks>

Benchmark Definitions (b). For executing verification runs, we need to know for each verifier, (i) which verification tasks need to be given to the verifier (derived from participation declaration) and (ii) which parameters need to be passed to the verifier (there are global parameters that are specified for all categories, and there are specific parameters such as the bit architecture and memory model). The benchmark definitions are XML files in the format that `BENCHEXEC` expects; they are collected in a specific repository for SV-COMP⁷, in which the PGP-signed tag ‘svcomp16’ points to the exact versions that were used in SV-COMP 2016.

Tool-Specific Information (c). In order to successfully execute a verifier and correctly interpret its results, a tool-info module needs to be provided to `BENCHEXEC`. First, the command-line to properly invoke the verifier (including source and property file as well as the options) is assembled from the parts specified in the benchmark definition (b). Second, the (tool-specific) information that the verifier produces needs to be interpreted and translated into the uniform SV-COMP result (`TRUE`, `FALSE(p)`, `UNKNOWN`). The tool-info modules that were used in SV-COMP 2016 are available in `BENCHEXEC` release 1.7.

Verifier Archive (d). The verifiers are provided in an archive containing a license (that permits academic use, use in SV-COMP, and reproducing the results) and all parts that are needed to execute the verifier (statically-linked executables, all components that are required in a certain version, or for which no standard Ubuntu package is available, are included). The verifiers and the above-mentioned components are provided on the systems-description page of the SV-COMP web site⁸, together with the SHA1 hashes for verification of consistency.

Violation Witnesses (f). SV-COMP counts answers `FALSE` only if a valid witness according to an exchangeable, machine-readable format is part of the result triple as `WITNESS`. This means that each verification run must be followed by a validation run that checks if the witness adheres to the exchange format and can be reproduced. The time limit for a validation run was set to 10 % of the CPU time for a verification run, i.e., the witness validation was limited to 90 s. The purpose of the tighter resource limit is to avoid delegating verification work to the validator. This ensures a high quality of assignment of scores: if a verifier claims a found bug but is not able to provide a witness, then no score is assigned. The witness format and the validation process is explained on the web page⁹. More details on witness validation is given in a related research article [6].

Correctness Witnesses (g). Although SV-COMP requires since its second edition (2013) that each result must be accompanied by a witness, this requirement was not enforced for the answer `TRUE`, mainly due to the lack of validators for correctness witnesses. This year, there was a demonstration category on validation of correctness witnesses, with the purpose to get prepared for witness validation for correctness results in the future.

⁷ <https://github.com/sosy-lab/sv-comp> ⁸ <http://sv-comp.sosy-lab.org/2016/systems.php>

⁹ <http://sv-comp.sosy-lab.org/2016/witnesses/>

5 Results and Discussion

For the fifth time, the competition experiments represent the state of the art in fully-automatic and publicly-available software-verification tools. The report shows the improvements of the last year, in terms of effectiveness (number of verification tasks that can be solved, correctness of the results, as accumulated in the score) and efficiency (resource consumption in terms of CPU time). The results that are presented in this article were approved by the participating teams.

Table 3: Competition candidates with their system-description references and representing jury members

Participant	Ref.	Jury member	Affiliation
2LS	[31]	Peter Schrammel	U Oxford, UK
APROVE	[33]	Jera Hensel	RWTH Aachen, Germany
BLAST	[32]	Vadim Mutilin	ISPRAS, Russia
CASCADE	[35]	Wei Wang	New York U, USA
CBMC	[22]	Michael Tautschnig	Queen Mary U London, UK
CEAGLE		Dexi Wang	Tsinghua U, China
CEAGLE-ABSREF		Guang Chen	Tsinghua U, China
CIVL	[36]	Stephen Siegel	U Delaware, USA
CPA-BAM	[14]	Karlheinz Friedberger	U Passau, Germany
CPA-KIND	[7]	Matthias Dangl	U Passau, Germany
CPA-REFSEL	[9]	Stefan Löwe	U Passau, Germany
CPA-SEQ	[12]	—	U Passau, Germany
DIVINE	[37]	Vladimír Štill	Masaryk U, Czech Republic
ESBMC	[24]	Mikhail Ramalho	U Southampton, UK
ESBMC+DEPTHK	[28]	Lucas Cordeiro	Federal U Amazonas, Brazil
FOREST	[13]	Pablo Sanchez	U Cantabria, Spain
FORESTER	[18]	Ondřej Lengál	Brno UT, Czech Republic
HIPREC	[23]	Quang Loc Le	National U, Singapore
IMPARA		Björn Wachter	U Oxford, UK
LAZY-CSEQ	[19]	Omar Inverso	Gran Sasso Sc. Inst., Italy
LCTD	[30]	Keijo Heljanko	Aalto U, Finland
LPI	[20]	George Karpenkov	VERIMAG, France
MAP2CHECK	[29]	Herbert Rocha	Federal U Roraima, Brazil
MU-CSEQ	[34]	Gennaro Parlato	U Southampton, UK
PAC-MAN	[11]	Ming-Hsien Tsai	Academia Sinica, Taiwan
PREDATORHP	[21]	Tomas Vojnar	Brno UT, Czech Republic
SEAHORN	[15]	Jorge Navas	NASA Ames, USA
SKINK		Franck Cassez	Macquarie U, Australia
SMACK+CORRAL	[27]	Zvonimir Rakamaric	U Utah, USA
SYMBIOTIC	[10]	Jan Strejček	Masaryk U, Czech Republic
SYMDIVINE	[1]	Jiří Barnat	Masaryk U, Czech Republic
UAUTOMIZER	[17]	Matthias Heizmann	U Freiburg, Germany
UKOJAK	[26]	Daniel Dietsch	U Freiburg, Germany
UL-CSEQ	[25]	Bernd Fischer	Stellenbosch U, ZA
VVT	[16]	Alfons Laarman	TU Vienna, Austria

Table 4: Technologies and features that the verification tools offer

Verifier	CEGAR	Predicate Abstraction	Symbolic Execution	Bounded Model Checking	k-Induction	Property-Directed Reach.	Explicit-Value Analysis	Numeric. Interval Analysis	Shape Analysis	Separation Logic	Bit-Precise Analysis	ARG-Based Analysis	Lazy Abstraction	Interpolation	Automata-Based Analysis	Concurrency Support	Ranking Functions
2LS			✓	✓	✓			✓			✓						✓
APROVE			✓	✓	✓		✓	✓			✓						✓
BLAST	✓	✓					✓					✓	✓	✓			
CASCADE			✓	✓							✓						
CBMC				✓	✓						✓					✓	
CEAGLE			✓	✓	✓						✓						
CEAGLE-ABSREF	✓	✓		✓							✓		✓	✓			
CIVL			✓	✓				✓								✓	
CPA-BAM	✓	✓					✓				✓	✓	✓	✓			
CPA-KIND	✓	✓		✓	✓		✓	✓			✓	✓	✓	✓			
CPA-REFSEL	✓	✓					✓				✓	✓	✓	✓			
CPA-SEQ	✓	✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓		✓	
DIVINE							✓				✓				✓	✓	✓
ESBMC				✓							✓					✓	✓
ESBMC+DEPTHK			✓	✓	✓						✓					✓	
FOREST			✓	✓							✓						
FORESTER	✓								✓						✓		
HIPREC									✓	✓							
IMPARA			✓				✓	✓			✓	✓	✓	✓		✓	
LAZY-CSEQ				✓							✓					✓	
LCTD	✓	✓	✓								✓						
LPI	✓				✓			✓				✓	✓				
MAP2CHECK				✓							✓						
MU-CSEQ				✓							✓					✓	
PAC-MAN			✓												✓		
PREDATORHP									✓								
SEAHORN				✓		✓		✓				✓	✓	✓			✓
SKINK	✓						✓						✓	✓	✓		
SMACK+CORRAL	✓			✓		✓					✓		✓			✓	
SYMBIOTIC			✓														
SYMDIVINE			✓				✓				✓				✓	✓	
UAUTOMIZER	✓	✓									✓		✓	✓	✓		✓
UKOJAK	✓	✓									✓		✓	✓			
UL-CSEQ	✓	✓										✓	✓	✓		✓	
VVT	✓	✓		✓		✓	✓						✓	✓		✓	

Participating Verifiers. Table 3 provides an overview of the participating competition candidates and Table 4 lists the features and technologies that are used in the verification tools.

Computing Resources. The resource limits were the same as last year [5]: Each verification run was limited to 8 processing units (cores), 15 GB of memory, and 15 min of CPU time. The witness validation was limited to 2 processing units, 7 GB of memory, and 1.5 min of CPU time. The machines for running the experiments were different from last year, because we had to use 24 machines instead of eight. Each machine had two Intel Xeon E5-2650 v2 CPUs, with 16 processing units each, a frequency of 3.4 GHz, 135 GB of RAM, and a GNU/Linux operating system (x86_64-linux, Ubuntu 14.04 with Linux kernel 4.2). All verification runs were executed on a dedicated CPU, i.e., 8 processing units were assigned to the verification run, while the other 8 processing units were reserved and left idle.

One complete verification execution of the competition consisted of 313 benchmarks (each verifier on each selected category according to the opt-outs), summing up to 115 761 verification runs. Witness validation required 524 benchmarks (combinations of verifier, category with witness validation, and two validators) summing up to 50 249 validation runs. The consumed total CPU time for one competition run for verification only required a total of 319 days of CPU time. Each tool was executed several times, in order to make sure no installation issues occur during the execution.

Quantitative Results. Table 5 presents the quantitative overview over all tools and all categories (HIPREC participated only in subcategory *Recursive* and LCTD only in subcategory *BitVectorsReach*). The format of the table is similar to those of previous SV-COMP editions [5], with the exception that due to the volume we now omit the CPU times. The tools are listed in alphabetical order; every table row lists the scores of one verifier for each category. We indicate the top-three candidates by formatting their scores in bold face and in larger font size. An empty table cell means that the verifier opted-out from the respective category. For the calculation of the score and for the ranking, the scoring schema in Table 2 was applied, the scores for the meta categories were computed using normalized scores as defined in the report for SV-COMP’13 [3]. There were two categories for which the winner was decided based on the run time: in category *Concurrency*, all top-three verifiers achieved the maximum score of 1240 points, but the run time differed considerably; in category *Floats* the first and second both achieved a score of 136 points. More information (including formatted interactive tables, quantile plots for every category, and also the raw data in XML format) is available on the competition web-site.¹⁰

Table 6 reports the top-three verifiers for each category. The run time (column ‘CPU Time’) refers to successfully solved verification tasks (column ‘Solved Tasks’). The columns ‘False Alarms’ and ‘Wrong Proofs’ report the number of verification tasks for which the tool reported wrong results: reporting an error path but the property holds (incorrect FALSE) and claiming that the program fulfills the property although it actually contains a bug (incorrect TRUE), respectively.

¹⁰ <http://sv-comp.sosy-lab.org/2016/results/>

Table 5: Quantitative overview over all results

Verifier	Arrays 316 points max. 183 tasks	Bit Vectors 92 points max. 60 tasks	Heap 382 points max. 239 tasks	Floats 140 points max. 81 tasks	IntegersControlFlow 3629 points max. 2331 tasks	Termination 1129 points max. 631 tasks	Concurrency 1240 points max. 1016 tasks	DeviceDriversLinux64 3977 points max. 2120 tasks	FalsificationOverall 2371 points max. 6030 tasks	Overall 10855 points max. 6661 tasks
2LS				136	1196				-2438	-38205
APROVE						909				
BLAST					-1653			2704		
CASCADE			197							
CBMC	62	46	8	134	-1239		882	1972	391	3386
CEAGLE				136						
CEAGLE-ABSREF				124						
CIVL							1240			
CPA-BAM	-57	28	-80	42	1822	0	0	2550	-1218	1939
CPA-KIND	3	77	161	76	2095	0	0	2350	707	4094
CPA-REFSEL				35	1539	0	0	3177	36	2157
CPA-SEQ	-61	87	234	75	2652	0	282	2801	496	4794
DIVINE							951			
ESBMC	190	84	163	-15	1217	0	742	1688	248	4145
ESBMC+DEPTHK	62	47	58	7	1111	0	877	2009	495	3110
FOREST	-970		-1263				-20613			
FORESTER			86							
HIPREC										
IMPARA		-592		132	-1524		42			
LAZY-CSEQ							1240			
LCTD										
LPI						1804		2107		
MAP2CHECK			-121							
MU-CSEQ							1240			
PAC-MAN				-449						
PREDATORHP			298							
SEAHORN	-301	-131	-257	0	1572	504	-24659	1694	-4333	-22393
SKINK					113					
SMACK+CORRAL	146	44	155	0	2013	0	999	2206	800	4223
SYMBIOTIC	101	-2	105	-18	633	0	0	980	-370	1223
SYMDIVINE							-135			
UAUTOMIZER	83	69	169	2	1865	895		2686	823	4843
UKOJAK	60	19	31	0	1096			937	339	1407
UL-CSEQ							856			
VVT					421		1029			

Table 6: Overview of the top-three verifiers for each category (CPU time in h, rounded to two significant digits)

Rank	Verifier	Score	CPU Time	Solved Tasks	False Alarms	Wrong Proofs
<i>Arrays</i>						
1	ESBMC	190	3.2	131	2	
2	SMACK+CORRAL	146	2.5	111		
3	SYMBIOTIC	101	.61	77		
<i>Bit Vectors</i>						
1	CPA-SEQ	87	1.1	55		
2	ESBMC	84	.61	51		
3	CPA-KIND	77	.67	47		
<i>Heap</i>						
1	PREDATORHP	298	.31	211	2	
2	CPA-SEQ	234	1.1	188	4	
3	CASCADE	197	2.7	140	2	
<i>Floats</i>						
1	2LS	136	.98	79		
2	CEAGLE	136	1.0	77		
3	CBMC	134	5.0	78		
<i>Integers/ControlFlow</i>						
1	CPA-SEQ	2652	35	1625	1	
2	CPA-KIND	2095	35	1278		
3	SMACK+CORRAL	2013	97	978		4
<i>Termination</i>						
1	APROVE	909	4.8	500		
2	UAUTOMIZER	895	3.2	503		
3	SEAHORN	504	.97	323		2
<i>Concurrency</i>						
1	MU-CSEQ	1240	.93	1016		
2	LAZY-CSEQ	1240	2.7	1016		
3	CIVL	1240	7.8	1016		
<i>Device Drivers Linux64</i>						
1	CPA-REFSEL	3177	24	1646	2	
2	CPA-SEQ	2801	23	1458	4	
3	BLAST	2704	5.9	1547	13	5
<i>Falsification Overall</i>						
1	UAUTOMIZER	823	7.0	381	1	
2	SMACK+CORRAL	800	17	1140	26	
3	CPA-KIND	707	14	479	2	
<i>Overall</i>						
1	UAUTOMIZER	4843	44	3138	1	5
2	CPA-SEQ	4794	65	3535	16	
3	SMACK+CORRAL	4223	160	3464	26	9

Discussion of Scoring Schema and Normalization. The SV-COMP community considers it more difficult to compute correctness proofs compared to computing error paths (cf. Table 2: TRUE yields 2 points, FALSE yields 1 point) [2]. This has consequences on the final ranking: For example, APROVE won the category *Termination* although UAUTOMIZER solved more verification tasks: APROVE solved 500, UAUTOMIZER solved 503 verification tasks. Both verifiers did not report any wrong results in this category. So the higher score of APROVE (score: 909) is due to its ability to compute more proofs than UAUTOMIZER (score: 895), while UAUTOMIZER found more violations. APROVE computed 409 proofs and found 91 property violations, while UAUTOMIZER computed 392 proofs and found 111 property violations. So in this case, the scoring schema provides a good mapping from the community’s intuition to the ranking.

A similar observation can be made on the score normalization. The community considers the value of solving a verification task in a large category (many verification tasks) less than the value of solving a verification task in a small category (only a few verification tasks) [3]. The values for category *Overall* in Table 6 illustrate the purpose of the score normalization: CPA-SEQ solved 3535 tasks, which is about 400 solved tasks more than the winner UAUTOMIZER could solve (3138). So why did CPA-SEQ not win the category? Because UAUTOMIZER is better in the intuitive sense of ‘overall’: UAUTOMIZER solved tasks more diversely, the ‘overall’ value of the verification work is higher. Most prominently, UAUTOMIZER solved many tasks in category *Termination* which is not supported by CPA-SEQ. Similarly, in category *FalsificationOverall*, SMACK+CORRAL solved more tasks than UAUTOMIZER, but produced also a lot of false alarms and the tasks that SMACK+CORRAL solved were considered of less value (i.e., from large categories with many tasks). In these cases, the score normalization correctly maps the community’s intuition.

Score-Based Quantile Functions for Quality Assessment. We use score-based quantile functions [3] because these visualizations make it easier to understand the results of the comparative evaluation. The competition web-site¹⁰ includes such a plot for each category; as example, we illustrate the category *Overall* (all verification tasks) in Fig. 3 and discuss the results below. A total of 13 verifiers participated in category *Overall* (only 6 the year before), for which the quantile plot shows the overall performance over all categories (scores for meta categories are normalized [3]).

Overall Quality Measured in Scores (Right End of Graph). UAUTOMIZER is the winner of this category: the x -coordinate of the right-most data point represents the highest total score (and thus, the total value) of the completed verification work (cf. Table 6; right-most x -coordinates match the score values in the table).

Amount of Incorrect Verification Work (Left End of Graph). The left-most data points of the quantile functions represent the total negative score of a verifier (x -coordinate), i.e., the amount of incorrect and misleading verification work. Verifiers should start with a score close to zero; the winner UAUTOMIZER is very good in this aspect, together with the second place CPA-KIND (the two right-most columns of category *Overall* in Table 6 report the concrete numbers:




Fig. 3: Quantile functions for category *Overall*. Each quantile function illustrates the quantile (x -coordinate) of the scores obtained by correct verification runs below a certain run time (y -coordinate). More details are given in a previous report [3]. A logarithmic scale is used for the time range from 1 s to 1000 s, and a linear scale is used for the time range between 0 s and 1 s.

only 1 and 16 false alarms, respectively, and 5 and 0 wrong proofs, for a total of 6 661 verification tasks).

Characteristics of the Verification Tools. Quantile plots also give hints on how a verification strategy works. For example, the horizontal lines show that some verifiers ‘solve’ a large quantity of verification tasks in the same run time, suggesting that an answer is given without the result being actually computed. A quick look at the wrapping execution scripts reveals that indeed a pre-mature answer is returned after 850 s or 880 s, respectively. This insight is one of the arguments for the community’s goal to have each result supported by evidence, e.g., in the form of a verification witness.

Robustness, Soundness, and Completeness. Table 6 shows in the last two columns that the best verifiers of each category report a low number of wrong verification results (compared to the large number of verification tasks), indicating the advancement of the state-of-the-art verification technology. In the three categories *BitVectors*, *Floats*, and *Concurrency*, the top-three verifiers did not report any wrong results.

Verifiable Witnesses. SV-COMP counts answers FALSE (bug reports) only if the result contains a violation witness, which represents directions through the state space to easily recover an error path. All verifiers in categories that required witness validation supported the common exchange format for error witnesses, and produced error paths in that format. For SV-COMP 2016, we used two completely different witness validators: CPACHECKER and UAUTOMIZER.

Table 7: Validation of Correctness Witnesses

	Verification	Validation	
		CPACHECKER	UAUTOMIZER
Total tasks	3 171	1 950	1 950
Results TRUE	1 950	1 029	1 149
Confirmed witnesses		53 %	59 %

Demonstration on Correctness Witnesses. The validation of the results for answers TRUE was not yet considered, but is identified as the next open problem that the community should solve. As part of SV-COMP 2016, a demonstration category (i.e., without ranking and scores) was announced to explore the possibilities of validating correctness witnesses. Two teams participated, and the results are reported in Table 7. The table lists the results of a verification with CPACHECKER (k-induction-based configuration) and the validation results of the correctness witnesses using the validators CPACHECKER and UAUTOMIZER. The first row reports the total number of verification tasks that were given as input. The verification was performed on an SV-COMP subset of 3 171 verification tasks from the categories *IntegersControlFlow* and *DeviceDriversLinux64*. The second row reports that for 1 950 verification tasks the expected and computed verification result was TRUE. Those 1 950 verification tasks were given as input to the two validators, together with the correctness witness that the verification produced. CPACHECKER was able to validate (i.e., re-verify with the given invariants from the witness) 1 029 verification tasks (53 %) and UAUTOMIZER was able to validate 1 149 verification tasks (59 %). More information is given in the detailed table on the web page.¹¹ Considering the short time slot for developing the validators (in the context of the demonstration category), the achieved results show the high potential of this technology, and that it should be further developed and used in practice.

6 Conclusion

SV-COMP 2016, the 5th edition of the Competition on Software Verification, attracted *35 participating teams* from 16 countries, which is so far the largest number of participants (2012: 10, 2013: 11, 2014: 15, 2015: 22). The repository of verification tasks was consolidated and the number of verification tasks was increased (from 5 803) to *6 661 verification tasks*. We used *verifiable witnesses* again to validate the bug reports, and the results FALSE were counted towards the score only if the witness was confirmed. The number of witness validators was increased from one to two, which contributed to the trust and neutrality of SV-COMP’s evaluation. SV-COMP 2016 is the so-far broadest overview of the state of the art in software verification. The large jury and the organizer made sure that the competition follows the high quality standards of the TACAS conference, in particular with respect to the important principles of fairness, community support, and transparency. Technical accuracy was ensured by using the benchmarking framework BENCHEXEC.

¹¹ <http://sv-comp.sosy-lab.org/2016/witnesses/correctness-demo.html>

References

1. P. Bauch, V. Havel, and J. Barnat. LTL model checking of LLVM bitcode with symbolic data. In *Proc. MEMICS*, pages 47–59. Springer, 2014.
2. D. Beyer. Competition on software verification (SV-COMP). In *Proc. TACAS*, LNCS 7214, pages 504–524. Springer, 2012.
3. D. Beyer. Second competition on software verification. In *Proc. TACAS*, LNCS 7795, pages 594–609. Springer, 2013.
4. D. Beyer. Status report on software verification. In *Proc. TACAS*, LNCS 8413, pages 373–388. Springer, 2014.
5. D. Beyer. Software verification and verifiable witnesses. In *Proc. TACAS*, LNCS 9035, pages 401–416. Springer, 2015.
6. D. Beyer, M. Dangl, D. Dietsch, M. Heizmann, and A. Stahlbauer. Witness validation and stepwise testification across software verifiers. In *Proc. FSE*, pages 721–733. ACM, 2015.
7. D. Beyer, M. Dangl, and P. Wendler. Boosting k-induction with continuously-refined invariants. In *Proc. CAV*, LNCS 9206, pages 622–640. Springer, 2015.
8. D. Beyer, S. Löwe, and P. Wendler. Benchmarking and resource measurement. In *Proc. SPIN*, LNCS 9232, pages 160–178. Springer, 2015.
9. D. Beyer, S. Löwe, and P. Wendler. Refinement selection. In *Proc. SPIN*, LNCS 9232, pages 20–38. Springer, 2015.
10. M. Chalupa, M. Jonáš, J. Slaby, J. Strejček, and M. Vitovská. Symbiotic 3: New slicer and error-witness generation (competition contribution). In *Proc. TACAS*. Springer, 2016.
11. Y.-F. Chen, C. Hsieh, O. Lengál, T.-J. Lii, M.-H. Tsai, B.-Y. Wang, and F. Wang. Learning-based verification and model synthesis. In *Proc. ICSE*, 2016.
12. M. Dangl, S. Löwe, and P. Wendler. CPACHECKER with support for recursive programs and floating-point arithmetic. In *Proc. TACAS*. Springer, 2015.
13. P. G. de Aledo and P. Sanchez. Framework for embedded system verification (competition contribution). In *Proc. TACAS*. Springer, 2015.
14. K. Friedberger. CPA-BAM: Block-abstraction memoization with value analysis and predicate analysis (competition contribution). In *Proc. TACAS*. Springer, 2016.
15. A. Gurfinkel, T. Kahsai, and J. A. Navas. SeaHorn: A framework for verifying C programs (competition contribution). In *Proc. TACAS*. Springer, 2015.
16. H. Günther, A. Laarman, and G. Weissenbacher. Vienna Verification Tool: IC3 for parallel software (competition contribution). In *Proc. TACAS*. Springer, 2016.
17. M. Heizmann, D. Dietsch, M. Greitschus, J. Leike, B. Musa, C. Schätzle, and A. Podelski. Ultimate Automizer with two-track proofs (competition contribution). In *Proc. TACAS*. Springer, 2016.
18. M. Hruška, L. Holík, O. Lengál, A. Rogalewicz, J. Šimáček, and T. Vojnar. Run forester, run backwards! (competition contribution). In *Proc. TACAS*. Springer, 2016.
19. O. Inverso, E. Tomasco, B. Fischer, S. La Torre, and G. Parlato. Bounded model checking of multi-threaded C programs via lazy sequentialization. In *Proc. CAV*, LNCS 8559, pages 585–602. Springer, 2014.
20. E. G. Karpenkov, D. Monniaux, and P. Wendler. Program analysis with local policy iteration. In *Proc. VMCAI*, pages 127–146. Springer, 2016.
21. M. Kotoun, P. Peringer, V. Šoková, and T. Vojnar. Optimized Predators and the SV-COMP heap and memory safety benchmark (competition contribution). In *Proc. TACAS*. Springer, 2016.

22. D. Kröning and M. Tautschnig. CBMC: C bounded model checker (competition contribution). In *Proc. TACAS*, LNCS 8413, pages 389–391. Springer, 2014.
23. Q. L. Le, M. Tran, and W.-N. Chin. HIPrec: Verifying recursive programs with a satisfiability solver. Technical report, 2016.
24. J. Morse, M. Ramalho, L. Cordeiro, D. Nicole, and B. Fischer. ESBMC 1.22 (competition contribution). In *Proc. TACAS*, LNCS 8413, pages 405–407. Springer, 2014.
25. T. L. Nguyen, B. Fischer, S. La Torre, and G. Parlato. Unbounded Lazy-CSeq: A lazy sequentialization tool for C programs with unbounded context switches (competition contribution). In *Proc. TACAS*. Springer, 2015.
26. A. Nutz, D. Dietsch, M. M. Mohamed, and A. Podelski. Ultimate Kojak with memory safety checks (competition contribution). In *Proc. TACAS*. Springer, 2015.
27. Z. Rakamarić and M. Emmi. SMACK: Decoupling source language details from verifier implementations. In *Proc. CAV*, LNCS 8559, pages 106–113. Springer, 2014.
28. H. Rocha, H. I. Ismail, L. C. Cordeiro, and R. S. Barreto. Model checking embedded C software using k-induction and invariants. In *Proc. SBESC*. IEEE, 2015.
29. H. O. Rocha, R. Barreto, and L. Cordeiro. Hunting memory bugs in c programs with Map2Check (competition contribution). In *Proc. TACAS*. Springer, 2016.
30. O. Saarikivi and K. Heljanko. LCTD: Tests-guided proofs for C programs on LLVM (competition contribution). In *Proc. TACAS*. Springer, 2016.
31. P. Schrammel and D. Kröning. 2LS for program analysis (competition contribution). In *Proc. TACAS*. Springer, 2016.
32. P. Shved, M. Mandrykin, and V. Mutilin. Predicate analysis with BLAST 2.7 (competition contribution). In *Proc. TACAS*. Springer, 2012.
33. T. Ströder, C. Aschermann, F. Frohn, J. Hensel, and J. Giesl. AProVE: Termination and memory safety of C programs (competition contribution). In *Proc. TACAS*. Springer, 2015.
34. E. Tomasco, T. N. Lam, O. Inverso, B. Fischer, S. L. Torre, and G. Parlato. MUCSeq 0.4: Individual memory location unwindings (competition contribution). In *Proc. TACAS*. Springer, 2016.
35. W. Wang and C. Barrett. Cascade (competition contribution). In *Proc. TACAS*. Springer, 2015.
36. M. Zheng, J. G. Edenhofner, Z. Luo, M. J. Gerrard, M. B. Dwyer, and S. F. Siegel. CIVL: Applying a general concurrency verification framework to C/Pthreads programs (competition contribution). In *Proc. TACAS*. Springer, 2016.
37. V. Štill, P. Ročkai, and J. Barnat. DIVINE: Explicit-state LTL model checker (competition contribution). In *Proc. TACAS*. Springer, 2016.