

CASE-V12

CASE-V12

CASE-V12

CASE-V12

Proceedings of the 12th IJCAR ATP System Competition (CASC-J12)

Geoff Sutcliffe

University of Miami, USA

Abstract

The CADE ATP System Competition (CASC) is the annual evaluation of fully automatic, classical logic, Automated Theorem Proving (ATP) systems - the world championship for such systems. CASC-J12 was the twenty-ninth competition in the CASC series. Nineteen ATP systems competed in the various divisions. These proceedings present the competition design and information about the competing systems.

1 Introduction

The CADE ATP System Competition (CASC) [105] is the annual evaluation of fully automatic, classical logic, ATP systems – the world championship for such systems. One purpose of CASC is to provide a public evaluation of the relative capabilities of ATP systems. Additionally, CASC aims to stimulate ATP research, motivate development and implementation of robust ATP systems that can be easily and usefully deployed in applications, provide an inspiring environment for personal interaction between ATP researchers, and expose ATP systems within and beyond the ATP community. CASC evaluates the performance of the ATP systems in terms of the number of problems solved, the number of acceptable solutions (proofs or models) output, and the average time taken for problems solved, in the context of a bounded number of eligible problems and specified time limits.

CASC is held at each CADE (the International Conference on Automated Deduction) and IJCAR (the International Joint Conference on Automated Reasoning) conference – the major forums for the presentation of new research in all aspects of automated deduction. CASC-J12 was held on 4th July 2024, as part of the 12th International Joint Conference on Automated Reasoning (IJCAR 2024)¹, in Nancy, France. It was the twenty-ninth competition in the CASC series; see [128, 134, 131, 74, 76, 127, 125, 126, 81, 83, 85, 87, 90, 92, 94, 96, 98, 100, 102, 133, 104, 107, 110, 113, 115, 120, 121, 122] and the CASC web site <http://tptp.org/CASC>, for information about previous CASCs. CASC-J12 was organized by Geoff Sutcliffe, and overseen by a panel consisting of Cláudia Nalon, Christoph Wernhard, and Christoph Weidenbach. The competition was run on computers provided by the StarExec project [69] at the University of Miami. The CASC-J12 web site provides access to all the resources used before, during, and after the event: <http://tptp.org/CASC/29>.

The design and organization of CASC has evolved over the years to a sophisticated state [128, 129, 124, 130, 72, 73, 75, 77, 78, 79, 80, 82, 84, 86, 89, 91, 93, 95, 97, 99, 101, 103, 106, 109, 111, 112, 114, 116, 117]. Important changes for CASC-J12 were (for readers already familiar with the general design of CASC):

- The FNT division went on hiatus.
- The ICU division was added.

¹CADE was a constituent conference of the 12th IJCAR, hence CASC-“J12”.

The CASC rules, specifications, and deadlines are absolute. Only the panel has the right to make exceptions. It is assumed that all entrants have read the documentation related to the competition, and have complied with the competition rules. Non-compliance with the rules can lead to disqualification. A “catch-all” rule is used to deal with any unforeseen circumstances: *No cheating is allowed*. The panel is allowed to disqualify entrants due to unfairness, and to adjust the competition rules in case of misuse.

These proceedings are organized as follows: Section 2 describes the competition divisions and the ATP systems that entered the various divisions. Sections 3 and 4 describe the competition infrastructure and the requirements for the ATP systems. Section 5 describes how the systems are evaluated. Sections 6 and 7 describe the practical steps for registering a system and checking that it has the required properties. Section 8 provides descriptions (written by the entrants) of the systems entered into CASC-J12. Section 9 concludes.

A Tense Note: Attentive readers will notice changes between the present and past tenses in this paper. Many parts of CASC are established and stable – they are described in the present tense (the rules are the rules). Aspects that were particular to CASC-J12 are described in the past tense so that they make sense when reading this after the event.

2 Divisions and Systems

CASC is divided into divisions according to problem and system characteristics, in a coarse version of the TPTP problem library’s Specialist Problem Classes (SPCs) [132]. Each division uses problems that have certain logical, language, and syntactic characteristics, so that the systems that compete in a division are, in principle, able to attempt all the problems in the division. Some divisions are further divided into problem categories that make it possible to analyze, at a more fine-grained level, which systems work well for what types of problems. Table 1 catalogs the divisions and problem categories of CASC-J12. The example problems can be viewed online at <http://tptp.org/cgi-bin/SeeTPTP?Category=Problems>. Sections 3.2 and 3.3 explain what problems are eligible for use in each division and category.

Systems that do not run in the competition divisions for any reason (e.g., the system requires special hardware, or the entrant is an organizer) can be entered into the demonstration division. The demonstration division uses the same problems as the competition divisions, and the entry specifies which competition divisions’ problems are to be used.

Nineteen ATP systems competed in the various divisions of CASC-J12. The division winners from the previous CASC (CASC-29) and the Prover9 1109a system are automatically entered into the corresponding demonstration divisions to provide benchmarks against which progress can be judged. The systems, the divisions in which they were entered, and their entrants, are listed in Tables 2 and 3. A division acronym in *italics* indicates the system was in the demonstration division. System descriptions are in the competition proceedings [118] and on the CASC-J12 web site.

3 Infrastructure

3.1 Computers

The competition computers had:

- An octa-core Intel(R) Xeon(R) E5-2667, 3.20GHz CPUs, without hyperthreading
- 128GB memory

Table 1: Divisions and Problem categories

Division	Problems	Problem categories
THF	Typed (monomorphic) H igher-order F orm theorems (axioms with a provable conjecture).	TNE – THF with N o E quality, e.g., NUM738 ¹ . TEQ – THF with E Quality, e.g., SET171 ³ .
TFA	Typed (monomorphic) F irst-order form theorems with A rithmetic (axioms with a provable conjecture).	TFI – TFA with only I nteger arithmetic, e.g., DAT016 ¹ . TFE – TFA with only rE al arithmetic, e.g., MSC022 ² .
TFN	Typed F irst-order form N on-theorems (axioms with a countersatisfiable conjecture, and satisfiable axiom sets without a conjecture) without arithmetic.	E.g., COM002 ²⁰ .
FOF	F irst- O rd E r F orm theorems (axioms with a provable conjecture).	FNE – FOF with N o E quality, e.g., COM003 ¹ . FEQ – FOF with E Quality, e.g., SEU147 ³ .
UEQ	U nit E Quality theorems in clause normal form (unsatisfiable clause sets).	E.g., RNG026 ⁷ .
SLH	Typed (monomorphic) higher-order theorems (axioms with a provable conjecture), generated by Isabelle’s S Ledge H ammer system [47].	The problems were generated from sessions in Isabelle’s Archive of Formal Proofs [14].
ICU	First-order theorems (axioms with a provable conjecture) provided by the entrants.	The problems supplied by an entrant were expected to be easy enough for that entrant’s ATP system, but difficult for the other entrants’ systems, i.e., each entrant is saying to the others: “I Challenge yoU!”.

- The CentOS Linux release 7.4.1708 (Core) operating system, with Linux kernel 3.10.0-693.el7.x86_64.

One ATP system runs on one CPU at a time. (Each StarExec computer has two sockets, i.e., two CPUs, and 256 GiB memory. StarExec uses Linux’s `sched_setaffinity` to restrict each system run to a single CPU, and `setrlimit` to limit memory use to 128 GiB.) Systems can use all the cores on the CPU, which can be advantageous in divisions where a wall clock time limit is used. StarExec copies the systems and problems to the compute nodes before starting execution, so that there are no network delays. The StarExec computers used for CASC are the same as are publicly available to the TPTP community, which allows system developers to test and tune their systems in exactly the same environment as is used for the competition.

Demonstration division systems can run on the competition computers, or the computers can be supplied by the entrant. The CASC-J12 demonstration division systems all used the competition computers.

ATP System	Divisions	Entrant (Associates)	Entrant's Affiliation
Connect++ 0.6.0			
CSE 1.7	FOF	Sean Holden	University of Cambridge
	FOF	Feng Cao (Yang Xu, Peiyao Liu, Jun Liu, Shuwei Chen, Guoyan Zeng, Jian Zhong, Guanfeng Wu, Xingxing He, Peng Xu)	JiangXi University of Science and Technology
CSELE 1.6	FOF	Peiyao Liu (Yang Xu, Peng Cao, Stephan Schulz, Jun Liu, Shuwei Chen, Guoyan Zeng)	Southwest Jiaotong University
CSGE 1.6	FOF	Peiyao Liu (Yang Xu, Feng Cao, Stephan Schulz, Jun Liu, Shuwei Chen, Guoyan Zeng)	Southwest Jiaotong University
CSLE 1.6	FOF	Peiyao Liu (Yang Xu, Feng Cao, Stephan Schulz, Jun Liu, Shuwei Chen, Guoyan Zeng)	Southwest Jiaotong University
cvc5 1.1.3	THF TFA TFN FOF	Andrew Reynolds (Haniel Barbosa, Cesare Tinelli, Clark Barrett)	University of Iowa
Drodi 3.6.0	FOF UEQ	Oscar Contreras	Amateur Programmer
E 3.1		CASC	CASC-29 winner
E 3.2.0	THF	Stephan Schulz	DHBW Stuttgart
GKC 0.8	FOF UEQ	Tanel Tammet	Tallinn University of Technology
iProver 3.9	TFA TFN FOF UEQ	Konstantin Korovin (Edward Holden)	University of Manchester

Table 2: The ATP systems and entrants

ATP System	Divisions	Entrant (Associates)	Entrant's Affiliation
LEO-II 1.7.0	THF	Alexander Steen (Christoph Benz Müller)	University of Greifswald
Leo-III 1.7.8	THF	Alexander Steen (Christoph Benz Müller)	University of Greifswald
Prover9 1109a		CASC (Bob Veroff, Bill McCune)	CASC fixed point
Twee 2.4.2		CASC	CASC-29 winner
Twee 2.5.0		Nick Smallbone	Chalmers University of Technology
Vampire 4.8	THF TFA TFN FOF	CASC	CASC-29 winner
Vampire 4.9	THF TFA TFN FOF UEQ SLH ICU	Michael Rawson (Ahmed Bhayat, Robin Coutelier, Márton Hajdú, Matthias Hetzenberger, Petra Hozzová, Laura Kovács, Jakob Rath, Giles Reger, Martin Reiner, Martin Suda, Johannes Schoiswohl, Andrei Voronkov)	TU Wien
Zipperp'n 2.1.9999	THF FOF	Jasmin Blanchette (Alexander Bentkamp, Simon Cruanes, Petar Vukmirović)	Ludwig-Maximilians-Universität München

Table 3: The ATP systems and entrants, continued

3.2 Problems for the TPTP-based Divisions

The problems for the THF, TFA, TFN, FOF, and UEQ divisions were taken from the Thousands of Problems for Theorem Provers (TPTP) problem library [108], v9.0.0. The TPTP version used for CASC is released after the competition has started, so that new problems in the release have not been seen by the entrants. The problems have to meet certain criteria to be eligible for use:

- The TPTP tags problems that are designed specifically to be suited or ill-suited to some ATP system, calculus, or control strategy as *biased*. They are excluded from the competition.
- The problems must be syntactically non-propositional.
- The TPTP uses system performance data in the Thousands of Solutions from Theorem Provers (TSTP) solution library to compute problem difficulty ratings in the range 0.00 (easy) to 1.00 (unsolved) [132] – Problems with ratings in the range 0.21 to 0.99 are eligible – the upper bound of 0.99 excludes problems that cannot be solved by any system and thus don't differentiate between systems; the lower bound of 0.21 was chosen (many years ago) to exclude almost-easy problems that would be solved by most of the systems and thus don't differentiate between systems (the choice of 0.21 has been fortuitously successful).

Problems of lesser and greater ratings are made eligible if there are not enough problems with ratings in that range. For the TFN division of CASC-J12, 59 problems with rating 0.00-0.20 and 48 problems with rating 1.00 were made eligible, because there were only 55 eligible problems with rating 0.21 to 0.99. The organizer considered making these additional 107 problems eligible to be acceptably useful: solving easy problems would be encouraging for weaker systems, and solving hard problems would be encouraging for stronger systems.

Systems can be submitted before the competition so that their performance data is used in computing the problem ratings – problems that are solved get a rating less than 1.00 and are thus eligible (unless the rating drops below 0.21). The computing of the ratings also uses performance data from ATP systems that are not entered into the competition, which can produce ratings that make some problems eligible for selection but easy or unsolvable for the systems in the competition. Using problems that are solved by all or none of the competition systems does not affect the competition rankings, has the benefit of placing the systems' performances in the context of the state-of-the-art in ATP, but does reduce the differentiation between the systems in the competition.

In order to ensure that no system receives an advantage or disadvantage due to the specific presentation of the problems in the TPTP, the problems are preprocessed to

- strip out all comment lines (in particular, the problem header)
- randomly reorder the formulae (`include` directives are left before the formulae, and type declarations are left before the symbols' uses)
- randomly swap the arguments of associative connectives and randomly reverse implications
- randomly reverse equalities

The numbers of problems used in each division and problem category are constrained by the numbers of eligible problems, the number of systems entered across the divisions, the number of CPUs available, the time limits, and the time available for running the competition live in one conference day, i.e., in about 6 hours. The numbers of problems used are set within these

constraints, according to the judgement of the organizers. The problems used are randomly selected from the eligible problems based on a seed supplied by the competition panel:

- The selection is constrained so that no division or category contains an excessive number of very similar problems, according to the “very similar problems” (VSP) lists distributed with the TPTP [74]: For each problem category in each division, if the category is going to use N problems and there are L VSP lists that have an intersection of at least $N/(L+1)$ with the eligible problems for the category, then maximally $N/(L+1)$ problems are taken from each of those VSP lists.
- In order to combat excessive tuning towards problems that were already in the preceding released TPTP version, the selection is biased to select problems that are new in the TPTP version used until 50% of the problems in each problem category have been selected or there are no more new problems to select, after which random selection from old and new problems continues. The number of new problems used depends on how many new problems are eligible and the limitation on very similar problems.
- Problems with rating 0.21 to 0.99 are selected before problems with other ratings.

The problems are given to the ATP systems as files in TPTP format, with `include` directives, in increasing order of TPTP difficulty rating.

3.3 Problems for the SLH Division

For the SLH division of CASC-29, Isabelle’s Sledgehammer system was used to generate 8400 problems that could be used, of which 1000 appropriately difficult problems were selected based on performance data [117, 122]. For the SLH division of CASC-J12, the same problem set was used, but 1000 different problems were selected. The same CPU limit was imposed. This was announced in advance of CASC-J12, so that developers could tune their systems using the CASC-29 problems. The problems are given in a roughly estimated increasing order of difficulty.

3.4 Problems for the ICU Division

Each entrant had to submit 5 to 10 FOF problems that have a conjecture. The problems had to be provided in decreasing order of desired use in the division, i.e., probably from hardest to easiest for other systems. The problems had to all be different, as assessed by the competition organizers. Problems from the TPTP problem library had their `include` directives expanded. At least five problems were taken from each entrant’s submission, in the order specified, and all systems attempted all the selected problems. The problems were given in the reverse of the decreasing order of desired use, so that the “easier” problems were used before “harder” ones.

It was expected that each entrant would submit problems that are easy enough for that entrant’s system, but difficult for the other entrants’ systems, i.e., each entrant is saying to the others: “I Challenge yoU!”.

3.5 Time Limits

In the TPTP-based divisions time limits are imposed for each problem. The minimal time limit for each problem is 120 s. The maximal time limit for each problem is constrained by the same factors that constrain the numbers of problems that are used. The time limit is chosen as a reasonable value within the range allowed according to the judgement of the organizers, and is announced at the competition. In CASC-J12 a wall clock time limit was imposed for each

problem, and no CPU time limits were imposed (so that it can be advantageous to use all the cores on the CPU).

In the SLH division a CPU time limit is imposed for each problem. The limit is between 15 s and 90 s, which is the range of CPU time that can be usefully allocated for a proof attempt in the Sledgehammer context.² The time limit is chosen as a reasonable value within the range allowed according to the judgement of the organizers, and is announced at the competition.

In the ICU division a wall clock time limit is imposed for each problem, and no CPU time limits are imposed (so that it can be advantageous to use all the cores on the CPU). The limit is between 300 s and 600 s, and is announced at the competition.

4 System Entry, Delivery, and Execution

Systems can be entered at only the division level, and can be entered into more than one division. A system that is not entered into a division is assumed to perform worse than the entered systems, for that type of problem – wimping out is not an option. Entering many similar versions of the same system is deprecated, and entrants might be required to limit the number of system versions that they enter. Systems that rely essentially on running other ATP systems without adding value are deprecated; such systems might be disallowed or moved to the demonstration division.

The ATP systems entered into CASC are delivered to the competition organizer as StarExec installation packages, which the organizer installs and tests on StarExec. Source code is delivered separately, under the trusting assumption that the installation package does correspond to the source code. After the competition all competition division systems’ StarExec and source code packages are made available on the CASC web site. This allows anyone to use the systems on StarExec, and to examine the source code. An open source license is encouraged, to allow the systems to be freely used, modified, and shared. Many of the StarExec packages include statically linked binaries that provide further portability and longevity of the systems.

For systems running on entrant supplied computers in the demonstration division, entrants must email a `.tgz` file containing the source code and any files required for building the executable system to the competition organizers by the system delivery deadline. In the demonstration division the entrant specifies whether or not the source code is placed on the CASC web site.

The ATP systems must be fully automatic. They are executed as black boxes, on one problem at a time. Any command line parameters have to be the same for all problems in each division. The ATP systems must be sound, and are tested for soundness by submitting non-theorems to the systems in the THF, TFA, FOF, UEQ, and SLH divisions, and theorems to the systems in the TFN division. Claiming to have found a proof of a non-theorem or a disproof of a theorem indicates unsoundness. If a system fails the soundness testing it must be repaired by the unsoundness repair deadline or be withdrawn.

5 System Evaluation

CASC ranks the ATP systems at only the division level. For each ATP system, for each problem, four items of data are recorded: whether or not the problem was solved, the CPU and wall clock times taken (as measured by StarExec’s `runsolver` utility [59], and prepended to each line of the system’s `stdout`), and whether or not a solution (proof or model) was output. The

²According to a personal communication from Jasmin Blanchette, and he should know.

systems are ranked according to the number of problems solved with an acceptable solution output. Ties are broken according to the average time taken over problems solved. Trophies are awarded to the competition divisions' winners.

The competition panel decides whether or not the systems' solutions are "acceptable". The criteria include:

- Derivations must be complete, starting at formulae from the problem, and ending at the conjecture (for axiomatic proofs) or a *false* formula (for proofs by contradiction, e.g., CNF refutations).
- For solutions that use translations from one form to another, e.g., translation of FOF problems to CNF, the translations must be adequately documented.
- Derivations must show only relevant inference steps.
- Inference steps must document the parent formulae, the inference rule used, and the inferred formula.
- Inference steps must be reasonably fine-grained.
- An unsatisfiable set of ground instances of clauses is acceptable for establishing the unsatisfiability of a set of clauses.
- Models must be complete, documenting the domain, function maps, and predicate maps. The domain, function maps, and predicate maps may be specified by explicit ground lists (of mappings), or by any clear, terminating algorithm.

In addition to the ranking data, other measures are made and presented in the results:

- The *state-of-the-art contribution* (SotAC) quantifies the unique abilities of each system. For each problem solved by a system, its SotAC for the problem is the fraction of systems that do not solve the problem. A system's overall SotAC is the average SotAC over the problems it solves but that are not solved by all the systems.
- The *efficiency* measure combines the number of problems solved with the time taken. It is the average solution rate over the problems solved multiplied by the fraction of problems solved (the solution rate for one problem is the reciprocal of the time taken to solve it). Efficiency is computed for both CPU time and wall clock time, to measure how efficiently the systems use one core and multiple cores respectively.
- The *core usage* measures the extent to which the systems take advantage of multiple cores. The raw core usage is the ratio of CPU time to wall clock time used, and the average core usage is over problems solved with a raw core usage greater than 1.0. The number of problems solved with a lesser raw core usage is also reported, because the ability to solve problems quickly before multi-core search is necessary is also of interest. The competition ran on octa-core computers, thus the maximal core usage was 8.0.

At some time after the competition all high ranking systems in each division are tested over the entire TPTP. This provides a final check for soundness. If a system is found to be unsound during or after the competition, but before the competition report is published, and it cannot be shown that the unsoundness did not manifest itself in the competition, then the system is retrospectively disqualified. At some time after the competition, the solutions from the winners are checked by the panel. If any of the solutions are unacceptable, i.e., they are sufficiently worse than the samples provided, then that system is retrospectively disqualified. All disqualifications are explained in the competition report.

The demonstration division results are presented along with the competition divisions' results, but might not be comparable with those results. The demonstration division is not ranked.

6 System Registration

ATP systems must be registered for the competition using the CASC system registration form, by the registration deadline. For each system an entrant must be nominated to handle all issues (e.g., installation and execution difficulties) arising before, during, and after the competition. The nominated entrant must formally register for CASC. It is not necessary for entrants to physically attend the competition.

6.1 System Descriptions

A system description has to be provided for each ATP system, using the HTML schema supplied on the CASC web site. The schema has the following sections:

- Architecture. This section introduces the ATP system, and describes the calculus and inference rules used.
- Strategies. This section describes the search strategies used, why they are effective, and how they are selected for given problems. Any strategy tuning that is based on specific problems' characteristics must be clearly described (and justified in light of the tuning restrictions described in Section 7).
- Implementation. This section describes the implementation of the ATP system, including the programming language used, important internal data structures, and any special code libraries used. The availability of the system is also given here.
- Expected competition performance. This section makes some predictions about the performance of the ATP system for each of the divisions and categories in which it is competing.
- References.

The system description has to be emailed to the competition organizers by the system description deadline. The system descriptions form part of the competition proceedings (Section 8).

6.2 Sample Solutions

For systems in the divisions that require solution output, representative sample solutions must be emailed to the competition organizers by the sample solutions deadline. Use of the TPTP format for proofs and finite interpretations is encouraged. The competition panel decides whether or not each system's solutions are acceptable (see Section 5).

Proof/model samples are required as follows:

- THF and SLH: SET014^4
- TFA: DAT013=1
- TFN: HWV042.1 and HWV042.3 (but as these problems are rather difficult, some systems were given two easier problems from TPTP v8.2.0: COM001.10 and DAT335.2).
- FOF and ICU : SEU140+2
- UEQ: B00001-1

An explanation must be provided for any non-obvious features.

7 System Requirements

Entrants must ensure that their systems execute in the competition environment, and have the following properties. Entrants are advised to finalize their installation packages and check these properties well in advance of the system delivery deadline. This gives the competition organizers time to help resolve any difficulties encountered.

Execution, Soundness, and Completeness

- Systems must be fully automatic, i.e., all command line switches have to be the same for all problems in each division.
- Systems' performances must be reproducible by running the system again.
- Systems must be sound.
- Systems do not have to be complete in any sense, including calculus, search control, implementation, or resource requirements.
- All techniques used must be general purpose, and expected to extend usefully to new unseen problems. The precomputation and storage of information about individual problems that might appear in the competition, or their solutions, is not allowed. Strategies and strategy selection based on individual problems or their solutions are not allowed. If machine learning procedures are used to tune a system, the learning must ensure that sufficient generalization is obtained so that there is no specialization to individual problems. The system description must explain any such tuning or training that has been done. The competition panel may disqualify any system that is deemed to be problem specific rather than general purpose.

Output

- All solution output must be to `stdout`.
- For each problem, the system must output a distinguished string indicating what solution has been found or that no conclusion has been reached. Systems must use the SZS ontology and standards [88] for this. For example

```
SZS status Theorem for SYN075+1
```

or

```
SZS status GaveUp for SYN075+1
```

- When outputting a solution, the start and end of the solution must be delimited by distinguished strings. Systems must use the SZS ontology and standards for this. For example

```
SZS output start CNFRefutation for SYN075-1.p
```

```
...
```

```
SZS output end CNFRefutation for SYN075-1.p
```

The string specifying the problem status must be output before the start of a solution. Use of the TPTP format for proofs and finite interpretations [123] is encouraged.

- Solutions may not have irrelevant output (e.g., from other threads running in parallel) interleaved in the solution.

Resource Usage

- Systems must be interruptible by a `SIGXCPU` signal so that CPU time limits can be imposed, and interruptible by a `SIGALRM` signal so that wall clock time limits can be imposed. For systems that create multiple processes the signal is sent first to the process at the top of the hierarchy, then one second later to all processes in the hierarchy. The default action on receiving these signals is to exit (thus complying with the time limit, as required), but systems may catch the signals and exit of their own accord. If a system runs past a time limit this is noticed in the timing data, and the system is considered to have not solved the problem.
- If a system terminates of its own accord it may not leave any temporary or intermediate output files. If a system is terminated by a `SIGXCPU` or `SIGALRM` it may not leave any temporary or intermediate files anywhere other than in `/tmp`.
- For practical reasons excessive output from an ATP system is not allowed. A limit, dependent on the disk space available, is imposed on the amount of output that can be produced.

8 The ATP Systems

These system descriptions were written by the entrants.

8.1 Connect++ 0.6.0

Sean Holden
University of Cambridge, United Kingdom

Architecture

Connect++ Version 0.6.0 is the current publicly-available release of the Connect++ prover for first-order logic, introduced in [32]. It is a connection prover using the same calculus and inference rules as leanCoP [46]. That is, it uses the connection calculus with regularity and (optionally) with lemmas.

Strategies

The (default) proof search is essentially the same as that used by leanCoP Version 2.1. That is, it employs a search trying left branches of extensions first, with restricted backtracking as described in [46], and using the leftmost literal of the relevant clause for all inference rules. It does not attempt to exactly reproduce leanCoP's search order. When not using a schedule, command-line options allow many of these choices to be altered; for example allowing backtracking restriction for extensions while still fully exploring left branches, or other modification of the backtracking restrictions. If run with its default schedule it uses one similar to that of leanCoP Version 2.1, including the various applications of definitional clause conversion. Alternatively it can read and apply arbitrary schedules if desired. At present Connect++ does not attempt to tune its proof search based on the characteristics of individual problems.

Implementation

Connect++ is implemented in C++ - minimally the 2017 standard - and built using cmake. Libraries from the Boost collection are used for parsing, hashing, some random number generation, and processing of command-line options. The system has a built-in proof checker for verifying its own output, but also includes a standalone checker implemented in SWI Prolog.

As substitutions need to apply to an entire proof tree the system only represents each variable once and shares the representation, simultaneously maintaining a stack of substitutions making removal of substitutions under backtracking trivial. It also creates subterms only once and shares them; these are indexed allowing constant-time lookup, and nothing is ever removed from the index, meaning that if a term is constructed again after its initial construction no new memory allocation takes place and the term itself is obtained in constant time. At the same time, fresh copies of variables are recycled under backtracking - these two design choices appear to interact very effectively, as the recycling of the variables seems to make it quite likely that subterms already in the index can be reused.

By default a standard recursive unification algorithm is used, but a polynomial-time version is optional.

If a schedule is used, it is assumed that different approaches to definitional clause conversion may be needed - typically all clauses, conjecture clauses only, or no clauses. As these choices

can lead to different matrices, and the conversion itself can be expensive, the system stores and switches between the different matrices rather than converting multiple times.

As the system was developed with two guiding aims - to provide a clear implementation easily modified by others, somewhat in the spirit of MiniSAT [22], and to support experiments in machine learning for guiding the proof search - the implementation avoids the use of direct recursion in favour of a pair of stacks and an iterative implementation based on these, as described in [32]. This allows complete and arbitrary control of backtracking restriction and other modifications to the proof search using typically quite simple modifications to the code.

The source and documentation are available at

<http://www.cl.cam.ac.uk/~sbh11/connect++.html>

Expected Competition Performance

The system remains at an early stage of development and is currently undergoing systematic profiling and improvement. It is not expected at this stage to be competitive.

8.2 CSE 1.7

Feng Cao
JiangXi University of Science and Technology, China

Architecture

CSE 1.7 is a developed prover based on the last version - CSE 1.6. It is an automated theorem prover for first-order logic without equality, based mainly on a novel inference mechanism called Contradiction Separation Based Dynamic Multi-Clause Synergized Automated Deduction (S-CS) [146]. S-CS is able to handle multiple (two or more) clauses dynamically in a synergized way in one deduction step, while binary resolution is a special case. CSE 1.7 also adopts conventional factoring, equality resolution (ER rule), and variable renaming. Some pre-processing techniques, including pure literal deletion and simplification based on the distance to the goal clause, and a number of standard redundancy criteria for pruning the search space: tautology deletion, subsumption (forward and backward), are applied as well. CSE 1.7 has been improved compared with CSE 1.6, mainly from the following aspects:

1. A multi-clause contradiction separation deduction algorithm based on unit clauses has been proposed, which can effectively enhance the reasoning ability of unit clauses for multi-clause deduction.
2. A multi-clause synergized deduction algorithm with full use of synergized clauses is proposed, which can make the unify substitution of candidate literals involved in deduction from simple to complex, thus effectively optimizing the multi-clause deduction search path.

Internally CSE 1.7 works only with clausal normal form. The E prover [65] is adopted with thanks for clausification of full first-order logic problems during preprocessing.

Strategies

CSE 1.7 inherited most of the strategies in CSE 1.6. The main new strategies are:

- By analyzing clause influence degree, the number of remaining literals, and literal deduction ability in the process of multi-clause dynamic deduction, a clause evaluation method based on the clause comprehensive weight is proposed, which can effectively control the number of literals in the contradiction separation clause.
- By analyzing the relationship and differences between variables, functions, and constants, a stability based on clause measurement and calculation method is proposed, which can effectively evaluate clauses with different term structures. The lower the stability of clauses, the easier it is to participate in deduction.

Implementation

CSE 1.7 is implemented mainly in C++, and Java is used for batch problem running implementation. A shared data structure is used for constants and shared variables storage. In addition, special data structure is designed for property description of clause, literal and term, so that it can support the multiple strategy mode. E prover is used for clausification of FOF problems, and then TPTP4X is applied to convert the CNF format into TPTP format.

Expected Competition Performance

CSE 1.7 has made some improvements compared to CSE 1.6, and so we expect a better performance in this year's competition.

Acknowledgement: Development of CSE 1.7 has been partially supported by the General Research Project of Jiangxi Education Department (Grant No. GJJ200818).

8.3 CSE_E 1.6

Peiyao Liu
Southwest Jiaotong University, China

Architecture

CSE_E 1.6 is an automated theorem prover for first-order logic by combining CSE 1.6 and E 3.1, where CSE 1.6 is based on the Contradiction Separation Based Dynamic Multi-Clause Synergized Automated Deduction (S-CS) [146] and E is mainly based on superposition. The combination mechanism is like this: E and CSE are applied to the given problem sequentially. If either prover solves the problem, then the proof process completes. If neither CSE nor E can solve the problem, some inferred clauses with no more than two literals, especially unit clauses, by CSE will be fed to E as lemmas, along with the original clauses, for further proof search.

This kind of combination is expected to take advantage of both CSE and E, and produce a better performance. Concretely, CSE is able to generate a good number of unit clauses, based on the fact that unit clauses are helpful for proof search and equality handling. On the other hand, E has a good ability on equality handling.

Strategies

The CSE part of CSE_E 1.6 takes almost the same strategies as in that in CSE 1.6 standalone, e.g., clause/literal selection, strategy selection, and CSC strategy. The only difference is that equality handling strategies of CSE part of the combined system are blocked. The main new strategies for the combined systems are:

- Lemma filtering mainly based on deduction weight of binary clauses.
- Fine-grained dynamic time allocation scheme in different run stages.

Implementation

CSE_E 1.6 is implemented mainly in C++, and Java is used for batch problem running implementation. The job dispatch between CSE and E is implemented in C++.

Expected Competition Performance

We expect CSE_E 1.6 to solve some hard problems that E cannot solve and have a satisfying performance.

Acknowledgement: Development of CSE_E 1.6 has been partially supported by the National Natural Science Foundation of China (NSFC) (Grant No. 61976130). Stephan Schulz for his kind permission on using his E prover that makes CSE_E possible.

8.4 CSG_E 1.0

Peiyao Liu
Southwest Jiaotong University, China

Architecture

CSG_E 1.0 is an automated theorem prover for first-order logic by combining CSG 1.0 and E 3.1, where CSG 1.0 is based on the Contradiction Separation Based Dynamic Multi-Clause Synergized Automated Deduction (S-CS) [146] and E is mainly based on superposition. CSG uses a new deduction calculus based on S-CS rule called gridle construction method. The main idea of the new method is to select N literals to construct a maximal standard contradiction (called a gridle), then use the gridle to match the clauses until the gridle is filled, and finally the literals that cannot be matched are composed into the new inferred clause. The combination mechanism is like this: E and CSG are applied to the given problem sequentially. If either prover solves the problem, then the proof process completes. If neither CSG nor E can solve the problem, some inferred clauses with no more than two literals, especially unit clauses, by CSG will be fed to E as lemmas, along with the original clauses, for further proof search. This kind of combination is expected to take advantage of both CSE and E, and produce a better performance. Concretely, CSE is able to generate a good number of unit clauses, based on the fact that unit clauses are helpful for proof search and equality handling. On the other hand, E has a good ability on equality handling.

Strategies

The CSG part of CSG_E 1.0 takes almost the new strategies as in that in CSG 1.0 standalone, e.g., clause/literal selection, gridle generation literal strategy, gridle-size adjustment strategy and CSC strategy. The only difference is that equality handling strategies of CSG part of the combined system are blocked.

Implementation

CSG_E 1.0 is implemented mainly in C++, and Java is used for batch problem running implementation. The job dispatch between CSG and E is implemented in C++.

Expected Competition Performance

We expect CSG_E 1.0 to solve some hard problems that E cannot solve and have a satisfying performance.

Acknowledgement: Development of CSG_E 1.0 has been partially supported by the National Natural Science Foundation of China (NSFC) (Grant No. 61976130). Stephan Schulz for his kind permission on using his E prover that makes CSE_E possible.

8.5 CSLE 1.0

Guoyan Zeng
Southwest Jiaotong University, China

Architecture

CSLE 1.0 is an automated theorem prover for first-order logic, combining CSI 1.0 and E, where CSI 1.0 is a multi-layer inverse and parallel prover based on the Contradiction Separation Based Dynamic Multi-Clause Synergized Automated Deduction (S-CS) [146] and E is mainly based on superposition. The combination mechanism is like this: E and CSI are applied to the given problem sequentially. If either prover solves the problem, then the proof process completes. If neither CSI nor E can solve the problem, some inferred clauses with no more than two literals, especially unit clauses, by CSI will be fed to E as lemmas, along with the original clauses, for further proof search. This kind of combination is expected to take advantage of both CSI and E, and produce a better performance. Concretely, CSI is able to generate a good number of unit clauses, based on the fact that unit clauses are helpful for proof search and equality handling. On the other hand, E has a good ability on equality handling.

Strategies

The CSI part of CSLE 1.0 takes almost the same strategies as in that in CSE 1.6 standalone, e.g., clause/literal selection, strategy selection, and CSC strategy. The only difference is that equality handling strategies of CSE part of the combined system are blocked. The main new strategies for the combined systems are:

- Complementary ratio strategy. A measure and calculation method of complementary relation between two clauses, which can guide the selection of clauses to participate in deduction and plan the deduction path effectively.

- Portfolio strategy clause selection schemes for different run stages.

Implementation

CSLE 1.0 is implemented mainly in C++, and Java is used for batch problem running implementation. The job dispatch between CSI and E is implemented in C++.

Expected Competition Performance

We expect CSLE 1.0 to solve some hard problems that E cannot solve and have a satisfying performance.

Acknowledgement: Development of CSE_E 1.6 has been partially supported by the National Natural Science Foundation of China (NSFC) (Grant No. 62106206, 62206227). Stephan Schulz for his kind permission on using his E prover that makes CSE_E possible.

8.6 cvc5 1.1.3

Andy Reynolds
University of Iowa, USA

Architecture

cvc5 [3] is the successor of CVC4 [6]. It is an SMT solver based on the CDCL(T) architecture [44] that includes built-in support for many theories, including linear arithmetic, arrays, bit vectors, datatypes, finite sets and strings. It incorporates approaches for handling universally quantified formulas. For problems involving free function and predicate symbols, cvc5 primarily uses heuristic approaches based on conflict-based instantiation and E-matching for theorems, and finite model finding approaches for non-theorems.

Like other SMT solvers, cvc5 treats quantified formulas using a two-tiered approach. First, quantified formulas are replaced by fresh Boolean predicates and the ground theory solver(s) are used in conjunction with the underlying SAT solver to determine satisfiability. If the problem is unsatisfiable at the ground level, then the solver answers “unsatisfiable”. Otherwise, the quantifier instantiation module is invoked, and will either add instances of quantified formulas to the problem, answer “satisfiable”, or return unknown. Finite model finding in cvc5 targets problems containing background theories whose quantification is limited to finite and uninterpreted sorts. In finite model finding mode, cvc5 uses a ground theory of finite cardinality constraints that minimizes the number of ground equivalence classes, as described in [57]. When the problem is satisfiable at the ground level, a candidate model is constructed that contains complete interpretations for all predicate and function symbols. It then adds instances of quantified formulas that are in conflict with the candidate model, as described in [57]. If no instances are added, it reports “satisfiable”.

cvc5 has native support for problems in higher-order logic, as described in [5]. It uses a pragmatic approach for HOL, where lambdas are eliminated eagerly via lambda lifting. The approach extends the theory solver for quantifier-free uninterpreted functions (UF) and E-matching. For the former, the theory solver for UF in cvc5 now handles equalities between functions using an extensionality inference. Partial applications of functions are handled using a (lazy) applicative encoding where some function applications are equated to the applicative encoding. For the latter, several of the data structures for E-matching have been modified to

incorporate matching in the presence of equalities between functions, function variables, and partial function applications.

Strategies

For handling theorems, `cvc5` primarily uses conflict-based quantifier instantiation [56, 4], enumerative instantiation [54] and E-matching. `cvc5` uses a handful of orthogonal trigger selection strategies for E-matching, and several orthogonal ordering heuristics for enumerative instantiation. For handling non-theorems, `cvc5` primarily uses finite model finding techniques. Since `cvc5` with finite model finding is also capable of establishing unsatisfiability, it is used as a strategy for theorems as well.

Implementation

`cvc5` is implemented in C++. The code is available from

<https://github.com/cvc5/cvc5>

Expected Competition Performance

The performance of `cvc5` will be comparable to last year. We continue to rely on a conversion from TPTP to `smt2` as a preprocess step. This year, we have added various new strategies to TFN and THF, including a new instantiation strategy that combines model-based quantifier instantiation [24] with enumerative syntax-guided synthesis [55].

8.7 Drodi 3.6.0

Oscar Contreras
Amateur Programmer, Spain

Architecture

Drodi 3.6.0 is a very basic and lightweight automated theorem prover. It implements the following main features:

- Ordered resolution and equality paramodulation inferences as well as demodulation and some other standard simplifications.
- A basic implementations of clausal normal form conversion as in [45].
- AVATAR architecture with a SAT solver [140].
- Limited Resource Strategy [58].
- Discrimination trees.
- KBO, non recursive and lexicographic reduction orderings.
- Improved literal selection including lookahead as in [28].
- SinE distance for clauses and symbols as in [70].
- Layered clause selection as in [25].
- Stochastic strategy inspired in [71].
- Drodi produces a (hopefully) verifiable proof in TPTP format.

Strategies

Drodi has a fair number of selectable strategies including but not limited to the following:

- Otter, Discount and Limited Resource Strategy [58] saturation algorithms.
- A basic implementation of AVATAR architecture [140].
- Several literal and term reduction orderings.
- Several literal selection options [28].
- Several layered clause selection heuristics with adjustable selection ratios [25].
- Classical clause relevancy pruning.
- Some strategies are run a second time with a previously applied randomization to the problem [71].
- Drodi can generate a learning file from successful proofs and use the file to guide clause selection strategy. It is based in the enhanced ENIGMA method. However, unlike ENIGMA, the learning file is completely general and can be used with any kind of problems. This generality allows the use of the same learning file in both FOF and UEQ CASC competition divisions. The learning file is generated over a set of TPTP problems before the CASC competition using built-in Drodi functions that include a L2 Support Vector Machine. Drodi integrated learning functions are a generalization of ENIGMA [33, 34]. Literals polarity, equality, skolem and variable occurrences are stored in clause feature vectors. Unlike ENIGMA, instead of storing the specific functions and predicates themselves only the SinE distance and arity of functions and non equality predicates are stored in clause feature vectors with different features assigned to predicates and functions.

Implementation

Drodi is implemented in C. It includes discrimination trees and hashing indexing. All the code is original, without special code libraries or code taken from other sources.

Expected Competition Performance

Drodi 3.6.0 is basically the same solver as last year with only minor improvements so performance will be similar to CASC-29.

8.8 E 3.1

Stephan Schulz
DHBW Stuttgart, Germany

Architecture

E [61, 65, 66] is a purely equational theorem prover for many-sorted first-order logic with equality, and for monomorphic higher-order logic. It consists of an (optional) clausifier for pre-processing full first-order formulae into clausal form, and a saturation algorithm implementing an instance of the superposition calculus with negative literal selection and a number of redundancy elimination techniques, optionally with higher-order extensions [141, 143]. E is based on the DISCOUNT-loop variant of the given-clause algorithm, i.e., a strict separation of active

and passive facts. No special rules for non-equational literals have been implemented. Resolution is effectively simulated by paramodulation and equality resolution. As of E 2.1, PicoSAT [13] can be used to periodically check the (on-the-fly grounded) proof state for propositional unsatisfiability.

Strategies

Proof search in E is primarily controlled by a literal selection strategy, a clause selection heuristic, and a simplification ordering. The prover supports a large number of pre-programmed literal selection strategies. Clause selection heuristics can be constructed on the fly by combining various parameterized primitive evaluation functions, or can be selected from a set of predefined heuristics. Clause evaluation heuristics are based on symbol-counting, but also take other clause properties into account. In particular, the search can prefer clauses from the set of support, or containing many symbols also present in the goal. Supported term orderings are several parameterized instances of Knuth-Bendix-Ordering (KBO) and Lexicographic Path Ordering (LPO), which can be lifted in different ways to literal orderings.

For CASC-29, E implements a two-stage multi-core strategy-scheduling automatic mode. The total CPU time available is broken into several (unequal) time slices. For each time slice, the problem is classified into one of several classes, based on a number of simple features (number of clauses, maximal symbol arity, presence of equality, presence of non-unit and non-Horn clauses, possibly presence of certain axiom patterns...). For each class, a schedule of strategies is greedily constructed from experimental data as follows: The first strategy assigned to a schedule is the one that solves the most problems from this class in the first time slice. Each subsequent strategy is selected based on the number of solutions on problems not already solved by a preceding strategy. The strategies are then scheduled onto the available cores and run in parallel.

About 140 different strategies have been thoroughly evaluated on all untyped first-order problems from TPTP 7.3.0. We have also explored some parts of the heuristic parameter space with a short time limit of 5 seconds. This allowed us to test about 650 strategies on all TPTP problems, and an extra 7000 strategies on UEQ problems from TPTP 7.2.0. About 100 of these strategies are used in the automatic mode, and about 450 are used in at least one schedule.

Implementation

E is build around perfectly shared terms, i.e. each distinct term is only represented once in a term bank. The whole set of terms thus consists of a number of interconnected directed acyclic graphs. Term memory is managed by a simple mark-and-sweep garbage collector. Unconditional (forward) rewriting using unit clauses is implemented using perfect discrimination trees with size and age constraints. Whenever a possible simplification is detected, it is added as a rewrite link in the term bank. As a result, not only terms, but also rewrite steps are shared. Subsumption and contextual literal cutting (also known as subsumption resolution) is supported using feature vector indexing [64]. Superposition and backward rewriting use fingerprint indexing [63], a new technique combining ideas from feature vector indexing and path indexing. Finally, LPO and KBO are implemented using the elegant and efficient algorithms developed by Bernd Löchner in [39, 40]. The prover and additional information are available at

<https://www.eprover.org>

Expected Competition Performance

E 3.1 is the CASC-29 SLH winner.

8.9 E 3.2.0

Stephan Schulz
DHBW Stuttgart, Germany

Architecture

E [61, 65, 66] is a purely equational theorem prover for many-sorted first-order logic with equality, and for monomorphic higher-order logic. It consists of an (optional) clausifier for pre-processing full first-order formulae into clausal form, and a saturation algorithm implementing an instance of the superposition calculus with negative literal selection and a number of redundancy elimination techniques, optionally with higher-order extensions [141, 143]. E is based on the DISCOUNT-loop variant of the given-clause algorithm, i.e., a strict separation of active and passive facts. No special rules for non-equational literals have been implemented. Resolution is effectively simulated by paramodulation and equality resolution. As of E 2.1, PicoSAT [13] can be used to periodically check the (on-the-fly grounded) proof state for propositional unsatisfiability.

Strategies

Proof search in E is primarily controlled by a literal selection strategy, a clause selection heuristic, and a simplification ordering. The prover supports a large number of pre-programmed literal selection strategies. Clause selection heuristics can be constructed on the fly by combining various parameterized primitive evaluation functions, or can be selected from a set of predefined heuristics. Clause evaluation heuristics are based on symbol-counting, but also take other clause properties into account. In particular, the search can prefer clauses from the set of support, or containing many symbols also present in the goal. Supported term orderings are several parameterized instances of Knuth-Bendix-Ordering (KBO) and Lexicographic Path Ordering (LPO), which can be lifted in different ways to literal orderings.

For CASC-J12, E implements a two-stage multi-core strategy-scheduling automatic mode. The total CPU time available is broken into several (unequal) time slices. For each time slice, the problem is classified into one of several classes, based on a number of simple features (number of clauses, maximal symbol arity, presence of equality, presence of non-unit and non-Horn clauses, possibly presence of certain axiom patterns, ...). For each class, a schedule of strategies is greedily constructed from experimental data as follows: The first strategy assigned to a schedule is the one that solves the most problems from this class in the first time slice. Each subsequent strategy is selected based on the number of solutions on problems not already solved by a preceding strategy. The strategies are then scheduled onto the available cores and run in parallel.

About 140 different strategies have been thoroughly evaluated on all untyped first-order problems from TPTP 7.3.0. We have also explored some parts of the heuristic parameter space with a short time limit of 5 seconds. This allowed us to test about 650 strategies on all TPTP problems, and an extra 7000 strategies on UEQ problems from TPTP 7.2.0. About 100 of these strategies are used in the automatic mode, and about 450 are used in at least one schedule.

Implementation

E is build around perfectly shared terms, i.e., each distinct term is only represented once in a term bank. The whole set of terms thus consists of a number of interconnected directed acyclic graphs. Term memory is managed by a simple mark-and-sweep garbage collector. Unconditional (forward) rewriting using unit clauses is implemented using perfect discrimination trees with size and age constraints. Whenever a possible simplification is detected, it is added as a rewrite link in the term bank. As a result, not only terms, but also rewrite steps are shared. Subsumption and contextual literal cutting (also known as subsumption resolution) is supported using feature vector indexing [64]. Superposition and backward rewriting use fingerprint indexing [63], a new technique combining ideas from feature vector indexing and path indexing. Finally, LPO and KBO are implemented using the elegant and efficient algorithms developed by Bernd Löchner in [39, 40]. The prover and additional information are available at

<https://www.eprover.org>

Expected Competition Performance

E 3.2.0 is basically E 3.1 made more robust and somewhat more efficient. We have not yet been able to evaluate and integrate new search strategies making full use of all new features. As a result, we expect performance to be similar to last year’s version. The system is expected to perform well in most proof classes, but will at best complement top systems in the disproof classes.

8.10 GKC 0.8

Tanel Tammet
Tallinn University of Technology, Estonia

Architecture

GKC [135] is a resolution prover optimized for search in large knowledge bases. The GKC version 0.8 running at CASC-29 is a marginally improved version of the GKC 0.7 running in two previous CASCs. Almost all of the GKC development effort this year has gone to the commonsense superstructure GK and the natural language reasoning pipeline [138].

GKC is used as a foundation (GK Core) for building a common-sense reasoner GK. In particular, GK can handle inconsistencies and perform probabilistic and nonmonotonic reasoning [136, 137].

The WASM version of the previous GKC 0.6 is used as the prover engine in the educational system. It can read and output proofs in the TPTP, simplified TPTP and JSON format, the latter compatible with JSON-LD [139].

GKC only looks for proofs and does not try to show non-provability. These standard inference rules have been implemented in GKC:

- Binary resolution with optionally the set of support strategy, negative or positive ordered resolution or unit restriction.
- Hyperresolution.
- Factorization.
- Paramodulation and demodulation with the Knuth-Bendix ordering.

GKC includes an experimental implementation of propositional inferencing and instance generation, which we do not plan to use during the current CASC.

Strategies

GKC uses multiple strategies run sequentially, with the time limit starting at 0.1 seconds for each, increased 10 or 5 times once the whole batch has been performed. The strategy selections takes into consideration the basic properties of the problem: the presence of equality and the approximate size of the problem.

We perform the selection of a given clause by using several queues in order to spread the selection relatively uniformly over these categories of derived clauses and their descendants: axioms, external axioms, assumptions and goals. The queues are organized in two layers. As a first layer we use the common ratio-based algorithm of alternating between selecting n clauses from a weight-ordered queue and one clause from the FIFO queue with the derivation order. As a second layer we use four separate queues based on the derivation history of a clause. Each queue in the second layer contains the two sub-queues of the first layer.

Implementation

GKC is implemented in C. The data representation machinery is built upon a shared memory graph database `WhiteDB` (<http://whitedb.org>), enabling it to solve multiple different queries in parallel processes without a need to repeatedly parse or load the large parsed knowledge base from the disk. An interesting aspect of GKC is the pervasive use of hash indexes, feature vectors and fingerprints, while no tree indexes are used.

GKC can be obtained from

<https://github.com/tammet/gkc/>

Expected Competition Performance

We expect GKC to be in the middle of the final ranking for FOF and below the average in UEQ. We expect GKC to perform well on very large problems.

8.11 iProver 3.9

Konstantin Korovin
University of Manchester, United Kingdom

Architecture

iProver [36, 18] is a theorem prover for quantified first-order logic with theories. iProver interleaves instantiation calculus Inst-Gen [37, 36, 23] with ordered resolution and superposition calculi [18]. iProver approximates first-order clauses using propositional abstractions that are solved using MiniSAT [22] or Z3 [17] and refined using model-guided instantiations. iProver also implements a general abstraction-refinement framework for under- and over-approximations of first-order clauses [26, 27]. First-order clauses are exchanged between calculi during the proof search.

Recent features in iProver include:

- Ground joinability and connectedness in the superposition calculus [20].
- Support for quantified reasoning with arithmetic and arrays.
- AC joinability and AC normalisation [19].
- Superposition calculus with simplifications including: demodulation, light normalisation, subsumption, subsumption resolution and global subsumption. iProver’s simplification set up [18] is tunable via command line options and generalises common architectures such as Discount or Otter.
- HOS-ML framework for learning heuristics using combination of hyper-parameter optimisation and dynamic clustering together with schedule optimisation using constraint solving [31, 30]

Strategies

iProver has around 100 options to control the proof search including options for literal selection, passive clause selection, frequency of calling the SAT/SMT solvers, simplifications, and options for combination of instantiation with resolution and superposition. For the competition HOS-ML [31] was used to build a multi-core schedule from heuristics learnt over a sample of FOF problems. Some theories and fragments are recognised such as EPR, UEQ, Horn, groups, rings and lattices for which options are adapted accordingly.

Implementation

iProver is implemented in OCaml. For the ground reasoning uses MiniSat [22] and Z3 [17]. iProver accepts FOF, TFF and CNF formats. Vampire [38, 50] and E prover [65] are used for proof-producing classification of FOF/TFF problems. Vampire is also used for SInE axiom selection [29] in the LTB division and for theory axioms in the TFA division. iProver is available at:

<https://gitlab.com/korovin/iprover>

Expected Competition Performance

iProver is regularly in the top three in FOF, UEQ, TFN, TFA. We expect an improved performance compared to the previous year due to polishing new methods for equational reasoning, and heuristic optimization.

8.12 LEO-II 1.7.0

Alexander Steen
University of Greifswald, Germany

Architecture

LEO-II [10], the successor of LEO [9], is a higher-order ATP system based on extensional higher-order resolution. More precisely, LEO-II employs a refinement of extensional higher-order RUE resolution [8]. LEO-II is designed to cooperate with specialist systems for fragments of higher-order logic. By default, LEO-II cooperates with the first-order ATP system E [60]. LEO-II is often too weak to find a refutation amongst the steadily growing set of clauses on its

own. However, some of the clauses in LEO-II’s search space attain a special status: they are first-order clauses modulo the application of an appropriate transformation function. Therefore, LEO-II launches a cooperating first-order ATP system every n iterations of its (standard) resolution proof search loop (e.g., 10). If the first-order ATP system finds a refutation, it communicates its success to LEO-II in the standard SZS format. Communication between LEO-II and the cooperating first-order ATP system uses the TPTP language and standards.

Strategies

LEO-II employs an adapted “Otter loop”. Moreover, LEO-II uses some basic strategy scheduling to try different search strategies or flag settings. These search strategies also include some different relevance filters.

Implementation

LEO-II is implemented in OCaml 4, and its problem representation language is the TPTP THF language [11]. In fact, the development of LEO-II has largely paralleled the development of the TPTP THF language and related infrastructure [119]. LEO-II’s parser supports the TPTP THF0 language and also the TPTP languages FOF and CNF.

Unfortunately the LEO-II system still uses only a very simple sequential collaboration model with first-order ATPs instead of using the more advanced, concurrent and resource-adaptive OANTS architecture [12] as exploited by its predecessor LEO.

The LEO-II system is distributed under a BSD style license, and it is available from

<http://www.leoprover.org>

Expected Competition Performance

LEO-II is not actively being developed anymore, hence there are no expected improvements to last year’s CASC results.

8.13 Leo-III 1.7.15

Alexander Steen
University of Greifswald, Germany

Architecture

Leo-III [68], the successor of LEO-II [10], is a higher-order ATP system based on extensional higher-order paramodulation with inference restrictions using a higher-order term ordering. The calculus contains dedicated extensionality rules and is augmented with equational simplification routines that have their intellectual roots in first-order superposition-based theorem proving. The saturation algorithm is a variant of the given clause loop procedure inspired by the first-order ATP system E.

Leo-III cooperates with external first-order ATPs that are called asynchronously during proof search; a focus is on cooperation with systems that support typed first-order (TFF) input. For this year’s CASC E [61, 65] is used as external system. However, cooperation is in general not limited to first-order systems. Further TPTP/TSTP-compliant external systems (such as

higher-order ATPs or counter model generators) may be included using simple command-line arguments. If the saturation procedure loop (or one of the external provers) finds a proof, the system stops, generates the proof certificate and returns the result.

Strategies

Leo-III comes with several configuration parameters that influence its proof search by applying different heuristics and/or restricting inferences. These parameters can be chosen manually by the user on start-up. Leo-III implements a very naive time slicing approach in which at most three manually fixed parameter configurations are used, one after each other. In practice, this hardly ever happens and Leo-III will just run with its default parameter setting.

Implementation

Leo-III utilizes and instantiates the associated LeoPARD system platform [145] for higher-order (HO) deduction systems implemented in Scala (currently using Scala 2.13 and running on a JVM with Java ≥ 8). The prover makes use of LeoPARD's data structures and implements its own reasoning logic on top. A hand-crafted parser is provided that supports all TPTP syntax dialects. It converts its produced concrete syntax tree to an internal TPTP AST data structure which is then transformed into polymorphically typed lambda terms. As of version 1.1, Leo-III supports all common TPTP dialects (CNF, FOF, TFF, THF) as well as their polymorphic variants [15, 35]. Since version 1.6.X ($X \geq 0$) Leo-III also accepts non-classical problem input represented in non-classical TPTP.

The term data structure of Leo-III uses a polymorphically typed spine term representation augmented with explicit substitutions and De Bruijn-indices. Furthermore, terms are perfectly shared during proof search, permitting constant-time equality checks between alpha-equivalent terms.

Leo-III's saturation procedure may at any point invoke external reasoning tools. To that end, Leo-III includes an encoding module which translates (polymorphic) higher-order clauses to polymorphic and monomorphic typed first-order clauses, whichever is supported by the external system. While LEO-II relied on cooperation with untyped first-order provers, Leo-III exploits the native type support in first-order provers (TFF logic) for removing clutter during translation and, in turn, higher effectivity of external cooperation.

Leo-III is available on GitHub:

<https://github.com/leoprover/Leo-III>

Expected Competition Performance

Version 1.7.15 is, for all intents and purposes of CASC, equivalent to the version from last year except that some minor bugs have been fixed, and the support for reasoning in various quantified non-classical logics (not relevant to CASC) was improved. We do not expect Leo-III to be strongly competitive against more recent higher-order provers as Leo-III does not implement several standard features of effective systems (including time slicing and proper axiom selection).

8.14 Prover9 1109a

Bob Veroff on behalf of William McCune
University of New Mexico, USA

Architecture

Prover9, Version 2009-11A, is a resolution/paramodulation prover for first-order logic with equality. Its overall architecture is very similar to that of Otter-3.3 [43]. It uses the “given clause algorithm”, in which not-yet-given clauses are available for rewriting and for other inference operations (sometimes called the “Otter loop”).

Prover9 has available positive ordered (and nonordered) resolution and paramodulation, negative ordered (and nonordered) resolution, factoring, positive and negative hyperresolution, UR-resolution, and demodulation (term rewriting). Terms can be ordered with LPO, RPO, or KBO. Selection of the “given clause” is by an age-weight ratio.

Proofs can be given at two levels of detail: (1) standard, in which each line of the proof is a stored clause with detailed justification, and (2) expanded, with a separate line for each operation. When FOF problems are input, proof of transformation to clauses is not given.

Completeness is not guaranteed, so termination does not indicate satisfiability.

Strategies

Prover9 has available many strategies; the following statements apply to CASC.

Given a problem, Prover9 adjusts its inference rules and strategy according to syntactic properties of the input clauses such as the presence of equality and non-Horn clauses. Prover9 also does some preprocessing, for example, to eliminate predicates.

For CASC Prover9 uses KBO to order terms for demodulation and for the inference rules, with a simple rule for determining symbol precedence.

For the FOF problems, a preprocessing step attempts to reduce the problem to independent subproblems by a miniscope transformation; if the problem reduction succeeds, each subproblem is classified and given to the ordinary search procedure; if the problem reduction fails, the original problem is classified and given to the search procedure.

Implementation

Prover9 is coded in C, and it uses the LADR libraries. Some of the code descended from EQP [42]. (LADR has some AC functions, but Prover9 does not use them). Term data structures are not shared (as they are in Otter). Term indexing is used extensively, with discrimination tree indexing for finding rewrite rules and subsuming units, FPA/Path indexing for finding subsumed units, rewritable terms, and resolvable literals. Feature vector indexing [62] is used for forward and backward nonunit subsumption. Prover9 is available from

<http://www.cs.unm.edu/~mccune/prover9/>

Expected Competition Performance

Prover9 is the CASC fixed point, against which progress can be judged. Each year it is expected do worse than the previous year, relative to the other systems.

8.15 Twee 2.4.2

Nick Smallbone
Chalmers University of Technology, Sweden

Architecture

Twee 2.4.2 [67] is a theorem prover for unit equality problems based on unfailing completion [2]. It implements a DISCOUNT loop, where the active set contains rewrite rules (and unorientable equations) and the passive set contains critical pairs. The basic calculus is not goal-directed, but Twee implements a transformation which improves goal direction for many problems.

Twee features ground joinability testing [41] and a connectedness test [1], which together eliminate many redundant inferences in the presence of unorientable equations. The ground joinability test performs case splits on the order of variables, in the style of [41], and discharges individual cases by rewriting modulo a variable ordering.

Strategies

Twee's strategy is simple and it does not tune its heuristics or strategy based on the input problem. The term ordering is always KBO; by default, functions are ordered by number of occurrences and have weight 1. The proof loop repeats the following steps:

- Select and normalise the lowest-scored critical pair, and if it is not redundant, add it as a rewrite rule to the active set.
- Normalise the active rules with respect to each other.
- Normalise the goal with respect to the active rules.

Each critical pair is scored using a weighted sum of the weight of both of its terms. Terms are treated as DAGs when computing weights, i.e., duplicate subterms are counted only once per term.

For CASC, to take advantage of multiple cores, several versions of Twee run in parallel using different parameters (e.g., with the goal-directed transformation on or off).

Implementation

Twee is written in Haskell. Terms are represented as array-based flatterms for efficient unification and matching. Rewriting uses a perfect discrimination tree.

The passive set is represented compactly (12 bytes per critical pair) by storing only the information needed to reconstruct the critical pair, not the critical pair itself. Because of this, Twee can run for an hour or more without exhausting memory.

Twee uses an LCF-style kernel: all rules in the active set come with a certified proof object which traces back to the input axioms. When a conjecture is proved, the proof object is transformed into a human-readable proof. Proof construction does not harm efficiency because the proof kernel is invoked only when a new rule is accepted. In particular, reasoning about the passive set does not invoke the kernel.

Twee can be downloaded as open source from:

<https://nick8325.github.io/twee/>

Expected Competition Performance

Twee 2.4.2 is the CASC-29 UEQ winner.

8.16 Twee 2.5.0

Nick Smallbone
Chalmers University of Technology, Sweden

Architecture

Twee 2.4.2 [67] is a theorem prover for unit equality problems based on unfailing completion [2]. It implements a DISCOUNT loop, where the active set contains rewrite rules (and unorientable equations) and the passive set contains critical pairs. The basic calculus is not goal-directed, but Twee implements a transformation which improves goal direction for many problems.

Twee features ground joinability testing [41] and a connectedness test [1], which together eliminate many redundant inferences in the presence of unorientable equations. The ground joinability test performs case splits on the order of variables, in the style of [41], and discharges individual cases by rewriting modulo a variable ordering. New this year is a mode which rewrites backwards from the goal instead of enumerating critical pairs, but it is still rather rough.

Strategies

Twee's strategy is simple and it does not tune its heuristics or strategy based on the input problem. The term ordering is always KBO; by default, functions are ordered by number of occurrences and have weight 1. The proof loop repeats the following steps:

- Select and normalise the lowest-scored critical pair, and if it is not redundant, add it as a rewrite rule to the active set.
- Normalise the active rules with respect to each other.
- Normalise the goal with respect to the active rules.

Each critical pair is scored using a weighted sum of the weight of both of its terms. Terms are treated as DAGs when computing weights, i.e., duplicate subterms are counted only once per term.

For CASC, to take advantage of multiple cores, several versions of Twee run in parallel using different parameters (e.g., with the goal-directed transformation on or off).

Implementation

Twee is written in Haskell. Terms are represented as array-based flatterms for efficient unification and matching. Rewriting uses a perfect discrimination tree.

The passive set is represented compactly (12 bytes per critical pair) by storing only the information needed to reconstruct the critical pair, not the critical pair itself. Because of this, Twee can run for an hour or more without exhausting memory.

Twee uses an LCF-style kernel: all rules in the active set come with a certified proof object which traces back to the input axioms. When a conjecture is proved, the proof object is transformed into a human-readable proof. Proof construction does not harm efficiency because

the proof kernel is invoked only when a new rule is accepted. In particular, reasoning about the passive set does not invoke the kernel.

Twee can be downloaded as open source from:

<https://nick8325.github.io/twee/>

Expected Competition Performance

Similar to last year, but we hope the new goal-directed mode might solve a few interesting problems.

8.17 Vampire 4.8

Michael Rawson
TU Wien, Austria

There have been a number of changes and improvements since Vampire 4.7, although it is still the same beast. Most significant from a competition point of view are long-awaited refreshed strategy schedules. As a result, several features present in previous competitions will now come into full force, including new rules for the evaluation and simplification of theory literals. A large number of completely new features and improvements also landed this year: highlights include a significant refactoring of the substitution tree implementation, the arrival of encompassment demodulation to Vampire, and support for parametric datatypes.

Vampire's higher-order support has also been re-implemented from the ground up. The new implementation is still at an early stage and its theoretical underpinnings are being developed. There is currently no documentation of either.

Architecture

Vampire [38] is an automatic theorem prover for first-order logic with extensions to theory-reasoning and higher-order logic. Vampire implements the calculi of ordered binary resolution, and superposition for handling equality. It also implements the Inst-gen calculus and a MACE-style finite model builder [51]. Splitting in resolution-based proof search is controlled by the AVATAR architecture which uses a SAT or SMT solver to make splitting decisions [140, 48]. A number of standard redundancy criteria and simplification techniques are used for pruning the search space: subsumption, tautology deletion, subsumption resolution and rewriting by ordered unit equalities. The reduction ordering is the Knuth-Bendix Ordering. Substitution tree and code tree indexes are used to implement all major operations on sets of terms, literals and clauses. Internally, Vampire works only with clausal normal form. Problems in the full first-order logic syntax are classified during preprocessing [52]. Vampire implements many useful preprocessing transformations including the SinE axiom selection algorithm. When a theorem is proved, the system produces a verifiable proof, which validates both the classification phase and the refutation of the CNF.

Strategies

Vampire 4.8 provides a very large number of options for strategy selection. The most important ones are:

- Choices of saturation algorithm:
 - Limited Resource Strategy [58]

- DISCOUNT loop
- Otter loop
- Instantiation using the Inst-Gen calculus
- MACE-style finite model building with sort inference
- Splitting via AVATAR [140]
- A variety of optional simplifications.
- Parameterized reduction orderings.
- A number of built-in literal selection functions and different modes of comparing literals [28].
- Age-weight ratio that specifies how strongly lighter clauses are preferred for inference selection. This has been extended with a layered clause selection approach [25].
- Set-of-support strategy with extensions for theory reasoning.
- For theory-reasoning:
 - Ground equational reasoning via congruence closure.
 - Addition of theory axioms and evaluation of interpreted functions [49].
 - Use of Z3 with AVATAR to restrict search to ground-theory-consistent splitting branches [48].
 - Specialised theory instantiation and unification [53].
 - Extensionality resolution with detection of extensionality axioms

The schedule for the new HOL implementation was developed using Snake, a strategy schedule construction tool described in more detail last year. The Snake schedule this year fully embraces Vampire randomisation support [71] and, in particular, every strategy independently shuffles the input problem, to nullify (in expectation) the effect of problem scrambling done by the organisers.

Implementation

Vampire 4.8 is implemented in C++. It makes use of fixed versions of Minisat and Z3. See the website for more information and access to the GitHub repository.

Expected Competition Performance

Vampire 4.8 is the CASC-29 THF, TFA, TFN, and FOF winner.

8.18 Vampire 4.9

Michael Rawson
TU Wien, Austria

There have been a number of improvements since Vampire 4.8, although it is still the same beast. For the first time this year, Vampire's schedules were constructed mostly using the Snake strategy selection tool, although a return of the traditional Spider is still possible in future. Improvements from the past year include:

- A runtime-specialised version of unidirectional term ordering checks
- Improvements to unification with abstraction
- Surprising improvements to Vampire's basic routines such as renaming and unification

- A simple interactive mode
- Revitalisation of code trees
- Experimental features not yet fully understood, mostly aimed at unit-equational reasoning.
- Portability: Vampire is much more standards-compliant and portable than previously, with much-reduced dependence on platform-specific APIs and hardware architectures, aided by C++17

Vampire’s higher-order support remains very similar to last year, although a re-implementation intended for mainline Vampire is already underway.

Architecture

Vampire [38] is an automatic theorem prover for first-order logic with extensions to theory-reasoning and higher-order logic. Vampire implements the calculi of ordered binary resolution, and superposition for handling equality. It also implements a MACE-style finite model builder for finding finite counter-examples [51]. Splitting in resolution-based proof search is controlled by the AVATAR architecture which uses a SAT or SMT solver to make splitting decisions [140, 48]. A number of standard redundancy criteria and simplification techniques are used for pruning the search space: subsumption, tautology deletion, subsumption resolution and rewriting by ordered unit equalities. The reduction ordering is the Knuth-Bendix Ordering. Substitution tree and code tree indexes are used to implement all major operations on sets of terms, literals and clauses. Internally, Vampire works only with clausal normal form. Problems in the full first-order logic syntax are clausified during preprocessing [52]. Vampire implements many useful preprocessing transformations including the SInE axiom selection algorithm. When a theorem is proved, the system produces a verifiable proof, which validates both the clausification phase and the refutation of the CNF.

Strategies

Vampire 4.9 provides a very large number of options for strategy selection. The most important ones are:

- Choices of saturation algorithm:
 - Limited Resource Strategy [58]
 - DISCOUNT loop
 - Otter loop
 - MACE-style finite model building with sort inference
- Splitting via AVATAR [140]
- A variety of optional simplifications.
- Parameterized reduction orderings.
- A number of built-in literal selection functions and different modes of comparing literals [28].
- Age-weight ratio that specifies how strongly lighter clauses are preferred for inference selection. This has been extended with a layered clause selection approach [25].
- Set-of-support strategy with extensions for theory reasoning.
- For theory-reasoning:
 - Ground equational reasoning via congruence closure.

- Addition of theory axioms and evaluation of interpreted functions [49].
- Use of Z3 with AVATAR to restrict search to ground-theory-consistent splitting branches [48].
- Specialised theory instantiation and unification [53].
- Extensionality resolution with detection of extensionality axioms

Implementation

Vampire 4.9 is implemented in C++. It makes use of fixed versions of Minisat and Z3. See the GitHub repository and associated wiki for more information.

Expected Competition Performance

Vampire 4.9 should be an improvement on the previous version. A reasonably strong performance across all divisions is therefore expected. In the higher-order divisions, performance should be the same as last year.

8.19 Zipperposition 2.1.9999

Jasmin Blanchette
Ludwig-Maximilians-Universität München, Germany

Architecture

Zipperposition is a superposition-based theorem prover for typed first-order logic with equality and for higher-order logic. It is a pragmatic implementation of a complete calculus for full higher-order logic [7]. It features a number of extensions that include polymorphic types, user-defined rewriting on terms and formulas (“deduction modulo theories”), a lightweight variant of AVATAR for case splitting [21], and Boolean reasoning [144]. The core architecture of the prover is based on saturation with an extensible set of rules for inferences and simplifications. Zipperposition uses a full higher-order unification algorithm that enables efficient integration of procedures for decidable fragments of higher-order unification [142]. The initial calculus and main loop were imitations of an earlier version of E [61]. With the implementation of higher-order superposition, the main loop had to be adapted to deal with possibly infinite sets of unifiers [141].

Strategies

The system uses various strategies in a portfolio. The strategies are run in parallel, making use of all CPU cores available. We designed the portfolio of strategies by manual inspection of TPTP problems. Zipperposition’s heuristics are inspired by efficient heuristics used in E. Various calculus extensions are used by the strategies [141]. The portfolio mode distinguishes between first-order and higher-order problems. If the problem is first-order, all higher-order prover features are turned off. In particular, the prover uses standard first-order superposition calculus and disables collaboration with the backend prover (described below). Other than that, the portfolio is static and does not depend on the syntactic properties of the problem.

Implementation

The prover is implemented in OCaml. Term indexing is done using fingerprints for unification, perfect discrimination trees for rewriting, and feature vectors for subsumption. Some inference rules such as contextual literal cutting make heavy use of subsumption. For higher-order problems, some strategies use the E prover as an end-game backend prover.

Zipperposition's code can be found at

`https://github.com/sneeuwballen/zipperposition`

and is entirely free software (BSD-licensed).

Zipperposition can also output graphic proofs using graphviz. Some tools to perform type inference and classification for typed formulas are also provided, as well as a separate library for dealing with terms and formulas [16].

Expected Competition Performance

The prover is expected to perform well on THF, about as well as last year's version. We expect to beat E.

9 Conclusion

The 12th IJCAR ATP System Competition was the twenty-ninth large scale competition for classical logic ATP systems. The organizers believe that CASC fulfills its main motivations: evaluation of relative capabilities of ATP systems, stimulation of research, motivation for improving implementations, and providing an exciting event. Through the continuity of the event and consistency in the reporting of the results, performance comparisons with previous and future years are easily possible. The competition provides exposure for system builders both within and outside of the community, and provides an overview of the implementation state of running, fully automatic, classical logic ATP systems.

References

- [1] L. Bachmair and N. Dershowitz. Critical Pair Criteria for Completion. *Journal of Symbolic Computation*, 6(1):1–18, 1988.
- [2] L. Bachmair, N. Dershowitz, and D.A. Plaisted. Completion Without Failure. In H. Ait-Kaci and M. Nivat, editors, *Resolution of Equations in Algebraic Structures*, pages 1–30. Academic Press, 1989.
- [3] H. Barbosa, C. Barrett, M. Brain, G. Kremer, H. Lachnitt, M. Mann, A. Mohamed, M. Mohamed, A. Niemetz, A. Nötzli, A. Ozdemir, M. Preiner, A. Reynolds, Y. Sheng, C. Tinelli, and Y. Zohar. cvc5: A Versatile and Industrial-Strength SMT Solver. In D. Fisman and G. Rosu, editors, *Proceedings of the 28th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 13243 in Lecture Notes in Computer Science, pages 415–442. Springer, 2022.
- [4] H. Barbosa, P. Fontaine, and A. Reynolds. Congruence Closure with Free Variables. In A. Legay and T. Margaria, editors, *Proceedings of the 23rd International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 10205 in Lecture Notes in Computer Science, pages 2134–230. Springer-Verlag, 2017.
- [5] H. Barbosa, A. Reynolds, D. El Ouraoui, C. Tinelli, and C. Barrett. Extending SMT Solvers to Higher-Order Logic. In P. Fontaine, editor, *Proceedings of the 27th International Conference on Automated Deduction*, number 11716 in Lecture Notes in Computer Science, pages 35–54. Springer-Verlag, 2019.
- [6] C. Barrett, C. Conway, M. Deters, L. Hadarean, D. Jovanovic, T. King, A. Reynolds, and C. Tinelli. CVC4. In G. Gopalakrishnan and S. Qadeer, editors, *Proceedings of the 23rd International Conference on Computer Aided Verification*, number 6806 in Lecture Notes in Computer Science, pages 171–177. Springer-Verlag, 2011.
- [7] A. Bentkamp, J. Blanchette, S. Tournet, and P. Vukmirović. Superposition for Full Higher-order Logic. In A. Platzer and G. Sutcliffe, editors, *Proceedings of the 28th International Conference on Automated Deduction*, number 12699 in Lecture Notes in Computer Science, pages 396–412. Springer-Verlag, 2021.
- [8] C. Benzmüller. Extensional Higher-order Paramodulation and RUE-Resolution. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction*, number 1632 in Lecture Notes in Artificial Intelligence, pages 399–413. Springer-Verlag, 1999.
- [9] C. Benzmüller and M. Kohlhase. LEO - A Higher-Order Theorem Prover. In C. Kirchner and H. Kirchner, editors, *Proceedings of the 15th International Conference on Automated Deduction*, number 1421 in Lecture Notes in Artificial Intelligence, pages 139–143. Springer-Verlag, 1998.
- [10] C. Benzmüller, L. Paulson, F. Theiss, and A. Fietzke. LEO-II - A Cooperative Automatic Theorem Prover for Higher-Order Logic. In P. Baumgartner, A. Armando, and G. Dowek, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 162–170. Springer-Verlag, 2008.

- [11] C. Benzmüller, F. Rabe, and G. Sutcliffe. THF0 - The Core TPTP Language for Classical Higher-Order Logic. In P. Baumgartner, A. Armando, and G. Dowek, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 491–506. Springer-Verlag, 2008.
- [12] C. Benzmüller, V. Sorge, M. Jamnik, and M. Kerber. Combined Reasoning by Automated Cooperation. *Journal of Applied Logic*, 6(3):318–342, 2008.
- [13] A. Biere. PicoSAT Essentials. *Journal on Satisfiability, Boolean Modeling and Computation*, 4:75–97, 2008.
- [14] J. Blanchette, M. Haslbeck, D. Matichuk, and T. Nipkow. Mining the Archive of Formal Proofs. In M. Kerber, J. Carette, C. Kaliszyk, F. Rabe, and V. Sorge, editors, *Proceedings of the 8th Conference on Intelligent Computer Mathematics*, number 9150 in Lecture Notes in Computer Science, pages 3–17, 2015.
- [15] J. Blanchette and A. Paskevich. TFF1: The TPTP Typed First-order Form with Rank-1 Polymorphism. In M.P. Bonacina, editor, *Proceedings of the 24th International Conference on Automated Deduction*, number 7898 in Lecture Notes in Artificial Intelligence, pages 414–420. Springer-Verlag, 2013.
- [16] S. Cruanes. *Extending Superposition with Integer Arithmetic, Structural Induction, and Beyond*. PhD thesis, Ecole Polytechnique, Paris, France, 2015.
- [17] L. de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In C. Ramakrishnan and J. Rehof, editors, *Proceedings of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 4963 in Lecture Notes in Artificial Intelligence, pages 337–340. Springer-Verlag, 2008.
- [18] A. Duarte and K. Korovin. Implementing Superposition in iProver. In N. Peltier and V. Sofronie-Stokkermans, editors, *Proceedings of the 10th International Joint Conference on Automated Reasoning*, number 12167 in Lecture Notes in Artificial Intelligence, pages 388–397, 2020.
- [19] A. Duarte and K. Korovin. AC Simplifications and Closure Redundancies in the Superposition Calculus. In A. Das and S. Negri, editors, *Proceedings of the 30th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, number 12842 in Lecture Notes in Artificial Intelligence, pages 200–217. Springer-Verlag, 2021.
- [20] A. Duarte and K. Korovin. Ground Joinability and Connectedness in the Superposition Calculus. In J. Blanchette, L. Kovacs, and D. Pattinson, editors, *Proceedings of the 11th International Joint Conference on Automated Reasoning*, number 13385 in Lecture Notes in Artificial Intelligence, pages 169–187, 2022.
- [21] G. Ebner, J. Blanchette, and S. Tourret. A Unifying Splitting Framework. In A. Platzer and G. Sutcliffe, editors, *Proceedings of the 28th International Conference on Automated Deduction*, number 12699 in Lecture Notes in Computer Science, pages 344–360. Springer-Verlag, 2021.
- [22] N. Eén and N. Sörensson. An Extensible SAT-solver. In E. Giunchiglia and A. Tacchella, editors, *Proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing*, number 2919 in Lecture Notes in Computer Science, pages 502–518. Springer-Verlag, 2004.
- [23] H. Ganzinger and K. Korovin. New Directions in Instantiation-based Theorem Proving. In P. Kolaitis, editor, *Proceedings of the 18th IEEE Symposium on Logic in Computer Science*, pages 55–64. IEEE Press, 2003.
- [24] Y. Ge and L. de Moura. Complete Instantiation for Quantified Formulas in Satisfiability Modulo Theories. In A. Bouajjani and O. Maler, editors, *Proceedings of the 21st International Conference on Computer Aided Verification*, number 5643 in Lecture Notes in Computer Science, pages 306–320. Springer-Verlag, 2009.
- [25] B. Gleiss and M. Suda. Layered Clause Selection for Theory Reasoning. In N. Peltier and V. Sofronie-Stokkermans, editors, *Proceedings of the 10th International Joint Conference on Automated Reasoning*, number 12166 in Lecture Notes in Computer Science, pages 402–409, 2020.

- [26] J. Hernandez and K. Korovin. An Abstraction-Refinement Framework for Reasoning with Large Theories. In D. Galmiche, S. Schulz, and R. Sebastiani, editors, *Proceedings of the 9th International Joint Conference on Automated Reasoning*, number 10900 in Lecture Notes in Computer Science, pages 663–679, 2018.
- [27] J. Hernandez and K. Korovin. Towards an Under-Approximation Abstraction-Refinement for Reasoning with Large Theories. In A. Bolotov and F. Kammüller, editors, *Proceedings of the 26th Automated Reasoning Workshop*, page To appear, 2019.
- [28] K. Hoder, G. Reger, M. Suda, and A. Voronkov. Selecting the Selection. In N. Olivetti and A. Tiwari, editors, *Proceedings of the 8th International Joint Conference on Automated Reasoning*, number 9706 in Lecture Notes in Artificial Intelligence, pages 313–329, 2016.
- [29] K. Hoder and A. Voronkov. Sine Qua Non for Large Theory Reasoning. In V. Sofronie-Stokkermans and N. Bjørner, editors, *Proceedings of the 23rd International Conference on Automated Deduction*, number 6803 in Lecture Notes in Artificial Intelligence, pages 299–314. Springer-Verlag, 2011.
- [30] E. Holden and K. Korovin. SMAC and XGBoost your Theorem Prover. In T. Hales, C. Kaliszyk, R. Kumar, S. Schulz, and J. Urban, editors, *Proceedings of the 4th Conference on Artificial Intelligence and Theorem Proving*, pages 93–95, 2019.
- [31] E. Holden and K. Korovin. Heterogeneous Heuristic Optimisation and Scheduling for First-Order Theorem Proving. In F. Kamareddine and C. Sacerdoti, editors, *Proceedings of the 14th International Conference on Intelligent Computer Mathematics*, number 12833 in Lecture Notes in Computer Science, pages 107–123. Springer-Verlag, 2021.
- [32] S. Holden. Connect++: A New Automated Theorem Prover Based on the Connection Calculus. In J. Otten and W. Bibel, editors, *Proceedings of the 1st International Workshop on Automated Reasoning with Connection Calculi*, number 3613 in CEUR Workshop Proceedings, pages 95–106, 2023.
- [33] J. Jakubuv and J. Urban. ENIGMA: Efficient Learning-based Inference Guiding Machine. In H. Geuvers, M. England, O. Hasan, F. Rabe, and O. Teschke, editors, *Proceedings of the 10th International Conference on Intelligent Computer Mathematics*, number 10383 in Lecture Notes in Artificial Intelligence, pages 292–302. Springer-Verlag, 2017.
- [34] J. Jakubuv and J. Urban. Enhancing ENIGMA Given Clause Guidance. In F. Rabe, W. Farmer, G. Passmore, and A. Youssef, editors, *Proceedings of the 11th International Conference on Intelligent Computer Mathematics*, number 11006 in Lecture Notes in Artificial Intelligence, pages 118–124. Springer-Verlag, 2018.
- [35] C. Kaliszyk, G. Sutcliffe, and F. Rabe. TH1: The TPTP Typed Higher-Order Form with Rank-1 Polymorphism. In P. Fontaine, S. Schulz, and J. Urban, editors, *Proceedings of the 5th Workshop on Practical Aspects of Automated Reasoning*, number 1635 in CEUR Workshop Proceedings, pages 41–55, 2016.
- [36] K. Korovin. iProver - An Instantiation-based Theorem Prover for First-order Logic (System Description). In P. Baumgartner, A. Armando, and G. Dowek, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 292–298, 2008.
- [37] K. Korovin. Inst-Gen - A Modular Approach to Instantiation-based Automated Reasoning. In A. Voronkov and C. Weidenbach, editors, *Programming Logics, Essays in Memory of Harald Ganzinger*, number 7797 in Lecture Notes in Computer Science, pages 239–270. Springer-Verlag, 2013.
- [38] L. Kovacs and A. Voronkov. First-Order Theorem Proving and Vampire. In N. Sharygina and H. Veith, editors, *Proceedings of the 25th International Conference on Computer Aided Verification*, number 8044 in Lecture Notes in Artificial Intelligence, pages 1–35. Springer-Verlag, 2013.
- [39] B. Loechner. Things to Know When Implementing KBO. *Journal of Automated Reasoning*,

- 36(4):289–310, 2006.
- [40] B. Loechner. Things to Know When Implementing LBO. *Journal of Artificial Intelligence Tools*, 15(1):53–80, 2006.
 - [41] U. Martin and T. Nipkow. Ordered Rewriting and Confluence. In M.E. Stickel, editor, *Proceedings of the 10th International Conference on Automated Deduction*, number 449 in Lecture Notes in Artificial Intelligence, pages 366–380. Springer-Verlag, 1990.
 - [42] W.W. McCune. Solution of the Robbins Problem. *Journal of Automated Reasoning*, 19(3):263–276, 1997.
 - [43] W.W. McCune. Otter 3.3 Reference Manual. Technical Report ANL/MS-C-263, Argonne National Laboratory, Argonne, USA, 2003.
 - [44] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: from an Abstract Davis-Putnam-Logemann-Loveland Procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, 2006.
 - [45] A. Nonnengart and C. Weidenbach. Computing Small Clause Normal Forms. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 335–367. Elsevier Science, 2001.
 - [46] J. Otten. Restricting Backtracking in Connection Calculi. *AI Communications*, 23(2-3):159–182, 2010.
 - [47] L. Paulson and J. Blanchette. Three Years of Experience with Sledgehammer, a Practical Link between Automatic and Interactive Theorem Provers. In G. Sutcliffe, E. Ternovska, and S. Schulz, editors, *Proceedings of the 8th International Workshop on the Implementation of Logics*, number 2 in EPiC Series in Computing, pages 1–11. EasyChair Publications, 2010.
 - [48] G. Reger, N. Bjørner, M. Suda, and A. Voronkov. AVATAR Modulo Theories. In C. Benzmüller, G. Sutcliffe, and R. Rojas, editors, *Proceedings of the 2nd Global Conference on Artificial Intelligence*, number 41 in EPiC Series in Computing, pages 39–52. EasyChair Publications, 2016.
 - [49] G. Reger, J. Schoisswohl, and A. Voronkov. Making Theory Reasoning Simpler. In J. Groote and K. Larsen, editors, *Proceedings of the 27th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 12652 in Lecture Notes in Computer Science, pages 164–180. Springer-Verlag, 2021.
 - [50] G. Reger, M. Suda, and A. Voronkov. Playing with AVATAR. In A. Felty and A. Middeldorp, editors, *Proceedings of the 25th International Conference on Automated Deduction*, number 9195 in Lecture Notes in Computer Science, pages 399–415. Springer-Verlag, 2015.
 - [51] G. Reger, M. Suda, and A. Voronkov. Finding Finite Models in Multi-Sorted First Order Logic. In N. Creignou and D. Le Berre, editors, *Proceedings of the 19th International Conference on Theory and Applications of Satisfiability Testing*, number 9710 in Lecture Notes in Computer Science, pages 323–341. Springer-Verlag, 2016.
 - [52] G. Reger, M. Suda, and A. Voronkov. New Techniques in Clausal Form Generation. In C. Benzmüller, G. Sutcliffe, and R. Rojas, editors, *Proceedings of the 2nd Global Conference on Artificial Intelligence*, number 41 in EPiC Series in Computing, pages 11–23. EasyChair Publications, 2016.
 - [53] G. Reger, M. Suda, and A. Voronkov. Unification with Abstraction and Theory Instantiation in Saturation-based Reasoning. In D. Beyer and M. Huisman, editors, *Proceedings of the 24th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 10805 in Lecture Notes in Computer Science, pages 3–22. Springer-Verlag, 2018.
 - [54] A. Reynolds, H. Barbosa, and P. Fontaine. Revisiting Enumerative Instantiation. In D. Beyer and M. Huisman, editors, *Proceedings of the 24th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 10805 in Lecture Notes in Computer Science, pages 112–131. Springer-Verlag, 2018.
 - [55] A. Reynolds, H. Barbosa, A. Nötzli, C. Barrett, and C. Tinelli. cvc4sy: Smart and Fast Term Enumeration for Syntax-Guided Synthesis. In I. Dillig and S. Tasiran, editors, *Proceedings of the*

- 31st International Conference on Computer Aided Verification*, number 11562 in Lecture Notes in Computer Science, pages 74–83. Springer-Verlag, 2019.
- [56] A. Reynolds, C. Tinelli, and L. de Moura. Finding Conflicting Instances of Quantified Formulas in SMT. In K. Claessen and V. Kuncak, editors, *Proceedings of the 14th Conference on Formal Methods in Computer-Aided Design*, pages 195–202, 2014.
- [57] A. Reynolds, C. Tinelli, A. Goel, S. Krstic, M. Deters, and C. Barrett. Quantifier Instantiation Techniques for Finite Model Finding in SMT. In M.P. Bonacina, editor, *Proceedings of the 24th International Conference on Automated Deduction*, number 7898 in Lecture Notes in Artificial Intelligence, pages 377–391. Springer-Verlag, 2013.
- [58] A. Riazanov and A. Voronkov. Limited Resource Strategy in Resolution Theorem Proving. *Journal of Symbolic Computation*, 36(1-2):101–115, 2003.
- [59] O. Roussel. Controlling a Solver Execution with the `runsolver` Tool. *Journal of Satisfiability, Boolean Modeling and Computation*, 7(4):139–144, 2011.
- [60] S. Schulz. A Comparison of Different Techniques for Grounding Near-Propositional CNF Formulae. In S. Haller and G. Simmons, editors, *Proceedings of the 15th International FLAIRS Conference*, pages 72–76. AAAI Press, 2002.
- [61] S. Schulz. E: A Brainiac Theorem Prover. *AI Communications*, 15(2-3):111–126, 2002.
- [62] S. Schulz. System Abstract: E 0.81. In M. Rusinowitch and D. Basin, editors, *Proceedings of the 2nd International Joint Conference on Automated Reasoning*, number 3097 in Lecture Notes in Artificial Intelligence, pages 223–228. Springer-Verlag, 2004.
- [63] S. Schulz. Fingerprint Indexing for Paramodulation and Rewriting. In B. Gramlich, D. Miller, and U. Sattler, editors, *Proceedings of the 6th International Joint Conference on Automated Reasoning*, number 7364 in Lecture Notes in Artificial Intelligence, pages 477–483. Springer-Verlag, 2012.
- [64] S. Schulz. Simple and Efficient Clause Subsumption with Feature Vector Indexing. In M.P. Bonacina and M. Stickel, editors, *Automated Reasoning and Mathematics: Essays in Memory of William W. McCune*, number 7788 in Lecture Notes in Artificial Intelligence, pages 45–67. Springer-Verlag, 2013.
- [65] S. Schulz. System Description: E 1.8. In K. McMillan, A. Middeldorp, and A. Voronkov, editors, *Proceedings of the 19th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning*, number 8312 in Lecture Notes in Computer Science, pages 477–483. Springer-Verlag, 2013.
- [66] S. Schulz, S. Cruanes, and P. Vukmirović. Faster, Higher, Stronger: E 2.3. In P. Fontaine, editor, *Proceedings of the 27th International Conference on Automated Deduction*, number 11716 in Lecture Notes in Computer Science, pages 495–507. Springer-Verlag, 2019.
- [67] N. Smallbone. Twee: An Equational Theorem Prover (System Description). In A. Platzer and G. Sutcliffe, editors, *Proceedings of the 28th International Conference on Automated Deduction*, number 12699 in Lecture Notes in Computer Science, pages 602–613. Springer-Verlag, 2021.
- [68] A. Steen and C. Benzmüller. Extensional Higher-Order Paramodulation in Leo-III. *Journal of Automated Reasoning*, 65(6):775–807, 2021.
- [69] A. Stump, G. Sutcliffe, and C. Tinelli. StarExec: a Cross-Community Infrastructure for Logic Solving. In S. Demri, D. Kapur, and C. Weidenbach, editors, *Proceedings of the 7th International Joint Conference on Automated Reasoning*, number 8562 in Lecture Notes in Artificial Intelligence, pages 367–373, 2014.
- [70] M. Suda. Aiming for the Goal with SInE. In C. Benzmüller, C. Lisetti, and M. Theobald, editors, *Proceedings of the 5th and 6th Vampire Workshops*, number 71 in EPiC Series in Computing, page 38–44. EasyChair Publications, 2019.
- [71] M. Suda. Vampire Getting Noisy: Will Random Bits Help Conquer Chaos? (System Description). In J. Blanchette, L. Kovacs, and D. Pattinson, editors, *Proceedings of the 11th International Joint Conference on Automated Reasoning*, number 13385 in Lecture Notes in Artificial Intelligence,

- pages 659–667, 2022.
- [72] G. Sutcliffe. Proceedings of the CADE-16 ATP System Competition. Trento, Italy, 1999.
 - [73] G. Sutcliffe. Proceedings of the CADE-17 ATP System Competition. Pittsburgh, USA, 2000.
 - [74] G. Sutcliffe. The CADE-16 ATP System Competition. *Journal of Automated Reasoning*, 24(3):371–396, 2000.
 - [75] G. Sutcliffe. Proceedings of the IJCAR ATP System Competition. Siena, Italy, 2001.
 - [76] G. Sutcliffe. The CADE-17 ATP System Competition. *Journal of Automated Reasoning*, 27(3):227–250, 2001.
 - [77] G. Sutcliffe. Proceedings of the CADE-18 ATP System Competition. Copenhagen, Denmark, 2002.
 - [78] G. Sutcliffe. Proceedings of the CADE-19 ATP System Competition. Miami, USA, 2003.
 - [79] G. Sutcliffe. Proceedings of the 2nd IJCAR ATP System Competition. Cork, Ireland, 2004.
 - [80] G. Sutcliffe. Proceedings of the CADE-20 ATP System Competition. Tallinn, Estonia, 2005.
 - [81] G. Sutcliffe. The IJCAR-2004 Automated Theorem Proving Competition. *AI Communications*, 18(1):33–40, 2005.
 - [82] G. Sutcliffe. Proceedings of the 3rd IJCAR ATP System Competition. Seattle, USA, 2006.
 - [83] G. Sutcliffe. The CADE-20 Automated Theorem Proving Competition. *AI Communications*, 19(2):173–181, 2006.
 - [84] G. Sutcliffe. Proceedings of the CADE-21 ATP System Competition. Bremen, Germany, 2007.
 - [85] G. Sutcliffe. The 3rd IJCAR Automated Theorem Proving Competition. *AI Communications*, 20(2):117–126, 2007.
 - [86] G. Sutcliffe. Proceedings of the 4th IJCAR ATP System Competition. Sydney, Australia, 2008.
 - [87] G. Sutcliffe. The CADE-21 Automated Theorem Proving System Competition. *AI Communications*, 21(1):71–82, 2008.
 - [88] G. Sutcliffe. The SZS Ontologies for Automated Reasoning Software. In G. Sutcliffe, P. Rudnicki, R. Schmidt, B. Konev, and S. Schulz, editors, *Proceedings of the LPAR Workshops: Knowledge Exchange: Automated Provers and Proof Assistants, and the 7th International Workshop on the Implementation of Logics*, number 418 in CEUR Workshop Proceedings, pages 38–49, 2008.
 - [89] G. Sutcliffe. Proceedings of the CADE-22 ATP System Competition. Montreal, Canada, 2009.
 - [90] G. Sutcliffe. The 4th IJCAR Automated Theorem Proving System Competition - CASC-J4. *AI Communications*, 22(1):59–72, 2009.
 - [91] G. Sutcliffe. Proceedings of the 5th IJCAR ATP System Competition. Edinburgh, United Kingdom, 2010.
 - [92] G. Sutcliffe. The CADE-22 Automated Theorem Proving System Competition - CASC-22. *AI Communications*, 23(1):47–60, 2010.
 - [93] G. Sutcliffe. Proceedings of the CADE-23 ATP System Competition. Wroclaw, Poland, 2011.
 - [94] G. Sutcliffe. The 5th IJCAR Automated Theorem Proving System Competition - CASC-J5. *AI Communications*, 24(1):75–89, 2011.
 - [95] G. Sutcliffe. Proceedings of the 6th IJCAR ATP System Competition. Manchester, England, 2012.
 - [96] G. Sutcliffe. The CADE-23 Automated Theorem Proving System Competition - CASC-23. *AI Communications*, 25(1):49–63, 2012.
 - [97] G. Sutcliffe. Proceedings of the 24th CADE ATP System Competition. Lake Placid, USA, 2013.
 - [98] G. Sutcliffe. The 6th IJCAR Automated Theorem Proving System Competition - CASC-J6. *AI Communications*, 26(2):211–223, 2013.
 - [99] G. Sutcliffe. Proceedings of the 7th IJCAR ATP System Competition. Vienna, Austria, 2014.
 - [100] G. Sutcliffe. The CADE-24 Automated Theorem Proving System Competition - CASC-24. *AI*

- Communications*, 27(4):405–416, 2014.
- [101] G. Sutcliffe. Proceedings of the CADE-25 ATP System Competition. Berlin, Germany, 2015. <http://tptp.org/CASC/25/Proceedings.pdf>.
 - [102] G. Sutcliffe. The 7th IJCAR Automated Theorem Proving System Competition - CASC-J7. *AI Communications*, 28(4):683–692, 2015.
 - [103] G. Sutcliffe. Proceedings of the 8th IJCAR ATP System Competition. Coimbra, Portugal, 2016. <http://tptp.org/CASC/J8/Proceedings.pdf>.
 - [104] G. Sutcliffe. The 8th IJCAR Automated Theorem Proving System Competition - CASC-J8. *AI Communications*, 29(5):607–619, 2016.
 - [105] G. Sutcliffe. The CADE ATP System Competition - CASC. *AI Magazine*, 37(2):99–101, 2016.
 - [106] G. Sutcliffe. Proceedings of the 26th CADE ATP System Competition. Gothenburg, Sweden, 2017. <http://tptp.org/CASC/26/Proceedings.pdf>.
 - [107] G. Sutcliffe. The CADE-26 Automated Theorem Proving System Competition - CASC-26. *AI Communications*, 30(6):419–432, 2017.
 - [108] G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure. From CNF to TH0, TPTP v6.4.0. *Journal of Automated Reasoning*, 59(4):483–502, 2017.
 - [109] G. Sutcliffe. Proceedings of the 9th IJCAR ATP System Competition. Oxford, United Kingdom, 2018. <http://tptp.org/CASC/J9/Proceedings.pdf>.
 - [110] G. Sutcliffe. The 9th IJCAR Automated Theorem Proving System Competition - CASC-J9. *AI Communications*, 31(6):495–507, 2018.
 - [111] G. Sutcliffe. Proceedings of the CADE-27 ATP System Competition. Natal, Brazil, 2019. <http://tptp.org/CASC/27/Proceedings.pdf>.
 - [112] G. Sutcliffe. Proceedings of the 10th IJCAR ATP System Competition. Online, 2020. <http://tptp.org/CASC/J10/Proceedings.pdf>.
 - [113] G. Sutcliffe. The CADE-27 Automated Theorem Proving System Competition - CASC-27. *AI Communications*, 32(5-6):373–389, 2020.
 - [114] G. Sutcliffe. Proceedings of the CADE-28 ATP System Competition. Online, 2021. <http://tptp.org/CASC/28/Proceedings.pdf>.
 - [115] G. Sutcliffe. The 10th IJCAR Automated Theorem Proving System Competition - CASC-J10. *AI Communications*, 34(2):163–177, 2021.
 - [116] G. Sutcliffe. Proceedings of the 11th IJCAR ATP System Competition. Online, 2022. <http://tptp.org/CASC/J11/Proceedings.pdf>.
 - [117] G. Sutcliffe. Proceedings of the CADE-29 ATP System Competition. Online, 2023. <http://tptp.org/CASC/29/Proceedings.pdf>.
 - [118] G. Sutcliffe. Proceedings of the 12th IJCAR ATP System Competition. Online, 2024. <http://tptp.org/CASC/J12/Proceedings.pdf>.
 - [119] G. Sutcliffe and C. Benz Müller. Automated Reasoning in Higher-Order Logic using the TPTP THF Infrastructure. *Journal of Formalized Reasoning*, 3(1):1–27, 2010.
 - [120] G. Sutcliffe and M. Desharnais. The CADE-28 Automated Theorem Proving System Competition - CASC-28. *AI Communications*, 34(4):259–276, 2022.
 - [121] G. Sutcliffe and M. Desharnais. The 11th IJCAR Automated Theorem Proving System Competition - CASC-J11. *AI Communications*, 36(2):73–91, 2023.
 - [122] G. Sutcliffe and M. Desharnais. The CADE-29 Automated Theorem Proving System Competition - CASC-29. *AI Communications*, pages To appear, 2024.
 - [123] G. Sutcliffe, S. Schulz, K. Claessen, and A. Van Gelder. Using the TPTP Language for Writing Derivations and Finite Interpretations. In U. Furbach and N. Shankar, editors, *Proceedings of the 3rd International Joint Conference on Automated Reasoning*, number 4130 in Lecture Notes in Artificial Intelligence, pages 67–81. Springer, 2006.

- [124] G. Sutcliffe and C. Suttner. The CADE-14 ATP System Competition. Technical Report 98/01, Department of Computer Science, James Cook University, Townsville, Australia, 1998.
- [125] G. Sutcliffe and C. Suttner. The CADE-18 ATP System Competition. *Journal of Automated Reasoning*, 31(1):23–32, 2003.
- [126] G. Sutcliffe and C. Suttner. The CADE-19 ATP System Competition. *AI Communications*, 17(3):103–182, 2004.
- [127] G. Sutcliffe, C. Suttner, and F.J. Pelletier. The IJCAR ATP System Competition. *Journal of Automated Reasoning*, 28(3):307–320, 2002.
- [128] G. Sutcliffe and C.B. Suttner. Special Issue: The CADE-13 ATP System Competition. *Journal of Automated Reasoning*, 18(2), 1997.
- [129] G. Sutcliffe and C.B. Suttner. The Procedures of the CADE-13 ATP System Competition. *Journal of Automated Reasoning*, 18(2):163–169, 1997.
- [130] G. Sutcliffe and C.B. Suttner. Proceedings of the CADE-15 ATP System Competition. Lindau, Germany, 1998.
- [131] G. Sutcliffe and C.B. Suttner. The CADE-15 ATP System Competition. *Journal of Automated Reasoning*, 23(1):1–23, 1999.
- [132] G. Sutcliffe and C.B. Suttner. Evaluating General Purpose Automated Theorem Proving Systems. *Artificial Intelligence*, 131(1-2):39–54, 2001.
- [133] G. Sutcliffe and J. Urban. The CADE-25 Automated Theorem Proving System Competition - CASC-25. *AI Communications*, 29(3):423–433, 2016.
- [134] C.B. Suttner and G. Sutcliffe. The CADE-14 ATP System Competition. *Journal of Automated Reasoning*, 21(1):99–134, 1998.
- [135] T. Tammet. GKC: a Reasoning System for Large Knowledge Bases. In P. Fontaine, editor, *Proceedings of the 27th International Conference on Automated Deduction*, number 11716 in Lecture Notes in Computer Science, pages 538–549. Springer-Verlag, 2019.
- [136] T. Tammet, D. Draheim, and P. Järvi. Confidences for Commonsense Reasoning. In A. Platzer and G. Sutcliffe, editors, *Proceedings of the 28th International Conference on Automated Deduction*, number 12699 in Lecture Notes in Computer Science, pages 507–524. Springer-Verlag, 2021.
- [137] T. Tammet, D. Draheim, and P. Järvi. Gk: Implementing Full First Order Default Logic for Commonsense Reasoning (System Description). In J. Blanchette, L. Kovacs, and D. Pattinson, editors, *Proceedings of the 11th International Joint Conference on Automated Reasoning*, number 13385 in Lecture Notes in Artificial Intelligence, pages 300–309, 2022.
- [138] T. Tammet, P. Järvi, M. Verrev, and D. Draheim. An Experimental Pipeline for Automated Reasoning in Natural Language. In B. Pientka and C. Tinelli, editors, *Proceedings of the 29th International Conference on Automated Deduction*, number 14132 in Lecture Notes in Computer Science, pages 509–521. Springer-Verlag, 2023.
- [139] T. Tammet and G. Sutcliffe. Combining JSON-LD with First Order Logic. In E. Marx and T. Soru, editors, *Proceedings of the 15th IEEE International Conference on Semantic Computing*, pages 256–261, 2021.
- [140] A. Voronkov. AVATAR: The New Architecture for First-Order Theorem Provers. In A. Biere and R. Bloem, editors, *Proceedings of the 26th International Conference on Computer Aided Verification*, number 8559 in Lecture Notes in Computer Science, pages 696–710, 2014.
- [141] P. Vukmirović, A. Bentkamp, J. Blanchette, S. Cruanes, V. Nummelin, and S. Tournet. Making Higher-order Superposition Work. In A. Platzer and G. Sutcliffe, editors, *Proceedings of the 28th International Conference on Automated Deduction*, number 12699 in Lecture Notes in Computer Science, pages 415–432. Springer-Verlag, 2021.
- [142] P. Vukmirović, A. Bentkamp, and V. Nummelin. Efficient Full Higher-order Unification. In Z.M. Ariola, editor, *Proceedings of the 5th International Conference on Formal Structures for Computation and Deduction*, number 167 in Leibniz International Proceedings in Informatics,

- pages 5:1–5:20. Dagstuhl Publishing, 2020.
- [143] P. Vukmirović, J. Blanchette, and S. Schulz. Extending a High-Performance Prover to Higher-Order Logic. In S. Sankaranarayanan and N. Sharygina, editors, *Proceedings of the 29th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 13994 in Lecture Notes in Computer Science, page 111–129. Springer-Verlag, 2023.
 - [144] P. Vukmirović and V. Nummelin. Boolean Reasoning in a Higher-Order Superposition Prover. In P. Fontaine, P. Rümmer, and S. Tourret, editors, *Proceedings of the 7th Workshop on Practical Aspects of Automated Reasoning*, number 2752 in CEUR Workshop Proceedings, pages 148–166, 2020.
 - [145] M. Wisniewski, A. Steen, and C. Benz Müller. LeoPARD - A Generic Platform for the Implementation of Higher-Order Reasoners. In M. Kerber, J. Carette, C. Kaliszyk, F. Rabe, and V. Sorge, editors, *Proceedings of the International Conference on Intelligent Computer Mathematics*, number 9150 in Lecture Notes in Computer Science, pages 325–330. Springer-Verlag, 2015.
 - [146] Y. Xu, J. Liu, S. Chen, X. Zhong, and X. He. Contradiction Separation Based Dynamic Multi-clause Synergized Automated Deduction. *Information Sciences*, 462:93–113, 2018.