

CASO-25

CASO-25

CASO-25

CASO-25



# Proceedings of the CADE-25 ATP System Competition CASC-25

Geoff Sutcliffe

University of Miami, USA

## Abstract

The CADE ATP System Competition (CASC) evaluates the performance of sound, fully automatic, classical logic, ATP systems. The evaluation is in terms of the number of problems solved, the number of acceptable proofs and models produced, and the average runtime for problems solved, in the context of a bounded number of eligible problems chosen from the TPTP problem library, and specified time limits on solution attempts. The CADE-25 ATP System Competition (CASC-25) was held on 4th August 2015. The design of the competition and its rules, and information regarding the competing systems, are provided in this report.

## 1 Introduction

The CADE and IJCAR conferences are the major forum for the presentation of new research in all aspects of automated deduction. In order to stimulate ATP research and system development, and to expose ATP systems within and beyond the ATP community, the CADE ATP System Competition (CASC) is held at each CADE and IJCAR conference. CASC-25 was held on 4th August 2015, as part of the 25th International Conference on Automated Deduction (CADE-25), in Berlin, Germany. It was the twentieth competition in the CASC series [122, 127, 125, 87, 89, 121, 119, 120, 94, 96, 98, 100, 103, 106, 108, 110, 112, 114, 115].

CASC evaluates the performance of sound, fully automatic, classical logic, ATP systems. The evaluation is in terms of:

- the number of problems solved,
- the number of acceptable proofs and models produced, and
- the average runtime for problems solved;

in the context of:

- a bounded number of eligible problems, chosen from the TPTP problem library [104], and
- specified time limits on solution attempts.

Twenty-seven ATP system versions, listed in Table 1, entered into the various competition and demonstration divisions. The winners of the CASC-J7 (the previous CASC) divisions were automatically entered into those divisions, to provide benchmarks against which progress can be judged (the competition archive provides access to the systems' executables and source code).<sup>1</sup>

The design and procedures of this CASC evolved from those of previous CASCs [122, 123, 118, 124, 85, 86, 88, 90, 91, 92, 93, 95, 97, 99, 102, 105, 107, 109, 111, 113]. Important changes for this CASC were:

- The THN and TFN divisions were added.

---

<sup>1</sup>As the LTB division had been suspended since CASC-24 in 2013, the CASC-24 UEQ winner was entered.

ATP System	Divisions	Entrant (Associates)	Entrant's Affiliation
Beagle 0.9.22	TFA TFN	Peter Baumgartner (Joshua Bax)	NICTA and ANU
CVG4 1.4	TFA	CASC	CASC-J7 TFA winner
CVG4 1.5	TFA TFN FOF FNT	Andrew Reynolds (Clark Barrett, Cesare Tinelli, Tim King)	EPFL
E 1.9.1	FOF FNT EPR LTB	Stephan Schulz	DHBW Stuttgart
ePrincess 1.0	FOF	Peter Backeman (Philipp Rümmer)	Uppsala University
ET 0.2	FOF LTB-demo	Josef Urban	Radboud University Nijmegen
Geo-III 2015E	FOF FNT EPR	Hans de Nivelle	University of Wrocław
iProver 0.9	EPR	CASC	CASC-J7 EPR winner
iProver 1.0	FNT	CASC	CASC-J7 FNT winner
iProver 2.0	FOF FNT EPR LTB	Konstantin Korovin	University of Manchester
iProverModulo 0.7-0.3	FOF	Guillaume Burel	ENSIEE/Cedric/Deducteam
Isabelle 2015	THF	Jasmin Blanchette (Lawrence Paulson, Tobias Nipkow, Makarius Wenzel)	Inria Nancy
leanCoP 2.2	FOF	Jens Otten	University of Potsdam
LEO-II 1.6.2	THF	Christoph Benz Müller	Freie Universität Berlin
MaLARea 0.5	LTB	CAS	CASC-24 LTB winner
Muscadet 4.5	FOF	Dominique Pastre	University Paris Descartes
Nitpick 2015	THN	Jasmin Blanchette	Inria Nancy
Princess 20150706	TFA TFN	Philipp Rümmer (Peter Backeman)	Uppsala University
Prover9 2009-11A	FOF	CASC (William McCune, Bob Veroff)	CASC fixed point
Refite 2015	THN	Jasmin Blanchette (Tjark Weber)	Inria Nancy
Satallax 2.8	THF THN	Nik Sultana (Chad Brown)	Cambridge University
Satallax-MaLeS 1.3	THF	CASC	CASC-J7 THF winner
SPASS+T 2.2.22	TFA	Uwe Waldmann	Max-Planck-Institut für Informatik
Vampire 2.6	FOF	CASC	CASC-J7 FOF winner
Vampire 4.0	TFA FOF FNT EPR LTB	Giles Reger (Andrei Voronkov, Martin Suda, Laura Kovacs)	University of Manchester
VampireZ3 1.0	TFA	Giles Reger (Andrei Voronkov, Martin Suda)	University of Manchester
ZenonArith 0.1.0	TFA	Guillaume Bury (David Delahaye)	Inria

Table 1: The ATP systems and entrants

- The LTB division returned from its one year hiatus.
- The UEQ division returned to its hiatus state.

The competition organizer was Geoff Sutcliffe, assisted in CASC-25 by Josef Urban (who was responsible for the LTB division). The competition is overseen by a panel of knowledgeable researchers who are not participating in the event. The panel members were Pascal Fontaine, Aart Middeldorp, and Neil Murray. The CASC rules, specifications, and deadlines are absolute. Only the panel has the right to make exceptions. The competition was run on computers provided by StarExec at the University of Iowa, and the Department of Computer Science, University of Manchester, United Kingdom. The CASC-25 web site provides access to resources used before, during, and after the event: <http://www.tptp.org/CASC/25>

It was assumed that all entrants had read the web pages related to the competition, and have complied with the competition rules. Non-compliance with the rules could lead to disqualification. A “catch-all” rule was used to deal with any unforeseen circumstances: *No cheating is allowed*. The panel was allowed to disqualify entrants due to unfairness, and to adjust the competition rules in case of misuse.

## 2 Divisions

CASC is divided into divisions according to problem and system characteristics. There are *competition divisions* in which systems are explicitly ranked, and a *demonstration division* in which systems demonstrate their abilities without being ranked. Some divisions are further divided into problem categories, which makes it possible to analyse, at a more fine grained level, which systems work well for what types of problems. The problem categories have no effect on the competition rankings, which are made at only the division level.

### 2.1 The Competition Divisions

The competition divisions are open to ATP systems that meet the required system properties, described in Section 6.1. Each division uses problems that have certain logical, language, and syntactic characteristics, so that the ATP systems that compete in the division are, in principle, able to attempt all the problems in the division.

The **THF** division: Typed Higher-order Form theorems (axioms with a provable conjecture). The THF division has two problem categories:

- The **TNE** category: THF with No Equality
- The **TEQ** category: THF with Equality

The **THN** division: Typed Higher-order form Non-theorems (axioms with a countersatisfiable conjecture, and satisfiable axiom sets). The THN division has two problem categories:

- The **TNN** category: THN with No equality
- The **TNQ** category: THN with Equality

The **TFA** division: Typed First-order with Arithmetic theorems (axioms with a provable conjecture). The TFA division has three problem categories:

- The **TFI** category: TFA with only Integer arithmetic
- The **TFR** category: TFA with only Rational arithmetic
- The **TFE** category: TFA with only Real arithmetic

The **TFN** division: Typed First-order with arithmetic Non-theorems (axioms with a provable conjecture). The TFN division has three problem categories:

- The **TIN** category: TFN with only Integer arithmetic
- The **TRN** category: TFN with only Rational arithmetic
- The **TEN** category: TFN with only Real arithmetic

The **FOF** division: First-Order Form theorems (syntactically non-propositional, axioms with a provable conjecture). The FOF division has two problem categories:

- The **FNE** category: FOF with No Equality
- The **FEQ** category: FOF with Equality

The **FNT** division: First-order form Non-Theorems (syntactically non-propositional, axioms with a countersatisfiable conjecture, and satisfiable axiom sets). The FNT division has two problem categories:

- The **FNN** category: FNT with No equality
- The **FNQ** category: FNT with Equality

The **EPR** division: Effectively PRopositional (but syntactically non-propositional) clause normal form theorems and non-theorems (clause sets). *Effectively propositional* means that the problems are known to be reducible to propositional problems, e.g., CNF problems that have no functions with arity greater than zero. The EPR division has two problem categories:

- The **EPT** category: Effectively Propositional Theorems (unsatisfiable clause sets)
- The **EPS** category: Effectively Propositional non-theorems (Satisfiable clause sets)

The **LTB** division: First-order form theorems (axioms with a provable conjecture) from Large Theories, presented in Batches. A large theory has many functors and predicates, has many axioms of which typically only a few are required for the proof of a theorem, and a common core set of axioms that are used in many problems. The batch presentation allows the ATP systems to load and preprocess the common core set of axioms just once, and to share logical and control results between proof searches. The LTB division's problem categories are accompanied by sets of training problems and their solutions, taken from the same exports as the competition problems, that can be used for tuning and training during (typically at the start of) the competition. The LTB division has four problem categories:

- The **HLL** category: Problems exported from HOL Light.
- The **HL4** category: Problems exported from HOL4.
- The **ISA** category: Problems exported from Isabelle.
- The **MZR** category: Problems exported from the Mizar Mathematical Library.

Section 3.2 explains what problems are eligible for use in each division and category. Section 4 explains how the systems are ranked in each division.

## 2.2 The Demonstration Division

ATP systems that cannot 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. Demonstration division systems can run on the competition computers, or the computers can be supplied by the entrant. Computers supplied by the entrant may be brought to CASC, or may be accessed via the internet. The demonstration division results are presented along with the competition divisions' results, but might not be comparable with those results. The systems are not ranked and no prizes are awarded.

## 3 Infrastructure

### 3.1 Computers

The non-LTB divisions' computers had

- Two quad-core Intel(R) Xeon(R) E5-2609, 2.40GHz CPUs
- 256GB memory
- The Red Hat Enterprise Linux Workstation release 6.3 (Santiago) operating system, kernel 2.6.32-431.1.2.el6.x86\_64

The LTB division's computers had:

- Four (one quad core chip) Intel(R) Xeon(R) L5410, 2.333GHz CPUs 12GB memory
- The Linux 2.6.29.4-167.fc11.x86\_64 operating system

Each ATP system ran one job on one computer at a time. Systems could use all the cores on the computers (although this did not necessarily help in the non-LTB divisions, because a CPU time limit was imposed).

### 3.2 Problems

#### 3.2.1 Problem Selection

Problems for the non-LTB divisions were taken from the TPTP Problem Library, version v6.2.0. The TPTP version used for CASC is released after the competition has started, so that new problems have not been seen by the entrants. The problems have to meet certain criteria to be eligible for selection. The problems used are randomly selected from the eligible problems based on a seed supplied by the competition panel.

- The TPTP uses system performance data to compute problem difficulty ratings [126]. Difficult problems with a rating in the range 0.21 to 0.99 are eligible. Problems of lesser and greater ratings might also be eligible in some divisions if there are not enough problems with the desired ratings. Performance data from systems submitted by the system submission deadline is used for computing the problem ratings for the TPTP version used for the competition.
- The TPTP distinguishes versions of problems as one of standard, incomplete, augmented, especial, or biased. All except biased problems are eligible.
- The selection is constrained so that no division or category contains an excessive number of very similar problems.
- The selection mechanism is biased to select problems that are new in the TPTP version used, until 50% of the problems in each category have been selected, after which random selection (from old and new problems) continues. The actual percentage of new problems used depends on how many new problems are eligible and the limitation on very similar problems.

Problems for the LTB division are taken from publicly available problem sets: the HLL problem category used the HH7150 problem set<sup>2</sup>; the HL4 problem category used the H4H13897 problem set<sup>3</sup>; the ISA problem category used the SH5795 problem set<sup>4</sup>; the MZR problem

<sup>2</sup><http://mizar.cs.ualberta.ca/~mptp/hh1/HH7150.tar.gz>

<sup>3</sup><https://github.com/JUrban/H4H13897>

<sup>4</sup><http://mws.cs.ru.nl/~urban/isaltb/SH5795.tar.gz>

category used the the MPTP2078 problem set<sup>5</sup>. The problems had to meet certain criteria to be eligible for selection. The problems used are randomly selected from the eligible problems based on a seed supplied by the competition panel.

- Problems in the training sets are not eligible.
- Problems that are solvable in 60s by at least one of the non-LTB versions of the systems are eligible. 70% of the selected problems will from this group.
- Problems that are not solvable in 60s by any of the non-LTB versions of the systems, but for which a reduced-axiom version of the problem is solvable by some non-LTB system, are eligible. 30% of the selected problems will from this group.

The LTB problems have consistent symbol usage between problems in each category, and almost always consistent axiom naming between problems.

### 3.2.2 Number of Problems

The minimal numbers of problems that must be used in each division and category, to ensure sufficient confidence in the competition results, are determined from the numbers of eligible problems in each division and category [34] (the competition organizers have to ensure that there are sufficient computers available to run the ATP systems on this minimal number of problems). The minimal numbers of problems is used in determining the time limits imposed on each solution attempt - see Section 3.3.

A lower bound on the total number of problems to be used is determined from the number of computers available, the time allocated to the competition, the number of ATP systems to be run on the competition computers over all the divisions, and the per-problem time limit, according to the following relationship:

$$NumberOfProblems = \frac{NumberOfComputers * TimeAllocated}{NumberOfATPSystems * TimeLimit}$$

It is a lower bound on the total number of problems because it assumes that every system uses all of the time limit for each problem. Since some solution attempts succeed before the time limit is reached, more problems can be used.

The numbers of problems used in the categories in the various divisions are (roughly) proportional to the numbers of eligible problems, after taking into account the limitation on very similar problems. The numbers of problems used in each division and category are determined according to the judgement of the competition organizers.

In the LTB division there will be at least 200 problems in each problem category.

### 3.2.3 Problem Preparation

The problems are in TPTP format, with `include` directives. The problems in each non-LTB division are given in increasing order of TPTP difficulty rating. The problems in the LTB division are given in the natural order of their export.

### 3.2.4 Batch Specification Files

The problems for each batch division and category are listed in a batch specification file, containing containing global data lines and one or more *batch specifications*. The global data lines are:

<sup>5</sup><http://wiki.mizar.org/twiki/bin/view/Mizar/MpTP2078>

- A problem category line of the form  
`division.category division_mnemonic.category_mnemonic`  
 For the LTB division it was  
`division.category LTB.category_mnemonic` where the category mnemonics were HLL, HL4, ISA, and MZR.
- The name of a directory that contains training data in the form of problems in TPTP format and one or more solutions to each problem in TSTP format, in a line of the form  
`division.category.training_directory directory_name` The **Axioms** directory in the training data contains all the axiom files that can be used in the competition problems.

Each batch specification consists of:

- A header line `% SZS start BatchConfiguration`
- A specification of whether or not the problems in the batch must be attempted in order is given, in a line of the form  
`execution.order ordered/unordered`  
 For the LTB division it was  
`execution.order ordered`  
 i.e., systems may not start any attempt on a problem, including reading the problem file, before ending the attempt on the preceding problem.
- A specification of what output is required from the ATP systems for each problem, in a line of the form  
`output.required space_separated_list`  
 where the available list values are the SZS values **Assurance**, **Proof**, **Model**, and **Answer**.  
 For the LTB division it was  
`output.required Proof.`
- The wall clock time limit per problem, in a line of the form  
`limit.time.problem.wc limit_in_seconds`  
 A value of zero indicates no per-problem limit.
- The overall wall clock time limit (for the batch), in a line of the form  
`limit.time.overall.wc limit_in_seconds`  
 A value of zero indicates no per-problem limit.
- A terminator line `% SZS end BatchConfiguration`
- A header line `% SZS start BatchIncludes`
- `include` directives that are used in every problem. Problems in the batch have all these `include` directives, and can also have other `include` directives that are not listed here.
- A terminator line `% SZS end BatchIncludes`
- A header line `% SZS start BatchProblems`
- Pairs of absolute problem file names, and absolute output file names where the output for the problem must be written.
- A terminator line `% SZS end BatchProblems`

### 3.3 Resource Limits

#### 3.3.1 Non-LTB divisions

CPU and wall clock time limits are imposed. The minimal CPU time limit per problem is 240s. The maximal CPU time limit per problem is determined using the relationship used for determining the number of problems, with the minimal number of problems as the

*NumberOfProblems.* The CPU time limit is chosen as a reasonable value within the range allowed, and is announced at the competition. The wall clock time limit is imposed in addition to the CPU time limit, to limit very high memory usage that causes swapping. The wall clock time limit per problem is double the CPU time limit. An additional memory limit is imposed, depending on the computers' memory. The time are imposed individually on each solution attempt.

In the demonstration division, each entrant can choose to use either a CPU or a wall clock time limit, whose value is the CPU time limit of the competition divisions.

### 3.3.2 LTB division

For each batch there is a wall clock time limit per problem, which is provided in the configuration section at the start of each batch. The minimal wall clock time limit per problem is 30s. For each problem category there is an overall wall clock time limit, which is provided in the configuration section at the start of each batch, and is also available as a command line parameter. The overall limit is the sum over the batches of the batch's per-problem limit multiplied by the number of problems in the batch. Time spent before starting the first problem of a batch (e.g., preloading and analysing the batch axioms), and times spent between ending a problem and starting the next (e.g., learning from a proof just found), are not part of the times taken on the individual problems, but are part of the overall time taken. There are no CPU time limits.

## 4 System Evaluation

For each ATP system, for each problem, four items of data are recorded: whether or not the problem was solved, the CPU time taken, the wall clock time taken, and whether or not a solution (proof or model) was output. In the LTB division, the wall clock time is the time from when the system reports starting on a problem and reports ending on a problem - the time spent before starting the first problem, and times spent between ending a problem and starting the next, are not part of the time taken on problems.

The systems are ranked in the competition divisions, from the performance data. The THF, TFA, FOF, FNT, and LTB divisions are ranked according to the number of problems solved with an acceptable proof/model output. The THN, TFN, and EPR divisions are ranked according to the number of problems solved, but not necessarily accompanied by a proof or model (but systems that do output proofs/models are highlighted in the presentation of results). Ties are broken according to the average time over problems solved. In the competition divisions winners are announced and prizes are awarded.

The competition panel decides whether or not the systems' proofs and models 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, including CNF refutations).
- For proofs of FOF problems by CNF refutation, the conversion from FOF to CNF 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 criteria, 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 inverse of the number of systems that solved the problem. A system's overall SOTAC is its average SOTAC over the problems it solves.
- The *efficiency measure* balances the number of problems solved with the CPU time taken. It is the average of the inverses of the times for problems solved (CPU times for the non-LTB divisions, wall clock times for the LTB division, with times less than the timing granularity rounded up to the granularity, to avoid skewing caused by very low times), multiplied by the fraction of problems solved. This can be interpreted intuitively as the average of the solution rates for problems solved, multiplied by the fraction of problems solved.
- The *core usage* is the average of the ratios of CPU time to wall clock time used, over the problems solved. This measures the extent to which the systems take advantage the multiple cores.

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 (see Section 6.1 regarding soundness checking before the competition). 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 proofs and models from the winners (of divisions ranked by the numbers of proofs/models output) are checked by the panel. If any of the proofs or models are unacceptable, i.e., they are significantly worse than the samples provided, then that system is retrospectively disqualified. All disqualifications are explained in the competition report.

## 5 System Entry

To be entered into CASC, systems must be registered using the CASC system registration form. No registrations are accepted after the registration deadline. For each system entered, an entrant has to be nominated to handle all issues (including execution difficulties) arising before and during the competition. The nominated entrant must formally register for CASC. It is not necessary for entrants to physically attend the competition.

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 competition 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 may 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; the competition panel may disallow or move such systems to the demonstration division.

The division winners from the previous CASC are automatically entered into their divisions, to provide benchmarks against which progress can be judged. Prover9 2009-11A is automatically entered into the FOF division, to provide a fixed-point against which progress can be judged.

## 5.1 System Description

A system description has to be provided for each ATP system entered, using the HTML schema supplied on the CASC web site. (See Section 7 for these descriptions.) 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 6.1).
- **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 in 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.

## 5.2 Sample Solutions

For systems in the proof/model classes, 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 proofs and models are acceptable.

Proof/model samples are required as follows:

- THF: SET014<sup>4</sup>
- TFA: DAT013=1
- FOF and LTB: SEU140+2
- FNT: NLP042+1 and SWV017+1

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

# 6 System Requirements

## 6.1 System Properties

Entrants must ensure that their systems execute in a competition-like 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.

## Soundness and Completeness

- Systems must be sound. At some time before the competition all the systems in the competition divisions are tested for soundness. Non-theorems are submitted to the systems in the THF, TFA, FOF, EPR, and LTB divisions, and theorems are submitted to the systems in the THN, TFN, FNT and EPR divisions. Finding 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. For systems running on entrant supplied computers in the demonstration division, the entrant must perform the soundness testing and report the results to the competition organizers.
- 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 TPTP problems and axiom sets is not allowed. Strategies and strategy selection based on individual TPTP problems is not allowed. If machine learning procedures are used, the learning must ensure that sufficient generalization is obtained so that there is no specialization to individual problems or their solutions.
- The LTB division's problem categories are accompanied by sets of training problems and their solutions (taken from the same exports as the competition problems), that can be used for tuning and training during (typically at the start of) the competition. The training problems are not used in the competition. There are at least twice as many training problems as competition problems in each problem category. The training problems and solutions may be used for producing generally useful strategies that extend to other problems in the problem sets. Such strategies can rely on the consistent naming of symbols and formulas in the problem sets, and may use techniques for memorization and generalization of problems and solutions in the training set. The system description must fully explain any such tuning or training that has been done. Precomputation and storage of information about other problems in the LTB problem sets, or their solutions, is not allowed.
- The competition panel may disqualify any system whose tuning or training is deemed to be problem specific rather than general purpose.
- The system's performance must be reproducible by running the system again.

## Execution

- Systems in the non-LTB divisions must run on StarExec, and systems in the LTB division must run on the specified competition computers (see Section 3.1). ATP systems that cannot run on StarExec/competition computers can be entered into the demonstration division.
- Systems must be fully automatic, i.e., all command line switches have to be the same for all problems in each division.
- In the LTB division the systems must attempt the problems in the order given in the batch specification file. Systems may not start any attempt on a problem, including reading the problem file, before ending the attempt on the preceding problem.

## Output

- In the non-LTB divisions all solution output must be to `stdout`. In the LTB division all solution output must be to the named output file for each problem.
- In the LTB division the systems must print SZS notification lines to `stdout` when starting and ending work on a problem (including any cleanup work, such as deleting temporary files). For example

```
% SZS status Started for /home/graph/tptp/TPTP/Problems/CSR/CSR075+2.p
... (system churns away, result and solution output to file)
% SZS status Theorem for /home/graph/tptp/TPTP/Problems/CSR/CSR075+2.p
% SZS status Ended for /home/graph/tptp/TPTP/Problems/CSR/CSR075+2.p
```

- 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 [101] for this. For example

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

or

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

In the LTB division this line must be the last line output before the ending notification line (the line must also be output to the output file).

- When outputting proofs/models, the start and end of the proof/model 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
...
SZS output end CNFRefutation for SYN075-1
```

The string specifying the problem status must be output before the start of a proof/model. Use of the TPTP format for proofs and finite interpretations is encouraged [117].

## Resource Usage

- Systems that run on the competition computers must be interruptible by a `SIGXCPU` signal, so that the CPU time limit can be imposed, and interruptible by a `SIGALRM` signal, so that the wall clock time limit 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 that problem.
- If an ATP system terminates of its own accord, it may not leave any temporary or intermediate output files. If an ATP 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.

## 6.2 System Delivery

For systems in the non-LTB divisions, entrants must email a StarExec installation package to the competition organizers by the system delivery deadline. The installation package must be a `.tgz` file containing only the components necessary for running the system (i.e., not including source code, etc.). The entrants must also email a `.tgz` file containing the source code and any files required for building the StarExec installation package to the competition organizers by the system delivery deadline.

For systems running in the LTB 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.

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.

After the competition all competition division systems' source code is made publicly available on the CASC web site. In the demonstration division, the entrant specifies whether or not the source code is placed on the CASC web site. An open source license is encouraged.

## 6.3 System Execution

Execution of the ATP systems in the non-LTB divisions is controlled by StarExec. The jobs are queued onto the computers so that each computer is running one job at a time. All attempts at the  $N$ th problems in all the divisions and categories are started before any attempts at the  $(N+1)$ th problems.

Execution of the ATP systems in the LTB division is controlled by a `perl` script, provided by the competition organizers. The jobs are queued onto the computers so that each computer is running one job at a time. All attempts in each category in the division are started before any attempts in the next category.

A system has solved a problem iff it outputs its termination string within the time limit, and a system has produced a proof/model iff it outputs its end-of-proof/model string within the time limit. The result and timing data is used to generate an HTML file, and a web browser is used to display the results.

The execution of the demonstration division systems is supervised by their entrants.

# 7 The ATP Systems

These system descriptions were written by the entrants.

## 7.1 Beagle 0.9.22

Peter Baumgartner  
NICTA and Australian National University, Australia

### Architecture

Beagle is an automated theorem prover for sorted first-order logic with equality over built-in theories. The theories currently supported are integer arithmetic, linear rational arithmetic and linear real arithmetic. It accepts formulas in the FOF and TFF formats of the TPTP syntax, and formulas in the SMT-LIB version 2 format.

Beagle first converts the input formulas into clause normal form. Pure arithmetic (sub)formulas are treated by eager application of quantifier elimination. The core reasoning component implements the Hierarchic Superposition Calculus with Weak Abstraction (HSPWA) [6]. Extensions are a splitting rule for clauses that can be divided into variable disjoint parts, and a chaining inference rule for reasoning with inequalities.

The HSPWA calculus generalizes the superposition calculus by integrating theory reasoning in a black-box style. For the theories mentioned above, Beagle combines quantifier elimination procedures and other solvers to dispatch proof obligations over these theories. The default solvers are an improved version of Cooper’s algorithm for linear integer arithmetic, and the Fourier-Motzkin algorithm for linear real/rational arithmetic. Non-linear integer arithmetic is treated by partial instantiation.

### Strategies

Beagles uses the Discount loop for saturating a clause set under the calculus’ inference rules. Simplification techniques include standard ones, such as subsumption deletion, demodulation by ordered unit equations, and tautology deletion. It also includes theory specific simplification rules for evaluating ground (sub)terms, and for exploiting cancellation laws and properties of neutral elements, among others. In the competition an aggressive form of arithmetic simplification is used, which seems to perform best in practice.

Beagle uses strategy scheduling by trying (at most) two flag settings sequentially.

### Implementation

Beagle is implemented in Scala. It is a full implementation of the HSPWA calculus. It uses a simple form of indexing, essentially top-symbol hashes, stored with each term and computed in a lazy way. Fairness is achieved through a combination of measuring clause weights and their derivation-age. It can be fine-tuned with a weight-age ratio parameter, as in other provers.

Beagle’s web site is

<https://bitbucket.org/peba123/beagle>

### Expected Competition Performance

Beagle is implemented in a straightforward way and would benefit from optimized data structures. We do not expect it to come in among the first.

## 7.2 CVC4 1.4

Andrew Reynolds  
EPFL, Switzerland

### Architecture

CVC4 [4] is an SMT solver based on the DPLL(T) architecture [54] that includes built-in support for many theories including linear arithmetic, arrays, bit vectors, datatypes and strings. It incorporates various approaches for handling universally quantified formulas. In particular, CVC4 uses primarily heuristic approaches based on E-matching for answering “unsatisfiable”, and finite model finding approaches for answering “satisfiable”.

Like other SMT solvers, CVC4 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.

The finite model finding has been developed to target problems containing background theories, whose quantification is limited to finite and uninterpreted sorts. In finite model finding mode, CVC4 uses a ground theory of finite cardinality constraints that minimizes the number of ground equivalence classes, as described in [75]. When the problem is satisfiable at the ground level, a candidate model is constructed that contains complete interpretations for all predicate and function symbols. Quantifier instantiation strategies are then invoked to add instances of quantified formulas that are in conflict with the candidate model, as described in [76]. If no instances are added, then the solver reports “satisfiable”.

### Strategies

For handling theorems, CVC4 primarily uses various configurations of E-matching. This year, CVC4 incorporates new methods for finding conflicting instances of quantified formulas [74], which have been shown to lead to improved performance on unsatisfiable TPTP benchmarks. CVC4 also incorporates a model-based heuristic for handling quantified formulas containing only pure arithmetic, which will be used in the TFA division.

For handling non-theorems, the finite model finding feature of CVC4 will use a number of orthogonal quantifier instantiation strategies. This year, it will incorporate several new features, including an optimized implementation of model-based quantifier instantiation which improves upon [76], as well as techniques for sort inference. Since CVC4 with finite model finding is also capable of answering “unsatisfiable”, it will be used as a strategy for theorems as well.

### Implementation

CVC4 is implemented in C++. The code is available from

<https://github.com/CVC4>

### Expected Competition Performance

CVC4 1.4 is the CASC-J7 TFA division winner.

## 7.3 CVC4 1.5

Andrew Reynolds  
EPFL, Switzerland

### Architecture

CVC4 [4] is an SMT solver based on the DPLL(T) architecture [54] that includes built-in support for many theories, including linear arithmetic, arrays, bit vectors, datatypes and strings. It incorporates approaches for handling universally quantified formulas. CVC4 primarily uses heuristic approaches based on E-matching for theorems, and finite model finding approaches for non-theorems.

Like other SMT solvers, CVC4 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 CVC4 targets problems containing background theories whose quantification is limited to finite and uninterpreted sorts. In finite model finding mode, CVC4 uses a ground theory of finite cardinality constraints that minimizes the number of ground equivalence classes, as described in [75]. 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 [76]. If no instances are added, it reports “satisfiable”.

### Strategies

For handling theorems, CVC4 primarily uses configurations that combine conflict-based quantifier instantiation [74] and E-matching. CVC4 uses a handful of orthogonal trigger selection strategies for E-matching.

For handling non-theorems, CVC4 primarily uses finite model finding techniques. These techniques can also be used for bounded integer quantification for non-theorems involving arithmetic [72]. Since CVC4 with finite model finding is also capable of establishing unsatisfiability, it is used as a strategy for theorems as well.

For problems in pure arithmetic, CVC4 uses techniques for counterexample-guided quantifier instantiation [73], which select relevant quantifier instantiations based on models for counterexamples to quantified formulas. CVC4 relies on this method both for theorems in TFA and non-theorems in TFN.

### Implementation

CVC4 is implemented in C++. The code is available from

<https://github.com/CVC4>

### Expected Competition Performance

CVC4 has undergone various performance improvements in the past year. It is expected to perform better than last year in FOF, due to improved trigger selection for E-matching, and higher reliance upon conflict-based quantifier instantiation. It is expected to perform better than last year in TFA, due to a revised implementation of counterexample-guided quantifier instantiation, and improvements to E-matching. It is expected to be competitive in TFN.

## 7.4 E 1.9.1

Stephan Schulz  
DHBW Stuttgart, Germany

### Architecture

E [80, 83] is a purely equational theorem prover for full first-order logic with equality. 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. 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. For the LTB divisions, a

control program uses a SInE-like analysis to extract reduced axiomatizations that are handed to several instances of E. E will not use the on-the-fly learning introduced this year.

### Strategies

Proof search in E is primarily controlled by a literal selection strategy, a clause evaluation heuristic, and a simplification ordering. The prover supports a large number of pre-programmed literal selection strategies. Clause evaluation 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).

For CASC-25, E implements a strategy-scheduling automatic mode. The total CPU time available is broken into 8 (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,...). For each class, a schedule of strategies is greedily constructed from experimental data as follows: The first strategy assigned to a schedule is the 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.

About 210 different strategies have been evaluated on all untyped first-order problems from TPTP 6.0.0, and about 180 of these strategies are used in the automatic mode. A few new strategies may be added.

### 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 [81]. Superposition and backward rewriting use fingerprint indexing [82], 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 [52, 51]. The prover and additional information are available at

<http://www.eprover.org>

### Expected Competition Performance

E 1.9.1 has only seen minor changes and inconsequential bug fixes compared to last years version. It can produce proof objects quite efficiently. The system is expected to perform reasonably well in most proof classes, but will at best complement top systems in the disproof classes.

## 7.5 ePrincess 1.0

Peter Backeman  
Uppsala University, Sweden

### Architecture

Princess [78, 79] is a theorem prover for first-order logic modulo linear integer arithmetic. The prover uses a combination of techniques from the areas of first-order reasoning and SMT solving. The main underlying calculus is a free-variable tableau calculus, which is extended with constraints to enable backtracking-free proof expansion, and positive unit hyper-resolution for lightweight instantiation of quantified formulae. Linear integer arithmetic is handled using a set of built-in proof rules resembling the Omega test, which altogether yields a calculus that is complete for full Presburger arithmetic, for first-order logic, and for a number of further fragments. In addition, some built-in procedures for nonlinear integer arithmetic are available. The internal calculus of Princess only supports uninterpreted predicates; uninterpreted functions are encoded as predicates, together with the usual axioms. Through appropriate translation of quantified formulae with functions, the e-matching technique common in SMT solvers can be simulated; triggers in quantified formulae are chosen based on heuristics similar to those in the Simplify prover.

ePrincess is an extension of Princess using a calculus presented in [2] to handle first-order formulas with equality. A restricted version of simultaneous rigid e-unification is utilised to handle equalities with free variables. The solving procedures presented in [1] have been implemented and constitute the foundation of the unification step.

### Strategies

ePrincess is using the same strategies as Princess, which means for CASC, ePrincess will run a fixed schedule of configurations for each problem (portfolio method). Configurations determine, among others, the mode of proof expansion (depth-first, breadth-first), selection of triggers in quantified formulae, clausification, and the handling of functions. The portfolio was chosen based on training with a random sample of problems from the TPTP library.

### Implementation

ePrincess is entirely written in Scala and runs on any recent Java virtual machine; besides the standard Scala and Java libraries, only the Cup parser library is used. ePrincess is available from

<http://user.it.uu.se/~petba168/breu/>

### Expected Competition Performance

ePrincess is mainly designed for first order formula (FOF), and based on performance measurements it should perform decently dealing with problems containing equality.

## 7.6 ET 0.2

Josef Urban  
Radboud University Nijmegen, The Netherlands

### Architecture

E.T. 0.2 is a metasytem using E prover with specific strategies [130, 43] and preprocessing tools [41, 39, 40] that are targeted mainly at problems with many redundant axioms. Its design is motivated by the recent experiments in the Large-Theory Batch division [44] and on the Flyspeck, Mizar and Isabelle datasets, however, E.T. does no learning from related proofs.

### Strategies

We characterize formulas by the symbols and terms that they contain, normalized in various ways. Then we run various algorithms that try to remove the redundant axioms and use special strategies on such problems.

### Implementation

The metasytem is implemented in ca. 1000 lines of Perl. It uses a number of external programs, some of them based on E's code base, some of them independently implemented in C++.

### Expected Competition Performance

ET 0.2 is expected to be slightly better than ET 0.1, provided I do not do too much windsurfing.

## 7.7 Geo-III 2015E

Hans de Nivelles  
University of Wrocław, Poland

### Architecture

Geo III is a theorem prover for Partial Classical Logic [28], based on reduction to Kleene Logic [29]. Currently, only Sections 4 and 5 of [29] are implemented. Since Kleene logic generalizes 2-valued logic, it is still possible to take part in CASC. Apart from being 3-valued, the main differences with earlier versions of Geo are (1) more sophisticated learning schemes, (2) improved proof logging, and (3) replacement of recursion by explicit use of a stack.

1. The Geo family of provers uses exhaustive backtracking, combined with learning after failure. Earlier versions learned only conflict formulas. Geo III learns disjunctions of arbitrary width. Experiments show that this often results in shorter proofs.
2. If Geo will be ever embedded in proof assistants, these assistants will require proofs. In order to be able to provide these at the required level of detail, Geo III contains a hierarchy of proof rules that is independent of the rest of the system, and that can be modified independently.
3. In order to be flexible in the main loop, recursion was replaced by an explicit stack. Using an explicit stack, it is easier to remove unused assumptions, or to rearrange the order of assumptions. Also, restarts are easier to implement with a stack.

### Strategies

Geo uses breadth-first, exhaustive model search, combined with learning. In case of branching, the branches are explored in random order. Specially for CASC, a restart strategy was added, which ensures that proof search is always restarted after 4 minutes. This was done because Geo III has no indexing. After some time, proof search becomes so inefficient that it makes no sense to continue, so that it is better to restart.

**Implementation**

Geo III is written in C++ 11. No features outside of the standard were used. It has been tested with g++ (version 4.8.4) and with clang. Version 2015E contains almost no technical optimizations, because the calculus (the geometric format, the learning schemes, and the restart strategies) are not yet mature.

**Expected Competition Performance**

Since we didn't have time to test Geo extensively, we are unable to make any predictions. We hope to learn from CASC about the relative strengths of Geo.

## 7.8 iProver 0.9

Konstantin Korovin  
University of Manchester, United Kingdom

### Architecture

iProver is an automated theorem prover based on an instantiation calculus Inst-Gen [32, 46] which is complete for first-order logic. One of the distinctive features of iProver is a modular combination of first-order reasoning with ground reasoning. In particular, iProver currently integrates MiniSat [31] for reasoning with ground abstractions of first-order clauses. In addition to instantiation, iProver implements ordered resolution calculus and a combination of instantiation and ordered resolution; see [45] for the implementation details. The saturation process is implemented as a modification of a given clause algorithm. iProver uses non-perfect discrimination trees for the unification indexes, priority queues for passive clauses, and a compressed vector index for subsumption and subsumption resolution (both forward and backward). The following redundancy eliminations are implemented: blocking non-proper instantiations; mismatching constraints [33, 45]; global subsumption [45]; resolution-based simplifications and propositional-based simplifications. A compressed feature vector index is used for efficient forward/backward subsumption and subsumption resolution. Equality is dealt with (internally) by adding the necessary axioms of equality with an option of using Brand's transformation. In the LTB division, iProver-SInE uses axiom selection based on the SInE algorithm [38] as implemented in Vampire [36], i.e., axiom selection is done by Vampire and proof attempts are done by iProver.

Major additions in the current version are:

- answer computation,
- several modes for model output using first-order definitions in term algebra,
- Brand's transformation.

### Strategies

iProver has around 40 options to control the proof search including options for literal selection, passive clause selection, frequency of calling the SAT solver, simplifications and options for combination of instantiation with resolution. At CASC iProver will execute a small number of fixed schedules of selected options depending on general syntactic properties such as Horn/non-Horn, equational/non-equational, and maximal term depth.

### Implementation

iProver is implemented in OCaml and for the ground reasoning uses MiniSat. iProver accepts FOF and CNF formats, where Vampire [36] is used for clausification of FOF problems.

iProver is available from

<http://www.cs.man.ac.uk/~korovink/iprover/>

### Expected Competition Performance

iProver 0.9 is the CASC-J7 EPR division winner.

## 7.9 iProver 1.0

Konstantin Korovin, Christoph Stickel  
University of Manchester, United Kingdom

### Architecture

iProver is an automated theorem prover based on an instantiation calculus Inst-Gen [32, 47] which is complete for first-order logic. iProver combines first-order reasoning with ground reasoning for which it uses MiniSat [31] and was recently extended with PicoSAT [13] and Lingeling [14] (only MiniSat will be used at this CASC). iProver also combines instantiation with ordered resolution; see [45] for the implementation details. The proof search is implemented using a saturation process based on the given clause algorithm. iProver uses non-perfect discrimination trees for the unification indexes, priority queues for passive clauses, and a compressed vector index for subsumption and subsumption resolution (both forward and backward). The following redundancy eliminations are implemented: blocking non-proper instantiations; dismatching constraints [33, 45]; global subsumption [45]; resolution-based simplifications and propositional-based simplifications. A compressed feature vector index is used for efficient forward/backward subsumption and subsumption resolution. Equality is dealt with (internally) by adding the necessary axioms of equality with an option of using Brand's transformation. In the LTB division, iProver uses axiom selection based on the SInE algorithm [38] as implemented in Vampire [37], i.e., axiom selection is done by Vampire and proof attempts are done by iProver.

Some of iProver features are summarised below.

- proof extraction for both instantiation and resolution,
- model representation, using first-order definitions in term algebra,
- answer substitutions,
- semantic filtering,
- type inference, monotonic [26] and non-cyclic types,
- Brand's transformation.

Type inference is targeted at improving finite model finding and symmetry breaking. Semantic filtering is used in preprocessing to eliminate irrelevant clauses. Proof extraction is challenging due to simplifications such as global subsumption which involve global reasoning with the whole clause set and can be computationally expensive.

### Strategies

iProver has around 60 options to control the proof search including options for literal selection, passive clause selection, frequency of calling the SAT solver, simplifications and options for combination of instantiation with resolution. At CASC iProver will execute a small number of fixed schedules of selected options depending on general syntactic properties such as Horn/non-Horn, equational/non-equational, and maximal term depth. The strategy for satisfiable problems (FNT division) includes finite model finding.

### Implementation

iProver is implemented in OCaml and for the ground reasoning uses MiniSat [31]. iProver accepts FOF and CNF formats. Vampire [37, 35] is used for proof-producing classification of FOF problems as well as for axiom selection [38] in the LTB division.

iProver is available from

<http://www.cs.man.ac.uk/~korovink/iprover/>

### Expected Competition Performance

iProver 1.0 is the CASC-J7 FNT division winner.

## 7.10 iProver 2.0

Konstantin Korovin

University of Manchester, United Kingdom

### Architecture

iProver is an automated theorem prover based on an instantiation calculus Inst-Gen [32, 47], which is complete for first-order logic. iProver combines first-order reasoning with ground reasoning for which it uses MiniSat [31] and optionally PicoSAT [13] and Lingeling [14] (only MiniSat will be used at this CASC).

iProver also combines instantiation with ordered resolution; see [45, 47] for the implementation details. The proof search is implemented using a saturation process based on the given clause algorithm. iProver uses non-perfect discrimination trees for the unification indexes, priority queues for passive clauses, and a compressed vector index for The following redundancy eliminations are implemented: blocking non-proper instantiations; mismatching constraints [33, 45]; global subsumption [45]; resolution-based simplifications and propositional-based simplifications. A compressed feature vector index is used for efficient forward/backward subsumption and subsumption resolution. Equality is dealt with (internally) by adding the necessary axioms of equality. Recent changes in iProver include improved preprocessing and incremental finite model finding; support of the AIG format for hardware verification and model-checking (implemented by Dmitry Tsarkov).

In the LTB division, iProver uses axiom selection based on the SInE algorithm [38] as implemented in Vampire [37], i.e., axiom selection is done by Vampire and proof attempts are done by iProver.

Some of iProver features are summarised below.

- proof extraction for both instantiation and resolution [49]
- model representation, using first-order definitions in term algebra [49]
- answer substitutions,
- semantic filtering
- incremental finite model finding,
- sort inference, monotonic [26] and non-cyclic [48] sorts.
- Brand's transformation.

Sort inference is targeted at improving finite model finding and symmetry breaking. Semantic filtering is used in preprocessing to eliminate irrelevant clauses. Proof extraction is challenging due to simplifications such global subsumption which involve global reasoning with the whole clause set and can be computationally expensive.

### Strategies

iProver has around 60 options to control the proof search including options for literal selection, passive clause selection, frequency of calling the SAT solver, simplifications and options for combination of instantiation with resolution. At CASC iProver will execute a small number of fixed schedules of selected options depending on general syntactic properties such as Horn/non-Horn, equational/non-equational, and maximal term depth. For the LTB and FNT divisions several strategies are run in parallel.

### Implementation

iProver is implemented in OCaml and for the ground reasoning uses MiniSat [31]. iProver accepts FOF and CNF formats. Vampire [37, 35] and E prover [84] are used for proof-producing clausification of FOF problems, Vampire is also used for axiom selection [38] in the LTB division.

iProver is available at:

<http://www.cs.man.ac.uk/~korovink/iprover/>

### Expected Competition Performance

Compared to the last year, iProver had several improvements in datastructures and finite model finding. We expect similar performance in the EPR division and improved performance on satisfiable problems (FNT division).

## 7.11 iProverModulo 0.7-0.3

Guillaume Burel  
ENSIIE/Cedric, France

### Architecture

iProverModulo [24] is an extension of iProver [45] to integrate Polarized resolution modulo [30]. Polarized resolution modulo consists in presenting the theory in which the problem has to be solved by means of polarized rewriting rules instead of axioms. It can also be seen as a combination of the set-of-support strategy and selection of literals.

iProverModulo consists of two tools: First, *autotheo* is a theory preprocessor that converts the axioms of the input into rewriting rules that can be used by Polarized resolution modulo. Second, these rewriting rules are handled by a patched version of iProver 0.7 which integrates Polarized resolution modulo. The integration of polarized resolution modulo in iProver only affects its ordered resolution calculus, so that the instantiation calculus is untouched.

iProverModulo 0.7-0.3 outputs a proof that is made of two parts: First, *autotheo* print a derivation of the transformation of the axioms into rewriting rules. This derivation is in TSTP format and includes the CNF conversions obtained from E. Second, the modified version of iProver outputs a Dedukti proof from this rewriting rules and the non-axiom formulas, following the ideas of [25].

### Strategies

*Autotheo* is first run to transform the formulas of the problem whose role is “axiom” into polarized rewriting rules. *Autotheo* offers a set of strategies to that purpose. For the competition, the *Equiv* and the *ClausalAll* strategies will be used. The former strategy orients formulas intuitively depending of their shape. It may be incomplete, so that the prover may give up in certain cases. However, it shows interesting results on some problems. The second strategy should be complete, at least when equality is not involved. The rewriting system for the first strategy is tried for half the time given for the problem, then the prover is restarted with the second strategy if no proof has been found.

The patched version of iProver is run on the remaining formulas modulo the rewriting rules produced by *autotheo*. No scheduling is performed. To be compatible with Polarized resolution modulo, literals are selected only when they are maximal w.r.t. a KBO ordering, and orphans are not eliminated. To take advantage of Polarized resolution modulo, the resolution calculus is triggered more often than the instantiation calculus, on the contrary to the original iProver.

Normalization of clauses w.r.t. the term rewriting system produced by `autotheo` is performed by transforming these rules into an OCaml program, compiling this program, and dynamically linking it with the prover.

### Implementation

`iProverModulo` is available as a patch to `iProver`. The most important additions are the plugin-based normalization engine and the handling of polarized rewriting rules. `iProverModulo` is available from

[http://www.ensiie.fr/~guillaume.burel/blackandwhite\\_iProverModulo.html.en](http://www.ensiie.fr/~guillaume.burel/blackandwhite_iProverModulo.html.en)

Since `iProverModulo` needs to compile rewriting rules, an OCaml compiler is also provided. `Autotheo` is available independently from `iProverModulo` from

[http://www.ensiie.fr/~guillaume.burel/blackandwhite\\_autotheo.html.en](http://www.ensiie.fr/~guillaume.burel/blackandwhite_autotheo.html.en)

`Autotheo` uses `E` to compute clausal normal form of formula. The version of `E` it uses is very slightly modified to make it print the CNF derivation even if no proof is found.

Both of `autotheo` and `iProver` are written in OCaml.

### Expected Competition Performance

The core of `iProverModulo` was untouched since last time it entered the competition in CASC-24. However, compilation of rewriting rules failed at the time, so a slight improvement is to be expected this year.

## 7.12 Isabelle 2015

Jasmin Blanchette  
Inria Nancy, France

### Architecture

`Isabelle/HOL 2015` [56] is the higher-order logic incarnation of the generic proof assistant `Isabelle2015`. `Isabelle/HOL` provides several automatic proof tactics, notably an equational reasoner [55], a classical reasoner [70], and a tableau prover [68]. It also integrates external first- and higher-order provers via its subsystem `Sledgehammer` [69, 15]. `Isabelle` includes a parser for the TPTP syntaxes CNF, FOF, TFF0, and THF0, due to Nik Sultana. It also includes TPTP versions of its popular tools, invocable on the command line as `isabelle tptp_tool max_secs file.p`. For example:

```
isabelle tptp_isabelle.hot 100 SEU/SEU8243.p
```

`Isabelle` is available in two versions. The `HOT` version (which is not participating in CASC-25) includes `LEO-II` [9] and `Satallax` [21] as `Sledgehammer` backends, whereas the competition version leaves them out.

### Strategies

The `Isabelle` tactic submitted to the competition simply tries the following tactics sequentially:

- `sledgehammer` - Invokes the following sequence of provers as oracles via `Sledgehammer`:
  - `satallax` - `Satallax 2.7` [21] (*HOT version only*);
  - `leo2` - `LEO-II 1.6.2` [9] (*HOT version only*);

- `spass` - SPASS 3.8ds [17];
- `vampire` - Vampire 3.0 (revision 1435) [77];
- `e` - E 1.8 [81];
- `nitpick` - For problems involving only the type `$o` of Booleans, checks whether a finite model exists using Nitpick [20].
- `simp` - Performs equational reasoning using rewrite rules [55].
- `blast` - Searches for a proof using a fast untyped tableau prover and then attempts to reconstruct the proof using Isabelle tactics [68].
- `auto+spass` - Combines simplification and classical reasoning [70] under one roof; then invoke Sledgehammer with SPASS on any subgoals that emerge.
- `z3` - Invokes the SMT solver Z3 4.4.0 [27].
- `cvc4` - Invokes the SMT solver CVC4 1.5pre [5].
- `fast` - Searches for a proof using sequent-style reasoning, performing a depth-first search [70]. Unlike `blast`, it constructs proofs directly in Isabelle. That makes it slower but enables it to work in the presence of the more unusual features of HOL, such as type classes and function unknowns.
- `best` - Similar to `fast`, except that it performs a best-first search.
- `force` - Similar to `auto`, but more exhaustive.
- `meson` - Implements Loveland's MESON procedure [53]. Constructs proofs directly in Isabelle.
- `fastforce` - Combines `fast` and `force`.

### Implementation

Isabelle is a generic theorem prover written in Standard ML. Its meta-logic, Isabelle/Pure, provides an intuitionistic fragment of higher-order logic. The HOL object logic extends pure with a more elaborate version of higher-order logic, complete with the familiar connectives and quantifiers. Other object logics are available, notably FOL (first-order logic) and ZF (Zermelo-Fraenkel set theory).

The implementation of Isabelle relies on a small LCF-style kernel, meaning that inferences are implemented as operations on an abstract `theorem` datatype. Assuming the kernel is correct, all values of type `theorem` are correct by construction.

Most of the code for Isabelle was written by the Isabelle teams at the University of Cambridge and the Technische Universität München. Isabelle/HOL is available for all major platforms under a BSD-style license from

<http://www.cl.cam.ac.uk/research/hvg/Isabelle/>

### Expected Competition Performance

Thanks to the addition of CVC4 and a new version of Vampire, Isabelle might have become now strong enough to take on Satallax and its various declensions. But we expect Isabelle to end in second or third place, to be honest.

### 7.13 leanCoP 2.2

Jens Otten  
University of Potsdam, Germany

#### Architecture

leanCoP [61, 57] is an automated theorem prover for classical first-order logic with equality. It is a very compact implementation of the connection (tableau) calculus [12, 50].

#### Strategies

The reduction rule of the connection calculus is applied before the extension rule. Open branches are selected in a depth-first way. Iterative deepening on the proof depth is performed in order to achieve completeness. Additional inference rules and techniques include regularity, lemmata, and restricted backtracking [58]. leanCoP uses an optimized structure-preserving transformation into clausal form [58] and a fixed strategy scheduling, which is controlled by a shell script.

#### Implementation

leanCoP is implemented in Prolog. The source code of the core prover consists only of a few lines of code. Prolog's built-in indexing mechanism is used to quickly find connections when the extension rule is applied.

leanCoP can read formulae in leanCoP syntax and in TPTP first-order syntax. Equality axioms and axioms to support distinct objects are automatically added if required. The leanCoP core prover returns a very compact connection proof, which is then translated into a more comprehensive output format, e.g., into a lean (TPTP-style) connection proof or into a readable text proof.

The source code of leanCoP 2.2 is available under the GNU general public license. It can be downloaded from the leanCoP website at:

<http://www.leancop.de>

The website also contains information about ileanCoP [57] and MleanCoP [59, 60], two versions of leanCoP for first-order intuitionistic logic and first-order modal logic, respectively.

#### Expected Competition Performance

As the core prover has not changed, the performance of leanCoP 2.2 is expected to be similar to the performance of leanCoP 2.1.

### 7.14 LEO-II 1.6.2

Christoph Benzmüller  
Freie Universität Berlin, Germany

#### Architecture

LEO-II [9], the successor of LEO [8], 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 [7]. 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 [80]. 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 [10]. In fact, the development of LEO-II has largely paralleled the development of the TPTP THF language and related infrastructure [116]. 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 [11] 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

There are only minor recent improvements on LEO-II. It is unclear whether they are sufficient for attacking LEO-II’s main competitors.

## 7.15 MaLAREa 0.5

Josef Urban

Radboud University Nijmegen, The Netherlands

### Architecture

MaLAREa 0.5 [129, 131] is a metasystem for ATP in large theories where symbol and formula names are used consistently. It uses several deductive systems (now E, SPASS, Vampire, Paradox, Mace), as well as complementary AI techniques like machine learning (the SNoW system) based on symbol-based similarity, model-based similarity, term-based similarity, and obviously previous successful proofs. The version for CASC will mainly use E prover with the BliStr [130] large-theory strategies, possibly also Prover9, Mace and Paradox. The premise selection methods will likely also use the distance-weighted k-nearest neighbor [42] and E’s implementation of SInE.

### Strategies

The basic strategy is to run ATPs on problems, then use the machine learner to learn axiom relevance for conjectures from solutions, and use the most relevant axioms for next ATP attempts. This is iterated, using different time limits and axiom limits. Various features are used for learning, and the learning is complemented by other criteria like model-based reasoning, symbol and term-based similarity, etc.

### Implementation

The metasystem is implemented in ca. 2500 lines of Perl. It uses many external programs - the above mentioned ATPs and machine learner, TPTP utilities, LADR utilities for work with models, and some standard Unix tools.

MaLAREa is available from

<https://github.com/JUrban/MPTP2/tree/master/MaLAREa>

The metasystem's Perl code is released under GPL2

### Expected Competition Performance

Thanks to machine learning, MaLAREa is strongest on batches of many related problems with many redundant axioms where some of the problems are easy to solve and can be used for learning the axiom relevance. MaLAREa is not very good when all problems are too difficult (nothing to learn from), or the problems (are few and) have nothing in common. Some of its techniques (selection by symbol and term-based similarity, model-based reasoning) could however make it even there slightly stronger than standard ATPs. MaLAREa has a very good performance on the MPTP Challenge, which is a predecessor of the LTB division, and it is the winner of the 2008 MZR LTB category. MaLAREa 0.4 came second in the 2012 MZR@Turing competition and solved most problems in the Assurance class.

## 7.16 Muscadet 4.5

Dominique Pastre  
University Paris Descartes, France

### Architecture

The Muscadet theorem prover is a knowledge-based system. It is based on Natural Deduction, following the terminology of [19] and [62], and uses methods which resembles those used by humans. It is composed of an inference engine, which interprets and executes rules, and of one or several bases of facts, which are the internal representation of "theorems to be proved". Rules are either universal and put into the system, or built by the system itself by metarules from data (definitions and lemmas). Rules may add new hypotheses, modify the conclusion, create objects, split theorems into two or more subtheorems or build new rules which are local for a (sub-)theorem.

### Strategies

There are specific strategies for existential, universal, conjunctive or disjunctive hypotheses and conclusions, and equalities. Functional symbols may be used, but an automatic creation of intermediate objects allows deep subformulae to be flattened and treated as if the concepts were defined by predicate symbols. The successive steps of a proof may be forward deduction (deduce new hypotheses from old ones), backward deduction (replace the conclusion by a new one), refutation (only if the conclusion is a negation), search for objects satisfying the conclusion or dynamic building of new rules.

The system is also able to work with second order statements. It may also receive knowledge and know-how for a specific domain from a human user; see [63] and [64]. These two possibilities are not used while working with the TPTP Library.

### Implementation

Muscadet [65] is implemented in SWI-Prolog. Rules are written as more or less declarative Prolog clauses. Metarules are written as sets of Prolog clauses. The inference engine includes the Prolog interpreter and some procedural Prolog clauses. A theorem may be split into several subtheorems, structured as a tree with “and” and ”or” nodes. All the proof search steps are memorized as facts including all the elements which will be necessary to extract later the useful steps (the name of the executed action or applied rule, the new facts added or rule dynamically built, the antecedents and a brief explanation).

Muscadet is available from

<http://www.normalesup.org/~pastre/muscadet/muscadet.html>

### Expected Competition Performance

The best performances of Muscadet will be for problems manipulating many concepts in which all statements (conjectures, definitions, axioms) are expressed in a manner similar to the practice of humans, especially of mathematicians [66, 67]. It will have poor performances for problems using few concepts but large and deep formulas leading to many splittings. Its best results will be in set theory, especially for functions and relations. It has not been not really improved from previous years. Nevertheless it is still maintained and slightly improved. Its originality is that proofs are given in natural style. It should now be combined with other provers.

## 7.17 Nitpick 2015

Jasmin Blanchette  
Inria Nancy, France

### Architecture

Nitpick [20] is an open source counterexample generator for Isabelle/HOL [56]. It builds on Kodkod [128], a highly optimized first-order relational model finder based on SAT. The name Nitpick is appropriated from a now retired Alloy precursor. In a case study, it was applied successfully to a formalization of the C++ memory model [18].

### Strategies

Nitpick employs Kodkod to find a finite model of the negated conjecture. The translation from HOL to Kodkod’s first-order relational logic (FORL) is parameterized by the cardinalities of the atomic types occurring in it. Nitpick enumerates the possible cardinalities for each atomic type, exploiting monotonicity to prune the search space [16]. If a formula has a finite counterexample, the tool eventually finds it, unless it runs out of resources.

SAT solvers are particularly sensitive to the encoding of problems, so special care is needed when translating HOL formulas. As a rule, HOL scalars are mapped to FORL singletons and functions are mapped to FORL relations accompanied by a constraint. More specifically, an  $n$ -ary first-order function (curried or not) can be coded as an  $(n + 1)$ -ary relation accompanied by a constraint. However, if the return type is the type of Booleans, the function is more efficiently coded as an unconstrained  $n$ -ary relation. Higher-order quantification and functions bring complications of their own. A function from  $\sigma$  to  $\tau$  cannot be directly passed as an argument in FORL; Nitpick’s workaround is to pass  $|\sigma|$  arguments of type  $\tau$  that encode a function table.

**Implementation**

itself, which adheres to the LCF small-kernel discipline, Nitpick does not certify its results and must be trusted.

Nitpick is available as part of Isabelle/HOL for all major platforms under a BSD-style license from

<http://www.cl.cam.ac.uk/research/hvg/Isabelle/>

**Expected Competition Performance**

Thanks to Kodkod's amazing power, we expect that Nitpick will beat both Satallax and Refute with its hands tied behind its back.

**7.18 Princess 20150706**

Philipp Rümmer  
Uppsala University, Sweden

**Architecture**

Princess [78, 79] is a theorem prover for first-order logic modulo linear integer arithmetic. The prover uses a combination of techniques from the areas of first-order reasoning and SMT solving. The main underlying calculus is a free-variable tableau calculus, which is extended with constraints to enable backtracking-free proof expansion, and positive unit hyper-resolution for lightweight instantiation of quantified formulae. Linear integer arithmetic is handled using a set of built-in proof rules resembling the Omega test, which altogether yields a calculus that is complete for full Presburger arithmetic, for first-order logic, and for a number of further fragments. In addition, some built-in procedures for nonlinear integer arithmetic are available. The internal calculus of Princess only supports uninterpreted predicates; uninterpreted functions are encoded as predicates, together with the usual axioms. Through appropriate translation of quantified formulae with functions, the e-matching technique common in SMT solvers can be simulated; triggers in quantified formulae are chosen based on heuristics similar to those in the Simplify prover.

**Strategies**

For CASC, Princess will run a fixed schedule of configurations for each problem (portfolio method). Configurations determine, among others, the mode of proof expansion (depth-first, breadth-first), selection of triggers in quantified formulae, clausification, and the handling of functions. The portfolio was chosen based on training with a random sample of problems from the TPTP library.

**Implementation**

Princess is entirely written in Scala and runs on any recent Java virtual machine; besides the standard Scala and Java libraries, only the Cup parser library is used. Princess is available from

<http://www.philipp.ruemmer.org/princess.shtml>

**Expected Competition Performance**

Princess is mainly designed for integer problems (TFI), and should perform reasonably well here. Compared to last year, only minor updates were done, so that performance should be similar as in 2014. This version of Princess does not output proofs, however, and will be ranked accordingly.

**7.19 Refute 2015**

Jasmin Blanchette  
Inria Nancy, France

**Architecture**

Refute [132] is an open source counterexample generator for Isabelle/HOL [56] based on a SAT solver, and Nitpick's [20] precursor.

**Strategies**

Refute employs a SAT solver to find a finite model of the negated conjecture. The translation from HOL to propositional logic is parameterized by the cardinalities of the atomic types occurring in the conjecture. Refute enumerates the possible cardinalities for each atomic type. If a formula has a finite counterexample, the tool eventually finds it, unless it runs out of resources.

**Implementation**

Refute, like most of Isabelle/HOL, is written in Standard ML. Unlike Isabelle itself, which adheres to the LCF small-kernel discipline, Refute does not certify its results and must be trusted.

Refute is available as part of Isabelle/HOL for all major platforms under a BSD-style license from

<http://www.cl.cam.ac.uk/research/hvg/Isabelle/>

**Expected Competition Performance**

We expect Refute to beat Satallax but also to be beaten by Nitpick.

**7.20 Satallax 2.8**

Nik Sultana  
Cambridge University, United Kingdom

**Architecture**

Satallax 2.8 [21] is an automated theorem prover for higher-order logic. The particular form of higher-order logic supported by Satallax is Church's simple type theory with extensionality and choice operators. The SAT solver MiniSat [31] is responsible for much of the search for a proof. The theoretical basis of search is a complete ground tableau calculus for higher-order logic [23] with a choice operator [3]. A problem is given in the THF format. A branch is formed from the axioms of the problem and the negation of the conjecture (if any is given). From this point on, Satallax tries to determine unsatisfiability or satisfiability of this branch. Satallax progressively generates higher-order formulae and corresponding propositional clauses [22]. These formulae

and propositional clauses correspond to instances of the tableau rules. Satallax uses the SAT solver MiniSat as an engine to test the current set of propositional clauses for unsatisfiability. If the clauses are unsatisfiable, then the original branch is unsatisfiable. Additionally, Satallax may optionally generate first-order formulas in addition to the propositional clauses. If this option is used, then Satallax periodically calls the first-order theorem prover E to test for first-order unsatisfiability. If the set of first-order formulas is unsatisfiable, then the original branch is unsatisfiable.

### Strategies

There are about a hundred flags that control the order in which formulas and instantiation terms are considered and propositional clauses are generated. Other flags activate some optional extensions to the basic proof procedure (such as whether or not to call the theorem prover E). A collection of flag settings is called a mode. Approximately 500 modes have been defined and tested so far. A strategy schedule is an ordered collection of modes with information about how much time the mode should be allotted. Satallax tries each of the modes for a certain amount of time sequentially. Satallax 2.7 has strategy schedule consisting of 68 modes. Each mode is tried for time limits ranging from 0.1 seconds to 54.9 seconds. The strategy schedule was determined through experimentation using the THF problems in version 5.4.0 of the TPTP library.

### Implementation

Satallax is implemented in OCaml. A foreign function interface is used to interact with MiniSat 2.2.0. Satallax is available from

<http://mathgate.info/cebrown/satallax/>

### Expected Competition Performance

Similar to last year, since the changes from v2.7 to v2.8 consist of improvements in proof output.

## 7.21 Satallax-MaLeS 1.3

Daniel Kuehlwein  
Radboud University Nijmegen, The Netherlands

### Architecture

Satallax-MaLeS 1.3 improves Satallax's automatic mode with machine learning techniques to predict which search strategy is most likely to find a proof.

### Strategies

Satallax-MaLeS 1.3 relies on Geoff Sutcliffe's MakeListStat to classify problems. It uses the same data as Satallax-MaLeS 1.2 as basis for the learning algorithms.

### Implementation

Satallax-MaLeS is based on Python, in particular the Sklearn library that contains many machine learning algorithms. After CASC Satallax-MaLeS will be available from

<http://www.cs.ru.nl/~kuehlwein/>

**Expected Competition Performance**

Satallax-MaLes 1.3 is the CASC-J7 THF division winner.

**7.22 SPASS+T 2.2.22**

Uwe Waldmann  
Max-Planck-Institut für Informatik, Germany

**Architecture**

SPASS+T is an extension of the superposition-based theorem prover SPASS that integrates algebraic knowledge into SPASS in three complementary ways: by passing derived formulas to an external SMT procedure (currently Yices or CVC3), by adding standard axioms, and by built-in arithmetic simplification and inference rules. A first version of the system has been described in [71]; later a much more sophisticated coupling of the SMT procedure has been added [133]. The latest version fixes some bugs related to memory management and to name clashes with the SMT procedure and provides improved support for non-arithmetic problems.

**Strategies**

Standard axioms and built-in arithmetic simplification and inference rules are integrated into the standard main loop of SPASS. Inferences between standard axioms are excluded, so the user-supplied formulas are taken as set of support. The external SMT procedure runs in parallel in a separate process, leading occasionally to non-deterministic behaviour.

**Implementation**

SPASS+T is implemented in C. The system is available from

<http://people.mpi-inf.mpg.de/~uwe/software/#TSPASS>

**Expected Competition Performance**

We expect SPASS+T to be among the leading systems in the TFA division at CASC-25.

**7.23 Vampire 2.6**

Krystof Hoder, Andrei Voronkov  
University of Manchester, England

**Architecture**

Vampire 2.6 is an automatic theorem prover for first-order classical logic. It consists of a shell and a kernel. The kernel implements the calculi of ordered binary resolution and superposition for handling equality. It also implements the Inst-gen calculus. The splitting rule in kernel adds propositional parts to clauses, which are being manipulated using binary decision diagrams and a SAT solver. 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. Although the kernel of the system works only with clausal normal form, the shell accepts a problem in the full first-order logic syntax, clausifies it and performs

a number of useful transformations before passing the result to the kernel. Also the axiom selection algorithm Sine [38] can be enabled as part of the preprocessing.

When a theorem is proved, the system produces a verifiable proof, which validates both the clausification phase and the refutation of the CNF.

### Strategies

The Vampire 2.6 kernel provides a fairly large number of options for strategy selection. The most important ones are:

- Choice of the main procedure:
  - Limited Resource Strategy
  - DISCOUNT loop
  - Otter loop
  - Goal oriented mode based on tabulation
  - Instantiation using the Inst-gen calculus
- A variety of optional simplifications.
- Parameterized reduction orderings.
- A number of built-in literal selection functions and different modes of comparing literals.
- Age-weight ratio that specifies how strongly lighter clauses are preferred for inference selection.
- Set-of-support strategy.

### Implementation

Vampire 2.6 is implemented in C++.

### Expected Competition Performance

Vampire 2.6 is the CASC-J7 FOF division winner.

## 7.24 Vampire 4.0

Giles Regeer

University of Manchester, United Kingdom

### Architecture

Vampire 4.0 is an automatic theorem prover for first-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. Splitting in resolution-based proof search is controlled by the AVATAR architecture, which uses a SAT solver to make splitting decisions. Both resolution and instantiation based proof search make use of global subsumption.

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

- Choices of saturation algorithm:
  - Limited Resource Strategy
  - DISCOUNT loop
  - Otter loop
  - Instantiation using the Inst-Gen calculus
  - MACE-style finite model building with sort inference
- Splitting via AVATAR
- A variety of optional simplifications.
- Parameterized reduction orderings.
- A number of built-in literal selection functions and different modes of comparing literals.
- Age-weight ratio that specifies how strongly lighter clauses are preferred for inference selection.
- Set-of-support strategy.
- Ground equational reasoning via congruence closure.
- Evaluation of interpreted functions.
- Extensionality resolution with detection of extensionality axioms

### Implementation

Vampire 4.0 is implemented in C++.

### Expected Competition Performance

Vampire 4.0 should be an improvement on Vampire 2.6, which is the CASC-J7 FOF division winner. Improvements have also been made to the parts of the prover relevant to satisfiability checking and theory reasoning.

## 7.25 VampireZ3 1.0

Giles Reger

University of Manchester, United Kingdom

### Architecture

VampireZ3 version 1.0 is a combination of Vampire version 4.0 (see the Vampire—4.0 description) and Z3 version 4.3.1. Vampire 4.0 uses the AVATAR architecture to make splitting decisions. Briefly, the first-order search space is represented in the SAT solver with propositional symbols consistently naming variable-disjoint components. A SAT solver is then used to (iteratively) select a subset of components to search. In VampireZ3 the Z3 SMT solver is used in place of a SAT solver and *ground* components are translated into Z3 terms. This means Z3's efficient methods for ground reasoning with equality and theories are exposed by AVATAR, as the SMT solver only produces theory-consistent models.

### Strategies

All strategies of Vampire 4.0 are available. Z3 is only used when splitting is selected.

**Implementation**

Vampire and Z3 are both implemented in C++.

**Expected Competition Performance**

VampireZ3 1.0 should perform better than Vampire 4.0 on theory problems. The combination of Z3 for reasoning with ground theories and Vampire for reasoning with quantifiers should be very powerful.

**7.26 ZenonArith 0.1.0**

Guillaume Bury  
Inria, France

**Architecture**

ZenonArith 0.1.0 [BD15] is an extension of the tableaux-based Zenon [BDD07], to linear arithmetic. This extension uses the simplex algorithm as a decision procedure to solve problems over rationals, and a branch-and-bound approach to solve problems over integers.

This extension is also able to handle real linear arithmetic, which actually coincides with the rational fragment of linear arithmetic due to the syntactical restrictions of the TPTP input format for reals.

**Strategies**

This extension consists of a smooth integration of arithmetic deductive rules to the basic tableau rules, so that there is a natural interleaving between arithmetic and regular analytic rules. This extension is able to deal with problems that are (exclusively) universally quantified, or existentially quantified formulas. Universally quantified formulas are handled by translating unsatisfiability explanation from the simplex algorithm into inference rules. Instantiation for existentially quantified formulas is achieved by trying to solve sets of formulas that cover all branches of a derivation tree.

**Implementation**

Zenon Arith is implemented in OCaml, and uses the Zarith (which is a binding for the gmp library) library to handle arbitrary precision integers and rationals. It uses an incremental implementation of the simplex algorithm in OCaml as a decision procedure over linear arithmetic. This version of Zenon comes with a built-in notion of types for terms, and inference rules that discriminate over the types of expressions. Zenon Arith can automatically output Coq [CH88] proof scripts, that can then be checked by the Coq proof assistant. The main developer is Guillaume Bury. Zenon Arith is available at:

<https://www.rocq.inria.fr/deducteam/ZenonArith/>

**Expected Competition Performance**

We do not expect to be among the best provers in terms of number of solved problems, but we aim to solve a large part of the problems, with the lowest average CPU time.

## 8 Conclusion

The CADE-25 ATP System Competition was the twentieth large scale competition for classical logic ATP systems. The organizer believes that CASC fulfills its main motivations: stimulation of research, motivation for improving implementations, evaluation of relative capabilities of ATP systems, 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] P. Backeman and P. Rümmer. Efficient Algorithms for Bounded Rigid E-Unification. In H. de Nivelle, editor, *Proceedings of the 24th Conference on Automated Reasoning with Analytic Tableaux and Related Methods*, Lecture Notes in Computer Science, page To appear. Springer-Verlag, 2015.
- [2] P. Backeman and P. Rümmer. Theorem Proving with Bounded Rigid E-Unification. In A. Felty and A. Middeldorp, editors, *Proceedings of the 25th International Conference on Automated Deduction*, Lecture Notes in Computer Science, page To appear. Springer-Verlag, 2015.
- [3] J. Backes and C.E. Brown. Analytic Tableaux for Higher-Order Logic with Choice. *Journal of Automated Reasoning*, 47(4):451–479, 2011.
- [4] 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.
- [5] C. Barrett and C. Tinelli. CVC3. In W. Damm and H. Hermanns, editors, *Proceedings of the 19th International Conference on Computer Aided Verification*, number 4590 in Lecture Notes in Computer Science, pages 298–302. Springer-Verlag, 2007.
- [6] P. Baumgartner and U. Waldmann. Hierarchic Superposition With Weak Abstraction. In M.P. Bonacina, editor, *Proceedings of the 24th International Conference on Automated Deduction*, number 7898 in Lecture Notes in Artificial Intelligence, pages 39–57. Springer-Verlag, 2013.
- [7] 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.
- [8] 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.
- [9] 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 D. Gilles, 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.
- [10] C. Benzmüller, F. Rabe, and G. Sutcliffe. THF0 - The Core TPTP Language for Classical Higher-Order Logic. In P. Baumgartner, A. Armando, and D. Gilles, 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.
- [11] C. Benzmüller, V. Sorge, M. Jamnik, and M. Kerber. Combined Reasoning by Automated Cooperation. *Journal of Applied Logic*, 6(3):318–342, 2008.
- [12] W. Bibel. *Automated Theorem Proving*. Vieweg and Sohn, 1987.

- [13] A. Biere. PicoSAT Essentials. *Journal on Satisfiability, Boolean Modeling and Computation*, 4:75–97, 2008.
- [14] A. Biere. Lingeling and Friends Entering the SAT Challenge 2012. In A. Balint, A. Belov, A. Diepold, S. Gerber, M. Järvisalo, and C. Sinz, editors, *Proceedings of SAT Challenge 2012: Solver and Benchmark Descriptions*, number B-2012-2 in Department of Computer Science Series of Publications B, pages 15–16. University of Helsinki, 2012.
- [15] J. Blanchette, S. Boehme, and L. Paulson. Extending Sledgehammer with SMT Solvers. In N. Bjorner and V. Sofronie-Stokkermans, editors, *Proceedings of the 23rd International Conference on Automated Deduction*, number 6803 in Lecture Notes in Artificial Intelligence, pages 116–130. Springer-Verlag, 2011.
- [16] J. Blanchette and A. Kraus. Monotonicity Inference for Higher-Order Formulas. *Journal of Automated Reasoning*, page To appear, 2011.
- [17] J. Blanchette, A. Popescu, D. Wand, and C. Weidenbach. More SPASS with Isabelle. In L. Beringer and A. Felty, editors, *Proceedings of Interactive Theorem Proving 2012*, number 7406 in Lecture Notes in Computer Science, pages 345–360. Springer-Verlag, 2012.
- [18] J. Blanchette, T. Weber, M. Batty, S. Owens, and S. Sarkar. Nitpicking C++ Concurrency. In M. Hanus, editor, *Proceedings of the 13th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming*, pages 113–124. ACM Press, 2011.
- [19] W.W. Bledsoe. Splitting and Reduction Heuristics in Automatic Theorem Proving. *Artificial Intelligence*, 2:55–77, 1971.
- [20] S. Böhme and T. Nipkow. Sledgehammer: Judgement Day. In J. Giesl and R. Haehnle, editors, *Proceedings of the 5th International Joint Conference on Automated Reasoning*, number 6173 in Lecture Notes in Artificial Intelligence, pages 107–121, 2010.
- [21] C.E. Brown. Satallax: An Automated Higher-Order Prover (System Description). 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 111–117, 2012.
- [22] C.E. Brown. Reducing Higher-Order Theorem Proving to a Sequence of SAT Problems. *Journal of Automated Reasoning*, 51(1):57–77, 2013.
- [23] C.E. Brown and G. Smolka. Analytic Tableaux for Simple Type Theory and its First-Order Fragment. *Logical Methods in Computer Science*, 6(2), 2010.
- [24] G. Burel. Experimenting with Deduction Modulo. In N. Bjorner and V. Sofronie-Stokkermans, editors, *Proceedings of the 23rd International Conference on Automated Deduction*, number 6803 in Lecture Notes in Artificial Intelligence, pages 162–176. Springer-Verlag, 2011.
- [25] G. Burel. A Shallow Embedding of Resolution and Superposition Proofs into the lambda-Pi-Calculus Modulo. In J. Blanchette and J. Urban, editors, *Proceedings of the 3rd International Workshop on Proof Exchange for Theorem Proving*, number 14 in EasyChair Proceedings in Computing, pages 43–57, 2013.
- [26] K. Claessen, A. Lilliestrom, and N. Smallbone. Sort It Out with Monotonicity - Translating between Many-Sorted and Unsorted First-Order Logic. In N. Bjorner and V. Sofronie-Stokkermans, editors, *Proceedings of the 23rd International Conference on Automated Deduction*, number 6803 in Lecture Notes in Artificial Intelligence, pages 207–221. Springer-Verlag, 2011.
- [27] L. de Moura and N. Bjorner. 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.
- [28] H. de Nivelle. Classical Logic with Partial Functions. *Journal of Automated Reasoning*, 47(4):399–425, 2011.
- [29] H. de Nivelle. Theorem Proving for Classical Logic with Partial Functions by Reduction to Kleene Logic. *Journal of Logic and Computation*, page To appear, 2014.

- [30] G. Dowek. Polarized Resolution Modulo. In C. Calude and V. Sassone, editors, *Theoretical Computer Science*, IFIP Advances in Information and Communication Technology, pages 182–196. Springer-Verlag, 2010.
- [31] 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.
- [32] 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.
- [33] H. Ganzinger and K. Korovin. Integrating Equational Reasoning into Instantiation-Based Theorem Proving. In J. Marcinkowski and A. Tarlecki, editors, *Proceedings of the 18th International Workshop on Computer Science Logic, 13th Annual Conference of the EACSL*, number 3210 in Lecture Notes in Computer Science, pages 71–84. Springer-Verlag, 2004.
- [34] M. Greiner and M. Schramm. A Probabilistic Stopping Criterion for the Evaluation of Benchmarks. Technical Report 19638, Institut für Informatik, Technische Universität München, München, Germany, 1996.
- [35] K. Hoder, Z. Khasidashvili, K. Korovin, and A. Voronkov. Preprocessing Techniques for First-Order Clausification. In G. Cabodi and S. Singh, editors, *Proceedings of the Formal Methods in Computer-Aided Design 2012*, pages 44–51. IEEE Press, 2012.
- [36] K. Hoder, L. Kovacs, and A. Voronkov. Interpolation and Symbol Elimination in Vampire. In J. Giesl and R. Haehnle, editors, *Proceedings of the 5th International Joint Conference on Automated Reasoning*, number 6173 in Lecture Notes in Artificial Intelligence, pages 188–195, 2010.
- [37] K. Hoder, L. Kovacs, and A. Voronkov. Invariant Generation in Vampire. In P. Abdulla and R. Leino, editors, *Proceedings of the 17th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 6605 in Lecture Notes in Computer Science, pages 60–64. Springer-Verlag, 2011.
- [38] 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.
- [39] C. Kaliszyk and J. Urban. Automated Reasoning Service for HOL Light. In J. Carette, D. Aspinall, C. Lange, P. Sojka, and W. Windsteiger, editors, *Intelligent Computer Mathematics*, number 7961 in Lecture Notes in Computer Science, pages 120–135. Springer-Verlag, 2013.
- [40] C. Kaliszyk and J. Urban. Lemma Mining over HOL Light. 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 503–517. Springer-Verlag, 2013.
- [41] C. Kaliszyk and J. Urban. MizAR 40 for Mizar 40. arXiv:1310.2805, 2013.
- [42] C. Kaliszyk and J. Urban. PRoCH: Proof Reconstruction for HOL Light. In M.P. Bonacina, editor, *Proceedings of the 24th International Conference on Automated Deduction*, number 7898 in Lecture Notes in Artificial Intelligence, pages 267–274. Springer-Verlag, 2013.
- [43] C. Kaliszyk and J. Urban. Stronger Automation for FLYSPECK by Feature Weighting and Strategy Evolution. In *Proceedings of the 3rd International Workshop on Proof Exchange for Theorem Proving*, page To appear. EasyChair Proceedings in Computing, 2013.
- [44] C. Kaliszyk, J. Urban, and J. Vyskocil. Machine Learner for Automated Reasoning 0.4 and 0.5. arXiv:1402.2359, 2014.
- [45] K. Korovin. iProver - An Instantiation-Based Theorem Prover for First-order Logic (System Description). In P. Baumgartner, A. Armando, and D. Gilles, editors, *Proceedings of the 4th In-*

- ternational Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 292–298, 2008.
- [46] K. Korovin. Instantiation-Based Reasoning: From Theory to Practice. In R. Schmidt, editor, *Proceedings of the 22nd International Conference on Automated Deduction*, number 5663 in Lecture Notes in Computer Science, pages 163–166. Springer-Verlag, 2009.
- [47] 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.
- [48] K. Korovin. Non-cyclic Sorts for First-order Satisfiability. In P. Fontaine, C. Ringeissen, and R. Schmidt, editors, *Proceedings of the International Symposium on Frontiers of Combining Systems*, number 8152 in Lecture Notes in Computer Science, pages 214–228, 2013.
- [49] K. Korovin and C. Stickel. A Note on Model Representation and Proof Extraction in the First-order Instantiation-based Calculus Inst-Gen. In R. Schmidt and F. Papacchini, editors, *Proceedings of the 19th Automated Reasoning Workshop*, pages 11–12, 2012.
- [50] R. Letz and G. Stenz. System Description: DCTP - A Disconnection Calculus Theorem Prover. In R. Gore, A. Leitsch, and T. Nipkow, editors, *Proceedings of the International Joint Conference on Automated Reasoning*, number 2083 in Lecture Notes in Artificial Intelligence, pages 381–385. Springer-Verlag, 2001.
- [51] B. Loechner. Things to Know When Implementing KBO. *Journal of Automated Reasoning*, 36(4):289–310, 2006.
- [52] B. Loechner. Things to Know When Implementing LBO. *Journal of Artificial Intelligence Tools*, 15(1):53–80, 2006.
- [53] D.W. Loveland. *Automated Theorem Proving : A Logical Basis*. Elsevier Science, 1978.
- [54] 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.
- [55] T. Nipkow. Equational Reasoning in Isabelle. *Science of Computer Programming*, 12(2):123–149, 1989.
- [56] T. Nipkow, L. Paulson, and M. Wenzel. *Isabelle/HOL: A Proof Assistant for Higher-Order Logic*. Number 2283 in Lecture Notes in Computer Science. Springer-Verlag, 2002.
- [57] J. Otten. leanCoP 2.0 and ileancop 1.2: High Performance Lean Theorem Proving in Classical and Intuitionistic Logic. In P. Baumgartner, A. Armando, and D. Gilles, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 283–291, 2008.
- [58] J. Otten. Restricting Backtracking in Connection Calculi. *AI Communications*, 23(2-3):159–182, 2010.
- [59] J. Otten. Implementing Connection Calculi for First-order Modal Logics. In K. Korovin and S. Schulz, editors, *Proceedings of the 9th International Workshop on the Implementation of Logics*, number 22 in EPIc, pages 18–32, 2012.
- [60] J. Otten. MleanCoP: A Connection Prover for First-Order Modal Logic. 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 269–276, 2014.
- [61] J. Otten and W. Bibel. leanCoP: Lean Connection-Based Theorem Proving. *Journal of Symbolic Computation*, 36(1-2):139–161, 2003.
- [62] D. Pastre. Automatic Theorem Proving in Set Theory. *Artificial Intelligence*, 10:1–27, 1978.
- [63] D. Pastre. Muscadet : An Automatic Theorem Proving System using Knowledge and Meta-knowledge in Mathematics. *Artificial Intelligence*, 38:257–318, 1989.
- [64] D. Pastre. Automated Theorem Proving in Mathematics. *Annals of Mathematics and Artificial*

- Intelligence*, 8:425–447, 1993.
- [65] D. Pastre. Muscadet version 2.3 : User’s Manual. <http://www.math-info.univ-paris5.fr/pastre/muscadet/manual-en.ps>, 2001.
  - [66] D. Pastre. Strong and Weak Points of the Muscadet Theorem Prover. *AI Communications*, 15(2-3):147–160, 2002.
  - [67] D. Pastre. Complementarity of a Natural Deduction Knowledge-based Prover and Resolution-based Provers in Automated Theorem Proving. <http://www.math-info.univ-paris5.fr/pastre/compl-NDKB-RB.pdf>, 2007.
  - [68] L. Paulson. A Generic Tableau Prover and its Integration with Isabelle. *Artificial Intelligence*, 5(3):73–87, 1999.
  - [69] 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, pages 1–11, 2010.
  - [70] L.C. Paulson and T. Nipkow. *Isabelle: A Generic Theorem Prover*. Number 828 in Lecture Notes in Computer Science. Springer-Verlag, 1994.
  - [71] V. Prevosto and U. Waldmann. SPASS+T. In G. Sutcliffe, R. Schmidt, and S. Schulz, editors, *Proceedings of the FLoC’06 Workshop on Empirically Successful Computerized Reasoning, 3rd International Joint Conference on Automated Reasoning*, number 192 in CEUR Workshop Proceedings, pages 19–33, 2006.
  - [72] A. Reynolds. *Finite Model Finding in Satisfiability Modulo Theories*. PhD thesis, The University of Iowa, Iowa City, USA, 2013.
  - [73] A. Reynolds, M. Deters, V. Kuncak, C. Barrett, and C. Tinelli. Counterexample Guided Quantifier Instantiation for Synthesis in CVC4. In D. Kroening and C. Pasareanu, editors, *Proceedings of the 27th International Conference on Computer Aided Verification*, Lecture Notes in Computer Science, page To appear. Springer-Verlag, 2015.
  - [74] 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.
  - [75] A. Reynolds, C. Tinelli, A. Goel, and S. Krstic. Finite Model Finding in SMT. In N. Sharygina and H. Veith, editors, *Proceedings of the 25th International Conference on Computer Aided Verification*, number 8044 in Lecture Notes in Computer Science, pages 640–655. Springer-Verlag, 2013.
  - [76] 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.
  - [77] A. Riazanov and A. Voronkov. The Design and Implementation of Vampire. *AI Communications*, 15(2-3):91–110, 2002.
  - [78] P. Rümmer. A Constraint Sequent Calculus for First-Order Logic with Linear Integer Arithmetic. In I. Cervesato, H. Veith, and A. Voronkov, editors, *Proceedings of the 15th International Conference on Logic for Programming Artificial Intelligence and Reasoning*, number 5330 in Lecture Notes in Artificial Intelligence, pages 274–289. Springer-Verlag, 2008.
  - [79] P. Rümmer. E-Matching with Free Variables. In N. Bjorner and A. Voronkov, editors, *Proceedings of the 18th International Conference on Logic for Programming Artificial Intelligence and Reasoning*, number 7180 in Lecture Notes in Artificial Intelligence, pages 359–374. Springer-Verlag, 2012.
  - [80] 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.

- [81] 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.
- [82] 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.
- [83] 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.
- [84] 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.
- [85] G. Sutcliffe. Proceedings of the CADE-16 ATP System Competition. Trento, Italy, 1999.
- [86] G. Sutcliffe. Proceedings of the CADE-17 ATP System Competition. Pittsburgh, USA, 2000.
- [87] G. Sutcliffe. The CADE-16 ATP System Competition. *Journal of Automated Reasoning*, 24(3):371–396, 2000.
- [88] G. Sutcliffe. Proceedings of the IJCAR ATP System Competition. Siena, Italy, 2001.
- [89] G. Sutcliffe. The CADE-17 ATP System Competition. *Journal of Automated Reasoning*, 27(3):227–250, 2001.
- [90] G. Sutcliffe. Proceedings of the CADE-18 ATP System Competition. Copenhagen, Denmark, 2002.
- [91] G. Sutcliffe. Proceedings of the CADE-19 ATP System Competition. Miami, USA, 2003.
- [92] G. Sutcliffe. Proceedings of the 2nd IJCAR ATP System Competition. Cork, Ireland, 2004.
- [93] G. Sutcliffe. Proceedings of the CADE-20 ATP System Competition. Tallinn, Estonia, 2005.
- [94] G. Sutcliffe. The IJCAR-2004 Automated Theorem Proving Competition. *AI Communications*, 18(1):33–40, 2005.
- [95] G. Sutcliffe. Proceedings of the 3rd IJCAR ATP System Competition. Seattle, USA, 2006.
- [96] G. Sutcliffe. The CADE-20 Automated Theorem Proving Competition. *AI Communications*, 19(2):173–181, 2006.
- [97] G. Sutcliffe. Proceedings of the CADE-21 ATP System Competition. Bremen, Germany, 2007.
- [98] G. Sutcliffe. The 3rd IJCAR Automated Theorem Proving Competition. *AI Communications*, 20(2):117–126, 2007.
- [99] G. Sutcliffe. Proceedings of the 4th IJCAR ATP System Competition. Sydney, Australia, 2008.
- [100] G. Sutcliffe. The CADE-21 Automated Theorem Proving System Competition. *AI Communications*, 21(1):71–82, 2008.
- [101] 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.
- [102] G. Sutcliffe. Proceedings of the CADE-22 ATP System Competition. Montreal, Canada, 2009.
- [103] G. Sutcliffe. The 4th IJCAR Automated Theorem Proving System Competition - CASC-J4. *AI Communications*, 22(1):59–72, 2009.
- [104] G. Sutcliffe. The TPTP Problem Library and Associated Infrastructure. The FOF and CNF Parts, v3.5.0. *Journal of Automated Reasoning*, 43(4):337–362, 2009.
- [105] G. Sutcliffe. Proceedings of the 5th IJCAR ATP System Competition. Edinburgh, United

- Kingdom, 2010.
- [106] G. Sutcliffe. The CADE-22 Automated Theorem Proving System Competition - CASC-22. *AI Communications*, 23(1):47–60, 2010.
  - [107] G. Sutcliffe. Proceedings of the CADE-23 ATP System Competition. Wroclaw, Poland, 2011.
  - [108] G. Sutcliffe. The 5th IJCAR Automated Theorem Proving System Competition - CASC-J5. *AI Communications*, 24(1):75–89, 2011.
  - [109] G. Sutcliffe. Proceedings of the 6th IJCAR ATP System Competition. Manchester, England, 2012.
  - [110] G. Sutcliffe. The CADE-23 Automated Theorem Proving System Competition - CASC-23. *AI Communications*, 25(1):49–63, 2012.
  - [111] G. Sutcliffe. Proceedings of the 24th CADE ATP System Competition. Lake Placid, USA, 2013.
  - [112] G. Sutcliffe. The 6th IJCAR Automated Theorem Proving System Competition - CASC-J6. *AI Communications*, 26(2):211–223, 2013.
  - [113] G. Sutcliffe. Proceedings of the 7th IJCAR ATP System Competition. Vienna, Austria, 2014.
  - [114] G. Sutcliffe. The CADE-24 Automated Theorem Proving System Competition - CASC-24. *AI Communications*, 27(4):405–416, 2014.
  - [115] G. Sutcliffe. The 7th IJCAR Automated Theorem Proving System Competition - CASC-J7. *AI Communications*, 28:To appear, 2015.
  - [116] G. Sutcliffe and C. Benzmüller. Automated Reasoning in Higher-Order Logic using the TPTP THF Infrastructure. *Journal of Formalized Reasoning*, 3(1):1–27, 2010.
  - [117] 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, 2006.
  - [118] 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.
  - [119] G. Sutcliffe and C. Suttner. The CADE-18 ATP System Competition. *Journal of Automated Reasoning*, 31(1):23–32, 2003.
  - [120] G. Sutcliffe and C. Suttner. The CADE-19 ATP System Competition. *AI Communications*, 17(3):103–182, 2004.
  - [121] G. Sutcliffe, C. Suttner, and F.J. Pelletier. The IJCAR ATP System Competition. *Journal of Automated Reasoning*, 28(3):307–320, 2002.
  - [122] G. Sutcliffe and C.B. Suttner, editors. *Special Issue: The CADE-13 ATP System Competition*, volume 18, 1997.
  - [123] G. Sutcliffe and C.B. Suttner. The Procedures of the CADE-13 ATP System Competition. *Journal of Automated Reasoning*, 18(2):163–169, 1997.
  - [124] G. Sutcliffe and C.B. Suttner. Proceedings of the CADE-15 ATP System Competition. Lindau, Germany, 1998.
  - [125] G. Sutcliffe and C.B. Suttner. The CADE-15 ATP System Competition. *Journal of Automated Reasoning*, 23(1):1–23, 1999.
  - [126] G. Sutcliffe and C.B. Suttner. Evaluating General Purpose Automated Theorem Proving Systems. *Artificial Intelligence*, 131(1-2):39–54, 2001.
  - [127] C.B. Suttner and G. Sutcliffe. The CADE-14 ATP System Competition. *Journal of Automated Reasoning*, 21(1):99–134, 1998.
  - [128] E. Torlak and D. Jackson. Kodkod: A Relational Model Finder. In O. Grumberg and M. Huth, editors, *Proceedings of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, number 4424 in Lecture Notes in Computer Science, pages 632–647. Springer-Verlag, 2007.

- [129] J. Urban. MaLAREa: a Metasystem for Automated Reasoning in Large Theories. In J. Urban, G. Sutcliffe, and S. Schulz, editors, *Proceedings of the CADE-21 Workshop on Empirically Successful Automated Reasoning in Large Theories*, number 257 in CEUR Workshop Proceedings, pages 45–58, 2007.
- [130] J. Urban. BliStr: The Blind Strategymaker. arXiv:1301.2683, 2013.
- [131] J. Urban, G. Sutcliffe, P. Pudlak, and J. Vyskocil. MaLAREa SG1: Machine Learner for Automated Reasoning with Semantic Guidance. In P. Baumgartner, A. Armando, and D. Gilles, editors, *Proceedings of the 4th International Joint Conference on Automated Reasoning*, number 5195 in Lecture Notes in Artificial Intelligence, pages 441–456. Springer-Verlag, 2008.
- [132] T. Weber. *SAT-based Finite Model Generation for Higher-Order Logic*. PhD thesis, Technische Universität München, Munich, Germany, 2008.
- [133] S. Zimmer. *Intelligent Combination of a First Order Theorem Prover and SMT Procedures*. Master’s thesis, Saarland University, Saarbruecken, Germany, 2007.