The TPTP Format for Clausal Connection Tableaux

Geoff Sutcliffe[1], Sean B. Holden[2] and Mantas Baksys[2]

[1]University of Miami, USA
[2]University of Cambridge, United Kingdom

Abstract

This paper describes the (new) TPTP format for writing clausal connection tableaux. The format builds on the existing infrastructure of the TPTP World, in particular the TPTP format for recording derivations. An ATP system that outputs tableaux in this format is described. Existing TPTP World tools for verifying and viewing derivations have been extended to verify and view tableaux in this format.

Keywords

TPTP, Tableaux, Proof

1. Introduction

Automated Theorem Proving (ATP) [1] is concerned with the development and use of software that automates sound reasoning: the derivation of conclusions that follow inevitably from known facts. ATP is at the heart of many computational tasks, including sensitive tasks such as software/hardware verification [2] and system security [3]. ATP systems are often used as components of more complex Artificial Intelligence (AI) systems, which means that the impact of ATP extends into many facets of society. In many of these applications the use of ATP systems is mission critical, in the sense that incorrect results from ATP might have nasty consequences. Facing the demand for error-free results from ATP systems is the reality that ATP systems are complex pieces of software, implementing complex calculi with complex data structures and algorithms [4]. Despite best intentions and efforts, incorrect results are possible. To counter incorrectness, an ATP system can be required to output a proof that serves as a certificate for the system's claim. To ensure that a proof is correct, proof verification can be required, which serves as a certification (but not a certificate) of the proof.

At one top level, proofs can be divided into two types: Hilbert-style proofs that start at axioms and derive theorems [5], and proofs-by-contradiction that negate the conjecture to be proved and derive a contradiction with the axioms [6]. ATP systems that search for Hilbert-style proofs are often called "natural deduction systems", e.g., THINKER [7] and Muscadet [8]. Most contemporary high-performance ATP systems that search for a proof-by-contradiction, called a refutation in this context, use a saturation based approach [4], e.g., E [9] and Vampire [10]. A complementary approach is taken in systems that build tableaux and closed connection matrices [11], e.g., leanCoP [12] and Connect++ [13]. The TPTP World (see Section 2) has an established format for writing refutations and Hilbert-style proofs [14], and has a tool for verifying proofs in that format (see Section 2.1). One early proposal for a TPTP style format for writing non-clausal tableaux never gained traction [15]. A more recent proposal for a TPTP style format for clausal connection tableaux [16] provided some inspiration for the design of the format described in this paper. This paper describes the (new) TPTP format for writing *clausal connection tableaux*.

This paper is structured as follows: Section 2 introduces the TPTP World, providing the necessary background to the TPTP language, explains the TPTP format for recording (non-tableau) derivations, and describes the GDV and IDV tools for verifying and viewing derivations. Section 3 describes the TPTP format for writing clausal connection tableaux, and explains how the format meets the requirements for

The 10th Workshop on Practical Aspects of Automated Reasoning, as part of the Federated Logic Conference (FLoC) 2026, Lisbon, Portugal, July 25, 2026

geoff@cs.miami.edu (G. Sutcliffe); sbh11@cl.cam.ac.uk (S. B. Holden); mb2412@cam.ac.uk (M. Baksys)
0000-0001-7116-9338 (G. Sutcliffe); 0000-0001-7979-1148 (S. B. Holden); 0000-0001-9532-1007 (M. Baksys)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-08-20

CEUR-WS.org/Vol-4241/short2.pdf

easy reconstruction of the tableau, and for semantic verification of the tableau inference steps. Section 4 describes an ATP system that can output clausal connection tableaux in the TPTP format, and explains how the GDV and IDV tools can be used to verify and view such tableaux. Section 5 concludes.

2. The TPTP World

The TPTP World [17] is the well-established infrastructure that supports research, development, and deployment of Automated Theorem Proving (ATP) systems. The TPTP World infrastructure includes the TPTP language [14], the TPTP problem library [18], the TSTP solution library [19], the SZS ontologies [20], the Specialist Problem Classes (SPCs) and problem difficulty ratings [21], SystemOnTPTP [22] and StarExec [23], and the CADE ATP System Competition (CASC) [24]. The problem library is a large collection of Thousands of Problems for Theorem Proving – hence the name. The problem library contains over 25000 problems from over 50 different domains, written in the TPTP language. The problems are categorized into Specialist Problem Classes according to their syntactic and logical status [20]. The TSTP solution library is the result of running numerous ATP systems on that library and collecting their output. The solutions are categorized according to their logical and output form [20]. The TPTP and TSTP libraries, with their categorizations, provide the basis for assigning a difficulty rating to each problem, according to the number of ATP systems that are able to solve it [21].

The most salient feature of the TPTP World for this work is the TPTP language. The TPTP language [25] is one of the keys to the success of the TPTP World. The TPTP language is used for writing both problems and solutions, which enables convenient communication between ATP systems and tools. Originally the TPTP World supported only first-order clause normal form (CNF) [26]. Over the years full first-order form (FOF) [18], typed first-order form (TFF) [27, 28], typed extended first-order form (TXF) [29], typed higher-order form (THF) [30, 31], and non-classical forms (NTF) [32] have been added. A general principle of the TPTP language is: "We provide the syntax, you provide the semantics". As such, there is no a priori commitment to any semantics for each of the language forms, although in almost all cases the intended logic and semantics are well known.

Problems and solutions are built from *annotated formulae* of the form:

language(*name*, *role*, *formula*, *source*, *useful_info*)

The *language*s supported are `cnf` (clause normal form), `fof` (first-order form), `tff` (typed first-order form), and `thf` (typed higher-order form). The *role*, e.g., `axiom`, `lemma`, `conjecture`, defines the use of the formula. In a *formula*, terms and atoms follow Prolog conventions – functions and predicates start with a lowercase letter or are 'single quoted', and variables start with an uppercase letter. The language also supports interpreted symbols that either start with a `$`, e.g., the truth constants `$true`, `$false`, the individual and boolean types `$i`, `$o`, or are composed of non-alphabetic characters, e.g., integer/rational/real numbers such as 27, 43/92, -99.66. The logical connectives in the TPTP language are `!>`, `?*`, `@+`, `@-`, `!`, `?`, `~`, `|`, `&`, `=>`, `<=`, `<=>`, and `<~>`, for the mathematical connectives Π, Σ, choice (indefinite description), definite description, $\forall$, $\exists$, $\neg$, $\vee$, $\wedge$, $\Rightarrow$, $\Leftarrow$, $\Leftrightarrow$, and $\oplus$ respectively. Equality and inequality are expressed as the infix operators `=` and `!=`. The *source* and *useful_info* are optional. Figure 1 shows an example first-order form problem.

```
%----------------------------------------------------------------------
fof(a1,axiom,         ! [Y] : ~( ~q(Y) & ? [X] : s(X) ) ).
fof(a2,axiom,         ( r & ? [Z] : q(Z) ) => ! [X] : ~p(X) ).
fof(a3,axiom,         p(c) | ( ~q(c) & q(a) ) ).
fof(a4,axiom,         ~q(c) => ! [W] : ~q(W) ).
fof(a5,axiom,         p(c) => r ).
fof(prove,conjecture, ! [X] : ( ~s(X) & ~q(b) & p(c) ) ).
%----------------------------------------------------------------------
```

Figure 1: An example first-order form problem in TPTP format

2.1. The TPTP Format for Derivations

A derivation written in the TPTP language is a list of annotated formulae. The leaves of a derivation typically have the role `axiom` or `conjecture`, and the inferred formulae typically have the role `plain` or `negated_conjecture` (also used for leaves in CNF). The source is either a `file` record for leaves or an `inference` record for inferred formulae. A `file` record contains the problem file name and the corresponding annotated formulae name in the problem file. An `inference` record contains the inference rule name, a list of useful inference information, and a list of the inference parents. The inference parents can be annotated formulae names, and nested `inference` records. Common types of useful inference information are the semantic relationship of the inferred formula to its parents as an SZS ontology value [20] in a `status` record, special information about recognized types of complex inference rules, e.g., Skolemization and splitting, and details of new symbols introduced in the inference. The use of SZS values is core to GDV's approach to verification, described below.

Figure 2 shows the transformation to CNF of the problem in Figure 1, which is used in the refutation discussed here, and the tableau proof in Section 3. Points of note:

- The leaf formulae are `a1`-`a5` and `prove`. In this example they are copies of the corresponding problem formulae.
- Many of the formulae are inferred with SZS status `thm`, i.e., they are logical consequences of their parents, e.g., clause `c1` is a logical consequence of `a1`.
- The negated conjecture `nc1` is inferred with SZS status `cth` – its negation is a logical consequence of its parent `prove`.
- The Skolemized formula `nc3` is inferred with SZS status `esa` – it is equisatisfiable with its parent `nc2`.

Figure 3 shows the final steps of the refutation found by E 3.2.5 for the problem in Figure 1, following the transformation to CNF shown in Figure 2.[1] Points of note:

- Many of the formulae are inferred with SZS status `thm`, i.e., they are logical consequences of their parents, e.g., `c_0_26` is a logical consequence of `c_0_23` and `c_0_24`.
- Several of the inferred formulae have nested `inference` records – the intermediate inferred formulae have not been output, e.g., in `c_0_23` the intermediate result of the `spm` inference between `c_0_18` and `c_0_19` is not shown, and only the subsequent final result of the `csr` inference between the intermediate result and `c_0_20` is given.

2.2. TPTP World Tools for Verifying and Viewing Derivations

The GDV derivation verifier [33] verifies proofs in the TPTP format. GDV's input is a TPTP format proof, and optionally (required for complete verification) the problem for which the proof was produced. GDV verifies a proof in four phases: structural verification, leaf verification, rule-specific verification, and inference verification. Many of the checks rely on a "check-by-ATP", which calls a trusted ATP system – either a theorem prover or a model finder (the systems have become trusted through the process described in [34]). Structural verification deals with non-logical aspects of a proof, checking whether the annotated formulae of the proof have the right format and relationships, e.g., the derivation is acyclic. Leaf verification deals with the leaves of the derivation, and their relationship with the problem annotated formulae, e.g., the leaves are copies of problem formulae or can be proved from problem formulae using a check-by-ATP. Rule specific verification deals with inference rules that require special treatment, e.g., Skolemization. Inference verification deals with the various types of inferences that are made by ATP systems in a proof. Examples of checks are: for inference steps with SZS status `thm` check-by-ATP that the inferred formula can be proved from the parent formulae, for inference steps with SZS status `cth` check-by-ATP that the negation of the inferred formula can be proved from the parent formulae. GDV is available online in SystemOnTSTP.[2]

The Interactive Derivation Viewer (IDV) [35] provides a graphical rendering of a proof DAG. It has features that allow the user to examine the proof structure in various ways, including identification of "interesting" steps in the proof [36]. Figure 4 shows the IDV rendering of (the full version of) the refutation in Figure 3. IDV is available online in SystemOnTSTP.[3]

[1] The full refutation can be generated in SystemOnTPTP, available at `tptp.org/cgi-bin/SystemOnTPTP`
[2] Available at `tptp.org/cgi-bin/SystemOnTSTP`
[3] Available at `tptp.org/cgi-bin/SystemOnTSTP`

```
%----------------------------------------------------------------------
fof(a1,axiom,               ! [Y] : ~( ~q(Y) & ? [X] : s(X) ),
    file('PaperFOF.p',a1) ).
fof(a2,axiom,               ( ( r & ? [Z] : q(Z) ) => ! [X] : ~p(X) ),
    file('PaperFOF.p',a2) ).
fof(a3,axiom,               ( p(c) | ( ~q(c) & q(a) ) ),
    file('PaperFOF.p',a3) ).
fof(a4,axiom,               ~q(c) => ! [W] : ~q(W),
    file('PaperFOF.p',a4) ).
fof(a5,axiom,               ( p(c) => r ),
    file('PaperFOF.p',a5) ).
fof(prove,conjecture,       ! [X] : ( ~s(X) & ~q(b) & p(c) ),
    file('PaperFOF.p',prove) ).

fof(nc1,negated_conjecture, ~! [X] : ( ~s(X) & ~q(b) & p(c) ),
    inference(negate,[status(cth)],[prove]) ).
fof(nc2,negated_conjecture, ? [X] : ~( ~s(X) & ~q(b) & p(c) ),
    inference(negate,[status(thm)],[nc1]) ).
fof(nc3,negated_conjecture, ~( ~s(sK1) & ~q(b) & p(c) ),
    inference(skolemize,[status(esa),new_symbols(skolem,[sK1]),
skolemized(X)],[nc2]) ).

cnf(c1,plain,               ( q(Y) | ~s(X) ),
    inference(clausify,[status(thm)],[a1]) ).
cnf(c2,plain,               ( ~q(Z) | ~p(X) | ~r ),
    inference(clausify,[status(thm)],[a2]) ).
cnf(c3,plain,               ( p(c) | ~q(c) ),
    inference(clausify,[status(thm)],[a3]) ).
cnf(c4,plain,               ( p(c) | q(a) ),
    inference(clausify,[status(thm)],[a3]) ).
cnf(c5,plain,               ( q(c) | ~q(W) ),
    inference(clausify,[status(thm)],[a4]) ).
cnf(c6,plain,               ( r | ~p(c) ),
    inference(clausify,[status(thm)],[a5]) ).
cnf(c7,negated_conjecture,  ( s(sK1) | q(b) | ~p(c) ),
    inference(clausify,[status(thm)],[nc3]) ).
%----------------------------------------------------------------------
```

Figure 2: Clausification for CNF-based proofs of the problem in Figure 1

3. The (new) TPTP Format for Clausal Connection Tableaux

Three primary requirements were observed for the design of the TPTP format for clausal connection tableaux:

1. Easy reconstruction of the tableau.
2. Sufficient information for structural verification of the closed tableau.
3. Sufficient information for semantic verification of the inference steps.

Additional beneficial features, adopted from [16], are:

4. Concise and simple enough for a natural representation of proofs.
5. Readable by humans as well as ATP tools.

One early proposal for a TPTP style format for non-clausal tableaux [15] met requirements 1 and 3, but fell short of requirement 2 and features 4 and 5 because it attempted to "squeeze" [16] connection proofs into a derivation style format, with a clumsy encoding of parent formulae details. It also did not provide an explicit way of recording lemmas. Another existing TPTP style format is the leanTPTP format for clausal connection tableaux that is output by the leanCoP ATP system. [37] The leanTPTP format definitely meets requirement 4, but omits the information required for requirements 1, 2, and 3. Most notably, the leanTPTP format does not record the tableau tree structure, does not explicitly close the tableau branches, and does not appear to have any way of recording lemmas. A more recent proposal for a TPTP style format for clausal connection tableaux [16]

```
%--------------------------------------------------------------------------
cnf(c_0_24,plain,               ( ~r | ~q(X1) ),
    inference(csr,[status(thm)],
[inference(spm,[status(thm)],[c_0_18,c_0_19]),c_0_20]) ).

cnf(c_0_25,negated_conjecture,  ( q(a) | q(b) ),
    inference(spm,[status(thm)],[c_0_21,c_0_22]) ).

cnf(c_0_26,plain,               ( r | ~p(c) ),
    inference(split_conjunct,[status(thm)],[c_0_23]) ).

cnf(c_0_27,negated_conjecture,  ~r,
    inference(csr,[status(thm)],
[inference(spm,[status(thm)],[c_0_24,c_0_25]),c_0_24]) ).

cnf(c_0_28,plain,               q(a),
    inference(sr,[status(thm)],
[inference(spm,[status(thm)],[c_0_26,c_0_22]),c_0_27]) ).

cnf(c_0_29,plain,               ( r | ~q(c) ),
    inference(spm,[status(thm)],[c_0_26,c_0_19]) ).

cnf(c_0_30,plain,               q(c),
    inference(spm,[status(thm)],[c_0_20,c_0_28]) ).

cnf(c_0_31,plain,               $false,
    inference(sr,[status(thm)],[inference(cn,[status(thm)],
[inference(rw,[status(thm)],[c_0_29,c_0_30])]),c_0_27]),
    [proof] ).
%--------------------------------------------------------------------------
```

Figure 3: The final steps of E 3.2.5's refutation of Figure 1

has weaknesses akin to the leanTPTP format. The (new) TPTP format for writing clausal connection tableaux meets all the requirements and provides both the beneficial features.

Figure 5 shows the clausal connection tableau for the clauses in Figure 2. The dotted arrow shows the one reduction step, the solid rightward arrows point to the lemmas that are created, and the dashed arrows show where the lemmas are used. Note that lemmas can be used only below the parent node of where they are created. The tableau in Figure 5 has the variables instantiated as they would be when the tableau is closed (but note that in general a tableau need not be ground). Figure 6 shows the TPTP format for recording the tableau in Figure 5. The TPTP format for clausal connection tableaux recognizes six inference rules, which are described below. The inference record of each annotated formula records the name of the rule used, the SZS status of the inferred formula with respect to its parents, the tableau parent node, and the inference parents. The SZS status and inference parent information makes it possible to use semantic verification of each inference, and the tableau parent information makes it easy to reconstruct the tableau. The labels in square brackets in Figure 5 identify the literals in the annotated formulae in Figure 6.

The Inference Rules

`start`: The initial clause below the root node. For example, in Figure 6 `t1` starts the tableau. The parent is recorded as `0:0`, indicating that the node above is the root node. The logical parent of `t1` is recorded as `[c1]`. The inference has status `thm`, i.e., `t1` is a logical consequence of its parent. `start` can be viewed as a special form of `extension`, described next.

`extension`: The standard tableau extension rule. For example, in Figure 6 `t2` extends from the first literal `q(Y)` of `t1` to the 1st literal `~q(Y)` of `c2`. The parent is recorded as `t1:1`. The logical parent of `t2` is recorded as `[c2]`. The inference has status `thm`, i.e., `t2` is a logical consequence of its parent.

`connection`: Explicitly close the branch of the contradiction of an extension. For example, in Figure 6 `t3` closes the branch of the contradiction between `q(Y)` and `~q(Y)` in the extension to `t2`. The parent is recorded as

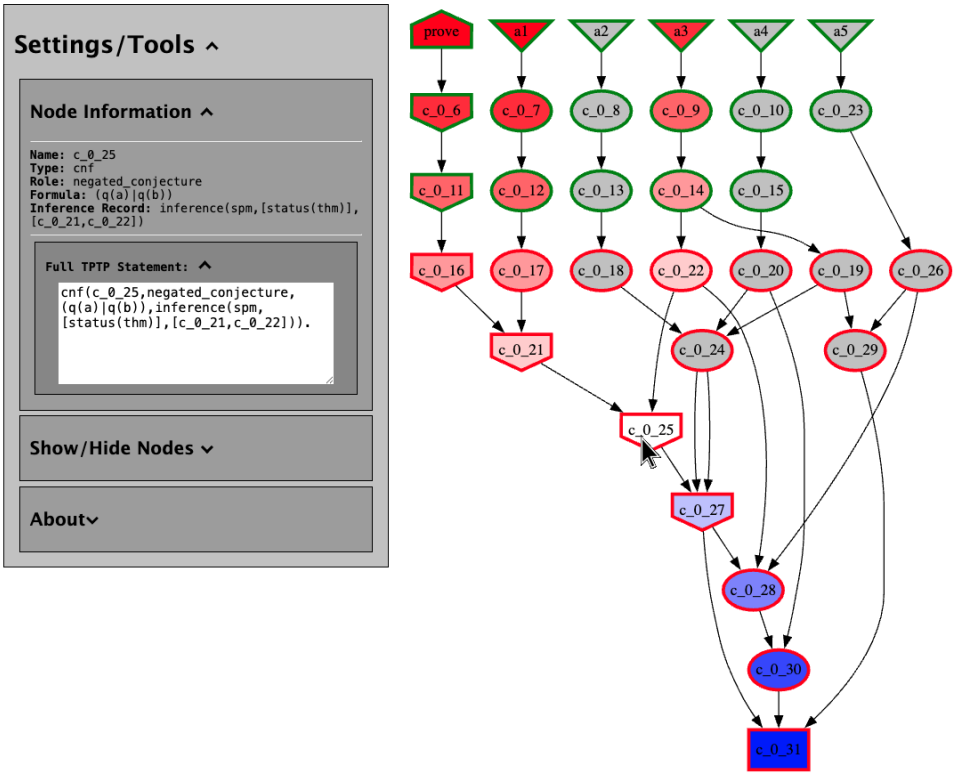

Figure 4: IDV rendering of the (full version of the) refutation in Figure 3

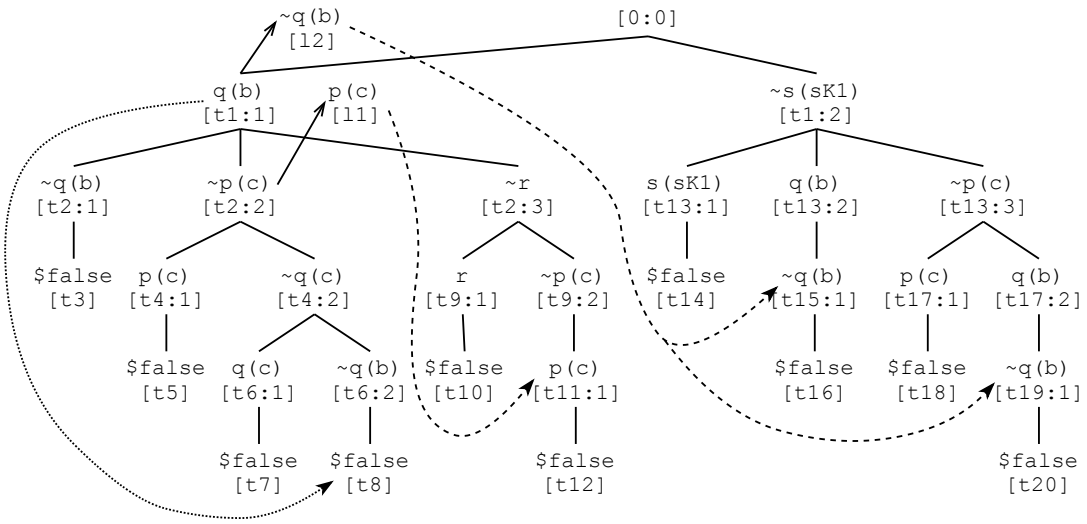

Figure 5: A tableau for the problem in Figure 1

`t2:1`. The logical parents of `t3` are recorded as `[t2:1,t1:1]`, meaning the 1st literal of `t2` and the 1st literal of `t1`. The inference has status `thm`, i.e., `t3` is a logical consequence of its parents. `connection` can be viewed as a degenerate form of `reduction`, described next.

`reduction`: The standard tableau reduction rule. For example, in Figure 6 `t8` closes the branch of the contradiction between the 2nd literal `~q(Y)` of `t6` and the 1st literal `q(Y)` of `t1`. In Figure 5 this is denoted by the dotted arrow from `q(Y)` to `t8`. The parent is recorded as `t6:2`. The logical parents of `t8` are recorded as `[t6:2,t1:1]`, meaning the 2nd literal of `t6` and the 1st literal of `t1`. The inference has status `thm`, i.e., `t8` is a logical consequence of its parents.

```
%----------------------------------------------------------------------------------------------
cnf(t1,plain,                q(b) | ~s(sK1),
    inference(start,[status(thm),parent(0:0)],[c1]) ).

cnf(t2,plain,                ~q(b) | ~p(c) | ~r,
    inference(extension,[status(thm),parent(t1:1)],[c2]) ).

cnf(t3,plain,                $false,
    inference(connection,[status(thm),parent(t2:1)],[t2:1,t1:1]) ).

cnf(t4,plain,                p(c) | ~q(c),
    inference(extension,[status(thm),parent(t2:2)],[c3]) ).

cnf(t5,plain,                $false,
    inference(connection,[status(thm),parent(t4:1)],[t4:1,t2:2]) ).

cnf(t6,plain,                q(c) | ~q(b),
    inference(extension,[status(thm),parent(t4:2)],[c5]) ).

cnf(t7,plain,                $false,
    inference(connection,[status(thm),parent(t6:1)],[t6:1,t4:2]) ).

cnf(t8,plain,                $false,
    inference(reduction,[status(thm),parent(t6:2)],[t6:2,t1:1]) ).

cnf(l1,lemma,                p(c),
    inference(lemma,[status(cth),parent(t2:2),below(t1:1)],[t2:2]) ).

cnf(t9,plain,                r | ~p(c),
    inference(extension,[status(thm),parent(t2:3)],[c6]) ).

cnf(t10,plain,               $false,
    inference(connection,[status(thm),parent(t9:1)],[t9:1,t2:3]) ).

cnf(t11,plain,               p(c),
    inference(lemma_extension,[status(thm),parent(t9:2)],[l1:1]) ).

cnf(t12,plain,               $false,
    inference(connection,[status(thm),parent(t11:1)],[t9:2,t11:1]) ).

cnf(l2,lemma,                ~q(b),
    inference(lemma,[status(cth),parent(t1:1),below(0:0)],[t1:1]) ).

cnf(t13,plain,               s(sK1) | q(b) | ~p(c),
    inference(extension,[status(thm),parent(t1:2)],[c7]) ).

cnf(t14,plain,               $false,
    inference(connection,[status(thm),parent(t13:1)],[t13:1,t1:2]) ).

cnf(t15,plain,               ~q(b),
    inference(lemma_extension,[status(thm),parent(t13:2)],[l2:1]) ).

cnf(t16,plain,               $false,
    inference(connection,[status(thm),parent(t15:1)],[t15:1,t13:2]) ).

cnf(t17,plain,               p(c) | q(b),
    inference(extension,[status(thm),parent(t13:3)],[c4]) ).

cnf(t18,plain,               $false,
    inference(connection,[status(thm),parent(t17:1)],[t17:1,t13:3]) ).

cnf(t19,plain,               ~q(b),
    inference(lemma_extension,[status(thm),parent(t17:2)],[l2:1]) ).

cnf(t20,plain,               $false,
    inference(connection,[status(thm),parent(t19:1)],[t19:1,t17:2]) ).
%----------------------------------------------------------------------------------------------
```

Figure 6: A tableau proof in TPTP format, for the problem in Figure 1

`lemma`**:** The creation of a unit lemma when a branch is closed. For example, in Figure 6 `l1` is the lemma `p(c)` created when the branch rooted at `t2:2` is closed. The lemma is the negation of the 2nd literal `~p(c)` of `t2`. The role of the annotated formula records it as a `lemma`. The parent is recorded as `t2:2`. In Figure 5 this is denoted by the solid arrow from `t2:2` to `l1`. As the node `t1:1` is used in a reduction in closing the branch, the lemma is available only below `t1:1`, recorded as `below(t1:1)`. The logical parent of `l1` is recorded as `[t2:2]`, meaning the 2nd literal `~p(c)` of `t2`. The inference has status `cth`, i.e., the negation of `l1` is a logical consequence of its parent.

`lemma_extension`**:** Use of a lemma to close a branch. For example, in Figure 6 `t11:1` is the lemma `l1`, and `t12` is the `connection` that closes the branch down to `t11:1`. In Figure 5 this is denoted by the dashed arrow from `l1` to `t11:1`. The lemma `l1` can be used here because `t9:2` is below `t1:1`. The parent is recorded as `t11:1`. The logical parent of `t11` is recorded as `[l1:1]`, meaning the 1st literal `p(c)` of `l1`. The inference has status `thm`, i.e., `t11:1` is a logical consequence of its parent. Lemmas can be used in multiple lemma extensions.

A new copy of the lemma is used each time, which provides fresh variables in the lemma. For example, in Figure 6 `12` is used in the lemma extension `t15`, after which the variable `Y` is instantiated to `b` in the connection `t16`. The lemma is used again in the lemma extension `t19`, after which the variable `Y` is instantiated to `a` in the connection `t20`. `lemma_extension` can be viewed as a special form of `extension`.

Additional tableau inference rules, e.g., factorization [38] (another view of lemma generation and use), non-unit lemma generation (also known as bottom-up lemma generation) [39, 40], lazy paramodulation [41] (to deal with equality), etc., can easily be added to this format.

The point at which a variable becomes instantiated can optionally be recorded with a `bind()` record, e.g.,

```
cnf(t4,plain,              p(c) | ~q(c),
    inference(extension,
              [status(thm),parent(t2:2),bind(X,$fot(c))],[c3]) ).
```

records that the variable `X` in `c3` is bound to the first-order term `c` (assuming variables have been renamed apart so there is no ambiguity about which `X` is bound).

4. ATP Systems and Tools

Connect++ is an ATP for first-order logic with equality, using the connection calculus to construct proofs. [13]. Connect++ implements most of the search methods and other heuristics developed for `leanCoP` versions 2 and later, including restricted backtracking (and an alternative version of backtracking restriction for extensions), iterative deepening by path length, various approaches to start clause selection, definitional clause conversion, deterministic or random re-ordering, and regularity testing. It also incorporates some further options such as miniscoping and polynomial-time unification. It accepts input in the TPTP `cnf` and `fof` formats. Connect++ incorporates a standard schedule similar to that of `leanCoP`, but also allows arbitrary schedules to be specified in a simple language. It produces certificates for proofs in a very simple format described in [13, 16]; these can be verified internally, or output to a file and checked independently by a short Prolog program. From version 0.7.0 Connect++ can also produce closed tableau in the TPTP format. Connect++ is implemented in C++ and freely available under the GNU General Public License (GPL) Version 3.[4]

The extension of GDV to verify tableau proofs requires extending the structural verification phase[5], and a minor change to inference verification. Leaf verification remains unchanged, and the rule-specific steps of derivation verification are naturally inapplicable. Structural verification of a tableau proof requires checking:

- All inference parents exist, in particular that the specified literals exist.
- The `parent` records define an acyclic DAG, without contradictions in the branches.
- All branches start at the root `0:0` node.
- The tableau is closed, i.e., every branch ends at a `$false` formula.
- All `lemma_extension` steps use lemmas that are below the specified node in the tableau.

Semantic verification is essentially the same as for derivations, with the small additional need to extract the specified literals from inference parents.

The IDV derivation viewer has been adapted to display tableau proofs, as the Interactive Tableau Viewer (ITV). Figure 7 shows the ITV rendering of the tableau in Figure 6. The input formulae and the steps of conversion to CNF are shown above the tableau, in the same format used for derivations in IDV (see Section 2.2). The root node of the tableau is the top `$true` box. Nodes from `extension` steps are in green-bordered ovals, `$false` nodes from `connection` and `reduction` steps are in red-bordered `$false` boxes, lemmas created in `lemma` steps are in blue-bordered triangles, and lemmas used in `lemma_extension` steps are in green-bordered triangles. When the cursor hovers over a node in the tableau its descendants in the tableau are highlighted in blue, its ancestors in the tableau are highlighted in red, and additionally its source in the input formula is also highlighted in red. This is shown in Figure 7 with the cursor hovering over `~q(c)` in the tableau. When the cursor is hovered over a `reduction` or `lemma_extension` node in the tableau, the relevant ancestor/lemma lights up in bright green. These browsing features provide useful support for understanding a tableau proof. ITV is available online in SystemOnTSTP.[6]

[4]Download from www.cl.cam.ac.uk/~sbh11/connect++.html
[5]This ongoing work that will be completed in June 2026
[6]Available at tptp.org/cgi-bin/SystemOnTSTP

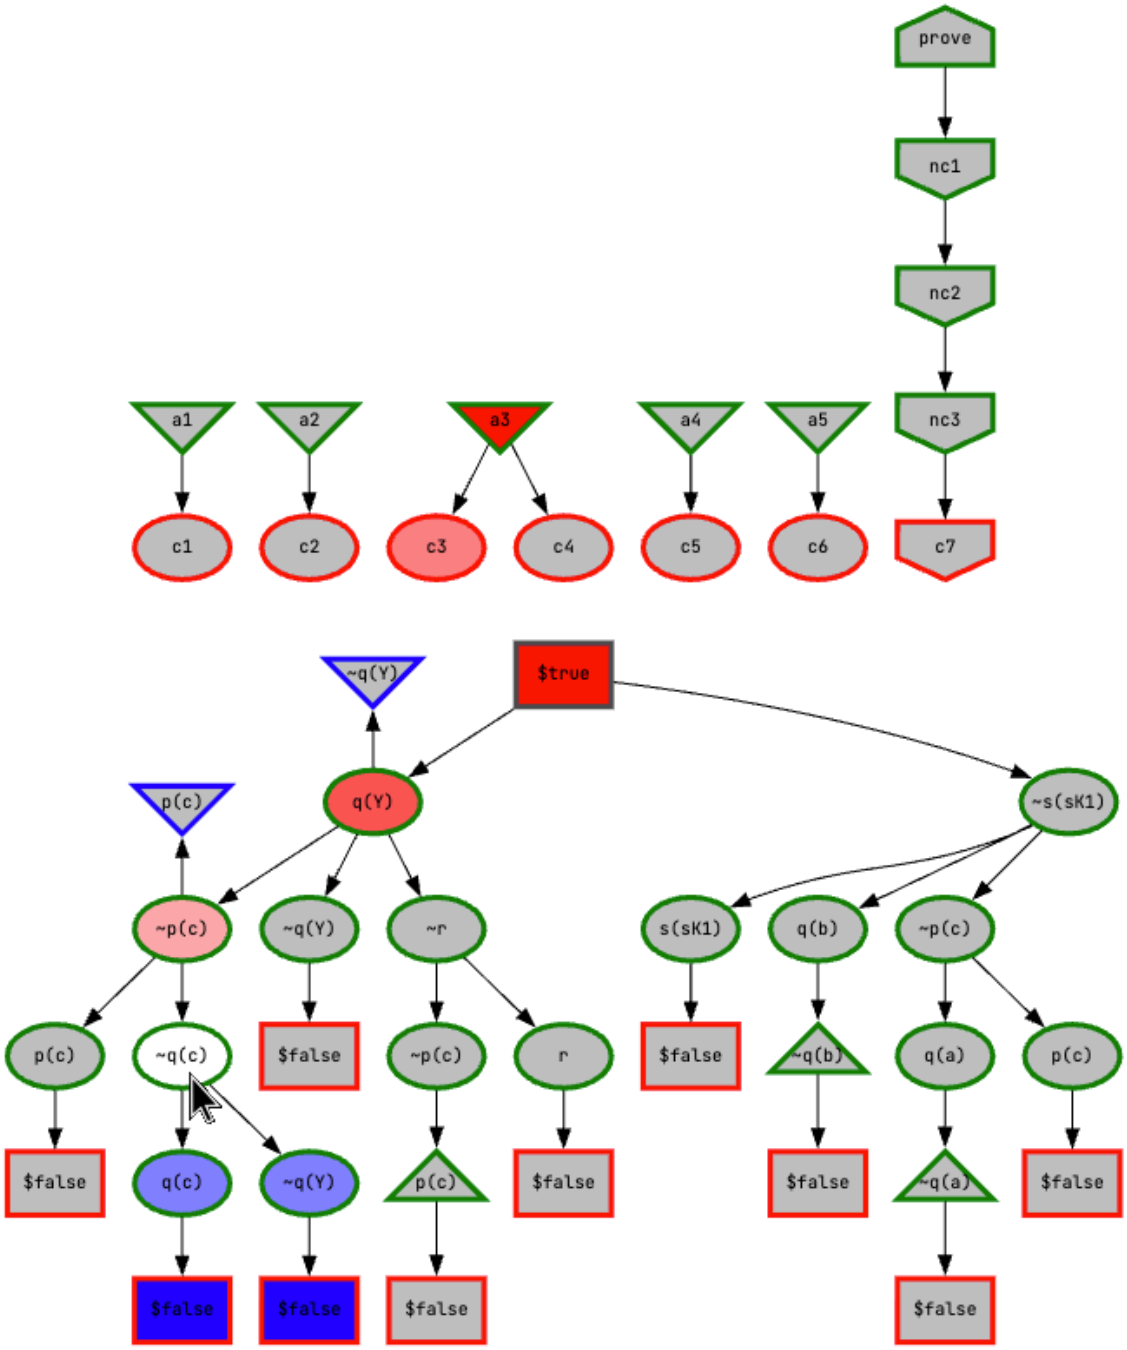

Figure 7: ITV's rendering of the tableau in Figure 6

5. Conclusion

This paper has described the TPTP format for writing clausal connection tableaux. The format builds on the existing infrastructure of the TPTP World, in particular the TPTP format for writing derivations. The format meets the requirements and beneficial features listed in Section 3:

1. Easy reconstruction of the tableau - this can be done directly from the `parents`.

2. Sufficient information for structural verification of the closed tableau - this is provided by the `parents`, the `below` record in lemmas, and the explicit closing of each branch with `connection` steps and after `reduction` steps.
3. Sufficient information for semantic verification of the inference steps - this is provided with SZS `status` information.
4. Concise and simple enough for a natural representation of proofs - there is no redundant information or unnecessary detail.
5. Readable by humans as well as ATP tools - the format uses the established TPTP syntax, which is both human and machine readable.

The format has been adopted in the Connect++ ATP system, and the existing GDV and IDV tools for verifying and viewing derivations have been adapted to verifying and viewing tableaux.

Short-term future work includes completing structural verification in GDV. Medium term efforts will focus on convincing developers of ATP systems that produce clausal connection tableau to output their tableaux in the TPTP format. In the long-term the format might be extended to record non-connection [11, 42] and non-clausal [43, 44] tableaux.

Acknowledgments:

- For the purpose of open access, the second author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.
- ITV was implemented at the University of Miami by Daniel Li and Esteban Morales, building on the IDV code of Jack McKeown.

Declaration on Generative AI: We didn't use any.

References

[1] A. Robinson, A. Voronkov, Handbook of Automated Reasoning, Elsevier Science, Amsterdam, The Netherlands, 2001.
[2] R. Hähnle, M. Huisman, Deductive Software Verification: From Pen-and-Paper Proofs to Industrial Tools, in: B. Steffen, G. Woeginger (Eds.), Computing and Software Science: State of the Art and Perspectives, number 10000 in Lecture Notes in Computer Science, Springer, Berlin, Germany, 2019, pp. 345–373.
[3] B. Cook, Formal Reasoning About the Security of Amazon Web Services, in: H. Chockler, G. Weissenbacher (Eds.), Proceedings of the 30th International Conference on Computer Aided Verification, number 10981 in Lecture Notes in Computer Science, Springer, Berlin, Germany, 2018, pp. 38–47.
[4] S. Schulz, Algorithms and Data Structures for First-Order Equational Deduction, in: C. Benzmüller, B. Fischer, G. Sutcliffe (Eds.), Proceedings of the 6th International Workshop on the Implementation of Logics, number 212 in CEUR Workshop Proceedings, 2006, pp. 1–6.
[5] H. Enderton, A Mathematical Introduction to Logic, Academic Press, 1972.
[6] E. Mendelson, Introduction to Mathematical Logic, 3 ed., Wadsworth and Brooks/Cole, 1987.
[7] F. Pelletier, Automated Natural Deduction in THINKER, Studia Logica 60 (1998) 3–43.
[8] D. Pastre, Muscadet 2.3 : A Knowledge-based Theorem Prover based on Natural Deduction, in: R. Gore, A. Leitsch, T. Nipkow (Eds.), Proceedings of the International Joint Conference on Automated Reasoning, number 2083 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2001, pp. 685–689.
[9] S. Schulz, S. Cruanes, P. Vukmirović, Faster, Higher, Stronger: E 2.3, in: P. Fontaine (Ed.), Proceedings of the 27th International Conference on Automated Deduction, number 11716 in Lecture Notes in Computer Science, Springer, Berlin, Germany, 2019, pp. 495–507.
[10] L. Kovács, A. Voronkov, First-Order Theorem Proving and Vampire, in: N. Sharygina, H. Veith (Eds.), Proceedings of the 25th International Conference on Computer Aided Verification, number 8044 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2013, pp. 1–35.
[11] U. Furbach, B. Beckert, R. Hähnle, R. Letz, P. Baumgartner, U. Egly, W. Bibel, S. Brüning, J. Otten, T. Rath, T. Schaub, Tableau and Connection Calculi, in: W. Bibel, P. Schmitt (Eds.), Automated Deduction - A Basis for Applications, Volume 1: Foundations - Calculi and Methods, Kluwer, 1998, pp. 3–179.
[12] J. Otten, 20 Years of leanCoP - An Overview of the Provers, in: J. Otten, W. Bibel (Eds.), Proceedings of the 1st International Workshop on Automated Reasoning with Connection Calculi, number 3613 in CEUR Workshop Proceedings, 2023, pp. 4–22.

[13] S. Holden, Connect++: A New Automated Theorem Prover Based on the Connection Calculus, in: J. Otten, W. Bibel (Eds.), Proceedings of the 1st International Workshop on Automated Reasoning with Connection Calculi, number 3613 in CEUR Workshop Proceedings, 2023, pp. 95–106.
[14] G. Sutcliffe, S. Schulz, K. Claessen, A. Van Gelder, Using the TPTP Language for Writing Derivations and Finite Interpretations, in: U. Furbach, N. Shankar (Eds.), Proceedings of the 3rd International Joint Conference on Automated Reasoning, number 4130 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2006, pp. 67–81.
[15] J. Otten, G. Sutcliffe, Using the TPTP Language for Representing Derivations in Tableau and Connection Calculi, in: B. Konev, R. Schmidt, S. Schulz (Eds.), Proceedings of the Workshop on Practical Aspects of Automated Reasoning, 2010, pp. 90–100.
[16] J. Otten, S. Holden, A Syntax for Connection Proofs, in: J. Otten, W. Bibel (Eds.), Proceedings of the 1st International Workshop on Automated Reasoning with Connection Calculi, number 3613 in CEUR Workshop Proceedings, 2023, pp. 84–94.
[17] G. Sutcliffe, Stepping Stones in the TPTP World, in: C. Benzmüller, M. Heule, R. Schmidt (Eds.), Proceedings of the 12th International Joint Conference on Automated Reasoning, number 14739 in Lecture Notes in Artificial Intelligence, 2024, pp. 30–50.
[18] G. Sutcliffe, The TPTP Problem Library and Associated Infrastructure. The FOF and CNF Parts, v3.5.0, Journal of Automated Reasoning 43 (2009) 337–362.
[19] G. Sutcliffe, The TPTP World - Infrastructure for Automated Reasoning, in: E. Clarke, A. Voronkov (Eds.), Proceedings of the 16th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning, number 6355 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2010, pp. 1–12.
[20] G. Sutcliffe, The SZS Ontologies for Automated Reasoning Software, in: G. Sutcliffe, P. Rudnicki, R. Schmidt, B. Konev, S. Schulz (Eds.), 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, 2008, pp. 38–49.
[21] G. Sutcliffe, C. Suttner, Evaluating General Purpose Automated Theorem Proving Systems, Artificial Intelligence 131 (2001) 39–54. doi:10.1016/S0004-3702(01)00113-8.
[22] G. Sutcliffe, SystemOnTPTP, in: D. McAllester (Ed.), Proceedings of the 17th International Conference on Automated Deduction, number 1831 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2000, pp. 406–410.
[23] A. Stump, G. Sutcliffe, C. Tinelli, StarExec: a Cross-Community Infrastructure for Logic Solving, in: S. Demri, D. Kapur, C. Weidenbach (Eds.), Proceedings of the 7th International Joint Conference on Automated Reasoning, number 8562 in Lecture Notes in Artificial Intelligence, 2014, pp. 367–373.
[24] G. Sutcliffe, The CADE ATP System Competition - CASC, AI Magazine 37 (2016) 99–101.
[25] G. Sutcliffe, The Logic Languages of the TPTP World, Logic Journal of the IGPL 31 (2023) 1153–1169.
[26] G. Sutcliffe, C. Suttner, The TPTP Problem Library: CNF Release v1.2.1, Journal of Automated Reasoning 21 (1998) 177–203.
[27] G. Sutcliffe, S. Schulz, K. Claessen, P. Baumgartner, The TPTP Typed First-order Form with Arithmetic, in: N. Bjørner, A. Voronkov (Eds.), Proceedings of the 18th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning, number 7180 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2012, pp. 406–419.
[28] J. Blanchette, A. Paskevich, TFF1: The TPTP Typed First-order Form with Rank-1 Polymorphism, in: M. Bonacina (Ed.), Proceedings of the 24th International Conference on Automated Deduction, number 7898 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2013, pp. 414–420.
[29] G. Sutcliffe, E. Kotelnikov, TFX: The TPTP Extended Typed First-order Form, in: B. Konev, J. Urban, S. Schulz (Eds.), Proceedings of the 6th Workshop on Practical Aspects of Automated Reasoning, number 2162 in CEUR Workshop Proceedings, 2018, pp. 72–87.
[30] G. Sutcliffe, C. Benzmüller, Automated Reasoning in Higher-Order Logic using the TPTP THF Infrastructure, Journal of Formalized Reasoning 3 (2010) 1–27.
[31] C. Kaliszyk, G. Sutcliffe, F. Rabe, TH1: The TPTP Typed Higher-Order Form with Rank-1 Polymorphism, in: P. Fontaine, S. Schulz, J. Urban (Eds.), Proceedings of the 5th Workshop on Practical Aspects of Automated Reasoning, number 1635 in CEUR Workshop Proceedings, 2016, pp. 41–55.
[32] A. Steen, D. Fuenmayor, T. Gleißner, G. Sutcliffe, C. Benzmüller, Automated Reasoning in Non-classical Logics in the TPTP World, in: B. Konev, C. Schon, A. Steen (Eds.), Proceedings of the 8th Workshop on Practical Aspects of Automated Reasoning, number 3201 in CEUR Workshop Proceedings, 2022, p. Online.
[33] G. Sutcliffe, Semantic Derivation Verification: Techniques and Implementation, International Journal on Artificial Intelligence Tools 15 (2006) 1053–1070. doi:10.1142/S0218213006003119.

[34] G. Sutcliffe, F. Blanqui, G. Burel, Proof Verification with GDV and LambdaPi - It's a Matter of Trust, in: I. Biskri, D. Talbert (Eds.), Proceedings of the 38th International FLAIRS Conference, 2025. doi:10.32473/flairs.38.1.138642.
[35] S. Trac, Y. Puzis, G. Sutcliffe, An Interactive Derivation Viewer, in: S. Autexier, C. Benzmüller (Eds.), Proceedings of the 7th Workshop on User Interfaces for Theorem Provers, volume 174 of *Electronic Notes in Theoretical Computer Science*, 2007, pp. 109–123.
[36] Y. Puzis, Y. Gao, G. Sutcliffe, Automated Generation of Interesting Theorems, in: G. Sutcliffe, R. Goebel (Eds.), Proceedings of the 19th International FLAIRS Conference, AAAI Press, Palo Alto, USA, 2006, pp. 49–54.
[37] J. Otten, W. Bibel, leanCoP: Lean Connection-based Theorem Proving, Journal of Symbolic Computation 36 (2003) 139–161.
[38] C. Johnson, Factorization and Circuit in the Connection Method, Journal of the ACM 40 (1993) 536–557.
[39] O. Astrachan, M. Stickel, Caching and Lemmaizing in Model Elimination Theorem Provers, in: D. Kapur (Ed.), Proceedings of the 11th International Conference on Automated Deduction, number 607 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 1992, pp. 224–238.
[40] J. Schumann, DELTA - A Bottom-up Preprocessor for Top-Down Theorem Provers, in: A. Bundy (Ed.), Proceedings of the 12th International Conference on Automated Deduction, number 814 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 1994, pp. 774–777.
[41] A. Paskevich, Connection Tableaux with Lazy Paramodulation, Journal of Automated Reasoning 40 (2008) 179–194.
[42] R. Hähnle, Tableaux and Related Methods, in: A. Robinson, A. Voronkov (Eds.), Handbook of Automated Reasoning, Elsevier Science, Amsterdam, The Netherlands, 2001, pp. 100–178.
[43] J. Otten, A Non-clausal Connection Calculus, in: K. Brünnler, G. Metcalfe (Eds.), Proceedings of the 20th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, number 6793 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2011, pp. 226–241.
[44] J. Otten, Non-clausal Connection Calculi for Non-classical Logics, in: C. Nalon, R. Schmidt (Eds.), Proceedings of the 26th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, number 10501 in Lecture Notes in Artificial Intelligence, Springer, Berlin, Germany, 2017, pp. 209–227.