# The VAMPIRE Diary

Filip Bártek<sup>1</sup>, Ahmed Bhayat<sup>4</sup>, Robin Coutelier<sup>3</sup>, Márton Hajdu<sup>3</sup>,
 Matthias Hetzenberger<sup>3</sup>, Petra Hozzová<sup>1</sup>, Laura Kovács<sup>3</sup>, <sup>(S)</sup>

Jakob Rath<sup>3</sup>, Michael Rawson<sup>5</sup> , Giles Reger<sup>4</sup>, Martin Suda<sup>1</sup> , Martin Suda<sup>1</sup>, Suda<sup>1</sup>, Suda<sup>1</sup>, Jakob Rath<sup>3</sup>, Michael Rawson<sup>5</sup>, Subara , S

Johannes Schoisswohl<sup>3</sup>, and Andrei Voronkov<sup>2,4</sup>  $(\boxtimes)$ 



<sup>1</sup> Czech Technical University in Prague, Czech Republic martin.suda@cvut.cz <sup>2</sup> EasyChair <sup>3</sup> TU Wien, Vienna, Austria laura.kovacs@tuwien.ac.at <sup>4</sup> University of Manchester, UK andrei@voronkov.com <sup>5</sup> University of Southampton, UK michael@rawsons.uk



**Abstract.** During the past decade of continuous development, the theorem prover VAMPIRE has become an automated solver for the combined theories of commonly-used data structures. VAMPIRE now supports arithmetic, induction, and higher-order logic. These advances have been made to meet the demands of software verification, enabling VAMPIRE to effectively complement SAT/SMT solvers and aid proof assistants. We explain how best to use VAMPIRE in practice and review the main changes VAMPIRE has undergone since its last tool presentation, focusing on the engineering principles and design choices we made during this process.

# 1 Introduction

Automated reasoning has become indispensable for certifying the correctness of software systems and services [55], from Boolean satisfiability (SAT) through satisfiability modulo theories (SMT) to automated theorem proving (ATP) in first-order and higher-order logic. This tool paper describes major developments in saturation-based theorem proving, bringing our VAMPIRE system to bear on modern software certification. VAMPIRE now reasons efficiently in a polymorphic first-order logic with theories, induction and quantifiers, which is realized through (i) combining satisfiability solving with first-order theorem proving using the AVATAR framework [68,47]; (ii) native support for quantified reasoning with mixed arithmetic using extensions of superposition with quantifier elimination [37,56]; and (iii) embedding second-order induction schemata as inference rules in proof search [22,40]. Furthermore, (iv) VAMPIRE has evolved to support higher-order logic [9], program synthesis [27], and finding counterexamples [49].

Our advances in saturation-based reasoning *proved to make a difference*. VAMPIRE outperforms or complements many other state-of-the-art reasoners,



Fig. 1. VAMPIRE timeline since our 2013 tutorial tool paper [41].

including leading SMT solvers and inductive theorem provers. For example, in CASC-J12, the most recent world championship in theorem proving, VAMPIRE proved more problems than any other system in every competition division [64]. We believe that the increasing demand on efficient reasoning with quantifiers, theories and induction turns VAMPIRE into a powerful solver in the automation of mathematics [17], verification of logic programs [42], ensuring system security [32], and many other areas.

This paper details the aforementioned advances in VAMPIRE. It aims to explain how to use VAMPIRE (Section 2) and to give an overview of the reasoning techniques used under the hood in order to allow for reasoning in more expressive logics (Section 4) and more efficient reasoning in general (Section 5). With these new features, a new permissive license, and unprecedented performance, we believe that a tool demonstration after more than a decade of continuous development is overdue. Our diary of improvements since our 2013 tutorial paper demonstration [41] is summarized in Figure 1. The present paper serves as a self-contained tool demonstration, describing the many new features VAMPIRE supports and users can exploit. Our paper contains typical usage guidance, properly instructing readers/users interested in VAMPIRE. For details on VAMPIRE's calculus, inferences processes and proof search algorithms we refer to [41].

# 2 User Guide

VAMPIRE ingests an input problem in either the TPTP [63] or SMT-LIB 2 [3] formats. It then attempts to show unsatisfiability<sup>6</sup> of the input by deriving falsum through application of its logical calculus. If it succeeds, VAMPIRE halts and prints a step-by-step refutation. In some cases VAMPIRE can also show satisfiability. This happens either by finding a finite model (Section 5.2) or by *saturation*: detecting that a refutation cannot be derived with a complete calculus.

<sup>&</sup>lt;sup>6</sup> TPTP allows the user to supply a conjecture: this formula is automatically negated.

**Licensing.** VAMPIRE is now open-source and available online<sup>7</sup> under a 3-clause BSD license. This has completely changed the Vampire team dynamics and coding culture: the development team grew, code became fully shared, and external contributors can improve VAMPIRE. The new license has increased the user base of Vampire in teaching, research and development.

VAMPIRE uses external code for specific tasks, including the MINISAT [18] and CADICAL [11] SAT solvers, the Z3 SMT solver [16] (optional), the VI-RAS quantifier elimination routine [56], and mini-gmp [20] for arbitrary-precision arithmetic.

**Installation.** VAMPIRE is written in C++17 [30] and uses the CMake [35] build system. Since the previous tool demonstration over a decade ago [41], a series of patches have improved portability and VAMPIRE now runs on a variety of modern architectures and UNIX-like operating systems. We provide pre-compiled binaries for UNIX-like systems, and it is generally straightforward to compile VAMPIRE from source on such systems. Support for other operating systems is more experimental, but users report success with compatibility layers such as the Windows Subsystem for Linux [43], Cygwin [45], or Cosmopolitan Libc [67].

**Quick Start.** It is possible, but *usually undesirable*, to invoke VAMPIRE directly on an input file. This will cause it to run a single proof attempt, which will likely not be well-suited to the input. Instead, users should *schedule* an execution of a *portfolio* of many different *strategies*, which is achieved with the invocation

vampire --mode portfolio --schedule <schedule> --cores 0 <problem>

where <problem> is the input problem, --cores 0 instructs VAMPIRE to use all available CPU cores, and <schedule> should be selected from VAMPIRE's list of built-in schedules (Section 5.4), depending on the input.

**Understanding the Output.** VAMPIRE prints status messages as new strategies are launched or old strategies fail, until either a strategy succeeds or the time limit is reached. When a strategy succeeds, by default VAMPIRE reports an SZS status [65] and then, assuming that the input is unsatisfiable, a human-readable proof. The most common SZS statuses VAMPIRE reports are:

Satisfiable: input does not contain a conjecture and is satisfiable Unsatisfiable: input does not contain a conjecture and is unsatisfiable CounterSatisfiable: conjecture present, after negation the input is satisfiable Theorem: conjecture present, after negation the input is unsatisfiable ContradictoryAxioms: conjecture present, but axioms alone are unsatisfiable

**Controlling Output.** VAMPIRE by default uses the SZS standards for output and reports a compact, human-readable proof. The VAMPIRE output can be changed with the *output mode* (-om) and *proof format* (-p) flags. For example, -om smtcomp produces a very terse output suitable for SMT-COMP [4], while -p tptp produces machine-readable TSTP [66] proofs. Some more exotic formats such as LATEX (--latex\_output) or Dedukti [1] are under development.

<sup>&</sup>lt;sup>7</sup> https://github.com/vprover/vampire/

Setting Resource Limits. VAMPIRE can be configured to limit the amount of time (-t <seconds>), memory (-m <MB>), and on recent Linux systems the number of userspace instructions retired (-i <MI><sup>8</sup>). The value 0 means no limit.

**Exploring Options.** VAMPIRE has many, many options. A complete list can be generated with --show\_options on, and a particular flag can be examined with --explain, e.g. --explain output\_mode. Most options only affect a single strategy's behavior, but some affect the global behavior of VAMPIRE, such as the proof format or the time limit. Single strategy options are not usually controlled by the user but automatically set by portfolio schedules (more on this in Section 5.4).

### 2.1 Looking Under the Hood

It is sometimes useful for users to inspect the internal state of VAMPIRE, such as when debugging or optimizing an encoding. Here we sketch the internal mechanisms of VAMPIRE and explain how to inspect them during operation.

VAMPIRE works in two phases. First, the input is parsed, the conjecture — if present — is negated, and the resulting formulas are converted to clause normal form (CNF) and preprocessed. Then VAMPIRE tries to derive the empty clause (witnessing contradiction) in order to refute the CNF of the input problem. New clauses are derived from old by applying generating inferences in VAMPIRE's proof calculus, *superposition* [41]. The search space is partitioned into three sets of clauses: *new* clauses have been freshly derived; *passive* clauses survived VAMPIRE's simplification efforts but have not yet participated in inference, and *active* clauses have themselves participated in inferences generating new clauses. VAMPIRE allows inspection of these processes. To show the clauses resulting from an input, use

#### vampire --mode clausify <problem>

which causes VAMPIRE to stop after preprocessing <problem> and print the resulting CNF. The CNF may be surprising at times as VAMPIRE's preprocessing will happily eliminate parts of the input that it can detect will not help reaching a refutation. To inspect the progress of the preprocessing pipeline users may --show\_preprocessing on. Once proof search begins, --show\_new on displays new clauses (analogously, show\_passive and show\_active). Not all new clauses will make it to *passive*: use --show\_reductions on to see the simplifications that VAMPIRE applies.

# 3 Demonstration: Arithmetic and Induction

Consider the proposition "the sum of two lists of real numbers is equal to the sum of their concatenation". While clearly true, a formal proof requires reasoning about arithmetic, algebraic datatypes, recursive functions, polymorphism, and at

<sup>&</sup>lt;sup>8</sup> Millions of instructions [61]. Instruction limits tend to more stable than time limits across hardware and operating system conditions.

least one inductive step. When given the natural first-order formalization of the problem (Fig. 2) encoded into version 2.7 of SMT-LIB, VAMPIRE is able to find a proof immediately. We show a distilled version of the proof in mathematical notation in Figure 3.

```
\begin{array}{c} \operatorname{sum}(\epsilon) = 0 \\ \forall x : \mathbb{R}. \ \forall xs : [\mathbb{R}]. & \operatorname{sum}(x \, \# \, xs) = x + \operatorname{sum}(xs) \\ \Lambda \alpha. \ \forall ys : [\alpha]. & \epsilon + ys = ys \\ \hline \Lambda \alpha. \ \forall x : \alpha. \ \forall xs, ys : [\alpha]. & (x \, \# \, xs) + ys = x \, \# \, (xs + ys) \\ \hline \forall xs, ys : [\mathbb{R}]. \ \operatorname{sum}(xs) + \operatorname{sum}(ys) = \operatorname{sum}(xs + ys) \end{array}
```

**Fig. 2.** Motivating example in polymorphic first-order logic with uninterpreted functions sum:  $[\mathbb{R}] \to \mathbb{R}, \# : A\alpha$ .  $[\alpha] \times [\alpha] \to [\alpha]$ , algebraic datatypes and real arithmetic.

The proof displays some new features of VAMPIRE, in particular the use of structural induction (Sect. 4.2) and superposition-based arithmetic reasoning via the ALASCA calculus (Sect. 4.1). These features are key to VAMPIRE's success on this problem. The need for our unique blend of induction, arithmetic and polymorphism is supported by the fact that other solvers such as CVC5 [2] or Z3 [16] cannot, to the best of our knowledge, yet process or prove problems such as this one.

# 4 New Capabilities

Here we present the most significant new capabilities implemented in VAMPIRE since 2013 [41]. Improvements to existing capabilities are in Section 5.

#### 4.1 Arithmetic Reasoning

Reasoning about arithmetic in the presence of quantifiers is highly desirable. To this end, VAMPIRE implements the Abstracting Linear Arithmetic Superposition Calculus (ALASCA) [37] that combines ideas like inequality chaining, unification with abstraction [51], and rewriting modulo linear arithmetic. Equality reasoning in ALASCA can be seen as applying the superposition calculus modulo the axioms of linear arithmetic. For example, take step 18 of Figure 3. There, the literal  $0 = x + \operatorname{sum}(y) - \operatorname{sum}(x \# y)$  is used to perform a rewrite  $\operatorname{sum}(x \# y) \rightsquigarrow x +$  $\operatorname{sum}(y)$  instead of a rewrite  $x + \operatorname{sum}(y) - \operatorname{sum}(x \# y) \rightsquigarrow 0$ , which would be the only permissible rewrite in standard superposition. In addition, ALASCA uses inequality chaining and dedicated factoring rules to deal with inequalities, and combines variable elimination rules with unification with abstraction to efficiently perform unification, modulo linear arithmetic.

**Non-Linear Reasoning.** ALASCA itself supports reasoning in linear real arithmetic with uninterpreted functions and quantifiers. Nonlinear problems are also

```
1. 0 = \operatorname{sum}(\epsilon)
                                                                                                                    assumption 1 (cnf)
  2. 0 = x + \operatorname{sum}(y) - \operatorname{sum}(x \# y)
                                                                                                                    assumption 2 (cnf)
                                                                                                                    assumption 3 (cnf)
  3. \epsilon + x = x
  4. y \# (z + x) = (y \# z) + x
                                                                                                                    assumption 4 (cnf)
  5. 0 \neq \operatorname{sum}(\sigma_1) + \operatorname{sum}(\sigma_0) - \operatorname{sum}(\sigma_1 + \sigma_0)
                                                                                                                    conjecture (cnf)
  6. 0 \neq \operatorname{sum}(\sigma_3 \# \sigma_4) + \operatorname{sum}(\sigma_0) - \operatorname{sum}((\sigma_3 \# \sigma_4) + \sigma_0)
                                                                                                                  structural induction 5
       \vee 0 \neq \operatorname{sum}(\epsilon) + \operatorname{sum}(\sigma_0) - \operatorname{sum}(\epsilon + \sigma_0)
  7. 0 = \operatorname{sum}(\sigma_4) + \operatorname{sum}(\sigma_0) - \operatorname{sum}(\sigma_4 + \sigma_0)
                                                                                                                   structural induction 5
       \vee 0 \neq \operatorname{sum}(\epsilon) + \operatorname{sum}(\sigma_0) - \operatorname{sum}(\epsilon + \sigma_0)
  8. 0 \neq sum(\sigma_0) - sum(\sigma_3 \# (\sigma_4 + \sigma_0)) + sum(\sigma_3 \# \sigma_4) forward demodulation 6,4
      \vee 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\epsilon) - \operatorname{sum}(\epsilon + \sigma_0)
 9. 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\epsilon) - \operatorname{sum}(\sigma_0)
                                                                                                                    forward demodulation 7,3
      \vee 0 = \operatorname{sum}(\sigma_0) - \operatorname{sum}(\sigma_4 + \sigma_0) + \operatorname{sum}(\sigma_4)
10. 0 \neq \operatorname{sum}(\epsilon)
                                                                                                                    ALASCA normalization 9
       \vee 0 = \operatorname{sum}(\sigma_0) - \operatorname{sum}(\sigma_4 + \sigma_0) + \operatorname{sum}(\sigma_4)
11. 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\epsilon) - \operatorname{sum}(\sigma_0)
                                                                                                                    forward demodulation 8,3
       \vee 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\sigma_3 \# \sigma_4) - \operatorname{sum}(\sigma_3 \# (\sigma_4 \# \sigma_0))
          0 \neq \operatorname{sum}(\epsilon)
                                                                                                                    ALASCA normalization 11
12.
       \vee 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\sigma_3 \# \sigma_4) - \operatorname{sum}(\sigma_3 \# (\sigma_4 + \sigma_0))
13. 0 = sum(\sigma_0) - sum(\sigma_4 + \sigma_0) + sum(\sigma_4)
                                                                                                                   subsumption resolution 10,1
14. 0 \neq \operatorname{sum}(\sigma_0) + \operatorname{sum}(\sigma_3 \# \sigma_4) - \operatorname{sum}(\sigma_3 \# (\sigma_4 + \sigma_0))
                                                                                                                   subsumption resolution 12,1
15. 0 \neq sum(\sigma_0) + sum(\sigma_3 \# \sigma_4) - (\sigma_3 + sum(\sigma_4 \# \sigma_0))
                                                                                                                    ALASCA superposition 2,14
16. 0 \neq \operatorname{sum}(\sigma_0) - \sigma_3 + \operatorname{sum}(\sigma_3 \# \sigma_4) - (\operatorname{sum}(\sigma_0) + \operatorname{sum}(\sigma_4))
                                                                                                                   ALASCA superposition 13,15
17. 0 \neq \sigma_3 - \operatorname{sum}(\sigma_3 \# \sigma_4) + \operatorname{sum}(\sigma_4)
                                                                                                                    ALASCA normalization 16
18. 0 \neq \sigma_3 - (\sigma_3 + \operatorname{sum}(\sigma_4)) + \operatorname{sum}(\sigma_4)
                                                                                                                    ALASCA superposition 2,17
19. 🗆
                                                                                                                    ALASCA normalization 18
```

Fig. 3. VAMPIRE's proof output of the problem from Fig. 2 in mathematical notation. The symbols  $\sigma_i$  are fresh Skolem constants. The induction steps are detailed in Sect. 4.2.

supported by treating nonlinear multiplications as uninterpreted functions, automatically adding the relevant axiomatization.

Mixed Integer-Real Arithmetic. While the original ALASCA work is limited to real arithmetic, our current implementation in VAMPIRE lifts these restrictions. We support reasoning in mixed integer-real arithmetic, using a tailored quantifier-elimination procedure [56], as well as various new inference rules<sup>9</sup> to handle the combination of mixed arithmetic and uninterpreted functions by natively supporting the rounding (floor) function.

Simplifications and Generalizations. In addition to ALASCA, VAMPIRE also provides lightweight arithmetic reasoning [48]. This includes arithmetic subterm generalization rules that complement ALASCA reasoning, and other simplification rules entailed by ALASCA. Although these simplification rules are not as widely applicable as ALASCA, they provide for more lightweight and therefore efficient arithmetic reasoning, sufficient for many practical problems. The generalization rules include transformations like turning  $\forall x, y : \mathbb{R}$ . P(3x+y) into the equivalent clause  $\forall x : \mathbb{R}$ . P(x).

**Integrating SMT Solvers.** VAMPIRE sometimes hands off ground, that is quantifier-free, arithmetic reasoning to the Z3 SMT solver [16]. This is done either by invoking AVATAR modulo theories [47] (see Sect. 5.2) or by theory instantiation [52]. Theory instantiation uses an SMT solver to find possible instantiations of clauses based on their purely arithmetical literals. To illustrate, consider the clause  $\forall x : \mathbb{Z}$ .  $P(x) \lor 0 > 3x \lor x \ge 1$ . The SMT solver is queried for a model satisfying  $\neg(0 > 3x \lor x \ge 1)$ , which is only the case for x = 0. The clause is instantiated with  $\{x \mapsto 0\}$  and simplified to P(0). Such integration of SMT solvers enables using state-of-the-art developments in SMT and is particularly beneficial for problem areas such as non-linear reasoning, for which VAMPIRE does not yet have dedicated calculi.

### 4.2 Inductive Reasoning

VAMPIRE supports inductive reasoning [40] over literals with up to one free variable.<sup>10</sup> It applies induction by generating theory lemmas, triggered by deriving an eligible induction goal. VAMPIRE supports structural induction over inductively-defined datatypes [53], induction over bounded intervals of integers [29], and well-founded induction principles generated from recursive function definitions [26]. Immediately after VAMPIRE generates the induction lemmas, it uses them to resolve their corresponding goals.

A distinctive feature of VAMPIRE is that it seamlessly interleaves induction with other inferences, efficiently handling hundreds of thousands of induction formulas. This makes it possible to use more explosive lemma generation techniques essential for solving some inductive problems. To synthesize lemmas, VAMPIRE

<sup>&</sup>lt;sup>9</sup> The work on this inference system has not been published yet.

<sup>&</sup>lt;sup>10</sup> This covers the cases of a universally-quantified conjecture, and a conjecture with any number of universally-quantified variables and one existentially-quantified variable.

can generalize over terms and term occurrences [23] or over multiple literals and clauses [24], use function definitions [26] and perform general rewriting [25].

Let us highlight some key steps of the automated induction in VAMPIRE using the proof from Figure 3. First, when VAMPIRE sees clause 5, it detects that induction might be in order, as clause 5 corresponds to a universally-quantified goal using an inductively-defined datatype. Therefore, VAMPIRE uses clause 5 to instantiate the structural induction axiom for lists,

$$L[\epsilon] \land \forall x : \alpha, y : [\alpha].(L[y] \to L[x \# y]) \to \forall z : [\alpha].L[z],$$

by setting  $\alpha := \mathbb{R}$  and  $L[t] := 0 = \operatorname{sum}(t) + \operatorname{sum}(\sigma_0) - \operatorname{sum}(t + \sigma_0)$ . Note that  $L[\sigma_1]$  is set to be complementary to the literal from clause 5.

The instantiated axiom concludes that L[z] holds for any  $z : [\mathbb{R}]$ , while clause 5 expresses that  $L[\sigma_1]$  does not hold. To use this, VAMPIRE converts the axiom into CNF, obtaining clauses  $\neg L[\epsilon] \lor L[\sigma_4] \lor L[z]$  and  $\neg L[\epsilon] \lor \neg L[\sigma_3 \# \sigma_4] \lor L[z]$ , where  $\sigma_3, \sigma_4$  are Skolem constants corresponding to x and y, respectively. These clauses together express that either the antecedent of the axiom does not hold (the base case  $L[\epsilon]$  does not hold, or for some  $\sigma_3, \sigma_4$  we have  $L[\sigma_4]$  but not  $L[\sigma_3 \# \sigma_4]$ ), or the conclusion that L[z] is true for any z must hold. Then VAM-PIRE applies binary resolution on these two clauses with clause 5, resolving away L[z], and deriving clauses 6 and 7. Clauses 6 and 7 are exactly  $\neg L[\epsilon] \lor L[\sigma_4]$  and  $\neg L[\epsilon] \lor \neg L[\sigma_3 \# \sigma_4]$ , spelled out in full in Figure 3. The rest of the proof then covers the refutation of these two clauses.

# 4.3 Polymorphic Logic

VAMPIRE now supports rank-1 polymorphic types [8] in the tradition of Standard ML [44]. This represents a trade-off between expressivity and ease of implementation. Note that VAMPIRE does *not* presently implement a sort inference routine and all sorts in non-variable terms must be explicitly given as sort arguments [12], which may themselves be variables. For example, the axiom  $sum(\epsilon) = 0$  is actually represented as sum(nil(\$real)) = 0, and  $\epsilon + ys = ys$  as concat(A, nil(A), Ys) = Ys. The original motivation for introducing polymorphic logic was supporting combinatory higher-order logic [7], but it has also proved useful for supporting polymorphic theories such as arrays, and for verifying programs that use parametric polymorphism.

#### 4.4 Beyond First-Order Logic

VAMPIRE supports extensions of first-order logic useful for software analysis and verification. In particular, VAMPIRE implements FOOL [38], a conservative extension of many-sorted first-order logic with *if-then-else* and *let-in* expressions, which can be used to capture the next-state relation of loop-free programs [39]. As such, VAMPIRE also supports first-class *Boolean sorts*, by encoding the axiom  $\forall x : o. \ x = 0 \lor x = 1$  as a new inference rule. The rule exploits the two-element domain property of the Boolean sort without blowing up proof search. **Higher-Order Logic.** In addition, a branch of VAMPIRE<sup>11</sup> implements a superposition-based calculus tailored for higher-order logic [10], while still using the general saturation framework from first-order logic. As higher-order unification is undecidable, our implementation bypasses eager unification by performing bounded-depth unification and introducing constraints for remaining unification terms. This technique of constraint introduction has also been used in pure first-order reasoning as delayed unification [9] and in arithmetic reasoning as unification with abstraction [37,6].

### 4.5 Synthesis

We further utilize VAMPIRE's powerful proving capabilities to extend it to a program synthesizer [28,27]. VAMPIRE works with a relational input-output specification expressed in first-order logic, capturing "for all inputs x there exists an output y such that a given relation between x and y holds". In parallel to proving this conjecture, VAMPIRE constructs a program which computes the value of yfor any given value of x. To switch on synthesis mode, use -qa synthesis.<sup>12</sup>

# 5 Making It Work

Taking the above extensions into account, VAMPIRE must now prove theorems in a substantially richer logic with a much greater number of possible inferences. We now describe improvements to VAMPIRE's core that we consider most important for meeting this challenge and maintaining good performance in practice. This adds to the observations of our previous tutorial paper [41], which remain valid. We hope to provide useful information here for those readers who develop their own reasoning systems.

#### 5.1 Preprocessing

Computing normal forms and preprocessing remain of vital importance: the right normal form can eliminate much search space or drastically shorten the required proof. To this end VAMPIRE has grown a new top-down clausal normal form routine [50], lifted the *blocked clause elimination* technique [34] from SAT [31], and adapted a highly effective goal-oriented rewriting technique from Twee [58]. As a general rule of thumb, preprocessing techniques have linear-time complexity, and avoid recursion to prevent stack overflow on large inputs.

# 5.2 Integrating SAT and SMT

One of the tricks for efficiently tackling real-life problems in rich formalisms such as first-order logic with theories is to look for sub-problems in simpler logics

<sup>&</sup>lt;sup>11</sup> https://github.com/vprover/vampire/tree/hol

<sup>&</sup>lt;sup>12</sup> While synthesis of recursion-free programs is available in the mainline VAMPIRE, synthesis of recursive programs is currently in the branch synthesis-recursive.

and offload them to dedicated tools. In this spirit, VAMPIRE implements the AVATAR architecture for clause splitting [68], which allows a prover to delegate the "propositional essence" of the given problem to a SAT solver. In AVATAR modulo theories [47], VAMPIRE uses a finer abstraction<sup>13</sup> and delegates ground *theory* sub-problems to an SMT solver.

SAT solving is also applied within VAMPIRE to find counterexamples to false conjectures. VAMPIRE now provides a MACE-style finite model building mode, using a translation to SAT [14,49]. This is often a useful complement to theoremproving modes (provided that a small counter-model exists), which helps terminate futile searches early and delivers useful insights in the form of bug traces.

# 5.3 Redundancy and Proof Search

Redundancy elimination is key to efficient proof search. Intuitively, a clause is redundant if it is a logical consequence of smaller clauses from the search space: checking whether a first-order clause is redundant is therefore undecidable in general. VAMPIRE implements cheap conditions for detecting some cases of redundancy. The central technique used to implement these checks efficiently is *term indexing* [46] and here in particular *substitution trees* [19] and *code trees* [54]. Code trees are used for rewriting clauses by unit equalities [26,25] and eliminating duplicate clauses, while substitution trees are used for other inferences. To efficiently solve term ordering constraints in redundancy elimination, VAMPIRE uses *term ordering diagrams*, which offer runtime-specialized implementations of simplification orderings [21]. Finally, VAMPIRE also uses SAT solving to check some redundancy conditions that can be modeled as at-most-one ground constraints over the Boolean structure of clause sets [15].

# 5.4 A Sea of Options, Strategies, and Schedules

By *strategy* we mean a particular configuration of VAMPIRE's option values. Since the behavior of VAMPIRE is controlled by more than 200 options, the number of available strategies is vast. Although expert users may sometimes have an idea of which options could be well suited to tackle a problem, the prover's behavior tends to be so chaotic [61] that even expert hunches often fail. For this reason, VAMPIRE provides *schedules* of pre-selected strategies executed sequentially, possibly adapting to the given problem's features.

Creating powerful schedules is a challenging problem. Since 2010, VAMPIRE has employed a dedicated support tool *Spider* [69] to construct schedules from a set of training problems. Spider trials random strategies to solve as many training problems as possible and eventually selects those strategies that *complement* each other particularly well and lead to a schedule with good coverage and a short overall runtime. Techniques have recently been developed to encourage the generalization of the constructed schedule to unseen problems [5].

<sup>&</sup>lt;sup>13</sup> We remark that quantifiers are always handled natively by VAMPIRE.

Actively maintained schedules include casc and casc\_sat, for general firstorder theorem proving and disproving, resp., referring to the famous championship [62]. Similarly, smtcomp has its origin in another competition [71] and is optimized to work well on problems requiring theory reasoning. The higher-order branch (Section 4.4) provides schedules for reasoning in higher-order logic and the *Sledgehammer* [17] use-case. Finally, there is also induction and more.<sup>14</sup>

### 5.5 Branches

Some extensions to VAMPIRE would have violent and extensive impact on the code base. This is true of VAMPIRE's higher-order logic (HOL), for example. Integrating the HOL extension into VAMPIRE would be a significant amount of work and impose a burden on all VAMPIRE developers: but we would like it to continue, as it is a world-leading system for higher-order logic. The way we are currently dealing with this tension is by keeping this kind of feature on git branches, which are periodically synchronised with mainline VAMPIRE. When a branch is widely-used enough, stable, and has a clear path to be integrated cleanly with mainline VAMPIRE, we may consider merging it: this has happened in the past with rank-1 polymorphism and a previous approach to HOL [7].

# 6 Related Work and Conclusion

We have explained how best to use VAMPIRE, discussed new features of VAM-PIRE that better align saturation-based first-order theorem proving with software verification, and described engineering required to make it work in practice.

As a first-order theorem prover with support for theories, induction and higher-order logic, VAMPIRE has been influenced by, competes with, and might be variously compared to: SMT solvers such as cvc5 [2] or Z3 [16]; first-order ATPs such as E [57], SPASS [72], iProver [36], or Twee [58]; inductive theorem provers such as ACL2 [33], HipSpec [13], or Zeno [59]; and higher-order ATPs such as Zipperposition [70] or Leo-III [60]. VAMPIRE distinguishes itself with its native support for quantifiers combined with calculus extensions to reason about theories, induction and higher-order logic, all tied together by highly-efficient adaptive data structures and algorithms. Naturally, VAMPIRE integrates SAT and SMT solving for ground reasoning tasks.

This paper overviewed the main reasoning engines and practices VAMPIRE offers in order to assist users in understanding the many ways VAMPIRE can be integrated in other technologies. The system is under continuous development, with new applications towards proof checking and extracting system code from formal proofs. Further advances in creating tailored VAMPIRE proof schedules for proof assistants, for example in Isabelle's *Sledgehammer* [17], are also under active development.

<sup>&</sup>lt;sup>14</sup> Use --explain\_option schedule to list schedules available from your VAMPIRE.

Acknowledgements. We would like to thank all users and prior developers who contributed to VAMPIRE. We acknowledge the valuable VAMPIRE contributions made by Daneshvar Amrollahi, Ioan Dragan, Bernhard Gleiss, Bernhard Kragl, Kryštof Hoder, Evgenii Kotelnikov, Alexandre Riazanov, Martin Riener, Simon Robillard, Boris Shminke, and Eva Maria Wagner.

This research was funded in whole or in part by the ERC Consolidator Grant ARTIST 101002685, the ERC Proof of Concept Grant LEARN 101213411, the TU Wien Doctoral College SecInt, the FWF SpyCoDe Grant 10.55776/F85, the WWTF grant [ForSmart Grant ID: 10.47379/ICT22007], and the Amazon Research Award 2023 QuAT. Martin Suda was supported by the project CORE-SENSE no. 101070254 under the Horizon Europe programme and by the Czech Ministry of Education, Youth and Sports under the ERC CZ project POST-MAN no. LL1902. Petra Hozzová was supported by the European Union under the project ROBOPROX (reg. no. CZ.02.01.01/00/22\_008/0004590).

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

### References

- Assaf, A., Burel, G., Cauderlier, R., Delahaye, D., Dowek, G., Dubois, C., Gilbert, F., Halmagrand, P., Hermant, O., Saillard, R.: Dedukti: a Logical Framework based on the λΠ-Calculus Modulo Theory. CoRR abs/2311.07185 (2023)
- Barbosa, H., Barrett, C.W., Brain, M., Kremer, G., Lachnitt, H., Mann, M., Mohamed, A., Mohamed, M., Niemetz, A., Nötzli, A., Ozdemir, A., Preiner, M., Reynolds, A., Sheng, Y., Tinelli, C., Zohar, Y.: cvc5: A Versatile and Industrial-Strength SMT Solver. In: TACAS. pp. 415–442 (2022)
- 3. Barrett, C., Fontaine, P., Tinelli, C.: The Satisfiability Modulo Theories Library (SMT-LIB). www.SMT-LIB.org (2016)
- Barrett, C., de Moura, L., Stump, A.: SMT-COMP: Satisfiability modulo Theories Competition. In: CAV. p. 20–23 (2005)
- Bártek, F., Chvalovský, K., Suda, M.: Regularization in Spider-Style Strategy Discovery and Schedule Construction. In: IJCAR. pp. 194–213 (2024)
- Bhayat, A., Korovin, K., Kovács, L., Schoisswohl, J.: Refining Unification with Abstraction. In: LPAR. pp. 36–47 (2023)
- Bhayat, A., Reger, G.: A Combinator-Based Superposition Calculus for Higher-Order Logic. In: IJCAR. pp. 278–296 (2020)
- Bhayat, A., Reger, G.: A Polymorphic VAMPIRE (Short Paper). In: IJCAR. pp. 361–368 (2020)
- Bhayat, A., Schoisswohl, J., Rawson, M.: Superposition with Delayed Unification. In: CADE. pp. 23–40 (2023)
- Bhayat, A., Suda, M.: A Higher-Order VAMPIRE (Short Paper). In: IJCAR. pp. 75–85 (2024)
- Biere, A., Faller, T., Fazekas, K., Fleury, M., Froleyks, N., Pollitt, F.: Cadical 2.0. In: CAV. pp. 133–152 (2024). https://doi.org/10.1007/978-3-031-65627-9\_7, https://doi.org/10.1007/978-3-031-65627-9\_7

- Blanchette, J.C., Paskevich, A.: TFF1: The TPTP Typed First-Order Form with Rank-1 Polymorphism. In: CADE. pp. 414–420 (2013)
- Claessen, K., Johansson, M., Rosén, D., Smallbone, N.: Automating Inductive Proofs Using Theory Exploration. In: CADE (2013)
- Claessen, K., Sörensson, N.: New Techniques that Improve MACE-style Model Finding. In: WS on Model Computation - Principles, Algorithms and Applications (2003)
- 15. Coutelier, R., Rath, J., Rawson, M., Biere, A., Kovács, L.: SAT Solving for Variants of First-Order Subsumption. Formal Methods in System Design (2024)
- De Moura, L., Bjørner, N.: Z3: An Efficient SMT Solver. In: TACAS. pp. 337–340 (2008)
- 17. Desharnais, M., Vukmirović, P., Blanchette, J., Wenzel, M.: Seventeen Provers Under the Hammer. In: ITP. pp. pp. 8:1–8:18 (2022)
- 18. Eén, N., Sörensson, N.: An Extensible SAT-solver. In: SAT. pp. 502–518 (2003)
- 19. Graf, P.: Substitution Tree Indexing. In: RTA. pp. 117–131 (1995)
- Granlund, T.: The GNU Multiple Precision Arithmetic Library (2023), https:// gmplib.org/gmp-man-6.3.0.pdf
- Hajdu, M., Coutelier, R., Kovács, L., Voronkov, A.: Term Ordering Diagrams. In: CADE (2025), to appear
- Hajdu, M., Hozzová, P., Kovács, L., Reger, G., Voronkov, A.: Getting Saturated with Induction. In: Principles of Systems Design. pp. 306–322 (2022)
- Hajdu, M., Hozzová, P., Kovács, L., Schoisswohl, J., Voronkov, A.: Induction with Generalization in Superposition Reasoning. In: CICM. pp. 123–137 (2020)
- Hajdu, M., Kovacs, L., Rawson, M., Voronkov, A.: The VAMPIRE Approach to Induction. EasyChair Preprint no. 9217 (EasyChair, 2022)
- Hajdu, M., Kovács, L., Rawson, M.: Rewriting and Inductive Reasoning. In: LPAR. pp. 278–294 (2024)
- Hajdu, M., Hozzová, P., Kovács, L., Voronkov, A.: Induction with Recursive Definitions in Superposition. In: FMCAD. pp. 1–10 (2021)
- Hozzová, P., Amrollahi, D., Hajdu, M., Kovács, L., Voronkov, A., Wagner, E.M.: Synthesis of Recursive Programs in Saturation. In: IJCAR. p. 154–171 (2024)
- Hozzová, P., Kovács, L., Norman, C., Voronkov, A.: Program synthesis in saturation. In: CADE. pp. 307–324 (2023)
- Hozzová, P., Kovács, L., Voronkov, A.: Integer Induction in Saturation. In: CADE. pp. 361–377 (2021)
- ISO: ISO/IEC 14882:2017: Programming languages C++. International Organization for Standardization, Geneva, Switzerland (Dec 2017)
- Järvisalo, M., Biere, A., Heule, M.: Blocked Clause Elimination. In: TACAS. pp. 129–144 (2010)
- Jeanteur, S., Kovács, L., Maffei, M., Rawson, M.: CryptoVampire: Automated Reasoning for the Complete Symbolic Attacker Cryptographic Model. In: SP. pp. 3165–3183 (2024)
- Kaufmann, M., Manolios, P., Moore, J.S.: Computer-Aided Reasoning: An Approach, vol. 3. Springer (06 2000). https://doi.org/10.1007/978-1-4615-4449-4
- Kiesl, B., Suda, M., Seidl, M., Tompits, H., Biere, A.: Blocked Clauses in First-Order Logic. In: LPAR. pp. 31–48 (2017)
- 35. Kitware, I.: CMake (2025), https://cmake.org/
- Korovin, K.: iProver An Instantiation-Based Theorem Prover for First-Order Logic (System Description). In: IJCAR. pp. 292–298 (2008)
- Korovin, K., Kovács, L., Reger, G., Schoisswohl, J., Voronkov, A.: ALASCA: Reasoning in Quantified Linear Arithmetic. In: TACAS. pp. 647–665 (2023)

- 14 The VAMPIRE Team
- Kotelnikov, E., Kovács, L., Reger, G., Voronkov, A.: The VAMPIRE and the FOOL. In: CPP. pp. 37–48 (2016)
- Kotelnikov, E., Kovács, L., Voronkov, A.: A FOOLish Encoding of the Next State Relations of Imperative Programs. In: IJCAR. pp. 405–421 (2018)
- Kovács, L., Hozzová, P., Hajdu, M., Voronkov, A.: Induction in Saturation. In: IJCAR. pp. 21–29 (2024)
- Kovács, L., Voronkov, A.: First-Order Theorem Proving and VAMPIRE. In: CAV. pp. 1–35 (2013)
- Lifschitz, V., Lühne, P., Schaub, T.: Towards Verifying Logic Programs in the Input Language of clingo. In: Fields of Logic and Computation III. pp. 190–209 (2020)
- 43. Microsoft: Windows Subsystem for Linux (WSL), https://ubuntu.com/desktop/ wsl
- 44. Milner, R.: The Definition of Standard ML: Revised. MIT press (1997)
- 45. Racine, J.: The Cygwin Tools: a GNU Toolkit for Windows (2000)
- Ramakrishnan, I.V., Sekar, R., Voronkov, A.: Term Indexing. In: Handbook of Automated Reasoning, pp. 1853–1964. Elsevier and MIT Press (2001)
- Reger, G., Bjørner, N.S., Suda, M., Voronkov, A.: AVATAR Modulo Theories. In: GCAI. pp. 39–52 (2016)
- Reger, G., Schoisswohl, J., Voronkov, A.: Making Theory Reasoning Simpler. In: TACAS. pp. 164–180 (2021)
- Reger, G., Suda, M., Voronkov, A.: Finding Finite Models in Multi-sorted First-Order Logic. In: SAT. pp. 323–341 (2016)
- Reger, G., Suda, M., Voronkov, A.: New Techniques in Clausal Form Generation. In: GCAI. pp. 11–23 (2016)
- Reger, G., Suda, M., Voronkov, A.: Unification with Abstraction and Theory Instantiation in Saturation-Based Reasoning. In: TACAS. pp. 3–22 (2018)
- Reger, G., Suda, M., Voronkov, A.: Unification with Abstraction and Theory Instantiation in Saturation-Based Reasoning. In: TACAS. pp. 3–22 (2018)
- Reger, G., Voronkov, A.: Induction in Saturation-Based Proof Search. In: CADE. pp. 477–494 (2019)
- Riazanov, A., Voronkov, A.: Partially Adaptive Code Trees. In: JELIA. pp. 209–223 (2000)
- Rungta, N.: A Billion SMT Queries a Day (Invited Paper). In: CAV. pp. 3–18 (2022)
- Schoisswohl, J., Kovács, L., Korovin, K.: VIRAS: Conflict-Driven Quantifier Elimination for Integer-Real Arithmetic. In: LPAR. pp. 147–164 (2024)
- Schulz, S., Cruanes, S., Vukmirović, P.: Faster, Higher, Stronger: E 2.3. In: CADE. pp. 495–507 (2019)
- Smallbone, N.: Twee: An Equational Theorem Prover. In: CADE. pp. 602–613 (2021)
- Sonnex, W., Drossopoulou, S., Eisenbach, S.: Zeno: An Automated Prover for Properties of Recursive Data Structures. In: TACAS. pp. 407–421 (2012)
- Steen, A., Benzmüller, C.: The Higher-Order Prover Leo-III. In: IJCAR. pp. 108– 116 (2018)
- Suda, M.: VAMPIRE Getting Noisy: Will Random Bits Help Conquer Chaos? (System Description). In: IJCAR. pp. 659–667 (2022)
- Sutcliffe, G.: The CADE ATP System Competition CASC. AI Magazine 37(2), 99–101 (2016)
- Sutcliffe, G.: The Logic Languages of the TPTP World. Logic Journal of the IGPL (2022). https://doi.org/10.1093/jigpal/jzac068

- 64. Sutcliffe, G.: The 12th IJCAR Automated Theorem Proving System Competition CASC-J12. The European Journal on Artificial Intelligence  $\mathbf{0}(0)$ , 30504554241305110 (0)
- Sutcliffe, G.: The SZS Ontologies for Automated Reasoning Software. In: LPAR Workshops (2008)
- 66. Sutcliffe, G.: The TPTP World Infrastructure for Automated Reasoning. In: LPAR. pp. 1–12 (2010)
- 67. Tunney, J.: Cosmopolitan Libc (2025), https://justine.lol/cosmopolitan/
- Voronkov, A.: AVATAR: The Architecture for First-Order Theorem Provers. In: CAV. pp. 696–710 (2014)
- Voronkov, A.: Spider: Learning in the Sea of Options (2023), https://easychair. org/smart-program/Vampire23/2023-07-05.html#talk:223833
- Vukmirovic, P., Bentkamp, A., Blanchette, J., Cruanes, S., Nummelin, V., Tourret, S.: Making Higher-Order Superposition Work. J. Autom. Reason. 66(4), 541–564 (2022)
- Weber, T., Conchon, S., Déharbe, D., Heizmann, M., Niemetz, A., Reger, G.: The SMT competition 2015-2018. J. Satisf. Boolean Model. Comput. 11(1), 221–259 (2019), https://doi.org/10.3233/SAT190123
- Weidenbach, C., Dimova, D., Fietzke, A., Kumar, R., Suda, M., Wischnewski, P.: SPASS Version 3.5. In: CADE. pp. 140–145 (2009)