



Interpolating bit-vector formulas using uninterpreted predicates and Presburger arithmetic

Peter Backeman¹ · Philipp Rümmer² · Aleksandar Zeljic³

Accepted: 25 March 2021 / Published online: 12 May 2021
© The Author(s) 2021

Abstract

The inference of program invariants over machine arithmetic, commonly called bit-vector arithmetic, is an important problem in verification. Techniques that have been successful for unbounded arithmetic, in particular Craig interpolation, have turned out to be difficult to generalise to machine arithmetic: existing bit-vector interpolation approaches are based either on eager translation from bit-vectors to unbounded arithmetic, resulting in complicated constraints that are hard to solve and interpolate, or on bit-blasting to propositional logic, in the process losing all arithmetic structure. We present a new approach to bit-vector interpolation, as well as bit-vector quantifier elimination (QE), that works by lazy translation of bit-vector constraints to unbounded arithmetic. Laziness enables us to fully utilise the information available during proof search (implied by decisions and propagation) in the encoding, and this way produce constraints that can be handled relatively easily by existing interpolation and QE procedures for Presburger arithmetic. The lazy encoding is complemented with a set of native proof rules for bit-vector equations and non-linear (polynomial) constraints, this way minimising the number of cases a solver has to consider. We also incorporate a method for handling concatenations and extractions of bit-vector efficiently.

Keywords Bit-vectors · Interpolation · Quantifier elimination · Presburger arithmetic

1 Introduction

Craig interpolation is a commonly used technique to infer invariants or contracts in verification. Over the last 15 years, efficient interpolation techniques have been developed for a variety of logics and theories, including propositional logic [1,2], uninterpreted functions [1,3,4], first-order logic [5–7], algebraic data-types [8,9], linear real arithmetic [1], non-linear real arithmetic [10], Presburger arithmetic [4,11,12], and arrays [13–15].

✉ Peter Backeman
peter@backeman.se

¹ Mälardalen University, Västerås, Sweden

² Department of Information Technology, Uppsala University, Uppsala, Sweden

³ Stanford University, Stanford, USA

A theory that has turned out notoriously difficult to handle in Craig interpolation is bounded machine arithmetic, commonly called bit-vector arithmetic. Decision procedures for bit-vectors are predominantly based on bit-blasting, in combination with sophisticated preprocessing and simplification methods, which implies that also extracted interpolants stay on the level of propositional logic and are difficult to map back to compact high-level bit-vector constraints. An alternative interpolation approach translates bit-vector constraints to unbounded integer arithmetic formulas [16], but is limited to linear constraints and tends to produce integer formulas that are hard to solve and interpolate, due to the necessary introduction of additional variables and large coefficients to model wrap-around semantics correctly.

In this article, we introduce a new Craig interpolation method for bit-vector arithmetic, initially focusing on arithmetic bit-vector operations including addition, multiplication, and division. Like [16], we compute interpolants by reducing bit-vectors to unbounded integers; unlike in earlier approaches, we define a calculus that carries out this reduction lazily, and can therefore dynamically choose between multiple possible encodings of the bit-vector operations. This is done by initially representing bit-vector operations as uninterpreted predicates, which are expanded and replaced by Presburger arithmetic expressions on demand. The calculus also includes native rules for non-linear constraints and bit-vector equations, so that formulas can often be proven without having to resort to a full encoding as integer constraints. Our approach gives rise to both Craig interpolation and quantifier elimination (QE) methods for bit-vector constraints, with both procedures displaying competitive performance in our experiments.

Reduction of bit-vectors to unbounded integers has the additional advantage that integer and bit-vector formulas can be combined efficiently, including the use of conversion functions between both theories, which are difficult to support using bit-blasting. This combination is of practical importance in software verification, since programs and specifications often mix machine arithmetic with arbitrary-precision numbers; tools might also want to switch between integer semantics (if it is known that no overflows can happen) and bit-vector semantics for each individual program instruction.

This is an extended version of a paper presented at FMCAD 2018 [17]. Compared to the conference version, this article considers an extended fragment of bit-vector logic, including also concatenation and extraction operations on bit-vectors, as well as bit-wise operators like `bvor` or `bvnot`. We show that the representation of concatenation and extraction using uninterpreted predicates is sufficient to obtain an interpolation procedure for the quantifier-free structural fragment of bit-vector logic, i.e., bit-vector constraints with only concatenation, extraction, and positive equations [18,19]. Bit-wise operations are handled via a direct translation to Presburger arithmetic akin to bit-blasting.

The contributions of the article are: a new calculus for non-linear integer arithmetic, which can eliminate quantifiers (in certain cases) and extract Craig interpolants (Sect. 3); a corresponding calculus for arithmetic bit-vector constraints (Sect. 4); the extension of the calculus to handle concatenation, extraction, and bit-wise operations (Sect. 5); an experimental evaluation using SMT-LIB and model checking benchmarks (Sect. 6).

1.1 Example 1: Interpolating arithmetic bit-vector operations

We start by considering one of the examples from [16], the interpolation problem $A \wedge B$ defined by

$$\begin{aligned} A &= \neg \text{bvule}_8(\text{bvadd}_8(y_4, 1), y_3) \wedge y_2 = \text{bvadd}_8(y_4, 1) \\ B &= \text{bvule}_8(\text{bvadd}_8(y_2, 1), y_3) \wedge y_7 = 3 \wedge y_7 = \text{bvadd}_8(y_2, 1) \end{aligned}$$

where all variables range over unsigned 8-bit bit-vectors. The function bvadd_8 represents addition of two bit-vectors, while the predicate bvule_8 is the unsigned \leq comparison. An interpolant for $A \wedge B$ is a formula I such that the implications $A \Rightarrow I$ and $B \Rightarrow \neg I$ hold, and such that only variables common to A and B occur in I .

An eager encoding into Presburger arithmetic (linear integer arithmetic, LIA) would typically add variables to handle wrap-around semantics, e.g., mapping $y'_4 = \text{bvadd}_8(y_4, 1)$ to $y'_4 = y_4 + 1 - 2^8 \sigma_1 \wedge 0 \leq y'_4 < 2^8 \wedge 0 \leq \sigma_1 \leq 1$. This yields a formula in Presburger arithmetic that exactly models the bit-vector semantics, and can be solved and interpolated using existing methods implemented in SMT solvers. Interpolants can be mapped back to a pure bit-vector formula if needed. However, additional variables and large coefficients tend to be hard both for solving and interpolation; the LIA interpolant presented in [16] for $A \wedge B$ is the somewhat complicated formula $I_{LIA} = -255 \leq y_2 - y_3 + 256 \lfloor -1 \frac{y_2}{256} \rfloor$.

Our approach translates bit-vector formulas to our *core language* — an extension of Presburger arithmetic with constructs to express bit-vector domains, wrap-around semantics and operations that can be simplified in different ways, such as bvmul . For example, domain predicate $\text{in}_w(x)$ expresses that variable x belongs to the value range of a bit-vector of width w . Similarly, predicate $\text{ubmod}_w(x, y)$ expresses the unsigned wrap-around semantics without explicitly encoding it. Translating A and B to the core language yields:

$$\begin{aligned} A_{\text{core}} &= \psi_A \wedge \text{ubmod}_8(y_4 + 1, c_1) \wedge c_1 > y_3 \wedge y_2 = c_1 \\ B_{\text{core}} &= \psi_B \wedge \text{ubmod}_8(y_2 + 1, c_2) \wedge c_2 \leq y_3 \wedge y_7 = 3 \wedge y_7 = c_2 \end{aligned}$$

where $\psi_A = \text{in}_8(y_2) \wedge \text{in}_8(y_3) \wedge \text{in}_8(y_4) \wedge \text{in}_8(c_1)$ and $\psi_B = \text{in}_8(y_2) \wedge \text{in}_8(y_3) \wedge \text{in}_8(y_7) \wedge \text{in}_8(c_2)$ capture the domain constraints.

The core language enables a layered calculus that encodes predicates on a case by case basis, preferring simpler encodings whenever possible. In our example, rule BMOD-SPLIT splits the $\text{ubmod}_8(y_2 + 1, c_2)$ into the only two relevant cases based on the bounds of y_2 implied by $A_{\text{core}}, B_{\text{core}}$:

$$\frac{\begin{array}{l} \dots, 0 \leq c_2 < 256, y_2 + 1 = c_2 \vdash \\ \dots, 0 \leq c_2 < 256, y_2 + 1 = c_2 + 256 \vdash \end{array}}{\dots, \text{ubmod}_8(y_2 + 1, c_2) \vdash} \text{BMOD-SPLIT}$$

Due to $y_7 = 3 \wedge y_7 = c_2$, the cases reduce to $y_2 = 2$ and $y_2 = 258$, and immediately contradict $A_{\text{core}}, B_{\text{core}}$.

When variable bounds are tight enough and there are only a few cases, case splits are more efficient than σ variables. However, that is not always the case and our calculus lazily decides how to handle each occurrence. Simpler proofs also lead to simpler and more compact interpolants; using our lazy approach, the final interpolant in the example is $I_{LAZY} = y_3 < y_2$, which is simple and avoids the division operator in I_{LIA} . We will revisit this example in Sect. 4.4 and explain in greater detail how this interpolant is obtained.

1.2 Example 2: Interpolating structural bit-vector operations

We continue with a (reduced) example taken from [19], a formula of equalities between (slices of) bit-vectors of length 8:

$$x[5 : 0] = 22 \wedge y[7 : 2] = 6 \wedge x = y$$

where $x[u : l]$ is the *extraction* of the slice of bits from u th down to l th (inclusive). In the previous example the bit-vector formula was translated to integer arithmetic, however this can

sometimes be inefficient when dealing with *structural* bit-vector operations, e.g., extractions and concatenations. A direct translation to integer arithmetic has a hard time to isolate the conflict, since integer operations cannot capture extractions in a natural way. Instead, it is possible to split the bit-vectors into segments

$x[5:0] = 22$	$y[7:2] = 6$	$x = y$
$x[5:2] = 5$	$y[7:6] = 0$	$x[7:6] = y[7:6]$
$x[1:0] = 2$	$y[5:2] = 6$	$x[5:2] = y[5:2]$
		$x[1:0] = y[1:0]$

Given this *decomposition* of the bit-vectors, it is easy see the conflict $x[5 : 2] = 5 \neq 6 = y[5 : 2]$ without a translation to integers. Interpolants can in this setting be extracted by referring to individual slices of bit-vectors, with the help of the extraction operator. In Sect. 5 we show how bit-vectors can be decomposed in this manner using an interpolating calculus.

1.3 Related work

Most SMT solvers handle bit-vectors using bit-blasting and SAT solving, and usually cannot extract interpolants for bit-vector problems. The exception is MATHSAT [20], which uses a layered approach [16] to compute interpolants: MATHSAT first tries to compute interpolants by keeping bit-vector operations uninterpreted; then using a restricted form of quantifier elimination; then by eager encoding into linear integer arithmetic (LIA); and finally through bit-blasting. Our approach has some similarities to the LIA encoding, but can choose simpler encodings thanks to laziness, and also covers non-linear arithmetic constraints.

A similarly layered approach, proposed in [21], can be used to compute function summaries in bounded model checking. When bounded model checking is able to prove (bounded) safety of a program, Craig interpolation can subsequently be used to extract function summaries; such summaries can later be useful to speed up other verification tasks. To handle bit-vector constraints in this context, [21] successively applies more and more precise over-approximations of bit-vectors: using uninterpreted functions, linear real arithmetic, and finally using precise bit-blasting. Interpolants are computed in the coarsest theory that was able to prove safety of a verification task.

Other related work has focused on interpolation for fragments of bit-vector logic. In [22], an algorithm is given for reconstructing bit-vector interpolants from bit-level interpolants, however restricted to the case of bit-vector equalities. An interpolation procedure based on a set of tailor-made (but incomplete) rewriting rules for bit-vectors is given in [23].

Looking more generally at model checking for finite-state systems formulated over the theory of bit-vectors (often called word-level model checking), lazy approaches to handle complex bit-vector operations have been proposed. In [24], an approximation method for model checking RTL designs is defined that instantiates complex bit-vector operations lazily. Initially, such operations are over-approximated by leaving the results unconstrained; when spurious counterexamples occur, the approximation is refined by adding additional constraints, or ultimately by precisely instantiating the operator. Such approaches are independent of the underlying finite-state model checking algorithm, and do not necessarily involve Craig interpolation, however.

The core logic of bit-vectors (formulas with only concatenation, extraction, and positive equations) was identified in [18] to be solvable in polynomial time.¹ Our work is inspired by the decomposition-based decision procedure for this fragment developed in [19], where the authors present an algorithm together with a data-structure designed for solving formulas over the core logic of bit-vectors efficiently. To the best of our knowledge, Craig interpolation for the structural fragment has not been considered previously.

2 Preliminaries: the base logic

We formulate our approach on top of a simple logic of Presburger arithmetic constraints combined with uninterpreted predicates, introduced in [25] and extended in [4,11] to support Craig interpolation. Let x range over an infinite set X of variables, c over an infinite set C of constants, p over a set P of predicate symbols with fixed arity, and α over the set \mathbb{Z} of integers. The syntax of terms and formulas is defined by the following grammar:

$$\begin{aligned}\phi &::= t = 0 \mid t \leq 0 \mid p(t, \dots, t) \mid \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \forall x. \phi \mid \exists x. \phi \\ t &::= \alpha \mid c \mid x \mid \alpha t + \dots + \alpha t\end{aligned}$$

The symbol t denotes terms of linear arithmetic. Substitution of a term t for a variable x in ϕ is denoted by $[x/t]\phi$; we assume that variable capture is avoided by renaming bound variables as necessary. For simplicity, we sometimes write $s = t$ as a shorthand of $s - t = 0$, inequalities $s \leq t$ and $t \geq s$ for $s - t \leq 0$, and $\forall c. \phi$ as a shorthand of $\forall x. [c/x]\phi$ if c is a constant. The abbreviation *true* (*false*) stands for equality $0 = 0$ ($1 = 0$), and the formula $\phi \rightarrow \psi$ abbreviates $\neg \phi \vee \psi$. Semantic notions such as structures, models, satisfiability, and validity are defined as is common (e.g., [26]), but we assume that evaluation always happens over the universe \mathbb{Z} of integers; bit-vectors will later be defined as a subset of the integers.

2.1 A sequent calculus for the base logic

For checking whether a formula in the base logic is satisfiable or valid, we work with the calculus presented in [25], a part of which is shown in Fig. 1. If Γ, Δ are finite sets of formulas, then $\Gamma \vdash \Delta$ is a *sequent*. A sequent is *valid* if the formula $\bigwedge \Gamma \rightarrow \bigvee \Delta$ is valid. Positions in Δ that are underneath an even/odd number of negations are called positive/negative; and vice versa for Γ . Proofs are trees growing upward, in which each node is labeled with a sequent, and each non-leaf node is related to the node(s) directly above it through an application of a calculus rule. A proof is *closed* if it is finite and all leaves are justified by an instance of a rule without premises. Soundness of the calculus implies that the root of a closed proof is a valid sequent.

In addition to propositional and quantifier rules in Fig. 1, the calculus in [25] also includes rules for equations and inequalities in Presburger arithmetic; the details of those rules are not relevant for this paper. The calculus is complete for quantifier-free formulas in the base logic, i.e., for every valid quantifier-free sequent a closed proof can be found. It is well-known that the base logic including quantifiers does not admit complete calculi [27], but as discussed in [25] the calculus can be made complete (by adding slightly more sophisticated quantifier

¹ To avoid confusing with our own “core” fragment introduction in Sect. 4, we call the logic from [18] the “structural fragment” in this article.

$\frac{\Gamma, \phi \vdash \Delta \quad \Gamma, \psi \vdash \Delta}{\Gamma, \phi \vee \psi \vdash \Delta} \vee\text{-LEFT}$	$\frac{\Gamma, \phi \vdash \Delta \quad \Gamma, \psi \vdash \Delta}{\Gamma \vdash \phi \wedge \psi, \Delta} \wedge\text{-RIGHT}$
$\frac{\Gamma, \phi, \psi \vdash \Delta}{\Gamma, \phi \wedge \psi \vdash \Delta} \wedge\text{-LEFT}$	$\frac{\Gamma \vdash \phi, \psi, \Delta}{\Gamma \vdash \phi \vee \psi, \Delta} \vee\text{-RIGHT}$
$\frac{\Gamma \vdash \phi, \Delta}{\Gamma, \neg \vdash \Delta} \neg\text{-LEFT}$	$\frac{\Gamma, \phi \vdash \Delta}{\Gamma \vdash \neg \phi, \Delta} \neg\text{-RIGHT}$
$\frac{*}{\Gamma, \phi \vdash \phi, \Delta} \text{CLOSE}$	

$\frac{\Gamma, [x/t]\phi, \forall x.\phi \vdash \Delta}{\Gamma, \forall x.\phi \vdash \Delta} \forall\text{-LEFT}$	$\frac{\Gamma, [x/c]\phi \vdash \Delta}{\Gamma, \exists x.\phi \vdash \Delta} \exists\text{-LEFT}$
$\frac{\Gamma \vdash [x/t]\phi, \exists x.\phi, \Delta}{\Gamma \vdash \exists x.\phi, \Delta} \exists\text{-RIGHT}$	$\frac{\Gamma \vdash [x/c]\phi, \Delta}{\Gamma \vdash \forall x.\phi, \Delta} \forall\text{-RIGHT}$

Fig. 1 A selection of the basic calculus rules for propositional logic (upper box) and quantifier rules (lower box). In the rules \exists -LEFT and \forall -RIGHT, c is a constant that does not occur in the conclusion

handling) for interesting undecidable fragments, for instance for sequents $\vdash \phi$ in which ϕ contains \exists/\forall only under an even/odd number of negations.

For quantifier-free input formulas, proof search can be implemented in depth-first style following the core concepts of DPLL(T) [28]: rules with multiple premises correspond to *decisions* and explore the branches one by one; rules with a single premise represent *propagation* or *rewriting*; and logging of rule applications is used in order to implement conflict-driven learning and proof extraction. For experiments, we use the implementation of the calculus in PRINCESS.²

2.2 Quantifier elimination in the base logic

The sequent calculus can eliminate quantifiers in Presburger arithmetic, i.e., in the base logic without uninterpreted predicates, since the arithmetic calculus rules are designed to systematically eliminate constants. To illustrate this use case, suppose ϕ is a formula without uninterpreted predicates ($P = \emptyset$) and without constants c , but possibly containing variables x . Formula ϕ furthermore only contains \forall/\exists under an even/odd number of negations, i.e., all quantifiers are effectively universal. To compute a quantifier-free formula ψ that is equivalent to ϕ , we can construct a proof with root sequent $\vdash \phi$, and keep applying rules until no further applications are possible in any of the remaining open goals $\{\Gamma_i \vdash \Delta_i \mid i = 1, \dots, n\}$. In this process, rules \exists -LEFT and \forall -RIGHT can introduce fresh constants, which are subsequently isolated and eliminated by the arithmetic rules. To find ψ , it is essentially enough to extract the constant-free formulas $\Gamma_i^v \subseteq \Gamma_i$, $\Delta_i^v \subseteq \Delta_i$ in the open goals, and construct $\psi = \bigwedge_{i=1}^n (\bigwedge \Gamma_i^v \rightarrow \bigvee \Delta_i^v)$.

The full calculus [25] is moreover able to eliminate arbitrarily nested quantifiers, and can be used similarly to prove validity of sequents with quantifiers. A recent independent evaluation [29] showed that the resulting proof procedure is competitive with state-of-the-art SMT solvers and theorem provers on a wide range of quantified integer problems.

² <http://www.philipp.ruemmer.org/princess.shtml>.

$\frac{\Gamma, [\phi]_L \vdash \Delta \triangleright I \quad \Gamma, [\psi]_L \vdash \Delta \triangleright J}{\Gamma, [\phi \vee \psi]_L \vdash \Delta \triangleright I \vee J} \vee\text{-LEFT}_L$	
$\frac{\Gamma, [\phi]_R \vdash \Delta \triangleright I \quad \Gamma, [\psi]_R \vdash \Delta \triangleright J}{\Gamma, [\phi \vee \psi]_R \vdash \Delta \triangleright I \wedge J} \vee\text{-LEFT}_R$	
$\frac{\Gamma, [\phi]_D, [\psi]_D \vdash \Delta \triangleright I}{\Gamma, [\phi \wedge \psi]_D \vdash \Delta \triangleright I} \wedge\text{-LEFT}_D$	$\frac{\Gamma \vdash [\phi]_D, \Delta \triangleright I}{\Gamma, [\neg\phi]_D \vdash \Delta \triangleright I} \neg\text{-LEFT}_D$
$\frac{*}{\Gamma, [\phi]_L \vdash [\phi]_L, \Delta \triangleright \text{false}} \text{CLOSE}_{LL}$	$\frac{*}{\Gamma, [\phi]_R \vdash [\phi]_R, \Delta \triangleright \text{true}} \text{CLOSE}_{RR}$
$\frac{*}{\Gamma, [\phi]_L \vdash [\phi]_R, \Delta \triangleright \phi} \text{CLOSE}_{LR}$	$\frac{*}{\Gamma, [\phi]_R \vdash [\phi]_L, \Delta \triangleright \neg\phi} \text{CLOSE}_{RL}$

$\frac{\Gamma, [x/t]_L, [\forall x.\phi]_L \vdash \Delta \triangleright I}{\Gamma, [\forall x.\phi]_L \vdash \Delta \triangleright \forall_{Rt} I} \forall\text{-LEFT}_L$	$\frac{\Gamma, [x/t]_R, [\forall x.\phi]_R \vdash \Delta \triangleright I}{\Gamma, [\forall x.\phi]_R \vdash \Delta \triangleright \exists_{Lt} I} \forall\text{-LEFT}_R$
$\frac{\Gamma, [x/c]_D \vdash \Delta \triangleright I}{\Gamma, [\exists x.\phi]_D \vdash \Delta \triangleright I} \exists\text{-LEFT}_D$	$\frac{\Gamma \vdash [x/c]_D, \Delta \triangleright I}{\Gamma \vdash [\forall x.\phi]_D, \Delta \triangleright I} \forall\text{-RIGHT}_D$

Fig. 2 The upper box presents a selection of interpolating rules for propositional logic, while the lower box shows rules for quantifiers. Parameter D stands for either L or R . The quantifier \forall_{Rt} denotes universal quantification over all constants occurring in t but not in $\Gamma_L \cup \Delta_L$; likewise, \exists_{Lt} denotes existential quantification over all constants occurring in t but not in $\Gamma_R \cup \Delta_R$. In $\exists\text{-LEFT}_D$, c is a constant that does not occur in the conclusion

2.3 Craig interpolation in the base logic

Given formulas A and B such that $A \wedge B$ is unsatisfiable, Craig interpolation can determine a formula I such that the implications $A \Rightarrow I$ and $B \Rightarrow \neg I$ hold, and non-logical symbols in I occur in both A and B [30]. An interpolating version of our sequent calculus has been presented in [4,11], and is summarised in Fig. 2. To keep track of the partitions A , B , the calculus operates on labeled formulas $[\phi]_L$ (with L for “left”) to indicate that ϕ is derived from A , and similarly formulas $[\phi]_R$ for ϕ derived from B . If Γ , Δ are finite sets of L/R -labeled formulas, and I is an unlabeled formula, then $\Gamma \vdash \Delta \triangleright I$ is an *interpolating sequent*.

Semantics of interpolating sequents is defined using the following projections: $\Gamma_L =_{\text{def}} \{\phi \mid [\phi]_L \in \Gamma\}$ and $\Gamma_R =_{\text{def}} \{\phi \mid [\phi]_R \in \Gamma\}$, which extract the L/R -parts of a set Γ of labeled formulas. A sequent $\Gamma \vdash \Delta \triangleright I$ is *valid* if 1. the sequent $\Gamma_L \vdash I$, Δ_L is valid, 2. the sequent $\Gamma_R, I \vdash \Delta_R$ is valid, and 3. constants and predicates in I occur in both $\Gamma_L \cup \Delta_L$ and $\Gamma_R \cup \Delta_R$. As a special case, note that the sequent $[A]_L, [B]_R \vdash \emptyset \triangleright I$ is valid iff I is an interpolant of $A \wedge B$. Soundness of the calculus guarantees that the root of a closed interpolating proof is a valid interpolating sequent.

To solve an interpolation problem $A \wedge B$, a prover typically first constructs a proof of $A, B \vdash \emptyset$ using the ordinary calculus from Sect. 2.1. Once a closed proof has been found, it can be lifted to an interpolating proof: this is done by replacing the root formulas A , B with $[A]_L$, $[B]_R$, respectively, and recursively assigning labels to all other formulas as defined by the rules from Fig. 2. Then, starting from the leaves, intermediate interpolants are computed and propagated back to the root, leading to an interpolating sequent $[A]_L, [B]_R \vdash \emptyset \triangleright I$.

3 Solving non-linear constraints

We extend the base logic in three steps: in this section, symbols and rules are added to solve *non-linear diophantine problems*; a second extension is then done in Sect. 4 to handle *arithmetic bit-vector constraints*; and, finally, additional symbols to express *structural bit-vector constraints* are introduced in Sect. 5. All constructions preserve the ability of the calculus to eliminate quantifiers (under certain assumptions) and derive Craig interpolants.

For non-linear constraints, we assume that the set P of predicates contains a distinguished ternary predicate \times , with the intended semantics that the third argument represents the result of multiplying the first two arguments, i.e., $\times(s, t, r) \Leftrightarrow s \cdot t = r$. The predicate \times is clearly sufficient to express arbitrary polynomial constraints by introducing a \times -literal for each product in a formula, at the cost of introducing a linear number of additional constants or existentially quantified variables. We make the simplifying assumption that \times only occurs in negative positions; that means, top-level occurrences will be on the left-hand side of sequents. Positive occurrences can be eliminated thanks to the equivalence $\neg \times(s, t, r) \Leftrightarrow \exists x. (\times(s, t, x) \wedge x \neq r)$.

3.1 Calculus rules for non-linear constraints

We now introduce classes of calculus rules to reason about the \times -predicate. The rules are necessarily incomplete for proving that a sequent is valid, but they are complete for finding counterexamples: if ϕ is a satisfiable quantifier-free formula with \times as the only predicate symbol, then it is possible to construct a proof for $\phi \vdash \emptyset$ that has an open and unprovable goal in pure Presburger arithmetic (by systematically splitting variable domains, Sect. 3.1.4). The rule classes are:

- *Deriving Implied Equalities with Gröbner Bases*: if implied linear equalities can be found using Buchberger’s algorithm these can be added to the proof goal.
- *Interval Constraint Propagation*: if new bounds for constants can be derived from existing bounds these can be added to the proof goal.
- *Cross-Multiplication of Inequalities*:. if two terms are known to be non-negative, then the non-negativity of their product can be added to the proof goal.
- *Interval Splitting*: as a last resort, the proof branch can be split by dividing the possible values for a constant or variable in half.
- \times -*Elimination*: if a occurrence of \times is implied by other literals, it can be eliminated from the proof goal.

3.1.1 Deriving implied equalities with Gröbner bases

The first rule applies standard algebra methods to infer new equalities from multiplication literals. To avoid the computation of more and more complex terms in this process, we restrict the calculus to the inference of *linear* equations that can be derived through computation of a Gröbner basis.³ Given a set $\{\times(s_i, t_i, r_i)\}_{i=1}^n$ of \times -literals and a set $\{e_j = 0\}_{j=1}^m$ of linear equations, the generated ideal $I = \text{Ideal}(\{s_i \cdot t_i - r_i\}_{i=1}^n \cup \{e_j\}_{j=1}^m)$ over rational numbers is the smallest set of rational polynomials that contains $\{s_i \cdot t_i - r_i\}_{i=1}^n \cup \{e_j\}_{j=1}^m$, is closed under addition, and closed under multiplication with arbitrary rational polynomials [31]. Any

³ The set of *all* linear equations implied by a set of \times -literals over integers is clearly not computable, by reduction of Hilbert’s 10th problem.

$f \in I$ corresponds to an equation $f = 0$ that logically follows from the literals, and can therefore be added to a proof goal:

$$\frac{\Gamma, \{\times(s_i, t_i, r_i)\}_{i=1}^n, \{e_j = 0\}_{j=1}^m, f = 0 \vdash \Delta}{\Gamma, \{\times(s_i, t_i, r_i)\}_{i=1}^n, \{e_j = 0\}_{j=1}^m \vdash \Delta} \times\text{-EQ}$$

if f is linear, has integer coefficients, and $f \in I$

To see how this rule can be applied practically, note that the subset of linear polynomials in I forms a rational vector space, and therefore has a finite basis. It is enough to apply $\times\text{-EQ}$ for terms f_1, \dots, f_k corresponding to any such basis, since linear arithmetic reasoning (in the base logic) will then be able to derive all other linear polynomials in I . To compute a basis f_1, \dots, f_k , we can transform $\{s_i \cdot t_i - r_i\}_{i=1}^n \cup \{e_j\}_{j=1}^m$ to a Gröbner basis using Buchberger's algorithm [32], and then apply Gaussian elimination to find linear basis polynomials (or directly by choosing a suitable monomial order).

Example 1 Consider the formula for the square of a sum: $(x + y)^2 = x^2 + 2xy + y^2$. We can show its validity by rewriting it to normal form and constructing a proof. Let $\Pi = \{\times(x, x, c_1), \times(x, y, c_2), \times(y, y, c_3), \times(x + y, x + y, c_4)\}$:

$$\frac{\begin{array}{c} * \\ \vdots \\ \Pi, c_1 + 2c_2 + c_3 - c_4 = 0 \vdash c_4 = c_1 + 2c_2 + c_3 \end{array}}{\Pi \vdash c_4 = c_1 + 2c_2 + c_3} \times\text{-EQ}$$

Here, the $\times\text{-EQ}$ -step is motivated by the fact that the Gröbner basis derived from Π contains the linear polynomial $c_1 + 2c_2 + c_3 - c_4$, from which the desired equation can be derived using linear reasoning (using calculus rules not presented in this paper, see Sect. 2.1).

3.1.2 Interval constraint propagation (ICP)

Our main technique for inequality reasoning in the presence of \times -predicates is interval constraint propagation (ICP) [33]. ICP is a fixed-point computation on the lattice \mathbb{I}^S of functions mapping constants and variables $S = C \cup X$ to intervals \mathbb{I} , and can efficiently approximate the value ranges of symbols. We define the lattice \mathbb{I} of intervals and the lattice \mathbb{I}^S of interval assignments as follows; $S \rightarrow \mathbb{I}$ represents the set of (total) functions from $S = C \cup X$ to \mathbb{I} , and \perp is the distinguished bottom element of \mathbb{I}^S :

$$\begin{aligned} \mathbb{I} &= \{[x, y] \mid x, y \in \mathbb{Z}, x \leq y\} \cup \{(-\infty, \infty)\} \cup \\ &\quad \{(-\infty, y) \mid y \in \mathbb{Z}\} \cup \{(x, \infty) \mid x \in \mathbb{Z}\} \\ \mathbb{I}^S &= (S \rightarrow \mathbb{I}) \cup \{\perp\} \end{aligned}$$

We denote the (point-wise) join and meet on \mathbb{I}^S with \sqcup, \sqcap , respectively.

To define the fixed-point computation, we then introduce abstraction and concretisation functions that connect the lattice \mathbb{I}^S with the powerset lattice $\mathcal{P}(S \rightarrow \mathbb{Z})$ of value assignments. The abstraction of a set $V \in \mathcal{P}(S \rightarrow \mathbb{Z})$ of value assignments is the least element $\alpha(V)$ of \mathbb{I}^S such that the interval $\alpha(V)(c)$ assigned to a symbol $c \in S$ contains all values of c in V (or $\alpha(V) = \perp$ if V is empty). The abstraction function $\alpha : \mathcal{P}(S \rightarrow \mathbb{Z}) \rightarrow \mathbb{I}^S$ is formally defined as follows:

$$\alpha(V) = \bigsqcup_{\beta \in V} \{c \mapsto [\beta(c), \beta(c)] \mid c \in S\} \quad \text{for } V \in \mathcal{P}(S \rightarrow \mathbb{Z}).$$

The concretisation $\gamma(I)$ of some interval assignment $I \in \mathbb{I}^S$ is the set V of all value assignments that stay within the intervals specified by I . More formally, $\gamma : \mathbb{I}^S \rightarrow \mathcal{P}(S \rightarrow \mathbb{Z})$ is defined by:

$$\gamma(I) = \begin{cases} \emptyset & \text{if } I = \perp \\ \{\beta : S \rightarrow \mathbb{Z} \mid \beta(c) \in I(c) \text{ for all } c \in S\} & \text{otherwise} \end{cases} \quad \text{for } I \in \mathbb{I}^S.$$

The result of ICP can then be defined as the greatest fixed-point of a monotonic propagation function $Prop : \mathbb{I}^S \rightarrow \mathbb{I}^S$ on the lattice \mathbb{I}^S . Propagation can be defined separately for each formula occurring in a sequent; in particular, propagation $Prop_\phi : \mathbb{I}^S \rightarrow \mathbb{I}^S$ for a multiplication literal $\phi = \times(s, t, r)$ is defined as:

$$Prop_{\times(s,t,r)}(I) = \alpha(\{\beta \in \gamma(I) \mid \beta \models s \cdot t = r\})$$

This means, propagation eliminates values from the intervals that are inconsistent with $\times(s, t, r)$. Propagation for equalities $t = 0$ and inequalities $t \leq 0$ is defined similarly; in practice, also any monotonic over-approximation of $Prop_\phi$ can be used instead of $Prop_\phi$, at the cost of more over-approximate results in the end.

Given a set $\{\phi_1, \dots, \phi_n\}$ of formulas, the overall propagation function $Prop = Prop_{\{\phi_1, \dots, \phi_n\}}$ is the meet of the individual propagators:

$$Prop_{\{\phi_1, \dots, \phi_n\}}(I) = \bigcap_{i=1}^n Prop_{\phi_i}(I)$$

The ICP rule assumes that a greatest fixed-point $\text{gfp } Prop_{\{\phi_1, \dots, \phi_n\}}$ for equality, inequality, and multiplication literals ϕ_1, \dots, ϕ_n in a sequent has been computed, and adds resulting bounds for a constant c :

$$\frac{\Gamma, \phi_1, \dots, \phi_n, l \leq c, c \leq u \vdash \Delta}{\Gamma, \phi_1, \dots, \phi_n \vdash \Delta} \times\text{-ICP} \\ \text{if } (\text{gfp } Prop_{\{\phi_1, \dots, \phi_n\}})(c) = [l, u]$$

Example 2 From two inequalities $x \geq 5$ and $y \geq 5$, the rule \times -ICP can derive $(x+y)^2 \geq 100$:

$$\frac{\times(x+y, x+y, c_4), x \geq 5, y \geq 5, c_4 \geq 100 \vdash}{\times(x+y, x+y, c_4), x \geq 5, y \geq 5 \vdash} \times\text{-EQ}$$

The slightly different problem $x + y \geq 10 \rightarrow (x + y)^2 \geq 100$ cannot be proven in the same way, since ICP will not be able to deduce bounds for x or y from $x + y \geq 10$.

3.1.3 Cross-multiplication of inequalities

While ICP is highly effective for approximating the range of constants, and quickly detecting inconsistencies, it is less useful for inferring relationships between multiple constants that follow from multiplication literals. We cover such inferences using a *cross-multiplication* rule that resembles procedures used in ACL2 [34]. The rule captures the fact that if s, t are both non-negative, then also the product $s \cdot t$ is non-negative.

Like in Sect. 3.1.1, we prefer to avoid the introduction of new multiplication literals during proof search. By disallowing non-linear terms, we avoid the introduction of more and more complex terms and thus only add $s \cdot t \geq 0$ if the term $s \cdot t$ can be expressed linearly. For this,

we again write $I = \text{Ideal}(\{s_i \cdot t_i - r_i\}_{i=1}^n \cup \{e_j\}_{j=1}^m)$ for the ideal induced by equations and \times -literals:

$$\frac{\Gamma, s \leq 0, t \leq 0, -f \leq 0 \vdash \Delta}{\Gamma, s \leq 0, t \leq 0 \vdash \Delta} \quad \times\text{-CROSS}$$

if f is linear, has integer coefficients, and $s \cdot t - f \in I$

The term f can practically be found by computing a Gröbner basis of I , and reducing the product $s \cdot t$ to check whether an equivalent linear term exists.

3.1.4 Interval splitting

If everything else fails, as last resort it can become necessary to systematically split over the possible values of a variable or constant $c \in C \cup X$:

$$\frac{\Gamma, c \leq \alpha - 1 \vdash \Delta \quad \Gamma, c \geq \alpha \vdash \Delta}{\Gamma \vdash \Delta} \quad \times\text{-SPLIT}$$

The $\alpha \in \mathbb{Z}$ can in principle be chosen arbitrarily in the rule, but in practice a useful strategy is to make use of the range information derived for \times -ICP: when no ranges can be tightened any further using \times -ICP, instead \times -SPLIT can be applied to split one of the intervals in half.

3.1.5 \times -Elimination

Finally, occurrences of \times can be eliminated whenever a formula is subsumed by other literals in a goal, again writing $I = \text{Ideal}(\{s_i \cdot t_i - r_i\}_{i=1}^n \cup \{e_j\}_{j=1}^m)$:

$$\frac{\Gamma \vdash \Delta}{\Gamma, \times(s, t, r) \vdash \Delta} \quad \times\text{-ELIM}$$

if $s \cdot t - r \in I$

Note that \times -ELIM only eliminates non-linear \times -literals, whereas \times -EQ only introduces linear equations, so that the application of the two rules cannot induce cycles.

3.2 Quantifier elimination for non-linear constraints

Due to necessary incompleteness of calculi for Peano arithmetic, quantifiers can in general not be eliminated in the presence of the \times predicate, even when considering formulas that do not contain uninterpreted predicates. By combining the QE approach in Sect. 2.2 with the rules for \times that we have introduced, it is nevertheless possible to reason about quantified non-linear constraints in many practical cases, and sometimes even get rid of quantifiers. This is possible because the rules in Sect. 3.1 are not only sound, but even *equivalence transformations*: in any application of the rules, the conjunction of the premises is equivalent to the conclusion.

Similarly as in [35], QE is always possible if sufficiently many constants or variables in a formula ϕ range over *bounded* domains: if there is a set $B \subseteq C \cup X$ of symbols with bounded domain such that in each literal $\times(s, t, r)$ either s or t contain only symbols from B . In this case, proof construction will terminate when applying the rule \times -SPLIT only to variables or constants with bounded domain. This guarantees that eventually every literal $\times(s, t, r)$ can be turned into a linear equation using \times -EQ, and then be eliminated using \times -ELIM, only

leaving proof goals with pure Presburger arithmetic constraints. The boundedness condition is naturally satisfied for bit-vector formulas.

3.3 Craig interpolation for non-linear constraints

To carry over the Craig interpolation approach from Sect. 2.3 to non-linear formulas, interpolating versions of the calculus rules for the \times -predicate are needed. For this, we follow the approach used in [4] (which in turn resembles the use of theory lemmas in SMT in general): when translating a proof to an interpolating proof, we replace applications of the \times -rules with instantiation of an equivalent theory axiom QAx . Suppose a non-interpolating proof contains a rule application

$$\frac{\begin{array}{c} \vdots \\ \Gamma, \Gamma', \Gamma_1 \vdash \Delta_1, \Delta', \Delta \end{array} \cdots \begin{array}{c} \vdots \\ \Gamma, \Gamma', \Gamma_n \vdash \Delta_n, \Delta', \Delta \end{array}}{\Gamma, \Gamma' \vdash \Delta', \Delta} R \quad (1)$$

in which Γ', Δ' are the formulas assumed by the rule application, Γ, Δ are side formulas not required or affected by the application, and $\Gamma_1, \Delta_1, \dots, \Gamma_n, \Delta_n$ are newly introduced formulas in the individual branches.

The (unquantified) theory axiom Ax corresponding to the rule application expresses that the conjunction of the premises has to imply the conclusion; the quantified theory axiom $QAx =_{\text{def}} \forall S. Ax$ in addition contains universal quantifiers for all constants $S \subseteq C$ occurring in Ax .

$$Ax =_{\text{def}} \bigwedge_{i=1}^n (\bigwedge \Gamma_i \rightarrow \bigvee \Delta_i) \rightarrow (\bigwedge \Gamma' \rightarrow \bigvee \Delta')$$

Ax and QAx are specific to the *application* of R : the axioms for two distinct applications of R will in general be different formulas. QAx is defined in such a way that it can simulate the effect of R (as in (1)). This is done by introducing QAx in the antecedent of a sequent, applying the rule \forall -LEFT to instantiate the axiom with the constants S and obtain Ax , and then applying propositional rules. The propositional rules \vee -LEFT and \neg -LEFT are used to eliminate implications \rightarrow (which are short-hand for \neg, \vee), and the rule \wedge -RIGHT to eliminate the conjunction $\bigwedge_{i=1}^n$:

$$\frac{\begin{array}{c} * \\ \vdots \\ \Gamma, \Gamma', \bigwedge \Gamma' \rightarrow \bigvee \Delta' \vdash \Delta', \Delta \end{array} \quad \frac{\Gamma, \Gamma', \Gamma_1 \vdash \Delta_1, \Delta', \Delta \cdots \Gamma, \Gamma', \Gamma_n \vdash \Delta_n, \Delta', \Delta}{\Gamma, \Gamma', Ax \vdash \Delta', \Delta} \quad \vee\text{-LEFT}, \neg\text{-LEFT}, \wedge\text{-RIGHT} *}{\Gamma, \Gamma', \forall S. Ax \vdash \Delta', \Delta} \forall\text{-LEFT} *$$

This construction leads to a proof using only the standard rules from Sect. 2.1, which can be interpolated as discussed earlier. Since QAx is a valid formula not containing any constants, it can be introduced in a proof at any point, and labelled $[QAx]_L$ or $[QAx]_R$ on demand.

The obvious downside of this approach is the possibility of quantifiers occurring in interpolants. The interpolating rules \forall -LEFT_{L/R} (Fig. 2) have to introduce quantifiers $\forall_{Rt}/\exists_{Lt}$ for local symbols occurring in the substituted term t ; whether such quantifiers actually occur in

the final interpolant depends on the applied \times -rules, and on the order of rule application. For instance, with \times -SPLIT it is always possible to choose the label of QAx so that no quantifiers are needed, whereas \times -EQ might mix symbols from left and right partitions in such a way that quantifiers become unavoidable. In our implementation we approach this issue pragmatically. We leave proof search unrestricted, and might thus sometimes get proofs that do not give rise to quantifier-free interpolants; when that happens, we afterwards apply QE to get rid of the quantifiers. QE is always possible for bit-vector constraints, see Sect. 4.4.⁴

4 Solving bit-vector constraints

We now define the extension of the base logic to bit-vector constraints. The main idea of the extension is to represent bit-vectors of width w as integers in the interval $\{0, \dots, 2^w - 1\}$, and to translate bit-vector operations to the corresponding operation in Presburger arithmetic (or possible the \times -predicate for non-linear formulas), followed by an integer remainder operation to map the result back to the correct bit-vector domain. Since the remainder operation tends to be a bottleneck for interpolation, we keep the operation symbolic and initially consider it as an uninterpreted predicate $bmod_a^b$. The predicate is only gradually reduced to Presburger arithmetic by applying the calculus rules introduced later in this section.

Formally, we introduce binary predicates $P_{bv} = \{bmod_a^b \mid a, b \in \mathbb{Z}, a < b\}$. The semantics of each predicate $bmod_a^b$ is to relate any whole number $s \in \mathbb{Z}$ to its remainder modulo $b - a$ in the interval $\{a, \dots, b - 1\}$:

$$\begin{aligned} bmod_a^b(s, r) &\Leftrightarrow a \leq r < b \wedge \exists z. r = s + (b - a) \cdot z \\ &\Leftrightarrow a \leq r < b \wedge r \equiv s \pmod{b - a} \end{aligned}$$

We also introduce short-hand notations for the casts to the unsigned and signed bit-vector domains:

$$ubmod_w \stackrel{\text{def}}{=} bmod_0^{2^w}, \quad sbmod_w \stackrel{\text{def}}{=} bmod_{-2^{w-1}}^{2^{w-1}}.$$

4.1 Translating bit-vector constraints to the core language

For the rest of the section, we use the base logic augmented with \times and $bmod_a^b$ -predicates as the *core language* to which bit-vector constraints are translated. For presentation, the translation focuses on a subset of the arithmetic bit-vector operations, $BVOP_a = \{bvadd_w, bvmul_w, bvdiv_w, bvneg_w, ze_{w+w'}, bvule_w, bvsle_w\}$. An extension to bit-vector concatenation, extraction, and bit-wise functions is presented in Sect. 5. All operations are sub-scripted with the bit-width of the operands; the zero-extend function $ze_{w+w'}$ maps bit-vectors of width w to width $w + w'$. Semantics follows the FixedSizeBitVectors⁵ theory of the SMT-LIB [36]. Other arithmetic operations, for instance $bvsdiv_w$ or $bvsmod_w$, can be handled in the same way as shown here, though sometimes the number of cases to be considered is larger.

The translation from bit-vector constraints ϕ to core formulas ϕ_{core} has two parts: first, $BVOP_a$ occurrences in a formula ϕ have to be replaced with equivalent expressions in the

⁴ Non-linear integer arithmetic in general does not admit quantifier-free interpolants. For instance, $(x > 1 \wedge x = y^2) \wedge x = z^2 + 1$ is unsatisfiable, but no quantifier-free interpolants exist, regardless of whether divisibility predicates $\alpha \mid t$ are allowed or not.

⁵ <http://www.smtlib.org/theories-FixedSizeBitVectors.shtml>.

$\text{bvadd}_w(s, t) = r$	\rightarrow	$\text{ubmod}_w(s + t, r)$
$\text{bvneg}_w(s) = r$	\rightarrow	$\text{ubmod}_w(-s, r)$
$\text{bvmul}_w(s, t) = r$	\rightarrow	$\exists x. (\times(s, t, x) \wedge \text{ubmod}_w(x, r))$
$\text{ze}_{w+w'}(s) = r$	\rightarrow	$s = r$
$\text{bvsle}_w(s, t)$	\rightarrow	$\exists x, y. (\text{sbmod}_w(s, x) \wedge \text{sbmod}_w(t, y) \wedge x \leq y)$
$\neg \text{bvsle}_w(s, t)$	\rightarrow	$\exists x, y. (\text{sbmod}_w(s, x) \wedge \text{sbmod}_w(t, y) \wedge x > y)$
$\text{bvule}_w(s, t)$	\rightarrow	$s \leq t$
$\neg \text{bvule}_w(s, t)$	\rightarrow	$s > t$
$\text{bvdiv}_w(s, t) = r$	\rightarrow	$(t = 0 \wedge r = 2^w - 1) \vee$ $(t \geq 1 \wedge \exists x. (\times(t, r, x) \wedge s - t < x \leq s))$

Fig. 3 Rules translating bit-vector operations into the core language. The rules only apply in negative positions

core language; second, since the core language only knows the sort of unbounded integers, type information has to be made explicit by adding domain constraints.

BVOP_a Elimination. Like in Sect. 3, we assume that the bit-vector formula ϕ has already been brought into a flat form by introducing additional constants or quantified variables: the operations in BVOP_a must not occur nested, and functions only occur in equations of the form $f(\bar{s}) = t$ in negative positions. The translation from ϕ to ϕ' is then defined by the rewriting rules in Fig. 3. Since the rules for the predicates bvsle_w and bvule_w distinguish between positive and negative occurrences, we assume that rules are only applied to formulas in negation normal-form, and only in negative positions.

The rules for bvadd_w , bvneg_w , $\text{ze}_{w+w'}$, and bvule_w simply translate to the corresponding Presburger term, if necessary followed by remainder ubmod_w . Multiplication bvmul_w is mapped similarly to the \times -predicate defined in Sect. 3, adding an existential quantifier to store the intermediate product. Since rules are only applied in negative positions, the quantified variable can later be replaced with a Skolem constant. An optimised rule could be defined for the case that one of the factors is constant, avoiding the use of the \times -predicate. Translation of bvsle_w maps the operands to a signed bit-vector domain $\{-2^{w-1}, \dots, 2^{w-1} - 1\}$, in which then the arithmetic inequality predicates $\leq, >$ can be used. The rule for unsigned division bvdiv_w distinguishes the cases that the divisor t is zero or positive (as required by SMT-LIB), and maps the latter case to standard integer division.

Domain constraints. Bit-vector variables/constants x of width w occurring in ϕ are interpreted as unbounded integer variables in ϕ_{core} , which therefore has to contain explicit assumptions about the ranges of bit-vector variables. We use the abbreviation $\text{in}_w(x) =_{\text{def}} (0 \leq x < 2^w)$ and define

$$\phi_{\text{core}} = \left(\bigwedge_{x \in S} \text{in}_{w_x}(x) \right) \rightarrow \phi'$$

where $S \subseteq C \cup X$ is the set of free variables and constants occurring in ϕ , w_x is the bit-width of $x \in S$, and ϕ' is the result of applying rules from Fig. 3 to ϕ . Similar constraints are used to express quantification over bit-vectors, for instance $\exists x. (\text{in}_w(x) \wedge \dots)$ and $\forall x. (\text{in}_w(x) \rightarrow \dots)$.

Example 3 Consider `challenge/multiplyOverflow.smt2`, a problem from SMT-LIB QF_BV containing a bit-vector formula that is known to be hard for most SMT solvers

$\frac{\Pi : \phi \rightarrow \phi' \quad \Pi : \psi \rightarrow \psi'}{\Pi : \phi \circ \psi \rightarrow \phi' \circ \psi'} \text{ } \circ\text{-RW}$	
$\frac{\Pi : Lit \rightarrow Lit' \quad \Pi, Lit' : \phi \rightarrow \phi'}{\Pi : Lit \wedge \phi \rightarrow Lit' \wedge \phi'} \text{ } \text{LIT-}\wedge\text{-RW}$	
$\frac{\Pi : Lit \rightarrow Lit' \quad \Pi, \neg Lit' : \phi \rightarrow \phi'}{\Pi : Lit \vee \phi \rightarrow Lit' \vee \phi'} \text{ } \text{LIT-}\vee\text{-RW}$	
$\frac{\Pi : \phi \rightarrow \phi'}{\Pi : \neg \phi \rightarrow \neg \phi'} \text{ } \neg\text{-RW}$	$\frac{\Pi : \phi \rightarrow \phi'}{\Pi : Qx.\phi \rightarrow Qx.\phi'} \text{ } Q\text{-RW}$
$\frac{\left\lfloor \frac{lbound(\Pi, s) - a}{b - a} \right\rfloor = k = \left\lfloor \frac{ubound(\Pi, s) - a}{b - a} \right\rfloor}{\Pi : bmod_a^b(s, r) \rightarrow s = r + k \cdot (b - a)} \text{ } \text{BOUND-RW}$	
$\frac{s + (b - a) \cdot t \prec s}{\Pi : bmod_a^b(s, r) \rightarrow bmod_a^b(s + (b - a) \cdot t, r) \quad bmod_a^{b'}(s', r') \in \Pi, \quad (b - a) \mid k \cdot (b' - a'),}$	
$\frac{s + k \cdot (s' - r') \prec s}{\Pi : bmod_a^b(s, r) \rightarrow bmod_a^b(s + k \cdot (s' - r'), r)} \text{ } \text{BMOD-RW}$	

Fig. 4 Simplification rules for bit-vector formulas. In \circ -RW, ϕ and ψ are not literals, and $\circ \in \{\wedge, \vee\}$. In LIT- \wedge -RW and LIT- \vee -RW, the formula Lit is a literal. In Q -RW, x must not occur in Π , and $Q \in \{\forall, \exists\}$. In COEFF-RW, all constants or variables in t also occur in s

since it contains both multiplication and division. In experiments, neither Z3 nor CVC4 could prove the formula within 10min. In our notation, the problem amounts to showing validity of the following implication, with a, b ranging over bit-vectors of width 32:

$$\text{bvule}_{32}(b, \text{bvdiv}_{32}(2^{32} - 1, a)) \rightarrow \text{bvule}_{64}(\text{bvmul}_{64}(\text{ze}_{32+32}(a), \text{ze}_{32+32}(b)), 2^{32} - 1)$$

As a flat formula, with additional constants c_1 of width 32 and c_2, c_3, c_4 of width 64, the implication takes the form:

$$\left(\text{bvdiv}_{32}(2^{32} - 1, a) = c_1 \wedge \text{bvmul}_{64}(c_3, c_4) = c_2 \wedge \left(\text{ze}_{32+32}(a) = c_3 \wedge \text{ze}_{32+32}(b) = c_4 \wedge \text{bvule}_{32}(b, c_1) \right) \right) \rightarrow \text{bvule}_{64}(c_2, 2^{32} - 1)$$

The final formula ϕ_{core} is obtained by application of the rules in Fig. 3, and adding domain constraints:

$$\left(\begin{aligned} & in_{32}(a) \wedge in_{32}(b) \wedge in_{32}(c_1) \wedge in_{64}(c_2) \wedge in_{64}(c_3) \wedge in_{64}(c_4) \wedge \\ & \left(\left((a = 0 \wedge c_1 = 2^{32} - 1) \vee \right. \right. \\ & \left. \left. (a \geq 1 \wedge \exists x. (\times(a, c_1, x) \wedge 2^{32} - 1 - a < x \leq 2^{32} - 1)) \right) \right) \wedge \\ & \left(\exists z. (\times(c_3, c_4, z) \wedge \text{ubmod}_{64}(z, c_2)) \wedge a = c_3 \wedge b = c_4 \wedge b \leq c_1 \right) \end{aligned} \right) \rightarrow c_2 \leq 2^{32} - 1$$

4.2 Preprocessing and simplification

An encoded formula ϕ_{core} tends to contain a lot of redundancy, in particular nested or unnecessary occurrences of the $bmod_a^b$ predicates. As an important component of our calculus, and in line with the approach in other bit-vector solvers, we therefore apply simplification rules both during preprocessing and during the solving phase (“inprocessing”). The most

important simplification rules are shown in Fig. 4. Our implementation in addition applies rules for Boolean and Presburger connectives, for instance to inline equations $x = t$ or to propagate inequalities, not shown here.

The notation $\Pi : \phi \rightarrow \phi'$ expresses that formula ϕ can be rewritten to ϕ' , given the set Π of formulas as context. The structural rules in the upper half of Fig. 4 define how formulas are traversed, and how the context Π is extended to Π, Lit' when encountering further literals. We apply the structural rules modulo associativity and commutativity of \wedge, \vee , and prioritise LIT- \vee -RW and LIT- \wedge -RW over the other rules. Simplification is iterated until a fixed-point is reached and no further rewriting is possible. The connection between rewriting rules and the sequent calculus is established by the following rules:

$$\frac{\Gamma, \phi' \vdash \Delta}{\Gamma, \phi \vdash \Delta} \text{RW-LEFT} \quad \frac{\Gamma \vdash \phi', \Delta}{\Gamma \vdash \phi, \Delta} \text{RW-RIGHT}$$

if $\Gamma \cup \{\neg\psi \mid \psi \in \Delta\} : \phi \rightarrow \phi'$

The lower half of Fig. 4 shows three of the bit-vector-specific rules. The BOUND- RW rule defines elimination of $bmod_a^b$ -predicates that do not require any case splits; the definition of the rule assumes functions $lbound(\Pi, s)$ and $ubound(\Pi, s)$ that derive lower and upper bounds of a term s , respectively, given the current context Π . The two functions can be implemented by collecting inequalities (and possibly type information available for predicates) in Π to obtain an over-approximation of the range of s .

Rule COEFF- RW reduces coefficients in $bmod_a^b(s, r)$ by adding a multiple of the modulus $b - a$ to s . The rule assumes a well-founded order $<$ on terms to prevent cycles during simplification. One way to define such an order is to choose a total well-founded order $<$ on the union $C \cup X$ of variables and constants, extend $<$ to expressions $\alpha \cdot x$ by sorting coefficients as $0 < 1 < -1 < 2 < \dots$, and finally extend $<$ to arbitrary terms $\alpha_1 t_1 + \dots + \alpha_n t_n$ as a multiset order [25].

The same order $<$ is used in BMOD- RW, defining how $bmod_a^b(s, r)$ can be rewritten in the context of a second literal $bmod_{a'}^{b'}(s', r')$. The rule is useful to optimise the translation of nested bit-vector operations. Assuming $bmod_{a'}^{b'}(s', r')$, the value of $s' - r'$ is known to be a multiple of $b' - a'$, and therefore $k \cdot (s' - r')$ is a multiple of $b - a$ provided that $b - a$ divides $k \cdot (b' - a')$. This implies that the truth value of $bmod_a^b(s, r)$ is not affected by adding $k \cdot (s' - r')$ to s .

Our implementation uses various further simplification rules, for instance to eliminate \times or $bmod_a^b$ whose result is never used; we skip those for lack of space.

Example 4 Consider $bvadd_{32}(bvadd_{32}(a, b), c)$, which corresponds to the expression $ubmod_{32}(a + b, r_1) \wedge ubmod_{32}(r_1 + c, r_2)$ in the core language. Using BMOD- RW, the formula can be rewritten to $ubmod_{32}(a + b, r_1) \wedge ubmod_{32}(a + b + c, r_2)$, provided that $a + b + c < r_1 + c$.

Example 5 We continue Example 3 and show that ϕ_{core} is valid, focusing on the $a \geq 1$ case of $bvdiv_{32}$. The proof (Fig. 5) consists of three core steps: 1. using \times -ICP, from the constraints $in_{32}(a)$, $in_{32}(b)$, $\times(a, b, d)$ the inequalities $0 \leq d$ and $d \leq 2^{64} - 2^{33} + 1$ can be derived; 2. therefore, using RW-LEFT and BOUND- RW, the literal $ubmod_{64}(d, c_2)$ can be rewritten to $d = c_2$, capturing the fact that 64-bit multiplication cannot overflow for unsigned 32-bit operands; 3. using \times -CROSS, from the inequalities $a \geq 1$ and $b \leq c_1$ we derive $(a - 1)(c_1 - b) = ac_1 - ab - c_1 + b \geq 0$. Using the products $\times(a, b, d)$ and $\times(a, c_1, e)$, we can express it linearly as $e - d - c_1 + b \geq 0$. The proof branch can then be

$$\begin{array}{c}
\vdots \\
\dots, a \geq 1, e < 2^{32}, b \leq c_1, d \geq 2^{32}, e - d - c_1 + b \geq 0 \vdash \\
\dots, \times(a, b, d), \times(a, c_1, e), \underline{a \geq 1}, \underline{e < 2^{32}}, \underline{b \leq c_1}, \underline{d \geq 2^{32}} \vdash \quad \times\text{-CROSS} \\
\dots, 0 \leq d, d \leq 2^{64} - 2^{33} + 1, d = c_2 \vdash \quad (b) \\
\dots, 0 \leq d, d \leq 2^{64} - 2^{33} + 1, \underline{ubmod_{64}(d, c_2)} \vdash \quad \text{RW-LEFT} \\
\dots, \underline{in_{32}(a)}, \underline{in_{32}(b)}, \times(a, b, d), \underline{ubmod_{64}(d, c_2)} \vdash \quad \times\text{-ICP} \\
\vdash \phi_{core} \quad (a)
\end{array}$$

Fig. 5 Proof tree for Example 5, with the sequences (a), (b) of rule applications not shown in detail

closed using standard arithmetic reasoning. The implementation of our procedure can easily find the outlined proof automatically.

4.3 Splitting rules for $bmod_a^b$

In general, formulas will of course also contain occurrences of $bmod_a^b$ that cannot be eliminated just by simplification. We introduce two calculus rules for reasoning about such general literals $bmod_a^b(s, r)$. The first rule makes the assumption that lower and upper bounds of s are available, and are reasonably tight, so that an explicit case analysis can be carried out; the rule generalises BOUND- RW to the situation in which the factors l, u do not coincide:

$$\frac{\left\{ \Gamma, a \leq r < b, s = r + i \cdot (b - a) \vdash \Delta \right\}_{i=l}^u}{\Gamma, bmod_a^b(s, r) \vdash \Delta} \quad \text{BMOD- SPLIT}$$

assuming the bounds $\lfloor \frac{lbound(\Pi, s) - a}{b - a} \rfloor = l$ and $\lfloor \frac{ubound(\Pi, s) - a}{b - a} \rfloor = u$ with $\Pi = \Gamma \cup \{\neg\psi \mid \psi \in \Delta\}$.

If the bounds l, u are too far apart, the number of cases created by BMOD- SPLIT would become unmanageable, and it is better to choose a direct encoding of the remainder operation in Presburger arithmetic:

$$\frac{\Gamma, a \leq r < b, s = r + (b - a) \cdot c \vdash \Delta}{\Gamma, bmod_a^b(s, r) \vdash \Delta} \quad \text{BMOD- CONST}$$

where c is assumed to be a fresh constant. Rule BMOD- CONST corresponds to the encoding chosen in [16].

In practice, it turns out to be advantageous to prioritise rule BMOD- SPLIT over BMOD- CONST, as long as the number of cases does not become too big. This is because each of the premises of BMOD- SPLIT tends to be significantly simpler to solve (and interpolate) than the conclusion; in addition, splitting one $bmod_a^b$ literal often allows subsequent simplifications that eliminate other $bmod_a^b$ occurrences. We investigate experimentally in Sect. 6.1 how many applications of the rules BMOD- SPLIT and BMOD- CONST are needed to prove formulas satisfiable or unsatisfiable, and show that the numbers are surprisingly low, in particular in the unsatisfiable case.

4.4 Quantifier elimination and Craig interpolation

Since the bit-vector rules in this section are all equivalence transformations, QE for bit-vectors can be done exactly as described in Sect. 3.2. As the ranges of all symbols are now bounded,

it is guaranteed that any formula will eventually be reduced to Presburger arithmetic, so that we obtain complete QE for (arithmetic) bit-vector constraints.

Similarly, the interpolation approach from Sect. 3.3 carries over to bit-vectors, with theory axioms being generated for each of the rules defined in this section. Since the translation of bit-vector formulas to the core language happens upfront, also interpolants are guaranteed to be in the core language, and can be mapped back to bit-vector formulas if necessary (e.g., as in [16]). Interpolants might contain quantifiers, in which case QE can be applied (as described in the first paragraph), so that we altogether obtain a complete procedure for quantifier-free interpolation of arithmetic bit-vector formulas.

In our implementation, we restrict the use of the simplification rules RW-LEFT and RW-RIGHT when computing proofs for the purpose of interpolation. Unrestricted use could quickly mix up the vocabularies of the individual partitions in an interpolation problem $A \wedge B$, and thus increase the likelihood of quantifiers in interpolants. Instead we simplify A, B separately upfront using rules in Fig. 4, and apply RW-LEFT, RW-RIGHT only when the modified formula ϕ is a literal.

Example 6 We recall the example from Sect. 1.1, and show how our calculus finds the simpler interpolant $I'_{LIA} = y_3 < y_2$ for the interpolation problem $A \wedge B$. The core step is to turn the application of BMOD-SPLIT into an explicit axiom; after slight simplifications, this axiom is:

$$Ax = \frac{(ubmod_w(y_2 + 1, c_2) \wedge 3 \leq y_2 < 256 \wedge in_8(c_2))}{(y_2 + 1 = c_2 \vee y_2 + 1 = c_2 + 256)}$$

The axiom mentions all assumptions made by the rule, including the bounds $3 \leq y_2 < 256$ that determine the number of resulting cases (or, alternatively, the formulas $c_1 > y_3, y_2 = c_1, c_2 \leq y_3, y_7 = 3, y_7 = c_2$ from which the bounds derive). The axiom also includes domain constraints like $in_8(c_2)$ for occurring symbols, which later ensures that possible quantifiers in interpolants range over bounded domains. The quantified axiom is $QAx = \forall y_2, c_2. Ax$, and can be used to construct an interpolating proof:

$$\frac{\begin{array}{c} \vdots \\ * \\ \vdots \\ \dots \frac{[c_1 > y_3]_L, [y_2 = c_1]_L, [c_2 \leq y_3]_R, \quad \vdash \emptyset \blacktriangleright y_3 < y_2 \quad \dots}{[y_7 = 3]_R, [y_7 = c_2]_R, [y_2 + 1 = c_2]_R} \end{array}}{\frac{[A_{core}]_L, [B_{core}]_R, [Ax]_R \vdash \emptyset \blacktriangleright y_3 < y_2}{[A_{core}]_L, [B_{core}]_R, [QAx]_R \vdash \emptyset \blacktriangleright y_3 < y_2} \forall\text{-LEFT}_R} \forall\text{-LEFT}_R$$

We only show one of the cases, \mathcal{P} , resulting from splitting the axiom $[Ax]_R$ using the rules from Fig. 2. The final interpolant $I_{LAZY} = y_3 < y_2$ records the information needed from A_{core} to derive a contradiction in the presence of $y_2 + 1 = c_2$; the branch is closed using standard arithmetic reasoning [11].

5 Interpolation in the presence of extract and concat

Two bit-vector operations which are more tricky to translate to integer arithmetic are the extraction of bits from a larger bit-vector, and the concatenation of two bit-vectors. This is formalised using the function $bvextract_{[u,l]}$, which cuts out a slice of $u - l + 1$ consecutive bits, and the function $bvconcat_{v+w}$ forming a bit-vector of length $v + w$. We call the fragment of bit-vector logic containing only these operations, and positive equalities, the *structural fragment* of bit-vector theory.

Fig. 6 Illustration of Example 8.

The conflict is due to the contradicting assignments to slices of x and y . The upper table shows the situation with bit-vectors of size 8, the lower one with size 16

	7	6	5	4	3	2	1	0
$x_{[8]}$	*	*	1	0	1	1	*	*
$y_{[8]}$	0	0	0	1	1	0	*	*

	15	14	13	12	11	10	9	...
$x_{[16]}$	*	*	1	0	1	1	*	...
$y_{[16]}$	0	0	0	1	1	0	*	...

Example 7 Consider the following bit-vectors:

$$\begin{aligned} c_1 &= [0, 0] = 0 & c_2 &= [1, 1] = 3 & c_3 &= [0, 1] = 1 \\ c_4 &= [0, 0, 1, 1] = 3 & c_5 &= [1, 1, 0, 0] = 12 \end{aligned}$$

where a bit-sequence $[b_n, \dots, b_0]$ is represented by the number $\sum_{i=0}^n b_i 2^i$. The following equations hold between those bit-vectors:

$$\begin{aligned} \text{bvconcat}_{2+2}(c_1, c_2) &= c_4 & \text{bvconcat}_{2+2}(c_2, c_1) &= c_5 \\ \text{bvextract}_{[3,2]}(c_4) &= c_1 & \text{bvextract}_{[1,0]}(c_4) &= c_2 \\ \text{bvextract}_{[2,1]}(c_4) &= c_3. \end{aligned}$$

While it is possible to translate extractions and concatenations to integer arithmetic, it is not always efficient. We show here that it can be more efficient to keep extractions abstract and only convert at need. This approach can also help to compute simpler interpolants.

Example 8 We consider the example from [19], a formula over the structural fragment, which we divide into two parts

$$\begin{aligned} A &= (x[5 : 0] = z \wedge z[5 : 2] = 11) \\ B &= (y[7 : 2] = 6 \wedge x = y) \end{aligned}$$

where the bit-vectors x, y are of width 8, and z is of width 6. The unsatisfiability of $A \wedge B$ could be proved, as before, by translating the bit-vector constraints to our core language:

$$\begin{aligned} A_{\text{core}} &= \text{in}_8(x) \wedge \text{in}_6(z) \wedge \text{ubmod}_6(x, z) \wedge \exists c_1. (\text{in}_2(c_1) \wedge z = 11 \cdot 2^2 + c_1) \\ B_{\text{core}} &= \text{in}_8(x) \wedge \text{in}_8(y) \wedge \exists c_2. (\text{in}_2(c_2) \wedge y = 6 \cdot 2^2 + c_2) \wedge x = y \end{aligned}$$

Using the interpolation procedure presented in Sect. 4, we can compute the following interpolant:

$$I = \exists c. (x = 44 + 64c \vee x = 45 + 64c \vee x = 46 + 64c \vee x = 46 + 64c)$$

The conjunction is unsatisfiable due to the conflicting assignment of a small slice in x and y , as illustrated in Fig. 6. However, when translating to integer arithmetic the overall structure is lost, and the interpolant contains a lot of redundancy. The situation gets worse with increasing bit-widths. Consider a formula where the width of x, y are doubled, while the widths of the extractions are kept constant:

$$\begin{aligned} A' &= (x[13 : 8] = z \wedge z[5 : 2] = 11) \\ B' &= (y[15 : 10] = 6 \wedge x = y) \end{aligned}$$

The same procedure now yields an interpolant of exponentially greater size:

$$I' = \exists c. (x = 4097 + 8192c \vee x = 4098 + 8192c \vee \dots \vee x = 5120 + 8192c)$$

$$\begin{aligned}
& \text{bvextract}_{[u,l]}(s) = r \rightarrow \text{extr}_l^u(s, r) \\
& \text{bvconcat}_{w+v}(s, t) = r \rightarrow \text{in}_{w+v}(r) \wedge \text{extr}_v^{w+v-1}(r, s) \wedge \text{extr}_0^{v-1}(r, t) \\
& \text{bvnot}_w(s) = r \rightarrow \text{in}_w(r) \wedge \bigwedge_{i=0}^{w-1} \exists x. (\text{extr}_i^i(s, x) \wedge \text{extr}_i^i(r, 1-x)) \\
& \text{bvand}_w(s, t) = r \rightarrow \text{in}_w(r) \wedge \bigwedge_{i=0}^{w-1} \exists s_b, t_b, r_b. \left(\text{extr}_i^i(s, s_b) \wedge \text{extr}_i^i(t, t_b) \wedge \text{extr}_i^i(r, r_b) \right. \\
& \quad \left. \wedge r_b \leq s_b \wedge r_b \leq t_b \wedge r_b \geq s_b + t_b - 1 \right) \\
& \text{bvor}_w(s, t) = r \rightarrow \text{in}_w(r) \wedge \bigwedge_{i=0}^{w-1} \exists s_b, t_b, r_b. \left(\text{extr}_i^i(s, s_b) \wedge \text{extr}_i^i(t, t_b) \wedge \text{extr}_i^i(r, r_b) \right. \\
& \quad \left. \wedge r_b \geq s_b \wedge r_b \geq t_b \wedge r_b \leq s_b + t_b \right) \\
& \text{bvxor}_w(s, t) = r \rightarrow \text{in}_w(r) \wedge \bigwedge_{i=0}^{w-1} \exists s_b, t_b, r_b. \left(\text{extr}_i^i(s, s_b) \wedge \text{extr}_i^i(t, t_b) \wedge \text{extr}_i^i(r, r_b) \right. \\
& \quad \left. \wedge \text{ubmod}_1(s_b + t_b, r_b) \right)
\end{aligned}$$

Fig. 7 Rules translating structural and bit-wise operations into the extended core language. As before, the rules are only applied in negative positions. Note that u, l, v, w are integer constants

We will explain in the next sections how more succinct interpolants can be computed by eliminating the extraction operation only lazily.

5.1 The structural fragment

In [19], a polynomial fragment of the bit-vector theory is identified, consisting of formulas that only contain extractions, concatenations, and positive equalities. The satisfiability of formulas in the structural fragment is decidable in polynomial time by a congruence closure procedure over decomposed bit-vectors [18].

Definition 1 The *structural fragment* consists of bit-vector formulas of the form $\phi_1 \wedge \phi_2 \wedge \dots \wedge \phi_n$ where each ϕ_i is an equation constructed using bit-vector variables, concrete bit-vectors, and the operators $\text{bvextract}_{[u,l]}$ and bvconcat_{v+w} .

Note that a formula containing bvconcat_{v+w} can be translated into an equi-satisfiable formula that only uses $\text{bvextract}_{[u,l]}$. We illustrate the translation with an example, and introduce the formal rule in Fig. 7. Consider the formula $\phi[\text{bvconcat}_{v+w}(s, t)]$ containing the concatenation of two bit-vectors. The concatenation can be eliminated by introducing an existentially quantified variable to represent the result of the concatenation; the relationship with the arguments is established using extraction terms:

$$\exists x. (\text{in}_{v+w}(x) \wedge \text{bvextract}_{[v+w-1,w]}(x) = s \wedge \text{bvextract}_{[w-1,0]}(x) = t \wedge \phi[x])$$

To define a formal calculus for the structural fragment, we introduce an *extended* core language by adding a further family $P_{ex} = \{\text{extr}_l^u \mid u, l \in \mathbb{N}, u \geq l\}$ of predicates representing extraction from bit-vectors. Semantically, $\text{extr}_l^u(s, t)$ relates s and t if t is the result of extracting the bits u through l from s . This is formally expressed as the existence of two bit-vectors x, y such that t can be made equal to s by prepending x and appending y :

$$\text{extr}_l^u(s, t) \Leftrightarrow \text{in}_{u-l+1}(t) \wedge \exists x, y. (\text{in}_l(y) \wedge s = x \cdot 2^{u+1} + t \cdot 2^l + y)$$

Note that the argument s is not bounded, which implies that the definition contains a bound for the lower slice y , but not for x .

Fig. 8 Cut point propagation of bit-vectors x , y and z . The top table contains each bit-vector and the corresponding cut points. Since all bit-vectors are related by some constraints, all cut points (within the width of the bit-vector) are propagated, e.g., the cut point for y between bit 2 and 1 is propagated to x

$x :$	7	6	5	4	3	2	1	0
$y :$	7	6	5	4	3	2	1	0
$z :$			5	4	3	2	1	0

↓

$x :$	7	6	5	4	3	2	1	0
$y :$	7	6	5	4	3	2	1	0
$z :$			5	4	3	2	1	0

The rewriting rules in Fig. 7 define how extraction and concatenation are translated to the P_{ex} predicates, following the same schema as in Sect. 4. As a side-effect of adding $bvextract_{[u,l]}$ and $bvconcat_{w+v}$, and moving beyond the structural fragment, the calculus can also reason about the bit-wise operators $BVOP_{bv} = \{bvnot, bvand, bvor, bvxor\}$, by extracting the individual bits of the operands and encoding the Boolean semantics using inequalities. The rewriting rules for such an encoding are given Fig. 7 as well, and are also used in our implementation.

5.2 Bit-vector decomposition

As demonstrated in Sect. 1.2, unsatisfiability of formulas can sometimes be proven by just focusing on the right slice of bit-vectors; the challenge lies in how to decompose the bit-vectors to find the conflicts. Intuitively, there is no need to split apart bits which are never constrained individually. We follow the procedure described in [19], and use the notion of *cut points* for this reason. Cut points of a bit-vector variable determine the slices that need to be considered, and the points at which the bit-vectors might have to be decomposed, and are determined by the boundaries of extraction operations.

More formally, given a formula ϕ in the structural fragment over set $S = C \cup X$ of constants and variables, a cut point configuration is a function $\mathcal{C} : S \rightarrow \mathcal{P}(\mathbb{N})$ satisfying the following properties:

- for each $extr_l^u(s, t)$ literal, it is the case that:
 - $\{l, u + 1\} \subseteq \mathcal{C}(s)$, and
 - $\{i - l \mid i \in \mathcal{C}(s) \text{ with } l \leq i \leq u + 1\} = \{i \in \mathcal{C}(t) \mid i \leq u - l + 1\}$.
- for each equality $s = t$ it is the case that $\mathcal{C}(s) = \mathcal{C}(t)$.

The set of cut points for all bit-vectors can be obtained by a fix-point computation, which begins with all cut points from extractions, and then propagates using the equalities until all conditions hold.

Example 9 Translated to our extended core language, the constraints from Example 8 are:

$$\begin{aligned}
 A'_{\text{core}} &= in_8(x) \wedge in_6(z) \wedge extr_0^5(x, z) \wedge extr_2^5(z, 11) \\
 B'_{\text{core}} &= in_8(x) \wedge in_8(y) \wedge extr_2^7(y, 6) \wedge x = y
 \end{aligned}$$

The extraction literals induce immediate cut points for x , y and z , respectively: $\{6, 0\}$, $\{8, 2\}$, and $\{6, 2\}$. Since x and z are related by the literal $extr_0^5(x, z)$, the cut point 2 needs to be added to the set for x as well; similarly, due to the equation $x = y$, also the cut points $\{6, 0\}$ have to be added for y , and 8 for x . The alignment of the bit-vectors is illustrated in Fig. 8, and the complete sets of cut points after propagation are $\{8, 6, 2, 0\}$, $\{8, 6, 2, 0\}$, and $\{6, 2, 0\}$.

$\frac{\Gamma, [\exists x_1. (extr_i^u(s, x_1) \wedge extr_{i-l}^{u-l}(r, x_1) \wedge in_{u-i+1}(x_1))]_D \vdash \Delta \triangleright I}{\Gamma, [extr_i^u(s, r)]_D \vdash \Delta \triangleright I} \text{ EXTR-SPLIT}_D$	
$\frac{\Gamma, [\exists x. (in_l(x) \wedge ubmod_{u+1}(s, r2^l + x))]_D \vdash \Delta \triangleright I}{\Gamma, [extr_i^u(s, r)]_D \vdash \Delta \triangleright I} \text{ EXTR-ARITH}_D$	
$\frac{\Gamma \vdash s_1 = s_2, \Delta \quad \Gamma, r_1 = r_2 \vdash \Delta}{\Gamma, extr_i^u(s_1, r_1), extr_i^u(s_2, r_2) \vdash \Delta} \text{ EXTR-CC}$	
$\frac{\Gamma \vdash [s_1 = s_2]_L, \Delta \triangleright I \quad \Gamma, [r_1 = r_2]_L \vdash \Delta \triangleright J}{\Gamma, [extr_i^u(s_1, r_1)]_L, [extr_i^u(s_2, r_2)]_L \vdash \Delta \triangleright I \vee J} \text{ EXTR-CC}_{LL}$	
$\frac{\Gamma \vdash [s_1 = s_2]_R, \Delta \triangleright I \quad \Gamma, [r_1 = r_2]_R \vdash \Delta \triangleright J}{\Gamma, [extr_i^u(s_1, r_1)]_R, [extr_i^u(s_2, r_2)]_R \vdash \Delta \triangleright I \wedge J} \text{ EXTR-CC}_{RR}$	
$\frac{\Gamma \vdash [s_1 = s_2]_R, \Delta \triangleright I \quad \Gamma, [r_1 = r_2]_R \vdash \Delta \triangleright J}{\Gamma, [extr_i^u(s_1, r_1)]_L, [extr_i^u(s_2, r_2)]_R \vdash \Delta \triangleright \exists_{Ls_1r_1}. (extr_i^u(s_1, r_1) \wedge I \wedge J)} \text{ EXTR-CC}_{LR}$	
$\frac{\Gamma \vdash [s_1 = s_2]_L, \Delta \triangleright I \quad \Gamma, [r_1 = r_2]_L \vdash \Delta \triangleright J}{\Gamma, [extr_i^u(s_1, r_1)]_R, [extr_i^u(s_2, r_2)]_L \vdash \Delta \triangleright \forall_{Rs_1r_1}. (\neg extr_i^u(s_1, r_1) \vee I \vee J)} \text{ EXTR-CC}_{RL}$	
$\frac{\Gamma, [\phi]_L \vdash \Delta \triangleright I \quad \Gamma \vdash [\phi]_R, \Delta \triangleright J}{\Gamma \vdash \Delta \triangleright (I \vee \neg \phi) \wedge J} \text{ CUT}_{LR}$	
$\frac{\Gamma, [\phi]_R \vdash \Delta \triangleright I \quad \Gamma \vdash [\phi]_L, \Delta \triangleright J}{\Gamma \vdash \Delta \triangleright \neg I \wedge \phi \vee J} \text{ CUT}_{RL}$	

Fig. 9 Rules for handling extraction operations in bit-vector formulas. In EXTR-SPLIT and EXTR-ARITH, $D \in \{L, R\}$. In EXTR-SPLIT, $l < i \leq u$. In EXTR-CC_{LR}, $\exists_{Ls_1r_1}$ denotes existential quantification over all constants occurring in s_1 or t_1 but not in $\Gamma_R \cup \Delta_R \cup \{s_2, r_2\}$. In EXTR-CC_{RL}, $\forall_{Rs_1r_1}$ denotes universal quantification over all constants occurring in s_1 or t_1 but not in $\Gamma_L \cup \Delta_L \cup \{s_2, r_2\}$. The rules CUT_{LR} and CUT_{RL} are only allowed for formulas ϕ in which all constants are common to both $\Gamma_L \cup \Delta_L$ and $\Gamma_R \cup \Delta_R$

If the bit-vectors x, y, z are decomposed according to their cut points, then simple reasoning reveals the inconsistency between $extr_2^5(x, 5)$, $extr_2^5(y, 6)$, and $x = y$.

5.3 An interpolating calculus for extractions

A decomposition according to the cut points yields a complete and polynomial procedure for the structural fragment [19]. We formalise the splitting of bit-vector extractions with an interpolating calculus rule EXTR-SPLIT in Fig. 9. Intuitively, we cut the bit-vector s at point i and introduce two existential variables x_1, x_2 corresponding to the two slices. By constraining the corresponding slices of r according to the decomposition, we ensure the original equality between the extraction of s and r . The rule can be generalised to split into several slices at once, allowing for shorter proofs.

After extractions have been split at the cut points, we rely on congruence closure to take care of $extr_i^u$ literals, in a similar way as the procedure in [18]. Congruence closure is in our setting expressed by an axiom schema for functional consistency of the $extr_i^u$ predicates:

$$\forall s_1, s_2, r_1, r_2. (extr_i^u(s_1, r_1) \wedge extr_i^u(s_2, r_2) \wedge s_1 = s_2 \rightarrow r_1 = r_2) \quad (2)$$

The axiom states that two *extr*-literals will yield the same result if the first arguments coincide, and if the same bits are extracted. Similar axioms for predicate consistency and functional consistency are used in [4]. To simplify presentation, we model the instantiation of (2) using the calculus rule EXTR-CC in Fig. 9; the figure also gives four interpolating versions of the rule, for the four possible combinations of *L/R*-labels. The *LR/RL* rules differ in the label used in their premises. The rules can be derived by instantiating (2) using either \forall -LEFT_{*L*} or \forall -LEFT_{*R*}, in a similar way as in Sect. 3.3. In practice, and in our implementation, the calculus rules are used as in classical SMT-style congruence closure: they are triggered when the first arguments of a pair of *extr*-literals have become equal.

Extraction operations can also be translated to arithmetic, which is needed to evaluate extraction from concrete numbers, and when constructing proofs for formulas that are not in the structural fragment (i.e., that combine extraction with other bit-vector operations). The rule EXTR-ARITH encodes an operation $\text{extr}_l^u(s, r)$ using a modulo constraint to eliminate bits above position *u*, and division to strip away bits below position *l*. We express the division by existentially quantifying the remainder. A more direct rule to perform evaluation is introduced in Sect. 5.4.

Example 10 We continue Example 9, and show how an interpolating proof can be constructed for the conjunction $A'_{\text{core}} \wedge B'_{\text{core}}$. The root sequent of the proof is $\lfloor A'_{\text{core}} \rfloor_L, \lfloor B'_{\text{core}} \rfloor_R \vdash \emptyset$. In the proof tree shown in Fig. 10, we first split the given formulas using Boolean rules. Then, the literal $\text{extr}_0^5(x, z)$ can be decomposed using our EXTR-SPLIT rule at the cut point 2, and congruence closure is applied to the literals $\text{extr}_2^5(z, c_1)$ and $\text{extr}_2^5(z, 11)$. We similarly decompose the literal $\text{extr}_2^7(y, 6)$ at cut point 6, and then use EXTR-ARITH to reduce $\text{extr}_0^3(6, c_2)$ to $c_2 = 6$. Finally, we need a second application of congruence closure, EXTR-CC_{*LR*}, to relate the extractions from *x* and *y*.

The resulting interpolant is the formula $\exists c. (\text{extr}_2^5(x, c) \wedge c = 11)$. The quantifier in the formula stems from the application of EXTR-CC_{*LR*}, which transfers the local symbol *c*₁ from *L* to *R*, and thus makes it shared. A quantifier is needed to eliminate the symbol again from the interpolant. However, as can be seen the quantifier naturally disappears when translating the interpolant back to functional notation, which yields the formula $\text{bvextract}_{[5,2]}(x) = 11$ in the structural fragment.

It is quite easy to see that our calculus is sound and complete for formulas in the structural fragment. The calculus does not guarantee, however, that interpolants computed for formulas in the structural fragment are again in the structural fragment, or that interpolants are quantifier-free. This leads to the question whether the structural fragment is actually closed under interpolation, i.e., whether every unsatisfiability conjunction has an interpolant that is again in the fragment. The answer to this question is positive, and it turns out that our calculus can also be used to compute such interpolants, if the right strategy is used to construct proofs.

Theorem 1 *The structural fragment is closed under interpolation.*

Proof We give a simple proof that follows from the fact that our calculus can model bit-blasting of bit-vector formulas, and builds on work on EUF interpolation in [3,4]. The existence of more compact interpolants, avoiding the need to blast to individual bits, can also be shown, but this is slightly more involved.

Suppose $A \wedge B$ is an unsatisfiable conjunction in the structural fragment, and $A_{\text{core}} \wedge B_{\text{core}}$ the translation to the (extended) core language. Further assume that x_1, \dots, x_m are the shared bit-vector constants of *A*, *B*, and that w_1, \dots, w_m are their bit-widths, respectively. This means that we are searching for an interpolant in the structural fragment that may only

with $I = \exists c.(c = 11 \wedge \text{extr}_2^5(x, c))$, or in the structural fragment $\text{bvextract}_{[5,2]}(x) = 11$.

contain the constants x_1, \dots, x_m . To model bit-blasting, we augment $A_{\text{core}}, B_{\text{core}}$ by adding fresh names $\{b_i^j, c_i^j\}_{i,j}$ for the individual bits of x_1, \dots, x_m :

This means that b_i^j/c_i^j is the name of the j th bit of x_i in the L/R partition. Note that a formula I_{core} is an interpolant of $A_{\text{core}} \wedge B_{\text{core}}$ iff it is an interpolant of $A'_{\text{core}} \wedge B'_{\text{core}}$, because no shared symbols are added, and we can therefore work with the latter conjunction.

We then construct a proof for the root sequent $[A'_{\text{core}}]_L, [B'_{\text{core}}]_R \vdash \emptyset$ in our interpolating calculus by systematically applying the calculus rules. Rules are applied likewise to the L - and the R -formulas, so that we only mention the L -versions at this point for sake of brevity:

-
- Springer

- EXTR- CC_{LL} whenever two extracts $\lfloor extr_j^i(s, r) \rfloor_L$, $\lfloor extr_j^i(t, r') \rfloor_L$ in combination with an equation $\lfloor s = t \rfloor_L$ occur, followed by an application of $CLOSE_{LL}$ to close the left premise; again, this generates an equation $\lfloor r = r' \rfloor_L$;
- EXTR- $ARITH_L$ whenever an extract $\lfloor extr_j^i(\alpha, r) \rfloor_L$ from a concrete number $\alpha \in \mathbb{Z}$ occurs, followed by arithmetic rules to simplify the generated formula to an equation $\lfloor r = 0 \rfloor_L$ or $\lfloor r = 1 \rfloor_L$;
- EXTR- $ARITH_L$ whenever an extract $\lfloor extr_0^0(s, r) \rfloor_L$ in combination with the domain constraint $in_1(s)$ occurs, i.e., when the single bit of a bit-vector of width 1 is extracted, followed by arithmetic rules to simplify the generated formula to an equation $\lfloor r = s \rfloor_L$.

After those rule applications, we can focus on the obtained bit-level equations in the proof goal: on equations $\lfloor s = t \rfloor_D$ or $\lfloor s = \alpha \rfloor_D$ in which s, t have width 1 (i.e., $in_1(s)$ and $in_1(t)$) and $\alpha \in \{0, 1\}$, with $D \in \{L, R\}$. By construction, the set of L -labelled bit-level equations is equi-satisfiable to the original formula A'_{core} , and the R -labelled equations are equi-satisfiable to B'_{core} , so that we can continue constructing a bit-level proof using the equations.

Since the conjunction $A'_{core} \wedge B'_{core}$ is by assumption unsatisfiable, there are three possible cases:

(i) the equations labelled with L are by themselves unsatisfiable, the interpolant is *false*, and the proof can be closed by applying arithmetic rules;

(ii) symmetrically, the equations labelled with R are unsatisfiable, and the interpolant is *true*; or

(iii) the equations are unsatisfiable only in combination.

In case (iii), there has to be a chain of equations, alternating between L - and R -equations, that witnesses unsatisfiability. There are several symmetric cases, of which we only consider one:

$$\underbrace{0 = \dots = b_{i_1}^{j_1}}_{L\text{-equations}} \doteq \underbrace{c_{i_1}^{j_1} = \dots = c_{i_2}^{j_2}}_{R\text{-equations}} \doteq \underbrace{b_{i_2}^{j_2} = \dots = b_{i_3}^{j_3}}_{L\text{-equations}} \doteq \dots \doteq \underbrace{c_{i_k}^{j_k} = \dots = 1}_{R\text{-equations}} \quad (3)$$

In the other cases, the chain can start in R and end in L , or start and end in the same partition. The dotted equations $b_i^j \doteq c_i^j$ are implied by the literal $\lfloor extr_j^j(x_i, b_i^j) \rfloor_L$, $\lfloor extr_j^j(x_i, c_i^j) \rfloor_R$, but do not exist explicitly in the proof goal.

From (3), it is easy to read off an interpolant for $A \wedge B$ in the structural fragment by summarising the L -chains:

$$I = \text{bvextract}_{[j_1, j_1]}(x_{i_1}) = 0 \wedge \bigwedge_{l=2}^{k-1} \text{bvextract}_{[j_l, j_l]}(x_{i_l}) = \text{bvextract}_{[j_{l+1}, j_{l+1}]}(x_{i_{l+1}})$$

Clearly, I follows from A , and $I \wedge B$ is unsatisfiable due to (3).

To construct an interpolant mechanically in our calculus, there are several strategies. The simplest one is to apply the interpolating cut-rule CUT_{RL} to each the formulas $extr_{j_1}^{j_1}(x_{i_1}, 0)$ and $\{\exists z. (extr_{j_l}^{j_l}(x_{i_l}, z) \wedge extr_{j_{l+1}}^{j_{l+1}}(x_{i_{l+1}}, z))\}_{l=2}^{k-1}$, which yields an interpolant I_{core} in the core language that corresponds to the structural interpolant I shown above. \square

5.4 A rewriting rule for constant extraction

Given an extraction $extr_{j_l}^{j_l}(s, r)$ and bounds on s , it is in some cases possible to determine value of extracted bits. For example, the longest prefix on which the lower and upper bound agree is guaranteed to be present in any consistent value of s . Therefore, extractions that

overlap with that prefix yield some bit values of the extraction without knowing the exact value of s . We allow rewriting if the extraction operator falls entirely within the common prefix:

$$\frac{\begin{array}{l} (\text{lbound}(\Pi, s) \text{ xor } \text{ubound}(\Pi, s)) < 2^l \wedge \\ c = (\text{lbound}(\Pi, s) \text{ rem } 2^u) \text{ div } 2^l \wedge 0 \leq c < 2 \end{array}}{\Pi : \text{extr}_l^u(s, r) \rightarrow r = c} \text{ EXTR-CONST}$$

where *rem* and *div* are the integer remainder and division, respectively. The rule EXTR-CONST allows in particular evaluation of extractions from constant bit-vectors.

5.5 Splitting of disequalities

As shown above, proofs can be closed by finding contradicting assignments to (a slice of) a bit-vector. In general, formulas can also contain bit-vector disequalities, i.e., negative equalities between bit-vectors. As an optimisation, disequalities can be split using the notion of cut points as well. Given a formula with a disequality $s \neq t$, we extend the notion of cut point configurations (Sect. 5.2) by also propagating between s and t . For a cut point $i \in \mathcal{C}(s) = \mathcal{C}(t)$, we can then replace the disequality with a disjunction of two disequalities, as expressed by the following rule:

$$\frac{\begin{array}{l} \Gamma, \text{in}_w(s), \text{in}_w(t), \text{in}_{w-i}(c), \text{in}_{w-i}(d), \text{extr}_i^{w-1}(s, c), \text{extr}_i^{w-1}(t, d) \vdash c = d, \Delta \\ \Gamma, \text{in}_w(s), \text{in}_w(t), \text{in}_i(c), \text{in}_i(d), \text{extr}_0^{i-1}(s, c), \text{extr}_0^{i-1}(t, d) \vdash c = d, \Delta \end{array}}{\Gamma, \text{in}_w(s), \text{in}_w(t) \vdash s = t, \Delta} \neq\text{-SPLIT}$$

The constants c, d must be fresh and not occur in the conclusion in this rule.

6 Experiments

To evaluate the effectiveness of the approach, the procedures described in this article have been implemented in the PRINCESS theorem prover⁶ [25]. The implementation of the full SMT-LIB theory of bit-vectors in PRINCESS is still an ongoing effort, and at this point includes fairly refined versions of the calculi for non-linear arithmetic (Sect. 3) and for arithmetic bit-vector operators (Sect. 4). The implementation of the calculus for the structural fragment (Sect. 5) has been added more recently, and still lacks many optimisations that could be applied. Support for bit-wise operations (like *bvand*) is also quite naïve at the moment, and simply bit-blasts each bit-wise operation separately by introducing *bvextract*_[i,i] terms for the individual bits, as shown in Fig. 7. A more refined encoding would choose, for each sub-expression, whether the arithmetic encoding or bit-blasting should be applied, but this refinement is left for future work. The implementation also supports the SMT-LIB shift operators, which are handled by splitting over the possible values of the second argument. The SMT-LIB rotation operators are not supported yet; those operators are over-approximated as uninterpreted functions, which means that it might be possible to prove problems involving the operators unsatisfiable, but not satisfiable.

All experiments were done on an AMD Opteron 2220 SE machine, running 64-bit Linux and Java 1.8. Runtime was limited to 10min wall clock time, and heap space to 2GB. We used PRINCESS version 2019-10-02 for all experiments. Where runtimes are reported, we use wall clock time.

⁶ <http://www.philipp.ruemmer.org/princess.shtml>.

We evaluate the performance of our approach in three different ways:

- Sect. 6.1: performance of satisfiability queries on quantifier-free bit-vector formulas (SMT-LIB QF_BV), in comparison to the state-of-the-art solvers Z3 4.8.0 [37] and CVC4 1.6 [38].
- Section 6.2: performance of satisfiability queries on bit-vector formulas with quantifiers (SMT-LIB BV), again with comparison to Z3 and CVC4.
- Section 6.3: applicability of the interpolation procedure for software model checking, using the integration in the Horn solver ELDARICA. We compare to the software model checker CPACHECKER 1.7 [39], which internally uses MATHSAT 5 [20] and the interpolation method from [16].

6.1 Satisfiability queries on quantifier-free formulas

While our procedure is not specifically designed for just checking satisfiability of formulas, it is nevertheless interesting to evaluate how the approach performs on problems from the QF_BV category of the SMT-LIB. Results for this category are given in Table 1, and show that our implementation can overall solve a decent number of benchmarks, but is not competitive with Z3 and CVC4 on most of the benchmark families. As a general trend, and unsurprisingly, it can be observed that our lazy arithmetic encoding works relatively well for problems that use arithmetic bit-vector operators, but does not pay off for problems that are mostly Boolean, or problems involving bit-wise operators.

Our implementation can solve the “challenge/multiplyOverflow.smt2” problem discussed in Example 3, and it performs particularly well on the “brummayerbiere4” and “pspace” families, which contain benchmarks with large bit-widths, with variables with up to 30 000 bits. This is to be expected, since our arithmetic encoding is essentially agnostic about the bounds of bit-vector variables, so that the complexity of a problem hardly changes when adding more bits.

The family “bruttomesso/core” consists of SMT-LIB problems in the structural fragment from Sect. 5 (with additional Boolean structure). Compared to the implementation described in the FMCAD 2018 paper [17], the performance on those problems has improved significantly as a result of adding the procedure described in Sect. 5. In 2018, only 8 benchmarks from the “core” family could be solved, compared to 142 in the new version. This is still lower than the results for Z3 and CVC4, which is likely due to more efficient Boolean reasoning.

We also investigated how often the rules \times -SPLIT, BMOD-SPLIT, and BMOD-CONST for splitting and eliminating predicates were applied on the benchmarks. We first determined how many of the problems required the application of the rules at all:

	Total solved	BMOD- SPLIT	BMOD- CONST	\times - SPLIT
SAT	7331	2699	268	334
UNSAT	15,148	647	130	44

The statistics show that a large number of the benchmarks can indeed be solved without those rules. This is in particular the case for unsatisfiable problems, for which it is apparently largely sufficient to work with the simplification rules from Fig. 4, in combination with the rules for Presburger arithmetic, (non-splitting) multiplication, and extraction. In Fig. 11 we compare the required number of applications of BMOD- SPLIT and BMOD- CONST for the

individual benchmarks; the scatter plot shows that often a small number of rule applications is sufficient.

The runtimes reported in Table 1 for PRINCESS are somewhat higher than those of Z3 and CVC4, which can partly be explained by the fact that PRINCESS is entirely implemented in Scala, and runs on a JVM. This results in repeated overhead for starting up the JVM and for just-in-time compilation. In actual applications, for instance software model checking as discussed in Sect. 6.3, normally many queries are handled without restarts in between, and the amortised overhead is smaller.

6.2 Satisfiability queries on formulas with quantifiers

We evaluate the effectiveness of our quantifier elimination approach on problems from the BV category of SMT-LIB. In order to check whether a quantified bit-vector formula is satisfiable, QE often does not have to be run to completion, instead the elimination approach from Sect. 2.2 can be stopped as soon as a statement about satisfiability of the resulting formula can be made. This incremental approach to solving quantified formulas has been implemented in PRINCESS for Presburger arithmetic, and in combination with our lazy encoding for bit-vectors also directly applies to quantified bit-vector formulas.

Results on the SMT-LIB BV benchmarks are given in Table 2. Our procedure can solve a similar number of problems as Z3 and CVC4 on many of the BV families, although the total number of problems solved is still lower than for Z3 and CVC4. Like for QF_BV, the results confirm that the encoding of bit-vectors into arithmetic is more effective for problems that are arithmetic in nature (e.g., the families “Automizer,” “model,” and “Heizmann”), than for more combinatorial problems (e.g., “psyco”). In general, quantified bit-vector problems tend to be smaller and harder than quantifier-free problems, which leads to a situation where it is essential to have the right heuristics and optimisations in place; in this respect our implementation is clearly still lagging behind Z3 and CVC4.

In the experiments, we used a simple portfolio mode enabled by the option `-portfolio=bv`. This mode is inspired by the observation that a closed bit-vector formula ϕ (i.e., a formula without free variables or uninterpreted predicates) can be shown to be satisfiable also by proving that the negation $\neg\phi$ is unsatisfiable, and vice versa. Experiments showed that often one of ϕ or $\neg\phi$ is significantly simpler to solve than the other, but that it is difficult to predict the easier one; in the portfolio mode the prover therefore simultaneously tries to solve ϕ and $\neg\phi$.

6.3 Interpolation and verification of C programs

The main purpose of our procedure is the computation of Craig interpolants for bit-vector formulas. Unfortunately, comparing and evaluating interpolation procedures is relatively tricky, since the properties that can be measured easily (e.g., the size, shape, or strength of interpolants, or the time required to extract interpolants) are ultimately only of limited importance in applications. The decisive property that makes interpolants useful is the ability to generalise, which is hard to measure syntactically. Adding to this, there is no standard set of interpolation benchmarks that could be used, and the interpolation queries that occur in model checking can differ from run to run. In model checking, moreover interpolation queries are interdependent: the results of earlier queries in a run will affect the later interpolation queries being generated.

Table 1 Performance on SMT-LIB QF_BV Problems. For each row, the first/second value gives sat/unsat problems. Experiments were done with PRINCESS 2019-10-02, default settings

Family	PRINCESS		Z3		CVC4	
	Solved	Time (s)	Solved	Time (s)	Solved	Time (s)
2.-BuchwaldFried			0/1	-/292	0/1	-/0.7
2.-Hansen-Check	1/2	0.9/0.7	1/2	0.0/0.0	1/2	0.0/0.0
asp	1/3	9.8/14.5	202/68	86.5/67.6	104/22	130.5/82.4
RWS			16/0	16.5/-	16/0	34.0/-
VS3			2/0	217/-		
bench_ab	284/0	2.2/-	285/0	0.0/-	285/0	0.0/-
bmc-bv	10/8	5.3/45.0	15/15	2.2/13.6	15/12	7.5/35.8
bmc-bv-svcomp14	0/17	-/65.5	8/56	1.5/6.0	8/54	67.1/47.2
brummayerbiere	0/6	-/40.6	0/40	-/30.7	0/37	-/37.6
brummayerbiere2	0/1	-/181	4/25	46.1/90.4	6/43	61.3/54.6
brummayerbiere3	1/0	574/-	5/37	112/97.5	6/12	0.7/85.9
brummayerbiere4	9/0	32.4/-	10/0	0.0/-		
bruttomesso						
.../simple_processor	0/1	-/38.1	0/64	-/53.1	0/64	-/1.2
.../core	0/142	-/56.3	0/672	-/1.2	0/672	-/16.1
.../lfsr			0/225	-/69.9	0/240	-/16.8
calypto	1/2	276/4.7	4/7	0.5/8.6	4/10	13.2/4.3
challenge	0/1	-/1.4				
check2	3/2	1.2/0.9	3/3	0.0/0.0	3/3	0.0/0.0
crafted	2/18	1.0/2.6	2/19	0.0/0.0	2/19	0.0/0.0
dwp_formulas	136/175	12.3/6.4	154/178	0.0/0.0	154/178	0.0/0.0
ecc			0/8	-/0.1	0/8	-/1.9
fft			5/4	40.8/212	1/0	5.7/-
float	0/1	-/270	59/53	116/89.6	38/28	75.6/131
galois			0/1	-/0.2	0/1	-/0.4
gulwani-pldi08			6/0	14.8/-	6/0	35.2/-
log-slicing			0/51	-/287	0/17	-/352
mcm			36/16	122/232	14/0	40.4/-
pspace	21/42	1.4/1.3	0/13	-/366	21/42	0.0/0.0
rubik			3/4	40.1/21.5	0/2	-/73.1
sage	6601/14658	8.2/23.6	8077/18530	0.3/0.1	8077/18530	0.3/1.7
spear						
.../openldap.v2.3.35			3/0	0.1/-	5/0	138/-
.../samba.v3.0.24	34/0	46.6/-	1373/13	8.8/2.4	1343/13	19.9/16
.../zebra.v0.95a	9/0	16.3/-	9/0	1.9/-	9/0	4.0/-
.../xinetd.v2.3.14	0/2	-/2.4	0/2	-/1.4	0/2	-/1.6
.../cvs.v1.11.22	0/5	-/5.3	24/5	4.9/7.4	24/5	14.4/4.5
.../wget.v1.10.2	5/0	8.1/-	38/4	60.7/3.7	36/4	74.8/9.1
.../inn.v2.4.3	163/0	14.2/-	219/0	13.7/-	204/0	29.1/-
stp			1/0	22.4/-	1/0	288/-
stp_samples	48/35	114/27.9	151/273	0.0/0.0	151/273	0.3/0.5
tacas07	2/0	6.9/-	3/0	1.0/-	3/2	22.1/186
uclid/catchconv	0/16	-/35.5	262/152	2.2/0.9	262/152	7.79/9.6
uclid/tcas	0/2	-/1.4	0/2	-/0.0	0/2	-/0.1
uclid.c.smtcomp09			0/6	-/218	0/6	-/383
uum			0/2	-/5.6	0/1	-/2.2
wienand-cav2008						
.../Booth			0/2	-/12.0	0/2	-/16.6
.../Distrib	0/6	-/2.9	0/6	-/0.0	0/6	-/0.0
.../Commute	0/3	-/4.0	0/6	-/0.0	0/6	-/0.0
Total						
SAT	7331	9.14	10980	4.69	10799	5.53
UNSAT	15148	23.7	20565	3.19	20471	3.42

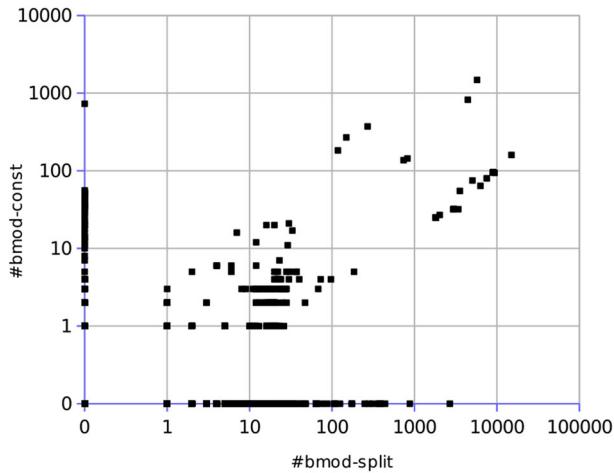


Fig. 11 Scatter plot comparing the number of applications of the rules BMOD-SPLIT and BMOD-CONST on QF_BV benchmarks

Table 2 Performance on SMT-LIB BV problems

Category	Problems	PRINCESS		Z3		CVC4	
		Solved	Time (s)	Solved	Time (s)	Solved	Time (s)
Automizer	16	16	5.7	16	0.1	14	0.1
	137	137	7.5	137	0.0	137	0.3
keymaera	4035	3	5.6	108	6.9	34	1.0
		3754	1.2	3923	0.3	3921	0.1
psyco	132	4	24.0	132	0.1	132	1.5
	62	3	90.8	62	0.2	62	0.5
tptp	17	16	1.2	17	0.0	17	0.0
	56	54	1.3	56	0.0	56	0.0
RND	100	5	5.2	40	6.9	24	39.5
		2	2.3	28	6.7	22	13.2
RNDPRE	130	13	30.6	20	19.0	22	26.9
		15	47.4	36	14.1	26	29.3
model	144	138	8.3	144	0.0	73	10.8
	0	0		0		0	
Heizmann	131	17	32.9	15	37.8	18	18.1
		45	37.7	17	50.7	108	8.3
ranking	60	9	8.4	34	4.4	32	1.5
		4	10.2	19	19.5	13	0.4
fixpoint	131	34	16.0	36	0.5	54	14.2
		30	55.8	73	0.6	75	2.3
Total	5151	255	11.6	562	3.9	420	8.8
		4044	2.5	4351	0.7	4420	0.6

For each family, the first/second row gives sat/unsat problems. Several of the families contains benchmarks with unknown status; for those families only the total number of benchmarks is specified. Experiments were done with PRINCESS 2019-10-02 and option `-portfolio=bv`

We therefore decided to evaluate in an application-oriented way, by integrating our interpolation procedure into a model checker and measuring its ability to verify safety properties of C programs with machine integer semantics. As model checker we use ELDARICA version 2.0.2⁷ [40], a Horn clause-based model checker that uses Cartesian predicate abstraction and the CEGAR algorithm. ELDARICA was already previously tightly integrated with PRINCESS, and was in the scope of this work extended to also handle Horn clauses over bit-vectors. Since ELDARICA internally uses the PRINCESS data-structures to store Horn clauses, we could implement the translation from bit-vectors to our core language (Sect. 4.1) as a preprocessing step that is applied to all Horn clauses upfront. This means that the actual model checking engine operates purely on expressions in the core language, and all interpolation queries and implication checks stay within the core language; the need to translate back and forth between bit-vector formulas and core language is eliminated. As an obvious downside of this approach, however, it is no longer easily possible to replace the interpolation procedure with other solvers.

Benchmarks. For the experiments, we used the built-in C parser of ELDARICA, and work with the benchmark set of 551 C programs already used in [41] for evaluating different predicate generation strategies. The programs stem from a variety of sources, including the SV-COMP 2016 categories “Integers and Control Flow” and “Loops,” and were selecting by taking all programs that do not include arrays or heap data structures (i.e., only arithmetic operations). The verification task consisted in showing that safety assertions included in the programs can never fail. For the experiment, we interpret the programs as operating either on the unbounded mathematical integers (*math*, ELDARICA option `-arithMode:math`), or on signed 32-bit bit-vectors (*ilp32*, ELDARICA option `-arithMode:ilp32`) with wrap-around semantics. ELDARICA was otherwise run with default settings, which means that it also applies the interpolation abstraction technique from [42].

Comparison *math* versus *ilp32*. The results for *math* and *ilp32* semantics are given in Table 3 and Fig. 12. It has to be pointed out that the status of the programs depends on the chosen semantics: for instance, the 46 HOLA programs [43] are all known to be safe in mathematical semantics, but several of the programs turn out to be unsafe in bit-vector semantics due to the possibility of overflow. ELDARICA can consistently verify safety of more programs in *math* than in *ilp32*, but it can disprove safety in more of the *ilp32* cases. The total number of solved cases is higher in *math* than in *ilp32*, but *ilp32* is quite close (403 vs. 337); given the higher complexity of the bit-vector semantics, this is an encouraging result. The scatter plot in Fig. 12 shows that the runtimes for the two semantics are strongly correlated, and while *ilp32* is on average slower than *math* the difference is relatively small.

Table 3 also shows that the number of CEGAR iterations is comparable for *math* and *ilp32*, while the size of interpolants (measured as the average number of sub-formulas of interpolants) is bigger for *ilp32* than for *math*, but usually by less than a factor of 2. The exception is the category “llreve/unsafe,” where drastically bigger interpolants are computed for *ilp32* than for *math*. Inspecting this case, we found that there was a single benchmark in “llreve” that was solved after 579 seconds with interpolants of size 2133; when removing this outlier, the average interpolant size for “llreve/unsafe” is only 1.1.

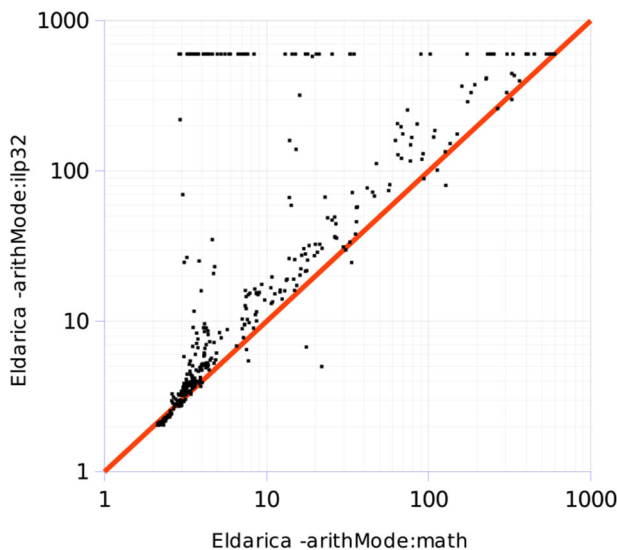
Comparison with CPACHECKER. As comparison, we also ran the model checker CPACHECKER 1.7 [39], using options `-predicateAnalysis -32` and MATHSAT 5 [20] as solver. MATHSAT 5 uses the interpolation method from [16]. The results are given in Table 4 and Fig. 13. Our method is competitive with CPACHECKER on all consid-

⁷ <https://github.com/uuverifiers/eldarica>.

Table 3 Comparison of ELDARICA configurations *math* and *ilp32*

Category	Programs	ELDARICA <i>math</i>				ELDARICA <i>ilp32</i>			
		Solved	Time (s)	Iter.	I. Size	Solved	Time (s)	Iter.	I. Size
HOLA	46	43 0	5.3	4.1	1.1	21 5	7.1 17.3	4.8 41.0	2.0 1.3
llreve	21	16 5	9.8 7.6	13.0 6.8	1.3 1.0	8 5	13.4 120.6	14.8 8.6	2.0 427.5*
VeriMAP	155	133 21	4.4 6.7	2.4 4.2	1.0 1.0	100 39	3.8 14.2	1.5 7.4	1.2 1.4
SVCOMP	329	107 78	53.7 93.5	22.3 41.0	1.0 1.0	89 70	65.7 66.7	21.6 30.7	1.4 1.2
Total	551	299	22.5	10.3	1.1	218	29.8	10.5	1.5
		104	71.8	31.9	1.0	119	49.7	21.2	23.7**

For each category, the table shows the number of safe/unsafe results, and for the solved cases the average time, the required number of CEGAR iterations, and the average size of computed interpolants. For *, after removing an outlier the number is 1.1, and ** becomes 1.3

**Fig. 12** Comparison of the ELDARICA runtime in seconds for *math* and *ilp32* semantics


ered categories: ELDARICA with *ilp32* can consistently prove more programs safe, whereas CPACHECKER can show more programs unsafe, with a lower number of CEGAR iterations. We suspect that the use of large-block encoding [44] in CPACHECKER is responsible for this phenomenon, and indeed makes CPACHECKER very effective for bug finding. The runtimes of the systems are on average close, but the scatter plot in Fig. 13 shows no clear trend.

Altogether, we remark that we are comparing different verification systems here: although both ELDARICA and CPACHECKER apply CEGAR and interpolation, there are many factors affecting the results. What the experiments do show, however, is that the interpolation method proposed in this paper can be used to create a software model checker that is competitive with the state of the art.

Table 4 Comparison of ELDARICA configuration *ilp32* and CPAchecker

Category	Programs	ELDARICA <i>ilp32</i>			CPAchecker -32		
		Solved	Time (s)	Iter.	Solved	Time (s)	Iter.
HOLA	46	21	7.1	4.8	12	80.8	20.1
		5	17.3	41.0	4	8.5	7.0
llreve	21	8	13.4	14.8	7	19.5	14.7
		5	120.6	8.6	5	33.8	7.0
VeriMAP	155	100	3.8	1.5	87	9.4	3.3
		39	14.2	7.4	33	20.9	6.1
SVCOMP	329	89	65.7	21.6	74	37.5	12.1
		70	66.7	30.7	126	49.0	9.0
Total	551	218	29.8	10.5	180	26.1	8.5
		119	49.7	21.2	168	42.1	8.3

For each category, the table shows the number of safe/unsafe results, and for the solved cases the average time, and the required average number of CEGAR iterations

**Fig. 13** ELDARICA versus CPAchecker runtime in seconds for *ilp32* semantics

7 Conclusions

We have presented a new calculus for Craig interpolation and quantifier elimination in bit-vector arithmetic. Furthermore, we have shown how to efficiently integrate reasoning over the structural fragment. While the experimental results in model checking are already promising, we believe that there is still a lot of room for extension and improvement of the approach. This includes more powerful propagation and simplification rules, and more sophisticated strategies to apply the splitting rules \times -SPLIT and BMOD-SPLIT. Future work also includes more efficient use of bounds, and a strategy to employ bit-blasting directly to whole sub-expressions when deemed more efficient.

Acknowledgements This research has been supported by the Swedish Research Council (VR) under the Grants 2014-5484 and 2018-04727, by the Swedish Foundation for Strategic Research (SSF) under the project WebSec (Ref. RIT17-0011), and by a grant from Microsoft. We also wish to thank the reviewers for their extensive and helpful feedback.

Funding Open access funding provided by Mälardalen University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- McMillan KL (2005) An interpolating theorem prover. *Theor Comput Sci* 345(1):101–121
- D'Silva V, Purandare M, Weissenbacher G, Kroening D (2010) Interpolant strength. In: *VMCAI, LNCS*. Springer
- Fuchs A, Goel A, Grundy J, Krstić S, Tinelli C (2009) Ground interpolation for the theory of equality. In: *TACAS, LNCS*. Springer
- Brillout A, Kroening D, Rümmer P, Wahl T (2011) Beyond quantifier-free interpolation in extensions of Presburger arithmetic. In: *VMCAI, LNCS*. Springer, pp 88–102
- McMillan KL (2008) Quantified invariant generation using an interpolating saturation prover. In: Ramakrishnan CR, Rehof J (eds) *Tools and algorithms for the construction and analysis of systems, TACAS 2008*, vol 4963 of *Lecture Notes in Computer Science*. Springer, pp 413–427
- Kovács L, Voronkov A (2009) Interpolation and symbol elimination. In: *CADE*, pp 199–213
- Bonacina MP, Johansson M (2015) On interpolation in automated theorem proving. *J. Autom. Reason.* 54(1):69–97. <https://doi.org/10.1007/s10817-014-9314-0>
- Kapur D, Majumdar R, Zarba CG (2006) Interpolation for data structures. In: *SIGSOFT'06/FSE-14, ACM, New York, NY, USA*, pp 105–116. <https://doi.org/10.1145/1181775.1181789>
- Hojjat H, Rümmer P (2017) Deciding and interpolating algebraic data types by reduction. In: Jelebean T, Negru V, Petcu D, Zaharie D, Ida T, Watt SM (eds) *19th international symposium on symbolic and numeric algorithms for scientific computing, SYNASC 2017, Timisoara, Romania, September 21–24, IEEE Computer Society, 2017*, pp 145–152. <https://doi.org/10.1109/SYNASC.2017.00033>
- Dai L, Xia B, Zhan N (2013) Generating non-linear interpolants by semidefinite programming. In: Sharygina N, Veith H (eds) *Computer aided verification—25th international conference, CAV 2013, Saint Petersburg, Russia, July 13–19, 2013. Proceedings*, vol 8044 of *Lecture Notes in Computer Science*. Springer, pp 364–380. https://doi.org/10.1007/978-3-642-39799-8_25
- Brillout A, Kroening D, Rümmer P, Wahl T (2011) An interpolating sequent calculus for quantifier-free Presburger arithmetic. *J Autom Reason* 47:341–367
- Griggio A, Le TTH, Sebastiani R (2010) Efficient interpolant generation in satisfiability modulo linear integer arithmetic. *Log Methods Comput Sci*. [https://doi.org/10.2168/LMCS-8\(3:3\)2012](https://doi.org/10.2168/LMCS-8(3:3)2012)
- Bruttomesso R, Ghilardi S, Ranise S (2012) Quantifier-free interpolation of a theory of arrays. *Log Methods Comput Sci*. [https://doi.org/10.2168/LMCS-8\(2:4\)2012](https://doi.org/10.2168/LMCS-8(2:4)2012)
- Totla N, Wies T (2016) Complete instantiation-based interpolation. *J. Autom Reason* 57(1):37–65. <https://doi.org/10.1007/s10817-016-9371-7>
- Hoenicke J, Schindler T (2018) Efficient interpolation for the theory of arrays. *CoRR abs/1804.07173*. [arXiv:1804.07173](https://arxiv.org/abs/1804.07173)
- Griggio A (2011) Effective word-level interpolation for software verification. In: Bjesse P, Slobodová A (eds) *International conference on formal methods in computer-aided design, FMCAD '11, Austin, TX, USA, October 30–November 02, 2011, FMCAD Inc.*, pp 28–36
- Backeman P, Rümmer P, Zeljic A Bit-vector interpolation and quantifier elimination by lazy reduction. In: Björner and Gurfinkel, vol 45, pp. 1–10. <https://doi.org/10.23919/FMCAD.2018.8603023>
- Cyrluk D, Möller O, Rueß H (1997) An efficient decision procedure for the theory of fixed-sized bit-vectors. In: Grumberg O (ed) *Computer aided verification*. Springer, Berlin, pp 60–71

19. Bruttomesso R, Sharygina N (2009) A scalable decision procedure for fixed-width bit-vectors. In: Proceedings of the 2009 international conference on computer-aided design, ICCAD '09, ACM, New York, pp 13–20. <https://doi.org/10.1145/1687399.1687403>
20. Cimatti A, Griggio A, Schaafsma BJ, Sebastiani R (2013) The MathSAT5 SMT solver. In: TACAS, vol 7795 of LNCS
21. Asadi S, Blichla M, Fedyukovich G, Hyvärinen AEJ, Even-Mendoza K, Sharygina N, Chockler H (2018) Function summarization modulo theories. In: Barthe G, Sutcliffe G, Veanes M (eds) LPAR-22. 22nd international conference on logic for programming, artificial intelligence and reasoning, Awassa, Ethiopia, 16–21 November 2018, vol 57 of EPIc Series in Computing, EasyChair, pp 56–75
22. Kroening D, Weissenbacher G (2007) Lifting propositional interpolants to the word-level. In: FMCAD, IEEE Computer Society, pp 85–89
23. Kroening D, Weissenbacher G (2009) An interpolating decision procedure for transitive relations with uninterpreted functions. In: Haifa verification conference, vol 6405 of Lecture Notes in Computer Science. Springer, pp 150–168
24. Ho Y, Chauhan P, Roy P, Mishchenko A, Brayton RK (2016) Efficient uninterpreted function abstraction and refinement for word-level model checking. In: Piskac R, Talupur M (eds), 2016 formal methods in computer-aided design, FMCAD 2016, Mountain View, CA, USA, October 3–6, IEEE, 2016, pp 65–72. <https://doi.org/10.1109/FMCAD.2016.7886662>
25. Rümmer P (2008) A constraint sequent calculus for first-order logic with linear integer arithmetic. In: LPAR, vol 5330 of LNCS. Springer, pp 274–289
26. Fitting MC (1996) First-order logic and automated theorem proving, 2nd edn. Springer, New York
27. Halpern JY (1991) Presburger arithmetic with unary predicates is Π_1^1 complete. J Symbol Log 56:637–642
28. Nieuwenhuis R, Oliveras A, Tinelli C (2006) Solving SAT and SAT modulo theories: from an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). J ACM 53(6):937–977
29. Reynolds A, King T, Kuncak V (2017) Solving quantified linear arithmetic by counterexample-guided instantiation. Formal Methods Syst Des 51(3):500–532
30. Craig W (1957) Linear reasoning: a new form of the Herbrand–Gentzen theorem. J Symbol Log 22(3):250–268
31. Lang SS (1993) Algebra, 3rd edn. Addison-Wesley, Reading
32. Buchberger B (3, 2006) An algorithm for finding the basis elements in the residue class ring modulo a zero dimensional polynomial ideal, Ph.D. thesis
33. Van Hentenryck P, McAllester D, Kapur D (1997) Solving polynomial systems using a branch and prune approach. SIAM J Numer Anal 34(2):797–827
34. Warren J, Hunt A, Krug RB, Moore JS (2003) Linear and nonlinear arithmetic in ACL2. In: Proceedings, correct hardware design and verification methods, 12th IFIP WG 10.5 advanced research working conference, vol 2860 of LNCS. Springer, pp 319–333
35. Borralleras C, Lucas S, Oliveras A, Rodríguez-Carbonell E, Rubio A (2012) SAT modulo linear arithmetic for solving polynomial constraints. J. Autom Reason 48(1):107–131
36. Barrett C, Fontaine P, Tinelli C (2017) The SMT-LIB Standard: Version 2.6, Technical report, Department of Computer Science, The University of Iowa. www.SMT-LIB.org
37. de Moura LM, Bjørner N (2008) Z3: an efficient SMT solver. In: TACAS, vol 4963 of LNCS. Springer, pp 337–340
38. Barrett C, Conway CL, Deters M, Hadarean L, Jovanovic D, King T, Reynolds A, Tinelli C (2011) CVC4. In: CAV, vol 6806 of LNCS. Springer, pp 171–177. https://doi.org/10.1007/978-3-642-22110-1_14
39. Beyer D, Keremoglu ME (2009) CPAchecker: a tool for configurable software verification. CoRR abs/0902.0019. [arXiv:0902.0019](https://arxiv.org/abs/0902.0019)
40. Hojjat H, Rümmer P The ELDARICA Horn solver. In: Bjørner and Gurfinkel, vol 45, pp 1–7. <https://doi.org/10.23919/FMCAD.2018.8603013>
41. Demyanova Y, Rümmer P, Zuleger F (2017) Systematic predicate abstraction using variable roles. In: Barrett C, Davies M, Kahsay T (eds) NASA formal methods—9th international symposium, NFM 2017, Moffett Field, CA, USA, May 16–18, 2017, Proceedings, vol 10227 of Lecture Notes in Computer Science, pp 265–281. https://doi.org/10.1007/978-3-319-57288-8_18
42. Leroux J, Rümmer P, Subotic P (2016) Guiding Craig interpolation with domain-specific abstractions. Acta Inf 53(4):387–424. <https://doi.org/10.1007/s00236-015-0236-z>
43. Dillig I, Dillig T, Li B, McMillan KL (2013) Inductive invariant generation via abductive inference. In: Hosking AL, Eugster PT, Lopes CV (eds), Proceedings of the 2013 ACM SIGPLAN international conference on object oriented programming systems languages & applications, OOPSLA. ACM, pp 443–456. <https://doi.org/10.1145/2509136.2509511>
44. Beyer D, Cimatti A, Griggio A, Keremoglu ME, Sebastiani R (2009) Software model checking via large-block encoding. In: Proceedings of 9th international conference on formal methods in computer-aided

design, FMCAD 2009, 15–18 November 2009, Austin, Texas, USA. IEEE, pp 25–32. <https://doi.org/10.1109/FMCAD.2009.5351147>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.