King's Research Portal

DOI:

[10.1145/3359986.3361214](https://doi.org/10.1145/3359986.3361214)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Even Mendoza, K., Hyyvarinen, A., Chockler, H., & Sharygina, N. (2019). Lattice-based SMT for Program Verification. In *MEMOCODE 2019 - 17th ACM-IEEE International Conference on Formal Methods and Models for System Design: 17th ACM-IEEE International Conference on Formal Methods and Models for System Design* (pp. 1-11). [16] ACM. <https://doi.org/10.1145/3359986.3361214>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Lattice-based SMT for Program Verification

Karine Even-Mendoza
King’s College London, UK
karine.even_mendoza@kcl.ac.uk

Hana Chockler
King’s College London, UK
hana.chockler@kcl.ac.uk

Antti E. J. Hyvärinen
Università della Svizzera italiana, Switzerland
antti.hyvaerinen@usi.ch

Natasha Sharygina
Università della Svizzera italiana, Switzerland
natasha.sharygina@usi.ch

Abstract—We present a lattice-based satisfiability modulo theory for verification of programs with library functions, for which the mathematical libraries supporting these functions contain a high number of equations and inequalities. Common strategies for dealing with library functions include treating them as uninterpreted functions or using the theories under which the functions are fully defined. The full definition could in most cases lead to instances that are too large to solve efficiently.

Our lightweight theory uses lattices for efficient representation of library functions by a subset of guarded literals. These lattices are constructed from equations and inequalities of properties of the library functions. These subsets are found during the lattice traversal. We generalise the method to a number of lattices for functions whose values depend on each other in the program, and we describe a simultaneous traversal algorithm of several lattices, so that a combination of guarded literals from all lattices does not lead to contradictory values of their variables.

We evaluate our approach on benchmarks taken from the robotics community, and our experimental results demonstrate that we are able to solve a number of instances that were previously unsolvable by existing SMT solvers.

I. INTRODUCTION

The satisfiability modulo theories (SMT) [1] reasoning framework is currently one of the most successful approaches to verifying software in a scalable way. The approach is based on modeling the software and its specifications in propositional logic, while expressing domain-specific knowledge with first-order theories connected to the logic through equalities. Successful verification of software relies on finding a model that is expressive enough to capture software behavior relevant to correctness, while being sufficiently high-level to prevent reasoning from becoming prohibitively expensive.

Finding a scalable way for verifying programs or systems which use library functions as a main part of their application (e.g., implementation of robots’ movements in the Robot Operating System (ROS) [2]) is a non-trivial task: the code may contain hundreds of interacting expressions of the properties of the library functions, whose truth values depend on each other. A straightforward solution would be to use increasingly precise theories. However, this approach results in prohibitively expensive computations.

Trigonometric functions serve as a good illustration of the problem outlined above, as many domains of application,

such as robotics, planning [3], and simulations for physics and engineering [4], rely on the computation of trigonometric functions. Verification of software using trigonometric library functions [5]–[10] either requires a large amount of numerical calculations of polynomials along with irrational numbers or uses large look-up trigonometric tables, which tend to be less precise and memory consuming [11]. The former technique usually replaces the irrational expressions with rational expressions with a defined error bound [5], [6], [12]–[14] in order to bound or to evaluate trigonometric expressions to some precision. A more precise approach relies on Taylor series representation of trigonometric functions over reals; as it leads to complex computations, the resulting instances are too large to solve efficiently for all, but very small programs.

Finally, the solver implemented in HiFROG [15] supports the addition of sets of equations and inequalities as *user-defined function summaries*. We can, therefore, extract the known properties of library functions from the external libraries and encode them as user-defined SMT summaries to pass to HiFROG, and, ultimately, to the SMT solver. However, this approach is not scalable either, as we do not know beforehand which properties are going to be relevant for solving a particular instance. Hence, for library functions with a large number of equations (such as trigonometric functions—there are many equations describing properties of these functions on some subdomains), the user-defined summaries will render the instance too large to be solvable efficiently (or at all).

In this paper, we present a novel approach to reasoning about programs whose correctness depends on the values of library functions. Our approach uses the concept of *subset lattices* to construct an efficient representation of known properties of these functions. Essentially, we order the set of subsets of equations describing properties of library functions in a lattice, where each element corresponds to a set of properties that hold for some subdomain of the inputs. At every iteration of the algorithm, we verify the program with only a subset of equations that corresponds to the current element in the lattice. If this subset is insufficient for the verification (that is, does not provide enough information about the library function), we *refine* it by traversing the lattice to a higher element, containing a superset of the equations.

This lattice-based counter-example-guided abstraction refinement algorithm (LB-CEGAR) is based on the traditional counter-example-guided abstraction refinement (CEGAR) [16], [17], but replaces the refinement of the theory by the refinement of the set of equations for the library function in the program. Our approach is similar to the traditional CEGAR approach in the sense that a SAT result may indicate a real counterexample (in which case there are concrete values of symbolic variables that show the existence of this execution), or a *spurious* counterexample, where the satisfying assignment provided by the solver is due to overapproximation in the representation of the program. In contrary to the traditional CEGAR, where an UNSAT result indicated that there are no counterexamples in an abstraction of the program and hence in the concrete program as well, in LB-CEGAR the UNSAT result merely means that there are no counterexamples in the current subdomain of the input to the library function. As we describe in the paper, the lattice is constructed so that every *lattice frontier* covers the whole domain of the input variables. Hence, in case of an UNSAT result, the LB-CEGAR algorithm attempts to construct a *frontier of unsatisfiability*. Such a frontier would indicate that there are no counterexamples in the current abstraction for each subdomain of the input, and hence for the whole domain as well.

In our previous work we described a simplified LB-CEGAR algorithm for the case of one library function in the program and for small lattices [18]. In this work, we extend LB-CEGAR to the general case, where the program may contain several library functions whose values can be interconnected (for example, $\sin x$ and $\cos x$). Furthermore, each function can appear in the program multiple times, thus inducing several instances of the lattice, which are traversed simultaneously. We describe the generalised LB-CEGAR algorithm and analyse its worst case complexity and heuristics in Sec. IV.

We implemented the generalized LB-CEGAR algorithm in the bounded model checker HiFROG [15] supporting a subset of the C language and using the SMT solvers OPENSM [19] and Z3 [20] and evaluated the implementation on a large set of benchmarks containing programs whose correctness depends on the values of trigonometric functions. The experimental results clearly demonstrate an advantage to LB-CEGAR over other approaches. We outline the implementation in Sec. V and the experimental results in Sec. VI.

Our results are based on the trigonometric functions being treated as uninterpreted functions in the encoding of the problem to the SMT solver and the encoding of the mathematical equations as user-defined function summaries in the semantics of reals. We assume the correctness of these equations (such as $\sin^2 x + \cos^2 x = 1$) over real numbers. The challenge of verifying problems over IEEE floating point semantics, stemming from the implementation of the trigonometric functions in the underlying architecture, is outside of the scope of this paper. There are clear advantages to pinpointing the subset of mathematical equations that are instrumental for the correctness of the program under verification (which is what we do in this paper) to the challenge of verification over

```

1  #include <math.h>
2
3  double nonlin(double x) {
4      double x_sin = sin(x);
5      double x_cos = cos(x);
6      return x_sin*x_sin + x_cos*x_cos;
7  }
8
9  void main() {
10     double y = nondet();
11     double z = nonlin(y);
12     assert(z == 1);
13 }
14

```

Figure 1. A program with two different library functions.

floating point semantics, and we leave the exploration of this direction to the future work.

The following example illustrates the motivation for LB-CEGAR on a small program with trigonometric functions.

Example 1: The program in Fig. 1 contains two library function calls: $\sin x$ and $\cos x$. The correctness of the program follows immediately from the following trigonometric identity:

$$\forall x \in \mathbb{R}. \sin^2 x + \cos^2 x = 1. \quad (1)$$

Clearly, verifying this program with $\sin x$ and $\cos x$ treated as uninterpreted functions (that is, having nondeterministic values) would result in numerous spurious counterexamples. LB-CEGAR overcomes this problem by representing some salient properties of these functions as lattices of equations, including, in particular, Eq. (1).

In this case, the elements of the lattices for $\sin x$ and for $\cos x$ at each iteration of LB-CEGAR are not independent, as Eq. (1) should hold for each combination of these elements. Moreover, having a lattice only for one function would not suffice for proving the correctness of this program, as then we would have Eq. 1 only for one of these functions, while leaving the other one as a non-deterministic variable. This would lead to spurious counter-examples, stemming from assigning illegal values to the non-deterministic variable (for example, if $\cos x$ is left as a non-deterministic variable, it could be assigned the value 2, hence falsifying the assertion). ■

Due to the lack of space, the proofs are omitted from this version, and can be found in the full version of the paper at [21]. The implementation, the set of benchmarks, and the experimental results are available at [21]–[23].

II. PRELIMINARIES

A. SMT-based bounded model checking

Let P be a *loop-free* program represented as a transition system, and t a *safety property*, that is, a formula over the variables of P . The bounded model-checking problem amounts to determining whether all states of P , reachable within a predefined bound, satisfy t . In other words, the task of a model-checker is to find a counterexample, that is, a bounded execution of P that falsifies t , or to prove absence of such executions.

In the SMT-based bounded model-checking approach followed in this paper, the model-checker encodes all bounded

executions of P as an SMT formula, conjoins it with the negation of t , and invokes an SMT solver to check the satisfiability of the resulting formula. If the formula is deemed unsatisfiable, the program is safe, that is, P satisfies t . Otherwise, a satisfying assignment found by the SMT solver is used to build a concrete counter-example. Depending on the theory used by the SMT solver, an abstract counterexample can also be *spurious*, that is, not corresponding to any concrete execution. This situation arises when the theory is too abstract, and hence the resulting overapproximation of the behaviors of the program is too coarse. In this case, the program is re-verified with a more refined theory.

B. Function summaries

The tool HiFROG allows to incorporate function summaries into the verification process [15]. These summaries can be *interpolants* [24] from one of the previous iterations of model-checking or *user-defined* summaries supplied by the user, based on their external knowledge of the system. Some examples of user-defined summaries are available on HiFROG's webpage [25] and in the full version of this paper. We exploit this functionality by providing HiFROG with the library of *user-defined* summaries derived from external libraries for the functions, whose values are critical for determining correctness of the program, and we organise them in lattices as we explain below. This allows us to verify programs in the most abstract theory of equality logic with uninterpreted functions (EUF).

C. A subset lattice

For a given set X , the family of all subsets of X , partially ordered by the inclusion operator, forms a *subset lattice* $SL(X)$. The $\sqcap$ and $\sqcup$ operators are defined on $SL(X)$ as *intersection* and *union*, respectively. The top element $\top$ is the set X , and the bottom element $\perp$ is the empty set $\emptyset$. We note that $SL(X)$ is a De-Morgan lattice [26], as meet and join distribute over each other.

A *meet-semilattice* $\langle L, \sqcap \rangle$ of a lattice L is a partially ordered set (poset) when the $\sqcap$ operator is defined for any subset of its elements (but not necessarily the $\sqcup$ operator). A *subposet* of a lattice is a subset of elements, which follow the same partial order as in the poset. A *chain* of a lattice is a subposet of a lattice where every two elements are ordered.

In this paper, we consider $SL(X)$ and $\langle SL(X), \sqcap \rangle$, for X being a finite set of guarded expressions, as defined in Sec. III.

III. LATTICES OF GUARDED LITERALS

In this section we describe the construction of lattices of expressions for external functions.

A *guarded literal* is a Boolean expression describing some property of the function in question, together with the guard that defines a *continuous subdomain* of the inputs for which this property holds. For example, the property expressing the fact that for $0 < x < 2$, the value of $\sin x$ is positive is described by the guarded literal

$$(assume(0 < x < 2)) \wedge (\sin x > 0),$$

where $(0 < x < 2)$ is a *guard* (denoted by G) of the *literal* $(\sin x > 0)$. Literals that hold for all x (such as, for example, $(\sin x \leq 1)$) are guarded with $assume(\mathbf{true})$. A guard cannot refer to a non-continuous domain. For example,

$$(assume(0 < x < 2) \vee (7 < x < 8)) \wedge (\sin x > 0)$$

is not a legal guarded literal in our framework.

Given a set of guarded literals F for a library function f , the subset lattice $SL(F)$ consists of all subsets of these literals. However, it is easy to see that some elements in $SL(F)$ contain literals with contradictory guards. For example, a lattice of all subsets of $\sin x$ would contain, for example, the element $(assume(0 < x < 2)) \wedge (\sin x > 0)$ and the element $(assume(x = 0)) \wedge (\sin x = 0)$, which do not intersect on any subdomain of x . To reduce the size of the lattice and avoid unnecessary calls to the SMT solver, we reduce $SL(F)$ to a *meet semi-lattice* $L = \langle SL(F), \sqcap \rangle$ by removing all elements that have contradictory guards (that is, the conjunction of their guards is false).

Note that after the removal of contradictory elements, the resulting set of subsets is no longer closed under union, but it is still closed under intersection, hence the resulting set is a meet semi-lattice. Note also that the resulting meet semi-lattice can have a set of maximal elements instead of the single maximal element. For brevity, in the rest of the paper we refer to the meet semi-lattice of guarded literals for a function f simply as a lattice.

A *frontier* of a lattice L is a set of elements $X(L)$ such that each chain from $\perp$ to a maximal element in L intersects $X(L)$ in at least one element. The LB-CEGAR algorithm described in the next section relies on the fact that the union of guards of each frontier of the lattice is the whole domain of the inputs. If this is not the case, we add elements to the lattice to cover the missing subdomains. For the example of $\sin x$ above, if we have only two guarded literals

$$(assume(0 < x < 2)) \wedge (\sin x > 0)$$

and

$$(assume(x = 0)) \wedge (\sin x = 0)$$

in our set, we add the guarded literals

$$((assume(x < 0)) \wedge \mathbf{true})$$

and

$$((assume(x \geq 2)) \wedge \mathbf{true})$$

to the set to cover the whole domain of x (recall that the guards should refer to continuous subdomains, hence we need to add two guarded literals).

Claim 1: If the union of guards of a given set of guarded literals S covers the whole domain of the input, then for each frontier $X(L_S)$ of the subset lattice L_S of S , the union of guards of $X(L_S)$ also covers the whole domain of the input. And conversely, if the union of guards of a subset $X(L_S)$ of the elements of L_S covers the whole domain of the inputs, then $X(L_S)$ is a frontier of L_S .

Informally, the claim follows from the structure of the subset lattice and the fact that the bottom element of L_S covers the whole domain (the reader is referred to the full version for the formal proof of this claim).

The procedure described in this section is done at the *preprocessing stage*, once for each library function, and the resulting lattices can be used in verification of multiple programs.

IV. THE LATTICE-BASED COUNTEREXAMPLE GUIDED ABSTRACTION REFINEMENT (LB-CEGAR) ALGORITHM

In this section we present the main contribution of the paper—the LB-CEGAR algorithm. We start with an informal overview and then present the formal description of the algorithm. We proceed with discussing its worst-case complexity and then present several heuristics that reduce the complexity for the majority of the cases.

A. Overview of the LB-CEGAR algorithm

The inputs to the Lattice-based Counterexample Guided Abstraction Refinement (LB-CEGAR) algorithm (Alg. 1) are a bounded loop-free program P that includes the function f and a safety property t . The algorithm follows the standard procedure of translating P and the negation of t to a first order formula φ and invoking an SMT solver in order to find a satisfying assignment. In contrast to the standard approach, in LB-CEGAR the SMT solver has access, in addition to φ , to the external lattice L_f of guarded literals for f constructed in Sec. III. At each iteration LB-CEGAR adds the set of guarded literals in the current element E of this lattice to φ before sending φ to the SMT solver.

The refinement loop in LB-CEGAR, invoked when a satisfying assignment does not correspond to a concrete counterexample, amounts to the traversal of L_f as described below in the procedure *traverse_{SAT}*.

The algorithm terminates when it either finds a satisfying assignment that corresponds to a concrete counterexample (and hence a bug in P), reaches all maximal elements of L_f without finding concrete counterexamples for any of the satisfying assignments (that is, the current set of properties of f encoded in L_f is insufficient to verify P), or finds a *frontier* of L_f such that φ is unsatisfiable with each element of the frontier separately. The latter case implies that there are no counterexamples in the overapproximation of P for the whole domain of the inputs, and hence P satisfies t .

An iteration of LB-CEGAR with a program P , a safety property t , and a current element e consisting of the set of guarded literals $S(e)$ of the lattice L_f results in one of the following (for one library function f and a single occurrence of f in the loop-free program P):

- An SMT solver finds a satisfying assignment for φ with $S(e)$, and there is a concrete counterexample corresponding to this assignment. The algorithm terminates, outputting the counterexample as an evidence of the negative result of model-checking P .

- An SMT solver finds a satisfying assignment for φ with $S(e)$, but there is no concrete counterexample corresponding to this assignment. The algorithm invokes a *refinement step* that amounts to traversing L_f to an element e' that refines e , that is, $S(e) \subset S(e')$. If no such element exists (in other words, e is a maximal element of L_f), the algorithm terminates with inconclusive results.
- An SMT solver returns the UNSAT result for φ with $S(e)$. In other words, there is no satisfying assignment to φ in the subdomain of inputs induced by e . The refinement step of LB-CEGAR is, then, to check satisfiability of φ with elements of L_f that complement the subdomain of e to the whole domain of the input (that is, with elements of L_f that together with e form a *frontier* of L_f).
- An SMT solver returns the UNSAT result for φ with $S(e)$, and e is a part of a frontier of L_f for which φ is unsatisfiable. This result implies that there is no satisfying assignment to φ over the whole domain of the inputs, and therefore P is safe with respect to t .

If the function f appears in P several times, an instance of L_f is created for each occurrence. Furthermore, if P contains more than one library function for which we have a lattice of guarded literals, all these lattices are incorporated in LB-CEGAR. For programs with trigonometric functions, which are the primary domain of application in this paper, it is often the case that an equation includes several functions—see, for example, the program in Ex. 1.

In the next section we present a pseudo-code for LB-CEGAR and discuss the general case of several functions and several occurrences of each function in the program.

B. The main LB-CEGAR algorithm

The pseudo-code of LB-CEGAR is presented below. The input to the algorithm is a loop-free program P , a safety property t , and a set of lattices *Lattices*.

The sub-procedures and notations in Alg. 1 are defined as follows.

- The sub-procedure *checkSAT*(x) determines the satisfiability of an input formula x .
- The sub-procedure *checkRealCE*(P, t, CE) returns **true** if CE can be concretised to a counterexample, demonstrating a behaviour of P that falsifies t .
- The set *Lattices* consists of all occurrences of lattices for all library functions in P .
- We denote by L_f^i a lattice for the i -th occurrence of f in P , and by e the current element in the lattice traversal. For an element e , we define *literals* as the *conjunction of guarded literals* of e .
- The sub-procedure *traverse_{UNSAT}*(*Lattices*) performs the traversal of the lattice from the current element e to the next element e' if the result of model-checking $\varphi \wedge \text{literals}$ is **UNSAT**. The next element e' in the same lattice as e is e 's 'sibling', that is, an element, whose set of literals corresponds to a different subdomain of the input. If there is already a frontier of elements

Algorithm 1: LB-CEGAR

Input : Program P , safety property t , and set $Lattices$
Output: $\langle \text{Safe} \rangle$, $\langle \text{Unsafe}, CE \rangle$, or $\langle \text{Unknown}, \perp \rangle$

```
1  $\varphi \leftarrow P \wedge \neg t$ 
2  $Query \leftarrow \varphi$ 
3  $\langle result, CE \rangle \leftarrow checkSAT(Query)$ 
4 if  $result$  is UNSAT  $\vee checkRealCE(\varphi, CE)$  then
5   | go to Exit // No lattice-based refinement needed
6 end
7  $\chi \leftarrow \text{true}$ 
8 repeat
9   |  $\chi' \leftarrow \chi$  // Formula from the previous iteration
10  | if  $result$  is UNSAT then
11    |  $traverse_{UNSAT}(Lattices)$ 
12  | end
13  | if  $result$  is SAT then
14    |  $traverse2refine_{SAT}(Lattices)$ 
15  | end
16  |  $\chi \leftarrow literals(\varphi, Lattices)$ 
17  | // Solve again if there are new literals
18  | if  $\chi \neq \chi'$  then
19    |  $Query \leftarrow \varphi \wedge \chi$ 
20    |  $\langle result, CE \rangle \leftarrow checkSAT(Query)$ 
21  | end
22 until  $(\chi == \chi') \vee$ 
     $\vee checkRealCE(Query, CE) \vee termination(result, Lattices)$ ;
23 End-LB-CEGAR:
24 if  $result$  is UNSAT then
25   | return  $\langle \text{Safe} \rangle$  // Safe
26 end
27 if  $checkRealCE(P, t, CE)$  then
28   | return  $\langle \text{Unsafe}, CE \rangle$  // Real counterexample
29 end
30 return  $\langle \text{Unknown} \rangle$  // Inconclusive results, further refinement
    needed
```

in each lattice such that model-checking $\varphi \wedge literals$ returns **UNSAT** for each element of these frontiers, the procedure $traverse_{UNSAT}(Lattices)$ does not change the current element e .

- The sub-procedure $traverse_{SAT}(Lattices)$ is invoked when there is a satisfying assignment for $\varphi \wedge literals$, but the counterexample induced by it is *spurious*, that is, it does not correspond to a behaviour of P falsifying t . The procedure traverses the lattice to an element e' that refines e , that is, $S(e) \subset S(e')$. If e is a maximal element, the procedure $traverse_{SAT}(Lattices)$ does not change the current element e .
- The sub-procedure $termination(result, Lattices)$ checks whether one of the termination conditions holds: either the current satisfying assignment induces a concrete counterexamples, or there is an **UNSAT** frontier for each lattice $L_f^i \in Lattices$, or there is a satisfying assignment for each maximal element in each lattice in $Lattices$ that does not induce a concrete counterexample.

Finally, we address the complexity resulting from having several functions in P , whose lattices refer to each other. This is illustrated by Ex. 1, where the correctness of the program depends on the guarded literal

$$(assume(\text{true})) \wedge (\sin^2 x + \cos^2 x) = 1.$$

In fact, this is quite common in programs with trigonometric functions, as trigonometric identities often refer to several

functions in the same identity. The algorithm identifies library functions used in the set $Lattices$ and assigns the same variable to all occurrences of the same function, hence connecting between the lattices of different functions.

C. Correctness and complexity of LB-CEGAR

It is easy to see that LB-CEGAR terminates. Indeed, the lattice traversal visits every combination of elements of lattices in $Lattices$ at most once, and for each combination of elements it invokes the model-checking procedure of a bounded loop-free program P with respect to t , which terminates. The number of possible combinations of elements in the lattices is exponential in the number of lattices, hence leading to the complexity result below.

Theorem 1: The worst-case running-time complexity of LB-CEGAR is $O(|L|^n \times MC(P, t))$, where $|L|$ is the bound on the size of each lattice in the set $Lattices$, n is the number of lattices in $Lattices$, and $MC(P, t)$ is the running-time complexity of model-checking P with respect to t using the guarded literals.

Moreover, the following theorem states that LB-CEGAR produces a correct result.

Theorem 2: The following holds for any bounded loop-free program P and a safety property t , assuming correctness of the guarded literals in $Lattices$:

- If LB-CEGAR outputs **Safe**, the program P is correct with respect to t .
- If LB-CEGAR outputs **Unsafe** with an accompanying CE , the CE demonstrates an execution of P that falsifies t .
- If LB-CEGAR outputs **Unknown**, the current theory and the set of guarded literals are insufficient to produce a conclusive result.

We observe that, while the worst-case complexity of LB-CEGAR is exponential in the number of lattices, in practice the algorithm is very efficient, as we show in Sec. VI. This is partly due to the *incrementality* of the calls to the SMT solver, as the formula φ representing $P \wedge \neg t$ stays the same for all iterations, and the next element e' differs from the current element e of the lattice only slightly. Another reason for the significantly lower complexity in practice is that our implementation of LB-CEGAR includes several heuristics, which we describe in the next section. The heuristics do not alter the correctness of the algorithm.

V. IMPLEMENTATION

The algorithms were implemented on top of the SMT-based function summarisation bounded model checker HiFROG [15] with OPENSMT [19] and Z3 [7], [20] solvers. The details of our initial implementation are described in [18]. Here we describe the extension of the implementation to support the full LB-CEGAR algorithm.

Fig. 2 presents a high-level view of the implementation of LB-CEGAR in HiFROG and a comparison between the implementation as a flat (non-hierarchical) set of user-defined summaries, our prototype implementation with one occurrence

tices simultaneously. The motivation for the ‘coordinated’ traversal is that all loop iterations, except, perhaps, for the last one, are similar, and hence there is a high probability that the same set of guarded literals would fit all these occurrences.

VI. EVALUATION

For the evaluation of LB-CEGAR, we constructed two lattices for $\sin$ and $\cos$ functions with 40 and 38 guarded literals, respectively. The validation test for these expressions contains a set of 144 benchmarks in C with a total of 365 assert statements. The scripts for the lattice construction, the benchmarks for the validation test, and the results of the validation test are available at [21], [22].

The set of benchmarks contained a mix of our crafted benchmarks, programs from the software verification competition SVCOMP [31], and HiFROG benchmarks [15], with a total of 141 C programs with at least one library function call, containing in total 194 calls for $\sin$ and 179 calls for $\cos$, with 279 claims (127 SAT and 152 UNSAT). In 42 benchmarks, the library function is called at least 4 times, and in 8 benchmarks, the library function call is in a loop. The crafted benchmarks either assert known properties of trigonometric functions or contain a small part of code that is typical to kinematic problems, mainly examining the ability of verifying code with multiplication between two library function calls; e.g., $\cos \phi \times \sin \theta$.

We used EUF with a semi-incremental solving mode in OPENSMT [19] and linear arithmetics with EUF with an incremental solving mode in Z3 solver [20].

The experiments were performed on a virtual machine (VM) with Ubuntu 16.04 Linux system, single core, 8GB RAM; the VM runs on a machine with 4-Intel i7-6600U CPUs clocked at 2.60GHz. The experimental results, the benchmarks, and the source code, are available at [21], [23].

A. Evaluation of LB-CEGAR with Real arithmetics

Figure 3 presents the comparison of LB-CEGAR with CBMC version 5.10 [32], HiFROG [15], and our previous implementation [18] supporting one library function at a time.

The total number of **solved instances** is the **blue** bar and the **orange** bar, for **Safe** and **Unsafe** instances respectively. The total number (as a negative number) of **unsolved instances** is the **gray** bar and the **yellow** bar, for **SAT** instances that are classified as **Unknown** (or **SAT** without a counterexample), and for the instances that timed out (TO) or were out-of-memory (OM), with the timeout set to 4000s and out-of-memory set to 3GB, respectively.

The four different colours of the bars are consistent across all six charts. Each chart represents the total solved instances for a particular tool or a variant of a tool. The tools at the clockwise order are, LB-CEGAR (top-left), CBMC [33] (top-right), HiFROG [15] LRA (middle-right), EUF (bottom-right) and with user defined summaries (bottom-left), and HiFROG [18] with a single lattice (middle-left). All approaches with summaries used EUF with LRA (UFLRA).

Figure 3. Comparison of the Number (#) of Solved Claims with Different Approaches and a Set of Trigonometric Benchmarks in C.

Verification with function summaries in HiFROG avoids processing the single static assignment (SSA) expression of the original function and uses only the SMT summary. Hence, complicated benchmarks, which contain library function calls in a loop or in a non-linear expression, were more likely to be successfully verified with summary-based approaches.

HiFROG with user-defined summaries used ~ 80 equations of trigonometric properties, loaded at once as unstructured data and solved a total of 227 instances (HiFrog - User-Defined, bottom-left, Fig. 3). HiFROG with a single lattice used ~ 40 equations of trigonometric properties, and solved a total of 185 instances (Single-Lattice - LRA, middle-left, Fig. 3).

LB-CEGAR with two lattices, constructed from a set of ~ 80 equations of trigonometric properties (Mix-lattice graph, top-left, Fig. 3), outperformed both variants of HiFROG and had the highest overall number of *solved instances*: over **260** instances.

Connecting lattices of different functions in LB-CEGAR algorithm allowed our approach to verify the highest number

of **Safe** instances, on many on which other tools failed (including HiFrog-User-Defined and Single-Lattice-LRA). In comparison with LB-CEGAR (with 261 solved instances, out of 279), other tools that do not use summaries or insights on trigonometric functions, managed to solve at most 136 instances (CBMC V. 5. 10, HiFrog-LRA, and HiFrog-EUF, in Fig. 3, at top-right, bottom-middle, and bottom-right, respectively), mainly because of the use of non-deterministic variables to represent trigonometric functions.

B. Evaluation of LB-CEGAR with different parameters of HiFROG

Table I presents a full comparison of our implementation of the LB-CEGAR algorithm with other tools. The comparison is performed using various parameters of HiFROG, even-though LRA with EUF is the most suitable theory-combination for trigonometric functions, based on our experience.

The physical files of SMT summaries for LRA and EUF were the same; HiFROG read an SMT summary file differently according to the theory in use. The SMT summaries for linear integer arithmetics were different to prevent any conversions to real arithmetics (e.g., *to_real* token) in the SMT query.

In Table I, the verification results of LB-CEGAR appear in white and are compared to CBMC [33] and HiFROG [15], [18] (grey and dark grey), with the best results underlined and marked in light-yellow.

The evaluation of HiFROG in greyscale is with a single lattice, and with and without user defined summaries, using: EUF, LRA with EUF, LIA with EUF, each of which is a different column in Table I (with the theory being EUF, LRA or LIA). The symbol # stands for the number of instances solved (first two lines) or unsolved (last two lines). Unsolved instances are false negative results (FN-SAT), inconclusive (also marked as FN-SAT), or timeout or out-of-memory (*TO + OM*).

The description of each column in Table I is as follows.

- The **white** columns in Table I (3 columns, **Lattice Mix col.**) contain the results the LB-CEGAR algorithm along with the modifications presented in Sec. IV with different sets of parameters (with the theory being EUF, LRA ,or LIA), against other tools in the gray scale columns.
- The **gray scale** columns in Table I (10 columns from the end) contain the results of other tools: the lattice-based refinement approach with a single lattice [18] in **Lattice Single col.**, HiFROG with a large set of user defined summaries [15] in **HiFrog UDS col.**, and HiFROG [15] and CBMC version 5.10 [32] in the **most right columns**.

The variant of HiFROG with theory refinement is omitted from the comparison, as its current implementation does not support trigonometric functions.

For the lattice-based refinement approach for verification of programs with multiple trigonometric functions, the best setting was with LRA with EUF (Lattice Mix, LRA col.), which had the highest overall number of *solved instances* over

260 instances, and performed almost as well as HiFROG without any summary (ran out of resources 4 times vs. 2 times HiFROG did). The other two configurations of parameters that we tried with the lattice-based refinement approach for programs with multiple library functions, were EUF (Lattice Mix, EUF col.) and LIA with EUF (Lattice Mix, LIA col.). While EUF performed poorly in general, LIA with EUF has shown limited potential in solving instances that required real arithmetics, which indicating the possibility of applying this method for code with library functions over significantly different input domains.

The comparison with a single lattice [18] (Lattice Single, LRA col.) used either a meet-semilattice for sin function or for cos function per benchmark, which led to a poor performance when both lattices were required; however, perhaps unsurprisingly, this did not result a poor performance when a single lattice was sufficient to prove safety of a claim (59 instance, Lattice Single, LRA col.). HiFROG with user-defined summaries (HiFrog UDS, LRA col.) could not solve around 50 safe instances that required a wider context regarding other library functions, for one or more expressions with a library function call.

VII. RELATED WORK

The problem of verification of programs with transcendental functions and, in particular, trigonometric functions is addressed by several verification tools, such as iSAT3 [34] via interval propagation, dReal [35] by using δ -satisfiability, where δ is associated with the numerical error, Coq interval [36] and Gappa [37] via interval propagation with Taylor series, and MathSAT5 [6], [12], [38] by using Taylor series with a partial set of trigonometric properties and for hardware verification [39]. In contrast to these approaches, our algorithm does not require a nonlinear arithmetic or a calculation of Taylor series, which is computationally expensive for large programs.

Computationally inexpensive theories can be used to over-approximate complex problems. This approach has been used in solving equations on non-linear real arithmetic and transcendental functions based on linear real arithmetic and equality logic with uninterpreted functions [6], [12], [40], [41], as well as on scaling up bit-vector solving [15], [42], [43]. Our work can be seen as a generalisation of these approaches as we support inclusion of lemmas from more descriptive logics to increase the expressiveness of computationally lighter logics.

Lattices are a useful mathematical structure in understanding the relationships between different abstractions and have been widely applied in program solving with Craig interpolation. [44] presents a semantic solver-independent framework for systematically exploring interpolant lattices using the notion of interpolation abstraction. A lattice-based system for interpolation in propositional structures is presented in [45], extended to consider size optimisation techniques in the context of function summaries in [46], [47], and further extended to partial variable assignments in [48]. Similar lattice-based reasoning has also been extended to interpolation in first

Table I
COMPARISON OF LB-CEGAR IN HiFrog WITH DIFFERENT PARAMETERS WITH CBMC AND HiFrog. THE COMPARISON IS ON THE NUMBER OF SOLVED AND UNSOLVED INSTANCES, WITH THE TIMEOUT (TO) SET TO 4,000s AND OUT-OF-MEMORY (OM) SET TO 3GB.

		Lattice Mix			Lattice Single			HiFrog UDS			HiFrog			CBMC
		LRA	LIA	EUF	LRA	LIA	EUF	LRA	LIA	EUF	LRA	LIA	EUF	
#Solved	UNSAT	134	42	8	59	13	6	100	24	9	2	3	1	9
	SAT	127	126	126	126	125	126	127	126	126	127	127	127	127
#Failed	FN	14	104	136	87	133	139	48	122	136	148	147	149	136
	TO+OM	4	7	9	7	8	8	4	7	8	2	2	2	7

order logic with other SMT theories [49], [50]. The approach presented in this work differs from the above in that we do not rely on interpolation and work in tight integration with the model-checker.

Lattices and posets are used in abstract interpretation [51] to model a sound approximation of the semantics of code, where completeness and partial completeness [52]–[55] refer to the no loss of precision during the approximation of the semantics of code. Giacobazzi et al. [54], [55] present the notation of backward and forward completeness and show the connection between iteratively computing the backward (forward)-complete shell to the general CEGAR framework [17]. The completeness of their algorithm depends on the properties of the abstraction, while our algorithm has no such requirements.

VIII. CONCLUSIONS AND FUTURE WORK

We presented a new algorithm LB-CEGAR that is used for verification of programs with library functions, for which

a number of equations, some of which are instrumental for verification of these programs, exist in external sources (the mathematical library, Coq proof assistant, etc.). The main idea of the algorithm is to organize the equations in subset lattices, and to replace the traditional CEGAR refinement loop with lattice traversal. The algorithm is general in the sense that it allows several occurrence of the same library function and/or several different library functions, some of which depend on each other, in the same program. While the theoretical worst-case complexity of LB-CEGAR is high due to an exponential number of combinations of elements of different lattices, our experimental results show that the algorithm is very efficient in practice and outperforms state-of-the-art model-checking tools on benchmarks with trigonometric functions.

We view the programs with trigonometric functions as the primary domain of application of LB-CEGAR. In the future, we plan to explore the domain of verification of programs describing robots' movements and kinematics in general.

REFERENCES

- [1] D. Detlefs, G. Nelson, and J. B. Saxe, “Simplify: a theorem prover for program checking,” *J. ACM*, vol. 52, no. 3, pp. 365–473, 2005.
- [2] “ROS homepage,” <http://www.ros.org/>.
- [3] J. Wittenburg, *Kinematics: Theory and Applications*. Springer, 2016.
- [4] “Open Source Physics page,” <http://www.compadre.org/osp/>.
- [5] B. Akbarpour and L. C. Paulson, “Metitarski: An automatic theorem prover for real-valued special functions,” *Journal of Automated Reasoning*, vol. 44, no. 3, pp. 175–205, 2010.
- [6] A. Cimatti, A. Griggio, A. Irfan, M. Roveri, and R. Sebastiani, “Satisfiability modulo transcendental functions via incremental linearization,” in *CADE*, ser. LNCS, vol. 10395. Springer, 2017, pp. 95–113.
- [7] L. De Moura and G. O. Passmore, “Computation in real closed infinitesimal and transcendental extensions of the rationals,” in *CADE*, ser. LNCS, vol. 7898. Springer, 2013, pp. 178–192.
- [8] W. Denman and C. Muñoz, “Automated real proving in pvs via metitarski,” in *FM*, ser. LNCS, vol. 8442. Springer, 2014, pp. 194–199.
- [9] N. Ge, E. Jenn, N. Breton, and Y. Fonteneau, “Integrated formal verification of safety-critical software,” *International Journal on Software Tools for Technology Transfer*, vol. 20, no. 4, pp. 423–440, 2018.
- [10] P. Trojanek and K. Eder, “Verification and testing of mobile robot navigation algorithms: A case study in spark,” in *IROS*. IEEE, 2014, pp. 1489–1494.
- [11] R. Kirner, M. Grössing, and P. P. Puschner, “Comparing WCET and resource demands of trigonometric functions implemented as iterative calculations vs. table-lookup,” in *WCET*, ser. OASICS. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006.
- [12] A. Cimatti, A. Griggio, A. Irfan, M. Roveri, and R. Sebastiani, “Incremental linearization for satisfiability and verification modulo nonlinear arithmetic and transcendental functions,” *ACM Trans. Comput. Logic*, vol. 19, no. 3, pp. 19:1–19:52, 2018.
- [13] M. Dumas, D. Lester, and C. Muñoz, “Verified real number calculations: A library for interval arithmetic,” *IEEE Trans. Comput.*, vol. 58, no. 2, pp. 226–237, 2009.
- [14] G. Melquiond and C. Munoz, “Guaranteed proofs using interval arithmetic,” in *ARITH*. IEEE, 2005, pp. 188–195.
- [15] L. Alt, S. Asadi, H. Chockler, K. Even Mendoza, G. Fedyukovich, A. E. J. Hyvärinen, and N. Sharygina, “HiFrog: SMT-based function summarization for software verification,” in *TACAS*, ser. LNCS, vol. 10206. Springer, 2017, pp. 207–213.
- [16] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith, “Counterexample-guided abstraction refinement,” in *CAV*, ser. LNCS, vol. 1855. Springer, 2000, pp. 154–169.
- [17] —, “Counterexample-guided abstraction refinement for symbolic model checking,” *J. ACM*, vol. 50, no. 5, pp. 752–794, 2003.
- [18] K. Even-Mendoza, S. Asadi, A. E. J. Hyvärinen, H. Chockler, and N. Sharygina, “Lattice-based refinement in bounded model checking,” in *VSTTE*, ser. LNCS, vol. 11294. Springer, 2018, pp. 50–68.
- [19] A. E. J. Hyvärinen, M. Marescotti, L. Alt, and N. Sharygina, “OpenSMT2: An SMT solver for multi-core and cloud computing,” in *SAT*, ser. LNCS, vol. 9710. Springer, 2016, pp. 547–553.
- [20] L. De Moura and N. Bjørner, “Z3: An efficient smt solver,” in *TACAS*, ser. LNCS, vol. 4963. Springer, 2008, pp. 337–340.
- [21] http://verify.inf.usi.ch/content/trig_refinement, 2019.
- [22] “Git repository of evaluation of lattice-based refinement approach,” <https://github.com/karineek/latticeref>, 2019.
- [23] “Git repository of HiFrog,” <https://scm.ti-edu.ch/projects/hifrog/>.
- [24] W. Craig, “Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory,” in *J. of Symbolic Logic*, 1957, pp. 269–285.
- [25] “HiFrog - tool usage page,” <http://verify.inf.usi.ch/hifrog/tool-usage>.
- [26] G. Birkhoff, *Lattice Theory*, 3rd ed. AMS, 1967.
- [27] I. Anderson, *Combinatorics of Finite Sets*. Clarendon Press, Oxford, 1987.
- [28] “The Coq proof assistant,” <https://coq.inria.fr/>.
- [29] “List of trigonometric identities, from Wikipedia, the free encyclopedia,” https://en.wikipedia.org/wiki/List_of_trigonometric_identities.
- [30] “Trigonometric tables, from Wikipedia, the free encyclopedia,” https://en.wikipedia.org/wiki/Trigonometric_tables.
- [31] “Competition on software verification (SV-COMP),” <https://sv-comp.sosy-lab.org/2018/>, 2018.
- [32] <http://www.cprover.org/cbmc/>.
- [33] E. Clarke, D. Kroening, and F. Lerda, “A tool for checking ANSI-C programs,” in *TACAS*, ser. LNCS, vol. 2988. Springer, 2004, pp. 168–176.
- [34] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert, “Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure,” *Journal on Satisfiability, Boolean Modeling and Computation*, vol. 1, pp. 209–236, 2007.
- [35] S. Gao, S. Kong, and E. M. Clarke, “dreal: An smt solver for nonlinear theories over the reals,” in *CADE*, ser. LNAI, vol. 7898. Springer, 2013, pp. 208–214.
- [36] G. Melquiond, “Floating-point arithmetic in the coq system,” *Inf. Comput.*, vol. 216, pp. 14 – 23, 2012, Special Issue: 8th Conference on Real Numbers and Computers.
- [37] F. de Dinechin, C. Lauter, and G. Melquiond, “Certifying the floating-point implementation of an elementary function using gappa,” *IEEE Trans. Comput.*, vol. 60, no. 2, pp. 242–253, 2011.
- [38] A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani, “The mathsat5 smt solver,” in *TACAS*, ser. LNCS, vol. 7795. Springer, 2013, pp. 93–107.
- [39] J. Harrison, “Formal verification of floating point trigonometric functions,” in *FMCAD*, ser. LNCS, vol. 1954. Springer, 2000, pp. 217–233.
- [40] A. Cimatti, A. Griggio, A. Irfan, M. Roveri, and R. Sebastiani, “Invariant checking of NRA transition systems via incremental reduction to LRA with EUF,” in *TACAS*, ser. LNCS, vol. 10205. Springer, 2017, pp. 58–75.
- [41] T. Kutsuna, Y. Ishii, and A. Yamamoto, “Abstraction and refinement of mathematical functions toward smt-based test-case generation,” *International Journal on Software Tools for Technology Transfer*, vol. 18, no. 1, pp. 109–120, 2016.
- [42] Y.-S. Ho, P. Chauhan, P. Roy, A. Mishchenko, and R. Brayton, “Efficient uninterpreted function abstraction and refinement for word-level model checking,” in *FMCAD*. ACM, 2016, pp. 65–72.
- [43] A. E. J. Hyvärinen, S. Asadi, K. Even-Mendoza, G. Fedyukovich, H. Chockler, and N. Sharygina, “Theory refinement for program verification,” in *SAT*, ser. LNCS, vol. 10491. Springer, 2017, pp. 347–363.
- [44] P. Rummer and P. Subotic, “Exploring interpolants,” in *FMCAD*. IEEE, 2013, pp. 69–76.
- [45] V. D’Silva, M. Purandare, G. Weissenbacher, and D. Kroening, “Interpolant strength,” in *VMCAI*, ser. LNCS, vol. 5944. Springer, 2010, pp. 129–145.
- [46] L. Alt, G. Fedyukovich, A. E. J. Hyvärinen, and N. Sharygina, “A Proof-Sensitive Approach for Small Propositional Interpolants,” in *VSTTE*, ser. LNCS, vol. 9593. Springer, 2015, pp. 1–18.
- [47] S. F. Rollini, L. Alt, G. Fedyukovich, A. E. J. Hyvärinen, and N. Sharygina, “PeRIPLO: A framework for producing effective interpolants in SAT-based software verification,” in *LPAR*, ser. LNCS, vol. 8312. Springer, 2013, pp. 683–693.
- [48] P. Jancík, L. Alt, G. Fedyukovich, A. E. J. Hyvärinen, J. Kofron, and N. Sharygina, “PVAIR: Partial Variable Assignment InterpolatoR,” in *FASE*, ser. LNCS, vol. 9633. Springer, 2016, pp. 419–434.
- [49] L. Alt, A. E. J. Hyvärinen, and N. Sharygina, “LRA interpolants from no man’s land,” in *HVC*, ser. LNCS, vol. 10629. Springer, 2017, pp. 195–210.
- [50] L. Alt, A. E. J. Hyvärinen, S. Asadi, and N. Sharygina, “Duality-based interpolation for quantifier-free equalities and uninterpreted functions,” in *FMCAD*. IEEE, 2017, pp. 39–46.
- [51] P. Cousot and R. Cousot, “Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints,” in *POPL*. ACM, 1977, pp. 238–252.
- [52] P. Cousot, “Partial completeness of abstract fixpoint checking,” in *SARA*, ser. LNAI, vol. 1864. Springer, 2000, pp. 1–25.
- [53] P. Cousot and R. Cousot, “Systematic design of program analysis frameworks,” in *POPL*. ACM, 1979, pp. 269–282.
- [54] R. Giacobazzi and E. Quintarelli, “Incompleteness, counterexamples, and refinements in abstract model-checking,” in *SAS*, ser. LNCS, vol. 2126. Springer, 2001, pp. 356–373.
- [55] R. Giacobazzi, F. Ranzato, and F. Scozzari, “Making abstract interpretations complete,” *J. ACM*, vol. 47, no. 2, pp. 361–416, 2000.