

Invariant Inference With Provable Complexity From the Monotone Theory

Yotam M. Y. Feldman and Sharon Shoham

Tel Aviv University

Abstract. Invariant inference algorithms such as interpolation-based inference and IC3/PDR show that it is feasible, in practice, to find inductive invariants for many interesting systems, but non-trivial upper bounds on the computational complexity of such algorithms are scarce, and limited to simple syntactic forms of invariants. In this paper we achieve invariant inference algorithms, in the domain of propositional transition systems, with *provable upper bounds* on the number of SAT calls. We do this by building on the *monotone theory*, developed by Bshouty for exact learning Boolean formulas. We prove results for two invariant inference frameworks: (i) *model-based interpolation*, where we show an algorithm that, under certain conditions about reachability, efficiently infers invariants when they have both short CNF and DNF representations (transcending previous results about monotone invariants); and (ii) *abstract interpretation* in a domain based on the monotone theory that was previously studied in relation to *property-directed reachability*, where we propose an efficient implementation of the best abstract transformer, leading to overall complexity bounds on the number of SAT calls. These results build on a novel procedure for computing least monotone overapproximations.

1 Introduction

In a fruitful, recent trend, many that aspire to innovate in verification seek advice from research in machine learning [e.g. 38, 37, 36, 21, 18, 11, 22, 12, 24, 13]. The topic of this paper is the application of the *monotone theory*, developed by Bshouty in exact concept learning, to improve theoretical complexity results for *inductive invariant inference*.

One of the *modi operandi* for automatically proving that a system is safe—that it cannot reach a state it should not—is an *inductive invariant*, which is an assertion that (i) holds for the initial states, (ii) does not hold in any bad state, and (iii) is closed under transitions of the system. These properties are reminiscent of a data classifier, separating good from bad points, prompting the adaptation of algorithms from classical classification to invariant inference [e.g. 38, 37, 36, 35, 18, 11, 19, 13]. In this paper we focus on inductive invariants for propositional transition systems, which are customary in hardware verification and also applicable to software systems through predicate abstraction [20, 17].

The *monotone theory* by Bshouty [5] is a celebrated achievement in learning theory (most notably in exact learning with queries [1]) that is the foundation for learning Boolean formulas with complex syntactic structures. At its core,

the monotone theory studies the *monotonization* $\mathcal{M}_b(\varphi)$ of a formula φ w.r.t. a valuation b , which is the smallest b -monotone formula that overapproximates φ . (This concept is explained in §3.) In Bshouty’s work, several monotonizations are used to efficiently reconstruct φ .

Recently, the monotone theory has been applied to theoretical studies of invariant inference in the context of two prominent SAT-based inference approaches. In this paper, we solve open problems in each, using a new efficient algorithm to compute monotonizations.

Efficient interpolation-based inference. The study [13] of interpolation-based invariant inference—a hugely influential approach pioneered by McMillan [26]—identified the fence condition as a property of systems and invariants under which the success of a model-based inference algorithm [6, 3] is guaranteed. (We explain the fence condition in §5.1.) Under this condition, the number of SAT calls (specifically, bounded model checking queries) of the original model-based interpolation algorithm was shown to be polynomial in the DNF size (the number of terms in the smallest DNF representation) of the invariant, but only when the invariant is *monotone* (containing no negated variables) [13]. Based on the monotone theory, the authors of [13] further introduced an algorithm that, under the same fence condition, could efficiently infer invariants that were *almost* monotone (containing $O(1)$ terms with negated variables). However, their techniques could not extend to mimic the pinnacle result of Bshouty’s paper: the CDNF algorithm [5], which can learn formulas in a number of queries that is polynomial in their DNF size, their CNF size (the number of clauses in their smallest CNF representation), and the number of variables. It was unclear whether an analogous result is possible in invariant inference without strengthening the fence condition, e.g. to assume that the fence condition holds both forwards and backwards (see §7).

We solve this question, and introduce an algorithm that can infer an invariant in a number of SAT queries (specifically, bounded model checking queries) that is polynomial in the invariant’s DNF size, CNF size, and the number of variables, under the assumption that the invariant satisfies the fence condition, without further restrictions (Thm. 5). In particular, this implies that invariants that are representable by a small decision tree can be inferred efficiently. The basic idea is to learn an invariant I as a conjunction of monotonizations $\mathcal{M}_\sigma(I)$ where σ are chosen as counterexamples to induction. The challenge is that I is unknown, and the relatively weak assumption on the transition system of the fence condition does not allow the use of several operations (e.g. membership queries) that learning algorithms rely on to efficiently generate such a representation.

Efficient abstract interpretation. The study of IC3/PDR [4, 9] revealed that part of the overapproximation this sophisticated algorithm performs is captured by an abstract interpretation procedure, in an abstract domain founded on the monotone theory [14]. In this procedure, dubbed \mathcal{A} -PDR, each iteration involves several monotonizations of the set of states reachable in one step from the value of the previous iteration. Upper bounds on the number of iterations in \mathcal{A} -PDR were investigated to shed light on the number of frames of PDR [14]. However, it

was unclear whether the abstract domain itself can be implemented in an efficient manner, and whether efficient complexity bounds on the number of SAT queries (and not just the number of iterations) in Λ -PDR can be obtained.

We solve this question, and show that Λ -PDR can be implemented to yield an overall upper bound on the number of SAT calls which is polynomial in the same quantity that was previously used to bound the number of iterations in Λ -PDR (Thm. 7). This is surprising because, until now, there was no way to compute monotonicizations of the post-image of the previous iteration that did not suffer from the fact that the exact post-image of a set of states may be much more complex to represent than its abstraction.

Super-efficient monotonicization. Bshouty [5] provided an algorithm to compute the monotonicization $\mathcal{M}_b(\varphi)$, but the complexity of this algorithm depends on the DNF size of the original formula φ . Our aforementioned results build on a new algorithm for the same task, whose complexity depends on the DNF size of the monotonicization $\mathcal{M}_b(\varphi)$ (Thm. 2), which may be much smaller (and never larger). This enables our efficient interpolation-based inference algorithm and our efficient implementation of abstract interpretation, although each result requires additional technical sophistication: For our efficient model-based interpolation result, the key idea is that the monotonicization of an invariant satisfying the fence condition can be computed through the monotonicization of the set of states reachable in at most s steps, and our new monotonicization algorithm allows to do this efficiently even when the latter set is complex to represent exactly. For our efficient abstract interpretation result, the key idea is that the DNF size of an abstract iterate is bounded by a quantity related to monotonicizations of the transition relation, and our new monotonicization algorithm allows to compute it efficiently w.r.t the same quantity even though the DNF size of the exact post-image of the previous iterate may be larger.

Overall, we make the following contributions:

- We introduce a new efficient algorithm to compute monotonicizations, whose complexity in terms of the number of SAT queries is proportional to the DNF size of its output (§4).
- We prove that an invariant that satisfies the fence condition can be inferred in a number of SAT (bounded model checking) queries that is polynomial in its CNF size, its DNF size, and the number of variables; in particular, invariants represented by small decision trees can be efficiently inferred (§5).
- We prove an efficient complexity upper bound for the number of SAT queries performed by abstract interpretation in a domain based on the monotone theory (§6).

§2 sets preliminary notation and §3 provides background on the monotone theory. §7 discusses related work and §8 concludes.

2 Preliminaries

We work with propositional transition systems defined over a vocabulary $\Sigma = \{p_1, \dots, p_n\}$ of n Boolean variables. We identify a formula with the set of its

valuations, and at times also identify a set of valuations with an arbitrary formula that represents it which is chosen arbitrarily (one always exists in propositional logic). $\varphi \implies \psi$ denotes the validity of the formula $\varphi \rightarrow \psi$.

States, transition systems, inductive invariants. A *state* is a *valuation* to Σ . If x is a state, $x[p]$ is the value (*true/false* or $1/0$) that x assigns to the variable $p \in \Sigma$. A *transition system* is a triple $(Init, \delta, Bad)$ where $Init, Bad$ are formulas over Σ denoting the set of initial and bad states respectively, and the *transition relation* δ is a formula over $\Sigma \uplus \Sigma'$, where $\Sigma' = \{x' \mid x \in \Sigma\}$ is a copy of the vocabulary used to describe the post-state of a transition. If $\tilde{\Sigma}, \tilde{\Sigma}'$ are distinct copies of Σ, Σ' , $\delta[\tilde{\Sigma}, \tilde{\Sigma}']$ denotes the substitution in δ of each $p \in \Sigma$ by its corresponding in $\tilde{\Sigma}$ and likewise for $\Sigma', \tilde{\Sigma}'$. Given a set of states S , the *post-image* of S is $\delta(S) = \{\sigma' \mid \exists \sigma \in S. (\sigma, \sigma') \models \delta\}$. A transition system is *safe* if all the states that are reachable from $Init$ via any number of steps of δ satisfy $\neg Bad$. An *inductive invariant* is a formula I over Σ such that (i) $Init \implies I$, (ii) $I \wedge \delta \implies I'$, and (iii) $I \implies \neg Bad$, where I' denotes the result of substituting each $x \in \Sigma$ for $x' \in \Sigma'$ in I . In the context of propositional logic, a transition system is safe iff it has an inductive invariant.

Use of SAT in invariant inference. Given a candidate, the requirements for being an inductive invariant can be verified using SAT; we refer to the SAT query that checks requirement ii by the name *inductiveness check*. When an inductiveness check fails, a SAT solver returns a *counterexample to induction*, which is a transition (σ, σ') with $\sigma \models I$ but $\sigma' \not\models I$. Another important check in invariant inference algorithms that can be implemented using SAT is *bounded model checking (BMC)* [2], which asks whether a set of states described by a formula ψ is forwards unreachable in a bounded number $s \in \mathbb{N}$ of steps; we write this as the check $\underline{\delta}^s(Init) \cap \psi \stackrel{?}{=} \emptyset$. Using SAT it is also possible to obtain a counterexample $\sigma \in \underline{\delta}^s(Init) \cap \psi$ if it exists.

We measure the *complexity* of a SAT-based inference algorithm by the number of SAT calls it performs (including inductiveness checks, BMC, and other SAT calls), and the number of other steps, when each SAT call is considered one step (an oracle call).

Literals, Cubes, Clauses, CNF, DNF. A *literal* ℓ is a variable p or its negation $\neg p$. A *clause* c is a disjunction of literals. The empty clause is *false*. A formula is in *conjunctive normal form (CNF)* if it is a conjunction of clauses. A *cube* or *term* d is a conjunction of a consistent set of literals; at times, we also refer directly to the set and write $\ell \in d$. The empty cube is *true*. A formula is in *disjunctive normal form (DNF)* if it is a disjunction of terms. The *domain*, $dom(d)$, of a cube d is the set of variables that appear in it (positively or negatively). Given a state σ , we use the state and the (full) cube that consists of all the literals that are satisfied in σ interchangeably. $|\varphi|_{\text{dnf}}$ is the minimal number of terms in any DNF representation of φ . $|\varphi|_{\text{cnf}}$ is the minimal number of clauses in any CNF representation of φ .

3 Background: The Monotone Theory

This section provides necessary definitions and results from the monotone theory by Bshouty [5] as used in this paper. Our presentation is based on [13, 14] (lemmas that are stated here slightly differently are proved in the extended version [15]).

Boolean functions which are *monotone* are special in many ways; one is that they are easier to learn [e.g. 41, 1]. Syntactically, a monotone function can be written in DNF so that all variables appear positively. This is easily generalized to *b*-monotone formulas, where each variable appears only at one polarity specified by *b* (Def. 2). The *monotone theory* aims to handle functions that are not monotone through the conjunction of *b*-monotone formulas. §3.1 considers the (over)approximation of a formula by a *b*-monotone formula, the “monotonization” of a formula; §3.2 studies the conjunction of several such monotoneizations through the *monotone hull* operator.

3.1 Least *b*-Monotone Overapproximations

Definition 1 (*b*-Monotone Order). *Let b be a cube. We define a partial order over states where $v \leq_b x$ when x, v agree on all variables not present in b , and x disagrees with b on all variables on which also v disagrees with b : $\forall p \in \Sigma. x[p] \neq v[p]$ implies $p \in \text{dom}(b) \wedge v[p] = b[p]$.*

Intuitively, $v \leq_b x$ when x can be obtained from v by flipping bits to the opposite of their value in b .

Definition 2 (*b*-Monotonicity). *A formula ψ is *b*-monotone for a cube b if $\forall v \leq_b x. v \models \psi$ implies $x \models \psi$.*

That is, if v satisfies ψ , so do all the states that are farther away from b than v . For example, if ψ is 000-monotone and $100 \models \psi$, then because $100 \leq_{000} 111$ (starting in 100 and moving away from 000 can reach 111), also $111 \models \psi$. In contrast, $100 \not\leq_{000} 011$ (the same process cannot flip the 1 bit that already disagrees with 000), so 011 does not necessarily belong to ψ . (000-monotonicity corresponds to the usual notion of monotone formulas.)

Definition 3 (Least *b*-Monotone Overapproximation). *For a formula φ and a cube b , the least *b*-monotone overapproximation of φ is a formula $\mathcal{M}_b(\varphi)$ defined by*

$$x \models \mathcal{M}_b(\varphi) \text{ iff } \exists v. v \leq_b x \wedge v \models \varphi.$$

For example, if $100 \models \varphi$, then $100 \models \mathcal{M}_{000}(\varphi)$ because $\mathcal{M}_{000}(\varphi)$ is an overapproximation, and hence $111 \models \mathcal{M}_{000}(\varphi)$ because it is 000-monotone, as above. Here, thanks to minimality, 011 does not belong to $\mathcal{M}_{000}(\varphi)$, unless 000, 001, 010, or 011 belong to φ .

The minimality property of $\mathcal{M}_b(\varphi)$ is formalized as follows:

Lemma 1. $\mathcal{M}_b(\varphi)$ (Def. 3) is the least b -monotone formula ψ (Def. 2) s.t. $\varphi \implies \psi$ (i.e., for every other b -monotone formula ψ , if $\varphi \implies \psi$ then $\mathcal{M}_b(\varphi) \implies \psi$).

An immediate but useful fact is that $\mathcal{M}_b(\cdot)$ is a monotone operator:

Lemma 2. If $\varphi_1 \implies \varphi_2$ then $\mathcal{M}_b(\varphi_1) \implies \mathcal{M}_b(\varphi_2)$.

Syntactic intuition. Ordinary monotone formulas are $\bar{0}$ -monotone; for general b , a formula ψ in DNF is b -monotone if interchanging $p, \neg p$ whenever $b[p] = \text{true}$ results in a formula that is monotone DNF per the standard definition.¹ When φ is not b -monotone, the monotonization $\mathcal{M}_b(\varphi)$ is the “closest thing”, in the sense that it is the smallest b -monotone ψ s.t. $\varphi \implies \psi$. As we shall see, $\mathcal{M}_b(\varphi)$ can be efficiently obtained from φ by deleting literals. The syntactic viewpoint is key for our results in §4 and §5.

Geometric intuition. Geometrically, ψ is b -monotone if $v \models \psi \implies x \models \psi$ for every states v, x s.t. $v \leq_b x$; the partial order \leq_b indicates that x is “farther away” from b in the Hamming cube than v from b , namely, that there is a shortest path w.r.t. Hamming distance from b to x (or from $\pi_b(x)$ —the projection of x onto b —to x , when b is not a full cube) that goes through v . A formula ψ is b -monotone when it is closed under this operation, of getting farther from b . In this way, $\mathcal{M}_b(\varphi)$ corresponds to the set of states x to which there is a shortest path from b that intersects φ .² The geometric viewpoint is key for our results in §4 and for the abstract domain in §6.

Disjunctive form. The monotone overapproximation can be obtained from a DNF representation of the original formula, a fact that is useful for algorithms that compute the monotone overapproximation. Starting with a DNF representation of φ , we can derive a DNF representation of $\mathcal{M}_b(\varphi)$ by dropping in each term the literals that agree with b . Intuitively, if ℓ agrees with b , the “constraint” that $\sigma \models \ell$ is dropped from $\mathcal{M}_b(t)$ because if $\sigma \models \mathcal{M}_b(t)$ then flipping the value of ℓ in σ results in a state $\tilde{\sigma}$ such that $\sigma \leq_b \tilde{\sigma}$ and hence also $\tilde{\sigma} \models \mathcal{M}_b(t)$.

Lemma 3. Let $\varphi = t_1 \vee \dots \vee t_m$ in DNF. Then the monotonization $\mathcal{M}_b(\varphi) \equiv \mathcal{M}_b(t_1) \vee \dots \vee \mathcal{M}_b(t_m)$ where $\mathcal{M}_b(t_i) \equiv t_i \setminus b = \bigwedge \{\ell \in t_i \wedge \ell \notin b\}$.

This fact has several useful corollaries. First, for the important special case of a state (full cube) v , the monotonization $\mathcal{M}_b(v)$ is the conjunction of all literals that hold in v except those that are present in b , written

$$\text{cube}_b(v) \stackrel{\text{def}}{=} \mathcal{M}_b(v) = \bigwedge \{p \mid v[p] = \text{true}, p \notin b\} \wedge \bigwedge \{\neg p_i \mid v[p_i] = \text{false}, \neg p_i \notin b\}.$$

¹ When b is a full cube, another way to say this is that ψ is b -monotone if it is monotone in the ordinary sense under the translation [42] specified by b .

² This is reminiscent of visibility in Euclidean geometry [e.g. 29]: picturing b as a guard, the source of visibility, then $\mathcal{M}_b(\varphi)$ is the set of states that are visible in $\neg\varphi$, that is, the set of states σ s.t. the “line segment” $[b, \sigma]$ is contained in $\neg\varphi$. Here $[b, \sigma]$ is the Hamming interval [e.g. 42] between b, σ , the union of all the multiple shortest paths between the states (each path corresponds to a different permutation of the variables on which the states disagree).

In particular, if $v \models \varphi$ then $\text{cube}_b(v) \implies \mathcal{M}_b(\varphi)$ (follows from Lemma 3 thinking about the representation $v \vee \varphi$). A similar property holds under the weaker premise that v is known to belong to the monotone hull:

Lemma 4. *If $v \models \mathcal{M}_b(\varphi)$ then $\text{cube}_b(v) \implies \mathcal{M}_b(\varphi)$.*

Another corollary is that the DNF size cannot increase from φ to $\mathcal{M}_b(\varphi)$:

Lemma 5. $|\mathcal{M}_b(\varphi)|_{\text{dnf}} \leq |\varphi|_{\text{dnf}}$.

3.2 Monotone Hull

We now define the monotone hull, which is a conjunction of b -monotone overapproximations over all the b 's from a fixed set of states B .

Definition 4 (Monotone Hull). *The monotone hull of a formula φ w.r.t. a set of states B is $\text{MHull}_B(\varphi) = \bigwedge_{b \in B} \mathcal{M}_b(\varphi)$.*

The monotone hull can be simplified to use a succinct DNF representation of the basis B instead of a conjunction over all states.

Lemma 6. *If $B \equiv b_1 \vee \dots \vee b_m$ where b_1, \dots, b_m are cubes, then $\text{MHull}_B(\varphi) \equiv \mathcal{M}_{b_1}(\varphi) \wedge \dots \wedge \mathcal{M}_{b_m}(\varphi)$.*

Note that when $B = b$ is a single cube, $\text{MHull}_b(\varphi) = \mathcal{M}_b(\varphi)$.

Similarly to $\mathcal{M}_b(\varphi)$, the monotone hull is an overapproximation:

Lemma 7. $\varphi \implies \text{MHull}_B(\varphi)$.

In general, $\text{MHull}_B(\varphi)$ is not equivalent to φ . However, we can always choose B so that $\text{MHull}_B(\varphi) \equiv \varphi$. A set B that suffices for this is called a basis:

Definition 5 (Monotone Basis). *A monotone basis is a set of states B . It is a basis for a formula φ if $\varphi \equiv \text{MHull}_B(\varphi)$.*

Conversely, given a set B , we are interested in the set of formulas for which B forms a basis:

Definition 6 (Monotone Span). $\text{MSpan}(B) = \{\text{MHull}_B(\varphi) \mid \varphi \text{ over } \Sigma\}$, the set of formulas for which B is a monotone basis.

The following theorem provides a syntactic characterization of $\text{MSpan}(B)$, as the set of all formulas that can be written in CNF using clauses that exclude states from the basis. The connection between CNF and the monotone basis is useful in §5 where a monotone basis is constructed automatically, and in §6 where it is used to define an abstract domain.

Theorem 1 ([5]). $\varphi \in \text{MSpan}(B)$ iff there exist clauses c_1, \dots, c_s such that $\varphi \equiv c_1 \wedge \dots \wedge c_s$ and for every $1 \leq i \leq s$ there exists $b_j \in B$ such that $b_j \not\models c_i$.

In particular, a basis B for φ can be constructed by writing a CNF representation of φ and choosing for B a state $b_j \not\models c_j$ for each clause c_j .

Exact learning using the monotone theory. The monotone theory was first developed by Bshouty for the purpose of exact learning formulas that are not monotone. Essentially, the idea is to reconstruct the formula φ by finding a monotone basis $B = \{b_1, \dots, b_t\}$ for it, and constructing $\text{MHull}_B(\varphi)$ while using equivalence and membership queries. The CDNF algorithm [5] achieves efficient learning in terms of the DNF & CNF size of the target formula, and its code is shown in the extended version [15]. Our CDNF invariant inference algorithm (§5) is inspired by it, although it departs from it in significant ways (see Remark 3).

4 Super-Efficient Monotonization

In this section we develop an efficient procedure to compute $\mathcal{M}_b(\varphi)$, which is a technical enabler of the results in following sections. The algorithm, presented in Alg. 1, satisfies the following:

Theorem 2. *Let φ be a formula and b a cube. The algorithm $\text{MONOTONIZE}(\varphi, b)$ computes $\mathcal{M}_b(\varphi)$ in $O(n^2|\mathcal{M}_b(\varphi)|_{\text{dnf}})$ SAT queries and time.*

(Throughout this paper, n denotes the number of propositional variables $|\Sigma|$.)

What distinguishes Thm. 2 is that the complexity bound depends on the DNF size of the *output*, the monotonization $\mathcal{M}_b(\varphi)$, and *not* on the size of the *input* φ , in contrast to the algorithm by Bshouty [5] (see Remark 1).

Algorithm 1 Super-Efficient Monotonization

1: procedure MONOTONIZE(φ, b) 2: $H \leftarrow false$ 3: while SAT($\varphi \wedge \neg H$) do 4: let $\sigma_r \models \varphi \wedge \neg H$ 5: $v \leftarrow \text{GENERALIZE}(\varphi, b, \sigma_r)$ 6: $H \leftarrow H \vee \text{cube}_b(v)$ 7: return H 8: 9: 10: 11:	12: procedure GENERALIZE(φ, b, σ_r) 13: $v \leftarrow \sigma_r$; walked $\leftarrow true$ 14: while walked do 15: walked $\leftarrow false$ 16: for $j = 1, \dots, n$ do 17: if $b[p_j] = v[p_j]$ then 18: continue 19: $x \leftarrow v[p_j \mapsto b[p_j]]$ 20: if SAT($\varphi \wedge \overline{x, \pi_b(x)}$) then 21: $v \leftarrow x$; walked $\leftarrow true$ 22: return v
---	---

Starting from the candidate $H = false$, the algorithm iteratively samples—through satisfying models of a SAT query—states that belong in φ but not yet included in H . Every such state σ_r generates a new term in H . Since H is supposed to be b -monotone, the minimal term to include is $\text{cube}_b(\sigma_r)$. To be efficient, the algorithm generalizes each example, trying to flip bits to find an

example v that also should be included in $\mathcal{M}_b(\varphi)$ and is closer in Hamming distance to b , which would result in a smaller term $\text{cube}_b(v)$, thereby including more states in each iteration and converging faster. The criterion for v is that a bit cannot be flipped if this would result in a state x where the Hamming interval $\overline{x, \pi_b(x)}$ does not intersect φ . Here, $\pi_b(x)$, the projection [e.g. 42] of x onto the (possibly partial) cube b is the state s.t.

$$\pi_b(x) = \begin{cases} b[p] & p \in \text{dom}(b) \\ x[p] & \text{otherwise} \end{cases},$$

and the Hamming interval $\overline{\sigma_1, \sigma_2}$ between two states σ_1, σ_2 is the smallest cube that contains both—the conjunction of the literals where these agree. In sum, $\overline{x, \pi_b(x)}$ is the conjunction of the literals where x, b agree and the literals of x over variables that are not present in b . As we will show, $\overline{x, \pi_b(x)}$ intersecting with φ is an indicator for x belonging to the monotoneization of φ .

The use of SAT queries in the algorithm does not necessarily assume that φ is given explicitly, and indeed in §5 we apply this algorithm with an implicit representation of φ (using additional copies of the vocabulary).

The rest of this section proves Thm. 2. First, the result v of generalization is so that when we disjoin the term $\text{cube}_b(v)$, we do not “overshoot” to include states that do not belong to the true monotoneization:

Lemma 8. *If $\sigma_r \models \varphi$, then $\text{GENERALIZE}(\varphi, b, \sigma_r)$ returns v s.t. $\text{cube}_b(v) \implies \mathcal{M}_b(\varphi)$.*

Proof. v is chosen s.t. $\overline{v, \pi_b(v)} \cap \varphi \neq \emptyset$ —note that this holds trivially in the initial choice of v which is $\sigma_r \models \varphi$. Let $\tilde{\sigma} \models \varphi$ s.t. $\tilde{\sigma} \models \overline{v, \pi_b(v)}$. The latter means that $\tilde{\sigma} \leq_b v$, because $\overline{v, \pi_b(v)}$ consists of all the literals in v except for those that disagree with b , so $\tilde{\sigma}$ agrees with v whenever v, b agree. In more detail, $\tilde{\sigma}$ agrees with v on all $p \notin \text{dom}(b)$ (because $\pi_b(v)[p] = v[p]$ on such variables), and for $p \in \text{dom}(p)$, if $\tilde{\sigma}[p] \neq v[p]$, if v, b agree on p then likewise $\tilde{\sigma}$ agrees with them (because then $v[p] = \pi_b(v)[p]$ and p is retained in the conjunction that forms the Hamming interval), which satisfies Def. 2. As also $\tilde{\sigma} \models \varphi$, this implies that $v \models \mathcal{M}_b(\varphi)$ per Def. 3. Hence $\text{cube}_b(v) \implies \mathcal{M}_b(\varphi)$, by Lemma 4. \square

This shows that it is reasonable to disjoin the term $\text{cube}_b(v)$ to H in the hope of eventually obtaining $H = \mathcal{M}_b(\varphi)$. The following lemma argues that the algorithm continues to sample states until it converges to the true monotoneization.

Lemma 9. *$\text{MONOTONIZE}(\varphi, b)$ terminates and returns $\mathcal{M}_b(\varphi)$.*

Proof. First we show that when it terminates, the result is correct. Always $H \subseteq \mathcal{M}_b(\varphi)$, because in each iteration we disjoin to H a formula that satisfies the same property, by Lemma 8. The algorithm terminates when $\varphi \subseteq H$, and H is always a b -monotone formula (by Lemma 3 the monotoneization of H is H itself, which is b -monotone as in Lemma 1). From the minimality of $\mathcal{M}_b(\varphi)$ (Lemma 1), necessarily also $\mathcal{M}_b(\varphi) \subseteq H$.

To show termination it suffices to show that H strictly increases in each iteration (because the number of non-equivalent propositional formulas is finite). To see this, note that $\text{GENERALIZE}(\varphi, b, \sigma_r)$ returns v s.t. $v \leq_b \sigma_r$, since the procedure starts with σ_r and only flips literals to agree with b . This implies that $\sigma_r \models \text{cube}_b(v)$, so after the iteration $\sigma_r \models H$ whereas previously $\sigma_r \not\models H$. \square

The novelty of the algorithm is its efficiency, which we now turn to establish. The crucial point is the generalization is able to produce, term by term, a minimal representation of $\mathcal{M}_b(\varphi)$. To this end, we first show that $\text{cube}_b(v)$ that the algorithm computes in lines 5 to 6 is a *prime implicant* of $\mathcal{M}_b(\varphi)$. Recall that a term t is an *implicant* of a formula ψ if $t \implies \psi$, and it is *prime* if this no longer holds after dropping a literal, that is, for every $\ell \in t$ (as a set of literals), $(\wedge (t \setminus \{\ell\})) \not\implies \psi$. It is *non-trivial* if $t \not\equiv \text{false}$ (not an empty set of literals).

Lemma 10. *If $\sigma_r \models \varphi$, then $\text{GENERALIZE}(\varphi, b, \sigma_r)$ returns v s.t. $\text{cube}_b(v)$ is a non-trivial prime implicant of $\mathcal{M}_b(\varphi)$.*

Proof. Lemma 8 shows that it is an implicant. It is non-trivial because σ_r is a model of it, as shown as part of the proof of Lemma 9. Suppose that $\text{cube}_b(v)$ is not prime. Then there a literal over some variable p that can be dropped. It is present in $\text{cube}_b(v)$, which means that $p \in \text{dom}(b)$ and $v[p] \neq b[p]$. Then the cube obtained from dropping the literal can be written as $\text{cube}_b(x)$ where $x = v[p \mapsto \neg v[p]]$. If this cube is an implicant of $\mathcal{M}_b(\varphi)$, then, because $x \models \text{cube}_b(x)$, in particular $x \models \mathcal{M}_b(\varphi)$. By Def. 3, there is $\tilde{\sigma} \models \varphi$ such that $\tilde{\sigma} \leq_b x$. But the latter implies that $\tilde{\sigma} \in \overline{x, \pi_b(x)}$, because, by Def. 2, for every $p \notin \text{dom}(b)$, $\tilde{\sigma}[p] = x[p] = \pi_b(x)[p]$ and for every $p \in \text{dom}(b)$ where $x, \pi_b(x)$ agree also $\tilde{\sigma}, b$ agree (because $\pi_b(x)[p] = b[p]$). Thus $\overline{x, \pi_b(x)} \cap \varphi \neq \emptyset$, in contradiction to the choice of v , according to the check in line 20. \square

Thanks to the fact that $\mathcal{M}_b(\varphi)$ is b -monotone, through one of the basic properties of monotone functions that dates back to Quine [30], a prime implicant reproduces a term from the (unique) minimal representation of $\mathcal{M}_b(\varphi)$. We use this to show that the monotonization is computed in few iterations:

Lemma 11. *The number of iterations of the loop in line 3 of $\text{MONOTONIZE}(\varphi, b)$ is at most $|\mathcal{M}_b(\varphi)|_{\text{dnf}}$.*

Proof. Lemma 10 shows that in each iteration we disjoin a prime implicant. It is a property of monotone functions that they have a unique DNF representations with irredundant, which consists of the disjunction of all non-trivial prime implicants [30], and this extends to b -monotone functions (through a simple renaming of variables to make the function monotone). Thus the non-trivial prime implicant we disjoin is a term of the minimal DNF representation of $\mathcal{M}_b(\varphi)$. Each additional σ_r produces a new term, as shown in Lemma 9. \square

We are now ready to prove that the algorithm overall is efficient.

Proof (Proof of Thm. 2). By Lemma 11 the number of iterations of the loop in line 3 is bounded by $|\mathcal{M}_b(\varphi)|_{\text{dnf}}$. Each iteration calls GENERALIZE, which performs at most n iterations of the loop in line 14 because the same variable is never flipped twice. Each iteration of this loop performs n SAT queries in line 20. Note that the cube $\overline{x, \pi_b(x)}$ is straightforward to compute in linear time. \square

Remark 1. Bshouty [5] used an algorithm for computing $\mathcal{M}_b(\varphi)$ whose complexity is bounded by the DNF input size $|\varphi|_{\text{dnf}}$, whereas Alg. 1’s complexity is bounded by the DNF *output* size, $|\mathcal{M}_b(\varphi)|_{\text{dnf}}$, which is never worse (Lemma 5), and sometimes significantly smaller. When considered as learning algorithms, the improved complexity of Alg. 1 comes at the expense of the need for richer queries: Bshouty’s algorithm is similar to Alg. 1 (using an equivalence query in line 3 that produces a positive example—see GENMQ of Alg. 6 in Appendix B), except that the condition in line 20 is replaced by checking whether $x \models \varphi$. This is a membership query to φ , whereas our check amounts to a disjointness query [1].

5 Efficient Inference of CDFN Invariants

Algorithm 2 Dual of model-based interpolation-based inference [6, 3]	Algorithm 3 Interpolation-based inference of CDFN invariants
1: procedure Dual-MB-ITP($Init, \delta, Bad, s$) 2: if $\underline{\delta}^s(Init) \cap Bad \neq \emptyset$ then 3: unsafe 4: $\varphi \leftarrow \neg Bad$ 5: while φ not inductive do 6: let $\sigma, \sigma' \models \varphi \wedge \delta \wedge \neg \varphi'$ 7: if $\underline{\delta}^s(Init) \cap \{\sigma\} \neq \emptyset$ then 8: restart with larger s 9: take minimal clause $c \subseteq \neg \sigma$ s.t. $\underline{\delta}^s(Init) \implies c$ 10: $\varphi \leftarrow \varphi \wedge c$ 11: return φ	1: procedure CDFN-ITP($Init, \delta, Bad, s$) 2: if $\underline{\delta}^s(Init) \cap Bad \neq \emptyset$ then 3: unsafe 4: $\varphi \leftarrow \neg Bad$ 5: while φ not inductive do 6: let $\sigma, \sigma' \models \varphi \wedge \delta \wedge \neg \varphi'$ 7: if $\underline{\delta}^s(Init) \cap \{\sigma\} \neq \emptyset$ then 8: restart with larger s 9: $H \leftarrow \text{MONOTONIZE}(\underline{\delta}^s(Init), \sigma)$ 10: $\varphi \leftarrow \varphi \wedge H$ 11: return I

In this section we build on the algorithm of §4 to devise a new model-based interpolation-based algorithm that can efficiently infer invariants that have poly-size CNF and DNF representations (dubbed “CDFN invariants”). We start with background on the theoretical condition that guarantees the success of the original model-based algorithm for simpler forms of invariants.

5.1 Background: Interpolation With the Fence Condition

The essence of interpolation-based invariant inference (ITP) is to generalize a proof of *bounded* unreachability—i.e., bounded model checking [2]—into a proof

of *unbounded* reachability, that is, a part of the inductive invariant. The reader is more likely to be familiar with the structure of the original algorithm by McMillan [26], which uses bounded unreachability *to* the *bad* states, where in each iteration the algorithm adds states to the candidate by *disjoining* a formula from which it is impossible to reach a bad state in s steps. However, for our purposes here, it is more convenient to consider the *dual* version, which uses bounded unreachability *from* the *initial* states, where in each iteration the algorithm excludes states from the candidate by *conjoining* a formula which does not exclude any state that the system can reach in s steps from an initial state (i.e., the candidate φ is updated by $\varphi \leftarrow \varphi \wedge H$ where H is a formula s.t. $\underline{\delta}^s(\text{Init}) \subseteq H$).

The original interpolation-based algorithm by McMillan uses a procedure that relies on the internals of the SAT solver [26]. Complexity bounds on interpolation-based algorithms analyze later approaches that exercise control on how interpolants are generated, and do this in a model-based fashion [6, 3] inspired by IC3/PDR [4, 9]. The dual version is presented in Alg. 2. After starting the candidate φ as $\neg \text{Bad}$, each iteration checks for a counterexample to induction (line 6), whose pre-state σ is excluded from φ at the end of the iteration (line 10). Many states are excluded in each iteration beyond the counterexample, by conjoining to the candidate a minimal *clause* that excludes σ but retains all the states that are reachable in the system in s steps (line 9—this involves up to n queries of s -BMC, each time dropping a literal and checking whether the clause is still valid). If the counterexample cannot be blocked, because it is in fact reachable in s steps, this is an indication that s needs to be larger (line 7) to find a proof or a safety violation. The algorithm detects that the transition system is unsafe in line 3 when s is enough to find an execution from *Init* to *Bad* with at most s transitions. (Our analysis of the algorithms in the paper focuses on the safe case, the complexity of finding an invariant.)

A condition that guarantees that s is large enough for Alg. 2 to successfully find an inductive invariant, called the *fence condition*, was recently put forward [13], involving the Hamming-geometric boundary of the invariant.

Definition 7 (Boundary). *Let I be a set of states. Then the (inner) boundary of I , denoted $\partial^+(I)$, is the set of states $\sigma^+ \models I$ s.t. there is a state σ^- that differs from σ in exactly one variable, and $\sigma^+ \models I, \sigma^- \not\models I$.*

Definition 8 (Fence Condition). *Let I be an inductive invariant for a transition system $(\text{Init}, \delta, \text{Bad})$ and $s \in \mathbb{N}$. Then I is s -forwards fenced if $\partial^+(I) \subseteq \underline{\delta}^s(\text{Init})$.*

Example 1. Let I be the set of all states where at least two bits are 0 *and* at least two bits are 1. Then $\partial^+(I)$ is the set where exactly two bits are 0 (and at least two bits are 1) or exactly two bits are 1 (and at least two bits are 0). Note that $I \setminus \partial^+(I)$ contains many (most) states—those where three or more bits are 0 and three or more bits are 1. The fence condition requires only from the states in $\partial^+(I)$ to be reachable in s steps.

The fence condition guarantees that throughout the algorithm’s execution, φ contains the fenced invariant, giving rise to the following correctness property:

Theorem 3 ([13]). *If there exists an s -forwards fenced invariant for $(Init, \delta, Bad)$, then $\text{Dual-MB-ITP}(Init, \delta, Bad, s)$ successfully finds an invariant.*

The idea is that the property that $I \implies \varphi$ is maintained because the fence condition ensures that it suffices to verify that the clause that GENERALIZE computes, starting from a state $\sigma \models \varphi$ that is outside of I , does not exclude a state from $\underline{\delta}^s(Init)$ to guarantee that it also does not exclude a state from I . (Note that the fence condition does not provide a way to know whether an *arbitrary* state belongs to I .)

The fence condition ensures the algorithm’s success, but not that it is efficient—the number of iterations until convergence may be large, even when there is a fenced inductive invariant that has a short representation in CNF. (Note that in Alg. 2, φ is always in CNF.) This was ameliorated in [13] by the assumption that the invariant is *monotone*:

Theorem 4 ([13]). *If I is an s -forwards fenced inductive invariant for $(Init, \delta, Bad)$ and I is monotone (can be written in CNF/DNF with all variables un-negated), then $\text{Dual-MB-ITP}(Init, \delta, Bad, s)$ successfully finds an inductive invariant in $O(|I|_{\text{cnf}})$ inductiveness checks, and $O(n \cdot |I|_{\text{cnf}})$ checks of s -BMC and time.³*

However, when I is *not* monotone, it is possible for the algorithm to require an exponential number of iterations even though $|I|_{\text{cnf}}$ is small (and in fact, even though *every* representation of I without redundant clauses is small).

The challenge that we address in this section is to create an invariant inference algorithm that efficiently infers inductive invariants that are not monotone, while relying only on the fence condition.

5.2 CDNF Inference With the Fence Condition

We now present our new invariant inference algorithm (Alg. 3), that is guaranteed to run in time polynomial in $n, |I|_{\text{cnf}}, |I|_{\text{dnf}}$ of a fenced invariant I :

Theorem 5. *Let I be a forwards s -fenced inductive invariant for $(Init, \delta, Bad)$. Then $\text{CDNF-ITP}(Init, \delta, Bad, s)$ finds an inductive invariant in at most $|I|_{\text{cnf}} \cdot |I|_{\text{dnf}} \cdot n^2$ of s -BMC checks, $|I|_{\text{cnf}}$ inductiveness checks, and $O(|I|_{\text{cnf}} \cdot |I|_{\text{dnf}} \cdot n^2)$ time.*

Example 2. I from Example 1 has poly-size representations in both CNF and DNF. We write:

³ A similar result applies when the invariant is *unate*, that is, can be written in CNF/DNF so that every variable is either always negated or always un-negated.

DNF: there is a choice of four bits with two 0's and two 1's,

$$I \equiv \bigvee_{1 \leq i_1 \neq i_2 \neq i_3 \neq i_4 \leq n} (x_{i_1} = 0 \wedge x_{i_2} = 0 \wedge x_{i_3} = 1 \wedge x_{i_4} = 1).$$

CNF: it is impossible that $n - 1$ bits or more are 1, likewise for 0,

$$I \equiv \left(\bigwedge_{i=1}^n \bigvee_{j \neq i} x_j = 0 \right) \wedge \left(\bigwedge_{i=1}^n \bigvee_{j \neq i} x_j = 1 \right).$$

The CNF formula has $2n$ clauses, and the DNF has $\binom{n}{4} = \Theta(n^4)$ terms. Thm. 5 shows that such I satisfying the fence condition can be inferred in a number of queries and time that is polynomial in n . Note that these formulas fall outside the classes that previous results can handle efficiently (see §7) as they are not monotone nor almost-monotone (the number of terms/clauses with negated variables is not constant).

As noted by Bshouty [5], the class of formulas with short DNF and CNF includes the formulas that can be expressed by a small *decision tree*: a binary tree in which every internal node is labeled by a variable and a leaf by *true/false*, and σ satisfies the formula if the path defined by starting from the root, turning left when the σ assigns *false* to the variable labeling the node and right otherwise, reaches a leaf *true*. The size of a decision tree is the number of leaves in the tree. We conclude that (see the proof in the extended version [15]):

Corollary 1. *Let I be a forwards s -fenced inductive invariant for $(Init, \delta, Bad)$, that can be expressed as a decision tree of size m . Then $CDNF\text{-ITP}(Init, \delta, Bad, s)$ finds an inductive invariant in at most $m^2 \cdot n^2$ of s -BMC checks, m inductiveness checks, and $O(m^2 \cdot n^2)$ time.*

The algorithm $CDNF\text{-ITP}$ which attains Thm. 5 is presented in Alg. 3. Its overall structure is similar to Alg. 2, except the formula used to block a counterexample is the monotonization of the s -reachable states. Specifically, starting from the candidate $\varphi = true$ (line 4, the algorithm iteratively samples counterexamples to induction (line 6) and blocks the pre-state σ from φ by conjoining $\mathcal{M}_\sigma(\underline{\delta}^s(Init))$, computed by invoking Alg. 1. The SAT queries of the form $SAT(\varphi \wedge \theta)$ that Alg. 1 performs (see §4) have $\varphi = \underline{\delta}^s(Init)$, and they amount to the BMC checks of whether $\underline{\delta}^s(Init) \cap \theta \stackrel{?}{=} \emptyset$.

It is important for the efficiency result that Alg. 3 uses Alg. 1 as a sub-procedure. Using Bshouty's procedure (see Remark 1) would yield a bound of $n \cdot |\underline{\delta}^s(Init)|_{\text{dnf}}$ checks of s -BMC, and it is likely that $\underline{\delta}^s(Init)$ is complex to capture in a formula when s is significant (as common for sets defined by exact reachability, such as the set of the reachable states).

We now proceed to prove the correctness and efficiency of the algorithm (Thm. 5). Throughout, assume that I is an inductive invariant for $(Init, \delta, Bad)$. I will be s -forwards fenced; we state this explicitly in the premise of lemmas where this assumption is used. The idea behind the correctness and efficiency

of Alg. 3 is that $\mathcal{M}_\sigma(I)$ is a stronger formula than the clauses that are produced in Alg. 2, causing the candidate to converge down to the invariant in fewer iterations, while never excluding states that belong to I (because $I \subseteq \mathcal{M}_\sigma(I)$, as used in Lemma 14). As we will show (in Lemma 15), this strategy results in a number of iterations that is bounded by the CNF size of I (without further assumptions on the syntactic structure of I). The trick, however, is to show (in Lemma 13) that what the algorithm computes in line 9 is indeed $\mathcal{M}_\sigma(I)$, even though I is unknown. The crucial observation is that under the fence condition, the monotonicization of the s -reachable states matches the monotonicization of the invariant (even though these are different sets!). Note that this holds for any invariant that satisfies the fence condition. To prove this we need to recall a fact about the monotonicization of the boundary of a set:

Lemma 12 ([14]). *Let I, S be sets of states s.t. $\partial^+(I) \subseteq S$ and σ a state s.t. $\sigma \not\models I$. Then $I \subseteq \mathcal{M}_\sigma(S)$.*

The idea is that for every $x \in I$, there is a state on the boundary $v \in \partial^+(I)$ s.t. $v \leq_\sigma x$ (where a shortest path between x, b crosses I), and because also $v \in S$ we would have that $x \in \mathcal{M}_\sigma(S)$.

We proceed to relate the monotonicizations of $\underline{\delta}^s(\text{Init}), I$:

Lemma 13. *If I is forwards s -fenced for $(\text{Init}, \delta, \text{Bad})$, and $\sigma \not\models I$, then we have that $\mathcal{M}_\sigma(\underline{\delta}^s(\text{Init})) = \mathcal{M}_\sigma(I)$.*

Proof. Since I is an inductive invariant, $\underline{\delta}^s(\text{Init}) \subseteq I$, so $\mathcal{M}_\sigma(\underline{\delta}^s(\text{Init})) \subseteq \mathcal{M}_s(I)$ from Lemma 2. For the other direction we use Lemma 12: by the fence condition, $\partial^+(I) \subseteq \underline{\delta}^s(\text{Init})$ and hence, as $\sigma \not\models I$, we obtain $I \subseteq \mathcal{M}_\sigma(\underline{\delta}^s(\text{Init}))$. By Lemma 1, this implies that $\mathcal{M}_\sigma(I) \subseteq \mathcal{M}_\sigma(\underline{\delta}^s(\text{Init}))$. \square

We use this to characterize the candidate invariant the algorithm constructs:

Lemma 14. *If I is forwards s -fenced for $(\text{Init}, \delta, \text{Bad})$, then in each step of CDNF-ITP($\text{Init}, \delta, \text{Bad}, k$), $\varphi = \text{MHull}_{\mathcal{C}_i}(I) \wedge \neg \text{Bad}$, where \mathcal{C}_i is the set of counterexamples σ the algorithm has observed so far. In particular, $I \subseteq \varphi$.*

Proof. First, $I \subseteq \varphi$ holds from the rest of the lemma because $I \subseteq \neg \text{Bad}$ (it is an inductive invariant), and $I \subseteq \text{MHull}_{\mathcal{C}_i}(I)$ by Lemma 7. The proof of $\varphi = \text{MHull}_{\mathcal{C}_i}(I) \wedge \neg \text{Bad}$ is by induction on iterations of the loop in line 5. Initially, $\mathcal{C} = \emptyset$ and indeed $\varphi = \neg \text{Bad}$. In each iteration, $I \subseteq \varphi$ using the argument above and the induction hypothesis. Hence, the counterexample to induction of line 6 has $\sigma \not\models I$ (otherwise $\sigma' \models I$ because I is an inductive invariant, and this would imply also $\sigma \models \varphi$, in contradiction). Then Lemma 9 ensures that $H = \mathcal{M}_\sigma(\underline{\delta}^s(\text{Init}))$. Lemma 13 shows that this is $\mathcal{M}_\sigma(I)$, as required. \square

Essentially, the algorithm gradually learns a monotone basis (Def. 5) for I from the counterexamples to induction, and constructs I via the monotone hull w.r.t. this basis. The next lemma shows that the size of the basis that the algorithm finds is bounded by $|I|_{\text{cnf}}$.

Lemma 15. *If I is forwards s -fenced for $(Init, \delta, Bad)$, then $CDNF\text{-ITP}(Init, \delta, Bad, k)$ successfully finds an inductive invariant. Further, the number of iterations of the loop in line 5 is at most $|I|_{\text{cnf}}$.*

Proof. Since $\sigma \not\models I$, also it is not a model of the monotonization w.r.t. to itself, $\sigma \not\models \mathcal{M}_\sigma(I)$ (because the only state $x \leq_\sigma \sigma$ is $x = \sigma$ —see Defs. 1 and 3). This shows, using Lemma 14, that at least one state is excluded from the candidate φ in each iteration. By the same lemma always $I \implies \varphi$, and the algorithm terminates when φ is inductive, so this shows that the algorithm successfully converges to an inductive invariant.

To see that this occurs in at most $|I|_{\text{cnf}}$ iterations, consider a minimal CNF representation of I , $I = c_1 \wedge \dots \wedge c_{|I|_{\text{cnf}}}$. We argue that in each iteration produces at least one new clause from that representation, in the sense that for some i , $\varphi \wedge \mathcal{M}_{\sigma_b}(I) \implies c_i$ whereas previously $\varphi \not\implies c_i$. Let c_i be the clause that $\sigma \not\models c_i$ (recall e.g. that $\sigma \not\models \varphi$ and $I \subseteq \varphi$). Then $\mathcal{M}_\sigma(I) \subseteq c_i$, since c_i is σ -monotone ($\mathcal{M}_\sigma(c_i) = c_i$, using Lemma 3, because all the literals disagree with σ) and $I \subseteq c_i$, and $\mathcal{M}_\sigma(I)$ is the smallest such (Lemma 1). Thus when we conjoin $H = \mathcal{M}_\sigma(I)$ to φ we conjoin at least one new c_i that was not present in a CNF representation of φ ; this can happen at most $|I|_{\text{cnf}}$ times. \square

Overall:

Proof (Proof of Thm. 5). The algorithm’s success in finding an invariant is established in Lemma 15. As for efficiency, by Lemma 15, there are at most $|I|_{\text{cnf}}$ iterations of the loop in $CDNF\text{-ITP}$, each performs a single inductiveness query, and calls MONOTONIZE . By Thm. 2 each such call performs at most $O(n^2 |\mathcal{M}_\sigma(I)|_{\text{dnf}})$ s -BMC queries. The claim follows because $|\mathcal{M}_\sigma(I)|_{\text{dnf}} \leq |I|_{\text{dnf}}$ (Lemma 5). \square

Remark 2 (Backwards fence condition). Our main theorem in this section, Thm. 5, also has a dual version that applies to a fence condition concerning backwards reachability. I is s -backwards fenced if every state in the outer boundary $\partial^-(I) \stackrel{\text{def}}{=} \partial^+(\neg I)$ can reach a state in Bad in at most s steps [13]. The dual of Thm. 5 is that there is an algorithm that achieves the same complexity bound under the assumption that I is s -backwards fenced (instead of s -forwards fenced). The dual- $CDNF$ algorithm is obtained by running our $CDNF$ algorithm on the dual transition system $(Bad, \delta^{-1}, Init)$ (see e.g. [12, Appendix A]) and negate the invariant; notice that the $CDNF$ class is closed under negation. This algorithm also achieves the same bound for decision trees as in Corollary 1, under the backwards fence assumption.

Remark 3 (Comparison to Bshouty’s $CDNF$ algorithm). Our $CDNF$ algorithm, Alg. 3, is inspired by Bshouty’s $CDNF$ algorithm [5], but diverges from it in several ways. The reason is the different queries available in each setting. (The code for Bshouty’s $CDNF$ algorithm is provided in the extended version [15].) Structurally, while the candidate in both algorithms is gradually constructed to be $\text{MHull}_{\mathcal{C}_i}(I) = \bigwedge_{\sigma \in \mathcal{C}_i} \mathcal{M}_\sigma(I)$ (I being the unknown invariant/formula, and \mathcal{C}_i the set of negative examples so far), Alg. 3 constructs each monotonization separately, one by one, whereas Bshouty’s algorithm increases all monotonizations

simultaneously. Bshouty’s design follows from having the source of examples—both positive and negative—equivalence queries, checking whether the candidate matches I . A membership query is necessary to decide whether the differentiating example is positive or negative for I in order to decide whether to add disjuncts to the existing monotizations or to add a new monotization, respectively. This procedure is problematic in invariant inference, because we cannot in general decide, for a counterexample (σ, σ') showing that our candidate is not inductive, whether $\sigma \not\models I$ (negative) or $\sigma' \models I$ (positive) [18, 12]. The solution in previous work [13] was to assume that the invariant satisfies both the forwards *and* backwards fence condition (see Remark 2). Under this assumption it is possible to decide whether $\sigma \models I$ for an arbitrary state σ . However, this condition is much stronger than a one-sided version of the fence condition. Instead, in our inference algorithm, the candidate is ensured to be an overapproximation of the true I , so each counterexample to induction in line 6 yields a negative example. Positive examples are obtained in line 4 from $\delta^s(\text{Init}) \subseteq I$; there is no obvious counterpart to that in exact learning, because in that setting we have no a-priori knowledge of some set S that underapproximates I , let alone one where we know—as the fence condition guarantees through Lemma 12—that covering S in the monotization is enough to cover I .

6 Efficient Implementation of Abstract Interpretation

In this section we build on the algorithm of §4 to prove a complexity upper bound on abstract interpretation in the domain based on the monotone theory (Thm. 7). We begin with background on this domain.

6.1 Background: Abstract Interpretation in the Monotone Theory

Recall that given a set of states B , the monotone span (Def. 6) of B , $\text{MSpan}(B)$, is the set of formulas φ s.t. $\text{MHull}_B(\varphi) \equiv \varphi$, or, equivalently, the set of formulas that can be written as conjunctions of clauses that exclude states from B (Thm. 1). The *abstract domain* $\mathbb{M}[B] = \langle \text{MSpan}(B), \implies, \sqcup_B, \text{false} \rangle$, introduced in [14], is a join-semilattice over the monotone span of B , ordered by logical implication, with bottom element *false*. The lub \sqcup_B exists because the domain is finite and closed under conjunction (follows from Thm. 1). A Galois connection $(2^{\text{States}[\Sigma]}, \subseteq) \xleftrightarrow[\alpha_B]{\gamma} (\text{MSpan}(B), \implies)$ with the concrete domain is obtained through the *concretization* $\gamma(\varphi) = \{\sigma \mid \sigma \models \varphi\}$ and the *abstraction* $\alpha_B(\psi) = \text{MHull}_B(\psi)$ [14].⁴

⁴ $\mathbb{M}[B]$ is parametrized by a choice of a monotone basis B . When B is large, the abstraction is more precise; it is precise enough to prove safety when there exists an inductive invariant that can be expressed in CNF such that each clause excludes at least one state from B (through Thm. 1). The fewer states B includes, the more extrapolation is performed in each abstraction step. However, since B also changes the available inductive invariants, the overall convergence might actually be faster

Given a transition system $(Init, \delta)$, iterations of abstract interpretation with the abstract transformer are given by $\xi_0 = \alpha_B(Init)$, $\xi_{i+1} = \alpha_B(\delta(\gamma(\xi_i)) \cup Init)$. Substituting γ, α_B yields the iterations as shown in Alg. 4.

Algorithm 4 Kleene Iterations in $\mathbb{M}[B]$

```

1: procedure AI- $\mathbb{M}[B](Init, \delta, Bad)$ 
2:    $i \leftarrow 0$ 
3:    $\xi_{-1} \leftarrow false$ 
4:    $\xi_0 \leftarrow MHull_B(Init)$ 
5:   while  $\xi_i \not\Rightarrow \xi_{i-1}$  do
6:      $\xi_{i+1} = MHull_B(\delta(\xi_i) \cup Init)$ 
7:      $i \leftarrow i + 1$ 
8:   return  $\xi_i$ 

```

Each iterate in Alg. 4 involves a monotone hull (lines 4 and 6), which is a conjunction of monotonicizations. Using Alg. 1 this can be computed efficiently. We follow on this idea to prove efficient complexity upper bounds on Alg. 4.

6.2 Complexity Upper Bound

To obtain a complexity upper bound on Alg. 4 we need to bound the time needed to compute each ξ_i as well as the number of ξ_i 's. A bound for the latter is provided by [14]:

Theorem 6 ([14]). *Let $(Init, \delta, Bad)$ be a transition system. Then the algorithm $AI\text{-}\mathbb{M}[B](Init, \delta, Bad)$ converges in iteration number at most*

$$\zeta \stackrel{\text{def}}{=} \prod_{i=1}^m \left(|\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta)|_{\text{dnf}} + |\mathcal{M}_{b_i}(Init)|_{\text{dnf}} \right),$$

where B can be written in DNF as $b_1 \vee \dots \vee b_m$, the cube \mathbb{B} consists of the literals that appear in all b_1, \dots, b_m (i.e., $\mathbb{B} = \bigcap_{i=1}^m b_i$ as sets of literals), and the reflection of a cube $d = \ell_1 \wedge \dots \wedge \ell_r$ is $Ref(d) = \neg \ell_1 \wedge \dots \wedge \neg \ell_r$.

We fix a DNF representation of $B = b_1 \vee \dots \vee b_m$. For brevity, we use ζ to refer to the bound in Thm. 6. When ζ is small, it reflects the benefit of using abstract interpretation in $\text{MSpan}(B)$ over exact reachability (even though ζ is not always a tight bound) [14]. An example of ζ for a simple system appears in Example 3.

In this section we prove that it is possible to implement Alg. 4 so that its overall complexity is polynomial in the same quantity ζ , the number of variables n , and the number m of terms in the representation of B :

with a larger (less extrapolating) B . Understanding how to choose B is an important direction for future work. (In [14], B was obtained from the states that reach a bad state in a fixed number of steps, mimicking PDR's scheme for generating proof obligations.)

Theorem 7. *Alg. 4 can be implemented to terminate in $O(n^2\zeta + (n+m)\zeta^2)$ SAT queries and time.*

Example 3. Let n be an odd number. Consider a transition system over $\bar{x} = x_1, \dots, x_n$, where $Init$ is $\bar{x} = 00\dots 00$ and the transition relation chooses an even number of variables that are 0 from the initial state and turns them into 1. If we take B to be the singleton set containing the state $\bar{x} = 11\dots 11$ (hence, B is a cube and $m = 1$), then Lemma 3 yields that $\mathcal{M}_{\bar{x}=00\dots 00 \wedge \bar{x}'=11\dots 11}(\delta) = \bigvee_{i=1}^n (x'_i = 0)$ (see Appendix C for details) so $\zeta = |\mathcal{M}_{\bar{x}=00\dots 00 \wedge \bar{x}'=11\dots 11}(\delta)|_{\text{dnf}} = O(n)$. Thm. 7 shows that an implementation of abstract interpretation in $\mathbb{M}[B]$ for this system terminates in $O(n^3)$ SAT queries and time. This is significant because a naive implementation of Alg. 4 would start, for the first iteration of line 6, by computing the exact post-image $\delta(Init)$; in our example this is the set of states where the parity of \bar{x} is 0, which cannot be represented in polynomial-size DNF nor CNF [e.g. 8]. Our implementation is able to compute the abstraction of the post-image without constructing the post-image and avoids the blowup in complexity. [14] contains other examples with small ζ .

At this point, the direct approach to implement Alg. 4 is to perform $\text{MHull}_B(\varphi)$ in lines 4 and 6 through $\bigwedge_{j=1}^m \text{MONOTONIZE}(\varphi, b_j)$, invoking Alg. 1 on φ . Indeed, this achieves a bound that is only slightly worse than Thm. 7 (see Remark 4). In what follows we provide an implementation that both explicates the connection to ζ , and achieves exactly the bound of Thm. 7.

Algorithm 5 Efficient Kleene Iterations in $\mathbb{M}[B]$

```

1: procedure AI- $\mathbb{M}[B](Init, \delta, Bad)$ 
2:    $i \leftarrow 0$ 
3:    $\xi_{-1} \leftarrow false$ 
4:    $\xi_0 \leftarrow \bigwedge_{j=1}^m \text{MONOTONIZE}(Init, b_j)$ 
5:   for  $j = 1..m$  do
6:      $\delta_j^* \leftarrow \text{MONOTONIZE}(\delta \vee Init', Ref(\mathbb{B}) \wedge b'_j)$ 
7:   while  $\xi_i \not\Rightarrow \xi_{i-1}$  do
8:      $\xi_{i+1} = \bigvee \{(t_1|_{\Sigma'}) \wedge \dots \wedge (t_m|_{\Sigma'}) \mid t_j \text{ a term of } \delta_j^*, \exists \sigma_j \in \xi_i. \sigma_j \models (t_j|_{\Sigma'})\}$ 
9:      $i \leftarrow i + 1$ 
10:  return  $\xi_i$ 

```

Our implementation is displayed in Alg. 5. The first iterate is computed as described above by invoking Alg. 1 on $Init$ (line 4). The SAT queries performed by Alg. 1 are in this case straightforward, with $\varphi = Init$.

To compute the next iterates, we first compute monotizations of the concrete transformer, $\delta \vee Init'$ (line 6). This is a two-vocabulary formula, and accordingly the monotizations are w.r.t. two-vocabulary cubes. The monotizations are computed in DNF form and stored in δ_j^* . The next iterate ξ_{i+1} is formed from the δ_j^* 's by taking all the combinations of terms from $\delta_1^*, \dots, \delta_m^*$ whose pre-state

part is satisfied by at least one state in ξ_i , and forming the conjunction of the post-state parts: for a term $t = \ell_1 \vee \dots \vee \ell_{i_1} \vee \ell'_{i_1+1} \vee \dots \vee \ell'_{i_2}$ over $\Sigma \uplus \Sigma'$, the restriction $t|_{\Sigma} = \ell_1 \vee \dots \vee \ell_{i_1}$ and $t|_{\Sigma'} = \ell'_{i_1+1} \vee \dots \vee \ell'_{i_2}$.

The intuition is that in the original algorithm, given a set of states ξ_i , we find the set of states in ξ_{i+1} by taking the result of the transformer $\delta \vee \text{Init}'$ on the specific ξ_i , then, for the monotonization of the result, adding also the states that are required by the b_j -monotone order, and this we do for every disjunct b_j in B , letting ξ_{i+1} be the conjunction of the said monotonizations. Here, instead, we monotinize $\delta \vee \text{Init}'$ itself w.r.t. every b_j , such that for every pre-state we have ready the monotonization of the corresponding post-state. We then form ξ_{i+1} by picking and conjoining the right monotonizations—the ones whose pre-state is in the previous frame. (The monotonization of the pre-state w.r.t. $\text{Ref}(\mathbb{B})$ is useful for decreasing ζ , and hence the obtained bound, without altering ξ_{i+1} ; the latter stems from the fact that the input ξ_i is also the result of such a procedure, so the presence of a pre-state in ξ_i indicates the presence all the states in its monotone hull w.r.t. $\text{Ref}(\mathbb{B})$ in ξ_i .)

The invocation of Alg. 1 in line 6 is on a double-vocabulary formula; still, the SAT queries to be performed in the invocation of Alg. 1 are simple SAT queries about two-vocabulary formulas (and a counterexample is a pair of states).

It is important for the efficiency result that Alg. 3 uses Alg. 1 as a subprocedure. Using Bshouty's procedure (see Remark 1) would yield a bound in terms of the DNF size of the original transition relation, which could be significantly larger, especially in cases where the abstract interpretation procedure terminates faster than exact forward reachability [14].

The rest of this section proves that Alg. 5 realizes Thm. 7. To show that it correctly implements Alg. 4, we need the following fact about Alg. 4:

Lemma 16 ([14]). *In Alg. 4, $\sigma' \models \xi_{i+1}$ iff there exist $\sigma_1, \dots, \sigma_m \models \xi_i$ s.t.*

$$(\sigma_1, \sigma') \models \mathcal{M}_{\text{Ref}(\mathbb{B}) \wedge b'_1}(\delta \vee \text{Init}') \wedge \dots \wedge (\sigma_m, \sigma') \models \mathcal{M}_{\text{Ref}(\mathbb{B}) \wedge b'_m}(\delta \vee \text{Init}'). \quad (1)$$

We use this to show the correctness of Alg. 5:

Lemma 17. *ξ_i in Alg. 5 is logically equivalent to ξ_i in Alg. 4.*

Proof. By induction over i . The correctness of ξ_0 follows from the correctness of Alg. 1 (Thm. 2). For the same reasons, δ_j^* of Alg. 5 is equivalent to

$\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_j}(\delta \vee Init')$. Now for some DNF manipulation: for every σ' ,

$$\begin{aligned}
 & \exists \sigma_1, \dots, \sigma_m. (\sigma_1, \sigma') \models \delta_1^* \wedge \dots \wedge (\sigma_m, \sigma') \models \delta_m^* \\
 \iff & \exists \sigma_1, \dots, \sigma_m. \exists t_1 \text{ term of } \delta_1^*, \dots, \exists t_m \text{ term of } \delta_m^*. \\
 & (\sigma_1, \sigma') \models t_1 \wedge \dots \wedge (\sigma_m, \sigma') \models t_m \\
 \iff & \exists \sigma_1, \dots, \sigma_m. \exists t_1 \text{ term of } \delta_1^*, \dots, \exists t_m \text{ term of } \delta_m^*. \\
 & \sigma_1 \models (t_1|_{\Sigma}) \wedge \sigma' \models (t_1|_{\Sigma'}) \wedge \dots \wedge \sigma_1 \models (t_m|_{\Sigma}) \wedge \sigma' \models (t_m|_{\Sigma'}) \\
 \iff & \exists \sigma_1, \exists t_1 \text{ term of } \delta_1^*. \sigma_1 \models (t_m|_{\Sigma}) \wedge \sigma' \models (t_m|_{\Sigma'}) \\
 & \wedge \dots \wedge \\
 & \exists \sigma_m, \exists t_m \text{ term of } \delta_m^*. \sigma_m \models (t_m|_{\Sigma}) \wedge \sigma' \models (t_m|_{\Sigma'}).
 \end{aligned}$$

Hence, $\exists \sigma_1, \dots, \sigma_m \in \xi_i$ of Alg. 5 that with σ' satisfy Eq. (1) iff $\sigma' \in \xi_{i+1}$ of Alg. 5. Lemma 16 and the induction hypothesis complete the proof. \square

We can now proceed to prove the complexity bound for Alg. 5.

Lemma 18. *Alg. 5 terminates in $O(n^2\zeta + (n+m)\zeta^2)$ SAT queries and time.*

Proof. By Thm. 2, each invocation of Alg. 1 in line 4 takes $O(n^2 \cdot |\mathcal{M}_{b_i}(Init)|_{\text{dnf}}) = O(n^2\zeta)$ queries and time. Similarly, each invocation in line 6 takes

$$O\left((2n)^2 \cdot |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta \vee Init')|_{\text{dnf}}\right)$$

queries and time. This quantity is $O(n^2\zeta)$, because by Lemma 3

$$\begin{aligned}
 |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta \vee Init')|_{\text{dnf}} & \leq |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta)|_{\text{dnf}} + |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(Init')|_{\text{dnf}} \\
 & = |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta)|_{\text{dnf}} + |\mathcal{M}_{b_i}(Init)|_{\text{dnf}}.
 \end{aligned}$$

In each iteration, the number of combinations of terms in line 8 is at most $\prod_{i=1}^m |\mathcal{M}_{Ref(\mathbb{B}) \wedge b'_i}(\delta \vee Init')|_{\text{dnf}}$. For each of the m terms in the combination, we split the term to Σ, Σ' parts in time linear term size which is at most n , and perform a SAT check for whether the term intersects ξ_i . Overall this step involves $O(m \cdot \zeta)$ queries and $O((n+m) \cdot \zeta)$ time. This is the cost of each iteration; the number of iterations is bounded by ζ by Thm. 6. The claim follows. \square

The proof of Thm. 7 follows from Lemmas 17 and 18.

Remark 4. Lemmas 17 and 18 have the consequence that in Alg. 4, $|\xi_i|_{\text{dnf}} \leq \zeta$ (interestingly, this is true in particular for the resulting inductive invariant). This is a proof that the direct implementation of the monotone hull by m calls to Alg. 1 amounts to $O(n^2m\zeta)$ SAT queries in each iteration, and $O(n^2m\zeta^2)$ time thanks to Thm. 6. Though asymptotically inferior, this implementation approach may be more efficient than Alg. 5 when $|\xi_i|_{\text{dnf}} \ll \zeta$.

7 Related Work

Complexity bounds for invariant inference. Conjunctive/disjunctive invariants can be inferred in a linear number of SAT calls [16, 25]. On the other hand, inferring CNF/DNF invariants for general transition systems is **NP**-hard with access to a SAT solver [25], even when the invariants are restricted to monotone formulas [12]. Complexity results for model-based interpolation-based invariant inference (stemming from the analysis of the algorithm by [6, 3]) were presented in [13], based on the (backwards) fence condition, which tames reachability enough to efficiently infer monotone and almost-monotone DNF invariants, or monotone and almost-monotone CNF invariants under the (forwards) fence condition with the dual algorithms. Our algorithm (Alg. 5) can achieve the same bounds if, similar to [13], the monotone basis (set of counterexamples) is fixed in advance. (This alteration is needed because some short monotone DNF formulas have large CNF size [27].) Our algorithm is more versatile because it applies to CDNF formulas which are not monotone or almost-monotone, and alleviates the need to know a monotone basis in advance. A CDNF complexity bound similar to ours was obtained in [13] under the much stronger assumption that both the backwards and the forwards fence condition hold simultaneously (see Remark 2). Property-directed reachability algorithms [4, 9] were shown efficient on several parametrized examples [34, 14] and the very special case of maximal transition systems for monotone invariants [12]. We show (§6) that abstract interpretation in the monotone theory, studied under the name \mathcal{A} -PDR [14] in relation to standard PDR, is efficient in broader circumstances, when the DNF size of certain monotoneizations of the transition relation are small, using the same quantity ζ that in previous work [14] was established as an upper bound on the number of iterations (without an overall complexity result).

Monotone theory in invariant inference. The monotone theory has been employed in previous works on invariant inference. The aforementioned previous results on inference under the fence condition [13] also employ the monotone theory, and are based on one-to-one translations of Bshouty’s algorithms to invariant inference, replacing equivalence queries by inductiveness checks and membership queries by bounded model checking. For Bshouty’s CDNF algorithm this only works under the stronger two-sided fence condition, which is why our algorithm differs significantly (see Remark 3). The one-sided fence condition suffices for the translation of Bshouty’s \mathcal{A} -algorithm, which is suitable when a monotone basis is known in advance, e.g. for almost-monotone invariants, whereas the CDNF algorithm learns a monotone basis on-the-fly. Another translation of Bshouty’s algorithm to invariant inference is by Jung et al. [23], who combine the CDNF algorithm of Bshouty [5] with predicate abstraction and templates to infer quantified invariants. They overcome the problem of membership queries in the original algorithm (see Remark 3) heuristically, using under- and over-approximations and sometimes random guesses, which could lead to the need to restart. The monotone theory is also used in a non-algorithmic way in [14] to

analyze overapproximation in IC3/PDR through abstract interpretation in the monotone theory, to which we prove a complexity upper bound.

Inferring decision tree invariants. In machine learning, decision trees are a popular representation of hypotheses and target concepts. Garg et al. [19] adapt an algorithm by Quinlan [31] to infer invariants (later, general Horn clauses [11]) in the form of decision trees over numerical and Boolean attributes, which is guaranteed to converge, but not necessarily efficiently overall (even though the algorithm efficiently generates the candidates, the number of candidates could be large). Similarly to our algorithm, the translation of the CDNF algorithm to invariant inference in [13] is applicable also to Boolean decision trees, but, as previously mentioned, requires the stronger two-sided fence condition, whereas our result is the first to do so under the (one-sided) fence condition.

Complexity of abstract interpretation. The efficiency of the abstract transformers is crucial to the overall success of abstract interpretation, which is often at odds with the domain accuracy; a famous example is the octagon abstract domain [28], whose motivation is the prohibitive cost of the expressive polyhedra domain [7]. We provide a way to compute abstract transformers in the monotone span domain that is efficient in terms of the DNF size of the result (see also Remark 4). The computation of the abstract transformer in Alg. 5 is inspired by works in symbolic abstraction [32, 40] about finding *representations* of the best abstract transformer, rather than computing it anew per input [10, 39, 33].

8 Conclusion

This work has accomplished invariant inference algorithms with efficient complexity guarantees in two settings—model-based interpolation and property-directed reachability—resolving open problems where the missing component (as it turns out) was a new way to compute monotone overapproximations. A common theme is the use of rich syntactic forms of the formulas that the algorithms maintain: in our model-based interpolation algorithm, the candidate invariant is a conjunction of DNFs, even though the target invariant has both short CNF and DNF representations. In our efficient implementation of abstract interpretation, each iterate is again a conjunction of DNFs, although the natural definition (Thm. 1, and as inspired by PDR) is with CNF formulas. We hope that these ideas could inspire new interpolation-based and property-directed reachability algorithms that would benefit in practice from richer hypotheses and techniques from the monotone theory.

Acknowledgements We thank the anonymous reviewers and Hila Peleg for insightful comments. The research leading to these results has received funding from the European Research Council under the European Union’s Horizon 2020 research and innovation programme (grant agreement No [759102-SVIS]). This research was partially supported by the Israeli Science Foundation (ISF) grant No. 1810/18.

Bibliography

- [1] Angluin, D.: Queries and concept learning. *Machine Learning* **2**(4), 319–342 (1987)
- [2] Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic model checking without bdds. In: *Tools and Algorithms for Construction and Analysis of Systems, 5th International Conference, TACAS '99, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'99, Amsterdam, The Netherlands, March 22-28, 1999, Proceedings*. pp. 193–207 (1999). https://doi.org/10.1007/3-540-49059-0_14, https://doi.org/10.1007/3-540-49059-0_14
- [3] Bjørner, N., Gurfinkel, A., Korovin, K., Lahav, O.: Instantiations, zippers and EPR interpolation. In: *LPAR 2013, 19th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, December 12-17, 2013, Stellenbosch, South Africa, Short papers proceedings*. pp. 35–41 (2013), <https://easychair.org/publications/paper/XtN>
- [4] Bradley, A.R.: Sat-based model checking without unrolling. In: *Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings*. pp. 70–87 (2011). https://doi.org/10.1007/978-3-642-18275-4_7, http://dx.doi.org/10.1007/978-3-642-18275-4_7
- [5] Bshouty, N.H.: Exact learning boolean function via the monotone theory. *Inf. Comput.* **123**(1), 146–153 (1995). <https://doi.org/10.1006/inco.1995.1164>, <https://doi.org/10.1006/inco.1995.1164>
- [6] Chockler, H., Ivrii, A., Matsliah, A.: Computing interpolants without proofs. In: *Hardware and Software: Verification and Testing - 8th International Haifa Verification Conference, HVC 2012, Haifa, Israel, November 6-8, 2012. Revised Selected Papers*. pp. 72–85 (2012). https://doi.org/10.1007/978-3-642-39611-3_12, https://doi.org/10.1007/978-3-642-39611-3_12
- [7] Cousot, P., Halbwachs, N.: Automatic discovery of linear restraints among variables of a program. In: Aho, A.V., Zilles, S.N., Szymanski, T.G. (eds.) *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages*, Tucson, Arizona, USA, January 1978. pp. 84–96. ACM Press (1978). <https://doi.org/10.1145/512760.512770>, <https://doi.org/10.1145/512760.512770>
- [8] Crama, Y., Hammer, P.L.: *Boolean Functions - Theory, Algorithms, and Applications*, Encyclopedia of mathematics and its applications, vol. 142. Cambridge University Press (2011), http://www.cambridge.org/gb/knowledge/isbn/item6222210/?site_locale=en_GB
- [9] Eén, N., Mishchenko, A., Brayton, R.K.: Efficient implementation of property directed reachability. In: *International Conference on Formal Methods in Computer-Aided Design, FMCAD '11, Austin, TX, USA, October 30 - November 02, 2011*. pp. 125–134 (2011), <http://dl.acm.org/citation.cfm?id=2157675>

- [10] Elder, M., Lim, J., Sharma, T., Andersen, T., Reps, T.W.: Abstract domains of affine relations. In: Yahav, E. (ed.) *Static Analysis - 18th International Symposium, SAS 2011, Venice, Italy, September 14-16, 2011. Proceedings*. Lecture Notes in Computer Science, vol. 6887, pp. 198–215. Springer (2011). https://doi.org/10.1007/978-3-642-23702-7_17, https://doi.org/10.1007/978-3-642-23702-7_17
- [11] Ezudheen, P., Neider, D., D’Souza, D., Garg, P., Madhusudan, P.: Horn-ice learning for synthesizing invariants and contracts. *PACMPL* **2**(OOPSLA), 131:1–131:25 (2018)
- [12] Feldman, Y.M.Y., Immerman, N., Sagiv, M., Shoham, S.: Complexity and information in invariant inference. *Proc. ACM Program. Lang.* **4**(POPL), 5:1–5:29 (2020). <https://doi.org/10.1145/3371073>, <https://doi.org/10.1145/3371073>
- [13] Feldman, Y.M.Y., Sagiv, M., Shoham, S., Wilcox, J.R.: Learning the boundary of inductive invariants. *Proc. ACM Program. Lang.* **5**(POPL), 1–30 (2021). <https://doi.org/10.1145/3434296>, <https://doi.org/10.1145/3434296>
- [14] Feldman, Y.M.Y., Sagiv, M., Shoham, S., Wilcox, J.R.: Property-directed reachability as abstract interpretation in the monotone theory. *Proc. ACM Program. Lang.* **6**(POPL), 1–31 (2022). <https://doi.org/10.1145/3498676>, <https://doi.org/10.1145/3498676>
- [15] Feldman, Y.M.Y., Shoham, S.: Invariant inference with provable complexity from the monotone theory. *CoRR* (2022), <https://arxiv.org/pdf/2208.07451.pdf>
- [16] Flanagan, C., Leino, K.R.M.: Houdini, an annotation assistant for `esc/java`. In: *FME 2001: Formal Methods for Increasing Software Productivity, International Symposium of Formal Methods Europe, Berlin, Germany, March 12-16, 2001, Proceedings*. pp. 500–517 (2001)
- [17] Flanagan, C., Qadeer, S.: Predicate abstraction for software verification. In: *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, OR, USA, January 16-18, 2002*. pp. 191–202 (2002). <https://doi.org/10.1145/503272.503291>, <http://doi.acm.org/10.1145/503272.503291>
- [18] Garg, P., Löding, C., Madhusudan, P., Neider, D.: Ice: A robust framework for learning invariants. In: *Computer Aided Verification*. pp. 69–87. Springer (2014)
- [19] Garg, P., Neider, D., Madhusudan, P., Roth, D.: Learning invariants using decision trees and implication counterexamples. In: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*. pp. 499–512 (2016). <https://doi.org/10.1145/2837614.2837664>, <https://doi.org/10.1145/2837614.2837664>
- [20] Graf, S., Saïdi, H.: Construction of abstract state graphs with PVS. In: *Computer Aided Verification, 9th International Conference, CAV ’97, Haifa, Israel, June 22-25, 1997, Proceedings*. pp. 72–83 (1997). https://doi.org/10.1007/3-540-63166-6_10, https://doi.org/10.1007/3-540-63166-6_10

- [21] Jha, S., Gulwani, S., Seshia, S.A., Tiwari, A.: Oracle-guided component-based program synthesis. In: Proceedings of the 32nd ACM/IEEE International Conference on Software Engineering - Volume 1, ICSE 2010, Cape Town, South Africa, 1-8 May 2010. pp. 215–224 (2010). <https://doi.org/10.1145/1806799.1806833>, <https://doi.org/10.1145/1806799.1806833>
- [22] Jha, S., Seshia, S.A.: A theory of formal synthesis via inductive learning. *Acta Inf.* **54**(7), 693–726 (2017). <https://doi.org/10.1007/s00236-017-0294-5>, <https://doi.org/10.1007/s00236-017-0294-5>
- [23] Jung, Y., Kong, S., David, C., Wang, B., Yi, K.: Automatically inferring loop invariants via algorithmic learning. *Math. Struct. Comput. Sci.* **25**(4), 892–915 (2015). <https://doi.org/10.1017/S0960129513000078>, <https://doi.org/10.1017/S0960129513000078>
- [24] Koenig, J.R., Padon, O., Immerman, N., Aiken, A.: First-order quantified separators. In: Donaldson, A.F., Torlak, E. (eds.) Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020. pp. 703–717. ACM (2020), <https://doi.org/10.1145/3385412.3386018>
- [25] Lahiri, S.K., Qadeer, S.: Complexity and algorithms for monomial and clausal predicate abstraction. In: Automated Deduction - CADE-22, 22nd International Conference on Automated Deduction, Montreal, Canada, August 2-7, 2009. Proceedings. pp. 214–229 (2009)
- [26] McMillan, K.L.: Interpolation and sat-based model checking. In: Computer Aided Verification, 15th International Conference, CAV 2003, Boulder, CO, USA, July 8-12, 2003, Proceedings. pp. 1–13 (2003)
- [27] Miltersen, P.B., Radhakrishnan, J., Wegener, I.: On converting cnf to dnf. *Theoretical Computer Science* **347**(1), 325–335 (2005). <https://doi.org/10.1016/j.tcs.2005.07.029>, <https://www.sciencedirect.com/science/article/pii/S0304397505004688>
- [28] Miné, A.: The octagon abstract domain. *High. Order Symb. Comput.* **19**(1), 31–100 (2006). <https://doi.org/10.1007/s10990-006-8609-1>, <https://doi.org/10.1007/s10990-006-8609-1>
- [29] O’Rourke, J.: Visibility. In: Goodman, J.E., O’Rourke, J. (eds.) Handbook of Discrete and Computational Geometry, Second Edition, pp. 643–663. Chapman and Hall/CRC (2004). <https://doi.org/10.1201/9781420035315.ch28>, <https://doi.org/10.1201/9781420035315.ch28>
- [30] Quine, W.: Two theorems about truth-functions. *Boletín de la Sociedad Matemática Mexicana* **10**(1–2), 64–70 (1954)
- [31] Quinlan, J.R.: Induction of decision trees. *Mach. Learn.* **1**(1), 81–106 (1986). <https://doi.org/10.1023/A:1022643204877>, <https://doi.org/10.1023/A:1022643204877>
- [32] Reps, T.W., Sagiv, S., Yorsh, G.: Symbolic implementation of the best transformer. In: Verification, Model Checking, and Abstract Interpretation, 5th International Conference, VMCAI 2004, Venice, Italy, January 11-13, 2004, Proceedings. pp. 252–266 (2004). https://doi.org/10.1007/978-3-540-24622-0_21, https://doi.org/10.1007/978-3-540-24622-0_21

- [33] Reps, T.W., Thakur, A.V.: Automating abstract interpretation. In: Jobstmann, B., Leino, K.R.M. (eds.) *Verification, Model Checking, and Abstract Interpretation - 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17-19, 2016*. Proceedings. Lecture Notes in Computer Science, vol. 9583, pp. 3–40. Springer (2016). https://doi.org/10.1007/978-3-662-49122-5_1, https://doi.org/10.1007/978-3-662-49122-5_1
- [34] Seufert, T., Scholl, C.: Sequential verification using reverse PDR. In: Große, D., Drechsler, R. (eds.) *Methoden und Beschreibungssprachen zur Modellierung und Verifikation von Schaltungen und Systemen, MBMV 2017, Bremen, Germany, February 8-9, 2017*. pp. 79–90. Shaker Verlag (2017)
- [35] Sharma, R., Aiken, A.: From invariant checking to invariant inference using randomized search. *Formal Methods in System Design* **48**(3), 235–256 (2016). <https://doi.org/10.1007/s10703-016-0248-5>, <https://doi.org/10.1007/s10703-016-0248-5>
- [36] Sharma, R., Gupta, S., Hariharan, B., Aiken, A., Liang, P., Nori, A.V.: A data driven approach for algebraic loop invariants. In: *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013*. Proceedings. pp. 574–592 (2013). https://doi.org/10.1007/978-3-642-37036-6_31, https://doi.org/10.1007/978-3-642-37036-6_31
- [37] Sharma, R., Gupta, S., Hariharan, B., Aiken, A., Nori, A.V.: Verification as learning geometric concepts. In: *Static Analysis - 20th International Symposium, SAS 2013, Seattle, WA, USA, June 20-22, 2013*. Proceedings. pp. 388–411 (2013)
- [38] Sharma, R., Nori, A.V., Aiken, A.: Interpolants as classifiers. In: *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012* Proceedings. pp. 71–87 (2012). https://doi.org/10.1007/978-3-642-31424-7_11, https://doi.org/10.1007/978-3-642-31424-7_11
- [39] Thakur, A.V., Elder, M., Reps, T.W.: Bilateral algorithms for symbolic abstraction. In: Miné, A., Schmidt, D. (eds.) *Static Analysis - 19th International Symposium, SAS 2012, Deauville, France, September 11-13, 2012*. Proceedings. Lecture Notes in Computer Science, vol. 7460, pp. 111–128. Springer (2012). https://doi.org/10.1007/978-3-642-33125-1_10, https://doi.org/10.1007/978-3-642-33125-1_10
- [40] Thakur, A.V., Lal, A., Lim, J., Reps, T.W.: Postthat and all that: Automating abstract interpretation. *Electr. Notes Theor. Comput. Sci.* **311**, 15–32 (2015). <https://doi.org/10.1016/j.entcs.2015.02.003>, <https://doi.org/10.1016/j.entcs.2015.02.003>
- [41] Valiant, L.G.: A theory of the learnable. *Commun. ACM* **27**(11), 1134–1142 (1984). <https://doi.org/10.1145/1968.1972>, <https://doi.org/10.1145/1968.1972>
- [42] Wiedemann, D.H.: Hamming geometry. Ph.D. thesis, University of Waterloo (1987)

A Missing Proofs

Proof (Proof of Lemma 1). That $\mathcal{M}_b(\varphi)$ is b -monotone overapproximation of φ is immediate from the definition. For minimality, let ψ be a b -monotone formula s.t. $\varphi \implies \psi$, we need to show that $\mathcal{M}_b(\varphi) \implies \psi$. Let $x \models \mathcal{M}_b(\varphi)$. Then, by definition, there is $v \models \varphi$ s.t. $v \leq_b x$. By the assumption that $\varphi \implies \psi$ also $v \models \psi$, and, because ψ is assumed to be b -monotone, it follows that $x \models \psi$. The claim follows. \square

Proof (Proof of Lemma 4). From the premise it follows that $\text{cube}_b(v) \implies \mathcal{M}_b(\mathcal{M}_b(\varphi))$, but by Lemma 3 and conversions to DNF, $\mathcal{M}_b(\mathcal{M}_b(\varphi)) \equiv \mathcal{M}_b(\varphi)$. \square

Proof (Proof of Corollary 1). A decision tree of size m has a DNF representation of m terms: a disjunction of terms representing the paths that reach a leaf *true*, each is the conjunction of the variables on the path with polarity according to left/right branch. Similarly, it has a CNF representation of m clauses: a conjunction of clauses which are the negations of paths that reach a leaf *false*. Now apply Thm. 5. \square

B Bshouty's CDNF Algorithm

Algorithm 6 CDNF-algorithm: exact learning CDNF formulas [5]

<pre> 1: procedure CDNF-ALGORITHM 2: $t \leftarrow 0$ 3: while $x \leftarrow \text{EQ}(\bigwedge_{i=1}^t H_i)$ is not \perp 4: do 5: if $x \models \bigwedge_{i=1}^t H_i$ then 6: $H_{t+1} \leftarrow \text{false}; b_{t+1} \leftarrow x$ 7: $t \leftarrow t + 1$ 8: else 9: $\sigma' \leftarrow x$ 10: for $i = 1, \dots, t$ do 11: if $\sigma' \not\models H_i$ then 12: $d \leftarrow \text{GENMQ}(\sigma', b_i)$ 13: $H_i \leftarrow H_i \vee d$ 13: return $\bigwedge_{i=1}^t H_i$ </pre>	<pre> 14: procedure GENMQ(σ, b) 15: $v \leftarrow \sigma$ 16: $\text{walked} \leftarrow \text{true}$ 17: while walked do 18: $\text{walked} \leftarrow \text{false}$ 19: for $j = 1, \dots, n$ do 20: if $b[p_j] = v[p_j]$ then 21: continue 22: $x \leftarrow v[p_j \mapsto b[p_j]]$ 23: if $\text{MQ}(x) = \text{true}$ then 24: $v \leftarrow x$ 25: $\text{walked} \leftarrow \text{true}$ 26: return $\text{cube}_b(v)$ </pre>
---	--

The original CDNF algorithm is presented in Alg. 6, aiming to identify an unknown formula ψ . $\text{EQ}(H)$ returns \perp if $H \equiv \psi$ and a differentiating counterexample otherwise. $\text{MQ}(x)$ returns *true* iff $x \models \psi$.

C Calculation of ζ in Example 3

In Example 3, δ can be written in DNF as

$$\begin{array}{l}
 \delta = (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 00 \dots 000) \\
 \hline
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 00 \dots 011) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 00 \dots 101) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 01 \dots 001) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 10 \dots 001) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 00 \dots 110) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 01 \dots 010) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 10 \dots 010) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 01 \dots 100) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 10 \dots 100) \\
 \vee (\bar{x} = 00 \dots 000 \wedge \bar{x}' = 11 \dots 000) \\
 \hline
 \vee \dots \\
 \hline
 \vee \boxed{(\bar{x} = 00 \dots 000 \wedge \bar{x}' = 11 \dots 110)} \\
 \vee \boxed{(\bar{x} = 00 \dots 000 \wedge \bar{x}' = 11 \dots 101)} \\
 \vee \boxed{\dots} \\
 \vee \boxed{(\bar{x} = 00 \dots 000 \wedge \bar{x}' = 11 \dots 011)} \\
 \vee \boxed{(\bar{x} = 00 \dots 000 \wedge \bar{x}' = 10 \dots 111)} \\
 \vee \boxed{(\bar{x} = 00 \dots 000 \wedge \bar{x}' = 01 \dots 111)}.
 \end{array}$$

The horizontal lines groups disjuncts according to the number of variables that are turned to 1—zero, two, etc., up to $n - 1$. The literals that appear in color are those that are dropped when computing a DNF for $\mathcal{M}_{\bar{x}=00\dots 000 \wedge \bar{x}'=11\dots 111}(\delta)$ according to Lemma 3. In this way we see that the monotonization of the last group of terms, indicated by boxes, subsumes the monotonization of the other terms, yielding

$$\mathcal{M}_{\bar{x}=00\dots 000 \wedge \bar{x}'=11\dots 111}(\delta) \equiv \bigvee_{i=1}^n (x'_i = 0),$$

and hence $\zeta = |\mathcal{M}_{\bar{x}=00\dots 000 \wedge \bar{x}'=11\dots 111}(\delta)|_{\text{dnf}} \leq n$.