

A Hierarchy of Proof Rules for Checking Differential Invariance of Algebraic Sets^{*}

Khalil Ghorbal¹, Andrew Sogokon², and André Platzer¹

¹ Carnegie Mellon University, Computer Science Department, Pittsburgh, PA, USA,
{kghorbal|aplatzer}@cs.cmu.edu

² University of Edinburgh, LFCS, School of Informatics, Edinburgh, Scotland, UK,
a.sogokon@sms.ed.ac.uk

Abstract This paper presents a theoretical and experimental comparison of sound proof rules for proving invariance of algebraic sets, that is, sets satisfying polynomial equalities, under the flow of polynomial ordinary differential equations. Problems of this nature arise in formal verification of continuous and hybrid dynamical systems, where there is an increasing need for methods to expedite formal proofs. We study the trade-off between proof rule generality and practical performance and evaluate our theoretical observations on a set of heterogeneous benchmarks. The relationship between increased deductive power and running time performance of the proof rules is far from obvious; we discuss and illustrate certain classes of problems where this relationship is interesting.

1 Introduction

In safety verification of dynamical systems, either purely continuous or hybrid [22,29], one is typically concerned with ensuring that by initializing a system in some set of states $X_0 \subseteq X$ (where X is the state space), the system will never evolve into an unsafe state (belonging to some $X_u \subseteq X$). When the system is given by ordinary differential equations $\dot{x} = p(x)$, one may attempt to solve this problem by showing that the solution to the initial value problem, for any initial value $x_0 \in X_0$, cannot enter the unsafe region; that is, $x(x_0, t) \notin X_u$ for all $t \geq 0$, where $x(x_0, t)$ is the state of the system at time t w.r.t. the initial value x_0 . This safety verification problem is equivalent to showing that the intersection of the reachable set $\{x(x_0, t) \in X \mid t \geq 0\}$ with the set of unsafe states is empty. However, solutions to ordinary differential equations will rarely be available in closed form; and even when they are, will often be much more complicated than the differential equations themselves. Instead, it is possible to work with the differential equations *directly* [26,21,23,29].

A fundamental notion in safety verification is that of an *invariant set*. In fact, exact reachable sets of any given state x_0 of the system are the *smallest* invariant sets one can hope to find that include x_0 . However, obtaining and working with exact descriptions of

^{*} This material is based upon work supported by the National Science Foundation (NSF) under NSF CAREER Award CNS-1054246, NSF EXPEDITION CNS-0926181, NSF CNS-0931985, by DARPA under agreement number FA8750-12-2-029, as well as the Engineering and Physical Sciences Research Council (UK) under grant EP/I010335/1.

reachable sets is not always practical or even possible. This does not mean that system safety cannot be established by other means - if one finds a *larger* invariant set, $I \subseteq X$, with a simpler (perhaps algebraic) description which contains the reachable set and does not itself intersect the set of unsafe states (i.e. $I \cap X_u = \emptyset$), then one can soundly conclude that the system is safe. In this paper, we focus on checking whether a given set is an *invariant region* from which no system trajectory can escape.

Hybrid systems verification completely reduces to questions about invariant regions [20,22]. We focus on the important case where the invariant regions are algebraic sets, i.e. can be defined by polynomial equations. Many sound proof rules already exist for deciding invariance properties of algebraic sets. However, in order to identify a good trade-off, it is crucial to study the relationship between the deductive power and the practical running time performance of these proof rules.

Contributions. (I) We theoretically compare the deductive power of 7 different proof rules for checking invariance properties of algebraic sets under the flow of polynomial ordinary differential equations. Further, we assess the practical utility of each of these rules in order to identify a good trade-off between generality and running time performance. (II) We investigate the effect of *square-free reduction* on both the deductive power and the computational complexity of the proof rules. (III) We assess the practical proof rule performance on a heterogeneous set of 75 benchmarks. We demonstrate the counter-intuitive fact that square-free reduction does not necessarily improve the computational efficiency of certain proof rules and explore interesting connections between the deductive power and the practical running time performance that we observe for the proof rules.

Content. In Section 2, we recall some basic definitions and concepts that will be used throughout the paper. We then introduce (in Section 3) two proof rules that serve as extensions of Lie’s criterion for equational invariants. In Section 4, we compare the deductive power of the proof rules. The benefits and drawbacks of performing square-free reduction as a pre-processing step are investigated in Section 5. In Section 6, we present the set of benchmarks and our experimental results. We finally discuss other related work in Section 7 before concluding. All proofs, as well as more detailed results from running our benchmarks, can be found in the companion technical report [10].

2 Preliminaries

We consider autonomous³ polynomial vector fields (see Def. 1 below).

Let $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{R}^n$, and $\mathbf{x}(t) = (x_1(t), \dots, x_n(t))$, where $x_i : \mathbb{R} \rightarrow \mathbb{R}$, $t \mapsto x_i(t)$. The ring of polynomials over the reals will be denoted by $\mathbb{R}[x_1, \dots, x_n]$.

Definition 1 (Polynomial Vector Field). Let p_i , $1 \leq i \leq n$, be multivariate polynomials of the polynomial ring $\mathbb{R}[\mathbf{x}]$. A polynomial vector field, \mathbf{p} , is an explicit system of ordinary differential equations with polynomial right-hand side:

$$\frac{dx_i}{dt} = \dot{x}_i = p_i(\mathbf{x}), \quad 1 \leq i \leq n . \quad (1)$$

³ That is, the rate of change of the system over time depends only on the system’s state, not on time. Non-autonomous systems with polynomial time-dependence can be made autonomous by adding an extra clock variable that reflects the progress of time.

Since polynomial functions are smooth (C^∞ , i.e. they have derivatives up to any order), they are locally Lipschitz-continuous. By the Cauchy-Lipschitz theorem (a.k.a. Picard-Lindelöf) [14], there exists a unique maximal solution to the initial value problem ($\dot{\mathbf{x}} = \mathbf{p}$, $\mathbf{x}(0) = \mathbf{x}_0$) defined for t in some nonempty open interval.

For $h \in \mathbb{R}[x_1 \dots, x_n]$, if $h(\mathbf{x}(t)) = 0$ for all $t \geq 0$, we say that the equation $h = 0$ is a (*positive*) *invariant* under the flow of \mathbf{p} . In differential dynamic logic [20], invariance of $h = 0$ is semantically equivalent to the validity of the following formula:

$$(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0) \quad (2)$$

Geometrically, the equation $h = 0$ represents the set of real roots of h . Such a set is known as *real algebraic set* or a *real variety* and will be henceforth denoted by $V_{\mathbb{R}}(h)$. Algebraic sets are intimately related to sets of polynomials with special algebraic properties called *ideals*. Ideals are closed under addition and external multiplication; that is, if I is an ideal, then for all $h_1, h_2 \in I$, the sum $h_1 + h_2 \in I$; and if $h \in I$, then, $qh \in I$, for all $q \in \mathbb{R}[x_1 \dots, x_n]$. To say that the real variety $V_{\mathbb{R}}(h)$ of the ideal *generated* by h is invariant under the flow of the vector field \mathbf{p} is equivalent to the statement that the equation $h = 0$ is invariant.

We will use ∇h to denote the gradient of $h : \mathbb{R}^n \rightarrow \mathbb{R}$, that is the vector of its partial derivatives $(\frac{\partial h}{\partial x_1}, \dots, \frac{\partial h}{\partial x_n})$. The *Lie derivative* of h along the vector field \mathbf{p} gives the rate of change of h along the flow of $\dot{\mathbf{x}} = \mathbf{p}(\mathbf{x})$ and is formally defined as the scalar product of ∇h and \mathbf{p} .

$$\mathfrak{L}_{\mathbf{p}}(h) \stackrel{\text{def}}{=} \nabla h \cdot \mathbf{p} . \quad (3)$$

Higher-order Lie derivatives are defined recursively as $\mathfrak{L}_{\mathbf{p}}^{(k+1)}(h) = \mathfrak{L}_{\mathbf{p}}(\mathfrak{L}_{\mathbf{p}}^{(k)}(h))$, with $\mathfrak{L}_{\mathbf{p}}^{(0)}(h) = h$.

We now recall five important proof rules for checking invariance of polynomial equalities, or equivalently the validity of Eq. (2). In Fig. 1, $\text{DI}_=$ shows the differential invariant [21] proof rule restricted to handling equalities. The condition imposed by the premise of $\text{DI}_=$ is sufficient, but not necessary; it characterizes polynomial invariant functions. The premise of the Polynomial-scale consecution proof rule [26], P-c in Fig. 1, requires $\mathfrak{L}_{\mathbf{p}}(h)$ to be in the ideal generated by h . The condition is also only sufficient (but is particularly suitable for *generating* invariant varieties [16]). We also consider the constant-scale consecution proof rule [26,29], denoted by C-c . The premise of proof rule C-c requires that $\mathfrak{L}_{\mathbf{p}}(h) = \lambda h$, where λ is a scalar, not a polynomial as in P-c . It is therefore a simple special case of P-c . When $\lambda = 0$, one obtains the premise of the proof rule $\text{DI}_=$. It is worth noting that P-c , including its special case C-c , was mentioned as early as 1878 [5] and used extensively in the study of integrability of dynamical systems, where they are known as *second integrals* [12, Chapter 2]. It serves as a natural extension to invariant functions, also known as *first integrals*, which are covered by the proof rule $\text{DI}_=$. The proof rule *Lie* gives Lie's criterion [13,19] for invariance of $h = 0$; this proof rule will be discussed in more depth and extended to handle tricky cases in Section 3. The last rule, DRI in Fig. 1, was recently introduced and characterizes (i.e. gives necessary and sufficient conditions for) invariant varieties under the flow of polynomial vector fields [9]. The number N in DRI is the maximum length of the ascending chain of polynomial ideals $\langle h \rangle \subset \langle h, \mathfrak{L}_{\mathbf{p}}(h) \rangle \subset \langle h, \mathfrak{L}_{\mathbf{p}}(h), \mathfrak{L}_{\mathbf{p}}^{(2)}(h) \rangle \subset \dots$, which is finite and computable [9].

$$\begin{array}{l}
(\text{DI}_=) \frac{\mathfrak{L}_{\mathbf{p}}(h) = 0}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)} \quad (\text{C-c}) \frac{\exists \lambda \in \mathbb{R}, \mathfrak{L}_{\mathbf{p}}(h) = \lambda h}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)} \\
(\text{Lie}) \frac{h = 0 \rightarrow (\mathfrak{L}_{\mathbf{p}}(h) = 0 \wedge \nabla h \neq \mathbf{0})}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)} \quad (\text{P-c}) \frac{\mathfrak{L}_{\mathbf{p}}(h) \in \langle h \rangle}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)} \\
(\text{DRI}) \frac{h = 0 \rightarrow \bigwedge_{i=0}^{N-1} \mathfrak{L}_{\mathbf{p}}^{(i)}(h) = 0}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)}
\end{array}$$

Figure 1: Proof rules for checking the invariance of $h = 0$ w.r.t. \mathbf{p} : $\text{DI}_=$ [23, Theorem 3], C-c and P-c [26, Lemma 2], Lie [19, Theorem 2.8], DRI [9, Theorem 2]

3 Lie's Criterion

One immediate (and somewhat embarrassing) deficiency of the proof rule Lie (Fig. 1) is its inability to prove invariance properties for isolated points (e.g. system equilibria) for the simple reason that a description of such a point $\mathbf{a} = (a_1, \dots, a_n) \in \mathbb{R}^n$ is (when $n > 1$) given by the sum-of-squares equation $h(\mathbf{x}) = (x_1 - a_1)^2 + \dots + (x_n - a_n)^2 = 0$. This sum-of-squares polynomial h is *positive-definite*, i.e. $h(\mathbf{a}) = 0$ and $h(\mathbf{x}) > 0$ for all $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{a}\}$. Positive definite functions have vanishing gradient at their minima, in this case \mathbf{a} , and thus the formula $h = 0 \rightarrow \nabla h = \mathbf{0}$ holds. This violates the regularity condition in the premise of the proof rule Lie, namely:

$$h = 0 \longrightarrow \nabla h \neq \mathbf{0} . \quad (4)$$

In fact, $h = 0 \rightarrow \mathfrak{L}_{\mathbf{p}}(h) = 0$ is a necessary condition when $h = 0$ is an invariant equation. Note that simply removing Eq. (4) from the premise of the proof rule Lie is unsound (see e.g. [23]); that is, the condition $h = 0 \rightarrow \mathfrak{L}_{\mathbf{p}}(h) = 0$ alone is insufficient to prove the invariance property for $h = 0$. Unsoundness in the above naïve attempt at a generalization is a consequence of *singularities* that may be present in the variety $V_{\mathbb{R}}(h)$. Singularities of $V_{\mathbb{R}}(h)$ are points $\mathbf{x} \in V_{\mathbb{R}}(h)$ where the gradient of h vanishes, i.e. $\nabla h(\mathbf{x}) = \mathbf{0}$.

Definition 2 (Singular Locus). *Let $h \in \mathbb{R}[x_1, \dots, x_n]$, the singular locus of $h = 0$, henceforth denoted $\text{SL}(h)$, is the set of singular points, that is, points \mathbf{x} satisfying*

$$h = 0 \wedge \frac{\partial h}{\partial x_1} = 0 \wedge \dots \wedge \frac{\partial h}{\partial x_n} = 0 .$$

Points that are not singular are called regular. At singular points, the Lie derivative of h along any vector field is $\mathbf{0} \cdot \mathbf{p} = 0$. To avoid these degenerate cases, the regularity condition (Eq. (4)) rules out singularities altogether. In the next section we present two extensions of Lie's criterion that, in a similar vein to [27], partially overcome the strong regularity condition by treating the points on the singular locus separately.

3.1 Handling Singularities

Equilibria are points in the state space where the vector field vanishes ($\mathbf{p} = \mathbf{0}$) so that there is no motion. However, as seen above, Lie’s criterion cannot generally be applied to prove invariance properties of isolated equilibria because their description involves singularities. One simple way to resolve this issue is to drop the non-vanishing gradient condition and replace it with the proviso that there be no flow (that is $\mathbf{p} = \mathbf{0}$) in the variables of the invariant candidate on the singular locus; this will allow singularities in the invariant candidate and will provide a *sound* proof method in which there is no need to check for non-vanishing gradient. Below we present two extensions to the proof rule Lie and justify their soundness after recalling some basic geometric notions.

Definition 3 (Lie[◦]: Lie + Equilibria).

$$(\text{Lie}^\circ) \frac{h = 0 \rightarrow (\mathfrak{L}_{\mathbf{p}}(h) = 0 \wedge (\text{SL}(h) \rightarrow \bigwedge_{x_i \in \text{vars}(h)} p_i = 0))}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)},$$

where $\text{vars}(h)$ denotes the set of state variables x_i occurring in the polynomial h .

The Lie[◦] proof rule can be generalized further at the expense of adding an extra variable by replacing the “no flow” condition ($p_i = 0$) for points on the singular locus with $\forall \lambda. h(\mathbf{x} + \lambda \mathbf{p}(\mathbf{x})) = 0$, where λ is a fresh symbol.

Definition 4 (Lie^{*}: Lie + Vanishing Sub-tangent).

$$(\text{Lie}^*) \frac{h = 0 \rightarrow (\mathfrak{L}_{\mathbf{p}}(h) = 0 \wedge (\text{SL}(h) \rightarrow h(\mathbf{x} + \lambda \mathbf{p}) = 0))}{(h = 0) \rightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)}.$$

To prove soundness of Lie[◦] and Lie^{*}, we use a result about positive invariance of closed sets under locally Lipschitz-continuous vector fields, known as the Nagumo theorem [18,30, Chapter 10, XV–XVI, pp. 117-119], which gives a powerful (but generally intractable) geometric characterization of positively invariant closed sets. The notion of positive invariance of the equation $h = 0$ from Section 2 generalizes to an arbitrary set.

Definition 5 (Invariant Sets). A set S is positively (negatively) invariant under the flow of $\dot{\mathbf{x}} = \mathbf{p}$ if for all $\mathbf{x}_0 \in S$ we have $\mathbf{x}(\mathbf{x}_0, t) \in S$ for all $t \geq 0$ ($t \leq 0$), where $\mathbf{x}(\mathbf{x}_0, t)$ is the solution of the initial value problem ($\dot{\mathbf{x}} = \mathbf{p}, \mathbf{x}(0) = \mathbf{x}_0$). A set S is bi-invariant if it is both positively and negatively invariant.

Nagumo’s theorem needs the geometric notion of sub-tangential vectors to a set.

Definition 6 (Sub-tangent vector). A vector $\mathbf{v} \in \mathbb{R}^n$ is sub-tangential to a set S at $\mathbf{x} \in S$ if

$$\liminf_{\lambda \rightarrow 0^+} \frac{\text{dist}(S, \mathbf{x} + \lambda \mathbf{v})}{\lambda} = 0,$$

where dist denotes the Euclidean set distance, i.e. $\text{dist}(S, \mathbf{x}) \equiv \inf_{\mathbf{y} \in S} \|\mathbf{x} - \mathbf{y}\|$.

Theorem 1 (Nagumo Theorem). Given a continuous system $\dot{\mathbf{x}} = \mathbf{p}(\mathbf{x})$ and assuming that solutions exist and are unique inside some open set $O \subseteq \mathbb{R}^n$, let $S \subset O$ be a closed set. Then, S is positively invariant under the flow of the system if and only if $\mathbf{p}(\mathbf{x})$ is sub-tangential to S for all $\mathbf{x} \in \partial S$, where ∂S is the boundary of S .

Let us observe that given $\mathbf{x} \in \partial S$, if $\mathbf{x} + \lambda \mathbf{p}(\mathbf{x}) \in S$ for all $\lambda \in \mathbb{R}$, then $\text{dist}(S, \mathbf{x} + \lambda \mathbf{p}(\mathbf{x})) = 0$ and $\mathbf{p}(\mathbf{x})$ is sub-tangential to S at \mathbf{x} . This observation is important for algebraic sets, for which $\partial S = S$, and the condition $\mathbf{x} + \lambda \mathbf{p}(\mathbf{x}) \in S$ translates to $h(\mathbf{x} + \lambda \mathbf{p}(\mathbf{x})) = 0$. This is the main idea behind the soundness of the proof rule Lie^* (see [10] for the detailed proof).

Proposition 1. *The proof rule Lie^* is sound.*

The case $\mathbf{p}(\mathbf{x}) = 0$ for all \mathbf{x} in the singular locus is a special case of the proof rule Lie^* . Therefore, the soundness of Lie° is an immediate corollary of Prop. 1.

Corollary 1. *The proof rule Lie° is sound.*

4 Proof Rules: Hierarchy and Complexity

In this section, we compare the deductive power of the existing (Fig. 1) and the newly-introduced proof rules (Lie° and Lie^* in Section 3) for checking the invariance of algebraic sets. This study should be complemented by another comparison that considers the interaction between the different proof rules in the context of a formal proof system in a similar vein to [24]. We leave this for future work.

Given two proof rules (let us call them R_1 and R_2) featuring the same conclusion ($(h = 0) \longrightarrow [\dot{\mathbf{x}} = \mathbf{p}](h = 0)$), we will say that R_2 generalizes R_1 and write $R_2 \succcurlyeq R_1$ (or $R_1 \preccurlyeq R_2$), if the premise of R_1 implies the premise of R_2 . That is, if R_1 proves that $h = 0$ is an invariant, then so does R_2 . If $R_1 \preccurlyeq R_2$ and $R_1 \succcurlyeq R_2$, we say that R_1 and R_2 are equivalent, and denote this by $R_1 \sim R_2$. Likewise, $R_1 \not\preccurlyeq R_2$ (or $R_2 \not\preccurlyeq R_1$) denotes that R_1 is not generalized by R_2 . We also write $R_1 \prec R_2$ when $R_1 \preccurlyeq R_2$ and $R_1 \not\sim R_2$. That is, the rule R_2 *increases* the deductive power of R_1 .

It is easy to see that the order \preccurlyeq is a partial order (with \sim acting as equality): it is reflexive, $R \preccurlyeq R$ (the premise of R implies itself); it is anti-symmetric (by definition), and transitive: if $R_1 \preccurlyeq R_2$ and $R_2 \preccurlyeq R_3$, then the premise of R_1 implies the premise of R_3 by transitivity of the implication, so $R_1 \preccurlyeq R_3$. Finally, If $R_1 \not\preccurlyeq R_2$ and $R_1 \not\preccurlyeq R_2$, we will write $R_1 \prec\!\succ R_2$ and say that the proof rules R_1 and R_2 are *incomparable*. This means that for both R_1 and R_2 there are problems that one rule can prove and the other cannot. In the sequel, we use the partial order \preccurlyeq to illustrate the lattice structure of the proof rules under consideration. In Section 4.2 we discuss the computational complexity of the conditions appearing in their premises.

4.1 Hierarchy

We use the partial order (\preccurlyeq) to compare the deductive power of all considered proof rules $\{\text{DI}_=, \text{C-c}, \text{P-c}, \text{Lie}, \text{Lie}^\circ, \text{Lie}^*, \text{DRI}\}$. For convenience, the propositions of this section are summarized in the comparison matrix (Fig. 3). For instance, Prop. 6 proves that $\text{DI}_= \prec\!\succ \text{Lie}$. Cells without numbers are proved by transitivity of the partial order. For instance, $\text{DI}_= \prec \text{DRI}$ can be proved using $\text{DI}_= \prec \text{C-c}$ (Prop. 2) and $\text{C-c} \prec \text{P-c}$ (Prop. 3) and $\text{P-c} \prec \text{DRI}$ (Prop. 5). The Hasse diagram (Fig. 2) gives the lattice structure where arrows represent strictly increasing deductive power; every missing edge in the graph represents $\prec\!\succ$, as shown in the comparison matrix.




Figure 2: Hasse diagram. An arrow $R_1 \rightarrow R_2$ means $R_1 \prec R_2$, all other non depicted links mean $(\prec \succ)$.

	DI=	C-c	P-c	Lie	Lie°	Lie*	DRI
DI=	\sim	\prec	\prec	$\prec \succ$	$\prec \succ$	$\prec \succ$	\prec
		2		6	8	7	
C-c	\succ	\sim	\prec	$\prec \succ$	$\prec \succ$	$\prec \succ$	\prec
	2		3	9	10	10	
P-c	\succ	\succ	\sim	$\prec \succ$	$\prec \succ$	$\prec \succ$	\prec
		3		9	10	10	5
Lie	$\prec \succ$	$\prec \succ$	$\prec \succ$	\sim	\prec	\prec	\prec
	6	9	9		4		
Lie°	$\prec \succ$	$\prec \succ$	$\prec \succ$	\succ	\sim	\prec	\prec
	8	10	10	4		4	
Lie*	$\prec \succ$	$\prec \succ$	$\prec \succ$	\succ	\succ	\sim	\prec
	7	10	10		4		5
DRI	\succ	\succ	\succ	\succ	\succ	\succ	\sim

Figure 3: Comparison matrix of the deductive power of $\{DI=, C-c, P-c, Lie, Lie^\circ, Lie^*, DRI\}$. The numbers refer to the propositions.

We begin by comparing Darboux-based proof rules, i.e. $\{DI=, C-c, P-c\}$ and then proceed to the Lie-based proof rule family, i.e. $\{Lie, Lie^\circ, Lie^*\}$. Next, we demonstrate the deductive superiority of the necessary and sufficient conditions in the premise of the proof rule DRI. Finally, we establish that Darboux-based proof rules and Lie-based proof rules form two *distinct* proof rule families; that is, any proof rule from one family is deductively incomparable to any proof rule from the other family.

Proposition 2. $DI= \prec C-c$.

Proof. The premise of the rule C-c requires the existence of some $\lambda \in \mathbb{R}$, such that $\mathfrak{L}_{\mathbf{p}}(h) = \lambda h$. In particular, $\lambda = 0$ gives the premise of $DI=$. Thus, $DI= \preceq C-c$. To see that $DI= \not\preceq C-c$, consider the one-dimensional vector field $\mathbf{p} = (x)$, we have $\mathfrak{L}_{\mathbf{p}}(x) = 1x$, and hence C-c ($\lambda = 1$) concludes that $x = 0$ is an invariant. However, $DI=$ cannot prove the invariance of $x = 0$ because x is not a conserved quantity in the system. \square

Proposition 3. $C-c \prec P-c$.

Proof. The premise of the rule P-c requires the existence of some $\alpha \in \mathbb{R}[x]$, such that $\mathfrak{L}_{\mathbf{p}}(h) = \alpha h$ (equivalently, $\mathfrak{L}_{\mathbf{p}}(h) \in \langle h \rangle$). In particular, the constant polynomial gives the premise of C-c. Thus, $C-c \preceq P-c$. To prove that $C-c \not\preceq P-c$, consider the two-dimensional vector field $\mathbf{p} = (xy, x)$, we have $\mathfrak{L}_{\mathbf{p}}(x) = xy$ (or equivalently $\mathfrak{L}_{\mathbf{p}}(x) \in \langle x \rangle \subset \mathbb{R}[x, y]$) and hence conclude, using P-c, that $x = 0$ is an invariant. However, C-c fails to prove this invariant as the required cofactor is not a scalar. \square

Proposition 4. $Lie \prec Lie^\circ$ and $Lie^\circ \prec Lie^*$.

Proof. We already established that $\text{Lie} \preceq \text{Lie}^\circ$ (Prop. 1) and $\text{Lie}^\circ \preceq \text{Lie}^*$ (Prop. 1); we give two counterexamples to establish the strict inclusion. **(I)** $\text{Lie} \not\preceq \text{Lie}^\circ$. Whenever the variety has a singularity, the proof rule Lie will fail. Lie° is tailored to prove invariance of equilibrium points in addition to regular points of the variety. For instance, for $\mathbf{p} = ((-1 + x_1)x_2, x_2(1 + x_2))$, Lie fails to prove that $h = (-1 + x_1)^2 + (1 + x_2)^2 = 0$ is invariant as the gradient ∇h vanishes at $(1, -1)$ and $h((1, -1)) = 0$. However, at $(1, -1)$ we also have $p_1 = p_2 = 0$, and hence the premise of Lie° is satisfied, and $h = 0$ is proved to be an invariant under the flow of \mathbf{p} . **(II)** $\text{Lie}^\circ \not\preceq \text{Lie}^*$. In addition to equilibria, Lie^* goes one step further and handles all singular points, \mathbf{x} , where the vector $\mathbf{x} + \lambda\mathbf{p}$ is in the variety $V_{\mathbb{R}}(h)$ for all $\lambda \in \mathbb{R}$ (that is $h(\mathbf{x} + \lambda\mathbf{p}) = 0$, for all λ). For instance, consider the polynomial $h = x_1x_2x_3$, its singular locus is given by the three axes $x_1 = x_2 = 0$, $x_1 = x_3 = 0$ and $x_2 = x_3 = 0$. For the vector field $\mathbf{p} = (x_1, x_2, x_3)$, the equilibrium point is at the origin $(0, 0, 0)$, which obviously does not contain the entire singular locus of h . Thus, Lie° fails but Lie^* succeeds because $h(\mathbf{x} + \lambda\mathbf{p}) = 0$ when \mathbf{x} is a point of one of the axes. \square

Proposition 5. $\text{P-c} \prec \text{DRI}$ and $\text{Lie}^* \prec \text{DRI}$.

Proof. DRI is both necessary and sufficient [9], so we know that $\text{P-c} \preceq \text{DRI}$ and $\text{Lie}^* \preceq \text{DRI}$. To prove the claim it is left to show that **(I)** $\text{P-c} \not\preceq \text{DRI}$. Consider the following two-dimensional vector field: $\mathbf{p} = ((-1 + x_1)(1 + x_1), (-1 + x_2)(1 + x_2))$. The candidate invariant (given by the roots of the Motzkin polynomial) $h = 1 - 3x_1^2x_2^2 + x_1^4x_2^2 + x_1^2x_2^4 = 0$ cannot be proved using P-c , as $\mathcal{L}_{\mathbf{p}}(h) \notin \langle h \rangle$. However, the invariance property may be proved using DRI. For this, we need to consider the second-order Lie derivative of h and we prove that $\mathcal{L}_{\mathbf{p}}^{(2)}(h) \in \langle h, \mathcal{L}_{\mathbf{p}}(h) \rangle$. Thus, the premise of DRI holds for $N = 2$. **(II)** $\text{Lie}^* \not\preceq \text{DRI}$. Consider the following three-dimensional vector field $\mathbf{p} = (-x_2 + x_1(1 - x_1^2 - x_2^2), x_1 + x_2(1 - x_1^2 - x_2^2), x_3)$. We want to prove that $h = (-1 + x_1^2 + x_2^2)^2 + x_3^2 = 0$ is an invariant. In this case, the variety $V_{\mathbb{R}}(h)$ is exactly equal to the singular locus of h which is the two-dimensional unit circle $-1 + x_1^2 + x_2^2 = 0$. However, at all points of this unit circle, the vector field \mathbf{p} is equal to $(-x_2, x_1, 0) \neq 0$, which prevents us from using Lie^* (because $h((x_1, x_2, 0) + \lambda(-x_2, x_1, 0)) \neq 0$ for some $\lambda \in \mathbb{R}$). The rule DRI proves the invariance of $h = 0$ with $N = 2$. \square

To appreciate the difference between $\text{DI}_=$ and Lie , let us note that while the condition in the premise of $\text{DI}_=$ may seem strong (i.e. too conservative), singularities in the invariant candidate do not present a problem for $\text{DI}_=$, whereas the premise of Lie rules out such candidates altogether (see Fig. 4). Indeed, the proof rule Lie cannot prove that $0 = 0$ (the whole space is invariant), whereas this is the most trivial case for $\text{DI}_=$.

Proposition 6 ($\text{DI}_=$ and Lie are incomparable.). $\text{DI}_= \prec \succ \text{Lie}$.

Proof. **(I)** $\text{DI}_= \not\preceq \text{Lie}$. For the vector field $\mathbf{p} = (-2x_2, -2x_1 - 3x_1^2)$, the equation $x_1^2 + x_1^3 - x_2^2 = 0$ is provable with $\text{DI}_=$ but not Lie , see Fig. 4 (left). **(II)** $\text{DI}_= \not\preceq \text{Lie}$. For the vector field $\mathbf{p} = (x_1 - x_1^3 - x_2 - x_1x_2^2, x_1 + x_2 - x_1^2x_2 - x_2^3)$, the invariance of the limiting cycle $x_1^2 + x_2^2 - 1 = 0$ is provable with Lie but not $\text{DI}_=$, see Fig. 4 (right). \square

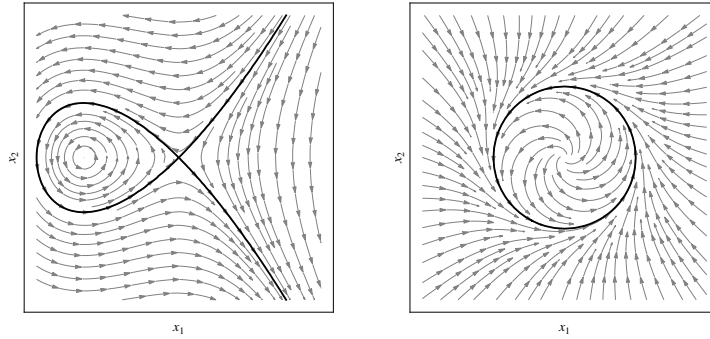


Figure 4: The invariance of the variety $V_{\mathbb{R}}(x_1^2 + x_1^3 - x_2^2)$ (**left**) provable using $\text{DI}_=$ (but not Lie since $(0, 0)$ is a singular point) and a smooth invariant limit cycle $V_{\mathbb{R}}(x_1^2 + x_2^2 - 1)$ (**right**) provable using Lie (but not $\text{DI}_=$ since it is not an invariant function).

We now prove that Lie -based proof rules $\{\text{Lie}, \text{Lie}^\circ, \text{Lie}^*\}$, and Darboux-based proof rules $\{\text{DI}_=, \text{C-c}, \text{P-c}\}$ are two distinct families of proof rules; that is, any Lie -based proof is deductively incomparable to any Darboux-based proof rule. The following lemma follows from the transitivity of the partial order.

Lemma 1. *If $R_1 \preceq R_2$ and $R_3 \prec \succ R_1$, then $R_2 \not\preceq R_3$.*

Proof. Consider three proof rules R_1, R_2 and R_3 . If $R_2 \preceq R_3$, using $R_1 \preceq R_2$, one gets by transitivity $R_1 \preceq R_3$, which contradicts the assumption $R_3 \prec \succ R_1$.

Proposition 7. $\text{DI}_= \prec \succ \text{Lie}^*$.

Proof. Since $\text{Lie} \preceq \text{Lie}^\circ$ (Prop. 1) and $\text{Lie}^\circ \preceq \text{Lie}^*$ (Prop. 1), then $\text{Lie} \preceq \text{Lie}^*$. By Lem. 1, from $\text{Lie} \preceq \text{Lie}^*$ and $\text{DI}_= \prec \succ \text{Lie}$ (Prop. 6), we get $\text{Lie}^* \not\preceq \text{DI}_=$. The following example proves that $\text{DI}_= \not\preceq \text{Lie}^*$: Consider the three-dimensional vector field $\mathbf{p} = (x_2, -x_1, 0)$. The invariance of the equation $x_3^2 + (-1 + x_1^2 + x_2^2 + x_3^2)^2 = 0$ cannot be established using Lie^* (the singular locus is a circle in \mathbb{R}^3), but is easily provable using $\text{DI}_=$ as $\mathcal{L}_{\mathbf{p}}(h)$ vanishes. \square

Proposition 8. $\text{DI}_= \prec \succ \text{Lie}^\circ$.

Proof. By Lem. 1, from $\text{Lie} \preceq \text{Lie}^\circ$ (Prop. 1) and $\text{DI}_= \prec \succ \text{Lie}$ (Prop. 6), we get $\text{Lie}^\circ \not\preceq \text{DI}_=$. On the other hand, if $\text{DI}_= \preceq \text{Lie}^\circ$ then, by transitivity $\text{DI}_= \preceq \text{Lie}^*$ (since $\text{Lie}^\circ \preceq \text{Lie}^*$ by Prop. 1), which contradicts $\text{DI}_= \prec \succ \text{Lie}^*$ (Prop. 7). Thus, $\text{DI}_= \not\preceq \text{Lie}^\circ$, and the proposition follows. \square

Similarly, by substituting $\text{DI}_=$ by Lie , Lie^* by P-c , and Lie° by C-c in Prop. 7 and Prop. 8 as well as their respective proofs, we show that:

Proposition 9. $\text{Lie} \prec \succ \text{P-c}$ and $\text{Lie} \prec \succ \text{C-c}$.

Proof. To complete the proof, we still need an example showing that $\text{Lie} \not\preceq \text{P-c}$. Consider the vector field $\mathbf{p} = (3(-4 + x^2), 3 + xy - y^2)$, the proof rule Lie fails to prove that the equation $h = -3 + x^2 + 2xy + 6y^2 + 2xy^3 + y^4 = 0$ is invariant as the singular locus of h contains $(-2, 1)$ and $(2, -1)$. However, $\mathcal{L}_{\mathbf{p}}(h) = (6x - 4y)h$ and therefore P-c proves that $h = 0$ is an invariant equation. \square

The remaining cases follow from the results established above.

Proposition 10. *For $d \in \{C\text{-c}, P\text{-c}\}$, $\ell \in \{\text{Lie}^\circ, \text{Lie}^*\}$, $d \prec \succ \ell$.*

Proof. Since $\text{DI}_= \prec d$, if $d \preceq \ell$, then $\text{DI}_= \preceq \ell$. However, $\text{DI}_= \prec \succ \ell$ (Prop. 7 and Prop. 8). Thus $d \not\preceq \ell$. Similarly, since $\ell \succ \text{Lie}$, if $d \succ \ell$, then $d \succ \text{Lie}$ which contradicts $d \prec \succ \text{Lie}$ (Prop. 9). Hence $d \not\succeq \ell$ and the proposition follows. \square

Remark 1. Provided the invariant candidate has no singular points, Lie’s criterion is known to be both necessary and sufficient to prove invariance properties of level sets [19, Theorem 2.8]. Also, $\text{DI}_=$ characterizes invariant functions [23] but not all invariant equations. On the other hand, for algebraic differential equations, the differential radical criterion in DRI fully characterizes all invariant algebraic sets [9]. Thus, as established in Prop. 5, DRI increases the deductive power of both Darboux-based rules $\{\text{DI}_=, C\text{-c}, P\text{-c}\}$ and Lie-based rules $\{\text{Lie}, \text{Lie}^\circ, \text{Lie}^*\}$, which form different families.

4.2 Complexity

While decidable [28], the complexity of real quantifier elimination is doubly exponential in the number of quantifier alternations [6]. Most existing implementations of real quantifier elimination procedures are based on cylindrical algebraic decomposition (CAD) [2,3], which has doubly-exponential running time in the number of variables.

The purely existential fragment of real quantifier elimination has been shown to exhibit singly exponential time complexity in the number of variables [1]. However, in practice this has not yet led to an efficient decision procedure, so typically it is much more efficient to use CAD. Theoretically, the best bound on the complexity of deciding a sentence in the universal theory of \mathbb{R} is given by $(sd)^{O(n)}$, where s is the number of polynomials in the formula, d their maximum degree and n the number of variables [1].

The premises of rules $\text{DI}_=$, Lie , Lie° , Lie^* are universally quantified sentences in the theory of real arithmetic. One sees from the expression for the complexity bound that it is important for these rules to keep the number of variables low and also that it is desirable to work with polynomials of low degree. In this respect, we would anticipate the rule Lie^* to incur a performance penalty from introducing a fresh variable.

$C\text{-c}$ and $P\text{-c}$ involve reasoning about multivariate polynomial ideal membership, which is an EXPSPACE -complete problem over \mathbb{Q} [17]. Gröbner basis algorithms allow us to perform membership checks in ideals generated by multivariate polynomials. Significant advances have been made in algorithms for computing Gröbner bases [8] which in practice can be expected to perform very well.

The premise of DRI may be decided using a real quantifier elimination procedure, like any other first-order sentence in the theory of real arithmetic. However, in order to obtain the bound N on the order of the Lie derivatives, one is also required to check for polynomial ideal membership at least $N - 1$ times.

5 Square-free Reduction

In this section we assess the utility of performing square-free reduction of invariant candidates as a means of (i) increasing the deductive power of Lie-based proof rules and (ii) simplifying problems passed to decision procedures for real arithmetic.

5.1 Square-free Reduction with Lie-based Proof Rules

While Lie uses a powerful criterion that captures a large class of practically relevant invariant sets, it will fail for some seemingly simple invariant candidates. For instance, the condition in the premise of Lie will not hold when the goal is to prove that $h = x^2 - 6x + 9 = 0$ is invariant, no matter what vector field one considers. The reason for this is simple: $x^2 - 6x + 9$ factorizes into $(x - 3)^2$. The problem here lies in the polynomial h itself, rather than the real variety $V_{\mathbb{R}}(h)$. In fact, $V_{\mathbb{R}}(h)$ is exactly the singular locus of h and the proof rule Lie fails because *all* points inside $V_{\mathbb{R}}(h)$ are singular points. More generally, the chain rule implies $\nabla h^k \cdot \mathbf{p} = kh^{k-1} \nabla h \cdot \mathbf{p}$, which has the consequence that any polynomial h which is not square-free will have vanishing gradient at the real roots of factors with multiplicity greater than 1.

One can eliminate such annoying instances by reducing h to square-free form, which is a basic pre-processing step used in computer algebra systems. The square-free reduction of a polynomial h may be computed efficiently as follows:

$$\text{SF}(h) = \frac{h}{\gcd\left(h, \frac{\partial h}{\partial x_1}, \dots, \frac{\partial h}{\partial x_n}\right)}. \quad (5)$$

Intuitively, in performing square-free reduction we hope to shrink the singular locus of the original polynomial. If $\text{SL}(\text{SF}(h))$ is the empty set (which is the case for $h = x^2 - 6x + 9$ in the example given above), the proof rule Lie applies to $\text{SF}(h)$ but not to h . In general, $\text{SF}(h)$ may satisfy the assumptions of the proof rules Lie^o or Lie^{*}, where h fails to do so. It is always sound to conclude that $h = 0$ is invariant from the knowledge that $\text{SF}(h) = 0$ is invariant, since real varieties remain unaltered under square-free reduction of their defining polynomials [4], i.e. $V_{\mathbb{R}}(h) \equiv V_{\mathbb{R}}(\text{SF}(h))$, thus replacing h with $\text{SF}(h)$ in the premise of Lie, Lie^o and Lie^{*} remains sound and enlarges the class of polynomials that these proof rules can work with.

Proposition 11. For all $\ell \in \{\text{Lie}, \text{Lie}^o, \text{Lie}^*\}$, $\ell \prec \text{SF } \ell$.

This result is unsurprising when one understands that Lie-based proof rules use geometric concepts to prove invariance properties of sets. In fact, the square-free reduction removes some purely algebraic oddities that prevent the geometric condition from holding true when checked syntactically by a machine.

In addition to increasing the deductive power, the square-free reduction reduces the total degree of the polynomial in the invariant candidate and hence serves to reduce the complexity of deciding the conditions in the premise (see Section 4.2). In our implementation, we adopt the convention that invariant candidates supplied to Lie and its generalizations are square-free reduced in a pre-processing step.

5.2 Square-free Reduction with Darboux-based proof rules

Unlike Lie-based proof rules, it is perhaps surprising that using square-free reduction as a pre-processing step for the proof rules DI₌ and C-c, denoted SFDI₌ and SFC-c respectively, does *not*, in general, increase the deductive power.

Proposition 12. DI₌ \prec SFDI₌.

Proof. (I) $\text{DI}_= \not\prec \text{SFDI}_=$. The polynomial $h = x^2y$ is an invariant function for the vector field $\mathbf{p} = (\frac{\partial h}{\partial y}, -\frac{\partial h}{\partial x}) = (x^2, -2xy)$, thus $\text{DI}_=$ proves the invariance of $h = 0$. However, $\text{SF}(h)$ is not an invariant function for the same vector field, since $\mathfrak{L}_{\mathbf{p}}(\text{SF}(h)) = \mathfrak{L}_{\mathbf{p}}(xy) = -x^2y \neq 0$, thus $\text{SFDI}_=$ fails to prove the invariance of $h = 0$. (II) $\text{SFDI}_= \not\prec \text{DI}_=$. Similarly, the polynomial $h = xy$ is an invariant function for the vector field $\mathbf{p} = (\frac{\partial h}{\partial y}, -\frac{\partial h}{\partial x}) = (x, -y)$, thus $\text{SFDI}_=$ proves the invariance of $x^2y = 0$, since $\text{SF}(x^2y) = h$. However, $\text{DI}_=$ fails to prove the invariance of $x^2y = 0$, because $\mathfrak{L}_{\mathbf{p}}(x^2y) = x^2y \neq 0$. \square

Prop. 12 may at first seem counter-intuitive. However, the criterion in the premise of $\text{DI}_=$ is different in that it proves that the candidate h is an *invariant function*. In performing square-free reduction on h , one in general obtains a different function, $\text{SF}(h)$, which need not be conserved in the system if h is conserved or, conversely, may be conserved even if h is not.

The same observation holds for C-c as the SF reduction does not preserve the constant rate exponential decrease (or increase).

Proposition 13. C-c $\prec \succ$ SFC-c.

Proof. (I) C-c $\not\prec$ SFC-c. The proof rule C-c proves the invariance of $h = x^2y = 0$ for the vector field $\mathbf{p} = (x^2, y(1-2x))$ as $\mathfrak{L}_{\mathbf{p}}(h) = 1h$. However, C-c cannot prove $\text{SF}(h) = 0$, since $\mathfrak{L}_{\mathbf{p}}(\text{SF}(h)) = \mathfrak{L}_{\mathbf{p}}(xy) = (1-x)\text{SF}(h)$. (II) SFC-c $\not\prec$ C-c. For the same h , C-c proves the invariance of $\text{SF}(h) = 0$ for the vector field $\mathbf{p} = (x^2, y(1-x))$ as $\mathfrak{L}_{\mathbf{p}}(\text{SF}(h)) = \mathfrak{L}_{\mathbf{p}}(xy) = 1\text{SF}(h)$. However, without the SF reduction C-c alone fails to prove the invariance of $h = 0$ for the considered \mathbf{p} , as $\mathfrak{L}_{\mathbf{p}}(h) = (x+1)h$. \square

After Prop. 12 and 13, one expects P-c to be incomparable with its square-free counterpart. Surprisingly, the proof rules P-c and SFP-c (which applies P-c after the square-free reduction) are in fact equivalent. This follows from the fact that a polynomial is Darboux for a vector field \mathbf{p} if and only if all its factors are also Darboux for the same vector field. Our findings are stated in Prop. 14 and its corollary Prop. 15 (both proofs are available in the report [10]).

Proposition 14. Let $h = q_1^{m_1} \cdots q_r^{m_r}$ denote the decomposition of the polynomial h into irreducible (over the reals) factors, q_i . Then, h is Darboux for \mathbf{p} if and only if, for all i , q_i is Darboux for \mathbf{p} .

Proposition 15. P-c \sim SFP-c.

Remark 2. The condition $\mathfrak{L}_{\mathbf{p}}(p) \in \langle \text{SF}(p) \rangle$ —which is weaker than $\mathfrak{L}_{\mathbf{p}}(p) \in \langle p \rangle$ —is not sufficient to prove the invariance of $p = 0$. It is therefore an unsound proof rule. Consider the polynomial $p = (-1 + x^2)^2$ and the 1-dimensional vector field $\dot{x} = x$. Although $\mathfrak{L}_{\mathbf{p}}(p) = 4(-1 + x^2)x^2 \in \langle -1 + x^2 \rangle = \langle \text{SF}(p) \rangle$, the equation $p = 0$ is not invariant, however, because $x(t) = \pm e^t$. Notice that the proof rule P-c (with or without the square-free reduction) is unable to prove or disprove the invariance of $p = 0$.

5.3 Square-free Reduction On Differential Radical Invariants (DRI)

Square-free reduction cannot increase the deductive power of the proof rule DRI because its premise is necessary and sufficient to prove invariance of real algebraic sets, which is unaffected by applying SF reduction. However, the computational impact of using square-free reduction with DRI remains an interesting question. Empirically, we observed a better performance of DRI when the SF reduction is applied first. In addition to lowering the degrees of the involved polynomials (as it did for Lie-based proof rules), we observed that the order N_{SF} for $\text{SF}(h)$ is always lower than the order N for h . We, therefore, conjecture $N_{\text{SF}} \leq N$. However, we identified an example for which square-free reduction resulted in a significant ($\times 100$) computational overhead (see [10, Section 5.3]) due to the ideal membership checking (which we perform using Gröbner bases with reverse lexicographic monomial ordering). In our implementation of DRI, called DRI_{opt} in the sequel, we use the square-free reduction only as a pre-processing step for the quantifier elimination problems in the premise of DRI.

Remark 3. Notice that Prop. 14 does not have an analogue for DRI. In other words, if a polynomial equation $h = 0$ is invariant for \mathbf{p} , its irreducible factors need not define invariant equations themselves. Geometrically, this means that if a variety is invariant under the flow of \mathbf{p} , its irreducible components need not be invariants under the flow of \mathbf{p} . For instance, consider the irreducible polynomials $q_1 = y - 1$ and $q_2 = x^2 + (y - 1)^2$. The equation $q_1 q_2 = 0$ which is equivalent to $y = 1$, is invariant for $\mathbf{p} = (1, 0)$, since the premise of the proof rule DRI holds true for $N = 3$. However, the equation $q_2 = 0$, which is equivalent to $x = 0 \wedge y = 1$, is not an invariant equation for \mathbf{p} .

6 Experimental Comparison

We empirically compare the running time performance of all the proof rules discussed in this paper on a heterogeneous collection of 76 invariant varieties (available in [10]). The examples we used originate from a number of sources—many come from textbooks on Dynamical Systems; some from the literature on formal verification of hybrid systems; others have been hand-crafted to exploit sweetspots of certain proof rules. In this section, the prefix SF is implicit for all Lie-based proof rules. We consider 4 equally sized classes of invariant sets: (1) 24 smooth invariants, where Lie is both necessary and sufficient, (2) 17 isolated equilibria as trivial (for humans, not machines) equational invariants for which both Lie° and Lie^* provide necessary and sufficient conditions, (3) 17 other singularities and high integrals, (4) 18 functional invariants, where $\text{DI}_=$ is necessary and sufficient. The most interesting experimental question we seek to address here is whether the greater generality of the more deductively powerful proof rules also comes at a substantially higher computational cost when assessed across the entire spectrum of examples. As a complement to the theoretical deductive power relationships between the different proof rules (Section 4), we also seek to identify some nuances in the complexity of the conditions in the premises, which the coarse-grained complexity bounds miss, being highly sensitive to the number of variables.

From our experiments it emerges that the proof rules exhibit different (and at times surprising) trade-offs between generality and efficiency. Figure 5 compares the number




Figure 5: Experimental performance of proof rules: problems solved per time (log scale)

of invariant varieties that each rule could prove within 60 seconds. The vertical axis shows cumulative time spent on the problems. All runs were performed on an Intel Core i5 1.7GHz machine with 4Gb RAM. Generally, we observe DRI performing very well across the entire spectrum of problem classes. This is very encouraging, but also at first sight appears to defy intuition since it implies that one does not necessarily sacrifice performance when opting to use a more deductively powerful rule. In this graph, we also see that overall Lie^o appears to offer an interesting compromise between deductive power and efficiency—it is able to prove a significant body of problems that are out of scope for Lie, while avoiding the complexity penalty which affects Lie^{*} (due to introducing an extra variable).

A more careful analysis of the benchmarks reveals interesting relationships that are obscured in the “big picture”; to see them, one needs to consider the individual classes of invariants for which some of the sufficient conditions in the rules are in fact *necessary and sufficient*. Together with DRI, this yields two *decision procedures* for each class and allows us to focus only on running time performance and assess practicality of proof rules. In Fig. 6, we observe the rules Lie^o and Lie^{*} performing very well in proving invariance of isolated equilibria. This is to be expected as Lie^o in particular was formulated with this problem class in mind. It is interesting that DRI remains highly competitive here; though its performance is slightly poorer in our set of benchmarks.

It is clear that because proof rules Lie^o and Lie^{*} generalize Lie, they will be able to prove every problem in the smooth invariant benchmarks. The running time performance of the three rules is almost identical, with Lie offering a slight speed-up over its generalizations. The premises of Lie^o and Lie^{*} impose conditions on states in the singular locus, which is the empty set for smooth invariants; this, in practice, appears to be slightly more expensive than checking an equivalent property that the gradient is non-vanishing on the variety (as in the premise of Lie).

The proof rules DI₌ and P-c, corresponding to conditions with historical origins in the study of integrability of dynamical systems, can be seen to perform very well in proving functional invariants, while performing very poorly in benchmarks for isolated equilibria. In proofs of smooth invariants their behaviour is radically different, with DI₌ proving only a handful of examples and P-c succeeding in proving most of the prob-




Figure 6: Number of problems solved per class (log scale).

lems very efficiently. This can be explained by the fact that P-c generalizes $DI_{=}$ and is therefore more deductively powerful. P-c appears slightly slower at proving functional invariants, but shows very impressive running time performance for some problems from the smooth invariant benchmarks, where it is the fastest proof rule for many of problems where it succeeds. Comparing running time performance with DRI, we see that DRI is only slightly slower at proving functional invariants than $DI_{=}$ and P-c. Again, the performance gap between DRI and the two rules appears to be insignificant for most problems. Theoretically, when P-c proves an invariant, DRI applies conditions that are identical to the premise of P-c. Hence, although DRI is a generalization, this does not come at a significant extra cost for the classes where P-c shows good running time performance. The slightly greater running time of DRI compared to that of P-c can be accounted for by the fact that in our implementation DRI computes the Gröbner basis for *every* order N including for $N = 1$ where such computation is unnecessary.

For functional invariants $DI_{=}$ benefits from the fact that the condition in its premise, which requires to show that the Lie derivative evaluates to zero everywhere, is equivalent to showing that the Lie derivative is the zero polynomial, which can be checked very efficiently by symbolic computation, without a decision procedure for real arithmetic.

In the examples featuring singularities and high integrals in the benchmarks we see DRI as the clear winner, simply because there was no other rule that was tailored to work on this class. Indeed, the structure of these invariant sets can be rather involved, making it difficult to characterize in a single proof rule; however, sometimes it is pos-

sible to exploit the structure of high integrals inside a proof system and arrive at very efficient proofs that outperform DRI [11].

It is not surprising that DRI should overtake all the other rules in terms of deductive power (it is, after all, necessary and sufficient); what is remarkable is that the performance we observe for DRI is often very competitive to that of the sufficient rules when they also succeed at a proof. This observation suggests a possible strategy for proof search in a proof system: give precedence to DRI and switch to other sufficient rules if DRI takes longer than some time-out value. The rationale behind this decision is our empirical observation that DRI performs consistently well on all problem classes we considered, but it is also sometimes possible to save time by using a proof rule which is less deductively powerful. It is important to note here that the overall proof system benefits from including the sufficient proof rules, rather than relying solely upon DRI.

7 Related Work

TALY & TIWARI in [27] investigate an approach to proving invariance properties of non-strict polynomial inequalities and closed semi-algebraic sets which inspired our formulation of the proof rules Lie° and Lie^* for real algebraic varieties; we employ the same ideas for reasoning about the singular locus separately and appealing to the Nagumo theorem for the proof of soundness. At least some of the difficulties encountered with inequalities in [27] can be eliminated for real algebraic sets by working only with square-free reduced polynomials; a reduction we perform as a pre-processing step. Indeed, in [27] the authors provide a simple example in which an invariant polynomial equality is encoded as a polynomial inequality of the form $h^2 \leq 0$ (over the reals this is equivalent to $h^2 = 0$) which falls out of scope of their proof rules. Square-free reduction may be extended to polynomial inequalities using order parity decomposition [7] and makes progress possible on similar problems.

The deductive power of the proof rule DI (which generalizes $\text{DI}_=$ to semi-algebraic sets) combined with other proof rules (such as differential cut or differential weakening) have been investigated in [24]. In this work, we focus on sound proof rules for checking invariance properties of algebraic sets and investigate their deductive power as well as their practical efficiency. To our knowledge, this is the first attempt to structure and empirically compare the performance of the proof rules we considered.

8 Conclusions and Future Work

We have theoretically and empirically compared proof rules for checking invariance properties of real algebraic sets in polynomial vector fields. Our work investigated an important aspect of deductive safety verification of continuous and hybrid dynamical systems. Namely, given the abundance of existing sufficient conditions for invariant equations ($\text{DI}_=$, C-c and P-c, Lie), in addition to the extensions of Lie’s criterion, Lie° and Lie^* , and the recently developed *necessary and sufficient* conditions for real algebraic invariants (DRI [9]), it is crucial to know whether the gains in deductive power come at the price of greater computational complexity and poor running time

performance that would hinder practical applications. The work presented in this paper leads us to arrive at the following conclusions:

- Empirically, we observe that the most deductively powerful rule (DRI) performs very well in checking invariance of polynomial equalities.
- P-c is made redundant by DRI (DRI strictly increases the deductive power of P-c while being equally efficient).
- Reducing polynomials to square-free form is always of benefit to the proof rule Lie and its generalizations, where it yields improvements in both the deductive power and the running time performance.
- We proved that combining SF with the proof rules $DI_{=}$ and C-c yields new *incomparable* proof rules, whereas SF with P-c is as powerful as P-c alone.
- Performing square-free reduction of an invariant candidate may introduce a performance penalty for DRI and therefore cannot be regarded as an optimization.

It is our hope to extend this work to similarly study proof methods for invariance of semi-algebraic sets in polynomial vector fields. This problem is of fundamental importance to verification of continuous and hybrid systems [20,22] and a better understanding of the factors affecting proof rule efficiency has the potential to be of considerable practical utility. There are currently three available methods that have been proposed for checking invariance of semi-algebraic sets: the method of differential invariants due to Platzer [25], a characterization of invariant semi-algebraic sets due to Liu et al. [15] and a method for closed semi-algebraic sets based on the Nagumo theorem proposed by Taly & Tiwari [27]. The latter approach can unfortunately be shown to be unsound (we identify the problem in [10, Appendix B]); however, this deficiency can be fixed. It would be very interesting to extend the work presented in this paper to investigate the relationship between deductive power and running time performance in the aforementioned methods. We leave this for future work.

Acknowledgments The authors would like to thank Dr. Ashish Tiwari at SRI International for his kind and informative response to our technical queries and extend special thanks to Dr. Paul B. Jackson at the LFCS, University of Edinburgh, for his valuable help in improving the manuscript.

References

1. Basu, S., Pollack, R., Roy, M.F.: On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM* 43(6), 1002–1045 (1996)
2. Collins, G.E.: Hauptvortrag: Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In: *Automata Theory and Formal Languages*. LNCS, vol. 33, pp. 134–183. Springer (1975)
3. Collins, G.E., Hong, H.: Partial cylindrical algebraic decomposition for quantifier elimination. *J. Symb. Comput.* 12(3), 299–328 (Sep 1991)
4. Cox, D.A., Little, J., O’Shea, D.: *Ideals, Varieties, and Algorithms - an introduction to computational algebraic geometry and commutative algebra* (2. ed.). Springer (1997)
5. Darboux, J.G.: Mémoire sur les équations différentielles algébriques du premier ordre et du premier degré. *Bulletin des Sciences Mathématiques et Astronomiques* 2(1), 151–200 (1878), <http://eudml.org/doc/84988>

6. Davenport, J.H., Heintz, J.: Real quantifier elimination is doubly exponential. *J. Symb. Comput.* 5(1/2), 29–35 (1988)
7. Dolzmann, A., Sturm, T.: Simplification of quantifier-free formulas over ordered fields. *Journal of Symbolic Computation* 24, 209–231 (1995)
8. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In: *ISSAC*. pp. 75–83. ACM, New York, NY, USA (2002)
9. Ghorbal, K., Platzer, A.: Characterizing algebraic invariants by differential radical invariants. In: *TACAS*. vol. 8413, pp. 279–294. Springer (2014)
10. Ghorbal, K., Sogokon, A., Platzer, A.: A hierarchy of proof rules for checking differential invariance of algebraic sets. Tech. Rep. CMU-CS-14-140, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA (11 2014), <http://reports-archive.adm.cs.cmu.edu/anon/2014/abstracts/14-140.html>
11. Ghorbal, K., Sogokon, A., Platzer, A.: Invariance of conjunctions of polynomial equalities for algebraic differential equations. In: *SAS*. LNCS, vol. 8723, pp. 151–167. Springer (2014)
12. Goriely, A.: *Integrability and Nonintegrability of Dynamical Systems*. Advanced series in nonlinear dynamics, World Scientific (2001)
13. Lie, S.: *Vorlesungen über kontinuierliche Gruppen mit Geometrischen und anderen Anwendungen*. Teubner, Leipzig (1893)
14. Lindelöf, E.: Sur l'application de la méthode des approximations successives aux équations différentielles ordinaires du premier ordre. *Comptes rendus hebdomadaires des séances de l'Académie des sciences* 116, 454–458 (1894)
15. Liu, J., Zhan, N., Zhao, H.: Computing semi-algebraic invariants for polynomial dynamical systems. In: *EMSOFT*. pp. 97–106. ACM (2011)
16. Matringe, N., Moura, A.V., Rebiha, R.: Generating invariants for non-linear hybrid systems by linear algebraic methods. In: *SAS*. LNCS, vol. 6337, pp. 373–389. Springer (2010)
17. Mayr, E.W.: Membership in polynomial ideals over \mathbb{Q} is exponential space complete. In: Monien, B., Cori, R. (eds.) *STACS*. LNCS, vol. 349, pp. 400–406. Springer (1989)
18. Nagumo, M.: Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen (in German). In: *Proceedings of the Physico-Mathematical Society of Japan*. vol. 24, pp. 551–559 (May 1942)
19. Olver, P.J.: *Applications of Lie Groups to Differential Equations*. Springer (2000)
20. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reasoning* 41(2), 143–189 (2008)
21. Platzer, A.: Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.* 20(1), 309–352 (2010)
22. Platzer, A.: The complete proof theory of hybrid systems. In: *LICS*. pp. 541–550. IEEE (2012)
23. Platzer, A.: A differential operator approach to equational differential invariants - (invited paper). In: *ITP*. LNCS, vol. 7406, pp. 28–48. Springer (2012)
24. Platzer, A.: The structure of differential invariants and differential cut elimination. *Logical Methods in Computer Science* 8(4), 1–38 (2012)
25. Platzer, A., Clarke, E.M.: Computing differential invariants of hybrid systems as fixedpoints. In: *CAV*. LNCS, vol. 5123, pp. 176–189. Springer (2008)
26. Sankaranarayanan, S., Sipma, H.B., Manna, Z.: Constructing invariants for hybrid systems. *Form. Methods Syst. Des.* 32(1), 25–55 (2008)
27. Taly, A., Tiwari, A.: Deductive verification of continuous dynamical systems. In: *FSTTCS*. LIPIcs, vol. 4, pp. 383–394 (2009)
28. Tarski, A.: A decision method for elementary algebra and geometry. *Bull. Amer. Math. Soc.* 59 (1951)
29. Tiwari, A.: Abstractions for hybrid systems. *Form. Methods Syst. Des.* 32(1), 57–83 (2008)
30. Walter, W.: *Ordinary Differential Equations*. Springer New York (1998)