

Bounded Quantifier Instantiation for Checking Inductive Invariants

Yotam M. Y. Feldman¹, Oded Padon¹, Neil Immerman², Mooly Sagiv¹, and Sharon Shoham¹

¹ Tel Aviv University, Tel Aviv, Israel

² UMass, Amherst, USA

Abstract. We consider the problem of checking whether a proposed invariant φ expressed in first-order logic with quantifier alternation is *inductive*, i.e. preserved by a piece of code. While the problem is undecidable, modern SMT solvers can sometimes solve it automatically. However they employ powerful quantifier instantiation methods that may diverge, especially when φ is not preserved. A notable difficulty arises due to counterexamples of infinite size.

This paper studies *Bounded-Horizon instantiation*, a natural method for guaranteeing the termination of SMT solvers. The method bounds the depth of terms used in the quantifier instantiation process. We show that this method is surprisingly powerful for checking quantified invariants in uninterpreted domains. Furthermore, by producing partial models it can help the user diagnose the case when φ is not inductive, especially when the underlying reason is the existence of infinite counterexamples.

Our main technical result is that Bounded-Horizon is at least as powerful as *instrumentation*, which is a manual method to guarantee convergence of the solver by modifying the program so that it admits a purely universal invariant. We show that with a bound of 1 we can simulate a natural class of instrumentations, without the need to modify the code and in a fully automatic way. We also report on a prototype implementation on top of Z3, which we used to verify several examples by Bounded-Horizon of bound 1.

1 Introduction

This paper addresses a fundamental problem in automatic program verification: how to prove that a piece of code preserves a given invariant. In Floyd-Hoare style verification this means that we want to automatically prove the validity of the Hoare triple $\{P\}C\{P\}$ where P is an assertion and C is a command. Alternatively, this can be shown by proving the unsatisfiability of the formula $P(V) \wedge \delta(V, V') \wedge \neg P(V')$ (the *verification condition*) where $P(V)$ denotes the assertion P before the command, $P(V')$ denotes the assertion P after the command, and $\delta(V, V')$ is a formula expressing the meaning of the command C as a transition relation between pre- and post-states. When C is a loop body, such a P is an inductive invariant and can be used to prove safety properties of the loop (if it also holds initially and implies the desired property).

For programs with infinite state space, proving the validity of $\{P\}C\{P\}$ is generally undecidable even when C does not include loops. Indeed, existing SMT solvers

can diverge even for simple assertions and simple commands. Recent attempts to apply program verification to prove the correctness of critical system’s design and code [15] identify this as the main hurdle for using program verification.

The difficulty is rooted in powerful constructs used in SMT-based verification of interesting programs. Prominent among these constructs are arithmetic and other program operations modeled using background theories, and logical quantifiers. In this paper we target the verification of applications in which the problem can be modeled without interpreted theories. This is in line with recent works that show that although reasoning about arithmetic is crucial for low-level code, in many cases the verification of high-level programs and designs can be performed by reasoning about quantification in uninterpreted theories. Specifically, the decidable Effectively Propositional logic (EPR) has been successfully applied to domains such as linked-list manipulation [20], Software-Defined Networks [5] and some distributed protocols [28]. Without interpreted theories it remains to address the complications induced by the use of quantifier alternation.

In the presence of quantifier alternation, the solver’s ability to check assertions is hindered by the following issues: (1) an infinite search space of proofs that needs to be explored for correct assertions, a problem which is sometimes manifested in matching loops [11], and (2) a difficulty of finding counterexamples for invalid assertions, notably when counterexamples may be of infinite size. Current SMT techniques often fail to produce models of satisfiable quantified formulas [14,33], which is somewhat unfortunate since one of the main values of program verification is early detection of flaws in designs and programs. The existence of infinite counterexamples is a major complication as they are difficult to find. In uninterpreted domains, infinite counterexamples usually do not indicate a real violation of the verification conditions and are counterintuitive to programmers, yet render assertions invalid in the context of general first-order logic (on which SMT proof techniques are based). Hence infinite counter-models pose a real problem in the verification process.

Previous work on EPR [20,5,28] used universally quantified invariants with programs expressed by $\exists^*\forall^*$ formulas³, in which case checking inductive invariants is decidable, hence problems (1) and (2) do not occur. In particular, EPR enjoys the finite-model property and so counterexamples are of finite size. EPR programs are in fact Turing-complete [28], but universal invariants are not always sufficient to express the required program properties.

For example, [15] describes a client server scenario where the invariant is “For every reply message sent by the server, there exists a corresponding request message sent by a client”. (See Example 1 for further details.) This invariant is $\forall^*\exists^*$ and thus leads to verification conditions with quantifier alternation. This kind of quantifier alternation may lead to divergence of the solver as problems (1) and (2) re-emerge.

The current work aims to expand the applicability of the EPR-based verification approach to invariants of more complex quantification. We focus on the class $\forall^*\exists^*$ invariants, which arise in interesting programs. As we show, checking inductiveness of invariants in this class is undecidable. We thus study problems (1),(2) above for this

³ Automated tools that extract EPR transition relation from code exist for C code manipulating linked lists [20,21,22] and for the modeling language RML [28] which is Turing-complete.

setting using the notion of *bounded quantifier instantiations*, which we term *Bounded-Horizon*.

Main results. This paper explores the utility of limited quantifier instantiations for checking $\forall^*\exists^*$ invariants, and for dealing with the problems that arise from quantifier alternation: divergence of the proof search and infinite counter-models.

We consider instantiations that are *bounded in depth* of terms. Bounded instantiations trivially prevent divergence while maintaining soundness. Although for a given bound the technique is not complete, i.e. unable to prove every correct invariant, we provide completeness guarantees by comparing bounded instantiations to the method of *instrumentation*, a powerful technique implicitly employed in previous works [20,22,28]. Instrumentation tackles a $\forall^*\exists^*$ invariant by transforming the program in a way that allows the invariant to be expressed in a universal form, and, accordingly, makes the verification conditions fall in EPR. We show that for invariants that can be proven using a typical form of instrumentation, bounded instantiations of a small bound are also complete. Namely, they are sufficiently powerful to prove the original program without modifications and in a fully automatic way. This is encouraging since instrumentation is labor-intensive and error-prone while bounded instantiations are completely automatic.

This result suggests that in many cases correct $\forall^*\exists^*$ invariants of EPR programs can be proven using a simple proof technique. Typically in such cases tools such as Z3 will also manage to automatically prove the verification conditions. However, bounded instantiations guarantee termination a-priori even when the invariant is not correct. When it terminates, the procedure returns a logical structure which is not necessarily a true counterexample but “approximates” it, as it satisfies all the bounded instantiations. Interestingly, this suggests a way to overcome the problem of infinite models. This problem arises when the user provides an invariant that is correct for finite models but is incorrect in general first-order logic. In such cases, state-of-the-art SMT solvers typically produce “unknown” or timeout since they fail to find infinite models. Thus the user is left with very little aid from the solver when attempting to make progress and successfully verify the program. In contrast, bounded quantifier instantiation can be used to find finite models with increasing sizes, potentially indicating the existence of an infinite model, and provide hints as to the source of the error. This information allows the user to modify the program or the invariant to exclude the problematic models. We demonstrate this approach on a real example in which such a scenario occurred in one of our verification attempts. We show that the provided models assist in identifying and fixing the error, allowing the user to successfully verify the program.

We also implemented a prototype tool that performs bounded instantiations of bound 1, and used it to verify several distributed protocols and heap manipulating programs. The implementation efficiently reduces the problem of checking inductiveness with bound 1 to a Z3 satisfiability check on which the solver always terminates, thereby taking advantage of Z3’s instantiation techniques while guaranteeing termination.

2 Preliminaries

In this section we provide background and explain our notation. Σ will always denote a relational first-order vocabulary, which may contain constant symbols, c_i , and relation

symbols, r_j , but no function symbols. For a formula φ we denote by $\text{const}[\varphi]$ the set of constants that appear in φ . We write that $\varphi \in \exists^*(\Sigma)$ to mean that φ is an *existential* formula defined over vocabulary Σ . Similarly, the class of *universal* formulas is denoted by $\forall^*(\Sigma)$. We say that φ is *quantifier-free*, denoted $\varphi \in \text{QF}(\Sigma)$ if it contains no quantifiers, and that it is *alternation free*, denoted $\varphi \in \text{AF}(\Sigma)$, if it can be written as a Boolean combination of formulas in $\exists^*(\Sigma)$. $\text{FOL}(\Sigma)$ stands for arbitrary first-order formulas over Σ . A *sentence* is a closed formula.

EPR. The effectively-propositional (EPR) fragment of first-order logic, also known as the Bernays-Schönfinkel-Ramsey class, consists of $\exists^*\forall^*(\Sigma)$ sentences. Such sentences enjoy the *small model property*, in fact a satisfiable EPR sentence has a model of size no larger than the number of its constants plus existential quantifiers. Thus satisfiability of EPR sentences is decidable [30].

EPR Transition Relation. We specify a transition relation via an EPR sentence, δ , over a vocabulary $\Sigma \uplus \Sigma'$ where Σ is a relational vocabulary used to describe the source state of a transition and $\Sigma' = \{a' \mid a \in \Sigma\}$ is used to describe the target state.

Inductive Invariants. A first-order sentence I over Σ is an *inductive invariant* for δ if $I \wedge \delta \rightarrow I'$ is valid, or, equivalently, if $I \wedge \delta \wedge \neg I'$ is unsatisfiable⁴, where I' results from substituting every constant and relation symbol in I by its primed version (i.e. $I' \in \text{FOL}(\Sigma')$).

Skolemization. Let $\varphi(z_1, \dots, z_n) \in \text{FOL}(\Sigma)$. The *Skolemization* of φ , denoted φ_S , is a universal formula over $\Sigma \uplus \Sigma_S$, where Σ_S consists of fresh constant symbols and function symbols, obtained as follows. We first convert φ to negation normal form (NNF) using the standard rules. For every existential quantifier $\exists y$ that appears under the scope of the universal quantifiers $\forall x_1, \dots, \forall x_m$, we introduce a fresh function symbol $f_y \in \Sigma_S$ of arity $n + m$. We replace each bound occurrence of y by $f_y(z_1, \dots, z_n, x_1, \dots, x_m)$, and remove the existential quantifier. If $n + m = 0$ (i.e., φ has no free variables and $\exists y$ does not appear in the scope of a universal quantifier) a fresh constant symbol is used to replace y . It is well known that $\varphi_S \rightarrow \varphi$ is valid and φ_S and φ are equi-satisfiable.

3 Bounded-Horizon

In this section, we define a systematic method of quantifier instantiation called *Bounded-Horizon* as a way of checking the inductiveness of first-order logic formulas, and explore some of its basic properties. We start with the undecidability of the problem.

Undecidability of inductiveness. For a universal formula $I \in \forall^*(\Sigma)$, checking inductiveness amounts to checking unsatisfiability of an EPR formula, and is therefore decidable. The same holds for $I \in \text{AF}(\Sigma)$. However, this is no longer true when quantifier alternation is introduced. For example, checking inductiveness of $I \in \forall^*\exists^*(\Sigma)$ amounts to checking unsatisfiability of a formula in a fragment for which satisfiability is undecidable. In fact we prove that:

Theorem 1. *The problem of determining on input $I \in \forall^*\exists^*(\Sigma)$ and $\delta \in \exists^*\forall^*(\Sigma, \Sigma')$, whether I is an inductive invariant for δ , is undecidable.*

⁴ In this paper, satisfiability and validity refer to general models and are not restricted to finite models. Note that for EPR formulas, finite satisfiability and general satisfiability coincide.

Proof Sketch. The proof is by reduction from the halting problem, which can be encoded using a $\forall^*\exists^*$ formula via tiling (see e.g. [19]). For the setting of checking invariants, we start with a Turing machine M , and construct $\delta \in \exists^*\forall^*(\Sigma, \Sigma')$, and $I \in \forall^*\exists^*(\Sigma)$ s.t. I is an inductive invariant for δ iff M halts on the empty tape. In case M does not halt, the counter-model that shows that I is not inductive is an infinite structure which encodes an infinite run of M . \square

Bounded-Horizon Instantiations. Let $\delta \in \exists^*\forall^*(\Sigma, \Sigma')$ be an EPR transition relation and $I \in \text{FOL}(\Sigma)$ a candidate invariant. We would like to check the satisfiability of $I \wedge \delta \wedge \neg I'$, and equivalently of $\text{Ind} = I_S \wedge \delta_S \wedge (\neg I')_S$. Recall that φ_S denotes the Skolemization of φ , and note that I_S and $(\neg I')_S$ possibly add Skolem functions to the vocabulary. Roughly speaking, for a given $k \in \mathbb{N}$, Bounded-Horizon instantiates the universal quantifiers in Ind , while restricting the instantiations to produce ground-terms of function nesting at most k . We then check if this (finite) set of instantiations is unsatisfiable; if it is already unsatisfiable then we have a proof that I is inductive. Otherwise we report that I is not known to be inductive. As we will show, this algorithm is sound but not necessarily complete for a given k .

Below we provide the formal definitions. We start with the notion of instantiations, and recall Herbrand's theorem which establishes completeness of proof by (unrestricted) instantiations. Suppose that some vocabulary $\tilde{\Sigma}$ including constants and function symbols is understood (e.g., $\tilde{\Sigma} = \Sigma \uplus \Sigma_S$, where Σ_S includes Skolem constants and function symbols).

Definition 1 (Instantiation). Let $\varphi(\bar{x}) \in \forall^*(\tilde{\Sigma})$ be a universal formula with n free variables and m universal quantifiers. An instantiation of φ by a tuple \bar{t} of $n + m$ ground terms, denoted by $\varphi[\bar{t}]$, is obtained by substituting \bar{t} for the free variables and the universally quantified variables, and then removing the universal quantifiers.

Note that an instantiation is a quantifier-free sentence.

Theorem 2 (Herbrand's Theorem). Let $\varphi \in \forall^*(\tilde{\Sigma})$. Then φ is satisfiable iff the (potentially infinite) set $\{\varphi[\bar{t}] \mid \bar{t} \text{ is a tuple of ground terms over } \tilde{\Sigma}\}$ is satisfiable.

We now turn to restrict the depth of terms used in instantiations.

Definition 2 (Bounded-Depth Terms). For every $k \in \mathbb{N}$, we define BHT_k to be the set of ground terms over $\tilde{\Sigma}$ with function symbols nested to depth at most k . BHT_k is defined by induction over k . Let C be the set of constants in $\tilde{\Sigma}$, F the set of functions, and for every $f \in F$ let Arity_f be the arity of f . Then

$$\text{BHT}_0 = C$$

$$\text{BHT}_k = \text{BHT}_{k-1} \cup \{f(t_1, \dots, t_m) \mid f \in F, m = \text{Arity}_f, t_1, \dots, t_m \in \text{BHT}_{k-1}\}.$$

We will also write $\bar{t} \in \text{BHT}_k$ for a tuple of terms \bar{t} , to mean that every entry of \bar{t} is in BHT_k (the number of elements in \bar{t} should be clear from the context). Note that the set of ground terms is $\text{BHT}_\infty = \bigcup_{k \in \mathbb{N}} \text{BHT}_k$.

Definition 3 (Depth of Instantiation). Let $\varphi \in \forall^*(\tilde{\Sigma})$ and $\bar{t} \in \text{BHT}_\infty$. The depth of instantiation, denoted $\text{depth}(\varphi[\bar{t}])$, is the smallest k such that all ground terms that appear in $\varphi[\bar{t}]$ are included in BHT_k .

Bounded-Horizon algorithm. Given a candidate invariant $I \in \text{FOL}(\Sigma)$, a transition relation δ over $\Sigma \uplus \Sigma'$, and $k \in \mathbb{N}$, the Bounded-Horizon algorithm constructs the formula $\text{Ind} = I_S \wedge \delta_S \wedge (\neg I')_S$, and checks if the set

$$\{\text{Ind}[\bar{t}] \mid \bar{t} \in \text{BHT}_k, \text{depth}(\text{Ind}[\bar{t}]) \leq k\} \quad (1)$$

is unsatisfiable. If it is, then I is provably inductive w.r.t. δ with Bounded-Horizon of bound k . Otherwise we report that I is not known to be inductive.

Note that the satisfiability check performed by Bounded-Horizon is decidable since the set of instantiations is finite, and each of them is a ground formula.

Bounded-Horizon for $\forall^*\exists^*$ Invariants. We illustrate the definition of Bounded-Horizon in the case that $I \in \forall^*\exists^*(\Sigma)$. Assume that $I = \forall \bar{x}. \exists \bar{y}. \alpha(\bar{x}, \bar{y})$ where $\alpha \in \text{QF}$. Then $I_S = \forall \bar{x}. \alpha(\bar{x}, \bar{f}(\bar{x}))$ where \bar{f} are new Skolem function symbols. δ_S introduces Skolem constants but no function symbols, and in this case so does $(\neg I')_S$. Bounded-Horizon check of bound k can be approximately understood as checking the satisfiability of

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_{k-1}} I_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_k} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_k} (\neg I')_S[\bar{t}] \right). \quad (2)$$

(In fact, it is possible that I_S contains sub-formulas for which instantiations of depth k do not increase the total depth of instantiations beyond k , and are thus also included.)

Lemma 1 (Soundness). For every $k \in \mathbb{N}$, Bounded-Horizon with bound k is sound, i.e., if Bounded-Horizon of bound k reports that $I \in \text{FOL}(\Sigma)$ is inductive w.r.t. δ , then I is indeed inductive.

Proof. Assume that I is not inductive w.r.t. δ , so there is a structure \mathcal{A} such that $\mathcal{A} \models I_S \wedge \delta_S \wedge (\neg I')_S$. In particular $\mathcal{A} \models \text{Ind}[\bar{t}]$ for every $\bar{t} \in \text{BHT}_\infty$ and in particular for every $\bar{t} \in \text{BHT}_k$ such that $\text{depth}(\text{Ind}[\bar{t}]) \leq k$. Hence, Bounded-Horizon of bound k will not report that I is inductive. \square

Example 1. Fig. 1 presents a simple model of the client server scenario described in [15]. The program induces an EPR transition relation, and its invariant is provable by Bounded-Horizon of bound 1.

We first explain this example ignoring the annotations denoted by “/@”. The system state is modeled using three binary relations. The *req* relation stores pairs of users and requests, representing requests sent by users. The *resp* relation similarly stores pairs of users and replies, representing replies sent back from the server. The *match* relation maintains the correspondence between a request and its reply.

The action `new_request` models an event where a user u sends a new request to the server. The action `respond` models an event where the server responds to a pending request by sending a reply to the user. The request and response are related by the *match* relation. The action `check` is used to verify the safety property that every response sent by the server has a matching request, by aborting the system if this does not hold.

```

req := ∅; resp := ∅; match := ∅;
action new_request(u) {
  q := new request;
  req := req ∪ {(u, q)}
  /@ r := r ∪ {(u, y) | match(q, y)}
}
action respond(u, q) {
  assume req(u, q);
  p := new response;
  match := match ∪ {(q, p)};
  /@ r := r ∪ {(x, p) | req(x, q)}
  resp := resp ∪ {(u, p)}
}

action check(u, p) {
  if resp(u, p) ∧ ∀q. req(u, q) → ¬match(q, p)
  /@ ⇐ if resp(u, p) ∧ ¬r(u, p)
  then abort
}

Invariant I = ∀u, p. resp(u, p) →
  ∃q. req(u, q) ∧ match(q, p)
/@ r(x, y) ≡ ∃z. req(x, z) ∧ match(z, y)
/@ Invariant  $\hat{I} = \forall u, p. \text{resp}(u, p) \rightarrow r(u, p)$ 

```

Fig. 1. Example demonstrating a $\forall^*\exists^*$ invariant that is provable with bound 1. The reader should first ignore the instrumentation code denoted by $/@$ (see Example 3). This example models a simple client-server scenario, with the safety property that every response sent by the server was triggered by a request from a client. Verification of this example requires a $\forall^*\exists^*$ invariant. This example is inspired by [15]. The complete program is provided in [1] (files `client.server.ae.ivy`, `client.server.instr.ivy`).

A natural inductive invariant for this system is

$$I = \forall u, p. \text{resp}(u, p) \rightarrow \exists q. \text{req}(u, q) \wedge \text{match}(q, p).$$

The invariant proves that the `then` branch in action `check` will never happen and thus the system will never abort. This invariant is preserved under execution of all actions, and is provable by Bounded Horizon of bound 1.

Lemma 2 (Completeness for some k). *For every $I \in \text{FOL}(\Sigma)$ and δ such that I is inductive w.r.t. δ there exists a finite $k \in \mathbb{N}$ s.t. I is provably inductive w.r.t. δ with Bounded-Horizon of bound k .*

Proof. From Theorem 2 and compactness there is a finite set S of instantiations that is unsatisfiable. Take k to be the maximal depth of instantiations in S . \square

For example, if $I \in \forall^*$ then Bounded-Horizon of bound 0 is complete. However, as expected due to the undecidability of checking inductiveness, for arbitrary invariants Bounded-Horizon is *not* necessarily complete for a given k :

Example 2. An example of a program and an inductive invariant for which a bound of 0 or 1 is insufficient appears in Fig. 2. In this example the server operates as a middleman between clients and the database (DB), and is used to anonymize user requests before they reach the database. The server performs a translation p between clients' identity and an anonymous unique id, sends a translated request to the DB, and forwards the DB's response to the clients. The safety property is that every response sent by the server was triggered by a request from a client. The inductive invariant states, in addition to the safety property, that every server request to the DB was triggered by a client's request from the server, and that every DB response was triggered by a server's request. Proving that the invariant is inductive under the action `server.recv_db_response` requires

```

req := ∅; resp := ∅;
db_req := ∅; db_resp := ∅; t := ∅;
action new_request(u) {
  q := new request;
  req := req ∪ {(u, q)}
}
action db_recv_request(id, q) {
  assume db_req(id, q);
  p := DB response ensuring db(q, p);
  db_resp := db_resp ∪ {(id, p)}
}
action check(u, p) {
  assume resp(u, p);
  if ∀q. req(u, q) → ¬db(q, p)
  then abort
}
action server_recv_request(u, q) {
  assume req(u, q);
  id := new DB request id;
  t := t ∪ {(id, u)};
  db_req := db_req ∪ {(id, q)}
}
action server_recv_db_response(id, p) {
  assume db_resp(id, p);
  resp := resp ∪ {(x, p) | t(id, x)}
}

```

$$\begin{aligned}
\text{Invariant } I = & \forall u, p. \text{resp}(u, p) \rightarrow \exists q. \text{req}(u, q) \wedge \text{db}(q, p) \wedge \\
& \forall id, q. \text{db_req}(id, q) \rightarrow \exists u. t(id, u) \wedge \text{req}(u, q) \wedge \\
& \forall id, p. \text{db_resp}(id, p) \rightarrow \exists q. \text{db_req}(id, q) \wedge \text{db}(q, p) \wedge \\
& \forall id, u_1, u_2. t(id, u_1) \wedge t(id, u_2) \rightarrow u_1 = u_2
\end{aligned}$$

Fig. 2. Example demonstrating a $\forall^*\exists^*$ invariant that is provable only with bound 2. The server anonymizes requests from clients to the database (DB) and forwards the answer to the client. The server performs a translation t between clients’ identity and an anonymous unique id. The safety property is that every response sent by the server to a client was triggered by a request from the client. The inductive invariant further states that every server request to the DB was triggered by a client’s request from the server, and that every DB response was triggered by a server’s request. The complete program corresponding to this Figure appears in [1] (file `client_server_db_ae.ivy`).

the prover to understand that for the response from the DB there is a matching request from the server to the DB, and that for this request to the DB there is a matching request from the client to the server. Every such translation requires another nested instantiation. In this example, a bound of 2 manages to prove inductiveness. This example can be lifted to require an even larger depth of instantiations by adding more translation entities similar to the server, and describing the invariant in a similar, modular, way.

Small Bounded-Horizon for $\forall^*\exists^*$ Invariants. Despite the incompleteness, we conjecture that a small depth of instantiations typically suffices to prove inductiveness. The intuition is that an EPR transition relation has a very limited “horizon” of the domain: it interacts only with a small fraction of the domain, namely elements pointed to by program variables (that correspond to logical constants in the vocabulary).

When performing the Bounded-Horizon check with bound 1 on a $\forall^*\exists^*$ invariant $I = \forall \bar{x}. \exists \bar{y}. \alpha(\bar{x}, \bar{y})$, we essentially assume that the existential part of the invariant $\psi(\bar{x}) = \exists \bar{y}. \alpha(\bar{x}, \bar{y})$ holds on all program variables — but not necessarily on all elements of the domain — and try to prove that it holds on all elements of the domain after the transition. We expect that for most elements of the domain, the correctness of ψ is maintained simply because they were not modified at all by the transition. For elements

that are modified by the transition, we expect the correctness after modifications to result from the fact that ψ holds for the elements of the domain that the transition directly interacts with. If this is indeed the reason that ψ is maintained, a bound of 1 sufficiently uses ψ in the pre-state to prove the invariant in the post-state, i.e. it is inductive.

This is the case in Example 1. Additional examples are listed in Section 6. The example of Fig. 2 itself also admits a different invariant that is provable by bound 1.

4 Power of Bounded-Horizon for Proving Inductiveness

We now turn to investigate the ability of Bounded-Horizon to verify inductiveness. In this section we provide sufficient conditions for its success by relating it to the notion of instrumentation (which we explain below). We show that Bounded-Horizon with a low bound of 1 or 2 is as powerful as a natural class of sound program instrumentations, those that do not add existential quantifiers. Section 6 demonstrates the method’s power on several interesting programs we verified using Bounded-Horizon of bound 1.

4.1 Instrumentation

We present our view of the instrumentation procedure used in previous works [20,22,28] to eliminate the need for quantifier-alternation, thus reducing the verification task to a decidable fragment. The procedure begins with a program that induces a transition relation $\delta \in \exists^*\forall^*(\Sigma \cup \Sigma')$. The purpose of instrumentation is to modify δ into another transition relation $\widehat{\delta}$ that admits an inductive invariant with simpler quantification (e.g., universal, in which case it is decidable to check). We note that instrumentation is generally a manual procedure. For simplicity, we describe the instrumentation process informally, but provide the semantic soundness requirement in Definition 4. The instrumentation procedure consists of the following three steps:

1. Identify a formula $\psi(\bar{x}) \in \text{FOL}(\Sigma)$ (usually ψ will be existential) that captures information that is needed in the inductive invariant. Extend the vocabulary with an instrumentation relation $r(\bar{x})$ that intentionally should capture the derived relation defined by $\psi(\bar{x})$. Let $\widehat{\Sigma} = \Sigma \cup \{r\}$ denote the extended vocabulary⁵.
2. Add update code that updates r when the original (“core”) relations are modified, and maintains the meaning of r as encoding ψ . The update code must not block executions of real code, and can possibly be a sound approximation. Sometimes it can be generated automatically via finite differencing [31].
3. Modify the program to use r . Often this is performed by rewriting some program conditions, keeping in mind that r encodes ψ . This means replacing some quantified expressions by uses of r .

Example 3. In the example of Fig. 1, to achieve a universal invariant we add an instrumentation relation r defined by $r(x, y) \equiv \exists z. req(x, z) \wedge match(z, y)$ (step 1). The simple form of ψ allows us to obtain precise update code, which appears as annotations marked with /@ in lines that mutate *req* and *match* (step 2). We also replace the

⁵ It is also possible to instrument the program with constants rather than relations. This can be emulated by adding a unary relation $c(x)$ representing the constant, and adding the assumption that c contains exactly one element to the invariant. This is also aligned with the conditions of Theorem 5.

if condition in the action check by an equivalent condition that uses r (step 3). The line marked with /@ \leftrightarrow in the check action replaces the line above it. The resulting program has the invariant $\widehat{I} = \forall u, p. \text{resp}(u, p) \rightarrow r(u, p)$, which is universal.

Let $\widehat{\delta} \in \exists^* \forall^* (\widehat{\Sigma} \cup \widehat{\Sigma}')$ denote the transition relation induced by the modified program (modifications occur in steps 2,3). The soundness of the instrumentation procedure is formalized in the following connection between ψ , δ , and $\widehat{\delta}$:

Definition 4 (Sound Instrumentation). $\widehat{\delta} \in \exists^* \forall^* (\widehat{\Sigma} \cup \widehat{\Sigma}')$ is a sound instrumentation for $\delta \in \exists^* \forall^* (\Sigma \cup \Sigma')$ and $\psi \in \text{FOL}(\Sigma)$ if $(\forall \bar{x}. r(\bar{x}) \leftrightarrow \psi(\bar{x}) \wedge \delta \wedge \forall \bar{x}. r'(\bar{x}) \leftrightarrow \psi'(\bar{x})) \rightarrow \widehat{\delta}$ is valid, or equivalently, $\delta \rightarrow \widehat{\delta}[\psi/r, \psi'/r']$ is valid.

Definition 4 ensures that the instrumented program includes at least all the behaviors of the original program, when r is interpreted according to ψ . Thus, if the instrumented program is safe, then it is sound to infer that the original program is safe.

Remark 1. In the expression $\widehat{\delta}[\psi/r, \psi'/r']$ the update code of r in $\widehat{\delta}$ becomes a constraint over the core relations in Σ . In a sound instrumentation this constraint is required to follow from the way the core relations are updated by δ , essentially meaning that r is updated in a way that is consistent with its interpretation as ψ .

The instrumentation procedure does not require the user to know an inductive invariant for the original program. However, if a sound instrumentation which leads to an invariant exists, then an inductive invariant for the original δ can be produced by substituting back the “meaning” of r as ψ (thus, safety of the original program is implied):

Lemma 3. Let $\widehat{\delta}$ be a sound instrumentation for δ and ψ , and $\widehat{I} \in \text{FOL}(\widehat{\Sigma})$ be an inductive invariant for $\widehat{\delta}$. Then $I = \widehat{I}[\psi/r]$ is inductive w.r.t. δ .

Proof. $\widehat{I} \wedge \widehat{\delta} \rightarrow \widehat{I}'$ is valid, thus, so is $(\widehat{I} \wedge \widehat{\delta} \rightarrow \widehat{I}')[\psi/r, \psi'/r']$. $\widehat{\delta}$ is a sound instrumentation for δ , so (using Definition 4) $I \wedge \delta \rightarrow I'$ is valid. \square

Note that typically the quantification structure of I is more complex than that of \widehat{I} .

Instrumentation without additional existential quantifiers. In order to relate instrumentation to Bounded-Horizon instantiations, we consider the typical case where the instrumentation process of δ does not add new existential quantifiers to $\widehat{\delta}$. This happens when the update code does not introduce additional existential quantifiers. Formally:

Definition 5 (Existential Naming). Let $\widehat{\delta} = \exists z_1, \dots, z_m. \varphi(z_1, \dots, z_m)$ where $\varphi \in \forall^* (\widehat{\Sigma}, \widehat{\Sigma}')$. An existential naming η for $(\widehat{\delta}, \delta)$ is a mapping $\eta : \{z_1, \dots, z_m\} \rightarrow \text{const}[\delta_S] \cup \text{const}[\widehat{\delta}_S]$. We define $\eta(\widehat{\delta})$ to be $\varphi[\eta(z_1)/z_1, \dots, \eta(z_m)/z_m]$.

An existential naming provides a Skolemization procedure which uses existing constants rather than fresh ones. If such η exists, it maps the (Skolemized) existential quantifiers in $\widehat{\delta}$ to their counterparts in δ . For example, the instrumentation in Fig. 1 results in $\widehat{\delta}$ that has an existential naming w.r.t. the original δ . Note that it is possible that $\widehat{\delta}$ has in fact *fewer* existential quantifiers than δ , for example due to the rewriting of conditions (as happens in the example of Fig. 1 — see the if statement in action check).

Definition 6 (Instrumentation Without Additional Existentials). $\widehat{\delta}$ is a sound instrumentation without additional existentials for δ if there exists an existential naming η such that $\delta_S \rightarrow \eta(\widehat{\delta})[\psi/r, \psi'/r']$ is valid.

4.2 From Instrumentation to Bounded-Horizon

The results described in this section show that if there is an instrumentation without additional existentials, then Bounded-Horizon with a low bound is able to prove the original invariant, without specific knowledge of the instrumentation and without manual assistance from the programmer. This is the case in the example of Fig. 1, which admits an instrumentation that transforms the invariant to a universal invariant (see Example 3) in a form that matches Theorem 3, and indeed the original invariant is provable by Bounded-Horizon of bound 1.

Interestingly, in case Bounded-Horizon with a small bound does not prove inductiveness, as in the example of Fig. 2⁶, the results imply that either the invariant is not inductive or *no instrumentation* that does not add existential quantifiers can be used to show that it is inductive (even with the programmer's manual assistance).

In the remainder of this section we will assume that $\widehat{\delta}$ is a sound instrumentation without additional existentials for δ , and η is the corresponding naming of existentials. Further, we assume that \widehat{I} is an inductive invariant for $\widehat{\delta}$ and denote $I = \widehat{I}[\psi/r]$.

The following theorems state our results for $I \in \forall^* \exists^*$.

Theorem 3. Let $\widehat{I} \in \forall^*$. Assume $\psi \in \exists^*$ and r appears only positively in \widehat{I} , or $\psi \in \forall^*$ and r appears only negatively in \widehat{I} . Then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 1. (Note that $I \in \forall^* \exists^*$.)

Proof Sketch. Let $I = \forall \bar{x}. \alpha(\bar{x})$ where $\alpha \in \exists^*$. Assume for the sake of contradiction that I is not inductive for δ with Bounded-Horizon of bound 1. By the assumptions on ψ and \widehat{I} , this means that there is a structure \mathcal{A} such that

$$\mathcal{A} \models (\bigwedge_{\bar{c}} \alpha(\bar{c})) \wedge \delta_S \wedge (\neg I')_S.$$

From the assumption (Definition 6) and properties of Skolemization, it follows that

$$\mathcal{A} \models (\bigwedge_{\bar{c}} \alpha(\bar{c})) \wedge (\eta(\widehat{\delta}))[\psi/r, \psi'/r'] \wedge ((\neg \widehat{I}')_S)[\psi/r, \psi'/r'].$$

From the assumptions on the way ψ appears in \widehat{I} , when we write $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}$ we have $\alpha = \widehat{\alpha}[\psi/r]$. Thus, from properties of substitution (interpreting r, r' according to ψ, ψ' in \mathcal{A}) it follows that there is a structure $\widehat{\mathcal{A}}$ such that

$$\widehat{\mathcal{A}} \models (\bigwedge_{\bar{c}} \widehat{\alpha}(\bar{c})) \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I}')_S.$$

By reducing $\widehat{\mathcal{A}}$'s domain to the constants we have that $(\forall \bar{x}. \widehat{\alpha}(\bar{x})) \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I}')_S$ is satisfiable. (This is a use of complete instantiation for universal formulas.)

This in turn implies (by properties of Skolemization) that $\widehat{I} \wedge \widehat{\delta} \wedge \neg \widehat{I}'$ is satisfiable, which is a contradiction to the assumption that \widehat{I} is inductive for $\widehat{\delta}$. \square

⁶ Strictly speaking this shows that there is no such instrumentation where the instrumentation relation appears only positively in the invariant, which is the most common case. Examples that require an even larger bound (only sketched for simplicity) do not have any instrumentation without additional existential quantifiers that transforms the invariant to a universal form.

Proof. Let $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}$. In both cases of the claim $\alpha = \widehat{\alpha}[\psi/r] \in \exists^*$, and so all the universal quantifiers in I are those of \widehat{I} . This implies that the satisfiability check of Lemma 4 is simply the Bounded-Horizon satisfiability check with bound 1, and it shows that the result must be unsatisfiable.

More formally, assume for the sake of contradiction that I is not inductive w.r.t. δ with Bounded-Horizon of bound 1. Let $\alpha(\bar{x}) = \exists y_1, \dots, y_m. \theta(\bar{x}, y_1, \dots, y_m)$ where $\theta \in \text{QF}$, and let

$$\alpha_S(\bar{x}) = \theta(\bar{x}, f_1(\bar{x}), \dots, f_m(\bar{x}))$$

be its Skolemization with fresh Skolem function symbols f_1, \dots, f_m (introduced for y_1, \dots, y_m , respectively). Then there is a structure \mathcal{A} satisfying

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\neg I')_S[\bar{t}] \right) \quad (3)$$

Since α_S has no universal quantifiers, the instantiation is just a substitution of the free variables, and \mathcal{A} satisfies

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha_S(\bar{t}) \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\neg I')_S[\bar{t}] \right) \quad (4)$$

By reducing \mathcal{A} to the elements pointed to by BHT_1 terms we have that

$$\mathcal{A}|_{\text{BHT}_1} \models \left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha_S(\bar{t}) \right) \wedge \delta_S \wedge (\neg I')_S \quad (5)$$

We now move from the Skolem functions back to existential quantifiers. By the valuation that to every existentially quantified variable y_i in α assigns the interpretation of $f_i(\bar{t})$ in $\mathcal{A}|_{\text{BHT}_1}$ (recall that $f_i(\bar{t})$ appears in α_S instead of the quantifier $\exists y_i$ in α), we know that

$$\mathcal{A}|_{\text{BHT}_1} \models \left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha(\bar{t}) \right) \wedge \delta_S \wedge (\neg I')_S. \quad (6)$$

Recall that $\text{BHT}_0 = \text{const}[\delta_S \wedge (\neg I')_S]$. Therefore, Equation (26) can be rewritten as

$$\mathcal{A}|_{\text{BHT}_1} \models \left(\bigwedge_{\bar{c} \in C^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S \quad (7)$$

where $C^n = \text{const}[\delta_S \wedge (\neg I')_S]^n$ and is the number of universal quantifiers in I (and \widehat{I}).

By Lemma 4 this is a contradiction to the assumption that \widehat{I} is inductive w.r.t. $\widehat{\delta}$, and the claim follows.

Theorem 4. *Let $\widehat{I} \in \forall^*$. If $\psi \in \text{AF}$ then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 2. (Note that $I \in \forall^* \exists^*$.)*

Proof. As before, Let $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}$. $\psi \in \text{AF}$ implies that $\alpha = \widehat{\alpha}[\psi/r] \in \text{AF}$. Let

$$\alpha(\bar{x}) = (\forall \bar{v}_1 \theta_{1,1}(\bar{x}, \bar{v}_1) \vee \exists \bar{z}_1 \theta_{1,2}(\bar{x}, \bar{z}_1)) \wedge \dots \wedge (\forall \bar{v}_r \theta_{r,1}(\bar{x}, \bar{v}_r) \vee \exists \bar{z}_r \theta_{r,2}(\bar{x}, \bar{z}_r)) \quad (8)$$

where $\theta_{1,1}, \theta_{1,2}, \dots, \theta_{r,1}, \theta_{r,2} \in \text{QF}$.

Assume for the sake of contradiction that I is not inductive w.r.t. δ with Bounded-Horizon of bound 2.

For brevity denote

$$\xi(k) = \left(\bigwedge_{\bar{t} \in \text{BHT}_k} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_k} (\neg I')_S[\bar{t}] \right)$$

and let $\bar{g}_i(\bar{x})$ denote the fresh Skolem function introduced for \bar{z}_i in α_S .

By the assumption that inductiveness is not provable using Bounded-Horizon of bound 2,

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_1} \alpha_S[\bar{t}] \right) \wedge \xi(2) \quad (9)$$

is satisfiable by a structure \mathcal{A} .

In particular

$$\begin{aligned} & \bigwedge \left((\theta_{1,1}(\bar{c}, \bar{d}_1) \vee \theta_{1,2}(\bar{c}, \bar{g}_1(\bar{c}))) \wedge \right. \\ & \left. \dots \wedge (\theta_{r,1}(\bar{c}, \bar{d}_r) \vee \theta_{r,2}(\bar{c}, \bar{g}_r(\bar{c}))) \right) \wedge \xi(1) \end{aligned} \quad (10)$$

is satisfied by \mathcal{A} , where the conjunction is over $\bar{c} \in \text{BHT}_0$ and $\bar{d}_1, \dots, \bar{d}_r \in \text{BHT}_1$ (with the relevant arity). Note that the difference from Equation (29) is that there are fewer conjuncts here, because (i) the full Bounded-Horizon check with bound 2 has conjuncts for each $\bar{c} \in \text{BHT}_1$ and not just BHT_0 , and (ii) we now have $\xi(1)$ instead of $\xi(2)$.

Reduce \mathcal{A} to the elements pointed by BHT_1 terms, let $\mathcal{A}^\downarrow = \mathcal{A}|_{\text{BHT}_1}$.

Now,

$$\begin{aligned} & \bigwedge_{\bar{c} \in \text{BHT}_0} \left((\forall \bar{v}_1 \theta_{1,1}(\bar{c}, \bar{v}_1) \vee \exists \bar{z}_1 \theta_{1,2}(\bar{c}, \bar{z}_1)) \wedge \dots \right. \\ & \left. \wedge (\forall \bar{v}_r \theta_{r,1}(\bar{c}, \bar{v}_r) \vee \exists \bar{z}_r \theta_{r,2}(\bar{c}, \bar{z}_r)) \right) \wedge \xi(1) \end{aligned} \quad (11)$$

is satisfied by \mathcal{A}^\downarrow . This is because:

- The universal quantifiers are semantically equivalent to a conjunction over all BHT_1 elements because the domain was reduced, and
- The existential quantifiers are justified by the following valuation: the valuation assigns every \bar{z}_i the interpretation of $\bar{g}_i(\bar{c})$.

With this valuation the conjunctions of formula 31 are all guaranteed by the conjunctions in formula 30.

Now formula 31 exactly means that

$$\mathcal{A}^\downarrow \models \left(\bigwedge_{\bar{c} \in \text{BHT}_0} \alpha(\bar{c}) \right) \wedge \xi(1). \quad (12)$$

As in the proof of Theorem 3, using Lemma 4, this is a contradiction to the assumption that \widehat{I} is inductive w.r.t. $\widehat{\delta}$, and the claim follows.

The following theorem generalizes the above result to *I-alternation invariants*. A formula is 1-alternation if it can be written as a Boolean combination of $\forall^* \exists^*$ formulas.

Theorem 5. *Let $\widehat{I} \in \text{AF}$. If $\psi \in \text{AF}$ then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 2. (Note that $I \in 1\text{-alternation}$.)*

The full proofs appear in Appendix A. The results of this section also apply when multiple instrumentation relations $\psi_1, \dots, \psi_t \in \text{FOL}(\Sigma)$ are simultaneously substituted instead of the relation symbols r_1, \dots, r_t in $\widehat{\delta}$ and \widehat{I} .

Instrumentations for higher bounds. While instrumentation that does not add existentials is at most as powerful as Bounded-Horizon with a low bound, sound instrumentations that do add existentials to the program (thereby not satisfying Definition 6) can be used to simulate quantifier instantiation of an arbitrary depth. A simple way is to add r as an instrumentation that tracks the existential part of a $\forall^* \exists^*$ invariant. Instantiations are performed by introducing existentially quantified variables to the program and using `assume` statements to make these variables function as witnesses for a tuple of variables that instantiate the universal quantifiers. Doing this recursively generates instantiations of an arbitrary depth. See Appendix B for further details.

5 Partial Models for Understanding Non-Inductiveness

When conducting SMT-based deductive verification (e.g., using Dafny [25]), the user constructs both the formal representation of the system and its invariants. In many cases, the invariant I is initially not inductive w.r.t. the given program, due to a bug in the program or in the invariant. Therefore, deductive verification is typically an iterative process in which the user attempts to prove inductiveness, and, when this fails, adapts the program, the invariant, or both.

In such scenarios, it is extremely desirable to present the user with a *counterexample to induction* in the form of a state that satisfies I but makes a transition to a state that violates it. Such a state can be obtained from a model of the formula $Ind = I \wedge \delta \wedge \neg I'$ which is used to check inductiveness. It explains the error, and guides the user towards fixing the program and/or the invariant [25,12]. However, in many cases where the check involves quantifier alternation, current SMT solvers are unable to produce counterexamples. Instead, SMT solvers usually diverge or report “unknown” [14,32]. In such cases, Bounded-Horizon instantiations can be used to present a concrete logical structure which is comprehensible to the user, and is obtained as a model of the (finite) instantiations of the formula Ind . While this structure is not a true counterexample (as it is only a model of a subset of the instantiations of the formula), it can still guide the user in the right direction towards fixing the program and/or the invariant.

We illustrate this using a simple leader-election protocol in a ring [9], whose model is presented in Fig. 3(a). The protocol assumes that nodes are organized in a directional ring topology with unique IDs, and elects the node with the highest ID as the leader. Each node sends its own ID to its successor, and forwards messages when they contain an ID higher than its own ID. A node that receives its own ID is elected as leader. We wish to prove a termination property which states that once all nodes have sent their ID, and there are no pending messages in the network, then there is an elected leader. To verify this we use a relational model of the protocol similar to [28], and specify the property via the following formula:

$$(\exists n. leader(n)) \vee (\exists n_1, n_2. \neg sent(n_1) \vee pending(n_1, n_2)) \quad (13)$$

A natural attempt of proving this using an inductive invariant is by conjoining Equation (13) (which is not inductive by itself) with the following property (this was the

```

pending := ∅;
... # ring topology

action send_packet(n) {
  assume ring_next(n, m)
  pending := pending ∪ {(n, m)}
  sent := sent ∪ {n}
}

action receive_packet(n, m) {
  assume pending(m, n)
  pending := pending \ {(m, n)}
  if m = n then
    leader := leader ∪ {n}
  else
    if n < m then
      assume ring_next(n, n_0)
      pending := pending ∪ {(m, n_0)}
    else # do not forward
  }
}

```

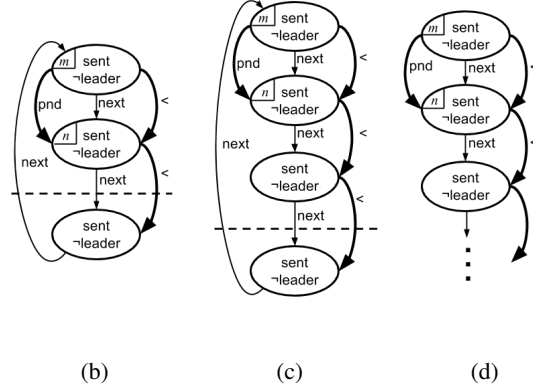


Fig. 3. Leader-election in a ring protocol as an illustration of the use of partial models for incorrect programs and invariants. (a) sketches the protocol (the complete program appears in [1], file `ring_leader_termination.ivy`). (b),(c) show partial models of bound 1 and 2, respectively, and (d) illustrates an infinite structure that explains the root cause of the non-inductiveness.

authors’ actual next step in proving this termination property):

$$\forall n_1. \text{sent}(n_1) \wedge \neg \text{leader}(n_1) \rightarrow ((\exists n_2. \text{pending}(n_1, n_2)) \vee (\exists n_2. n_1 < n_2)) \quad (14)$$

meaning that if a node has sent its own ID but has not (yet) become leader, then there is either a message pending in the network with the node’s ID, or a node with a higher ID.

Alas, the conjunction of Equations (13) and (14) is still not an inductive invariant for the protocol (as we explain below). Since Equation (14) contains $\forall^* \exists^*$ quantification, the associated inductiveness check is outside of the decidable EPR fragment. Indeed, Z3 diverges when it is used to check *Ind*. This is not surprising since the formula has no satisfying finite structures, but has an infinite model (a scenario that is not unusual for $\forall^* \exists^*$ formulas).

On the other hand, applying Bounded-Horizon (with any bound) to *Ind* results in a formula that has finite models. These concrete models are *partial models* of *Ind*. Figs. 3(b) and (c) show partial models (restricted to the pre-states) obtained with bounds of 1 and 2, respectively, on this example.

These models are not true counterexamples to induction: the sub-formula of Equation (14) residing under the universal quantifier does not hold for all the elements of the domain. It does, however, hold for all elements with which the quantifier was instantiated, which are the elements above the dashed line. These elements have all sent their own ID, which was blocked by their successor that has a higher ID, so none of them is the leader. In a finite model, this has to end somewhere, because one of the nodes must have the highest ID. Hence, no finite counter-model exists. However, extrapolating from Fig. 3(b) and (c), we can obtain the infinite model depicted in Fig. 3(d). This model represents an infinite (“open”) ring in which each node has a lower ID than its successor. This model is a true model of the formula *Ind* generated by the invariant in Equations (13) and (14), but the fact that it is infinite prevented Z3 from producing it.

Since we use tools that check general (un)satisfiability, which is not limited to finite structures, the only way to prove that an invariant is inductive is to exclude infinite counterexamples to induction as well. Using Bounded-Horizon instantiations, we are able to obtain meaningful partial models that provide hints to the user about what is missing. In this case, the solution is to add an axiom to the system model which states that there is a node with maximal ID: $\exists n_1. \forall n_2. n_2 \leq n_1$. With this additional assumption, the formula *Ind* is unsatisfiable so the invariant is inductive, and this is proven both by Z3’s instantiation heuristics and by Bounded-Horizon with a bound of 1. This illustrates the usefulness of Bounded-Horizon when the invariant is not inductive.

6 Implementation and Initial Evaluation

We implemented a prototype of Bounded-Horizon of bound 1 on top of Z3 [10] and used it within Ivy [28] and the framework of [20]. We applied the procedure to the incorrect example of Section 5, and successfully verified several correct programs and invariants using bound 1. These examples are (the examples’ code can be found in [1]):

- The client-server example of Fig. 1.
- List reverse [20], where the invariant states that the n edges (“next” pointers) are reversed. The invariant is $\forall^* \exists^*$ due to the encoding of n via n^* as explained in [20].
- Learning switch [5], where the invariant states every routing node has a successor.
- Hole-punching firewall [5], where the invariant states that every allowed external node was contacted by some internal node. We explored two modeling alternatives: using a ghost history relation, or existentially quantifying over time.
- Leader election in a ring [9,28] with the invariant discussed in Section 5.

Table 1. Experimental results.

Program	# \forall	#Func	#Consts	# \forall^\downarrow	B1 Total	B1 Solve	Baseline Z3
Client-server	14	1	15	2	58 ms	3 ms	3 ms
List reverse	47	3	15	4	319 ms	211 ms	50 ms
Learning switch	70	1	7	37	245 ms	65 ms	33 ms
Hole-punching firewall with ghost	15	1	18	3	75 ms	4 ms	4 ms
Hole-punching firewall \exists time	32	2	21	3	102 ms	4 ms	4 ms
Leader-election in a ring (correct)	41	1	21	1	113 ms	36 ms	27 ms
Leader-election in a ring (incorrect)	40	1	20	1	1112 ms	1008 ms	—

B1 Total is the time in milliseconds for the bound 1 implementation. It is compared to **Baseline Z3** which is the solving time in milliseconds of *Ind* as is (with quantifier alternation) by Z3. **B1 Solve** measures the solving time of the formula restricted to bound 1, which demonstrates that most of the overhead occurs when constructing the formula. **# \forall** is the number of universal quantifiers in *Ind*, **#Func** the number of different Skolem function symbols, and **#Consts** the number of constants. **# \forall^\downarrow** is the number of universally quantified variables that were restricted in the bound 1 check. Measurements were performed on a 3.5GHz Intel i5-4690 CPU with 8GB RAM running Linux 3.13 x86_64.

An initial evaluation of the method’s performance appears in Table 1.

Our implementation works by adding “guards” that restrict the range of universal quantifiers to the set of constants where necessary. Technically, recall that we are

considering the satisfiability of $Ind = I_S \wedge \delta_S \wedge (\neg I')$.⁷ Let $\forall x. \theta$ be a subformula of Ind . If θ contains function symbol applications⁸, we transform the subformula to $\forall x. (\bigvee_c x = c) \rightarrow \theta$ where c ranges over $\text{const}[Ind]$. The resulting formula is then dispatched to the solver. This is a simple way to encode the termination criterion of bound 1 while leaving room for the solver to perform the necessary instantiations cleverly. The translation enlarges the formula by $O(\#\text{Consts} \cdot \#\forall)$ although the number of bounded instantiations grows exponentially with $\#\forall$. The exponential explosion is due to combinations of constants in the instantiation, a problem we defer to the solver.

Z3 terminates on the class of formulas because during the Model-Based Quantifier Instantiation process every instantiation of a universally quantified formula has the same truth value in the model as an instantiation using one of the existing ground terms (constants and then BHT_1 terms). Z3’s instantiation engine will produce instantiations using existing terms rather than create superfluous new terms [7].

The results are encouraging because they suggest that the termination strategy of Bounded-Horizon, at least for bound 1, can be combined with existing instantiation techniques to assure termination with only a slight performance penalty. Most encouraging is the satisfiable example of Section 5. On this instance, Z3 was able to return “sat” within seconds, although to do so, in theory, the solver must exhaust the entire set of bounded instantiations. This suggests that the Bounded-Horizon termination criterion might indeed be useful for “sat” instances on which the solver may diverge.

A different approach to the implementation is to integrate the termination criterion of the bound with the solver’s heuristics more closely (see [6]).

7 Related Work

Quantifier instantiation. The importance of formulas with quantifier-alternations for program verification has led to many developments in the SMT and theorem-proving communities that aim to allow automated reasoning with quantifier-alternations. The Simplify system [11] promoted the practical usage of quantifier triggers, which let the user affect the quantifier instantiation in a clever way. Similar methods are integrated into modern SMT solvers such as Z3 [10]. Recently, a method for annotating the source code with triggers has been developed for Dafny [26]. The notion of instantiation depth is related to the notions of matching-depth [11] and instantiation-level [13] which are used for prioritization within the trigger-based instantiation procedure.

In addition to user-provided triggers, many automated heuristics for quantifier instantiation have been developed, such as Model-Based Quantifier Instantiation [14]. Even when quantifier instantiation is refutation-complete, it is still important and challenging to handle the SAT cases, which are especially important for program verification. Accordingly, many works (e.g., [32]) consider the problem of model finding.

Local Theory Extensions and Psi-Local Theories [35,18,6] identify settings in which limited quantifier instantiations are complete. They show that completeness is achieved

⁷ Skolemization is performed via Z3, taking advantage of heuristics that reduce the number of different Skolem functions.

⁸ This in fact implements the approximation as of Equation (2). The exact bound 1 per Equation (1) can be implemented by a more careful consideration of which universally quantified variables should be restricted, but this was not necessary for our examples.

exactly when every partial model can be extended to a (total) model. In such settings Bounded-Horizon instantiations are complete for invariant checking. However, Bounded-Horizon can also be useful when completeness cannot be guaranteed.

Classes of SMT formulas that are decidable by complete instantiations have been studied by [14]. In the uninterpreted fragment, a refined version of Herbrand’s Theorem generates a finite set of instantiations when the dependencies are stratified. Bounded-Horizon is a way to bound unstratified dependencies.

Natural Proofs. Natural proofs [29] provide a sound and incomplete proof technique for deductive verification. The key idea is to instantiate recursive definitions over the terms appearing in the program. Bounded-Horizon is motivated by a similar intuition, but focuses on instantiating quantifiers in a way that is appropriate for the EPR setting.

Decidable logics. Different decidable logics can be used to check inductive invariants. For example, Monadic second-order logic [16] obtains decidability by limiting the underlying domain to consist of trees only, and in particular does not allow arbitrary relations, which are useful to describe properties of programs. There are also many decidable fragments of first-order logic [8]. Our work aims to transcend the class of invariants checkable by a reduction to the decidable logic EPR. We note that the example of Section 5 does not fall under the Loosely-Guarded Fragment of first-order logic [17] due to a use of a transitivity axiom, and does not enjoy the finite-model property.

Abstractions for verification of infinite-state systems. Our work is closely related to abstractions of infinite state systems. These abstractions aim at automatically inferring inductive invariants in a sound way. We are interested in checking if a given invariant is inductive either for automatic and semi-automatic verification.

The View-Abstraction approach [3,4,2] defines a useful abstraction for the verification of parameterized systems. This abstraction is closely related to universally quantified invariants. An extension of this approach [4] adds contexts to the abstraction, which are used to capture $\forall^*\exists^*$ invariants in a restricted setting where nodes have finite-state and are only related by specific topologies. Our work is in line with the need to use $\forall^*\exists^*$ invariants for verification, but applies in a more general setting (with unrestricted high-arity relations) at the cost of losing completeness of invariant checking.

Our work is related to the TVLA system [27,34] which allows the programmers to define instrumentation relations. TVLA also employs finite differencing to infer sound update code for updating instrumentation relations [31], but generates non-EPR formulas and does not guarantee completeness. The focus operation in TVLA implements materialization which resembles quantifier-instantiation. TVLA shows that very few built-in instrumentation relations can be used to verify many different programs.

Instrumentation and update formulas. The idea of using instrumentation relations and generating update formulas is not limited to TVLA and was also used for more predictable SMT verification [23,24].

8 Conclusion

We have provided an initial study of the power of bounded instantiations for tackling quantifier alternation. This paper shows that quantifier instantiation with small bounds can simulate instrumentation. This is a step in order to eliminate the need for instrumenting the program, which can be error-prone. The other direction, i.e. simulating

quantifier instantiation with instrumentation, is also possible but is less appealing from a practical point of view, and is presented in Appendix B.

We are encouraged by our initial experience that shows that various protocols can be proven with small instantiation bounds, and that partial models are useful for understanding the failures of the solver to prove inductiveness. Some of these failures correspond to non-inductive claims, especially those due to infinite counterexamples. In the future we hope to leverage this in effective deductive verification tools, and explore meaningful ways to display infinite counterexamples to the user.

Acknowledgments. We would like to thank Nikolaj Bjørner, Shachar Itzhaky, and Bryan Parno for helpful discussions, and Gilit Zohar-Oren for help and feedback. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) /ERC grant agreement no [321174]. This research was partially supported by BSF grant no. 2012259, and by Len Blavatnik and the Blavatnik Family foundation.

References

1. Full code materials. http://www.cs.tau.ac.il/research/yotam.feldman/papers/tacas17/examples_code.zip.
2. P. Abdulla, F. Haziza, and L. Holík. Parameterized verification through view abstraction. *International Journal on Software Tools for Technology Transfer*, pages 1–22, 2015.
3. P. A. Abdulla, F. Haziza, and L. Holík. All for the price of few. In *Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings*, pages 476–495, 2013.
4. P. A. Abdulla, F. Haziza, and L. Holík. Block me if you can! In *International Static Analysis Symposium*, pages 1–17. Springer, 2014.
5. T. Ball, N. Bjørner, A. Gember, S. Itzhaky, A. Karbyshev, M. Sagiv, M. Schapira, and A. Valadarsky. Vericon: towards verifying controller programs in software-defined networks. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*, page 31, 2014.
6. K. Bansal, A. Reynolds, T. King, C. W. Barrett, and T. Wies. Deciding local theory extensions via e-matching. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II*, pages 87–105, 2015.
7. N. Bjørner. personal communication.
8. E. Börger, E. Grädel, and Y. Gurevich. *The classical decision problem*. Springer Science & Business Media, 2001.
9. E. Chang and R. Roberts. An improved algorithm for decentralized extrema-finding in circular configurations of processes. *Communications of the ACM*, 22(5):281–283, 1979.
10. L. De Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.
11. D. Detlefs, G. Nelson, and J. B. Saxe. Simplify: a theorem prover for program checking. *J. ACM*, 52(3):365–473, 2005.
12. C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany, June 17-19, 2002*, pages 234–245, 2002.
13. Y. Ge, C. W. Barrett, and C. Tinelli. Solving quantified verification conditions using satisfiability modulo theories. *Ann. Math. Artif. Intell.*, 55(1-2):101–122, 2009.
14. Y. Ge and L. D. Moura. Complete instantiation for quantified formulas in satisfiability modulo theories. In *International Conference on Computer Aided Verification*, pages 306–320. Springer, 2009.
15. C. Hawblitzel, J. Howell, M. Kapritsos, J. R. Lorch, B. Parno, M. L. Roberts, S. T. V. Setty, and B. Zill. Ironfleet: proving practical distributed systems correct. In *Proceedings of the 25th Symposium on Operating Systems Principles, SOSP*, pages 1–17, 2015.
16. J. G. Henriksen, J. L. Jensen, M. E. Jørgensen, N. Klarlund, R. Paige, T. Rauhe, and A. Sandholm. Mona: Monadic second-order logic in practice. In *Tools and Algorithms for Construction and Analysis of Systems, First International Workshop, TACAS '95, Aarhus, Denmark, May 19-20, 1995, Proceedings*, pages 89–110, 1995.
17. I. Hodkinson. Loosely guarded fragment of first-order logic has the finite model property. *Studia Logica*, 70(2):205–240, 2002.
18. C. Ihlemann, S. Jacobs, and V. Sofronie-Stokkermans. On local reasoning in verification. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, pages 265–281, 2008.

19. N. Immerman, A. M. Rabinovich, T. W. Reps, S. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *CSL*, 2004.
20. S. Itzhaky, A. Banerjee, N. Immerman, A. Nanevski, and M. Sagiv. Effectively-propositional reasoning about reachability in linked data structures. In *CAV*, volume 8044 of *LNCS*, pages 756–772, 2013.
21. S. Itzhaky, N. Bjørner, T. W. Reps, M. Sagiv, and A. V. Thakur. Property-directed shape analysis. In *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, pages 35–51, 2014.
22. A. Karbyshev, N. Bjørner, S. Itzhaky, N. Rinetzky, and S. Shoham. Property-directed inference of universal invariants or proving their absence. In *CAV*, 2015.
23. S. K. Lahiri and S. Qadeer. Verifying properties of well-founded linked lists. In *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, Charleston, South Carolina, USA, January 11-13, 2006*, pages 115–126, 2006.
24. S. K. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using SMT solvers. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 171–182, 2008.
25. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In *Logic for Programming, Artificial Intelligence, and Reasoning*, pages 348–370. Springer, 2010.
26. K. R. M. Leino and C. Pit-Claudel. Trigger selection strategies to stabilize program verifiers. In *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part I*, pages 361–381, 2016.
27. T. Lev-Ami and S. Sagiv. TVLA: A system for implementing static analyses. In *Static Analysis, 7th International Symposium, SAS 2000, Santa Barbara, CA, USA, June 29 - July 1, 2000, Proceedings*, pages 280–301, 2000.
28. O. Padon, K. L. McMillan, A. Panda, M. Sagiv, and S. Shoham. Ivy: safety verification by interactive generalization. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, pages 614–630, 2016.
29. X. Qiu, P. Garg, A. Stefanescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 231–242, 2013.
30. F. P. Ramsey. On a problem of formal logic. *Proceedings of the London Mathematical Society*, s2-30(1):264–286, 1930.
31. T. W. Reps, M. Sagiv, and A. Loginov. Finite differencing of logical formulas for static analysis. *ACM Trans. Program. Lang. Syst.*, 32(6), 2010.
32. A. Reynolds, C. Tinelli, A. Goel, and S. Krstic. Finite model finding in SMT. In *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, pages 640–655, 2013.
33. A. Reynolds, C. Tinelli, A. Goel, S. Krstic, M. Deters, and C. Barrett. Quantifier instantiation techniques for finite model finding in SMT. In *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, pages 377–391, 2013.
34. S. Sagiv, T. W. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *ACM Trans. Program. Lang. Syst.*, 24(3):217–298, 2002.
35. V. Sofronie-Stokkermans. Hierarchic reasoning in local theory extensions. In *Automated Deduction - CADE-20, 20th International Conference on Automated Deduction, Tallinn, Estonia, July 22-27, 2005, Proceedings*, pages 219–234, 2005.

A Bounded-Horizon vs. Instrumentation

In the following we include the full proofs of the lemmas formalized in Section 4. Throughout this section, we fix a relational vocabulary Σ and an EPR transition relation $\delta \in \exists^*\forall^*(\Sigma, \Sigma')$. We use $\psi \in \text{FOL}(\Sigma)$ to denote an instrumentation-defining formula with free variables x_1, \dots, x_a . Further, r denotes a fresh relation symbol of arity a , and $\widehat{\Sigma} = \Sigma \cup \{r\}$ is the instrumented vocabulary used to define an instrumented transition relation $\widehat{\delta} \in \exists^*\forall^*(\widehat{\Sigma}, \widehat{\Sigma}')$. Finally, \widehat{I} is an inductive invariant for $\widehat{\delta}$ (over $\widehat{\Sigma}$), and $I = \widehat{I}[\psi/r]$. We will later further restrict the quantifier structure of I and \widehat{I} in different ways.

A.1 Power for $\forall^*\exists^*$ Invariants

We now establish that low bounds are sufficient for the Bounded-Horizon check when $\widehat{I} \in \forall^*(\widehat{\Sigma})$ and $I \in \forall^*\exists^*$. To do so, we first prove the following lemma.

Lemma 4. *Let $\widehat{\delta}$ be a sound instrumentation of δ without new existentials and with naming η . Write $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}(\widehat{\Sigma})$ and let $\alpha = \widehat{\alpha}[\psi/r]$. Then,*

$$\left(\bigwedge_{\bar{c} \in C^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S \quad (15)$$

is unsatisfiable, where $C = \text{const}[\delta_S \wedge (\neg I')_S]$ and n is the number of universal quantifiers in \widehat{I} .

Proof. Assume not, i.e., there exists a structure \mathcal{A}_0 such that,

$$\mathcal{A}_0 \models \left(\bigwedge_{\bar{c} \in C^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S. \quad (16)$$

We will show that \widehat{I} is not inductive for $\widehat{\delta}$. Let $\widehat{C} = \text{const}[\eta(\widehat{\delta}) \wedge (\neg \widehat{I})_S]$. Then,

$$\mathcal{A}_1 \models \left(\bigwedge_{\bar{c} \in (C \cup \widehat{C})^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S \quad (17)$$

where \mathcal{A}_1 is the same as \mathcal{A}_0 but also interprets any constant in $\widehat{C} \setminus C$ as some arbitrary constant in C . Thus $\alpha(\bar{c})$ holds for the new constants as well.

Removing some conjuncts from Equation (17), we get,

$$\mathcal{A}_1 \models \left(\bigwedge_{\bar{c} \in \widehat{C}^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S. \quad (18)$$

By assumption (Definition 6), it follows that,

$$\mathcal{A}_1 \models \left(\bigwedge_{\bar{c} \in \widehat{C}^n} \alpha(\bar{c}) \right) \wedge \eta(\widehat{\delta})[\psi/r, \psi'/r'] \wedge (\neg I')_S. \quad (19)$$

Recall that $I' = \widehat{I}'[\psi'/r']$. Since $\mathcal{A}_1 \models (\neg \widehat{I}'[\psi'/r'])_S$, it follows that $\mathcal{A}_1 \models (\neg \widehat{I})_S[\psi'/r']$. In the latter formula, some existentially quantified variables from ψ or $\neg\psi$ may remain, whereas in the former formula they were replaced by Skolem constants. Thus this is just a corollary of the fact that $\gamma_S \rightarrow \gamma$ is valid for any γ .

Thus we have shown,

$$\mathcal{A}_1 \models \left(\left(\bigwedge_{\bar{c} \in \widehat{C}^n} \widehat{\alpha}(\bar{c}) \right) \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I})_S \right) [\psi/r, \psi'/r']. \quad (20)$$

Now, consider the structure $\widehat{\mathcal{A}}$ that expands \mathcal{A}_1 by interpreting r and r' the way that \mathcal{A}_1 interprets ψ and ψ' , respectively. Then,

$$\widehat{\mathcal{A}} \models \left(\bigwedge_{\bar{c} \in \widehat{C}^n} \widehat{\alpha}(\bar{c}) \right) \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I}')_S. \quad (21)$$

Since the formula in Equation (21) is universal, it is also satisfied by $\widehat{\mathcal{A}}|_{\widehat{C}}$, the substructure of $\widehat{\mathcal{A}}$ with universe $\widehat{C}^{\widehat{\mathcal{A}}}$, i.e., $\widehat{\mathcal{A}}$'s interpretation of the constant symbols \widehat{C} (recall that $\widehat{C} = \text{const}[\eta(\widehat{\delta}) \wedge (\neg \widehat{I}')_S]$). Thus,

$$\widehat{\mathcal{A}}|_{\widehat{C}} \models (\forall \bar{x}. \widehat{\alpha}(\bar{x})) \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I}')_S. \quad (22)$$

Finally, since $\gamma_S \rightarrow \gamma$ is valid and so is $\eta(\widehat{\delta}) \rightarrow \widehat{\delta}$ (for the same reasons), we know,

$$\widehat{\mathcal{A}}|_{\widehat{C}} \models \widehat{I} \wedge \widehat{\delta} \wedge \neg \widehat{I}'. \quad (23)$$

But this contradicts the fact that \widehat{I} is inductive for $\widehat{\delta}$. \square

The following results are corollaries of Lemma 4.

Theorem (3). *If $\psi \in \exists^*$ and r appears only positively in \widehat{I} , or $\psi \in \forall^*$ and r appears only negatively in \widehat{I} , then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 1. (Note that $I \in \forall^* \exists^*$.)*

Proof. Let $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}$. In both cases of the claim $\alpha = \widehat{\alpha}[\psi/r] \in \exists^*$, and so all the universal quantifiers in I are those of \widehat{I} . This implies that the satisfiability check of Lemma 4 is simply the Bounded-Horizon satisfiability check with bound 1, and it shows that the result must be unsatisfiable.

More formally, assume for the sake of contradiction that I is not inductive w.r.t. δ with Bounded-Horizon of bound 1. Let $\alpha(\bar{x}) = \exists y_1, \dots, y_m. \theta(\bar{x}, y_1, \dots, y_m)$ where $\theta \in \text{QF}$, and let

$$\alpha_S(\bar{x}) = \theta(\bar{x}, f_1(\bar{x}), \dots, f_m(\bar{x}))$$

be its Skolemization with fresh Skolem function symbols f_1, \dots, f_m (introduced for y_1, \dots, y_m , respectively). Then there is a structure \mathcal{A} satisfying

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\neg I')_S[\bar{t}] \right) \quad (24)$$

Since α_S has no universal quantifiers, the instantiation is just a substitution of the free variables, and \mathcal{A} satisfies

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha_S(\bar{t}) \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\neg I')_S[\bar{t}] \right) \quad (25)$$

Reduce \mathcal{A} to the elements pointed to by BHT_1 terms, and remove the interpretation of the function symbols (since the interpretation now is of a partial function only). Consider the valuation that to every existentially quantified variable y_i in $\alpha(\bar{t})$ assigns the interpretation of $f_i(\bar{t})$ in \mathcal{A} (recall that f_i appears in α_S instead of the quantifier $\exists y_i$ in α). This is well defined in $\mathcal{A}|_{\text{BHT}_1}$ because $\bar{t} \in \text{BHT}_0$ and so $f_i(\bar{t}) \in \text{BHT}_1$. By this valuation we know that

$$\mathcal{A}|_{\text{BHT}_1} \models \left(\bigwedge_{\bar{t} \in \text{BHT}_0} \alpha(\bar{t}) \right) \wedge \delta_S \wedge (\neg I')_S. \quad (26)$$

Note that the universal quantifiers are semantically equivalent to a conjunction over all BHT_1 elements because the domain was reduced, and thus it is correct to replace the conjunction over BHT_1 terms in Equation (25) by universal quantifiers.

Recall that $\text{BHT}_0 = \text{const}[\delta_S \wedge (\neg I')_S]$. Therefore, Equation (26) can be rewritten as

$$\mathcal{A}_{|\text{BHT}_1} \models \left(\bigwedge_{\bar{c} \in C^n} \alpha(\bar{c}) \right) \wedge \delta_S \wedge (\neg I')_S \quad (27)$$

where $C^n = \text{const}[\delta_S \wedge (\neg I')_S]^n$ and is the number of universal quantifiers in I (and \widehat{I}).

By Lemma 4 this is a contradiction to the assumption that \widehat{I} is inductive w.r.t. $\widehat{\delta}$, and the claim follows. \square

Theorem (4). *If $\psi \in \text{AF}$ then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 2. (Note that $I \in \forall^* \exists^*$.)*

Proof. As before, Let $\widehat{I} = \forall \bar{x}. \widehat{\alpha}(\bar{x})$ where $\widehat{\alpha} \in \text{QF}$. $\psi \in \text{AF}$ implies that $\alpha = \widehat{\alpha}[\psi/r] \in \text{AF}$. Let

$$\begin{aligned} \alpha(\bar{x}) = & (\forall \bar{v}_1 \theta_{1,1}(\bar{x}, \bar{v}_1) \vee \exists \bar{z}_1 \theta_{1,2}(\bar{x}, \bar{z}_1)) \wedge \dots \\ & \wedge (\forall \bar{v}_r \theta_{r,1}(\bar{x}, \bar{v}_r) \vee \exists \bar{z}_r \theta_{r,2}(\bar{x}, \bar{z}_r)) \end{aligned} \quad (28)$$

where $\theta_{1,1}, \theta_{1,2}, \dots, \theta_{r,1}, \theta_{r,2} \in \text{QF}$.

Assume for the sake of contradiction that I is not inductive w.r.t. δ with Bounded-Horizon of bound 2.

For brevity denote

$$\xi(k) = \left(\bigwedge_{\bar{t} \in \text{BHT}_k} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_k} (\neg I')_S[\bar{t}] \right)$$

and let $\bar{g}_i(\bar{x})$ denote the fresh Skolem function introduced for \bar{z}_i in α_S .

By the assumption that inductiveness is not provable using Bounded-Horizon of bound 2,

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_1} \alpha_S[\bar{t}] \right) \wedge \xi(2) \quad (29)$$

is satisfiable by a structure \mathcal{A} .

In particular

$$\begin{aligned} \bigwedge \left((\theta_{1,1}(\bar{c}, \bar{d}_1) \vee \theta_{1,2}(\bar{c}, \bar{g}_1(\bar{c}))) \wedge \right. \\ \left. \dots \wedge (\theta_{r,1}(\bar{c}, \bar{d}_r) \vee \theta_{r,2}(\bar{c}, \bar{g}_r(\bar{c}))) \right) \wedge \xi(1) \end{aligned} \quad (30)$$

is satisfied by \mathcal{A} , where the conjunction is over $\bar{c} \in \text{BHT}_0$ and $\bar{d}_1, \dots, \bar{d}_r \in \text{BHT}_1$ (with the relevant arity). Note that the difference from Equation (29) is that there are fewer conjuncts here, because (i) the full Bounded-Horizon check with bound 2 has conjuncts for each $\bar{c} \in \text{BHT}_1$ and not just BHT_0 , and (ii) we now have $\xi(1)$ instead of $\xi(2)$.

Reduce \mathcal{A} to the elements pointed by BHT_1 terms, let $\mathcal{A}^\downarrow = \mathcal{A}_{|\text{BHT}_1}$. Remove the interpretations of the function symbols (which after the reduction are only partial functions, but will not be used in the rest of the proof).

Now,

$$\bigwedge_{\bar{c} \in \text{BHT}_0} \left((\forall \bar{v}_1 \theta_{1,1}(\bar{c}, \bar{v}_1) \vee \exists \bar{z}_1 \theta_{1,2}(\bar{c}, \bar{z}_1)) \wedge \dots \right. \\ \left. \wedge (\forall \bar{v}_r \theta_{r,1}(\bar{c}, \bar{v}_r) \vee \exists \bar{z}_r \theta_{r,2}(\bar{c}, \bar{z}_r)) \right) \wedge \xi(1) \quad (31)$$

is satisfied by \mathcal{A}^\downarrow . This is because:

- The universal quantifiers are semantically equivalent to a conjunction over all BHT_1 elements because the domain was reduced, and
- The existential quantifiers are justified by the following valuation: the valuation assigns every \bar{z}_i the interpretation of $\bar{g}_i(\bar{c})$.

With this valuation the conjunctions of formula 31 are all guaranteed by the conjunctions in formula 30.

Now formula 31 exactly means that

$$\mathcal{A}^\downarrow \models \left(\bigwedge_{\bar{c} \in \text{BHT}_0} \alpha(\bar{c}) \right) \wedge \xi(1). \quad (32)$$

As in the proof of Theorem 3, using Lemma 4, this is a contradiction to the assumption that \hat{I} is inductive w.r.t. $\hat{\delta}$, and the claim follows. \square

A.2 Generalization to 1-Alternation Invariants

We now generalize the results of Appendix A.1 to *1-alternation invariants*. A formula is 1-alternation if it can be written as a Boolean combination of $\forall^* \exists^*$ formulas. In the sequel, $\hat{I} \in \text{AF}(\hat{\Sigma})$ and $I = \hat{I}[\psi/r] \in 1\text{-alternation}(\Sigma)$.

Lemma 5. *Let $\psi \in \text{FOL}(\Sigma)$. Let $\hat{I} \in \text{AF}(\hat{\Sigma})$ be an inductive invariant for $\hat{\delta} \in \exists^* \forall^*(\hat{\Sigma})$. Write $\hat{I}_S = \forall \bar{x}. \hat{\alpha}_1(\bar{x})$ and $(\neg \hat{I})_S = \forall \bar{x}. \hat{\alpha}_2(\bar{x})$, where $\hat{\alpha}_1, \hat{\alpha}_2 \in \text{QF}(\hat{\Sigma})$. Let $\alpha_1 = \hat{\alpha}_1[\psi/r]$ and $\alpha_2 = \hat{\alpha}_2[\psi/r]$. Let $\delta \in \exists^* \forall^*(\Sigma)$ and let $\hat{\delta} \in \exists^* \forall^*(\hat{\Sigma}, \hat{\Sigma}')$ be a sound instrumentation of δ without new existentials and with naming η . Then,*

$$\left(\bigwedge_{\bar{c}_1 \in C^n} \alpha_1(\bar{c}_1) \right) \wedge \delta_S \wedge \left(\bigwedge_{\bar{c}_2 \in C^m} \alpha_2(\bar{c}_2) \right) \quad (33)$$

is unsatisfiable, where $C = \text{const}[\hat{I}_S \wedge \delta_S \wedge (\neg \hat{I})_S]$, n is the number of universal quantifiers in \hat{I}_S and m is the number of universal quantifiers in $(\neg \hat{I})_S$.

Proof. The proof is similar to that of Lemma 4, with the transformations applied to the conjunction are performed both on the invariant in the pre-state and in the post-state.

Assume not, i.e., there exists a structure \mathcal{A}_0 such that,

$$\mathcal{A}_0 \models \left(\bigwedge_{\bar{c}_1 \in C^n} \alpha_1(\bar{c}_1) \right) \wedge \delta_S \wedge \left(\bigwedge_{\bar{c}_2 \in C^m} \alpha_2(\bar{c}_2) \right). \quad (34)$$

We will show that \hat{I} is not inductive for $\hat{\delta}$. Let $\hat{C} = \text{const}[\hat{I}_S \wedge \eta(\hat{\delta}) \wedge (\neg \hat{I})_S]$. Then,

$$\mathcal{A}_1 \models \left(\bigwedge_{\bar{c}_1 \in \hat{C}^n} \alpha_1(\bar{c}_1) \right) \wedge \delta_S \wedge \left(\bigwedge_{\bar{c}_2 \in \hat{C}^m} \alpha_2(\bar{c}_2) \right). \quad (35)$$

where \mathcal{A}_1 is the same as \mathcal{A}_0 but also interprets any constant in $\widehat{C} \setminus C$ as some arbitrary constant in C .

By the assumption (Definition 6), it follows that,

$$\mathcal{A}_1 \models \left(\bigwedge_{\bar{c}_1 \in \widehat{C}^n} \alpha_1(\bar{c}_1) \right) \wedge \eta(\widehat{\delta})[\psi/r, \psi'/r'] \wedge \left(\bigwedge_{\bar{c}_2 \in \widehat{C}^m} \alpha_2(\bar{c}_2) \right). \quad (36)$$

Thus we have shown,

$$\mathcal{A}_1 \models \left(\left(\bigwedge_{\bar{c}_1 \in \widehat{C}^n} \widehat{\alpha}_1(\bar{c}_1) \right) \wedge \eta(\widehat{\delta}) \wedge \left(\bigwedge_{\bar{c}_2 \in \widehat{C}^m} \widehat{\alpha}_2(\bar{c}_2) \right) \right) [\psi/r, \psi'/r']. \quad (37)$$

Now, consider the structure $\widehat{\mathcal{A}}$ that expands \mathcal{A}_1 by interpreting r and r' the way that \mathcal{A}_1 interprets ψ and ψ' , respectively. Then,

$$\widehat{\mathcal{A}} \models \left(\bigwedge_{\bar{c}_1 \in \widehat{C}^n} \widehat{\alpha}_1(\bar{c}_1) \right) \wedge \eta(\widehat{\delta}) \wedge \left(\bigwedge_{\bar{c}_2 \in \widehat{C}^m} \widehat{\alpha}_2(\bar{c}_2) \right) \quad (38)$$

Since the formula in Equation (38) is universal, it is also satisfied by $\widehat{\mathcal{A}}_{|\widehat{C}}$, the sub-structure of $\widehat{\mathcal{A}}$ with universe $\widehat{C}^{\widehat{\mathcal{A}}}$, i.e., $\widehat{\mathcal{A}}$'s interpretation of the constant symbols \widehat{C} . Thus,

$$\widehat{\mathcal{A}}_{|\widehat{C}} \models (\forall \bar{x}. \widehat{\alpha}_1(\bar{x})) \wedge \eta(\widehat{\delta}) \wedge (\forall \bar{x}. \widehat{\alpha}_2(\bar{x})) \quad (39)$$

Recall that \widehat{C} was defined as $\widehat{C} = \text{const}[\widehat{I} \wedge \eta(\widehat{\delta}) \wedge (\neg \widehat{I})_S]$.

Finally, since $\gamma_S \rightarrow \gamma$ is valid and for the same reasons $\eta(\widehat{\delta}) \rightarrow \widehat{\delta}$ is valid, we know,

$$\widehat{\mathcal{A}}_{|\widehat{C}} \models \widehat{I} \wedge \widehat{\delta} \wedge \neg \widehat{I}'. \quad (40)$$

But this contradicts the fact that \widehat{I} is inductive for $\widehat{\delta}$. \square

The following result is a corollary of Lemma 5.

Theorem (5). *If $\psi \in \text{AF}$ then $I = \widehat{I}[\psi/r]$ is inductive for δ with Bounded-Horizon of bound 2.*

Proof. $\psi \in \text{AF}$ implies that $\alpha_1(\bar{x}) = \widehat{\alpha}_1[\psi/r] \in \text{AF}$ and $\alpha_2(\bar{x}) = \widehat{\alpha}_2[\psi/r] \in \text{AF}$ (recall that $\widehat{\alpha}_1, \widehat{\alpha}_2 \in \text{QF}$).

By the assumption that inductiveness is not provable using Bounded-Horizon of bound 2,

$$\left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\alpha_1)_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_2} \delta_S[\bar{t}] \right) \wedge \left(\bigwedge_{\bar{t} \in \text{BHT}_1} (\alpha_2)_S[\bar{t}] \right) \quad (41)$$

is satisfiable, where $(\alpha_1)_S, (\alpha_2)_S$ introduce Skolem functions. Let \mathcal{A} be such a satisfying structure, and $\mathcal{A}^\downarrow = \mathcal{A}_{|\text{BHT}_1}$.

Because $\alpha_1 \in \text{AF}$, in the same way as in the proof of Theorem 4,

$$\mathcal{A}^\downarrow \models \bigwedge_{\bar{c} \in \text{BHT}_0} \alpha_1(\bar{c}) \quad (42)$$

and in the same way, since $\alpha_2 \in \text{AF}$ as well, the same structure has

$$\mathcal{A}^\downarrow \models \bigwedge_{\bar{c} \in \text{BHT}_0} \alpha_2(\bar{c}). \quad (43)$$

Note that from Equation (41), $\mathcal{A}^\downarrow \models \widehat{\delta}$. Overall we have

$$\mathcal{A}^\downarrow \models \left(\bigwedge_{\bar{c} \in \text{BHT}_0} \alpha_1(\bar{c}) \right) \wedge \delta_S \wedge \left(\bigwedge_{\bar{c} \in \text{BHT}_0} \alpha_2(\bar{c}) \right) \quad (44)$$

and by Lemma 5 this is a contradiction to the assumption that \widehat{I} is inductive w.r.t. $\widehat{\delta}$. \square

B Instrumentation for High Depth Instantiations

In this section we discuss the connection between quantifier instantiation and program instrumentation in the converse direction, i.e. simulating quantifier instantiation with instrumentation. We show how to systematically construct instrumented programs in a way that corresponds to quantifier instantiation for $\forall^* \exists^*$ -invariants: performing the required instantiations within the program allows to make the verification conditions decidable to check.

As before, the instrumentation process begins with a program that does not have a universal inductive invariant. We would like to transform the program in a sound way to a program that has a universal inductive invariant \widehat{I} . As before, we start by identifying some existential formula $\psi(\bar{x}) \in \exists^*(\Sigma)$ that expresses needed information, and encode it using an instrumentation relation r , with the meaning that “ $r(\bar{x}) \equiv \psi(\bar{x})$ ”. As before, we present the procedure for a single formula and instrumentation relation, but in practice it can be used with multiple formulas and instrumentation relations (see ??).

Adding the instrumentation relation r lets $\widehat{I} \in \forall^*(\widehat{\Sigma})$ express existential information by referring to r . The modifications to the program must encode enough of “ $r(\bar{x}) \equiv \psi(\bar{x})$ ” to make \widehat{I} inductive. As opposed to the instrumentation described in Section 4.1, this will be done by instantiating the correspondence between $\psi(\bar{t})$ and $r(\bar{t})$ for specific variables in the program using `assume` statements. This replaces the rewriting of program conditions that use ψ to use r . This approach is also extended to obtain deep instantiations which simulate Bounded-Horizon with arbitrary bound.

B.1 Local Instantiations

Rather than rewriting $\psi(\bar{t})$ to $r(\bar{t})$ in the program, we would like to enforce r to be interpreted according to ψ in the pre-state, i.e., to enforce $\forall \bar{x}. \psi(\bar{x}) \leftrightarrow r(\bar{x})$. The direction $\forall \bar{x}. \psi(\bar{x}) \rightarrow r(\bar{x})$ is an EPR formula (since $\psi \in \exists^*$), and thus we can simply conjoin it to the verification condition while keeping decidability.

The converse implication, $\forall \bar{x}. r(\bar{x}) \rightarrow \psi(\bar{x})$, is a $\forall^* \exists^*$ formula, and adding it to the verification condition will lead to undecidability. Note that Bounded-Horizon with bound 1 is analogous to enforcing $r(\bar{t}) \rightarrow \psi(\bar{t})$ for every t that is a program variable. Inspired by this, we define the following instrumentation that lets the user locally enforce the definition of r for program variables.

Definition 7 (Local Instantiation). *Let $\psi(\bar{x}) = \exists \bar{y}. \varphi(\bar{x}, \bar{y})$. To generate an instantiation of $\forall \bar{x}. r(\bar{x}) \rightarrow \exists \bar{y}. \varphi(\bar{x}, \bar{y})$ on some tuple of program variables \bar{t} , we instrument*

the program by adding new program variables \bar{c} , and inserting the following code:

$$\bar{c} := *; \text{assumer } r(\bar{t}) \rightarrow \varphi(\bar{t}, \bar{c})$$

This code uses the havoc statement, which sets the value of \bar{c} to arbitrary values, followed by an assume statement that restricts the execution such that if $r(\bar{t})$ holds, then $\varphi(\bar{t}, \bar{c})$ holds. Thus, this code realizes the restriction that $r(\bar{t}) \rightarrow \exists \bar{y}. \varphi(\bar{t}, \bar{y})$, and also assigns to the new program variables \bar{c} , the witnesses for the existential quantifiers. We call this addition to the program a local instantiation, as it imposes the connection between r and ψ locally for some program variables \bar{t} .

Lemma 6 (Soundness of Local Instantiations). *If $\widehat{\delta}$ is obtained from δ by a local instantiation then $\widehat{\delta}$ is a sound instrumentation by Definition 4.*

Proof. The code added by a local instantiation for \bar{t} that uses new variables \bar{c} translates a new constraint in $\widehat{\delta}$ of the form $\gamma = \exists \bar{c}. r(\bar{t}) \rightarrow \varphi(\bar{t}, \bar{c})$. Since $\forall \bar{x}. r(\bar{x}) \leftrightarrow \psi(\bar{x}) \Rightarrow \gamma$, we have $\forall \bar{x}. r(\bar{x}) \leftrightarrow \psi(\bar{x}) \wedge \delta \Rightarrow \widehat{\delta}$, which implies the condition of Definition 4. \square

Remark 2. The combination of adding $\forall \bar{x}. \psi(\bar{x}) \rightarrow r(\bar{x})$ to the verification condition and allowing the user to perform local instantiations on the program variables is at least as powerful as rewriting program conditions, since any rewrite of $\psi(\bar{t})$ to $r(\bar{t})$ can be simulated by a local instantiation on \bar{t} .

Instantiations for the Invariant The mechanism of local instantiations is designed to support instantiations of $\forall \bar{x}. \psi(\bar{x}) \leftrightarrow r(\bar{x})$ required to prove that the invariant $I = \widehat{I}[\psi/r]$ is preserved by the program. This proof is carried out by showing that $(\forall \bar{x}. \psi(\bar{x}) \leftrightarrow r(\bar{x})) \wedge \widehat{I} \wedge \delta \wedge \neg(\widehat{I}[\psi/r])'$ is unsatisfiable. This may require instantiating the definition of r on Skolem constants that come from the negation of the invariant in the post-state. Thus, we extend our instrumentation method by adding new “program variables” that represent the elements of the domain on which the invariant is potentially violated in the post-state. For every universally quantified variable x in \widehat{I} , we add a special program variable sk_x which can be used in local instantiations, enhancing their power to prove that \widehat{I} is inductive.

Obtaining Deep Instantiations Applying local instrumentation on a tuple \bar{t} that consists of original program variables, or variables that represent Skolem constants, corresponds to instantiations of Bounded-Horizon with bound 1. However, once a local instantiation is performed, new program variables \bar{c} are added. Performing a local instantiation on these new variables now corresponds to instantiation from depth 2. By iteratively applying local instantiations, where each iteration adds new program variables, we can thus obtain quantifier instantiations of arbitrary depth.

B.2 Illustrating Example

Fig. 4 illustrates the local instantiation procedure on the example of Fig. 2. Recall that the $\forall^* \exists^*$ invariant I of Fig. 2 is not provable using Bounded-Horizon of bound 1, but is provable using Bounded-Horizon of bound 2. The instrumentation presented in Fig. 4

```

/@  $r_1(x, y) \equiv \exists z. req(x, z) \wedge db(z, y)$ 
/@  $r_2(x, y) \equiv \exists z. t(x, z) \wedge req(z, y)$ 
/@  $r_3(x, y) \equiv \exists z. db\_req(x, z) \wedge db(z, y)$ 
/@ Invariant  $\widehat{I} = \forall u, p. resp(u, p) \rightarrow r_1(u, p) \wedge$ 
/@  $\forall id, q. db\_req(id, q) \rightarrow r_2(id, q) \wedge$ 
/@  $\forall id, p. db\_resp(id, p) \rightarrow r_3(id, p) \wedge$ 
/@  $\forall id, u_1, u_2. t(id, u_1) \wedge t(id, u_2) \rightarrow u_1 = u_2$ 
action server_process_db_response( $id, p$ ) {
  # instantiate  $r_3$  on ( $id, sk_p$ ) (depth 1)
  /@  $c_1 := *$ ;
  /@ assume  $r_3(id, sk_p) \rightarrow db\_req(i, c_1) \wedge db(c_1, sk_p)$ ;
  # instantiate  $r_2$  on ( $id, c_1$ ) (depth 2)
  /@  $c_2 := *$ ;
  /@ assume  $r_2(id, c_1) \rightarrow t(id, c_2) \wedge req(c_2, c_1)$ ;
  ...
}
action check( $u, p$ ) {
  # instantiate  $r_1$  on ( $u, p$ ) (depth 1)
  /@  $c_3 := *$ ;
  /@ assume  $r_1(u, p) \rightarrow req(u, c_3) \wedge db(c_3, p)$ ;
  ...
}

```

Fig. 4. An illustration of instrumentation by local instantiations for the example of Fig. 2. The instrumentation adds three instrumentation relations r_1, r_2, r_3 , and performs three local instantiations in order to prove that the invariant is inductive. Note that an instantiation depth of 2 is used, in accordance with the fact that the original invariant is provable using bound 2 but not bound 1. The complete model corresponding to this Figure appears in [1] (file `client_server_db_instr.ivy`).

introduces three instrumentation relations to encode the existential parts of I , thereby producing the instrumented universal invariant \widehat{I} .

To prove the inductiveness of \widehat{I} , we use local instantiations in both the `server_process_db_response` action and the `check` action. In `server_process_db_response`, we instantiate the definition of r_3 on (id, sk_p) , and assign the existential witness to c_1 . Intentionally, c_1 gets the request that was sent from the server to the DB that led to the response sk_p being sent from the DB to the server. sk_p is the response that supposedly causes a violation of the invariant when the action `server_process_db_response` is executed (the instantiations are used to prove that a violation does not occur). This instantiation is of depth 1. Next, we make an instantiation of the definition of r_2 on (id, c_1) and obtain a new existential witness c_2 . The use of c_1 here makes this instantiation depth 2. The `check` action includes another instantiation of depth 1, which is simply used to prove that the `abort` cannot happen (similarly to rewriting a program condition).

The reader can observe that the local instantiations introduced during the instrumentation process closely correspond to the instantiations required to prove that the original $\forall^* \exists^*$ invariant I (see Section 3).

We note that the process of instrumentation by local instantiations discussed here is different in spirit from those of Section 4: instrumentation by local instantiation consists of almost nothing but adding existential quantifiers to the transition relation, as opposed to the condition in Definition 6 where we do not allow the instrumentation to add new existential quantifiers.