

Logical Characterizations of Heap Abstractions

GRETA YORSH

School of Comp. Sci., Tel-Aviv University

THOMAS REPS

Comp. Sci. Dept., University of Wisconsin

MOOLY SAGIV

School of Comp. Sci., Tel-Aviv University

and

REINHARD WILHELM

Informatik, Univ. des Saarlandes

Shape analysis concerns the problem of determining “shape invariants” for programs that perform destructive updating on dynamically allocated storage. In recent work, we have shown how shape analysis can be performed using an abstract interpretation based on 3-valued first-order logic. In that work, concrete stores are finite 2-valued logical structures, and the sets of stores that can possibly arise during execution are represented (conservatively) using a certain family of finite 3-valued logical structures. In this paper, we show how 3-valued structures that arise in shape analysis can be characterized using formulas in first-order logic with transitive closure. We also define a non-standard (“supervaluational”) semantics for 3-valued first-order logic that is more precise than a conventional 3-valued semantics, and demonstrate that the supervaluational semantics can be implemented using existing theorem provers.

Categories and Subject Descriptors: D.2.4 [Software/Program Verification]: General

General Terms: Verification

Additional Key Words and Phrases: logic, characterization, canonical abstraction, shape analysis

1. INTRODUCTION

Abstraction and abstract interpretation [Cousot and Cousot 1977] are key tools for automatically verifying properties of systems, both for hardware systems [Clarke et al. 1994; Dams 1996] and software systems [Nielson et al. 1999]. In abstract interpretation, sets of concrete stores are represented in a conservative manner by abstract values (as explained below). Each transition of the system is given an interpretation over abstract values that is conservative with respect to its interpretation over corresponding sets of concrete stores; that is, the result of “executing” a transition must be an abstract value that describes a superset of the concrete stores that actually arise. This methodology guarantees that the results of abstract interpretation overapproximate the sets of concrete stores that actually

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2006 ACM 1529-3785/06/0700-0001 \$5.00

arise at each point in the system.

One issue that arises when abstraction is employed concerns the *expressiveness* of the abstraction method: “What collections of concrete states can be expressed exactly using the given abstraction method?” A second issue that arises when abstraction is employed is how to *extract information* from an abstract value. For instance, this is a fundamental problem for clients of abstract interpretation, such as verification tools, program optimizers, program-understanding tools, etc., which need to be able to interpret what an abstract value means. An abstract value a represents a set of concrete stores X ; ideally, a query φ should return an answer that summarizes the result of posing φ against each concrete store $S \in X$:

- If φ is true for each S , the summary answer should be “true”.
- If φ is false for each S , the summary answer should be “false”.
- If φ is true for some $S \in X$ but false for some $S' \in X$, the summary answer should be “unknown”.

This paper presents results on both of these questions, for a class of abstractions that originally arose in work on the problem of shape analysis [Jones and Muchnick 1981; Chase et al. 1990; Sagiv et al. 2002]. Shape analysis concerns the problem of finding “shape descriptors” that characterize the shapes of the data structures that a program’s pointer variables point to. Shape analysis is one of the most challenging problems in abstract interpretation because it generally deals with programs written in languages like C, C++, and Java, which allow (i) dynamic allocation and deallocation of cells from the heap, (ii) destructive updating of structure fields, and, in the case of Java, (iii) dynamic creation and destruction of threads. This combination of features creates considerable difficulties for any abstract-interpretation method.

The motivation for the present paper was to understand the expressiveness of the shape abstractions defined in [Sagiv et al. 2002]. In that work, concrete stores are finite 2-valued logical structures, and the sets of stores that can possibly arise during execution are represented (conservatively) using a certain family of finite 3-valued logical structures. In this setting, an abstract value is a set of 3-valued logical structures.

Because the notion of abstraction used in [Sagiv et al. 2002] is based on logical structures, our results are actually more broadly applicable than shape-analysis problems. For example, it was applied to verification of sorting algorithms [Lev-Ami et al. 2000]; showing absence of concurrent modification exception [Ramalingam et al. 2002]; correct usage of JDBC, I/O streams, Java collections and iterators [Yahav and Ramalingam 2004]; correctness of concurrent queue algorithms [Yahav and Sagiv 2003]; modelling concurrency in Java programs, which contain dynamic creation of objects and threads [Yahav 2001]; analyzing processes in ambient calculus [Nielson et al. 2000]; and reducing space consumption of Java programs via compile-time memory management, with application to JavaCard programs [Shaham et al. 2003].

In fact, our results apply to any abstraction in which concrete states of a system are represented by finite 2-value logical structure and abstraction is performed via the mechanisms described in Sections 2 and 3. The approach taken in the paper should also be relevant for addressing expressibility issues for a number of other abstractions that are related to [Sagiv et al. 2002], including [McMillan 1999; Kuncak et al. 2002; Godefroid and Jagadeesan 2003; Huth et al. 2001; Clarke et al. 1994; Clarke et al. 2000], as well as for

the *allocation-site* abstraction—often used in points-to analysis [Andersen 1993; Steensgaard 1996; Shapiro and Horwitz 1997; Fähndrich et al. 1998; Su et al. 2000; Das 2000; Heintze and Tardieu 2001]—in which all objects allocated at a single statement are represented by a single “abstract memory object” [Jones and Muchnick 1982; Chase et al. 1990]. Throughout the paper, however, we use shape-analysis examples to illustrate the concepts discussed.

The paper investigates the expressiveness of finite 3-valued structures by giving a logical characterization of these structures; that is, we examine the question

For a given 3-valued structure S , under what circumstances is it possible to create a formula $\hat{\gamma}(S)$, such that S^{\natural} satisfies $\hat{\gamma}(S)$ exactly when S^{\natural} is a 2-valued structure that S represents? I.e., $S^{\natural} \models \hat{\gamma}(S)$ iff S represents S^{\natural} .

This paper presents two results concerning this question:

- It is not possible to give a formula $\hat{\gamma}(S)$ written in first-order logic with transitive closure for an arbitrary structure S (unless $NL = NP$, see Section 3). However, it is always possible for a well-defined class of 3-valued structures. (This class includes all the 3-valued structures that have been shown to be useful for shape analysis [Sagiv et al. 2002].)
- Moreover, it is always possible to give a $\hat{\gamma}(S)$ in general, using a more powerful formalism, namely, monadic second-order formulas.

The ability to write a formula $\hat{\gamma}(S)$ that exactly captures what S represents provides a fundamental tool for improving TVLA [Lev-Ami and Sagiv 2000] by the use of symbolic methods. The current TVLA system performs iterative fixed-point computations and yields at every program point a set of 3-valued structures, which represent a superset of all possible stores that can arise at this point in any execution. However, TVLA suffers from two limitations: (i) it is not always as precise as possible (as explained below); (ii) it does not scale to handle large programs, because the worst-case complexity of the algorithm is doubly-exponential in certain parameters (typically, the number of program variables).

The contributions of this paper lay the required groundwork for using symbolic techniques to address both of these limitations. The ability to characterize a 3-valued structure S by a formula $\hat{\gamma}(S)$ is a key step toward harnessing a standard (2-valued) theorem prover to aid in abstract interpretation:

- Computing the effect of a program statement on an abstract value in the most-precise way possible for a given shape-analysis abstraction.
- Developing a modular shape-analysis by using *assume-guarantee* reasoning. The idea is to allow arbitrary first-order formulas to be used to express pre- and post-conditions, thereby enabling the code of each procedure to be analyzed once for all potential contexts. This allows to use shape analysis for applications in which not all the source code is available. This becomes specifically profitable for recursive procedures since it saves the need to iterate shape analysis.

These methods are the subject of [Yorsh et al. 2004; Lam et al. 2005].

Another contribution of this paper directly addresses the first of the aforementioned limitations of TVLA’s current technique. We give a procedure for extracting information from a 3-valued logical structure S in the most-precise way possible. That is, we give a nonstandard way to check if a formula φ holds in S :

- If $\hat{\gamma}(S) \Rightarrow \varphi$ is valid, i.e., holds in all 2-valued structures, we know that φ evaluates to 1 in all the 2-valued structures represented by S .
- If $\hat{\gamma}(S) \Rightarrow \neg\varphi$ is valid, we know that φ evaluates to 0 in all the 2-valued structures represented by S .
- Otherwise we know that there exists a 2-valued structure represented by S where φ evaluates to 1, and there exists another 2-valued structure represented by S where φ evaluates to 0.

This method represents the most-precise way of extracting information from a 3-valued logical structure; in particular, whenever this method returns 1/2 (standing for “unknown”), any sound method for extracting information from S must also return 1/2. This is in contrast with the techniques used in [Sagiv et al. 2002], which can return 1/2 even when all the 2-valued structures represented by S have the value 1 (or all have the value 0).

For practical purposes, the success of using symbolic methods depends on having a terminating theorem prover. Although the validity question is undecidable for first-order logic with transitive closure, several theorem provers for first-order logic have been created. In this paper, we report on two experiments in which we used these tools to implement symbolic procedures for extracting information from a 3-valued structure in the most-precise way possible. We also performed several successful experiments with other symbolic operations [Yorsh et al. 2004; Erez 2004]. Although these experiments are rather preliminary, we believe that this approach can be made to work in practice. For example, there has been some progress recently in using SPASS, including the use of transitive closure [Lev-Ami et al. 2005]. Also, in [Immerman et al. 2004a], we have identified a decidable subset of first-order logic with transitive closure that is useful for shape analysis. We define conditions under which $\hat{\gamma}$ can be expressed in that logic (Section 5.2). We are also investigating other decidable logics, as well.

The remainder of the paper is organized as follows. Section 2 defines our terminology, and explains the use of 3-valued structures as abstractions of 2-valued structures. Section 3 presents the results on the expressiveness of 3-valued structures, and gives an algorithm for generating $\hat{\gamma}$ for certain families of 3-valued structures. Section 4 discusses the problem of reading out information from a 3-valued structure in the most-precise way possible. Section 5 discusses the applications of $\hat{\gamma}$ to program analysis and some implementation issues. Section 6 discusses related work. Appendix A defines an alternative abstract domain for shape analysis, based on canonical abstraction, and the $\hat{\gamma}$ operation for that domain. Appendix B shows how to characterize general 3-valued structures. Appendix C contains the details for one of the paper’s examples. The proofs appear in Appendix D.

2. PRELIMINARIES

Section 2.1 defines the syntax and standard Tarskian semantics of first-order logic with transitive closure and equality. Section 2.2 introduces *integrity formulas*, which exclude structures that do not represent a potential store. Section 2.3 introduces 3-valued logical structures, which extend ordinary logical structures with an extra value, 1/2, which represents “unknown” values that arise when several concrete nodes are represented by a single abstract node. The powerset of 3-valued structures forms an abstract domain, which is related to the concrete domain consisting of the powerset of 2-valued structures via *embedding*, as described in Section 2.4.

```

/* list.h */
typedef struct node {
    struct node *n;
    int data;
} *List;

/* insert.c */
#include "list.h"
void insert(List x, int d) {
    List y, t, e;
    assert(acyclic_list(x) && x != NULL);
    y = x;
    while (y->n != NULL && ...)
        y = y->n;
    t = malloc();
    t->data = d;
    e = y->n;
    t->n = e;
    y->n = t;
}

```

(a) (b)

Fig. 1. (a) Declaration of a linked-list data type in C. (b) A C function that searches a list pointed to by parameter x , and splices in a new element.

Fig. 1(a) shows the declaration of a linked-list data type in C, and Fig. 1(b) shows a C program that searches a list and splices a new element into the list. This program will be used as a running example throughout this paper.

2.1 Syntax and Semantics of First-Order Formulas with Transitive Closure

We represent concrete stores by ordinary 2-valued logical structures over a fixed finite set of predicate symbols $\mathcal{P} = \{eq, p_1, \dots, p_n\}$, where eq is a designated binary predicate, denoting equality of nodes. We also use $maxR$ to denote the maximal arity of the predicates in \mathcal{P} . Without loss of generality we exclude constant and function symbols from the logic.¹

EXAMPLE 2.1. *Table 1 lists the set of predicates used in the running example. The unary predicates x , y , t , and e correspond to the program variables x , y , t , and e , respectively. The binary predicate n corresponds to the n fields of `List` elements. The unary predicate is (“is shared”) captures “heap sharing”, i.e., `List` elements pointed to by more than one field. (It was introduced in [Chase et al. 1990] to capture list and tree data structures.) The unary predicates r_x , r_y , r_t , and r_e hold for heap nodes reachable from the program variables x , y , t , and e , respectively. A heap node u is said to be reachable from a program variable if the variable points to a heap node u' , and it is possible to go from u' to u by following zero or more n -links. Reachability is defined in term of the reflexive transitive closure of the predicate n .*

The notion of reachability plays a crucial role in defining abstractions that are useful for proving program properties in practice. For instance, it may have the effect of preventing disjoint lists from being collapsed in the abstract representation. This may significantly improve the precision of the answers obtained by a program analysis.

¹Constant symbols can be encoded via unary predicates, and n -ary functions via $(n + 1)$ -ary predicates.

Predicate	Intended Meaning
$eq(v_1, v_2)$	Do v_1 and v_2 denote the same heap node?
$q(v)$	Does pointer variable \mathfrak{q} point to node v ?
$n(v_1, v_2)$	Does the n field of v_1 point to v_2 ?
$is(v)$	Is v pointed to by more than one field ?
$r_q(v)$	Is the node v reachable from \mathfrak{q} ?

Table I. The set of predicates for representing the stores manipulated by programs that use the `List` data-type from Fig. 1(a). q denotes an arbitrary predicate in the set $PVar$, which contains a predicate for each program variable of type `List`. In the case of `insert`, $PVar = \{x, y, t, e\}$.

We define first-order formulas inductively over the **vocabulary** \mathcal{P} using the logical connectives \vee and \neg , the quantifier \exists , and the operator ‘ TC ’ in the standard way:

$$\varphi ::= \mathbf{0} \mid \mathbf{1} \mid p(v_1, \dots, v_k) \mid (\neg\varphi_1) \mid (\varphi_1 \vee \varphi_2) \mid (\exists v_1 : \varphi_1) \mid (TC\ v_1, v_2 : \varphi_1)(v_3, v_4)$$

where $p \in \mathcal{P}$; v_i are variables; φ, φ_i are formulas

The set of free variables of a formula is defined as usual. A formula is **closed** when it has no free variables. The operator ‘ TC ’ denotes transitive closure. If φ_1 is a formula with free variables V , then $(TC\ v_1, v_2 : \varphi_1)(v_3, v_4)$ is a formula with free variables $(V - \{v_1, v_2\}) \cup \{v_3, v_4\}$.

We use several shorthand notations: $\varphi_1 \Rightarrow \varphi_2 \stackrel{\text{def}}{=} (\neg\varphi_1 \vee \varphi_2)$; $\varphi_1 \wedge \varphi_2 \stackrel{\text{def}}{=} \neg(\neg\varphi_1 \vee \neg\varphi_2)$; $\varphi_1 \Leftrightarrow \varphi_2 \stackrel{\text{def}}{=} (\varphi_1 \Rightarrow \varphi_2) \wedge (\varphi_2 \Rightarrow \varphi_1)$; and $\forall v : \varphi \stackrel{\text{def}}{=} \neg\exists v : \neg\varphi$. The transitive closure of a binary predicate p is $p^+(v_3, v_4) \stackrel{\text{def}}{=} (TC\ v_1, v_2 : p(v_1, v_2))(v_3, v_4)$. The *reflexive* transitive closure of a binary predicate p is $p^*(v_3, v_4) \stackrel{\text{def}}{=} ((TC\ v_1, v_2 : p(v_1, v_2))(v_3, v_4)) \vee eq(v_3, v_4)$. The order of precedence among the connectives, from highest to lowest, is as follows: \neg , \wedge , \vee , ‘ TC ’, \forall , and \exists . We drop parentheses wherever possible, except for emphasis.

Definition 2.1. 2-valued Logical Structures Let \mathcal{P}_i denote the set of predicate symbols with arity i . A **logical structure over** \mathcal{P} is a pair $S = \langle U, \iota \rangle$ in which

— U is a (possibly infinite) set of nodes.

— ι is the interpretation of predicate symbols, i.e., for every predicate symbol $p \in \mathcal{P}_i$, $\iota(p) : U^i \rightarrow \{0, 1\}$ determines the tuples for which p holds. Also, $\iota(eq)$ is the interpretation of equality, i.e., $\iota(eq)(u_1, u_2) = 1$ iff $u_1 = u_2$.

Below we define the standard Tarskian semantics for first-order logic.

Definition 2.2. Semantics of First-Order Logical Formulas Consider a logical structure $S = \langle U, \iota \rangle$. An **assignment** Z is a function that maps free variables to nodes (i.e., an assignment has the functionality $Z : \{v_1, v_2, \dots\} \rightarrow U$). An assignment that is defined on all free variables of a formula φ is called **complete** for φ . In the sequel, we assume that every assignment Z that arises in connection with the discussion of some formula φ is complete for φ . We say that S and Z **satisfy** a formula φ (denoted by $S, Z \models \varphi$) when one of the following holds:

— $\varphi \equiv \mathbf{1}$

— $\varphi \equiv p(v_1, v_2, \dots, v_i)$ and $\iota(p)(Z(v_1), Z(v_2), \dots, Z(v_i)) = 1$.

— $\varphi \equiv \neg\varphi_0$ and $S, Z \models \varphi_0$ does not hold.

- $\varphi \equiv \varphi_1 \vee \varphi_2$, and either $S, Z \models \varphi_1$ or $S, Z \models \varphi_2$.
- $\varphi \equiv \exists v_1 : \varphi_1$ and there exists a node $u \in U$, $m \geq 2$, such that $S, Z[v_1 \mapsto u] \models \varphi_1$.
- $\varphi \equiv (TC \ v_1, v_2 : \varphi_1)(v_3, v_4)$ and there exists $u_1, u_2, \dots, u_m \in U$, $m \geq 2$, such that $Z(v_3) = u_1$, $Z(v_4) = u_m$ and for all $1 \leq i < m$, $S, Z[v_1 \mapsto u_i, v_2 \mapsto u_{i+1}] \models \varphi_1$.

For a closed formula φ , we will omit the assignment in the satisfaction relation, and merely write $S \models \varphi$.

2.2 Integrity Formula

Because not all logical structures represent stores, we use a designated closed formula F , called the *integrity formula*,² to exclude structures that are not of interest; in our application, such structures are ones that do not correspond to possible stores. This allows us to restrict the set of structures to the ones satisfying F .

Definition 2.3. A structure S is **admissible** if $S \models F$.

In the rest of the paper, we assume that we work with a fixed integrity formula F . All our notations are parameterized by \mathcal{P} and F .

EXAMPLE 2.2. For the `List` data type, there are four conditions that define the admissible structures. At any time during execution,

- (a). each program variable can point to at most one heap node.
- (b). the `n` field of a heap node can point to at most one heap node.
- (c). predicate *is* (“is shared”) holds for exactly those nodes that have two or more predecessors.
- (d). the reachability predicate for each variable α holds for exactly those nodes that are reachable from program variable α .

The set $PVar$ contains a predicate for each program variable of type `List`; in the case of `insert`, $PVar = \{x, y, t, e\}$. Thus, the integrity formula F_{List} for the `List` data-type is:

$$\begin{aligned}
 & \bigwedge_{p \in PVar} \forall v_1, v_2 : p(v_1) \wedge p(v_2) \Rightarrow eq(v_1, v_2) & (a) \\
 \wedge & \quad \forall v, v_1, v_2 : n(v, v_1) \wedge n(v, v_2) \Rightarrow eq(v_1, v_2) & (b) \\
 \wedge & \quad \forall v : is(v) \iff \exists v_1, v_2 : \neg eq(v_1, v_2) \wedge n(v_1, v) \wedge n(v_2, v) & (c) \\
 \wedge & \quad \bigwedge_{q \in PVar} \forall v : r_q(v) \iff \exists v_1 : q(v_1) \wedge n^*(v_1, v) & (d)
 \end{aligned}$$

2.3 3-Valued Logical Structures and Embedding

In this section, we define 3-valued logical structures, which provide a way to represent a set of 2-valued logical structures in a compact and conservative way.

We say that the values 0 and 1 are *definite values* and that $1/2$ is an *indefinite value*, and define a partial order \sqsubseteq on truth values to reflect information content. $l_1 \sqsubseteq l_2$ denotes that l_1 possibly has more definite information than l_2 :

Definition 2.4. [Information Order]. For $l_1, l_2 \in \{0, 1/2, 1\}$, we define the **information order** on truth values as follows: $l_1 \sqsubseteq l_2$ if $l_1 = l_2$ or $l_2 = 1/2$.

²In [Sagiv et al. 2002] these are called “hygiene conditions”.




Fig. 2. (a),(b),(c) Examples of 2-valued structures representing linked-lists that are pointed to by program variable x , of length 2, 3, and 4, respectively. (d) S represents all lists that are pointed to by program variable x and that have at least two elements, including the lists represented by (a)-(c).

Definition 2.5. A **3-valued logical structure** over \mathcal{P} is the generalization of 2-valued structures given in Definition 2.1, in that predicates may have the value $1/2$. This means that $S = \langle U, \iota \rangle$ where for $p \in \mathcal{P}_i$, $\iota(p): (U^S)^i \rightarrow \{0, 1, 1/2\}$. In addition, (i) for all $u \in U^S$, $\iota^S(eq)(u, u) \sqsupseteq 1$, and (ii) for all $u_1, u_2 \in U^S$ such that u_1 and u_2 are distinct nodes, $\iota^S(eq)(u_1, u_2) = 0$.

A node $u \in U$ having $\iota^S(eq)(u, u) = 1/2$ is called a **summary node**. As we shall see, such a node may represent more than one node from a given 2-valued structure.

We denote the set of 2-valued logical structures by $2\text{-STRUCT}[\mathcal{P}]$. The set of 3-valued logical structures is denoted by $3\text{-STRUCT}[\mathcal{P}]$.

A 3-valued structure can be depicted as a directed graph, with nodes as graph nodes. A unary predicate p is represented in the graph by having a solid arrow from the predicate name p to node u for each node u for which $\iota(p)(u) = 1$. An arrow between two nodes indicates whether a binary predicate holds for the corresponding pair of nodes. An indefinite value of a predicate is shown by a dotted arrow; the value 1 is shown by a solid arrow; and the value 0 is shown by the absence of an arrow.

EXAMPLE 2.3. Fig. 2(d) shows a 3-valued structure that represents possible inputs of the *insert* program. This structure represents all lists that are pointed to by program variable x and have at least two elements. The structure has 2 nodes, u_1 and u_2 , where u_1 is the head of the list pointed to by x , and u_2 is a summary node (drawn as a double circle), which represents the tail of the list. Predicate r_x holds for u_1 and u_2 , indicating that all elements of the list are reachable from x . Other unary predicates are not shown, indicating that their values are 0 for all nodes, i.e., the program variables y , e , and t are NULL, and there is no sharing in the list. The dotted edge from u_1 to u_2 indicates that there may be n -links from the head of the list to some elements in the tail. In fact, the (u_1, u_2) -edge represents exactly one n -link that points to exactly one list element, because of conjunct (b) of the integrity formula Example 2.2. In contrast, the dotted self-loop on u_2 represents all n -links that may occur in the tail.

2.4 Embedding Order

We define the *embedding ordering* on structures as follows:

Definition 2.6. Let $S = \langle U^S, \iota^S \rangle$ and $S' = \langle U^{S'}, \iota^{S'} \rangle$ be two logical structures, and let $f: U^S \rightarrow U^{S'}$ be a surjective. We say that f **embeds** S in S' (denoted by $S \sqsubseteq^f S'$) if for every predicate symbol $p \in \mathcal{P}_i$ and all $u_1, \dots, u_i \in U^S$,

$$\iota^S(p)(u_1, \dots, u_i) \sqsubseteq \iota^{S'}(p)(f(u_1), \dots, f(u_i)) \quad (1)$$

We say that S **can be embedded in** S' (denoted by $S \sqsubseteq S'$) if there exists a function f such that $S \sqsubseteq^f S'$.

EXAMPLE 2.4. Fig. 2(a)-(c) show some of the 2-valued structures that can be embedded into the 3-valued structure S shown in Fig. 2(d). The function that embeds S_a into S maps the node $u_i^{\natural} \in U^{S_a}$ to $u_i \in U^S$, for $i = 1, 2$. The function that embeds S_b into S maps the node $u_1^{\natural} \in U^{S_b}$ to $u_1 \in U^S$, and both $u_2^{\natural}, u_3^{\natural} \in U^{S_b}$ to $u_2 \in U^S$. Also, Eq. (1) holds, because whenever a predicate has a definite value in S , the corresponding predicate in S_b has the same value. For example, $\iota^S(x)(u_2)$ is 0 and $f(u_2^{\natural}) = f(u_3^{\natural}) = u_2$, and both $\iota^{S_b}(x)(u_2^{\natural})$ and $\iota^{S_b}(x)(u_3^{\natural})$ are 0. Similarly, $\iota^S(r_x)(u_2) = 1$, and both $\iota^{S_b}(r_x)(u_2^{\natural})$ and $\iota^{S_b}(r_x)(u_3^{\natural})$ are 1. For a binary predicate, $\iota^S(n)(u_2, u_1) = 0$, and both $\iota^{S_b}(n)(u_2^{\natural}, u_1^{\natural})$ and $\iota^{S_b}(n)(u_3^{\natural}, u_1^{\natural})$ are 0.

Remark. Embedding can be viewed as a variant of homomorphism [Hell and Nesetril 2004]. In cases where S is a 2-valued structure (i.e., all predicates in S have definite values, including eq , which is interpreted as standard equality), checking whether a 2-valued structure S' embeds into S is equivalent to checking whether there is an isomorphism between S' and S . In cases where all nodes in S are summary nodes (i.e., for all $u \in U^S$, $\iota^S(eq)(u, u) = 1/2$), and all other values of predicates are definite, embedding is equivalent to strong homomorphism. In cases where all nodes in S are summary nodes and all other values of predicates are either 0 or $1/2$, embedding is equivalent to homomorphism. In all other cases, i.e. when a predicate value for some tuple in S is 1, embedding generalizes the notion of homomorphism.

Remark. In Definition 2.6, we require that f be surjective in order to guarantee that a quantified formula, such as $\exists v : \varphi$, has consistent values in two 3-valued structures S and S' related by embedding. For example, if f were not surjective, then there could exist an individual $u' \in U^{S'}$, not in the range of f , such that the value of S' on φ is 1 when v is assigned to u' . This would permit there to be structures S and S' for which the value of $\exists v : \varphi$ on S is 0 but its value on S' is 1.

Concretization of 3-Valued Structures. Embedding allows us to define the (potentially infinite) set of concrete structures that a set of 3-valued structures represents:

Definition 2.7. Concretization of 3-Valued Structures For a set of structures $X \subseteq 3\text{-STRUCT}[\mathcal{P}]$, we denote by $\gamma(X)$ the set of 2-valued structures that X represents, i.e.,

$$\gamma(X) = \{S^{\natural} \in 2\text{-STRUCT}[\mathcal{P}] \mid \text{exists } S \in X \text{ such that } S^{\natural} \sqsubseteq S \text{ and } S^{\natural} \models F\} \quad (2)$$

Also, for a singleton set $X = \{S\}$ we write $\gamma(S)$ instead of $\gamma(X)$.

EXAMPLE 2.5. Example 2.4 shows that $S_a \sqsubseteq S$, $S_b \sqsubseteq S$, and $S_c \sqsubseteq S$ for the 2-valued structures in Figs. 2(a-c); also, the integrity formula is satisfied for S_a , S_b , and S_c . Therefore, S_a , S_b , and S_c are in the concretization of 3-valued structure S : $S_a, S_b, S_c \in \gamma(S)$. Note that the indefinite values of predicates in S allow the corresponding values in S_b to be either 0 or 1. In particular, $\iota^S(eq)(u_2, u_2) = 1/2$ reflects the fact that the abstract

node u_2 may represent more than one concrete node. Indeed, S_b contains two nodes, u_2^b and u_3^b , that are represented by $u_2 \in S$. Also, $\iota^S(eq)(u_2^b, u_3^b) = 0$, but $\iota^S(eq)(u_2^b, u_2^b) = 1$.

The abstract domain we consider is the powerset of 3-valued structures, where the ordering relation \sqsubseteq is defined as follows: for every two sets of 3-valued structures X_1 and X_2 , $X_1 \sqsubseteq X_2$ iff for all $S_1 \in X_1$ there exists $S_2 \in X_2$ such that S_1 is embedded into S_2 .

2.4.1 The Analysis Technique. The TVLA ([Lev-Ami and Sagiv 2000]) system carries out an abstract interpretation [Cousot and Cousot 1977] to collect a set of structures at each program point P . This involves finding the least fixed point of a certain set of equations. To ensure termination, the analysis is carried out with respect to a finite abstract domain, that is, the set of different structures is finite. When the fixed point is reached, the structures that have been collected at program point p describe a superset of all the concrete stores that can occur at p . To determine whether a query is always satisfied at p , one checks whether it holds in all of the structures that were collected there. Instantiations of this framework are capable of establishing nontrivial properties of programs that perform complex pointer-based manipulations of *a priori* unbounded-size heap-allocated data structures.

3. CHARACTERIZING 3-VALUED STRUCTURES BY FIRST-ORDER FORMULAS

This section presents our results on characterizing 3-valued structures using first-order formulas. Given a 3-valued structure S , the question that we wish to answer is whether it is possible to give a formula $\hat{\gamma}(S)$ that accepts exactly the set of 2-valued structures that S represents, i.e., $S^b \models \hat{\gamma}(S)$ iff $S^b \in \gamma(S)$.

This question has different answers depending on what assumptions are made. The task of generating a characteristic formula for a 3-valued structure S is challenging because we have to find a formula that identifies when embedding is possible, i.e., that is satisfied by exactly those 2-valued structures that embed into S . It is not always possible to characterize an *arbitrary* 3-valued structure by a first-order formula, i.e., there exists a 3-valued structure S for which there is no first-order formula with transitive closure that accepts exactly the set of 2-valued structures $\gamma(S)$.

For example, consider the 3-valued structure S shown in Fig. 3. The absence of a self loop on any of the three summary nodes implies that a 2-valued structure can be embedded into this structure if and only if it can be colored using 3 colors (Lemma D.1 in the appendix). It is well-known that there exists no first-order formula, even with transitive closure, that expresses 3-colorability of undirected graphs, unless $P = NP$ (e.g., see [Immerman 1999; Courcelle 1996]). Therefore, there is no first-order formula that accepts exactly the set $\gamma(S)$.

Remark. In fact, the condition is even stronger. First-order logic with transitive closure can only express non-deterministic logspace (NL) computations, thus, the NP-complete problem of 3-colorability is not expressible in first-order logic, unless $NL = NP$. It is shown in [Immerman 1999] using an ordering relation on the nodes. In our context, without the ordering, the logic is less expressive. Thus, the condition under which 3-colorability is expressible is even stronger than $NL = NP$. We believe that there is an example of a 3-valued structure that is not expressible in the logic, independently of the question whether $P = NP$. However, it is not the main focus of the current paper.




Fig. 3. A 3-valued structure that represents 3-colorable undirected graphs. A 2-valued structure can be embedded into this structure if and only if it can be colored using 3 colors.

3.1 FO-Identifiable Structures

Intuitively, the difficulty in characterizing 3-valued structures is how to uniquely identify the correspondence between concrete and abstract nodes using a first-order formula. Fortunately, as we will see, for the subclass of 3-valued structures used in shape analysis (also known as “bounded structures”), the correspondence can be easily defined using first-order formulas. The bounded structures are a subclass of the 3-valued structures in which it is possible to identify uniquely each node using a first-order formula.

Definition 3.1. A 3-valued structure S is called **FO-identifiable** if for every node $u \in U^S$ there exists a first-order formula $\text{node}_u^S(w)$ with designated free variable w such that for every 2-valued structure S^\natural that embeds into S using a function f , for every concrete node $u^\natural \in U^{S^\natural}$ and for every node $u_i \in U^S$:

$$f(u^\natural) = u_i \iff S^\natural, [w \mapsto u^\natural] \models \text{node}_{u_i}^S(w) \quad (3)$$

The idea behind this definition is to have a formula that uniquely identifies each node u of the 3-valued structure S . This will be used to identify the set of nodes of a 2-valued structure that are mapped to u by embedding. In other words, a concrete node u^\natural satisfies the *node* formula of at most one abstract node, as formalized by the lemma:

LEMMA 3.2. *Let S be an FO-identifiable structure, and let $u_1, u_2 \in S$ be distinct nodes. Let S^\natural be a 2-valued structure that embeds into S and let $u^\natural \in S^\natural$. At most one of the following hold:*

- (1) $S^\natural, [w \mapsto u^\natural] \models \text{node}_{u_1}^S(w)$
- (2) $S^\natural, [w \mapsto u^\natural] \models \text{node}_{u_2}^S(w)$

Remark. Definition 3.1 can be generalized to handle arbitrary 2-valued structures, by also allowing extra designated free variables for every concrete node and using equality to check if the concrete node is equal to the designated variable: $\text{node}_{u_i}^S(w, v_1, \dots, v_n) \stackrel{\text{def}}{=} w = v_i$. However, the equality formula cannot be used to identify nodes in a 3-valued structure because equality evaluates to $1/2$ on summary nodes.

We now introduce a standard concept for turning valuations into formulas.

Definition 3.3. For a predicate p of arity k and truth value $B \in \{0, 1, 1/2\}$, we define the formula $p^B(v_1, v_2, \dots, v_k)$ to be the **characteristic formula of B for p** , by

$$\begin{aligned} p^0(v_1, v_2, \dots, v_k) &\stackrel{\text{def}}{=} \neg p(v_1, v_2, \dots, v_k) \\ p^1(v_1, v_2, \dots, v_k) &\stackrel{\text{def}}{=} p(v_1, v_2, \dots, v_k) \\ p^{1/2}(v_1, v_2, \dots, v_k) &\stackrel{\text{def}}{=} 1 \end{aligned}$$

The main idea in the above definition is that, for $B \in \{0, 1\}$, p^B holds when the value of p is B , and for $B = 1/2$ the value of p is unrestricted. This is formalized by the following lemma:

LEMMA 3.4. *For every 2-valued structure S^\natural and assignment Z*

$$S^\natural, Z \models p^B(v_1, \dots, v_k) \text{ iff } \iota^{S^\natural}(p)(Z(v_1), \dots, Z(v_k)) \sqsubseteq B$$

Definition 3.1 is not a constructive definition, because the premises range over arbitrary 2-valued structures and arbitrary embedding functions. For this reason, we now introduce a testable condition that implies FO-identifiability.

Bounded Structures. The following subclass of 3-values structures was defined in [Sagiv et al. 1999];³ the motivation there was to guarantee that shape analysis was carried out with respect to a finite set of abstract structures, and hence that the analysis would always terminate.

Definition 3.5. A **bounded structure** over vocabulary \mathcal{P} is a structure $S = \langle U^S, \iota^S \rangle$ such that for every $u_1, u_2 \in U^S$, where $u_1 \neq u_2$, there exists a predicate symbol $p \in \mathcal{P}_1$ such that (i) $\iota^S(p)(u_1) \neq \iota^S(p)(u_2)$ and (ii) both $\iota^S(p)(u_1)$ and $\iota^S(p)(u_2)$ are not $1/2$, i.e., $\iota^S(p)(u_1), \iota^S(p)(u_2) \in \{0, 1\}$.

Intuitively, for each pair of nodes in a bounded structure, there is at least one predicate that has different definite values for these nodes. Thus, there is a finite number of different bounded structures (up to isomorphism).

The following lemma shows that bounded structures are FO-identifiable using formulas over unary predicates only (denoted by \mathcal{P}_1):

LEMMA 3.6. *Every bounded 3-valued structure S is FO-identifiable, where*

$$node_{u_i}^S(w) \stackrel{\text{def}}{=} \bigwedge_{p \in \mathcal{P}_1} p^{\iota^S(p)(u_i)}(w) \quad (4)$$

EXAMPLE 3.1. *The first-order node formulas for the structure S shown in Fig. 2, are:*

$$\begin{aligned} node_{u_1}^S(w) &= x(w) \wedge r_x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ node_{u_2}^S(w) &= \neg x(w) \wedge r_x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \end{aligned}$$

Remark. In the case that S is a bounded 2-valued structure, the definition of a bounded structure becomes trivial. The reason is that every node in S can be named by a quantifier-free formula built from unary predicates. This is essentially the same as saying that every node can be named by a constant. If structure S' embeds into S , then S' must be isomorphic to S , therefore it is possible to name all nodes of S' by the same constants. However, this restricted case is not of particular interest for us, because, to guarantee termination, shape analysis operates on structures that contain summary nodes and indefinite values.

³This definition of bounded structures was given in [Sagiv et al. 1999]; it is slightly more restrictive than the one given in [Sagiv et al. 2002; Lev-Ami 2000], which did not impose requirement 3.5(ii). However, it does not limit the set of problems handled by our method, if the structure that is bounded in the weak sense is also FO-identifiable.




Fig. 4. A 3-valued structure S' is FO-identifiable, but not bounded.

In the case that S contains a summary node, a structure S' that embeds into S may have an unbounded number of nodes; hence the nodes of S' cannot be named by a finite set of constants in the language.

We already know of interesting cases of FO-identifiable structures that are not bounded, which can be used to generalize the abstraction defined in [Sagiv et al. 1999]:

EXAMPLE 3.2. *The 3-valued structure S' in Fig. 4 is FO-identifiable by:*

$$\begin{aligned}
 \text{node}_{u_1}^{S'}(w) &\stackrel{\text{def}}{=} x(w) \wedge r_x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\
 &\quad \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\
 \text{node}_{u_2}^{S'}(w) &\stackrel{\text{def}}{=} \underline{\exists w_1 : x(w_1) \wedge n(w_1, w)} \wedge \neg x(w) \wedge r_x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\
 &\quad \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\
 \text{node}_{u_3}^{S'}(w) &\stackrel{\text{def}}{=} \neg(\exists w_1 : x(w_1) \wedge n(w_1, w)) \wedge \neg x(w) \wedge r_x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\
 &\quad \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w)
 \end{aligned}$$

However, S' is not a bounded structure because nodes u_2 and u_3 have the same values of unary predicates. To distinguish between these nodes, we extended $\text{node}_{u_2}^{S'}(w)$ with the underlined subformula, which captures the fact that only u_2 is directly pointed to by an n -edge from u_1 .

It can be shown that every FO-identifiable structure can be converted into a bounded structure by introducing more instrumentation predicates. For methodological reasons, we use the notion of FO-identifiable which directly capture the ability to uniquely identify embedding via (FO) formulas.⁴ One of the interesting features of FO-identifiable structures is that the structures generated by a common TVLA operation “focus”, defined in [Lev-Ami 2000], are all FO-identifiable (see Lemma D.2 in Appendix D). For example, Fig. 4 shows the structure S' , which is one of the structures resulting from applying the “focus” operation to the structure S from Fig. 2(d) with the formula $\exists v_1, v_2 : x(v_1) \wedge n(v_1, v_2)$. S' is FO-identifiable, but not bounded. However, structures like the one shown in Fig. 3 are not FO-identifiable unless $P = NP$.

3.2 Characterizing FO-identifiable structures

To characterize an FO-identifiable 3-valued structure, we must ensure

- (1) the existence of a surjective embedding function.
- (2) that every concrete node is represented by some abstract node.
- (3) that corresponding concrete and abstract predicate values meet the embedding condition of Eq. (1).

⁴In subsequent sections, we redefine this notion to capture other classes of structures.

Definition 3.7. First-order Characteristic Formula Let $S = \langle U = \{u_1, u_2, \dots, u_n\}, \iota \rangle$ be an FO-identifiable 3-valued structure.

We define the **totality characteristic formula** to be the closed formula:

$$\xi_{total}^S \stackrel{\text{def}}{=} \forall w : \bigvee_{i=1}^n \text{node}_{u_i}^S(w) \quad (5)$$

We define the **nullary characteristic formula** to be the closed formula:

$$\xi_{nullary}^S \stackrel{\text{def}}{=} \bigwedge_{p \in \mathcal{P}_0} p^{\iota^S(p)}() \quad (6)$$

For a predicate p of arity $r \geq 1$, we define the **predicate characteristic formula** to be the closed formula:

$$\begin{aligned} \xi^S[p] \stackrel{\text{def}}{=} \forall w_1, \dots, w_r : & \bigwedge_{\{u'_1, \dots, u'_r\} \in U} \\ & \bigwedge_{j=1}^r \text{node}_{u'_j}^S(w_j) \Rightarrow p^{\iota^S(p)(u'_1, \dots, u'_r)}(w_1, \dots, w_r) \end{aligned} \quad (7)$$

The **characteristic formula of S** is defined by:

$$\begin{aligned} \xi^S \stackrel{\text{def}}{=} & \bigwedge_{i=1}^n (\exists v : \text{node}_{u_i}^S(v)) \\ & \wedge \xi_{total}^S \\ & \wedge \xi_{nullary}^S \\ & \wedge \bigwedge_{r=1}^{maxR} \bigwedge_{p \in \mathcal{P}_r} \xi^S[p] \end{aligned} \quad (8)$$

The **characteristic formula of set $X \subseteq \mathbf{3}\text{-STRUCT}[\mathcal{P}]$** is defined by:

$$\hat{\gamma}(X) = F \wedge \left(\bigvee_{S \in X} \xi^S \right) \quad (9)$$

Finally, for a singleton set $X = \{S\}$ we write $\hat{\gamma}(S)$ instead of $\hat{\gamma}(X)$.

The main ideas behind the four conjuncts of Eq. (8) are:

- The existential quantification in the first conjunct requires that the 2-valued structures have at least n distinct nodes. For each abstract node in S , the first sub-formula locates the corresponding concrete node. Overall, this conjunct guarantees that embedding is surjective.
- The totality formula ensures that every concrete node is represented by some abstract node. It guarantees that the embedding function is well-defined.
- The nullary characteristic formula ensures that the values of nullary predicates in the 2-valued structures are at least as precise as the values of the corresponding nullary predicates in S .
- The predicate characteristic formulas guarantee that predicate values in the 2-valued structures obey the requirements imposed by an embedding into S .⁵

⁵Definition 3.7 relates to all FO-identifiable structures, not only to bounded structures. For bounded structures, it can be simplified by omitting $\xi^S[p]$ for all unary predicates p , because it is implied by ξ_{total}^S . In fact, it can be omitted only for the abstraction predicates, as defined in [Sagiv et al. 2002]; however throughout this paper we consider all unary predicates to be abstraction predicates.

EXAMPLE 3.3. *After a small amount of simplification, the characteristic formula $\widehat{\gamma}(S)$ for the structure S shown in Fig. 2 is $F_{List} \wedge \xi^S$, where ξ^S is:*

$$\begin{aligned} & \exists v : \text{node}_{u_1}^S(v) \wedge \exists v : \text{node}_{u_2}^S(v) \\ & \wedge \forall w : \text{node}_{u_1}^S(w) \vee \text{node}_{u_2}^S(w) \\ & \wedge \bigwedge_{p \in \mathcal{P}_1} \forall w_1 : \bigwedge_{i=1,2} (\text{node}_{u_i}^S(w_1) \Rightarrow p^{t^S(p)(u_i)}(w_1)) \\ & \wedge \forall w_1, w_2 : (\text{node}_{u_1}^S(w_1) \wedge \text{node}_{u_1}^S(w_2) \Rightarrow \text{eq}(w_1, w_2) \wedge \neg n(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ & \quad \wedge (\text{node}_{u_1}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \Rightarrow \neg \text{eq}(w_1, w_2) \wedge \neg n(w_2, w_1)) \end{aligned}$$

The node formulas are given in Example 3.1, and the predicates for the `insert` program in Fig. 1(b) are shown in Table I. Above, we simplified the formula from Eq. (8) by combining implications that had the same premises. The integrity formula F_{List} is given in Example 2.2. Note that it uses transitive closure to define the reachability predicates; consequently, $\widehat{\gamma}(S)$ is a formula in first-order logic with transitive closure.

When a predicate has an indefinite value on some node tuple, a corresponding conjunct of Eq. (7) can be omitted, because it simplifies to **1**.

Thus, the size of this simplified version of ξ^S is linear in the number of definite values of predicates in S . Assuming that the node^S formulas contain no quantifiers or transitive-closure operator, e.g., when S is bounded, the ξ^S formula has no quantifier alternation, and does not contain any occurrences of the transitive-closure operator. Thus, the formula $\widehat{\gamma}$ is in Existential-Universal normal form (and thus decidable for satisfiability) whenever F is in Existential-Universal normal form and does not contain transitive closure.⁶ Moreover, if the maximal arity of the predicate in \mathcal{P} is 2, then $\widehat{\gamma}$ is in the two-variable fragment of first-order logic [Mortimer 1975], wherever F is. In Section 5, we discuss other conditions under which $\widehat{\gamma}$ can be expressed in a decidable logic.

The following theorem shows that for every FO-identifiable structure S , the formula $\widehat{\gamma}(S)$ accepts exactly the set of 2-valued structures represented by S .

THEOREM 3.8. *For every FO-identifiable 3-valued structure S , and 2-valued structure S^\sharp , $S^\sharp \in \gamma(S)$ iff $S^\sharp \models \widehat{\gamma}(S)$.*

4. SUPERVALUATIONAL SEMANTICS FOR FIRST-ORDER FORMULAS

In this section, we consider the problem of how to extract information from a 3-valued structure by evaluating a query. A compositional semantics for 3-valued first-order logic is defined in [Sagiv et al. 2002]; however, that semantics is not as precise as the one defined here. The semantics given in this section can be seen as providing the limit of obtainable precision.

The Notion of Supervaluational Semantics defined below, has been used in [van Fraassen 1966; Bruns and Godefroid 2000].

Definition 4.1. Supervaluational Semantics of First-Order Formulas Let X be a set of 3-valued structures and φ be a closed formula. The **supervaluational semantics of φ in X** , denoted by $\langle\langle \varphi \rangle\rangle(X)$, is defined to be the join of the values of φ obtained from each

⁶For practical reasons, we often replace the *node* formula by a new (definable) predicate, and add its definition to the integrity formula.

```

procedure Supervaluation( $\varphi$ : Formula,
                        X: Set of 3-valued structures): Value
  if ( $\hat{\gamma}(X) \Rightarrow \varphi$  is valid) return 1;
  else if ( $\hat{\gamma}(X) \Rightarrow \neg\varphi$  is valid) return 0;
  otherwise return 1/2;

```

Fig. 5. A procedure for computing the supervaluational value of a formula φ that encodes a query on a 3-valued structures S .

of the 2-valued structures that X represents, i.e., the most-precise conservative value that can be reported for the value of formula φ in the 2-valued structures represented by X is

$$\llbracket \varphi \rrbracket (X) = \begin{cases} 1 & \text{if } S^{\sharp} \models \varphi \text{ for all } S^{\sharp} \in \gamma(X) \\ 0 & \text{if } S^{\sharp} \not\models \varphi \text{ for all } S^{\sharp} \in \gamma(X) \\ 1/2 & \text{otherwise} \end{cases} \quad (10)$$

The compositional semantics given in [Sagiv et al. 2002] and used in TVLA can yield $1/2$ for φ , even when the value of φ is 1 for all the 2-valued structures S^{\sharp} that S represents (or when the value of φ is 0 for all the S^{\sharp}). In contrast, when the supervaluational semantics yields $1/2$, we know that any sound extraction of information from S must return $1/2$.

EXAMPLE 4.1. *We demonstrate now that the supervaluational semantics of the formula $\varphi_{x \rightarrow \text{next} \neq \text{NULL}} \stackrel{\text{def}}{=} \exists v_1, v_2 : x(v_1) \wedge n(v_1, v_2)$ on the structure S from Fig. 2(d) is 1. That is, we wish to argue that for all of the 2-valued structures that structure S from Fig. 2(d) represents, the value of the formula $\varphi_{x \rightarrow \text{next} \neq \text{NULL}}$ must be 1.*

We reason as follows: S represents a list with at least two nodes; i.e., all 2-valued structures represented by S have at least two nodes. One node, u_1^{\sharp} , corresponding to u_1 in S , is pointed to by program variable x . The other node, corresponding to the summary node u_2 , must be reachable from x . Consider the sequence of nodes reachable from x , starting with u_1^{\sharp} . Denote the first node in the sequence that embeds into u_2 by u_2^{\sharp} . By the definition of reachability, there must be an n -link to u_2^{\sharp} from a node embedded into u_1 . But the integrity rules guarantee that there is exactly one node that embeds into u_1 , namely, u_1^{\sharp} . Therefore, the formula $x(v_1) \wedge n(v_1, v_2)$ holds for $[v_1 \mapsto u_1^{\sharp}, v_2 \mapsto u_2^{\sharp}]$.

Note that this formula will be evaluated to $1/2$ by TVLA, because $x(v_1) \wedge n(v_1, v_2)$ evaluates to $1/2$ under the assignment $[v_1 \mapsto u_1, v_2 \mapsto u_2]$: the compositional semantics yields $x(u_1) \wedge n(u_1, u_2) = 1 \wedge 1/2 = 1/2$.

Notice that Definition 4.1 does not provide a constructive way to compute $\llbracket \varphi \rrbracket (X)$ because $\gamma(X)$ is usually an infinite set.

Computing Supervaluational Semantics using Theorem Provers. If an appropriate theorem prover is at hand, $\llbracket \varphi \rrbracket (S)$ can be computed with the procedure shown in Fig. 5. This procedure is not an algorithm, because the theorem prover might not terminate. Termination can be assured by using standard techniques (e.g., having the theorem prover return a safe answer if a time-out threshold is exceeded) at the cost of losing the ability to guarantee that a most-precise result is obtained. If the queries posed by operation Supervaluation can be expressed in a decidable logic, the algorithm for computing supervaluation can be implemented using a decision procedure for that logic. In Section 5, we discuss such decidable logics that are useful for shape analysis.

5. APPLICATIONS

The experiments discussed in this section demonstrate how the $\hat{\gamma}$ operation can be harnessed in the context of program analysis: the results described below go beyond what previous systems were capable of. In Section 5.1, we discuss the use existing theorem provers and their limitations. In Section 5.2, we suggest a way to overcome these limitations, using decidable logic.

We present two examples that use $\hat{\gamma}$ to read out information from 3-valued structures in a conservative, but rather precise way. The first example demonstrates how supervaluational semantics allows us to obtain more precise information from a 3-valued structure than we would have using compositional semantics. The second example demonstrates how to use the 3-valued structures obtained from a TVLA analysis to construct a loop invariant; this is then used to show that certain properties of a linked data structure hold on each loop iteration. In addition, we briefly describe how $\hat{\gamma}$ can be used in algorithms for computing most-precise abstraction operations for shape analysis. Finally, we report on other work that employs $\hat{\gamma}$ to generate a concrete counter-example for shape analysis.

Remark. The $\hat{\gamma}$ operation defines a symbolic concretization with respect to a given abstract domain. In Section 3, we defined $\hat{\gamma}$ for the abstract domain of sets of 3-valued structures. In Appendix A, we describe a related abstract domain and define $\hat{\gamma}$ for it. The applications described in this section can be used with any domain for which $\hat{\gamma}$ is defined in some logic and a theorem prover for that logic exists. In our examples, we use $\hat{\gamma}$ defined in Section 3 and the first-order logic with transitive closure.

5.1 Using the First-Order Theorem Prover SPASS

The TVLA ([Lev-Ami and Sagiv 2000]) system performs an iterative fixed-point computation, which yields at every program point p a set X_p of bounded structures. It guarantees that $\gamma(X_p)$ is a superset of the 2-valued structures that can arise at p in any execution. We have implemented the $\hat{\gamma}$ operation in TVLA, and employed SPASS [Weidenbach] to check, using the formula $\hat{\gamma}(X_p)$, that certain properties of the heap hold at program point p . Also, we implemented the supervaluational procedure described in Section 4, employing SPASS. The enhanced version of TVLA generates the formula $\hat{\gamma}(S)$ and makes at most two calls to SPASS to compute the supervaluational value of a query φ in structure S . In this section, we report on our experience in using SPASS and the problems we have encountered.

First, calls to SPASS theorem prover need not terminate, because first-order logic is undecidable in general. However, in the examples described below, SPASS always terminated.

EXAMPLE 5.1. *In Example 4.1 we (manually) proved that the supervaluational value of the formula $\varphi_{x \rightarrow \text{next} \neq \text{NULL}}$ on the structure S from Fig. 2(d) is 1. To check this automatically, we used SPASS to determine the validity of $\hat{\gamma}(S) \Rightarrow \varphi_{x \rightarrow \text{next} \neq \text{NULL}}$; SPASS indicated that the formula is valid. This guarantees that the formula $\varphi_{x \rightarrow \text{next} \neq \text{NULL}}$ evaluates to 1 on all of the 2-valued structures that embed into S .*

In contrast, TVLA uses Kleene semantics for 3-valued formulas, and will evaluate the formula $\varphi_{x \rightarrow \text{next} \neq \text{NULL}}$ to $1/2$: under the assignment $[v_1 \mapsto u_1, v_2 \mapsto u_2]$, $x(v_1) \wedge n(v_1, v_2)$ evaluates to $1 \wedge 1/2$, which equals $1/2$.

5.1.1 Generating and Querying a Loop Invariant. We used TVLA to compute, for each program point p , a set X_p of bounded structures that overapproximate the set of

stores that may occur at that point. We then generated $\hat{\gamma}(X_p)$. Because TVLA is sound, $\hat{\gamma}(X_p)$ must be an invariant that holds at program point p , according to Theorem 3.8. In particular, when p is a program point that begins a loop, $\hat{\gamma}(X_p)$ is a loop invariant.

EXAMPLE 5.2. Let $X = \{S_i \mid i = 1, \dots, 5\}$ denote the set of five 3-valued structures that TVLA found at the beginning of the loop in the `insert` program from Fig. 2. Table II and Table III of Appendix C show the S_i and their characteristic formulas. The loop invariant is defined by

$$\hat{\gamma}(X) = F_{List} \wedge \left(\bigvee_{i=1}^5 \xi^{S_i} \right)$$

Using SPASS, this formula was then used to check that in every structure that can occur at the beginning of the loop, x points to a valid list, i.e., one that is acyclic and unshared. This property is defined by the following formulas:

$$\begin{aligned} acyc_x &\stackrel{\text{def}}{=} \forall v_1, v_2 : r_x(v_1) \wedge n^+(v_1, v_2) \Rightarrow \neg n^+(v_2, v_1) \\ uns_x &\stackrel{\text{def}}{=} \forall v : r_x(v) \Rightarrow \neg(\exists w_1, w_2 : \neg eq(w_1, w_2) \wedge n(w_1, v) \wedge n(w_2, v)) \\ list_x &\stackrel{\text{def}}{=} acyc_x \wedge uns_x \end{aligned}$$

We applied SPASS to check the validity of $\hat{\gamma}(S) \Rightarrow list_x$; SPASS indicated that the formula is valid.⁷

In addition to the termination issue, a second obstacle is that SPASS considers infinite structures, which are not allowed in our setting.⁸ As a consequence, SPASS can fail to verify that a formula is valid for our intended set of structures; however, the opposite can never happen: whenever SPASS indicates that a formula is valid, it is indeed valid for our intended set of structures.

EXAMPLE 5.3. We tried to verify that every concrete linked-list represented by the 3-valued structure S from Fig. 2(d) has a last element. This condition is expressed by the formula $\varphi_{last} \stackrel{\text{def}}{=} \exists v_1 \forall v_2 : \neg n(v_1, v_2)$. The supervaluational value of φ_{last} on a structure S is $\langle\langle \varphi \rangle\rangle(S) = 1$, for the following reasons. Because r_x has the definite value 1 on u_2 in S , all concrete nodes represented by the summary node u_2 must be reachable from x . Thus, these nodes must form a linked list, i.e., each of these concrete nodes, except for one node that is the “last”, has an n -edge to another concrete node represented by u_2 . The last node does not have an n -edge back to any of the nodes represented by u_2 , because that would create sharing, whereas the value of predicate is in S is 0 on u_2 . Also, the last node cannot have an n -edge to the concrete node represented by u_1 , because the value of predicate n on the pair $\langle u_2, u_1 \rangle$ in S is 0. Therefore, the last element cannot have an outgoing n -edge.

We used SPASS to determine the validity of $\hat{\gamma}(S) \Rightarrow \varphi_{last}$; SPASS indicated that the formula is not valid, because it considered a structure that has infinitely many concrete nodes, all represented by u_2 . Each of these concrete nodes has an n -edge to the next node.

The validity test of the formula $\hat{\gamma}(S) \Rightarrow \neg \varphi_{last}$ failed, of course, because there exists a finite structure that is represented by S (and thus satisfies $\hat{\gamma}(S)$) and has a last element.

⁷SPASS input is available from www.cs.tau.ac.il/~gretay.

⁸Our intended structures are finite, because they represent memory configurations, which are guaranteed to be finite, although their size is not bounded.




Fig. 6. SPASS takes into account structures in which the $t[n]$ predicate overapproximates the n^+ predicate, such as the structure shown in this figure.

For example, the structure in Fig. 2(a) that represents a list of size 2. Therefore, the procedure $\text{Supervaluation}(\varphi_{last}, S)$ implemented using SPASS returns $1/2$, even though the supervaluational value of φ_{last} on S is 1.

The third, and most severe, problem that we face is that SPASS does not support transitive closure. Because transitive closure is not expressible in first-order logic, we could only partially model transitive closure in SPASS, as described below.

SPASS follows other theorem provers in allowing axioms to express requirements on the set of structures considered. We used SPASS axioms to model integrity rules. To partially model transitive closure, we replaced uses of $n^+(v_1, v_2)$ by uses of a new designated predicate $t[n](v_1, v_2)$. Therefore, SPASS will consider some structures that do not represent possible stores. As a consequence, SPASS can fail to verify that a formula is valid for our intended set of structures; however, the opposite can never happen: whenever SPASS indicates that a formula is valid, it is indeed valid for our intended set of structures. To avoid some of the spurious failures to prove validity, we added axioms to guarantee that (i) $t[n](v_1, v_2)$ is transitive and (ii) $t[n](v_1, v_2)$ includes all of $n(v_1, v_2)$; thus, $t[n](v_1, v_2)$ includes all of $n^+(v_1, v_2)$. Because transitive closure requires a minimal set, which is not expressible in first-order logic, this approach provides a looser set of integrity rules than we would like. However, it is still the case that whenever SPASS indicates that a formula is valid, it is indeed valid for the set of structures in which $t[n](v_1, v_2)$ is exactly $n^+(v_1, v_2)$.

EXAMPLE 5.4. *SPASS takes into account the structure shown in Fig. 6, in which the value of $t[n](u_1, u_3)$ is 1, but the value of $n^+(u_1, u_3)$ is 0 because there is no n -edge from u_2 to u_3 .*

Remark. For practical purposes, the success of using symbolic methods depends on having a terminating theorem prover. We have successfully used SPASS as part of a prototype implementation of the *assume* operation (Section 5.3), and the path-pruning optimization for counter-example generation (Section 5.4). Although these experiments are rather preliminary, we believe that this approach can be made to work in practice. For example, there has been some recent progress in using SPASS, including the use of transitive closure [Lev-Ami et al. 2005]. Also, we have investigated a complementary approach, discussed in Section 5.2.

5.2 Decidable Logic

The obstacles mentioned in Section 5.1 are not specific to SPASS. They occur in all theorem provers for first-order logic that we are aware of. To address these obstacles, we are investigating the use of a decidable logic. To reason about linked data structures, we need

a notion of reachability to be expressible, for example, using transitive closure. However, a logic that is both decidable and includes reachability must be limited in other aspects.

One such example is the decidable second-order theory of two successors *WS2S* [Rabin 1969]; its decision procedure is implemented in a tool called MONA [Henriksen et al. 1996]. Second-order quantification suffices to express reachability, but there are still two problems. First, the decision procedure for *WS2S* is necessarily non-elementary [Meyer 1975]. Second, *WS2S* only applies to trees, or, equivalently, to function graphs (graphs with at most one edge leaving any vertex).

Another example is $EA(TC, f^1)$, which is a subset of first-order logic with transitive closure, in which the following restriction are imposed on formulas: (i) they must be in existential-universal form, and (ii) they must use at most a single unary function f , but can use an arbitrary number of unary predicates. [Immerman et al. 2004a] shows that the decision procedure for satisfiability of $EA(TC, f^1)$ is NEXPTIME-complete.

In spite of their limitations, both *WS2S* and $EA(TC, f^1)$ can be useful for reasoning about shape invariants and mutation operations on data structures, such as singly and doubly linked lists, (shared) trees, and graph types [Klarlund and Schwartzbach 1993]. The key is the *simulation technique* [Immerman et al. 2004b], which encodes complex data-structures using *tractable* structures, e.g., function graphs or simple trees, where we can reason with decidable logics.

For example, given a suitable simulation, $\hat{\gamma}$ formula can be expressed in *WS2S* and $EA(TC, f^1)$ if the integrity formula F can. This follows from the definition of $\hat{\gamma}$ in Eq. (9) and the fact that ξ^S does not contain quantifier alternation. This makes $EA(TC, f^1)$ and *WS2S* candidate implementations for the decision procedure used in the supervaluational semantics and in the algorithms described below.

5.3 Assume-Guarantee Shape Analysis

The $\hat{\gamma}$ operation is useful beyond computing supervaluational semantics: it is a necessary operation used in the algorithms described in [Yorsh et al. 2004; Reps et al. 2004]. These algorithms perform abstract operations symbolically by representing abstract values as logical formulas, and use a theorem prover to check validity of these formulas. These algorithms improve on existing shape-analysis techniques by:

- conducting abstract interpretation in the most-precise fashion, improving the technique used in the TVLA system [Lev-Ami and Sagiv 2000; Sagiv et al. 2002], which provides no guarantees about the precision of its basic mechanisms.
- performing modular verification using assume-guarantee reasoning and procedure specifications. This is perhaps the most-exciting potential application of $\hat{\gamma}$ (and $EA(TC, f^1)$ logic), because existing mechanisms for shape analysis, including TVLA, do not support assume-guarantee reasoning.

5.4 Counter-example Generation

Some preliminary work to use the techniques presented in this paper to improve the applicability of TVLA has been carried out. The tool described in [Erez et al. 2003; Erez 2004] uses the $\hat{\gamma}$ operation to generate a concrete counter-example for a potential error message produced by TVLA for an intermediate 3-valued structure S at a program point p . Such a tool is useful to check if a reported error is a real error or a false-alarm, i.e., it never occurs on any concrete store.

Generation of concrete counter-examples from S proceeds as follows. First, S is converted to the formula $\hat{\gamma}(S)$. Then, the tool uses weakest precondition to generate a formula that represents the stores at the entry point that lead to an execution trace that reaches p with a store that satisfies $\hat{\gamma}(S)$. Finally, a separate tool [McCune 2001] generates a concrete store that satisfies the formula for the entry point.

6. RELATED WORK

There is a sizeable literature on *structure-description formalisms* for describing properties of linked data structures (see [Benedikt et al. 1999; Sagiv et al. 2002] for references). The motivation for the present paper was to understand the expressive power of the shape abstractions defined in [Sagiv et al. 2002].

In previous work, Benedikt et al. [Benedikt et al. 1999] showed how to translate two kinds of shape descriptors, “path matrices” [Hendren 1990; Hendren and Nicolau 1990] and the variant of shape graphs discussed in [Sagiv et al. 1998], into a logic called L_r (“logic of reachability expressions”). The shape graphs from [Sagiv et al. 1998] are also amenable to the techniques presented in the present paper: the characteristic formula defined in Eq. (8) is much simpler than the translation to L_r given in [Benedikt et al. 1999]; moreover, Eq. (8) applies to a more general class of shape descriptors. However, the logic used in [Benedikt et al. 1999] is decidable, which guarantees that terminating procedures can be given for problems that can be addressed using L_r .

The Pointer Analysis Logic Engine (PALE) [Møller and Schwartzbach 2001] provides a structure-description formalism that serves as an assertion language; assertions are translated to second-order monadic logic and fed to MONA. PALE does not handle all data structures, but can handle data structures describable as graph types [Klarlund and Schwartzbach 1993]. Because the logic used by MONA is decidable, PALE is guaranteed to terminate.

One point of contrast between the shape abstractions based on 3-valued structures studied in this paper and both L_r and the PALE assertion language is that the powerset of 3-valued structures forms an abstract domain. This means that 3-valued structures can be used for program analysis by setting up an appropriate set of equations and finding its fixed point [Sagiv et al. 2002]. In contrast, when PALE is used for program analysis, an invariant must be supplied for each loop.

Other structure-description formalisms in the literature include ADDS [Hendren et al. 1992] and shape types [Fradet and Metayer 1997].

The supervaluational semantics for first-order logic discussed in Section 4 is related to a number of other supervaluational semantics for partial logics and 3-valued logics discussed in the literature [van Fraassen 1966; Blamey 2002; Bruns and Godefroid 2000]. Compared to previous work, an innovation of Fig. 5 is the use of $\hat{\gamma}$ to translate a 3-valued structure to a formula. In fact, Fig. 5 is an example of a general reductionist strategy for providing a supervaluational evaluation procedure for abstract domains by using existing logics and theorem-provers/decision-procedures.

A recent work [Kuncak and Rinard 2003a], which is an abbreviated version of a more extensive presentation of the results reported in [Kuncak and Rinard 2003b], provides an alternative characterization of 3-valued structures using logical formulas, equivalent to the characterization presented in the present paper. The present paper, which extends and elaborates on the results of [Yorsh 2003], unlike [Kuncak and Rinard 2003a; 2003b], reports on experience and algorithmic issues in using logical characterization of structures for shape

analysis; this material is important because shape analysis is the primary motivation and the intended application of this paper, as well as [Kuncak and Rinard 2003a; 2003b]. Also, Section A.4 of the present paper gives a simple semantic argument for the property of closure under negation, shown in [Kuncak and Rinard 2003b] using a different formalism. The technical similarities and differences between the two works are described in a note available from www.cs.tau.ac.il/~gretay.

7. FINAL REMARKS

In [Reps et al. 2004], we discuss how to perform all operations required for abstract interpretation in the most-precise way possible (relative to the abstraction in use), if certain primitive operations can be carried out, and if a sufficiently powerful theorem prover is at hand. Chief among the primitive operations that must be available is $\hat{\gamma}$; thus, the material that has been presented in this paper shows how to fulfill the requirements of [Reps et al. 2004] for a family of abstractions based on 3-valued structures (essentially those used in our past work [Sagiv et al. 2002] and in the TVLA system [Lev-Ami and Sagiv 2000]).

In ongoing work, we are investigating the feasibility of actually applying the techniques from [Reps et al. 2004] to perform abstract interpretation for abstractions based on 3-valued structures. This approach could be more precise than TVLA because, for instance, it would take into account in a first-class way the integrity formula of the abstraction. In contrast, in TVLA some operations temporarily ignore the integrity formula, and rely on later clean-up steps to rectify matters.

Another step can be taken in this direction, which is to eliminate the use of 3-valued structures, and directly carry out fixed-point computations over logical formulas.

We are also investigating the feasibility of using the results from this paper to develop a more precise and modular version of TVLA by using *assume-guarantee* reasoning [Yorsh et al. 2004]. The idea is to allow arbitrary first-order formulas with transitive closure to be used to express pre- and post-conditions, and to analyze the code for each procedure separately.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: <http://www.acm.org/pubs/citations/journals/tocl/2006-V-N/p1-URL>

ACKNOWLEDGMENTS

We thank Neil Immerman, Viktor Kuncak, Tal Lev-Ami, and Alexander Rabinovich for their contributions to this paper.

REFERENCES

- ANDERSEN, L. O. 1993. Binding-time analysis and the taming of C pointers. In *Proc. of ACM Symposium on Partial Evaluation and Semantics-Based Program Manipulation, PEPM'93*, D. Schmidt, Ed. ACM Press, New York, NY, 47–58.
- BENEDIKT, M., REPS, T., AND SAGIV, M. 1999. A decidable logic for describing linked data structures. In *Proceedings of the 1999 European Symposium On Programming*. 2–19.
- BLAMEY, S. 2002. Partial logic. In *Handbook of Phil. Logic, 2nd. Ed., Vol. 5*, D. Gabbay and F. Guentner, Eds. Kluwer Academic Publishers, 261–353.
- BRUNS, G. AND GODEFROID, P. 2000. Generalized model checking: Reasoning about partial state spaces. In *Proc. CONCUR*. Springer-Verlag, 168–182.

- CHASE, D., WEGMAN, M., AND ZADECK, F. 1990. Analysis of pointers and structures. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* ACM Press, New York, NY, 296–310.
- CLARKE, E., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. 2000. Counterexample-guided abstraction refinement. In *Proc. Computer-Aided Verif.* 154–169.
- CLARKE, E., GRUMBERG, O., AND LONG, D. 1994. Model checking and abstraction. *Trans. on Prog. Lang. and Syst.* 16, 5, 1512–1542.
- COURCELLE, B. 1996. On the expression of graph properties in some fragments of monadic second-order logic. In *Descriptive Complexity and Finite Models: Proceedings of a DIAMCS Workshop*, N. Immerman and P. Kolaitis, Eds. American Mathematical Society, Chapter 2, 33–57.
- COUSOT, P. AND COUSOT, R. 1977. Abstract interpretation: A unified lattice model for static analysis of programs by construction of approximation of fixed points. In *Symp. on Princ. of Prog. Lang.* ACM Press, New York, NY, 238–252.
- DAMS, D. 1996. Abstract interpretation and partial refinement for model checking. Ph.D. thesis, Technical Univ. of Eindhoven, Eindhoven, The Netherlands.
- DAS, M. 2000. Unification-based pointer analysis with directional assignments. In *Conf. on Prog. Lang. Design and Impl.* 35–46.
- EREZ, G. 2004. Generating concrete counter examples for arbitrary abstract domains. M.S. thesis, Tel-Aviv University, Tel-Aviv, Israel. In Preparation.
- EREZ, G., SAGIV, M., AND YAHAV, E. 2003. Generating concrete counter examples for arbitrary abstract domains. Unpublished Manuscript.
- FAGIN, R. 1975. Monadic generalized spectra. *Z. Math. Logik* 21, 89–96.
- FÄHNDRICH, M., FOSTER, J., SU, Z., AND AIKEN, A. 1998. Partial online cycle elimination in inclusion constraint graphs. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* ACM Press, New York, NY, 85–96.
- FRADET, P. AND METAYER, D. L. 1997. Shape types. In *Symp. on Princ. of Prog. Lang.* ACM Press, New York, NY, 27–39.
- GODEFROID, P. AND JAGADEESAN, R. 2003. On the expressiveness of 3-valued models. In *VMCAI*. 206–222.
- HEINTZE, N. AND TARDIEU, O. 2001. Ultra-fast aliasing analysis using CLA: A million lines of C code in a second. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* ACM Press, New York, NY.
- HELL, P. AND NESETRIL, J. 2004. *Graphs and Homomorphisms*. Oxford University Press.
- HENDREN, L. 1990. Parallelizing programs with recursive data structures. Ph.D. thesis, Cornell Univ., Ithaca, NY.
- HENDREN, L., HUMMEL, J., AND NICOLAU, A. 1992. Abstractions for recursive pointer data structures: Improving the analysis and the transformation of imperative programs. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* ACM Press, New York, NY, 249–260.
- HENDREN, L. AND NICOLAU, A. 1990. Parallelizing programs with recursive data structures. *IEEE Trans. on Par. and Dist. Syst.* 1, 1 (January), 35–47.
- HENRIKSEN, J., JENSEN, J., JØRGENSEN, M., KLARLUND, N., PAIGE, B., RAUHE, T., AND SANDHOLM, A. 1996. Mona: Monadic second-order logic in practice. In *Proc. of TACAS 95*. 89–110.
- HUTH, M., JAGADEESAN, R., AND SCHMIDT, D. A. 2001. Modal transition systems: A foundation for three-valued program analysis. In *ESOP*. 155–169.
- IMMERMAN, N. 1999. *Descriptive Complexity*. Springer-Verlag.
- IMMERMAN, N., RABINOVICH, A., REPS, T., SAGIV, M., AND YORSH, G. 2004a. The boundary between decidability and undecidability for transitive-closure logics. In *CSL*.
- IMMERMAN, N., RABINOVICH, A., REPS, T., SAGIV, M., AND YORSH, G. 2004b. Verification via structure simulation. In *CAV*.
- JONES, N. AND MUCHNICK, S. 1981. Flow analysis and optimization of Lisp-like structures. In *Program Flow Analysis: Theory and Applications*, S. Muchnick and N. Jones, Eds. Prentice-Hall, Englewood Cliffs, NJ, Chapter 4, 102–131.
- JONES, N. AND MUCHNICK, S. 1982. A flexible approach to interprocedural data flow analysis and programs with recursive data structures. In *Symp. on Princ. of Prog. Lang.* ACM Press, New York, NY, 66–74.
- KLARLUND, N. AND SCHWARTZBACH, M. 1993. Graph types. In *Symp. on Princ. of Prog. Lang.* ACM Press, New York, NY, 196–205.
- KUNCAK, V., LAM, P., AND RINARD, M. C. 2002. Role analysis. In *POPL*. 17–32.

- KUNCAK, V. AND RINARD, M. 2003a. Boolean algebra of shape analysis constraints. In *VMCAI*. 59–72.
- KUNCAK, V. AND RINARD, M. 2003b. On Boolean algebra of shape analysis constraints. Tech. rep., MIT, CSAIL. Available at “<http://www.mit.edu/~vkuncak/papers/index.html>”.
- LAM, P., KUNCAK, V., AND RINARD, M. 2005. Hob: A tool for verifying data structure consistency. In *Conf. on Compiler Construction (tool demo)*.
- LEV-AMI, T. 2000. TVLA: A framework for Kleene based static analysis. M.S. thesis, Tel-Aviv University, Tel-Aviv, Israel.
- LEV-AMI, T., IMMERMANN, N., REPS, T., SAGIV, M., SRIVASTAVA, S., AND YORSH, G. 2005. Simulating reachability using first-order logic with applications to verification of linked data structures. Submitted for publication.
- LEV-AMI, T., REPS, T., SAGIV, M., AND WILHELM, R. 2000. Putting static analysis to work for verification: A case study. In *Proc. of the Int. Symp. on Software Testing and Analysis*. 26–38.
- LEV-AMI, T. AND SAGIV, M. 2000. TVLA: A system for implementing static analyses. In *Static Analysis Symp.* 280–301.
- MCCUNE, W. 2001. Mace 2.0 reference manual and guide. Available at “<http://www-unix.mcs.anl.gov/AR/mace/>”.
- MCMILLAN, K. L. 1999. Verification of infinite state systems by compositional model checking. In *CHARME*. 219–234.
- MEYER, A. R. 1975. Weak monadic second-order theory of successor is not elementary recursive. In *Logic Colloquium, (Proc. Symposium on Logic, Boston, 1972)*. 132–154.
- MØLLER, A. AND SCHWARTZBACH, M. 2001. The pointer assertion logic engine. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* 221–231.
- MORTIMER, M. 1975. On languages with two variables. *Zeitschr. f. math. Logik u. Grundlagen d. Math* 21, 135–140.
- NIELSON, F., NIELSON, H., AND HANKIN, C. 1999. *Principles of Program Analysis*. Springer-Verlag.
- NIELSON, F., NIELSON, H., AND SAGIV, M. 2000. A Kleene analysis of mobile ambients. In *Proc. of ESOP 2000*, G. Smolka, Ed. LNCS, vol. 1782. Springer, 305–319.
- RABIN, M. 1969. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.* 141, 1–35.
- RAMALINGAM, G., VARSHAVSKY, A., FIELD, J., GOYAL, D., AND SAGIV, M. 2002. Deriving specialized program analyses for certifying component-client conformance. In *PLDI*. 83–94.
- REPS, T., SAGIV, M., AND YORSH, G. 2004. Symbolic implementation of the best transformer. In *VMCAI*. 252–266.
- SAGIV, M., REPS, T., AND WILHELM, R. 1998. Solving shape-analysis problems in languages with destructive updating. *Trans. on Prog. Lang. and Syst.* 20, 1 (Jan.), 1–50.
- SAGIV, M., REPS, T., AND WILHELM, R. 1999. Parametric shape analysis via 3-valued logic. In *Symp. on Princ. of Prog. Lang.* ACM Press, New York, NY, 105–118.
- SAGIV, M., REPS, T., AND WILHELM, R. 2002. Parametric shape analysis via 3-valued logic. *Trans. on Prog. Lang. and Syst.*
- SHAHAM, R., YAHAV, E., KOLODNER, E., AND SAGIV, M. 2003. Establishing local temporal heap safety properties with applications to compile-time memory management. In *Proc. of Static Analysis Symposium (SAS’03)*. LNCS, vol. 2694. Springer, 483–503.
- SHAPIRO, M. AND HORWITZ, S. 1997. Fast and accurate flow-insensitive points-to analysis. In *Symp. on Princ. of Prog. Lang.* 1–14.
- STEENSGAARD, B. 1996. Points-to analysis in almost-linear time. In *Symp. on Princ. of Prog. Lang.* 32–41.
- SU, Z., FÄHNDRICH, M., AND AIKEN, A. 2000. Projection merging: Reducing redundancies in inclusion constraint graphs. In *Symp. on Princ. of Prog. Lang.*, T. Reps, Ed. ACM Press, New York, NY, 81–95.
- VAN FRAASSEN, B. 1966. Singular terms, truth-value gaps, and free logic. *J. Phil* 63, 17, 481–495.
- WEIDENBACH, C. SPASS: An automated theorem prover for first-order logic with equality. Available at “<http://spass.mpi-sb.mpg.de/index.html>”.
- YAHAV, E. 2001. Verifying safety properties of concurrent Java programs using 3-valued logic. *Symp. on Princ. of Prog. Lang.* 36, 3, 27–40.

- YAHAV, E. AND RAMALINGAM, G. 2004. Verifying safety properties using separation and heterogeneous abstractions. In *Proceedings of the ACM SIGPLAN 2004 conference on Programming language design and implementation*. ACM Press, 25–34.
- YAHAV, E. AND SAGIV, M. 2003. Automatically verifying concurrent queue algorithms. In *Electronic Notes in Theoretical Computer Science*, B. Cook, S. Stoller, and W. Visser, Eds. Vol. 89. Elsevier.
- YORSH, G. 2003. Logical characterizations of heap abstractions. M.S. thesis, Tel-Aviv University, Tel-Aviv, Israel. Available at “<http://www.math.tau.ac.il/~gretay>”.
- YORSH, G., REPS, T. W., AND SAGIV, M. 2004. Symbolically computing most-precise abstract operations for shape analysis. In *TACAS*. 530–545.

Received April 2004; revised March 2005; accepted April 2005

THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Logical Characterizations of Heap Abstractions

GRETA YORSH

School of Comp. Sci., Tel-Aviv University

THOMAS REPS

Comp. Sci. Dept., University of Wisconsin

MOOLY SAGIV

School of Comp. Sci., Tel-Aviv University

and

REINHARD WILHELM

Informatik, Univ. des Saarlandes

ACM Transactions on Computational Logic, Vol. V, No. N, July 2006, Pages 1–24.

A. CHARACTERIZING CANONICAL ABSTRACTION BY FIRST-ORDER FORMULAS

This section defines an alternative abstract domain for use in shape analysis (and other logic-based analyses). This domain keeps more explicit information than the one in Section 2.4 and enjoys nice closure properties (see Section A.4). This domain uses a particular class of embedding functions that are defined by a simple operation, called *canonical abstraction*, which maps 2-valued structures into a limited subset of bounded structures.

A.1 Canonical Abstraction

Canonical abstraction was defined in [Sagiv et al. 1999] as an abstraction with the following properties:

- It provides a uniform way to obtain 3-valued structures of a priori bounded size. This is important to automatically derive properties of programs with loops by employing iterative fixed-point algorithms. Canonical abstraction maps concrete nodes into abstract nodes according to the definite values of the unary predicates.
- The information loss is minimized when multiple nodes of S are mapped to the same node in S' ,

This is formalized by the following definition:

Definition A.1. A structure $S' = \langle U^{S'}, \iota^{S'} \rangle$ is a **canonical abstraction** of a structure S , if $S \sqsubseteq^{canonical} S'$, where $canonical: U^S \rightarrow U^{S'}$ is the following surjective mapping:

$$canonical(u) = u_{\{p \in \mathcal{P}_1 | \iota^S(p)(u)=1\}, \{p \in \mathcal{P}_1 | \iota^S(p)(u)=0\}} \quad (11)$$

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2006 ACM 1529-3785/06/0700-0001 \$5.00

and, for every $p \in \mathcal{P}_k$ of arity k ,

$$\iota^{S'}(p)(u'_1, \dots, u'_k) = \bigsqcup_{\substack{u_i \in U^S, \text{ s.t.} \\ \text{canonical}(u_i) = u'_i \in U^{S'}, \\ 1 \leq i \leq k}} \iota^S(p)(u_1, \dots, u_k) \quad (12)$$

We say that $S' = \text{canonical}(S)$.

The name “ $u_{\{p \in \mathcal{P}_1 | \iota^S(p)(u)=1\}, \{p \in \mathcal{P}_1 | \iota^S(p)(u)=0\}}$ ” is known as the **canonical name** of node u . The subscript on the canonical name of u involves two sets of unary predicate symbols: (i) those that are true at u , and (ii) those that are false at u .

EXAMPLE A.1. In structure S from Fig. 2, the canonical names of the nodes are as follows:

Node	Canonical Name
u_1	$u_{\{x, r_x\}, \{y, t, e, is, r_y, r_t, r_e\}}$
u_2	$u_{\{r_x\}, \{x, y, t, e, is, r_y, r_t, r_e\}}$

In the context of canonical abstraction, S shown in Fig. 2 represents S_b and S_c , but not S_a ; i.e., S represents lists that are pointed to by x that have at least three nodes, but it does not represent a list with just two nodes. The reason is that predicates n and eq have indefinite values in S , but a list with only two nodes cannot have both 0 and 1 values for the corresponding entries, as required for minimizing information loss as defined in Eq. (12).⁹ In contrast, according to the abstraction that relies on embedding, defined in Section 2.4, S represents lists with two or more elements.

To characterize canonical abstraction, we define the set of 3-valued structures that are “images of canonical abstraction” (ICA), i.e., the results of applying canonical abstraction to 2-valued structures.

Definition A.2. **Image of canonical abstraction (ICA)** Structure S is an ICA if there exists a 2-valued structure S^\natural such that S is the canonical abstraction of S^\natural .

Concretization of 3-Valued Structures. Canonical abstraction allows us to define the (potentially infinite) set of 2-valued structures represented by a set of 3-valued structures, that are ICA

Definition A.3. **Concretization of ICA Structures** For a set of structures $X \subseteq 3\text{-STRUCT}[\mathcal{P}]$, that are ICA structures, we denote by $\gamma_c(X)$ the set of 2-valued structures that X represents, i.e.,

$$\gamma_c(X) = \left\{ S^\natural \in 2\text{-STRUCT}[\mathcal{P}] \mid \text{exists } S \in X \text{ such that } \begin{array}{l} S \text{ is the canonical abstraction of } S^\natural \\ \text{and } S^\natural \models F \end{array} \right\} \quad (13)$$

Also, for a singleton set $X = \{S\}$ we write $\gamma_c(S)$ instead of $\gamma_c(X)$.

The abstract domain is the powerset of ICA structures, where the order relation is set inclusion. Note that this abstract domain is finite, because there is a finite number of different ICA structures (up to isomorphism). Denote by α_c the extension of the abstraction function *canonical* to sets. This defines a Galois connection $\langle \alpha_c, \gamma_c \rangle$ between sets of 2-valued structures and sets of ICA structures.

⁹Eq. (12) is called the *tight-embedding* condition in [Sagiv et al. 2002].

A.2 Canonical-FO-Identifiable Structures

We define the notion of canonical-FO-identifiable nodes using canonical abstraction rather than embedding, which was used for the notion of FO-identifiable nodes in Definition 3.1.

Definition A.4. We say that a node u in a 3-valued structure S is **canonical-FO-identifiable** if there exists a formula $\text{node}_u^S(w)$ with designated free variable w , such that for every 2-valued structure S^\natural , if S is the canonical abstraction of S^\natural , i.e., $S^\natural \in \gamma_c(S)$, then for every concrete node $u^\natural \in U^{S^\natural}$:

$$\text{canonical}(u^\natural) = u \iff S^\natural, [w \mapsto u^\natural] \models \text{node}_u^S(w) \quad (14)$$

S is called canonical-FO-identifiable if all the nodes in S are canonical-FO-identifiable.

We can also prove Lemma 3.2 for the case of canonical abstraction rather than embedding.

A.3 Characterizing Canonical Abstraction

An ICA structure is always a bounded structure, in which all nullary and unary predicates have definite values.¹⁰ This is formalized by the following lemma:

LEMMA A.5. *If 3-valued structure $S = \langle U^S, \iota^S \rangle$ over vocabulary \mathcal{P} is ICA then:*

- (i). S is a bounded structure.
- (ii). For each nullary predicate p , $\iota^S(p)() \in \{0, 1\}$.
- (iii). For each element $u \in U$ and each unary predicate p , $\iota^S(p)(u) \in \{0, 1\}$.

The following lemma shows that ICA structures are canonical-FO-identifiable:

LEMMA A.6. *Every 3-valued structure S that is an ICA is canonical-FO-identifiable, where*

$$\text{node}_{u_i}^S(w) \stackrel{\text{def}}{=} \bigwedge_{p \in \mathcal{P}_1} p^{\iota^S(p)(u_i)}(w) \quad (15)$$

Using this fact, we can define a formula τ^S that accepts exactly the set of 2-valued structures represented by S under canonical abstraction. The formula τ^S is merely ξ^S with additional conjuncts to ensure that the information loss is minimized, i.e., for every predicate p and every 1/2 entry of p , the 2-valued structure has both a corresponding 1 entry and a corresponding 0 entry.

Definition A.7. First-Order Characteristic Formula for Canonical Abstraction Let 3-valued structure $S = \langle U^S, \iota \rangle$ be an ICA.

For a predicate p of arity r , we define the closed formula for p :

$$\tau^S[p] \stackrel{\text{def}}{=} \bigwedge_{\substack{\{u'_1, \dots, u'_r\} \subseteq U^S \\ \text{s.t. } \iota^S(p)(u'_1, \dots, u'_r) = 1/2}} \left(\begin{array}{l} \exists w_1, \dots, w_r : \bigwedge_{j=1}^r \text{node}_{u'_j}^S(w_j) \wedge p(w_1, \dots, w_r) \\ \wedge \exists w_1, \dots, w_r : \bigwedge_{j=1}^r \text{node}_{u'_j}^S(w_j) \wedge \neg p(w_1, \dots, w_r) \end{array} \right) \quad (16)$$

¹⁰If not all unary predicates are defined as abstraction predicates, then the result may be a bounded structure of the less restrictive kind mentioned in Section 3.1. Also, unary predicates that are not abstraction predicates may have indefinite values.

The formula of S is defined by:

$$\tau^S \stackrel{\text{def}}{=} \xi^S \wedge \bigwedge_{r=2}^{\max R} \bigwedge_{p \in \mathcal{P}_r} \tau^S[p] \quad (17)$$

The **characteristic formula for canonical abstraction of a set of ICA structures** $X \subseteq 3\text{-STRUCT}[\mathcal{P}]$ is defined by

$$\widehat{\gamma}_c(X) = F \wedge \left(\bigvee_{S \in X} \tau^S \right) \quad (18)$$

Also, for a singleton set $X = \{S\}$, where S is an ICA structure, we write $\widehat{\gamma}_c(S)$ instead of $\widehat{\gamma}_c(X)$.

EXAMPLE A.2. *The characteristic formula for canonical abstraction of the structure S shown in Fig. 2(d) is:*

$$\begin{aligned} \widehat{\gamma}_c(S) = & \widehat{\gamma}(S) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_1}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge n(w_1, w_2) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_1}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge \neg n(w_1, w_2) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_2}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge n(w_1, w_2) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_2}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge \neg n(w_1, w_2) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_2}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge \text{eq}(w_1, w_2) \\ & \wedge \exists w_1, w_2 : \text{node}_{u_2}^S(w_1) \wedge \text{node}_{u_2}^S(w_2) \wedge \neg \text{eq}(w_1, w_2) \end{aligned} \quad (19)$$

where $\widehat{\gamma}(S)$ is given in Example 3.3. As explained in Example A.1, S does not represent a list of two nodes; the corresponding 2-valued structure S_a , shown in Fig. 2(a), does not satisfy Eq. (19), because the last four lines cannot be satisfied by any assignment in S_a .

Remark. The formula τ^S does not contain quantifier alternation and transitive closure. Therefore, $\widehat{\gamma}_c$ is in Existential-Universal normal form (and thus decidable) whenever F is in Existential-Universal form and does not contain transitive closure.

THEOREM A.8. *For every 3-valued structure S that is an ICA and 2-valued structure S^\natural*

$$S^\natural \in \gamma_c(S) \text{ iff } S^\natural \models \widehat{\gamma}_c(S)$$

A.4 Closure Properties of ICA Structures

This section gives a simple semantic proof that the class of formulas that characterize ICA structures is closed under negation. This result was shown in [Kuncak and Rinard 2003b] using a different formalism.

From Eq. (12) it follows that for two distinct ICA structures S_1 and S_2 , $\gamma_c(S_1) \cap \gamma_c(S_2) = \emptyset$. Intuitively, each 2-valued structure can be represented by exactly one ICA structure. This implies that the complement of the concretization of an ICA structure can be represented precisely by a finite set of ICA structures.

Denote by \mathcal{D} the set of all 2-valued structures that satisfy the integrity formula F : $\mathcal{D} \stackrel{\text{def}}{=} \{S^\natural \in 2\text{-STRUCT}[\mathcal{P}] \mid S^\natural \models F\}$.

LEMMA A.9. *Let S be an ICA structure. There exists a set of ICA structures X such that $\gamma_c(X) = \mathcal{D} \setminus \gamma_c(S)$.*

This can be reformulated using Theorem A.8 in terms of characteristic formulas for ICA structures. This shows that the class of formulas that characterize canonical abstraction is closed under negation, in the following sense:

LEMMA A.10. *Consider the formula τ^S from Eq. (17), for some ICA structure S . There exists a set of ICA structures X , such that the formula $F \wedge \neg\tau^S$ is equivalent to the formula $\hat{\gamma}_c(X)$.*

Remark. Note that Lemma A.9 and Lemma A.10 do not hold for bounded structures using γ , described in Section 3.1, instead of γ_c . The reason, intuitively, is that some 2-valued structures can be represented by more than one bounded structure.

For example, consider the 2-valued structure S_a from Fig. 2, which denotes a linked-list of length exactly 2. It is in the concretization of two different 3-valued structures: the first is S_a itself, considered as a 3-valued structure S' (that represents a single 2-valued structure: $\gamma(S') = \{S_a\}$); the second is the structure S from Fig. 2.

For the purpose of this example, assume that the integrity formula F (that defines \mathcal{D}) requires that all elements be reachable from x , in addition to the integrity formula F_{List} from Example 2.2. The complement $\mathcal{C} \stackrel{\text{def}}{=} \mathcal{D} \setminus \gamma(S') = \mathcal{D} \setminus S_a$ is the set that contains an empty linked list, a linked list of length 1, and linked lists of length 3 or more. The representation of \mathcal{C} is a set X of bounded structures. To capture linked lists of length 3 or more, X must contain a 3-valued structure S from Fig. 2. However, $\gamma(S)$ includes a list of length 2 as well, denoted by S_a , which is not in \mathcal{C} . Therefore, X cannot contain S , and a contradiction is obtained.

B. CHARACTERIZING GENERAL 3-VALUED STRUCTURES BY NP FORMULAS

In this section, we show how to characterize general 3-valued structures.

B.1 Motivating Example

If the input structure is FO-identifiable, Theorem 3.8 ensures that the result of operation $\hat{\gamma}$ precisely captures the concretization of the input structure. The purpose of this example is to show what happens if we apply the $\hat{\gamma}$ operation, as defined in Section 3, to a structure that is not FO-Identifiable. When S is not FO-identifiable, $\hat{\gamma}(S)$ only provides a sufficient test for the embedding of 2-valued structures into S .

EXAMPLE B.1. *The 3-valued structure S shown in Fig. 3 describes undirected graphs. We draw undirected edges as two-way directed edges. This structure uses a set of predicates $\mathcal{P} = \{eq, f, b\}$, where $f(v_1, v_2)$ and $b(v_2, v_1)$ denote the forward and backward directions of an edge between nodes v_1 and v_2 .*

When Eq. (8) is applied to the 3-valued structure S shown in Fig. 3, we get

$$\begin{aligned}
 & \bigwedge_{i=1}^3 \exists v : \text{node}_{u_i}^S(v) \\
 & \wedge \forall w : \bigvee_{i=1}^3 \text{node}_{u_i}^S(w) \\
 & \wedge \forall w_1, w_2 : \bigwedge_{k \neq j} (\text{node}_{u_k}^S(w_1) \wedge \text{node}_{u_j}^S(w_2) \Rightarrow f^{1/2}(w_1, w_2)) \\
 & \wedge \forall w_1, w_2 : \bigwedge_{k \neq j} (\text{node}_{u_k}^S(w_1) \wedge \text{node}_{u_j}^S(w_2) \Rightarrow b^{1/2}(w_1, w_2)) \\
 & \wedge \forall w_1, w_2 : \bigwedge_{i=1}^3 (\text{node}_{u_i}^S(w_1) \wedge \text{node}_{u_i}^S(w_2) \Rightarrow b^0(w_1, w_2)) \\
 & \wedge \forall w_1, w_2 : \bigwedge_{i=1}^3 (\text{node}_{u_i}^S(w_1) \wedge \text{node}_{u_i}^S(w_2) \Rightarrow f^0(w_1, w_2))
 \end{aligned} \tag{20}$$

Because this example does not include unary predicates, the node formula given in Lemma 3.6 evaluates to **1** on all elements. Hence, Eq. (20) can be simplified to:

$$\begin{aligned}
& \bigwedge_{i=1}^3 \exists v : \mathbf{1} \\
& \wedge \quad \forall w : \bigvee_{i=1}^3 \mathbf{1} \\
& \wedge \quad \forall w_1, w_2 : \bigwedge_{k \neq j} (\mathbf{1} \wedge \mathbf{1} \Rightarrow \mathbf{1}) \\
& \wedge \quad \forall w_1, w_2 : \bigwedge_{k \neq j} (\mathbf{1} \wedge \mathbf{1} \Rightarrow \mathbf{1}) \\
& \wedge \quad \forall w_1, w_2 : \bigwedge_{i=1}^3 (\mathbf{1} \wedge \mathbf{1} \Rightarrow \neg b(w_1, w_2)) \\
& \wedge \quad \forall w_1, w_2 : \bigwedge_{i=1}^3 (\mathbf{1} \wedge \mathbf{1} \Rightarrow \neg f(w_1, w_2))
\end{aligned}$$

After further simplification, we get the formula $\forall w_1, w_2 : \neg f(w_1, w_2) \wedge \forall w_1, w_2 : \neg b(w_1, w_2)$. The simplification is due to the fact that the implication in Eq. (7) unconditionally holds for all pairs of distinct nodes, because f and b evaluate to $1/2$ on those pairs, except for the requirement imposed by the absence of self-loops in S .

This formula is only fulfilled by graphs with no edges, which are obviously 3-colorable. But this formula is too restrictive: it does not capture some 3-colorable graphs.

B.2 Characterizing General 3-Valued Structures

Existential monadic second-order formulas are a subset of Fagin's second-order formulas [Fagin 1975], named NP formulas, which capture NP computations. A formula in existential monadic second-order logic has the form:

$$\exists V_1, V_2, \dots, V_n : \varphi$$

where the V_i are set variables, and φ is a first-order formula that can use membership tests in V_i . We show that in this subset of second-order logic, the characteristic formula from Definition 3.7 can be generalized to handle arbitrary 3-valued structures using existential quantification over set variables (with one set variable for each abstract node).¹¹

Definition B.1. NP Characteristic Formula Let $S = \langle U = \{u_1, u_2, \dots, u_n\}, \iota \rangle$ be a 3-valued structure.

We define the following formula to ensure that the sets are non_empty:

$$\xi_{non_empty}^S[i] \stackrel{\text{def}}{=} \exists w_i : \text{node}_{u_i}^S(w_i) \quad (21)$$

We define the following formula to ensure that the sets V_k, V_j are disjoint:

$$\xi_{disjoint}^S[k, j] \stackrel{\text{def}}{=} \forall w_1, w_2 : \text{node}_{u_k}^S(w_1) \wedge \text{node}_{u_j}^S(w_2) \Rightarrow \neg eq(w_1, w_2) \quad (22)$$

The **NP characteristic formula of S** is defined by:

$$\begin{aligned}
\xi^S \stackrel{\text{def}}{=} & \exists V_1, \dots, V_n : \bigwedge_{i=1}^n \xi_{non_empty}^S[i] \wedge \bigwedge_{k \neq j} \xi_{disjoint}^S[k, j] \\
& \wedge \xi_{total}^S \\
& \wedge \xi_{nullary}^S \\
& \wedge \bigwedge_{r=1}^{maxR} \bigwedge_{p \in \mathcal{P}_r} \xi^S[p]
\end{aligned} \quad (23)$$

where $\xi_{total}^S, \xi_{nullary}^S, \xi^S[p]$ are defined as in Definition 3.7, except that $\text{node}_{u_i}^S$ is the NP formula $\text{node}_{u_i}^S(w) \stackrel{\text{def}}{=} (w \in V_i)$. (Here, we abuse notation slightly by referring to V_i in

¹¹This result is mostly theoretical. In principle, this encoding falls into monadic-second order logic, which is decidable if we restrict the concrete structures of interest to trees. However, we have not investigated this direction further.

node $_{u_i}^S(w)$. This could have been formalized by passing V_1, \dots, V_n as extra parameters to node $_{u_i}^S(\cdot)$.

The **NP characteristic formula of a finite set** $X \subseteq \mathbf{3}\text{-STRUCT}[\mathcal{P}]$ is defined by:

$$\widehat{\gamma}_{NP}(X) = F \wedge \left(\bigvee_{S \in X} \xi^S \right) \quad (24)$$

Finally, for a singleton set $X = \{S\}$ we write $\widehat{\gamma}_{NP}(S)$ instead of $\widehat{\gamma}_{NP}(X)$.

EXAMPLE B.2. *After a small amount of simplification, the NP characteristic formula ξ^S for the graph shown in Fig. 3 is:*

$$\begin{aligned} \exists V_1, V_2, V_3 : \bigwedge_{i=1}^3 (\exists w : w \in V_i) & \quad (i) \\ \wedge \bigwedge_{k \neq j} \forall w_1, w_2 : (w_1 \in V_k \wedge w_2 \in V_j \Rightarrow \neg eq(w_1, w_2)) & \quad (ii) \\ \wedge \forall w : \bigvee_{i=1}^3 w \in V_i & \quad (iii) \\ \wedge \forall w_1, w_2 : \bigwedge_{i=1}^3 (\bigwedge_{j=1,2} w_j \in V_i \Rightarrow \neg e(w_1, w_2) \wedge \neg e(w_2, w_1)) & \quad (iv) \end{aligned}$$

In this formula, V_1, V_2 , and V_3 represent the three color classes. Line by line, the formula says: (i) each color class has at least one member; (ii) the color classes are pairwise disjoint; (iii) every node is in a color class; (iv) nodes in the same color class are not connected by an undirected edge.

The following theorem generalizes the result in Theorem 3.8 for an arbitrary 3-valued structure S , using NP-formula $\widehat{\gamma}_{NP}(S)$ to accept exactly the set of 2-valued structures represented by S .

THEOREM B.2. *For every 3-valued structure S , and 2-valued structure S^\natural :*

$$S^\natural \in \gamma(S) \text{ iff } S^\natural \models \widehat{\gamma}_{NP}(S)$$

C. GENERATING AND QUERYING A LOOP INVARIANT

Table II and Table III show the structures and the characteristic formulas for the experiment described in Example 5.2.

It is interesting to note that the size of ξ^{S_2} is bigger than the size of ξ^{S_1} . This is natural because S_2 has more definite values, which impose more restrictions than are imposed by S_1 .

D. PROOFS

LEMMA D.1. *Consider the 3-valued structure S shown in Fig. 3. For all 2-valued structures C , C can be embedded into S if and only if C can be colored using 3 colors.*

Proof of the if direction: Suppose that C is 3-colorable, let c be a mapping from the nodes of C to the colors $\{1, 2, 3\}$. We define embedding function f from C to S as follows: $f(u) = u_{c(u)}$, i.e., a node $u \in C$ that has color i is mapped to $u_i \in S$. It is easy to see that f preserves predicate values in S , because the only definite values in S indicate the absence of self-loops. It is preserved, because there are no edges in C with both endpoints in the same color.

Proof of the only-if direction: Suppose that C is embedded into S using f . We show that C is 3-colorable. For each node $u \in C$, let the color of u , $c(u)$, be the name of the corresponding node in S , i.e., $c(u) = f(u)$. The absence of self loops on any of the three summary nodes guarantees that a pair of adjacent nodes in C cannot be mapped by f to

Structure	Characteristic Formula
<p>S_1</p>	$\begin{aligned} \text{node}_{u_1}^{S_1}(w) &= x(w) \wedge y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_2}^{S_1}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \xi^{S_1} &= \bigwedge_{i=1,2} (\exists v : \text{node}_{u_i}^{S_1}(v)) \\ &\quad \wedge \forall w : \bigvee_{i=1,2} \text{node}_{u_i}^{S_1}(w) \\ &\quad \wedge \forall w_1, w_2 : \bigwedge_{i=1,2} \text{node}_{u_i}^{S_1}(w_i) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \\ &\quad \wedge \forall w_1, w_2 : \bigwedge_{i=1,2} \text{node}_{u_i}^{S_1}(w_i) \Rightarrow \\ &\quad \quad \wedge eq(w_1, w_2) \wedge \neg n(w_1, w_2) \end{aligned}$
<p>S_2</p>	$\begin{aligned} \text{node}_{u_1}^{S_2}(w) &= x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_2}^{S_2}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \xi^{S_2} &= \bigwedge_{i=1,2} (\exists v : \text{node}_{u_i}^{S_2}(v)) \\ &\quad \wedge \forall w : \bigvee_{i=1,2} \text{node}_{u_i}^{S_2}(w) \\ &\quad \wedge \forall w_1, w_2 : \bigwedge_{i=1,2} \text{node}_{u_i}^{S_2}(w_i) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge n(w_1, w_2) \\ &\quad \wedge \forall w_1, w_2 : \bigwedge_{i=1,2} \text{node}_{u_i}^{S_2}(w_i) \Rightarrow \\ &\quad \quad \wedge eq(w_1, w_2) \wedge \neg n(w_1, w_2) \\ &\quad \wedge \forall w_1, w_2 : \bigwedge_{i=1,2} \text{node}_{u_i}^{S_2}(w_i) \Rightarrow \\ &\quad \quad \wedge eq(w_1, w_2) \wedge \neg n(w_1, w_2) \end{aligned}$
<p>S_3</p>	$\begin{aligned} \text{node}_{u_1}^{S_3}(w) &= x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_2}^{S_3}(w) &= \neg x(w) \wedge y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_3}^{S_3}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\quad \wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \xi^{S_3} &= \bigwedge_{i=1,2,3} (\exists v : \text{node}_{u_i}^{S_3}(v)) \\ &\quad \wedge \forall w : \bigvee_{i=1,2,3} \text{node}_{u_i}^{S_3}(w) \\ &\quad \wedge \forall w_1, w_2 : (\bigwedge_{i=1,2} \text{node}_{u_i}^{S_3}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\bigwedge_{i=1,2} \text{node}_{u_i}^{S_3}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_1}^{S_3}(w_1) \wedge \text{node}_{u_2}^{S_3}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_2}^{S_3}(w_1) \wedge \text{node}_{u_3}^{S_3}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_1}^{S_3}(w_1) \wedge \text{node}_{u_3}^{S_3}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge \neg n(w_1, w_2)) \end{aligned}$

Table II. (Continued in Table III.) The left column shows the structures that arise at the beginning of the loop in the `insert` program from Fig. 1(b). The right column shows the characteristic formula for each structure. Note that we omit the redundant sub-formulas $\xi^S[p]$, for $p \in \mathcal{P}_1$, that are part of ξ_{total}^S and $\text{node}_{u_i}^{S_i}(w)$ definitions.

the same summary node. That is, for any edge in C the endpoints must be mapped by f to different summary nodes, thus they have different colors.

Lemma 3.2 *Let S be an FO-identifiable structure and let $u_1, u_2 \in S$ be distinct individuals. Let S^{\natural} be a 2-valued structure that embeds into S and let $u^{\natural} \in S^{\natural}$. At most one of the following can hold, but not both:*

- (I) $S^{\natural}, [w \mapsto u^{\natural}] \models \text{node}_{u_1}^S(w)$

Structure	CharacteristicFormula
<p style="text-align: center;">S_4</p>	$\begin{aligned} \text{node}_{u_1}^{S_4}(w) &= x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_2}^{S_4}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_3}^{S_4}(w) &= \neg x(w) \wedge y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_4}^{S_4}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \end{aligned}$ $\begin{aligned} \xi^{S_4} &= \bigwedge_{i=1,\dots,4} (\exists v : \text{node}_{u_i}^{S_4}(v)) \\ &\wedge \forall w : \bigvee_{i=1,\dots,4} \text{node}_{u_i}^{S_4}(w) \\ &\wedge \forall w_1, w_2 : \\ &\quad (\bigwedge_{i=1,2} \text{node}_{u_i}^{S_4}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\bigwedge_{i=1,2} \text{node}_{u_3}^{S_4}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_1}^{S_4}(w_1) \wedge \text{node}_{u_2}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_2}^{S_4}(w_1) \wedge \text{node}_{u_3}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_1}^{S_4}(w_1) \wedge \text{node}_{u_3}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_3}^{S_4}(w_1) \wedge \text{node}_{u_4}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_1}^{S_4}(w_1) \wedge \text{node}_{u_4}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_2}^{S_4}(w_1) \wedge \text{node}_{u_4}^{S_4}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge \neg n(w_1, w_2)) \end{aligned}$
<p style="text-align: center;">S_5</p>	$\begin{aligned} \text{node}_{u_1}^{S_5}(w) &= x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_2}^{S_5}(w) &= \neg x(w) \wedge \neg y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge \neg r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \\ \text{node}_{u_3}^{S_5}(w) &= \neg x(w) \wedge y(w) \wedge \neg t(w) \wedge \neg e(w) \\ &\wedge r_x(w) \wedge r_y(w) \wedge \neg r_t(w) \wedge \neg r_e(w) \wedge \neg is(w) \end{aligned}$ $\begin{aligned} \xi^{S_5} &= \bigwedge_{i=1,2,3} (\exists v : \text{node}_{u_i}^{S_5}(v)) \\ &\wedge \forall w : \bigvee_{i=1,2,3} \text{node}_{u_i}^{S_5}(w) \\ &\wedge \forall w_1, w_2 : \\ &\quad (\bigwedge_{i=1,2} \text{node}_{u_i}^{S_5}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\bigwedge_{i=1,2} \text{node}_{u_3}^{S_5}(w_i) \Rightarrow \\ &\quad \quad eq(w_1, w_2) \wedge \neg n(w_1, w_2)) \\ &\quad \wedge (\text{node}_{u_1}^{S_5}(w_1) \wedge \text{node}_{u_2}^{S_5}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_2}^{S_5}(w_1) \wedge \text{node}_{u_3}^{S_5}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1)) \\ &\quad \wedge (\text{node}_{u_1}^{S_5}(w_1) \wedge \text{node}_{u_3}^{S_5}(w_2) \Rightarrow \\ &\quad \quad \neg eq(w_1, w_2) \wedge \neg n(w_2, w_1) \wedge \neg n(w_1, w_2)) \end{aligned}$

Table III. Table II continued.

(2) $S^\natural, [w \mapsto u^\natural] \models \text{node}_{u_2}^S(w)$

PROOF. Because S^\natural embeds into S , there exists an embedding function f , such that $S^\natural \sqsubseteq^f S$. For the sake of argument, assume that both claims hold. By Definition 3.1, we get that $f(u^\natural) = u_1$ and $f(u^\natural) = u_2$; because f is a function, we get that $u_1 = u_2$. This yields a contradiction to the assumption that u_1 and u_2 are distinct individuals. \square

Lemma 3.4 For every 2-valued structure S^\natural and assignment Z

$$S^\natural, Z \models p^B(v_1, v_2, \dots, v_k) \text{ iff } \iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) \sqsubseteq B$$

Proof of the if direction: Suppose that $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) \sqsubseteq B$. There are two cases to consider: (i) $B = 1/2$ or (ii) $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) = B$. If $B = 1/2$, then by Definition 3.3, $p^B(v_1, v_2, \dots, v_k) = 1$ and thus $S^\natural, Z \models p^B(v_1, v_2, \dots, v_k)$ for all Z . If $B = 1$, then $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) = 1$, thus $S^\natural, Z \models p(v_1, v_2, \dots, v_k)$ which is $S^\natural, Z \models p^1(v_1, v_2, \dots, v_k)$ by Definition 3.3. Similarly, if $B = 0$, then $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) = 0$ implies that $S^\natural, Z \models \neg p(v_1, v_2, \dots, v_k) = p^0(v_1, v_2, \dots, v_k)$.

Proof of the only-if direction: Assume that $S^\natural, Z \models p^B(v_1, v_2, \dots, v_k)$. If $B = 1/2$, then $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) \sqsubseteq B$ trivially holds. If $B = 0$, apply Definition 3.3 to the assumption to get $S^\natural, Z \models \neg p(v_1, v_2, \dots, v_k)$, which implies

$\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) = 0 = B$. Similarly, if $B = 1$, the assumption implies $\iota^{S^\natural}(p)(Z(v_1), Z(v_2), \dots, Z(v_k)) = 1 = B$.

Lemma 3.6 *Every bounded 3-valued structure S is FO-identifiable, where*

$$\text{node}_{u_i}^S(w) \stackrel{\text{def}}{=} \bigwedge_{p \in \mathcal{P}_1} p^{\iota^S(p)(u_i)}(w)$$

Proof: Consider a bounded 3-valued structure $S = \{U, \iota^S\}$. We shall show that every element $u \in U$ is FO-identifiable using the formula defined in Eq. (4). Let S^\natural be a 2-valued structure that embeds into S using a function f , and let u^\natural be a concrete element in U^{S^\natural} . By Definition 3.1, we have to show that the following holds:

$$f(u^\natural) = u \iff S^\natural, [w \mapsto u^\natural] \models \text{node}_u^S(w)$$

Proof of the if direction: Suppose that $S^\natural, [w \mapsto u^\natural] \models \text{node}_u^S(w)$. In particular, each conjunct of node_u^S must hold, i.e., for each predicate $p \in \mathcal{P}_1$, $S^\natural, [w \mapsto u^\natural] \models p^{\iota^S(p)(u)}(w)$. Using Lemma 3.4 we get that $\iota^{S^\natural}(p)(u^\natural) \sqsubseteq \iota^S(p)(u)$. In addition, the embedding condition in Eq. (1), requires, in particular, that for each unary predicate p $\iota^{S^\natural}(p)(u^\natural) \sqsubseteq \iota^S(p)(f(u^\natural))$ holds. Let $u_1 = f(u^\natural)$. For the sake of argument, assume that $u_1 \neq u$. Recall that S is a bounded structure, in which every individual must have a unique combination of definite values of unary predicates. As a consequence, there must be a unary predicate p such that $\iota^S(p)(u_1) \neq \iota^S(p)(u)$ and the value of p on both u_1 and u is definite. This yields a contradiction, because \sqsubseteq on definite values implies equality; however $\iota^{S^\natural}(p)(u^\natural) = \iota^S(p)(u)$ and $\iota^{S^\natural}(p)(u^\natural) = \iota^S(p)(f(u^\natural)) = \iota^S(p)(u_1)$ can not hold simultaneously, by the assumption.

Proof of the only-if direction: Suppose that $f(u^\natural) = u$. Using Eq. (1), the embedding function f guarantees that for each unary predicate p , $\iota^{S^\natural}(p)(u^\natural) \sqsubseteq \iota^S(p)(f(u^\natural))$. This means that $S^\natural, [w \mapsto u^\natural] \models p^{\iota^S(p)(f(u^\natural))}(w)$ by Lemma 3.4, or $S^\natural, [w \mapsto u^\natural] \models p^{\iota^S(p)(u)}(w)$ by the assumption. This holds for all unary predicates, and thus holds for their conjunction as well, namely, for the formula node_u^S .

LEMMA D.2. *Given a set of formulas F and a 3-valued structure S , if the “focus” algorithm [Lev-Ami 2000, Sec.6] terminates, it returns a set of structures X such that $\gamma(S) = \gamma(X)$ and every formula $\varphi \in F$ evaluates, using the compositional semantics, to a definite value in every structure in X , for every assignment. If the input structure S is FO-Identifiable, then all structures in X are FO-Identifiable.*

Proof: By induction on the iterations of the loop in the “focus” algorithm, it is sufficient to show that the structures returned by the procedure `FocusAssignment` from [Lev-Ami

2000, Fig.17] are FO-Identifiable. The only interesting case is when the input literal of `FocusAssignment` is of the form $p(u_1, \dots, u_k)$. The resulting set of structures X is $\{S_0, S_1, S''\}$ where S_0 and S_1 are copies of S with $p(u_1, \dots, u_k)$ set to 0 and 1, respectively. Thus, if S is FO-identifiable, then S_0 and S_1 are FO-identifiable. S'' is a result of splitting a node $u_i \in S$ into $u.0$ and $u.1$, and setting $p(u_1, \dots, u_k)$ to 0 on one of the copies, and to 1 on the other. To simplify the exposition, suppose that the first node u_1 is split. Then S'' is FO-identifiable using the formulas $\text{node}_u^S(w)$ for all u except $u.0, u.1$, and

$$\begin{aligned} \text{node}_{u.0}^{S''}(w) &\stackrel{\text{def}}{=} \exists v_2, \dots, v_k. \neg p(w, v_2, \dots, v_k) \wedge \text{node}_u^S(w) \wedge \bigwedge_{j=2, \dots, k} \text{node}_{u_j}^S(v_j) \\ \text{node}_{u.1}^{S''}(w) &\stackrel{\text{def}}{=} \exists v_2, \dots, v_k. p(w, v_2, \dots, v_k) \wedge \text{node}_u^S(w) \wedge \bigwedge_{j=2, \dots, k} \text{node}_{u_j}^S(v_j) \end{aligned}$$

Theorem 3.8 *For every FO-identifiable 3-valued structure S , and 2-valued structure S^\natural*

$$S^\natural \in \gamma(S) \text{ iff } S^\natural \models \widehat{\gamma}(S)$$

Proof: In Lemma D.3, we show that the if-direction holds, even when S is not FO-identifiable, i.e., every concrete structure satisfying the characteristic formula $\widehat{\gamma}(S)$ is indeed in $\gamma(S)$. In Lemma D.4 we show the only-if part, i.e., for an FO-identifiable structure, the other direction is also true.

LEMMA D.3. *Let S be a first-order structure with set of individuals $U = \{u_1, u_2, \dots, u_n\}$. Let $\text{node}_{u_i}^S(w)$ used in $\widehat{\gamma}(S)$ be an arbitrary first-order formula free in w , such that Lemma 3.2 holds. Then, for all S^\natural such that $S^\natural \models \widehat{\gamma}(S)$, $S^\natural \in \gamma(S)$.*

Proof: Let $S^\natural = \langle U^\natural, \iota^\natural \rangle$ be a concrete structure such that $S^\natural \models \widehat{\gamma}(S)$. We shall construct a surjective function $f: U^\natural \rightarrow U$ such that $S^\natural \sqsubseteq^f S$. Let Z^\natural be an assignment over v_1, \dots, v_n such that $S^\natural, Z^\natural \models \varphi$, where $\varphi \stackrel{\text{def}}{=} \bigwedge_{i=1}^n \text{node}_{u_i}^S(v_i)$, i.e., φ is the first line of Eq. (8) without the existential quantification. Note that all $Z^\natural(v_i)$ are distinct, according to Lemma 3.2. Define the function $f: U^\natural \rightarrow U$ by:

$$f(u^\natural) = \begin{cases} u_i & \text{if } Z^\natural(v_i) = u^\natural \\ u_j & \text{if for all } i, Z^\natural(v_i) \neq u^\natural \text{ and } u_j \text{ is an arbitrary element such that} \\ & S^\natural, [w \mapsto u^\natural] \models \text{node}_{u_j}^S(w) \end{cases} \quad (25)$$

Let us show that every concrete element is mapped to some element in U . In the case that $Z(v_i) = u^\natural$, the concrete element u^\natural is mapped to $u_i \in U$ by f . Otherwise, because $S^\natural \models \xi^S[\text{total}]$ holds, at least one of its disjuncts must be satisfied by each u^\natural , i.e. $S^\natural, [w \mapsto u^\natural]$ must satisfy $\text{node}_{u_j}^S(w)$ for some u_j ; thus f 's definition will map u^\natural to this u_j . Therefore, $f(u^\natural)$ is well-defined.

In addition, every element $u_i \in U$ is assigned by f to some concrete element $u_i^\natural \in U^\natural$ such that $Z(v_i) = u_i^\natural$. According to Lemma 3.2, all such elements u_i^\natural are different. Therefore, $f(u^\natural)$ is surjective.

Let p be a nullary predicate. Because S^\natural satisfies ξ_{nullary}^S , it must satisfy each conjunct, in particular $S^\natural \models p^{\iota^S(p)}()$. Using Lemma 3.4 we get that $\iota^{S^\natural}(p)() \sqsubseteq \iota^S(p)()$.

Let $p \in P$ be a predicate of arity $r \geq 1$. Let $u_1^\natural, u_2^\natural, \dots, u_r^\natural \in U^\natural$ and let us show that

$$\iota^{S^\natural}(p)(u_1^\natural, u_2^\natural, \dots, u_r^\natural) \sqsubseteq \iota^S(p)(f(u_1^\natural), f(u_2^\natural), \dots, f(u_r^\natural)) \quad (26)$$

Let Z be an assignment such that $Z(w_i) = u_i^\natural$ for $i = 1, \dots, r$. Because $S^\natural \models \xi^S[p]$, we conclude that S^\natural, Z satisfies the body of Eq. (7). Consider the conjunct of the body with

premise $\bigwedge_{j=1}^r \text{node}_{f(u_j^{\natural})}^S(w_j)$. By definition of f , $S^{\natural}, w_j \mapsto u_j^{\natural}$ satisfies $\text{node}_{f(u_j^{\natural})}^S(w_j)$ for all $j = 1, \dots, r$, which means that the premise is satisfied by S^{\natural}, Z . Therefore, the conclusion must hold: $S^{\natural}, Z \models p^{\iota^S(p)}(f(u_1^{\natural}), \dots, f(u_r^{\natural})) (w_1, \dots, w_r)$ and the result follows from Lemma 3.4.

LEMMA D.4. *For every 3-valued FO-identifiable structure S , and 2-valued structure S^{\natural} such that $S^{\natural} \models F$ and $S^{\natural} \sqsubseteq S$, $S^{\natural} \models \xi^S$.*

Proof: Let $f: S^{\natural} \rightarrow S$ be a surjective function such that $S^{\natural} \sqsubseteq^f S$. Let u_i^{\natural} be an arbitrary element such that $f(u_i^{\natural}) = u_i$. Define an assignment Z^{\natural} such that $Z^{\natural}(v_i) = u_i^{\natural}$; u_i^{\natural} must exist because f is surjective. Because S is FO-identifiable, by Definition 3.1 we conclude that for every $1 \leq i \leq n$, $S^{\natural}, Z^{\natural} \models \text{node}_{u_i}^S(v_i)$. Because f is a function, all u_i^{\natural} are distinct elements, according to Lemma 3.2.

Because f is a function, for every u^{\natural} there is u such that $f(u^{\natural}) = u$. Then, by Definition 3.1, $S^{\natural}, [w \mapsto u^{\natural}] \models \text{node}_u^S(w)$, i.e., every assignment to w in S^{\natural} satisfies some disjunct of ξ_{total}^S . That is S^{\natural} satisfies ξ_{total}^S .

For every nullary predicate $p \in \mathcal{P}_0$, using Eq. (1) and Lemma 3.4, we conclude that S^{\natural} satisfies $p^{\iota^S(p)}$. Therefore, S^{\natural} satisfies $\xi_{nullary}^S$.

Let $p \in \mathcal{P}$ be a predicate of arity r . Let $u_1^{\natural}, \dots, u_r^{\natural} \in U^{\natural}$ and let Z^{\natural} be an assignment such that $Z^{\natural}(w_i) = u_i^{\natural}$. We shall show that $S^{\natural}, Z^{\natural}$ satisfy the body of Eq. (7). If the premise of the implication is not satisfied then the formula vacuously holds. Otherwise, $S^{\natural}, Z^{\natural} \models \text{node}_{u_i}^S(w_i)$ for all $i = 1, \dots, r$. Then, by Definition 3.1, $f(u_i^{\natural}) = u_i$. Using Eq. (1) on f , we get $\iota^{S^{\natural}}(p)(u_1^{\natural}, \dots, u_r^{\natural}) \sqsubseteq \iota^S(p)(f(u_1^{\natural}), \dots, f(u_r^{\natural}))$, which means that $\iota^{S^{\natural}}(p)(u_1^{\natural}, \dots, u_r^{\natural}) \sqsubseteq \iota^S(p)(u_1, \dots, u_r)$ holds. By Lemma 3.4, we conclude that $S^{\natural}, Z^{\natural}$ satisfies $p^{\iota^S(p)}(u_1, \dots, u_r)(w_1, \dots, w_r)$.

Lemma A.5 *If 3-valued structure $S = \langle U, \iota^S \rangle$ over vocabulary \mathcal{P} is ICA then:*

- (i). S is a bounded structure.
- (ii). For each nullary predicate p , $\iota^S(p)() \in \{0, 1\}$.
- (iii). For each element $u \in U$, and each unary predicate p , $\iota^S(p)(u) \in \{0, 1\}$.

Proof: Let $S^{\natural} = \{U^{\natural}, \iota^{S^{\natural}}\}$ be a 2-valued structure, such that S is the canonical abstraction of S^{\natural} . Let *canonical*: $U^{\natural} \rightarrow U$ be the mapping that identifies S as the canonical abstraction of S^{\natural} .

(i). Show that S is a bounded structure. By Eq. (11), every abstract element represents concrete elements with the same canonical name. Thus, for two distinct abstract elements $u_0, u_1 \in U^S$, the canonical name of concrete elements represented by u_0 is different from the canonical name of concrete elements represented by u_1 . Without loss of generality, assume that the canonical names differ in a unary predicate p , such that p evaluates to 0 on all concrete elements represented by u_0 , and p evaluates to 1 on all concrete elements represented by u_1 . From the join operation in Eq. (12), it follows that the value of p on u_0 must be 0 and the value of p on u_1 must be 1. This shows that, in general, every pair of distinct elements in S differs in a definite value of some unary predicate, proving that S is a bounded structure.

(ii). Let p be a nullary predicate. Show that $\iota^S(p)() \in \{0, 1\}$. By Eq. (12), $\iota^S(p)() = \sqcup \{\iota^{S^{\natural}}(p)()\} = \iota^{S^{\natural}}(p)()$. This means that p has the same value in S and S^{\natural} . Because S^{\natural} is

a concrete structure, the value of p must be definite.

(iii). Let p be a unary predicate and let $u \in U$. Show that $\iota^S(p)(u) \in \{0, 1\}$. Suppose that the opposite holds: $\iota^S(p)(u) = 1/2$. By Eq. (12), there exist two concrete elements, denoted by u_0 and u_1 , such that $\text{canonical}(u_0) = u$ and $\text{canonical}(u_1) = u$, and p evaluates to 0 on u_0 and to 1 on u_1 . Hence, these concrete elements have different canonical names and by Eq. (11) they cannot be mapped by canonical to the same abstract element; this contradicts the supposition and hence $\iota^S(p)(u) \in \{0, 1\}$.

Lemma A.6 *Every 3-valued structure S that is an ICA is canonical-FO-identifiable, where*

$$\text{node}_{u_i}^S(w) \stackrel{\text{def}}{=} \bigwedge_{p \in \mathcal{P}_1} p^{\iota^S(p)(u_i)}(w) \quad (27)$$

Proof: Let $S = \{U, \iota^S\}$ be a 3-valued structure that is ICA. We shall show that every element $u \in U$ is canonical-FO-identifiable using the formula defined in Eq. (15). Let $S^\natural = \{U^\natural, \iota^{S^\natural}\}$ be a 2-valued structure, such that S is the canonical abstraction of S^\natural , induced by a function canonical , and let $u^\natural \in U^{S^\natural}$. By Definition A.4, we have to show that the following holds:

$$\text{canonical}(u^\natural) = u \iff S^\natural, [w \mapsto u^\natural] \models \text{node}_u^S(w)$$

Proof of the if direction: Suppose that $S^\natural, [w \mapsto u^\natural] \models \text{node}_u^S(w)$. Let $u_1 = \text{canonical}(u^\natural)$. For the sake of argument, assume that $u_1 \neq u$. S is an ICA and using Lemma A.5(i) we get that S is a bounded structure. By Definition 3.5, there exists a unary predicate p that evaluates to different definite values on u and u_1 . Without loss of generality, suppose that p evaluates to 0 on u and to 1 on u_1 . This implies the following two facts. First, from property Eq. (12) of the definition of canonical abstraction, p also evaluates to 1 on all concrete values mapped to u_1 by canonical ; in particular, p must evaluate to 1 on u^\natural . Second, recall that by assumption, each conjunct of node_u^S must hold, i.e., for each predicate $p \in \mathcal{P}_1$, $S^\natural, [w \mapsto u^\natural] \models p^{\iota^{S^\natural}(p)(u^\natural)}(w)$. Because p evaluates to 0 on u , we get from Definition 3.3 that $S^\natural, [w \mapsto u^\natural] \models p^0(w)$, which means $\iota^{S^\natural}(p)(u^\natural) = 0$ and a contradiction is obtained.

Proof of the only-if direction: Suppose that $\text{canonical}(u^\natural) = u$. Because S is an ICA by Lemma A.5(iii) we know that all unary predicates have definite values in S . Let p be a unary predicate. Let $B \in \{1, 0\}$ be such that $\iota^S(p)(u) = B$. Because p has definite value B on u in S , by Eq. (12) it must have the same definite value B on all concrete nodes in S^\natural that are mapped to u by canonical ; in particular, on u^\natural : $\iota^{S^\natural}(p)(u^\natural) = B$. Therefore, using Definition 3.3, $S^\natural, [w \mapsto u^\natural] \models p^B(w)$, in other words, $S^\natural, [w \mapsto u^\natural] \models p^{\iota^{S^\natural}(p)(u^\natural)}(w)$. This holds for all unary predicates, and thus holds for their conjunction as well, i.e., for the formula node_u^S .

Theorem A.8 *For every 3-valued structure S that is an ICA and 2-valued structure S^\natural*

$$S^\natural \in \gamma_c(S) \text{ iff } S^\natural \models \widehat{\gamma}_c(S)$$

Proof: In Lemma D.5, we show that the if-direction holds, i.e., a 3-valued structure S is the canonical abstraction of every concrete structure satisfying the characteristic formula $\widehat{\gamma}_c(S)$; in Lemma D.6 we show the other direction.

LEMMA D.5. *Let S be an ICA with set of individuals $U = \{u_1, u_2, \dots, u_n\}$. Let $\text{node}_{u_i}^S(w)$ be an arbitrary formula free in w , used in $\widehat{\gamma}_c$, such that Lemma 3.2 holds. Then, for all S^{\natural} such that $S^{\natural} \models \widehat{\gamma}_c(S)$, S is a canonical abstraction of S^{\natural} .*

Proof: Let $S^{\natural} = \langle U^{\natural}, \iota^{\natural} \rangle$ be a concrete structure such that $S^{\natural} \models \widehat{\gamma}_c(S)$. We shall construct a surjective function *canonical*: $U^{\natural} \rightarrow U$ such that S^{\natural} is a canonical abstraction of S . From Definition A.7 it follows, in particular, that $S^{\natural} \models \xi^S$. Let Z^{\natural} be an assignment over v_1, \dots, v_n such that $S^{\natural}, Z^{\natural} \models \varphi$, where $\varphi \stackrel{\text{def}}{=} \bigwedge_{i=1}^n \text{node}_{u_i}^S(v_i)$, i.e., φ is the first line of Eq. (8) without the existential quantification). Note that all $Z^{\natural}(v_i)$ are distinct, according to Lemma 3.2. Define the function *canonical*: $U^{\natural} \rightarrow U$ by:

$$\text{canonical}(u^{\natural}) = \begin{cases} u_i & \text{if } Z^{\natural}(v_i) = u^{\natural} \\ u_j & \text{if for all } i, Z^{\natural}(v_i) \neq u^{\natural} \text{ and } u_j \text{ is an arbitrary element such that} \\ & S^{\natural}, [w \mapsto u^{\natural}] \models \text{node}_{u_j}^S(w) \end{cases} \quad (28)$$

Let us show that every concrete element is mapped to some element in U . In the case that $Z(v_i) = u^{\natural}$, the concrete element u^{\natural} is mapped to $u_i \in U$ by *canonical*. Otherwise, because $S^{\natural} \models \xi^S[\text{total}]$ holds, at least one of its disjuncts must be satisfied by each u^{\natural} , i.e., $S^{\natural}, [w \mapsto u^{\natural}]$ must satisfy $\text{node}_{u_i}^S(w)$ for some u_i ; thus *canonical*'s definition will map u^{\natural} to this u_i . Therefore, *canonical*(u^{\natural}) is well-defined.

In addition, every element $u_i \in U$ is assigned by *canonical* to some concrete element $u_i^{\natural} \in U^{\natural}$ such that $Z(v_i) = u_i^{\natural}$. According to Lemma 3.2, all such elements u_i^{\natural} are different. Therefore, *canonical*(u^{\natural}) is surjective.

We shall show that *canonical* satisfies Eq. (11) and Eq. (12); that is, *canonical* identifies S as the canonical abstraction of S^{\natural} .

First, let us show that Eq. (12) holds for the abstraction imposed by *canonical*, namely that a predicate p in S has the most precise abstract value w.r.t. the concrete values that it represents, as is imposed by *canonical*.

Because S is an ICA, all nullary predicates in S must have definite values, by Lemma A.5(ii). S^{\natural} satisfies ξ_{nullary}^S ; therefore, by Definition 3.3, nullary predicates in S^{\natural} must have the same definite values as in S ; this shows that Eq. (12) holds for nullary predicates.

Because S is an ICA, all unary predicates in S must have definite values, by Lemma A.5(iii).

Let p be a unary predicate and let $u \in U$ be an individual of S such that $\iota^S(p)(u) = b$. We shall show that p has the same definite value b on all concrete elements mapped to u by *canonical*. Because the join of these values is also b , we will get that Eq. (12) holds for p and u . Recall that S^{\natural} satisfies formula $\xi^S[p]$, hence each assignment to w satisfies the conjunct $\text{node}_u^S(w) \Rightarrow p^b(w)$ of $\xi^S[p]$. Let $u^{\natural} \in U^{\natural}$ be an individual of U^{\natural} such that *canonical*(u^{\natural}) = u and consider an assignment in which w is mapped to u^{\natural} . By the definition of *canonical*, this assignment satisfies $\text{node}_u^S(w)$, the premise of the conjunct. Therefore, it satisfies the conclusion, i.e., $S^{\natural}, [w \mapsto u^{\natural}]$ satisfies $p^b(w)$. Using Definition 3.3 we get that $\iota^{S^{\natural}}(p)(u^{\natural}) = b$.

Let p be a predicate of arity $r > 1$. If p has a definite value b in S on a tuple u_1, \dots, u_r , $\xi^S[p]$ requires that p evaluates to the same definite value b on every concrete tuple $u_1^{\natural}, \dots, u_r^{\natural}$ such that *canonical*(u_i^{\natural}) = u_i (by the same argument as for unary predicates). Therefore, the join operation returns b as the most precise abstract value of p for these concrete tuples. Otherwise, if p evaluates to $1/2$ on $u_1, \dots, u_r \in U$, there must be two tuples of elements in U^{\natural} , say $u_{01}^{\natural}, \dots, u_{0r}^{\natural}$ and $u_{11}^{\natural}, \dots, u_{1r}^{\natural}$, such that

$S^{\natural}, [w_1 \mapsto u_{01}^{\natural}, \dots, w_r \mapsto u_{0r}^{\natural}] \models \neg p(w_1, \dots, w_r)$ and $S^{\natural}, [w_1 \mapsto u_{11}^{\natural}, \dots, w_1 \mapsto u_{1r}^{\natural}] \models p(w_1, \dots, w_r)$, because $S^{\natural} \models \tau^S[p]$. Thus, p evaluates to 0 on the first tuple and to 1 on the second tuple of the concrete structure; therefore, the most precise value obtained by the join operation on these values is $1/2$.

We shall show that *canonical* satisfies Eq. (11), i.e., it maps elements according to their canonical names. This involves showing two directions:

1.. For the sake of contradiction, assume that there are two distinct elements $u_0^{\natural}, u_1^{\natural} \in U^{\natural}$ that have the same canonical name (meaning that for all $p \in \mathcal{P}_1$, $\iota^{S^{\natural}}(p)(u_0^{\natural}) = \iota^{S^{\natural}}(p)(u_1^{\natural})$), but $\text{canonical}(u_0^{\natural}) \neq \text{canonical}(u_1^{\natural})$. Because S is a bounded structure, there must be unary predicate p that evaluates to 0 on $\text{canonical}(u_0^{\natural})$ and to 1 on $\text{canonical}(u_1^{\natural})$. As shown above, p evaluates to the same definite values in the concrete structure S^{\natural} : $\iota^{S^{\natural}}(p)(u_0^{\natural}) = 0$, and $\iota^{S^{\natural}}(p)(u_1^{\natural}) = 1$ and a contradiction is obtained.

2.. For the sake of contradiction, assume that two concrete elements, denoted by $u_0^{\natural}, u_1^{\natural} \in U^{\natural}$, have different canonical names, but are mapped by *canonical* to the same element in U : $\text{canonical}(u_0^{\natural}) = \text{canonical}(u_1^{\natural})$, denoted by u . By definition of *canonical*, $S^{\natural}, [w \mapsto u_i^{\natural}]$ satisfies $\text{node}_{\text{canonical}(u_i^{\natural})}^S(w)$, for $i = 0, 1$, in other words $S^{\natural}, [w \mapsto u_i^{\natural}]$ satisfies $\text{node}_u^S(w)$. Therefore, it satisfies each conjunct of *node* formula, i.e., for all p , $S^{\natural}, [w \mapsto u_i^{\natural}]$ satisfies $p^{\iota^S(p)(u)}(w)$. From this and the fact that all unary predicates in S have definite values because S is an ICA, we conclude by Definition 3.3, that $\iota^{S^{\natural}}(p)(u_i^{\natural}) = \iota^S(p)(u)$. Therefore, $\iota^{S^{\natural}}(p)(u_0^{\natural}) = \iota^S(p)(u)$ and $\iota^{S^{\natural}}(p)(u_1^{\natural}) = \iota^S(p)(u)$, for all $p \in \mathcal{P}_1$. Therefore, u_0^{\natural} and u_1^{\natural} have the same canonical name and a contradiction is obtained.

LEMMA D.6. *For every 3-valued structure S that is an ICA and 2-valued structure S^{\natural} such that $S^{\natural} \models F$, such that S is the canonical abstraction of S^{\natural} , $S^{\natural} \models \tau^S$.*

Proof: Let *canonical*: $U^{\natural} \rightarrow U$ be the mapping that identifies S as the canonical abstraction of S^{\natural} . *canonical* is a surjective function and possesses the properties in Eq. (11) and Eq. (12).

First, we show that $S^{\natural} \models \xi^S$. Let u_i^{\natural} be an arbitrary element such that $\text{canonical}(u_i^{\natural}) = u_i$. Define an assignment Z^{\natural} such that $Z^{\natural}(v_i) = u_i^{\natural}$; u_i^{\natural} must exist because *canonical* is surjective. Because S is canonical-FO-identifiable, by Lemma A.6 we conclude that for every $1 \leq i \leq n$, $S^{\natural}, Z^{\natural} \models \text{node}_{u_i}^S(v_i)$. According to Lemma 3.2, all the u_i^{\natural} are distinct elements.

Because *canonical* is a function, for every u^{\natural} there is a u such that $\text{canonical}(u^{\natural}) = u$. Then, by Definition A.4, $S^{\natural}, [w \mapsto u^{\natural}] \models \text{node}_u^S(w)$, i.e., every assignment to w in S^{\natural} satisfies some disjunct of ξ_{total}^S . That is, S^{\natural} satisfies ξ_{total}^S .

Because S is an ICA, nullary predicates have the same definite values in S and in S^{\natural} , by Lemma A.5(ii). Therefore, by Definition 3.3, S^{\natural} satisfies $p^{\iota^S(p)(\cdot)}$, for every nullary predicate $p \in \mathcal{P}_0$, which means that S^{\natural} satisfies $\xi_{nullary}^S$.

Let $p \in P$ be a predicate of arity r . Let $u_1^{\natural}, \dots, u_r^{\natural} \in U^{\natural}$ and let Z^{\natural} be an assignment such that $Z^{\natural}(w_i) = u_i^{\natural}$. We shall show that $S^{\natural}, Z^{\natural}$ satisfies the body of Eq. (7). Consider a conjunct of the body. If the premise of the implication in this conjunct is not satisfied, then the conjunct vacuously holds. Otherwise, $S^{\natural}, Z^{\natural} \models \text{node}_{u_i}^S(w_i)$ for all $i = 1, \dots, r$. Then, by Lemma A.6, $\text{canonical}(u_i^{\natural}) = u_i$. We have two cases to consider: (i) if $\iota^S(p)(u_1, \dots, u_r) = b \in \{1, 0\}$ then by Eq. (12) $\iota^{S^{\natural}}(p)(u_1^{\natural}, \dots, u_r^{\natural}) = b$,

in other words, $S^{\natural}, Z^{\natural}$ satisfies $p^b(w_1, \dots, w_r)$. (ii) if $t^S(p)(u_1, \dots, u_r) = 1/2$ then by Definition 3.3, $p^{t^S(p)}(u_1, \dots, u_r)(w_1, \dots, w_r) = p^{1/2}(w_1, \dots, w_r) = \mathbf{1}$, which holds for any assignment.

To complete the proof, we show that for every $p \in \mathcal{P}_r$ of arity $r > 1$, $\tau^S[p]$ holds. Let p be a predicate that evaluates to $1/2$ on a tuple $u_1, \dots, u_r \in S$. Because S is an ICA $t^S(p)(u_1, \dots, u_r) = 1/2$ means that the join operation in Eq. (12) yields $1/2$. By the definition of join as the least upper bound, and using the information order in Definition 2.4, we conclude that (i) S^{\natural} must contain at least two distinct tuples; denoted by $u_{01}^{\natural}, \dots, u_{0r}^{\natural}$ and $u_{11}^{\natural}, \dots, u_{1r}^{\natural}$. Because $\text{canonical}(u_{ij}^{\natural}) = u_j$ for $i = 0, 1$ and $j = 1, \dots, r$, by Lemma A.6 we get that $S^{\natural}, [w \mapsto u_{ij}^{\natural}] \models \text{node}_{u_j}^S(w)$. Therefore, each tuple satisfies $\bigwedge_{j=1}^r \text{node}_{u_j}^S(w_j)$. (ii) p evaluates to 0 on the first tuple and 1 on the second tuple. This shows that $S^{\natural} \models \tau^S[p]$.

Lemma A.9 Denote by \mathcal{D} the set of all 2-valued structures that satisfy the integrity formula $F: \mathcal{D} \stackrel{\text{def}}{=} \{S^{\natural} \in 2\text{-STRUCT}[\mathcal{P}] \mid S^{\natural} \models F\}$. Let S be an ICA structure. There exists a set of ICA structures X such that $\gamma_c(X) = \mathcal{D} \setminus \gamma_c(S)$.

Proof: Denote by Y the set of all ICA structures over a fixed vocabulary \mathcal{P} , i.e., $\gamma_c(Y) = \mathcal{D}$. We claim that X is defined by $Y \setminus S$. By definition, $\gamma_c(X) = \gamma_c(Y \setminus S)$, and we show that $\gamma_c(Y \setminus S) = \gamma_c(Y) \setminus \gamma_c(S)$. By the definitions of Y and γ_c in Eq. (13), $\gamma_c(Y \setminus S) \supseteq \mathcal{D} \setminus \gamma_c(S)$ holds. To complete the proof, we show that the other direction of inclusion holds as well. For the sake of argument, assume that there exists a 2-valued structure S^{\natural} that belongs to both $\gamma_c(S)$ and $\gamma_c(Y \setminus S)$. Thus, by Definition A.3, there exists an ICA structure S' such that $\text{canonical}(S^{\natural}) = S'$, and S' is different from S . From Eq. (12), it follows that $\text{canonical}(S^{\natural}) \neq S$, which contradicts the assumption that $S^{\natural} \in \gamma_c(S)$.

Lemma A.10 Consider the formula τ^S from Eq. (17), for some ICA structure S . There exists a set of ICA structures X , such that the formula $F \wedge \neg \tau^S$ is equivalent to the formula $\widehat{\gamma}_c(X)$.

Proof: Let \mathcal{D} be the set of all 2-valued structures that satisfy the integrity formula F . Let X be the set of ICA structures that describes the complement of $\gamma_c(S)$, as given by Lemma A.9. Let S^{\natural} be a 2-valued structure such that $S^{\natural} \in \gamma_c(X)$ if and only if $S^{\natural} \in \mathcal{D} \setminus \gamma_c(S)$. The right-hand side simplifies to $S^{\natural} \in \mathcal{D}$ and $S^{\natural} \notin \gamma_c(S)$. Applying Theorem A.8, we get that $S^{\natural} \models \widehat{\gamma}_c(X)$ if and only if S^{\natural} satisfies F but does not satisfy $\widehat{\gamma}_c(S)$. Using Eq. (18), this is equivalent to $S^{\natural} \models F \wedge \neg \tau^S$.

Theorem B.2 For every 3-valued structure S , and a 2-valued structure S^{\natural} :

$$S^{\natural} \in \gamma(S) \text{ iff } S^{\natural} \models \widehat{\gamma}_{NP}(S)$$

Proof: In Lemma D.7, we show that the if-direction holds, i.e., every concrete structure satisfying the NP-characteristic formula $\widehat{\gamma}_{NP}$ is indeed in $\gamma(S)$. In Lemma D.8 we show the only-if part.

LEMMA D.7. Let S be a logical structure with set of individuals $U = \{u_1, u_2, \dots, u_n\}$. Then, for all S^{\natural} such that $S^{\natural} \models \widehat{\gamma}_{NP}(S)$, $S^{\natural} \in \gamma(S)$.

Proof: Let $S^{\natural} = \langle U^{\natural}, t^{\natural} \rangle$ be a concrete structure such that $S^{\natural} \models \widehat{\gamma}(S)$. We shall construct a surjective function $f: U^{\natural} \rightarrow U$ such that $S^{\natural} \sqsubseteq^f S$. Let Z^{\natural} be an assignment such that $S^{\natural}, Z^{\natural} \models \varphi$ where φ is the body of ξ^S without the existential quantifiers on sets. Let $Z^{\natural}(V_i) = U_i \subseteq U^{\natural}$. Consider the following definition:

$$f(u^{\natural}) = \{u_i \mid u^{\natural} \in U_i\} \quad (29)$$

$f(u^{\natural})$ is a set of size at most 1 because the pair $S^{\natural}, Z^{\natural}$ satisfies the sub-formula $\xi_{disjoint}^S$. This insures that the sets U_1, \dots, U_n are disjoint, i.e., each concrete element belongs to at most one set. For simplicity, we say that $f(u^{\natural}) = u_i$, whenever $f(u^{\natural}) = \{u_i\}$.

We shall show that every concrete element is mapped by f to some element in U . Because $S^{\natural}, Z^{\natural}$ satisfies ξ_{total}^S , we conclude that every concrete element satisfies the formula $\text{node}_{u_i}^S(w)$ for some u_i . Also, $\text{node}_{u_i}^S(w)$ given in Definition B.1 is a membership test in the set V_i ; therefore, every concrete element must be a member of some set U_i . Thus, u^{\natural} is mapped to $u_i \in U$, by the definition of f in Eq. (29). This shows that f is well-defined.

Because $S^{\natural}, Z^{\natural}$ satisfies $\models \xi_{non_empty}^S[i]$ for $i = 1, \dots, n$, it must be that every U_i contains at least one element, say u_i^{\natural} , that is mapped to u_i by f . Because the sets are disjoint, all such elements u_i^{\natural} are different. Therefore, f is surjective.

Let p be a nullary predicate. Because S^{\natural} satisfies $\xi_{nullary}^S$, it must satisfy each conjunct, in particular $S^{\natural} \models p^{\iota^S(p)()}$. Using Lemma 3.4 we get that $\iota^{S^{\natural}}(p)() \sqsubseteq \iota^S(p)()$.

Let $p \in P$ be a predicate of arity $r \geq 1$. Let $u_1^{\natural}, u_2^{\natural}, \dots, u_r^{\natural} \in U^{\natural}$ and let us show that

$$\iota^{S^{\natural}}(p)(u_1^{\natural}, u_2^{\natural}, \dots, u_r^{\natural}) \sqsubseteq \iota^S(p)(f(u_1^{\natural}), f(u_2^{\natural}), \dots, f(u_r^{\natural})) \quad (30)$$

Let Z_1^{\natural} be an extension of assignment Z^{\natural} such that $Z_1^{\natural}(w_i) = u_i^{\natural}$ for $i = 1, \dots, r$. Because $S^{\natural}, Z^{\natural} \models \xi^S[p]$, we conclude that $S^{\natural}, Z_1^{\natural}$ satisfies the body of Eq. (7). Consider the conjunct of the body with premise $\bigwedge_{j=1}^r \text{node}_{f(u_j^{\natural})}^S(w_j)$. By definition of f , $S^{\natural}, w_j \mapsto u_j^{\natural}$ satisfies $\text{node}_{f(u_j^{\natural})}^S(w_j)$ for all $j = 1, \dots, r$, which means that the premise is satisfied by $S^{\natural}, Z_1^{\natural}$. Therefore, the conclusion must hold:

$S^{\natural}, Z_1^{\natural} \models p^{\iota^S(p)(f(u_1^{\natural}), \dots, f(u_r^{\natural}))}(w_1, \dots, w_r)$ and the result follows from Lemma 3.4.

LEMMA D.8. *For every 3-valued structure S , and 2-valued structure S^{\natural} such that $S^{\natural} \models F$ and $S^{\natural} \sqsubseteq S$, $S^{\natural} \models \xi^S$.*

Proof: Let $f: S^{\natural} \rightarrow S$ be a surjective function such that $S^{\natural} \sqsubseteq^f S$. Define an assignment Z^{\natural} such that $Z^{\natural}(V_i) = U_i \subseteq U^{\natural}$ and $U_i = \{u_i^{\natural} \mid f(u_i^{\natural}) = u_i\}$.

Because f is a surjective function, there must exist at least one concrete element that is mapped to u_i by f . This element belongs to the set U_i . Therefore, $S^{\natural}, Z^{\natural} \models \bigwedge_{i=1}^n \xi_{non_empty}^S[i]$.

Because f is a well-defined function, it maps each concrete element to exactly one element $u_i \in U$, which induces the set U_i . Therefore, a concrete element cannot belong to more than one set; hence $S^{\natural}, Z^{\natural} \models \bigwedge_{k \neq j} \xi_{disjoint}^S[k, j]$.

Because f is a function, f maps every concrete element to some element in U . Therefore, every concrete element belongs to some set, i.e., satisfies some disjunct of ξ_{total}^S . That is $S^{\natural}, Z^{\natural} \models \xi_{total}^S$.

For every nullary predicate $p \in \mathcal{P}_0$, using Eq. (1) and Lemma 3.4, we conclude that $S^{\natural}, Z^{\natural}$ satisfies $p^{\iota^S(p)()}$. Therefore, $S^{\natural}, Z^{\natural} \models \xi_{nullary}^S$.

Let $p \in P$ be a predicate of arity r . Let $u_1^{\natural}, \dots, u_r^{\natural} \in U^{\natural}$ and let Z_1^{\natural} be an extension of assignment Z^{\natural} such that $Z_1^{\natural}(w_i) = u_i^{\natural}$. We shall show that $S^{\natural}, Z_1^{\natural}$ satisfy the body of Eq. (7). If the premise of the implication is not satisfied, then the formula vacuously holds. Otherwise, $S^{\natural}, Z_1^{\natural} \models \text{node}_{u_i}^S(w_i)$ for all $i = 1, \dots, r$. Then, by Definition B.1, u_i^{\natural} belongs to the set U_i . The definition of U_i implies that $f(u_i^{\natural}) = u_i$. Using Eq. (1), we get $\iota^{S^{\natural}}(p)(u_1^{\natural}, \dots, u_r^{\natural}) \sqsubseteq \iota^S(p)(f(u_1^{\natural}), \dots, f(u_r^{\natural}))$ which means $\iota^{S^{\natural}}(p)(u_1^{\natural}, \dots, u_r^{\natural}) \sqsubseteq \iota^S(p)(u_1, \dots, u_r)$. By Lemma 3.4 we conclude that $S^{\natural}, Z^{\natural}$ satisfies $p^{\iota^S(p)(u_1, \dots, u_r)}(w_1, \dots, w_r)$.