Finite Differencing of Logical Formulas for Static Analysis

THOMAS REPS
University of Wisconsin and GrammaTech, Inc.
MOOLY SAGIV
Tel Aviv University
and
ALEXEY LOGINOV
GrammaTech, Inc.

This paper concerns mechanisms for maintaining the value of an instrumentation relation (also known as a *derived relation* or *view*), defined via a logical formula over core relations, in response to changes in the values of the core relations. It presents an algorithm for transforming the instrumentation relation's defining formula into a *relation-maintenance formula* that captures what the instrumentation relation's new value should be.

The technique applies to program-analysis problems in which the semantics of statements is expressed using logical formulas that describe changes to core-relation values, and provides a way to reflect those changes in the values of the instrumentation relations.

Categories and Subject Descriptors: D.2.4 [**Software Engineering**]: Software/Program Verification—*formal methods*; D.3.3 [**Programming Languages**]: Language Constructs and Features—*data types and structures; dynamic storage management*; E.1 [**Data**]: Data Structures—*graphs and networks; lists, stacks, and queues; records; trees*; E.2 [**Data**]: Data Storage Representations—*composite structures; linked representations*; F.3.1 [**Logics and Meanings of Programs**]: Specifying and Verifying and Reasoning about Programs—*assertions; invariants; mechanical verification*; F.3.2 [**Logics and Meanings of Programs**]: Semantics of Programming Languages—*program analysis*

General Terms: Algorithms, Languages, Theory, Verification

Additional Key Words and Phrases: Abstract interpretation, finite differencing, materialized view, shape analysis, static analysis, 3-valued logic

1. INTRODUCTION

This paper addresses an instance of the following fundamental challenge in abstract interpretation:

Authors' addresses: T. Reps, Comp. Sci. Dept., University of Wisconsin, and GrammaTech, Inc., reps@cs.wisc.edu. M. Sagiv, School of Comp. Sci., Tel Aviv University, msagiv@post.tau.ac.il. A. Loginov, GrammaTech, Inc., alexey@grammatech.com. At the time the research reported in the paper was carried out, A. Loginov was affiliated with the Univ. of Wisconsin.

The work was supported in part by NSF under grants CCR-{9619219,9986308}, and CCF-{0540955,0524051}, by the U.S.-Israel BSF under grant 96-00337, by ONR under contracts N00014-01-1-{0708,0796}, and by the von Humboldt and Guggenheim Foundations. Portions of the work appeared in the 12th European Symp. on Programming [Reps et al. 2003], R. Wilhelm's 60^{th}-birthday Festschrift [Loginov et al. 2007], and A. Loginov's Ph.D. dissertation [Loginov 2006].

> Given the concrete semantics for a language and a desired abstraction, how does one create the associated abstract transformers?

The problem that we address arises in program-analysis problems in which the semantics of statements is expressed using logical formulas that describe changes to core-relation values. When instrumentation relations (defined via logical formulas over the core relations) have been introduced to refine an abstraction, the challenge is to reflect the changes in core-relation values in the values of the instrumentation relations [Graf and Saïdi 1997; Das et al. 1999; McMillan 1999; Sagiv et al. 2002; Ball et al. 2001]. The algorithm presented in this paper provides a way to create formulas that maintain correct values for the instrumentation relations, and thereby provides a way to generate, completely automatically, the part of the transformers of an abstract semantics that deals with instrumentation relations. The algorithm runs in time linear in the size of the instrumentation relation's defining formula.

This research was motivated by our work on static analysis based on 3-valued logic [Sagiv et al. 2002]; however, any analysis method that relies on logic—2-valued or 3-valued—to express a program's semantics may be able to benefit from these techniques.

In our setting, two related logics come into play: an ordinary 2-valued logic, as well as a related 3-valued logic. A memory configuration, or store, is modeled by what logicians call a *logical structure*; an individual of the structure's universe either models a single memory element or, in the case of a *summary individual*, it models a collection of memory elements. A run of the analyzer carries out an abstract interpretation to collect a set of structures at each program point P. This involves finding the least fixed point of a certain set of equations. When the fixed point is reached, the structures that have been collected at program point P describe a superset of all the execution states that can occur at P. To determine whether a property always holds 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. The TVLA system (**T**hree-**V**alued-**L**ogic **A**nalyzer) implements this approach [Lev-Ami and Sagiv 2000; TVLA].

Summary individuals play a crucial role. They are used to ensure that abstract descriptors have an *a priori* bounded size, which guarantees that a fixed-point is always reached. However, the constraint of working with limited-size descriptors implies a loss of information about the store. Intuitively, certain properties of concrete individuals are lost due to abstraction, which groups together multiple individuals into summary individuals: a property can be true for some concrete individuals of the group but false for other individuals. It is for this reason that 3-valued logic is used; uncertainty about a property's value is captured by means of the third truth value, 1/2.

An advantage of using 2- and 3-valued logic as the basis for static analysis is that the language used for extracting information from the concrete world and the abstract world is identical: *every* syntactic expression—i.e., every logical formula—can be interpreted either in the 2-valued world or the 3-valued world. The consistency of the 2-valued and 3-valued viewpoints is ensured by a basic theorem that relates the two logics [Sagiv et al. 2002, Theorem 4.9]. This provides a partial answer to the fundamental challenge posed above: formulas that define the concrete semantics, when interpreted in 2-valued logic, define a sound abstract semantics when interpreted in 3-valued logic [Sagiv et al. 2002].

Unfortunately, unless some care is taken in the design of an analysis, there is a danger that as abstract interpretation proceeds, the indefinite value 1/2 will become pervasive. This can destroy the ability to recover interesting information from the 3-valued structures collected (although soundness is maintained). A key role in combating indefiniteness is played by *instrumentation relations*, which record auxiliary information in a logical structure. The benefit of introducing instrumentation relations was annunciated as the Instrumentation Principle:

OBSERVATION 1.1. **(Instrumentation Principle [Sagiv et al. 2002, Observation 2.8])**. *Suppose that S is a 3-valued structure that represents the 2-valued structure $S^\natural$. By explicitly "storing" in S the values that a formula φ has in $S^\natural$, it is sometimes possible to extract more precise information from S than can be obtained just by evaluating φ in S.*

Instrumentation relations provide a mechanism to fine-tune an abstraction: an instrumentation relation, which is defined by a logical formula φ over the core relation symbols, captures a property that may or may not be possessed by a structure, an individual memory cell, or a tuple of memory cells (according to whether φ is a nullary, unary, or k-ary formula, respectively). In general, the introduction of additional instrumentation relations refines an abstraction into one that is prepared to track finer distinctions among stores. For reasons discussed in §3, the values of instrumentation relations are stored and maintained in response to the store transformations performed by program statements. In general, this allows more precise properties of the program's stores to be established.

From the standpoint of the concrete semantics, instrumentation relations represent cached information that could always be recomputed by reevaluating the instrumentation relation's defining formula in the local state. From the standpoint of the abstract semantics, however, reevaluating a formula in the local (3-valued) state can lead to a drastic loss of precision. To gain maximum benefit from instrumentation relations, an abstract-interpretation algorithm must obtain their values in some other way.

This problem, *instrumentation-relation maintenance*, is solved by incremental computation. After a transition via transformer τ from abstract state σ to abstract state σ', the new value that instrumentation relation p should have is computed from the stored value of p in σ. The contributions of the work reported in the paper can be summarized as follows:

- We give an algorithm for the relation-maintenance problem; it creates a relation-maintenance formula by applying a finite-differencing transformation to p's defining formula. The algorithm runs in time linear in the size of the defining formula.
- We present experimental evidence that our technique is an effective one, at least for the analysis of programs that manipulate (cyclic and acyclic) singly-linked lists, doubly-linked lists, and binary trees, and for certain sorting programs. In particular, the relation-maintenance formulas produced automatically using our

approach are as effective for maintaining precision as the best available hand-crafted ones.

—This work is related to the view-maintenance problem in databases [Gupta and Mumick 1999; Dong and Su 1995; Patnaik and Immerman 1997]. Compared with that work, the novelty is the ability to create relation-maintenance formulas that are suitable for use when abstraction has been performed.

The remainder of the paper is organized as follows: §2 introduces terminology and notation. §3 defines the relation-maintenance problem. §4 presents a method for generating maintenance formulas for instrumentation relations. §5 discusses extensions to handle instrumentation relations that use transitive closure. §6 presents experimental results. §7 discusses related work. §8 presents some concluding remarks. Finally, the Appendix presents a proof of the correctness of our solution to the relation-maintenance problem.

2. BACKGROUND

This section introduces terminology and notation; it presents the logic that we employ and describes the use of logical structures for representing memory stores.

The first half of §2.1 introduces 2-valued first-order logic with transitive closure. These concepts are standard in logic. The second half of §2.1 presents a straightforward extension of the logic to the 3-valued setting, in which a third truth value—1/2—is introduced to denote uncertainty. §2.2 summarizes the program-analysis framework described in [Sagiv et al. 2002]. In that approach, memory configurations are encoded as logical structures, the semantics of programs, as well as the properties of memory configurations, is encoded as logical formulas, and abstract interpretation computes the set of logical structures that describe the memory configurations that can arise at each point in the program being analyzed.

2.1 First-Order Logic with Transitive Closure

2-*Valued First-Order Logic with Transitive Closure.* The syntax of first-order formulas with equality and reflexive transitive closure is defined as follows:

Definition 2.1. Let R_i denote a set of arity-i relation symbols, with $eq \in R_2$. A *formula* over the *vocabulary* $\mathcal{R} = \bigcup_i R_i$ is defined by

$$\begin{array}{lll}
p \in \mathcal{R}_k & \varphi ::= & \mathbf{0} \mid \mathbf{1} \mid p(v_1, \ldots, v_k) \\
\varphi \in \mathit{Formulas} & & \mid (\neg\varphi_1) \mid (\varphi_1 \wedge \varphi_2) \mid (\varphi_1 \vee \varphi_2) \mid (\exists v \colon \varphi_1) \mid (\forall v \colon \varphi_1) \\
v \in \mathit{Variables} & & \mid (\mathbf{RTC}\ v_1', v_2' \colon \varphi_1)(v_1, v_2)
\end{array}$$

A formula of the form $\mathbf{0}$, $\mathbf{1}$, or $p(v_1, \ldots, v_k)$ is called an *atomic formula*.

The set of *free variables* of a formula is defined as usual. "**RTC**" stands for reflexive transitive closure. In $\varphi \equiv (\mathbf{RTC}\ v_1', v_2' \colon \varphi_1)(v_1, v_2)$, if φ_1's free-variable set is V, we require $v_1, v_2 \notin V$. The free variables of φ are $(V - \{v_1', v_2'\}) \cup \{v_1, v_2\}$.

We use several shorthand notations: $(v_1 = v_2) \overset{\text{def}}{=} eq(v_1, v_2)$; $(v_1 \neq v_2) \overset{\text{def}}{=} \neg eq(v_1, v_2)$; and for a binary relation p, $p^*(v_1, v_2) \overset{\text{def}}{=} (\mathbf{RTC}\ v_1', v_2' \colon p(v_1', v_2'))(v_1, v_2)$. We also use a C-like syntax for conditional expressions: $\varphi_1\ ?\ \varphi_2 : \varphi_3$.[1] The order of precedence among the connectives, from highest to lowest, is as follows: $\neg$, $\wedge$, $\vee$, $\forall$, and $\exists$. We drop parentheses wherever possible, except for emphasis.

Definition 2.2. A *2-valued interpretation* over $\mathcal{R}$ is a *2-valued logical structure* $S = \langle U^S, \iota^S \rangle$, where U^S is a set of *individuals* and ι^S maps each relation symbol p of arity k to a truth-valued function: $\iota^S(p)\colon (U^S)^k \to \{0, 1\}$. In addition, (i) for all $u \in U^S$, $\iota^S(eq)(u, u) = 1$, and (ii) for all $u_1, u_2 \in U^S$ such that u_1 and u_2 are distinct individuals, $\iota^S(eq)(u_1, u_2) = 0$.

An *assignment* Z is a function that maps variables to individuals (i.e., it has the functionality $Z\colon \{v_1, v_2, \ldots\} \to U^S$). When Z is defined on all free variables of a formula φ, we say that Z is *complete* for φ. (We generally assume that every assignment that arises in connection with the discussion of some formula φ is complete for φ.)

The *(2-valued) meaning* of a formula φ, denoted by $\llbracket \varphi \rrbracket_2^S(Z)$, yields a truth value in $\{0, 1\}$; it is defined inductively as follows:

$$
\begin{array}{rclrcl}
\llbracket \mathbf{0} \rrbracket_2^S(Z) &=& 0 & \llbracket \varphi_1 \wedge \varphi_2 \rrbracket_2^S(Z) &=& \min(\llbracket \varphi_1 \rrbracket_2^S(Z), \llbracket \varphi_2 \rrbracket_2^S(Z)) \\
\llbracket \mathbf{1} \rrbracket_2^S(Z) &=& 1 & \llbracket \varphi_1 \vee \varphi_2 \rrbracket_2^S(Z) &=& \max(\llbracket \varphi_1 \rrbracket_2^S(Z), \llbracket \varphi_2 \rrbracket_2^S(Z)) \\
\llbracket p(v_1, \ldots, v_k) \rrbracket_2^S(Z) &=& \iota^S(p)(Z(v_1), \ldots, Z(v_k)) & \llbracket \exists v\colon \varphi_1 \rrbracket_2^S(Z) &=& \max\limits_{u \in U^S} \llbracket \varphi_1 \rrbracket_2^S(Z[v_1 \mapsto u]) \\
\llbracket \neg \varphi_1 \rrbracket_2^S(Z) &=& 1 - \llbracket \varphi_1 \rrbracket_2^S(Z) & \llbracket \forall v\colon \varphi_1 \rrbracket_2^S(Z) &=& \min\limits_{u \in U^S} \llbracket \varphi_1 \rrbracket_2^S(Z[v_1 \mapsto u])
\end{array}
$$

$$
\begin{aligned}
&\llbracket (\mathbf{RTC}\ v_1', v_2' \colon \varphi_1)(v_1, v_2) \rrbracket_2^S(Z) \\
&= \begin{cases} 1 & \text{if} Z(v_1) = Z(v_2) \\ \max\limits_{\substack{n \geq 1, \\ u_1, \ldots, u_{n+1} \in U, \\ Z(v_1) = u_1, \\ Z(v_2) = u_{n+1}}} \ \min\limits_{i=1}^{n} \llbracket \varphi_1 \rrbracket_2^S(Z[v_1' \mapsto u_i, v_2' \mapsto u_{i+1}]) & \text{otherwise} \end{cases}
\end{aligned}
$$

S and Z *satisfy* φ if $\llbracket \varphi \rrbracket_2^S(Z) = 1$. The set of 2-valued structures is denoted by $\mathcal{S}_2[\mathcal{R}]$.

3-*Valued Logic and Embedding.* In 3-valued logic, the formulas that we work with are identical to the ones used in 2-valued logic. At the semantic level, a third truth value—1/2—is introduced to denote uncertainty.

Definition 2.3. The truth values 0 and 1 are *definite values*; 1/2 is an *indefinite value.* For $l_1, l_2 \in \{0, 1/2, 1\}$, the *information order* is defined as follows: $l_1 \sqsubseteq l_2$ iff $l_1 = l_2$ or $l_2 = 1/2$. $l_1 \sqsubseteq l_2$ denotes that l_1 is at least as definite as l_2. We use $l_1 \sqsubset l_2$ when $l_1 \sqsubseteq l_2$ and $l_1 \neq l_2$. The symbol $\sqcup$ denotes the least-upper-bound operation with respect to $\sqsubseteq$.

As shown in Fig. 1, we place two orderings on 0, 1, and 1/2: (i) the *information order*, denoted by $\sqsubseteq$ and illustrated in Fig. 1(a), captures "(un)certainty"; (ii) the *logical order*, shown in Fig. 1(b), defines the meaning of $\wedge$ and $\vee$; that is, $\wedge$ and $\vee$ are meet and join in the logical order. 3-valued logic retains a number of properties that are familiar from 2-valued logic, such as De Morgan's laws, associativity of $\wedge$ and $\vee$, and distributivity of $\wedge$ over $\vee$ (and vice versa). Because $\varphi_1\ ?\ \varphi_2 : \varphi_3$ is treated as a shorthand for $(\varphi_1 \wedge \varphi_2) \vee (\neg\varphi_1 \wedge \varphi_3) \vee (\varphi_2 \wedge \varphi_3)$ in 3-valued logic [Reps

[1] In 2-valued logic, one can think of $\varphi_1\ ?\ \varphi_2 : \varphi_3$ as a shorthand for $(\varphi_1 \wedge \varphi_2) \vee (\neg\varphi_1 \wedge \varphi_3)$. In 3-valued logic, it becomes a shorthand for $(\varphi_1 \wedge \varphi_2) \vee (\neg\varphi_1 \wedge \varphi_3) \vee (\varphi_2 \wedge \varphi_3)$ [Reps et al. 2002].

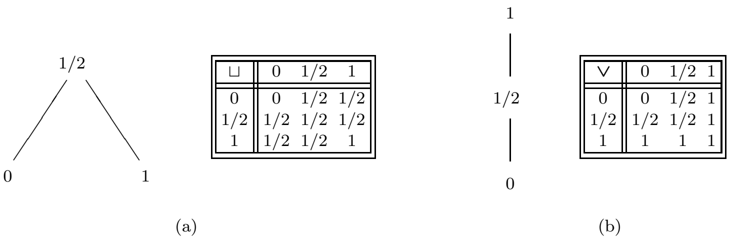

$\sqcup$	0	1/2	1
0	0	1/2	1/2
1/2	1/2	1/2	1/2
1	1/2	1/2	1

$\vee$	0	1/2	1
0	0	1/2	1
1/2	1/2	1/2	1
1	1	1	1

Fig. 1. (a) The information order ($\sqsubseteq$) and its join operation ($\sqcup$). (b) The logical order and its join operation ($\vee$).

et al. 2002], the value of $1/2 \mathbin{?} V_1 : V_2$ equals $V_1 \sqcup V_2$. We now generalize Defn. 2.2 to define the meaning of a formula with respect to a 3-valued structure.

Definition 2.4. A *3-valued interpretation* over $\mathcal{R}$ is a *3-valued logical structure* $S = \langle U^S, \iota^S \rangle$, where U^S is a set of individuals and ι^S maps each relation symbol p of arity k to a truth-valued function: $\iota^S(p)\colon (U^S)^k \to \{0, 1/2, 1\}$. 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 individuals, $\iota^S(eq)(u_1, u_2) = 0$.

For an assignment Z, the *(3-valued) meaning* of a formula φ, denoted by $[\![\varphi]\!]_3^S(Z)$, yields a truth value in $\{0, 1/2, 1\}$. The meaning of φ is defined exactly as in Defn. 2.2, but interpreted over $\{0, 1/2, 1\}$. S and Z *potentially satisfy* φ if $[\![\varphi]\!]_3^S(Z) \sqsupseteq 1$. The set of 3-valued structures is denoted by $\mathcal{S}_3[\mathcal{R}]$.

Defn. 2.4 requires that for each individual u, the value of $\iota^S(eq)(u,u)$ is 1 or 1/2. An individual for which $\iota^S(eq)(u,u) = 1/2$ is called a *summary individual*. In the program-analysis framework of [Sagiv et al. 2002], a summary individual abstracts one or more nodes of a data structure, and hence can represent more than one concrete memory cell.

The embedding ordering on structures is defined as follows:

Definition 2.5. Let $S = \langle U^S, \iota^S \rangle$ and $S' = \langle U^{S'}, \iota^{S'} \rangle$ be two structures, and let $f\colon U^S \to U^{S'}$ be a surjective function. We say that f *embeds* S in S' (denoted by $S \sqsubseteq^f S'$) if for every relation symbol $p \in \mathcal{R}$ of arity k and for all $u_1, \ldots, u_k \in U^S$, $\iota^S(p)(u_1, \ldots, u_k) \sqsubseteq \iota^{S'}(p)(f(u_1), \ldots, f(u_k))$. We say that S *can be embedded in* S' (denoted by $S \sqsubseteq S'$) if there exists a function f such that $S \sqsubseteq^f S'$.

The Embedding Theorem says that if $S \sqsubseteq^f S'$, then every piece of information extracted from S' via a formula φ is a conservative approximation of the information extracted from S via φ. To formalize this, we extend mappings on individuals to operate on assignments: if $f\colon U^S \to U^{S'}$ is a function and $Z\colon Var \to U^S$ is an assignment, $f \circ Z$ denotes the assignment $f \circ Z\colon Var \to U^{S'}$ such that $(f \circ Z)(v) = f(Z(v))$.

Theorem 2.6. **(Embedding Theorem [Sagiv et al. 2002, Theorem 4.9])**. *Let $S = \langle U^S, \iota^S \rangle$ and $S' = \langle U^{S'}, \iota^{S'} \rangle$ be two structures, and let $f\colon U^S \to U^{S'}$ be a function such that $S \sqsubseteq^f S'$. Then, for every formula φ and complete assignment Z for φ, $[\![\varphi]\!]_3^S(Z) \sqsubseteq [\![\varphi]\!]_3^{S'}(f \circ Z)$.*

2.2 Stores as Logical Structures and their Abstractions

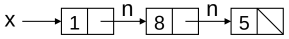

Fig. 2. A possible store for a linked list.

```
typedef struct node {
  struct node *n;
  int data;
} *List;
```

(a)

Relation	Intended Meaning
$eq(v_1, v_2)$	Do v_1 and v_2 denote the same memory cell?
$x(v)$	Does pointer variable `x` point to memory cell v?
$n(v_1, v_2)$	Does the `n` field of v_1 point to v_2?
$dle(v_1, v_2)$	Is the `data` field of v_1 less than or equal to that of v_2?

(b)

Fig. 3. (a) Declaration of a linked-list datatype in C. (b) Core relations used for representing the stores manipulated by programs that use type `List`.

Program Analysis Via 3-Valued Logic. The remainder of this section summarizes the program-analysis framework described in [Sagiv et al. 2002]. In that approach, concrete memory configurations (i.e., *stores*) are encoded as logical structures (associated with a *vocabulary* of relation symbols with given arities) in terms of a fixed collection of *core relations*, $\mathcal{C}$. Core relations are part of the underlying semantics of the language to be analyzed; they record atomic properties of stores. For instance, Fig. 3 gives the definition of a C linked-list datatype, and lists the relations that would be used to represent the stores manipulated by programs that use type `List`, such as the store in Fig. 2. (The core relations are fixed for a given combination of language and datatype; in general, different languages and datatypes require different collections of core relations.) 2-valued logical structures then represent memory configurations: the individuals of the structure are the set of memory cells; a nullary relation represents a Boolean variable of the program; a unary relation represents either a pointer variable or a Boolean-valued field of a record; and a binary relation represents a pointer field of a record. In this example, unary relations represent pointer variables and binary relation n represents the `n`-field of a `List` cell. Numeric-valued variables and numeric-valued fields (such as `data`) can be modeled by introducing other relations, such as the binary relation *dle* (which stands for "`data` less-than-or-equal-to") listed in Fig. 3; *dle* captures the relative order of two nodes' `data` values. (Alternatively, numeric-valued entities can be handled by combining abstractions of logical structures with previously known techniques for creating numeric abstractions [Gopan et al. 2004].) Fig. 4 shows 2-valued structure S_4, which represents the store of Fig. 2 using the relations of Fig. 3. S_4 has three individuals, u_1, u_2, and u_3, which represent the three list elements.

Information about a concrete memory configuration encoded as a logical structure can be extracted from the logical structure by evaluating formulas.

A concrete operational semantics is defined by specifying a *structure transformer* $st_{(n_1,n_2)}$ for each outgoing control-flow graph (CFG) edge (n_1, n_2). (Ordinarily (n_1, n_2) is understood, and we just write *st*.) A structure transformer is specified by providing a collection of *relation-transfer formulas*, $\tau_{c,st}$, one for each core relation c. These formulas define how the core relations of a 2-valued logical structure S

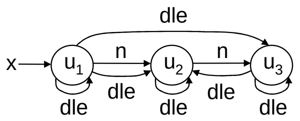

	x
u_1	1
u_2	0
u_3	0

n	u_1	u_2	u_3
u_1	0	1	0
u_2	0	0	1
u_3	0	0	0

dle	u_1	u_2	u_3
u_1	1	1	1
u_2	0	1	0
u_3	0	1	1

Fig. 4. A logical structure S_4 that represents the store shown in Fig. 2 in graphical and tabular forms using the relations of Fig. 3 (Relation *eq* is not shown explicitly; each node has an *eq* self-loop, and the relation in tabular form is the identity matrix.)

that arises at n_1 are transformed by $st_{(n_1,n_2)}$ to create a 2-valued logical structure S' at n_2; typically, they define the value of relation c in S' as a function of c's value in S and the values of other core relations in S. Transformer *st* may optionally have a *precondition formula*, which filters out structures that should not follow the transition along (n_1, n_2). The postcondition operator *post* for edge (n_1, n_2) is defined by lifting (n_1, n_2)'s structure transformer to sets of structures.

Abstract stores are 3-valued logical structures. Concrete stores are abstracted to abstract stores by means of *embedding functions*—onto functions that map individuals of a 2-valued structure $S^\natural$ to those of a 3-valued structure S. The Embedding Theorem ensures that every piece of information extracted from S by evaluating a formula φ is a conservative approximation ($\sqsupseteq$) of the information extracted from $S^\natural$ by evaluating φ.

To obtain a computable abstract domain, we ensure that the size of the 3-valued structures used to represent memory configurations is always bounded. We do this by defining an equivalence relation on individuals and considering the (bounded-size) quotient structure with respect to this equivalence relation; in particular, each individual of a 2-valued logical structure (representing a concrete memory cell) is mapped to an individual of a 3-valued logical structure according to the vector of values that the concrete individual has for a user-chosen collection of unary abstraction relations:

Definition (Canonical Abstraction). Let $S \in \mathcal{S}_2[\mathcal{R}]$, and let $\mathcal{A} \subseteq \mathcal{R}_1$ be some chosen (nonempty) subset of the unary relation symbols. The relations in $\mathcal{A}$ are called *abstraction relations*; they define the following equivalence relation $\simeq_\mathcal{A}$ on U^S:

$$u_1 \simeq_\mathcal{A} u_2 \iff \text{for all } p \in \mathcal{A}, p^S(u_1) = p^S(u_2),$$

and the surjective function $f_\mathcal{A} : U^S \to U^S / \simeq_\mathcal{A}$, such that $f_\mathcal{A}(u) = [u]_{\simeq_\mathcal{A}}$, which maps an individual to its equivalence class. The *canonical abstraction* of S with respect to $\mathcal{A}$ (denoted by $f_\mathcal{A}(S)$) performs the join (in the information order) of relation values, thereby introducing 1/2's.

Intuitively, canonical abstraction maps a group of individuals that are indistinguishable according to the set of (unary) abstraction relations $\mathcal{A}$ to a single individual.

If $\mathcal{A} = \{x\}$, the canonical abstraction of 2-valued logical structure S_4 is S_5, shown

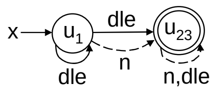

	x
u_1	1
u_{23}	0

n	u_1	u_{23}
u_1	0	1/2
u_{23}	0	1/2

dle	u_1	u_{23}
u_1	1	1
u_{23}	0	1/2

Fig. 5. A 3-valued structure S_5 that is the canonical abstraction of structure S_4.

in Fig. 5, with $f_{\mathcal{A}}(u_1) = u_1$ and $f_{\mathcal{A}}(u_2) = f_{\mathcal{A}}(u_3) = u_{23}$. In addition to S_4, S_5 represents any list with two or more elements that is pointed to by program variable `x`, and in which the first element's `data` value is (definitely) lower than the `data` values in the rest of the list (note the absence of either a 1-valued or 1/2-valued *dle* edge from individual u_{23} to individual u_1). The following graphical notation is used for depicting 3-valued logical structures:

—Individuals are represented by circles containing their names and (non-0) values for unary relations. Summary individuals are represented by double circles.

—A unary relation p corresponding to a pointer-valued program variable is represented by a solid arrow from p to the individual u for which $p(u) = 1$, and by the absence of a p-arrow to each node u' for which $p(u') = 0$. (If $p = 0$ for all individuals, the relation name p is not shown.)

—A binary relation q is represented by a solid arrow labeled q between each pair of individuals u_i and u_j for which $q(u_i, u_j) = 1$, and by the absence of a q-arrow between pairs u'_i and u'_j for which $q(u'_i, u'_j) = 0$.

—Relations with value 1/2 are represented by dotted arrows.

Canonical abstraction ensures that each 3-valued structure is no larger than some fixed size, known *a priori*.

2.2.1 *Instrumentation Relations.* The abstraction function on which an analysis is based, and hence the precision of the analysis defined, can be tuned by (i) choosing to equip structures with additional *instrumentation relations* to record derived properties, and (ii) varying which of the unary core and unary instrumentation relations are used as the set of abstraction relations. The set of instrumentation relations is denoted by $\mathcal{I}$. Each arity-k relation symbol $p \in \mathcal{I}$ is defined by an *instrumentation-relation definition formula* $\psi_p(v_1, \ldots, v_k)$. Instrumentation relations may appear in the defining formulas of other instrumentation relations as long as there are no circular dependences.

The introduction of unary instrumentation relations that are used as abstraction relations provides a way to control which concrete individuals are merged together into an abstract individual, and thereby control the amount of information lost by abstraction. Instrumentation relations that involve reachability properties, which can be defined using **RTC**, often play a crucial role in the definitions of abstractions. For instance, in program-analysis applications, reachability properties from specific pointer variables have the effect of keeping disjoint sublists summarized separately. Fig. 6 lists some instrumentation relations that are important for the analysis of programs that use type `List`.

p	**Intended Meaning**	ψ_p
$is_n(v)$	Do `n` fields of two or more list nodes point to v?	$\exists\, v_1, v_2 \colon n(v_1, v) \wedge n(v_2, v) \wedge v_1 {\neq} v_2$
$t_n(v_1, v_2)$	Is v_2 reachable from v_1 along zero or more `n` fields?	$n^*(v_1, v_2)$
$r_{n,x}(v)$	Is v reachable from pointer variable `x` along zero or more `n` fields?	$\exists\, v_1 \colon x(v_1) \wedge t_n(v_1, v)$
$c_n(v)$	Is v on a directed cycle of `n` fields?	$\exists\, v_1 \colon n(v_1, v) \wedge t_n(v, v_1)$

Fig. 6. Defining formulas of some commonly used instrumentation relations. The relation name is_n abbreviates "is-shared". There is a separate reachability relation $r_{n,x}$ for every program variable `x`. (Recall that $v_1 {\neq} v_2$ is a shorthand for $\neg eq(v_1, v_2)$, and $n^*(v_1, v_2)$ is a shorthand for $(\mathbf{RTC}\ v_1', v_2' \colon n(v_1', v_2'))(v_1, v_2)$.)

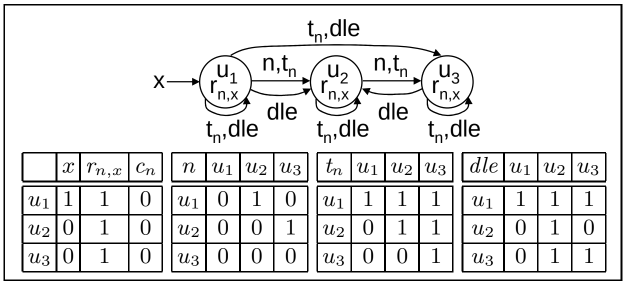

	x	$r_{n,x}$	c_n
u_1	1	1	0
u_2	0	1	0
u_3	0	1	0

n	u_1	u_2	u_3
u_1	0	1	0
u_2	0	0	1
u_3	0	0	0

t_n	u_1	u_2	u_3
u_1	1	1	1
u_2	0	1	1
u_3	0	0	1

dle	u_1	u_2	u_3
u_1	1	1	1
u_2	0	1	0
u_3	0	1	1

Fig. 7. A logical structure S_7 that represents the store shown in Fig. 2 in graphical and tabular forms using the relations of Figs. 3 and 6

.

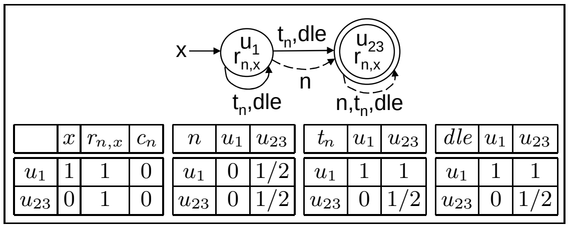

	x	$r_{n,x}$	c_n
u_1	1	1	0
u_{23}	0	1	0

n	u_1	u_{23}
u_1	0	1/2
u_{23}	0	1/2

t_n	u_1	u_{23}
u_1	1	1
u_{23}	0	1/2

dle	u_1	u_{23}
u_1	1	1
u_{23}	0	1/2

Fig. 8. A 3-valued structure S_8 that is the canonical abstraction of structure S_7.

Fig. 7 shows 2-valued structure S_7, which represents the store of Fig. 2 using the core relations of Fig. 3, as well as the instrumentation relations of Fig. 6. If all unary relations are abstraction relations ($\mathcal{A} = \mathcal{R}_1$), the canonical abstraction of 2-valued logical structure S_7 is S_8, shown in Fig. 8, with $f_{\mathcal{A}}(u_1) = u_1$ and $f_{\mathcal{A}}(u_2) = f_{\mathcal{A}}(u_3) = u_{23}$.

2.2.2 *Abstract Interpretation.* For each kind of statement in the programming language, the abstract semantics is again defined by a collection of formulas: the *same* relation-transfer formula that defines the concrete semantics, in the case of a core relation, and, in the case of an instrumentation relation p, by a *relation-maintenance formula* $\mu_{p,st}$.[2]

[2] In [Sagiv et al. 2002], relation-transfer formulas and relation-maintenance formulas are both called "relation-update formulas". Here we use separate terms so that we can refer easily to relation-

<table>
<tr><td>Structure before</td><td><b>unary rels.</b><table><tr><th>indiv.</th><th>x</th><th>y</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>0</td></tr></table></td><td><b>binary rels.</b><table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table></td><td><table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table></td><td>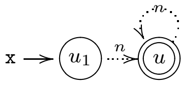
</td></tr>
<tr><td>Statement</td><td colspan="4"><code>y = x</code></td></tr>
<tr><td>Relation-transfer formulas</td><td colspan="4">$\tau_{x,\mathtt{y=x}}(v) = x(v)$
$\tau_{y,\mathtt{y=x}}(v) = x(v)$
$\tau_{n,\mathtt{y=x}}(v_1, v_2) = n(v_1, v_2)$
$\tau_{eq,\mathtt{y=x}}(v_1, v_2) = eq(v_1, v_2)$</td></tr>
<tr><td>Structure after</td><td><b>unary rels.</b><table><tr><th>indiv.</th><th>x</th><th>y</th></tr><tr><td>u_1</td><td>1</td><td>1</td></tr><tr><td>u</td><td>0</td><td>0</td></tr></table></td><td><b>binary rels.</b><table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table></td><td><table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table></td><td>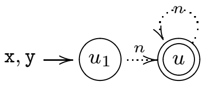
</td></tr>
</table>

Fig. 9. The relation-transfer formulas for x, y, and n express a transformation on logical structures that corresponds to the semantics of $\mathtt{y} = \mathtt{x}$.

Abstract interpretation collects a set of 3-valued structures at each program point. It can be implemented as an iterative procedure that finds the least fixed point of a certain set of equations [Sagiv et al. 2002]. (It is important to understand that although the analysis framework is based on logic, it is model theoretic, not proof theoretic: the abstract interpretation collects sets of 3-valued logical structures—i.e., abstracted models; its actions do not rely on deduction or theorem proving.) When the fixed point is reached, the structures that have been collected at program point P describe a superset of all the execution states that can occur at P. To determine whether a property always holds at P, one checks whether it holds in all of the structures that were collected there.

Fig. 9 illustrates the abstract execution of the statement $\mathtt{y} = \mathtt{x}$ on a 3-valued logical structure that represents concrete lists of length 2 or more. Instrumentation relations and relation-maintenance formulas have been omitted from the figure. The abstract execution of the statement $\mathtt{y} = \mathtt{x}$ is revisited in Ex. 3.2 of §3, which discusses relation-maintenance formulas.

Other Operations on Logical Structures. *focus*$[\varphi]$ is a heuristic that elaborates a 3-valued structure—causing it to be replaced by a collection of more precise structures that, taken together, represent the same set of concrete stores;[3] the criterion for refinement is to ensure that the formula φ evaluates to a definite value for all complete assignments to φ's free variables. The operation thus brings φ "into focus".

By invoking *focus* before applying each structure transformer, focusing is used to reduce the number of indefinite values that arise when relation-transfer and relation-maintenance formulas are evaluated in 3-valued structures. The *focus* formulas aim to *sharpen* the values of relations when applied to the individuals that are affected by the transformer. (This often involves the *materialization* of a concrete individual

maintenance formulas, which are the main subject of this paper. The term "relation-maintenance formula" emphasizes the connection to work in the database community on *view maintenance* (see §7). ("View updating" is something different: an update is made to the value of a view relation and changes are propagated back to the base relations.)

[3] This operation can be viewed as a partial concretization.

out of a summary individual.) For program-analysis applications, it was proposed in [Sagiv et al. 2002] that for a statement of the form *lhs* = *rhs*, the focus formula should identify the memory cells that correspond to the L-value of *lhs* and the R-value of *rhs*. This ensures that the application of an abstract transformer performs a *strong update* of the values of core relations that represent pointer variables and fields that are updated by the statement, i.e., does not set those values to 1/2.

Not all logical structures represent admissible stores. To exclude structures that do not, we impose integrity constraints. For instance, relation $x(v)$ of Fig. 3 captures whether pointer variable `x` points to memory cell v; x would be given the attribute "unique", which imposes the integrity constraint that x can hold for at most one individual in any structure: $\forall v_1, v_2 \colon x(v_1) \wedge x(v_2) \Rightarrow v_1 = v_2$. This formula evaluates to 1 in any 2-valued logical structure that corresponds to an admissible store. Integrity constraints contribute to the concretization function (γ) for our abstraction [Yorsh et al. 2007]. Integrity constraints are enforced by *coerce*, a clean-up operation that may "sharpen" a 3-valued logical structure by setting an indefinite value (1/2) to a definite value (0 or 1), or discard a structure entirely if an integrity constraint is definitely violated by the structure (e.g., if it cannot represent any admissible store). To help prevent an analysis from losing precision, *coerce* is applied at certain steps of the algorithm, e.g., after the application of an abstract transformer.

In addition, most of the operations described in this section are not constrained to manipulate 3-valued structures that are images of canonical abstraction; they rely on the Embedding Theorem, which applies to any pair of structures for which one can be embedded into the other. Thus, it is not necessary to perform canonical abstraction after the application of each abstract structure transformer. To ensure that abstract interpretation terminates, it is only necessary that canonical abstraction be applied somewhere in each loop, e.g., at the target of each backedge in the CFG.

3. THE PROBLEM: MAINTAINING INSTRUMENTATION RELATIONS

The execution of a statement *st* transforms a logical structure S, which represents a store that arises just before *st*, into a new structure S', which represents the corresponding store just after *st* executes. The structure that consists of just the core relations of S' is called a *proto-structure*, denoted by S'_{proto}. The creation of S'_{proto} from S, denoted by $S'_{proto} := [\![st]\!]_3(S)$, can be expressed as

$$\begin{aligned} &\text{for each } c \in \mathcal{C} \text{ and } u_1, \ldots, u_k \in U^S, \\ &\quad \iota^{S'_{proto}}(c)(u_1, \ldots, u_k) := [\![\tau_{c,st}(v_1, \ldots, v_k)]\!]_3^S([v_1 \mapsto u_1, \ldots, v_k \mapsto u_k]). \end{aligned} \tag{1}$$

In general, if we compare the various relations of S'_{proto} with those of S, some tuples will have been added and others will have been deleted.

We now come to the crux of the matter: Suppose that ψ_p defines instrumentation relation p; how should the static-analysis engine obtain the value of p in S'?

An instrumentation relation whose defining formula is expressed solely in terms of core relations is said to be in *core normal form*. Because there are no circular dependences, an instrumentation relation's defining formula can always be put in core normal form by repeated substitution until only core relations remain. When

ψ_p is in core normal form, or has been converted to core normal form, it is possible to determine the value of each instrumentation relation p by evaluating ψ_p in structure S'_{proto}:

$$\begin{array}{l}\text{for each } u_1, \ldots, u_k \in U^S, \\ \quad \iota^{S'}(p)(u_1, \ldots, u_k) := [\![\psi_p(v_1, \ldots, v_k)]\!]_3^{S'_{proto}}([v_1 \mapsto u_1, \ldots, v_k \mapsto u_k]).\end{array} \tag{2}$$

Thus, in principle it is possible to maintain the values of instrumentation relations via Eqn. (2). In practice, however, this approach does not work very well. As observed elsewhere [Sagiv et al. 2002], when working in 3-valued logic, it is usually possible to retain more precision by defining a special *instrumentation-relation maintenance formula*, $\mu_{p,st}(v_1, \ldots, v_k)$, and evaluating $\mu_{p,st}(v_1, \ldots, v_k)$ in structure S:

$$\begin{array}{l}\text{for each } u_1, \ldots, u_k \in U^S, \\ \quad \iota^{S'}(p)(u_1, \ldots, u_k) := [\![\mu_{p,st}(v_1, \ldots, v_k)]\!]_3^{S}([v_1 \mapsto u_1, \ldots, v_k \mapsto u_k]).\end{array} \tag{3}$$

The advantage of the relation-maintenance approach is that the results of program analysis can be more accurate. In 3-valued logic, when $\mu_{p,st}$ is defined appropriately, the relation-maintenance strategy can generate a definite value (0 or 1) when the evaluation of ψ_p on S'_{proto} generates the indefinite value 1/2.

To ensure that an analysis is conservative, however, one must also show that the following property holds:

Definition 3.1. Suppose that p is an instrumentation relation defined by formula ψ_p. Relation-maintenance formula $\mu_{p,st}$ *maintains* p *correctly for statement* st if, for all $S \in \mathcal{S}_2[\mathcal{R}]$ and all Z, $[\![\mu_{p,st}]\!]_2^S(Z) = [\![\psi_p]\!]_2^{[\![st]\!]_2(S)}(Z)$.

For an instrumentation relation in core normal form, it is always possible to provide a relation-maintenance formula that satisfies Defn. 3.1 by defining $\mu_{p,st}$ as

$$\mu_{p,st} \overset{\text{def}}{=} \psi_p[c \hookleftarrow \tau_{c,st} \mid c \in \mathcal{C}], \tag{4}$$

where $\varphi[q \hookleftarrow \varphi']$ denotes the formula obtained from φ by replacing each relation occurrence $q(w_1, \ldots, w_k)$ by $\varphi'\{w_1, \ldots, w_k\}$, and $\varphi'\{w_1, \ldots, w_k\}$ denotes the formula obtained from $\varphi'(v_1, \ldots, v_k)$ by replacing each free occurrence of variable v_i by w_i.

The formula $\mu_{p,st}$ defined in Eqn. (4) maintains p correctly for statement st because, by the 2-valued version of Eqn. (1), $[\![\tau_{c,st}]\!]_2^S(Z) = [\![c]\!]_2^{S'_{proto}}(Z)$; consequently, when $\mu_{p,st}$ of Eqn. (4) is evaluated in structure S, the use of $\tau_{c,st}$ in place of c is equivalent to using the value of c when ψ_p is evaluated in S'_{proto}; i.e., for all Z, $[\![\psi_p[c \hookleftarrow \tau_{c,st} \mid c \in \mathcal{C}]]\!]_2^S(Z) = [\![\psi_p]\!]_2^{S'_{proto}}(Z)$. However—and this is precisely the drawback of using Eqn. (4) to obtain the $\mu_{p,st}$—the steps of evaluating $[\![\psi_p[c \hookleftarrow \tau_{c,st} \mid c \in \mathcal{C}]]\!]_2^S(Z)$ *mimic exactly* those of evaluating $[\![\psi_p]\!]_2^{S'_{proto}}(Z)$. Consequently, when we pass to 3-valued logic, for all Z, $[\![\psi_p[c \hookleftarrow \tau_{c,st} \mid c \in \mathcal{C}]]\!]_3^S(Z)$ yields exactly the same value as $[\![\psi_p]\!]_3^{S'_{proto}}(Z)$ (i.e., as evaluating Eqn. (2)). Thus, although $\mu_{p,st}$ that satisfy Defn. 3.1 can be obtained automatically via Eqn. (4), this approach does not provide a satisfactory solution to the relation-maintenance problem.

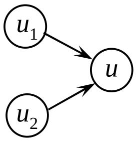

Fig. 10. A store in which u is shared; i.e., $is_n(u) = 1$.

Example 3.2. Eqn. (5) shows the defining formula for the instrumentation relation is_n ("is-shared using n fields"),

$$is_n(v) \stackrel{\text{def}}{=} \exists\, v_1, v_2 \colon n(v_1, v) \wedge n(v_2, v) \wedge v_1 \neq v_2, \tag{5}$$

which captures whether a memory cell is pointed to by two or more pointer fields of memory cells, e.g., see Fig. 10.

Fig. 11 illustrates how execution of the statement $\texttt{y} = \texttt{x}$ causes the value of is_n to lose precision when its relation-maintenance formula is created according to Eqn. (4). The initial 3-valued structure represents all singly-linked lists of length 2 or more in which all memory cells are unshared. Because execution of $\texttt{y} = \texttt{x}$ does not change the value of core relation n, $\tau_{n,\texttt{y}=\texttt{x}}(v_1, v_2)$ is $n(v_1, v_2)$, and hence the formula $\mu_{is_n,\texttt{y}=\texttt{x}}(v)$ created according to Eqn. (4) is $\exists\, v_1, v_2 \colon n(v_1, v) \wedge n(v_2, v) \wedge v_1 \neq v_2$. As shown in Fig. 11, the structure created using this maintenance formula is not as precise as we would like. In particular, $is_n(u) = 1/2$, which means that u can represent a shared cell. Thus, the final 3-valued structure also represents certain cyclic linked lists, such as

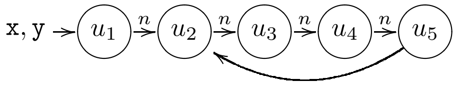

□

This sort of imprecision can usually be avoided by devising better relation-maintenance formulas. For instance, when $\mu_{is_n,\texttt{y}=\texttt{x}}(v)$ is defined to be the formula $is_n(v)$—meaning that $\texttt{y} = \texttt{x}$ does not change the value of $is_n(v)$—the imprecision illustrated in Fig. 11 is avoided (see Fig. 12). Hand-crafted relation-maintenance formulas for a variety of instrumentation relations are given in [Sagiv et al. 2002; Lev-Ami and Sagiv 2000; TVLA]; however, those formulas were created by *ad hoc* methods.

To sum up, prior to the work presented in this paper, the user needed to supply a formula $\mu_{p,st}$ for each instrumentation relation p and each statement *st*. In effect, the user needed to write down two *separate* characterizations of each instrumentation relation p: (i) ψ_p, which defines p directly; and (ii) $\mu_{p,st}$, which specifies how execution of each kind of statement in the language affects p. Moreover, it was the user's responsibility to ensure that the two characterizations were mutually consistent. In contrast, with the method for automatically creating relation-maintenance formulas presented in §4 and §5, the user's responsibility is dramatically reduced: he only needs to give a *single* characterization of each instrumentation relation p—namely, by defining ψ_p. (In separate work, we have developed ways to use inductive logic programming to *discover* an appropriate set of instrumentation relations that define a suitable abstraction for checking whether a given program has a given property [Loginov et al. 2005; Loginov 2006; Loginov et al. 2007].)

<table>
<tr><td>Structure before</td><td>
<b>unary rels.</b>
<table><tr><th>indiv.</th><th>x</th><th>y</th><th>is_n</th></tr><tr><td>u_1</td><td>1</td><td>0</td><td>0</td></tr><tr><td>u</td><td>0</td><td>0</td><td>0</td></tr></table>
<b>binary rels.</b>
<table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
<table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
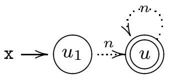

</td></tr>
<tr><td>Statement</td><td>y = x</td></tr>
<tr><td>Relation-transfer formulas</td><td>$\tau_{x,\mathtt{y=x}}(v) = x(v)$
$\tau_{y,\mathtt{y=x}}(v) = x(v)$
$\tau_{n,\mathtt{y=x}}(v_1, v_2) = n(v_1, v_2)$
$\tau_{eq,\mathtt{y=x}}(v_1, v_2) = eq(v_1, v_2)$</td></tr>
<tr><td>Relation-maintenance formula</td><td>$\mu_{is_n,\mathtt{y=x}}(v) = \exists v_1, v_2 : n(v_1, v) \wedge n(v_2, v) \wedge v_1 \neq v_2$</td></tr>
<tr><td>Structure after</td><td>
<b>unary rels.</b>
<table><tr><th>indiv.</th><th>x</th><th>y</th><th>is_n</th></tr><tr><td>u_1</td><td>1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>0</td><td>1/2</td></tr></table>
<b>binary rels.</b>
<table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
<table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
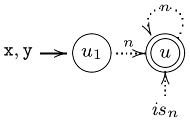

</td></tr>
</table>

Fig. 11. An illustration of the loss of precision in the value of is_n when its relation-maintenance formula is defined by $\exists v_1, v_2 : n(v_1, v) \wedge n(v_2, v) \wedge v_1 \neq v_2$. The use of this relation-maintenance formula causes a structure to be created in which the individual u may represent a shared memory cell.

<table>
<tr><td>Structure before</td><td>
<b>unary rels.</b>
<table><tr><th>indiv.</th><th>x</th><th>y</th><th>is_n</th></tr><tr><td>u_1</td><td>1</td><td>0</td><td>0</td></tr><tr><td>u</td><td>0</td><td>0</td><td>0</td></tr></table>
<b>binary rels.</b>
<table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
<table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
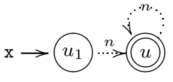

</td></tr>
<tr><td>Statement</td><td>y = x</td></tr>
<tr><td>Relation-transfer formulas</td><td>$\tau_{x,\mathtt{y=x}}(v) = x(v)$
$\tau_{y,\mathtt{y=x}}(v) = x(v)$
$\tau_{n,\mathtt{y=x}}(v_1, v_2) = n(v_1, v_2)$
$\tau_{eq,\mathtt{y=x}}(v_1, v_2) = eq(v_1, v_2)$</td></tr>
<tr><td>Relation-maintenance formula</td><td>$\mu_{is_n,\mathtt{y=x}}(v) = is_n(v)$</td></tr>
<tr><td>Structure after</td><td>
<b>unary rels.</b>
<table><tr><th>indiv.</th><th>x</th><th>y</th><th>is_n</th></tr><tr><td>u_1</td><td>1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>0</td><td>0</td></tr></table>
<b>binary rels.</b>
<table><tr><th>n</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>0</td><td>1/2</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
<table><tr><th>eq</th><th>u_1</th><th>u</th></tr><tr><td>u_1</td><td>1</td><td>0</td></tr><tr><td>u</td><td>0</td><td>1/2</td></tr></table>
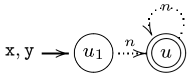

</td></tr>
</table>

Fig. 12. Example showing how the imprecision that was illustrated in Fig. 11 is avoided with the relation-maintenance formula $\mu_{is_n,\mathtt{y=x}}(v) = is_n(v)$. (Ex. 4.1 shows how this is generated automatically.)

4. A FINITE-DIFFERENCING SCHEME FOR 2-VALUED (AND 3-VALUED) FIRST-ORDER LOGIC

This section presents a finite-differencing scheme for creating relation-maintenance formulas. The discussion will be couched in terms of 2-valued logic; however, by the Embedding Theorem (Theorem 2.6, [Sagiv et al. 2002, Theorem 4.9]), the relation-maintenance formulas that we derive provide sound results when interpreted in 3-valued logic. In 3-valued logic, as demonstrated in Fig. 12 (and discussed further in Ex. 4.1), the resulting formula can lead to a strictly more precise result than merely reevaluating an instrumentation relation's defining formula.

A relation-maintenance formula $\mu_{p,st}$ for $p \in \mathcal{I}$ is defined in terms of two finite-differencing operators, denoted by $\Delta^-_{st}[\cdot]$ and $\Delta^+_{st}[\cdot]$, which capture the negative and positive changes, respectively, that execution of structure transformer st induces in an instrumentation relation's value. The formula $\mu_{p,st}$ is created by combining p

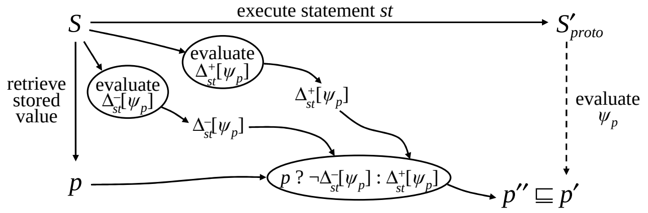

Fig. 13. How to maintain the value of ψ_p in 3-valued logic in response to changes in the values of core relations caused by the execution of structure transformer *st*.

φ	$\Delta^+_{st}[\varphi]$	$\Delta^-_{st}[\varphi]$
$\mathbf{1}$	$\mathbf{0}$	$\mathbf{0}$
$\mathbf{0}$	$\mathbf{0}$	$\mathbf{0}$
$p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p \,?\, \neg\delta^-_{p,st} : \delta^+_{p,st}$	$(\delta^+_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}$	$(\delta^-_{p,st} \wedge p)\{w_1,\ldots,w_k\}$
$p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p \vee \delta_{p,st}$ or $\delta_{p,st} \vee p$	$(\delta_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}$	$\mathbf{0}$
$p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p \wedge \delta_{p,st}$ or $\delta_{p,st} \wedge p$	$\mathbf{0}$	$(\neg\delta_{p,st} \wedge p)\{w_1,\ldots,w_k\}$
$p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, but $\tau_{p,st}$ is not of the above forms	$(\tau_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}$	$(p \wedge \neg\tau_{p,st})\{w_1,\ldots,w_k\}$
$p(w_1,\ldots,w_k)$, $p \in \mathcal{I}$	$((\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge \neg p)\{w_1,\ldots,w_k\}$ if$\psi_p \equiv \exists\, v\colon \varphi_1$ $\Delta^+_{st}[\psi_p]\{w_1,\ldots,w_k\}$ otherwise	$((\exists\, v\colon \Delta^-_{st}[\varphi_1]) \wedge p)\{w_1,\ldots,w_k\}$ if$\psi_p \equiv \forall\, v\colon \varphi_1$ $\Delta^-_{st}[\psi_p]\{w_1,\ldots,w_k\}$ otherwise
$\neg\varphi_1$	$\Delta^-_{st}[\varphi_1]$	$\Delta^+_{st}[\varphi_1]$
$\varphi_1 \vee \varphi_2$	$(\Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2) \vee (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2])$	$(\Delta^-_{st}[\varphi_1] \wedge \neg\mathbf{F}_{st}[\varphi_2]) \vee (\neg\mathbf{F}_{st}[\varphi_1] \wedge \Delta^-_{st}[\varphi_2])$
$\varphi_1 \wedge \varphi_2$	$(\Delta^+_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_2]) \vee (\mathbf{F}_{st}[\varphi_1] \wedge \Delta^+_{st}[\varphi_2])$	$(\Delta^-_{st}[\varphi_1] \wedge \varphi_2) \vee (\varphi_1 \wedge \Delta^-_{st}[\varphi_2])$
$\exists\, v\colon \varphi_1$	$(\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v\colon \varphi_1)$	$(\exists\, v\colon \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v\colon \mathbf{F}_{st}[\varphi_1])$
$\forall\, v\colon \varphi_1$	$(\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge (\forall\, v\colon \mathbf{F}_{st}[\varphi_1])$	$(\exists\, v\colon \Delta^-_{st}[\varphi_1]) \wedge (\forall\, v\colon \varphi_1)$

Fig. 14. Finite-difference formulas for first-order formulas.

with $\Delta^-_{st}[\psi_p]$ and $\Delta^+_{st}[\psi_p]$ as follows: $\mu_{p,st} = p \,?\, \neg\Delta^-_{st}[\psi_p] : \Delta^+_{st}[\psi_p]$. The formula $\mu_{p,st}$ states the conditions under which the new value of p (i.e., its value in S') is 1. These conditions are specified in terms of the old values of p, $\Delta^-_{st}[\psi_p]$, and $\Delta^+_{st}[\psi_p]$ (i.e., their values in S). The formula $\mu_{p,st}$ states that if p's old value is 1, then its new value is 1 unless there is a negative change; if p's old value is 0, then its new value is 1 if there is a positive change.

Fig. 13 depicts how the static-analysis engine evaluates $\Delta^-_{st}[\psi_p]$ and $\Delta^+_{st}[\psi_p]$ in S and combines these values with the old value p to obtain the desired new value p''. The operators $\Delta^-_{st}[\cdot]$ and $\Delta^+_{st}[\cdot]$ are defined recursively, as shown in Fig. 14. The definitions in Fig. 14 make use of the operator $\mathbf{F}_{st}[\varphi]$ (standing for "Future"), defined as follows:

$$\mathbf{F}_{st}[\varphi] \overset{\text{def}}{=} \varphi \,?\, \neg\Delta^-_{st}[\varphi] : \Delta^+_{st}[\varphi]. \qquad (6)$$

Thus, maintenance formula $\mu_{p,st}$ can also be expressed as $\mu_{p,st} = \mathbf{F}_{st}[p]$.

Formula (6) and Fig. 14 define a syntax-directed translation scheme that can be implemented via a recursive walk over a formula φ. The operators $\Delta^-_{st}[\cdot]$ and $\Delta^+_{st}[\cdot]$ are mutually recursive. For instance, $\Delta^+_{st}[\neg\varphi_1] = \Delta^-_{st}[\varphi_1]$ and $\Delta^-_{st}[\neg\varphi_1] = \Delta^+_{st}[\varphi_1]$. Moreover, each occurrence of $\mathbf{F}_{st}[\varphi_i]$ contains additional occurrences of $\Delta^-_{st}[\varphi_i]$ and $\Delta^+_{st}[\varphi_i]$.

Note how $\Delta^-_{st}[\cdot]$ and $\Delta^+_{st}[\cdot]$ for $\varphi_1 \vee \varphi_2$ and $\varphi_1 \wedge \varphi_2$ exhibit the "convolution" pattern characteristic of differentiation, finite differencing, and divided differencing.

Continuing the analogy with differentiation, it helps to bear in mind that the "independent variables" are the core relations—which are being changed by the $\tau_{c,st}$ formulas; the dependent variable is the value of φ. A formal justification of Fig. 14 is stated later (Theorem 4.3 and Cor. 4.4); here we merely explain informally a few of the cases from Fig. 14:

$\Delta^+_{st}[\mathbf{1}] = \mathbf{0}$, $\Delta^-_{st}[\mathbf{1}] = \mathbf{0}$. The value of atomic formula $\mathbf{1}$ does not depend on any core relations; hence its value is unaffected by changes in them.

$\Delta^-_{st}[\varphi_1 \wedge \varphi_2] = (\Delta^-_{st}[\varphi_1] \wedge \varphi_2) \vee (\varphi_1 \wedge \Delta^-_{st}[\varphi_2])$. Tuples of individuals removed from $\varphi_1 \wedge \varphi_2$ are either tuples of individuals removed from φ_1 for which φ_2 also holds (i.e., $(\Delta^-_{st}[\varphi_1] \wedge \varphi_2)$), or they are tuples of individuals removed from φ_2 for which φ_1 also holds, (i.e., $(\varphi_1 \wedge \Delta^-_{st}[\varphi_2])$.

$\Delta^+_{st}[\exists\, v\colon \varphi_1] = (\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v\colon \varphi_1)$. For $\exists\, v\colon \varphi_1$ to change value from 0 to 1, there must be at least one individual for which φ_1 changes value from 0 to 1 (i.e., $\exists\, v\colon \Delta^+_{st}[\varphi_1]$ holds), and $\exists\, v\colon \varphi_1$ must not already hold (i.e., $\neg(\exists\, v\colon \varphi_1)$ holds).

$\Delta^+_{st}[p(w_1, \ldots, w_k)] = (\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge \neg p$, if $p \in \mathcal{I}$ and $\psi_p \equiv \exists\, v\colon \varphi_1$. This is similar to the previous case, except that the term to ensure that $\exists\, v\colon \varphi_1$ does not already hold (i.e., $\neg(\exists\, v\colon \varphi_1)$) is replaced by the formula $\neg p$. Thus, when $(\exists\, v\colon \Delta^+_{st}[\varphi_1]) \wedge \neg p$ is evaluated, the stored value of $\exists\, v\colon \varphi_1$, i.e., p, will be used instead of the value obtained by reevaluating $\exists\, v\colon \varphi_1$.

$\Delta^+_{st}[p(w_1, \ldots, w_k)] = \Delta^+_{st}[\psi_p\{w_1, \ldots, w_k\}]$, if $p \in \mathcal{I}$ and $\psi_p \not\equiv \exists\, v\colon \varphi_1$. To characterize the positive changes to p, apply Δ^+_{st} to p's defining formula ψ_p.

One special case is also worth noting: $\Delta^+_{st}[v_1 = v_2] = \mathbf{0}$ and $\Delta^-_{st}[v_1 = v_2] = \mathbf{0}$ because the value of the atomic formula $(v_1 = v_2)$ (shorthand for $eq(v_1, v_2)$) does not depend on any core relations; hence, its value is unaffected by changes in them.[4]

Example 4.1. Consider the instrumentation relation is_n ("is-shared using `n` fields"), defined in Eqn. (5). Fig. 15 shows the formulas obtained for $\Delta^+_{st}[is_n(v)]$ and $\Delta^-_{st}[is_n(v)]$.

For a particular statement, the formulas in Fig. 15 can usually be simplified. For instance, for `y = x`, the relation-transfer formula $\tau_{n,\mathtt{y=x}}(v_1, v_2)$ is $n(v_1, v_2)$; see Fig. 11. Thus, by Fig. 14, the formulas for $\Delta^-_{\mathtt{y=x}}[n(v_1, v)]$ and $\Delta^+_{\mathtt{y=x}}[n(v_1, v)]$ are both $n(v_1, v) \wedge \neg n(v_1, v)$, which simplifies to $\mathbf{0}$. (In our implementation, simplifications are performed greedily at formula-construction time; e.g., the constructor for $\wedge$ rewrites $\mathbf{0} \wedge p$ to $\mathbf{0}$, $\mathbf{1} \wedge p$ to p, $p \wedge \neg p$ to $\mathbf{0}$, etc.) The formulas in Fig. 15 simplify to $\Delta^+_{\mathtt{y=x}}[is_n(v)] = \mathbf{0}$ and $\Delta^-_{\mathtt{y=x}}[is_n(v)] = \mathbf{0}$. Consequently,

[4] We avoid issues that could arise due to changes in a structure's universe of individuals by modeling storage allocation and deallocation via a free-storage list. We describe our solution in more detail at the end of this section.

$\Delta^+_{st}[is_n(v)]$	$=$	$\left(\exists\, v_1, v_2\colon \left(\begin{array}{rl} & (\Delta^+_{st}[n(v_1,v)] \land \mathbf{F}_{st}[n(v_2,v)]) \\ \lor & (\mathbf{F}_{st}[n(v_1,v)] \land \Delta^+_{st}[n(v_2,v)]) \end{array}\right) \land v_1 \neq v_2\right) \land \neg is_n(v)$
$\Delta^-_{st}[is_n(v)]$	$=$	$\left\{\begin{array}{l} \left(\exists\, v_1, v_2\colon \left(\begin{array}{rl} & (\Delta^-_{st}[n(v_1,v)] \land n(v_2,v)) \\ \lor & (n(v_1,v) \land \Delta^-_{st}[n(v_2,v)]) \end{array}\right) \land v_1 \neq v_2\right) \\ \land \\ \neg\left(\exists\, v_1, v_2\colon \left(\begin{array}{l} (n(v_1,v) \land n(v_2,v) \land v_1 \neq v_2) \\ ?\ \neg\left(\left(\begin{array}{rl} & (\Delta^-_{st}[n(v_1,v)] \land n(v_2,v)) \\ \lor & (n(v_1,v) \land \Delta^-_{st}[n(v_2,v)]) \end{array}\right) \land v_1 \neq v_2\right) \\ :\ \left(\begin{array}{rl} & (\Delta^+_{st}[n(v_1,v)] \land \mathbf{F}_{st}[n(v_2,v)]) \\ \lor & (\mathbf{F}_{st}[n(v_1,v)] \land \Delta^+_{st}[n(v_2,v)]) \end{array}\right) \land v_1 \neq v_2 \end{array}\right)\right) \end{array}\right.$

Fig. 15. Finite-difference formulas for the instrumentation relation $is_n(v)$.

$\mu_{is_n,\mathtt{y=x}}(v) = \mathbf{F}_{\mathtt{y=x}}[is_n(v)] = is_n(v)\ ?\ \neg\mathbf{0} : \mathbf{0} = is_n(v)$. As shown in Fig. 12, this definition of $\mu_{is_n,\mathtt{y=x}}(v)$ avoids the imprecision that was illustrated in Ex. 3.2. □

Correctness of the Finite-Differencing Scheme

The correctness of the finite-differencing scheme given above is established with the help of the following lemma:

Lemma 4.2. *For every formula φ, φ_1, φ_2 and structure transformer st, the following properties hold:*[5]

(i). $\Delta^+_{st}[\varphi] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi] \land \neg\varphi$

(ii). $\Delta^-_{st}[\varphi] \overset{\text{meta}}{\Longleftrightarrow} \varphi \land \neg\boldsymbol{F}_{st}[\varphi]$

(iii). (a). $\boldsymbol{F}_{st}[\neg\varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \neg\boldsymbol{F}_{st}[\varphi_1]$

(b). $\boldsymbol{F}_{st}[\varphi_1 \lor \varphi_2] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi_1] \lor \boldsymbol{F}_{st}[\varphi_2]$

(c). $\boldsymbol{F}_{st}[\varphi_1 \land \varphi_2] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi_1] \land \boldsymbol{F}_{st}[\varphi_2]$

(d). $\boldsymbol{F}_{st}[\exists\, v\colon \varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \exists\, v\colon \boldsymbol{F}_{st}[\varphi_1]$

(e). $\boldsymbol{F}_{st}[\forall\, v\colon \varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \forall\, v\colon \boldsymbol{F}_{st}[\varphi_1]$

Proof. *See App. A.* □

Lemma 4.2 shows that for structures in $\mathcal{S}_2$, $\Delta^+_{st}[\varphi]$ specifies the tuples that are not in the relation defined by φ, but need to be added in response to the execution of *st*, and that $\Delta^-_{st}[\varphi]$ specifies the tuples that are in the relation defined by φ that need to be removed. This lemma is used in the proof of the following theorem, which ensures the correctness of the finite-differencing transformation given in Fig. 14:

Theorem 4.3. *Let S be a structure in $\mathcal{S}_2$, and let S'_{proto} be the proto-structure obtained from S using structure transformer st. Let S' be the structure obtained by using S'_{proto} as the first approximation to S' and then filling in instrumentation relations in a topological ordering of the dependences among them: for each arity-k relation $p \in \mathcal{I}$, $\iota^{S'}(p)$ is obtained by evaluating $[\![\psi_p(v_1,\ldots,v_k)]\!]^{S'}_2([v_1 \mapsto u'_1,\ldots,v_k \mapsto u'_k])$ for all tuples $(u'_1,\ldots,u'_k) \in (U^{S'})^k$.*

[5] To simplify the presentation, we use $lhs \overset{\text{meta}}{\Longleftrightarrow} rhs$ and $lhs \overset{\text{meta}}{\Longrightarrow} rhs$ as shorthands for $[\![lhs]\!]^S_2(Z) = [\![rhs]\!]^S_2(Z)$ and $[\![lhs]\!]^S_2(Z) \leq [\![rhs]\!]^S_2(Z)$, respectively, for any $S \in \mathcal{S}_2$ and assignment Z that is complete for *lhs* and *rhs*.

φ	$\mathbf{F}_{st}[\varphi]$
$\mathbf{1}$	$\mathbf{1}$
$\mathbf{0}$	$\mathbf{0}$
$p(w_1, \ldots, w_k), p \in \mathcal{C}$	$\tau_{p,st}\{w_1, \ldots, w_k\}$
$p(w_1, \ldots, w_k), p \in \mathcal{I}$	$p(w_1, \ldots, w_k)\ ?\ \neg\Delta^-_{st}[p(w_1, \ldots, w_k)] : \Delta^+_{st}[p(w_1, \ldots, w_k)]$
$\neg\varphi_1$	$\neg\mathbf{F}_{st}[\varphi_1]$
$\varphi_1 \vee \varphi_2$	$\mathbf{F}_{st}[\varphi_1] \vee \mathbf{F}_{st}[\varphi_2]$
$\varphi_1 \wedge \varphi_2$	$\mathbf{F}_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_2]$
$\exists\, v \colon \varphi_1$	$\exists\, v \colon \mathbf{F}_{st}[\varphi_1]$
$\forall\, v \colon \varphi_1$	$\forall\, v \colon \mathbf{F}_{st}[\varphi_1]$

Fig. 16. Optimized formulas for the operator $\mathbf{F}_{st}[\varphi]$.

Then for every formula $\varphi(v_1, \ldots, v_k)$ and complete assignment Z for $\varphi(v_1, \ldots, v_k)$, $[\![\boldsymbol{F}_{st}[\varphi(v_1, \ldots, v_k)]]\!]^S_2(Z) = [\![\varphi(v_1, \ldots, v_k)]\!]^{S'}_2(Z)$.

Proof. *See App. A.* □

For structures in $\mathcal{S}_3$, the soundness of the finite-differencing transformation given in Fig. 14 follows from Theorem 4.3 by the Embedding Theorem (Theorem 2.6):

Corollary 4.4. *Let $S, S' \in \mathcal{S}_2$ be defined as in Theorem 4.3. Let $S^\sharp \in \mathcal{S}_3$ be such that $f \colon U^S \to U^{S^\sharp}$ embeds S in $S^\sharp$, i.e., $S \sqsubseteq^f S^\sharp$. Then for every formula $\varphi(v_1, \ldots, v_k)$ and complete assignment Z for $\varphi(v_1, \ldots, v_k)$, $[\![\boldsymbol{F}_{st}[\varphi(v_1, \ldots, v_k)]]\!]^{S^\sharp}_3(f \circ Z) \sqsupseteq [\![\varphi(v_1, \ldots, v_k)]\!]^{S'}_2(Z)$.*

Optimized Formulas for $\mathbf{F}_{st}[\varphi]$

For a non-atomic formula φ, the operator $\mathbf{F}_{st}[\varphi]$ defined in Formula (6) introduces a copy of φ, because it has no way, in general, to refer to a relation that holds the stored value of φ. The reevaluation inherent in the version of $\mathbf{F}_{st}[\cdot]$ from Formula (6) may cause a substantial loss of precision. One way to retain higher precision is to propagate $\mathbf{F}_{st}[\cdot]$ into the subformulas of φ, down to the level of atomic formulas—either core-relation symbols or instrumentation-relation symbols—as shown in Fig. 16.

Suppose that φ' is the result of $\mathbf{F}_{st}[\varphi]$ by the method of Fig. 16. An evaluation of φ' will evaluate (copies of) the operators of φ, down to the level of each atomic subformula $p(w_1, \ldots, w_k)$ in φ. At that level, if $p \in \mathcal{I}$, $\mathbf{F}_{st}[\cdot]$ will have introduced an occurrence of p in φ':

$$\mathbf{F}_{st}[p(w_1, \ldots, w_k)] \stackrel{\text{def}}{=} p(w_1, \ldots, w_k)\ ?\ \neg\Delta^-_{st}[p(w_1, \ldots, w_k)] : \Delta^+_{st}[p(w_1, \ldots, w_k)]. \quad (7)$$

The occurrence of p in the test refers to the stored ("pre-state") value of instrumentation-relation p; consequently, the stored tuples of relation p will be used when evaluating φ'.

Note that $\Delta^+_{st}[p]$ and $\Delta^-_{st}[p]$ in Formula (7) dispatch according to the case for $p \in \mathcal{I}$ in Fig. 14. In particular, because $\mathbf{F}_{st}[\cdot]$ occurs in four of the eight cases for $\Delta^+_{st}[\cdot]$ and $\Delta^-_{st}[\cdot]$ in Fig. 14—i.e., for $\vee$, $\wedge$, $\exists$, and $\forall$—the optimized $\mathbf{F}_{st}[\cdot]$ is invoked recursively on various subterms of ψ_p.

The correctness of the version of $\mathbf{F}_{st}[\cdot]$ defined in Fig. 16 follows from Lemma 4.2.

The method described above also usually produces smaller instrumentation-relation maintenance formulas, and hence creates abstract transformers that generally can be evaluated more quickly. This technique is incorporated into our implementation.

Discussion

Because earlier in the paper we touted the advantages of being able to apply related 2-valued and 3-valued interpretation functions to a single formula, it may seem somewhat inconsistent for us to make use of a transformation-based approach to maintaining instrumentation relations, in lieu of an approach based on overloading. The reason that we use a transformation-based approach is that it gives us an opportunity to simplify the resulting formulas (either on the fly, or in a post-processing phase after finite differencing).

In the context of evaluation in 3-valued logic, simplification is important because even formulas that are tautologies in 2-valued logic may evaluate to 1/2 in 3-valued logic. For instance, $p \vee \neg p$ yields 1/2 when p has the value 1/2, even when p is a nullary relation symbol. The finite-differencing transformation that we implemented uses a formula-minimization procedure for 3-valued logic that we developed [Reps et al. 2002]. The minimization procedure applies to propositional logic; for propositional logic, it is guaranteed to return an answer that captures the formula's "supervaluational meaning" [van Fraassen 1966]. This procedure is used as a subroutine in a heuristic method for minimizing first-order formulas; the method works on a formula bottom-up, applying the propositional minimizer to the body of each non-propositional operator (i.e., each quantifier or transitive-closure operator).

A relation-maintenance formula that has been simplified in this way can sometimes yield a definite value in situations where the evaluation of the unsimplified relation-maintenance formula—or, equivalently, an overloaded evaluation of the relation's defining formula—yields 1/2. (For instance, minimizing $p \vee \neg p$ yields 1, which evaluates to 1 even when p has the value 1/2.) Consequently, the formula-transformation approach to the relation-maintenance problem leads to more precise static-analysis algorithms.

Malloc and Free

In [Sagiv et al. 2002], the modeling of storage-allocation/deallocation operations is carried out with a two-stage structure transformer, the first stage of which changes the number of individuals in the structure. This creates some problems for the finite-differencing approach in establishing appropriate, mutually consistent values for relation tuples that involve the newly allocated individual. Such relation values are needed for the second stage, in which relation-transfer formulas for core relations and relation-maintenance formulas for instrumentation relations are applied in the usual fashion, using Eqns. (1) and (3).

However, there is a simple way to sidestep this problem, which is to model the free-storage list explicitly, making the following substructure part of every 3-valued

p	**IntendedMeaning**	ψ_p
$t_n(v_1, v_2)$	Is v_2 reachable from v_1 along `n` fields?	$n^*(v_1, v_2)$
$r_{n,z}(v)$	Is v reachable from pointer variable `z` along `n` fields?	$\exists\, v_1 : z(v_1) \wedge t_n(v_1, v)$
$c_n(v)$	Is v on a directed cycle of `n` fields?	$\exists\, v_1 : n(v_1, v) \wedge t_n(v, v_1)$

Fig. 17. Defining formulas of some instrumentation relations that depend on **RTC**. (Recall that $n^*(v_1, v_2)$ is a shorthand for $(\mathbf{RTC}\ v_1', v_2' \colon n(v_1', v_2'))(v_1, v_2)$.)

structure:

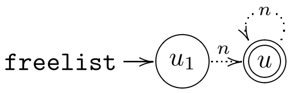

(8)

A `malloc` is modeled by advancing the pointer `freelist` into the list, and returning the memory cell that it formerly pointed to. A `free` is modeled by inserting, at the head of `freelist`'s list, the cell being deallocated. This approach models limits on available storage naturally, while the introduction of one integrity constraint enables it to model unbounded storage.[6]

It is true that the use of structure (8) to model storage-allocation/deallocation operations also causes the number of individuals in a 3-valued structure to change; however, because the new individual is materialized using the usual mechanisms from [Sagiv et al. 2002] (namely, the *focus* and *coerce* operations), values for relation tuples that involve the newly materialized individual will always have safe, mutually consistent values.

5. MAINTENANCE FORMULAS FOR REACHABILITY AND TRANSITIVE CLOSURE

Several instrumentation relations that depend on **RTC** are shown in Fig. 17. Unfortunately, finding a good way to maintain instrumentation relations defined using **RTC** is challenging because the evaluation of a formula that uses the **RTC** operator in a 3-valued structure generally produces many tuples with the value 1/2. This happens because in an *abstracted* binary relation, tuples ("edges") that involve summary individuals often have the value 1/2. (For instance, see the dashed n edges incident on u in Eqn. (8)).) Because the semantics of a tuple (u_1, u_2) computed via **RTC** is defined to be the "max over all paths P from u_1 to u_2 of the minimum value of an edge along P" (see Defns. 2.2 and 2.4), the presence of indefinite edge-tuples often causes the path-tuple computed for a pair (u_1, u_2) to have the value 1/2. Moreover, it is not known, in general, whether it is possible to write a first-order formula (i.e., without using a transitive-closure operator) that specifies how to maintain the closure of a directed graph in response to edge insertions and deletions. Thus, our strategy has been to investigate special cases for classes of instrumentation relations for which first-order maintenance formulas do exist. Whenever these do not apply, the system falls back on safe maintenance formulas (which themselves use **RTC**).

[6] Instead of a free-storage list, one could use a (bounded or unbounded) *set* of memory locations to model storage allocation and deallocation. In that approach, instead of reachability from the pointer `freelist`, a core unary relation would mark free cells, thus distinguishing them from allocated cells that have been leaked.

In this section, we confine ourselves to important special cases for the maintenance of instrumentation relations specified via the **RTC** of a binary formula $\varphi_1(v_1, v_2)$. In §5.1, we consider the case that $\varphi_1(v_1, v_2)$ defines a directed acyclic graph. In §5.2, we consider the case that $\varphi_1(v_1, v_2)$ defines a tree-shaped graph. Finally, in §5.3, we consider the case that $\varphi_1(v_1, v_2)$ defines a *deterministic graph*—i.e., a possibly-cyclic graph, in which every node has outdegree at most one (this class of graphs corresponds to possibly-cyclic linked lists). This collection of techniques allows us to handle most common data structures, such as lists (singly- and doubly-linked; cyclic and acyclic) and trees. The precision of all of these techniques is due to the fact that maintenance of **RTC** after unit-size changes (single-edge additions or deletions)[7] is performed via first-order logical formulas only. However, maintaining **RTC** of an arbitrary directed graph, as well as maintaining **RTC** of restricted classes of graphs with arbitrary-size changes, is not known to be first-order expressible. In such cases, our algorithm returns a formula that uses the **RTC** operator; the evaluation of such a formula may yield more indefinite answers than necessary.

To specify that the maintenance of binary relation $p(v_1, v_2)$ defined as the **RTC** of binary formula $\varphi_1(v_1, v_2)$ should rely on one of the special cases, the user annotates formula φ_1 with attributes. To state that $\varphi_1(v_1, v_2)$ defines a directed acyclic graph, the user gives φ_1 attribute “acyclic”; to state that $\varphi_1(v_1, v_2)$ defines a tree-shaped graph, the user gives φ_1 attribute “tree”; to state that $\varphi_1(v_1, v_2)$ defines a deterministic graph, the user gives φ_1 attribute “function”.

The analysis uses those attributes to generate integrity constraints to be enforced by the coerce operation. For instance, when relation $p(v_1, v_2)$ is defined as the **RTC** of formula $\varphi_1(v_1, v_2)$ that is annotated with the attribute “acyclic”, the analysis generates the following two constraints:

$$\forall\, v_1, v_2 \colon p(v_1, v_2) \land p(v_2, v_1) \;\Rightarrow\; v_1 = v_2$$
$$\forall\, v_1, v_2 \colon p(v_1, v_2) \land v_1 \neq v_2 \;\Rightarrow\; \neg p(v_2, v_1).$$

When $\varphi_1(v_1, v_2)$ is annotated with the attribute “tree”, the analysis generates the above acyclicity constraints, together with constraints that ensure that $\varphi_1(v_1, v_2)$ is an inverse partial function.

Whenever coerce determines that a constraint is (possibly) not satisfied after the application of a transformer, a warning is generated.

5.1 Transitive-Closure Maintenance in Directed Acyclic Graphs

Consider a binary instrumentation relation p, defined by $\psi_p(v_1, v_2) \equiv (\mathbf{RTC}\ v_1', v_2' \colon \varphi_1)(v_1, v_2)$. If the graph defined by φ_1 is acyclic, it is possible to give a first-order formula that maintains p after the addition or deletion of a single φ_1-edge. The method we use is a minor modification of a method for maintaining non-reflexive transitive closure in a directed acyclic graph, due to Dong and Su [Dong and Su 2000].

In the case of an insertion of a single φ_1-edge, the maintenance formula is

$$\mathbf{F}_{st}[p](v_1, v_2) = p(v_1, v_2) \lor (\exists\, v_1', v_2' \colon p(v_1, v_1') \land \Delta^+_{st}[\varphi_1](v_1', v_2') \land p(v_2', v_2)). \quad (9)$$

[7] These techniques can be extended to handle bounded-size addition and deletion sets.

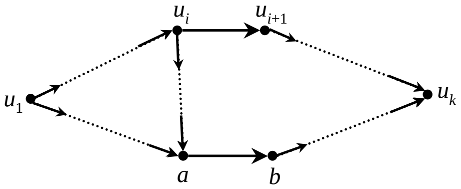

Fig. 18. Edge (a, b) is being deleted; u_i is the last node along path $u_1, \ldots, u_i, u_{i+1}, \ldots, u_k$ from which a is reachable.

The new value of p contains the old tuples of p, as well as those that represent two old paths (i.e., $p(v_1, v_1')$ and $p(v_2', v_2)$) connected with the new φ_1-edge (i.e., $\Delta^+_{st}[\varphi_1](v_1', v_2')$).

The maintenance formula to handle the deletion of a single φ_1-edge is a bit more complicated. We first identify the tuples of p that represent paths that might rely on the edge to be deleted, and thus may need to be removed from p (S stands for *suspicious*):

$$S[p, \varphi_1](v_1, v_2) = \exists\, v_1', v_2' \colon p(v_1, v_1') \wedge \Delta^-_{st}[\varphi_1](v_1', v_2') \wedge p(v_2', v_2).$$

We next collect a set of p-tuples that definitely remain in p (T stands for *trusted*):

$$T[p, \varphi_1](v_1, v_2) = (p(v_1, v_2) \wedge \neg S[p, \varphi_1](v_1, v_2)) \vee \mathbf{F}_{st}[\varphi_1](v_1, v_2). \tag{10}$$

Finally, the maintenance formula for p for a single φ_1-edge deletion is

$$\mathbf{F}_{st}[p](v_1, v_2) = \exists\, v_1', v_2' \colon T[p, \varphi_1](v_1, v_1') \wedge T[p, \varphi_1](v_1', v_2') \wedge T[p, \varphi_1](v_2', v_2). \tag{11}$$

Maintenance formulas (9) and (11) maintain p when two conditions hold: the graph defined by φ_1 is acyclic, and the change to the graph is a single edge addition or deletion (but not both). To see that under these assumptions the maintenance formula for a φ_1-edge deletion is correct, suppose that there is a suspicious tuple $p(u_1, u_k)$, i.e., $S[p, \varphi_1](u_1, u_k) = 1$, but there is a φ_1-path $u_1, \ldots, u_k$ that does not use the deleted φ_1-edge. We need to show that $\mathbf{F}_{st}[p](u_1, u_k)$ has the value 1. Suppose that (a, b) is the φ_1-edge being deleted; because the graph defined by φ_1 is acyclic, there is a $u_i \neq u_k$ that is the last node along path $u_1, \ldots, u_i, u_{i+1}, \ldots, u_k$ from which a is reachable (see Fig. 18). Because $p(u_1, u_i)$ and $p(u_{i+1}, u_k)$ both hold, and because u_i cannot be reachable from b (by acyclicity), neither tuple is suspicious; consequently, $T[p, \varphi_1](u_1, u_i) = 1$ and $T[p, \varphi_1](u_{i+1}, u_k) = 1$. Because (u_i, u_{i+1}) is an edge in the new (as well as the old) graph defined by φ_1, we have $\mathbf{F}_{st}[\varphi_1](u_i, u_{i+1}) = 1$, which means that $T[p, \varphi_1](u_i, u_{i+1}) = 1$ as well, yielding $\mathbf{F}_{st}[p](u_1, u_k) = 1$ by Eqn. (11).

Fig. 19 extends the method for generating relation-maintenance formulas to handle instrumentation relations specified via the **RTC** of a binary formula that defines a directed acyclic graph. Fig. 19 makes use of the operator $T[p, \varphi_1](v, v')$ (Eqn. (10)), but recasts Eqns. (9) and (11) as finite-difference expressions $\Delta^+_{st}[\psi_p]$ and $\Delta^-_{st}[\psi_p]$, respectively.

Figs. 20 and 21 show the formulas obtained via the finite-differencing scheme given in Figs. 14 and 19 for positive and negative changes, respectively, for instrumentation relations defined in Fig. 17.

φ	$\Delta^+_{st}[\varphi]$	
$p(w_1,\ldots,w_k)$, $p \in \mathcal{I}$	$((\exists\, v\colon \Delta^+_{st}[\varphi_1]) \land \neg p)\{w_1,\ldots,w_k\}$	if$\psi_p \equiv \exists\, v\colon \varphi_1$
	$\begin{array}{l} (\exists\, v_1', v_2'\colon \Delta^+_{st}[\varphi_1](v_1',v_2')) \\ \land \left(\left(\exists\, v_1', v_2'\colon \begin{array}{c} p(v_1,v_1') \\ \land\, \Delta^+_{st}[\varphi_1](v_1',v_2') \\ \land\, p(v_2',v_2) \end{array} \right) \land \neg p(v_1,v_2) \right) \{w_1,w_2\} \end{array}$	if$\psi_p \equiv$ $(\mathbf{RTC}\ v_1',v_2'\colon \varphi_1)(v_1,v_2)$
	$\Delta^+_{st}[\psi_p]\{w_1,\ldots,w_k\}$	otherwise

φ	$\Delta^-_{st}[\varphi]$	
$p(w_1,\ldots,w_k)$, $p \in \mathcal{I}$	$((\exists\, v\colon \Delta^-_{st}[\varphi_1]) \land p)\{w_1,\ldots,w_k\}$	if$\psi_p \equiv \forall\, v\colon \varphi_1$
	$\begin{array}{l} (\exists\, v_1', v_2'\colon \Delta^-_{st}[\varphi_1](v_1',v_2')) \\ \land \left(\neg \left(\exists\, v_1', v_2'\colon \begin{array}{c} T[p,\varphi_1](v_1,v_1') \\ \land\, T[p,\varphi_1](v_1',v_2') \\ \land\, T[p,\varphi_1](v_2',v_2) \end{array} \right) \land p(v_1,v_2) \right) \{w_1,w_2\} \end{array}$	if$\psi_p \equiv$ $(\mathbf{RTC}\ v_1',v_2'\colon \varphi_1)(v_1,v_2)$
	$\Delta^-_{st}[\psi_p]\{w_1,\ldots,w_k\}$	otherwise

Fig. 19. Extension of the finite-differencing method from Fig. 14 to cover **RTC** formulas, for unit-sized changes to a directed acyclic graph defined by φ_1.

relation p		$\Delta^+_{st}[\psi_p]$
$t_n(v_3,v_4)$		$\Delta^+_{st}[t_n(v_3,v_4)]$
	$=$	$(t_n(v_3,v_4) \lor (\exists\, v_1,v_2\colon t_n(v_3,v_1) \land \Delta^+_{st}[n(v_1,v_2)] \land t_n(v_2,v_4))) \land \neg t_n(v_3,v_4)$
$r_{n,z}(v)$		$\Delta^+_{st}[r_{n,z}(v)]$
	$=$	$(\exists\, v_1 : \Delta^+_{st}[z(v_1) \land t_n(v_1,v)]) \land \neg r_{n,z}(v)$
	$=$	$(\exists\, v_1 : (\Delta^+_{st}[z(v_1)] \land \mathbf{F}_{st}[t_n(v_1,v)]) \lor (\mathbf{F}_{st}[z(v_1)] \land \Delta^+_{st}[t_n(v_1,v)])) \land \neg r_{n,z}(v)$
$c_n(v)$		$\Delta^+_{st}[c_n(v)]$
	$=$	$(\exists\, v_1 : \Delta^+_{st}[n(v_1,v) \land t_n(v,v_1)]) \land \neg c_n(v)$
	$=$	$(\exists\, v_1 : (\Delta^+_{st}[n(v_1,v)] \land \mathbf{F}_{st}[t_n(v,v_1)]) \lor (\mathbf{F}_{st}[n(v_1,v)] \land \Delta^+_{st}[t_n(v,v_1)])) \land \neg c_n(v)$

Fig. 20. The formulas obtained via the finite-differencing scheme given in Figs. 14 and 19 for the positive changes in the values of the instrumentation relations defined in Fig. 17.

5.1.1 *Testing the Unit-Size-Change Assumption.* To know whether this special-case maintenance strategy can be applied, for each statement *st* we need to know at analysis-generation time whether the change performed at *st*, to the graph defined by φ_1, always results in a single edge addition or deletion. If in any admissible structure in $\mathcal{S}_2[\mathcal{R}]$ there is a unique satisfying assignment to the two free variables of $\Delta^+_{st}[\varphi_1]$ and no assignment satisfies $\Delta^-_{st}[\varphi_1]$, then the pair $\Delta^+_{st}[\varphi_1]$, $\Delta^-_{st}[\varphi_1]$ defines a change that adds exactly one edge to the graph. Similarly, if in any admissible structure in $\mathcal{S}_2[\mathcal{R}]$ there is a unique satisfying assignment to the two free variables of $\Delta^-_{st}[\varphi_1]$ and no assignment satisfies $\Delta^+_{st}[\varphi_1]$, then the change is a deletion of exactly one edge from the graph.

Because answering (unique-)satisfiability questions in this logic is in general undecidable, we employ a conservative approximation based on a syntactic analysis of logical formulas. The analysis uses a heuristic to determine a set of variables V such that for each admissible structure, the variables in V have a single possible binding in the formula's satisfying assignments. We refer to such variables as *anchored* variables. For instance, if relation q has the attribute "unique", for each admissible structure there is a single possible binding for variable v in any assignment that

relation p	$\Delta^-_{st}[p]$
$t_n(v_3,v_4)$	$\Delta^-_{st}[t_n(v_3,v_4)]$ $= (\exists v_1,v_2 : T[t_n,n](v_3,v_1) \wedge T[t_n,n](v_1,v_2) \wedge T[t_n,n](v_2,v_4)) \wedge t_n(v_3,v_4)$ $= \begin{pmatrix} & & (t_n(v_3,v_1) \wedge \neg S[t_n,n](v_3,v_1) \vee \mathbf{F}_{st}[n](v_3,v_1)) \\ \exists v_1,v_2 : & \wedge & (t_n(v_1,v_2) \wedge \neg S[t_n,n](v_1,v_2) \vee \mathbf{F}_{st}[n](v_1,v_2)) \\ & \wedge & (t_n(v_2,v_4) \wedge \neg S[t_n,n](v_2,v_4) \vee \mathbf{F}_{st}[n](v_2,v_4)) \end{pmatrix} \wedge t_n(v_3,v_4)$ $= \begin{pmatrix} & & (t_n(v_3,v_1) \wedge \neg(\exists v_1',v_2' : t_n(v_3,v_1') \wedge \Delta^-_{st}[n](v_1',v_2') \wedge t_n(v_2',v_1)) \vee \mathbf{F}_{st}[n](v_3,v_1)) \\ \exists v_1,v_2 : & \wedge & (t_n(v_1,v_2) \wedge \neg(\exists v_1',v_2' : t_n(v_1,v_1') \wedge \Delta^-_{st}[n](v_1',v_2') \wedge t_n(v_2',v_2)) \vee \mathbf{F}_{st}[n](v_1,v_2)) \\ & \wedge & (t_n(v_2,v_4) \wedge \neg(\exists v_1',v_2' : t_n(v_2,v_1') \wedge \Delta^-_{st}[n](v_1',v_2') \wedge t_n(v_2',v_4)) \vee \mathbf{F}_{st}[n](v_2,v_4)) \end{pmatrix}$ $\wedge\, t_n(v_3,v_4)$
$r_{n,z}(v)$	$\Delta^-_{st}[r_{n,z}(v)]$ $= \Delta^-_{st}[\exists v_1 : x(v_1) \wedge t_n(v_1,v)]$ $= (\exists v_1 : \Delta^-_{st}[z(v_1) \wedge t_n(v_1,v)]) \wedge \neg(\exists v_1 \mathbf{F}_{st}[z(v_1) \wedge t_n(v_1,v)])$ $= \begin{cases} (\exists v_1 : ((\Delta^-_{st}[z(v_1)] \wedge t_n(v_1,v)) \vee (z(v_1) \wedge \Delta^-_{st}[t_n(v_1,v)]))) \\ \wedge \\ \neg\left(\exists v_1 : \begin{pmatrix} (z(v_1) \wedge t_n(v_1,v)) \\ ?\ \neg\Delta^-_{st}[z(v_1) \wedge t_n(v_1,v)] \\ :\ \Delta^+_{st}[z(v_1) \wedge t_n(v_1,v)] \end{pmatrix}\right) \end{cases}$ $= \begin{cases} (\exists v_1 : ((\Delta^-_{st}[z(v_1)] \wedge t_n(v_1,v)) \vee (z(v_1) \wedge \Delta^-_{st}[t_n(v_1,v)]))) \\ \wedge \\ \neg\left(\exists v_1 : \begin{pmatrix} (z(v_1) \wedge t_n(v_1,v)) \\ ?\ \neg((\Delta^-_{st}[z(v_1)] \wedge t_n(v_1,v)) \vee (z(v_1) \wedge \Delta^-_{st}[t_n(v_1,v)])) \\ :\ ((\Delta^+_{st}[z(v_1)] \wedge \mathbf{F}_{st}[t_n(v_1,v)]) \vee (\mathbf{F}_{st}[z(v_1)] \wedge \Delta^+_{st}[t_n(v_1,v)])) \end{pmatrix}\right) \end{cases}$
$c_n(v)$	$\Delta^-_{st}[c_n(v)]$ $= \Delta^-_{st}[\exists v_1 : n(v_1,v) \wedge t_n(v,v_1)]$ $= (\exists v_1 : \Delta^-_{st}[n(v_1,v) \wedge t_n(v,v_1)]) \wedge \neg\mathbf{F}_{st}[\exists v_1 : n(v_1,v) \wedge t_n(v,v_1)]$ $= (\exists v_1 : (\Delta^-_{st}[n(v_1,v)] \wedge t_n(v,v_1)) \vee (n(v_1,v) \wedge \Delta^-_{st}[t_n(v,v_1)])) \wedge \neg\mathbf{F}_{st}[\exists v_1 : n(v_1,v) \wedge t_n(v,v_1)]$

Fig. 21. The formulas obtained via the finite-differencing scheme given in Figs. 14 and 19 for the negative changes in the values of the instrumentation relations defined in Fig. 17.

satisfies $q(v)$; in a formula that contains an occurrence of $q(v)$, v is an anchored variable.

If both free variables of $\Delta^+_{st}[\varphi_1]$ are anchored and $\Delta^-_{st}[\varphi_1] = 0$, then the change adds one edge to the graph defined by φ_1. Similarly, if both free variables of $\Delta^-_{st}[\varphi_1]$ are anchored and $\Delta^+_{st}[\varphi_1] = 0$, then the change removes one edge from the graph. In these cases, the reflexive transitive closure of φ_1 can be updated using the method discussed above.

A Test for Anchored Variables. Function *Anchored*, shown in Fig. 22, conservatively identifies anchored variables in a formula φ. It is invoked as *Anchored*$(\varphi, \emptyset)$. (In our application, at top-level φ is always either $\Delta^+_{st}[\varphi_1]$ or $\Delta^-_{st}[\varphi_1]$.) *Anchored* uses a handful of patterns to identify anchored variables. For example, if variable v_1 is anchored and binary relation p has the attribute "function",[8] then v_2 is anchored as well. In essence, negations are handled by pushing the negation deeper into the formula. In a disjunction, an anchored variable must be anchored in both subformulas. The conjunction rule accumulates anchored variables in A by a process of successive approximation, during which variables anchored in the left subformula are used to identify new anchored variables in the right subformula and vice versa; this process is iterated until a fixed point is reached. The rules for $\exists v \colon \varphi_1$ and $\forall v \colon \varphi_1$ contain recursive calls on *Anchored* with v removed from the second argument (because bound variable v refers to a different occurrence of v

[8] For instance, in program-analysis applications a relation $n(v_1, v_2)$ that records whether field `n` of v_1 points to v_2 has the "function" attribute.

φ	$\mathit{Anchored}(\varphi, A_0)$
$\mathbf{0}, \mathbf{1}$	A_0
$v_1 = v_2$	$v_1 \in A_0 \to A_0 \cup \{v_2\} \talloblong v_2 \in A_0 \to A_0 \cup \{v_1\} \talloblong A_0$
$p()$	A_0
$p(v)$	$\mathit{unique}(p) \to A_0 \cup \{v\} \talloblong A_0$
$p(v_1, v_2)$	$\mathit{function}(p) \wedge v_1 \in A_0 \to A_0 \cup \{v_2\}$ $\talloblong\ \mathit{invfunction}(p) \wedge v_2 \in A_0 \to A_0 \cup \{v_1\}$ $\talloblong\ A_0$
$\neg\varphi_1$	$\varphi_1 \equiv \neg\varphi_2 \to \mathit{Anchored}(\varphi_2, A_0)$ $\talloblong\ \varphi_1 \equiv \varphi_2 \vee \varphi_3 \to \mathit{Anchored}(\neg\varphi_2 \wedge \neg\varphi_3), A_0)$ $\talloblong\ \varphi_1 \equiv \varphi_2 \wedge \varphi_3 \to \mathit{Anchored}(\neg\varphi_2 \vee \neg\varphi_3), A_0)$ $\talloblong\ \varphi_1 \equiv \forall\, v\colon \varphi_2 \to \mathit{Anchored}(\exists\, v\colon \neg\varphi_2), A_0)$ $\talloblong\ \varphi_1 \equiv \exists\, v\colon \varphi_2 \to \mathit{Anchored}(\forall\, v\colon \neg\varphi_2), A_0)$ $\talloblong\ A_0$
$\varphi_1 \vee \varphi_2$	$\mathit{Anchored}(\varphi_1, A_0) \cap \mathit{Anchored}(\varphi_2, A_0)$
$\varphi_1 \wedge \varphi_2$	$\mu A.(\mathit{Anchored}(\varphi_1, A \cup A_0) \cup \mathit{Anchored}(\varphi_2, A \cup A_0))$
$\exists\, v\colon \varphi_1, \forall\, v\colon \varphi_1$	$(\mathit{Anchored}(\varphi_1, A_0 - \{v\}) - \{v\}) \cup A_0$
$(\mathbf{RTC}\ v_1', v_2'\colon \varphi_1)(v_1, v_2)$	$(\mathit{Anchored}(\varphi_1, A_0 - \{v_1', v_2'\}) - \{v_1', v_2'\}) \cup A_0$

Fig. 22. Function *Anchored* conservatively identifies anchored variables in φ. A_0 contains variables known to be anchored due to the surrounding context.

from an identically named v in A_0). If v is anchored in φ_1, it needs to be removed before this call returns, to avoid confusion with a v in the outer scope (note the second subtraction of $\{v\}$). Finally, the union of A_0 is performed because v may be in A_0, in which case it has to be included in the answer. $(\mathbf{RTC}\ v_1', v_2'\colon \varphi_1)(v_1, v_2)$ is handled similarly to $\exists\, v\colon \varphi_1$ and $\forall\, v\colon \varphi_1$.

5.2 Transitive-Closure Maintenance in Tree-Shaped Graphs

Consider a binary instrumentation relation p, defined by $\psi_p(v_1, v_2) \equiv (\mathbf{RTC}\ v_1', v_2'\colon \varphi_1)(v_1, v_2)$. If the graph defined by φ_1 is not only acyclic but is tree-shaped, it is possible to take advantage of this fact.[9] This fact has no bearing on the maintenance formula that updates the values of relation p after a positive unit-size change $\Delta^+[\varphi_1]$ to the relation φ_1 (see Formula (9)). However, it allows the values of p to be updated in a more efficient manner after a negative unit-size change $\Delta^-[\varphi_1]$ to φ_1. In a tree-shaped graph, there exists at most one path between a pair of nodes; if that path goes through the φ_1 edge to be deleted, it should be removed (cf. Formula (11)):

$$\mathbf{F}_{st}[p](v_1, v_2) = p(v_1, v_2) \wedge \neg(\exists\, v_1', v_2'\colon p(v_1, v_1') \wedge \Delta^-_{st}[\varphi_1](v_1', v_2') \wedge p(v_2', v_2)). \quad (12)$$

Fig. 23 extends the method for generating relation-maintenance formulas to handle instrumentation relations specified via the **RTC** of a binary formula that defines a tree-shaped graph. Fig. 23 recasts Eqn. (12) as a finite-difference expression $\Delta^-_{st}[\psi_p]$.

When comparing the techniques of §5.1 for the maintenance of the **RTC** of a

[9] The special-case maintenance strategy that we describe in this subsection also applies only in the case that the change to the graph is a single edge addition or deletion (but not both). We rely on the test described in §5.1.1 to ensure that this is the case.

φ	$\Delta^-_{st}[\varphi]$	
$p(w_1, \ldots, w_k)$, $p \in \mathcal{I}$	$((\exists\, v\colon \Delta^-_{st}[\varphi_1]) \wedge p)\{w_1, \ldots, w_k\}$	if$\psi_p \equiv \forall\, v\colon \varphi_1$
	$(\exists\, v'_1, v'_2\colon p(v_1, v'_1) \wedge \Delta^-_{st}[\varphi_1](v'_1, v'_2) \wedge p(v'_2, v_2))\{w_1, w_2\}$	if$\psi_p \equiv$ $(\mathbf{RTC}\ v'_1, v'_2\colon \varphi_1)(v_1, v_2)$
	$\Delta^-_{st}[\psi_p]\{w_1, \ldots, w_k\}$	otherwise

Fig. 23. Extension of the finite-differencing method from Fig. 14 to cover **RTC** formulas, for unit-sized changes to a tree-shaped graph defined by φ_1. The finite-difference expression $\Delta^+_{st}[\psi_p]$ is as defined in Fig. 19.

binary formula φ_1 with those presented in this subsection, we will refer to the method of §5.1 as *acyclic-*φ_1 *maintenance* and the method of this subsection as *tree-shaped-*φ_1 *maintenance*.

5.3 Reachability Maintenance in Deterministic Graphs

A *deterministic graph* is a graph in which every node has outdegree at most one. If the graph defined by φ_1 is deterministic, it is possible to give first-order formulas that maintain reachability information in the graph in response to the addition or deletion of a single φ_1-edge.

The class of deterministic graphs corresponds exactly to the set of possibly-cyclic linked lists. Our solution to the problem of reachability maintenance in possibly-cyclic linked lists can be summarized as follows:

(1) A binary instrumentation relation sfe_n(for **s**panning-**f**orest **e**dge) is introduced to maintain a spanning forest of the (possibly-cyclic) graph defined by the n edges. Thus, we have two types of edges: possibly-cyclic n edges and acyclic sfe_n edges.

(2) We introduce an instrumentation relation sfp_n (for **s**panning-**f**orest **p**ath) that captures reachability along the (acyclic) sfe_n edges. sfp_n is the **RTC** of sfe_n, but because sfe_n is acyclic and tree-shaped, sfp_n can be maintained via the techniques described in §5.1 and §5.2.

(3) We introduce an instrumentation relation t_n to capture reachability along n edges. Instead of defining t_n as n^*, as done in Fig. 17, we express t_n using first-order logic, based on sfp_n (cf. Fig. 28). Thus, t_n can be maintained in terms of sfp_n, via the techniques described in §4.

(4) A unary core relation roc_n (for **r**epresentative **o**f the **c**ycle) is introduced to identify a distinguished node of each cycle; the outgoing n edge from a roc_n node is a cycle-breaking edge that is *not* used in the construction of the spanning forest.

In other words, we have a two-level scheme: reachability in the induced, acyclic spanning forest (sfp_n) is maintained via the rules from §5.1 and §5.2; reachability in the underlying, possibly-cyclic graph (t_n) is then maintained via the rules from §4.

5.3.1 *Abstractions of Possibly-Cyclic Linked Lists.* We will illustrate our techniques on *panhandle lists*, i.e., linked lists that contain a cycle but in which at least the head of the list is not part of the cycle. (The lists shown in Fig. 24 are examples of panhandle lists.) Fig. 3 gives the definition of a C linked-list datatype, and

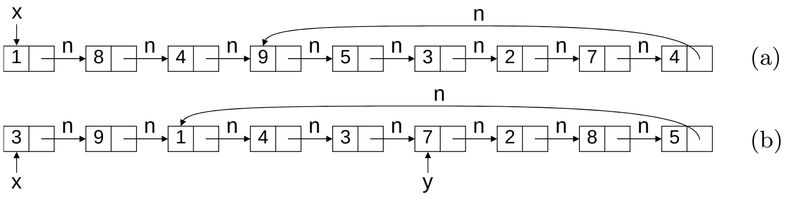

Fig. 24. Possible stores for *panhandle* linked lists. (a) A panhandle list pointed to by `x`. We will refer to lists of this shape as type-X lists. (b) A panhandle list pointed to by `x` with `y` pointing into the middle of the cycle. We will refer to lists of this shape as type-XY lists.

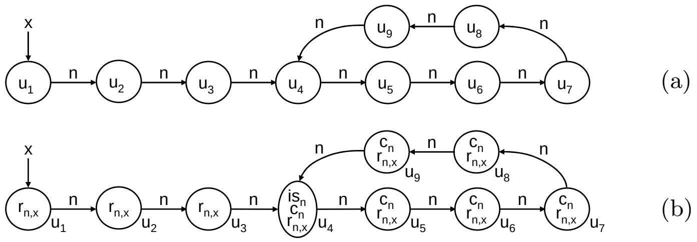

Fig. 25. A logical structure S_{25} that represents the store shown in Fig. 24(a) in graphical form: (a) S_{25} with relations from Fig. 3; (b) S_{25} with relations from Figs. 3 and 6. (Transitive-closure relation t_n has been omitted to reduce clutter.)

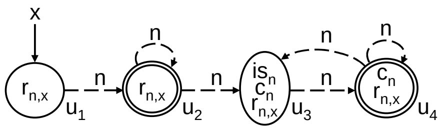

Fig. 26. A 3-valued structure S_{26} that is the canonical abstraction of structure S_{25}. In addition to S_{25}, S_{26} represents any type-X panhandle list with at least two nodes in the panhandle and at least two nodes in the cycle.

lists the core relations that would be used to represent the stores manipulated by programs that use type `List`, such as the stores in Fig. 24.

Fig. 25 shows two versions of 2-valued structure S_{25}, which represents the store shown in Fig. 24(a): Fig. 25(a) shows the relations from Fig. 3.[10] Fig. 25(b) shows the relations from Fig. 3, as well as the instrumentation relations from Fig. 6.

If all unary relations are abstraction relations ($\mathcal{A} = \mathcal{R}_1$), the canonical abstraction of 2-valued logical structure S_{25} is S_{26}, shown in Fig. 26, with list nodes corresponding to u_2 and u_3 in S_{25} represented by the summary individual u_2 of S_{26} and list nodes corresponding to u_5 and u_6 in S_{25} represented by the summary individual u_4 of S_{26}. S_{26} represents any type-X panhandle list with at least two nodes in the panhandle and at least two nodes in the cycle.

5.3.2 *Reachability Maintenance in Possibly-Cyclic Linked Lists.* Unfortunately, the relations defined in Figs. 3 and 6 do not permit precise maintenance of reach-

[10] We will not show the *dle* relation in the rest of this section because it is not relevant to the problem of reachability maintenance.

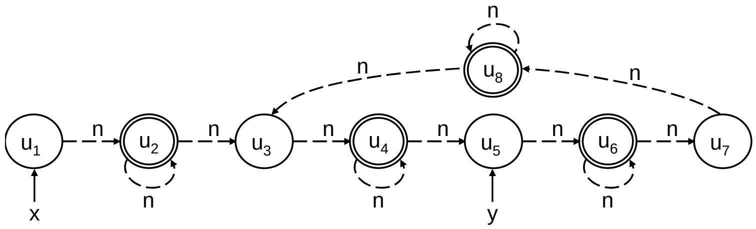

Fig. 27. Logical structure S_{27} that represents type-XY panhandle lists, such as the store depicted in Fig. 24(b). The relations from Fig. 6 are omitted to reduce clutter. Their values are as expected for a type-XY list: $r_{n,x}$ holds for all nodes, $r_{n,y}$ and c_n hold for all nodes on the cycle, and is_n holds for u_3.

ability information, such as relation $r_{n,x}$, in possibly-cyclic lists. A difficulty arises when reachability information has to be updated after the deletion of an n edge on a cycle (e.g., as a result of statement `y->n = NULL`). With the relations defined in Figs. 3 and 6, such an update requires the reevaluation of a transitive-closure formula, which generally results in a drastic loss of precision in the presence of abstraction.

We demonstrate the issue on panhandle lists represented by the abstract structure S_{27} shown in Fig. 27, i.e., lists of type XY. (Note that although the store depicted in Fig. 24(b) embeds into structure S_{27}, S_{27} is not the canonical abstraction of the store from Fig. 24(b); in particular, nodes u_4, u_6, u_7, and u_8 are all indistinguishable according to the instrumentation relations discussed thus far. However, this embedding gives Fig. 27 a shape similar to figures that appear later in the section, which will help in illustrating our solution.) Statement `y->n = NULL` has the effect of deleting the `n` edge leaving u_5, thus making the nodes represented by u_6, u_7, and u_8 unreachable from `x`.[11] Note that a first-order-logic formula over the relations of Figs. 3 and 6 cannot distinguish the list nodes represented by u_4 from those represented by u_6, u_7, and u_8: all of those nodes are reachable from both `x` and `y`, none of those nodes are shared, and all of them lie on a cycle. Our inability to characterize the group of nodes represented by u_4 via a first-order formula requires the maintenance formula for the reachability relation $r_{n,x}$ to recompute some transitive-closure information, e.g., the transitive-closure subformula of the definition of $r_{n,x}$, namely, $n^*(v_1, v)$. However, in the presence of abstraction, reevaluating transitive-closure formulas often yields 1/2. For instance, in S_{27}, formula $n^*(v_1, v)$ evaluates to 1/2 under the assignment $[v_1 \mapsto u_1, v \mapsto u_4]$ because of the many 1/2 values of relation n (see the dashed edges connecting u_1 with u_2, for example).

The essence of a solution that enables maintaining reachability relations for possibly-cyclic lists in first-order logic is to find a way to break the symmetry of each cycle. The basic idea for a solution was suggested to us by W. Hesse and N. Immerman. As discussed at the beginning of §5.3, it consists of maintaining a spanning-tree representation of a possibly-cyclic list. Reachability in such a representation can be maintained using first-order-logic formulas. Reachability in the actual list can be expressed in first-order logic based on the spanning-tree representation. We now explain our approach and highlight some differences with the

[11] Clearly, all nodes except u_5 also become unreachable from `y`.

p	**Intended Meaning**	ψ_p
$is_n(v)$	Do `n` fields of two or more list nodes point to v?	$\exists\, v_1, v_2 : n(v_1, v) \land n(v_2, v) \land v_1 \neq v_2$
$sfe_n(v_1, v_2)$	Is there an `n` edge from v_2 to v_1 (assuming that v_2 is not a roc_n node)	$n(v_2, v_1) \land \neg roc_n(v_2)$
$sfp_n(v_1, v_2)$	Is v_2 reachable from v_1 along sfe_n edges?	$sfe_n^*(v_1, v_2)$
$t_n(v_1, v_2)$	Is v_2 reachable from v_1 along `n` fields?	$sfp_n(v_2, v_1) \lor \exists\, u, w : \begin{pmatrix} sfp_n(u, v_1) \land \\ roc_n(u) \land n(u, w) \\ \land\, sfp_n(v_2, w) \end{pmatrix}$
$r_{n,x}(v)$	Is v reachable from pointer variable `x` along `n` fields?	$\exists\, v_1 : x(v_1) \land t_n(v_1, v)$
$c_n(v)$	Is v on a directed cycle of `n` fields?	$\exists\, v_1, v_2 : \; roc_n(v_1) \land n(v_1, v_2) \land \; sfp_n(v, v_2)$
$pr_x(v)$	Does v lie on an sfe_n path from `x` (does v **pr**ecede `x` on an `n`-path to a roc_n node)?	$\exists\, v_1 : x(v_1) \land sfp_n(v_1, v)$
$pr_{is}(v)$	Does v lie on an sfe_n path from a shared node (does v **pr**ecede a shared node on an `n`-path to a roc_n node)?	$\exists\, v_1 : is_n(v_1) \land sfp_n(v_1, v)$

Fig. 28. Defining formulas of instrumentation relations. The sharing relation is_n is defined as in Fig. 6. Relations t_n, $r_{n,x}$, and c_n are redefined via first-order-logic formulas in terms of other relations. The exact meaning and purpose of relations sfe_n, sfp_n, pr_x, and pr_{is} will be explained later in the section. Their names stand for **s**panning-**f**orest **e**dge, **s**panning-**f**orest **p**ath, **pr**ecedes `x` (along a certain path in a cycle), and **pr**ecedes a shared node (along a certain path in a cycle), respectively.

approach taken by Hesse [Hesse 2003].

Our approach relies on the introduction of additional core and instrumentation relations. We extend the set of core relations (Fig. 3) with unary relation roc_n, which designates one node on each cycle to be the **r**epresentative **o**f the **c**ycle. (We refer to such a node as a roc_n node.) Relation roc_n is used for tracking a unique *cut edge* on each cycle, which allows the maintenance of a spanning tree. Fig. 29(a) shows 2-valued structure S_{29}, which represents the store of Fig. 24(a) using the extended set of core relations. Here, we let u_7 be the roc_n node. In general, we simply require that exactly one node on each cycle be designated as a roc_n node. Later in this section we describe how we ensure this.[12]

Fig. 28 lists the extended set of instrumentation relations. Note that relation roc_n is not part of the semantics of the language. A natural question is whether $roc_n(v)$ can be defined as an instrumentation relation. For instance, we could try to define it using the following defining formula:

$$c_n(v) \land \exists\, v_1 : n(v_1, v) \land \neg c_n(v_1) \qquad (13)$$

Formula (13) identifies nodes that lie on a cycle but have a predecessor that does not lie on the cycle. There are three problems with this approach. First, this definition

[12] With the relation-transfer formulas that we use for relation roc_n in this paper, the roc_n node for a cycle is the source of the n edge that was inserted to complete the cycle. Note that with this approach, the node that receives the roc_n designation in a given cycle depends on the order of operations that the program performs to construct the cycle.

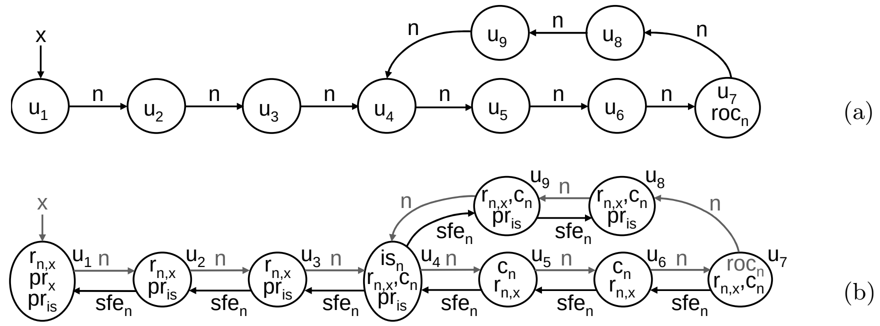

Fig. 29. A logical structure S_{29} that represents the store shown in Fig. 24(a) in graphical form: (a) S_{29} with the extended set of core relations.(b) S_{29} with the extended set of core and instrumentation relations (core relations appear in grey). Transitive-closure relations sfp_n and t_n have been omitted to reduce clutter. The values of the transitive-closure relations can be readily seen from the graphical representation of relations sfe_n and n. For instance, node u_5 is related via the sfp_n relation to itself and all nodes appearing to the left or above it in the pictorial representation.

works for panhandle lists but not for cyclic lists without a panhandle. (In general, no other definition can work for cyclic lists without a panhandle because if one existed, it would need to choose one list node among identical-looking nodes that lie on each cycle.) Second, because the cyclicity relation c_n is defined in terms of roc_n (and sfp_n), the definition of roc_n has a circular dependence, which is disallowed. (This circularity cannot be avoided, if we want all reachability relations to benefit from the precise maintenance of one transitive-closure relation—here, sfp_n.) The third problem with introducing roc_n as an instrumentation relation is discussed later in the section (see footnote 14).

We divide our description of the abstraction based on the new set of relations into three parts, which describe (i) how the relations of Fig. 28 define *directed* spanning forests, (ii) how we maintain precision on a cycle in the presence of abstraction, and (iii) how we generate maintenance formulas for instrumentation relations *automatically.* The three parts highlight the differences between our approach and that of Hesse.

Defining Directed Spanning Forests. Recall from §5.3.2 that the core relations are extended with unary relation roc_n, which designates one node on each cycle to be the **r**epresentative **o**f the **c**ycle. The roc_n nodes can be used to define a (directed) spanning forest of the n edges. Instrumentation relation sfe_n—*sfe* stands for **s**panning-**f**orest **e**dge—is used to maintain the set of edges that forms the spanning forest. In Hesse's work, the spanning-forest edges are directed in the same direction as the n edges; as a result, he maintains a directed spanning forest in which each edge is directed towards the root of a spanning-forest tree. In our work, we define sfe_n to be directed in the direction *opposite* to that of the n edges. The graph defined by the sfe_n relation then defines a directed spanning forest with roc_n nodes as spanning-forest roots, and with each spanning-forest edge directed away from the root of a spanning-forest tree (see Fig. 29(b)).

Instrumentation relation sfp_n—*sfp* stands for **s**panning-**f**orest **p**ath—is used to maintain the set of paths in the spanning forest of list nodes. Binary reachability in the actual lists (see relation t_n in Fig. 28) can be defined in terms of n, roc_n, and sfp_n using a first-order-logic formula: v_2 is reachable from v_1 if there is a spanning-forest path from v_2 to v_1 or there is a pair of spanning-forest paths, one from the source of a cut edge (a roc_n node) to v_1 and the other from v_2 to the target of the cut edge (the n-successor of the same roc_n node).

Unary reachability relations $r_{n,x}$ and the cyclicity relation c_n can be defined via first-order formulas, as well. We defined $r_{n,x}$ in terms of binary reachability relation t_n. While we could define c_n in terms of t_n, as well, we chose another simple definition by observing that a node lies on a cycle if and only if there is a spanning-forest path from it to the target of a cut edge (the n-successor of a roc_n node).

Fig. 29(b) shows 2-valued structure S_{29}, which represents the store of Fig. 24(a) using the extended set of core and instrumentation relations. The relations pr_x and pr_{is} will be explained shortly.

Preserving Node Ordering on a Cycle in the Presence of Abstraction. The fact that our techniques need to be applicable in the presence of abstraction introduces a complication that is not present in the setting studied by Hesse. His concern was with the expressibility of certain properties within the confines of a logic with certain syntactic restrictions. Our concern is with the ability to maintain precision in the framework of canonical abstraction.

Unary reachability relations $r_{n,x}$ (one for every program variable x) play a crucial role in the analysis of programs that manipulate acyclic linked lists. In addition to keeping disjoint lists summarized separately, they keep list nodes that have been visited during a traversal summarized separately from nodes that have not been visited: if x is the pointer used to traverse the list, then the nodes that have been visited will have value 0 for relation $r_{n,x}$, while the nodes that have not been visited will have value 1. If a list contains a cycle, then all nodes on the cycle are reachable from the same set of variables, namely, all variables that point to any node in that list. As a result, the instrumentation relations discussed thus far cannot prevent nodes u_4, u_6, and u_8 of S_{27} shown in Fig. 27 from being summarized together. Thus, assuming that u_7 is the roc_n node, the canonical abstraction of S_{27} is the 3-valued structure S_{30} shown in Fig. 30. The nodes represented by u_4, u_6, and u_8 of S_{27} are represented by the single summary individual u_6 in S_{30}. The symmetry hides all information about the order of traversal via pointer variable y. Moreover, the values of the sfp_n relation (not shown in Fig. 30) lose precision because ancestors of the shared node

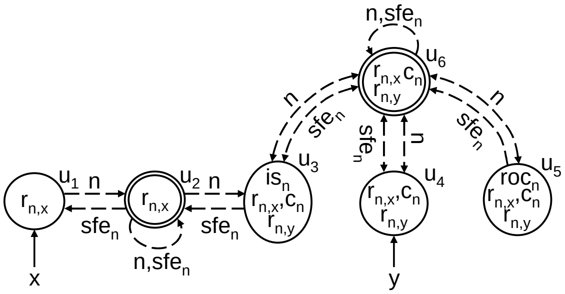

n,sfe$_n$
u$_6$
r$_{n,x}$ c$_n$ r$_{n,y}$
n
sfe$_n$
sfe$_n$
n
n
sfe$_n$
u$_3$
is$_n$ r$_{n,x}$,c$_n$ r$_{n,y}$
u$_4$
r$_{n,x}$,c$_n$ r$_{n,y}$
u$_5$
roc$_n$ r$_{n,x}$,c$_n$ r$_{n,y}$
u$_1$ n
u$_2$ n
r$_{n,x}$
r$_{n,x}$
sfe$_n$
sfe$_n$
n,sfe$_n$
x
y

Fig. 30. A 3-valued structure S_{30} that is the canonical abstraction of structure S_{27} if relations pr_x and pr_{is} are not added to $\mathcal{A}$ and node u_7 is the roc_n node.

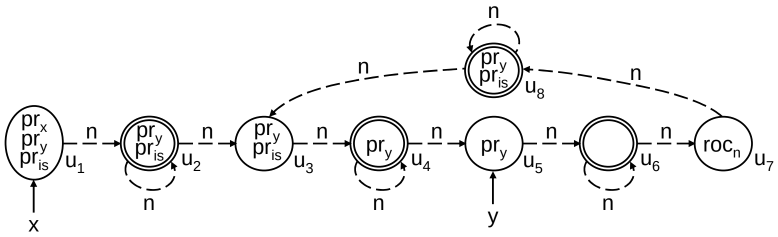

Fig. 31. A 3-valued structure S_{31} that is the canonical abstraction of structure S_{27} if node u_7 is the roc_n node. S_{31} represents panhandle lists of type XY, such as the store of Fig. 24(b). The only instrumentation relations shown in the figure are pr_x, pr_y, and pr_{is}. As in structure S_{27} shown in Fig. 27, $r_{n,x}$ holds for all nodes, $r_{n,y}$ and c_n hold for all nodes on the cycle, and is_n holds for u_3.

in the spanning tree are summarized together with its descendants in the spanning tree.

We break the symmetry of the nodes on a cycle using a general mechanism via unary properties akin to unary reachability relations $r_{n,x}$. In the definitions of relations pr_x of Fig. 28, full reachability (relation t_n) has been replaced with spanning-forest reachability (relation sfp_n). The relations pr_x distinguish nodes according to whether or not they are reachable from program variable x along spanning-forest edges. The relation pr_{is} is defined similarly but using instrumentation relation is_n; pr_{is} partitions the nodes of a panhandle list into ancestors and descendants of the shared node in the spanning tree. Fig. 31 shows structure S_{31} that is the canonical abstraction of S_{27} of Fig. 27, assuming that u_7 is the roc_n node. In S_{31}, each of the nodes u_4, u_6, and u_8 has a distinct vector of values for the relations pr_y and pr_{is}, thus breaking the symmetry.[13]

Automatic Generation of Maintenance Formulas for Instrumentation Relations. In his thesis, Hesse gives hand-specified relation-maintenance formulas for a collection of relations that are used for maintaining a spanning-forest representation of possibly-cyclic linked lists. Instead of specifying relation-maintenance formulas by hand, we rely on finite differencing, as described in previous sections of this paper, to generate maintenance formulas for all instrumentation relations. Finite-differencing-generated maintenance formulas have been effective in maintaining all relations defined via first-order-logic formulas, i.e., all relations of Fig. 28 except sfp_n. Additionally, under certain conditions, finite-differencing-generated maintenance formulas have been effective in maintaining relations defined via the reflexive transitive closure of binary relations. Two conditions are necessary for this technique to be applicable for maintaining relation sfp_n:

Graph-shape condition. The graph defined by sfe_n must be acyclic or tree-shaped.

Unit-size-change condition. Each program statement must only change the graph

[13] Note that in the presence of multiple panhandles, we may be required to introduce finer distinctions to account for the possibility of multiple shared nodes on a cycle. These distinctions could take the form of a family of *is-shared* relations—one for each variable—to capture the panhandle that contributes to sharing. We do not discuss a detailed solution to this problem here, as it will not provide significant further insight.

of n edges by adding a single edge or deleting a single edge (but not both).

The graph-shape condition applies in our setting because the graph defined by sfe_n defines a spanning forest (which is both acyclic and tree-shaped). The unit-size-change condition requires some discussion.

The relation sfe_n is defined in terms of n and roc_n. While we have not yet discussed the relation-update formulas for core relation roc_n, it should be clear that the value of the relation roc_n should only change in response to a change in the value of a node's `n` field. There are two types of statements that change the value of the `n` field and thus may have an effect on the value of the sfe_n relation—namely, statements of the forms `x->n = NULL` and `x->n = y`. The former destroys the n edge leaving the node pointed to by `x`, and the latter creates a new n-connection from the node pointed to by `x` to the node pointed to by `y`. While both of these statements add or remove a single edge of the n relation, it is not necessarily the case that they add or remove a single edge of the sfe_n relation. When interpreted on logical structure S_{31} of Fig. 31, statement `y->n = NULL` has the effect of deleting the n edge leaving u_5, an action that should result in the deletion of the sfe_n edge entering u_5 (not shown in the figure). However, to preserve the spanning-forest representation, we need to ensure that roc_n holds only for nodes that lie on a cycle, and that sfe_n represents spanning-forest edges. This requires setting the value of roc_n for u_7 to 0 and adding an sfe_n edge from u_8 to u_7. Because, as this example illustrates, a language statement may result in the deletion of one sfe_n edge and the addition of another, neither of the techniques from §5.1 and §5.2 for maintaining instrumentation relations defined via **RTC** applies.

To work around this problem, we apply each transformer associated with statements `x->n = NULL` and `x->n = y` in two phases. In the first phase, we apply the part of the transformer that corresponds to the relation n, and update the values of all instrumentation relations. In the second phase, we apply the part of the transformer that corresponds to the relation roc_n, and update the values of all instrumentation relations. As we explain below, each phase of the two transformers satisfies the requirement that the change adds a single edge or removes a single edge of the sfe_n relation.[14] Additionally, by paying attention to the order of phases, we ensure that the graph defined by the relation sfe_n remains acyclic and tree-shaped throughout the application of the transformers.

To preserve the graph-shape condition in the case of statement `x->n = NULL`, we apply the part of the transformer that corresponds to the relation n first:

$$\tau_{n,\texttt{x->n = NULL}}(v_1, v_2) = n(v_1, v_2) \wedge \neg x(v_1). \tag{14}$$

Unless `x` points to a roc_n node (or `x->n` is `NULL`), this phase results in the deletion of the sfe_n edge that enters the node pointed to by `x`. In the second phase, we apply the part of the transformer that corresponds to the relation roc_n:

$$\tau_{roc_n,\texttt{x->n = NULL}}(v) = roc_n(v) \wedge \exists\, v_1 : n(v, v_1) \wedge \mathit{sfp}_n(v, v_1). \tag{15}$$

[14] The third problem with defining roc_n as an instrumentation relation (alluded to earlier in the section) is that we would lose the ability to apply the two parts of a transformer separately: the change in the values of n would immediately trigger a change in the values of roc_n. The resulting transformer would not be able to satisfy the unit-size-change condition.

This phase sets the roc_n property of the source n_s of a cut edge to 0, if there is no longer a spanning-forest path from n_s to the target n_t of the same cut edge. When this happens and x does not point to n_s, i.e., the cut edge is not being deleted, this phase results in the addition of an sfe_n edge from n_t to n_s.

To preserve the graph-shape condition in the case of statement `x->n = y`, we apply the part of the transformer that corresponds to the relation roc_n first:

$$\tau_{roc_n,\mathtt{x->n\ =\ y}}(v) = roc_n(v) \vee (x(v) \wedge \exists\, v_1 : y(v_1) \wedge sfp_n(v, v_1)). \qquad (16)$$

If there is a spanning-forest path from node n_x, pointed to by x, to node n_y, pointed to by y, the statement creates a new cycle in the data structure. The update of Formula (16) sets the roc_n property of n_x to 1, thus making n_x the source of a new cut edge and n_y the target of the cut edge. Because there was no n edge from n_x to n_y prior to the execution of this statement,[15] this phase results in no change to the sfe_n relation. In the second phase, we apply the part of the transformer that corresponds to the relation n:

$$\tau_{n,\mathtt{x->n\ =\ y}}(v_1, v_2) = n(v_1, v_2) \vee (x(v_1) \wedge y(v_2)). \qquad (17)$$

Unless the node pointed to by x became a roc_n node in the first phase, this phase results in the addition of an sfe_n edge from n_y to n_x.

The break-up of the transformers corresponding to statements `x->n = NULL` and `x->n = y` into two phases, as described above, ensures that the sfe_n relation remains acyclic and tree-shaped throughout the analysis (the graph-shape condition) and that the change to the sfe_n relation that results from each phase is a unit-size change (the unit-size-change condition).[16] Thus, it is sound to maintain sfp_n $(= sfe_n^*)$ via the techniques described in either §5.1 or §5.2. Additionally, it is also sound to maintain the remaining instrumentation relations via the techniques of §4 because the remaining relations are defined by first-order-logic formulas. Soundness guarantees that the stored values of instrumentation relations agree with the relations' defining formulas throughout the analysis. However, the stored values may not agree with the relations' *intended meanings*. For instance, if the n-transfer phase of the transformer for statement `x->n = NULL` removes a non-cut n edge on a cycle, the sfe_n relation will temporarily not span the entire list. However, we do not query the results of abstract interpretation *in between* the phases of a two-phase transformer. Thus, the stored values of instrumentation relations agree with the relations' intended meanings, as well as their defining formulas, at all points in the program's control-flow graph.

Optimized Maintenance of Relation sfp_n. By demonstrating that the acyclicity and unit-size-change conditions hold for relation sfe_n, we were able to rely on the techniques of §5.1 to maintain the relation sfp_n. Note, however, that the definition of sfe_n ensures that the graph defined by sfe_n is not only acyclic but is tree-shaped. This fact has no bearing on the maintenance formula that updates the values of relation sfp_n after a positive unit-size change $\Delta^+[sfe_n]$ to sfe_n (see Formula (9)). However, it allows the values of sfp_n to be updated in a more efficient manner after

[15] By normalizing procedures to include a statement of the form `x->n = NULL` prior to a statement of the form `x->n = y`, we ensure that `x->n` is always `NULL` prior to the latter assignment.

[16] The test described in §5.1.1 confirms that the unit-size-change condition holds for each phase.

Category	Test Program	# of non-identity maintenance formulas				Performance				
						Analysis Time (sec.)			% increase	
		schemas			# inst.	Ref.	FD		FD	
		total	TC	non-TC			acyc.	tree	acyc.	tree
SLL Shape Analysis	Search	2	0	2	2	0.30	0.30	0.31	1.10	1.90
	GetLast	3	0	3	4	0.31	0.32	0.32	2.23	2.22
	DeleteAll	11	2	9	15	0.30	0.32	0.30	4.97	-0.13
	Reverse	12	2	10	16	0.43	0.49	0.44	12.69	1.99
	Create	11	2	9	21	0.28	0.31	0.28	9.61	-0.60
	Delete	12	2	10	39	1.13	2.13	1.23	87.90	7.76
	Merge	11	2	9	64	1.77	3.67	1.96	107.27	10.42
	Insert	12	2	10	72	1.19	2.03	1.31	70.43	9.67
DLL Shape Analysis	Append	15	2	13	50	1.76	1.78	1.77	1.13	0.57
	Delete	16	2	14	74	8.35	8.78	8.38	5.15	0.36
	Splice	15	2	13	96	1.06	1.69	1.10	59.70	3.79
Binary Tree Shape Analysis	InsertSorted	13	2	11	43	1.25	1.28	1.28	1.97	1.54
	Lindstrom	10	2	8	43	40.44	82.29	41.48	103.47	2.57
	DSW	10	2	8	52	101.30	180.20	109.51	77.89	8.15
	DeleteSorted	13	2	11	554	75.26	409.31	97.71	443.85	29.69
SLL Sorting	ReverseSorted	18	2	16	23	0.47	0.54	0.49	13.05	2.58
	BubbleSort	18	2	16	80	5.74	8.91	6.42	55.32	11.77
	BubbleSortBug	18	2	16	80	5.41	7.61	6.01	40.75	11.14
	InsertSortBug2	18	2	16	87	5.19	17.57	6.09	238.55	17.04
	InsertSort	18	2	16	88	5.65	18.55	6.66	228.26	17.95
	InsertSortBug1	18	2	16	88	18.94	32.93	20.25	73.84	7.27
	MergeSorted	18	2	16	91	2.26	4.22	2.53	86.35	11.46
Information Flow	Good Flow	12	2	10	66	13.59	23.28	15.37	71.30	13.59
	Bad Flow	12	2	10	86	78.05	180.85	94.92	131.70	21.79

Fig. 32. Results from using hand-crafted vs. automatically-generated maintenance formulas for instrumentation relations.

a negative unit-size change $\Delta^-[sfe_n]$ to sfe_n. In a tree-shaped graph, there exists at most one path between a pair of nodes; if that path goes through the sfe_n edge to be deleted, the corresponding sfp_n edge should be removed (cf. Formula (11)):

$$\begin{aligned} sfp'_n(v_1, v_2) = \; & sfp_n(v_1, v_2) \\ \wedge \; & \neg(\exists\, v'_1, v'_2 \colon sfp_n(v_1, v'_1) \wedge \Delta^-[sfe_n](v'_1, v'_2) \wedge sfp_n(v'_2, v_2)). \end{aligned} \tag{18}$$

6. EXPERIMENTAL EVALUATION

To evaluate the techniques presented in §4 and §5, we extended TVLA to generate relation-maintenance formulas, and applied it to a test suite of five existing analysis specifications, involving twenty-four programs, along with a variety of different abstractions and properties to check (see Fig. 32). The experiment was designed to determine what penalty is incurred when the relation-maintenance formulas generated by our finite-differencing-based algorithm are used in place of hand-crafted relation-maintenance formulas. In the experiment, we used the set of hand-crafted relation-maintenance formulas that had been built up during several years of experience with TVLA. The idea was that if the penalty is low for the programs, abstractions, and program properties in the test suite, that would be evidence that the penalty will be low for other examples—and that one can afford to forgo the effort of hand-crafting maintenance formulas for other analysis examples.

The test programs consisted of various operations on acyclic singly-linked lists, doubly-linked lists, binary trees, and binary-search trees, plus several sorting programs [Lev-Ami et al. 2000]. The system was used to verify some partial-correctness properties of the test programs. For instance, Reverse, an in-situ list-reversal program, must preserve list properties and lose no elements; InsertSorted and DeleteSorted must preserve binary-search-tree properties; InsertSort must return a sorted list; Good Flow must not allow high-security input data to flow to a low-security output channel. (Loginov et al. discuss the verification of stronger properties, such as the *partial correctness* of several of the algorithms [Loginov et al. 2005; Loginov 2006; Loginov et al. 2007].)

Lindstrom and DSW are two variants of Deutsch-Schorr-Waite, a constant-space tree-traversal algorithm that uses destructive pointer rotation. For Lindstrom and DSW, we verified that the algorithms have no unsafe pointer operations or memory leaks, and that the data structure produced at the end is, in fact, a binary tree. (Loginov et al. discuss the verification of the *total correctness* of Deutsch-Schorr-Waite—i.e., that the binary tree produced at the end is identical to the input tree and that the algorithm terminates [Loginov et al. 2006].)

A few of the programs contained bugs: for instance, InsertSortBug2 is an insert-sort program that ignores the first element of the list; BubbleBug is a bubble-sort program with an incorrect condition for swapping elements, which causes an infinite loop if the input list contains duplicate data values. (See [Lev-Ami et al. 2000; Dor et al. 2000; Lev-Ami and Sagiv 2000] for more details.)

In TVLA, the operational semantics of a programming language is defined by specifying, for each kind of statement, an *action schema* to be used on outgoing CFG edges. Action schemas are instantiated according to a program's statement instances to create the CFG. For each combination of action schema and instrumentation relation, a *maintenance-formula schema* must be provided. The number of non-identity maintenance-formula schemas is reported in columns 3–5 of Fig. 32, broken down in columns 4–5 into those whose defining formula contains an occurrence of **RTC**, and those that do not. Relation-maintenance formulas produced by finite differencing are generally larger than the hand-crafted ones. Because this affects analysis time, the number of instances of non-identity maintenance-formula schemas is a meaningful size measure for our experiments. These numbers appear in column 6. The number of instances of non-identity schemas for `DeleteSorted` is high because `DeleteSorted` includes three inline expansions of the routine that finds the tree node that takes the place of the deleted node.[17]

The data structures manipulated by all programs in our test suite are acyclic and tree-shaped, thus acyclic reachability maintenance (i.e., the techniques of §5.1), as well as tree-shaped reachability maintenance (i.e., the techniques of §5.2), apply for the maintenance of reachability relations. In the absence of hand-crafted maintenance formulas for reachability relations in possibly-cyclic linked lists, we could not extend our experiments to cover the techniques of §5.3.2. Instead, we validated those techniques as part of the verification of properties of `Reverse` when applied to possibly-cyclic linked lists (see [Loginov et al. 2007]).

For each program in the test suite, we first ran the analysis using hand-crafted maintenance formulas, to obtain a reference answer in which CFG nodes were annotated with their final sets of logical structures. We then ran the analysis using au-

[17] Work on interprocedural shape analysis provides a solution that does not require inline-expanded programs [Rinetzky and Sagiv 2001; Rinetzky et al. 2005; Jeannet et al. 2004].

tomatically generated maintenance formulas with acyclic reachability maintenance and compared the result against the reference answer. For all 24 test programs, the analysis using automatically generated formulas yielded answers identical to the reference answers. Finally, we ran the analysis using automatically generated maintenance formulas with tree-shaped reachability maintenance and compared the result against the reference answer. Again, for all 24 test programs, the analysis using automatically generated formulas yielded answers identical to the reference answers.

Columns 7–11 show performance data, which were collected on a 3GHz PC with 3.7GB of RAM running CentOS 4 Linux. The column labeled "Ref." gives the references times. Columns labeled "acyc." give the data for the analyses that used automatically generated maintenance formulas with acyclic reachability maintenance. Columns labeled "tree" give the data for the analyses that used automatically generated maintenance formulas with tree-shaped reachability maintenance. In each case, five runs were made; the longest and shortest times were discarded from each set, and the remaining three averaged. The geometric mean of the slowdowns when using the automatically generated formulas with acyclic reachability maintenance was approximately 60%, with a median of 55%, mainly due to the fact that the automatically generated formulas are larger than the hand-crafted ones. The maximum slowdown was 444%. The highest slowdowns occurred in analyses of programs that involved deletions of edges in a data structure's graph.

Because the edge-deletion maintenance formulas produced by the tree-shaped reachability-maintenance technique are much smaller than those that are produced by acyclic reachability maintenance, our expectation was that the use of tree-shaped reachability-maintenance formulas would cause a much smaller slowdown. This expectation was confirmed: the geometric mean of the slowdowns when using the automatically generated formulas with tree-shaped reachability maintenance was approximately 8%, with a median of 7%. The maximum slowdown was 30%.[18] A few analyses were actually faster with the automatically generated formulas; these speedups are either due to random variation or are accidental benefits of subformula orderings that are advantageous for short-circuit evaluation.

These results are encouraging. At least for abstractions of several common data structures, they suggest that the algorithm for generating relation-maintenance formulas from §4 and §5 is capable of automatically generating formulas that (i) are as precise as the hand-crafted ones, and (ii) have a tolerable effect on runtime performance.

The extended version of TVLA also uncovered several bugs in the hand-crafted formulas. A maintenance formula of the form $\mu_{p,st}(v_1, \ldots, v_k) = p(v_1, \ldots, v_k)$ is called an *identity relation-maintenance formula*. For each identity relation-maintenance formula in the hand-crafted specification, we checked that (after simplification) the corresponding generated relation-maintenance formula was also an identity formula. Each inconsistency turned out to be an error in the hand-crafted specification. We also found one instance of an incorrect non-identity hand-crafted maintenance formula. (The measurements reported in Fig. 32 are based on cor-

[18] We expect that some simple optimizations, such as caching the results from evaluating subformulas, could reduce the slowdown further.

rected hand-crafted specifications.)

7. RELATED WORK

A weakness of the original formulation of TVLA [Sagiv et al. 2002; Lev-Ami and Sagiv 2000] was that the user needed to define relation-maintenance formulas by hand to specify how each structure transformer affects each instrumentation relation. Past criticisms of TVLA based on this deficiency [Ball et al. 2001; Møller and Schwartzbach 2001] are no longer valid, at least for analyses that can be defined using formulas that define acyclic relations (and also for some classes of formulas that define cyclic relations). With the algorithm presented in §4 and §5, the user's responsibility is merely to write the ψ_p formulas; appropriate relation-maintenance formulas are created automatically.

Graf and Saïdi [Graf and Saïdi 1997] showed that theorem provers can be used to generate best abstract transformers [Cousot and Cousot 1979] for abstract domains that are fixed, finite, Cartesian products of Boolean values. (The use of such domains is known as *predicate abstraction*; predicate abstraction is also used in SLAM [Ball et al. 2001] and other systems [Das et al. 1999].) In contrast, the abstract transformers created using the algorithm described in §4 and §5 are not best transformers; however, this algorithm uses only very simple, linear-time, recursive tree-traversal procedures, whereas the theorem provers used in predicate abstraction are not even guaranteed to terminate. Moreover, our setting makes available much richer abstract domains than the ones offered by predicate abstraction, and experience to date has been that very little precision is lost (using only *good* abstract transformers) once the right instrumentation relations have been identified.

Paige studied how finite-differencing transformations of applicative set-former expressions could be exploited to optimize loops in very-high-level languages, such as SETL [Paige and Koenig 1982]. Liu et al. used related program-transformation methods in the setting of a functional programming language to derive incremental algorithms for various problems from the specifications of exhaustive algorithms [Liu and Teitelbaum 1995; Liu et al. 1996]. In their work, the goal is to maintain the value of a function $F(x)$ as the input x undergoes small changes. The methods described in §4 and §5 address a similar kind of incremental-computation problem, except that the language in which the exhaustive and incremental versions of the problem are expressed is first-order logic with reflexive transitive closure.

The finite-differencing operators defined in §4 and §5 are most closely related to a number of previous papers on logic and databases: finite-difference operators for the propositional case were studied by Akers [Akers 1959] and Sharir [Sharir 1982]. Previous work on incrementally maintaining materialized views in databases [Gupta and Mumick 1999], "first-order incremental evaluation schemes (FOIES)" [Dong and Su 1995], and "dynamic descriptive complexity" [Patnaik and Immerman 1997] has also addressed the problem of maintaining one or more auxiliary relations after new tuples are inserted into or deleted from the base relations. In databases, view maintenance is solely an optimization; the correct information can always be obtained by reevaluating the formula. In the abstract-interpretation context, where abstraction has been performed, this is no longer true: reevaluating a formula in the local (3-valued) state can lead to a drastic loss of precision. Thus, one aspect that sets our work apart from previous work is the goal of developing a finite-differencing

φ	$\Delta_{st}[\varphi]$
$\mathbf{1}$	$\mathbf{0}$
$\mathbf{0}$	$\mathbf{0}$
$p(w_1, \ldots, w_k), p \in \mathcal{C}$	$(\tau_{p,st} \oplus p)\{w_1, \ldots, w_k\}$
$p(w_1, \ldots, w_k), p \in \mathcal{I}$	$\Delta_{st}[\psi_p]\{w_1, \ldots, w_k\}$
$\varphi_1 \oplus \varphi_2$	$\Delta_{st}[\varphi_1] \oplus \Delta_{st}[\varphi_2]$
$\varphi_1 \wedge \varphi_2$	$(\Delta_{st}[\varphi_1] \wedge \varphi_2) \oplus (\varphi_1 \wedge \Delta_{st}[\varphi_2]) \oplus (\Delta_{st}[\varphi_1] \wedge \Delta_{st}[\varphi_2])$
$\forall v\colon \varphi_1$	$(\forall v\colon \varphi_1)\ ?\ (\exists v\colon \Delta_{st}[\varphi_1]) : (\forall v\colon \varphi_1 \oplus \Delta_{st}[\varphi_1])$

Fig. 33. An alternative finite-differencing scheme for first-order formulas.

transformation suitable for use when abstraction has been performed.

Not all finite-differencing transformations that are correct in 2-valued logic (i.e., satisfy Theorem 4.3), are appropriate for use in 3-valued logic. For instance, Fig. 33 presents an alternative finite-differencing scheme for first-order formulas. In this scheme, $\Delta_{st}[\varphi]$ captures both the negative and positive changes to φ's value. With Fig. 33, the maintenance formula for instrumentation relation p is

$$\mu_{p,st} \stackrel{\text{def}}{=} p \oplus \Delta_{st}[\psi_p], \tag{19}$$

where $\oplus$ denotes exclusive-or. However, in 3-valued logic, we have $1/2 \oplus V = 1/2$, regardless of whether V is 0, 1, or 1/2. Consequently, Eqn. (19) has the unfortunate property that if $p(u) = 1/2$, then $\mu_{p,st}$ evaluates to 1/2 on u, and $p(u)$ becomes "pinned" to the indefinite value 1/2; it will have the value 1/2 in all successor structures S', in all successors of S', and so on. With Eqn. (19), $p(u)$ can *never* reacquire a definite value.

In contrast, the maintenance formulas created using the finite-differencing scheme of Fig. 14 do not have this trouble because they have the form $p ? \neg\Delta^-_{st}[\psi_p] : \Delta^+_{st}[\psi_p]$. The use of if-then-else allows $p(u)$ to reacquire a definite value after it has been set to 1/2: if $p(u)$ is 1/2, $\mu_{p,st}$ evaluates to a definite value on u if $[\![\Delta^-_{st}[\psi_p(v)]]\!]^S_3([v \mapsto u])$ is 1 and $[\![\Delta^+_{st}[\psi_p(v)]]\!]^S_3([v \mapsto u])$ is 0, or vice versa.

In §5.3.2, we compared our work to that of W. Hesse, which is closest in spirit to our techniques for maintaining reachability information in possibly-cyclic linked lists. Below, we discuss a few approaches that bear resemblance to ours in that they attempt to translate or simulate a data structure that cannot be handled by some core techniques into one that can.

The idea of using spanning-tree representations for specifying or reasoning about data structures that are "close to trees" is not new. Klarlund and Schwartzbach introduced *graph types*, which can be used to specify some common non-tree-shaped data structures in terms of a spanning-tree *backbone* and regular expressions that specify where non-backbone edges occur within the backbone [Klarlund and Schwartzbach 1993]. Examples of data structures that can be specified by graph types are doubly-linked lists and threaded trees. A panhandle list cannot be specified by a graph type because in a graph type the location of each non-backbone edge has to be defined in terms of the backbone using a regular expression, and a regular expression cannot be used to specify the existence of a backedge to *some* node that occurs earlier in the list. In the PALE project [Møller and Schwartzbach 2001], which incorporates work on graph types, automated reasoning about programs that manipulate data structures specified as graph types can be carried out using a decision procedure for monadic second-order logic. Unfortunately, the decision procedure has non-elementary complexity. An advantage of our approach over that of PALE is that we do not rely on the use of a decision procedure.

Immerman et al. presented *structure simulation*, a technique that broadens the applicability of decision procedures to a larger class of data structures [Immerman et al. 2004]. Under certain conditions, it allows data structures that cannot be reasoned about using decidable logics to be translated into data structures that can, with the translation expressed as a first-order-logic formula. Unlike graph types, structure simulation is capable of specifying panhandle lists. However, this technique shares a limitation of graph types because it relies on decision procedures for automated reasoning about programs.

Manevich et al. specified abstractions (in canonical-abstraction and predicate-abstraction forms) for showing safety properties of programs that manipulate possibly-cyclic linked lists [Manevich et al. 2005]. By maintaining reachability within list segments that are not *interrupted* by nodes that are shared or pointed to by a variable, they are able to break the symmetry of a cycle. The definition of several key instrumentation relations in that work makes use of transitive-closure formulas that cannot be handled precisely by finite differencing. As a result, a drawback of that work is the need to define some relation-maintenance formulas by hand. Another drawback is the difficulty of reasoning about reachability (in a list) from a program variable (see reachability relations $r_{n,x}$ of Fig. 28). Because in [Manevich et al. 2005] reachability in a list has to be expressed in terms of reachability over a sequence of uninterrupted segments, a formula that expresses the reachability of node v from program variable `x` in a list has to enumerate all permutations of other program variables that may act as interruptions on a path from `x` to v in the list.

A number of past approaches to the analysis of programs that manipulate linked lists relied on first-order axiomatizations of reachability information. All of these approaches involved the use of first-order-logic decision procedures. While our approach does not have this limitation, it is instructive to compare our work with those approaches that included mechanisms for breaking the symmetry on a cycle. Nelson defined a set of first-order axioms that describe the ternary reachability relation $r_n(u, v, w)$, which has the meaning: w is reachable from u along n edges without encountering v [Nelson 1983]. The use of this relation alone is not sufficient in our setting because in the presence of abstraction we require unary distinctions (such as the relations pr_x and pr_{is} of Fig. 28) to break the symmetry. Additionally, the maintenance of ternary relations is more expensive than the maintenance of binary relations. Lahiri and Qadeer specify a collection of first-order axioms that are sufficient to verify properties of procedures that perform a single change to a cyclic list, e.g., the removal of an element [Lahiri and Qadeer 2006]. They also verify properties of in-situ list reversal, albeit under the assumption that the input list is acyclic. In a recent publication, we describe a case study in which we use the techniques developed in §5.3.2 to verify the *total correctness* (partial correctness and termination) of `Reverse` when applied to any linked list, including cyclic and panhandle lists [Loginov et al. 2007]. Lahiri and Qadeer break the symmetry of cycles in a similar fashion to how it is done in [Manevich et al. 2005]: the *blocking cells* of [Lahiri and Qadeer 2006] are a subset of the interruptions of [Manevich et al. 2005]. The blocking cells include only the set of *head variables*—program variables that act as heads of lists used in the program. This set has to be maintained carefully by the user to (i) satisfy the system's definition of acceptable (*well-founded*) lists, (ii) allow the system to verify useful postconditions, and (iii) avoid falling prey to the difficulty—that arises in [Manevich et al. 2005]—of expressing reachability in the list. The current mechanism of [Lahiri and Qadeer 2006] is insufficient for reasoning about panhandle lists because the set of blocking cells does not include shared nodes. This limitation can be partially addressed by generalizing the set of blocking cells to mimic interruptions of [Manevich et al. 2005] more faithfully. However, this may make it more difficult to satisfy points (ii) and (iii) above. As in our work, Lahiri and Qadeer rely on the insight that reachability information can be maintained in first-order logic. They use a collection of manually-specified update formulas that define how their relations are affected by the statements of the language and the (user-inserted) statements that manage the set of head variables.

8. CONCLUSIONS

This paper addresses a fundamental challenge that arises in abstract interpretation:

> Given the concrete semantics for a language and a desired abstraction,
> how does one create the associated abstract transformers?

This challenge arises in program-analysis problems in which the semantics of statements is expressed using logical formulas that describe changes to core-relation values. When instrumentation relations have been introduced to refine an abstraction, the challenge is to reflect the changes in core-relation values in the values of the instrumentation relations. The algorithm presented in this paper provides a way to create formulas that maintain correct values for the instrumentation relations, and thereby provides a way to generate—*completely automatically*—the part of the transformers of an abstract semantics that deals with instrumentation relations.

The work described in this paper opened the way for TVLA to be extended to support *automatic abstraction refinement* [Loginov et al. 2005; Loginov 2006]. The idea is to start the analyzer with a crude abstraction, and use the results of analysis runs that fail to establish a definite answer (about whether the property of interest does or does not hold) as feedback about how the abstraction should be refined.

Abstraction refinement had previously been used in the model-checking community [Kurshan 1994; Clarke et al. 2000; Ball and Rajamani 2001]; however, finding an analog of this that was suitable for TVLA was a challenging problem because canonical abstraction is considerably more sophisticated than the abstractions used by the model-checking community: in particular, predicate abstraction [Graf and Saïdi 1997] can be viewed as the degenerate case of canonical abstraction in which only nullary relations are retained [Reps et al. 2004, §4]. Because of this difference, we had to use completely different mechanisms than those used in tools such as SLAM [Ball and Rajamani 2001], BLAST [Henzinger et al. 2002], and Magic [Chaki et al. 2003]. Instead, we developed ways to use inductive logic programming to discover an appropriate set of instrumentation relations that refine the abstraction in use [Loginov et al. 2005; Loginov 2006]. Finite-differencing is a crucial enabling technique in this approach because it provides the ability to create relation-maintenance formulas *automatically* after refinement has been performed.

Finally, although the work described in the paper was motivated by a problem that arose in work on static analysis based on 3-valued logic, any method in which systems are described as evolving (2-valued or 3-valued) logical structures—e.g., Alloy [Jackson 2006] or Abstract State Machines [Boerger and Staerk 2003]—may be able to benefit from these techniques.

Acknowledgments. We are grateful to W. Hesse, N. Immerman, T. Lev-Ami and R. Wilhelm for their comments and suggestions concerning this work. R. Manevich provided invaluable help with TVLA.

REFERENCES

Akers, Jr., S. 1959. On a theory of Boolean functions. *J. Soc. Indust. Appl. Math. 7,* 4 (December), 487–498.

Ball, T., Majumdar, R., Millstein, T., and Rajamani, S. 2001. Automatic predicate abstraction of C programs. In *Prog. Lang. Design and Impl.* 203–213.

Ball, T. and Rajamani, S. 2001. Automatically validating temporal safety properties of interfaces. In *Spin Workshop.* 103–122.

Boerger, E. and Staerk, R. 2003. *Abstract State Machines: A Method for High-Level System Design and Analysis.* Springer.

Chaki, S., Clarke, E., Groce, A., Jha, S., and Veith, H. 2003. Modular verification of software components in C. In *Int. Conf. on Softw. Eng.* 385–395.

Clarke, E., Grumberg, O., Jha, S., Lu, Y., and Veith, H. 2000. Counterexample-guided abstraction refinement. In *Computer Aided Verif.* 154–169.

Cousot, P. and Cousot, R. 1979. Systematic design of program analysis frameworks. In *Princ. of Prog. Lang.* 269–282.

Das, S., Dill, D., and Park, S. 1999. Experience with predicate abstraction. In *Computer Aided Verif.* Springer-Verlag, 160–171.

Dong, G. and Su, J. 1995. Incremental and decremental evaluation of transitive closure by first-order queries. *Inf. and Comp. 120*, 101–106.

Dong, G. and Su, J. 2000. Incremental maintenance of recursive views using relational calculus/SQL. *SIGMOD Record 29,* 1, 44–51.

Dor, N., Rodeh, M., and Sagiv, M. 2000. Checking cleanness in linked lists. In *Static Analysis Symp.* 115–134.

Gopan, D., DiMaio, F., Dor, N., Reps, T., and Sagiv, M. 2004. Numeric domains with summarized dimensions. In *Tools and Algs. for the Construct. and Anal. of Syst.* 512–529.

Graf, S. and Saïdi, H. 1997. Construction of abstract state graphs with PVS. In *Computer Aided Verif.* 72–83.

Gupta, A. and Mumick, I., Eds. 1999. *Materialized Views: Techniques, Implementations, and Applications.* The M.I.T. Press, Cambridge, MA.

Henzinger, T., Jhala, R., Majumdar, R., and Sutre, G. 2002. Lazy abstraction. In *Princ. of Prog. Lang.* 58–70.

Hesse, W. 2003. Dynamic computational complexity. Ph.D. thesis, Dept. of Computer Science, University of Massachusetts.

Immerman, N., Rabinovich, A., Reps, T., Sagiv, M., and Yorsh, G. 2004. Verification via structure simulation. In *Computer Aided Verif.* 281–294.

Jackson, D. 2006. *Software Abstractions: Logic, Language, and Analysis.* The M.I.T. Press.

Jeannet, B., Loginov, A., Reps, T., and Sagiv, M. 2004. A relational approach to interprocedural shape analysis. In *Static Analysis Symp.* 246–264.

Klarlund, N. and Schwartzbach, M. 1993. Graph types. In *Princ. of Prog. Lang.* 196–205.

Kurshan, R. 1994. *Computer-Aided Verification of Coordinating Processes.* Princeton Univ. Press.

Lahiri, S. and Qadeer, S. 2006. Verifying properties of well-founded linked lists. In *Princ. of Prog. Lang.* 115–126.

Lev-Ami, T., Reps, T., Sagiv, M., and Wilhelm, R. 2000. Putting static analysis to work for verification: A case study. In *Int. Symp. on Softw. 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.

Liu, Y., Stoller, S., and Teitelbaum, T. 1996. Discovering auxiliary information for incremental computation. In *Symp. on Princ. of Prog. Lang.* 157–170.

Liu, Y. and Teitelbaum, T. 1995. Systematic derivation of incremental programs. *Sci. of Comp. Program. 24*, 1–39.

Loginov, A. 2006. Refinement-based program verification via three-valued-logic analysis. Ph.D. thesis, Comp. Sci. Dept., Univ. of Wisconsin, Madison, WI. Tech. Rep. 1574.

Loginov, A., Reps, T., and Sagiv, M. 2005. Abstraction refinement via inductive learning. In *Computer Aided Verif.* 519–533.

Loginov, A., Reps, T., and Sagiv, M. 2006. Automated verification of the Deutsch-Schorr-Waite tree-traversal algorithm. In *Static Analysis Symp.* 261–279.

Loginov, A., Reps, T., and Sagiv, M. 2007. Refinement-based verification for possibly-cyclic lists. In *Program Analysis and Compilation, Theory and Practice: Essays Dedicated to Reinhard Wilhelm*. 247–272.

Manevich, R., Yahav, E., Ramalingam, G., and Sagiv, M. 2005. Predicate abstraction and canonical abstraction for singly-linked lists. In *Verif., Model Checking, and Abs. Interp.* 181–198.

McMillan, K. 1999. Verification of infinite state systems by compositional model checking. In *Correct Hardware Design and Verification Methods (CHARME)*. 219–234.

Møller, A. and Schwartzbach, M. 2001. The pointer assertion logic engine. In *Prog. Lang. Design and Impl.* 221–231.

Nelson, G. 1983. Verifying reachability invariants of linked structures. In *Princ. of Prog. Lang.* 38–47.

Paige, R. and Koenig, S. 1982. Finite differencing of computable expressions. *Trans. on Prog. Lang. and Syst. 4,* 3 (July), 402–454.

Patnaik, S. and Immerman, N. 1997. Dyn-FO: A parallel, dynamic complexity class. *J. Comput. Syst. Sci. 55,* 2 (Oct.), 199–209.

Reps, T., Loginov, A., and Sagiv, M. 2002. Semantic minimization of 3-valued propositional formulae. In *Proc. Symp. on Logic in Comp. Sci.* 40–54.

Reps, T., Sagiv, M., and Loginov, A. 2003. Finite differencing of logical formulas for static analysis. In *European Symp. On Programming*. 380–398.

Reps, T., Sagiv, M., and Wilhelm, R. 2004. Static program analysis via 3-valued logic. In *Computer Aided Verif.* 15–30.

Rinetzky, N., Bauer, J., Reps, T., Sagiv, M., and Wilhelm, R. 2005. A semantics for procedure local heaps and its abstractions. In *Princ. of Prog. Lang.* 296–309.

Rinetzky, N. and Sagiv, M. 2001. Interprocedural shape analysis for recursive programs. In *Comp. Construct.* 133–149.

Sagiv, M., Reps, T., and Wilhelm, R. 2002. Parametric shape analysis via 3-valued logic. *Trans. on Prog. Lang. and Syst. 24,* 3, 217–298.

Sharir, M. 1982. Some observations concerning formal differentiation of set theoretic expressions. *Trans. on Prog. Lang. and Syst. 4,* 2 (April), 196–225.

TVLA. TVLA system. "www.cs.tau.ac.il/∼tvla/".

van Fraassen, B. 1966. Singular terms, truth-value gaps, and free logic. *J. Phil. 63,* 17 (Sept.), 481–495.

Yorsh, G., Reps, T., Sagiv, M., and Wilhelm, R. 2007. Logical characterizations of heap abstractions. *Trans. on Computational Logic (TOCL) 8(1).*

A. CORRECTNESS OF THE FINITE-DIFFERENCING SCHEME OF §4

The proofs in this section are by induction, using a size measure for formulas based on the process of putting φ in core normal form. Because of the assumption of no circular dependences among the definitions of instrumentation relations, φ can always be put in core normal form by repeated substitution until only core relations remain. The size measure is basically the size of φ when put in core normal form, except that each occurrence of an instrumentation relation $p(w_1, \ldots, w_k)$, $p \in \mathcal{I}$, encountered during the process is counted as being 1 larger than the size measure of $\psi_p\{w_1, \ldots, w_k\}$, the defining formula for relation p with $w_1, \ldots, w_k$ substituted for ψ_p's formal parameters. The proofs, therefore, look like standard structural-induction proofs, except that in the case for $p(w_1, \ldots, w_k)$, $p \in \mathcal{I}$, we are permitted to assume that the induction hypothesis holds for $\psi_p\{w_1, \ldots, w_k\}$.

Recall from §4 that our results are couched in terms of 2-valued logic, but by the Embedding Theorem (Theorem 2.6, [Sagiv et al. 2002, Theorem 4.9]), the relation-maintenance formulas that we define provide sound results when interpreted in 3-valued logic.

We only consider first-order formulas because the correctness of the extension of the Finite-Differencing Scheme for Reachability and Transitive Closure of has been argued in §5.

Lemma 4.2. *For every formula φ, φ_1, φ_2 and structure transformer st, the following properties hold:*[19]

(i). $\Delta^+_{st}[\varphi] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi] \wedge \neg\varphi$

(ii). $\Delta^-_{st}[\varphi] \overset{\text{meta}}{\Longleftrightarrow} \varphi \wedge \neg \boldsymbol{F}_{st}[\varphi]$

(iii). (a). $\boldsymbol{F}_{st}[\neg\varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \neg\boldsymbol{F}_{st}[\varphi_1]$
(b). $\boldsymbol{F}_{st}[\varphi_1 \vee \varphi_2] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi_1] \vee \boldsymbol{F}_{st}[\varphi_2]$
(c). $\boldsymbol{F}_{st}[\varphi_1 \wedge \varphi_2] \overset{\text{meta}}{\Longleftrightarrow} \boldsymbol{F}_{st}[\varphi_1] \wedge \boldsymbol{F}_{st}[\varphi_2]$
(d). $\boldsymbol{F}_{st}[\exists\, v\colon \varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \exists\, v\colon \boldsymbol{F}_{st}[\varphi_1]$
(e). $\boldsymbol{F}_{st}[\forall\, v\colon \varphi_1] \overset{\text{meta}}{\Longleftrightarrow} \forall\, v\colon \boldsymbol{F}_{st}[\varphi_1]$

Proof:

Atomic. For the cases $\varphi \equiv \boldsymbol{l}$, where $\boldsymbol{l} \in \{\mathbf{0}, \mathbf{1}\}$, and $\varphi \equiv (v_1 = v_2)$, $\Delta^+_{st}[\varphi] = \Delta^-_{st}[\varphi] = 0$, and (i) and (ii) follow immediately.

For $\varphi \equiv p(w_1, \ldots, w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p\ ?\ \neg\delta^-_{p,st} : \delta^+_{p,st}$

$$
\begin{aligned}
\text{(i) } \Delta^+_{st}[p(w_1, \ldots, w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (\delta^+_{p,st} \wedge \neg p)\{w_1, \ldots, w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{pmatrix} p(w_1, \ldots, w_k) \\ \quad ?\ \neg(\delta^-_{p,st} \wedge p)\{w_1, \ldots, w_k\} \\ \quad :\ (\delta^+_{p,st} \wedge \neg p)\{w_1, \ldots, w_k\} \end{pmatrix} \wedge \neg p(w_1, \ldots, w_k) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[p] \wedge \neg p)\{w_1, \ldots, w_k\} \\
&\qquad \text{(by the definitions of } \mathbf{F}_{st}[\cdot],\ \Delta^+_{st}[\cdot], \text{ and } \Delta^-_{st}[\cdot])
\end{aligned}
$$

[19] To simplify the presentation, we use $lhs \overset{\text{meta}}{\Longleftrightarrow} rhs$ and $lhs \overset{\text{meta}}{\Longrightarrow} rhs$ as shorthands for $[\![lhs]\!]^S_2(Z) = [\![rhs]\!]^S_2(Z)$ and $[\![lhs]\!]^S_2(Z) \le [\![rhs]\!]^S_2(Z)$, respectively, for any $S \in \mathcal{S}_2$ and assignment Z that is complete for lhs and rhs.

$$\begin{aligned}
\text{(ii)}\ \Delta^{-}_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (\delta^{-}_{p,st} \wedge p)\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge \begin{pmatrix} p(w_1,\ldots,w_k) \\ \quad ?\ (\delta^{-}_{p,st} \wedge p)\{w_1,\ldots,w_k\} \\ \quad :\ \neg(\delta^{+}_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \end{pmatrix} \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge \neg \begin{pmatrix} p(w_1,\ldots,w_k) \\ \quad ?\ \neg(\delta^{-}_{p,st} \wedge p)\{w_1,\ldots,w_k\} \\ \quad :\ (\delta^{+}_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \end{pmatrix} \\
&\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg \mathbf{F}_{st}[p])\{w_1,\ldots,w_k\}
\end{aligned}$$

(by the definitions of $\mathbf{F}_{st}[\cdot]$, $\Delta^{+}_{st}[\cdot]$, and $\Delta^{-}_{st}[\cdot]$)

For $\varphi \equiv p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p \vee \delta_{p,st}$ or $\delta_{p,st} \vee p$

$$\begin{aligned}
\text{(i)}\ \Delta^{+}_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (\delta_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{aligned}[t] &\quad (p(w_1,\ldots,w_k)\ ?\ \neg 0 : (\delta_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}) \\ &\wedge\ \neg p(w_1,\ldots,w_k) \end{aligned} \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[p] \wedge \neg p)\{w_1,\ldots,w_k\}
\end{aligned}$$

(by the definitions of $\mathbf{F}_{st}[\cdot]$, $\Delta^{+}_{st}[\cdot]$, and $\Delta^{-}_{st}[\cdot]$)

$$\begin{aligned}
\text{(ii)}\ \Delta^{-}_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} 0 \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge \neg p\{w_1,\ldots,w_k\} \wedge (\neg \delta_{p,st} \vee p)\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{aligned}[t] &\quad p\{w_1,\ldots,w_k\} \\ &\wedge\ \neg(p\{w_1,\ldots,w_k\} \vee (\delta_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}) \end{aligned} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{aligned}[t] &\quad p\{w_1,\ldots,w_k\} \\ &\wedge\ \neg(p(w_1,\ldots,w_k)\ ?\ \neg 0 : (\delta_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\}) \end{aligned} \\
&\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg \mathbf{F}_{st}[p])\{w_1,\ldots,w_k\}
\end{aligned}$$

(by the definitions of $\mathbf{F}_{st}[\cdot]$, $\Delta^{+}_{st}[\cdot]$, and $\Delta^{-}_{st}[\cdot]$)

For $\varphi \equiv p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, and $\tau_{p,st}$ is of the form $p \wedge \delta_{p,st}$ or $\delta_{p,st} \wedge p$

$$\begin{aligned}
\text{(i)}\ \Delta^{+}_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} 0 \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge (\delta_{p,st} \vee \neg p)\{w_1,\ldots,w_k\} \wedge \neg p\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{aligned}[t] &\quad (p(w_1,\ldots,w_k)\ ?\ (\delta_{p,st} \vee \neg p)\{w_1,\ldots,w_k\} : 0) \\ &\wedge\ \neg p\{w_1,\ldots,w_k\} \end{aligned} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{aligned}[t] &\quad (p(w_1,\ldots,w_k)\ ?\ \neg(\neg \delta_{p,st} \wedge p)\{w_1,\ldots,w_k\} : 0) \\ &\wedge\ \neg p\{w_1,\ldots,w_k\} \end{aligned} \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[p] \wedge \neg p)\{w_1,\ldots,w_k\}
\end{aligned}$$

(by the definitions of $\mathbf{F}_{st}[\cdot]$, $\Delta^{+}_{st}[\cdot]$, and $\Delta^{-}_{st}[\cdot]$)

$$\begin{aligned}
\text{(ii)}\ \Delta^-_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (\neg\delta_{p,st} \wedge p)\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{array}{l} \quad p\{w_1,\ldots,w_k\} \\ \wedge\ (p(w_1,\ldots,w_k)\ ?\ (\neg\delta_{p,st} \wedge p)\{w_1,\ldots,w_k\} : 1) \end{array} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{array}{l} \quad p\{w_1,\ldots,w_k\} \\ \wedge\ \neg(p(w_1,\ldots,w_k)\ ?\ \neg(\neg\delta_{p,st} \wedge p)\{w_1,\ldots,w_k\} : 0) \end{array} \\
&\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg\mathbf{F}_{st}[p])\{w_1,\ldots,w_k\} \\
&\qquad \text{(by the definitions of } \mathbf{F}_{st}[\cdot],\ \Delta^+_{st}[\cdot], \text{ and } \Delta^-_{st}[\cdot])
\end{aligned}$$

For $\varphi \equiv p(w_1,\ldots,w_k)$, $p \in \mathcal{C}$, but $\tau_{p,st}$ is not of the above forms

$$\begin{aligned}
\text{(i)}\ \Delta^+_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (\tau_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{pmatrix} p(w_1,\ldots,w_k) \\ \quad ?\ \neg(p \wedge \tau_{p,st})\{w_1,\ldots,w_k\} \\ \quad :\ (\tau_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \end{pmatrix} \wedge \neg p(w_1,\ldots,w_k) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[p] \wedge \neg p)\{w_1,\ldots,w_k\} \\
&\qquad \text{(by the definitions of } \mathbf{F}_{st}[\cdot],\ \Delta^+_{st}[\cdot], \text{ and } \Delta^-_{st}[\cdot])
\end{aligned}$$

$$\begin{aligned}
\text{(ii)}\ \Delta^-_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg\tau_{p,st})\{w_1,\ldots,w_k\} \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge \begin{pmatrix} p(w_1,\ldots,w_k) \\ \quad ?\ (p \wedge \neg\tau_{p,st})\{w_1,\ldots,w_k\} \\ \quad :\ \neg(\tau_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \end{pmatrix} \\
&\overset{\text{meta}}{\Longleftrightarrow} p\{w_1,\ldots,w_k\} \wedge \neg \begin{pmatrix} p(w_1,\ldots,w_k) \\ \quad ?\ \neg(p \wedge \neg\tau_{p,st})\{w_1,\ldots,w_k\} \\ \quad :\ (\tau_{p,st} \wedge \neg p)\{w_1,\ldots,w_k\} \end{pmatrix} \\
&\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg\mathbf{F}_{st}[p])\{w_1,\ldots,w_k\} \\
&\qquad \text{(by the definitions of } \mathbf{F}_{st}[\cdot],\ \Delta^+_{st}[\cdot], \text{ and } \Delta^-_{st}[\cdot])
\end{aligned}$$

For $p(w_1,\ldots,w_k)$, $p \in \mathcal{I}$,

$$\begin{aligned}
\text{(i)}\ \Delta^+_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^+_{st}[\psi_p\{w_1,\ldots,w_k\}] \\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\psi_p\{w_1,\ldots,w_k\}] \wedge \neg\psi_p\{w_1,\ldots,w_k\} \\
&\qquad \text{(by inductive hypothesis (i) for } \psi_p) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[p] \wedge \neg p)\{w_1,\ldots,w_k\} \\
&\qquad (\psi_p \text{ is the defining formula for } p)
\end{aligned}$$

$$\begin{aligned}
\text{(ii)}\ \Delta^-_{st}[p(w_1,\ldots,w_k)] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^-_{st}[\psi_p\{w_1,\ldots,w_k\}] \\
&\overset{\text{meta}}{\Longleftrightarrow} \psi_p\{w_1,\ldots,w_k\} \wedge \neg\mathbf{F}_{st}[\psi_p\{w_1,\ldots,w_k\}] \\
&\qquad \text{(by inductive hypothesis (ii) for } \psi_p) \\
&\overset{\text{meta}}{\Longleftrightarrow} (p \wedge \neg\mathbf{F}_{st}[p])\{w_1,\ldots,w_k\} \\
&\qquad (\psi_p \text{ is the defining formula for } p)
\end{aligned}$$

Not. $\varphi \equiv \neg\varphi_1$.

$$\begin{aligned}
\text{(i)}\ \Delta^{+}_{st}[\neg\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^{-}_{st}[\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} \varphi_1 \land \neg\mathbf{F}_{st}[\varphi_1] && \text{(by inductive hypothesis (ii) for } \varphi_1) \\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\neg\varphi_1] \land \neg(\neg\varphi_1) && \text{(by inductive hypothesis (iii) for } \varphi_1)
\end{aligned}$$

$$\begin{aligned}
\text{(ii)}\ \Delta^{-}_{st}[\neg\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^{+}_{st}[\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\varphi_1] \land \neg\varphi_1 && \text{(by inductive hypothesis (i) for } \varphi_1) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\neg\varphi_1) \land \neg\neg\mathbf{F}_{st}[\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} (\neg\varphi_1) \land \neg\mathbf{F}_{st}[\neg\varphi_1] && \text{(by inductive hypothesis (iii) for } \varphi_1)
\end{aligned}$$

$$\begin{aligned}
\text{(iii)}\ \mathbf{F}_{st}[\neg\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} (\neg\varphi_1)\ ?\ \neg\Delta^{-}_{st}[\neg\varphi_1] : \Delta^{+}_{st}[\neg\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} \varphi_1\ ?\ \Delta^{+}_{st}[\neg\varphi_1] : \neg\Delta^{-}_{st}[\neg\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} \varphi_1\ ?\ \Delta^{-}_{st}[\varphi_1] : \neg\Delta^{+}_{st}[\varphi_1] && \text{(by the definitions of } \Delta^{+}_{st}[\cdot] \text{ and } \Delta^{-}_{st}[\cdot]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \neg(\varphi_1\ ?\ \neg\Delta^{-}_{st}[\varphi_1] : \Delta^{+}_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \neg\mathbf{F}_{st}[\varphi_1]
\end{aligned}$$

Or. $\varphi \equiv \varphi_1 \lor \varphi_2$.

$$\begin{aligned}
\text{(i)}\ \Delta^{+}_{st}[\varphi_1 \lor \varphi_2] &\overset{\text{meta}}{\Longleftrightarrow} (\Delta^{+}_{st}[\varphi_1] \land \neg\varphi_2) \lor (\neg\varphi_1 \land \Delta^{+}_{st}[\varphi_2]) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[\varphi_1] \land \neg\varphi_1 \land \neg\varphi_2) \lor (\neg\varphi_1 \land \mathbf{F}_{st}[\varphi_2] \land \neg\varphi_2) \\
&\qquad \text{(by inductive hypothesis (i) for } \varphi_1 \text{ and } \varphi_2) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[\varphi_1] \lor \mathbf{F}_{st}[\varphi_2]) \land (\neg\varphi_1 \land \neg\varphi_2) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\mathbf{F}_{st}[\varphi_1 \lor \varphi_2]) \land \neg(\varphi_1 \lor \varphi_2) \\
&\qquad \text{(by part (iii) for } \varphi_1 \lor \varphi_2\text{, proved independently below)}
\end{aligned}$$

$$\begin{aligned}
\text{(ii)}\ \Delta^{-}_{st}[\varphi_1 \lor \varphi_2] &\overset{\text{meta}}{\Longleftrightarrow} (\Delta^{-}_{st}[\varphi_1] \land \neg\mathbf{F}_{st}[\varphi_2]) \lor (\neg\mathbf{F}_{st}[\varphi_1] \land \Delta^{-}_{st}[\varphi_2]) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\varphi_1 \land \neg\mathbf{F}_{st}[\varphi_1] \land \neg\mathbf{F}_{st}[\varphi_2]) \lor (\neg\mathbf{F}_{st}[\varphi_1] \land \neg\mathbf{F}_{st}[\varphi_2] \land \varphi_2) \\
&\qquad \text{(by inductive hypothesis (ii) for } \varphi_1 \text{ and } \varphi_2) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\varphi_1 \lor \varphi_2) \land (\neg\mathbf{F}_{st}[\varphi_1] \land \neg\mathbf{F}_{st}[\varphi_2]) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\varphi_1 \lor \varphi_2) \land \neg(\mathbf{F}_{st}[\varphi_1] \lor \mathbf{F}_{st}[\varphi_2]) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\varphi_1 \lor \varphi_2) \land \neg\mathbf{F}_{st}[\varphi_1 \lor \varphi_2] \\
&\qquad \text{(by part (iii) for } \varphi_1 \lor \varphi_2\text{, proved independently below)}
\end{aligned}$$

$$
\begin{array}{rl}
\text{(iii)}\ \mathbf{F}_{st}[\varphi_1 \vee \varphi_2] \overset{\text{meta}}{\Longleftrightarrow} & (\varphi_1 \vee \varphi_2)\ ?\ \neg\Delta^-_{st}[\varphi_1 \vee \varphi_2] : \Delta^+_{st}[\varphi_1 \vee \varphi_2] \\[1ex]
\overset{\text{meta}}{\Longleftrightarrow} & \begin{array}{l} (\varphi_1 \vee \varphi_2) \\ \ ?\ \neg \left[(\Delta^-_{st}[\varphi_1] \wedge \neg\mathbf{F}_{st}[\varphi_2]) \vee (\neg\mathbf{F}_{st}[\varphi_1] \wedge \Delta^-_{st}[\varphi_2])\right] \\ \ :\ (\Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2) \vee (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2]) \end{array} \\
& \quad \text{(by the definitions of } \Delta^+_{st}[\cdot] \text{ and } \Delta^-_{st}[\cdot]) \\[1ex]
\overset{\text{meta}}{\Longleftrightarrow} & \begin{array}{l} \left(\begin{array}{ll} & (\varphi_1 \vee \varphi_2) \\ \wedge & \neg \left[(\Delta^-_{st}[\varphi_1] \wedge \neg\mathbf{F}_{st}[\varphi_2]) \vee (\neg\mathbf{F}_{st}[\varphi_1] \wedge \Delta^-_{st}[\varphi_2])\right] \end{array} \right) \\ \vee \left(\begin{array}{ll} & \neg(\varphi_1 \vee \varphi_2) \\ \wedge & \left[(\Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2) \vee (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2])\right] \end{array} \right) \end{array} \\[1ex]
\overset{\text{meta}}{\Longleftrightarrow} & \begin{array}{l} \left(\begin{array}{ll} & (\varphi_1 \vee \varphi_2) \\ \wedge & \neg \left[(\Delta^-_{st}[\varphi_1] \wedge \neg\mathbf{F}_{st}[\varphi_2]) \vee (\neg\mathbf{F}_{st}[\varphi_1] \wedge \Delta^-_{st}[\varphi_2])\right] \end{array} \right) \\ \vee\ (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2) \vee (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2] \wedge \neg\varphi_2) \end{array} \\[1ex]
\overset{\text{meta}}{\Longleftrightarrow} & \left\{ \begin{array}{l} \quad \left(\begin{array}{ll} & (\varphi_1 \vee \varphi_2) \\ \wedge & (\neg\Delta^-_{st}[\varphi_1] \vee \mathbf{F}_{st}[\varphi_2]) \\ \wedge & (\mathbf{F}_{st}[\varphi_1] \vee \neg\Delta^-_{st}[\varphi_2]) \end{array} \right) \\ \vee\ (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2) \vee (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2] \wedge \neg\varphi_2) \end{array} \right. \\[1ex]
\overset{\text{meta}}{\Longleftrightarrow} & \left\{ \begin{array}{ll} & \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_1] \\ \vee & \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \neg\Delta^-_{st}[\varphi_2] \\ \vee & \varphi_1 \wedge \mathbf{F}_{st}[\varphi_2] \wedge \mathbf{F}_{st}[\varphi_1] \\ \vee & \varphi_1 \wedge \mathbf{F}_{st}[\varphi_2] \wedge \neg\Delta^-_{st}[\varphi_2] \\ \vee & \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_1] \\ \vee & \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \neg\Delta^-_{st}[\varphi_2] \\ \vee & \varphi_2 \wedge \mathbf{F}_{st}[\varphi_2] \wedge \mathbf{F}_{st}[\varphi_1] \\ \vee & \varphi_2 \wedge \mathbf{F}_{st}[\varphi_2] \wedge \neg\Delta^-_{st}[\varphi_2] \\ \vee & \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2 \\ \vee & \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2] \wedge \neg\varphi_2 \end{array} \right.
\end{array}
\tag{20}
$$

We consider the direction $\mathbf{F}_{st}[\varphi_1 \vee \varphi_2] \overset{\text{meta}}{\Longrightarrow} \mathbf{F}_{st}[\varphi_1] \vee \mathbf{F}_{st}[\varphi_2]$ first. We consider the ten cases that correspond to the cases that (at least) one of the ten disjuncts of Formula (20) holds. Each case that concerns a disjunct that contains $\mathbf{F}_{st}[\varphi_1]$ or $\mathbf{F}_{st}[\varphi_2]$ as a conjunct trivially implies that $\mathbf{F}_{st}[\varphi_1] \vee \mathbf{F}_{st}[\varphi_2]$ holds. We consider the remaining four cases.

$$
\begin{array}{rl}
\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \neg\Delta^-_{st}[\varphi_2] & \overset{\text{meta}}{\Longrightarrow} \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \\
& \overset{\text{meta}}{\Longrightarrow} \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \vee \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \\
& \overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{array}
$$

$$
\begin{array}{rl}
\varphi_2 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \neg\Delta^-_{st}[\varphi_2] & \overset{\text{meta}}{\Longrightarrow} \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_2] \\
& \overset{\text{meta}}{\Longrightarrow} \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_2] \vee \neg\varphi_2 \wedge \Delta^+_{st}[\varphi_2] \\
& \overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\varphi_2] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{array}
$$

$$\begin{aligned}
\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2 &\overset{\text{meta}}{\Longrightarrow} \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \\
&\overset{\text{meta}}{\Longrightarrow} \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \vee \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot]\text{)}
\end{aligned}$$

$$\begin{aligned}
\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_2] \wedge \neg\varphi_2 &\overset{\text{meta}}{\Longrightarrow} \neg\varphi_2 \wedge \Delta^+_{st}[\varphi_2] \\
&\overset{\text{meta}}{\Longrightarrow} \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_2] \vee \neg\varphi_2 \wedge \Delta^+_{st}[\varphi_2] \\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\varphi_2] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot]\text{)}
\end{aligned}$$

We consider the direction $\mathbf{F}_{st}[\varphi_1] \vee \mathbf{F}_{st}[\varphi_2] \overset{\text{meta}}{\Longrightarrow} \mathbf{F}_{st}[\varphi_1 \vee \varphi_2]$ next. Without loss of generality, assume that $\mathbf{F}_{st}[\varphi_1]$ holds. We consider two cases: $\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1]$ holds; $\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]$ holds. We show that both cases imply that a disjunct of Formula (20) holds. If a disjunct of Formula (20) holds, then $\mathbf{F}_{st}[\varphi_1 \vee \varphi_2]$ must hold because the latter holds if and only if Formula (20) holds. First, assume that $\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1]$ holds.

$$\begin{aligned}
\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_1] \qquad \text{(by inductive hypothesis (ii) for } \varphi_1\text{)} \\
&\overset{\text{meta}}{\Longrightarrow} \mathbf{F}_{st}[\varphi_1 \vee \varphi_2] \qquad \text{(the RHS above is a disjunct of Formula (20))}
\end{aligned}$$

Now, assume that $\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]$ holds. We consider two subcases: φ_2 holds; $\neg\varphi_2$ holds. Assume that φ_2 holds.

$$\begin{aligned}
\varphi_2 \wedge \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \\
&\qquad (\neg\varphi_1 \overset{\text{meta}}{\Longrightarrow} \neg\Delta^-_{st}[\varphi_1] \;\; \text{by inductive hypothesis (ii))} \\
&\overset{\text{meta}}{\Longrightarrow} \varphi_2 \wedge \neg\Delta^-_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot]\text{)} \\
&\overset{\text{meta}}{\Longrightarrow} \mathbf{F}_{st}[\varphi_1 \vee \varphi_2] \quad \text{(the RHS above is a disjunct of Formula (20))}
\end{aligned}$$

If $\neg\varphi_2$ holds (the second subcase), the result is immediate; it implies that the following disjunct of Formula (20) holds: $\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1] \wedge \neg\varphi_2$.

And. $\varphi \equiv \varphi_1 \wedge \varphi_2$. The entries for $\Delta^+_{st}[\varphi_1 \wedge \varphi_2]$ and $\Delta^-_{st}[\varphi_1 \wedge \varphi_2]$ can be derived from those for $\Delta^+_{st}[\varphi_1 \vee \varphi_2]$, $\Delta^-_{st}[\varphi_1 \vee \varphi_2]$, $\Delta^+_{st}[\neg\varphi_1]$, and $\Delta^-_{st}[\neg\varphi_1]$.

$$\begin{aligned}
\Delta^+_{st}[\varphi_1 \wedge \varphi_2] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^+_{st}[\neg(\neg\varphi_1 \vee \neg\varphi_2)] \\
&\overset{\text{meta}}{\Longleftrightarrow} \Delta^-_{st}[\neg\varphi_1 \vee \neg\varphi_2] \qquad \text{(by the definition of } \Delta^+_{st}[\cdot]\text{)} \\
&\overset{\text{meta}}{\Longleftrightarrow} (\Delta^-_{st}[\neg\varphi_1] \wedge \neg\mathbf{F}_{st}[\neg\varphi_2]) \vee (\neg\mathbf{F}_{st}[\neg\varphi_1] \wedge \Delta^-_{st}[\neg\varphi_2]) \\
&\qquad \text{(by the definition of } \Delta^-_{st}[\cdot]\text{)} \\
&\overset{\text{meta}}{\Longleftrightarrow} (\Delta^+_{st}[\varphi_1] \wedge \mathbf{F}_{st}[\varphi_2]) \vee (\mathbf{F}_{st}[\varphi_1] \wedge \Delta^+_{st}[\varphi_2]) \\
&\qquad \text{(by the definition of } \Delta^-_{st}[\cdot] \text{ and inductive hypothesis (iii))}
\end{aligned}$$

$$\begin{aligned}
\Delta^-_{st}[\varphi_1 \wedge \varphi_2] &\stackrel{\text{meta}}{\Longleftrightarrow} \Delta^-_{st}[\neg(\neg\varphi_1 \vee \neg\varphi_2)]\\
&\stackrel{\text{meta}}{\Longleftrightarrow} \Delta^+_{st}[\neg\varphi_1 \vee \neg\varphi_2] \qquad \text{(by the definition of } \Delta^-_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\Delta^+_{st}[\neg\varphi_1] \wedge \neg(\neg\varphi_2)) \vee (\neg(\neg\varphi_1) \wedge \Delta^+_{st}[\neg\varphi_2])\\
&\qquad \text{(by the definition of } \Delta^+_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\Delta^-_{st}[\varphi_1] \wedge \varphi_2) \vee (\varphi_1 \wedge \Delta^-_{st}[\varphi_2]) \qquad \text{(by the definition of } \Delta^+_{st}[\cdot])
\end{aligned}$$

Exists. $\varphi \equiv \exists\, v_1 : \varphi_1$.

$$\begin{aligned}
\text{(i)}\ \Delta^+_{st}[\exists\, v_1 : \varphi_1] &\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \varphi_1) \qquad \text{(by the definition of } \Delta^+_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \varphi_1) \wedge \neg(\exists\, v_1 : \varphi_1)\\
&\stackrel{\text{meta}}{\Longleftrightarrow} \left(\begin{array}{l} (\exists\, v_1 : \varphi_1)\\ \quad ?\ \neg\left[(\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \mathbf{F}_{st}[\varphi_1])\right]\\ \quad :\ (\exists\, v_1 : \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \varphi_1) \end{array}\right) \wedge \neg(\exists\, v_1 : \varphi_1)\\
&\stackrel{\text{meta}}{\Longleftrightarrow} \big[(\exists\, v_1 : \varphi_1)\ ?\ \neg\Delta^-_{st}[(\exists\, v_1 : \varphi_1)] : \Delta^+_{st}[(\exists\, v_1 : \varphi_1)]\big] \wedge \neg(\exists\, v_1 : \varphi_1)\\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\exists\, v_1 : \varphi_1] \wedge \neg(\exists\, v_1 : \varphi_1) \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{aligned}$$

$$\begin{aligned}
\text{(ii)}\ \Delta^-_{st}[\exists\, v_1 : \varphi_1] &\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \mathbf{F}_{st}[\varphi_1]) \qquad \text{(by the definition of } \Delta^-_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \varphi_1) \wedge (\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \mathbf{F}_{st}[\varphi_1])\\
&\qquad ((\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \stackrel{\text{meta}}{\Longrightarrow} (\exists\, v_1 : \varphi_1) \ \text{ by inductive hypothesis (ii))}\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \varphi_1) \wedge \left(\begin{array}{l} (\exists\, v_1 : \varphi_1)\\ \quad ?\ (\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \mathbf{F}_{st}[\varphi_1])\\ \quad :\ \neg\left[(\exists\, v_1 : \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \varphi_1)\right] \end{array}\right)\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \varphi_1) \wedge \neg \left(\begin{array}{l} (\exists\, v_1 : \varphi_1)\\ \quad ?\ \neg\left[(\exists\, v_1 : \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \mathbf{F}_{st}[\varphi_1])\right]\\ \quad :\ (\exists\, v_1 : \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 : \varphi_1) \end{array}\right)\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \varphi_1) \wedge \neg\big[(\exists\, v_1 : \varphi_1)\ ?\ \neg\Delta^-_{st}[(\exists\, v_1 : \varphi_1)] : \Delta^+_{st}[(\exists\, v_1 : \varphi_1)]\big]\\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot])\\
&\stackrel{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 : \varphi_1) \wedge \neg\mathbf{F}_{st}[\exists\, v_1 : \varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{aligned}$$

(iii) We consider the direction $\mathbf{F}_{st}[\exists\, v_1 \colon \varphi_1] \overset{\text{meta}}{\Longrightarrow} \exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1]$ first.

$$\begin{aligned}
\mathbf{F}_{st}[\exists\, v_1 \colon \varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 \colon \varphi_1)\ ?\ \neg\Delta^-_{st}[(\exists\, v_1 \colon \varphi_1)] : \Delta^+_{st}[(\exists\, v_1 \colon \varphi_1)] \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{cases} (\exists\, v_1 \colon \varphi_1) \\ \quad ?\ \neg\left[(\exists\, v_1 \colon \Delta^-_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1])\right] \\ \quad :\ (\exists\, v_1 \colon \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1 \colon \varphi_1) \end{cases} \\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \begin{cases} \quad (\exists\, v_1 \colon \varphi_1) \wedge \neg(\exists\, v_1 \colon \Delta^-_{st}[\varphi_1]) \\ \vee\ (\exists\, v_1 \colon \varphi_1) \wedge (\exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1]) \\ \vee\ \neg(\exists\, v_1 \colon \varphi_1) \wedge (\exists\, v_1 \colon \Delta^+_{st}[\varphi_1]) \end{cases} \qquad (21)
\end{aligned}$$

We consider the three cases that correspond to the cases that (at least) one of the three disjuncts of Formula (21) holds. The case that concerns the middle disjunct, which contains $(\exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1])$ as a conjunct, is immediate. We consider the remaining two cases. First, assume that $(\exists\, v_1 \colon \varphi_1) \wedge \neg(\exists\, v_1 \colon \Delta^-_{st}[\varphi_1])$ holds.

$$\begin{aligned}
(\exists\, v_1 \colon \varphi_1) \wedge \neg(\exists\, v_1 \colon \Delta^-_{st}[\varphi_1]) &\overset{\text{meta}}{\Longrightarrow} \exists\, v_1 \colon (\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longrightarrow} \exists\, v_1 \colon (\varphi_1\ ?\ \neg\Delta^-_{st}[\varphi_1] : \Delta^+_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{aligned}$$

Now, assume that $\neg(\exists\, v_1 \colon \varphi_1) \wedge (\exists\, v_1 \colon \Delta^+_{st}[\varphi_1])$ holds.

$$\begin{aligned}
\neg(\exists\, v_1 \colon \varphi_1) \wedge (\exists\, v_1 \colon \Delta^+_{st}[\varphi_1]) &\overset{\text{meta}}{\Longrightarrow} \exists\, v_1 \colon (\neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longrightarrow} \exists\, v_1 \colon (\varphi_1\ ?\ \neg\Delta^-_{st}[\varphi_1] : \Delta^+_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1] \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])
\end{aligned}$$

We consider the direction $\exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1] \overset{\text{meta}}{\Longrightarrow} \mathbf{F}_{st}[\exists\, v_1 \colon \varphi_1]$ next.

$$\begin{aligned}
\exists\, v_1 \colon \mathbf{F}_{st}[\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \exists\, v_1 \colon (\varphi_1\ ?\ \neg\Delta^-_{st}[\varphi_1] : \Delta^+_{st}[\varphi_1]) \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot]) \\
&\overset{\text{meta}}{\Longleftrightarrow} \exists\, v_1 \colon (\varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] \vee \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]) \\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1 \colon \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1]) \vee (\exists\, v_1 \colon \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]) \qquad (22)
\end{aligned}$$

We consider the two cases that correspond to the cases that (at least) one of the two disjuncts of Formula (22) holds. First, assume that $(\exists v_1\colon \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1])$ holds.

$$\begin{aligned}
\exists v_1\colon \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon \varphi_1) \wedge (\exists v_1\colon \varphi_1 \wedge \neg\Delta^-_{st}[\varphi_1])\\
&\overset{\text{meta}}{\Longrightarrow} (\exists v_1\colon \varphi_1)\wedge(\exists v_1\colon \varphi_1 \mathbin{?} \neg\Delta^-_{st}[\varphi_1] : \Delta^+_{st}[\varphi_1])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon\varphi_1)\wedge(\exists v_1\colon \mathbf{F}_{st}[\varphi_1]) \qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])\\
&\overset{\text{meta}}{\Longrightarrow} (\exists v_1\colon\varphi_1)\wedge\left[\neg(\exists v_1\colon\Delta^-_{st}[\varphi_1])\vee(\exists v_1\colon\mathbf{F}_{st}[\varphi_1])\right]\\
&\overset{\text{meta}}{\Longrightarrow} \begin{array}{l}(\exists v_1\colon\varphi_1)\\ \;\;?\;\; \neg\left[(\exists v_1\colon\Delta^-_{st}[\varphi_1])\wedge\neg(\exists v_1\colon \mathbf{F}_{st}[\varphi_1])\right]\\ \;\;:\;\; (\exists v_1\colon\Delta^+_{st}[\varphi_1])\wedge\neg(\exists v_1\colon\varphi_1)\end{array}\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon\varphi_1) \mathbin{?} \neg\Delta^-_{st}[(\exists v_1\colon\varphi_1)] : \Delta^+_{st}[(\exists v_1\colon\varphi_1)]\\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\exists v_1\colon\varphi_1] \qquad \text{(by the definition of }\mathbf{F}_{st}[\cdot])
\end{aligned}$$

Now, assume that $(\exists v_1\colon \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1])$ holds. We consider two subcases: $(\exists v_1\colon \varphi_1)$ holds; $\neg(\exists v_1\colon \varphi_1)$ holds. Assume that $(\exists v_1\colon \varphi_1)$ holds.

$$\begin{aligned}
(\exists v_1\colon \varphi_1) \wedge (\exists v_1\colon \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]) &\overset{\text{meta}}{\Longrightarrow} (\exists v_1\colon \varphi_1)\wedge(\exists v_1\colon \varphi_1 \mathbin{?} \neg\Delta^-_{st}[\varphi_1] : \Delta^+_{st}[\varphi_1])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon\varphi_1)\wedge(\exists v_1\colon \mathbf{F}_{st}[\varphi_1])\\
&\qquad \text{(by the definition of } \mathbf{F}_{st}[\cdot])\\
&\overset{\text{meta}}{\Longrightarrow} (\exists v_1\colon\varphi_1)\wedge\left[\neg(\exists v_1\colon\Delta^-_{st}[\varphi_1])\vee(\exists v_1\colon\mathbf{F}_{st}[\varphi_1])\right]\\
&\overset{\text{meta}}{\Longrightarrow} \begin{array}{l}(\exists v_1\colon\varphi_1)\\ \;\;?\;\; \neg\left[(\exists v_1\colon\Delta^-_{st}[\varphi_1])\wedge\neg(\exists v_1\colon \mathbf{F}_{st}[\varphi_1])\right]\\ \;\;:\;\; (\exists v_1\colon\Delta^+_{st}[\varphi_1])\wedge\neg(\exists v_1\colon\varphi_1)\end{array}\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon\varphi_1) \mathbin{?} \neg\Delta^-_{st}[(\exists v_1\colon\varphi_1)] : \Delta^+_{st}[(\exists v_1\colon\varphi_1)]\\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\exists v_1\colon\varphi_1] \qquad \text{(by the definition of }\mathbf{F}_{st}[\cdot])
\end{aligned}$$

Assume that $\neg(\exists v_1\colon \varphi_1)$ holds (the second subcase).

$$\begin{aligned}
\neg(\exists v_1\colon \varphi_1) \wedge (\exists v_1\colon \neg\varphi_1 \wedge \Delta^+_{st}[\varphi_1]) &\overset{\text{meta}}{\Longrightarrow} \neg(\exists v_1\colon \varphi_1)\wedge(\exists v_1\colon \Delta^+_{st}[\varphi_1])\\
&\overset{\text{meta}}{\Longleftrightarrow} \neg(\exists v_1\colon\varphi_1)\wedge(\exists v_1\colon \Delta^+_{st}[\varphi_1])\wedge\neg(\exists v_1\colon\varphi_1)\\
&\overset{\text{meta}}{\Longrightarrow} \begin{array}{l}(\exists v_1\colon\varphi_1)\\ \;\;?\;\; \neg\left[(\exists v_1\colon\Delta^-_{st}[\varphi_1])\wedge\neg(\exists v_1\colon \mathbf{F}_{st}[\varphi_1])\right]\\ \;\;:\;\; (\exists v_1\colon\Delta^+_{st}[\varphi_1])\wedge\neg(\exists v_1\colon\varphi_1)\end{array}\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists v_1\colon\varphi_1) \mathbin{?} \neg\Delta^-_{st}[(\exists v_1\colon\varphi_1)] : \Delta^+_{st}[(\exists v_1\colon\varphi_1)]\\
&\qquad \text{(by the definitions of } \Delta^-_{st}[\cdot] \text{ and } \Delta^+_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} \mathbf{F}_{st}[\exists v_1\colon\varphi_1] \qquad \text{(by the definition of }\mathbf{F}_{st}[\cdot])
\end{aligned}$$

Forall. $\varphi \equiv \forall v_1\colon \varphi_1$. The entries for $\Delta^+_{st}[\forall v_1\colon \varphi_1]$ and $\Delta^-_{st}[\forall v_1\colon \varphi_1]$ can be derived from those for $\Delta^+_{st}[\exists\, v_1\colon \varphi_1]$, $\Delta^-_{st}[\exists\, v_1\colon \varphi_1]$, $\Delta^+_{st}[\neg\varphi_1]$, and $\Delta^-_{st}[\neg\varphi_1]$.

$$\begin{aligned}
\Delta^+_{st}[\forall\, v_1\colon \varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^+_{st}[\neg(\exists\, v_1\colon \neg\varphi_1)]\\
&\overset{\text{meta}}{\Longleftrightarrow} \Delta^-_{st}[\exists\, v_1\colon \neg\varphi_1] \qquad \text{(by the definition of } \Delta^+_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1\colon \Delta^-_{st}[\neg\varphi_1]) \wedge \neg(\exists\, v_1\colon \mathbf{F}_{st}[\neg\varphi_1]) \qquad \text{(by the definition of } \Delta^-_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1\colon \Delta^+_{st}[\varphi_1]) \wedge \neg(\exists\, v_1\colon \neg\mathbf{F}_{st}[\varphi_1])\\
&\qquad \text{(by the definition of } \Delta^-_{st}[\cdot] \text{ and inductive hypothesis (iii))}\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1\colon \Delta^+_{st}[\varphi_1]) \wedge (\forall\, v_1\colon \mathbf{F}_{st}[\varphi_1])
\end{aligned}$$

$$\begin{aligned}
\Delta^-_{st}[\forall\, v_1\colon \varphi_1] &\overset{\text{meta}}{\Longleftrightarrow} \Delta^-_{st}[\neg(\exists\, v_1\colon \neg\varphi_1)]\\
&\overset{\text{meta}}{\Longleftrightarrow} \Delta^+_{st}[\exists\, v_1\colon \neg\varphi_1] \qquad \text{(by the definition of } \Delta^-_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1\colon \Delta^+_{st}[\neg\varphi_1]) \wedge \neg(\exists\, v_1\colon \neg\varphi_1) \qquad \text{(by the definition of } \Delta^+_{st}[\cdot])\\
&\overset{\text{meta}}{\Longleftrightarrow} (\exists\, v_1\colon \Delta^-_{st}[\varphi_1]) \wedge (\forall\, v_1\colon \varphi_1) \qquad \text{(by the definition of } \Delta^+_{st}[\cdot])
\end{aligned}$$

□

THEOREM 4.3. *Let S be a structure in $\mathcal{S}_2$, and let S'_{proto} be the proto-structure for statement st obtained from S. Let S' be the structure obtained by using S'_{proto} as the first approximation to S' and then filling in instrumentation relations in a topological ordering of the dependences among them: for each arity-k relation $p \in \mathcal{I}$, $\iota^{S'}(p)$ is obtained by evaluating $[\![\psi_p(v_1, \ldots, v_k)]\!]^{S'}_2([v_1 \mapsto u'_1, \ldots, v_k \mapsto u'_k])$ for all tuples $(u'_1, \ldots, u'_k) \in (U^{S'})^k$. Then for every formula $\varphi(v_1, \ldots, v_k)$ and complete assignment Z for $\varphi(v_1, \ldots, v_k)$,*

$$[\![\boldsymbol{F}_{st}[\varphi(v_1, \ldots, v_k)]]\!]^S_2(Z) = [\![\varphi(v_1, \ldots, v_k)]\!]^{S'}_2(Z)$$

PROOF. *The proof is by induction on the size of φ. Let Z be $[v_1 \mapsto u_1, \ldots, v_k \mapsto u_k]$. By Lemma 4.2(iii) and the induction hypothesis, we need only consider the cases for atomic formulas.*

(1) For $\varphi \equiv \boldsymbol{l}$, where $\boldsymbol{l} \in \{\mathbf{0}, \mathbf{1}\}$,

$$\begin{aligned}
[\![\boldsymbol{F}_{st}[\boldsymbol{l}]]\!]^S_2(Z) &= [\![\boldsymbol{l} \;?\; \neg\Delta^-_{st}[\boldsymbol{l}] : \Delta^+_{st}[\boldsymbol{l}]]\!]^S_2(Z)\\
&= [\![\boldsymbol{l} \;?\; \neg\mathbf{0} : \mathbf{0}]\!]^S_2(Z)\\
&= [\![\boldsymbol{l}]\!]^S_2(Z)\\
&= l\\
&= [\![\boldsymbol{l}]\!]^{S'}_2(Z)
\end{aligned}$$

(2) For $\varphi \equiv (v_{i_1} = v_{i_2})$,

$$\begin{aligned}
[\![\boldsymbol{F}_{st}[v_{i_1} = v_{i_2}]]\!]^S_2(Z) &= [\![v_{i_1} = v_{i_2} \;?\; \neg\Delta^-_{st}[v_{i_1} = v_{i_2}] : \Delta^+_{st}[v_{i_1} = v_{i_2}]]\!]^S_2(Z)\\
&= [\![v_{i_1} = v_{i_2} \;?\; \neg\mathbf{0} : \mathbf{0}]\!]^S_2(Z)\\
&= [\![v_{i_1} = v_{i_2}]\!]^S_2(Z)\\
&= Z(v_{i_1}) = Z(v_{i_2})\\
&= [\![v_{i_1} = v_{i_2}]\!]^{S'}_2(Z)
\end{aligned}$$

(*3*) *For* $\varphi \equiv p(v_{i_1}, \dots, v_{i_k})$, $p \in \mathcal{C}$, *and* $\tau_{p,st}$ *is of the form* $p \mathbin{?} \neg\delta^-_{p,st} : \delta^+_{p,st}$

$$\begin{aligned}
\llbracket \boldsymbol{F}_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) &= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg\Delta^-_{st}[p(v_{i_1}, \dots, v_{i_k})] : \Delta^+_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) \\
&= \left\llbracket \begin{array}{l} p(v_{i_1}, \dots, v_{i_k}) \\ \quad ? \; \neg(\delta^-_{p,st} \wedge p)\{v_{i_1}, \dots, v_{i_k}\} \\ \quad : \; (\delta^+_{p,st} \wedge \neg p)\{v_{i_1}, \dots, v_{i_k}\} \end{array} \right\rrbracket_2^S (Z) \\
&= \llbracket (p \mathbin{?} \neg\delta^-_{p,st} : \delta^+_{p,st})\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket \tau_{p,st}(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^{S'}(Z)
\end{aligned}$$

(*4*) *For* $\varphi \equiv p(w_1, \dots, w_k)$, $p \in \mathcal{C}$, *and* $\tau_{p,st}$ *is of the form* $p \vee \delta_{p,st}$ *or* $\delta_{p,st} \vee p$

$$\begin{aligned}
\llbracket \boldsymbol{F}_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) &= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg\Delta^-_{st}[p(v_{i_1}, \dots, v_{i_k})] : \Delta^+_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg 0 : (\delta_{p,st} \wedge \neg p)\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket (p \vee \neg p \wedge \delta_{p,st})\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket (p \vee \delta_{p,st})\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket \tau_{p,st}(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^{S'}(Z)
\end{aligned}$$

(*5*) *For* $\varphi \equiv p(w_1, \dots, w_k)$, $p \in \mathcal{C}$, *and* $\tau_{p,st}$ *is of the form* $p \wedge \delta_{p,st}$ *or* $\delta_{p,st} \wedge p$

$$\begin{aligned}
\llbracket \boldsymbol{F}_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) &= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg\Delta^-_{st}[p(v_{i_1}, \dots, v_{i_k})] : \Delta^+_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg(\neg\delta_{p,st} \wedge p)\{v_{i_1}, \dots, v_{i_k}\} : 0 \rrbracket_2^S(Z) \\
&= \llbracket (p \wedge \neg(\neg\delta_{p,st} \wedge p))\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket (p \wedge (\delta_{p,st} \vee \neg p))\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket (p \wedge \delta_{p,st})\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket \tau_{p,st}(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^{S'}(Z)
\end{aligned}$$

(*6*) *For* $\varphi \equiv p(v_{i_1}, \dots, v_{i_k})$, $p \in \mathcal{C}$, *but* $\tau_{p,st}$ *is not of the above forms*

$$\begin{aligned}
\llbracket \boldsymbol{F}_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) &= \llbracket p(v_{i_1}, \dots, v_{i_k}) \mathbin{?} \neg\Delta^-_{st}[p(v_{i_1}, \dots, v_{i_k})] : \Delta^+_{st}[p(v_{i_1}, \dots, v_{i_k})] \rrbracket_2^S(Z) \\
&= \left\llbracket \begin{array}{l} p(v_{i_1}, \dots, v_{i_k}) \\ \quad ? \; \neg(p \wedge \neg\tau_{p,st})\{v_{i_1}, \dots, v_{i_k}\} \\ \quad : \; (\tau_{p,st} \wedge \neg p)\{v_{i_1}, \dots, v_{i_k}\} \end{array} \right\rrbracket_2^S (Z) \\
&= \llbracket ((p \wedge \neg p) \vee (p \wedge \tau_{p,st}) \vee (\tau_{p,st} \wedge \neg p))\{v_{i_1}, \dots, v_{i_k}\} \rrbracket_2^S(Z) \\
&= \llbracket \tau_{p,st}(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^S(Z) \\
&= \llbracket p(v_{i_1}, \dots, v_{i_k}) \rrbracket_2^{S'}(Z)
\end{aligned}$$

(7) *For* $\varphi \equiv p(v_{i_1}, \ldots, v_{i_k})$, $p \in \mathcal{I}$,

$$\begin{aligned}
[\![\boldsymbol{F}_{st}[p(v_{i_1}, \ldots, v_{i_k})]]\!]_2^S(Z) &= [\![p(v_{i_1}, \ldots, v_{i_k}) \mathbin{?} \neg\Delta^-_{st}[p(v_{i_1}, \ldots, v_{i_k})] : \Delta^+_{st}[p(v_{i_1}, \ldots, v_{i_k})]]\!]_2^S(Z) \\
&= \left[\begin{matrix} p(v_{i_1}, \ldots, v_{i_k}) \\ \quad ?\neg\Delta^-_{st}[\psi_p]\{v_{i_1}, \ldots, v_{i_k}\} : \Delta^+_{st}[\psi_p]\{v_{i_1}, \ldots, v_{i_k}\} \end{matrix}\right]_2^S(Z) \\
&= \left[\begin{matrix} \psi_p\{v_{i_1}, \ldots, v_{i_k}\} \\ \quad ?\neg\Delta^-_{st}[\psi_p]\{v_{i_1}, \ldots, v_{i_k}\} : \Delta^+_{st}[\psi_p]\{v_{i_1}, \ldots, v_{i_k}\} \end{matrix}\right]_2^S(Z) \\
&= [\![\boldsymbol{F}_{st}[\psi_p]\{v_{i_1}, \ldots, v_{i_k}\}]\!]_2^S(Z) \\
&= [\![\psi_p\{v_{i_1}, \ldots, v_{i_k}\}]\!]_2^{S'}(Z) \\
&= [\![p(v_{i_1}, \ldots, v_{i_k})]\!]_2^{S'}(Z)
\end{aligned}$$

□