

Graphs and Decidable Transductions based on Edge Constraints (Extended Abstract)

Nils Klarlund* & Michael I. Schwartzbach**

Aarhus University, Department of Computer Science,
Ny Munkegade, DK-8000 Århus, Denmark
{klarlund,mis}@daimi.aau.dk

Abstract. We give examples to show that not even **c-edNCE**, the most general known notion of context-free graph grammar, is suited for the specification of some common data structures.

To overcome this problem, we use monadic second-order logic and introduce *edge constraints* as a new means of specifying a large class of graph families. Our notion stems from a natural dichotomy found in programming practice between ordinary pointers forming spanning trees and auxiliary pointers cutting across.

Our main result is that for certain transformations of graphs definable in monadic second-order logic, the question of whether a graph family given by a specification \mathcal{A} is mapped to a family given by a specification \mathcal{B} is decidable. Thus a decidable Hoare logic arises.

1 Introduction

Graphs are complicated objects to describe. Thus various grammars and logics have emerged for their representation, see the chapter by Courcelle [1]. The *monadic second-order logic of graphs* (M2L-G) allows a very large class of graph families to be described. The first-order terms of the logic denote nodes. The second-order terms denote sets of nodes. Nodes and edges are related by built-in predicates. The M2L-G formalism is very well-suited for describing properties of some common data structures, see our earlier paper [5].

Some authors consider logics that comprise quantification over edges. For these logics, a fundamental result is that a family of graphs allows a decidable M2L if and only if the family is specified by a *hyperedge-replacement grammar* [2]. Such grammars constitute a natural generalization of context-free grammars for string languages.

* The author is supported by a fellowship from the Danish Research Council.

** The author is partially supported by the BRICS Center under the Danish Research Foundation.

An even larger class of context-free grammars is known as **c-edNCE**. The monadic logic of graph families thus given is undecidable, but certain other questions, such as non-emptiness of a specification, are decidable, see [4].

For programming purposes, we would like to describe common data structures found in the store such as trees and doubly-linked lists. Indeed, this is possible within the framework of decidable formalisms as e.g. hyperedge-replacement grammars. Many other graph shapes are not representable. But whatever specification formalism we choose, we should be able to represent trees with additional, unconstrained pointers—reflecting a situation where almost nothing is said about the store, as is the case with type systems of most imperative programming languages.

We show in this paper that not even **c-edNCE** grammars are able to define such families of graphs.

To reason about data structures, it is vital to model the execution of programs. Therefore, we must formulate ways of transforming graphs corresponding to statements in a programming language. For program correctness, we would use Hoare logic to show that the store transformations leave the graph specifications satisfied.

In this paper we consider restricted graph transformations, called *transductions*, which are based on the method of *semantic interpretation* [7] and studied in [3]. Given logical graph specifications \mathcal{A} and \mathcal{B} and a transduction, we address the problem of verifying what we call *transductional correctness*: for any graph satisfying \mathcal{A} , any graph resulting from the transduction satisfies \mathcal{B} . This informal definition omits the difficulty of having shared logical variables in \mathcal{A} and \mathcal{B} —a problem that is explicitly solved in this paper. Decidability of transductional correctness amounts to decidability of the corresponding Hoare logic.

Contributions of this paper

We devise a class of graph specifications

- that may model loosely restrained edges, and
- for which transductional correctness is decidable.

Our graphs consist of *ordinary edges* constituting an underlying spanning forest, called the *backbone*, and *auxiliary edges* cutting across the backbone.

These notions stem from a natural dichotomy found in programming practice between ordinary pointers forming spanning trees and auxiliary pointers cutting across as used for short-cuts (such as extra links pointing backward to previous elements) or for indexing into other data structures using unrestrained pointers.

Our graph specifications are based on combining the full M2L in form of a *backbone formula* for specifying ordinary edges together with a special M2L syntax, called *edge constraints*, for specifying auxiliary edges. The formulas in an edge constraint involve only the backbone to specify the sources and destinations of auxiliary edges. The resulting class of graph families thus definable is called **EC**. We show that the classes **c-edNCE** and **EC** are incomparable.

We next introduce a class of transductions. They are formulated in M2L and are similar to the ones considered in [3]. We use extra logical variables to model edges that are followed, deleted, or added during the transformation of the graph.

Our main result is that the transduction problem is decidable for **EC**. This result is based on a rather complicated encoding of the effects of the transduction within M2L on the backbone alone. The obstacle that we overcome is that it is impossible to directly represent all auxiliary edges in the logic of the backbone. The key idea is to distinguish between the bounded number of auxiliary edges that are explicitly manipulated by the transduction and the others, which are represented by a universal quantification in the logic.

Our other work

In an accompanying paper [6], we outline a typing system for data structures and define a programming language. The typing information is expressed in a logic on the underlying recursive data types. The programming language provides assignment, dereference, allocation, deallocation, and limited forms of iterations based on regular walks. We show in [6] that the operational semantics is captured by transductions and that by the results in this paper the resulting Hoare logic on data structures is decidable.

In [5], we also used monadic second-order logic to reason about data structures as graphs, but we restricted ourselves to trees with auxiliary edges that are functionally determined by the backbone in terms of regular walks.

2 Rooted Graphs

A *graph alphabet* Λ consists of a finite set $\Lambda^{\mathbf{V}}$ of node labels (which include a special label **spare**) and a finite set $\Lambda^{\mathbf{E}}$ of edge labels. Usually, we denote a node label by v . There are two kinds of edge labels: *ordinary* and *auxiliary*. Usually, an ordinary edge label is denoted f and an auxiliary edge label is denoted a . An edge label that is either ordinary or auxiliary is denoted n .

A *rooted graph* G over Λ consists of a finite set $G^{\mathbf{V}}$ of labeled nodes; a finite set $G^{\mathbf{E}}$ of labeled edges; and a finite set of node variables x , called *roots*, denoting nodes in G . The label of node $v \in G^{\mathbf{V}}$ is denoted $G^{\mathbf{L}}(v)$. Nodes are either *ordinary* or *spare* according to their label. An edge from v to w labeled n is denoted (v, n, w) . For each v and n , there is at most one such edge. Loops are allowed. The edges of G are divided into *ordinary* and *auxiliary* ones according to their label. The node denoted by root x is written x^G .

The set of all graphs over Λ is denoted $\mathbf{GR}(\Lambda)$. An *edge set* \mathcal{E} is a set of edges such that $(v, n, w) \in \mathcal{E}$ and $(v, n, u) \in \mathcal{E}$ implies $w = u$.

We sometimes view G as consisting of \overline{G} , called the *backbone*, which is all of G except for the auxiliary edges, and \underline{G} , which is the edge set of auxiliary edges in G . Thus, G may be written as $(\overline{G}, \underline{G})$.

The spare nodes model free memory cells in programming language applications. They are essential to allow addition and deletion of nodes by transductions.

Figure 1 shows a sketch of a rooted graph. The ordinary edges are drawn as solid arrows, whereas the auxiliary edges are dashed; spare nodes are black; the roots are called x_1 , x_2 , and x_3 .

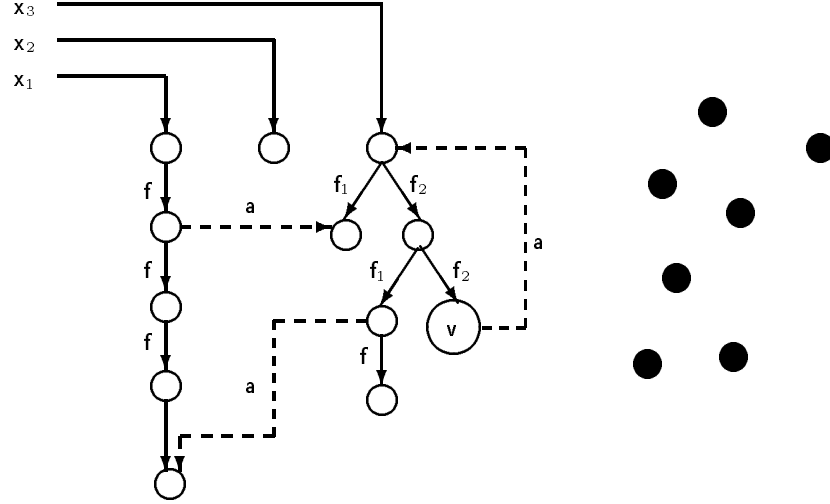


Fig. 1. A rooted graph.

3 The Logic M2L-BB

The key to specifying data structures is the *Monadic Second-Order of Backbones*, abbreviated *M2L-BB*. First-order terms range over nodes in the graph. Second-order terms range over *sets* of nodes.

Syntax

Assume a graph alphabet A . The logic of rooted graphs over A is denoted $M2L-BB(A)$. Its syntax is as follows.

Address terms \mathbf{A} denote nodes in the graph.

$\mathbf{A} ::=$	\mathbf{x}	root
	\mathbf{src}	source
	\mathbf{dst}	destination
	α, β, \dots	first-order variable

The terms \mathbf{src} and \mathbf{dst} are special variables used in certain assertions. *Address set terms* Σ denote sets of nodes.

$\Sigma ::= \emptyset$	empty set
$\Sigma_1 \cup \Sigma_2$	set union
$\Sigma_1 \setminus \Sigma_2$	set difference
S, T, \dots	second-order variable

Formulas Φ denote **true** or **false**.

$\Phi ::= A_1 = A_2$	equality
$A \in \Sigma$	set membership
$\Sigma_1 \subseteq \Sigma_2$	set inclusion
$A_1 \xrightarrow{f} A_2$	successor relation, where $f \in A^{\mathbf{E}}$ is ordinary
$v?A$	test for node label, where $v \in A^{\mathbf{V}}$
$\neg \Phi$	negation
$\Phi_1 \wedge \Phi_2$	conjunction
$\exists^0 \alpha : \Phi$	first-order quantification over all nodes
$\exists^0 S : \Phi$	second-order quantification over all nodes

Note that the syntax does not allow references to auxiliary edges. We also use unmarked quantifiers that range only over ordinary nodes. They can be viewed as abbreviations according to the following.

$$\begin{aligned} \exists \alpha : \Phi &\equiv \exists^0 \alpha : \neg \mathbf{spare?} \alpha \wedge \Phi \\ \exists S : \Phi &\equiv \exists^0 S : (\neg \exists^0 \alpha : \alpha \in S \wedge \mathbf{spare?} \alpha) \wedge \Phi \end{aligned}$$

We also assume abbreviations $\forall, \Rightarrow, \vee$, etc.

Semantics

M2L-BB is interpreted relative to a backbone \overline{G} . The interpretation of x is given by \overline{G} as $x^{\overline{G}}$. The constants **dst** and **src** are used as variables. The semantics of variables is formulated below by substitution for values in $\overline{G}^{\mathbf{V}}$. A value v is interpreted as itself, i.e. $v^{\overline{G}} = v$. A non-variable address set term Σ is interpreted as follows.

$$\begin{aligned} \emptyset^{\overline{G}} &= \emptyset \\ (\Sigma_1 \cup \Sigma_2)^{\overline{G}} &= \Sigma_1^{\overline{G}} \cup \Sigma_2^{\overline{G}} \\ (\Sigma_1 \setminus \Sigma_2)^{\overline{G}} &= \Sigma_1^{\overline{G}} \setminus \Sigma_2^{\overline{G}} \end{aligned}$$

The semantics of formulas is as follows.

$$\begin{aligned}
\overline{G} \models A_1 = A_2 & \text{ if } A_1^{\overline{G}} = A_2^{\overline{G}} \\
\overline{G} \models A \in \Sigma & \text{ if } A^{\overline{G}} \in \Sigma^{\overline{G}} \\
\overline{G} \models \Sigma_1 \subseteq \Sigma_2 & \text{ if } \Sigma_1^{\overline{G}} \subseteq \Sigma_2^{\overline{G}} \\
\overline{G} \models A_1 \xrightarrow{f} A_2 & \text{ if } (A_1^{\overline{G}}, f, A_2^{\overline{G}}) \in \overline{G}^E \\
\overline{G} \models v?A & \text{ if } \overline{G}^L(A^{\overline{G}}) = v \\
\overline{G} \models \neg \Phi & \text{ if not } \overline{G} \models \Phi \\
\overline{G} \models \Phi_1 \wedge \Phi_2 & \text{ if } \overline{G} \models \Phi_1 \text{ and } \overline{G} \models \Phi_2 \\
\overline{G} \models \exists^\circ \alpha : \Phi & \text{ if there is } v \in \overline{G}^V \text{ such that } \overline{G} \models \Phi(\alpha \mapsto v) \\
\overline{G} \models \exists^\circ S : \Phi & \text{ if there is } V \subseteq \overline{G}^V \text{ such that } \overline{G} \models \Phi(S \mapsto V),
\end{aligned}$$

If Φ has free variables \mathfrak{F} and \mathfrak{I} is an interpretation of these variables in \overline{G}^V , then

$$\overline{G}, \mathfrak{I} \models \Phi \text{ if } \overline{G} \models \Phi(\mathfrak{F} \mapsto \mathfrak{I}).$$

If $\overline{G} \models \Phi$ holds for all \overline{G} , then we say that Φ is *valid* and we write $\models \Phi$. A graph G is *tree-formed* if

- all edges are between ordinary nodes; and
- the graph induced by ordinary nodes and ordinary edges is a directed forest such that each root is the value of some root variable.

Note that the graph depicted in Figure 1 is tree-formed.

Lemma 1. *There is a formula Φ such that G is tree-formed if and only if $G \models \Phi$.*

Proof Among other conditions, acyclicity and reachability can be encoded in M2L-BB. \square

We say that Φ is *tree-valid* and we write $\models \Phi$ if $\overline{G} \models \Phi$ holds for all tree-formed \overline{G} .

Theorem 2. *Validity is undecidable, but tree-validity is decidable.*

Proof The first result follows from the undecidability of the first-order logic of finite graphs. The second result follows from the decidability of the monadic second-order logic of finite trees. \square

Edge Constraints and Assertions

Constraints on auxiliary edges cannot just be formulas, since the logic refers only to ordinary edges. Instead, an *edge constraint* is of the form $[\sigma \xrightarrow{a} \delta]$, where σ is a formula involving **src** as a free variable, and δ is a formula with free variables **src** and **dst**. The edge constraint is *valid* for a given graph if whenever σ is

valid with a node v in place of **src**, then there is an **a**-edge (which is unique by definition of a rooted graph) from v to some node w and δ is valid with v and w in place of **src** and **dst**. Note that the edge constraint does not describe any **a**-edges outside where σ holds.

Formally, let $[\sigma \xrightarrow{a} \delta]$ be an edge constraint with free variables \mathfrak{F} . We say that G and \mathfrak{A} *satisfy* $[\sigma \xrightarrow{a} \delta]$, and we write $G, \mathfrak{A} \models [\sigma \xrightarrow{a} \delta]$ if:

for all $v \in G^V$, $G, \mathfrak{A} \models \sigma(\mathbf{src} \mapsto v)$ implies
for some $(v, a, w) \in \overline{G}$, $G, \mathfrak{A} \models \delta(\mathbf{src} \mapsto v, \mathbf{dst} \mapsto w)$.

An *assertion* $\mathcal{A} = \Phi[\sigma_1 \xrightarrow{a_1} \delta_1] \dots [\sigma_n \xrightarrow{a_n} \delta_n]$ consists of a formula Φ , called the *backbone formula*, and a number of edge constraints $[\sigma_i \xrightarrow{a_i} \delta_i]$. These components are connected through free variables, which are implicitly existentially quantified.

Let \mathfrak{F} be a list containing the free variables and let \mathfrak{A} be a value assignment to these variables. An assertion \mathcal{A} is *satisfied* in G with \mathfrak{A} , and we write $G, \mathfrak{A} \models \mathcal{A}$, if $\overline{G}, \mathfrak{A} \models \Phi$ and for all i , $G, \mathfrak{A} \models [\sigma_i \xrightarrow{a_i} \delta_i]$.

An assertion \mathcal{A} specifies the language of graphs

$$\{G \mid G \text{ is tree-formed and for some } \mathfrak{A}, G, \mathfrak{A} \models \mathcal{A}\}$$

The class of such graph languages is called **EC**.

Example

Consider the common data structure, shown in Figure 2, of linked lists with a head node that points both to the first element of the list and to some designated element. The **f**- and **n**-edges are ordinary; the **s**-edge is auxiliary.




Fig. 2. A list structure

The corresponding backbone formula contains these clauses.

$H?x$	<i>The head node has label H</i>
$\exists \alpha : x \xrightarrow{f} \alpha$	<i>and an outgoing f-edge;</i>
$\forall \beta, \beta' : \beta \xrightarrow{f} \beta' \Rightarrow \beta = x$	<i>no other node has an outgoing f-edge;</i>
$\forall \beta : \neg \beta = x \Rightarrow L?\beta$	<i>all other nodes have label L;</i>
$\forall \beta, \beta' : \beta \xrightarrow{n} \beta' \Rightarrow \beta \neq x$	<i>the head node has no outgoing n-edge;</i>
$L?\gamma$	<i>and there is a designated L-node...</i>

Note that we quantify only over ordinary nodes. There is only a single edge constraint.

$$[H?src \xrightarrow{s} \gamma = dst] \quad \text{that is the destination of the } s\text{-edge.}$$

Here the free variable γ connects the backbone formula and the edge constraint. In conjunction with the general requirement of tree-formedness, this assertion describes backbones that are lists with a head node. Note that the assertion does not eliminate extraneous s -edges from nodes other than the one marked H . In a programming language application these are avoided through elementary type-checking of the transductions that build graphs [6].

4 Relations to Other Formalisms

It is interesting to compare the expressive power of this graph specification formalism with those of other proposals. In particular we show in this section that the set of trees with unrestrained auxiliary edges is not representable as a context-free graph grammar.

We look at the most general class known of context-free graphs languages: **c-edNCE**, which stands for “confluent edge and node labeled, directed graphs given by Neighborhood Controlled Embedding.” The grammars that define such languages are complicated. Instead we shall use a result by Engelfriet that these languages are exactly the images of trees under functions definable in monadic second-order logic [4]. The following definition is from [4] (but changed as to allow loops in graphs):

Let Λ_1 and Λ_2 be alphabets. An *M2L-definable function* $f : \mathbf{GR}(\Lambda_1) \rightarrow \mathbf{GR}(\Lambda_2)$ is given by the following formulas in $\mathbf{M2L-BB}(\Lambda_1)$:

- a closed formula ϕ_{dom} , called the *domain formula*;
- for every $v \in \Lambda_2^V$, a formula ϕ_v , called a *node formula*, with one free variable **src**; and
- for every $n \in \Lambda_2^E$, a formula ϕ_n , called an *edge formula*, with two free variables **src** and **dst**.

The domain of f is $\{G \in \mathbf{GR}(\Lambda_1) \mid G \models \phi_{dom}\}$. For every $G \in \text{dom}(f)$, the graph $G' = f(G) \in \mathbf{GR}(\Lambda_2)$ is given by

$$\begin{aligned} G'^V &= \{v \in G^V \mid \text{there is exactly one } v \in \Lambda_1^V \text{ such that } G \models \phi_v(\mathbf{src} \mapsto v)\} \\ G'^E &= \{(v, n, w) \mid v, w \in G^V \text{ and } G \models \phi_n(\mathbf{src} \mapsto v, \mathbf{dst} \mapsto w)\}. \end{aligned}$$

(For simplicity, we ignore roots in this section.)

Theorem 3. [4] *A language of graphs is c-edNCE if and only if it is the image of an M2L-definable function $f : \mathbf{GR}(\Lambda_1) \rightarrow \mathbf{GR}(\Lambda_2)$ applied to the set of directed trees over Λ_1 .*

Such a language is then said to be *f-definable*.

Theorem 4. [4] *It is decidable whether a function f defines a finite language of graphs.*

Lemma 5. [4] *The class of M2L-definable functions is closed under composition.*

Now fix $\Lambda_T^V = \{v\}$, $\Lambda_T^E = \{f_1, f_2, a\}$. A *tree with equi-level edges* is a graph G over Λ_T such that G restricted to f -edges is a directed tree and such that $(v, a, w) \in G^E$ if and only if w is the left-most node to the right of v at the same level as v , as shown in Figure 3.




Fig. 3. A tree with equi-level edges.

Lemma 6. *The set of trees over Λ_T with equi-level edges is not c-edNCE.*

Proof Suppose for a contradiction that the set is c-edNCE by means of an M2L-definable function f . Then there would be a uniform way of obtaining an M2L-definable function f_i whose graph language represents all finite sequences of configurations that TM (Turing Machine) i may produce with an empty input tape. In fact we may choose $\Lambda^V = \{0, 1, \#\}$ and construct f'_i such that it maps trees with equi-level edges into trees whose Λ^V labels at level k encode the configuration of TM i after the k 'th step (details are omitted). By Lemma 5, the set of graphs representing finite configuration sequences is then definable by a function $f_i = f'_i \circ f$. But then the Halting Problem would be decidable by Theorem 4, which is a contradiction. \square

Lemma 7. *The set of trees over Λ_T with unrestrained a -edges is not c-edNCE.*

Proof If it was we could use Lemmas 5 and 6 to show that also the set of trees with equi-level edges is **c-edNCE**. (We would construct a domain formula checking, among other things, that whenever (v, \mathbf{a}, w) and (v', \mathbf{a}, w') are edges and v' is a child of v , then w' is a child of w .) \square

Theorem 8. *c-edNCE and EC are incomparable.*

Proof **EC** $\not\subseteq$ **c-edNCE**: The set of trees with unrestrained **a**-edges is certainly **EC**, but not **c-edNCE** by Lemma 7.

c-edNCE $\not\subseteq$ **EC**: The set of cyclic graphs over singleton node and edge alphabets is **c-edNCE**, but not **EC** (in fact, since the edge label determines whether an edge is ordinary or auxiliary, only list-like structures and certain degenerate structures can be described with singleton edge alphabets). \square

5 Transductions

We are interested in graph transformations that model pointer manipulations in programs. These can be specified through a *transduction*, which is defined to be of the form $\mathcal{T} = \langle L, \mathcal{E}, \rho \rangle$. The component L is a list of labeled *entries*. An entry t defines one or two first-order variables, called *transduction variables*, according to its label as follows.

- **add-n**: this indicates the creation of an **n**-edge between two nodes denoted by first-order terms **src**(t) and **dst**(t); an existing **n**-edge from the source is deleted.
- **del-n**: this indicates the deletion of the **n**-edge whose origin is denoted by the first-order term **src**(t).
- **foll-a**: this indicates the existence of an **a**-edge which has been followed between two cells denoted by first-order terms **src**(t) and **dst**(t); this makes for an explicit representation of auxiliary edges that are followed and, therefore, known to exist in the original graph.
- **v**: this indicates that a node denoted by the first-order logical variable **src**(t) is marked with label **v** (which may be **spare**); if an ordinary node is marked **spare**, then its outgoing and incoming edges are deleted.

The component \mathcal{E} is an environment, which maps root variables to address terms denoting their values. The component ρ is a formula which must hold in order for the free variables in L and \mathcal{E} to denote a transformation. The formula ρ may contain other transduction variables than those defined by L . Together they are designated μ .

The formula ρ must ensure that the entries are consistent with each other. Thus if a graph G and a value assignment μ are such that $G, \mu \models \rho$, then some examples of technical relationships that must hold are:

- given any v and \mathbf{a} , there are at most one **fol** \mathbf{a} entry t such that $G, \underline{\mu} \models \mathbf{src}(t) = v$; and
- given any (v, \mathbf{a}, w) that is marked by a **del** \mathbf{a} entry before any **add** \mathbf{a} entry, there is a **fol** \mathbf{a} entry, which makes explicit the assumption that (v, \mathbf{a}, w) is an edge in G .

6 Predicate Transformers

Each transduction \mathcal{T} determines a predicate transformer $\mathbf{Tr}_{\mathcal{T}}$. A formula Φ is translated into $\mathbf{Tr}_{\mathcal{T}}\Phi$ according to the following rules.

$$\begin{aligned}
\mathbf{Tr}_{\mathcal{T}}(x) &= \mathcal{T}.\mathcal{E}(x) \\
\mathbf{Tr}_{\mathcal{T}}(\alpha) &= \alpha \\
\mathbf{Tr}_{\mathcal{T}}(\mathbf{A}_1 = \mathbf{A}_2) &= \mathbf{Tr}_{\mathcal{T}}(\mathbf{A}_1) = \mathbf{Tr}_{\mathcal{T}}(\mathbf{A}_2) \\
\mathbf{Tr}_{\mathcal{T}}(\alpha \xrightarrow{\mathbf{f}} \beta) &= \begin{cases} \beta = \mathbf{dst}(t) & \text{if } t \text{ is an } \mathbf{add}\text{-}\mathbf{f} \text{ entry in } \mathcal{T}.L, \\ & \alpha = \mathbf{src}(t), t \text{ is the last such entry, and no later } \mathbf{spare} \text{ entry } t' \text{ is} \\ & \text{such that } \mathbf{src}(t') \in \{\alpha, \beta\} \text{ and no} \\ & \text{later } \mathbf{del}\text{-}\mathbf{f} \text{ entry } t' \text{ is such that} \\ & \mathbf{src}(t') = \alpha \\ \mathbf{false} & \text{if there is a } \mathbf{spare} \text{ entry } t \text{ with} \\ & \mathbf{src}(t) \in \{\alpha, \beta\} \text{ or there is a } \mathbf{del}\text{-}\mathbf{f} \\ & \text{entry } t \text{ with } \mathbf{src}(t) = \alpha, \text{ and no} \\ & \text{later } \mathbf{add}\text{-}\mathbf{f} \text{ entry } t' \text{ is such that} \\ & \mathbf{src}(t') = \alpha \} \\ \alpha \xrightarrow{\mathbf{f}} \beta & \text{otherwise} \end{cases} \\
\mathbf{Tr}_{\mathcal{T}}(\mathbf{v}?\alpha) &= \begin{cases} \mathbf{true} & \text{if there is an } \mathbf{v}\text{-entry } t \text{ in } \mathcal{T}.L \\ & \text{such that } \mathbf{src}(t) = \alpha \text{ and no later} \\ & \mathbf{v}'\text{-entry } t' \text{ is such that } \mathbf{src}(t') = \\ & \alpha \\ \mathbf{v}?\alpha & \text{otherwise} \end{cases} \\
\mathbf{Tr}_{\mathcal{T}}(\mathbf{A} \in \Sigma) &= \mathbf{Tr}_{\mathcal{T}}(\mathbf{A}) \in \Sigma \\
\mathbf{Tr}_{\mathcal{T}}(\Sigma_1 \subseteq \Sigma_2) &= \Sigma_1 \subseteq \Sigma_2 \\
\mathbf{Tr}_{\mathcal{T}}(\neg\Phi) &= \neg\mathbf{Tr}_{\mathcal{T}}\Phi \\
\mathbf{Tr}_{\mathcal{T}}(\Phi_1 \wedge \Phi_2) &= \mathbf{Tr}_{\mathcal{T}}(\Phi_1) \wedge \mathbf{Tr}_{\mathcal{T}}(\Phi_2) \\
\mathbf{Tr}_{\mathcal{T}}(\exists^\circ \alpha : \Phi) &= \exists^\circ \alpha : \mathbf{Tr}_{\mathcal{T}}\Phi \\
\mathbf{Tr}_{\mathcal{T}}(\exists^\circ S : \Phi) &= \exists^\circ S : \mathbf{Tr}_{\mathcal{T}}\Phi
\end{aligned}$$

The *transformed backbone*, denoted $\mathbf{BB}_{\mathcal{T}}(\overline{G}, \underline{\mu})$, according to \mathcal{T} on \overline{G} with transduction values $\underline{\mu}$ is the graph \overline{G}' defined as follows.

- $\overline{G}'^{\mathbf{V}} = \overline{G}^{\mathbf{V}}$;
- $(v, \mathbf{f}, w) \in \overline{G}'^{\mathbf{E}}$ iff $\overline{G}, \underline{\mu} \models \mathbf{Tr}_{\mathcal{T}}(v \xrightarrow{\mathbf{f}} w)$;

- $\overline{G'}^L(v) = v$ iff $\overline{G}, \underline{\mu} \models \mathbf{Tr}_T(v?v)$; and
- $x^{\overline{G'}}$ is the node v such that $\overline{G}, \underline{\mu} \models v = \mathbf{Tr}_T(\mathcal{T}.\mathcal{E}(x))$.

Lemma 9. (Faithfulness) Let $\overline{G'} = \mathbf{BB}_T(\overline{G}, \underline{\mu})$ and let $\underline{\mathfrak{X}}$ be a value assignment to the free variables of Φ . Then,

$$\begin{aligned} & \overline{G'}, \underline{\mathfrak{X}} \models \Phi \\ \text{if and only if} & \\ & \overline{G}, \underline{\mathfrak{X}}, \underline{\mu} \models \mathbf{Tr}_T \Phi \end{aligned}$$

Proof (Sketch) By a straightforward structural induction. \square

We say that \overline{G} , $\underline{\mu}$, and \mathcal{T} determine a *transformation*. In addition to the transformed backbone, the transformation also determines:

- $\mathbf{Foll}_T\text{-a}(\overline{G}, \underline{\mu})$, the set of **a**-edges in the old graph G that were followed;
- $\mathbf{Del}_T\text{-a}(\overline{G}, \underline{\mu})$, the set of **a**-edges in the old graph G that were both followed and deleted; and
- $\mathbf{Add}_T\text{-a}(\overline{G}, \underline{\mu})$, the set of **a**-edges in the new graph G' that were added.

To specify $\mathbf{Foll}_T\text{-a}(\overline{G}, \underline{\mu})$, we define a predicate $\mathbf{Foll}_T\text{-a}$ with free variables **src** and **dst** expressing that an **a**-edge from **src** to **dst** was followed. Informally,

$$\mathbf{Foll}_T\text{-a} \equiv \text{“for some } \mathbf{foll}\text{-a} \text{ entry in } \mathcal{T}.L, \mathbf{src} = \mathbf{src}(t) \text{ and } \mathbf{dst} = \mathbf{dst}(t),\text{”}$$

which can be encoded as a formula. Now,

$$\mathbf{Foll}_T\text{-a}(\overline{G}, \underline{\mu}) = \{(v, a, w) \mid \overline{G}, \underline{\mu}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \mathbf{Foll}_T\text{-a}\}.$$

Similarly, we define the two other sets by defining predicates $\mathbf{Del}_T\text{-a}$ and $\mathbf{Add}_T\text{-a}$:

$$\mathbf{Del}_T\text{-a} \equiv \text{“}\mathbf{Foll}_T\text{-a} \text{ and there is some } \mathbf{spare} \text{ entry with } \mathbf{src} = \mathbf{src}(t) \text{ or } \mathbf{dst} = \mathbf{src}(t), \text{ or some } \mathbf{del}\text{-a} \text{ or } \mathbf{add}\text{-a} \text{ entry } t \text{ with } \mathbf{src} = \mathbf{src}(t).\text{”}$$

$$\mathbf{Add}_T\text{-a} \equiv \text{“if there is an } \mathbf{add}\text{-a} \text{ entry } t \text{ such that } \mathbf{src}(t) = \mathbf{src} \text{ and } \mathbf{dst}(t) = \mathbf{dst}, \text{ and no later entries delete this edge.”}$$

Lemma 10. $\mathbf{Del}_T\text{-a}(\overline{G}, \underline{\mu}) \subseteq \mathbf{Foll}_T\text{-a}(\overline{G}, \underline{\mu})$ if $G, \underline{\mu} \models \rho$.

Proof By the definitions and imposed technical relationships. \square

The *transformation relation* induced by \mathcal{T} is:

$$\begin{aligned} & G \longrightarrow_T G' \\ \text{if and only if} & \\ & \text{for some } \underline{\mu} : \\ & \quad \overline{G}, \underline{\mu} \models \mathcal{T}.\rho, \\ & \quad \mathbf{Foll}\text{-a}_T(\overline{G}, \underline{\mu}) \subseteq \overline{G}, \\ & \quad \overline{G'} = \mathbf{BB}_T(\overline{G}, \underline{\mu}), \text{ and} \\ & \quad \overline{G'} = (\overline{G} \setminus \mathbf{Del}\text{-a}_T(\overline{G}, \underline{\mu})) \cup \mathbf{Add}\text{-a}_T(\overline{G}, \underline{\mu}) \end{aligned}$$

Example (continued)

Consider the linked list with a designated element from Section 4. A common transduction on such structures is the insertion of a new element just before the head. This is realized by the following transduction.

$$\begin{aligned} L: & \text{L}(\mu').\mathbf{del}\text{-f}(\mathbf{x}, \mu).\mathbf{add}\text{-f}(\mathbf{x}, \mu').\mathbf{add}\text{-n}(\mu', \mu) \\ \mathcal{E}: & \mathbf{x} \mapsto \mathbf{x} \\ \rho: & \mathbf{x} \xrightarrow{\mathbf{f}} \mu \wedge \mathbf{spare}?\mu' \end{aligned}$$

Notice how this closely mimics the code that one would write in a conventional programming language. The expressive power of transductions goes beyond mere straight-line code, since regular control structures can be encoded in formulas [5].

7 Transductional Correctnesse

Let \mathfrak{A} be the free variables in the assertion \mathcal{A} and let \mathfrak{B} be the free variables in the assertion \mathcal{B} that are not already free in \mathcal{A} . The problem of transductional correctness is:

Given assertions \mathcal{A} , \mathcal{B} , and a transduction \mathcal{T} . Does it hold for all G , G' , and $\underline{\mathfrak{A}}$ that if G is tree-formed and satisfies \mathcal{A} with $\underline{\mathfrak{A}}$, and if $G \longrightarrow_{\mathcal{T}} G'$, then G' is tree-formed and satisfies \mathcal{B} for some $\underline{\mathfrak{B}}$?

Since tree-formedness by Lemma 1 can be encoded as a backbone formula, we can without loss of generality rephrase the question as follows. We say that the triple $\mathcal{A}\{\mathcal{T}\}\mathcal{B}$ is *tree-valid*, and write $\vdash \mathcal{A}\{\mathcal{T}\}\mathcal{B}$, if:

for all tree-formed G , all G' , and all $\underline{\mathfrak{A}}$, $G, \underline{\mathfrak{A}} \models \mathcal{A}$ and $G \longrightarrow_{\mathcal{T}} G'$
implies there is $\underline{\mathfrak{B}}$ such that $G', \underline{\mathfrak{B}} \models \mathcal{B}$

Note that triple tree-validity concerns only transformations of tree-formed graphs.

Our main result is to demonstrate that tree triple validity can be encoded in M2L-BB. For simplicity we assume in what follows that an assertion now contains only one edge constraint, and that $\mathcal{A} = \Phi[\sigma \xrightarrow{\mathbf{a}} \delta]$ and $\mathcal{B} = \Phi'[\sigma' \xrightarrow{\mathbf{a}} \delta']$. Then we say that triple $\mathcal{A}\{\mathcal{T}\}\mathcal{B}$ is *provable* and write $\vdash \mathcal{A}\{\mathcal{T}\}\mathcal{B}$ if

$$\begin{aligned} \vdash \forall^{\circ} \mathfrak{A} : \forall^{\circ} \mu : \\ & (\Phi \wedge \rho \wedge \forall^{\circ} \mathbf{src} \exists^{\circ} \mathbf{dst} : (\sigma \Rightarrow (\delta \wedge (\neg \mathbf{Foll}_{\mathcal{T}} \Rightarrow (\forall^{\circ} \mathbf{dst} : \neg \mathbf{Foll}_{\mathcal{T}})))) \\ & \Rightarrow \exists^{\circ} \mathfrak{B} : (\mathbf{Tr}_{\mathcal{T}} \Phi' \\ & \quad \wedge \forall^{\circ} \mathbf{src} : \mathbf{Tr}_{\mathcal{T}} \sigma' \Rightarrow \\ & \quad ((\exists^{\circ} \mathbf{dst} : \mathbf{Add}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta') \\ & \quad \vee (\exists^{\circ} \mathbf{dst} : \mathbf{Foll}_{\mathcal{T}} \wedge \neg \mathbf{Del}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta') \\ & \quad \vee (\sigma \wedge \forall^{\circ} \mathbf{dst} : \neg \mathbf{Add}_{\mathcal{T}} \wedge \neg \mathbf{Foll}_{\mathcal{T}} \wedge (\delta \Rightarrow \mathbf{Tr}_{\mathcal{T}} \delta')))) \end{aligned}$$

8 Soundness, Completeness, and Decidability

Theorem 11. (*Soundness*) $\vdash \mathcal{A}\{T\}\mathcal{B}$ implies $\Vdash \mathcal{A}\{T\}\mathcal{B}$.

Proof Assume

$$(1) \vdash \mathcal{A}\{T\}\mathcal{B}.$$

Fix a tree-formed G , a G' , and a value assignment $\underline{\mathcal{A}}$ to the free variables \mathcal{A} of \mathcal{A} such that

$$(2) G, \underline{\mathcal{A}} \models \mathcal{A}, \text{ and}$$

$$(3) G \longrightarrow_T G'.$$

To establish $\Vdash \mathcal{A}\{T\}\mathcal{B}$, we only need to find a value assignment $\underline{\mathcal{B}}$ to the remaining free variables \mathcal{B} such that

$$(4) G', \underline{\mathcal{A}}, \underline{\mathcal{B}} \models \mathcal{B}.$$

Now by (3) and the definition of transductions, there is a value assignment $\underline{\mu}$ to the transduction variables μ of T such that

$$(5) \overline{G}, \underline{\mu} \models T.\rho$$

$$(6) \mathbf{Foll}_T(\overline{S}, \underline{\mu}) \subseteq \overline{G},$$

$$(7) \overline{G}' = \mathbf{BB}_T(\overline{G}, \underline{\mu}), \text{ and}$$

$$(8) \overline{G}' = (\overline{G} \setminus \mathbf{Del}_T(\overline{G}, \underline{\mu})) \cup \mathbf{Add}_T(\overline{G}, \underline{\mu})$$

In order to apply (1), we would like to show that

$$(9) \overline{G}, \underline{\mathcal{A}}, \underline{\mu} \models \Phi \wedge \rho \wedge \forall^\circ \mathbf{src} \exists^\circ \mathbf{dst} : \sigma \Rightarrow (\delta \wedge (\neg \mathbf{Foll}_T \Rightarrow (\forall^\circ \mathbf{dst} : \neg \mathbf{Foll}_T)))$$

holds. Now by (2), we have $\overline{G}, \underline{\mathcal{A}} \models \Phi$ and $\overline{G}, \underline{\mathcal{A}} \models [\sigma \xrightarrow{a} \delta]$. Thus it is sufficient to find for each v such that $\overline{G}, \underline{\mathcal{A}}, \mathbf{src} \mapsto v \models \sigma$ some w satisfying

$$(10) \overline{G}, \underline{\mathcal{A}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \delta \wedge (\neg \mathbf{Foll}_T \Rightarrow (\forall^\circ \mathbf{dst} : \neg \mathbf{Foll}_T))$$

The w we choose is the one such that $(v, \mathbf{a}, w) \in \overline{G}$. This w exists by virtue of (2) and the definition of edge constraint satisfaction. Moreover, $\overline{G}, \underline{\mathcal{A}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \delta$. Thus in order to establish (10), it suffices to suppose that

$$(11) \overline{G}, \underline{\mathcal{A}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \neg \mathbf{Foll}_T$$

and to prove that no u exists such that

$$(12) \overline{G}, \underline{\mathcal{A}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto u \models \mathbf{Foll}_T.$$

For a contradiction, assume that some u does satisfy (12). Then $(v, \mathbf{a}, u) \in \mathbf{Foll}_{\mathcal{T}}(\overline{G}, \underline{\mu})$. But by (5), $\mathbf{Foll}_{\mathcal{T}}(\overline{G}, \underline{\mu}) \subseteq \overline{G}$, and thus $u = w$, which contradicts our supposition (11). It follows that (9) holds, and by (1) we then obtain a $\underline{\mathfrak{B}}$ such that

$$(13) \quad \begin{aligned} \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu} \models \mathbf{Tr}_{\mathcal{T}} \Phi' \\ \wedge \forall^{\circ} \mathbf{src} : \mathbf{Tr}_{\mathcal{T}} \sigma' \Rightarrow \\ ((\exists^{\circ} \mathbf{dst} : \mathbf{Add}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta') \\ \vee (\exists^{\circ} \mathbf{dst} : \mathbf{Foll}_{\mathcal{T}} \wedge \neg \mathbf{Del}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta') \\ \vee (\sigma \wedge \forall^{\circ} \mathbf{dst} : \neg \mathbf{Add}_{\mathcal{T}} \wedge \neg \mathbf{Foll}_{\mathcal{T}} \wedge (\delta \Rightarrow \mathbf{Tr}_{\mathcal{T}} \delta')))) \end{aligned}$$

holds. From (13) and Lemma 9 (Faithfulness), it follows that

$$(14) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}} \models \Phi'$$

We thus only need to show that also the edge constraint $[\sigma' \xrightarrow{\mathbf{a}} \delta']$ holds. To do this, we consider $v \in \overline{G}'$ such that

$$(15) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v \models \sigma'.$$

We must then prove that there is w such that $(v, \mathbf{a}, w) \in \overline{G}'$ and

$$(16) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \delta'.$$

Now by (15) and Lemma 9 (Faithfulness), we have

$$(17) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu}, \mathbf{src} \mapsto v \models \mathbf{Tr}_{\mathcal{T}} \sigma'.$$

Discharging the hypothesis in (13) by means of (17) gives us three cases:

$$(18) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu}, \mathbf{src} \mapsto v \models \exists^{\circ} \mathbf{dst} : \mathbf{Add}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta'$$

$$(19) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu}, \mathbf{src} \mapsto v \models \exists^{\circ} \mathbf{dst} : \mathbf{Foll}_{\mathcal{T}} \wedge \neg \mathbf{Del}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta'$$

$$(20) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu}, \mathbf{src} \mapsto v \models \sigma \wedge \forall^{\circ} \mathbf{dst} : \neg \mathbf{Add}_{\mathcal{T}} \wedge \neg \mathbf{Foll}_{\mathcal{T}} \wedge (\delta \Rightarrow \mathbf{Tr}_{\mathcal{T}} \delta')$$

In case (18) there is a w such that

$$(21) \overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \underline{\mu}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \mathbf{Add}_{\mathcal{T}} \wedge \mathbf{Tr}_{\mathcal{T}} \delta'$$

By (8), $(v, \mathbf{a}, w) \in \overline{G}'$, and by Lemma 9 (Faithfulness) (16) holds. Case (19) is handled by a similar argument. Finally, in Case (20) we have by Lemma 9 (Faithfulness) that $\overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v \models \sigma$ and $\overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \neg \mathbf{Add}_{\mathcal{T}} \wedge \neg \mathbf{Foll}_{\mathcal{T}} \wedge (\delta \Rightarrow \mathbf{Tr}_{\mathcal{T}} \delta')$, where w is the node such that $(v, \mathbf{a}, w) \in \overline{G}$ (this node exists by virtue of (2)). By (8), (20), and Lemma 10, we infer that $(v, \mathbf{a}, w) \in \overline{G}'$ and by (2) that $\overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \mathbf{Tr}_{\mathcal{T}} \delta$. Thus $\overline{G}, \underline{\mathfrak{A}}, \underline{\mathfrak{B}}, \mathbf{src} \mapsto v, \mathbf{dst} \mapsto w \models \mathbf{Tr}_{\mathcal{T}} \delta'$ holds, whence (16) holds by Lemma 9 (Faithfulness). \square

Theorem 12. (Completeness) $\Vdash \mathcal{A}\{T\}\mathcal{B}$ implies $\vdash \mathcal{A}\{T\}\mathcal{B}$.

Proof Proof can be found in full paper.

□

Theorem 13. *Transductional correctness is decidable for EC.*

Proof By Theorems 2, 11, and 12.

□

References

1. B. Courcelle. Graph rewriting: an algebraic and logic approach. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 193–242. Elsevier Science Publishers, 1990.
2. B. Courcelle. The monadic second-order logic of graphs I. Recognizable sets of finite graphs. *Information and computation*, 85:12–75, 1990.
3. B. Courcelle. Monadic second-order definable graph transductions. In J.C. Raoult, editor, *CAAP '92, Colloquium on Trees in Algebra and Programming, LNCS 581*, pages 124–144. Springer Verlag, 1992.
4. J. Engelfriet. A characterization of context-free NCE graph languages by monadic second-order logic on trees. In H. Ehrig, H.J. Kreowski, and G. Rozenberg, editors, *Graph grammars and their applications to computer science, 4th International Workshop, LNCS 532*, pages 311–327. Springer Verlag, 1990.
5. N. Klarlund and M. Schwartzbach. Graph types. In *Proc. 20th Symp. on Princ. of Prog. Lang.*, pages 196–205. ACM, 1993.
6. N. Klarlund and M. Schwartzbach. Invariants as data types. Unpublished, 1993.
7. M. Rabin. A simple method for undecidability proofs and some applications. In *Logic, Methodology and Philosophy of Science II*, pages 58–68. North-Holland, 1965.