



Selected Revised Papers from the
4th International Workshop on
Graph Computation Models
(GCM 2012)

Verifying Total Correctness of Graph Programs

Christopher M. Poskitt and Detlef Plump

20 pages

Verifying Total Correctness of Graph Programs

Christopher M. Poskitt¹ and Detlef Plump²

¹ ETH Zürich, Switzerland

² The University of York, UK

Abstract: GP 2 is an experimental nondeterministic programming language based on graph transformation rules, allowing for visual programming and the solving of graph problems at a high-level of abstraction. In previous work we demonstrated how to verify graph programs using a Hoare-style proof calculus, but only partial correctness was considered. In this paper, we add new proof rules and termination functions, which allow for proofs to additionally guarantee that program executions always terminate (weak total correctness), or that programs always terminate and do so without failure (total correctness). We show that the new proof rules are sound with respect to the operational semantics of GP 2, complete for termination, and demonstrate their use on some example programs.

Keywords: graph programs, verification, Hoare logic, total correctness, termination

1 Introduction

The verification of graph transformation systems is an area of active and growing interest, motivated by the many applications of graph transformation to specification and programming. Whilst much of the research in this area has focused on sets of rules or graph grammars (see e.g. [BCK08, BHE09, KE10, CR12]), the challenge of verifying graph-based programming languages is also beginning to be addressed. In particular, Habel, Pennemann, and Rensink [HPR06, HP09] contributed a verification framework – based on weakest preconditions – for a simple graph transformation language, expressing graph properties with nested conditions (a formalism based on graph morphisms). Their language however does not support important practical features such as computation on labels, and their weakest precondition calculus generates infinite preconditions for loops.

In [PP12a] we considered the verification of GP [Plu09], a nondeterministic programming language based on graph transformation. The states are directed labelled graphs, which are manipulated via the application of (conditional) rule schemata. These generalise double-pushout rules with relabelling and expressions. The verification framework of the paper is a Hoare calculus for partial correctness, with which one can prove that programs executed on graphs satisfying given preconditions will only ever result in graphs satisfying given postconditions. However, the calculus cannot be used to prove that such programs do eventually terminate, and cannot be used to prove the absence of failing executions. Addressing these two issues is the focus of this paper.

We define two notions of total correctness: a weaker one accounting for termination, and a stronger one accounting for termination as well as for absence of failures. We define two calculi for these notions of total correctness, using termination functions (that map graphs to natural

numbers) in the new proof rule for the iteration command. We demonstrate the proof calculi on programs that have loops and potential failure points, before proving the calculi to be sound as well as complete for termination.

In contrast to our previous papers, we present the work here in the setting of GP 2 [Plu12] (henceforth referred to as simply GP). This extended version of the language has an improved type system, a marking (shading) mechanism for nodes and edges, a new conditional construct, and a simplified semantics for branching and iteration to support a more efficient implementation. Our previous verification work has been updated in [Pos13] to support these new features, but due to space limitations we cannot present all of the revised definitions here. We attempt to make the intuition behind each concept clear, but refer the interested reader to [Pos13] for the full technical details and further explanations.

Section 2 reviews some technical preliminaries. Section 3 is an informal introduction to graph programs. Section 4 reviews our assertion language and the partial correctness proof rules of our previous calculus. Section 5 formalises the notion of (weak) total correctness and presents new proof rules which allow one to prove these properties. Section 6 demonstrates the use of the new calculi on some example programs. Section 7 presents a proof that the new calculi are sound for (weak) total correctness, and also a proof that the calculi are complete for termination. Finally, we conclude in Section 8.

This paper is a revised and extended version of [PP12b].

2 Preliminaries

Graph transformation in GP is based on the double-pushout (DPO) approach with relabelling [HP02], i.e. an approach in which both node and edge labels can be relabelled. This framework deals with rules containing partially labelled graphs, the definition of which we recall below. In this section we treat the label alphabet as a parameter because we will require two different alphabets for two classes of graphs: “syntactic” graphs labelled with expressions, and “semantic” graphs labelled with lists composed of integers and strings. We also introduce assignments which translate syntactic graphs into semantic graphs.

A *graph* over a label alphabet \mathcal{C} is a system $G = (V_G, E_G, s_G, t_G, l_G, m_G)$, where V_G and E_G are finite sets of *nodes* (or *vertices*) and *edges*, $s_G, t_G: E_G \rightarrow V_G$ are the *source* and *target* functions for edges, $l_G: V_G \rightarrow \mathcal{C}$ is the partial¹ node labelling function and $m_G: E_G \rightarrow \mathcal{C}$ is the (total) edge labelling function (edges can be relabelled by deletion and re-insertion, hence unlabelled edges are not necessary). Given a node v , we write $l_G(v) = \perp$ to express that $l_G(v)$ is undefined. Graph G is *totally labelled* if l_G is a total function. We write $\mathcal{G}(\mathcal{C})$ for the set of all totally labelled graphs over \mathcal{C} , and $\mathcal{G}(\mathcal{C}_\perp)$ for the set of all graphs over \mathcal{C} . The *empty graph*, denoted by \emptyset , has empty node and edge sets.

A *graph morphism* $g: G \rightarrow H$ between graphs G, H in $\mathcal{G}(\mathcal{C}_\perp)$ consists of two functions $g_V: V_G \rightarrow V_H$ and $g_E: E_G \rightarrow E_H$ that preserve sources, targets and labels; that is, $s_H \circ g_E = g_V \circ s_G$, $t_H \circ g_E = g_V \circ t_G$, $m_H \circ g_E = m_G$, and $l_H(g(v)) = l_G(v)$ for all v such that $l_G(v) \neq \perp$. Morphism g is an *inclusion* if $g(x) = x$ for all nodes and edges x . It is *injective* (*surjective*) if g_V and g_E are injective (surjective). It is an *isomorphism* if it is injective, surjective, and satisfies

¹ Unlabelled nodes appear in the interfaces of rule schemata to allow relabelling, see [PP12a, Pos13].

$l_H(g_V(v)) = \perp$ for all nodes v with $l_G(v) = \perp$. In this case G and H are *isomorphic*, which is denoted by $G \cong H$.

We consider graphs over two distinct label alphabets. Graph programs and our assertion language contain graphs labelled with expressions, while the graphs on which programs operate are labelled with lists composed of integers and character strings. In both cases nodes and edges can be marked; marked nodes are displayed as shaded, and marked edges are displayed as dashed (see Figure 2). We consider graphs of the first type as syntactic objects and graphs of the second type as semantic objects, and aim to clearly separate these levels of syntax and semantics.

Let \mathbb{Z} denote the set of integers and Char a finite set of characters. We fix the label alphabet:

$$\mathcal{L} = (\mathbb{Z} \cup \text{Char}^*)^* \times \mathbb{B}$$

where $\mathbb{B} = \{\text{true}, \text{false}\}$, i.e. all sequences over integers and character strings, along with a Boolean value indicating whether the node or edge is marked or not. Occasionally we will refer only to the list component $(\mathbb{Z} \cup \text{Char}^*)^*$, which shall be denoted by \mathbb{L} .

The other label alphabet we are using, `Label`, consists of a mark component and (colon delimited) sequences of arithmetical expressions and strings. These may contain variables from a set denoted by `VarId`. Variables represent values in \mathbb{L} , i.e. lists, and can be constrained in rule schemata to represent integers, strings, or atoms (an integer or a string). These types correspond to the semantic domains in Figure 1, in which we identify atoms and unit-length lists to establish a subtype hierarchy.



Figure 1: Subtype hierarchy for lists

We write $\mathcal{G}(\text{Label})$ to denote the set of all graphs labelled over `Label` (grammars defining the label alphabet are given in [Plu12, Pos13]). Examples of list components of labels in $\mathcal{G}(\text{Label})$ include `x*5` and `"root":y` (the variable `x` may only be instantiated to integers, whereas `y` be instantiated to any value in \mathbb{L} , unless otherwise constrained).

Each graph in $\mathcal{G}(\text{Label})$ represents a possibly infinite set of graphs in $\mathcal{G}(\mathcal{L})$. The latter are obtained by instantiating variables with values from \mathbb{L} and evaluating expressions. An *assignment* is a partial function $\alpha: \text{VarId} \rightarrow \mathbb{L}$. For a rule schema (see the next section), α must satisfy for all variables `x` with type `int` (resp. `string`, `atom`, `list`), $\alpha(x) \in \mathbb{Z}$ (resp. `Char*`, `ℤ ∪ Char*`, \mathbb{L}). For assertions (see Section 4), we require that α is *well-typed* for the expressions to which it is applied, i.e. it assigns values to variables of types determined by their contexts. For example, a well-typed assignment for `x+y:z` (with `+` interpreted as addition) must map `x`, `y` to integers.

Given a (well-typed) assignment α and label $(e\ b)$ with e a list and $b \in \mathbb{B}$, we define $(e\ b)^\alpha = (e^\alpha, b)$ where the value $e^\alpha \in \mathbb{L}$ is inductively defined as follows. If e is the empty list, then e^α is the empty sequence. If e is a numeral or a sequence of characters, then e^α is the integer or character string represented by e . (Note that the empty list and empty character string are distinct values.) If e is a variable identifier, then $e^\alpha = \alpha(e)$. For arithmetic and string expressions, e^α is defined inductively in the usual way. Finally, if e has the form $e_1 : e_2$ with e_1, e_2 list expressions, then $e^\alpha = e_1^\alpha e_2^\alpha$ (the concatenation of the sequences e_1^α and e_2^α). Given a graph G in $\mathcal{G}(\text{Label})$ and an assignment α well-typed for all expressions in G , we write G^α for the graph in $\mathcal{G}(\mathcal{L})$ that is obtained from G by replacing each label l with l^α (note that G^α has the same nodes, edges, source and target functions as G). If $g : G \rightarrow H$ is a graph morphism with $G, H \in \mathcal{G}(\text{Label})$, then g^α denotes the morphism $\langle g_V, g_E \rangle : G^\alpha \rightarrow H^\alpha$.

Remark 1 In [Plu12], variables belong to one of four distinct sets of variables – one set for each type – and assignments are families of mappings from the these sets to the appropriate semantic domains. We use a different definition in this paper, taking all variables to be members of the set `VarId`, and interpreting type declarations in rule schemata as constraints on possible assignments. This definition allows us to treat variables in rule schemata and our assertion language more uniformly, simplifying the presentation of this paper. Additionally, though GP 2 introduced indegree and outdegree functions in expressions, we do not consider them in this paper, as applicability properties of rule schemata that use them cannot be expressed by our assertion language. \square

3 Graph Programs

We introduce graph programs informally and by example in this section. For technical details, further examples, and more discussion on the operational semantics, refer to [Plu12].

The “building blocks” of graph programs are (conditional) rule schemata: a program is essentially a list of declarations of (conditional) rule schemata together with a command sequence for controlling their application. Rule schemata generalise graph transformation rules, in that labels contain (sequences of) expressions and relabelling is supported. Expressions may contain variables, which in rule schemata are associated with types integer, string, atom, or list, constraining the possible mappings for assignments. Conditional rule schemata further constrain assignments with a condition: one use is in requiring relations between expressions (e.g. $x < y + z$), but they can also be used to require (the absence of) edges between nodes in a match (e.g. `not edge (1, 2)`). As the values of variables at execution are determined by graph matching, we require that expressions in the left graph have a simple shape: (1) expressions contain no arithmetic operators; (2) expressions contain at most one occurrence of a list variable; and (3) each string expression contains at most one occurrence of a string variable.

In Figure 2 we give an example of a conditional rule schema and a possible application of it. The first row of the diagram (in the box) contains the conditional rule schema. There is an identifier, here `bridge`, and a declaration of variables with their types. The left- and right-hand graphs describe the rule with the small numbers indicating which nodes correspond to each other. Here, `bridge` is applied to a path of length two across nodes in which only the first is marked, and across unmarked string-labelled edges, provided that there is not already a direct edge from

the first node to the third (as per the condition). The effect of applying `bridge` is to add a marked edge from the first node to the third, removing the mark from the former whilst adding a mark to the latter, and taking the composition of their list components for the new edge. The conditional rule schema describes an infinite number of “concrete” graph transformation rules with labels fitting the pattern described by the schema. The second row of the diagram shows one such rule, obtained by evaluating expressions according to the assignment:

$$\alpha = \{a \mapsto \text{“AB”}, b \mapsto \text{“BC”}, x \mapsto 0 : 1 : 2, y \mapsto 3, z \mapsto 4\}$$

which adheres to the constraints of the type declaration. The bottom row shows an application of `bridge` to a graph via the same assignment α and an injective morphism g . It is applied in the double-pushout approach with relabelling [HP02], which intuitively means that nodes can be relabelled in an arbitrary context (edges can simply be deleted and reinserted with the new labels), and that the application is side-effect free (i.e. it is forbidden to delete a node unless the rule schema explicitly deletes all edges it is incident to).




Figure 2: A conditional rule schema and a possible application of it

The application of a rule schema r to a graph $G \in \mathcal{G}(\mathcal{L})$ that yields a graph isomorphic to

$H \in \mathcal{G}(\mathcal{L})$, denoted $G \Rightarrow_r H$, proceeds roughly as follows:

1. Match the left graph L of r with a subgraph of G , ignoring labels, by means of a so-called premorphism $g: L \rightarrow G$. (A premorphism is a graph morphism that does not need to preserve labels.)
2. Check whether there is an assignment α of values to all the variables in r (adhering to the declared types) such that after evaluating the expressions in L , g is label-preserving.
3. If r is a conditional rule schema, check that the condition evaluates to true with respect to α and g (conditions are evaluated in the obvious way, with $\text{edge}(m, n)^{\alpha, g} = \text{true}$ for node identifiers m, n if and only if there is an edge with source $g_V(m)$ and target $g_V(n)$).
4. Apply the rule r^α (obtained from r by evaluating² all expressions in the left and right graph) to G with match g via the double-pushout approach with relabelling.

We also write $G \Rightarrow_{\mathcal{R}} H$ for a set of (conditional) rule schemata \mathcal{R} if there is some $r \in \mathcal{R}$ such that $G \Rightarrow_r H$.

Declarations of (conditional) rule schemata are, in graph programs, applied according to a number of simple control constructs. GP provides non-deterministic choice, sequential composition, conditional constructs, and as-long-as-possible iteration. We demonstrate these informally with two example programs.

The program `colouring` in Figure 3 produces a colouring (assignment of integers to nodes such that adjacent nodes have different colours), provided that the input graph consists of unmarked items only, and the list components of nodes are atomic. Colours are recorded as so-called tags, i.e. information stored in a label by extending the list component.

The program initially colours each node with 0 by applying the rule schema `init` as long as possible, using the iteration operator `!`. It then repeatedly increments the target colour of edges with the same colour at both ends. Note that this process is nondeterministic: Figure 3 shows two possible executions.

The program `reachable?` in Figure 4 checks whether or not there is a path from one distinguished node (tagged with 1) to another (tagged with 2), again provided that the input graph contains unmarked items only and the list components of nodes are atoms (except for the distinguished nodes). An execution of `reachable?` returns the original input graph if there exists such a path, otherwise it returns the same graph but with an additional edge linking the distinguished nodes. With `propagate!`, the program iteratively tags nodes with 0 that are reachable from the 1-tagged node. An if-then-else conditional is then encountered: if its “guard” `reachable` can be applied (to a copy of the graph), then `skip` is executed; otherwise `addlink`. The idea is as follows: if `reachable` can be applied, then there must be a tagged node adjacent to the second distinguished node, indicating the existence of a path. In this case, `skip` is applied which does not change the graph. If `reachable` cannot be applied, then there must not exist a path, and so `addlink` is applied to add an edge directly between the distinguished nodes. In both cases, the 0-tags used in the computation are removed by the iteration of `undo`.

² The evaluation of full GP 2 expressions (with in- and outdegree functions) depends on g as well as α .

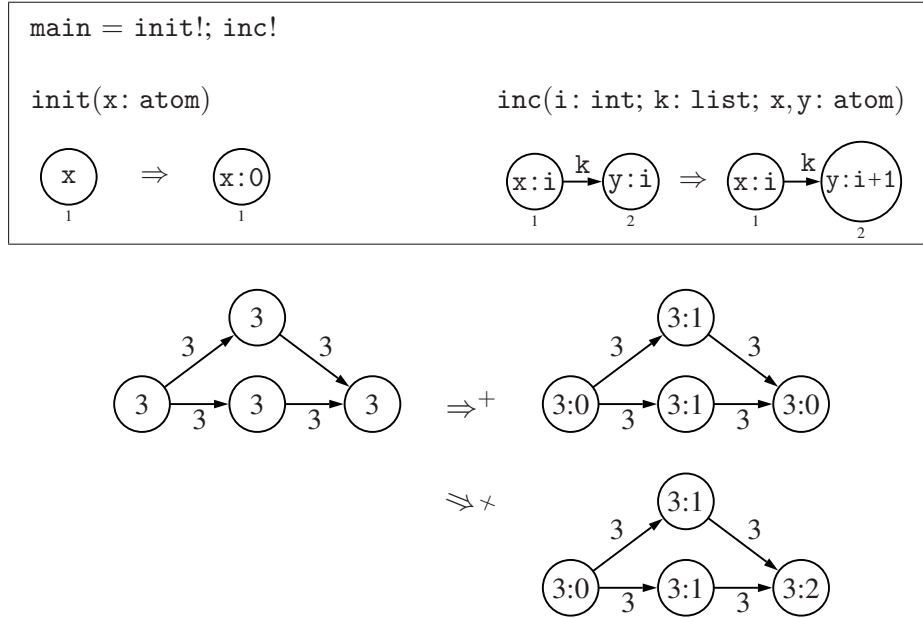


Figure 3: The program colouring and two of its executions

The formal semantics of GP is given in the style of structural operational semantics. Inference rules inductively define a small-step transition relation \rightarrow on *configurations*. In our setting, a configuration is either a command sequence (ComSeq) together with a graph (i.e. an unfinished computation), just a graph, or the special element fail (representing a failure state). The meaning of graph programs is summarised by a semantic function $\llbracket _ \rrbracket$, which assigns to every program P the function $\llbracket P \rrbracket$ mapping an input graph G to the set of all possible results of running P on G . The result set may contain, besides proper results in the form of graphs, the special values fail and \perp . The value fail indicates a failed program run whilst \perp indicates a non-terminating or stuck computation. The *semantic function* $\llbracket _ \rrbracket: \text{ComSeq} \rightarrow (\mathcal{G}(\mathcal{L}) \rightarrow 2^{\mathcal{G}(\mathcal{L}) \cup \{\text{fail}, \perp\}})$ is defined by:

$$\llbracket P \rrbracket G = \{X \in (\mathcal{G}(\mathcal{L}) \cup \{\text{fail}\}) \mid \langle P, G \rangle \xrightarrow{\pm} X\} \cup \{\perp \mid P \text{ can diverge or get stuck from } G\}$$

where P can diverge from G if there is an infinite sequence $\langle P, G \rangle \rightarrow \langle P_1, G_1 \rangle \rightarrow \langle P_2, G_2 \rangle \rightarrow \dots$, and P can get stuck from G if there is a terminal configuration $\langle Q, H \rangle$ such that $\langle P, G \rangle \rightarrow^* \langle Q, H \rangle$ where the rest program Q cannot be executed because no inference rule is applicable. A program can get stuck if it contains a non-terminating subprogram in a loop or in a conditional.

Figure 5 shows the inference rules of commands in GP 2. Each rule consists of a premise and a conclusion separated by a horizontal bar. Both parts contain meta-variables for command sequences and graphs, where \mathcal{R} stands for a rule schema set call, C, P, P', Q stand for command sequences, and G, H stand for graphs in $\mathcal{G}(\mathcal{L})$. Meta-variables are considered to be universally quantified. The notation $G \not\Rightarrow_{\mathcal{R}} H$ expresses that for graph G in $\mathcal{G}(\mathcal{L})$ there is no graph H such that $G \Rightarrow_{\mathcal{R}} H$. Derived commands such as skip can be expressed by semantically equivalent

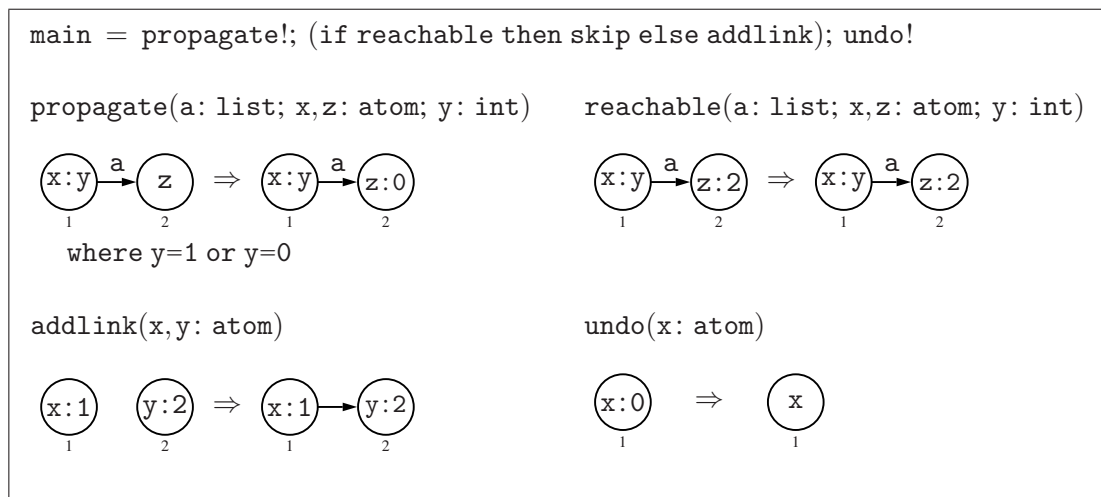


Figure 4: The program reachable?

programs made up of core commands only (in this case, the rule schema $\emptyset \Rightarrow \emptyset$). We refer to [Plu12] for details.

4 Proving Partial Correctness

In this section we first review E-conditions, the assertion language of our proof calculi. Then, we review the partial correctness proof calculus presented in previous work (updated for GP 2, e.g. the new [try] proof rule).

Nested graph conditions with expressions (or *E-conditions*) are a morphism-based formalism for expressing both structural properties of graphs and properties about their labels. E-conditions [PP12a] extend the nested conditions of [HPR06] with expressions for labels and *assignment constraints*, which are simple Boolean expressions used to restrict the instantiations of variables to values of particular types, or values that hold in particular relations. A simple example of an E-condition is:

$$\neg \exists (\textcircled{x} \mid \text{int}(x))$$

which is satisfied by graphs that do not have any unmarked integer-labelled nodes. The formalism combines logical quantifiers with graph structure and constraints on labels: E-conditions demand the (non-)existence of particular subgraphs, subject to some constraint on the labels (the vertical bar can be read as “such that”). More generally, the formalism exploits nesting to allow universally quantified expressions. For example,

$$\forall (\textcircled{x}_1 \mid \text{atom}(x), \exists (\textcircled{x}_1) \vee \exists (\textcircled{x}_1))$$

which expresses that every unmarked atom-labelled node is incident to a loop. Here, the number

$$\begin{array}{ll}
 [\text{call}_1] \frac{G \Rightarrow_{\mathcal{R}} H}{\langle \mathcal{R}, G \rangle \rightarrow H} & [\text{call}_2] \frac{G \not\Rightarrow_{\mathcal{R}}}{\langle \mathcal{R}, G \rangle \rightarrow \text{fail}} \\
 [\text{seq}_1] \frac{\langle P, G \rangle \rightarrow \langle P', H \rangle}{\langle P; Q, G \rangle \rightarrow \langle P'; Q, H \rangle} & [\text{seq}_2] \frac{\langle P, G \rangle \rightarrow H}{\langle P; Q, G \rangle \rightarrow \langle Q, H \rangle} \\
 [\text{seq}_3] \frac{\langle P, G \rangle \rightarrow \text{fail}}{\langle P; Q, G \rangle \rightarrow \text{fail}} & \\
 [\text{if}_1] \frac{\langle C, G \rangle \rightarrow^+ H}{\langle \text{if } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle P, G \rangle} & [\text{if}_2] \frac{\langle C, G \rangle \rightarrow^+ \text{fail}}{\langle \text{if } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle Q, G \rangle} \\
 [\text{try}_1] \frac{\langle C, G \rangle \rightarrow^+ H}{\langle \text{try } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle P, H \rangle} & [\text{try}_2] \frac{\langle C, G \rangle \rightarrow^+ \text{fail}}{\langle \text{try } C \text{ then } P \text{ else } Q, G \rangle \rightarrow \langle Q, G \rangle} \\
 [\text{alap}_1] \frac{\langle P, G \rangle \rightarrow^+ H}{\langle P!, G \rangle \rightarrow \langle P!, H \rangle} & [\text{alap}_2] \frac{\langle P, G \rangle \rightarrow^+ \text{fail}}{\langle P!, G \rangle \rightarrow G}
 \end{array}$$

Figure 5: Inference rules for core commands

identifies the nodes as being the same; the nesting adds more detail about the required context of the particular subgraph.

Similarly to rule schemata, in checking whether a graph satisfies a property described by an E-condition, a suitable assignment α must be found for the label expressions and assignment constraint. Note however that unlike in rule schemata, we are not declaring types for variables, but rather using predicates about types within assignment constraints. For example, we could write not `int(x)` as an assignment constraint, or even omit type predicates completely.

Due to space limitations we do not give a formal syntax or semantics of assignment constraints (we refer the reader to [PP12a, Pos13]) – there are several examples in this paper however. Example 1 includes a simple assignment constraint, `x > y`. An assignment is well-typed for this if it maps both `x` and `y` to integers. Such an assignment constraint γ is evaluated with respect to a well-typed assignment α , denoted γ^α , by instantiating variables with the values given by α and then replacing function and relation symbols with the obvious functions and relations.

In our formal definition of E-conditions, the part of the formalism immediately after each quantifier is a morphism – not simply a graph. In our examples, we draw only the codomains; in general there are chains of morphisms along the nesting (starting from the empty graph in this paper). We discuss this aspect of E-conditions in more detail shortly.

Definition 1 (E-condition) *An E-condition c over a graph P is of the form `true` or $\exists(a \mid \gamma, c')$, where $a: P \hookrightarrow C$ is an injective graph morphism with $P, C \in \mathcal{G}(\text{Label})$, γ is an assignment constraint, and c' is an E-condition over C . Boolean formulae over E-conditions over P yield E-conditions over P , that is, $\neg c$ and $c_1 \wedge c_2$ are E-conditions over P if c, c_1, c_2 are E-conditions over P . \square*

In order to define the satisfaction relation for E-conditions, we first define substitutions to allow the replacement of variables with lists (this is used to enforce equal assignment of variables in the nesting of E-conditions). A *substitution* is a partial function $\sigma: \text{VarId} \rightarrow \text{List}$. Given a label $(e\ b)$ with e a list and b a mark, σ is *well-typed* for e if it does not replace variables in arithmetic (resp. string) expressions with string (resp. arithmetic) expressions. In this case, the list e^σ is obtained from e by replacing every variable x for which σ is defined with $\sigma(x)$ (if σ is not defined for a variable x , then $x^\sigma = x$). Given a graph G in $\mathcal{G}(\text{Label})$ such that σ is well-typed for all lists in G , we write G^σ for the graph in $\mathcal{G}(\text{Label})$ that is obtained by replacing each list e with e^σ . If $g: G \rightarrow H$ is a graph morphism between graphs in $\mathcal{G}(\text{Label})$, then g^σ denotes the morphism $\langle g_V, g_E \rangle: G^\sigma \rightarrow H^\sigma$. A substitution $\sigma: \text{VarId} \rightarrow \text{List}$ can be applied to an assignment constraint γ , if it is well-typed for all expressions in γ . The resulting assignment constraint, denoted by γ^σ , is simply γ with each expression e replaced by e^σ .

Given an assignment $\alpha: \text{VarId} \rightarrow \mathbb{L}$, the *substitution* $\sigma_\alpha: \text{VarId} \rightarrow \text{List}$ induced by α maps every variable x to the expression that is obtained from $\alpha(x)$ by replacing integers and strings with their syntactic counterparts. For example, if $\alpha(x)$ is the integer 23, then $\sigma_\alpha(x)$ is 23 (the syntactic digits). Consider another example: if $\alpha(x)$ is the sequence $56: "a": "bc"$, where 56 is an integer and "a" and "bc" are strings, then $\sigma_\alpha(x) = 56: "a": "bc"$.

The *satisfaction* of E-conditions by injective graph morphisms between graphs in $\mathcal{G}(\mathcal{L})$ is defined inductively. Every such morphism satisfies the E-condition true . An injective graph morphism $s: S \hookrightarrow G$ with $S, G \in \mathcal{G}(\mathcal{L})$ satisfies the E-condition $c = \exists(a: P \hookrightarrow C \mid \gamma, c')$, denoted $s \models c$, if there exists an assignment α that is well-typed for all expressions in P, C, γ and is undefined for variables present only in c' , such that $S = P^\alpha$, and such that there is an injective graph morphism $q: C^\alpha \hookrightarrow G$ with $q \circ a^\alpha = s$, $\gamma^\alpha = \text{true}$, and $q \models (c')^{\sigma_\alpha}$. Here, σ_α is the substitution induced by α , which we require to be well-typed for all expressions in c' . If such an assignment α and morphism q exist, we say that s satisfies c by α , and write $s \models_\alpha c$. The satisfaction of Boolean formulae over E-conditions is defined inductively, in the obvious way.

For brevity, we write false for $\neg \text{true}$, $\exists(a \mid \gamma)$ for $\exists(a \mid \gamma, \text{true})$, $\exists(a, c')$ for $\exists(a \mid \text{true}, c')$, and $\forall(a \mid \gamma, c')$ for $\neg \exists(a \mid \gamma, \neg c')$. In our examples, when the domain of morphism $a: P \hookrightarrow C$ can unambiguously be inferred, we write only the codomain C . For instance, an E-condition $\exists(\emptyset \hookrightarrow C, \exists(C \hookrightarrow C'))$ can be written as $\exists(C, \exists(C'))$, where the domain of the outermost morphism is the empty graph, and the domain of the nested morphism is the codomain of the encapsulating E-condition's morphism.

An E-condition over a graph morphism whose domain is the empty graph is referred to as an *E-constraint*.

Example 1 The E-constraint $\forall(\textcircled{x}_1 \xrightarrow{k} \textcircled{y}_2 \mid x > y, \exists(\textcircled{x}_1 \xrightarrow{k} \textcircled{y}_2))$ expresses that every pair of unmarked integer-labelled nodes linked by an unmarked edge with the source label greater than the target label, has an unmarked loop incident to the source node. The (fully) unabbreviated version of the E-constraint is as follows:

$$\neg \exists(\emptyset \hookrightarrow \textcircled{x}_1 \xrightarrow{k} \textcircled{y}_2 \mid x > y, \neg \exists(\textcircled{x}_1 \xrightarrow{k} \textcircled{y}_2 \hookrightarrow \textcircled{x}_1 \xrightarrow{k} \textcircled{y}_2 \mid \text{true}, \text{true})). \quad \square$$

A graph G in $\mathcal{G}(\mathcal{L})$ satisfies an E-constraint c , denoted $G \models c$, if the morphism $\emptyset \hookrightarrow G$ satisfies c .

Definition 2 (Partial correctness) A graph program P is *partially correct* with respect to a precondition c and postcondition d (both of which are E-constraints), denoted $\models_{\text{par}} \{c\} P \{d\}$, if for every graph $G \in \mathcal{G}(\mathcal{L})$, $G \models c$ implies $H \models d$ for every graph H in $\llbracket P \rrbracket G$. \square

In [PP12a] we defined axioms and inference rules for proving partial correctness specifications about graph programs. These are given in Figure 6 (with [try] new for GP 2), where r (resp. \mathcal{R}) ranges over conditional rule schemata (resp. sets of conditional rule schemata), $c, c', d, d', e, \text{inv}$ over E-constraints, and P, Q over graph programs. Together, the axioms and rules define a proof system for partial correctness. If a Hoare triple $\{c\} P \{d\}$ can be proved via the axioms and inference rules (by constructing a proof tree, as in Section 6), we write $\vdash_{\text{par}} \{c\} P \{d\}$. The proof system is sound in the sense of partial correctness, that is, $\vdash_{\text{par}} \{c\} P \{d\}$ implies $\models_{\text{par}} \{c\} P \{d\}$ (see [PP12a] for GP 1, and [Pos13] for an analogous GP 2 proof).

$$\begin{array}{c}
 \text{[ruleapp]} \frac{}{\{\text{Pre}(r, c)\} r \{c\}} \qquad \qquad \qquad \text{[nonapp]} \frac{}{\{\neg \text{App}(\mathcal{R})\} \mathcal{R} \{\text{false}\}} \\
 \\
 \text{[ruleset]} \frac{\{c\} r \{d\} \text{ for each } r \in \mathcal{R}}{\{c\} \mathcal{R} \{d\}} \qquad \qquad \qquad \text{[!]} \frac{\{\text{inv}\} \mathcal{R} \{\text{inv}\}}{\{\text{inv}\} \mathcal{R}! \{\text{inv} \wedge \neg \text{App}(\mathcal{R})\}} \\
 \\
 \text{[comp]} \frac{\{c\} P \{e\}, \{e\} Q \{d\}}{\{c\} P; Q \{d\}} \qquad \qquad \qquad \text{[cons]} \frac{c \Rightarrow c', \{c'\} P \{d'\}, d' \Rightarrow d}{\{c\} P \{d\}} \\
 \\
 \text{[if]} \frac{\{c \wedge \text{App}(\mathcal{R})\} P \{d\}, \{c \wedge \neg \text{App}(\mathcal{R})\} Q \{d\}}{\{c\} \text{ if } \mathcal{R} \text{ then } P \text{ else } Q \{d\}} \\
 \\
 \text{[try]} \frac{\{c \wedge \text{App}(\mathcal{R})\} \mathcal{R}; P \{d\}, \{c \wedge \neg \text{App}(\mathcal{R})\} Q \{d\}}{\{c\} \text{ try } \mathcal{R} \text{ then } P \text{ else } Q \{d\}}
 \end{array}$$

Figure 6: Partial correctness proof rules for core commands

Two transformations – App and Pre – appear in the axioms and rules. Intuitively, App takes as input a set \mathcal{R} of conditional rule schemata, and transforms it into an E-condition satisfied only by graphs for which at least one rule schema in \mathcal{R} is applicable. Pre on the other hand constructs an E-condition such that if $G \models \text{Pre}(r, c)$, and the application of r to G results in a graph H , then $H \models c$. Formal constructions of the transformations are omitted from this paper, but can be found in [PP12a] for GP 1 (and for GP 2, in [Pos13]).

We remark that the proof system is for a strict subset of graph programs. Specifically, as-long-as-possible iteration can only be applied to sets of rule schemata, and the guards of conditionals are restricted to sets of rule schemata (in both cases the semantics of GP allows arbitrary pro-

grams). Without this restriction, the proof rules would require an assertion language able to express that an arbitrary program will not fail.

5 Proving Total Correctness

If $\vdash_{\text{par}} \{c\} P \{d\}$, then if P is executed on a graph G satisfying c , we can be sure that any graph resulting will satisfy d . What we cannot be sure about is whether an execution of P will ever terminate (i.e. whether an execution might diverge or not). Moreover, if an execution of P does in fact terminate, we cannot be sure that it does so without failure. When referring to *total correctness*, we follow [Apt84] in meaning both absence of divergence and failure; and when referring to *weak total correctness*, we mean only absence of divergence.

Definition 3 (Weak total correctness) A graph program P is *weakly totally correct* with respect to a precondition c and postcondition d (both of which are E-constraints), denoted $\models_{\text{wtot}} \{c\} P \{d\}$, if $\vdash_{\text{par}} \{c\} P \{d\}$ and for every graph $G \in \mathcal{G}(\mathcal{L})$ such that $G \models c$, we have $\perp \notin \llbracket P \rrbracket G$. \square

Definition 4 (Total correctness) A graph program P is *totally correct* with respect to a precondition c and postcondition d (both of which are E-constraints), denoted $\models_{\text{tot}} \{c\} P \{d\}$, if $\vdash_{\text{wtot}} \{c\} P \{d\}$, and for every graph $G \in \mathcal{G}(\mathcal{L})$ such that $G \models c$, we have $\text{fail} \notin \llbracket P \rrbracket G$. \square

A weakly totally correct program executed on a graph satisfying the precondition will either produce an output graph or terminate with failure (it cannot diverge or get stuck). A totally correct program however has the additional guarantee that it will not fail, that is, a graph will eventually result from its execution.

Our proof system for weak total correctness is formed from the proof rules of Figure 6, but with $[\!]_{\text{tot}}$ in Figure 7 substituted for $[\!]$. If a triple $\{c\} P \{d\}$ can be obtained from this proof system, we write $\vdash_{\text{wtot}} \{c\} P \{d\}$. The issue of termination is localised to the proof rule for as-long-as-possible iteration: $[\!]_{\text{tot}}$ has an additional premise to $[\!]$ which handles this. It requires, for a given rule schema set, that there is a function assigning natural numbers to graphs such that these naturals are decreasing along rule applications. Such a function $\#$ is called a *termination function*. If the $\#$ -values of graphs satisfying the invariant inv decrease under applications of \mathcal{R} , we say that \mathcal{R} is *$\#$ -decreasing under inv* . These definitions are given more precisely below.

$$[\!]_{\text{tot}} \frac{\vdash_{\text{par}} \{inv\} \mathcal{R} \{inv\}, \quad \mathcal{R} \text{ is } \# \text{-decreasing under } inv}{\{inv\} \mathcal{R}! \{inv \wedge \neg \text{App}(\mathcal{R})\}}$$

$$[\text{ruleset}]_{\text{tot}} \frac{c \Rightarrow \text{App}(\mathcal{R}), \quad \vdash_{\text{par}} \{c\} r \{d\} \text{ for each } r \in \mathcal{R}}{\{c\} \mathcal{R} \{d\}}$$

Figure 7: Total correctness proof rules for two core GP commands

Definition 5 (Termination function; #-decreasing) A *termination function* is a mapping $\# : \mathcal{G}(\mathcal{L}) \rightarrow \mathbb{N}$ from (semantic) graphs to natural numbers. Given an E-constraint c , a set of conditional rule schemata \mathcal{R} is *#-decreasing under c* if for all graphs G, H in $\mathcal{G}(\mathcal{L})$ with $G \models c$ and $H \models c$,

$$G \Rightarrow_{\mathcal{R}} H \text{ implies } \#G > \#H.$$

□

In an application of $[\!]_{\text{tot}}$, one must find a suitable termination function $\#$ that returns smaller natural numbers along the graphs of direct derivations. The problem of deciding whether a set of rule schemata has a termination function or not is undecidable in general [Plu98]. Often however simple termination functions will suffice in several contexts. For example, a useful, intuitive termination function would be one that maps a graph to its size (e.g. total number of nodes and edges). If a rule schemata set is reducing the size of a graph upon each application, then clearly the iteration cannot continue indefinitely, and this is reflected by the output of $\#$ tending towards zero. However, in cases when rule schemata are not decreasing the size of the graph, less obvious termination functions may be needed (one such example will be discussed in Section 6). Note that the rule $[\!]_{\text{tot}}$ requires only that $\#$ is decreasing for graphs that satisfy the invariant *inv*, i.e. it need not be decreasing for other graphs.

Our proof system for total correctness is formed of [comp], [cons], [if], [try], and the proof rules of Figure 7. If a triple $\{c\} P \{d\}$ can be derived from this proof system, we write $\vdash_{\text{tot}} \{c\} P \{d\}$. (We do not include a proof rule for a program that is just a single rule schema r , because this case is captured by proving $\vdash_{\text{tot}} \{c\} \{r\} \{d\}$.) This proof system allows one to prove that all program executions terminate without failure. Essentially, this is achieved by ensuring that the preconditions of rule schema sets imply their applicability, when they are applied outside of iterations or the guards of conditionals. Hence if graphs satisfy the preconditions, by implication the rule schema sets are applicable to those graphs, and thus we can be certain that no execution will fail.

The proof rule $[\text{ruleset}]_{\text{tot}}$ separates the issues of failure and partial correctness. In using the proof rule, one must show (outside the calculus) that the applicability of \mathcal{R} is logically implied by the precondition c . In showing that this premise holds, we can be sure that at least one rule schema in \mathcal{R} can be applied to a graph satisfying c , hence no execution on that graph will fail. Separately, it must be shown that $\vdash_{\text{par}} \{c\} r \{d\}$ for each $r \in \mathcal{R}$, that is, each rule schema in the set is partially correct with respect to the pre- and postcondition. Together, we derive that every execution of \mathcal{R} will yield a graph, and that the graph will satisfy the postcondition.

The axiom [nonapp] is excluded from our proof system for total correctness, as the specification $\{\neg \text{App}(\mathcal{R})\} \mathcal{R} \{\text{false}\}$ does not hold in the sense of total correctness. Suppose that it did. Then \mathcal{R} would never fail on graphs satisfying the precondition. But satisfying $\neg \text{App}(\mathcal{R})$ implies that \mathcal{R} fails on that graph – a contradiction.

6 Example Proofs

In this section, we return to the example graph programs from Section 3, and demonstrate how to prove (weak) total correctness properties using our new proof calculi.

First, we revisit the program `colouring` of Figure 3. Though the program contains no failure points (since if a rule schema under as-long-as-possible iteration cannot be applied, the execution simply moves on to the next command), the iteration operator can introduce non-termination. In [PP12a] we proved that `colouring` is partially correct, in the sense that any graph resulting is properly coloured. In Figure 8, we strengthen this to $\vdash_{\text{tot}} \{c\} \text{colouring } \{d \wedge \neg \text{App}(\{\text{inc}\})\}$, i.e. if the program is executed on a graph containing only atom-labelled nodes, then (1) a graph will eventually be returned; (2) it will be properly coloured; and (3) for any colour n in the graph, every colour k with $0 \leq k < n$ is also in the graph. (The specification ignores marked nodes and edges for simplicity.) Note that the E-conditions resulting from Pre, implications in instances of [cons], and their justifications, are omitted to preserve space – but can be found in [Pos13].

$$\begin{array}{c}
\text{[ruleapp]} \frac{}{\text{[cons]} \frac{\{\text{Pre}(\text{init}, e)\} \text{init } \{e\}}{\vdash_{\text{par}} \{e\} \text{init } \{e\}} X} \\
\text{[!]}_{\text{tot}} \frac{\{e\} \text{init! } \{e \wedge \neg \text{App}(\{\text{init}\})\}}{\text{[cons]} \frac{\{c\} \text{init! } \{d\}}{\vdash_{\text{tot}} \{c\} \text{init!; inc! } \{d \wedge \neg \text{App}(\{\text{inc}\})\}}} \\
\text{[ruleapp]} \frac{}{\text{[cons]} \frac{\{\text{Pre}(\text{inc}, d)\} \text{inc } \{d\}}{\vdash_{\text{par}} \{d\} \text{inc } \{d\}} Y} \\
\text{[!]}_{\text{tot}} \frac{\{d\} \text{inc! } \{d \wedge \neg \text{App}(\{\text{inc}\})\}}{\vdash_{\text{tot}} \{c\} \text{init!; inc! } \{d \wedge \neg \text{App}(\{\text{inc}\})\}} \\
\text{[comp]} \frac{}{\vdash_{\text{tot}} \{c\} \text{init!; inc! } \{d \wedge \neg \text{App}(\{\text{inc}\})\}}
\end{array}$$

$X : \text{init}$ is $\#_{\text{init}}$ -decreasing under e ; $Y : \text{inc}$ is $\#_{\text{inc}}$ -decreasing under d

$$\begin{aligned}
c &= \forall (\textcircled{a}_1, \exists (\textcircled{a}_1 \mid \text{atom}(a))) \\
d &= \forall (\textcircled{a}_1, \exists (\textcircled{a}_1 \mid a = b:c \text{ and } \text{atom}(b) \text{ and } c \geq 0)) \\
&\quad \wedge \forall (\textcircled{b:c}_1 \mid \text{atom}(b), \exists (\textcircled{b:c}_1 \mid c = 0)) \\
&\quad \vee \exists (\textcircled{b:c}_1, \textcircled{d:c-1} \mid \text{atom}(d)) \\
e &= \forall (\textcircled{a}_1, \exists (\textcircled{a}_1 \mid \text{atom}(a))) \\
&\quad \vee \exists (\textcircled{a}_1 \mid a = b:c \text{ and } \text{atom}(b) \text{ and } c \geq 0)) \\
&\quad \wedge \forall (\textcircled{b:c}_1 \mid \text{atom}(b), \exists (\textcircled{b:c}_1 \mid c = 0)) \\
&\quad \vee \exists (\textcircled{b:c}_1, \textcircled{d:c-1} \mid \text{atom}(d)) \\
\neg \text{App}(\{\text{init}\}) &= \neg \exists (\textcircled{x} \mid \text{atom}(x)) \\
\neg \text{App}(\{\text{inc}\}) &= \neg \exists (\textcircled{x:i} \xrightarrow{k} \textcircled{y:i} \mid \text{atom}(x, y) \text{ and } \text{int}(i))
\end{aligned}$$

Figure 8: A proof tree for the program `colouring` of Figure 3

The key revision in the proof tree is in the two uses of $[\!]_{\text{tot}}$, which unlike its partial correctness counterpart requires the definition of termination functions. For `init`, we define $\#_{\text{init}}: \mathcal{G}(\mathcal{L}) \rightarrow \mathbb{N}$ to map graphs to the number of their nodes labelled by an atom. The rule schema is clearly $\#_{\text{init}}$ -decreasing under e , since every application of `init` reduces by one the number of nodes with such labels. The rule schema `inc` however requires a less obvious termination function

$\#_{\text{inc}} : \mathcal{G}(\mathcal{L}) \rightarrow \mathbb{N}$. For a graph $G \in \mathcal{G}(\mathcal{L})$, we define:

$$\#_{\text{inc}}(G) = \sum_{i=0}^{|V_G|} i - \sum_{v \in V_G} \text{colour}(v)$$

where $\text{colour}(v)$ for a node $v \in V_G$ is defined:

$$\text{colour}(v) = \begin{cases} i & \text{if } l_G(v) = x:i \text{ with } x \in \mathbb{Z} \cup \text{Char}^*, i \in \mathbb{N}; \\ 0 & \text{otherwise.} \end{cases}$$

We show that inc is $\#_{\text{inc}}$ -decreasing under inv . Observe that if G is a graph with $\text{colour}(v) = 0$ for every node v in V_G , then for every derivation $G \Rightarrow_{\text{inc}}^* H$ there is some $0 \leq k < |V_H|$ such that k is the largest tag in V_H . We obtain an upper bound for the second summation:

$$\sum_{v \in V_H} \text{colour}(v) \leq 0 + 1 + \dots + (|V_H| - 1) = 0 + 1 + \dots + (|V_G| - 1) < \sum_{i=0}^{|V_G|} i.$$

Since $\sum_{v \in V_H} \text{colour}(v)$ equals the number of rule schema applications in $G \Rightarrow_{\text{inc}}^* H$, it follows that inc must eventually terminate (as it approaches the upper bound). By subtracting the summation from the upper bound, we instead have a number decreasing towards 0 after every application of inc . Hence $\#_{\text{inc}}$ is a suitable termination function, and inc is $\#_{\text{inc}}$ -decreasing under inv .

We remark that $\text{inc}!$ will terminate on any graph – not just those satisfying inv . A termination function however is harder to write without the assumptions the invariant allows us to make about the graphs.

Now, we return to the program reachable? of Figure 4, which unlike earlier, can fail on some input graphs (in particular, those graphs omitting the pair of 1- and 2-tagged nodes). We give a proof tree³ for $\vdash_{\text{tot}} \{c\} \text{reachable? } \{c\}$ in Figure 9, where the E-constraints are as in Figure 10. For clarity, we let $\text{visited}(p)$ abbreviate:

$$p = a:0 \text{ and } \text{atom}(a)$$

where a is a fresh variable.

If the program is executed on a graph that contains only atom-labelled nodes but with one tagged 1 and another tagged 2, then (1) the program is guaranteed to return a graph eventually; and (2) that graph will satisfy the same condition (i.e. an invariant). Again, due to space limitations, we have omitted the implications in instances of $[\text{cons}]$ and their justifications. Moreover, we have only provided one of the E-constraints generated by Pre .

In this proof tree, there are simple suitable termination functions $\#_p, \#_u$. We define the termination function $\#_p : \mathcal{G}(\mathcal{L}) \rightarrow \mathbb{N}$ (resp. $\#_u$) to return the number of nodes in a graph that are labelled by an atom (resp. number of atom-labelled nodes tagged with a 0). That is, both termination functions exploit that each application of their respective rule schema explicitly reduces the number of remaining matches.

³ For simplicity we use an obvious additional axiom $[\text{skip}] : \vdash_{\text{tot}} \{c\} \text{skip } \{c\}$.

Let $P = \text{if reachable then skip else addlink}$

$$\begin{array}{c}
 \text{[ruleapp]} \frac{}{\{\text{Pre}(\text{propagate}, e)\} \text{propagate } \{e\}} \\
 \text{[cons]} \frac{}{\{\text{Pre}(\text{propagate}, e)\} \text{propagate } \{e\}} \\
 \text{[!]}_{\text{tot}} \frac{\text{[par]} \frac{}{\{e\} \text{propagate } \{e\}} \quad \text{propagate is } \#_p\text{-decreasing under } e}{\text{[cons]} \frac{}{\{e\} \text{propagate! } \{e \wedge \neg \text{App}(\{\text{propagate}\})\}}} \\
 \text{[comp]} \frac{\text{[cons]} \frac{}{\{c\} \text{propagate! } \{e\}} \quad \text{Subtree } X}{\text{[!]}_{\text{tot}} \frac{}{\{c\} \text{propagate!}; P; \text{undo! } \{c\}}} \\
 \\
 \text{Subtree } X: \\
 \\
 \text{[ruleapp]} \frac{}{\{\text{Pre}(\text{undo}, e)\} \text{undo } \{e\}} \\
 \text{[cons]} \frac{}{\{\text{Pre}(\text{undo}, e)\} \text{undo } \{e\}} \\
 \text{[!]}_{\text{tot}} \frac{\text{[par]} \frac{}{\{e\} \text{undo } \{e\}} \quad \text{undo is } \#_u\text{-decreasing under } e}{\text{[cons]} \frac{}{\{e\} \text{undo! } \{e \wedge \neg \text{App}(\{\text{undo}\})\}}} \\
 \text{[comp]} \frac{\text{[skip]} \frac{}{\{e\} \text{skip } \{e\}} \quad \text{Subtree } Y \quad \text{[cons]} \frac{}{\{e\} \text{undo! } \{e \wedge \neg \text{App}(\{\text{undo}\})\}}}{\text{[if]} \frac{}{\{e \wedge \text{App}(\{\text{reachable}\})\} \text{skip } \{e\}} \quad \text{[!]}_{\text{tot}} \frac{}{\{e\} \text{undo! } \{c\}}} \\
 \text{[comp]} \frac{}{\{e\} P \{e\}} \quad \text{[!]}_{\text{tot}} \frac{}{\{e\} \text{undo! } \{c\}} \\
 \text{[!]}_{\text{tot}} \frac{}{\{e\} P; \text{undo! } \{c\}} \\
 \\
 \text{Subtree } Y: \\
 \\
 \text{[ruleapp]} \frac{}{\{\text{Pre}(\text{addlink}, e)\} \text{addlink } \{e\}} \\
 \text{[cons]} \frac{}{\{\text{Pre}(\text{addlink}, e)\} \text{addlink } \{e\}} \\
 \text{[!]}_{\text{tot}} \frac{\text{[par]} \frac{}{\{e \wedge \neg \text{App}(\{\text{reachable}\})\} \text{addlink } \{e\}}}{\text{[ruleset]}_{\text{tot}} \frac{}{\{e \wedge \neg \text{App}(\{\text{reachable}\}) \Rightarrow \text{App}(\{\text{addlink}\})\} \text{addlink } \{e\}}} \\
 \text{[!]}_{\text{tot}} \frac{}{\{e \wedge \neg \text{App}(\{\text{reachable}\})\} \text{addlink } \{e\}}
 \end{array}$$

Figure 9: Total correctness proof tree for the program `reachable?` of Figure 4

$$\begin{aligned}
 c &= \exists(\textcircled{x:1} \textcircled{y:2} \mid \text{atom}(x, y), \\
 &\quad \neg \exists(\textcircled{x:1} \textcircled{y:2} \textcircled{p} \mid \text{not atom}(p))) \\
 e &= \exists(\textcircled{x:1} \textcircled{y:2} \mid \text{atom}(x, y), \\
 &\quad \neg \exists(\textcircled{x:1} \textcircled{y:2} \textcircled{p} \mid \text{not atom}(p) \\
 &\quad \text{and not visited}(p))) \\
 \text{App}(\{\text{reachable}\}) &= \exists(\textcircled{x:y} \xrightarrow{a} \textcircled{z:2} \mid \text{atom}(x, z) \text{ and int}(y)) \\
 \text{App}(\{\text{addlink}\}) &= \exists(\textcircled{x:1} \textcircled{y:2} \mid \text{atom}(x, y)) \\
 \neg \text{App}(\{\text{propagate}\}) &= \neg \exists(\textcircled{x:y} \xrightarrow{a} \textcircled{z} \mid \text{atom}(x, z), \\
 &\quad \exists(\textcircled{x:y} \xrightarrow{a} \textcircled{z} \mid y = 1) \\
 &\quad \vee \exists(\textcircled{x:y} \xrightarrow{a} \textcircled{z} \mid y = 0)) \\
 \neg \text{App}(\{\text{undo}\}) &= \neg \exists(\textcircled{x:0} \mid \text{atom}(x)) \\
 \text{Pre}(\text{undo}, e) &\equiv \forall(\textcircled{x:0} \mid \text{atom}(x), \exists(\textcircled{x:0} \textcircled{y:1} \textcircled{z:2} \mid \text{atom}(y, z), \\
 &\quad \neg \exists(\textcircled{x:0} \textcircled{y:1} \textcircled{z:2} \textcircled{p} \mid \text{not atom}(p) \\
 &\quad \text{and not visited}(p))))
 \end{aligned}$$

Figure 10: Partial list of E-constraints for Figure 9

The rule schema `addlink` is the only potential failure point of the program, and is addressed in the proof tree by the application of $[\text{ruleset}]_{\text{tot}}$. It must be shown that the precondition at that point implies the applicability of `addlink`. From Figure 10, it is clear that satisfying e is sufficient to deduce the applicability of `addlink`.

7 Soundness and Completeness for Termination

In this section we revise our soundness proof from [PP12a] to account for (weak) total correctness, before showing that any iterating rule schemata set that terminates can be proven to terminate by the rule $[\text{!}]_{\text{tot}}$. Soundness is relative to the operational semantics of the language. An updated version of the soundness proof for partial correctness with regards to the GP 2 semantics can be found in [Pos13].

Theorem 1 (Soundness of \vdash_{wtot}) For all graph programs P and E-constraints c, d , we have that $\vdash_{\text{wtot}} \{c\} P \{d\}$ implies $\models_{\text{wtot}} \{c\} P \{d\}$. \square

Proof. For all weak total correctness proof rules except $[\text{!}]_{\text{tot}}$, this follows from (1) the soundness result for partial correctness in [PP12a], and (2) the semantics of graph programs, from which it is clear that only as-long-as-possible iteration can introduce divergence.

Let \mathcal{R} be a set of (conditional) rule schemata, inv an E-constraint, and $\#$ a termination function. Assume $\vdash_{\text{par}} \{inv\} \mathcal{R} \{inv\}$. By soundness for partial correctness, we have $\models_{\text{par}} \{inv\} \mathcal{R} \{inv \wedge \neg \text{App}(\mathcal{R})\}$. Assume also that \mathcal{R} is $\#$ -decreasing under inv . By Definition 5, for all graphs

$G, H \in \mathcal{G}(\mathcal{L})$ with $G \models \text{inv}$ and $H \models \text{inv}$, $G \Rightarrow_{\mathcal{R}} H$ implies $\#G > \#H$. Assume that $\mathcal{R}!$ diverges for any such G . Since \mathcal{R} is $\#$ -decreasing under inv , every derivation step yields a graph for which $\#$ returns a smaller natural number. Since $\mathcal{R}!$ diverges, there are infinitely many derivation steps. But from any natural n , there are only finitely many smaller numbers. A contradiction. It cannot be the case that $\mathcal{R}!$ diverges from any such G . Hence $\models_{\text{wtot}} \{\text{inv}\} \mathcal{R}! \{\text{inv} \wedge \neg \text{App}(\mathcal{R})\}$. \square

Theorem 2 (Soundness of \vdash_{tot}) For all graph programs P and E-constraints c, d , we have that $\vdash_{\text{tot}} \{c\} P \{d\}$ implies $\models_{\text{tot}} \{c\} P \{d\}$. \square

Proof. For the proof rules [comp], [cons], [if], [try], [!]_{tot}, this follows from (1) the soundness of \vdash_{wtot} (see Theorem 1), and (2) the semantics of graph programs, from which it is clear that these proof rules are sound in the sense of total correctness. What remains to be shown is the soundness of [ruleset]_{tot} in the sense of total correctness.

Let \mathcal{R} denote a set of (conditional) rule schemata and c, d denote E-constraints. Assume that $\vdash_{\text{par}} \{c\} r \{d\}$ for each $r \in \mathcal{R}$. Then by soundness for partial correctness, we have $\models_{\text{par}} \{c\} \mathcal{R} \{d\}$. Now assume the validity of $c \Rightarrow \text{App}(\mathcal{R})$. Then if a graph $G \in \mathcal{G}(\mathcal{L})$ satisfies c , by assumption it will satisfy $\text{App}(\mathcal{R})$. By Proposition 7.1 of [PP12a] (updated for the GP 2 syntax in [Pos13]), there is a graph H such that $G \Rightarrow_{\mathcal{R}} H$. Then the semantic rule [call₁] will be applied (and in particular, [call₂] will not be), hence a graph is guaranteed from the execution and failure is avoided. We yield $\models_{\text{tot}} \{c\} \mathcal{R} \{d\}$. \square

Now, we show that every iterating set of rule schemata that terminates can be proven to terminate using [!]_{tot}, by showing that there always exists a termination function for which the rule schemata set is decreasing under its invariant.

Theorem 3 (Completeness of [!]_{tot} for termination) Let \mathcal{R} be a set of conditional rule schemata and c be an E-constraint such that for every graph G in $\mathcal{G}(\mathcal{L})$, $G \models c$ implies that $\mathcal{R}!$ cannot diverge from G . Then there exists a termination function $\#$ such that \mathcal{R} is $\#$ -decreasing under c . \square

Proof. Let G be a graph such that $G \models c$. Then there cannot exist an infinite sequence $G \Rightarrow_{\mathcal{R}} G_1 \Rightarrow_{\mathcal{R}} G_2 \Rightarrow_{\mathcal{R}} \dots$ as otherwise, by the semantics of GP, there would be an infinite sequence $\langle \mathcal{R}!, G \rangle \rightarrow \langle \mathcal{R}!, G_1 \rangle \rightarrow \langle \mathcal{R}!, G_2 \rangle \dots$. To define the termination function $\#$, we show that the length of $\Rightarrow_{\mathcal{R}}$ -derivations starting from G is bounded. (Note that, in general, a terminating relation need not be bounded.)

We exploit that $\Rightarrow_{\mathcal{R}}$ is closed under isomorphism in the following sense: given graphs M, M', N and N' such that $M \cong M'$ and $N \cong N'$, then $M \Rightarrow_{\mathcal{R}} N$ implies $M' \Rightarrow_{\mathcal{R}} N'$. Hence we can lift $\Rightarrow_{\mathcal{R}}$ to a relation on isomorphism classes of graphs by defining: $[M] \Rightarrow_{\mathcal{R}} [N]$ if $M \Rightarrow_{\mathcal{R}} N$. Then, since \mathcal{R} is finite, for every isomorphism class $[M]$ the set $\{[N] \mid [M] \Rightarrow_{\mathcal{R}} [N]\}$ is finite.

Now, since there is no infinite sequence of $\Rightarrow_{\mathcal{R}}$ -steps starting from $[G]$, it follows from König's lemma [Kön36] that the length of $\Rightarrow_{\mathcal{R}}$ -derivations starting from $[G]$ is bounded. (In the tree of all derivations starting from $[G]$, all nodes have a finite degree. Hence the tree cannot be infinite, as otherwise it would contain an infinite derivation.) Hence the length of $\Rightarrow_{\mathcal{R}}$ -derivations starting from G is bounded as well. In general, given any graph M in $\mathcal{G}(\mathcal{L})$, let $\#M$ be the length of a longest $\Rightarrow_{\mathcal{R}}$ -derivation starting from M if $M \models c$, and $\#M = 0$ otherwise. Then if $M, N \models c$ and

$M \Rightarrow_{\mathcal{R}} N$, we have $\#M > \#N$. Thus \mathcal{R} is $\#$ -decreasing under c . □

8 Conclusion

In this paper we have presented two Hoare calculi which allow one to prove (weak) total correctness. Both proof systems have been shown to be sound. We have shown how to reason about termination via termination functions, and shown that the proof rule for termination is complete in the sense that all terminating loops (having a set of conditional rule schemata as the body) can be proven to be terminating. Finally, we have demonstrated the use of the proof systems on two non-trivial graph programs, showing how to prove the absence of divergence and failure.

Future work will explore how to implement the proof calculi in an interactive proof system. A first step towards this is made in [Pos13], in which translations from E-conditions to many-sorted formulae (and back) are defined, providing a suitable front-end logic for an implemented verification system. Future work will also address the question of whether or not the calculi are (relatively) complete. It would also be worthwhile to integrate a stronger assertion language into the calculi that can express non-local properties, such as the hyperedge-replacement conditions of [HR10].

Acknowledgements: We thank the anonymous referees for their comments which helped to improve the final version of this paper. Most of this work was completed whilst the first author was a Ph.D. student at the University of York, where he was supported by a scholarship of the Engineering and Physical Sciences Research Council. Since January 2013 he has been based at ETH Zürich, where he is supported with funding from the European Research Council (ERC grant agreement no. 291389).

Bibliography

- [Apt84] K. R. Apt. Ten Years of Hoare's Logic: A Survey Part II: Nondeterminism. *Theoretical Computer Science* 28:83–109, 1984.
- [BCK08] P. Baldan, A. Corradini, B. König. A framework for the verification of infinite-state graph transformation systems. *Information and Computation* 206(7):869–907, 2008.
- [BHE09] D. Bisztray, R. Heckel, H. Ehrig. Compositional Verification of Architectural Refactorings. In *Proc. Architecting Dependable Systems VI (WADS 2008)*. Volume 5835, pp. 308–333. Springer-Verlag, 2009.
- [CR12] S. A. da Costa, L. Ribeiro. Verification of graph grammars using a logical approach. *Science of Computer Programming* 77(4):480–504, 2012.
- [HP02] A. Habel, D. Plump. Relabelling in Graph Transformation. In *Proc. International Conference on Graph Transformation (ICGT 2002)*. Volume 2505, pp. 135–147. Springer-Verlag, 2002.

- [HP09] A. Habel, K.-H. Pennemann. Correctness of high-level transformation systems relative to nested conditions. *Mathematical Structures in Computer Science* 19(2):245–296, 2009.
- [HPR06] A. Habel, K.-H. Pennemann, A. Rensink. Weakest Preconditions for High-Level Programs. In *Proc. International Conference on Graph Transformation (ICGT 2006)*. Volume 4178, pp. 445–460. Springer-Verlag, 2006.
- [HR10] A. Habel, H. Radke. Expressiveness of graph conditions with variables. In *Proc. Colloquium on Graph and Model Transformation on the Occasion of the 65th Birthday of Hartmut Ehrig*. Electronic Communications of the EASST 30. 2010.
- [KE10] B. König, J. Esparza. Verification of Graph Transformation Systems with Context-Free Specifications. In *Proc. International Conference on Graph Transformation (ICGT 2010)*. Volume 6372, pp. 107–122. Springer-Verlag, 2010.
- [Kön36] D. König. Sur les correspondances multivoques des ensembles. *Fundamenta Mathematicae* 8:114–134, 1936.
- [Plu98] D. Plump. Termination of Graph Rewriting is Undecidable. *Fundamenta Informaticae* 33(2):201–209, 1998.
- [Plu09] D. Plump. The Graph Programming Language GP. In *Proc. Algebraic Informatics (CAI 2009)*. Volume 5725, pp. 99–122. Springer-Verlag, 2009.
- [Plu12] D. Plump. The Design of GP 2. In *Proc. International Workshop on Reduction Strategies in Rewriting and Programming (WRS 2011)*. Electronic Proceedings in Theoretical Computer Science 82, pp. 1–16. 2012.
- [Pos13] C. M. Poskitt. *Verification of Graph Programs*. PhD thesis, The University of York, 2013. To appear.
- [PP12a] C. M. Poskitt, D. Plump. Hoare-Style Verification of Graph Programs. *Fundamenta Informaticae* 118(1-2):135–175, 2012.
- [PP12b] C. M. Poskitt, D. Plump. Verifying Total Correctness of Graph Programs. In *Proc. International Workshop on Graph Computation Models (GCM 2012)*. 2012. Available from <http://gcm2012.imag.fr/>.