

## RULE-BASED GRAPH PROGRAMS MATCHING THE TIME COMPLEXITY OF IMPERATIVE ALGORITHMS

ZIAD ISMAILI ALAOUI <sup>a,b</sup> AND DETLEF PLUMP <sup>c</sup>

<sup>a</sup> University of Liverpool, United Kingdom  
*e-mail address:* [ziad.ismaili-alaoui@liverpool.ac.uk](mailto:ziad.ismaili-alaoui@liverpool.ac.uk)

<sup>b</sup> Most of this author’s work was done while he was affiliated with the University of York.

<sup>c</sup> University of York, United Kingdom  
*e-mail address:* [detlef.plump@york.ac.uk](mailto:detlef.plump@york.ac.uk)

**ABSTRACT.** We report on recent advances in rule-based graph programming, which allow us to match the time complexity of some fundamental imperative graph algorithms. In general, achieving the time complexity of graph algorithms implemented in conventional languages using a rule-based graph-transformation language is challenging due to the cost of graph matching. Previous work demonstrated that with *rooted* rules, certain algorithms can be implemented in the graph programming language GP2 such that their runtime matches the time complexity of imperative implementations. However, this required input graphs to have a bounded node degree and (for some algorithms) to be connected. In this paper, we overcome these limitations by enhancing the graph data structure generated by the GP2 compiler and exploiting the new structure in programs. We present three case studies: the first program checks whether input graphs are connected, the second program checks whether input graphs are acyclic, and the third program solves the single-source shortest-path problem for graphs with integer edge-weights. The first two programs run in linear time on (possibly disconnected) input graphs with arbitrary node degrees. The third program runs in time  $O(nm)$  on arbitrary input graphs, matching the time complexity of imperative implementations of the Bellman-Ford algorithm. For each program, we prove its correctness and time complexity, and provide runtime experiments on various graph classes.

### 1. INTRODUCTION

Designing and implementing languages for rule-based graph rewriting, such as GReAT [AKN<sup>+</sup>06], GROOVE [GdMR<sup>+</sup>12], GrGen.Net [JBK10], Henshin [SBG<sup>+</sup>17], and PORGY [FKP19], poses significant performance challenges. Typically, programs written in these languages do not achieve the same asymptotic runtime efficiency as those written in conventional imperative languages such as C or Java. The primary obstacle is the cost of graph matching, where matching the left-hand graph  $L$  of a rule within a host graph  $G$  generally requires time  $|G|^{|L|}$ , with  $|X|$  denoting the size of graph  $X$ . (Since  $L$  is fixed, this is a polynomial.) As a consequence, standard imperative graph algorithms running in linear time (see, for example, [CLRS22, Ski20]) may exhibit non-linear, polynomial runtimes when recast as rule-based graph programs.

To address this issue, the graph programming language GP 2 [Plu12] supports *rooted* graph transformation rules, initially proposed by Dörr [Dör95]. This approach involves designating certain nodes as *roots* and matching them with roots in the host graphs. Consequently, only the neighbourhoods of host graph roots need to be searched for matches, which can often be done in constant time under mild conditions. The GP 2 compiler [Bak15] maintains a list of pointers to roots in the host graph, facilitating constant-time access to roots if their number remains bounded throughout the program’s execution. In [BP12], *fast* rules were identified as a class of rooted rules that can be applied in constant time, provided host graphs contain a bounded number of roots and have a bounded node degree.

The first linear-time graph problem implemented by a GP 2 program was 2-colouring. In [BP12, Bak15], it is shown that this program colours connected graphs of bounded node degree in linear time. Since then, the GP 2 compiler has received some major improvements, particularly related to the runtime graph data structure used by the compiled programs [CRP20]. These improvements made a linear-time worst-case performance possible for a wider class of programs on input graph classes of bounded degree, in some cases even on those of unbounded degree. See [CCP22] for an overview.

Prior to the improvements presented in this paper, besides the issue of input graph classes with unbounded node degree, some programs, such as the aforementioned 2-colouring program, needed a non-linear runtime on disconnected input graphs. In this paper, we present two updates to the GP 2 compiler which allow lifting the constraints that input graphs (1) must be connected and (2) have a bounded node degree. In short, the solution is to improve the graph data structure generated by the compiler. Nodes are now stored in separate linked lists based on their *marks* (red, green, blue, grey or unmarked), and each node comes with a two-dimensional array of linked lists storing all incident edges based on their marks (red, green, blue, dashed or unmarked) and orientation (incoming, outgoing or looping). This enables the matching algorithm to find in constant time a node with a specific mark or an edge with a specific mark and orientation. For instance, if a red node is needed, a single access to the list of red nodes will either locate such a node or confirm its absence. Similarly, if an outgoing green edge is required, a single access to the corresponding linked list will either find such an edge or determine that none exists.

In addition to the new graph data structure, a programming technique is needed to take advantage of the improved storage. In our first case study, we demonstrate how the new approach allows checking in linear time whether arbitrary input graphs are connected. In the second case study, we provide a program that recognises acyclic graphs in linear time. Again, input graphs may be disconnected and have arbitrary node degrees. Finally, our third case study shows how to implement the Bellman-Ford single-source shortest-path algorithm in GP 2 such that it runs in time  $O(nm)$  on arbitrary input graphs (where  $n$  and  $m$  are the numbers of nodes and edges, respectively), thus matching the complexity of imperative versions of the algorithm. Each case study comes with proofs of the program’s correctness and time complexity. In addition, we provide runtime experiments on various graph classes.

Building upon the improvements to the compiler presented in this paper, this work further expands previous studies on fast rule-based graph programs [CCP22, AP24, AP25] with the following contributions: (1) We present a GP 2 program checking connectedness in linear time on arbitrary input graphs, eliminating the constraint of [CCP22] that input graphs must have a bounded node degree. We prove that the program is correct and always runs in linear time. In addition, we present empirical runtime measurements on six classes of input graphs. (2) We prove the correctness and time complexity of the linear-time

cycle checking program from [AP25]. Again, the proofs are complemented by runtime measurements on various classes of input graphs. (3) For the first time, we present a rule-based implementation of Bellman-Ford’s single-source shortest-paths algorithm such that the imperative time complexity of  $O(nm)$  is matched. We give full proofs of the program’s correctness and time complexity, and also provide runtime experiments confirming the complexity.

The rest of this paper is organised as follows. In Section 2, we recall the basic features of the graph programming language GP 2. In the next section, we briefly describe the GP 2-to-C compiler and state assumptions on the time complexity of elementary operations of the generated code. In Section 4, we explain the problem caused by classes of input graphs with an unbounded node degree, while Section 5 introduces the compiler improvements. Section 6, Section 7 and Section 8 present our three case studies. Finally, Section 9 concludes the paper and mentions some topics for future work. In addition, Appendix A lists the abstract syntax of GP 2 labels and application conditions.

## 2. THE GRAPH PROGRAMMING LANGUAGE GP 2

This section (largely adapted from [CCP22]) introduces GP 2, an experimental programming language based on graph-transformation rules. The language was defined in [Plu12], an up-to-date version of its concrete syntax can be found in [Bak15].

**2.1. Graphs and Rules.** GP 2 programs transform input graphs into output graphs, as depicted in Figure 1. Graphs are directed and may contain parallel edges and loops. Both nodes and edges are labelled.




Figure 1: GP 2 computational model.

**Definition 2.1** (Graph). Let  $\mathcal{L}$  be a set of labels. A *graph* over  $\mathcal{L}$  is a system  $G = \langle V_G, E_G, s_G, t_G, l_G, m_G \rangle$ , where  $V_G$  and  $E_G$  are finite sets of nodes (or vertices) and edges,  $s_G: E_G \rightarrow V_G$  and  $t_G: E_G \rightarrow V_G$  are source and target functions for edges,  $l_G: V_G \rightarrow \mathcal{L}$  is the partial node labelling function, and  $m_G: E_G \rightarrow \mathcal{L}$  is the (total) edge labelling function.

A graph  $G$  is *totally labelled* if  $l_G$  is a total function. Unlabelled nodes only occur in the interfaces of rules and are used to relabel nodes (see Definition 2.3 below). There is no need to relabel edges as they can always be deleted and reinserted with different labels.

We distinguish two types of graphs, *host graphs* and *rule graphs*. The former are the graphs on which GP 2 programs are executed while the latter are the graphs occurring in rules. Figure 37 and Figure 38 in Appendix A give grammars defining the abstract syntax of labels in host graphs and rule graphs, respectively.

Both nodes and edges are labelled with list expressions denoting integers and character strings. This includes the empty list as a special case; by convention, we draw such items in pictures without labels. Labels are of type `list`, `atom`, `int`, `string` or `char`, where `atom` is




Figure 2: GP 2 label subtype hierarchy.

the union of `int` and `string`, and `list` is the type of a list of atoms. Lists of length one and strings of length one are equated with their entry. As a consequence, every expression of category Atom, Integer, String or Char as defined by the grammar in Figure 38 can be considered as a list expression. The label subtype hierarchy of GP 2 is shown in Figure 2.

Labels in host graphs as defined by the grammar in Figure 37 form a subset of the labels allowed in rule graphs. They represent constant values in that they do not contain variables or unevaluated expressions.

Besides carrying expressions, nodes and edges can be *marked* red, green or blue. In addition, nodes can be marked grey and edges can be dashed. When we draw graphs, marks are represented graphically. The wildcard mark `any`, represented by the magenta colour, can only be used in rules and matches arbitrary marks in host graphs. Marks are convenient, among other things, to record visited items during a graph traversal and to encode auxiliary structures in graphs. The programs in the case studies of this paper use marks extensively.

The principal programming construct in GP 2 are conditional graph transformation rules labelled with expressions. The formal definition of rules makes use of the following subset of list expressions, which specifies what expressions may occur in the left-hand graph of a rule.

**Definition 2.2** (Simple Expression). A list expression  $e$  is *simple* if (1)  $e$  contains no arithmetic, degree or length operators (with the possible exception of a unary minus preceding a sequence of digits), (2)  $e$  contains at most one occurrence of a list variable, and (3) each occurrence of a string expression in  $e$  contains at most one occurrence of a string variable.

We call a label *simple* if its list expression is simple. We can now define GP 2's conditional graph transformation rules.

**Definition 2.3** (Rule). A *rule*  $\langle L \leftarrow K \rightarrow R, c \rangle$  consists of totally labelled rule graphs  $L$  and  $R$ , a graph  $K$  consisting of unlabelled nodes only, where  $V_K \subseteq V_L$  and  $V_K \subseteq V_R$ , and a condition  $c$  according to the grammar of Figure 39. Graph  $K$  is called the *interface* of the rule. We require that all labels in  $L$  are simple and that all variables occurring in  $R$  also occur in  $L$ . In addition, all variables occurring in  $c$  must occur in  $L$ .

When a rule is represented graphically, the interface  $K$  is given implicitly by the node identifiers that occur in  $L$  and  $R$  (which must coincide). We use small numbers attached to nodes as identifiers. For example, in the rule `link` of Figure 5,  $K$  consists of three unlabelled nodes with identifiers 1, 2 and 3. In general, interface nodes are to be preserved, nodes in  $L - K$  are to be deleted and nodes in  $R - K$  are to be created. Node labels such as `x`, `y` and `z` are written inside nodes, whereas small integers below nodes are their identifiers. The

constraint that all variables in  $R$  must also occur in  $L$  ensures that for a given match of  $L$  in a host graph, applying the rule produces a unique graph (up to isomorphism). Similarly, the evaluation of the condition  $c$  has a unique result if all its variables occur in  $L$ .

Next we introduce premorphisms which locate a rule's left-hand graph in a host graph.

**Definition 2.4** (Premorphism). Given graphs  $L$  and  $G$ , a *premorphisms*  $g: L \rightarrow G$  consists of two functions  $g_V: V_L \rightarrow V_G$  and  $g_E: E_L \rightarrow E_G$  that preserve sources and targets, that is,  $s_G \circ g_E = g_V \circ s_L$  and  $t_G \circ g_E = g_V \circ t_L$ .

We also need to introduce assignments, which instantiate variables in labels with host graph values of the same type.

**Definition 2.5** (Assignment). An *assignment* is a family of mappings

$$\alpha = (\alpha_X)_{X \in \{L, A, I, S, C\}}$$

where  $\alpha_L: \text{LVar} \rightarrow \text{HostList}$ ,  $\alpha_A: \text{AVar} \rightarrow \text{HostAtom}$ ,  $\alpha_I: \text{IVar} \rightarrow \text{HostInteger}$ ,  $\alpha_S: \text{SVar} \rightarrow \text{HostString}$ , and  $\alpha_C: \text{CVar} \rightarrow \text{HostChar}$ . We may omit the subscripts of these mappings as exactly one of them is applicable to a given variable.

Given an expression  $e \in \text{List}$ , an assignment  $\alpha$  and a premorphism  $g: L \rightarrow G$  from a rule graph  $L$  to a host graph  $G$ , the *instance* of  $e$  with respect to  $g$  and  $\alpha$  is obtained by (1) replacing each variable  $x$  with  $\alpha(x)$ , (2) replacing each node identifier  $n$  with  $g_V(n)$ , and (3) evaluating the resulting expression according to the meaning of the operators in the grammar of Figure 38. The result is an expression in  $\text{HostList}$  which we denote by  $e^{g, \alpha}$ .

For example, if  $e = \mathbf{x} : 1 : \mathbf{x}/2$  and  $\alpha(\mathbf{x}) = 3$  then  $e^{g, \alpha} = 3 : 1 : 1$  (note that  $'/'$  is integer division). If  $e = \mathbf{indeg}(1) + \mathbf{outdeg}(1)$  then  $e^{g, \alpha}$  is the sum of the indegree and outdegree of  $g_V(1)$  in  $G$ .

The evaluation of a condition  $c$  with respect to  $g$  and  $\alpha$  proceeds analogously to the evaluation of list expressions but results in a value  $c^{g, \alpha}$  which is either 'true' or 'false'. For example, if  $c = \mathbf{int}(\mathbf{w})$  or  $\mathbf{edge}(1, 2)$  and  $\alpha(\mathbf{w}) = 1 : 2$  then  $c^{g, \alpha} = \text{true}$  if and only if there is an edge in  $G$  from  $g_V(1)$  to  $g_V(2)$  (because  $\mathbf{int}(\mathbf{w})^{g, \alpha} = \text{false}$ ).

In the definitions below, given a node or edge  $u$  in a graph  $M$ , we write  $\text{mark}_M(u)$  for the mark of  $u$  (if it exists).

**Definition 2.6** (Compatibility). A premorphism  $g: L \rightarrow G$  and an assignment  $\alpha$  are *compatible* if for all nodes and edges  $u$  in  $L$ , (1)  $l_L(u) = e \in \text{List}$  implies  $l_G(g(u)) = e^{g, \alpha}$ , (2)  $l_L(u) = em$  with  $e \in \text{List}$  and  $m \in \text{HostMark}$  implies  $l_G(g(u)) = e^{g, \alpha}m$ , and (3)  $l_L(u) = e$  **any** with  $e \in \text{List}$  implies  $l_G(g(u)) = e^{g, \alpha}m$  for some  $m \in \text{HostMark}$ .

**Definition 2.7** (Rule Application). A rule  $r = \langle L \leftarrow K \rightarrow R, c \rangle$  is applied to a host graph  $G$  as follows:

- (1) Find an injective premorphism  $g: L \rightarrow G$  and an assignment  $\alpha$  such that  $g$  and  $\alpha$  are compatible.
- (2) Check that  $g$  satisfies the *dangling condition*: no node in  $g_V(V_L - V_K)$  is incident to an edge in  $E_G - g_E(E_L)$ .
- (3) Check that  $c^{g, \alpha} = \text{true}$ .
- (4) Construct a host graph  $H$  from  $G$  as follows:
  - (a) Remove all edges in  $g_E(E_L)$  and all nodes in  $g_V(V_L) - g_V(V_K)$ , obtaining a subgraph  $D$  of  $G$ .
  - (b)  $V_H = V_D + (V_R - V_K)$  and  $E_H = E_D + E_R$ .

- (c) For each edge  $e \in E_D$ ,  $s_H(e) = s_D(e)$ . For each edge  $e \in E_R$ ,  $s_H(e) = s_R(e)$  if  $s_R(e) \in V_R - V_K$ , otherwise  $s_H(e) = g_V(s_R(e))$ . Targets are defined analogously.
- (d) For each edge  $e \in E_D$ ,  $m_H(e) = m_D(e)$ . For each edge  $e \in E_R$ ,  $m_H(e) = \underline{\text{if}} m_R(e) \in \text{List}$  then  $m_R(e)^{g,\alpha}$  else if  $m_R(e) = dm$  with  $m \in \text{Hostmark}$  then  $d^{g,\alpha}m$ .
- (e) For each node  $v \in V_D - g_V(V_K)$ ,  $l_H(v) = l_D(v)$ . For each node  $v \in V_R - V_K$ ,  $l_H(v) = \underline{\text{if}} l_R(v) \in \text{List}$  then  $l_R(v)^{g,\alpha}$  else if  $l_R(v) = dm$  with  $m \in \text{Hostmark}$  then  $d^{g,\alpha}m$ . For each node  $v \in g_V(V_K)$  with  $g_V(\bar{v}) = v$ ,  $l_H(v) = \underline{\text{if}} l_R(\bar{v}) \in \text{List}$  then  $l_R(\bar{v})^{g,\alpha}$  else if  $l_R(\bar{v}) = dm$  with  $m \in \text{Hostmark}$  then  $d^{g,\alpha}m$  else if  $l_R(\bar{v}) = d$  any then  $d^{g,\alpha}\text{mark}_G(v)$ .

We write  $G \Rightarrow_{r,g} H$  to express that  $H$  results from  $G$  by applying  $r$  with *match*  $g$ .

Alternatively, rule application can be defined as a two-stage process in which first the rule  $\langle L \leftarrow K \rightarrow R, c \rangle$  is instantiated to  $\langle L^{g,\alpha} \leftarrow K \rightarrow R^{g,\alpha}, c^{g,\alpha} \rangle$  which is a conditional rule in the double-pushout approach with relabelling [HP02]. In the second stage, the instantiated rule is applied to the host graph by constructing two suitable pushouts in the category of partially labelled graphs. See [Cou23] for details.

**2.2. Programs.** Figure 3 defines the abstract syntax of GP 2 programs. We omit the syntax of rule declarations, rule identifiers and procedure identifiers, which can be found in [Bak15]. A program consists of declarations of conditional rules and procedures, and exactly one declaration of a main command sequence. The category `RuleId` refers to declarations of conditional rules in `RuleDecl`. Procedures must be non-recursive, they can be seen as macros with local declarations.

```

Prog      ::= Decl {Decl}
Decl      ::= RuleDecl | ProcDecl | MainDecl
ProcDecl  ::= ProcId '=' [ '[' LocalDecl ']' ] ComSeq
LocalDecl ::= (RuleDecl | ProcDecl) {LocalDecl}
MainDecl  ::= Main '=' ComSeq
ComSeq    ::= Com { ';' Com }
Com       ::= RuleSetCall | ProcCall
           | if ComSeq then ComSeq [else ComSeq]
           | try ComSeq [then ComSeq] [else ComSeq]
           | ComSeq '!' | ComSeq or ComSeq
           | '(' ComSeq ')' | break | skip | fail
RuleSetCall ::= RuleId | '{' [RuleId {',' RuleId}] '}'
ProcCall    ::= ProcId

```

Figure 3: Abstract syntax of GP 2 programs.

Next, we describe the effect of GP 2's main control constructs informally.

The call of a rule set  $\{r_1, \dots, r_n\}$  non-deterministically applies one of the rules whose left-hand graph matches a subgraph of the *host graph* such that the dangling condition and the rule's application condition are satisfied. The call *fails* if none of the rules is applicable to the host graph.

The command `if  $C$  then  $P$  else  $Q$`  is executed on a host graph  $G$  by first executing  $C$  on a copy of  $G$ . If this results in a graph,  $P$  is executed on the original graph  $G$ ; otherwise, if  $C$  fails,  $Q$  is executed on  $G$ . The `try` command has a similar effect, except that  $P$  is executed on the result of  $C$ 's execution.

The loop command  $P!$  executes the body  $P$  repeatedly until it fails. When this is the case,  $P!$  terminates with the graph on which the body was entered for the last time. The **break** command inside a loop terminates that loop and transfers control to the command following the loop.

In general, the execution of a program on a host graph may result in different graphs, fail, or diverge. For a formal account of GP 2’s operational semantics, see [CP21].

**2.3. Rooted Programs.** The bottleneck for efficiently implementing a language based on graph transformation rules is the cost of graph matching. In general, to match the left-hand graph  $L$  of a rule within a host graph  $G$  requires time polynomial in the size of  $L$  [BP12, BP16]. As a consequence, linear-time graph algorithms in imperative languages may be slowed down to polynomial time when they are recast as rule-based programs.

To speed up matching, GP 2 supports *rooted* graph transformation where graphs in rules and host graphs are equipped with so-called root nodes. Roots in rules must match roots in the host graph so that matches are restricted to the neighbourhood of the host graph’s roots. We draw root nodes using double circles.

A rule  $\langle L \leftarrow K \rightarrow R, c \rangle$  is *fast* if (1) each node in  $L$  is undirectedly reachable from some root, (2) neither  $L$  nor  $R$  contain repeated occurrences of list, string or atom variables, and (3) the condition  $c$  contains neither an **edge** predicate nor a test  $e_1=e_2$  or  $e_1!\neq e_2$  where both  $e_1$  and  $e_2$  contain a list, string or atom variable.

In [BP12, BP16], it is shown that rooted graph matching can be implemented to run in constant time for fast rules, provided there are upper bounds on the maximal node degree and the number of roots in host graphs. However, in Section 5 we show how to overcome the restriction that node degrees must be bounded.

### 3. THE GP 2-TO-C COMPILER

We now introduce the GP 2-to-C compiler which translates GP 2 programs into C code. A description of the compiler’s implementation is beyond the scope of this paper; the interested reader may consult [Bak15, BP16, CRP20]. We also give a list of assumptions on the performance of some basic operations generated by the compiler. These assumptions are needed to analyse the time complexity of programs.

**3.1. Execution Pipeline.** Figure 4 shows that GP 2 programs are first compiled to C programs using the GP 2-to-C compiler and then transformed by GCC into platform-dependent binaries which are executed by the runtime system. The execution entails reading an input graph from the disc.




Figure 4: GP 2 execution pipeline.

Executing a program results in one of three possible outcomes: the program (1) produces an output graph, (2) evaluates to **fail**, or (3) diverges (does not terminate).

For example, consider the GP 2 program in Figure 5 which computes the transitive closure of a graph.<sup>1</sup> The C code generated for this program (1) reads and parses the input host graph, (2) applies the `link` rule as long as possible, and (3) returns the resulting host graph as output.




Figure 5: GP 2 program `transitive-closure`.

Applying the `link` rule as long as possible requires to find a match satisfying the application condition, and if one is found, adding a fresh edge according to the rule and the match. This is repeated zero or more times until a match cannot be found anymore. Finding a match involves finding a grey node corresponding to node 1, then finding an outgoing edge to another grey node, corresponding to node 2, then finding an outgoing edge to yet another grey node, corresponding to node 3, and then checking the application condition. This search is controlled by the matching algorithm and is described in general in Bak’s thesis [Bak15]. Updates to the algorithm are described in [CCP19]).

We conclude this subsection by stressing that the GP 2-to-C compiler does not attempt to avoid program failure or to generate more than one output graph. This would require backtracking, which is only used in the rule matching algorithm and in the implementation of rule sets (a rule set only fails if none of its rules can be matched). The compiler has been deliberately designed not to backtrack in program execution, which is necessary to achieve the fast run times reported in this paper.

**3.2. Reasoning about Time Complexity.** When analysing the time complexity of programs, we assume that these are fixed. This is customary in algorithm analysis where programs are fixed and runtime is measured in terms of input size [AHU74, Ski20]. In our setting, the input size is the *size* of a host graph, which we define to be the total number of nodes and edges.

When we discuss the time complexity of rule applications, it suffices to reason about the complexity of finding a match because all the programs in this paper satisfy the assumption of the following lemma.

**Lemma 3.1** Constant Time Graph Construction [CCP22]. *If a match has been found for a rule that does not modify labels, other than possibly changing marks or changing non-host-graph labels to host-graph labels, the resulting graph can be constructed in constant time.*

To reason about a program’s time complexity, we need to make assumptions on the complexity of certain elementary operations generated by the compiler. Figure 6 shows the complexity

<sup>1</sup>The concrete syntax of an equivalent program with unmarked nodes can be found at <https://github.com/UoYCS-plasma/GP2/blob/master/programs/transitive-closure.gp2>.

of the basic procedures of the rule matching algorithm, adapted from [CCP22]. The grey rows indicate procedures updated by the changes introduced in Section 5. The runtimes reported in this paper’s case studies and in [AP24, AP25] are consistent with the complexity assumptions of Figure 6.

Operation	Description	Complexity
<code>alreadyMatched</code>	Test if the given item has been matched in the host graph.	$O(1)$
<code>clearMatched</code>	Clear the <code>is matched</code> flag for a given item.	$O(1)$
<code>setMatched</code>	Set the <code>is matched</code> flag for a given item.	$O(1)$
<code>firstHostNode(m)</code>	<b>Fetch the first node of mark <math>m</math> in the host graph.</b>	$O(1)$
<code>nextHostNode(m)</code>	<b>Given a node of mark <math>m</math>, fetch the next node of mark <math>m</math> in the host graph.</b>	$O(1)$
<code>firstHostRootNode</code>	Fetch the first root node in the host graph.	$O(1)$
<code>nextHostRootNode</code>	Given a root node, fetch the next root node in the host graph.	$O(1)$
<code>firstInEdge(m)</code>	<b>Given a node, fetch the first inedge of mark <math>m</math>.</b>	$O(1)$
<code>nextInEdge(m)</code>	<b>Given a node and an edge of mark <math>m</math>, fetch the next inedge of mark <math>m</math>.</b>	$O(1)$
<code>firstOutEdge(m)</code>	<b>Given a node, fetch the first outedge of mark <math>m</math>.</b>	$O(1)$
<code>nextOutEdge(m)</code>	<b>Given a node and an edge of mark <math>m</math>, fetch the next outedge of mark <math>m</math>.</b>	$O(1)$
<code>firstLoop(m)</code>	<b>Given a node, fetch the first loop edge of mark <math>m</math>.</b>	$O(1)$
<code>nextLoop(m)</code>	<b>Given a node and an edge of mark <math>m</math>, fetch the next loop edge of mark <math>m</math>.</b>	$O(1)$
<code>getInDegree</code>	Given a node, fetch its indegree.	$O(1)$
<code>getOutDegree</code>	Given a node, fetch its outdegree.	$O(1)$
<code>getMark</code>	Given a node or edge, fetch its mark.	$O(1)$
<code>isRooted</code>	Given a node, determine if it is rooted.	$O(1)$
<code>getSource</code>	Given an edge, fetch the source node.	$O(1)$
<code>getTarget</code>	Given an edge, fetch the target node.	$O(1)$
<code>parseInputGraph</code>	Parse and load the input graph into memory.	$O(n)$
<code>printHostGraph</code>	Write the current host graph as output.	$O(n)$

Figure 6: Updated complexity assumptions on elementary operations generated by the compiler, where  $n$  is the size of the input. Modified operations are highlighted in grey.

#### 4. THE PROBLEM WITH UNBOUNDED-DEGREE GRAPHS

This section explains the problem caused by classes of input graphs with an unbounded node degree, while the next section introduces the compiler improvements.

Previously, non-destructive GP 2 programs based on depth-first search ran in linear time on graph classes of bounded node degree but in non-linear time on graph classes of unbounded degree [CCP22]. For example, consider the program `is-connected` in Figure 7 which checks whether a graph is connected.<sup>2</sup> Input graphs are arbitrary GP 2 host graphs with grey, non-rooted nodes and unmarked edges. If the graph is connected, the program returns a graph that is isomorphic up to marks. On disconnected graphs, the program fails.

The rule `init` picks an arbitrary grey node as a root (if the input graph is non-empty) and then the loop `DFS!` performs a depth-first search of the connected component of the node chosen by `init`. Rule `forward` marks each newly visited node blue, and `back` unmarks it once it has been processed. The procedure `DFS` ends when `back` fails to match, indicating that the search is complete. Rule `match` checks whether a grey-marked node still exists in the graph following the execution of `DFS!`. This is the case if and only if the input graph contains more than one connected component. In this situation the program invokes the command `fail`, otherwise it terminates by returning the graph resulting from the depth-first search.

<sup>2</sup>A (possibly dashed) line without arrowhead between two nodes is a *bidirectional* edge which matches host graph edges of either direction.




Figure 7: The old program `is-connected`.

It can be shown that the program `is-connected` runs in linear time on classes of graphs with bounded node degree [CCP22]. However, as the following example shows, the program may require non-linear time on unbounded-degree graph classes. Figure 8 shows an execution of `is-connected` on a star graph with 8 edges. Below each graph, we record the range of attempts that the matching algorithm may make. For instance, in the second graph of the top row, either a match is found immediately among the edges that connect the central node with the grey nodes, or the dashed edge is unsuccessfully tried first. In order to find

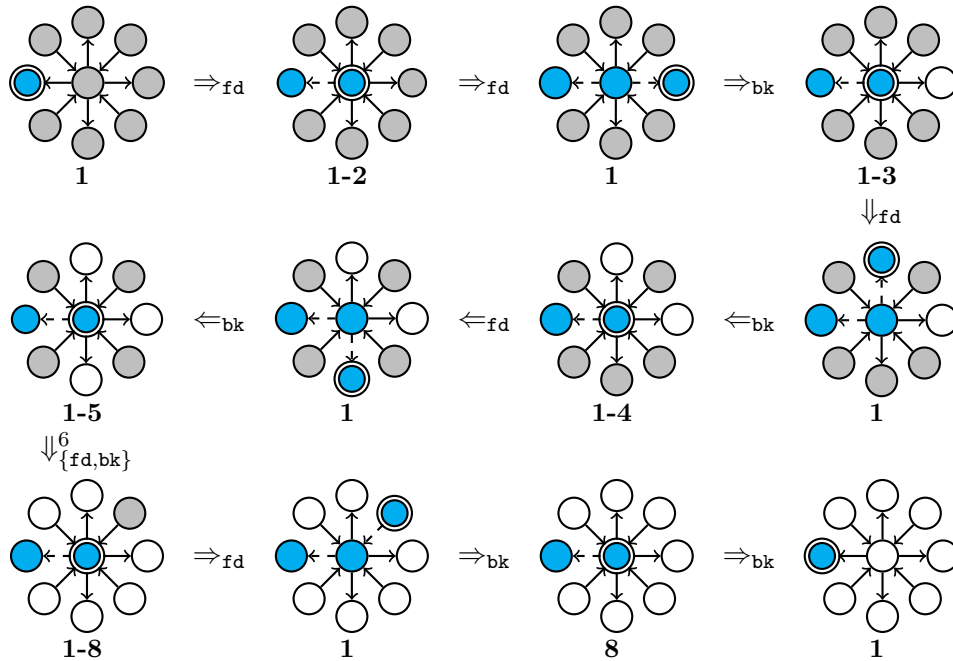


Figure 8: Matching attempts with the `forward` rule. `fd` and `bk` denote `forward` and `back`, respectively.

a match for the rule `forward`, the matching algorithm considers, in the worst case, every edge incident with the root. When the node central to the graph is rooted and the rule

**forward** is called, the matching algorithm may first attempt a match with the dashed back edge and all edges incident with an unmarked node. Therefore, the maximum number of matching attempts for **forward** grows as the root moves back to the central node. As can be seen from this example, the worst-case complexity of matching **forward** throughout the program’s execution is  $2|E| + \sum_{i=1}^{|E|} i = O(|E|^2)$  where  $E$  is the set of edges.

## 5. COMPILER ENHANCEMENTS

To overcome the limitation that depth-first search requires bounded-degree graph classes to run in linear time, we changed the GP 2 compiler described in [CRP20]. We refer to that compiler as the *2020 compiler* and call the version introduced in this paper the *new compiler*<sup>3</sup>.

**5.1. First Enhancement.** The 2020 compiler stored the host graph’s structure as one linked list containing every node in the graph, with each node storing two linked lists of edges: one for incoming edges and one for outgoing edges. Searching through these lists to find an edge with a particular mark could involve looking up a large number of edges with incompatible marks, slowing down the search. In particular, the number of unmatchable edges could grow during program execution if the marks changed.

When trying to match an edge of a given rule, host graph edges with incorrect orientation or mark can be ignored. Thus, by organising the edges incident to a node into linked lists based on their mark and orientation, the matching algorithm can select edges of correct mark and orientation. More precisely, the new compiler generates for each node a two-dimensional array such that each entry is a linked list of all incident edges with a particular combination of mark and orientation. We also consider loops to be a distinct type of orientation. The array, therefore, consists of five rows (unmarked, dashed, red, green, blue) and three columns (incoming, outgoing, loop), totalling 15 cells. See Figure 9 for an illustration. We mention that the program in Figure 7 still does not run in linear time for input graph classes of unbounded degree; see Section 6 for an adjusted program which exploits the aforementioned changes and runs in linear time.

Note that, given a rule with a node  $v$  and an incident edge with a particular mark and orientation, the matching algorithm can access in constant time the corresponding list in  $v$ ’s array. The first element of that list is a host graph edge with matching mark and orientation; in case the list is empty, the algorithm has determined that no matching edge exists in the host graph.

	in	out	loop
unmarked	...	...	...
dashed	...	...	...
red	...	...	...
green	...	...	...
blue	...	...	...

Figure 9: Two-dimensional array of linked lists of edges.

<sup>3</sup>Available at: <https://github.com/UoYCS-plasma/GP2>.

**5.2. Finding Nodes in Constant Time.** The previous subsection implies that the code generated by the new compiler can match edges labelled with variables of type `list` in constant time (or detect that there are no matching edges in the host graph). The corresponding problem of finding nodes labelled with `list` variables in constant time could not be solved with the 2020 compiler either. This is explained below.

Consider the program `is-discrete` in Figure 10 which checks whether a graph is discrete, that is, contains no edges. Given an unmarked input graph, the program fails if and only if the input graph contains any edges. The program repeatedly picks some unmarked node, turns it into a red root node, and checks whether the node is isolated by applying the rule `isolated`. The left-hand node of the rule is not an interface node and hence, by the dangling condition, the rule is applicable if and only if the root node is isolated. If `isolated` is applied, the root node is unrooted and the loop continues. If `isolated` is not applicable, a non-isolated node has been found and the loop breaks. In this case rule `root` finds that a root exists and the `fail` command is invoked.




Figure 10: The program `is-discrete`.

The code generated by the 2020 compiler matched the rule `mark` with a complexity of  $O(n)$ , where  $n$  is the number of nodes in the host graph. This is because finding an unmarked node required iterating through a unique linked list holding all nodes in the host graph. As a result, the overall time complexity of `is-discrete` was  $O(n^2)$ , as empirically evidenced by the timing in Figure 11.

**5.3. Second Enhancement.** As explained above, the 2020 compiler generated code storing all host graph nodes in a single linked list. Hence the matching algorithm had to carry out unnecessary lookups to find a match for a node in the host graph when there existed nodes with marks that did not match. The solution in the new compiler is to store nodes in five separate linked lists: one list for each node mark and one list for unmarked nodes. The linked lists are stored in an array, providing constant-time access to the first element of each list. See Figure 12 for an illustration. This structure allows to match the rule `mark` in constant time as the matching algorithm immediately inspects the linked list holding all unmarked nodes in the host graph. Since the left-hand node of `mark` is labelled with a variable of type `list`, the first node found in the linked list by the matching algorithm matches successfully. As a result of the improved data structure, the time complexity of the program `is-discrete` under the new compiler is  $O(n)$ . Figure 11 highlights the difference in runtimes of the program `is-discrete` run under both compilers.

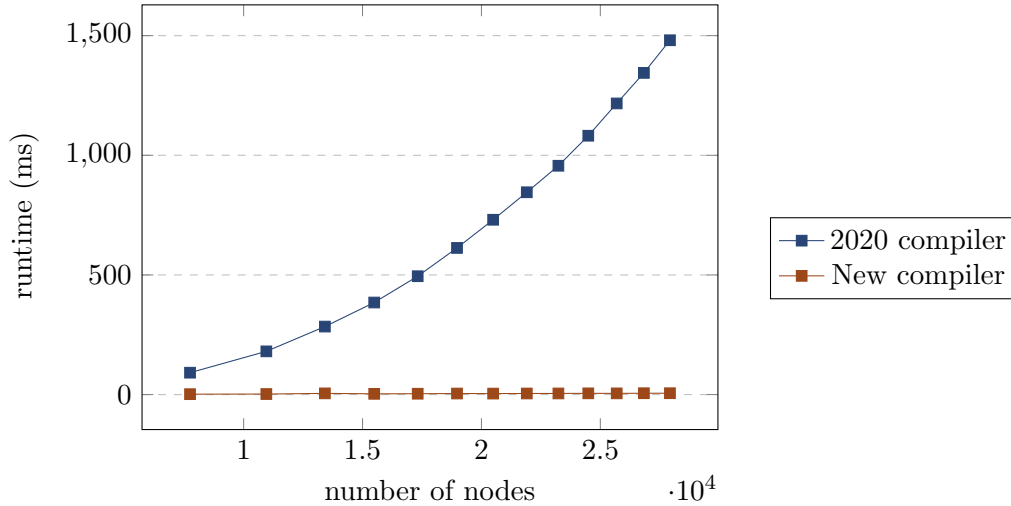


Figure 11: Measured performance of the program `is-discrete` on discrete graphs with the 2020 compiler and the new compiler.

unmarked	...
grey	...
red	...
green	...
blue	...

Figure 12: Array of linked lists of nodes.

## 6. CASE STUDY: RECOGNISING CONNECTED GRAPHS

Checking for the connectedness of a graph is a fundamental application of depth-first search (DFS). A GP 2 program solving this problem is given in [Bak15, CCP22], but to run in linear time it requires input graphs of bounded node degree.

**6.1. Program.** The program `is-connected`<sup>4</sup> in Figure 13 recognises connected graphs, in that each pair of nodes are connected by a sequence of edges of arbitrary orientation, with respect to the following specification.

**Input:** An arbitrary GP 2 host graph such that

- (1) each node is non-rooted and marked grey, and
- (2) each edge is unmarked.

**Output:** If the input graph is connected, a host graph that is isomorphic to the input graph up to marks. Otherwise *failure*.

This can be achieved by conducting a DFS from an arbitrarily chosen node while marking newly visited nodes. Since the DFS cannot propagate beyond the connected component it started in, the presence of an unmarked node following the DFS traversal indicates that the host graph is not connected. This program improves implementations of the same problem

<sup>4</sup>The concrete syntax of the program available at: <https://gist.github.com/ismaili-ziad/7b419a12a8d156abf5ad8ac8c352713a>.

in [Bak15, CCP22], which is discussed in Section 4, by achieving linear-time complexity on arbitrary input graphs (including star graphs).

We stress that the old `is-connected` program (Figure 7) still does not run in linear time for input graph classes of unbounded degree. This means that the new compiler does not necessarily improve the time complexity of existing programs. Instead, we need a new programming technique to exploit the improved graph data structure. To explain this technique, consider the program in Figure 13. The rule `next_edge` selects in constant time an adjacent node that shares an unmarked edge with the root. Subsequently, the ruleset `{move, ignore}` acts like a case statement, distinguishing the cases where the adjacent node is blue or grey. We use a similar technique for the case studies of Section 7 and Section 8.

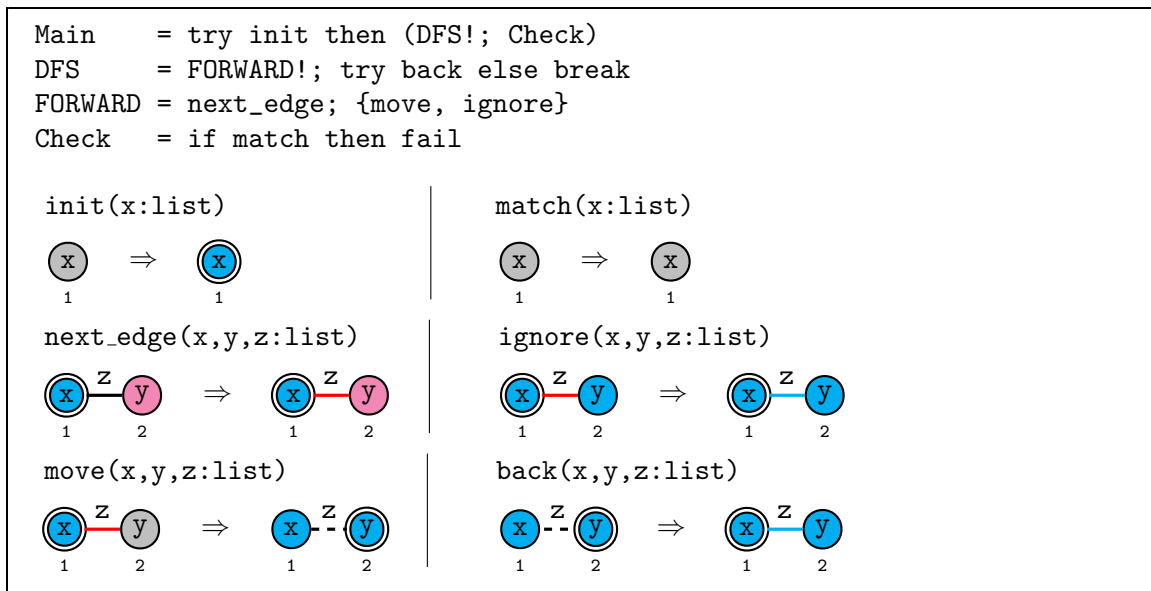


Figure 13: The (new) program `is-connected`.

**6.2. Proof of Correctness.** Let us examine the correctness of `is-connected`. We write  $|X|$  for the cardinality of a finite set  $X$ . We use the notation  $A \Rightarrow_r B$  to indicate that  $B$  results from applying  $r$  to  $A$ .

First of all, we establish what we mean by an input graph in the context of the `is-connected` program.

**Definition 6.1** (Connectedness). A graph  $G$  is *connected* if and only if, for any pair of distinct vertices  $u$  and  $v$  in  $G$ , there exists a path of edges (regardless of direction) from  $u$  to  $v$ .

From now on, we call graphs satisfying the assumptions of the specification given at the beginning of the subsection as *input graphs*.

The choice of having input nodes grey is motivated by the fact that nodes with the wildcard colour magenta (corresponding to `any`) cannot match unmarked nodes. We overcome this problem by requiring that all nodes in input graphs are marked grey.

**Proposition 6.2.** *Throughout the execution of `is-connected` on an input graph all non-grey nodes in the host graph share a connected component.*

*Proof.* All nodes in the input graph are initially grey. Upon inspection of the rules, `init` and `move` are the only rules altering the mark of a grey node, where `init` is only called once at the start of the program. Let us show inductively that the invariant holds.

If `init` is applied at the start of the program, the host graph will contain exactly one non-grey node, which does not violate the invariant. If `init` is not applied, `move` is never called, and the invariant remains true.

Now, assume that the invariant is true at some stage of the execution. If `move` is applied, the non-grey node  $v$  that turns blue has to be adjacent to a non-grey node. Thus, if all non-grey nodes shared a connected component prior to the application of `move`,  $v$  has to be part of that component as well. Thus, the invariant remains satisfied.  $\square$

**Lemma 6.3** (Termination of `is-connected`). *The program `is-connected` terminates on any input graph.*

*Proof.* The program `is-connected` contains two looping bodies: `DFS!` and `FORWARD!`. We first show termination of `FORWARD!`. Consider the measure  $\#_1 X$  consisting of the number of unmarked edges in the host graph  $X$ . The rule `next_edge` is invoked at the beginning of `FORWARD`, and if it fails to match, the loop breaks. Clearly, an application of `next_edge` reduces the measure  $\#_1$  because the edge in the left-hand side is unmarked, and is marked red in the right-hand side. Since no rule in `FORWARD` creates or unmarks an edge, `FORWARD!` terminates.

We now show that the loop `DFS!` terminates. This time, define  $\#_2 X$  to be the number of non-blue edges in the host graph  $X$ . Suppose that  $G \Rightarrow_{\text{FORWARD!}} H$ . Since no rule in `FORWARD!` increases the number of non-blue edges, we have  $\#_2 G \geq \#_2 H$ .

Let now  $H \Rightarrow_{\text{back}} M$ . Clearly,  $\#_2 H > \#_2 M$  since the rule marks a dashed edge blue. As we know that `FORWARD!` is terminating, it follows that `DFS!` is also terminating.  $\square$

**Proposition 6.4.** *Throughout the execution of `DFS!`, (1) there is exactly one root node in the host graph, and (2) all dashed edges in the host graph form a directed path such that every node on the path is blue, and the endpoint of the path is the root node.*

*Proof.* (1) Initially, every node of the host graph is non-rooted. Assuming that the input graph is non-empty, the main program applies the rule `init` to root exactly one node. Then, observe that every rule in the procedure `DFS` preserves the number of roots in the host graph. Therefore, (1) follows.

(2) Upon inspection of the rules, `move` is the only rule dashing edges in the program. We proceed by induction on the execution steps of the program.

Initially, before `move` is applied, no dashed edges exist in the host graph. The rule `init` creates a non-grey root node, and `back` cannot be applied. Thus, the invariant holds vacuously. Now, assume the invariant holds at some stage of execution. Clearly, the rule `move` extends the dashed path by adding one edge, and `back` reduces the dashed path by removing one edge. Both operations preserve the invariant and ensure the dashed path remains unique, its nodes remain blue, and the endpoint remains the root. Therefore, the invariant is maintained throughout the execution.  $\square$

**Proposition 6.5.** *Let  $G$  be a non-empty input graph to which `init` is applied. Then, the graph resulting from `DFS!` contains a connected component in which every node is non-grey.*

*Proof.* The proposition is trivially true for an input graph consisting of a single node: `init` marks the node blue, and no rule marks a non-grey node grey. Assume now that the input graph contains at least two nodes. By Proposition 6.2, all non-grey nodes belong to the same connected component.

*Claim.* Each connected component containing at least one non-grey node does not contain any grey node.

*Proof of Claim.* For the sake of contradiction, suppose the connected component containing all non-grey nodes in the output graph also includes at least one grey node. Let  $v$  and  $u$  be such that  $v$  is grey,  $u$  is non-grey, and  $v$  is adjacent to  $u$ , which must occur if they share a connected component.

Upon inspection of the rules, `init` and `move` are the only rules altering the mark of a grey node; both mark a grey node blue and root it. Thus,  $v$  must have been rooted at some point during the execution. Let us consider two cases based on whether  $v$  remains rooted in the output graph.

*Case 1:*  $v$  is rooted in the output graph. The procedure `DFS!` terminates only if `back` fails to apply, which occurs when `next_edge` fails to find an unmarked edge. However,  $u$ , being non-grey, shares an unmarked edge with  $v$ . Since  $u$  is blue and  $v$  is rooted, the rules `next_edge` and `move` must have applied, which contradicts the assumption that  $v$  is grey in the output graph.

*Case 2:*  $v$  is not rooted in the output graph. Throughout the execution of `DFS!`, by Proposition 6.4(1), there is exactly one root node at any time. Since  $v$  is not rooted, it must have been unrooted by either `move`, which dashed an edge incident with  $v$ , or `back`. Two scenarios arise:

- (1) If `back` was applied to  $v$ , it would have unrooted  $v$  only after no rules in `FORWARD` were applicable. However, since  $v$  is adjacent to  $u$ , and  $u$  shares an unmarked edge, `next_edge` and `move` would have applied, turning  $u$  non-grey. This contradicts the assumption.
- (2) If `back` was not applied to  $v$ ,  $v$  must be part of a path of dashed edges with a root node as the endpoint (by Proposition 6.4). The termination of `DFS!` implies that the rule `back` was no longer applicable, which is, again, a contradiction to the fact that there exists some root node adjacent to a blue node sharing a dashed edge (which follows from the fact that  $v$  must be incident with at least one dashed edge).

Given that both cases contradict the assumption, at the termination of `try init then DFS!` on a non-empty graph, there must be a connected component in which every node is non-grey. □

**Theorem 6.6** (Correctness of `is-connected`). *The program `is-connected` is totally correct<sup>5</sup> with respect to the specification at the beginning of Subsection 6.1.*

*Proof.* Observe that each rule can only change marks or rootedness, and therefore the resulting graph is isomorphic to the input graph up to marks.

Termination follows from Lemma 6.3. Let  $G$  be the input graph. First, assume that the input graph  $G$  is empty. Since the rule `init` is not applicable, the program immediately terminates and no `fail` command is invoked.

<sup>5</sup>A program is *totally correct* if it terminates on all input graphs and produces an output graph satisfying the program's specification.

Now, assume  $G$  contains at least one node. Consider the graph  $H$  as the output graph of `try init then DFS!` on  $G$ , i.e.  $G \Rightarrow_{\text{init}} I \Rightarrow_{\text{DFS!}} H$  for some intermediate graph  $I$ . By Proposition 6.5, there exists some connected component in  $H$  where every node is non-grey, let us call it  $C$ , and by Proposition 6.2, it follows that  $H - C$  contains no non-grey node. Therefore,  $G$  is not connected if and only if  $H - C$  is non-empty. The procedure `Check` is executed after `DFS!`, and verifies that no grey node, which must belong to  $H - C$  should one exist, exists. If a grey node exists,  $H - C$  is non-empty and the program fails, indicating that  $G$  is not connected. If none exists, the program terminates without invoking the `fail` command, indicating that  $G$  is connected.  $\square$

### 6.3. Proof of Complexity.

**Proposition 6.7.** *Throughout the execution of `is-connected` on an input graph, there is at most one red edge in the host graph.*

*Proof.* Upon inspection of the rules, `next_edge` is the only rule that marks an edge red in the host graph (note that the program is structure-preserving). It is easy to see that when `next_edge` is applied, either `move` or `ignore` has to apply, as node 2 of `next_edge` can either be blue or grey. Therefore, it follows that the number of red edges in the host graph is bounded to 1.  $\square$

**Theorem 6.8** (Complexity of `is-connected`). *The program `is-connected` terminates in time linear in the size of the input graph.*

*Proof.* We now prove that the program `is-connected` terminates in time linear in the size of the input graph. The argument proceeds by bounding both the number of matching attempts and the number of rule applications. More specifically, we show (1) that for every rule, there is at most one matching attempt; then, (2) we derive an expression for the number of calls during the execution of the program for each rule; finally, (3) we combine those results to derive the time complexity of `is-connected`.

(1) We consider the matching behaviour of each rule. The rule `init` attempts to match a single grey node. If the input graph is non-empty, this match succeeds immediately on the first node; otherwise, it fails without making further attempts. Hence, there is at most one matching attempt for `init`. The rule `match` also attempts to find a grey node using the compiler procedure `firstHostNode(m)`, which operates in constant time. As a result, `match` also makes only one matching attempt.

Since there is at most one root in the host graph (`init` is the only rule creating or deleting a root), node 1 of `next_edge`, `move` and `ignore`, and node 2 of `back match` in constant time. The rule `next_edge` matches any unmarked edge incident to the root and incident to an any-marked node. Since any node adjacent to the root is marked, and non-loop edges are stored in distinct lists with respect to their marks, there is at most one matching attempt for `next_edge`. An analogous argument can be made for all rules beside `init` and `match`, as it is known from Proposition 6.7 that there is at most one red edge in the host graph, and Proposition 6.4 establishes that a root node is incident to at most one dashed edge, hence limiting the number of possible matches to a single edge. The rule `match` requires the matching algorithm to find a grey node, which the procedure `firstHostNode(m)` of the updated compiler (Figure 6) executes in constant time.

- (2) We now analyse the number of times each rule is called during execution. For the purpose of this proof, let  $n$  and  $m$  be the number of nodes and edges, respectively.

The rule `init` is invoked once at the beginning of execution and either succeeds (if the graph is non-empty) or fails. The rule `match` is also called exactly once in the procedure `Check`.

The rule `back` marks an edge blue and is called within the depth-first search loop. Since blue edges retain their mark, each edge can only be marked blue once. Thus, `back` may be successfully applied at most  $m$  times. The loop containing `back` terminates when it fails to apply, so there is at most one unsuccessful application, giving us a total of at most  $m + 1$  calls.

The rules `move` and `ignore` are called only after a successful application of `next_edge`. The rule `move` marks grey nodes blue, and since non-grey nodes do not revert to grey, at most  $n - 1$  applications are needed (i.e. one for each node except the one initially marked by `init`). The rule `ignore`, like `back`, marks edges blue, and thus can be applied at most  $m$  times.

The rule `next_edge` selects and marks an unmarked edge. Since edges cannot revert to being unmarked, it can be successfully applied at most  $m$  times. Unsuccessful applications occur when no matching edge remains, which leads to a call to `back`. As previously noted, there are at most  $m + 1$  such calls, resulting in at most  $m$  unsuccessful applications of `next_edge`. Therefore, the total number of calls to `next_edge` is bounded by  $2m \in O(m)$ .

- (3) Since each rule application and `match` is performed in constant time, and the total number of applications across all rules is  $O(n + m)$ , it follows that the overall time complexity of `is-connected` is linear in the size of the input graph (i.e. the number of nodes and edges).  $\square$

Figure 21 shows the empirical timings of `is-connected` on the graph classes of Figure 14 to Figure 20. The measured runtimes do not account for graph parsing, building and printing, as these operations have a linear-time complexity with respect to the input size (Figure 6). Compilation time is also not included. Note that on discrete graphs, the program runs in constant time because no rule in `DFS` is applicable, and both `init` and `match` are applied only once.




Figure 14: Square grid. Figure 15: Binary tree. Figure 16: Star graph. Figure 17: Cycle graph.




Figure 18: Complete graph.

Figure 19: Linked list.

Figure 20: Discrete graph.




Figure 21: Measured performance of the program `is-connected`.

## 7. CASE STUDY: RECOGNISING ACYCLIC GRAPHS

Checking whether a given graph contains a directed cycle is a basic problem in the area of graph algorithms [Ski20]. A GP 2 program solving this problem is given in [CCP22], but to run in linear time it requires input graphs of bounded node degree. The same paper contains a program for the related problem of recognising binary DAGs, which are acyclic graphs in which each node has at most two outgoing edges. This program has a linear runtime on arbitrary input graphs but is destructive in that the input graph is partially or totally deleted.

**7.1. Program.** The program `is-dag`<sup>6</sup> in Figure 22 recognises acyclic graphs with respect to the following specification.

**Input:** An arbitrary GP 2 host graph such that

- (1) each node is non-rooted and marked grey, and
- (2) each edge is unmarked.

**Output:** If the input graph is acyclic, a host graph that is isomorphic to the input graph up to marks. Otherwise *failure*.

Figure 23 illustrates an execution of `is-dag` on a cyclic input graph, and Figure 24 on an acyclic graph. The program implements a directed DFS (depth-first search) of the host graph that marks the visited nodes red or blue, where the red nodes are currently being visited. All the red nodes are connected by a directed path of dashed edges, where the last node is rooted.

It is an invariant of `is-dag` that there is at most one root in the host graph throughout the program’s execution (see Proposition 7.2). The graph contains a cycle if and only if the search finds an edge from the root to a red node (see Proposition 7.4).

Consider the loop `(init; DFS!; try unroot else break)!` of `is-dag`’s main procedure. Rule `init` selects an arbitrary grey node as a root to start a directed DFS. The loop

<sup>6</sup>The concrete syntax of the program available at: <https://gist.github.com/ismaili-ziad/959f26e188b210821d08eb5ed2965404>.

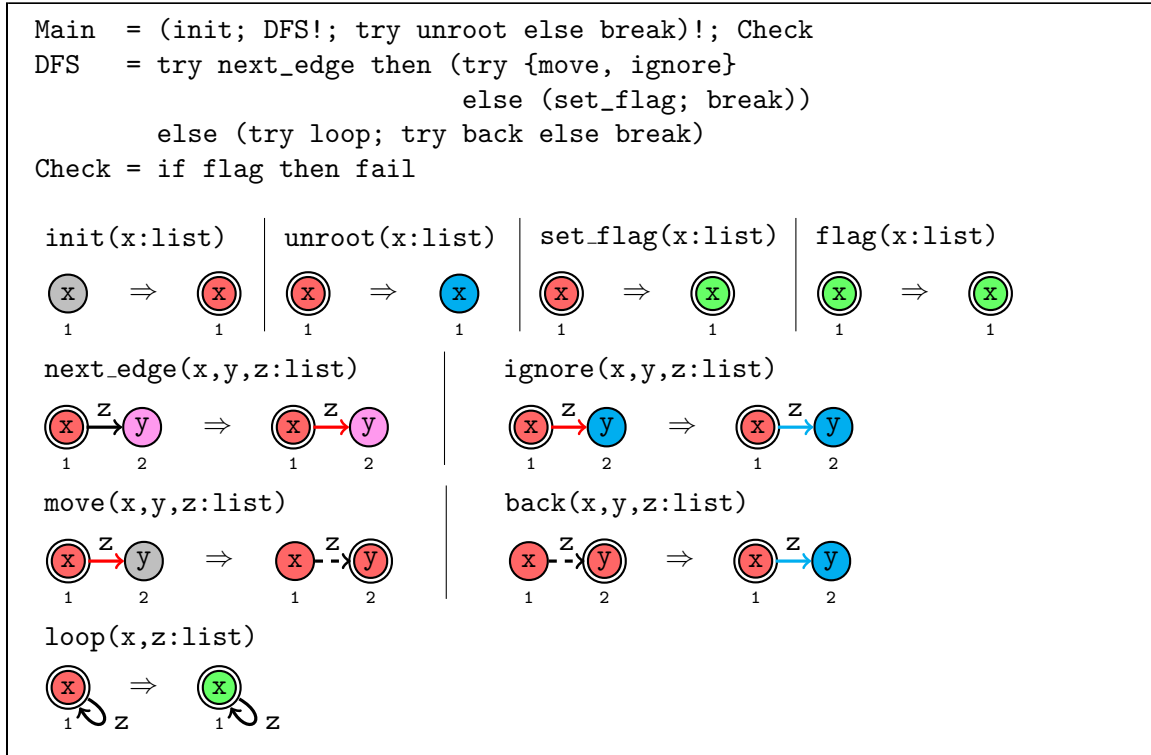


Figure 22: The program is-dag.

DFS! moves the root in depth-first fashion through the host graph. The procedure uses a `try-else` command to find any unprocessed (that is, unmarked) edge outgoing from the root. It does this by calling `next_edge`. If there is such an edge, the rule marks it red so that it can be uniquely identified by the rest of the procedure. If no such edge exists, the root can no longer move forward and the `else` statement is invoked instead.

After a successful application of `next_edge`, the root is adjacent to either (1) a grey node, (2) a blue node, or (3) a red node. In case (1), the rule `move` moves the root to the grey node, marks it red and dashes the traversed edge. Dashed edges represent the path followed by the directed DFS. In case (2), the red edge is marked blue by the rule `ignore` so that it can no longer be matched by `next_edge`. In case (3), neither `move` nor `ignore` is applicable so that `set_flag` marks the root green, indicating the existence of a cycle.

If `next_edge` is not applicable, the command sequence `(try loop; try back else break)` is executed. Rule `loop` checks whether there is a loop attached to the root. If this is the case, the rule marks the root green, similar to `set_flag`. Then rule `back` is tried which implements the *pop* operation in the above mentioned stack model. The rule moves the root backwards along an incoming dashed edge. If no incoming dashed edge is present, the root must be the only element on the stack so that the `break` command terminates the loop DFS!.

Upon termination of DFS!, the rule `unroot` attempts to turn the root into an unrooted blue node. If this is not possible, the root must have been marked green by `set_flag` or `loop`. This implies the existence of a cycle and hence the outer loop of `is-dag` is terminated.

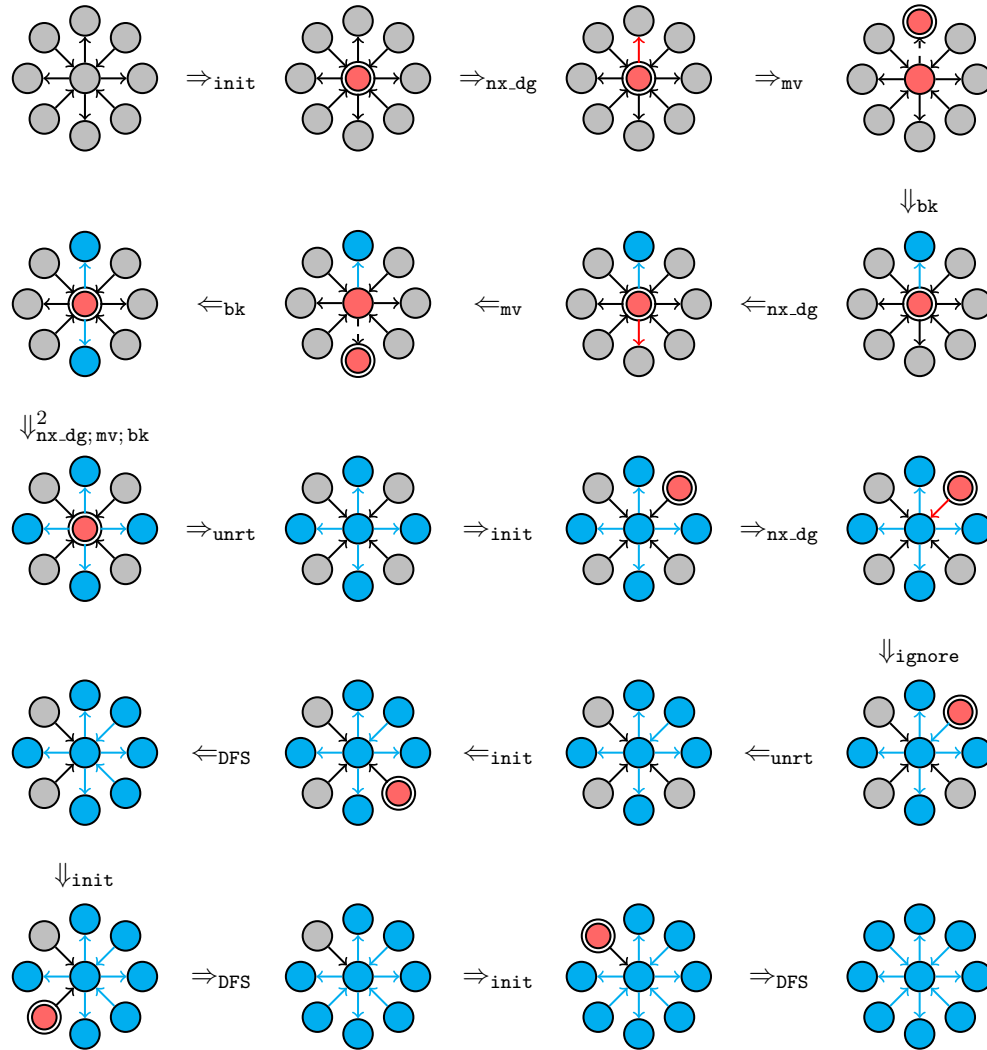



Figure 23: Sample execution of `is-dag` on an acyclic graph (`nx.dg`, `mv`, `bk` and `unrt` correspond to `next_edge`, `move`, `back` and `unroot`, respectively).

If rule `unroot` could be applied, there may still be nodes that have not been visited by the DFS. These are nodes that are not directly reachable from the initial nodes chosen so far. In this case the execution of the outer loop is continued until `init` is no longer applicable or `unroot` fails.

## 7.2. Proof of Correctness.

**Definition 7.1** (Input Graph). An *input graph*, in the context of the acyclicity problem, is an arbitrary GP 2 host graph such that:

- (1) every node is marked grey and non-rooted, and
- (2) every edge is unmarked.

Figure 24: Sample execution of `is-dag` on a cyclic graph.

**Proposition 7.2.** *Throughout the execution of `is-dag` on an input graph, there is at most one root node in the host graph.*

*Proof.* Upon inspection of the rules, `init` and `unroot` are the only rules adding or deleting root nodes in the host graph, both invoked in the looping sequence (`init`; `DFS!`; `try unroot else break`)!.

The input graph initially contains no root nodes according to the specification at the beginning of this section. On the first invocation of the parenthetical looping sequence in `Main`, `init` adds one root node to the host graph. The looping procedure `DFS!` does not add nor delete root nodes. The rule `unroot` deletes the node added by `init` should it apply, or exits the loop.  $\square$

**Proposition 7.3.** *Let  $G$  be a non-empty graph where all nodes are either grey or blue, at least one node is grey, all edges outgoing from blue nodes are marked blue, and all other edges are unmarked. Let  $v$  be the grey node marked blue and rooted by the rule `init` on  $G$ . Throughout the execution of `DFS!` on  $G$ , all dashed edges in the host graph form a directed path such that all red nodes in the graph are in the path,  $v$  is the starting node, a root node (marked either red or green) is the endpoint, and all nodes on the path, except possibly the last node, are red.*

*Proof.* The invariant holds in  $G$ , i.e. before `DFS!` is invoked; here, the host graph contains exactly one red (rooted) node and no dashed edges. Now, assume the invariant holds at some stage of execution. Clearly, the rule `move` extends the dashed path by adding one dashed edge and one red node, and `back` reduces the dashed path by removing one edge. These two rules are the only ones creating dashed edges in the host graph, and no other rule deletes or adds root nodes (Proposition 7.2). Furthermore, `set_flag` and `loop` are the only rules altering the mark of the root node, and mark it green. Therefore, the invariant is maintained throughout the execution.  $\square$

**Proposition 7.4.** *Let  $G$  be a host graph and  $v$  be a node as specified in Proposition 7.3. Upon the execution of `DFS!` on  $G$ , either `set_flag` or `loop` is applied if and only if there exists some directed cycle all nodes of which are reachable from  $v$ .*

*Proof.* As established in Proposition 7.3, there exists exactly one path of red nodes starting from  $v$ , where the endpoint is marked either green or red. Initially, this path consists of exactly one root node marked red,  $v$ .

Observe that the rule `move` moves the root to a node directly reachable from the current root while leaving the previously rooted node marked red. The only rule that turns a red node blue is `back`, which corresponds to the `pop` operation in the DFS. Once a node is turned blue, it will no longer be revisited (this can be verified by inspecting the marks of the rules). The nodes that remain red are precisely those ancestral to the current rooted endpoint in the DFS.

If the rules `move` or `ignore` fail to apply, it implies that the second node (`node 2`) in these rules is already marked red. This condition indicates the presence of an edge from the current endpoint to a node already on the path, hence the existence of a directed cycle. Consequently, the rule `set_flag` is applied, and the looping procedure `DFS!` terminates.

Now, assume that some node  $u$ , reachable from  $v$ , has a looping edge. In this case, the rule `loop` applies before the invocation of `back`. The rule `loop` is called when the endpoint of the current path is  $u$ , which will eventually occur since all nodes reachable from  $v$  will be rooted during the DFS traversal, provided there is no edge-cycle (i.e. a cycle that is not a looping edge).  $\square$

**Lemma 7.5** (Termination of `is-dag`). *The program `is-dag` terminates on any input graph.*

*Proof.* Consider the following weight assignment to edge marks: `unmarked`  $\mapsto$  3, `red`  $\mapsto$  2, `dashed`  $\mapsto$  1, and `blue`  $\mapsto$  0.

Given a host graph  $X$ , define  $\#_1 X$  to be the sum of all the edge weights, as defined above. Then, `DFS!` terminates because its loop body reduces  $\#_1 X$ : rules `next_edge`, `move`, `ignore`, and `back` all decrease the total weight, while no other rule increases it.

Now, define  $\#_2 X$  to be the number of grey nodes in  $X$ . Then, `init` decreases  $\#_2 X$  while `DFS!` and `unroot` do not increase the measure.

Hence, `is-dag` terminates.  $\square$

**Theorem 7.6** (Correctness of `is-dag`). *The program `is-dag` is totally correct with respect to the specification at the beginning of Subsection 7.1.*

*Proof.* Termination follows from Lemma 7.5. Let  $G$  be the input graph. If  $G$  is empty, `init` fails to match, `flag` fails to match, and the program terminates, outputting the empty graph and satisfying the specification. Suppose now that  $G$  consists of at least one node. We then split the remainder of this proof into two cases and define the parenthetical sequence (`init`; `DFS!`; `try unroot else break`) as  $R$ .

*Case 1:*  $G$  is acyclic. Observe that  $R$  terminates under two conditions: (1) `init` fails to apply, or (2) `unroot` fails to apply.

The rule `init` fails to apply if and only if no grey node exists in the host graph, which indicates that all nodes have been visited in the DFS implemented by `DFS!`. By Proposition 7.4, the inapplicability of `init` implies that no directed cycle was found during any execution of `DFS!`, as `init` cannot be invoked after `unroot` fails to match. Consequently,  $R$  terminates, and the rule `flag` in `Check` fails to match. Furthermore, the rule `unroot` can never fail to apply in this context because the only rules capable of turning the root node green in `DFS!` are `set_flag` and `loop`. Both of these rules apply if and only if  $G$  contains a cycle (Proposition 7.4), which it does not. Therefore, `unroot` will always succeed, leaving no green root for `flag` to apply successfully.

*Case 2:*  $G$  is cyclic. Let  $v$  be some node to which `init` is applied.

*Subcase 2.1:* A cycle is reachable from  $v$ . Then either the rule `set_flag` or `loop` will eventually apply during the subsequent execution of `DFS!` (Proposition 7.4). Once those rules are applied, `break` is immediately invoked, the rule `unroot` is not applicable, and therefore,  $R$  terminates. The rule `flag` will apply, causing the program to terminate with the execution of the `fail` command.

*Subcase 2.2:* A cycle is not reachable from  $v$ . Then a cycle must exist within the set of nodes that are unreachable from  $v$ .

*Claim.* All the nodes on the cycle are grey.

*Proof of Claim.* Suppose by contradiction that there is at least one non-grey node on the cycle. Then, by Proposition 7.4, at least one node that can reach this cycle had `init` applied to it. However, by Proposition 7.4, this means that the program `is-dag` would have terminated with failure prior to the application of `init` on  $v$ , a contradiction.  $\square$

By a straightforward induction on the number of remaining grey nodes, it follows that a node that can reach a cycle will eventually have `init` applied to it. By Proposition 7.4, this implies that the program will eventually terminate by invocation of a `fail` command.  $\square$

**7.3. Proof of Complexity.** Let us examine the time complexity of `is-dag`.

**Proposition 7.7.** *Throughout the execution of `is-dag`, there is at most one red edge in the host graph.*

*Proof.* Upon inspection of the rules, `next_edge` is the only rule that marks an edge red in the host graph (again, note that the program is structure-preserving). It is easy to see that when `next_edge` is applied, either `move` or `ignore` has to apply, or both fail to apply, triggering the looping sequence to terminate. Therefore, it follows that the number of red edges in the host graph is bounded to 1.  $\square$

**Theorem 7.8** (Complexity of `is-dag`). *The program `is-dag` terminates in time linear in the size of the input graph.*

*Proof.* We first show that, for every rule, there is at most one matching attempt with respect to the complexity assumptions of the updated compiler (Figure 6).

An analogous argument to the proof of Theorem 6.8 can be made with respect to the rules `init`, `next_edge` (Proposition 7.7), `ignore`, `move`, and `back`. Since there is at most one root node in the host graph, the rules `flag`, `loop` and `set_flag` only require at most one matching attempt.

Similarly, analogous arguments can be made to prove that the rules `next_edge`, `ignore` and `move` are called at most a number of times linear to the size of the graph. It is easy to see that the rule `init` can be called at most  $n$  times, where  $n$  is the number of nodes in the input graph (which does not vary). The rule `set_flag` is only called once as its application results in the termination of `(init; DFS!; try unroot else break)!`. The rule `loop` precedes the invocation of `back` and is therefore called as many times as the latter, thus a number of times linear to the size of the input graph. Finally, the rule `flag` is only called once in `Check`. Therefore, the program `is-dag` runs in time linear to the size of the input graph.  $\square$




Figure 25: Measured performance of the program `is-dag`.

Empirical runtimes are shown in Figure 25 for various graph classes of bounded degree (Figures 14, 15, 17, 19 and 20) and unbounded degree (Figure 16).

## 8. CASE STUDY: THE BELLMAN-FORD ALGORITHM

In this section, we present the first implementation of a single-source shortest-path graph algorithm in GP 2 with a time complexity comparable to that of conventional imperative programming languages:  $O(nm)$ , where  $n$  and  $m$  denote the numbers of nodes and edges, respectively. It computes the shortest distances from one node to all other nodes in the graph. The Bellman-Ford algorithm has been chosen for its simplistic elegance and its detection of negative-weight cycles, in contrast to other popular alternatives, such as Dijkstra’s shortest-path algorithm, which only accepts graphs without negative-weight cycles.

We will first examine a pseudo-code representation of the Bellman-Ford algorithm.

Algorithm 1, adapted from [Sed02], provides a high-level overview of the classical Bellman-Ford algorithm. The chief idea is to iteratively decrease the shortest distances  $\text{dist}[\cdot]$  in the graph until the threshold of  $n - 1$  iterations is reached, where  $n$  is the number of nodes in the graph. The following gives a descriptive account of each section of the code.

**Lines 1 to 3.** The algorithm begins by initialising an array  $\text{dist}[\cdot]$ .  $\text{dist}[v]$  defines the lowest possible accumulated cost to go to  $v$  from the source node  $s$  in  $G$ . The upper bound distance for each node is initialised to  $\infty$ . Naturally, the distance from the source  $s$  to itself is 0.

**Lines 4 to 8.** This stage is called the *relaxation phase*. At each iteration, the algorithm iterates through all the edges in  $E$  and updates their upper bound whenever it finds a shorter distance. Let us consider the edge  $e = (u, v)$ . At the  $i^{\text{th}}$  iteration ( $1 \leq i \leq n - 1$ ), the algorithm checks whether  $\text{dist}[v]$  from the  $(i - 1)^{\text{th}}$  iteration (we can consider the initialisation phase to be the  $0^{\text{th}}$  iteration) can be further improved if there exists a shorter path to  $v$  via  $u$  with a weighted cost of  $\text{dist}[u] + w(e)$ . That process is repeated  $n - 1$  times.

---

**Algorithm 1** Classic Bellman-Ford Algorithm
 

---

```

1: Input: A weighted digraph  $G = (V, E)$  of  $n$  nodes and a source node  $s \in V$ 
2: Set  $\text{dist}[v] = \infty$  for all  $v$  in  $V$ 
3: Set  $\text{dist}[s] = 0$ 
4: for  $n - 1$  iterations do
5:   for each edge  $e$  in  $E$  do
6:      $\text{dist}[t(e)] = \min(\text{dist}[s(e)] + w(e), \text{dist}[t(e)])$ 
7:   end for
8: end for
9: for each edge  $e$  in  $E$  do
10:  if  $\text{dist}[s(e)] + w(e) < \text{dist}[t(e)]$  then
11:    Return FALSE
12:  end if
13: end for
14: Return  $\text{dist}[\cdot]$ 

```

---

**Lines 9 to 13.** This section of the code detects negative cycles. The previous phase guarantees the computation of all shortest paths from  $s$  to all other nodes  $v$  in  $G$ , given that no negative cycle exists. If a negative cycle exists, it follows that there always exists a more optimal solution. The algorithm iterates through all edges and verifies whether a relaxation can still decrease at least one upper bound (which should have all been previously minimised). Should that be possible, this implies that a negative cycle exists and hence triggers the code to return FALSE.

**Line 14.** If the algorithm did not terminate during the negative cycle detection, it then successfully returns TRUE and exits.

**8.1. Program.** The program `bellman-ford`<sup>7</sup> (Figure 26) expects an input graph whose edges are labelled with integers. It calculates the shortest distances from a unique root node to all nodes reachable via directed paths. For each such node, the program appends the shortest distance from the root (an integer value) to the node's label. The exact specification is as follows:

**Input:** An arbitrary GP 2 host graph such that

- (1) every node is marked grey,
- (2) there is exactly one root node,
- (3) every edge is unmarked and labelled with an integer, and
- (4) there are no looping edges.<sup>8</sup>

**Output:** If the input graph contains a directed cycle whose sum of edge labels is negative and there is a directed path from the root to that cycle, *failure*.

Otherwise, a graph obtained from the input graph as follows:

- (1) for each node reachable from the root, append its shortest distance to the label,
- (2) append the string `f` to the label of all other nodes, and
- (3) mark all edges blue.

---

<sup>7</sup>The concrete syntax of the program is available at: <https://gist.github.com/ismaili-ziad/adc54b05aaff9ef6b3c57e7ea8a68543>.

<sup>8</sup>We make this assumption for simplicity. It is straightforward to extend the program to handle looping edges.

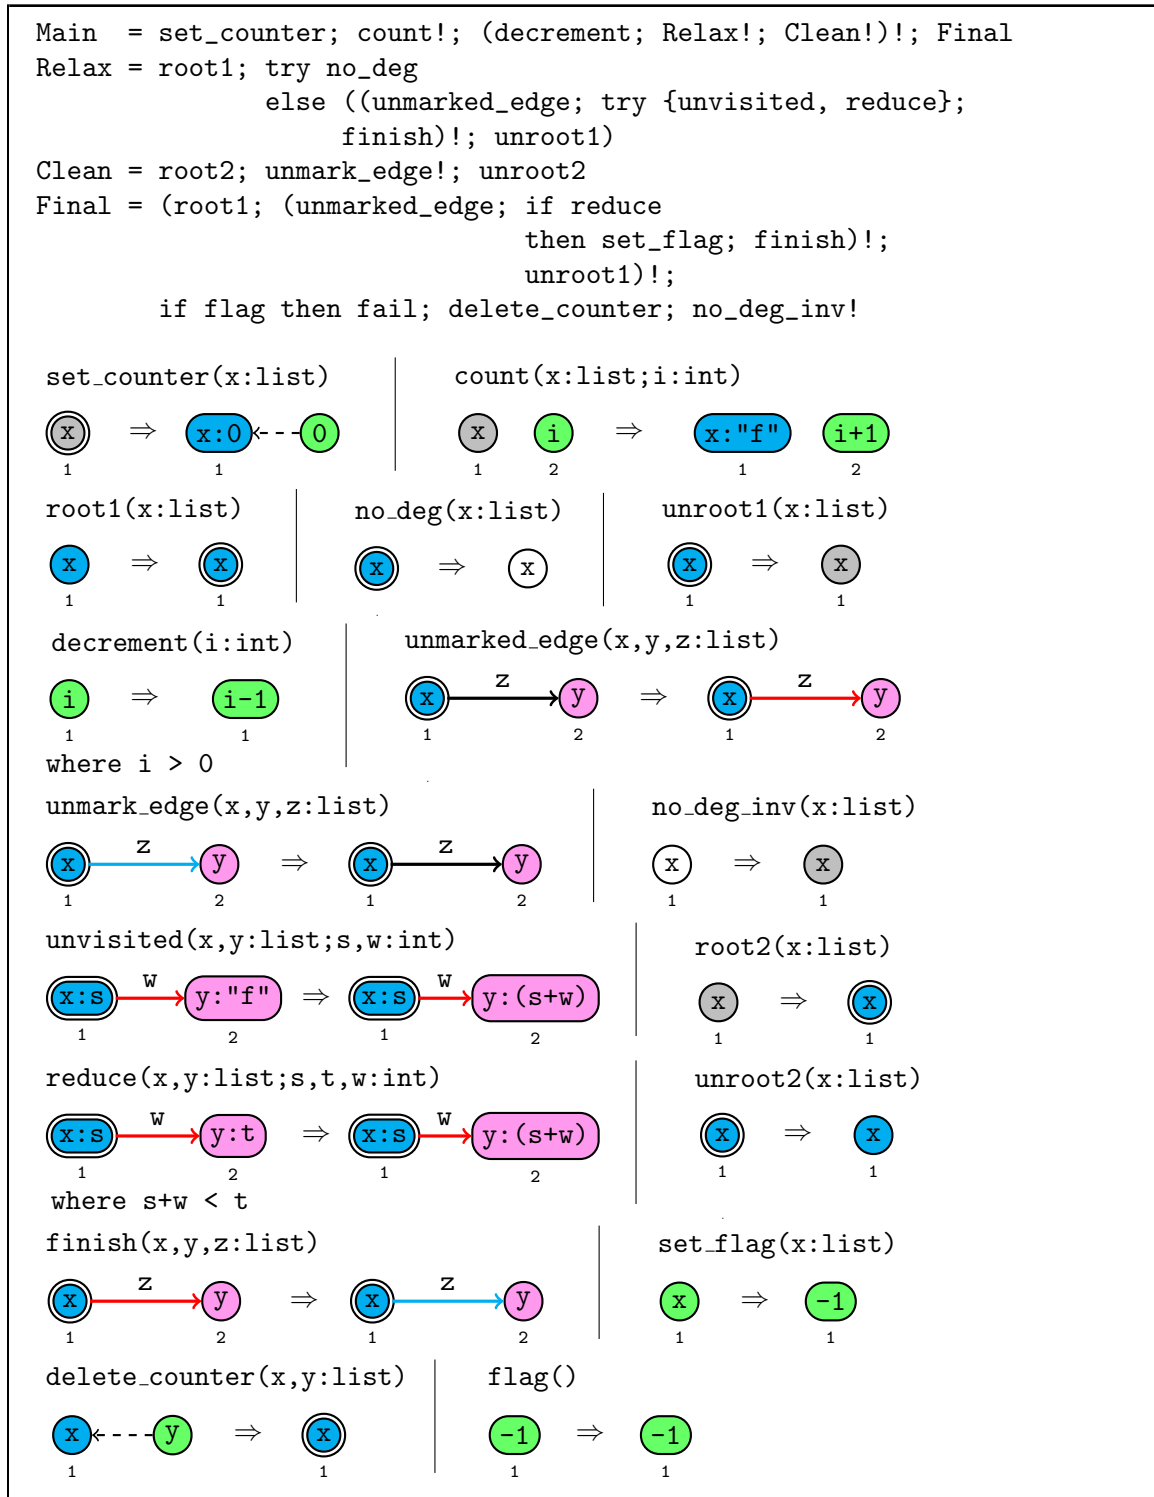


Figure 26: The program bellman-ford.

Note that the input graph may contain negative-weight cycles reachable from the root, as the program is designed to detect them. The program first creates a *counter* node labelled with the integer  $n - 1$ , where  $n$  is the number of nodes in the input graph. The node is marked green, a mark unique to that node to allow for fast, constant-time access. The rule `set_counter` initialises this process by creating a green node labelled 0 and marking the root node blue to indicate it has been counted. The rule also unroots the node, creates a dashed edge from the counter to store the location of the original root, and appends the value 0 to the previously rooted node's label to denote that its distance to itself is 0.

The rule `count` is then applied iteratively to increment the counter for each grey node in the graph. As nodes are counted, they are marked blue to prevent duplicate counting. Additionally, the rule appends the string `f` to each node's label. Here, `f` serves as a representation of infinity, in practice commonly represented by a large positive integer in Bellman-Ford implementations, but is implemented as a string to avoid dependency on the integer size limitations of the GP 2 compiler. Define  $l(v)$  as the value appended to the list label of node  $v$  by `set_counter` or `count`.

The sequence `(decrement; Relax!; Clean!)` implements the relaxation phase of the Bellman-Ford algorithm. The rule `decrement` decrements the value in the counter by 1 should that value be non-zero; this is to ensure that the relaxation phase happens exactly  $n - 1$  times. Initially, all nodes incident to at least one edge are marked blue in the host graph, no root exists, and all edges are unmarked. A single execution of `Relax` accomplishes the following:

- (1) The rule `root1` selects a blue node and roots it. At this stage, there is exactly one root in the host graph.
- (2) The rule `no_deg` checks whether the root node is isolated (notice that the node is not in the interface). If it is, it unmarks and unroots it to prevent the procedure `Relax` from considering it again in a future iteration. Isolated nodes do not have outgoing edges, so they need not be considered in the relaxation phase.
- (3) In case the node is not isolated, the sequence `(unmarked_edge; try {unvisited, reduce})` is invoked as long as possible. `unmarked_edge` selects an unmarked edge outgoing from the root node and marks it red; let  $u$  and  $v$  be the source and target of the edge, respectively, and  $w$  the weight of the edge; then two possible cases arise:
  - (a)  $l(u)$  is an integer value (i.e. not the string `f`). If  $l(v) = \mathbf{f}$ , then the rule `unvisited` applies and sets  $l(v)$  to  $l(u) + w$ . Otherwise, if  $l(u) + w$  is strictly lesser than  $l(v)$ , the rule `reduce` applies and sets  $l(v)$  to  $l(u) + w$ ; else, none of the rules in the rule set `{unvisited, reduce}` applies.
  - (b)  $l(u)$  is the string `f`. None of the rules in the rule set `{unvisited, reduce}` applies.
- (4) The rule `finish` marks the selected edge blue to denote that it has been processed, so that it is no longer considered for relaxation until all other unmarked edges have been relaxed.
- (5) The rule `unroot1` unroots the selected root node and unmarks it, so that it is no longer considered as a source node for relaxation until all other grey nodes have been considered. After the application of this rule, there is no root node in the host graph.

Following an execution of `Relax!`, all edges in the host graph are marked blue (due to `finish`), and all non-isolated nodes (i.e. nodes whose degrees are greater than 0) are marked grey (due to `unroot1`). The looping procedure `Clean!` prepares the host graph for another round of edge relaxations by unmarking all edges and turning all grey nodes blue.

The procedure **Final** checks that no negative-weight cycle reachable from the source exists in the host graph by running a modified version of the procedure **Relax!** once more. Note that the counter node is labelled 0 at this stage. Should the rule **reduce** be applicable, the counter node's label is set to  $-1$  by the rule **set\_flag**.

Finally, the program checks the applicability of the rule **flag**, which can only apply if and only if **set\_flag** had been applied. If **flag** is applicable, it entails that a negative cycle was found, and the program fails. Otherwise, no negative cycle exists in the host graph and **no\_deg\_inv!** marks all unmarked nodes (i.e. isolated nodes) grey, as they were in the input graph.

**8.2. Proof of Correctness.** We now examine the correctness of the program **bellman-ford**.

From now on, we call graphs satisfying the assumptions of the specification given at the beginning of the subsection as *input graphs*.

**Proposition 8.1.** *Let  $G$ ,  $I$  and  $H$  be graphs such that  $G$  is an input graph, and  $G \Rightarrow_{\text{set\_counter}} I \Rightarrow_{\text{count!}} H$ . Then  $H$  is obtained from  $G$  as follows:*

- (1) *All nodes are marked blue.*
- (2) *The integer 0 is appended to the list label of the root node.*
- (3) *The string  $f$  is appended to the list labels of other nodes (i.e. non-root nodes).*
- (4) *A new green-marked node is added, labelled with the integer  $n - 1$ , where  $n$  is the number of nodes in  $G$  (prior to the addition of the node).*
- (5) *A new dashed-marked edge is added, with the green-marked node as the source and the root node in  $G$  as the target.*
- (6) *The root node is unrooted.*

*Proof.* Let us first look at the application of the rule **set\_counter**; that is,  $G \Rightarrow_{\text{set\_counter}} I$ . Initially, all nodes in  $G$  are grey and one node is rooted. Thus, the application of the rule **set\_counter** marks the root node blue, appends 0 to its list label, creates a green node labelled 0 and generates a dashed edge from that newly created green node pointing at the previously rooted node. Now, let us examine the application of **count!**; that is,  $I \Rightarrow_{\text{count!}} H$ .  $I$  consists of  $n + 1$  nodes and  $m$  edges, where  $n$  and  $m$  are the numbers of nodes and edges in  $G$ , with one green node, one blue node and  $n - 1$  grey nodes. The rule **count** marks one grey node blue, appends the string  $f$  to its list label, and increments the counter by 1. Since the rule is applied as many times as there are grey nodes in  $I$ , the counter is labelled  $n - 1$  in  $H$ . Hence,  $H$  consists of one green node labelled 0,  $n$  blue nodes,  $m$  unmarked edges retaining their source and target nodes (notice that none of the rules involves edges), and one dashed edge from the green node to the node rooted in  $G$ . In addition, the node rooted in  $G$  has 0 appended to its list label in  $H$ , and all other nodes have the string  $f$  appended to their list labels.  $\square$

**Lemma 8.2** (Termination of **Relax!**). *The looping procedure **Relax!** terminates.*

*Proof.* Let us show that the loop (**unmarked\_edge**; **try\_unvisited**, **reduce**; **finish**)! terminates first. To show termination, consider the measure  $\#X$  consisting of the number of unmarked edges in the host graph  $X$ . The rule **unmarked\_edge** is invoked at the beginning of the loop, and if it fails to match, the loop breaks. Clearly, an application of **unmarked\_edge** reduces the measure  $\#$ . Since the number of edges is finite and no rule in the loop unmarks an edge, the rule eventually fails and the loop terminates.

Coming back to **Relax!**, now consider  $\#X$  to be the measure consisting of the number of blue nodes in the host graph  $X$ . Clearly, an application of **root1** reduces the measure  $\#$ . Since there is a finite number of nodes in the graph and no other rule in the procedure marks a node blue, the rule eventually fails and **Relax!** terminates.  $\square$

**Lemma 8.3** (Termination of **Clean!**). *The looping procedure **Clean!** terminates.*

*Proof.* Define  $\#X$  as the measure consisting of the number of grey nodes in the host graph  $X$ . An application of **root2** clearly reduces that measure, and since there is a finite number of nodes in the host graph and no rule in **Clean** marks a node grey, the rule eventually fails and the loop terminates.  $\square$

**Proposition 8.4.** *Upon the execution of the program **bellman-ford** on an input graph  $G$ , the looping sequence (**decrement**; **Relax!**; **Clean!**)! terminates after exactly  $n - 1$  iterations, where  $n$  is the number of nodes in  $G$ .*

*Proof.* Termination of **Relax!** and **Count!** follow from Lemmata 8.2 and 8.3.

Recall from Proposition 8.1 that the host graph contains exactly one green node labelled  $n - 1$  after the execution of the loop **count!**. The looping sequence terminates if and only if the rule **decrement** fails to apply. However, since the rule **decrement** is the only rule in the sequence to involve the counter node, then the sequence must be invoked exactly  $n - 1$  times before **decrement** can no longer apply and the sequence terminates.  $\square$

**Proposition 8.5.** *Upon termination of **Relax!**, all non-counter and non-isolated nodes in the host graph are grey.*

*Proof.* We first show that if **Relax!** has terminated, then the last call to **root1** failed.

For the sake of contradiction, suppose that the last call to **root1** succeeded. We show that **root1** is called again, contradicting our assumption.

Upon inspection of the program sequence, a successful application of **root1** leads to a call to **no\_deg**. If the latter succeeds, then **Relax** will be called, implying that **root1** will be called again. If **no\_deg** fails to apply, then (**unmarked\_edge**; **try\_unvisited**, **reduce**; **finish**)! is executed. If this loop terminates, then **unroot1** will be applied because the root node is blue, implying that **Relax** and hence **root1** will be applied again. If the loop does not terminate, then **Relax!** does not terminate, contradicting Lemma 8.2.

Since **Relax!** terminates, the last call to **root1** failed. This implies that there is no blue node in the host graph. It is easy to see that unmarked nodes can only be produced by **no\_deg**, which applies to isolated nodes only. Hence, there are neither blue nor unmarked nodes in the graph, and thus all non-counter and non-isolated nodes must be grey.  $\square$

**Proposition 8.6.** *On a host graph following the application of **decrement** during the execution of **bellman-ford** on an input graph, the looping procedure **Relax!** relaxes each edge in the host graph exactly once.*

*Proof.* Initially, all non-isolated nodes are blue; recall that isolated nodes need not be considered in the relaxation phase as they do not admit of any edge incident to them.

Once a blue node is rooted by **root1**, either **no\_deg** applies and **Relax!** moves on to its next iteration, or it does not apply and the looping sequence (**unmarked\_edge**; **try\_unvisited**, **reduce**}; **finish**)! is executed instead. Given that all edges outgoing from non-counter nodes are unmarked initially, the rule **unmarked\_edge** applies for any edge outgoing from the root node following the application of **root1** on that node. After such

a rule is applied, if the root node's distance written on its list label is infinity (i.e. the string `f`), then the edge cannot be used to reduce the distance inscribed on the target of the edge and none of the rules in the rule set `{unvisited, reduce}` applies, and the program subsequently applies `finish` to mark the edge blue, to denote that it has been processed in the current iteration of `Relax!`. Otherwise, the last variable on the list label of the root is an integer; in that case, either the edge is used to reduce the distance of the target node via the application of either `unvisited` in case the distance of the target node is still infinity, or the application of `reduce` should the weight added to the distance of the root be strictly smaller than the distance inscribed on the target of the edge, or none of those applies (the edge cannot be used to improve the distance on the target node) and the program applies `finish` to mark the edge blue to denote it has been processed in the iteration of `Relax!`.

The loop `(unmarked_edge; try {unvisited, reduce}; finish)!` terminates if and only if `unmarked_edge` is inapplicable (observe that `finish` will always apply following an application of `unmarked_edge`). Therefore, the loop will terminate when all outgoing edges are processed and turned blue by `finish`. Since, by Proposition 8.5, the loop `Relax!` terminates once all blue nodes are processed by `root1`, all unmarked edges are incident to some blue node in the host graph and `(unmarked_edge; try {unvisited, reduce}; finish)!` terminates once all unmarked outgoing edges are processed by `unmarked_edge` and relaxed, then the looping procedure `Relax!` relaxes all edges in the host graph.

Finally, observe that no rule in the procedure `Relax` changes the mark of a blue edge; therefore, every edge is relaxed exactly once.  $\square$

The next proposition establishes a graph-theoretical property that is needed in the subsequent proof.

**Proposition 8.7** (Path Size). *Let  $G$  be an input graph where  $s$  the root node in  $G$ , no negative-weight cycle is reachable from  $s$ , and  $t$  any node reachable from  $s$  in  $G$ . Let  $P$  be the shortest path of edges from  $s$  to  $t$  such that the sum of the weights of the edges in the path is minimal among all paths from  $s$  to  $t$ . Then, the number of edges in  $P$  is at most  $n - 1$ , where  $n$  is the number of nodes in  $G$ .*

*Proof.* For the sake of contradiction, assume that  $P$  contains more than  $n - 1$  edges. Since  $G$  contains  $n$  nodes, this implies that  $P$  must revisit at least one node, forming a cycle in the path. Let this cycle be denoted as  $C$ .

If the total weight of the edges in  $C$  is positive, removing the cycle would result in a path with a smaller total weight, contradicting the assumption that the sum of the weights of the edges in  $P$  is minimal among all paths from  $s$  to  $t$ . If the total weight of the edges in  $C$  is zero, removing the cycle does not change the total weight of the path. In this case, we can remove  $C$  and obtain a shorter path (in terms of the number of edges) with the same weight, contradicting the minimality of  $P$  as a shortest path.

In all cases, we arrive at a contradiction. Therefore,  $P$  cannot contain more than  $n - 1$  edges. Since the number of edges in any simple path (i.e. a path that does not contain cycles) in  $G$  is at most  $n - 1$ , the shortest path  $P$  from  $s$  to  $t$  must also satisfy this condition.  $\square$

For the remainder of the proof, define  $l(v)$  as the value appended to the list label of node  $v$  by `set_counter` or `count` in the program `bellman-ford`.

**Proposition 8.8** (Computation of Shortest Distances). *After  $k$  iterations of the loop `(decrement; Relax!; Clean!)!`, for every non-counter node  $v$  in the host graph reachable*

from the root node  $r$  in the input graph by a path of at most  $k$  edges,  $l(v)$  is set to the shortest weighted distance from  $r$  to  $v$  by a path of at most  $k$  edges.

*Proof.* Let us proceed by induction on the number of iterations of the looping sequence **(decrement; Relax!; Clean!)!**.

When  $k = 0$  (base case), the program set  $l(r)$  to 0, and for all other nodes  $v \neq r$ ,  $l(v)$  was assigned the value  $\mathbf{f}$ , representing infinity. At this stage, no edges were considered by the procedure **Relax**, so the shortest path to  $r$  from  $r$ , trivially of length 0, is correctly set to 0, and the shortest path to all other nodes remains undefined (i.e. infinity, the string  $\mathbf{f}$ ). Thus, the lemma holds for  $k = 0$ .

Now, assume the property holds for some  $k$ ; we show it still holds for  $k + 1$ . Consider the  $(k + 1)^{\text{th}}$  iteration of **(decrement; Relax!; Clean!)!**. During this iteration, the procedure **Relax!** processes all edges in the host graph (Proposition 8.6). Indeed, for each edge  $(u, v)$  with weight  $w$ , the program checks whether the distance to  $v$  can be improved by relaxing the edge  $(u, v)$  from a node  $u$  whose current shortest distance  $l(u)$  has already been computed in prior iterations (by the induction hypothesis). The relaxation phase **Relax!** updates<sup>9</sup>  $l(v)$  as follows:

$$l(v) \leftarrow \min(l(v), l(u) + w).$$

If there exists a path from  $r$  to  $v$  consisting of at most  $k + 1$  edges, it must either:

- (1) pass through a node  $u$  such that the shortest path to  $u$  uses at most  $k$  edges, followed by the edge  $(u, v)$ ; or
- (2) be a path already computed during the first  $k$  iterations, using at most  $k$  edges.

Therefore, during the  $(k + 1)^{\text{th}}$  iteration, the looping procedure **Relax!** correctly updates  $l(v)$  to the shortest distance from  $r$  to  $v$  using at most  $k + 1$  edges.  $\square$

**Proposition 8.9.** *Upon the execution of **bellman-ford** on an input graph  $G$ , the rule **set\_flag** is applied if and only if  $G$  contains a negative-weight cycle reachable from the root node.*

*Proof.* Let  $r$  be the root node in  $G$ . From Proposition 8.7, it is established that the shortest path of edges from  $r$  to some arbitrary non-counter node  $v$  in the host graph such that the sum of the weights of the edges in the path is minimal among all paths from  $r$  to  $v$  consists of at most  $n - 1$  edges (where  $n$  is the number of nodes in  $G$ ), provided no negative-weight cycle reachable from  $r$  exists. Proposition 8.8 states that  $k$  iterations of the loop **(decrement; Relax!; Clean!)!**, called prior to **Final**, compute the smallest weighted distances from  $r$  to any non-counter node  $v$  in the host graph by a path of at most  $k$  edges. Furthermore, Proposition 8.4 establishes that this loop terminates after  $n - 1$  iterations. Therefore, upon termination of **(decrement; Relax!; Clean!)!**, the program has computed the smallest distances from  $r$  to all other non-counter nodes in the host graph by a path of at most  $n - 1$  edges.

The procedure **Final** executes one iteration of **Relax**, during which only the rule **reduce** is considered for relaxation. If the rule **reduce** is applicable during this iteration, it implies that there exists some node  $u$  with  $l(u)$  set to the shortest distance from  $r$  to  $u$  (by a path of at most  $n - 1$  edges), and an edge  $(u, v)$  with weight  $w$  such that  $l(u) + w < l(v)$ . In this case, the rule **reduce** would update  $l(v)$  to  $l(u) + w$ .

For such a relaxation to occur after  $n - 1$  iterations, there must exist a path from  $r$  to  $v$  consisting of more than  $n - 1$  edges that has a smaller total weight than any path of at

<sup>9</sup> $\min(l(v), l(u) + w) = l(u) + w$  if  $l(v)$  is not an integer.

most  $n - 1$  edges. This is only possible if there exists a negative-weight cycle reachable from  $r$ . Indeed, by repeatedly traversing the negative-weight cycle, the total path weight can decrease indefinitely, which would for a smaller value of  $l(v)$  to be computed.

When the rule `reduce` applies, the rule `set_flag` subsequently sets the counter node's label to  $-1$ . This serves as an indicator that a negative-weight cycle was detected. Finally, conversely, if no negative-weight cycle is reachable from  $r$ , then no edge relaxation can occur during the `Final` procedure, and the rule `set_flag` is never applied.  $\square$

**Theorem 8.10** (Correctness of `bellman-ford`). *The program `bellman-ford` is totally correct with respect to the specification at the beginning of Subsection 8.1.*

*Proof.* Termination easily follows from Lemmata 8.2 and 8.3, and Proposition 8.4. An argument analogous to the proof of Lemma 8.2 can be made with respect to the termination of `Final`.

Let  $G$  be the input graph. Let us split the remainder of this proof into two cases.

*Case 1:*  $G$  contains a negative-weight cycle reachable from the source. Upon the termination of `count!`, the counter node is labelled  $n - 1$ , where  $n$  is the number of nodes in  $G$ . By Proposition 8.4, the rule `decrement` is applied exactly  $n - 1$ . Therefore, upon termination of the loop `(decrement; Relax!; Clean!)`, the counter node is labelled 0. As per Proposition 8.9, `set_flag` will apply during the execution of `Final`, and given that `set_flag` is the only rule tempering with the list label of the counter node, an application of `flag` would entail that `set_flag` was applied, and a negative-weight cycle reachable from the root was found. The program would subsequently invoke the `fail` command and fail, satisfying the specifications of the program.

*Case 2:*  $G$  does not contain a negative-weight cycle reachable from the source. Analogously to the previous case, the counter node is labelled  $n - 1$  upon termination of the loop `(decrement; Relax!; Clean!)`. Given that no negative-weight cycle reachable from the source node exists in  $G$ , the rule `set_flag` is never applied and the `fail` command is never invoked. The rule `delete_counter` returns the root to the source node and deletes the counter (observe that no rule beside `set_counter` and `delete_counter` involves a dashed edge). Finally, as all shortest distances from the source can be achieved by a path of at most  $n - 1$  edges (Proposition 8.7), the loop `(decrement; Relax!; Clean!)` terminates after  $n - 1$  iterations (Proposition 8.4), and  $k$  iterations of the loop label the shortest weighted distances from the root to all other nodes by a path of at most  $k$  edges (Lemma 8.8), all nodes reachable from the source have an integer appended to their list label denoting the smallest weighted distance from the root. Nodes that are not reachable from the source retained the string `f`.  $\square$

**8.3. Proof of Complexity.** Let us now analyse the time complexity of the program `bellman-ford`.

**Proposition 8.11.** *Every rule in the program `bellman-ford` matches or fails to match in constant time with respect to the complexity assumptions of the updated compiler (Figure 6).*

*Proof.* Since the input graph contains exactly one root node, a call to the rule `set_counter` requires only a constant-time lookup in the set of root nodes.

The rule `count` matches node 1 in constant time since the list label of the node can be arbitrary, and node 2 in constant time as well since there exists at most one green node in

the host graph. Node 1 of `decrement`, `set_flag` and `flag` is also matched in constant time due to the bound on the number of green nodes.

The rule `root1` matches any blue node in the host graph of arbitrary list label, and thus in constant time. An analogous argument can be made for `root2`. Since there is exactly one root node after an application of `root1`, `no_deg` and `unroot1` match in constant time. An analogous argument can be made for `unroot2` with respect to `root2`.

The rule `unmarked_edge` matches node 1 in constant time due to the bound on the number of roots in the host graph, and matches any unmarked edge of arbitrary list label outgoing from the root in constant time. The rules `finish`, `unvisited` and `reduce` match node 1 and the red edge in constant time, since there is at most one red edge throughout the execution of the program `bellman-ford`. The rule `unmark_edge` matches node 1 and the blue edge in constant time, given that the latter can have an arbitrary list label.

Finally, `delete_counter` matches node 1 and the dashed edge in constant time, as there is at most one dashed edge throughout the execution of the program. And the rule `no_deg_inv` matches in constant time given that the list label of the node can be arbitrary.  $\square$

The following proposition is a simple graph-theoretic property that we subsequently need.

**Proposition 8.12.** *For every graph  $G$ , the number of non-isolated nodes is at most  $2|E_G|$ .*

*Proof.* Given a node  $v$ , we write  $\deg(v)$  for its degree.<sup>10</sup> It is a basic fact that every graph  $G$  satisfies  $\sum_{v \in V_G} \deg(v) = 2|E_G|$  (see, for example, [BJG09]). Clearly, if  $v$  is not isolated, we have  $\deg(v) \geq 1$ . Thus there can be at most  $2|E_G|$  non-isolated nodes.  $\square$

**Theorem 8.13** (Complexity of `bellman-ford`). *On any class of input graphs, the program `bellman-ford` terminates in time  $O(nm)$ , where  $n$  and  $m$  are the numbers of nodes and edges in the input graph, respectively.*

*Proof.* We established in Proposition 8.11 that, for every rule, there is at most one matching attempt with respect to the complexity assumptions of the updated compiler (Figure 6). Let us now examine how many times each rule is called.

Clearly, the rules `set_counter`, `delete_counter` and `flag` are called at most once. Following an application of `set_counter`, there are  $n - 1$  grey nodes in the host graph, where  $n$  is the number of nodes in the input graph. The loop terminates when there is no grey node left, thus it is called  $n$  times. Proposition 8.4 states that the looping sequence (`decrement`; `Relax!`; `Clean!`) terminates after  $n - 1$  iterations, thus the rule `decrement` is called  $n$  times. Let that sequence be referred to as `R`. Prior to the very first iteration of `R`, during the execution of the program, isolated nodes are not unmarked. Following the iteration, every isolated node is unmarked. Therefore, the rule `root1` is called  $n + 1$  times in the first iteration of `R` in the execution of the program, and  $n' + 1$  times on each iteration afterward, where  $n'$  is the number of non-isolated non-counter nodes in the host graph. The rule `no_deg` is called after each application of `root1`, thus  $n$  times on the first iteration of `R`, and  $n'$  times afterward. The loop (`unmarked_edge`; `try unvisited`, `reduce`; `finish`)! is called as many times as there are unmarked edges in the host graph and non-isolated nodes, thus  $m + n'$  times where  $m$  is the number of edges in the input graph. The rule `unroot1` is called  $n'$  times, that is, as many times as there are non-isolated nodes.

<sup>10</sup>The degree of a node is the sum of its indegree and outdegree.

In each execution of `Clean!`, the rule `root2` is called  $n' + 1$  times, and the rule `unroot2`,  $n'$  times. The rule `unmark_edge` is called  $m + n'$  times ( $m$  successful applications and  $n'$  failed attempts to terminate the loop on each non-isolated node). Finally, the rule `root1` in `Final` is called at most  $n' + 1$  times, `unmarked_edge` is called at most  $m + n'$  times, `reduce` at most  $m$  times and `finish`, at most  $m$  times.

By Proposition 8.12,  $n'$  is at most  $2m$ . Therefore, on all iterations of `(decrement; Relax!; Clean!)!` but the first one, a number of calls at most linear to the number of edges are called. On the first iteration, a number of calls linear to the size of the graph are called.

Therefore, the overall time complexity of the program `bellman-ford` is  $O(nm)$ .  $\square$




Figure 27: Rooted square grid.




Figure 28: Rooted binary tree.




Figure 29: Rooted star graph.




Figure 30: Rooted cycle graph.




Figure 31: Rooted complete graph.




Figure 32: Rooted linked list.



Figure 33: Rooted discrete graph.

We supplement the proof with empirical measurements on the graph classes of Figures 27, 28, 29, 30, 31, 32 and 33. In Figures 34 and 35, all input graph edges are labelled with an integer chosen uniformly at random in the interval  $[-100, 100]$ ; the allocation of weights was done such that cycle graphs do not contain negative-weight cycles. And finally, Figure 36 shows the runtimes on cycle graphs whose weights alternate between  $-2$  and  $1$  such that the total weight is negative. Hence, the program fails on these inputs.

The low runtime of complete graphs, compared with the other graph classes, is due to the fact that the number of edges is quadratic in the number of nodes, and hence the runtime bound  $nm$  grows at a slower rate (note that the scales of the plots are different).

In Figure 35, we include the runtimes for discrete graphs to stress that even for these simple graphs, the time complexity is not constant but linear. That is because the loop `count!` does require time linear in the number of nodes.

## 9. CONCLUSION

We have demonstrated by case studies how to implement three fundamental graph algorithms in GP 2 such that their imperative time complexities are matched, even if input graphs have an unbounded node degree or are possibly disconnected. Addressing the issues of unbounded degree and disconnectedness has been an open problem since the publication of the first paper on rooted graph transformation [BP12]. Up to now, only certain graph reduction programs that destroy their input graphs could be constructed to run in linear time on graph classes that have an unbounded node degree or contain disconnected graphs [CCP22].




Figure 34: Measured performance of the program `bellman-ford` on rooted graph classes with random weights.




Figure 35: Measured performance of the program `bellman-ford` on rooted discrete graphs, and rooted complete graphs with random weights.

In Subsections 5.1 and 5.3, we show how the graph data structure generated by the new compiler enables linear runtimes on arbitrary GP 2 host graphs.

Our case studies include an implementation of the Bellman-Ford single-source shortest-path algorithm that matches the  $O(nm)$  time complexity of imperative implementations. This is, to the best of our knowledge, the first time that a rule-based shortest-path algorithm achieves the complexity of a corresponding algorithm in an imperative programming language.

Our approach involves both enhancing the graph data structure generated by the GP 2 compiler and developing a programming technique that leverages the new graph representation. Previously, the graph data structure in the C program generated by the

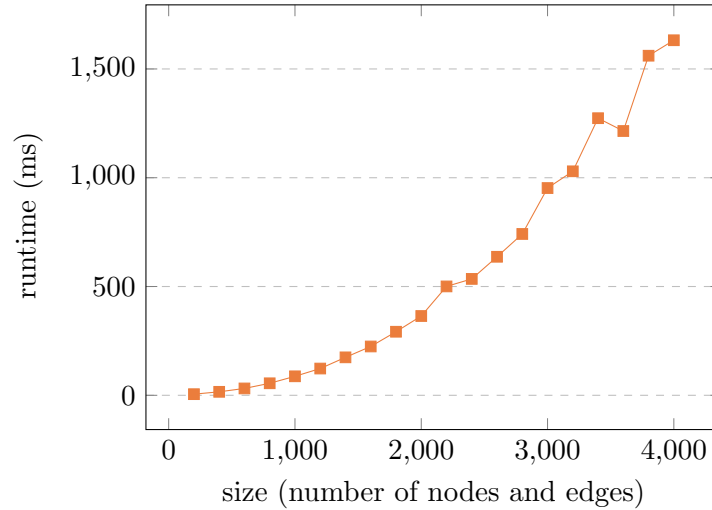


Figure 36: Measured performance of the program `bellman-ford` on rooted cycle graphs of negative total weight.

compiler stored in each host-graph node two linked lists of incoming and outgoing edges. In contrast, the new two-dimensional array of linked lists of edges allows to find an edge with a given mark and orientation in constant time. Also, nodes are now stored in multiple linked lists, one for each mark. Thus, a node of a given mark can be found in constant time.

Comparing our proofs of correctness and complexity in the case studies with corresponding proofs for imperative programs in standard textbooks such as [CLRS22], the latter may appear more concise. However, one should note that those proofs are at the pseudo-code level, whereas ours address all details at the level of concrete syntax, including the specifics of the compiler implementation.

We speculate that many more graph algorithms can be implemented as GP 2 programs running under time bounds achievable with conventional programming languages. To make progress in developing such algorithms, we intend to implement a GP 2 library of advanced data structures, such as priority queues, Fibonacci heaps, or AVL trees, supporting programmers in constructing GP 2 versions of efficient imperative graph algorithms.

## APPENDIX A. ABSTRACT SYNTAX OF LABELS AND CONDITIONS

This appendix gives grammars in Extended Backus-Naur Form defining the abstract syntax of labels and application conditions. There are some context constraints for these definitions which can be found in [Bak15].

The class of *host graphs* consists of all totally labelled graphs over the set `HostLabel` of host graph labels as defined by the grammar in Figure 37.

Figure 38 defines the labels that may occur in the graphs contained in rules. (In the case studies of this paper, however, we only need labels that are variables of type `list` or `int`.) We obtain the class of *rule graphs* by replacing the label set  $\mathcal{L}$  of Definition 2.1 with the set `Label` of rule graph labels.

The concatenation of two lists  $x$  and  $y$  is written  $x:y$  (not to be confused with Haskell’s colon operator which adds an element to the beginning of a list). The binary arithmetic

```

HostLabel ::= HostList [HostMark]
HostList  ::= empty | HostAtom {':' HostAtom}
HostAtom  ::= HostInteger | HostString
HostInteger ::= ['-'] Digit {Digit}
HostString ::= "'{HostChar}'"
HostChar   ::= "'Character'"
HostMark   ::= red | green | blue | grey | dashed

```

Figure 37: Abstract syntax of host graph labels.

```

Label ::= List [Mark]
List  ::= LVar | empty | Atom | List ':' List
Atom  ::= AVar | Integer | String
Integer ::= IVar | ['-'] Digit {Digit} | ('Integer') |
          Integer ('+' | '-' | '*' | '/') Integer |
          (indeg | outdeg) ('Node') |
          length ('(LVar | AVar | SVar)')
String ::= SVar | Char | "'{Character}'" |
          String '.' String
Char    ::= CVar | "'Character'"
Mark    ::= red | green | blue | grey | dashed | any

```

Figure 38: Abstract syntax of rule graph labels.

operators '+', '-', '\*', and '/' expect integer expressions as arguments while '.' is string concatenation. The categories LVar, AVar, IVar, SVar and CVar consist of variables of type List, Atom, Integer, String and Char, respectively. Variables can be used as expressions of supertypes. For example, integer variables can be used as lists of length one. The category Node consists of identifiers for nodes in the left-hand graphs of rules. The operators `indeg` and `outdeg` denote the number of ingoing respectively outgoing edges of a node. The `length` operator returns the length of a list or string represented by a variable. The category Character consists of the printable ASCII characters except "'". Character strings are enclosed in double quotes.

The grammar in Figure 39 defines the abstract syntax of rule application conditions. Conditions are Boolean combinations of subtype assertions, applications of the `edge` predicate to left-hand node identifiers, or relational comparisons of expressions (where = and != can be used for arbitrary expressions). A condition of the form `edge(v1, v2, l)` requires the existence of an edge from node  $v_1$  to node  $v_2$  with label  $l$ .

```

Condition ::= (int | char | string | atom) '('(LVar | AVar | SVar)')
           | List ('=' | '!=') List
           | Integer ('>' | '>=' | '<' | '<=') Integer
           | edge '(' Node ',' Node [',' Label] ')
           | not Condition
           | Condition (and | or) Condition
           | '(' Condition ')

```

Figure 39: Abstract syntax of application conditions.

## ACKNOWLEDGEMENTS

The authors thank Brian Courtehoue and Robert Söldner for discussions on the topics of this paper, and the anonymous reviewers for their helpful comments.

## REFERENCES

- [AHU74] Alfred Aho, John Hopcroft, and Jeffrey Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.
- [AKN<sup>+</sup>06] Aditya Agrawal, Gabor Karsai, Sandeep Neema, Feng Shi, and Attila Vizhanyo. The design of a language for model transformations. *Software & Systems Modeling*, 5(3):261–288, 2006. doi:10.1007/s10270-006-0027-7.
- [AP24] Ziad Ismaili Alaoui and Detlef Plump. Linear-time graph programs for unbounded-degree graphs. In *Proc. 17th International Conference on Graph Transformation (ICGT 2024)*, volume 14774 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2024. doi:10.1007/978-3-031-64285-2\_1.
- [AP25] Ziad Ismaili Alaoui and Detlef Plump. Linear-time graph programs without preconditions. In *Proc. 14th and 15th International Workshop on Graph Computation Models (GCM 2023 and GCM 2024)*, volume 417 of *Electronic Proceedings in Theoretical Computer Science*, pages 39–54. Open Publishing Association, 2025. doi:10.4204/EPTCS.417.3.
- [Bak15] Christopher Bak. *GP 2: Efficient Implementation of a Graph Programming Language*. PhD thesis, Department of Computer Science, University of York, UK, 2015. URL: <https://etheses.whiterose.ac.uk/12586/>.
- [BJG09] Jørgen Bang-Jensen and Gregory Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer, second edition, 2009. doi:10.1007/978-1-84800-998-1.
- [BP12] Christopher Bak and Detlef Plump. Rooted graph programs. In *Proc. 7th International Workshop on Graph Based Tools (GraBaTs 2012)*, volume 54 of *Electronic Communications of the EASST*, 2012. doi:10.14279/tuj.eceasst.54.780.
- [BP16] Christopher Bak and Detlef Plump. Compiling graph programs to C. In *Proc. 9th International Conference on Graph Transformation (ICGT 2016)*, volume 9761 of *Lecture Notes in Computer Science*, pages 102–117. Springer, 2016. doi:10.1007/978-3-319-40530-8\_7.
- [CCP19] Graham Campbell, Brian Courtehoue, and Detlef Plump. Linear-time graph algorithms in GP 2. In *Proc. 8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019)*, volume 139 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:23. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.CALCO.2019.16.
- [CCP22] Graham Campbell, Brian Courtehoue, and Detlef Plump. Fast rule-based graph programs. *Science of Computer Programming*, 214:102727, 2022.
- [CLRS22] Thomas H. Cormen, Charles E. Leiserson, Robert L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, fourth edition, 2022.
- [Cou23] Brian Coutehoute. *Time and Space Complexity of Rule-Based Graph Programs*. PhD thesis, Department of Computer Science, University of York, 2023. URL: <https://etheses.whiterose.ac.uk/33407/>.

- [CP21] Brian Courtehoue and Detlef Plump. A small-step operational semantics for GP 2. In *12th International Workshop on Graph Computation Models (GCM 2021), Revised Selected Papers*, volume 350 of *Electronic Proceedings in Theoretical Computer Science*, pages 89–110. Open Publishing Association, 2021. doi:10.4204/EPTCS.350.6.
- [CRP20] Graham Campbell, Jack Romö, and Detlef Plump. The improved GP 2 compiler. Technical report, Department of Computer Science, University of York, UK, 2020. URL: <https://arxiv.org/abs/2010.03993>.
- [Dör95] Heiko Dörr. *Efficient Graph Rewriting and its Implementation*, volume 922 of *Lecture Notes in Computer Science*. Springer, 1995. doi:10.1007/BFb0031909.
- [FKP19] Maribel Fernández, Hélène Kirchner, and Bruno Pinaud. Strategic port graph rewriting: an interactive modelling framework. *Mathematical Structures in Computer Science*, 29(5):615–662, 2019. doi:10.1017/S0960129518000270.
- [GdMR<sup>+</sup>12] Amir Ghamarian, Maarten de Mol, Arend Rensink, Eduardo Zambon, and Maria Zimakova. Modelling and analysis using GROOVE. *International Journal on Software Tools for Technology Transfer*, 14(1):15–40, 2012. doi:10.1007/s10009-011-0186-x.
- [HP02] Annegret Habel and Detlef Plump. Relabelling in graph transformation. In *Proc. First International Conference on Graph Transformation (ICGT 2002)*, volume 2505 of *Lecture Notes in Computer Science*, pages 135–147. Springer, 2002. doi:10.1007/3-540-45832-8\_12.
- [JBK10] Edgar Jakumeit, Sebastian Buchwald, and Moritz Kroll. GrGen.NET – the expressive, convenient and fast graph rewrite system. *International Journal on Software Tools for Technology Transfer*, 12(3–4):263–271, 2010. doi:10.1007/s10009-010-0148-8.
- [Plu12] Detlef Plump. The design of GP 2. In *Proc. 10th International Workshop on Reduction Strategies in Rewriting and Programming (WRS 2011)*, volume 82 of *Electronic Proceedings in Theoretical Computer Science*, pages 1–16, 2012. doi:10.4204/EPTCS.82.1.
- [SBG<sup>+</sup>17] Daniel Strüber, Kristopher Born, Kanwal Daud Gill, Raffaella Groner, Timo Kehrer, Manuel Ohrndorf, and Matthias Tichy. Henshin: A usability-focused framework for EMF model transformation development. In *Proc. 10th International Conference on Graph Transformation (ICGT 2017)*, volume 10373 of *Lecture Notes in Computer Science*, pages 196–208. Springer, 2017. doi:10.1007/978-3-319-61470-0\_12.
- [Sed02] Robert Sedgewick. *Algorithms in C. Part 5: Graph Algorithms*. Addison-Wesley, third edition, 2002.
- [Ski20] Steven S. Skiena. *The Algorithm Design Manual*. Springer, third edition, 2020. doi:10.1007/978-3-030-54256-6.