


# 1 Linear-Time Graph Algorithms in GP 2

2 **Graham Campbell** 

3 Department of Computer Science, University of York, United Kingdom

4 <https://gjcampbell.co.uk/>

5 gjc510@york.ac.uk

6 **Brian Courtehoue** 

7 Department of Computer Science, University of York, United Kingdom

8 <https://www.cs.york.ac.uk/people/brianc>

9 bc956@york.ac.uk

10 **Detlef Plump** 

11 Department of Computer Science, University of York, United Kingdom

12 <https://www-users.cs.york.ac.uk/det/>

13 detlef.plump@york.ac.uk

## 14 — Abstract —

15 GP 2 is an experimental programming language based on graph transformation rules which aims to  
16 facilitate program analysis and verification. However, implementing graph algorithms efficiently in a  
17 rule-based language is challenging because graph pattern matching is expensive. GP 2 mitigates this  
18 problem by providing *rooted* rules which, under mild conditions, can be matched in constant time.  
19 In this paper, we present linear-time GP 2 programs for three problems: tree recognition, binary  
20 DAG recognition, and topological sorting. In each case, we show the correctness of the program,  
21 prove its linear time complexity, and also give empirical evidence for the linear run time. For DAG  
22 recognition and topological sorting, the linear behaviour is achieved by implementing depth-first  
23 search strategies based on an encoding of stacks in graphs.

24 **2012 ACM Subject Classification** Theory of computation → Design and analysis of algorithms

25 **Keywords and phrases** Graph transformation; rooted graph programs; GP 2; linear-time algorithms;  
26 depth-first search; topological sorting

27 **Digital Object Identifier** 10.4230/LIPIcs.CALCO.2019.23

## 28 **1** Introduction

29 Rule-based graph transformation was established as a research field in the 1970s and has  
30 since then been the subject of countless articles. While many of these contributions have a  
31 theoretical nature (see the monograph [8] for a recent overview), there has also been work on  
32 languages and tools for executing and analysing graph transformation systems.

33 Languages based on graph transformation rules include AGG [17], GReAT [1], GROOVE  
34 [10], GrGen.Net [12], Henshin [3] and PORGY [9]. This paper focuses on GP 2 [13], an  
35 experimental graph programming language which aims to facilitate formal reasoning on  
36 programs. The language has a simple formal semantics and is computationally complete  
37 in that every computable function on graphs can be programmed [14]. Research on graph  
38 programs has provided, for example, a Hoare-calculus for program verification [15, 16] and a  
39 static analysis for confluence checking [11].

40 A challenge for the design and implementation of graph transformation languages is to  
41 narrow the performance gap between imperative and rule-based graph programming. The  
42 bottleneck for achieving fast graph transformation is the cost of graph matching. In general,  
43 matching the left-hand graph  $L$  of a rule within a host graph  $G$  requires time  $\text{size}(G)^{\text{size}(L)}$   
44 (which is polynomial since  $L$  is fixed). As a consequence, linear-time imperative graph  
45 algorithms may be slowed down to polynomial time when they are recast as rule-based graph



© Graham Campbell, Brian Courtehoue and Detlef Plump;

licensed under Creative Commons License CC-BY

8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019).

Editors: Markus Roggenbach and Ana Sokolova; Article No. 23; pp. 23:1–23:25

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 programs. To mitigate this problem, GP 2 allows to distinguish nodes as *roots* and to match  
 47 roots in rules with roots in host graphs. Then only the neighbourhood of host graph roots  
 48 needs to be searched for matches, allowing, under mild conditions, to match rules in constant  
 49 time.

50 In [5], *fast* rules were identified as a class of rooted rules that can be applied in constant  
 51 time if host graphs have a bounded node degree and contain a bounded number of roots.  
 52 A graph program with fast rules was shown in [6] to 2-colour graphs of bounded degree in  
 53 linear time. The compiled program matches the speed of Sedgewick’s textbook C program  
 54 [18] on grid graphs of up to 100,000 nodes.

55 In this paper, we continue this line of research with case studies on three linear-time  
 56 graph algorithms: recognition of trees, recognition of binary DAGs, and topological sorting.  
 57 In each case, we present a GP 2 program with fast rules, show its correctness, and prove its  
 58 linear time complexity on graphs of bounded node degree. We also give empirical evidence  
 59 for the linear run time by presenting benchmark results for graphs of up to 100,000 nodes in  
 60 various graph classes. For DAG recognition and topological sorting, the linear behaviour is  
 61 achieved by implementing depth-first search strategies based on an encoding of stacks.

## 62 **2 The Graph Programming Language GP 2**

63 This section briefly introduces GP 2, a non-deterministic language based on graph-transfor-  
 64 mation rules, first defined in [13]. Up-to-date versions of the syntax and semantics of GP 2  
 65 can be found in [4]. The language is implemented by a compiler generating C code [6].

### 66 **2.1 Graphs, Rules and Programs**

67 GP 2 programs transform input graphs into output graphs, where graphs are directed and  
 68 may contain parallel edges and loops. Both nodes and edges are labelled with lists consisting  
 69 of integers and character strings. This includes the special case of items labelled with the  
 70 empty list which may be considered as “unlabelled”.

71 The principal programming construct in GP 2 are conditional graph transformation rules  
 72 labelled with expressions. For example, the rule `one_of_one` in Figure 10 has four formal  
 73 parameters of type `list`, a left-hand graph and a right-hand graph which are specified  
 74 graphically, and a textual condition starting with the keyword `where`.

75 The small numbers attached to nodes are identifiers, all other text in the graphs are  
 76 labels. Parameters are typed but in this paper we only need the most general type `list`  
 77 which represents arbitrary lists.

78 Besides carrying expressions, nodes and edges can be *marked* red, green or blue. In  
 79 addition, nodes can be marked grey and edges can be dashed. For example, rule `one_of_one`  
 80 in Figure 10 contains red and blue nodes and a blue edge. Marks are convenient, among other  
 81 things, to record visited items during a graph traversal and to encode auxiliary structures in  
 82 graphs. The programs in the following sections use marks extensively.

83 Rules operate on *host graphs* which are labelled with constant values (lists containing  
 84 integer and string constants). Applying a rule  $L \Rightarrow R$  to a host graph  $G$  works roughly  
 85 as follows: (1) Replace the variables in  $L$  and  $R$  with constant values and evaluate the  
 86 expressions in  $L$  and  $R$ , to obtain an instantiated rule  $\hat{L} \Rightarrow \hat{R}$ . (2) Choose a subgraph  $S$  of  
 87  $G$  isomorphic to  $\hat{L}$  such that the dangling condition and the rule’s application condition are  
 88 satisfied (see below). (3) Replace  $S$  with  $\hat{R}$  as follows: numbered nodes stay in place (possibly  
 89 relabelled), edges and unnumbered nodes of  $\hat{L}$  are deleted, and edges and unnumbered nodes  
 90 of  $\hat{R}$  are inserted.

91 In this construction, the *dangling condition* requires that nodes in  $S$  corresponding to  
 92 unnumbered nodes in  $\hat{L}$  (which should be deleted) must not be incident with edges outside  
 93  $S$ . The rule's application condition is evaluated after variables have been replaced with the  
 94 corresponding values of  $\hat{L}$ , and node identifiers of  $L$  with the corresponding identifiers of  $S$ .  
 95 For example, the condition  $\text{indeg}(1) = 1$  of rule `one_of_one` in Figure 10 requires that node  
 96  $g(1)$  has exactly one incoming edge, where  $g(1)$  is the node in  $S$  corresponding to 1.

97 A program consists of declarations of conditional rules and procedures, and exactly one  
 98 declaration of a main command sequence. Procedures must be non-recursive, they can be  
 99 seen as macros. We describe GP 2's main control constructs.

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

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

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

112 In general, the execution of a program on a host graph may result in different graphs,  
 113 fail, or diverge. The operational semantics of GP 2 defines a semantic function which maps  
 114 each host graph to the set of all possible outcomes. See, for example, [14].

## 115 2.2 Rooted Programs

116 The bottleneck for efficiently implementing algorithms in a language based on graph trans-  
 117 formation rules is the cost of graph matching. In general, to match the left-hand graph  $L$  of a  
 118 rule within a host graph  $G$  requires time polynomial in the size of  $L$  [5, 6]. As a consequence,  
 119 linear-time graph algorithms in imperative languages may be slowed down to polynomial  
 120 time when they are recast as rule-based programs.

121 To speed up matching, GP 2 supports *rooted* graph transformation where graphs in rules  
 122 and host graphs are equipped with so-called root nodes. Roots in rules must match roots in  
 123 the host graph so that matches are restricted to the neighbourhood of the host graph's roots.  
 124 We draw root nodes using double circles. For example, in the rule `prune` of Figure 2, the  
 125 node labelled  $y$  in the left-hand side and the single node in the right-hand side are roots.

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

130 Conditions (2) and (3) will be satisfied by all rules occurring in the following sections; in  
 131 particular, we neither use the `edge` predicate nor the equality tests. For example, the rules  
 132 `prune` and `push` in Figure 2 are fast rules.

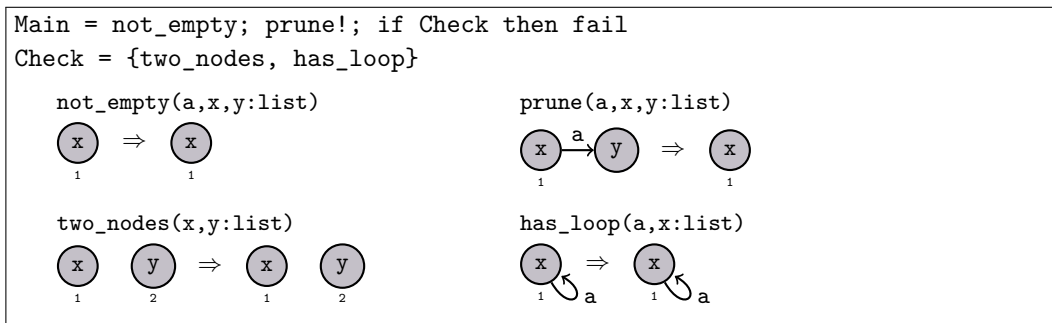
133 ► **Theorem 1** (Complexity of matching fast rules [5]). *Rooted graph matching can be imple-*  
 134 *mented to run in constant time for fast rules, provided there are upper bounds on the maximal*  
 135 *node degree and the number of roots in host graphs.*

## 23:4 Linear-Time Graph Algorithms in GP 2

136 When analysing the time complexity of rules and programs, we assume that these are  
 137 fixed. This is customary in algorithm analysis where programs are fixed and running time is  
 138 measured in terms of input size [2, 19]. In our setting, the input size is the size of a host  
 139 graph. The implementation of GP 2 does match fast rooted rules in constant time [6].

### 140 **3** Recognising Trees

141 A tree is a connected graph without undirected cycles such that every node has at most one  
 142 incoming edge. It is easy to see that it is possible to generate the collection of all non-empty  
 143 trees by inductively adding new leaf nodes onto the discrete graph of size one. Thus, given  
 144 an input graph, if we prune leaf nodes as long as possible and end up with the discrete graph  
 145 of size one, then the start graph must have been a tree. Figure 1 is implementation of this  
 146 idea in GP 2.



147 **Figure 1** The GP 2 program `is-tree-slow`

147 **► Definition 2** (Tree recognition specification). *The tree recognition specification is as follows.*

148 **■** Input: An arbitrary labelled graph with every node coloured grey, no root nodes, and no  
 149 other marks.

150 **■** Output: Fail if and only if the input is not a non-empty tree.

151 **► Theorem 3** (Correctness of `is-tree-slow`). *The program `is-slow-tree` fulfills the tree  
 152 recognition specification.*

153 **Proof.** Similar to the proof of Theorem 6. ◀

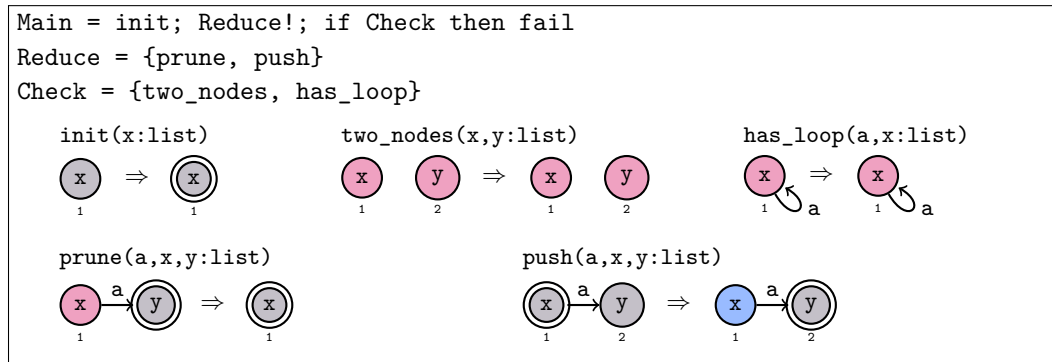
154 **► Proposition 4** (Termination of `prune!`). *`prune!` terminates after at most  $|V_G|$  steps.*

155 **Proof.** If  $G \Rightarrow H$ , then  $|V_G| > |V_H|$ . Suppose there were an infinite sequence of derivations  
 156  $G_0 \Rightarrow G_1 \Rightarrow G_2 \Rightarrow \dots$ , then there would be an infinite descending chain of natural numbers  
 157  $|V_{G_0}| > |V_{G_1}| > |V_{G_2}| > \dots$ , which contradicts the well-ordering of  $\mathbb{N}$ . The last part is  
 158 immediate since there are only  $V_G$  natural numbers less than  $V_G$ . ◀

159 **► Theorem 5** (Complexity of `is-tree-slow`). *Given an input graph of bounded degree,  
 160 `is-tree-slow` will terminate in quadratic time with respect to the number of nodes in the  
 161 input graph.*

162 **Proof.** Clearly `not_empty` and `Check` run in linear time. Unfortunately `prune` is not a fast  
 163 rule schema, and so it takes linear time to find a match. Finding a match for `prune` takes  
 164 linear time and so by Proposition 4, `prune!` terminates in quadratic time. ◀

165 Unfortunately, our program does not run in linear time due to our rules not being such  
 166 that we have constant time matching. We need to modify the program so that we can always  
 167 perform a match in constant time. Figure 2 is a refined implementation, using root nodes.  
 168 We will see that this program is not only correct, but always terminates in linear time.




■ **Figure 2** The GP 2 program *is-tree*

169 ► **Theorem 6** (Correctness of *is-tree*). *The program is-tree fulfills the tree recognition*  
 170 *specification.*

171 **Proof.** The *init* rule will fail if the input graph is empty, otherwise, it will make exactly one  
 172 node rooted. The *Reduce!* step derives the singleton discrete graph if and only if the input  
 173 was a tree (Lemmata 16 and 20). Finally, by Lemma 17, *Reduce!* cannot derive the empty  
 174 graph, so it is sufficient for *Check* to test if there is more than one node, or a loop edge. ◀

175 ► **Theorem 7** (Complexity of *is-tree*). *Given an input graph of bounded degree, is-tree*  
 176 *will terminate in linear time with respect to the number of nodes in the input graph.*

177 **Proof.** Clearly *init* and *Check* run in linear time. Since *push* and *prune* are fast rules, they  
 178 take only constant time (Theorem 1), and then by Lemma 15, *Reduce* can only be applied a  
 179 linear number of times. Thus, *Reduce!* terminates in linear time too. ◀

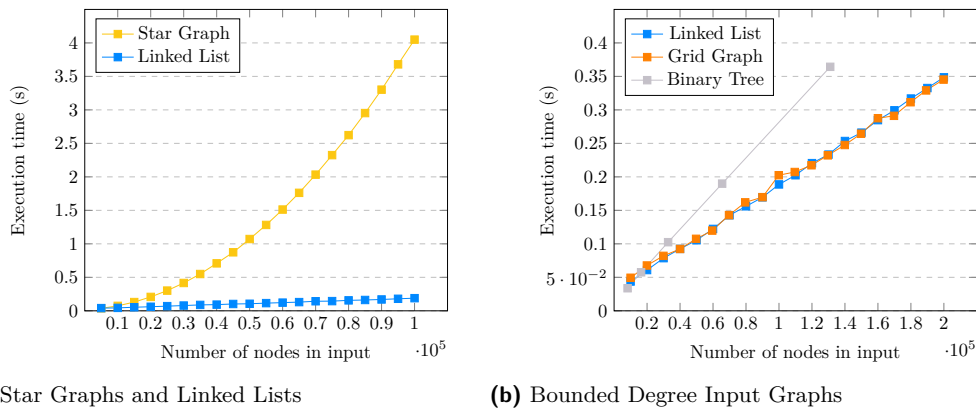


■ **Figure 3** Types of Graph

180 We have performed empirical benchmarking to verify the complexity of the program,  
 181 testing it with Linked Lists, Binary Trees, Grid Graphs, and Star Graphs (Figure 3). Star

## 23:6 Linear-Time Graph Algorithms in GP 2

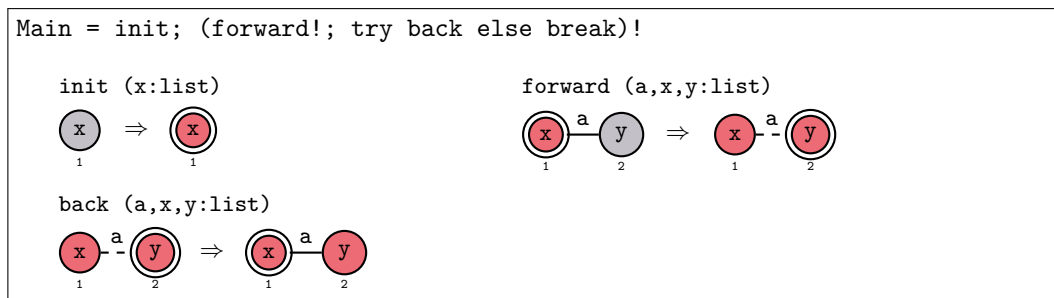
182 Graphs are not of bounded degree, so we saw quadratic time complexity as expected. The  
 183 other graphs are of bounded degree, thus we observed linear time complexity (Figure 4).



■ **Figure 4** Measured performance of `is-tree`

## 184 4 Implementing Depth First Search

185 The depth first search (DFS) seen in Figure 5 is based on the graph traversal done during  
 186 the GP 2 2-colouring program [6]. Due to the nature of GP 2, it differs from commonly used  
 187 implementation.




■ **Figure 5** The GP 2 program `general-dfs`

188 Specifically, standard implementations of DFS [7, 19] loop over the nodes of the input  
 189 graph, and if the current node has not been visited yet, a recursive function is called on it.  
 190 Said function marks the node it is called on as visited, and then calls itself on an unvisited  
 191 neighbour on an outgoing edge. If a program does not loop over all the nodes, and just  
 192 applies the recursive function to some node  $v$ , only the nodes reachable from  $v$  are visited,  
 193 since only outgoing edges are considered at each step. In GP 2 however, there is no obvious  
 194 linear-time way to loop over all nodes. So instead, the GP 2 program does not require the  
 195 edges to be outgoing, and applies a command sequence analogous to the recursive function  
 196 on an arbitrary node.

197 This approach implements DFS in an undirected graph, but not DFS in a directed  
 198 graph. So `general-dfs` is guaranteed to visit all nodes of the input graph in linear time,  
 199 but not necessarily in the order one might expect from a DFS in a directed graph. A DFS  
 200 implementation that visits the nodes of the input graph in the expected order can be found  
 201 at the core of the GP 2 program in Section 6.

202 The program starts by rooting and marking an arbitrary node red. **forward!** moves the  
 203 root along a path through the unvisited grey nodes until it reaches a node with no unvisited  
 204 neighbours. This path is marked with dashed edges. Then **back** is applied, bringing the root  
 205 one step back in the marked path. The program loops until **back** cannot be applied, which  
 206 is when there is no marked path left, i.e. the root is back at the initially rooted node.




■ **Figure 6** Example execution of `general-dfs`

207 ► **Theorem 8** (Correctness and Complexity of `general-dfs`). *Given a connected input graph*  
 208 *G of bounded degree with grey unrooted nodes and unmarked edges, `general-dfs` marks all*  
 209 *nodes red in linear time.*

210 **Proof.** Correctness follows from Lemma 22 and complexity from Corollary 24. They are  
 211 partially adapted from Bak’s [4] proofs of correctness and complexity for a different DFS  
 212 program. ◀

## 213 5 Recognising Binary DAGs

214 A *directed acyclic graph* (DAG) is a graph containing no directed cycles. A DAG is *binary* if  
 215 each of its nodes has an outdegree of at most two.



■ **Figure 7** The GP2 Program `is-bin-dag`

216 The idea behind recognising connected binary DAGs is as follows. First, all indegree-0  
 217 nodes of the input graph are identified. Then, if any indegree-0 nodes have been found, one  
 218 of them is deleted, and all of its children that become a new indegree-0 node get designated  
 219 as such. This is repeated until no indegree-0 nodes are left. Every time an indegree-0 node is  
 220 checked, the number of its children are checked as well. If there are any leftover nodes (i.e.  
 221 nodes that never had indegree-0 in the execution), then there were no directed cycles, and  
 222 the input graph is a DAG.

223 ► **Theorem 9** (Correctness of `is-bin-dag`). *The program `is-bin-dag` fulfills the following*  
 224 *specification.*

- 225 ■ Input: A connected graph *G* with grey unrooted nodes and unmarked edges.
- 226 ■ Output: The empty graph if *G* minus the blue edges was a binary DAG, and failure  
 227 otherwise.

228 **Proof.** If *G* is the empty graph, `SearchIndeg0Nodes` fails by Proposition 10, anything does  
 229 not match, and the output is the empty graph. So assume *G* is nonempty.

230 If  $G$  has no indegree-0 nodes, `nonempty_stack` will not match, and `fail` will be invoked.  
 231 So assume  $G$  has indegree-0 nodes.

232 Then Propositions 10 and 11 can be applied to deduce the following. `SearchIndeg0Nodes`  
 233 succeeds, `nonempty_stack` matches, then `ReduceIndeg0Nodes` gets applied. If  $G$  is a binary  
 234 DAG, the host graph becomes the empty graph, `anything` will not match, and the output  
 235 is the empty graph. If  $G$  is not a binary DAG, there's failure, or a nonempty graph which  
 236 results in failure since `anything` is matched. ◀

### 237 5.1 Correctness of Procedures

238 The proof of Theorem 9 depends upon the correctness of the procedures `SearchIndeg0Nodes`  
 239 and `ReduceIndeg0Nodes`. We will now give their definitions and prove their correctness.

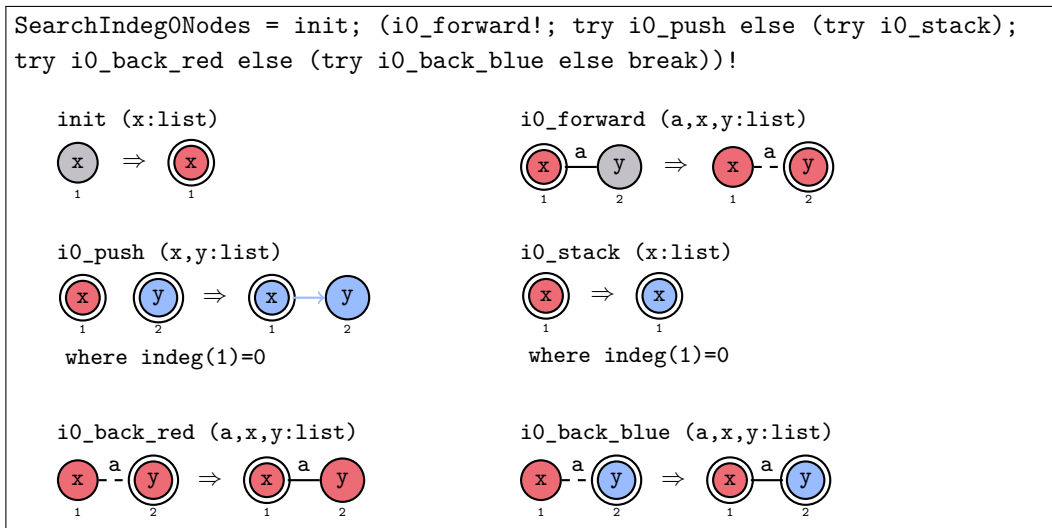
240 ▶ **Proposition 10** (Correctness of `SearchIndeg0Nodes`). *The procedure `SearchIndeg0Nodes`*  
 241 *fulfills the following specification.*

- 242 ■ Input: A connected graph  $G$  with grey unrooted nodes and unmarked edges.
- 243 ■ Output: If  $G$  is the empty graph, then failure. Otherwise  $G$  with red non-indegree-0 nodes  
 244 containing at most one root, and blue indegree-0 nodes that are connected with blue edges  
 245 forming a path graph (i.e. a linked list). The blue node with no incoming blue edges is a  
 246 root.

247 **Proof.** If  $G$  is empty, `init` cannot match, causing failure. Otherwise, the output conditions  
 248 are satisfied by Lemmata 26 and 27. ◀

249 The absence of a red root in the output is an edge case caused by `init` being applied to  
 250 an indegree-0 node. Because then, either `i0_stack` or `i0_push` will be the last rule that is  
 251 applied, and the red root becomes a blue root.

252 The blue nodes linked with blue edges are a GP 2 implementation of stacks. The top of  
 253 the stack is the only blue root, making it accessible in constant time.




■ **Figure 8** The `SearchIndeg0Nodes` procedure

254 `SearchIndeg0Nodes`, as seen in Figure 8, is based on the DFS implementation from  
 255 Section 4, with a few key differences. Using DFS ensures that each node is visited.



256 Between the forward and back steps lies the command sequence `try i0_push else (try`  
 257 `i0_stack)`. Its purpose is to push the node currently visited by the DFS if it has indegree-0.  
 258 If the stack is nonexistent, there are no blue nodes, and `i0_push` fails. So the program tries  
 259 to apply `i0_stack`, turning the node into the initial stack element (if its indegree is 0). After  
 260 the stack has been created, `i0_push` will always be applicable for indegree-0 nodes.

261 Since the current node may be marked blue by the stack operations after the previous  
 262 command sequence has been executed, the back step needs to account for that. Hence the  
 263 program first tries to apply the back rule from the previous DFS program, and if that fails,  
 264 it tries to apply an alternate version considering a blue current node. In the latter case, the  
 265 blue node is rooted since we want to keep accessing the top of the stack in constant time.



■ **Figure 9** Example execution of `SearchIndeg0Nodes`

266 ► **Proposition 11** (Correctness of `ReduceIndeg0Nodes`). *The procedure `ReduceIndeg0Nodes`*  
 267 *fulfills the following specification.*

- 268 ■ *Input: A connected graph  $G$  with red non-indegree-0 nodes containing at most one root,*  
 269 *and blue indegree-0 nodes that are connected with blue edges forming a path graph. The*  
 270 *blue node with no incoming blue edges is a root.*
- 271 ■ *Output: The empty graph if  $G$  minus the blue edges was a binary DAG, and a nonempty*  
 272 *one or failure otherwise.*

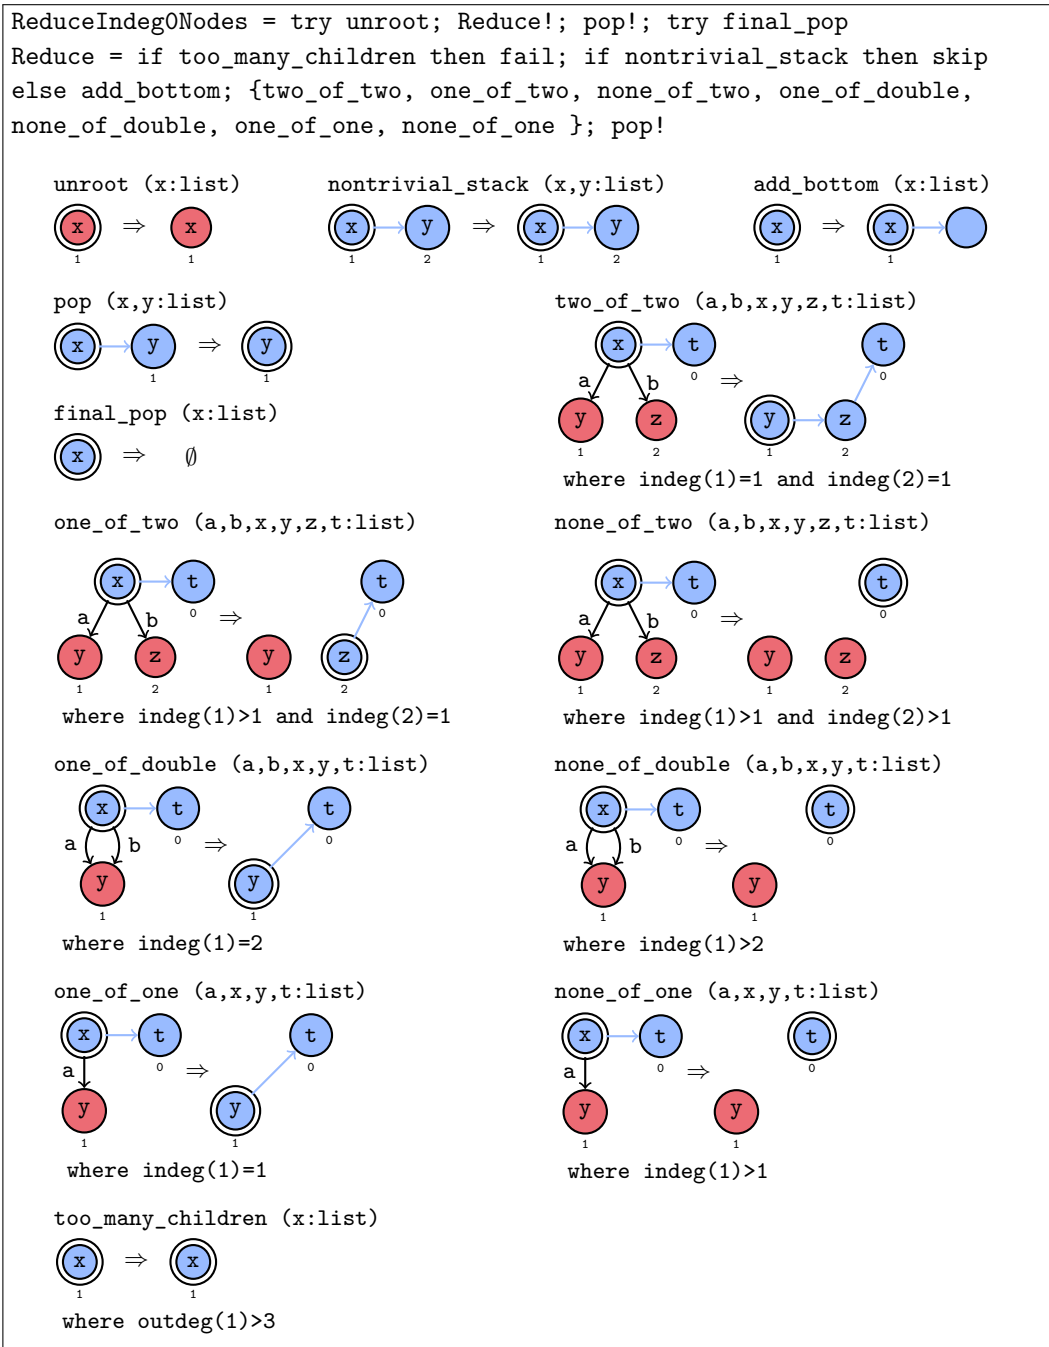
273 **Proof.** This result follows directly from Lemmata 31 and 32. ◀

274 The procedure starts by trying to apply `unroot` to get rid of any red roots left over by  
 275 `SearchIndeg0Nodes`. Then it enters the loop `Reduce!`. The blue root in each iteration shall  
 276 be called the “top root”. First, the program checks whether the top root has more than two  
 277 children, i.e. whether its outdegree is greater than three, since the blue stack edge needs to  
 278 be taken into account. If there are too many, the `fail` statement is invoked.


279 `nontrivial_stack` checks whether the stack has more than one element. If it does not,  
 280 `add_bottom` artificially adds a node to the bottom of the stack, in order for the following  
 281 rules to still match.

282 Next is a non-deterministic choice of rules that cover every case of the number of children  
 283 the top root has, and how many of those are indegree-0 nodes. In each case, they pop the  
 284 top root, and push the children that would have indegree 0 after the deletion. `pop!` serves  
 285 to pop childless indegree-0 nodes for as long as there are any.

23:10 Linear-Time Graph Algorithms in GP 2



■ **Figure 10** The ReduceIndeg0Nodes procedure



■ **Figure 11** Example execution of ReduceIndeg0Nodes


286 **5.2 Performance**

287 We will show that our binary DAG recognition program always terminates in linear time,  
 288 given an input graph of bounded degree. We have also included empirical evidence for this.

289 ► **Theorem 12** (Complexity of `is-bin-dag`). *Given a connected input graph of bounded*  
 290 *degree, the program `is-bin-dag` terminates in linear time.*

291 **Proof.** The Main procedure of `is-bin-dag` contains no loops. `SearchIndeg0Nodes` and  
 292 `ReduceIndeg0Nodes` terminate in linear time by Lemmata 25 and 33. `nonempty_stack`  
 293 matches in constant time by Theorem 1 since it is a fast rule schema. `anything` also matches  
 294 in constant time since any node is a valid match. ◀


295 In order to support the linear time complexity of `is-bin-dag`, performance will be  
 296 measured on two graph classes, one consisting of binary DAGs, and the other of non-DAGs.



■ **Figure 12** Input Graph Classes

297 Consider the following class of binary DAGs. For  $n \geq 1$ , the *grid chain*  $GC(n)$  consists of  
 298  $n$  grids of size  $n \times n$ , joint by the nodes of indegree and outdegree 1 in order to form a chain.  
 299 This class was chosen for having an unbounded number of indegree-0 nodes, meaning that  
 300 the implemented stack is relatively large.

301 Now consider the following class of non-DAGs. For  $n \geq 3$ , the *sun graph*  $SG(n)$  consists  
 302 of a directed cycle of  $n$  nodes, each of which has an an additional neighbour connected by an  
 303 incoming edge. The reason for using this class is, in addition to half the nodes having indegree  
 304 0, the other half are part of the cycle, and therefore never get deleted by `ReduceIndeg0Nodes`.  
 305 This causes an unbounded amount of nodes to be left over.



■ **Figure 13** Measured performance of `is-bin-dag`

## 306 **6** Topological Sorting

307 Given a DAG  $G$ , a *topological sorting* is a total order (an antisymmetric, transitive, and  
 308 connex binary relation)  $\leq$  on  $V_G$ , the set of nodes of  $G$ , such that for each edge of source  
 309  $u$  and target  $v$ ,  $u \leq v$  (*topological property*). Topological sortings cannot exist for graphs  
 310 containing directed cycles, since there is no way to define a total order on the nodes of a cycle  
 311 such that the topological property is satisfied. Furthermore, every DAG has a topological  
 312 sorting.

313 There are two commonly used linear-time algorithms for finding a topological sorting  
 314 [19, 18]. One seeks out indegree-0 nodes, adds them to the total order, deletes them, and  
 315 repeats this process until all nodes have been added to the order. The other, which is used  
 316 as the basis for the algorithm in this paper, conducts a DFS. Upon completion of a node  
 317 in the DFS, that node is added as the new minimum element of the linear order. However,  
 318 unlike the program `general-dfs` from Section 4, the direction of the edges needs to be  
 319 respected in order to get a topological sorting in the end. Simply turning the bidirectional  
 320 edges of `general-dfs` into directed edges is not enough since that would only visit the nodes  
 321 reachable from the initially rooted node, which is not necessarily the entire input graph.  
 322 Traditional algorithms solve this problem by skipping to the next unvisited node in the data  
 323 structure representing the graph, and continuing the DFS from there. Similarly, the program  
 324 `top-sort` uses a DFS implementation with directed edges (`SortNodes`), and once it runs out  
 325 of unvisited nodes, it uses a DFS similar to `general-dfs` (`SearchUnsortedNodes`) to find a  
 326 node that has not been visited yet, and to continue the `SortNodes` DFS.

### 327 **6.1** The Program

328 We give the GP 2 implementation of topological sorting in Figure 14 and show its correctness.  
 329 We have added the restriction that the input graph must be connected since in the current  
 330 version of GP 2, there is no known way to implement a DFS that is linear-time for graphs  
 331 with an unbounded number of connected components. We have also included an example  
 332 execution of the program in Figure 15.

333 ► **Theorem 13** (Correctness of `top-sort`). *The program `top-sort` fulfills the following*  
 334 *specification.*

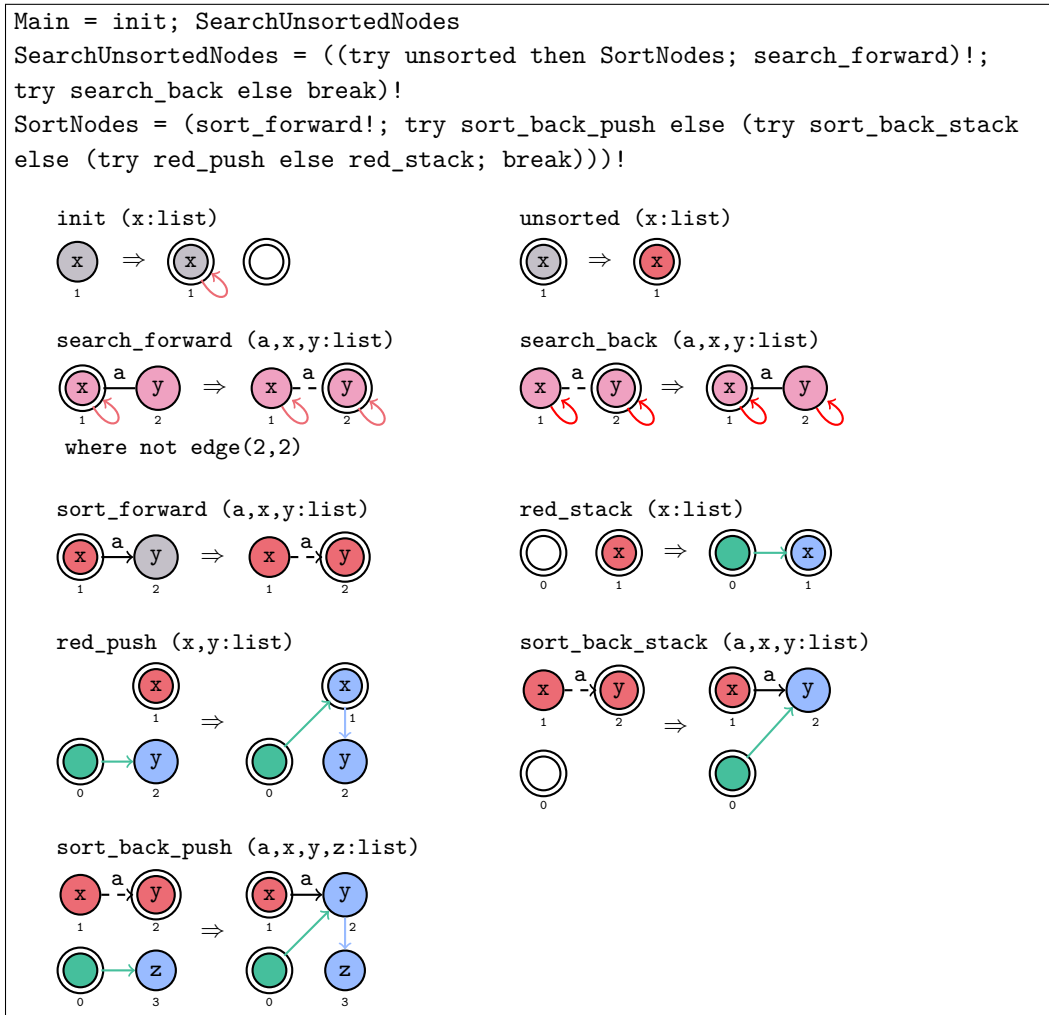
- 335 ■ *Input: A connected DAG  $G$  with no roots whose nodes are all marked grey, and whose*  
 336 *edges are unmarked.*
- 337 ■ *Output:  $G$  with additional blue edges that define a topological ordering on  $V_G$ . The nodes*  
 338 *of  $G$  are marked blue and each have a red loop. One of these nodes is rooted. Furthermore,*  
 339 *there is an additional unlabelled green root node with an outgoing green edge pointing to a*  
 340 *node with no incoming blue edges.*

341 **Proof.** This theorem follows from Lemma 37. ◀

342 The additional constructs in the output graph, apart from the blue edges, are needed for  
 343 the execution of the program. One could define a linear-time cleanup procedure to remove  
 344 these constructs. The green root and its outgoing edge can be deleted in constant time, since  
 345 access to roots is constant. Similarly, the blue rooted node can be unrooted in constant time.  
 346 A DFS similar to the one in Section 4 can be used to remove the red loops or unmark all the  
 347 nodes in linear time.

348 The subgraph induced by the blue edges is a path graph containing all the nodes from  
 349 the input graph. So the binary relation  $\leq$  on the set of nodes defined by  $u \leq v$  if there is a

350 path of blue edges from  $u$  to  $v$  is a total order, which is a necessary property for a topological  
 351 sorting. Similarly to the `SearchIndeg0Nodes` procedure of Subsection 5.1, the blue nodes  
 352 and edges implement a stack. However, this time the top of the stack is denoted with a  
 353 green root pointing towards it with a green edge in order not to interfere with a DFS in  
 354 `SortNodes`.




■ **Figure 14** The GP 2 program top-sort

355 The program starts by rooting an input node and endowing it with a red loop, as well as  
 356 creating an unmarked, unlabelled root that is disconnected from the rest of the graph. This  
 357 root will point to the top of the stack, and shall hence be called the “pointer”.

358 `SearchUnsortedNodes` is a DFS implementation similar to `general-dfs` from Section 4  
 359 that seeks out a node that has not been visited by `SortNodes` yet. Instead of using a red  
 360 mark to designate a node as visited, it uses a red loop. Since the input is assumed to be a  
 361 DAG, it has no loops. This leaves the use of marks to the DFS in `SortNodes`. So in order for  
 362 the forward step to only match unvisited neighbours of the root, a predicate to forbid loops is  
 363 needed. The “any” mark ensures that colour does not matter. Right before each application  
 364 of the forward step, `unsorted` tests whether the current root has been visited by `SortNodes`  
 365 yet, i.e. whether it is grey. At the same time, said root is initialised for `SortNodes` by being

366 marked red.

367 Next, `SortNodes` is applied. It performs a DFS with directed edges. Similarly to  
 368 `SearchIndeg0Nodes` from Section 5, it pushes the current root onto the stack during its  
 369 back step. `sort_back_push` is applied when the stack has at least one element, otherwise  
 370 `sort_back_stack` creates the stack. The pointer being green represents the stack being  
 371 nonempty. The break statement is preceded by `try red_push else red_stack`, since when  
 372 the back step can no longer be applied, the current root is still pushed onto the stack. Again,  
 373 two rules are needed to cover the cases of the stack being empty or not. Because of the  
 374 repeated application of the back step, the root ends up where it was at the beginning of  
 375 `SortNodes`, meaning that the DFS of `SearchUnsortedNodes` can resume undisturbed.



■ Figure 15 Example execution of top-sort


## 376 6.2 Performance

377 Finally, we show that, given a valid input graph of bounded degree, our topological sorting  
 378 program will always terminate in linear time.

379 ► **Theorem 14** (Complexity of top-sort). *Given a connected DAG of bounded degree with*  
 380 *only grey unrooted nodes whose edges are unmarked as an input, the program top-sort*  
 381 *terminates in linear time.*

382 **Proof.** This theorem follows from Lemma 38. ◀

383 In order to support the linear time complexity of top-sort, we make use of the grid  
 384 chains from Subsection 5.2. They are DAGs, the type of graph top-sort is meant to be used  
 385 on. Furthermore, they have an unbounded number of indegree-0 nodes. Since indegree-0  
 386 nodes are unreachable from any other node, and SortNodes can only visit nodes reachable  
 387 from the red root it is called on, SortNodes will have to be applied at least once for each  
 388 indegree-0 node, i.e. an unbounded number of times. Thus these input graphs can adequately  
 389 illustrate the linearity of top-sort. Figure 16 is a plot of the program timings, demonstrating  
 390 linear time complexity.



■ **Figure 16** Measured performance of `top-sort` on grid chains

## 7 Conclusion

391

392 The polynomial cost of graph matching is the performance bottleneck for languages based on  
 393 standard graph transformation rules. GP 2 mitigates this problem by providing rooted rules  
 394 which under mild conditions can be matched in constant time. We presented rooted GP 2  
 395 programs for three graph algorithms: tree recognition, connected binary DAG recognition,  
 396 and topological sorting. The programs were proved to be correct and to run in linear time  
 397 on graphs of bounded node degree. The proofs demonstrate that graph transformation  
 398 rules provide a convenient and intuitive abstraction level for formal reasoning on graph  
 399 programs. We also gave empirical evidence for the linear run time of the programs, by  
 400 presenting benchmark results for graphs of up to 100,000 nodes in various graph classes. For  
 401 DAG recognition and topological sorting, the linear behaviour was achieved by implementing  
 402 depth-first search strategies based on an encoding of stacks in graphs.

403 In future work, we intend to investigate for more graph algorithms whether and under  
 404 what conditions their time complexity in conventional programming languages can be reached  
 405 in GP 2. The more involved the data structures of those algorithms are, the more challenging  
 406 will be the implementation task. This is because in GP 2, the internal graph data structure  
 407 is (intentionally) hidden from the programmer and hence any data structures used by an  
 408 algorithm need to be encoded in host graphs. A simple example for this is the encoding of  
 409 stacks as linked lists in the programs for DAG recognition and topological sorting.

410 The three programs in this paper and also the 2-colouring program of [6] need host graphs  
 411 of bounded node degree in order to run in linear time. A topic for future work is therefore to  
 412 find a mechanism that allows to overcome this restriction. Clearly, such a mechanism will  
 413 require to modify GP 2 and its implementation.

## References

414

- 415 1 Aditya Agrawal, Gabor Karsai, Sandeep Neema, Feng Shi, and Attila Vizhanyo. The design  
 416 of a language for model transformations. *Software and System Modeling*, 5(3):261–288, 2006.  
 417 doi:10.1007/s10270-006-0027-7.
- 418 2 Alfred V. Aho, John E. Hopcroft, and Jeffrey D. Ullman. *The Design and Analysis of Computer*  
 419 *Algorithms*. Addison-Wesley, 1974.
- 420 3 Thorsten Arendt, Enrico Biermann, Stefan Jurack, Christian Krause, and Gabriele Taentzer.  
 421 Henshin: Advanced concepts and tools for in-place EMF model transformations. In *Model*

- 422 *Driven Engineering Languages and Systems (MODELS 2010)*, volume 6394 of *Lecture Notes*  
423 *in Computer Science*, pages 121–135. Springer, 2010. doi:10.1007/978-3-642-16145-2\_9.
- 424 **4** Christopher Bak. *GP 2: Efficient Implementation of a Graph Programming Language*. PhD  
425 thesis, Department of Computer Science, University of York, 2015. URL: [http://etheses.](http://etheses.whiterose.ac.uk/12586/)  
426 [whiterose.ac.uk/12586/](http://etheses.whiterose.ac.uk/12586/).
- 427 **5** Christopher Bak and Detlef Plump. Rooted graph programs. In *Proc. International Workshop*  
428 *on Graph Based Tools (GraBaTs 2012)*, volume 54 of *Electronic Communications of the EASST*,  
429 2012. doi:10.14279/tuj.eceasst.54.780.
- 430 **6** Christopher Bak and Detlef Plump. Compiling graph programs to C. In *Proc. International*  
431 *Conference on Graph Transformation (ICGT 2016)*, volume 9761 of *LNCS*, pages 102–117.  
432 Springer, 2016. doi:10.1007/978-3-319-40530-8\_7.
- 433 **7** Thomas H. Cormen, Charles E. Leiserson, Robert L. Rivest, and Clifford Stein. *Introduction*  
434 *to Algorithms*. The MIT Press, third edition, 2009.
- 435 **8** Hartmut Ehrig, Claudia Ermel, Ulrike Golas, and Frank Hermann. *Graph and Model*  
436 *Transformation*. Monographs in Theoretical Computer Science. Springer, 2015. doi:  
437 10.1007/978-3-662-47980-3.
- 438 **9** Maribel Fernández, Hélène Kirchner, Ian Mackie, and Bruno Pinaud. Visual modelling of  
439 complex systems: Towards an abstract machine for PORGY. In *Proc. Computability in Europe*  
440 *(CiE 2014)*, volume 8493 of *Lecture Notes in Computer Science*, pages 183–193. Springer, 2014.  
441 doi:10.1007/978-3-319-08019-2\_19.
- 442 **10** Amir Hossein Ghamarian, Maarten de Mol, Arend Rensink, Eduardo Zambon, and Maria  
443 Zimakova. Modelling and analysis using GROOVE. *International Journal on Software Tools*  
444 *for Technology Transfer*, 14(1):15–40, 2012. doi:10.1007/s10009-011-0186-x.
- 445 **11** Ivaylo Hristakiev and Detlef Plump. Checking graph programs for confluence. In *Software*  
446 *Technologies: Applications and Foundations – STAF 2017 Collocated Workshops, Revised*  
447 *Selected Papers*, volume 10748 of *Lecture Notes in Computer Science*, pages 92–108. Springer,  
448 2018. doi:10.1007/978-3-319-74730-9\_8.
- 449 **12** Edgar Jakumeit, Sebastian Buchwald, and Moritz Kroll. GrGen.NET - the expressive,  
450 convenient and fast graph rewrite system. *International Journal on Software Tools for*  
451 *Technology Transfer*, 12(3–4):263–271, 2010. doi:10.1007/s10009-010-0148-8.
- 452 **13** Detlef Plump. The design of GP 2. In *Proc. Workshop on Reduction Strategies in Rewriting*  
453 *and Programming (WRS 2011)*, volume 82 of *Electronic Proceedings in Theoretical Computer*  
454 *Science*, pages 1–16, 2012. doi:10.4204/EPTCS.82.1.
- 455 **14** Detlef Plump. From imperative to rule-based graph programs. *Journal of Logical and Algebraic*  
456 *Methods in Programming*, 88:154–173, 2017. doi:10.1016/j.jlamp.2016.12.001.
- 457 **15** Christopher M. Poskitt and Detlef Plump. Hoare-style verification of graph programs. *Funda-*  
458 *menta Informaticae*, 118(1-2):135–175, 2012. doi:10.3233/FI-2012-708.
- 459 **16** Christopher M. Poskitt and Detlef Plump. Verifying monadic second-order properties of  
460 graph programs. In *Proc. International Conference on Graph Transformation (ICGT 2014)*,  
461 volume 8571 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2014. doi:  
462 10.1007/978-3-319-09108-2\_3.
- 463 **17** Olga Runge, Claudia Ermel, and Gabriele Taentzer. AGG 2.0 — new features for specifying  
464 and analyzing algebraic graph transformations. In *Proc. Applications of Graph Transformations*  
465 *with Industrial Relevance (AGTIVE 2011)*, volume 7233 of *Lecture Notes in Computer Science*,  
466 pages 81–88. Springer, 2012. doi:10.1007/978-3-642-34176-2\_8.
- 467 **18** Robert Sedgewick. *Algorithms in C. Part 5: Graph Algorithms*. Addison-Wesley, third edition,  
468 2002.
- 469 **19** Steven Skiena. *The Algorithm Design Manual*. Springer, second edition, 2008. doi:10.1007/  
470 978-1-84800-070-4.



## A Appendix: Proofs

472 This appendix consists of lemmata and proofs omitted from the main sections.

### 473 A.1 Tree Recognition Lemmata

474 By *rooted input graph*, we mean an arbitrary labelled GP 2 input graph with every node  
 475 coloured grey, exactly one *root* node, and no additional marks. That is, a valid input graph  
 476 after `init` has been applied. By *rooted input tree*, we mean an rooted input graph that is  
 477 a non-empty tree. In this appendix, we give the proofs of the lemmata needed to support  
 478 Theorems 6 and 7 from Section 3.

479 ► **Lemma 15.** *Reduce! terminates after at most  $2|V_G|$  steps.*

480 **Proof.** Let  $\#G$  be the number of nodes, and  $\square G$  be the number of grey nodes. If  $G \Rightarrow_{prune} H$ ,  
 481 then  $\#G > \#H$  and  $\square G > \square H$ . If  $G \Rightarrow_{push} H$  then  $\#G = \#H$  and  $\square G > \square H$ . Suppose  
 482 there were an infinite sequence of derivations  $G_0 \Rightarrow G_1 \Rightarrow G_2 \Rightarrow \dots$ , then there would be an  
 483 infinite descending chain of natural numbers  $\#G_0 + \square G_0 > \#G_1 + \square G_1 > \#G_2 + \square G_2 > \dots$ ,  
 484 which contradicts the well-ordering of  $\mathbb{N}$ . To see the last part, notice that  $\square G \leq \#G$  for all  
 485 graphs  $G$ , so the result is immediate since there are only  $2\#G$  natural numbers less than  
 486  $2\#G$ . ◀

487 ► **Lemma 16.** *If  $G$  is a tree and  $G \Rightarrow_{Reduce}^* H$ , then  $H$  is a tree. If  $G$  is not a tree and  
 488  $G \Rightarrow_{Reduce!} H$ , then  $H$  is not a tree.*

489 **Proof.** Clearly, the application of `push` preserves structure. Suppose  $G$  is a tree. `prune` is  
 490 applicable if and only the second node is matched against a leaf node, due to the dangling  
 491 condition. Upon application, the leaf node and its incoming edge is removed. Clearly the  
 492 result graph is still a tree. If  $G$  is not a tree and `prune` is applicable, then we can see the  
 493 properties of not being a tree are preserved. That is, if  $G$  is not connected,  $H$  is certainly  
 494 not connected. If  $G$  had parallel edges, due to the dangling condition, they must exist in  
 495  $G \setminus g(L)$ , so  $H$  has parallel edges. Similarly, cycles are preserved. Finally, if  $G$  had a node  
 496 with incoming degree greater than one, then  $H$  must too, since the node in  $G$  that is deleted  
 497 in  $H$  had incoming degree one, and the degree of all other nodes is preserved. So, we have  
 498 shown `Reduce` is structure preserving, and then by induction, so is `Reduce!`. ◀

499 ► **Lemma 17.** *If  $G$  is a rooted input graph and  $G \Rightarrow_{Reduce}^* H$ , then  $H$  has exactly one root  
 500 node. Moreover, there is no derivation sequence that derives the empty graph.*

501 **Proof.** In each application of `prune` or `push`, the number of root nodes is invariant since  
 502 the LHS of each rule must be matched against a root node in the host graph, so the other  
 503 non-roots can only be matched against non-roots, and so the result holds by induction. To  
 504 see that the empty graph cannot be derived, notice that each derivation reduces  $\#G$  by at  
 505 most one, and no rules are applicable when  $\#G = 1$ . ◀

506 ► **Lemma 18.** *If  $G$  is a rooted input graph and  $G \Rightarrow_{Reduce}^* H$ . Then, every blue node in  $H$   
 507 either has a blue child or a root-node child.*

508 **Proof.** Clearly  $G$  satisfies this, as there are no blue nodes. We now proceed by induction.  
 509 Suppose  $G \Rightarrow_{Reduce}^* H \Rightarrow_{Reduce} H'$  where  $H$  satisfies the condition. If `prune` is applicable,  
 510 we introduce no new blue nodes. Additionally, any blue parents of the node 1 are preserved.  
 511 So  $H'$  satisfies the condition. Finally, if `push` is applied, then the new blue node has a

## 23:18 Linear-Time Graph Algorithms in GP 2

512 root-node child, and the blue nodes in  $H' \setminus h(R)$  have the same children, so  $H'$  satisfies the  
513 condition. ◀

514 ▶ **Corollary 19.** *Let  $G$  be a rooted input tree and  $G \Rightarrow_{Reduce}^* H$ . Then the root-node in  $H$   
515 has no blue children.*

516 **Proof.** By Lemma 17,  $H$  has exactly one root node, and by Lemma 18, all chains of blue  
517 nodes terminate with a root-node. If said root-node were to have a blue child, then we would  
518 have a cycle, which contradicts that  $H$  is a tree (Lemma 16). ◀

519 ▶ **Lemma 20.** *Let  $G$  be a rooted input tree and  $G \Rightarrow_{Reduce}^* H$ . Then, either  $|V_H| = 1$  or  $H$   
520 is not in normal form.*

521 **Proof.** By Lemma 17,  $|V_H| \geq 1$ . If  $|V_G| = 1$ , then  $G$  is in normal form. Otherwise, either  
522 the root node has no children, or it has at least one grey child. In the first case, **prune** must  
523 be applicable, and in the second, **push**. Suppose  $G \Rightarrow_{Reduce}^* H$ . If  $|V_H| = 1$ , then  $H$  is in  
524 normal form by the proof to Lemma 17. Otherwise, by Lemma 16  $H$  is a tree and  $|V_H| > 1$ .  
525 Now, the root-node in  $H$  (Lemma 17) must have a non-empty neighbourhood. If it has  
526 no children, then **prune** must be applicable. Otherwise, **push** must be applicable, since by  
527 Corollary 19, there must be a grey node child. So  $H$  is not in normal form. ◀

## 528 A.2 DFS Lemmata

529 In this appendix, we give the proofs of the lemmata needed to support Theorem 8 from  
530 Section 4.

531 ▶ **Lemma 21 (Termination of `general-dfs`).** *Given an input graph  $G$  with grey nodes and  
532 unmarked edges, `general-dfs` terminates.*

533 **Proof.** Consider the following (total) lexicographical ordering  $>$  on graphs, defined as  
534  $H_1 > H_2$  if  $H_1$  has more grey nodes than  $H_2$ , or  $H_1$  and  $H_2$  have the same number of grey  
535 nodes and  $H_1$  has more dashed edges than  $H_2$ . If **forward** is applied on a graph  $H_1$ , yielding  
536  $H_2$ , then  $H_1 > H_2$  since the rule strictly decreases the number of grey nodes. In particular,  
537 **forward!** terminates since eventually, there is no possible match for the left hand side of  
538 **forward**. Similarly, applying **back** on a graph  $H_1$  to obtain  $H_2$  strictly decreases the number  
539 of dashed edges, and keeps the number of grey nodes constant. So  $H_1 > H_2$ . Since  $G$  has  
540 only grey nodes and unmarked edges, there are at most  $|V_G| \cdot |E_G|$  (i.e. finitely many) graphs  
541  $H$  such that  $G > H$ . So at some point, neither **forward**, nor **back** are applicable. Since **back**  
542 is not applicable, the **break** statement is invoked, causing the program to exit the loop. ◀

543 ▶ **Lemma 22 (Correctness of `general-dfs`).** *Given a connected input graph  $G$  with grey  
544 unrooted nodes and unmarked edges, the subgraph  $H$  induced by the edges that have been dashed  
545 during the execution of `general-dfs` is a spanning tree, all of whose nodes are marked red.*

546 **Proof.**  $H$  is a tree if and only if it is connected and has  $|V_H| - 1$  edges. This property shall  
547 be used to show that  $H$  is a tree. The program starts by applying **init**, which roots an  
548 input node  $v$  and marks it red. If  $G$  only has one node, the program terminates and  $H$   
549 is the empty graph, which is a tree. Otherwise, **forward** gets applied to a node  $w$  (since  $G$  is  
550 connected) and dashes the edge between  $v$  and  $w$ , making them part of  $H$ . So  $H$  has at  
551 least two nodes and one edge. The only rule dashing edges is **forward**, which, apart from  
552 **init**, is the only rule marking grey nodes red. So the nodes of  $H$  are exactly the nodes  
553 that are marked red during execution. **forward** matches a red node, a grey node, and the

554 edge between them, ensuring they are all part of  $H$ . Since only the red node was part of  $H$   
 555 already, exactly one node and one edge are added to  $H$ . Hence, if **forward** is applied  $n \geq 0$   
 556 times, we have  $|V_H| = 2 + n$  and  $|E_H| = 1 + n = |V_H| - 1$ . Since **forward** adds an edge and  
 557 a node to what is known to be part of  $H$  already, starting from a connected graph,  $H$  is  
 558 connected.

559 It remains to show that  $H$  is a spanning tree. If  $|V_G| \leq 2$ , applying **init** followed by  
 560 **forward** ensures  $H$  being a spanning tree. Otherwise, since  $H$  is a subgraph of  $G$  as well as  
 561 a tree, it can be extended to a spanning tree  $T$ . Assume for the sake of a contradiction that  
 562  $G \setminus H$  is nonempty. Let  $w$  be a (grey) node of  $G \setminus H$  adjacent to a (red) node  $u$  of  $H$ . It  
 563 exists since  $G$  is connected. Each leaf  $v$  of  $H$  marks the termination of **forward!** (which  
 564 terminates by Lemma 21), because if **forward** could have been applied again, it would have,  
 565 and  $v$  would have two adjacent dashed edges. If  $u$  is a leaf, **forward** would be applicable to  $u$   
 566 and  $w$ , marking  $w$  red and causing a contradiction. If  $u$  is not a leaf, after all its descendants  
 567 have been marked red, the root would be moved back to  $u$  using applications of the **back**  
 568 rule. Subsequently,  $u$  and  $w$  would have been matched by **forward**, causing a contradiction  
 569 again. ◀

570 ▶ **Corollary 23.** *Given a connected input graph  $G$  with grey unrooted nodes and unmarked*  
 571 *edges, where  $|V_G| \geq 2$ , **general-dfs** applies **forward** and **back** exactly  $|V_G| - 1$  times each.*  
 572 *For  $|V_G| < 2$ , they are not applied at all.*

573 **Proof.** If  $|V_G| < 2$ , **forward** and **back** do not have enough vertices to match.

574 Otherwise, since  $H$  as in Lemma 22 is a spanning tree, it has  $|V_G| - 1$  edges. As seen  
 575 in the proof of that Lemma,  $H$  is constructed a red root by applying **forward** a number of  
 576 times, which adds an edge and a node each time. So **forward** must have been applied at  
 577 least  $|V_G| - 1$  times. It must also have been applied at most  $|V_G| - 1$  times, since that's the  
 578 number of grey nodes after application of **init**, since it marks every added grey node red,  
 579 and since no other rule introduces grey nodes.

580 **back** matches at most  $|V_G| - 1$  times since it can only match a dashed edge. It matches  
 581 at least  $|V_G| - 1$  times since its left hand side is matchable throughout the spanning tree. ◀

582 ▶ **Corollary 24 (Complexity of **general-dfs**).** *Given a connected input graph  $G$  of bounded*  
 583 *degree with grey unrooted nodes and unmarked edges, **general-dfs** terminates in linear time.*

584 **Proof.** **init** is matched exactly once at the beginning of the program. Since the input has  
 585 grey nodes only and hence valid matches, the rule matches in constant time. The other  
 586 rules are fast rules, and hence match in constant time on graphs of bounded degree with a  
 587 bounded number of roots by Theorem 1. There is always exactly one root since the input has  
 588 none, **init** introduces one, and the other rules conserve the number of roots. By Corollary  
 589 23, **forward** and **back** are applied a number of times linear in the size of the input graph.  
 590 Since Lemma 21 guarantees the program to terminate, **general-dfs** terminates in linear  
 591 time. ◀

### 592 A.3 Binary DAG Recognition Lemmata

593 In this appendix, we give the proofs of the lemmata needed to support Propositions 10 and  
 594 11 and Theorem 12 from Section 5.

595 ▶ **Lemma 25 (Complexity and Partial Correctness of **SearchIndeg0Nodes**).** *Given a connected*  
 596 *input graph  $G$  with grey unrooted nodes and unmarked edges, **SearchIndeg0Nodes** terminates,*  
 597 *and the subgraph  $H$  induced by the edges that have been dashed during the execution is a*  
 598 *spanning tree. Furthermore, if  $G$  has bounded degree, the procedure terminates in linear time.*

599 **Proof.** Similar to the proofs in Appendix A.2. ◀

600 ▶ **Lemma 26.** *Given a nonempty connected input graph  $G$  with grey unrooted nodes and unmarked edges, at any point of the execution of `SearchIndeg0Nodes`, there is at most one red root.*

603 **Proof.** `init` introduces a red root, and is only applied once and in the beginning. The other rules that do not preserve red roots are `i0_push`, `i0_stack` and `i0_back_blue`. If either `i0_push` or `i0_stack` are applied, the red root vanishes. Subsequently, `i0_back_red` cannot be applied. If `i0_back_blue` then gets applied the red root is reintroduced, conserving the existence of a red root within the iteration of the loop. If `i0_back_blue` does not get applied, the break statement is invoked and the procedure terminates. ◀

609 ▶ **Lemma 27.** *Given a nonempty connected input graph  $G$  with grey unrooted nodes and unmarked edges, `SearchIndeg0Nodes` outputs  $G$  where all the indegree-0 nodes (and only those) are marked blue and connected with blue edges forming a path graph. The blue node with no incoming blue edge is rooted.*

613 **Proof.** If  $G$  has no indegree-0 nodes, then the lemma is trivially satisfied. So assume  $G$  has at least one.

615 By Lemma 25, `SearchIndeg0Nodes` visits all nodes. Since the right hand side of each rule only contains red and blue nodes, every node is marked either red or blue. The only rules that introduce a blue mark are `i0_push` and `i0_back_blue`, and they turn a red root into a blue root. These rules only get applied if the indegree of said red node is 0. Furthermore, the only edges introduced by `SearchIndeg0Nodes` are blue edges between two blue nodes (in `i0_push`), hence the indegree of a red node is the same as its indegree in the input graph. So only indegree-0 nodes are marked blue. Furthermore, since `SearchIndeg0Nodes` visits, i.e. roots every node of the input graph at some point, all indegree-0 nodes are marked blue, and all non-indegree-0 nodes red.

624 All rules apart from `i0_push` and `i0_stack` preserve the structure of the subgraph consisting of blue nodes and edges. `i0_stack` only is applied only if `i0_push` is not applicable. But the left hand side of `i0_push` contains a blue root, which can only be created by itself or `i0_stack`. So `i0_push` cannot be applied until `i0_stack` is applied. Since  $G$  cannot consist of only indegree-0 nodes (which would mean  $G$  is disconnected), `i0_push` can always be matched if the red root has indegree 0. If the red root does not have indegree 0, `i0_stack` cannot be matched either. So the only way for these two rules to match is for `i0_stack` to be matched first and only once, followed by `i0_push` being matched any number of times. Thus, a blue root is created, and then, repeatedly, a new blue node gets connected to the blue root with an outgoing blue edge, while the root moves to the newly added blue node. This construction results in the blue nodes and edges forming a path graph where the node with no incoming edges is a root. ◀

636 ▶ **Lemma 28 (Termination of `ReduceIndeg0Nodes`).** *Let  $G$  be a connected graph with red non-indegree-0 nodes containing at most one root, and blue indegree-0 nodes that are connected with blue edges forming a path graph. The blue node with no incoming blue edges is a root. Given a  $G$  as an input, `ReduceIndeg0Nodes` terminates.*

640 **Proof.** `pop` can only be applied a finite number of times since it reduces the number of nodes in the host graph. So `pop!` terminates.

642 *Claim:* During the execution of `ReduceIndeg0Nodes`, `add_bottom` gets applied at most twice.

644 *Proof of Claim:* Assume `add_bottom` has already been applied twice, creating the  
 645 additional nodes  $u$  and  $v$ . It only gets applied when `nontrivial_stack` cannot be matched,  
 646 i.e. when the stack consists of only one element. So it adds a blue node to the bottom of the  
 647 stack. Since every other rule that modifies the stack only does so at the top,  $u$  and  $v$  must  
 648 be consecutive and at the bottom. So right after the second application of `add_bottom`, the  
 649 stack consists of only  $u$  and  $v$ . Neither  $u$  nor  $v$  can ever have red neighbours, since there is  
 650 no rule with an edge incident to a red node in its right hand side. Hence none of the rules in  
 651 the rule set call between curly braces is applicable, causing `Reduce` to fail, and `add_bottom`  
 652 never to be applied again.

653 The rules in the rule set call and `pop` reduce the number of nodes in the host graph by  
 654 exactly one. So by the claim, they can be applied at most  $|V_G| + 2$  times each. So at some  
 655 point in the loop, they will no longer be applicable. Neither will `add_bottom` since it can  
 656 only be applied twice. So `Reduce!` terminates. ◀

657 ▶ **Lemma 29.** *Given an input graph  $G$  as described in Lemma 28, every node that has no*  
 658 *incoming unmarked edges (called quasi-indegree-0 node) in some host graph of the execution*  
 659 *of `ReduceIndeg0Nodes` gets marked blue.*

660 **Proof.** Indeed, the input graph has all quasi-indegree-0 nodes marked blue already. The  
 661 only rules deleting edges are those from the rule set call (`pop` and `final_pop` cannot delete  
 662 unmarked edges incident to the node they delete because the dangling condition needs to be  
 663 satisfied for them to match). So these are the only rules that can create new quasi-indegree-0  
 664 nodes. If one of said nodes has indegree 0, it gets detected by the condition of a rule and  
 665 marked blue. These rules cover each case of how many children their quasi-indegree-0 parent  
 666 can have in a binary DAG, namely one, one with two parallel edges, and two. The case of  
 667 no children is covered by `pop` afterwards. They also cover all cases of how many of these  
 668 children are quasi-indegree-0. So at each execution step, the newly created quasi-indegree-0  
 669 nodes get marked blue, proving this lemma. ◀

670 ▶ **Lemma 30.** *Given an input graph  $G$  as described in Lemma 28, every node that is marked*  
 671 *blue during execution of `ReduceIndeg0Nodes` is not present in the output.*

672 **Proof.** Nodes can only be marked blue if an already existing blue node is matched. So it is  
 673 enough to show that, at some point of the execution, there will be no blue nodes. There  
 674 are three potential ways to exit the loop `Reduce!`. The first is through the `fail` statement  
 675 after matching `too_many_children`. This will never happen since the input minus the blue  
 676 edges is binary, and every rule conserves the blue root having exactly one outgoing blue edge.  
 677 The second way is for `add_bottom` to fail. This can only happen when there is no blue root.  
 678 The only rule deleting a blue root is `final_pop`, which is only called after termination of  
 679 `Reduce!`. Since furthermore, the input is assumed to have a blue root, and every other rule  
 680 conserves the existence of a blue root, `add_bottom` is always applicable. The third and final  
 681 way to exit the loop is when none of the rules in the rule set call are applicable. The blue  
 682 root not having an element below it in the stack cannot be a reason for that, since in that  
 683 case, `add_bottom` would have been applied. So the current blue root  $v$  does not have red  
 684 neighbours. Since `pop!` has been applied in the previous iteration of `Reduce!`,  $v$  was the  
 685 only blue node in the previous iteration, otherwise it would have been popped. Hence in the  
 686 current iteration, `add_bottom` was applied, and so the only blue nodes are  $v$  and the node  
 687 created by `add_bottom`, say  $w$ . By Lemma 28, `Reduce!` terminates, so this always happens  
 688 for the given input. As established,  $v$  has no children. Neither does  $w$  since it was created  
 689 by `add_bottom` and there is no rule with edges incident to red nodes in its right hand side.

## 23:22 Linear-Time Graph Algorithms in GP 2

690 Thus `pop` deletes  $v$ , then `final_pop` deletes  $w$ , causing all previously blue marked nodes to  
691 be deleted. ◀

692 ▶ **Lemma 31** (Correctness of `ReduceIndeg0Nodes` for Binary DAGs). *Given an input graph  $G$   
693 as described in Lemma 28, if  $G$  minus the blue edges is a binary DAG, `ReduceIndeg0Nodes`  
694 yields the empty graph.*

695 **Proof.** Assume, for the sake of a contradiction, that the output of `ReduceIndeg0Nodes`  
696 contains a node  $v$ . By Lemmata 29 and 30,  $v$  cannot have been a quasi-indegree-0 node (i.e.  
697 an indegree-0 node when ignoring blue edges) at any point during execution. Furthermore,  $v$   
698 must have a parent that never was a quasi-indegree-0 node, because otherwise it would have  
699 been marked blue by one of the rule set call rules. The same argument can then be applied  
700 to the parent's parent, and so on indefinitely. Since the input is finite however, two of these  
701 parents must be equal, meaning that there is a cycle. This contradicts the input minus the  
702 blue edges being a DAG. ◀

703 ▶ **Lemma 32** (Correctness of `ReduceIndeg0Nodes` for Non-Binary Graphs or Non-DAGs).  
704 *Given an input graph  $G$  as described in Lemma 28, if  $G$  minus the blue edges is either not  
705 binary, or not a DAG, then `ReduceIndeg0Nodes` yields a nonempty graph.*

706 **Proof.** Assume  $G$  is not a DAG. Then it has a directed cycle consisting of consecutive  
707 nodes  $v_1, v_2, \dots, v_n$ . None of these nodes have indegree 0 ignoring blue edges, so they are  
708 never matched by the rule set call rules that would mark them blue. Since there are no  
709 rules that delete red nodes (only rules that mark them blue),  $v_1, v_2, \dots, v_n$  never get deleted.  
710 Thus the output is nonempty. Failure cannot occur since every rule and procedure of  
711 `ReduceIndeg0Nodes` is either preceded by `try` or followed by `!`.

712 Now assume that  $G$  is a DAG but is not binary. Consider an arbitrary node  $v$  of  $G$ . The  
713 aim is to show that, if  $v$  has more than two children (excluding blue edges), then the output  
714 is nonempty. By Lemma 29,  $v$  gets marked blue at some point of the execution. This can  
715 only happen in the rule set call rules. Assume  $v$  has just been marked blue by one of these  
716 rules. We can also assume that  $v$  is rooted since, by Lemma 30, every blue node gets deleted  
717 at some point, which can only happen in one of the rule set call rules or in `pop`. The case of  
718 it happening in `pop` shall be discarded since that would mean  $v$  has no children (disregarding  
719 blue edges). Back in the execution right after execution of one of the rule set call rules, since  
720 `pop!` cannot fail, the loop `Reduce!` enters its next iteration. The procedure tries to apply  
721 `too_many_children` to the blue root. If  $v$  has more than two children (disregarding blue  
722 edges), it succeeds, and the `fail` statement is invoked, terminating the loop `Reduce!`. Since  
723  $v$  has children, both `pop` and `final_pop` do not get applied, for the dangling condition is  
724 not satisfied. So the output contains  $v$  and is therefore nonempty. ◀

725 ▶ **Lemma 33** (Complexity of `ReduceIndeg0Nodes`). *Given an input graph  $G$  as described in  
726 Lemma 28 with bounded degree, `ReduceIndeg0Nodes` terminates in linear time.*

727 **Proof.** By Lemma 28, the procedure terminates. Every rule is a fast rule schema, and is  
728 hence applied in constant time by Theorem 1 (the input is assumed to have bounded degree,  
729 and from the input specification, the fact that `unroot` removes a red root if it is present, and  
730 the fact that all the other rules conserve the number of roots, there are at most two roots  
731 in the host graph at any given point of the execution). So it is enough to show that each  
732 of the constantly many rules gets applied a linear number of times. `unroot` and `final_pop`  
733 get applied at most once. By the proof of Lemma 28, `add_bottom` gets applied at most  
734 twice, and each rule set call rule as well as `pop` at most  $|V_G| + 2$  times. `too_many_children`

735 and `nontrivial_stack` can only get reapplied if the rule set call does not fail, which can  
 736 only happen at most  $|V_G| + 2$  times. Hence `ReduceIndeg0Nodes` terminates in linear time.  
 737 terminates in linear time. ◀

#### 738 A.4 Topological Sorting Lemmata

739 In this appendix, we give the proofs of the lemmata needed to support Theorems 13 and 14  
 740 from Section 6.

741 ▶ **Lemma 34** (Termination of `top-sort`). *Given a connected DAG  $G$  with no roots, grey*  
 742 *nodes, and unmarked edges as an input, `top-sort` terminates.*

743 **Proof.** `sort_forward!` terminates since in each iteration, the number of grey nodes de-  
 744 creases.

745 For the termination of `SortNodes`, consider the following lexicographical ordering  $>$ .  
 746  $H_1 > H_2$  if one of the following three statements are satisfied.  $H_1$  has more grey nodes than  
 747  $H_2$ , or they have the same number of grey nodes but  $H_1$  has more dashed edges, or they have  
 748 the same number of grey nodes and dashed edges but  $H_1$  has more red nodes. Let  $H_1$  be the  
 749 input of an arbitrary iteration of `SortNodes`, and  $H_2$  its output. If `sort_forward` is applied  
 750 any number of times,  $H_1 > H_2$  since the number of grey nodes are reduced. Otherwise, if  
 751 either `sort_back_push` or `sort_back_stack` is applied,  $H_1 > H_2$  since the number of grey  
 752 nodes is conserved and the number of dashed edges decreases in both rules. Otherwise, either  
 753 `red_push` or `red_stack` have to be applied, which conserve the number of grey nodes and  
 754 dashed edges, but decreases the number of red nodes. So in any case,  $H_1 > H_2$ . For a given  
 755 graph  $H_1$  consider how many graphs  $H_2$  satisfy  $H_1 > H_2$ . By definition of  $>$ ,  $H_1$  gives a  
 756 (finite) upper bound on the number of grey nodes, dashed edges, and red nodes. Hence there  
 757 are only finitely many possible  $H_2$ s. Since `sort_forward!` terminates, and each iteration of  
 758 the loop reduces the host graph with respect to  $<$ , `SortNodes` terminates.

759 Consider `(try unsorted then SortNodes; search_forward)!`. If `search_forward`  
 760 cannot be applied, the loop terminates. It is the only rule in this loop that increases the  
 761 number of looped edges in the graph. Due to its predicate, it can only add looped edge to a  
 762 node if it does not already have one. Furthermore, no rule decreases the number of looped  
 763 edges. So for an arbitrary input graph  $H$  for the loop, at most  $|V_H|$  looped edges can be  
 764 added before `search_forward` fails. Hence the loop terminates (knowing that `SortNodes`  
 765 also terminates).

766 Finally, consider the loop that `SearchUnsortedNodes` consists of. Furthermore, consider  
 767 the lexicographic ordering  $>$  defined by  $H_1 > H_2$  if  $H_2$  has more nodes with looped edges  
 768 than  $H_1$ , or they have the same number of nodes with looped edges but  $H_2$  has less dashed  
 769 edges than  $H_1$ . By an argument similar to that of Lemma 21, `SearchUnsortedNodes`  
 770 terminates. ◀

771 For the correctness of `SortNodes`, the following concepts needs to be defined. In a graph  
 772  $G$ , a *directed path* from a node  $v$  to a node  $w$  is a sequence of distinct nodes  $v_1, v_2, \dots, v_n$   
 773 such that  $v_1 = v$  and  $v_n = w$ , and for each  $i$  where  $1 \leq i \leq n - 1$ , there is an edge of source  
 774  $v_i$  and of target  $v_{i+1}$ .

775 A directed path from  $v$  to  $w$  is called *grey-noded* if all the nodes it consists of, except  
 776 possibly  $v$ , are marked grey.

Given a node  $v$  in a DAG  $G$ , let its *descendants*  $\text{Desc}_G(v)$  be defined as the subgraph of  
 $G$  induced by the

$$\{w \in V_G \mid \text{there is a directed grey-noded path from } v \text{ to } w\} \cup \{v\}.$$

777 ► **Lemma 35** (Correctness of `SortNodes`). *Assume the input graph of `top-sort` has no blue*  
 778 *edges. Let  $G$  be a connected DAG with a single red root  $v$ , where the nodes of  $\text{Desc}_G(v)$  are*  
 779 *unrooted. Furthermore, let  $G$  have an additional root that is either unmarked and disconnected,*  
 780 *or green and connected to the rest of the graph with an outgoing green edge. Let  $H$  be the*  
 781 *output of `SortNodes` applied on  $G$ . Consider the binary relation  $\leq$  on nodes of  $\text{Desc}_H(v)$*   
 782 *defined by  $u \leq w$  if there is a directed path from  $u$  to  $w$ , or if  $u = w$ , such that all of the*  
 783 *involved edges are blue. Then  $\leq$  defines a topological sorting on  $\text{Desc}_H(v)$  minus the blue*  
 784 *edges.*

785 **Proof.** Since the input graph of `top-sort` has no blue edges, any that are present in the  
 786 host graph were created by rules. Whenever these rules create blue edges, they mark the  
 787 incident nodes blue. No rule removes a blue mark, so the subgraph of the host graph induced  
 788 by the blue edges always exclusively consists of blue nodes. Furthermore, every time a node  
 789 gets marked blue, the green root points towards it. And when a new blue edge gets created,  
 790 the target node must also have the green root pointing towards it, and the source node must  
 791 be a red root. So the procedure only adds a blue edge from a non-blue to the node that has  
 792 most recently been marked blue. From this construction, we can infer that the graph induced  
 793 by the blue edges is a path graph. Furthermore, no blue looped edges are introduced. So  
 794 there can be no path from a node  $u$  to a node  $w$  and vice versa. Hence if  $u \leq w$  and  $w \leq u$ ,  
 795  $u$  and  $w$  must be equal by definition of  $\leq$ , i.e.  $\leq$  is antisymmetric.

796 From the definition of  $\leq$ , it is clear that transitivity holds due to path concatenation  
 797 resulting in paths.

798 With a proof similar to that of Lemma 22, one can show that `SortNodes` turns every  
 799 node of  $\text{Desc}_G(v)$  into a red root. Furthermore, all the red roots become blue nodes incident  
 800 to blue edges. So  $\leq$  is connex.

801 To show that the topological property holds, consider two nodes  $u$  and  $w$  of  $\text{Desc}_H(v)$ ,  
 802 both of which being distinct from  $v$  ( $v$  itself will be handled later). So by definition, there  
 803 is path of non-blue edges from  $v$  to  $u$ , and one from  $v$  to  $w$ . We can assume without loss  
 804 of generality that  $u$  becomes a red root before  $w$ . If there is no edge between  $u$  and  $w$ , the  
 805 topological property imposes no constraint on said pair of nodes. If there is an edge from  
 806  $u$  to  $w$ , `sort_forward` gets applied again, dashing said edge and turning  $w$  into a red root.  
 807 Hence later in the execution,  $w$  gets pushed before  $u$ , ensuring that the topological property  
 808 is satisfied. If there is an edge from  $w$  to  $u$ , there can be no non-blue path from  $u$  to  $w$  since  
 809 the input is a DAG. Hence  $u$  will be pushed before  $w$ , satisfying the topological property  
 810 again. As for  $v$ , any condition involving it must have it as the source node by definition of  
 811  $\text{Desc}_H(v)$ . Since  $v$  is pushed last, the topological property is satisfied.

812 ◀

813 ► **Lemma 36.** *Given an input  $G$  as described in Lemma 35, the output of `SortNodes` has*  
 814 *the same dashed edges, and the red root in the same place as  $G$ .*

815 **Proof.** Let  $v$  be the red root of  $G$ . During the execution of `SortNodes`, there is always a  
 816 path of dashed edges from  $v$  to the current red root, since `sort_forward` is the only rule of  
 817 `SortNodes` with dashed edges in its right hand side and generates a path graph of nodes  
 818 and dashed edges, and since `sort_back_stack` and `sort_back_push` only remove the latest  
 819 node from that path graph. The only way for their encompassing loop to end is for both of  
 820 these rules not to be applicable. By the previous argument, this means that there are no  
 821 dashed edges in said path graph left, and  $v$  is the red root when `SortNodes` terminates. ◀



822 ► **Lemma 37** (Correctness of `SearchUnsortedNodes`). *Let  $G$  be a connected acyclic graph*  
 823 *with grey nodes and unmarked edges, except for a disconnected unmarked root and a red looped*  
 824 *edge on a unique grey root. Given  $G$  as an input, `SearchUnsortedNodes` yields a graph such*  
 825 *that the ordering from Lemma 37 extended to the entire output graph is a topological sorting.*

826 **Proof.** `SearchUnsortedNodes` is similar to `general-dfs` from Section 4, but instead of read  
 827 marks, red looped edges are used. Furthermore, `try_unsorted` then `SortNodes` is added,  
 828 none of whose rules modify red looped edges. Also, after application of `SortNodes`, the red  
 829 root remains at the same place, and the same edges remain dashed. So a reasoning similar  
 830 to that in the proof of Lemma 22 can be used to justify that `SearchUnsortedNodes` visits  
 831 every node of its input graph.

832 `SearchUnsortedNodes` applies `SortNodes` to each of these visited nodes that are marked  
 833 grey, say  $v$ , and implements a stack on `Desc(v)` defining a topological sorting. Clearly, the  
 834 subgraph induced by the union of all these descendant graphs is just the output graph. So  
 835 the concatenation of their topological sortings is a topological sorting of the entire output  
 836 graph. ◀

837 ► **Lemma 38** (Complexity of `top-sort`). *Given a connected acyclic graph of bounded de-*  
 838 *gree with grey unrooted nodes and unmarked edges  $G$  as an input, `SearchUnsortedNodes`*  
 839 *terminates in linear time.*

840 **Proof.** First, let us give an upper bound to the number of applications of each rule. `init`  
 841 is applied exactly once. Since `init` is the only rule having an unmarked root in its right  
 842 hand side, and the input has no unmarked roots, `red_stack` and `sort_back_stack` can  
 843 be matched at most once (in total). `unsorted` and `sort_forward` reduce the number of  
 844 grey nodes by one. Since all the other rules conserve the number of grey nodes, and the  
 845 input graph has  $|V_G|$  grey nodes, they can be applied at most  $|V_G|$  times in total. Similarly,  
 846 `search_forward` (and `init`) reduce the number of nodes with no red looped edge by one. So  
 847 they can also only be applied at most  $|V_G|$  times in total. `red_push` and `sort_back_push` (as  
 848 well as `red_stack` and `sort_back_stack`) are the only rules not to conserve the number of  
 849 blue nodes, and reduce the number of non-blue nodes by exactly one. Since the input graph  
 850 has no blue nodes, they can be applied at most  $|V_G|$  times in total. As for `search_back`, a  
 851 reasoning as in Corollary 23 can be used to justify `search_back` is applied an at most linear  
 852 amount of times, since `SortNodes` conserves the number of dashed edges by Lemma 36.

853 `init` is the only rule to increase the number of roots, specifically by two. All the other  
 854 rules conserve the number of roots. So since the input graph has no roots, there is a constant  
 855 number of roots at any point during the execution of `top-sort`.

856 The only rules that are not fast are `init` due to the lack of roots, and `search_forward`  
 857 due to the `edge` predicate. So by Theorem 1, all the other rules can be matched in constant  
 858 time since the input has bounded degree. `init` is matched in constant time since it matches  
 859 any input node. As for `search_forward`, since the input has bounded degree and the  
 860 rules cannot create an unbounded number of edges incident to a single node, the predicate  
 861 `edge(2,2)` only has to check a constant number of incident edges.

862 Since each rule is matched a linear number of times in constant time, and the program  
 863 terminates by Lemma 34, `top-sort` terminates in linear time. ◀