

Correctness of a Set-based Algorithm for Computing Strongly Connected Components of a Graph

Stephan Merz and Vincent Trélat

March 19, 2025

Abstract

We prove the correctness of a sequential algorithm for computing maximal strongly connected components (SCCs) of a graph due to Vincent Bloemen.

Contents

1	Overview	2
2	Auxiliary lemmas about lists	3
3	Finite directed graphs	5
4	Strongly connected components	7
5	Algorithm for computing strongly connected components	7
6	Definition of the predicates used in the correctness proof	9
6.1	Main invariant	9
6.2	Consequences of the invariant	11
6.3	Pre- and post-conditions of function <i>dfs</i>	12
6.4	Pre- and post-conditions of function <i>dfss</i>	14
7	Proof of partial correctness	15
7.1	Lemmas about function <i>unite</i>	15
7.2	Lemmas establishing the pre-conditions	16
7.3	Lemmas establishing the post-conditions	17
8	Proof of termination and total correctness	18

1 Overview

Computing the maximal strongly connected components (SCCs) of a finite directed graph is a celebrated problem in the theory of graph algorithms. Although Tarjan’s algorithm [5] is perhaps the best-known solution, there are many others. In his PhD thesis, Bloemen [1] presents an algorithm that is itself based on earlier algorithms by Munro [4] and Dijkstra [2]. Just like these algorithms, Bloemen’s solution is based on enumerating SCCs in a depth-first traversal of the graph. Gabow’s algorithm that has already been formalized in Isabelle [3] also falls into this category of solutions. Nevertheless, Bloemen goes on to present a parallel variant of the algorithm suitable for execution on multi-core processors, based on clever data structures that minimize locking.

In the following, we encode the sequential version of the algorithm in the proof assistant Isabelle/HOL, and prove its correctness. Bloemen’s thesis briefly and informally explains why the algorithm is correct. Our proof expands on these arguments, making them completely formal. The encoding is based on a direct representation of the algorithm as a pair of mutually recursive functions; we are not aiming at extracting executable code.

```
theory SCC-Bloemen-Sequential  
imports Main  
begin
```

The record below represents the variables of the algorithm. Most variables correspond to those used in Bloemen’s presentation. Thus, the variable \mathcal{S} associates to every node the set of nodes that have already been determined to be part of the same SCC. A core invariant of the algorithm will be that this mapping represents equivalence classes of nodes: for all nodes v and w , we maintain the relationship

$$v \in \mathcal{S} w \iff \mathcal{S} v = \mathcal{S} w.$$

In an actual implementation of this algorithm, this variable could conveniently be represented by a union-find structure. Variable *stack* holds the list of roots of these (not yet maximal) SCCs, in depth-first order, *visited* and *explored* represent the nodes that have already been seen, respectively that have been completely explored, by the algorithm, and *sccs* is the set of maximal SCCs that the algorithm has found so far.

Additionally, the record holds some auxiliary variables that are used in the proof of correctness. In particular, *root* denotes the node on which the algorithm was called, *cstack* represents the call stack of the recursion of function *dfs*, and *vsuccs* stores the successors of each node that have already been visited by the function *dfss* that loops over all successors of a given node.

```
record 'v env =
```

```

root :: 'v
S :: 'v ⇒ 'v set
explored :: 'v set
visited :: 'v set
vsuccs :: 'v ⇒ 'v set
sccs :: 'v set set
stack :: 'v list
cstack :: 'v list

```

The algorithm is initially called with an environment that initializes the root node and trivializes all other components.

definition *init-env* **where**

```

init-env v = (
  root = v,
  S = (λu. {u}),
  explored = {},
  visited = {},
  vsuccs = (λu. {}),
  sccs = {},
  stack = [],
  cstack = []
)

```

— Make the simplifier expand let-constructions automatically.

declare *Let-def*[simp]

2 Auxiliary lemmas about lists

We use the precedence order on the elements that appear in a list. In particular, stacks are represented as lists, and a node x precedes another node y on the stack if x was pushed on the stack later than y .

definition *precedes* (\prec - \preceq - *in* -> [100,100,100] 39) **where**

```

x ≼ y in xs ≡ ∃ l r. xs = l @ (x # r) ∧ y ∈ set (x # r)

```

lemma *precedes-mem*:

```

assumes x ≼ y in xs
shows x ∈ set xs y ∈ set xs
⟨proof⟩

```

lemma *head-precedes*:

```

assumes y ∈ set (x # xs)
shows x ≼ y in (x # xs)
⟨proof⟩

```

lemma *precedes-in-tail*:

```

assumes x ≠ z
shows x ≼ y in (z # xs) ⟷ x ≼ y in xs

```

<proof>

lemma *tail-not-precedes*:

assumes $y \preceq x$ in $(x \# xs)$ $x \notin \text{set } xs$

shows $x = y$

<proof>

lemma *split-list-precedes*:

assumes $y \in \text{set } (ys @ [x])$

shows $y \preceq x$ in $(ys @ x \# xs)$

<proof>

lemma *precedes-refl* [simp]: $(x \preceq x \text{ in } xs) = (x \in \text{set } xs)$

<proof>

lemma *precedes-append-left*:

assumes $x \preceq y$ in xs

shows $x \preceq y$ in $(ys @ xs)$

<proof>

lemma *precedes-append-left-iff*:

assumes $x \notin \text{set } ys$

shows $x \preceq y$ in $(ys @ xs) \longleftrightarrow x \preceq y$ in xs (**is** ?lhs = ?rhs)

<proof>

lemma *precedes-append-right*:

assumes $x \preceq y$ in xs

shows $x \preceq y$ in $(xs @ ys)$

<proof>

lemma *precedes-append-right-iff*:

assumes $y \notin \text{set } xs$

shows $x \preceq y$ in $(xs @ ys) \longleftrightarrow x \preceq y$ in xs (**is** ?lhs = ?rhs)

<proof>

Precedence determines an order on the elements of a list, provided elements have unique occurrences. However, consider a list such as $[2,3,1,2]$: then 1 precedes 2 and 2 precedes 3, but 1 does not precede 3.

lemma *precedes-trans*:

assumes $x \preceq y$ in xs **and** $y \preceq z$ in xs **and** *distinct* xs

shows $x \preceq z$ in xs

<proof>

lemma *precedes-antisym*:

assumes $x \preceq y$ in xs **and** $y \preceq x$ in xs **and** *distinct* xs

shows $x = y$

<proof>

3 Finite directed graphs

We represent a graph as an Isabelle locale that identifies a finite set of vertices (of some base type $'v$) and associates to each vertex its set of successor vertices.

```
locale graph =  
  fixes vertices :: 'v set  
  and successors :: 'v  $\Rightarrow$  'v set  
  assumes vfin: finite vertices  
  and sclosed:  $\forall x \in$  vertices. successors  $x \subseteq$  vertices  
  
context graph  
begin  
  
abbreviation edge where  
  edge  $x y \equiv y \in$  successors  $x$ 
```

We inductively define reachability of nodes in the graph.

```
inductive reachable where  
  reachable-refl[iff]: reachable  $x x$   
| reachable-succ[elim]:  $\llbracket$ edge  $x y$ ; reachable  $y z$  $\rrbracket \Longrightarrow$  reachable  $x z$ 
```

```
lemma reachable-edge: edge  $x y \Longrightarrow$  reachable  $x y$   
  <proof>
```

```
lemma succ-reachable:  
  assumes reachable  $x y$  and edge  $y z$   
  shows reachable  $x z$   
  <proof>
```

```
lemma reachable-trans:  
  assumes  $y$ : reachable  $x y$  and  $z$ : reachable  $y z$   
  shows reachable  $x z$   
  <proof>
```

In order to derive a “reverse” induction rule for *reachable*, we define an alternative reachability predicate and prove that the two coincide.

```
inductive reachable-end where  
  re-refl[iff]: reachable-end  $x x$   
| re-succ:  $\llbracket$ reachable-end  $x y$ ; edge  $y z$  $\rrbracket \Longrightarrow$  reachable-end  $x z$ 
```

```
lemma succ-re:  
  assumes  $y$ : edge  $x y$  and  $z$ : reachable-end  $y z$   
  shows reachable-end  $x z$   
  <proof>
```

```
lemma reachable-re:  
  assumes reachable  $x y$ 
```

shows *reachable-end* $x y$
 ⟨*proof*⟩

lemma *re-reachable*:
assumes *reachable-end* $x y$
shows *reachable* $x y$
 ⟨*proof*⟩

lemma *reachable-end-induct*:
assumes r : *reachable* $x y$
and *base*: $\bigwedge x. P x x$
and *step*: $\bigwedge x y z. \llbracket P x y; \text{edge } y z \rrbracket \implies P x z$
shows $P x y$
 ⟨*proof*⟩

We also need the following variant of reachability avoiding certain edges. More precisely, y is reachable from x avoiding a set E of edges if there exists a path such that no edge from E appears along the path.

inductive *reachable-avoiding* **where**
ra-refl[*iff*]: *reachable-avoiding* $x x E$
| *ra-succ*[*elim*]: $\llbracket \text{reachable-avoiding } x y E; \text{edge } y z; (y,z) \notin E \rrbracket \implies \text{reachable-avoiding } x z E$

lemma *edge-ra*:
assumes *edge* $x y$ **and** $(x,y) \notin E$
shows *reachable-avoiding* $x y E$
 ⟨*proof*⟩

lemma *ra-trans*:
assumes 1: *reachable-avoiding* $x y E$ **and** 2: *reachable-avoiding* $y z E$
shows *reachable-avoiding* $x z E$
 ⟨*proof*⟩

lemma *ra-cases*:
assumes *reachable-avoiding* $x y E$
shows $x=y \vee (\exists z. z \in \text{successors } x \wedge (x,z) \notin E \wedge \text{reachable-avoiding } z y E)$
 ⟨*proof*⟩

lemma *ra-mono*:
assumes *reachable-avoiding* $x y E$ **and** $E' \subseteq E$
shows *reachable-avoiding* $x y E'$
 ⟨*proof*⟩

lemma *ra-add-edge*:
assumes *reachable-avoiding* $x y E$
shows *reachable-avoiding* $x y (E \cup \{(v,w)\})$
 $\vee (\text{reachable-avoiding } x v (E \cup \{(v,w)\}) \wedge \text{reachable-avoiding } w y (E \cup \{(v,w)\}))$
 ⟨*proof*⟩

Reachability avoiding some edges obviously implies reachability. Conversely, reachability implies reachability avoiding the empty set.

lemma *ra-reachable*:

reachable-avoiding $x\ y\ E \implies \text{reachable}\ x\ y$
 ⟨*proof*⟩

lemma *ra-empty*:

reachable-avoiding $x\ y\ \{\} = \text{reachable}\ x\ y$
 ⟨*proof*⟩

4 Strongly connected components

A strongly connected component is a set S of nodes such that any two nodes in S are reachable from each other. This concept is represented by the predicate *is-subsc* below. We are ultimately interested in non-empty, maximal strongly connected components, represented by the predicate *is-scc*.

definition *is-subsc* **where**

is-subsc $S \equiv \forall x \in S. \forall y \in S. \text{reachable}\ x\ y$

definition *is-scc* **where**

is-scc $S \equiv S \neq \{\} \wedge \text{is-subsc}\ S \wedge (\forall S'. S \subseteq S' \wedge \text{is-subsc}\ S' \longrightarrow S' = S)$

lemma *subsc-add*:

assumes *is-subsc* S **and** $x \in S$
and *reachable* $x\ y$ **and** *reachable* $y\ x$
shows *is-subsc* (*insert* $y\ S$)
 ⟨*proof*⟩

lemma *sccE*:

— Two nodes that are reachable from each other are in the same SCC.
assumes *is-scc* S **and** $x \in S$
and *reachable* $x\ y$ **and** *reachable* $y\ x$
shows $y \in S$
 ⟨*proof*⟩

lemma *scc-partition*:

— Two SCCs that contain a common element are identical.
assumes *is-scc* S **and** *is-scc* S' **and** $x \in S \cap S'$
shows $S = S'$
 ⟨*proof*⟩

5 Algorithm for computing strongly connected components

We now introduce our representation of Bloemen's algorithm in Isabelle/HOL. The auxiliary function *unite* corresponds to the inner while loop in Bloemen's

pseudo-code [1, p.32]. It is applied to two nodes v and w (and the environment e holding the current values of the program variables) when a loop is found, i.e. when w is a successor of v in the graph that has already been visited in the depth-first search. In that case, the root of the SCC of node w determined so far must appear below the root of v 's SCC in the *stack* maintained by the algorithm. The effect of the function is to merge the SCCs of all nodes on the top of the stack above (and including) w . Node w 's root will be the root of the merged SCC.

definition $unite :: 'v \Rightarrow 'v \Rightarrow 'v \text{ env} \Rightarrow 'v \text{ env}$ **where**
 $unite\ v\ w\ e \equiv$
 $let\ pfx = takeWhile\ (\lambda x. w \notin \mathcal{S}\ e\ x)\ (stack\ e);$
 $sfx = dropWhile\ (\lambda x. w \notin \mathcal{S}\ e\ x)\ (stack\ e);$
 $cc = \bigcup \{ \mathcal{S}\ e\ x \mid x. x \in set\ pfx \cup \{hd\ sfx\} \}$
 $in\ e(\mathcal{S} := \lambda x. if\ x \in cc\ then\ cc\ else\ \mathcal{S}\ e\ x,$
 $stack := sfx)$

We now represent the algorithm as two mutually recursive functions dfs and $dfss$ in Isabelle/HOL. The function dfs corresponds to Bloemen's function `SetBased`, whereas $dfss$ corresponds to the `forall` loop over the successors of the node on which dfs was called. Instead of using global program variables in imperative style, our functions explicitly pass environments that hold the current values of these variables.

A technical complication in the development of the algorithm in Isabelle is the fact that the functions need not terminate when their pre-conditions (introduced below) are violated, for example when dfs is called for a node that was already visited previously. We therefore cannot prove termination at this point, but will later show that the explicitly given pre-conditions ensure termination.

function (*domintros*) $dfs :: 'v \Rightarrow 'v \text{ env} \Rightarrow 'v \text{ env}$
and $dfss :: 'v \Rightarrow 'v \text{ env} \Rightarrow 'v \text{ env}$ **where**
 $dfs\ v\ e =$
 $(let\ e1 = e(\mathit{visited} := \mathit{visited}\ e \cup \{v\},$
 $stack := (v \# stack\ e),$
 $cstack := (v \# cstack\ e));$
 $e' = dfss\ v\ e1$
 $in\ if\ v = hd(stack\ e')$
 $then\ e'(\mathit{sccs} := \mathit{sccs}\ e' \cup \{\mathcal{S}\ e'\ v\},$
 $explored := explored\ e' \cup (\mathcal{S}\ e'\ v),$
 $stack := tl(stack\ e'),$
 $cstack := tl(cstack\ e'))$
 $else\ e'(\mathit{cstack} := tl(cstack\ e')))$
 $| dfss\ v\ e =$
 $(let\ vs = successors\ v - vsuccs\ e\ v$
 $in\ if\ vs = \{\}$ then e
 $else\ let\ w = SOME\ x. x \in vs;$
 $e' = (if\ w \in explored\ e\ then\ e$


```

      else if  $w \notin \text{visited } e$ 
      then  $\text{dfs } w \ e$ 
      else  $\text{unite } v \ w \ e$ ;
 $e'' = (e'(\text{vsucss} :=$ 
      ( $\lambda x. \text{if } x=v \text{ then } \text{vsucss } e' \ v \cup \{w\}$ 
      else  $\text{vsucss } e' \ x$ ))
  in  $\text{dfss } v \ e''$ )
<proof>

```

6 Definition of the predicates used in the correctness proof

Environments are partially ordered according to the following definition.

definition *sub-env* where

```

 $\text{sub-env } e \ e' \equiv$ 
   $\text{root } e' = \text{root } e$ 
 $\wedge \text{visited } e \subseteq \text{visited } e'$ 
 $\wedge \text{explored } e \subseteq \text{explored } e'$ 
 $\wedge (\forall v. \text{vsucss } e \ v \subseteq \text{vsucss } e' \ v)$ 
 $\wedge (\forall v. \mathcal{S} \ e \ v \subseteq \mathcal{S} \ e' \ v)$ 
 $\wedge (\bigcup \{\mathcal{S} \ e \ v \mid v. v \in \text{set } (\text{stack } e)\}$ 
   $\subseteq \bigcup \{\mathcal{S} \ e' \ v \mid v. v \in \text{set } (\text{stack } e')\})$ 

```

lemma *sub-env-trans*:

```

assumes  $\text{sub-env } e \ e'$  and  $\text{sub-env } e' \ e''$ 
shows  $\text{sub-env } e \ e''$ 
<proof>

```

The set *unvisited* $e \ u$ contains all edges (a,b) such that node a is in the same SCC as node u and the edge has not yet been followed, in the sense represented by variable *vsucss*.

definition *unvisited* where

```

 $\text{unvisited } e \ u \equiv$ 
 $\{(a,b) \mid a \ b. a \in \mathcal{S} \ e \ u \wedge b \in \text{successors } a - \text{vsucss } e \ a\}$ 

```

6.1 Main invariant

The following definition characterizes well-formed environments. This predicate will be shown to hold throughout the execution of the algorithm. In words, it asserts the following facts:

- Only nodes reachable from the root (for which the algorithm was originally called) are visited.
- The two stacks *stack* and *cstack* do not contain duplicate nodes, and *stack* contains a subset of the nodes on *cstack*, in the same order.

- Any node higher on the *stack* (i.e., that was pushed later) is reachable from nodes lower in the *stack*. This property also holds for nodes on the call stack, but this is not needed for the correctness proof.
- Every explored node, and every node on the call stack, has been visited.
- Nodes reachable from fully explored nodes have themselves been fully explored.
- The set $vsuccs\ e\ n$, for any node n , is a subset of n 's successors, and all these nodes are in *visited*. The set is empty if $n \notin visited$, and it contains all successors if n has been fully explored or if n has been visited, but is no longer on the call stack.
- The sets $\mathcal{S}\ e\ n$ represent an equivalence relation. The equivalence classes of nodes that have not yet been visited are singletons. Also, equivalence classes for two distinct nodes on the *stack* are disjoint because the stack only stores roots of SCCs, and the union of the equivalence classes for these root nodes corresponds to the set of live nodes, i.e. those nodes that have already been visited but not yet fully explored.
- More precisely, an equivalence class is represented on the stack by the oldest node in the sense of the call order: any node in the class that is still on the call stack precedes the representative on the call stack and was therefore pushed later.
- Equivalence classes represent the maximal available information about strong connectedness: nodes represented by some node n on the *stack* can reach some node m that is lower in the stack only by taking an edge from some node in n 's equivalence class that has not yet been followed. (Remember that m can reach n by one of the previous conjuncts.)
- Equivalence classes represent partial SCCs in the sense of the predicate *is-subsc*. Variable *sccs* holds maximal SCCs in the sense of the predicate *is-scc*, and their union corresponds to the set of explored nodes.

definition *wf-env* where

$$\begin{aligned}
wf\text{-env}\ e &\equiv \\
&(\forall n \in visited\ e.\ reachable\ (root\ e)\ n) \\
&\wedge\ distinct\ (stack\ e) \\
&\wedge\ distinct\ (cstack\ e) \\
&\wedge\ (\forall n\ m.\ n \preceq m\ in\ stack\ e \longrightarrow n \preceq m\ in\ cstack\ e) \\
&\wedge\ (\forall n\ m.\ n \preceq m\ in\ stack\ e \longrightarrow reachable\ m\ n) \\
&\wedge\ explored\ e \subseteq visited\ e \\
&\wedge\ set\ (cstack\ e) \subseteq visited\ e
\end{aligned}$$

$$\begin{aligned}
& \wedge (\forall n \in \text{explored } e. \forall m. \text{reachable } n \ m \longrightarrow m \in \text{explored } e) \\
& \wedge (\forall n. \text{vsuccs } e \ n \subseteq \text{successors } n \cap \text{visited } e) \\
& \wedge (\forall n. n \notin \text{visited } e \longrightarrow \text{vsuccs } e \ n = \{\}) \\
& \wedge (\forall n \in \text{explored } e. \text{vsuccs } e \ n = \text{successors } n) \\
& \wedge (\forall n \in \text{visited } e - \text{set } (\text{cstack } e). \text{vsuccs } e \ n = \text{successors } n) \\
& \wedge (\forall n \ m. m \in \mathcal{S} \ e \ n \longleftrightarrow (\mathcal{S} \ e \ n = \mathcal{S} \ e \ m)) \\
& \wedge (\forall n. n \notin \text{visited } e \longrightarrow \mathcal{S} \ e \ n = \{n\}) \\
& \wedge (\forall n \in \text{set } (\text{stack } e). \forall m \in \text{set } (\text{stack } e). n \neq m \longrightarrow \mathcal{S} \ e \ n \cap \mathcal{S} \ e \ m = \{\}) \\
& \wedge \bigcup \{\mathcal{S} \ e \ n \mid n. n \in \text{set } (\text{stack } e)\} = \text{visited } e - \text{explored } e \\
& \wedge (\forall n \in \text{set } (\text{stack } e). \forall m \in \mathcal{S} \ e \ n. m \in \text{set } (\text{cstack } e) \longrightarrow m \preceq n \text{ in } \text{cstack } e) \\
& \wedge (\forall n \ m. n \preceq m \text{ in } \text{stack } e \wedge n \neq m \longrightarrow \\
& \quad (\forall u \in \mathcal{S} \ e \ n. \neg \text{reachable-avoiding } u \ m \ (\text{unvisited } e \ n))) \\
& \wedge (\forall n. \text{is-subsc} (\mathcal{S} \ e \ n)) \\
& \wedge (\forall S \in \text{sccs } e. \text{is-scc } S) \\
& \wedge \bigcup (\text{sccs } e) = \text{explored } e
\end{aligned}$$

6.2 Consequences of the invariant

Since every node on the call stack is an element of *visited* and every node on the *stack* also appears on *cstack*, all these nodes are also in *visited*.

lemma *stack-visited*:

assumes *wf-env* $e \ n \in \text{set } (\text{stack } e)$

shows $n \in \text{visited } e$

<proof>

Classes represented on the stack consist of visited nodes that have not yet been fully explored.

lemma *stack-class*:

assumes *wf-env* $e \ n \in \text{set } (\text{stack } e) \ m \in \mathcal{S} \ e \ n$

shows $m \in \text{visited } e - \text{explored } e$

<proof>

Conversely, every such node belongs to some class represented on the stack.

lemma *visited-unexplored*:

assumes *wf-env* $e \ m \in \text{visited } e \ m \notin \text{explored } e$

obtains n **where** $n \in \text{set } (\text{stack } e) \ m \in \mathcal{S} \ e \ n$

<proof>

Every node belongs to its own equivalence class.

lemma *S-reflexive*:

assumes *wf-env* e

shows $n \in \mathcal{S} \ e \ n$

<proof>

No node on the stack has been fully explored.

lemma *stack-unexplored*:

assumes 1 : *wf-env* e

and 2: $n \in \text{set } (\text{stack } e)$
and 3: $n \in \text{explored } e$
shows P
 $\langle \text{proof} \rangle$

If w is reachable from visited node v , but no unvisited successor of a node reachable from v can reach w , then w must be visited.

lemma *reachable-visited*:
assumes $e: \text{wf-env } e$
and $v: v \in \text{visited } e$
and $w: \text{reachable } v w$
and $s: \forall n \in \text{visited } e. \forall m \in \text{successors } n - \text{vsuccs } e n.$
 $\text{reachable } v n \longrightarrow \neg \text{reachable } m w$
shows $w \in \text{visited } e$
 $\langle \text{proof} \rangle$

Edges towards explored nodes do not contribute to reachability of unexplored nodes avoiding some set of edges.

lemma *avoiding-explored*:
assumes $e: \text{wf-env } e$
and $xy: \text{reachable-avoiding } x y E$
and $y: y \notin \text{explored } e$
and $w: w \in \text{explored } e$
shows $\text{reachable-avoiding } x y (E \cup \{(v,w)\})$
 $\langle \text{proof} \rangle$

6.3 Pre- and post-conditions of function *dfs*

Function *dfs* should be called for a well-formed environment and a node v that has not yet been visited and that is reachable from the root node, as well as from all nodes in the stack. No outgoing edges from node v have yet been followed.

definition *pre-dfs* **where**
 $\text{pre-dfs } v e \equiv$
 $\text{wf-env } e$
 $\wedge v \notin \text{visited } e$
 $\wedge \text{reachable } (\text{root } e) v$
 $\wedge \text{vsuccs } e v = \{\}$
 $\wedge (\forall n \in \text{set } (\text{stack } e). \text{reachable } n v)$

Function *dfs* maintains the invariant *wf-env* and returns an environment e' that extends the input environment e . Node v has been visited and all its outgoing edges have been followed. Because the algorithm works in depth-first fashion, no new outgoing edges of nodes that had already been visited in the input environment have been followed, and the stack of e' is a suffix of the one of e such that v is still reachable from all nodes on the stack. The stack may have been shortened because SCCs represented at the top of the

stack may have been merged. The call stack is reestablished as it was in e . There are two possible outcomes of the algorithm:

- Either v has been fully explored, in which case the stacks of e and e' are the same, and the equivalence classes of all nodes represented on the stack are unchanged. This corresponds to the case where v is the root node of its (maximal) SCC.
- Alternatively, the stack of e' must be non-empty and v must be represented by the node at the top of the stack. The SCCs of the nodes lower on the stack are unchanged. This corresponds to the case where v is not the root node of its SCC, but some SCCs at the top of the stack may have been merged.

definition *post-dfs* where

$$\begin{aligned}
\text{post-dfs } v \ e \ e' &\equiv \\
&wf\text{-env } e' \\
&\wedge v \in \text{visited } e' \\
&\wedge \text{sub-env } e \ e' \\
&\wedge \text{vsuccs } e' \ v = \text{successors } v \\
&\wedge (\forall w \in \text{visited } e. \text{vsuccs } e' \ w = \text{vsuccs } e \ w) \\
&\wedge (\forall n \in \text{set } (\text{stack } e'). \text{reachable } n \ v) \\
&\wedge (\exists ns. \text{stack } e = ns \ @ \ (\text{stack } e')) \\
&\wedge ((v \in \text{explored } e' \wedge \text{stack } e' = \text{stack } e \\
&\quad \wedge (\forall n \in \text{set } (\text{stack } e'). \mathcal{S} \ e' \ n = \mathcal{S} \ e \ n)) \\
&\quad \vee (\text{stack } e' \neq [] \wedge v \in \mathcal{S} \ e' \ (\text{hd } (\text{stack } e')) \\
&\quad \wedge (\forall n \in \text{set } (\text{tl } (\text{stack } e')). \mathcal{S} \ e' \ n = \mathcal{S} \ e \ n))) \\
&\wedge \text{cstack } e' = \text{cstack } e
\end{aligned}$$

The initial environment is easily seen to satisfy *dfs*'s pre-condition.

lemma *init-env-pre-dfs*: $\text{pre-dfs } v \ (\text{init-env } v)$
 $\langle \text{proof} \rangle$

Any node represented by the top stack element of the input environment is still represented by the top element of the output stack.

lemma *dfs-S-hd-stack*:

assumes wf : $wf\text{-env } e$
and $post$: $\text{post-dfs } v \ e \ e'$
and n : $\text{stack } e \neq [] \wedge n \in \mathcal{S} \ e \ (\text{hd } (\text{stack } e))$
shows $\text{stack } e' \neq [] \wedge n \in \mathcal{S} \ e' \ (\text{hd } (\text{stack } e'))$
 $\langle \text{proof} \rangle$

Function *dfs* leaves the SCCs represented by elements in the (new) tail of the *stack* unchanged.

lemma *dfs-S-tl-stack*:

assumes $post$: $\text{post-dfs } v \ e \ e'$
and $nempty$: $\text{stack } e \neq []$
shows $\text{stack } e' \neq [] \wedge \forall n \in \text{set } (\text{tl } (\text{stack } e')). \mathcal{S} \ e' \ n = \mathcal{S} \ e \ n$
 $\langle \text{proof} \rangle$

6.4 Pre- and post-conditions of function *dfss*

The pre- and post-conditions of function *dfss* correspond to the invariant of the loop over all outgoing edges from node *v*. The environment must be well-formed, node *v* must be visited and represented by the top element of the (non-empty) stack. Node *v* must be reachable from all nodes on the stack, and it must be the top node on the call stack. All outgoing edges of node *v* that have already been followed must either lead to completely explored nodes (that are no longer represented on the stack) or to nodes that are part of the same SCC as *v*.

definition *pre-dfss* where

$$\begin{aligned}
 \text{pre-dfss } v \ e \equiv & \\
 & \text{wf-env } e \\
 & \wedge v \in \text{visited } e \\
 & \wedge (\text{stack } e \neq []) \\
 & \wedge (v \in \mathcal{S} \ e \ (\text{hd} \ (\text{stack } e))) \\
 & \wedge (\forall w \in \text{vsuccs } e \ v. w \in \text{explored } e \cup \mathcal{S} \ e \ (\text{hd} \ (\text{stack } e))) \\
 & \wedge (\forall n \in \text{set} \ (\text{stack } e). \text{reachable } n \ v) \\
 & \wedge (\exists ns. \text{cstack } e = v \ \# \ ns)
 \end{aligned}$$

The post-condition establishes that all outgoing edges of node *v* have been followed. As for function *dfs*, no new outgoing edges of previously visited nodes have been followed. Also as before, the new stack is a suffix of the old one, and the call stack is restored. In case node *v* is still on the stack (and therefore is the root node of its SCC), no node that is lower on the stack can be reachable from *v*. This condition guarantees the maximality of the computed SCCs.

definition *post-dfss* where

$$\begin{aligned}
 \text{post-dfss } v \ e \ e' \equiv & \\
 & \text{wf-env } e' \\
 & \wedge \text{vsuccs } e' \ v = \text{successors } v \\
 & \wedge (\forall w \in \text{visited } e - \{v\}. \text{vsuccs } e' \ w = \text{vsuccs } e \ w) \\
 & \wedge \text{sub-env } e \ e' \\
 & \wedge (\forall w \in \text{successors } v. w \in \text{explored } e' \cup \mathcal{S} \ e' \ (\text{hd} \ (\text{stack } e'))) \\
 & \wedge (\forall n \in \text{set} \ (\text{stack } e'). \text{reachable } n \ v) \\
 & \wedge (\text{stack } e' \neq []) \\
 & \wedge (\exists ns. \text{stack } e = ns \ @ \ (\text{stack } e')) \\
 & \wedge v \in \mathcal{S} \ e' \ (\text{hd} \ (\text{stack } e')) \\
 & \wedge (\forall n \in \text{set} \ (\text{tl} \ (\text{stack } e')). \mathcal{S} \ e' \ n = \mathcal{S} \ e \ n) \\
 & \wedge (\text{hd} \ (\text{stack } e') = v \longrightarrow (\forall n \in \text{set} \ (\text{tl} \ (\text{stack } e')). \neg \text{reachable } v \ n)) \\
 & \wedge \text{cstack } e' = \text{cstack } e
 \end{aligned}$$

7 Proof of partial correctness

7.1 Lemmas about function *unite*

We start by establishing a few lemmas about function *unite* in the context where it is called.

lemma *unite-stack*:

fixes $e v w$

defines $e' \equiv \text{unite } v w e$

assumes $wf: wf\text{-env } e$

and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$

obtains px **where** $\text{stack } e = px \text{ @ } (\text{stack } e')$

$\text{stack } e' \neq []$

$\text{let } cc = \bigcup \{ \mathcal{S} e n \mid n. n \in \text{set } px \cup \{hd (\text{stack } e')\} \}$

$\text{in } \mathcal{S} e' = (\lambda x. \text{if } x \in cc \text{ then } cc \text{ else } \mathcal{S} e x)$

$w \in \mathcal{S} e' (hd (\text{stack } e'))$

<proof>

Function *unite* leaves intact the equivalence classes represented by the tail of the new stack.

lemma *unite-S-tl*:

fixes $e v w$

defines $e' \equiv \text{unite } v w e$

assumes $wf: wf\text{-env } e$

and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$

and $n: n \in \text{set } (tl (\text{stack } e'))$

shows $\mathcal{S} e' n = \mathcal{S} e n$

<proof>

The stack of the result of *unite* represents the same vertices as the input stack, potentially in fewer equivalence classes.

lemma *unite-S-equal*:

fixes $e v w$

defines $e' \equiv \text{unite } v w e$

assumes $wf: wf\text{-env } e$

and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$

shows $(\bigcup \{ \mathcal{S} e' n \mid n. n \in \text{set } (\text{stack } e') \}) = (\bigcup \{ \mathcal{S} e n \mid n. n \in \text{set } (\text{stack } e) \})$

<proof>

The head of the stack represents a (not necessarily maximal) SCC.

lemma *unite-subsc*:

fixes $e v w$

defines $e' \equiv \text{unite } v w e$

assumes $pre: pre\text{-dfss } v e$

and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$

shows $is\text{-subsc} (\mathcal{S} e' (hd (\text{stack } e')))$

<proof>

The environment returned by function *unite* extends the input environment.

lemma *unite-sub-env*:
fixes $e v w$
defines $e' \equiv \text{unite } v w e$
assumes $\text{pre: pre-dfss } v e$
and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$
shows $\text{sub-env } e e'$
 $\langle \text{proof} \rangle$

The environment returned by function *unite* is well-formed.

lemma *unite-wf-env*:
fixes $e v w$
defines $e' \equiv \text{unite } v w e$
assumes $\text{pre: pre-dfss } v e$
and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$
shows $\text{wf-env } e'$
 $\langle \text{proof} \rangle$

7.2 Lemmas establishing the pre-conditions

The precondition of function *dfs* ensures the precondition of *dfss* at the call of that function.

lemma *pre-dfs-pre-dfss*:
assumes $\text{pre-dfs } v e$
shows $\text{pre-dfss } v (e(\text{visited} := \text{visited } e \cup \{v\},$
 $\text{stack} := v \# \text{stack } e,$
 $\text{cstack} := v \# \text{cstack } e))$
 $(\text{is pre-dfss } v ?e')$
 $\langle \text{proof} \rangle$

Similarly, we now show that the pre-conditions of the different function calls in the body of function *dfss* are satisfied. First, it is very easy to see that the pre-condition of *dfs* holds at the call of that function.

lemma *pre-dfss-pre-dfs*:
assumes $\text{pre-dfss } v e$ **and** $w \notin \text{visited } e$ **and** $w \in \text{successors } v$
shows $\text{pre-dfs } w e$
 $\langle \text{proof} \rangle$

The pre-condition of *dfss* holds when the successor considered in the current iteration has already been explored.

lemma *pre-dfss-explored-pre-dfss*:
fixes $e v w$
defines $e'' \equiv e(\text{vsuccs} := (\lambda x. \text{if } x=v \text{ then vsuccs } e v \cup \{w\} \text{ else vsuccs } e x))$
assumes $1: \text{pre-dfss } v e$ **and** $2: w \in \text{successors } v$ **and** $3: w \in \text{explored } e$
shows $\text{pre-dfss } v e''$
 $\langle \text{proof} \rangle$

The call to *dfs* establishes the pre-condition for the recursive call to *dfss* in the body of *dfss*.

lemma *pre-dfss-post-dfs-pre-dfss*:
fixes $e v w$
defines $e' \equiv \text{dfs } w e$
defines $e'' \equiv e'(\backslash \text{vsuccs} := (\lambda x. \text{if } x=v \text{ then } \text{vsuccs } e' v \cup \{w\} \text{ else } \text{vsuccs } e' x))$
assumes *pre*: *pre-dfss* $v e$
and $w: w \in \text{successors } v w \notin \text{visited } e$
and *post*: *post-dfs* $w e e'$
shows *pre-dfss* $v e''$
 $\langle \text{proof} \rangle$

Finally, the pre-condition for the recursive call to *dfs* at the end of the body of function *dfss* also holds if *unite* was applied.

lemma *pre-dfss-unite-pre-dfss*:
fixes $e v w$
defines $e' \equiv \text{unite } v w e$
defines $e'' \equiv e'(\backslash \text{vsuccs} := (\lambda x. \text{if } x=v \text{ then } \text{vsuccs } e' v \cup \{w\} \text{ else } \text{vsuccs } e' x))$
assumes *pre*: *pre-dfss* $v e$
and $w: w \in \text{successors } v w \notin \text{vsuccs } e v w \in \text{visited } e w \notin \text{explored } e$
shows *pre-dfss* $v e''$
 $\langle \text{proof} \rangle$

7.3 Lemmas establishing the post-conditions

Assuming the pre-condition of function *dfs* and the post-condition of the call to *dfss* in the body of that function, the post-condition of *dfs* is established.

lemma *pre-dfs-implies-post-dfs*:
fixes $v e$
defines $e1 \equiv e(\backslash \text{visited} := \text{visited } e \cup \{v\},$
 $\text{stack} := (v \# \text{stack } e),$
 $\text{cstack} := (v \# \text{cstack } e))$
defines $e' \equiv \text{dfs } v e1$
defines $e'' \equiv e'(\backslash \text{cstack} := \text{tl}(\text{cstack } e'))$
assumes *1*: *pre-dfs* $v e$
and *2*: *dfs-dfss-dom* $(\text{Inl}(v, e))$
and *3*: *post-dfss* $v e1 e'$
shows *post-dfs* $v e (\text{dfs } v e)$
 $\langle \text{proof} \rangle$

The following lemma is central for proving partial correctness: assuming termination (represented by the predicate *dfs-dfss-dom*) and the pre-condition of the functions, both *dfs* and *dfss* establish their post-conditions. The first part of the theorem follows directly from the preceding lemma and the computational induction rule generated by Isabelle, the second part is proved directly, distinguishing the different cases in the definition of function *dfss*.

lemma *pre-post*:
shows
 $\llbracket \text{dfs-dfss-dom } (\text{Inl}(v, e)); \text{pre-dfs } v e \rrbracket \implies \text{post-dfs } v e (\text{dfs } v e)$

$\llbracket \text{dfs-dfss-dom } (Inr(v,e)); \text{pre-dfss } v \ e \rrbracket \implies \text{post-dfss } v \ e \ (\text{dfss } v \ e)$
 $\langle \text{proof} \rangle$

We can now show partial correctness of the algorithm: applied to some node v and the empty environment, it computes the set of strongly connected components in the subgraph reachable from node v . In particular, if v is a root of the graph, the algorithm computes the set of SCCs of the graph.

theorem *partial-correctness*:

fixes v

defines $e \equiv \text{dfs } v \ (\text{init-env } v)$

assumes $\text{dfs-dfss-dom } (Inl(v, \text{init-env } v))$

shows $\text{sccs } e = \{S \ . \ \text{is-scc } S \ \wedge \ (\forall n \in S. \ \text{reachable } v \ n)\}$

(**is** $- = ?\text{rhs}$)

$\langle \text{proof} \rangle$

8 Proof of termination and total correctness

We define a binary relation on the arguments of functions dfs and dfss , and prove that this relation is well-founded and that all calls within the function bodies respect the relation, assuming that the pre-conditions of the initial function call are satisfied. By well-founded induction, we conclude that the pre-conditions of the functions are sufficient to ensure termination.

Following the internal representation of the two mutually recursive functions in Isabelle as a single function on the disjoint sum of the types of arguments, our relation is defined as a set of argument pairs injected into the sum type. The left injection Inl takes arguments of function dfs , the right injection Inr takes arguments of function dfss .¹ The conditions on the arguments in the definition of the relation overapproximate the arguments in the actual calls.

definition $\text{dfs-dfss-term}::((\text{'v} \times \text{'v env} + \text{'v} \times \text{'v env}) \times (\text{'v} \times \text{'v env} + \text{'v} \times \text{'v env})) \ \text{set}$ **where**

$\text{dfs-dfss-term} \equiv$

$\{ (Inr(v, e1), Inl(v, e)) \mid v \ e \ e1. \}$

$v \in \text{vertices} - \text{visited } e \ \wedge \ \text{visited } e1 = \text{visited } e \cup \{v\} \}$

$\cup \{ (Inl(w, e), Inr(v, e)) \mid v \ w \ e. \ v \in \text{vertices} \}$

$\cup \{ (Inr(v, e''), Inr(v, e)) \mid v \ e \ e''. \}$

$v \in \text{vertices} \ \wedge \ \text{sub-env } e \ e''$

$\wedge \ (\exists w \in \text{vertices}. \ w \notin \text{vsuccs } e \ v \ \wedge \ w \in \text{vsuccs } e'' \ v) \}$

Informally, termination is ensured because at each call, either a new vertex is visited (hence the complement of the set of visited nodes w.r.t. the finite set of vertices decreases) or a new successor is added to the set $\text{vsuccs } e \ v$ of some vertex v .

¹Note that the types of the arguments of dfs and dfss are actually identical. We nevertheless use the sum type in order to remember the function that was called.

In order to make this argument formal, we inject the argument tuples that appear in our relation into tuples consisting of the sets mentioned in the informal argument. However, there is one added complication because the call of *dfs* from *dfss* does not immediately add the vertex to the set of visited nodes (this happens only at the beginning of function *dfs*). We therefore add a third component of 0 or 1 to these tuples, reflecting the fact that there can only be one call of *dfs* from *dfss* for a given vertex *v*.

fun *dfs-dfss-to-tuple* **where**

$$\begin{aligned} & \text{dfs-dfss-to-tuple } (Inl(v::'v, e::'v \text{ env})) = \\ & \quad (\text{vertices} - \text{visited } e, \text{vertices} \times \text{vertices} - \{(u, u') \mid u \ u'. \ u' \in \text{vsuccs } e \ u\}, 0) \\ | & \text{dfs-dfss-to-tuple } (Inr(v::'v, e::'v \text{ env})) = \\ & \quad (\text{vertices} - \text{visited } e, \text{vertices} \times \text{vertices} - \{(u, u') \mid u \ u'. \ u' \in \text{vsuccs } e \ u\}, 1::\text{nat}) \end{aligned}$$

The triples defined in this way can be ordered lexicographically (with the first two components ordered as finite subsets and the third one following the predecessor relation on natural numbers). We prove that the injection of the above relation into sets of triples respects the lexicographic ordering and conclude that our relation is well-founded.

lemma *wf-term: wf dfs-dfss-term*

<proof>

The following theorem establishes sufficient conditions that ensure termination of the two functions *dfs* and *dfss*. The proof proceeds by well-founded induction using the relation *dfs-dfss-term*. Isabelle represents the termination domains of the functions by the predicate *dfs-dfss-dom* and generates a theorem *dfs-dfss.domintros* for proving membership of arguments in the termination domains. The actual formulation is a little technical because the mutual induction must again be encoded in a single induction argument over the sum type representing the arguments of both functions.

theorem *dfs-dfss-termination:*

$$\begin{aligned} & \llbracket v \in \text{vertices} ; \text{pre-dfs } v \ e \rrbracket \implies \text{dfs-dfss-dom}(Inl(v, e)) \\ & \llbracket v \in \text{vertices} ; \text{pre-dfss } v \ e \rrbracket \implies \text{dfs-dfss-dom}(Inr(v, e)) \end{aligned}$$

<proof>

Putting everything together, we prove the total correctness of the algorithm when applied to some (root) vertex.

theorem *correctness:*

assumes $v \in \text{vertices}$

shows $\text{sccs } (\text{dfs } v \ (\text{init-env } v)) = \{S . \text{is-scc } S \wedge (\forall n \in S. \text{reachable } v \ n)\}$

<proof>

end

end

References

- [1] V. Bloemen. *Strong Connectivity and Shortest Paths for Checking Models*. PhD thesis, University of Twente, Enschede, The Netherlands, 2019.
- [2] E. W. Dijkstra. Finding the maximum strong components in a directed graph. In *Selected Writings in Computing: A Personal Perspective*, Texts and Monographs in Computer Science, pages 22–30. Springer, 1982.
- [3] P. Lammich. Verified efficient implementation of gabow’s strongly connected components algorithm. *Archive of Formal Proofs*, May 2014. https://isa-afp.org/entries/Gabow_SCC.html, Formal proof development.
- [4] J. Munro. Efficient determination of the transitive closure of a directed graph. *Information Processing Letters*, 1(2):56–58, 1971.
- [5] R. Tarjan. Depth first search and linear graph algorithms. *SIAM Journal on Computing*, 1972.