

MACHINE-CHECKED C IMPLEMENTATION OF DIJKSTRA

Anshuman Mohan Aquinas Hobor

National University of Singapore

Overview

We report on a machine-checked proof of correctness for Dijkstra’s one-to-all shortest path algorithm. Unlike previous work, we use classic textbook code written in executable C [2].

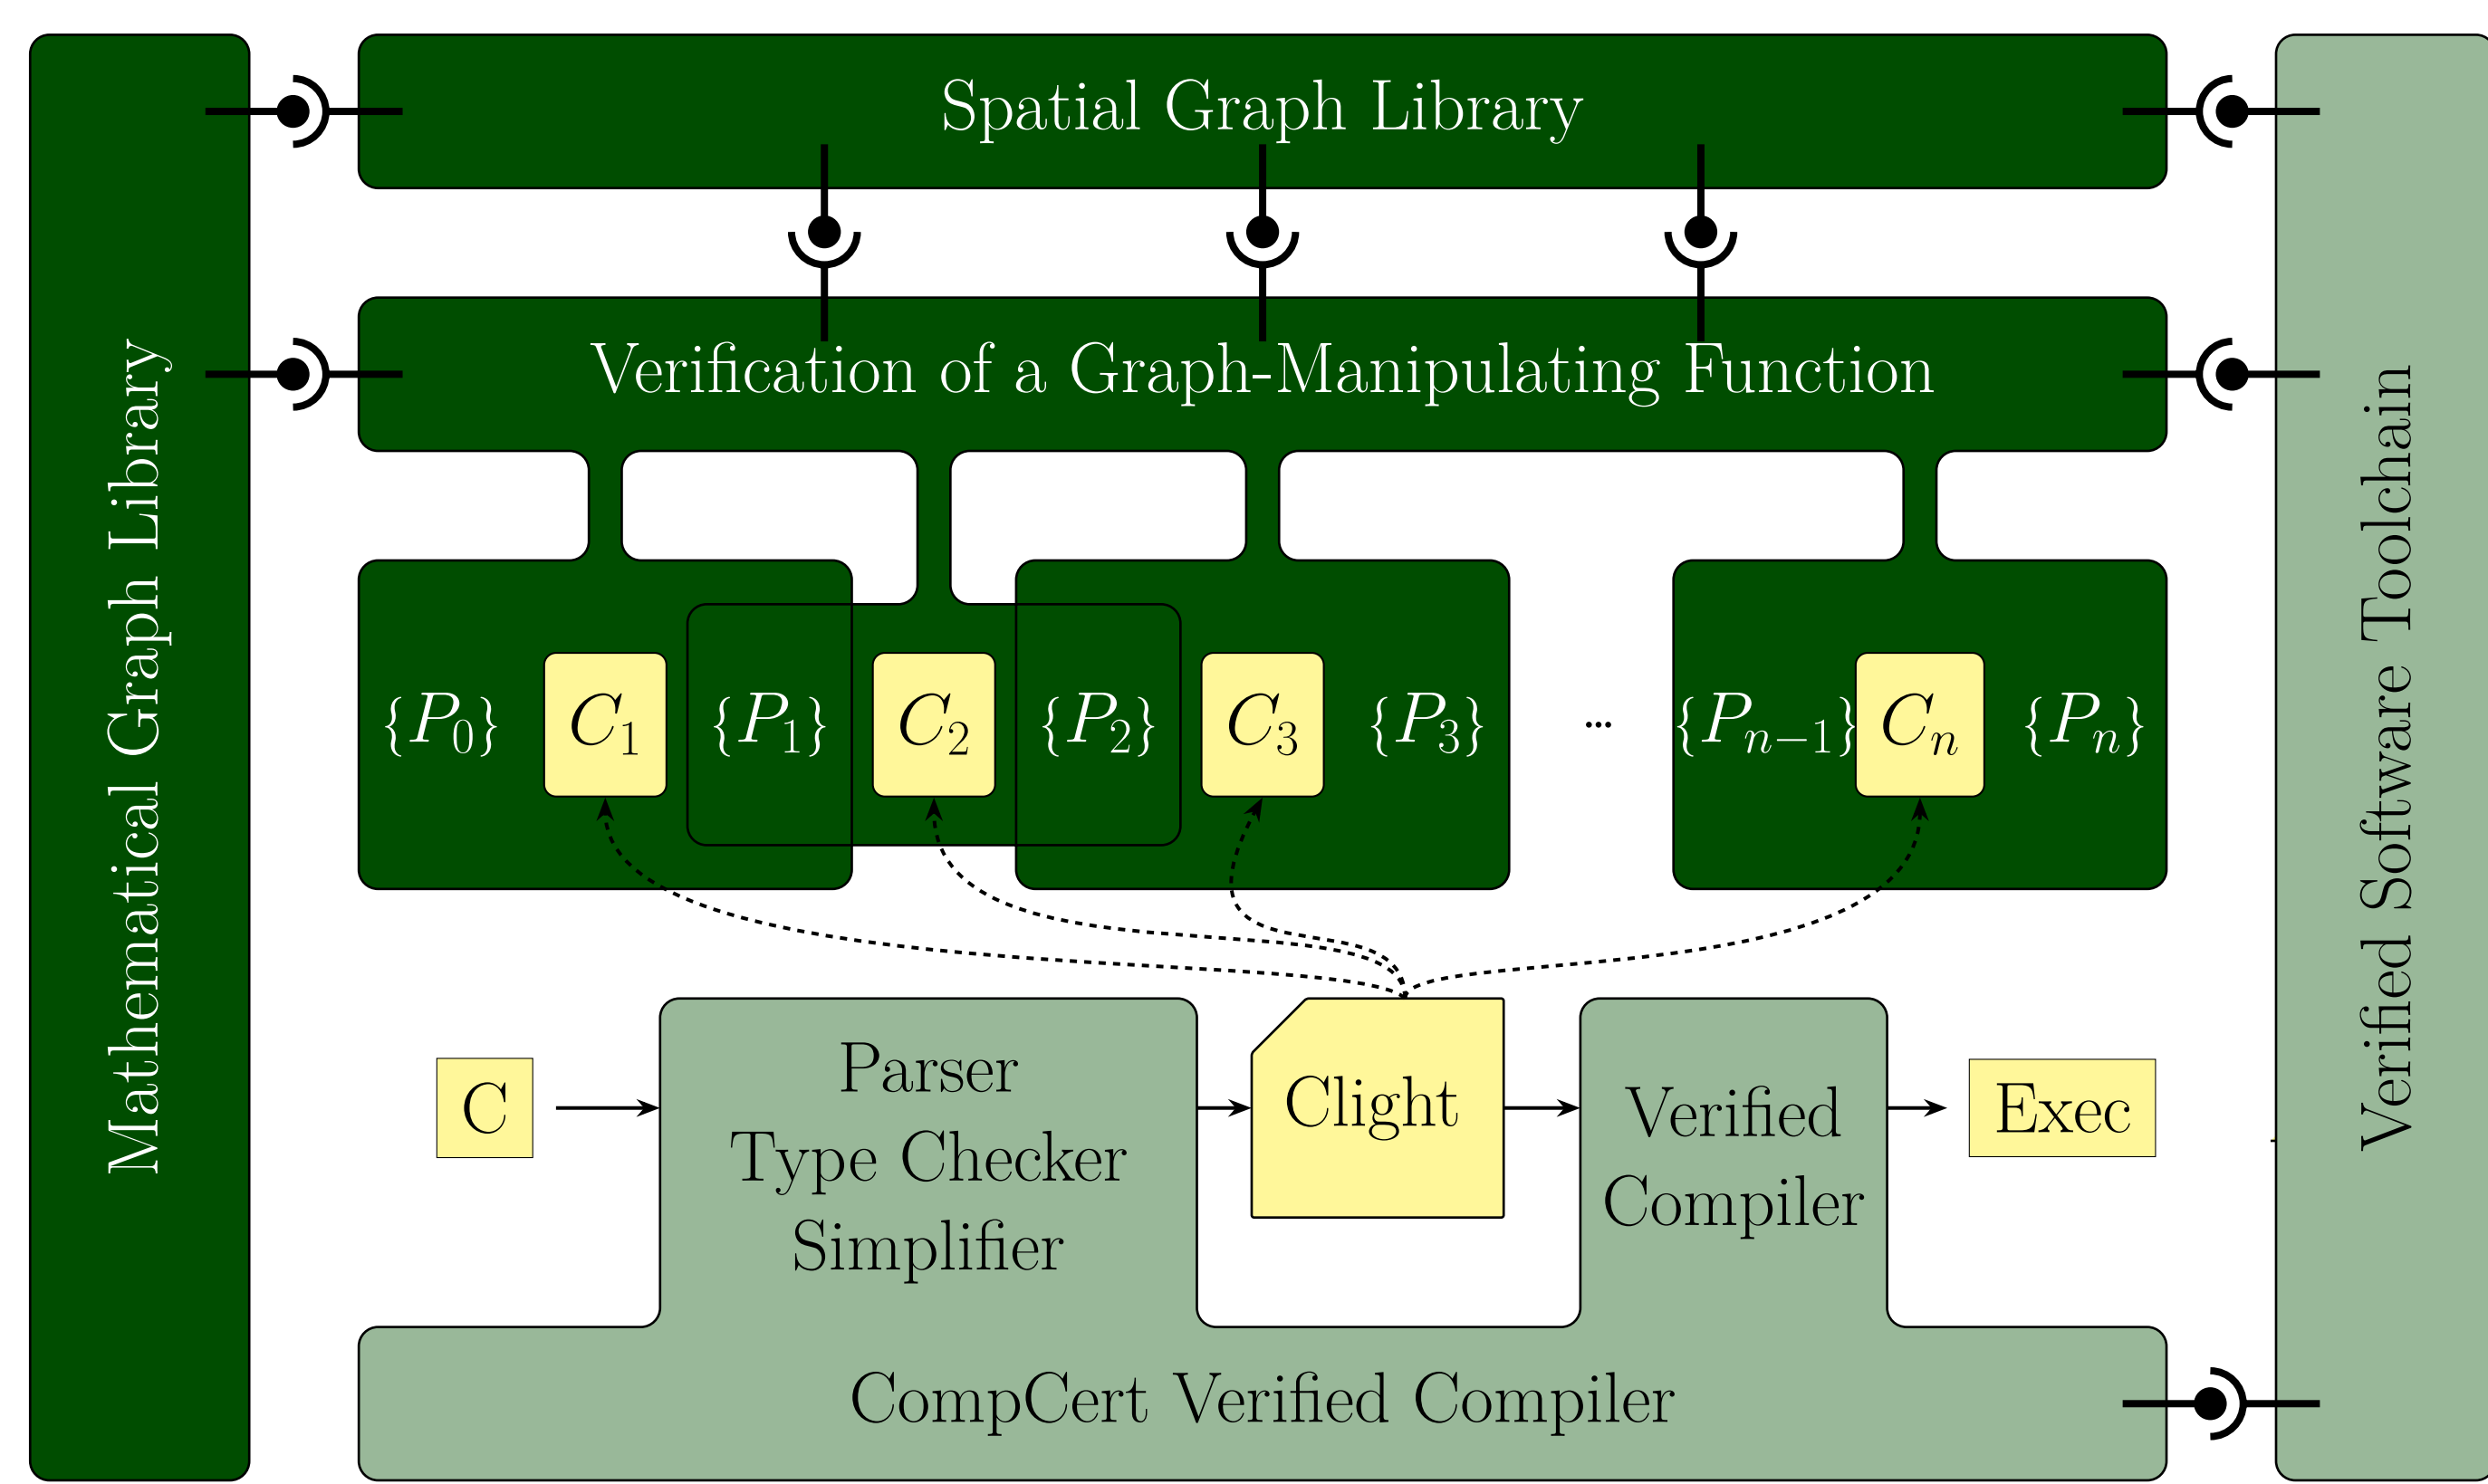
Challenges: We use CompCert C, which is executable and realistic but also has real-world complications. We prove full functional correctness, and not just program safety.

Solution: We use the Verified Software Toolchain [1] and our Mathematical and Spatial Graph Libraries [4] to establish machine-checked correctness at the C level (3 files, 3k LOC). The CompCert compiler [3] then guarantees this correctness on the executable code.

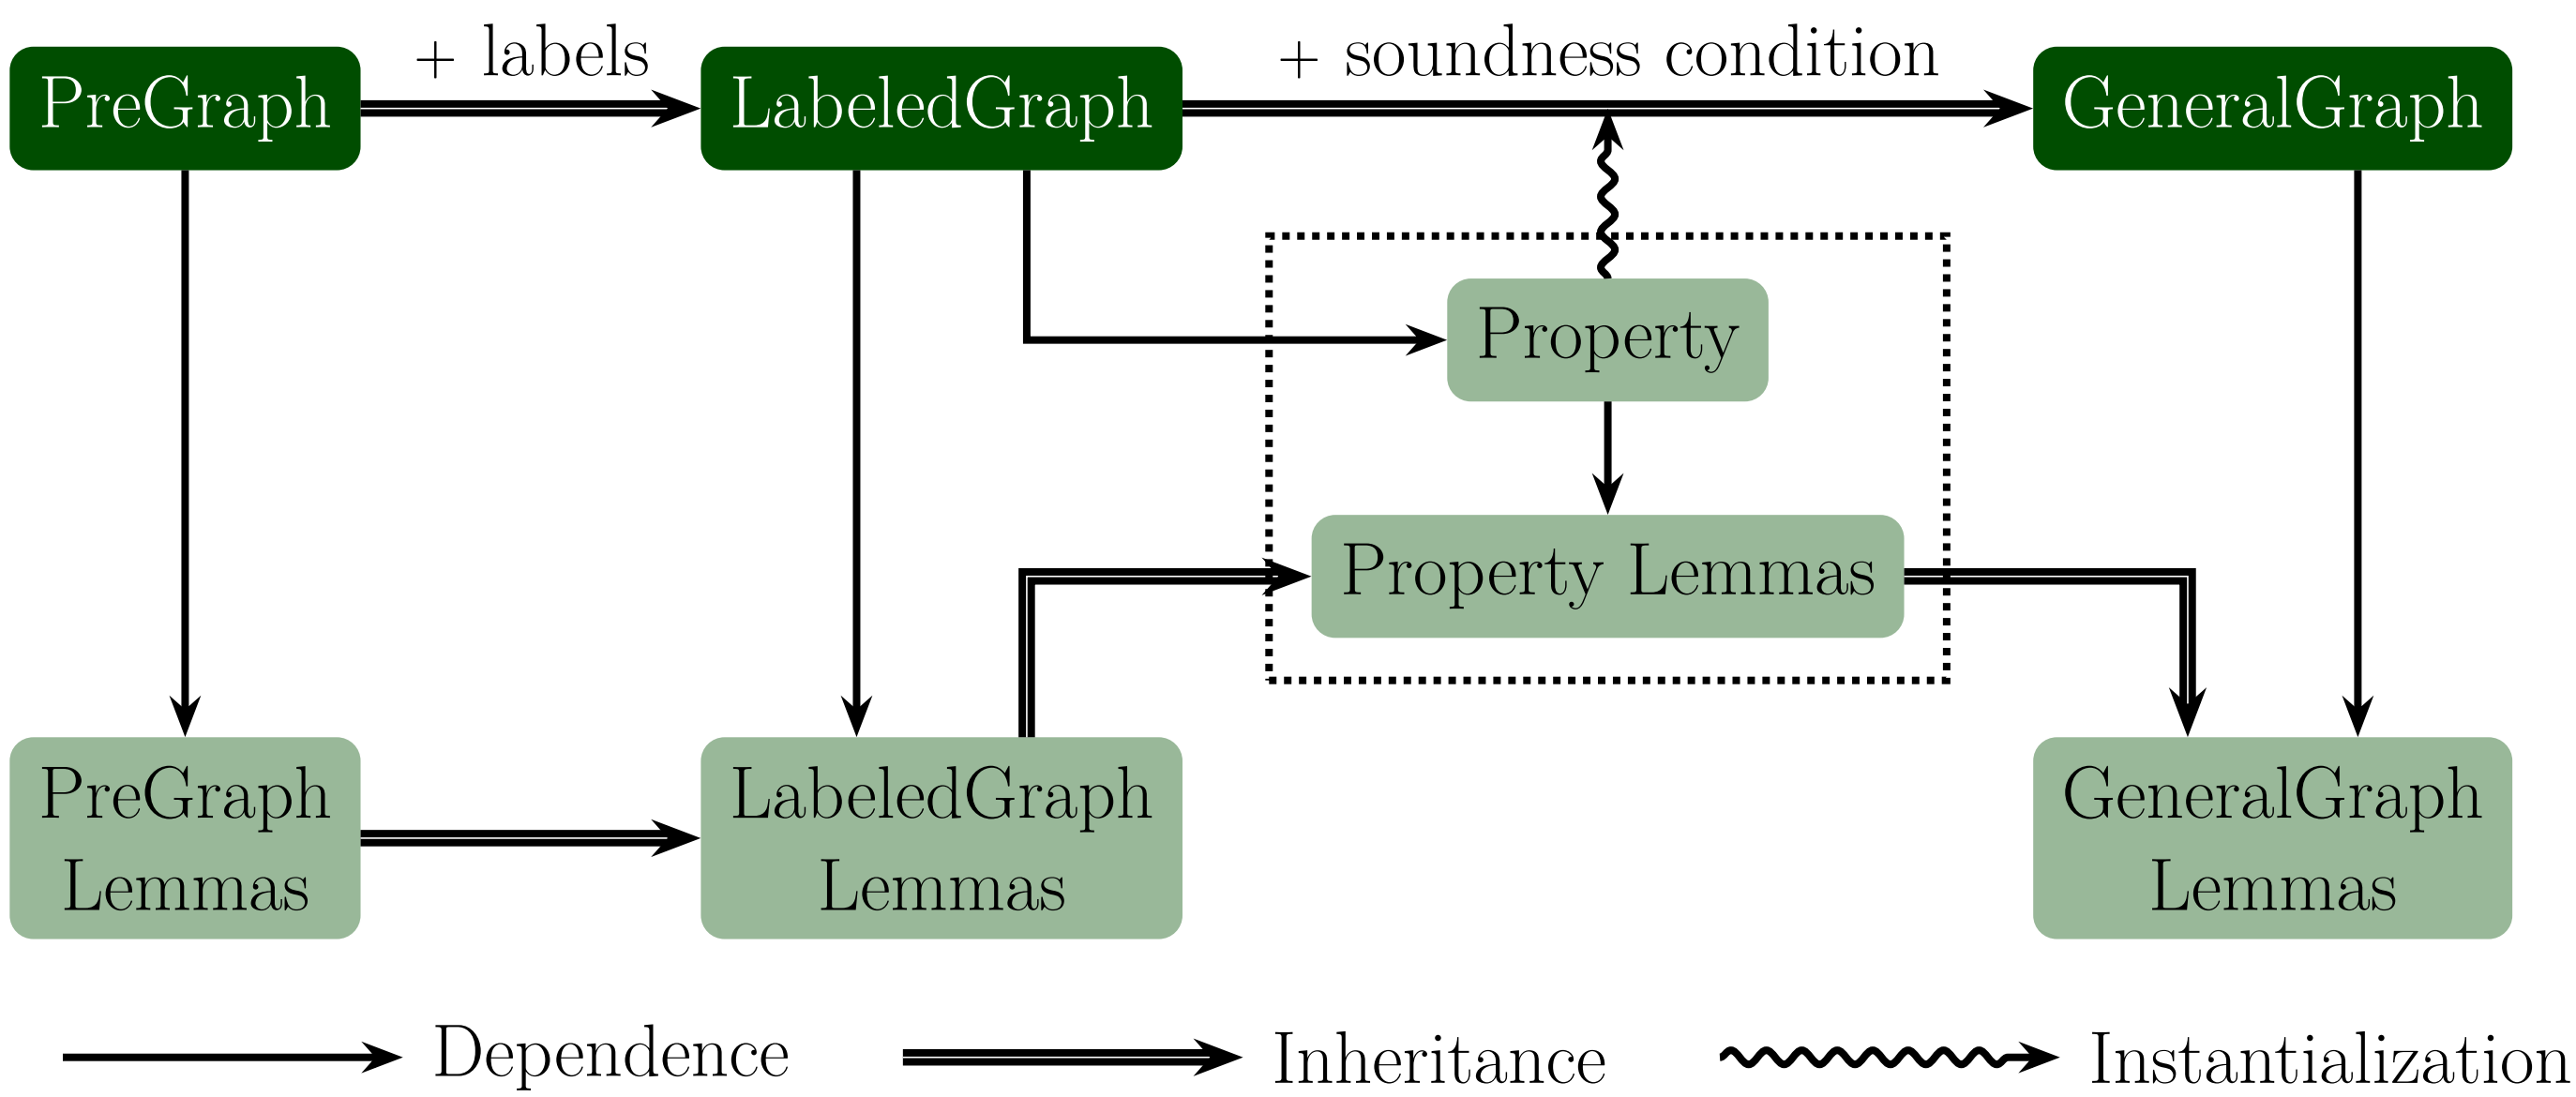
Key findings: The algorithm suffers from potential overflow issues. The precise bound is nontrivial: we show that the intuitive guess fails, and provide a workable refinement.

Workflow

A sketch of how we verify C programs. Note where we integrate with other projects.



The structure of our Mathematical Graph Library. The soundness condition is entirely customizable. Lemmas and properties can be composed, and are automatically inherited.



Instantiating DijkGraph


PreGraph: $VType := Z$, $EType := Z * Z$, $src := fst$, $dst := snd$,
 $\forall v. vvalid(\gamma, v) \Leftrightarrow 0 \leq v < SIZE$,
 $\forall s, d. evalid(\gamma, (s, d)) \Leftrightarrow vvalid(\gamma, s) \wedge vvalid(\gamma, d)$

LabeledGraph: $ELType := Z$, $VLType := list ELType$, $GLType := unit$

GeneralGraph: $FiniteGraph(\gamma) \wedge$
 $\forall i, j. vvalid(\gamma, i) \wedge vvalid(\gamma, j) \Rightarrow$
 $i = j \Rightarrow elabel(\gamma, (i, j)) = 0 \wedge$
 $i \neq j \Rightarrow 0 \leq elabel(\gamma, (i, j)) \leq MAX/SIZE$

Upper Bound on Path Cost

The longest optimal path has $SIZE-1$ links, so say we set $elabel$ ’s upper bound to $\lfloor MAX/(SIZE-1) \rfloor$. Consider the following example, where $MAX = 7$ and $SIZE = 3$, and so $0 \leq elabel(\gamma, e) \leq 3$. The check on line 20 overflows when scanning C ’s neighbors: $(\lfloor MAX/(SIZE-1) \rfloor * (SIZE-1)) + \lfloor MAX/(SIZE-1) \rfloor > MAX$.



Solution: Conservatively set the upper bound for an individual edge to $\lfloor MAX/SIZE \rfloor$. New worst case: line 20 calculates $(\lfloor MAX/SIZE \rfloor * (SIZE-1)) + \lfloor MAX/SIZE \rfloor = MAX$. The true maximum path cost is $\lfloor MAX/SIZE \rfloor * (SIZE-1) = MAX - \lfloor MAX/SIZE \rfloor$.

Code and Specification

```

1 void dijkstra (int graph[SIZE][SIZE], int src,
2               int *dist, int *prev) {
3 // { DijkGraph( $\gamma$ ) }
4   int pq[SIZE];
5   int i, j, u, cost;
6   for (i = 0; i < SIZE; i++) {
7     dist[i] = INF;
8     prev[i] = INF;
9     pq[i] = INF;
10  }
11  dist[src] = 0;
12  pq[src] = 0;
13  prev[src] = src;
14 // { DijkGraph( $\gamma$ )  $\wedge$ 
15     { dijk_correct( $\gamma, src, prev, dist, priq$ ) }
16   while (!pq_emp(pq)) {
17     u = popMin(pq);
18     for (i = 0; i < SIZE; i++) {
19       cost = graph[u][i];
20       if (cost < INF) {
21         if (dist[i] > dist[u] + cost) {
22           dist[i] = dist[u] + cost;
23           prev[i] = u;
24           pq[i] = dist[i];
25         }
26       }
27     }
28 // { DijkGraph( $\gamma$ )  $\wedge$ 
29     {  $\forall dst \in priq. priq[dst] = INF \wedge$ 
30     { dijk_correct( $\gamma, src, prev, dist, priq$ ) }
31     }
32   }
33   return;
34 }


```

$list_rep(\gamma, i) \triangleq data_at\ array\ graph2mat(\gamma)[i]\ list_addr(\gamma, i)$

$graph_rep(\gamma) \triangleq \ast_{vvalid(\gamma, v)} v \mapsto list_rep(\gamma, v)$

$dijk_correct(\gamma, src, prev, dist, priq) \triangleq$
 $\forall dst. dst \in popped(priq) \Rightarrow \exists path. path_correct(\gamma, prev, dist, path) \wedge$
 $path_glob_optimal(\gamma, dist, path) \wedge$
 $path_entirely_in_popped(\gamma, prev, path) \wedge$
 $priq[dst] < INF \Rightarrow let\ m := prev[dst]\ in\ m \in popped(priq) \wedge$
 $\forall m' \in popped(priq). cost(path2m+ :: (m, dst)) \leq$
 $cost(path2m'+ :: (m', dst)) \wedge$
 $priq[dst] = INF \Rightarrow \forall m \in popped(priq). cost(path2m+ :: (m, dst)) = INF$

Key Transformation: Growing the Subgraph



To begin, vertices dst_1 , dst_2 , and dst_3 obey the first, second, and third clauses of the invariant $dijk_correct$ respectively. Vertex u obeys the second clause with minimal cost.

The invariant is broken when relaxing u ’s neighbors, and reestablished thereafter: u now obeys the first clause. Eventually no vertices obey the second clause, and we are done.

References

- [1] Andrew Appel. “Verified Software Toolchain”. In: *NFM 2012, Norfolk, USA, April 3-5, 2012*, p. 2.
- [2] Thomas Cormen et al. *Introduction to Algorithms*. MIT Press, 2009.
- [3] Xavier Leroy. “Formal certification of a compiler back-end or: programming a compiler with a proof assistant”. In: *POPL 2006, Charleston, USA, January 11-13, 2006*, pp. 42–54.
- [4] Shengyi Wang et al. “Certifying graph-manipulating C programs via localizations within data structures”. In: *OOPSLA 2019, Athens, Greece, October 20-25, 2019*, pp. 171:1–171:30.