

## VERIFIED APPROXIMATION ALGORITHMS

ROBIN ESSMANN<sup>a</sup>, TOBIAS NIPKOW<sup>a</sup>, SIMON ROBILLARD<sup>b</sup>, AND UJKAN SULEJMANI<sup>a</sup>

<sup>a</sup> Technische Universität München, Germany

<sup>b</sup> LIRMM, Université de Montpellier & CNRS, Montpellier, France

**ABSTRACT.** We present the first formal verification of approximation algorithms for NP-complete optimization problems: vertex cover, independent set, set cover, center selection, load balancing, and bin packing. We uncover incompletenesses in existing proofs and improve the approximation ratio in one case. All proofs are uniformly invariant based.

### 1. INTRODUCTION

Approximation algorithms for NP-complete problems [Vaz03] are a rich area of research untouched by automated verification. We present the first formal verifications of five classical and one lesser known approximation algorithm. Three of these algorithms had been verified on paper by program verification experts [BM03, BR03]. We found that their claimed invariants need additional conjuncts before they are strong enough to be real invariants. That is, their proofs are incomplete. The other three algorithms only come with informal textbook proofs [KT06].

To put an end to this situation we formalized the correctness proofs of six approximation algorithms for fundamental NP-complete problems in the theorem prover Isabelle/HOL [NK14, NPW02]. We verified (all proofs are online [ENRS20]) the following algorithms and approximation ratios:

- the classic  $k$ -approximation algorithm for minimal vertex covers of rank  $k$  hypergraphs;
- Wei’s [Wei81]  $\Delta$ -approximation algorithm for maximal independent sets of graphs with maximum degree  $\Delta$ ;
- the greedy  $\frac{3}{2}$  (resp. 2) approximation algorithm for the load balancing problem of sorted (resp. unsorted) job loads;
- the greedy 2-approximation algorithm for the center selection problem;
- the greedy  $H(d^*)$ -approximation algorithm for set covers where  $d^*$  is the size of the largest set and  $H(n)$  is the  $n$ th harmonic number;
- the  $\frac{3}{2}$ -approximation algorithm for bin packing by Berghammer and Reuter [BR03].

---

This is a revised version of the conference publication [ENR20]; the sections on Center Selection and Set Cover have been added.

All algorithms are expressed in a simple imperative *WHILE*-language. In each case we show that the approximation algorithm computes a valid solution that is at most a certain factor worse than an optimum solution. The polynomial running time of the approximation algorithm is easy to see in each case.

Isabelle not only helped finding mistakes in pen-and-paper proofs but also encouraged proof refactoring that led to simpler proofs, and in one case, to a stronger result: The invariant given by Berghammer and Müller for Wei’s algorithm [BM03] is sufficient to show that the algorithm has an approximation ratio of  $\Delta + 1$ . We managed to simplify their argument significantly which led to an improved approximation ratio of  $\Delta$ .

Last but not least all our proofs are uniformly invariant-based. This is in contrast to the three textbook proofs from [KT06] that rely on special *ad hoc* arguments about the programs at hand. We show that this is completely unnecessary and that invariant based proofs can even be simpler.

## 2. ISABELLE/HOL AND IMPERATIVE PROGRAMS

Isabelle/HOL is largely based on standard mathematical notation but with some differences and extensions.

Type variables are denoted by  $'a$ ,  $'b$ , etc. The notation  $t :: \tau$  means that term  $t$  has type  $\tau$ . Except for function types  $'a \Rightarrow 'b$ , type constructors follow postfix syntax, e.g.,  $'a \text{ set}$  is the type of sets of elements of type  $'a$ . The image of a function  $f$  over a set  $S$  is denoted by  $f \text{ ` } S$ . Function  $\text{some} :: 'a \text{ set} \Rightarrow 'a$  picks an arbitrary element from a set; the result is unspecified if the set is empty.

The types *nat* and *real* represent the sets  $\mathbb{N}$  and  $\mathbb{R}$ . In this paper we drop the coercion function  $\text{real} :: \text{nat} \Rightarrow \text{real}$ . The set  $\{m..n\}$  is the closed interval  $[m, n]$ .

The Isabelle/HOL distribution comes with a simple implementation of Hoare logic where programs are annotated with pre- and post-conditions and invariants as in this example, where all variables are of type *nat*:

$$\begin{array}{l} \{m = 0 \wedge p = 0\} \\ \text{WHILE } m \neq a \text{ INV } \{p = m*b\} \text{ DO } p := p+b; m := m+1 \text{ OD} \\ \{p = a*b\} \end{array}$$

The box around the program means that it has been verified. All our proofs employ a verification condition generator and essentially reduce to showing the preservation of the invariants.

Behind the scenes, a Hoare triple is just a HOL formula based on a formalized program semantics.  $\{P\} c \{Q\}$  is syntactic sugar for the following partial correctness formula: if execution of  $c$  starting in a state that satisfies  $P$  terminates, then the final state satisfies  $Q$ . All expressions in a Hoare triple (pre- and post-conditions, invariants, tests in loops and conditionals, and right-hand sides of assignments) are arbitrary HOL expressions and can talk about both program variables and elements from the context.

## 3. VERTEX COVER

We verify the proof in [BM03] that the classic greedy algorithm for vertex cover is a 2-approximation algorithm. In fact, we generalize the setup from graphs to hypergraphs. A

hypergraph is simply a set of edges  $E$ , where an edge is a set of vertices of type  $'a$ . A vertex cover for  $E$  is a set of vertices  $C$  that intersects with every edge of  $E$ :

$vc :: 'a \text{ set set} \Rightarrow 'a \text{ set} \Rightarrow bool$   
 $vc \ E \ C = (\forall e \in E. e \cap C \neq \emptyset)$

A matching ( $matching :: 'a \text{ set set} \Rightarrow bool$ ) is a set of pairwise disjoint sets. The following is a key property that relates  $vc$  and  $matching$ :

$finite \ C \wedge matching \ M \wedge M \subseteq E \wedge vc \ E \ C \longrightarrow |M| \leq |C|$

We fix a rank- $k$  hypergraph  $E :: 'a \text{ set set}$  with the assumptions  $\emptyset \notin E$ ,  $finite \ E$  and  $e \in E \longrightarrow finite \ e \wedge |e| \leq k$ . (Via an Isabelle “locale”. We use the same mechanism in all of our proofs.)

We have verified the well known greedy algorithm that computes a vertex cover  $C$  for  $E$ . It keeps picking an arbitrary edge that is not covered by  $C$  yet until all vertices are covered. The final  $C$  has at most  $k$  times as many vertices as any vertex cover of  $E$  (which is essentially optimal [BK10]).

```

{True}
C := ∅; F := E;
WHILE F ≠ ∅ INV {invar C F}
DO C := C ∪ some F; F := F - {e' ∈ F | some F ∩ e' ≠ ∅} OD
{vc E C ∧ (∀ C'. finite C' ∧ vc E C' → |C| ≤ k * |C'|)}

```

where  $invar$  is the following invariant:

$invar :: 'a \text{ set} \Rightarrow 'a \text{ set set} \Rightarrow bool$

$invar \ C \ F =$

$(F \subseteq E \wedge vc \ (E - F) \ C \wedge finite \ C \wedge (\exists M. inv\_matching \ C \ F \ M))$

$inv\_matching \ C \ F \ M =$

$(matching \ M \wedge M \subseteq E \wedge |C| \leq k * |M| \wedge (\forall e \in M. \forall f \in F. e \cap f = \emptyset))$

The key step in the program proof is that the invariant is invariant:

**Lemma 3.1.**  $F \neq \emptyset \wedge invar \ C \ F \longrightarrow$

$invar \ (C \cup some \ F) \ (F - \{e' \in F \mid some \ F \cap e' \neq \emptyset\})$

Our invariant is stronger than the one in [BM03] which lacks  $F \subseteq E$ . Indeed, without this property, the claimed invariant is not invariant (as acknowledged by Müller-Olm).

#### 4. INDEPENDENT SET

As in the previous section, a graph is a set of edges. An independent set of a graph  $E$  is a subset of its vertices such that no two vertices are adjacent.

$iv :: 'a \text{ set set} \Rightarrow 'a \text{ set} \Rightarrow bool$

$iv \ E \ S = (S \subseteq \bigcup E \wedge (\forall v_1 \ v_2. v_1 \in S \wedge v_2 \in S \longrightarrow \{v_1, v_2\} \notin E))$

We fix a finite graph  $E :: 'a \text{ set set}$  such that all edges of  $E$  are sets of cardinality 2. The set of vertices  $\bigcup E$  is denoted  $V$ , and the maximum degree (number of neighbors) of any vertex in  $V$  is denoted  $\Delta$ . We show that the greedy algorithm proposed by Wei is a  $\Delta$ -approximation algorithm. The proof is inspired by one given in [BM03]. In particular, the proof relies on an auxiliary variable  $P$ , which is not needed for the execution of the

algorithm, but is used for bookkeeping in the proof. In [BM03],  $P$  is initially a program variable and is later removed from the program and turned into an existentially quantified variable in the invariant. We directly use the latter representation.

```

{ True }
S := ∅; X := ∅;
WHILE X ≠ V INV { ∃ P. inv_partition S X P }
DO x := some (V - X); S := S ∪ {x}; X := X ∪ neighbors x ∪ {x} OD
{ iv E S ∧ (∀ S'. iv E S' → |S'| ≤ |S| * Δ) }

```

To keep the size of definitions manageable, we split the invariant in two. The first part is not concerned with  $P$ , but suffices to prove the functional correctness of the algorithm, i.e., that it outputs an independent set of the graph:

$inv\_iv :: 'a\ set \Rightarrow 'a\ set \Rightarrow bool$

$inv\_iv\ S\ X =$

$(iv\ E\ S \wedge X \subseteq V \wedge (\forall v_1 \in V - X. \forall v_2 \in S. \{v_1, v_2\} \notin E) \wedge S \subseteq X)$

This invariant is taken almost verbatim from [BM03], except that in [BM03] it says that  $S$  is an independent set of the subgraph generated by  $X$ . This is later used to show that the  $x$  picked at each iteration from  $V - X$  is not already in  $S$ . Defining subgraphs adds unnecessary complexity to the invariant. We simply state  $S \subseteq X$ , together with the fact that  $S$  is an independent set of the whole graph.

We now extend the invariant with properties of the auxiliary variable  $P$ .

$inv\_partition :: 'a\ set \Rightarrow 'a\ set \Rightarrow 'a\ set\ set \Rightarrow bool$

$inv\_partition\ S\ X\ P =$

$(inv\_iv\ S\ X \wedge$

$\bigcup P = X \wedge (\forall p \in P. \exists s \in V. p = \{s\} \cup neighbors\ s) \wedge |P| = |S| \wedge finite\ P)$

We can view the set  $P$  as an auxiliary program variable. In order to satisfy the invariant,  $P$  would be initially empty and the loop body would include the assignment  $P := P \cup \{neighbors\ x \cup \{x\}\}$ . Intuitively,  $P$  contains the sets of vertices that are added to  $X$  at each iteration (or more precisely, an over-approximation, since some vertices in  $neighbors\ x$  may have been added to  $X$  in a previous iteration). Instead of adding an unnecessary variable to the program, we only use the existentially quantified invariant. The assignments described above correspond directly to instantiations of the quantifier that are needed to solve proof obligations. This is illustrated with the following lemma, which corresponds to the preservation of the invariant:

**Lemma 4.1.**  $(\exists P. inv\_partition\ S\ X\ P) \wedge x \in V - X \longrightarrow$

$(\exists P'. inv\_partition\ (S \cup \{x\})\ (X \cup neighbors\ x \cup \{x\})\ P')$

The existential quantifier in the antecedent yields a witness  $P$ . After instantiating the quantifier in the succedent with  $P \cup \{neighbors\ x \cup \{x\}\}$ , the goal can be solved straightforwardly. Finally, the following lemma combines the invariant and the negated post-condition to prove the approximation ratio:

**Lemma 4.2.**  $inv\_partition\ S\ V\ P \longrightarrow (\forall S'. iv\ E\ S' \longrightarrow |S'| \leq |S| * \Delta)$

To prove it, we observe that any set  $p \in P$  consists of a vertex  $x$  and its neighbors, therefore an independent set  $S'$  can contain at most  $\Delta$  of the vertices in  $p$ , thus  $|S'| \leq |P| * \Delta$ . Furthermore, as indicated by the invariant,  $|P| = |S|$ .

Compared to the proof in [BM03], our invariant describes the contents of the set  $P$  more precisely, and thus yields a better approximation ratio. In [BM03], the invariant merely indicates that  $X = \bigcup P$ , together with two cardinality properties:  $\forall p \in P. |p| \leq \Delta + 1$  and  $|P| \leq |S|$ . Taken with the negated post-condition, this invariant can be used to show that for any independent set  $S'$ , we have  $|S'| \leq |S| * (\Delta + 1)$ . The proof of this lemma makes use of the following (in)equalities:  $|S'| \leq |V|$ ,  $|V| = |\bigcup P|$ ,  $|\bigcup P| \leq |P| * (\Delta + 1)$  and finally  $|P| * (\Delta + 1) \leq |S| * (\Delta + 1)$ . Note that this only relies on the trivial fact that an independent set cannot contain more vertices than the graph. By contrast, our own argument takes into account information regarding the edges of the graph.

Although this proof results in a weaker approximation ratio than our own, it yields a useful insight: an approximation ratio is given by the cardinality of the largest set  $p \in P$  (i.e., the largest number of vertices added to  $X$  during any given iteration). In the worst case, this is equal to  $\Delta + 1$ , but in practice the number may be smaller. This suggests a variant of the algorithm that stores that value in a variable  $r$ , as described in [BM03]. At every iteration, the variable  $r$  is assigned the value  $\max r \{ |x| \cup \text{neighbors } x - X |$ . Ultimately, the algorithm returns both the independent set  $S$  and the value  $r$ , with the guarantee that  $|S'| \leq |S| * r$  for any independent set  $S'$ .

We also formalized this variant and proved the aforementioned property. The proof follows the idea outlined above, but does away with the variable  $P$  entirely: instead, the invariant simply maintains that  $\text{inv. } |S| * r \wedge |X| \leq |S| * r$ , and the proof of preservation is adapted accordingly. Indeed, this demonstrates that the argument used in [BM03] does not require an auxiliary variable nor an existentially quantified invariant. For the proof of the approximation ratio  $\Delta$ , a similar simplification is not as easy to obtain, because the argument relies on a global property of the graph (a constraint that edges impose on independent sets) that is not easy to summarize in an inductive invariant.

So far, we have only considered an algorithm where the vertex  $x$  is picked non-deterministically. An obvious heuristic is to pick, at every iteration, the vertex with the smallest number of neighbors among  $V - X$ . Halldórsson and Radhakrishnan [HR97] prove that this heuristic achieves an approximation ratio of  $(\Delta + 2) / 3$ . They proceed by considering the sequence of graph reductions corresponding to the execution of the algorithm. Formalizing this argument would require an inductive invariant to keep track of the sequences of values of  $X$  and  $x$  at each iteration and the relations between those values, providing a much more precise record than the auxiliary variable  $P$ . In addition, the proof itself is far more complex than the arguments presented here, relying on case analysis for different types of graphs. This is beyond the scope of our paper.

## 5. LOAD BALANCING

Our starting point for the load balancing problem is [KT06, Chapter 11.1]. We need to distribute  $n :: \text{nat}$  jobs on  $m :: \text{nat}$  machines with  $0 < m$ . A job  $j \in \{1..n\}$  has a load  $t(j) :: \text{nat}$ . Variables  $m$ ,  $n$ , and  $t$  are fixed throughout this section. A solution is described by a function  $A$  that maps machines to sets of jobs:  $k \in \{1..m\}$  has job  $j$  assigned to it iff  $j \in A(k)$ . The sum of job loads on a machine is given by a function  $T$  that is derived from  $t$  and  $A$ :  $T k = (\sum_{j \in A k} t j)$ . Predicate  $lb$  defines when  $T$  and  $A$  are a partial solution for  $j \leq n$  jobs:

$lb :: (nat \Rightarrow nat) \Rightarrow (nat \Rightarrow nat\ set) \Rightarrow nat \Rightarrow bool$

$lb\ T\ A\ j =$

$((\forall x \in \{1..m\}. \forall y \in \{1..m\}. x \neq y \longrightarrow A\ x \cap A\ y = \emptyset) \wedge$   
 $(\bigcup_{x \in \{1..m\}} A\ x) = \{1..j\} \wedge (\forall x \in \{1..m\}. (\sum_{y \in A\ x} t\ y) = T\ x))$

It consists of three conjuncts. The first ensures that the sets returned by  $A$  are pairwise disjoint, thus, no job appears in more than one machine. The second conjunct ensures that every job  $x \in \{1..j\}$  is contained in at least one machine. It also ensures that only jobs  $\{1..j\}$  have been added. The final conjunct ensures that  $T$  is correctly defined to be the total load on a machine. To ensure that jobs are distributed evenly, we need to consider the machine with maximum load. This load is referred to as the *makespan* of a solution:

$makespan :: (nat \Rightarrow nat) \Rightarrow nat$

$makespan\ T = Max\ (T\ ' \{1..m\})$

The greedy approximation algorithm outlined in [KT06] relies on the ability to determine the machine  $k \in \{1..m\}$  that has a minimum combined load. As the goal is to approximate the optimum in polynomial time, a linear scan through  $T$  suffices to find the machine with minimum load. However, other methods may be considered to further improve time complexity. To determine the machine with minimum load, we will use the following function:

$min\_arg :: (nat \Rightarrow nat) \Rightarrow nat \Rightarrow nat$

$min\_arg\ T\ 0 = 1$

$min\_arg\ T\ (x + 1) =$

$(let\ k = min\_arg\ T\ x\ in\ if\ T\ (x + 1) < T\ k\ then\ x + 1\ else\ k)$

We will focus on the approximation factor of  $\frac{3}{2}$ , which can be proved if the job loads are assumed to be sorted in descending order. The proof for the approximation factor of 2 if jobs are unsorted is very similar and we describe the differences at the end. We say that  $j$  jobs are sorted in descending order if *sorted* holds:

$sorted :: nat \Rightarrow bool$

$sorted\ j = (\forall x \in \{1..j\}. \forall y \in \{1..x\}. t\ x \leq t\ y)$

Below we prove the following conditional Hoare triple that expresses the approximation factor and functional correctness of the algorithm given in [KT06]:

$sorted\ n \longrightarrow$   
 $\{True\}$   
 $T := (\lambda_. 0); A := (\lambda_. \emptyset); j := 0;$   
 $WHILE\ j < n\ INV\ \{inv_2\ T\ A\ j\}$   
 $DO\ i := min\_arg\ T\ m; j := j + 1;$   
 $\quad A := A(i := A(i) \cup \{j\}); T := T(i := T(i) + t(j))$   
 $OD$   
 $\{lb\ T\ A\ n \wedge$   
 $(\forall T'\ A'. lb\ T'\ A'\ n \longrightarrow makespan\ T \leq 3 / 2 * makespan\ T')\}$

Property *sorted*  $n$  does not need to be part of the precondition because it does not mention any program variable. Therefore we can make *sorted*  $n$  an assumption of the whole Hoare triple, which simplifies the proof. The notation  $f(a := b)$  denotes an updated version of function  $f$  that maps  $a$  to  $b$  and behaves like  $f$  otherwise. Thus an assignment  $f := f(i := b)$  is nothing but the conventional imperative array update notation  $f[i] := b$ .

Functional correctness follows because each iteration extends a partial solution for  $j$  jobs to one for  $j + 1$  jobs:

**Lemma 5.1.**  $lb\ T\ A\ j \wedge x \in \{1..m\} \longrightarrow$   
 $lb\ (T(x := T\ x + t\ (j + 1)))\ (A(x := A\ x \cup \{j + 1\}))\ (j + 1)$

Moreover, it is easy to see that the initialization establishes  $lb\ T\ A\ j$ .

To prove the approximation factor in both the sorted and unsorted case, the following lower bound is important:

**Lemma 5.2.**  $lb\ T\ A\ j \longrightarrow (\sum_{x=1}^j t\ x) / m \leq makespan\ T$

This is a result of  $\sum_{x=1}^m T(x) = \sum_{x=1}^j t(x)$  together with this general property of sums: *finite*  
 $A \wedge A \neq \emptyset \longrightarrow (\sum_{a \in A} f\ a) \leq |A| * Max\ (f\ 'A)$ .

A similar observation applies to individual jobs. Any job must be a lower bound on some machine, as it is assigned to one and, by extension, it must also be a lower bound of the makespan:

**Lemma 5.3.**  $lb\ T\ A\ j \longrightarrow Max_0\ (t\ '\{1..j\}) \leq makespan\ T$

As any job load is a lower bound on the makespan over the machines, the job with maximum load must also be a lower bound. Note that  $Max_0$  returns 0 for the empty set.

When jobs are sorted in descending order, a stricter lower bound for an individual job can be established. We observe that an added job is at most as large as the jobs preceding it. Therefore, if a machine contains at least two jobs, this added job is only at most *half* as large as the makespan. We can use this observation by assuming the machines to be filled with more than  $m$  jobs, as this will ensure that some machine must contain at least two jobs.

**Lemma 5.4.**  $lb\ T\ A\ j \wedge m < j \wedge sorted\ j \longrightarrow 2 * t\ j \leq makespan\ T$

Note that this lower bound only holds if there are strictly more jobs than machines. One must, however, also consider how the algorithm behaves in the other case. One may intuitively see that the algorithm will be able to distribute the jobs such that every machine will only have at most one job assigned to it, making the algorithm trivially optimal. To prove this, we need to show the following behavior of *min\_arg*:

**Lemma 5.5.**

(1)  $x \in \{1..m\} \wedge T\ x = 0 \longrightarrow T\ (min\_arg\ T\ m) = 0$

(2)  $x \in \{1..m\} \wedge T\ x = 0 \longrightarrow min\_arg\ T\ m \leq x$

Both properties can be shown by induction on the number of machines  $m$ .

As the proof in [KT06] is only informal, Kleinberg and Tardos do not provide any loop invariant. We propose the following invariant for sorted jobs:

$inv_2 :: (nat \Rightarrow nat) \Rightarrow (nat \Rightarrow nat\ set) \Rightarrow nat \Rightarrow bool$

$inv_2\ T\ A\ j =$

$(lb\ T\ A\ j \wedge j \leq n \wedge$

$(\forall T'\ A'. lb\ T'\ A'\ j \longrightarrow makespan\ T \leq 3 / 2 * makespan\ T') \wedge$

$(\forall x > j. T\ x = 0) \wedge (j \leq m \longrightarrow makespan\ T = Max_0\ (t\ '\{1..j\})))$

The final two conjuncts relate to the trivially optimal behavior of the algorithm if  $j \leq m$ . The penultimate conjunct shows that only as many machines can be occupied as there are available jobs. The final conjunct ensures that every job is distributed on its own machine, making the makespan equivalent to the job with maximum load.

It should be noted that if the makespan is sufficiently large, an added job may not increase the makespan at all, as the machine with minimum load combined with the job may not exceed the previous makespan. As such, we will also consider the possibility that an added job can simply be ignored without affecting the overall makespan.

**Lemma 5.6.**  $\text{makespan } (T(x := T x + y)) \neq T x + y \longrightarrow \text{makespan } (T(x := T x + y)) = \text{makespan } T$

To make use of this observation, we need to be able to relate the makespan of a solution with the added job to the makespan of a solution without it. One can easily show the following by removing  $j + 1$  from the solution:

**Lemma 5.7.**  $\text{lb } T A (j + 1) \longrightarrow (\exists T' A'. \text{lb } T' A' j \wedge \text{makespan } T' \leq \text{makespan } T)$

We can now prove the preservation of  $\text{inv}_2$ . Let  $i = \text{min\_arg } T m$  be the machine with minimum load. We define:

$$T_g := T (i := T(i) + t(j + 1)) \quad A_g := A (i := A(i) \cup \{j + 1\})$$

We begin with a case distinction. If  $j + 1 \leq m$ , we can make use of the additional conjuncts to prove the trivially optimal behavior. We first note *in-range*:  $j + 1 \in \{1..m\}$ . Moreover, from the penultimate conjunct,  $T(j + 1) = 0$ . Combining this with Lemma 5.5.1, we can see that  $T(i) = 0$ . Therefore  $T_g(i) = t(j + 1)$  and with the final conjunct of the assumed invariant, the makespan of  $T_g$  remains equivalent to the job with maximum load. To prove that the penultimate conjunct is preserved, we again use *in-range*,  $T(j + 1) = 0$ , and Lemma 5.5.2 to prove that  $i \leq j + 1$ . Moreover,  $T_g$  only differs from  $T$  by the modification of machine  $i$ . Thus, the penultimate conjunct for  $j + 1$  jobs is preserved as well. From Lemma 5.3 we can then see that, as the makespan of  $T_g$  is equivalent to the job with maximum load, it must be trivially optimal. Functional correctness can be shown using Lemma 5.1, and proving the preservation of the remaining conjunct is trivial. We now come to the case  $j + 1 > m$ . We first show that the penultimate conjunct is preserved (the final conjunct can be ignored, as  $\neg j + 1 \leq m$ ). This follows from the correctness of *min\_arg*, as the index returned by it has to be in  $\{1..m\}$  as long as  $m > 0$ . Therefore, we can simply show this from the penultimate conjunct of the assumed invariant. We now come to the proof of the approximation factor:

$$\forall T' A'. \text{lb } T' A' (j + 1) \longrightarrow \text{makespan } T_g \leq 3 / 2 * \text{makespan } T'$$

To prove it, we fix  $T_1$  and  $A_1$  such that  $\text{lb } T_1 A_1 (j + 1)$ . Using Lemma 5.7, one can now obtain  $T_0$  and  $A_0$  such that  $\text{lb } T_0 A_0 j$  and *MK*:  $\text{makespan } T_0 \leq \text{makespan } T_1$ . From the assumed loop invariant, we can now show:

$$\begin{aligned} \text{makespan } T &\leq \frac{3}{2} \text{makespan } T_0 && \text{by } \text{inv}_2\text{-def} \\ &\leq \frac{3}{2} \text{makespan } T_1 && \text{by } \text{MK} \end{aligned}$$

To prove the makespan for  $j + 1$  jobs, there are now two cases to consider: The added job  $j + 1$  contributes to the makespan or it does not. The case in which it does not can be shown by combining the previous calculation with Lemma 5.6. For the first case, we may then assume that  $\text{makespan } T_g = T(i) + t(j + 1)$ . Like in Lemma 5.2, we note that *sum-eq*:  $(\sum_{x=1}^m T x) = (\sum_{x=1}^j t x)$ . Moreover, *min-avg*:  $m * T (\text{min\_arg } T m) \leq (\sum_{i=1}^m T i)$ .



This allows us to calculate the following lower bound for  $T(i)$ :

$$\begin{aligned}
m * T(i) &\leq \sum_{i=1}^m T(i) = \sum_{i=1}^j t(i) && \text{by } \textit{min-avg} \text{ and } \textit{sum-eq} \\
\iff T(i) &\leq \frac{\sum_{i=1}^j t(i)}{m} && \text{because } m > 0 \\
&\leq \textit{makespan } T_0 \leq \textit{makespan } T_1 && \text{by Lemma 5.2 and } \textit{MK}
\end{aligned}$$

From Lemma 5.4 we can also show that  $t(j + 1)$  is a lower bound for  $\frac{1}{2}$  of the makespan of  $T_1$ . Therefore:

$$\begin{aligned}
\textit{makespan } T_g = T(i) + t(j + 1) &\leq \textit{makespan } T_1 + \frac{\textit{makespan } T_1}{2} \\
&= \frac{3}{2} \textit{makespan } T_1
\end{aligned}$$

The proof of functional correctness and remaining conjuncts is again trivial.

Let us now consider the unsorted case where one can still show an approximation factor of 2. The algorithm is identical but the invariant is simpler:

$\textit{inv}_1 \ T \ A \ j =$

$$(\textit{lb } T \ A \ j \wedge j \leq n \wedge (\forall T' \ A'. \textit{lb } T' \ A' \ j \longrightarrow \textit{makespan } T \leq 2 * \textit{makespan } T'))$$

The proof for this invariant is a simpler version of the proof above: We do not need the initial case distinction (case  $j + 1 \leq m$  need not be considered separately), and we use Lemma 5.3 instead of Lemma 5.4 to obtain a bound for  $t(j + 1)$ .

## 6. CENTER SELECTION

This section is based on [KT06, Chapter 11.2]. Given a finite, non-empty set of sites  $S :: 'a \textit{ set}$  (i.e., points) in a metric space, our objective is to select an optimal set of centers  $C \subseteq S$  of size  $k > 0$ , such that

$\textit{radius } C = \textit{Max } (\textit{distance } C \ ' \ S)$

is minimized, where  $\textit{distance } C \ s = \textit{Min } (\textit{dist } s \ ' \ C)$ .

We call a site  $s$  a *candidate* (with respect to some  $r$ ) if  $\textit{distance } C \ s > 2r$  and call it *included* if  $\textit{distance } C \ s \leq 2r$  (i.e., if it already lies within  $2r$  of some center). Now consider the following two observations for the optimal radius  $r^*$ :

- (1) Selecting as a center a candidate site  $s$  w.r.t.  $r^*$  guarantees that at least one more site will now lie within  $2r^*$  of a center—namely  $s$  itself—and rightfully discards sites which are already included.
- (2) Should there exist a candidate site w.r.t.  $r^*$ , then a furthest site will also be a candidate (w.r.t.  $r^*$ ).

The strategy is then to construct  $C$  by repeatedly selecting a furthest site. Of course, it's not immediately clear why this approach should work for the  $2r^*$  boundary but not, say, for  $1.5r^*$ . We resolve this matter in the proof's details.

Before we proceed, we note that our formal proof diverges from that of Kleinberg and Tardos. While they begin by constructing an algorithm around the first observation and proceed to establish, rather informally, a semantic correspondence between it and the

algorithm presented in this paper, we follow a more direct approach—that is, we only reason about the presented algorithm. This reduces the first algorithm to a pedagogical tool.

The algorithm repeatedly selects a site that lies furthest from the set of all hitherto selected centers. The following Hoare triple expresses functional correctness and the approximation factor of 2.

$$\boxed{
\begin{array}{l}
\{k \leq |S|\} \\
C := (\text{SOME } s. s \in S); \\
\text{WHILE } |C| < k \text{ INV } \{\text{invar } C\} \text{ DO} \\
\quad C := C \cup \{\text{furthest\_from } C\}; \\
\text{OD} \\
\{|C| = k \wedge (\forall C'. 0 < |C'| \wedge |C'| \leq k \longrightarrow \text{radius } C \leq 2 * \text{radius } C')\}
\end{array}
}$$

where  $\text{furthest\_from } C = (\text{SOME } s. s \in S \wedge \text{distance } C s = \text{Max } (\text{distance } C ' S))$ .

We now present the proof, beginning with the invariant. To motivate the choice of invariant, consider the special case when there exists a candidate site (w.r.t. some fixed  $r$ ) in each loop iteration. Let  $r$  be the radius of an arbitrary set,  $C$  be the set of previously selected centers,  $s \in S$  the site to be added to  $C$ , and  $s' \in S$  a candidate site w.r.t.  $r$ . (Recall: candidate means  $\text{distance } C s' > 2 * r$ ). By choice of  $s$ , we have  $\text{distance } C s \geq \text{distance } C s'$  and therefore  $\text{distance } C s > 2 * r$ . Since  $\text{distance } C s = \text{Min } (\text{dist } C ' S)$ , we have  $\forall c \in C. \text{dist } c s \geq \text{distance } C s > 2 * r$ . Inductively, it follows (which we prove later) that  $\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow \text{dist } c_1 c_2 > 2 * r$ .

Now consider the case when in some loop iteration no more candidate sites exist, i.e., all sites are included. As before, let  $C$  be the set of previously selected centers. By assumption, we have  $\forall s \in S. \text{distance } C s \leq 2 * r$ .

This case distinction forms the invariant, which is defined as:

$$\begin{array}{l}
\text{invar } C = \\
(C \neq \emptyset \wedge |C| \leq k \wedge C \subseteq S \wedge \\
(\forall C'. (\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow 2 * \text{radius } C' < \text{dist } c_1 c_2) \vee \\
(\forall s \in S. \text{distance } C s \leq 2 * \text{radius } C')))
\end{array}$$

It is obvious that the invariant holds initially for  $C = \{\text{SOME } s. s \in S\}$ . Before we prove that the invariant holds in each iteration, we first prove a useful lemma.

**Lemma 6.1.**  $(\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow x < \text{dist } c_1 c_2) \wedge$   
 $x < \text{distance } C s \wedge \text{finite } C \wedge C \neq \emptyset \longrightarrow$   
 $(\forall c_1 \in C \cup \{s\}. \forall c_2 \in C \cup \{s\}. c_1 \neq c_2 \longrightarrow x < \text{dist } c_1 c_2)$

*Proof.* The case  $c_1 \in C \wedge c_2 \in C$  is true by assumption and the case  $c_1 = s \wedge c_2 = s$  is vacuously true. Assume  $c_1 \in C \wedge c_2 = s$ . Then

$$\begin{array}{ll}
x < \text{distance } C c_2 & \text{by } \text{distance } C s > x \\
\leq \text{dist } c_2 c_1 & \text{by definition and } C \neq \{\} \text{ and } \text{finite } C \\
= \text{dist } c_1 c_2 & \text{by } \text{dist\_commute}
\end{array}$$

By symmetry, the case  $c_1 = s \wedge c_2 \in C$  is also true.  $\square$

Now we prove the preservation of the invariant. Assuming  $\text{invar } C$ , we show  $\text{invar } (C \cup \{s\})$  where  $s = \text{furthest\_from } C$ . The first three conjuncts hold trivially.

*Proof.* Let  $C'$  be an arbitrary set and let  $r = \text{radius } C'$ . To prove the fourth conjunct, we distinguish two cases:

**Case 1**  $s$  is a candidate site w.r.t.  $r$ . This negates  $\forall s' \in S. \text{distance } C s' \leq 2 * r$ . Consequently,  $\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow 2 * r < \text{dist } c_1 c_2$  must hold, since one of the two must be true by assumption *invar*  $C$ . By substituting  $2r$  for  $x$  in Lemma 6.1, we conclude  $\forall c_1 \in (C \cup \{s\}). \forall c_2 \in (C \cup \{s\}). c_1 \neq c_2 \longrightarrow 2 * r < \text{dist } c_1 c_2$ , hence *invar*  $(C \cup \{s\})$ .

**Case 2**  $s$  is not a candidate site w.r.t.  $r$ , i.e.,  $\text{distance } C s \leq 2 * r$ . Then, for an arbitrary  $s' \in S$

$$\begin{aligned} \text{distance } (C \cup \{s\}) s' &\leq \text{distance } C s' && \text{by monotonicity of } \textit{Min} \\ &\leq \text{distance } C s && \text{by choice of } s \\ &\leq 2 * \text{radius } C' && \text{by assumption} \end{aligned}$$

Hence *invar*  $(C \cup \{s\})$ , i.e., the invariant is preserved.  $\square$

We now prove a lemma concerning the cardinality of the selected centers  $C$  that will be useful in showing the postcondition. It is furthermore in this lemma that the question raised above about the factor 2 is answered. Concretely, we will show that the  $\text{dist } c_1 c_2 > 2 * r$  for all  $c_1 \neq c_2$  in  $C$  condition implies that any smaller set of centers must have radius larger than  $r$ .

**Lemma 6.2.**

$$\begin{aligned} &\textit{finite } C \wedge |C| > n \wedge C \subseteq S \wedge \\ &(\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow \text{dist } c_1 c_2 > 2 * r) \longrightarrow \\ &(\forall C'. 0 < |C'| \wedge |C'| \leq n \longrightarrow \text{radius } C' > r) \end{aligned}$$

*Proof.* Assume, to the contrary, that there exists a set of centers  $C'$  with  $0 < |C'| \leq n$  and  $\text{radius } C' \leq r$ . We show that  $|C'| > n$ , a contradiction. Since the selected centers in  $C$  are themselves sites in  $S$ , each must be by definition within  $\text{radius } C' \leq r$  of a center in  $C'$ , but the centers in  $C$  are all more than  $2 * r$  apart from each other. There cannot be two distinct  $c_1, c_2 \in C$  with  $\text{dist } c_i c' \leq r$ , otherwise we would have:

$$\begin{aligned} 2 * r &< \text{dist } c_1 c_2 && \text{by assumption} \\ &\leq \text{dist } c_1 c' + \text{dist } c' c_2 && \text{by the triangle inequality} \\ &\leq 2 * r && \text{by } \text{dist } c_i c' \leq r \end{aligned}$$

(Note that this argument would not work for any  $\alpha * r$  with  $\alpha < 2$ .) Therefore there is a unique center in  $C'$  for each center in  $C$ , and hence  $n = |C'| \geq |C|$ . But we assume  $|C| > n$ , a contradiction.  $\square$

Finally we prove that the invariant and  $k \leq |C|$  imply the postcondition:

$$|C| = k \wedge (\forall C'. |C'| > 0 \wedge |C'| \leq k \longrightarrow \text{radius } C \leq 2 * \text{radius } C')$$

*Proof.* Here  $|C| = k$  follows immediately. Let  $C'$  be arbitrary and assume  $0 < |C'| \leq k$ . We define  $r = \text{radius } C'$ . By the invariant,  $(\forall s \in S. \text{distance } C s \leq 2 * r) \vee (\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow \text{dist } c_1 c_2 > 2 * r)$  holds for  $C'$ . We prove  $\text{radius } C \leq 2 * r$  by case distinction on this disjunction.

**Case 1** Assumption:  $\forall s \in S. \text{distance } C s \leq 2 * r$ . Then it trivially follows that  $\text{Max}(\text{distance } C' S) = \text{radius } C \leq 2 * r$ .

**Case 2 Assumption:**  $\forall c_1 \in C. \forall c_2 \in C. c_1 \neq c_2 \longrightarrow \text{dist } c_1 \ c_2 > 2 * r$ . We prove the conclusion by contradiction. Assume, to the contrary, that  $\text{radius } C > 2 * r$ . This implies that the site  $s = \text{furthest\_from } C$  is a candidate site w.r.t.  $r$  since  $\text{distance } C \ s = \text{radius } C > 2 * r$ . Furthermore,  $s$  cannot be in  $C$ , for if it were, then  $0 = \text{distance } C \ s = \text{radius } C > 2 * r$ , which is impossible. We can therefore insert it into  $C$  to obtain a set satisfying  $\forall c_1 \in C \cup \{s\}. c_2 \in C \cup \{s\}. c_1 \neq c_2 \longrightarrow \text{dist } c_1 \ c_2 > 2 * r$  by Lemma 6.1. We thus have, by Lemma 6.2, that  $\forall C. |C| \leq k \wedge |C| > 0 \longrightarrow \text{radius } C > r$ . In particular,  $\text{radius } C' > r = \text{radius } C'$ , a contradiction.  $\square$

## 7. SET COVER

As in the two previous sections, we base our formalization on [KT06], but this time Chapter 11.3. We are given a finite set of elements  $U :: 'a \text{ set}$  and an indexed collection  $S :: \text{nat} \Rightarrow 'a \text{ set}$  of subsets of  $U$ , defined over a segment  $\{1..m\}$  of the natural numbers, such that  $U = \bigcup (S \ ' \ \{1..m\})$ . A function  $w :: \text{nat} \Rightarrow \text{real}$  associates a non-negative weight to each (index of a) subset. Our goal is to find a set cover  $C$  of  $U$

$sc :: \text{nat set} \Rightarrow 'a \text{ set} \Rightarrow \text{bool}$

$sc \ C \ U = (C \subseteq \{1..m\} \wedge \bigcup (S \ ' \ C) = U)$

that minimizes the total weight  $W \ C = (\sum_{i \in C}. w \ i)$ . The greedy approach takes the subset that covers the most elements relative to its weight, i.e., the subset  $S \ i$  that minimizes

$$\text{cost } R \ i = \frac{w \ i}{|S \ i \cap R|}$$

where  $R$  is the set of elements yet to be covered. We will prove that this greedy algorithm has an approximation factor of  $H \ d^*$  where  $H \ n$  is the  $n$ -th harmonic number and  $d^*$  the cardinality of the largest subset. The algorithm is described by the following Hoare triple:

```

{True}
R := U; C := ∅;
WHILE R ≠ ∅ INV {inv C R}
DO i := min_arg R m;
   R := R - S i;
   C := C ∪ {i};
OD
{sc C U ∧ (∀ C'. sc C' U → W C ≤ H d* * W C')}

```

The invariant  $inv$  is explained later. The function  $min\_arg$  (a variant of the load balancing  $min\_arg$ ) selects a subset with minimum cost while taking care that the overlap with  $R$  is non-empty (unless there is no alternative):

$min\_arg :: 'a \text{ set} \Rightarrow \text{nat} \Rightarrow \text{nat}$

$min\_arg \ R \ 0 = 1$

$min\_arg \ R \ (x + 1) =$

$(\text{let } j = min\_arg \ R \ x$

$\text{in if } S \ j \cap R = \emptyset \vee S \ (x + 1) \cap R \neq \emptyset \wedge \text{cost } R \ (x + 1) < \text{cost } R \ j$

$\text{then } x + 1 \text{ else } j)$

Like with the load balancing problem, a linear scan suffices here. The following properties of  $\text{min\_arg}$  can be shown by induction on  $m$ :

**Lemma 7.1.**

- (1)  $0 < m \longrightarrow \text{min\_arg } R \ m \in \{1..m\}$
- (2)  $S (\text{min\_arg } R \ m) \cap R = \emptyset \longrightarrow (\forall i \in \{1..m\}. S \ i \cap R = \emptyset)$
- (3)  $k \in \{1..m\} \wedge S \ k \cap R \neq \emptyset \longrightarrow \text{cost } R (\text{min\_arg } R \ m) \leq \text{cost } R \ k$

Unlike the proof of the load balancing problem, the desired approximation factor does not arise from general bounds. Instead we employ a function  $c$  that returns the cost an individual element contributes to the set cover. Namely, at every iteration of the algorithm let  $c \ s = \text{cost } R \ i$  for all  $s \in S \ i \cap R$  where  $i$  is the subset picked by  $\text{min\_arg}$ . As this function is only relevant to the proof of the algorithm, we can use existential quantification to lift it into the invariant. This function satisfies two important properties upon termination of the algorithm:

- $W \ C = (\sum s \in U. c \ s)$
- $\forall k \in \{1..m\}. (\sum s \in S \ k. c \ s) \leq H \ |S \ k| * w \ k$

The former may be rather intuitive, but reasoning about the latter is more difficult. Kleinberg and Tardos provide an informal proof that argues about the sum of all element costs in the subset at an arbitrary step of the algorithm where an element  $j$  of the subset is being covered. Doing so requires indexing and reordering the elements from 1 to  $d$  where  $d$  is the cardinality of the subset. We can then show the bound by splitting the subset in such a way that elements already covered come before the element(s) that will be covered in this step, and elements that have not yet been covered come after. Assuming  $k \in \{1..m\}$ , we can see:

- (1) At least the elements from index  $j$  to  $d$  (cardinality) are not covered yet  
(i.e.,  $|S \ k \cap R| \geq d - j + 1$ ).
- (2) Therefore  $\text{cost } R \ k = \frac{w \ k}{|S \ k \cap R|} \leq \frac{w \ k}{d - j + 1}$
- (3) Covered elements up to index  $j$  were covered by subsets picked by  $\text{min\_arg}$
- (4) Therefore  $\text{cost } R \ i \leq \text{cost } R \ k$  because of Lemma 7.1.3
- (5)  $\sum_{s \in S \ k} c \ s = \sum_{j=1}^d c \ s_{k_j} \leq \sum_{j=1}^d \frac{w \ k}{d - j + 1} = \frac{w \ k}{d} + \frac{w \ k}{d-1} + \dots + \frac{w \ k}{1} = H \ d * w \ k$

This proof is concise and the final statement readily implies the desired approximation factor. This is a nice textbook proof, but a direct formalization is very unpleasant: it requires indexing elements in the order in which they were covered by the algorithm. We give an invariant-based proof that builds on the same proof idea but does not require any indexing. This is the invariant:

$$\begin{aligned}
& \text{inv} :: \text{nat set} \Rightarrow 'a \text{ set} \Rightarrow \text{bool} \\
& \text{inv } C \ R = (\text{sc } C \ (U - R) \wedge R \subseteq U \wedge c\_exists \ C \ R) \\
& c\_exists :: \text{nat set} \Rightarrow 'a \text{ set} \Rightarrow \text{bool} \\
& c\_exists \ C \ R = \\
& \quad (\exists c. W \ C = (\sum s \in U - R. c \ s) \wedge \\
& \quad \quad (\forall i. 0 \leq c \ i) \wedge \\
& \quad \quad (\forall k \in \{1..m\}. \\
& \quad \quad \quad (\sum s \in S \ k \cap (U - R). c \ s) \leq (\sum_{j=1}^{|S \ k|} \frac{1}{|S \ k \cap R| + 1 - j}) * w \ k))
\end{aligned}$$

Using  $U - R$  we can argue about the elements that have already been covered by the algorithm. The relation between the proof above and the upper bound in the last conjunct of  $c\_exists$  is less apparent, but we can perform an index shift that makes the relation more

obvious. Let  $d = |S k|$ , then:

$$\sum_{j=|S k \cap R| + 1}^d \frac{w k}{j} = \sum_{j=1}^{d - |S k \cap R|} \frac{w k}{d - j + 1} = \sum_{j=1}^{|S k \cap (U - R)|} \frac{w k}{d - j + 1}$$

This is closer to the sum we see in the proof by Kleinberg and Tardos, but complicates the invariant proof as one now has to argue about the bounds of the sum as well as the content of it, while our definition in *c\_exists* leaves the content of the sum as  $\frac{1}{j}$ .

We now come to the proof of the invariant. Initialization is trivial if we pick  $c = (\lambda. 0)$ . We may now assume the invariant and  $R \neq \emptyset$ . Let  $i$  be the subset picked by *min\_arg*, moreover  $R_g := R - S i$  and  $C_g := C \cup \{i\}$ . First note  $0 < m$ , as  $R = \emptyset$  otherwise ( $R \subseteq U$  and  $U = \bigcup (S \{1..m\})$ ), hence  $i \in \{1..m\}$  (Lemma 7.1.1). Correctness is preserved as  $S i$  is a subset of  $U$ , thus  $sc C_g (U - R_g)$ . Since we are only removing elements from  $R$ ,  $R_g$  remains a subset of  $U$  as well. Before we prove that *c\_exists*  $C_g R_g$  holds, note that the set cover actually strictly grows in an iteration of the algorithm. We know that  $\exists k \in \{1..m\}. S k \cap R \neq \emptyset$  because  $R \neq \emptyset$  and  $R \subseteq U$ . Combined with Lemma 7.1.2 we know that  $S i \cap R \neq \emptyset$ . Therefore new elements are covered in this iteration and by extension a subset was picked that was not in the cover before ( $i \notin C$ ), thus *weight-eq*:  $W C_g = W C + w i$ . We define the cost function  $c_g = (\lambda x. \text{if } x \in S i \cap R \text{ then } cost R i \text{ else } c x)$ , where  $c$  is the cost function obtained using *c\_exists*, and see that

$$\begin{aligned} \sum_{s \in U - R_g} c_g s &= \sum_{s \in U - R} c_g s + \sum_{s \in S i \cap R} c_g s && \text{by } R \subseteq U \\ &= \sum_{s \in U - R} c s + |S i \cap R| cost R i && \text{by } c_g\text{-def} \\ &= \sum_{s \in U - R} c s + |S i \cap R| \frac{w i}{|S i \cap R|} && \text{by } cost\text{-def} \\ &= \sum_{s \in U - R} c s + w i && \text{by } S i \cap R \neq \emptyset \\ &= W C + w i = W C_g && \text{by } inv\text{-def and } weight\text{-eq} \end{aligned}$$

As *cost R i* is always positive, we know that  $\forall i. 0 \leq c_g i$  is preserved. Finally we come to the last conjunct of *c\_exists*. Assume  $k \in \{1..m\}$  and let  $A = S k \cap S i \cap R$  be the elements of subset  $S k$  that will be covered in this iteration. We can transform the lower bound of the right sum as follows *lbr*:  $|S k \cap R_g| + 1 = |S k \cap R| - |A| + 1$ . This allows for the following transformation:

$$\begin{aligned} \sum_{s \in S k \cap (U - R_g)} c_g s &= \sum_{s \in S k \cap (U - R)} c_g s + \sum_{s \in A} c_g s && \text{by } R \subseteq U \\ &= \sum_{s \in S k \cap (U - R)} c s + |A| cost R i && \text{by } c_g\text{-def} \\ &\leq w k \sum_{j=|S k \cap R| + 1}^{|S k|} \frac{1}{j} + |A| cost R k && \text{by } inv\text{-def and Lem. 7.1.3} \end{aligned}$$

$$\begin{aligned}
&= w k \left( \sum_{j=|S k \cap R|+1}^{|S k|} \frac{1}{j} + |A| \frac{1}{|S k \cap R|} \right) && \text{by } \textit{cost-def} \\
&= w k \left( \sum_{j=|S k \cap R|+1}^{|S k|} \frac{1}{j} + \sum_{j=|S k \cap R| - |A| + 1}^{|S k \cap R|} \frac{1}{|S k \cap R|} \right) \\
&\leq w k \left( \sum_{j=|S k \cap R|+1}^{|S k|} \frac{1}{j} + \sum_{j=|S k \cap R| - |A| + 1}^{|S k \cap R|} \frac{1}{j} \right) \\
&= w k \sum_{j=|S k \cap R| - |A| + 1}^{|S k|} \frac{1}{j} = w k \sum_{j=|S k \cap R_g| + 1}^{|S k|} \frac{1}{j} && \text{by } \textit{lbr}
\end{aligned}$$

Unlike the proof in [KT06] we only had to argue about the cost of  $A$  here, as the cost of previous elements is already proved by the assumed invariant.

Finally we show how the invariant implies the approximation factor upon termination of the algorithm:

**Lemma 7.2.**  $\textit{inv } C \emptyset \wedge \textit{sc } C' U \longrightarrow W C \leq H d^* * W C'$

From  $\textit{inv } C \emptyset$ ,  $\textit{c\_exists}$  and  $R = \emptyset$  we can infer  $\textit{cost-eq}$ :  $W C = (\sum_{s \in U}. c s)$  and  $\textit{h-bound}$ :  $\forall k \in \{1..m\}. (\sum_{s \in S k}. c s) \leq H |S k| * w k$  and can derive the claim:

$$\begin{aligned}
W C &= \sum_{s \in U} c s && \text{by } \textit{cost-eq} \\
&\leq \sum_{k \in C'} \sum_{s \in S k} c s && \text{by } \forall i. 0 \leq c i \\
&\leq \sum_{k \in C'} H |S k| * w k && \text{by } \textit{h-bound} \\
&\leq \sum_{k \in C'} H d^* * w k && \text{by } \textit{d^*-def} \\
&= H d^* \sum_{k \in C'} w k = H d^* * W C'
\end{aligned}$$

## 8. BIN PACKING

We consider the linear time  $\frac{3}{2}$ -approximation algorithm for the bin packing problem proposed by Berghammer and Reuter [BR03]. We are given a finite, non-empty set of objects  $U :: 'a \textit{ set}$ , whose *weights* are given by a function  $w :: 'a \Rightarrow \textit{real}$ . Note that in this paper *nats* are implicitly converted to *reals* if needed. The weight of an object in  $U$  is strictly greater than zero, but bounded by a maximum capacity  $c :: \textit{nat}$ . The abbreviation  $W(B) \equiv \sum_{u \in B} w(u)$  denotes the weight of a bin  $B \subseteq U$ . The set  $U$  can be separated into *small* and *large* objects. An object  $u$  is considered small if  $w(u) \leq \frac{c}{2}$ , and large otherwise. We assume that all small objects in  $U$  can be found in a set  $S$ , and large objects in  $U$  can be found in a set  $L$ , such that  $S \cup L = U$  and  $S \cap L = \emptyset$ . Of course  $L$  and  $S$  could also be computed from  $U$  in linear time. Variables  $U$ ,  $w$ ,  $c$ ,  $L$ , and  $S$  are fixed throughout this section.

A solution  $P$  to the bin packing problem is defined as follows:

$bp :: 'a \text{ set set} \Rightarrow bool$   
 $bp P = (partition\_on U P \wedge (\forall B \in P. W B \leq c))$

The predicate  $partition\_on :: 'a \text{ set} \Rightarrow 'a \text{ set set} \Rightarrow bool$  is defined in the Isabelle/HOL library:  $partition\_on U P$  means that  $P$  is a partition of the set  $U$ . Viewing every element of  $P$  as a bin,  $bp P$  expresses that all objects are contained in exactly one bin and the weight of every bin is bounded by  $c$ .

The idea behind the algorithm proposed by Berghammer and Reuter is to split the solution  $P$  into two partial solutions  $P_1$  and  $P_2$ . At every step of the algorithm we consider two bins  $B_1$  and  $B_2$  which we try to fill with remaining objects from  $V \subseteq U$  that have not been assigned yet. If adding the object to  $B_1$  or  $B_2$  would cause it to exceed its maximum capacity, the bin is moved into the partial solution  $P_1$  or  $P_2$  respectively and cleared. Once there are no small objects left, the solution is the union of the partial solutions  $P_1$  and  $P_2$ , the bins  $B_1$  and  $B_2$  (if they still contain objects), and the remaining large objects, which each receive their own bin, as no two large objects can fit into a single bin. To ensure that no empty bins are added to the solution, we define:

$\llbracket \cdot \rrbracket :: 'a \text{ set} \Rightarrow 'a \text{ set set}$   
 $\llbracket B \rrbracket = (if B = \emptyset \text{ then } \emptyset \text{ else } \{B\})$

The final union can now be written as  $P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V\}$  where  $V$  contains the remaining large elements. The algorithm can be specified by the following Hoare triple:

```

{ True }
P1 := ∅; P2 := ∅; B1 := ∅; B2 := ∅; V := U;
WHILE V ∩ S ≠ ∅ INV { inv3 P1 P2 B1 B2 V } DO
IF B1 ≠ ∅ THEN u := some (V ∩ S)
ELSE IF V ∩ L ≠ ∅ THEN u := some (V ∩ L)
ELSE u := some (V ∩ S) FI FI;
V := V - {u};
IF W(B1) + w(u) ≤ c THEN B1 := B1 ∪ {u}
ELSE IF W(B2) + w(u) ≤ c THEN B2 := B2 ∪ {u}
ELSE P2 := P2 ∪ ⌊B2⌋; B2 := {u} FI;
P1 := P1 ∪ ⌊B1⌋; B1 := ∅ FI
OD;
P := P1 ∪ ⌊B1⌋ ∪ P2 ∪ ⌊B2⌋ ∪ {{v} | v ∈ V}
{ bp P ∧ (∀ Q. bp Q → |P| ≤ 3 / 2 * |Q|) }

```

Berghammer and Reuter prove functional correctness using a simplified version of this algorithm where an arbitrary element of  $V$  is assigned to  $u$ . This allows for fewer case distinctions, as the first *IF-THEN-ELSE* block can be ignored. One needs to find a loop invariant that implies functional correctness and prove that it is preserved in the following cases:

**Case 1** The object fits into  $B_1$ :

$inv_1 P_1 P_2 B_1 B_2 V \wedge u \in V \wedge W B_1 + w u \leq c \rightarrow$   
 $inv_1 P_1 P_2 (B_1 \cup \{u\}) B_2 (V - \{u\})$

**Case 2** The object fits into  $B_2$ :



$$\begin{aligned} & inv_1 P_1 P_2 B_1 B_2 V \wedge u \in V \wedge W B_2 + w u \leq c \longrightarrow \\ & inv_1 (P_1 \cup \llbracket B_1 \rrbracket) P_2 \emptyset (B_2 \cup \{u\}) (V - \{u\}) \end{aligned}$$

**Case 3** The object fits into neither bin:

$$\begin{aligned} & inv_1 P_1 P_2 B_1 B_2 V \wedge u \in V \longrightarrow \\ & inv_1 (P_1 \cup \llbracket B_1 \rrbracket) (P_2 \cup \llbracket B_2 \rrbracket) \emptyset \{u\} (V - \{u\}) \end{aligned}$$

Berghammer and Reuter [BR03] define the following predicate as their loop invariant:

$$inv_1 P_1 P_2 B_1 B_2 V = bp (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V\})$$

As it turns out, this invariant is too weak. Assume  $inv_1 P_1 P_2 B_1 B_2 V$ . Suppose  $P_1$  (alternatively  $P_2$ ) already contains the non-empty bin  $B_1$ . Note that this does not violate the invariant because  $P_1 \cup \llbracket B_1 \rrbracket = P_1$ . Now, if the algorithm modifies  $B_1$  by adding an element from  $V$  such that  $B_1$  becomes some  $B_1'$  then  $B_1 \cap B_1' \neq \emptyset$  and  $B_1 \in P_1$ , i.e.,  $B_1'$  is no longer disjoint from the elements of  $P$ . The same issue arises with the added object  $u \in V$ , if  $\{u\}$  is already in  $P_1$  or  $P_2$ . To account for such cases, we will require additional conjuncts:

$$\begin{aligned} & inv_1 :: 'a \text{ set } set \Rightarrow 'a \text{ set } set \Rightarrow 'a \text{ set } \Rightarrow 'a \text{ set } \Rightarrow 'a \text{ set } \Rightarrow bool \\ & inv_1 P_1 P_2 B_1 B_2 V = \\ & \quad (bp (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V\}) \wedge \\ & \quad \quad \bigcup (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket) = U - V \wedge \\ & \quad \quad B_1 \notin P_1 \cup P_2 \cup \llbracket B_2 \rrbracket \wedge \\ & \quad \quad B_2 \notin P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \wedge \\ & \quad \quad (P_1 \cup \llbracket B_1 \rrbracket) \cap (P_2 \cup \llbracket B_2 \rrbracket) = \emptyset) \end{aligned}$$

There are different ways to strengthen the original  $inv_1$ . We use the above additional conjuncts as they can be inserted in existing proofs with little modification, and their preservation in the invariant can be proved quite trivially. The first additional conjunct ensures that no element still in  $V$  is already in a bin or partial solution. The second and third additional conjuncts ensure distinctness of the bins  $B_1$  and  $B_2$  with the remaining solution. The final conjunct ensures that the partial solutions with their added bins are disjoint from each other. It should be noted that the last conjunct is *not* necessary to prove functional correctness. It will, however, be needed in later proofs, and as its preservation in this invariant for the simplified algorithm can be used in the proof of the full algorithm, one can save redundant case distinctions by proving it now. Another advantage of proving it now is that later invariants can remain identical to the invariants proposed in the paper.

We now prove the preservation of  $inv_1$  in all three cases. As we assume the invariant to hold before the execution of the loop body, we can see from the first additional conjunct  $\bigcup (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket) = U - V$  and the assumption  $u \in V$  that *not-in*:  $\forall B \in P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket. u \notin B$  holds. This will be needed for all three cases. Now, we can begin with Case 1. We first show

$$bp (P_1 \cup \llbracket B_1 \cup \{u\} \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V - \{u\}\})$$

One can see that this union does not contain the empty set. The object  $u$  is now moved from a singleton set into  $B_1$ . Therefore, the union of all bins will again return  $U$ . To show that this union remains pairwise disjoint, we can use *not-in* and the second additional conjunct of  $inv_1$  to show that  $u$  is not yet contained in the partial solution and  $B_1$  is *distinct* from any other bin. Therefore, combined with the assumption that the union was pairwise disjoint before the modification, the union remains pairwise disjoint. To prove the preservation of the second conjunct of  $bp$ , we need to show that the bin weights do not exceed their maximum

capacity  $c$ . The only bin that was changed in this step is  $B_1$ , which has increased its weight by  $w(u)$ . As we are in Case 1, we can assume that  $u$  fits into  $B_1$ ,  $W(B_1) + w(u) \leq c$ . Therefore, this conjunct holds as well. Now, one only needs to show that the additional conjuncts are preserved. For the first additional conjunct, we can again use *not-in* to show:

$$\begin{aligned} & \bigcup (P_1 \cup \llbracket B_1 \cup \{u\} \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket) = U - (V - \{u\}) \\ \iff & \bigcup (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket) \cup \{u\} = U - (V - \{u\}) && \text{by } \textit{not-in} \\ \iff & \bigcup (P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket) \cup \{u\} = U - V \cup \{u\} && \text{by } u \in U \end{aligned}$$

Using the first additional conjunct of the assumed invariant, one can see that this must hold. The remaining conjuncts

$$\begin{aligned} & B_1 \cup \{u\} \notin P_1 \cup P_2 \cup \llbracket B_2 \rrbracket \\ & B_2 \notin P_1 \cup \llbracket B_1 \cup \{u\} \rrbracket \cup P_2 \\ & (P_1 \cup \llbracket B_1 \cup \{u\} \rrbracket) \cap (P_2 \cup \llbracket B_2 \rrbracket) = \emptyset \end{aligned}$$

can be automatically proved in Isabelle using *not-in* and the assumption that the conjuncts of  $\textit{inv}_1$   $P_1$   $P_2$   $B_1$   $B_2$   $V$  held before the modification. The proof for Case 2 is almost identical to that of Case 1. The main difference is that the focus now lies on  $B_2$  and the fact that  $B_1$  is now emptied and the previous contents added to the partial solution  $P_1$ . One therefore has to show that

$$\textit{bp} (P_1 \cup \llbracket B_1 \rrbracket \cup \llbracket \emptyset \rrbracket \cup P_2 \cup \llbracket B_2 \cup \{u\} \rrbracket \cup \{\{v\} \mid v \in V - \{u\}\})$$

holds. As  $\llbracket \emptyset \rrbracket$  can be ignored, one can see that the act of emptying  $B_1$  and adding it to the partial solution will not otherwise affect the proof. The proof of *bp* in Case 3 is trivial, as the modifications made in this step can simply be undone by applying the following steps:

$$\begin{aligned} & P_1 \cup \llbracket B_1 \rrbracket \cup \llbracket \emptyset \rrbracket \cup (P_2 \cup \llbracket B_2 \rrbracket) \cup \llbracket \{u\} \rrbracket \cup \{\{v\} \mid v \in V - \{u\}\} \\ & = P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{u\}\} \cup \{\{v\} \mid v \in V - \{u\}\} && \text{by } \llbracket \cdot \rrbracket\text{-def} \\ & = P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V\} && \text{by } u \in V \end{aligned}$$

Now, one only needs to show that the remaining additional conjuncts hold. This can again be shown automatically using *not-in* and the fact that  $\textit{inv}_1$   $P_1$   $P_2$   $B_1$   $B_2$   $V$  held before the modifications. Therefore,  $\textit{inv}_1$  is preserved in all three cases.

To prove the approximation factor, we proceed as in [BR03] and establish suitable lower bounds. The first lower bound

**Lemma 8.1.**  $\textit{bp} P \longrightarrow |L| \leq |P|$

holds because a bin can only contain at most one large object, and every large object needs to be in the solution. To prove this in Isabelle, we first make the observation that for every large object there exists a bin in  $P$  in which it is contained. Therefore, we may obtain a function  $f$  that returns this bin for every  $u \in L$ . Using the fact that any bin can hold at most one large object, we can show that this function has to be injective, as every large object must map to a unique bin. Hence, the number of large objects is equal to the number of bins  $f$  maps to. Moreover, the image of  $f$  has to be a subset of  $P$ . Thus, the number of large objects has to be a lower bound on the number of bins in  $P$ .

As it turns out, the algorithm will ensure that there is always at least one large object in a bin for the first partial solution as long as large objects are available. This means we can assume that  $V \cap L \neq \emptyset \longrightarrow (\forall B \in P_1 \cup \llbracket B_1 \rrbracket. B \cap L \neq \emptyset)$ . Therefore, we can use the previous lower bound to show the following:

**Lemma 8.2.**  $bp\ P \wedge inv_1\ P_1\ P_2\ B_1\ B_2\ V \wedge (\forall B \in P_1 \cup \llbracket B_1 \rrbracket. B \cap L \neq \emptyset) \longrightarrow |P_1 \cup \llbracket B_1 \rrbracket \cup \{\{v\} \mid v \in V \cap L\}| \leq |P|$

Another easy lower bound is this one:

**Lemma 8.3.**  $bp\ P \longrightarrow (\sum_{u \in U} w\ u) \leq c * |P|$

The next lower bound arises from the fact that an object is only ever put into  $B_2$ , and therefore  $P_2$ , if it would have caused  $B_1$  to overflow. As a result of this, we can define a bijective function  $f$  that maps every bin of  $P_1$  to the object in  $P_2 \cup \llbracket B_2 \rrbracket$  that would have caused the bin to overflow. We define:

$bij\_exists :: 'a\ set\ set \Rightarrow 'a\ set \Rightarrow bool$

$bij\_exists\ P\ V = (\exists f. bij\_betw\ f\ P\ V \wedge (\forall B \in P. c < W\ B + w\ (f\ B)))$

From this, we can make the observation that the number of bins in  $P_1$  is a *strict* lower bound on the number of bins of any correct bin packing  $P$ :

**Lemma 8.4.**  $bp\ P \wedge inv_1\ P_1\ P_2\ B_1\ B_2\ V \wedge bij\_exists\ P_1\ (\bigcup (P_2 \cup \llbracket B_2 \rrbracket)) \longrightarrow |P_1| + 1 \leq |P|$

Unlike the proof outlined in [BR03], we begin with a case distinction on  $P_1$ . The reasoning behind this is that if  $P_1$  is empty, the strict nature of the lower bound cannot be shown from the calculation that Berghammer and Reuter make. Therefore, we consider the case where  $P_1$  is empty separately. If  $P_1$  is empty, our goal is to prove that 1 is a lower bound on the number of bins in  $P$ . This follows from the fact that  $U$  is non-empty, and therefore any correct bin packing must contain at least one bin. For the other case, we may now assume that  $P_1$  is non-empty. In the following proof, we will need the final conjunct of  $inv_1$ ,  $(P_1 \cup \llbracket B_1 \rrbracket) \cap (P_2 \cup \llbracket B_2 \rrbracket) = \emptyset$ , which we can transform into *disjoint*:  $P_1 \cap (P_2 \cup \llbracket B_2 \rrbracket) = \emptyset$ . We also obtain the bijective function  $f$  and observe that, as the object obtained from  $f$  for a bin  $B \in P_1$  caused  $B$  to exceed its capacity, *exceed*:  $c < W(B) + w(f(B))$  must hold. We can now perform the following calculation:

$$\begin{aligned}
c|P_1| &= \sum_{B \in P_1} c \\
&< \sum_{B \in P_1} W(B) + \sum_{B \in P_1} w(f(B)) && \text{by } P_1 \neq \emptyset \text{ and } exceed \\
&= \sum_{B \in P_1} W(B) + \sum_{B \in P_2 \cup \llbracket B_2 \rrbracket} W(B) && \text{by } f \text{ bijective} \\
&= \sum_{B \in P_1 \cup P_2 \cup \llbracket B_2 \rrbracket} W(B) && \text{by } disjoint \\
&\leq \sum_{u \in U} w(u) \leq c|P| && \text{by } inv_1 \text{ and Lemma 8.3}
\end{aligned}$$

Therefore  $|P_1| < |P|$  and, by extension,  $|P_1| + 1 \leq |P|$ .

We only sketch the rest of the proof because it is almost identical to that in [BR03]. First we need two extensions of  $inv_1$  to show the approximation ratio:

$inv_2\ P_1\ P_2\ B_1\ B_2\ V =$   
 $(inv_1\ P_1\ P_2\ B_1\ B_2\ V \wedge$   
 $(V \cap L \neq \emptyset \longrightarrow (\forall B \in P_1 \cup \llbracket B_1 \rrbracket. B \cap L \neq \emptyset)) \wedge$   
 $bij\_exists\ P_1\ (\bigcup (P_2 \cup \llbracket B_2 \rrbracket))) \wedge 2 * |P_2| \leq |\bigcup P_2|)$

$inv_3 P_1 P_2 B_1 B_2 V = (inv_2 P_1 P_2 B_1 B_2 V \wedge B_2 \subseteq S)$

The motivation for the last conjunct in  $inv_2$  is the following lower bound:

$inv_1 P_1 P_2 B_1 B_2 V \wedge 2 * |P_2| \leq |\bigcup P_2| \wedge bij\_exists P_1 (\bigcup (P_2 \cup \llbracket B_2 \rrbracket)) \longrightarrow$   
 $2 * |P_2 \cup \llbracket B_2 \rrbracket| \leq |P_1| + 1$

The main lower bound lemma (Theorem 4.1 in [BR03]) is the following:

**Lemma 8.5.**  $V \cap S = \emptyset \wedge inv_2 P_1 P_2 B_1 B_2 V \wedge bp P \longrightarrow$   
 $|P_1 \cup \llbracket B_1 \rrbracket \cup P_2 \cup \llbracket B_2 \rrbracket \cup \{\{v\} \mid v \in V\}| \leq 3 / 2 * |P|$

From this lower bound the postcondition of the algorithm follows easily under the assumption that  $inv_2$  holds at the end of the loop. This in turn follows because  $inv_3$  can be shown to be a loop invariant.

## 9. CONCLUSION AND FUTURE WORK

In this first application of theorem proving to approximation algorithms we have verified a number of classical approximation algorithms for fundamental NP-complete problems, have corrected purported invariants from the literature and could even strengthen the approximation ratio in one case. Moreover we have given simple invariant-based proofs for algorithms where the proofs in the literature (or their direct formalization) are more complicated.

Although we have demonstrated the benefits of formal verification of approximation algorithms, we have only scratched the surface of this rich theory. The next step is to explore the subject more systematically. As a large fraction of the theory of approximation algorithms is based on linear programming, this is a promising and challenging direction to explore. Some linear programming theory has been formalized in Isabelle already [BHT19, PK19]. Approximation algorithms can also be formulated as relational programs, and verified accordingly. This approach was explored in [BHS16], with some support from theorem provers, but has yet to be fully formalized.

**Acknowledgement.** Tobias Nipkow is supported by DFG grant NI 491/16-1.

## REFERENCES

- [BHS16] Rudolf Berghammer, Peter Höfner, and Insa Stucke. Cardinality of relations and relational approximation algorithms. *Journal of Logical and Algebraic Methods in Programming*, 85(2):269–286, 2016.
- [BHT19] Ralph Bottesch, Max W. Haslbeck, and René Thiemann. Verifying an incremental theory solver for linear arithmetic in Isabelle/HOL. In Andreas Herzig and Andrei Popescu, editors, *Frontiers of Combining Systems, FroCoS 2019*, volume 11715 of *Lecture Notes in Computer Science*, pages 223–239. Springer, 2019. doi:10.1007/978-3-030-29007-8\_13.
- [BK10] Nikhil Bansal and Subhash Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and Programming, ICALP 2010, Part I*, volume 6198 of *LNCS*, pages 250–261. Springer, 2010. doi:10.1007/978-3-642-14165-2\_22.
- [BM03] Rudolf Berghammer and Markus Müller-Olm. Formal development and verification of approximation algorithms using auxiliary variables. In Maurice Bruynooghe, editor, *Logic Based Program Synthesis and Transformation, LOPSTR 2003*, volume 3018 of *LNCS*, pages 59–74. Springer, 2003. doi:10.1007/978-3-540-25938-1\_6.

- [BR03] Rudolf Berghammer and Florian Reuter. A linear approximation algorithm for bin packing with absolute approximation factor  $3/2$ . *Sci. Comput. Program.*, 48(1):67–80, 2003. doi:10.1016/S0167-6423(03)00011-X.
- [ENR20] Robin Eßmann, Tobias Nipkow, and Simon Robillard. Verified approximation algorithms. In Nicolas Peltier and Viorica Sofronie-Stokkermans, editors, *Automated Reasoning, IJCAR 2020, Part II*, volume 12167 of *Lecture Notes in Computer Science*, pages 291–306. Springer, 2020. doi:10.1007/978-3-030-51054-1\_17.
- [ENRS20] Robin Eßmann, Tobias Nipkow, Simon Robillard, and Ujkan Sulejmani. Verified approximation algorithms. *Archive of Formal Proofs*, January 2020. [http://isa-afp.org/entries/Approximation\\_Algorithms.html](http://isa-afp.org/entries/Approximation_Algorithms.html), Formal proof development.
- [HR97] Magnús M Halldórsson and Jaikumar Radhakrishnan. Greed is good: Approximating independent sets in sparse and bounded-degree graphs. *Algorithmica*, 18(1):145–163, 1997.
- [KT06] Jon M. Kleinberg and Éva Tardos. *Algorithm Design*. Addison-Wesley, 2006.
- [NK14] Tobias Nipkow and Gerwin Klein. *Concrete Semantics with Isabelle/HOL*. Springer, 2014. <http://concrete-semantics.org>.
- [NPW02] Tobias Nipkow, Lawrence Paulson, and Markus Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer, 2002.
- [PK19] Julian Parsert and Cezary Kaliszyk. Linear programming. *Archive of Formal Proofs*, August 2019. [http://isa-afp.org/entries/Linear\\_Programming.html](http://isa-afp.org/entries/Linear_Programming.html), Formal proof development.
- [Vaz03] Vijai Vazirani. *Approximation Algorithms*. Springer, 2003.
- [Wei81] V.K. Wei. A lower bound for the stability number of a simple graph. Technical Memorandum 81-11217-9, Bell Laboratories, 1981.