

A FRAMEWORK FOR CERTIFIED SELF-STABILIZATION*

KARINE ALTISEN, PIERRE CORBINEAU, AND STÉPHANE DEVISMES

Univ. Grenoble Alpes, CNRS, Grenoble INP[†], VERIMAG, 38000 Grenoble, France
e-mail address: Firstname.Lastname@univ-grenoble-alpes.fr

ABSTRACT. We propose a general framework to build certified proofs of distributed self-stabilizing algorithms with the proof assistant Coq. We first define in Coq the locally shared memory model with composite atomicity, the most commonly used model in the self-stabilizing area. We then validate our framework by certifying a non trivial part of an existing silent self-stabilizing algorithm which builds a k -clustering of the network. We also certify a quantitative property related to the output of this algorithm. Precisely, we show that the computed k -clustering contains at most $\lfloor \frac{n-1}{k+1} \rfloor + 1$ clusterheads, where n is the number of nodes in the network. To obtain these results, we also developed a library which contains general tools related to potential functions and cardinality of sets.

1. INTRODUCTION

In 1974, Dijkstra introduced the notion of *self-stabilizing* algorithm [24] as any distributed algorithm that resumes correct behavior within finite time, regardless of the initial configuration of the system. A self-stabilizing algorithm can withstand *any* finite number of transient faults. Indeed, after transient faults hit the system and place it in some arbitrary configuration — where, for example, the values of some variables have been arbitrarily modified — a self-stabilizing algorithm is guaranteed to resume correct behavior without external (*e.g.*, human) intervention within finite time. Thus, self-stabilization makes no hypothesis on the nature or extent of transient faults that could hit the system, and recovers from the effects of those faults in a unified manner.

For more than 40 years, a vast literature on self-stabilizing algorithms has been developed. Self-stabilizing solutions have been proposed for many kinds of classical distributed problems, *e.g.*, token circulation [31], spanning tree construction [13], clustering [9], routing [25], propagation of information with feedback [7], clock synchronization [18], *etc.* Moreover, self-stabilizing algorithms have been designed to handle various environments, *e.g.*, wired networks [31, 13, 9, 25, 7, 18], WSNs [4, 38], peer-to-peer systems [10, 8], *etc.*

Progress in self-stabilization has led to consider more and more adversarial environments. As an illustrative example, the three first algorithms proposed by Dijkstra in 1974 [24] were designed for oriented ring topologies and assuming sequential executions only, while

Key words and phrases: Self-stabilization, Proof assistant, Coq, Silent algorithms, Potential functions.

* This work is an extended version of [1] and has been partially supported by ANR projects ESTATE (ANR-16-CE25-0009) and DESCARTES (ANR-16-CE40-0023).

[†] Institute of Engineering Univ. Grenoble Alpes

nowadays most self-stabilizing algorithms are designed for fully asynchronous arbitrary connected networks, *e.g.*, [31, 9, 19].

Consequently, the design of self-stabilizing algorithms becomes more and more intricate, and accordingly, the proofs of their respective correctness and complexity are now often tricky to establish. However, proofs in distributed computing, in particular in self-stabilization, are commonly written by hand, based on informal reasoning. This potentially leads to errors when arguments are not perfectly clear, as explained by Lamport in its position paper [35]. So, in the current context, such methods are clearly pushed to their limits, since the question on confidence in proofs naturally arises.

This justifies the use of a *proof assistant*, a tool which allows to develop certified proofs interactively and check them mechanically. In this paper, we use Coq [39], recipient of the ACM 2013 Software system Award. Coq has been successfully employed for various tasks such as mathematical developments as involved in the Feit-Thompson theorem [30], formalization of the correctness of a C compiler [36, 37], certified numerical libraries [28], and verification of cryptographic protocols [3, 14].

1.1. Contribution. We propose a general framework to build certified proofs of self-stabilizing algorithms for wired networks with the tool Coq. We first define in Coq the *locally shared memory model with composite atomicity*, introduced by Dijkstra [24]. This model is the most commonly used in the self-stabilizing area. Our modeling is versatile, *e.g.*, it supports any class of network topologies (including arbitrary ones), the diversity of anonymity levels (from fully anonymous to fully identified), and various levels of asynchrony (*e.g.*, sequential, synchronous, fully asynchronous).

We show how to use and validate our framework by certifying a non trivial part of an existing silent self-stabilizing algorithm proposed in [19] which builds a k -clustering of the network. Starting from an arbitrary configuration, a silent algorithm converges within finite time to a configuration from which all communication variables are constant. This class of self-stabilizing algorithms is important, as self-stabilizing algorithms building distributed data structures (such as spanning tree or clustering) often achieve the silent property, and these silent self-stabilizing data structures are widely used as basic building blocks for more complex self-stabilizing solutions, *e.g.*, [19, 20].

Using a usual proof scheme, the certified proof consists of two main parts, one dealing with termination and the other with partial correctness.

For the termination part, we developed tools on potential functions and termination at a fine-grained level. Precisely, we define a potential function as a multiset containing a local potential per node. We then exploit two criteria that are sufficient to meet the conditions for using the Dershowitz-Manna well-founded ordering on multisets [22]. These two criteria, and the associated proof scheme, are versatile enough to be applied to prove the termination many other (silent) algorithms, whether using our framework, or by hand. We also provide tools to build termination proofs of algorithms consisting of prioritized sets of actions. These tools use a lexicographical order on multisets of local potentials. Notice that the termination proof we propose for the case study assumes a distributed unfair daemon, the most general scheduling assumption of the model. By contrast, the proof given in [19] uses a stronger scheduling hypothesis, namely, a distributed weakly fair daemon.

The partial correctness part consists of showing that (1) a k -clustering is defined in the network whenever the algorithm has terminated, and (2) a quantitative property related to this k -clustering; namely the computed k -clustering contains at most $\lfloor \frac{n-1}{k+1} \rfloor + 1$ clusterheads,

where n is the number of nodes in the network. To obtain this latter result, we provide a library dealing with cardinality of sets in general and properties on cardinals of finite sets *w.r.t.* basic set operations, *i.e.*, Cartesian product, disjoint union and subsets.

This work is an extended version of [1]; it represents about 17,560 lines of code (as computed by `coqwc`: 5k lines of specifications, 10k lines of proofs) written in Coq 8.6 compiled with OCaml 4.04.1.

1.2. Related Work. Many formal approaches have been used in the context of distributed computing. There exist tools to validate a given distributed algorithm, such as the tools embedded with TLA+ (a model checker and a proof assistant) [34]. Constructive approaches aim at synthesizing algorithms based on a given specification and a fixed topology; many of them are now based on SMT-solvers, see [6, 27]. Note that model-checking as well as synthesis are fully automated, but require to fix the topology and sometime the scheduling (*e.g.*, synchronous execution). Moreover, these techniques usually succeed with small topologies, due to computation limits. For example, model checking has been also successfully used to prove impossibility results applying on small-scale distributed systems [23]. In contrast, a proof assistant may validate a given algorithm for arbitrary-sized topologies, but is only semi-automated and requires heavy development for each algorithm. We now focus on works related to certification of distributed algorithms, most of them using Coq.

Several works have shown that proof assistants (in particular Coq) are well-suited to certify the correction of algorithms as well as impossibility results in various kinds of distributed systems. Certification of non fault-tolerant (consequently non self-stabilizing) distributed algorithms in Coq is addressed in [11, 12, 17]. In [16], an impossibility proof for the gathering problem is certified. Notice that [16, 17] consider *mobile* distributed systems. Precisely, these works are dedicated to swarms of robots that are endowed with motion actuators and visibility sensors and deployed in the Euclidean plane. These robots are weak, *i.e.*, they are anonymous, uniform, unable to explicitly communicate, and oblivious (they have no persistent memory).

Certification in the context of fault-tolerant, yet non self-stabilizing, distributed computing is addressed in [32, 2]. Kufner *et al.* [32] propose to certify (using the proof assistant *Isabelle*) fault-tolerant distributed algorithms. Their framework deals with masking fault-tolerance whereas self-stabilization is non-masking by essence. Moreover the network topology is restricted to fully connected graphs. Bouzid *et al.* [2] certify impossibility results for swarms of robots that are subjected to Byzantine faults using a model based on the one described in [16].

To the best of our knowledge, only three works deal with certification of self-stabilizing algorithms [15, 21, 33]. First, [21] proposes to certify in Coq self-stabilizing population protocols. Population protocols are used as a theoretical model for a collection (or population) of tiny mobile agents that interact with one another to carry out a computation. The movement pattern of the agents is unpredictable, and communication is implicit between close agents (there is no notion of communication network). A formal correctness proof of Dijkstra's seminal self-stabilizing algorithm [24] is conducted with the PVS proof assistant [33], where only sequential executions are considered. In [15], Courtieu proposes a setting for reasoning on self-stabilization in Coq. He restricts his study to very simple self-stabilizing algorithms (*e.g.*, the 4-states algorithm of Ghosh [29]) working on networks of restrictive topologies (lines and rings).

1.3. Roadmap. The rest of the paper is organized into two parts. The first one, from Section 2 to Section 5, describes the general framework. The case study is given in Sections 6-8. Section 9 is dedicated to concluding remarks and perspectives.

In the next section, we describe how we define the locally shared memory model with composite atomicity in Coq. In Section 3, we express the definitions of self-stabilization and silence in Coq, moreover we certify a sufficient condition to show that an algorithm is silent and self-stabilizing. We present tools for proving termination in Section 4 and tools for proving quantitative properties in Section 5. In Section 6, we present an algorithm called $\mathcal{C}(k)$ (the case study), its specification, and the assumptions under which it will be proven. In Section 7, we present the termination proof of $\mathcal{C}(k)$. Section 8 deals with the partial correctness of $\mathcal{C}(k)$.

Along this paper, we present our work together with few pieces of Coq code that we simplify in order to make them readable. In particular, we intend to use notations as defined in the model or in the algorithm. The Coq definitions, lemmas, theorems, and documentation related to this paper are available as an online browsing at <http://www-verimag.imag.fr/~altisen/PADEC/>. All source codes are also available at this URL. We encourage the reader to visit this web page for a deeper understanding of our work.

2. LOCALLY SHARED MEMORY MODEL WITH COMPOSITE ATOMICITY

In this section, we explain how we model the *locally shared memory model with composite atomicity* in Coq. This model has been introduced by Dijkstra [24], and since then it is the most commonly used in the self-stabilizing area.

2.1. Distributed Systems. We define a *distributed system* as a finite set of interconnected nodes. Each node has its own private memory and runs its own code. It can also interact with other nodes in the network *via* interconnections. The model in Coq reflects this by defining two independent types:

- A **Network** is equipped with a type **Node**, representing nodes of the network. A **Network** defines functions and properties that depict its topology, *i.e.*, interconnections between nodes. Those interconnections are specified using the type **Channel**.
- The **Algorithm** of a node **p** is equipped with a type **State**, which describes the memory state of **p**. Its main function, **run**, specifies how **p** executes and interacts with other nodes through channels (type **Channel**).

2.2. Network and Topology. Nodes in a distributed system can directly communicate with a subset of other nodes. As commonly done in the literature, we view the communication *network* as a simple directed graph $G = (V, E)$, where V is set of vertices representing nodes and $E \subseteq V \times V$ is a set of edges representing direct communication between distinct nodes. We write n to denote the numbers of nodes: $n = |V|$.

Two distinct nodes p and q are said to be *neighbors* if $(p, q) \in E$. From a computational point of view, p uses a distinct channel $c_{p,q}$ to communicate with each of its neighbors q : it does not have direct access to q . In the type **Network**, the topology is defined using this narrow point of view, *i.e.*, interconnections (edges of the graph) are represented using channels only. In particular, the neighborhood of p is encoded with the set \mathcal{N}_p which contains all channels

$c_{p,q}$ outgoing from p . The sets \mathcal{N}_p , for all p , are modeled in Coq as lists, using the function (`peers: Node → list Channel`). The function (`peer: Node → Channel → option Node`) returns the destination neighbor for a given channel name, *i.e.*, (`peer p cp,q`) returns (`Some q`), or \perp ¹ if the name is unused. We also define the shortcut ternary relation (`is_channel p c p'`) as (`peer p c`) equals (`Some p'`) where p and p' are nodes and c a channel.

Communications can be made bidirectional, assuming a property called `sym_net`, which states that for all nodes p_1 and p_2 , the network defines a channel from p_1 to p_2 if and only if it also defines a channel from p_2 to p_1 . In case of bidirectional links (p, q) and (q, p) in E , p can access its channel name at q using the function (`ρp: Channel → Channel`). Thus, we have the following identities: $\rho_p(c_{p,q})$ equals $c_{q,p} \in \mathcal{N}_q$ and $\rho_q(c_{q,p})$ equals $c_{p,q} \in \mathcal{N}_p$. In our Coq model, the role of ρ_p is assigned to the function (`reply_to: Node → Channel → Channel`).

As a last requirement, we suppose that, since the number of nodes in the network is finite, we have a list, called `all_nodes`, containing all the nodes. In particular, this assumption makes the *emptiness test* decidable: this test states that for any function (`f: Node → option A`) (with A , some type), one can compute whether f always returns \perp for any parameter. This test is used in the framework to detect termination of the algorithm.

As a means of checking actual usability of the `Network` type definition, we have defined a function that can build any finite `Network` from a description of its topology given by a list of lists of neighbors.

2.3. Computational Model. In the *locally shared memory model with composite atomicity*, nodes communicate with their neighbors using finite sets of locally shared registers, called *variables*. A node can read its own variables and those of its neighbors, but can only write to its own variables.

2.3.1. Distributed Algorithm. Each node operates according to its local *program*. A *distributed algorithm* \mathcal{A} is defined as a collection of n *programs*, each operating on a single node. The *state* of a node in \mathcal{A} is defined by the values of its local variables and is represented using an abstract immutable Coq datatype `State`. Such a datatype is usually implemented as a record containing the values of the algorithm variables. A node p can access the states of its neighbors using the corresponding channels: we call this the *local configuration* of p , and model it as a function typed (`Local_Env := Channel → option State`) which returns the current state of a neighbor, given the name of the corresponding channel (or \perp for an invalid name).

The program of each node p in \mathcal{A} consists of a finite set of guarded actions:

$$\langle guard \rangle \hookrightarrow \langle statement \rangle$$

The *guard* is a Boolean expression involving variables of p and its neighbors. The *statement* updates some variables of p . An action can be executed only if its guard evaluates to *true*; in this case, the action is said to be *enabled*. A node is said to be *enabled* if at least one of its actions is enabled. The local program at node p is modeled by a function `run` of type (`State → list Channel → Local_Env → (Channel → Channel) → option State`).

This function accesses the local topology and states around p . It takes as first argument the current state of p . The two other arguments are \mathcal{N}_p and ρ_p . These arguments allow the

¹Option type is used for partial functions which, by convention, return (`Some _`) when defined, and `None` otherwise. `None` is denoted by \perp in this paper.

function to access the local states of p 's neighbors. The returned value is the next state of node p if p is enabled, \perp otherwise. Note that `run` provides a functional view of the algorithm: it includes the whole set of possible actions, but returns a single result; this model is thus restricted to *deterministic algorithms*.²

2.3.2. Semantics. A *configuration* g of the system is defined as an instance of the states of all nodes in the system, *i.e.*, a function typed $(\text{Env} := \text{Node} \rightarrow \text{State})$. For a given node p and a configuration $(g: \text{Env})$, the term $(g\ p)$ represents the state of p in configuration g . Thanks to this encoding, we easily obtain the local configuration (type `Local_Env`) of node p by composing g and `peer` as a function $(\text{local_env } g\ p) := (\text{fun } (c: \text{Channel}) \Rightarrow \text{option_map } g\ (\text{peer } p\ c))$, where `option_map` $g\ (\text{peer } p\ c)$ returns $(g\ p')$ when $(\text{peer } p\ c)$ returns `Some` p' , and \perp otherwise. Hence, the execution of the algorithm on node p in the current configuration g is obtained by: $(\text{run } (g\ p)\ \mathcal{N}_p\ (\text{local_env } g\ p)\ \rho_p)$; it returns either \perp if the node is disabled or $(\text{Some } s)$ where $(s: \text{State})$ is the next state of p . We define $(\text{enabled_b } g\ p)$ as the Boolean value (type `bool`) `true` if node p is enabled in configuration g and `false` otherwise.

Assume the system is in some configuration g . If there exist some enabled nodes, a *daemon*³ selects a non-empty set of them; every chosen node *atomically* executes its algorithm, leading to a new configuration g' . The transition from g to g' is called a *step*. To model steps in Coq, we use functions with type $(\text{Diff} := \text{Node} \rightarrow \text{option State})$. We simply call *difference* a variable of type `Diff`. A difference contains the updated states of the nodes that actually execute some action during the step, and maps any other node to \perp . Steps are defined as a binary relation \mapsto over configurations expressed in Coq by the relation `Step`: $(\text{Step } g'\ g)$ holds for $g \mapsto g'$.⁴ It requires that there exists a difference d such that

- at least one node actually changes its state,,
- every update in d corresponds to the execution of the algorithm, namely, `run`
- and the next configuration, g' , is obtained applying the function $(\text{diff_eval } d\ g)$ given by: $\forall(p: \text{Node}), (g'\ p) = (d\ p)$ if $(d\ p) \neq \perp$, and $(g'\ p) = (g\ p)$ otherwise.

An *execution* of \mathcal{A} is a sequence of configurations $g_0\ g_1\ \dots\ g_i\ \dots$ such that $g_{i-1} \mapsto g_i$ for all $i > 0$. Executions may be finite or infinite and are modeled in Coq with the type

```
CoInductive Exec: Type :=
| e_one: Env → Exec
| e_cons: Env → Exec → Exec
```

and the predicate

```
CoInductive is_exec: Exec -> Prop :=
| i_one: ∀(g: Env), terminal g → is_exec (e_one g)
| i_cons: ∀(e: Exec) (g: Env),
          is_exec e → Step (Fst e) g → is_exec (e_cons g e)
```

where the keyword **CoInductive** generates a greatest fixed point capturing potentially infinite constructions⁵. Considering first the constructor `i_cons`, function $(\text{Fst } e)$ returns the first

²Finite non-determinism could be handled by having `run` output (list State) instead of (option State) .

³The daemon achieves the asynchrony of the system.

⁴Note the inverse order of the parameters in `Step`.

⁵As opposed to this, the keyword **Inductive** only captures finite constructions.

configuration of execution e . Thus, a variable (e : `Exec`) actually represents an execution of \mathcal{A} when (`is_exec e`) holds, since each pair of consecutive configurations g, g' in e satisfies (`Step g' g`).

Considering now the constructor `i_one`, proposition (`terminal g`) means that g is a *terminal* configuration, namely, no action of \mathcal{A} is enabled at any node in g . In our framework, a terminal configuration is any configuration g where `run` returns \perp for every node. Note that this predicate is decidable thanks to the emptiness test. The predicate `is_exec` requires that only executions that end with a terminal configuration are finite; every other execution is infinite.

As previously stated, each step from a configuration to another is driven by a *daemon*. In our case study, we assume that the daemon is *distributed* and *unfair*. *Distributed* means that while the configuration is not terminal, the daemon should select at least one enabled node, maybe more. *Unfair* means that there is no fairness constraint, *i.e.*, the daemon might never select an enabled node unless it is the only one enabled. Notice that the propositions `Step` and `is_exec` are sufficient to handle the distributed unfair daemon.

2.3.3. Read-Only Variables. We allow a part of a node state to be read-only: this is modeled with the type `R0State` and by the function (`R0_part: State → R0State`) which typically represents a subset of the variables handled in the `State` of the node. The projection `R0_part` is extended to configurations by the function (`R0Env_part g := (fun (p: Node) => R0_part (g p))`), which returns a value of type `R0Env`.

We add the property `R0_stable` to express the fact that those variables are actually read-only, namely no execution of `run` can change their values. From the assumption `R0_stable`, we show that any property defined on the read-only variables of a configuration is indeed preserved during steps.

The introduction of Read-Only variables has been motivated by the fact that we want to encompass the diversity of anonymity levels from the distributing computing literature, *e.g.*, fully anonymous, semi-anonymous, rooted, fully identified networks, *etc.* By default, our Coq model defines fully anonymous network thanks to the distinction between nodes (type `Node`) and channels (type `Channel`). We enriched our model to reflect other assumptions.

For example, consider the fully identified assumption. Identifiers are typically constant data, stored in the node states. In our model, they would be stored in the read-only part of the state. Furthermore, identifiers should be constant and unique all along the execution of the algorithm (see the assumption in the case study). This means they should be unique in the initial configuration and kept constant during the whole execution.

We define a predicate `Assume_R0` on `R0Env` (in the case of fully identified assumption, `Assume_R0` would express uniqueness of identifiers) that will be assumed in each initial configuration. From `R0_stable`, this property will remain true all along any execution. Furthermore, the predicate `Assume_R0` can express other assumptions on the network such as connected networks or tree networks (for this latter, see again the case study). As a shortcut, for any configuration (g : `Env`), we use notation `Assume g := (Assume_R0 (R0Env_part g))`.

2.4. Setoids. When using Coq function types to represent configurations and differences, we need to state pointwise function equality, which equates functions having equal values (extensional equality). The Coq default equality is inadequate for functions since it asserts equality of implementations (intensional equality). So, instead we chose to use the setoid

paradigm: we endow every base type with an *equivalence relation*. Setoids are commonly used in Coq for subsets, function sets, and to represent set-theoretic quotient sets (such as rational numbers or real numbers); in particular we make use of libraries `Coq.Setoids.Setoid` and `Coq.Lists.SetoidList`.

Consequently, every function type is endowed with a *partial equivalence relation* (*i.e.*, symmetric and transitive) which states that, given equivalent inputs, the outputs of two equivalent functions are equivalent. However, we also need reflexivity to reason about it, *i.e.*, functions are equivalent to themselves. In the context of partial equivalence relations, objects that are equal to themselves are said to be *proper* elements (in Coq: `Proper R x := R x x`). For example, all elements of base types are proper since we use equivalence relations. Proper functions are also called *compatible* functions or *relation morphisms*: they return equivalent results when executed with equivalent parameters. Through all the framework, we assume *compatible configurations and differences only*. We also prove compatibility (properness) for every function and predicate defined in the sequel. Additionally, we assume that equivalence relations on base types are decidable.

As an example, let us consider configurations. The equality for type `Node` is noted (`eqN: relation Node`) and assumes

```
(eqN_equiv: Equivalence eqN; eqN_dec: Decider eqN).
```

Note that (`relation Node`) stands for (`Node → Node → Prop`); (`Equivalence eqN`) defines the conjunction of reflexivity, symmetry, and transitivity of the relation `eqN`; (`Decider eqN`) expresses that the relation is decidable by $\forall(p\ p': \text{Node}), \{\text{eqN } p\ p'\} + \{\neg \text{eqN } p\ p'\}$, where $\{A\} + \{B\}$ is the standard Coq notation for computational disjunction between A and B , *i.e.*, Booleans carrying proofs of A or B .

The decidable equivalence relation on type `State`, noted `eqS` is defined similarly. Now, the equality between configurations, which are functions of type (`Env: Node → State`), is defined by `eqE := (eqN ==> eqS)`. This means that, for any two configurations (`g1 g2: Env`), (`eqE g1 g2`) is defined by

$$\forall(p1\ p2: \text{Node}), \text{eqN } p1\ p2 \rightarrow \text{eqS } (g1\ p1)\ (g2\ p2)$$

(in this model, (`eqN p1 p2`) means that $p1$ and $p2$ represent the same node in the network). Note that `eqE` is not reflexive *a priori*. We enforce reflexivity assuming compatible configurations only: any compatible configuration (`g: Env`) will satisfy

$$\text{Proper eqE } g := \forall (p1\ p2: \text{Node}), \text{eqN } p1\ p2 \rightarrow \text{eqS } (g\ p1)\ (g\ p2)$$

This means that for any two equivalent nodes $p1$ and $p2$, *i.e.*, such that (`eqN p1 p2`), we expect that (`g p1`) and (`g p2`) produce the same result with respect to `eqS`: (`eqS (g p1) (g p2)`).

3. SELF-STABILIZATION AND SILENCE

In this section, we express self-stabilization [24] in the locally shared memory model with composite atomicity using Coq properties.

3.1. Self-Stabilization. Consider a distributed algorithm \mathcal{A} . Let $\$$ be a predicate on executions (type $(\text{Exec} \rightarrow \text{Prop})$). \mathcal{A} is *self-stabilizing w.r.t. specification* $\$$ if there exists a predicate \mathbb{P} on configurations (type $(\text{Env} \rightarrow \text{Prop})$) such that:

- \mathbb{P} is *closed* under \mathcal{A} , *i.e.*, for each possible step $g \mapsto g'$, $(\mathbb{P} \ g)$ implies $(\mathbb{P} \ g')$:

$$\text{closure } \mathbb{P} := \forall (g \ g' : \text{Env}), \text{ Assume } g \rightarrow \text{ Step } g' \ g \rightarrow \mathbb{P} \ g'$$
- \mathcal{A} *converges* to \mathbb{P} , *i.e.*, every execution contains a configuration which satisfies \mathbb{P} :

$$\text{convergence } \mathbb{P} := \forall (e : \text{Exec}),$$

$$\text{Assume } (\text{Fst } e) \rightarrow \text{is_exec } e \rightarrow$$

$$\text{safe_suffix } (\text{fun } \text{suf} : \text{Exec} \Rightarrow \mathbb{P} \ (\text{Fst } \text{suf})) \ e$$

where $(\text{safe_suffix } \mathbb{P} \ e)$ inductively checks that execution e contains a suffix that satisfies \mathbb{P} .
- \mathcal{A} *meets* $\$$ from \mathbb{P} , *i.e.*, every execution which starts from a configuration where \mathbb{P} holds, satisfies $\$$:

$$\text{spec_ok } \$ \ \mathbb{P} := \forall (e : \text{Exec}),$$

$$\text{Assume } (\text{Fst } e) \rightarrow \text{is_exec } e \rightarrow \mathbb{P} \ (\text{Fst } e) \rightarrow \$ \ e.$$

The configurations which satisfy the predicate \mathbb{P} are called *legitimate configurations*. The following predicate characterizes the property of being self-stabilizing for an algorithm:

$$\text{self_stab } \$:= \exists \mathbb{P}, \text{ closure } \mathbb{P} \wedge \text{convergence } \mathbb{P} \wedge \text{spec_ok } \$ \ \mathbb{P}.$$

3.2. Silence. An algorithm is *silent* if the communication between the nodes is fixed from some point of the execution [26]. This latter definition can be transposed in the locally shared memory model as follows: \mathcal{A} is *silent* if all its executions are finite.

Inductive $\text{finite_exec} : \text{Exec} \rightarrow \text{Prop} :=$
 $| \text{f_one} : \forall (g : \text{Env}), \text{finite_exec } (e_one \ g)$
 $| \text{f_cons} : \forall (e : \text{Exec}) (g : \text{Env}),$
 $\quad \text{finite_exec } e \rightarrow \text{finite_exec } (e_cons \ g \ e).$
 $\text{silence} := \forall (e : \text{Exec}), \text{Assume } (\text{Fst } e) \rightarrow \text{is_exec } e \rightarrow \text{finite_exec } e.$

By definition, executions of a *silent and self-stabilizing* algorithm *w.r.t* some specification $\$$ end in configurations which are usually used as legitimate configurations, *i.e.*, satisfying \mathbb{P} . In this case, $\$$ can only allow constrained executions made of a single configuration which is legitimate; $\$$ is then noted $\$_{\mathbb{P}}$. To prove that \mathcal{A} is both silent and self-stabilizing *w.r.t.* $\$_{\mathbb{P}}$, we use, as commonly done, a sufficient condition which requires to prove that

- all terminal configurations of \mathcal{A} satisfy \mathbb{P} :

$$\text{P_correctness } \mathbb{P} := \forall (g : \text{Env}), \text{Assume } g \rightarrow \text{terminal } g \rightarrow \mathbb{P} \ g$$
- and all executions of \mathcal{A} are finite:

$$\text{termination} := \forall (g : \text{Env}), \text{Assume } g \rightarrow \text{Acc } \text{Step } g.$$

The latter property is expressed with $(\text{Acc } \text{Step } g)$ for every configuration g . The inductive proposition Acc is taken from `Library Coq.Init.Wf` which provides tools on well-founded induction. The accessibility predicate $(\text{Acc } \text{Step } g)$ is translated into

$$(\forall (g' : \text{Env}), \text{Step } g' \ g \rightarrow \text{Acc } \text{Step } g') \rightarrow \text{Acc } \text{Step } g$$

Namely, the base case of induction holds when no step is possible from current configuration g and then, inductively, any configuration g' that eventually reaches such a terminal configuration satisfies (**Acc Step** g').

The sufficient condition, used to prove that an algorithm is both silent and self-stabilizing, is expressed and proven by:

Lemma `silent_self_stab` ($P: \text{Env} \rightarrow \text{Prop}$):
`P_correctness P \wedge termination \rightarrow silence \wedge self_stab $P`.

4. TOOLS FOR PROVING TERMINATION

Usual termination proofs are based on some global potential built from local ones. For example, local potentials can be integers and the global potential can be the sum of them. In this case, the argument for termination may be, for example, the fact that the global potential is lower bounded and strictly decreases at each step of the algorithm. Global potential decrease is due to the modification of local states at some nodes, however studying aggregators such as sums may hide scenarios, making the proof more complex. Instead, we build here a global potential as the multiset containing the local potential of each node and provide a sufficient condition for termination on this multiset. Our method is based on two criteria that are sufficient to meet the conditions for using the Dershowitz-Manna well-founded ordering on multisets [22]. Given those criteria, we can show that the multiset of (local) potentials globally decreases at each step. Note that instead of developing our own library, we have built specialized termination theorems on top of existing work, namely the Coq Standard Library for the lexicographic product ordering and the CoLoR library [5] for multisets and the Dershowitz-Manna ordering.

We also provide tools for algorithms that have (local) priorities on actions, *e.g.*, an action i is enabled at node p only if every action j , with $j < i$, is disabled at p . The overall idea is to ease the proof by considering a given set of actions separately from the others and to prove the termination of the algorithm assuming that only this set of actions is executed. Once termination is proved for each set of actions separately, we use tuples of multisets ordered with the lexicographical order to prove the termination of the whole algorithm.

4.1. Steps. One difficulty we faced, when trying to apply our method straightforwardly, is that we cannot always define the local potential function at a node without assuming some properties on its local state, and so on the associated configuration. Thus, we had to assume the existence of some stable set of configurations in which the local potential function can be defined. When necessary, we use our technique to prove termination of a subrelation of the relation **Step**, provided that the algorithm has been initialized in the required stable set of configurations. This point is modeled by a predicate on configurations, (**safe**: $\text{Env} \rightarrow \text{Prop}$), and a type `safeEnv := { g | safe g }` which represents the set of *safe configurations* into which we restrict the termination proof. Precisely, `safeEnv` is a type whose values are ordered pairs containing a term g and a proof of (**safe** g). Safe configurations should be stable, *i.e.*, it is assumed that no step can exit from the set using the proposition:

`stable_safe := $\forall(g\ g': \text{Env}), \text{safe } g \rightarrow \text{Step } g' \ g \rightarrow \text{safe } g'$` .

Steps of the algorithm for which termination will be proven is defined by

`safeStep sg2 sg1 := Step (getEnv sg2) (getEnv sg1)`

with `(sg1 sg2: safeEnv)` two safe configurations and where `(getEnv sg1)` (resp. `(getEnv sg2)`) accesses the actual configuration, of type `Env`, of `sg1` (resp. `sg2`). We aim at proving that this relation is well-founded. Since we know that property `safe` is stable from `stable_safe`, we have the following lemma (proven by induction, starting from the assumption):

Lemma `Acc_Algo_Multiset`:

`well_founded safeStep → ∀(g: Env), safe g → Acc Step g.`

Note that `(well_founded R := ∀a, Acc R a)`, like `Acc`, is taken from the standard Coq Library `Coq.Init.Wf`. Hence, to prove termination of the algorithm as defined in Section 3, we prove that `safeStep` is well-founded and use the above lemma to guarantee that the whole algorithm terminates when initiated from any safe configuration.

We also allow restrictions on the kind of steps that will be handled in termination proofs because, in some cases, it is easier to partition steps and to prove termination of some kind of steps separately from the others. We introduce the predicate `QTrans: safeEnv → safeEnv → Prop` for that purpose. The relation for which termination will be proven is then defined by

`safeQStep sg2 sg1 := safeStep sg2 sg1 ∧ QTrans sg2 sg1`

for any two safe configuration `(sg1 sg2: safeEnv)`. Note that proving termination for all safe steps just consists in applying the method with `QTrans` defined as a tautology.

4.2. Potential. We assume that within safe configurations, each node can be endowed with a potential value obtained using function `(pot: safeEnv → Node → Mnat)`. Notice that `Mnat` simply represents natural numbers⁶ encoded using the type from Library `CoLoR.MultisetNat` [5]; it is equipped with the usual equivalence relation, noted `=P`, and the usual well-founded order on natural numbers, noted `<P`.

4.3. Multiset Ordering. We recall that a multiset of elements in the setoid P endowed with its equivalence relation $=_P$, is defined as a set containing *finite numbers of occurrences (w.r.t. $=_P$) of elements of P* . Such a multiset is usually formally defined as a multiplicity function $m : P \rightarrow \mathbb{N}_{\geq 1}$ which maps any element to its number of occurrences in the multiset. We focus here on *finite multisets*, namely, multisets whose multiplicity function has finite support. We define equality between multisets, noted \approx , as the equality between multiplicity functions. Now, we assume that P is also ordered using relation $<_P$, compatible with $=_P$. We use the Dershowitz-Manna order on finite multisets [22] defined as follows: the multiset N is smaller than the multiset M , noted $N \prec M$, if and only if there are three multisets X , Y and Z such that

- N is obtained from M by removing all elements in X and adding all elements in Y . Elements in Z are present in both M and N , and '+' between multisets means adding multiplicities, namely

$$M \approx Z + X \wedge N \approx Z + Y$$

- at least one element is removed, *i.e.*,

$$X \not\approx \emptyset$$

⁶Natural numbers cover many cases and we expect the same results when further extending to other types of potential.

- each element that is added (*i.e.* in Y) must be smaller (*w.r.t.* $<_P$) than some removed element (*i.e.* in X), that is:

$$\forall y \in Y, \exists x \in X, y <_P x$$

It is shown [22] that if $<_P$ is a well-founded order, then the corresponding order $<$ is also well-founded.

In our context, we consider finite multisets over \mathbf{Mnat} , (*i.e.*, $=_P$ is $=_{\mathbf{P}}$ and $<_P$ stands for $<_{\mathbf{P}}$). We have chosen to model them as lists of elements of \mathbf{Mnat} and we build the potential of a configuration as the multiset of the potentials of all nodes, *i.e.*, a multiset of (local) potentials of a configuration (`sg: safeEnv`) is defined by

```
Pot sg := List.map (pot sg) all_nodes
```

where `all_nodes` is the list of all nodes in the network (see Section 2) and `(List.map f l)` is the standard operation that returns the list of values obtained by applying `f` to each element of `l`. The corresponding Dershowitz-Manna order is defined using the library `CoLoR` [5]. The library also contains the proof that $(\text{well_founded } <_P) \rightarrow (\text{well_founded } <)$.

Using this latter result and the standard result which proves $(\text{well_founded } <_P)$, we easily deduce $(\text{well_founded } <)$.

4.4. Termination Theorem. Proving the termination of a set of safe steps then consists in showing that for any such a step, the corresponding global potential decreases *w.r.t.* the Dershowitz-Manna order $<$. We call this proof goal *safe inclusion*:

```
safe_incl :=  $\forall$ (sg1 sg2: safeEnv),
             safeQStep sg2 sg1  $\rightarrow$  (Pot sg2) < (Pot sg1).
```

We establish a sufficient condition made of two criteria on node potentials which validates `safe_incl`. The *Local Criterion* finds for any node `p` whose potential has increased, a witness node `p'` whose potential has decreased from a value that is even higher than the new potential of `p`:

```
Hypothesis local_crit:  $\forall$ (sg1 sg2: safeEnv), safeQStep sg2 sg1  $\rightarrow$ 
 $\forall$ (p: Node), (pot sg1 p) <_P (pot sg2 p)  $\rightarrow$ 
 $\exists$ (p': Node), (pot sg1 p')  $\not\leq_P$  (pot sg2 p')  $\wedge$ 
                (pot sg2 p) <_P (pot sg1 p').
```

The *Global Criterion* exhibits, at any step, a node whose potential has changed:

```
Hypothesis global_crit:  $\forall$ (sg1 sg2: safeEnv), safeQStep sg2 sg1  $\rightarrow$ 
 $\exists$ (p: Node), (pot sg2 p)  $\not\leq_P$  (pot sg1 p).
```

Assuming both hypotheses (see Section 7 for the instantiation of these criteria), we are able to prove `safe_incl` as follows: we define Z as the multiset of local potentials that did not change, and X (resp. Y) as the complement of Z in the multiset of local potentials $(\text{Pot } \text{sg1})$ (resp. $(\text{Pot } \text{sg2})$). Global criterion is used to show that $X \neq \emptyset$, and local criterion is used to show that $\forall y \in Y, \exists x \in X, y <_P x$. Since any relation included in a well-founded order is also well-founded, we get that relation `safeQStep` is well-founded.

4.5. Lexicographical Order. We now provide tools to divide a termination proof according to given subsets of safe steps. Usually, algorithms made of several actions enforce priority between them. For example, actions can be prioritized so that only one action is enabled at a given node at one time; namely, the second action can only be enabled at a node if the first action is disabled, and so on. In such a case, it is often more convenient to consider each action separately, *i.e.*, show that when nodes execute a particular action only, the algorithm converges and then generalize by gradually incorporating the other ones. To that goal, we consider a partition of safe steps; we detail here the case for two subsets, one having priority over the other.

We consider two relations over safe configurations, noted (**Trans1 Trans2**: $\text{safeEnv} \rightarrow \text{safeEnv} \rightarrow \mathbf{Prop}$). This may represent steps induced by two different actions. We assume for each a local potential (pot1 pot2 : $\text{safeEnv} \rightarrow \text{Node} \rightarrow \text{Mnat}$) and as in Section 4.3, we build the corresponding multisets of local potentials, for any safe configuration sg , by ($\text{Pot}i \text{ sg} := \text{List.map } (\text{pot}i \text{ sg}) \text{ all_nodes}$) with $i \in \{1, 2\}$. Here we expect that **Trans1** has priority on **Trans2**. To encode this priority, we require that when a step from **Trans2** occurs, the multisets of potentials measured by the action from **Trans1** (**Pot1**) is left unchanged:

$$\begin{aligned} \text{Hdisjoint} &:= \forall(\text{sg sg}': \text{safeEnv}), \\ &\quad \text{Trans2 sg}' \text{ sg} \rightarrow (\text{Pot1 sg}) \approx (\text{Pot1 sg}'). \end{aligned}$$

The idea is to prove that steps from **Trans1** and **Trans2** taken together, namely the union of the relations **Trans1** and **Trans2**, converge, provided that steps from **Trans1** (resp. **Trans2**) terminate when taken separately.

We use the following ordering relation: for any two safe configurations ($\text{sg sg}'$: safeEnv), $\text{sg} <_{lex} \text{sg}'$ is defined by

$$\begin{aligned} &(\text{Pot1 sg}) \prec (\text{Pot1 sg}') \\ \vee &(\text{Pot1 sg}) \approx (\text{Pot1 sg}') \wedge (\text{Pot2 sg}) \prec (\text{Pot2 sg}'). \end{aligned}$$

$<_{lex}$ is built using the lexicographical order from the Library `ColoR`, applied on pairs of multisets of local potentials. We also use results from this library to show that $<_{lex}$ is well founded, as far as the order on local potentials, $<_p$, is.

Now, the argument is the same as for the Termination Theorem (see 4.4): any order included in a well-founded order is also well-founded. We aim at showing that ($\text{Trans1} \cup \text{Trans2}$) is well-founded using the following argument: ($\text{Trans1} \cup \text{Trans2}$) should be included into relation $<_{lex}$. To obtain this, we first use the assumption **Hdisjoint** that ensures priorities between **Trans1** and **Trans2**. Second, we require safe inclusion for both relations **Trans1** and **Trans2**, namely:

$$\begin{aligned} \text{safe_incl1} &:= \forall(\text{sg sg}': \text{safeEnv}), \\ &\quad \text{Trans1 sg}' \text{ sg} \rightarrow (\text{Pot1 sg}') \prec (\text{Pot1 sg}). \\ \text{safe_incl2} &:= \forall(\text{sg sg}': \text{safeEnv}), \\ &\quad \text{Trans2 sg}' \text{ sg} \rightarrow (\text{Pot2 sg}') \prec (\text{Pot2 sg}). \end{aligned}$$

This ensures that, at any safe step of **Trans1** (resp. **Trans2**), the corresponding global potential decreases. From this we obtain our goal:

$$\mathbf{Lemma} \text{ union_lex_wf2: well_founded } (\text{Trans1} \cup \text{Trans2}).$$

Our framework also contains the same results for three relations; the lemma corresponding to the above one is called `union_lex_wf3`.

Proving the two assumptions are satisfied (namely, the priorities between `Trans1` and `Trans2`, and safe inclusion for both relations `Trans1` and `Trans2`) implies the termination of the algorithm, while considering both relations `Trans1` and `Trans2` separately. In particular, we can use the Termination Theorem given in 4.4 to show the safe inclusion of each relation, by instantiating the local and global criteria for each relation.

5. TOOLS FOR QUANTITATIVE PROPERTIES

To handle some quantitative properties of an algorithm, we have to set up a library dealing with cardinality of sets in general and also cardinals of *finite* sets. The need for a new library arises from the absence of setoid-compatible formalization of set cardinality. For example, the `Ensembles` and `Finite_Sets` module from the Coq standard library relates subset (predicates) over a fixed universe type `U`, and elements are considered up to Leibniz equality (see definition of `Singleton`). Instead, we build a theory allowing to compare the cardinality of arbitrary setoids (*i.e.* of sets of equivalence classes) built on top of distinct types instead of subsets over the same type.

The library contains basic properties about set operations such as Cartesian product, disjoint union, and subset. Proofs are conducted using standard techniques.

5.1. Cardinality on Setoids. To be able to order cardinalities, we define a property, called `Inj`, on a pair of setoids $(A, =_A)$ and $(B, =_B)$ which requires the existence of an injective and compatible function, `inj`, from A to B whose domain is A . Namely:

- `Inj_compat`: `inj` is compatible (see Section 2.4),
- `Inj_left_total`: domain of `inj` is A , *i.e.*, any element in A is related to at least one element in B ,
- `Inj_left_unique`: `inj` is injective, *i.e.*, any element in B is related to at most one (w.r.t. $=_A$) element in A .

Relation `Inj` is proven reflexive and transitive. We model cardinality ordering using the three-valued type `(Card Prop := Smaller | Same | Larger)` and the following property `Card`. `Card` distinguishes the different ways `Inj` can apply to pairs of setoids:

- `(Card Smaller A B)`⁷ is defined by `(Inj A B)` which expresses that A has a cardinal smaller or equal to that of B , w.r.t. equalities $=_A$ and $=_B$;
- Similarly, `(Card Larger A B)` is defined by `(Inj B A)`
- and `(Card Same A B)` by `(Inj B A ∧ Inj A B)`.

`(Card prop)` is reflexive and transitive for any value of `prop` in `Card_prop`. It is also antisymmetric in the sense that `(Card Smaller)` and `(Card Larger)` implies `(Card Same)` for a given pair of setoids (trivial from the definitions).

⁷We omit parameters $=_A$ and $=_B$ for better readability.

5.2. Finite Cardinalities. We now focus on finite setoids and define tools to express their cardinalities. We first define, for a given natural number N , the setoid

$$\mathcal{M}_N := \{i: \text{nat} \mid i < N\}.$$

\mathcal{M}_N simply models the set of natural numbers $\{0, 1, \dots, N - 1\}$.⁸ We first proved that `Inj` captures finite cardinality ordering

$$\text{Lemma Inj_le_iff: } \forall (m \ n: \text{nat}), \text{ Inj } \mathcal{M}_m \ \mathcal{M}_n \leftrightarrow m \leq n.$$

and the corresponding corollaries with `Card`, *e.g.*,

$$\forall (m \ n: \text{nat}), \text{ Card Smaller } \mathcal{M}_m \ \mathcal{M}_n \leftrightarrow m \leq n.$$

(similar corollaries exist for `Larger` and `Same`). The following predicate `Num_Card` is then used to express that a setoid A has cardinality at least (resp. at most, resp. equal to) some natural number n with `Num_Card prop A n := (Card prop A \mathcal{M}_n)` where `prop` is any `Card_Prop`. For instance, `(Num_Card Smaller A n)` means that A contains at most n elements w.r.t. $=_A$.

5.3. Cartesian Products. We developed results about Cartesian products. First, the Cartesian product is monotonic *w.r.t.* cardinality:

$$\text{Lemma Inj_prod: } \forall \text{prop}, \text{ Card prop } A_1 \ A_2 \rightarrow \text{ Card prop } B_1 \ B_2 \rightarrow \\ \text{Card prop } (A_1 \times B_1) \ (A_2 \times B_2).$$

where $(A_1, =_{A_1})$, $(A_2, =_{A_2})$, $(B_1, =_{B_1})$, $(B_2, =_{B_2})$ are any setoids. Now, we showed that:

$$\forall \ n \ m: \text{nat}, \text{ Card Same } (\mathcal{M}_n \times \mathcal{M}_m) \ \mathcal{M}_{n \times m}$$

namely, the Cartesian product of $\mathcal{M}_n = \{0, \dots, n - 1\}$ and $\mathcal{M}_m = \{0, \dots, m - 1\}$ contains the same number of elements as $\mathcal{M}_{n \times m} = \{0, \dots, n \times m - 1\}$. This latter result is shown using encoding functions from $\mathcal{M}_n \times \mathcal{M}_m$ to $\mathcal{M}_{n \times m}$ and from $\mathcal{M}_{n \times m}$ to $\mathcal{M}_n \times \mathcal{M}_m$. This intermediate result allows one to easily deduce that the cardinality of a Cartesian product is the product of cardinalities:

$$\forall \text{prop } (n \ m: \text{nat}), \text{ Num_Card prop } A \ n \rightarrow \text{ Num_Card prop } B \ m \rightarrow \\ \text{Num_Card prop } (A \times B) \ (n \times m)$$

5.4. Disjoint Unions. We developed similar lemmas about the disjoint union of sets, noted `+`. The main results is:

$$\forall \text{prop } (n \ m: \text{nat}), \text{ Num_Card prop } A \ n \rightarrow \text{ Num_Card prop } B \ m \rightarrow \\ \text{Num_Card prop } (A + B) \ (n + m)$$

⁸In Coq, the values of \mathcal{M}_N are ordered pairs containing a natural number i and a proof of $i < N$ and \mathcal{M}_N is equipped with the standard equality on type `nat`, wrapped to be able to compare values of type \mathcal{M}_N .

5.5. Subsets. We proved many toolbox results, about subsets, which are expressed using `Card` as well as `Num_Card`. For instance,

- any subset of a set A has `Smaller` cardinality than that of A ,
- a set is one of its subsets with `Same` cardinality,
- the empty subset contains 0 element,
- a non-empty set contains at least 1 element,
- a singleton contains exactly one element.

5.6. Number of Elements in Lists. To prove the existence of finite cardinality for finite setoids, we use lists, since, for example, the setoid of nodes of the network is encoded in our framework as the list `all_nodes`. We now consider a setoid A , whose equality $=_A$ satisfies the classical excluded middle property ($\forall a1\ a2: A, a1 =_A a2 \vee a1 \neq_A a2$) and a predicate function ($P: A \rightarrow \mathbf{Prop}$), which also satisfies the classical excluded middle property ($\forall a: A, P\ a \vee \neg P\ a$). Under these conditions, we can prove:

$$\forall (l: \text{list } A), \exists (n: \text{nat}), \text{Num_Card Same } \{a: A \mid P\ a \wedge a \in_{=A} l\} n$$

namely, for any list l , the set of elements in l (w.r.t. $=_A$) which satisfies predicate P has finite cardinality n . Or, equivalently, assuming the existence of a list l which contains every element of type A , we get that the number of elements which satisfy P is finite:

$$\forall (l: \text{list } A), (\forall (a: A), a \in_{=A} l) \rightarrow \exists (n: \text{nat}), \text{Num_Card Same } \{a: A \mid P\ a\} n.$$

When predicate function P returns `True` for all argument values, this provides the number of elements of list l , up to $=_A$.

6. k -CLUSTERING ALGORITHM $\mathcal{C}(k)$

We have certified a non trivial part of the silent self-stabilizing algorithm proposed in [19]. Given a non-negative integer k , this algorithm builds a k -clustering of a bidirectional connected network $G = (V, E)$ containing at most $\lfloor \frac{n-1}{k+1} \rfloor + 1$ k -clusters, where n is the number of nodes. A k -cluster of G is a set $C \subseteq V$, together with a designated node $h \in C$, such that each member of C is within distance k of h .⁹ A k -clustering is then a partition of V into distinct k -clusters.

The algorithm proposed in [19] is actually a hierarchical collateral composition [20] of two silent self-stabilizing sub-algorithms: the former builds a rooted spanning tree, the latter is a k -clustering construction which stabilizes once a rooted spanning tree is available in the network. In this paper, we focus in the certification of the second part, namely, the construction, in a self-stabilizing and silent way, of a k -clustering on a rooted spanning tree containing at most $\lfloor \frac{n-1}{k+1} \rfloor + 1$ clusterheads. The k -clustering is actually organized as a spanning forest. Each k -cluster is an in-tree of height at most k rooted at its clusterhead. Moreover, each k -cluster is colored with the identifier of its clusterhead. Hence, each node p should compute the identifier of its clusterhead c and the channel corresponding to its *parent link*, that is, the link from p to its *parent* in the k -cluster, that is, the unique neighbor of p on the shortest path from p to c in the k -cluster.

The code of the algorithm, called $\mathcal{C}(k)$, is given in Algorithm 1. We have used our framework to encode $\mathcal{C}(k)$, its assumptions and specification, and to certify its correctness.

⁹The distance $\|p, q\|$ between two nodes p and q is the length of a shortest path linking p to q in G .

Algorithm 1 $\mathcal{C}(k)$, code for each process p

Constant Input: $\text{Id}(p) \in \text{Ids}$; $\text{Par}(p) \in \mathcal{N}_p \cup \{\perp\}$
Variable: $p.\alpha \in \mathbb{Z}$; $p.\text{parC} \in \mathcal{N}_p \cup \{\perp\}$; $p.\text{headC} \in \text{Ids}$
Predicates:

$$\begin{aligned} \text{IsShort}(p) &\equiv p.\alpha < k \\ \text{IsTall}(p) &\equiv p.\alpha \geq k \\ k\text{Dominator}(p) &\equiv (p.\alpha = k) \vee (\text{IsShort}(p) \wedge \text{Par}(p) = \perp) \end{aligned}$$

Macros:

$$\begin{aligned} \text{ShortChildren}(p) &= \{q \in \mathcal{N}_p \mid \text{Par}(q) = \rho_p(q) \wedge \text{IsShort}(q)\} \\ \text{TallChildren}(p) &= \{q \in \mathcal{N}_p \mid \text{Par}(q) = \rho_p(q) \wedge \text{IsTall}(q)\} \\ \text{MaxAShort}(p) &= \max(\{q.\alpha \mid q \in \text{ShortChildren}(p)\} \cup \{-1\}) \\ \text{MinATall}(p) &= \min(\{q.\alpha \mid q \in \text{TallChildren}(p)\} \cup \{2k + 1\}) \\ \text{MinCMinATall}(p) &= \text{if } \text{TallChildren}(p) = \emptyset \text{ then } \perp \\ &\quad \text{else } \min_{<_C} \{q \in \text{TallChildren}(p) \mid q.\alpha = \text{MinATall}(p)\} \\ \text{Alpha}(p) &= \text{if } \text{MaxAShort}(p) + \text{MinATall}(p) \leq 2k - 2 \text{ then } \text{MinATall}(p) + 1 \\ &\quad \text{else } \text{MaxAShort}(p) + 1 \\ \text{ParC}(p) &= \text{if } p.\alpha = k \text{ then } \perp \\ &\quad \text{else if } \text{IsShort}(p) \text{ then } \text{Par}(p) \text{ else } \text{MinCMinATall}(p) \\ \text{HeadC}(p) &= \text{if } k\text{Dominator}(p) \text{ then } \text{Id}(p) \text{ else} \\ &\quad \text{else if } p.\text{parC} \notin \mathcal{N}_p \text{ then } p.\text{headC} \text{ else } p.\text{parC}.\text{headC} \end{aligned}$$

Actions:

$$\begin{aligned} p.\alpha \neq \text{Alpha}(p) &\quad \hookrightarrow \quad p.\alpha \leftarrow \text{Alpha}(p) \\ p.\alpha = \text{Alpha}(p) \wedge p.\text{parC} \neq \text{ParC}(p) &\quad \hookrightarrow \quad p.\text{parC} \leftarrow \text{ParC}(p) \\ p.\alpha = \text{Alpha}(p) \wedge p.\text{parC} = \text{ParC}(p) \wedge p.\text{headC} \neq \text{HeadC}(p) &\quad \hookrightarrow \quad p.\text{headC} \leftarrow \text{HeadC}(p) \end{aligned}$$

6.1. Assumptions.

6.1.1. Unique Identifiers. The algorithm $\mathcal{C}(k)$ requires that nodes are uniquely identified: we assume a datatype for identifiers, noted Ids , which is endowed with a decidable equivalence relation noted eqId . Each node p is equipped with a constant input $\text{Id}(p)$ of type Ids that represents its identifier. We use the predicate uniqueId to represent uniqueness of the identifiers as follows:

$$\text{uniqueID } \text{Id} := \forall (p1 \ p2: \text{Node}), \text{eqId } \text{Id}(p) \ \text{Id}(p2) \rightarrow \text{eqN } p1 \ p2$$

6.1.2. Spanning Tree. We denote the directed spanning tree and its root by T and r , respectively: the knowledge of T is locally distributed at each node p using the constant input $\text{Par}(p) \in \mathcal{N}_p \cup \{\perp\}$. When $p \neq r$, $\text{Par}(p) \in \mathcal{N}_p$ and designates its parent in the tree (precisely, the channel outgoing to its parent). Otherwise, p is the root and $\text{Par}(p) = \perp$.

We express the assumption about the spanning tree using predicate `(span_tree r Par)`. This predicate checks that the graph T induced by `Par` is a subgraph of G which actually encodes a spanning tree rooted at r by the conjunction of

- r is the unique node such that $\text{Par}(r) = \perp$,
- $\text{Par}(p)$, for every non-root node p , is an existing channel outgoing from p ,
- T contains no loop.

From the last point, we show that, since the number of nodes is finite, the relation extracted from `Par` between nodes and their parents (resp. children) in T is well-founded. We call this result `WF_par` (resp. `WF_child`) and express it using `well_founded`.

6.1.3. Predicate `Assume_cl`. We instantiate the predicate `Assume_R0` to express that in any configuration $(g: \text{Env})$, G is bidirectional, identifiers are unique, and a rooted spanning tree is available in G (*n.b.*, this latter also implies that G is connected):

$$\text{Assume}_{cl} \ g := \text{sym_net} \wedge \text{uniqueId Id} \wedge \exists r, \text{span_tree } r \ \text{Par}.$$

6.2. Specification. The goal of algorithm $\mathcal{C}(k)$ is to compute a k -clustering using the spanning tree T . We consider any positive parameter k , (here, k is taken in \mathbb{Z} , as for other numbers, and assumed to be positive) and we model the k -clustering, *i.e.*, the output of the algorithm, using the predicate `kCluster`: for a given terminal configuration $(g: \text{Env})$, the proposition $(\text{kCluster } g \ h \ p)$ means that node p is in the k -cluster of node h and h is a clusterhead; precisely, h is the clusterhead of the k -cluster $\{p \mid \text{kCluster } g \ h \ p\}$. Note that using this definition, the fact that h is a clusterhead is given by the predicate $(\text{clusterHead } g \ h := \text{kCluster } g \ h \ h)$. The predicate `kCluster` actually designates a k -clustering when

- for any clusterhead h , the set $\{p \mid \text{kCluster } g \ h \ p\}$ actually represents a k -cluster, namely, for every node p in this set, there exists a path in this k -cluster (*i.e.*, the path is made of nodes q such that $\text{kCluster } g \ h \ q$) of length smaller than or equal to k from p to h , and
- the set of k -clusters is a partition of the set of nodes, or equivalently, every node belongs to a k -cluster and the intersection of any two distinct k -clusters is empty.

The complete check for k -clustering is performed using the conjunction of the two following predicates on configuration $(g: \text{Env})$:

$$\begin{aligned} \text{kCluster_OK } g := & \forall (h: \text{Node}), \text{clusterHead } g \ h \rightarrow \\ & \forall (p: \text{Node}), \text{kCluster } g \ h \ p \rightarrow \\ & \exists (\text{path}: \text{list Node}), \text{is_path } h \ \text{path } p \wedge \\ & (\forall (q: \text{Node}), q \in \text{path} \rightarrow \text{kCluster } g \ h \ q) \wedge \\ & (\text{length path}) \leq k \end{aligned}$$

$$\begin{aligned} \text{partition_OK } g := & \forall (p: \text{Node}), \\ & (\exists (h: \text{Node}), \text{kCluster } g \ h \ p) \wedge \\ & (\forall (h \ h': \text{Node}), \text{kCluster } g \ h \ p \rightarrow \text{kCluster } g \ h' \ p \rightarrow \text{eqN } h \ h') \end{aligned}$$

where predicate `is_path` detects if the list of nodes `path` actually represents a path in the network between the nodes h and p , and `length` computes the length of the path.

Actually, the algorithm computes a stronger specification. First, it ensures that there are no more than $\lfloor \frac{n-1}{k+1} \rfloor + 1$ clusterheads in any terminal configuration $(g: \text{Env})$:

$$\text{count_OK } g := (n - 1) \geq (k + 1)(|\text{CH}| - 1)$$

where CH is the set of clusterheads, given by $\{ h: \text{Node} \mid \text{clusterHead } g \ h \}$. Second, in a terminal configuration g , each node knows the identifier of its clusterhead and the channel corresponding to its parent link in the k -cluster:

- there exists a local function (on state of a node), ($\text{clusterHeadID}: \text{State} \rightarrow \text{Ids}$), which provides the identifier of the clusterhead of the node, and
- there exists a local function, ($\text{clusterParent}: \text{State} \rightarrow \text{Channel}$), which returns the channel outgoing to the parent of the node in the k -cluster.

This provides the third part of the specification, for a configuration ($g: \text{Env}$):

$$\begin{aligned} \text{kCluster_strong } g := & \\ & \forall (h \ p: \text{Node}), \text{ kCluster } g \ h \ p \leftrightarrow \text{eqId } \text{Id}(h) \ (\text{clusterHeadID } (g \ p)) \\ & \wedge \\ & \forall (h: \text{Node}), \text{ clusterHead } g \ h \rightarrow \\ & \quad \forall (p: \text{Node}), \text{ kCluster } g \ h \ p \rightarrow \\ & \quad \exists (\text{path}: \text{list } \text{Node}), \text{ agreed_cluster_path } g \ h \ \text{path } p \wedge \\ & \quad \quad (\text{length } \text{path}) \leq k \end{aligned}$$

where ($\text{agreed_cluster_path } g \ h \ \text{path } p$) is true if

- ($\text{cluster_path } g \ h \ \text{path } p$) holds, meaning that path is a *cluster path* in the network linking h to p , *i.e.*, the path described by the values of the clusterParent pointers in g , and
- every node in path declares the same clusterHeadID in g .

Note that for any configuration ($g: \text{Env}$), ($\text{kCluster_strong } g$) enforces ($\text{kCluster_OK } g$). Hence, the complete specification is given by the conjunction

$$\mathbb{P}_{cl} \ g := (\text{kCluster_strong} \wedge \text{partition_OK } g \wedge \text{count_OK } g)$$

for a configuration ($g: \text{Env}$).

6.3. Algorithm $\mathcal{C}(k)$ in Coq. We translate $\mathcal{C}(k)$ into the type `Algorithm`. First, the *state* of each node p contains, in addition to $\text{Id}(p)$ and $\text{Par}(p)$,

- an integer variable $p.\alpha$,
- a variable $p.\text{parC}$ which is either a channel or \perp , and
- a variable $p.\text{headC}$ of type Ids .

Hence, we have instantiated the `State` of a node as a record containing fields ($\text{Par}: \text{option Channel}$), ($\text{Id}: \text{Ids}$), ($\alpha: \mathbb{Z}$), ($\text{headC}: \text{Ids}$) and ($\text{parC}: \text{option Channel}$). Par and Id are declared as *read-only* variables.

Note that to be able to compute a path from each node to its clusterhead, the algorithm requires that channels are totally ordered (to be able to compute the minimum value on a set of channels). Hence we assume a strict total order $<_C$ on the type `Channel`. Furthermore, we chose to encode every number in the algorithm as integers in \mathbb{Z} , as α is, since some of them may be negative (see *MaxAShort*) and computations use minus (see *Alpha*).

Now, *every predicate and macro* of Algorithm 1 can be directly encoded in Coq: for a node p and a current configuration g , mainly all of them depend on \mathcal{N}_p , ρ_p , ($g \ p$), and ($\text{local_env } g \ p$); then the translation is quasi-syntactic (see `Library KClustering_algo` in the online browsing) and provides a definition of `run`. The definition of $\mathcal{C}(k)$, of type `Algorithm`, comes with a proof that `run` is compatible, as a composition of compatible functions, and

also with a straightforward proof of `RO_stable` which asserts that the read-only parts of the state, `Par` and `Id`, are constant during steps, when applying `run`.

6.4. Overview of $\mathcal{C}(k)$. A *k-hop dominating set* of a graph is a subset Dom of nodes such that every node of the graph is within distance k from at least one node of Dom . The k -clustering problem is related to the notion of k -hop dominating set, since the set of clusterheads of any k -clustering is, by definition, a k -hop dominating set.

Algorithm $\mathcal{C}(k)$ builds a k -clustering in two phases. During the first phase, $\mathcal{C}(k)$ computes the set of clusterheads as a k -hop dominating set of the spanning tree T (and so of G), using the variables α and the first action. The second phase consists of building (using variables $parC$ and $headC$, and the two other actions) a spanning forest : Algorithm $\mathcal{C}(k)$ computes each k -cluster as an in-tree of height at most k rooted at one of the already computed clusterhead. Moreover, each k -cluster will be colored with the identifier of its clusterhead.

6.4.1. Building Dom . Dom is constructed in a bottom-up fashion starting from the leaves of T , using the values of $p.\alpha$ for all p . Precisely, Dom is defined as the set of nodes p such that the predicate $kDominator(p)$ holds, namely, when $p.\alpha = k$, or $p.\alpha < k$ and $p = r$ (*i.e.*, p is the root). The goal of variable $p.\alpha$ at each node p is twofold. First, it allows to determine a path of length at most k from p to a particular node q of Dom which acts as a *witness* for guaranteeing the k -hop domination of Dom . Consequently, q will be denoted as $Witness(p)$ in the following. Second, once correctly evaluated, the value $p.\alpha$ is equal to $\|p, x\|$, where x is the furthest node in $T(p)$, the subtree of T rooted at p , that has the same witness as p .

The algorithm divides processes into *short* and *tall* according to the value of their α -variable: if p satisfies $IsShort(p)$, *i.e.*, $p.\alpha < k$, then p is said to be *short*; otherwise, p satisfies $IsTall(p)$ and is said to be *tall*. In a terminal configuration, the meaning of $p.\alpha$ depends on whether p is *short* or *tall*.

If p is *short*, we have two cases: $p \neq r$ or $p = r$. In the former case, $Witness(p) \in Dom$ is outside of $T(p)$, that is, the path from p to $Witness(p)$ goes through the parent link of p in the tree, and the distance from p to $Witness(p)$ is at most $k - p.\alpha$. See, for example, in Configuration (I) of Figure 1, $k = 2$ and $m.\alpha = 0$ mean that $Witness(m)$ is at most at distance $k - 0 = 2$, now its witness g is at distance 2.

In the latter case, $p (= r)$ may not be k -hop dominated by any process of Dom inside its subtree and, by definition, there is no process outside its subtree, indeed $T(p) = T$, see the root a in Configuration (I) of Figure 1. Thus, p must be placed in Dom .

If p is *tall*, there is at least one process q at $p.\alpha - k$ hops below p such that $q.\alpha = k$. Any such a process q belongs to Dom and k -hop dominates p . Hence, p can select any of them as *witness* (in the algorithm, we break ties using the order on channels).

The path from p to $Witness(p)$ goes through a tall child with minimum α -value. See, for example, in Configuration (II) of Figure 1, $k = 2$ and $a.\alpha = 4$ mean that $Witness(a)$, here c , is $4 - k = 2$ hops below a . In Configuration (I), remark that there are two possible witnesses for c ($c.\alpha = 4$): g and h , both are $c.\alpha - k = 2$ hops below c .

Note that, if $p.\alpha = k$, then $p.\alpha - k = 0$, that is, $p = q = Witness(p)$ and p belongs to Dom .




Figure 1: Two examples of 2-clustering computed by $\mathcal{C}(2)$. We only draw the spanning tree, other edges are omitted. The root of each tree is the rightmost node. α -values are given inside the nodes. Bold circles represent clusterheads. Identifiers are given next to the nodes. $headC$ -values are given in brackets. Arrows represent $parC$ pointers. We colored the different k -clusters.

6.4.2. Constructing the k -Clustering. The second phase of $\mathcal{C}(k)$ partitions the nodes into distinct k -clusters, each of which contains one clusterhead. Each k -cluster is actually built as a k -cluster spanning tree, a tree containing all the nodes of that k -cluster. Each k -cluster spanning tree is a subgraph of T rooted at the clusterhead, possibly with the directions of some edges reversed. Furthermore, the height of the k -cluster spanning tree is at most k .

Each node p of Dom designates itself as clusterhead by setting $p.parC$ and $p.headC$ to \perp and $Id(p)$, respectively (see the second and third actions in Algorithm 1).

Other nodes q designate their parent in the k -cluster, $q.parC$, using second action as follows: (1) if q is *short*, then $q.parC$ should be its parent in the tree T ; (2) if q is *tall*, then q selects $q.parC$ as its *tall* child in the tree of minimum α value; as explained before we use channel order to break ties (see $MinCMinATall(q)$ and $\min_{<C}$ which selects the minimum channel in a set).

Finally, identifiers of clusterheads are propagated in the $headC$ variables top-down in the k -clusters using the third action (see macro $HeadC$).

6.4.3. Examples. Two examples of 2-clustering computed by $\mathcal{C}(2)$ are given in Figure 1. In Subfigure 1.(I), the root is a *short* process, consequently it belongs to Dom . In Subfigure 1.(II), the root is a *tall* process, consequently it does not belong to Dom .

7. TERMINATION OF $\mathcal{C}(k)$

The goal of this section is to build a certified proof of the **termination** property of $\mathcal{C}(k)$. Namely, we show that $\mathcal{C}(k)$ converges, as far as \mathbf{Assume}_{cl} is satisfied.

Algorithm $\mathcal{C}(k)$ is made of three actions which are locally prioritized at each node (see the guard of every action). The first action (which computes α) has priority over the second and last actions. The second action (which computes $parC$) can only be enabled when the first action is disabled, but has priority over the last one. This latter action can be enabled only if the two others are disabled. We use the method given in Section 4.5 to show the termination of $\mathcal{C}(k)$: it consists in proving the termination of the three actions separately (using results from Section 4.4). Then, we use the lexicographic order result from Section 4.5 to conclude.

First, we assume `sym_net` and a root node r . For the definition of safe configuration, we instantiate `safe` as every configuration in which read-only Par-variables satisfy: (`span_tree r Par`). This assumption on the existence of the spanning tree T rooted at r is mandatory, since, as we will see below, the local potentials we use in proofs are based on the tree T . Note that it is easy to prove that `safe` is stable, since it only depends on read-only variables.

7.1. α _SafeStep. For the first action, we build safe steps, called α _SafeStep, by instantiating the predicate `QTrans` (see Section 4.1) as follows: for any two safe configurations (`sg sg': safeEnv`), an α _SafeStep occurs between `sg` and `sg'` if (`Step sg' sg`) holds and if the following condition is satisfied:

$$\text{QTrans}_\alpha \text{ sg' sg} := \exists(p: \text{Node}), \alpha_enabled \text{ sg } p \wedge \text{has_moved } \text{sg}' \text{ sg } p$$

where $\alpha_enabled$ is exactly the guard of the first action and (`has_moved sg' sg p`) means that p has executed its local program during the step (this is done by checking that the state of node p in safe configuration `sg'` is the same as the result of `run` on `sg` and p). We manage to have Boolean versions of the above predicates; this is made possible due to the fact that all nodes in the network are stored in the list `all_nodes`. When ($\alpha_enabled \text{ sg } p$), we say that node p is *α -enabled in the safe configuration `sg`* and when (`has_moved sg sg' p`) holds additionally, we say that p *has α -moved from `sg` to `sg'`*.

In an α _SafeStep, we require that at least a node executes its first action; note that this gives no guarantee on other nodes, which may or may not execute an action. We use then the method explained in Section 4.4 to prove the safe inclusion of α _SafeStep:

Theorem $\alpha_safe_inclusion$:

$$\forall(\text{sg1 sg2: safeEnv}), \alpha_SafeStep \text{ sg2 sg1} \rightarrow (\text{Pot}_\alpha \text{ sg2}) \prec (\text{Pot}_\alpha \text{ sg1}).$$

In other words, when an α _SafeStep occurs, the global potential Pot_α decreases. Pot_α is built from α , as the list of local potentials, α_pot , at every node. In the following, we explain how we compute this α -potential, α_pot , at each node.

7.1.1. α -Potential. We define the *depth* of a node as one plus the distance from the root r to the node in the tree T . For a given safe configuration `sg` and a node p , (`depth sg p`) returns 1 (natural number, type `nat`) if p is the root r and, otherwise, $(1 + (\text{depth } \text{sg } q))$ where q the parent of p in the tree T ; the definition relies on structural induction on (`WF_par p`). We define the α -potential of a node p in a safe configuration `sg`, ($\alpha_pot \text{ sg } p$), as 0 if p is not α -enabled in `sg` and (`depth sg p`), otherwise.

7.1.2. Local Criterion for α _SafeSteps. Let `sg1` and `sg2` be two safe configurations where ($\alpha_safeStep \text{ sg2 sg1}$) holds. Consider a node p whose α -potential has increased during the step, *i.e.*, ($\alpha_pot \text{ sg1 } p$) $<_p$ ($\alpha_pot \text{ sg2 } p$). This means, by definition of α_pot , that p is disabled in `sg1` (its potential is 0) and becomes enabled in `sg2` (its potential becomes (`depth sg2 p`) > 0).

To show the local criterion, we exhibit a down-path in the tree T from p to some leaf, which contains a witness node that is α -enabled in `sg1` and α -disabled in next configuration `sg2`. We prove the result in two steps. First, we exhibit a child of node p , `child`, which necessarily executes the first action of its algorithm during the step. This is proven by

induction on the neighbors of p using the fact that `run` only depends on the states of the children of p in the tree T . Next, we prove the following lemma:

Lemma `moving_node_has_disabled_desc`: $\forall(\text{child: Node}),$
 $\alpha((\text{getEnv sg1}) \text{ child}) \neq \alpha((\text{getEnv sg2}) \text{ child}) \rightarrow$
 $\exists(\text{desc: Node}),$
 $(\exists(\text{path: list Node}), \text{directed_tree_path child path desc}) \wedge$
 $\alpha_enabled \text{ sg1 desc} \wedge \neg \alpha_enabled \text{ sg2 desc}$

where `directed_tree_path` checks that `path` is actually a path from `child` to `desc` in the directed spanning tree T . The lemma states that when the node `child` α -moves, it is down-linked in T to a node which was α -enabled and becomes α -disabled, during the step. Hence, the lemma provides the witness node required to prove the local criterion.

The lemma is proven by induction on (WF_child child) , *i.e.*, on the down-paths from `child` in T . Consider a node in such a path which is enabled in `sg1` and that α -moves during the step from `sg1` to `sg2`. We have two cases.

- Either it becomes disabled in `sg2`: this is the base case of the induction, taking `desc` as `child` and `path` empty.
- Or it is still enabled in `sg2`: for this case, we prove that any node that executed the first action of the algorithm in `sg1` but is still α -enabled in `sg2` has a child in T which has also α -moved (the proof is based on induction on the children of the node). This result provides the induction step of the proof.

7.1.3. Global Criterion for α -SafeSteps. The global criterion requires to find a witness node whose α -potential differs between `sg1` and `sg2`. We show that there exists a node p with α -potential $(\text{depth sg1 } p)$ in `sg1` (such a potential is necessarily greater than 0), and α -potential 0 in `sg2`. Namely, p is α -enabled in `sg1`, but α -disabled in `sg2`. The proof uses the fact that at least one node, say q , has α -moved during the step (see definition of QTrans_α). Then, we use Lemma `moving_node_has_disabled_desc` again to exhibit a witness node p (on a given down-path of T from q) which is α -enabled in `sg1`, but α -disabled in `sg2`.

7.1.4. Conclusion for α -SafeSteps. Local and global criteria being proven, we directly obtain Theorem `α _safe_inclusion` from Section 4.4. For safe steps involving the second action (predicate `parC_SafeStep`) and the third action (predicate `headC_SafeStep`), we use exactly the same method, see in the next sections.

7.2. `parC_SafeStep`. For the second action, we build safe steps, called `parC_SafeStep`, by instantiating the predicate `QTrans` as follows: for any two safe configurations `sg` and `sg'`,

$$\text{QTrans}_{\text{parC}} \text{ sg}' \text{ sg} := \neg \text{QTrans}_\alpha \text{ sg}' \text{ sg} \wedge$$

$$\exists(p: \text{Node}), \text{parC_enabled sg } p \wedge \text{has_moved sg}' \text{ sg } p$$

where `parC_enabled` is exactly the guard of the second action. As before, `QTransparC` is proven decidable and when $(\text{parC_enabled sg } p)$, we say that node p is *parC-enabled in the safe configuration `sg`*. In a `parC_SafeStep`, we require that no node executes its first action (nodes can be α -enabled, but do not move) and at least a node executes its second action (other nodes can execute or not their third action). We then show the theorem

Theorem *parC_safe_inclusion*: $\forall(\text{sg1 } \text{sg2}: \text{safeEnv}),$
 $\text{parC_SafeStep } \text{sg2 } \text{sg1} \rightarrow (\text{Pot}_{\text{parC}} \text{sg2}) \prec (\text{Pot}_{\text{parC}} \text{sg1})$

using exactly the same method as for α .

During any *parC_SafeStep* between safe configurations *sg1* and *sg2*, no node executes its first action, as expressed in $\text{QTrans}_{\text{parC}}$. Hence, all the values of α stay unchanged. As a consequence, for every node *p*, the macro *ParC(p)* from Algorithm 1 (which only depends on the values of α) has the same result when evaluated at both safe configurations *sg1* and *sg2*. This ensures that

- (1) a node which executes its second action during the *parC_SafeStep* from *sg1* to *sg2* is no longer *parC*-enabled in *sg2*;
- (2) a *parC*-disabled node in *sg1* remains *parC*-disabled in *sg2*.

From those observations, we define the *parC*-potential of a node *p* in safe configuration *sg*, (*parC_pot sg p*), as 1 if *p* is *parC*-enabled in *sg* and 0, otherwise. Now, criteria from Section 4.4 are straightforward. Indeed, we have

- the *local criterion*, since the *parC*-potential of a node cannot increase (*i.e.* switch from 0 to 1, see (2) above);
- the *global criterion*, since, due to $\text{QTrans}_{\text{parC}}$, there exists a node which executes its second action: from (1) above, its *parC*-potential differs between *sg1* and *sg2* since it switches from 1 to 0.

7.3. headC_SafeStep. For the third action, we build safe steps, called *headC_SafeStep*, by instantiating the predicate QTrans as follows: for any two safe configurations *sg* and *sg'*,

$$\text{QTrans}_{\text{headC}} \text{sg}' \text{sg} := \neg \text{QTrans}_{\alpha} \text{sg}' \text{sg} \wedge \neg \text{QTrans}_{\text{parC}}$$

As before, $\text{QTrans}_{\text{headC}}$ is proven decidable. We also denote by *headC_enabled* the guard of the third action and when (*headC_enabled sg p*), we say that node *p* is *headC-enabled in the safe configuration sg*. In a *headC_SafeStep* between safe configurations, we require that no node executes its first or second action (nodes can be α -enabled or *parC*-enabled, but do not move in that case); furthermore, as enforced by the predicate *Step*, at least one node executes: therefore it executes its third action. We then show the theorem:

Theorem *headC_safe_inclusion*: $\forall(\text{sg1 } \text{sg2}: \text{safeEnv}),$
 $\text{headC_SafeStep } \text{sg2 } \text{sg1} \rightarrow (\text{Pot}_{\text{headC}} \text{sg2}) \prec (\text{Pot}_{\text{headC}} \text{sg1})$

using exactly the same method as for α .

7.3.1. parC-path. The fact that third actions terminate is due to the fact that *headC* values are computed along the paths made of *parC* pointers. We call those paths *parC-paths* and we define them using the relation *pcl_rel sg*, where *sg* is a safe configuration: two nodes *p* and *q* are related *via* (*pcl_rel sg*) (*pcl_rel sg q p*) when *p* is α -disabled and *parC*-disabled in *sg*, and when its *parC*-pointer points to *q* in *sg*. A *parC*-path between two nodes *P* and *Q*, if exists, is then the list of nodes, from *P* to *Q*, built using the transitive closure of (*pcl_rel sg*).

We show that the relation (*pcl_rel sg*) is well-founded, for any safe configuration *sg*. First, we observe from the algorithm that when two nodes *p* and *q* are related, *i.e.* (*pcl_rel sg q p*), either *p* is short and *q* is its parent in *T*, or *p* is tall and *q* is a child of *p* in *T*. Then, we split the proof into two parts. *For tall nodes*, we prove that for any

two related nodes p and q , such that $(\text{pcl_rel } \text{sg } q \ p)$, we have: q is tall whenever p is. Hence, for a given node p , we can prove $(\text{Acc } (\text{pcl_rel } \text{sg}) \ p)$ directly from induction on $(\text{WF_child } p)$. For *short nodes*, a short node can be linked using $(\text{pcl_rel } \text{sg})$ to a short (like m in Configuration (I) of Figure 1) or a tall node (like i in Configuration (I) of Figure 1). To prove $(\text{Acc } (\text{pcl_rel } \text{sg}) \ p)$ for a given node p , we proceed again by induction, following the *parC*-path of p , on $(\text{WF_par } p)$. But it may occur that the path reaches a tall node, in which case, we use the previous result for tall nodes to be able to conclude the induction case.

From the well-foundedness of $(\text{pcl_rel } \text{sg})$ for any safe configuration sg , we can inductively define $(\text{dist_hd } \text{sg } p)$ as the length of the *parC*-path of p in sg . Since *parC*-pointers are constant in any *headC_SafeStep*, dist_hd has the same value in any two safe configurations sg1 and sg2 as far as they are linked by a *headC_SafeStep*, *i.e.*, $(\text{headC_SafeStep } \text{sg2 } \text{sg1})$ holds. From this result, we can show that

- (1) If a node p is *headC*-enabled in sg1 , executes, and remains *headC*-enabled in sg2 , then this means that p has a successor in its *parC*-path and this successor is also *headC*-enabled in sg1 and executes during the step.
- (2) As the *parC*-path is finite, this proves (by a structural induction on $(\text{Acc } (\text{pcl_rel } \text{sg1}) \ p)$) that p is necessarily linked in its *parC*-path to a node q , which is *headC*-enabled in sg1 , executes, and becomes *headC*-disabled in sg2 ; furthermore we have that

$$\text{dist_hd } \text{sg1 } q < \text{dist_hd } \text{sg1 } p$$

- (3) Therefore, from (1) and (2), a node P , which is *headC*-enabled in sg1 and executes during the step, is necessarily linked *via* its *parC*-path to a node Q such that Q is *headC*-enabled in sg1 , *headC*-disabled in sg2 , and

$$\text{dist_hd } \text{sg1 } Q \leq \text{dist_hd } \text{sg1 } P$$

This node can be P itself if P becomes *headC*-disabled in sg2 , or a node which is further along in the *parC*-path.

7.3.2. *headC*-Potential. We could have used dist_hd to build the *headC*-potential but results from Section 4 assume decreasing potential, whereas this one would have been increasing. Instead, we prove the existence of a natural number NN such that proposition

$$\text{HNN} := \forall(\text{sg}: \text{safeEnv}) (p: \text{Node}), \text{NN} > \text{dist_hd } \text{sg } p$$

is true. We use the tools about quantitative properties (see Section 5) to achieve the proof. We set NN as $(n + 1)$, where n is the number of nodes in the network, using the list `all_nodes` to show the existence of n . We prove HNN using the fact that in a *parC*-path, each node occurs at most once; this comes from the well-foundedness of $(\text{pcl_rel } \text{sg})$ from which we can infer that a *parC*-path contains no loop.

Finally, for a given safe configuration sg and a given node p , we pick its *headC*-potential to be $(\text{NN} - \text{dist_hd } \text{sg } p)$ if p is *headC*-enabled in sg and 0 otherwise.

7.3.3. Global Criterion for $headC_SafeSteps$. The global criterion requires to exhibit a node whose potential has changed from $sg1$ to $sg2$: we look for a node which is $headC$ -enabled in $sg1$ (potential is > 0) and $headC$ -disabled in $sg2$ (potential is 0). From **Step**, there exists a node p which executes during the step and $QTrans_{headC}$ guarantees that it uses its third action. We directly use the result (3) above: there exists a node q in the $parC$ -path of p which is $headC$ -enabled in $sg1$ and $headC$ -disabled in $sg2$, hence its $headC$ -potential changes during the step.

7.3.4. Local Criterion for $headC_SafeSteps$. To show the local criterion, we consider a node p whose $headC$ -potential increases during the step. Specifically, a node which is $headC$ -disabled in $sg1$ and $headC$ -enabled in $sg2$. We prove that this situation is possible only if p has a successor p' in its $parC$ -path such that p' $headC$ -executes during the step. Note that $(dist_hd\ sg1\ p)$ is greater than $(dist_hd\ sg1\ p')$. We use the result (3) above, which ensures the existence of a node q such that

- q is $headC$ -enabled in $sg1$ and $headC$ -disabled in $sg2$ (hence its $headC$ -potential changes during the step), and
- $(dist_hd\ sg1\ q \leq dist_hd\ sg1\ p' = (dist_hd\ sg1\ p) - 1)$. Hence,
 $(pot_hd\ sg2\ p = NN - dist_hd\ sg1\ p < pot_hd\ sg1\ q = NN - dist_hd\ sg1\ q)$.

7.4. Termination. We proved the safe inclusions for the three kinds of safe steps, so we can apply the lexicographical order method with three dimensions. It requires to show that assumptions about priorities, as encoded by the lexicographical order, conform to the algorithm and to the definitions of α -steps, $parC$ -steps, and $headC$ -steps; this is translated into the assumption $Hdisjoint$ instantiated at two levels, namely, we need to verify that

$$\begin{aligned} Hdisjoint_{cl} &:= \forall (sg' sg: safeEnv), \\ &\quad parC_SafeStep\ sg' sg \vee headC_SafeStep\ sg' sg \rightarrow \\ &\quad (Pot_{\alpha}\ sg') \approx (Pot_{\alpha}\ sg) \\ &\quad \wedge \\ &\quad headC_SafeStep\ sg' sg \rightarrow (Pot_{parC}\ sg') \approx (Pot_{parC}\ sg) \end{aligned}$$

(when a $parC$ -step or a $headC$ -step occurs, no α -potential changes and when a $headC$ -step occurs, no $parC$ -potential changes); the validity of the condition comes directly from the values of $QTrans_{parC}$ and $QTrans_{headC}$.

From $Hdisjoint_{cl}$, $\alpha_safe_inclusion$, $parC_safe_inclusion$ and $headC_safe_inclusion$, we apply Lemma `union_lex_wf3` and obtain that

$$well_founded\ (\alpha_SafeStep \cup parC_SafeStep \cup headC_SafeStep).$$

Finally, we proved the equivalence between the relations `SafeStep` and $(\alpha_SafeStep \cup parC_SafeStep \cup headC_SafeStep)$: this directly comes from the values of the predicates $QTrans_{\alpha}$, $QTrans_{parC}$ and $QTrans_{headC}$. This ends the proof and concludes that $(well_founded\ SafeStep)$. Using Lemma `Acc_Algo_Multiset`, we obtain the desired property:

$$\mathbf{Theorem}\ \mathcal{C}(k)_termination: \forall (g: Env),\ Assume_{cl}\ g \rightarrow Acc\ Step\ g.$$

8. PARTIAL CORRECTNESS OF $\mathcal{C}(k)$

We develop a certified proof of the `P_correctness` property of $\mathcal{C}(k)$. Namely, we show the partial correctness of $\mathcal{C}(k)$, as far as `Assumecl` is satisfied:

Theorem `C(k)_at_terminal`: $\forall (g: \text{Env}), \text{Assume}_{cl} g \rightarrow \text{terminal } g \rightarrow \text{P}_{cl} g$.

This goal is divided into three subgoals. First, we prove the partial correctness of the first actions; this is achieved by proving that once the algorithm has converged, the α -values allow to define a k -hop dominating set. Second, we prove, that after termination, the strong k -clustering specification holds (*i.e.*, each node knows the identifier of its clusterhead and the channel corresponding to its parent link in its k -cluster tree). Third, we show that any terminal configuration contains at most $\lfloor \frac{n-1}{k+1} \rfloor + 1$ k -clusters.

8.1. Proof for a k -hop Dominating Set.

8.1.1. **Values of α are in range $\{0, \dots, 2k\}$.** As a preliminary result, the value of α at a node p is in range $\{0, \dots, 2k\}$ after p participates in any step and also when the system is in a terminal configuration. The proof shows that the value returned by macro `Alpha(p)` is in range $\{0, \dots, 2k\}$: this is proven using a case analysis on $\text{MaxAShort}(p) + \text{MinATall}(p) > 2k - 2$ and the fact that, by definition, $-1 \leq \text{MaxAShort}(p) \leq k - 1$ and $k \leq \text{MinATall}(p) \leq 2k + 1$.

8.1.2. **Proof for `Dom`.** We prove that the set `Dom`, made of all the nodes p such that `kDominator(p)` holds, is a k -hop dominating set in any terminal configuration, namely we need to check the existence of a path in G between any node p and any node `kdom` of `Dom`, such that this path is of length at most k . To be usable by the rest of the proof, we show a bit more. Actually, once the α -values allow to define a k -hop dominating set, they also exhibit routing paths between each node and one of its witnesses in `Dom` by choosing one of the possible path: this choice is made using the ordering on channels and the $\min_{<_C}$ operator. Therefore, in this part, we prove that *any* such possible routing path has length at most k .

Tree Paths. To achieve this property, the algorithm builds tree paths of particular shape: those paths use edges of T in both directions. Precisely, these edges are defined using relation `(is_kDom_edge g)`, in a given configuration g , which depends on α -values: for any short node s , we select the edge from p to s , where p is the parent of s in T (`Par`), *i.e.*, `(is_kDom_edge g p s)` holds; while for any tall node t which is not in `Dom`, we select every edge from c to t , where c is a child of t in T such that $(\alpha(g\ c) = \alpha(g\ t) - 1)$, *i.e.*, `(is_kDom_edge g c t)` holds. The relation `(is_kDom_edge g)` defines a subgraph of G called the `kdom`-graph of g . (Remark that all directed edges in graphs of Figure 1 appear in their associated `kdom`-graph, yet in opposite sense.)

The rest of the analysis is conducted assuming a terminal configuration $(g: \text{Env})$ which contains a rooted spanning tree built upon a bidirectional graph, namely such that `(Assumecl g)` and `(terminal g)` hold. We aim at proving the following result:

```

OK_dom g p :=
  (∃(kdom: Node), (kDominator g kdom) ∧
    ∃(path: list Node), is_kDom_path g path kdom p)
  ∧

```

$$\forall(\text{kdom} : \text{Node}) (\text{path} : \text{list Node}),$$

$$\text{is_kDom_path } g \text{ path kdom } p \rightarrow (\text{length path}) \leq k.$$

Theorem `kDom_correctness`: $\forall(p : \text{Node}), \text{OK_dom } g \text{ } p.$

where `is_kDom_path` checks that its parameter `path` is a path in the `kdom`-graph of configuration `g` between `kdom` and `p`. The proof of $(\text{OK_dom } g \text{ } p)$ for any node `p` is split into two cases, depending on whether `p` is tall or short. Actually, we prove by straightforward induction on `i` that, for any node `p` and any natural `i`, such that $(\alpha (g \text{ } p) = k + i)$ (resp. $(\alpha (g \text{ } p) = k - i)$), the property $(\text{OK_dom } g \text{ } p)$ holds – where the length of `path` is at most `i`.

Proof for Tall Nodes. For case $(i = 0)$, `p` satisfies $k\text{Dominator}(p)$ and any path from `p` to `p` in the `kdom`-graph of `g` has length 0. For case $(i = j + 1)$, as $(\alpha (g \text{ } p) = k + i)$ is positive, we can prove using a case analysis on $\text{MaxAShort}(p) + \text{MinATall}(p) \leq 2k - 2$, that there exists a child `q` of `p` with $(\alpha (g \text{ } q) = k + j)$ on which we can apply the induction hypothesis. This exhibits a path in the `kdom`-graph of `g` from some k -hop dominator `kdom` to `q`. Since `p` is the parent of `q` in T (i.e., `q` is a child of `p`), we obtain a path from `kdom` to `p` at `g`.

Looking at the second part of the result, we have to prove that any path in the `kdom`-graph of `g` has length at most `i`. Then, either `path` is empty and its length is 0, or we can decompose it into a `kdom`-path `path'` followed by some node `q` and then `p`, such that there is a `kdom`-edge between `q` and `p`. Using the definition of `kdom_edge`, we obtain that $(\alpha (g \text{ } q) = \alpha (g \text{ } p) - 1 = k + j)$: hence we apply again the induction hypothesis to `q` and obtain that `path'` has length at most `j`; hence `path` has length at most $(j + 1)$.

Proof for Short Nodes. The case $(i = 0)$ is already proven by the above result for tall nodes. We now look at case $(i = j + 1)$. When `p` is the root of T , then $k\text{Dominator}(p)$ holds and any path in the `kdom`-graph of `g` whose terminal extremity is `p` is empty. We now assume that `p` is non-root.

For the first part of the property (looking for a witness `kdom` $\in \text{Dom}$ and a path from `kdom` to `p`), we pick the parent `q` of `p` in T . We can show that $(\alpha (g \text{ } p) \leq \alpha (g \text{ } q) + 1)$ and that $(\text{is_kDom_edge } g \text{ } q \text{ } p)$. If `q` is also short, the induction hypothesis applies directly, otherwise (`q` is tall), the above property $(\text{OK_dom } g \text{ } q)$ holds. In both cases, this provides a witness node `kdom` in Dom and a path from `kdom` to `q` in the `kdom`-graph of `g`; we add to `path` the `kdom`-edge from `q` to `p` to build a path from `kdom` to `p` in `kdom`-graph.

For the second part of the property, we consider a `kdom`-path `path` from some node to `p`. Either this path is empty, in which case, the property trivially holds, or we can decompose it into a sub-path `path'` and an edge from some node `q` to `p` in `kdom`-graph. From the definition of `kdom_edge`, `q` is the parent of `p` in T . If `q` is short, we apply the induction hypothesis to obtain that the length of `path'` is at most `j`. Otherwise `q` is tall and we have two cases:

- If $\text{MaxAShort}(q) + \text{MinATall}(q) > 2k - 2$, then we have

$$(\alpha (g \text{ } q) = \text{MaxAShort}(q) + 1 \leq k).$$

The fact that `q` is tall implies $(\alpha (g \text{ } q) = k)$. Hence, the `path'` has length 1.

- Otherwise, since `q` is tall, from the result $(\text{OK_dom } g \text{ } q)$ above, `path'` has length at most $(\alpha (g \text{ } q) - k)$. Now, since `p` is a short child of `q` in T , we have that

$$(\alpha (g \text{ } p) \leq \text{MaxAShort}(q)).$$

We also have $(\alpha (g \ q) = \text{MinATall}(q) + 1)$. Combining all, we obtain that

$$(\alpha (g \ q) - k \leq k - \text{MaxAShort}(q) - 1 \leq j).$$

Hence, path' has a length at most j and path at most $(j + 1)$.

8.2. Proof for k -Clustering. We explain here the proofs of the parts `kCluster_strong` and `partition_OK` of the specification \mathbb{P}_{cl} . We instantiate

- `clusterHeadID` as `headC`,
- `clusterParent` is set to `parC` and
- `(kCluster g h p)` is defined by `(eqId Id(h) (clusterHeadID (g p)))` for any configuration `(g: Env)` and nodes `h p: Node`.

For the rest of the analysis, we fix a configuration `(g: Env)` such that `(Assumecl g)` and `(terminal g)` hold. As a preliminary remark, we prove that every node in a cluster path declares the same clusterhead in `g`:

Lemma same_hd: $\forall (b \ e: \text{Node}) (\text{path}: \text{list Node}),$
 $\text{cluster_path } g \ b \ \text{path } e \rightarrow \text{eqId } (\text{headC } (g \ b)) (\text{headC } (g \ e)).$

The proof is an easy induction on the list `path` using the expression of macro `ParC`. This lemma ensures that the predicates `(cluster_path g)` and `(agreed_cluster_path g)` are equivalent (see Section 6.2).

8.2.1. Relation `is_cluster_parent`. We first study the relation `(is_cluster_parent g)`. We show that it is included into the `kDom`-edges of the `kDom`-graph:

$\text{inclusion } (\text{is_cluster_parent } g) (\text{is_kDom_edge } g)$

The proof is a (quite long) case analysis based on the fact that `g` is terminal and on the macros `Alpha` and `ParC`. We also show that this relation is well-founded:

$\text{well_founded } (\text{is_cluster_parent } g)$

Actually, we proved it during the termination proof (see the part about `headC_SafeSteps` in Section 7.3). Precisely, we showed that `(pcl_re1 g')` is well-founded for any configuration `g'`. Now, since `g` is terminal, `(is_cluster_parent g)` is included in `(pcl_re1 g)`, and we obtain the well-foundedness.

8.2.2. `kDominator` and `clusterHead`. We manage to prove the equivalence between `(kDominator g)` and `(clusterHead g)`:

$\forall (p: \text{Node}), \text{kDominator } g \ p \leftrightarrow \text{clusterHead } g \ p.$

First, we transform this goal into

$\forall (p: \text{Node}), \text{eqoptionA eqC } (\text{parC } (g \ p)) \perp \leftrightarrow \text{clusterHead } g \ p.$

since we can prove that `(kDominator g p \leftrightarrow eqoptionA eqC (parC (g p)) \perp)`. Indeed, this latter result is based on the fact the `(parC (g p))` is equal to `ParC(p)` in the terminal configuration `g`; the proof uses a case analysis which treats the case when `p` is short easily. For the case when `p` is a non-root tall node, it requires to prove that `MinCMinATall` actually returns a tall child of `p` which is a quite tricky intermediate result, based on the definition of `MinCMinATall`.

Back to the goal above, the direct part of the equivalence comes directly from the fact that $(parC (g p))$ is \perp , using the expression of $ParC$ and the fact that g is terminal. Now, we focus on the reverse part of the equivalence and we assume that $Id(p)$ and $(headC (g p))$ are equal.

As $(is_cluster_parent g)$ is a well-founded and decidable relation, and as there exists a finite number of nodes in the network, for any node p , we can build the maximal cluster path called $(path: list Node)$ from p : it reaches some node called $(h: Node)$ which has no cluster parent pointer; hence $path$ and h satisfy:

$$cluster_path h path p \wedge \forall(x: Node), \neg cluster_parent x h.$$

When $path$ is empty, proof is done, since p and h are the same node with no cluster parent pointer. Otherwise $path$ is not empty: we show that this case is not possible since it yields a contradiction. Indeed, nodes in a non-empty cluster path are all different (we prove that $path$ contains no loop since $(is_cluster_parent g)$ is a well-founded relation over a finite set), hence have all different identifiers (since $uniqueId$ is assumed). This ensures that $Id(h)$ and $Id(p)$ are different. From Lemma `same_hd`, we obtain that $(headC (g h))$ and $(headC (g p))$ are equal and since $(parC (g h))$ is \perp , the expression of $ParC$ ensures that $(headC (g h))$ is $Id(h)$. Those three equations yield a contradiction with the fact that $Id(p)$ and $(headC (g p))$ are equal, as assumed at the beginning of the proof.

8.2.3. Proof of `kCluster_strong`. The first line of `kCluster_strong` (see page 19, for the definition) is exactly given by the instantiation of `kCluster`. For the second line, assuming $(kCluster g b e)$, for any two nodes b and e , we define a function that builds the maximal cluster path $path$ from e to b and we prove that b is a clusterhead in g : $(clusterHead g b)$. As $(cluster_path g)$ is included into $(kdom_path g)$, we are able to use the theorem which asserts that $(OK_dom g e)$ and prove that any cluster path has length at most k .

8.2.4. Proof of `partition_OK`. (see page 18 for the definition) Let p be a node. As before, we can build the maximal cluster path, called $path$ from p to its clusterhead h . From Lemma `same_hd`, p declares h as its clusterhead and so does h . Hence, $(kCluster g h p)$ holds. Now to prove uniqueness of clusterheads, we assume $(kCluster g h p)$ and $(kCluster g h' p)$ and the goal is to prove that nodes h and h' are equal. This goal is transformed into $(eqId Id(h) Id(h'))$ using the uniqueness assumption `uniqueId`. As the assumptions on `kCluster` both expands into $(eqId Id(b) (headC (g n)))$ and $(eqId Id(b') (headC (g n)))$ we are done by transitivity.

8.3. Proof for counting. In this section, we formally prove the property `count_OK` which states that $(n - 1) \geq (k + 1)(|CH| - 1)$ where `CH` has been defined as the set of clusterheads. Intuitively, this means that all but one element of `CH` have been chosen as clusterheads by at least $k + 1$ distinct nodes each. Actually, we prove that Dom has this property: $(n - 1) \geq (k + 1)(|Dom| - 1)$ and then use the equivalence between `clusterHead` and `kDominator` to conclude.

The proof outline is the following. First, we assume a terminal configuration $(g: Env)$, $(terminal g)$, such that $(Assume_{cl} g)$ holds. The existence of the natural number n (number of nodes) is given using the results in Section 5.6 about the number of elements in the list `all_nodes`. Similarly, the existence of the natural number $|Dom|$ (number of nodes

in Dom) is given using the above results applied to list `all_nodes` and predicate function (`fun p: Node => (kDominator (g p) = true)`).

We define as *regular head* each node h such that α equals k in g

`RegHead h := (α (g h) = k)`

and the set of regular heads as `RegHeads := { h: Node | RegHead h }`. Note that by definition, `RegHeads` is included in Dom . Again, we prove the existence of the natural number rh which represents the number of nodes in `RegHeads` using list `all_nodes` and predicate `RegHead`.

We also define a *regular node* as a node which declares a regular head as clusterhead. In practice, regular nodes are tall or have a tall ancestor in T . In case the root is a short clusterhead, they cannot be in its cluster. We define predicate `HasTallAncestor` as an inductive predicate which selects any node which has a tall ancestor in T (*i.e.* such that there is a directed path in T from p to the root r that contains a node with α at least k). The set of regular nodes is defined by: `RegNodes := { p: Node | HasTallAncestor p }`. Again, we prove the existence of the natural number rn which is the number of nodes in `RegNodes`. Now, we prove the following theorem:

Theorem `simple_counting`: $rn \geq (k + 1)rh$.

Using results from the library on cardinality of sets and lists, this theorem is reduced to

`Card Smaller (Mk+1 × RegHeads) RegNodes`.

This latter proposition is proven by constructing a relation `Rcount` from pairs of natural numbers $i \in \{0, \dots, k\}$ and *regular heads* to regular nodes, such that: for a regular head h , some $i \in \{0, \dots, k\}$ and a regular node p_i , (`Rcount (i, h) pi`) holds if and only if $(\alpha (g p_i) = i)$ and p_i designates h as clusterhead (*i.e.*, there is a maximal path from p_i to h in the `kdom`-graph of g). We show that `Rcount` is actually an injection of domain $(\mathcal{M}_{k+1} \times \text{RegHeads})$. Indeed, for any pair (i, h) , there is a node p_i such that $(\alpha (g p_i) = i)$ which designates h as clusterhead; the proof is carried out by induction on values of i . Intuitively, this implies that there is a path of length $k + 1$ in the `kdom`-graph of g linking p_0 to h . We then group each regular head with the regular nodes that designate it as clusterhead: each contains at least $k + 1$ regular nodes, *i.e.*, $rn \geq (k + 1)rh$.

Now, we have two cases. If the root is tall, with $(\alpha (g r) \geq k)$, every node in Dom is a regular head, *i.e.* is in `RegHeads` and every node is regular, in `RegNodes`. Otherwise, the root is short and every clusterhead is a regular head except the root and at least one node is not regular, namely the root. These two cases yield the following lemma:

Lemma `split_counting_cases`:

$|Dom| = rh \wedge n = rn \vee |Dom| = 1 + rh \wedge n \geq 1 + rn$.

The proof of the lemma first uses the results on cardinalities, in particular disjoint union between the singleton containing the root and the set of regular nodes (*resp.* regular heads)) and then the above case analysis. The main theorem that proves `count_OK` is then just a case analysis from this lemma.

9. CONCLUSION

We proposed a general framework to build certified proofs of self-stabilizing algorithms. To achieve our goals, we developed general tools about potential functions, which are commonly used in termination proofs of self-stabilizing algorithms. We also proposed a library dealing

with cardinality of sets. We apply our framework to prove that an existing algorithm is silent self-stabilizing for its specification and we show a quantitative property on the output of this case study.

In future work, we expect to certify more complex self-stabilizing algorithms. Such algorithms are usually designed by composing more basic blocks. In this line of thought, we envision to certify general theorems related to classic composition techniques such as collateral or fair compositions.

Finally, we expect to use our experience on quantitative properties to tackle the certification of time complexity of stabilizing algorithms, *aka.* the stabilization time.

REFERENCES

- [1] Karine Altisen, Corbineau Pierre, and Stéphane Devismes. A framework for certified self-stabilization. In *FORTE'2016, the 36th IFIP International Conference on Formal Techniques for Distributed Objects, Components and System*, volume 9688 of *Lecture Notes in Computer Science*, pages 36–51. Springer, 2016.
- [2] Cédric Auger, Zohir Bouzid, Pierre Courtieu, Sébastien Tixeuil, and Xavier Urbain. Certified impossibility results for byzantine-tolerant mobile robots. In Teruo Higashino, Yoshiaki Katayama, Toshimitsu Masuzawa, Maria Potop-Butucaru, and Masafumi Yamashita, editors, *Stabilization, Safety, and Security of Distributed Systems - 15th International Symposium, SSS 2013, Osaka, Japan, November 13-16, 2013. Proceedings*, volume 8255 of *Lecture Notes in Computer Science*, pages 178–190. Springer, 2013.
- [3] José Bacelar Almeida, Manuel Barbosa, Endre Bangerter, Gilles Barthe, Stephan Krenn, and Santiago Zanella Béguelin. Full proof cryptography: Verifiable compilation of efficient zero-knowledge protocols. In *ACM Conference on Computer and Communications Security*, pages 488–500, 2012.
- [4] Jalel Ben-Othman, Karim Bessaoud, Alain Bui, and Laurence Pilard. Self-stabilizing algorithm for efficient topology control in wireless sensor networks. *Journal of Computational Science*, 4(4):199 – 208, 2013.
- [5] Frédéric Blanqui and Adam Koprowski. Color: a coq library on well-founded rewrite relations and its application to the automated verification of termination certificates. *Mathematical Structures in Computer Science*, 21(4):827–859, 2011.
- [6] Roderick Bloem, Nicolas Braud-Santoni, and Swen Jacobs. Synthesis of self-stabilising and byzantine-resilient distributed systems. In *Computer Aided Verification - 28th International Conference, CAV 2016*, 2016.
- [7] Alain Bui, Ajoy Kumar Datta, Franck Petit, and Vincent Villain. Optimal PIF in tree networks. In Yuri Breitbart, Sajal K. Das, Nicola Santoro, and Peter Widmayer, editors, *Distributed Data & Structures 2, Records of the 2nd International Meeting (WDAS 1999), Princeton, USA, May 10-11, 1999*, volume 6 of *Proceedings in Informatics*, pages 1–16. Carleton Scientific, 1999.
- [8] Eddy Caron, Florent Chuffart, and Cdric Tedeschi. When self-stabilization meets real platforms: An experimental study of a peer-to-peer service discovery system. *Future Generation Computer Systems*, 29(6):1533 – 1543, 2013.
- [9] Eddy Caron, Ajoy Kumar Datta, Benjamin Depardon, and Lawrence L. Larmore. A self-stabilizing k-clustering algorithm for weighted graphs. *J. Parallel Distrib. Comput.*, 70(11):1159–1173, 2010.
- [10] Eddy Caron, Frédéric Desprez, Franck Petit, and Cédric Tedeschi. Snap-stabilizing prefix tree for peer-to-peer systems. *Parallel Processing Letters*, 20(1):15–30, 2010.
- [11] Pierre Castéran, Vincent Filou, and Mohamed Mosbah. Certifying distributed algorithms by embedding local computation systems in the coq proof assistant. In *Symbolic Computation in Software Science (SCSS'09)*, 2009.
- [12] Meixian Chen and Jean-François Monin. Formal Verification of Netlog Protocols. In Tiziana Margaria, Zongyan Qiu, and Hongli Yang, editors, *Sixth International Symposium on Theoretical Aspects of Software Engineering, TASE 2012, 4-6 July 2012, Beijing, China*, pages 43–50. IEEE, 2012.
- [13] Nian-Shing Chen, Hwey-Pyng Yu, and Shing-Tsaan Huang. A self-stabilizing algorithm for constructing spanning trees. *Inf. Process. Lett.*, 39(3):147–151, 1991.

- [14] Pierre Corbineau, Mathilde Duclos, and Yassine Lakhnech. Certified security proofs of cryptographic protocols in the computational model: An application to intrusion resilience. In Jean-Pierre Jouannaud and Zhong Shao, editors, *Certified Programs and Proofs - First International Conference, CPP 2011, Kenting, Taiwan, December 7-9, 2011. Proceedings*, volume 7086 of *Lecture Notes in Computer Science*, pages 378–393. Springer, 2011.
- [15] Pierre Courtieu. Proving self-stabilization with a proof assistant. In *16th International Parallel and Distributed Processing Symposium (IPDPS 2002), 15-19 April 2002, Fort Lauderdale, FL, USA, CD-ROM/Abstracts Proceedings*. IEEE Computer Society, 2002.
- [16] Pierre Courtieu, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. Impossibility of gathering, a certification. *Inf. Process. Lett.*, 115(3):447–452, 2015.
- [17] Pierre Courtieu, Lionel Rieg, Sébastien Tixeuil, and Xavier Urbain. Certified universal gathering in R^2 for oblivious mobile robots. In Cyril Gavoille and David Ilcinkas, editors, *Distributed Computing - 30th International Symposium, DISC 2016, Paris, France, September 27-29, 2016. Proceedings*, volume 9888 of *Lecture Notes in Computer Science*, pages 187–200. Springer, 2016.
- [18] Jean-Michel Couvreur, Nissim Francez, and Mohamed G. Gouda. Asynchronous unison (extended abstract). In *Proceedings of the 12th International Conference on Distributed Computing Systems, Yokohama, Japan, June 9-12, 1992*, pages 486–493. IEEE Computer Society, 1992.
- [19] Ajoy Kumar Datta, Stéphane Devismes, Karel Heurtefeux, Lawrence L. Larmore, and Yvan Rivierre. Competitive self-stabilizing k-clustering. *Theor. Comput. Sci.*, 626:110–133, 2016.
- [20] Ajoy Kumar Datta, Lawrence L. Larmore, Stéphane Devismes, Karel Heurtefeux, and Yvan Rivierre. Self-stabilizing small k-dominating sets. *IJNC*, 3(1):116–136, 2013.
- [21] Yuxin Deng and Jean-François Monin. Verifying self-stabilizing population protocols with coq. In Wei-Ngan Chin and Shengchao Qin, editors, *TASE 2009, Third IEEE International Symposium on Theoretical Aspects of Software Engineering, 29-31 July 2009, Tianjin, China*, pages 201–208. IEEE Computer Society, 2009.
- [22] Nachum Dershowitz and Zohar Manna. Proving termination with multiset orderings. *Commun. ACM*, 22(8):465–476, 1979.
- [23] Stéphane Devismes, Anissa Lamani, Franck Petit, Pascal Raymond, and Sébastien Tixeuil. Optimal grid exploration by asynchronous oblivious robots. In *Stabilization, Safety, and Security of Distributed Systems - 14th International Symposium, SSS 2012, Toronto, Canada, October 1-4, 2012. Proceedings*, volume 7596 of *Lecture Notes in Computer Science*, pages 64–76. Springer, 2012.
- [24] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.
- [25] Shlomi Dolev. Self-stabilizing routing and related protocols. *Journal of Parallel and Distributed Computing*, 42(2):122 – 127, 1997.
- [26] Shlomi Dolev, Mohamed G. Gouda, and Marco Schneider. Memory requirements for silent stabilization. *Acta Inf.*, 36(6):447–462, 1999.
- [27] Fathiyeh Faghih, Borzoo Bonakdarpour, Sébastien Tixeuil, and Sandeep S. Kulkarni. Specification-based synthesis of distributed self-stabilizing protocols. In *Formal Techniques for Distributed Objects, Components, and Systems - 36th IFIP WG 6.1 International Conference, FORTE 2016*, 2016.
- [28] Alexis Fouilhé, David Monniaux, and Michaël Périn. Efficient generation of correctness certificates for the abstract domain of polyhedra. In *Static Analysis Symposium (SAS)*, volume 7935 of *Lecture Notes in Computer Science*, pages 345–365. Springer, 2013.
- [29] Sukumar Ghosh. An alternative solution to a problem on self-stabilization. *ACM Trans. Program. Lang. Syst.*, 15(4):735–742, 1993.
- [30] Georges Gonthier, Andrea Asperti, Jeremy Avigad, Yves Bertot, Cyril Cohen, François Garillot, Stéphane Le Roux, Assia Mahboubi, Russell O’Connor, Sidi Ould Biha, Ioana Pasca, Laurence Rideau, Alexey Solovyev, Enrico Tassi, and Laurent Théry. A machine-checked proof of the odd order theorem. In Sandrine Blazy, Christine Paulin-Mohring, and David Pichardie, editors, *Interactive Theorem Proving - 4th International Conference, ITP 2013, Rennes, France, July 22-26, 2013. Proceedings*, volume 7998 of *Lecture Notes in Computer Science*, pages 163–179. Springer, 2013.
- [31] Shing-Tsaan Huang and Nian-Shing Chen. Self-stabilizing depth-first token circulation on networks. *Distributed Computing*, 7(1):61–66, 1993.

- [32] Philipp Küfner, Uwe Nestmann, and Christina Rickmann. Formal verification of distributed algorithms. In JosC.M. Baeten, Tom Ball, and FrankS. de Boer, editors, *Theoretical Computer Science*, volume 7604 of *Lecture Notes in Computer Science*, pages 209–224. Springer Berlin Heidelberg, 2012.
- [33] Sandeep S. Kulkarni, John M. Rushby, and Natarajan Shankar. A case-study in component-based mechanical verification of fault-tolerant programs. In Anish Arora, editor, *1999 ICDCS Workshop on Self-stabilizing Systems, Austin, Texas, June 5, 1999, Proceedings*, pages 33–40. IEEE Computer Society, 1999.
- [34] Leslie Lamport. *Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2002.
- [35] Leslie Lamport. How to write a 21st century proof. *Journal of Fixed Point Theory and Applications*, 11(1):43–63, 2012.
- [36] Xavier Leroy. A Formally Verified Compiler Back-End. *Journal of Automated Reasoning*, 43(4):363–446, 2009.
- [37] John McCarthy and James Painter. Correctness of a Compiler for Arithmetic Expressions. In *Applied Mathematica*, volume 19 of *Mathematical Aspects of Computer Science*, pages 33–41, 1967.
- [38] Gerry Siegemund, Volker Turau, Christoph Weyer, Stefan Lobs, and Jrg Nolte. Brief announcement: Agile and stable neighborhood protocol for wsns. In *Proceedings of the 15th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS'13)*, pages 376–378, November 2013.
- [39] The Coq Development Team. *The Coq Proof Assistant Documentation*, June 2012.