Kalpa Publications in Computing

Volume 1, 2017, Pages 73–85

LPAR-21S: IWIL Workshop
and LPAR Short Presentations

Decidability of Fair Termination of Gossip Protocols

Krzysztof R. Apt[1,2] and Dominik Wojtczak[3]

[1] CWI, Amsterdam, the Netherlands
[2] University of Warsaw, Poland
[3] University of Liverpool, UK

Abstract

Gossip protocols deal with a group of communicating agents, each holding some private information, and aim at arriving at a situation in which all the agents know each other secrets. Distributed epistemic gossip protocols are particularly simple distributed programs that use as guards formulas from an epistemic logic. We showed recently that the implementability of these distributed gossip protocols and the problems of their partial correctness and termination are decidable, but the problem of decidability of their fair termination was left open. We study here rule-fair and agent-fair termination of these protocols and show that both properties are decidable.

1 Introduction

1.1 Background and motivation

The aim of this paper is to study fair termination questions for gossip protocols. The set up of these protocols is the following. Each agent holds a secret initially known only to him. During the communications (for example phone calls) the participating agents share some of the secrets they know. The aim of the gossip protocols is to arrive at a situation in which all agents know all secrets. One of the early results established by a number of authors in the seventies, for instance [16], is that under the assumption that all secrets during every phone calls are shared, for $n \geq 4$ agents at least $2n - 4$ phone calls are sufficient and necessary to reach the situation in which all agents know all secrets.

To see the sufficiency of this number of calls assume that the set of agents is $\{a, b, c, d, e_1, \ldots, e_{n-4}\}$, where $n \geq 4$. Then take the call sequence $(a, e_1), (a, e_2), \ldots, (a, e_{n-4})$.. After it agent a knows all the secrets of $e_1, \ldots, e_{n-4}$. We follow it by the call sequence $(a, b), (c, d), (a, d), (b, c)$. After it agents a, b, c, d know all the secrets. Finally, by appending the call sequence $(a, e_1), (a, e_2), \ldots, (a, e_{n-4})$ at the end, we reach the situation in which all agents know all the secrets.

The above protocol is centralised in the sense that it requires a centralised scheduler. We are concerned here with specific distributed gossip protocols that were introduced in [5] and further studied with different type of calls in [2]. These protocols use as guards epistemic formulas and thus are examples of ***knowledge based programs*** introduced in [8].

The formulation of distributed gossip protocols as knowledge-based programs considerably simplifies the task of their verification. The reason is that these protocols are defined simply as a parallel composition of simple loops in which the agents repeatedly evaluate a guard, which is an epistemic formula, and subsequently perform the corresponding call. Consequently partial correctness of the protocol can be reduced to the problem of deciding the truth of formulas of the underlying epistemic language.

We established in [3] that such distributed epistemic gossip protocols (in short, gossip protocols) are implementable (i.e., the problem of evaluating a guard after a sequence of calls is decidable) and that the problems of their partial correctness and termination are decidable in the setting when during each call the participating agents exchange all their secrets and the underlying topology of the network is a clique. However, the decidability of their fair termination was left open.

In the paper, we analyse the fair termination problem for such gossip protocols. In distributed systems fairness can be defined in a number of ways, see [1]. Here we analyse two such possibilities: rule-fairness and agent-fairness. First, we stipulate that each finite computation is rule-fair and agent-fair. We say that an infinite computation is ***rule-fair*** if every call that is infinitely often enabled is also infinitely often executed. We also say that an infinite computation is ***agent-fair*** if every agent that is infinitely often enabled is also infinitely often selected. Then we say that a gossip protocol ***rule-fairly*** (resp. ***agent-fairly***) ***terminates*** if all rule-fair (resp. agent-fair) computations are finite. Agent-fairness was introduced in [2], where it was simply called fairness. We show that both types of fair termination for distributed epistemic gossip protocols that employ various type of calls are decidable.

1.2 Related work

Gossip protocols have been studied for more than forty years and have been successful in various domains, e.g., communication networks [9], computation of aggregate information [12], and data replication [14]. A more recent account is given in the book [11] and in [13]. In these references gossip protocols are viewed as parallel, probabilistic and/or distributed programs.

Epistemic gossip protocols were studied in a number of recent publications. In [4] a tool is discussed that given a high level description of an epistemic protocol in the setting of [5] generates the characteristics of the protocol. The calls considered there differ from the ones considered here, so their approach is not applicable to our setting. In turn, in [18] dynamic distributed gossip protocols are studied in which the calls allow the agents not only to share the secrets but also to transmit the links. The purpose of the paper is to characterise such protocols in terms of the class of graphs for which they terminate. Consequently these protocols differ from the ones considered here, which are static. Next, [10] and [7] consider gossip protocols that aim at achieving higher-order shared knowledge. Finally, in [6] gossip protocols are expressed as an instance of multi-agent epistemic planning and subsequently translated into the classical planning language PDDL.

1.3 Plan

The paper is organised as follows. In Sections 2 and 3, we recall the syntax and semantics of the considered epistemic logic, originally introduced in [2], though we provide here a more general framework for defining calls. Then, in Section 4, we recall the definition of distributed epistemic gossip protocols and illustrate their power on some examples. In Section 5 we recall the decidability results of checking whether a formula is true after a given sequence of calls and the definition of epistemic views from [3] that we rely upon. Section 6 contains the main results, namely decidability of determining whether a given gossip protocol rule-fairly or agent-fairly terminates. Finally, in Section 7 we discuss some open problems and explain that in the context of gossip protocols the notions of fairness and justice (a notion introduced in [15]) coincide.

2 Syntax

2.1 Calls and call types

Throughout the paper we assume a fixed finite set A of ***agents***. We assume that at the beginning each agent holds exactly one ***secret*** and that there exists a bijection between the set of agents and the set of secrets. We denote by S the set of all secrets. Our aim is to analyse what the agents know after a sequence of calls took place. We will first introduce different type of calls and then consider an epistemic language allowing us to refer to agents' knowledge.

A ***call*** concerns two agents, this call's caller x and callee y. In this paper we assume that a call is possible between any two agents; in other words, the communication graph is a clique. However, this is not a strong restriction and, as we shall see in Example 3, gossip protocols themselves can restrict the set of agents a given agent can call. Calls are denoted by c, d. Abusing notation we write $x \in \mathsf{c}$ to denote that agent x is one of the two agents involved in the call c. The ***type of a call*** is a function $\bowtie : 2^{\mathsf{S}} \times 2^{\mathsf{S}} \to 2^{\mathsf{S}} \times 2^{\mathsf{S}}$, which specifies the outcome of this call given the sets of secrets the caller and callee are familiar with. A call of type $\bowtie$ between caller x and callee y is written as $x \bowtie y$.

Following [2] we study here the following three types of calls.

- *Push-pull* calls between agents x and y, written as $x \circ y$ or simply xy, where agents exchange all their secrets. In this case we define $\circ(X, Y) := (X \cup Y, X \cup Y)$, where X and Y are, respectively, the set of secrets the caller and callee are familiar with before this call takes place.

- *Push* calls, written as $x \triangleright y$, where only the caller x passes his secrets to the callee y. In this case we define $\triangleright(X, Y) := (X, X \cup Y)$.

- *Pull* calls, written as $x \triangleleft y$, where only the caller x learns the secrets of the callee y. In this case we define $\triangleleft(X, Y) := (X \cup Y, Y)$.

2.2 Epistemic logic

We consider formulas in a simple epistemic language $\mathcal{L}$ defined by the following grammar:

$$\phi ::= F_a p \mid \neg\phi \mid \phi \wedge \phi \mid K_a \phi,$$

where $p \in \mathsf{S}$ and $a \in \mathsf{A}$. Each secret is viewed as a distinct constant. We denote the secret of agent a by A, the secret of agent b by B and so on.

We read $F_a p$ as 'agent a is familiar with the secret p' and $K_a \phi$ as 'agent a knows that formula ϕ is true'. So $F_a p$ is an atomic formula, while $K_a \phi$ is a compound formula. In fact, all atomic formulas of $\mathcal{L}$ have form $F_a p$.

In what follows we shall distinguish the following two sublanguages of $\mathcal{L}$:

- $\mathcal{L}_1$, which consists of the formulas without the nested use of the K_a modalities,

- $\mathcal{L}_1^a$, where $a \in \mathsf{A}$ is a fixed agent, which consists of the formulas from $\mathcal{L}_1$ in which the only modality allowed is K_a.

3 Semantics

We now recall from [2] semantics of the epistemic formulas. To this end we recall first the concept of a gossip situation.

3.1 Gossip situations

A ***gossip situation*** (in short a ***situation***) is a sequence $\mathsf{s} = (\mathsf{Q}_a)_{a \in \mathsf{A}}$, where $\mathsf{Q}_a \subseteq \mathsf{S}$ for each agent a. Intuitively, Q_a is the set of secrets agent a is familiar with in situation s. The ***initial gossip situation*** is the one in which each Q_a equals $\{A\}$ and is denoted by root. This situation reflects the fact that initially each agent is familiar only with his own secret. We say that an agent a is an ***expert*** in a situation s if he is familiar in s with all the secrets, i.e., if $\mathsf{Q}_a = \mathsf{S}$. We denote by G the set of all gossip situations.

We will use the following concise notation for gossip situations. Sets of secrets will be written down as lists. e.g., the set $\{A, B, C\}$ will be written as ABC. Gossip situations will be written down as lists of lists of secrets separated by dots. E.g., if there are three agents, then $\mathsf{root} = A.B.C$ while the gossip situation $(\{A, B\}, \{A, B\}, \{C\})$ will be written as $AB.AB.C$.

Each call transforms the current gossip situation by modifying the set of secrets the agents involved in the call are familiar with. Consider a gossip situation $\mathsf{s} := (\mathsf{Q}_d)_{d \in \mathsf{A}} \in \mathsf{G}$. Then $a \bowtie b(\mathsf{s}) := (Q'_d)_{d \in \mathsf{A}}$, where $(\mathsf{Q}'_a, \mathsf{Q}'_b) = \bowtie(\mathsf{Q}_a, \mathsf{Q}_b)$, and $\mathsf{Q}'_c = \mathsf{Q}_c$, for $c \neq a, b$. This simply says that a call $a \bowtie b$ only affects the secrets of the involved agents, a and b, and they are shared according to the semantics of $\bowtie$.

3.2 Call sequences

In [2] computations of the gossip protocols were studied, so both finite and infinite call sequences were used. For brevity, unless explicitly stated, a ***call sequence*** is assumed to be finite.

The empty call sequence is denoted by ϵ. We use $\mathbf{c}$ to denote a call sequence and $\mathbf{C}$ to denote the set of all finite call sequences. Given call sequences $\mathbf{c}$ and $\mathbf{d}$ and a call c we denote by $\mathbf{c}.\mathsf{c}$ the outcome of adding c at the end of the sequence $\mathbf{c}$ and by $\mathbf{c}.\mathbf{d}$ the concatenation of $\mathbf{c}$ and $\mathbf{d}$. Further, $\mathbf{c}^\omega$ denotes the infinite call sequence consisting of the infinite repetition of $\mathbf{c}$.

The result of applying a call sequence to a situation s is defined inductively by putting $\epsilon(\mathsf{s}) := \mathsf{s}$ and $(\mathsf{c}.\mathbf{c})(\mathsf{s}) := \mathbf{c}(\mathsf{c}(\mathsf{s}))$.

Example 1. Let $\mathsf{A} = \{a, b, c\}$. Consider the call sequence $(b \rhd c, a \lhd c, ac)$. It generates the following successive gossip situations starting from root: $A.B.C \xrightarrow{b \rhd c} A.B.BC \xrightarrow{a \lhd c} ABC.B.BC \xrightarrow{ac} ABC.B.ABC$. Hence $(b \rhd c, a \lhd c, ac)(\mathsf{root}) = (ABC.B.ABC)$. □

3.3 Gossip models and truth

A gossip situation is a set of possible distributions of secrets among the agents. As calls progress in sequence from the initial gossip situation, agents may be uncertain about which one of such secrets distributions is the actual one. This uncertainty is captured by the appropriate equivalence relations on the call sequences.

Definition 1. *A **gossip model** is a tuple $\mathcal{M} := (\mathbf{C}, \{\sim_a\}_{a \in \mathsf{A}})$, where each $\sim_a \subseteq \mathbf{C} \times \mathbf{C}$ is the minimal relation satisfying the following conditions:*

- $\epsilon \sim_a \epsilon$,

- *Suppose* $\mathbf{c} \sim_a \mathbf{d}$.

  (i) *If* $a \notin \mathsf{c}$, *then* $\mathbf{c}.\mathsf{c} \sim_a \mathbf{d}$ *and* $\mathbf{c} \sim_a \mathbf{d}.\mathsf{c}$.

  (ii) *If* $a \in \mathsf{c}$ *and* $\mathbf{c}.c(\mathsf{root})_a = \mathbf{d}.c(\mathsf{root})_a$, *then* $\mathbf{c}.\mathsf{c} \sim_a \mathbf{d}.\mathsf{c}$.

A gossip model with a designated call sequence is called a ***pointed gossip model****.*

For instance, by *(i)* we have $ab, bc \sim_a ab, bd$. But we do not have $bc, ab \sim_a bd, ab$ since $(bc, ab)(\mathsf{root})_a = ABC \neq ABD = (bd, ab)(\mathsf{root})_a$. Clearly, each $\sim_a$ is an equivalence relation.

Finally, we recall the definition of truth.

Definition 2. *Let* $(\mathcal{M}, \mathbf{c})$ *be a pointed gossip model with* $\mathcal{M} := (\mathbf{C}, (\sim_a)_{a \in \mathsf{A}})$ *and* $\mathbf{c} \in \mathbf{C}$. *We define the satisfaction relation* $\models$ *inductively as follows (clauses for Boolean connectives are as usual and omitted):*

$$\begin{array}{lll} (\mathcal{M}, \mathbf{c}) \models F_a p & \textit{iff} & p \in \mathbf{c}(\mathsf{root})_a, \\ (\mathcal{M}, \mathbf{c}) \models K_a \phi & \textit{iff} & \forall \mathbf{d} \textit{ s.t. } \mathbf{c} \sim_a \mathbf{d},\ (\mathcal{M}, \mathbf{d}) \models \phi. \end{array}$$

Further

$$\mathcal{M} \models \phi \quad \textit{iff} \quad \forall \mathbf{c}\ (\mathcal{M}, \mathbf{c}) \models \phi.$$

When $\mathcal{M} \models \phi$ *we say that* ϕ **is true***.* □

So a formula $F_a p$ is true whenever secret p belongs to the set of secrets agent a is familiar with in the situation generated by the designated call sequence $\mathbf{c}$ applied to the initial situation root. In turn, the knowledge operator is interpreted as is customary in epistemic logic, using the equivalence relations $\sim_a$.

4 Gossip Protocols

In [2], as a follow up on [5] we studied distributed epistemic gossip protocols. Their goal is to reach a gossip situation in which each agent is an expert. In other words, their goal is to transform a gossip situation in which the formula $\bigwedge_{a \in \mathsf{A}} F_a A \wedge \bigwedge_{a,b \in \mathsf{A}, a \neq b} \neg F_a B$ is true into one in which the formula $\bigwedge_{a,b \in \mathsf{A}} F_a B$ is true. Let us recall their definition.

By a ***component program***, in short a ***program***, for an agent a we mean a statement of the form $*[[]_{j=1}^{m}\ \psi_j \to \mathsf{c}_j]$, where $m > 0$ and each $\psi_j \to \mathsf{c}_j$ is such that $\psi_j \in \mathcal{L}_1^a$ and $a \in \mathsf{c}_j$.

Given a formula $\psi \in \mathcal{L}_1^a$ and a call c, we call the construct $\psi \to \mathsf{c}$ a ***rule*** and call ψ its ***guard***. A rule is ***enabled*** after a generated call sequence $\mathbf{c}$ if its guard is true after $\mathbf{c}$. Given a gossip protocol an agent is ***enabled*** after a call sequence $\mathbf{c}$ if one of the rules in its program is enabled. Intuitively, $*$ denotes a repeated execution, one at a time, of the enabled rules. Finally, by a ***distributed epistemic gossip protocol***, in short a ***gossip protocol***, we mean a parallel composition of component programs, one for each agent.

Assume now a gossip protocol P that is a parallel composition of the component programs $*[[]_{j=1}^{m_a}\ \psi_j^a \to \mathsf{c}_j^a]$, one for each agent $a \in \mathsf{A}$.

The ***computation tree*** of P is defined as the (possibly infinite) set $\mathbf{C}^P$ of (possibly infinite) call sequences $\mathbf{c} = \mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_n, \ldots$ such that:

- $\mathbf{C}^P$ is closed under prefixes,

- for any call sequence $(\mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_i, \mathsf{c}_{i+1})$ in it: for some a and $j \in \{1, \ldots, m_a\}$ we have $(\mathcal{M}, (\mathsf{c}_0, \ldots, \mathsf{c}_i)) \models \psi^a_j$ and $\mathsf{c}^a_j = \mathsf{c}_{i+1}$.

  In this case we say that a transition between $(\mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_i)$ and $(\mathsf{c}_0, \mathsf{c}_1, \ldots, \mathsf{c}_i, \mathsf{c}_{i+1})$ took place due to the ***selection*** of the rule $\psi^a_j \to \mathsf{c}^a_j$ and that agent a was ***selected***.

By a ***computation*** of a gossip protocol we mean a maximal rooted path in its computation tree. We stipulate that each finite computation is ***rule-fair*** and ***agent-fair***. An infinite computation is ***rule-fair*** (resp. ***agent-fair***) if all rules (resp. agents) that are enabled after infinitely many prefixes (in short, infinitely often) are selected infinitely often. We say that the gossip protocol P is ***partially correct***, in short ***correct***, if for all finite computations $\mathbf{c}$ that are leaves of the computation tree of P, the following holds:

$$(\mathcal{M}, \mathbf{c}) \models \bigwedge_{a,b \in \mathsf{A}} F_a B,$$

i.e., if for all call sequences $\mathbf{c}$ that are leaves of the computation tree of P, each agent is an expert in the gossip situation $\mathbf{c}(\mathsf{root})$. Note that $\mathbf{c}$ is a finite computation iff

$$(\mathcal{M}, \mathbf{c}) \models \bigwedge_{a \in \mathsf{A}} \bigwedge_{j=1}^{m_a} \neg\psi^a_j.$$

We call the formula

$$\bigwedge_{a \in \mathsf{A}} \bigwedge_{j=1}^{m_a} \neg\psi^a_j$$

the ***exit condition*** of the gossip protocol P. So P is partially correct iff the implication

$$\bigwedge_{a \in \mathsf{A}} \bigwedge_{j=1}^{m_a} \neg\psi^a_j \to \bigwedge_{a,b \in \mathsf{A}} F_a B \tag{1}$$

is true. We say furthermore that gossip protocol P ***terminates*** if all its computations are finite. In turn, P ***rule-fairly terminates*** (resp. ***agent-fairly terminates***) if all rule-fair (agent-fair) computations are finite.

Agent-fair termination was defined in [2], where it was called fair termination, but its decidability was not studied so far, while rule-fair termination was not considered. Notice that agent-fair termination implies rule-fair termination because any rule-fair computation is also an agent-fair computation, but as Example 4 below shows not the other way around.

To illustrate the power of the gossip protocols and various aspects of their behaviour we now present some examples. We begin with the following simple example considered in [5] and [2].

Example 2. *Consider a gossip protocol with the following program for each agent $i \in \mathsf{A}$:*

$$*[[]_{j \in \mathsf{A}} \neg F_i J \to i \lhd j].$$

Informally, agent i can make a pull call to agent j if he is not familiar with his secret. Recall that this protocol is obviously partially correct since its exit condition is $\bigwedge_{i,j \in \mathsf{A}} \neg\neg F_i J$ and that it always terminates because the cardinality of the set $\{(i,j) \in \mathsf{A} \times \mathsf{A} \mid \neg F_i J\}$ decreases by one after each call.

Our next example, taken from [2], illustrates that even if a gossip protocol may not always terminate, it may still always rule-fairly or agent-fairly terminate. It also shows how the communication between agents can induce any given graph topology (in this case a ring). First, let us define $i \oplus 1 = i \texttt{ mod } k + 1$ and $i \ominus 1 = (i-2) \texttt{ mod } k + 1$.

Example 3. *Let $\mathsf{A} = \{1, \ldots, k\}$ where $k \geq 3$. The secret of agent $i \in \{1, \ldots, k\}$ is denoted by I. In particular the secret of agent $i \ominus 1$ is denoted by $I \ominus 1$.*

Consider a gossip protocol with the following program for each agent $i \in \mathsf{A}$:

$$*[(\neg \bigwedge_{a \in \mathsf{A}} F_i A) \vee \neg K_i F_{i \oplus 1} I \ominus 1 \to i \circ (i \oplus 1)].$$

Informally, agent i calls his successor, agent $i \oplus 1$, if i is not familiar with all the secrets or i does not know that his successor is familiar with the secret of his predecessor, agent $i \ominus 1$.

In this protocol each agent has just one rule, so agent-fairness and rule-fairness coincide. It was shown in [2] that this gossip protocol is partially correct, but it does not always terminate. However, it was shown there that it always agent-fairly terminates, and a fortiori *also rule-fairly terminates.*

The next example shows that rule-fair termination and agent-fair termination may differ.

Example 4. *Consider a gossip protocol with the following program for each agent $i \in \mathsf{A}$:*

$$*[[]_{j \in \mathsf{A}} \neg \bigwedge_{a \in \mathsf{A}} F_i A \to i \lhd j].$$

Intuitively, agent i can make a pull call to any other agent as long as i is not an expert. This protocol is partially correct, since the implication (1) clearly holds. However, it may not terminate or even agent-fairly terminate, because if $\mathsf{A} = \{a_1, \ldots, a_k\}$, where $k \geq 3$, then $(a_1 \lhd a_2, a_2 \lhd a_1, a_3 \lhd a_1, \ldots, a_k \lhd a_1)^\omega$ is an infinite agent-fair computation of this protocol. Indeed, in this sequence all agents are infinitely often selected and all of them learn only the secrets of agents a_1 and a_2. So prior to each call in the above sequence the corresponding guard is true and consequently this sequence is a legal computation.

On the other hand, this protocol rule-fairly terminates. Indeed, consider an infinite computation χ. Some agent, say i, is then infinitely selected in χ, so it never becomes an expert and hence by the form of the protocol all the rules of i are always enabled. In χ agent i never becomes familiar with the secret of some agent, say j. So the rule $\neg \bigwedge_{a \in \mathsf{A}} F_i A \to i \lhd j$ is never selected in χ. Thus χ is not rule-fair.

Finally, we exhibit a protocol which is partially correct but does not even rule-fairly terminate.

Example 5. *Consider the following gossip protocol with the following program for agent i:*

$$*[[]_{j \in \mathsf{A}} \neg K_i F_j I \to i \lhd j].$$

Informally, agent i can make a pull call to agent j if agent i does not know whether agent j is familiar with his secret. We explained in [2] that this protocol is partially correct but may not agent-fairly terminate. We show here that it does not even rule-fairly terminate.

Let $\mathsf{A} = \{a_1, \ldots, a_k\}$. Consider the following sequence of calls $\mathbf{c}_1 = a_1 \lhd a_2, a_1 \lhd a_3, a_1 \lhd a_3, \ldots, a_1 \lhd a_k$. After $\mathbf{c}_1$ each agent a_i knows that agent a_1 is familiar with his secret, so no agent a_i can call a_1 anymore. Next, consider $\mathbf{c}_2 = a_2 \lhd a_3, a_2 \lhd a_4, \ldots, a_2 \lhd a_k$. After $\mathbf{c}_1.\mathbf{c}_2$ each *agent a_i knows that both agents a_1 and a_2 are familiar with his secret, so no agent a_i can call a_1 or a_2 anymore.*

In general, let $\mathbf{c}_l = a_l \lhd a_{l+1}, a_l \lhd a_{l+2}, \ldots, a_l \lhd a_k$ for $l \in \{1, \ldots, k-1\}$ and consider $\mathbf{c} = \mathbf{c}_1.\mathbf{c}_2.\ldots.\mathbf{c}_{k-1}$. After $\mathbf{c}$ each agent a_i knows that all agents except a_k are familiar with his secret, so no agent a_i can call $a_1, \ldots, a_{k-1}$ anymore. In other words, only agent a_k can be called after $\mathbf{c}$. Moreover, after $\mathbf{c}$ each rule $\neg K_i F_j I \rightarrow i \lhd j$, where $i \neq a_k$, remains enabled as long as agent a_k is not a caller. It follows that the infinite call sequence $\mathbf{c}.(a_1 \lhd a_k, \ldots, a_{k-1} \lhd a_k)^\omega$ is a rule-fair computation.

5 Decidability of Semantics and Epistemic Views

In this section we summarise the main definitions and results from [3] that we shall need in the next section to prove decidability of rule-fair and agent-fair termination of gossip protocols. We start by recalling the decidability of semantics, which implies the implementability of gossiping protocols.

Theorem 1 (Decidability of Semantics)**.** *For each call sequence $\mathbf{c}$, it is decidable whether for a formula $\phi \in \mathcal{L}_1$, $(\mathcal{M}, \mathbf{c}) \models \phi$ holds.*

We now recall the key notion of ***epistemic view*** used in [3]. It is a function of a call sequence $\mathbf{c}$, denoted by $\mathsf{E}V(\mathbf{c}) : \mathsf{A} \cup \{*\} \rightarrow 2^{\mathsf{G}}$, and defined by

- putting for each agent $a \in \mathsf{A}$, $\mathsf{E}V(\mathbf{c})(a) = \{\mathbf{d}(\mathsf{root}) \mid \mathbf{c} \sim_a \mathbf{d}\}$, and setting
- $\mathsf{E}V(\mathbf{c})(*) = \mathbf{c}(\mathsf{root})$.

So $\mathsf{E}V(\mathbf{c})(a)$ is the set of all gossip situations consistent with agent a's observations made throughout $\mathbf{c}$ and $\mathsf{E}V(\mathbf{c})(*)$ is the actual gossip situation after $\mathbf{c}$ takes place.

Lemma 1. *For each call sequence $\mathbf{c}$ and agent a the set $\mathsf{E}V(\mathbf{c})(a)$ is finite and can be effectively constructed.*

Our interest in epistemic views stems from the following result.

Lemma 2. *Suppose that $\mathsf{E}V(\mathbf{c}) = \mathsf{E}V(\mathbf{d})$. Then for all formulas $\phi \in \mathcal{L}_1$, $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.*

The above lemma is useful because the set of epistemic views is finite, in contrast to the set of call sequences.

Finally, we recall the following crucial concept. Consider a call sequence $\mathbf{c}$. If for some prefix $\mathbf{c}_1.\mathbf{c}_2$ of $\mathbf{c}$, we have $\mathsf{E}V(\mathbf{c}_1) = \mathsf{E}V(\mathbf{c}_1.\mathbf{c}_2)$, then we say that the call subsequence $\mathbf{c}_2$ is ***epistemically redundant*** in $\mathbf{c}$ and that $\mathbf{c}$ is ***epistemically redundant***.

We say that $\mathbf{c}$ is ***epistemically non-redundant*** if it is not epistemically redundant. Equivalently, a call sequence $\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k$ is epistemically non-redundant if the set

$$\{\mathsf{E}V(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_i) \mid i \in \{1, \ldots, k\}\}$$

has k elements.

Example 6. *Let us consider a model with three agents $\mathsf{A} = \{a, b, c\}$ and look at all epistemic views along the call sequence ab.ac.ab.ac.*

$$\mathsf{E}V(\epsilon)(*) = A.B.C$$

$$
\begin{aligned}
&\mathsf{E}V(\epsilon)(a) = \{A.B.C, A.BC.BC\}\\
&\mathsf{E}V(\epsilon)(b) = \{A.B.C, AC.B.AC\}\\
&\mathsf{E}V(\epsilon)(c) = \{A.B.C, AB.AB.C\}\\
&\mathsf{E}V(ab)(*) = AB.AB.C\\
&\mathsf{E}V(ab)(a) = \{AB.AB.C, AB.ABC.ABC\}\\
&\mathsf{E}V(ab)(b) = \{AB.AB.C, ABC.AB.ABC\}\\
&\mathsf{E}V(ab)(c) = \{A.B.C, AB.AB.C\}\\
&\mathsf{E}V(ab.ac)(*) = ABC.AB.ABC\\
&\mathsf{E}V(ab.ac)(a) = \{ABC.AB.ABC, ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac)(b) = \{AB.AB.C, ABC.AB.ABC\}\\
&\mathsf{E}V(ab.ac)(c) = \{ABC.AB.ABC, ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab)(*) = ABC.ABC.ABC\\
&\mathsf{E}V(ab.ac.ab)(a) = \{ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab)(b) = \{ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab)(c) = \{ABC.AB.ABC, ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab.ac)(*) = ABC.ABC.ABC\\
&\mathsf{E}V(ab.ac.ab.ac)(a) = \{ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab.ac)(b) = \{ABC.ABC.ABC\}\\
&\mathsf{E}V(ab.ac.ab.ac)(c) = \{ABC.AB.ABC, ABC.ABC.ABC\}
\end{aligned}
$$

This shows that the second call ac in the call sequence ab.ac.ab.ac is epistemically redundant and no other call is epistemically redundant in this call sequence.

The following states that epistemically redundant calls can be removed without affecting the consequent epistemic views.

Lemma 3 (Epistemic Stuttering)**.** *Suppose that* $\mathbf{c} := \mathbf{c}_1.\mathbf{c}_2.\mathbf{c}_3$ *and* $\mathbf{d} := \mathbf{c}_1.\mathbf{c}_3$*, where* $\mathbf{c}_2$ *is epistemically redundant in* $\mathbf{c}$*. Then* $\mathsf{E}V(\mathbf{c}) = \mathsf{E}V(\mathbf{d})$*.*

Next, we state the following crucial lemma.

Lemma 4. *For every given gossip model* $\mathcal{M}$*, there are only finitely many epistemically non-redundant call sequences.*

6 Decidability of Fair Termination

We proved in [3] that it is decidable to determine whether a given gossip protocol terminates. We now show that it is also decidable to determine whether a given gossip protocol agent-fairly or rule-fairly terminates.

We shall need the following consequence of the results listed in the previous section.

Lemma 5. *Suppose that* $\overline{\mathbf{c}} = \mathsf{c}_1.\mathsf{c}_2.\ldots$ *is a (possibly infinite) computation of a gossip protocol* P *such that a call* c_i *is epistemically redundant in the prefix* $\mathsf{c}_1.\ldots.\mathsf{c}_i$*. Then* $\overline{\mathbf{c}}$ *with the call* c_i *removed is also a computation of* P*.*

Proof. By definition for every $k \geq i$ the call c_i is epistemically redundant in $\mathsf{c}_1.\ldots.\mathsf{c}_k$, so by Lemma 2 for every $k \geq i$ we have $\mathsf{E}V(\mathsf{c}_1.\ldots.\mathsf{c}_k) = \mathsf{E}V(\mathsf{c}_1.\ldots.\mathsf{c}_{i-1}.\mathsf{c}_{i+1}.\ldots.\mathsf{c}_k)$. Thus by the Epistemic Stuttering Lemma 3 for all formulas $\phi \in \mathcal{L}_1$

$$(\mathcal{M}, \mathsf{c}_1.\ldots.\mathsf{c}_k) \models \phi \text{ iff } (\mathcal{M}, \mathsf{c}_1.\ldots.\mathsf{c}_{i-1}.\mathsf{c}_{i+1}.\ldots.\mathsf{c}_k \models \phi.$$

This implies the claim. □

In what follows we adjust the approach used in [3] to deal with customary termination. It relies on establishing an appropriate form of monotonicity of epistemic views with respect to the call sequence extensions. Informally, as the call sequence gets longer each agent acquires more information. This information is captured by the current epistemic view.

This approach required an introduction of suitable partial orderings $\leq_s$ and $\leq_{ev}$ over gossip situations and epistemic views that we now recall.

Definition 3. *For any two gossip situations* s, s' *we write* $\mathsf{s} \leq_s \mathsf{s}'$ *if for all* $a \in \mathsf{A}$ *we have* $\mathsf{s}_a \subseteq \mathsf{s}'_a$.

Note 1. *For all call sequences* $\mathbf{c}$ *and* $\mathbf{d}$ *such that* $\mathbf{c} \preceq \mathbf{d}$ *we have* $\mathbf{c}(\mathsf{root}) \leq_s \mathbf{d}(\mathsf{root})$.

Definition 4. *For any two epistemic views* $V, V' \in \widetilde{\mathsf{E}V}$ *we write* $V \leq_{ev} V'$ *if for all* $a \in \mathsf{A}$ *there exists* $X \subseteq V(a)$ *and an surjective (onto) function* $g : X \to V'(a)$ *such that for all* $\mathsf{s} \in X$ *we have* $\mathsf{s} \leq_s g(\mathsf{s})$.

Lemma 6. $\leq_{ev}$ *is a partial order.*

We provide the proof of this crucial claim as it was omitted in [3].

Proof.
(Reflexivity) For any epistemic view V, we have $V \leq_{ev} V$, because for each $a \in \mathsf{A}$ we can pick $V(a)$ as X and the identity function on $V(a)$ as g.

(Transitivity) Suppose V, V', V'' are three epistemic views such that $V \leq_{ev} V'$ and $V' \leq_{ev} V''$. Then, from the definition of $\leq_{ev}$, for any $a \in \mathsf{A}$ there exist $X \subseteq V(a)$, $Y \subseteq V'(a)$, and surjective functions $g : X \to V'(a)$ and $h : Y \to V''(a)$. Let $Z = \{\mathsf{s} \in X \mid g(\mathsf{s}) \in Y\}$. Note that $g|_Z : Z \to Y$, i.e. the restriction of g to Z, is surjective. The composition $g|_Z \circ h : Z \to V''(a)$ is also surjective and for any gossip situation $\mathsf{s} \in Z$ the following holds $\mathsf{s} \leq_s g|_Z(\mathsf{s}) \leq_s h(g|_Z(\mathsf{s})) = (g|_Z \circ h)(\mathsf{s})$.

(Antisymmetry) Suppose V, V' are two epistemic views such that $V \leq_{ev} V'$ and $V' \leq_{ev} V$. Then, from the definition of $\leq_{ev}$, for any $a \in \mathsf{A}$ there exist $X \subseteq V(a)$, $Y \subseteq V'(a)$, and surjective functions $g : X \to V'(a)$ and $h : Y \to V(a)$. Let $Z = \{\mathsf{s} \in X \mid g(\mathsf{s}) \in Y\}$. Note that $g|_Z : Z \to Y$, i.e. the restriction of g to Z, is surjective. Moreover, $g|_Z \circ h : Z \to V(a)$ is also surjective, and because $Z \subseteq V(a)$ is finite, $Z = V(a)$, $g|_Z = g$, and $g \circ h$ is a permutation on $V(a)$. Similarly we can show that $Y = V'(a)$. Since $(g \circ h)$ is a permutation on a finite set, there exists k such that $(g \circ h)^k$ is the identity function on $V(a)$. Note that for any $\mathsf{s} \in V(a)$, we have $\mathsf{s} \leq_s (g \circ h)(\mathsf{s})$, because $\mathsf{s} \leq_s g(\mathsf{s}) \leq_s h(g(\mathsf{s})))$. Now consider the sequence: $\mathsf{s} \leq_s (g \circ h)(\mathsf{s}) \leq_s (g \circ h)^2(\mathsf{s}) \leq_s \ldots \leq_s (g \circ h)^k(\mathsf{s}) = \mathsf{s}$. In fact, all of the elements in this sequence have to be the same, because $\leq_s$ is a partial order. In particular, this shows that $(g \circ h)(\mathsf{s}) = \mathsf{s}$. Therefore, $g \circ h$ is the identity function on $V(a)$. Now, for any $\mathsf{s} \in V(a)$ we have that $\mathsf{s} \leq_s g(\mathsf{s}) \leq_s h(g(\mathsf{s})) = (g \circ h)(\mathsf{s}) = \mathsf{s}$, so g is the identity function as well. This shows that $V(a) = V'(a)$ for all $a \in \mathsf{A}$. □

Finally, we recall the following lemma from [3] which formalises the intuition that epistemic information grows along a call sequence.

Lemma 7. *For all two call sequences such that* $\mathbf{c} \preceq \mathbf{d}$ *we have* $\mathsf{E}V(\mathbf{c}) \leq_{ev} \mathsf{E}V(\mathbf{d})$.

We are now ready to establish the decidability of fair termination for gossip protocols. We start with the rule-fair termination.

Theorem 2 (Decidability of Rule-Fair Termination)**.** *Given a gossip protocol that does not use nested modalities, it is decidable to determine whether it rule-fairly terminates.*

Proof. We first show that a gossip protocol fails to rule-fairly terminate iff it can generate an epistemically non-redundant call sequence $\mathbf{c}$ such that for every call c, which is part of an enabled rule after the call sequence $\mathbf{c}$, we have that $\mathsf{E}V(\mathbf{c}.\mathsf{c}) = \mathsf{E}V(\mathbf{c})$.

$(\Rightarrow)$ Consider an infinite rule-fair computation $\overline{\mathbf{d}} = \mathsf{d}_1.\mathsf{d}_2.\ldots$ of the considered gossip protocol. By Lemma 7 the sequence $\mathsf{E}V(\mathsf{d}_1)$, $\mathsf{E}V(\mathsf{d}_1.\mathsf{d}_2), \ldots,$ is weakly increasing w.r.t. the partial order $\leq_{ev}$. As there are only finitely many epistemic views, at some point this sequence stabilises, i.e., for some l we have $\mathsf{E}V(\mathsf{d}_1.\ldots.\mathsf{d}_l) = \mathsf{E}V(\mathsf{d}_1.\ldots.\mathsf{d}_l.\mathsf{d}_{l+1}.\ldots.\mathsf{d}_{l+i})$ for all $i > 0$. Pick the smallest such l and let $\mathbf{d} = \mathsf{d}_1.\ldots.\mathsf{d}_l$. By Lemma 5 we can repeatedly remove the epistemically redundant calls from $\mathbf{d}$ without destroying the property that it is a prefix of an infinite computation.

Moreover, the resulting infinite computation $\overline{\mathbf{c}} = \mathsf{c}_1.\mathsf{c}_2.\ldots$ of the protocol is rule-fair, as well, for some k the call sequence $\mathbf{c} = \mathsf{c}_1.\ldots.\mathsf{c}_k$ (resulting from the repeated removal of epistemically redundant calls from $\mathbf{d}$) is epistemically non-redundant, and by the above choice of l and the Epistemic Stuttering Lemma 3 $\mathsf{E}V(\mathbf{c}) = \mathsf{E}V(\mathbf{c}.\mathsf{c}_{k+1}.\ldots.\mathsf{c}_{k+i})$ for all $i > 0$.

Take a rule $\psi \to \mathsf{c}$ that is enabled after $\mathbf{c}$, i.e., such that $(\mathcal{M}, \mathbf{c}) \models \psi$. By Lemma 2 and the choice of $\mathbf{c}$, this rule is enabled after each call sequence $\mathbf{c}.\mathsf{c}_{k+1}.\ldots.\mathsf{c}_{k+i}$, where $i > 0$, that is, it is enabled infinitely often. By the rule-fairness of $\overline{\mathbf{c}}$ this rule $\psi \to \mathsf{c}$ is infinitely often selected in it. So for some $i > 1$ we have $\mathsf{c} = \mathsf{c}_{k+i}$.

By the choice of k the call sequence $\mathsf{c}_{k+1}.\ldots.\mathsf{c}_{k+i-1}$ is epistemically redundant in $\mathsf{c}_1.\ldots.\mathsf{c}_{k+i}$, so by the above equality and the Epistemic Stuttering Lemma 3

$$\mathsf{E}V(\mathsf{c}_1.\ldots.\mathsf{c}_k) = \mathsf{E}V(\mathsf{c}_1.\ldots.\mathsf{c}_k.\mathsf{c}_{k+1}.\ldots.\mathsf{c}_{k+i}) = \mathsf{E}V(\mathsf{c}_1.\ldots.\mathsf{c}_k.\mathsf{c}_{k+i}),$$

i.e., $\mathsf{E}V(\mathbf{c}.\mathsf{c}) = \mathsf{E}V(\mathbf{c})$ as required.

$(\Leftarrow)$ Suppose that the protocol generates a sequence of calls $\mathbf{c}$ such that $\mathbf{c}$ is epistemically non-redundant and $\mathsf{E}V(\mathbf{c}.\mathsf{c}) = \mathsf{E}V(\mathbf{c})$ for every call c which is part of a enabled rule after the call sequence $\mathbf{c}$ takes place.

Let $R = \{\phi_1 \to \mathsf{c}_1, \phi_2 \to \mathsf{c}_2, \ldots, \phi_k \to \mathsf{c}_k\}$ be the set of all enabled rules after the call sequence $\mathbf{c}$. We claim that $\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^\omega$ is a rule-fair infinite computation of this protocol.

First, due to Epistemic Stuttering Lemma 3 for every $1 \leq j \leq k$ and $0 \leq i$ we have $\mathsf{E}V(\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^i.\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_j) = \mathsf{E}V(\mathbf{c})$. This and Lemma 2 implies that all rules in R are enabled after any call sequence of the form $\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^i.\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_j$ for any $j \in \{1, \ldots, k\}$ and $i \geq 0$. This shows that $\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^\omega$ is an infinite computation of this protocol. Also, we know that no other rule can be enabled after $\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^i.\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_j$, because otherwise such a rule would already be enabled after $\mathbf{c}$ and so would belong to R. This shows that $\mathbf{c}.(\mathsf{c}_1.\mathsf{c}_2.\ldots.\mathsf{c}_k)^\omega$ is a rule-fair infinite computation of this protocol, because every rule enabled infinitely many times is executed infinitely many times.

Now, due to Lemma 4 there are only finitely many epistemically non-redundant call sequences to try as candidates for $\mathbf{c}$. For each such call sequence, by the Decidability of Semantics Theorem 1 it is decidable to determine whether it can be generated by the protocol.

For each such call sequence $\mathbf{c}$ we then check which rules, $\psi \rightarrow \mathsf{c}$, are enabled after $\mathbf{c}$. For each such a call c we subsequently compute $\mathsf{E}V(\mathbf{c})$ and $\mathsf{E}V(\mathbf{c}.\mathsf{c})$ using Lemma 1 and check whether they are all equal. By the above equivalence the considered gossip protocol does not rule-fairly terminate iff for some such call sequence $\mathbf{c}$ all just mentioned equalities hold. □

Finally, we show that agent-fair termination is decidable as well.

Theorem 3 (Decidability of Agent-Fair Termination)**.** *Given a gossip protocol that do not use nested modalities, it is decidable to determine whether it agent-fairly terminates.*

Proof. First we show that a gossip protocol may fail to agent-fairly terminate iff it can generate an epistemically non-redundant call sequence $\mathbf{c}$ such that each agent a enabled after $\mathbf{c}$ has an enabled rule $\psi \rightarrow \mathsf{c}$ such that $\mathsf{E}V(\mathbf{c}.\mathsf{c}) = \mathsf{E}V(\mathbf{c})$ holds. The reasoning is completely analogous to the one presented in the proof of the previous theorem, so we omit the details.

The rest of the proof is a small modification of the reasoning used in the above proof. As before there are only finitely many epistemically non-redundant call sequences $\mathbf{c}$ and for each such call sequence it is decidable to determine whether it can be generated by the protocol.

For each such call sequence $\mathbf{c}$ we then check which agents are enabled after $\mathbf{c}$. For each such agent we then check whether it has a rule $\psi \rightarrow \mathsf{c}$ that is enabled after $\mathbf{c}$ and such that $\mathsf{E}V(\mathbf{c}) = \mathsf{E}V(\mathbf{c}.\mathsf{c})$, where, again, this test is decidable by Lemma 1. By the initial equivalence the considered gossip protocol does not agent-fairly terminate iff for some such call sequence $\mathbf{c}$ it is possible to choose the rules in such a way that all the equalities hold. □

7 Conclusions

In this paper we established the decidability of rule-fair and agent-fair termination of gossip protocols. An interesting future work would be to study the same problems for gossip protocols with nested modalities or with a common knowledge operator. Another interesting issue is to study the synthesis of a distributed epistemic gossip protocol from epistemic specifications (see, e.g., [17]). Finally, it would be interesting to establish the exact computational complexity of implementability of a gossip protocol and of checking its partial correctness, termination, and fair termination.

We conclude by mentioning a notion related to fairness, called ***justice*** (or ***weak fairness***), see [15]. An infinite computation is ***rule-just*** (resp. ***agent-just***) if all rules (resp. agents) that from a certain moment on are always enabled (in short, eventually always enabled) are selected infinitely often. The notion of infinite just and fair computations differ in the context of nondeterministic programs. However, this is not the case for the gossip protocols.

Indeed, it is straightforward to see that an infinite rule-fair computation is also rule-just. Take now an infinite rule-just computation $\bar{\mathbf{c}} = \mathsf{c}_1.\mathsf{c}_2....$ of a gossip protocol. As in the proof of Theorem 2, on the account of Lemma 7 and the fact that there are only finitely many epistemic views, for some l we have $\mathsf{E}V(\mathsf{c}_1.\dots.\mathsf{c}_l) = \mathsf{E}V(\mathsf{c}_1.\dots.\mathsf{c}_l.\mathsf{c}_{l+1}.\dots.\mathsf{c}_{l+i})$ for all $i > 0$.

Suppose now that a rule, say $\psi \rightarrow \mathsf{c}$, is infinitely often enabled. By Lemma 2 for all $i > 0$

$$(\mathcal{M}, \mathsf{c}_1.\dots.\mathsf{c}_l) \models \psi \text{ iff } (\mathcal{M}, \mathsf{c}_1.\dots.\mathsf{c}_l.\mathsf{c}_{l+1}.\dots.\mathsf{c}_{l+i}) \models \psi,$$

so $\psi \rightarrow \mathsf{c}$ is eventually always enabled. Since $\bar{\mathbf{c}}$ is rule-just, this rule is selected infinitely often.

An analogous argument shows that infinite agent-just and agent-fair computations coincide.

Acknowledgements. The first author was partially supported by the NCN grant nr 2014/13/B/ST6/01807. The second author was partially supported by the EPSRC through grants EP/M027287/1 and EP/P020909/1.

References

[1] K. R. Apt, N. Francez, and S. Katz. Appraising fairness in languages for distributed programming. *Distributed Computing*, 2(4):226–241, 1988.

[2] K. R. Apt, D. Grossi, and W. Van der Hoek. Epistemic protocols for distributed gossiping. In *Proc. of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015)*, volume 215 of *EPTCS*, pages 51–66, 2016.

[3] K. R. Apt and D. Wojtczak. On decidability of a logic of gossips. In *Proc. of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016)*, volume 10021 of *Lecture Notes in Computer Science*, pages 18–33. Springer, 2016.

[4] M. Attamah, H. van Ditmarsch, D. Grossi, and W. Van der Hoek. A framework for epistemic gossip protocols. In *Proc. of the 12th European Conference on Multi-Agent Systems (EUMAS 2014)*, pages 193–209, 2014.

[5] M. Attamah, H. van Ditmarsch, D. Grossi, and W. Van der Hoek. Knowledge and gossip. In *Proc. of ECAI'14*. IOS Press, 2014.

[6] M. C. Cooper, A. Herzig, F. Maffre, F. Maris, and P. Régnier. A simple account of multiagent epistemic planning. In *Proc. of ECAI'16*, pages 193–201. IOS Press, 2016.

[7] M. C. Cooper, A. Herzig, F. Maffre, F. Maris, and P. Regnier. Simple Epistemic Planning: Generalised Gossiping. In *Proc. of ECAI'16*, pages 1563–1564. IOS Press, 2016.

[8] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. Knowledge-based programs. *Distributed Computing*, 10(4):199–225, 1997.

[9] S. M. Hedetniemi, S. T. Hedetniemi, and A. L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

[10] A. Herzig and F. Maffre. How to share knowledge by gossiping. *AI Communications*, 30(1):1–17, 2017.

[11] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka, and W. Unger. *Dissemination of Information in Communication Networks - Broadcasting, Gossiping, Leader Election, and Fault-Tolerance*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2005.

[12] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proc. of 44th Symposium on Foundations of Computer Science (FOCS'03)*, pages 482–491. IEEE, 2003.

[13] A. Kermarrec and M. van Steen. Gossiping in distributed systems. *Operating Systems Review*, 41(5):2–7, 2007.

[14] R. Ladin, B. Liskov, L. Shrira, and S. Ghemawat. Providing high availability using lazy replication. *ACM Transactions on Computer Systems (TOCS)*, 10(4):360–391, 1992.

[15] D. J. Lehmann, A. Pnueli, and J. Stavi. Impartiality, justice, and fairness: the ethics of concurrent termination. In *Proc. of International Colloquium on Automata Languages and Programming (ICALP '81)*, pages 264–277. Springer-Verlag, 1981.

[16] R. Tijdeman. On a telephone problem. *Nieuw Archief voor Wiskunde*, 3(XIX):188–192, 1971.

[17] R. van der Meyden and T. Wilke. Synthesis of distributed systems from knowledge-based specifications. In *Proc. of 16th International Conference on Concurrency Theory (CONCUR'05)*, volume 3653 of *Lecture Notes in Computer Science*, pages 562–576. Springer, 2005.

[18] H. van Ditmarsch, J. van Eijck, P. Pardo, R. Ramezanian, and F. Schwarzentruber. Epistemic protocols for dynamic gossip. *Journal of Applied Logic*, 20:1–31, 2017.