

On Decidability of a Logic of Gossips

Krzysztof R. Apt¹ and Dominik Wojtczak²(✉)

¹ CWI, Amsterdam, The Netherlands

² University of Liverpool, Liverpool, UK
d.wojtczak@liverpool.ac.uk

Abstract. Gossip protocols aim at arriving, by means of point-to-point or group communications, at a situation in which all the agents know each other secrets, see, e.g., [11]. In [1], building upon [3], we studied distributed epistemic gossip protocols, which are examples of *knowledge based programs* introduced in [6]. These protocols use as guards formulas from a simple epistemic logic. We show here that these protocols are implementable by proving that it is decidable to determine whether a formula with no nested modalities is true after a sequence of calls. Building upon this result we further show that the problems of partial correctness and of termination of such protocols are decidable, as well.

1 Introduction

1.1 Background and Motivation

Knowledge-based programs were introduced in [6]—these are programs that use tests for knowledge. Examples are protocols for the sequence transmission problem, such as the alternating bit protocol, studied in [7]. A more recent example are the distributed epistemic gossip protocols introduced in [3] and further studied in a slightly different setting in [1].

In *gossip protocols* each agent holds a secret initially known only to him. The secrets spread by means of communications. During them, e.g., point-to-point or group communications, the participating agents exchange all secrets they know. The aim of the gossip protocols is to arrive at a situation in which all the agents know each other secrets, see, e.g., the early survey [8], the book coverage [10] or a more recent paper [11].

As shown in [1], the formulation of distributed gossip protocols as knowledge-based programs considerably simplifies the task of their verification. The reason is that these protocols are strikingly simple in their syntax based on epistemic logic (though not semantics)—they are just parallel compositions of loops in which the agents repeatedly perform a call assuming the corresponding epistemic guard evaluates to true. One issue ignored in [1] was the natural question: are these gossip protocols implementable?

In this paper we provide a positive answer to this question. More precisely, we show that it is decidable to determine whether a formula with no nested modalities is true after a sequence of calls. All gossip protocols studied in [3] use only such formulas as guards.

We also study correctness and termination of these protocols. Building upon the just mentioned result we show that it is decidable to determine whether a given distributed epistemic gossip protocol is correct. Namely, the formula that expresses its correctness is with no nested modalities and we show that for such formulas truth is decidable. The final result allows us to solve the halting problem for these protocols. This shows that the distributed epistemic gossip protocols are very specific programs that in particular do not have the full power of the Turing machines.

The obtained results, while sufficient for a study of the considered protocols, do not address more general questions concerning both the logic itself and the protocols, which remain open and to which we return in the conclusions.

Finally, let us mention here some recent works on gossip protocols. In [2] a tool is presented that given a high level description of an epistemic protocol in the setting of [3] generates the characteristics of the protocol. The calls considered there differ from ours, so this approach is not applicable to our setting. Further, [13] presents a study of dynamic distributed gossip protocols 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 characterize such protocols in terms of the class of graphs for which they terminate. Such protocols then differ from the ones here considered, which are static. Next, in [9] gossip protocols are studied that aim at achieving higher-order shared knowledge. Finally, in [4] gossip protocols are studied as an instance of multi-agent epistemic planning that is subsequently translated into the classical planning language PDDL.

1.2 Plan

The paper is organized as follows. In the next two sections we recall the syntax and semantics introduced in [1]. Then, in Sect. 4 we introduce an alternative, equivalent, semantics, which helps us to prove the desired decidability results. In Sect. 5 we prove the decidability of checking whether a formula with no nested modalities is true after a given sequence of calls, and in Sect. 6 we show how to extend this result to checking whether such a formula is true (so true after any sequence of calls). In turn, in Sect. 7 we show that it is also decidable to determine whether a given gossip protocol terminates. Then, in the final section, we list some related open problems and clarify the difference between the type of calls studied in [1,3].

2 Syntax

Throughout the paper we assume a fixed finite set A of at least three *agents*. We assume that 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 P the set of all secrets. Our aim is to analyze what the agents know after a sequence of calls took place. So first we introduce the calls and then consider an epistemic language allowing us to refer to agents' knowledge.

Assume a fixed ordering on the agents. Each **call** concerns two different agents, say a and b , and is written as ab , where agent a precedes agent b in the assumed ordering.

Calls are denoted by c, d . Abusing notation we write $a \in c$ to denote that agent a is one of the two agents involved in the call c (e.g., for $c := ab$ we have $a \in c$ and $b \in c$).

We consider formulas in a simple epistemic language defined by the following grammar:

$$\phi ::= F_ap \mid \neg\phi \mid \phi \wedge \phi \mid K_a\phi,$$

where $p \in P$ and $a \in A$. Each secret is viewed a distinct constant. We denote the secret of agent a by A , the secret of agent b by B and so on. We denote the set of so defined formulas by \mathcal{L} and we refer to its members as epistemic formulas.

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

In [1], as a follow up on [3], we also introduced distributed epistemic gossip protocols. We do not discuss them here and only mention that formulas of \mathcal{L} are used in them as guards. All guards used in [1] are built from the formulas F_aB and K_aF_bC , where a and b are different agents, by means of the Boolean connectives. Thus no nested modalities are used in the guards.

3 Semantics

We now recall from [1] semantics of the epistemic formulas. To this end we recall first the concept of a gossip situation.

3.1 Gossip Situations and Their Modifications

A **gossip situation** (in short a **situation**) is a sequence $s = (Q_a)_{a \in A}$, where $Q_a \subseteq P$ for each agent a . Intuitively, Q_a is the set of secrets 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**. 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 $Q_a = P$. The initial gossip situation reflects the fact that initially each agent is familiar only with his own secret.

In this paper we do not study particular gossip protocols. We mention only that their goal is to reach a gossip situation in which each agent is an expert.

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 **root** = $A.B.C$ and 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 $s := (Q_d)_{d \in A}$. Then $ab(s) := (Q'_d)_{d \in A}$, where $Q'_a = Q'_b = Q_a \cup Q_b$, $Q'_c = Q_c$,

for $c \neq a, b$. This simply says that the only effect of a call is that the secrets are shared between the two agents involved in it.

3.2 Call Sequences

In [1] computations of the gossip protocols were studied, so both finite and infinite call sequences were used. Here we limit ourselves to the finite call sequences as we are only interested in the semantics of epistemic formulas.

So in this paper, in contrast to [1], a **call sequence** is a *finite* sequence of calls. The empty sequence is denoted by ϵ . We use \mathbf{c} to denote a call sequence and \mathbf{C} to denote the set of all call sequences. Given call sequences \mathbf{c} and \mathbf{d} and a call c we denote by $\mathbf{c}.c$ the outcome of adding c at the end of the sequence \mathbf{c} and by $\mathbf{c}.\mathbf{d}$ the outcome of appending the sequences \mathbf{c} and \mathbf{d} . We write $\mathbf{c} \preceq \mathbf{d}$ to denote the fact that \mathbf{d} extends \mathbf{c} , i.e., that for some \mathbf{c}' we have $\mathbf{c}.\mathbf{c}' = \mathbf{d}$.

The result of applying a call sequence to a situation \mathbf{s} is defined inductively as follows:

$$\begin{aligned} [\text{Base}] \quad \epsilon(\mathbf{s}) &:= \mathbf{s}, \\ [\text{Step}] \quad (\mathbf{c}.\mathbf{c})(\mathbf{s}) &:= \mathbf{c}(\mathbf{c}(\mathbf{s})). \end{aligned}$$

Example 1. Let $\mathbf{A} = \{a, b, c\}$. Consider the call sequence (ac, bc, ac) . It generates the following successive gossip situations starting from **root**:

$$A.B.C \xrightarrow{ac} AC.B.AC \xrightarrow{bc} AC.ABC.ABC \xrightarrow{ac} ABC.ABC.ABC.$$

Hence $(ac, bc, ac)(\text{root}) = (ABC.ABC.ABC)$. \square

3.3 Gossip Models and Truth

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

Definition 1. A **gossip model** is a tuple $\mathcal{M} := (\mathbf{C}, \{\sim_a\}_{a \in \mathbf{A}})$, where each $\sim_a \subseteq \mathbf{C} \times \mathbf{C}$ is defined inductively as follows.

$$\begin{aligned} [\text{Base}] \quad \epsilon &\sim_a \epsilon; \\ [\text{Step}] \quad &\text{Suppose } \mathbf{c} \sim_a \mathbf{d}. \\ (i) \quad &\text{If } a \notin \mathbf{c}, \text{ then } \mathbf{c}.c \sim_a \mathbf{d} \text{ and } \mathbf{c} \sim_a \mathbf{d}.c. \\ (ii) \quad &\text{If } a \in \mathbf{c} \text{ and } \mathbf{c}.c(\text{root})_a = \mathbf{d}.c(\text{root})_a, \text{ then } \mathbf{c}.c \sim_a \mathbf{d}.c. \end{aligned}$$

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 $(bc, ab)(\text{root})_a = ABC \neq ABD = (bd, ab)(\text{root})_a$.

We recall now from [1] the following two properties of \sim_a .

Fact 1

- (i) Each \sim_a is an equivalence relation;
- (ii) For all $\mathbf{c}, \mathbf{d} \in \mathbf{C}$ if $\mathbf{c} \sim_a \mathbf{d}$, then $\mathbf{c}(\text{root})_a = \mathbf{d}(\text{root})_a$.

Finally, we recall the definition of truth.

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

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

Further

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

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

So 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 . The knowledge operator is interpreted as customary in epistemic logic using the equivalence relations \sim_a .

4 An Alternative Equivalence Relation

In this section we provide an alternative equivalence relation between the call sequences that is easier to work with. To this end we introduce a **view** of agent a of a call sequence \mathbf{c} , written as \mathbf{c}_a , and defined by induction as follows.

[Base]

$$\epsilon_a := \text{root},$$

[Step]

$$(\mathbf{c}.\mathbf{c})_a := \begin{cases} \mathbf{c}_a \xrightarrow{\mathbf{c}} \mathbf{s} & \text{if } a \in \mathbf{c} \\ \mathbf{c}_a & \text{otherwise} \end{cases}$$

where for $d \in A$

$$\mathbf{s}_d := \begin{cases} \mathbf{c}.\mathbf{c}(\text{root})_d & \text{if } d \in \mathbf{c} \\ \mathbf{s}'_d & \text{otherwise} \end{cases}$$

where \mathbf{s}' is the last gossip situation in \mathbf{c}_a .

Intuitively, a view of agent a of a call sequence \mathbf{c} is the information he acquires by means of the calls in \mathbf{c} he is involved in. It consists of a sequence of gossip situations connected by the calls in which a is involved in. After each such call, say ab , agent a updates the set of gossips he and b are currently familiar with.

Example 2. Let us return to Example 1. So $A = \{a, b, c\}$ and we consider the call sequence (ac, bc, ac) . We noticed there that it generates the following successive gossip situations starting from **root**:

$$A.B.C \xrightarrow{ac} AC.B.AC \xrightarrow{bc} AC.ABC.ABC \xrightarrow{ac} ABC.ABC.ABC.$$

We now compare it with the view of agent a of the sequence (ac, bc, ac) , which is

$$A.B.C \xrightarrow{ac} AC.B.AC \xrightarrow{ac} ABC.B.ABC.$$

Thus, in the final gossip situation of this view, agent b is familiar with neither the secret A nor C . However, the final gossip situation of a view does not reflect agents' knowledge. In fact, as we shall see, according to the semantics, after the above sequence of calls, agent a knows that agent b is familiar both with A and C . \square

We now introduce for each agent a an equivalence relation \equiv_a between the call sequences, defined as follows:

$$\mathbf{c} \equiv_a \mathbf{d} \text{ iff } \mathbf{c}_a = \mathbf{d}_a.$$

So according to this definition two call sequences are equivalent for agent a if his views of them are the same. The following result shows that the equivalence relations \sim_a and \equiv_a coincide.

Theorem 2 (Equivalence). *For each agent a the relations \sim_a and \equiv_a coincide.*

Proof. Omitted. \square

So two call sequences are \sim_a equivalent iff their views by agent a coincide. This alternative definition of the equivalence relation between the call sequences makes it simpler to determine various properties of our semantics.

Below, given a call \mathbf{c} , we denote by \mathbf{c}^* a sequence consisting of zero or more calls \mathbf{c} and by \mathbf{c}^+ a sequence consisting of one or more calls \mathbf{c} .

Example 3. Note that we have $(\mathcal{M}, (ac, bc, ac)) \models K_a F_b A$. To see this recall from Example 2 that the view of agent a of the sequence (ac, bc, ac) is

$$A.B.C \xrightarrow{ac} AC.B.AC \xrightarrow{ac} ABC.B.ABC.$$

So if $(ac, bc, ac) \equiv_a \mathbf{d}$, then \mathbf{d} is of the form $ac, (bc)^+, ac, (bc)^*$, which implies that $(\mathcal{M}, \mathbf{d}) \models F_b A$.

We conclude that it is possible that an agent, here a , knows that another agent, here b , is familiar with his (so a 's) secret even though no communication took place between them. The same argument shows that $(\mathcal{M}, (ac, bc, ac)) \models K_a F_b C$, as claimed in Example 2. \square

In the examples and proofs below we use the \equiv_a relation instead of \sim_a and repeatedly appeal to the Equivalence Theorem 2. First we show that an immediate repetition of a call has no effect on the truth of the formulas. More precisely, the following holds.

Theorem 3 (Stuttering). *Suppose that $\mathbf{c} := \mathbf{c}_1, \mathbf{c}, \mathbf{c}_2$ and $\mathbf{d} := \mathbf{c}_1, \mathbf{c}, \mathbf{c}, \mathbf{c}_2$. Then for all formulas ϕ , $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.*

Proof. We proceed by induction of the structure of ϕ . For the formulas of the form $F_a p$ it suffices to note that $\mathbf{c}(\text{root}) = \mathbf{d}(\text{root})$. The only induction step of interest is for the formulas of the form $K_a \phi$. Suppose first that $a \notin \mathbf{c}$. Then $\mathbf{c} \equiv_a \mathbf{d}$, so $(\mathcal{M}, \mathbf{c}) \models K_a \phi$ iff $(\mathcal{M}, \mathbf{d}) \models K_a \phi$.

Assume now that $a \in \mathbf{c}$. Suppose that $(\mathcal{M}, \mathbf{c}) \models K_a \phi$. Take \mathbf{d}' such that $\mathbf{d} \equiv_a \mathbf{d}'$. Then \mathbf{d}' is of the form $\mathbf{d}'_1, \mathbf{c}, \mathbf{c}, \mathbf{d}'_2$. Let $\mathbf{c}' := \mathbf{d}'_1, \mathbf{c}, \mathbf{d}'_2$. By the induction hypothesis $(\mathcal{M}, \mathbf{d}') \models \phi$ iff $(\mathcal{M}, \mathbf{c}') \models \phi$. Further, $\mathbf{d} \equiv_a \mathbf{d}'$ implies that $\mathbf{c} \equiv_a \mathbf{c}'$. So $(\mathcal{M}, \mathbf{c}') \models \phi$. Hence $(\mathcal{M}, \mathbf{d}') \models \phi$ and consequently $(\mathcal{M}, \mathbf{d}) \models K_a \phi$.

The proof in the other direction is analogous. \square

The above result cannot be extended to a repetition of the call sequences. Indeed, we have $(\mathcal{M}, (ab, bc)) \models \neg F_a C$, and $(\mathcal{M}, (ab, bc, ab, bc)) \models F_a C$. On the other hand a monotonicity result holds for positive formulas.

Theorem 4 (Monotonicity). *Suppose that ϕ is a formula that does not contain the \neg symbol. Then*

$$\mathbf{c} \preceq \mathbf{d} \text{ and } (\mathcal{M}, \mathbf{c}) \models \phi \text{ implies } (\mathcal{M}, \mathbf{d}) \models \phi.$$

Proof. We proceed by induction on the structure of ϕ . The only case of interest is when ϕ is of the form $K_a \psi$. Suppose that $\mathbf{c} \preceq \mathbf{d}$ and $(\mathcal{M}, \mathbf{c}) \models \phi$. Take some call sequence \mathbf{d}' such that $\mathbf{d} \equiv_a \mathbf{d}'$. Then for some call sequences \mathbf{d}_1 and \mathbf{d}'_1 such that $\mathbf{d}_1, \mathbf{d}'_1 = \mathbf{d}'$ we have $\mathbf{c} \equiv_a \mathbf{d}_1$.

We have by the assumption $(\mathcal{M}, \mathbf{d}_1) \models \psi$, so by the induction hypothesis $(\mathcal{M}, \mathbf{d}') \models \psi$. As \mathbf{d}' was arbitrarily chosen we conclude that $(\mathcal{M}, \mathbf{d}) \models \phi$. \square

Here and below we say that a call is a *b-call* if agent b is involved in it. Before we deal with the decidability matters consider the formula $K_a F_b C$ for pairwise different agents a, b, c . The following example reveals that it can be true in some subtle ways.

Example 4

- (i) First, note that a can learn (that is, know) that agent b is familiar with the secret C through a direct communication with b .

Indeed, we have $(\mathcal{M}, (bc, ab)) \models K_a F_b C$. Namely the view of agent a of the sequence (bc, ab) is

$$A.B.C \xrightarrow{ab} ABC.ABC.C.$$

So if $(bc, ab) \equiv_a \mathbf{d}$, then \mathbf{d} is of the form $(bc)^+, ab, (bc)^*$, which implies that $(\mathcal{M}, \mathbf{d}) \models F_b C$.

- (ii) Further, it is also possible that a learns that b is familiar with the secret C through a direct communication with c .

Indeed, we have $(\mathcal{M}, (bc, ac)) \models K_a F_b C$. To see this note that the view of agent a of the sequence (bc, ac) is

$$A.B.C \xrightarrow{ac} ABC.B.ABC.$$

So if $(bc, ac) \equiv_a \mathbf{d}$, then \mathbf{d} is of the form $(bc)^+, ac, (bc)^*$, which implies that $(\mathcal{M}, \mathbf{d}) \models F_b C$.

(iii) Also, it is possible that a learns that b is familiar with the secret C without ever communicating with b or c .

Namely, we have $(\mathcal{M}, (cd, ad, bd, ad)) \models K_a F_b C$. Indeed, the view of agent a of the sequence (cd, ad, bd, ad) is

$$A.B.C.D \xrightarrow{ad} ACD.B.C.ACD \xrightarrow{ad} ABCD.B.C.ABCD.$$

So if $(cd, ad, bd, ad) \equiv_a \mathbf{d}$, then \mathbf{d} is of the form $(cd)^+, (bc)^*, ad, \mathbf{d}', ad, \mathbf{d}''$, where in \mathbf{d}' a call bd took place or a call bc followed by a call cd took place, and in \mathbf{d}' and \mathbf{d}'' no a -call took place. This implies that $(\mathcal{M}, \mathbf{d}) \models F_b C$.

(iv) In (iii) agent a learned that b is familiar with c by communicating with agent d twice. But it is also possible that a learns that b is familiar with the secret C without communicating with any agent twice.

To see this note that $(\mathcal{M}, (cd, ad, bc, ac)) \models K_a F_b C$. Indeed, the view of agent a of the sequence (cd, ad, bc, ac) is

$$A.B.C.D \xrightarrow{ad} ACD.B.C.ACD \xrightarrow{ac} ABCD.B.ABCD.ACD.$$

So if $(cd, ad, bc, ac) \equiv_a \mathbf{d}$, then \mathbf{d} is of the form $(cd)^+, ad, \mathbf{d}', ac, \mathbf{d}''$, where in \mathbf{d}' a call bc took place or a call bd followed by a call cd took place, and in \mathbf{d}' and \mathbf{d}'' no a -call took place. This implies that $(\mathcal{M}, \mathbf{d}) \models F_b C$. \square

We conclude by noting that the Monotonicity Theorem 4 does not hold when we extend the call sequences to the left. Indeed, as observed in Example 4 (ii), $(\mathcal{M}, (bc, ac)) \models K_a F_b C$. However, $(\mathcal{M}, (cd, bc, ac)) \models \neg K_a F_b C$, since $(cd, bc, ac) \equiv_a (bd, cd, ac)$ and $(\mathcal{M}, (bd, cd, ac)) \models \neg F_b C$.

5 Decidability of Semantics

In this section we show that the definition of semantics given in Definition 2 is decidable for formulas that do not use nested modalities.

Consider a call sequence \mathbf{c} . If for some prefix $\mathbf{c}_1.c$ of \mathbf{c} , $\mathbf{c}_1(\text{root}) = \mathbf{c}_1.c(\text{root})$, then we say that c is **redundant** in \mathbf{c} . First note the following observation.

Lemma 1 (Semantic Stuttering). *Suppose that $\mathbf{c} := \mathbf{c}_1.c, \mathbf{c}_2$ and $\mathbf{d} := \mathbf{c}_1, \mathbf{c}_2$, where c is redundant in \mathbf{c} . Then for all propositional formulas ϕ , $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.*

Proof. We proceed by induction on the structure of ϕ . The only case of interest is when ϕ is of the form F_ap . The redundancy of \mathbf{c} implies that $\mathbf{c}(\text{root}) = \mathbf{d}(\text{root})$. Hence $(\mathcal{M}, \mathbf{c}) \models F_ap$ iff $p \in \mathbf{c}(\text{root})_a$ iff $p \in \mathbf{d}(\text{root})_a$ iff $(\mathcal{M}, \mathbf{d}) \models F_ap$. \square

The following example shows that Lemma 1 does not extend to arbitrary formulas of \mathcal{L} .

Example 5. In the call sequence ab, ac, bc, ab the second call ab is redundant since $(ab, ac, bc, ab)(\text{root}) = (ab, ac, bc)(\text{root}) = ABC.ABC.ABC$.

However, $(\mathcal{M}, (ab, ac, bc, ab)) \models K_a F_b C$, because if $\mathbf{d} \equiv_a (ab, ac, bc, ab)$ then \mathbf{d} is of the form (ab, ac, bc^+, ab, bc^*) . At the same time, $(\mathcal{M}, (ab, ac, bc)) \models \neg K_a F_b C$ since $(ab, ac, bc) \equiv_a (ab, ac)$. \square

Now, consider an agent a and a call sequence \mathbf{c} . Starting from \mathbf{c} we repeatedly remove from the current call sequence a redundant call that does not involve agent a . We call each outcome of such an iteration an *a -reduction* of \mathbf{c} .

Corollary 1. *Let \mathbf{d} be an a -reduction of \mathbf{c} . Then*

- (i) $\mathbf{c} \equiv_a \mathbf{d}$,
- (ii) for all propositional formulas ϕ , $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.

Proof

- (i) It suffices to note that a removal of a redundant call that does not involve agent a does not affect his view of the call sequence.
- (ii) By the repeated use of the Semantic Stuttering Lemma 1. \square

Given an agent a we now say that a call sequence \mathbf{c} is *a -redundant free* if no call \mathbf{c} from \mathbf{c} such that $a \notin \mathbf{c}$ is redundant in it. Clearly each a -reduction is a -redundant free.

We now prove the following crucial lemma.

Lemma 2. *For each agent a and a call sequence \mathbf{c} the set of a -redundant free call sequences \mathbf{d} such that $\mathbf{c} \equiv_a \mathbf{d}$ is finite.*

Proof. Consider an a -redundant free call sequence \mathbf{d} such that $\mathbf{c} \equiv_a \mathbf{d}$. Then \mathbf{d} has the same number, say k , of a -calls as \mathbf{c} .

Associate with \mathbf{d} the sequence of gossip situations $\mathbf{d}^0(\text{root}), \mathbf{d}^1(\text{root}), \dots, \mathbf{d}^m(\text{root})$, where m is the length of \mathbf{d} , $\mathbf{d}^0 = \epsilon$, and $\mathbf{d}^k = d_1, d_2, \dots, d_k$ for $k = 1, \dots, m$. This sequence monotonically grows, where we interpret the inclusion relation component wise. Moreover, for all calls d_i such that $a \notin d_i$ the corresponding inclusion is strict. Consequently, m , the length of \mathbf{d} , is bounded by $k + |\mathbf{A}|^2$, the sum of the number of a -calls in \mathbf{c} and of the total number of secrets in the gossip situation in which each agent is an expert.

But for each m there are only finitely many call sequences of length at most m . This concludes the proof. \square

We can now state and prove the desired result.

Theorem 5 (Decidability of Semantics). *For each call sequence \mathbf{c} it is decidable whether for a formula ϕ with no nested modalities $(\mathcal{M}, \mathbf{c}) \models \phi$ holds.*

Proof. We use the definition of semantics as the algorithm. We only need to show that the case of the formulas of the form $K_a\phi$, where ϕ is a propositional formula, can be rewritten by referring to a finite set of call sequences \mathbf{d} that can be explicitly constructed. Thanks to the Equivalence Theorem 2 and Corollary 1 we can rewrite the clause for $K_a\phi$ as:

$$(\mathcal{M}, \mathbf{c}) \models K_a\phi \text{ iff } \forall \mathbf{d} \text{ s.t. } \mathbf{c} \sim_a \mathbf{d} \text{ and } \mathbf{d} \text{ is } a\text{-redundant free, } (\mathcal{M}, \mathbf{d}) \models \phi,$$

and according to Lemma 2 this definition indeed refers to an explicitly constructed finite set of call sequences \mathbf{d} . \square

6 Decidability of Truth

Next, we show that truth for formulas that do not use nested modalities is decidable. This implies that the verification problem of gossip protocols, i.e., the problem of determining whether upon protocol's termination every agent is an expert, is decidable for protocols that do not use nested modalities. These include all protocols discussed in [1].

The key notion in our approach is that of an *epistemic view*. It is a function of a call sequence \mathbf{c} , denoted by $EV(\mathbf{c})$, defined by

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

So $EV(\mathbf{c})(a)$ is the set of all gossip situations consistent with agent a 's observations made throughout \mathbf{c} and $EV(\mathbf{c})(*)$ is the actual gossip situation after \mathbf{c} takes place. Note that if $\mathbf{c} \sim_a \mathbf{d}$ then $EV(\mathbf{c})(a) = EV(\mathbf{d})(a)$.

Lemma 3. *For each call sequence \mathbf{c} and agent a the set $EV(\mathbf{c})(a)$ is finite and can be effectively constructed.*

Proof. Fix an agent a . By Corollary 1, Equivalence Theorem 2, and Fact 1(ii) to construct the set $EV(\mathbf{c})(a)$ it suffices to consider a -redundant free call sequences \mathbf{d} and by Lemma 2 there are only finitely many such call sequences \mathbf{d} for which $\mathbf{d} \sim_a \mathbf{c}$. \square

Our interest in epistemic views stems from the following result.

Lemma 4. *Suppose that $EV(\mathbf{c}) = EV(\mathbf{d})$. Then for all epistemic formulas with no nested modalities ϕ , $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.*

Proof. A simple proof by induction shows that for a propositional formula ψ and arbitrary call sequences \mathbf{c}' and \mathbf{d}' , $\mathbf{c}'(\text{root}) = \mathbf{d}'(\text{root})$ implies that $(\mathcal{M}, \mathbf{c}') \models \psi$ iff $(\mathcal{M}, \mathbf{d}') \models \psi$. Since $EV(\mathbf{c})(*) = \mathbf{c}(\text{root})$ and $EV(\mathbf{d})(*) = \mathbf{d}(\text{root})$, this settles the case for $\phi = F_{ap}$.

The above observation also implies that for a propositional formula ψ and an agent a ,

$$(\mathcal{M}, \mathbf{c}) \models K_a \psi \text{ iff } \forall \mathbf{c}' \text{ s.t. } \mathbf{c}'(\text{root}) \in EV(\mathbf{c})(a), (\mathcal{M}, \mathbf{c}') \models \psi.$$

Since $EV(\mathbf{c})(a) = EV(\mathbf{d})(a)$, this settles the case for $\phi = K_a \psi$.

The remaining cases of negation and conjunction follow directly by the induction. \square

The above lemma is useful because the set of epistemic views is finite, in contrast to the set of call sequences. Next, we provide an inductive definition of $EV(\mathbf{c.c})(a)$ the importance of which will become clear in a moment.

Lemma 5. *For any call sequence \mathbf{c} , call c , and agent a such that $a \in c$*

$$EV(\mathbf{c.c})(a) = \{c(s) \mid s \in EV(\mathbf{c})(a) \text{ and } c(s)_a = c(\mathbf{c}(\text{root}))_a\}.$$

Proof. Intuitively the condition $c(s)_a = c(\mathbf{c}(\text{root}))_a$ states that s is consistent with the observation agent a gets after call c is made in the gossip situation $\mathbf{c}(\text{root})$.

(\subseteq) Take $s' \in EV(\mathbf{c.c})(a)$. By the definition of $EV(\mathbf{c.c})(a)$ there exists a call sequence $\mathbf{d.c}$ such that $\mathbf{d.c} \sim_a \mathbf{c.c}$ and $s' = \mathbf{d.c}(\text{root})$. So $s' = c(s)$, where $s = \mathbf{d}(\text{root})$. We also have $\mathbf{d} \sim_a \mathbf{c}$, so $\mathbf{d}(\text{root}) \in EV(\mathbf{c})(a)$. Moreover, $c(\mathbf{d}(\text{root}))_a = c(\mathbf{c}(\text{root}))_a$, because $\mathbf{d.c} \sim_a \mathbf{c.c}$.

(\supseteq) Take $s' \in \{c(s) \mid s \in EV(\mathbf{c})(a) \text{ and } c(s)_a = c(\mathbf{c}(\text{root}))_a\}$. So for some gossip situation s we have $s' = c(s)$, $s \in EV(\mathbf{c})(a)$ and $c(s)_a = c(\mathbf{c}(\text{root}))_a$. The second fact implies that there exists a call sequence \mathbf{d} such that $\mathbf{d} \sim_a \mathbf{c}$ and $s = \mathbf{d}(\text{root})$. Now, this and the third fact imply that $\mathbf{d.c} \sim_a \mathbf{c.c}$. So $\mathbf{d.c}(\text{root}) \in EV(\mathbf{c.c})(a)$. Consequently also $s' \in EV(\mathbf{c.c})(a)$, since $s' = c(s) = \mathbf{d.c}(\text{root})$. \square

This brings us to the following important conclusion stating that $EV(\mathbf{c.c})$ can be computed using $EV(\mathbf{c})$ and c only, i.e., without referring to \mathbf{c} . Denote the set of epistemic views by \widetilde{EV} and recall that C denotes the set of calls.

Corollary 2. *There exists a function $f : \widetilde{EV} \times C \rightarrow \widetilde{EV}$ such that for any call sequence \mathbf{c} and call c*

$$EV(\mathbf{c.c})(a) = f(EV(\mathbf{c}), c).$$

Proof. By the definition of \sim_a we have $EV(\mathbf{c.c})(a) = EV(\mathbf{c})(a)$ if $a \notin c$, $EV(\mathbf{c.c})(*) = c(EV(\mathbf{c})(*))$. This in conjunction with the above lemma implies the claim. \square

Consider a call sequence \mathbf{c} . If for some prefix $\mathbf{c}_1.\mathbf{c}_2$ of \mathbf{c} , we have $EV(\mathbf{c}_1) = EV(\mathbf{c}_1.\mathbf{c}_2)$, 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 $\mathbf{c}_1.\mathbf{c}_2.\dots.\mathbf{c}_k$ is epistemically non-redundant if the set

$$\{\text{EV}(\mathbf{c}_1.\mathbf{c}_2.\dots.\mathbf{c}_i) \mid i \in \{1, \dots, k\}\}$$

has k elements.

We now show a counterpart of the Semantic Stuttering Lemma 1 for epistemic views.

Lemma 6 (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 $\text{EV}(\mathbf{c}) = \text{EV}(\mathbf{d})$.*

Proof. Let $\mathbf{c}_3 = \mathbf{c}_1.\mathbf{c}_2.\dots.\mathbf{c}_k$. First note that thanks to Corollary 2 we have $\text{EV}(\mathbf{c}_1.\mathbf{c}_2.\mathbf{c}_1) = \text{EV}(\mathbf{c}_1.\mathbf{c}_1)$, since $\text{EV}(\mathbf{c}_1.\mathbf{c}_2.\mathbf{c}_1) = f(\text{EV}(\mathbf{c}_1.\mathbf{c}_2), \mathbf{c}_1) = f(\text{EV}(\mathbf{c}_1), \mathbf{c}_1) = \text{EV}(\mathbf{c}_1.\mathbf{c}_1)$ due to the epistemic redundancy of \mathbf{c}_2 in \mathbf{c} . Repeating this argument for all $i \in \{1, \dots, k\}$ we get that $\text{EV}(\mathbf{c}_1.\mathbf{c}_2.\mathbf{c}_1.\mathbf{c}_2.\dots.\mathbf{c}_i) = \text{EV}(\mathbf{c}_1.\mathbf{c}_1.\mathbf{c}_2.\dots.\mathbf{c}_i)$.

In particular $\text{EV}(\mathbf{c}) = \text{EV}(\mathbf{d})$. \square

Corollary 3. *For every call sequence \mathbf{c} there exists an epistemically non-redundant call sequence \mathbf{d} such that for all epistemic formulas with no nested modalities ϕ , $(\mathcal{M}, \mathbf{c}) \models \phi$ iff $(\mathcal{M}, \mathbf{d}) \models \phi$.*

Proof. By the repeated use of the Epistemic Stuttering Lemmas 4 and 6. \square

Next, we prove the following crucial lemma.

Lemma 7. *For any given model \mathcal{M} , there are only finitely many epistemically non-redundant call sequences.*

Proof. Note that each epistemic view is a function from $A \cup \{*\}$ to the set of functions from A to $2^{|P|}$ (this is an overestimation because for $*$ this set has only one element). There are $k = 2^{(|A|+1) \cdot 2^{|A| \cdot |P|}}$ such functions, so any call sequence longer than k has an epistemically redundant call subsequence. But there are only finitely many call sequences of length at most k . This concludes the proof. \square

Finally, we can establish the announced result.

Theorem 6 (Decidability of Truth). *For any formula ϕ with no nested modalities, it is decidable whether $\mathcal{M} \models \phi$ holds.*

Proof. Recall that $\mathcal{M} \models \phi$ iff $\forall \mathbf{c} (\mathcal{M}, \mathbf{c}) \models \phi$. Thanks to Corollary 3 we can rewrite the latter as

$$\forall \mathbf{c} \text{ s.t. } \mathbf{c} \text{ is epistemically non-redundant, } (\mathcal{M}, \mathbf{c}) \models \phi.$$

But according to Lemma 7 there are only finitely many epistemically non-redundant call sequences and by Lemma 3 their set can be explicitly constructed. \square

As an easy consequence we obtain the following.

Corollary 4. *It is decidable to determine whether a given gossip situation can be an outcome of a call sequence.*

Proof. Each gossip situation $\mathbf{s} = (Q_d)_{d \in \mathbf{A}}$ can be encoded as a conjunction

$$\phi(\mathbf{s}) = \bigwedge_{a \in \mathbf{A}} \left(\bigwedge_{B \in Q_a} F_a B \wedge \bigwedge_{B \notin Q_a} \neg F_a B \right).$$

Then $\exists \mathbf{c}(\mathbf{c}(\text{root}) = \mathbf{s})$ iff $\exists \mathbf{c}((\mathcal{M}, \mathbf{c}) \models \phi(\mathbf{s}))$ iff $\neg(\mathcal{M} \models \neg\phi(\mathbf{s}))$. \square

7 Decidability of Termination

Finally, we show that it is decidable to determine whether a gossip protocol terminates. First, we establish monotonicity of gossip situations and epistemic views with respect to call sequence extensions. Intuitively, we claim that as the call sequence gets longer each agent acquires more information. This can be seen as a counterpart of the Monotonicity Theorem 4. First we need to define suitable partial orderings \leq_s and \leq_{ev} over gossip situations and epistemic views, respectively.

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

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

Proof. For any gossip situation \mathbf{s} and call \mathbf{c} we have by definition $\mathbf{s} \leq_s \mathbf{c}(\mathbf{s})$. By induction this implies that for any call sequence \mathbf{c}' we have $\mathbf{s} \leq_s \mathbf{c}'(\mathbf{s})$. Now $\mathbf{c} \preceq \mathbf{d}$ implies that $\mathbf{d} = \mathbf{c}.\mathbf{c}'$ for some \mathbf{c}' . Therefore, $\mathbf{c}(\text{root}) \leq_s \mathbf{c}'(\mathbf{c}(\text{root})) = \mathbf{d}(\text{root})$. \square

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

Lemma 8. \leq_{ev} is a partial order.

Proof. Omitted. \square

The next lemma formalizes the intuition that epistemic information grows along a call sequence.

Lemma 9. *For all two call sequences such that $\mathbf{c} \preceq \mathbf{d}$ we have $\mathbf{EV}(\mathbf{c}) \leq_{ev} \mathbf{EV}(\mathbf{d})$.*

Proof. Let $\mathbf{d} = \mathbf{c}.\mathbf{c}'$. Take $a \in \mathbf{A}$. Note that by a repeated application of Lemma 5 we can show that $\mathbf{EV}(\mathbf{c}.\mathbf{c}')(a) = \{\mathbf{c}'(\mathbf{s}) \mid \mathbf{s} \in \mathbf{EV}(\mathbf{c})(a) \text{ and } \forall \mathbf{c}'' \preceq \mathbf{c}' \mathbf{c}''(\mathbf{s})_a = \mathbf{c}''(\mathbf{c}(\text{root}))_a\}$. It suffices then to pick $X = \{\mathbf{s} \in \mathbf{EV}(\mathbf{c})(a) \mid \forall \mathbf{c}'' \preceq \mathbf{c}' \mathbf{c}''(\mathbf{s})_a = \mathbf{c}''(\mathbf{c}(\text{root}))_a\}$, and set $g(\mathbf{s}) = \mathbf{c}'(\mathbf{s})$ for all $\mathbf{s} \in X$. It is easy to check that such $g : X \rightarrow \mathbf{EV}(\mathbf{d})$ is surjective, so $\mathbf{EV}(\mathbf{c}) \leq_{ev} \mathbf{EV}(\mathbf{d})$. \square

We can now draw the following useful conclusion.

Lemma 10. *Suppose that \mathbf{c} is epistemically redundant. Then a prefix $\mathbf{c}_1.\mathbf{c}$ of it exists such that \mathbf{c}_1 is epistemically non-redundant and $\text{EV}(\mathbf{c}_1.\mathbf{c}) = \text{EV}(\mathbf{c}_1)$.*

Proof. Let $\mathbf{c}_1.\mathbf{c}_2$ be the shortest prefix of \mathbf{c} such that $\text{EV}(\mathbf{c}_1) = \text{EV}(\mathbf{c}_1.\mathbf{c}_2)$. Then \mathbf{c}_1 is epistemically non-redundant. Let $\mathbf{c}_2 = \mathbf{c}_1 \dots \mathbf{c}_l$. By Lemma 9 we have $\text{EV}(\mathbf{c}_1) \leq_s \text{EV}(\mathbf{c}_1.\mathbf{c}_1) \leq_s \text{EV}(\mathbf{c}_1.\mathbf{c}_1.\mathbf{c}_2) \leq_s \dots \leq_s \text{EV}(\mathbf{c}_1.\mathbf{c}_1.\mathbf{c}_2 \dots \mathbf{c}_l) = \text{EV}(\mathbf{c}_1.\mathbf{c}_2) = \text{EV}(\mathbf{c}_1)$. Since \leq_s is a partial order, $\text{EV}(\mathbf{c}_1.\mathbf{c}_1) = \text{EV}(\mathbf{c}_1)$ holds. \square

Finally we can establish the desired result.

Theorem 7 (Decidability of Termination). *Given a gossip protocol it is decidable to determine whether it always terminates.*

Proof. We first prove that a gossip protocol may fail to terminate iff it can generate a call sequence $\mathbf{c}.\mathbf{c}$ such that \mathbf{c} is epistemically non-redundant and $\text{EV}(\mathbf{c}.\mathbf{c}) = \text{EV}(\mathbf{c})$.

(\Rightarrow) Let $\bar{\mathbf{c}}$ be an infinite sequence of calls generated by the protocol. There are only finitely many epistemic views, so some prefix \mathbf{c} of $\bar{\mathbf{c}}$ is epistemically redundant. The claim now follows by Lemma 10.

(\Leftarrow) Suppose that the protocol generates a sequence of calls $\mathbf{c}.\mathbf{c}$ such that \mathbf{c} is epistemically non-redundant and $\text{EV}(\mathbf{c}.\mathbf{c}) = \text{EV}(\mathbf{c})$.

Let ϕ be the guard associated with the call \mathbf{c} . By assumption $(\mathcal{M}, \mathbf{c}) \models \phi$. By the assumption about the gossip protocols the formula ϕ is without nested modalities, so by Lemma 4 $(\mathcal{M}, \mathbf{c}.\mathbf{c}) \models \phi$. Hence by the repeated use of the Stuttering Theorem 3, for all $i \geq 1$, $(\mathcal{M}, \mathbf{c}.\mathbf{c}^i) \models \phi$. Consequently, $\mathbf{c}.\mathbf{c}^\omega$ is an infinite sequence of calls that can be generated by the protocol.

The above equivalence shows that determining whether the protocol always terminates is equivalent to checking that it cannot generate a call sequence $\mathbf{c}.\mathbf{c}$ such that \mathbf{c} is epistemically non-redundant and $\text{EV}(\mathbf{c}.\mathbf{c}) = \text{EV}(\mathbf{c})$.

But given a call sequence, by the Decidability of Semantics Theorem 5, it is decidable to determine whether it can be generated by the protocol and by Lemma 3 it is decidable to determine whether a call sequence is epistemically non-redundant. Further, by Lemma 7 there are only finitely many epistemically non-redundant call sequences, so the claim follows. \square

8 Conclusions

In this paper we studied decidability questions concerning a natural epistemic logic appropriate for expressing gossip protocols. One of our aims was to show that the gossip protocols considered in [1] are executable. A self-contained summary is that the semantics of the introduced epistemic language \mathcal{L} is decidable for formulas with no nested modalities. Another aim was to prove that partial correctness of the gossip protocols studied in [1] is decidable. To this end we showed that truth of formulas of \mathcal{L} with no nested modalities is decidable. This implies the former

since partial correctness of such a gossip protocol means that a specific epistemic formula, namely the conjunction of the negation of all guards implies that each agent is an expert, is true and such a formula has no nested modalities. Finally, we showed the problem of determining termination of a gossip protocol is decidable. An interesting open question is whether all of these results can be extended to arbitrary formulas of the language \mathcal{L} . The main stumbling block in generalizing our proofs is that, as Example 5 shows, the crucial Semantic Stuttering Lemma 1 cannot be extended to arbitrary formulas of \mathcal{L} .

These considerations lead to another interesting open problem. Gossip protocols studied in [1] are parametric in the sense that they are formulated in such a way that they do not depend on the underlying graph (for instance a ring). The results we proved allow us only to consider each specific gossip protocol (for example for a ring formed by 5 agents) separately. What is needed is a decision procedure that would allow us to consider all instances of a protocol (for example for all rings) simultaneously. We conjecture that this decision problem is undecidable both for partial correctness and for termination.

The semantics we introduced in Sect. 3 stipulates through the definition of $\mathbf{c}(s)$ that a call ab is not noted by any agent $c \neq a, b$. In [3] different type of calls were studied, namely

- ab^- , which stipulates that every agent $c \neq a, b$ noted that a called b ,
- ab^0 , which stipulates that every agent $c \neq a, b$ noted that some call took place, though not between whom,
- ab^+ which stipulates that every agent $c \neq a, b$ noted that possibly some call took place, though not between whom.

It would be interesting to check whether our results hold for these types of calls, as well.

Another issue interesting to study is the synthesis of a distributed epistemic gossip protocol from epistemic specifications. For a related work on a synthesis of a knowledge-based programs see, e.g. [12]. Finally, it would be interesting to study the decidability of the problems considered here for a variant of our logic in which the only modal operator is the common knowledge operator $C_G\phi$. This operator states that the formula ϕ is commonly known among the group of agents G . The standard semantics of this operator is given in [5].

Acknowledgments. We thank the reviewers for helpful comments. First author is also a Visiting Professor at the University of Warsaw. He was partially supported by the NCN grant nr 2014/13/B/ST6/01807. The second author was partially supported by EPSRC grant EP/M027287/1.

References

1. Apt, K.R., Grossi, D., Van der Hoek, W.: Epistemic protocols for distributed gossiping. In: Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015). EPTCS, vol. 215, pp. 51–66 (2016)
2. Attamah, M., Ditmarsch, H., Grossi, D., Hoek, W.: A framework for epistemic gossip protocols. In: Bulling, N. (ed.) EUMAS 2014. LNCS (LNAI), vol. 8953, pp. 193–209. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-17130-2_13](https://doi.org/10.1007/978-3-319-17130-2_13)
3. Attamah, M., van Ditmarsch, H., Grossi, D., Van der Hoek, W.: Knowledge and gossip. In Proceedings of ECAI 2014. IOS Press (2014)
4. Cooper, M.C., Herzig, A., Maffre, F., Maris, F., Régnier, P.: A simple account of multiagent epistemic planning. In: Proceedings of ECAI 2016, pp. 193–201. IOS Press (2016)
5. Fagin, R., Halpern, J., Vardi, M., Moses, Y.: Reasoning About Knowledge. MIT Press, Cambridge (1995)
6. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Knowledge-based programs. Distrib. Comput. **10**(4), 199–225 (1997)
7. Halpern, J.Y., Zuck, L.D.: A little knowledge goes a long way: knowledge-based derivations and correctness proofs for a family of protocols. J. ACM **39**(3), 449–478 (1992)
8. Hedetniemi, S.M., Hedetniemi, S.T., Liestman, A.L.: A survey of gossiping and broadcasting in communication networks. Networks **18**(4), 319–349 (1988)
9. Herzig, A., Maffre, F.: How to share knowledge by gossiping. In: Rovatsos, M., Vouros, G., Julian, V. (eds.) EUMAS/AT -2015. LNCS (LNAI), vol. 9571, pp. 249–263. Springer, Heidelberg (2016). doi:[10.1007/978-3-319-33509-4_20](https://doi.org/10.1007/978-3-319-33509-4_20)
10. Hromkovic, J., Klasing, R., Pelc, A., Ruzicka, P., Unger, W.: Dissemination of Information in Communication Networks - Broadcasting, Gossiping, Leader Election, and Fault-Tolerance. Texts in Theoretical Computer Science. An EATCS Series. Springer, New York (2005)
11. Kermarrec, A., van Steen, M.: Gossiping in distributed systems. Oper. Syst. Rev. **41**(5), 2–7 (2007)
12. Meyden, R., Wilke, T.: Synthesis of distributed systems from knowledge-based specifications. In: Abadi, M., Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 562–576. Springer, Heidelberg (2005). doi:[10.1007/11539452_42](https://doi.org/10.1007/11539452_42)
13. van Ditmarsch, H., van Eijck, J., Pardo, P., Ramezani, R., Schwarzentruher, F.: Dynamic gossip (2015). CoRR, abs/1511.00867

Logics in Artificial Intelligence

15th European Conference, JELIA 2016, Larnaca,

Cyprus, November 9-11, 2016, Proceedings

Michael, L.; Kakas, A. (Eds.)

2016, XXII, 586 p. 70 illus., Softcover

ISBN: 978-3-319-48757-1