

Open Problems in a Logic of Gossips

Krzysztof R. Apt

Centrum Wiskunde & Informatica
Amsterdam, The Netherlands

University of Warsaw
Warsaw, Poland

k.r.apt@cwi.nl

Dominik Wojtczak

University of Liverpool
Liverpool, UK

d.wojtczak@liv.ac.uk

Gossip protocols are programs used in a setting in which each agent holds a secret and the aim is to reach a situation in which all agents know all secrets. Such protocols rely on a point-to-point or group communication. Distributed epistemic gossip protocols use epistemic formulas in the component programs for the agents. The advantage of the use of epistemic logic is that the resulting protocols are very concise and amenable for a simple verification.

Recently, we introduced a natural modal logic that allows one to express distributed epistemic gossip protocols and to reason about their correctness. We proved that the resulting protocols are implementable and that all aspects of their correctness, including termination, are decidable. To establish these results we showed that both the definition of semantics and of truth of the underlying logic are decidable. We also showed that the analogous results hold for an extension of this logic with the ‘common knowledge’ operator.

However, several, often deceptively simple, questions about this logic and the corresponding gossip protocols remain open. The purpose of this paper is to list and elucidate these questions and provide for them an appropriate background information in the form of partial of related results.

1 Introduction

Gossip protocols concern a set up in which each agent holds initially a secret and the aim is to arrive, by means of point-to-point or group communications (called *calls*), at a situation in which all agents know each other secrets. During the calls the agents exchange some, possibly all, secrets they know.

These protocols were successfully used in a number of domains, for instance communication networks [20], computation of aggregate information [25], and data replication [27]. For a more recent account see [24] and [26].

In [10] a dynamic epistemic logic was introduced in which gossip protocols could be expressed as formulas. These protocols rely on agents’ knowledge and are distributed, so they are distributed epistemic gossip protocols. This means that they can be seen as special cases of knowledge-based programs introduced in [18].

In [3] a simpler modal logic was introduced that is sufficient to define these protocols and to reason about their correctness. This logic is interesting in its own rights and was subsequently studied in a number of papers. In particular, in [5], and in the full version in [8], we established decidability of its semantics and truth for a limited fragment. Building upon these results we then proved that the distributed gossip protocols, the guards of which are defined in this logic, are implementable, that their partial correctness is decidable, and that termination and two forms of fair termination of these protocols are decidable, as well. Further, in [4] the computational complexity of this fragment was studied and in [6] we considered its extension with the common knowledge operator for which we established analogous decidability results.

In spite of the simplicity of this modal logic several natural questions about it and the gossip protocols defined using this logic remain open. In what follows we discuss these problems. For each of them we provide the relevant background information and establish some partial results.

Among these partial results let us mention the following ones:

- When the agents form a star graph, each correct distributed epistemic gossip protocol in the framework of [3] has to rely on guards with the modal operators (Theorem 2 in Section 7). This is relevant since the complexity of determining truth of guards is higher in presence of modal operators (see Section 4).
- It is well known (see, e.g., [31] discussed in Section 7) that for 4 agents 4 calls are both needed and sufficient to reach a situation in which all agents know all secrets. The resulting protocol is centralized. We show that when distributed epistemic gossip protocols are used, for 4 agents 5 calls are both needed and sufficient (Theorems 6 and 7 in Section 7).
- In the literature on distributed computing, e.g., on Calculus of Communicating systems (CCS) of [28], there is a wealth of literature on various ways of comparing behaviour of two processes (see, e.g., [29] for an extensive overview of the fundamental concept of bisimulation). Distributed epistemic gossip protocols can be naturally compared by means of a notion that one of them can simulate another. We show that checking it can be done in exponential time (Theorem 8 in Section 8).

Further, the arguments used to prove the last result imply that all computations of a terminating gossip protocol are of length $< n^4$.

The paper is organized as follows. In the next section we discuss related work. Then, in Section 3, introduce the already mentioned logic, originally defined in [3], and in Section 4 discuss some natural open problems about it. In Section 5 we recall an extension of this logic with the common knowledge operator and introduce two open problems concerning it. Next, in Section 6, we recall the distributed epistemic gossip protocols considered in [3]. Then, in Sections 7 and 8 we discuss natural open problems about these protocols. For several open problems we provide some partial results.

2 Related work

As already mentioned, distributed epistemic gossip protocols were introduced in [10]. In [9] a tool was presented that given a high level description of an epistemic protocol in this setting generates the characteristics of the protocol. In both papers three types of calls were considered but the ones considered here (initially studied in [3]) differ from them in that we assume that agents not participating in the call are not aware of it. In [3] also two other modes of communication were considered, in which only one agent (the caller or the called one) learns new secrets. The assumptions about the calls used in [10] and [3] were presented in a uniform framework in [2], where in total 18 types of communication were introduced and compared w.r.t. their epistemic strength.

In [21] and [22] centralized gossip protocols were studied the aim of which is to achieve higher-order shared knowledge, for example knowledge of level 2 which stipulates that everybody knows that everybody knows all secrets. In particular, a protocol was presented and proved correct that achieves in $(k+1)(n-2)$ steps shared knowledge of level k . These matters were further investigated in [12], where optimal protocols for various versions of such a generalized gossip problem were presented. These protocols depend on various parameters, for example type of the underlying graph or the type of communication. Further, different gossip problems were also studied in which some negative goals, for example that certain agents

must not know certain secrets, are supposed to be achieved. In [13] such problem were studied further in the presence of temporal constraints, i.e., a given call can only (or has to) be made within a given time interval.

Then in [11] gossip protocols were analyzed as an instance of multi-agent epistemic planning that was subsequently translated into a planning language.

The underlying framework was analyzed from a number of views. In [16] gossip problems were considered in an epistemic framework that provides several parameters allowing us to capture various aspects of it, for example the initial knowledge of the agents, the type of communication used, and the desired type of the protocol (for example, a symmetric one). For some of the combinations of the parameters the minimum number of calls needed to reach the final situation was established. The expected time of termination of several gossip protocols for complete graphs was studied by [17].

Next, in [14] dynamic distributed gossip protocols were studied in which the calls allow the agents not only to share the secrets but also to transmit the links. These protocols were characterized in terms of the class of graphs for which they terminate. They differ from the ones here considered, which are static. This set up was further investigated in [15] where various dynamic gossip protocols were proposed and analyzed. In [19] such protocols were analyzed by embedding them in a network programming language NetKAT proposed in [1].

3 Logic: a recall

We recall here the framework of [3]. We assume a fixed set A of $n \geq 3$ *agents* and stipulate that each agent holds exactly one *secret*, and that there exists a bijection between the set of agents and the set of secrets. We use it implicitly by denoting the secret of agent a by A , of agent b by B , etc. We denote by Sec the set of all secrets.

The language \mathcal{L} is defined by the following grammar:

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

where $S \in \text{Sec}$ and $a \in A$.

So $F_a S$ is an atomic formula, while $K_a \phi$ is a compound formula. We read $F_a S$ as ‘agent a is familiar with the secret S ’ (or ‘agent a holds secret S ’) and $K_a \phi$ as ‘agent a knows the formula ϕ ’. Below we shall freely use other Boolean connectives that can be defined using \neg and \wedge in a standard way.

In the sequel we shall use the following formula

$$\text{Exp}_a \equiv \bigwedge_{S \in \text{Sec}} F_a S,$$

that denotes the fact that agent a is familiar with all the secrets (is an ‘expert’).

In the paper we shall use the following sublanguages of \mathcal{L} :

- \mathcal{L}_0 , its propositional part, which consists of the formulas that do not use the K_a modalities;
- \mathcal{L}_1 , which consists of the formulas without the nested use of the K_a modalities;
- \mathcal{L}_1^a , where $a \in A$ is a fixed agent, which consists of the formulas from \mathcal{L}_1 where the only modality is K_a .

Each *call*, written as ab , concerns two different agents, the *caller*, a , and the *callee*, b . After the call the caller and the callee learn each others secrets. 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 .

Following [3] we stipulate that agents not involved in the call are not aware of it. This will be addressed in Definition 1 below.

In what follows we focus on call sequences. Unless explicitly stated each call sequence is assumed to be finite. The empty sequence is denoted by ε . We use c to denote a call sequence and C to denote the set of all finite call sequences. Given call sequences c and d and a call c we denote by $c.c$ the outcome of adding c at the end of the sequence c and by $c.d$ the outcome of appending the sequences c and d . We say that c_2 is a **subsequence** of a call sequence c if for some call sequences c_1 and c_3 we have $c = c_1.c_2.c_3$.

To describe what secrets the agents are familiar with, we use the concept of a **gossip situation**. It is a sequence $s = (Q_a)_{a \in A}$, where $\{A\} \subseteq Q_a \subseteq \text{Sec}$ for each agent a . Intuitively, Q_a is the set of secrets a is familiar with in the gossip situation s . The **initial gossip situation** is the one in which each Q_a equals $\{A\}$ and is denoted by root . It reflects the fact that initially each agent a is familiar only with his own secret, A . We say that an agent a is an **expert** in a gossip situation s if a is familiar in s with all the secrets, i.e., if $Q_a = \text{Sec}$.

Each call transforms the current gossip situation by modifying the sets of secrets the agents involved in the call are familiar with as follows. Consider a gossip situation $s := (Q_d)_{d \in A}$ and a call ab .

Then

$$ab(s) := (Q'_d)_{d \in A},$$

where $Q'_a = Q'_b = Q_a \cup Q_b$, and for $c \notin \{a, b\}$, $Q'_c = Q_c$.

So the effect of a call is that the caller and the callee share the secrets they are familiar with.

The result of applying a call sequence to a gossip situation s is defined inductively as follows:

$$\varepsilon(s) := s, (c.c)(s) := c(c(s)).$$

Example 1 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, a, b and c , then $\text{root} = A.B.C$ and the gossip situation $(\{A, B\}, \{A, B\}, \{C\})$ will be written as $AB.AB.C$.

Let $A = \{a, b, c\}$. 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 $(ac, bc, ac)(\text{root}) = (ABC.ABC.ABC)$. □

As calls progress in sequence from the initial situation, agents may be uncertain about which call sequence took place. This uncertainty is captured by the appropriate equivalence relations on the call sequences.

Definition 1 Fix an agent a . We define $\sim_a \subseteq C \times C$ as the smallest equivalence relation satisfying the following conditions:

[Base] $\varepsilon \sim_a \varepsilon$,

[Step] Suppose $c \sim_a d$.

(i) If $a \notin c$, then $c.c \sim_a d$ and $c \sim_a d.c$.

(ii) If there exists $b \in A$ and $c \in \{ab, ba\}$ such that $c.c(\text{root})_a = d.c(\text{root})_a$, then $c.c \sim_a d.c$.

In (i) we formalize the assumption that the agents are not aware of the calls they do not participate in. In turn, in (ii) we capture the intuition that two call sequences are indistinguishable for an agent if the sets of his calls in both sequences are the same and in each sequence he observes the same set of secrets. For instance, by (i) we have $ab, bc \sim_a ab, bd$. But we do not have $bc, ab \sim_a bd, ab$ since $C \in (bc, ab)(\text{root})_a$, while $C \notin (bd, ab)(\text{root})_a$.

Next, we recall the definition of truth.

Definition 2 Consider a call sequence $c \in C$. We define the satisfaction relation \models inductively as follows (clauses for Boolean connectives are as usual and omitted):

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

Further, we say that ϕ is **true**, and write $\models \phi$, when $\forall c c \models \phi$. Also, we say that ϕ and ψ are **equivalent** if $\forall c c \models \phi \leftrightarrow \psi$.

So a formula $F_a S$ is true after the call sequence c whenever secret S belongs to the set of secrets agent a is familiar with in the situation generated by the call sequence c applied to the initial situation root . Hence $c \models \text{Exp}_a$ iff agent a is an expert in $c(\text{root})$.

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

4 Open problems about the logic

The first problem concerns the propositional part \mathcal{L}_0 of the language \mathcal{L} . In [4] we established that the problem of determining the complexity of the semantics of the language of \mathcal{L}_0 (i.e, determining whether $c \models \phi$ for a given call sequence c) is in P, while the problem of determining truth in the language of \mathcal{L}_0 (i.e, determining whether $c \models \phi$ for all call sequences c) is CO-NP-complete. Note that the latter implies that checking whether there exists a call sequence c such that $c \models \phi$ is NP-complete.

Problem 1 Find an axiomatization of \mathcal{L}_0 .

Comments

We stipulate that the following formula should be an axiom:

$$\bigwedge_{\substack{a,b \in A \\ a \neq b}} (F_a B \rightarrow \bigvee_{k=2}^n \bigvee_{\substack{a_1, \dots, a_k \in A \\ a_1 = b, a_k = a \\ a_i \neq a_j \text{ for all } i \neq j}} \bigwedge_{h=1}^{k-1} (F_{a_h} A_{h+1} \wedge F_{a_{h+1}} A_h)). \quad (1)$$

It formalizes the fact that communication is bidirectional and that agents learn secrets through a chain of calls, and is easily seen to be true. Here and elsewhere we make use of the fact that there is a bijection between secrets and agents, which allows us to use agents (here b, a_h and a_{h+1}) when referring to the corresponding secrets (here B, A_h and A_{h+1}). Note that the second disjunction is finite, since we postulate that the agents a_1, \dots, a_k are pairwise different.

Here is the version of (1) when A has three agents that is more readable:

$$\bigwedge_{\substack{a,b,c \in A \\ \{a,b,c\} = A}} (F_a B \rightarrow (F_b A \wedge F_a B) \vee (F_b C \wedge F_c B \wedge F_c A \wedge F_a C)).$$

Assuming a sound and complete proof system for Boolean formulas formula (1) implies the following natural formula of \mathcal{L}_0 stating that each agent can learn a new secret only by revealing its own:

$$\bigwedge_{\substack{a,b \in A \\ a \neq b}} (F_a B \rightarrow \bigvee_{\substack{c,a \in A \\ c \neq a}} F_c A). \quad (2)$$

To see the claim it suffices to note that (1) implies

$$\bigwedge_{\substack{a,b \in A \\ a \neq b}} (F_a B \rightarrow \bigvee_{k=2}^n \bigvee_{\substack{a_1, \dots, a_k \in A \\ a_1 = b, a_k = a \\ a_i \neq a_j \text{ for all } i \neq j}} F_{a_{k-1}} A),$$

from which (2) follows.

In [2] it was observed that the following formula is true:

$$\bigwedge_{\substack{a,b \in A \\ a \neq b}} \left(F_a B \wedge \bigwedge_{\substack{i \in A \\ i \neq a, i \neq b}} \neg F_i B \right) \rightarrow F_b A. \quad (3)$$

It states that if agent a is the only agent (different from b) familiar with the secret of b , then agent b is familiar with the secret of a .

It is easy to see that assuming a sound and complete proof system for Boolean formulas (1) implies (3). To this end it suffices to note that (1) implies

$$\bigwedge_{\substack{a,b \in A \\ a \neq b}} (F_a B \rightarrow F_b A \vee \bigvee_{k=3}^n \bigvee_{\substack{a_1, \dots, a_k \in A \\ a_1 = b, a_k = a \\ a_i \neq a_j \text{ for all } i \neq j}} F_{a_2} B),$$

from which (3) follows.

A more ambitious problem is the following one.

Problem 2 Find an axiomatization of \mathcal{L} . Determine the complexity of the semantics and of truth in the language of \mathcal{L} .

Comments

In [5] and in the full version in [8], we showed that for the sublanguage \mathcal{L}_1 of \mathcal{L} both the semantics and truth are decidable. In [4] we sharpened these results by showing that the first problem is $\text{P}_{\parallel}^{\text{NP}}$ -complete, while the complexity of determining the truth is in coNP^{NP} . It is not clear how to extend these results to larger fragments of \mathcal{L} . Recall that the complexity class $\text{P}_{\parallel}^{\text{NP}}$, defined in [32], corresponds to the class of problems solvable by a deterministic polynomial-time Turing machine that has a parallel access to an NP oracle (i.e., no query to this oracle can depend on the outcome of any other). In [32, 30] many natural problems were shown to be complete for this class. One of them is checking for two Boolean formulas ϕ, ϕ' whether the maximum number of variables assigned **true** in a satisfying assignment is greater for ϕ than for ϕ' . On the other hand, the complexity class coNP^{NP} , commonly denoted by Π_2^{P} , corresponds to a non-deterministic polynomial-time Turing machines with an access to an NP oracle that accepts a given input if and only if all its non-deterministic branches accept it. An example problem complete for this class is the satisfiability for quantified Boolean formulas with two alternations of quantifiers $\forall x_1, \dots, x_k \exists y_1, \dots, y_l \phi$, where ϕ is a Boolean formula over the variables $x_1, \dots, x_k, y_1, \dots, y_l$.

5 Open problems concerning common knowledge

In [6] an extension of the language \mathcal{L} was considered that involves the common knowledge operator. The resulting language \mathcal{L}_{ck} is defined by the following grammar:

$$\phi ::= F_a S \mid \neg \phi \mid \phi \wedge \phi \mid C_G \phi,$$

where $S \in \text{Sec}$ and $a \in A$ and $G \subseteq A$.

Recall that the semantics of the C_G operator is defined as follows. Given a set A let A^* be the set of all finite sequences formed from the elements of A . For $t = a_1, \dots, a_k$ let $K_t = K_{a_1} \dots K_{a_k}$. Then we stipulate that

$$C_G \phi \equiv \bigwedge_{t \in G^*} K_t \phi.$$

Note that the formula on the right-hand side is an infinite conjunction.

When G is a singleton, say $G = \{a\}$, the formula $C_G \phi$ has the same semantics as $K_a \phi$, since $K_a K_a \phi$ is equivalent to $K_a \phi$. So the language \mathcal{L}_{ck} can be viewed as an extension of \mathcal{L} .

Problem 3 *Determine whether common knowledge is equivalent to a nested knowledge. More precisely, determine whether for every G and ϕ there exists $t \in G^*$ such that the formulas $C_G \phi$ and $K_t \phi$ are equivalent.*

If the answer to Problem 3 is positive, then it is natural to ask whether $t \in G^*$ can be found independently of the formula ϕ , that is, whether for some $t \in G^*$ the equivalence

$$C_G \phi \leftrightarrow K_t \phi$$

holds for all formulas ϕ in \mathcal{L}_{ck} or in some fragment of it.

Problem 4 *Find an axiomatization of \mathcal{L}_{ck} . Determine the complexity of the semantics and of truth in the language of \mathcal{L}_{ck} .*

Comments

We do not know the answer to Problem 3 even for the formulas of the form $C_G \phi$, where $\phi \in \mathcal{L}_0$. In [7] we succeeded to establish the following limited results.

THEOREM

- (i) *Suppose that $|G| \geq 3$. Then for all call sequences c and formulas $\phi \in \mathcal{L}_{ck}$*

$$c \models C_G \phi \text{ iff } c \models \phi.$$

Consequently, the formulas $C_G \phi$ and ϕ are equivalent.

- (ii) *For all call sequences c that do not contain the call ab and all formulas $\phi \in \mathcal{L}_{ck}$ that do not contain the \neg symbol*

$$c \models C_{\{a,b\}} \phi \text{ iff } c \models K_{abab} \phi.$$

Part (i) states that the formulas commonly known by a group of at least three agents are precisely the true formulas. Part (ii) states that for a group of two agents common knowledge of negation-free formulas is equivalent to the 4th fold iterated knowledge w.r.t. a call sequence in which no calls between these two agents were made. Examples in [7] show that this claim does not hold for arbitrary formulas and that the restriction on the call sequence cannot be dropped.

In [7] we showed that both the semantics and truth in the sublanguage of \mathcal{L}_{ck} that consists of the formulas with no nested C_G modalities are decidable.

6 Distributed epistemic gossip protocols: a recall

In [3], as a follow up on [10], 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 $\bigwedge_{a \in A} (F_a A \wedge \bigwedge_{b \in A, b \neq a} \neg F_a B)$ is true into one in which the formula $\bigwedge_{a, b \in A} F_a B$ is true. As explained in the introduction, in [3] a different syntax of the gossip protocols than in [10] was used.

Let us recall the definition. The adopted syntax follows the syntax of the CSP language (Communicating Sequential Processes) of [23], in which ‘*’ denotes a repetition and ‘[]’ a nondeterministic choice. By a **component program**, in short a **program**, for an agent a we mean a statement of the form

$$*[\big[\big]_{j=1}^m \psi_j \rightarrow c_j],$$

where $m \geq 0$ and each $\psi_j \rightarrow c_j$ is such that

- a is the caller in the call c_j ,
- $\psi_j \in \mathcal{L}_1^a$ and all atomic formulas used in ψ start with F_a .

If $m = 0$, the program is empty.

We call each such construct $\psi \rightarrow c$ a **rule** and refer in this context to ψ as a **guard**. Intuitively, * denotes a repeated execution of the rules, one at a time, where each time non-deterministically a rule is selected whose guard is true.

By a **distributed epistemic gossip protocol**, from now on just a **gossip protocol**, we mean a parallel composition of component programs, one for each agent. We call a gossip protocol **propositional** if all guards in it are propositional, i.e., are from the language \mathcal{L}_p .

We presuppose that in each gossip protocol the agents are the nodes of a directed graph (digraph) and that each call ab is allowed only if $a \rightarrow b$ is an edge in the digraph. A minimal digraph that satisfies this assumption is uniquely determined by the syntax of the protocol. Given that the aim of each gossip protocol is that all agents become experts it is natural to assume that this digraph is connected.

Here are two examples of gossip protocols to which we shall return later.

Example 2 In [10] the following correct gossip protocol, called Learn New Secrets (LNS in short), for complete graphs was proposed. In the syntax of [3] used here it is propositional, as it has the following program for agent i :

$$*[\big[\big]_{j \in A} \neg F_i J \rightarrow ij].$$

Informally, agent i calls agent j if agent i is not familiar with j 's secret. □

Example 3 In [10] also the following correct gossip protocol, called Hear My Secret (HMS in short), for complete graphs was proposed. In the syntax of [3] it has the following program for agent i :

$$*[\big[\big]_{j \in A} \neg K_i F_j I \rightarrow ij].$$

Informally, agent i calls agent j if agent i does not know whether j is familiar with his secret. □

Consider a gossip protocol P that is a parallel composition of the component programs $*[\big[\big]_{j=1}^{m_a} \psi_j^a \rightarrow c_j^a]$, one for each agent $a \in A$.

The **computation tree** of P is a directed tree defined inductively as follows. Its nodes are call sequences and its root is the empty call sequence ε . Further, if c is a node and for some rule $\psi_j^a \rightarrow c_j^a$ we have $c \models \psi_j^a$, then $c.c_j^a$ is a node that is a direct descendant of c . Intuitively, the arc from c to $c.c_j^a$ records the effect of the execution of the rule $\psi_j^a \rightarrow c_j^a$ performed after the call sequence c took place.

By a **computation** of a gossip protocol we mean a maximal rooted path in its computation tree. In what follows we identify each computation with the unique call sequence it generates. We say that the gossip protocol P is **partially correct** if for all leaf c of the computation tree of P

$$c \models \bigwedge_{a \in A, S \in \text{Sec}} F_a S, \quad (4)$$

i.e., if each agent is an expert in the gossip situation $c(\text{root})$.

We say furthermore that P **terminates** if all its computations are finite and say that P is **correct** if it is partially correct and terminates.

We also consider two variants of termination. To define them we need a subsidiary notion. We call a rule **enabled** after a call sequence c if its guard is true after c . Given a gossip protocol we say that an agent is **enabled** after a call sequence c if one of the rules in its program is enabled.

We now 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 a gossip protocol P **rule-fairly terminates** (resp. **agent-fairly terminates**) if all its rule-fair (resp. agent-fair computations) are finite. Agent-fairness was introduced in [3], where it was simply called fairness, while rule-fairness was introduced in [7] and further studied in [8].

7 Open problems about the gossip protocols

We begin with the following problem.

Problem 5 *Characterize the class of graphs for which correct propositional gossip protocols exist.*

Note that the LNS protocol from Example 2 shows that this class of graphs include all complete digraphs. However, as we will show in Theorem 2, star graphs do not belong to this class. We conjecture that any digraph whose complement of its set of edges contains a directed cycle does not have this property.

Comments

For further discussion it is useful to introduce some terminology. We say that a gossip protocol is **for arbitrary connected digraphs** if each agent a can only call the agents from the set N_a of its in-neighbours. In such gossip protocols the underlying digraph is a parameter that can be uncovered from the sets N_a . If the underlying digraph is supposed to be complete, we say that the gossip protocol is **parametric**. So in both cases we actually deal with a ‘parametrized’ gossip protocol, which is a template for an infinite set of gossip protocols.

For example, in [7] and [8] a correct gossip protocol with the following program for agent i was considered:

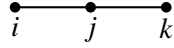
$$*[\bigwedge_{j \in N_i, S \in \text{Sec}} F_i S \wedge \neg K_i F_j S \rightarrow ij].$$

Informally, agent i calls a neighbour j if i is familiar with some secret (here S) and he does not know whether j is familiar with it. This gossip protocol is for arbitrary connected digraphs. However, it is not propositional.

In Example 2 the LNS gossip protocol was introduced. It is parametric and is both correct and propositional. However, its counterpart for arbitrary connected digraphs, so with the program

$$*[\bigwedge_{j \in N_i} \neg F_i J \rightarrow ij].$$

for agent i , is obviously incorrect. Indeed, for the graph



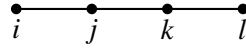
it terminates after the calls ij , jk with the agent i not being an expert.

In turn, a natural gossip protocol Exp for arbitrary connected digraphs, with the program

$$*[\bigwedge_{j \in N_i} \neg \text{Exp}_i \rightarrow ij]$$

for agent i , is obviously partially correct but it does not terminate, as initially a fixed call ij , with $j \in N_i$, can be repeated indefinitely.

In [8] we proved that the Exp gossip protocol agent-fairly terminates in the case of rings. However, this is not the case in general. Indeed take the graph



Then $(ij, ji, kl, lk)^*$ is an infinite agent-fair computation.

On the other hand, the following result holds. It generalizes the above result of [8] since for the case of rings the program for each agent has just one guard and consequently the notions of agent-fairness and rule-fairness coincide.

Theorem 1 *The gossip protocol Exp for arbitrary connected digraphs rule-fairly terminates.*

Proof. Suppose otherwise. Consider an infinite rule-fair computation ξ . We say that an agent i becomes an expert in ξ if for some element c of ξ we have $c \models \text{Exp}_i$.

As some agent i does not become an expert in ξ , there is a secret J that i does not learn in ξ . Let $i = i_1, i_2, \dots, i_h = j$ be a path connecting agent i with j . By rule-fairness the call $i_1 i_2$ takes place infinitely often. Hence agent i_2 does not become an expert in ξ and consequently by rule-fairness the call $i_2 i_3$ takes place infinitely often. Repeating this argument we conclude that each call $i_g i_{g+1}$, where $g \in \{1, \dots, h-1\}$, takes place infinitely often in ξ . So the call sequence $i_{h-1} i_h, i_{h-2} i_{h-1}, \dots, i_1 i_2$, possibly interspersed with other calls, exists in ξ . After the last call agent i learns the secret J , which yields a contradiction. \square

We now show that for a natural class of connected graphs no correct propositional gossip protocol exists.

Theorem 2 *Suppose that the agents form a star graph, so a graph in which some agent, say a , is present in all edges. Then no correct propositional gossip protocol exists.*

Proof. We begin by making two simple observations.

Claim 1 *Consider a propositional gossip protocol P . Suppose that c, c, d is a prefix of a computation of P such that some agent a*

- *is involved in the call c ,*
- *is not involved in any call in d , and*
- *is a caller in d .*

Then also c, c, d is a prefix of a computation of P .

Proof. Consider the guard ϕ associated with the call d . By assumption on P , ϕ is a propositional formula built out of the atomic formulas of the form $F_a S$. By assumption on d the truth of these atomic formulas is not affected by any call in d . So the truth value of ϕ before and after the call sequence d is the same. But ϕ is true just after this call sequence, so it is also true just before it. This shows that d can be performed immediately after c . \square

Claim 2 *Let ξ be a computation of a terminating propositional gossip protocol. If an agent, say a , becomes an expert after the call c , then a is not a caller in any call that follows c in ξ .*

Proof. Suppose otherwise. Let d be such a call in ξ and let ϕ be the corresponding guard. This call does not affect the set of gossips agent a is familiar with since this set as of the call c equals Sec . By assumption ϕ is a propositional formula built out of the atomic formulas of the form $F_a S$, so after the call d the formula ϕ remains true. Hence the call d can be indefinitely repeated, which shows that the considered protocol does not terminate. \square

Let now G be the considered star graph with an agent a present in all edges. Suppose by contradiction that a correct propositional gossip protocol P for G exists. Let ξ be a computation of P . Then agent a is involved in all calls in ξ . Let c' be the call after which a became an expert and let c be the call that precedes c' in ξ . By assumption the call c concerns agent a and some agent b not involved in the call c' . After the call c' agent b is not yet an expert, hence it is involved in ξ in another call. Let d be the first such call.

So for some call sequences c and d , we have that c, c, d, d is a prefix of ξ . Agent b is not involved in any calls in d and by Claim 2 it is also the caller in d . So by Claim 1 also c, c, d is a prefix of a computation, say χ , of P . Since both c and d involve the same pair of agents, after the second call the set of gossips of agent b does not change. Hence in χ the guard ϕ remains true after d and consequently this call can be indefinitely repeated. So P does not terminate. \square

One of the early results, see for instance [31], is that for $n \geq 4$ agents at least $2n - 4$ phone calls are needed to reach a situation in which each agent is an expert. Further, it is easy and well-known that this final situation can be reached using $2n - 4$ calls. Indeed, assume that the set of agents is $\{a, b, c, d, i_1, \dots, i_{n-4}\}$, where $n \geq 4$, (if $n = 4$ then there are no i_j agents) and take the following call sequence

$$\begin{aligned} &(a, i_1), (a, i_2), \dots, (a, i_{n-4}), \\ &(a, b), (c, d), (a, d), (b, c), \\ &(a, i_1), (a, i_2), \dots, (a, i_{n-4}). \end{aligned}$$

Problem 6 *Prove that the lower bound $2n - 4$ cannot be achieved for the gossip protocols. In other words, prove that every correct gossip protocol for $n \geq 4$ agents generates computations of length $> 2n - 4$.*

Comments

We show that this problem can be solved for $n = 4$. To prove it we need the following observations.

Lemma 3 *For all agents a, b, c and all call sequences c and all formulas ϕ*

$$c \models K_a \phi \text{ iff } c, bc \models K_a \phi.$$

Consequently, the same equivalence holds for all formulas that are Boolean combinations of formulas of the form $K_a \phi$, so in particular for each guard ψ used in a program for agent a .

Proof. By definition $c \sim_a c, bc$, which implies the claim. \square

This note states that the calls in which agent a is not involved have no effect on the truth of the guards used in the programs for agent a . This clarifies the syntax of the guards. If we allowed in guards for agent a formulas of the form $F_b C$ as conjuncts, this natural and desired property would not hold anymore. Indeed, for $c = \varepsilon$ we have $c, bc \models F_b C$, while $c \not\models F_b C$.

We also need the following observation that confirms the intuition that two calls involving different pairs of agents can be executed in an arbitrary order.

Lemma 4 Consider a protocol P . Let a, b, c, d be four agents such that for some call sequences c and d we have that c, ab, cd, d is a computation of P . Then also c, cd, ab, d is a computation of P .

Proof. Let ϕ_a be the guard that precedes the call ab in the program for agent a and ϕ_c the guard that precedes the call cd in the program for agent c . We have $c \models \phi_a$ and $c, ab \models \phi_c$. By Lemma 3

$$c \models \phi_c \text{ iff } c, ab \models \phi_c$$

and

$$c, cd \models \phi_a \text{ iff } c \models \phi_a.$$

Hence $c \models \phi_c$ and $c, cd \models \phi_a$. This means that c, cd, ab is a prefix of a computation of P .

Further, by the definition of the \sim_i relations, for all agents i and all call sequences d' we have $c, ab, cd, d' \sim_i c, cd, ab, d'$. Hence for all formulas of the form $K_i\phi$ we have

$$c, ab, cd, d' \models K_i\phi \text{ iff } c, cd, ab, d' \models K_i\phi$$

and consequently for all guards ϕ preceding the calls in d we have

$$c, ab, cd, d \models \phi \text{ iff } c, cd, ab, d \models \phi.$$

This concludes the proof. □

The above observation also holds for infinite computations but we shall not need it in the sequel.

Finally, we need to reason about the following notion. We call a computation of a gossip protocol **bad** if some agent is involved in the first two calls of it.

Lemma 5 Every gossip protocol admits bad computations.

Proof. Fix a gossip protocol P and consider its computation ξ that is not bad. By appropriate renaming we can assume that it begins with ab, cd . Let c be the third call in ξ and ψ the guard associated in P with the call c . By assumption we have $(ab, cd) \models \psi$. Four cases arise.

Case 1. Agent a is the caller in c .

We have $ab, cd \sim_a ab$, so by Lemma 4

$$(ab, cd) \models \psi \text{ iff } ab \models \psi$$

and hence $ab \models \psi$. In other words, the call c can be executed in P directly after the call ab , that is, ab, c is a prefix of a computation of the protocol P .

Case 2. Agent b is the caller in c .

We have $ab, cd \sim_b ab$, so, as in Case 1, $ab \models \psi$. Hence the call c can be executed in P directly after the call ab , that is, ab, c is a prefix of a computation of the protocol P .

Case 3. Agent c is the caller in c .

By Lemma 4 the calls ab and cd can be reversed in ξ , hence a computation of P exists that begins with cd, ab, c . So $(cd, ab) \models \psi$. We have $cd, ab \sim_c cd$, so, as in Case 1, $cd \models \psi$. Hence the call c can be executed in P directly after the call cd , that is, cd, c is a prefix of a computation of the protocol P .

Case 4. Agent d is the caller in c .

As in Case 3 a computation of P exists that begins with cd, ab, c . We have $(cd, ab) \models \psi$ and $cd, ab \sim_d cd$, so, as in Case 3, $cd \models \psi$. Hence the call c can be executed in P directly after the call cd , that is, cd, c is a prefix of a computation of the protocol P . □

Theorem 6 *Every correct gossip protocol for 4 agents generates computations of length > 4 .*

Proof. Suppose that $A = \{a, b, c, d\}$. Consider a correct protocol P .

We first show that each bad computation ξ is of length > 4 . There are 4 agents, so some agent, say a , is involved neither in the first nor the second call of ξ . So after these two calls none of the agents b, c, d is familiar with the secret of agent a . Each call increases the number of agents familiar with the secret of agent a by at most 1. So at least three more calls are needed to obtain a situation which all agents are familiar with the secret of a . We conclude that ξ is of length at least 5.

The conclusion now follows by Lemma 5. \square

Problem 6 concerned the lower bound $2n - 4$. It is useful to note that the lower bound $2n - 3$ can be achieved. So for $n = 4$ the precise bound is 5.

Theorem 7 *Suppose that $n \geq 4$. There exists a correct gossip protocol for n agents all computations of which are of length $2n - 3$.*

Proof. Assume that $A = \{a_1, \dots, a_n\}$. We use the observation due to [9] that the following call sequence results in all agents being experts:

$$\begin{aligned} & a_1 a_2, a_1 a_3, \dots, a_1 a_n, \\ & a_1 a_2, a_1 a_3, \dots, a_1 a_{n-1}. \end{aligned}$$

In this call sequence first agent a_1 calls all other agents and subsequently calls them again except agent a_{n-1} . But the exact order of the calls in each phase is not important. This allows us to use a gossip protocol with a simple structure.

We choose an arbitrary agent $a \in A$ and select for it the following program:

$$\begin{aligned} & * [\bigwedge_{i \in A \setminus \{a\}} \neg F_a I \rightarrow ai \\ & \quad \bigwedge_{i \in A \setminus \{a\}} \text{Exp}_a \wedge \neg K_a \text{Exp}_i \rightarrow ai]. \end{aligned}$$

The programs for the other agents are empty.

Note that after each call the corresponding guard becomes false. Further, by the form of the above program in every computation of the protocol first $n - 1$ calls of a take place, each to a different agent. These calls correspond to the first set of guards and are executed in an arbitrary order.

After the last of these $n - 1$ calls, say ab , agents a and b both become experts and agent a knows it. So from this moment on the guard $\text{Exp}_a \wedge \neg K_a \text{Exp}_b$ is false and the second call ab does not take place. Hence the remaining part of the computation consists of $n - 2$ calls of agent a , each to a different agent from $A \setminus \{a, b\}$. These calls correspond to the second set of guards and are executed in an arbitrary order, as well.

After these $2n - 3$ calls all agents become experts and the protocol terminates. \square

Problem 7 *Prove that the lower bound $2n - 3$ cannot be achieved by a correct propositional gossip protocol.*

8 Open problems about verification of gossip protocols

We studied in [4] the computational complexity of partial correctness and termination of gossip protocols and showed that when in guards only formulas from the sublanguage \mathcal{L}_1 are used both problems are in coNP^{NP} . This brings us naturally to the following problems.

Problem 8 *What is the exact computational complexity of checking whether a given gossip protocol*

- *is partially correct,*
- *terminates?*

We conjecture that both problems are coNP^{NP} -complete. This problem can be generalized as follows. In *generalized gossiping* the aim is to reach a gossip situation in which some given formula $\phi \in \mathcal{L}$ is true, possibly other than $\bigwedge_{a \in A} \text{Exp}_a$. We say that a gossip protocol P is ϕ -**partially correct** if for all leafs c of the computation tree of P

$$c \models \phi.$$

For example the HMS gossip protocol from Example 3 is $\bigwedge_{i,j \in A} K_i F_j I$ -partially correct, since upon termination all the guards are false. In contrast, this protocol is not $\bigwedge_{i,j \in A} K_i \text{Exp}_j$ -partially correct. Indeed, assume three agents, a, b, c , and the call sequence (ab, ac, bc) after which this gossip protocol terminates. However, $(ab, ac, bc) \models K_a \text{Exp}_b$ does not hold.

Problem 9 *What is the complexity of checking whether for a given formula $\phi \in \mathcal{L}$ a given gossip protocol is ϕ -partially correct?*

These questions can also be asked for parametrized gossip protocols.

Problem 10 *What is the complexity of checking whether a given gossip protocol for arbitrary connected digraphs*

- *is partially correct,*
- *terminates?*

Next problem concerns relation between two gossip protocols. To define it we need to discuss the computation trees.

We say that two unordered trees are **equal** if they become identical after some possible rearrangement of the children of each node in these trees.¹ Given a tree T , any removal of branches from T yields a **subtree** of T .

Consider now two gossip protocols P and P' . We say that P **can simulate** P' if the computational tree of P' is equal to some subtree of P . If both P can simulate P' and P' can simulate P then we say that they are **bisimilar**. Clearly the computational trees of two bisimilar gossip protocols are equal.

Example 4 *Consider the LNS and HMS gossip protocols introduced in Examples 2 and 3.*

Assume that there are only three agents, a, b, c . There are only four maximal call sequences that can be generated by LNS beginning with the call ab , namely:

$$ab.bc.ac, ab.cb.ac, ab.ac.bc, ab.ca.bc.$$

For every other starting call $ac, ba, bc, ca, or cb$, there are analogous four such maximal call sequences. It is easy to check that all these call sequences can be generated by HMS. So for three agents HMS can simulate LNS.

However, HMS can also generate the call sequence $ab.bc.ca$ that LNS cannot, so these two protocols are not bisimilar. \square

¹This rearrangement is allowed because computational trees are unordered. If we fix an order of children, e.g., subject to lexicographic order on the last call, then a computational tree for a given protocol is unique and this problem when checking for the equality of two trees no longer occurs.

We actually conjecture that HMS can simulate LNS for any number of agents. This naturally suggests the following two problems.

Problem 11 *What is the exact complexity of checking for two given parametric protocols whether*

- *one simulates the other,*
- *they are bisimilar?*

Problem 12 *What is the exact computational complexity of checking for two gossip protocols P and P' whether*

- *P simulates P' ,*
- *P and P' are bisimilar?*

Comments

The following result provides some insights into the last problem.

Theorem 8 *Checking whether a gossip protocol P can simulate another protocol P' can be done in exponential time.*

To prove it we need some preparatory results. We call the second call c in a call sequence $c_1.c.c_2.c.c_3$ **epistemically redundant** if $c_1.c(\text{root}) = c_1.c.c_2.c(\text{root})$. In [4] we established the following result (as Lemma 6) showing that removing an epistemically redundant call does not affect the truth of any formula from \mathcal{L}_1 .

Lemma 9 *If $c_1.c.c_2.c.c_3$ is a call sequence where the second call c is epistemically redundant, then for any formula $\psi \in \mathcal{L}_1$:*

$$c_1.c.c_2.c.c_3 \models \psi \text{ iff } c_1.c.c_2.c_3 \models \psi.$$

We say that a call c that appears in a call sequence c is **productive** if $c_1(\text{root}) \neq c_1.c(\text{root})$, where for some call sequences c_1 and c_2 , $c = c_1.c.c_2$.

Further, we call a subsequence c_2 of a call sequence c **stationary** if

$$c_1.c_1(\text{root}) = c_1.c_1.c_2(\text{root}) = \dots = c_1.c_2(\text{root}),$$

where $c = c_1.c_2.c_3$ for some call sequences c_1 and c_3 , and for some $k \geq 1$, $c_2 = c_1.c_2 \dots c_k$.

Recall that n is the number of agents.

Lemma 10 *Every call sequence contains at most $n^2 - n$ productive calls.*

Proof. The minimal total number of secrets in a gossip situation is n , which is achieved in the initial gossip situation root . In turn, the maximal total number of secrets is n^2 , which is achieved in the gossip situation in which every agent is an expert. After each productive call at least one agent learns a new secret. So in a given call sequence there can be at most $n^2 - n$ productive calls. \square

Lemma 11 *Every call sequence c of length $\geq n^4$ has a stationary subsequence of length $\geq n^2$.*

Proof. Split c into a sequence of n^2 subsequences, each of length at least n^2 . Suppose none of these subsequences is stationary. Then each of them contains a productive call. So c contains at least n^2 productive calls, which contradicts Lemma 10. \square

Corollary 12 *Every call sequence c of length $\geq n^4$ contains an epistemically redundant call.*

Proof. Take a stationary subsequence c_1 of c guaranteed by Lemma 11. There are at most $n^2 - n$ different calls, so some call c appears in c_1 at least twice. Its second occurrence in c_1 is then epistemically redundant in c . \square

Next lemma shows that lack of simulation entails existence of a specific call sequence of a bounded length.

Lemma 13 *If P cannot simulate P' then there exist a call sequence c of length $\leq n^4$ that can be generated by P' , but not by P .*

Proof. If P cannot simulate P' then take a shortest rooted path in the computation tree of P' that does not exist in the computation tree of P . This path corresponds to some call sequence c .

Suppose by contradiction that the length of c exceeds n^4 . By Corollary 12, c can be partitioned into $c_1.c.c_2.c.c_3.c'$ such that $c_1.c(\text{root}) = c_1.c.c_2.c(\text{root})$.

By Lemma 9 for any formula $\phi \in \mathcal{L}_1$ and a subsequence c_4 of c_3 we have

$$c_1.c.c_2.c.c_4 \models \phi \text{ iff } c_1.c.c_2.c_4 \models \phi. \quad (5)$$

By the definition of c the call sequence $c_1.c.c_2.c.c_3$ can be generated by P , while the call sequence c cannot. So P does not have a rule $\phi \rightarrow c'$ such that $c_1.c.c_2.c.c_3 \models \phi$.

Further, by (5) the call sequence $c_1.c.c_2.c_3$ can be generated by P , while, again by (5), the call sequence $c_1.c.c_2.c_3.c'$ (so c with the second indicated occurrence of c omitted) cannot.

On the other hand, by assumption c can be generated by P' , so P' has a rule $\phi' \rightarrow c'$ such that $c_1.c.c_2.c.c_3 \models \phi'$. Together with (5) this implies that also $c_1.c.c_2.c_3.c'$ can be generated by P' . This yields a contradiction with the definition of c . \square

Proof of Theorem 8. To show that P can simulate P' it suffices by Lemma 13 to check that each call sequence of length $\leq n^4$ which can be generated by P' can also be generated by P .

We showed in [4] that checking for a given call sequence c and a formula $\phi \in \mathcal{L}_1$ whether $c \models \phi$ can be done in exponential time (actually in coNP^{NP} time). This implies that checking whether a given call sequence c of length $\leq n^4$ can be generated by a gossip protocol can be done in exponential time.

Consequently, checking whether P can simulate P' can be done in exponential time, as well, since there are exponentially many call sequences of length $\leq n^4$. \square

Corollary 14 *Checking whether gossip protocols P and P' are bisimilar can be done in exponential time.*

As a side remark of independent interest we conclude the paper with the following consequence of Corollary 12.

Corollary 15 *If a gossip protocol terminates then all its computations are of length $< n^4$.*

To prove we use the following result from [8] that for the case of formulas from \mathcal{L}_1 is actually a special case of Lemma 9.

Lemma 16 (Stuttering) *Suppose that $c := c_1.c.c_2$ and $d := c_1.c.c.c_2$. Then for all formulas $\phi \in \mathcal{L}$, $c \models \phi$ iff $d \models \phi$.*

Proof of Corollary 15. Consider a finite computation of length $\geq n^4$ and the corresponding call sequence c . By Corollary 12, c begins with a call sequence $c_1.c.c_2.c$, where $c_1.c(\text{root}) = c_1.c.c_2.c(\text{root})$.

Let $\psi \rightarrow c$ be the rule used in the considered gossip protocol to generate the second occurrence of the call c . Then $c_1.c.c_2 \models \psi$ and by Lemma 9 for $c_3 = \varepsilon$ we have $c_1.c.c_2.c \models \psi$.

Hence by the repeated use of the Stuttering Lemma 16 for $c_2 = \varepsilon$ and all $i \geq 1$, $c_1.c.c_2.c^i \models \psi$. Consequently, after the call sequence $c_1.c.c_2$ is generated, the rule $\psi \rightarrow c$ can be repeatedly applied. Hence $c_1.c.c_2.c^\omega$ is an infinite sequence of calls that corresponds to an infinite computation. \square

Acknowledgements

We thank reviewers for useful suggestions concerning the presentation. The second author was partially supported by EPSRC grants EP/M027287/1 and EP/P020909/1.

References

- [1] Carolyn Jane Anderson, Nate Foster, Arjun Guha, Jean-Baptiste Jeannin, Dexter Kozen, Cole Schlesinger & David Walker (2014): *NetKAT: semantic foundations for networks*. In: *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14*, ACM, pp. 113–126, doi:10.1145/2535838.2535862.
- [2] Krzysztof R. Apt, Davide Grossi & Wiebe van der Hoek (2018): *When Are Two Gossips the Same?* In Gilles Barthe, Geoff Sutcliffe & Margus Veanes, editors: *LPAR-22. 22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, EPIc Series in Computing 57*, EasyChair, pp. 36–55, doi:10.29007/ww65.
- [3] Krzysztof R. Apt, Davide Grossi & Wiebe van der Hoek (2016): *Epistemic Protocols for Distributed Gossiping*. In: *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2015)*, *EPTCS 215*, pp. 51–66, doi:10.4204/EPTCS.215.5.
- [4] Krzysztof R. Apt, Eryk Kopczyński & Dominik Wojtczak (2017): *On the Computational Complexity of Gossip Protocols*. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017*, pp. 765–771, doi:10.24963/ijcai.2017/106.
- [5] Krzysztof R. Apt & Dominik Wojtczak (2016): *On Decidability of a Logic of Gossips*. In: *Proceedings of the 15th European Conference, JELIA 2016, Lecture Notes in Computer Science 10021*, Springer, pp. 18–33, doi:10.1007/978-3-319-48758-8_2.
- [6] Krzysztof R. Apt & Dominik Wojtczak (2017): *Common Knowledge in a Logic of Gossips*. In: *Proc. of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2017)*, *EPTCS 251*, pp. 10–27, doi:10.4204/EPTCS.251.2.
- [7] Krzysztof R. Apt & Dominik Wojtczak (2017): *Decidability of Fair Termination of Gossip Protocols*. In: *Proc. of the 21st International Conference on Logic for Programming, Artificial Intelligence and Reasoning (LPAR 21)*, *Kalpa Publications in Computing 1*, pp. 73–85, doi:10.29007/62s4.
- [8] Krzysztof R. Apt & Dominik Wojtczak (2018): *Verification of Distributed Epistemic Gossip Protocols*. *J. Artif. Intell. Res. (JAIR) 62*, pp. 101–132, doi:10.1613/jair.1.11204.
- [9] Maduka Attamah, Hans Van Ditmarsch, Davide Grossi & Wiebe van der Hoek (2014): *A Framework for Epistemic Gossip Protocols*. In: *Proceedings of the 12th European Conference on Multi-Agent Systems (EUMAS 2014), Revised Selected Papers*, 8953, Springer, pp. 193–209, doi:10.1007/978-3-319-17130-2_13.
- [10] Maduka Attamah, Hans Van Ditmarsch, Davide Grossi & Wiebe van der Hoek (2014): *Knowledge and Gossip*. In: *Proceedings of ECAI'14*, IOS Press, pp. 21–26, doi:10.3233/978-1-61499-419-0-21.
- [11] Martin C. Cooper, Andreas Herzig, Faustine Maffre, Frédéric Maris & Pierre Régnier (2016): *A simple account of multiagent epistemic planning*. In: *Proceedings of ECAI 2016*, IOS Press, pp. 193–201, doi:10.3233/978-1-61499-672-9-193.
- [12] Martin C. Cooper, Andreas Herzig, Faustine Maffre, Frédéric Maris & Pierre Régnier (2016): *Simple Epistemic Planning: Generalised Gossiping*. In: *Proceedings of ECAI 2016, Frontiers in Artificial Intelligence and Applications 285*, IOS Press, pp. 1563–1564, doi:10.3233/978-1-61499-672-9-1563.
- [13] Martin C. Cooper, Andreas Herzig, Frédéric Maris & Julien Vianey (2018): *Temporal Epistemic Gossip Problems*. In: *European Conference on Multi-Agent Systems*, Springer, pp. 1–14, doi:10.1007/978-3-030-14174-5_1.

- [14] Hans van Ditmarsch, Jan van Eijck, Pere Pardo, Rahim Ramezani & François Schwarzentruber (2017): *Epistemic Protocols for Dynamic Gossip*. *J. of Applied Logic* 20(C), pp. 1–31, doi:10.1016/j.jal.2016.12.001.
- [15] Hans van Ditmarsch, Jan van Eijck, Pere Pardo, Rahim Ramezani & François Schwarzentruber (2018): *Dynamic Gossip*. *Bull. Iran. Math. Soc.*, pp. 1–28, doi:10.1007/s41980-018-0160-4.
- [16] Hans van Ditmarsch, Davide Grossi, Andreas Herzig, Wiebe van der Hoek & Louwe B. Kuijter (2016): *Parameters for Epistemic Gossip Problems*. In: *Proceedings of the 12th Conference on Logic and the Foundations of Game and Decision Theory (LOFT 2016)*. Available at <https://pdfs.semanticscholar.org/74b5/2c025f335ba487cac612019e39ce6c818448.pdf>.
- [17] Hans van Ditmarsch, Ioannis Kokkinis & Anders Stockmarr (2017): *Reachability and Expectation in Gossiping*. In Bo An, Ana Bazzan, João Leite, Serena Villata & Leendert van der Torre, editors: *PRIMA 2017: Principles and Practice of Multi-Agent Systems*, Springer International Publishing, Cham, pp. 93–109, doi:10.1007/978-3-319-69131-2_6.
- [18] Ronald Fagin, Joseph Y. Halpern, Yoram Moses & Moshe Y. Vardi (1997): *Knowledge-Based Programs*. *Distributed Computing* 10(4), pp. 199–225, doi:10.1007/s004460050038.
- [19] Malvin Gattinger & Jana Wagemaker (2018): *Towards an Analysis of Dynamic Gossip in NetKAT*. In: *International Conference on Relational and Algebraic Methods in Computer Science*, Springer, pp. 280–297, doi:10.1007/978-3-030-02149-8_17.
- [20] Sandra M. Hedetniemi, Stephen T. Hedetniemi & Arthur L. Liestman (1988): *A survey of gossiping and broadcasting in communication networks*. *Networks* 18(4), pp. 319–349, doi:10.1002/net.3230180406.
- [21] Andreas Herzig & Faustine Maffre (2015): *How to Share Knowledge by Gossiping*. In: *Proc of the 13th European Conference on Multi-Agent Systems (EUMAS 2015), Revised Selected Papers*, 9571, Springer, pp. 249–263, doi:10.1007/978-3-319-33509-4_20.
- [22] Andreas Herzig & Faustine Maffre (2017): *How to Share Knowledge by Gossiping*. *AI Communications* 30(1), pp. 1–17, doi:10.3233/AIC-170723.
- [23] Charles A. R. Hoare (1978): *Communicating Sequential Processes*. *Commun. ACM* 21(8), pp. 666–677, doi:10.1145/359576.359585.
- [24] Juraj Hromkovič, Ralf Klasing, Andrzej Pelc, Peter Ruzicka & Walter Unger (2005): *Dissemination of Information in Communication Networks - Broadcasting, Gossiping, Leader Election, and Fault-Tolerance*. Texts in Theoretical Computer Science. An EATCS Series, Springer, doi:10.1007/b137871.
- [25] David Kempe, Alin Dobra & Johannes Gehrke (2003): *Gossip-based computation of aggregate information*. In: *Proc. of the 44th Annual IEEE Symposium on Foundations of Computer Science, FOCS '03*, IEEE, pp. 482–491, doi:10.1109/SFCS.2003.1238221.
- [26] Anne-Marie Kermarrec & Maarten van Steen (2007): *Gossiping in distributed systems*. *Operating Systems Review* 41(5), pp. 2–7, doi:10.1145/1317379.1317381.
- [27] Rivka Ladin, Barbara Liskov, Liuba Shrira & Sanjay Ghemawat (1992): *Providing high availability using lazy replication*. *ACM Transactions on Computer Systems (TOCS)* 10(4), pp. 360–391, doi:10.1145/138873.138877.
- [28] Robin Milner (1980): *A Calculus of Communicating Systems*. *Lecture Notes in Computer Science* 92, Springer, doi:10.1007/3-540-10235-3.
- [29] Davide Sangiorgi (2009): *On the origins of bisimulation and coinduction*. *ACM Trans. Program. Lang. Syst.* 31(4), pp. 15:1–15:41, doi:10.1145/1516507.1516510.
- [30] Holger Spakowski & Jörg Vogel (2000): Θ_2^P -Completeness: A Classical Approach for New Results. In: *International Conference on Foundations of Software Technology and Theoretical Computer Science*, Springer, pp. 348–360, doi:10.1007/3-540-44450-5_28.
- [31] Robert Tijdeman (1971): *On a telephone problem*. *Nieuw Archief voor Wiskunde* 3(XIX), pp. 188–192.
- [32] Klaus W Wagner (1987): *More complicated questions about maxima and minima, and some closures of NP*. *Theoretical Computer Science* 51(1-2), pp. 53–80, doi:10.1016/0304-3975(87)90049-1.