

Common Knowledge 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 aim at arriving, by means of point-to-point or group communications, at a situation in which all the agents know each other secrets. Recently a number of authors studied distributed epistemic gossip protocols. These protocols use as guards formulas from a simple epistemic logic, which makes their analysis and verification substantially easier.

We study here common knowledge in the context of such a logic. First, we analyze when it can be reduced to iterated knowledge. Then we show that the semantics and truth for formulas without nested common knowledge operator are decidable. This implies that implementability, partial correctness and termination of distributed epistemic gossip protocols that use non-nested common knowledge operator is decidable, as well. Given that common knowledge is equivalent to an infinite conjunction of nested knowledge, these results are non-trivial generalizations of the corresponding decidability results for the original epistemic logic, established in [2].

1 Introduction

Common knowledge is a fundamental notion in epistemic reasoning. It has its origins in the book of the philosopher David Lewis, [19], and the article of the sociologist Morris Friedell, [15]. By now this concept was applied in many other fields, including artificial intelligence, psychology, computer science, game theory, and logic. An early work on this subject in computer science and logic is discussed in [14]. For more recent accounts and surveys see e.g., [11] and [20].

Study and use of various logics equipped with the common knowledge operator is a rich field. As example of recent publications let us just mention [7], where an update logic augmented with common knowledge is investigated, and [21], where the correctness of epistemic protocols that rely on common knowledge is studied.

The purpose of this article is to investigate common knowledge in the context of a simple epistemic logic proposed in [1] to express and analyze distributed epistemic gossip protocols. 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., the early survey [16] or the book coverage [18]. Distributed epistemic gossip protocols were introduced in [6], and further studied in [5, 17, 1, 12, 9, 10, 8], where in particular various distributed gossiping protocols, their types, epistemic aspects and objectives, and their interpretation as planning problems were analyzed. Such protocols are strikingly simple in their syntax based on epistemic logic (though not semantics), which makes it easier to reason about them.

In [2] we showed that the distributed epistemic gossip protocols introduced in [1] are implementable and proved that the problems of partial correctness and termination of such protocols are decidable, as well. In [3] we built upon these results and showed that the implementability of a distributed epistemic gossip protocol is a $P_{\parallel}^{\text{NP}}$ -complete problem, while the problems of its partial correctness and termina-

tion are in coNP^{NP} . We also established in [4] that fair termination of the distributed epistemic gossip protocols is decidable, as well.

In this paper we extend the results of [2] to the language that includes the common knowledge operator. Given that common knowledge is equivalent to an infinite conjunction of nested knowledge, these results are non-trivial generalizations of the previous results.

The obtained results clarify when and how common knowledge can arise in the context of gossiping. We prove that three or more agents can have common knowledge only of true statements. This is not the case for two agents, even if they do not communicate. We also show that under some assumptions common knowledge of two agents coincides with the 4th fold iterated knowledge. The main open problem is whether in this context common knowledge can always be reduced to iterated knowledge.

2 Syntax

The purpose of this paper is to analyze common knowledge in the context of gossip protocols. To describe it we use a simple modal language introduced in [1], though we allow now the common knowledge operator instead of the agent related knowledge operator.

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.

Assume a fixed ordering on the agents. Each *call* concerns two different agents, say a and b , and is written as ab or (a, b) , 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 an epistemic language \mathcal{L}^{ck} defined by the following grammar:

$$\phi ::= F_{ap} \mid \neg\phi \mid \phi \wedge \phi \mid C_G\phi,$$

where $p \in P$ and $a \in A$ and $G \subseteq A$. Each secret is viewed a distinct constant. We denote the secret of agent a by A , the secret of agent b by B , where $a, b \in A$ and $A, B \in P$, and so on. When G is a singleton, say $G = \{a\}$, then we write C_G as K_a , which is the knowledge operator used and studied in the context of this logic in [1] and [2].

We read F_{ap} as ‘agent a is familiar with the secret p ’, $K_a\phi$ as ‘agent a knows that formula ϕ is true’, and $C_G\phi$ as ‘the group of agents G commonly knows that formula ϕ is true’.

So F_{ap} is an atomic formula, while $K_a\phi$ and $C_G\phi$ are compound formulas. In what follows we shall distinguish the following sublanguages of \mathcal{L}^{ck} :

- \mathcal{L}_{pr} , its propositional part, which consists of the formulas that do not use the C_G modalities,
- \mathcal{L}_{wn} , which consists of the formulas without the nested use of the C_G modalities.

3 Semantics

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

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.

Below 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 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.

In [1] computations of the gossip protocols were studied, so both finite and infinite call sequences were used. Here we focus on the finite call sequences as we are only interested in the semantics of epistemic formulas. So to be brief, unless explicitly stated, a *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 write $c \sqsubseteq d$ to denote the fact that d extends c , i.e., that for some c' we have $c.c' = d$.

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

$$\begin{aligned} \varepsilon(s) &:= s, \\ (c.c)(s) &:= c(c(s)). \end{aligned}$$

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 $\mathcal{M} := (C, \{\sim_a\}_{a \in A})$, where each $\sim_a \subseteq C \times C$ is the smallest relation such that $\varepsilon \sim_a \varepsilon$ and the following conditions hold. Suppose $c \sim_a d$.

- (i) If $a \notin c$, then $c.c \sim_a d$ and $c \sim_a d.c$.
- (ii) If $a \in c$ and $c.c(\text{root})_a = d.c(\text{root})_a$, then $c.c \sim_a d.c$.

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

So for each set of agents there is exactly one gossip model. To illustrate the definition of \sim_a note for instance that by (i) we have $ab, bc \sim_a ab, bd$. But we do not have $bc, ab \sim_a bd, ab$ since $(bc, ab)(\text{root})_a = ABC \neq ABD = (bd, ab)(\text{root})_a$.

To define semantics of the C_G operator we use the relation $\sim_G \subseteq C \times C$ defined by

$$\sim_G = (\bigcup_{a \in G} \sim_a)^*,$$

where $*$ stands for the transitive reflexive closure of a binary relation. As stated in [1], each \sim_a is an equivalence relation. As a result each \sim_G is an equivalence relation, as well.

Finally, we recall the definition of truth.

Definition 2 Let (\mathcal{M}, c) be a pointed gossip model with $\mathcal{M} := (C, (\sim_a)_{a \in A})$ and $c \in C$. We define the satisfaction relation \models inductively as follows. For convenience we also include the special case of K_a (i.e., $G_{\{a\}}$). The clauses for Boolean connectives are as usual and omitted.

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

Further

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

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

So a formula $F_a p$ is true whenever secret p belongs to the set of secrets agent a is familiar with in the situation generated by the designated call sequence c applied to the initial situation root . The knowledge operator K_a is interpreted as customary in epistemic logic, using the equivalence relations \sim_a , and the C_G operator is defined as in [14].

While \mathcal{L}^{ck} is a pretty standard epistemic language, its semantics is not. Indeed, it describes the truth of formulas after a sequence of calls took place, by analyzing the statements of the form $(\mathcal{M}, c) \models \phi$. So we actually study here a limited version of a dynamic epistemic logic. To put it differently, we actually consider statements of the form $[c]\phi$, where [...] is the standard dynamic logic operator, see, e.g., [13]. This explains why the study of the logic \mathcal{L}^{ck} cannot be reduced to a study of a routine epistemic logic.

4 An alternative equivalence relation

To reason about the \sim_a and \sim_G relations it is easier to use an alternative equivalence relation between the call sequences that was introduced in [2]. It is based on a concept of a *view* of agent a of a call sequence c , written as c_a , and defined by induction as follows.

[Base]

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

[Step]

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

where the gossip situation s is defined by putting for $d \in A$

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

where s' is the last gossip situation in c_a .

Intuitively, a view of agent a of a call sequence c is the information he acquires by means of the calls in 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 3 Let $A = \{a, b, c\}$ and consider the call sequence (ac, bc, ac) . It generates the following successive gossip situations starting from root:

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

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 . □

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

$$c \equiv_a d \text{ iff } c_a = d_a.$$

So according to this definition two call sequences are equivalent for agent a if his views of them are the same. Below we shall rely on the following result from [2].

Theorem 4 (Equivalence) *For each agent a the relations \sim_a and \equiv_a coincide.*

5 Semantic matters

5.1 General considerations

We shall need below an alternative definition of truth of the $C_G\phi$ formulas. Given a sequence a_1, \dots, a_k of elements of G we abbreviate $K_{a_1} \dots K_{a_k} \phi$ to $K_{a_1 \dots a_k} \phi$. We also denote the set of finite sequences of elements of G by G^* .

Note 5 ([14]) *For all call sequences c and formulas $C_G\phi \in \mathcal{L}^{ck}$*

$$(\mathcal{M}, c) \models C_G\phi \text{ iff for all } t \in G^* (\mathcal{M}, c) \models K_t\phi.$$

In other words, the formula $C_G\phi$ is equivalent to the infinite conjunction $\bigwedge_{t \in G^*} K_t\phi$. We shall also need the following generalization of the corresponding result from [2] to the logic here studied.

Theorem 6 (Monotonicity) *Suppose that $\phi \in \mathcal{L}^{ck}$ is a formula that does not contain the \neg symbol. Then*

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

Proof. By Note 5 and Monotonicity Theorem 4 of [2]. □

Let us focus now on the case of ≥ 3 agents. The following result holds.

Theorem 7 *Suppose that $|G| \geq 3$. Then for all call sequences c and formulas $\phi \in \mathcal{L}^{ck}$*

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

Proof. First we prove that for all c and d

$$c \sim_G d. \tag{1}$$

By the transitivity of \sim_G it suffices to prove that $c \sim_G \varepsilon$. We prove it by induction on the length of c . By definition $\varepsilon \sim_G \varepsilon$. Suppose that for some c we have $c \sim_G \varepsilon$ and consider a call c . Take $a \in G$ such that $a \notin c$ (it exists since $|G| \geq 3$). Then $c.c \sim_a c$, so $c.c \sim_G \varepsilon$.

By (1) we have $(\mathcal{M}, c) \models C_G\phi$ iff $\forall d (\mathcal{M}, d) \models \phi$, which concludes the proof. □

Theorem 7 states that the formulas commonly known by the agents in a group of at least three agents are precisely the true formulas. An example of such a statement is that each agent is familiar with his secret, i.e., $\bigwedge_{a \in A} F_a A$. In contrast, a statement that an agent is familiar with the secret of another agent, i.e. $F_a B$, where $a \neq b$, is not always true, so for all call sequences c we have $(\mathcal{M}, c) \not\models C_G F_a B$, when $|G| \geq 3$.

5.2 The case of two agents

The situation changes when the group consists of two agents. In what follows we abbreviate $C_{\{a,b\}}$ to C_{ab} .

Example 8

(i) Consider the formula

$$\phi := \neg F_a B \vee \bigvee_{c \in A \setminus \{a,b\}} F_c B.$$

It states that if a is familiar with the secret of b , then also another agent different from a and b is familiar with this secret. Note that $(\mathcal{M}, ab) \models \neg \phi$, i.e., ϕ is not always true. We claim that $(\mathcal{M}, \varepsilon) \models C_{ab} \phi$.

First note that if $c \sim_a d$ or $c \sim_b d$ and the call ab does not appear in c , then it does not appear in d either. Consequently, if $c \sim_{\{a,b\}} \varepsilon$, then the call ab does not appear in c .

Conversely, take a call sequence c such that the call ab does not appear in it. We prove by induction on the length of c that $c \sim_{\{a,b\}} \varepsilon$. By definition $\varepsilon \sim_{\{a,b\}} \varepsilon$. Suppose that for some c we have $c \sim_{\{a,b\}} \varepsilon$ and consider a call c . Either $a \notin c$ or $b \notin c$, so either $c.c \sim_a c$ or $c.c \sim_b c$. Consequently $c.c \sim_{\{a,b\}} \varepsilon$.

We conclude that $c \sim_{\{a,b\}} \varepsilon$ iff the call ab does not appear in c . But for any such c we have $(\mathcal{M}, c) \models \phi$. This proves that $(\mathcal{M}, \varepsilon) \models C_{ab} \phi$.

This shows that even without any call two agents can commonly know a formula that is not always true.

(ii) We also have $(\mathcal{M}, ab) \models C_{ab}(F_a B \wedge F_b A)$, i.e., two agents can commonly know some non-trivial information about themselves.

Indeed, if the call ab appears in c , then $(\mathcal{M}, c) \models F_a B \wedge F_b A$ and that if the call ab appears in c and $c \sim_a d$ or $c \sim_b d$, then it also appears in d . \square

On the other hand for two agents the following partial analogue of Theorem 7 holds.

Theorem 9 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

$$(\mathcal{M}, c) \models C_{ab} \phi \text{ iff } \mathcal{M} \models \phi.$$

Proof. Suppose $(\mathcal{M}, c) \models C_{ab} \phi$. As noticed in Example 8(i) $c \sim_{\{a,b\}} \varepsilon$, so $(\mathcal{M}, \varepsilon) \models \phi$. By the Monotonicity Theorem 6 for all call sequences d we have $(\mathcal{M}, d) \models \phi$. So $\models \phi$. Further, $\models \phi$ implies $(\mathcal{M}, c) \models C_{ab} \phi$ for arbitrary call sequences c and formulas $C_{ab} \phi$. \square

Example 8(ii) shows that the restriction that the call ab does not appear in c cannot be dropped and Example 8(i) shows that the claim does not hold for formulas that do contain the \neg symbol.

Next, we show that for formulas that do not contain the \neg symbol common knowledge for the group of two agents coincides with the 4th fold iterated knowledge.

Consider an agent a and a call sequence c . We say that a call is *a-irrelevant* in c if its removal does not affect the view (in the sense of Section 4) of agent a of the call sequence. Starting from c we repeatedly remove from the current call sequence the first not yet analyzed call if it is *a-irrelevant* and otherwise we keep it. We call the outcome of such an iteration the *a-simplification* of c .

Example 10 Suppose,

$$c = bf.cd.bc.ce.df.ef.bh.\underline{af}.bg.\underline{ag}.\underline{ah},$$

where for the visibility we underlined the a -calls. Then the a -simplification of c results in the deletion of the calls bf and cd and equals

$$bc.ce.df.ef.bh.\underline{af}.bg.\underline{ag}.\underline{ah}.$$

The views of agent a of both call sequences are as follows.

$$\begin{aligned}
& A.B.C.D.E.F.D.H \xrightarrow{af} \\
& ABCDEF.B.C.D.E.ABCDEF.G.H \xrightarrow{ag} \\
& ABCDEFGH.B.C.D.E.ABCDEF.ABCDEFGH.H \xrightarrow{ah} \\
& ABCDEFGH.B.C.D.E.ABCDEF.ABCDEFGH.ABCDEFGH
\end{aligned}$$

□

Below we say that two calls are **linked** if exactly one agent participates in both of them. Consider now a call sequence c with no a -irrelevant calls that does not contain the call ab . We focus on the b -calls in c . By the assumption about c for each b -call c in c there is a sequence of calls c_1, \dots, c_k in c such that

- $c = c_1$,
- $b \in c_1$,
- for $i \in \{1, \dots, k-1\}$ $a \notin c_i$,
- for $i \in \{1, \dots, k-1\}$ the calls c_i and c_{i+1} are linked.
- $a \in c_k$.

We say then that c_1 **leads to** c_k . Further, we call each b -call in c that leads to the earliest possible a -call in c **b -essential for a** and call each other b -call in c **b -inessential for a** . Intuitively, agent a learns the secret of b only through the b -essential calls. In contrast, he can learn other secrets both through the b -essential and the b -inessential calls.

Example 11 Consider the call sequence

$$bc.ce.df.ef.bh.\underline{af}.bg.\underline{ag}.\underline{ah}$$

from Example 10. The only b -essential call for a is bc as it leads to af . In turn, in the call sequence $bh.ce.df.ef.\underline{af}.bg.bc.\underline{ag}$ the b -essential calls for a are bh and bg as both of them lead to ag and no b -call leads to the earlier a -call af . □

Consider now the first b -call in c , call it bc , that is b -inessential for a and suppose that ad is the first a -call in c to which bc leads to, where c and d may coincide. If $c = d$, then we delete bc from c and otherwise we replace bc by cd . Intuitively, we ‘reroute’ the information collected by agent b in the b -inessential calls leading to ad using agent d . This does not affect agent’s a view of the call sequence since he does not learn the secret of b through the b -inessential calls.

We repeat this operation, starting with c , for all b -calls that are b -inessential for a and denote the resulting call sequence by $R_{ab}(c)$.

Example 12 Consider the call sequence

$$bc.ce.df.ef.bh.\underline{af}.bg.\underline{ag}.\underline{ah}$$

from Example 10. The b -inessential calls are bh and bg and both lead to ag . So

$$R_{ab}(bc.ce.df.ef.bh.\underline{af}.bg.\underline{ag}.\underline{ah}) = bc.ce.df.ef.g.\underline{af}.\underline{ag}.\underline{ah}.$$

□

The following lemma establishes the relevant property of $R_{ab}(c)$.

Lemma 13 *Consider a call sequence c with no a -irrelevant calls that does not contain the call ab . Then $c \equiv_a R_{ab}(c)$.*

Proof. Since c does not contain the call ab , the views of a of c and $R_{ab}(c)$ have the same sequences of the a -calls.

By definition agent a does not learn the secret of b through the b -inessential calls for a . So the replacement (or possibly deletion) of b in all b -inessential calls has no effect on the status of this secret in the views of a of both sequences.

Let bc be the first b -inessential for a call in c and suppose that ad is the first a -call in c it leads to. Let c' be the outcome of the first step of producing $R_{ab}(c)$. So it is obtained from c by replacing bc by cd if $c \neq d$, and by deleting bc otherwise. We show that the views of the agent a of the call sequences c and c' are the same. First, note that no a -call earlier than ad can be effected by the change in c , because ad is the first a -call bc leads to.

Consider the $c \neq d$ case first. The set of secrets an agent is familiar with at the same point in c and c' may differ as a result of bc being replaced by cd . However, we argue that it is impossible for the agent a to notice this difference. Let us consider a call ax in c , where $x \in A$ and the sets of secrets, S and S' , agent e is familiar before ax is made in c and c' , respectively.

- If ax takes place before bc then S and S' are the same.
- If ax is in-between the calls bc and ad then again the sets S and S' are the same, because ad is the first a -call that bc leads to.
- If ax is the call ad then we have the following. First, just before this call a is already familiar with all the secrets b is familiar with before the call bc is made in c , because ad is the first a -call bc leads to. So a is still familiar with all these secrets in c' . Thus the only secrets that may be lost by replacing bc by cd are the ones that c is familiar with at that point. However, these secrets are passed to d and a still learns them all in c' through the call ad .
- If ax takes place after the call ad then the difference between S and S' could be at most in the set of secrets a learned through the call bc . However, a already knows these secrets after the call ad is made, so $S = S'$.¹

The reasoning for the $c = d$ case is completely analogous and omitted.

By iterating the above argument, starting with c , we obtain $R_{ab}(c)$ without affecting the view of the agent a . □

The following consequence of the above lemma is crucial. Here \equiv_{abab} stands for the the composition of the relations \equiv_a , \equiv_b , \equiv_a and \equiv_b , i.e.,

$$\equiv_{abab} = \equiv_a \circ \equiv_b \circ \equiv_a \circ \equiv_b,$$

where \circ is the composition of two binary relations.

Theorem 14 *Consider a call sequence c that does not contain the call ab . Then*

$$c \equiv_{abab} \mathcal{E}.$$

¹Note that this crucially depends on the fact that from any call each involved agent learns the union of the sets of secrets the callers are familiar with and not the set of secrets the other caller is familiar with.

Proof. Let c_1 be the a -simplification of c and $c_2 = R_{ab}(c_1)$. Then $c \equiv_a c_1$ and by Lemma 13 $c_1 \equiv_a c_2$, so $c \equiv_a c_2$.

If there are no b -calls in c_2 , then $c_2 \equiv_b \varepsilon$, so $c \equiv_{ab} \varepsilon$, from which the conclusion follows since $\varepsilon \equiv_{ab} \varepsilon$. Otherwise let c_3 be the prefix of c_2 that ends with the last b -call. Then $c_2 \equiv_b c_3$.

All the b -calls in c_3 are b -essential for a in the call sequence c_2 and the a -call through which agent a learns in c_2 the secret of B is located after all these b -essential calls. It follows that $c_3 \equiv_a c_4$, where c_4 is the result of removing all b -calls from c_3 , because no b -call in c_3 can possibly lead to an a -call in c_3 .

Now, c_4 does not contain any b -calls, so $c_4 \equiv_b \varepsilon$. This proves the claim. \square

Example 15 The following example illustrates the call sequences generated in the above proof. Let

$$c = \underline{ah}.cd.bc.bd.be.\underline{ad}.bf.bg.\underline{af}.$$

Then the a -simplification of c results in removing the calls cd and bg and equals

$$c_1 = \underline{ah}.bc.bd.be.\underline{ad}.bf.\underline{af}.$$

Subsequently $R_{ab}(c_1)$ equals

$$c_2 = \underline{ah}.bc.bd.ef.\underline{ad}.\underline{af}.$$

Next, the prefix of c_2 that ends with the last b -call is

$$c_3 = \underline{ah}.bc.bd.$$

Finally, the result of removing all b -calls from c_3 equals

$$c_4 = \underline{ah}$$

and $c_4 \equiv_b \varepsilon$. \square

This brings us to the following conclusion.

Theorem 16 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

$$(\mathcal{M}, c) \models K_{abab}\phi \text{ iff } \mathcal{M} \models \phi$$

and

$$(\mathcal{M}, c) \models K_{baba}\phi \text{ iff } \mathcal{M} \models \phi.$$

Proof. By symmetry it suffices to prove the first equivalence. By Theorem 14 $(\mathcal{M}, c) \models K_{abab}\phi$ implies $(\mathcal{M}, \varepsilon) \models \phi$, so by the Monotonicity Theorem 6 for all call sequences d we have $(\mathcal{M}, d) \models \phi$, i.e., $\mathcal{M} \models \phi$.

Further, for arbitrary call sequences c and formulas $\phi \in \mathcal{L}^{ck}$, $\mathcal{M} \models \phi$ implies $(\mathcal{M}, c) \models K_{abab}\phi$. \square

Corollary 17 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

$$(\mathcal{M}, c) \models C_{ab}\phi \text{ iff } (\mathcal{M}, c) \models K_{abab}\phi.$$

Proof. By Theorems 16 and 9. \square

This together with Note 5 shows that under the conditions of the above corollary the infinite conjunction $\bigwedge_{t \in \mathbb{G}^*} K_t \phi$ is equivalent to the 4th fold iterated knowledge $K_{abab}\phi$. To conclude this analysis we now show that common knowledge for two agents is not equivalent to any shorter knowledge iteration.

Corollary 18 *Let $c = ac, bc, ac$. Then $(\mathcal{M}, c) \models K_{aba}F_cA$, $(\mathcal{M}, c) \models K_{ab}F_cA$ and $(\mathcal{M}, c) \models K_aF_cA$. Further, $(\mathcal{M}, c) \not\models C_{ab}F_cA$.*

Proof. Take d such that $c \sim_a d$. Then d is of the form $c_1, ac, c_2, bc, c_3, ac, c_4$. Next take d_1 such that $d_1 \sim_b d$. Then d_1 is of the form c_5, ac, c_6, bc, c_7 . Next take d_2 such that $d_2 \sim_a d_1$. Then d_2 is of the form c_8, ac, c_9 . So $(\mathcal{M}, d_2) \models F_cA$, which implies the claim.

The next two claims follow since for all formulas $\phi \in \mathcal{L}^{ck}$, $K_{aba}\phi$ implies both $K_{ab}\phi$ and $K_a\phi$.

Now note that for all sequences $t \in \{a, b\}^*$ that extend $abab$ or $baba$ and all formulas $\phi \in \mathcal{L}^{ck}$, $K_t\phi$ implies $K_{abab}\phi$ or $K_{baba}\phi$. So Theorem 16 implies that for all such sequences s and all call sequences c that do not contain the call ab

$$(\mathcal{M}, c) \not\models K_sF_cA.$$

Indeed, the formula F_cA is not always true.

Further, by Note 5 we also have for such call sequences c

$$(\mathcal{M}, c) \not\models C_{ab}F_cA.$$

In particular, this holds for the above call sequence $c = ac, bc, ac$. □

6 Decidability issues

6.1 Decidability of semantics

We begin with the decidability of the semantics. First we establish some properties of the semantics of common knowledge of two agents.

Consider a call ab and a call sequence c . Starting from c we repeatedly remove from the current call sequence a redundant call that differs from ab . We call each outcome of such an iteration an ***ab-reduction*** of c . Further, we say that a call sequence c is ***ab-redundant free*** if no call from c that differs from ab is redundant in it. Clearly each *ab-reduction* is *ab-redundant free*.

Corollary 19 *Let d be an *ab-reduction* of c . Then*

- (i) $c \sim_{\{a,b\}} d$,
- (ii) for all formulas $\phi \in \mathcal{L}_{pr}$, $(\mathcal{M}, c) \models \phi$ iff $(\mathcal{M}, d) \models \phi$.

Lemma 20 *For each call ab and a call sequence c the set of *ab-redundant free* call sequences d such that $c \sim_{\{a,b\}} d$ is finite.*

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

Associate with d the sequence of gossip situations

$$d^0(\text{root}), d^1(\text{root}), \dots, d^m(\text{root}),$$

where m is the length of d , $d^0 = \varepsilon$, and $d^k = d_1, d_2, \dots, d_k$ for $k = 1, \dots, m$. This sequence monotonically grows, where we interpret the inclusion relation componentwise. Moreover, for all calls d_i different from ab the corresponding inclusion is strict. Consequently, m , the length of d , is bounded by $k + |A|^{|A|}$, the sum of the number of calls ab in 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. □

We can now prove the desired result.

Theorem 21 (Decidability of Semantics) *For each call sequence c it is decidable whether for a formula $\phi \in \mathcal{L}_{wn}$, $(\mathcal{M}, c) \models \phi$ holds.*

Proof. We use the definition of semantics as the algorithm. We only need to consider the case of the formulas of the form $C_G\phi$, where $\phi \in \mathcal{L}_{pr}$.

If $|G| = 1$, then this is the contents of the Decidability of Semantics Theorem 5 of [2].

If $|G| = 2$, say $A = \{a, b\}$, then according to Corollary 19 we can rewrite the semantics of $C_G\phi$ as follows:

$$(\mathcal{M}, c) \models C_G\phi \quad \text{iff} \quad \forall d \text{ s.t. } c \sim_G d \text{ and } d \text{ is } ab\text{-redundant free, } (\mathcal{M}, d) \models \phi,$$

and according to Lemma 20 this definition refers to a finite set of call sequences d .

If $|G| \geq 3$, then the decidability follows from Theorem 7 and the Decidability of Truth Theorem 6 established in [2]. \square

This result implies that the gossip protocols that use guards with non-nested common knowledge operator are implementable.

6.2 Decidability of truth

Next, we show that truth definition for the formulas of the language \mathcal{L}_{wn} is decidable. Since partial correctness of gossip protocols with common knowledge operator can be expressed as a formula of \mathcal{L}_{wn} , this implies that the problem of determining partial correctness of such protocols is decidable.

The key notion is that of an *epistemic pair-view*. It is a function of a call sequence c , denoted by $EPV(c)$, defined by

- putting for any pair of agents a, b :
 $EPV(c)(a, b) = \{d(\text{root}) \mid c \sim_{\{a,b\}} d\}$, and setting
- $EPV(c)(*) = c(\text{root})$.

So $EPV(c)(a, b)$ is the set of all gossip situations obtained by means of call sequences that are $\sim_{\{a,b\}}$ -equivalent to c . Further, as $\sim_{\{a,a\}} = \sim_a$, $EPV(c)(a, a)$ is the set of all gossip situations consistent with agent a 's observations made throughout c . Finally, $EPV(c)(*)$ is the actual gossip situation after c takes place. Note that for any $a, b \in A$, if $c \sim_{\{a,b\}} d$ then $EPV(c)(b, a) = EPV(c)(a, b) = EPV(d)(a, b) = EPV(d)(b, a)$.

The following holds.

Lemma 22 *For each call sequence c and agents a, b , the set $EPV(c)(a, b)$ is finite and can be effectively constructed.*

Proof. For any $a \in A$, $EPV(c)(a, a)$ coincides with the epistemic view $EV(c)(a)$, as defined in [2]. Hence, we can compute $EPV(c)(a, a)$ as in Lemma 3 of [2].

Consider now a pair of agents a, b such that $a \neq b$. To construct the set $EPV(c)(a, b)$ it suffices by Corollary 19 to consider the ab -redundant free call sequences d and by Lemma 20 there are only finitely many such call sequences d for which $d \sim_{\{a,b\}} c$. \square

Our interest in epistemic pair-views stems from the following important observation.

Lemma 23 *Suppose that $EPV(c) = EPV(d)$. Then for all formulas $\phi \in \mathcal{L}_{wn}$, $(\mathcal{M}, c) \models \phi$ iff $(\mathcal{M}, d) \models \phi$.*

Proof. A straightforward proof by induction shows that for a formula $\psi \in \mathcal{L}_{pr}$ and arbitrary call sequences c' and d' ,

$$c'(\text{root}) = d'(\text{root}) \text{ implies that } (\mathcal{M}, c') \models \psi \text{ iff } (\mathcal{M}, d') \models \psi. \quad (2)$$

Since $\text{EPV}(c)(*) = c(\text{root})$ and $\text{EPV}(d)(*) = d(\text{root})$, this settles the case for $\phi = F_a p$.

Next, consider the case of the formulas of the form $C_G \phi$, where $\phi \in \mathcal{L}_{wn}$.

If $|G| = 1$, then $G = \{a\}$ for some $a \in A$ and C_G is the same as K_a . Since $\text{EPV}(c) = \text{EPV}(d)$ implies $\text{EV}(c) = \text{EV}(d)$, where the epistemic view $\text{EV}()$ is defined as in [2], the claim follows by Lemma 4 of [2].

If $|G| = 2$, then $G = \{a, b\}$ for some $a, b \in A$. By (2) and the definition of $\text{EPV}(c)$

$$(\mathcal{M}, c) \models C_G \phi \quad \text{iff} \quad \forall c' \text{ s.t. } c'(\text{root}) \in \text{EPV}(c)(a, b), (\mathcal{M}, c') \models \phi.$$

So the claim follows since $\text{EPV}(c)(a, b) = \text{EPV}(d)(a, b)$.

If $|G| \geq 3$, then by Theorem 7 both $(\mathcal{M}, c') \models C_G \phi$ and $(\mathcal{M}, d') \models C_G \phi$ are equivalent to $\models \phi$.

This settles the case for $C_G \phi$. The remaining cases of negation and conjunction follow directly by the induction. \square

The above lemma is useful because the epistemic pair-view of each call sequence is finite, in contrast to the set of call sequences. Next, we provide an inductive definition of $\text{EPV}(c.c)(a, b)$ the importance of which will become clear in a moment.

Lemma 24 *For any call sequence c and call $c = ab$ for agents $a, b \in A$*

$$\text{EPV}(c.c)(a, b) = \{c(s) \mid s \in \text{EPV}(c)(a, b), c(s)_a = c(c(\text{root}))_a \text{ and } c(s)_b = c(c(\text{root}))_b\}.$$

Proof.

(\subseteq) Take $s' \in \text{EPV}(c.c)(a, b)$. By the definition of $\text{EPV}(c.c)(a, b)$ there exists a call sequence d such that $d.c \sim_{\{a, b\}} c.c$ and $s' = d.c(\text{root})$. So $s' = c(s)$, where $s = d(\text{root})$. We prove now that $d \sim_{\{a, b\}} c$ and as a result $s = d(\text{root}) \in \text{EPV}(c)(a, b)$.

Note that $d.c \sim_{\{a, b\}} c.c$ implies that there exists a sequence $t_1 \dots t_k \in \{a, b\}^*$ such that for some call sequences d_1, d_2, \dots, d_{k-1} we have $d.c \sim_{t_1} d_1.c \sim_{t_2} d_2.c \sim_{t_3} \dots \sim_{t_{k-1}} d_{k-1}.c \sim_{t_k} c.c$. Note that for any $t \in \{a, b\}$ and call sequences c', d' we have that $c'.c \sim_t d'.c$ implies $c' \sim_t d'$, because \sim_t is the minimal relation satisfying the conditions stated in Definition 1. It follows that $d \sim_{t_1} d_1 \sim_{t_2} d_2 \sim_{t_3} \dots \sim_{t_{k-1}} d_{k-1} \sim_{t_k} c$, so by definition $d \sim_{\{a, b\}} c$.

Further, by the definition of the \sim_c relations, we also have that for $i \in \{0, \dots, k-1\}$ both $d_i.c(\text{root})_a = d_{i+1}.c(\text{root})_a$ and $d_i.c(\text{root})_b = d_{i+1}.c(\text{root})_b$, where $d_0 = d$ and $d_k = c$. So $d.c(\text{root})_a = c.c(\text{root})_a$ and $d.c(\text{root})_b = c.c(\text{root})_b$.

But $s = d(\text{root})$, so we get that $c(s)_a = c(c(\text{root}))_a$ and $c(s)_b = c(c(\text{root}))_b$.

(\supseteq) Take $s' \in \{c(s) \mid s \in \text{EPV}(c)(a, b), c(s)_a = c(c(\text{root}))_a \text{ and } c(s)_b = c(c(\text{root}))_b\}$. So for some gossip situation s we have $s' = c(s)$, $s \in \text{EPV}(c)(a, b)$, $c(s)_a = c(c(\text{root}))_a$, and $c(s)_b = c(c(\text{root}))_b$. The fact that $s \in \text{EPV}(c)(a, b)$ implies that there exists a call sequence d such that $d \sim_{\{a, b\}} c$ and $s = d(\text{root})$. Now, this and $c(d(\text{root}))_a = c(s)_a = c(c(\text{root}))_a$ imply by definition that $d.c \sim_a c.c$, so a fortiori $d.c \sim_{\{a, b\}} c.c$. So $d.c(\text{root}) \in \text{EPV}(c.c)(a, b)$. Consequently also $s' \in \text{EPV}(c.c)(a, b)$, because $s' = c(s) = d.c(\text{root})$. \square

This allows us to conclude that $\text{EPV}(c.c)$ can be computed using $\text{EPV}(c)$ and c only, i.e., without referring to c . More precisely, denote the set of epistemic pair-views by $\widetilde{\text{EPV}}$ and recall that C denotes the set of calls. Then the following holds.

Corollary 25 *There exists a function $f : \widetilde{\text{EPV}} \times C \rightarrow \widetilde{\text{EPV}}$ such that for any call sequence c , call c , and pair of agents $a, b \in A$*

$$\text{EPV}(c.c)(a, b) = f(\text{EPV}(c), c).$$

Proof. First note that $\text{EPV}(c.c)(*) = c(\text{EPV}(c)(*))$.

Suppose now that $a \neq b$. If $c = ab$, then by Lemma 24 $\text{EPV}(c.c)(a,b)$ is a function of $\text{EPV}(c)(a,b)$ and c . If $c \neq ab$, say $a \notin c$, then $c.c \sim_a c$ and hence $c.c \sim_{\{a,b\}} c$, which implies $\text{EPV}(c.c)(a,b) = \text{EPV}(c)(a,b)$.

Suppose next that $a = b$. By the definition of \sim_a for all d we have $\text{EPV}(d)(a,a) = \text{EV}(d)(a)$, so by Corollary 2 of [2] $\text{EPV}(c.c)(a,a)$ is a function of $\text{EPV}(c)(a,a)$ and c . \square

Consider a call sequence c . If for some prefix $c_1.c_2$ of c , we have $\text{EPV}(c_1) = \text{EPV}(c_1.c_2)$, then we say that the call subsequence c_2 is **pair-epistemically redundant** in c and that c is **pair-epistemically redundant**.

We say that c is **pair-epistemically non-redundant** if it is not pair-epistemically redundant. Equivalently, a call sequence $c_1.c_2 \dots .c_k$ is pair-epistemically non-redundant if the set

$$\{\text{EPV}(c_1.c_2 \dots .c_i) \mid i \in \{1, \dots, k\}\}$$

has k elements.

Lemma 26 (Pair-Epistemic Stuttering) *Suppose that $c := c_1.c_2.c_3$ and $d := c_1.c_3$, where c_2 is pair-epistemically redundant in c . Then $\text{EPV}(c) = \text{EPV}(d)$.*

Proof. Let $c_3 = c_1.c_2 \dots .c_k$. First note that thanks to Corollary 2 of [2] we have $\text{EPV}(c_1.c_2.c_1) = \text{EPV}(c_1.c_1)$, since $\text{EPV}(c_1.c_2.c_1) = f(\text{EPV}(c_1.c_2), c_1) = f(\text{EPV}(c_1), c_1) = \text{EPV}(c_1.c_1)$ due to the pair-epistemic redundancy of c_2 in c . Repeating this argument for all $i \in \{1, \dots, k\}$ we get that

$$\text{EPV}(c_1.c_2.c_1.c_2 \dots .c_i) = \text{EPV}(c_1.c_1.c_2 \dots .c_i).$$

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

Corollary 27 *For every call sequence c there exists a pair-epistemically non-redundant call sequence d such that for all formulas $\phi \in \mathcal{L}_{wn}$, $(\mathcal{M}, c) \models \phi$ iff $(\mathcal{M}, d) \models \phi$.*

Proof. By the repeated use of the Pair-Epistemic Stuttering Lemma 26 and Lemma 23. \square

Next, we prove the following crucial lemma.

Lemma 28 *For any given model \mathcal{M} , there are only finitely many pair-epistemically non-redundant call sequences.*

Proof. Note that each epistemic pair-view is a function from $A \times A \cup \{*\}$ to the set of functions from A to $2^{|P|}$ (this is an overestimation because for $*$ this set has only one element). There are $k = 2^{(|A|^2+1) \cdot 2^{|A| \cdot |P|}}$ such functions, so any call sequence longer than k has a pair-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 29 (Decidability of Truth) *For any formula $\phi \in \mathcal{L}_{wn}$, it is decidable whether $\mathcal{M} \models \phi$ holds.*

Proof. Recall that $\mathcal{M} \models \phi$ iff $\forall c (\mathcal{M}, c) \models \phi$. By Corollary 27 we can rewrite the latter as

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

But according to Lemma 28 there are only finitely many pair-epistemically non-redundant call sequences and by Lemma 23 their set can be explicitly constructed. \square

6.3 Decidability of termination with common knowledge operator

Finally, we show that it is decidable to determine whether a gossip protocol that uses guard with non-nested common knowledge operator (in short: a common knowledge protocol) terminates. For an example of such a protocol see Appendix A.

First, we establish monotonicity of gossip situations and epistemic pair-views with respect to call sequence extensions, w.r.t. suitable partial orderings. Intuitively, we claim that as the call sequence gets longer each agent acquires more information.

Definition 30 For any two gossip situations s, s' we write $s \leq_A s'$ if for all $a \in A$ we have $s_a \subseteq s'_a$.

Note 31 (Note 1 of [2]) For all call sequences c and d such that $c \sqsubseteq d$ we have $c(\text{root}) \leq_A d(\text{root})$.

Definition 32 For any two epistemic pair-views $V, V' \in \widetilde{\text{EPV}}$ we write $V \leq_{\text{EPV}} V'$ if for all $a, b \in A$ there exists $X \subseteq V(a, b)$ and an surjective (onto) function $g : X \rightarrow V'(a, b)$ such that for all $s \in X$ we have $s \leq_A g(s)$.

Lemma 33 \leq_{EPV} is a partial order.

Proof.

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

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

(Antisymmetry) Suppose V, V' are two epistemic pair-views such that $V \leq_{\text{EPV}} V'$ and $V' \leq_{\text{EPV}} V$. Then, from the definition of \leq_{EPV} , for any $a, b \in A$ there exist $X \subseteq V(a, b)$, $Y \subseteq V'(a, b)$, and surjective functions $g : X \rightarrow V'(a, b)$ and $h : Y \rightarrow V(a, b)$. Let $Z = \{s \in X \mid g(s) \in Y\}$. Note that $g|_Z : Z \rightarrow Y$, i.e. the restriction of g to Z , is surjective. Moreover, $g|_Z \circ h : Z \rightarrow V(a, b)$ is also surjective, and because $Z \subseteq V(a, b)$ is finite, $Z = V(a, b)$, $g|_Z = g$, and $g \circ h$ is a permutation on $V(a, b)$. Similarly we can show that $Y = V'(a, b)$. Since $(g \circ h)$ is a permutation on a finite set, there exists k such that $(g \circ h)^k$ is the identity function on $V(a, b)$.

Note that for any $s \in V(a, b)$, we have $s \leq_A (g \circ h)(s)$, because $s \leq_A g(s) \leq_A h(g(s))$. Now consider the sequence: $s \leq_A (g \circ h)(s) \leq_A (g \circ h)^2(s) \leq_A \dots \leq_A (g \circ h)^k(s) = s$. In fact, all of the elements in this sequence have to be the same, because \leq_A is a partial order. In particular, this shows that $(g \circ h)(s) = s$. Therefore, $g \circ h$ is the identity function on $V(a, b)$. Now, for any $s \in V(a, b)$ we have that $s \leq_A g(s) \leq_A h(g(s)) = (g \circ h)(s) = s$, so g is the identity function as well. This shows that $V(a, b) = V'(a, b)$ for all $a, b \in A$. \square

The next lemma formalizes the intuition that information captured by the epistemic pair-view grows along a call sequence.

Lemma 34 For all two call sequences such that $c \sqsubseteq d$ we have $\text{EPV}(c) \leq_{\text{EPV}} \text{EPV}(d)$.

Proof. Let $d = c.c'$. Take $a, b \in A$. By a repeated application of Lemma 24 we get $\text{EPV}(c.c')(a, b) = \{c'(s) \mid s \in \text{EPV}(c)(a, b) \text{ and } \forall c'' \sqsubseteq c' (c''(s)_a = c''(c(\text{root}))_a \wedge c''(s)_b = c''(c(\text{root}))_b)\}$. It suffices then to pick $X = \{s \in \text{EPV}(c)(a, b) \mid \forall c'' \sqsubseteq c' (c''(s)_a = c''(c(\text{root}))_a \wedge c''(s)_b = c''(c(\text{root}))_b)\}$ and set $g(s) =$

$c'(s)$ for all $s \in X$. It is easy to check that such $g : X \rightarrow \text{EPV}(d)$ is surjective, so $\text{EPV}(c) \leq_{\text{EPV}} \text{EPV}(d)$, as claimed. \square

We can now draw the following useful conclusion.

Lemma 35 *Suppose that c is pair-epistemically redundant. Then a prefix $c_1.c$ of it exists such that c_1 is pair-epistemically non-redundant and $\text{EPV}(c_1.c) = \text{EPV}(c_1)$.*

Proof. Let $c_1.c_2$ be the shortest prefix of c such that $\text{EPV}(c_1) = \text{EPV}(c_1.c_2)$. Then c_1 is pair-epistemically non-redundant. Let $c_2 = c_1 \dots c_l$. By Lemma 34 we have

$$\begin{aligned} \text{EPV}(c_1) &\leq_{\text{EPV}} \text{EPV}(c_1.c_1) \leq_{\text{EPV}} \text{EPV}(c_1.c_1.c_2) \leq_{\text{EPV}} \dots \leq_{\text{EPV}} \\ \text{EPV}(c_1.c_1.c_2 \dots c_l) &= \text{EPV}(c_1.c_2) = \text{EPV}(c_1). \end{aligned}$$

Since \leq_{EPV} is a partial order, $\text{EPV}(c_1.c_1) = \text{EPV}(c_1)$ holds. \square

Finally we can establish the desired result. In the proof we shall use the following observation.

Theorem 36 (Stuttering) *Suppose that $c := c_1.c.c_2$ and $d := c_1.c.c.c_2$. Then for all formulas $\phi \in \mathcal{L}^{ck}$, $(\mathcal{M}, c) \models \phi$ iff $(\mathcal{M}, d) \models \phi$.*

Proof. This is a direct consequence of the corresponding Stuttering Theorem 3 from [2] and Note 5. \square

Theorem 37 (Decidability of Termination) *Given a common knowledge 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 $c.c$ such that c is pair-epistemically non-redundant and $\text{EPV}(c.c) = \text{EPV}(c)$.

(\Rightarrow) Let \bar{c} be an infinite sequence of calls generated by the protocol. There are only finitely many pair-epistemic views, so some prefix c of \bar{c} is pair-epistemically redundant. The claim now follows by Lemma 35.

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

Let ϕ be the guard associated with the call c , i.e., $\phi \rightarrow c$ is a rule used in the considered protocol. By assumption $(\mathcal{M}, c) \models \phi$, so by Lemma 23 $(\mathcal{M}, c.c) \models \phi$. By the repeated application of the Stuttering Theorem 36 we get that for all $i \geq 1$, $(\mathcal{M}, c.c^i) \models \phi$. Consequently, $c.c^\omega$ 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 $c.c$ such that c is pair-epistemically non-redundant and $\text{EPV}(c.c) = \text{EPV}(c)$.

But given a call sequence, by the Decidability of Semantics Theorem 21, it is decidable to determine whether it can be generated by the protocol and by Lemma 22 it is decidable to determine whether a call sequence is pair-epistemically non-redundant. Further, by Lemma 28 there are only finitely many pair-epistemically non-redundant call sequences, so the claim follows. \square

7 Conclusions

We studied here various aspects of common knowledge in the context of a natural epistemic logic used to express and reason about distributed epistemic gossip protocols. We showed that the semantics and truth in this logic are decidable in the absence of nested modalities. The first result implies that the gossip

protocols relying on such a use of the common knowledge operator are implementable and the second one that their partial correctness is decidable, since partial correctness of these gossip protocols can be expressed as a formula of the considered language. Further, we proved that the termination of these gossip protocols is decidable, as well.

There are a number of interesting open problems related to this work. An obvious question is whether our results can be extended to formulas that admit nested modalities.

In Corollary 17 we showed that under certain conditions common knowledge for two agents is equivalent to the 4th fold iterated knowledge. An intriguing question is whether this result holds for arbitrary call sequences and arbitrary formulas. If not, is then common knowledge always equivalent to some finite iterated knowledge?

Finally, it would be interesting to clarify which formulas two agents can commonly know, i.e., given a call sequence c to characterize the formulas ϕ for which $(\mathcal{M}, c) \models C_{ab}\phi$ holds. Example 8 indicates that this problem is non-trivial even without any call being performed.

Acknowledgments

We thank the referees for helpful comments. First author he was partially supported by NCN grant 2014/13/B/ST6/01807. The second author was partially supported by EPSRC grants EP/M027287/1 and EP/P020909/1.

References

- [1] K. R. Apt, D. Grossi & W. 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.
- [2] K. R. Apt & D. 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.
- [3] K.R. Apt, E. Kopczyński & D. Wojtczak (2017): *On the Computational Complexity of Gossip Protocols*. In: *Proceedings of 26th IJCAI*. To appear.
- [4] K.R. Apt & D. Wojtczak (2017): *Decidability of Fair Termination of Gossip Protocols*. In: *Proceedings of the IWIL Workshop and LPAR Short Presentations*, Kalpa Publications, pp. 73–85.
- [5] M. Attamah, H. van Ditmarsch, D. Grossi & W. 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.
- [6] M. Attamah, H. van Ditmarsch, D. Grossi & W. 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.
- [7] J. van Benthem, J. van Eijck & B. Kooi (2005): *Common Knowledge in Update Logics*. In: *Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2005)*, pp. 253–261, doi:10.1145/1089933.1089960.
- [8] M.C. Cooper, A. Herzig, F. Maffre, F. Maris & P. Régnier (2016): *A simple account of multi-agent epistemic planning*. In: *Proceedings of ECAI 2016*, IOS Press, pp. 193–201, doi:10.3233/978-1-61499-672-9-193.
- [9] M.C. Cooper, A. Herzig, F. Maffre, F. Maris & P. Regnier (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.

- [10] H. van Ditmarsch, J. van Eijck, P. Pardo, R. Ramezani & F. Schwarzentruber (2017): *Epistemic Protocols for Dynamic Gossip*. *J. of Applied Logic* 20(C), pp. 1–31, doi:10.1016/j.jal.2016.12.001.
- [11] H. van Ditmarsch, J. van Eijck & R. Verbrugge (2009): *Common Knowledge and Common Belief*. In J. van Eijck & R. Verbrugge, editors: *Discourses on Social Software*, Amsterdam University Press, pp. 99–122.
- [12] H. van Ditmarsch, D. Grossi, A. Herzig, W. van der Hoek & L.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>.
- [13] H. van Ditmarsch, W. van der Hoek & B. Kooi (2007): *Dynamic Epistemic Logic*. *Synthese Library* 337, Springer, doi:10.1007/978-1-4020-5839-4.
- [14] R. Fagin, J. Halpern, M. Vardi & Y. Moses (1995): *Reasoning about knowledge*. MIT Press, Cambridge, Massachusetts.
- [15] M.F. Friedell (1969): *On the structure of shared awareness*. *Behavioral Science* 14(1), pp. 28–39, doi:10.1002/bs.3830140105.
- [16] S.M. Hedetniemi, S.T. Hedetniemi & A.L. Liestman (1988): *A survey of gossiping and broadcasting in communication networks*. *Networks* 18(4), pp. 319–349, doi:10.1002/net.3230180406.
- [17] A. Herzig & F. Maffre (2017): *How to Share Knowledge by Gossiping*. *AI Communications* 30(1), pp. 1–17, doi:10.3233/AIC-170723. Available at <http://content.iospress.com/articles/ai-communications/aic723>.
- [18] J. Hromkovic, R. Klasing, A. Pelc, P. Ruzicka & W. 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.
- [19] D.K. Lewis (1969): *Convention, a Philosophical Study*. Harvard University Press, Cambridge (MA).
- [20] P. Vanderschraaf & G. Sillari (2014): *Common Knowledge*. In: *The Stanford Encyclopedia of Philosophy*. Available at <https://plato.stanford.edu/entries/common-knowledge>.
- [21] Y. Wang, L. Kuppusamy & J. van Eijck (2009): *Verifying epistemic protocols under common knowledge*. In: *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 2009)*, ACM, pp. 257–266, doi:10.1145/1562814.1562848.

A Example: a common knowledge protocol

To illustrate gossip protocols that employ the common knowledge operator assume that the agents are nodes of an undirected connected graph (V, E) and that the calls can take place only between pairs of agents connected by an edge. Let N_i denote the set of neighbours of node i .

Consider a gossip protocol with the following program for agent i (we use here the syntax introduced in [1]):

$$*[\bigwedge_{j \in N_i, B \in \mathcal{P}} F_i B \wedge \neg C_{ij} F_j B \rightarrow (i, j)].$$

Informally, agent i calls a neighbour j if i is familiar with some secret (here B) and there is no common knowledge between i and j that j is familiar with this secret.

Partial correctness of a protocol states that upon its termination the formula $\bigwedge_{i,j \in A} F_i J$ holds.

To prove partial correctness of the above protocol consider the exit condition

$$\bigwedge_{(i,j) \in E} \bigwedge_{B \in \mathcal{P}} (F_i B \rightarrow C_{ij} F_j B).$$

For all agents i and j and secrets B , the formula $C_{ij}F_jB \rightarrow F_jB$ is true, so the exit condition implies

$$\bigwedge_{(i,j) \in E} \bigwedge_{B \in \mathcal{P}} (F_iB \rightarrow F_jB).$$

Consider now an agent i and the secret J of agent j . Let $j = i_1, \dots, i_h = i$ be a path that connects j with i . The above formula implies that for $g \in \{1, \dots, h-1\}$ we have $\bigwedge_{B \in \mathcal{P}} (F_{i_g}B \rightarrow F_{i_{g+1}}B)$. By combining these $h-1$ formulas we get $\bigwedge_{B \in \mathcal{P}} (F_jB \rightarrow F_iB)$. But F_jJ is true, so we conclude F_iJ . Consequently $\bigwedge_{i,j \in \mathcal{A}} F_iJ$, as desired.

To prove termination we need the following observation.

Lemma 38 *For all call sequences $c.(i, j)$ and secrets B*

$$(\mathcal{M}, c.(i, j)) \models F_iB \wedge F_jB \text{ implies } (\mathcal{M}, c.(i, j)) \models C_{ij}(F_iB \wedge F_jB).$$

Proof. Suppose that $(\mathcal{M}, c.(i, j)) \models F_iB \wedge F_jB$. Take some d such that $c.(i, j) \sim_{\{i,j\}} d$. So there exists a sequence of call sequences c_1, \dots, c_k such that

$$c.(i, j) = c_1 \sim_{m_1} c_2 \sim_{m_2} \dots \sim_{m_k} c_k = d,$$

where each m_h is i or j .

By the repeated use of the Equivalence Theorem 4 each of the call sequences c_1, \dots, c_k contains the call (i, j) that corresponds with the last call of $c.(i, j)$. Let $d_1.(i, j), \dots, d_k.(i, j)$ be the corresponding prefixes of c_1, \dots, c_k .

By the Equivalence Theorem 4 and the definition of a view given there we also have

$$d_1.(i, j)_i = d_1.(i, j)_j = d_2.(i, j)_i = d_2.(i, j)_j.$$

By the assumption this implies $(\mathcal{M}, d_2.(i, j)) \models F_iB \wedge F_jB$, since $d_1.(i, j) = c.(i, j)$.

Repeating this procedure we conclude that $(\mathcal{M}, d_k.(i, j)) \models F_iB \wedge F_jB$ and hence $(\mathcal{M}, d) \models F_iB \wedge F_jB$. This implies the claim. \square

Now, by the definition of semantics for all call sequences $c.(i, j)$ and secrets B , $(\mathcal{M}, c) \models F_iB$ implies $(\mathcal{M}, c.(i, j)) \models F_iB \wedge F_jB$, which implies by Lemma 38 that $(\mathcal{M}, c.(i, j)) \models C_{ij}(F_iB \wedge F_jB)$ and hence $(\mathcal{M}, c.(i, j)) \models C_{ij}F_jB$. This shows that after each call (i, j) the size of the set $\{(i, j, B) \mid \neg C_{ij}F_jB\}$ decreases.