Temporal Epistemic Gossip Problems

Martin Cooper, Andreas Herzig, Frédéric Maris, Julien Vianey

Official URL

DOI : https://doi.org/10.1007/978-3-030-14174-5_1

Temporal Epistemic Gossip Problems

Martin C. Cooper(✉), Andreas Herzig, Frédéric Maris, and Julien Vianey

IRIT, CNRS, Univ. Toulouse, 31062 Toulouse Cedex 9, France
{martin.cooper,andreas.herzig,frederic.maris,julien.vianey}@irit.fr

Abstract. Gossip problems are planning problems where several agents have to share information ('secrets') by means of phone calls between two agents. In epistemic gossip problems the goal can be to achieve higher-order knowledge, i.e., knowledge about other agents' knowledge; to that end, in a call agents communicate not only secrets, but also agents' knowledge of secrets, agents' knowledge about other agents' knowledge about secrets, etc. Temporal epistemic gossip problems moreover impose constraints on the times of calls. These constraints are of two kinds: either they stipulate that a call between two agents must necessarily be made at some time point, or they stipulate that a call can be made within some possible (set of) interval(s). In the non-temporal version, calls between two agents are either always possible or always impossible. We investigate the complexity of the plan existence problem in this gen-eral setting. Concerning the upper bound, we prove that it is in NP in the general case, and that it is in P when the problem is non-temporal and the goal is a positive epistemic formula. As for the lower bound, we prove NP-completeness for two fragments: problems with possibly neg-ative goals even in the non-temporal case, and problems with temporal constraints even if the goal is a set of positive atoms.

Keywords: Epistemic planning · Temporal planning · Gossip problem · Complexity · Epistemic logic

1 Introduction

The epistemic gossip problem defined in [11,12,21] is a problem in which n agents each have a secret. Agents communicate by calling other agents: during a call the two agents share all their knowledge, not only the secrets they have learned but also epistemic information concerning which agents know which information. The goal of this problem concerns agents' knowledge about other agents' secrets at various epistemic depths. For example, the goal may be shared knowledge of depth 2: all agents know that all agents know all secrets. Such goals can be described as logical formulas in Dynamic Epistemic Logic of Propositional Assignments and Observation DEL-PAO [17,18]. This setting generalises the well-known gossip problem [12,16] which has recently been analysed in the framework of dynamic epistemic logics [6]. We consider it to be an exemplary

https://doi.org/10.1007/978-3-030-14174-5_1

case of epistemic planning [7,8,20,22] in which communication actions are used to spread knowledge among a network of agents.

Here, we enrich this setting by allowing to limit any communication to a set of instants (such as an interval during which the two agents involved in the call are both available). For an example, one can think of satellites on different orbits which can only communicate when they 'see' each other. In a more down-to-earth example, the interval during which a mobile communication is available is often limited by the charge capacity of the battery. Another variant we study is when certain calls must occur at given instants (for example for maintenance or security reasons).

What follows applies to either one-way or two-way communication and to either sequential or parallel communication. During a one-way call (such as a letter or email) information only passes in one direction, whereas during a two-way call (such as a telephone conversation) information passes in both directions. In the case of parallel communication, several calls between distinct pairs of agents may take place simultaneously, but an agent can only call one other agent at the same instant. The sequential version, in which only one call can take place at the same time, is of interest when the aim is to minimise the total number of calls.

We show that the temporal epistemic gossip problem is in NP even for a complex goal given in the form of a CNF and in the presence of constraints on the instants when calls can or must take place. This positive result, when compared to classical planning which is PSPACE-complete [9], follows from the reasonable assumption that knowledge is never destroyed. Moreover, we show that in the absence of temporal constraints and negative goals, temporal epistemic gossiping is in P. We then show maximality of this tractable subproblem in the sense that the problem becomes NP-complete in the following cases: in the presence of temporal constraints (even as weak as a simple upper bound on the execution time of a plan) and in the presence of negative goals (such as agent i should not learn the secret of agent j).

2 Definitions

First of all, we introduce a general framework for epistemic planning before focussing on the specific subproblem of epistemic gossiping.

Let *Prop* be a countable set of *propositional variables* and *Agt* a finite set of *agents*. A *knowledge operator* is of the form K_i with $i \in Agt$. An *atom of depth d* is any sequence of knowledge operators K_i of length d followed by a propositional variable. (So when the depth is 0 then the atom is just a propositional variable.) Atoms are noted α, α', etc. The atom $K_i p$ is of depth 1 and reads "agent i knows that p"; the atom $K_j K_i p$ is of depth 2 and reads "j knows that i knows that p"; and so on. The set of all atoms of depth at most d is noted $ATM^{\leq d}$. Observe that if the depth of atom $\alpha \in ATM^{\leq d}$ is strictly less than d then $K_i\alpha$ also belongs to $ATM^{\leq d}$. The set of atoms that are about the mutual knowledge of α by agents i and j up to depth d is:

$$ATM_{i,j}^{\leq d}(\alpha) = \{K_{k_1} \dots K_{k_r}\alpha \mid k_1, \dots, k_r \in \{i, j\} \text{ and } r + depth(\alpha) \leq d\}$$

Finally, the set of *boolean formulas* Fml_{bool} is comprised of formulas with the following grammar, where $\alpha \in ATM^{\leq d}$:

$$\varphi ::= \alpha \mid \neg\varphi \mid (\varphi \wedge \varphi)$$

A *state* is an assignment of truth values to all atoms in $ATM^{\leq d}$ and is represented by the set of atoms which are assigned the value true. Satisfaction of a boolean formula ϕ in a state s, noted $s \models \phi$, is defined in the usual way.

A *conditional action* is a pair $\mathsf{a} = \langle pre(\mathsf{a}), \mathit{eff}(\mathsf{a})\rangle$ where:

- $pre(\mathsf{a}) \in Fml_{bool}$ is a boolean formula: the *precondition* of a;
- $\mathit{eff}(\mathsf{a}) \subseteq Fml_{bool} \times 2^{ATM^{\leq d}} \times 2^{ATM^{\leq d}}$ is a set of triples ce of the form

  $$\langle cnd(ce), \mathit{ceff}^+(ce), \mathit{ceff}^-(ce)\rangle,$$

  the *conditional effects* of a, where $cnd(ce)$ is a boolean formula (the condition) and $\mathit{ceff}^+(ce)$ and $\mathit{ceff}^-(ce)$ are sets of atoms (added and deleted atoms respectively).

The result of executing action a in a state s is the state $(s \cup e^+) \setminus e^-$, where $e^+ = \bigcup_{ce \in \mathit{eff}(\mathsf{a}), s \models cnd(ce)} \mathit{ceff}^+(ce)$ and $e^- = \bigcup_{ce \in \mathit{eff}(\mathsf{a}), s \models cnd(ce)} \mathit{ceff}^-(ce)$.

In the case of the gossip problem, $Prop = \{s_{ij} \mid i, j \in Agt\}$ where s_{ij} reads "i knows j's secret". So $K_k s_{ij}$ means that agent k knows that i knows j's secret. The actions in the gossip problem are calls between two agents leading to an update of the two agents' knowledge. In one-way calls, after a call from agent i to agent j, agent j knows everything that agent i knew before the call and both know that they know these atoms. Indeed, they both know that they both know that they know these atoms, and so on up to the maximum epistemic depth d.[1] More formally, $call(i,j) = \langle pre(call(i,j)), \mathit{eff}(call(i,j))\rangle$ with $pre(call(i,j)) = \top$ and $\mathit{eff}(call(i,j)) = \{\langle K_i\alpha, ATM_{ij}^{\leq d}(\alpha), \emptyset\rangle \mid \alpha \in ATM^{\leq d}\} \cup \{\langle \top, ATM_{ij}^{\leq d}(s_{ji}), \emptyset\rangle\}$. Two-way calls have the same effect as two simultaneous one-way calls.

An instance of the depth-d temporal epistemic gossip problem (TEGP) is given by a tuple $\Pi = \langle Init, Goal, Agt, I_p, I_n\rangle$:

$$\begin{aligned}
Init &\subseteq ATM^{\leq d} \text{ such that } Init \text{ contains every } s_{ii}, \text{ for } i \in Agt\\
Goal &\in Fml_{bool} \text{ is a conjunction of clauses}\\
I_p &\subseteq \mathbb{N} \times (\mathbb{N} \cup \{\infty\}) \times Agt \times Agt\\
I_n &\subseteq \mathbb{N} \times Agt \times Agt
\end{aligned}$$

where *Init* is the initial state; *Goal* is the goal we want to achieve in the form of a CNF formula (that we identify with a set of clauses); I_p is the set of intervals

[1] More generally, the caller's knowledge becomes common knowledge between i and j. We however have no common knowledge operator in our framework.

during which two agents can call each other and I_n is the set of instants when two agents must call each other. The set I_n of necessary calls may correspond to calls that have been programmed in the network for some other purpose. We suppose that I_n is included in I_p, in the sense that for every $\langle t, i, j\rangle \in I_n$ there is a $\langle t_1, t_2, i, j\rangle \in I_p$ such that $t_1 \leq t \leq t_2$. In this paper we always consider the initial state $\mathit{Init} = \{s_{ii} \mid i \in \mathit{Agt}\}$ in which all agents know their own secrets.

A set of calls A between agents induces a partial function between states, i.e. from $2^{ATM^{\leq d}}$ to $2^{ATM^{\leq d}}$. For a state $s \in 2^{ATM^{\leq d}}$:

$$\mathsf{A}(s) = \begin{cases} \bot & \text{if } \exists \mathsf{a} \in \mathsf{A} : s \not\models \mathit{pre}(\mathsf{a}), \text{ or } \ \exists \mathsf{a}_1, \mathsf{a}_2 \in \mathsf{A} : \\ & \mathsf{a}_1 = \mathit{call}(i_1, j_1) \text{ and } \mathsf{a}_2 = \mathit{call}(i_2, j_2) \\ & \quad \text{with } \{i_1, j_1\} \cap \{i_2, j_2\} \neq \emptyset \\ s \cup \bigcup\limits_{\substack{\mathsf{a} \in \mathsf{A}, \\ ce \in \mathit{eff}(\mathsf{a}), \\ \text{and } s \models \mathit{cnd}(ce)}} \mathit{ceff}^{+}(ce) & \text{otherwise} \end{cases}$$

Note that $\bot$ is not a state: a result $\bot$ represents a failure of the simultaneous execution of the set of actions A, either because a precondition does not hold or because one agent would be participating in two calls at the same time (which we assume to be impossible).

A *plan* is a relation $P \subseteq \mathbb{N} \times \mathit{Agt} \times \mathit{Agt}$. Given a plan P and a natural number t, the set of calls happening at instant t is $P(t) = \{(i, j) : (t, i, j) \in P\}$. We use $|P|$ to denote the number of distinct instants t for which $P(t) \neq \emptyset$. We use $T_P(k)$ to denote the k-th instant (in strictly increasing order of time) at which a call happens in P: i.e. $T_P(1) < \ldots < T_P(|P|)$ and $\forall t,\ P(t) \neq \emptyset \Leftrightarrow \exists k \in \{1, \ldots, |P|\}$, $T(k) = t$. Our modelling of time by the natural numbers implicitly imposes a fixed duration of one time unit for each call.

Given a TEGP $\mathit{\Pi} = \langle \mathit{Init}, \mathit{Goal}, \mathit{Agt}, I_p, I_n\rangle$, a plan P *satisfies the temporal constraints* of $\mathit{\Pi}$ if and only if all the necessary calls are in P and every call in P is possible; formally: $I_n \subseteq P$ and for every $\langle t, i, j\rangle \in P$ there is a $\langle t_1, t_2, i, j\rangle \in I_p$ such that $t_1 \leq t \leq t_2$. Moreover, P *solves* the TEGP if and only if it satisfies the temporal constraints and there is a sequence of states $\langle s_0, \cdots, s_{|P|}\rangle$ such that

- $s_0 = \mathit{Init}$
- $s_{|P|} \models \mathit{Goal}$
- $s_{k+1} = P(T_P(k+1))(s_k)$ for every k with $0 \leq k < |P|$

where $P(T_p(k+1))$ is the set of actions at instant $T_p(k+1)$ and $P(T_p(k+1))(s_k)$ is the result of executing these actions in state s_k. By the definition above of $\mathsf{A}(s)$, the set of actions $P(T_p(k+1))$ at instant $T_p(k+1)$ cannot contain two calls involving the same agent. In the sequential version of the TEGP, a solution plan P must also satisfy $\forall t, \mathit{card}(P(t)) \leq 1$.

A TEGP defines in a natural way a call digraph G in which the vertices are the agents and the directed edges the possible calls. In the two-way version, G is a graph.

Example 1. Consider a network of five servers (which we call a, b, c, d and e) where each server can only communicate with a subset of the others. Note that all calls are assumed to be one-way in this example. As part of the maintenance program, a, b, c and d send a backup of their data to e every night and these backups can be sent to any server during the day (between 8:00 and 18:00). The others servers can communicate with each other at any moment if there is a communication link between them. The communication graph is depicted in Fig. 1.

There is some information on the server b that needs to be transferred to a, and c must know that the transfer is done. As the servers have different access rights, the information on server a should not be communicated to c. In the TEGP this can be represented by $Goal = K_c s_{ab} \wedge \neg s_{ca}$. There is a family of solution plans for this problem: $call(b, a)$ at instant t_1, $call(b, d)$ at instant t_2, $call(d, c)$ at instant t_3, where $t_1 < t_2 < t_3$ (together with the necessary calls $call(a, e)$ at instant 2, $call(b, e)$ at instant 4, $call(c, e)$ at instant 20 and $call(d, e)$ at instant 22).

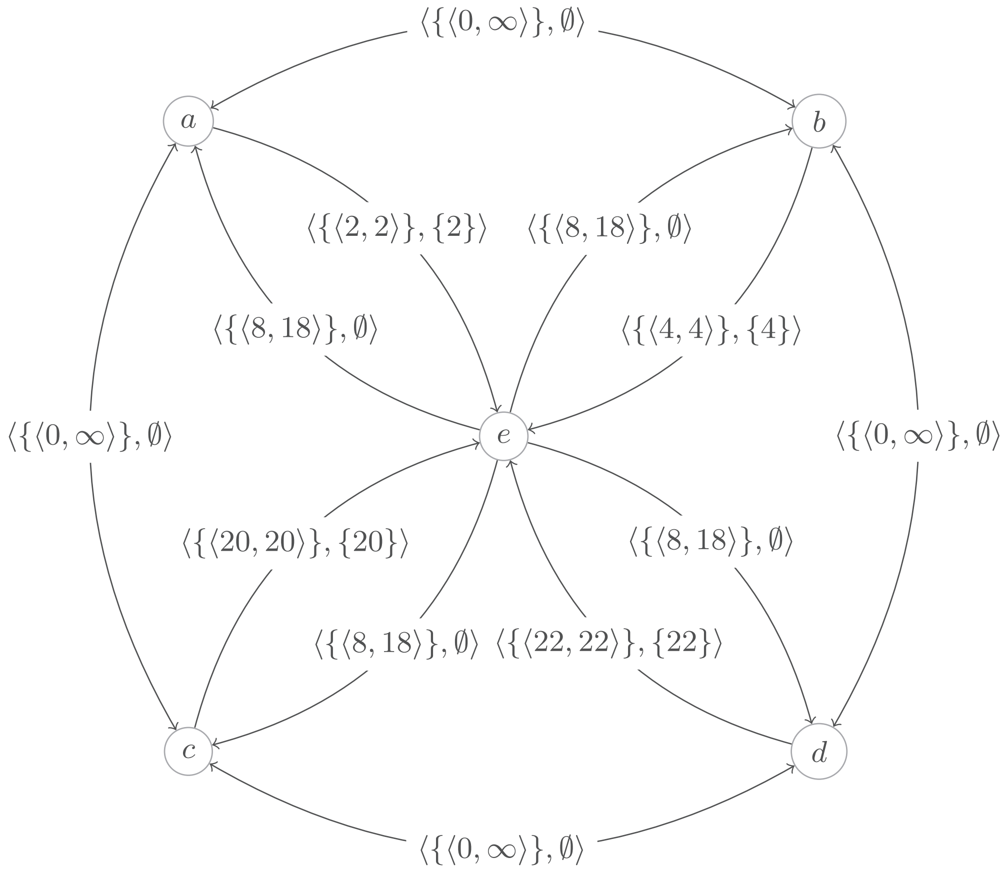

Fig. 1. Call graph for Example 1 involving necessary calls with e. A double-ended arrow represents two directed edges (i.e., the possibility of one-way calls in both directions).

On the same network, another question that we can ask is whether c can know a's data without a being aware of this. In this case, the goal is $s_{ca} \wedge \neg K_a s_{ca}$.

The answer is 'yes' since the following plan establishes the goal: the necessary $call(a, e)$ at instant 2, followed by $call(e, c)$ at an instant $t \in [8, 18]$ (together with the other necessary calls $call(b, e)$ at instant 4, $call(c, e)$ at instant 20 and $call(d, e)$ at instant 22).

3 Membership in NP

Proposition 1. *Let m be the number of clauses in the CNF of the goal. Let d be the depth of atoms for the problem. If a plan for an instance of a TEGP exists, then there is a plan with $md(n-1) + |I_n|$ calls or less*

Proof. Let $\alpha \in ATM^{\leq d}$ be an atom. $K_i\alpha$ can only be true if there is a path in the graph G between i and some agent who knows α. Without loss of generality, we can assume that this path is cycle-free. In the worst case the length of this path is $n-1$. Then, for any atom with an epistemic depth d, at most $d(n-1)$ calls are needed for this atom to be true, by the concatenation of d paths of length $n-1$.

The number of calls needed for a disjunction of formulas to be true is the maximum of the number of calls needed for each formula. In a CNF, there are only disjunctions over atoms, so the number of calls needed for a disjunction is $d(n-1)$.

The number of calls needed for a conjunction of formulas is the sum of the number of calls needed for every formula, which here is at most $d(n-1)$. So, with m being the number of conjunctions in the CNF of the goal, at most $md(n-1)$ calls are needed for a problem with only possible calls.

Thus, if a plan P exists, then P contains a subset Q of $md(n-1)$ calls which are sufficient to establish all positive atoms in the goal. For a problem with necessary calls, it can happen that the plan Q does not contain all the necessary calls I_n of P; but adding these necessary calls to Q to form a plan Q' cannot destroy positive goals. All negative atoms in the goal are also valid after execution of Q' since all calls, and in particular the ones in P but not in Q', only establish positive atoms. Thus Q' is a plan of at most $md(n-1) + |I_n|$ calls.

4 A Subproblem of the Temporal Gossip Problem in P

We say that a TEGP instance is *positive* if its goal is a CNF containing only positive atoms. A special case of TEGP is the class of positive non-temporally-constrained epistemic gossip problems $\Pi = \langle Init, Goal, Agt, I_p, I_n \rangle$ where $I_p = \{\langle 0, \infty, i, j \rangle : (i, j) \in E\}$, for some E, and $Goal$ is a positive CNF. In this case, E is the set of edges in the call digraph: if a call is possible (as specified by E), it is possible at any instant. On the other hand, there is no restriction on the set of necessary calls I_n.

Proposition 2. *The class of positive non-temporally-constrained epistemic gossip problems can be solved in polynomial time.*

Proof. There is a simple polynomial-time algorithm for positive non-temporally-constrained epistemic gossip problems: make all possible calls in some fixed order and repeat this operation $md(n-1)+|I_n|$ times. Call this sequential plan Q. By the proof of Proposition 1, if a solution plan exists, there is a sequential solution plan P of length at most $md(n-1)+|I_n|$. The actions of P necessarily appear as a subsequence of Q. Since the goal and preconditions of actions contain only positive atoms, the extra actions of Q cannot destroy any goals or preconditions. It follows that if a solution plan exists, then Q is also a solution plan. Thus this simple algorithm solves the class of positive non-temporally-constrained epistemic gossip problems in polynomial time.

Given an arbitrary instance of TEGP, we can construct a positive non-temporally-constrained instance by ignoring negative goals and temporal constraints (specified by I_p). This is a polynomial-time solvable relaxation of the original TEGP instance. This provides a relaxation which is inspired by the well-known delete-free relaxation of classical planning problems and is orthogonal to the relaxation of temporal planning problems based on establisher-uniqueness and monotonic fluents [13].

5 NP-completeness When Execution Time Is Bounded

The simplest temporal constraint is just a time limit on the execution of a plan. In the case of sequential plans this simply corresponds to placing a bound on plan length (which is equal to the number of calls) whereas in the parallel case execution time corresponds to the number of steps. We show in this section that this single constraint (a time limit on plan execution) is sufficient to render the epistemic gossip problem NP-complete. It is worth noting that the PSPACE complexity of classical planning is not affected by the possibility of placing an arbitrary limit on plan length, but the special case of delete-free planning passes from P to NP-hard when a bound is placed on plan length [9]. We show that this remains true for the specific case of gossiping problems.

We begin by studying the sequential case of TEGP.

Proposition 3. *The epistemic gossip problem with no temporal constraints but with a bound on the number of calls is NP-complete, even when the goal is a conjunction of positive atoms.*

Proof. We will exhibit a polynomial-time reduction from the well-known NP-complete problem SAT to the version of the epistemic gossip problem whose question is whether there is a sequential solution plan of length at most L. To do so, for a given set of clauses $\{C_1, \ldots, C_m\}$ we need the following agents:

- an agent S (the source),
- literal agents, i.e., agents for every variable and every negation of a variable (which we name, respectively, x^+ and x^-) for each variable x of the SAT instance,

– clause agents, i.e., agents for every clause (which we name C_i for the ith clause of the SAT instance).

Before performing this construction, we first add a dummy clause $(x \vee \neg x)$ for each SAT variable x. This clearly does not change the semantics of the instance but it does force us to specify the truth value of each variable in a solution of the SAT instance.

The source agent S and clause agents can only communicate with literal agents. The source agent S can communicate with every literal agent. A literal agent can only communicate with S and those clauses it is a member of. The graph G of communications is shown in Fig. 2 for a particular SAT instance. In this example, $C_1 = (\neg x \vee y \vee z)$, $C_2 = (\neg y \vee z)$ and the clauses C_3, C_4, C_5 are the dummy clauses $(x \vee \neg x)$, $(y \vee \neg y)$, $(z \vee \neg z)$.

A variable x is considered to be true (false) if S's secret passes through x^+ (respectively, x^-) in the solution plan on its way to the agent representing the dummy clause $(x \vee \neg x)$. The bound on the number of actions will prevent the possibility of S's secret passing through both x^+ and x^-. So the choice of whether S's secret passes through x^+ or x^- determines an assignment to the variable x in the SAT problem.

The goal of this instance of TEGP is that every clause agent knows the secret of S ($Goal = \bigwedge_{C_i} s_{C_i S}$). Now set the bound on plan length to be $L = 2n + m$, where n is the number of variables in the SAT instance and m the number of clauses in the original instance. With the new dummy clauses, the total number of clauses is $n + m$.

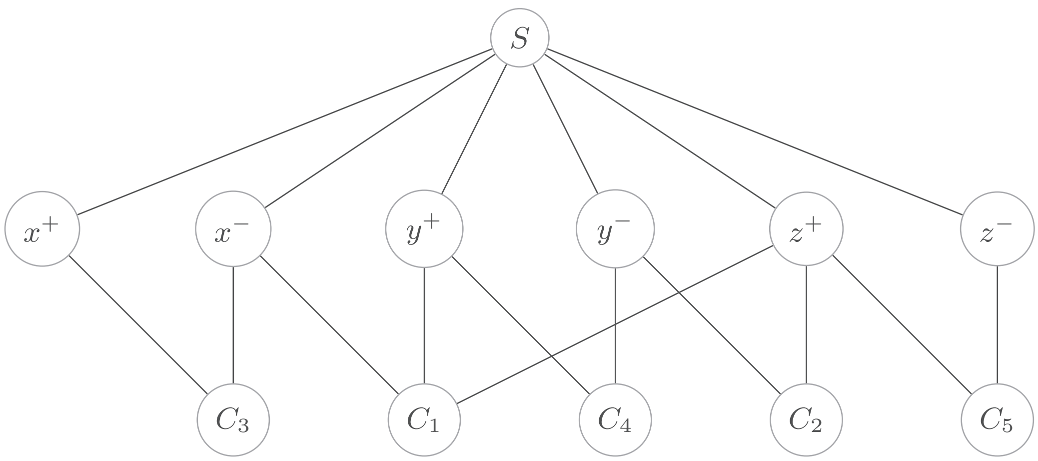

Fig. 2. Representation of the formula $(\neg x \vee y \vee z) \wedge (\neg y \vee z)$ as a temporal epistemic gossip problem in which the question is whether there is a plan using no more than 8 calls.

In a solution plan P we require at least n calls, one to either x^+ or x^-, for the n SAT variables x, in order for S's secret to be able to reach the agent corresponding to the dummy clause $x \vee \neg x$. P must also contain at least $n + m$ calls to the clause agents C_i (including the dummy clauses) to establish the goals $s_{C_i S}$. A solution plan of length precisely $2n+m$ corresponds to a solution of the corresponding SAT instance since such a plan defines a unique assignment to all variables that satisfies all clauses. For example, the solution $x = false$, $y = false$, $z = true$ to the SAT instance of Fig. 2 corresponds to the following solution plan of length 8: S calls x^-; S calls y^-; S calls z^+; x^- calls C_1; z^+ calls C_2; x^- calls C_3; y^- calls C_4; z^+ calls C_5. This reduction from SAT is clearly polynomial.

Proposition 1 proves the existence of a polynomial-length certificate for positive instances of the decision version of TEGP. Such certificates (solutions) can be verified in polynomial time. Thus TEGP $\in$ NP. Since the epistemic gossip problem with no temporal constraints but with a bound on the number of calls is clearly still in NP, this completes the proof of NP-completeness.

The proof of Proposition 3 was given for the case of two-way communications. It is trivial to adapt it to the case of one-way communications (for example, by only allowing calls from S to literal agents and from literal agents to clause agents).

We now consider the parallel version of the TEGP. Recall that in the parallel version of the TEGP, several calls may take place at each step, provided no agent is concerned by more than one call at each step.

Proposition 4. *The parallel version of the epistemic gossip problem with no temporal constraints except for a bound on the number of steps is NP-complete even when the goal is a conjunction of positive atoms.*

Proof. By the same argument as in the proof of Proposition 1, the problem is in NP. We complete the proof by exhibiting a polynomial reduction from 3SAT which is well known to be NP-complete. Given an instance I_{3SAT} of 3SAT, by introducing sufficiently many new variables x' which are copies of old variables x (together with the clauses $x \vee \neg x'$, $\neg x \vee x'$ to impose equality of x and x') we can transform I_{3SAT} into an equivalent instance in which each literal does not occur in more than three clauses. This is a polynomial reduction since we need to introduce at most one copy of each variable x per clause in which it occurs in I_{3SAT}. Therefore, from now on, we suppose that each literal occurs in at most two clauses in I_{3SAT}.

We construct an instance I of the epistemic gossip problem which has a parallel solution plan of length $2p$ if and only if I_{3SAT} is satisfiable. We choose the value of p to be strictly greater than $n+3$, where n is the number of variables in I_{3SAT}. To be concrete, we can choose $p = n+4$. We add to I_{3SAT} $p-n$ new dummy variables $x_{n+1}, \ldots, x_p$ none of which occur in the clauses of I_{3SAT}. In I there is an agent S (the source), literal agents x_i^+, x_i^- for each variable x_i $(i = 1, \ldots, p)$, and a clause agent C_j for each of the clauses C_j $(j = 1, \ldots, m)$ of I_{3SAT}. For each variable x_i $(i = 1, \ldots, p)$, we also add a dummy-clause agent D_i which we can consider as representing the dummy clause $x_i \vee \neg x_i$. Instead of linking these basic agents directly, we place paths of new agents between these basic agents. Between agent S and agent x_i^+ we add a path of length $p+1-i$. Similarly, we add a new path of the same length between S and agent x_i^-. For

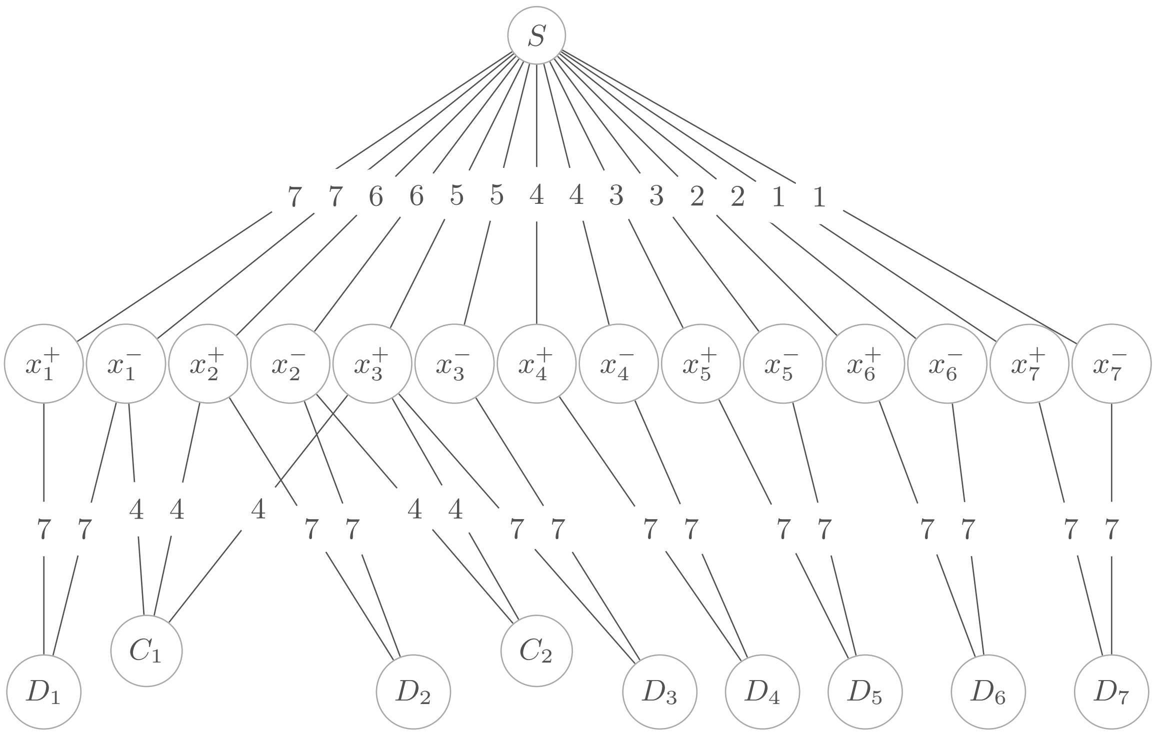

Fig. 3. Representation of the formula $(\neg x \vee y \vee z) \wedge (\neg y \vee z)$ as a temporal epistemic gossip problem in which the question is whether there is a parallel plan using no more than 14 steps.

$i = 1, \ldots, p$, we add two new paths both of length p between the literal agents x_i^+ and x_i^- and the dummy-clause agent D_i. For each clause C_j of I_{3SAT}, we also add three new paths of length q from the agents corresponding to the literals of C_j to the agent C_j, where $q = p - 2 = n + 1$. The resulting network is shown in Fig. 3 for an example instance. The numbers on edges in this figure represent the length of the corresponding path. For example, there are 6 intermediate agents (not shown so as not to clutter up the figure) between the agents S and x_1^+. The goal of I is

$$\left(\bigwedge_{i=1}^{p} s_{D_i S}\right) \wedge \left(\bigwedge_{j=1}^{m} s_{C_j S}\right)$$

In order to establish the goal $s_{D_i S}$, the secret of S has to follow a path from S to D_i. The shortest paths from S to D_1 are of length $2p$ and pass through either x_1^+ or x_1^-. Recall that our aim is to find a plan whose execution requires at most $2p$ steps. Thus, to establish $s_{D_1 S}$ in $2p$ steps, during the first step, S must call the first agent on the path to x_1^+ or the first agent on the path to x_1^-. The shortest path from S to D_2 is of length $2p - 1$, so during the second step, S must call the first agent either on the path to x_2^+ or the first agent on the path to x_2^-. By a simple inductive argument, we can see that during step i $(i = 1, \ldots, p)$, S must call the first agent on the path to x_i^+ or x_i^-. We can consider that the choice of whether S's secret passes through x_i^+ or x_i^- determines an assignment to the variables x_i. Due to the diminishing lengths of these paths as i increases, S's secret arrives simultaneously at the literal agents, either x_i^+ or x_i^-, for $i = 1, \ldots, p$. Another p steps are then required to send in parallel this secret to the dummy-clause agents D_j, for a total number of steps of $2p$. Almost simultaneously (within two time units), S' secret arrives at the clause agents C_j, provided it has passed through one of the agents corresponding to the literals of C_j. The length of paths from literal agents (x_i^+ or x_i^-) to clause agents C_j is $q = p - 3$ which is slightly less than p to allow for the fact that a literal agent, say x_i^+, may have to send S's secret along at most four paths: first towards D_i, then towards the (at most) three clauses in which x_i occurs.

It is important to note that S is necessarily occupied during the first p steps, as described above, so if S were to try to send its secret both to x_i^+ and x_i^- the secret could not arrive via the second of these paths at a clause agent C_j in less than $2p - n + q = 3p - n - 3$ steps which is greater than the upper bound of $2p$ steps (since $p = n + 4$). By our construction, the goal $s_{C_j S}$ is established only if the assignment to the variables x_i determined by the solution plan satisfies the clause C_j. Hence, parallel solution plans of length $2p$ steps correspond precisely to solutions of I_{3SAT}. We have therefore demonstrated a polynomial reduction from 3SAT to the parallel version of the epistemic gossip problem with a bound on the number of steps.

The following corollary follows from the fact that we can place an upper bound L on the number of steps in a plan by simply imposing via I_p an interval of possible instants $[1, L]$ for all calls.

Corollary 1. *TEGP is NP-complete.*

6 NP-completeness of Gossiping with Negative Goals

We show in this section that even without temporal constraints or a bound on plan length, when we allow negative goals the problem of deciding the existence of a solution plan is NP-complete.

Proposition 5. *The epistemic gossip problem with possibly negative goals is NP-complete even in the absence of any temporal constraints or bound on plan length.*

Proof. The same argument as in the proof of Proposition 1 shows that the problem belongs to NP since it is a subproblem of TEGP.

To complete the proof, it suffices to give a polynomial reduction from SAT. Let I_{SAT} be an instance of SAT. We will construct a call graph G and a set of goals such that the corresponding instance I_{Gossip} of the epistemic gossip problem is equivalent to I_{SAT}. Recall that the nodes of the call graph G are the agents and the edges of G the communication links between agents.

For each propositional variable x in I_{SAT}, we add four nodes x^+, x^-, b_x, d_x to G joined by the edges shown in Fig. 4(b). There is a source node S in G and edges (S, x^+), (S, x^-) for each variable x in I_{SAT}. For each clause C_j in I_{SAT}, we add a node C_j joined to the nodes corresponding to the literals of C_j. This is illustrated in Fig. 4(a) for the clause $C_j = \neg x \vee y \vee z$. The solution plan to I_{Gossip} will make S's secret transit through x^+ (on its way from S to some clause node C_j) if and only if $x = true$ in the corresponding solution to I_{SAT}.

For each clause C_j in I_{SAT}, G contains a clause gadget as illustrated in Fig. 4(a) for the clause $\neg x \vee y \vee z$. We also add $s_{C_j S}$ to the set of goals. Clearly, S's secret must transit through one of the nodes corresponding to the literals of C_j (x^-, y^+ or z^+ in the example of Fig. 4) to achieve the goal $s_{C_j S}$.

To complete the reduction, it only remains to impose the constraint that a's secret transits through at most one of the nodes x^+, x^-, for each variable x of I_{SAT}. This is achieved by the negation gadget shown in Fig. 4(b) for each variable x. We add the goals $s_{d_x b_x}$ and $\neg s_{d_x S}$ for each variable x, and the goal $\neg s_{C_j b_x}$ for each variable x and each clause C_j (containing the literal x or $\neg x$). The goal $s_{d_x b_x}$ ensures that b_x's secret transits through x or $\neg x$. Now, recall that we assume that during a call, agents communicate all their knowledge. Suppose that b_x's secret transits through x^+: then S's secret cannot transit through x^+ before b_x's secret (because of the negative goal $\neg s_{d_x S}$) and cannot transit through x^+ after b_x's secret (because of the negative goal $\neg s_{C_j b_x}$). By a similar argument, if b_x's secret transits through x^-, then S's secret cannot transit through x^+. Thus, this gadget imposes that S's secret transits through exactly one of the nodes x^+, x^-.

We have shown that I_{SAT} has a solution if and only if I_{Gossip} has a solution. Since the reduction is clearly polynomial, this completes the proof.

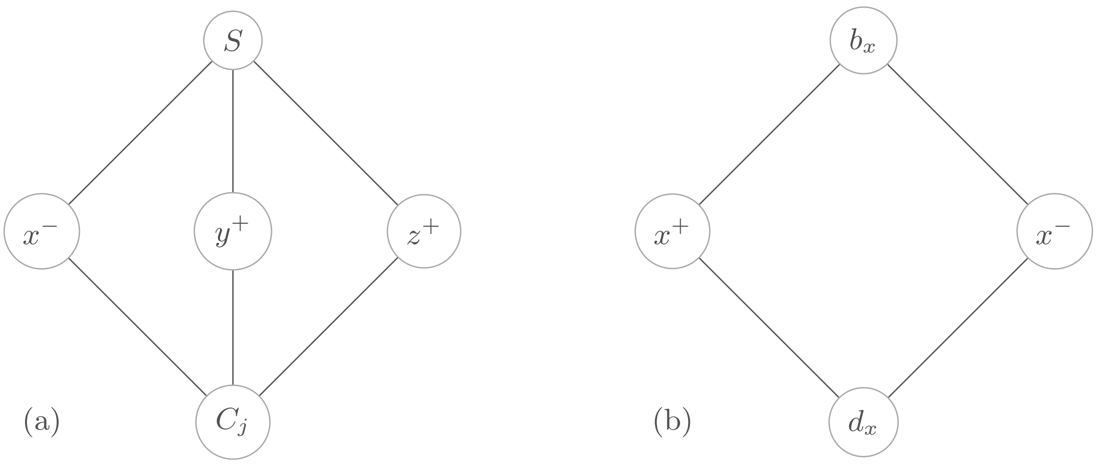

Fig. 4. (a) Gadget imposing the clause $C_j = \neg x \vee y \vee z$; (b) Gadget imposing the choice between x and $\neg x$.

The NP-completeness shown in the proof of Proposition 5 for two-way communication, would not be affected by a restriction to one-way communication. Similarly NP-completeness holds for both the sequential and parallel versions of the gossip problem.

7 Discussion and Conclusion

We have defined temporal epistemic gossip problems and have investigated their complexity. Our results are in line with previous results concerning epistemic planning: it is possible to add an epistemic dimension to planning, thus increasing expressibility, without increasing complexity [10].

In our approach agents are not introspective: s_{ij} does not imply $K_i s_{ij}$, and $K_k s_{ij}$ does not imply $K_k K_k s_{ij}$. This only concerns positive introspection: negative introspection cannot be expressed. Positive introspection can however be enforced by adding axioms $s_{ij} \rightarrow K_i s_{ij}$, and $K_k s_{ij} \rightarrow K_k K_k s_{ij}$. We however did not do so in order to simplify presentation.

We have assumed a centralized approach in which a centralized planner decides the actions of all agents. Several other researchers have recently studied distributed versions of the classical gossip problem where the agents have to decide themselves whom to call, based on the knowledge (and ignorance) they have [1–5,14,15]. An interesting avenue of future research would be to consider the epistemic gossip problem in this framework.

Several other variants of our centralized model could also be investigated, including the precondition that i has to know the telephone number of j in order to call j and telephone numbers are communicated in the same way as secrets. In another variant, the secrets can be passwords which are no longer constants since each agent i can change their own password [12].

References

1. Apt, K.R., Grossi, D., van der Hoek, W.: Epistemic protocols for distributed gossiping. In: Ramanujam, R. (ed.) Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015. EPTCS, vol. 215, pp. 51–66 (2015). https://doi.org/10.4204/EPTCS.215.5
2. Apt, K.R., Grossi, D., van der Hoek, W.: When are two gossips the same? Types of communication in epistemic gossip protocols. CoRR abs/1807.05283 (2018). http://arxiv.org/abs/1807.05283
3. Apt, K.R., Kopczynski, E., Wojtczak, D.: On the computational complexity of gossip protocols. In: Sierra, C. (ed.) Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, pp. 765–771. ijcai.org (2017). https://doi.org/10.24963/ijcai.2017/106
4. Apt, K.R., Wojtczak, D.: Common knowledge in a logic of gossips. In: Lang, J. (ed.) Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017. EPTCS, vol. 251, pp. 10–27 (2017). https://doi.org/10.4204/EPTCS.251.2
5. Apt, K.R., Wojtczak, D.: Decidability of fair termination of gossip protocols. In: Eiter, T., Sands, D., Sutcliffe, G., Voronkov, A. (eds.) IWIL@LPAR 2017 Workshop and LPAR-21 Short Presentations, vol. 1. Kalpa Publications in Computing, EasyChair (2017). http://www.easychair.org/publications/paper/342983
6. Attamah, M., van Ditmarsch, H., Grossi, D., van der Hoek, W.: A framework for epistemic gossip protocols. In: Bulling, N. (ed.) EUMAS 2014. LNCS (LNAI), vol. 8953, pp. 193–209. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-17130-2_13

7. Aucher, G., Bolander, T.: Undecidability in epistemic planning. In: Rossi, F. (ed.) Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI 2013, 3–9 August 2013, Beijing, China, pp. 27–33. IJCAI/AAAI (2013). http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6903
8. Bolander, T., Andersen, M.B.: Epistemic planning for single and multi-agent systems. J. Appl. Non-Classical Logics **21**(1), 9–34 (2011). https://doi.org/10.3166/jancl.21.9-34
9. Bylander, T.: The computational complexity of propositional STRIPS planning. Artif. Intell. **69**(1–2), 165–204 (1994). https://doi.org/10.1016/0004-3702(94)90081-7
10. Cooper, M.C., Herzig, A., Maffre, F., Maris, F., Régnier, P.: A simple account of multi-agent epistemic planning. In: Kaminka et al. [19], pp. 193–201. https://doi.org/10.3233/978-1-61499-672-9-193
11. Cooper, M.C., Herzig, A., Maffre, F., Maris, F., Régnier, P.: Simple epistemic planning: generalised gossiping. In: Kaminka et al. [19], pp. 1563–1564. https://doi.org/10.3233/978-1-61499-672-9-1563
12. Cooper, M.C., Herzig, A., Maffre, F., Maris, F., Régnier, P.: Simple epistemic planning: generalised gossiping. CoRR abs/1606.03244 (2016). http://arxiv.org/abs/1606.03244
13. Cooper, M.C., Maris, F., Régnier, P.: Monotone temporal planning: tractability, extensions and applications. J. Artif. Intell. Res. **50**, 447–485 (2014). https://doi.org/10.1613/jair.4358
14. van Ditmarsch, H., van Eijck, J., Pardo, P., Ramezanian, R., Schwarzentruber, F.: Epistemic protocols for dynamic gossip. J. Appl. Logic **20**, 1–31 (2017). https://doi.org/10.1016/j.jal.2016.12.001
15. van Ditmarsch, H., Grossi, D., Herzig, A., van der Hoek, W., Kuijer, L.B.: Parameters for epistemic gossip problems. In: Proceedings of LOFT 2016 (2016)
16. Hedetniemi, S.M., Hedetniemi, S.T., Liestman, A.L.: A survey of gossiping and broadcasting in communication networks. Networks **18**(4), 319–349 (1988). https://doi.org/10.1002/net.3230180406
17. Herzig, A., Lorini, E., Maffre, F.: A poor man's epistemic logic based on propositional assignment and higher-order observation. In: van der Hoek, W., Holliday, W.H., Wang, W. (eds.) LORI 2015. LNCS, vol. 9394, pp. 156–168. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48561-3_13
18. Herzig, A., Maffre, F.: How to share knowledge by gossiping. AI Commun. **30**(1), 1–17 (2017). https://doi.org/10.3233/AIC-170723
19. Kaminka, G.A., et al. (eds.): 22nd European Conference on Artificial Intelligence, Frontiers in Artificial Intelligence and Applications, ECAI 2016, vol. 285. IOS Press (2016)
20. Löwe, B., Pacuit, E., Witzel, A.: DEL planning and some tractable cases. In: van Ditmarsch, H., Lang, J., Ju, S. (eds.) LORI 2011. LNCS (LNAI), vol. 6953, pp. 179–192. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24130-7_13
21. Maffre, F.: Ignorance is bliss: observability-based dynamic epistemic logics and their applications. Ph.D. thesis, Paul Sabatier University, Toulouse, France (2016). https://tel.archives-ouvertes.fr/tel-01488408
22. Muise, C.J., et al.: Planning over multi-agent epistemic states: a classical planning approach. In: Bonet, B., Koenig, S. (eds.) Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, 25–30 January 2015, Austin, Texas, USA, pp. 3327–3334. AAAI Press (2015). http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9974