Epistemic protocols for dynamic gossip

Hans van Ditmarsch [a], Jan van Eijck [b], Pere Pardo [c,*], Rahim Ramezanian [d],
François Schwarzentruber [e]

[a] *LORIA, CNRS, University of Lorraine, France*
[b] *CWI and ILLC, University of Amsterdam, Netherlands*
[c] *Institut für Philosophie II, Ruhr-Universität Bochum, Germany*
[d] *Sharif University of Technology, Iran*
[e] *ENS Rennes, IRISA, France*

A R T I C L E   I N F O

A B S T R A C T

A gossip protocol is a procedure for spreading secrets among a group of agents, using a connection graph. In each call between a pair of connected agents, the two agents share all the secrets they have learnt. In dynamic gossip problems, dynamic connection graphs are enabled by permitting agents to spread as well the telephone numbers of other agents they know. This paper characterizes different distributed epistemic protocols in terms of the (largest) class of graphs where each protocol is successful, i.e. where the protocol necessarily ends up with all agents knowing all secrets.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Gossip protocols are procedures for spreading secrets among a group of agents, using a connection graph.[1] This connection graph denotes which agents know the telephone numbers of who, and so which telephone calls can be made. During each call, say ab from a to b, the two agents a and b share all the secrets they know at the time of the call (initially, each agent only knows its own secret). In the original set-up, where a totally connected graph was assumed, the main question was to find a sequence of calls making all agents knew all secrets (henceforth, successful) with a minimal number of calls; this minimum length turns out to be $2n - 4$ in a totally connected graph with $n > 3$ agents, see Tijdeman [22] or Hurkens [17]. Consider the totally connected graph with four agents, and the successful call sequence ab; cd; ac; bd in the following

---

* Corresponding author.
E-mail addresses: hans.van-ditmarsch@loria.fr (H. van Ditmarsch), jan.van.eijck@cwi.nl (J. van Eijck),
pere.pardoventura@mail.ruhr-uni-bochum.de (P. Pardo), rahim_ramezanian@mehr.sharif.ir (R. Ramezanian),
francois.schwarzentruber@ens-rennes.fr (F. Schwarzentruber).
[1] For an example of application, consider the problem of *scheduling a meeting* for a group of agents. The *secret* of each agent is its time availability for possible meetings.

picture (where dashed arrows denote knowledge of telephone number, and solid arrows denote who knows whose secrets; note we do not draw the reflexive edges corresponding to the fact that agents always know their own secret and telephone number).

Graphs that are not totally connected were also studied by different authors; [14] is a survey of results on the gossip problem, including the following minimum bounds on the length of successful sequences: for any (undirected) connected graph, there is a successful call sequence of length $2n - 3$ for $n > 3$ agents; moreover, if the graph contains a 4-cycle (as in totally connected graphs), then the minimum is again $2n - 4$ calls. Consider for example the following connection graph:

In the connection graph above, the call sequence $ac; bd; ab; bd; ac$ can be improved if we permit the agents to exchange, during each call, both the secrets and the telephone numbers. In this setting, called *dynamic gossip* in some previous work [24], we can recover the $2n - 4$ minimum length for the above gossip graph.


In all of the examples above, a way to ensure that minimum length solutions will be found is to assume a central authority who, knowing the initial network topology, designs the desired call sequence and communicates it to the agents. In distributed computing, in contrast, one is interested in procedures that do not need outside regulation, even if this comes at the price of non-optimal call sequences. Among distributed protocols that might find the above solutions one can mention

Any Call While not every agent knows all secrets, randomly select a pair xy such that x knows y's number and let x call y.

or, with less redundancy (since the sequence $ac; ac; bd; ab; cd$ is not permitted), the protocol

Learn New Secrets While not every agent knows all secrets, randomly select a pair xy such that x knows y's number but not her secret and let x call y.

The following execution $bc; ab$ of the Learn New Secrets protocol is unsuccessful in the next graph:

One can observe that after the second call ab, in the above sequence $bc; ab$, agents a and b know that it surely makes sense to make a new call ac or bc, since agent c would then learn something, namely, the initial secret from a. One protocol to address this is

Known Information Growth While not every agent knows all secrets, randomly select a pair xy such that x knows y's number and there is a z such that x knows that either x or y doesn't know z's secret; let x call y.

Assume (for now) that the calls are synchronous (i.e. it is common knowledge that one call is made after each clock-tick and the number of calls made is common knowledge[2]). The above situation can be solved by the Known Information Growth protocol with an extra call bc:

This protocol is based on high-order knowledge, namely on what b knows that c knows (beyond mere possession of her own secrets). Knowledge of other agent's knowledge cannot be as straightforwardly related to a graph property. In the same example, initially agent a does not know whether someone has her telephone number, so from a's perspective the graph could be either of the following, after the initial bc call (the call history is written on top of each graph):


We say in this case that a *cannot distinguish* the two graphs, and for such an indistinguishability relation we write $\approx_a$ (or $\sim_a$ in the case of asynchronous calls, see below).

After the first two calls $bc; ab$, agents a and b do not consider any graph possible other than the actual one. So the $\approx_a$ and $\approx_b$ possible worlds simply consist of:

An agent knows something about a graph if it is true in all indistinguishable graphs. Agents a and b therefore now know that c does not know a's secret; therefore either of them can call c according to the Known Information Growth protocol above, and thus successfully terminate the protocol.

The above assumption on the synchronicity of calls is crucial, since without it (denoted $\sim$), the protocol does not successfully terminate. Informally, we say that calls are asynchronous if at each time there is no upper bound on the number of calls that could already have taken place according to any agent. Under this assumption, we have that after $bc; ab$, the agents a and b consider other possibilities besides the actual world. The following picture shows some of those possibilities from the point of view of the actual world (depicted in the center):

This paper is a follow-up study of [24]. The main difference lies in the gossip protocols, which are knowledge-based, and the new semantic models which endow gossip scenarios with notions of knowledge for synchronous and asynchronous calls. Our choice of protocols for making calls is mainly motivated by simplicity and the immediate goal of learning secrets: typically, a protocol will require that the caller knows (or at least considers it possible) that the next call will make herself or the callee to learn some new secret; these expectations after the call can be expressed as distributions of secrets before the call. Further distinctions among our protocols are: the direction of the learning of secrets (the caller intends to learn new secrets from the callee, or to transmit her new secrets, or both), the origin of the secrets aimed to be

---

[2] Therefore, if there are only 3 agents, the two agents in a call are correctly identified by the remaining agent.

learnt (the secret of caller or callee vs. of any agent), and finally, the *de dicto/de re* distinction w.r.t. the identification of the secrets that will be learnt in the next call. In all these protocols, the preconditions for calls can be expressed by simple epistemic formulas of modal depth one, namely propositional formulas in the scope of a modality for the caller, effectively making these protocols distributed.

All in all, we aimed to study protocols that are simple and intuitive to grasp for human gossip agents. Our main goal is the characterization of the protocols under study, that is, to identify the class of graphs where all (or some) executions of a given protocol successfully terminate. Of course, our framework is expressive enough to capture many other protocols with higher-order epistemic conditions.

Outline of the paper　　The paper is divided as follows: in Section 2, we introduce the epistemic framework – gossip graphs, knowledge in (a)synchronous calls, protocols – and recall some general results used throughout the paper. After this, in Section 3 we review results on some basic protocols from [24] to be used later for comparison. The epistemic protocols studied in this paper start in Section 4, with some distributed protocols where learning a secret is possible, according to the caller of the next call. Then in Section 5 we study non-redundant distributed protocols: learning a secret is known necessary, before the call, according to the caller. Finally, in Section 6 we compare these protocols in terms of their sets of executions in given gossip scenarios. The paper concludes with a review of the related work, in Section 7, a summary of contributions and a list of directions for further research, in Section 8.

2. Basic concepts

Consider a finite set of agents $A = \{a, b, \ldots\}$, each with access to a set of other agents, and each carrying a unique secret, called the agent's secret. For these agents to reach the goal of knowing (all of them) all the secrets, they will call each other and exchange all the telephone numbers and secrets they know. The two call modes considered in this paper are: synchronous, where agents not involved in the current call at least know that a call is being made, and asynchronous, where those agents are not even aware of the current existence of this call. Despite this ignorance on the present, future calls might help an agent to identify the current call: the secrets that one receives at each call provide information on possible call histories. All this will be made precise in the epistemic accessibility relation corresponding to each call mode.

Rather than letting the agents randomly call each other, one can define distributed protocols executing fewer redundant calls. Sometimes this comes at a price, namely that the protocol is not successful in all possible initial configurations. The main results in this paper consist in characterizing the class of gossip graphs where each protocol is necessarily (or possibly) successful. The protocols studied here are definable in terms of the knowledge of agents at the time they are initiating a call.

2.1. Gossip graphs

Given a finite set of agents (or nodes) $A = \{a, b, \ldots\}$, we represent a gossip graph G with telephone numbers and secrets as a triple (A, N, S) with $N, S \subseteq A \times A$. That is, the agents A are the vertices and N, S are binary relations on A, with Nxy (for $(x, y) \in N$) expressing that x knows the (telephone) number of y, and Sxy expressing that x knows the secret of y.

Let us first introduce standard graph terminology, given a carrier set A and binary relations like N and S. We let $I_A = \{(a, a)\}_{a \in A}$ be the *identity* relation on A, and *converse* relation $N^{-1} = \{(x, y) \mid Nyx\}$. We write N_x for $\{y \in A \mid Nxy\}$. (We may further write $\neg Nxy$ for $(x, y) \notin N$ and anyway write xy for a pair (x, y), such as for the calls defined below.) Relation $N \circ S = \{(x, y) \mid$ there is a z such that Nxz and $Szy\}$ is the *composition* of N and S, and using that we define $N^1 = N$, $N^{i+1} = N^i \circ N$, and $N^* = \bigcup_{i \geq 1} N^i$. Relation N is *complete* iff $N = A^2$, it is *weakly connected* if for all $x, y \in A$ there is an $(N \cup N^{-1})$-path from x to y, and it is *strongly connected* if for all $x, y \in A$ there is an N-path from x to y.

Definition 1 *(Gossip graph).* A *gossip graph* is a triple $G = (A, N, S)$ with $N \subseteq A^2$ and $S \subseteq A^2$. An *initial gossip graph* is a gossip graph with $S = I_A \subseteq N$. Agent x is an *expert* if $S_x = A$. An agent or node is *terminal* iff $N_x = \{x\}$. Gossip graph G is *complete, weakly (strongly) connected* if the graph (A, N) is, respectively, complete, weakly (strongly) connected. The complete initial gossip graph with n agents is denoted K_n.

In an initial gossip graph each agent only knows its own secret and at least knows its own number. When we employ common graph terminology when referring to a gossip graph, this applies to the number relation, not to the secret relation, as above.

A *call* from x to y is a pair (x, y) for which we write xy. The call xy in G merges the secrets and the numbers of x and y. The expression $\overline{xy}$ stands for 'xy or yx'.

Definition 2 *(Call; call-induced gossip graph).* Let $G = (A, N, S)$ and $x, y \in A$. We say that the *call* xy is valid in G if Nxy, and in such case the call xy maps G to the *call-induced gossip graph* $G^{xy} = (A, N^{xy}, S^{xy})$, defined by

$$N_z^{xy} = \begin{cases} N_x \cup N_y & \text{if } z \in \{x, y\} \\ N_z & \text{otherwise} \end{cases} \quad \text{and} \quad S_z^{xy} = \begin{cases} S_x \cup S_y & \text{if } z \in \{x, y\} \\ S_z & \text{otherwise} \end{cases}$$

Definition 3 *(Call sequence).* A *call sequence* σ is a finite or infinite sequence of calls. The empty sequence is ϵ. We write $\sigma; \tau$ for the concatenation of a finite call sequence σ and a call sequence τ.

Given a finite call sequence σ, we define G^σ by induction on σ: $G^\epsilon = G$ and $G^{xy;\sigma} = (G^{xy})^\sigma$.

Given a sequence σ with at least n calls, we denote by σ_n the n-th call in σ (starting from 1), and $\sigma|n$ is the prefix of σ consisting of the first n calls, so that we have that $\sigma|0 = \epsilon$ and $\sigma|n = \sigma_1; \ldots; \sigma_n$. We will denote that σ is a (proper) prefix of τ by $\sigma \sqsubseteq \tau$ (resp. $\sigma \sqsubset \tau$). A call sequence σ is *valid* in G iff for each n, σ_{n+1} is valid in $G^{\sigma|n}$.

For a valid infinite call sequence σ, we define $G^\sigma = (A, \bigcup_n N^{\sigma|n}, \bigcup_n S^{\sigma|n})$.[3]

Finally, for $x \in A$, σ_x is the subsequence of calls containing x, i.e., $\epsilon_x = \epsilon$, $(\sigma; zw)_x = \sigma_x$ if $x \notin \{z, w\}$, and $(\sigma; zw)_x = \sigma_x; zw$ otherwise. If σ is finite, we denote by σ^n the sequence $\sigma; \ldots; \sigma$ (n times). If σ is finite, we denote by σ^ω the sequence $\sigma; \sigma; \ldots$ (σ is repeated infinitely many times).

Fact 4. *[24] For any initial gossip graph G and any finite valid call sequence σ, the resulting gossip graph $G^\sigma = (A, N^\sigma, S^\sigma)$ satisfies $S^\sigma \subseteq N^\sigma$.*

Therefore, provided that $S \subseteq N$ (as in our case: $S = I_A \subseteq N$), it does not matter whether the goal is to learn all the secrets, or both all secrets and all numbers. By Fact 4, the first is not possible without the second.

Lemma 5. *[24] If $G = (A, N, S)$ is an initial gossip graph and σ is a finite valid call sequence for G, then $S^\sigma \circ N \subseteq N^\sigma$ where $G^\sigma = (A, N^\sigma, S^\sigma)$.*

It is shown in [24] that for any valid call sequence σ in a gossip graph G,

$$G \text{ is weakly connected} \quad \text{iff} \quad G^\sigma \text{ is weakly connected}.$$

---

[3] Note that a protocol execution (see Definition 13) defines how to generate the next call, which gives all of its finite call sequences. For the characterization results, though, the protocol extensions need to contain the infinite sequences as well, i.e. the result of executing a protocol indefinitely. The restriction to finite sequences is only applied to the following concepts: gossip states (Definition 7), accessibility relations (Definitions 9–10), gossip models (Definition 12), the semantics (Definition 15) and model-checking.

Fig. 1. Update of the indistinguishable gossip graphs for agent c after a call $\overline{ab}$.

As a consequence, the goal of all agents being expert can never be reached in graphs that are not weakly connected.

Corollary 6. *[24] If G is not weakly connected, then for any finite valid call sequence σ, no one is expert in G^σ.*

2.2. Knowledge

In both the asynchronous and synchronous call modes, the agents may not know what secrets other agents know at a given moment. Formally, this is encoded with the help of possible call histories in initial gossip graphs.

Definition 7 *(Gossip state).* A *gossip state* is a pair (G, σ) where G is an initial gossip graph and σ is a finite valid call sequence in G.

A gossip state (G, σ) contains more information than the corresponding induced gossip graph G^σ — see Example 8 below. With the help of gossip states, the knowledge of an agent will be represented as a set of possible call histories, i.e. a set of initial gossip states and call sequences (upon them) that are compatible with both this agent's initial knowledge (contact list) and her learning through the calls made or received so far.

Example 8. We consider a synchronous scenario with three agents a, b and c, from the viewpoint of agent c. The number relation for c is given by: $\neg Nca$, Ncb, Ncc. Fig. 1 (left) depicts the possible initial gossip graphs for agent c (the knowledge of agents a, b is not represented). (Right) depicts the possible gossip graphs for agent c after a call $\overline{ab}$ occurs. With 3 agents only, as c is not involved in the first call, c will know that Nab or Nba and that the call is in fact an $\overline{ab}$ call.

Note that in Fig. 1 some initial possibilities, namely G_2, G_3, G_8, G_9, are discarded after the transition $\epsilon \mapsto \overline{ab}$. After a call $\overline{ab}$ only two gossip states are possible, namely G_0', G_1':

$$G_0' = G_0^{ab} = G_0^{ba} = G_1^{ab} = G_1^{ba} = G_4^{ba} = G_5^{ba} = G_6^{ba} = G_7^{ab} = G_{10}^{ab} = G_{10}^{ba} = G_{12}^{ba}$$
$$G_1' = G_{11}^{ab} = G_{11}^{ba} = G_{13}^{ba}$$

even if as gossip states these are pairwise different (but indistinguishable by c in the synchronous call mode):

$$(G_0, ab) \neq (G_0, ba) \neq (G_1, ab) \neq \cdots \neq (G_{12}, ba) \quad \text{and} \quad (G_{11}, ab) \neq (G_{11}, ba) \neq (G_{13}, ba)$$

In comparison, under the asynchronous call mode, the update with a call $\overline{ab}$ of Fig. 1 would result in the union of all the gossip graphs in this figure, as agent c would consider it also possible that no call was made.

The sequence σ is not explicitly represented in call-induced gossip graph G^σ, the latter is merely the gossip graph G' resulting from executing σ in G.

The first call mode considered in this paper is asynchronous: if an agent is not involved in a call (by making or receiving it), then she is not aware that a call has been made. Since the agents do not have internal clocks, any finite number of calls can be imagined to take place between any pair of actual calls.

Definition 9 *(Asynchronous accessibility relation).* Let $G = (A, N, S)$ and $H = (A, O, T)$ be initial gossip graphs, $x \in A$. We define $(G, \sigma) \sim_x (H, \tau)$ on the set of gossip states for A by induction on the sum of the lengths of finite call sequences σ and τ, valid in respectively G and H,

1. $(G, \epsilon) \sim_x (H, \epsilon)$ if $N_x = O_x$;
2. For all $(\sigma, \tau) \neq (\epsilon, \epsilon)$, $(G, \sigma) \sim_x (H, \tau)$ if
   either (a) $\sigma = \sigma'; yz$, $x \notin \{y, z\}$ and $(G, \sigma') \sim_x (H, \tau)$;
   or (b) $\tau = \tau'; yz$, $x \notin \{y, z\}$ and $(G, \sigma) \sim_x (H, \tau')$;
   or (c) $x \in \{y, z\}$, $\sigma = \sigma'; yz$, $\tau = \tau'; yz$, $S_y^{\sigma'} = T_y^{\tau'}$, $S_z^{\sigma'} = T_z^{\tau'}$, $N_y^{\sigma'} = O_y^{\tau'}$, $N_z^{\sigma'} = O_z^{\tau'}$ and $(G, \sigma') \sim_x (H, \tau')$.

(Note that in Definition 9(c) the conditions for agent $x \in \{y, z\}$, i.e. $S_x^{\sigma'} = T_x^{\tau'}$ and $N_x^{\sigma'} = O_x^{\tau'}$, already follow from $(G, \sigma') \sim_x (H, \tau')$. The same comment applies to Definition 10.)

In the synchronous mode of calls, every time a call is made the agents know that (but only that, not the identity of the two agents on the telephone). This call mode can be described by telephone networks that only allow for a single call at any given time, so the availability of the telephone system is common knowledge. Although less realistic than the asynchronous mode, it will be very instructive to compare each gossip protocol under the two call modes.

Definition 10 *(Synchronous accessibility relation).* Let $G = (A, N, S)$ and $H = (A, O, T)$ be initial gossip graphs, and $x \in A$. We define $\approx_x$ by induction on finite call sequences σ and τ valid, resp., in G and H:

1. $(G, \epsilon) \approx_x (H, \epsilon)$ if $N_x = O_x$;
2. For all $(\sigma, \tau) \neq (\epsilon, \epsilon)$, $(G, \sigma) \approx_x (H, \tau)$ if $\sigma = \sigma'; yz$, $\tau = \tau'; uv$ and
   (a) either $x \notin \{y, z, u, v\}$, $N^{\sigma'} yz$, $O^{\tau'} uv$ and $(G, \sigma') \approx_x (H, \tau')$;
   (b) or $x \in \{y, z\}$, $y = u$, $z = v$, $S_y^{\sigma'} = T_y^{\tau'}$, $S_z^{\sigma'} = T_z^{\tau'}$, $N_y^{\sigma'} = O_y^{\tau'}$ and $N_z^{\sigma'} = O_z^{\tau'}$ and $(G, \sigma) \approx_x (H, \tau)$.

The accessibility relations in Definitions 9 and 10 are equivalent to those found in Attamah et al. [2] for the original (non-dynamic) gossip problem. All these definitions are based on the *learn-then-merge* model of calls, where the two agents in a call first *learn* what the other knows, and then *merge* this with their own knowledge. An alternative model of calls is that of *merge-then-learn*, in Apt et al. [1], where the information from the two agents is first sent to a server and *merged* into a set, and then the agents *learn* this new set. This model of calls induces less fine-grained accessibility relations, see the Related Work for more details.

The following is easily proved by induction on the length of σ and τ.

Proposition 11. *For each pair of gossip graphs $G = (A, N, S)$ and $H = (A, O, T)$, and each pair of finite call sequences σ and τ, the following implications hold for any $x \in A$:*

(1) $(G, \sigma) \approx_x (H, \tau) \Rightarrow (G, \sigma) \sim_x (H, \tau) \Rightarrow \sigma_x = \tau_x$,
(2) *the accessibility relations $\sim_x$ and $\approx_x$ are equivalence relations.*

Definition 12 *(Gossip model).* Given a set of agents A, the *asynchronous gossip model* and the *synchronous gossip model* are respectively the triples

$$\mathcal{G}^\sim = (\mathcal{G}, \langle \sim_a \rangle_{a \in A}, \langle \xrightarrow{ab} \rangle_{a,b \in A}) \quad \text{and} \quad \mathcal{G}^\approx = (\mathcal{G}, \langle \approx_a \rangle_{a \in A}, \langle \xrightarrow{ab} \rangle_{a,b \in A})$$

where:

- $\mathcal{G}$ is the set of all gossip states (G, σ) for the set of agents A.
- $\sim_a$ and $\approx_a$ are relations defined in Definitions 9 and 10;
- each $\xrightarrow{ab}$: $\mathcal{G} \to \mathcal{G}$ is a partial function representing the dynamic transitions between $\mathcal{G}$-states $(G, \sigma) \xrightarrow{ab} (G, \sigma; ab)$ defined on the set of gossip states in $\mathcal{G}$ where the call ab is valid.

Under this definition, each gossip model $\mathcal{G}^\sim$ or $\mathcal{G}^\approx$ (with at least two agents) has an infinite number of states: $(G, \epsilon) \neq (G, ab) \neq (G, ab; ab) \neq \dots$.

2.3. Protocols

In a distributed setting for gossip, each agent u executes an algorithm of the form:

> **repeat** forever
>   **select** agent $v \in A$ such that condition $\varphi(u, v)$ is satisfied
>   **execute** call uv

where $\varphi(u, v)$ is a condition for the call uv to be executed. Under the joint execution of these agent-based gossip algorithms, the agents will keep calling each other forever. However, for our characterization results, we are only interested in the executions of gossip protocols up to the point where all agents become expert (even if they do not know this). Thus, we decide to use a global algorithm that artificially stops once the agents are experts, defined next.

Definition 13 *(Gossip protocol).* A *gossip protocol* P is a non-deterministic algorithm of the form

> Let an initial gossip graph $G = (A, N, S)$ be given. As long as not all agents are experts and there are $u, v \in A$ such that $\varphi(u, v)$, choose $u, v \in A$ such that the condition $\varphi(u, v)$ is satisfied and the call uv is valid, and execute this call uv.

where, given a call sequence σ, uv is valid in G^σ iff $N^\sigma uv$, and $\varphi(x, y)$ is a property such that for any $u, v \in A$, $\varphi(u, v)$ is a formula in the language of Definition 14.

Note that due to the *as long as not all agents are experts* condition, we need three agents in $\mathcal{G}$ (not two) in order to have infinitely-many gossip states reachable in the execution of a protocol P (like ANY): $(G, ab), (G, ab; ab), \dots$; for the two-agent case, if ab is valid, then ab is a successful sequence (see Definition 17)

and so $ab; ab$ is not P-permitted: the last call ab cannot be made because agents are already experts after the first call ab (see Definition 16).

Definition 14 *(Language $\mathcal{L}$)*. For a given finite set of agents A, the language $\mathcal{L}$ of (instantiated) protocol conditions for epistemic gossip is given by the following BNF:

$$\varphi := \mathtt{S}(a,b) \mid \mathtt{N}(a,b) \mid \mathtt{C}(ab,c) \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid K_a\varphi$$

where a, b, c range over A. The abbreviations $\varphi_1 \rightarrow \varphi_2 := \neg(\varphi_1 \wedge \neg\varphi_2)$, $\varphi_1 \vee \varphi_2 := \neg(\neg\varphi_1 \wedge \neg\varphi_2)$, $\varphi_1 \leftrightarrow \varphi_2 := (\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_2 \rightarrow \varphi_1)$ and $\widehat{K}_a\varphi := \neg K_a\neg\varphi$ will also be used.

The construction $\mathtt{S}(a,b)$ reads *agent a has the secret of b*. $\mathtt{N}(a,b)$ reads *agent a has the number of agent b*. $\mathtt{C}(ab,c)$ reads *ab is a call in the list of calls involving agent c*. Thus, if $c \notin \{a,b\}$, $\mathtt{C}(ab,c)$ is always false, while formulas $\mathtt{C}(ab,a)$ and $\mathtt{C}(ab,b)$ are true in (G,σ) if a call ab has already been made, i.e. if ab is in σ. $K_a\varphi$ reads *agent a knows that ϕ is true*, and finally $\widehat{K}_a\varphi$ reads *agent a considers it possible that ϕ is true*.

Definition 15 *(Semantics)*. Let $\mathcal{G}^{\sim}$ be the asynchronous gossip model for a set of agents A. For any $(G,\sigma) \in \mathcal{G}^{\sim}$ and any formula φ in $\mathcal{L}$, we define $\mathcal{G}^{\sim}, (G,\sigma) \models \varphi$ by induction on φ as follows:

$$
\begin{array}{lll}
\mathcal{G}^{\sim}, (G,\sigma) \models \top & \text{iff} & \text{always} \\
\mathcal{G}^{\sim}, (G,\sigma) \models \mathtt{N}(a,b) & \text{iff} & N^\sigma ab \\
\mathcal{G}^{\sim}, (G,\sigma) \models \mathtt{S}(a,b) & \text{iff} & S^\sigma ab \\
\mathcal{G}^{\sim}, (G,\sigma) \models \mathtt{C}(ab,c) & \text{iff} & ab \in \sigma_c \\
\mathcal{G}^{\sim}, (G,\sigma) \models \neg\varphi & \text{iff} & \text{not } \mathcal{G}^{\sim}, (G,\sigma) \models \varphi \\
\mathcal{G}^{\sim}, (G,\sigma) \models (\varphi_1 \wedge \varphi_2) & \text{iff} & \mathcal{G}^{\sim}, (G,\sigma) \models \varphi_1 \text{ and } \mathcal{G}^{\sim}, (G,\sigma) \models \varphi_2 \\
\mathcal{G}^{\sim}, (G,\sigma) \models K_a\varphi & \text{iff} & \text{for all } (H,\tau) \in \mathcal{G}, \text{ if } (G,\sigma) \sim_a (H,\tau) \text{ then } \mathcal{G}^{\sim}, (H,\tau) \models \varphi
\end{array}
$$

We define $\mathcal{G}^{\approx}, (G,\sigma) \models \varphi$ similarly: just replace $\sim_a$ with $\approx_a$ in the clause for $K_a\varphi$.

The corresponding condition for $\widehat{K}_a$ is the following: $\mathcal{G}^{\sim}, (G,\sigma) \models \widehat{K}_a\varphi$ iff there is $(H,\tau) \in \mathcal{G}$ such that $(G,\sigma) \sim_a (H,\tau)$ and $\mathcal{G}^{\sim}, (H,\tau) \models \varphi$.

The condition $\phi(x,y)$ in a gossip protocol is of the form $K_x\psi(x,y)$ or $\widehat{K}_x\psi(x,y)$. Note that the agent always considers the actual state (G,σ) as possible but she need not know which is the actual state. Informally, each of the agents executing a protocol P stores a list of secrets already learnt, a list of numbers known and the list of calls she was involved in. In the synchronous case, the agent also stores the number of clock ticks. This information enables each agent a to identify her actual $\approx_a$ or $\sim_a$-equivalence class of possible gossip states. Given a gossip state (G,σ) in this equivalence class, before attempting a call ab to agent b, agent a checks if $(G,\sigma) \models \varphi(a,b)$.

Definition 16 *(Permitted call; extension)*. Let P be a protocol given by condition $\varphi(x,y)$ and G be an initial gossip graph.

- A call ab is $\mathsf{P}^{\sim}$-*permitted* in (G,σ) if $\mathcal{G}^{\sim}, (G,\sigma) \models \varphi(a,b)$, call ab is valid in G^σ and not all agents are expert in G^σ.
- A call sequence σ is $\mathsf{P}^{\sim}$-permitted in G if each call σ_{n+1} is $\mathsf{P}^{\sim}$-permitted in $(G,\sigma|n)$.
- The *extension* $\mathsf{P}^{\sim}_G$ (of protocol P in G) is the set of $\mathsf{P}^{\sim}$-permitted sequences in G.

The definitions of $\mathsf{P}^{\approx}$-permitted calls and the extension $\mathsf{P}^{\approx}_G$ for the synchronous case are analogous. For readability, we omit the superscript $\approx$ and $\sim$: we write P_G for $\mathsf{P}^{\sim}_G$ or $\mathsf{P}^{\approx}_G$.

We say that σ is a P-*maximal* sequence on G iff it is a maximal sequence in P_G (that is, a sequence σ in P_G such that either it is infinite or for all calls ab, the sequence $\sigma; ab$ is not in P_G).

For a fixed set of agents A, we define the extensions of a protocol P to be the sets of pairs (including gossip states but also pairs with infinite sequences):

$$\mathsf{P}^\sim = \left\{ (G, \sigma) : \begin{array}{l} G = (A, N, S) \text{ is an initial gossip graph} \\ \text{and } \sigma \text{ is } \mathsf{P}^\sim\text{-permitted in } G \end{array} \right\} \text{ and}$$

$$\mathsf{P}^\approx = \left\{ (G, \sigma) : \begin{array}{l} G = (A, N, S) \text{ is an initial gossip graph} \\ \text{and } \sigma \text{ is } \mathsf{P}^\approx\text{-permitted in } G \end{array} \right\}.$$

Definition 17 *(Successful, fair).* Given an initial gossip graph G and protocol P, a finite call sequence $\sigma \in \mathsf{P}_G$ is *successful* if in G^σ all agents are experts. Call sequence $\sigma \in \mathsf{P}_G$ is *fair* iff σ is finite or, whenever σ is infinite, then for all xy: if for all $i \in \mathbb{N}$ there is a $j \geq i$ such that xy is P-permitted on $G^{\sigma|j}$, then for all $i \in \mathbb{N}$ there is a $j \geq i$ such that $\sigma_j = xy$.

- P is *strongly successful* on G if all maximal $\sigma \in \mathsf{P}_G$ are successful.
- P is *fairly successful* on G if it is strongly successful on the restriction of P_G to fair call sequences.
- P is *weakly successful* on G if there is a sequence $\sigma \in \mathsf{P}_G$ that is maximal and successful.
- P is *unsuccessful* on G if there is no sequence $\sigma \in \mathsf{P}_G$ that is successful.

Given a collection $\mathcal{I}$ of initial gossip graphs, P is (strongly, fairly, weakly, un-) successful on $\mathcal{I}$ iff P is (strongly, fairly, weakly, un-) successful on every $G \in \mathcal{I}$.

A finite call sequence is fair by definition. If a protocol is fairly successful, then all fair sequences in the extension are finite. Let us observe that strongly successful implies fairly successful and the latter implies weakly successful (see [24] for more details).

Definition 18 *(Gossip problem).* Given a collection $\mathcal{I}$ of gossip graphs and a protocol P, the *gossip problem* is: is P (not, weakly, fairly, strongly) successful on $\mathcal{I}$?

2.4. Assumptions on knowledge

The results proved in this paper assume several properties on the knowledge or ignorance of the gossiping agents. These properties, listed next, are encoded in Definitions 9 and 10.

- Common knowledge of the number of agents $|A|$. That is, all possible gossip states contain the actual number of agents. This prevents agents from stopping or pausing the protocol execution because of a belief that there are no more agents to gossip with.[4]
- Common knowledge of the reliability of the network. It is common knowledge that after a call ab is initialized, it is always completed, that is, all known telephone numbers and secrets are exchanged (i.e. equations $S_a^{ab} = S_a \cup S_b$ and $N_a^{ab} = N_a \cup N_b$ are common knowledge).
- Common knowledge of group ignorance of others' initial telephone contacts. It is common knowledge that each agent does not know what numbers are initially known by any other agent.

---

[4] For simplicity, we also assume common knowledge of the agents' identities, i.e. of A. This is not strictly needed though: each agent could learn each others' names together with their numbers, and assign random names before that. Either option gives rise to the same gossip graphs, if in the latter option we label the gossip graphs with agent variables $\{\mathbf{a}, \mathbf{b}, \ldots\}$, e.g. reading $\mathbf{a} = $ "the agent with name a, whatever is called by each agent".

- Common knowledge that agents only know their own secret initially.
- Common knowledge of group ignorance of others' protocols. It is common knowledge that any agent a considers that the actual sequence σ could be any valid sequence σ' such that $\sigma_a = \sigma'_a$. That is, agents do follow the same protocol P but they do not know it (Definitions 9 and 10 do not depend on a specific protocol).

2.5. Impact of conditions on extensions

In this subsection, we study the relations between syntactic restrictions on the conditions of protocols and their corresponding extensions.

Lemma 19. *For any gossip state $(G, \sigma) \in \mathcal{G}$ and any formula ψ in the language $\mathcal{L}$, if ψ contains no modal operator K_a (for $a \in A$), we have:*

$$\mathcal{G}^{\approx}, (G, \sigma) \models \psi \quad \text{iff} \quad \mathcal{G}^{\sim}, (G, \sigma) \models \psi.$$

Proof. The proof is by induction on ψ. (Base Case) By Definition 15, the claim clearly holds for formulas $N(u, v)$, $S(u, v)$ and $C(uv, w)$. (Ind. Case) Assuming that the claim holds for ψ and ψ', it will also hold for the formulas $\neg\psi$ and $(\psi \wedge \psi')$, again by Definition 15. $\square$

Corollary 20. *Let P be an epistemic protocol defined by condition $\varphi(x, y)$ and ψ a modal-free formula. Then,*

$$
\begin{array}{llll}
(i) & \varphi(x, y) = \psi & \text{implies} & \mathsf{P}_G^{\approx} = \mathsf{P}_G^{\sim} \\
(ii) & \varphi(x, y) = \widehat{K}_x\psi & \text{implies} & \mathsf{P}_G^{\approx} \subseteq \mathsf{P}_G^{\sim} \\
(iii) & \varphi(x, y) = K_x\psi & \text{implies} & \mathsf{P}_G^{\sim} \subseteq \mathsf{P}_G^{\approx}.
\end{array}
$$

Proof. For finite call sequences σ, this follows from Lemma 19, together with the fact that $\approx_x \subseteq \sim_x$ for any agent $x \in A$ (Proposition 11). Using this and Definition 16, the statement also holds for infinite sequences σ. $\square$

3. Basic gossip protocols

In this section, we review some gossip protocols already studied in [24]. These protocols are *basic* in the sense that they can be defined by conditions $\varphi(x, y)$ not containing knowledge modalities. Their characterization results, i.e. the classes of graphs where they are strongly, fairly or weakly successful will be used later on during the study of epistemic protocols.

For all these basic protocols, their extensions do not depend on the call mode, as shown in Corollary 20. We write P_G instead of $\mathsf{P}_G^{\sim}$ or $\mathsf{P}_G^{\approx}$.

The Any Call protocol is defined as follows.

Protocol ANY *(Any Call).* $\varphi(x, y) := \top$
While not every agent knows all secrets, randomly select a pair xy such that x knows y's number and let x call y.

The infinite sequence $ab; ab; ab; \ldots$ is ANY-permitted and does not lead to the complete gossip graph where all agents are expert. Thus, ANY is not successful in general (that is, in weakly connected graphs). Still, it is fairly successful, as shown next.

Theorem 21. *[24] For any initial gossip graph G, ANY is fairly successful in G iff G is weakly connected.*

The Learn New Secrets protocol, introduced in Attamah et al. [2] (also studied in Apt et al. [1]) is motivated by the aim of preventing highly redundant calls in call sequences, like those in $ab; ab; ab; \dots$.

Protocol LNS *(Learn New Secrets).* $\varphi(x, y) := \neg \mathsf{S}(x, y)$
While not every agent knows all secrets, select a pair xy such that x knows y's number but not y's secret and perform call xy.

The *skin* of a graph $G = (A, N, S)$ is the set of its terminal nodes T_G. Let $s(G)$ be the result of skinning graph G, i.e. removing all terminal nodes from G. That is, $s(G) = (B, N', S')$ where $B = A \setminus T_G$, with $N' = N \cap B^2$ and $S' = S \cap B^2$.

Definition 22 *(Sun graph).* A *sun graph* is a graph G such that $s(G)$ contains a single strongly connected component.

The previous graph G is a sun graph. Observe that $B = \{a, b, c, d\}$ is a strongly connected component and nodes e, f do not have outgoing edges, so $s(G) = (B, N', S')$.

Sun graphs describe natural scenarios where a core B of agents, more or less well-connected to each other, try to coordinate also with a group of outsiders $A \setminus B$ known by some core member. Gossiping is useful to coordinate an expansion of the initial core B. For an example, a start-up company B in need for funds wants to have a joint meeting with possible investors $A \setminus B$ each profiled and invited by some member of B.

Theorem 23. *[24] For any initial gossip graph G, LNS is strongly successful on G iff G is a sun graph.*

An initial gossip graph $G = (A, N, S)$ is a *tree* iff $(A, (N \setminus I_A)^{-1})$ is a directed rooted tree, i.e., from the perspective of the relation $N \setminus I_A$, there is unique node r called *root* such that for any node x there is a unique directed path from x to r. Any node not accessible from other nodes is a *leaf*.

Definition 24 *(Bush, double bush).* A *bush* is a tree where the root r has at least two predecessors, that is there exists $x \neq y$ such that Nxr and Nyr.

A *double bush* is a graph (A, N) such that there exists a bush $G_b = (A_b, B)$ with root b, a bush $G_d = (A_d, D)$ with root d, and an agent $c \in A$ such that:

- $A = A_b \cup A_d$;
- $A_b \cap A_d = \{c\}$;
- $N = B \cup D$;
- c is a leaf in both G_b and G_d
- $(c, b) \in B$ and $(c, d) \in D$.

A gossip graph $G = (A, N, S)$ is a *(double) bush* if the graph (A, N) is a (double) bush.

Fig. 2 (right) depicts the general shape of a double bush.

Theorem 25. *[24] LNS is weakly successful on a weakly connected gossip graph G iff G is neither a bush nor a double bush.*

Recall that LNS is strongly successful only in sun graphs (Theorem 23). The gossip graph with Nab, Nbc depicted after LNS in the Introduction is neither a sun graph, nor a bush or a double bush. Hence, LNS is weakly successful on it (indeed, consider $ab; ac; bc$), but not strongly successful ($bc; ab$).

Fig. 2. (Left) A bush with root a and leaves e, f, c. (Right) A double bush consists of two bushes sharing a leaf c directly connected to both roots.

Finally, in the Call Me Once protocol, any two agents cannot call each other more than once during the execution of the protocol.

Protocol CO *(Call Me Once).* $\varphi(x, y) := \neg C(xy, x) \wedge \neg C(yx, x)$
Agent x may call agent y if x knows y's number and there was no prior call between x and y.

Theorem 26. *[24] The* CO *protocol is strongly successful on an initial gossip graph G iff G is weakly connected.*

Conditions $\varphi(x, y)$ of the protocols ANY, LNS, CO are Boolean formulas. Checking condition $\varphi(x, y)$ can be implemented in linear-time in the size of $\varphi(x, y)$.

Since our epistemic agents are introspective, the basic protocols ANY, LNS and CO are equivalent to their epistemic counterparts: replacing $\varphi(x, y)$ by $K_x \varphi(x, y)$ does not change their extensions:

$$
\begin{array}{llcl}
\text{(ANY)} & \top & \equiv & K_x \top \\
\text{(LNS)} & \neg S(x, y) & \equiv & K_x \neg S(x, y) \\
\text{(CO)} & \neg C(xy, x) \wedge \neg C(yx, x) & \equiv & K_x (\neg C(xy, x) \wedge \neg C(yx, x))
\end{array}
$$

4. Possible learning protocols

In the protocols considered in this section, an agent can perform a call if he considers it possible that either himself or the agent being called will learn at least one new secret. Conditions $\varphi(x, y)$ are of the form $\widehat{K}_x \psi$.

4.1. Protocol PIG — Possible Information Growth

The protocol Possible Information Growth simply asks the caller not to make calls that he knows would be redundant for both agents.

Protocol PIG *(Possible Information Growth).* $\varphi(x, y) := \widehat{K}_x \bigvee_{z \in A} (S(x, z) \leftrightarrow \neg S(y, z))$
Call xy can be made if x knows y's number and if x considers it possible that there is a secret known by one of x, y but not the other.

Example 27. In [2], the infinite PIG$^{\approx}$-execution $ab; cd; ab; cd; \dots$ (also in PIG$^{\sim}$) is described for the complete gossip graph for 4 agents:

Every time a considers it possible that a call bd was just made (instead of the actual cd), a new call ab can be made in $\mathsf{PIG}^{\approx}$; similarly, every time c considers it possible that a call bd was just made (instead of ab), a new call cd can be made; and so on ad infinitum.

Proposition 28. *The following hold:*

1. $\mathsf{PIG}^{\sim}$ *is strongly successful on G iff G has at most 2 agents and is weakly connected;*
2. $\mathsf{PIG}^{\approx}$ *is strongly successful on G iff G has at most 3 agents and is weakly connected.*

Proof. Right to left of 1. and 2. are easy (for 2., note that any agent c not involved in a call $\overline{ab}$ knows that $\overline{ab}$ occurred).

Conversely, if G is not weakly connected, then by Corollary 6 both $\mathsf{PIG}^{\sim}$ and $\mathsf{PIG}^{\approx}$ are not strongly successful. Now, we prove left to right for G weakly connected.

1. If G has more than 2 agents, let a, b and c be distinct agents such that Nab and either Nbc or Ncb. The call sequence $\sigma = ab; ab; \ldots$ is unsuccessful and in $\mathsf{PIG}_G^{\sim}$: indeed for each $\sigma|n = (ab)^n$ we have $(G, (ab)^n) \sim_a (G, (ab)^{n-1}; \overline{bc})$, which makes $\sigma|(n+1) \in \mathsf{PIG}_G^{\sim}$.

2. If G has more than 3 agents and is weakly connected, we can prove that G contains one of the following patterns[5]:

$$
\begin{array}{cc}
a \text{-----} b \text{-----} c \text{-----} d & \quad \begin{array}{c} d \\ \vert \\ a \text{-----} b \text{-----} c \end{array} \\[1em]
\text{pattern of type 1} & \quad \text{pattern of type 2}
\end{array}
$$

where dashed lines represent edges in $N \cup N^{-1}$.

To show that $\mathsf{PIG}^{\approx}$ is not strongly successful in those G, we prove by cases that $\mathsf{PIG}_G^{\approx}$ contains an infinite sequence σ (similar to that in Example 27):

(Case pattern of type 1) Consider $\sigma = \overline{ab}; \overline{cd}; (ab; dc)^{\omega}$. After the first calls $\overline{ab}; \overline{cd}$, agent a keeps calling b forever (resp. d keeps calling c) with the (false) hope that the previous call was $\overline{bc}$, instead of an actual dc call (resp. an actual ab call).

(Case pattern of type 2) Consider $\sigma = \overline{ab}; \overline{bc}; (ab; cb)^{\omega}$. After the first calls $\overline{ab}; \overline{bc}$, agents a and c happen to alternate calls to b with the (false) hope that the previous call was $\overline{bd}$ instead of the actual cb or, resp., ab call.

Let us provide the formal details for the case pattern of type 1 (we omit the details for the case pattern of type 2 since the proof is analogous). In the first place, the first two calls are clearly in $\mathsf{PIG}_G^{\approx}$. Let now $n > 2$ be arbitrary and assume as inductive hypothesis that $\sigma|n \in \mathsf{PIG}_G^{\approx}$. If n is even, say $n = 2k$, then $\sigma|n$ is of the form $\sigma|2k = \overline{ab}; \overline{cd}; (ab; dc)^{k-1}$ and so we have

$$
(G, \overline{ab}; \overline{cd}; (ab; dc)^{k-2}; ab; dc) \quad \approx_a \quad (G, \overline{ab}; \overline{cd}; (ab; dc)^{k-2}; ab; \overline{bc})
$$

which makes the call sequence $\sigma|(n+1) = \overline{ab}; \overline{cd}; (ab; dc)^{k-1}; ab$ to be in $\mathsf{PIG}_G^{\approx}$ as well. If n is odd, say $n = 2k+1$, then $\sigma|n$ is of the form $\sigma|(2k+1) = \overline{ab}; \overline{cd}; (ab; dc)^{k-1}; ab$ and so

$$
(G, \overline{ab}; \overline{cd}; (ab; dc)^{k-1}; ab) \quad \approx_d \quad (G, \overline{ab}; \overline{cd}; (ab; dc)^{k-1}; \overline{bc})
$$

which again gives $\sigma|(n+1) = \overline{ab}; \overline{cd}; (ab; dc)^{k-1}; ab; dc$ to be in $\mathsf{PIG}_G^{\approx}$. $\quad\square$

---

[5] In order to see it, let us consider the execution of Depth-First-Search [5] on the undirected graph corresponding to G from an agent r. Let us consider the tree corresponding to that Depth First Search. Either there is a 3-chain from r (and we get a pattern of type 1), or a 2-chain from r and a 1-chain from r (and we still get a pattern of type 1) or the root r has 3 children (and we get a pattern of type 2).

Theorem 29. *Let G be an initial gossip graph. Then,*

$$\textsf{PIG} \ \textit{is fairly successful in } G \quad \textit{iff} \quad G \textit{ is weakly connected}$$

Proof. ($\Rightarrow$) This is just Corollary 6. ($\Leftarrow$) Let σ be a maximal $\textsf{PIG}$-permitted sequence.

(Case: σ is infinite) It suffices to prove that σ is not fair. There is then a prefix $\tau \sqsubset \sigma$ such that for all $\tau \sqsubset \tau' \sqsubset \sigma$ we have: $S^{\tau'} = S^\tau \neq A^2$ and $N^{\tau'} = N^\tau$. There exist $x, y \in A$ such that $N^\tau xy$ and $S_x^\tau \neq S_y^\tau$, otherwise we would have $S^\tau = A^2$. The call xy is permitted after τ but will not be made after it. Thus, σ is not fair.

(Case: σ is finite) It is enough to prove that there are $x, y \in A$ such that $N^\sigma xy$ implies $S_x^\sigma \neq S_y^\sigma$. (Since then σ would not be $\textsf{PIG}$-maximal.) Towards a contradiction, assume the contrary: for all x, y if $N^\sigma xy$ then $S_x^\sigma = S_y^\sigma$. Since G is weakly connected, by the transitivity of $(N^\sigma \cup (N^\sigma)^{-1})^*$ and of equality, the latter assumption implies that $S_x^\sigma = S_y^\sigma$ for all $x, y \in A$. From this, together with the fact that $S^\sigma xx$ for all $x \in A$, we conclude that $S_x^\sigma = A = S_y^\sigma$ for all $x, y \in A$, in contradiction with the assumption that σ is not successful.

It only remains to be shown that the set of successful $\textsf{PIG}$-permitted sequences in G is non-empty. To show that there exists at least one such successful $\textsf{PIG}$-permitted sequence, let σ be a minimum length successful $\textsf{ANY}$-permitted sequence, see e.g. [24]. Note that since σ has minimum length, it does not contain calls between expert agents. We claim that σ is a $\textsf{PIG}$-sequence. Trivially $\sigma|1$ is $\textsf{PIG}$-permitted. Let $\sigma|n$ be $\textsf{PIG}$-permitted and $\sigma_{n+1} = xy$. There are two cases.

$(S_x^{\sigma|n} = A)$: By the assumption $S_y^{\sigma|n} \neq A$ and so $\sigma|n; xy$ is $\textsf{PIG}$-permitted.
$(S_x^{\sigma|n} \neq A)$: Let $z \in (A \setminus S_x^{\sigma|n})$, then $(G, \sigma|n) \sim_x (G \cdot \sigma|n; zy)$ and $S_x^{\sigma|n} \neq S_y^{\sigma|n;zy}$. Thus, $\sigma|n; xy$ is $\textsf{PIG}$-permitted. $\square$

The condition defining the $\textsf{PIG}$ protocol, namely that *it is possible that either agent involved in the call learns some secret*, can be split into two conditions: that the caller might learn something; that the callee might learn something. Each of these conditions defines a new protocol.

4.2. Protocol $\textsf{TSS}$ — *Tell (me) Some Secret*

First we consider the protocol where an agent will make a call only if she might learn something.

Protocol $\textsf{TSS}$ *(Tell me Some Secret protocol).* $\varphi(x, y) := \widehat{K}_x \bigvee_{z \in A}(\textsf{S}(y, z) \wedge \neg\textsf{S}(x, z))$
Call xy can be made if x knows y's number and x considers it possible to learn some new secret from y.

For an example of an infinite $\textsf{TSS}$ call sequence, consider again the sequence $ab; cd; ab; cd; \ldots$ as in Example 27. Thus, the characterization will be in terms of fairly successful sequences.

Theorem 30. *Let G be an initial gossip graph. Then,*

$$\textsf{TSS}^\approx \ \textit{is fairly successful in } G \ \textit{iff} \ G \textit{ is weakly connected.}$$

Proof. ($\Rightarrow$) This is just Corollary 6. ($\Leftarrow$) We prove that all fair maximal $\textsf{TSS}^\approx$-sequences are successful, so let $\sigma \in \textsf{TSS}_G^{\approx}$ be fair and maximal. Also, let $\tau \sqsubseteq \sigma$ be a finite prefix of σ satisfying: $S^\tau = S^{\tau'}$ and $N^\tau = N^{\tau'}$ for any finite τ' with $\tau \sqsubseteq \tau' \sqsubseteq \sigma$. The proof is by cases.

(Case: there is an expert in G^τ) Then let $B \subseteq A$ be the set of experts in G^τ. If $B \neq A$, then there are two agents $x \in B$ and $y \in A \setminus B$ such that $\overline{xy}$ is a call in τ (otherwise, no expert agent would exist at all). Hence, the call yx is both valid ($\overline{xy} \in \tau$ implies $N^\tau yx$) and $\mathsf{TSS}^\approx$-permitted in any finite τ' with $\tau \sqsubseteq \tau' \sqsubseteq \sigma$. By fairness, the call yx is made in some such τ', contradicting the above assumption $B \neq A$.

(Case: there is no expert in G^τ) Then, select $x, y \in A$ with minimum $(N \cup N^{-1})$-distance and such that $\neg S^\tau xy$. So there is an agent z (the predecessor of y in that minimum distance path, e.g. $z = x$ if the path has length one) such that $S^\tau xz$ and Nzy or Nyz.

If Nzy, then by Lemma 5 $N^\tau xy$, and so the call xy is $\mathsf{TSS}^\approx$-permitted in any finite τ' with $\tau \sqsubseteq \tau' \sqsubseteq \sigma$, so the call xy will be made at some such τ', contradicting the selection of τ.

If Nyz, then since y is not expert, there is an agent u such that $\neg S^\tau yu$.

(Case 1) There is a finite call sequence, say of the form $\tau'; \nu$ with $\tau \sqsubseteq \tau'; \nu \sqsubseteq \sigma$ such that y does not occur in ν. Then, $G^{\tau'; \nu} \approx_y G^{\tau'; uz}$, so for any finite τ'' with $\tau'; \nu \sqsubseteq \tau'' \sqsubseteq \sigma$, we have that yz is $\mathsf{TSS}^\approx$-permitted. By fairness, this call yz will be made after $\tau'; \nu$ in σ. If x is in the call yz (that is, $x = z$), then we reach a contradiction with the selection of τ. Otherwise, if $x \notin \{y, z\}$, then by the same reasoning the call xz is in $\mathsf{TSS}^\approx$-permitted for any suffix of τ containing the call yz; by fairness, the call xz is eventually made in σ (where x learns the secret of y), so again we reach a contradiction with the selection of τ.

(Case 2) Each call after τ (in σ) contains agent y. For each finite $\tau'; \nu$ with $\tau \sqsubseteq \tau'; \nu \sqsubseteq \sigma$ we have $(G, \tau'; \nu) \approx_x (G, \tau'; yz)$, so the call xz is $\mathsf{TSS}^\approx$-permitted after τ, but this call xz is never made after τ (in σ), contradicting the assumption of fairness. $\square$

Corollary 31. *For any initial gossip graph G,*

$$\mathsf{TSS}^\sim \text{ is fairly successful in } G \quad iff \quad G \text{ is weakly connected.}$$

Proof. The same proof of Theorem 30 works for the asynchronous case: we only have to replace $\approx$ by $\sim$, and replace the paragraph for the sub-case Nyz with the following: if Nyz, then since y is not expert, there is an agent u such that $\neg S^\tau yu$. Now for each finite τ' with $\tau \sqsubseteq \tau' \sqsubseteq \sigma$, we have $G^{\tau'} \sim_y G^{\tau'; uz}$ so yz is $\mathsf{TSS}^\sim$-permitted after any such finite τ' with $\tau \sqsubseteq \tau' \sqsubseteq \sigma$. $\square$

4.3. Protocol HSS — Hear Some Secret

The other protocol defined from PIG consists in calls made under the pretext of (possibly) teaching some secrets to the other agents. In other words, the Hear Some Secret protocol (HSS) is an instruction for calling in which an agent x will make a call to y if x knows the secret of some agent z, and x does not know that y has the secret of z.

Protocol HSS *(Hear Some Secret protocol).* $\varphi(x, y) := \widehat{K}_x \bigvee_{z \in A} (\mathsf{S}(x, z) \wedge \neg \mathsf{S}(y, z))$
Call xy can be made if x knows y's number and considers it possible that she knows some secret that y does not know.

Besides the inclusion $\mathsf{HSS}^\approx \subseteq \mathsf{HSS}^\sim$ (using Corollary 20), the execution of HSS in asynchronous and synchronous gossip models gives rise to different extensions, so $\mathsf{HSS}^\approx \subset \mathsf{HSS}^\sim$, as shown next.

Example 32. Let $G = (A, N, S)$ be the initial gossip graph given by Nab and Nbc; see the graph in the top-left corner of the next picture.
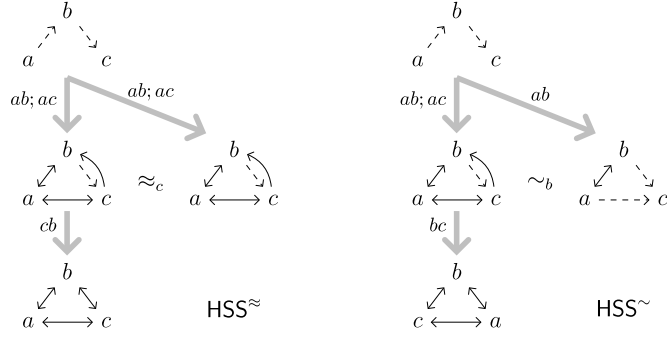
(Left) After executing the $\mathsf{HSS}_G^{\approx}$-sequence $ab; ac$, agent c knows that b could learn her secret, so the call cb is $\mathsf{HSS}_G^{\approx}$-permitted, leading to a successful execution in $G^{ab;ac;cb}$.

(Right) After executing the $\mathsf{HSS}_G^{\sim}$-sequence $ab; ac$, agent b considers it possible that only the call ab has been made; thus, a call bc is $\mathsf{HSS}_G^{\sim}$-permitted after $ab; ac$, since in the state G^{ab} agent c would learn b's secret. This learning will in fact not occur: the call bc will only lead to agent b learning c's secret, and thus to the successful state $G^{ab;ac;bc}$. The same sequence $ab; ac; bc$, though, is not $\mathsf{HSS}_G^{\approx}$-permitted, since b would be aware that $ab; \overline{ac}$ already took place.

According to both protocols $\mathsf{HSS}^{\approx}$ and $\mathsf{HSS}^{\sim}$, if there are n agents, then each of those can call any of the $n-1$ other agents in order to inform of at most $n-1$ secrets that they might not know. Therefore, the length of any HSS sequence is less than n^3, and so any HSS-maximal sequence σ is finite. The characterization result, then, is in terms of strong success.

Theorem 33. *For any initial gossip graph G,*

$$\mathsf{HSS}^{\approx} \text{ is strongly successful in } G \quad \text{iff} \quad G \text{ is weakly connected.}$$

Proof. Let $G = (A, N, S)$ be a weakly connected graph. For the $\Rightarrow$ direction, we use Corollary 6. For the $\Leftarrow$ direction, let σ be a maximal $\mathsf{HSS}^{\approx}$-permitted sequence in G and assume, towards a contradiction, that $S^{\sigma} \neq A^2$. Select agents a and b with minimum $(N \cup N^{-1})$-distance satisfying $\neg S^{\sigma}ab$; and let π be such a minimum distance path between a and b.

By the maximality of σ, we must also have that $\neg N^{\sigma}ab$. Again by the maximality of σ, we must have $\neg N^{\sigma}ba$ as well (otherwise $\sigma; ba$ would be $\mathsf{HSS}_G^{\approx}$-permitted), and so we also have that $\neg S^{\sigma}ba$.

Select any c in π different from a and b (such c must exist in the weakly connected gossip graph G: the previous fact $\neg N^{\sigma}ab$ implies $\neg Nab$, and similarly $\neg N^{\sigma}ba$ implies $\neg Nba$).

By the minimality of π, we must have that both $S^{\sigma}ca$ and $S^{\sigma}cb$. Recall that $(G, \sigma) \approx_c (G, \sigma)$ and in the latter $b \in S_c^{\sigma} \setminus S_a^{\sigma}$. Thus, $\sigma; ca$ is $\mathsf{HSS}^{\approx}$-permitted in G, contradicting the maximality of σ. □

The same proof can be used in asynchronous gossip models, just replacing $\mathsf{HSS}^{\approx}$ by $\mathsf{HSS}^{\sim}$.

Corollary 34. *For any initial gossip graph G,*

$$\mathsf{HSS}^{\sim} \text{ is strongly successful in } G \text{ iff } G \text{ is weakly connected.}$$

4.4. Protocol HMS — Hear My Secret

The Hear My Secret protocol (HMS) is proposed in Apt et al. [1], where it is proved (for the *merge-and-learn* call mode) that it is successful in any complete graph K_n, i.e. provided $N = A^2$. In this protocol, an

agent x will make a call to y if x does not know that y has the secret of x. Hence, HMS can be seen as the dual of LNS, and also as a particular case of HSS.

Protocol HMS *(Hear My Secret Protocol).* $\varphi(x,y) := \widehat{K}_x \neg \mathsf{S}(y,x)$
Call xy is possible if x knows y's number and x considers it possible that y does not know x's secret.

The following results are similarly proved for synchronous and asynchronous gossip models $\mathcal{G}^{\approx}$ and $\mathcal{G}^{\sim}$. Instead of $\mathsf{HMS}^{\approx}_G$ and $\mathsf{HMS}^{\sim}_G$, we simply write HMS in the proofs. Note that only finitely-many calls can be made in any HMS sequence, so in particular all HMS-maximal sequences are finite.

Lemma 35. *For any initial gossip graph $G = (A, N, S)$, if σ is an HMS-maximal call sequence on G, then $S^\sigma = N^\sigma$.*

Proof. Assume that σ is HMS-maximal on G. By Fact 4, we only need to show that $N^\sigma \subseteq S^\sigma$, so let $N^\sigma ab$ be arbitrary. By HMS-maximality, we must have that $S^\sigma ba$ (as otherwise $(G, \sigma) \approx_a (G, \sigma)$ and $(G, \sigma) \sim_a (G, \sigma)$ would make $\sigma; ab$ permitted); again by Fact 4, $N^\sigma ba$, and finally again by HMS-maximality, we must have $S^\sigma ab$. $\square$

Corollary 36. *If σ is an HMS-maximal call sequence on G, then $S^\sigma \circ N^* = S^\sigma$.*

Proof. We have that $S^\sigma \subseteq S^\sigma \circ N^*$ by definition of N^*. We now prove that $S^\sigma \circ N^* \subseteq S^\sigma$: let $(x, y) \in S^\sigma \circ N^*$. Then for some $k \in \mathbb{N}$, $(x, y) \in S^\sigma \circ N^k$. We get from Lemma 5 plus Lemma 35 that $S^\sigma \circ N \subseteq S^\sigma$. Applying this fact k times yields $(x, y) \in S^\sigma$. $\square$

Corollary 37. *If σ is an HMS-maximal call sequence on G, then S^σ and N^σ are symmetrical.*

Proof. Clearly, $(G, \sigma) \approx_a (G, \sigma)$ (resp. with $\sim_c$), and from $S^\sigma ab$ and $\neg S^\sigma ba$, we obtain $N^\sigma ab$. Now the graph G^σ satisfies $\neg S^\sigma ba$, so $\sigma; ab$ is also in $\mathsf{HMS}^{\approx}_G$ (resp. in $\mathsf{HMS}^{\sim}_G$), contradicting the assumption that σ is maximal. From $S^\sigma = (S^\sigma)^{-1}$ and Lemma 35, we conclude that N^σ is also symmetrical: $N^\sigma = S^\sigma = (S^\sigma)^{-1} = (N^\sigma)^{-1}$. $\square$

Definition 38 *(Sun* graph).* An initial gossip graph $G = (A, N, S)$ is a *sun** iff there is a strongly connected component $B \subseteq A$ such that for each $x \in B$ and each $y \in A$, we have $N^* xy$.

Sun* graphs obtain as a natural generalization of sun graphs by closure under transitivity: the outsiders can also bring to the group other outsiders not directly known by the core group B, who can bring more outsiders, and so on. For example, some activist group of vegans B try to strengthen their power by coordinating a dinner with other vegans they know about, and making this invitation extensible to others.

Theorem 39. *The HMS protocol is strongly successful in any sun* graph.*

Proof. Let G be a sun* graph, and let B be some strongly connected component in G such that its elements b satisfy $N^* by$ for any $y \in A$. Also, let σ be an HMS-maximal call sequence on G. Let $x, y \in A$. We show that $S^\sigma xy$.

(Case $x \in B$) Then, $N^* xy$, and because Sxx, and also $S^\sigma xx$, we obtain that $(x, y) \in S^\sigma \circ N^*$. By Corollary 36 it follows that $S^\sigma xy$.

(Case $x \notin B$) If $y \in B$, then from the previous case we have $S^\sigma yx$ and so $N^\sigma yx$. If $\neg S^\sigma xy$, then yx is HMS-permitted after σ. This contradicts the maximality of σ. Hence, x knows all the secrets in B. If $y \notin B$, let $u \in B$. Thus, $N^* uy$, and so $S^\sigma xu$. These two facts jointly imply that $S^\sigma xy$ using Corollary 36. $\square$

But the HMS protocol is also strongly successful on other graphs:

Example 40. Let G be the initial gossip graph given by: Nab and Ncb. It is not a sun* but HMS is strongly successful on G: its extension consists in $ab; cb; ca$ and $cb; ab; ac$.

The characterization of gossip graphs where the HMS protocol is weakly successful can be more easily established.

Theorem 41. *Let $G = (A, N, S)$ be an initial gossip graph. Then,*

$$\text{HMS}^{\approx} \text{ is weakly successful in } G \quad \text{iff} \quad G \text{ is weakly connected.}$$

Proof. ($\Rightarrow$) This is just Corollary 6. ($\Leftarrow$) The proof is based on induction on the number of agents $n = |A|$. (Base case $n = 2$) Let $A = \{a, b\}$ and without loss of generality suppose Nab. The sequence ab is HMS$^{\approx}$-permitted in G, and after it a and b are expert. (Ind. Step $n \mapsto n+1$) Assume (Ind. Hyp.) that for any weakly connected gossip graph with n agents there is an HMS$^{\approx}_G$ successful execution in this graph. Let $G = (A, N, S)$ be a weakly connected graph with $n + 1$ agents. There exists an agent x such that, for $A_{\bar{x}} = A \setminus \{x\}$, the restriction of G to $A_{\bar{x}}$, namely, the graph

$$G' = \left(A_{\bar{x}}, \ N \cap A_{\bar{x}}^2, \ S \cap A_{\bar{x}}^2 \right) \quad \text{is weakly connected.}^{[6]}$$

By the Ind. Hyp. there is an HMS$^{\approx}$-maximal sequence σ' on G' which is successful. Since G is weakly connected, a call involving x after σ' is possible, so let $\overline{xy}$ be such a call. After σ', agent y is an expert among $A_{\bar{x}}$ and the call $\overline{xy}$ makes x expert among A. After $\sigma'; \overline{xy}$ let agent x call all the agents in $A \setminus \{x, y\}$. The new sequence σ is HMS$^{\approx}$-permitted and every one is expert after it. $\square$

Combining this result with Proposition 11, we obtain the characterization of weak success for HMS under asynchronous gossip models.

Corollary 42. *Let $G = (A, N, S)$ be an initial gossip graph. Then,*

$$\text{HMS}^{\sim} \text{ is weakly successful in } G \text{ iff } G \text{ is weakly connected.}$$

5. Necessary learning protocols

Contrary to the previous section, where conditions were of the form $\widehat{K}_x \psi$, here we consider protocols whose conditions rely on the K_x modality. A protocol first considered in Attamah et al. [2] is the so-called Known Information Growth KIG, where one agent x can perform a call xy if she knows that either herself or the agent called y will learn at least one new secret. This informal condition has two readings: *de dicto* and *de re*, giving rise to two protocols.[7]

Protocol KIGd *(Known Information Growth de* dicto*).* $\varphi(x, y) := K_x \bigvee_{z \in A} (\text{S}(x, z) \leftrightarrow \neg\text{S}(y, z))$
Call xy can be made if x knows y's number and if x knows that some secret is to be learnt either by x or y; but x might not know whose secret will that be or who will learn it.

---

[6] For proving the existence of such an x, consider a spanning tree of the undirected graph corresponding to G. Then let x be a leaf of this tree.

[7] The *de dicto/de re* distinction is well-known in linguistics [20] and philosophical logic [9]: the famous example *every man loves a woman* can either mean that there is a woman every man loves (*de re* or $\exists\forall$-reading) or that for every man there exists a woman (typically but not necessarily different for different men) that he loves (*de dicto* or $\forall\exists$-reading). This distinction has also been used in epistemic logical settings, see e.g. [18].

Protocol KIGr *(Known Information Growth* de re*).* $\varphi(x,y) := \bigvee_{z \in A} K_x(\mathsf{S}(x,z) \leftrightarrow \neg\mathsf{S}(y,z))$
Call xy can be made if x knows y's number and there is some agent z such that x knows that its secret will be learnt by either x or y.

Let us observe that the *de dicto/de re* distinction is superfluous in Section 4. This is due to those protocols' conditions being of the form $\varphi(x,y) = \widehat{K}_x\psi$: the $\widehat{K}_x$-modality commutes with the disjunction $\bigvee_{z \in A}$; that is, $\widehat{K}_x \bigvee_{z \in A} \psi$ is equivalent to $\bigvee_{z \in A} \widehat{K}_x\psi$. Also, note that only finitely-many calls can be made in any KIGd or KIGr execution, so all maximal call sequences σ are finite.

5.1. Synchronous calls in KIG protocols

Under the synchronous mode of telephone calls, the protocols $\mathsf{KIGr}_G^{\eqapprox}$ and $\mathsf{KIGd}_G^{\eqapprox}$ are different from each other (and also different from LNS; compare with the asynchronous case below). Two properties of both protocols in synchronous gossip models $\mathsf{KIGr}_G^{\eqapprox}$ and $\mathsf{KIGd}_G^{\eqapprox}$ are expressed by the following Propositions 43 and 44.

Proposition 43. *If $G = (A, N, S)$ is an initial gossip graph and σ is $\mathsf{KIGd}^{\eqapprox}$-maximal (resp. $\mathsf{KIGr}^{\eqapprox}$-maximal) on G, then $S^\sigma = N^\sigma$.*

Proof. Let G be an initial gossip graph, and let there be x, y with $N^\sigma xy$ and $\neg S^\sigma xy$. Then the call xy is $\mathsf{KIGd}^{\eqapprox}$-permitted (resp. $\mathsf{KIGr}^{\eqapprox}$-permitted) in G^σ, which is in contradiction with the maximality of σ. This shows $N^\sigma \subseteq S^\sigma$. The property $S^\sigma \subseteq N^\sigma$ follows from Fact 4. Together, this gives $S^\sigma = N^\sigma$. □

Proposition 44. *If σ is a $\mathsf{KIGd}^{\eqapprox}$-maximal (resp. a $\mathsf{KIGr}^{\eqapprox}$-maximal) call sequence on an initial gossip graph G, then $S^\sigma \circ N^* = S^\sigma$.*

Proof. The proof is exactly as in Corollary 36, simply replace Lemma 35 by Proposition 43. □

Theorem 45. *The $\mathsf{KIGd}^{\eqapprox}$ and $\mathsf{KIGr}^{\eqapprox}$ protocols are strongly successful on any initial gossip graph G that is a sun graph.*

Proof. Let $G = (A, N, S)$ be a sun graph. Let σ be a $\mathsf{KIGd}^{\eqapprox}$-maximal (resp. $\mathsf{KIGr}^{\eqapprox}$-maximal) call sequence on G. Let $x, y \in A$. We first show that $S^\sigma xy$.

If x is in $s(G)$, then N^*xy. Because of Sxx, also $S^\sigma xx$, and therefore $(x, y) \in S^\sigma \circ N^*$. By Proposition 44 it follows that $S^\sigma xy$.

If x is a terminal node, then by maximality of σ, there is some u with $ux \in \sigma$. This means that $N^\sigma uz$ for some z with Nzx (where possibly $u = z$), because u must have learnt x's number from some such z. Thus, after the call ux, x has the number of some z with Nzx, that is, $N^\sigma xz$. By maximality of σ it follows that $S^\sigma xz$. Since $z \in s(G)$ it follows that N^*zy. From $S^\sigma xz$ and N^*zy we get $(x, y) \in S^\sigma \circ N^*$. By Proposition 44 we then get $S^\sigma xy$, and we are done. □

The previous result is not a full characterization: $\mathsf{KIGr}^{\eqapprox}$ and $\mathsf{KIG}^{\sim}$ are strongly successful in the non-sun gossip graph $G = (\{a, b, c, d\}, N, S)$ defined by Nad, Nbd, Ncd.

Theorem 46. *For any initial gossip graph $G = (A, N, S)$,*

$$\mathsf{KIGr}^{\eqapprox} \text{ and } \mathsf{KIGd}^{\eqapprox} \text{ are weakly successful in } G \text{ iff } G \text{ is weakly connected.}$$

Proof. The proof for $\mathsf{KIGr}^{\eqapprox}$ is exactly the same as that for $\mathsf{HMS}^{\eqapprox}$ (Theorem 41), just replace $\mathsf{HMS}^{\eqapprox}$ by $\mathsf{KIGr}^{\eqapprox}$. The proof for $\mathsf{KIGd}^{\eqapprox}$ simply follows from that for $\mathsf{KIGr}^{\eqapprox}$ together with the fact $\mathsf{KIGr}^{\eqapprox} \subseteq \mathsf{KIGd}^{\eqapprox}$. □

5.2. Asynchronous calls in KIG *protocols*

The situation in asynchronous gossip models $\mathcal{G}^\sim$ is different, as we show next, since for any gossip graph the three protocols LNS, KIGr$^\sim$ and KIGd$^\sim$ collapse into the same extension.

A first observation is that under asynchronous calls, if an agent a does not have a given secret z, then a considers it possible that any agent b does not have it either.

Lemma 47. *Let a gossip graph $G = (A, N, S)$ for three or more agents be given and σ a* KIG$^\sim$ *sequence in G. Let a, b, c be three of the agents in A such that agent a does not know the secret of agent c in G^σ. Then, agent a considers a sequence σ' possible after which agent b does not know the secret of agent c, and knows fewer secrets than under σ, i.e.*

$$(G, \sigma) \sim_a (G, \sigma') \quad \text{and} \quad c \notin S_b^{\sigma'} \subseteq S_b^\sigma \quad \text{(for some } \sigma'\text{)}.$$

Proof. If b does not know the secret of c in G^σ, just let $\sigma' = \sigma$. Therefore, suppose that b knows the secret of c. Let us consider all minimal ways in which b could have learnt the secret of c. This is the set Θ defined by:

$$\Theta = \left\{ \tau \text{ a subsequence of } \sigma \ : \ \begin{array}{l} \tau = \overline{cx_1}, \overline{x_1 x_2}; \ldots; \overline{x_k b} \text{ with } x_i \neq x_j \\ \text{and } b \neq x_i \neq c \text{ for each } 1 \leq i, j \leq k \end{array} \right\}$$

Now define the following procedure for deleting calls from each $\tau \in \Theta$. Select some such $\tau \in \Theta$, and denote by $A(\tau)$ the set of agents occurring in the calls in τ. Observe first that

(i) for any $x_i \in A(\tau)$ no subsequence of calls $\overline{x_i y_1}; \overline{y_1 y_2}; \ldots; \overline{y_k a}$ can exist in σ, after the first call involving x_i is made in τ.

The reason is that otherwise a would learn the secret of c. As a consequence of (i), if we let τ' (resp. σ'') be the result of removing the last call $\overline{x_k b}$ from τ (resp. σ), removing this call $\overline{x_k b}$ will not affect any of the calls $\overline{za}$ involving a in σ. That is, we claim that

(ii) $(G, \sigma) \sim_a (G, \sigma'')$

We show this claim by induction on the calls $\overline{za}$ in σ. We only show the inductive step: let $\overline{za} = \sigma_{l+1}$. Consider the cases:

(Case: $\overline{za}$ occurs before $\overline{x_k b}$) Then (G, σ_l) and (G, σ'_l) are identical and so indistinguishable by both z and a; thus, $(G, \sigma_l; \overline{za}) \sim_a (G, \sigma'_l; \overline{za})$.

(Case: $\overline{za}$ occurs after $\overline{x_k b}$) That is, $\sigma'_l = \sigma_{l+1}$. We can apply (i) to conclude that $(G, \sigma'_l) \sim_z (G, \sigma_{l+1})$ and $(G, \sigma'_l) \sim_a (G, \sigma_{l+1})$, since otherwise σ would contain a subsequence $\overline{x_k b}; \ldots; \overline{yz}; \overline{za}$ transmitting the secret of c to a through x_k, b and z.

Now removing the last call $\overline{x_k b}$ from τ may cause a call to be removed from some other τ'' in Θ. However, because τ'' is (also) a minimal subsequence transmitting the secret of c to b, this call $\overline{x_k b}$ can only be the last call of τ'' too. Thus, we can now repeat the procedure for all remaining sequences τ in Θ that still have their last call in place, and reason as in (ii) at each step of this procedure. Call the resulting sequence σ'. The sequence σ' is indistinguishable for agent a from the sequence σ. In addition, agent b does not know the secret of c after σ'. Finally, since σ' is a subsequence of σ, we also have $S_b^{\sigma'} \subseteq S_b^\sigma$. $\quad \square$

Lemma 48. *Let $G = (A, N, S)$ be any initial graph with $|A| \geq 3$. For any finite valid sequence σ of asynchronous calls in G and two agents $a, b \in A$, there exists a sequence σ' such that*

$$(G, \sigma) \sim_a (G, \sigma') \quad \text{and} \quad S_b^{\sigma'} \subseteq S_a^{\sigma'}$$

Proof. Let $z_1, z_2, \ldots, z_n$ be an enumeration of all agents in A such that agent a does not know their secrets after σ,

$$\forall 1 \leq i \leq n \quad z_i \notin S_a^\sigma$$

We inductively define a sequence σ' such that

$$(G, \sigma) \sim_a (G, \sigma') \quad \text{and} \quad S_b^{\sigma'} \subseteq S_a^{\sigma'}$$

Let $\sigma_0 = \sigma$ and define σ_{k+1} as the sequence of calls where

$$(G, \sigma_k) \sim_a (G, \sigma_{k+1}) \quad \text{and} \quad z_{k+1} \notin S_b^{\sigma_{k+1}} \subseteq S_b^{\sigma_k}$$

According to Lemma 47 these sequences exist. The sequence $\sigma' = \sigma_n$ is our desired sequence. $\quad\square$

As a consequence of the previous results for asynchronous calls, the two Known Information Growth protocols KIGd and KIGr have the same extension as the Learn New Secret protocol LNS.

Theorem 49. *The following equalities hold:*

$$\mathsf{KIGr}^\sim = \mathsf{KIGd}^\sim = \mathsf{LNS}.$$

Proof. We show the following chain of inclusions between the extensions

$$\mathsf{LNS} \quad \subseteq \quad \mathsf{KIGr}^\sim \quad \subseteq \quad \mathsf{KIGd}^\sim \quad \subseteq \quad \mathsf{LNS}$$

The first two inclusions $\mathsf{LNS} \subseteq \mathsf{KIGr}^\sim \subseteq \mathsf{KIGd}^\sim$ are obvious from the conditions of these protocols and Definition 9.

We prove that $\mathsf{KIGd}^\sim \subseteq \mathsf{LNS}$. For this, let σ be an arbitrary $\mathsf{KIGd}^\sim$-permitted sequence of calls in some G. Towards a contradiction, assume there is a call in σ that is not LNS-permitted in G and let ab be the first such call in σ. This means that a knows the secret of b before the call ab. Let $\sigma = \sigma_1; ab; \sigma_2$, hence σ_1 is LNS-permitted in G. By Lemma 48 there is a sequence σ_1' such that

$$(G, \sigma_1) \sim_a (G, \sigma_1') \quad \text{and} \quad S_b^{\sigma_1'} \subseteq S_a^{\sigma_1'}$$

Let $\sigma_1' = \pi_1; X; \pi_2$ where $X = \overline{ca}$ is the last call in σ_1' which contains a.
By Definition 9, $\qquad\qquad\qquad (G, \pi_1; X) \sim_a (G, \pi_1; X; cb)$
hence $\qquad\qquad (G, \sigma_1) \sim_a (G, \sigma_1') = (G, \pi_1; X; \pi_2) \sim_a (G, \pi_1; X; cb; \pi_2)$
therefore $\qquad\qquad\qquad (G, \sigma_1) \sim_a (G, \pi_1; X; cb; \pi_2).$
In $G^{\pi_1; X; cb; \pi_2}$ the secrets known by b and by a are the same. By this fact the call ab is not $\mathsf{KIGd}_G^\sim$-permitted after σ_1 (contradiction). Thus, $\mathsf{KIGd}_G^\sim \subseteq \mathsf{LNS}_G$. $\quad\square$

This result can be combined with Theorems 23 and 25 to obtain the following characterization results:

Corollary 50. *For each* $\mathsf{KIG}^\sim \in \{\mathsf{KIGr}^\sim, \mathsf{KIGd}^\sim\}$ *and any gossip graph* G,

(i) $\mathsf{KIG}^\sim$ *is strongly successful in* G *iff* G *is a sun graph,*
(ii) $\mathsf{KIG}^\sim$ *is weakly successful in* G *iff* G *is neither a bush nor a double bush.*

6. Comparison of protocol extensions

In addition to the characterization results presented in Sections 4 and 5, it is also of interest to study which of these protocols coincide with, or subsume, other protocols in terms of their extensions.

Let us first compare the asynchronous extension of each protocol with its synchronous counterpart.

Corollary 51. *The following inclusions hold:*

$$P_G^{\approx} \subset P_G^{\sim} \quad \text{for each } P \in \{\text{HMS}, \text{HSS}, \text{TSS}, \text{PIG}\}, \text{ and}$$
$$P_G^{\sim} \subset P_G^{\approx} \quad \text{for each } P \in \{\text{KIGr}, \text{KIGd}\}.$$

Proof. The inclusions follow from Corollary 20 together with the definition of each protocol. To see that these inclusions are proper, consider the following examples:

$$(K_3, ab; ab) \in \text{PIG}^{\sim} \setminus \text{PIG}^{\approx} \qquad (K_3, ab; ac; cb) \in \text{KIGr}^{\approx} \setminus \text{KIGr}^{\sim}$$
$$(K_3, ab; ab) \in \text{TSS}^{\sim} \setminus \text{TSS}^{\approx} \qquad (K_3, ab; ac; cb) \in \text{KIGd}^{\approx} \setminus \text{KIGd}^{\sim}$$
$$(K_3, ab; ac; bc; ab) \in \text{HSS}^{\sim} \setminus \text{HSS}^{\approx}$$
$$(K_3, ab; bc; ac) \in \text{HMS}^{\sim} \setminus \text{HMS}^{\approx} \qquad \qquad \square$$

Second, we can compare the basic protocols LNS, CO and ANY from [24] with those studied in this paper. For the PIG protocol:

Proposition 52. *The following inclusions hold:* $\text{CO} \subset \text{PIG}^{\approx} \subset \text{PIG}^{\sim} \subset \text{ANY}$.

Proof. The inclusion $\text{CO} \subset \text{PIG}^{\approx}$ follows from the following characterization of the sets $\text{PIG}^{\sim}$ and $\text{PIG}^{\approx}$:

$$\text{PIG}^{\sim} = \text{ANY} \setminus \{(G, \sigma) \mid \sigma = \tau; \overline{xy}; \ldots; \overline{xy} \text{ and } S_x^{\tau; \overline{xy}} = A\}$$
$$\text{PIG}^{\approx} = \text{ANY} \setminus (\text{PIG}^{\sim} \cup \{(G, \sigma) \mid \sigma = \ldots; xy; \overline{xy}; \ldots\})$$

From this it follows that if $\sigma \notin \text{PIG}_G^{\approx}$ then σ contains two calls between the same pair of agents, so $\sigma \notin \text{CO}_G$. Inclusions $\text{PIG}^{\approx} \subset \text{PIG}^{\sim} \subseteq \text{ANY}$ follow from Corollary 51 and Corollary 20.

The proper inclusions $\text{CO} \subset \text{PIG}^{\approx}$ and $\text{PIG}^{\sim} \subset \text{ANY}$ are shown next:

$$(K_3, ab; bc; ab) \in \text{PIG}^{\approx} \setminus \text{CO} \qquad (K_3, ab; bc; bc) \in \text{ANY} \setminus \text{PIG}^{\sim}. \qquad \square$$

The other non-basic protocols are also included in ANY, but none of them extends CO.

Proposition 53. $\text{CO} \not\subseteq P$, *for any* $P \in \{\text{LNS}, \text{HMS}, \text{HSS}, \text{TSS}, \text{KIGr}, \text{KIGd}\}$.

Proof. For an example of a CO-execution not in the extension of the other protocols, consider the sequence $ab; ca; da; cb; cd$. $\square$

More detailed comparisons between the extensions of the different protocols within each call mode are presented next.

6.1. Comparison of protocols in asynchronous gossip models

In the asynchronous gossip models, the protocols can be arranged according to the inclusions listed next. Fig. 3 illustrates these inclusions and also contains examples showing that the inclusions are proper.

Fig. 3. Inclusions of protocol extensions in asynchronous gossip models.

Corollary 54. *The following inclusions and equalities hold among the protocols in asynchronous gossip models:*

(i) $\mathsf{LNS} = \mathsf{KIGr}^\sim = \mathsf{KIGd}^\sim \subset \mathsf{HMS}^\sim \cap \mathsf{TSS}^\sim$
(ii) $\mathsf{HMS}^\sim \subset \mathsf{CO} \cap \mathsf{HSS}^\sim$
(iii) $\mathsf{TSS}^\sim \cup \mathsf{CO}^\sim \cup \mathsf{HSS}^\sim \subset \mathsf{PIG}^\sim$

Proof. The equalities in (i) were proved in Theorem 49. For the inclusion $\mathsf{LNS} \subset \mathsf{HMS}^\sim$ see Corollary 57 and Example 55.4 below. For the inclusion $\mathsf{LNS} \subseteq \mathsf{TSS}^\sim$, on the one hand we have that the LNS condition $\neg\mathsf{S}(x,y)$ is equivalent to both $K_x\neg\mathsf{S}(x,y)$ and $\widehat{K}_x\neg\mathsf{S}(x,y)$; but the latter is equivalent to $\widehat{K}_x(\mathsf{S}(y,y) \wedge \neg\mathsf{S}(x,y))$ (since $\mathsf{S}(y,y)$ is true in any gossip state); finally, if we let $z = y$, the latter is a particular case of $\widehat{K}_x \bigvee_{z \in A}(\mathsf{S}(y,z) \wedge \neg\mathsf{S}(x,z))$. Example 55.2 shows that this inclusion is proper.

For (ii), we show first $\mathsf{HMS}^\sim \subseteq \mathsf{CO}$. Note that after each call in any $\mathsf{HMS}^\sim$ sequence σ, both caller and callee know the other agent has their secret, and so no other call between the two will be made, so $\sigma \in \mathsf{CO}_G$. The inclusion $\mathsf{HMS}^\sim \subseteq \mathsf{HSS}^\sim$ holds by definition of the corresponding protocol conditions. With detail, let $\sigma \in \mathsf{HMS}^\sim_G$ and let (H, τ) be any gossip state. We have $(H, \tau) \models \mathsf{S}(x,x)$, so if $(H, \tau) \models \neg\mathsf{S}(y,x)$ also holds, then $(H, \tau) \models (\mathsf{S}(x,x) \leftrightarrow \neg\mathsf{S}(y,x))$, so $(H, \tau) \models \bigvee_{z \in A}(\mathsf{S}(x,z) \leftrightarrow \neg\mathsf{S}(y,z))$. Given a gossip state (G, σ) with $(G, \sigma) \models \widehat{K}_x\neg\mathsf{S}(y,x)$, we can apply the previous implication to those (H, τ) such that $(G, \sigma) \sim_x (H, \tau)$, and obtain $(G, \sigma) \models \widehat{K}_x \bigvee_{z \in A}(\mathsf{S}(x,z) \leftrightarrow \neg\mathsf{S}(y,z))$. Thus, $\sigma \in \mathsf{HSS}^\sim_G$. Example 55.5 shows these inclusions are proper.

For (iii), the inclusions $\mathsf{TSS}^\sim, \mathsf{HSS}^\sim \subseteq \mathsf{PIG}^\sim$ are obvious from the definitions of the protocol conditions. The inclusion $\mathsf{CO} \subseteq \mathsf{PIG}^\sim$ was proved in Proposition 52. Example 55.1 shows that these inclusions are proper. $\square$

The possible inclusion $\mathsf{CO} \subseteq \mathsf{HSS}^\sim$ remains an open problem. See Fig. 3 and Example 55 for an illustration of Corollary 54, together with examples showing that those inclusions are proper. The two question marks "?" correspond to the open problem whether $\mathsf{CO} \subseteq \mathsf{HSS}^\sim$ holds.

Example 55. The following sequences are examples used in Fig. 3:

1. the sequence $(ab; cd)^{32}; ab; bc; ba$ is in $\mathsf{PIG}_{\mathsf{K}_4}$ but not in $\mathsf{HSS}_{\mathsf{K}_4}$ (given the upper bound $4^3 = 64$ for HSS-sequences) and also not in $\mathsf{TSS}_{\mathsf{K}_4}$ since before the last call b becomes expert; clearly, it is not in $\mathsf{CO}_{\mathsf{K}_4}$ either.
2. $ab; ac; db; cb; ad$ is in $\mathsf{HSS}^{\sim}_{\mathsf{K}_4}, \mathsf{TSS}^{\sim}_{\mathsf{K}_4}, \mathsf{CO}_{\mathsf{K}_4}$ but not in $\mathsf{HMS}^{\sim}_{\mathsf{K}_4}$.
3. $ab; ac; db; cb$ is in $\mathsf{HMS}^{\sim}_{\mathsf{K}_4}, \mathsf{TSS}^{\sim}_{\mathsf{K}_4}$ and $\mathsf{CO}_{\mathsf{K}_4}$.
4. $bc; ab; ac$ is in $\mathsf{HMS}^{\sim}_{G}$ but not in $\mathsf{TSS}^{\sim}_{G}$, where G is the gossip graph: $a \dashrightarrow b \dashrightarrow c$.
5. $\sigma = ab; bc; bd; ce; ad; ae; ac$ is in $\mathsf{CO}_{\mathsf{K}_5}$ but not in $\mathsf{HMS}^{\sim}_{\mathsf{K}_5}$; the reason is that after the calls $\sigma|6 = ab; \ldots; ae$, agent a knows that c must know her own secret, so the last call $\sigma_7 = ac$ cannot be made in $\mathsf{HMS}^{\sim}_{\mathsf{K}_5}$. Note also that $\sigma \in \mathsf{HSS}^{\sim}_{\mathsf{K}_5} \setminus \mathsf{TSS}^{\sim}_{\mathsf{K}_5}$.

The rest of this section is devoted to the proof that $\mathsf{LNS} \subseteq \mathsf{HMS}^{\sim}$. Let us show first that for any $\sigma \in \mathsf{HMS}^{\sim}_{G}$, if agent a does not have the secret of b, then it is also possible (for a) that b does not have the secret of a; in other words

$$(G, \sigma) \models K_a \, \mathsf{S}(b, a) \text{ implies } (G, \sigma) \models \mathsf{S}(a, b).$$

Lemma 56. *Let $G = (A, N, S)$ be an initial gossip graph, and let σ be an $\mathsf{HMS}^{\sim}_{G}$-sequence. If $\neg S^{\sigma} ab$, then there is τ such that $(G, \sigma) \sim_a (G, \tau)$ and $\neg S^{\tau} ba$.*

Proof. Let $|A| = n$, we first define *the tree of calls (ultimately) to a in σ*. For this we define $tree^{\sigma}_k(a)$ inductively as:

$$tree^{\sigma}_0(a) = \{\overline{xa} : \overline{xa} \in \sigma\}$$

$$tree^{\sigma}_{k+1}(a) = \left\{ \overline{xy}; \pi : \begin{array}{l} \exists \pi = \overline{yz}; \ldots; \overline{wa} \in tree^{\sigma}_k(a) \\ \text{for } \sigma = \ldots; \overline{xy}; \ldots; \overline{yz}; \ldots \end{array} \right\}$$

Observe that in $\mathsf{HMS}^{\sim}_{G}$ there cannot be two calls between the same pair of agents. Thus, $n \cdot (n-1)$ is effectively the maximum length of sequences in $\mathsf{HMS}^{\sim}_{G}$ and hence we can define the tree of calls to a in σ as:

$$tree^{\sigma}(a) \quad = \quad \bigcup_{k \leq n \cdot (n-1)} tree^{\sigma}_k(a)$$

We proceed with the proof. Assume $\neg S^{\sigma} ab$. Let τ be the subsequence of σ consisting of all the calls in $tree^{\sigma}(a)$ (i.e. in increasing order w.r.t. that in σ).

Since τ contains no call involving b, we have that $S^{\tau}_b = \{b\}$ and in particular $\neg S^{\tau} ba$. Thus, it suffices to show that $(G, \sigma) \sim_a (G, \tau)$.

Towards a contradiction, let $\tau' = \tau''; cd$ be the smallest prefix such that $(G, \sigma') \not\sim_a (G, \tau')$, where $\sigma' = \sigma''; cd$ is the smallest prefix of σ such that τ' is a subsequence of σ'. Thus, we have $(G, \sigma'') \sim_a (G, \tau'')$ but $(G, \sigma') \not\sim_a (G, \tau')$. By Definition 9, $\overline{cd}$ must be a call of the form $\overline{ad}$ and moreover satisfying either of the following: $S^{\sigma''}_d \neq S^{\tau''}_d$ or $N^{\sigma''}_d \neq N^{\tau''}_d$. (The other two possibilities $S^{\sigma''}_a \neq S^{\tau''}_a$ or $N^{\sigma''}_a \neq N^{\tau''}_a$ are ruled out by the assumption $(G, \sigma'') \sim_a (G, \tau'')$.) In either case there must be a call $\overline{de}$ in σ'' before $\overline{ad}$ which is not in τ''. But this contradicts the definition of τ from the tree of calls $tree^{\sigma}(a)$. $\quad \square$

As a consequence, we have that the $\varphi(x, y)$ condition for LNS implies the condition for $\mathsf{HMS}^{\sim}$. From this we infer the following.

Corollary 57. *The following inclusion holds:* $\mathsf{LNS} \subseteq \mathsf{HMS}^{\sim}$.

Thus, every LNS sequence is in $\mathsf{HMS}^{\sim}$, but not every LNS sequence is in $\mathsf{HMS}^{\approx}$ (the next section contains more results within the synchronous case):
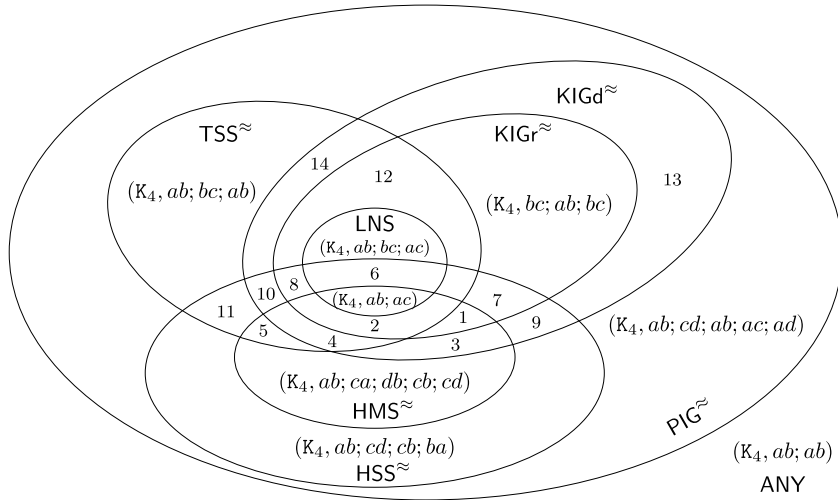
HMS$^\approx$ ⊆ HSS$^\approx$		LNS ⊆ KIGr$^\approx$ ⊆ KIGd$^\approx$			∉ TSS$^\approx$	∈ TSS$^\approx$
1	1	1	1	1	none	$(K_4, ab; ac)$
1	1	0	1	1	$(K_4, bc; ab; ac)$ (1)	$(K_5, ab; ed; bc; ca)$ (2)
1	1	0	0	1	$(G_3, ba; cb; cf; da; ca; cd)$ (3)	$(G_4, ab; bc; da; ca)$ (4)
1	1	0	0	0	$(K_4, ab; ca; db; cb; cd)$	$(K_4, ab; ac; db; cb)$ (5)
0	1	1	1	1	none	$(K_4, ab; cb; ad; ac)$ (6)
0	1	0	1	1	$(K_4, ab; bc; ba)$ (7)	$(K_4, ab; ac; bd; ba; ad)$ (8)
0	1	0	0	1	$(K_5, ab; ae; cd; bc; ad; ab)$ (9)	$(G_{10}, ab; cb; ea; db; ab)$ (10)
0	1	0	0	0	$(K_4, ab; cd; bc; ba)$	$(K_4, ab; bc; ad; ba)$ (11)
0	0	1	1	1	none	$(K_4, ab; bc; ac)$
0	0	0	1	1	$(K_4, bc; ab; bc)$	$(K_3, ab; ac; ba)$ (12)
0	0	0	0	1	$(G_{13}, ab; cb; ab; db; da)$ (13)	$(K_4, ab; ac; ba; bd; ad)$ (14)
0	0	0	0	0	$(K_4, ab; ac; ba; bd; ad)$	$(K_4, ab; bc; ab)$

Fig. 4. Inclusions of protocol extensions in synchronous gossip models.

Example 58. Let $K_{n+1} = (A, N, S)$ with, say, $A = \{a_1, \ldots, a_n, b\}$. Then

$$a_1 a_2; \ a_2 b; \ a_1 a_3; \ \ldots \ ; \ a_1 a_n; \ a_1 b$$

is in LNS, but the last call $a_1 b$ cannot be made in $\mathsf{HMS}_G^{\approx}$, since before this call a_1 knows that the second call could only be of the form $\overline{a_2 b}$.

6.2. Comparison of protocols in synchronous gossip models

Under synchronous calls, we obtain more fine-grained distinctions among protocol extensions than in the asynchronous case (Fig. 3). The inclusions among protocol extensions under synchronous calls are listed in Corollary 59 and depicted in Fig. 4.

Corollary 59. *The following inclusions hold:*

(i) LNS ⊂ KIGr$^\approx$ ⊂ KIGd$^\approx$ (ii) LNS ⊂ CO ∩ TSS$^\approx$

(iii) HMS$^\approx$ ⊂ CO ∩ HSS$^\approx$ (iv) KIGd$^\approx$ ∪ HSS$^\approx$ ∪ TSS$^\approx$ ⊂ PIG$^\approx$

Fig. 5. Graphs for some examples of Fig. 4.

Proof. The proof of the $\subseteq$-inclusions in (i) is analogous to that of Theorem 49 for asynchronous gossip models, just replace $\sim_x$ by $\approx_x$ and $\mathcal{G}^{\approx}$ by $\mathcal{G}^{\sim}$. For (ii), the claim LNS $\subseteq$ CO was shown in [24]; for the inclusion HMS$^{\approx} \subseteq$ CO in (ii), combine the proof in Corollary 54 with the fact HMS$^{\approx} \subseteq$ HMS$^{\sim}$. For (iii)–(iv), the proofs are the same than in Corollary 54 for the asynchronous case, again replacing $\sim_x$ by $\approx_x$. Finally, see Fig. 4 for examples showing that these inclusions are proper. □

In the table of Fig. 4, the two rightmost columns contain respectively sequences not in TSS$_{K_4}^{\approx}$ (left) and in TSS$_{K_4}^{\approx}$ (right). These sequences are (resp. are not) in the extension of each of the protocols listed in the left if they are labeled with 1 (resp. 0). Combinations labeled by 'none' are not possible due to LNS $\subseteq$ TSS$^{\approx}$.

Graphs G_3, G_4, G_{10}, G_{13} are given in Fig. 5. Examples 3, 4, 10 and 13 (in grey) describe sequences that are actually not in KIGd$_G^{\approx} \smallsetminus$ KIGr$_G^{\approx}$ because we assume for these four examples that all agents have knowledge of the initial topology N. It is unknown whether examples for the four cells in grey exist without this assumption.

7. Related work

The original gossip problem of spreading secrets on permanent networks has been extensively studied, first in the case of complete graphs, i.e. when any agent can call any other agent (Tijdeman [22], or Hurkens [17] for a modern reference) and later for other network topologies, e.g. Golumbic [10] and Harary et al. [13]. Hedetniemi et al. [14] is a survey of results in many different variants of the gossip problem. The focus is on the minimum number of calls leading to all agents knowing all secrets; and the different proofs describe centralized algorithms for call scheduling among a given network of agents. These papers, then, do not aim at the study of distributed protocols.

Distributed epistemic protocols for gossip have been introduced recently in Attamah et al. [2] and Apt et al. [1], where each agent decides to call depending on its knowledge. [2] studied LNS, PIG, KIGd and KIGr in complete graphs, whereas [1] studied LNS and HMS under the *merge-then-learn* call mode (again in complete graphs) and protocols designed for ring networks. As we said near the end of Section 2.1, the accessibility relations under the *merge-then-learn* model of calls are less fine-grained than under *learn-then-merge* (and presumably less natural as models of communication for standard gossip applications among humans and also computers). More formally,

Fact 60. *For any initial gossip state $G = (A, N, S)$ and any agent $x \in A$,*

$$\sim_x \subseteq \sim_x^{\texttt{ML}} \quad and \quad \approx_x \subseteq \approx_x^{\texttt{ML}}$$

where $\sim_x^{\texttt{ML}}$ and $\approx_x^{\texttt{ML}}$ denote the accessibility relations under merge-then-learn.

The subset part ($\subseteq$) of these two claims can be easily proved by translating these telephone calls into secret announcements in the framework of Dynamic Epistemic Logic [23]: for *learn-then-merge*, announce what secrets are known by each agent, e.g. $S(a,c) \wedge S(b,d)$; for *merge-then-learn*, announce what either of them know, e.g. $(S(a,c) \vee S(b,c)) \wedge (S(a,d) \vee S(b,d))$. Clearly, the former are at least as informative as the latter. Moreover, the above inclusions $\sim_x \subseteq \sim_x^{ML}$ and $\approx_x \subseteq \approx_x^{ML}$ can be proper, as shown in Example 61.

Example 61 *(Merge-then-learn).* Under asynchronous *merge-then-learn*, if a knows c's secret before a call ab (center gossip states), after the call ab she cannot tell if b also knew it before (bottom), or just learnt it from her (top). That is, $(G, ac; ab) \sim_a^{ML} (G, ac; bc; ab)$, while $(G, ac; ab) \not\sim_a (G, ac; bc; ab)$.

For the synchronous case, consider a similar gossip graph G' with 5 agents; then,

$$(G', ac; cb; ab) \approx_a^{ML} (G', ac; de; ab) \quad \text{while} \quad (G', ac; cb; ab) \not\approx_a (G', ac; de; ab).$$

The related problem of epidemic information (e.g. for distributed database maintenance among processes) has also been studied in dynamic topologies, either randomly generated at each round of calls in Karp et al. [19], or when processes (with bounded contact lists) exchange contacts in messages — see Eugster et al. [6] for a survey. Haeupler et al. [11] propose an efficient algorithm for the gossiping of agents' numbers only (i.e. addresses), also called the *resource discovery problem*. Other distributed algorithms for dynamic gossip include LNS, CO and other token-passing protocols, which have been studied in previous work [24].

The current paper is a follow-up of [24] with the goal of understanding distributed knowledge-based protocols. Thus, the conditions for a call xy in the protocols studied here are not just those definable only from the facts available to the caller agent x (i.e. the secrets and numbers learnt, the list of previous calls made or received by x). In epistemic protocols, the condition for call xy can also depend as well on the reasoning with higher-order information, e.g. the secrets that other agents knew in the past. From this, agent x can figure out which global call histories might or might not have occurred, and which secrets an agent might know.

Indeed, the language used here for protocol conditions (Definition 14) is the standard language of epistemic logic of Hinktikka [16] with multiple agents and where atomic propositions are grounded. Our language does not contain any dynamic operators, as in Baltag et al. [4] or temporal operators, in the style of Halpern et al. [12] or Parikh et al. [21]. The approach of [21] employs memory-based agents. Our agents have memory, but only indirectly so (and this indirect usage is common in the area): this is the role of the $C(ab, c)$ variables in our purely epistemic language, that become (and then remain) true once agent c has been involved (as a or as b) in a call ab. Attamah et al. [2] propose another similar language but with action models for calls and with protocols explicitly written in PDL dialect [8] with knowledge operators. In Apt et al. [1], they consider epistemic guards for knowledge-based programs like in our setting.

Our gossip models could be reformulated as specific interpreted systems, in the style of Halpern et al. [12] or Fagin et al. [7]. In [2], possible worlds are distributions of secrets (gossip graphs) whereas in our settings they encode both the initial gossip graph and the history of calls executed so far. In [1], possible worlds only contain the call sequences because the topology (complete graph or ring) is common knowledge.

Our protocols can be seen as specific knowledge-based programs (see chap. 7 in [7]): the conditions for calls by some agent x are of the form $\neg K_x \psi$ or $K_x \psi$, as explained at the end of Section 3.

Finally, Herzig et al. [15] proposes an epistemic logic where calls are modeled by assignments of visibility atoms. This logic is then used to model knowledge after telephone calls in classical gossip.

8. Conclusions and future work

In this paper, we studied some distributed protocols for dynamic gossip that are defined by epistemic conditions, following the work started in [24]. As in classical gossip, the goal is that all agents learn all secrets using telephone calls. But in dynamic gossip both secrets and telephone numbers of agents are communicated. As a consequence, any given protocol P requires less or equal calls than in the classical problem. We focused on two modes of calls, leading to the corresponding asynchronous and synchronous notions of knowledge (an agent obtains knowledge from the contents of each call with other agents). As usual in modal logic, each notion of knowledge is represented by an accessibility relation between possible states, here described by a pair consisting of an initial network and a call history.

The epistemic protocols studied in this paper are based on the idea that each call (possibly or necessarily) results in the learning of some secret by either agent involved in this call. After presenting each protocol, we proved the corresponding characterization results, i.e. that some class of graphs captures all (or contains some of) the gossip scenarios where this protocol is necessarily (or possibly) successful. The characterization results are summarized in the following table, where "..." denotes that the class of gossip graphs is the same as in the cell on the left.

	Strong success	Fair success	Weak success
ANY [24]	no	weakly connected	...
LNS [24]	sun graphs	...	not (bush or double bush)
CO [24]	weakly connected	...	...
PIG	no	weakly connected (*Theorem 29*)	...
TSS	no	weakly connected (*Theorem 30, Corollary 31*)	...
HSS	weakly connected (*Theorem 33, Corollary 34*)	...	...
HMS	at least for sun* graphs (*Theorem 39*)	...	weakly connected (*Theorem 41, Corollary 42*)
KIGd$^\approx$ KIGr$^\approx$	at least for sun graphs (*Theorem 45*)	...	weakly connected (*Theorem 46*)
KIGd$^\sim$ KIGr$^\sim$	sun graphs (*Corollary 50*)	...	not (bush or double bush) (*Corollary 50*)

Finally, we studied the protocols from a comparative point of view to find out whether a protocol is subsumed or equal to some other protocol, in terms of their sets of executions.

Still, some open problems remain: a full characterization of the protocols HMS, KIGd$^\approx$ and KIGr$^\approx$; a proof or refutation of the inclusion CO $\subseteq$ HSS$^\sim$; and a few examples of couples (G, σ) in KIGd$^\sim$ \ KIGr$^\sim$ that do not assume knowledge of the network topology. Solving these problems is left for future work.

Other interesting directions to be explored are: protocol execution length and other knowledge assumptions.

For execution length, it would be useful to characterize the class of graphs where each protocol P can be optimal, i.e. the class of gossip graphs where a successful P-execution exists using the minimum number

of calls $2n - 4$, for gossip graphs with $n \geq 4$ agents. A comparative study in terms of maximum execution length would permit to select a protocol minimizing the maximum number of calls in a given network.

With respect to the assumptions on agents' initial knowledge, one might consider strengthening the assumptions made in this paper. Presumably, this should result in the protocols being successful in wider classes of graphs. The main directions towards this might consist of assuming that agents know who are the contacts of their contacts; or, moreover, to assume common knowledge of the initial network, or common knowledge that the agents are executing the same protocol. All these open problems and directions must be left for future work.

Another open problem concerns the following model checking problem: given a gossip state (G, σ) and a formula $\varphi(a, b)$, determine whether $(G, \sigma) \models \varphi(a, b)$. Certainly it is decidable in the synchronous case. Indeed, the exploration of the model $\mathcal{G}^{\approx}$ from a given (G, σ) is confined to the set of gossip states (H, σ') where $|\sigma'| = |\sigma|$, which is finite. In contrast, the decidability of model checking in the asynchronous case is not obvious, since a model checker may explore an infinite number of gossip states. We conjecture its decidability, though, and that this can be proved by applying techniques developed in [3].

Acknowledgements

References

[1] K.R. Apt, D. Grossi, W. van der Hoek, Epistemic protocols for distributed gossiping, in: Proceedings Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2015, June 4–6, 2015, Carnegie Mellon University, Pittsburgh, USA, 2015, pp. 51–66.

[2] M. Attamah, H. van Ditmarsch, D. Grossi, W. van der Hoek, Knowledge and gossip, in: ECAI 2014 – 21st European Conference on Artificial Intelligence, 18–22 August 2014, Prague, Czech Republic – Including Prestigious Applications of Intelligent Systems (PAIS 2014), 2014, pp. 21–26.

[3] G. Aucher, B. Maubert, S. Pinchinat, Automata techniques for epistemic protocol synthesis, arXiv preprint arXiv:1404.0844.

[4] A. Baltag, L.S. Moss, S. Solecki, The logic of public announcements and common knowledge and private suspicions, in: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge, TARK-98, 1998, pp. 43–56.

[5] S. Dasgupta, C.H. Papadimitriou, U. Vazirani, Algorithms, McGraw–Hill, Inc., 2006.

[6] P.T. Eugster, R. Guerraoui, A. Kermarrec, L. Massoulié, Epidemic information dissemination in distributed systems, IEEE Comput. 37 (5) (2004) 60–67.

[7] R. Fagin, Y. Moses, J.Y. Halpern, M.Y. Vardi, Reasoning About Knowledge, The MIT Press, 2003.

[8] M.J. Fischer, R.E. Ladner, Propositional dynamic logic of regular programs, J. Comput. Syst. Sci. 18 (2) (1979) 194–211.

[9] M. Fitting, R.L. Mendelsohn, First-Order Modal Logic, vol. 277, Kluwer Academic Publishers, 1988.

[10] M.C. Golumbic, The general gossip problem, Technical Report, IBM Research Report RC4977, IBM.

[11] B. Haeupler, D. Malkhi, Distributed resource discovery in sub-logarithmic time, in: Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing, PODC 2015, 2015, pp. 413–419.

[12] J.Y. Halpern, Y. Moses, Knowledge and common knowledge in a distributed environment, J. ACM 37 (3) (1990) 549–587.

[13] F. Harary, A.J. Schwenk, The communication problem on graphs and digraphs, J. Franklin Inst. 297 (1974) 491–495.

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

[15] A. Herzig, E. Lorini, F. Maffre, A poor man's epistemic logic based on propositional assignment and higher-order observation, in: Proceedings of Logic, Rationality, and Interaction – 5th International Workshop, LORI 2015, 2015, pp. 156–168.

[16] J. Hintikka, Knowledge and Belief: An Introduction to the Logic of the Two Notions, Cornell University Press, Ithaca NY, 1962.

[17] C.A.J. Hurkens, Spreading gossip efficiently, Nieuw Arch. Wiskd. 5 (1) (2000) 208–210.

[18] W. Jamroga, W. van der Hoek, Agents that know how to play, Fundam. Inform. 63 (2004) 185–219.

[19] R.M. Karp, C. Schindelhauer, S. Shenker, B. Vöcking, Randomized rumor spreading, in: Proceedings of 41st Foundations of Computer Science, FOCS 2000, 2000, pp. 565–574.

[20] R. Montague, The proper treatment of quantification in ordinary English, in: P. Suppes, J. Moravcsik, J. Hintikka (Eds.), Approaches to Natural Language, Springer, Dordrecht, 1973, pp. 221–242.

[21] R. Parikh, R. Ramanujam, A knowledge based semantics of messages, J. Log. Lang. Inf. 12 (4) (2003) 453–467.
[22] R. Tijdeman, On a telephone problem, Nieuw Arch. Wiskd. 3 (19) (1971) 188–192.
[23] H. van Ditmarsch, W. van der Hoek, B. Kooi, Dynamic Epistemic Logic, Synthese Library Series, vol. 337, Springer, 2007.
[24] H. van Ditmarsch, J. van Eijck, P. Pardo, R. Ramezanian, F. Schwarzentruber, Dynamic gossip, http://arxiv.org/abs/1511.00867, 2015.