Reasoning about Gossip

Hans van Ditmarsch
CNRS

- Epistemic Goals

Materials found on https://reasoningaboutgossip.eu

What is epistemic?

Gossip Protocol: *Until all agents know all secrets (termination condition), select a and b satisfying the condition to make a call (call condition), and let a call b (message).*

A gossip protocol can be epistemic in different ways.

- The call condition (precondition) is epistemic.
  [Attamah, vD, vdHoek 2014] [Apt, Wojtczak 2018] [Apt, Grossi, vhHoek 2015] [vD, vEijck, Pardo, Ramezanian, Schw. 2017]
- The termination condition (goal) is epistemic.
  [vD et al. *Everyone knows that everyone knows.* Stud.Log. 2023]
  [vD, Gattinger. *You can only be lucky once.* MSCS 2024]
- The message (information exchanged) is epistemic.
  [Herzig, Maffre. *How to share knowledge by gossiping.* 2017]
  [Cooper et al. *The epistemic gossip problem.* 2019]
  [*various chapters of Reasoning about Gossip*]

We now focus on epistemic termination conditions. Complete graphs.

The termination goal is epistemic

The usual goal is that everyone knows all secrets (all are experts). Consider the goal that everyone knows that everyone knows all secrets. An agent who knows that all agents are experts is a super expert. The new goal is that all are super experts. A call sequence satisfying that is super-successful. *Example for 4 agents:*

$ab.cd.ac.bd.$	all agents know all secrets
$ab.ad.$	agent a knows that all agents know all secrets
$bc.$	agent b knows that all agents know all secrets
cd	agents c, d know that all agents know all secrets

If only secrets are exchanged, super-successful is the limit, and $n - 2 + \binom{n}{2}$ calls are optimal (for $n \geq 4$).

[vD, Gatt., Ramezanian. *Everyone knows that everyone knows*. 2023]
[vD, Gattinger. *The Limits to Gossip*, ... 2022]

If not, we can reach arbitrary higher-order epistemic goals, and faster:

Epistemic messages (and epistemic goal)

If agents can only communicate secrets, we got: $\mathcal{O}(n^2)$

$ab.cd.ac.bd.$	all agents know all secrets
$ab.ad.$	agent a knows that all agents know all secrets
$bc.$	agent b knows that all agents know all secrets
cd	agents c, d know that all agents know all secrets

If agents may communicate knowledge about secrets, we get: $\mathcal{O}(n)$ opt.

$ab.cd.ac.bd.$	all agents know all secrets
$ab.$	agent a informs b that a, c know all secrets
	agent b informs a that b, d know all secrets
	agents a, b know that all agents know all secrets
cd	agent c informs d that a, c know all secrets
	agent d informs c that b, d know all secrets
	agents c, d know that all agents know all secrets

[Herzig, Maffre. *How to share knowledge by gossiping.* AIComm 2017]
[Cooper *et al.* *The epistemic gossip problem.* Discrete Math. 2019]
[Chapter 10 *Epistemic Goals* of *Reasoning about Gossip*]

Everyone knows that everyone knows — different format

Reconsider call sequence $ab.cd.ac.bd.ab.ad.bc.cd$.

$ab.cd.ac.bd.$	all agents know all secrets
$ab.ad.$	agent a knows that all agents know all secrets
$bc.$	agent b knows that all agents know all secrets
cd	agents c, d know that all agents know all secrets

Same, with enhanced notation for ordered tuples:
— lower case y for agent x means that x knows the secret of y;
— upper case Y for agent x means that x knows that y is expert.

$$
\begin{aligned}
&a|b|c|d \overset{ab}{\to} ab|ab|c|d \overset{cd}{\to} ab|ab|cd|cd \overset{ac}{\to} AbCd|ab|AbCd|cd \\
&\overset{bd}{\to} AbCd|aBcD|AbCd|aBcD \overset{ab}{\to} ABCd|ABcD|AbCd|aBcD \\
&\overset{ad}{\to} ABCD|ABcD|AbCd|ABcD \overset{bc}{\to} ABCD|ABCD|ABCd|ABcD \\
&\overset{cd}{\to} ABCD|ABCD|ABCD|ABCD
\end{aligned}
$$

Notation assumes asynchrony. Synchronously an agent may know another agent is an expert without knowing that agent's secret.

The Limits to Gossip

- An expert is an agent who knows all secrets.
  A call sequence is successful if all agents are experts.
  (all agents know all secrets)
- A super expert knows that all agents are experts.
  A sequence is super-successful if all agents are super experts.
  (all agents know that all agents know all secrets)

Super-successful is the limit. We cannot have that:

> *all agents know that all agents are super experts.*
> *(all agents know that all agents know that all agents know all secrets.)*

This assumes that there is no global clock, no common knowledge of the protocol, and that all calls are permitted. Otherwise there is no limit, even when only secrets are exchanged.

Language and semantics for secret distributions

Epistemic language (a_b means that b knows the secret of a)

$$\varphi \ ::= \ a_b \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_a\varphi$$

Epistemic (observation) relation

- $\epsilon \sim_a \epsilon$
- if $\sigma \sim_a \tau$ and $a \notin \{b, c\}$, then $\sigma.bc \sim_a \tau$ (asynchronous!)
- if $\sigma \sim_a \tau$, and for all c, $\sigma \models c_b$ iff $\tau \models c_b$, then $\sigma.ab \sim_a \tau.ab$

Semantics *where $a \neq b$ in the clauses $\sigma.ab$ below*

$$\begin{array}{lll}
\epsilon \models a_b & \text{iff} & a = b \\
\sigma.ab \models c_a & \text{iff} & \sigma \models c_a \text{ or } \sigma \models c_b \qquad \text{for all } c \in A \\
\sigma.ab \models c_b & \text{iff} & \sigma \models c_a \text{ or } \sigma \models c_b \qquad \text{for all } c \in A \\
\sigma.ab \models c_d & \text{iff} & \sigma \models c_d \qquad \text{for all } c, d \in A \text{ with } d \notin \{a, b\} \\
\sigma \models \neg\varphi & \text{iff} & \sigma \not\models \varphi \\
\sigma \models \varphi \wedge \psi & \text{iff} & \sigma \models \varphi \text{ and } \sigma \models \psi \\
\sigma \models K_a\varphi & \text{iff} & \tau \models \varphi \text{ for all } \tau \text{ such that } \sigma \sim_a \tau
\end{array}$$

Second-order mutual knowledge is unsatisfiable

$Exp_B(C)$	$:=$	$\bigwedge_{b \in B}^{c \in C} c_b$	everyone in B knows all secrets in C
Exp_A	$:=$	$Exp_A(A)$	everyone knows all secrets
$E\varphi$	$:=$	$\bigwedge_{a \in A} K_a \varphi$	everyone knows φ (mutual knowledge)
$K_a Exp_A$			agent a is a super expert
$EExp_A$			everyone is a super expert

Gossip protocol termination goals. What is satisfiable and what not:

- success: Exp_A. Yes. (zeroth-order mutual knowledge)
- super success: $EExp_A$. Yes. (first-order mutual knowledge)
- $>$ super success: $EEExp_A$? No! (second-order mutual knowl.)

The notion of lucky call is important in obtaining this result and for optimality. Even when agents are only aware of calls they are involved in, and when there is no global clock, they can learn in a call that an agent not involved in the call is an expert: a lucky call.

Lucky calls

Given agents $\{a, b, c, d\}$ consider call sequence $ac.ad.ac.bc.ac$. Agent a learns in the final lucky call ac that a, b, c are experts.

$$a|b|c|d \stackrel{ac}{\to} ac|b|ac|d \stackrel{ad}{\to} acd|b|ac|acd \stackrel{ac}{\to} acd|b|acd|acd$$
$$\stackrel{bc}{\to} acd|aBCd|aBCd|acd \stackrel{ac}{\to} ABCd|aBCd|ABCd|acd$$

After $ac.ad.ac$, a knows that a, c, d know the secrets of a, c, d. Agent a is a super expert of the subset $\{a, c, d\}$ of all. This allows agent a to learn indirectly that b is an expert in her final call.

Experts can also be lucky:

$$a|b|c|d \stackrel{ab}{\to} ab|ab|c|d \stackrel{cd}{\to} ab|ab|cd|cd \stackrel{bc}{\to} ab|aBCd|aBCd|cd$$
$$\stackrel{ab}{\to} ABcd|ABCd|aBCd|cd \stackrel{ad}{\to} ABCD|ABCd|aBCd|AbcD$$

In the second magic call ab, agent a learns that c or d is an expert. Call ad resolves this uncertainty (but not if the call had been ac).

Second-order mutual knowledge is unsatisfiable

Definition: a is lucky about c if $\sigma \not\models K_a Exp_c$ and $\sigma.ab \models K_a Exp_c$.

- When two agents become experts they cannot be lucky.
- When two agents become experts they do not become super experts.
- An agent becoming an expert can only be lucky if she is a super expert for all agents but one:
  If $\sigma \not\models Exp_a$, $\sigma \not\models \neg K_a Exp_b$, and $\sigma \models Exp_c$, then
  $\sigma.ac \models K_a Exp_b$ iff $\sigma \models K_a Exp_{A-b}(A-b)$.
- An agent cannot become an expert and a super expert in the same call. (This is proved with the previous item.)
- An agent becoming a super expert considers possible that the other agent involved in the call did not become a super expert:
  $\sigma.ab \models \hat{K}_a \neg K_b Exp_A$.

We continue the proof on the next slide.

Second-order mutual knowledge is unsatisfiable — cont.

- An agent becoming a super expert considers possible that the other agent involved in the call did not become a super expert: $\sigma.ab \models \hat{K}_a \neg K_b Exp_A$.
- This property persists after any extension of the call sequence: $\sigma.ab.\tau \models \hat{K}_a \neg K_b Exp_A$.
- In other words: $\sigma.ab.\tau \not\models K_a K_b Exp_A$.
- In fact, for any ρ there are a, b such that: $\rho \not\models K_a K_b Exp_A$.
- There are a, b such that $K_a K_b Exp_A$ is unsatisfiable.
- Theorem: $EEExp_A$ is unsatisfiable. ($\neg EEExp_A$ is valid)

[vD & Gattinger. The Limits To Gossip. WoLLIC 2022]
[vD & Gattinger. You can only be lucky once. MSCS 2024]

Everyone knows that everyone knows: optimality

Recall that $2n - 4$ is optimal for all to become experts.
We show $n - 2 + \binom{n}{2}$ is optimal for all to become super experts.

We recall the example for $n = 4$: *ab.cd.ac.bd.ab.ad.bc.cd*.
After two calls noone is an expert yet: *ab.cd*.
After the remaining six calls callers are experts: *ac.bd.ab.ad.bc.cd*.

There are even simpler schedules. First, one agent a calls all other agents $(n - 1)$, then again except the last one $(n - 2)$, and finally the remaining calls between experts: *ab.ac.ad.ab.ac.bc.bd.cd*.

So if agents could not be lucky the proof would have been simple:
To create an expert you need $n - 1$ calls. So $n - 2$ is just one short.
To determine that another agent is expert, call the agent: $\binom{n}{2}$ calls.

A lucky call appears to permit shorter call sequences.
The proof consists of showing that it does not. We show:
If a is lucky about b, then b has to call a to learn a is an expert.

Everyone knows that everyone knows: optimality

Assumptions				Concl.: $\sigma.ac \models K_a Exp_b$ iff . . .
$\sigma \not\models K_a Exp_b$	$\sigma \not\models Exp_a$	$\sigma \not\models Exp_c$		false
		$\sigma \models Exp_c$		$\sigma \models K_a Exp_{A-b}(A-b)$
	$\sigma \models Exp_a$	$\sigma \not\models Exp_c$	$\neg K_a$magic	false
			$K_a$magic	$(*)$
		$\sigma \models Exp_c$		false

We characterize lucky calls by mutually exclusive cases, given all different $a, b, c \in A$, and call sequence σ. Property $K_a$magic is the formula $K_a \bigvee_{d \in A}(Exp_d \wedge \neg K_a Exp_d)$. It is established in the magic call, wherein a becomes expert by calling an agent who already is an expert. Condition $(*)$ is: $\sigma \models Exp_b \wedge K_a(Exp_b \vee Exp_c)$ and b, c are distant. An agent is distant if it is not close. Set D of agents is close if there is an agent e such that $\tau \models K_a Exp_D(A-e)$, where τ is the prefix of σ before the magic call. (Here, e must be b or c.)

Proof analyzes call sequence following the magic call. Luck requires magic. Magic is preserved except when calling a distant agent.

Reaching epistemic goals with known protocols

Consider a logical language consisting of formulas and programs.

- Formula $K_a^{\mathsf{P}}\varphi$ stands for "agent a knows φ given protocol P," where "given protocol P" means that the agents have common knowledge that they all execute protocol P.
- Protocol P is a program of shape "until all agents are super experts, select agents a, b such that protocol condition P_{ab} is satisfied, and execute call ab," where P_{ab} is a formula.

The formulas and the programs should therefore be defined by simultaneous recursion. This is well-defined. Formula $K_a^{\mathsf{P}}\varphi$ can be seen as an inductive construct with $\binom{n}{2}+1$ arguments, namely φ and all $\binom{n}{2}$ protocol conditions P_{bc} (for $b \neq c$) for the protocol P.

Dually, $K_a^{\mathsf{P}}\varphi$ is true after call sequence σ ($\sigma \models K_a^{\mathsf{P}}\varphi$) iff φ is true after all indistinguishable P-permitted call sequences τ ($\sigma \sim_a^{\mathsf{P}} \tau$), where τ is P-permitted iff for all bc occurring in τ, P_{bc} was true prior to the execution of call bc.

[vD, Gattinger, Kuijer, Pardo. *Strengthening Gossip Protocols*, 2019.]

Protocol knowledge — CMO (Call Me Once)

In CMO the maximum number of calls between n agents is $\binom{n}{2}$. All maximal CMO-permitted sequences are successful. Given agents a, b, c, d, a maximal CMO-permitted sequence is

$$\sigma := ab.bc.cd.ad.bd.ac$$

If CMO is common knowledge and synchronous, all agents are now super experts. (It is even common knowledge that all are experts.) Otherwise, not. E.g., σ is indistinguishable for agent a from

$$\tau := ab.bc.cd.ad.cd.ac$$

after which agent b does not know the secret of d and is not an expert. Call sequence τ is not CMO-permitted. But agent a does not know that agents c and d only make CMO-permitted calls. Let time not be known, but CMO known. Then $\sigma \sim_a \tau'$ for

$$\tau' := ab.bc.cd.ad.ac$$

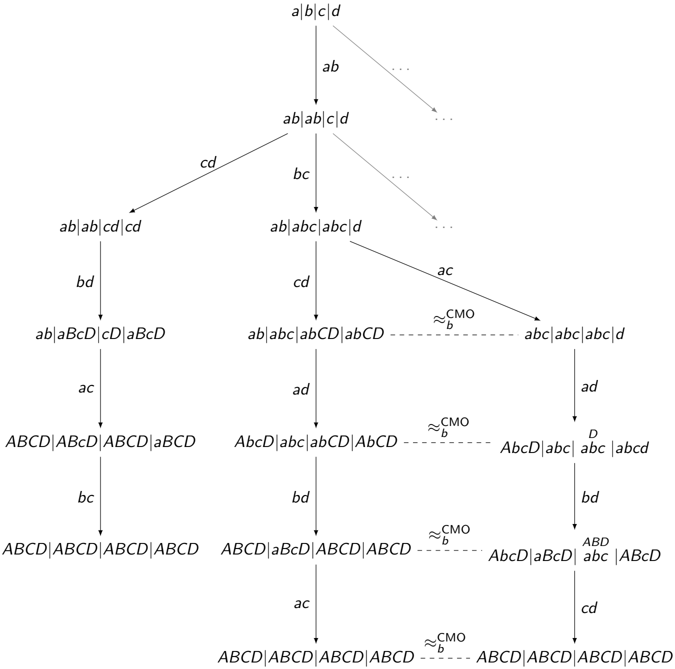

$a|b|c|d$

ab

$\dots$

$ab|ab|c|d$

cd

bc

$\dots$

$ab|ab|cd|cd$

$ab|abc|abc|d$

$\dots$

bd

cd

ac

$ab|aBcD|cD|aBcD$

$ab|abc|abCD|abCD \approx_b^{\mathsf{CMO}} abc|abc|abc|d$

ac

ad

ad

$ABCD|ABcD|ABCD|aBCD$

$AbcD|abc|abCD|AbCD \approx_b^{\mathsf{CMO}} AbcD|abc|\overset{D}{abc}|abcd$

bc

bd

bd

$ABCD|ABCD|ABCD|ABCD$

$ABCD|aBcD|ABCD|ABCD \approx_b^{\mathsf{CMO}} AbcD|aBcD|\overset{ABD}{abc}|ABcD$

ac

cd

$ABCD|ABCD|ABCD|ABCD \approx_b^{\mathsf{CMO}} ABCD|ABCD|ABCD|ABCD$

Results for super-successful gossip protocols

- ANY is super-successful (i.e., all fair executions are s-s.)
- PIG is super-successful
- synchronous known CMO is super-successful
- optimal synchronous ANY is less than optimal asynchr. ANY. *ab.ac.ab.cb* is asynchr. s-s, but prefix *ab.ac.ab* is synchr. s-s.
- many of these results require the model checker GoMoChe `https://github.com/m4lvin/gossip`

Other results are for semantics with engaged agents that do not make and do not answer calls once they are experts. When you call someone who does not answer, this now signals she is an expert. [vD et al. *Everyone knows that everyone knows.* Stud.Log. 2023]

Stronger messages (higher-order knowledge of secrets, full info.) permit reaching arbitrary higher-order epistemic goals $E^n Exp_A$ for any $n \in \mathbb{N}$. $CExp_A$ is unreachable without protocol knowledge. [*Reasoning about Gossip*, Chapter 10 *Epistemic Goals*, last section]