Reasoning about Gossip

Hans van Ditmarsch
CNRS

- Call, call sequence, semantics of calls

Materials found on http://reasoningaboutgossip.eu

Gossip terminology — call

A is the set of agents or callers.

A gossip graph G is a triple (A, N, S) where $N \subseteq A \times A$ is the neighbour relation and $S \subseteq A \times A$ is the secret relation.

If $N = A \times A$ (all agents can call each other), the gossip graph is *complete*. We then call the gossip graph a secret distribution denoted S. The initial secret distribution I is the triple (A, A^2, I).

A call or telephone call is a pair from $A \times A$. For $(a, b) \in A \times A$ we write ab, and we require $a \neq b$. We say that a and b are *involved* in call ab, that a is the *caller*, and b the *callee*.

Write S_a for $\{b \in A \mid (a, b) \in S\}$. If $S_a = A$, agent a is an expert.

Simplified notation for secret distributions: $a|b|c|d$, $abc|ab|abc|d$, $\ldots$

Gossip terminology — semantics of a call

Given gossip graph (A, N, S).

- **pushpull:** The result of applying call ab is the gossip graph (A, N, S^{ab}), where $S^{ab} = S \cup (\{(a,b),(b,a)\} \circ S)$.
  a and b learn each other's secrets

Alternatively, $S^{ab}_a = S^{ab}_b = S_a \cup S_b$ and $S^{ab}_c = S_c$ for $c \neq a, b$.

Variants (not varying the notation)

- **push:** The result of making call ab is the gossip graph (A, N, S^{ab}), where $S^{ab} = S \cup (\{(b,a)\} \circ S)$.
  b learns the secrets of a
- **pull:** The result of making call ab is the gossip graph (A, N, S^{ab}), where $S^{ab} = S \cup (\{(a,b)\} \circ S)$.
  a learns the secrets of b
- **dynamic pushpull:** The result of call ab is the gossip graph (A, N^{ab}, S^{ab}), where $N^{ab} = N \cup (\{(a,b),(b,a)\} \circ N)$ and $S^{ab} = S \cup (\{(a,b),(b,a)\} \circ S)$.
  a and b learn each other's secrets and neighbours/numbers

Gossip terminology — call sequence

A call sequence is inductively defined as: ϵ is a call sequence, if σ is a call sequence and ab is a call, then $\sigma.ab$ is a call sequence.

We write (all with obvious inductive definitions) :

- $|\sigma|$ to denote the length of a call sequence
- $\sigma[i]$ for the ith call of the sequence
- $\sigma|i$ for the first i calls of the sequence

Applying σ to a secret relation S: $S^{\epsilon} = S$; and $S^{\sigma.ab} = (S^{\sigma})^{ab}$. Same for N. By G^{σ}, where $G = (A, N, S)$, we mean (A, N, S^{σ}). Given secret distribution $I^{\sigma} = (A, A^2, I^{\sigma})$ we write σ_a for I^{σ}_a.

Executing a call sequence in the initial secret distribution $a|b|c|d$:

$$a|b|c|d \stackrel{ab}{\to} ab|ab|c|d \stackrel{cd}{\to} ab|ab|cd|cd \stackrel{ac}{\to}$$
$$abcd|ab|abcd|cd \stackrel{bd}{\to} abcd|abcd|abcd|abcd$$

Gossip terminology — full information

local view $\mathsf{v}_a^{=}(\sigma)$ for agent a of call sequence σ:

$$\begin{array}{lcl} \mathsf{v}_a^{=}(\epsilon) & := & \epsilon \\ \mathsf{v}_a^{=}(\sigma.bc) & := & \mathsf{v}_a^{=}(\sigma) \\ \mathsf{v}_a^{=}(\sigma.ab) & := & \mathsf{v}_a^{=}(\sigma).ab \\ \mathsf{v}_a^{=}(\sigma.ba) & := & \mathsf{v}_b^{=}(\sigma).ba \end{array}$$

full view $\mathsf{v}_a^{\sim}(\sigma)$ for agent a of call sequence σ: a dag!

$$\begin{array}{lcl} \mathsf{v}_a^{\sim}(\epsilon) & := & \epsilon \\ \mathsf{v}_a^{\sim}(\sigma.bc) & := & \mathsf{v}_a^{\sim}(\sigma) \\ \mathsf{v}_a^{\sim}(\sigma.ab) & := & (\mathsf{v}_a^{\sim}(\sigma), \mathsf{v}_b^{\sim}(\sigma)).ab \\ \mathsf{v}_a^{\sim}(\sigma.ba) & := & (\mathsf{v}_b^{\sim}(\sigma), \mathsf{v}_a^{\sim}(\sigma)).ba \end{array}$$

synchronous full view $\mathsf{v}_a^{\approx}(\sigma)$ for agent a of call sequence σ:

$$\begin{array}{lcl} \mathsf{v}_a^{\approx}(\epsilon) & := & \epsilon \\ \mathsf{v}_a^{\approx}(\sigma.bc) & := & \mathsf{v}_a^{\approx}(\sigma).\bullet \\ \mathsf{v}_a^{\approx}(\sigma.ab) & := & (\mathsf{v}_a^{\approx}(\sigma), \mathsf{v}_b^{\approx}(\sigma)).ab \\ \mathsf{v}_a^{\approx}(\sigma.ba) & := & (\mathsf{v}_b^{\approx}(\sigma), \mathsf{v}_a^{\approx}(\sigma)).ba \end{array}$$

Gossip terminology — example of full information

Let call sequence $\sigma = ab.bc.ab$ be given.

- $\mathsf{v}_a^{=}(\sigma) = ab.ab$, $\mathsf{v}_b^{=}(\sigma) = ab.bc.ab$ and $\mathsf{v}_c^{=}(\sigma) = bc$
- $\mathsf{v}_a^{\sim}(\sigma) = \mathsf{v}_b^{\sim}(\sigma) = (ab, ab.bc).ab$, $\mathsf{v}_c^{\sim}(\sigma) = ab.bc$
- $\mathsf{v}_a^{\approx}(\sigma) = \mathsf{v}_b^{\approx}(\sigma) = (ab.\bullet, (ab, \bullet).bc).ab$, $\mathsf{v}_c^{\approx}(\sigma) = (ab, \bullet).bc.\bullet$

Pairing the empty call sequence ϵ with a sequence σ delivers σ!
A picture says more than a thousand symbols ...

(asynchronous) full view

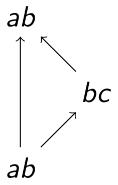

synchronous full view

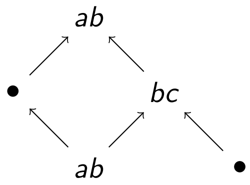

Gossip terminology — observation relation

Given $a \in A$ and gossip graphs $G = (A, N, S)$, $H = (A, O, T)$.
The *asynchronous observation relation* $\sim_a$ is the smallest equivalence relation such that:

- $(G, \epsilon) \sim_a (H, \epsilon)$ iff $N_a = O_a$ and $S_a = T_a$
- $(G, \sigma.bc) \sim_a (H, \tau)$ iff $(G, \sigma) \sim_a (H, \tau)$ and $a \notin \{b, c\}$
- $(G, \sigma.ab) \sim_a (H, \tau.ab)$ and $(G, \sigma.ba) \sim_a (H, \tau.ba)$ iff $(G, \sigma) \sim_a (H, \tau)$ and $S_b^\sigma = T_b^\tau$

The *synchronous observation relation* $\approx_a$ is the smallest ...s.t.:

- $(G, \epsilon) \approx_a (H, \epsilon)$ iff $N_a = O_a$ and $S_a = T_a$
- $(G, \sigma.bc) \approx_a (H, \tau.de)$ iff $(G, \sigma) \approx_a (H, \tau)$ and $a \notin \{b, c, d, e\}$
- $(G, \sigma.ab) \approx_a (H, \tau.ab)$ and $(G, \sigma.ba) \approx_a (H, \tau.ba)$ iff $(G, \sigma) \approx_a (H, \tau)$ and $S_b^\sigma = T_b^\tau$

$(G, \sigma) \sim_a (H, \tau)$ implies $S_a^\sigma = T_a^\tau$; $(G, \sigma) \approx_a (H, \tau)$ implies $S_a^\sigma = T_a^\tau$.

Note that $\approx_a \subseteq \sim_a$!

Observation relation for secret distributions

Recalling crucial clauses from the (a)synchronous relation:

- $(G, \sigma.ab) \sim_a (H, \tau.ab)$ iff $(G, \sigma) \sim_a (H, \tau)$ and $S^\sigma_b = T^\tau_b$
- $(G, \sigma.bc) \approx_a (H, \tau.de)$ iff $(G, \sigma) \approx_a (H, \tau)$ and $a \notin \{b, c, d, e\}$

Given (G, σ), agent a knows a proposition if it is true for all (H, τ) such that $(G, \sigma) \sim_a (H, \tau)$. Same for $\approx_a$. Precise but not formal!

If agents are not uncertain about a set of initial gossip graphs, but certain about the initial secret distribution, we get:

- $\epsilon \sim_a \epsilon$
- $\sigma.bc \sim_a \tau.de$ iff $\sigma \sim_a \tau$ and $a \notin \{b, c, d, e\}$
- $\sigma.ab \sim_a \tau.ab$ and $\sigma.ba \sim_a \tau.ba$ iff $\sigma \sim_a \tau$ and $\sigma_b = \tau_b$

where for the synchronous relation we write $\approx_a$ instead of $\sim_a$ and then get as the second clause

- $\sigma.bc \approx_a \tau.de$ iff $\sigma \approx_a \tau$ and $a \notin \{b, c, d, e\}$

Other observation relations

Recalling crucial clauses from the (a)synchronous relation:

- $(G, \sigma.ab) \sim_a (H, \tau.ab)$ iff $(G, \sigma) \sim_a (H, \tau)$ and $S_b^\sigma = T_b^\tau$
- $(G, \sigma.bc) \approx_a (H, \tau.de)$ iff $(G, \sigma) \approx_a (H, \tau)$ and $a \notin \{b, c, d, e\}$

Other observation relations ($\sim_a$ also used as **arbitrary** obs. rel.)

Agents observe all calls (e.g. cup phones; a synchronous relation)

- $(G, \sigma.bc) \approx_a (H, \tau.bc)$ iff $(G, \sigma) \approx_a (H, \tau)$, for any $b, c \in A$

Merge and inspect (agents see the output but not the input)

- $(G, \sigma.ab) \sim_a (H, \tau.ab)$ iff $(G, \sigma) \sim_a (H, \tau)$ and $S_b^\sigma \cup S_a^\sigma = T_b^\tau \cup T_a^\tau$

Asymmetric observation (a sees caller b but not the callee c)

- $(G, \sigma.bc) \approx_a (H, \tau.bd)$ iff $(G, \sigma) \approx_a (H, \tau)$ and $a \notin \{b, c, d\}$

All you know (full-information protocol in distributed computing)

- $(G, \sigma) \sim_a^{\mathsf{v}} (H, \tau)$ iff $N_a = O_a$, $S_a = T_a$, and $\mathsf{v}_a^\sim(\sigma) = \mathsf{v}_a^\sim(\tau)$

Note: $(G, \sigma.ab) \sim_a^{\mathsf{v}} (H, \tau.ab)$ iff $(G, \sigma) \sim_a^{\mathsf{v}} (H, \tau)$ and $(G, \sigma) \sim_b^{\mathsf{v}} (H, \tau)$!