Reasoning about Gossip

Hans van Ditmarsch
CNRS

- Reachability

Materials found on http://reasoningaboutgossip.eu

Reachability of secret distributions

Suppose you don't know what communication protocol is executed but you can register the communications actually taking place. Can you determine the communication protocol producing them?

In terms of gossip protocols and secret distributions: You know what secret distributions have been reached by the agents when executing the protocol. Can you determine the gossip protocol?

Let us consider three agents a, b, c. What secret distributions are reachable from the initial secret distribution $a|b|c$?

- some are unreachable by any protocol: $ab|b|c$, $abc|bc|bc$, ...
- $ab|ab|c$, $a|bc|bc$ are both reachable but similar (isomorphic)
- there are four reachable secret distributions that are not isomorphic: $a|b|c$, $ab|ab|c$, $abc|ab|abc$, $abc|abc|abc$

We determine the reachability hierarchy for the protocols ANY, CMO, LNS, TOK, SPI.

Reachability of secret distributions

- Given gossip protocol P, secret distribution S is P-reachable if there is a P-permitted call sequence σ such that $I^\sigma = S$.
- Secret distributions S and T are isomorphic if there is a permutation ι of agents such that $\iota(S) = T$.

We now determine the set $PC(S)$ of potential calls for secret distribution S, and the order $<_S$ between those potential calls. Secret distribution S can only be obtained by a call sequence in which calls in $PC(S)$ occur and that respects the order $<_S$.

- The set of potential calls $PC(S)$ for secret distribution S consists of all ab with $a \in S_b$ and $b \in S_a$.
- Given $ab, cd \in PV(S)$:

$$ab <_S cd \quad \text{if} \quad \begin{array}{l} a = c \text{ and } d \notin S_b \text{ or} \\ b = c \text{ and } d \notin S_a \text{ or} \\ a = d \text{ and } c \notin S_b \text{ or} \\ b = d \text{ and } c \notin S_a. \end{array}$$

S determines $(PC(S), <_S)$ but $(PC(S), <_S)$ does not determine S.

Potential call sequences

- $PC(S) = \{ab \mid a \in S_b \text{ and } b \in S_a\}$

$$ab <_S cd \quad \text{if} \quad \begin{array}{l} a = c \text{ and } d \notin S_b \text{ or} \\ b = c \text{ and } d \notin S_a \text{ or} \\ a = d \text{ and } c \notin S_b \text{ or} \\ b = d \text{ and } c \notin S_a \end{array}$$

Consider $S = abcd|abcd|abcd|abcd$. $PC(S)$ consists of all 12 calls $\{ab, ac, ad, bc, bd, cd, ba, ca, da, cb, db, dc\}$, and $<_S = \emptyset$: no order is imposed. Reaching S are $ab.cd.ac.bd$ and $ac.bd.ab.cd$. Call sequence $ab.ac.bd.cd$ does not reach S, but respects $<_S$.

Consider $S' = abc|abc|bcd|bcd$. $PC(S) = \ldots, <_S = \ldots$

Consider $T = ab|abc|bc$. Then $PC(T) = \{ab, bc, ba, cb\}$, and $ab, ba <_T bc, cb$ and $bc, cb <_T ab, ba$. No call sequence respecting $<_T$ reaches T. T is not ANY-reachable.

Reachability hierarchy

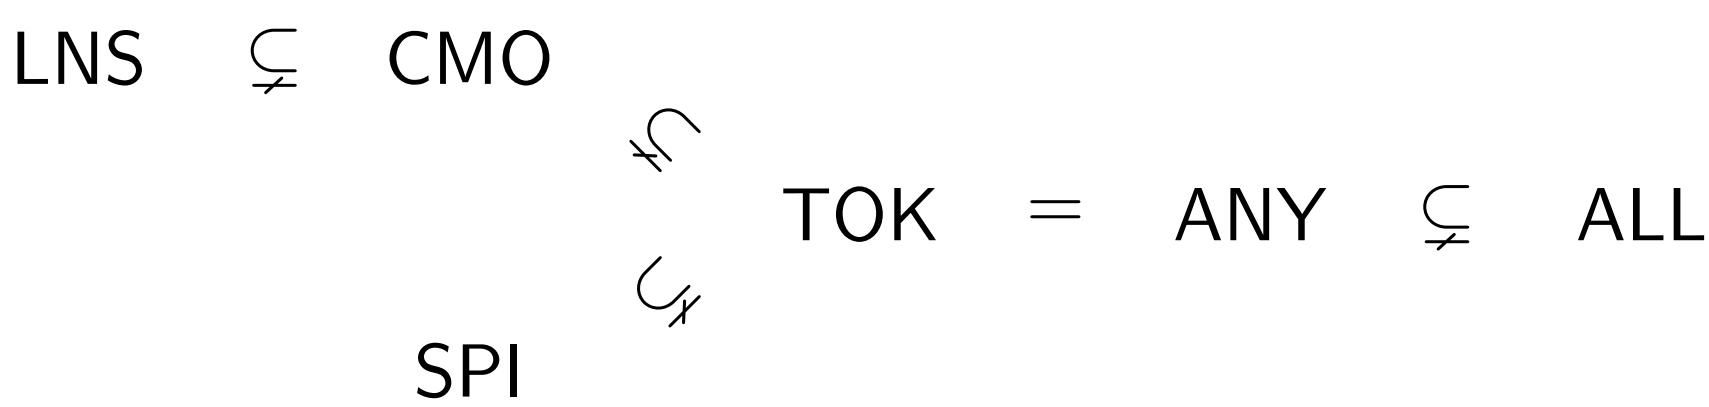

$\mathsf{P} \subseteq \mathsf{P}'$ means: for all σ, if σ is P-reachable then σ is P′-reachable.

ALL is not a gossip protocol but represents the set of all secret distributions (the set of all subsets of A^2 that are reflexive).

What again are the other gossip protocols?

Reachability hierarchy — details

ANY-reachable is TOK-reachable

This can be shown by induction on the length of ANY call sequences using the token density lemma:

> *Any* TOK-*reachable distribution* S *can be reached by a* TOK-*permitted call sequence* σ *such that after the execution of* σ *at least one of any two given agents* a *and* b *have a token.*

Proof:

If S is reached by a single call, it is ANY and TOK reachable. Now assume S is reached by $\sigma.ab$. By induction there is a TOK-perm. sequence τ with $I^\sigma = I^\tau$. By the token density lemma there is a τ' where a or b have a token with $I^\tau = I^{\tau'}$. So either $\tau'.ab$ or $\tau'.ba$ is TOK-permitted, and $I^{\sigma.ab} = I^{\tau'.ab} = I^{\tau'.ba}$.

Reachability hierarchy — details

There is a SPI-reachable distribution that is not CMO-reachable.

Secret distribution $abcd|abcd|abc|abd$ is a counter example.
Call sequence $ab.ad.cb.ab$ is SPI-permitted but not CMO-perm.

$$a|b|c|d \xrightarrow{ab} ab|ab|c|d \xrightarrow{ad} abd|ab|c|abd \xrightarrow{cb}$$
$$abd|abc|abc|abd \xrightarrow{ab} abcd|abcd|abc|abd$$

It is only a little bit more work to show that this is 'without loss of generality', that is, that no other call sequence exists also reaching that secret distribution but that is after all CMO-permitted.

Reachability of secret distributions

There is a SPI-/CMO-reachable distribution that is not LNS-reachable.

Counter-example: $abcdef|abc|abcde|abcdef|def|abdef$

There is a LNS-reachable distribution that is not SPI-reachable.

Counter-example for 16 agents and 22 calls! It is not know whether 16 is optimal. Showing off giving it without details:

$\sigma := \sigma_1.\sigma_2.\sigma_3$, where

$$\begin{aligned} \sigma_1 &= && 12.34.56.78.ab.cd.ef.gh \\ \sigma_2 &= && 23.45.67.81.bc.de.fg.ha \\ \sigma_3 &= && 1a.4c.7h.6f \end{aligned}$$

There is . . . (combining obtained results completes hierarchy)

Useful Resource: Encyclopedia of Integer Sequences

Number of non-isomorphic reachable distributions for ≤ 5 agents

n	LNS	CMO	SPI	TOK = ANY
2	2	2	2	2
3	4	4	4	4
4	15	15	16	16
5	97	97	111	111

This table allows us to conclude that some counterexamples are optimal. The 4-distribution $abcd|abcd|abc|abd$ that is SPI and not CMO reachable must be optimal, because SPI and CMO have the same number of 3-distributions. (Various such conclusions.)

See On-Line Encyclopedia of Integer Sequences for LNS and ANY: `https://oeis.org/A307085` and `https://oeis.org/A318154`. By Ioannis Kokkinis; 2025 additions for $n = 8$ by Bert Dobbelaere.

Subreachability of Secret Distributions

Secret distribution $ab|b$ is unreachable. Given agents a, b, c and call sequence $bc.ca$ we reach distribution $abc|bc|abc$. Restrict the distribution to agents a and b only. We obtain $ab|b$. This means: Secret distribution $ab|b$ is *subreachable*.

Subreachability is relevant when agents only know their neighbours are are even unaware of other agents.

All secret distributions are subreachable by all five protocols:
The proof is by induction on the number of secrets known by the agents in a secret distribution S for agents A. Base case: simple. Inductive case: If some agent a only knows its own secret, use induction on $S|(A\backslash\{a\})$. Otherwise, some agent a also knows some secret $b \neq a$. Consider S' which is as S except that a does not know the secret of b. By induction, S' is subreachable by some σ. Consider $bc.\sigma.ca$ for agents $A \cup \{c\}$ reaching distribution S''. The A-restriction of S'' (forget c) is S... (works for ANY, CMO, LNS)

Parallel Reachability of Secret Distributions

Let us now permit rounds of parallel calls. We can now define (where the old notion is *sequential* P-*reachable*:

Distribution S is *parallel* P-*reachable* if it can be obtained by a sequence of rounds of P-permitted calls.

Distribution $ab|abc|bc$ is parallel ANY-reachable by the round of simultaneous calls $\{ab, cb\}$, wherein agent b simultaneously receives the secret of a and the secret of c.

Distribution $abcdef|abc|abcde|abcdef|def|abdef$ is parallel LNS-reachable by the sequence consisting of three rounds of calls $\{ab, cb, de, fe\}.\{ca, dc, fd, af\}.\{da\}$. We successively get:

$ab|abc|bc|de|def|ef \rightarrow abcef|abc|abcde|bcdef|def|abdef \rightarrow abcdef|abc|abcde|abcdef|def|abdef$

But it was not sequential LNS-reachable! (See a previous slide.)
To be continued . . .