

k-set agreement in communication networks with omission faults

Emmanuel Godard Eloi Perdereau



Laboratoire d'Informatique Fondamentale

02/10/2017

DESCARTES

2-4 octobre 2017 - Chasseneuil-du-Poitou

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

Distributed model

The distributed model :

- Synchronous;
- Message passing;
- Underlying *connected* communication graph $G = (V, E)$ fixed;
- Dynamic network;
- f omission faults,
i.e. at each round : f messages can be lost.

The k -set agreement problem

The k -set agreement problem :

- **k -Agreement** The set of output values contains at most k elements.
- **Validity** The decided values are ones of the initial values.
- **Termination** All processes must decide.

Message adversaries

First introduced by Afek & Gafni (2013)

- Message adversary : set of infinite sequences of digraphs (*instant graphs*) defining the messages received in each round.
- *Oblivious* message adversary : the set of potential graphs in each round remains constant all along the execution.
⇒ Set of instant digraphs.

The f -omission message adversary : $\mathbf{O}_f(G)$

For example, the message adversary that allows at most f faults :

Definition (f -omissions message adversary)

$$\mathbf{O}_f(G) = \{G' = (V, A') \mid A' \subseteq A \wedge |A| - |A'| \leq f\}$$

Example of f -omission message adversary : $\mathbf{O}_1(P_3)$





The f -omission message adversary : $\mathbf{O}_f(G)$

For example, the message adversary that allows at most f faults :

Definition (f -omissions message adversary)

$$\mathbf{O}_f(G) = \{G' = (V, A') \mid A' \subseteq A \wedge |A| - |A'| \leq f\}$$

Example of f -omission message adversary : $\mathbf{O}_1(P_3)$



Our contribution : computability of k -set agreement

How many lost messages (omissions) the k -set agreement is tolerant to ?

Theorem

Let $k \in \mathbb{N}$ and $G = (V, E)$ be any communication network. The k -set agreement problem is solvable despite f omission faults if and only if $f \leq c_k(G)$.

$c_k(G)$ is the maximum number of (*undirected*) edges that can be removed without disconnecting G in $k + 1$ components.

i.e. removing $c_k(G)$ edges from G keeps at most k connected components on G .

\Rightarrow *the standard connectivity is $c_1(G) + 1$.*

Our contribution : computability of k -set agreement

How many lost messages (omissions) the k -set agreement is tolerant to ?

Theorem

Let $k \in \mathbb{N}$ and $G = (V, E)$ be any communication network. The k -set agreement problem is solvable despite f omission faults if and only if $f \leq c_k(G)$.

$c_k(G)$ is the maximum number of (*undirected*) edges that can be removed without disconnecting G in $k + 1$ components.

i.e. removing $c_k(G)$ edges from G keeps at most k connected components on G .

\Rightarrow the standard connectivity is $c_1(G) + 1$.

Our contribution : computability of k -set agreement

How many lost messages (omissions) the k -set agreement is tolerant to ?


Theorem


Let $k \in \mathbb{N}$ and $G = (V, E)$ be any communication network. The k -set agreement problem is solvable despite f omission faults if and only if $f \leq c_k(G)$.


$c_k(G)$ is the maximum number of (*undirected*) edges that can be removed without disconnecting G in $k + 1$ components.


i.e. removing $c_k(G)$ edges from G keeps at most k connected components on G .


\Rightarrow *the standard connectivity is $c_1(G) + 1$.*


Example of $c_k(G)$ 


Example of $c_k(G)$ 


Example of $c_k(G)$ 

Example of $c_k(G)$ 


Example of $c_k(G)$ 

Example of $c_k(G)$ 

Example of $c_k(G)$ 

Example of $c_k(G)$ 

Example of $c_k(G)$




$$c_1(G) = 0$$

$$c_2(G) = 3$$

$$c_3(G) = 5$$

Example of $c_k(G)$




$$c_1(G) = 0$$

$$c_2(G) = 3$$

$$c_3(G) = 5$$

Example of $c_k(G)$



$$c_1(G) = 0$$

$$c_2(G) = 3$$

$$c_3(G) = 5$$

Outline

- 1 Model and problem
- 2 Impossibility proofs
 - Impossibility of 2-set in K_3 and P_3 ($|G| = k + 1 = 3$)
 - Reduction from the case $|G| > k + 1$
 - Impossibility of set-agreement for generalized tournaments
- 3 Possibility : a priority-based algorithm

Outline



- 1 Model and problem
- 2 Impossibility proofs
 - Impossibility of 2-set in K_3 and P_3 ($|G| = k + 1 = 3$)
 - Reduction from the case $|G| > k + 1$
 - Impossibility of set-agreement for generalized tournaments
- 3 Possibility : a priority-based algorithm

Outline

- 1 Model and problem
- 2 Impossibility proofs
 - Impossibility of 2-set in K_3 and P_3 ($|G| = k + 1 = 3$)
 - Reduction from the case $|G| > k + 1$
 - Impossibility of set-agreement for generalized tournaments
- 3 Possibility : a priority-based algorithm

Outline


- 1 Model and problem
- 2 Impossibility proofs
 - Impossibility of 2-set in K_3 and P_3 ($|G| = k + 1 = 3$)
 - Reduction from the case $|G| > k + 1$
 - Impossibility of set-agreement for generalized tournaments
- 3 Possibility : a priority-based algorithm

Results for K_3 and P_3 (a) K_3 (b) P_3


$f \backslash k$	1	2	3
1	yes	yes	yes
2	no	yes	yes
3	no	no	yes

$f \backslash k$	1	2	3
1	no	yes	yes
2	no	no	yes
3	no	no	yes

Classical impossibility proof in $\mathbf{O}_3(K_3)$



Classical impossibility proof in $\mathbf{O}_3(K_3)$



Half-duplex message adversary

Remark : only “half-duplex” graphs are depicted, *i.e.* we allow only one arc to be removed between two nodes, **not both**.

Definition (Half-Duplex graphs)

$$\mathbf{HD}_f(G) = \{G' = (V, A') \mid G' \in \mathbf{O}_f(G) \wedge \\ \forall p, q \in V \quad \{p, q\} \in E \wedge (p, q) \notin A' \Rightarrow (q, p) \in A'\}$$

$$\Rightarrow \mathbf{HD}_f(G) \subseteq \mathbf{O}_f(G)$$

Proposition

The 2-set agreement problem is impossible in $\mathbf{HD}_2(P_3)$.

Half-duplex message adversary

Remark : only “half-duplex” graphs are depicted, *i.e.* we allow only one arc to be removed between two nodes, **not both**.

Definition (Half-Duplex graphs)

$$\mathbf{HD}_f(G) = \{G' = (V, A') \mid G' \in \mathbf{O}_f(G) \wedge \\ \forall p, q \in V \quad \{p, q\} \in E \wedge (p, q) \notin A' \Rightarrow (q, p) \in A'\}$$

$$\Rightarrow \mathbf{HD}_f(G) \subseteq \mathbf{O}_f(G)$$

Proposition

The 2-set agreement problem is impossible in $\mathbf{HD}_2(P_3)$.

Half-duplex message adversary

Remark : only “half-duplex” graphs are depicted, *i.e.* we allow only one arc to be removed between two nodes, **not both**.

Definition (Half-Duplex graphs)


$$\mathbf{HD}_f(G) = \{G' = (V, A') \mid G' \in \mathbf{O}_f(G) \wedge \\ \forall p, q \in V \quad \{p, q\} \in E \wedge (p, q) \notin A' \Rightarrow (q, p) \in A'\}$$

$$\Rightarrow \mathbf{HD}_f(G) \subseteq \mathbf{O}_f(G)$$


Proposition

The 2-set agreement problem is impossible in $\mathbf{HD}_2(P_3)$.

Impossibility of (2-)set agreement in P_3




Impossibility of (2-)set agreement in P_3




Reduction to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ for the 2-set (1)

For arbitrary $G = (V, E)$:


- Suppose $f > c_2(G)$;
- Partition V in 3 sets V_1, V_2 and V_3 ;
- An algorithm \mathcal{A} for G would solve the 2-set agreement for all $\mathbf{M} \subseteq \mathbf{O}_f(G)$;
- In particular in $\mathbf{S} \subseteq \mathbf{HD}_f(G)$ where omissions are synchronized between V_i s;
- Reduce to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ by syncing omissions.

Reduction to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ for the 2-set (2)


\Rightarrow Generalizable for all k : k -set reduced to set agreement in H
 ($|H| = k + 1$).

Reduction to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ for the 2-set (2)

\Rightarrow Generalizable for all k : k -set reduced to set agreement in H
 ($|H| = k + 1$).

Reduction to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ for the 2-set (2)

\Rightarrow Generalizable for all k : k -set reduced to set agreement in H
 $(|H| = k + 1)$.


Reduction to $\mathbf{HD}_3(K_3)$ or $\mathbf{HD}_2(P_3)$ for the 2-set (2)

\Rightarrow Generalizable for all k : k -set reduced to set agreement in H
 ($|H| = k + 1$).


How to generalize

- we had to consider both $\mathbf{HD}(K_3, 3)$ and $\mathbf{HD}(P_3, 2)$
- we need a way to handle all $\mathbf{HD}(G, f)$ we are reducing to (generalized tournaments)
- \Rightarrow subdivision **diagram**

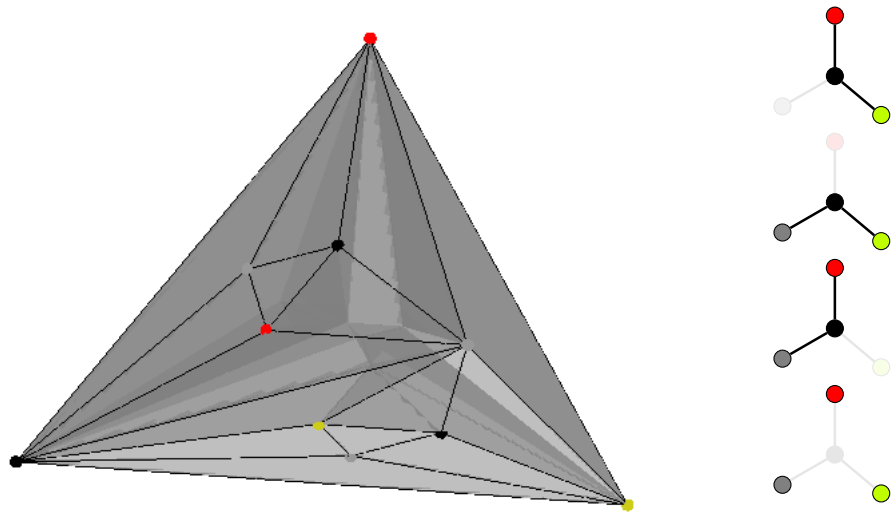
Construction of the subdivision diagram of S_3




Construction of the subdivision diagram of S_3



Construction of the subdivision diagram of S_3



Construction of the subdivision diagram of S_3



Outline

- 1 Model and problem
- 2 Impossibility proofs
 - Impossibility of 2-set in K_3 and P_3 ($|G| = k + 1 = 3$)
 - Reduction from the case $|G| > k + 1$
 - Impossibility of set-agreement for generalized tournaments
- 3 Possibility : a priority-based algorithm

An algorithm for the k -set when $f \leq c_k(G)$

Priority order π on $\{p_1, p_2, \dots, p_n\} : \pi(p_i) = i$

Algorithm: Algorithm for the k -set agreement in G for process p_i

$known \leftarrow i ;$

for T rounds **do**

 send $known$ to all neighbors ;

$known \leftarrow known \cup$ received ids ;

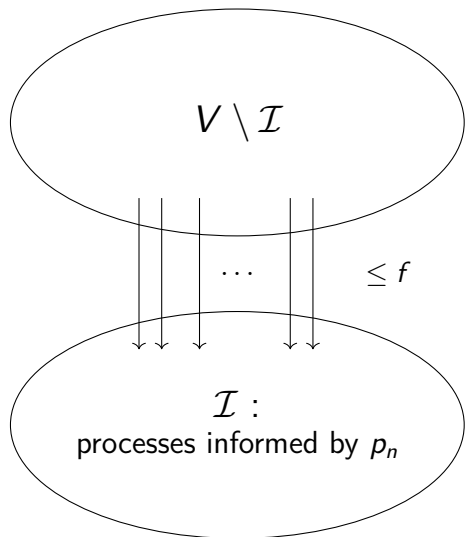
end

decide $\max(known) ;$

Proof of possibility when $f = c_2(G)$ (1)

p_n is the process with the highest priority.

- Let \mathcal{I} : the set of process informed by p_n ;
- Consider the processes eventually in \mathcal{I} ;
- If $\mathcal{I} = V$, the consensus is solved;
- Otherwise, $\bar{\mathcal{I}} = V \setminus \mathcal{I}$ form a strongly connected component because $f = c_2(G)$;
- After enough rounds, processes in $\bar{\mathcal{I}}$ have the same information.

Proof of possibility when $f = c_2(G)$ (2)

Conclusion and perspectives

Complete characterization of k -set agreement in oblivious message adversaries with f omission faults : solvable iff $f \leq c_k(G)$.

⇒ New **proof technique** for impossibility with **subdivision diagrams**.

Perspectives of generalization :

- Managing *partitioning vs uncertainty*,
- Find a generic method to construct a subdivision diagram for arbitrary oblivious message adversaries;

Conclusion and perspectives

Complete characterization of k -set agreement in oblivious message adversaries with f omission faults : solvable iff $f \leq c_k(G)$.

⇒ New **proof technique** for impossibility with **subdivision diagrams**.

Perspectives of generalization :

- Managing *partitioning vs uncertainty*,
- Find a generic method to construct a subdivision diagram for arbitrary oblivious message adversaries;

Contribution to DESCARTES

$HD(G, f)$ can be simulated by $\mathbf{HD}(K_\kappa, \kappa(\kappa - 1)/2)$ which is equivalent to the wait-free RW model for κ processes where κ is the smallest integer such that $c_\kappa(G) \geq f$.

Colorless Computability

$$RW_{waitfree}(\kappa) \subseteq HD(G, f) \subsetneq RW_{waitfree}(\kappa - 1)$$

Thank you