


Characterizing consensus in the Heard-Of model

Igor Walukiewicz

joint work with A.R. Balasubramanian



Consensus problem:

Every process starts with the same value of its input variable. We require:

- **termination:** every process eventually sets its dec value,
- **agreement:** all dec variables have the same value,
- every dec variable is set at most once,
- the value of dec variables is one of the initial values of input variables.

FLP theorem: consensus is impossible in fully asynchronous systems in the presence of faults

2/3 algorithm

the multiset of values received from other processes

```
send (inp)
| if uni(H)  $\wedge$  |H|  $>$   $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;
```

```
| if mult(H)  $\wedge$  |H|  $>$   $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;
```


```
send  $x_1$ 
```

```
| if uni(H)  $\wedge$  |H|  $>$   $\frac{2}{3}n$  then  $dec := \text{smor}(H)$ ;
```

*smallest most frequent value
in H*

Communication predicate: eventually $\psi^1 = (\varphi = \wedge \varphi_{\frac{2}{3}}, \text{true})$ and later
 $\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$

round:



2/3 algorithm

send (*inp*)

if $\text{uni}(H) \wedge |H| > \frac{2}{3}n$ then $x_1 := \text{inp} := \text{smor}(H)$;

if $\text{mult}(H) \wedge |H| > \frac{2}{3}n$ then $x_1 := \text{inp} := \text{smor}(H)$;

send x_1

if $\text{uni}(H) \wedge |H| > \frac{2}{3}n$ then $\text{dec} := \text{smor}(H)$;

Communication predicate: eventually $\psi^1 = (\varphi = \wedge \varphi_{\frac{2}{3}}, \text{true})$ and later

$\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$

the multiset of values received from other processes

smallest most frequent value in H

every process receives the same multiset


H contains values from $\geq \frac{2}{3}$ of processes

round:

process 1: $x_i = v_i^1$

⋮

process n: $x_i = v_i^n$



$\dots \rightarrow x_{i+1} = v_{i+1}^1$

⋮



$\dots \rightarrow x_{i+1} = v_{i+1}^n$

2/3 algorithm

send (*inp*)

if $\text{uni}(H) \wedge |H| > \frac{2}{3}n$ then $x_1 := \text{inp} := \text{smor}(H)$;

if $\text{mult}(H) \wedge |H| > \frac{2}{3}n$ then $x_1 := \text{inp} := \text{smor}(H)$;

send x_1

if $\text{uni}(H) \wedge |H| > \frac{2}{3}n$ then $\text{dec} := \text{smor}(H)$;

Communication predicate: eventually $\psi^1 = (\varphi = \wedge \varphi_{\frac{2}{3}}, \text{true})$ and later

$\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$

the multiset of values received from other processes

smallest most frequent value
in H

every process receives the same multiset

H contains values from $\geq \frac{2}{3}$ of processes

- At $(\varphi = \wedge \varphi_{\frac{2}{3}}, \text{true})$ phase, every process sets *inp* to the same value
- At $(\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$ phase later, every process sets *dec* to this value

Q : can we change $\frac{2}{3}$ to $\frac{1}{2}$?

2/3 algorithm

the multiset of values received from other processes

```

send (inp)
| if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;
| if mult(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;

```

send x_1

```

| if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $dec := \text{smor}(H)$ ;

```

smallest most frequent value in H

Communication predicate: eventually $\psi^1 = (\varphi = \wedge \varphi_{\frac{2}{3}}, \text{true})$ and later

$$\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$$

every process receives the same multiset

H contains values from $\geq \frac{2}{3}$ of processes

Q : can we change $\frac{2}{3}$ to $\frac{1}{2}$?

$$f = \{ \overbrace{a_1, \dots, a_1}^{\frac{1}{2}}, \overbrace{b_1, \dots, b_1}^{\frac{1}{2}} \}$$

H_a - more a's

H_b - more b's

$$f' = \{ \overbrace{a_1, \dots, a_1}^{\frac{1}{2} + \epsilon}, \overbrace{b_1, \dots, b_1}^{\frac{1}{2} - \epsilon} \}$$

send $H = \{ \overbrace{a_1, \dots, a_1}^{\frac{1}{2} + \epsilon} \}$ to p_1

and nothing to others. So p_1 dec on a.

$$\epsilon \quad \frac{1}{2} - \epsilon$$

In the next phase send $H_{b1} = \{ \overbrace{a_1, \dots, a_1}^{\epsilon}, \overbrace{b_1, \dots, b_1}^{\frac{1}{2} - \epsilon} \}$ to every process. They decide on b.

Heard-off model

Introduced by Bernadette Charron-Bost · André Schiper in 2009

A round based model for non synchronous computing.

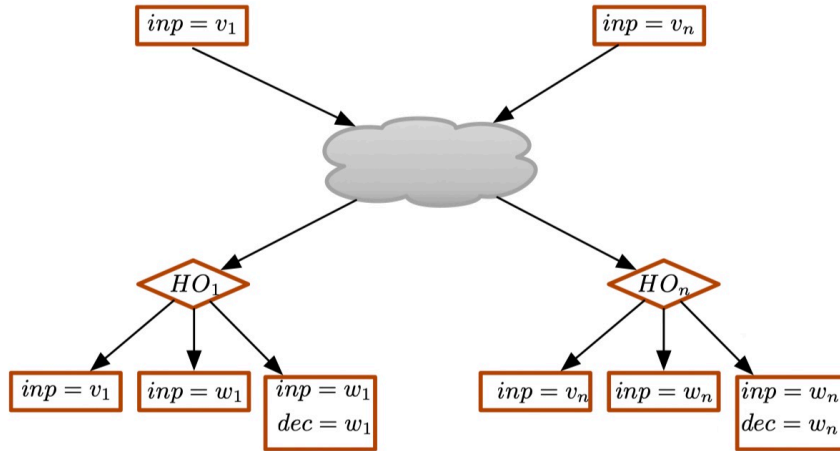
Unified treatment of different types of faults through transmission faults.

A model is relatively simple and concise:

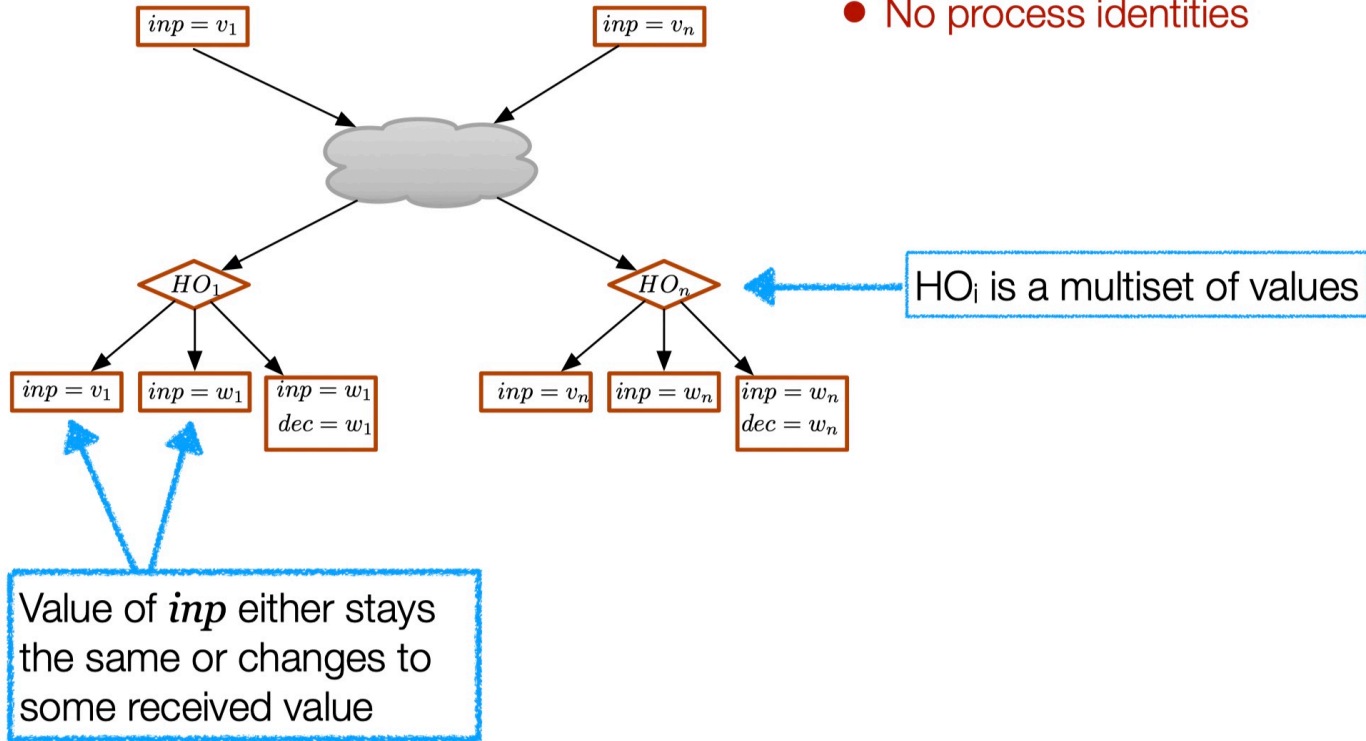
a good candidate to develop verification methods

- [Charron-Bost, Stefan Merz,..] Efficient encoding the model in Isabelle, and TLA
- [Drăgoi, Henzinger, Zufferey,..] A semi-automatic proof method, a domain-specific language based on HO-model.
- [Ognjen Maric, Christoph Sprenger, David Basin, *Cut-off Bounds for Consensus Algorithms*], see later
- [R. Bloem, S. Jacobs, A. Khalimov, I. Konnov, S. Rubin, H. Veith, and J. Widder. *Decidability of Parameterized Verification*], a book, 2015

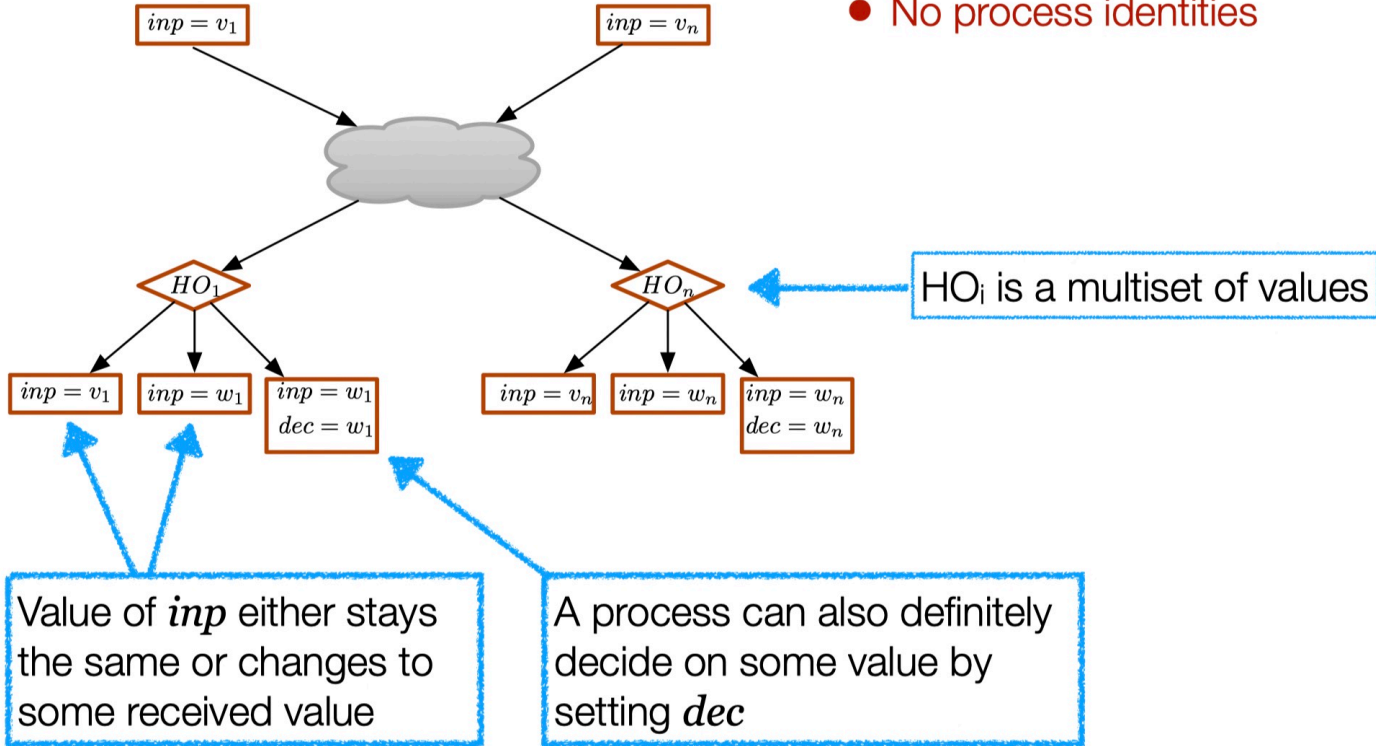
- No operations on variables
- No failure of components
- No process identities



- No operations on variables
- No failure of components
- No process identities



- No operations on variables
- No failure of components
- No process identities




Value of inp either stays the same or changes to some received value

A process can also definitely decide on some value by setting dec


HO_i is a multiset of values

Heard-of algorithm


round:



phase:



phase:



send *inp*

```

if uni(H) ∧ |H| > thru1 · n then x1 := op01(H);
:
if mult(H) ∧ |H| > thrm1,k · n then x1 := opk1(H);

```

input sent, x₁ set

send *x_{i-1}*

```

if uni(H) ∧ |H| > thrui · n then xi := op0i(H);
:
if mult(H) ∧ |H| > thrmi,k · n then xi := opki(H);

```

x_{i-1} sent, x_i set

send *x_{ir-1}*

```

if uni(H) ∧ |H| > thruir · n then xir := inp := op0ir(H);
:
if mult(H) ∧ |H| > thrmir,k · n then xir := inp := opkir(H);

```

round r_i where inp is set [update]

send *x_{l-1}*


```

if uni(H) ∧ |H| > thrul · n then dec := op0l(H);
:
if mult(H) ∧ |H| > thrml,k · n then dec := opkl(H);

```

last round where dec is set [decision]

phase:



send inp

```

if uni(H)  $\wedge$  |H| >  $thr_u^1 \cdot n$  then  $x_1 := op_0^1(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{1,k} \cdot n$  then  $x_1 := op_k^1(H)$ ;

```

send x_{i-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^i \cdot n$  then  $x_i := op_0^i(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{i,k} \cdot n$  then  $x_i := op_k^i(H)$ ;

```

send x_{ir-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^{ir} \cdot n$  then  $x_{ir} := inp := op_0^{ir}(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{ir,k} \cdot n$  then  $x_{ir} := inp := op_k^{ir}(H)$ ;

```

send x_{l-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^l \cdot n$  then  $dec := op_0^l(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{l,k} \cdot n$  then  $dec := op_k^l(H)$ ;

```

send (inp)

```

if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := smor(H)$ ;
if mult(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := smor(H)$ ;

```


send x_1

```

if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $dec := smor(H)$ ;

```

Communication predicate: eventually $\psi^1 = (\varphi = \wedge \varphi_{\frac{2}{3}}, true)$ and later
 $\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$



send inp

```

if uni(H)  $\wedge$   $|H| > thr_u^1 \cdot n$  then  $x_1 := op_0^1(H)$ ;
:
if mult(H)  $\wedge$   $|H| > thr_m^{1,k} \cdot n$  then  $x_1 := op_k^1(H)$ ;

```

send x_{i-1}

```

if uni(H)  $\wedge$   $|H| > thr_u^i \cdot n$  then  $x_i := op_0^i(H)$ ;
:
if mult(H)  $\wedge$   $|H| > thr_m^{i,k} \cdot n$  then  $x_i := op_k^i(H)$ ;

```

send x_{ir-1}

```

if uni(H)  $\wedge$   $|H| > thr_u^{ir} \cdot n$  then  $x_{ir} := inp := op_0^{ir}(H)$ ;
:
if mult(H)  $\wedge$   $|H| > thr_m^{ir,k} \cdot n$  then  $x_{ir} := inp := op_k^{ir}(H)$ ;

```

send x_{l-1}

```


if uni(H)  $\wedge$   $|H| > thr_u^l \cdot n$  then  $dec := op_0^l(H)$ ;
:
if mult(H)  $\wedge$   $|H| > thr_m^{l,k} \cdot n$  then  $dec := op_k^l(H)$ ;

```

Wanted:

- Given an algorithm over a fixed set of values decide if it solves the consensus
- What operations and tests are allowed?
- Do we have a cut-of principle? restriction to k -proc
- Do we have a 0/1 principle? restriction to $\{0,1\}$ values

phase:



send inp

```

if uni(H)  $\wedge$  |H| >  $thr_u^1 \cdot n$  then  $x_1 := op_0^1(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{1,k} \cdot n$  then  $x_1 := op_k^1(H)$ ;

```

Semantics:

phase round
 $(f, d) \xrightarrow{\psi} (f', d')$ and $f \xrightarrow{\varphi}_i f'$.

$f: \{1, \dots, n\} \rightarrow \text{Dvt}??$ $d: \{1, \dots, n\} \rightarrow \text{Dvt}??$
 \mathcal{D} finite and linearly ordered

send x_{i-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^i \cdot n$  then  $x_i := op_0^i(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{i,k} \cdot n$  then  $x_i := op_k^i(H)$ ;

```

send x_{ir-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^{ir} \cdot n$  then  $x_{ir} := inp := op_0^{ir}(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{ir,k} \cdot n$  then  $x_{ir} := inp := op_k^{ir}(H)$ ;

```


send x_{l-1}

```

if uni(H)  $\wedge$  |H| >  $thr_u^l \cdot n$  then  $dec := op_0^l(H)$ ;
:
if mult(H)  $\wedge$  |H| >  $thr_m^{l,k} \cdot n$  then  $dec := op_k^l(H)$ ;

```

phase:



send inp

```

if uni(H) ∧ |H| > thru1 · n then x1 := op01(H);
⋮
if mult(H) ∧ |H| > thrm1,k · n then x1 := opk1(H);

```

Semantics:

phase round
 $(f, d) \xrightarrow{\psi} (f', d')$ and $f \xrightarrow{\varphi}_i f'$.

• $f: \{1, \dots, n\} \rightarrow D \cup \{?\}$ $d: \{1, \dots, n\} \rightarrow D \cup \{?\}$
 D finite and linearly ordered

send x_{i-1}

```

if uni(H) ∧ |H| > thrui · n then xi := op0i(H);
⋮
if mult(H) ∧ |H| > thrmi,k · n then xi := opki(H);

```

• $f \xrightarrow{e}_i f'$ if $\exists (H_1, \dots, H_n) \models e \wedge \forall p \ H_p \in \text{mset}(f)$
 $f'(p) = \text{update}_i(H_p)$

$\min(H)$ min value in H - ?

$\text{smor}(H)$ smallest most frequent value in H - ?

$f'(p) = ?$ if no test holds for H_p

Rem: tests count value but operations ignore it

send x_{ir-1}

```

if uni(H) ∧ |H| > thruir · n then xir := inp := op0ir(H);
⋮
if mult(H) ∧ |H| > thrmir,k · n then xir := inp := opkir(H);


```

send x_{l-1}

```

if uni(H) ∧ |H| > thrul · n then dec := op0l(H);
⋮
if mult(H) ∧ |H| > thrml,k · n then dec := opkl(H);

```

send inp

```

if uni(H) ∧ |H| > thr_u^1 · n then x_1 := op_0^1(H);
⋮
if mult(H) ∧ |H| > thr_m^{1,k} · n then x_1 := op_k^1(H);

```

Semantics:

phase round
 $(f, d) \xrightarrow{\psi} (f', d')$ and $f \xrightarrow{\varphi}_i f'$.

$f: \{1, \dots, n\} \rightarrow \text{Dvt}^?$ $d: \{1, \dots, n\} \rightarrow \text{Dvt}^?$
 \emptyset finite and linearly ordered

send x_{i-1}

```

if uni(H) ∧ |H| > thr_u^i · n then x_i := op_0^i(H);
⋮
if mult(H) ∧ |H| > thr_m^{i,k} · n then x_i := op_k^i(H);

```

$f \xrightarrow{u}_i f'$ if $\exists (H_1, \dots, H_n) \neq \emptyset \forall p \ H_p \in \text{mset}(f)$
 $f'(p) = \text{update}_i(H_p)$

$\min(H)$ min value in H -? ?
 $\text{smor}(H)$ smallest most frequent value in H -? ?
 $f'(p) = ?$ if no test holds for H_p

send x_{ir-1}

```

if uni(H) ∧ |H| > thr_u^{ir} · n then x_{ir} := inp := op_0^{ir}(H);
⋮
if mult(H) ∧ |H| > thr_m^{ir,k} · n then x_{ir} := inp := op_k^{ir}(H);

```

$(f, d) \xrightarrow{\varphi} (f', d')$ if $f_0 \xrightarrow{u_1} f_1 \xrightarrow{u_2} \dots \xrightarrow{u_c} f_c$

$f_0 = f$

$f'(p) = f_{ir}(p)$ if $f_{ir}(p) \neq ?$ and $f(p) = f_{ir}(p)$ otherwise

$d'(p) = d(p)$ if $d(p) \neq ?$ and $d'(p) = f_c(p)$ otherwise


send x_{l-1}

```

if uni(H) ∧ |H| > thr_u^l · n then dec := op_0^l(H);
⋮
if mult(H) ∧ |H| > thr_m^{l,k} · n then dec := op_k^l(H);

```

phase:



send inp

```

if uni(H) ∧ |H| > thr_u^1 · n then x_1 := op_0^1(H);
⋮
if mult(H) ∧ |H| > thr_m^{1,k} · n then x_1 := op_k^1(H);
    
```

Semantics:

phase round
 $(f, d) \xrightarrow{\psi} (f', d')$ and $f \xrightarrow{\varphi}_i f'$.

$f: \{1, \dots, n\} \rightarrow \text{Dvt}??$ $d: \{1, \dots, n\} \rightarrow \text{Dvt}??$
 \emptyset finite and linearly ordered

send x_{i-1}

```

if uni(H) ∧ |H| > thr_u^i · n then x_i := op_0^i(H);
⋮
if mult(H) ∧ |H| > thr_m^{i,k} · n then x_i := op_k^i(H);
    
```

$f \xrightarrow{\varphi}_i f'$ if $\exists (H_1, \dots, H_n) \models \varphi \wedge \forall p \ H_p \in \text{mset}(f)$
 $f(p) = \text{update}_i(H_p)$

send x_{ir-1}

```

if uni(H) ∧ |H| > thr_u^{ir} · n then x_{ir} := inp := op_0^{ir}(H);
⋮
if mult(H) ∧ |H| > thr_m^{ir,k} · n then x_{ir} := inp := op_k^{ir}(H);
    
```

Communication predicate

$(G\bar{\psi}) \wedge (F(\psi_1 \wedge F(\psi_2 \wedge \dots (F\psi_k) \dots)))$

$\psi = (\varphi_1, \dots, \varphi_n)$
 ↑ round ↑ predicates

← phase predicates ↑

φ_i is $(|H| > thr \cdot n)$ or $(\text{eq } n \wedge |H| > thr \cdot n)$

send x_{l-1}

```

if uni(H) ∧ |H| > thr_u^l · n then dec := op_0^l(H);
⋮
if mult(H) ∧ |H| > thr_m^{l,k} · n then dec := op_k^l(H);
    
```

A Characterization

► **Theorem 22.** An algorithm solves consensus iff it is:

- syntactically safe,
- there are $i \leq j$ with ψ_i a unifier and ψ_j a decider.

$$(G\bar{\psi}) \wedge (F(\psi_1 \wedge F(\psi_2 \wedge \dots (F\psi_k) \dots)))$$

We fix $D = \{a, b\}$ (we show 0/1-principle with our proof)

Some notation for $f: \{?, \dots, n\} \rightarrow D \cup \{?\}$ $\text{bias}(\theta) = (a, \dots, a, \underbrace{b, \dots, b}_{\theta \cdot n})$ $\text{bias}^2(\theta) = (?, \dots, ?, \underbrace{b, \dots, b}_{\theta \cdot n})$
 $\text{solo} = (b, \dots, b)$ $\text{solo}^? = (?, \dots, ?)$

A round i is solo safe wrt. ψ if $\text{thr}_a^i \leq \text{thr}_i(\psi)$ ($\text{solo} \stackrel{\psi}{=} \text{solo}$)

A round i is preserving wrt. ψ if either:

- no uni instruction, or
- no mult instruction, or
- $\text{thr}_i(\psi) < \max(\text{thr}_a^i, \text{thr}_m^{i,\psi})$

Rem $\text{bias}(\theta) \stackrel{\psi}{=} \text{solo}^?$ is possible

A Characterization

► **Theorem 22.** An algorithm solves consensus iff it is:

- syntactically safe,
- there are $i \leq j$ with ψ_i a unifier and ψ_j a decider.

$$(\overline{G\bar{\psi}}) \wedge (F(\psi_1 \wedge F(\psi_2 \wedge \dots (F\psi_k) \dots)))$$

A round i is solo safe wrt. ψ if $\text{thr}_u^i \leq \text{thr}_i(\psi)$


A round i is preserving wrt. ψ if either:

- no uni instruction, or
- no mult instruction, or
- $\text{thr}_i(\psi) < \max(\text{thr}_u^i, \text{thr}_m^{i,k})$

Rem $\text{bites}(\theta) \stackrel{\psi}{\implies}_i$, solo? is possible

ψ is a decider if all rounds are solo safe wrt. ψ

ψ is a unifier if \exists round i s.t



- $\text{thr}_i(\psi) \geq \text{thr}_m^{i,k}$ and either $\text{thr}_i(\psi) \geq \text{thr}_u^i$ or $\text{thr}_i(\psi) \geq \overline{\text{thr}}$
 where $\overline{\text{thr}} = \max(1 - \text{thr}_u^i, 1 - \text{thr}_m^{i,k}/2)$

A Characterization

► **Theorem 22.** *An algorithm solves consensus iff it is:*

- *syntactically safe,*
- *there are $i \leq j$ with ψ_i a unifier and ψ_j a decider.*

► **Definition 9.** *An algorithm is syntactically safe when:*

1. *First round has a mult instruction.*
2. *Every round has a uni instruction.*
3. *In the first round the operation in every mult instruction is smor.*
4. *$thr_m^{1,k}/2 \geq 1 - thr_u^{ir+1}$, and $thr_u^1 \geq 1 - thr_u^{ir+1}$.*

```
send inp
| if uni(H) ∧ |H| > thr_u^1 · n then x_1 := op_0^1(H);
  :
  if mult(H) ∧ |H| > thr_m^{1,k} · n then x_1 := op_k^1(H);
```

```
send x_{i-1}
| if uni(H) ∧ |H| > thr_u^i · n then x_i := op_0^i(H);
  :
  if mult(H) ∧ |H| > thr_m^{i,k} · n then x_i := op_k^i(H);
```

```
send x_{ir-1}
| if uni(H) ∧ |H| > thr_u^{ir} · n then x_{ir} := inp := op_0^{ir}(H);
  :
  if mult(H) ∧ |H| > thr_m^{ir,k} · n then x_{ir} := inp := op_k^{ir}(H);
```

```
send x_{l-1}
| if uni(H) ∧ |H| > thr_u^l · n then dec := op_0^l(H);
  :
  if mult(H) ∧ |H| > thr_m^{l,k} · n then dec := op_k^l(H);
```

A Characterization

► **Theorem 22.** *An algorithm solves consensus iff it is:*

- *syntactically safe,*
- *there are $i \leq j$ with ψ_i a unifier and ψ_j a decider.*


► **Definition 9.** *An algorithm is syntactically safe when:*

1. *First round has a mult instruction.*
2. *Every round has a uni instruction.*
3. *In the first round the operation in every mult instruction is smor.*
4. *$\text{thr}_m^{1,k} / 2 \geq 1 - \text{thr}_u^{\text{ir}+1}$, and $\text{thr}_u^1 \geq 1 - \text{thr}_u^{\text{ir}+1}$.*

► **Lemma 23.** *An algorithm that is not syntactically safe cannot solve consensus. A syntactically safe algorithm has the agreement property.*

► **Lemma 24.** *A syntactically safe algorithm has the termination property iff it satisfies the second condition from Theorem 10.*

Elements of the proof



Lemma 0 If $f \stackrel{e}{=} f'$ and $a, b \in f'$ then $f \stackrel{e}{=} \text{bias}(\theta)$ for every θ .

Proof

θ not an equalizer if $(H_a, \dots, H_b) \models e$ then $(H_a', \dots, H_b') \models e$
for all $H_i' \in \{H_a, H_b\} \sqcup$

Elements of the proof

Lemma 2.9 If round i has mult instruction then there is $\theta \geq 1/2$ s.t.
 $\text{bias}(\theta) \stackrel{4}{=} \text{bias}(\theta')$ for arb. θ' .

Proof

If the operation is smor then take $\theta = 1/2$.

Construct $H_a > \text{thr}_m^i$ with more a than b. Similarly for H_b

If the operation is min, take $\theta > \max(\text{thr}_i(4), \text{thr}_{u_i}^i, 1/2)$

Complete H gives a

When H contains only b's it is big enough to give b. □

Elements of the proof

► **Lemma 23.** An algorithm that is not syntactically safe cannot solve consensus. A syntactically safe algorithm has the agreement property.

Proof of the second statement:

Consider $(bias(\theta), ?) \xrightarrow{\Psi} \dots \xrightarrow{\Psi} (bias(\theta'), d')$ first time some process, say p , decides

Recall $1 - thr_a^{i+1} \leq thr_m^{i+1}$ and $1 - thr_a^{i+1} \leq thr_a^i$.

First we show that round $i+1$ cannot have mult instruction.

Suppose $d(p) = a$.

This implies $\theta' < 1 - thr_a^{i+1}$ and $d(q) \in \{a, ?\}$, for all q .

Hence $\theta' < thr_a^i$, so b cannot be later obtained by uni instruction.

$\theta' < thr_m^{i+1}/2$ so b cannot be later obtained by mult instruction with smov,

(this is required for safe alg) \square

Extensions

Timestamps

```
send (inp, ts)
| if  $\text{cond}_1^1(H)$  then  $x_1 := \text{maxts}(H)$ ;
| :
| if  $\text{cond}_1^l(H)$  then  $x_1 := \text{maxts}(H)$ ;
```

Smallest of the values of the most recent timestamp

► **Definition 13.** An algorithm is syntactically t-safe when:

1. Every round has a uni instruction.
2. First round has a mult instruction.
3. $\text{thr}_m^{1,k} \geq 1 - \text{thr}_u^{\text{ir}+1}$ and $\text{thr}_u^1 \geq 1 - \text{thr}_u^{\text{ir}+1}$.

► **Definition 14.** A predicate ψ is a strong unifier ψ if it is a unifier in a sense of Definition 8 and $\text{thr}_u^1 \leq \text{thr}_1(\psi)$.

► **Theorem 25.** An algorithm satisfies consensus iff it is syntactically t-safe according to Definition 13, and it satisfies:

sT There are $i \leq j$ such that ψ^i is a strong unifier and ψ^j is a decider.

Extensions

Coordinators

- Three types of rounds
- lr (leader receive)
 - ls (leader send)
 - *every* (as before)

► **Theorem 26.** An algorithm satisfies consensus iff the first round and the $(ir + 1)^{th}$ round are not of type ls , it is syntactically safe according to Definition 9, and it satisfies the condition:

cT There are $i \leq j$ such that ψ^i is a c -unifier and ψ^j is a c -decider.

Coordinators + timestamps

► **Theorem 27.** An algorithm satisfies consensus iff the first round and the $(ir + 1)^{th}$ round are not of type ls , it has the structural properties from Definition 13, and it satisfies:

scT There are $i \leq j$ such that ψ^i is a strong c -unifier and ψ^j is a c -decider.

Examples

```
send (inp)
| if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;
| if mult(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $x_1 := inp := \text{smor}(H)$ ;
send  $x_1$ 
| if uni(H)  $\wedge$  |H| >  $\frac{2}{3}n$  then  $dec := \text{smor}(H)$ ;
Communication predicate: eventually  $\psi^1 = (\varphi_ = \wedge \varphi_{\frac{2}{3}}, \text{true})$  and later
 $\psi^2 = (\varphi_{\frac{2}{3}}, \varphi_{\frac{2}{3}})$ 
```

Because of the condition
 $\text{thr}_{n/2}^{i,k} \geq 1 - \text{thr}_u^{i+1}$

it is not possible to have
 $1/2$ threshold.

With timestamps the condition is weakened to $\text{thr}_n^{i,k} \geq 1 - \text{thr}_u^{i+1}$

```
send (inp, ts)
| if uni(H)  $\wedge$  |H| >  $1/2 \cdot |\Pi|$  then  $x_1 := \text{maxts}(H)$ ;
| if mult(H)  $\wedge$  |H| >  $1/2 \cdot |\Pi|$  then  $x_1 := \text{maxts}(H)$ ;
send  $x_1$ 
| if uni(H)  $\wedge$  |H| >  $1/2 \cdot |\Pi|$  then  $x_2 := inp := \text{smor}(H)$ ;
send  $x_2$ 
| if uni(H)  $\wedge$  |H| >  $1/2 \cdot |\Pi|$  then  $dec := \text{smor}(H)$ ;
Communication predicate:  $F(\psi^1)$  where  $\psi^1 := (\varphi_ = \wedge \varphi_{1/2}, \varphi_{1/2}, \varphi_{1/2})$ 
```

Examples

```
send (inp, ts)
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
| if mult(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
send  $x_1$ 
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_2 := \text{inp} := \text{smor}(H)$ ;
send  $x_2$ 
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $\text{dec} := \text{smor}(H)$ ;
Communication predicate:  $F(\psi^1)$  where  $\psi^1 := (\varphi = \wedge \varphi_{1/2}, \varphi_{1/2}, \varphi_{1/2})$ 
```

Another example with coordinators. Equalizer in round 2.

```
send (inp, ts)
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
| if mult(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
send  $x_1$ 
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_2 := \text{smor}(H)$ ;
| if mult(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_2 := \text{smor}(H)$ ;
send  $x_2$ 
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_3 := \text{inp} := \text{smor}(H)$ ;
send  $x_3$ 
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $\text{dec} := \text{smor}(H)$ ;
Communication predicate:  $F(\psi^1 \wedge F\psi^2)$ 
where:  $\psi^1 = (\varphi_{1/2}, \varphi = \wedge \varphi_{1/2}, \varphi_{1/2}, \text{true})$  and  $\psi^2 = (\varphi_{1/2}, \varphi_{1/2}, \varphi_{1/2}, \varphi_{1/2})$ 
```

Examples

Paxos

```
send (inp, ts) lr
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
| if mult(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
send  $x_1$  ls
| if uni(H) then  $x_2 := \text{inp} := \text{smor}(H)$ ;
send  $x_2$  lr
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_3 := \text{smor}(H)$ ;
send  $x_3$  ls
| if uni(H) then  $\text{dec} := \text{smor}(H)$ ;
Communication predicate:  $F(\psi^1)$  where  $\psi^1 := (\varphi_{1/2}, \varphi_{1s}, \varphi_{1/2}, \varphi_{1s})$ 
```

3-round Paxos

```
send (inp, ts) lr
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
| if mult(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $x_1 := \text{maxts}(H)$ ;
send  $x_1$  ls
| if uni(H) then  $x_2 := \text{inp} := \text{smor}(H)$ ;
send  $x_2$  every
| if uni(H)  $\wedge$  |H| > 1/2 · | $\Pi$ | then  $\text{dec} := \text{smor}(H)$ ;
Communication predicate:  $F(\psi^1)$  where  $\psi^1 := (\varphi_{1/2}, \varphi_{1s}, \varphi_{1/2})$ 
```

Conclusions

- We wanted to do verification but arrived at a characterization
- Distributed algorithms are not algorithms:
 - for a given setting there is very little freedom for a solution
 - consensus must happen in two rounds: unifier followed by decider

Challenges

- We still do not know how to do verification, or even specify properties.
- Finding the best algorithms.
- What are reasonable extensions of this model?
Ben-Or's algorithm.