Tolerating Corrupted Communication

Martin Biely Bernadette Charron-Bost Antoine Gaillard
Martin Hutle André Schiper Josef Widder

Technische Universität Wien

École Polytechnique

École Polytechnique Fédérale de Lausanne

PODC 2007, Portland, Oregon, August 12–15

Some lower bounds for consensus

Santoro & Widmayer, 1989 less than $\lfloor \frac{n}{2} \rfloor$ faulty transmissions per round

Some lower bounds for consensus

Santoro & Widmayer, 1989 less than $\lfloor \frac{n}{2} \rfloor$ faulty transmissions per round

Martin & Alvisi, 2006 $n > 5f$ for fast Byzantine consensus

Some lower bounds for consensus

Santoro & Widmayer, 1989 less than $\lfloor \frac{n}{2} \rfloor$ faulty transmissions per round

Martin & Alvisi, 2006 $n > 5f$ for fast Byzantine consensus

Faulty transmissions per round:

	Safety + Liveness	
[SW89]	$< \frac{n}{2}$	
[MA06]	$< \frac{n^2}{5}$	

Some lower bounds for consensus

Santoro & Widmayer, 1989 less than $\lfloor \frac{n}{2} \rfloor$ faulty transmissions per round

Martin & Alvisi, 2006 $n > 5f$ for fast Byzantine consensus

Faulty transmissions per round:

	Safety + Liveness	Safety	Liveness
[SW89]	$< \frac{n}{2}$	?	?
[MA06]	$< \frac{n^2}{5}$	?	?

Do these bounds also hold, if we consider safety and liveness separately?

Some lower bounds for consensus

Santoro & Widmayer, 1989 less than $\lfloor \frac{n}{2} \rfloor$ faulty transmissions per round

Martin & Alvisi, 2006 $n > 5f$ for fast Byzantine consensus

Faulty transmissions per round:

	Safety + Liveness	Safety	Liveness
[SW89]	$< \frac{n}{2}$	?	?
[MA06]	$< \frac{n^2}{5}$	?	?

Do these bounds also hold, if we consider safety and liveness separately?

- Having lower requirements for safety than for liveness makes only sense in the context of dynamic and transient faults.

Outline

1 The HO model for value faults

2 Algorithms
- The $\mathcal{A}_{T,E}$ consensus algorithm
- The $\mathcal{U}_{T,E,\alpha}$ consensus algorithm ($\rightarrow$ paper)

3 Discussion
- Relation to lower bounds
- Conclusion

Outline

1 The HO model for value faults

2 Algorithms
- The $\mathcal{A}_{T,E}$ consensus algorithm
- The $\mathcal{U}_{T,E,\alpha}$ consensus algorithm ($\rightarrow$ paper)

3 Discussion
- Relation to lower bounds
- Conclusion

The HO model for value faults

- round based model of computation

Round based model of computation

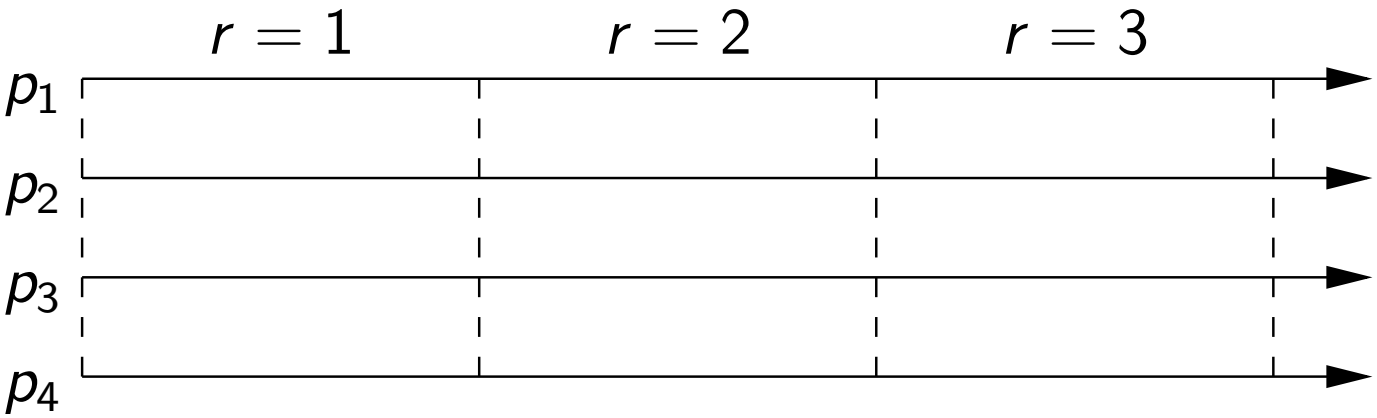

- computation proceeds in rounds

The HO model for value faults

- round based model of computation
- follows the *transmission fault* approach of Santoro & Widmayer

Value fault during a transmission of a message

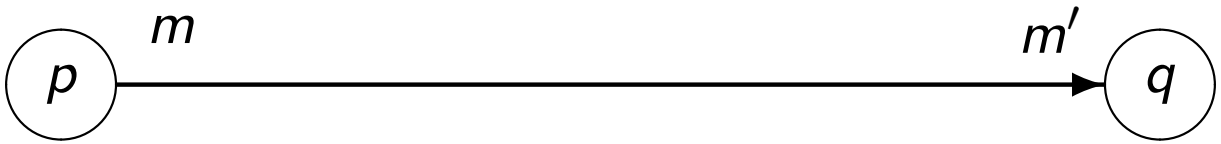

Value fault during a transmission of a message

send value fault at p

Value fault during a transmission of a message

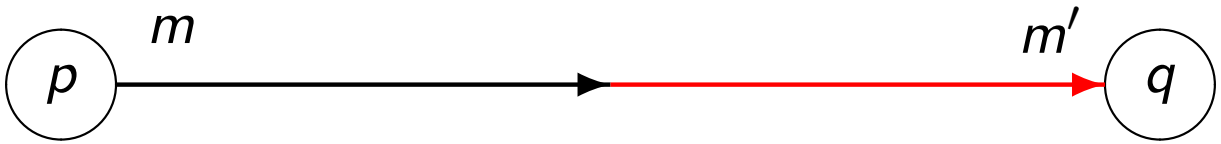

send value fault at p link faulty

Value fault during a transmission of a message

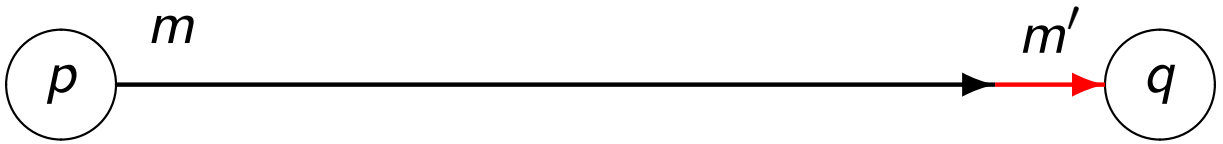

send value fault at p | link faulty | receive value fault at q

Value fault during a transmission of a message

Why distinguish these cases?

Value fault during a transmission of a message

Why distinguish these cases?

Transmission faults [SW89]:

- discrepancy between *what should have been sent* and *what has been received*

The HO model for value faults

- round based model of computation
- follows the *transmission fault* approach of Santoro & Widmayer
- generalizes the HO model for benign faults [CBS06]
  - HO = Heard-Of

The HO model for value faults

- round based model of computation
- follows the *transmission fault* approach of Santoro & Widmayer
- generalizes the HO model for benign faults [CBS06]
  - HO = Heard-Of
- able to deal with *transient* and *dynamic* faults
  - transient fault = non-permanent fault
  - dynamic fault = any component can be affected in a run by fault

Round based model of computation

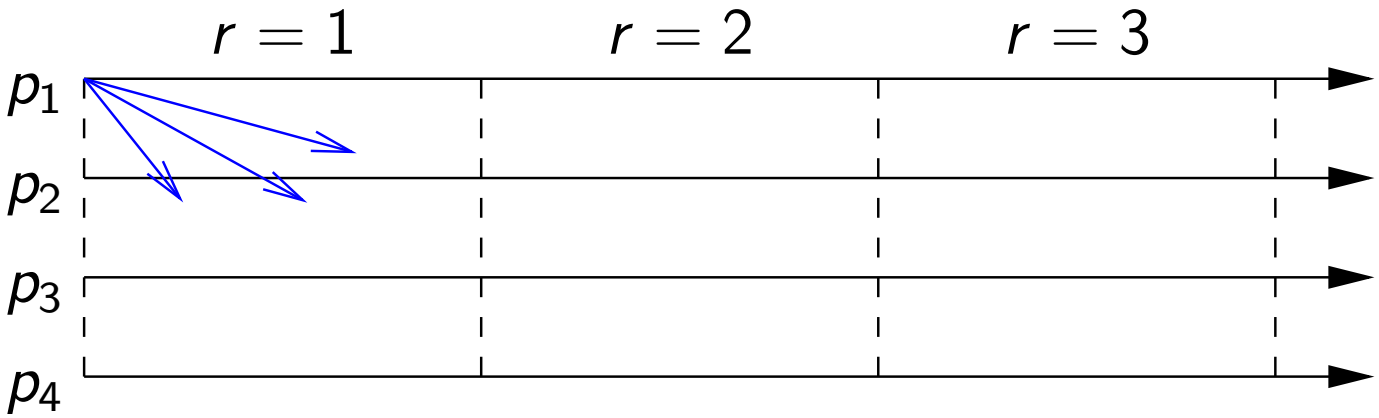

- computation proceeds in rounds
- send according to a sending function

Round based model of computation

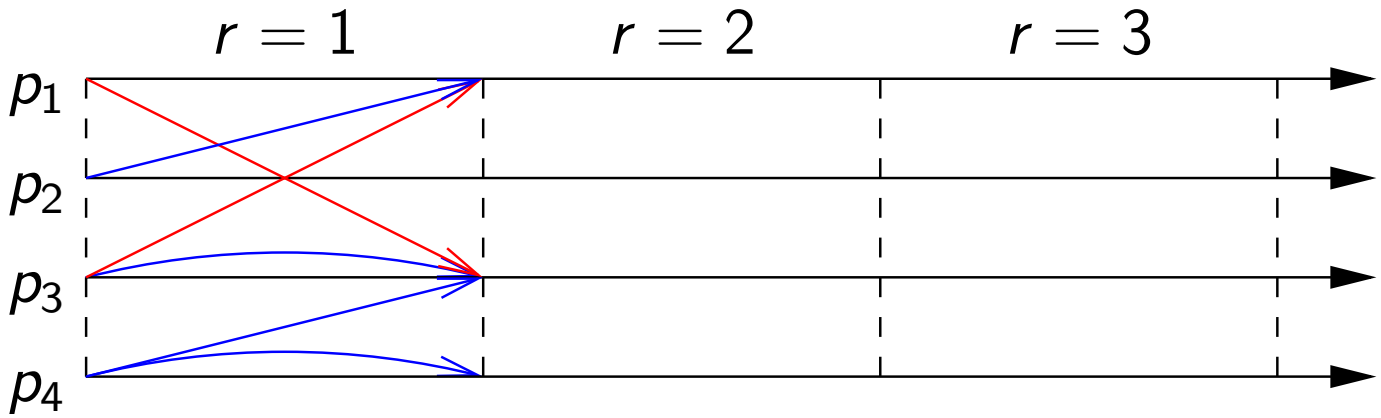

- computation proceeds in rounds
- send according to a sending function
- make state transition according to received messages

Round based model of computation

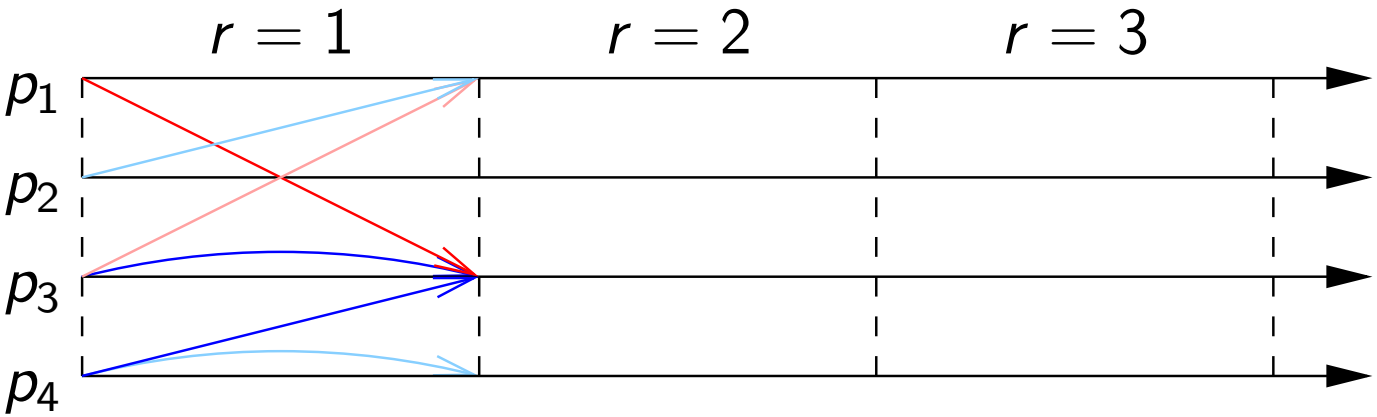

$HO(p_3, 1) = \{p_1, p_3, p_4\}$ $SHO(p_3, 1) = \{p_3, p_4\}$

- computation proceeds in rounds
- send according to a sending function
- make state transition according to received messages
- $HO(p, r) = \{q \mid p \text{ received a message from } q \text{ in round } r\}$
- $SHO(p, r) = \{q \mid \text{msg from } q \text{ in round } r \text{ is not corrupted}\}$

Communication predicates

A communication predicate is a predicate over the collection of all $HO(p, r)$, $SHO(p, r)$.

Communication predicates

A communication predicate is a predicate over the collection of all $HO(p, r)$, $SHO(p, r)$.

Useful abbreviations:

- altered HO set: $AHO(p, r) = HO(p, r) \setminus SHO(p, r)$

Communication predicates

A communication predicate is a predicate over the collection of all $HO(p, r)$, $SHO(p, r)$.

Useful abbreviations:

- altered HO set: $AHO(p, r) = HO(p, r) \setminus SHO(p, r)$
- altered span: $AS = \bigcup_{r>0, p \in \Pi} AHO(p, r)$

Communication predicates

A communication predicate is a predicate over the collection of all $HO(p, r)$, $SHO(p, r)$.

Useful abbreviations:

- altered HO set: $AHO(p, r) = HO(p, r) \setminus SHO(p, r)$
- altered span: $AS = \bigcup_{r>0, p \in \Pi} AHO(p, r)$

Example 1 (permanent static faults)

No more than f processes send corrupted information:

$$|AS| \leq f$$

Communication predicates

A communication predicate is a predicate over the collection of all $HO(p, r)$, $SHO(p, r)$.

Useful abbreviations:

- altered HO set: $AHO(p, r) = HO(p, r) \setminus SHO(p, r)$
- altered span: $AS = \bigcup_{r>0, p \in \Pi} AHO(p, r)$

Example 1 (permanent static faults)

No more than f processes send corrupted information:

$$|AS| \leq f$$

Not needed by our algorithms!

Outline

1 The HO model for value faults

2 Algorithms
- The $\mathcal{A}_{T,E}$ consensus algorithm
- The $\mathcal{U}_{T,E,\alpha}$ consensus algorithm ($\rightarrow$ paper)

3 Discussion
- Relation to lower bounds
- Conclusion

The $\mathcal{A}_{T,E}$ consensus algorithm

- parametrization of the *OneThirdRule* algorithm (benign case)

The $\mathcal{A}_{T,E}$ consensus algorithm

- parametrization of the *OneThirdRule* algorithm (benign case)
- the algorithm is always safe, independent of the number of benign faults

The $\mathcal{A}_{T,E}$ consensus algorithm

- parametrization of the *OneThirdRule* algorithm (benign case)
- the algorithm is always safe, independent of the number of benign faults
- the algorithm is *fast*

The $\mathcal{A}_{T,E}$ consensus algorithm

- parametrization of the *OneThirdRule* algorithm (benign case)
- the algorithm is always safe, independent of the number of benign faults
- the algorithm is *fast*
- allows up to $\frac{n-1}{4}$ corrupted messages per round and per process

The $\mathcal{A}_{T,E}$ consensus algorithm

round r, process p:

state

$x_p \in V$, initially v_p

The $\mathcal{A}_{T,E}$ consensus algorithm

round r, process p:

state

$x_p \in V$, initially v_p

sending function

send x_p to all processes

The $\mathcal{A}_{T,E}$ consensus algorithm

round r, process p:

state

$x_p \in V$, initially v_p

sending function

send x_p to all processes

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

The $\mathcal{A}_{T,E}$ consensus algorithm

round r, process p:

state

$x_p \in V$, initially v_p

sending function

send x_p to all processes

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

The $\mathcal{A}_{T,E}$ consensus algorithm

round r, process p:

state

$x_p \in V$, initially v_p

sending function

send x_p to all processes

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

α bounds value faults

T adopt threshold

E decision threshold

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

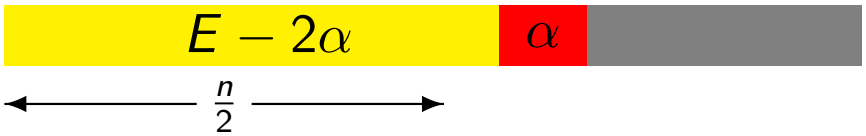

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

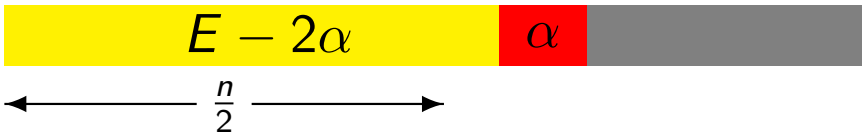

adopt values in round r_1:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

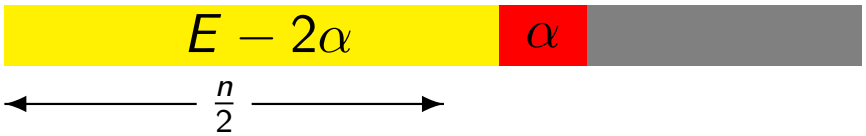

adopt values in round r_1:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

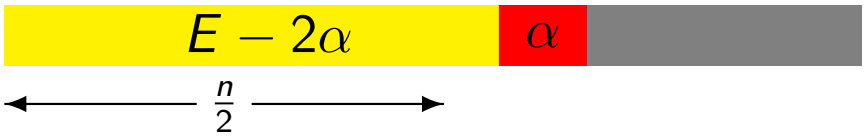

adopt values in round r_1:

estimates in round $r > r_1$:

The $\mathcal{A}_{T,E}$ consensus algorithm — safety

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

predicate for safety

$$|AHO(p,r)| \leq \alpha$$

parameters T, E, and α

$$n > E$$
$$n > T \geq 2(n + 2\alpha - E)$$

p decides v in round r_1:

estimates in round r_1:

q receives values in round r_1:

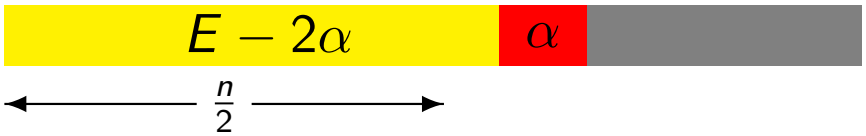

adopt values in round r_1:

estimates in round $r > r_1$:

similar for integrity

The $\mathcal{A}_{T,E}$ consensus algorithm — liveness

predicate for liveness = predicate for safety +:

Univalence: $|\Pi_1| > E - \alpha, |\Pi_2| > T : \forall p \in \Pi_1 : HO(p, r) = SHO(p, r) = \Pi_2$
Learning: $\forall p : |HO(p, r')| > T$
Deciding: $\forall p : |SHO(p, r'')| > E$

transition function

if $|HO(p, r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  $\text{DECIDE}(v)$

The $\mathcal{A}_{T,E}$ consensus algorithm — liveness

predicate for liveness = predicate for safety +:

Univalence: $|\Pi_1| > E-\alpha, |\Pi_2| > T : \forall p \in \Pi_1 : HO(p,r) = SHO(p,r) = \Pi_2$
Learning: $\forall p : |HO(p,r')| > T$
Deciding: $\forall p : |SHO(p,r'')| > E$

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  DECIDE(v)

- all processes in Π_1 have same estimate after r

The $\mathcal{A}_{T,E}$ consensus algorithm — liveness

predicate for liveness = predicate for safety +:

Univalence: $|\Pi_1| > E - \alpha, |\Pi_2| > T : \forall p \in \Pi_1 : HO(p,r) = SHO(p,r) = \Pi_2$
Learning: $\forall p : |HO(p,r')| > T$
Deciding: $\forall p : |SHO(p,r'')| > E$

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  DECIDE(v)

- all processes in Π_1 have same estimate after r
- all processes update estimate

The $\mathcal{A}_{T,E}$ consensus algorithm — liveness

predicate for liveness = predicate for safety +:

Univalence: $|\Pi_1| > E - \alpha, |\Pi_2| > T : \forall p \in \Pi_1 : HO(p,r) = SHO(p,r) = \Pi_2$
Learning: $\forall p : |HO(p,r')| > T$
Deciding: $\forall p : |SHO(p,r'')| > E$

transition function

if $|HO(p,r)| > T$ **then**
  $x_p :=$ the smallest most often received value in r
if more than E values are v **then**
  DECIDE(v)

- all processes in Π_1 have same estimate after r
- all processes update estimate
- all processes decide

Choice of E, T, and α

$$n > E$$

$$n > T \geq 2(n + 2\alpha - E)$$

Choice of E, T, and α

$$n > E$$

$$n > T \geq 2(n + 2\alpha - E)$$

Observation

If $\alpha < \frac{n}{4}$, there are E and T that fulfill the equations.

Choice of E, T, and α

$$n > E$$

$$n > T \geq 2(n + 2\alpha - E)$$

Observation

If $\alpha < \frac{n}{4}$, there are E and T that fulfill the equations.

E and T are inversely connected, but:

Choosing $E = T$

$T = E = \frac{2}{3}(n + 2\alpha)$ is a solution.

Choice of E, T, and α

$$n > E$$

$$n > T \geq 2(n + 2\alpha - E)$$

Observation

If $\alpha < \frac{n}{4}$, there are E and T that fulfill the equations.

E and T are inversely connected, but:

Choosing $E = T$

$T = E = \frac{2}{3}(n + 2\alpha)$ is a solution.

If we choose $\alpha = 0$, we get the OneThirdRule algorithm for the benign case [CBS06, HS07].

Outline

1 The HO model for value faults

2 Algorithms
- The $\mathcal{A}_{T,E}$ consensus algorithm
- The $\mathcal{U}_{T,E,\alpha}$ consensus algorithm ($\rightarrow$ paper)

3 Discussion
- Relation to lower bounds
- Conclusion

Some lower bounds for consensus

Faulty transmissions per round:

	Safety + Liveness	Safety	Liveness
[SW89]	$< \frac{n}{2}$	?	?
[MA06]	$< \frac{n^2}{5}$	?	?
	Lower bounds	Upper bound	

Some lower bounds for consensus

Faulty transmissions per round:

	Safety + Liveness	Safety	Liveness
[SW89]	$< \frac{n}{2}$	$< \frac{n^2}{4}$	0
[MA06]	$< \frac{n^2}{5}$	$< \frac{n^2}{4}$	0
	Lower bounds	Upper bound	

Conjectured lower bound of Lamport

Lower bound

$$N > 2Q + F + 2M$$

N # acceptors

M # byzantine acceptors despite which safety is ensured

F # byzantine acceptors despite which liveness is ensured

Q # byzantine acceptors despite which the protocol is fast

Conjectured lower bound of Lamport

Lower bound

$$N > 2Q + F + 2M$$

N # acceptors

M # byzantine acceptors despite which safety is ensured

F # byzantine acceptors despite which liveness is ensured

Q # byzantine acceptors despite which the protocol is fast

$\mathcal{A}_{T,E}$: attains bound with $F = 0$ in the case $\alpha = (n-1)/4$

Conjectured lower bound of Lamport

Lower bound

$$N > 2Q + F + 2M$$

N # acceptors

M # byzantine acceptors despite which safety is ensured

F # byzantine acceptors despite which liveness is ensured

Q # byzantine acceptors despite which the protocol is fast

$\mathcal{A}_{T,E}$: attains bound with $F = 0$ in the case $\alpha = (n-1)/4$
$\mathcal{U}_{T,E,\alpha}$: attains bound with $F = Q = 0$ in the case $\alpha = (n-1)/2$

Conclusion

In this work we:

- solved consensus in the presence of *transient* and *dynamic value faults*
  - separate liveness and safety conditions
  - liveness conditions hold only sporadically
  - lower bounds for static faults do not necessarily apply
  - our algorithm attains the lower bounds of Lamport

Conclusion

In this work we:

- solved consensus in the presence of *transient* and *dynamic value faults*
  - separate liveness and safety conditions
  - liveness conditions hold only sporadically
  - lower bounds for static faults do not necessarily apply
  - our algorithm attains the lower bounds of Lamport
- introduced the *HO model for value faults*
  - can express permanent fault assumptions (Byzantine processes) as well as transient and dynamic fault models
  - all assumptions are expressed by a predicate (comparison!)

Thank you for your attention

References I


Bernadette Charron-Bost and André Schiper.
The heard-of model: Unifying all benign failures.
Technical Report LSR-REPORT-2006-004, EPFL, 2006.


Martin Hutle and André Schiper.
Communication predicates: A high-level abstraction for coping with transient and dynamic faults.
In *Dependable Systems and Networks (DSN 2007)*, pages 92–10. IEEE, June 2007.


J.-P. Martin and L. Alvisi.
Fast byzantine consensus.
Transactions on Dependable and Secure Computing, 3(3):202–214, 2006.


Nicola Santoro and Peter Widmayer.
Time is not a healer.
In *Proc. 6th Annual Symposium on Theor. Aspects of Computer Science (STACS'89)*, volume 349 of *LNCS*, pages 304–313, Paderborn, Germany, February 1989. Springer-Verlag.