

# Reduction Revisited: Verifying Round-Based Distributed Algorithms

Stephan Merz

INRIA Nancy & LORIA

INSTITUT NATIONAL  
DE RECHERCHE  
EN INFORMATIQUE  
ET EN AUTOMATIQUE



INRIA

centre de recherche NANCY - GRAND-EST



joint work with Bernadette Charron-Bost, LIX & CNRS

MPC 2010  
June 23, 2010

# Example: mutual exclusion algorithms

```
integer  turn = 0;
boolean req0, req1 = false;

process P0
loop
  nc0: skip;
  rq0: req0 := true;
  ps0: turn := 1;
  wt0: await ¬req1 ∨ turn = 0;
  cs0: skip;
  ex0: req0 := false;
endloop

||

process P1
loop
  nc1: skip;
  rq1: req1 := true;
  ps1: turn := 0;
  wt1: await ¬req0 ∨ turn = 1;
  cs1: skip;
  ex1: req1 := false;
endloop
```

- Critical section can be abstracted to atomic step

# Example: mutual exclusion algorithms

```
integer  turn = 0;
boolean  req0, req1 = false;

process P0
loop
  nc0: skip;
  rq0: ⟨req0 := true;
        turn := 1;⟩
  wt0: await ¬req1 ∨ turn = 0;
  cs0: skip;
  ex0: req0 := false;
endloop

||

process P1
loop
  nc1: skip;
  rq1: ⟨req1 := true;
        turn := 0;⟩
  wt1: await ¬req0 ∨ turn = 1;
  cs1: skip;
  ex1: req1 := false;
endloop
```

- Critical section can be abstracted to atomic step
- Is it okay to combine the following actions into an atomic step?
  - 1 statements  $rq_i$  and  $ps_i$

# Example: mutual exclusion algorithms

```
integer  turn = 0;
boolean  req0, req1 = false;

process P0
loop
  nc0: skip;
  rq0: ⟨req0 := true;
        turn := 1;
        await ¬req1 ∨ turn = 0;⟩
  cs0: skip;
  ex0: req0 := false;
endloop

||

process P1
loop
  nc1: skip;
  rq1: ⟨req1 := true;
        turn := 0;
        await ¬req0 ∨ turn = 1;⟩
  cs1: skip;
  ex1: req1 := false;
endloop
```

- Critical section can be abstracted to atomic step
- Is it okay to combine the following actions into an atomic step?
  - 1 statements  $rq_i$  and  $ps_i$
  - 2 statements  $rq_i, ps_i,$  and  $wt_i$

# Example: mutual exclusion algorithms

```
integer  turn = 0;
boolean req0, req1 = false;

process P0
loop
  nc0: skip;
  rq0: req0 := true;
  ps0: turn := 1;
  wt0: await ¬req1 ∨ turn = 0;
  cs0: ⟨skip;
        req0 := false;⟩
endloop

process P1
loop
  nc1: skip;
  rq1: req1 := true;
  ps1: turn := 0;
  wt1: await ¬req0 ∨ turn = 1;
  cs1: ⟨skip;
        req1 := false;⟩
endloop
```

- Critical section can be abstracted to atomic step
- Is it okay to combine the following actions into an atomic step?
  - 1 statements  $rq_i$  and  $ps_i$
  - 2 statements  $rq_i$ ,  $ps_i$ , and  $wt_i$
  - 3 statements  $cs_i$  and  $ex_i$

# Outline

- 1 Reduction Theorems for the Verification of Concurrent Programs
- 2 Fault-Tolerant Distributed Computing
- 3 Reduction for Round-Based Distributed Algorithms
- 4 Experiments: Verification of Consensus Algorithms
- 5 Conclusion

# Reduction: overall idea

- Justify combining subsequent operations into an atomic step
- Fewer atomic steps  $\rightsquigarrow$  simpler verification

## Theorem (folklore)

*One can pretend that a sequence of statements is executed atomically if it contains at most one access to a shared variable.*

- Folk theorem justifies combining  $cs_i$  and  $ex_i$  (previous example)
- Folk theorem does not justify combining  $rq_i$  and  $ps_i$

# Reduction: overall idea

- Justify combining subsequent operations into an atomic step
- Fewer atomic steps  $\rightsquigarrow$  simpler verification

## Theorem (folklore)

*One can pretend that a sequence of statements is executed atomically if it contains at most one access to a shared variable.*

- Folk theorem justifies combining  $cs_i$  and  $ex_i$  (previous example)
- Folk theorem does not justify combining  $rq_i$  and  $ps_i$
- Consider the single-process program where initially  $x = y$

$$y := x + 1; x := y$$

Since no variable is shared, it should be equivalent to

$$\langle y := x + 1; x := y \rangle$$



# Reduction: overall idea

- Justify combining subsequent operations into an atomic step
- Fewer atomic steps  $\rightsquigarrow$  simpler verification

## Theorem (folklore)

*One can pretend that a sequence of statements is executed atomically if it contains at most one access to a shared variable.*

- Folk theorem justifies combining  $cs_i$  and  $ex_i$  (previous example)
- Folk theorem does not justify combining  $rq_i$  and  $ps_i$
- Consider the single-process program where initially  $x = y$

$$y := x + 1; x := y$$

Since no variable is shared, it should be equivalent to

$$\langle y := x + 1; x := y \rangle$$

**But the latter program satisfies  $\square(x = y)$  !**

# Left and right movers

## Definition (Lipton 1975)

An action  $a$  is a **right mover** if whenever  $\alpha ab$  is a computation where  $a$  and  $b$  are performed by different processes then  $\alpha ba$  is also a computation and these computations result in the same state. The definition of a **left mover** is symmetrical.

- **Right mover**  $s \xrightarrow{ab} t \Rightarrow s \xrightarrow{ba} t$  for all  $b$ 
  - ▶ right commutes with every action of different processes
  - ▶ example: acquisitions of resources (e.g., semaphores)
- **Left mover**  $s \xrightarrow{ba} t \Rightarrow s \xrightarrow{ab} t$  for all  $b$ 
  - ▶ left commutes with every action of different processes
  - ▶ example: releases of resources

*R.J. Lipton. Reduction: A Method of Proving Properties of Parallel Programs. CACM 18(12):717-721, 1975.*

# Left and right movers in example

```
integer  turn = 0;  
boolean req0, req1 = false;
```

```
process P0
```

```
loop
```

```
  nc0: skip;
```

```
  rq0: req0 := true;
```

```
  ps0: turn := 1;
```

```
  wt0: await ¬req1 ∨ turn = 0;
```

```
  cs0: skip;
```

```
  ex0: req0 := false;
```

```
endloop
```

```
process P1
```

```
loop
```

```
  nc1: skip;
```

```
  rq1: req1 := true;
```

```
  ps1: turn := 0;
```

```
  wt1: await ¬req0 ∨ turn = 1;
```

```
  cs1: skip;
```

```
  ex1: req1 := false;
```

```
endloop
```

```
||
```

- Actions  $rq_i$  are right movers

- ▶ in particular, cannot make **await** condition of other process true
- ▶ formally,  $s \xrightarrow{rq_0 \ wt_1} t$  implies  $s \xrightarrow{wt_1 \ rq_0} t$

# Left and right movers in example

```
integer  turn = 0;  
boolean req0, req1 = false;
```

```
process P0
```

```
loop
```

```
  nc0: skip;
```

```
  rq0: req0 := true;
```

```
  ps0: turn := 1;
```

```
  wt0: await ¬req1 ∨ turn = 0;
```

```
  cs0: skip;
```

```
  ex0: req0 := false;
```

```
endloop
```

```
process P1
```

```
loop
```

```
  nc1: skip;
```

```
  rq1: req1 := true;
```

```
  ps1: turn := 0;
```

```
  wt1: await ¬req0 ∨ turn = 1;
```

```
  cs1: skip;
```

```
  ex1: req1 := false;
```

```
endloop
```

```
||
```

- Actions  $rq_i$  are right movers

- ▶ in particular, cannot make **await** condition of other process true
- ▶ formally,  $s \xrightarrow{rq_0 \ wt_1} t$  implies  $s \xrightarrow{wt_1 \ rq_0} t$

- Actions  $cs_i$  and  $ex_i$  are left movers

# Left and right movers in example

```
integer  turn = 0;  
boolean req0, req1 = false;
```

```
process P0
```

```
loop
```

```
  nc0: skip;
```

```
  rq0: req0 := true;
```

```
  ps0: turn := 1;
```

```
  wt0: await ¬req1 ∨ turn = 0;
```

```
  cs0: skip;
```

```
  ex0: req0 := false;
```

```
endloop
```

```
process P1
```

```
loop
```

```
  nc1: skip;
```

```
  rq1: req1 := true;
```

```
  ps1: turn := 0;
```

```
  wt1: await ¬req0 ∨ turn = 1;
```

```
  cs1: skip;
```

```
  ex1: req1 := false;
```

```
endloop
```

```
||
```

- Actions  $rq_i$  are right movers

- ▶ in particular, cannot make **await** condition of other process true
- ▶ formally,  $s \xrightarrow{rq_0 \ wt_1} t$  implies  $s \xrightarrow{wt_1 \ rq_0} t$

- Actions  $cs_i$  and  $ex_i$  are left movers

- Actions  $ps_i$  and  $wt_i$  are neither left nor right movers

# Lipton's reduction theorem

## Theorem (Lipton 1975)

Suppose that  $A = A_1; \dots; A_k$  is such that for some  $i$ :

- $A_1, \dots, A_{i-1}$  are right movers,
- $A_{i+1}, \dots, A_k$  are left movers,
- and each  $A_2, \dots, A_k$  can always execute.

and let  $P/A$  denote the program obtained from  $P$  by replacing  $A_1; \dots; A_k$  by  $\langle A_1; \dots; A_k \rangle$ .

Then  $P$  halts iff  $P/A$  halts and the final states of  $P$  equal the final states of  $P/A$ .

- Preservation of deadlock-freedom and partial correctness

# Application to example

Lipton's theorem justifies reduction to

```
integer  turn = 0;
boolean req0, req1 = false;

process P0
loop
  nc0: skip;
  rq0: ⟨req0 := true;
        turn := 1;⟩
  wt0: ⟨await ¬req1 ∨ turn = 0;
        skip;
        req0 := false;⟩
endloop

||

process P1
loop
  nc1: skip;
  rq1: ⟨req1 := true;
        turn := 0;⟩
  wt1: ⟨await ¬req0 ∨ turn = 1;
        skip;
        req1 := false;⟩
endloop
```

... but only for proving absence of deadlock

# Doeppner's reduction theorem

## Theorem

Let  $\Pi$  be a program and  $S$  have the form  $R; \langle A \rangle; L$  where

- all actions in  $R$  are right movers and
- all actions in  $L$  are left movers.

Let  $in(S)$  be true iff control resides inside  $S$  and  $Q$  be an arbitrary predicate.

Then  $Q$  is an invariant of  $\Pi/S$  iff  $Q \vee in(S)$  is an invariant of  $\Pi$ .

- Generalization of Lipton's theorem to invariant reasoning
- Can be used for proving mutual exclusion of example program

*T.W. Doeppner. Parallel program correctness through refinement. POPL 1977 (ACM), pp. 155-169.*




# Other reduction theorems

- R. Back: *Refining atomicity in parallel algorithms* (1988)
  - ▶ first reduction theorem for total correctness
  - ▶ needs commutativity hypotheses for actions outside reduced block
- L. Lamport, F. Schneider: *Pretending Atomicity* (1989)
  - ▶ generalization of Doeppner's theorem
  - ▶ preservation of invariants  $Q$  of  $\Pi$  by reduction  
(explicit reasoning about control being external to reduced block)
- E. Cohen, L. Lamport: *Reduction in TLA* (1998)
  - ▶ reformulation of Lamport & Schneider in TLA
  - ▶ extension to (certain) liveness properties

# Outline

- 1 Reduction Theorems for the Verification of Concurrent Programs
- 2 Fault-Tolerant Distributed Computing**
- 3 Reduction for Round-Based Distributed Algorithms
- 4 Experiments: Verification of Consensus Algorithms
- 5 Conclusion

# Fault-tolerant distributed algorithms



- local computation of nodes
- asynchronous communication over network
- components may fail: replication & fault-tolerance
- **precisely state and prove correctness properties**

# Representative problem: consensus

- $N$  nodes (processes) agree on a value
  - ▶ each node proposes a value initially
  - ▶ eventually nodes decide a common value
  - ▶ nodes or communication links may fail
- Formal definition: conjunction of four properties
  - integrity**      decided value is among the initial proposals
  - irrevocability**      decisions cannot be undone
  - agreement**      any two nodes decide same value
  - termination**      all (non-failed) nodes decide eventually
- Fundamental problem in fault-tolerant distributed computing

# Why is this hard?

## Theorem (Fischer, Lynch, Paterson 1985)

*The Consensus problem cannot be solved in an asynchronous system where at least one process may fail (by crashing).*

- But: many consensus algorithms exist (and work well in practice)

# Why is this hard?

## Theorem (Fischer, Lynch, Paterson 1985)

*The Consensus problem cannot be solved in an asynchronous system where at least one process may fail (by crashing).*


- But: many consensus algorithms exist (and work well in practice)
- Basis: relax some assumption of FLP theorem
  - ▶ introduce timeouts: being late is a failure
  - ▶ assume reliable (broadcast) communication
  - ▶ augment system by an oracle to detect failures
- Verification of consensus algorithms
  - ▶ difficult proofs ... often absent or informal
  - ▶ DiskPaxos: careful paper proof (30 pages for 0.5 page algorithm)
- **Can we help make verification simpler?**

# Heard-Of Model (Charron-Bost & Schiper, 2006)

- Algorithmic model for fault-tolerant distributed algorithms
  - ▶ uniform treatment of all (benign) errors
  - ▶ do not identify “culprit” or “type” of failure

# Heard-Of Model (Charron-Bost & Schiper, 2006)


- Algorithmic model for fault-tolerant distributed algorithms
  - ▶ uniform treatment of all (benign) errors
  - ▶ do not identify “culprit” or “type” of failure
- Round-based computation model





# Heard-Of Model (Charron-Bost & Schiper, 2006)

- Algorithmic model for fault-tolerant distributed algorithms
  - ▶ uniform treatment of all (benign) errors
  - ▶ do not identify “culprit” or “type” of failure
- Round-based computation model



- ▶ rounds: local structure of process computation
- ▶ state  $s'$  computed from  $s$  and received messages
- ▶ heard-of set  $HO(p, r)$ : processes from which messages are received
- ▶ **communication-closed rounds: discard late messages**

# Formal representation of HO algorithms

- Collection of processes  $(State_p, s_{0,p}, S_p^r, T_p^r)_{p \in Proc, r \in \mathbb{N}}$

- ▶ process states: sets  $State_p$  with initial states  $s_{0,p} \in State_p$
- ▶ message sending and state transition

$$S_p^r : State_p \times Proc \rightarrow Msg$$

$$T_p^r : State_p \times (Proc \rightarrow Msg) \rightarrow State_p$$

- ▶ domain of second argument of  $T_p^r$ : heard-of set  $HO(p, r)$

- For simplicity: deterministic processes

- ▶ algorithm behavior determined by collection of heard-of sets
- ▶ extension to non-deterministic processes straightforward

# Communication predicates

- Algorithms do not work in presence of arbitrary failures
  - ▶ safety: restrict number or extent of errors
  - ▶ liveness: assume eventual functioning of components

- Sample communication predicates

**non-split rounds**  $\forall p, q, r : HO(p, r) \cap HO(q, r) \neq \emptyset$

**$\leq f$  failures**  $\forall p, r : |HO(p, r)| \geq N - f$

**event. uniform**  $\exists r_0 \in \mathbb{N}, P \subseteq Proc : \forall r \geq r_0, q \in Proc : HO(q, r) = P$

- Observations (Charron-Bost & Schiper)

- ▶ standard failure assumptions can be expressed in terms of  $HO$  sets

# HO Consensus Algorithm: One-Third Rule

## Initialization

$x_p := v_p, decide_p := null$       ( $v_p$  : initial value of  $p$ )

**For each round**  $r \geq 0$

$S_p^r$  : send  $x_p$  to all processes

$T_p^r$  : **if**  $|HO(p, r)| > 2N/3$  **then**

    set  $x_p$  to smallest among the most frequently received values

**if** more than  $2N/3$  values received are equal to  $x_p$  **then**

$decide_p := x_p$

## Simple but efficient consensus algorithm

- no coordinator needed
- quick convergence if few errors

# Representing executions of HO algorithms

- Fine-grained execution for HO collection  $(HO(p, r))_{p \in Proc, r \in \mathbb{N}}$ 
  - ▶ message receptions, local transitions, message sending
  - ▶ verify correctness for all HO collections

```
process Node( $p \in Proc$ )  
  state  $st = s_{0,p}$ ;  
  integer  $r = 0$ ;  
  for  $q \in Proc$  do  $send(p, q, r, S_p^r(st, q))$  enddo;  
  loop  
    array  $rcvd = [q \in Proc \mapsto null]$ ;  
    for  $q \in HO(p, r)$  do  $rcvd[q] := receive(q, p, r)$  enddo;  
     $st, r := T_p^r(st, rcvd), r + 1$ ;  
    for  $q \in Proc$  do  $send(p, q, r, S_p^r(st, q))$  enddo;  
  end loop  
end process
```

- Formally: infinite sequence  $\xi = c_0 c_1 \dots$  of configurations

# Representing executions of HO algorithms

- Fine-grained execution for HO collection  $(HO(p, r))_{p \in Proc, r \in \mathbb{N}}$ 
  - ▶ message receptions, local transitions, message sending
  - ▶ verify correctness for all HO collections

**process** Node( $p \in Proc$ )

**state**  $st = s_{0,p}$ ;

**integer**  $r = 0$ ;

**for**  $q \in Proc$  **do**  $send(p, q, r, S_p^r(st, q))$  **enddo**;

**loop**

**array**  $rcvd = [q \in Proc \mapsto null]$ ;

**for**  $q \in HO(p, r)$  **do**  $rcvd[q] := receive(q, p, r)$  **enddo**;

$st, r := T_p^r(st, rcvd), r + 1$ ;

**for**  $q \in Proc$  **do**  $send(p, q, r, S_p^r(st, q))$  **enddo**;

**end loop**

**end process**

- Formally: infinite sequence  $\xi = c_0 c_1 \dots$  of configurations
- **Infinite-state model, due to round numbers**

# Outline

- 1 Reduction Theorems for the Verification of Concurrent Programs
- 2 Fault-Tolerant Distributed Computing
- 3 Reduction for Round-Based Distributed Algorithms**
- 4 Experiments: Verification of Consensus Algorithms
- 5 Conclusion

# First reduction

- Remember left and right movers?
  - ▶ send actions are left movers
  - ▶ receive actions are right movers

(assuming infinite  
network capacity)



# First reduction

- Remember left and right movers?
  - send actions are left movers
  - receive actions are right movers


(assuming infinite  
network capacity)

- This motivates the following reduction:

```
process Node( $p \in Proc$ )  
   $\langle$ state  $st = s_{0,p}$ ;  
  integer  $r = 0$ ;  
  for  $q \in Proc$  do  $send(p, q, r, S_p^r(st, q))$  enddo;  
  loop  
     $\langle$ array  $rcvd = [q \in Proc \mapsto null]$ ;  
    for  $q \in HO(p, r)$  do  $rcvd[q] := receive(q, p, r)$  enddo;  
     $st, r := T_p^r(st, rcvd), r + 1$ ;  
    for  $q \in Proc$  do  $send(p, q, r, S_p^r(st, q))$  enddo;  
  end loop  
end process
```

# More reduction

- Processes execute rounds atomically




- Can we do any better?

# More reduction

- Processes execute rounds atomically



- Can we do any better?
- Remember communication-closed rounds
  - round  $rnd_p^m$  right-commutes with  $rnd_q^n$  if  $m > n$
  - messages sent during  $rnd_q^n$  did not influence  $rnd_p^m$
- Rearrange execution so that executions of same round are adjacent




# More reduction

- Processes execute rounds atomically



- Can we do any better?
- Remember communication-closed rounds
  - round  $rnd_p^m$  right-commutes with  $rnd_q^n$  if  $m > n$
  - messages sent during  $rnd_q^n$  did not influence  $rnd_p^m$
- Rearrange execution so that executions of same round are adjacent



- Executions of same round by different processes are independent

# Coarse-grained model of executions

- Unit of atomicity: entire system rounds
  - ▶ all processes simultaneously perform transition for same round
  - ▶ corresponds to “nice” executions in the fine-grained model
- Coarse-grained execution  $\sigma_0\sigma_1\dots$  ( $\sigma_i : Proc \rightarrow State$ )
  - ▶  $\sigma_0(p) = s_{0,p}$
  - ▶  $\sigma_{r+1}(p) = T_p^r(\sigma_r(p), rcd(p, r))$   
where  $rcd(p, r) = [q \in HO(p, r) \mapsto S_q^r(\sigma_r(q), p)]$
- Coarse abstraction of distributed execution
  - ▶ no need for explicit representation of network
  - ▶ no round numbers: “synchronized” processes

# Coarse-grained model of executions

- Unit of atomicity: entire system rounds
  - ▶ all processes simultaneously perform transition for same round
  - ▶ corresponds to “nice” executions in the fine-grained model
- Coarse-grained execution  $\sigma_0\sigma_1\dots$  ( $\sigma_i : Proc \rightarrow State$ )
  - ▶  $\sigma_0(p) = s_{0,p}$
  - ▶  $\sigma_{r+1}(p) = T_p^r(\sigma_r(p), rcd(p, r))$   
where  $rcd(p, r) = [q \in HO(p, r) \mapsto S_q^r(\sigma_r(q), p)]$
- Coarse abstraction of distributed execution
  - ▶ no need for explicit representation of network
  - ▶ no round numbers: “synchronized” processes

⇒ How exactly does the reduced model relate to the original one?

# Relating fine- and coarse-grained executions

- Fine-grained model contains more detail
- Compare executions w.r.t. the “local views” of processes

- ▶  $p$ -view of fine-grained execution  $\zeta = c_0c_1 \dots$

$$\zeta^p = c_0.st(p), c_1.st(p), \dots$$

- ▶  $p$ -view of coarse-grained execution  $\sigma = \sigma_0\sigma_1 \dots$

$$\sigma^p = \sigma_0(p), \sigma_1(p), \dots$$

- ▶  $p$ -views are sequences of states of  $p$  and can be compared

- Executions equivalent iff indistinguishable by any process

$$\zeta \approx \sigma \quad \text{iff} \quad \mathfrak{h}(\zeta^p) = \mathfrak{h}(\sigma^p) \quad \text{for every } p \in Proc$$

- ▶ local views equal up to stuttering, for every process

# Reduction theorem

## Theorem (Reduction)

*Given a HO collection  $(HO(p, r))$  and a fine-grained execution  $\xi$  there exists a coarse-grained execution  $\sigma$  for the same HO collection such that  $\sigma \approx \xi$ .*

**Proof.** For  $\xi = c_0 c_1 \dots$ , define sequence  $\sigma = ([p \in Proc \mapsto c_{\ell_r^p}.st(p)])_{r \in \mathbb{N}}$

where  $\begin{cases} \ell_0^p &= 0 \\ \ell_{r+1}^p &= k+1 \text{ if } (c_k, c_{k+1}) \text{ is } (r+1)\text{st local transition of } p. \end{cases}$

Then  $\sigma$  is a coarse-grained execution for the same HO collection.

Moreover,  $\mathfrak{h}(\sigma^p) = \mathfrak{h}(\xi^p)$  for all  $p \in Proc$ .

Q.E.D.

- Converse theorem is trivially true



# “Local” properties

- Application of reduction theorem to verification
  - ▶ many properties depend only on local views
  - ▶ these can be verified by considering only coarse-grained executions
- Local properties  $P$  of executions

$$\rho_1 \models P \quad \text{iff} \quad \rho_2 \models P \quad \text{whenever} \quad \rho_1 \approx \rho_2$$

# “Local” properties

- Application of reduction theorem to verification
  - ▶ many properties depend only on local views
  - ▶ these can be verified by considering only coarse-grained executions
- Local properties  $P$  of executions

$$\rho_1 \models P \quad \text{iff} \quad \rho_2 \models P \quad \text{whenever} \quad \rho_1 \approx \rho_2$$

- The following LTL-X properties are local
  - ▶ formulas  $Q(p)$  built solely from  $p$ 's state variables
  - ▶ arbitrary first-order combinations of local properties
  - ▶ **but**: temporal combinations need not be local, consider:

$$\bigwedge_{p,q \in Proc} \square(rnd_p = rnd_q) \quad (\text{where } rnd_p \text{ is the current round of } p)$$

# Consensus as a local property

- Integrity

$$\bigwedge_{p \in Proc} \forall v \neq null : \left( \diamond(decide_p = v) \Rightarrow \bigvee_{q \in Proc} x_q = v \right)$$

- Irrevocability

$$\bigwedge_{p \in Proc} \forall v \neq null : \square(decide_p = v \Rightarrow \square(decide_p = v))$$

- Agreement

$$\bigwedge_{p, q \in Proc} \forall v, w \neq null : \diamond(decide_p = v) \wedge \diamond(decide_q = w) \Rightarrow v = w$$

- Termination

$$\bigwedge_{p \in Proc} \diamond(decide_p \neq null)$$

# Outline

- 1 Reduction Theorems for the Verification of Concurrent Programs
- 2 Fault-Tolerant Distributed Computing
- 3 Reduction for Round-Based Distributed Algorithms
- 4 Experiments: Verification of Consensus Algorithms**
- 5 Conclusion

# Finite-state model checking

- Verification of finite instances of algorithms
  - ▶ model coarse-grained executions for fixed number of processes
  - ▶ non-deterministic choice of HO sets at every transition
  - ▶ resulting model is finite-state
- Generic TLA<sup>+</sup> module *HeardOf*
  - ▶ high-level definition of coarse-grained HO semantics
  - ▶ pre-define useful communication predicates
  - ▶ concrete algorithms obtained later as instances
- Here: favor clarity over efficiency

# Generic TLA<sup>+</sup> module

MODULE *HeardOf*

EXTENDS *Naturals*

CONSTANTS *Proc, State, Msg, nPhases, IniSt(-), Send(-, -, -, -), Trans(-, -, -, -)*

VARIABLES *phase, state, heardof*

*Init*  $\triangleq \wedge \text{phase} = 0$

$\wedge \text{state} = [p \in \text{Proc} \mapsto \text{IniSt}(p)]$

$\wedge \text{heardof} = [p \in \text{Proc} \mapsto \{\}]$

*Step*(*HO*)  $\triangleq$  LET *rcvd*(*p*)  $\triangleq \{\langle q, \text{Send}(q, \text{phase}, \text{state}[q], p) \rangle : q \in \text{HO}[p]\}$

IN  $\wedge \text{phase}' = (\text{phase} + 1) \% n\text{Phases}$

$\wedge \text{state}' = [p \in \text{Proc} \mapsto \text{Trans}(p, \text{phase}, \text{state}[p], \text{rcvd}(p))]$

$\wedge \text{heardof}' = \text{HO}$

*Next*  $\triangleq \exists \text{HO} \in [\text{Proc} \rightarrow \text{SUBSET Proc}] : \text{Step}(\text{HO})$

*NoSplit*(*HO*)  $\triangleq \forall p, q \in \text{Proc} : \text{HO}[p] \cap \text{HO}[q] \neq \{\}$

*NextNoSplit*  $\triangleq \exists \text{HO} \in [\text{Proc} \rightarrow \text{SUBSET Proc}] : \text{NoSplit}(\text{HO}) \wedge \text{Step}(\text{HO})$

*Uniform*(*HO*)  $\triangleq \exists S \in \text{SUBSET Proc} : \text{HO} = [q \in \text{Proc} \mapsto S]$

*InfiniteUniform*  $\triangleq \square \diamond \text{Uniform}(\text{heardof})$

- Definitions closely parallels “paper” version
  - ▶ expressiveness of TLA<sup>+</sup> leads to perspicuous formulation
  - ▶ (auxiliary) variable *heardof* records HO sets during a run
  - ▶ mainly used for debugging and printing counter-examples
- Formulation of communication predicates
  - ▶ safety predicates: add to next-state relation
  - ▶ liveness predicates: natural expression in temporal logic
  - ▶ used to express correctness properties

# One-Third Rule in TLA<sup>+</sup> (1/3)

## MODULE *OneThirdRule*

EXTENDS *Naturals, FiniteSets*

CONSTANT *N*

VARIABLES *phase, state, heardof*

<i>nPhases</i>	$\triangleq$	1
<i>Proc</i>	$\triangleq$	1.. <i>N</i>
<i>InitValue(p)</i>	$\triangleq$	10 * <i>p</i>
<i>Value</i>	$\triangleq$	{ <i>InitValue(p) : p</i> ∈ <i>Proc</i> }
<i>Msg</i>	$\triangleq$	<i>Value</i>
<i>null</i>	$\triangleq$	0
<i>ValueOrNull</i>	$\triangleq$	<i>Value</i> ∪ { <i>null</i> }
<i>State</i>	$\triangleq$	[ <i>x</i> : <i>Value, decide</i> : <i>ValueOrNull</i> ]

- definition of constant parameters for *OneThirdRule* algorithm
- arbitrary definition of (initial) values of a process



# One-Third Rule in TLA<sup>+</sup> (2/3)

```
IniSt(p)            $\triangleq$  [x  $\mapsto$  InitValue(p), decide  $\mapsto$  null]
Send(p, ph, s, q)   $\triangleq$  s.x
Trans(p, ph, s, rcvd)  $\triangleq$ 
  IF Cardinality(rcvd) > (2 * N)  $\div$  3
  THEN LET Freq(v)  $\triangleq$  Cardinality({q  $\in$  Proc : <q, v>  $\in$  rcvd})
        MFR(v)  $\triangleq$   $\forall w \in$  Value : Freq(w)  $\leq$  Freq(v)
        min  $\triangleq$  CHOOSE v  $\in$  Value : MFR(v)  $\wedge$   $\forall w \in$  Value : MFR(w)  $\Rightarrow$  v  $\leq$  w
        IN [x  $\mapsto$  min,
            decide  $\mapsto$  IF Freq(min) > (2 * N)  $\div$  3 THEN min ELSE s.decide]
  ELSE s
INSTANCE HeardOf
```

- definition of the send and state transition functions
- instantiation of generic module

# One-Third Rule in TLA<sup>+</sup> (3/3)

<i>Safety</i>	$\triangleq$	$Init \wedge \square[Next]_{vars}$
<i>Liveness</i>	$\triangleq$	$\square\Diamond(\text{Uniform}(\text{heardof}) \wedge \text{Cardinality}(\text{heardof}) > (2 * N) \div 3)$
<i>Integrity</i>	$\triangleq$	$\forall p \in Proc : \text{state}[p].\text{decide} \in \text{ValueOrNull}$
<i>Irrevocability</i>	$\triangleq$	$\forall p \in Proc : \square[\text{state}[p].\text{decide} = \text{null}]_{\text{state}[p].\text{decide}}$
<i>Agreement</i>	$\triangleq$	$\forall p, q \in Proc : (\text{state}[p].\text{decide} \neq \text{null} \wedge \text{state}[q].\text{decide} \neq \text{null} \\ \Rightarrow \text{state}[p].\text{decide} = \text{state}[q].\text{decide})$
<i>Termination</i>	$\triangleq$	$\forall p \in Proc : \Diamond(\text{state}[p].\text{decide} \neq \text{null})$

THEOREM  $\text{Safety} \Rightarrow \square(\text{Integrity} \wedge \text{Agreement}) \wedge \text{Irrevocability}$

THEOREM  $\text{Safety} \wedge \text{Liveness} \Rightarrow \text{Termination}$

- definition of correctness properties
- formulation of correctness theorems, under precise hypotheses

# Results of verification

	OneThirdRule		UniformVoting	
	$N = 3$	$N = 4$	$N = 3$	$N = 4$
states	5633	9,830,401	21,351	15,865,770
distinct	11	150	122	887
time (s)	1.87	939	13.8	1330

- Model checking feasible for small instances

- ▶ high branching factor: exploration of all HO collections
- ▶ many redundant states generated

- Symbolic model checking can be more efficient

- ▶ more complicated encodings necessary for tools like NuSMV
- ▶ cf. work by Tsuchiya and Schiper: Paxos for 10 processes

# Verification in Isabelle/HOL

## Similar overall model

- main difference: introduction of types
- generic *HeardOf* module represented as an Isabelle locale

```
locale HOAlgorithm =  
fixes  
  nPhases :: nat and  
  iniSt :: 'proc → 'pst and  
  send :: 'proc → nat → 'pst → 'proc → 'msg and  
  trans :: 'proc → nat → 'pst → ('proc → 'msg) → 'pst  
assumes  
  nSteps : 0 < nPhases and  
  finiteProc : finite(UNIV :: 'procset)
```

- defines generic behavior of HO algorithms
- proves useful rules, such as induction over executions

# Proof of correctness

- Validity: standard invariance proof
- Irrevocability and agreement via sequence of lemmas
  - 1 if process decides on value  $v$  then more than  $2N/3$  processes contain  $v$  in their  $x$  field
  - 2 if more than  $2N/3$  processes send  $v$  and process  $p$  hears from more than  $2N/3$  processes then  $p$  updates its  $x$  field to  $v$
  - 3 whenever process has decided on  $v$  then more than  $2N/3$  processes contain  $v$  in their  $x$  field
  - 4 hence, processes cannot decide on different values
- Liveness: symbolically execute uniform rounds

# Proof of correctness

- Validity: standard invariance proof
- Irrevocability and agreement via sequence of lemmas
  - 1 if process decides on value  $v$  then more than  $2N/3$  processes contain  $v$  in their  $x$  field
  - 2 if more than  $2N/3$  processes send  $v$  and process  $p$  hears from more than  $2N/3$  processes then  $p$  updates its  $x$  field to  $v$
  - 3 whenever process has decided on  $v$  then more than  $2N/3$  processes contain  $v$  in their  $x$  field
  - 4 hence, processes cannot decide on different values
- Liveness: symbolically execute uniform rounds
- Proof lengths in Isar (including model and explanations)
  - ▶ 8 pages for generic module and lemmas
  - ▶ 8 pages for *OneThirdRule*
  - ▶ 25 pages for *LastVoting* (cf. 130 pages for fine-grained model!)

# Outline

- 1 Reduction Theorems for the Verification of Concurrent Programs
- 2 Fault-Tolerant Distributed Computing
- 3 Reduction for Round-Based Distributed Algorithms
- 4 Experiments: Verification of Consensus Algorithms
- 5 Conclusion**

# Reduction: a revival?

- Recast of classical theorems
  - ▶ identify left and right movers for coarser unit of atomicity
  - ▶ distributed algorithms present interesting opportunities
  - ▶ substantial reduction of verification effort possible
- Transcend historical formulations
  - ▶ beyond programming-language based presentations
  - ▶ wide interpretation of “processes” (e.g., set of rounds)
  - ▶ verify safety *and* liveness properties
- Ongoing / future work
  - ▶ establish more general reduction theorems
  - ▶ better syntactic characterization of local properties
  - ▶ implementation of reduction in verification tools