


TLA⁺ Video Course – Lecture 7

Leslie Lamport

PAXOS COMMIT

This video should be viewed in conjunction with a Web page.
To find that page, search the Web for *TLA+ Video Course*.

The TLA⁺ Video Course
Lecture 7
Paxos Commit



In this lecture, we study a specification of Paxos Commit – a fault-tolerant distributed algorithm that implements transaction commit. The spec illustrates most of the TLA+ constructs you don't already know that you will use in writing specs.

I hope you'll also study the algorithm itself. I think it's neat, but then I'm prejudiced, since Jim Gray and I invented it. But that's up to you. These lectures are about TLA+, not distributed algorithms.

THE ALGORITHM

The Paxos Commit algorithm.

The problem with two-phase commit:

There's an obvious problem with two-phase commit:

The problem with two-phase commit:

It can hang forever if the TM fails.

There's an obvious problem with two-phase commit: It can hang forever if the transaction manager fails.

The problem with two-phase commit:

It can hang forever if the TM fails.

A simple engineering solution:

There's an obvious problem with two-phase commit: It can hang forever if the transaction manager fails.

There's a simple engineering solution.

The problem with two-phase commit:

It can hang forever if the TM fails.

A simple engineering solution:

Have a backup TM take over if the TM fails.

There's an obvious problem with two-phase commit: It can hang forever if the transaction manager fails.

There's a simple engineering solution.

Have a backup transaction manager take over if the primary transaction manager fails.

The problem with two-phase commit:

It can hang forever if the TM fails.

A simple engineering solution:

Have a backup TM take over if the TM fails.

You can find it in textbooks.

You can find this solution in database textbooks.

The problem with two-phase commit:

It can hang forever if the TM fails.

A simple engineering solution:

Have a backup TM take over if the TM fails.

You can find it in textbooks.

It's straightforward to implement

You can find this solution in database textbooks.

It's straightforward to implement

The problem with two-phase commit:

It can hang forever if the TM fails.

A simple engineering solution:

Have a backup TM take over if the TM fails.

You can find it in textbooks.

**It's straightforward to implement
and test that it works.**

You can find this solution in database textbooks.

It's straightforward to implement and to test that it works.

It's deployed and works fine

The system is deployed and works fine, and everyone's happy

It's deployed and works fine until one day:

The system is deployed and works fine, and everyone's happy until one day:

It's deployed and works fine until one day:

The primary TM decides to commit

The system is deployed and works fine, and everyone's happy until one day:

The primary transaction manager decides to commit

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The system is deployed and works fine, and everyone's happy until one day:

The primary transaction manager decides to commit and then pauses for some reason.

Perhaps it's pre-empted by a higher priority task.

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The backup TM thinks the primary failed and it decides to take over.

The backup transaction manager thinks the primary failed and it decides to take over.

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The backup TM thinks the primary failed and it decides to take over.

The backup TM broadcasts an *Abort* message.

The backup transaction manager thinks the primary failed and it decides to take over.

The backup transaction manager broadcasts an *Abort* message.

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The backup TM thinks the primary failed and it decides to take over.

The backup TM broadcasts an *Abort* message.

The primary TM resumes and broadcasts a *Commit* message.

The backup transaction manager thinks the primary failed and it decides to take over.

The backup transaction manager broadcasts an *Abort* message.

Meanwhile, the primary transaction manager resumes and broadcasts a *Commit* message.

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The backup TM thinks the primary failed and it decides to take over.

The backup TM broadcasts an *Abort* message.

The primary TM resumes and broadcasts a *Commit* message.

Some RMs abort and others commit.

The backup transaction manager thinks the primary failed and it decides to take over.

The backup transaction manager broadcasts an *Abort* message.

Meanwhile, the primary transaction manager resumes and broadcasts a *Commit* message.

This causes some resource managers to abort and others to commit.

It's deployed and works fine until one day:

The primary TM decides to commit and then pauses.

The backup TM thinks the primary failed and it decides to take over.

SYSTEM FAILURE

The backup TM broadcasts an *Abort* message.

The primary TM resumes and broadcasts a *Commit* message.

Some RMs abort and others commit.

Which constitutes a system failure.

Finding fault-tolerant distributed algorithms is hard.

Finding fault-tolerant distributed algorithms is hard.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong, and hard to find errors by testing.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong , and it's hard to find that they're wrong by testing.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong, and hard to find errors by testing.

We should get the algorithm right before we code.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong , and it's hard to find that they're wrong by testing.

It's important to get the algorithm right before we code it.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong, and hard to find errors by testing.

We should get the algorithm right before we code.

Writing and checking a TLA⁺ spec is the best way I know to do that.

Finding fault-tolerant distributed algorithms is hard.

They're easy to get wrong , and it's hard to find that they're wrong by testing.

It's important to get the algorithm right before we code it.

Writing and checking a TLA⁺ spec is the best way I know to do that.

Paxos Commit is a fault-tolerant transaction-commit algorithm described in this paper:

Consensus on Transaction Commit

Jim Gray and Leslie Lamport

ACM Transactions on Database Systems (TODS)

Volume 31, issue 1 (March 2006), pages 133–160

Paxos Commit is a fault-tolerant transaction-commit algorithm described in this 2006 paper by Jim Gray and me.

Paxos Commit is a fault-tolerant transaction-commit algorithm described in this paper:

Consensus on Transaction Commit

Jim Gray and Leslie Lamport

ACM Transactions on Database Systems (TODS)

Volume 31, issue 1 (March 2006), pages 133–160

The paper explains the algorithm and specifies it in module *PaxosCommit*.

Paxos Commit is a fault-tolerant transaction-commit algorithm described in this 2006 paper by Jim Gray and me.

The paper explains the algorithm and specifies it in a TLA⁺ module named *PaxosCommit*.

We're looking at this module for two reasons:

We're looking at this module for two reasons:

We're looking at this module for two reasons:

- To see what a real spec looks like.

We're looking at this module for two reasons:

The first is to see what a real spec looks like.

We're looking at this module for two reasons:

- To see what a real spec looks like.
- To learn some more TLA⁺.

We're looking at this module for two reasons:

The first is to see what a real spec looks like.

The second is to learn some more TLA⁺.

We're looking at this module for two reasons:

- To see what a real spec looks like.
- To learn some more TLA⁺.

You can read the paper if you want to understand the algorithm.

We're looking at this module for two reasons:

The first is to see what a real spec looks like.

The second is to learn some more TLA⁺.

You should read the paper if you want to understand the algorithm.

We're looking at this module for two reasons:

- To see what a real spec looks like.
- To learn some more TLA⁺.

You can read the paper if you want to understand the algorithm.

This lecture explains only the TLA⁺ operators you haven't seen yet that are used in the spec.

This lecture explains only the TLA⁺ operators you haven't seen yet that are used in the spec.

Stop the video now and:

Download module *PaxosCommit* to the same folder as *TCommit*.

Download the paper.

Stop the video now and download module *PaxosCommit* to the same folder as module *TCommit*; and download the paper if you want to read it.

Stop the video now and:

Download module *PaxosCommit* to the same folder as *TCommit*.

Download the paper.

Modules *TCommit*, *TwoPhase*, *PaxosCommit* used in these lectures differ slightly from the ones in the paper.

Stop the video now and download module *PaxosCommit* to the same folder as module *TCommit*; and download the paper if you want to read it.

The module *PaxosCommit* that we use here, as well as modules *TCommit* and *TwoPhase* used in previous lectures, differ slightly from the ones in the paper.

THE SPECIFICATION

The Paxos Commit algorithm's specification

EXTENDS *Integers*

The module begins with an EXTENDS statement that imports the definition of arithmetic operators from the standard *Integers* module.

$Maximum(S) \triangleq$

The module then defines $Maximum(S)$

$Maximum(S) \triangleq$

The largest element of the finite set S of integers

The module then defines $Maximum(S)$ to be the largest element of S if S is a finite set of integers,

$Maximum(S) \triangleq$

The largest element of the finite set S of integers,
or -1 if S is the empty set.

The module then defines $Maximum(S)$ to be the largest element of S if S is a finite set of integers, and to equal -1 if it's the empty set.

We don't care what it equals if S is infinite or not a set of numbers.

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE **smallest number in S**

The definition has this form

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE **smallest number in S**

The definition has this form

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE **smallest number in S**

The definition has this form

The smallest number in S is written this way

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE CHOOSE $n \in S : n \geq$ every element in S

where the CHOOSE expression

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE CHOOSE $n \in S$: $n \geq$ every element in S

Equals an arbitrarily chosen n in S

where the CHOOSE expression equals an arbitrarily chosen value n in S

$Maximum(S) \triangleq$

IF $S = \{\}$ THEN -1

ELSE CHOOSE $n \in S$: $n \geq$ every element in S

Equals an arbitrarily chosen n in S satisfying ...

where the CHOOSE expression equals an arbitrarily chosen value n in S satisfying the condition that n is greater-than or equal to every element in S .
If n is finite and nonempty, then there is exactly one such n .

$Maximum(S) \triangleq$

IF $S = \{\}$ THEN -1

ELSE CHOOSE $n \in S : \forall m \in S : n \geq m$

where the CHOOSE expression equals an arbitrarily chosen value n in S satisfying the condition that n is greater-than or equal to every element in S . If n is finite and nonempty, then there is exactly one such n .

That condition on n is written this way.

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE CHOOSE $n \in S : (\forall m \in S : n \geq m)$

It's a little easier to read with parentheses.

$Maximum(S) \triangleq$
IF $S = \{\}$ THEN -1
ELSE CHOOSE $n \in S : (\forall m \in S : n \geq m)$

It's a little easier to read with parentheses.

This formula states that for every m in S , n is greater than or equal to m .

CHOOSE $v \in S : P$ equals

In general, the expression CHOOSE variable v in expression S colon formula P equals

CHOOSE $v \in S : P$ equals

if there is a v in S for which P is true

In general, the expression **CHOOSE variable v in expression S colon formula P equals**

If there is at least one value v in the set S for which formula P is true

CHOOSE $v \in S : P$ equals

if there is a v in S for which P is true
then some such v

In general, the expression CHOOSE variable v in expression S colon formula P equals

If there is at least one value v in the set S for which formula P is true
then the expression equals some such v .

If there's more than one, then the semantics of TLA+ don't specify which one.

CHOOSE $v \in S : P$ equals

if there is a v in S for which P is true

then some such v

else a completely unspecified value.

Else, If there is no such v , then the value of the CHOOSE expression is completely unspecified.

And TLC will report an error if that's the case when it tries to evaluate the expression.

CHOOSE $v \in S : P$ equals

if there is a v in S for which P is true

then some such v

else a completely unspecified value.

CHOOSE $i \in 1..99 : \text{TRUE}$

Is an unspecified integer between 1 and 99.

Else, If there is no such v , then the value of the CHOOSE expression is completely unspecified.

And TLC will report an error if that's the case when it tries to evaluate the expression.

For example: this expression equals an unspecified integer between 1 and 99. We don't know which one.

CHOOSE $v \in S : P$ equals

if there is a v in S for which P is true

then some such v

else a completely unspecified value.

CHOOSE $i \in 1..99 : \text{TRUE}$

Is an unspecified integer between 1 and 99.

It might or might not equal 37.

It might equal 37, or it might not; the semantics of TLA+ don't say.

In math, any expression equals itself.

In math, any expression always equals itself.

In math, any expression equals itself.

`(CHOOSE $i \in 1..99$: TRUE) = (CHOOSE $i \in 1..99$: TRUE)`

In math, any expression always equals itself.

So this CHOOSE expression always equals itself

In math, any expression equals itself.

$(\text{CHOOSE } i \in 1..99 : \text{TRUE}) = (\text{CHOOSE } i \in 1..99 : \text{TRUE})$

There is no nondeterminism in a mathematical expression.

In math, any expression always equals itself.

So this CHOOSE expression always equals itself

There is no nondeterminism in any mathematical expression, including a CHOOSE expression.

In math, any expression equals itself.

$(\text{CHOOSE } i \in 1..99 : \text{TRUE}) = (\text{CHOOSE } i \in 1..99 : \text{TRUE})$

There is no nondeterminism in a mathematical expression.

If $\text{CHOOSE } i \in 1..99 : \text{TRUE}$ equals 37 today;
it will equal 37 next week.

If this CHOOSE expression equals 37 today, it will still equal 37 next week.

In math, any expression equals itself.

$(\text{CHOOSE } i \in 1..99 : \text{TRUE}) = (\text{CHOOSE } i \in 1..99 : \text{TRUE})$

There is no nondeterminism in a mathematical expression.

If $\text{CHOOSE } i \in 1..99 : \text{TRUE}$ equals 37 today;
it will equal 37 next week.

TLC will always get the same number when it evaluates it.

If this CHOOSE expression equals 37 today, it will still equal 37 next week.

TLC will always get the same number when it evaluates this expression.

You shouldn't care what number.

$$x' \in 1 \dots 99$$

Allows the value of x in the next state to be any number in $1 \dots 99$.

The formula x prime in the set $1 \dots 99$ allows the value of x in the next state to be any of the 99 numbers from 1 to 99.

$$x' \in 1 \dots 99$$

Allows the value of x in the next state to be any number in $1 \dots 99$.

$$x' = \text{CHOOSE } i \in 1 \dots 99 : \text{TRUE}$$

Allows the value of x in the next state to be one particular number.

The formula x' in the set $1 \dots 99$ allows the value of x in the next state to be any of the 99 numbers from 1 to 99.

The formula x' equals this CHOOSE expression allows the value of x in the next state to be some particular number between 1 and 99 — perhaps 37.

$$x' \in 1 \dots 99$$

Allows the value of x in the next state to be any number in $1 \dots 99$.

~~$$x' = \text{CHOOSE } i \in 1 \dots 99 : \text{TRUE}$$~~

Allows the value of x in the next state to be one particular number.

The formula x' in the set $1 \dots 99$ allows the value of x in the next state to be any of the 99 numbers from 1 to 99.

The formula x' equals this CHOOSE expression allows the value of x in the next state to be some particular number between 1 and 99 — perhaps 37.

There's no reason why you'd ever want to write something like this.

You should write `CHOOSE $v \in S : P$`

Only when there's exactly one v in S satisfying P .

You should write this `CHOOSE` expression only when there's exactly one value v in S satisfying formula P .

You should write $\text{CHOOSE } v \in S : P$

Only when there's exactly one v in S satisfying P .

As in the definition of $\text{Maximum}(S)$.

You should write this CHOOSE expression only when there's exactly one value v in S satisfying formula P .

For example, the way it was used in the definition of Maximum of S .

You should write $\text{CHOOSE } v \in S : P$

Only when there's exactly one v in S satisfying P .

Or when it's part of a larger expression whose value doesn't depend on which v is chosen.

You should write this CHOOSE expression only when there's exactly one value v in S satisfying formula P .

For example, the way it was used in the definition of *Maximum* of S .

Or when it's part of a larger expression whose value doesn't depend on which possible value of v is chosen.

We'll see an example of that later.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

After defining *Maximum*, the module contains a CONSTANTS statement declaring these four constants

CONSTANTS RM , *Acceptor*, *Majority*, *Ballot*
the set of resource managers

After defining *Maximum*, the module contains a CONSTANTS statement declaring these four constants

RM is again the set of resource managers

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*
the set of acceptor processes

After defining *Maximum*, the module contains a CONSTANTS statement declaring these four constants

RM is again the set of resource managers and *Acceptor* is another a set of processes called acceptors.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

After defining *Maximum*, the module contains a CONSTANTS statement declaring these four constants

RM is again the set of resource managers and *Acceptor* is another a set of processes called acceptors.

The constants *Majority* and *Ballot* are sets described in the following statement.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

This ASSUME statement asserts assumptions being made about the constants.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \textit{Ballot} \subseteq \textit{Nat}$

$\wedge 0 \in \textit{Ballot}$

$\wedge \textit{Majority} \subseteq \text{SUBSET } \textit{Acceptor}$

$\wedge \forall \textit{MS1}, \textit{MS2} \in \textit{Majority} : \textit{MS1} \cap \textit{MS2} \neq \{\}$

This ASSUME statement asserts assumptions being made about the constants.

CONSTANTS RM , $Acceptor$, $Majority$, $Ballot$

ASSUME

$\wedge Ballot \subseteq Nat$

$\wedge 0 \in Ballot$

$\wedge Majority \subseteq SUBSET Acceptor$

$\wedge \forall MS1, MS2 \in Majority : MS1 \cap MS2 \neq \{\}$

This ASSUME statement asserts assumptions being made about the constants.

For example, the second conjunct asserts the assumption that zero is an element of the set $Ballot$.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \textit{Ballot} \subseteq \textit{Nat}$

$\wedge 0 \in \textit{Ballot}$

$\wedge \textit{Majority} \subseteq \text{SUBSET } \textit{Acceptor}$

$\wedge \forall \textit{MS1}, \textit{MS2} \in \textit{Majority} : \textit{MS1} \cap \textit{MS2} \neq \{\}$

These assumptions use some TLA+ notation that you haven't seen yet.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \textit{Ballot} \subseteq \boxed{\textit{Nat}}$

the set of natural numbers

These assumptions use some TLA+ notation that you haven't seen yet.

Nat is defined in the imported *Integers* module to be the set of natural numbers (that is, the non-negative integers).

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \textit{Ballot} \subseteq \textit{Nat}$

Ballot is a subset of *Nat*.

These assumptions use some TLA+ notation that you haven't seen yet.

Nat is defined in the imported *Integers* module to be the set of natural numbers (that is, the non-negative integers).

The first conjunct asserts that *Ballot* is a subset of *Nat*, meaning that every element of *Ballot* is an element of the set *Nat* of natural numbers.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \textit{Ballot} \subseteq \textit{Nat}$
`\subsetq`

The subset symbol is typed `backslash subset e-q`.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \text{Majority} \subseteq \text{SUBSET } \text{Acceptor}$
the set of all subsets of *Acceptor*

The subset symbol is typed `backslash subset e-q`.

`SUBSET Acceptor` is the set of all subsets of the set *Acceptor*.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \text{Majority} \subseteq \text{SUBSET } \textit{Acceptor}$
the set of all subsets of *Acceptor*

Also called the *powerset of Acceptor*, written $\mathcal{P}(\textit{Acceptor})$

The subset symbol is typed `backslash subset e-q`.

`SUBSET Acceptor` is the set of all subsets of the set *Acceptor*.

Mathematicians call it the powerset of *Acceptor* and write it *P* of *Acceptor*.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

\wedge *Majority* \subseteq SUBSET *Acceptor*

The elements of *Majority* are subsets of *Acceptor*.

The subset symbol is typed `backslash subset e-q`.

SUBSET *Acceptor* is the set of all subsets of the set *Acceptor*.

Mathematicians call it the powerset of *Acceptor* and write it *P* of *Acceptor*.

The conjunct asserts the assumption that every element of *Majority* is a subset of the set *Acceptor*.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \forall MS1, MS2 \in Majority : MS1 \cap MS2 \neq \{\}$
intersection of *MS1* and *MS2*

This subexpression is the intersection of the sets *MS1* and *MS2*.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \forall MS1, MS2 \in Majority : MS1 \cap MS2 \neq \{\}$

the set of elements in both *MS1* and *MS2*

This subexpression is the intersection of the sets *MS1* and *MS2*.

It's the set consisting of all elements in both *MS1* and *MS2*.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$$\wedge \forall MS1, MS2 \in \textit{Majority} : MS1 \cap MS2 \neq \{\}$$

`\intersect`
`\cap`

This subexpression is the intersection of the sets *MS1* and *MS2*.

It's the set consisting of all elements in both *MS1* and *MS2*.

The intersection symbol is typed either `backslash intersect` or `backslash cap`.

CONSTANTS *RM*, *Acceptor*, *Majority*, *Ballot*

ASSUME

$\wedge \forall MS1, MS2 \in Majority : MS1 \cap MS2 \neq \{\}$

Any two elements of *Majority* have an element in common.

This subexpression is the intersection of the sets *MS1* and *MS2*.

It's the set consisting of all elements in both *MS1* and *MS2*.

The intersection symbol is typed either backslash intersect or backslash cap.

The conjunct asserts that every two elements of the set *Majority* are sets having at least one element in common.

CONSTANTS RM , $Acceptor$, $Majority$, $Ballot$

ASSUME

$\wedge Ballot \subseteq Nat$

$\wedge 0 \in Ballot$

$\wedge Majority \subseteq SUBSET Acceptor$

$\wedge \forall MS1, MS2 \in Majority : MS1 \cap MS2 \neq \{\}$

TLC will check these assumptions.

TLC will check all these assumptions.

$Messages \triangleq$

$[type : \{ "phase1a" \}, ins : RM, bal : Ballot \setminus \{0\}]$

\cup

$[type : \{ "phase1b" \}, ins : RM, mbal : Ballot, bal : Ballot \cup \{-1\},$
 $val : \{ "prepared", "aborted", "none" \}, acc : Acceptor]$

\cup

$[type : \{ "phase2a" \}, ins : RM, bal : Ballot, val : \{ "prepared", "aborted" \}]$

\cup

$[type : \{ "phase2b" \}, acc : Acceptor, ins : RM, bal : Ballot,$
 $val : \{ "prepared", "aborted" \}]$

\cup

$[type : \{ "Commit", "Abort" \}]$

The module next defines *Messages* to be a set consisting of several kinds of records.

$Messages \triangleq$

$[type : \{ "phase1a" \}, ins : RM, bal : Ballot \setminus \{0\}]$

\cup

$[type : \{ "phase1b" \}, ins : RM, mbal : Ballot, bal : Ballot \cup \{-1\},$
 $val : \{ "prepared", "aborted", "none" \}, acc : Acceptor]$

\cup

$[type : \{ "phase2a" \}, ins : RM, bal : Ballot, val : \{ "prepared", "aborted" \}]$

\cup

$[type : \{ "phase2b" \}, acc : Acceptor, ins : RM, bal : Ballot,$
 $val : \{ "prepared", "aborted" \}]$

\cup

$[type : \{ "Commit", "Abort" \}]$

The module next defines *Messages* to be a set consisting of several kinds of records. The definition contains this expression.

$Ballot \setminus \{0\}$

The module next defines *Messages* to be a set consisting of several kinds of records. The definition contains this expression.

$Ballot \setminus \{0\}$

The module next defines *Messages* to be a set consisting of several kinds of records. The definition contains this expression.

This *set minus* operator is defined as follows. For any sets S and T ,

$Ballot \setminus \{0\}$

$S \setminus T$ is the set of elements in S not in T .

The module next defines *Messages* to be a set consisting of several kinds of records. The definition contains this expression.

This *set minus* operator is defined as follows. For any sets S and T , S set-minus T is the set of all elements in S that are not in T .

$Ballot \setminus \{0\}$

$S \setminus T$ is the set of elements in S not in T .

$$(10..20) \setminus (1..14) = 15..20$$

For example, the integers from 10 to 20 set-minus the integers from 1 to 14 equals the set of integers from 15 to 20.

$Ballot \setminus \{0\}$

The set of non-0 elements in *Ballot*.

For example, the integers from 10 to 20 set-minus the integers from 1 to 14 equals the set of integers from 15 to 20.

So, *Ballot* set-minus the set containing only 0 is the set of non-zero elements in *Ballot*.

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

The module next declares its variables and defines the type-correctness invariant $PCTypeOK$.

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

The module next declares its variables and defines the type-correctness invariant $PCTypeOK$.

As in the two-phase commit spec, there is a variable $m-s-g-s$ whose value is a set of messages.

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

$PCTypeOK$ also asserts that the value of the variable $aState$

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

$PCTypeOK$ also asserts that the value of the variable $aState$ is a function with domain RM

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow \boxed{Acceptor} \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

$r \in RM$ implies $aState[r]$ is a function with domain $Acceptor$.

$PCTypeOK$ also asserts that the value of the variable $aState$ is a function with domain RM such that for every r in RM , $aState[r]$ is a function with domain $Acceptor$

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

$r \in RM$ implies $aState[r]$ is a function with domain $Acceptor$.

$a \in Acceptor$ implies $aState[r][a]$ is a record with three fields.

$PCTypeOK$ also asserts that the value of the variable $aState$ is a function with domain RM such that for every r in RM , $aState[r]$ is a function with domain $Acceptor$ such that for every a in the set $Acceptor$, $aState[r][a]$ is a record these three fields

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

$r \in RM$ implies $aState[r]$ is a function with domain $Acceptor$.

$a \in Acceptor$ implies $aState[r][a]$ is a record with three fields.

$aState[r][a].bal$ is in $Ballot$ or equals -1 .

$PCTypeOK$ also asserts that the value of the variable $aState$ is a function with domain RM such that for every r in RM , $aState[r]$ is a function with domain $Acceptor$ such that for every a in the set $Acceptor$, $aState[r][a]$ is a record these three fields

And, for example, $aState[r][a].bal$ is in the set $Ballot$ or equals -1 .

VARIABLES $rmState$, $aState$, $msgs$

$PCTypeOK \triangleq$

$\wedge rmState \in [RM \rightarrow \{\text{"working"}, \text{"prepared"}, \text{"committed"}, \text{"aborted"}\}]$

$\wedge aState \in [RM \rightarrow [Acceptor \rightarrow [mbal : Ballot,$
 $bal : Ballot \cup \{-1\},$
 $val : \{\text{"prepared"}, \text{"aborted"}, \text{"none"}\}]]]$

$\wedge msgs \subseteq Messages$

There's nothing new here.

There's nothing new here; it's just a little more complicated than the formulas you've seen so far.

That's true for what follows in the module, up until

```

Phase2a(bal, r)  $\triangleq$ 
   $\wedge \neg \exists m \in msgs : \wedge m.type = \text{"phase2a"}$ 
     $\wedge m.bal = bal$ 
     $\wedge m.ins = r$ 
   $\wedge \exists MS \in Majority :$ 
    LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"}$ 
       $\wedge m.ins = r$ 
       $\wedge m.mbal = bal$ 
       $\wedge m.acc \in MS \}$ 
       $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
       $val \triangleq$  IF  $maxbal = -1$ 
        THEN "aborted"
        ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
    IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
       $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 
   $\wedge \text{UNCHANGED} \langle rmState, aState \rangle$ 

```

This definition of *Phase2a*, which introduces several new features of TLA+.

```

LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"} \\
\wedge m.ins = r \\
\wedge m.mbal = bal \\
\wedge m.acc \in MS \}$ 
 $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
 $val \triangleq \text{IF } maxbal = -1$ 
    THEN "aborted"
    ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
 $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 

```

This definition of *Phase2a*, which introduces several new features of TLA+.
The first is this LET-IN expression.

```

LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"} \\
\wedge m.ins = r \\
\wedge m.mbal = bal \\
\wedge m.acc \in MS \}$ 
 $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
 $val \triangleq \text{IF } maxbal = -1$ 
    THEN "aborted"
    ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
 $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 

```

This definition of *Phase2a*, which introduces several new features of TLA+.

The first is this LET-IN expression.

$$\begin{array}{l}
\text{LET } \boxed{mset \triangleq} \{ m \in msgs : \wedge m.type = \text{"phase1b"} \\
\qquad \qquad \qquad \wedge m.ins = r \\
\qquad \qquad \qquad \wedge m.mbal = bal \\
\qquad \qquad \qquad \wedge m.acc \in MS \qquad \qquad \} \\
\boxed{maxbal \triangleq} \text{Maximum}(\{m.bal : m \in mset\}) \\
\boxed{val \triangleq} \text{IF } maxbal = -1 \\
\qquad \qquad \text{THEN "aborted"} \\
\qquad \qquad \text{ELSE (CHOOSE } m \in mset : m.bal = maxbal).val \\
\text{IN } \wedge \forall ac \in MS : \exists m \in mset : m.acc = ac \\
\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])
\end{array}$$

This definition of *Phase2a*, which introduces several new features of TLA+.

The first is this LET-IN expression.

The LET clause makes three definitions local to the let-in expression.

$$\begin{array}{l}
\text{LET } mset \triangleq \{ m \in msgs : \wedge m.type = \text{"phase1b"} \\
\qquad \qquad \qquad \wedge m.ins = r \\
\qquad \qquad \qquad \wedge m.mbal = bal \\
\qquad \qquad \qquad \wedge m.acc \in MS \qquad \qquad \qquad \} \\
maxbal \triangleq \text{Maximum}(\{ m.bal : m \in mset \}) \\
val \triangleq \text{IF } maxbal = -1 \\
\qquad \qquad \text{THEN "aborted"} \\
\qquad \qquad \text{ELSE (CHOOSE } m \in mset : m.bal = maxbal).val \\
\text{IN } \wedge \forall ac \in MS : \exists m \in mset : m.acc = ac \\
\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])
\end{array}$$

This definition of *Phase2a*, which introduces several new features of TLA+.

The first is this LET-IN expression.

The LET clause makes three definitions local to the let-in expression.

The defined identifiers can be used only in the expression.

```

LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"} \\
\wedge m.ins = r \\
\wedge m.mbal = bal \\
\wedge m.acc \in MS \quad \}$ 
 $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
 $val \triangleq \text{IF } maxbal = -1$ 
    THEN "aborted"
    ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
 $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 

```

The next TLA+ notation introduced here is

$$\{m \in msgs : \wedge m.type = \text{"phase1b"} \\ \wedge m.ins = r \\ \wedge m.mbal = bal \\ \wedge m.acc \in MS \quad \}$$

The next TLA+ notation introduced here is **this set expression**. It equals

$$\{m \in \boxed{msgs} : \wedge m.type = \text{"phase1b"} \\ \wedge m.ins = r \\ \wedge m.mbal = bal \\ \wedge m.acc \in MS \quad \}$$

The subset of *msgs*

The next TLA+ notation introduced here is this set expression. It equals **The subset of *msgs***

$$\{m \in msgs : \wedge m.type = \text{"phase1b"} \\ \wedge m.ins = r \\ \wedge m.mbal = bal \\ \wedge m.acc \in MS \quad \}$$

The subset of *msgs* containing all elements *m*

The next TLA+ notation introduced here is this set expression. It equals The subset of *msgs* consisting of all its elements *m*

$$\{m \in msgs : \left. \begin{array}{l} \wedge m.type = \text{"phase1b"} \\ \wedge m.ins = r \\ \wedge m.mbal = bal \\ \wedge m.acc \in MS \end{array} \right\}$$

The subset of *msgs* containing all elements *m* satisfying this formula.

The next TLA+ notation introduced here is this set expression. It equals The subset of *msgs* consisting of all its elements *m* satisfying this formula.

```

LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"}
\wedge m.ins = r
\wedge m.mbal = bal
\wedge m.acc \in MS \}$ 
 $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
 $val \triangleq \text{IF } maxbal = -1$ 
    THEN "aborted"
    ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
 $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 

```

The LET-IN expression also introduces another set notation.

$$\{m.bal : m \in mset\}$$

The LET-IN expression also introduces another set notation.

$$\{m.bal : m \in mset\}$$

The set of all *m.bal*

The LET-IN expression also introduces another set notation.

This expression equals the set of all elements of the form *m.bal*

$$\{m.bal : m \in mset\}$$

The set of all *m.bal* with *m* in *mset*.

The LET-IN expression also introduces another set notation.

This expression equals the set of all elements of the form *m.bal* for all *m* in the set *mset*.

Two Set Constructors

These are two different set constructors.

Two Set Constructors

$$\{v \in S : P\}$$

These are two different set constructors.

The first has the form variable v in set S colon formula P .

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

These are two different set constructors.

The first has the form variable v in set S colon formula P .

It's the subset of S consisting of all values v for which the formula P is true.

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{n \in \text{Nat} : n > 17\}$$

These are two different set constructors.

The first has the form variable v in set S colon formula P .

It's the subset of S consisting of all values v for which the formula P is true.

For example, this expression

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{n \in \text{Nat} : n > 17\} = \{18, 19, 20, \dots\}$$

the set of all natural numbers greater than 17

These are two different set constructors.

The first has the form variable v in set S colon formula P .

It's the subset of S consisting of all values v for which the formula P is true.

For example, this expression equals the set of all natural numbers greater than 17.

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{e : v \in S\}$$

The second constructor has the form expression e colon variable v in set S .

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{e : v \in S\}$$

the set of all e for v in S

The second constructor has the form expression e colon variable v in set S .

It's the set consisting of all values assumed by the expression e when v is an element of S .

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{e : v \in S\}$$

the set of all e for v in S

$$\{n^2 : n \in \text{Nat}\}$$

The second constructor has the form expression e colon variable v in set S .

It's the set consisting of all values assumed by the expression e when v is an element of S .

For example, this expression

Two Set Constructors

$$\{v \in S : P\}$$

the subset of S consisting of all v satisfying P

$$\{e : v \in S\}$$

the set of all e for v in S

$$\{n^2 : n \in \text{Nat}\} = \{0, 1, 4, 9, \dots\}$$

the set of all squares of natural numbers

The second constructor has the form expression e colon variable v in set S .

It's the set consisting of all values assumed by the expression e when v is an element of S .

For example, this expression equals the set of all squares of natural numbers.

```

LET  $mset \triangleq \{m \in msgs : \wedge m.type = \text{"phase1b"} \\
\wedge m.ins = r \\
\wedge m.mbal = bal \\
\wedge m.acc \in MS \quad \}$ 
 $maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\})$ 
 $val \triangleq \text{IF } maxbal = -1$ 
    THEN "aborted"
    ELSE (CHOOSE  $m \in mset : m.bal = maxbal$ ). $val$ 
IN  $\wedge \forall ac \in MS : \exists m \in mset : m.acc = ac$ 
 $\wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val])$ 

```

There's one more thing I'd like to point out about this expression.

$$(\text{CHOOSE } m \in mset : m.bal = maxbal).val$$

There's one more thing I'd like to point out about this expression.

$(\text{CHOOSE } m \in mset : m.bal = maxbal).val$

Choice of m need not be unique.

There's one more thing I'd like to point out about this expression.

This CHOOSE expression can allow more than one possible choice for m .

(CHOOSE $m \in mset : m.bal = maxbal$).*val*

All choices of m have same value of $m.val$.

There's one more thing I'd like to point out about this expression.

This CHOOSE expression can allow more than one possible choice for m .

In any reachable state of the algorithm, all possible choices of m have the same value of $m.val$.

$$\begin{aligned}
& \text{Phase2a}(bal, r) \triangleq \\
& \quad \wedge \neg \exists m \in msgs : \wedge m.type = \text{"phase2a"} \\
& \quad \quad \quad \wedge m.bal = bal \\
& \quad \quad \quad \wedge m.ins = r \\
& \quad \wedge \exists MS \in \text{Majority} : \\
& \quad \quad \text{LET } mset \triangleq \{ m \in msgs : \wedge m.type = \text{"phase1b"} \\
& \quad \quad \quad \quad \quad \quad \wedge m.ins = r \\
& \quad \quad \quad \quad \quad \quad \wedge m.mbal = bal \\
& \quad \quad \quad \quad \quad \quad \wedge m.acc \in MS \quad \quad \quad \} \\
& \quad \quad \quad maxbal \triangleq \text{Maximum}(\{m.bal : m \in mset\}) \\
& \quad \quad \quad val \triangleq \text{IF } maxbal = -1 \\
& \quad \quad \quad \quad \quad \text{THEN "aborted"} \\
& \quad \quad \quad \quad \quad \text{ELSE (CHOOSE } m \in mset : m.bal = maxbal).val \\
& \quad \quad \text{IN} \quad \quad \wedge \forall ac \in MS : \exists m \in mset : m.acc = ac \\
& \quad \quad \quad \quad \wedge \text{Send}([type \mapsto \text{"phase2a"}, ins \mapsto r, bal \mapsto bal, val \mapsto val]) \\
& \quad \wedge \text{UNCHANGED } \langle rmState, aState \rangle
\end{aligned}$$

Paxos Commit is not an easy algorithm to understand, and this is probably its most subtle part.

I don't know how to write a clearer precise description of this step of the algorithm.

If you understand the algorithm, then when you get used to the math, I think you'll find this definition as elegant as I do.

$$\begin{aligned}
\text{Phase1b}(acc) &\triangleq \\
&\exists m \in \text{msgs} : \\
&\quad \wedge m.type = \text{"phase1a"} \\
&\quad \wedge aState[m.ins][acc].mbal < m.bal \\
&\quad \wedge aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal] \\
&\quad \wedge \text{Send}([type \mapsto \text{"phase1b"}, \\
&\quad \quad \quad ins \mapsto m.ins, \\
&\quad \quad \quad mbal \mapsto m.bal, \\
&\quad \quad \quad bal \mapsto aState[m.ins][acc].bal, \\
&\quad \quad \quad val \mapsto aState[m.ins][acc].val, \\
&\quad \quad \quad acc \mapsto acc]) \\
&\quad \wedge \text{UNCHANGED } rmState
\end{aligned}$$

The next new construct is in this definition.

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$$

The next new construct is in this definition.

In this subformula.

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$$

The next new construct is in this definition.

In this subformula.

$$[aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$$

The next new construct is in this definition.

In this subformula. you haven't seen this form of EXCEPT expression. It's an abbreviation for

[slide 130]

$$[aState \text{ EXCEPT } ![m.ins]]$$
$$[aState \text{ EXCEPT } ![m.ins]] =$$

aState EXCEPT its value on *m.ins* equals

$[aState \text{ EXCEPT } ![m.ins]][acc]$

$[aState \text{ EXCEPT } ![m.ins]] =$

$[aState[m.ins] \text{ EXCEPT } ![acc]] =$

$aState$ EXCEPT its value on $m.ins$ equals

$aState$ of $m.ins$ EXCEPT its value on a-c-c equals

$$\begin{aligned} & [aState \text{ EXCEPT } ![m.ins][acc].mbal \\ & [aState \text{ EXCEPT } ![m.ins] = \\ & \quad [aState[m.ins] \text{ EXCEPT } ![acc] = \\ & \quad \quad [aState[m.ins][acc] \text{ EXCEPT }!.mbal = \end{aligned}$$

aState EXCEPT its value on *m.ins* equals

aState of *m.ins* EXCEPT its value on a-c-c equals

aState of *m.ins* of a-c-c EXCEPT its m-bal component equals

$[aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$

$[aState \text{ EXCEPT } ![m.ins] =$

$[aState[m.ins] \text{ EXCEPT } ![acc] =$

$[aState[m.ins][acc] \text{ EXCEPT }!.mbal = m.bal]]]$

$aState$ EXCEPT its value on $m.ins$ equals

$aState$ of $m.ins$ EXCEPT its value on a-c-c equals

$aState$ of $m.ins$ of a-c-c EXCEPT its m-bal component equals m dot bal.

Whew.

```
[aState EXCEPT ![m.ins][acc].mbal = m.bal]
[aState EXCEPT ![m.ins] =
  [aState[m.ins] EXCEPT ![acc] =
    [aState[m.ins][acc] EXCEPT !.mbal = m.bal]]]
```

If you stop and decipher this, you'll see that

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$$

If you stop and decipher this, you'll see that
this formula corresponds to

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal]$$

```
aState[m.ins][acc].mbal = m.bal
```

If you stop and decipher this, you'll see that this formula corresponds to **this programming-language statement**.

So you just have to remember this idiom and not try to figure out the **EXCEPT** expression. That's what I do.

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal, \\ ![m.ins][acc].bal = m.bal, \\ ![m.ins][acc].val = m.val]$$

This definition contains another generalization of the EXCEPT construct.
no pause

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal, \\ ![m.ins][acc].bal = m.bal, \\ ![m.ins][acc].val = m.val]$$

This definition contains another generalization of the EXCEPT construct.
no pause

```
[aState EXCEPT ![m.ins][acc].mbal = m.bal,  
                  ![m.ins][acc].bal   = m.bal,  
                  ![m.ins][acc].val   = m.val]
```

This definition contains another generalization of the EXCEPT construct.
no pause

If you want, you can try to figure out what this EXCEPT expression means
when I tell you that

$$\begin{aligned} aState' = [aState \text{ EXCEPT } & ![m.ins][acc].mbal = m.bal, \\ & ![m.ins][acc].bal = m.bal, \\ & ![m.ins][acc].val = m.val] \end{aligned}$$

This definition contains another generalization of the EXCEPT construct.
no pause

If you want, you can try to figure out what this EXCEPT expression means when I tell you that
this subformula describes the same change to $aState$ as

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal, \\ ![m.ins][acc].bal = m.bal, \\ ![m.ins][acc].val = m.val]$$

```
aState[m.ins][acc].mbal = m.bal ;  
aState[m.ins][acc].bal   = m.bal ;  
aState[m.ins][acc].val   = m.val ;
```

executing this sequence of three program statements.

Notice the correspondence between the parts of the EXCEPT expression

$$aState' = [aState \text{ EXCEPT } \begin{array}{l} ![m.ins][acc].mbal = m.bal, \\ ![m.ins][acc].bal = m.bal, \\ ![m.ins][acc].val = m.val \end{array}]$$

```
aState[m.ins][acc].mbal = m.bal ;  
aState[m.ins][acc].bal = m.bal ;  
aState[m.ins][acc].val = m.val ;
```

executing this sequence of three program statements.

Notice the correspondence between the parts of the EXCEPT expression and the program statements.

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal,$$
$$![m.ins][acc].bal = m.bal,$$
$$![m.ins][acc].val = m.val]$$

```
aState[m.ins][acc].mbal = m.bal ;
```

```
aState[m.ins][acc].bal = m.bal ;
```

```
aState[m.ins][acc].val = m.val ;
```

executing this sequence of three program statements.

Notice the correspondence between the parts of the EXCEPT expression and the program statements.

$$aState' = [aState \text{ EXCEPT } ![m.ins][acc].mbal = m.bal,$$
$$![m.ins][acc].bal = m.bal,$$
$$![m.ins][acc].val = m.val]$$

```
aState[m.ins][acc].mbal = m.bal ;  
aState[m.ins][acc].bal   = m.bal ;  
aState[m.ins][acc].val   = m.val ;
```

executing this sequence of three program statements.

Notice the correspondence between the parts of the EXCEPT expression and the program statements.

CHECKING THE SPEC

Checking the Specification

Open *PaxosCommit* in the Toolbox

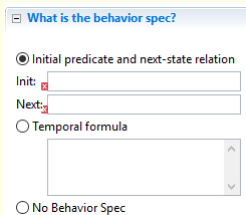
Open module *PaxosCommit* in the Toolbox

Open *PaxosCommit* in the Toolbox and create a new model.

Open module *PaxosCommit* in the Toolbox and create a new model.

Open *PaxosCommit* in the Toolbox and create a new model.

You have to enter



What is the behavior spec?

Initial predicate and next-state relation

Init:

Next:

Temporal formula

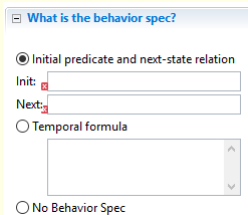
No Behavior Spec

Open module *PaxosCommit* in the Toolbox and create a new model.

You have to enter the initial and next-state formulas

Open *PaxosCommit* in the Toolbox and create a new model.

You have to enter



What is the behavior spec?

Initial predicate and next-state relation

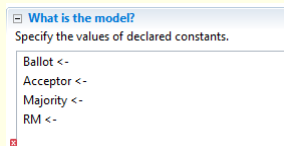
Init:

Next:

Temporal formula

No Behavior Spec

and



What is the model?

Specify the values of declared constants.

Ballot <-

Acceptor <-

Majority <-

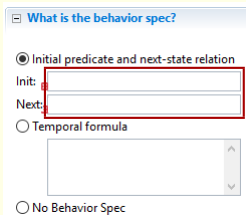
RM <-

Open module *PaxosCommit* in the Toolbox and create a new model.

You have to enter the initial and next-state formulas and the values of the constants.

Open *PaxosCommit* in the Toolbox and create a new model.

You have to enter



What is the behavior spec?

Initial predicate and next-state relation

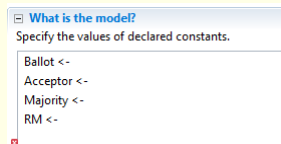
Init:

Next:

Temporal formula

No Behavior Spec

and



What is the model?

Specify the values of declared constants.

Ballot <-

Acceptor <-

Majority <-

RM <-

The initial and next-state formulas are named

Open *PaxosCommit* in the Toolbox and create a new model.

You have to enter

What is the behavior spec?

Initial predicate and next-state relation

Init:

Next:

Temporal formula

No Behavior Spec

and

What is the model?

Specify the values of declared constants.

Ballot <-

Acceptor <-

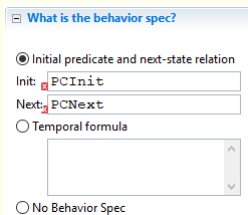
Majority <-

RM <-

The initial and next-state formulas are named *PCInit* and *PCNext*.

Open *PaxosCommit* in the Toolbox and create a new model.

You have to enter



What is the behavior spec?

Initial predicate and next-state relation

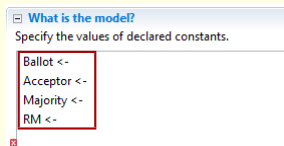
Init:

Next:

Temporal formula

No Behavior Spec

and



What is the model?

Specify the values of declared constants.

Ballot <-

Acceptor <-

Majority <-

RM <-

The initial and next-state formulas are named *PCInit* and *PCNext*.

Now for the values assigned to the constants.

Ballot ←
Acceptor ←
Majority ←
RM ←

We normally start with a tiny model

We normally start with a tiny model

Ballot ←
Acceptor ←
Majority ←
RM ←

We normally start with a tiny model, but we'll skip that.

We normally start with a tiny model but we'll skip that.

Ballot ←
Acceptor ←
Majority ←
RM ←

We normally start with a tiny model, but we'll skip that.

Instead, we'll use the smallest model that could reveal an error in the algorithm.

We normally start with a tiny model but we'll skip that.

Instead, we'll use a model which, if you understand the algorithm, you'll see is the smallest one that could reveal a non-trivial error.

Ballot ←
Acceptor ← {a1, a2, a3} a set of model values
Majority ←
RM ←

We normally start with a tiny model but we'll skip that.

Instead, we'll use a model which, if you understand the algorithm, you'll see is the smallest one that could reveal a non-trivial error.

We assign a set of three model values to *Acceptor*,

Ballot \leftarrow
Acceptor $\leftarrow \{a1, a2, a3\}$
Majority \leftarrow
RM $\leftarrow \{r1, r2\}$ a set of model values

We normally start with a tiny model but we'll skip that.

Instead, we'll use a model which, if you understand the algorithm, you'll see is the smallest one that could reveal a non-trivial error.

We assign a set of three model values to *Acceptor*, and a set of two model values to *RM*.

[slide 160]

Ballot $\leftarrow \{0, 1\}$
Acceptor $\leftarrow \{a1, a2, a3\}$
Majority \leftarrow
RM $\leftarrow \{r1, r2\}$

We assign this set of two numbers to *Ballot*,

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

We assign this set of two numbers to *Ballot*, and this set of sets of acceptors to *Majority*.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

What is the model?
Specify the values of declared constants.

Majority <- |{a1, a2}, {a1, a3}, {a2, a3}|

Ordinary assignment
 Model value
 Set of model values
 Symmetry set

We assign this set of two numbers to *Ballot*, and this set of sets of acceptors to *Majority*.

This is an ordinary assignment,

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

What is the model?
Specify the values of declared constants.

Majority <- |{a1, a2}, {a1, a3}, {a2, a3}|

Ordinary assignment
 Model value
 Set of model values
 Symmetry set

We assign this set of two numbers to *Ballot*, and this set of sets of acceptors to *Majority*.

This is an ordinary assignment, because the model values *a1*, *a2*, and *a3* are declared in the assignment of a set of model values to *Acceptor*.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

We assign this set of two numbers to *Ballot*, and this set of sets of acceptors to *Majority*.

This is an ordinary assignment, because the model values *a1*, *a2*, and *a3* are declared in the assignment of a set of model values to *Acceptor*.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

This can be a symmetry set, because
 $r1$ and $r2$ aren't used elsewhere.

The set we assigned to RM can be a symmetry set because its elements aren't used elsewhere.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

But what about this set?

The set we assigned to *RM* can be a symmetry set because its elements aren't used elsewhere.

But what about the set we assigned to *Acceptor*?

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

But what about this set?

$a1, a2, a3$ are used here.

The set we assigned to *RM* can be a symmetry set because its elements aren't used elsewhere.

But what about the set we assigned to *Acceptor*?

Its elements are used in the value assigned to *Majority*.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

This use is OK because the expression is symmetric in $a1, a2, a3$.

The set we assigned to *RM* can be a symmetry set because its elements aren't used elsewhere.

But what about the set we assigned to *Acceptor*?

Its elements are used in the value assigned to *Majority*.

But this use is OK because the expression they appear in is symmetric in the elements of the set we assigned to *Acceptor*.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a_1, a_2, a_3\}$

Majority $\leftarrow \{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

RM $\leftarrow \{r_1, r_2\}$

This use is OK because the expression is symmetric in a_1, a_2, a_3 .

Interchanging any two of these elements leaves the expression unchanged.

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a_1, a_2, a_3\}$

Majority $\leftarrow \{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

RM $\leftarrow \{r_1, r_2\}$

For example, interchanging $a_1 \leftrightarrow a_3$ in

$\{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange a_1 and a_3 in the expression,

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a_1, a_2, a_3\}$

Majority $\leftarrow \{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

RM $\leftarrow \{r_1, r_2\}$

For example, interchanging $a_1 \leftrightarrow a_3$ in

$\{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

produces

\downarrow

\downarrow

\downarrow

\downarrow

$\{\{a_3, a_2\}, \{a_3, a_1\}, \{a_2, a_1\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange a_1 and a_3 in the expression, we get this expression.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

These two sets are equal:

$\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

$\{\{a3, a2\}, \{a3, a1\}, \{a2, a1\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange $a1$ and $a3$ in the expression, we get this expression.

And these two expressions are equal because they describe sets with the same three elements

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

These two sets are equal:

$\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

$\{\{a3, a2\}, \{a3, a1\}, \{a2, a1\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange $a1$ and $a3$ in the expression, we get this expression.

And these two expressions are equal because they describe sets with the same three elements **one**

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

These two sets are equal:

$\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

$\{\{a3, a2\}, \{a3, a1\}, \{a2, a1\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange $a1$ and $a3$ in the expression, we get this expression.

And these two expressions are equal because they describe sets with the same three elements one two

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

These two sets are equal:

$\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

$\{\{a3, a2\}, \{a3, a1\}, \{a2, a1\}\}$

Remember, this means that interchanging any two elements of that set leaves the expression unchanged.

For example, if we interchange $a1$ and $a3$ in the expression, we get this expression.

And these two expressions are equal because they describe sets with the same three elements one two three

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

It's OK to use elements of a symmetry set

In general, it's OK to use elements of a symmetry set

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

It's OK to use elements of a symmetry set
in an expression assigned to another constant

In general, it's OK to use elements of a symmetry set
in an expression assigned to another constant

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

It's OK to use elements of a symmetry set
in an expression assigned to another constant
if the expression is symmetric

In general, it's OK to use elements of a symmetry set
in an expression assigned to another constant
if the expression is symmetric

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

It's OK to use elements of a symmetry set
in an expression assigned to another constant
if the expression is symmetric in the elements
of the symmetry set.

In general, it's OK to use elements of a symmetry set
in an expression assigned to another constant
if the expression is symmetric
in the elements of the symmetry set.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

There's just one additional condition a symmetry set must satisfy that I can now explain.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a_1, a_2, a_3\}$

Majority $\leftarrow \{\{a_1, a_2\}, \{a_1, a_3\}, \{a_2, a_3\}\}$

RM $\leftarrow \{r_1, r_2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set

There's just one additional condition a symmetry set must satisfy that I can now explain.

Elements of a symmetry set,

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set , or
a constant assigned elements of a symmetry set

There's just one additional condition a symmetry set must satisfy that I can now explain.

Elements of a symmetry set,
or a constant that's assigned elements of a symmetry set

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set, or
a constant assigned elements of a symmetry set
may not appear in a CHOOSE expression.

There's just one additional condition a symmetry set must satisfy that I can
now explain.

Elements of a symmetry set,
or a constant that's assigned elements of a symmetry set
may not appear in a CHOOSE expression.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set, or
a constant assigned elements of a symmetry set
may not appear in a CHOOSE expression.

In the *Paxos Commit* spec, elements of a symmetry set don't appear in a CHOOSE because

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set, or
a constant assigned elements of a symmetry set
may not appear in a CHOOSE expression.

In the *Paxos Commit* spec, elements of a symmetry set don't appear in a CHOOSE because they can appear only in these assignments and there's no CHOOSE there.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

There's one additional condition for symmetry sets.

Elements of a symmetry set, or
a constant assigned elements of a symmetry set
may not appear in a CHOOSE expression.

In the *PaxosCommit* spec, elements of a symmetry set don't appear in a CHOOSE because they can appear only in these assignments and there's no CHOOSE there.

To verify that a constant which is assigned elements of a symmetry set doesn't appear in a CHOOSE expression,

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

We must check that these constants don't appear
in a CHOOSE expression of the spec.

In the *PaxosCommit* spec, elements of a symmetry set don't appear in a CHOOSE because they can appear only in these assignments and there's no CHOOSE there.

To verify that a constant which is assigned elements of a symmetry set doesn't appear in a CHOOSE expression, we must check that these constants don't appear in any CHOOSE expression in the spec.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

We must check that these constants don't appear
in a CHOOSE expression of the spec.

They don't.

You can check that they don't.

Ballot $\leftarrow \{0, 1\}$

Acceptor $\leftarrow \{a1, a2, a3\}$

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

RM $\leftarrow \{r1, r2\}$

Assign these values in the model, with *Acceptor* and *RM* being symmetry sets.

You can check that they don't.

Assign these values in the model, letting *Acceptor* and *RM* be symmetry sets.

What to Check

We should check that the algorithm is correct.

We'll see in a later video, how to check that it implements transaction commit.

What to Check

There are two invariants we can check:

We should check that the algorithm is correct.

We'll see in a later video, how to check that it implements transaction commit.

For now, there are two invariants we can check:

What to Check

There are two invariants we can check:

- The type correctness invariant $PCTypeOK$

The type correctness invariant $PCTypeOK$ that we looked at earlier

What to Check

There are two invariants we can check:

- The type correctness invariant *PCTypeOK*
- Invariant *TCConsistent* imported from module *TCommit*

The type correctness invariant *PCTypeOK* that we looked at earlier and the invariant *TCConsistent*, which is imported with an `INSTANCE` statement from module *TCommit*.

What to Check

There are two invariants we can check:

- The type correctness invariant *PCTypeOK*
- Invariant *TCConsistent* imported from module *TCommit*

Add them and run TLC on the model.

The type correctness invariant *PCTypeOK* that we looked at earlier and the invariant *TCConsistent*, which is imported with an `INSTANCE` statement from module *TCommit*.

Add these invariants to the *What to check* part of the model and run TLC on the model.

TLC takes 30 seconds to run the model on my laptop using two cores.

TLC takes about 30 seconds to run the model on my laptop using two cores.

TLC takes 30 seconds to run the model on my laptop using two cores. It reports no error and finds 120 thousand distinct states.

TLC takes about 30 seconds to run the model on my laptop using two cores. It reports no error and finds about 120 thousand distinct states.

TLC takes 30 seconds to run the model on my laptop using two cores. It reports no error and finds 120 thousand distinct states.

If we change the model to assign *Ballot* the set $\{0, 1, 2\}$ instead of $\{0, 1\}$

TLC takes about 30 seconds to run the model on my laptop using two cores. It reports no error and finds about 120 thousand distinct states.

If we change the model to assign *Ballot* a set of three numbers instead of two,

TLC takes 30 seconds to run the model on my laptop using two cores. It reports no error and finds 120 thousand distinct states.

If we change the model to assign *Ballot* the set $\{0, 1, 2\}$ instead of $\{0, 1\}$, TLC runs for $1\frac{1}{2}$ hours on a 128 core machine and finds 220 million states.

TLC runs for about one and a half hours on a 128 core machine and finds about 220 million states.

We use very small models because

TLC takes 30 seconds to run the model on my laptop using two cores. It reports no error and finds 120 thousand distinct states.

If we change the model to assign *Ballot* the set $\{0, 1, 2\}$ instead of $\{0, 1\}$, TLC runs for $1\frac{1}{2}$ hours on a 128 core machine and finds 220 million states.

Execution time and space grow exponentially with the size of the model.

TLC runs for about one and a half hours on a 128 core machine and finds about 220 million states.

We use very small models because

execution time and space grow exponentially with the size of the model.

What good is checking such small models?

What good is checking such small models?

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

What good is checking such small models?

To answer that question, make this change to value the model assigns to *Majority*.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a2, a3\}\}$

What good is checking such small models?

To answer that question, make this change to value the model assigns to *Majority*.

Delete this element of an element of the set.

What good is checking such small models?

Make this change to the model.

Majority $\leftarrow \{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

What good is checking such small models?

To answer that question, make this change to value the model assigns to *Majority*.

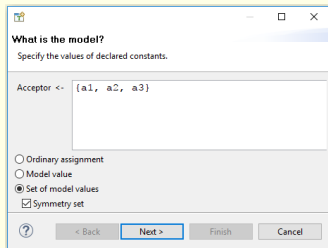
Delete this element of an element of the set.

The expression is no longer symmetric in $a1$, $a2$, and $a3$.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$




So we have to change the assignment to *Acceptor*

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$



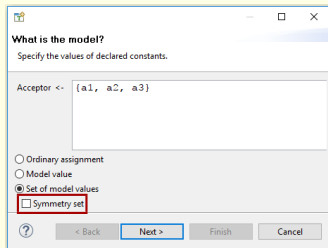
So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$



So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

If you run TLC on the model, it will complain that the assumption is violated.

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

ASSUME

```
/\ Ballot \subseteq Nat
/\ 0 \in Ballot
/\ Majority \subseteq SUBSET Acceptor
/\ \A MS1, MS2 \in Majority : MS1 \cap MS2 # {}
```

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

Because

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

ASSUME

```
/\ Ballot \subseteqq Nat
/\ 0 \in Ballot
/\ Majority \subseteqq SUBSET Acceptor
/\ \A MS1, MS2 \in Majority : MS1 \cap MS2 # {}
```

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

Because this assertion is no longer true. So, we have to comment it out.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

ASSUME

```
/\ Ballot \subseteqq Nat
/\ 0 \in Ballot
/\ Majority \subseteqq SUBSET Acceptor
/\ \A MS1, MS2 \in Majority : MS1 \cap MS2 # {}
```

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

Because this assertion is no longer true. So, we have to comment it out.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

```
ASSUME
/\ Ballot \subseteqq Nat
/\ 0 \in Ballot
/\ Majority \subseteqq SUBSET Acceptor
\ /\ \A MS1, MS2 \in Majority : MS1 \cap MS2 # {}
```

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

Because this assertion is no longer true. So, we have to comment it out.

What good is checking such small models?

Make this change to the model.

Majority \leftarrow $\{\{a1, a2\}, \{a1, a3\}, \{a3\}\}$

```
ASSUME
/\ Ballot \subseteqq Nat
/\ 0 \in Ballot
/\ Majority \subseteqq SUBSET Acceptor
\* /\ \A MS1, MS2 \in Majority : MS1 \cap MS2 # {}
```

So we have to change the assignment to *Acceptor*

So it's no longer a symmetry set.

Now if you run TLC on the model, it will complain that this assumption is violated

Because this assertion is no longer true. So, we have to comment it out.

Run TLC on the model.

Run TLC on the model.

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this value of *Majority*.

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this changed value of *Majority*.

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this value of *Majority*.

TLC reports that invariant *TCConsistent* is violated

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this changed value of *Majority*.

TLC reports that invariant *TCConsistent* is violated

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this value of *Majority*.

TLC reports that invariant *TCConsistent* is violated, and it produces a 14-state error trace.

Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this changed value of *Majority*.

TLC reports that invariant *TCConsistent* is violated and it produces a minimal-length 14-state error trace.

The Paxos commit algorithm is correct.


Run TLC on the model.

Because the assumption is not satisfied, the algorithm is incorrect for this value of *Majority*.

TLC reports that invariant *TCConsistent* is violated, and it produces a 14-state error trace.

Even a very small model can catch an error in an algorithm.

But this example shows that even a very small model can catch an error in a real algorithm.



You've now learned enough of the TLA+ language to start writing your own specs. However, before you do that, you should know more about what TLA+ specs *mean*. In particular, you should understand what it means for the Paxos Commit algorithm to implement the transaction-commit spec. That's the topic of the next lecture.

End of Lecture 7

PAXOS COMMIT