

Everything About Distributed Systems is Terrible

Hillel Wayne
hillelwayne.com
@hillelogram

Designing Distributed Systems with TLA+

Hillel Wayne
hillelwayne.com
@hillelogram

[hillelwayne.com/talks/
designing-distributed-
systems](https://hillelwayne.com/talks/designing-distributed-systems)

"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."

Leslie Lamport

Things Lamport did

- TLA+
- Paxos
- Bakery Algorithm
- Byzantine Generals
- LaTeX
- Many hats



@hillelogram

Distributed System

- Multiple agents
- Global properties
- Localized information
- Partial Failure

Computer 1

```
tmp = serverdb.get(x)  
serverdb.set(x, tmp + 1)
```

Computer 2

```
tmp = serverdb.get(x)  
serverdb.set(x, tmp + 1)
```

Thread 1

tmp = x

x = tmp + 1

Thread 2

tmp = x

x = tmp + 1

Threads = Computers

Temporal Logic


global $x = 1$


agent 1


$x = x + 1$

agent 2

$x = x * 2$







Behavior

global x = 1

agent 1

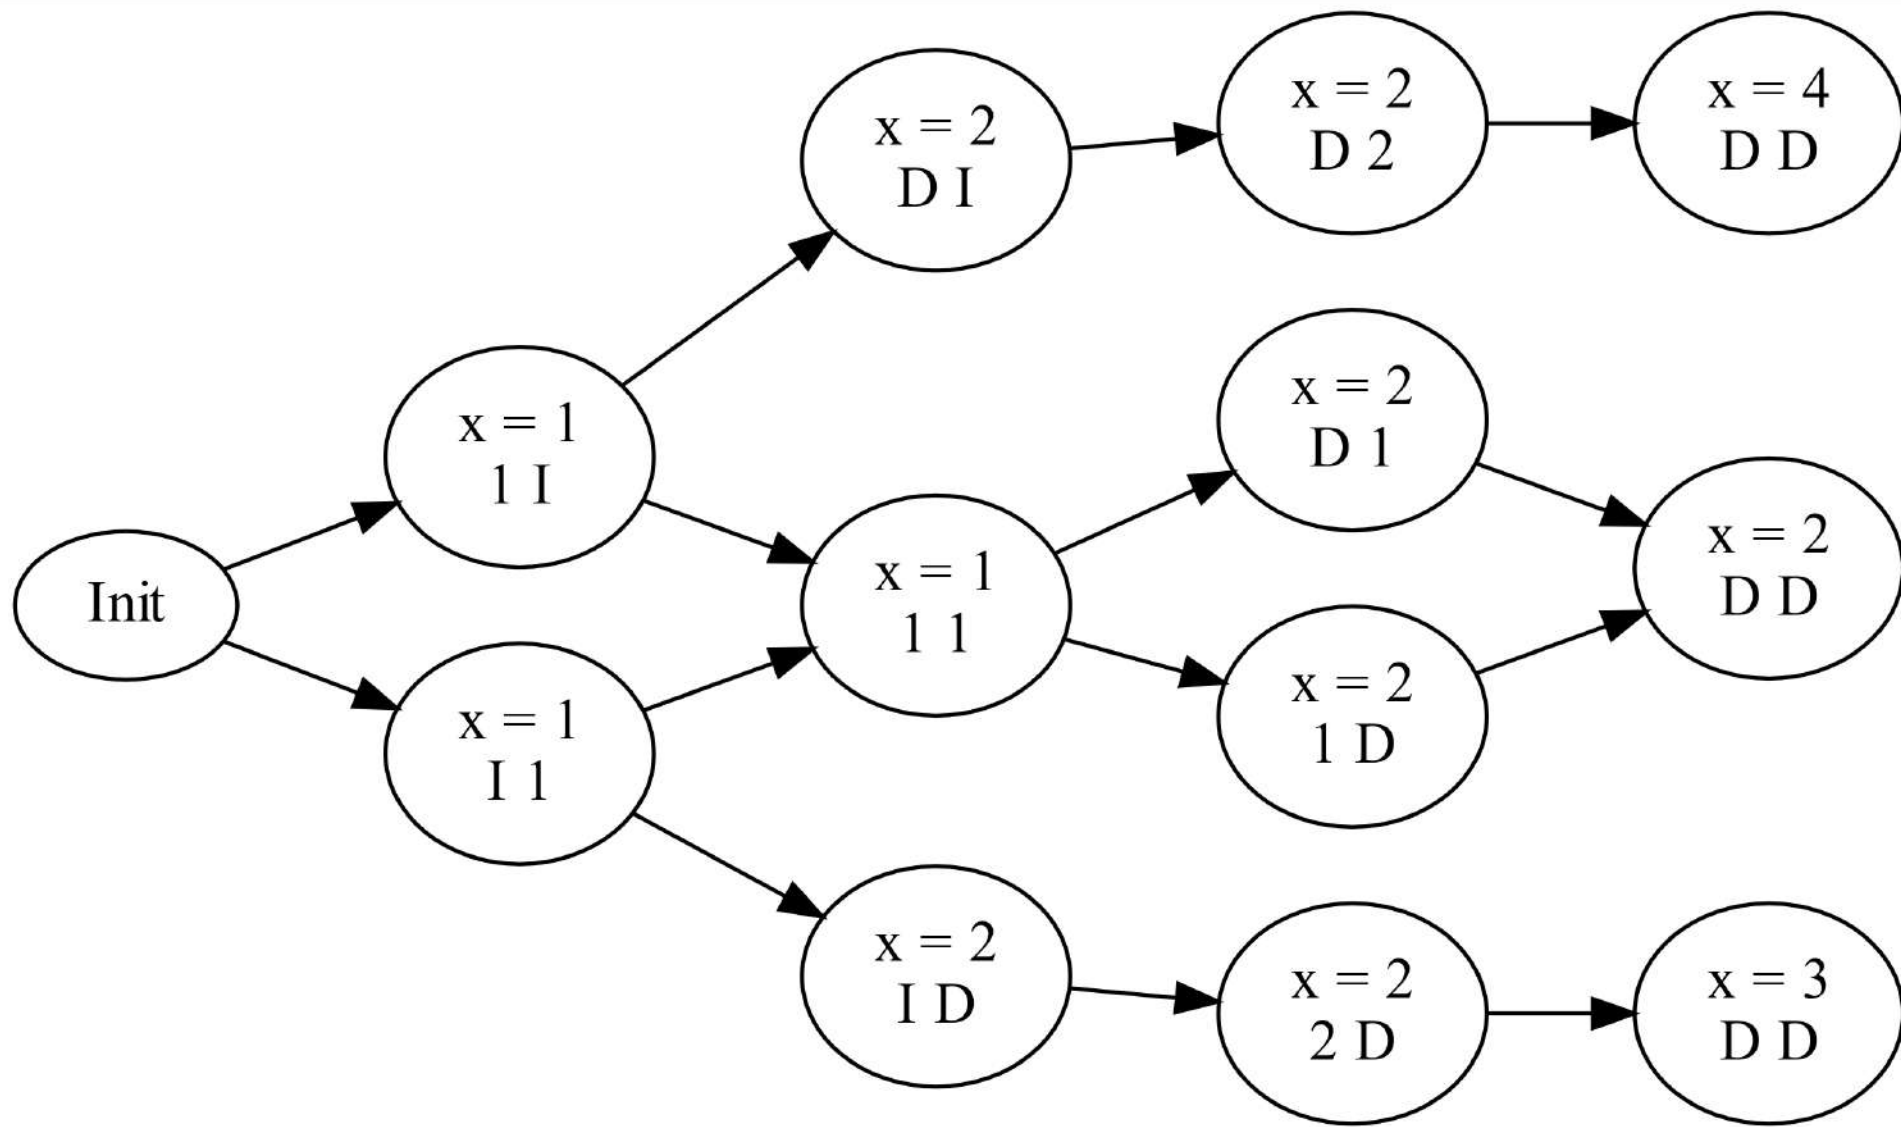
local tmp = x

x = tmp + 1

agent 2

local tmp = x


x = tmp * 2



$$(m \cdot n)(m \cdot n)! / m! \uparrow n$$

n = num agents

m = num steps



https://www3.hhu.de/stups/prob/index.php/State_space_visualization_examples

The number of states grows fast

$x = 0$


$x = x + 1$

or

$x = x + 2$

or


whoops crash






```
x = 0
```

```
while true:  
    x = x + 1  
or  
    x = x + 2
```

```
while true:  
    if x > 0:  
        x = x - 1
```











https://www3.hhu.de/stups/prob/index.php/State_space_visualization_examples


The number of behaviors grows
very fast













Systems may have invalid
behaviors and states

0.0000001%



0.0000001%
~ 3 months



Over a long enough time, a system will do everything.

Anything that can go wrong, will
go wrong.

Actors!

Types!


TDD!


Purity!


Drugs!


- CAS
- Transactions
- STM
- Locks
- Semaphores
- Promises
- Nurseries
- Async/await

- Monads
- Futures
- Supervisors
- Actors
- Goroutines









Code is not

Design

Formal Specification

πλΑ+

Max(set) ==
 CHOOSE x \in set:
 \A y \in set:
 x >= y

INIT $x = 1$

NEXT $x' > x$

OR $x \neq 0$ AND $x' = 0$

$1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \dots$

$1 \rightarrow 1.5 \rightarrow 0 \rightarrow 1$

$1 \rightarrow 2 \rightarrow 1$

State Invariants

"At least one server is online"

Action Invariants

"We do not remove servers if below capacity"

Behavior Invariants

"Eventually we have enough servers"

explicit

TLC Errors


Model_1

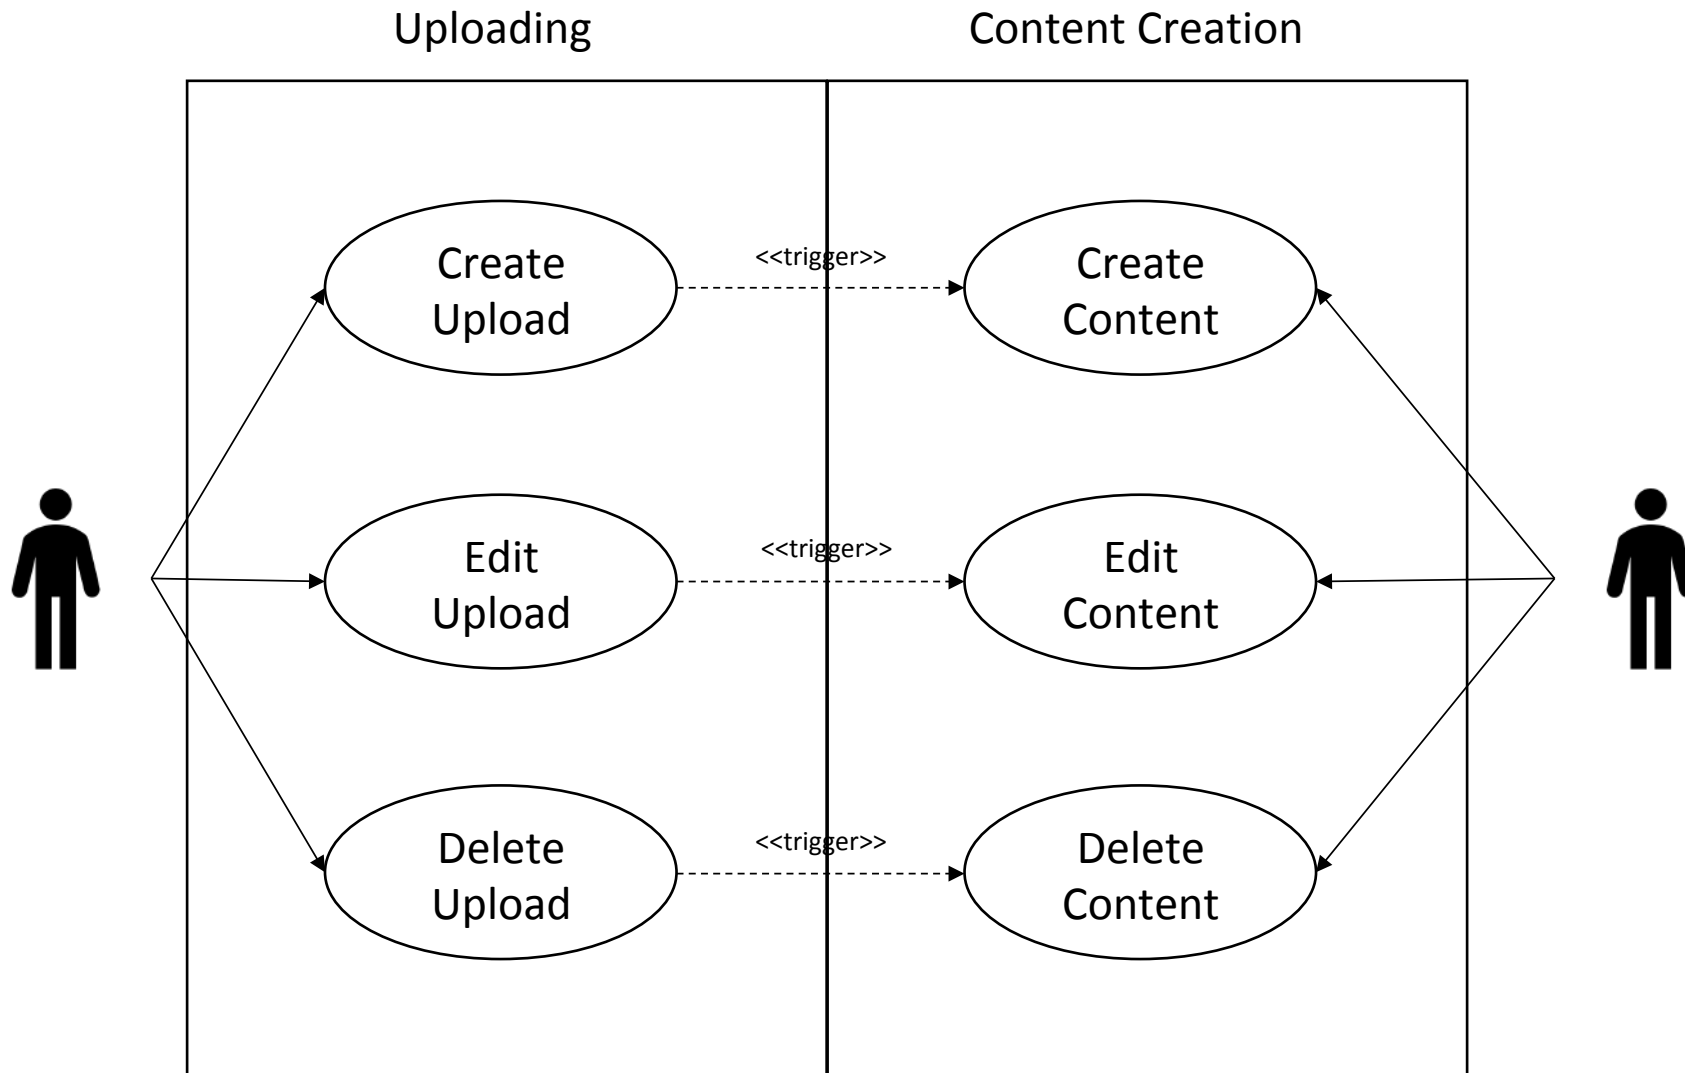
Temporal properties were violated.

Error-Trace Exploration

Error-Trace

Name	Value
▶ ▲ <Initial predicate>	State (num = 1)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 2)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 3)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 4)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 5)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 6)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 7)
▶ ▲ <Action line 119, col 20 to line...>	State (num = 8)
▶ ▲ <Action line 128, col 24 to line...>	State (num = 9)
▶ ▲ <Action line 128, col 24 to line...>	State (num = 10)
▶ ▲ <Action line 128, col 24 to line...>	State (num = 11)
▶ ▲ <Action line 119, col 20 to line...>	State (num = 12)
▶ ▲ <Action line 128, col 24 to line...>	State (num = 13)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 14)
▶ ▲ <Action line 128, col 24 to line...>	State (num = 15)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 16)
▶ ▲ <Action line 111, col 18 to line...>	State (num = 17)
▶ ▲ <Back to state 8>	State (num = 8)





Specs do not
check code

Original

$[a \rightarrow b, c \rightarrow d]$

$\langle\langle 1, 2, 5 \rangle\rangle$

\exists

Slides

$\{a: b, c: d\}$

$[1, 2, 5]$

\exists Exists

[hillelwayne.com/talks/
designing-distributed-
systems](https://hillelwayne.com/talks/designing-distributed-systems)

EXTENDS Sequences, Integers, TLC
PARAMS Users, Workers, NULL


```
(*--algorithm uploader
```

```
variables
```

```
  uploads = {};
```

```
  content = {};
```

```
  queue = [];
```

```
  next_id = 1;
```

```
macro send_msg(id, action) begin
  queue := Append(queue,
    {id: id,
     action: action});
end macro;
```

```
process user \in Users
begin
  User:
    while TRUE do
      either
        \* Create
      or
        \* Edit
      or
        \* Delete
      end either;
    end while;
end process;
```


either

`x := x + 1;`

or

`x := x * 2;`

end either;



Create

with

```
create = {from: self,  
          version: 1,  
          id: next_id}
```

do

```
uploads := uploads ++ create;  
send_msg(create.id, "create");  
next_id := next_id + 1;
```

end with;

Delete

with

```
delete \in {u \in uploads:  
          u.from = self}
```

do

```
uploads := uploads -- delete;  
send_msg(delete.id, "delete");
```

end with;


with

```
  y \in {1, 2, 3}
```

do

```
  x := x + y;
```

end with;



Edit

with

```
upload \in {u \in uploads: u.from = self},  
edit = {from: upload.from,  
        id: upload.id,  
        version: upload.version + 1}
```

do

```
uploads := uploads -- upload ++ edit;  
send_msg(edit.id, "edit");
```

end with;


```
process worker \in Workers
variables
  msg = NULL; local = NULL;
begin
  Receive:
  while TRUE do
    await Nonempty(queue);
    msg := Head(queue);
    queue := Tail(queue);
    \* process message code
  end while;
end process;
```

Process Message Code

```
if msg.action = "create" then
    Create:
elseif msg.action = "edit" then
    Edit:
    PushEdit:
elseif msg.action = "delete" then
    Delete:
else \* wtf
    assert FALSE;
end if;
```

A:

$x := x + 1$

B:


either

$x := 0$

or

$x := 10$

end either;



Create

```
with upload = UploadWith(msg.id) do  
  content := content ++ upload;  
end with;
```

```
def UploadWith(id):  
    CHOOSE msg \in uploads:  
        msg.id = id
```

Edit

with

```
    upload = UploadWith(msg.id),
```

```
    exists = ContentWith(msg.id)
```

do

```
    content := content -- exists;
```

```
    local := upload;
```

end with;

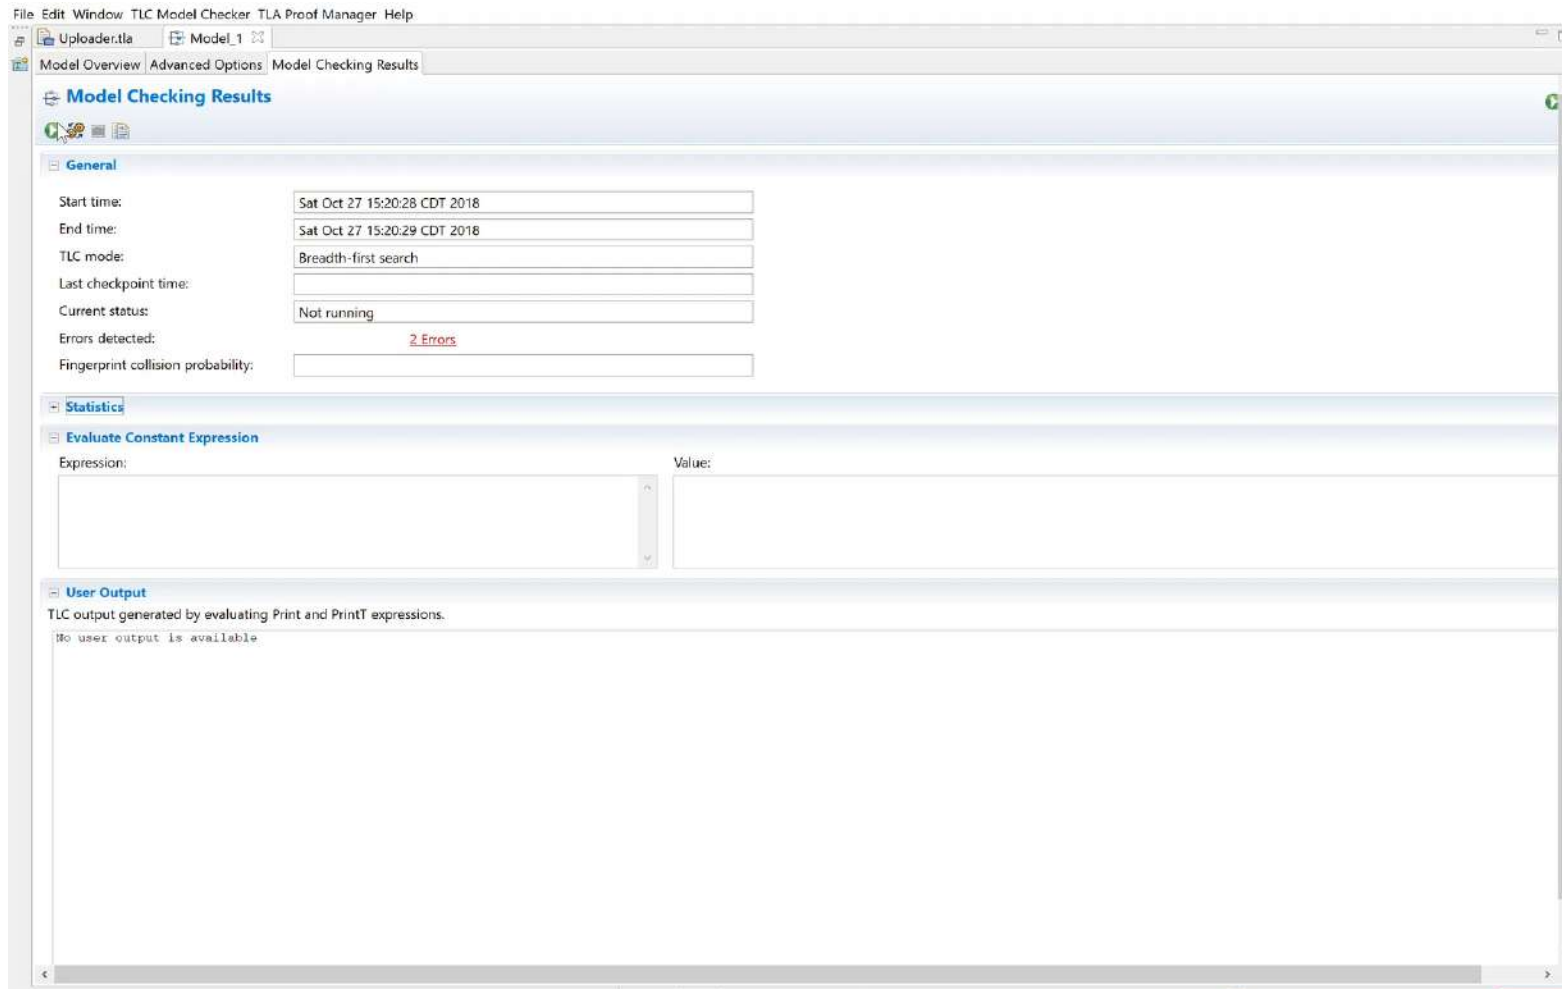
PushEdit

```
content := content ++ local;  
local := NULL;
```

PARAMS

```
Workers <- {w1}
```

```
Users <- {u1}
```

Model Checking Results [State space exploration incomplete](#)

General

Start time: Fri Oct 05 00:40:50 CDT 2018
 End time: Fri Oct 05 00:40:52 CDT 2018
 TLC mode: Breadth-first search
 Last checkpoint time:
 Current status: Not running
 Errors detected: **2 Errors**
 Fingerprint collision probability:

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:02	5	27	14	5
00:-59:-57	0	1	1	1

Coverage at 2018-10-05 00:40:52

Module	Location
Uploader	line 218, col 29 to
Uploader	line 219, col 29 to
Uploader	line 220, col 29 to
Uploader	line 223, col 31 to
Uploader	line 224, col 31 to

Evaluate Constant Expression

Expression: Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

TLC threw an unexpected exception. This was probably caused by an error in the spec or model. See the User Output or TLC Console for clues to what happened. The exception was a java.lang.RuntimeException: Attempted to compute the value of an expression of form CHOOSE x \in S: P, but no element of S satisfied P. [line 157, col 19 to line 157, col 53 of module Uploader](#)
 The error occurred when TLC was evaluating the nested expressions at the following positions:

Error-Trace Exploration

Error-Trace

Name	Value
<Initial predicate>	State (num = 1)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	1
pc	(w1 := "Receive" @@ u1 := "User")
queue	<< >>
uploads	{}
<User line 216, col 15 to line 231, col 52 of moc>	State (num = 2)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	2
pc	(w1 := "Receive" @@ u1 := "User")
queue	<<[id -> 1, action -> "create"]>>
uploads	{[from -> u1, version -> 1, id -> 1]}
<User line 216, col 15 to line 231, col 52 of moc>	State (num = 3)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	2
pc	(w1 := "Receive" @@ u1 := "User")
queue	<<[id -> 1, action -> "create"], [id -> 1, action -> "...]>>
uploads	{}
<Receive line 235, col 18 to line 248, col 68 of r>	State (num = 4)
content	{}
local	(w1 := NULL)
msg	(w1 := [id -> 1, action -> "create"])
next_id	2
pc	(w1 := "Create" @@ u1 := "User")

Select line in Error Trace to show its value here.

Error

What happened

Model_1

TLC threw an unexpected exception.
This was probably caused by an error in the spec or model.
See the User Output or TLC Console for clues to what happened.
The exception was a java.lang.RuntimeException
: Attempted to compute the value of an expression of form
CHOOSE x \in S: P, but no element of S satisfied P.
[line 157, col 19 to line 157, col 53 of module Uploader](#)
The error occurred when TLC was evaluating the nested
expressions at the following positions:

Error-Trace Exploration

Error-Trace

Name	Value
▼ ▲ <Initial predicate>	State (num = 1)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	1
pc	(w1 := "Receive" @@ u1 := "User")
queue	<< >>
uploads	{}
▼ ▲ <User line 216, col 15 to line 231, col 52 of moc	State (num = 2)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	2
pc	(w1 := "Receive" @@ u1 := "User")
queue	<<[id -> 1, action -> "create"]>>
uploads	{[from -> u1, version -> 1, id -> 1]}
▼ ▲ <User line 216, col 15 to line 231, col 52 of moc	State (num = 3)
content	{}
local	(w1 := NULL)
msg	(w1 := NULL)
next_id	2
pc	(w1 := "Receive" @@ u1 := "User")
queue	<<[id -> 1, action -> "create"], [id -> 1, action -> "...]
uploads	{}
▼ ▲ <Receive line 235, col 18 to line 248, col 68 of r	State (num = 4)
content	{}
local	(w1 := NULL)
msg	(w1 := [id -> 1, action -> "create"])
next_id	2
pc	(w1 := "Create" @@ u1 := "User")

Select line in Error Trace to show its value here.

Spec Status : parsed

What happened?

- User creates upload 1
- User deletes upload 1
- Worker receives "create 1"
- Worker can't find upload 1

TLA+ wants us to be explicit.

Create

```
if \Exists x \in uploads:  
    x.id = msg.id;  
then  
with upload = UploadWith(msg.id) do  
    content := content ++ upload;  
end with;  
end if
```

Create

```
if ExistsUploadWith(msg.id) then  
with upload = UploadWith(msg.id) do  
    content := content ++ upload;  
end with;  
end if
```

Model Checking Results



General

Start time:

End time:

TLC mode:

Last checkpoint time:

Current status:

Errors detected:

Fingerprint collision probability:

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:12	17	217	78	0
00:-59:-56	0	1	1	1

Coverage at

Module	Location	Count
Uploader	line 225, col 29 to line 225, col 56	78
Uploader	line 226, col 29 to line 226, col 88	78
Uploader	line 227, col 29 to line 227, col 50	78
Uploader	line 230, col 31 to line 230, col 66	37
Uploader	line 231, col 31 to line 231, col 86	37

Evaluate Constant Expression

Expression:

Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

PROPERTY NoOrphanContent:

[] < > (

\forall c \in content:

\exists u \in uploads:

c = u

)

Model Checking Results

General

Start time:	Fri Oct 05 00:31:35 CDT 2018
End time:	Fri Oct 05 00:31:38 CDT 2018
TLC mode:	Breadth-first search
Last checkpoint time:	
Current status:	Not running
Errors detected:	1 Error
Fingerprint collision probability:	

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:03	17	217	78	0
00:-59:-57	0	1	1	1

Coverage at 2018-10-05 00:31:37

Module	Location
Uploader	line 225, col 29 to
Uploader	line 226, col 29 to
Uploader	line 227, col 29 to
Uploader	line 230, col 31 to
Uploader	line 231, col 31 to

Evaluate Constant Expression

Expression:

Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

Model_1

Temporal properties were violated.

Error-Trace Exploration

Error-Trace

Name	Value
> msg	(w1 :-> [id -> 1, action -> "create"])
> next_id	2
> pc	(w1 :-> "Repeat" @@ u1 :-> "User")
> queue	<< >>
> uploads	{{from -> u1, version -> 1, id -> 1}}
▲ <Repeat line 289, col 17 to line 292, col 74 of m	State (num = 5)
> content	{{from -> u1, version -> 1, id -> 1}}
> local	(w1 :-> NULL)
> msg	(w1 :-> NULL)
> next_id	2
> pc	(w1 :-> "Receive" @@ u1 :-> "User")
> queue	<< >>
> uploads	{{from -> u1, version -> 1, id -> 1}}
▲ <User line 223, col 15 to line 238, col 52 of moc	State (num = 6)
> content	{{from -> u1, version -> 1, id -> 1}}
> local	(w1 :-> NULL)
> msg	(w1 :-> NULL)
> next_id	2
> pc	(w1 :-> "Receive" @@ u1 :-> "User")
> queue	<<[id -> 1, action -> "delete"]>>
> uploads	{}
▲ <Receive line 242, col 18 to line 255, col 68 of r	State (num = 7)
> content	{{from -> u1, version -> 1, id -> 1}}
> local	(w1 :-> NULL)
> msg	(w1 :-> [id -> 1, action -> "delete"])
> next_id	2
> pc	(w1 :-> "Delete" @@ u1 :-> "User")
> queue	<< >>
> uploads	{}
▲ <Stuttering>	State (num = 8)

Select line in Error Trace to show its value here.

▼ 🏠 <Receive line 242, col 18 to line 255, col 68 of r	State (num = 7)
> 📄 content	{[from -> u1, version -> 1, id -> 1]}
> 📄 local	(w1 :> NULL)
> 📄 msg	(w1 :> [id -> 1, action -> "delete"])
📄 next_id	2
> 📄 pc	(w1 :> "Delete" @@ u1 :> "User")
📄 queue	<< >>
📄 uploads	{}
🏠 <Stuttering>	State (num = 8)

▼ 🏠 <Receive line 242, col 18 to line 255, col 68 of r	State (num = 7)
> 📄 content	{[from -> u1, version -> 1, id -> 1]}
> 📄 local	(w1 :> NULL)
> 📄 msg	(w1 :> [id -> 1, action -> "delete"])
📄 next_id	2
> 📄 pc	(w1 :> "Delete" @@ u1 :> "User")
📄 queue	<< >>
📄 uploads	{}
🏠 <Stuttering>	State (num = 8)

What happened?

- User creates upload 1
- Worker receives "create 1"
- Worker creates content 1
- User deletes upload 1
- Worker receives "delete 1"
- Worker crashes

Fix

fair process worker \in Workers

Model Checking Results



General

Start time: Sat Oct 27 15:15:39 CDT 2018
 End time: Sat Oct 27 15:15:46 CDT 2018
 TLC mode: Breadth-first search
 Last checkpoint time:
 Current status: Not running
 Errors detected: No errors
 Fingerprint collision probability: calculated: 1.6E-16, observed: 3.1E-16

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:07	11	113	43	0
00:00:01	0	1	1	1

Coverage at 2018-10-27 15:15:46

Module	Location	Count
Uploader	line 226, col 29 to line 226, col 56	43
Uploader	line 227, col 29 to line 227, col 88	43
Uploader	line 228, col 29 to line 228, col 50	43
Uploader	line 231, col 31 to line 231, col 66	15
Uploader	line 232, col 31 to line 232, col 86	15

Evaluate Constant Expression

Expression: Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

PARAMS

Workers <- {w1, w2}

Users <- {u1}

Model Checking Results

General

Start time: Thu Oct 04 17:55:00 CDT 2018

End time: Thu Oct 04 17:55:02 CDT 2018

TLC mode: Breadth-first search

Last checkpoint time:

Current status: Not running

Errors detected: 1 Error

Fingerprint collision probability:

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:02	19	2379	752	0
00:00:01	0	1	1	1

Coverage at 2018-10-04 17:55:02

Module	Location
Uploader	line 229, col 12 to
Uploader	line 229, col 21 to
Uploader	line 229, col 30 to
Uploader	line 229, col 37 to
Uploader	line 229, col 46 to

Evaluate Constant Expression

Expression:

Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

Model_1

Temporal properties were violated.

















Error-Trace Exploration

Error-Trace

Name	Value
> ▲ <Initial predicate>	State (num = 1)
> ▲ <User line 244, col 15 to line 263, col 52 of mod>	State (num = 2)
> ▲ <Receive line 267, col 18 to line 280, col 68 of m>	State (num = 3)
> ▲ <Create line 282, col 17 to line 289, col 70 of mo>	State (num = 4)
> ▲ <User line 244, col 15 to line 263, col 52 of mod>	State (num = 5)
> ▲ <Receive line 267, col 18 to line 280, col 68 of m>	State (num = 6)
> ▲ <Edit line 291, col 15 to line 301, col 61 of modul>	State (num = 7)
> ▲ <Repeat line 318, col 17 to line 321, col 74 of mo>	State (num = 8)
> ▲ <User line 244, col 15 to line 263, col 52 of mod>	State (num = 9)
> ▲ <Receive line 267, col 18 to line 280, col 68 of m>	State (num = 10)
> ▲ <Edit line 291, col 15 to line 301, col 61 of modul>	State (num = 11)
> ▲ <PushEdit line 303, col 19 to line 307, col 65 of r>	State (num = 12)
> ▲ <Repeat line 318, col 17 to line 321, col 74 of mo>	State (num = 13)
> ▲ <User line 244, col 15 to line 263, col 52 of mod>	State (num = 14)
> ▲ <Repeat line 318, col 17 to line 321, col 74 of mo>	State (num = 15)
▲ <Stuttering>	State (num = 16)

/\ content = {}

Error-Trace

Name	Value
>  <Initial predicate>	State (num = 1)
>  <User line 244, col 15 to line 263, col 52 of modu	State (num = 2)
>  <Receive line 267, col 18 to line 280, col 68 of m	State (num = 3)
>  <Create line 282, col 17 to line 289, col 70 of mo	State (num = 4)
>  <User line 244, col 15 to line 263, col 52 of modu	State (num = 5)
>  <Receive line 267, col 18 to line 280, col 68 of m	State (num = 6)
>  <Edit line 291, col 15 to line 301, col 61 of modul	State (num = 7)
>  <Repeat line 318, col 17 to line 321, col 74 of mo	State (num = 8)
>  <User line 244, col 15 to line 263, col 52 of modu	State (num = 9)
>  <Receive line 267, col 18 to line 280, col 68 of m	State (num = 10)
>  <Edit line 291, col 15 to line 301, col 61 of modul	State (num = 11)
>  <PushEdit line 303, col 19 to line 307, col 65 of r	State (num = 12)
>  <Repeat line 318, col 17 to line 321, col 74 of mo	State (num = 13)
>  <User line 244, col 15 to line 263, col 52 of modu	State (num = 14)
>  <Repeat line 318, col 17 to line 321, col 74 of mo	State (num = 15)
 <Stuttering>	State (num = 16)

What happened?

- Oh boy...

What happened?

- User creates 1
- Worker creates 1
- User edits 1
- W1 edits 1
- User edits 1
- W2 edits 1
- W2 skips edit
- W1 pushes edit

```
if
  ExistsUploadWith(local.id)
  and UploadWith(local.id).version
    > local.version
then
  local := NULL;
  goto Edit;
else
  content := content ++ local;
  local := NULL;
end if;
```

Model Checking Results

General

Start time: Thu Oct 04 18:05:03 CDT 2018
 End time: Thu Oct 04 18:05:05 CDT 2018
 TLC mode: Breadth-first search
 Last checkpoint time:
 Current status: Not running
 Errors detected: **1 Error**
 Fingerprint collision probability:

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:02	18	1435	398	0
00:00:02	0	1	1	1

Coverage at 2018-10-04 18:05:05

Module	Location
Uploader	line 254, col 29 to
Uploader	line 255, col 29 to
Uploader	line 256, col 29 to
Uploader	line 259, col 31 to
Uploader	line 260, col 31 to

Evaluate Constant Expression

Expression: Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

TLC Errors Model_1

Temporal properties were violated.

Error-Trace Exploration

Error-Trace

Name	Value
msg	(w1 :-> [id -> 1, action -> "edit"] @@ w2 :-> [id -> 1,...
next_id	2
pc	(w1 :-> "PushEdit" @@ w2 :-> "Repeat" @@ u1 :-> "Us...
queue	<< >>
uploads	{}
<Repeat line 327, col 17 to line 330, col 74 of m	State (num = 12)
content	{}
local	(w1 :-> [from -> u1, version -> 2, id -> 1] @@ w2 :-...
msg	(w1 :-> [id -> 1, action -> "edit"] @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "PushEdit" @@ w2 :-> "Receive" @@ u1 :-> "Us...
queue	<< >>
uploads	{}
<PushEdit line 307, col 19 to line 316, col 65 of	State (num = 13)
content	{[from -> u1, version -> 2, id -> 1]}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> [id -> 1, action -> "edit"] @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "Repeat" @@ w2 :-> "Receive" @@ u1 :-> "User...
queue	<< >>
uploads	{}
<Repeat line 327, col 17 to line 330, col 74 of m	State (num = 14)
content	{[from -> u1, version -> 2, id -> 1]}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> NULL @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "Receive" @@ w2 :-> "Receive" @@ u1 :-> "Use...
queue	<< >>
uploads	{}
<Stuttering>	State (num = 15)
(w1 :-> [id -> 1, action -> "edit"] @@ w2 :-> [id -> 1, action ->	

```
fair process cleaner = "cleaner"
begin
  Clean:
    while TRUE do
      with
        id \in
          LET
            upload_ids == {u.id: u \in uploads}
            content_ids == {c.id: c \in content}
          IN
            content_ids \ upload_ids

            , exists = ContentWith(id)
          do
            content := content -- exists;
          end with;
        end while;
      end process;
end process;
```

Model Checking Results



General

Start time: Thu Oct 04 18:05:50 CDT 2018

End time: Thu Oct 04 18:05:57 CDT 2018

TLC mode: Breadth-first search

Last checkpoint time:

Current status: Not running

Errors detected: No errors

Fingerprint collision probability: calculated: 2.5E-14, observed: 2.4E-15

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
00:00:07	18	1560	398	0
00:00:01	0	1	1	1

Coverage at 2018-10-04 18:05:57

Module	Location	Count
Uploader	line 255, col 29 to line 255, col 56	398
Uploader	line 256, col 29 to line 256, col 88	398
Uploader	line 257, col 29 to line 257, col 50	398
Uploader	line 260, col 31 to line 260, col 66	161
Uploader	line 261, col 31 to line 261, col 86	161

Evaluate Constant Expression

Expression:

Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

Model Overview | Advanced Options | Model Checking Results

Model Checking Results

General

Start time: Thu Oct 04 17:56:07 CDT 2018
 End time: Thu Oct 04 17:56:10 CDT 2018
 TLC mode: Breadth-first search
 Last checkpoint time:
 Current status: Not running
 Errors detected: **1 Error**
 Fingerprint collision probability:

Statistics

State space progress (click column header for graph) Coverage at 2018-10-04 17:56:10

Time	Diameter	States Found	Distinct States	Queue Size	Module	Location
00:00:03	18	1454	376	0	Uploader	line 247, col 29 to
00:00:02	0	1	1	1	Uploader	line 248, col 29 to
					Uploader	line 249, col 29 to
					Uploader	line 252, col 31 to
					Uploader	line 253, col 31 to

Evaluate Constant Expression

Expression: Value:

User Output

TLC output generated by evaluating Print and PrintT expressions.

No user output is available

TLC Errors

Model_1

Temporal properties were violated.

Error-Trace Explanation

Error-Trace

Name	Value
<Initial predicate>	State (num = 1)
content	{}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> NULL @@ w2 :-> NULL)
next_id	1
pc	(w1 :-> "Receive" @@ w2 :-> "Receive" @@ u1 :-> "Use...
queue	<< >>
uploads	{}
<User line 245, col 15 to line 260, col 52 of moc	State (num = 2)
content	{}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> NULL @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "Receive" @@ w2 :-> "Receive" @@ u1 :-> "Use...
queue	<<[id -> 1, action -> "create"]>>
uploads	{[from -> u1, version -> 1, id -> 1]}
<Receive line 264, col 18 to line 277, col 68 of r	State (num = 3)
content	{}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> [id -> 1, action -> "create"] @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "Create" @@ w2 :-> "Receive" @@ u1 :-> "User...
queue	<< >>
uploads	{[from -> u1, version -> 1, id -> 1]}
<Create line 279, col 17 to line 286, col 70 of m	State (num = 4)
content	{[from -> u1, version -> 1, id -> 1]}
local	(w1 :-> NULL @@ w2 :-> NULL)
msg	(w1 :-> [id -> 1, action -> "create"] @@ w2 :-> NULL)
next_id	2
pc	(w1 :-> "Repeat" @@ w2 :-> "Receive" @@ u1 :-> "User...
	{[from -> u1, version -> 2, id -> 1]}

Spec Status : **parsed**

- Duplicate messages
- Dropped messages
- Permissions
- Partial failure
- Webhooks

Does it work?

Companies using TLA+


- AWS
- Amazon
- Azure
- Xbox
- eSpark Learning
- Sutori
- Elastic
- Mongo
- ING
- OSOCO
- OpenStack
- Several clients under NDA
- Like 80 blockchain companies

Conclusions

- Distributed Systems are Hard
- Specification is good
- TLA+ is good


Practical TLA+

- Me
- Beginner level
- Tons of practical examples
 - MapReduce!
 - Dekker's Algorithm!
 - Cache Invalidation!
- Out now!



Specifying Systems

- Leslie Lamport
- Intermediate level
- Canonical Text
- Covers theory



Just Hire Me Lol

- I do workshops!
- Public: Jan 14-18, <http://www.dabeaz.com/tla.html>
- Corporate: talk to me after

[hillelwayne.com/talks/
designing-distributed-
systems](https://hillelwayne.com/talks/designing-distributed-systems)

Hillel Wayne
hillelwayne.com
@hillelogram