

Symbolic search and test generation with Apalache & LLMs

igor@konnov.phd



D-CON 26 — Berlin, Mar 26, 2026

Igor Konnov



Independent security and formal methods researcher  Vienna, Austria

Verifying consensus protocols and smart contracts

Improving Apache, working on new tools

Principal research scientist at Informal Systems (Web3 – blockchains) 

Leading development of Apache and Quint

Postdoc, assistant professor & permanent researcher at TU Wien & Inria






2024-...

2019-2023

2011-2019

This talk

1. A bit of history 
2. Symbolic search with Apalache 
3. Testing loop with Apalache and LLMs 

[github.com/apalache-mc/apalache]

[apalache-mc.org]



APALACHE

A symbolic model checker for TLA+

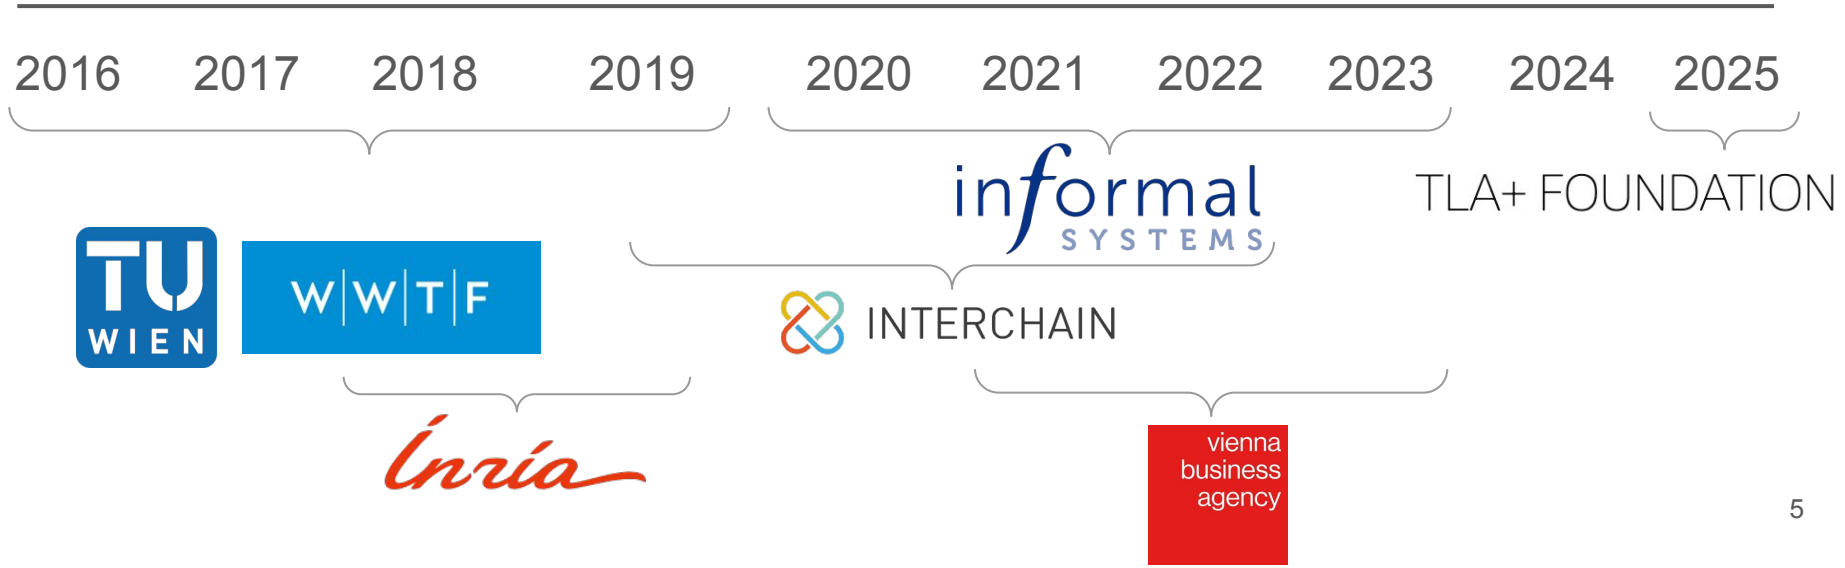
 build passing

Apalache translates [TLA+](#) into the logic supported by SMT solvers such as [Microsoft Z3](#). Apalache can check inductive invariants (for fixed or bounded parameters) and check safety of bounded executions (bounded model checking). To see the list of supported TLA+ constructs, check the [supported features](#). In general, Apalache runs under the same assumptions as [TLC](#). However, Apalache benefits from constraint solving and can handle potentially larger state-spaces, e.g., involving integer clocks and Byzantine faults.

To learn more about TLA+, visit [Leslie Lamport's page on TLA+](#) and his [Video course](#).



APALACHE



Epoch 1

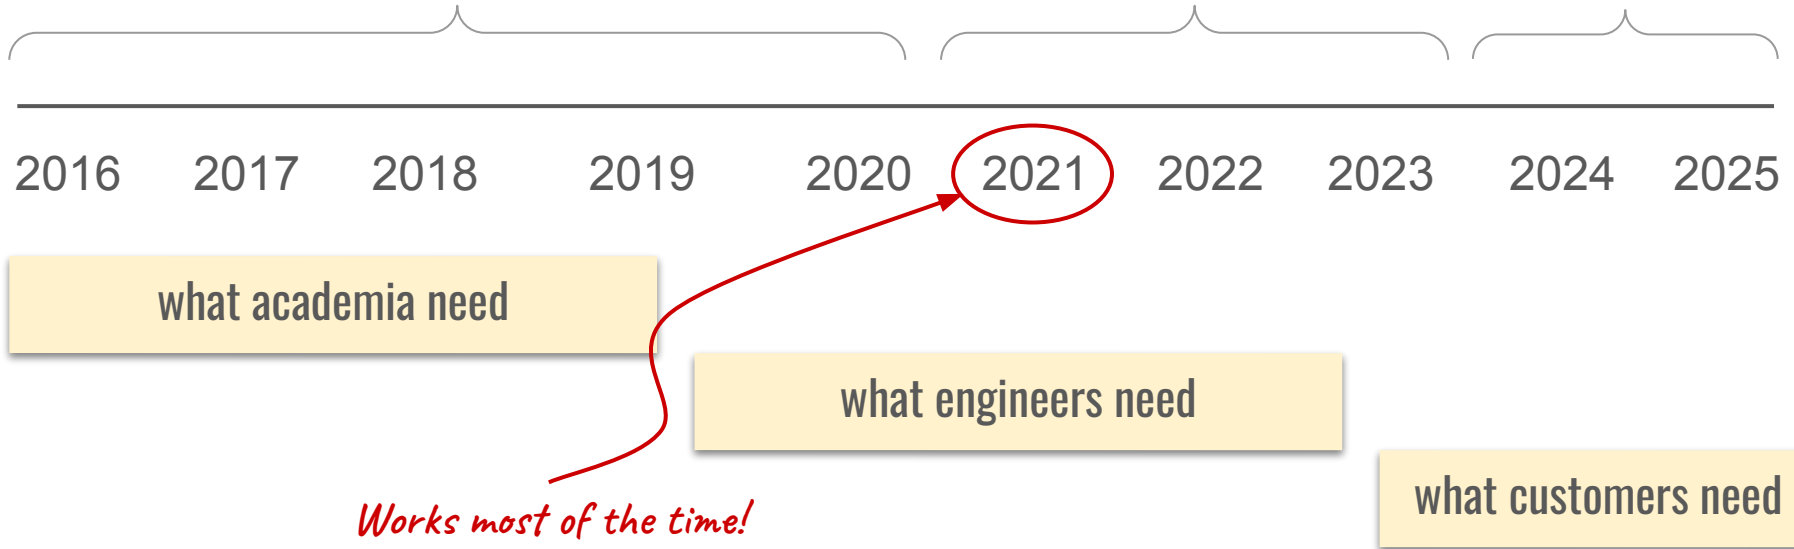
1. SMT transpiler
2. Type checker
3. Bounded model checker

Epoch 2

1. Random. symbolic execution
2. Data generators

Epoch 3

1. Parallelization
2. DIY search
3. Test harnesses



W|W|T|F



Thanh-Hai Tran



Marijana Lazić



Josef Widder



vienna
business
agency



Jure Kukovec



Shon Feder



Gabriela Moreira
@bugarela



Igor Konnov
@konnov



Thomas Pani



Andrey Kupriyanov



Philip Offtermatt



Rodrigo Otoni

informal
SYSTEMS



INTERCHAIN

Inria

Università
della
Svizzera
italiana



Governance Protocol (2024)

ChonkyBFT consensus (2024-2025)




Aztec Governance Protocol (2025)



Accountability in Ethereum 3-slot finality (2024)



TetraBFT consensus (2024)



Tendermint BFT consensus (2020)

Tendermint light client (2020)



Ben-Or's consensus (2024)

Offshoot projects

[quint-lang.org]

[github.com/informalsystems/atomkraft]

Quint

A modern and executable specification language

Quint helps you write precise specifications and check them automatically. Find subtle bugs before they reach production.

Get Started 5-min guide



✓ Executable	
Quint	English & Markdown
✓ checked names and types	✗ not checked
✓ executable	✗ not executable
✓ Abstract	
Quint	Programming Languages
✓ define only what matters	✗ define how things happen, in detail
✓ Modern	
Quint	Existing Spec Languages

informal
SYSTEMS

since 2021

Atomkraft: E2E testing for Cosmos blockchains

Below we describe what Atomkraft is about, and explain the concepts behind the tool. In case you skip that, and jump directly into action, please read our [Installation guide](#) to install the tool; afterwards following either our [Cosmos SDK Token Transfer tutorial](#), or [CosmWasm tutorial](#).

We cover the following topics in this file:

```
Modelator: model-based testing framework for TLA+
```

```
Options
```

```
--log-level TEXT [default: None]
--help Show this message and exit.
```

```
Commands
```

```
apache Check whether the JAR file
check Check that the invariants hold in the
info Display information on the loaded model
load Load a TLA+ model file and parse it
reload Reload model and configuration files
reset Removes any loaded model.
sample Generate execution traces that reach
simulate
typecheck
version
```

no LLMs back then

Quick glance at TLA⁺

<https://protocols-made-fun.com>

Another Advantage of Free Choice: Completely Asynchronous Agreement Protocols

(Extended Abstract)

Michael Ben-Or †

Focus on safety when:

$$N > 5 \cdot t \wedge t \geq f$$

Termination needs probabilities

out of scope

B — Byzantine Protocol

Process P : Initial value x_P .

step 0: set $r := 1$.

step 1: Send the message $(1, r, x_P)$ to all the processes.

step 2: Wait till messages of type $(1, r, *)$ are received from $N - t$ processes. If more than $(N + t)/2$ messages have the same value v , then send the message $(2, r, v, D)$ to all processes. Else send the message $(2, r, ?)$ to all processes.

step 3: Wait till messages of type $(2, r, *)$ arrive from $N - t$ processes.

(a) If there are at least $t + 1$ D -messages $(2, r, v, D)$, then set $x_P := v$.

(b) If there are more than $(N + t)/2$ D -messages then **decide** v .

(c) Else set $x_P = 1$ or 0 each with probability $\frac{1}{2}$.

step 4: Set $r := r + 1$ and go to step 1.

```

78 ∨ Init ≜
79   \* non-deterministically choose the initial values
80   ∧ value ∈ [ CORRECT → VALUES ]
81   ∧ decision = [ r ∈ CORRECT ↦ NO_DECISION ]
82   ∧ round = [ r ∈ CORRECT ↦ 1 ]
83   ∧ step = [ r ∈ CORRECT ↦ S1 ]
84   ∧ msgs1 = [ r ∈ ROUNDS ↦ {} ]
85   ∧ msgs2 = [ r ∈ ROUNDS ↦ {} ]

```

Process P : Initial value x_P .

step 0: set $r := 1$.

step 1: Send the message $(1, r, x_P)$ to all the processes.

```

103 \* @type: REPLICA => Bool;
104 Step1(id) ≜
105   LET r ≜ round[id] IN
106   ∧ step[id] = S1
107   \* "send the message (1, r, x_P) to all the processes"
108   ∧ msgs1' = [msgs1 EXCEPT ![r] = @ u { M1(id, r, value[id]) } ]
109   ∧ step' = [step EXCEPT ![id] = S2]
110   ∧ UNCHANGED ( value, decision, round, msgs2 )

```

step 2: Wait till messages of type $(1, r, *)$ are received from $N - t$ processes. If more than $(N + t)/2$ messages have the same value v , then send the message $(2, r, v, D)$ to all processes. Else send the message $(2, r, ?)$ to all processes.

```

112 Step2(id)  $\triangleq$ 
113   LET r  $\triangleq$  round[id] IN
114    $\wedge$  step[id] = S2
115    $\wedge$   $\exists$  received  $\in$  SUBSET msgs1[r]:
116     \* "wait till messages of type (1, r, *) are received from N - T processes"
117      $\wedge$  Cardinality(Senders1(received))  $\geq$  N - T
118      $\wedge$  LET Weights  $\triangleq$  [ v  $\in$  VALUES  $\mapsto$ 
119       Cardinality(Senders1({ m  $\in$  received: m.v = v }) ) ]
120     IN
121      $\exists$  v  $\in$  VALUES:
122       \* "if more than (N + T)/2 messages have the same value v..."
123        $\wedge$  2 * Weights[v] > N + T
124       \* "...then send the message (2, r, v, D) to all processes"
125        $\wedge$  msgs2' = [msgs2 EXCEPT ![r] = @ U { D2(id, r, v) } ]
126        $\vee$   $\forall$  v  $\in$  VALUES: 2 * Weights[v]  $\leq$  N
127       \* "Else send the message (2, r, ?) to all processes"
128        $\wedge$  msgs2' = [msgs2 EXCEPT ![r] = @ U { Q2(id, r) } ]
129    $\wedge$  step' = [ step EXCEPT ![id] = S3 ]
130    $\wedge$  UNCHANGED ( value, decision, round, msgs1 )

```

step 2: Wait till messages of type $(1, r, *)$ are received from $N - t$ processes. If more than $(N + t)/2$ messages have the same value v , then send the message $(2, r, v, D)$ to all processes. Else send the message $(2, r, ?)$ to all processes.

```

112 Step2(id)  $\triangleq$ 
113   LET  $r \triangleq$  round[id] IN
114    $\wedge$  step[id] = S2
115    $\wedge$   $\exists$  received  $\in$  SUBSET msgs1[r]:
116     \* "wait till messages of type (1, r, *) are received from N - T processes"
117      $\wedge$  Cardinality(Senders1(received))  $\geq$  N - T
118      $\wedge$  LET Weights  $\triangleq$  [  $v \in$  VALUES  $\mapsto$ 
119       Cardinality(Senders1({ m  $\in$  received: m.v = v }))) ]
120     IN
121      $\boxed{\forall}$   $\exists v \in$  VALUES:
122       \* "if more than (N + T)/2 messages have the same value v..."
123        $\wedge$  2 * Weights[v] > N + T
124       \* "...then send the message (2, r, v, D) to all processes"
125        $\wedge$  msgs2' = [msgs2 EXCEPT ![r] = @  $\boxed{U}$  { D2(id, r, v) }]
126      $\vee \wedge \forall v \in$  VALUES: 2 * Weights[v]  $\leq$  N
127       \* "Else send the message (2, r, ?) to all processes"
128        $\wedge$  msgs2' = [msgs2 EXCEPT ![r] = @  $\boxed{U}$  { Q2(id, r) } ]
129    $\wedge$  step' = [ step EXCEPT ![id] = S3 ]
130    $\wedge$  UNCHANGED ( value, decision, round, msgs1 )

```

Modeling faults

Works well for bounded model checking (Apalache)

```
87 ∨ InitWithFaults ≐
88   \* non-deterministically choose the initial values
89   ∧ value ∈ [ CORRECT → VALUES ]
90   ∧ decision = [ r ∈ CORRECT ↦ NO_DECISION ]
91   ∧ round = [ r ∈ CORRECT ↦ 1 ]
92   ∧ step = [ r ∈ CORRECT ↦ S1 ]
93   \* non-deterministically initialize the messages with faults
94 ∨ ∧ ∃ F1 ∈ SUBSET [ src: FAULTY, r: ROUNDS, v: VALUES ]:
95     | msgs1 = [ r ∈ ROUNDS ↦ { m ∈ F1: m.r = r } ]
96 ∨ ∧ ∃ F1D ∈ SUBSET [ src: FAULTY, r: ROUNDS, v: VALUES ],
97     | F1Q ∈ SUBSET [ src: FAULTY, r: ROUNDS ]:
98 ∨     | msgs2 = [ r ∈ ROUNDS ↦
99 ∨     |     { D2(mm.src, r, mm.v): mm ∈ { m ∈ F1D: m.r = r } }
100     |     ∪ { Q2(mm.src, r): mm ∈ { m ∈ F1Q: m.r = r } }
101     | ]
```

Transition predicate

```
174 CorrectStep  $\triangleq$   
175    $\exists$  id  $\in$  CORRECT:  
176      $\forall$  Step1(id)  
177      $\forall$  Step2(id)  
178      $\forall$  Step3(id)
```

State invariant

```
186  $\backslash*$  No two correct replicas decide on different values  
187 AgreementInv  $\triangleq$   
188    $\forall$  id1, id2  $\in$  CORRECT:  
189      $\forall$  decision[id1] = NO_DECISION  
190      $\forall$  decision[id2] = NO_DECISION  
191      $\forall$  decision[id1] = decision[id2]
```

Apalache as a Solver

Type checker

apalache-mc.org/docs/tutorials/snowcat-tutorial.html

Damas & Milner type inference + row types (no inductive types)

Resolving type imprecision between function-like types

May require type annotations for records, tuples, functions, and sequences

<code>Int</code>	<code>Bool</code>	<code>Str</code>
<code>UNINTERPRETED</code>	<code>Set(a)</code>	<code>Seq(a)</code>
<code>a -> b</code>	<code><<a, b, c>></code>	<code>{ f1: a, f2: b, f3: c }</code>
<code>(a, b, c) => d</code>		<code>Tag1(a) Tag2(b) Tag3(c)</code>

Translation to SMT

TLA⁺ Model Checking Made Symbolic [OOPSLA'19]

Mimic the semantics implemented by TLC – explicit model checker


Compute **layout of data structures**, constrain **contents with SMT**

Define operational semantics via reduction rules – **for bounded data structures**

Trade efficiency for expressivity

Static picture of TLA⁺ values and relations between them

Arena:



SMT:

integer

sort Int

Boolean

sort Bool


name, e.g., "abc", uninterpreted sort

finite set:

- a constant c of uninterpreted sort set_{τ}
- propositional constants for members

$in_{\langle c_1, c \rangle}, \dots, in_{\langle c_n, c \rangle}$

Arenas for sets: $\{\{1, 2\}, \{2, 3\}\}$



SMT defines the contents, e.g., to get $\{\{1\}, \{2\}\}$:


$$in_{\langle c_1, c_4 \rangle} \wedge \neg in_{\langle c_2, c_4 \rangle} \wedge in_{\langle c_2, c_5 \rangle} \wedge \neg in_{\langle c_3, c_5 \rangle}$$

Rewriting the set construction

$\{1, 2\} \rightsquigarrow \{c_1, 2\} \rightsquigarrow \{c_1, c_2\} \rightsquigarrow c_3$

Corresponding arena

$(\text{empty}) \rightsquigarrow c_1 : \text{Int} \rightsquigarrow c_1 : \text{Int} \quad c_2 : \text{Int} \rightsquigarrow$




Rewriting the set construction

$\{1, 2\} \rightsquigarrow \{c_1, 2\} \rightsquigarrow \{c_1, c_2\} \rightsquigarrow c_3$

Corresponding arena

$(\text{empty}) \rightsquigarrow c_1 : \text{Int} \rightsquigarrow c_1 : \text{Int} \quad c_2 : \text{Int} \rightsquigarrow$

$c_3 : \text{Set}[\text{Int}]$



`(declare-const C5 Int)`

`(assert (= C5 1))`

`(declare-const C6 Int)`

`(assert (= C6 2))`

`(declare-sort Cell_Si 0)`

`(declare-const C7 Cell_Si)`

`(declare-const in_i5_Si7 Bool)`

`(declare-const in_i6_Si7 Bool)`


`(assert in_i5_Si7)`

`(assert in_i6_Si7)`


Quantifier-free encoding!

Symbolic search with Apache

Example: Circular buffer




Example: Circular buffer



1. GET \rightarrow 33, $\text{head}' = (\text{head} + 1) \% N$

Example: Circular buffer



1. GET \rightarrow 33, $\text{head}' = (\text{head} + 1) \% N$
2. PUT(55), $\text{tail}' = (\text{tail} + 1) \% N$

```

9  CONSTANTS
10  \* Size of the circular buffer.
11  \* @type: Int;
12  BUFFER_SIZE,
13  \* The set of possible buffer elements.
14  \* @type: Set(Int);
15  BUFFER_ELEMS
16
17  ASSUME BUFFER_SIZE > 0
18
19  ▾ VARIABLES
20  \* The integer buffer of size BUFFER_SIZE.
21  \* @type: Int -> Int;
22  buffer,
23  \* Index of the next element to POP.
24  \* @type: Int;
25  head,
26  \* Index of the next free slot for PUSH.
27  \* @type: Int;
28  tail,
29  \* Number of elements currently stored.
30  \* @type: Int;
31  count

```

```

41  Put(x) ==
42  Put::
43  LET nextTail == (tail + 1) % BUFFER_SIZE IN
44  /\ buffer' = [buffer EXCEPT ![tail] = x]
45  /\ head' = head
46  /\ tail' = nextTail
47  /\ count' = count + 1

```

```

50  Get ==
51  Get::
52  LET nextHead == (head + 1) % BUFFER_SIZE IN
53  /\ count > 0
54  /\ UNCHANGED buffer
55  /\ head' = nextHead
56  /\ tail' = tail
57  /\ count' = count - 1

```

```

41 Put(x) ==
42   Put::
43   LET nextTail == (tail + 1) % BUFFER_SIZE IN
44   /\ buffer' = [buffer EXCEPT ![tail] = x]
45   /\ head' = head
46   /\ tail' = nextTail
47   /\ count' = count + 1

50 Get ==
51   Get::
52   LET nextHead == (head + 1) % BUFFER_SIZE IN
53   /\ count > 0
54   /\ UNCHANGED buffer
55   /\ head' = nextHead
56   /\ tail' = tail
57   /\ count' = count - 1

33  /* Initial state
34  Init ==
35   /\ buffer = [ i \in 0..(BUFFER_SIZE - 1) |-> 0 ]
36   /\ head = 0
37   /\ tail = 0
38   /\ count = 0

59  /* Either Put or Get may happen in any step.
60  Next ==
61   \/ \E x \in BUFFER_ELEMS:
62   |   Put(x)
63   \/ Get

```

```

70  /* Safety property we *intend* to hold, but it is violated:
71  /* count must never exceed the buffer capacity.
72  SafeInv == count <= BUFFER_SIZE

```

```

41 Put(x) ==
42   Put::
43   LET nextTail == (tail + 1) % BUFFER_SIZE IN
44   /\ buffer' = [buffer EXCEPT ![tail] = x]
45   /\ head' = head
46   /\ tail' = nextTail
47   /\ count' = count + 1

50 Get ==
51   Get::
52   LET nextHead == (head + 1) % BUFFER_SIZE IN
53   /\ count > 0
54   /\ UNCHANGED buffer
55   /\ head' = nextHead
56   /\ tail' = tail
57   /\ count' = count - 1

33  /* Initial state
34  Init ==
35   /\ buffer = [ i \in 0..(BUFFER_SIZE - 1) |-> 0 ]
36   /\ head = 0
37   /\ tail = 0
38   /\ count = 0

59  /* Either Put or Get may happen in any step.
60  Next ==
61   \/ \E x \in BUFFER_ELEMS:
62   |   Put(x)
63   \/ Get

70  /* Safety property we *intend* to hold, but it is violated:
71  /* count must never exceed the buffer capacity.
72  SafeInv == count <= BUFFER_SIZE

```

How do we find this violation?

Instance `MC10u8_BuggyCircularBuffer`:

- `BUFFER_SIZE` = 10
- `BUFFER_ELEMS` = 0..255

Explicit-state model checker TLC:

over 40 min on 30 CPU cores

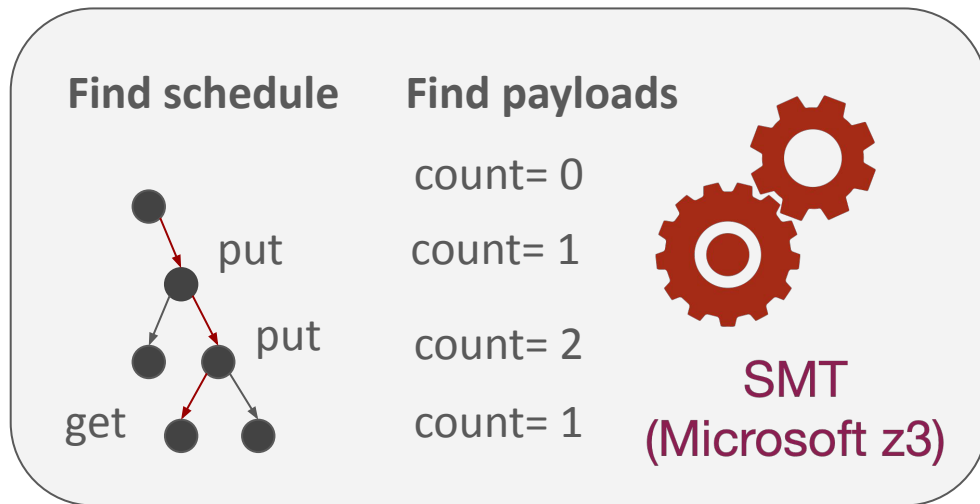
over 3 billion states

400G disk space ⇒ “No space left on device”


Can we analyze it without
enumerating all states?

Epoch #1: Bounded model checking (symbolic)

Explore all paths up to k steps, e.g., 20



Kind of symbolic
breadth-first search



Apalache is fast?

Bounded model checking with Apalache

```
$ apalache-mc check
```

3 sec to invariant violation (11 steps)

Explicit-state model checker TLC:

over 40 min on 30 CPU cores

over 3 billion states

400G disk space ⇒ “No space left on device”

(we could do trivial abstraction)

General perception: It is slow

Increasing the number of checkable steps would improve usability #1571

danwt started this conversation in General



danwt on Mar 28, 2022

Problem description

In my experience it can be hard to check invariants, and medium/large models can be quite limited.

ITF JSON as init TLA s

Open



rnbguy opened on Mar 3, 2023

Is your feature request related to a problem? Please describe

I want to restart model exploration from an ITF trace. I have a ITF JSON state, but I don't have the `init` operator in the specification.

Describe the solution you'd like

I would like a checker argument which I can use to pass the ITF JSON string - representing the concrete value assignment of the model state.

Use Z3 in parallel mode #2315

Closed

#2990



rnbguy opened on

Is your feature request related to a problem? Please describe

It would be nice to

For one, I want to

Describe the solution you'd like

Open



rnbguy opened on Feb 16, 2023 · edited by rnbguy

Edits Contributor

Is your feature request related to a problem? Please describe

Apache simulate still slow if the model is too big.

Describe the solution you'd like

When I use Apache to simulate random traces, but I don't care about an invariant being broken, it makes sense to split the traces into multiple short lengths and solve them separately.

We only have to make sure that the `init` operator of each run is the last state operator of the previous run.

Category



General

splitting trace length to speed up apache simulate #2419

But it works for inductive invariants

1. $Init \Rightarrow IndInv$
2. $IndInv \Rightarrow StateInv$
3. $IndInv \wedge Next \Rightarrow IndInv'$


TLA⁺ is well-suited for that

Aztec Governance Protocol (2025)

ChonkyBFT consensus (2024-2025)

Ben-Or's consensus (2024)

Time vs BUFFER_SIZE (log scale)



In 2022, after plenty of discussions with engineers:

I wrote a randomized simulator in TypeScript à la PBT

...and everybody was: “It’s really fast!”

Some of my peers use exclusively this simulator



```
$ quint run --main=mc10u8_buggy_circular_buffer --invariant=safeInv \  
  --max-steps=200 buggy_circular_buffer.qnt  
...  
[violation] Found an issue (16ms at 63 traces/second).
```

[quint-lang.org]

Random search in TLC

TLC `-simulate -depth 100`:

Finds a counterexample in **0.25 sec!**

Needs to explore only 174K states

Under 1 sec for `BUFFER_SIZE` in { **20, 30, 100** }

All you need is a simulator?

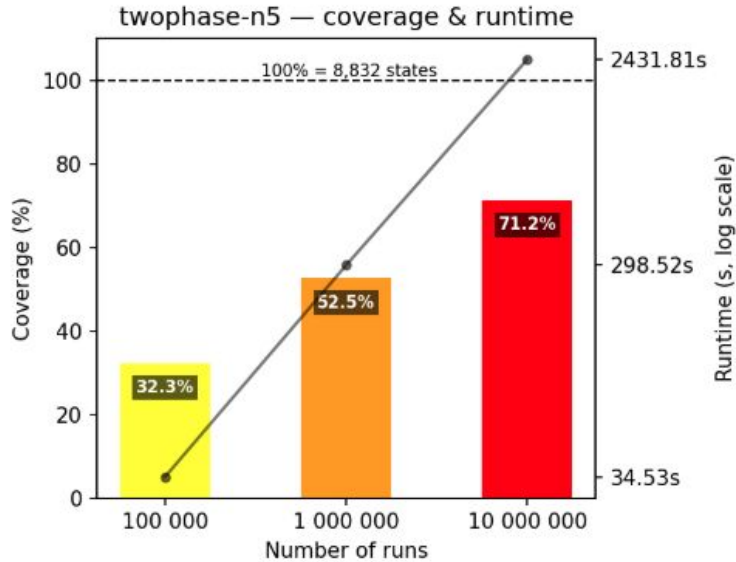


Figure 2.a: Two-phase commit, 5 RMs.

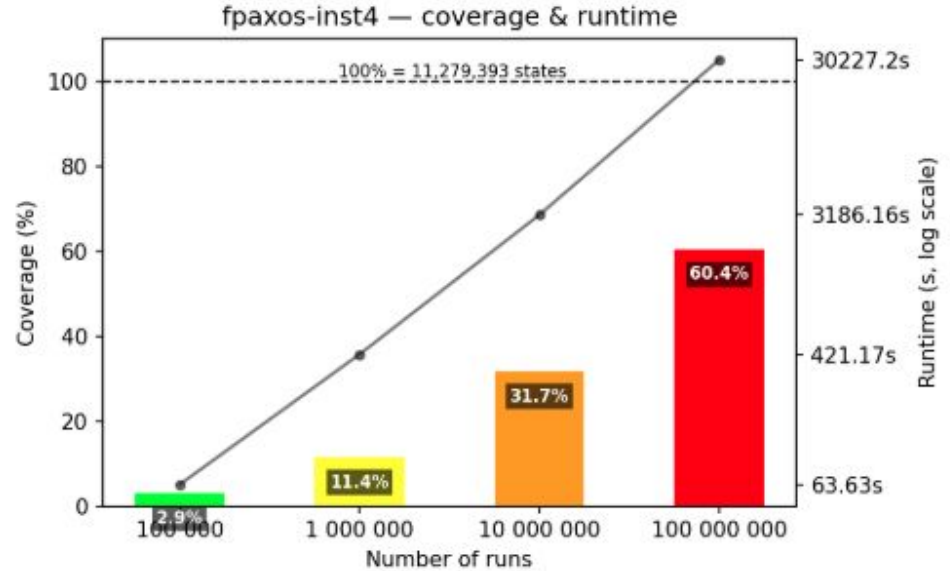



Figure 2.d: FPaxos, 4 acceptors.

<https://protocols-made-fun.com>

By fixing the search procedure,
we are missing easy bugs!

By fixing the search procedure, we are missing easy bugs!




Bounded model checking

Epoch #2: Randomized symbolic execution

Find schedule

Find payloads



42

22

42




SMT
(Microsoft z3)

Bounded model checking

Generate
schedule


Find payloads



42

22

42




SMT
(Microsoft z3)

Randomized symbolic execution

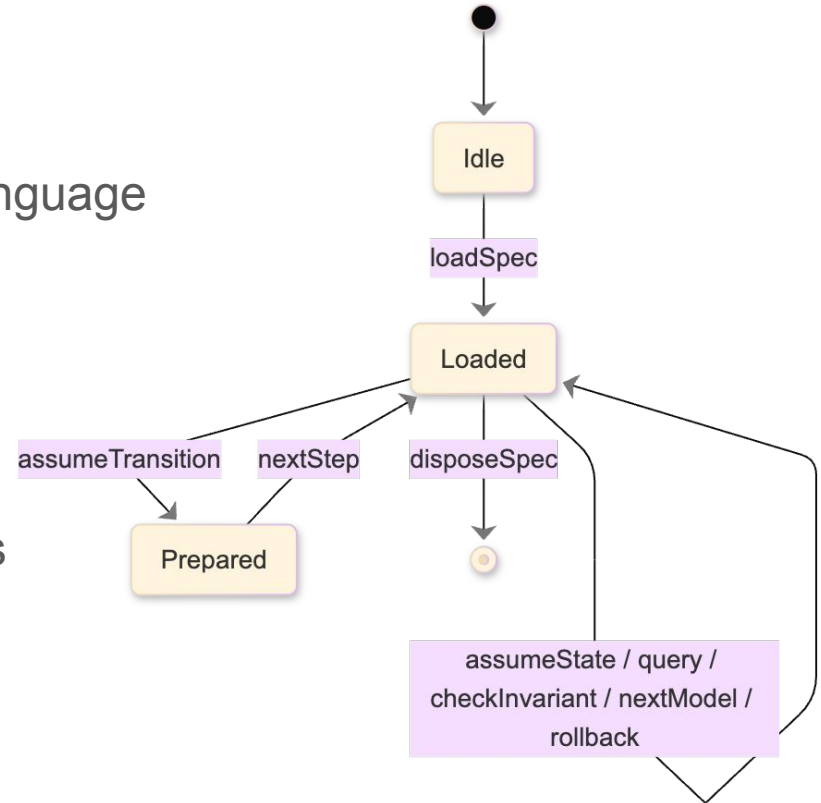
apache-mc simulate

- ✓ Finds bugs faster than bounded model checker
- ✓ Easier for Z3, as it is mostly doing propagation
- ✓ Trivially parallelizable
- ✗ Hard to customize



Epoch #3: DIY search

- ✓ Write your own search scripts in any language
- ✓ Prioritize schedules as you like
- ✓ Evaluate them with Apalache and Z3
- ✓ Query for traces and enumerate models



Testing TFTP

TFTP

Simple protocol to send and receive files over UDP since 1992

Network booting, transferring firmware to network devices

Network Working Group
Request For Comments: 1350
STD: 33
Obsoletes: RFC [783](#)

K. Sollins
MIT
July 1992

THE TFTP PROTOCOL (REVISION 2)

Status of this Memo

This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Standards" for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Summary

TFTP is a very simple protocol used to transfer files. It is from this that its name comes, Trivial File Transfer Protocol or TFTP. Each nonterminal packet is acknowledged separately. This document describes the protocol and its types of packets. The document also explains the reasons behind some of the design decisions.

Network Working Group
Request for Comments: 2347
Updates: [1350](#)
Obsoletes: [1782](#)
Category:

Status of

This d
Intern

Network Working Group
Request for Comments: 2348
Updates: [1350](#)
Obsoletes: [1783](#)
Category: Standards Track

Status of this Memo

TFTP Blocksize Option

Copyri

Abstract

The Tr
transf
remote
allow

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

The Trivial File Transfer Protocol [1] is a simple, lock-step, file transfer protocol which allows a client to get or put a file onto a remote host. One of its primary uses is the booting of diskless nodes on a Local Area Network. TFTP is used because it is very simple to implement in a small node's limited ROM space. However, the choice of a 512-octet blocksize is not the most efficient for use on a LAN whose MTU may 1500 octets or greater.

Network Working Group
Request for Comments: 2349
Updates: [1350](#)
Obsoletes: [1784](#)
Category: Standards Track

G. Malkin
Bay Networks
A. Hazkin
Hewlett Packard Co.
May 1998

TFTP Timeout Interval and Transfer Size Options

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

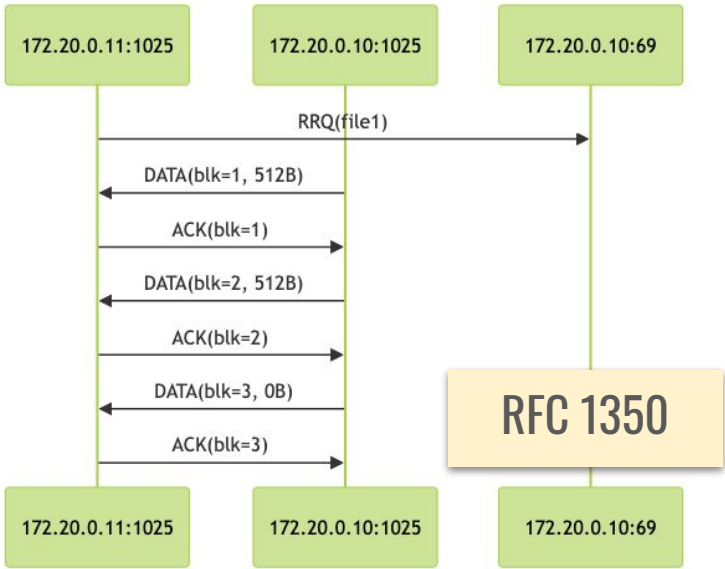
Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

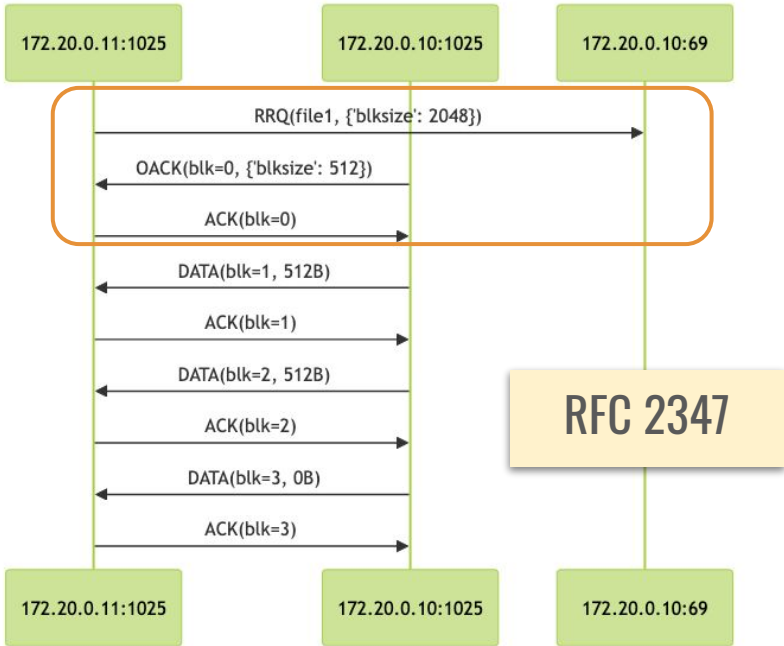
The Trivial File Transfer Protocol [1] is a simple, lock-step, file transfer protocol which allows a client to get or put a file onto a remote host.

RFCs 1350, 2347-2349, 1123

Happy paths are easy



only read requests in this talk



TFTP challenges for verification tools

1. **UDP**: packet loss, reordering, retransmission
2. **Large parameter space**: 64K of ports, block sizes, file sizes
3. Retransmission **timeouts**
4. Some operations are **input/output**
5. Some are **one-sided** (errors, termination)

TFTP is only called trivial

Specification in TLA⁺

Initial spec

Client: 5 actions

Server: 5 actions

Environment: 1 action

Size: about 800 LOC

```
Next ==
  \* the actions by the clients
  \V \E srcIp \in CLIENT_IPS, srcPort \in PORTS:
    \E filename \in DOMAIN FILES, timeout \in 1..255:
      \* "man tftpd": 65464 is the theoretical maximum for block size
      \* https://linux.die.net/man/8/tftpd
      \E tsize \in 0..FILES[filename], blksize \in 0..65464:
        \* choose a subset of the options to request
        \E optionKeys \in SUBSET OPTIONS_RFC2349:|
          LET options ==
            mk_options(optionKeys, blksize, tsize, timeout)
          IN
            ClientSendRRQ(srcIp, srcPort, filename, options)
  \V \E udp \in packets:
    \V ClientRecvDATA(udp)
    \V ClientRecvOACK(udp)
    \V ClientRecvErrorAndCloseConn(udp)
  \V \E ipPort \in DOMAIN clientTransfers:
    ClientTimeout(ipPort)
  \* the server
  \V \E udp \in packets:
    \V ServerRecvRRQ(udp)
    \V ServerSendDATA(udp)
    \V ServerRecvAckAndCloseConn(udp)
    \V ServerRecvErrorAndCloseConn(udp)
  \V \E ipPort \in DOMAIN serverTransfers:
    ServerTimeout(ipPort)
  \* handle the clock and timeouts
  \V \E delta \in 1..255:
    AdvanceClock(delta)
```

Example: ClientSendRRQ

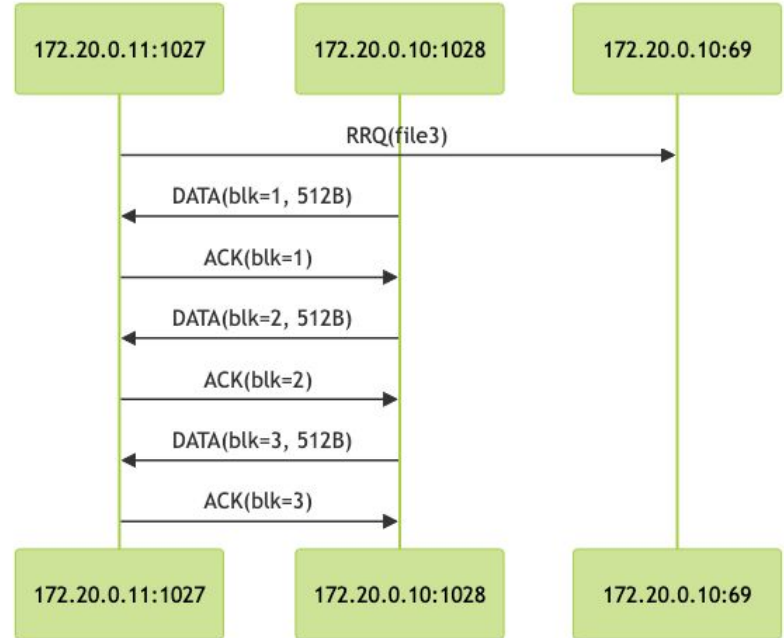
```
\* A client sends a read request to the server.
\* @type: (Str, Int, Str, Str -> Int) => Bool;
ClientSendRRQ(_srcIp, _srcPort, _filename, _options) ==
  \* We only specify the "octet" mode:
  \* "mail" mode is obsolete as per RFC 1350, what is "net
ClientSendRRQ::
LET rrq == RRQ(_filename, "octet", _options)
  udp == [srcIp |-> _srcIp,
          srcPort |-> _srcPort,
          destIp |-> SERVER_IP,
          destPort |-> 69, \* the port is fixed as per
          payload |-> rrq]
IN
/\ <<_srcIp, _srcPort>> \notin DOMAIN clientTransfers
\* RFC-2349: In Read Request packets, a size of "0" is s
/\ "tsize" \in DOMAIN _options => _options["tsize"] = 0
```

```
/\ clientTransfers' = [
  p \in DOMAIN clientTransfers \union {<<_srcIp, _srcPort>>} |->
  IF p = <<_srcIp, _srcPort>>
  THEN [ port |-> 69, \* initial port before negotiation
        tsize |-> 0,
        blksize |-> get_or_else(_options, "blksize", 512),
        timeout |-> get_or_else(_options, "timeout", 1),
        blocks |-> <<>>,
        blockNum |-> 0,
        timestamp |-> clock,
        transferred |-> 0,
        proto |->
          IF DOMAIN _options = {}
          THEN PROTO_OPTIONS_NO ELSE PROTO_OPTIONS_YES
        ]
  ELSE clientTransfers[p]
]
/\ packets' = packets \union {udp}
/\ lastAction' = ActionClientSendRRQ(udp)
/\ UNCHANGED <<serverTransfers, clock>>
```

Make sure the spec does something

```
$ apache-mc check \  
  --inv=RecvThreeDataBlocksEx MC2_ftp.tla  
...  
State 7: state invariant 0 violated.  
Total time: 24.813 sec
```

```
/* Check this falsy invariant to see an example  
RecvThreeDataBlocksEx ==  
  ~(\E p \in DOMAIN clientTransfers:  
    Len(clientTransfers[p].blocks) >= 3)
```



Initial specification effort

1. Relatively straightforward with my experience
2. Most of my time went into switching between RFCs
3. Generating repetitive type definitions with LLMs
4. Playing with falsy invariants to produce examples



[\[github.com/konnov/tftp-symbolic-testing\]](https://github.com/konnov/tftp-symbolic-testing)

13 hours \approx 2 days

Question #1 from every engineer:

**How does it connect to the
implementation?**

Tester-implementation game

Player 1 (tester)

1. Guess next transition with inputs
3. Evaluate the response. Repeat 1.

Player 2 (SUT)

2. Execute the transition. Respond.


Multishot

input-output conformance testing

Challenges

1. Completely random generation is really low quality
2. Produce high-quality inputs
3. Develop the harness and testing infrastructure

1. Completely random generation is really low quality
2. Produce high-quality inputs



3. Develop the harness and testing infrastructure



Give me a framework
in Golang, Rust, etc.

I don't want to learn
TLA⁺ and Python

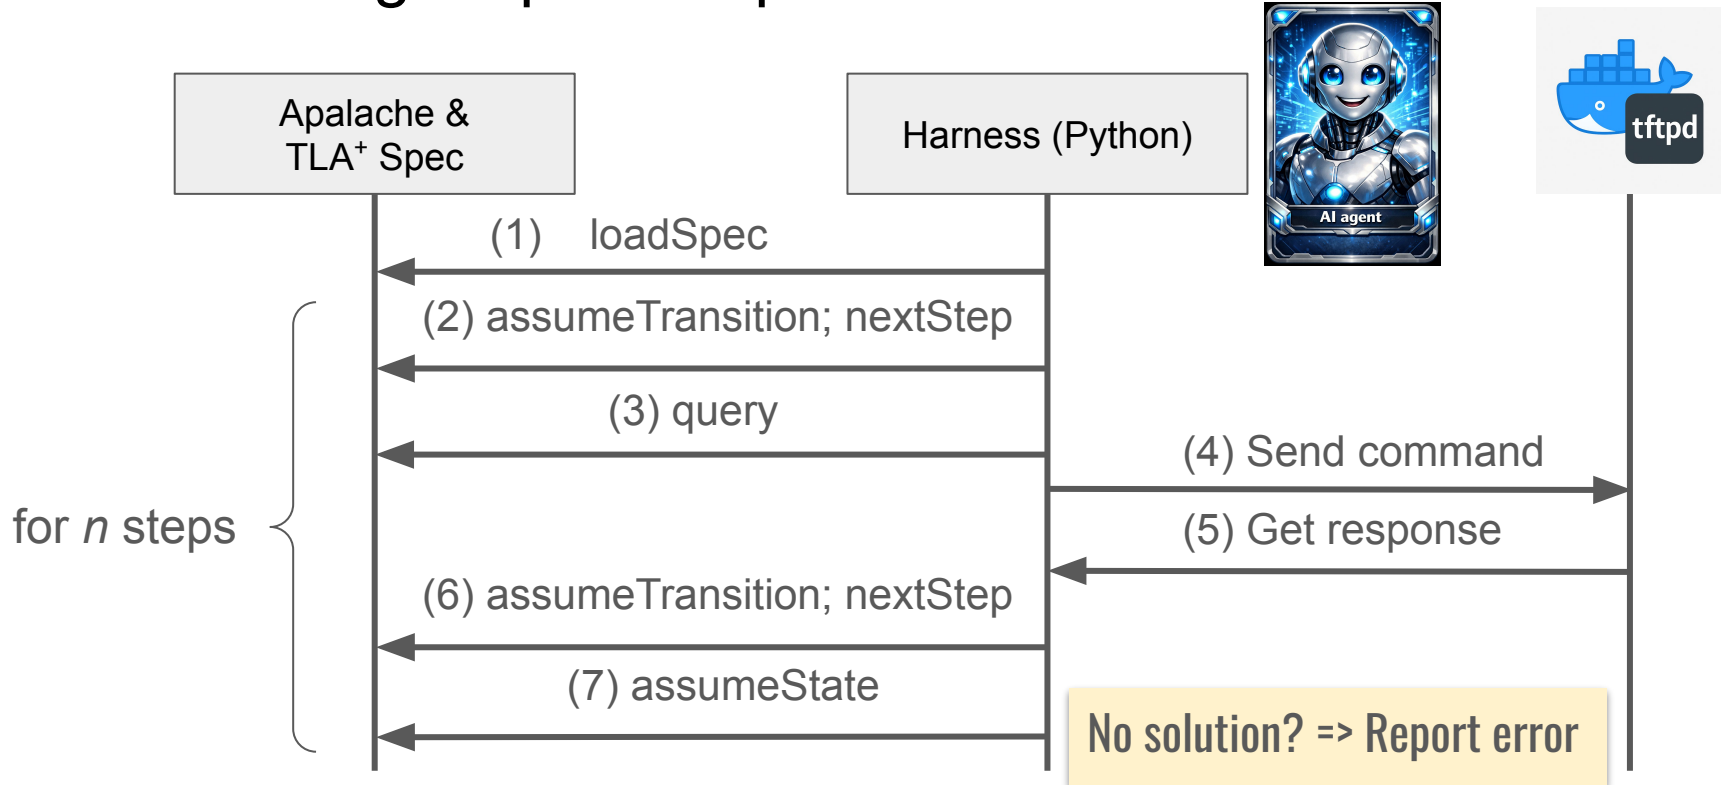
I'm coding...



Sure, I will help you to
write the test harness!

(You have to fix my
bugs later. He-he)

New testing loop with Apalache



prompt-test-harness.md > ## Requirement 2

1 You are a protocol testing engineer. Your goal is to use the TLA+ specification
2 of TFTP to create a test harness that produces new tests by analyzing symbolic
3 executions of the protocol via JSON RPC of Apache (see [client.py]).

5 Below are the requirements of this project. Do not change this text unless
6 explicitly instructed to do so.

8 ## Requirement 0

10 The TLA+ specification of the TFTP protocol is provided in the [spec] file.
11 The test harness should utilize this specification to generate symbolic
12 execution traces using Apache. You are not allowed to modify
13 specification, unless you are explicitly instructed to do so.


15 ## Requirement 1

17 The test harness must be implemented in Python and should interact
18 using JSON RPC to retrieve symbolic execution traces of the TFTP
19 Python script is located at [test-harness](./test-harness). The harness
20 should use Poetry to pull in dependencies.

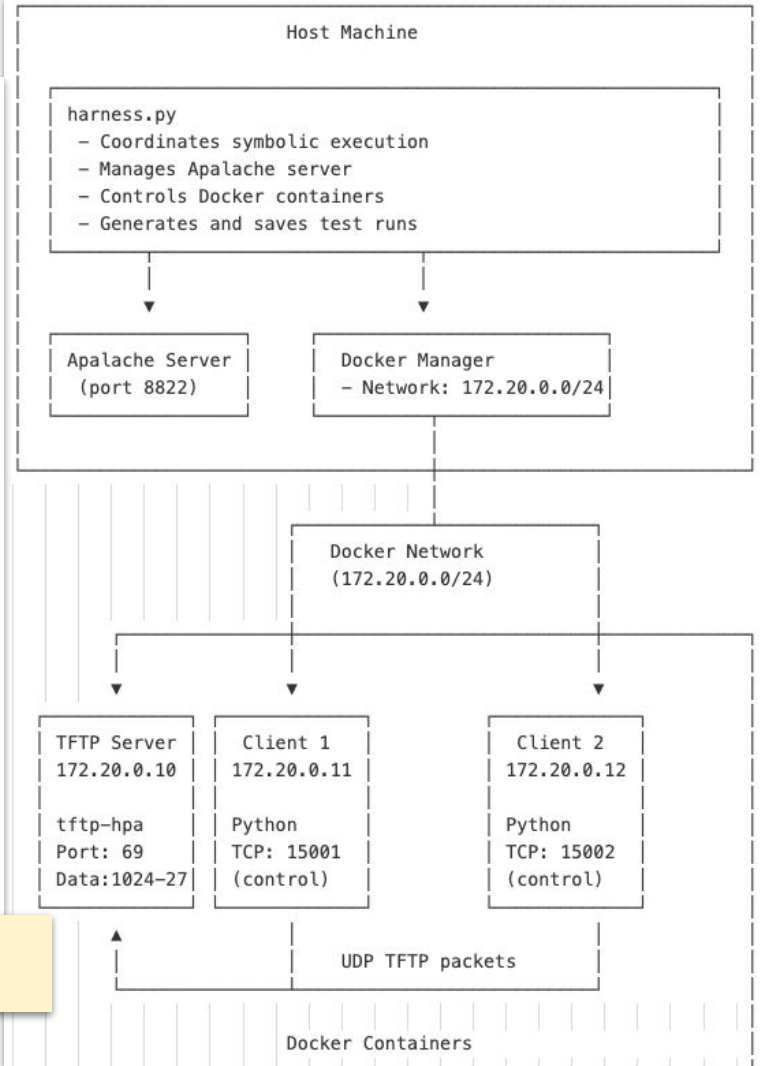
22 ## Requirement 2

24 The communication with Apache must be done using JSON RPC calls. The API is
25 implemented in the [client.py] file, which provides functions to send requests
26 to Apache and receive responses. The harness should utilize these functions to
27 obtain symbolic execution traces for the TFTP protocol.


29 ## Requirement 3



Overengineered



Debugging the harness



≈ 1 week

≈ 200 premium requests (Copilot)

1. Lots of debugging
2. Explaining Claude how to handle packets
3. Analyzing logs
4. Fixing bugs in the generated harness code
5. Claude lost in abstractions (e.g., clients in docker vs. clients in spec)

Lesson: design proper architecture upfront!

How it looks like

python_harness.log

tftpd_server.log

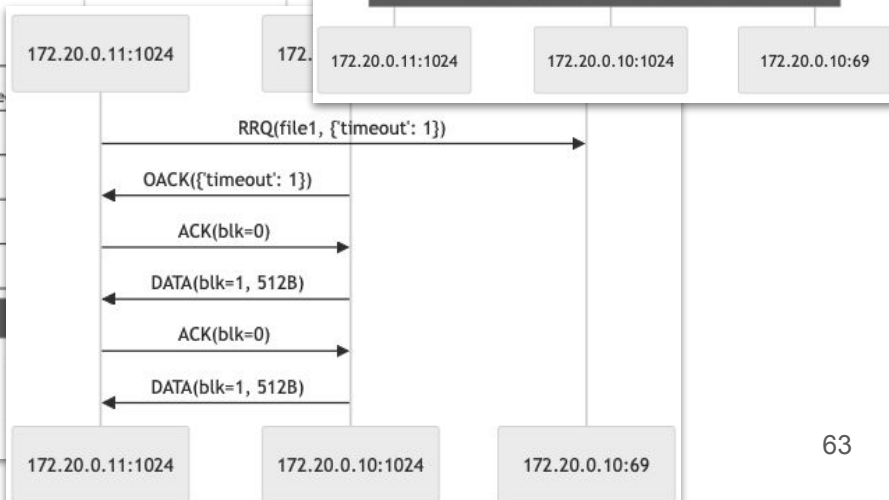
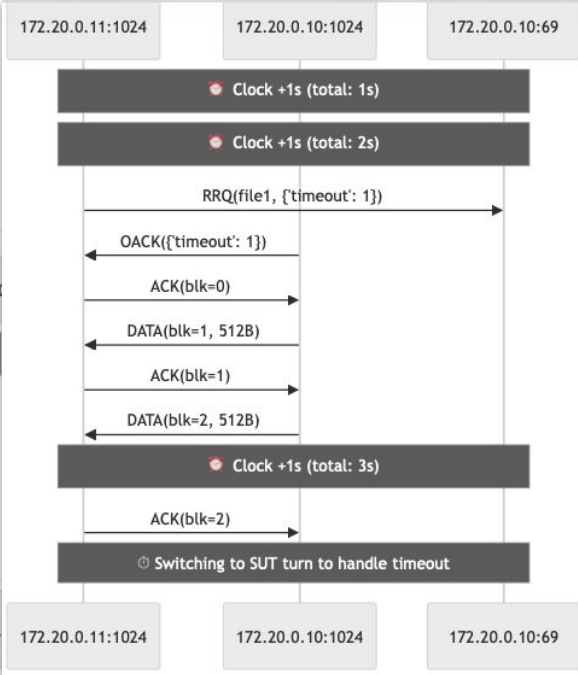
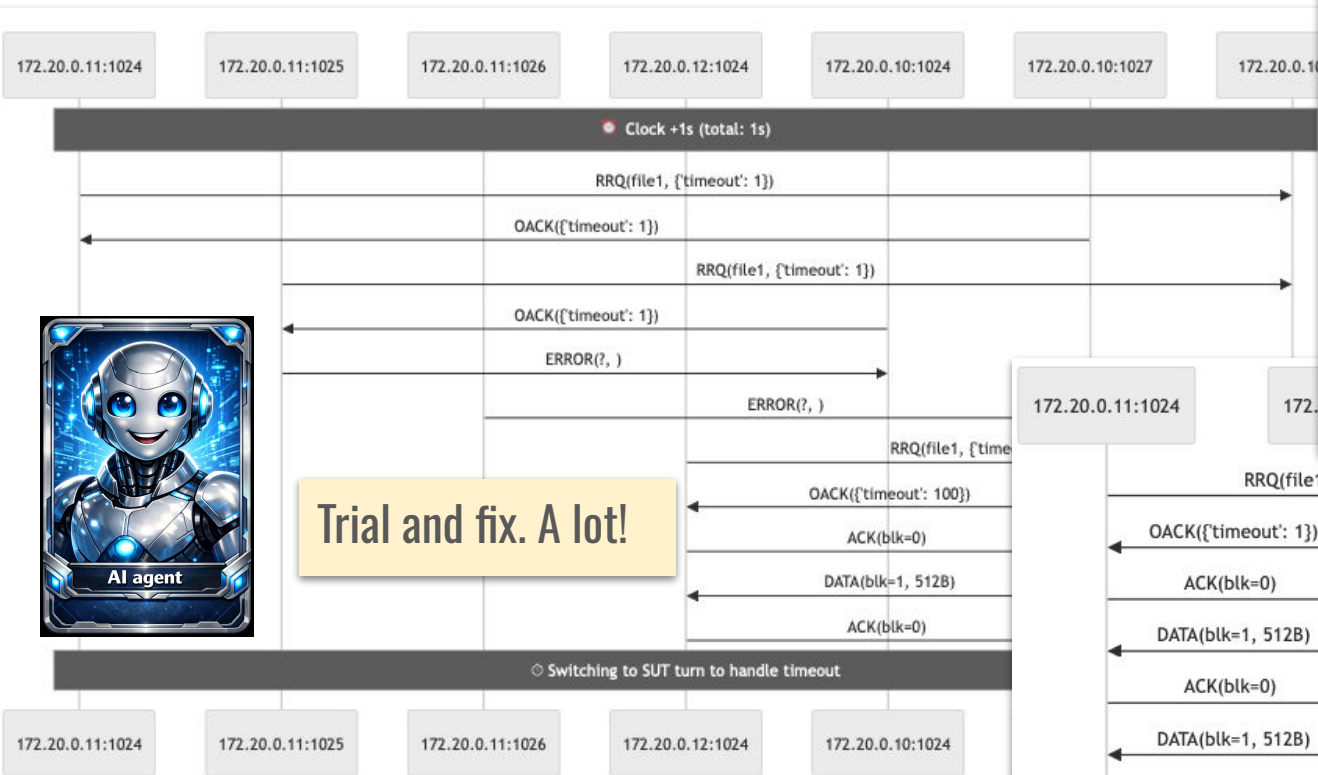
```
2025-11-12 22:21:03,033 - __main__ - INFO - === Starting test run generation 2 (max 100 steps) ===
2025-11-12 22:21:03,033 - __main__ - INFO - Trying transition 0 [Init]...
2025-11-12 22:21:03,035 - __client__ - INFO - Transition 0: ENABLED
2025-11-12 22:21:03,035 - __main__ - INFO - Transition 0 is ENABLED
2025-11-12 22:21:03,036 - __client__ - INFO - Moved to step 1
2025-11-12 22:21:03,036 - __main__ - INFO -
--- Step 1/100 ---
2025-11-12 22:21:03,036 - __main__ - INFO - Trying transition 4 [ClientRecvDATA]...
2025-11-12 22:21:03,038 - __client__ - INFO - Transition 4: DISABLED
2025-11-12 22:21:03,038 - __main__ - INFO - Transition 4 is DISABLED
2025-11-12 22:21:03,038 - __main__ - INFO - Rollback to snapshot 3
2025-11-12 22:21:03,039 - __main__ - INFO - Trying transition 5 [ClientRecvOACKthenSendAck]...
2025-11-12 22:21:03,040 - __client__ - INFO - Transition 5: DISABLED
2025-11-12 22:21:03,040 - __main__ - INFO - Transition 5 is DISABLED
2025-11-12 22:21:03,040 - __main__ - INFO - Rollback to snapshot 3
2025-11-12 22:21:03,041 - __main__ - INFO - Trying transition 3 [AdvanceClock]...
2025-11-12 22:21:03,042 - __client__ - INFO - Transition 3: ENABLED
2025-11-12 22:21:03,042 - __main__ - INFO - Transition 3 is ENABLED
2025-11-12 22:21:03,042 - __client__ - INFO - Moved to step 2
2025-11-12 22:21:03,042 - __main__ - INFO - Executing TFTP operation for transition 3
2025-11-12 22:21:03,044 - __main__ - INFO - Retrieved trace with 2 states
2025-11-12 22:21:03,044 - __main__ - INFO - Action: Advance Clock by 1
2025-11-12 22:21:03,044 - __main__ - INFO - Sleeping for 1 seconds to advance clock...
2025-11-12 22:21:04,049 - __main__ - INFO - ✓ Clock advanced by 1 seconds
2025-11-12 22:21:04,049 - __main__ - INFO -
--- Step 2/100 ---
2025-11-12 22:21:04,050 - __main__ - INFO - Trying transition 4 [ClientRecvDATA]...
2025-11-12 22:21:04,053 - __client__ - INFO - Transition 4: DISABLED
2025-11-12 22:21:04,053 - __main__ - INFO - Transition 4 is DISABLED
2025-11-12 22:21:04,053 - __main__ - INFO - Rollback to snapshot 6
2025-11-12 22:21:04,055 - __main__ - INFO - Trying transition 6 [ClientRecvOACKthenSendError]...
2025-11-12 22:21:04,057 - __client__ - INFO - Transition 6: DISABLED
2025-11-12 22:21:04,057 - __main__ - INFO - Transition 6 is DISABLED
2025-11-12 22:21:04,057 - __main__ - INFO - Rollback to snapshot 6
2025-11-12 22:21:04,059 - __main__ - INFO - Trying transition 0 [ClientSendRRQ]...
2025-11-12 22:21:04,066 - __client__ - INFO - Transition 0: ENABLED
2025-11-12 22:21:04,066 - __main__ - INFO - Transition 0 is ENABLED
2025-11-12 22:21:04,067 - __client__ - INFO - Moved to step 3
```

```
TFTP server watchdog started with PID 7
Started rsyslog daemon
Starting TFTP server (attempt 1/1)
Command: in.tftpd -vvv --foreground --address 172.20.0.10:69 --user nobody -
4:1027 --secure /var/tftp
Started tcpdump with PID 11 (unbuffered output)
tcpdump: data link type LINUX_SLL2
tcpdump: listening on any, link-type LINUX_SLL2 (Linux cooked v2), snapshot
ytes
21:21:00.532140 eth0 In IP (tos 0x0, ttl 64, id 11761, offset 0, flags [DF
7), length 52)
172.20.0.11.1024 > 172.20.0.10.69: [bad udp cksum 0x586f -> 0x514d!] TFTP
RQ "file1" octet timeout 1
21:21:00.533368 eth0 Out IP (tos 0x0, ttl 64, id 23847, offset 0, flags [no
(17), length 40)
172.20.0.10.1027 > 172.20.0.11.1024: [bad udp cksum 0x5863 -> 0xa93a!] U
21:21:00.582224 eth0 In IP (tos 0x0, ttl 64, id 52516, offset 0, flags [DF
7), length 64)
172.20.0.12.1025 > 172.20.0.10.69: [bad udp cksum 0x587c -> 0x4ba7!] TFTP
RQ "file1" octet blksize 513 timeout 2
21:21:00.583130 eth0 Out IP (tos 0x0, ttl 64, id 33392, offset 0, flags [no
(17), length 52)
172.20.0.10.1024 > 172.20.0.12.1025: [bad udp cksum 0x5870 -> 0xa397!] U
21:21:00.628954 eth0 In IP (tos 0x0, ttl 64, id 11822, offset 0, flags [DF
7), length 32)
172.20.0.11.1024 > 172.20.0.10.1027: [bad udp cksum 0x585b -> 0x9f91!] U
21:21:00.629025 eth0 Out IP (tos 0x0, ttl 64, id 23848, offset 0, flags [no
(17), length 544)
172.20.0.10.1027 > 172.20.0.11.1024: [bad udp cksum 0x5a5b -> 0xa8be!] U
21:21:00.686941 eth0 In IP (tos 0x0, ttl 64, id 52596, offset 0, flags [DF
7), length 58)
```

Generating sequence diagrams



Trial and fix. A lot!



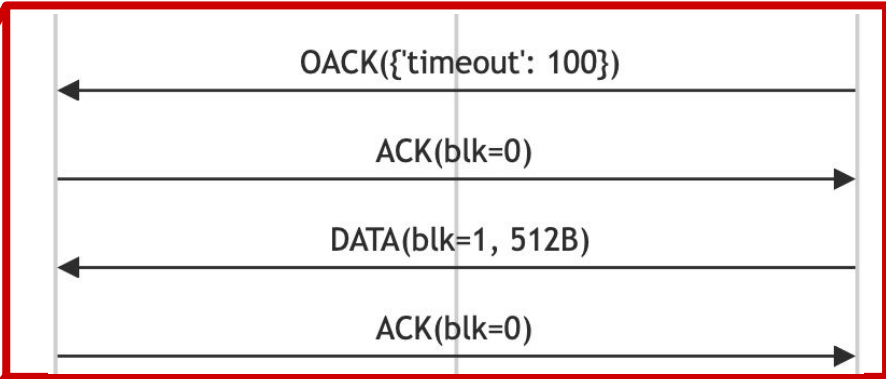
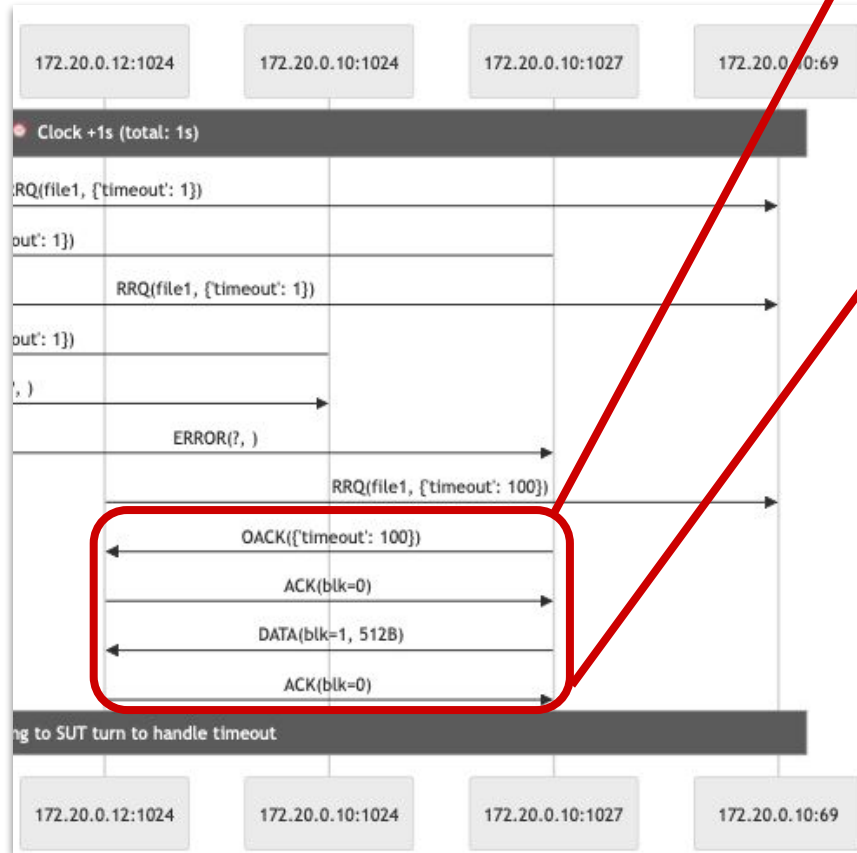
Spec vs. Code (1)

RFC 2347: “...the server should simply omit the option from the OACK, respond with an alternate value, or send an ERROR packet, with error code 8, to terminate the transfer.”

The server sends ERROR from the control port (69). Nope.

```
↑... @@ -207,7 +207,7 @@ _ServerSendErrorOnRrq(_rrq, _clientIpAndPort, _newServerPort, _rcvdPacket) ==
207 207     ServerRecvRRQthenSendError::
208 208     \E errorCode \in DOMAIN ALL_ERRORS:
209 209     LET errorPacket == [
210 210     -         srcIp |-> SERVER_IP, srcPort |-> 69,
210 210     +         srcIp |-> SERVER_IP, srcPort |-> _newServerPort,
211 211         destIp |-> _clientIpAndPort[1], destPort |-> _clientIpAndPort[2],
212 212         payload |-> ERROR(errorCode)
213 213     ]
```

Harness vs. Code (2)



The client has not received DATA(1) yet and sends ACK(0) again

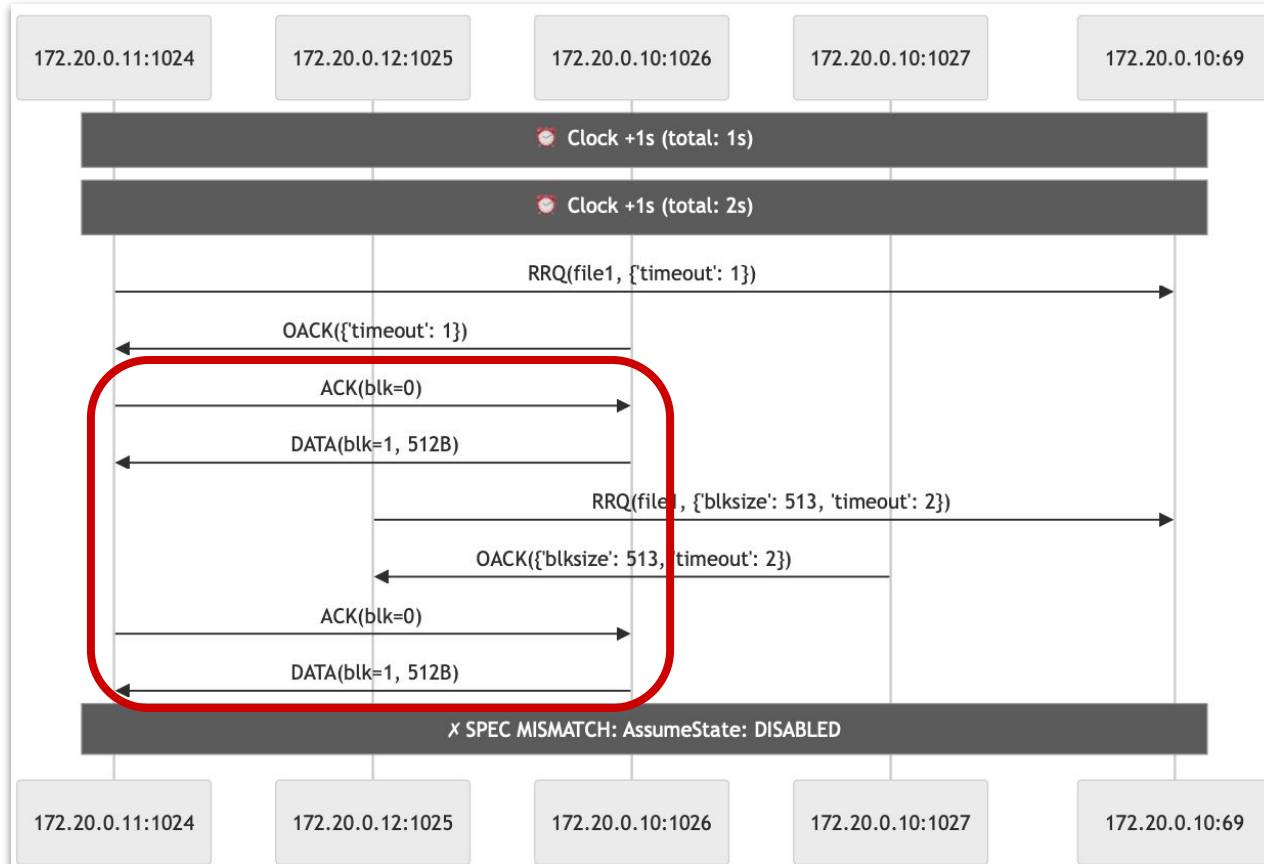
The harness expects a timeout

Fix #2: Give the client a chance to receive DATA after timeout (UDP is hard!)

Spec vs. Code (3)

Our server spec did not
send duplicate packets!

UDP is hard 🙄



```

\* The server receives an ACK packet and resends DATA that it sent in the past.
\* This is to fix the mismatch found by the test harness.
\* @type: $udpPacket => Bool;
ServerResendDATA(_udp) ==
  ServerResendDATA::
  LET ipPort == <<_udp.srcIp, _udp.srcPort>> IN
  /\ IsACK(_udp.payload)
  /\ _udp.destIp = SERVER_IP
  /\ ipPort \in DOMAIN serverTransfers
  /\ \E dataPacket \in packets:
    LET ack == AsACK(_udp.payload)
      data == AsDATA(dataPacket.payload)
      transfer == serverTransfers[ipPort]
      \* update the timestamp of the last transfer
      newTransfer == [ transfer EXCEPT !.timestamp = clock ]
    IN
  \* make sure that we receive from the correct port
  /\ _udp.destPort = transfer.port
  \* The DATA packet is sent in response to the ACK.
  /\ ack.blockNum + 1 = data.blockNum
  /\ dataPacket.srcIp = SERVER_IP
  /\ dataPacket.srcPort = _udp.destPort
  /\ dataPacket.destIp = _udp.srcIp
  /\ dataPacket.destPort = _udp.srcPort
  \* do not receive packets if the connection must timeout
  /\ clock <= transfer.timestamp + transfer.timeout
  \* either we have more data to send, or we send exactly 0 bytes in the last block
  /\ serverTransfers' = [ serverTransfers EXCEPT ![ipPort] = newTransfer ]
  /\ lastAction' = ActionRecvSend(_udp, dataPacket)
  /\ UNCHANGED <<packets, clientTransfers, clock>>

```

FIX #3:

Let the server re-send DATA

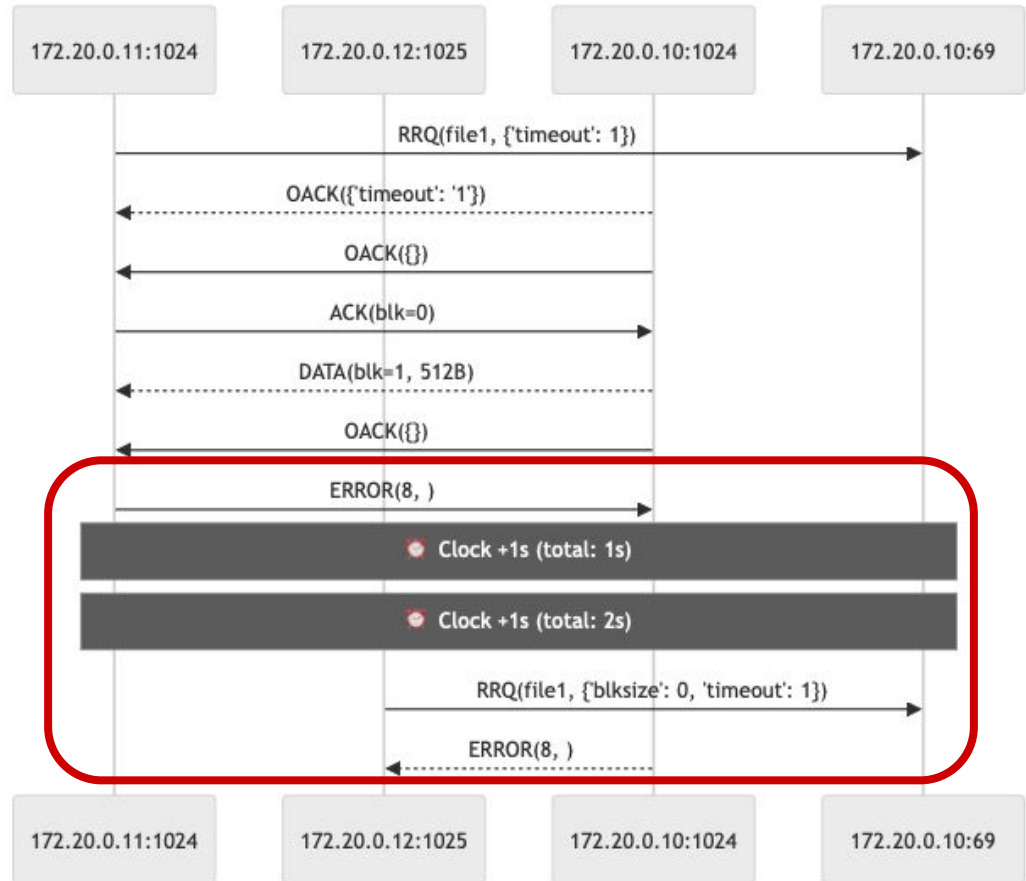
Spec vs. Code (4)

Per RFC 1350, the server does not have to reply on ERROR!

There is input, but no output!

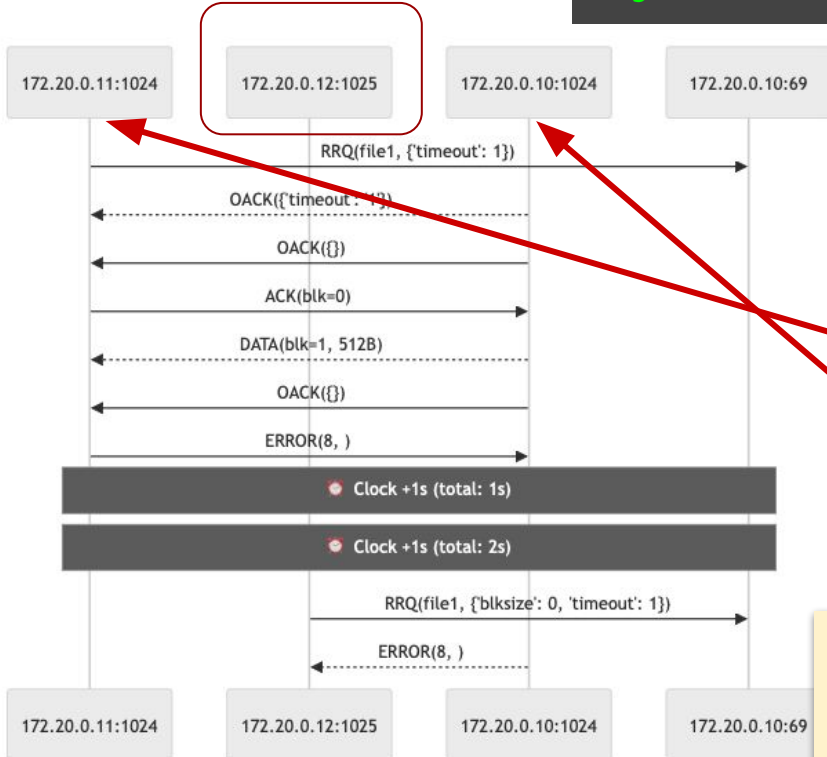
The server in the spec had not received ERROR so far

Oops. Our generated harness was supposed to work this way



How do we debug this divergence?

```
$ jless test-results/run_0002/divergence_trace.itf.json
```



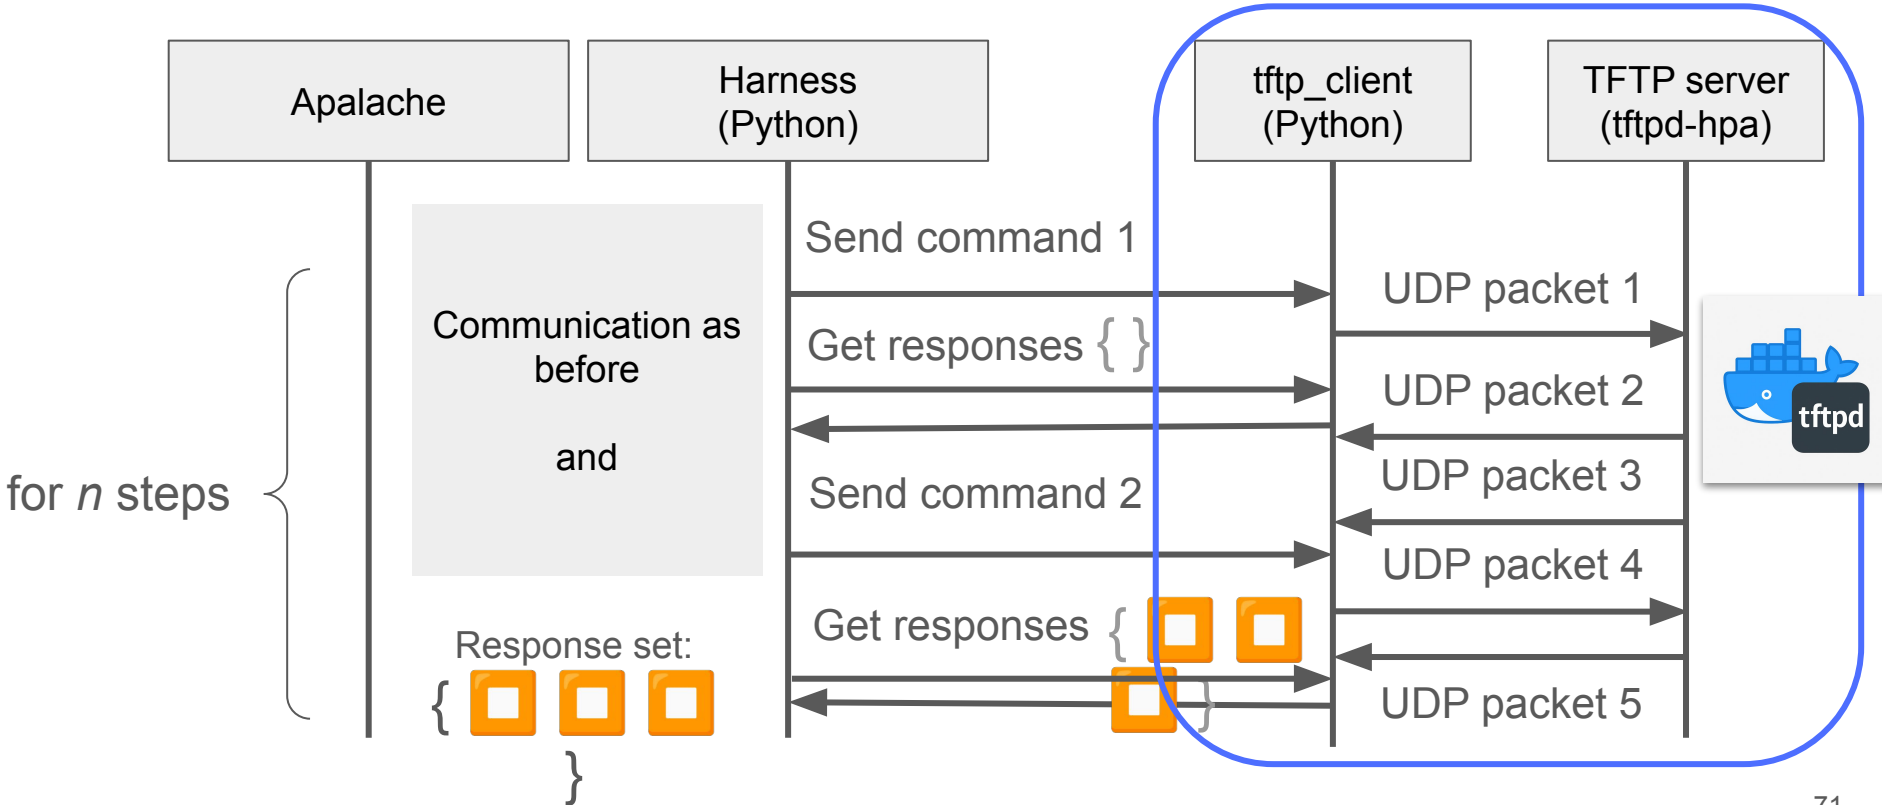
```
    "#bigint": "2"  
  ▽ transferred: (1) {"#bigint": "0"}  
    "#bigint": "0"  
  ▽ tsize: (1) {"#bigint": "0"}  
    "#bigint": "0"  
  ▽ clock: (1) {"#bigint": "2"}  
  ▽ lastAction: (2) {tag: "ActionClientSendRRQ", value: {...}}  
  ▽ packets: (1) {"#set": [...], {...}, {...}, {...}, {...}, {...}}  
  ▽ serverTransfers: (1) {"#map": [[{...}, {...}]]}  
    ▽ "#map": (1) [[{...}, {...}]]  
      ▽ [0]: (2) [...], {...}  
        ▽ [0]: (1) {"#tup": ["172.20.0.11", {...]}  
          ▽ "#tup": (2) ["172.20.0.11", {...]}  
            [0]: "172.20.0.11"  
            ▽ [1]: (1) {"#bigint": "1024"}  
              "#bigint": "1024"  
        ▽ [1]: (8) {blksize: {...}, blockNum: {...}, blocks: [], port: {...}, ti  
          ▽ blksize: (1) {"#bigint": "512"}  
          ▽ blockNum: (1) {"#bigint": "1"}  
          blocks: []  
          ▽ port: (1) {"#bigint": "1024"}  
            ▽ "#bigint": "1024"  
          ▽ timeout: (1) {"#bigint": "1"}  
            timeout: (1) {"#bigint": "10"}  
            "#bigint": "10"
```

The spec still has the transfer
172.20.0.11:1024 and 172.20.0.10:1024

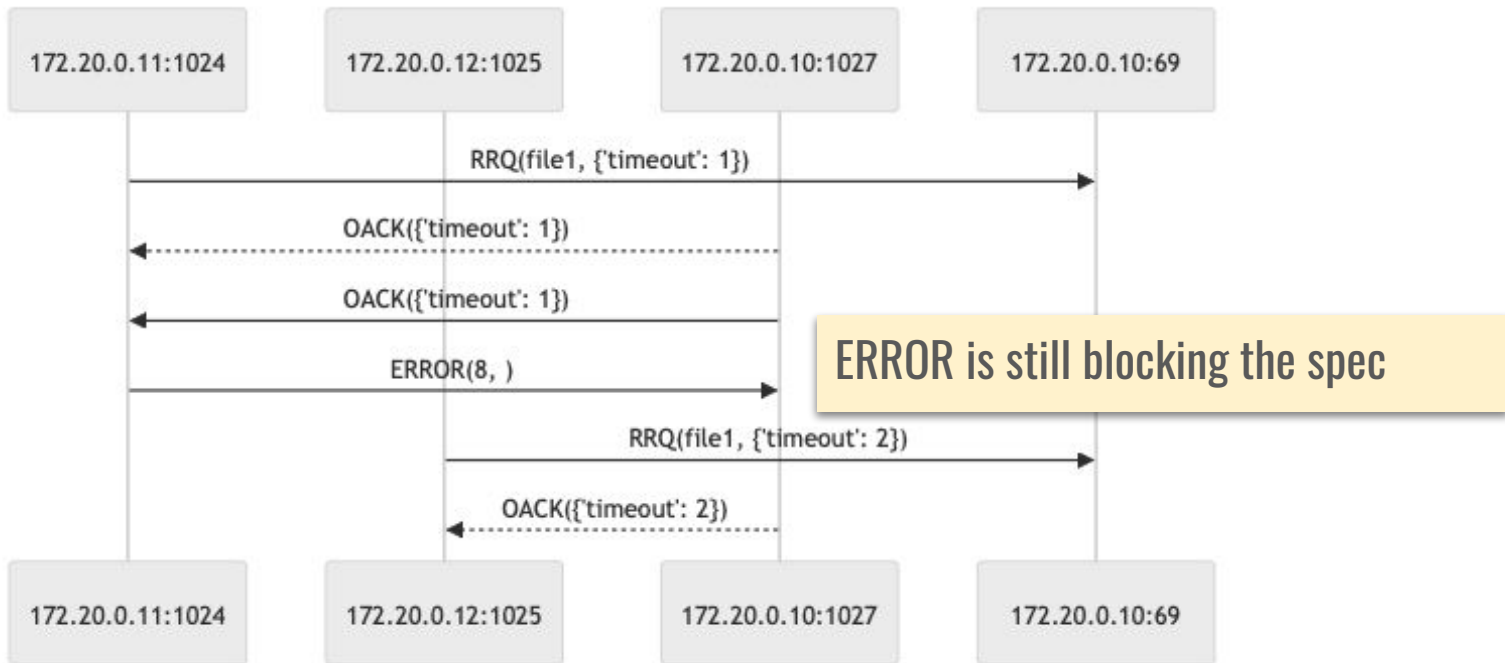
RFC 1350: *“In order to create a connection, each end of the connection chooses a TID for itself, to be used for the duration of that connection. The TID's chosen for a connection should be randomly chosen, so that the probability that the same number is chosen twice in immediate succession is very low.”*

We have hit the case of low probability?

Revised testing loop with Apache



Does this work?



FIX #5: the server may reuse the port if there was an ERROR

```
\* The server receives RRQ and sends one of: DATA, OACK, or ERROR.
\* @type: $udpPacket => Bool;
ServerRecvRRQ(_udp) ==
  /\ IsRRQ(_udp.payload)
  /\ _udp.destIp = SERVER_IP
  /\ _udp.destPort = 69
  /\ LET rrq == AsRRQ(_udp.payload)
     clientIpAndPort == <<_udp.srcIp, _udp.srcPort>> IN
  \* A client can open multiple connections from different ports.
```

```
\* the server allocates a new port for the connection, if it can find one
```

```
/\ \E newServerPort \in PORTS:
```

```
  /\ \/ \A p \in DOMAIN serverTransfers:
```

```
    serverTransfers[p].port /= newServerPort
```

```
  \* Or, there was an ERROR packet that cancelled the active transfer.
```

```
  \* This has a bad smell, but is needed to conform tftpd-hpa.
```

```
  \/ \E packet \in packets:
```

```
    /\ IsERROR(packet.payload)
```

```
    /\ packet.destIp = SERVER_IP
```

```
    /\ packet.destPort = newServerPort
```

```
  \E newServerPort
  \* A client can open multiple connections from different ports.
```

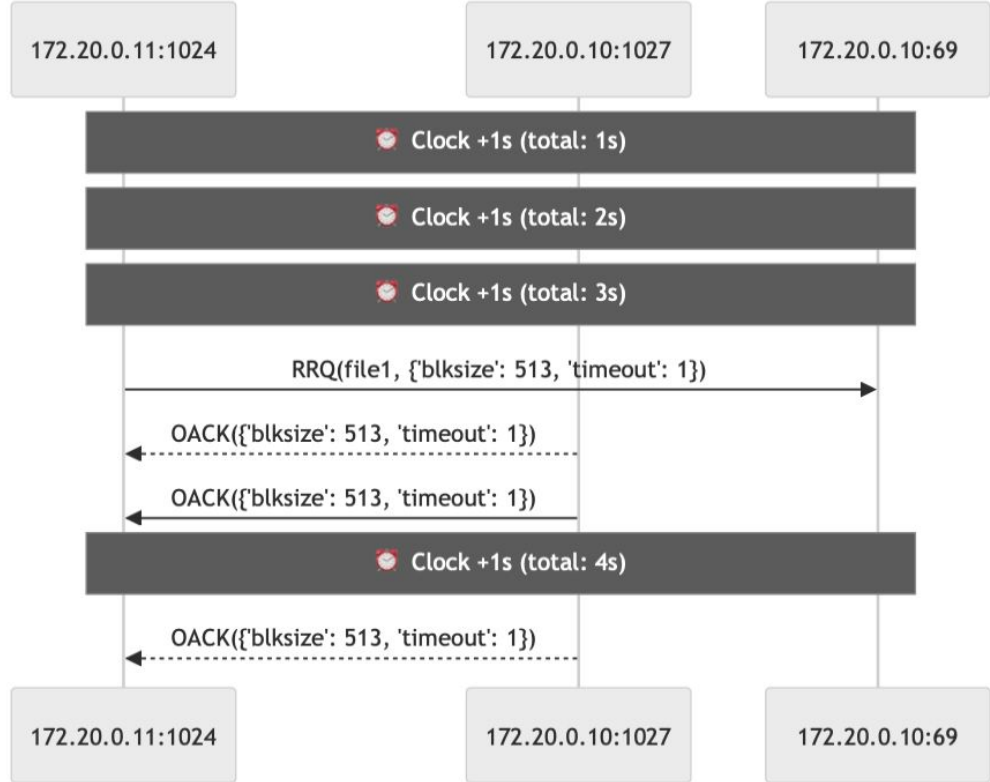
```
  \E newServerPort
  \* Or, there was an ERROR packet that cancelled the active transfer.
```

Easy in TLA+, hard in stricter specification languages!

```
\* According to RFC 2347, the server may respond with DATA or OACK
/\ \/ _ServerSendDataOnRrq(rrq, clientIpAndPort, newServerPort, _udp)
  \/ _ServerSendOackOnRrq(rrq, clientIpAndPort, newServerPort, _udp)
  \/ _ServerSendErrorOnRrq(rrq, clientIpAndPort, newServerPort, _udp)
```

Spec vs. Code (5)

Duplicate and outdated packets!

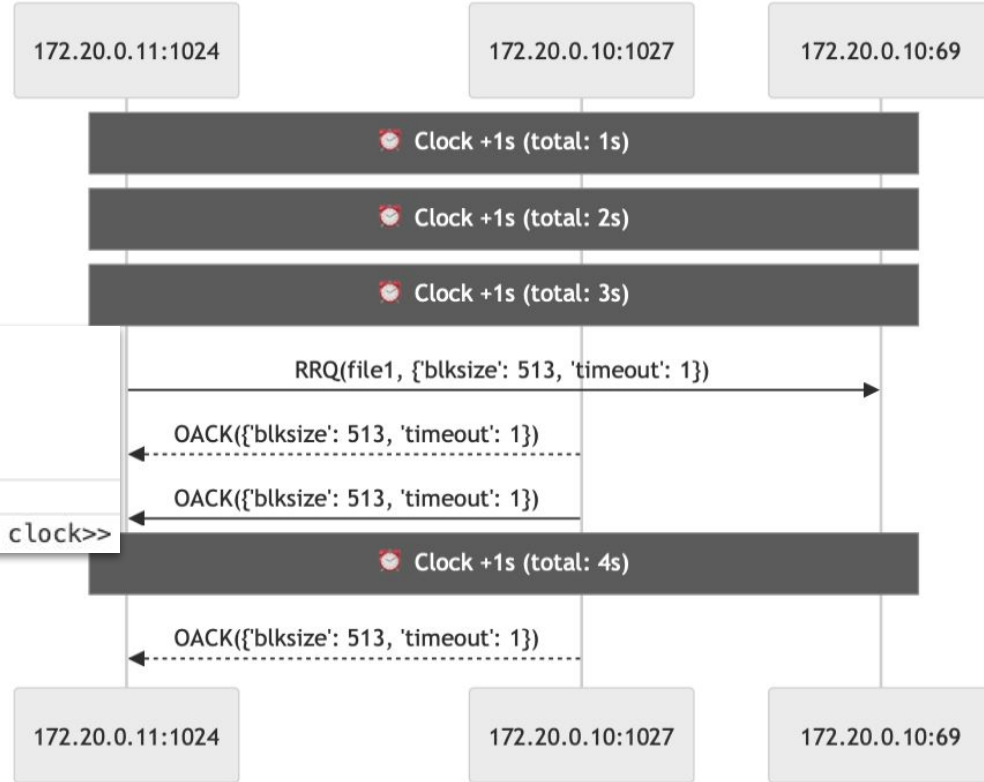


Spec vs. Code (5)

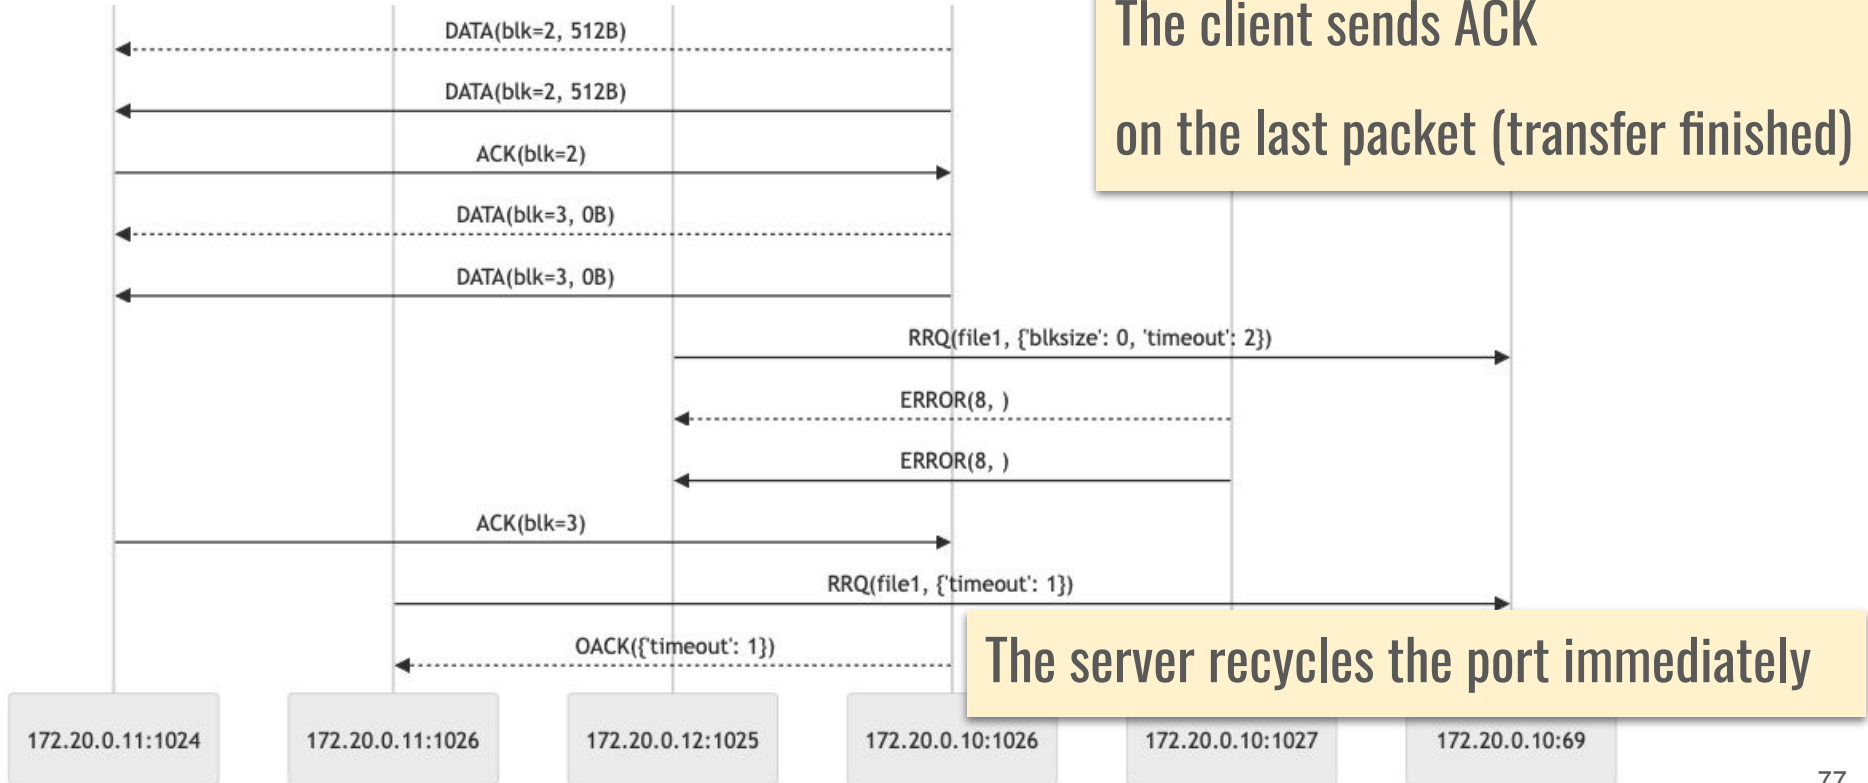
Duplicate and outdated packets!

```
\* @type: $udpPacket => Bool;  
ServerSendDup(_udp) ==  
  ServerSendDup::  
  /\ _udp.srcIp = SERVER_IP  
  /\ lastAction' = ActionRecvSend(_udp)  
  /\ UNCHANGED <<packets, serverTransfers, clientTransfers, clock>>
```

In TLA⁺ theory, this is just a stuttering step



Spec vs. Code (5)



RFC 2347: *“The end of a transfer is marked by a DATA packet that contains between 0 and 511 bytes of data (i.e., Datagram length < 516). This packet is acknowledged by an ACK packet like all other DATA packets. The host acknowledging the final DATA packet may terminate its side of the connection on sending the final ACK. On the other hand, dallying is encouraged. This means that the host sending the final ACK will wait for a while before terminating in order to retransmit the final ACK if it has been lost. The acknowledger will know that the ACK has been lost if it receives the final DATA packet again. The host sending the last DATA must retransmit it until the packet is acknowledged or the sending host times out. If the response is an ACK, the transmission was completed successfully. If the sender of the data times out and is not prepared to retransmit any more, the transfer may still have been completed successfully, after which the acknowledger or network may have experienced a problem. It is also possible in this case that the transfer was unsuccessful. In any case, the connection has been closed.”*

FIX #6: reuse the ports from completed transfers

```
diff --git a/spec/tftp.tla b/spec/tftp.tla
index 57265a6..ae20c73 100644
--- a/spec/tftp.tla
+++ b/spec/tftp.tla
@@ -233,7 +233,9 @@ ServerRecvRRQ(_udp) ==
    \* the server allocates a new port for the connection, if it can find one
    /\ \E newServerPort \in PORTS:
        /\ \/\ \A p \in DOMAIN serverTransfers:
-           serverTransfers[p].port /= newServerPort
+           \/\ serverTransfers[p].port /= newServerPort
+           \* FIX #6: allow reusing ports from completed transfers
+           \/\ serverTransfers[p].transferred = serverTransfers[p].tsize
    \* Or, there was an ERROR packet that cancelled the active transfer.
    \* This has a bad smell, but is needed to conform tftpd-hpa.
    \/\ \E packet \in packets:
```

DATA(blk=2, 512B)

DATA(blk=2, 512B)

172.20.0.11:1024

172.20.0.11:1026

172.20.0.12:1025

172.20.0.10:1026

172.20.0.10:1027


172.20.0.10:69

Test traces get longer

Ran 100 test runs 100 steps each: 37 of them diverge

Hard to analyze these traces by hand

LLM hallucinated too much in identifying the root cause



TFTP Test Divergence Analysis

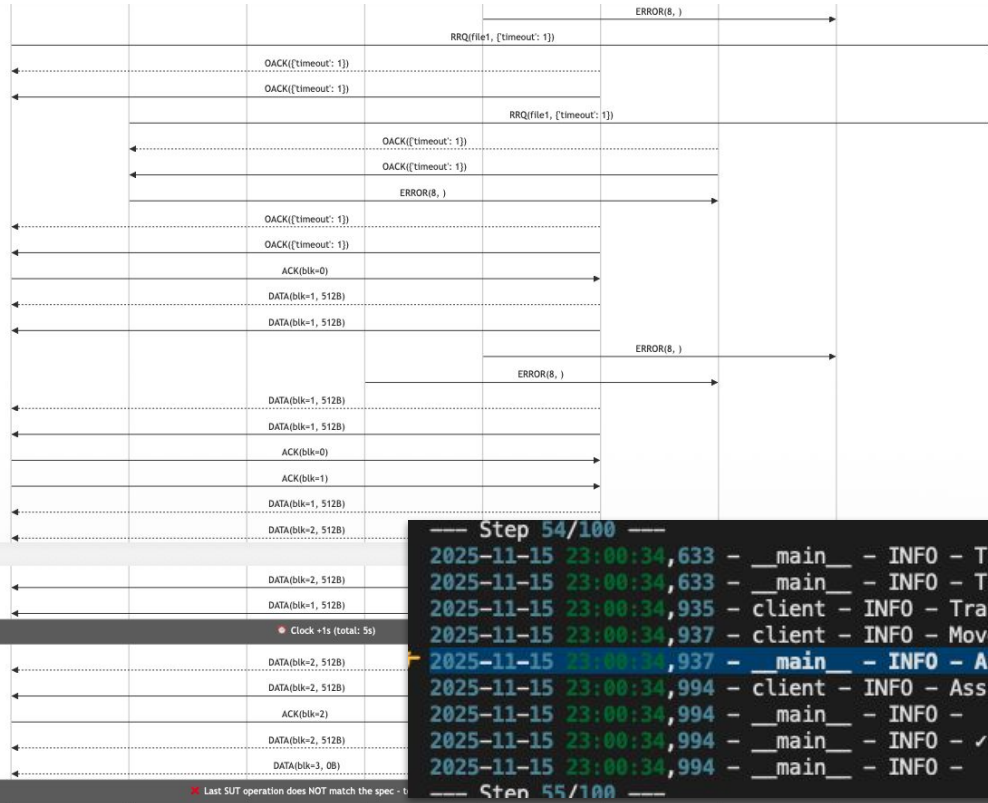
Based on my analysis of all diverging test runs in `test-results/`, here are the key findings:

Summary

37 diverging test runs with **3 main patterns**:

- 1. Duplicate DATA packets (26 cases - 81%):** The SUT sends DATA packets with the correct block number, but these are retransmissions/duplicates
- 2. Duplicate OACK packets (5 cases - 16%):** The SUT sends duplicate OACK packets during negotiation
- 3. Unexpected ACK (1 case - 3%):** Rare case of unexpected ACK packet

Spec vs. Code (7)



The server starts data exchange

In the middle, the server spec interprets an earlier RRQ as no-negotiation case, and sends DATA(1) again!

```
--- Step 54/100 ---
2025-11-15 23:00:34,633 - __main__ - INFO - Turn: sut. 4 transitions to try
2025-11-15 23:00:34,633 - __main__ - INFO - Trying transition 8 [ServerRecvRRQthenSendData]...
2025-11-15 23:00:34,935 - client - INFO - Transition 8: ENABLED
2025-11-15 23:00:34,937 - client - INFO - Moved to step 55
2025-11-15 23:00:34,937 - __main__ - INFO - Assume lastAction: ActionRecvSend(sent=UdpPacket(srcIp='172.20.0.10',
2025-11-15 23:00:34,994 - client - INFO - AssumeState: ENABLED
2025-11-15 23:00:34,994 - __main__ - INFO - EXECUTE ACTION: ActionRecvSend(sent=UdpPacket(srcIp='172.20.0.10',
2025-11-15 23:00:34,994 - __main__ - INFO - ✓ Received packet matches the spec
2025-11-15 23:00:34,994 - __main__ - INFO -
--- Step 55/100 ---
```

• Last SUT operation does NOT match the spec - 1

```

diff --git a/spec/tftp.tla b/spec/tftp.tla
index c48f936..df41aaa 100644
--- a/spec/tftp.tla
+++ b/spec/tftp.tla
@@ -145,6 +145,8 @@ _ServerSendDataOnRrq(_rrq, _clientIpAndPort, _newServerf
    transferred |-> dataSize
    ]
    IN
+   \* The no-negotiation case of RFC 2347. No options were requested.
+   /\ DOMAIN _rrq.options = {}
    /\ packets' = packets \union { dataPacket }
    /\ serverTransfers' = [
        p \in DOMAIN serverTransfers \union {_clientIpAndPort} |->
@@ -163,6 +165,8 @@ _ServerSendOackOnRrq(_rrq, _clientIpAndPort, _newServerf
    ServerRecvRRQthenSendOack:
    \E optionsSubset \in SUBSET DOMAIN _rrq.options,
        blksize \in 0..65464, timeout \in 1..255:
+   \* RFC 2349: option negotiation
+   /\ DOMAIN _rrq.options /= {}
    \* RFC 2349, Section 3.1: "If the server is willing to accept
    \* the blocksize option, it sends an Option Acknowledgment
    \* (OACK) to the client. The specified value must be less

```

FIX #7:

1. Apply RFC 2347 case only when options is empty
2. Apply RFC 2349 case only when options are non-empty

Spec vs. Code (8)

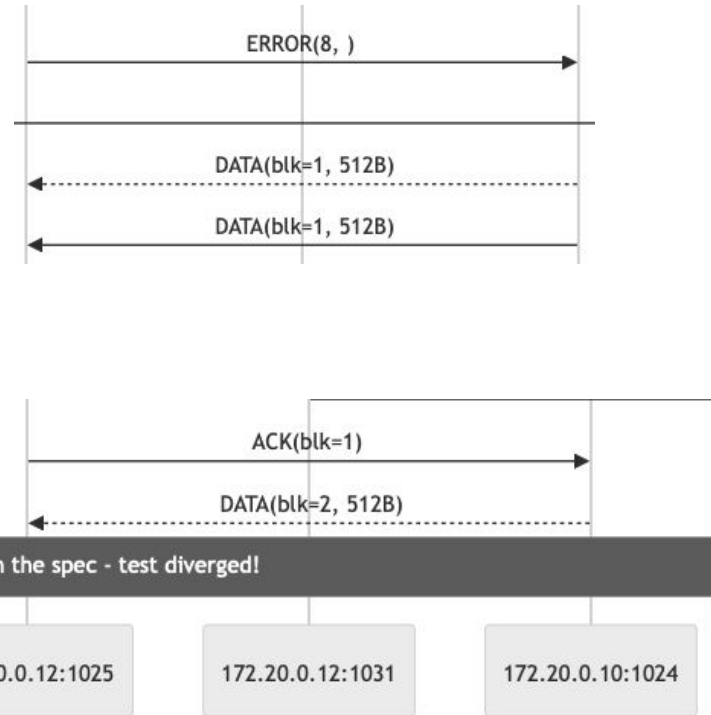
The client receives DATA and ERRORS

The server keeps sending DATA

The client interprets DATA(1) as the new session without negotiation

✘ Last SUT operation does NOT match the spec - test diverged!

FIX #8: Save the protocol version



FIX #9: The client must send tsize = 0 in RRQ

FIX #10: The server should send default timeout if it's not specified in the options

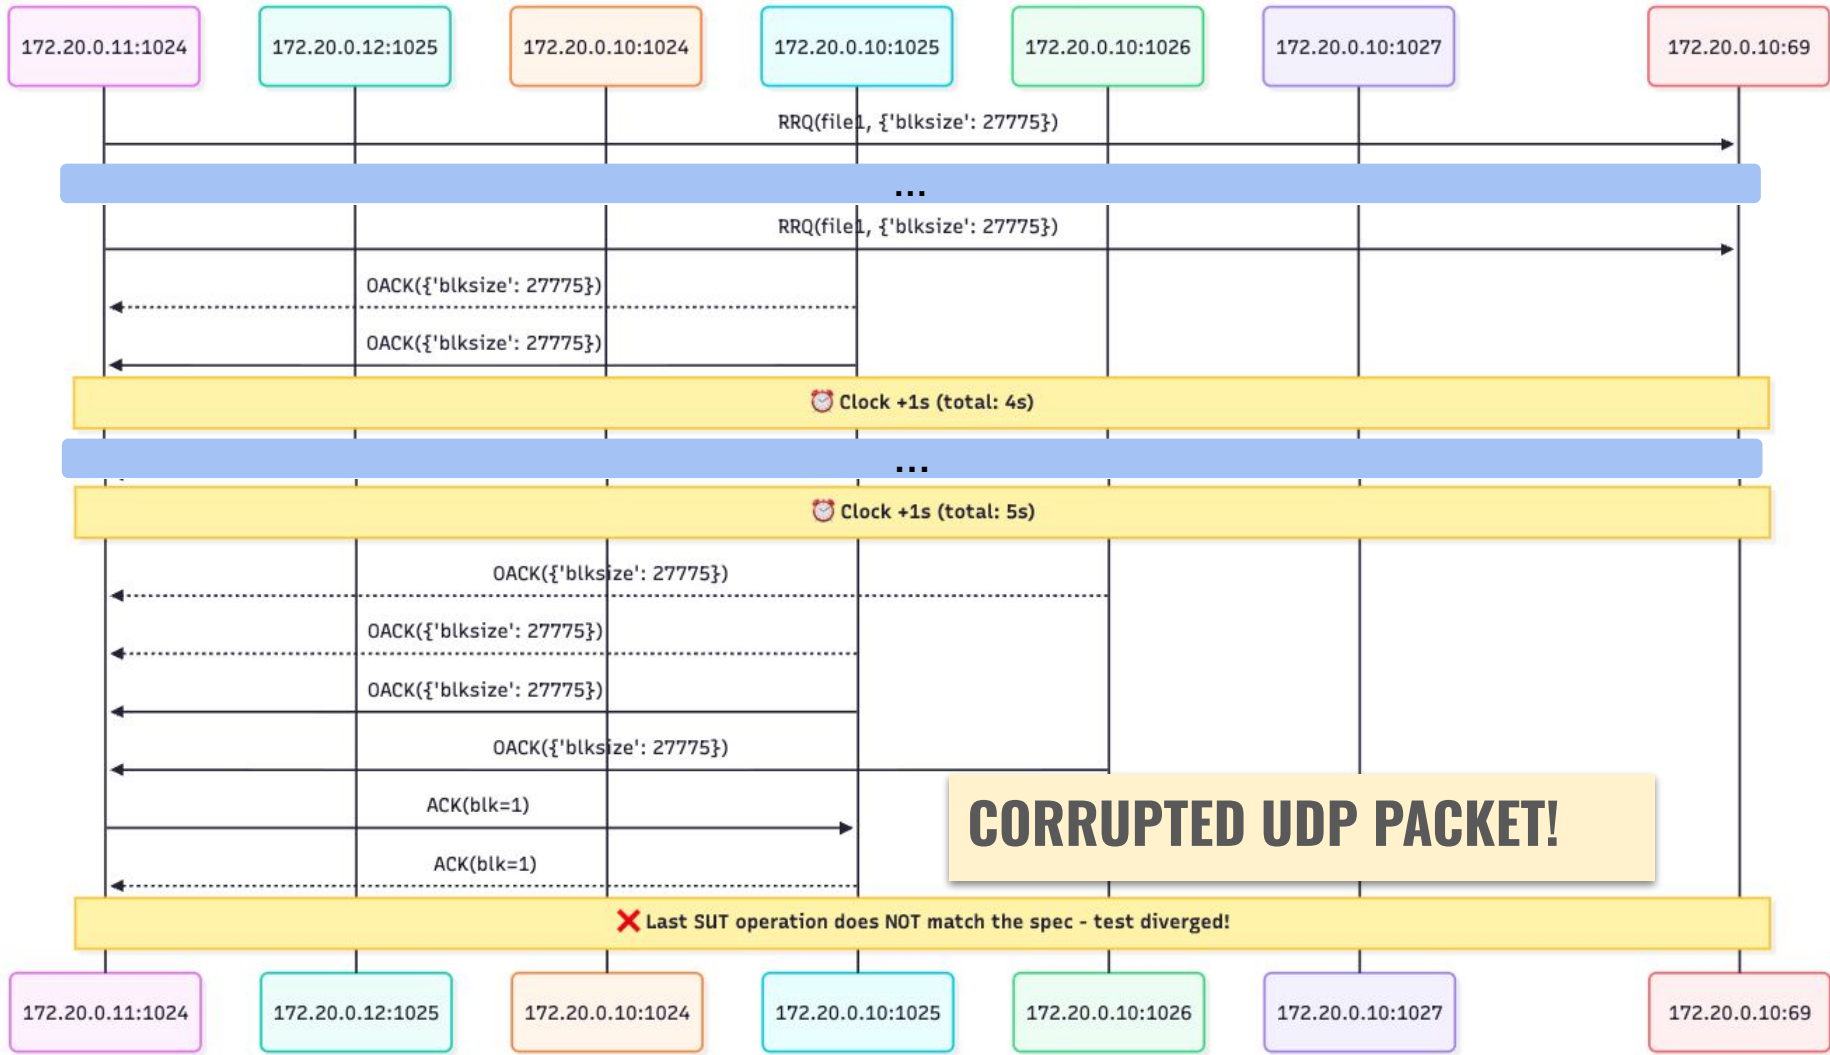
FIX #11: The server may send invalid (e.g., outdated) packets

FIX #12: My understanding of TFTP timeouts was wrong

FIX #13: Handle ERROR packets

FIX #14: Receive only one OACK message

FIX #15: Use clientIP-clientPort-serverPort triplets



Are all of our specifications broken?

The verification engineer's mindset:

The specification usually **overapproximates** the implementation

Reachability: if the impl. reaches a state s , then the spec. reaches $\alpha(s)$

Conformance testing (roughly):

If the spec. executes action A , the impl. must be able to execute A

If the impl. executes action A , the spec. must be able to execute A

Plenty of research in the 1990es

October 2025

Copilot + Sonnet 4.5

Lessons from LLM-generated harnesses




1. Surprisingly, **it works** – not effortlessly though
2. **Do not trust** the generated harness – TODOs and bugs inside
3. **Do not let** an LLM define **its own formats** – it goes wild
4. Log yourself and **define your log format** – or, face wild regexes
5. **Generate visualizations** that suit your needs – it's **amazing!**

Better, use structured formats – JSON

Produce 100 episodes, 100 steps each

```
$ ./harness.py --docker --steps=100 --tests=100
```



Improvements?

Docker restarts and `time.sleep` are **slow** (seconds)

Deterministic simulation instead?


Guided search = **fuzzing** + **symbolic execution**

Pros of our approach

Black-box testing: communication over the network, no instrumentation

Minimal mapping: define the labels of TESTER and SUT

Modular: test components in isolation, the rest is the spec



Early 2026 LLMs

Quality jump Q4 2025 – Q1 2026

Claude Code Opus 4.6 + Sonnet 4.6

Codex + GPT 5.3 / 5.4

```
→ zookeeper-testing git:(main) X claude
```



```
Claude Code v2.1.81
```

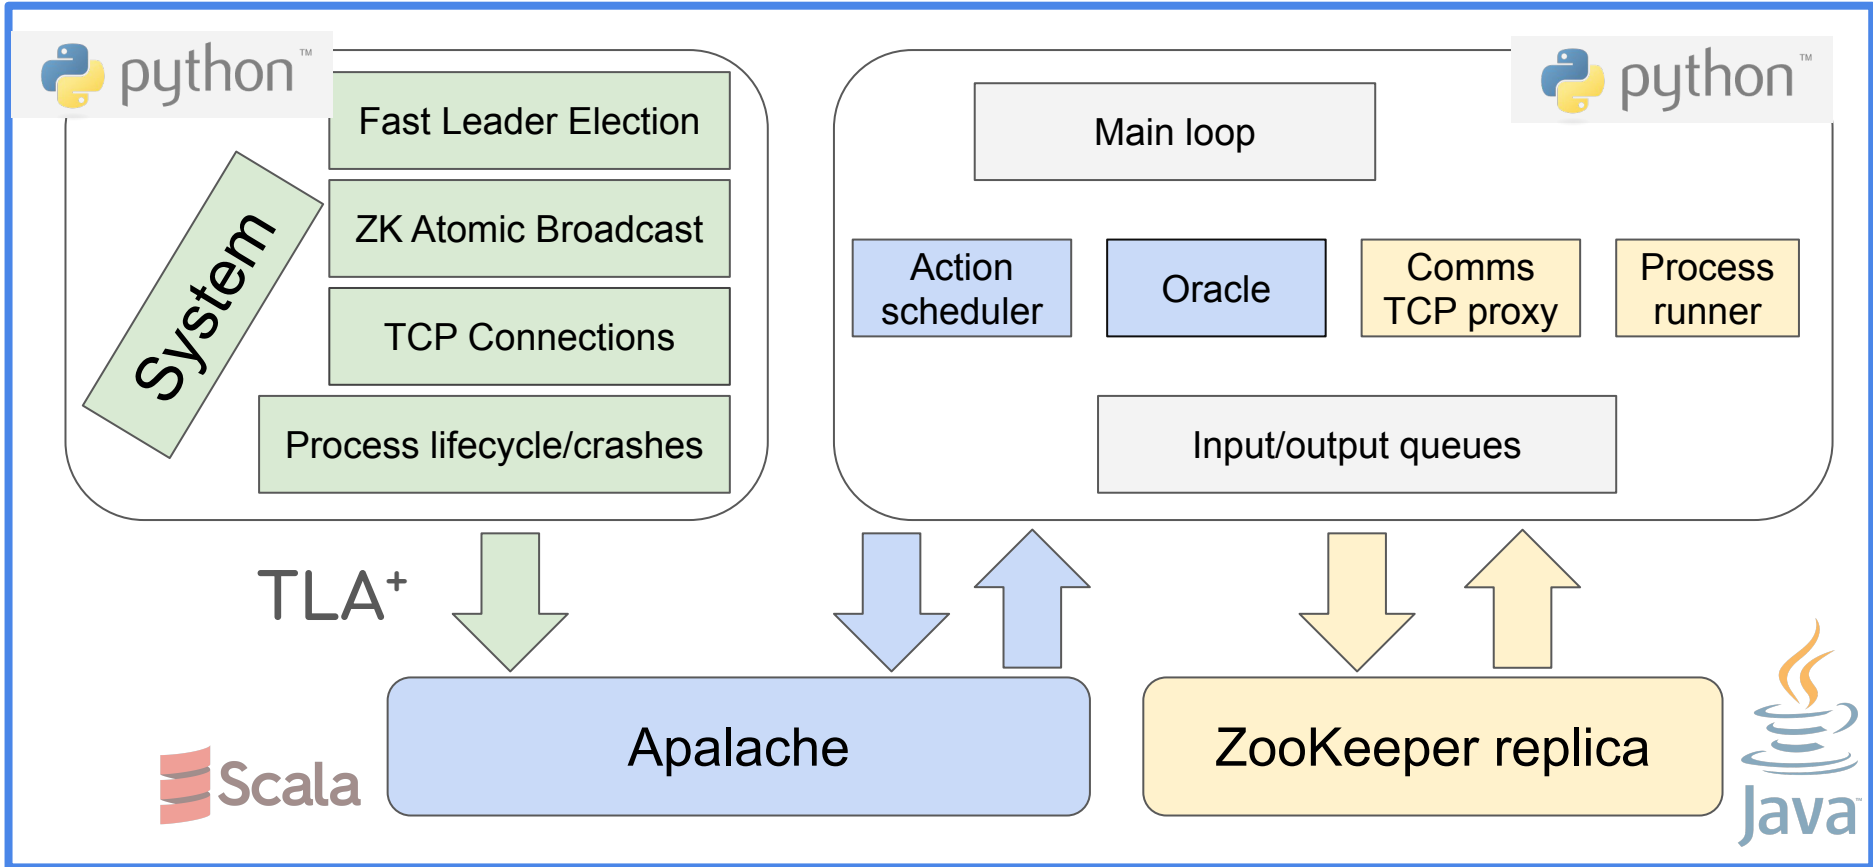
```
Sonnet 4.6 · Claude Pro
```

```
~/dev1/zookeeper-testing
```



APACHE
ZooKeeper™

```
> Extract TLA+ specifications from the source code of ZooKeeper, write a test harness, run it
```



Recommendation by ChatGPT



E
T
E
R
R
E
P
O
R
T

ETR 049

October 1992

Source: ETSI TC-ATM

Reference: DTR/ATM-1006

ICS: 33.020

Key words: Test specification methods

Advanced Testing Methods (ATM); State of research in the area of formal test specification methods

ETSI

European Telecommunications Standards Institute

ETSI Secretariat

Postal address: F-06921 Sophia Antipolis CEDEX - FRANCE

Office address: 650 Route des Lucioles - Sophia Antipolis - Valbonne - FRANCE

X.400: c=fr, a=atlas, p=etsi, s=secretariat - **Internet:** secretariat@etsi.fr

Tel.: +33 92 94 42 00 - Fax: +33 93 65 47 16

Copyright Notification: No part may be reproduced except as authorized by written permission. The copyright and the foregoing restriction extend to reproduction in all media.

© European Telecommunications Standards Institute 1992. All rights reserved.

Work in progress

Still no magic. Human guidance is needed


More autonomy. Sometimes, LLMs work for 1-2 hours

Better efficiency. 5 protocols + harness after 7 days of work



ZooMonkey
ZooKeeper Testing Framework

Takeaways



1. Write your own **symbolic search** scripts!
2. **Connect your spec with the code**, as you like it
3. **LLMs can help** you with automation
4. You know better **what matters to your project**
5. **Put it in your CI**
6. Easy to **parallelize**



APALACHE



Q & A

Detailed blog posts and links:

[protocols-made-fun.com]

e-mail: igor@konnov.phd telegram: [igor_konnov_phd](https://t.me/igor_konnov_phd)

Testing with TLA⁺

1. Nagendra et. al. ***Model guided fuzzing of distributed systems*** (2025)
2. Cirstea, Kuppe, Merz, Loillier. *Validating Traces of Distributed Systems Against TLA+ Specifications* (2024)
3. Chamayou et. al. *Validating System Executions with the TLA+ Tools* (2024)
4. Jordan Halterman. *Verifiability Gap: Why We Need More From Our Specs and How We Can Get It* (2020)
5. Jessie Davis et al. *eXtreme Modelling in Practice* (2020)
6. Kupriyanov, Konnov. *Model-based testing with TLA+ and Apache* (2020)
7. Pressler. *Verifying Software Traces Against a Formal Specification with TLA+ and TLC* (2018)

Except [6], all use TLC

Select fitter transitions

Filter tests with views

```
  \* Compute the measure as the max of block numbers of all transfers.
  \* @type: (<<Str, Int>> -> $transfer) => Int;
Measure(t) ==
  LET MaxS(S) == ApaFoldSet(LAMBDA x, y: Max(x, y), 0, S) IN
  MaxS({ t[ipPort].blockNum : ipPort \in DOMAIN t })
```

```
  \* Use the number of transferred blocks as the measure for computing
  \* the fitness function. The rest of the view is used for filtering
  \* similar states.
  \* @type: << Int, <<Str, Int>> -> $absTransfer, <<Str, Int>> -> $absTransfer >>;
MeasureView == <<
  (1 + Measure(serverTransfers) + Measure(clientTransfers)) * 10,
  [ p \in DOMAIN serverTransfers |-> AbsTransfer(serverTransfers[p]) ],
  [ p \in DOMAIN clientTransfers |-> AbsTransfer(clientTransfers[p]) ]
>>
```

```
AbsTransfer(t) ==
  [ port |-> t.port,
    \* -1, 0, [1, ...]
    tsize |-> Min(t.tsize, 1),
    \* abstract [0, 512), [512, 1024), [1024, 1536), ...
    blksize |-> t.blksize \div 512,
    \* interval abstraction [0, 10], [11, 255]
    timeout |-> IF t.timeout <= 10 THEN 10 ELSE 255,
    \* keep exact blockNum
    blockNum |-> t.blockNum,
    \* simply whether we are below or above the timeout threshold
    timestamp |-> IF clock <= t.timestamp + t.timeout THEN -1 ELSE 1
  ]
```

This is pure TLA+.

No tool extension needed

ITF JSON Traces

[apalache-mc.org/docs/adr/015adr-trace.html]

```
{
  "#meta": {
    "index": 1
  },
  "clientTransfers": {
    "#map": [
      [
        {
          "#tup": [
            "172.20.0.11",
            {
              "#bigint": "1024"
            }
          ]
        }
      ],
      {
        "blksize": {
          "#bigint": "512"
        },
        "blockNum": {
          "#bigint": "0"
        },
        "blocks": [],
        "port": {
          "#bigint": "69"
        },
        "proto": "options_yes",
        "timeout": {
          "#bigint": "1"
        }
      }
    ]
  }
}
```

```
trace_json = {
  "#meta": {"id": 23},
  "params": ["N"],
  "vars": ["pc", "x"],
  "loop": 0,
  "states": [
    {
      "#meta": {"no": 0},
      "N": {"#bigint": "3"},
      "pc": "idle",
      "x": {"#bigint": "42"},
    },
    {
      "#meta": {"no": 1},
      "pc": "lock",
      "x": {"#bigint": "43"},
    }
  ],
}
```

Apalache produces ITF traces

Trivial to parse:

see github.com/konnov/itf-py