

Mysteries of Dropbox

Property-Based Testing of a Distributed
Synchronization Service

John Hughes, Benjamin Pierce,
Thomas Arts, Ulf Norell



Synchronization Services



400 million (June 2015)



240 million (Oct 2014)




250 million (Nov 2014)

Are they trustworthy?

(exactly!)

What do they do?

Can we test them?



TEST ING



Writing test cases by hand

(especially for testing
distributed systems!)



Generate test cases from a model

Our Goals

- Develop a **precise specification** of the core behavior of a synchronization service
 - Phrased from the perspective of *users*
 - Applicable to a variety of different synchronizers
- Use **property-based random testing** to validate it against Dropbox's observed behavior

Why Generate Tests?

- Much wider variety!
 - Crucial for effective testing of distributed services
 - Subtle edge cases, timing dependencies,
...
- More confidence!

QuickCheck




1999—invented by Koen Claessen and John Hughes,
for Haskell

2006—Quviq founded marketing Erlang version

Many extensions

Finding deep bugs for Ericsson, Volvo Cars, Basho, etc...

QuickCheck



Test = list of *operations*

System
under test

Op_1

Op_2

Op_3

Each operation gives
rise to an *observation*

Obs_1

Obs_2

Obs_3

Model

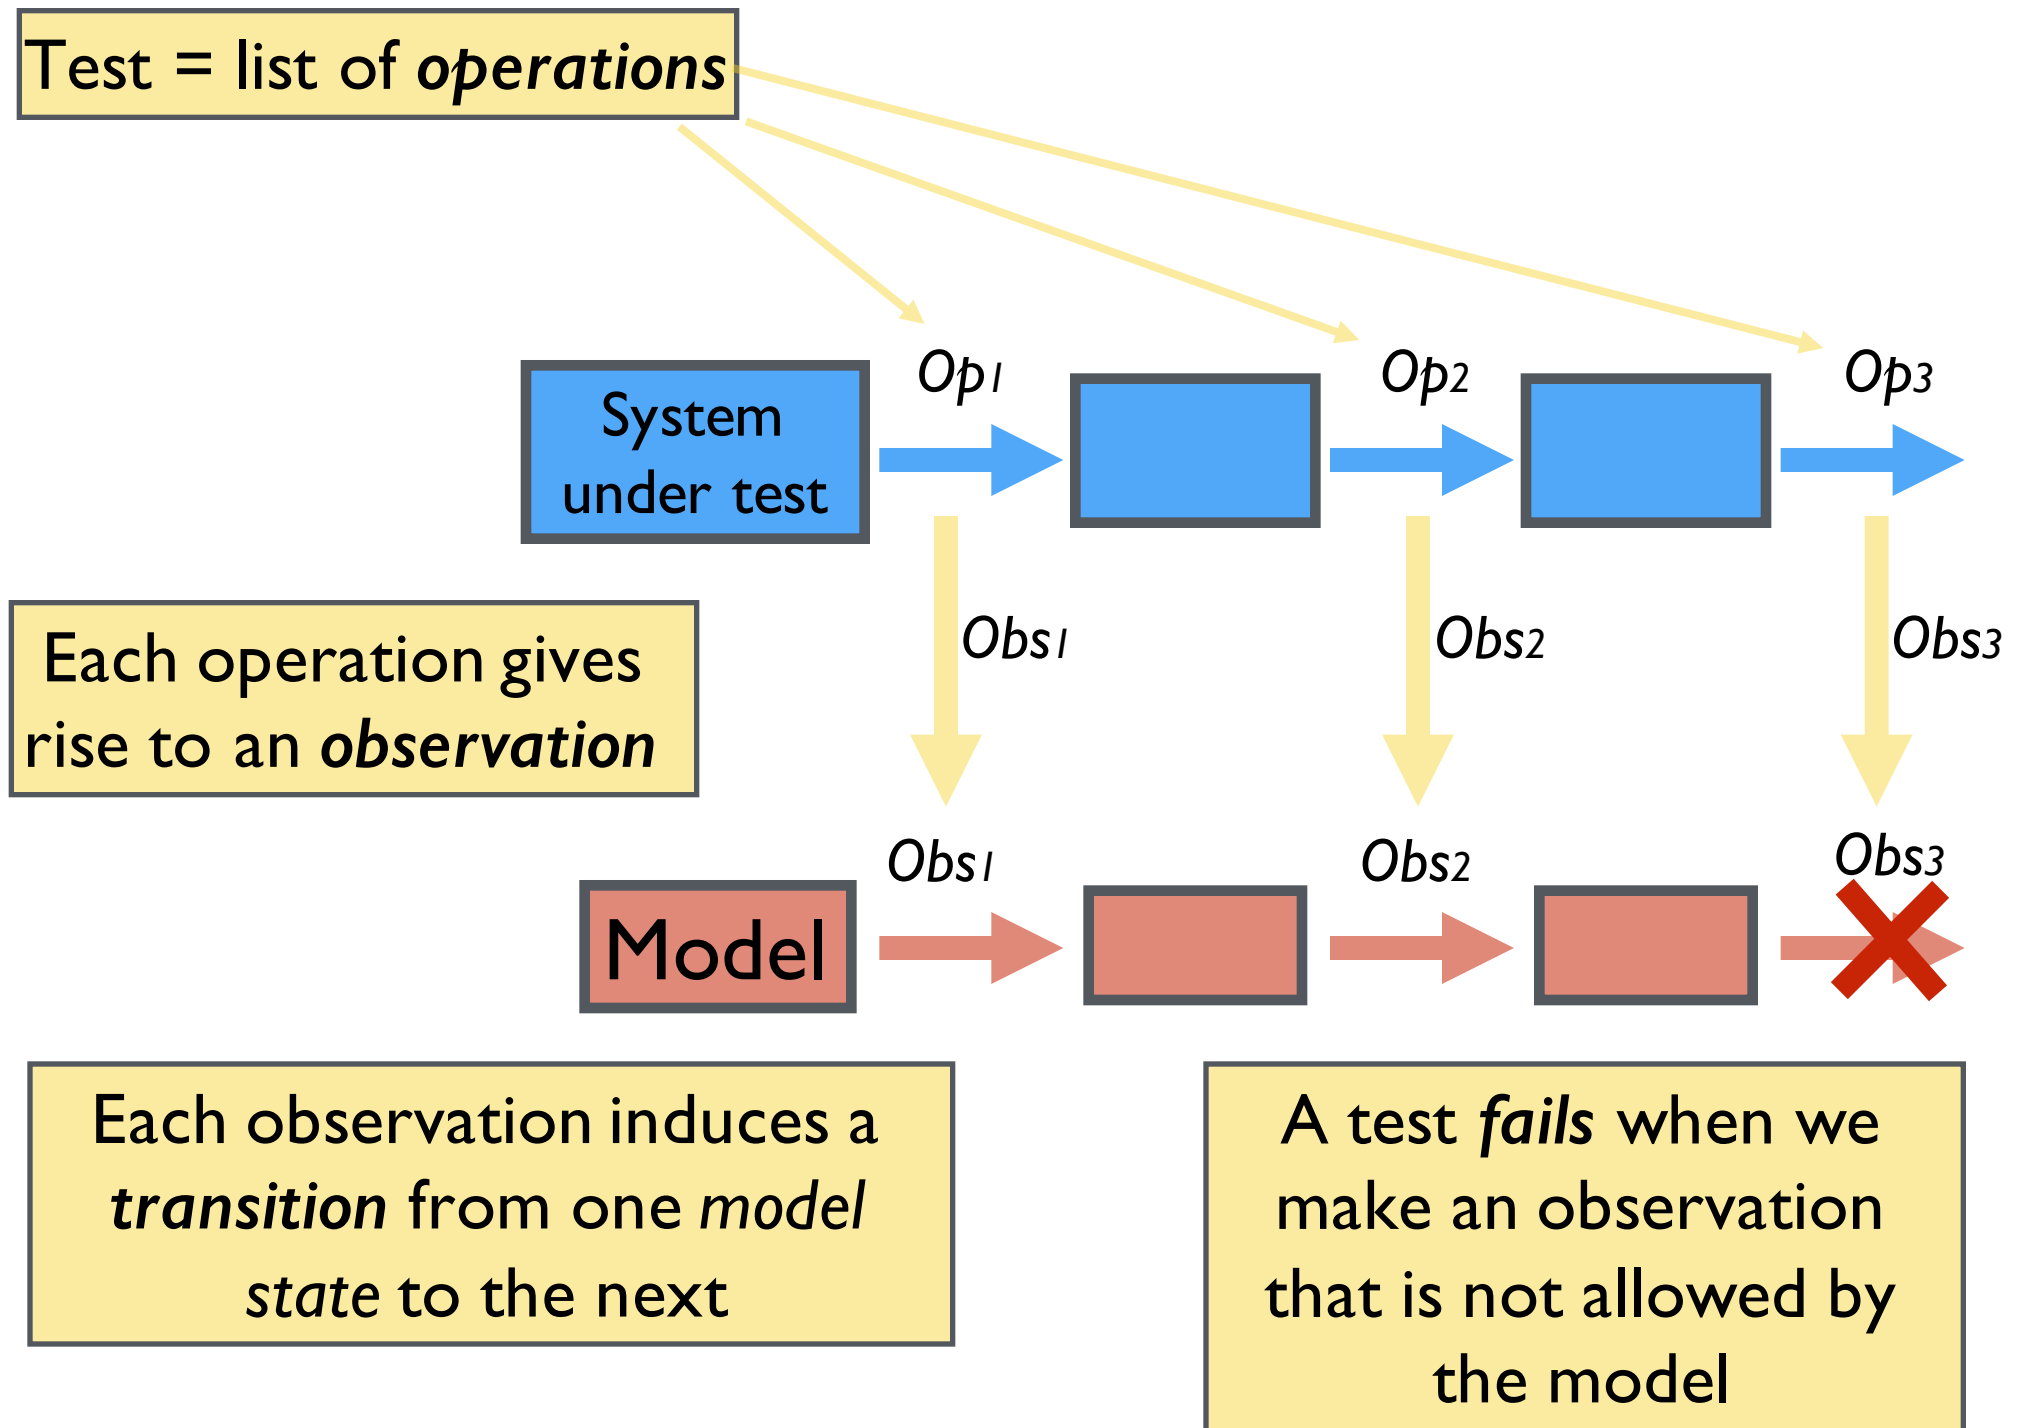
Obs_1

Obs_2


Obs_3

Each observation induces a
transition from one *model*
state to the next

A test *fails* when we
make an observation
that is not allowed by
the model



Test Harness for Dropbox



What operations and
observations do we
need?

One Simplification...

- Real filesystem APIs are complex
 - Files, directories, timestamps, permissions, extended attributes, symlinks, hard links, ...
- We make a small restriction...

Filesystem = single file

Operations	Observations
$READ_N$	$READ_N \rightarrow \text{"current value"}$
$WRITEN(\text{"new value"})$	$WRITEN(\text{"new value"}) \rightarrow \text{"old value"}$

Use special value \perp for “no file”

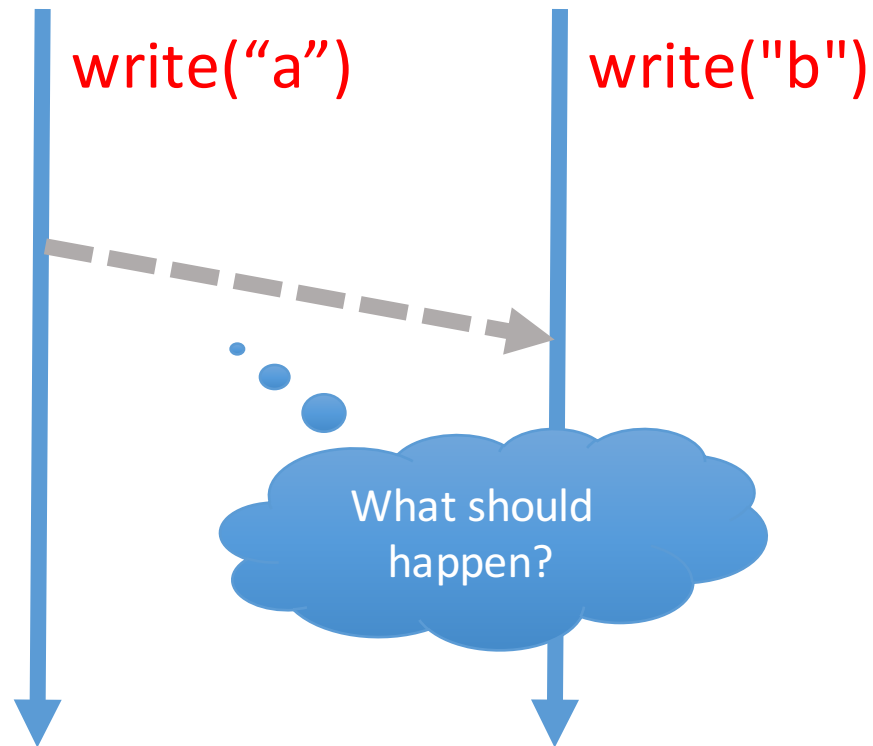
$READ_N \rightarrow \perp$

$WRITEN(\perp)$

means that the file is missing

means delete the file

Challenge #1: Conflicts



Dropbox's answer:

The first value to reach the server wins; other values are moved to *conflict files* in the same directory.

However, these conflict files may not appear for a little while!


Second try...

Operations	Observations
READ _N	READ _N → “current value”
WRITEN (“new value”)	WRITEN (“new value”) → “old value”
STABILIZE	STABILIZE → (“value”, {“conflict values”})

Same value in the file on all clients



Same set of values in conflict files on all clients



Challenge #2: Background operations

- The Dropbox client communicates with the test harness via the filesystem.

But...

- The Dropbox client *also* communicates with the Dropbox servers!
 - Timing of these communications is unpredictable

Invisible, unpredictable activity → Nondeterminism!


Approach

- Model the whole system state *including the (invisible) state of the server*
- Add "conjectured observations" to the ones we actually observe when running tests...

Operations	Observations
$READ_N$	$READ_N \rightarrow \text{"current value"}$
$WRITE_N(\text{"new value"})$	$WRITE_N(\text{"new value"}) \rightarrow \text{"old value"}$
STABILIZE	$STABILIZE \rightarrow (\text{"value"}, \{\text{"conflict values"}\})$
	UP_N
	$DOWN_N$

node N uploads its value to the server

node N is refreshed by the server



starting state

all possible sequences of Up/Downs

hypothetical states

*real observation (invalid in most
hypothetical states)*

etc.

Explanation

**No explanation
= failing test**

Example:


Test

Client 1	Client 2
WRITE 'a'	WRITE 'b'
READ	WRITE 'c'
STABILIZE	

Test

Client 1	Client 2
WRITE 'a'	WRITE 'b'
READ	WRITE 'c'
STABILIZE	

Observations



Client 1	Client 2
WRITE 'a' $\rightarrow \perp$	WRITE 'b' \rightarrow 'a'
READ \rightarrow 'b'	WRITE 'c' \rightarrow 'b'
STABILIZE \rightarrow ('c', \emptyset)	

Test

Client 1	Client 2
WRITE 'a'	WRITE 'b'
READ	WRITE 'c'
STABILIZE	

Explanation

Client 1	Client 2
WRITE 'a' $\rightarrow \perp$ UP	DOWN WRITE 'b' \rightarrow 'a' UP
DOWN	WRITE 'c' \rightarrow 'b' UP
READ \rightarrow 'b' DOWN	
STABILIZE \rightarrow ('c', \emptyset)	


Observations

Client 1	Client 2
WRITE 'a' $\rightarrow \perp$	WRITE 'b' \rightarrow 'a'
READ \rightarrow 'b'	WRITE 'c' \rightarrow 'b'
STABILIZE \rightarrow ('c', \emptyset)	

Example:

Using the model for testing

1. Generate a random sequence of operations $Op_1 \dots Op_n$
2. Apply them to the system under test, yielding observations $Obs_1 \dots Obs_n$
3. Calculate *all* ways of interleaving Up and Down observations with $Obs_1 \dots Obs_n$
4. For each of these, check whether



is a valid sequence of transitions of the model

4. If the answer is “no” for every possible interleaving, we have found a failing test; otherwise, repeat

Model states

- *Stable value* (i.e., the one on the server)
- *Conflict set* (only ever grows)
- For each node:
 - Current *local value*
 - "FRESH" or "STALE"
 - "CLEAN" or "DIRTY"

i.e., has the global value changed since this node's last communication with the server

i.e., has the local value been written since this node was last refreshed by the server

Modeling the operations

READ $\rightarrow V$

Precondition: $LocalVal_N = V$

Effect: none

WRITE $V_{new} \rightarrow V_{old}$

Precondition: $LocalVal_N = V_{old}$

Effect: $LocalVal_N \leftarrow V_{new}$

$Clean?_N \leftarrow \text{DIRTY}$

Modeling the operations

STABILIZE $\rightarrow (V, C)$

Precondition: $ServerVal = V$

$Conflicts = C$

for all N , $Fresh?_N = \text{FRESH}$

$Clean?_N = \text{CLEAN}$

Effect: none

Modeling the operations

DOWN

Precondition: $Fresh?_N = \text{STALE}$

$Clean?_N = \text{CLEAN}$

Effect: $LocalVal_N \leftarrow ServerVal$

$Fresh?_N \leftarrow \text{FRESH}$

Modeling the operations

UP

Precondition: $Clean?_N = \text{DIRTY}$

Effect: $Clean?_N \leftarrow \text{CLEAN}$

if $Fresh?_N = \text{FRESH}$ then

if $LocalVal_N \neq ServerVal$ then

$Fresh?_{N'} \leftarrow \text{STALE}$ for all $N' \neq N$

$ServerVal \leftarrow LocalVal_N$

else

if $LocalVal_N \neq ServerVal$ then

$Conflicts \leftarrow Conflicts \cup \{LocalVal_N\}$

Dealing with deletion

- Deletion can easily be added to the model:

DELETE_N just means $\text{WRITE}_N \perp$

- Try adding this and run some tests...

Still not quite right...

Client 1	Client 2	Client 3
<p>WRITE 'a' $\rightarrow \perp$</p> <p>WRITE $\perp \rightarrow$ 'a'</p> <p>READ \rightarrow 'b'</p>	<p>READ \rightarrow 'a'</p> <p>READ $\rightarrow \perp$</p>	<p>WRITE 'b' \rightarrow 'a'</p>

Write "a" on client 1

Delete the file

We now observe "b", so the stable value on the server must have been overwritten, despite the fact that 'b' was in conflict

Client 2 sees 1's value

2 sees "missing" (so stable value at server is "missing")

Now client 3 writes "b". Observes previous value 'a' (n.b.: not \perp).

Refining the specification...

- Add special cases for “missing” in Up and Down actions:
 - When “missing” encounters another value during an up or down, the other value always wins
 - I.e., when a write and a delete conflict, the delete gets undone

UP

Precondition: $Clean?_N = \text{DIRTY}$

Effect: $Clean?_N \leftarrow \text{CLEAN}$

if $Fresh?_N = \text{FRESH}$ then

if $LocalVal_N \neq ServerVal$ then

$Fresh?_{N'} \leftarrow \text{STALE}$ for all $N' \neq N$

$ServerVal \leftarrow LocalVal_N$

else


if $LocalVal_N \notin \{ServerVal, \perp\}$ then

$Conflicts \leftarrow Conflicts \cup \{LocalVal_N\}$

Surprises...

Surprise: Dropbox can (briefly) delete a newly created file...

	Client 1	Client 2	
Create file Delete it	WRITE 'a' $\rightarrow \perp$ WRITE $\perp \rightarrow$ 'a'	WRITE 'b' \rightarrow 'a'	Observe creation
Create it again File is gone!	WRITE 'c' $\rightarrow \perp$ READ $\rightarrow \perp$		

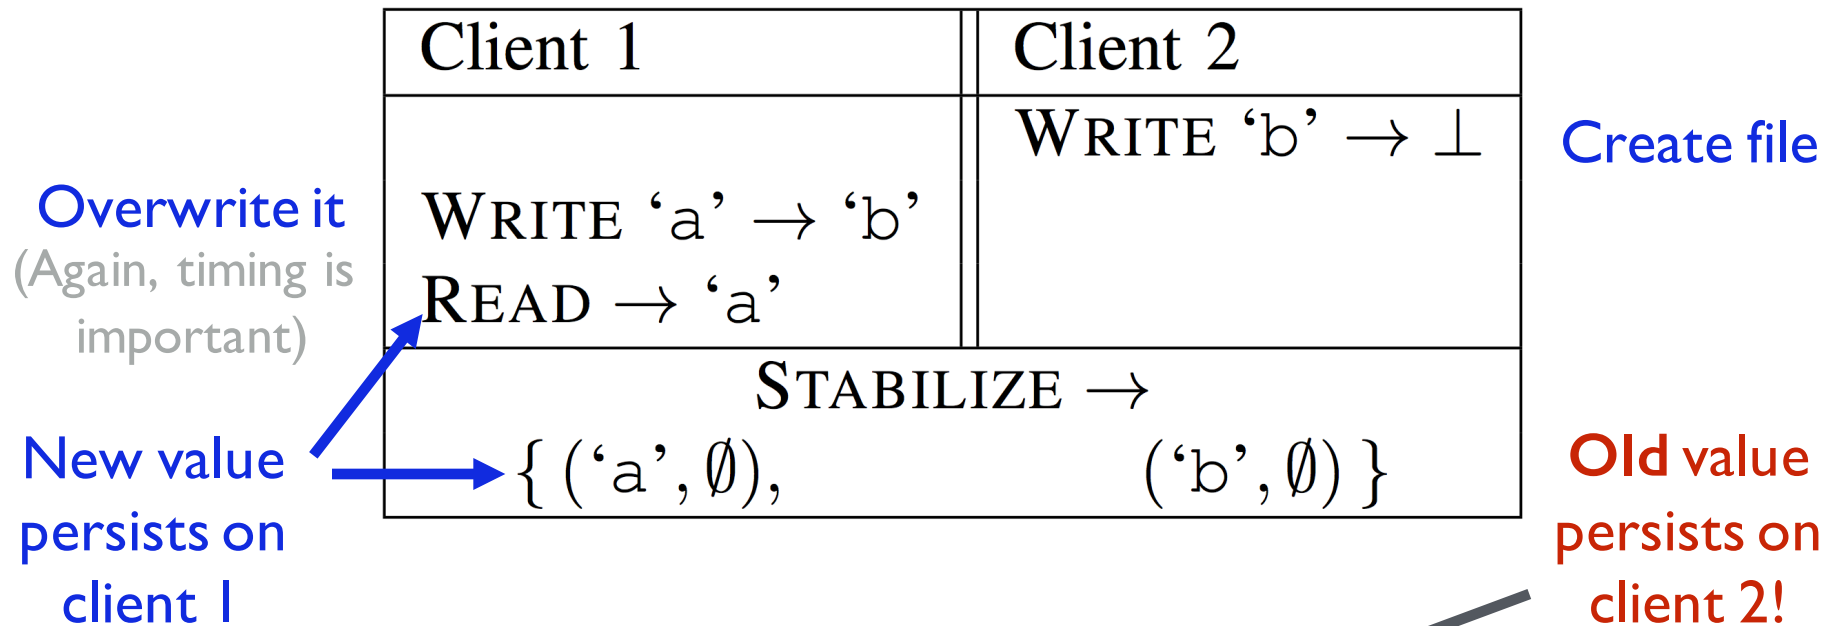
Timing is critical!  Add **SLEEP** operations
in tests

Surprise: Dropbox can (permanently)
re-create a deleted file...

	Client 1	(other clients idle)
Create file	WRITE 'b' $\rightarrow \perp$	
Delete it	WRITE $\perp \rightarrow$ 'b'	
File is back!	READ \rightarrow 'b'	

(Again, timing is critical)

Surprise: Dropbox can lose data



Client 1 believes it is still Fresh, so if we later write a new value on client 2, it will silently overwrite client 1's value and no conflict file will be created

Wrapping up...

What did we do?

- Tested a non-deterministic system by *searching for explanations* using a model with hidden actions
- Used QuickCheck's minimal failing tests to *refine* the model, until it matched the intended behaviour
- Now minimal failing tests reveal *unintended* system behaviour

What do Dropbox say?

- The synchronization team has reproduced the buggy behaviours
- They're *rare* failures which occur under very special circumstances
- They're developing fixes

Synchronization is subtle!

There's much more to do...

- Add directories!
 - Directories and files with the same names
 - Conflicts between deleting a directory and writing a file in it
 - ...
- More file synchronizers!

Thank you!

(Any questions?)

