

Weeks of Debugging can save you hours of TLA⁺

...building a ~~distributed~~ cloud system

Markus A. Kuppe – Principal Engineer - Microsoft Research

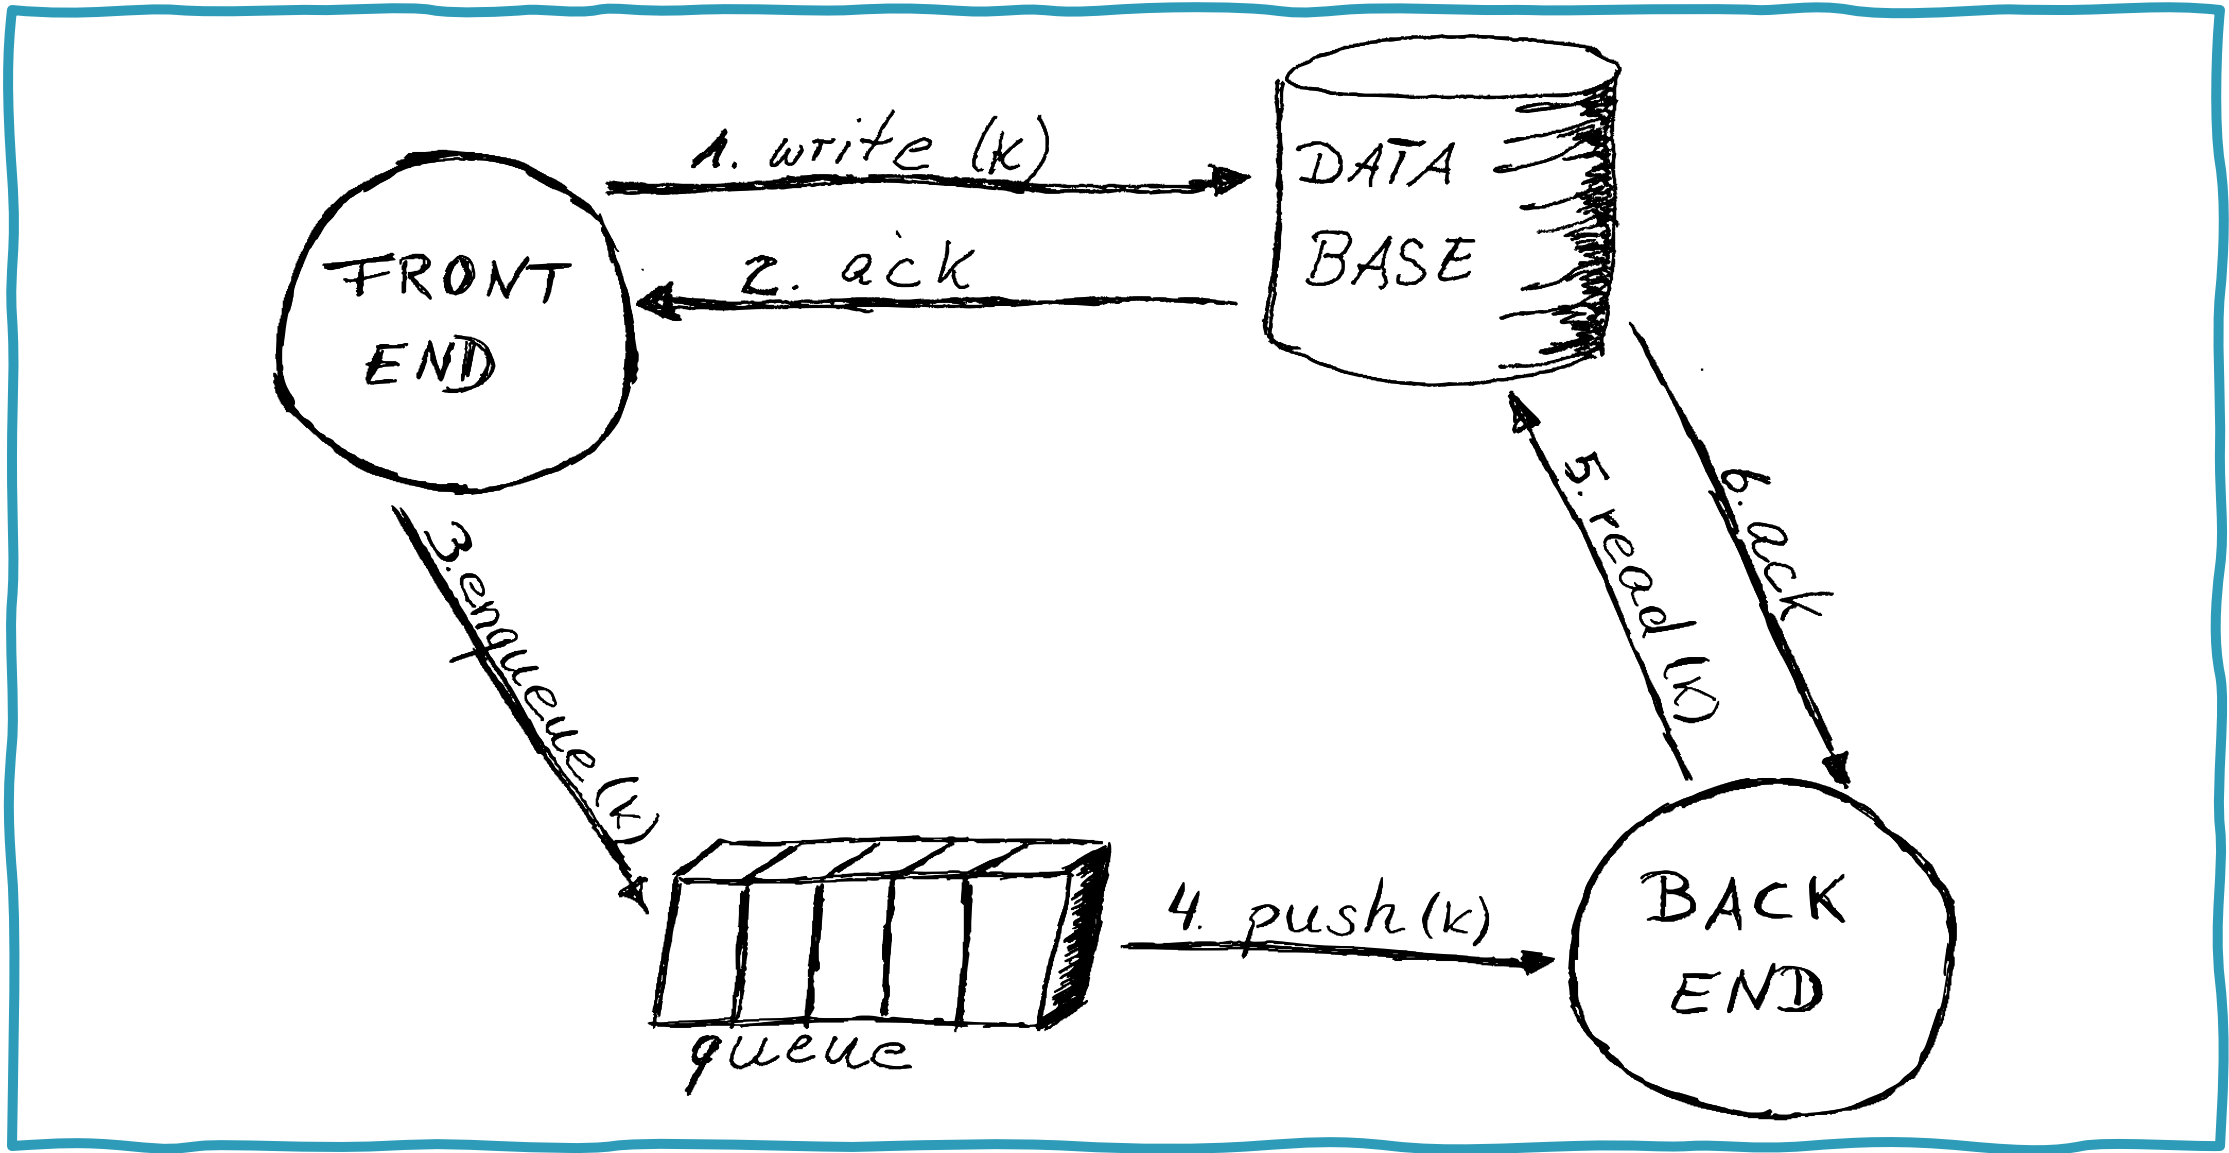
Your average cloud system

- A client (frontend) issuing user requests
- Some database system for persistence
- A message queue/broker to decouple frontend and backend
- A backend system

Your average workflow

1. User places an order
2. Frontend writes the order to the database
3. ACKs the order to the user
4. Frontend publishes a message
5. Backend receives message, reads order from database, and fulfills it





Let's build our first prototype




Get a subscription for your favorite cloud provider



Set up infrastructure: frontend, database, queue, backend server



Implement and deploy the workflow



Test!!!

VARIABLES

* database

log,

* frontend

requests, future,

* backend

backend,

* message queue

queue

Consider this to be akin to a transaction; it either happens or it doesn't happen at all. The next state of our system is defined by the primed variables!

12FrontendWrite \triangleq

* Precondition: No pending wr
 \wedge future = Nil

* Write request to database

$\wedge \exists o \in \text{requests}$:

DB!WriteInit(o.key, o.val, LAMBDA t: future' = t)

* Nothing else happens.

$\wedge \text{requests}' = \text{requests} \wedge \text{queue}' = \text{queue} \wedge \text{backend}' = \text{backend}$

(No worries, log' defined inside of
WriteInit)

34FrontendEnqueue \triangleq

* Precondition: Pending write that succeeded
 \wedge future \neq Nil \wedge DB!WriteSucceeded(future)

* Acknowledge write to user
 \wedge requests' =
requests \ { [key \mapsto future.key, val \mapsto future.value] }

* Enqueue msg containing key to backend
 \wedge queue' = Append(queue, [k \mapsto future.key, t \mapsto Nil])

* Clear pending write
 \wedge future' = Nil

* Nothing else happens.
 \wedge UNCHANGED \langle log, backend \rangle

56BackendRead \triangleq

* Precondition: Queue contains a msg
 \wedge queue $\neq \langle \rangle$

* Receive msg from queue
 \wedge LET msg \triangleq Head(queue) IN

* Read from database (with the given key)
 $\exists o \in \text{DB!EventualConsistencyRead}(msg.k):$
backend' = backend $\cup \{o.value\}$

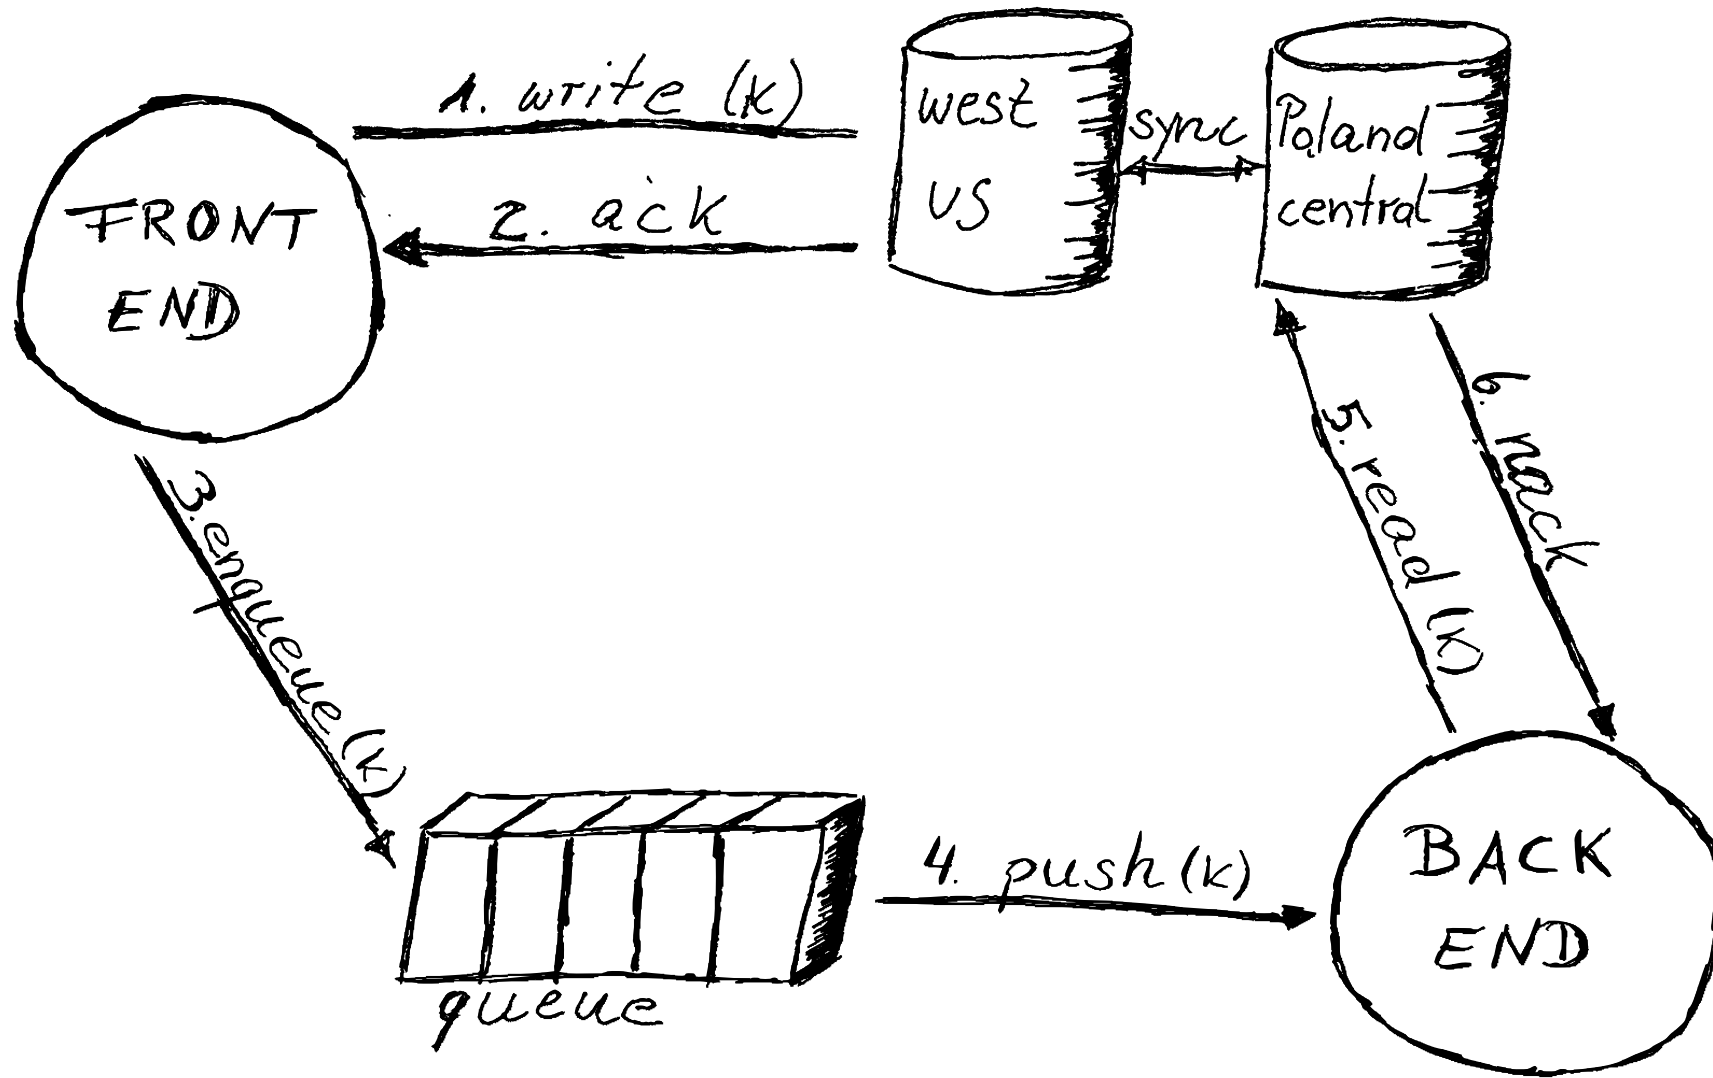
* Remove msg from queue
 \wedge queue' = Tail(queue)

\wedge UNCHANGED $\langle \log, requests, future \rangle$

The “Test”

The diamond says that req.val will be **eventually** an element of backend. We don't have to define something like timeouts that leads to slow and brittle tests.


BackendReadsCorrectValue \triangleq
 $\forall \text{ req} \in (\text{AllRequests} \setminus \text{requests})$:
 $\diamond(\text{req.val} \in \text{backend})$



Demo Part II

But at scale!

- More frontend & backend servers
- ...
- More queues and databases?
 - Sharding & Partitioning
 - Relax database consistency



Demo C

Pants Down!

Where is the database code?

A database is a just log on steroids!

Database code just 200 LoCs, four variables, and most of the “magic” in the *Read* operator.

```

GeneralRead(ke
LET maxCandi
  ∧ log[i].ke
  ∧ i ≤ index
allIndices ≜
  ∧ allowDir
  ∧ log[i].ke
  ∧ i > index
IN { [logIndex
  : i ∈
  maxCandi
WHEN {NotF
ELSE {}))

```

Understanding Inconsistency in Azure Cosmos DB with TLA+

Finn Hackett
University of British Columbia
Vancouver, Canada
fhackett@cs.ubc.ca

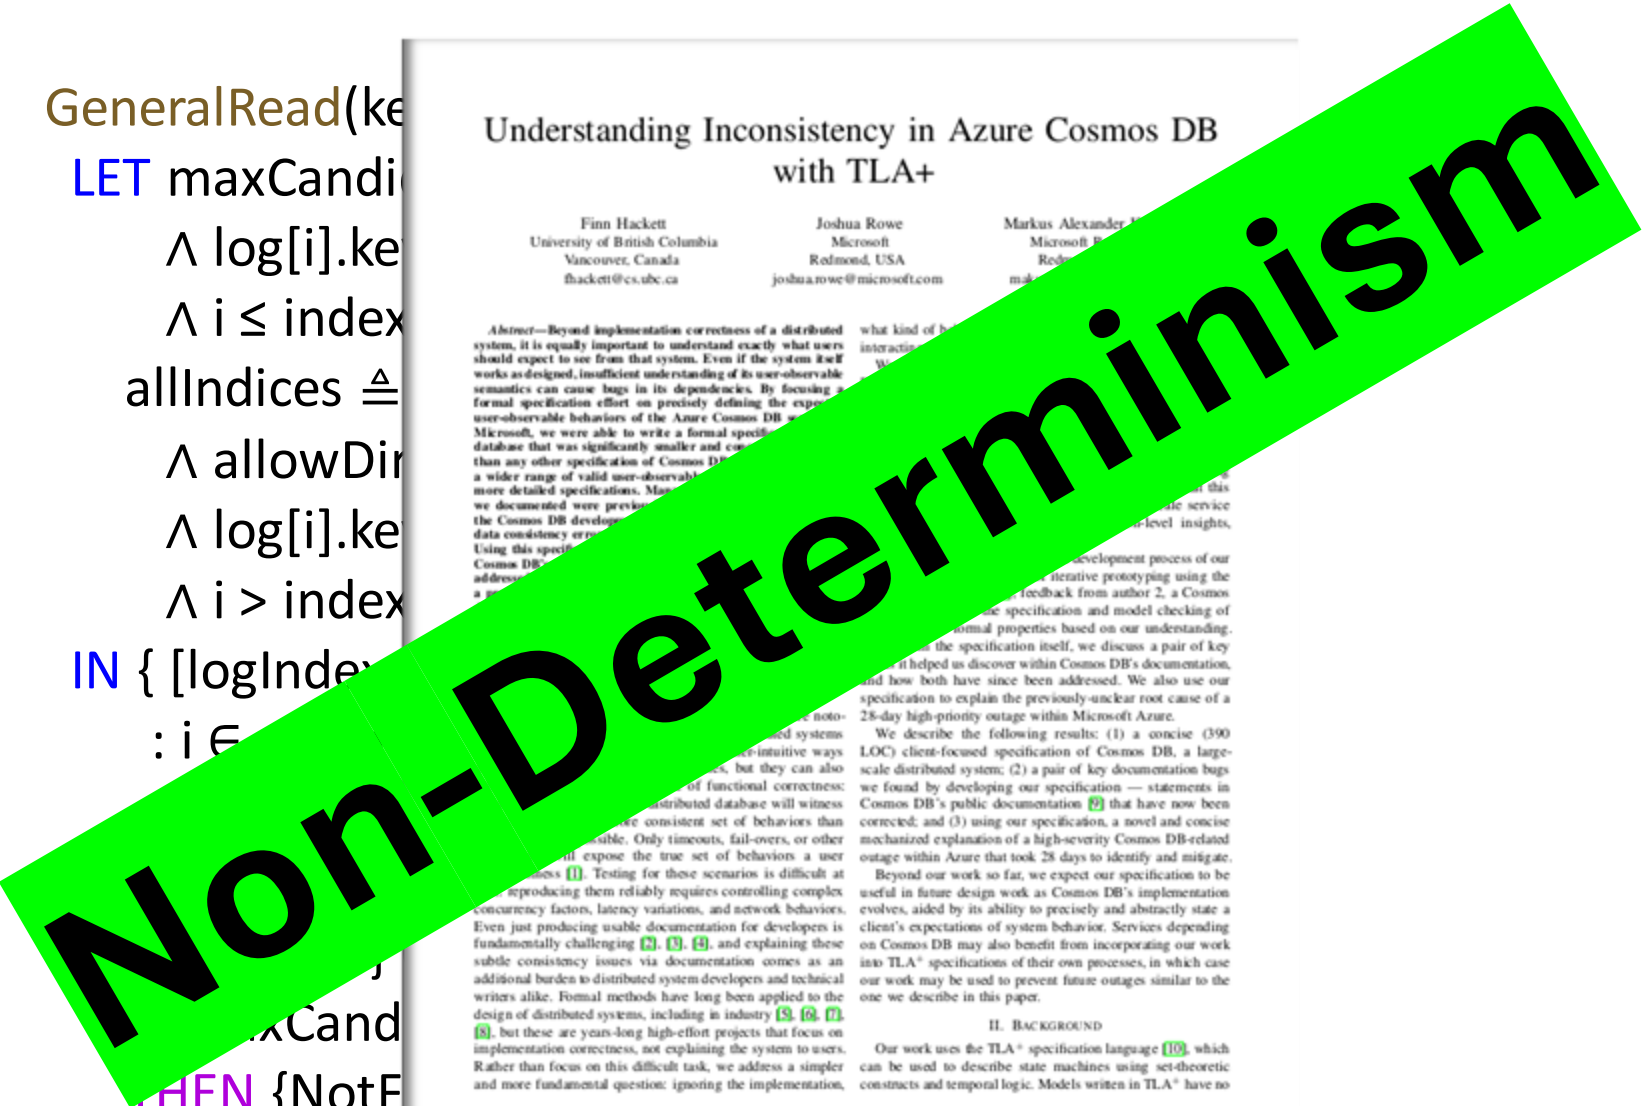
Joshua Rowe
Microsoft
Redmond, USA
joshua.rowe@microsoft.com

Markus Alexander
Microsoft
Redmond, USA
markus.alexander@microsoft.com

Abstract—Beyond implementation correctness of a distributed system, it is equally important to understand exactly what users should expect to see from that system. Even if the system itself works as designed, insufficient understanding of its user-observable semantics can cause bugs in its dependencies. By focusing a formal specification effort on precisely defining the expected user-observable behaviors of the Azure Cosmos DB service at Microsoft, we were able to write a formal specification of the database that was significantly smaller and easier to understand than any other specification of Cosmos DB. This specification covers a wider range of valid user-observable behaviors than any of the more detailed specifications, Microsoft has ever written. In this paper, we document the development process of this specification, the Cosmos DB development process, and the impact of this data consistency error on the development process of our service. Using this specification, we discovered several bugs in the Cosmos DB development process of our service, including a 28-day high-priority outage within Microsoft Azure. We describe the following results: (1) a concise (390 LOC) client-focused specification of Cosmos DB, a large-scale distributed system; (2) a pair of key documentation bugs we found by developing our specification — statements in Cosmos DB’s public documentation that have now been corrected; and (3) using our specification, a novel and concise mechanized explanation of a high-severity Cosmos DB-related outage within Azure that took 28 days to identify and mitigate. Beyond our work so far, we expect our specification to be useful in future design work as Cosmos DB’s implementation evolves, aided by its ability to precisely and abstractly state a client’s expectations of system behavior. Services depending on Cosmos DB may also benefit from incorporating our work into TLA+ specifications of their own processes, in which case our work may be used to prevent future outages similar to the one we describe in this paper.

H. BACKGROUND

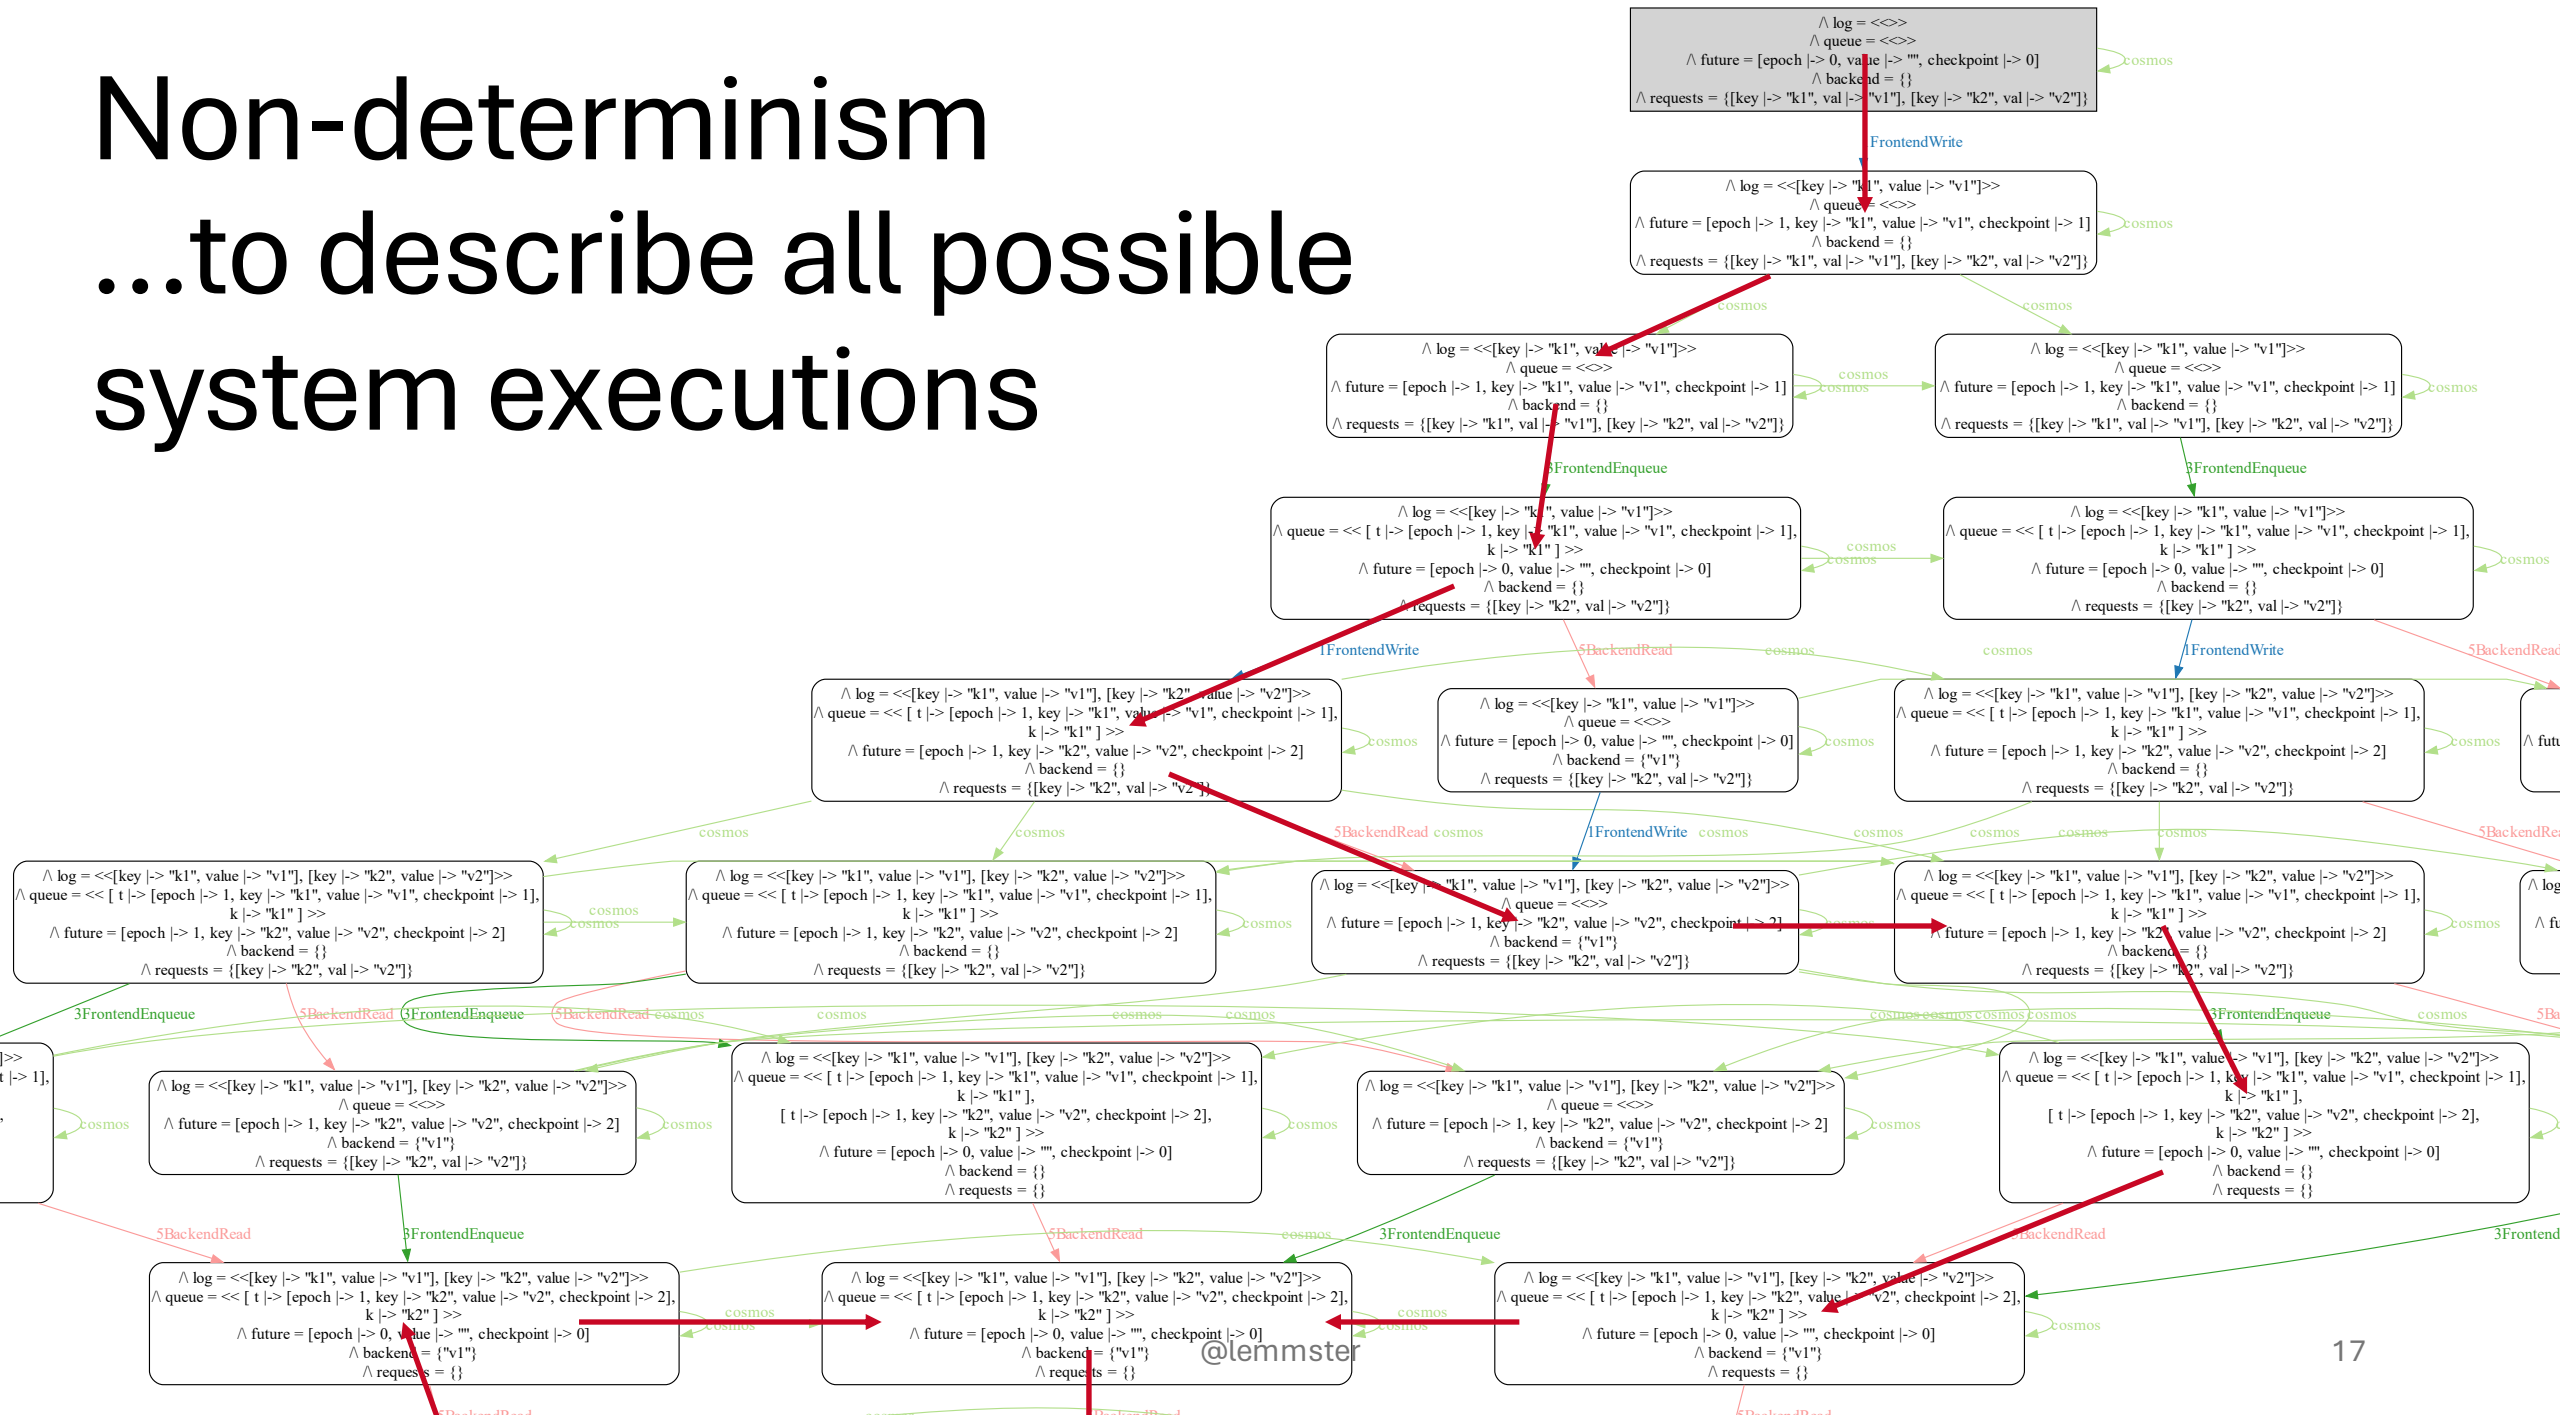
Our work uses the TLA+ specification language [10], which can be used to describe state machines using set-theoretic constructs and temporal logic. Models written in TLA+ have no



<https://doi.org/10.48550/arXiv.2210.13661>

@lemmster

Non-determinism ...to describe all possible system executions



Application of TLA+

- Consensus (Raft, Paxos, ...)
- Distributed Databases (Cosmos DB, DynamoDB, Mongo DB, ...)
- Hardware (cache coherence, ...)
- Network protocols (p2p, ...)
- Concurrency primitives (queues, critical section, glibc condition variable, ...)
- Cloud Security (Software Defined *)
- Dogfood
- ...
- Polymerase Chain Reaction (PCR)




Temporal Logic of Actions+

TLA+ is a *specification* language to design, document, and verify reactive systems.

Learn more at <http://tlapl.us> & <http://examples.tlapl.us>

[TLA+ Webinar June 14th](#)





Question & Answers

