

# An Epistemic Perspective on Consistency of Concurrent Computations

Klaus v. Gleissenthall<sup>1</sup> and Andrey Rybalchenko<sup>1,2</sup>

<sup>1</sup> TUM


<sup>2</sup> MSR Cambridge

# Consistency Properties

- ▶ Whenever read or write **shared data** at **different location**
- ▶ Need to arrive **consensus**
- ▶ Specify how much **inconsistency** is tolerated before

# Consistency Properties

## Memory



Processor  
Cache



Sequential Consistency/  
TSO/  
PSO ...


# Consistency Properties

## Concurrent Data-structures

Non-  
blocking

Cached  
Values

Queue/  
Set/...



`java.util.concurrent`

Linearizability


# Consistency Properties

## Geo-replicated Databases

Different  
revisions



NoSQL




Eventual Consistency

# Consistency Properties



Sequential Consistency/  
TSO ...



Linearizability




Eventual Consistency

Different formalisms:

permutation / partial-order / operational

# Knowledge Perspective



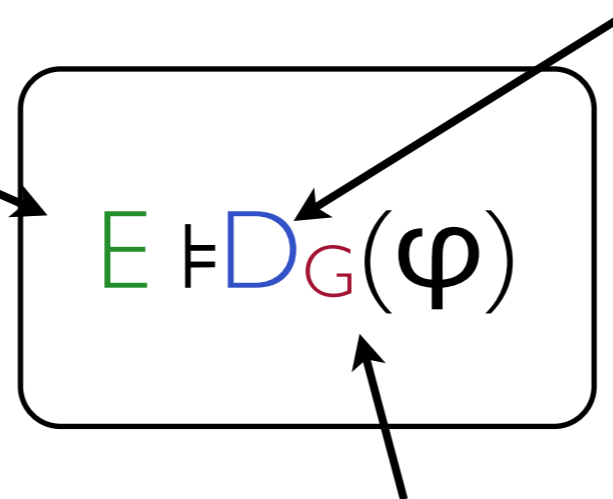
Distributed systems

own conference series (TARK)

# Distributed Knowledge

Distributed  
Knowledge

trace  $E$



Share everything  
you know

group  $G$  of  
participants

Halpern and Moses,  
JACM, 1990



# What do we get?

- ▶ Very similar form:

For memory:  
read last write

$$E \models \neg D_G(\neg \text{correct})$$

don't get caught  
cheating

Sequential  
specification  
of shared data-structure

# What do we get?

## ► Compare Conditions:

Sequential consistency:  $E \models \neg D_{\text{THREADS}}(\neg \text{correct})$

Eventual consistency:  $E \models \neg D_{\text{THREADS}}(\neg \text{correctEVC})$

Linearizability:  $E \models \neg D_{\text{THREADS} \cup \{\text{obs}\}}(\neg \text{correct})$

# What do we get?

## ► Compare Conditions:

reduction

Sequential consistency:  $E \models \neg D_{\text{THREADS}}(\neg \text{correct})$




Eventual consistency:  $E \models \neg D_{\text{THREADS}}(\neg \text{correctEVC})$



Linearizability:  $E \models \neg D_{\text{THREADS}}(\text{u}\{\text{obs}\})(\neg \text{correct})$

additional  
knowledge



# Outline

- ▶ Sequential Consistency
- ▶ Epistemic Knowledge
- ▶ Eventual Consistency
- ▶ A Theorem

# Sequential-Consistency

# Sequential Consistency

arranged by order  
of return

reason about  
traces



$E := (t_2, ld(0)) (t_2, ld(1)) (t_1, st(1))$

is this consistent ?

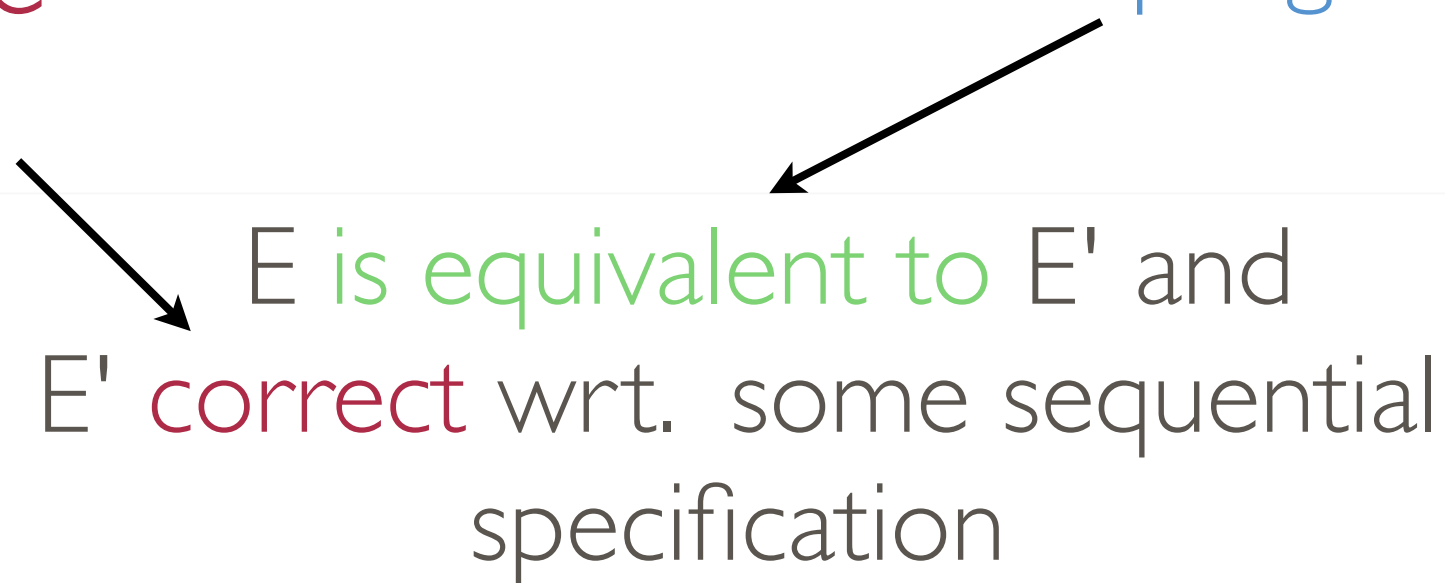
# Sequential Consistency

The **result of any execution** is the same as if the operations of all the processors were executed **in some sequential order**, and the operations of each individual processor appear in this sequence **in the order specified by its program**.

# Sequential Consistency

reads return  
last write

permutations preserving  
program order



$E$  is equivalent to  $E'$  and  
 $E'$  correct wrt. some sequential  
specification

Alternative def. from literature



# Sequential Consistency

$E := (t_2, ld(0)) (t_2, ld(1)) (t_1, st(1))$



reads return  
last write

$E' := (t_2, ld(0)) (t_1, st(1)) (t_2, ld(1))$



# Knowledge Perspective

t1 “*knows*” that it  
stored 1

t2 “*knows*” that it  
loaded 0 and then 1

$$E := (t2, ld(0)) (t2, ld(1)) (t1, st(1))$$

but: *nothing* about the  
other thread

# Knowledge Perspective

$\{t1, t2\}$  know the conjunction of  
these facts

$E := (t2, Id(0)) (t2, Id(1)) (t1, st(1))$

i.e., which operations were performed  
but not their order

# Knowledge Perspective

read  
last write

$$E \models \neg D_{\{t_1, t_2\}}(\neg \text{correctMem})$$

sequentially  
consistent

$$E := (t_2, \text{ld}(0)) (t_2, \text{ld}(1)) (t_1, \text{st}(1))$$

# Local views

$E := (t2, Id(0)) (t2, Id(1)) (t1, st(1))$

↓ ↓t1

$(t1, st(1))$

↓ ↓t2

$(t2, Id(0)) (t2, Id(1))$

projection  
onto  
t2's view

# Indistinguishability

$$E \sim_t E' \text{ iff: } E \downarrow t = E' \downarrow t$$

$$E := (t_2, \text{ld}(0)) (t_2, \text{ld}(1)) (t_1, \text{st}(1))$$

$\sim_{t_1}$  ✓

$\sim_{t_2}$  ✗

$$E' := (t_2, \text{ld}(0)) (t_1, \text{st}(1))$$

# Group Indistinguishability

$$\sim_G := \left( \bigcap_{a \in G} \sim_a \right)$$

$$E := (t_2, \text{Id}(0)) (t_2, \text{Id}(1)) (t_1, \text{st}(1))$$

$$\sim \{t_1, t_2\} \times$$

$$E' := (t_2, \text{Id}(0)) (t_1, \text{st}(1))$$

# Group Indistinguishability

$$\sim_G := \left( \bigcap_{a \in G} \sim_a \right)$$

$$E := (t_2, \text{Id}(0)) (t_2, \text{Id}(1)) (t_1, \text{st}(1))$$

$$\sim \{t_1, t_2\} \quad \checkmark$$

$$E' := (t_2, \text{Id}(0)) (t_1, \text{st}(1)) (t_2, \text{Id}(1))$$



# Distributed Knowledge

threads **can't tell** which trace  
they really saw

$E \models D_G(\varphi)$  :iff for all  $E'$  s.t.  $E \sim_G E'$  :  $E' \models \varphi$

if  $\varphi$  holds **for all** those traces,  
they know  $\varphi$

# Sequential Consistency

$E \models D_G(\varphi)$  :iff for all  $E'$  s.t.  $E \sim_G E'$ :  $E' \models \varphi$



$E \models \neg D_{\text{THREADS}}(\neg \text{correct})$   
:iff exists  $E'$  s.t.  $E \sim_{\text{Threads}} E'$  and  $E' \models \text{correct}$

The threads **can't tell**  $E$   
from a correct trace

A **good reason** to accept it!

# Sequential Consistency

$E$  is equivalent to  $E'$  and  
 $E'$  correct wrt. some sequential  
specification

$E \models \neg D_{\text{THREADS}}(\neg \text{correct})$   
:iff exists  $E'$  s.t.  $E \sim_{\text{Threads}} E'$  and  $E' \not\models \text{correct}$

# Eventual-Consistency

# Eventual Consistency

- ▶ **Geo-replicated** database systems  
( Google/Facebook ...)
- ▶ Different location need to maintain consistent view of data
- ▶ However must be **highly available**
- ▶ Minimize synchronization, allow updates **any time**

# Original Definition

## Existence of two orders

**Definition 4 (Eventual Consistency).** We adapt the definition presented in [4] to our notation. A trace  $E \in \mathcal{S}^\infty$  is eventually consistent ( $ecCons(E)$ ) if and only if there exist a partial order  $<_v$  (visibility order), and a total order  $<_a$  (arbitration order) on the events in  $set(E)$  such that:

- $<_v \subseteq <_a$  (arbitration extends visibility).
- $<_p \subseteq <_v$  (visibility is compatible with program-order).
- for each  $e_q = (t, qu(id, q, r)) \in E$ , we have  $r = apply(\{e \mid e <_v e_q\}, <_a, s_0)$  (consistent query results).
- $<_a$  and  $<_v$  factor over  $\equiv_t$  (atomic revisions).
- if  $(t, com(id)) \notin E$  and  $(t, -(id, -)) <_v (t', -)$  then  $t = t'$  (uncommitted updates).
- if  $e = (t, com(id)) \in E$  then there are only finitely many  $e' := (t', com(id'))$  such that  $e' \in E$  and  $e \not<_v e'$  (eventual visibility).

# Our Version

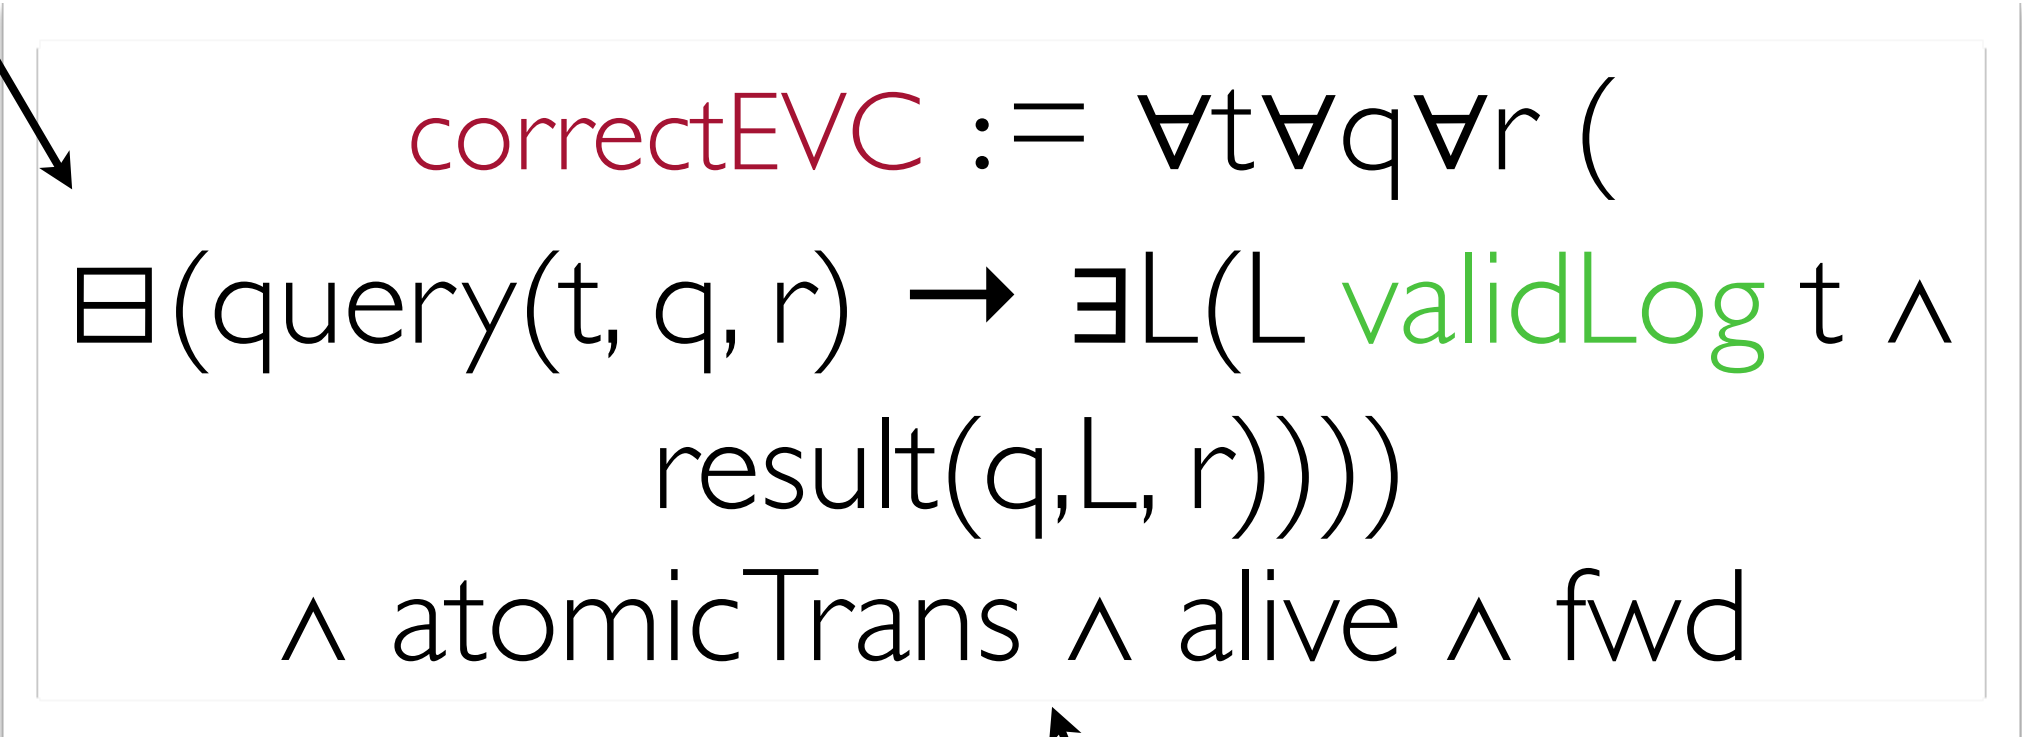

Participants don't know it  
violates


$$E \text{ F } \neg D_{\text{THREADS}}(\neg \text{correctEVC})$$

temporal/sequential specification

# Eventual Consistency

“so far”


$$\text{correctEVC} := \forall t \forall q \forall r ( \\ \exists L (L \text{ validLog } t \wedge \\ \text{result}(q, L, r))) \\ \wedge \text{atomicTrans} \wedge \text{alive} \wedge \text{fwd}$$



results are justified  
by consistent logs



# Eventual Consistency

store/load

t knows a



$L \text{ validLog } t := \forall a(a \text{ in } L \leftrightarrow t \text{ knows } a)$   
 $\wedge \text{ consistent } (L)$

log order matches  
consensus

# Eventual Consistency

$(E,i) \models t \text{ klog } a$  :iff  
there is  $j \leq i$  :  $(E@j = (t,a)$  or  
 $((E,j) \models \text{forward}(t',t,\text{id})$  and  
there is  $l < j$  :  $(E,l) \models \text{commit}(t',\text{id})$  and  
 $(E,l) \models t' \text{ klog } a$ ))

A theorem

# Axioms of knowledge

$$(T) := \vDash D_G(\varphi) \rightarrow \varphi$$

(Truth axiom)

$$(4) := \vDash D_G(\varphi) \rightarrow D_G(D_G(\varphi))$$

(Pos. Introspection)

$$(5) := \vDash \neg D_G(\varphi) \rightarrow D_G(\neg D_G(\varphi))$$

(Neg. Introspection)

# Knowledge about Consistency

$$\text{seqCons} := \neg D_{\text{THREADS}}(\neg \text{correct})$$
$$\text{lin} := \neg D_{\text{THREADS}_{\text{u}\{\text{obs}\}}}(\neg \text{correct})$$

# Knowledge about Consistency

$$\models (\text{seqCons} \leftrightarrow D_{\text{Threads}}(\text{seqCons})) \wedge \\ (\neg \text{seqCons} \leftrightarrow D_{\text{Threads}}(\neg \text{seqCons})).$$

# Knowledge about Consistency

$$\models \neg \text{Lin} \leftrightarrow D_{\text{Threads} \uplus \{\text{obs}\}}(\neg \text{Lin})$$

but

$$\text{exists } E: E \models \text{Lin} \wedge \neg D_{\text{Threads} \uplus \{\text{obs}\}}(\text{Lin})$$

# Conclusion

- ▶ Ramification of consistency guarantees are **notoriously oblique**
- ▶ We provide **declarative spec**
- ▶ Uncover non-trivial **relations between** properties
- ▶ Previously unstudied perspective on consistency



Thank you!

# Future Work

- ▶ Observational refinement
- ▶ Exploit Epistemic Logic Theory
- ▶ Other properties in Epistemic Logic
- ▶ Model-check logic to check arbitrary properties

# Eventual Consistency

$$\begin{aligned} \text{alive} &:= \forall t \forall t' \forall id \\ &(\exists (\text{commit}(t, id) \wedge \square \diamond (\exists id' \\ &(\text{commit}(t', id')) \rightarrow \\ &\diamond \text{forward}(t, t', id))) \end{aligned}$$

# The observer's indistinguishability relation

$$\text{obs}(E) = \{(r,c) \in \text{RET} \times \text{CALL} \mid \text{pos}(r,E) < \text{pos}(c,E)\}$$

$$E \sim_{\text{obs}} E' \text{ :iff } \text{obs}(E) \subseteq \text{obs}(E')$$

# The observer

$$E := (t_2, \text{call } \text{ld}()) (t_2, \text{ret } \text{ld}(l)) (t_1, \text{call } \text{st}(l)) (t_1, \text{ret } \text{st}(\text{TRUE}))$$
$$\text{obs}(E) = \{ ( (t_2, \text{ret } \text{ld}(l)) , (t_1, \text{call } \text{st}(l)) ) \}$$

The observer's view is the **order of non-overlapping** method calls

# The observer


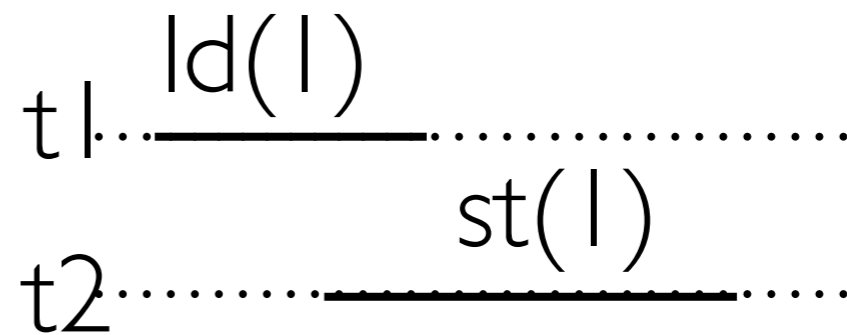

$$E := (t2, \text{call } \text{ld}()) (t1, \text{call } \text{st}(l)) (t2, \text{ret } \text{ld}(l)) (t1, \text{ret } \text{st}(\text{TRUE}))$$
$$\text{obs}(E) = \{ \}$$

The observer's view is the **order of non-overlapping** method calls

# Linearizability

$$E \models \neg D_{\text{THREADS} \cup \{\text{obs}\}}(\neg \text{correct})$$

:iff exists  $E'$  s.t.  $E \sim_{\text{Threads} \cup \{\text{obs}\}} E'$  and  $E' \models \text{correct}$



# Syntax

In the last time-step

proposition

Since

Until

knowledge

$\Phi \ni$

$\varphi ::= p \mid \varphi \wedge \varphi \mid \neg \varphi \mid \Theta \varphi \mid \varphi S \varphi \mid \varphi U \varphi \mid D_G \varphi \mid \forall x(\varphi)$

basic

temporal

quantification



# Semantics

$(E, i) \models \varphi \wedge \psi$  :iff  $(E, i) \models \varphi$  and  $(E, i) \models \psi$

$(E, i) \models \neg \varphi$  :iff not  $(E, i) \models \varphi$

$(E, i) \models \ominus \varphi$  :iff  $i > 0$  and  $(E, i-1) \models \varphi$

$(E, i) \models \varphi S \psi$  :iff  $(E, i) \models \varphi \cup \psi$  :iff

there is  $j \leq i$  s.t.  $(E, j) \models \psi$  and for all  $j < k \leq i$  :  $(E, k) \models \varphi$

$(E, i) \models D_G \varphi$  :iff for all  $(E', i')$ : if  $(E, i) \sim_G (E', i')$  then  
 $(E', i') \models \varphi$

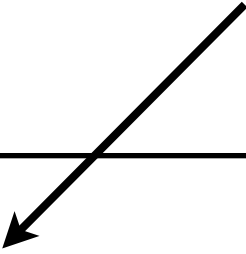
$(E, i) \models \forall x(\varphi)$  :iff for all  $d \in D$  :  $(E, i) \models \varphi[d/x]$

# TSO Consistency

$$E \models \neg D_{\text{THREADS}}(\neg \text{correctTSO})$$

# Read from the buffer

If you read from the store buffer, you need to read the **latest value** stored



$\text{locallyLatest}(t, a, v) := \neg(\exists v' \text{store}(t, a, v')) \wedge \text{store}(t, a, v)$

# TSO Consistency

correctTSO := flush in **Fifo-order**  
FlushOrder  $\wedge$   $\forall t \forall a \forall v$   
 $(\exists (\text{load}(t, a, v) \rightarrow$   
 $(\text{IdBuff}(t, a, v) \vee \text{IdMem}(t, a, v))))$

“so far”

load from  
buffer

load from **memory**

# Basic Predicates

E inspected at  
position  $i \in \mathbb{N}$

Event at  
position  $i$

$(E, i) \models \text{store}(t, a, v) \text{ :iff } E@i = (t, \text{st}(a, v))$   
 $(E, i) \models \text{load}(t, a, v) \text{ :iff } E@i = (t, \text{ld}(a, v))$   
 $(E, i) \models \text{flush}(t, a, v) \text{ :iff } E@i = (m_{\text{sys}}, \text{fl}(t, a, v))$

We need  
positions for  
time

The **memory system**  
flushed a value

Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. Journal of the ACM, 37(3):549-587, 1990.