

Relaxing the Consistency Condition

Local Linearizability
(CONCUR16)

Local Linearizability

main idea

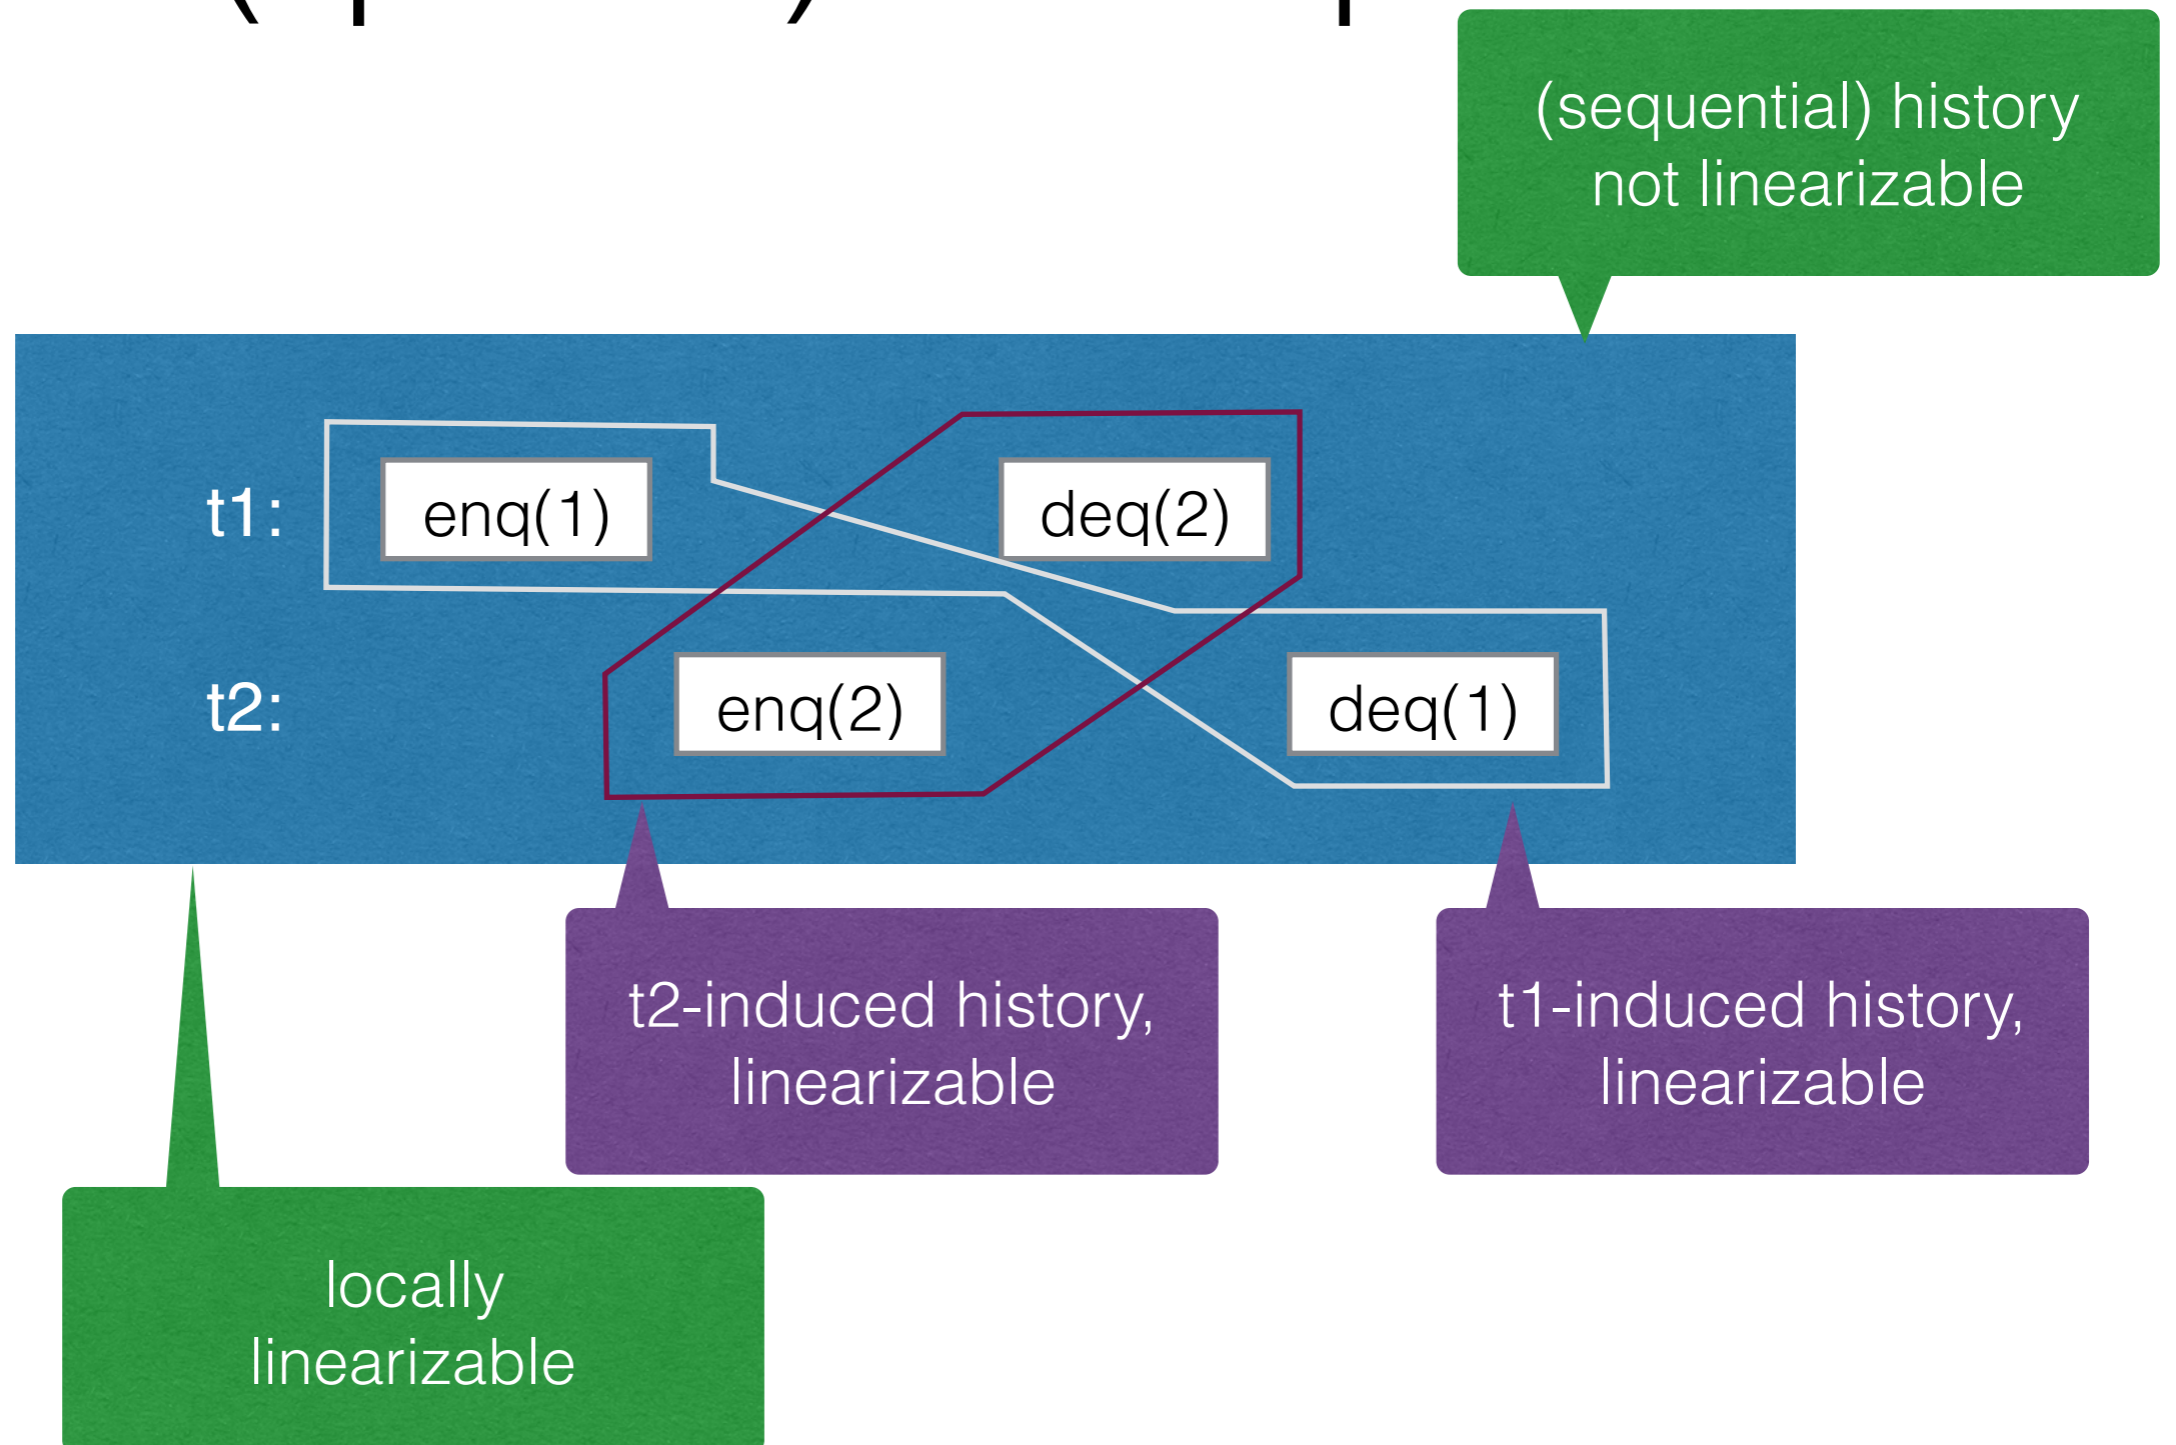
Already present in some shared-memory consistency conditions
(not in our form of choice)

- **Partition** a history into a set of local histories
- Require **linearizability per local history**

no global witness

Local sequential consistency... is also possible

Local Linearizability (queue) example



Local Linearizability (queue) definition

Queue signature $\Sigma = \{\text{enq}(x) \mid x \in V\} \cup \{\text{deq}(x) \mid x \in V\} \cup \{\text{deq}(\text{empty})\}$

For a history \mathbf{h} with a thread T , we put

$$I_T = \{\text{enq}(x)^T \in \mathbf{h} \mid x \in V\}$$

$$O_T = \{\text{deq}(x)^T \in \mathbf{h} \mid \text{enq}(x)^T \in I_T\} \cup \{\text{deq}(\text{empty})\}$$

in-methods of thread T
are
enqueuees performed
by thread T

out-methods of thread T
are dequeuees
(performed by any thread)
corresponding to enqueuees that
are in-methods

\mathbf{h} is locally linearizable iff every thread-induced history
 $\mathbf{h}_T = \mathbf{h} \mid (I_T \cup O_T)$
is linearizable.

Local Linearizability for Container-Type DS

Signature $\Sigma = \text{Ins} \cup \text{Rem} \cup \text{SOB} \cup \text{DOb}$

For a history \mathbf{h} with a thread T , we put

$$I_T = \{m^T \in \mathbf{h} \mid m \in \text{Ins}\}$$

$$O_T = \{m(a) \in \mathbf{h} \cap \text{Rem} \mid i(a)^T \in I_T\} \cup \{m(e) \mid e \in \text{Emp}\} \\ \cup \{m(a) \in \mathbf{h} \cap \text{DOb} \mid i(a)^T \in I_T\}$$

in-methods of thread T
are
inserts performed by
thread T

out-methods of thread T
are removes and data-observations
(performed by any thread)
in-methods

\mathbf{h} is locally linearizable iff every thread-induced history
 $\mathbf{h}_T = \mathbf{h} \mid (I_T \cup O_T)$
is linearizable.

Generalizations of Local Linearizability

Signature Σ

For a history \mathbf{h} with n threads, choose

$\text{In}_{\mathbf{h}}(i)$

$\text{Out}_{\mathbf{h}}(i)$

in-methods of thread i ,
methods that go in \mathbf{h}_i

by increasing the
in-methods,
LL gradually moves to
linearizability

out-methods of thread i ,
dependent methods
on the methods in $\text{In}_{\mathbf{h}}(i)$
(performed by any thread)

\mathbf{h} is locally linearizable iff every thread-induced history


$$\mathbf{h}_i = \mathbf{h} \mid (\text{In}_{\mathbf{h}}(i) \cup \text{Out}_{\mathbf{h}}(i))$$

is linearizable.

Where do we stand?

In general

Local Linearizability



Linearizability




Sequential Consistency

Where do we stand?

For queues (and most container-type data structures)

Local Linearizability



Linearizability



Sequential Consistency

Properties

Local linearizability is compositional

like linearizability
unlike sequential consistency

h (over multiple objects) is locally linearizable
iff
each per-object subhistory of h is locally linearizable

Local linearizability is modular /
“decompositional”

uses decomposition into smaller
histories, by definition

may allow for modular verification