

Linearizability via Order Extension Theorems

Ana Sokolova  UNIVERSITY
of SALZBURG

IRIF, 28.5.2018

- Part I: Concurrent data structures
correctness and performance



structure and power



via semantic
relaxations

- Part II: Order extension results for
verifying linearizability

Concurrent Data Structures

Correctness and Relaxations



Hannes Payer
Google



Tom Henzinger
IST AUSTRIA



Christoph Kirsch
UNIVERSITY
of SALZBURG



Ali Sezgin
UNIVERSITY OF
CAMBRIDGE



Andreas Haas
Google



Michael Lippautz




Andreas Holzer
Google




Helmut Veith
TU
WIEN

Data structures


- Queue FIFO



- Stack LIFO




- Pool unordered

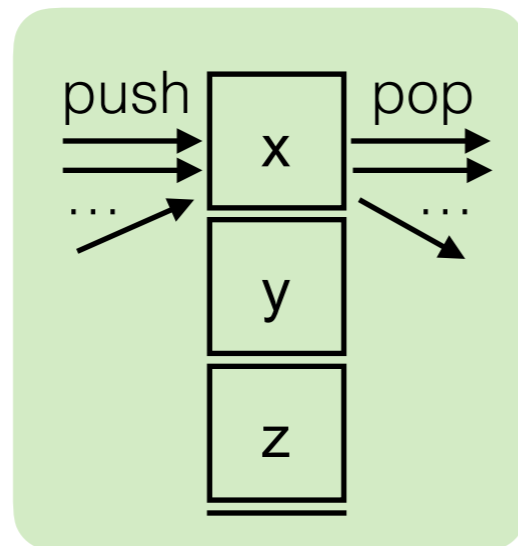


Concurrent data structures

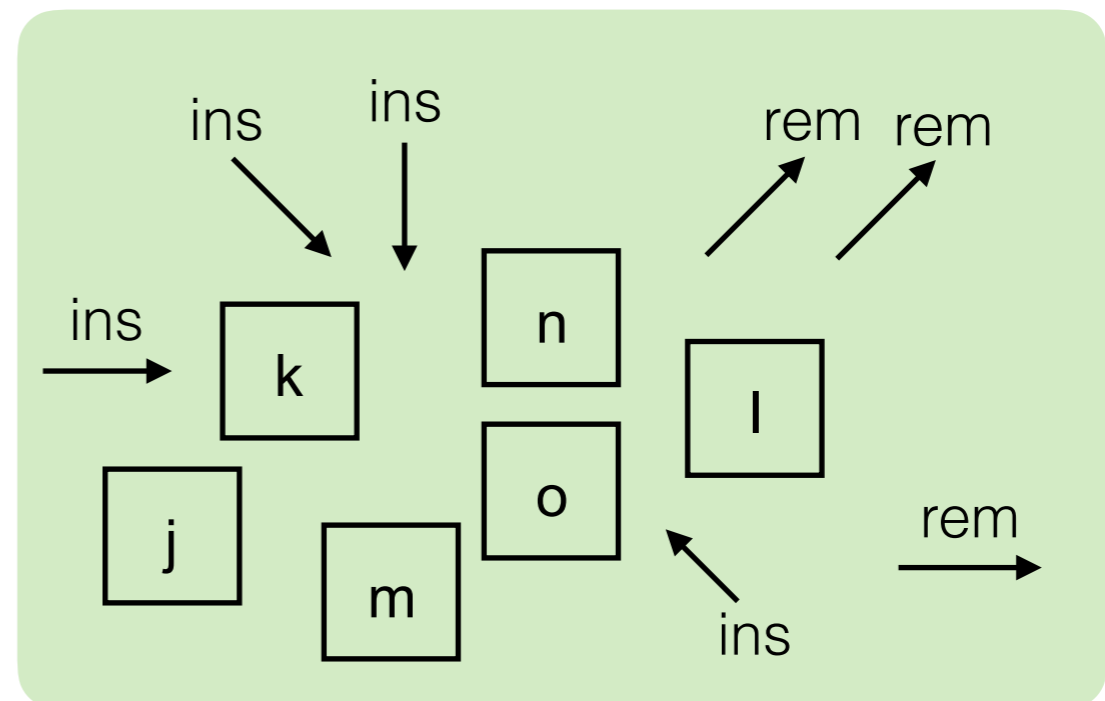
- Queue FIFO



- Stack LIFO



- Pool unordered



Semantics of concurrent data structures

t1: enq(2) deq(1)
t2: enq(1) deq(2)

e.g. queues

- Sequential specification = set of legal sequences

e.g. queue legal sequence
enq(1)enq(2)deq(1)deq(2)

- Consistency condition = e.g. linearizability / sequential consistency

e.g. the concurrent history above is a linearizable queue concurrent history

Consistency conditions

there exists a legal sequence that preserves precedence order

Linearizability [Herlihy, Wing '90]

consistency is about extending partial orders to total orders




t1: enq(2)² — deq(1)³
t2: ¹enq(1) — deq(2)⁴

Sequential Consistency [Lamport'79]

there exists a legal sequence that preserves per-thread precedence (program order)

t1: ¹enq(1) — deq(2)⁴
t2: deq(1)² — enq(2)³

Performance and scalability



Relaxations allow trading

correctness
for
performance

provide the **potential**
for better-performing
implementations

Relaxing the Semantics

Quantitative relaxations
Henzinger, Kirsch, Payer, Sezgin, S. POPL13

- **Sequential specification** = set of legal sequences
- **Consistency condition** = e.g. linearizability / sequential consistency

Local linearizability
Haas, Henzinger, Holzer, ..., S, Veith CONCUR16

Relaxing the sequential specification

Quantitative
relaxations
(POPL13)

Goal

Stack - incorrect behavior

```
push(a)push(b)push(c)pop(a)pop(b)
```

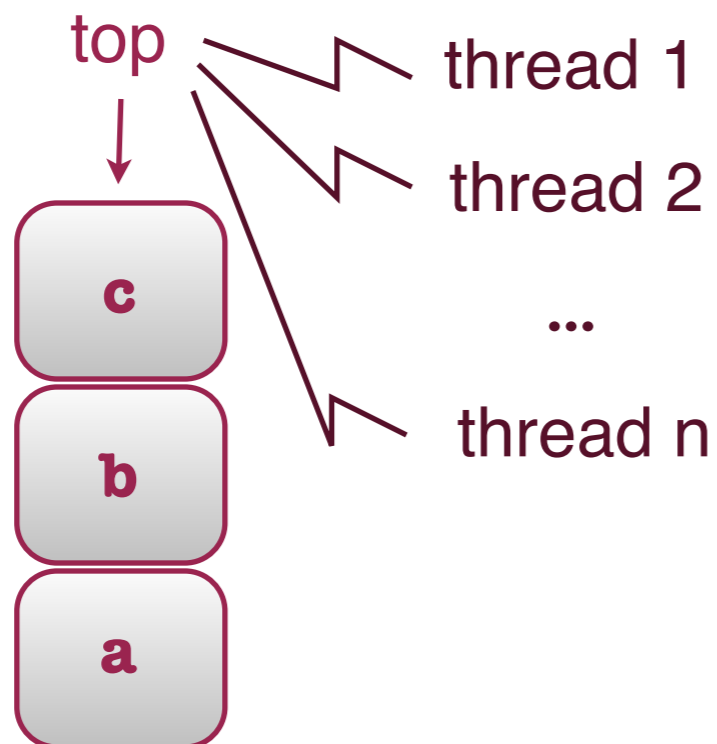
- trade correctness for performance
- in a controlled way with quantitative bounds

correct in a relaxed stack
... 2-relaxed? 3-relaxed?


measure the
error from correct
behaviour

How can relaxing help?

Stack



k-Relaxed stack



What we have

- Framework

for semantic relaxations

- Generic examples

out-of-order /
stuttering

- Concrete relaxation examples

stacks, queues,
priority queues,.. /
CAS, shared counter

- Efficient concurrent implementations

of relaxation instances

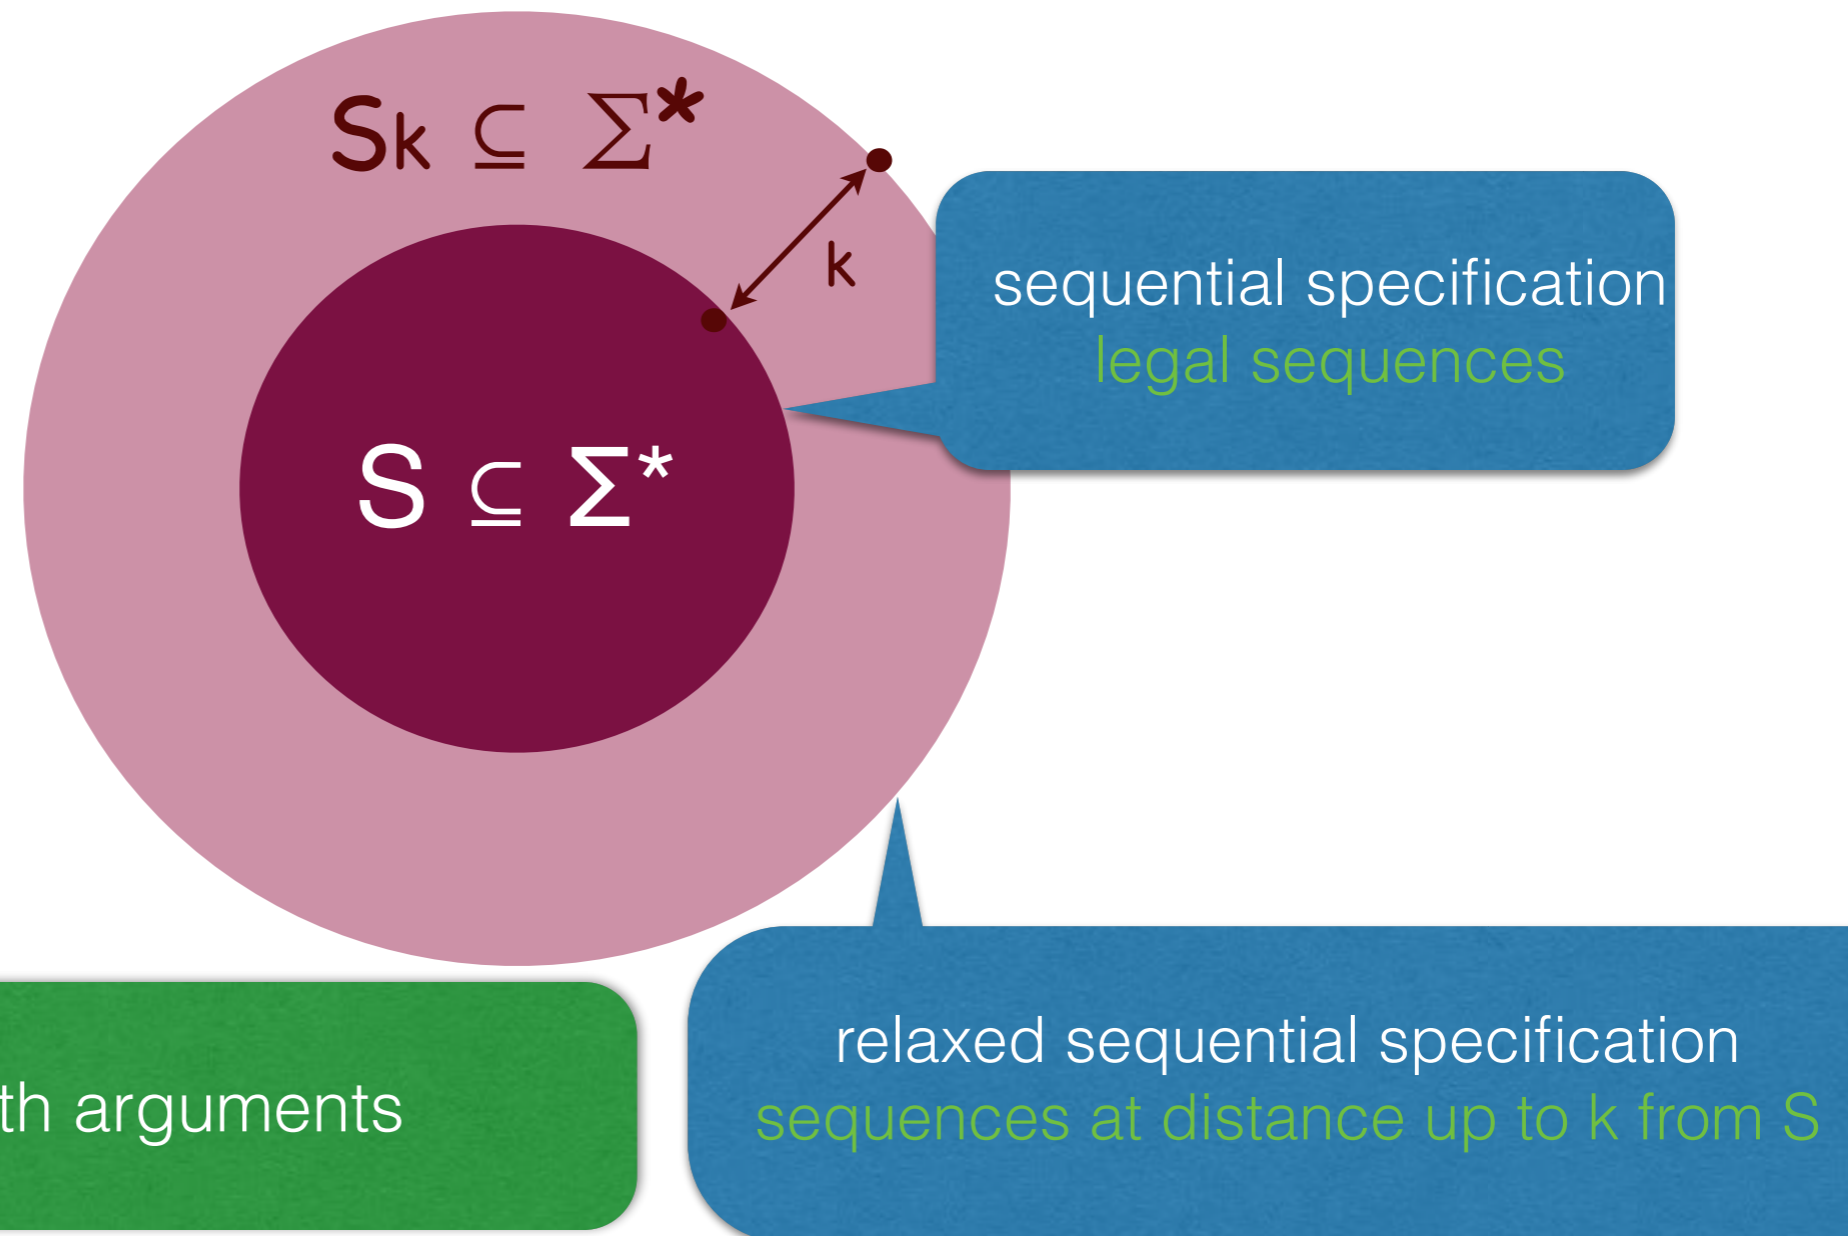
The big picture

$$S \subseteq \Sigma^*$$

sequential specification
legal sequences

Σ - methods with arguments

The big picture



Σ - methods with arguments

sequential specification
legal sequences

relaxed sequential specification
sequences at distance up to k from S

Relaxing the Consistency Condition

Local Linearizability
(CONCUR16)

Local Linearizability

main idea

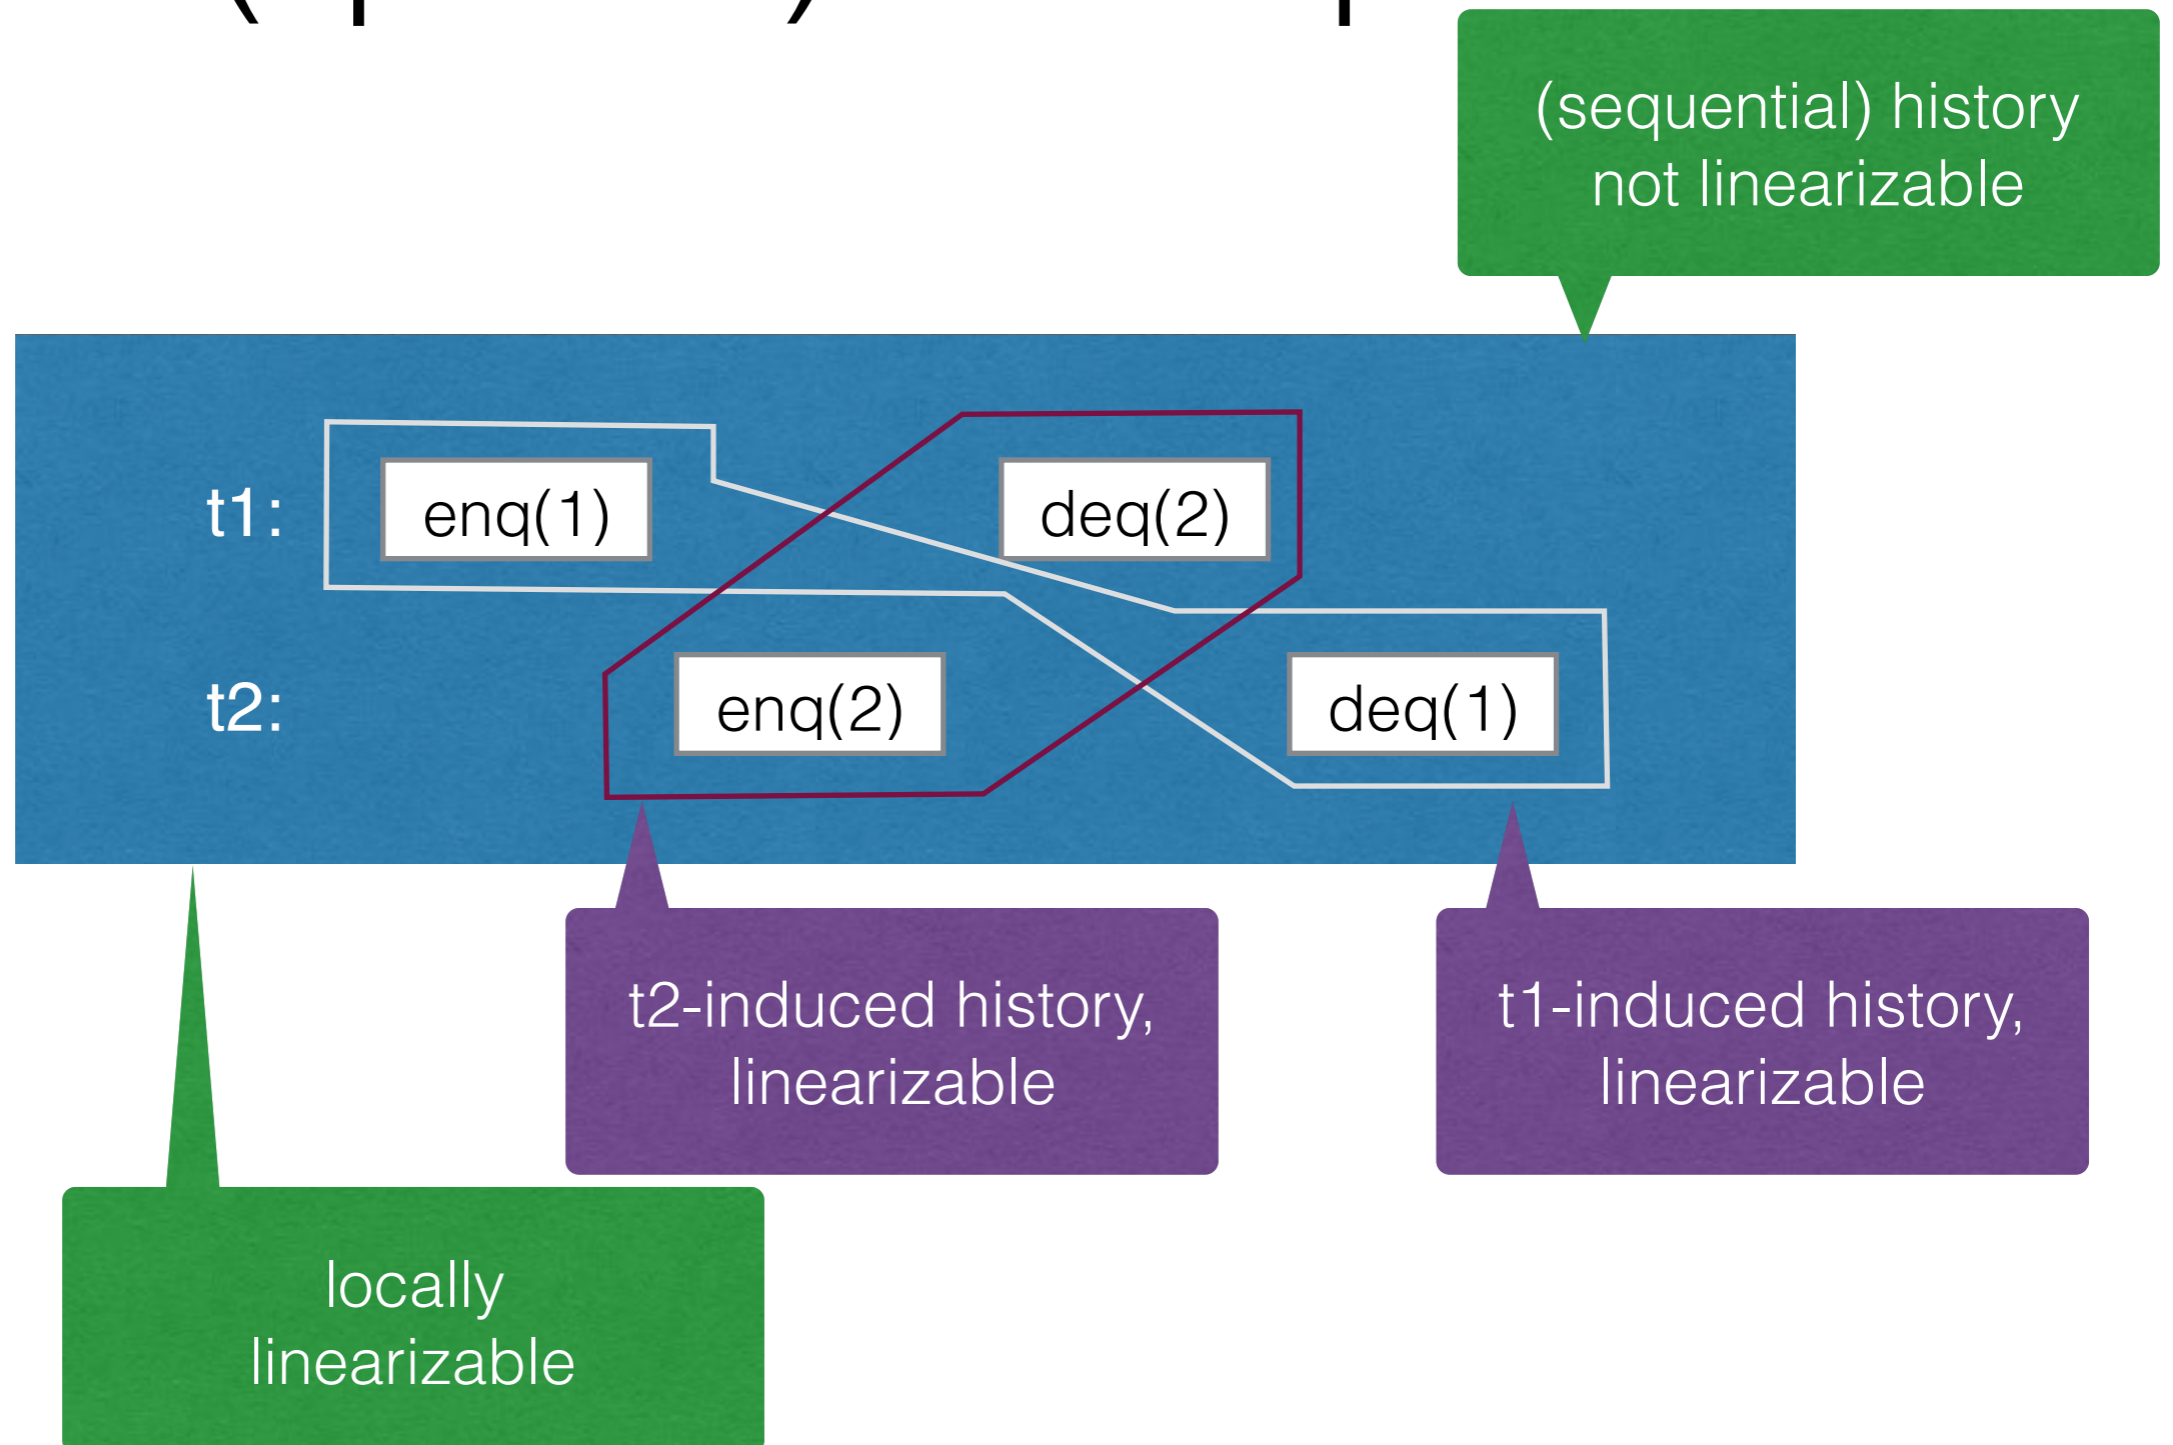
Already present in some shared-memory consistency conditions
(not in our form of choice)

- **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.

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


Lead to scalable implementations

e.g. k-FIFO, k-Stack



k-out-of-order queue


locally linearizable distributed implementation



local inserts / global removes

LLD Φ
LL+D Φ

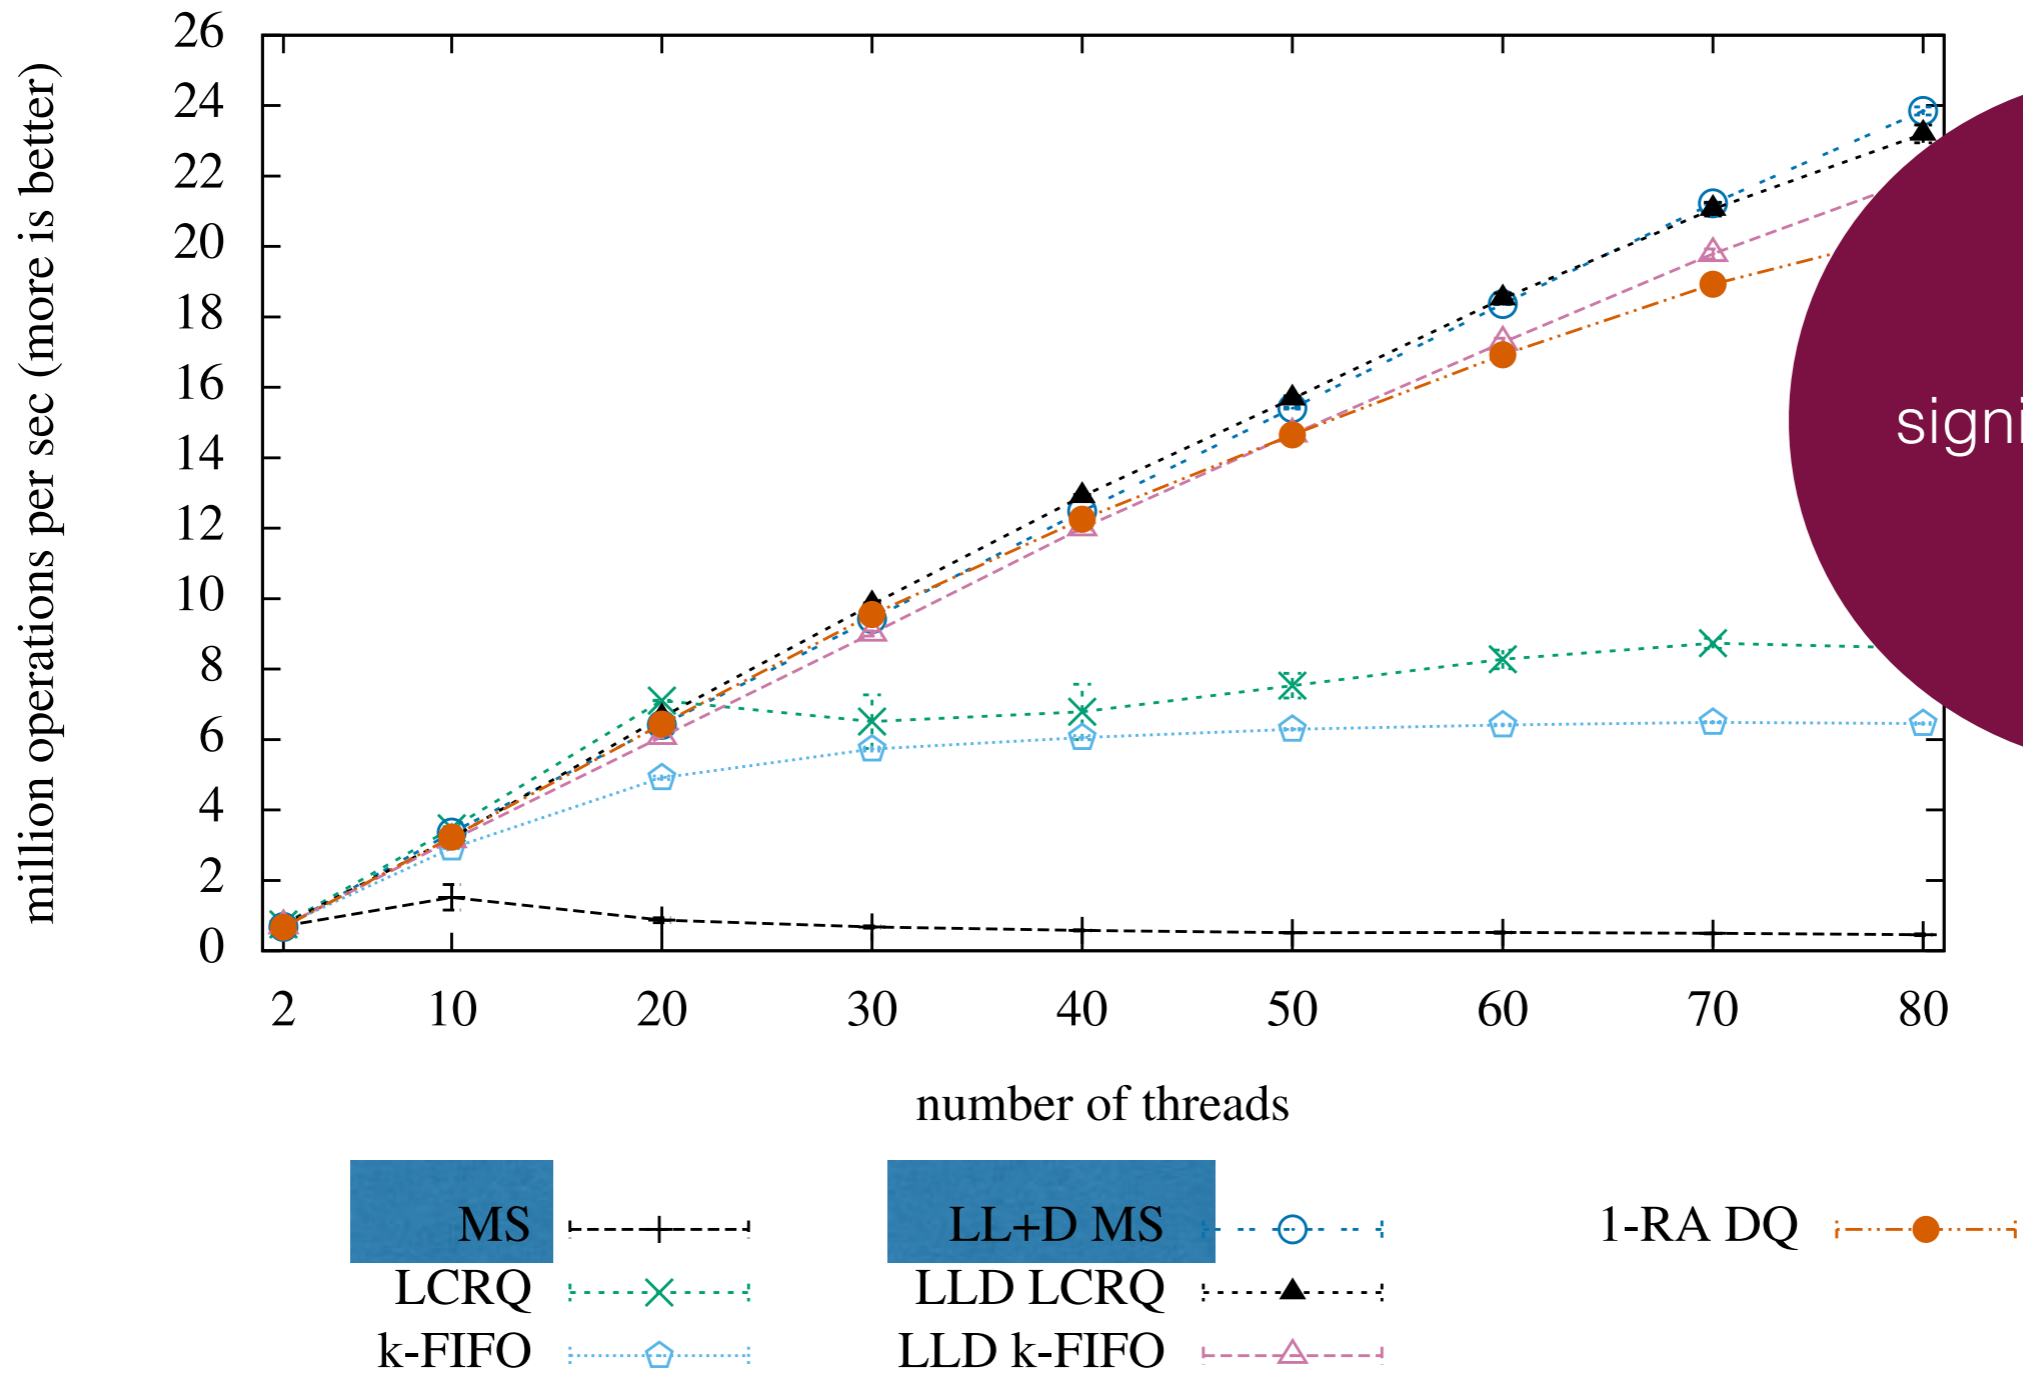
Performance



LL+D MS queue performs significantly better than MS queue

(a) Queues, LL queues, and “queue-like” pools


Performance



LLD ϕ
performs
significantly better
than
 ϕ

(a) Queues, LL queues, and “queue-like” pools

Performance



LL+D MS queue performs better than the best known pools

(a) Queues, LL queues, and “queue-like” pools

Linearizability via Order Extension Theorems

joint work with



Harald Woracek



foundational results
for
verifying linearizability

Inspiration

As well as
Reducing Linearizability to
State Reachability
[Bouajjani, Emmi, Enea, Hamza]
ICALP15 + ...

Queue sequential specification (axiomatic)

s is a legal queue sequence
iff

1. **s** is a legal pool sequence, and

2. $\text{enq}(x) <_{\mathbf{s}} \text{enq}(y) \wedge \text{deq}(y) \in \mathbf{s} \Rightarrow \text{deq}(x) \in \mathbf{s} \wedge \text{deq}(x) <_{\mathbf{s}} \text{deq}(y)$

Queue linearizability (axiomatic)

Henzinger, Sezgin, Vafeiadis CONCUR13

h is queue linearizable
iff

1. **h** is pool linearizable, and

2. $\text{enq}(x) <_{\mathbf{h}} \text{enq}(y) \wedge \text{deq}(y) \in \mathbf{h} \Rightarrow \text{deq}(x) \in \mathbf{h} \wedge \text{deq}(y) \not<_{\mathbf{h}} \text{deq}(x)$

precedence order

Concurrent Queues

Data independence => verifying executions where each value is enqueued at most once is sound

Reduction to **assertion checking** = exclusion of "bad patterns"

Value v dequeued without being enqueued

$\text{deq} \Rightarrow v$




Value v dequeued before being enqueued

$\text{deq} \Rightarrow v$ $\text{enq}(v)$




Value v dequeued twice

$\text{deq} \Rightarrow v$ $\text{deq} \Rightarrow v$





Value v_1 and v_2 dequeued in the wrong order

$\text{enq}(v_1)$ $\text{enq}(v_2)$ $\text{deq} \Rightarrow v_2$ $\text{deq} \Rightarrow v_1$



Dequeue wrongfully returns empty

$\text{deq} \Rightarrow \text{empty}$



Linearizability verification

Data structure

- signature Σ - set of method calls including data values
- sequential specification $S \subseteq \Sigma^*$, prefix closed

identify sequences with total orders

Sequential specification via violations

Extract a set of violations V , relations on Σ , such that $\mathbf{s} \in S$ iff \mathbf{s} has no violations

it is easy to find a large CV,
but difficult to find a small representative

$$\mathcal{P}(\mathbf{s}) \cap V = \emptyset$$

Linearizability verification

Find a set of violations CV such that: every interval order with no CV violations extends to a total order with no V violations.

we build
CV iteratively
from V

legal sequence

concurrent history

Pool without empty removals

Pool sequential specification (axiomatic)

s is a legal pool (without empty removals) sequence

iff
1. $\text{rem}(x) \in \mathbf{s} \Rightarrow \text{ins}(x) \in \mathbf{s} \wedge \text{ins}(x) <_{\mathbf{s}} \text{rem}(x)$

V violations
 $\text{rem}(x) <_{\mathbf{s}} \text{ins}(x)$

Pool linearizability (axiomatic)

h is pool (without empty removals) linearizable

iff
1. $\text{rem}(x) \in \mathbf{h} \Rightarrow \text{ins}(x) \in \mathbf{h} \wedge \text{rem}(x) \not<_{\mathbf{h}} \text{ins}(x)$

CV violations
= V violations

Queue without empty removals

Queue sequential specification (axiomatic)

s is a legal queue (without empty removals) sequence

iff

$$1. \text{deq}(x) \in \mathbf{s} \Rightarrow \text{enq}(x) \in \mathbf{s} \wedge \text{enq}(x) <_{\mathbf{s}} \text{deq}(x)$$

$$2. \text{enq}(x) <_{\mathbf{s}} \text{enq}(y) \wedge \text{deq}(y) \in \mathbf{s} \Rightarrow \text{deq}(x) \in \mathbf{s} \wedge \text{deq}(x) <_{\mathbf{s}} \text{deq}(y)$$

V violations
 $\text{deq}(x) <_{\mathbf{s}} \text{enq}(x)$
and
 $\text{enq}(x) <_{\mathbf{s}} \text{enq}(y) \wedge$
 $\text{deq}(y) <_{\mathbf{s}} \text{deq}(x)$

Queue linearizability (axiomatic)

h is queue (without empty removals) linearizable

iff

$$1. \text{rem}(x) \in \mathbf{h} \Rightarrow \text{ins}(x) \in \mathbf{h} \wedge \text{rem}(x) \not<_{\mathbf{h}} \text{ins}(x)$$

$$2. \text{enq}(x) <_{\mathbf{h}} \text{enq}(y) \wedge \text{deq}(y) \in \mathbf{h} \Rightarrow \text{deq}(x) \in \mathbf{h} \wedge \text{deq}(y) \not<_{\mathbf{h}} \text{deq}(x)$$

CV violations
= V violations

Pool

infinite inductive violations

Pool sequential specification (axiomatic)

s is a legal pool (with empty removals) sequence

iff

$$1. \text{rem}(x) \in \mathbf{s} \Rightarrow \text{ins}(x) \in \mathbf{s} \wedge \text{ins}(x) <_{\mathbf{s}} \text{rem}(x)$$

$$2. \text{rem}(\perp) <_{\mathbf{s}} \text{rem}(x) \Rightarrow \text{rem}(\perp) <_{\mathbf{s}} \text{ins}(x) \wedge \text{ins}(x) <_{\mathbf{s}} \text{rem}(\perp) \Rightarrow \text{rem}(x) <_{\mathbf{s}} \text{rem}(\perp)$$

\forall violations
 $\text{rem}(x) <_{\mathbf{s}} \text{ins}(x)$
and
 $\text{ins}(x) <_{\mathbf{s}} \text{rem}(\perp) <_{\mathbf{s}} \text{rem}(x)$

Pool linearizability (axiomatic)

h is pool (with empty removals) linearizable

iff

$$1. \text{rem}(x) \in \mathbf{h} \Rightarrow \text{ins}(x) \in \mathbf{h} \wedge \text{rem}(x) \not<_{\mathbf{h}} \text{ins}(x)$$

2.

infinitely many CV violations

$$\text{ins}(x_1) <_{\mathbf{h}} \text{rem}(\perp) \wedge \text{ins}(x_2) <_{\mathbf{h}} \text{rem}(x_1) \wedge \dots \wedge \text{ins}(x_{n+1}) <_{\mathbf{h}} \text{rem}(x_n) \wedge \text{rem}(\perp) <_{\mathbf{h}} \text{rem}(x_{n+1})$$

It works for

- Pool without empty removals
- Queue without empty removals
- Priority queue without empty removals
- Pool
- Queue
- Priority queue

infinite
inductive
violations

But not yet for Stack:
infinite CV violations
without clear
inductive structure

Exploring the space of
data structures
as well as new ideas
for problematic cases

How does it work?

The basics

$$\text{PO}[\mathcal{V}] = \{R \in \text{PO} \mid \mathcal{P}(R) \cap \mathcal{V} = \emptyset\}$$

$$\text{IO}[\mathcal{V}] = \{R \in \text{IO} \mid \mathcal{P}(R) \cap \mathcal{V} = \emptyset\}$$

$$\text{TO}[\mathcal{V}] = \{R \in \text{TO} \mid \mathcal{P}(R) \cap \mathcal{V} = \emptyset\}$$

partial orders

interval orders

total orders

$$\forall (a, b), (c, d) \in R. (a, d) \in R \vee (c, b) \in R$$

The problem

Given a set of violations \mathcal{V} , find a “small” set of violations \mathcal{V}' such that

$$\forall R \in \text{IO}[\mathcal{V}']. \exists \bar{R} \in \text{TO}[\mathcal{V}]. \bar{R} \supseteq R$$

Theorem (singleton violations)

Let \mathcal{V} consist only of singletons, and let $V = \bigcup \mathcal{V}$.

If V is transitive and not a cycle, then the problem is solved with $\mathcal{V}' = \mathcal{V}$.

this solves the case of
pool without empty removals

The closures

$$\text{Clos}_O(\mathcal{V}) = \bigcap_{S \in O[\mathcal{V}]} \mathcal{P}(S)^c$$

O-closure of a set
of violations

monotone, extensive, idempotent

Proposition

$$\forall R \in IO[\mathcal{V}']. \exists \bar{R} \in TO[\mathcal{V}]. \bar{R} \supseteq R$$

iff

$$\text{Clos}_{TO}(\mathcal{V}) = \text{Clos}_{IO}(\mathcal{V}')$$

How does it work ?

Theorem

Let \mathcal{V} consist only of finite sets and assume

(1) ★

(2) $\forall N, M \in \mathcal{V}. \forall (a_1, a_2) \in N. |\{(b_1, b_2) \in M \mid a_2 = b_1\}| \leq 1$

then the problem is solved

we provide an algorithm that produces a set of violations such that ★ holds

if we are lucky, (2) holds too

if we manage to construct such a set of violations, we are done

The algorithm

Take two violations $N_1, N_2 \in \mathcal{V}$ and an element $x \in X$ and produce a new violation

$$\begin{aligned} & \{(a, b) \mid (a, x) \in N_1, (x, b) \in N_2\} \\ & \cup \{(a, b) \in N_1 \mid b \neq x\} \\ & \cup \{(a, b) \in N_2 \mid a \neq x\} \end{aligned}$$

Take two violations $N_1, N_2 \in \mathcal{V}$ and a pair $(x, y) \in X \times X$ and produce a new violation

$$\begin{aligned} & \{(a, y) \mid (a, x) \in N_2\} \\ & \cup \{(x, b) \mid (y, b) \in N_2\} \\ & \cup \{(a, b) \in N_2 \mid b \neq x \wedge a \neq y\} \\ & \cup N_1 \setminus \{(x, y)\} \end{aligned}$$

until no new
violations are
produced

It works for

- Pool without empty removals
- Queue without empty removals
- Priority queue without empty removals
- Pool
- Queue
- Priority queue

Thank You !

But not yet for Stack:
infinite CV violations
without clear
inductive structure

Exploring the space of
data structures
as well as new ideas
for problematic cases