

# Local Linearizability

Ana Sokolova  UNIVERSITY  
of SALZBURG

joint work with:

Andreas Haas  UNIVERSITY  
of SALZBURG

Andreas Holzer  UNIVERSITY OF  
TORONTO

Michael Lippautz  UNIVERSITY  
of SALZBURG

Ali Sezgin  UNIVERSITY OF  
CAMBRIDGE

Tom Henzinger  IST AUSTRIA

Christoph Kirsch  UNIVERSITY  
of SALZBURG

Hannes Payer  Google

Helmut Veith  TU  
WIEN

# Concurrent Data Structures

## Correctness and Performance

# Semantics of concurrent data structures

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

e.g. pools, queues, stacks

- 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 sequential witness that preserves precedence

Linearizability [Herlihy, Wing '90]

t1: enq(2)<sup>2</sup> deq(1)<sup>3</sup>  
t2: <sup>1</sup>enq(1) deq(2)<sup>4</sup>




Sequential Consistency [Lamport '79]

there exists a sequential witness that preserves per-thread precedence (program order)

t1: <sup>1</sup>enq(1) deq(2)<sup>4</sup>  
t2: deq(1)<sup>2</sup> enq(2)<sup>3</sup>

# Performance and scalability



# Relaxations allow trading

correctness  
for  
performance

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

# Relaxing the Semantics

not  
“sequentially  
correct”

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

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

for queues only  
(feel free to ask for more)

Local linearizability  
in this talk

too weak

# Local Linearizability

## main idea

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

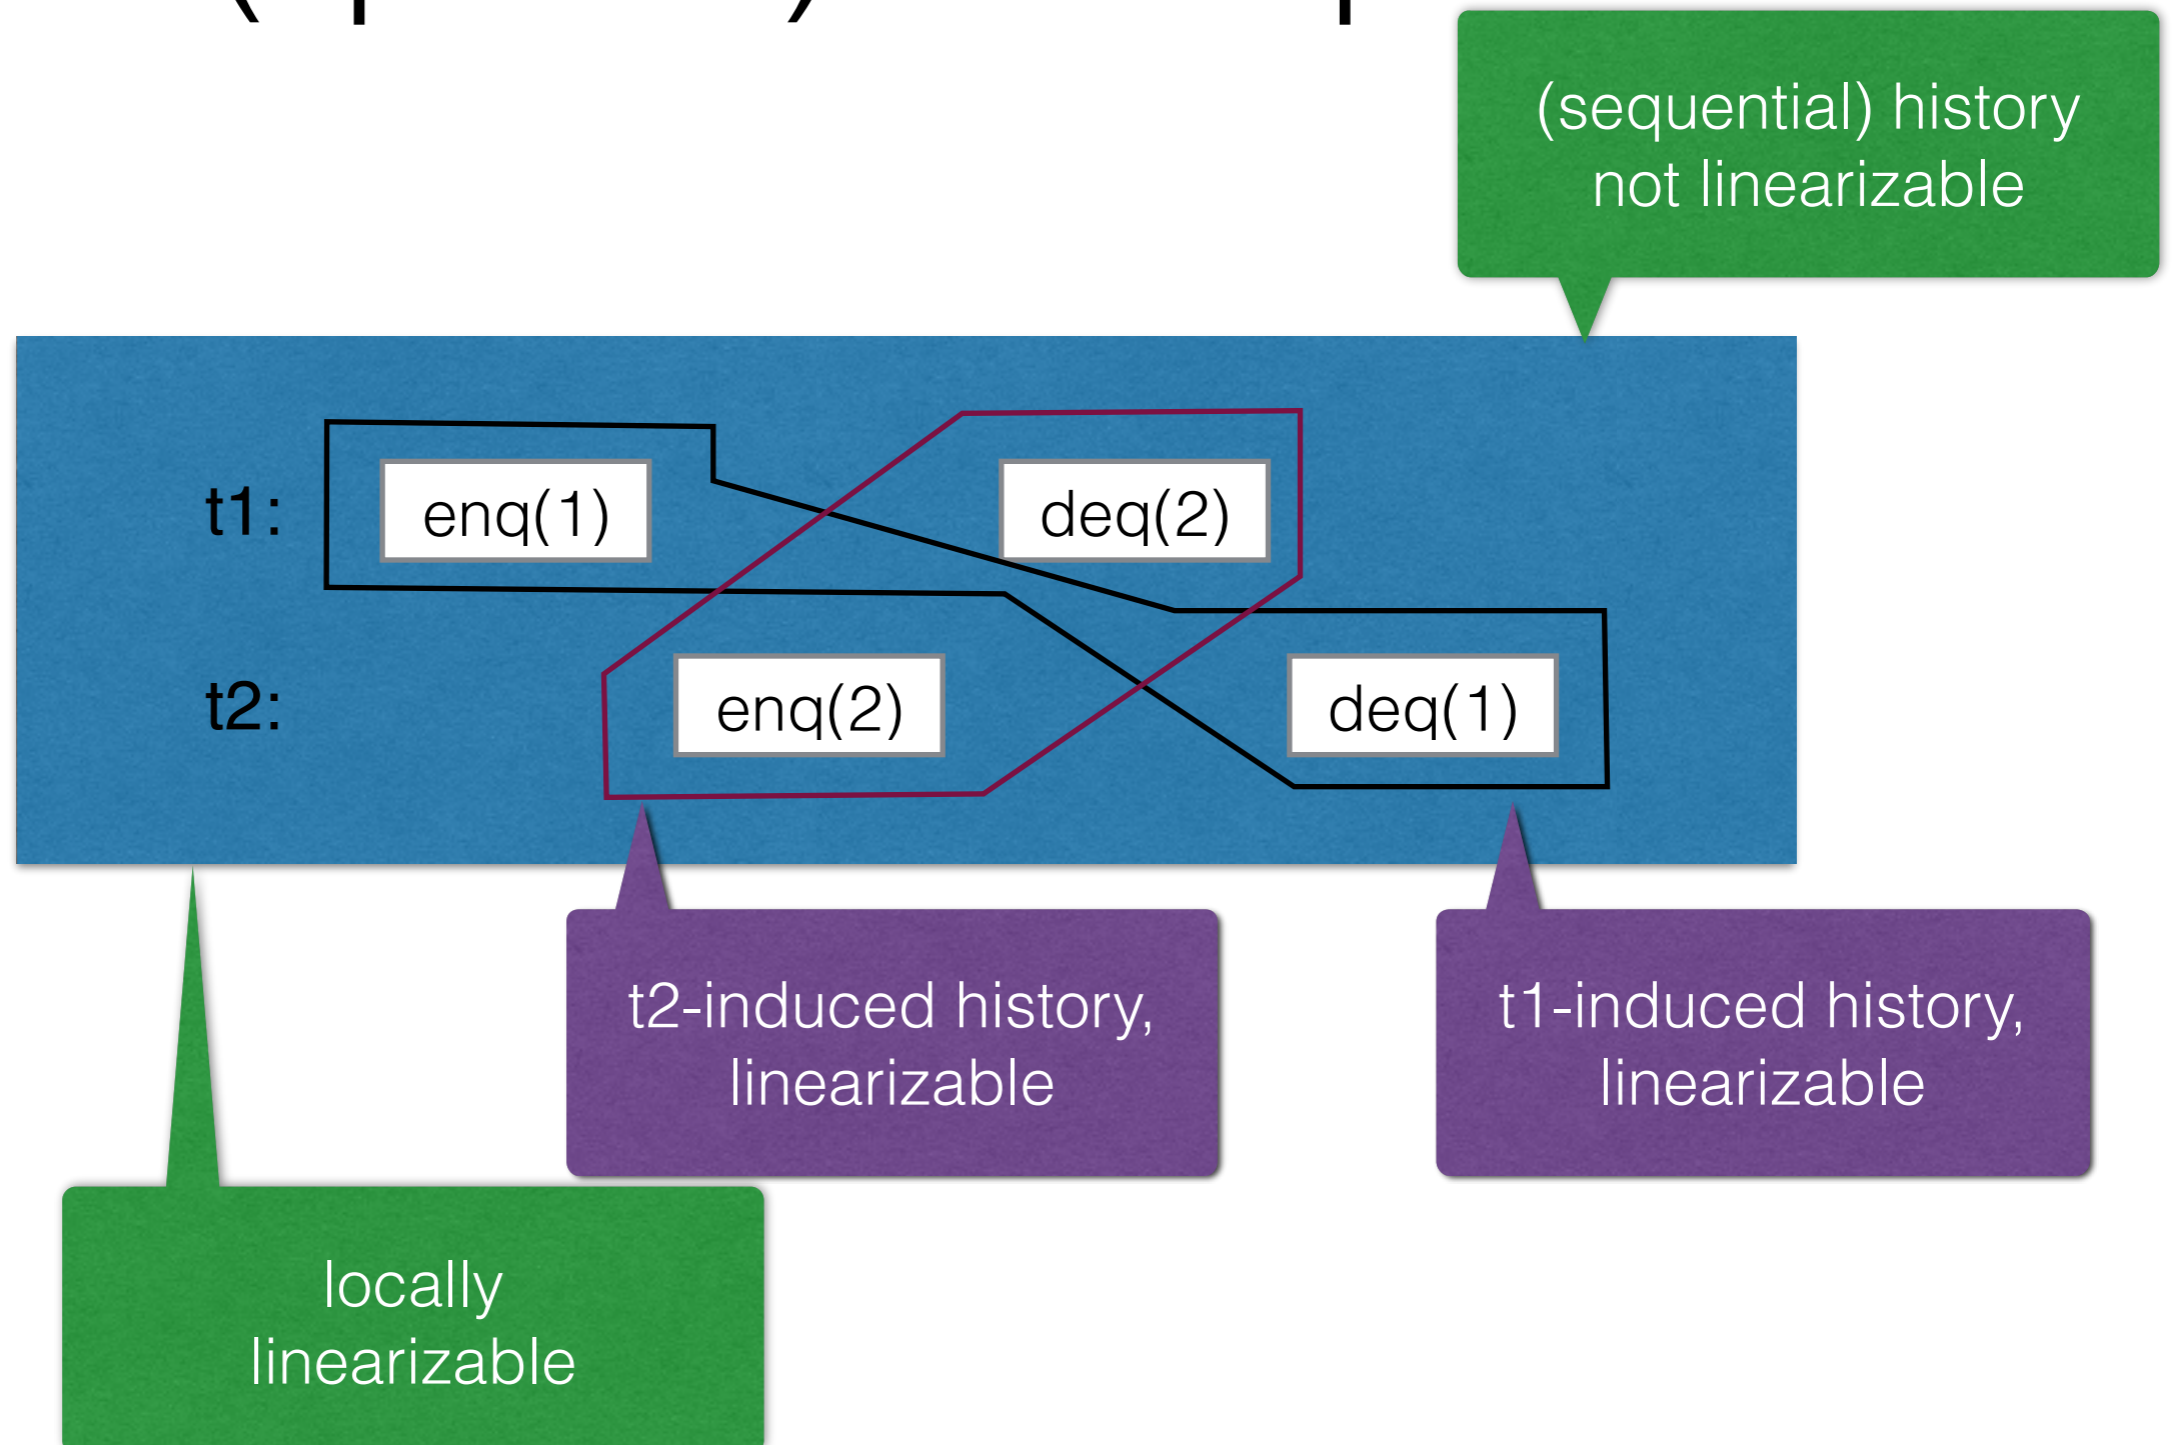
- **Partition** a history into a set of local histories
- Require **linearizability per local history**

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  $n$  threads, we put

$$\text{In}_{\mathbf{h}}(i) = \{\text{enq}(x)^i \in \mathbf{h} \mid x \in V\}$$

in-methods of thread  $i$   
enqueues performed  
by thread  $i$

$$\text{Out}_{\mathbf{h}}(i) = \{\text{deq}(x)^j \in \mathbf{h} \mid \text{enq}(x)^i \in \text{In}_{\mathbf{h}}(i)\} \cup \{\text{deq}(\text{empty})\}$$

out-methods of thread  $i$   
dequeues  
(performed by any thread)  
corresponding to enqueues that  
are in-methods

$\mathbf{h}$  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.

# Generalizations of Local Linearizability

Signature  $\Sigma$

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 all pool-like data structures)

Local Linearizability



Linearizability





Sequential Consistency

# Where do we stand?

C: For queues

Local Linearizability  
& Pool-seq.cons.

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

# Verification (queue)

## 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)

**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



# Verification (queue)

## 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 local linearizability (axiomatic)

**h** is queue locally linearizable

iff

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

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

thread-local  
precedence order

# Generic Implementations

Your favorite linearizable data structure implementation

$\Phi$

turns into a locally linearizable implementation by:

LLD  $\Phi$   
(locally linearizable)

$t_1$   $t_2$  ...  $t_n$

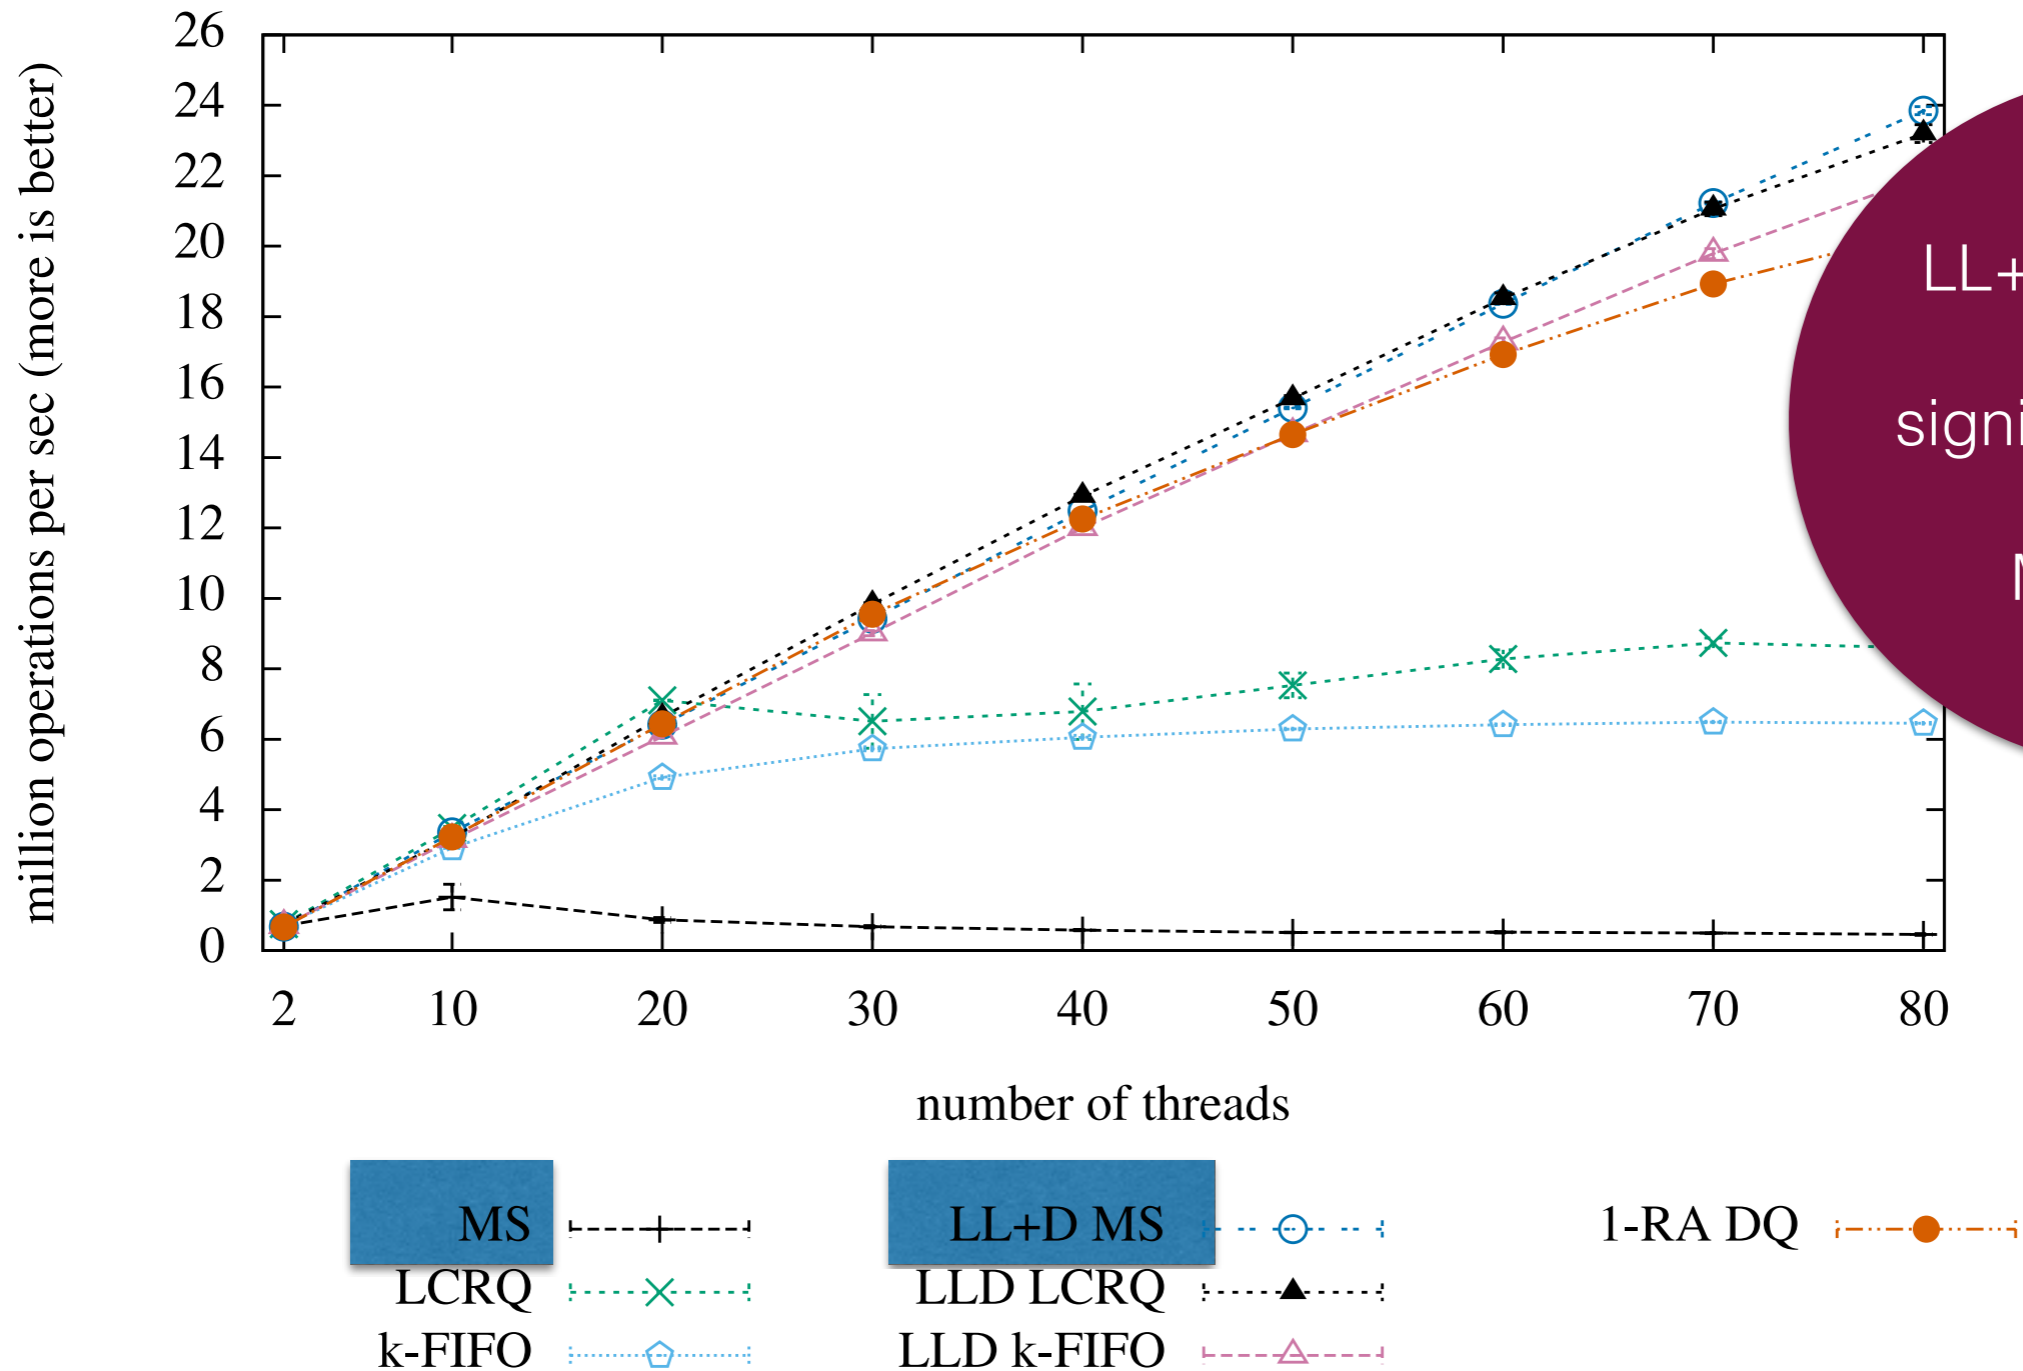
segment of dynamic size (n)

$\Phi$   $\Phi$   $\Phi$

LL+D  $\Phi$   
(also pool linearizable)

local inserts / global (randomly distributed) removes

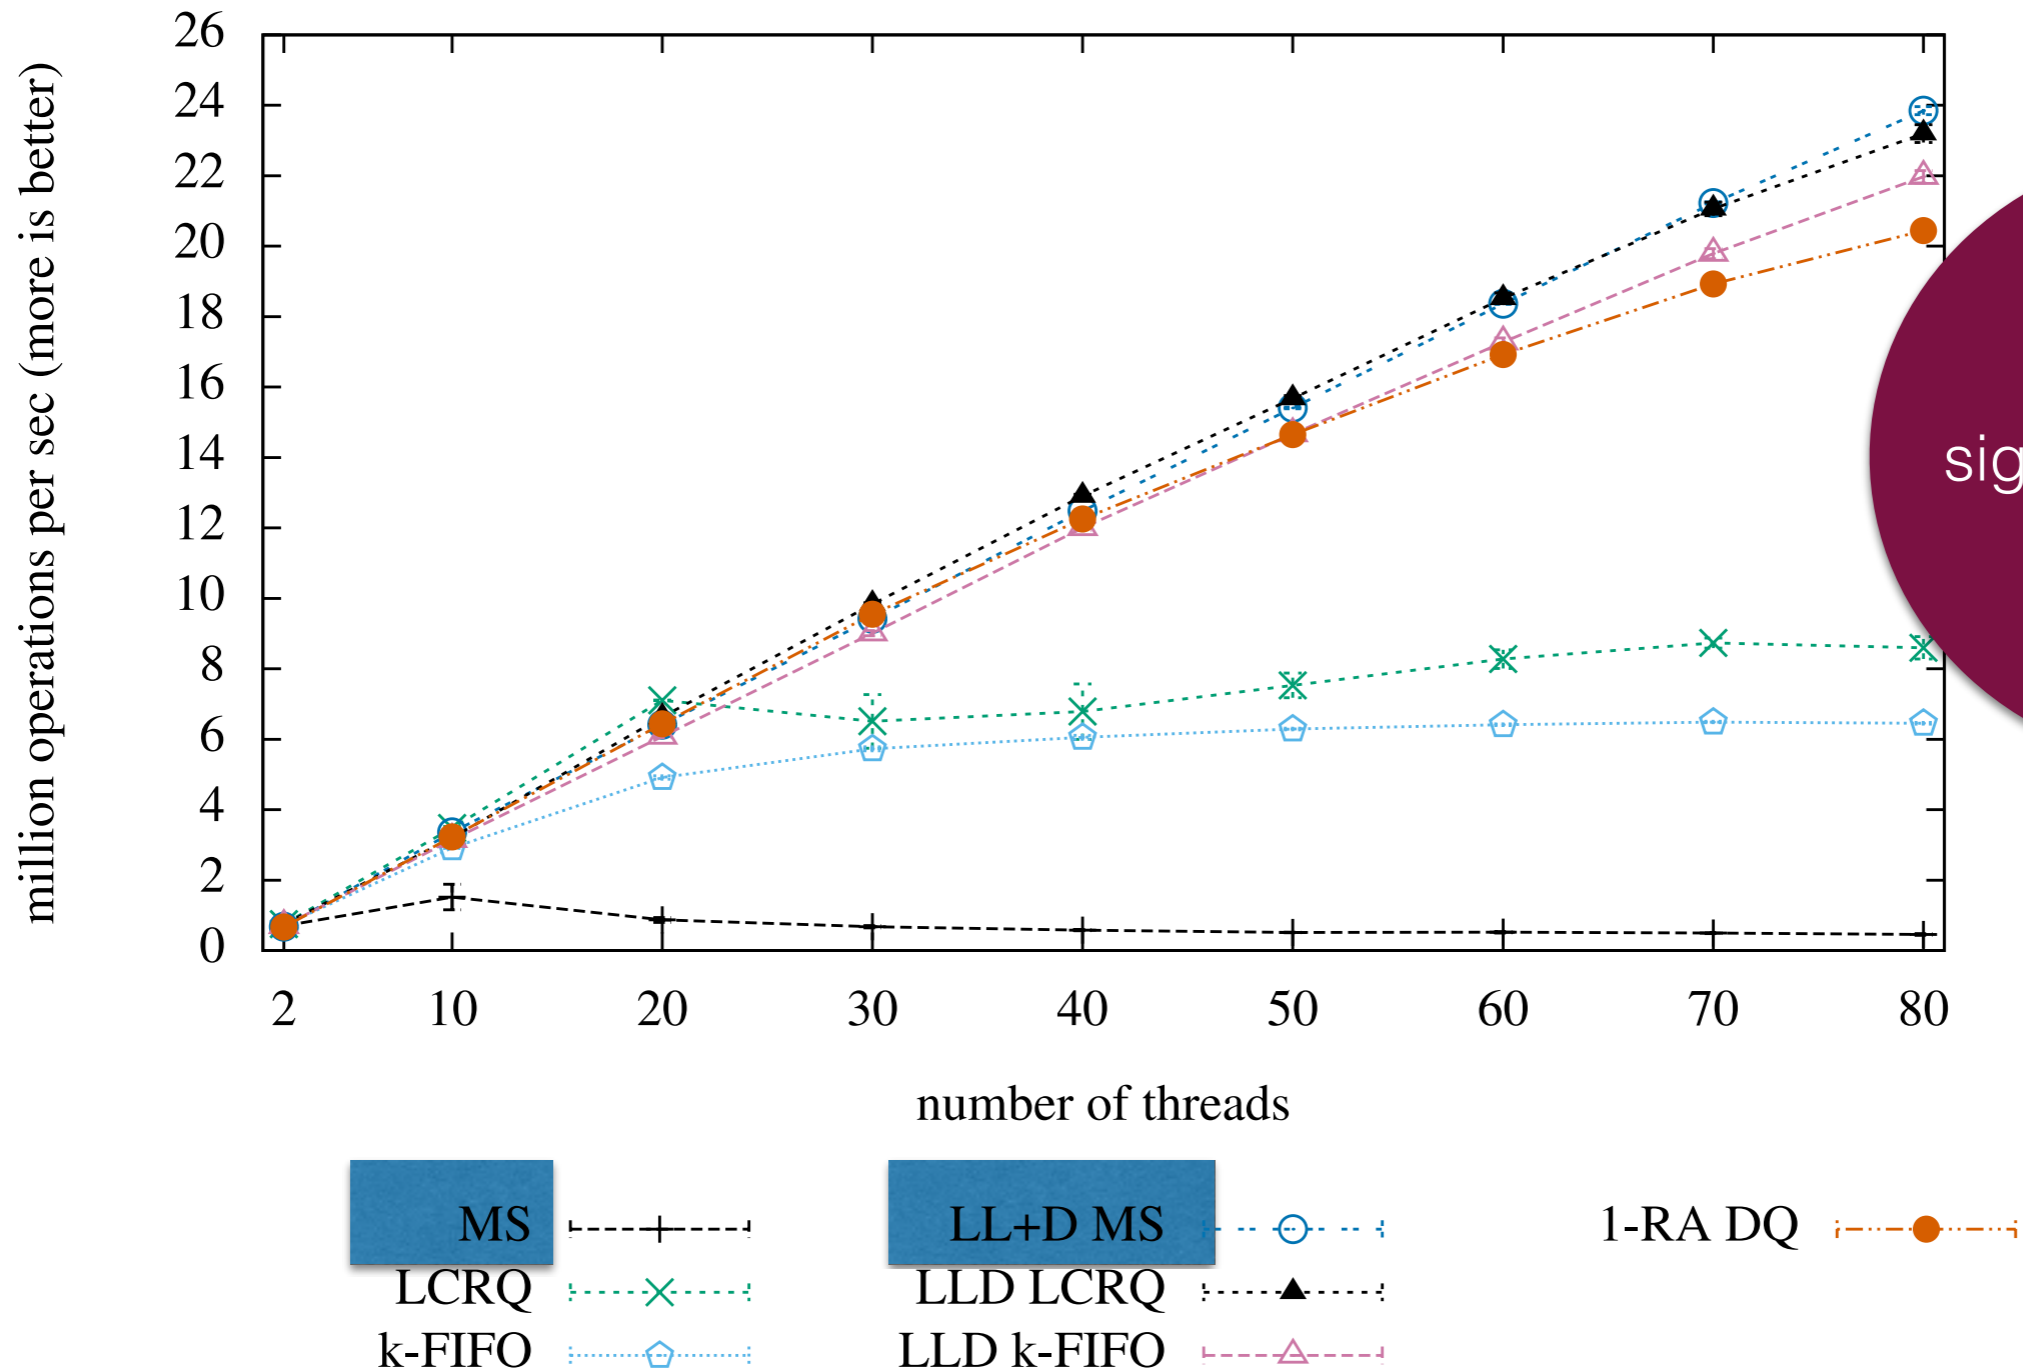
# Performance



LL+D MS queue performs significantly better than MS queue

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

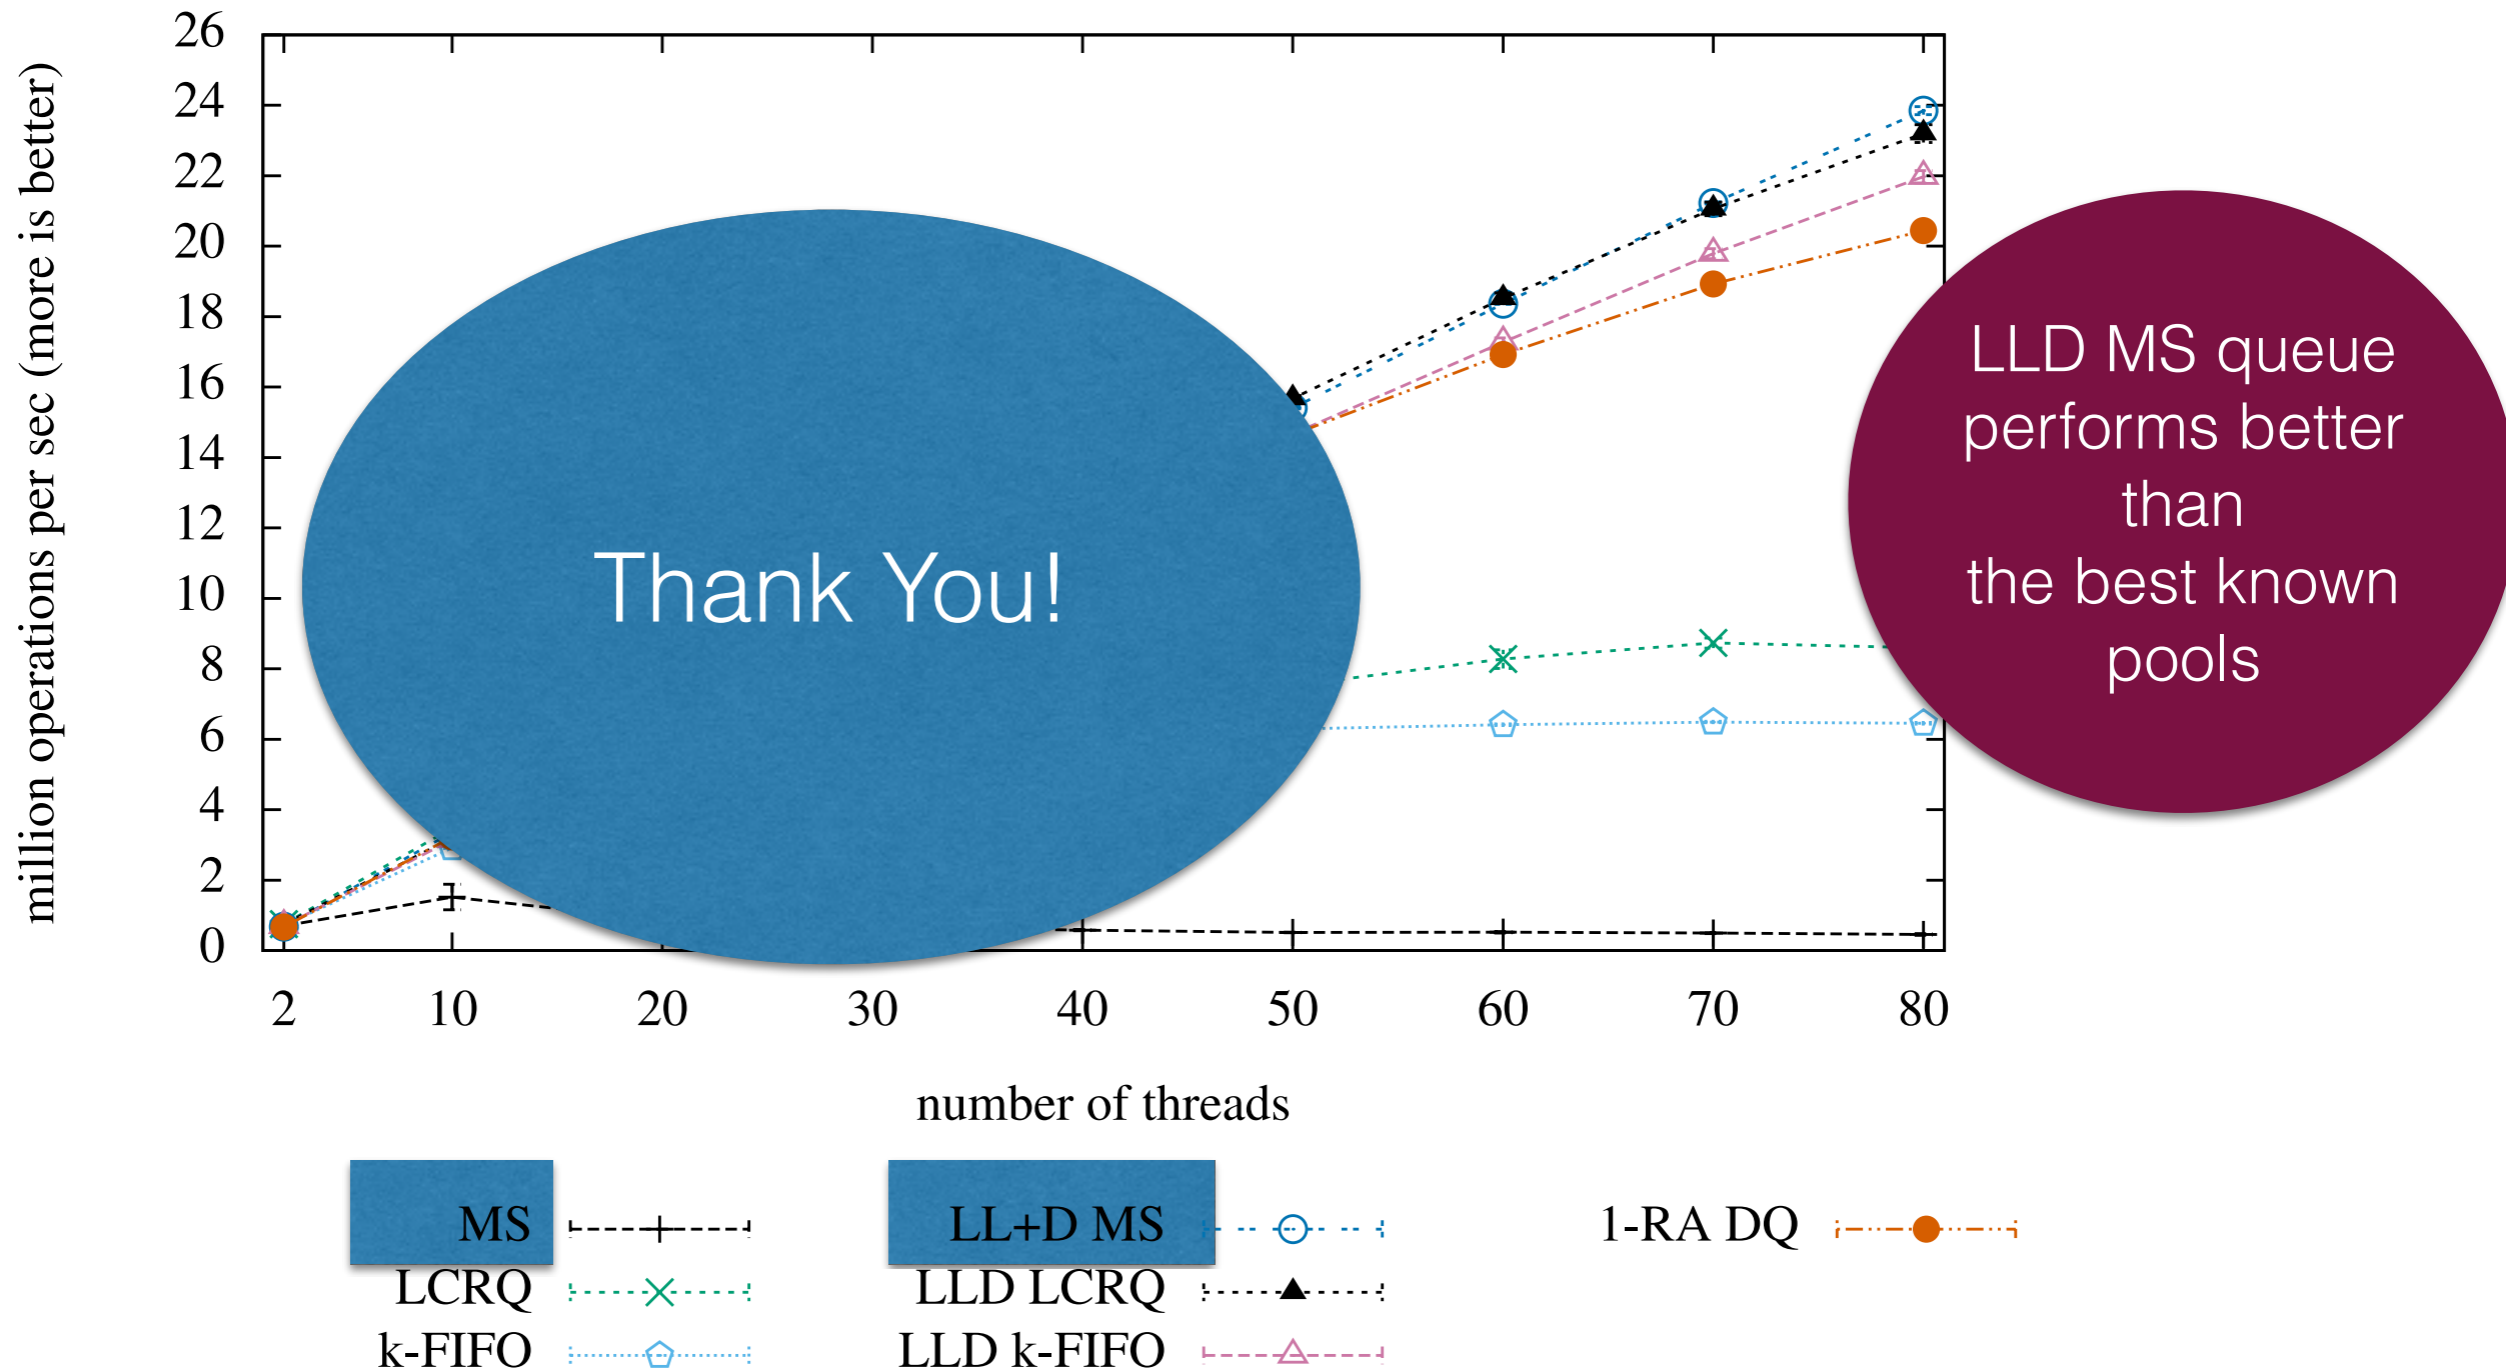
# Performance



LLD  $\phi$   
performs  
significantly better  
than  
 $\phi$

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

# Performance



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