

Concurrent Data Structures: Semantics and Relaxations

Ana Sokolova  UNIVERSITY
of SALZBURG

Background big picture

Background big picture

Computer Science

Background big picture

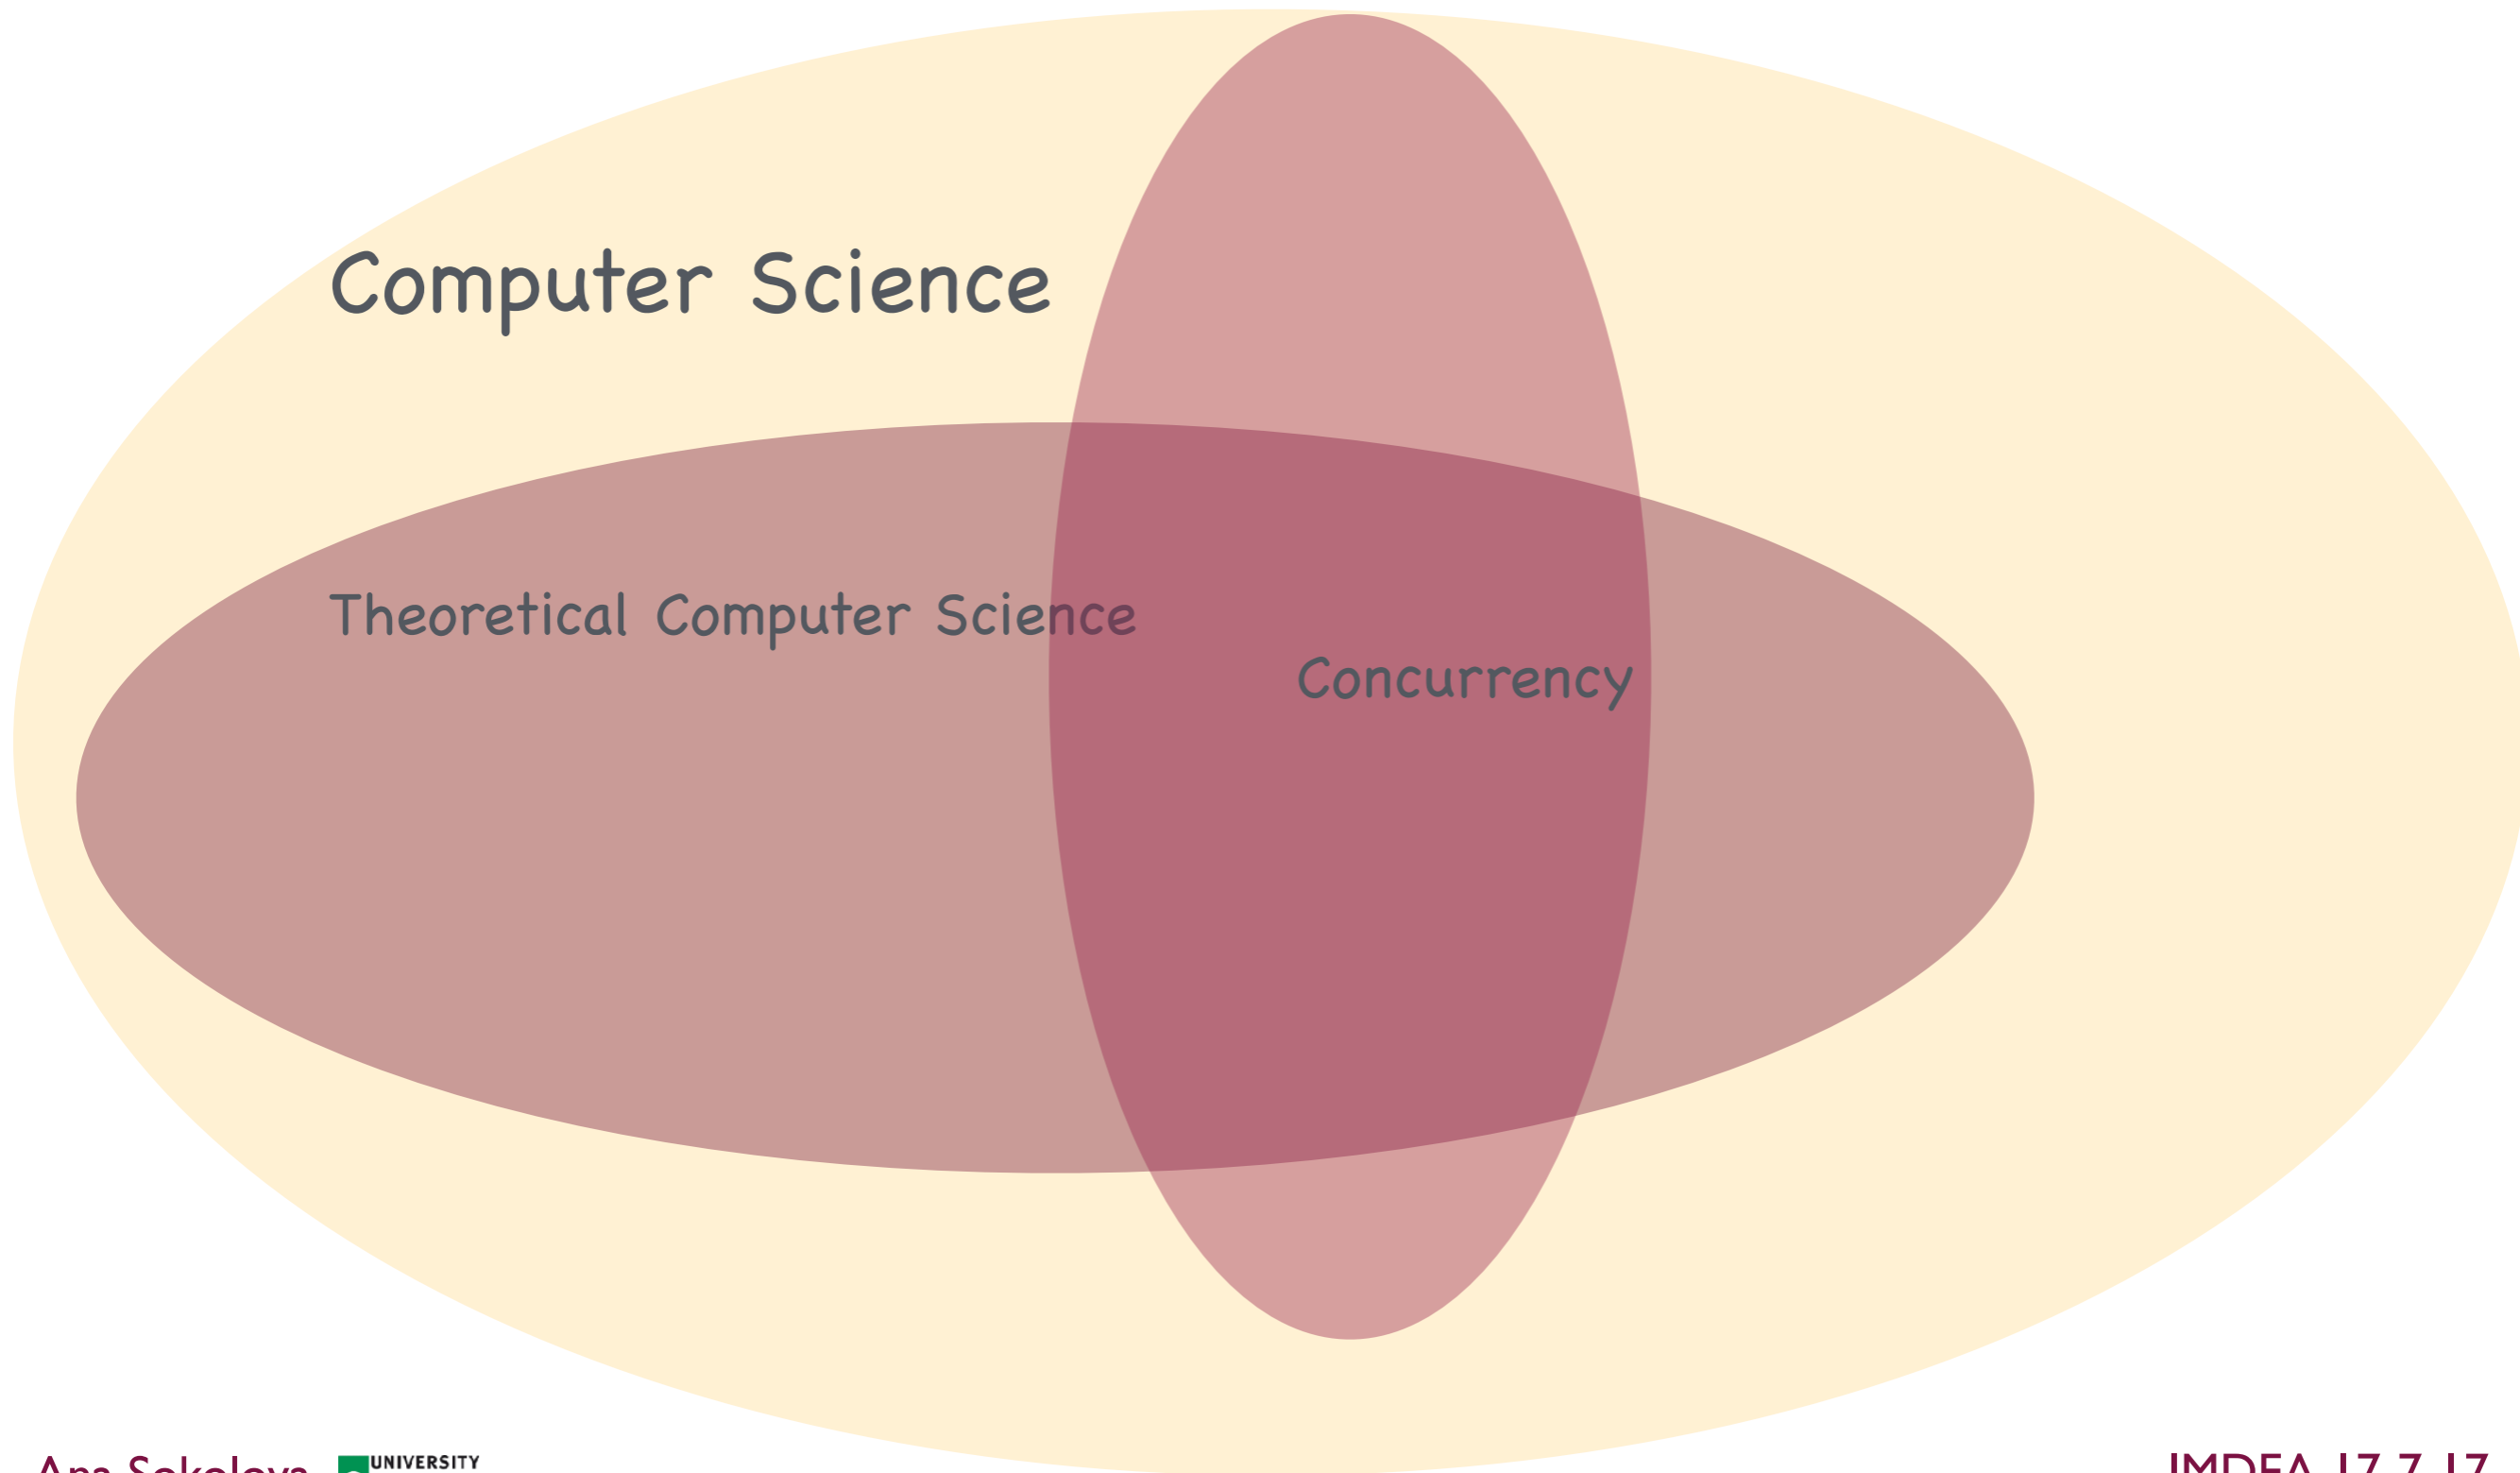


Computer Science

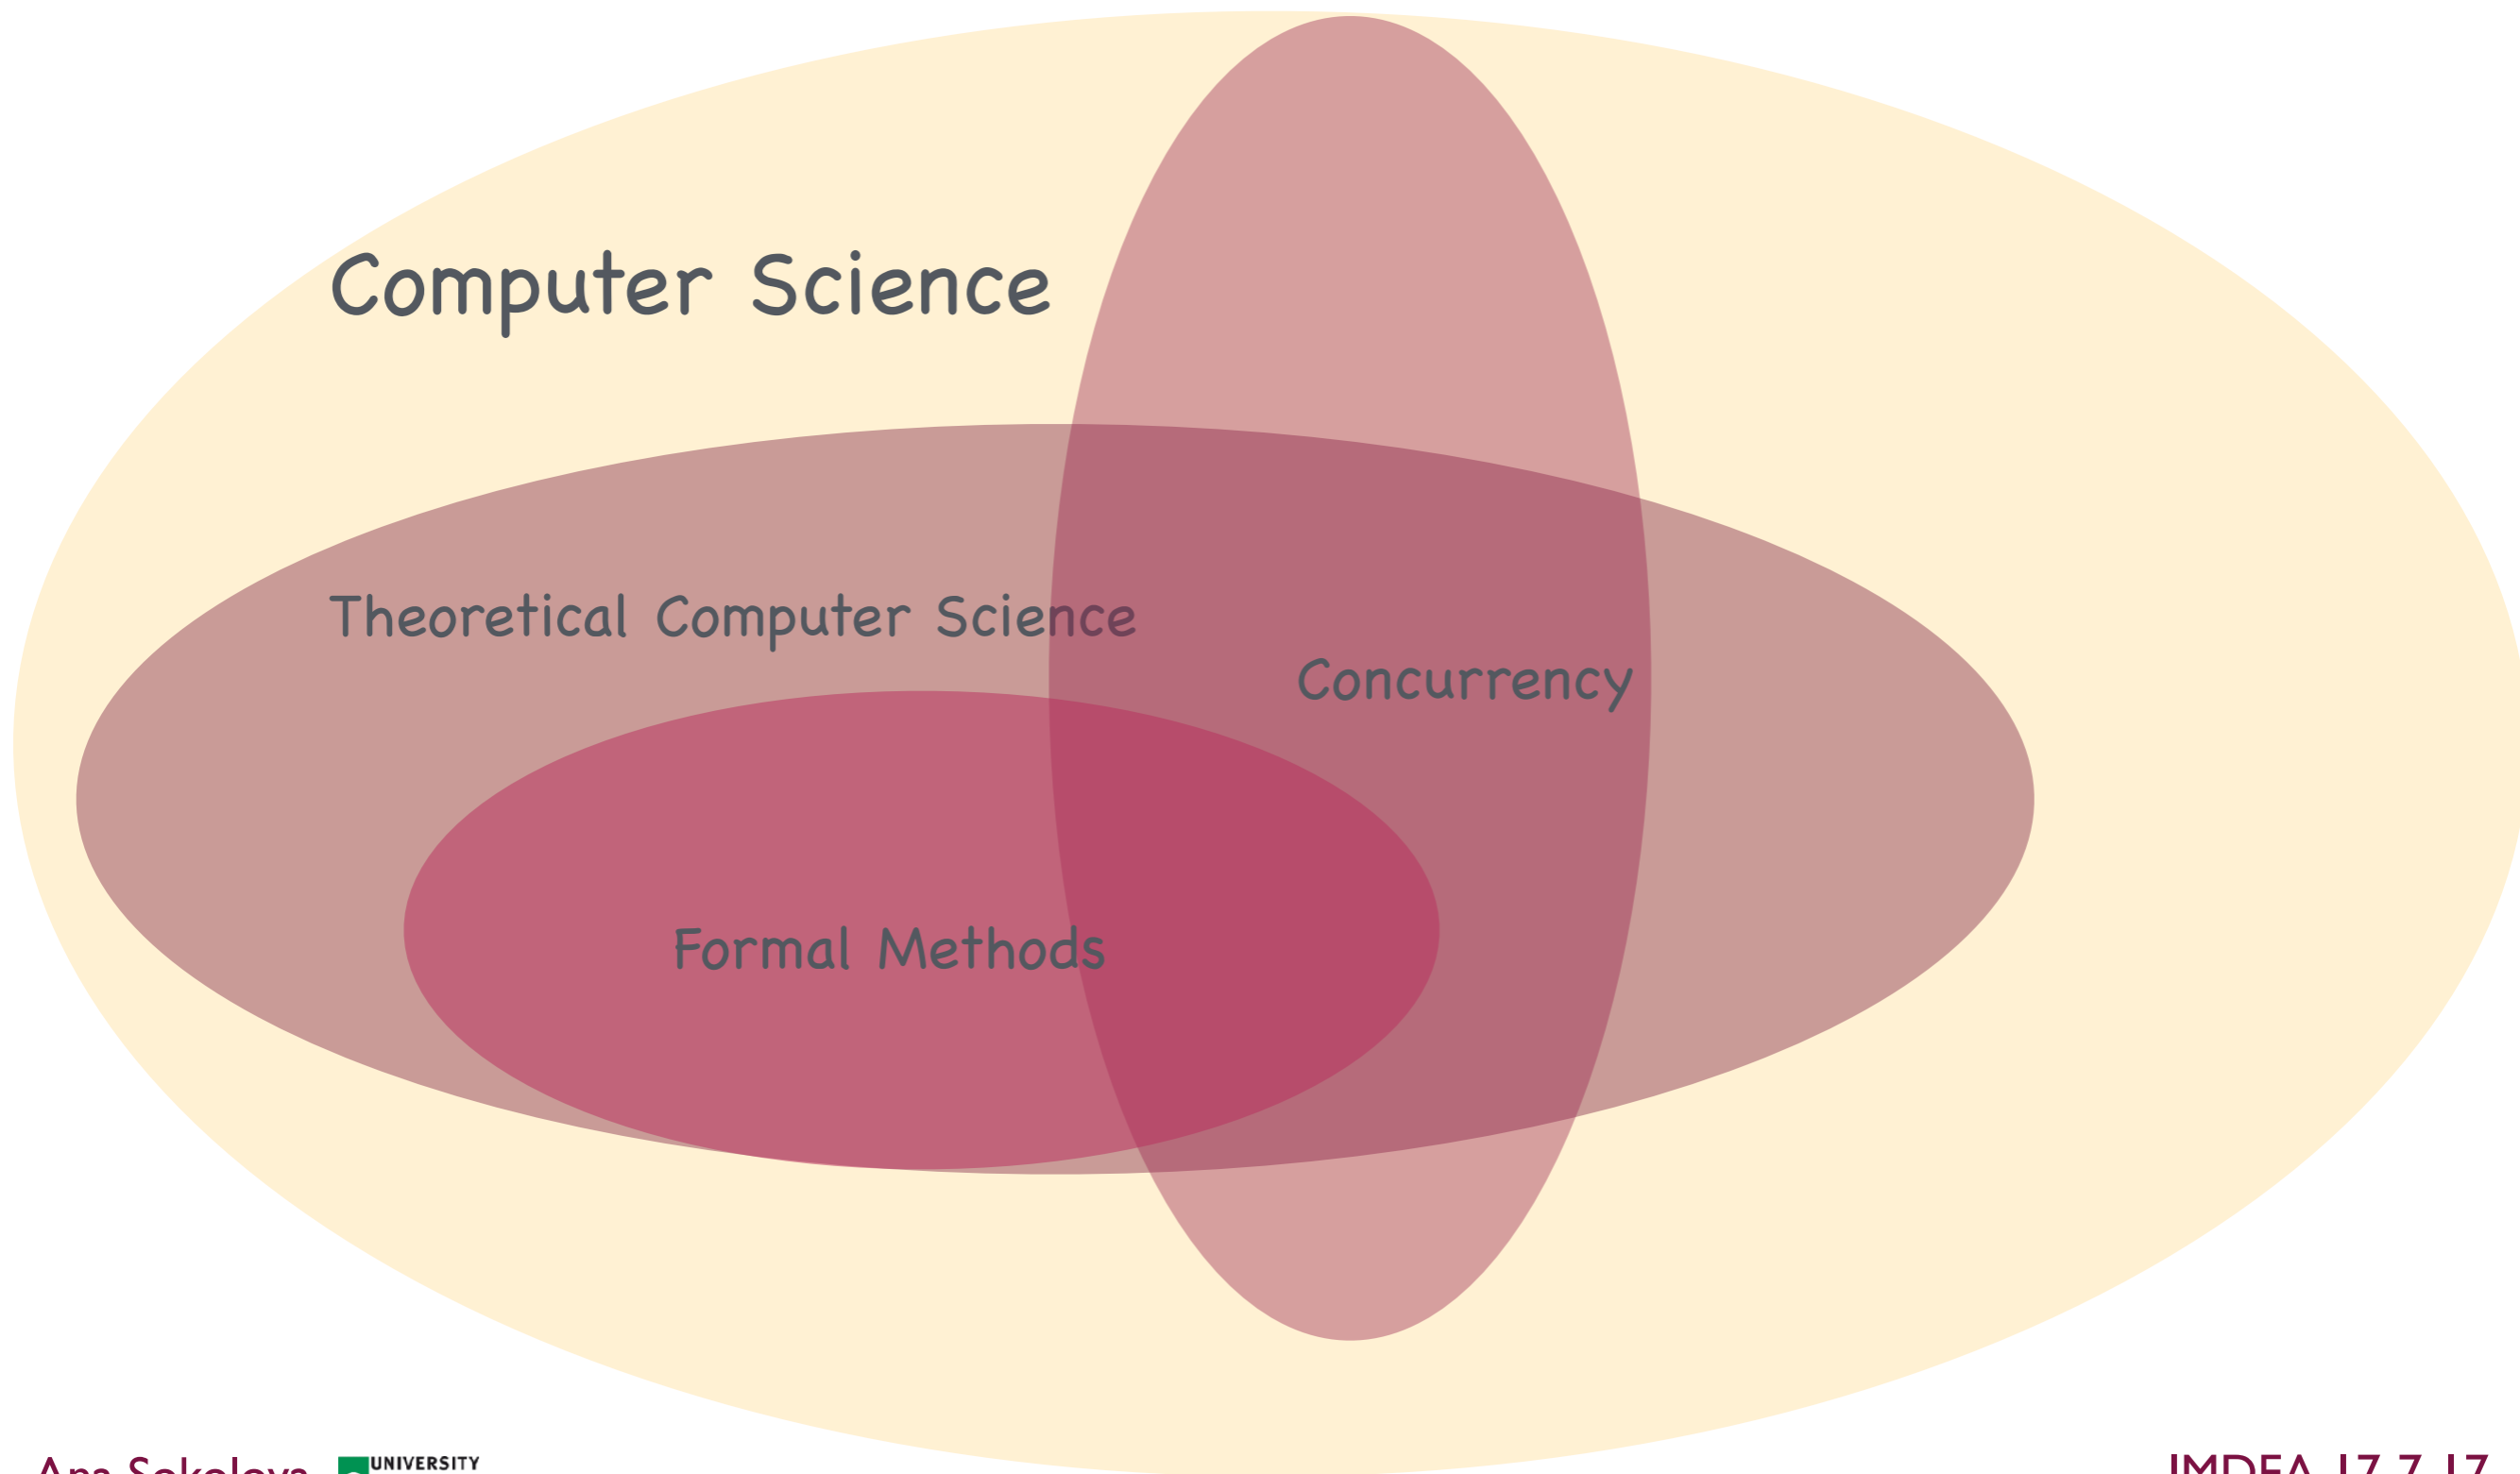
The diagram consists of two nested, horizontally-oriented ovals. The outer oval is light yellow and contains the text 'Computer Science'. The inner oval is a darker, reddish-brown color and contains the text 'Theoretical Computer Science'. The inner oval is centered within the outer oval, illustrating that theoretical computer science is a subset of the broader field of computer science.

Theoretical Computer Science

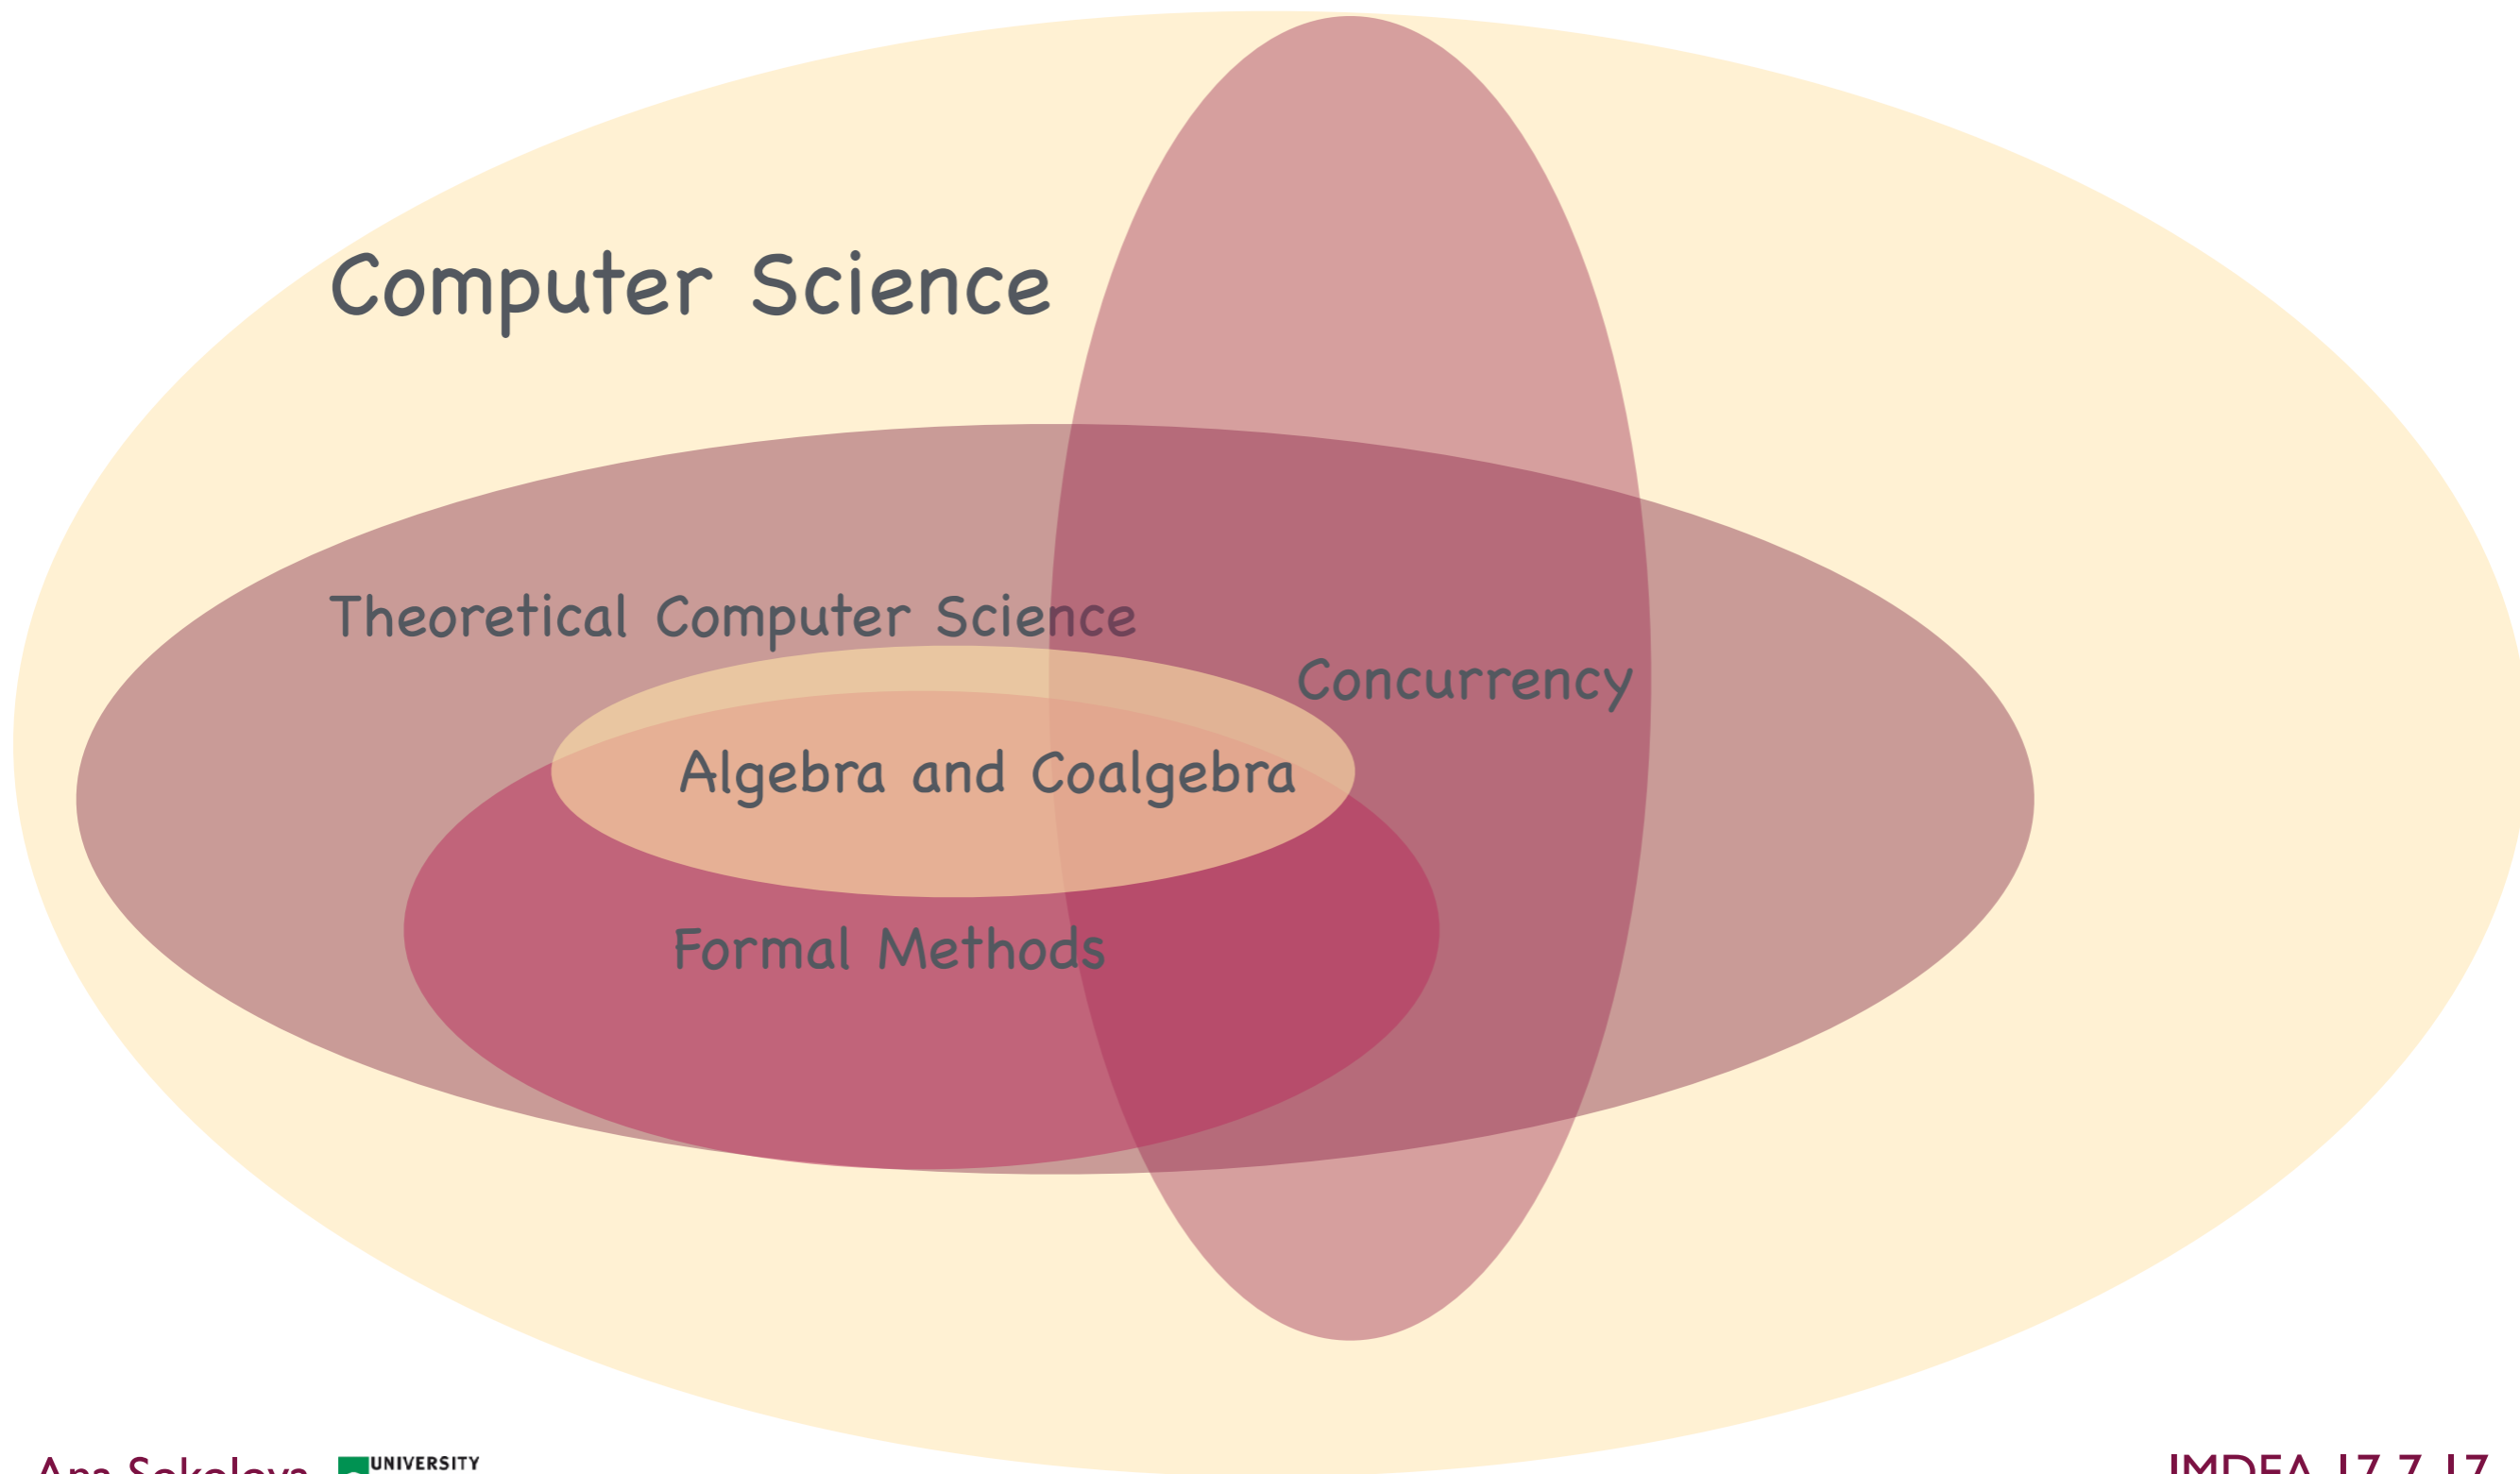
Background big picture



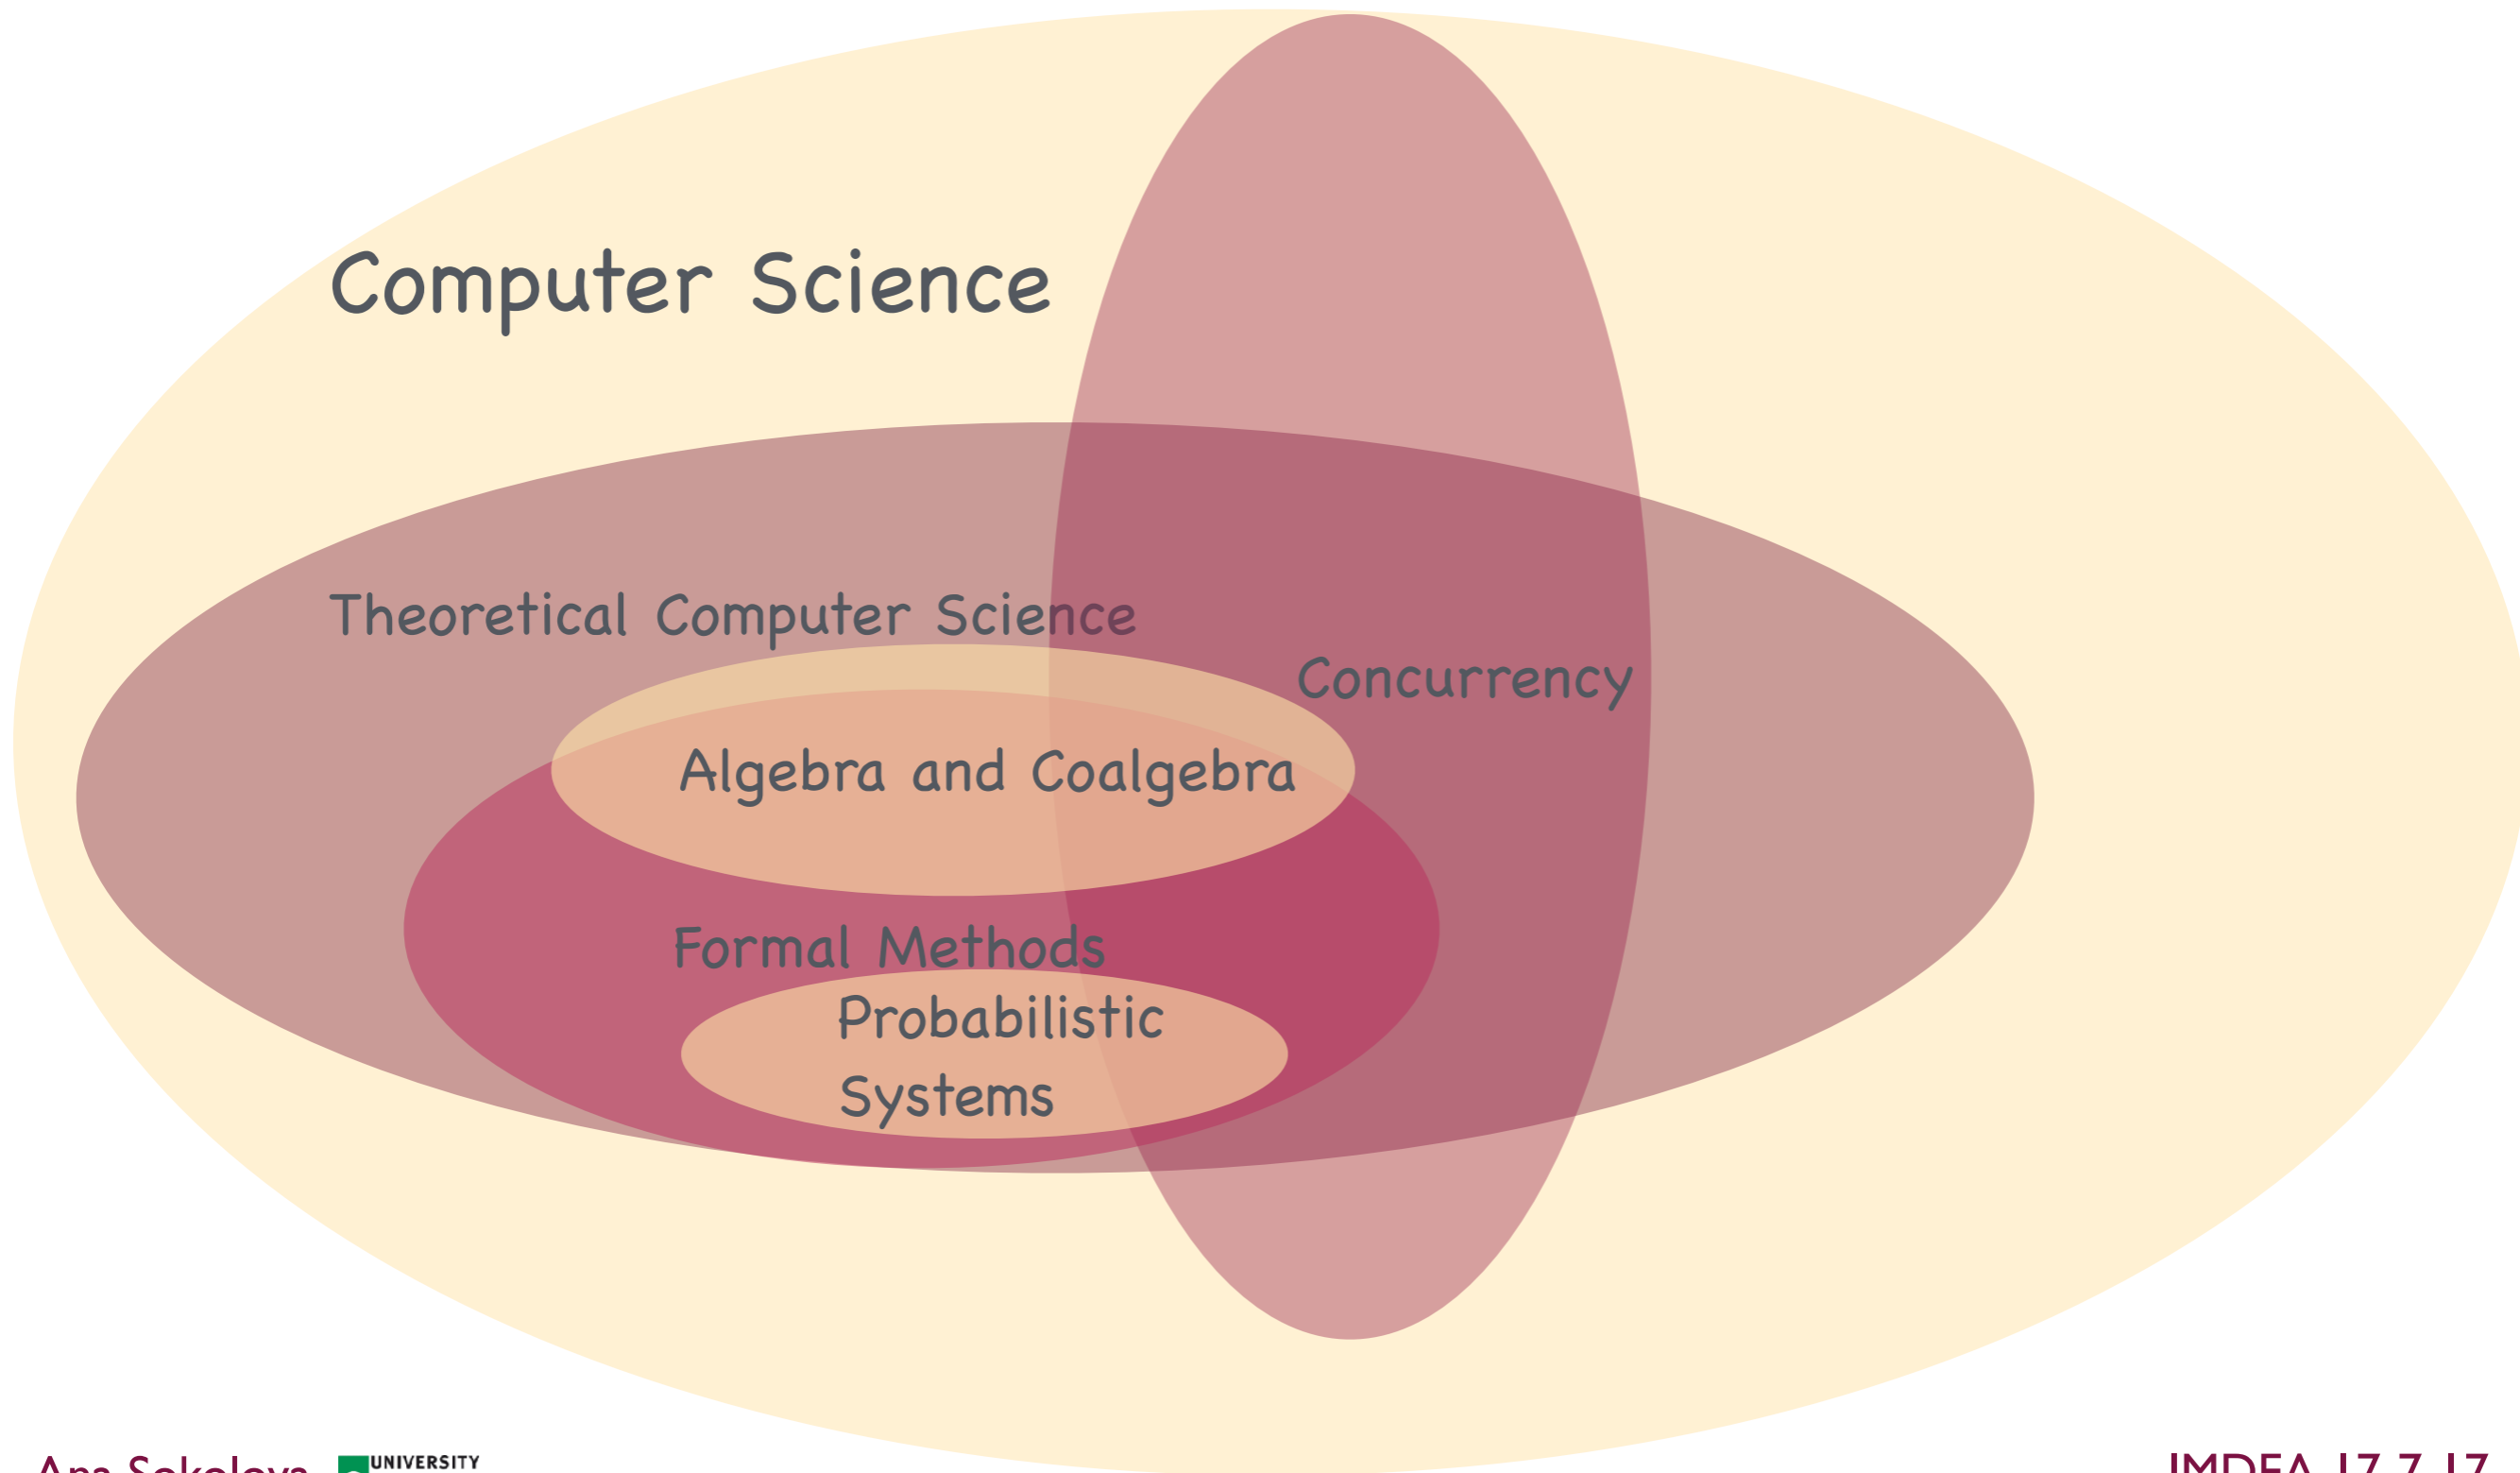
Background big picture



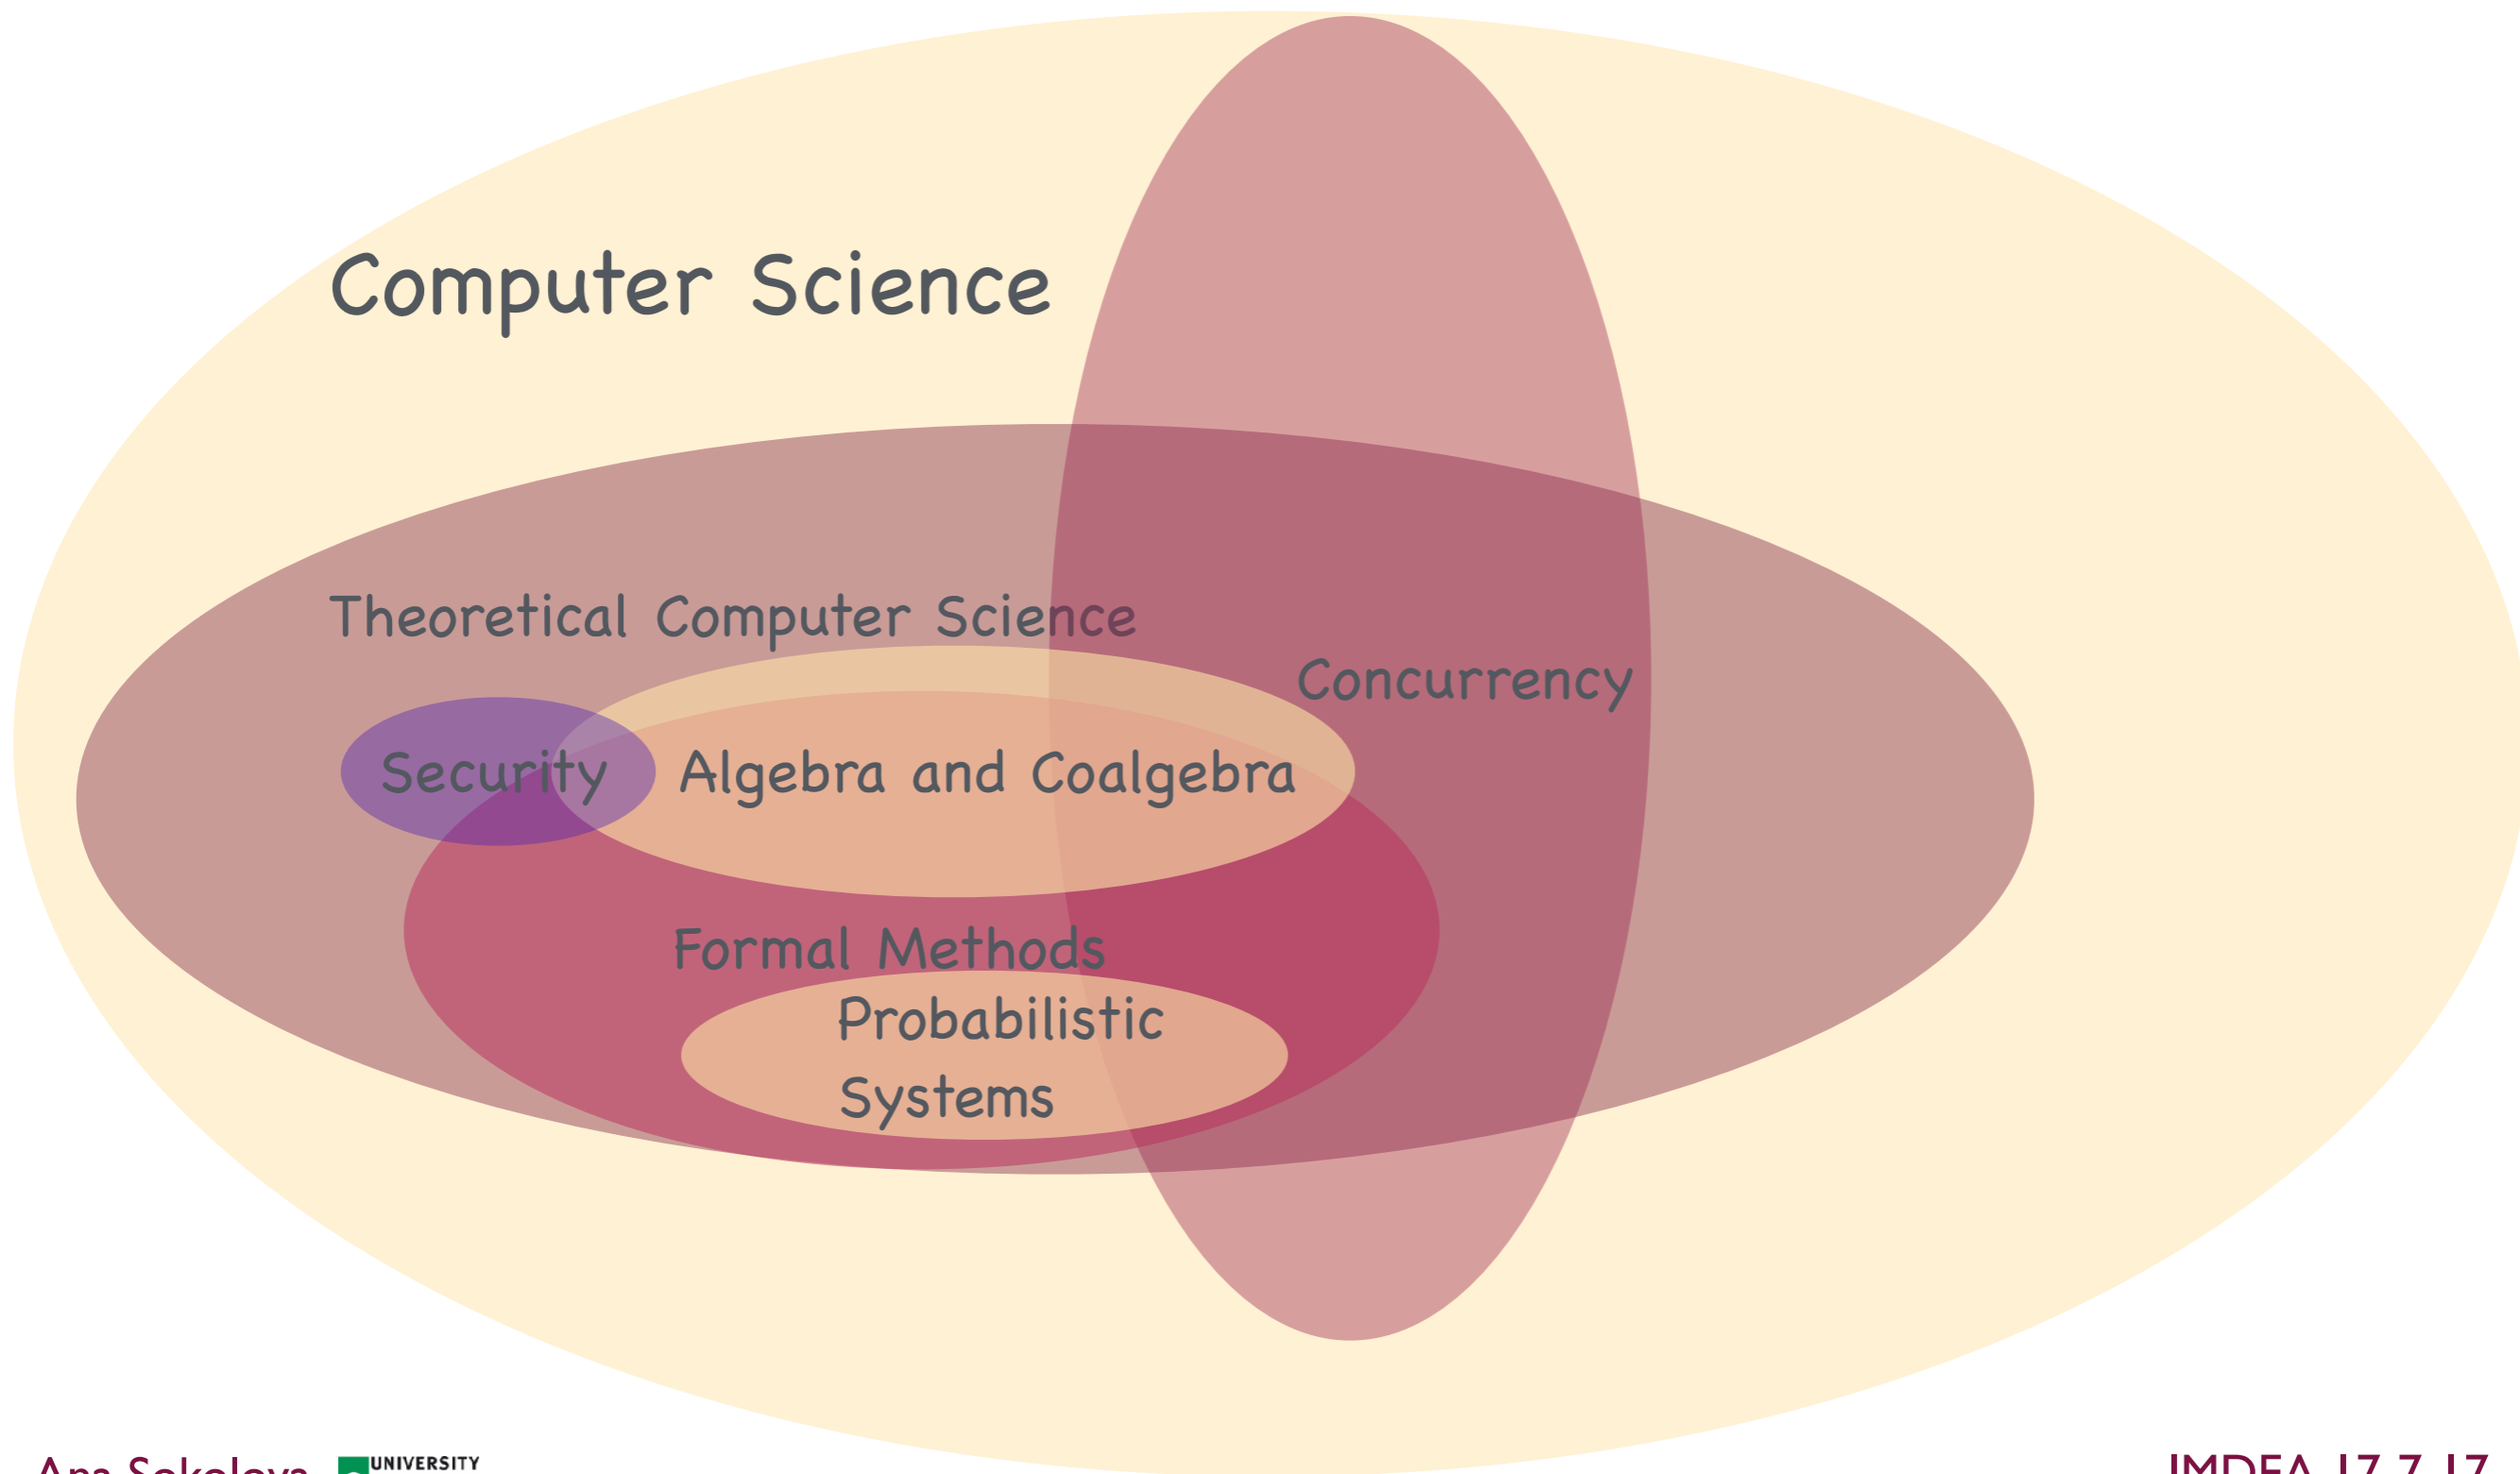
Background big picture



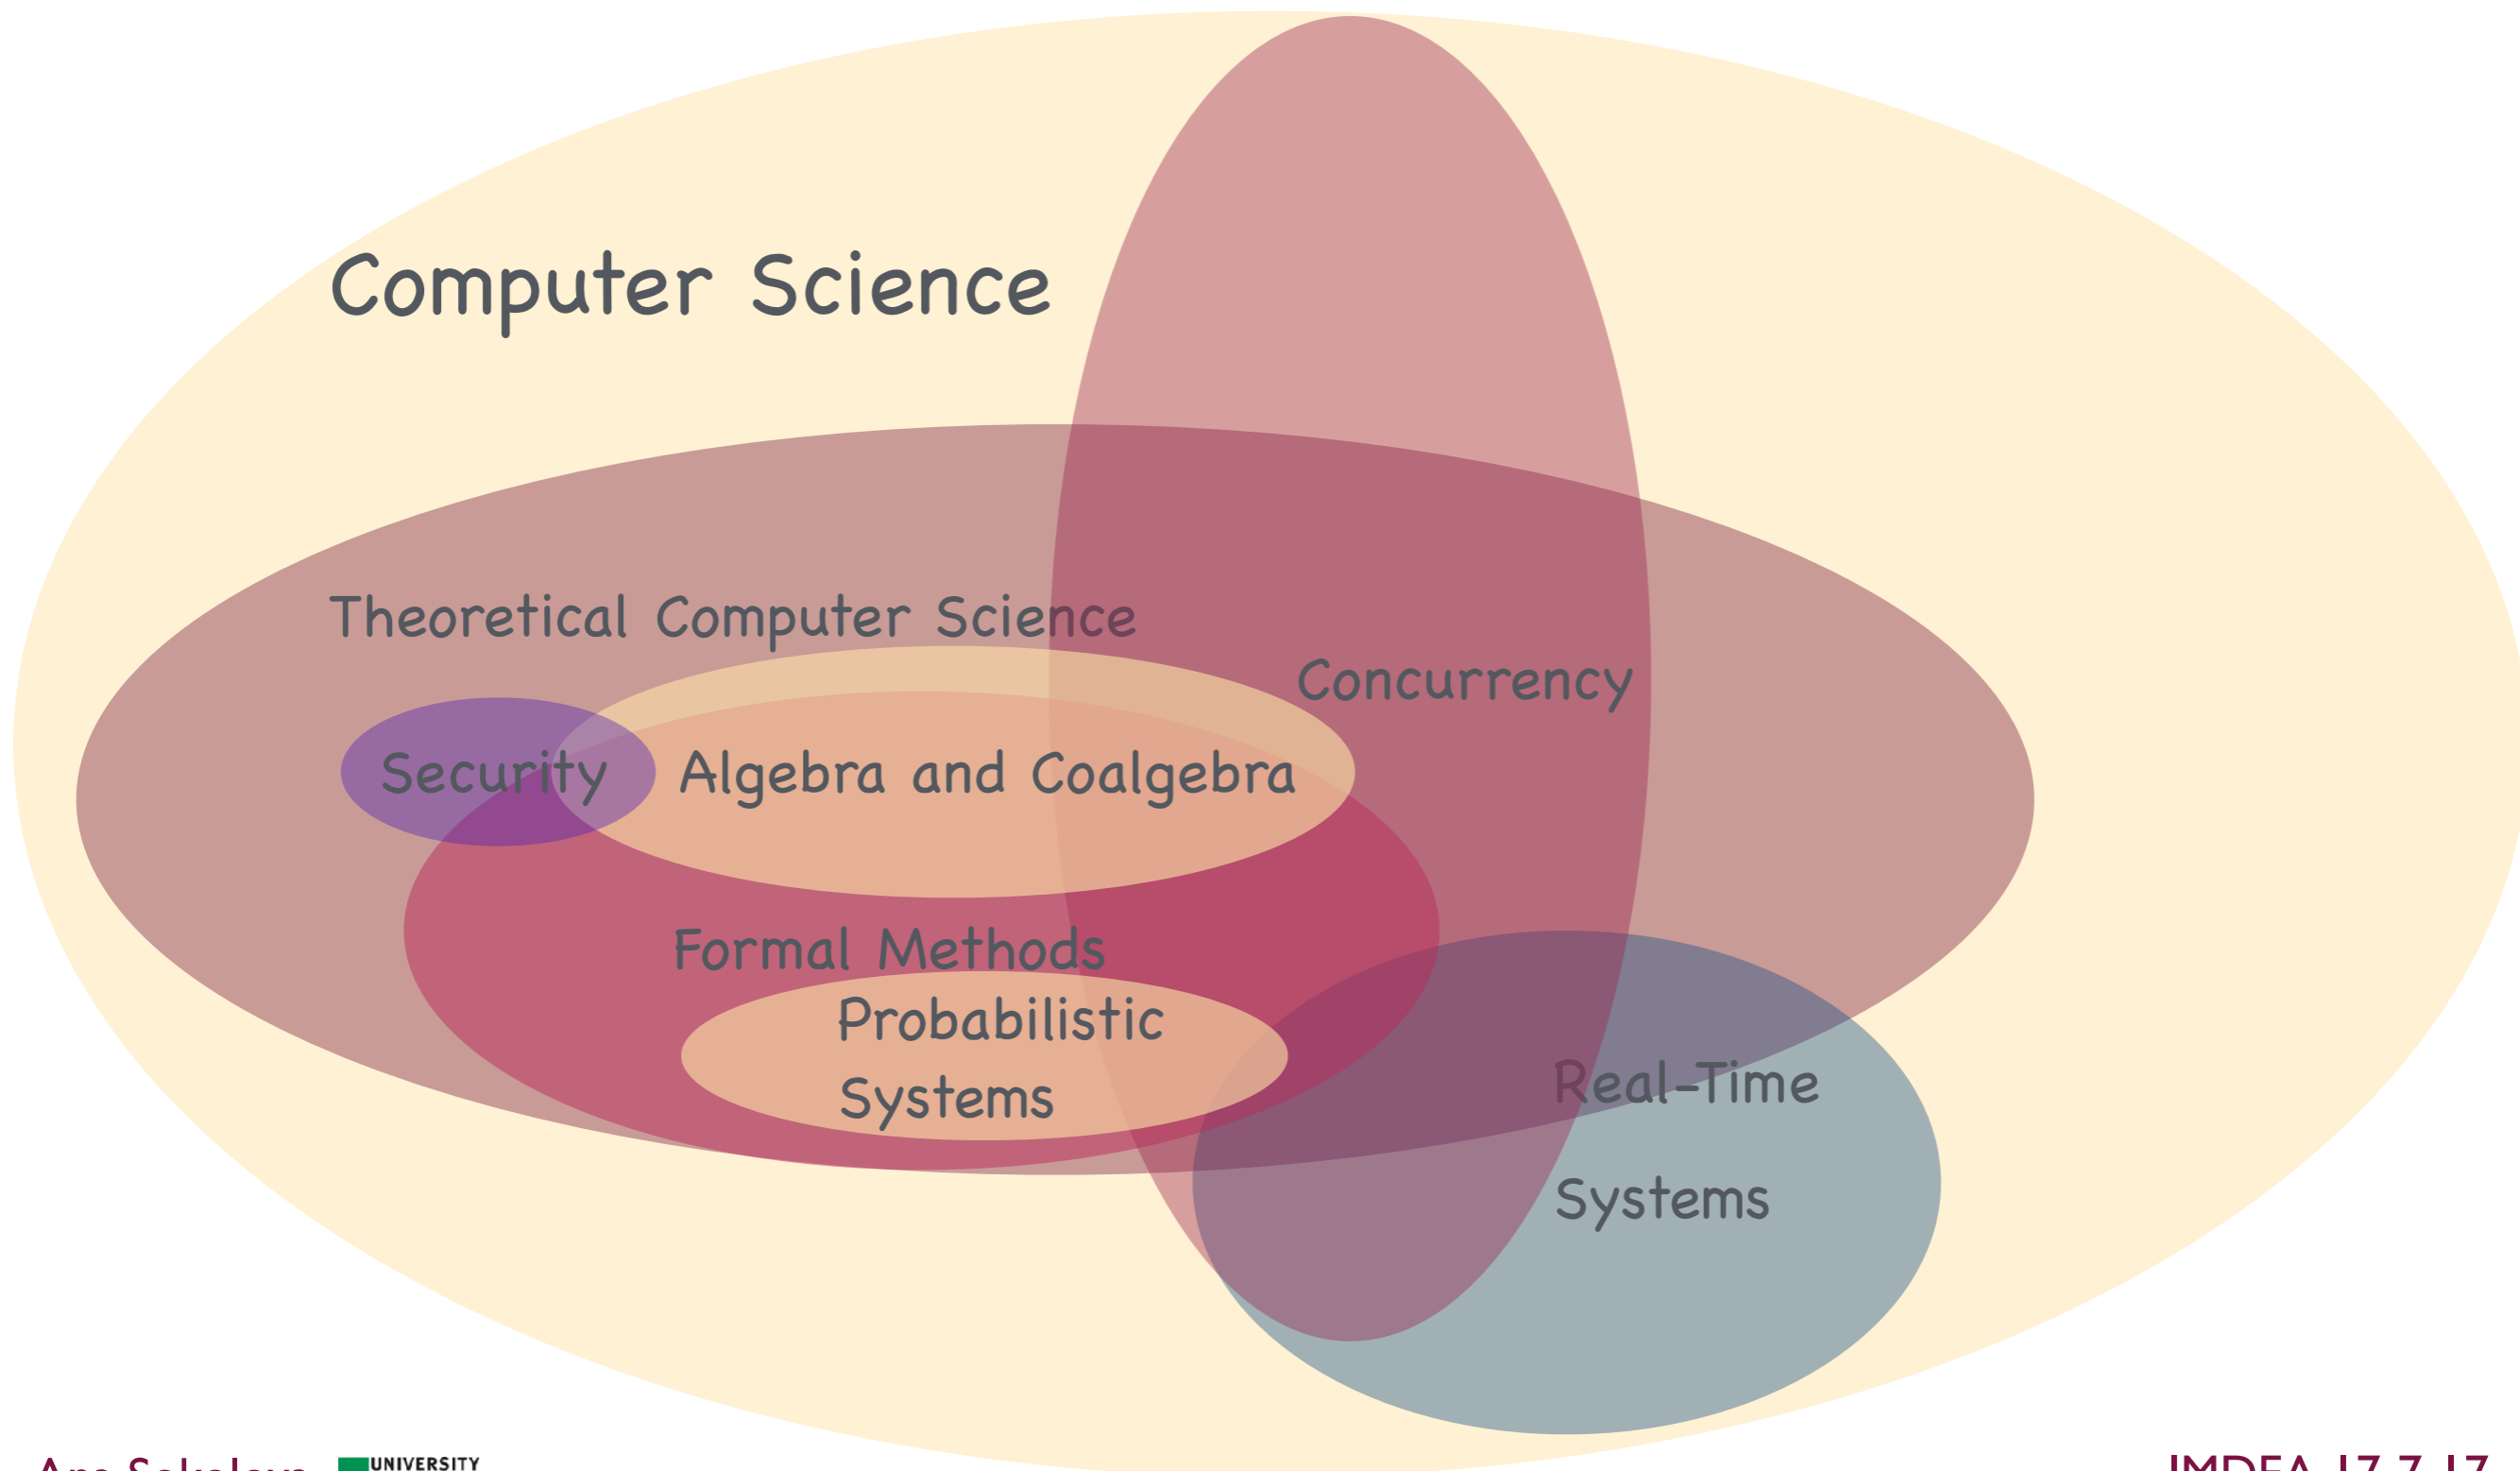
Background big picture



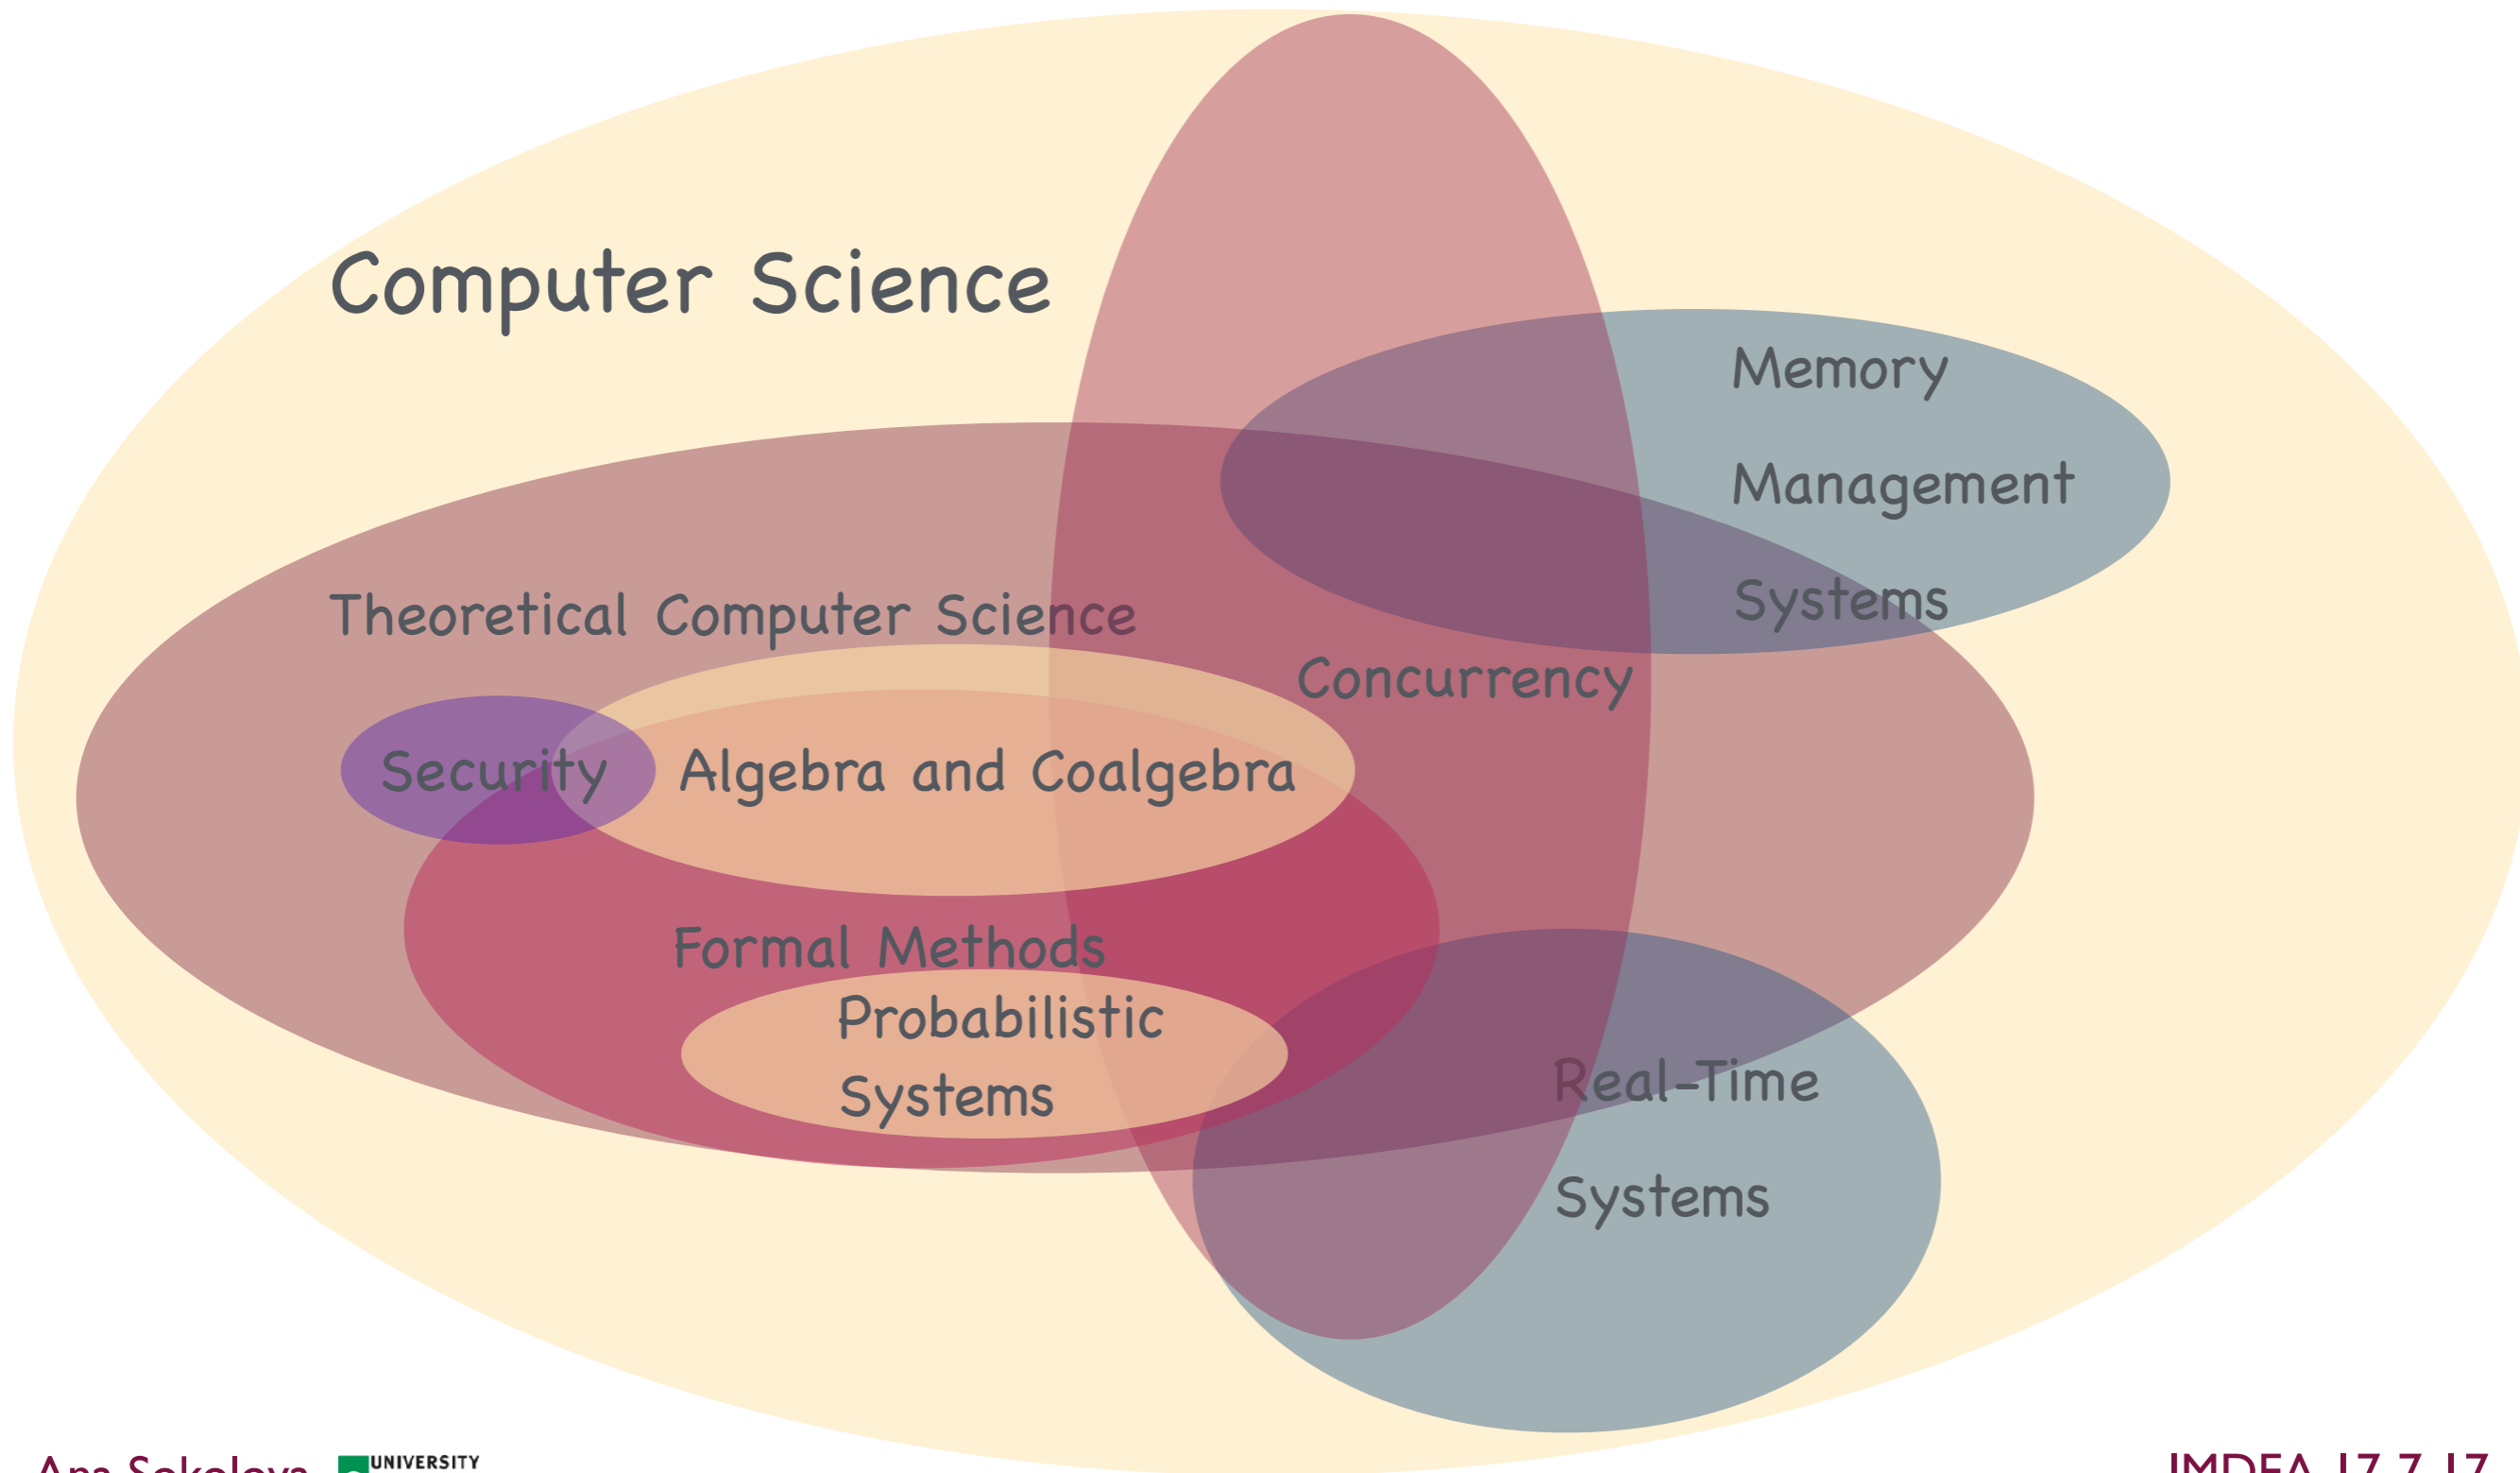
Background big picture



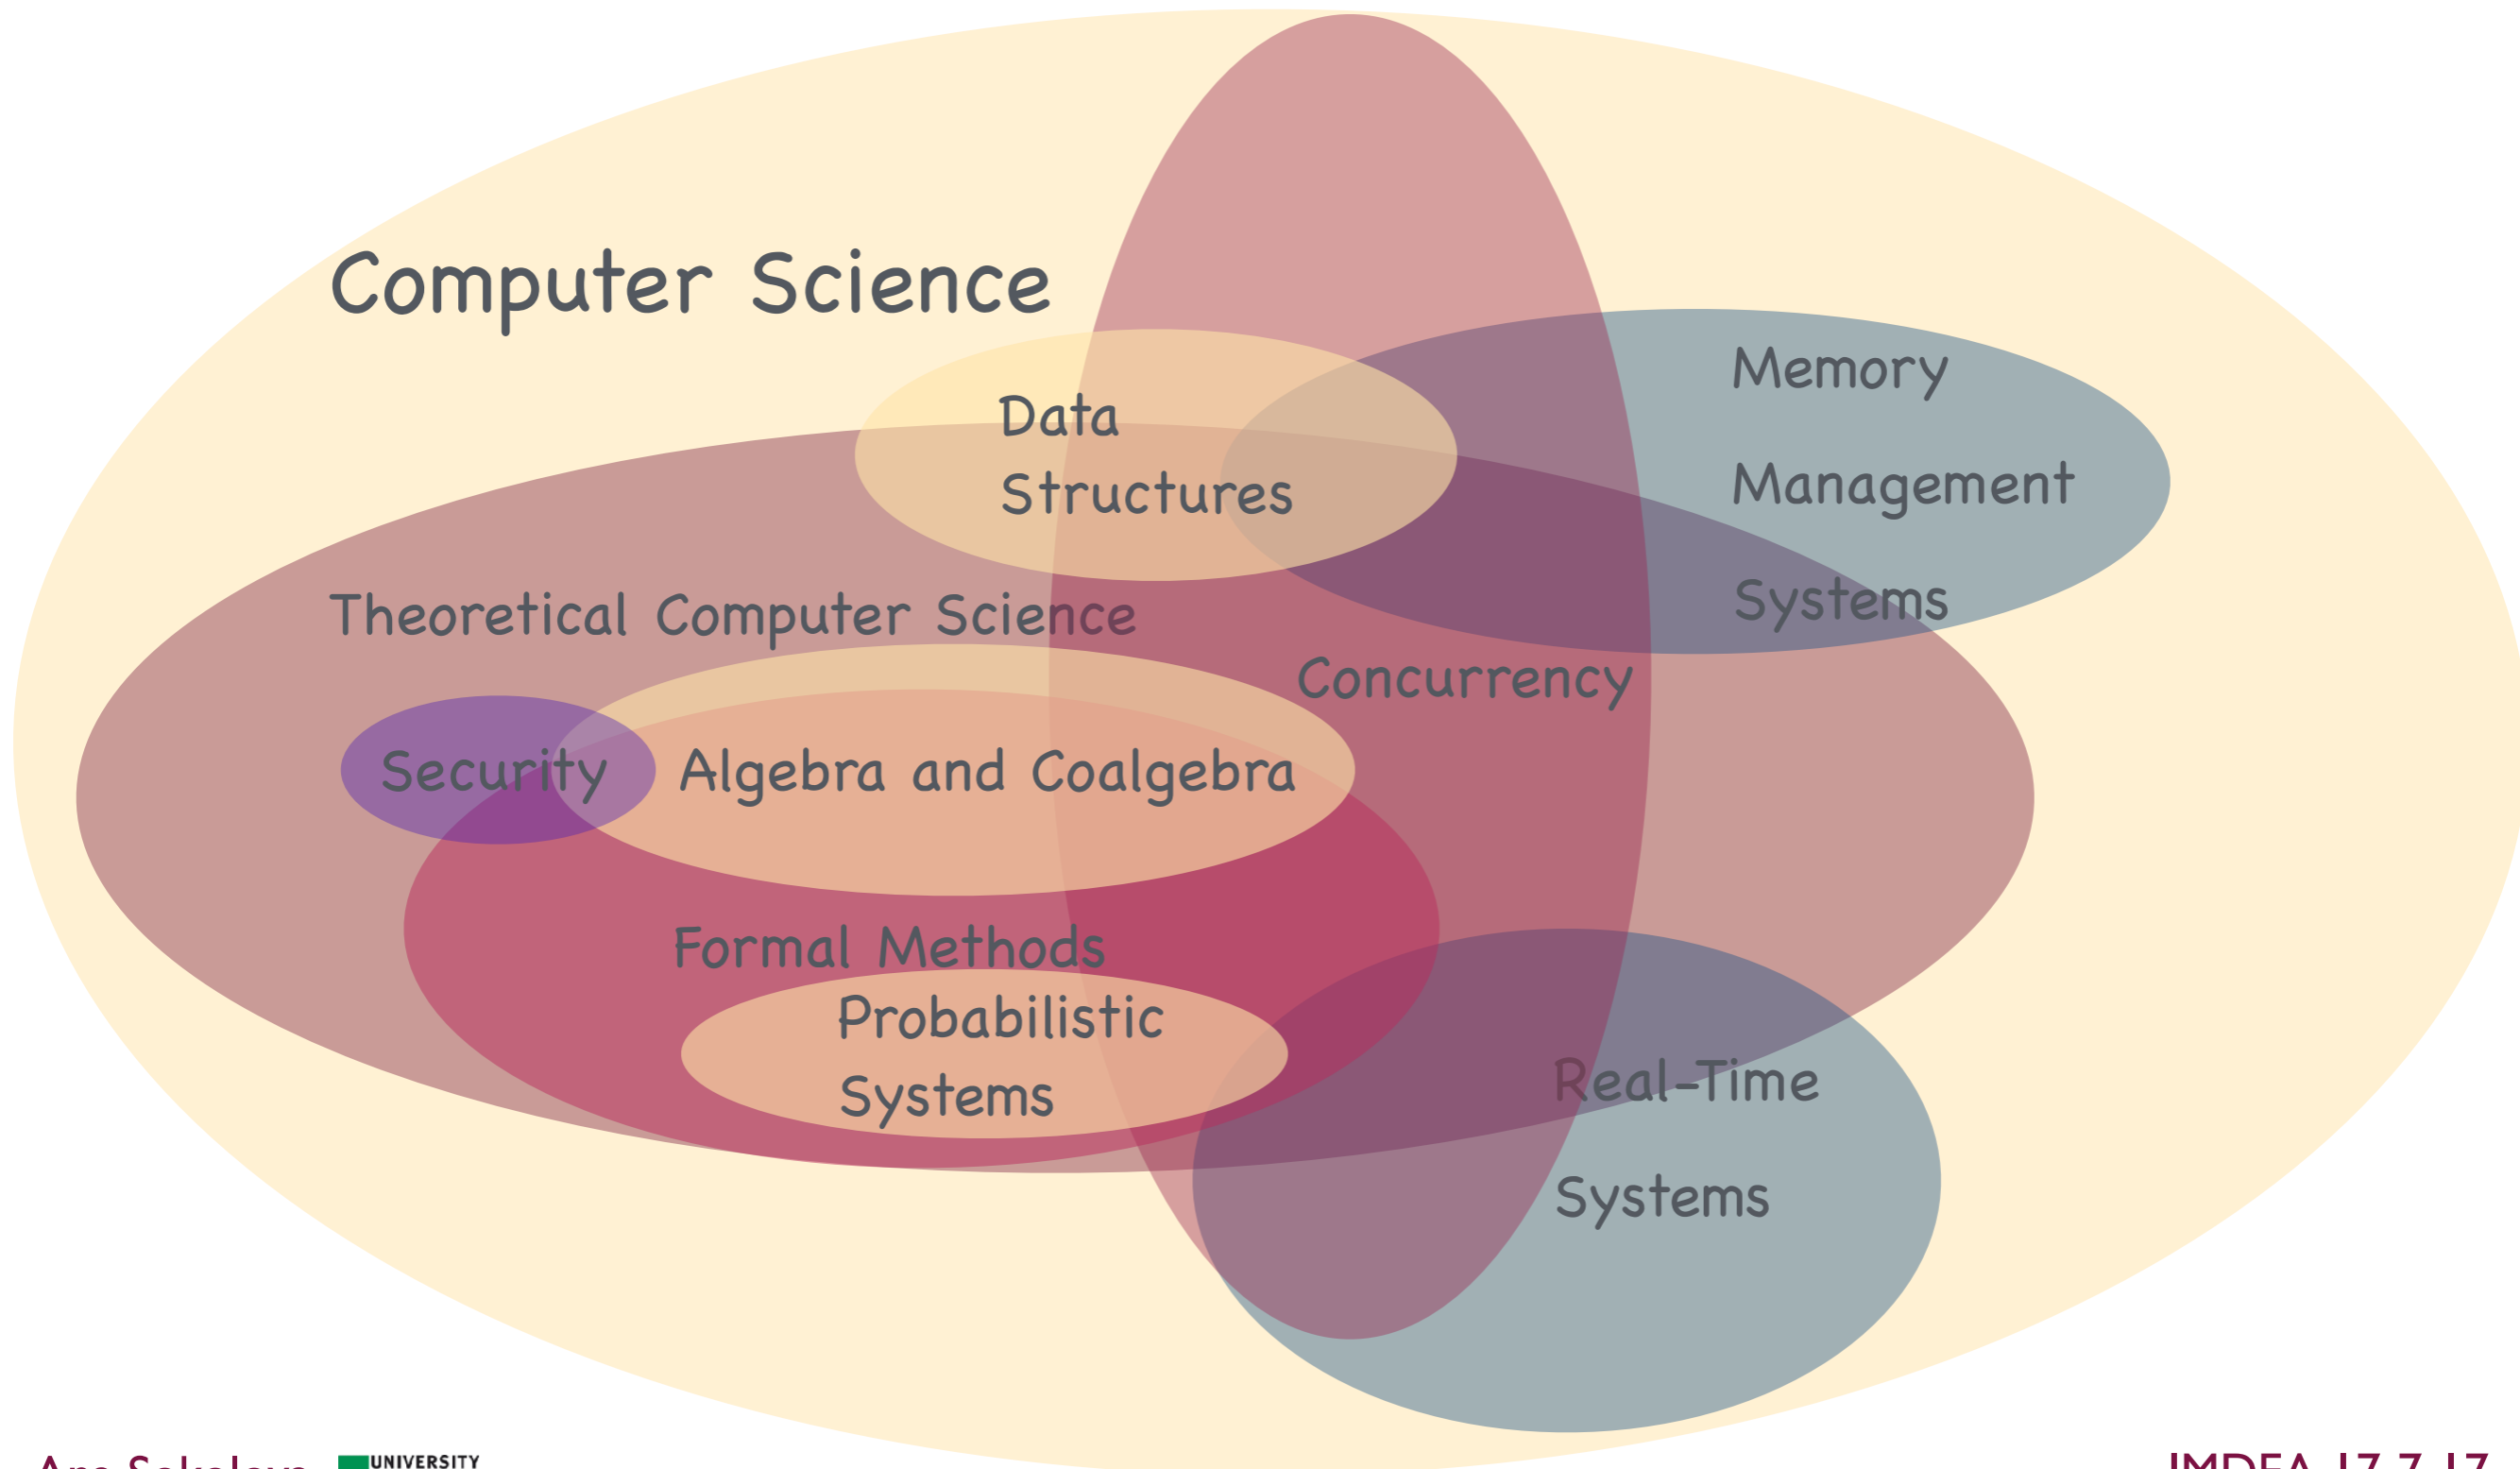
Background big picture



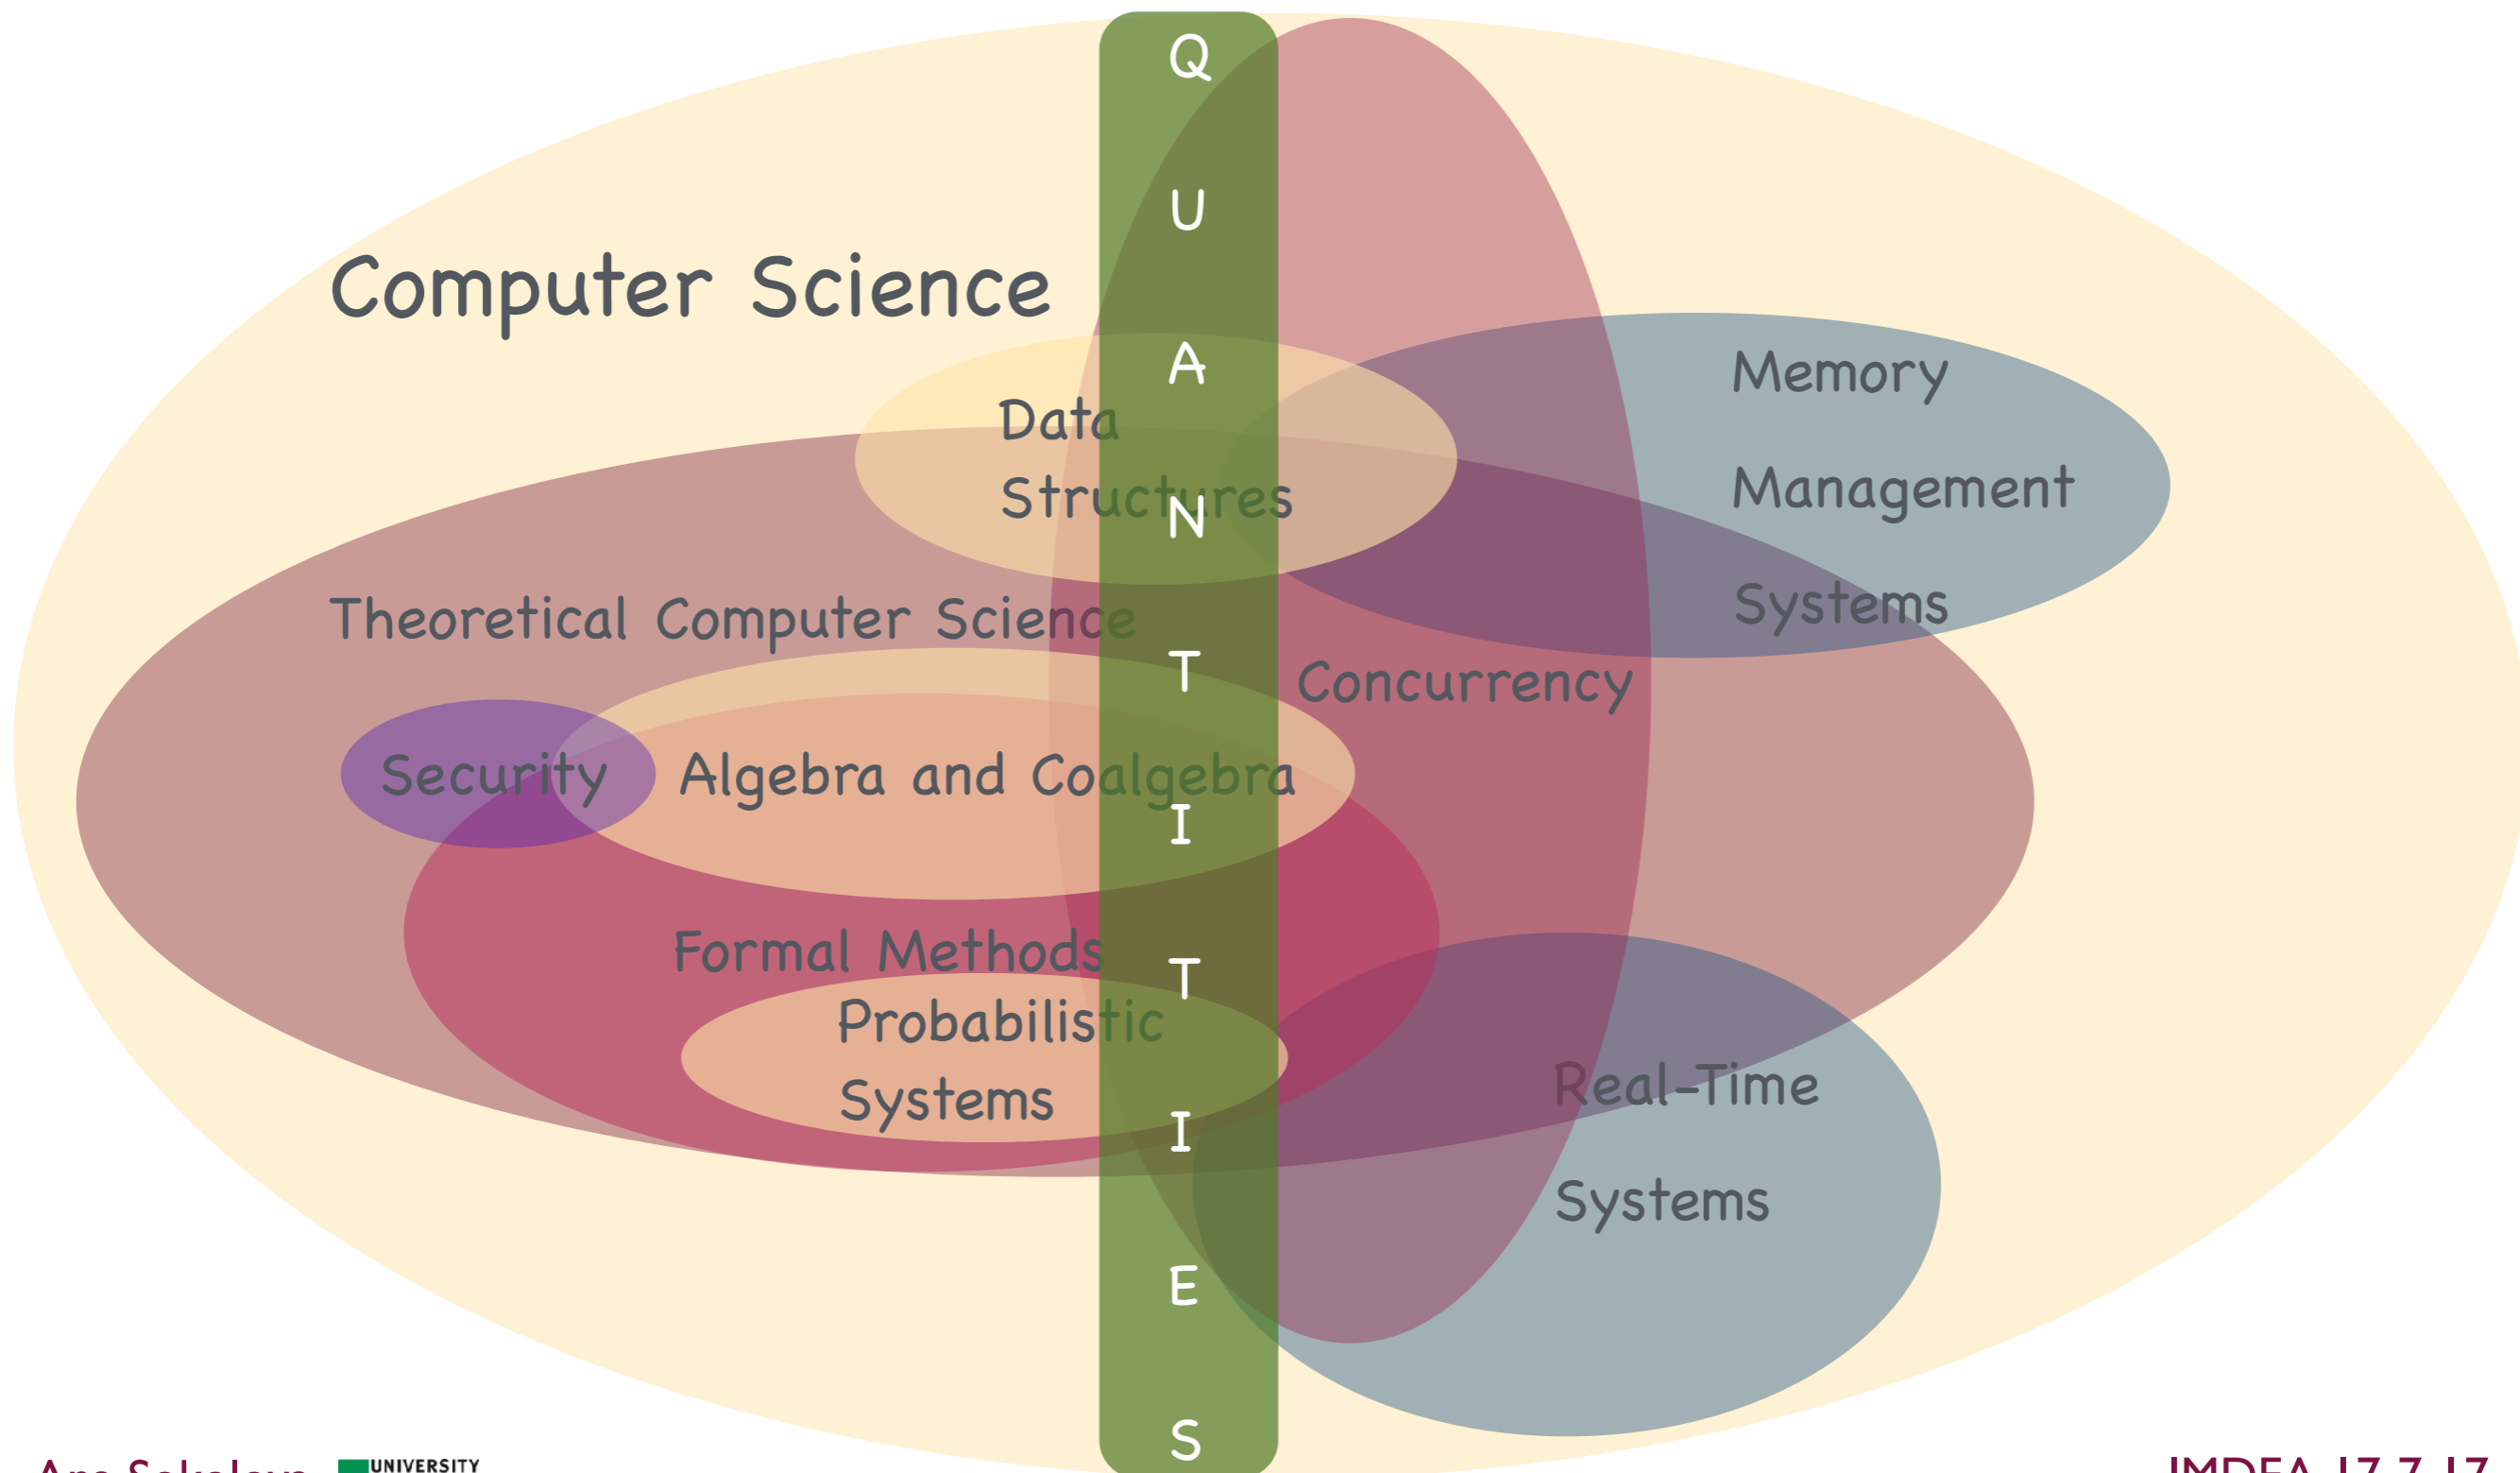
Background big picture



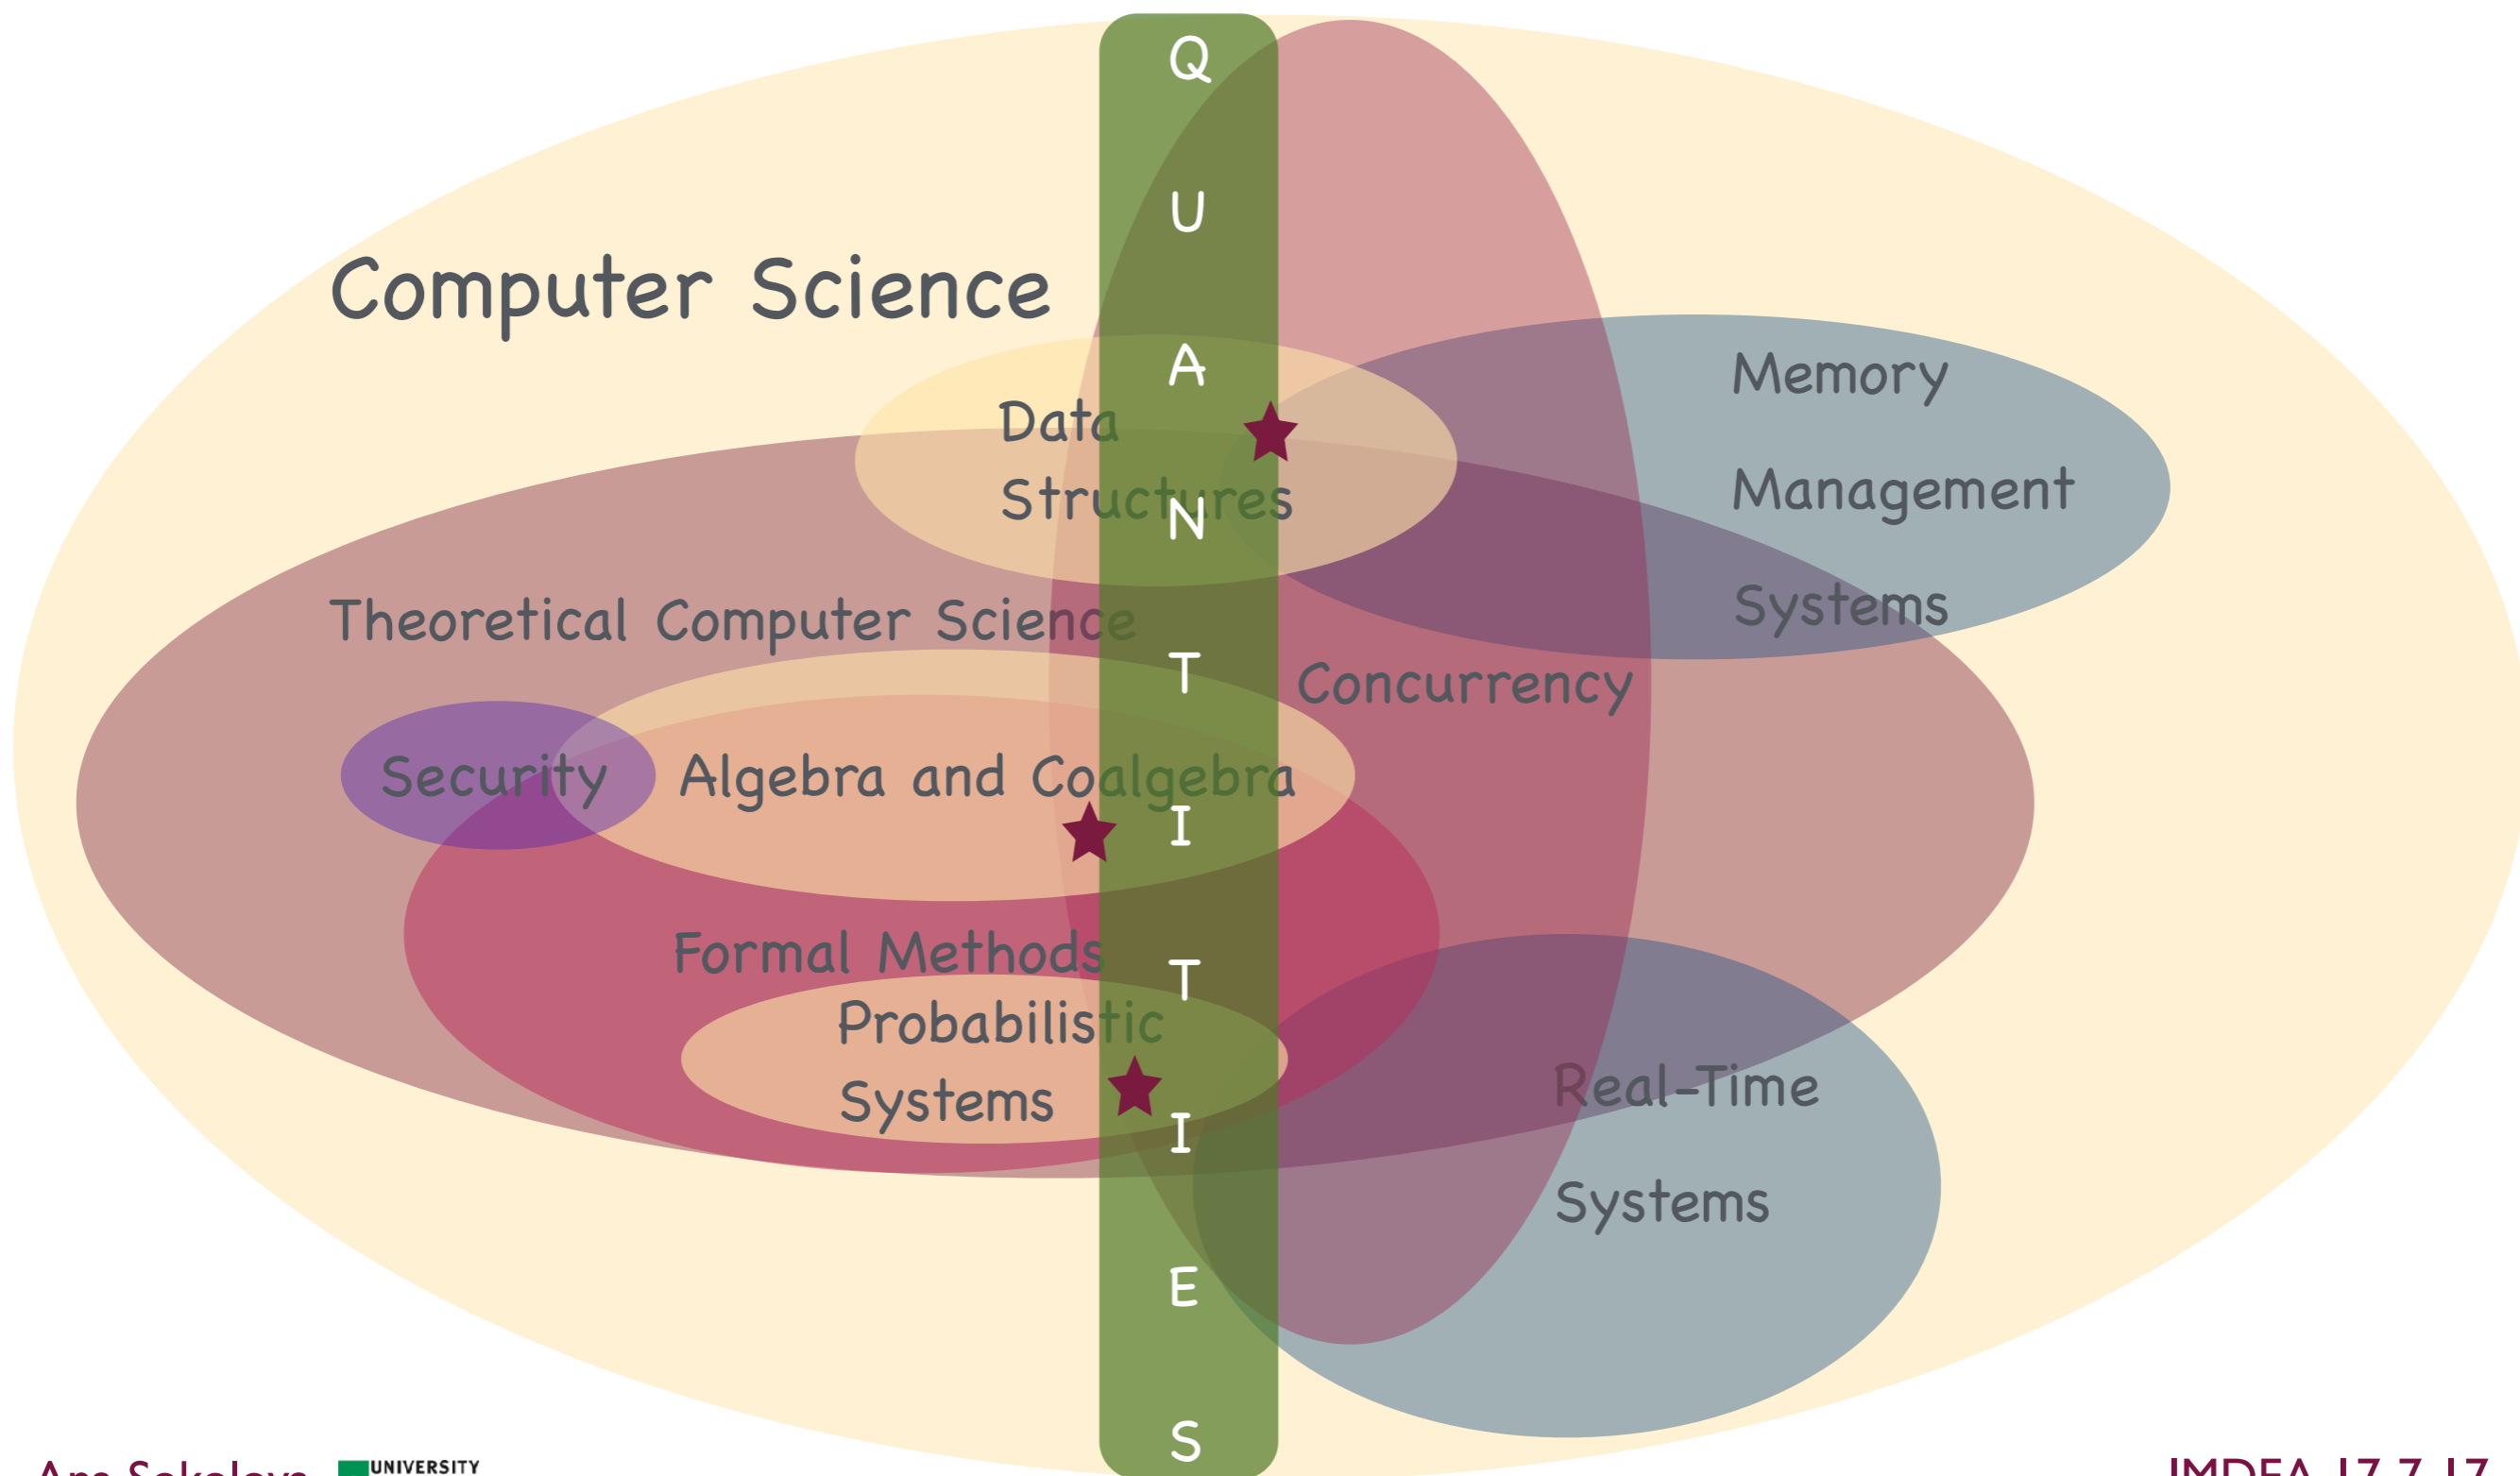
Background big picture



Background big picture



Favourites



Concurrent Data Structures: Semantics and Relaxations

Ana Sokolova  UNIVERSITY
of SALZBURG

Concurrent Data Structures: Correctness and Performance

Semantics of concurrent data structures

Semantics of concurrent data structures

e.g. pools, queues, stacks

Semantics of concurrent data structures

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

e.g. pools, queues, stacks

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

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

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

Linearizability [Herlihy, Wing '90]



Sequential Consistency [Lamport'79]

Consistency conditions

there exists a legal
sequence that preserves
precedence

Linearizability [Herlihy, Wing '90]



Sequential Consistency [Lamport'79]

Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]

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

↓

Sequential Consistency [Lamport'79]

Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]



Sequential Consistency [Lamport'79]

Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]

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

Sequential Consistency [Lamport'79]

Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]

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)

Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]

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

Consistency conditions


there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]



Sequential Consistency [Lamport'79]

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



Consistency conditions

there exists a legal sequence that preserves precedence

Linearizability [Herlihy, Wing '90]

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

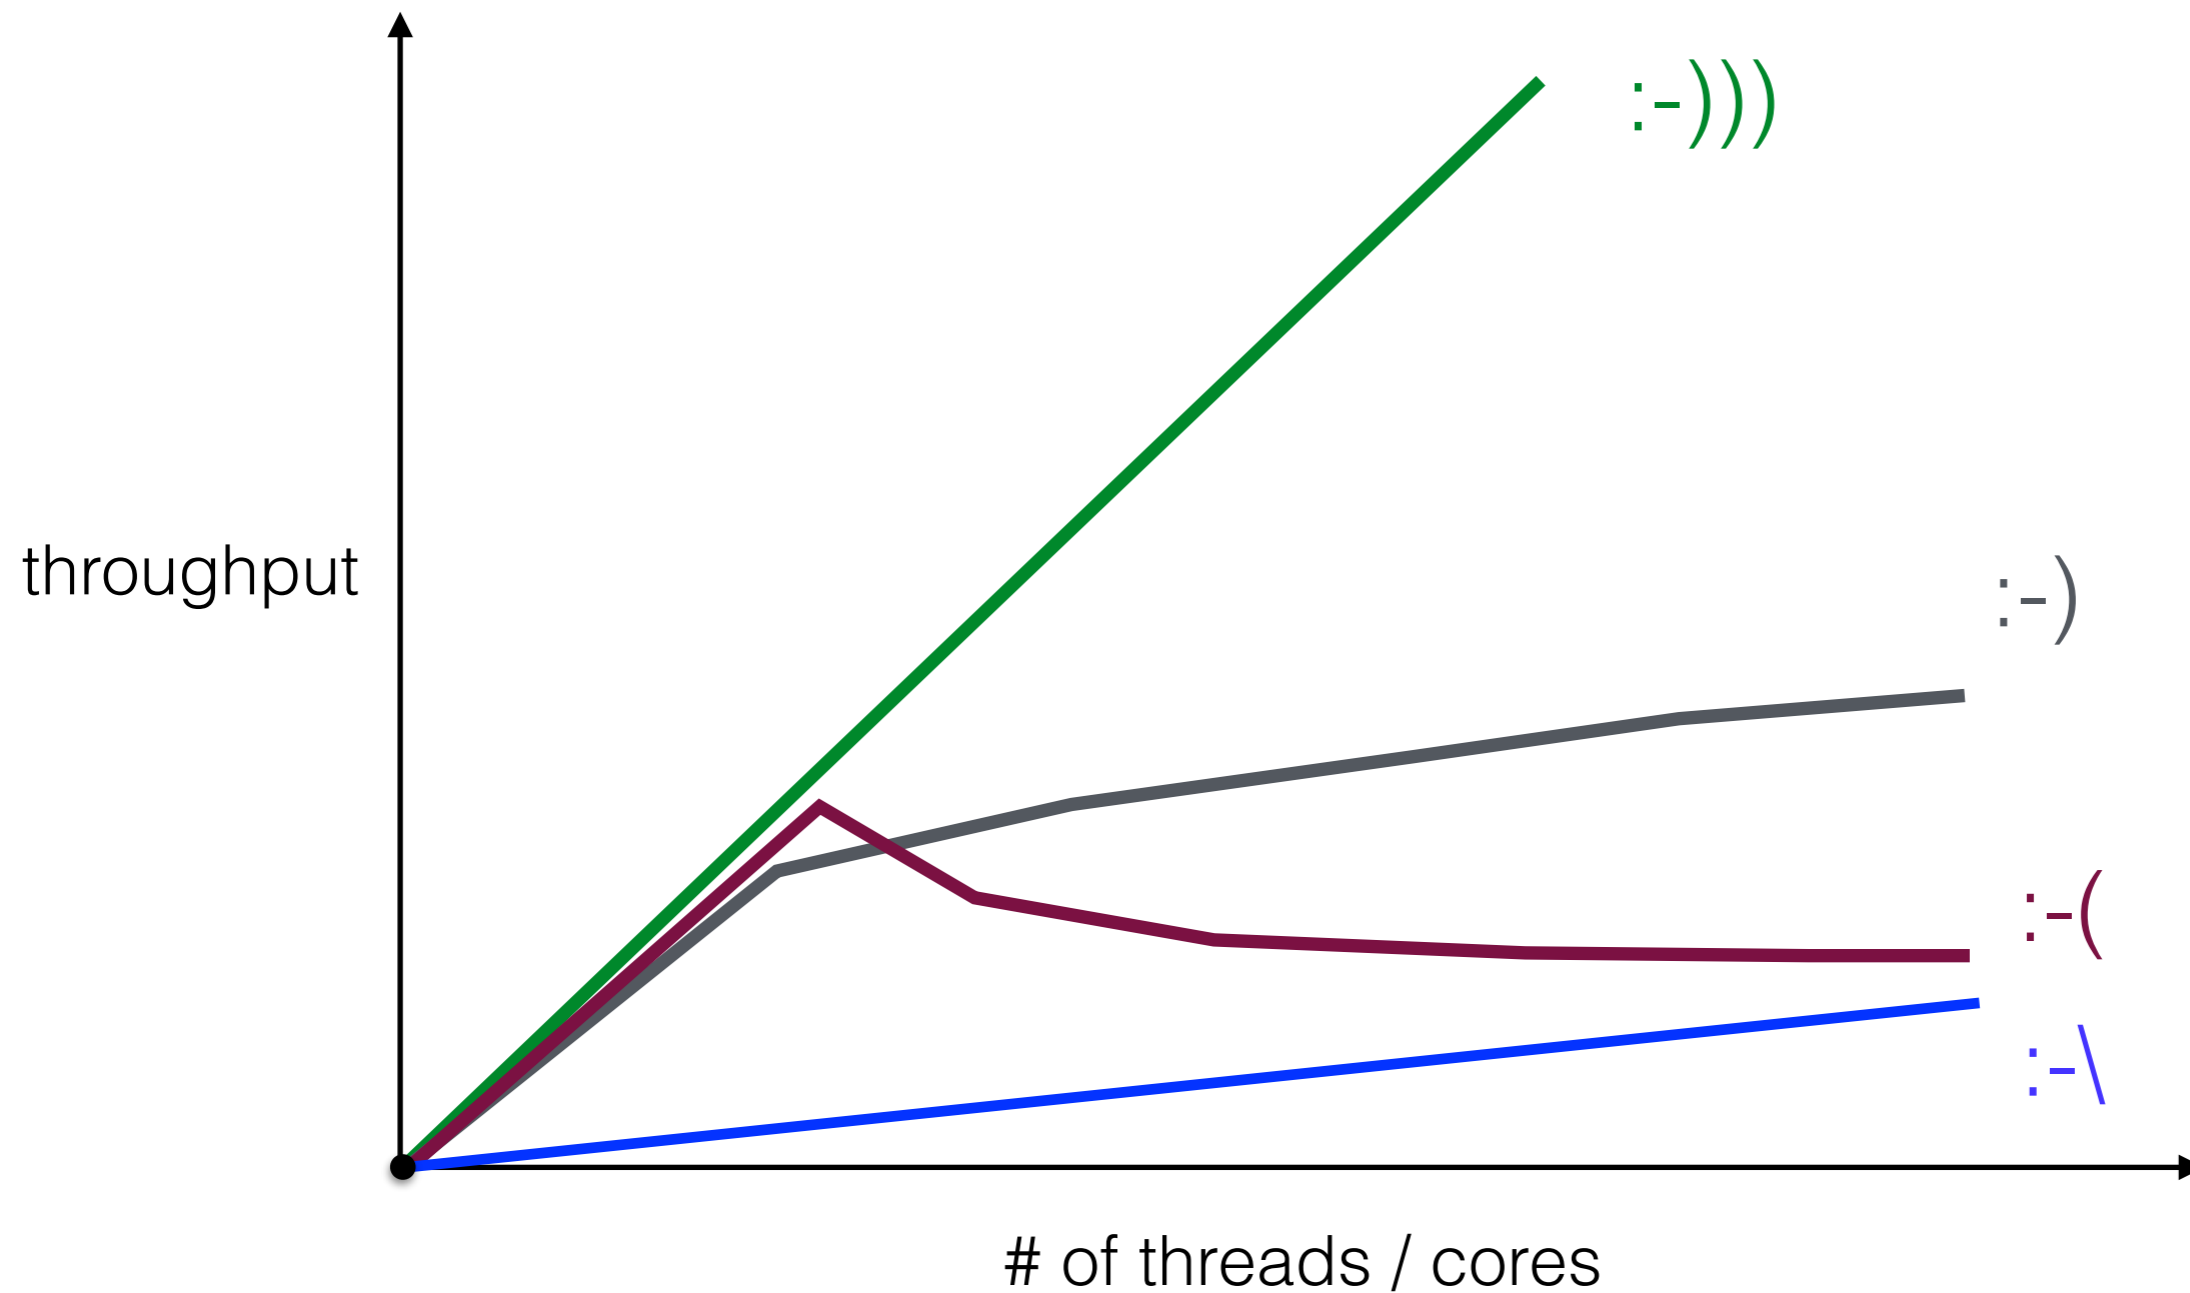


Sequential Consistency [Lamport'79]

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

Relaxations allow trading

correctness
for
performance

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

Relaxing the semantics

Relaxing the semantics

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

Relaxing the semantics

Quantitative relaxations
POPL13

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

Relaxing the semantics

not
“sequentially
correct”

Quantitative relaxations
POPL13

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

Relaxing the semantics

not
“sequentially
correct”

Quantitative relaxations
POPL13

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

Local linearizability
CONCUR16

Relaxing the semantics

not
“sequentially
correct”

Quantitative relaxations
POPL13

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

Local linearizability
CONCUR16

too weak

Relaxing the semantics

not
“sequentially
correct”

Quantitative relaxations
POPL13

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

for queues/stacks only
(feel free to ask for more)

Local linearizability
CONCUR16

too weak

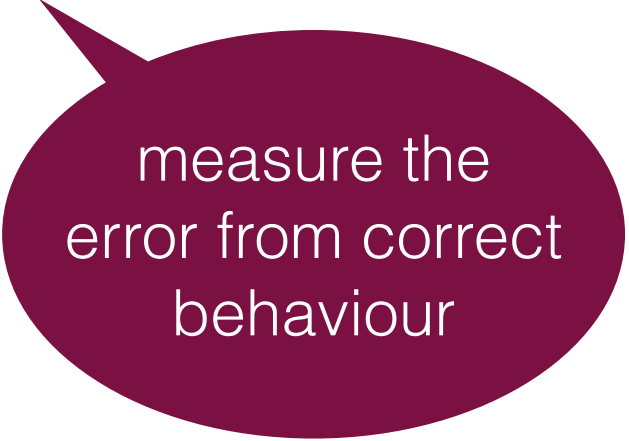
Relaxing the sequential specification

Relaxing the sequential specification

Quantitative
relaxations
(POPL13)

Goal

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



measure the
error from correct
behaviour

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

measure the
error from correct
behaviour

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

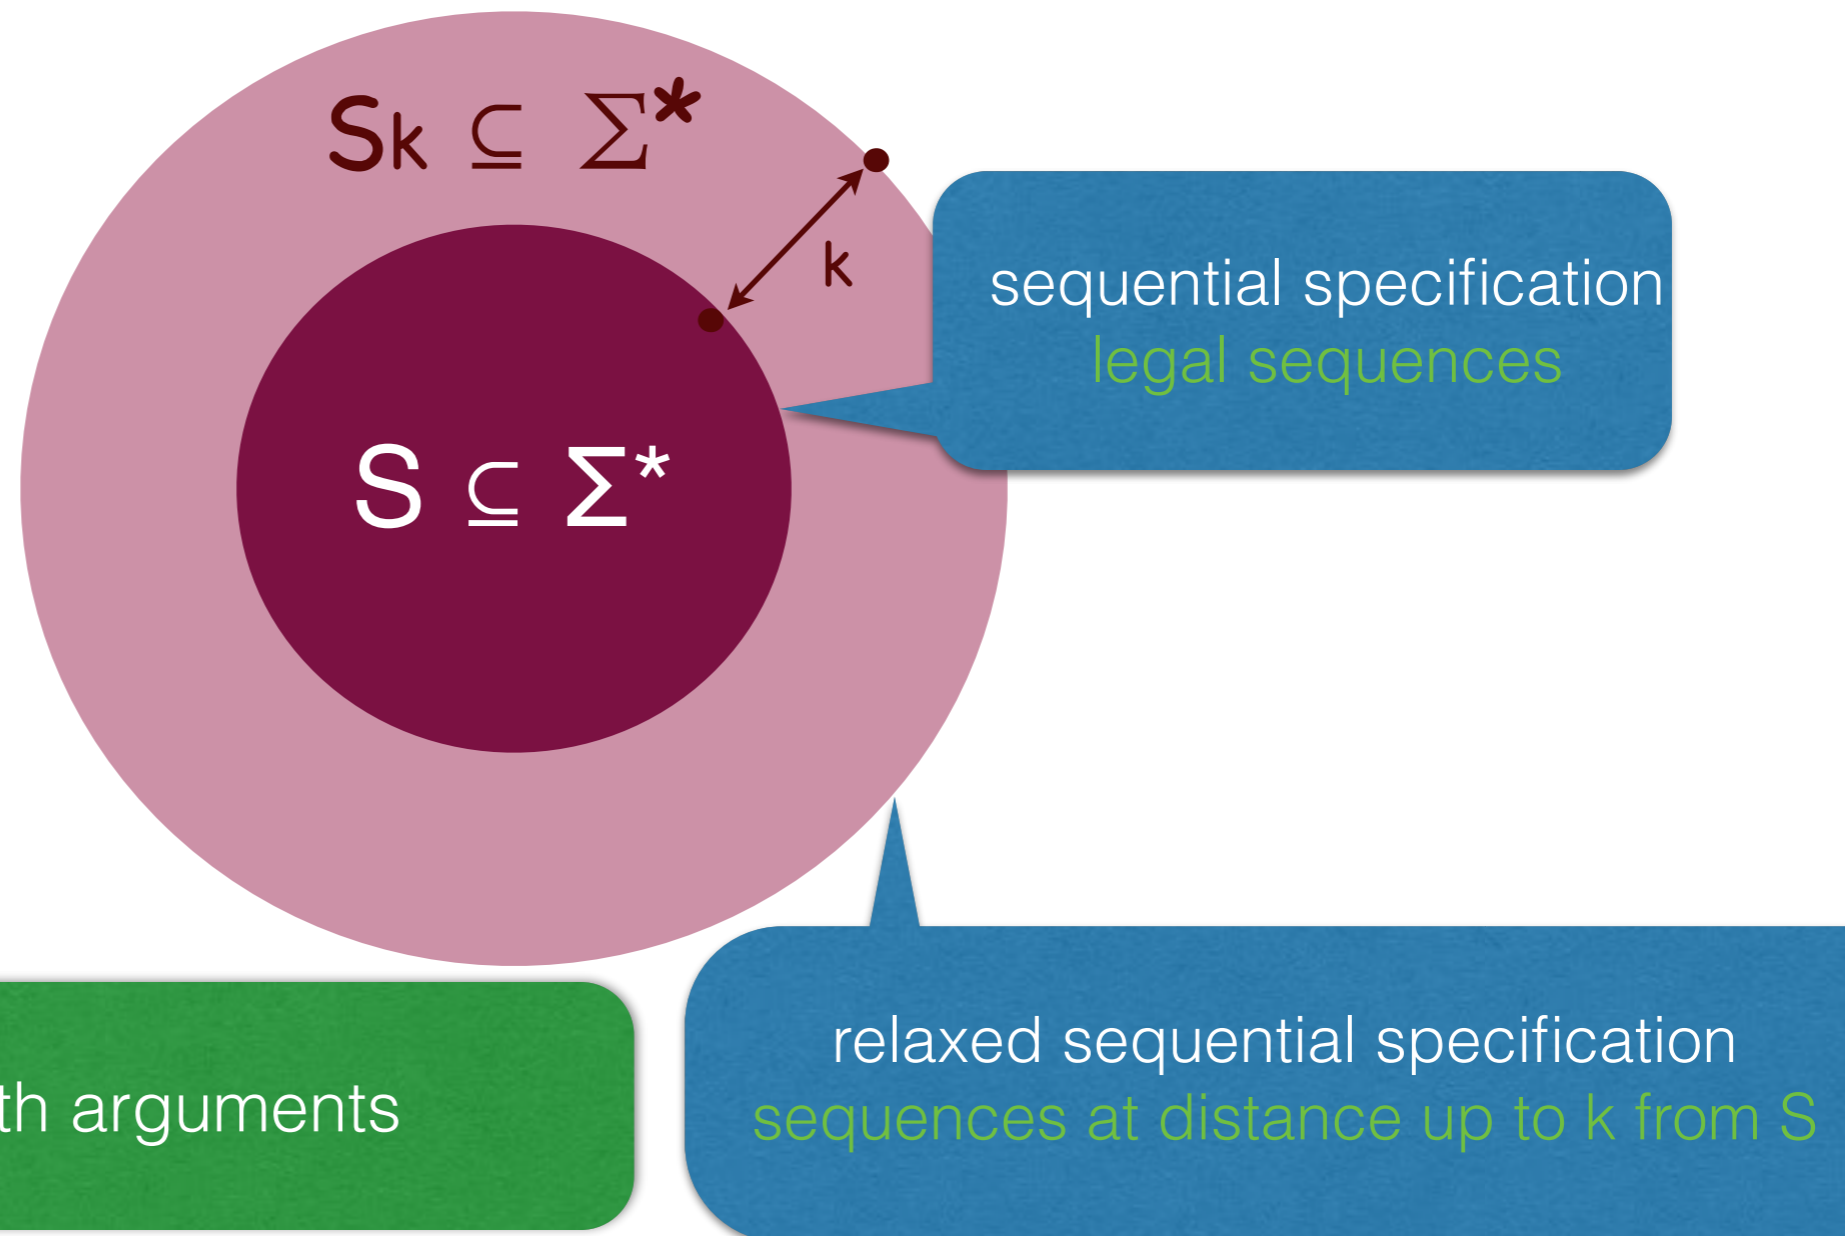
The big picture

$$S \subseteq \Sigma^*$$

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

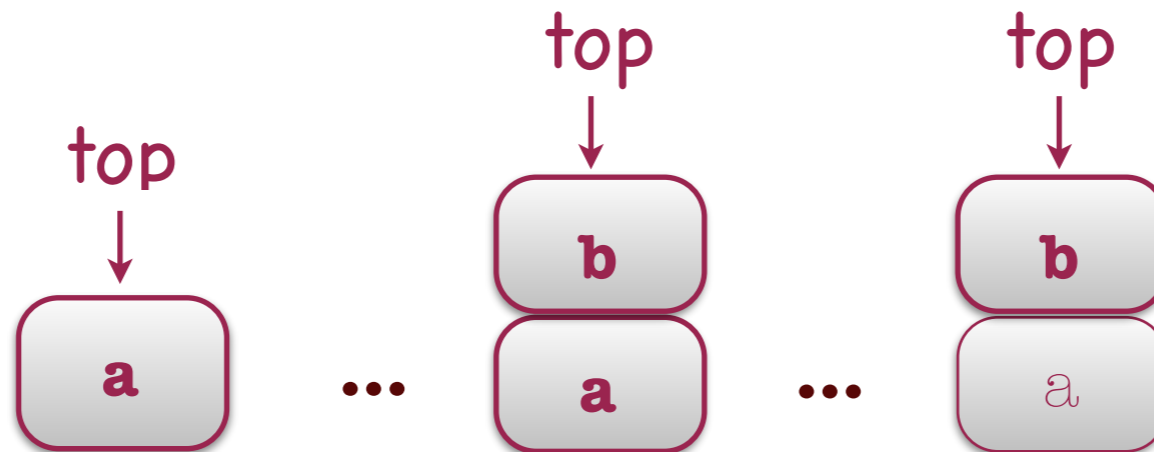
Syntactic distances do not help

$\text{push}(a)[\text{push}(i)\text{pop}(i)]^n\text{push}(b)[\text{push}(j)\text{pop}(j)]^m\text{pop}(a)$

Syntactic distances do not help

$\text{push}(a)[\text{push}(i)\text{pop}(i)]^n\text{push}(b)[\text{push}(j)\text{pop}(j)]^m\text{pop}(a)$


is a 1-out-of-order stack sequence



Syntactic distances do not help

$\text{push}(a)[\text{push}(i)\text{pop}(i)]^n\text{push}(b)[\text{push}(j)\text{pop}(j)]^m\text{pop}(a)$

is a 1-out-of-order stack sequence



its permutation distance is $\min(2n, 2m)$

Semantic distances need a notion of state

- States are equivalence classes of sequences in S
- Two sequences in S are equivalent iff they have an indistinguishable future

Semantic distances need a notion of state

- States are equivalence classes of sequences in S
- Two sequences in S are equivalent iff they have an indistinguishable future

$$x \equiv y \iff \forall u \in \Sigma^*. (xu \in S \iff yu \in S)$$

Semantic distances need a notion of state

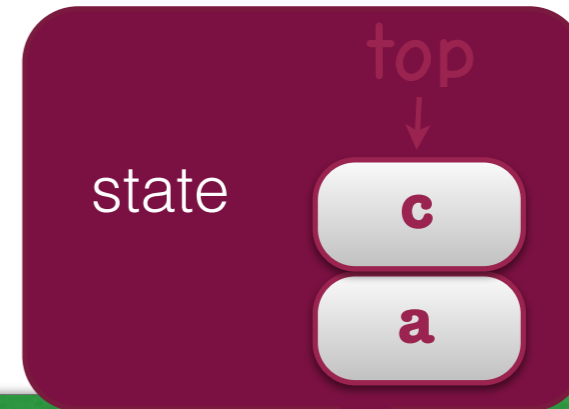
- States are equivalence classes of sequences in S

example: for stack

$\text{push}(a)\text{push}(b)\text{pop}(b)\text{push}(c) \equiv \text{push}(a)\text{push}(c)$

- Two sequences in S are equivalent iff they have an indistinguishable future

$$x \equiv y \iff \forall u \in \Sigma^*. (xu \in S \iff yu \in S)$$



Semantics goes operational

$S \subseteq \Sigma^*$ is the sequential specification

states

labels

initial state

$LTS(S) = (S/\equiv, \Sigma, \rightarrow, [\varepsilon]_{\equiv})$ with

transition relation

$$[s]_{\equiv} \xrightarrow{m} [sm]_{\equiv} \iff sm \in S$$

Semantics goes operational

$S \subseteq \Sigma^*$ is the sequential specification

states

labels

initial state

$LTS(S) = (S/\equiv, \Sigma, \rightarrow, [\varepsilon]_{\equiv})$ with

transition relation

$$[s]_{\equiv} \xrightarrow{m} [sm]_{\equiv} \iff sm \in S$$

Stack

top

a

push(c)

top

c


a

The relaxation framework

- Start from LTS(S)
- Add transitions with transition costs
- Fix a path cost function


The relaxation framework

- Start from LTS(S)
- Add transitions with transition costs
- Fix a path cost function




The relaxation framework

- Start from LTS(S)
- Add transitions with transition costs
- Fix a path cost function




The relaxation framework

- Start from LTS(S)
- Add transitions with transition costs
- Fix a path cost function



The relaxation framework

- Start from LTS(S)
- Add transitions with transition costs
- Fix a path cost function



distance - minimal cost on all paths labelled by the sequence

Generic out-of-order

$$\text{segment_cost}(q \xrightarrow{m} q') = |\mathbf{v}|$$

transition cost

Where \mathbf{v} is a sequence of minimal length s.t.

removing \mathbf{v} enables a transition

or

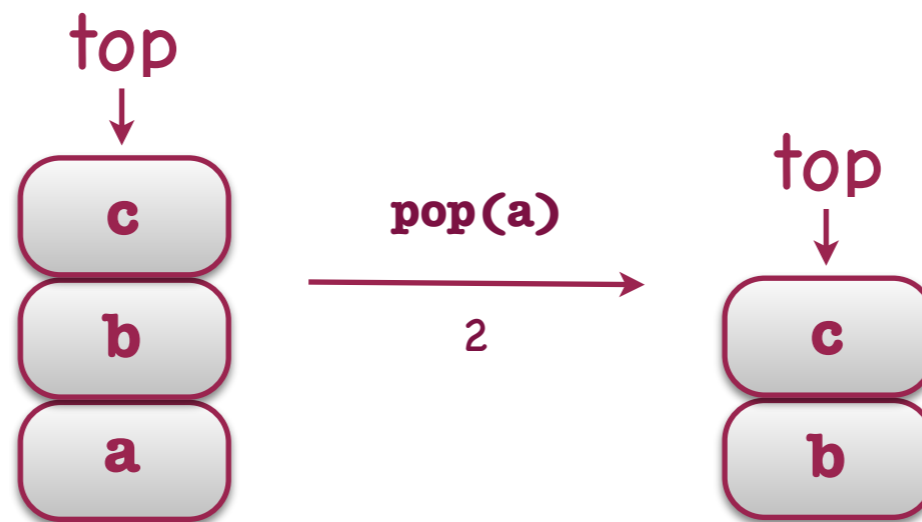
inserting \mathbf{v} enables a transition

goes with different path costs

Out-of-order stack

Sequence of push's with no matching pop

- Canonical representative of a state
- Add incorrect transitions with **segment-costs**



- Possible path cost functions **max**, **sum**,...

also more advanced

Relaxing the Consistency Condition

Relaxing the Consistency Condition

Local Linearizability
(CONCUR16)

Local Linearizability

main idea

Local Linearizability

main idea

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

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**

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**

Local sequential consistency... is also possible

Local Linearizability

main idea

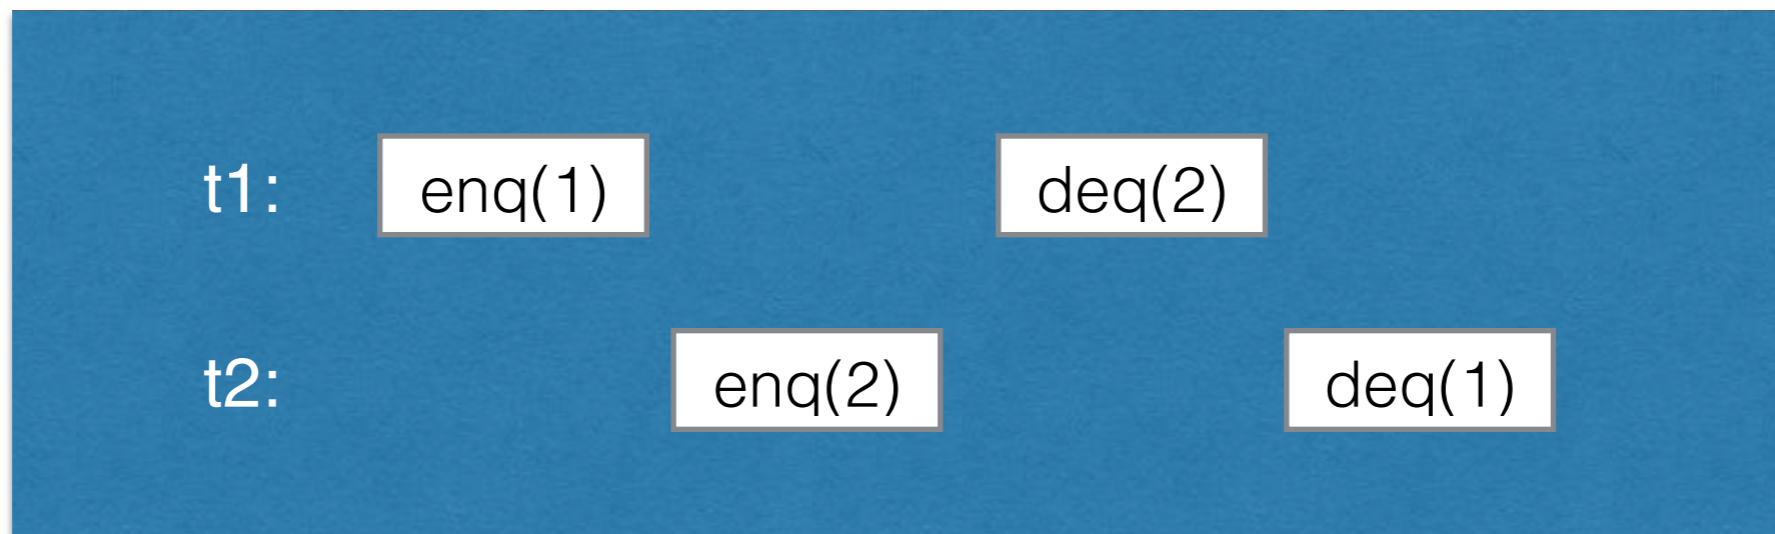
Already present in some shared-memory consistency conditions
(not in our form of choice)

- **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) example

(sequential) history
not linearizable

t1: enq(1)


deq(2)

t2: enq(2)

deq(1)


Local Linearizability (queue) example

(sequential) history
not linearizable



Local Linearizability (queue) example


(sequential) history
not linearizable



t1-induced history,
linearizable

Local Linearizability (queue) example


(sequential) history
not linearizable



t1-induced history,
linearizable

Local Linearizability (queue) example

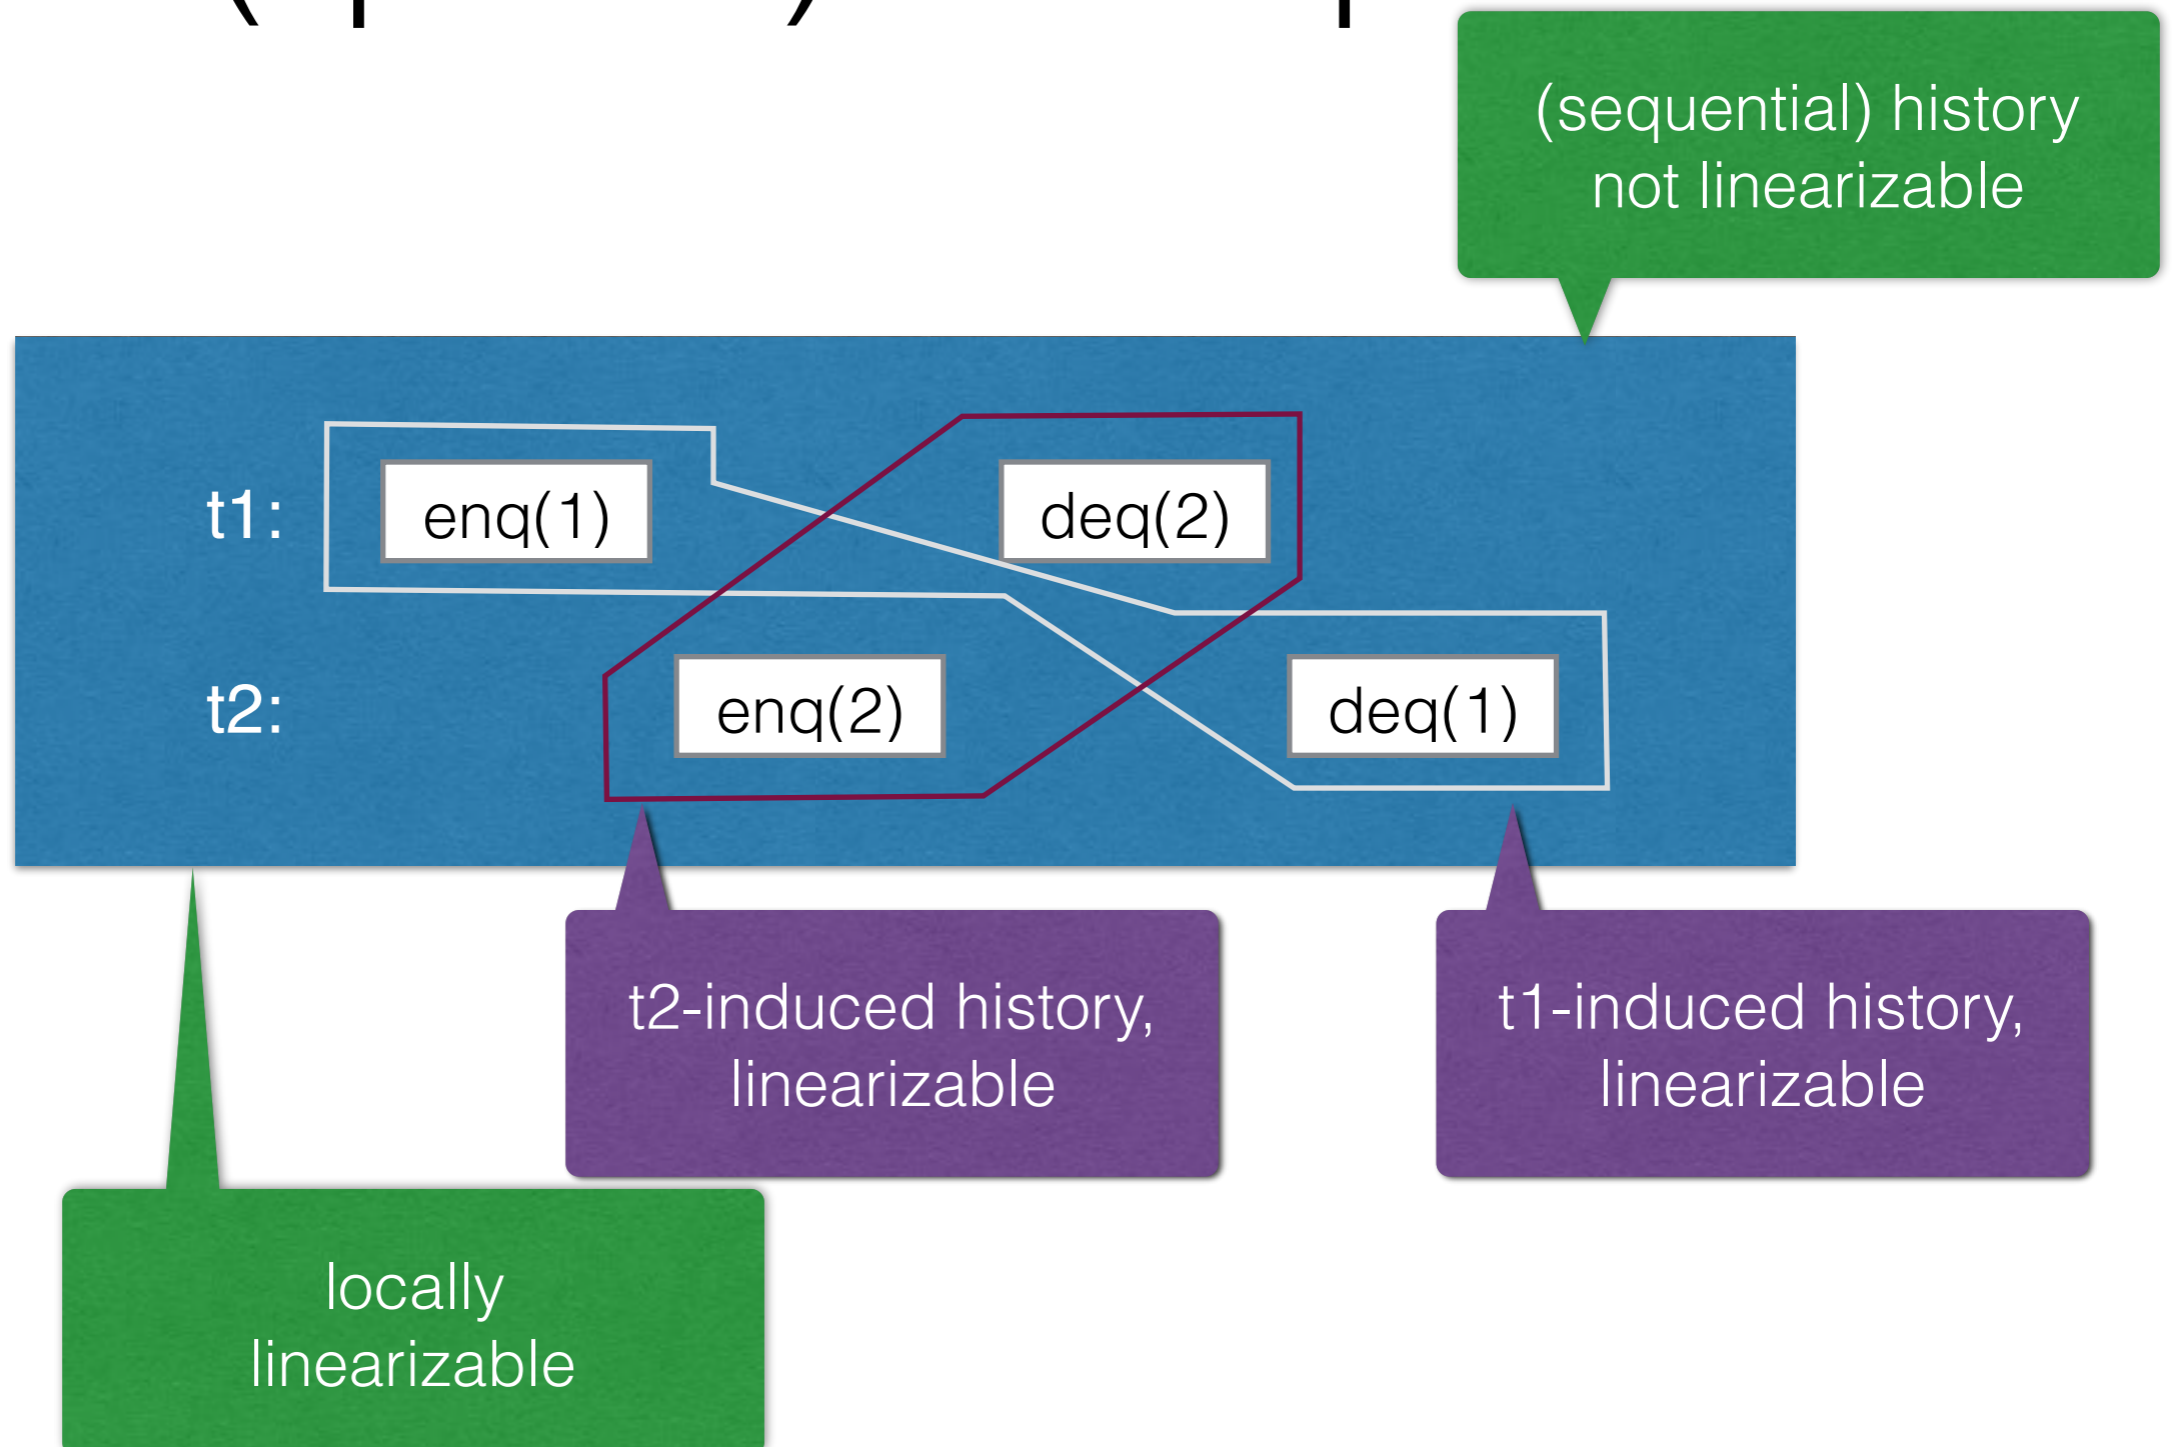
(sequential) history
not linearizable



t2-induced history,
linearizable

t1-induced history,
linearizable

Local Linearizability (queue) example



Local Linearizability (queue) definition

Local Linearizability (queue) definition

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

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

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

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

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

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

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?

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

Properties

Local linearizability is compositional

Properties

Local linearizability is compositional

like linearizability
unlike 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

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”

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

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

Generic Implementations

Generic Implementations

Your favorite linearizable data structure implementation

Generic Implementations

Your favorite linearizable data structure implementation

Φ

Generic Implementations

Your favorite linearizable data structure implementation

Φ


turns into a locally linearizable implementation by:

Generic Implementations

Your favorite linearizable data structure implementation

Φ

turns into a locally linearizable implementation by:




Generic Implementations

Your favorite linearizable data structure implementation

Φ

turns into a locally linearizable implementation by:




segment of possibly dynamic size (n)

Generic Implementations

Your favorite linearizable data structure implementation

Φ

turns into a locally linearizable implementation by:



segment of possibly dynamic size (n)

local inserts / global (randomly distributed) removes


Generic Implementations

Your favorite linearizable data structure implementation

Φ

turns into a locally linearizable implementation by:

LLD Φ
(locally
linearizable)



segment of possibly
dynamic size (n)

local inserts / global (randomly distributed) removes

Generic Implementations

Your favorite linearizable data structure implementation

Φ

turns into a locally linearizable implementation by:

LLD Φ
(locally linearizable)

t_1 t_2 ... t_n


segment of possibly dynamic size (n)

Φ Φ Φ

LL+D Φ
(also pool linearizable)

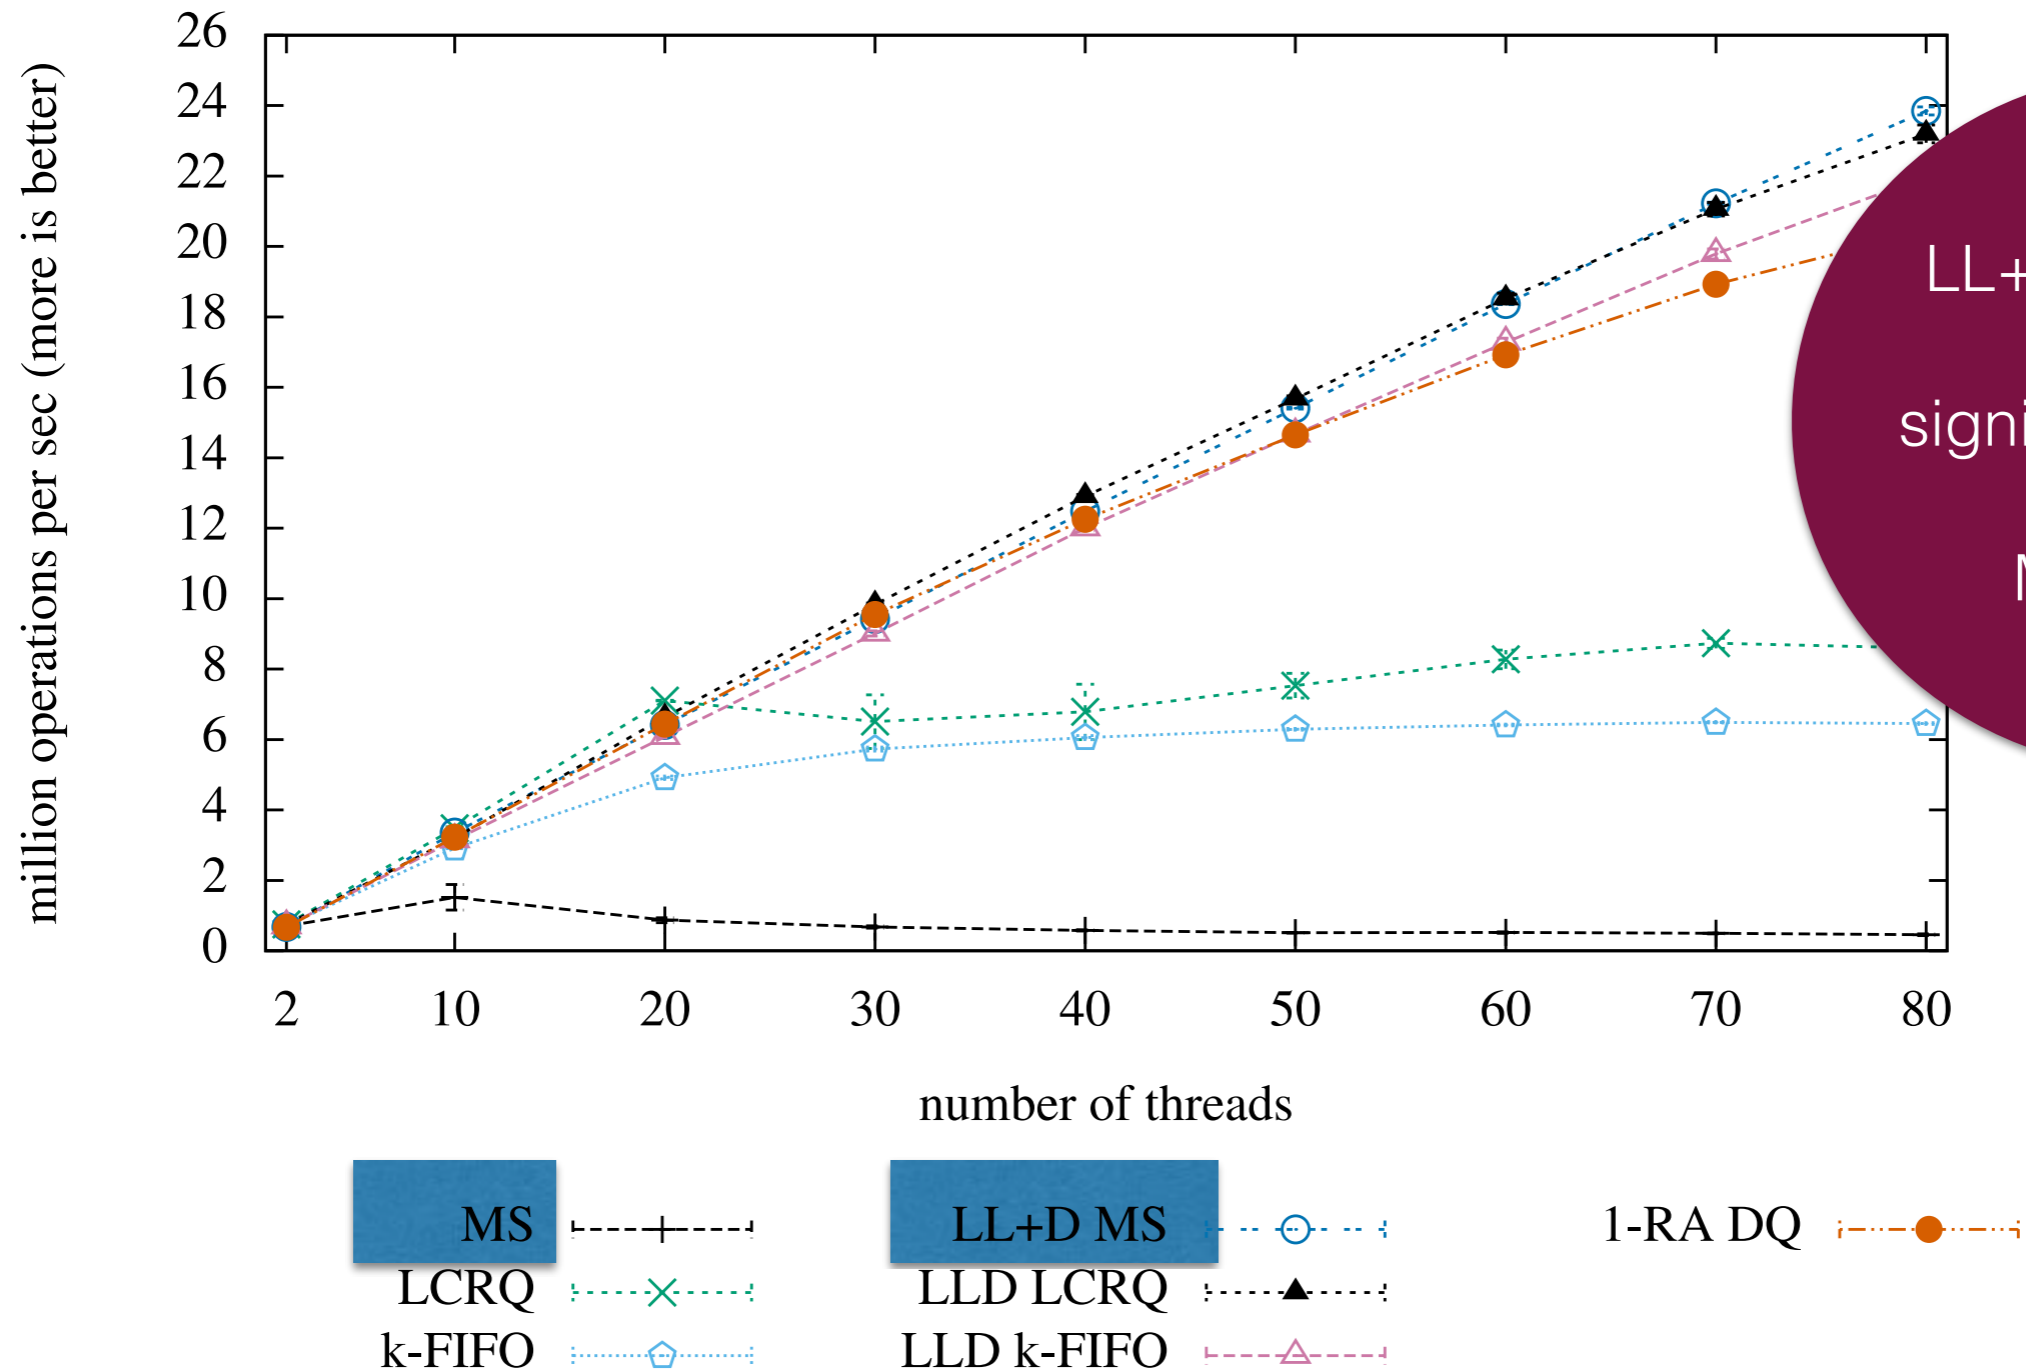
local inserts / global (randomly distributed) removes

Performance



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


Performance



LL+D MS queue performs significantly better than MS queue

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

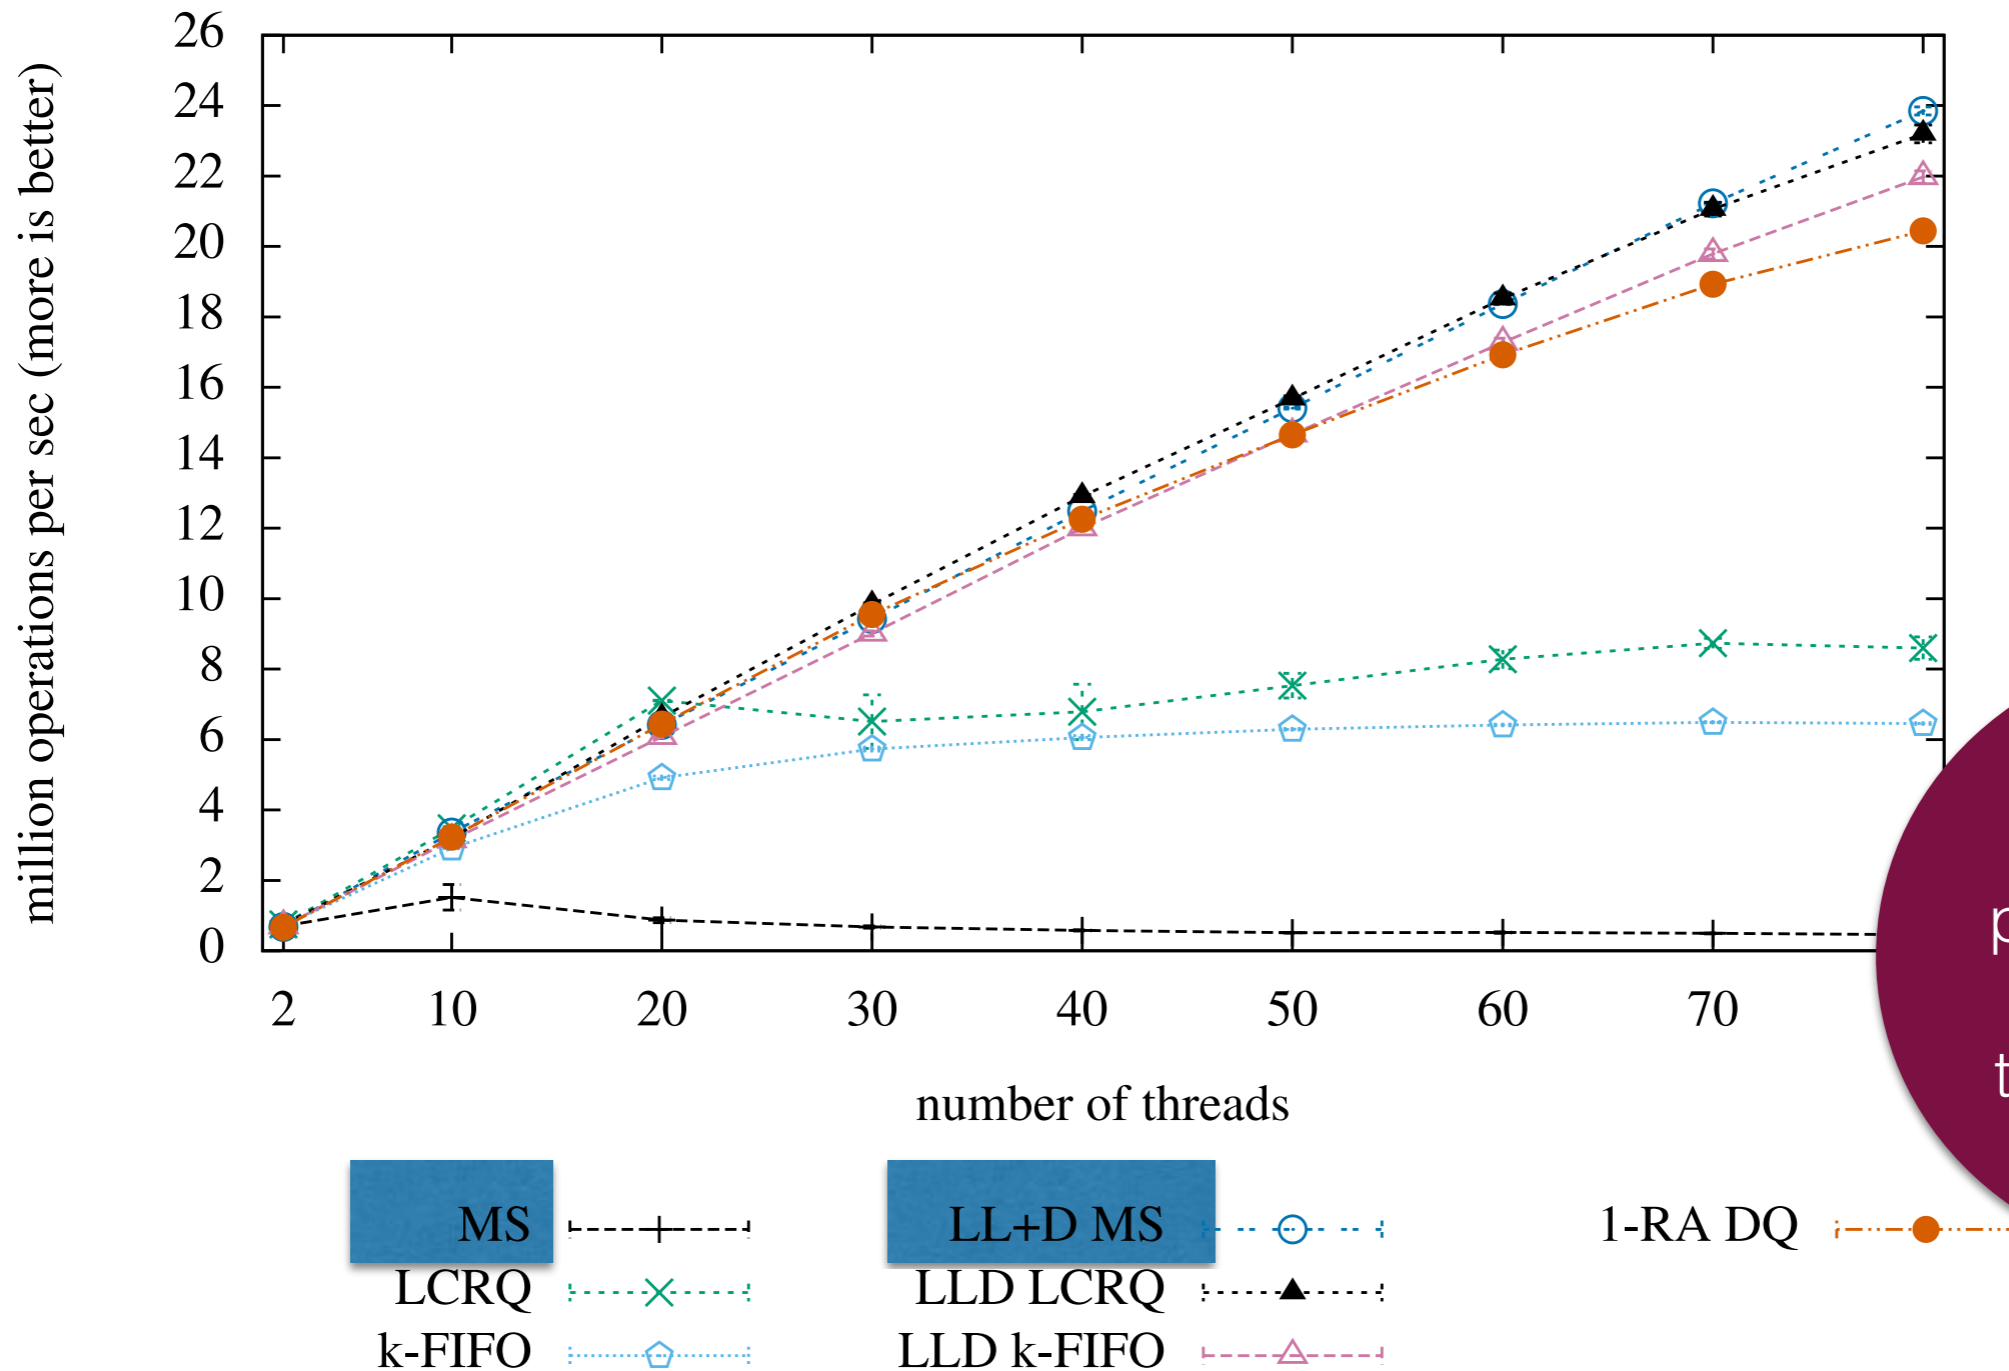
Performance



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

Thank You!

and many thanks to
my dear coauthors





Ali Sezgin 



Hannes Payer




Andreas Holzer 





Michael Lippautz




Tom Henzinger




Helmut Veith 



Andreas Haas 



Christoph Kirsch
