

Local Linearizability

Ana Sokolova  UNIVERSITY
of SALZBURG

joint work with:

Andreas Haas 

Andreas Holzer  UNIVERSITY OF
TORONTO

Michael Lippautz 

Ali Sezgin  UNIVERSITY OF
CAMBRIDGE

Tom Henzinger  IST AUSTRIA

Christoph Kirsch  UNIVERSITY
of SALZBURG

Hannes Payer 

Helmut Veith  TU
WIEN

Rigorous methods
for
engineering of
and
reasoning about
reactive systems

Rigorous methods
for
engineering of
and
reasoning about
reactive systems

concurrent

Background big picture

Background big picture

Computer Science

Background big picture

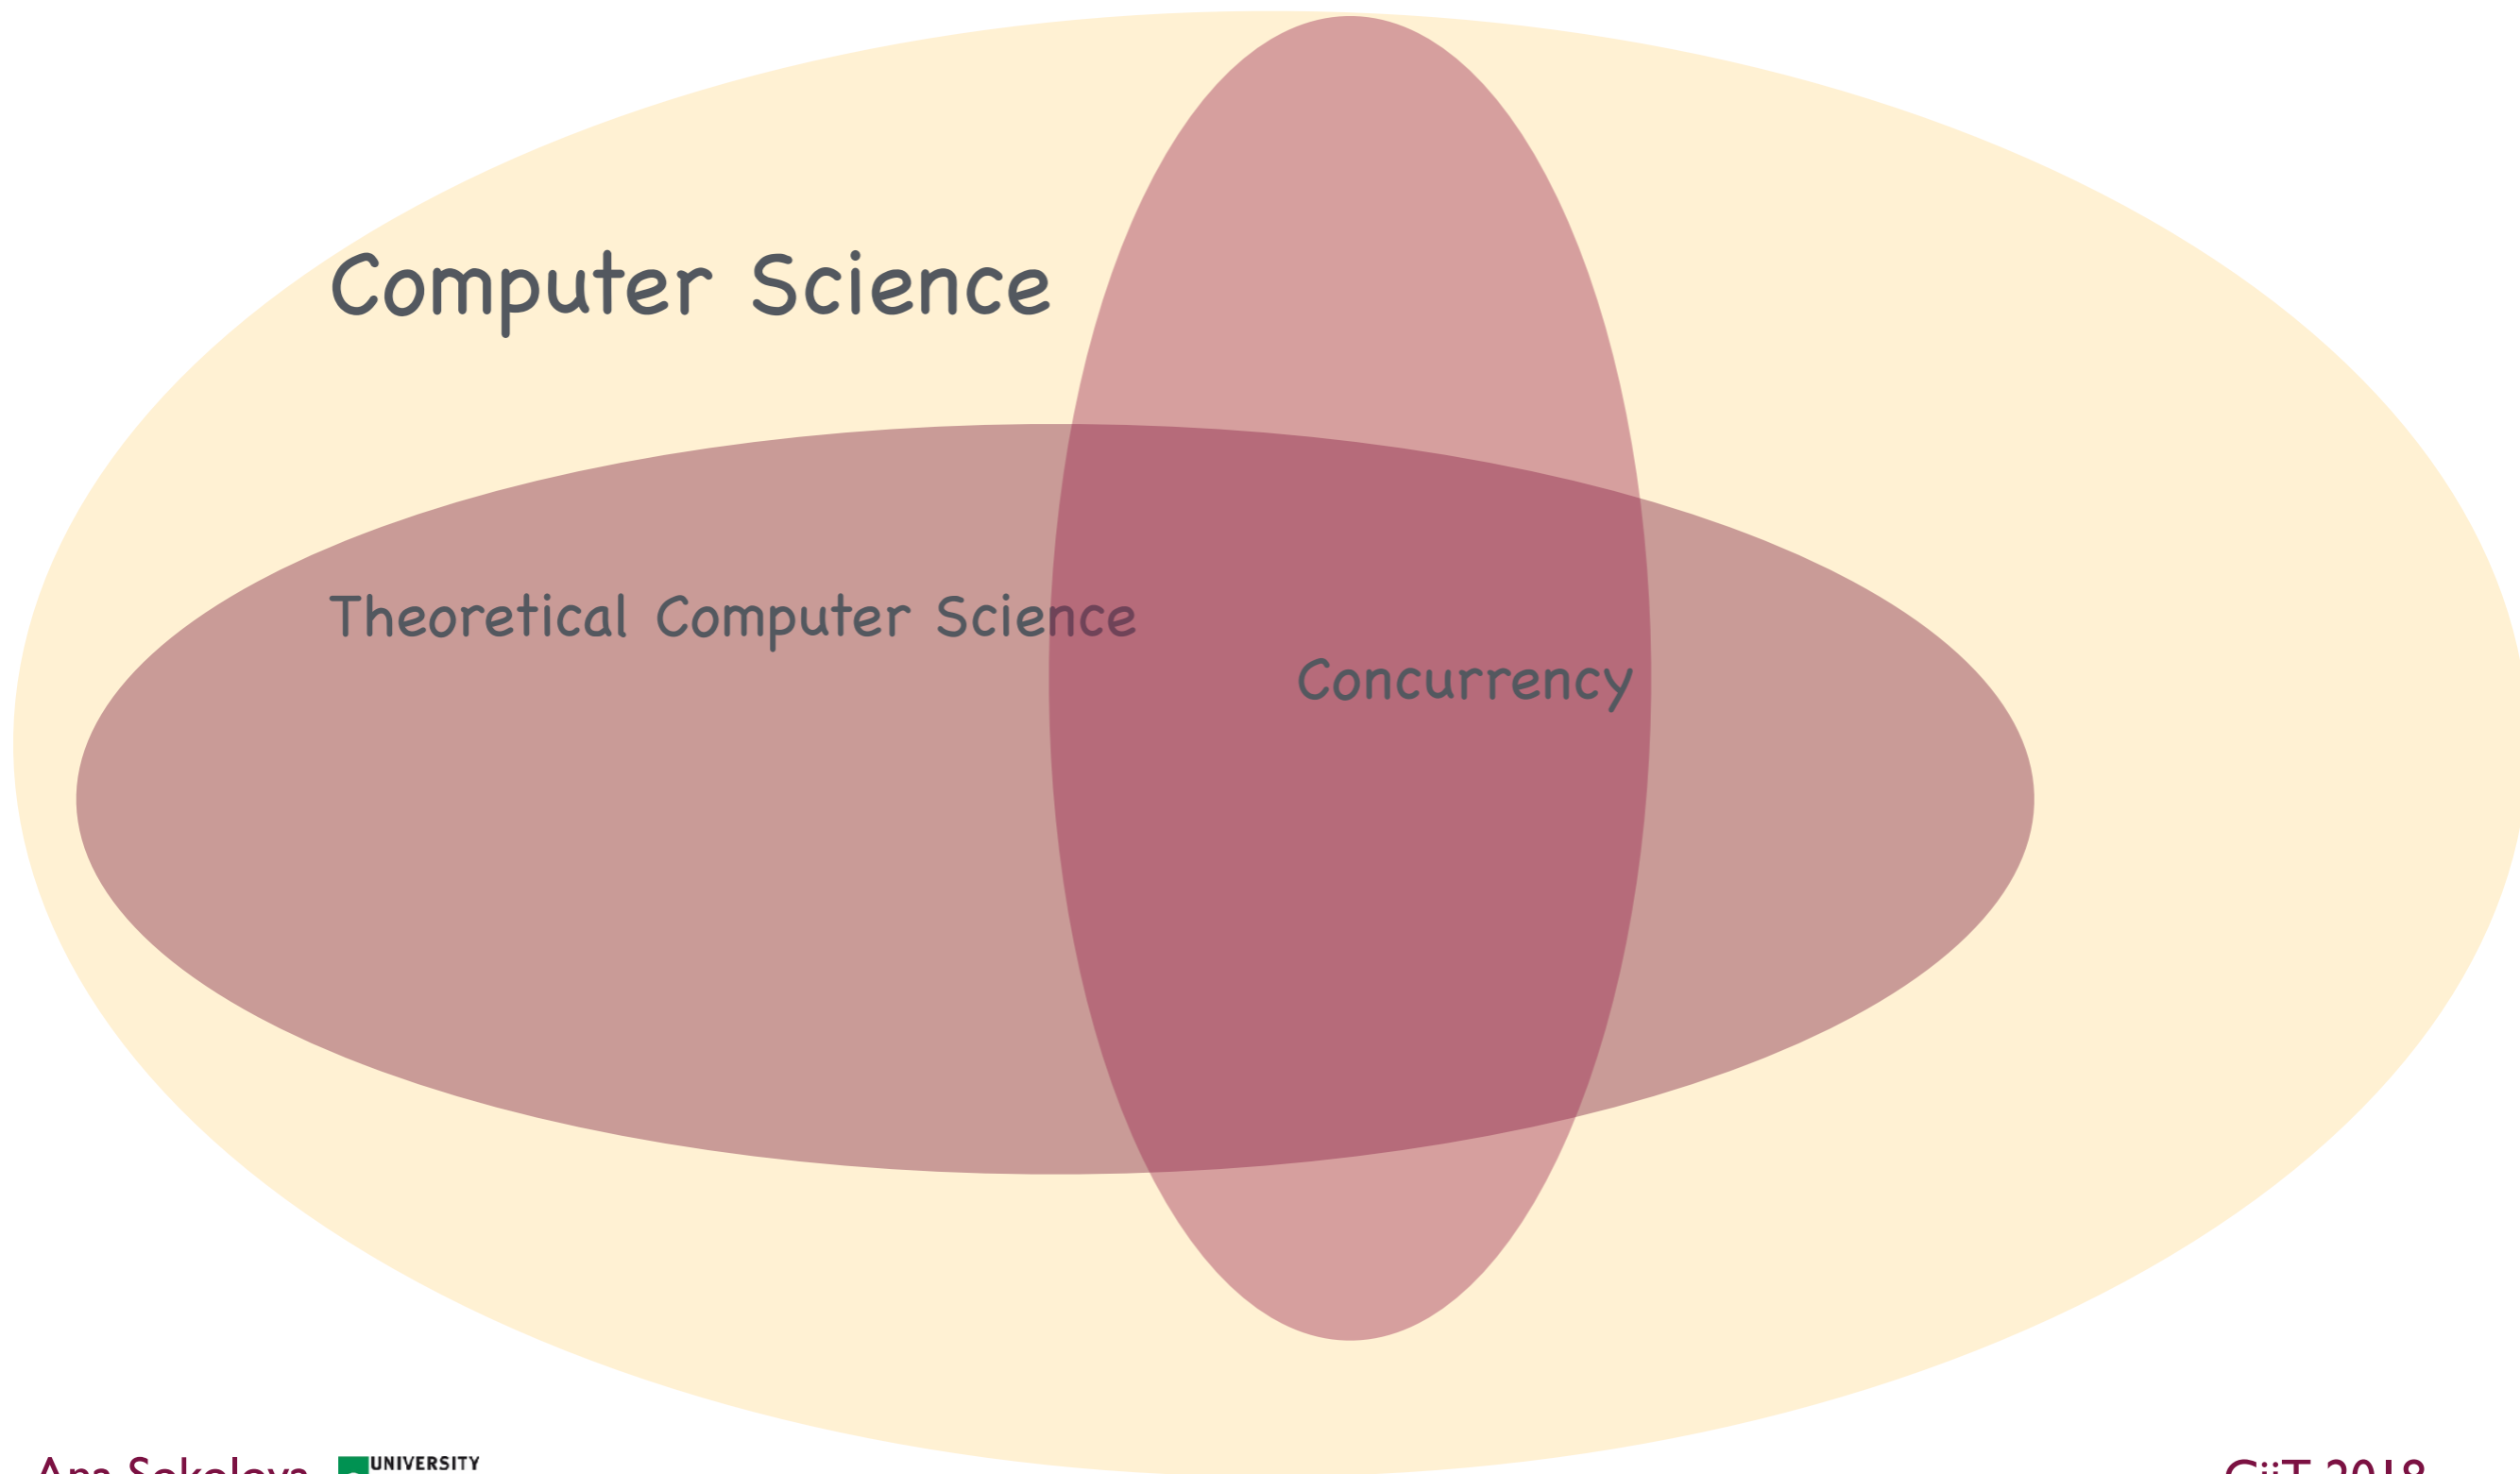


Computer Science

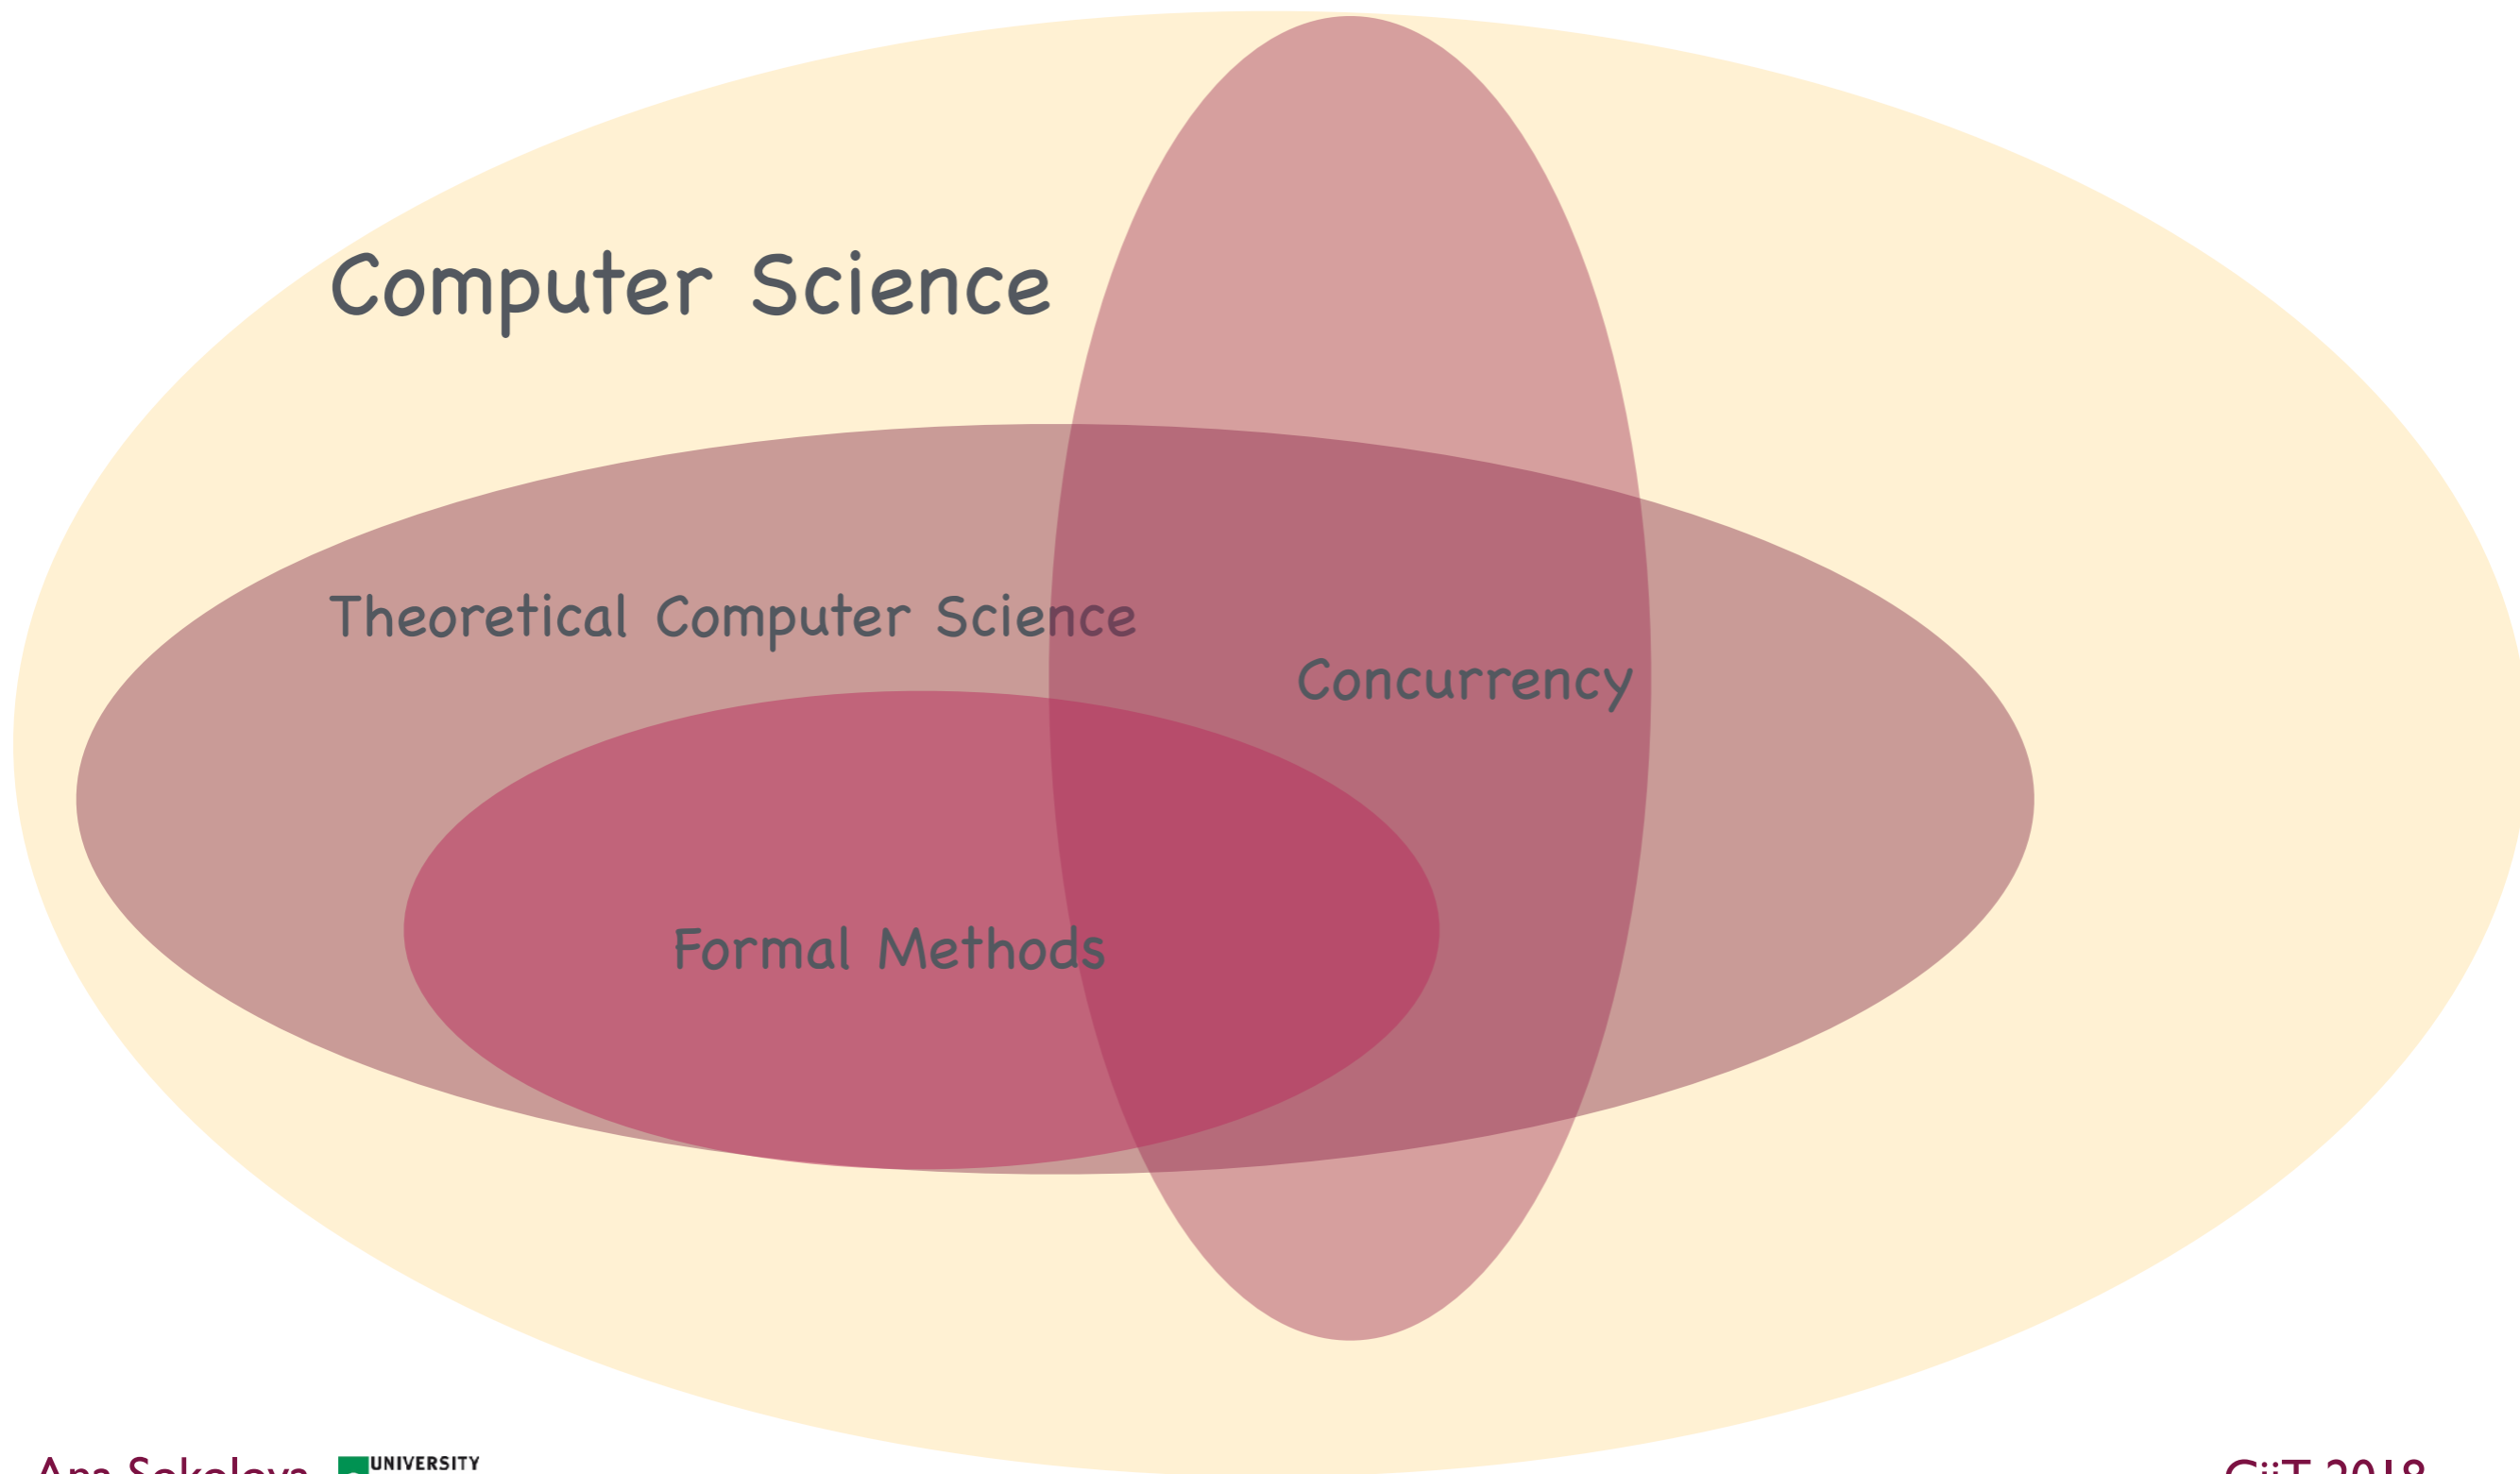
The diagram consists of two nested, horizontally-oriented ovals. The outer oval is light yellow and contains the text 'Computer Science'. The inner oval is a darker, reddish-brown color and contains the text 'Theoretical Computer Science'. The inner oval is centered within the outer oval, illustrating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

Theoretical Computer Science

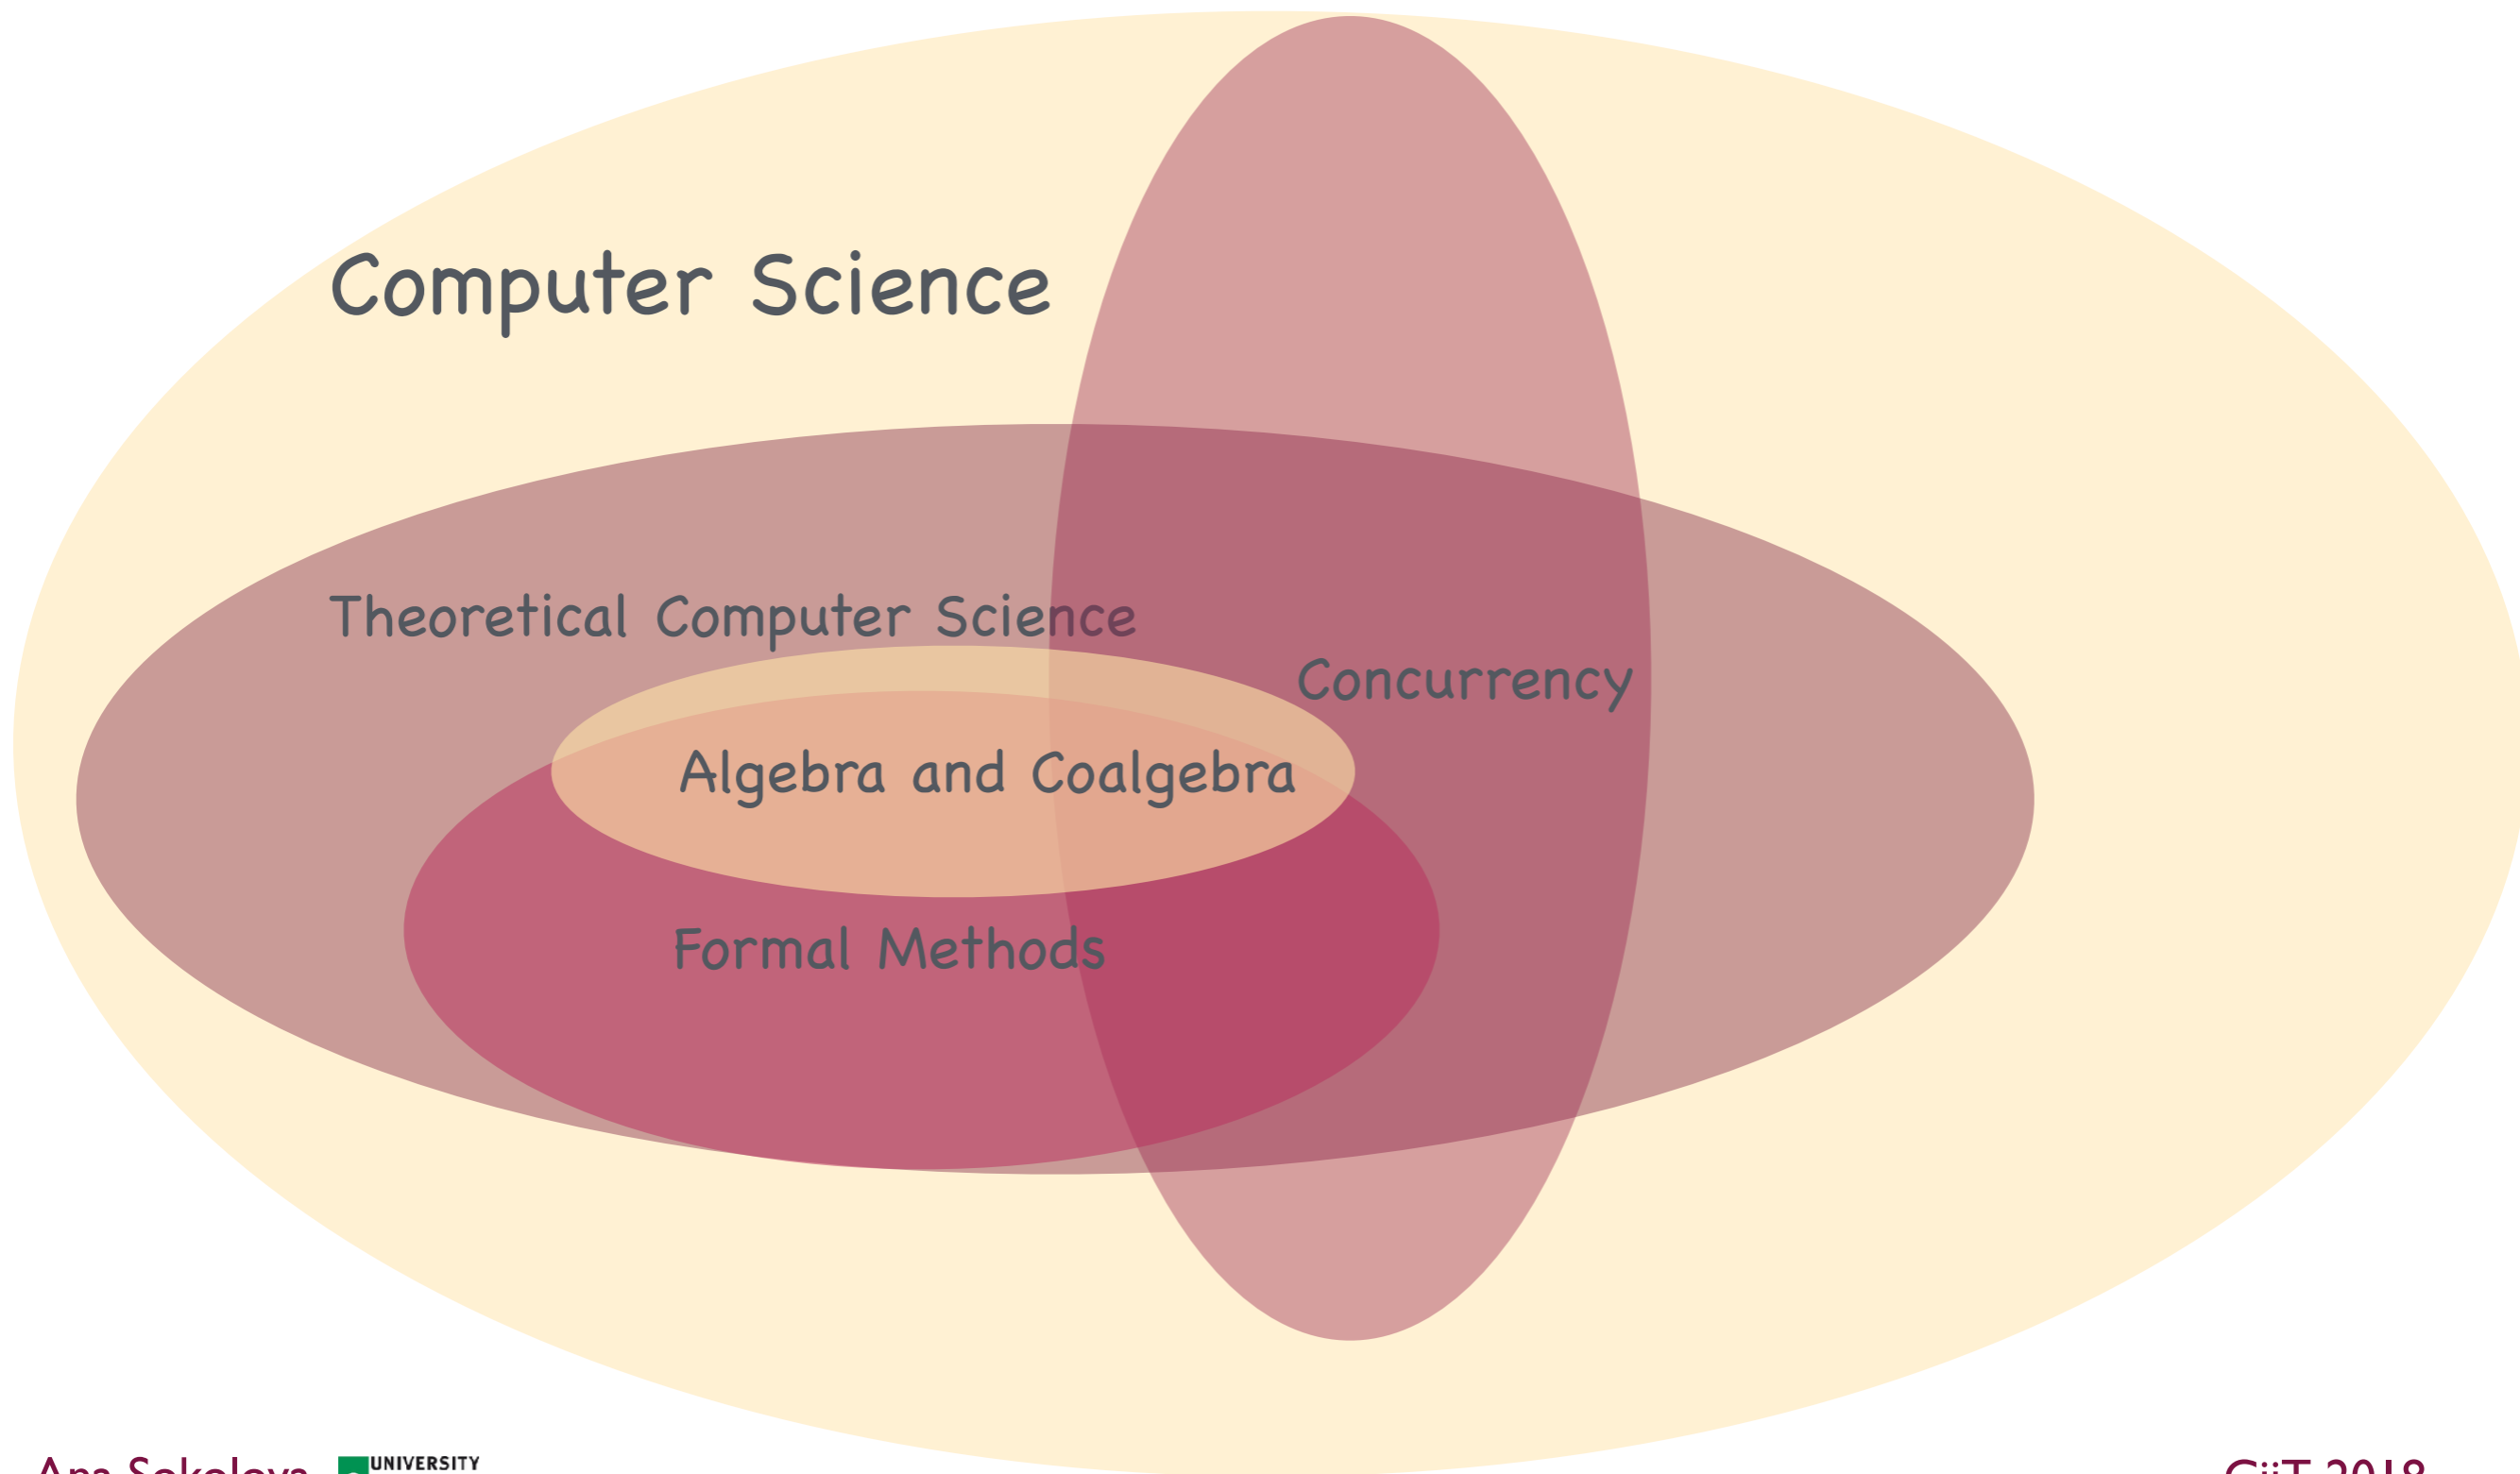
Background big picture



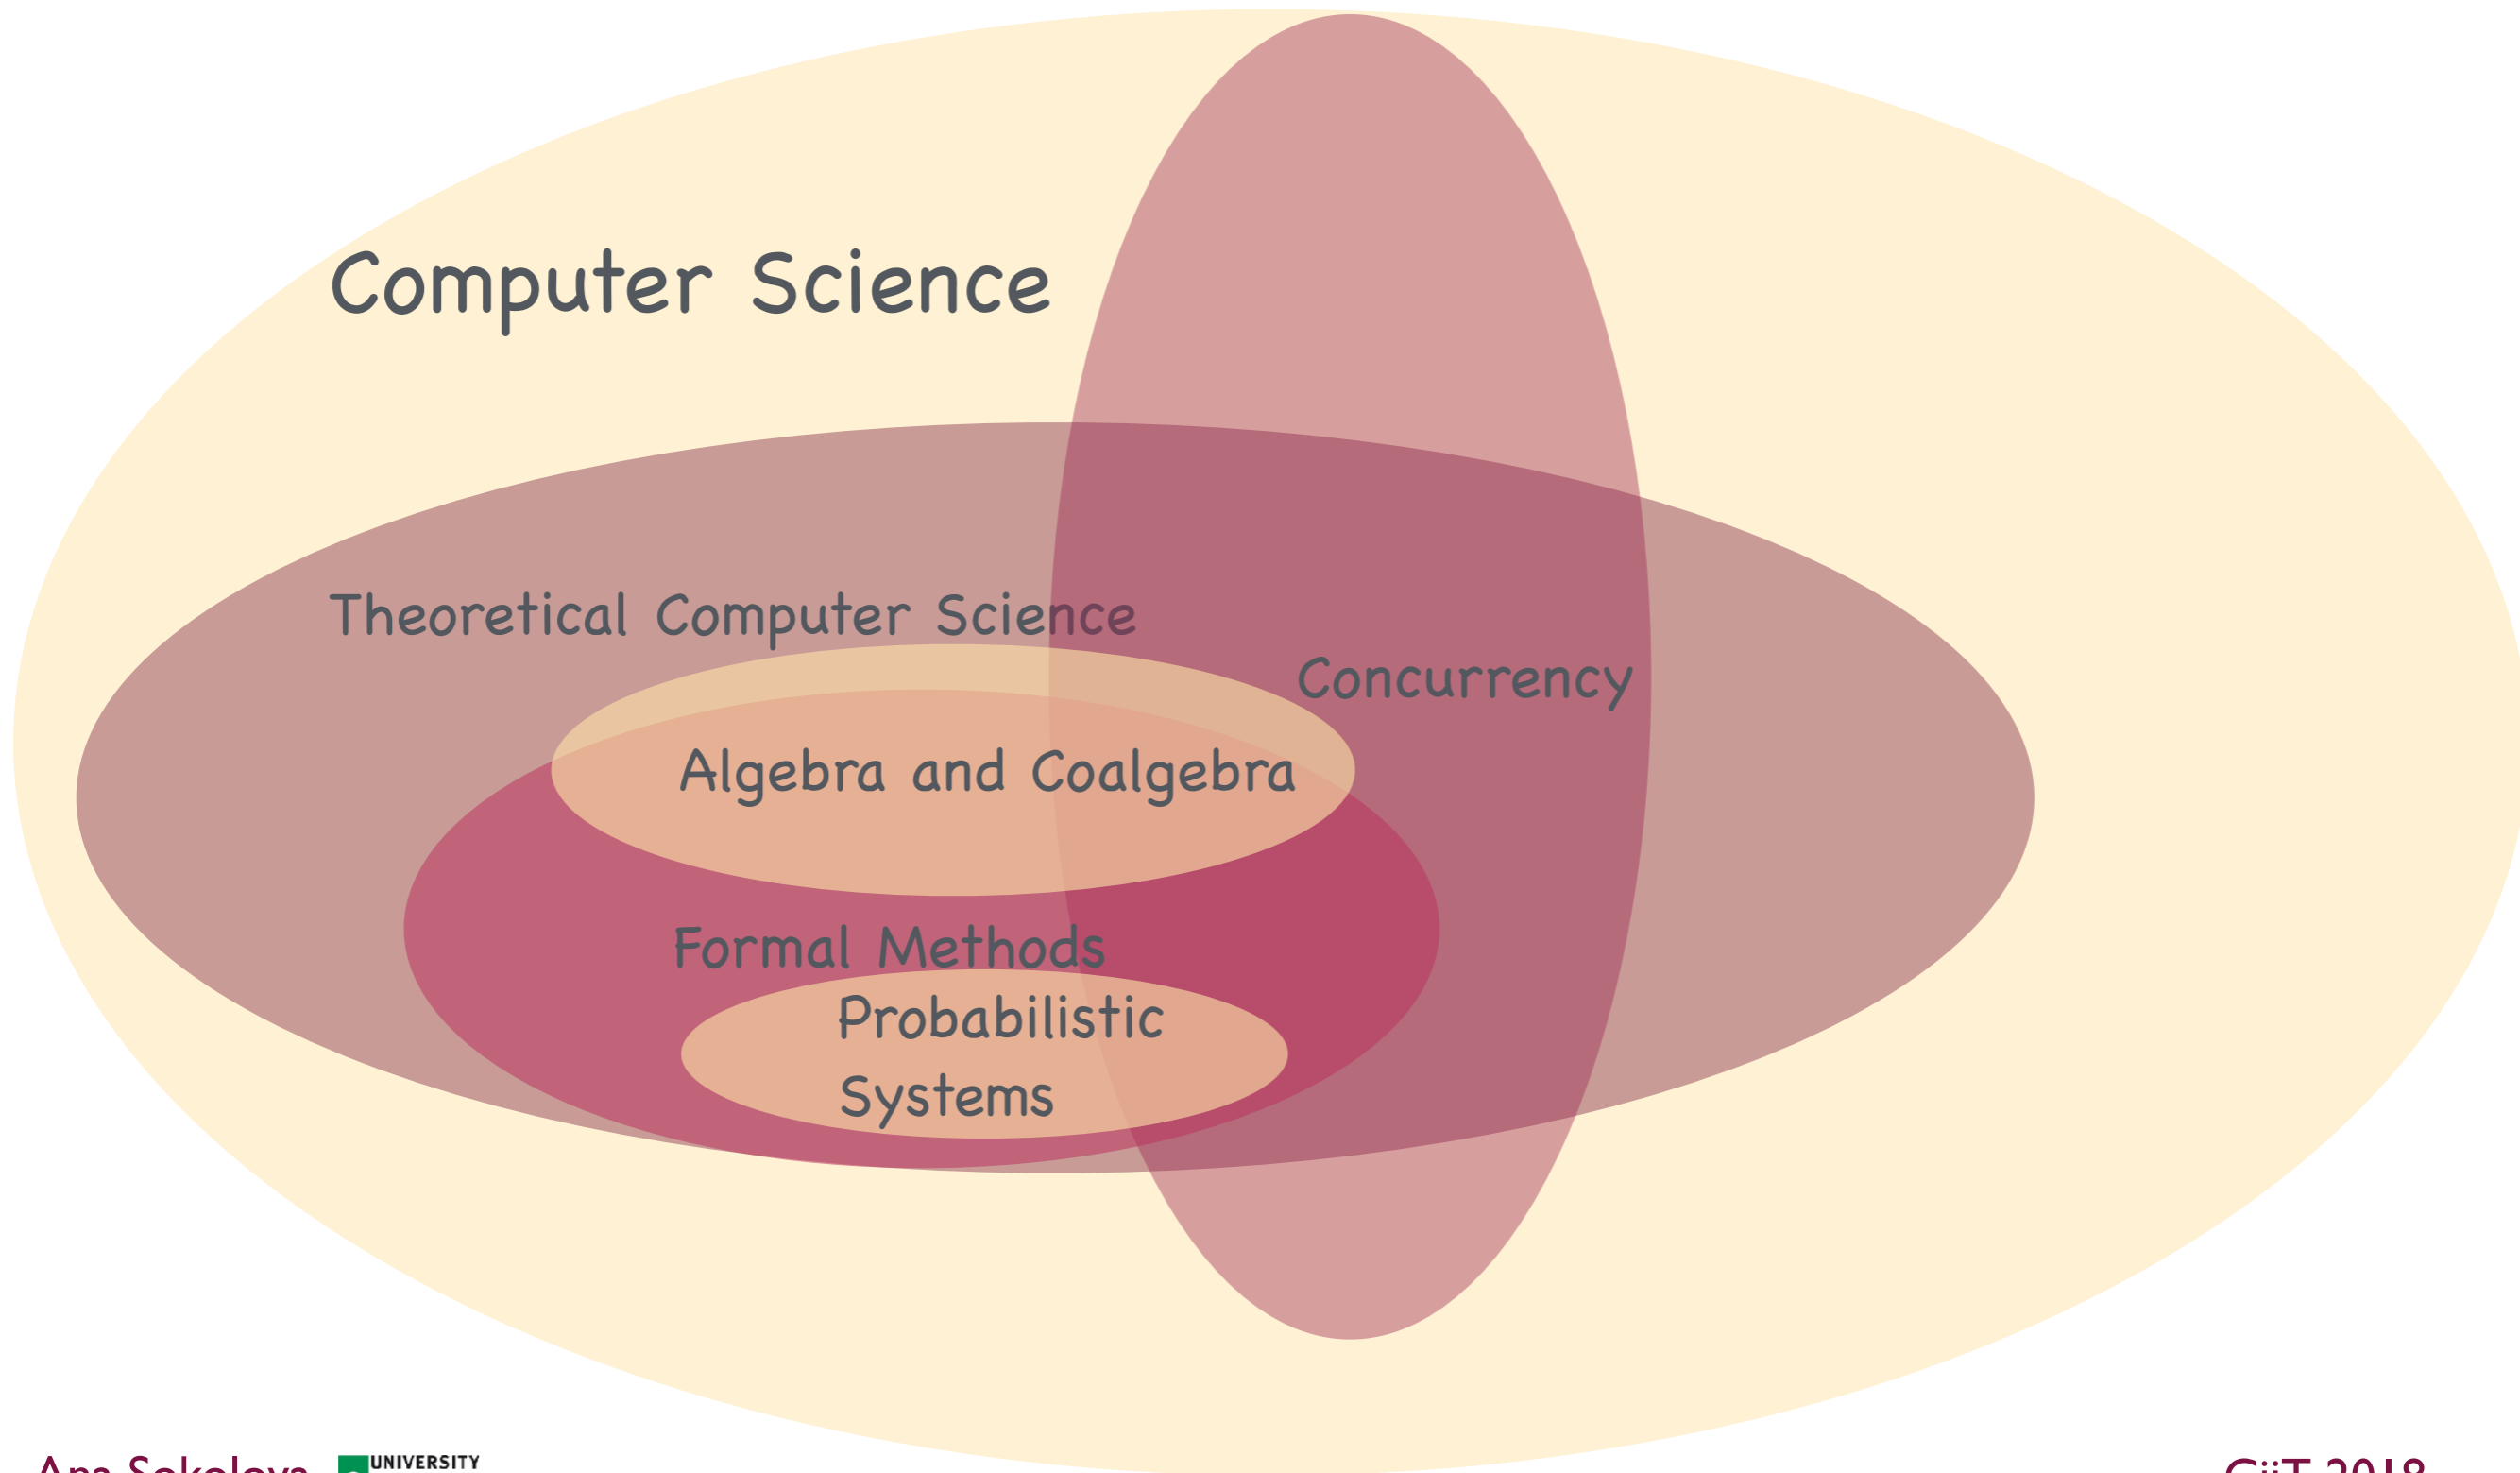
Background big picture



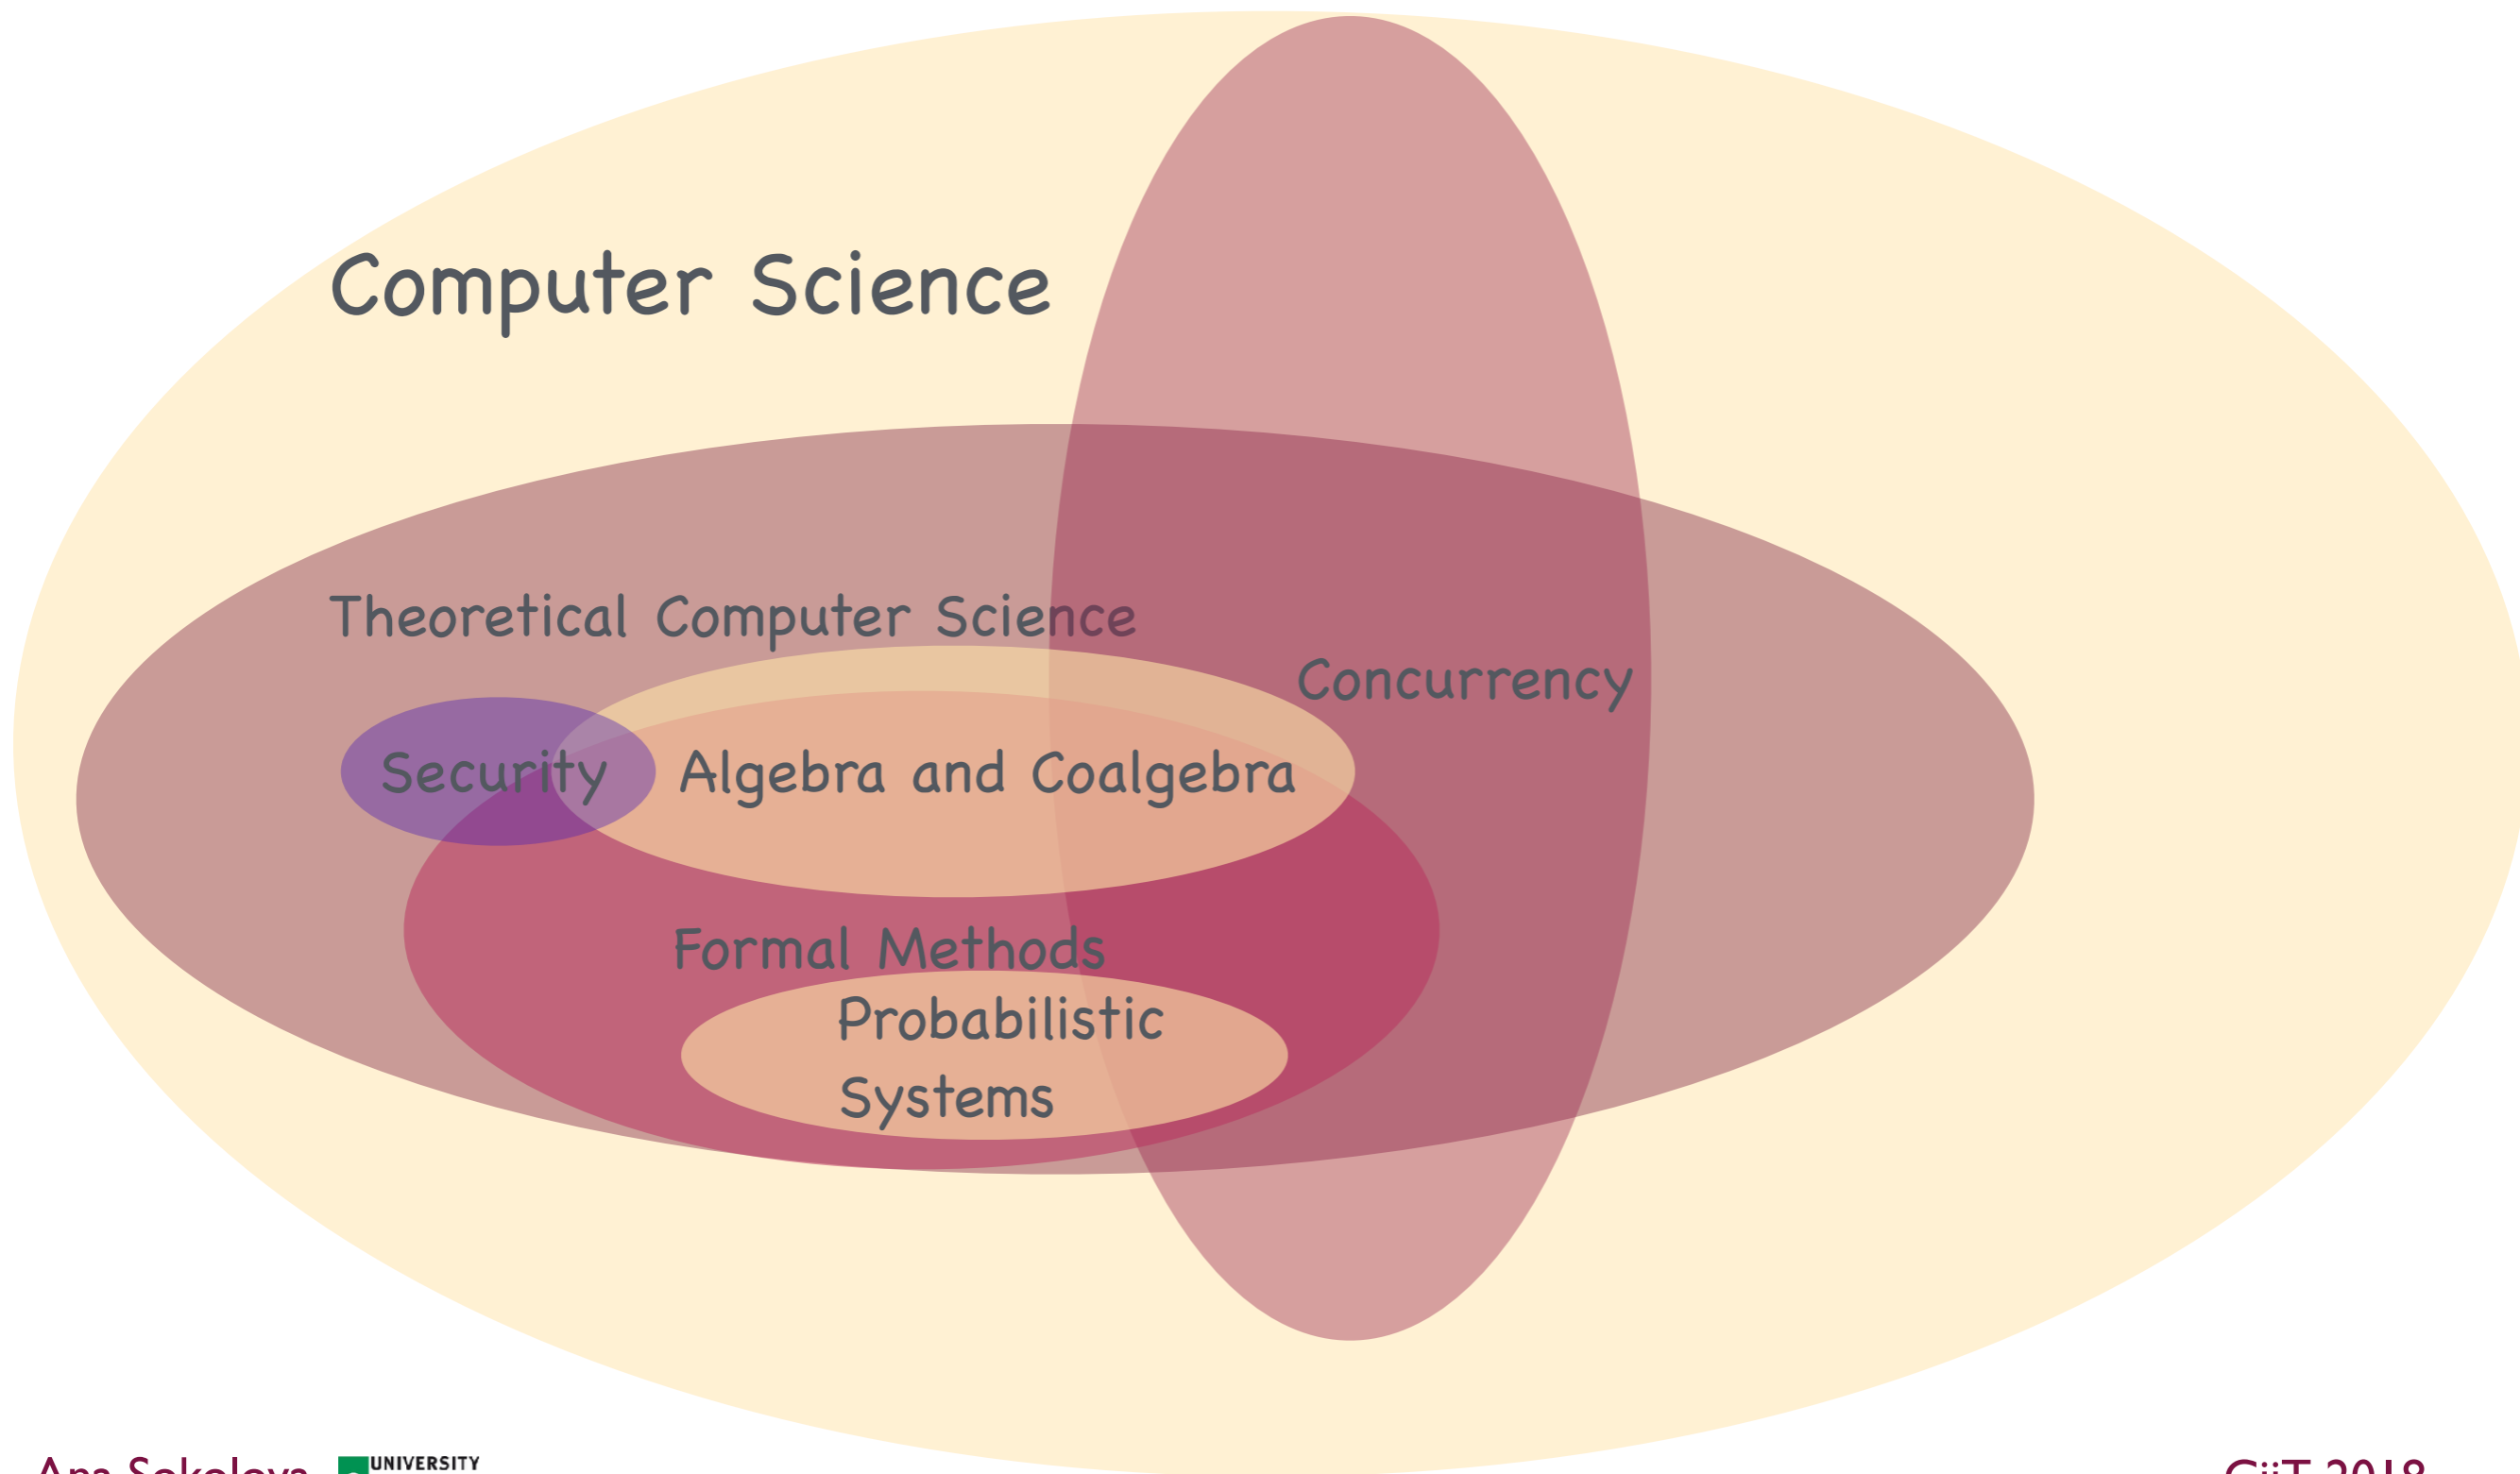
Background big picture



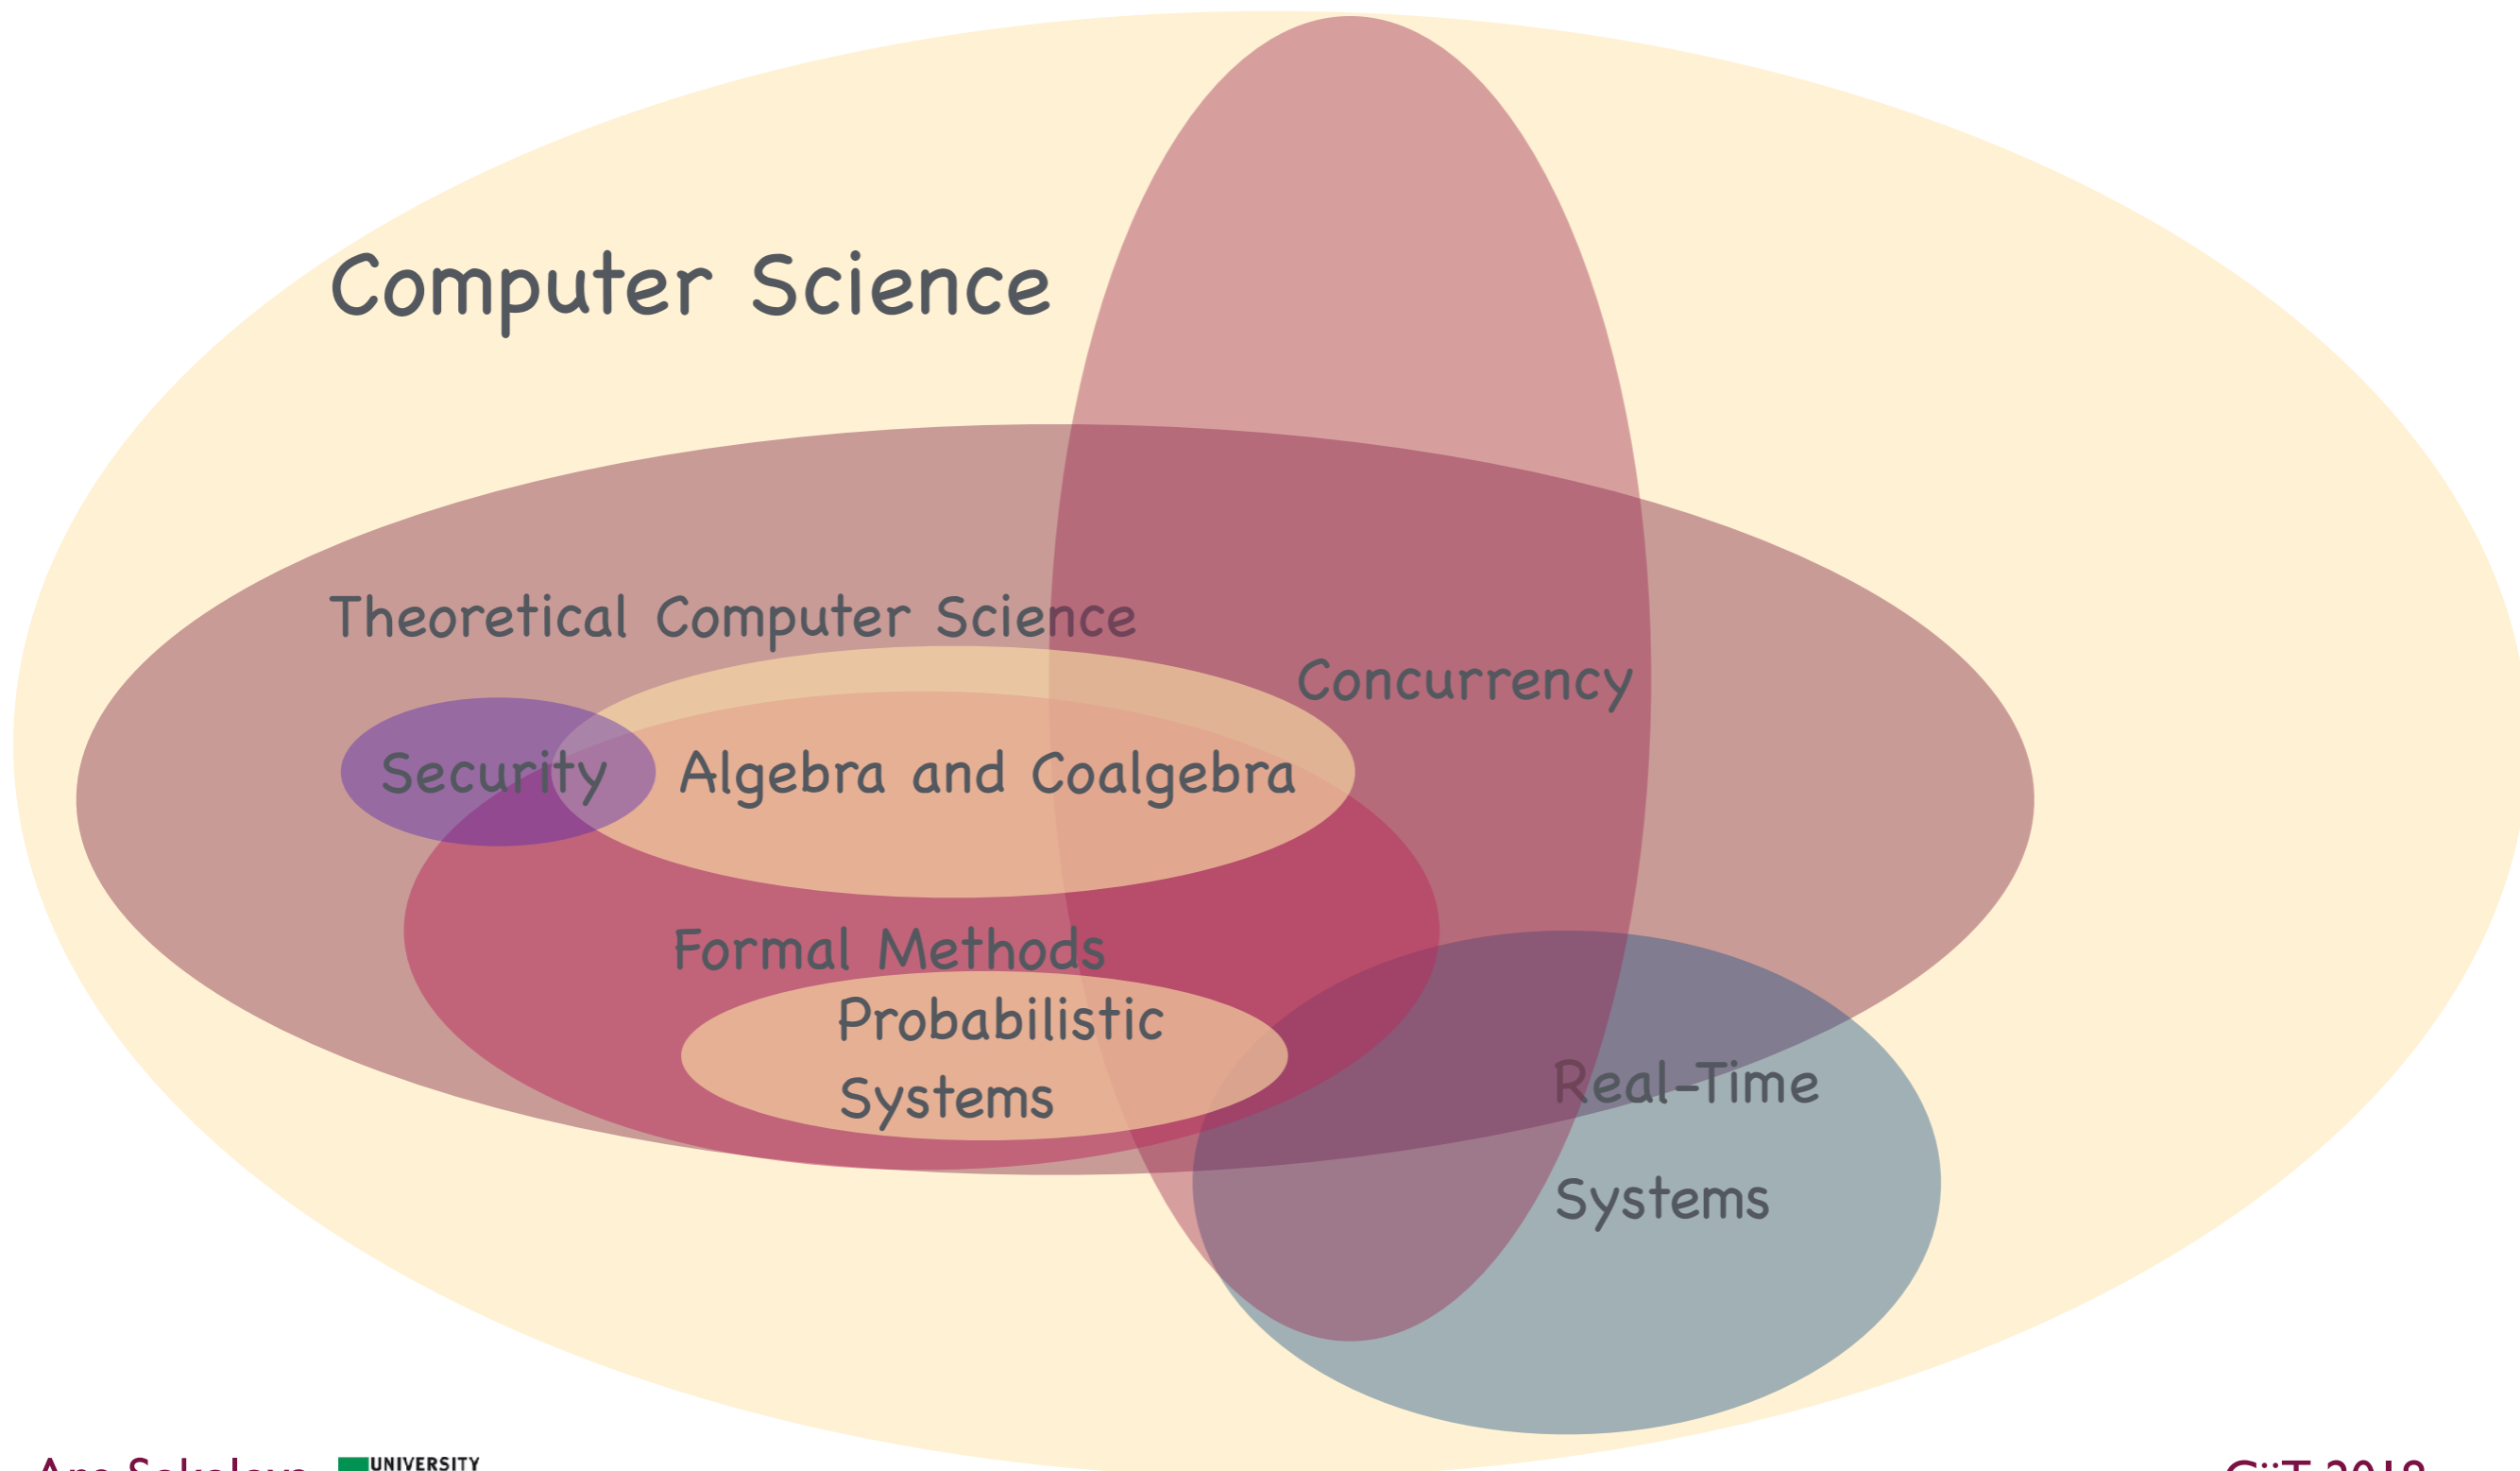
Background big picture



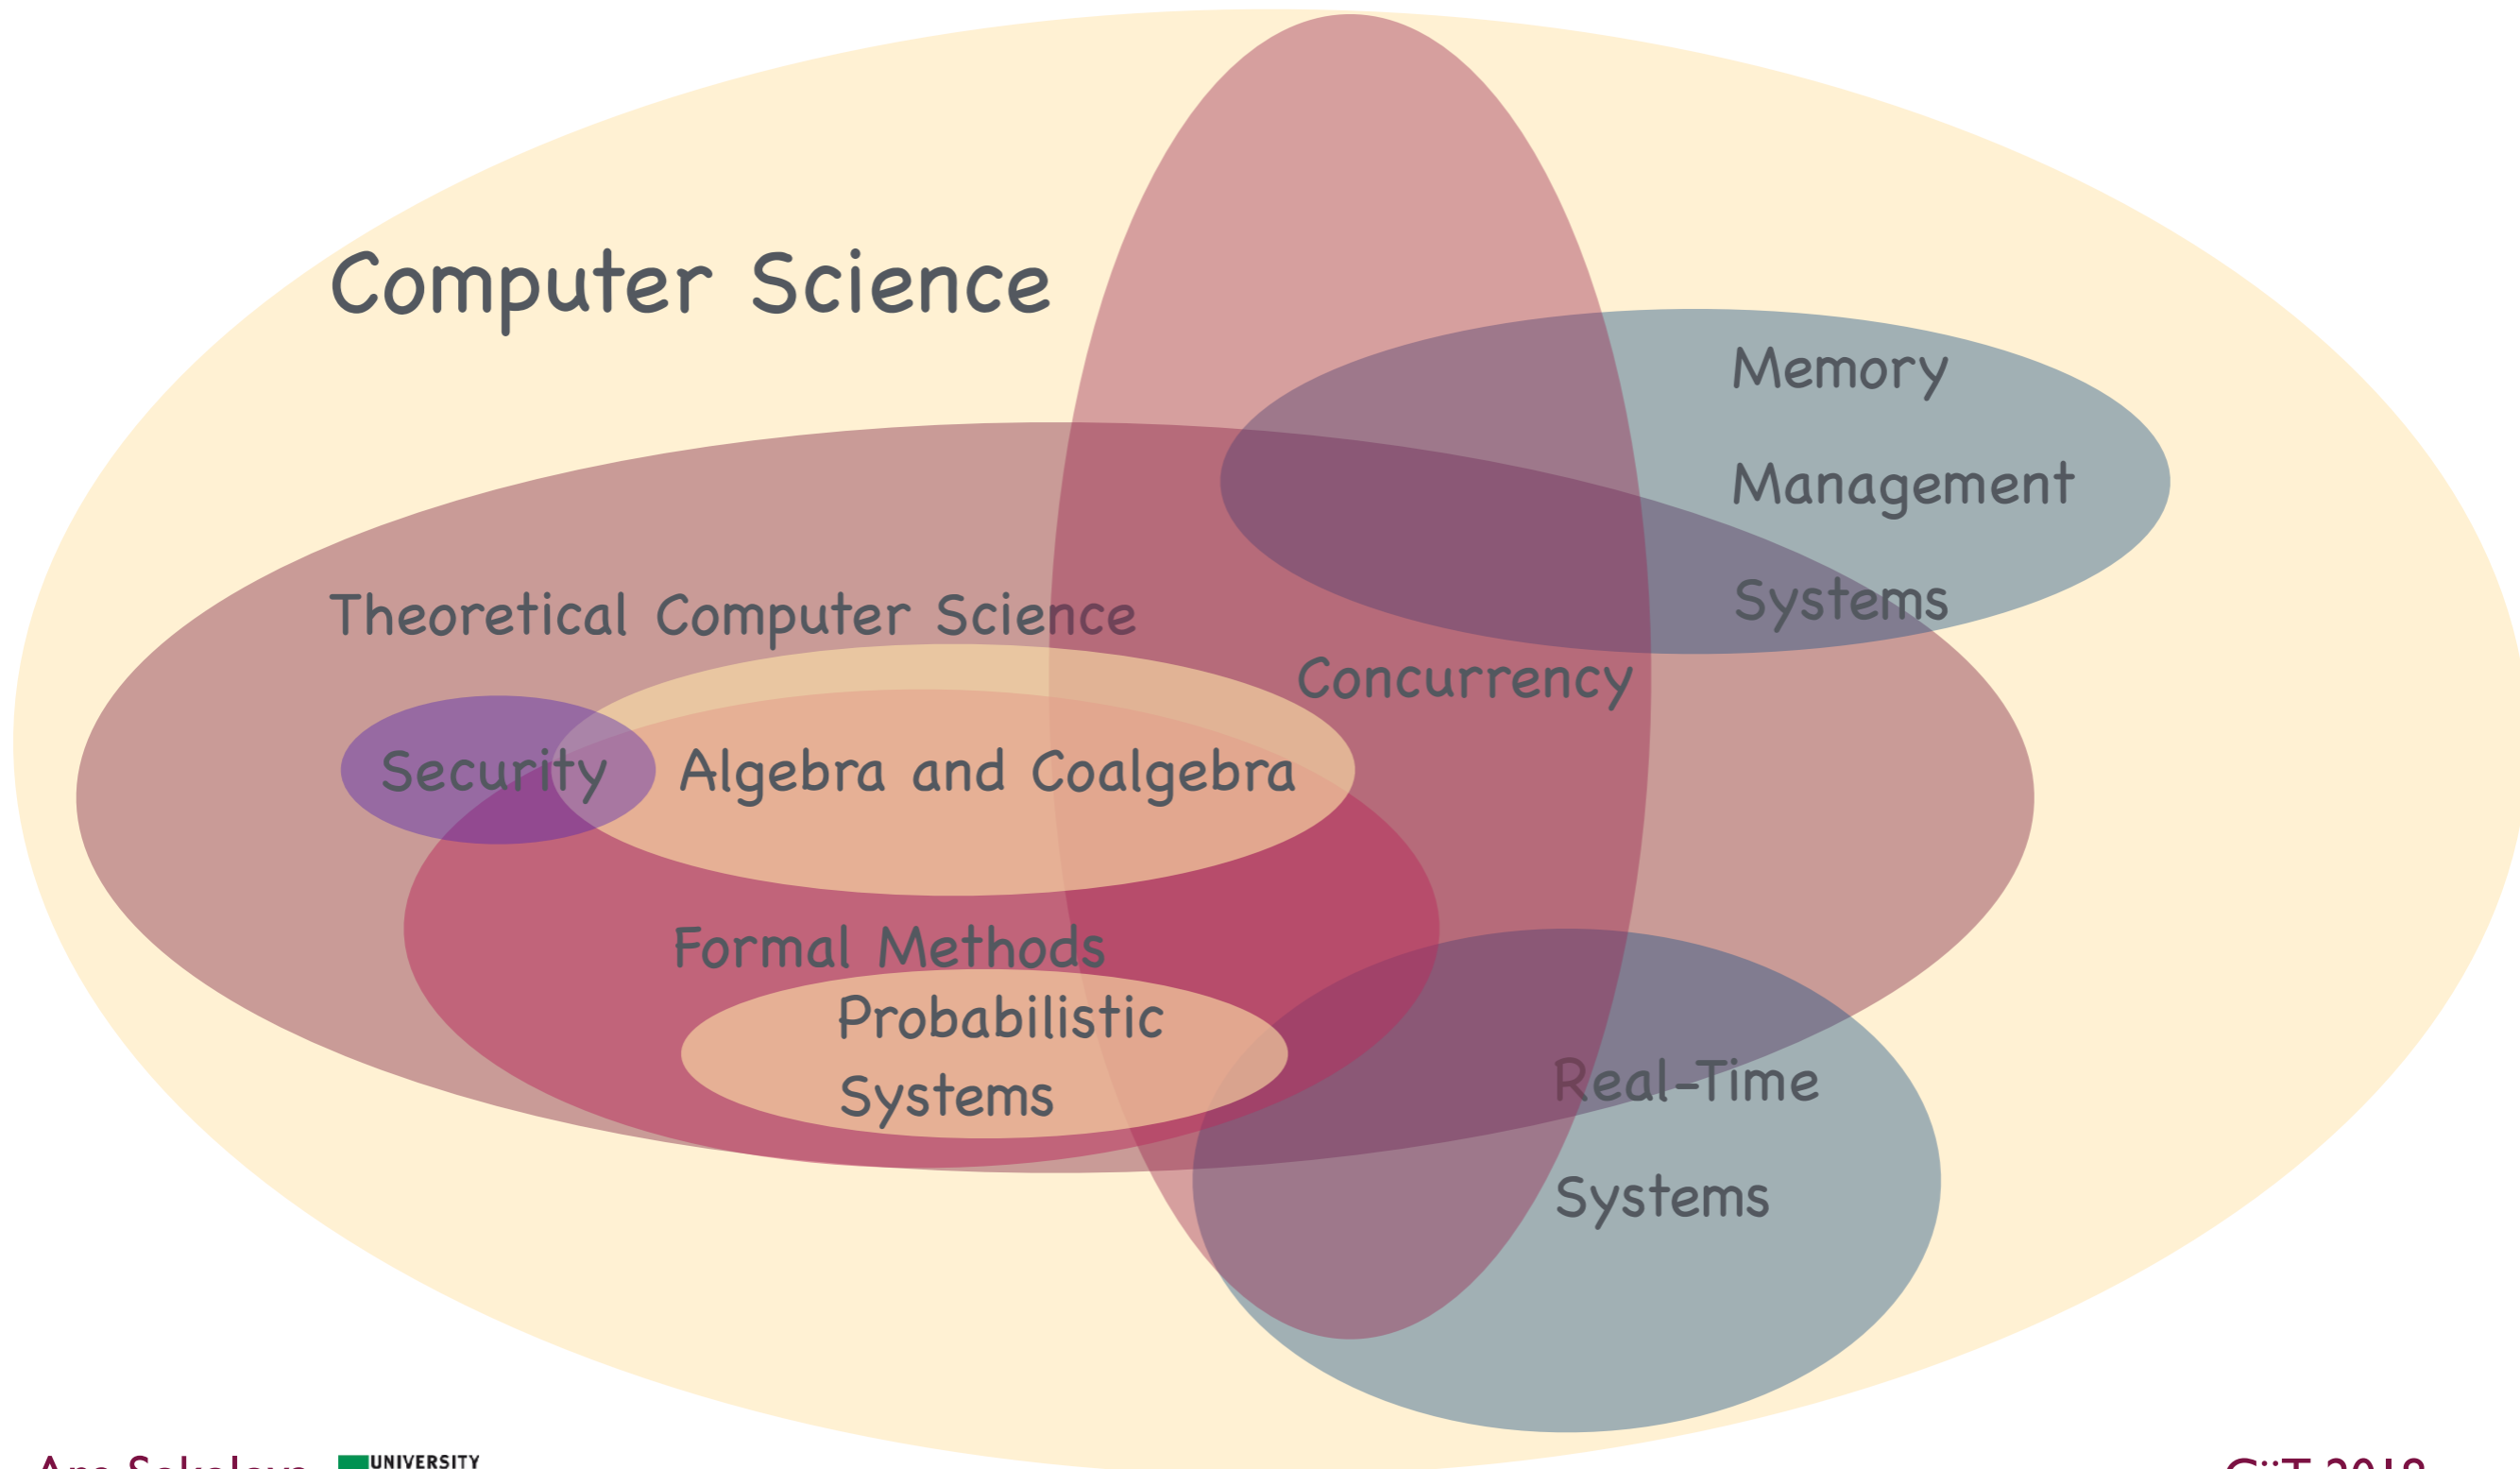
Background big picture



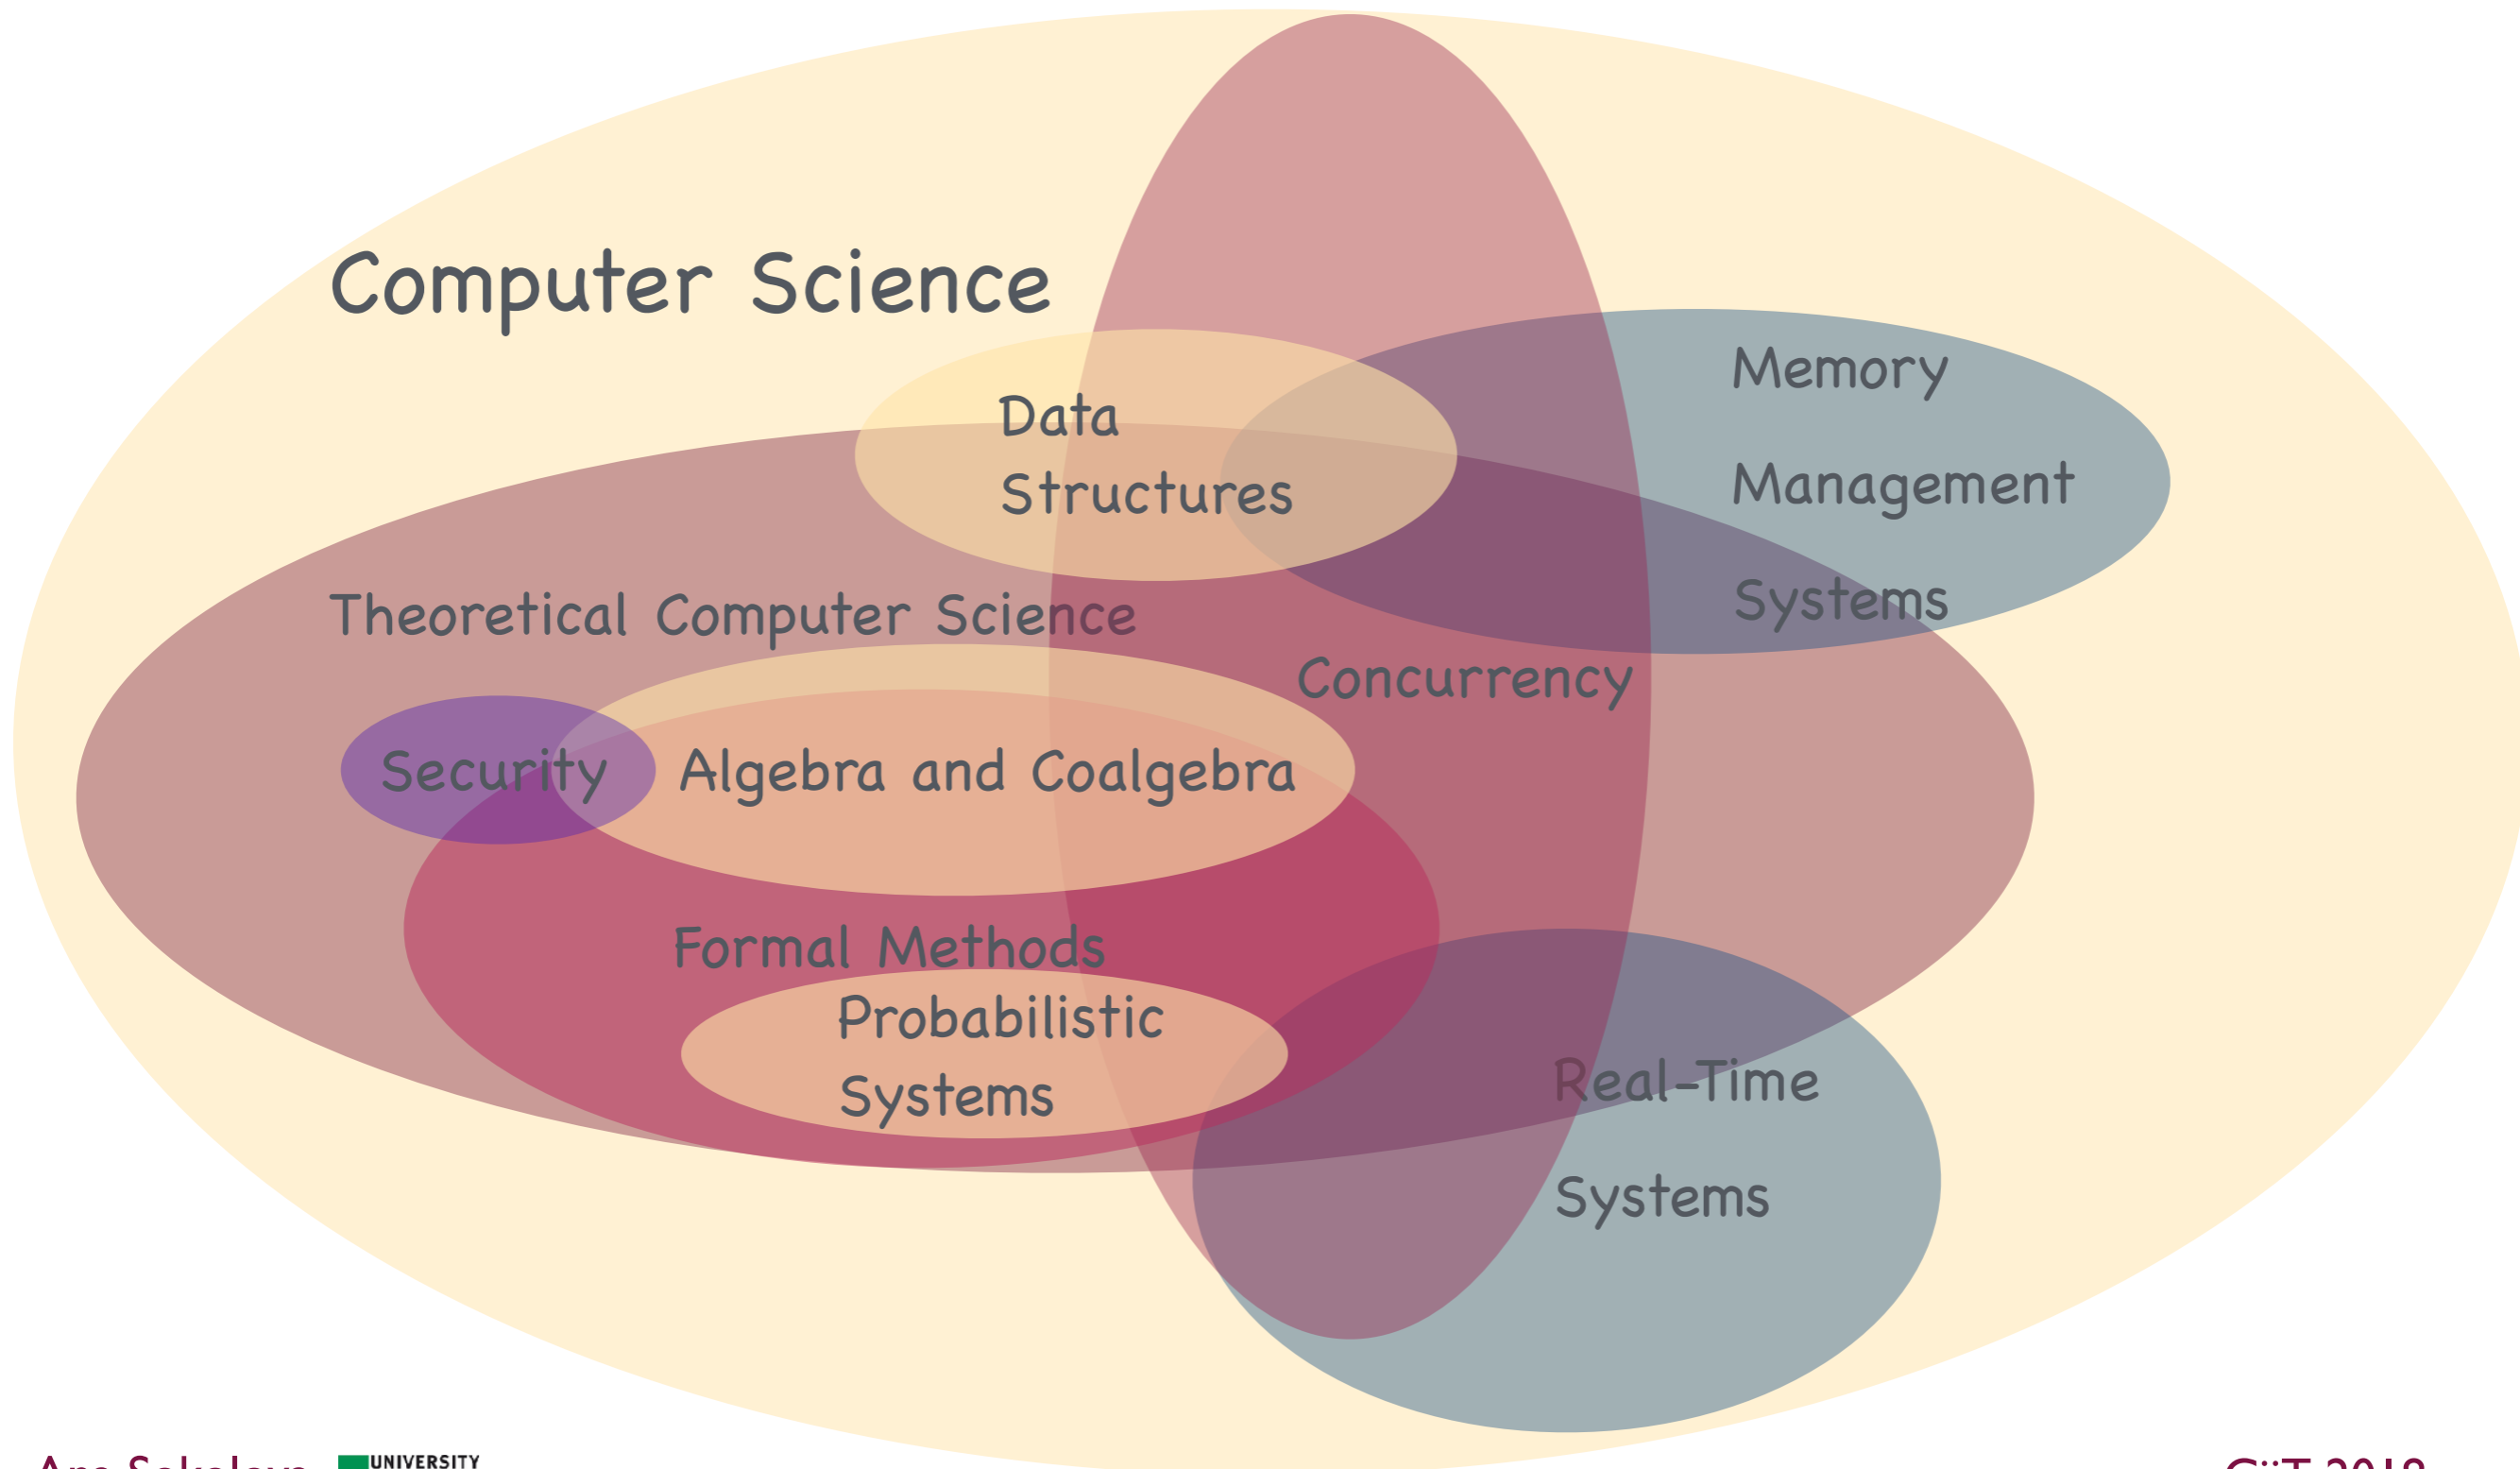
Background big picture



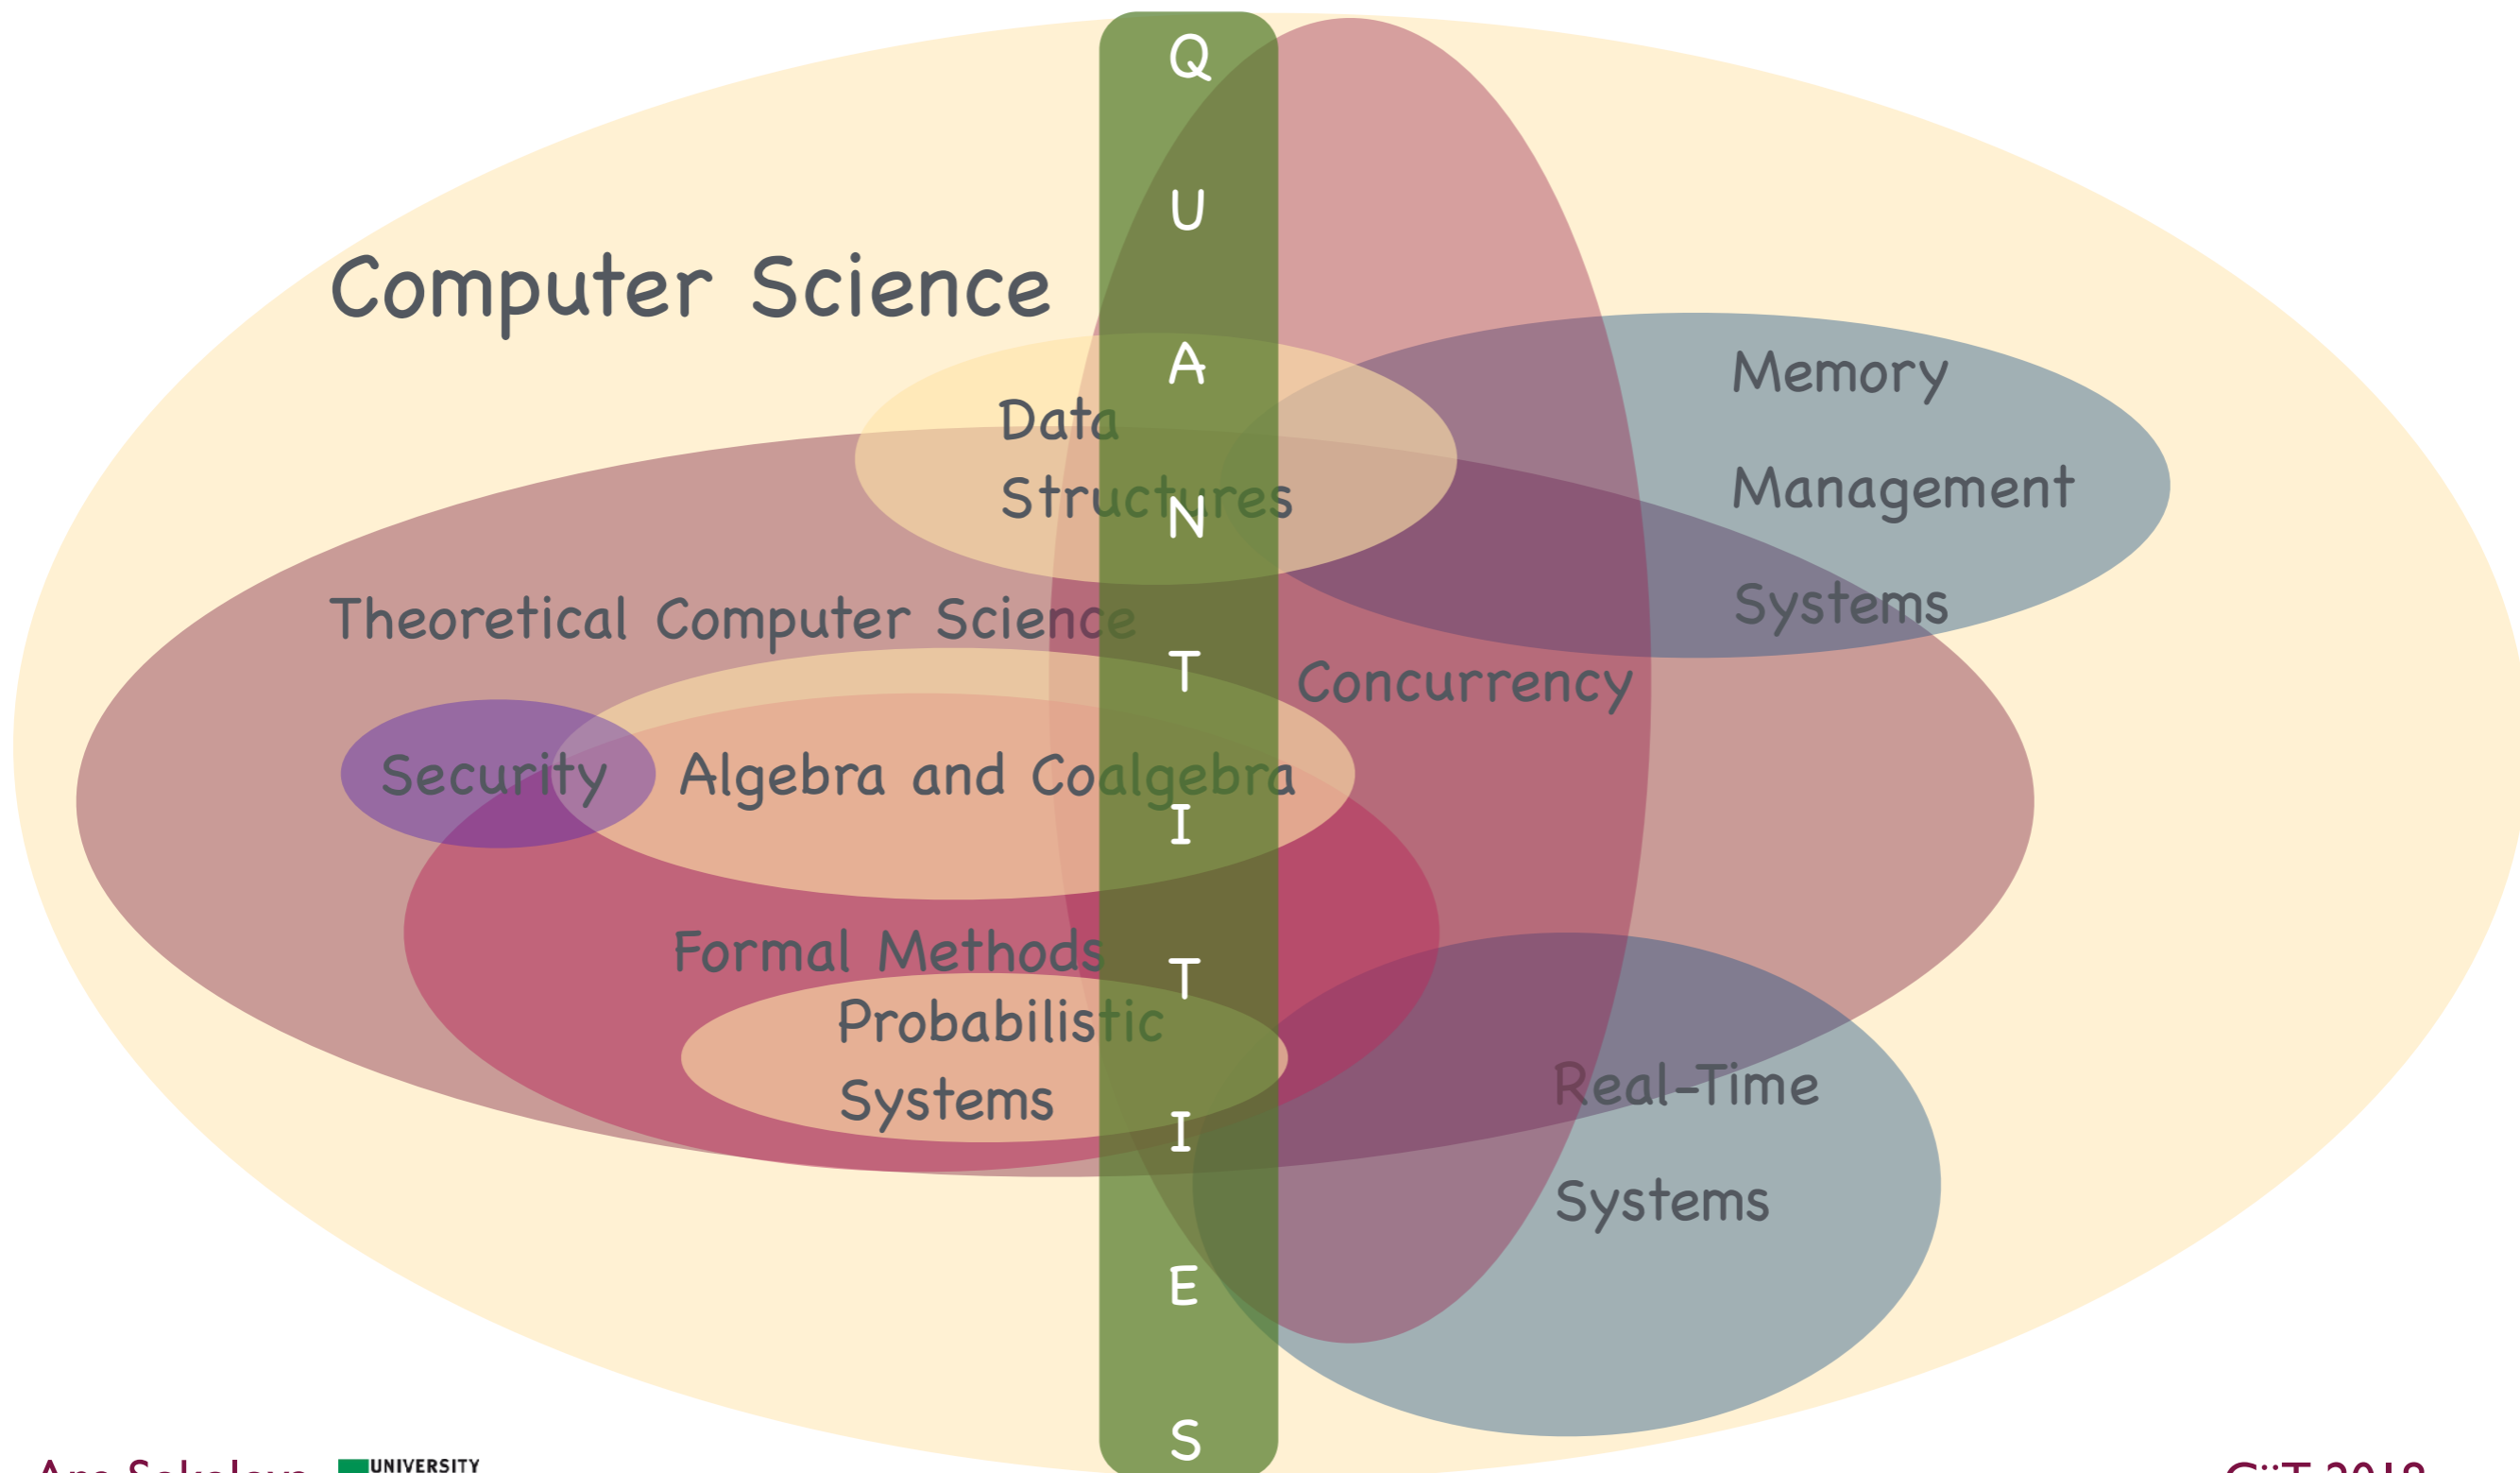
Background big picture



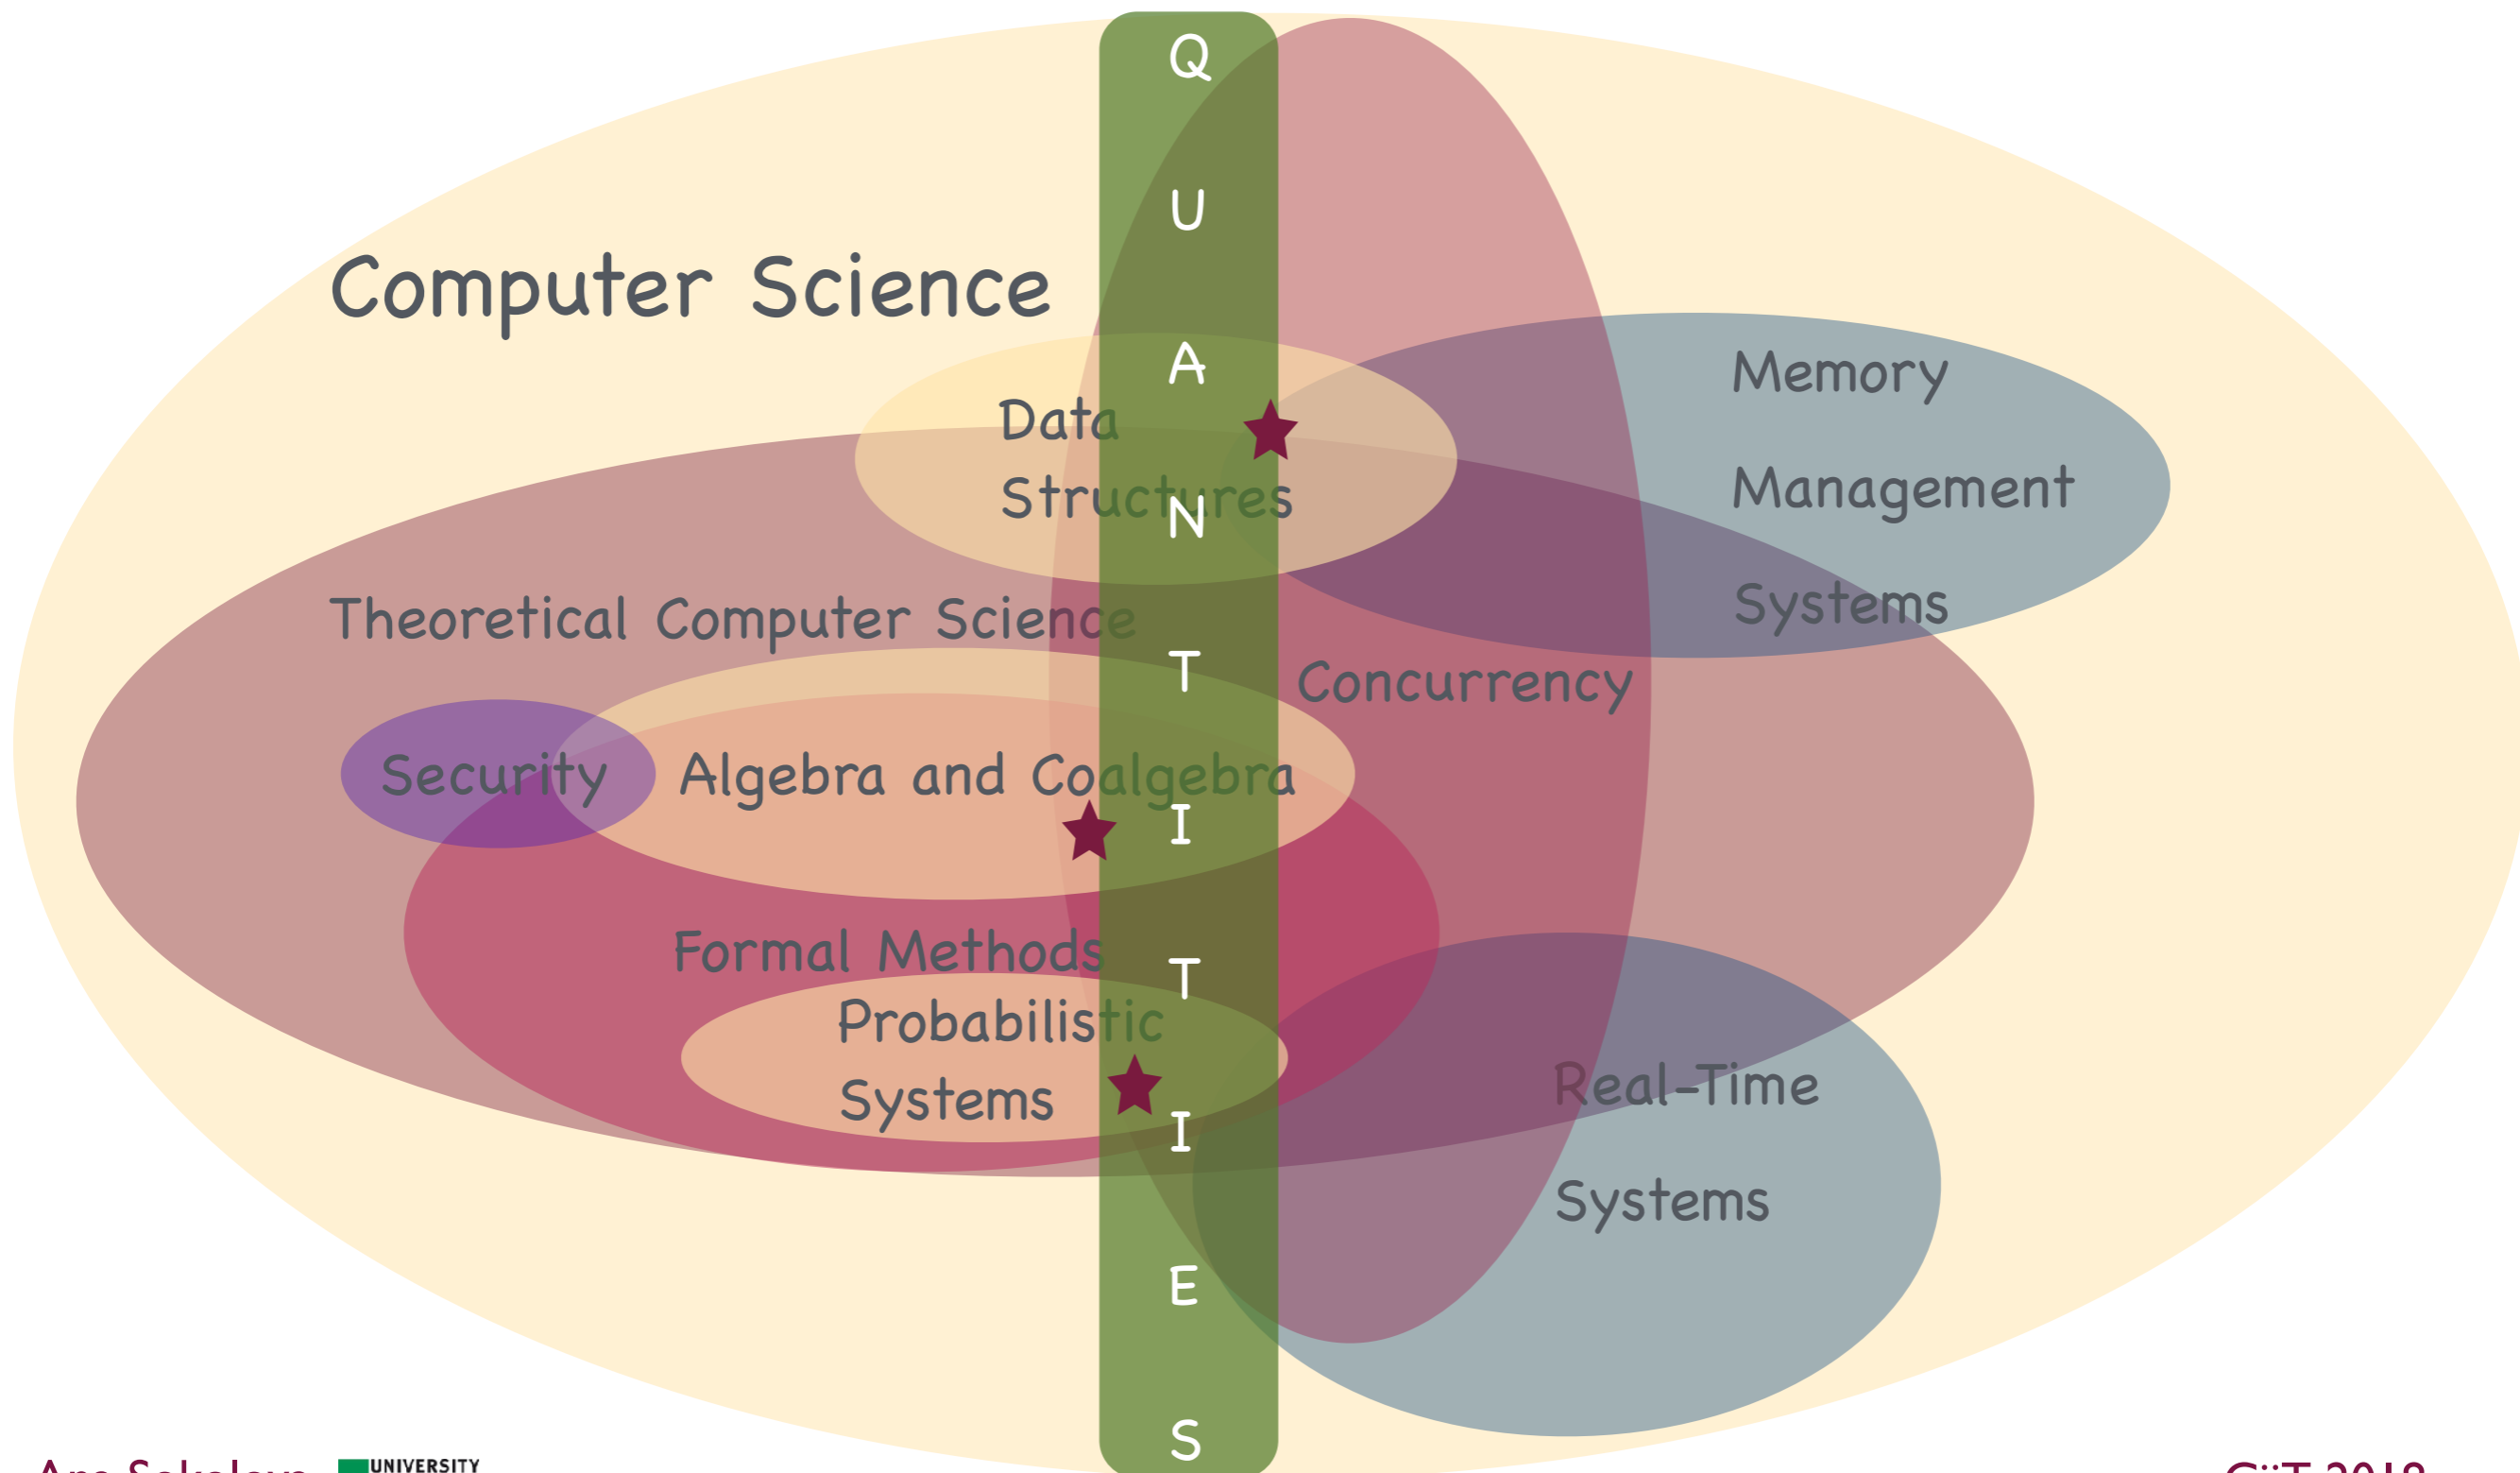
Background big picture



Background big picture



Current favourites



Local Linearizability

Ana Sokolova  UNIVERSITY
of SALZBURG

joint work with:

Andreas Haas 

Andreas Holzer  UNIVERSITY OF
TORONTO

Michael Lippautz 

Ali Sezgin  UNIVERSITY OF
CAMBRIDGE

Tom Henzinger  IST AUSTRIA

Christoph Kirsch  UNIVERSITY
of SALZBURG

Hannes Payer 

Helmut Veith  TU
WIEN

Concurrent Data Structures: Semantics and Relaxations

Concurrent Data Structures

Correctness and Performance

Concurrent Data Structures

Correctness and Performance



structure and power

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]



Sequential Consistency [Lamport'79]

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]



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]



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]




Sequential Consistency [Lamport'79]

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



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
Henzinger, Kirsch, Payer, Sezgin, S. POPL13

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

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

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

Local linearizability
in this talk

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

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

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

t1:

enq(1)

deq(2)

t2:

enq(2)

deq(1)

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.

Generalizations of Local Linearizability

Generalizations of Local Linearizability

Signature Σ

Generalizations of Local Linearizability

Signature Σ

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

$In_{\mathbf{h}}(i)$

$Out_{\mathbf{h}}(i)$

Generalizations of Local Linearizability

Signature Σ

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

$In_{\mathbf{h}}(i)$

$Out_{\mathbf{h}}(i)$

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

Generalizations of Local Linearizability

Signature Σ

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

$In_{\mathbf{h}}(i)$

$Out_{\mathbf{h}}(i)$

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

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

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

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.

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

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

Verification (queue)

Queue sequential specification (axiomatic)

s is a legal queue sequence

iff

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

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

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)

Henzinger, Sezgin, Vafeiadis CONCUR13

h is queue linearizable

iff

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

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

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)

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

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

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

thread-local
precedence order

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)




segment of possibly dynamic size (n)

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!

Local Linearizability

Ana Sokolova  UNIVERSITY
of SALZBURG

joint work with:

Andreas Haas 

Andreas Holzer  UNIVERSITY OF
TORONTO

Michael Lippautz 

Ali Sezgin  UNIVERSITY OF
CAMBRIDGE

Tom Henzinger  IST AUSTRIA

Christoph Kirsch  UNIVERSITY
of SALZBURG

Hannes Payer 

Helmut Veith  TU
WIEN

Thank You!

Local Linearizability

Ana Sokolova  UNIVERSITY
of SALZBURG

joint work with:

Andreas Haas 

Andreas Holzer  UNIVERSITY OF
TORONTO

Michael Lippautz 


Ali Sezgin  UNIVERSITY OF
CAMBRIDGE

Tom Henzinger  IST AUSTRIA

Christoph Kirsch  UNIVERSITY
of SALZBURG

Hannes Payer 

Helmut Veith  TU
WIEN



Ali Sezgin 



Hannes Payer




Andreas Holzer 




Michael Lippautz




Tom Henzinger




Helmut Veith 



Andreas Haas 



Christoph Kirsch
