

The Push/Pull model of serializability

Eric Koskinen

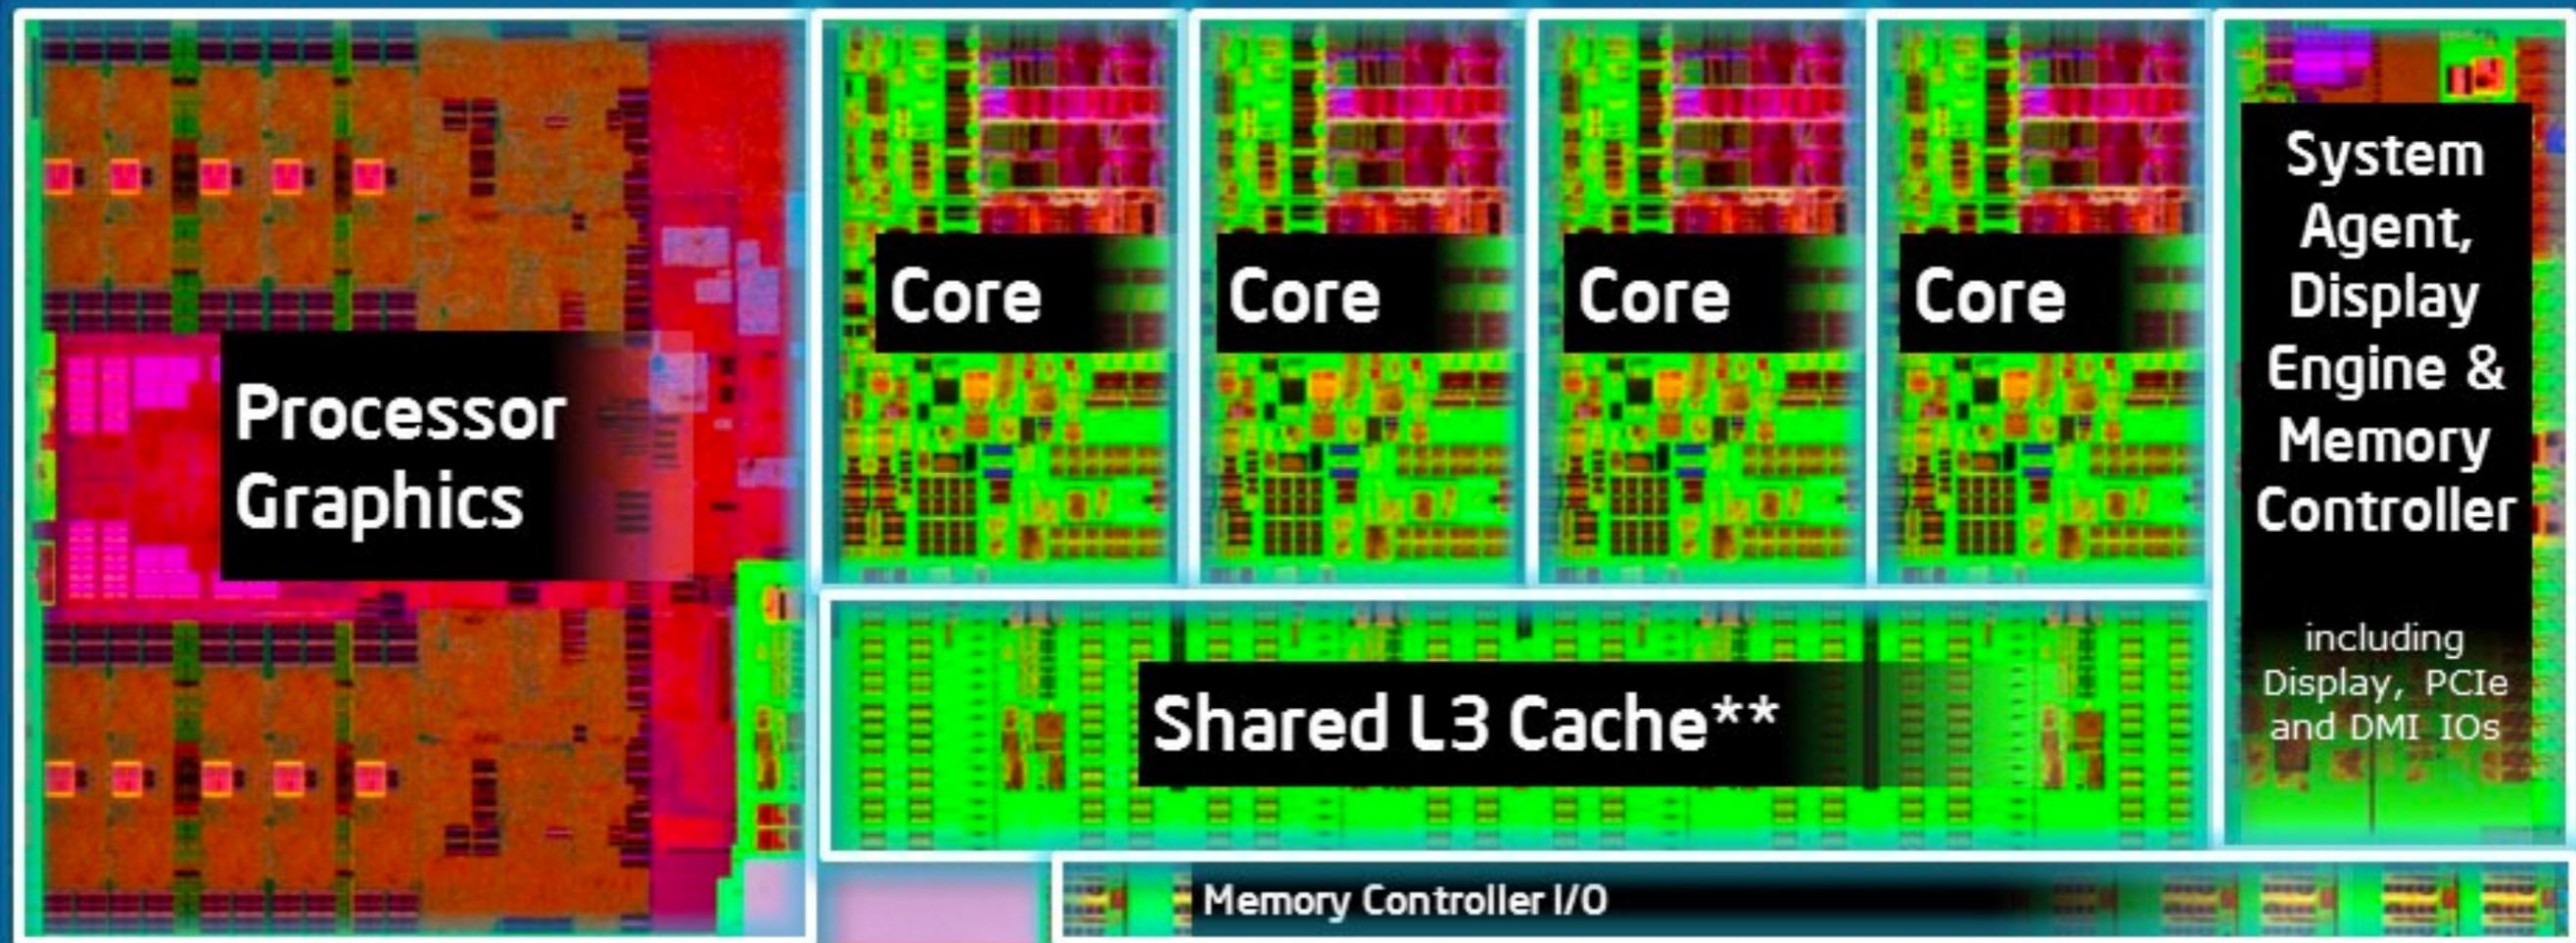
Research Staff Member
IBM TJ Watson Research Center

Joint work with Matthew Parkinson (Microsoft)



4th Generation Intel® Core™ Processor Die Map

22nm Haswell Tri-Gate 3-D Transistors



Quad core die shown above | Transistor count: 1.4Billion | Die size: 177mm²

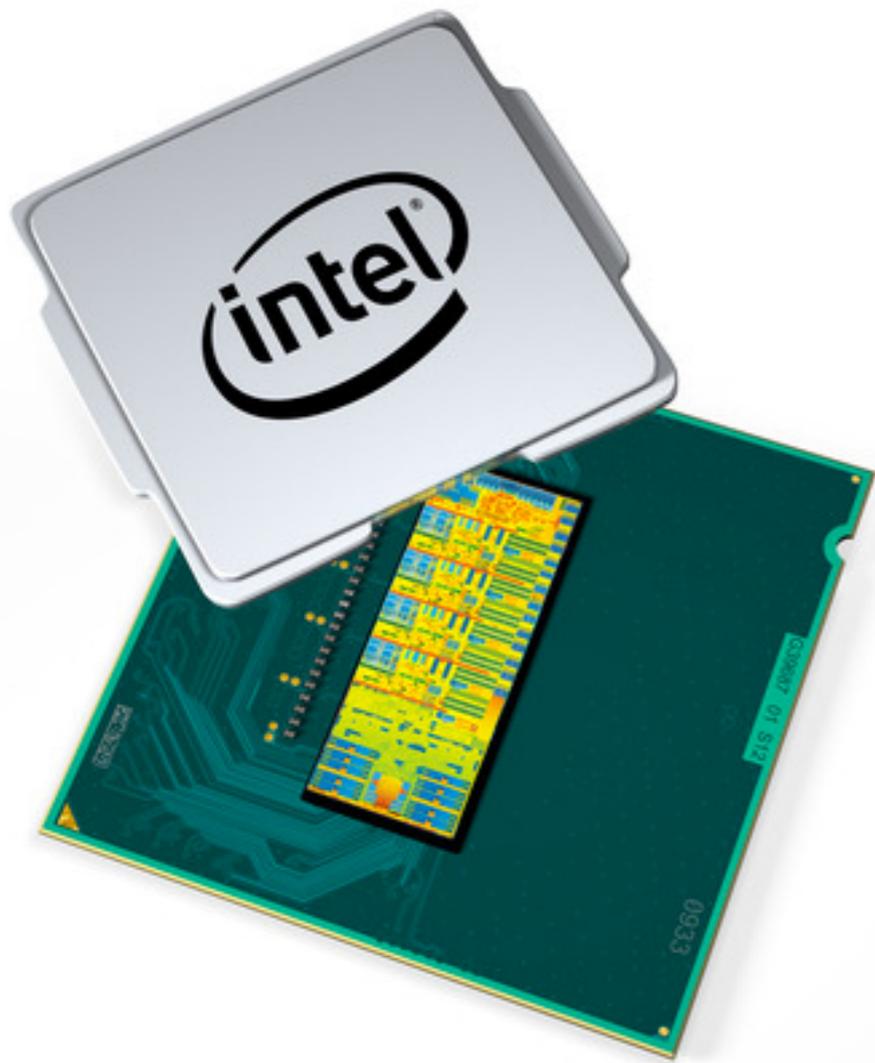
** Cache is shared across all 4 cores and processor graphics

All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

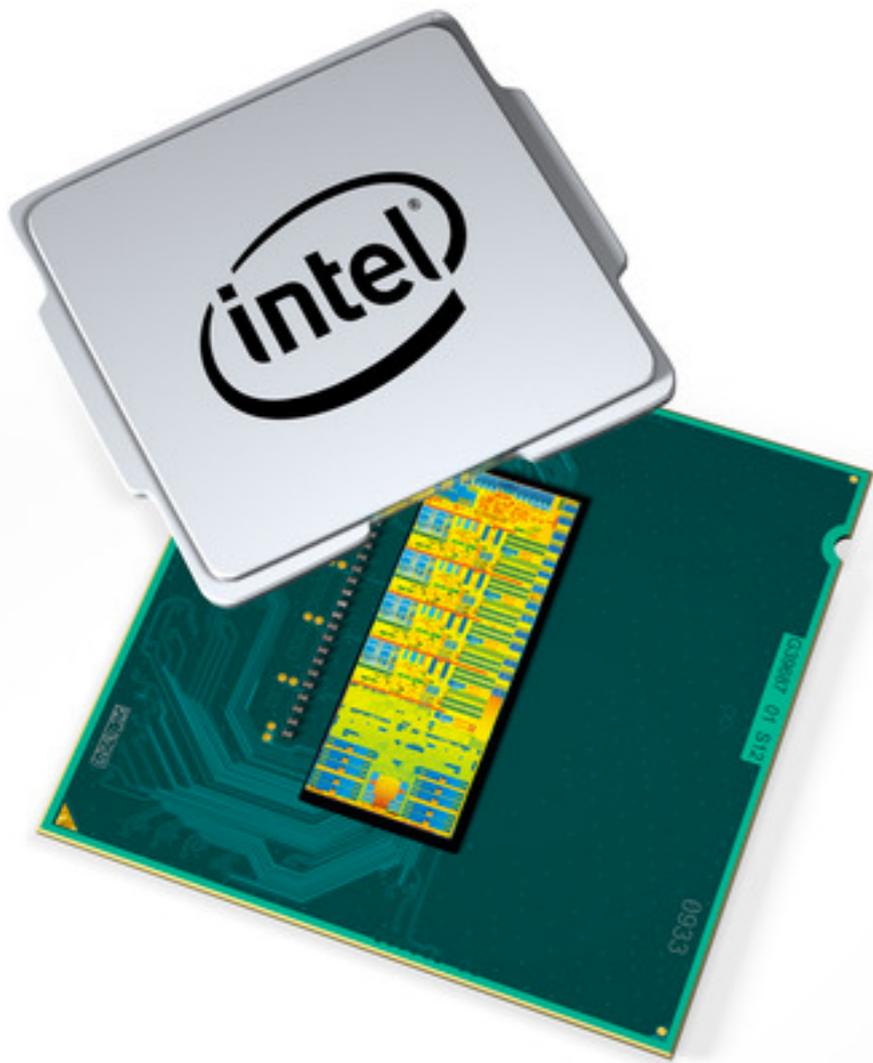
UNDER EMBARGO UNTIL FURTHER NOTICE

INTEL CONFIDENTIAL





XBEGIN , **XEND** and **XABORT**



XBEGIN , **XEND** and **XABORT**



Version 4.7

Transactions

```
x := 0; y := 0;
```

Thread 1

```
atomic {  
    x := 3;  
    y := 9;  
}
```

Thread 2

```
atomic {  
    if (x > y) {  
        //shdnt happen  
    } else {  
        ...  
    }  
}
```



Pessimism!

Thread

```
atomic {  
  x := 3;  
  y := 9;  
}
```

x changed?

y changed?

Pessimism!

Thread

```
atomic {  
  x := 3;  
  y := 9;  
}
```

x changed?

y changed?

Optimism!

Thread

```
atomic {  
  x := 3;  
  y := 9;  
}
```

all ok?

Pessimism!

Serializability!

Interleaved execution
equivalent to some
serial execution.

Optimism!

Pessimism!

Opacity!

Interleaved execution
equivalent to some
serial execution
and
cannot observe
intermediate state.

Serializability!

Interleaved execution
equivalent to some
serial execution.

Optimism!

Pessimism!

Opacity!

Serializability!

Optimism!

```
Thread  
atomic {  
  x := 3;  
  y := 9;  
}
```

Memory Ops

ADT Ops
(Boosting)

```
Thread  
atomic {  
  stk.push(4);  
  ht.get('a');  
}
```

Pessimism!

Opacity!

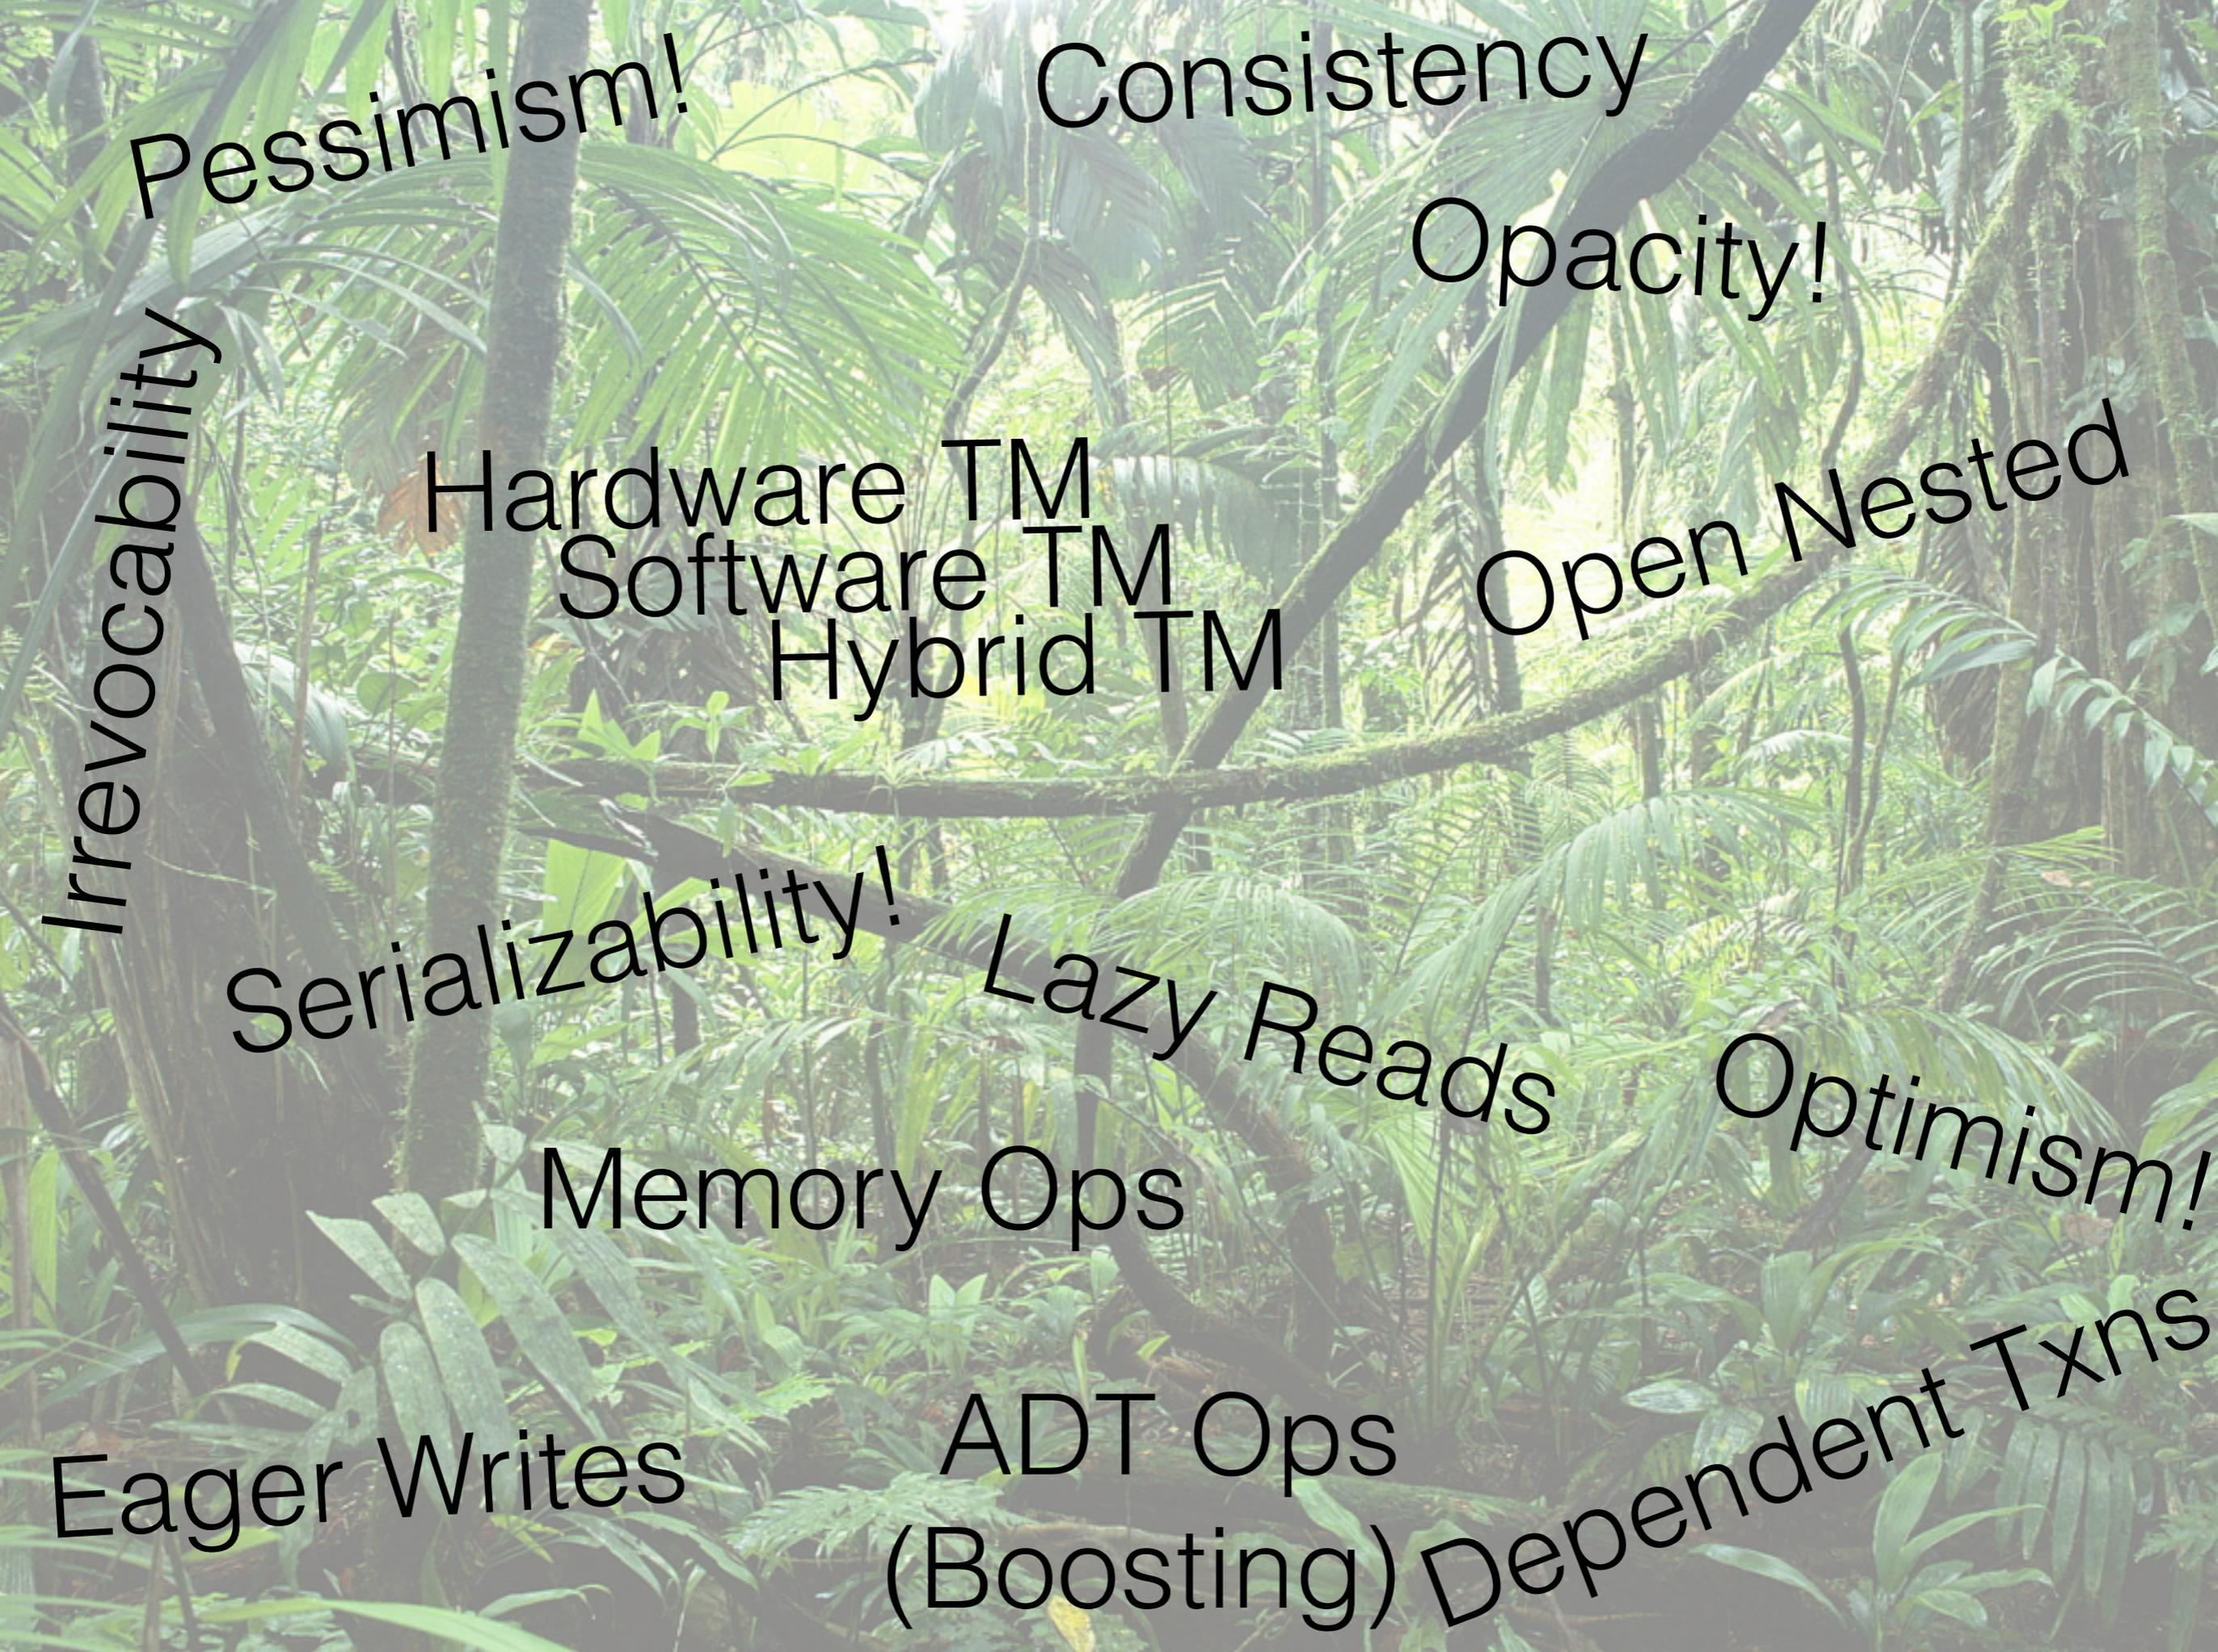
Hardware TM
Software TM
Hybrid TM

Serializability!

Optimism!

Memory Ops

ADT Ops
(Boosting)



Pessimism!

Consistency

Opacity!

Irrevocability

Hardware TM

Software TM

Hybrid TM

Open Nested

Serializability!

Lazy Reads

Optimism!

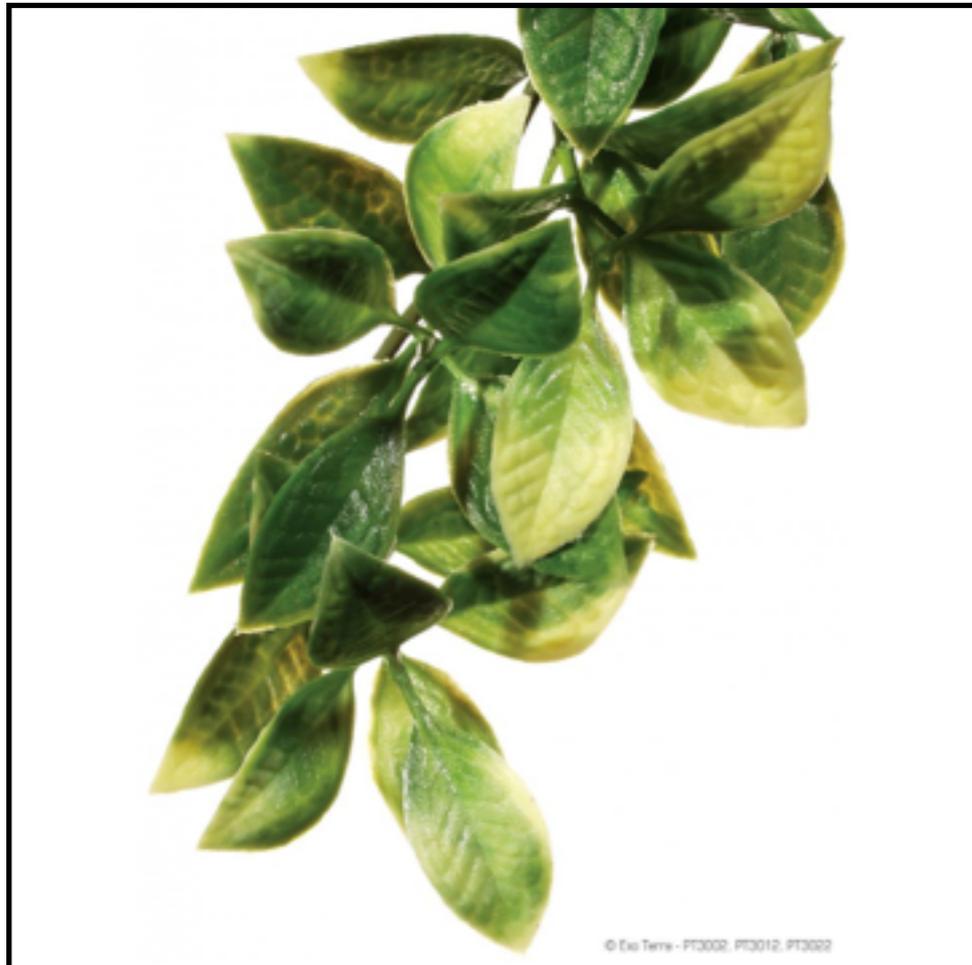
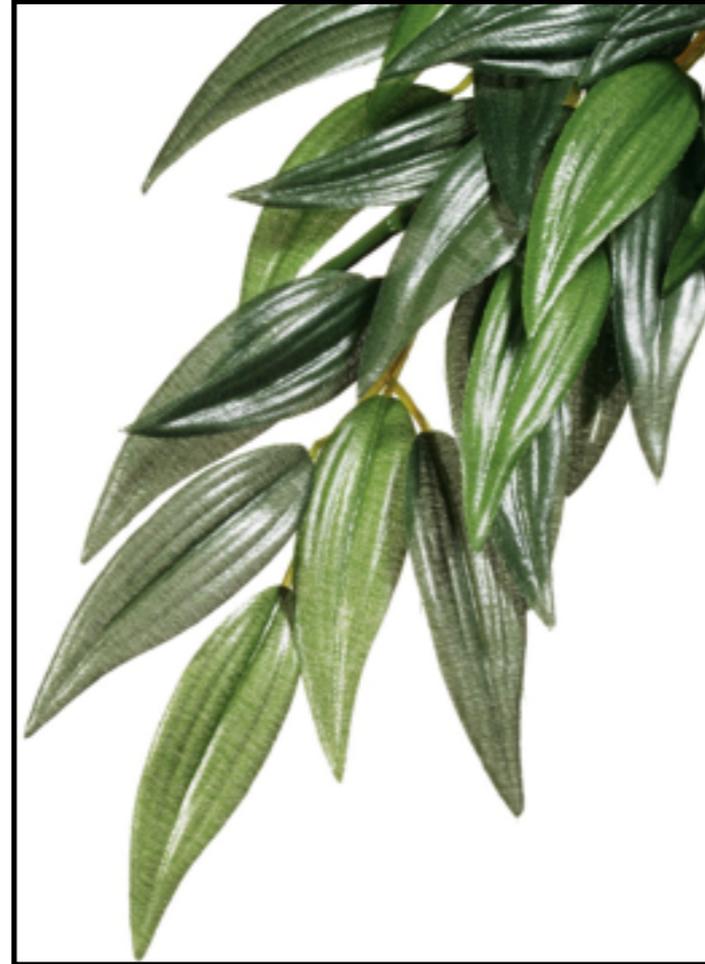
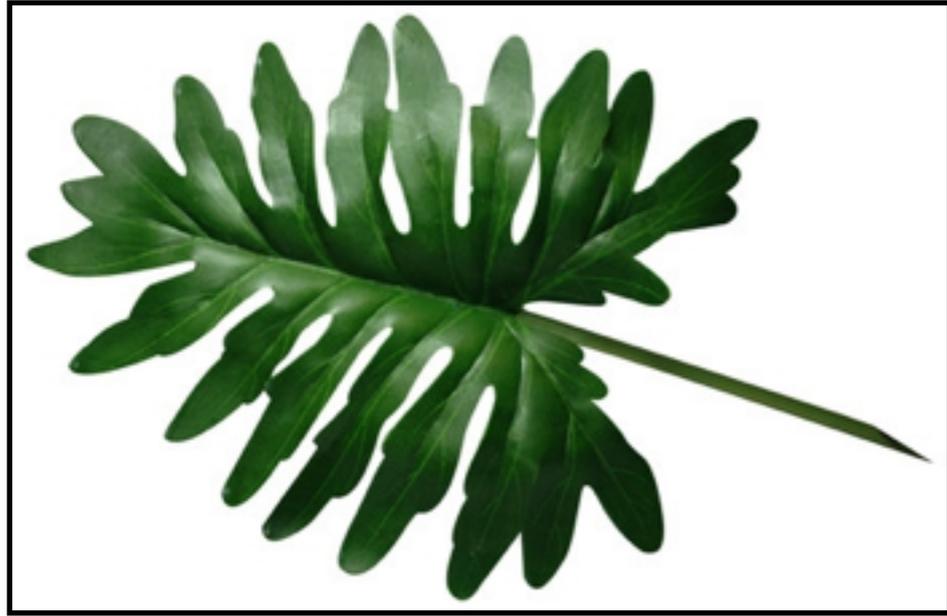
Memory Ops

Eager Writes

ADT Ops

(Boosting) Dependent Txns

Unified theory



Unified theory

The Push/Pull model of transactions

Eric Koskinen

Matthew Parkinson

Abstract

We present a general theory of serializability, unifying a wide range of transactional algorithms, including some that are yet to come. To this end, we provide a compact semantics in which concurrent transactions *push* their effects into the shared view (or *unpush* to recall effects) and *pull* the effects of potentially uncommitted concurrent transactions into their local view (or *unpull* to detangle). Each operation comes with simple side-conditions given in terms of commutativity (Lipton's left-movers and right-movers [24]). The benefit of this model is that most of the elaborate reasoning (coinduction, simulation, subtle invariants, etc.) necessary for proving the serializability of a transactional algorithm is already proved within the semantic model. Thus, proving serializability (or opacity) amounts simply to mapping the algorithm on to our rules, and showing that it satisfies the rules' side-conditions.

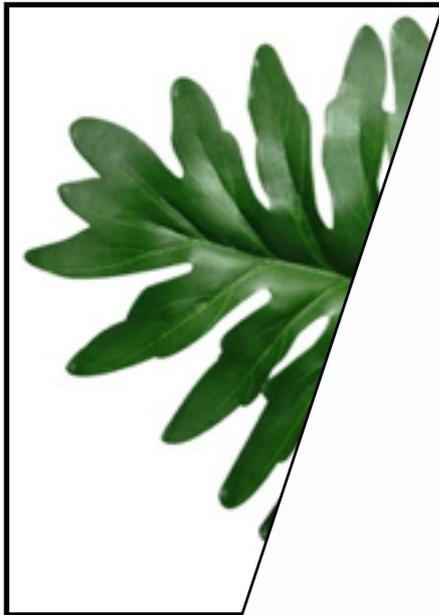
1. Introduction

Recent years have seen an explosion of research on methods of providing atomic sections in modern programming languages, typically implemented via transactional memory (TM). The atomic keyword provides programmers with a powerful concurrent programming building block: the ability to specify when a thread's operations on shared memory should appear to take place instantly when viewed by another thread. To support such a construct, we must be able to reason about detecting conflicts between concurrent threads. This can be done atomically. Implementations typically achieve this by dynamically tracking memory operations in hardware [15-17] or software [4, 6, 8, 14, 25]. Meanwhile, an alternate approach exploits abstract-level notions of conflict over linearizable data-structure operations such as commutativity [11, 20, 21, 28]. Both levels of abstraction choose between optimistic execution, pessimistic execution, and circumstances under which one or the other is appropriate. Unfortunately, we lack a unified theory that leads to

transactions (e.g. transactional boosting [11]). Today, at best, we have two custom semantics for reasoning about the models individually, but no unified view.

We present a simple calculus that illuminates the core of transactional memory systems. In our model concurrent transactions *push* their effects into the shared log (or *unpush* to roll-back) and *pull* in the effects of potentially uncommitted concurrent transactions (or *unpull* to detangle). Moreover, transactions can push or pull operations in non-chronological orders, provided certain commutativity (Lipton left/right-movers [24]) conditions hold. The benefit of this semantic model is that most of the elaborate reasoning (coinduction, simulation relations, subtle invariants, etc.) necessary for proving the correctness of a transactional algorithm is contained within the semantic model, and need only be proved once. Our work formulates an expressive class of transactions and we have applied it to a wide range of TM systems including: optimistic read/write software TMs [6, 8], hardware transactional memories (Intel [17], IBM [16]), pessimistic TMs [4, 11, 25], hybrid optimistic/pessimistic TMs such as irrevocability [34], open-nested transactions [28], and abstract-level techniques such as boosting [11].

Our choice of expressiveness includes transactions that are not opaque [10]: transactions may share their uncommitted effects. This choice carves out a design space for implementations (e.g. actions [30]) and is relatively unrestricted. However, despite the advantage of the full spectrum of possibilities (e.g. opacity gives the appropriate criteria), while, we can

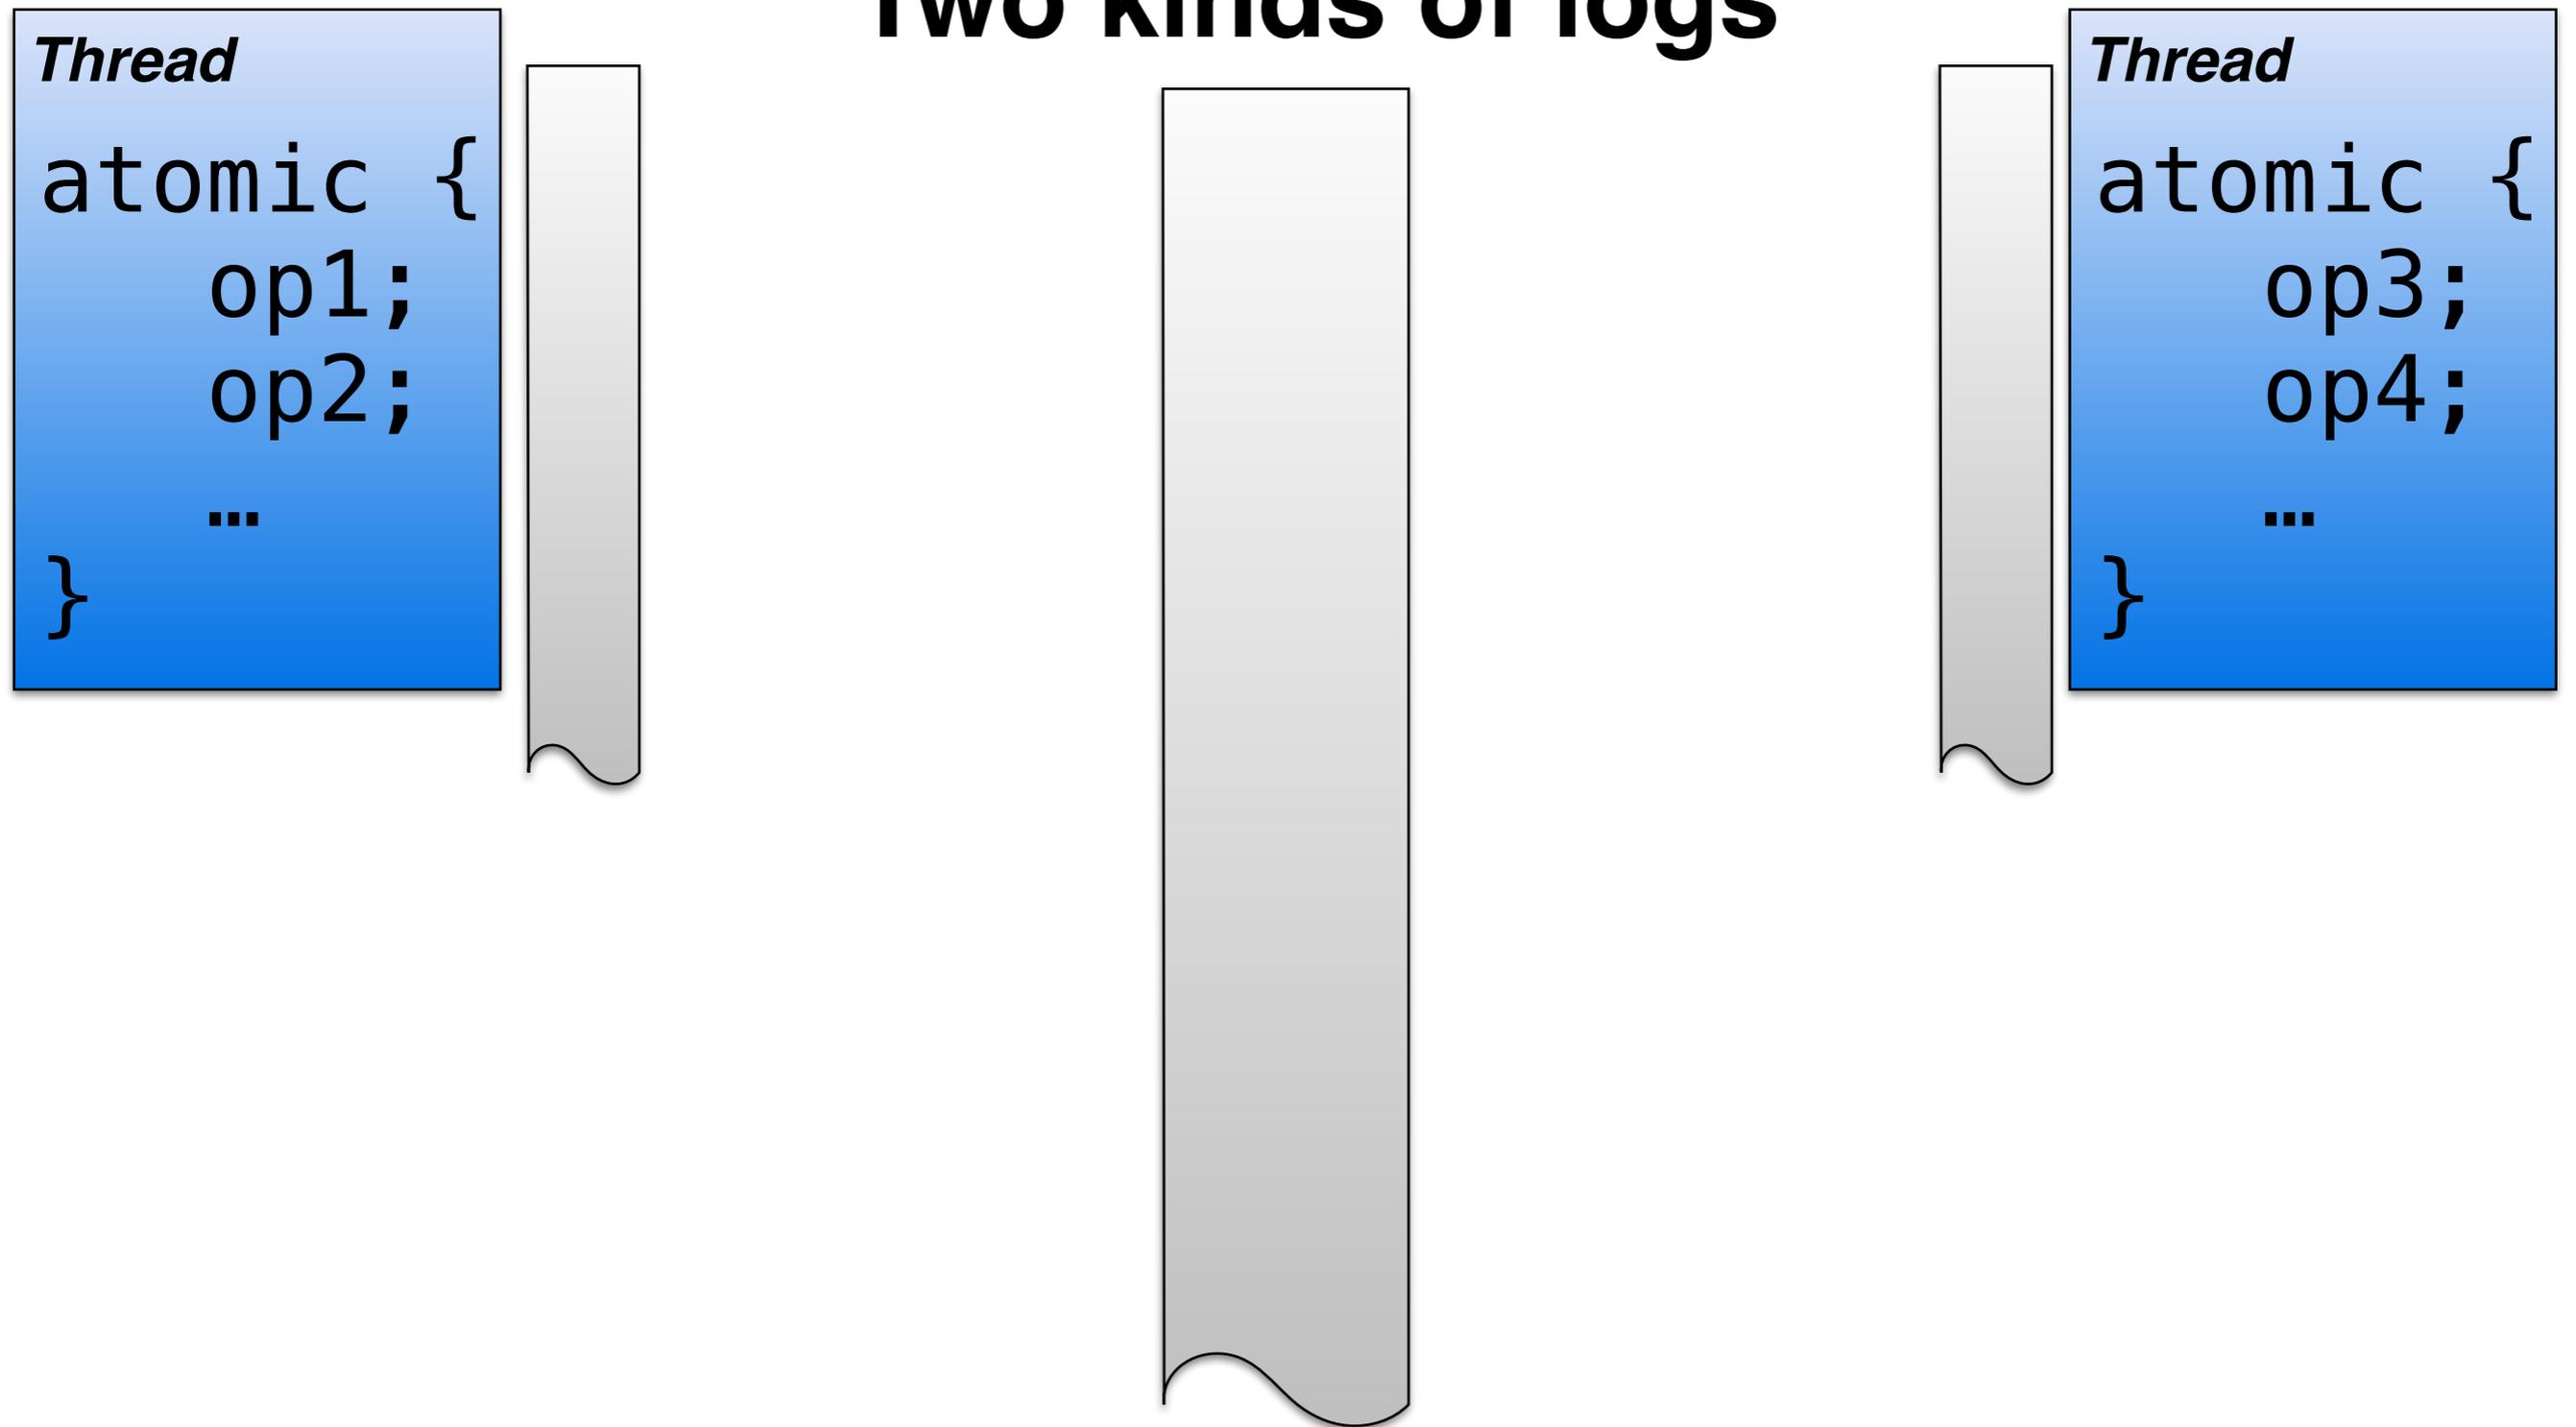


One model to rule them all.

There is no state.

There is no state.

Two kinds of logs



There is no state.

$\{\text{tx } c_1, \sigma_1\}$

Thread

L_1

$\langle m_1, \sigma_1, \sigma_1, l_1 \rangle$

G

$\{\text{tx } c_2, \sigma_2\}$

L_2

Thread

$\langle m_2, \sigma_2, \sigma_2, l_2 \rangle$

There is no state.

$\{\text{tx } c_1, \sigma_1\}$

Thread

L_1

$\langle m_1, \sigma_1, \sigma_1, l_1 \rangle$

Parameterized by predicate **allowed**

G

$\{\text{tx } c_2, \sigma_2\}$

L_2

Thread

$\langle m_2, \sigma_2, \sigma_2, l_2 \rangle$

There is no state.

$\{\text{tx } c_1, \sigma_1\}$

Thread

L_1

$\langle m_1, \sigma_1, \sigma_1, l_1 \rangle$

Parameterized by predicate **allowed**

G

Closed under log prefix

$\{\text{tx } c_2, \sigma_2\}$

Thread

L_2

$\langle m_2, \sigma_2, \sigma_2, l_2 \rangle$

There is no state.

$\{tx\ c_1, \sigma_1\}$

Thread

L_1

$\langle m_1, \sigma_1, \sigma_1, l_1 \rangle$

Parameterized by predicate **allowed**

G

Closed under log prefix

$\{tx\ c_2, \sigma_2\}$

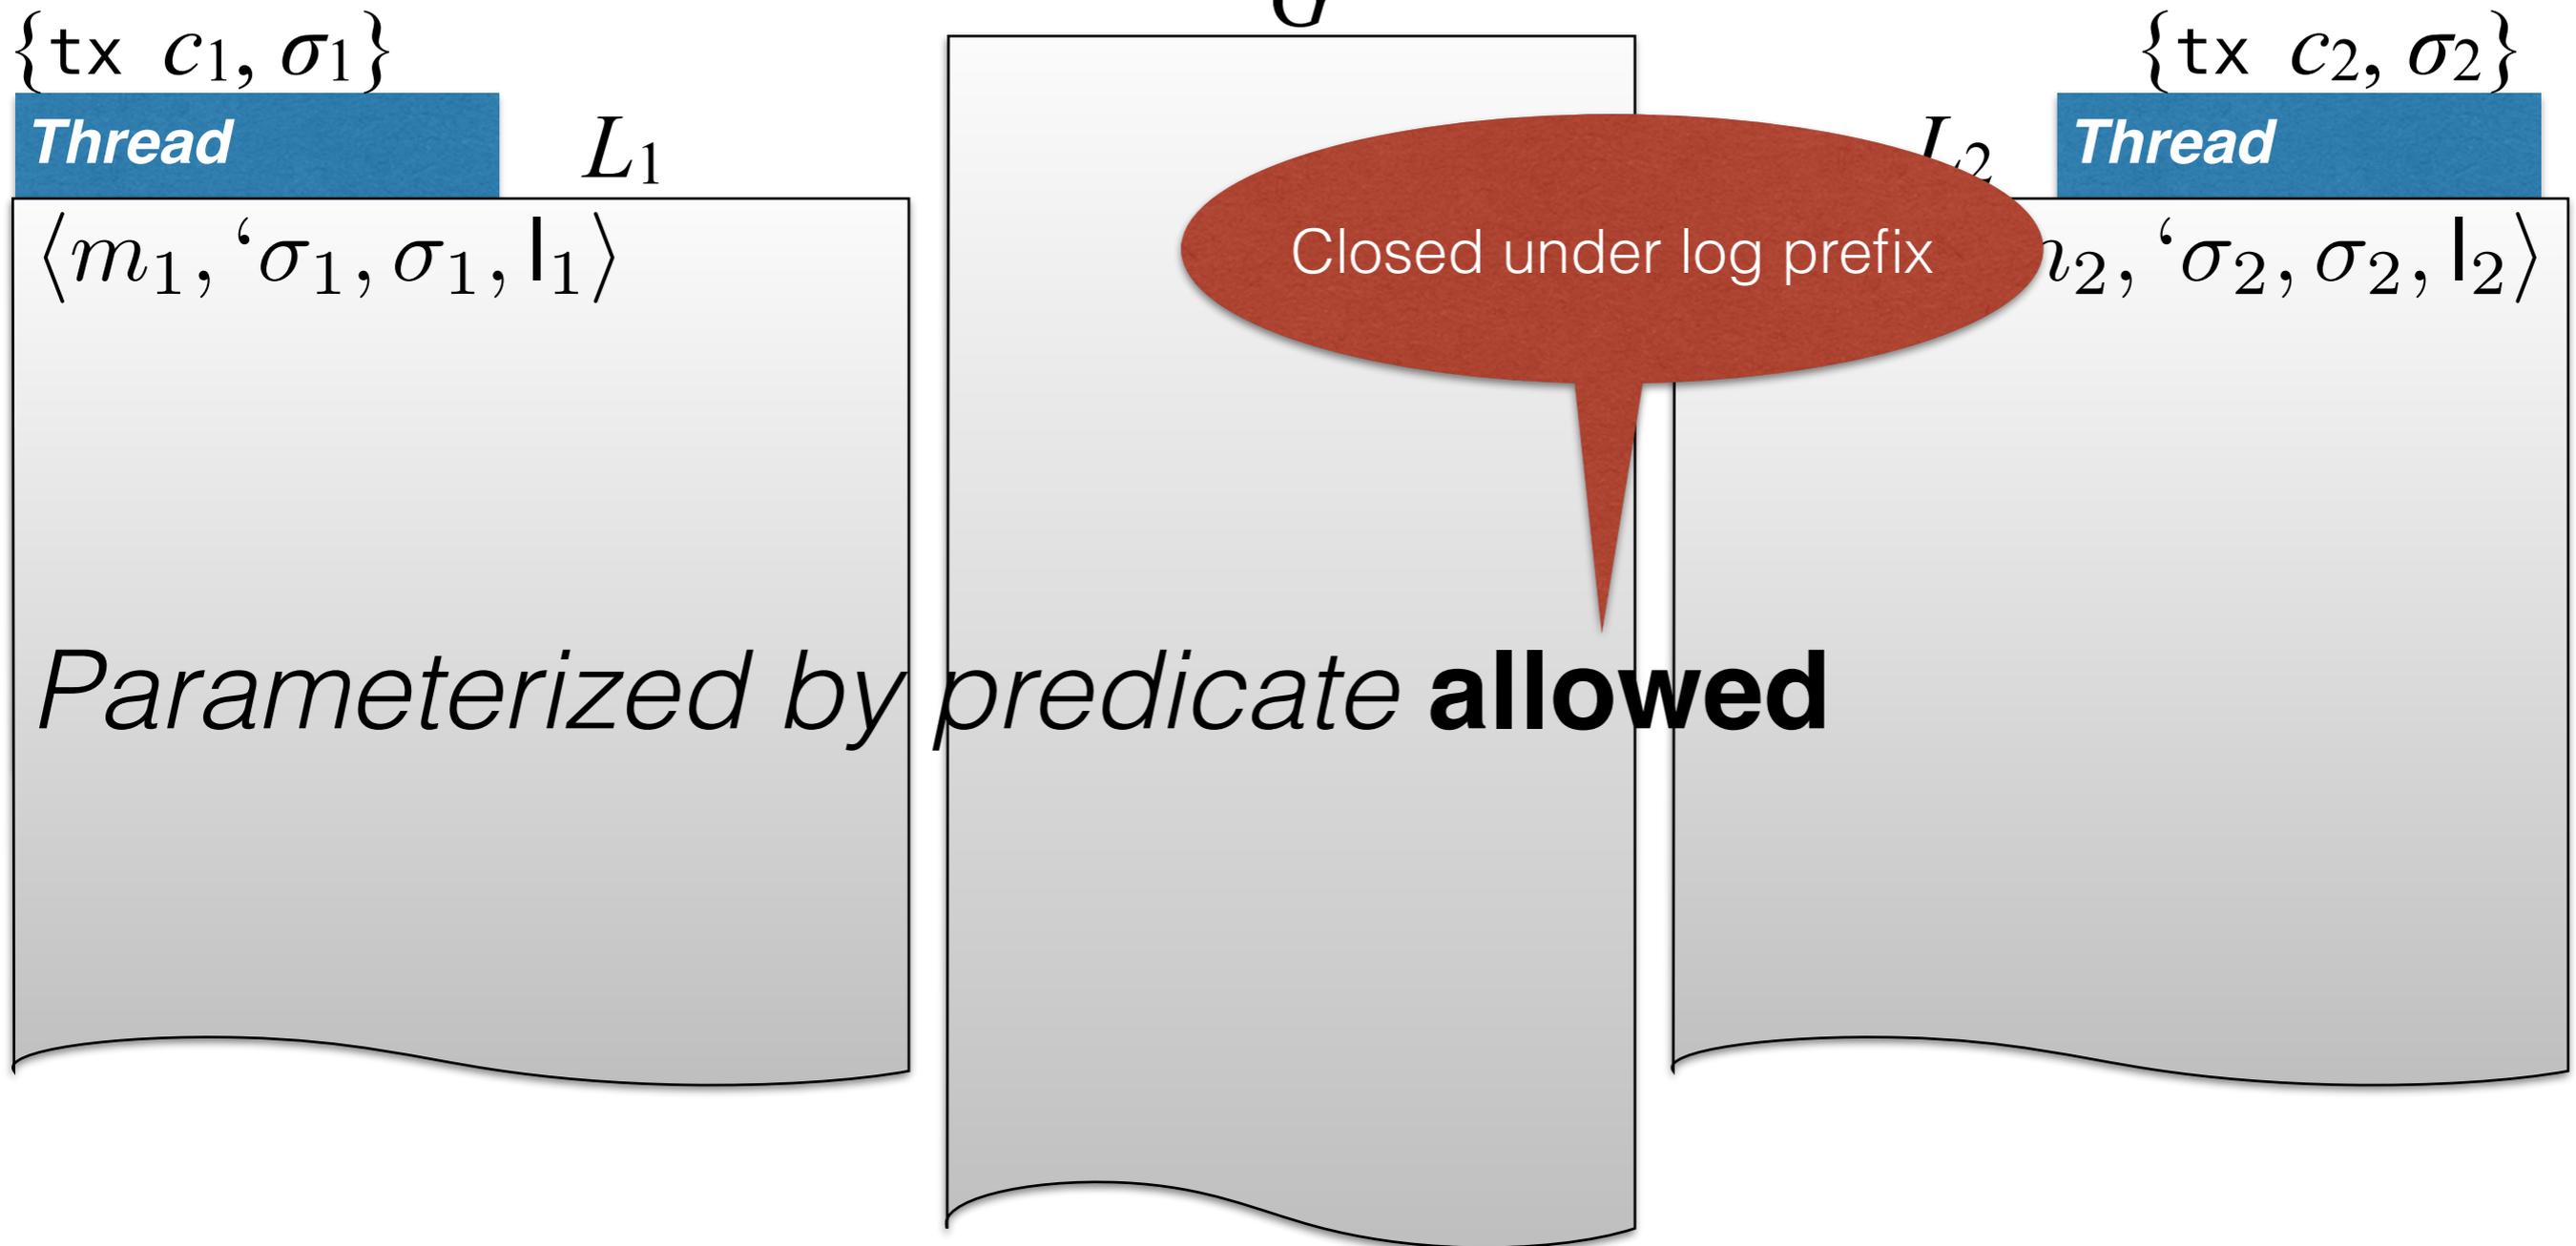
Thread

L_2

$\langle m_2, \sigma_2, \sigma_2, l_2 \rangle$

Log Equality

There is no state.



Log Equality

$$\frac{\text{allowed } l_1 \Rightarrow \text{allowed } l_2 \quad \forall op. (l_1 \cdot op) \preceq (l_2 \cdot op)}{l_1 \preceq l_2}$$

There is no state.

$\{tx\ c_1, \sigma_1\}$

Thread

L_1

$\langle m_1, \sigma_1, \sigma_1, l_1 \rangle$

$\{tx\ c_2, \sigma_2\}$

L_2

Thread

$\langle m_2, \sigma_2, \sigma_2, l_2 \rangle$

7 Simple Rules:

**Apply Push Pull Unpull Unpush Unapply
Commit**

The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

Thread

L_1

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

$\langle q.enq('a'), \sigma', \sigma'' \rangle$

G

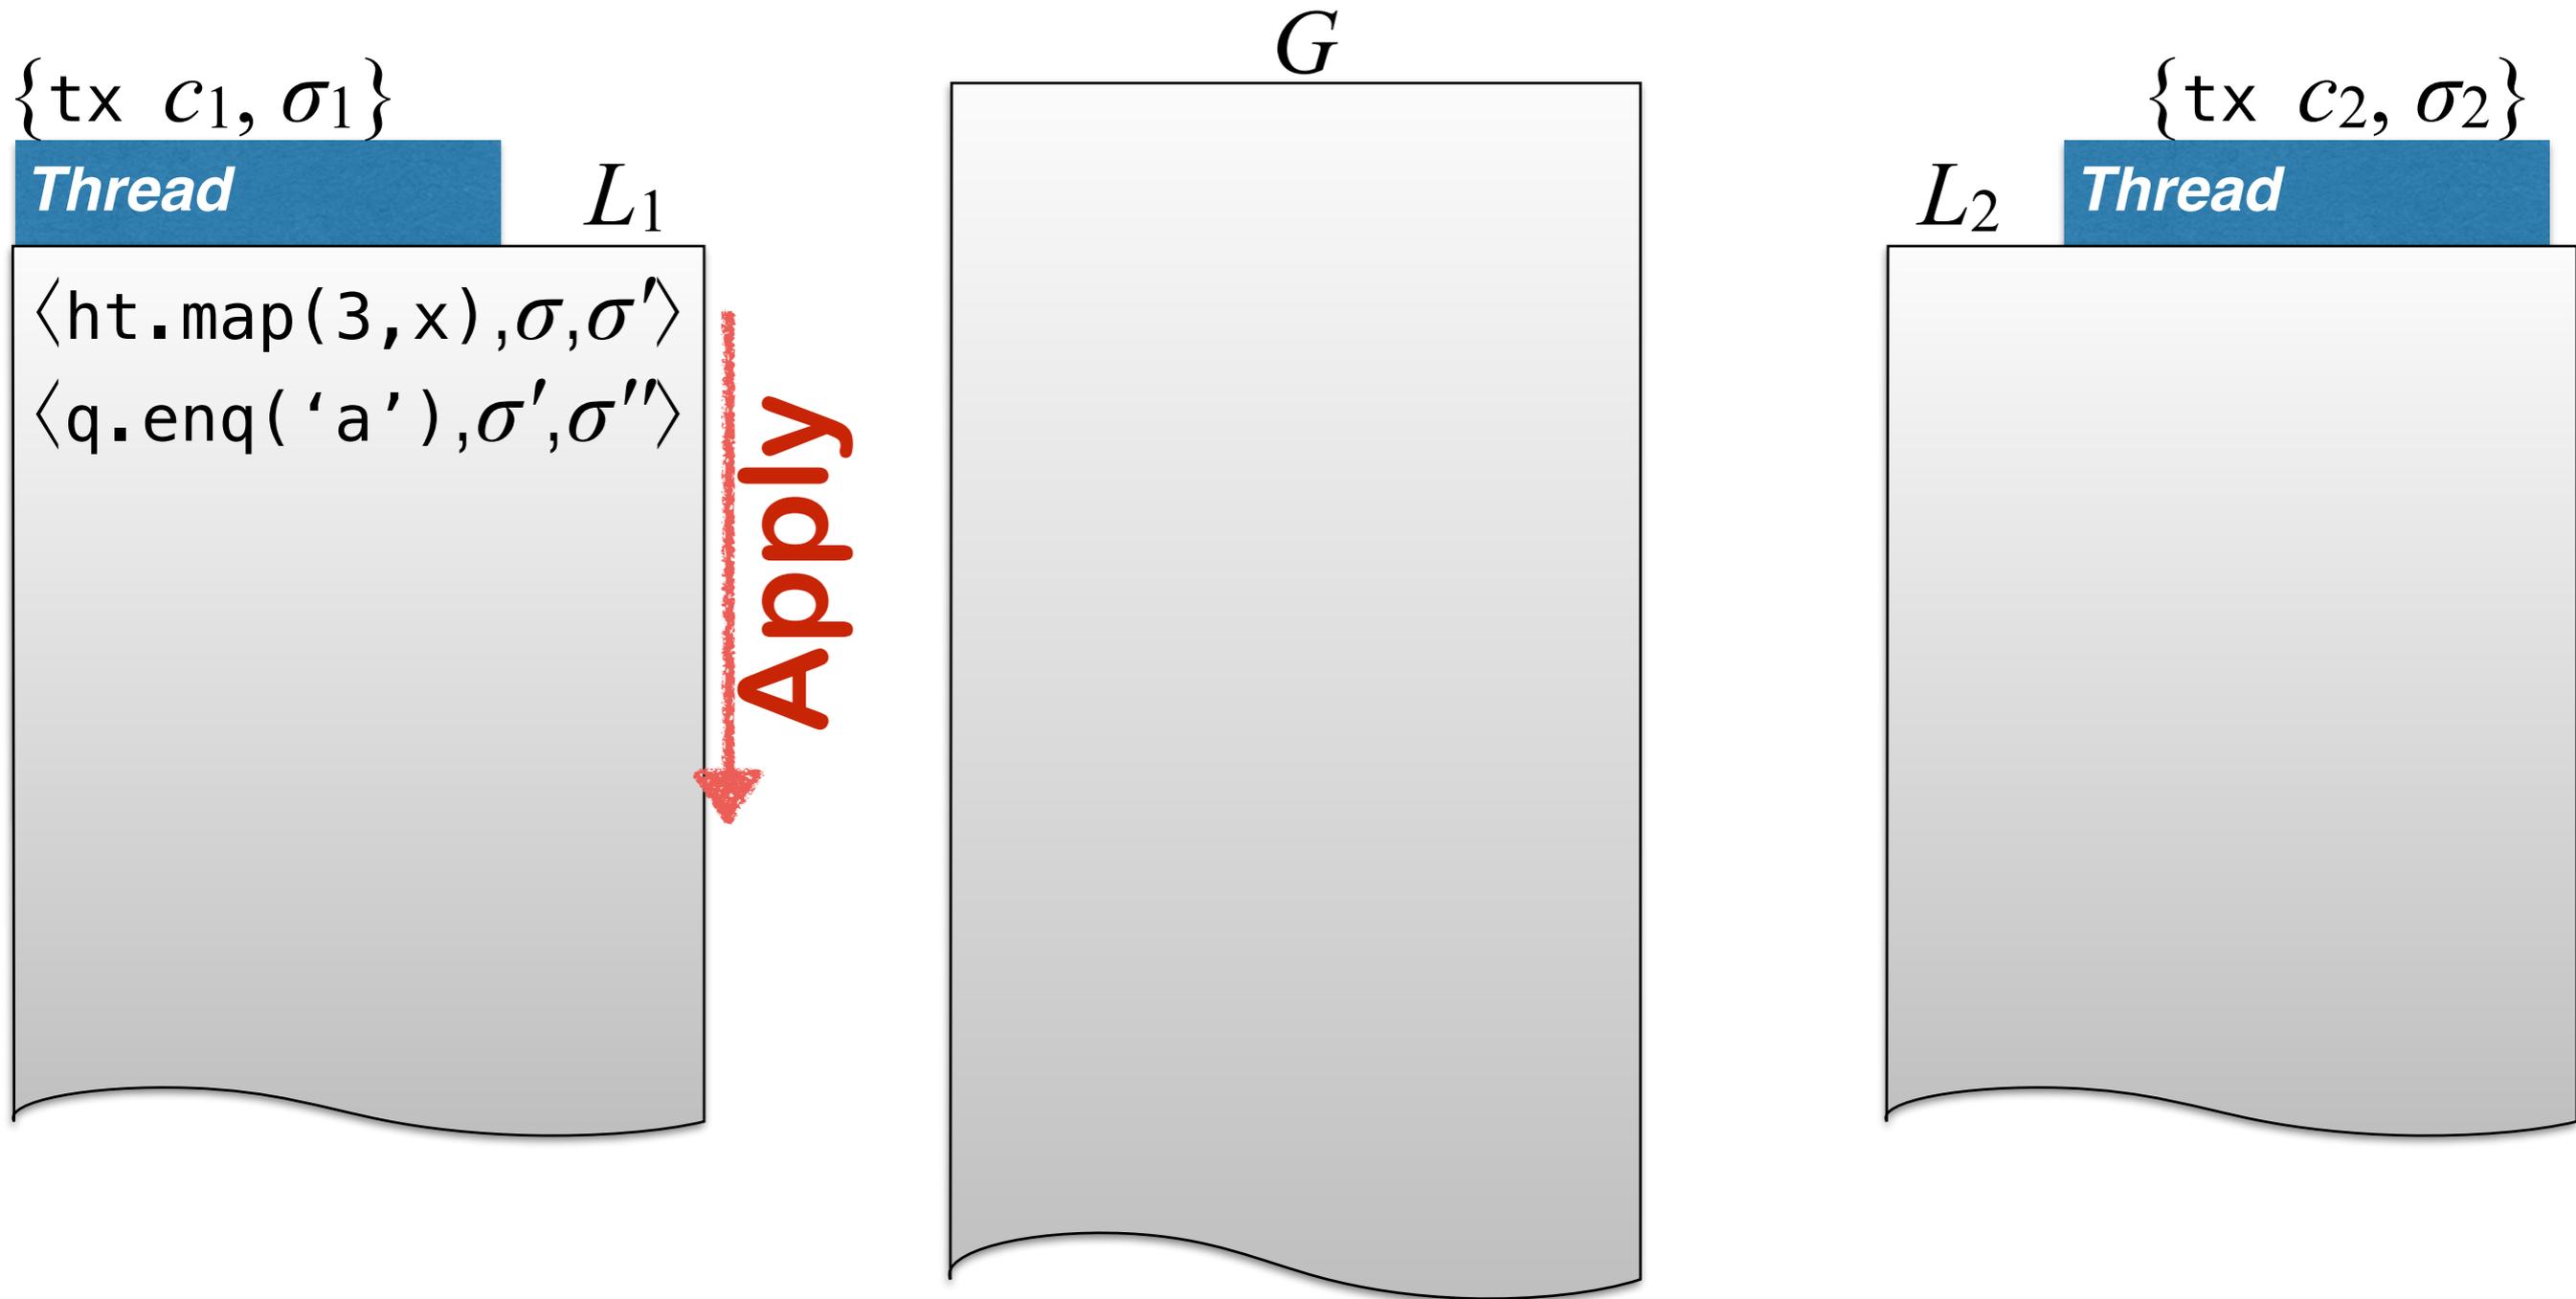
$\{tx\ c_2, \sigma_2\}$

Thread

L_2

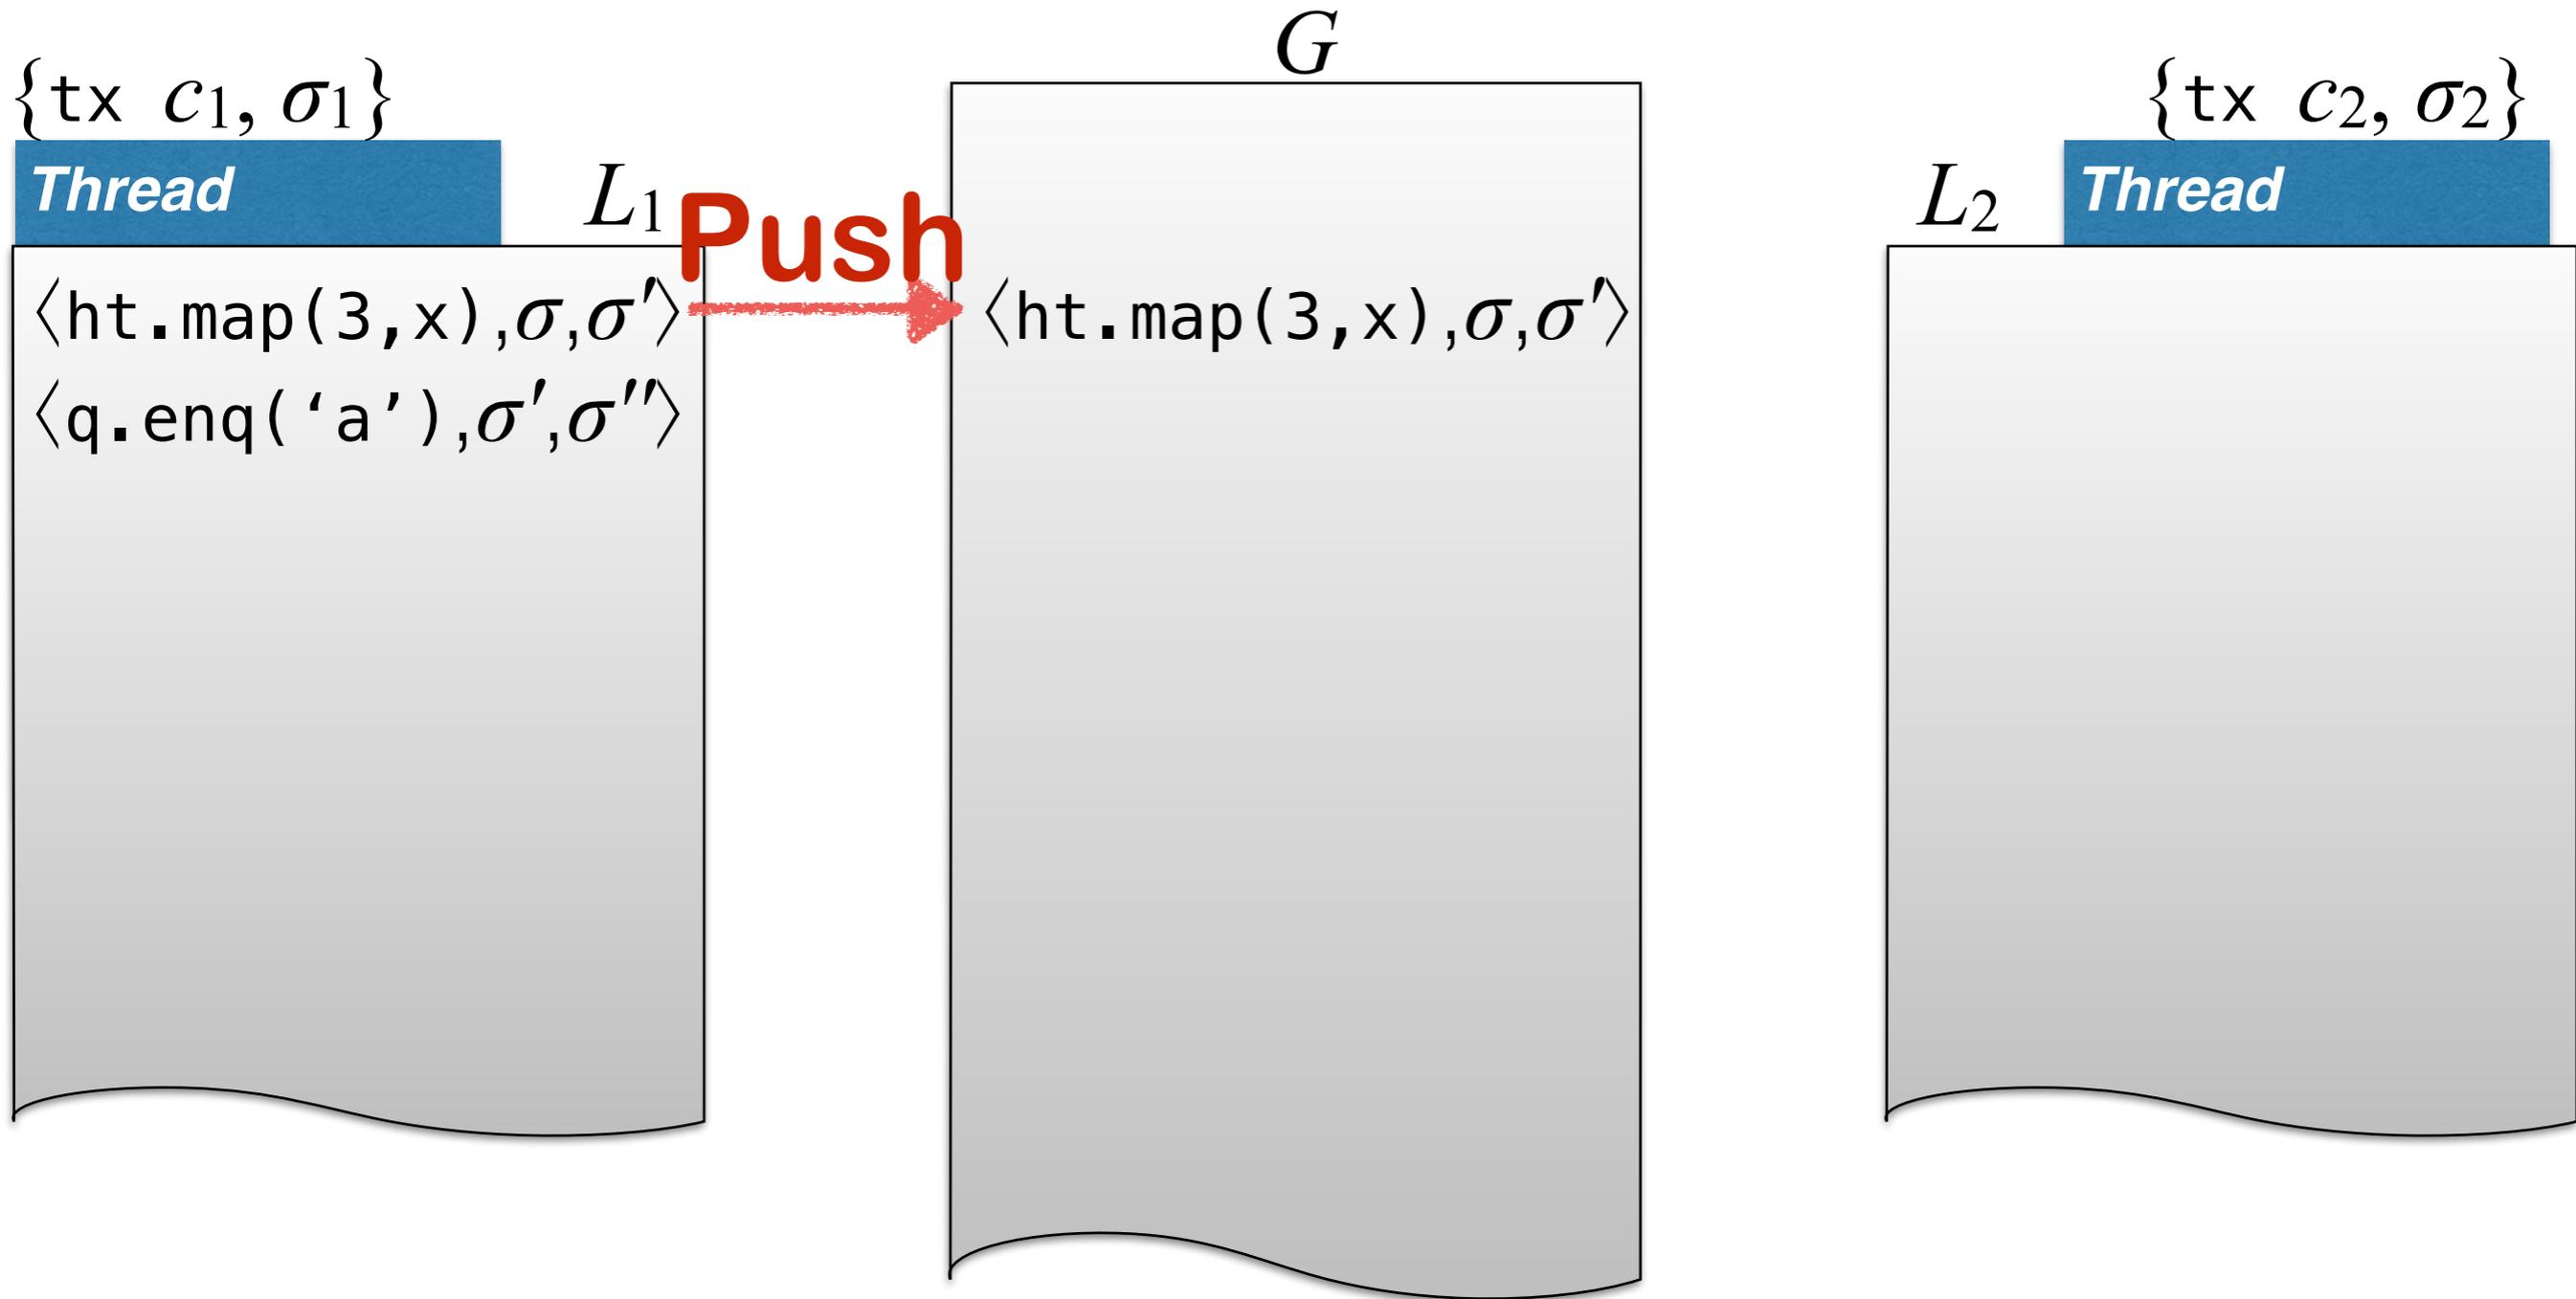
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



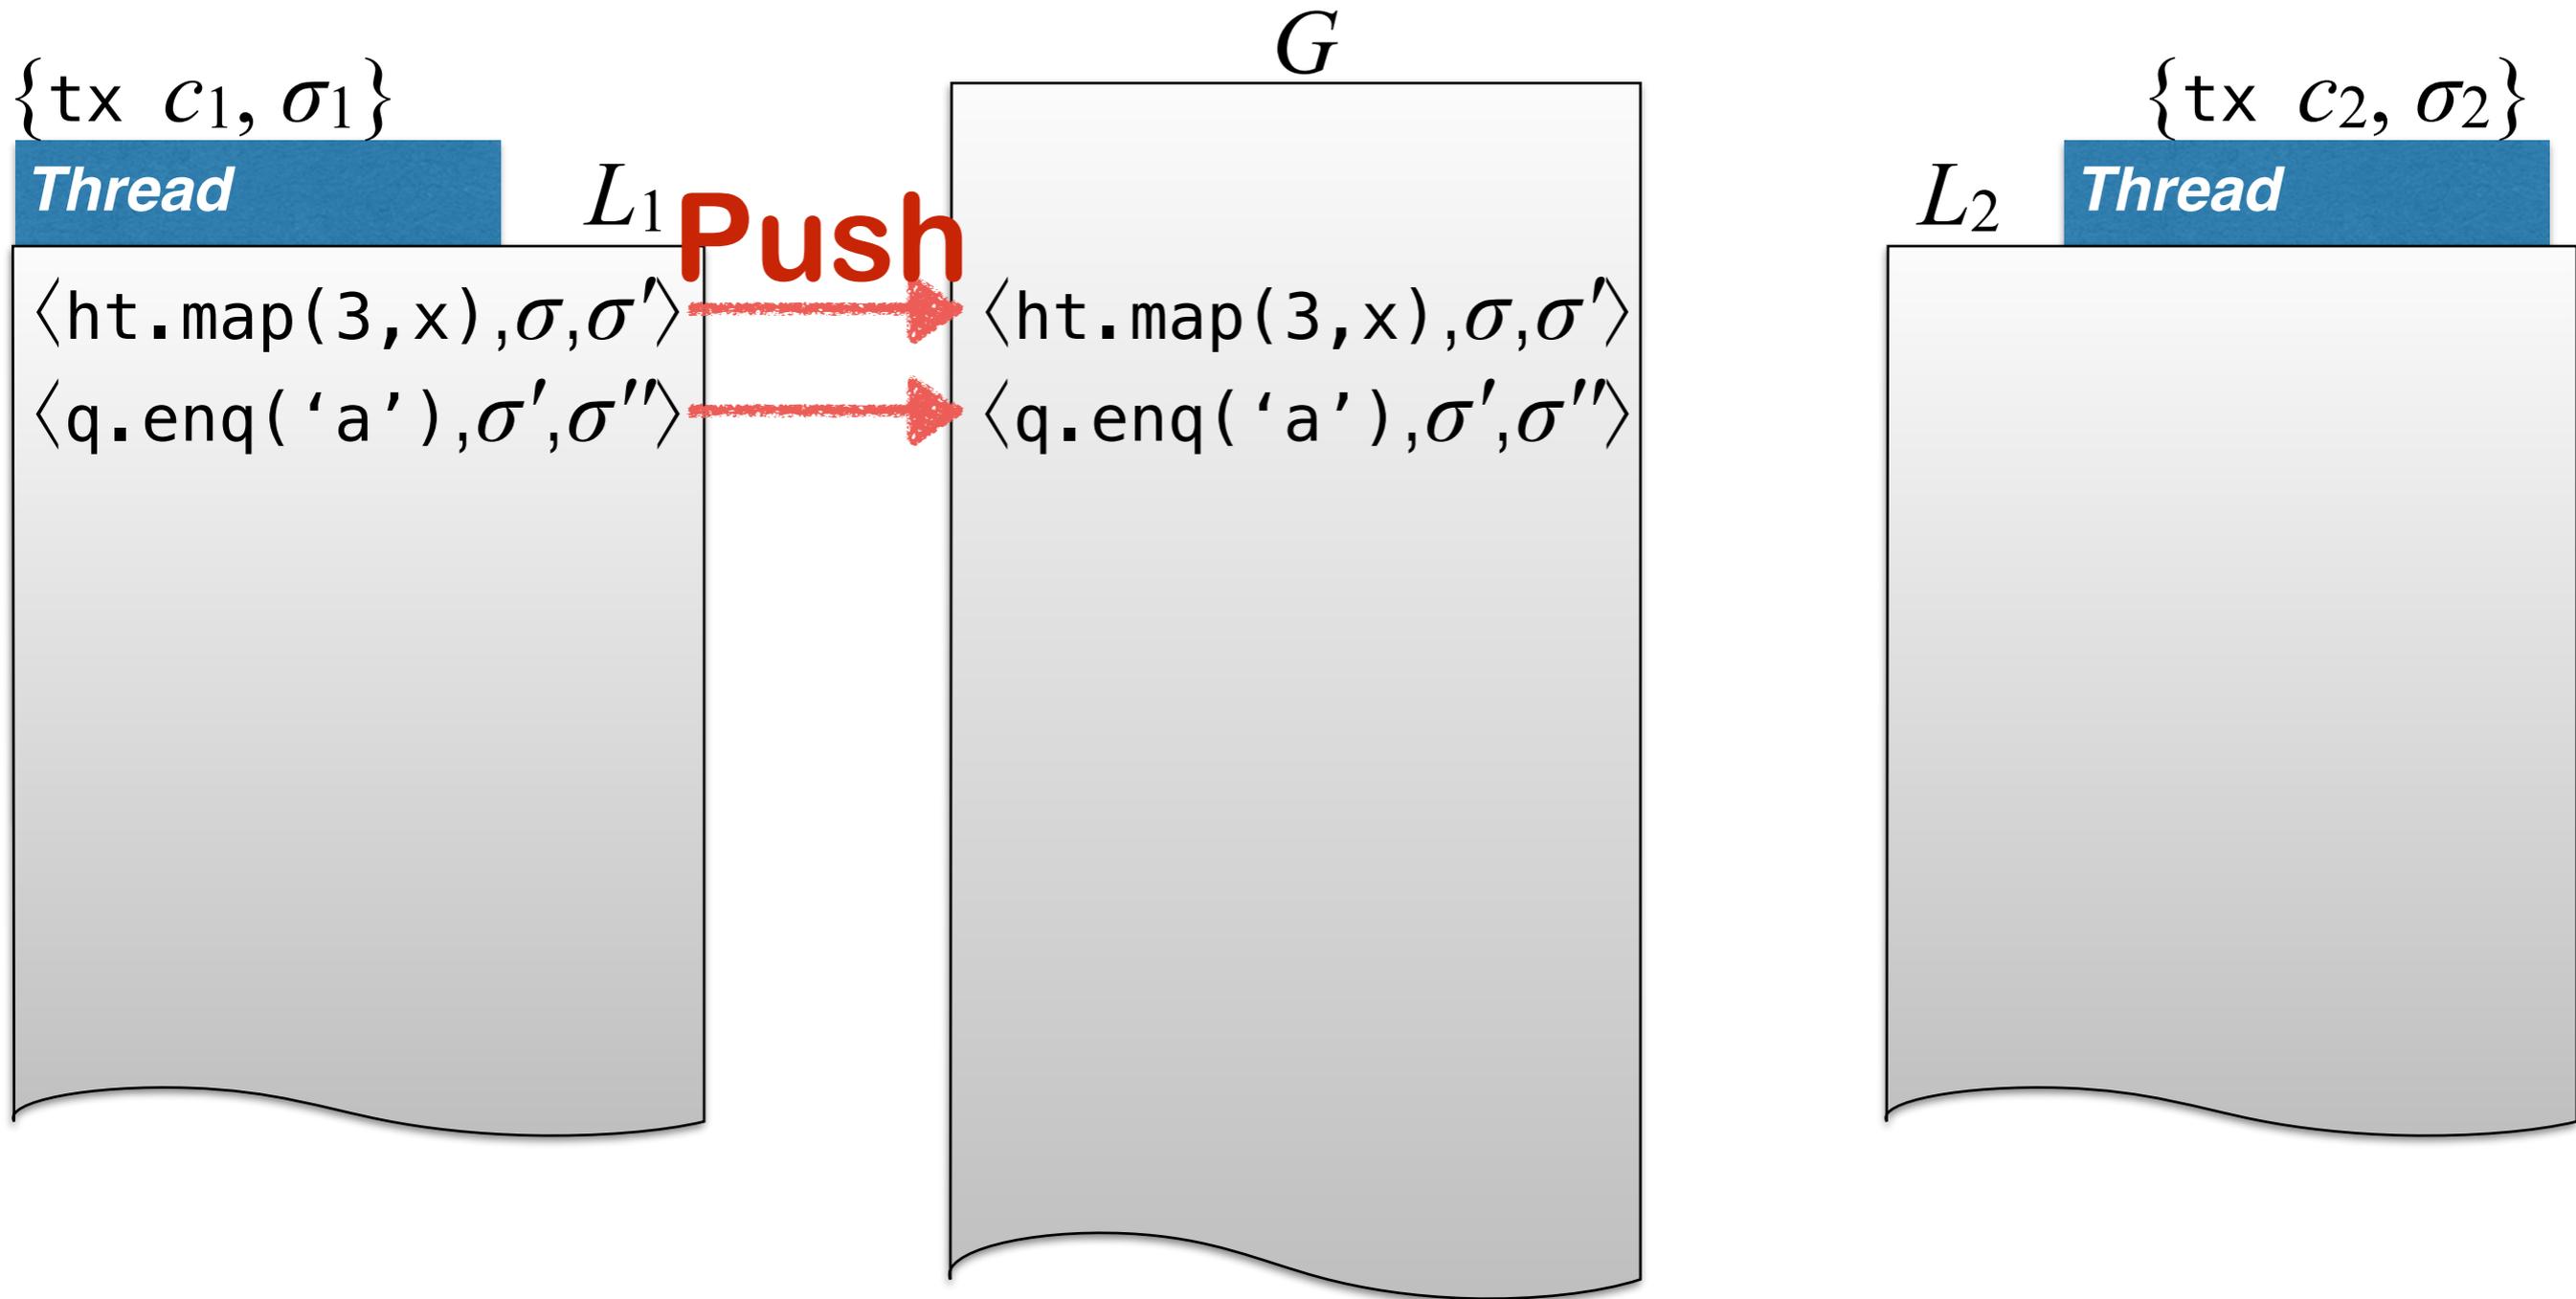
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



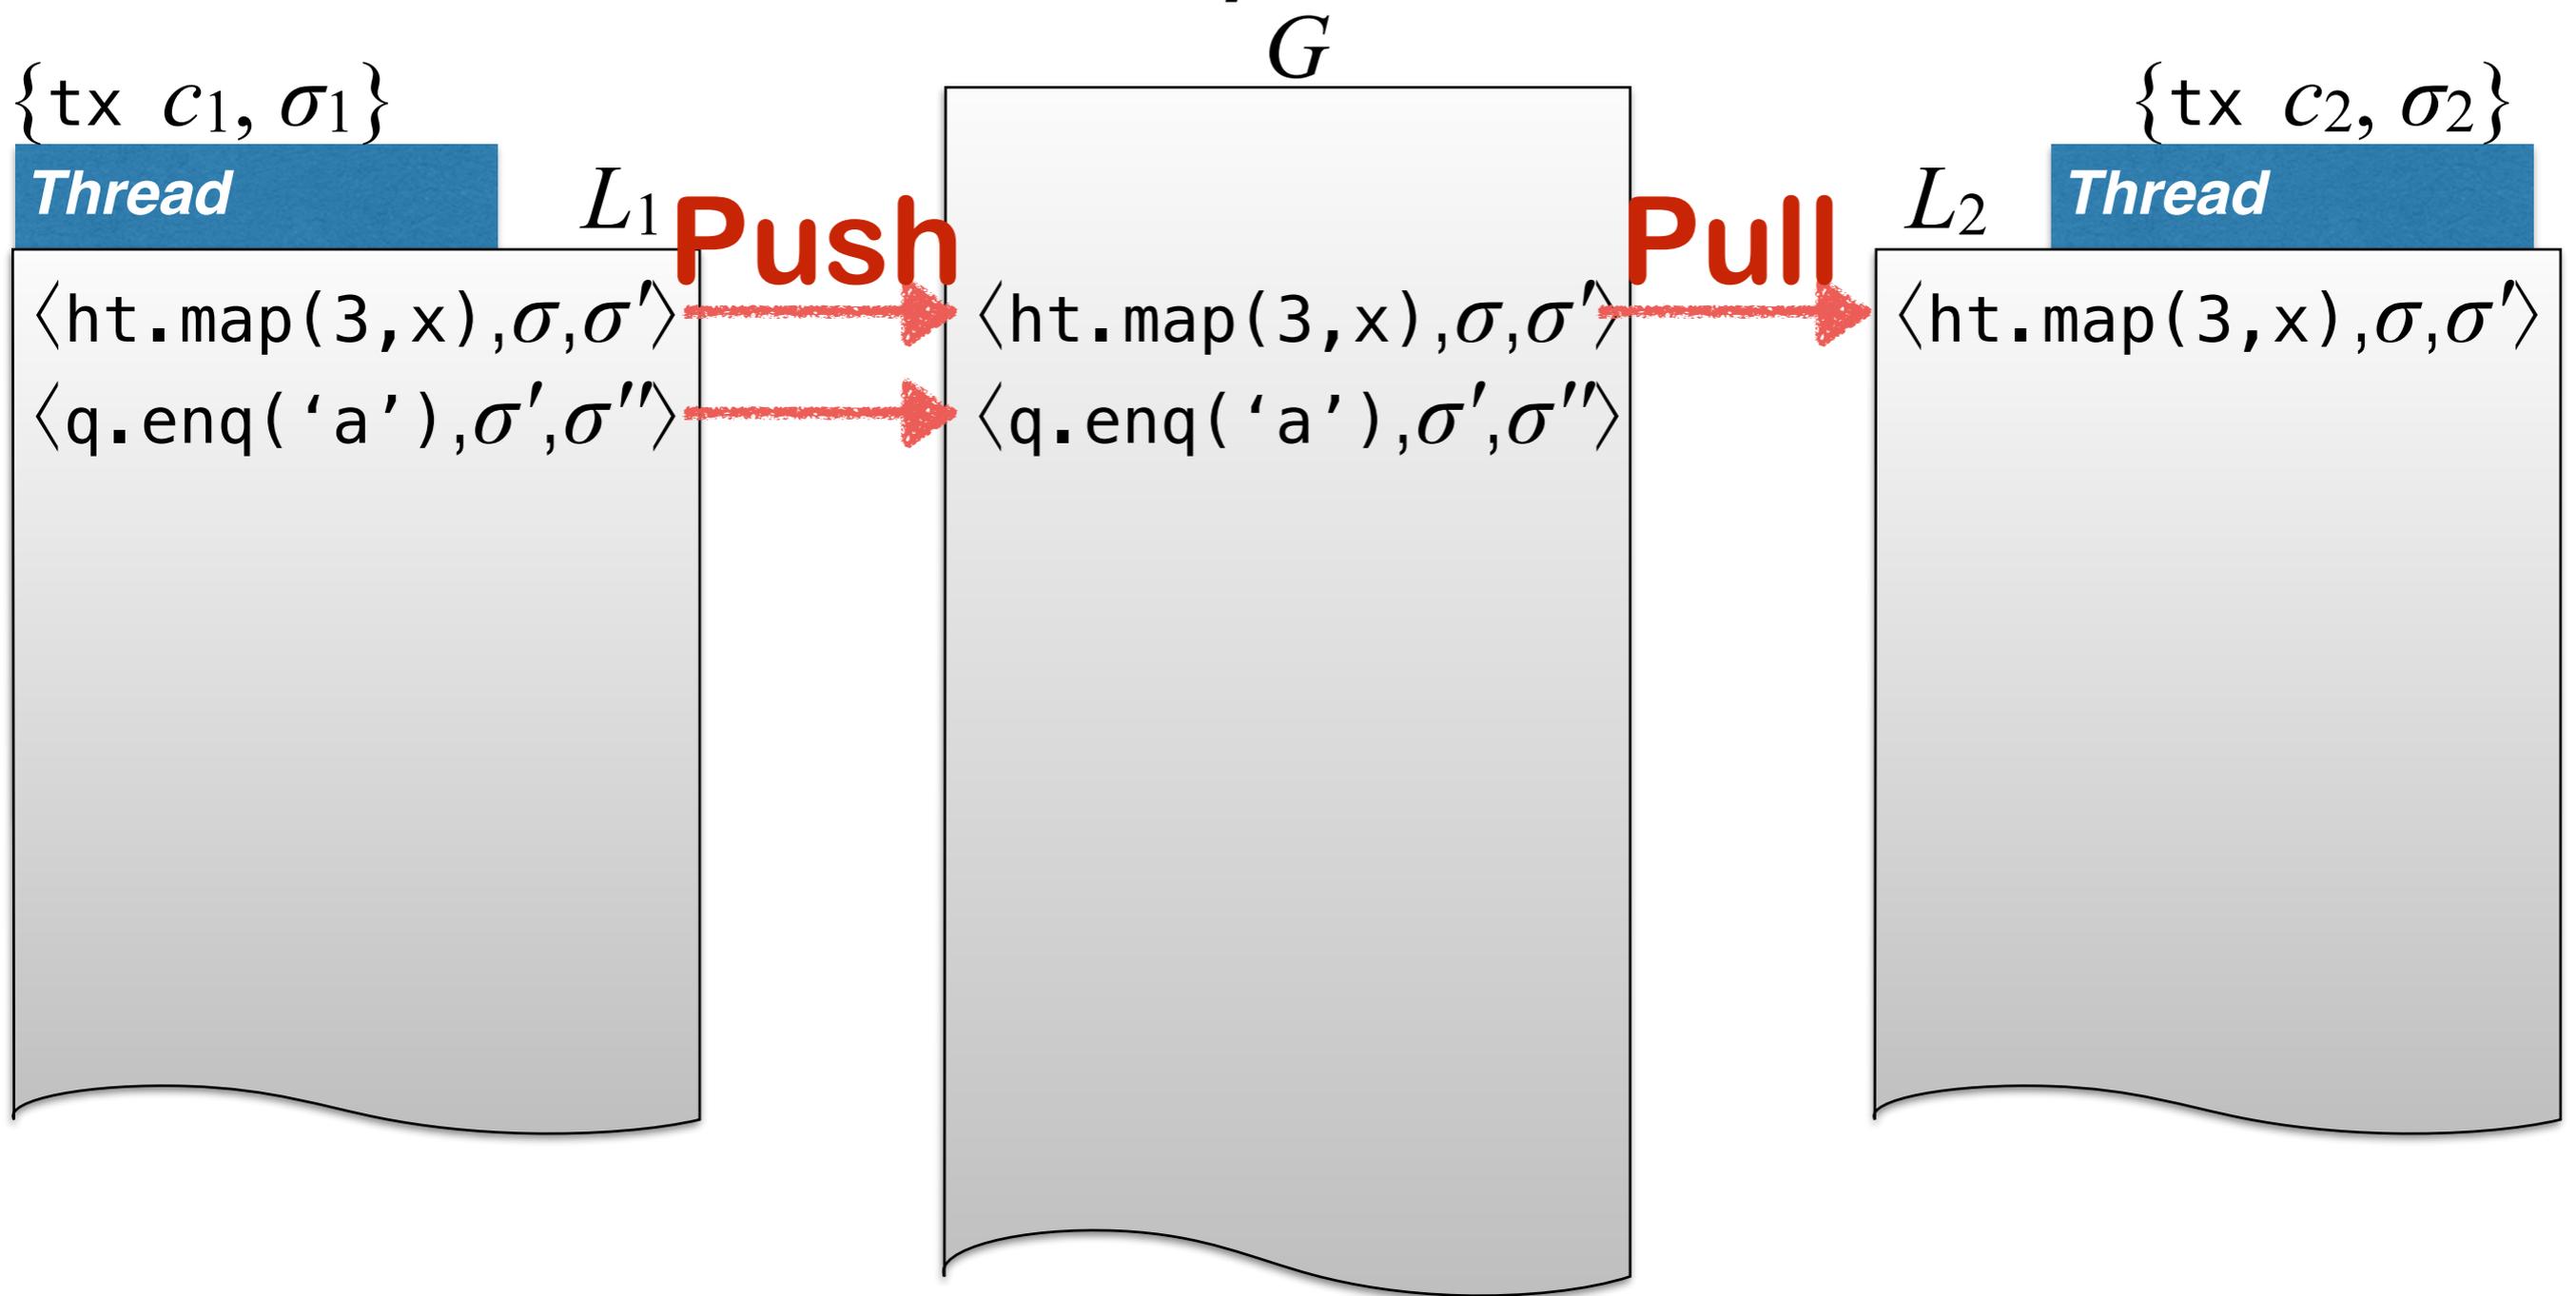
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



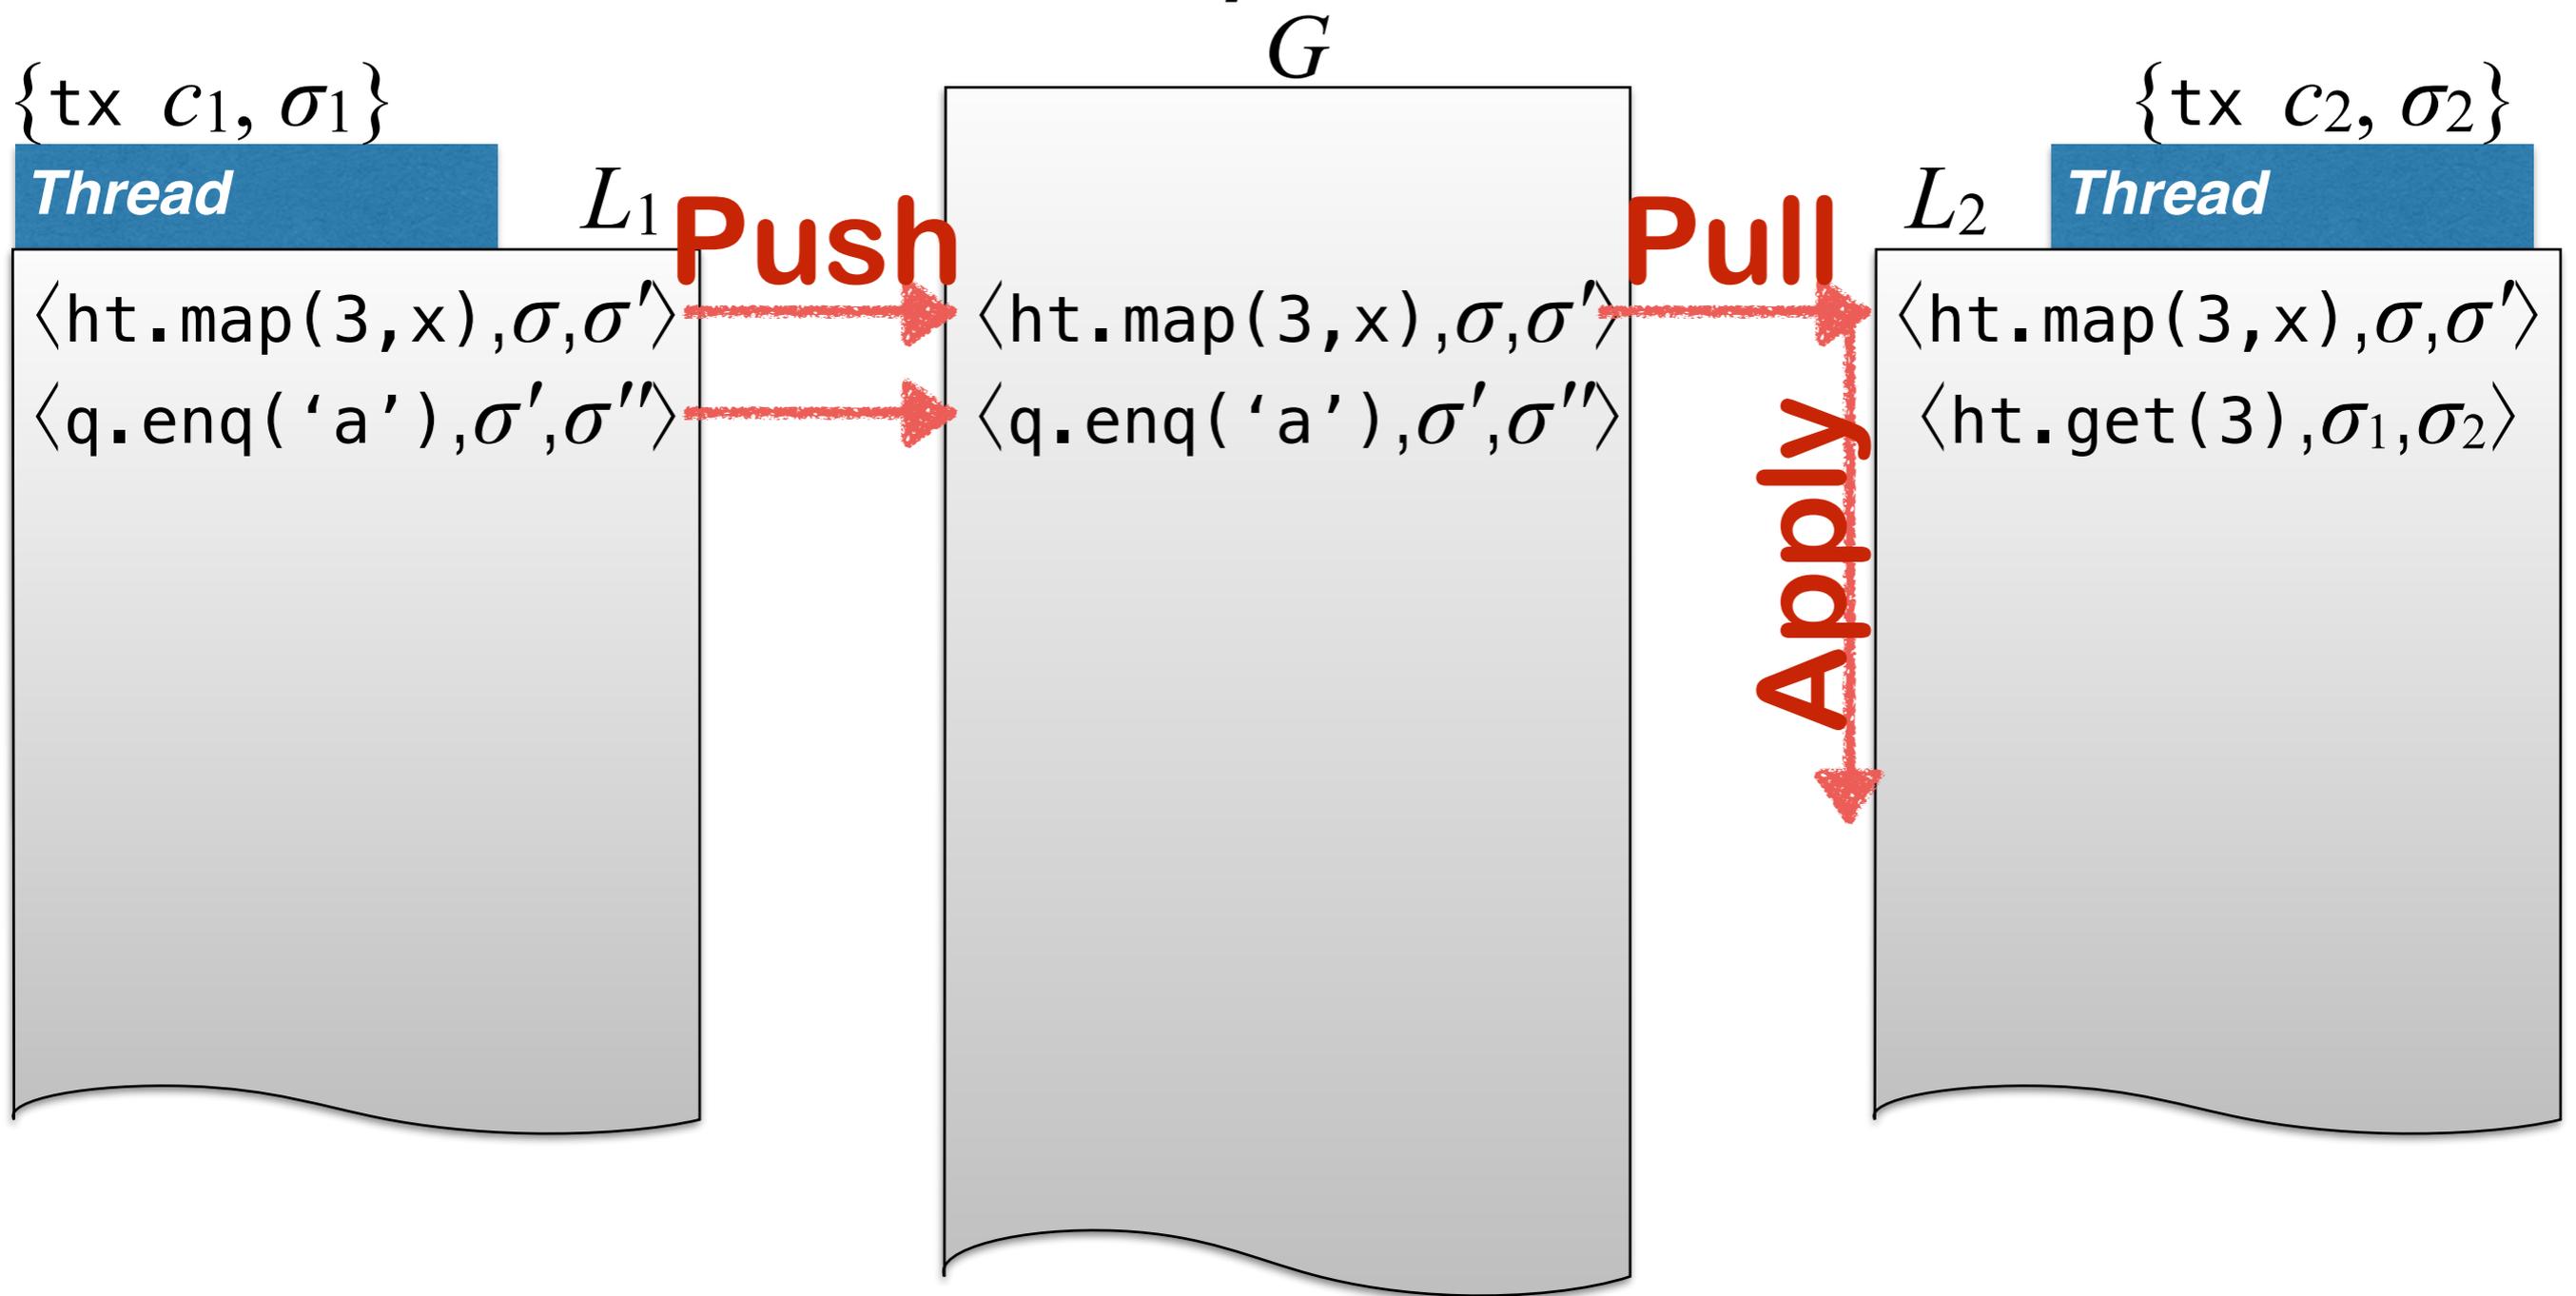
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

Thread L_1

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$

G

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$
 $\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\{tx\ c_2, \sigma_2\}$

L_2 **Thread**

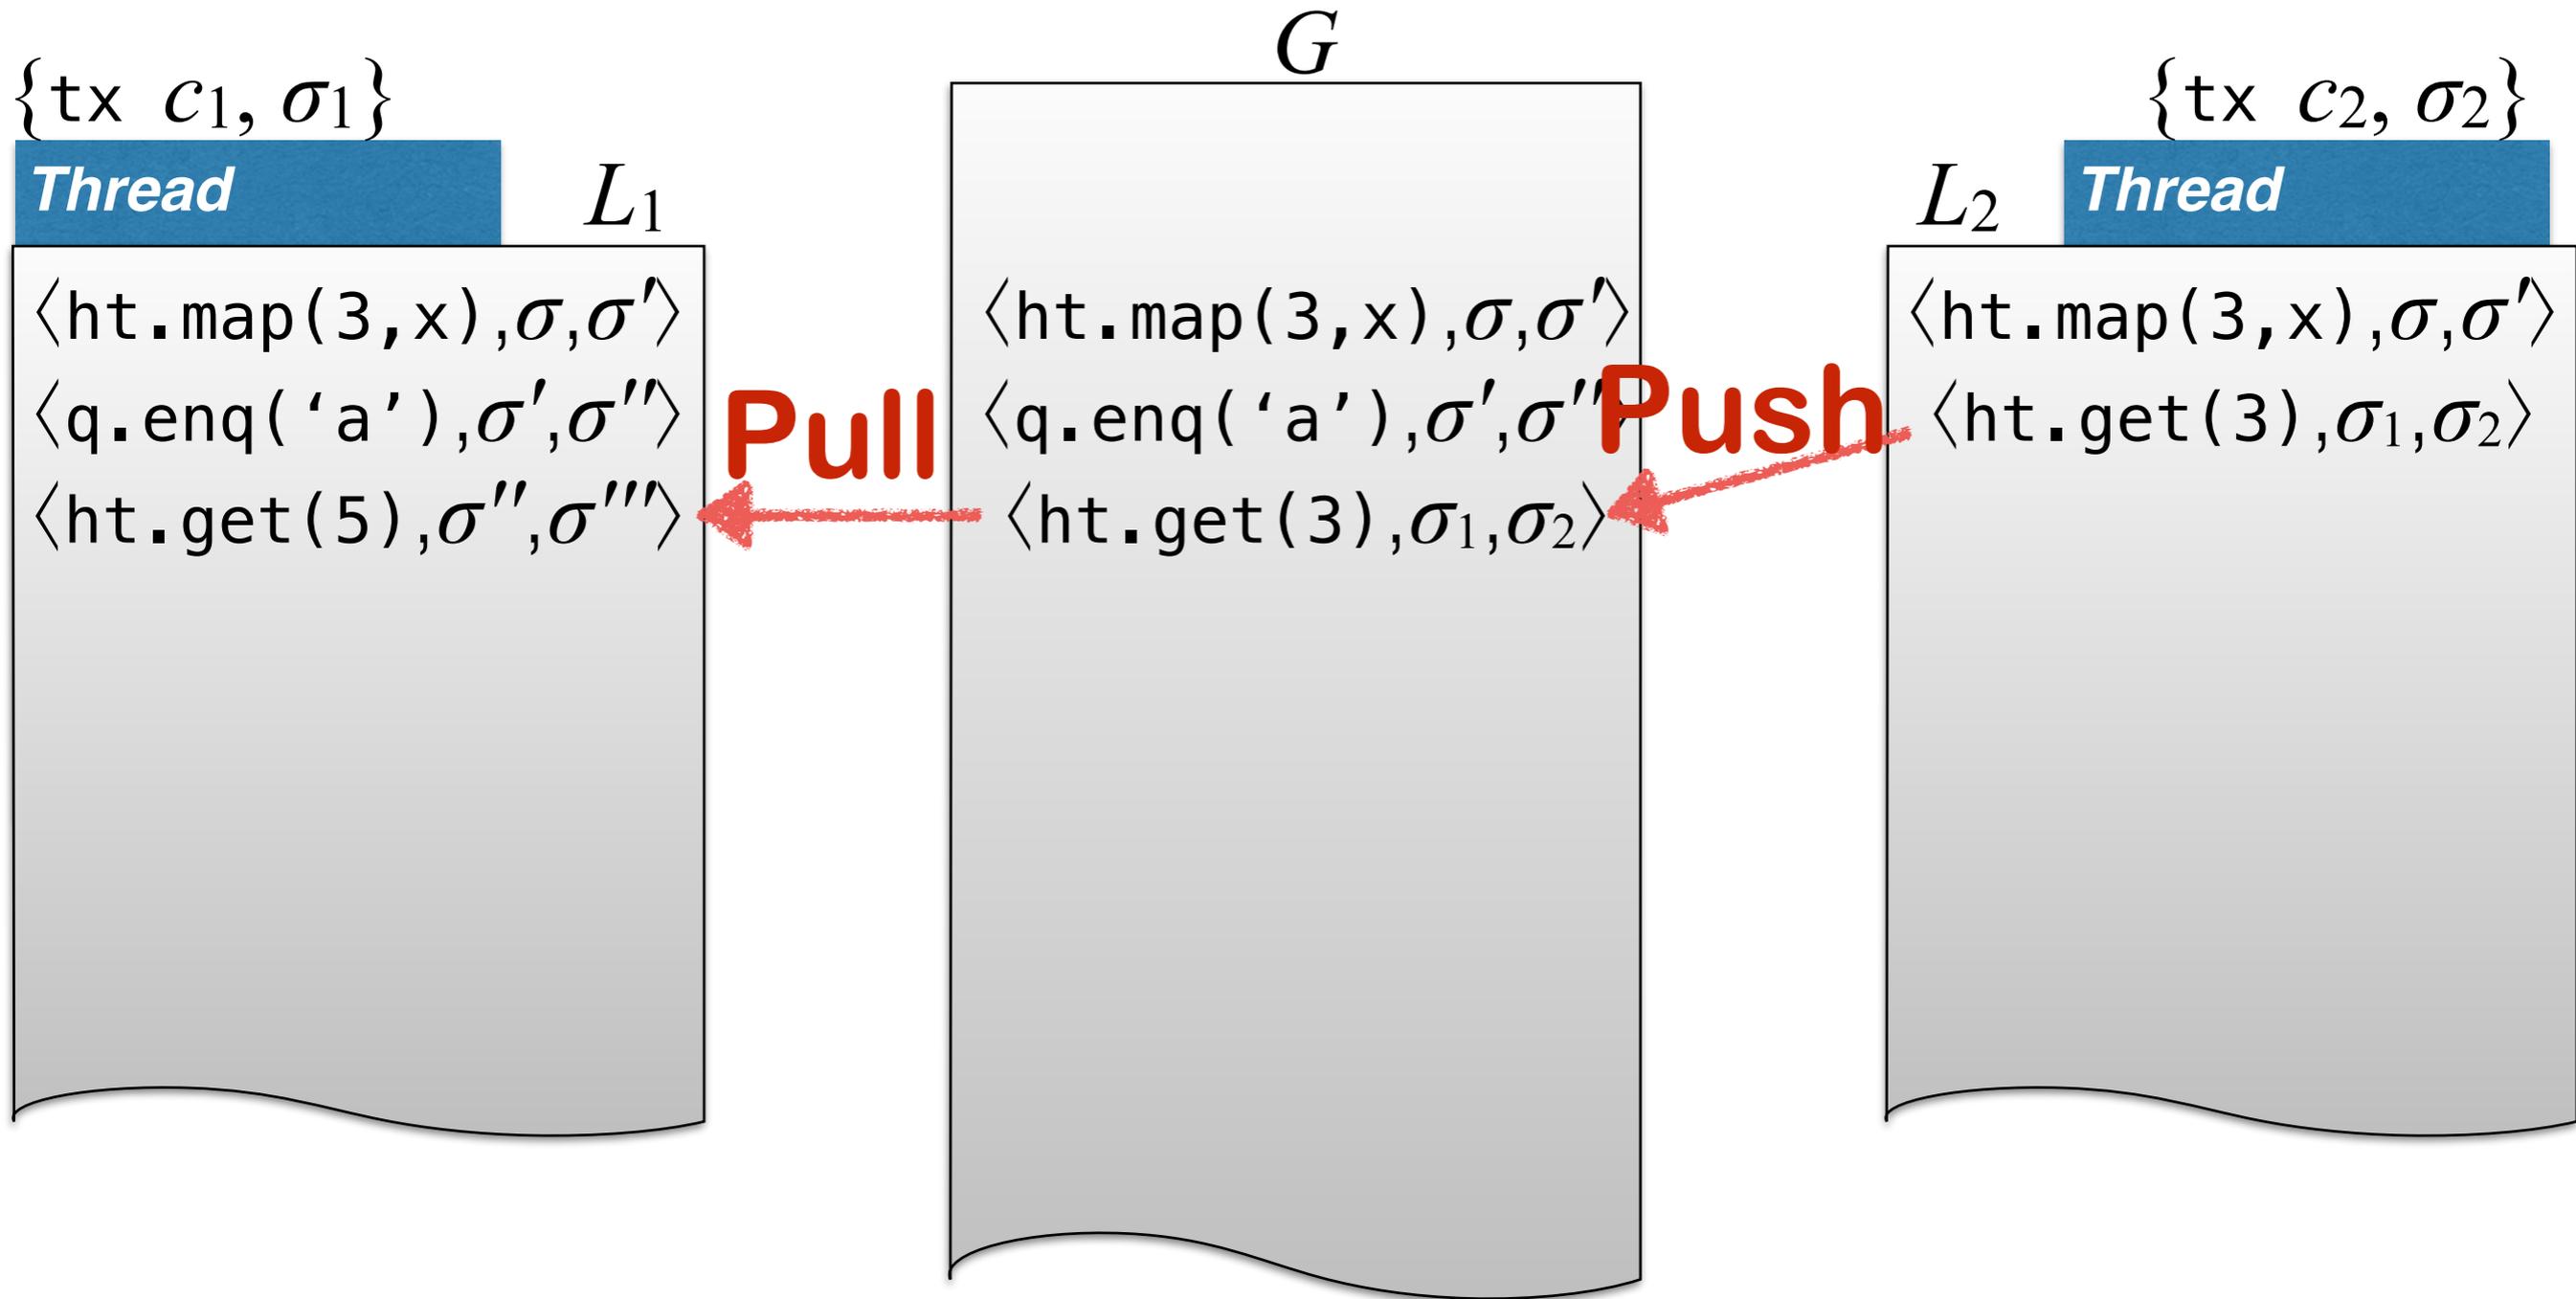
$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

Push



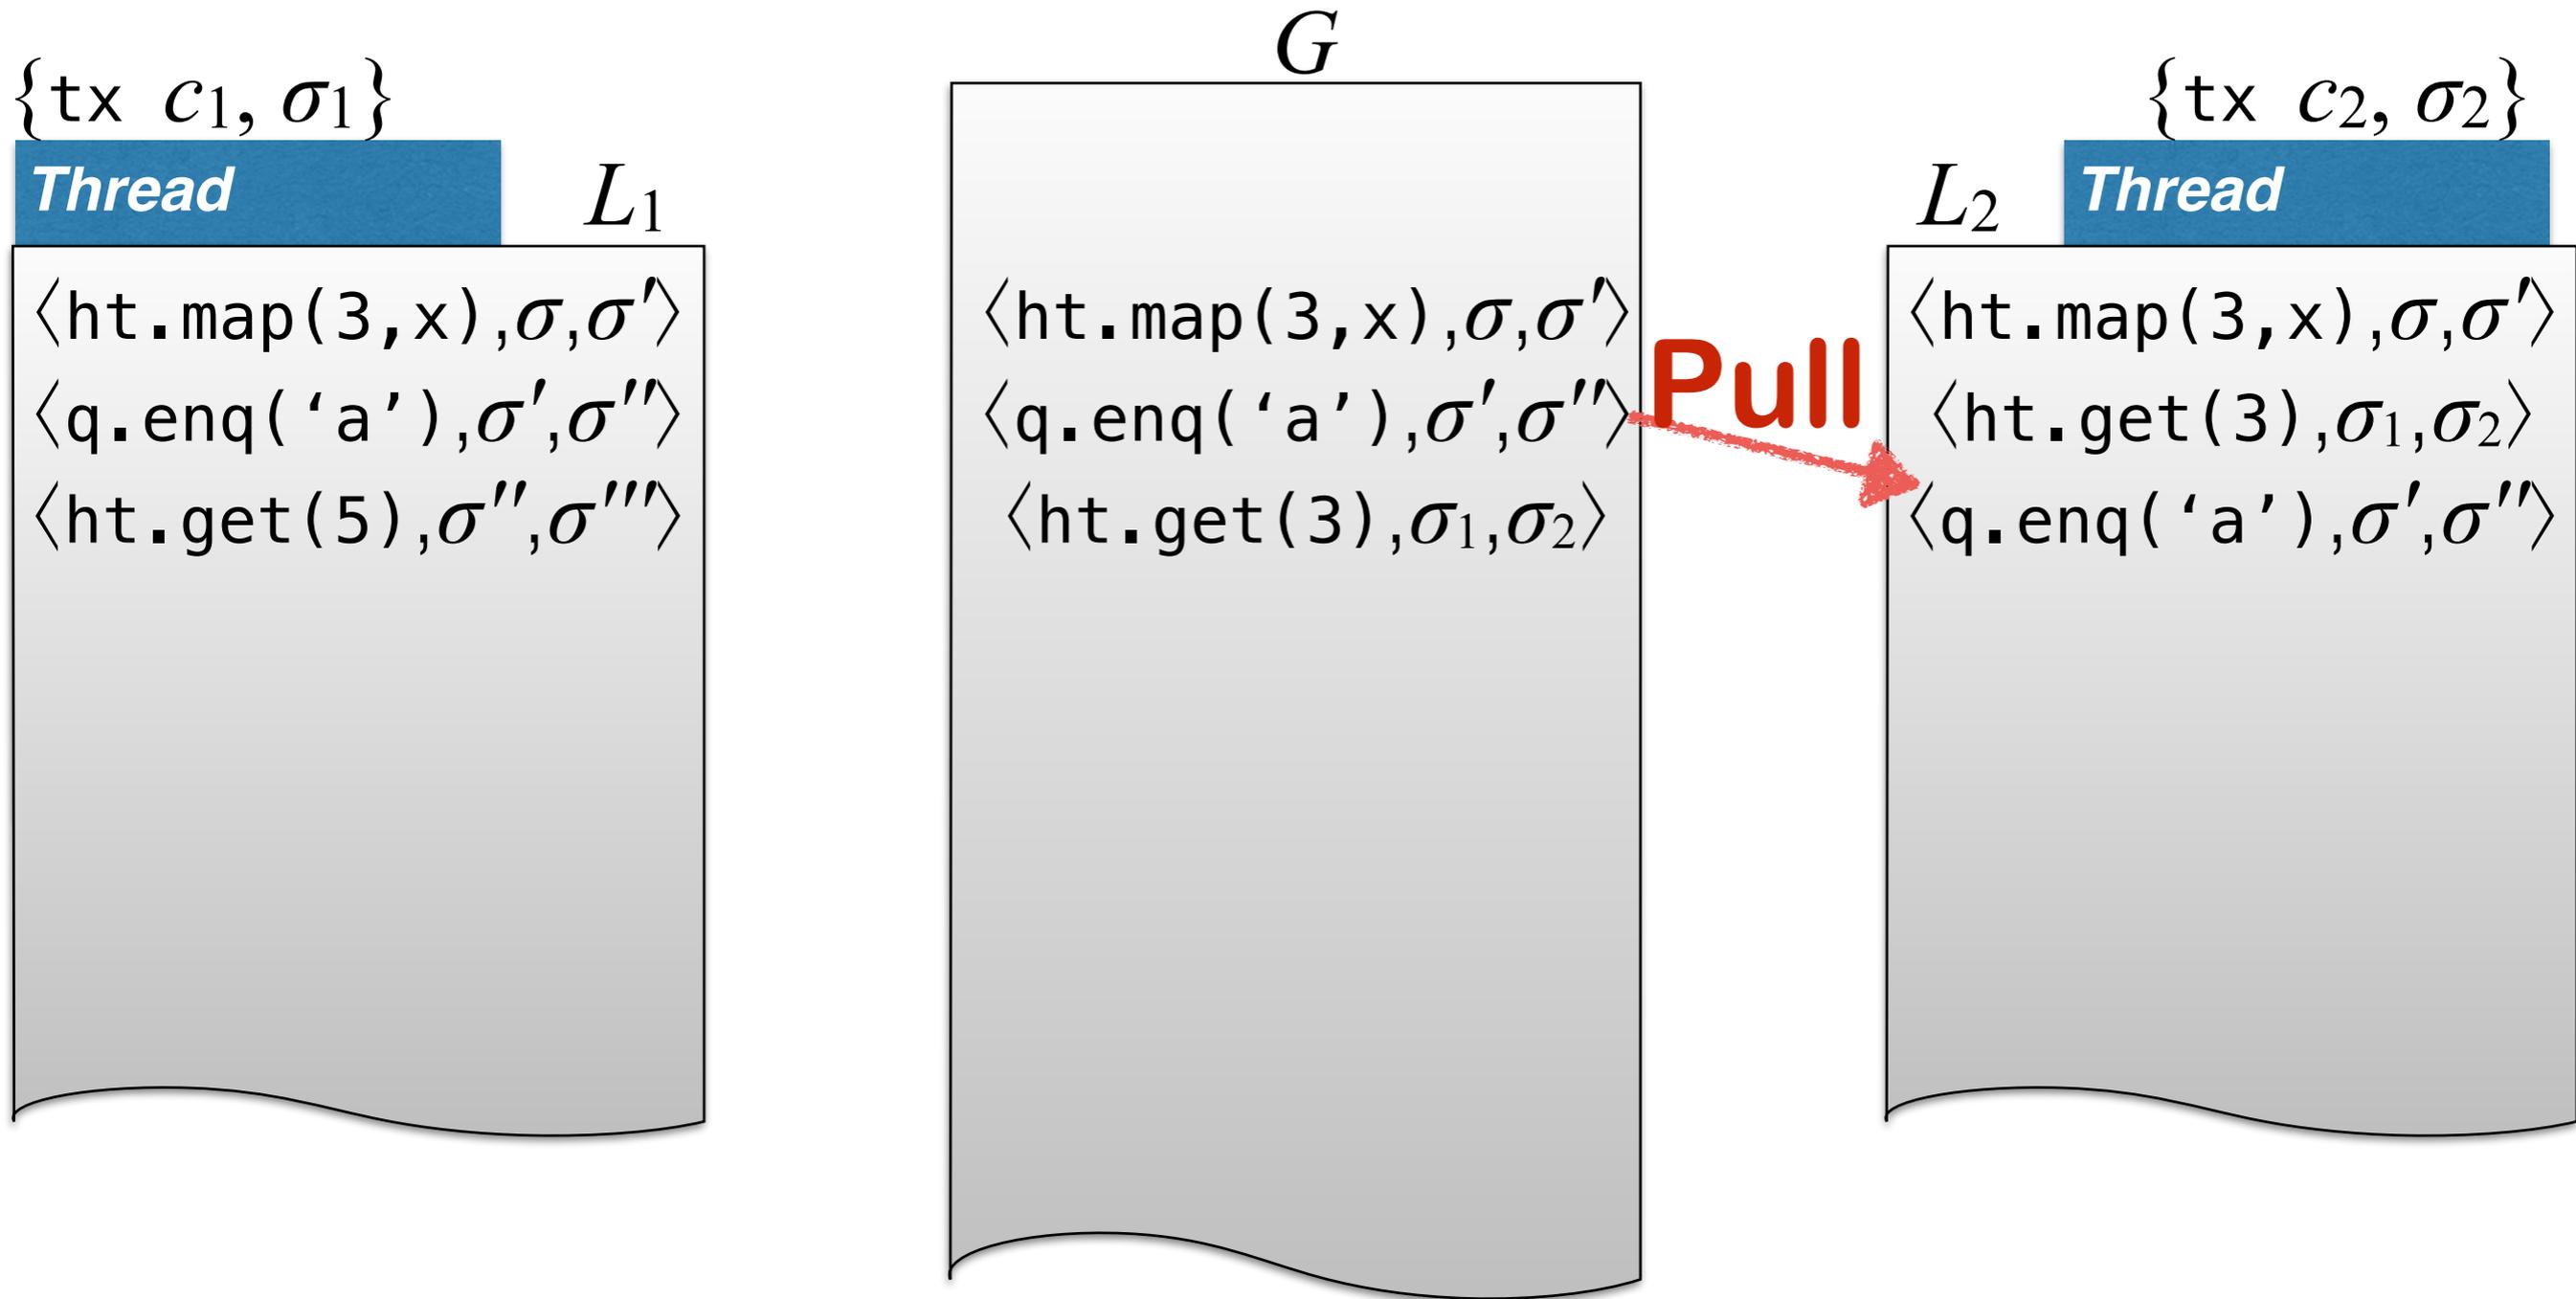
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model

{tx c_1 , σ_1 }

Thread L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$
 $\langle \text{ht.get}(5), \sigma'', \sigma''' \rangle$

Commit

G

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle$

{tx c_2 , σ_2 }

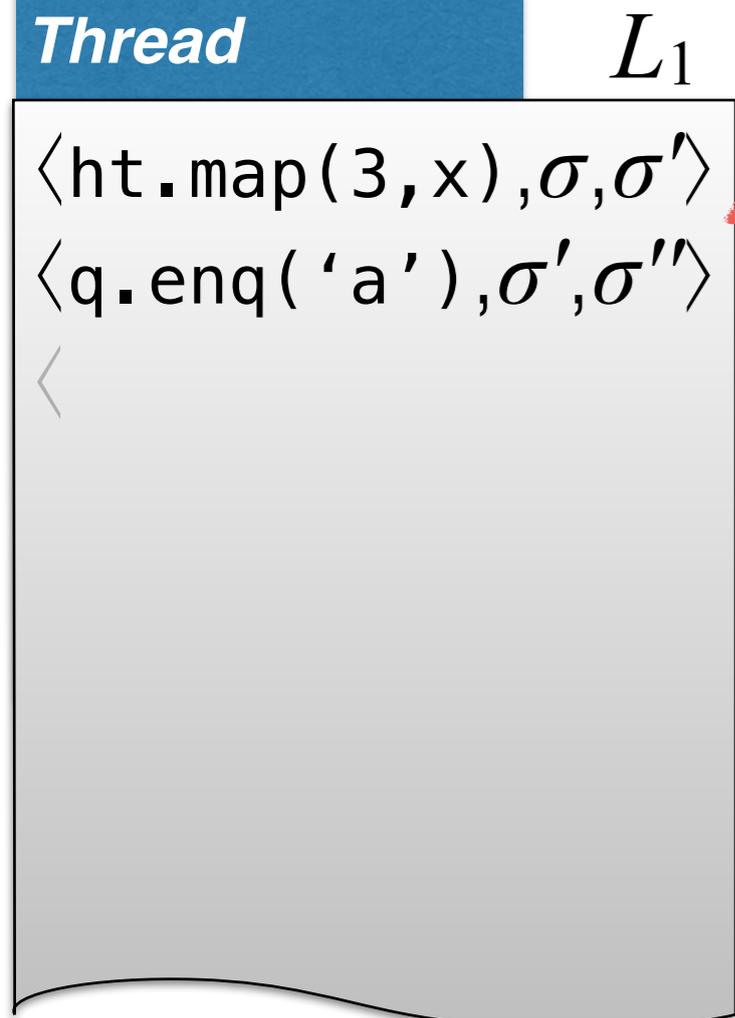
L_2 **Thread**

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply Push Pull Unpull Unpush Unapply Commit

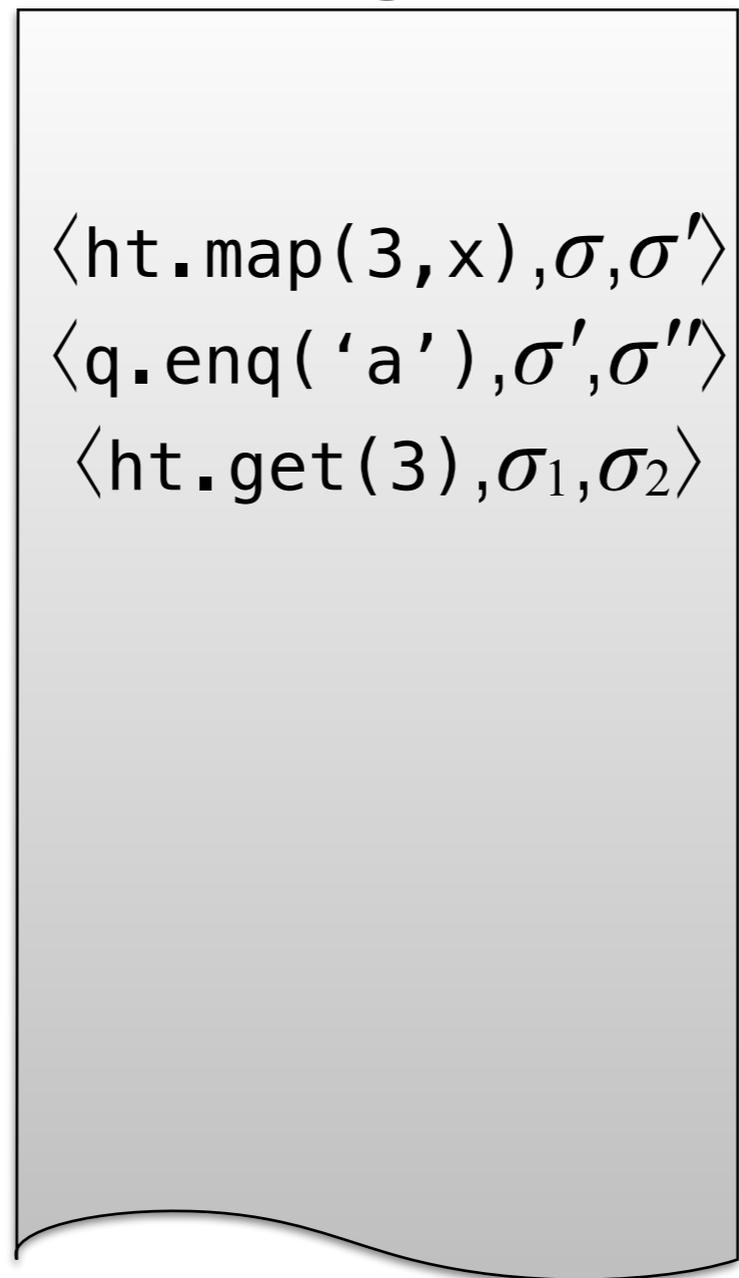
The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

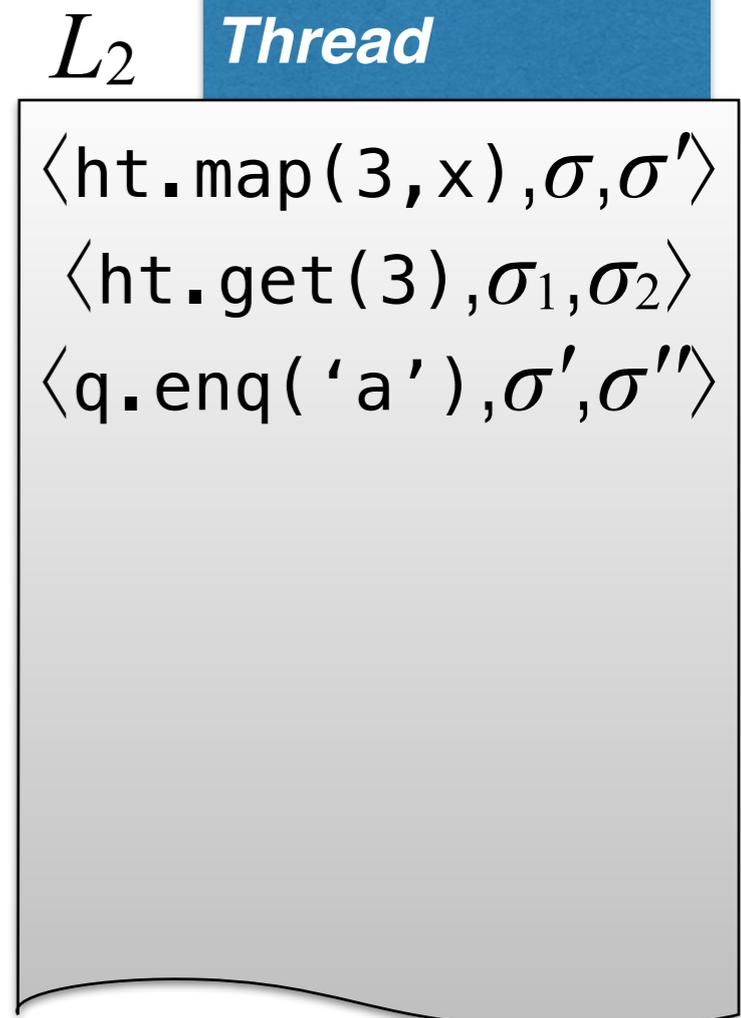


Unpull

G

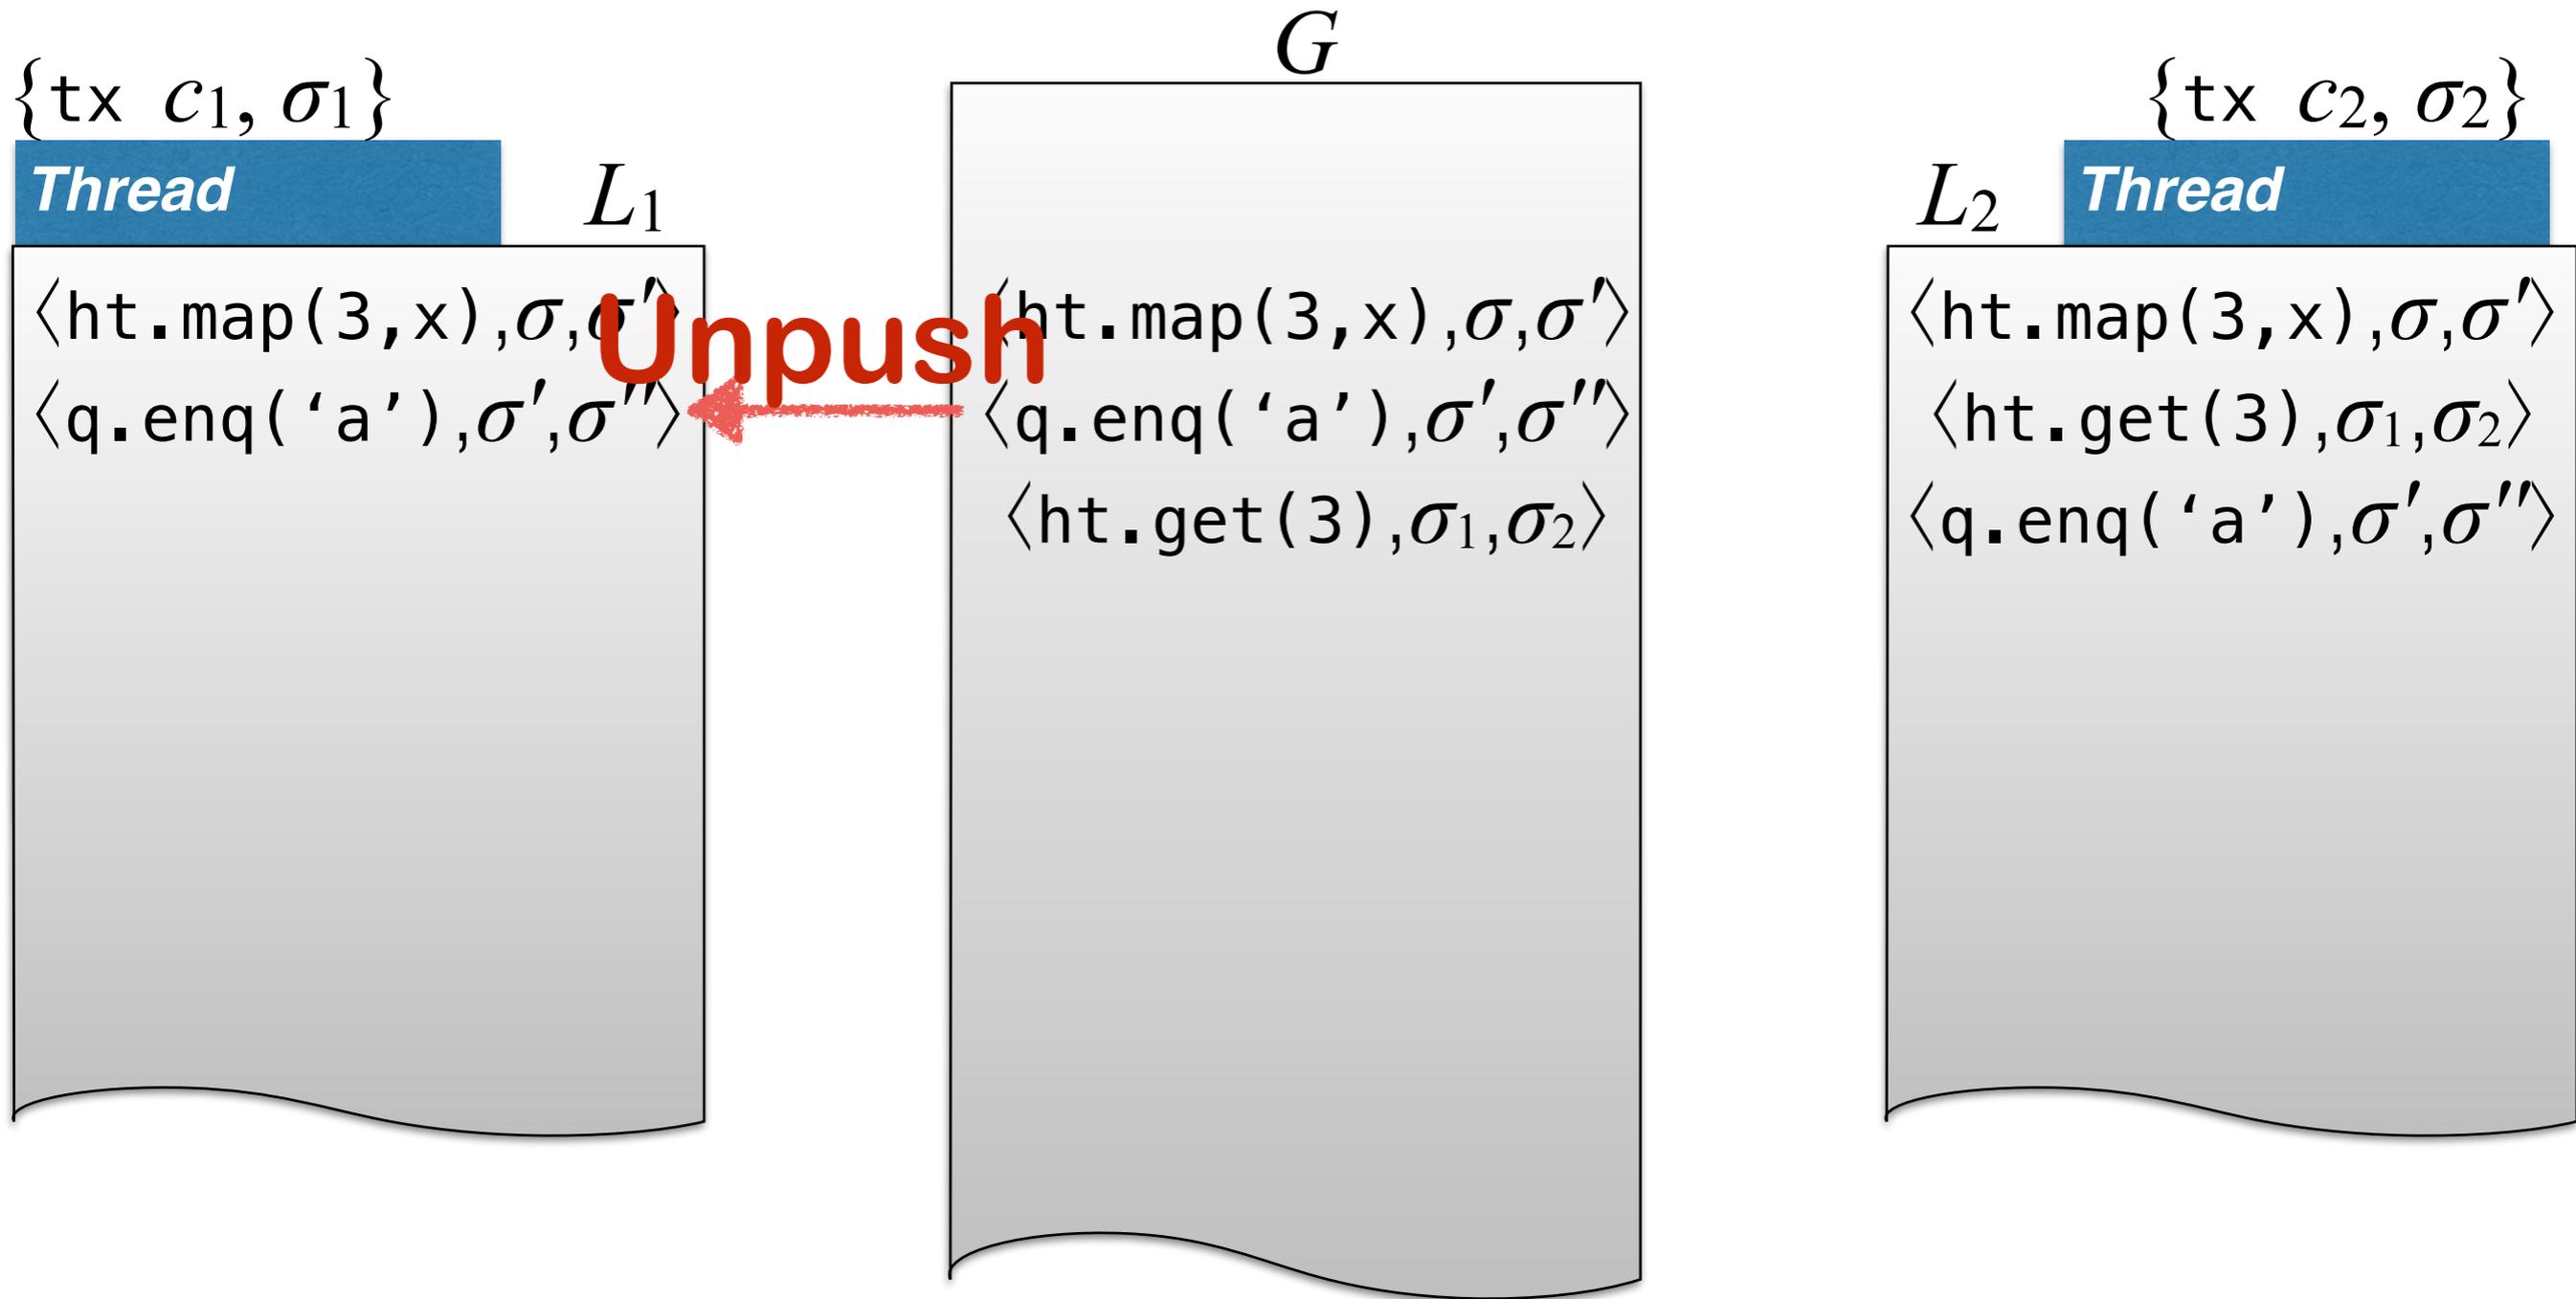


$\{tx\ c_2, \sigma_2\}$



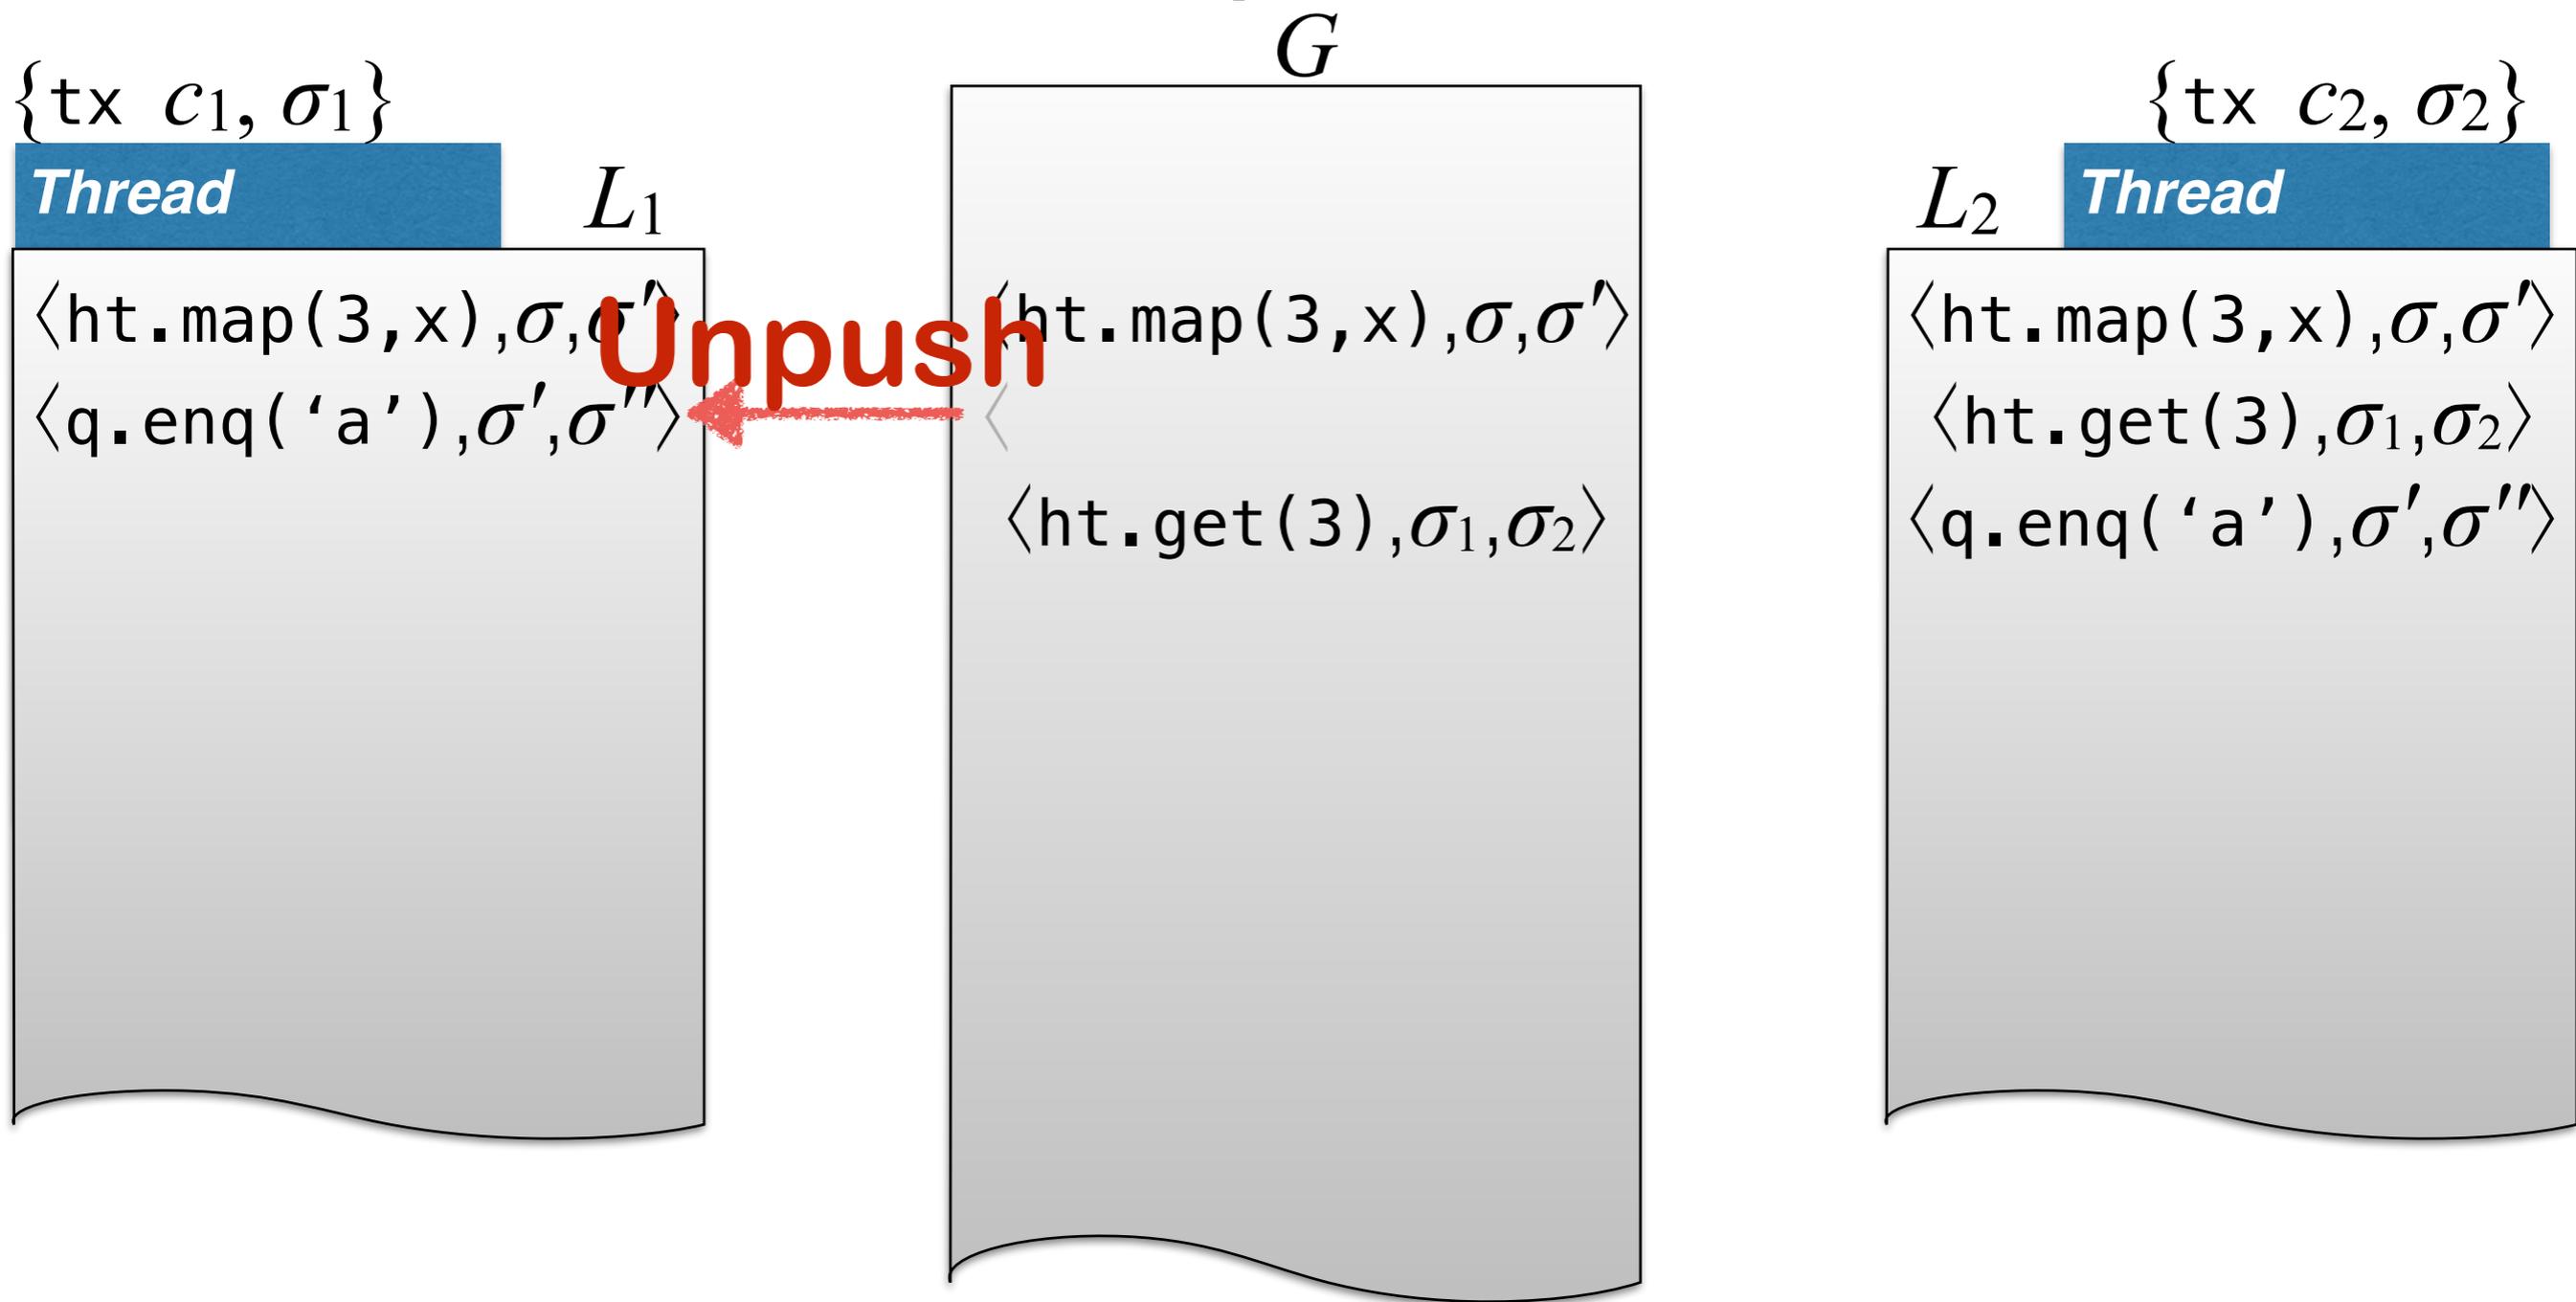
Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

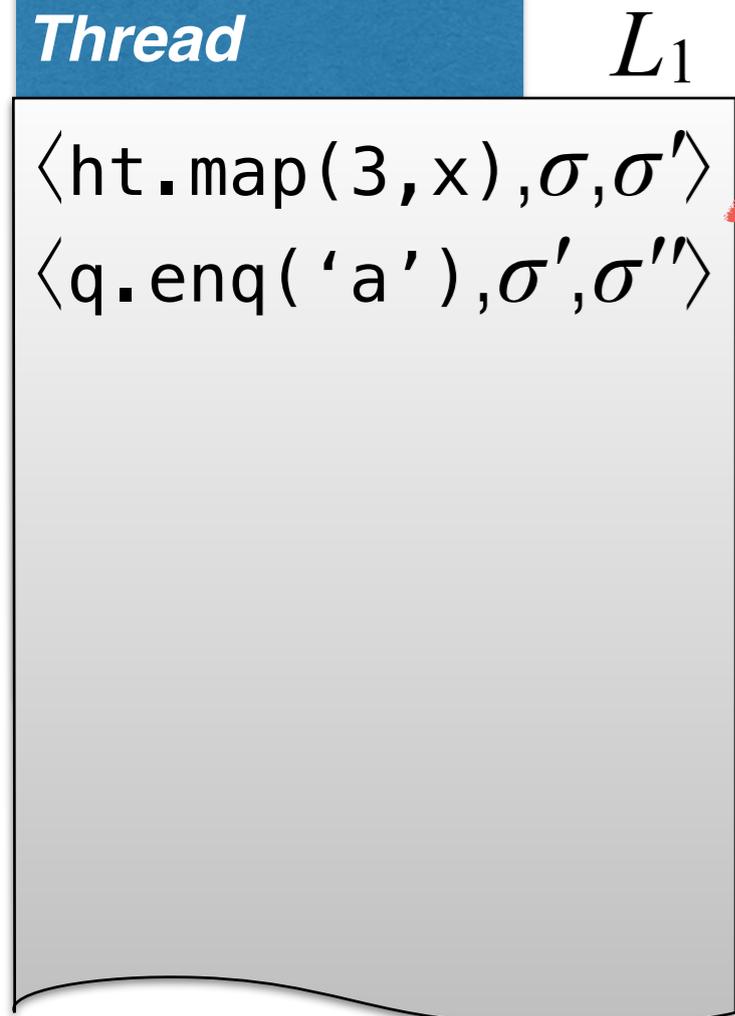
The Push/Pull Model



Apply Push Pull Unpull Unpush Unapply Commit

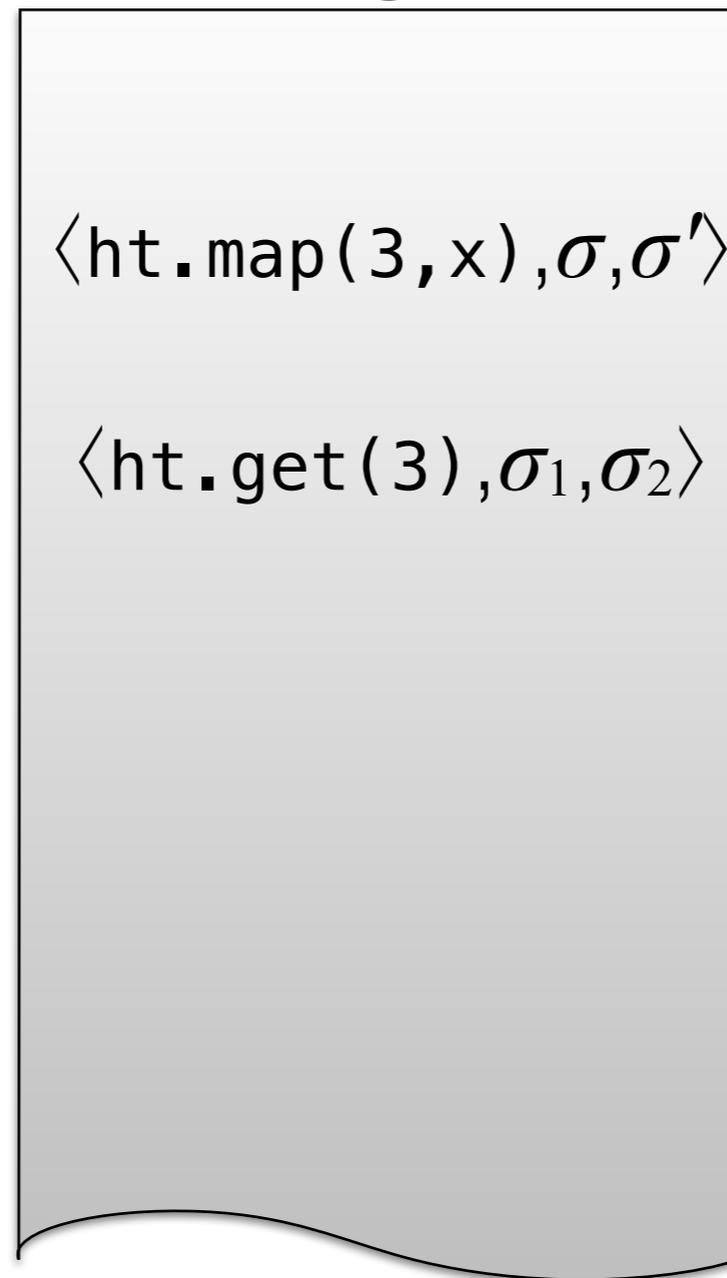
The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

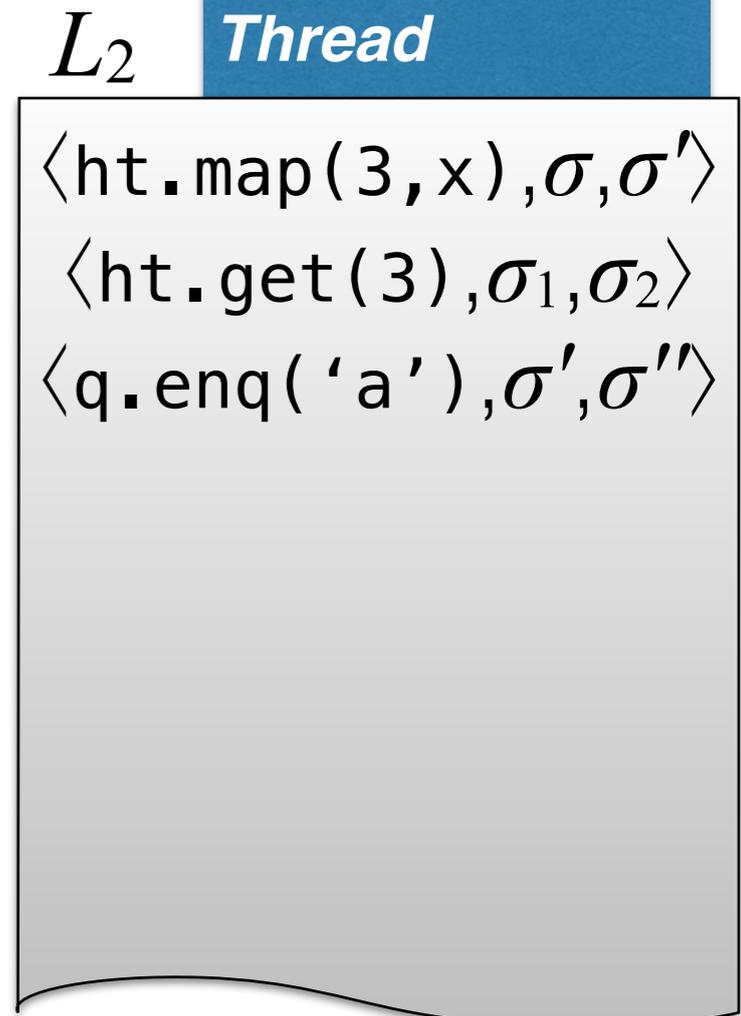


Unapply

G

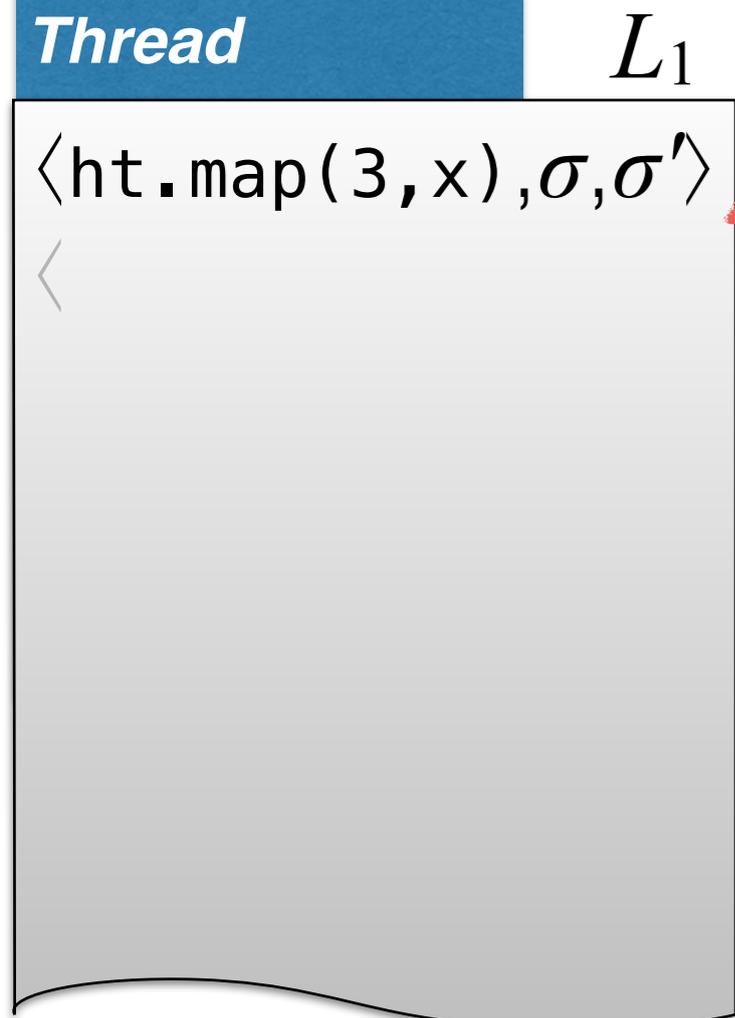


$\{tx\ c_2, \sigma_2\}$



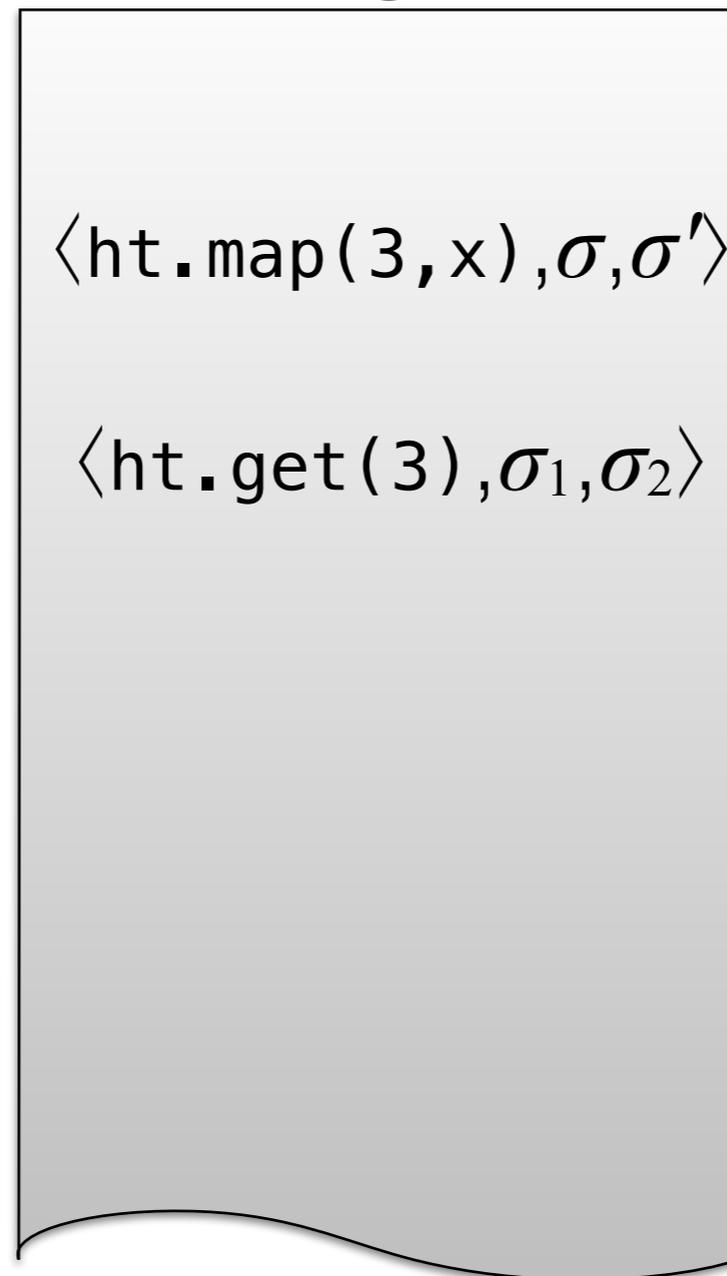
The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

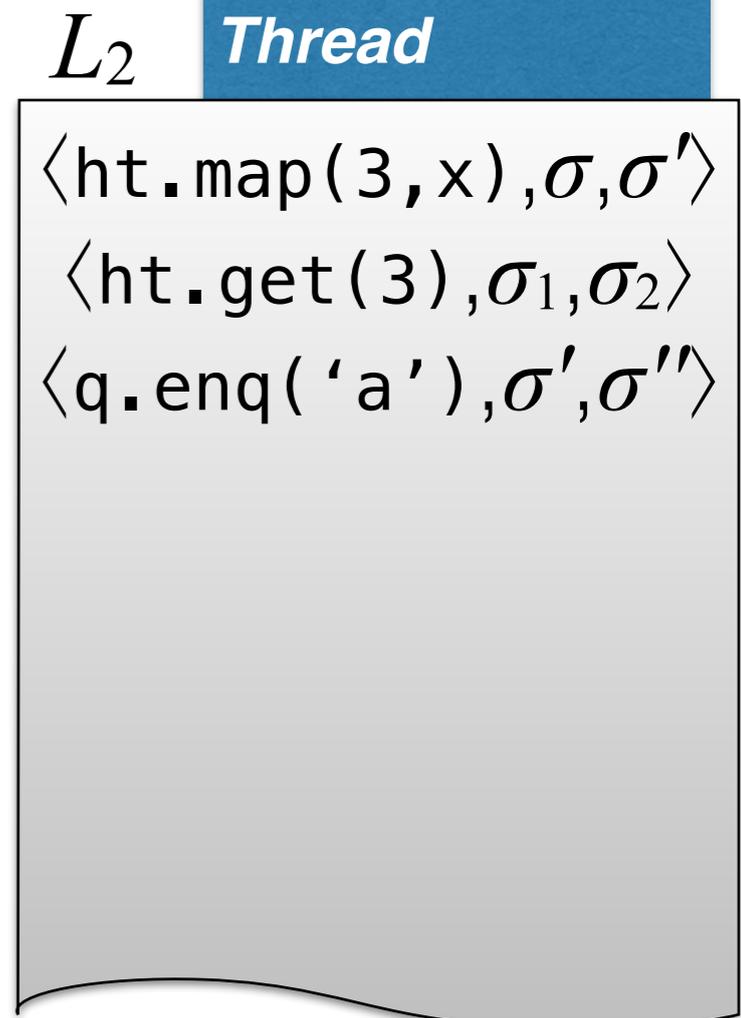


Unapply

G



$\{tx\ c_2, \sigma_2\}$



The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

Thread L_1

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$

G

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\{tx\ c_2, \sigma_2\}$

L_2 **Thread**

Unapply
 $\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle ht.get(3), \sigma_1, \sigma_2 \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$

Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

Thread L_1

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

\langle

G

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

$\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\{tx\ c_2, \sigma_2\}$

L_2 **Thread**

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

$\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\langle q.enq('a'), \sigma', \sigma'' \rangle$

Unapply

Apply Push Pull Unpull Unpush Unapply Commit

The Push/Pull Model

$\{tx\ c_1, \sigma_1\}$

Thread L_1

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

\langle

G

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

$\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\{tx\ c_2, \sigma_2\}$

L_2 **Thread**

$\langle ht.map(3, x), \sigma, \sigma' \rangle$

$\langle ht.get(3), \sigma_1, \sigma_2 \rangle$

$\langle q.enqueue(), \sigma', \sigma'' \rangle$

q.deq()

Unapply

Apply Push Pull Unpull Unpush Unapply Commit

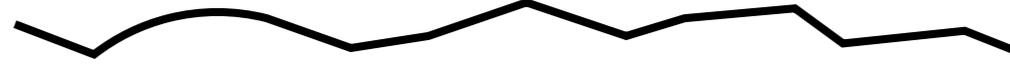
Simple (Intuitive?) Model

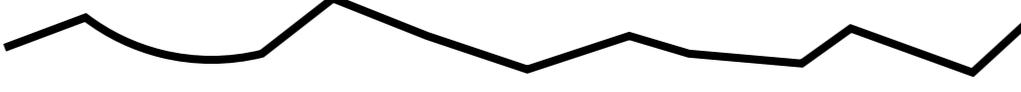
Simple (Intuitive?) Model

Getting it all to work . . .

Push/Pull Rule

Criterion (i): 

Criterion (ii): 

Criterion (iii): 

$$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c', \sigma', L'\}, G'$$

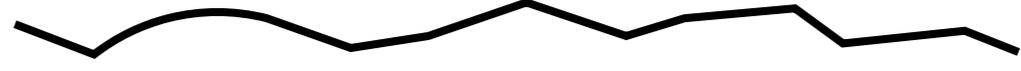
APPLY

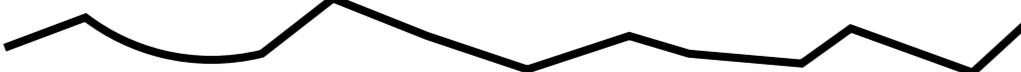
Machine Rule

$$T_1 \cdot \{\text{tx } c, \sigma_1, L_1\} \cdot T_2, G \rightarrow T_1 \cdot \{\text{tx } c', \sigma', L'\} \cdot T_2, G'$$

Push/Pull Rule

Criterion (i): 

Criterion (ii): 

Criterion (iii): 

$$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c', \sigma', L'\}, G'$$

APPLY

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply

Criterion (i):

Criterion (ii):

Criterion (iii):

$$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c', \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G \quad \text{APPLY}$$

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply

Criterion (i): $(m_1, c_2) \in \mathbf{step}(c)$

Criterion (ii):

Criterion (iii):

$$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G \quad \text{APPLY}$$

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply

step(c): The pair $(m, c') \in \text{step}(c)$ if m is a next reachable method call in the reduction of c , with remaining code c' .
fin(c): This predicate is true provided that there is a reduction of c to **skip** that does not encounter a method call.

Criterion (i): $(m_1, c_2) \in \text{step}(c)$

Criterion (ii):

Criterion (iii):

$$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G \quad \text{APPLY}$$

Thread

$c ::= c_1 + c_2 \mid c_1 ; c_2 \mid (c)^* \mid skip \mid tx\ c \mid m$

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$

Apply

step(c): The pair $(m, c') \in \text{step}(c)$ if m is a next reachable method call in the reduction of c , with remaining code c' .
fin(c): This predicate is true provided that there is a reduction of c to **skip** that does not encounter a method call.

Criterion (i): $(m_1, c_2) \in \text{step}(c)$

Criterion (ii):

Criterion (iii):

APPLY

$\{tx\ c, \sigma_1, L_1\}, G \rightarrow \{tx\ c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G$

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply

$c ::= c_1 + c_2 \mid c_1 ; c_2 \mid (c)^* \mid \text{skip} \mid \text{tx } c \mid m$

step(c): The pair $(m, c') \in \text{step}(c)$ if m is a next reachable method call in the reduction of c , with remaining code c' .
fin(c): This predicate is true provided that there is a reduction of c to **skip** that does not encounter a method call.

Criterion (i): $(m_1, c_2) \in \text{step}(c)$

Criterion (ii): L_1 allows $\langle m, \sigma_1, \sigma_2, id \rangle$

Criterion (iii):

APPLY

$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G$

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle$

Apply

$c ::= c_1 + c_2 \mid c_1 ; c_2 \mid (c)^* \mid \text{skip} \mid \text{tx } c \mid m$

step(c): The pair $(m, c') \in \text{step}(c)$ if m is a next reachable method call in the reduction of c , with remaining code c' .
fin(c): This predicate is true provided that there is a reduction of c to **skip** that does not encounter a method call.

Criterion (i): $(m_1, c_2) \in \text{step}(c)$

Criterion (ii): L_1 allows $\langle m, \sigma_1, \sigma_2, id \rangle$

Criterion (iii): $\text{fresh}(id)$

APPLY

$\{\text{tx } c, \sigma_1, L_1\}, G \rightarrow \{\text{tx } c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G$

Thread

$c ::= c_1 + c_2 \mid c_1 ; c_2 \mid (c)^* \mid skip \mid tx\ c \mid m$

$\langle ht.map(3, x), \sigma, \sigma' \rangle$
 $\langle q.enq('a'), \sigma', \sigma'' \rangle$

Apply

step(c): The pair $(m, c') \in \text{step}(c)$ if m is a next reachable method call in the reduction of c , with remaining code c' .
fin(c): This predicate is true provided that there is a reduction of c to **skip** that does not encounter a method call.

- Criterion (i): $(m_1, c_2) \in \text{step}(c)$
- Criterion (ii): L_1 allows $\langle m, \sigma_1, \sigma_2, id \rangle$
- Criterion (iii): $\text{fresh}(id)$

APPLY

$$\{tx\ c, \sigma_1, L_1\}, G \rightarrow \{tx\ c_2, \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G$$

append

Thread

L_1

↑
Unapply

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

UNAPPLY

$$\{\text{tx } c', \sigma_2, L_1 \cdot [\langle m, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c]\}, G \rightarrow \{\text{tx } c, \sigma_1, L_1\}, G$$

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push →

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [\text{op}, \text{unpushed } c] \cdot L_2 \cdot \}, G$

→

$\{\text{tx } c, \sigma, L_1 \cdot [\text{op}, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [\text{op}, \text{gUC}]$

Uncommitted

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push →

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

Criterion (i): $op \triangleleft [L_1]_{\text{unpushed}}$

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{unpushed } c] \cdot L_2 \cdot \}, G$

→

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [op, \text{gUC}]$

Uncommitted

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push 

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

Left-mover over logs

$$op_1 \blacktriangleleft op_2 \equiv \forall l. l \cdot \{op_1, op_2\} \preceq l \cdot \{op_2, op_1\}.$$

Criterion (i): $op \blacktriangleleft [L_1]_{\text{unpushed}}$

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{unpushed } c] \cdot L_2 \cdot \}, G$
 \rightarrow
 $\{\text{tx } c, \sigma, L_1 \cdot [op, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [op, \text{gUC}]$

Uncommitted

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push →

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

Criterion (i): $op \triangleleft [L_1]_{\text{unpushed}}$

Act as if op happens next

Criterion (ii): $[G]_{\text{gUC}} \setminus [L_1 \cdot L_2]_{\text{pushed}} \triangleleft op$

No conflict with other uncmtd

Criterion (iii): G allows op

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{unpushed } c] \cdot L_2 \cdot \}, G$
→
 $\{\text{tx } c, \sigma, L_1 \cdot [op, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [op, \text{gUC}]$

Application: Optimism vs. Pessimism

Apply Push Pull Unpull Unpush Unapply Commit

Thread

L_1
 $\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push →

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

Handwritten: $[op \text{ is } \text{unpushed}] \subseteq [L_1 \text{ is } \text{unpushed}; op]$

Criterion (i): $op \blacktriangleleft [L_1]_{\text{unpushed}}$ Act as if op happens next

Criterion (ii): $[G]_{\text{gUC}} \setminus [L_1 \cdot L_2]_{\text{pushed}} \blacktriangleleft op$ No conflict with other uncmtd

Criterion (iii): G allows op

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{unpushed } c] \cdot L_2 \cdot \}, G$
 \rightarrow
 $\{\text{tx } c, \sigma, L_1 \cdot [op, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [op, \text{gUC}]$

Application: Optimism vs. Pessimism

Apply Push Pull Unpull Unpush Unapply Commit

Thread

L_1
 $\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Push →

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$

Handwritten: $[op \text{ is pushed}] \subseteq [L_1 \text{ up to } op]$

Criterion (i): $op \blacktriangleleft [L_1]_{\text{unpushed}}$ Act as if op happens next

Criterion (ii): $[G]_{\text{gUC}} \setminus [L_1 \cdot L_2]_{\text{pushed}} \blacktriangleleft op$ No conflict with other uncmtd

Criterion (iii): G allows op

PUSH

$\{\text{tx } c, \sigma, L_1 \cdot [op, \text{unpushed } c] \cdot L_2 \cdot \}, G$
 \rightarrow
 $\{\text{tx } c, \sigma, L_1 \cdot [op, \text{pushed } c] \cdot L_2 \cdot \}, G \cdot [op, \text{gUC}]$

Application: Optimism vs. Pessimism

Apply Push Pull Unpull Unpush Unapply Commit

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$

Pull ←

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, g$

Criterion (i): $op \notin L$

← Didn't pull already

Criterion (ii): L allows op

← Args/RVs abide seq. spec

Criterion (iii): $op \triangleleft [L]_{\text{pushed}} \cup [L]_{\text{unpushed}}$

← Can act as if op happened earlier

PULL

$\{\text{tx } c, \sigma, L\}, G_1 \cdot [op, g] \cdot G_2$
→

$\{\text{tx } c, \sigma, L \cdot [op, \text{pulled}]\}, G_1 \cdot [op, g] \cdot G_2$

Application: Opacity [GK'08] and dependent transactions [RRHW'09]

Apply Push Pull Unpull Unpush Unapply Commit

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, \text{pulled}$

Pull ←

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, gUC$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, g$

Criterion (i): $op \notin L$

← Didn't pull already

Criterion (ii): L allows op

← Args/RVs abide seq. spec

Criterion (iii): $op \triangleleft [L]_{\text{pushed}} \cup [L]_{\text{unpushed}}$

← Can act as if op happened earlier

PULL

$\{\text{tx } c, \sigma, L\}, G_1 \cdot [op, g] \cdot G_2$

→

$\{\text{tx } c, \sigma, L \cdot [op, \text{pulled}]\}, G_1 \cdot [op, g] \cdot G_2$

Application: Opacity [GK'08] and dependent transactions [RRHW'09]

Apply Push Pull Unpull Unpush Unapply Commit

Thread

L_1

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } 'c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } 'c_1$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, \text{pulled}$

Pull ←

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, gUC$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, g$

Criterion (i): $op \notin L$

Criterion (ii): L allows op

Criterion (iii): op

← Didn't pull already

← Args/RVs abide seq. spec

← Can act as if op happened earlier

PULL

$\{\text{tx } c, \sigma, L\}, G_1 \cdot [op, g] \cdot G_2$

→

$\{\text{tx } c, \sigma, L \cdot [op, \text{pulled}]\}, G_1 \cdot [op, g] \cdot G_2$

Application: Opacity [GK'08] and dependent transactions [RRHW'09]

Apply Push Pull Unpull Unpush Unapply Commit

Thread

$\langle \text{ht.map}(3, x), \sigma, \sigma' \rangle, \text{unpushed } c$
 $\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{pushed } c_1$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, \text{pulled}$



L_1

Commit

$\langle \text{q.enq}('a'), \sigma', \sigma'' \rangle, \text{gUC}$
 $\langle \text{ht.get}(3), \sigma_1, \sigma_2 \rangle, g$

Criterion (i): **fin**(c)

Criterion (ii): $L \subseteq G_1$

Criterion (iii): $[L]_{\text{pulled}} \subseteq [G_1]_{\text{gC}}$

Criterion (iv): **cmt**(G_1, L, G_2)

Pushed all my stuff

Pulled all committed stuff

Swap my flags from **gUC** to **gC**

COMMIT

$\{\text{tx } c, \sigma, L\}, G_1$
 \rightarrow
 $\{\text{skip}, \sigma, []\}, G_2$

$$\begin{array}{l}
(i) - c_1 \not\ll (m_1, c_2) \\
(ii) - L_1 \text{ allows } \langle m_1, \sigma_1, \sigma_2 \rangle \\
(iii) - \text{fresh}(id) \\
\hline
\{\text{tx } c_1, \sigma_1, L_1\}, G_1 \xrightarrow{\text{fwd}} \{\text{tx } c_2, \sigma_2, L_1 \cdot [\langle m_1, \sigma_1, \sigma_2, id \rangle, \text{unpushed } c_1]\}, G_1 \quad \text{APP}
\end{array}$$

$$\frac{}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [\langle m_1, \sigma_2, \sigma_3, id \rangle, \text{unpushed } c_2]\}, G_1 \xrightarrow{\text{bwd}} \{\text{tx } c_2, \sigma_2, L_1\}, G_1 \quad \text{UNAPP}}$$

$$\begin{array}{l}
(i) - op \blacktriangleleft [L_1]_{\text{unpushed}} \\
(ii) - [G_1]_{\text{gUCmt}} \setminus [L_1 \cdot L_2]_{\text{pushed}} \blacktriangleleft op \\
(iii) - G_1 \text{ allows } op \\
\hline
\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{unpushed } c_2] \cdot L_2\}, G_1 \xrightarrow{\text{fwd}} \{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pushed } c_2] \cdot L_2\}, G_1 \cdot [op, \text{gUCmt}] \quad \text{PUSH}
\end{array}$$

$$\begin{array}{l}
(i) - \text{allowed } G_1 \cdot G_2 \\
(ii) - [L_2]_{\text{pushed}} \blacktriangleleft op \\
\hline
\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pushed } c_2] \cdot L_2\}, G_1 \cdot [op, g] \cdot G_2 \xrightarrow{\text{bwd}} \{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{unpushed } c_2] \cdot L_2\}, G_1 \cdot G_2 \quad \text{UNPUSH}
\end{array}$$

$$\begin{array}{l}
(i) - op \notin L \\
(ii) - L \text{ allows } op \\
(iii) - op \blacktriangleleft [L]_{\text{pushed}} \cup [L]_{\text{unpushed}} \\
\hline
\{\text{tx } c_1, \sigma_1, L\}, G_1 \cdot [op, g] \cdot G_2 \xrightarrow{\text{fwd}} \{\text{tx } c_1, \sigma_1, L \cdot [op, \text{pulled}]\}, G_1 \cdot [op, g] \cdot G_2 \quad \text{PULL}
\end{array}$$

$$\frac{(i) - \text{allowed } L_1 \cdot L_2}{\{\text{tx } c_1, \sigma_1, L_1 \cdot [op, \text{pulled}] \cdot L_2\}, G \xrightarrow{\text{bwd}} \{\text{tx } c_1, \sigma_1, L_1 \cdot L_2\}, G \quad \text{UNPULL}}$$

Apply Push Pull Unpull Unpush Unapply Commit

Theorem. The Push/Pull model is serializable.

Apply Push Pull Unpull Unpush Unapply Commit

Theorem. The Push/Pull model is serializable.

Push/Pull model guarantees that transactions are executed equivalently to just each being done atomically.

Apply Push Pull Unpull Unpush Unapply Commit

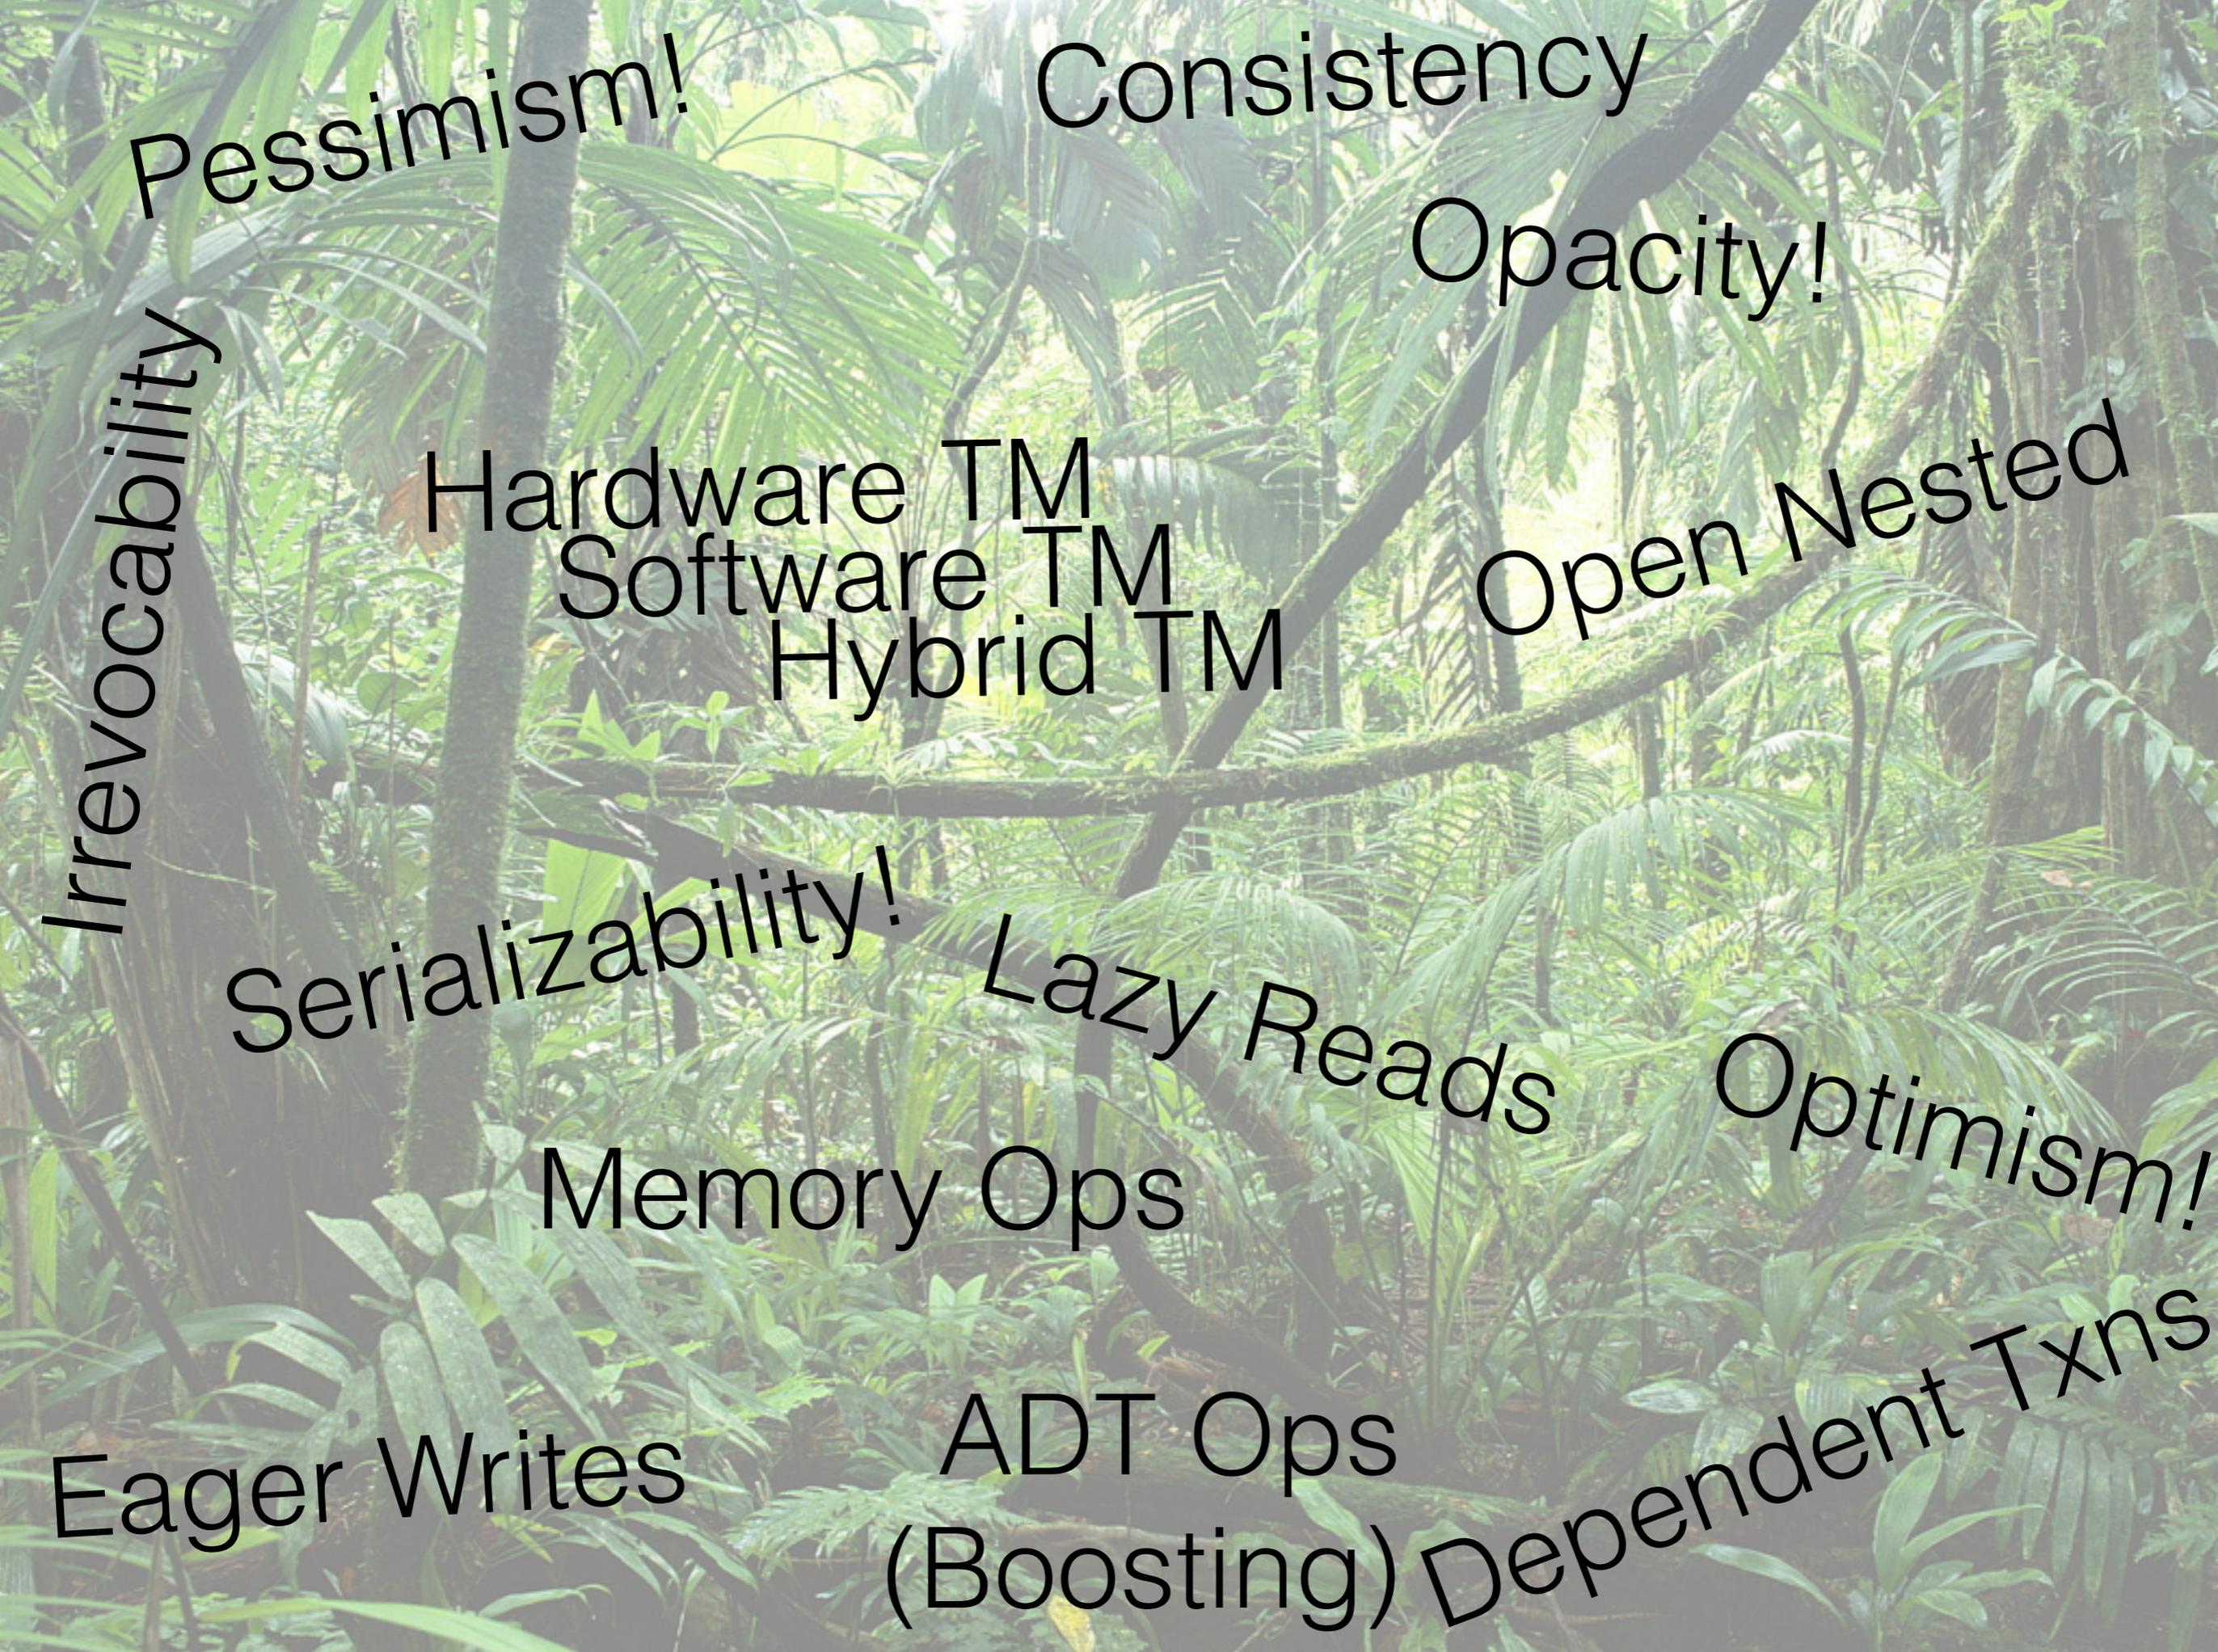
Theorem. The Push/Pull model is serializable.

Push/Pull model guarantees that transactions are executed equivalently to just each being done atomically.

Proof via simulation with atomic machine.

Apply Push Pull Unpull Unpush Unapply Commit

Evaluation



Pessimism!

Consistency

Opacity!

Irrevocability

Hardware TM

Software TM

Hybrid TM

Open Nested

Serializability!

Lazy Reads

Optimism!

Memory Ops

Eager Writes

ADT Ops

(Boosting) Dependent Txns

Conclusions

- Can encode numerous TM schemes into this
- Informal tool for thinking about TMs
- Formally provides serializability

Questions?



Apply Push Pull Unpull Unpush Unapply Commit