

Reasoning and Programming with Commutativity

Eric Koskinen

Stevens Institute of Technology
www.erickoskinen.com

MIT • Monday, May 9, 2022

Also Available →

[Top
Movies](#)[Photo
Galleries](#)[Video/DVD](#)[Browse
IMDb](#)[Independent
Film](#)

The Internet Movie Database

Visited by over 13 million movie lovers each month!

IMDbPro.com, the new website for the entertainment industry, now features over 22,000 agent/contact listings. [Click here for a free trial.](#)



Contact Information for Over 22,000 Professionals on IMDbPro.com

If you're looking for agent contact information for industry professionals, [IMDbPro.com](#) has contact data for over 22,000 names and the information is being added to and updated daily. If you're looking for company contact

information, international, domestic and daily box-office numbers, [IMDbPro.com](#) is the site for you. If you're looking for news and film reviews from the *Hollywood Reporter*, [IMDbPro.com](#) is where you need to go. Try [IMDbPro.com for free for two weeks](#). [Click here for details.](#)

Today's IMDb Poll Question Is:



It bothers me the most that.... (Vote for your favorite actors and films of 2002 with our new and improved [2002 poll](#).) ([vote](#))

SAG Nominees in [Road to the Oscars](#)



It's getting hotter and hotter in the Windy City, as the razzle-dazzle musical [Chicago](#) made off with five nominations from the [Screen Actors Guild Awards](#), including Best Performance by a Cast in a Motion Picture. Read more about the SAG awards in our [Road to the Oscars®](#) section. You'll also find other awards news and Oscar® trivia and quotes from Oscar®-nominated films. Best of all, have your say

and cast your ballot in our new and improved [Best of 2002 Poll](#). It's all on

Search the database for

All

Go!

[More searches](#) | [Tips](#)
[IMDbPro.com free trial](#)



Tops at the Box Office

- [Darkness Falls](#)
 - [Kangaroo Jack](#)
 - [Chicago](#)
 - [National Security](#)
 - [Just Married](#)
- [more](#)

Opening this Week

- [The Recruit](#)
 - [Final Destination](#)
 - [Biker Boyz](#)
 - [Lost in La Mancha](#)
 - [May](#)
- [more](#)

Coming Soon

- [How to Lose a Guy](#)
- [Shanghai Knights](#)
- [Deliver Us from Eva](#)
- [Daredevil](#)
- [The Jungle Book 2](#)
- [Max](#)
- [The Life of David Gale](#)

Movie and TV News

Wed January 29, 2003:
Celebrity News

- [Townshend: Email Will Clear Me](#)
- [Britney Dumps Durst for Farrell](#)
- [Crowe to Miss BAFTAs](#)

Studio Briefing

- [SAG Nods Go to 'Chicago,' 'The Hours'](#)
- [Eisner Getting \\$5-Million Stock Bonus](#)
- [O'Toole Rejects Oscar; Academy Says He Earned It](#)

Celebrity Interviews/Articles

- [Adaptation Filmmakers](#)
- [Glen Keane - Treasure Planet](#)

Cool Feature!

Happy Birthday to:

Thursday, January 30, 2003:

- [Christian Bale](#) (29)
- [Gene Hackman](#) (73)
- [Wilmer Valderrama](#) (23)
- [Vanessa Redgrave](#) (66)

[more birthdays](#)

Cool Services

- [Daily Newsletter](#)
- [Contact IMDb.com](#)



Can we serve web pages faster?



Can we serve web pages faster?

thttpd - tiny/turbo/throttling HTTP server



[Fetch the software.](#) [Release notes.](#)

thttpd is a simple, small, portable, fast, and secure HTTP server.

IPv6-Ready

Simple:

It handles only the minimum necessary to implement HTTP/1.1. Well, maybe a little more than the minimum.

Small:

See the [comparison chart](#). It also has a very small run-time size, since it does not fork and is very careful about memory allocation.

Portable:

It compiles cleanly on most any Unix-like OS, specifically including FreeBSD, SunOS 4, Solaris 2, BSD/OS, Linux, OSF.

Fast:

In typical use it's about as fast as the best full-featured servers (Apache, NCSA, Netscape). Under extreme load it's much faster.

Secure:

It goes to great lengths to protect the web server machine against attacks and breakins from other sites.

It also has one extremely useful feature ([URL-traffic-based throttling](#)) that no other server currently has. Plus, it supports [IPv6](#) out of the box, no patching required.

More specific info:

- [HTMLized man page](#)
- [thttpd notes](#)

Let me try to write something faster myself.

Wow, concurrency is hard!

Why does writing concurrent programs have to be so hard?

Transactional Boosting: A Methodology for Highly-Concurrent Transactional Objects

Maurice Herlihy Eric Koskinen

Computer Science Department, Brown University
{mph,ejk}@cs.brown.edu

Abstract

We describe a methodology for transforming a large class of highly-concurrent linearizable objects into highly-concurrent transactional objects. As long as the linearizable implementation satisfies certain regularity properties (informally, that every method has an inverse), we define a simple wrapper for the linearizable implementation that guarantees that concurrent transactions without inherent conflicts can synchronize at the same granularity as the original linearizable implementation.

Categories and Subject Descriptors D.1.3 [Programming Techniques]: Concurrent Programming – Parallel Programming; D.3.3 [Programming Languages]: Language Constructs and Features – Frameworks; Concurrent programming structures; E.1 [Data Structures]: Distributed data structures; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms Algorithms, Languages, Theory

Synchronizing via read/write conflicts has one substantial advantage: it can be done automatically without programmer participation. It also has a substantial disadvantage: it can severely and unnecessarily restrict concurrency for certain shared objects. If these objects are subject to high levels of contention (that is, they are “hot-spots”), then the performance of the system as a whole may suffer.

Here is a simple example. Consider a mutable set of integers that provides $\text{add}(x)$, $\text{remove}(x)$ and $\text{contains}(x)$ methods with the obvious meanings. Suppose we implement the set as a sorted linked list in the usual way. Each list node has two fields, an integer value and a node reference next. List nodes are sorted by value, and values are not duplicated. Integer x is in the set if and only if a list node has value field x . The $\text{add}(x)$ method reads along the list until it encounters the largest value less than x . Assuming x is absent, it creates a node to hold x , and links it to the next node.

Consider a set whose state is $\{1, 3, 5\}$ and a transaction A to add 2 to the set and transaction B is about to remove 3.

Transactional Boosting: A Methodology for Highly-Concurrent Transactional Objects

Coarse-Grained Transactions

Eric Koskinen
University of Cambridge

Matthew Parkinson
University of Cambridge

Maurice Herlihy
Brown University

Koskinen
University

POPL 2010

fer from overly con-
d false conflicts, be-
read/write conflicts.
In response, the recent trend has been toward integrating various
abstract data-type libraries using ad-hoc methods of high-level con-
flict detection. These proposals have led to improved performance
but a lack of a unified theory has led to confusion in the literature.
We clarify these recent proposals by defining a generaliza-
tion of transactional memory in which a transaction consists of
coarse-grained (abstract data-type) operations rather than simple
memory read/write operations. We provide semantics for both pes-
simistic (e.g. transactional boosting) and optimistic (e.g. traditional

the locations
it wrote. Two
intersects the
conflict to be
easy to class
reads or writ
Nevertheless
conflicts sup
ject to conten
is conservati
flict even tho
threads inser
neither one s
location type

The Push/Pull Model of Transactions

Eric Koskinen *

IBM TJ Watson Research Center, USA

Matthew Parkinson

Microsoft Research Cambridge, U

PLDI 2015

PODC 2017

Adding Concurrency to Smart Contracts

Thomas Dickerson
Brown University
thomas_dickerson@brown.edu

Maurice Herlihy

Paul Gazzillo
Yale University
paul.gazzillo@yale.edu

Eric Koskinen

shared memory should appear to
by another thread.
such a construct, we must be
implementations typically achieve
cts between concurrent threads
y operations in hardware [14]
hwhile, an alternate approach e
ict over linearizable data-struct
ty [11, 20, 21, 30]. Both level
optimistic execution, pessimist
. Finally, there are multiple no

Still have to write concurrent programs.

Oh the dream of parallelizing compilers!

Still have to write concurrent programs.

Oh the dream of parallelizing compilers!

Commutativity Analysis: A New Analysis Technique for Parallelizing Compilers

MARTIN C. RINARD

Massachusetts Institute of Technology

and

PEDRO C. DINIZ

University of Southern California / Information Sciences Institute

This article presents a new analysis technique, commutativity analysis, for automatically parallelizing computations that manipulate dynamic, pointer-based data structures. Commutativity analysis views the computation as composed of operations on objects. It then analyzes the pro-

Wanted to understand program analysis better...

- Moved to Cambridge, UK
- Dissertation on **program analysis for temporal logic** verification
- Abstraction-refinement, automata, ...
- Started to think more and more about how



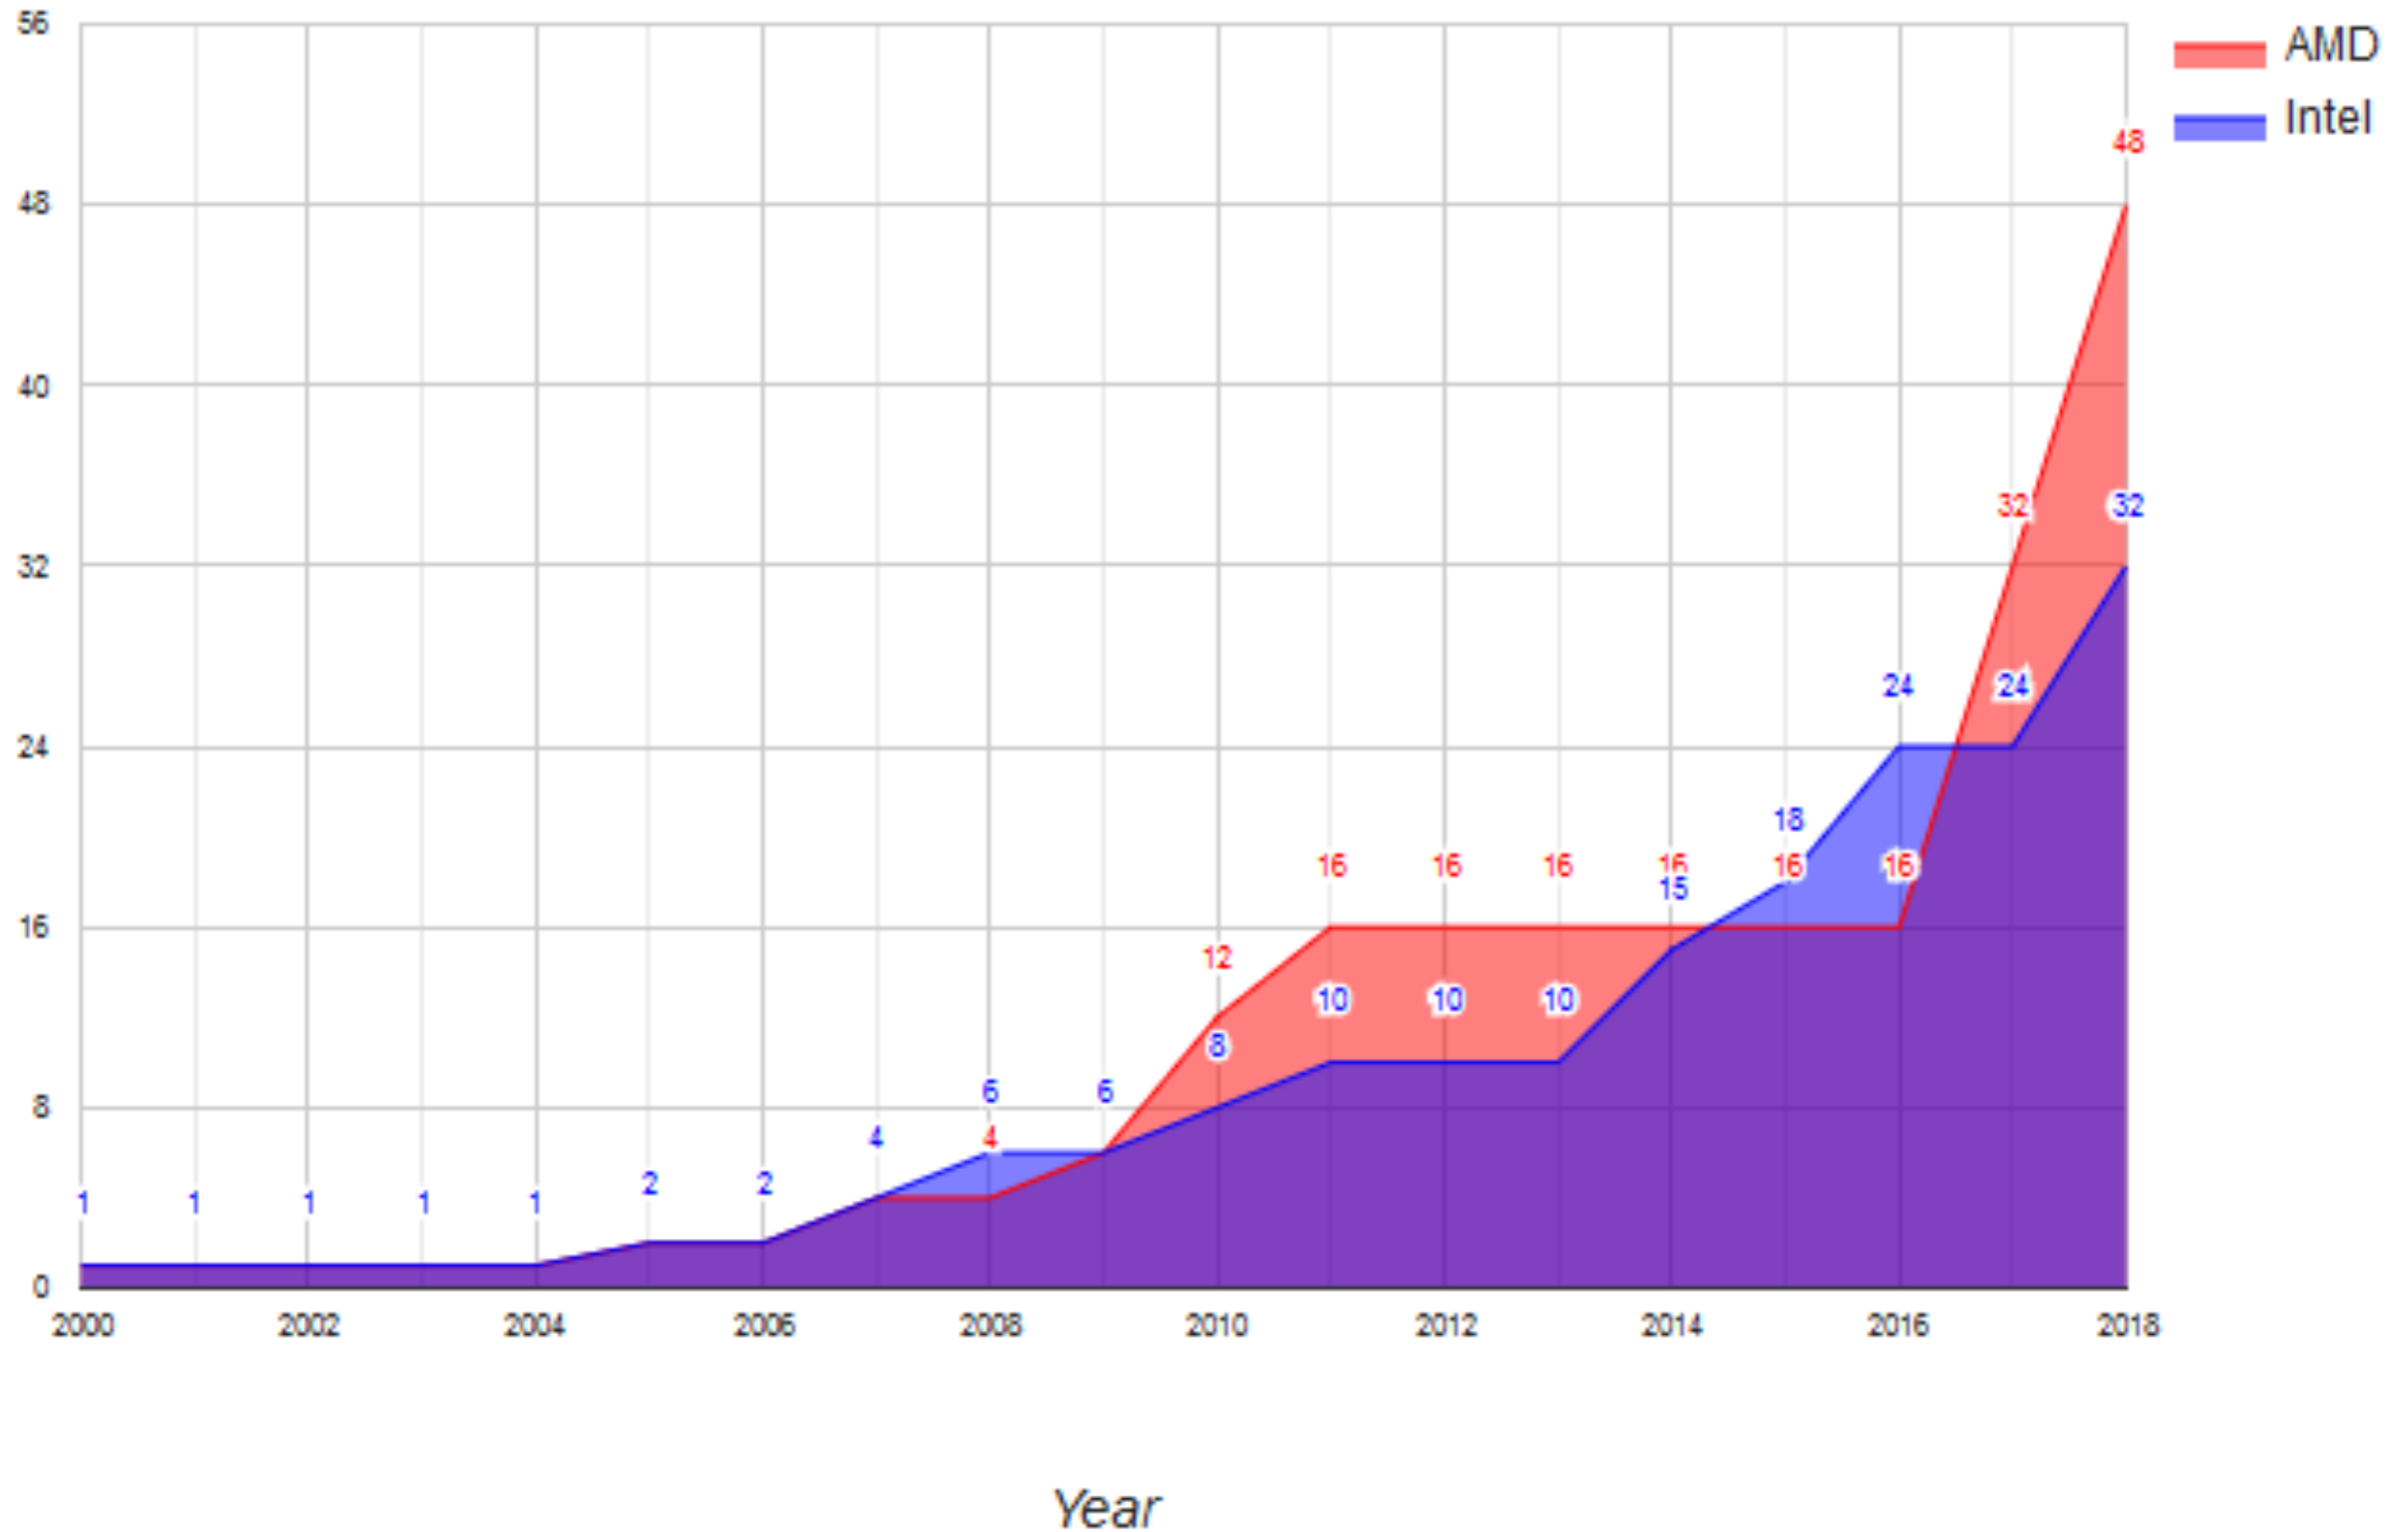
Wanted to understand program analysis better...

- Moved to Cambridge, UK
- Dissertation on **program analysis for temporal logic** verification
- Abstraction-refinement, automata, ...
- Started to think more and more about how



Symbolic program analysis could enable **parallelization**


Highest amount of cores per CPU (AMD vs Intel year by year)



8-core CPU

The highest-performance CPU
we've ever built.

Up to
3.5x
faster CPU
performance¹



Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Automatic parallelization

```
sum = 0;  
#pragma omp parallel for shared(sum, a) reduction(+: sum)  
for (auto i = 0; i < 10; i++)  
{  
    sum += a[i];  
}
```

Exploits
Commutativity

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

***Commutativity** is a well-known strategy for concurrency.*

- **Databases.** e.g. Weihl 1988.
- **Parallelizing compilers.** e.g. Rinard & Diniz, TOPLAS 1997.
- **Parallel graph algorithms.** e.g. Kulkarni et al, PLDI 2007.
- **Transactional memory.** e.g. Ni et al., PPOPP 2007. Herlihy & Koskinen, PPOPP 2008. Bronson et al., PODC 2010. Hassan et al., PPOPP 2014. Koskinen & Parkinson, PLDI 2015. Dickerson et al., APLAS 2019.
- **Runtime systems.** e.g. Tripp et al, OOPSLA 2011.
- **Software scalability.** e.g. Clements et al., TOCS 2015.
- **Layered concurrent programs.** e.g. Kragl & Qadeer, CAV 2018.

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Weihl 1988.

compilers. e.g. Rinard & Diniz, TOPLAS 1997.

algorithms. e.g. Kulkarni et al, PLDI 2007.

- **Transactional memory.** e.g. Ni et al., PPOPP 2007. Herlihy & Koskinen, PPOPP 2008. Bronson et al., PODC 2010. Hassan et al., PPOPP 2014. Koskinen & Parkinson, PLDI 2015. Dickerson et al., APLAS 2019.
- **Runtime systems.** e.g. Tripp et al, OOPSLA 2011.
- **Software scalability.** e.g. Clements et al., TOCS 2015.
- **Layered concurrent programs.** e.g. Kragl & Qadeer, CAV 2018.

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Weihl 1988.

Compilers. e
algorithm

Commutativity Analysis: A New Analysis Technique for Parallelizing Compilers

MARTIN C. RINARD

Massachusetts Institute of Technology

and

PEDRO C. DINIZ

University of Southern California / Information Sciences Institute

- **Transactional memory.** e

Bronson et al., PODC 2010. Hass

Dickerson et al., APLAS 2019.

- **Runtime systems.** e.g. Trip

- **Software scalability.** e.g. C

- **Layered concurrent prog**

This article presents a new analysis technique, commutativity analysis, for automatically parallelizing computations that manipulate dynamic, pointer-based data structures. Commutativity analysis views the computation as composed of operations on objects. It then analyzes the program at this granularity to discover when operations commute (i.e., generate the same final result).

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Weihl 1988.

Compilers. e
algorithm

Commutativity Analysis: A New Analysis Technique
for Parallelizing Compilers

MARTIN C. RINARD

- **Transactional memory.** e

The Scalable Commutativity Rule: Designing Scalable Software for Multicore Processors

AUSTIN T. CLEMENTS, M. FRANS KAASHOEK, NICKOLAI ZELDOVICH,
and ROBERT T. MORRIS, MIT CSAIL
EDDIE KOHLER, Harvard University

What opportunities for multicore scalability are latent in software interfaces, such as system call APIs? Can scalability challenges and opportunities be identified even before any implementation exists, simply by considering interface specifications? To answer these questions, we introduce the scalable commutativity rule: *whenever interface operations commute, they can be implemented in a way that scales*. This rule is useful throughout the development process for scalable multicore software, from the interface design through

of Technology

California / Information Sciences Institute

analysis technique, commutativity analysis, for automatically paral-
manipulate dynamic, pointer-based data structures. Commutativity
tion as composed of operations on objects. It then analyzes the pro-
discover when operations commute (i.e., generate the same final result

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Wehl 1988.

A Revised and Verified Proof of the Scalable Commutativity Rule

Lillian Tsai[†], Eddie Kohler^{*}, M. Frans Kaashoek[†], and Nikolai Zeldovich[†]
[†] MIT CSAIL ^{*} Harvard University

1 Introduction

This paper explains a flaw in the published proof of the Scalable Commutativity Rule (SCR) [1], presents a revised and formally verified proof of the SCR in the Coq proof assistant, and discusses the insights and open questions raised from our experience proving the SCR.

2 The Scalable Commutativity Rule

In order to explore the connection between commutativity and scalability in practical systems, Clements et al. [1] defined a new type of commutativity called *SIM commutativity*,¹ a property that can hold of certain interface specifications. This was used to state and prove the Scalable Commutativity Rule (SCR), which claims that every SIM-

A *specification* models an interface's behavior as a prefix-closed set of well-formed histories. A system execution is "correct" according to the specification if its trace is included in the specification. For instance, if \mathcal{S} corresponded to the POSIX specification, then $[\text{getpid}_\alpha, \overline{92}_\alpha] \in \mathcal{S}$ (a process may have PID 92) but $[\text{getpid}_\alpha, \overline{\text{ENOENT}}_\alpha] \notin \mathcal{S}$ (the `getpid()` system call may not return that error). A specification constrains both invocations and responses: $[\text{NtAddAtom}_\alpha]$ is not in the POSIX specification because `NtAddAtom` is not a POSIX system call.

An *implementation* is an abstract machine that takes invocations and calculates responses. The original proof of the SCR by Clements et al. [1] (also presented in Section 2.4) uses a class of machines on which conflict-freedom is defined: a good analogy is a Turing-type machine with a

Analysis Technique

s Institute

y analysis, for automatically paral-
d data structures. Commutativity
objects. It then analyzes the pro-
(i.e., generate the same final result

• Transac

The Scalable Comm for Multicore Proces

AUSTIN T. CLEMENTS, I
and ROBERT T. MORRIS
EDDIE KOHLER, Harvard

What opportunities for multic
Can scalability challenges an
by considering interface specifi
rule: *whenever interface oper*
useful throughout the develop

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Wehl 1988.

A Revised and Verified Proof of the Scalable Commutativity Rule

Lillian Tsai[†], Eddie Kohler^{*}, M. Frans Kaashoek[†], and Nikolai Zeldovich[†]
[†] MIT CSAIL ^{*} Harvard University

Analysis Technique

- **Transac**

The Scalable Comm

for Multicore Proces

AUSTIN T. CLEMENTS, I
and ROBERT T. MORRIS
EDDIE KOHLER, Harvard

What opportunities for multie
Can scalability challenges an
by considering interface specif
rule: *whenever interface oper*
useful throughout the develop

1 Introduction

This paper explains
able Commutativity
formally verified pro
and discusses the ins
experience proving t

2 The Scalable

In order to explore
and scalability in pr
fined a new type of
tivity,¹ a property th
ifications. This was
Commutativity Rule

ScaleFS: A Multicore-Scalable File System

by

Rasha Eqbal

B.Tech.(Hons.) in Computer Science and Engineering
Indian Institute of Technology Kharagpur (2011)

cally paral-
mutativity
es the pro-
final result

Automatic parallelization

```
sum = 0;
#pragma omp parallel for shared(sum, a) reduction(+: sum)
for (auto i = 0; i < 10; i++)
{
    sum += a[i];
}
```

Commutativity is a well-known strategy for concurrency.



. Wehl 1988.

A Revised and Verified Proof of the Scalable Commutativity Rule

Lillian Tsai[†], Eddie Kohler^{*}, M. Frans Kaashoek[†], and Nikolai Zeldovich[†]
[†] MIT CSAIL ^{*} Harvard University

Analysis Technique

• Transac

The Scalable Comm

AUSTIN T. CLEMENTS, I
and ROBERT T. MORRIS
EDDIE KOHLER, Harvard

What opportunities for multie
Can scalability challenges an
by considering interface specif
rule: *whenever interface oper*
useful throughout the develop

1 Introduction

This paper explains
able Commutativity
formally verified pro
and discusses the ins
experience proving t

2 The Scalable

In order to explore
and scalability in pr
fined a new type o
tivity,¹ a property t
ifications. This was
Commutativity Rule

ScaleFS: A Multicore-Scalable File System

Verifying concurrent software using movers in CSPEC

Tej Chajed, M. Frans Kaashoek, Butler Lampson[†], and Nikolai Zeldovich
 MIT CSAIL and [†]Microsoft Research

Automatic parallelization

But programs don't always commute ...

```
x = calc1(a);  
c = c + (x*x);  
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```



```
x = calc1(a);  
c = c + (x*x);  
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Expensive, pure
computation

```
x = calc1(a);  
c = c + (x*x);  
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Expensive, pure
computation

Another expensive,
pure computation

```
x = calc1(a);  
c = c + (x*x);  


---

if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Expensive, pure
computation

Another expensive,
pure computation

Can we run them in parallel?

Let's try separating splitting the code.

Wouldn't it be nice if we could parallelize these two blocks?

```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Can we run them in parallel?

```
x = calc1(a);  
c = c + (x*x);
```

Observation on c

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Can we run them in parallel?

Mutation of c

```
x = calc1(a);  
c = c + (x*x);
```

Observation on c

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Can we run them in parallel?

Mutation of c

```
x = calc1(a);  
c = c + (x*x);
```

Observation on c

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Can we run them in parallel?

A simple dataflow analysis cannot parallelize them.

Dataflow dependency prevents naive parallelization.

Splitting differently doesn't help; x conflicts.

```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```



Idea: Making commutativity *explicit* can allow us to **weaken** the dataflow dependency.

Consider: what if **c>0** initially?

Then these blocks are semantically **independent**.

(with some atomicity assumptions)


```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Observes if $c > 0$
Decreases c by 1



Idea: Making commutativity *explicit* can allow us to **weaken** the dataflow dependency.

Consider: what if $c > 0$ initially?

Then these blocks are semantically **independent**.

(with some atomicity assumptions)

Increases c

```
x = calc1(a);  
c = c + (x*x);
```

Observes if c>0
Decreases c by 1

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```



Idea: Making commutativity *explicit* can allow us to **weaken** the dataflow dependency.

Consider: what if **c>0** initially?

Then these blocks are semantically **independent**.

(with some atomicity assumptions)

```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```



Idea: Making commutativity *explicit* can allow us to **weaken** the dataflow dependency.


Allow the programmer to explicitly express conditional, sequential commutativity.

- Introduce the **commute** keyword.
- Programmer only has to reason **sequentially**.
- Verification tools need only reason **sequentially**.
- Obtain **speedup from parallel** execution.

commute ($c > 0$) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```



Idea: Making commutativity *explicit* can allow us to **weaken** the dataflow dependency.

Allow the programmer to explicitly express conditional, sequential commutativity.

- Introduce the **commute** keyword.
- Programmer only has to reason **sequentially**.
- Verification tools need only reason **sequentially**.
- Obtain **speedup from parallel** execution.

```
x = calc1(a);  
c = c + (x*x);
```

if(c>0)

then →

else →

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

||

```
x = calc1(a);  
c = c + (x*x);  
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

if (c>0)

then

else

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

||

```
x = calc1(a);  
c = c + (x*x);
```

```
x = calc1(a);  
c = c + (x*x);  
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

commute (c>0) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```

```
}
```

commute (c>0) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```

```
}
```

Semantics?

commute (c>0) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

Semantics?

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```

Synchronization?

}

commute (c>0) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

Semantics?

Correct?

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```

Synchronization?

}

commute (c>0) {

```
{  
  x = calc1(a);  
  c = c + (x*x);  
}
```

```
{  
  if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
  } else {  
    z = calc3(y);  
  }  
}
```

}

Semantics?

Correct?

Synchronization?

Speedup?

✓ Introducing **commute** blocks

2 Semantic Implications & Correctness Criteria *“Scoped Serializability”
and lock synthesis*


3 Demo of the Veracity language veracity-lang.org

4 Speedup

Part B:

5 Symbolic Commutativity Reasoning *TACAS’18, JAR’20, VMCAI’21*

6 Inferring/verifying **commute** blocks *Translation to logic*



2. Semantic Implications & Correctness Criteria

Want to parallelize a program s .

We can do so soundly when the parallel behavior matches that of its equivalent straight-line code:

$$\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$$

Outcome: Get gains without changing the way we write sequential programs.

Sequential Behavior: non-deterministic

$$\langle \mathbf{commute(true)}\{s_1\}\{s_2\}, \sigma \rangle \begin{array}{l} \rightsquigarrow_{nd} \langle s_1; s_2, \sigma \rangle \\ \rightsquigarrow_{nd} \langle s_2; s_1, \sigma \rangle \end{array}$$

Sequential Behavior: non-deterministic

$$\begin{aligned}
 \langle \mathbf{commute}(\mathbf{true})\{s_1\}\{s_2\}, \sigma \rangle &\rightsquigarrow_{nd} \langle s_1; s_2, \sigma \rangle \\
 &\rightsquigarrow_{nd} \langle s_2; s_1, \sigma \rangle
 \end{aligned}$$

Parallel Behavior: Interleaved, as expected:

$$\langle \mathbf{commute}(\mathbf{true})\{s_1\}\{s_2\}, \sigma \rangle \rightsquigarrow_{par} \langle (\langle s_1, \emptyset \rangle, \langle s_2, \emptyset \rangle,), \mathbf{skip}, \sigma \rangle$$

$ \mathfrak{C}_0 \oplus \sigma \rightsquigarrow_{par} \mathfrak{C}'_0 \oplus \sigma' $	Left-Proj
<hr style="border: 0.5px solid black;"/> $ \langle (\mathfrak{C}_0, \mathfrak{C}_1), s, \sigma \rangle \rightsquigarrow_{par} \langle (\mathfrak{C}'_0, \mathfrak{C}_1), s, \sigma' \rangle $	
(mut. mut.) R-Proj	

(Full semantics in the paper)

Sequential Behavior: non-deterministic

$$\langle \mathbf{commute}(\mathbf{true})\{s_1\}\{s_2\}, \sigma \rangle \begin{cases} \rightsquigarrow_{nd} \langle s_1 \mathbin{;} s_2, \sigma \rangle \\ \rightsquigarrow_{nd} \langle s_2 \mathbin{;} s_1, \sigma \rangle \end{cases}$$

Want equivalence
(as state fns)

Parallel Behavior: Interleaved, as expected:

$$\langle \mathbf{commute}(\mathbf{true})\{s_1\}\{s_2\}, \sigma \rangle \rightsquigarrow_{par} \langle (\langle s_1, \emptyset \rangle, \langle s_2, \emptyset \rangle,), \mathbf{skip}, \sigma \rangle$$

$\mathcal{C}_0 \oplus \sigma \rightsquigarrow_{par} \mathcal{C}'_0 \oplus \sigma'$	Left-Proj
<hr style="border: 0.5px solid black;"/> $\langle (\mathcal{C}_0, \mathcal{C}_1), s, \sigma \rangle \rightsquigarrow_{par} \langle (\mathcal{C}'_0, \mathcal{C}_1), s, \sigma' \rangle$	
(mut. mut.) R-Proj	

(Full semantics in the paper)

Goal: $[[s]]_{nd} = [[s]]_{par}$

How to ensure equivalence.

Serializability?

2. Semantic Implications & Correctness Criteria

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

Serializability?

```
y = 0; x = 1;
commute(true) {
  { commute(true)
    f1:{ x = 0; }
    f2:{ x = x*2; }}
  f3:{ if(x>=2) y = 1; }
}
```

2. Semantic Implications & Correctness Criteria

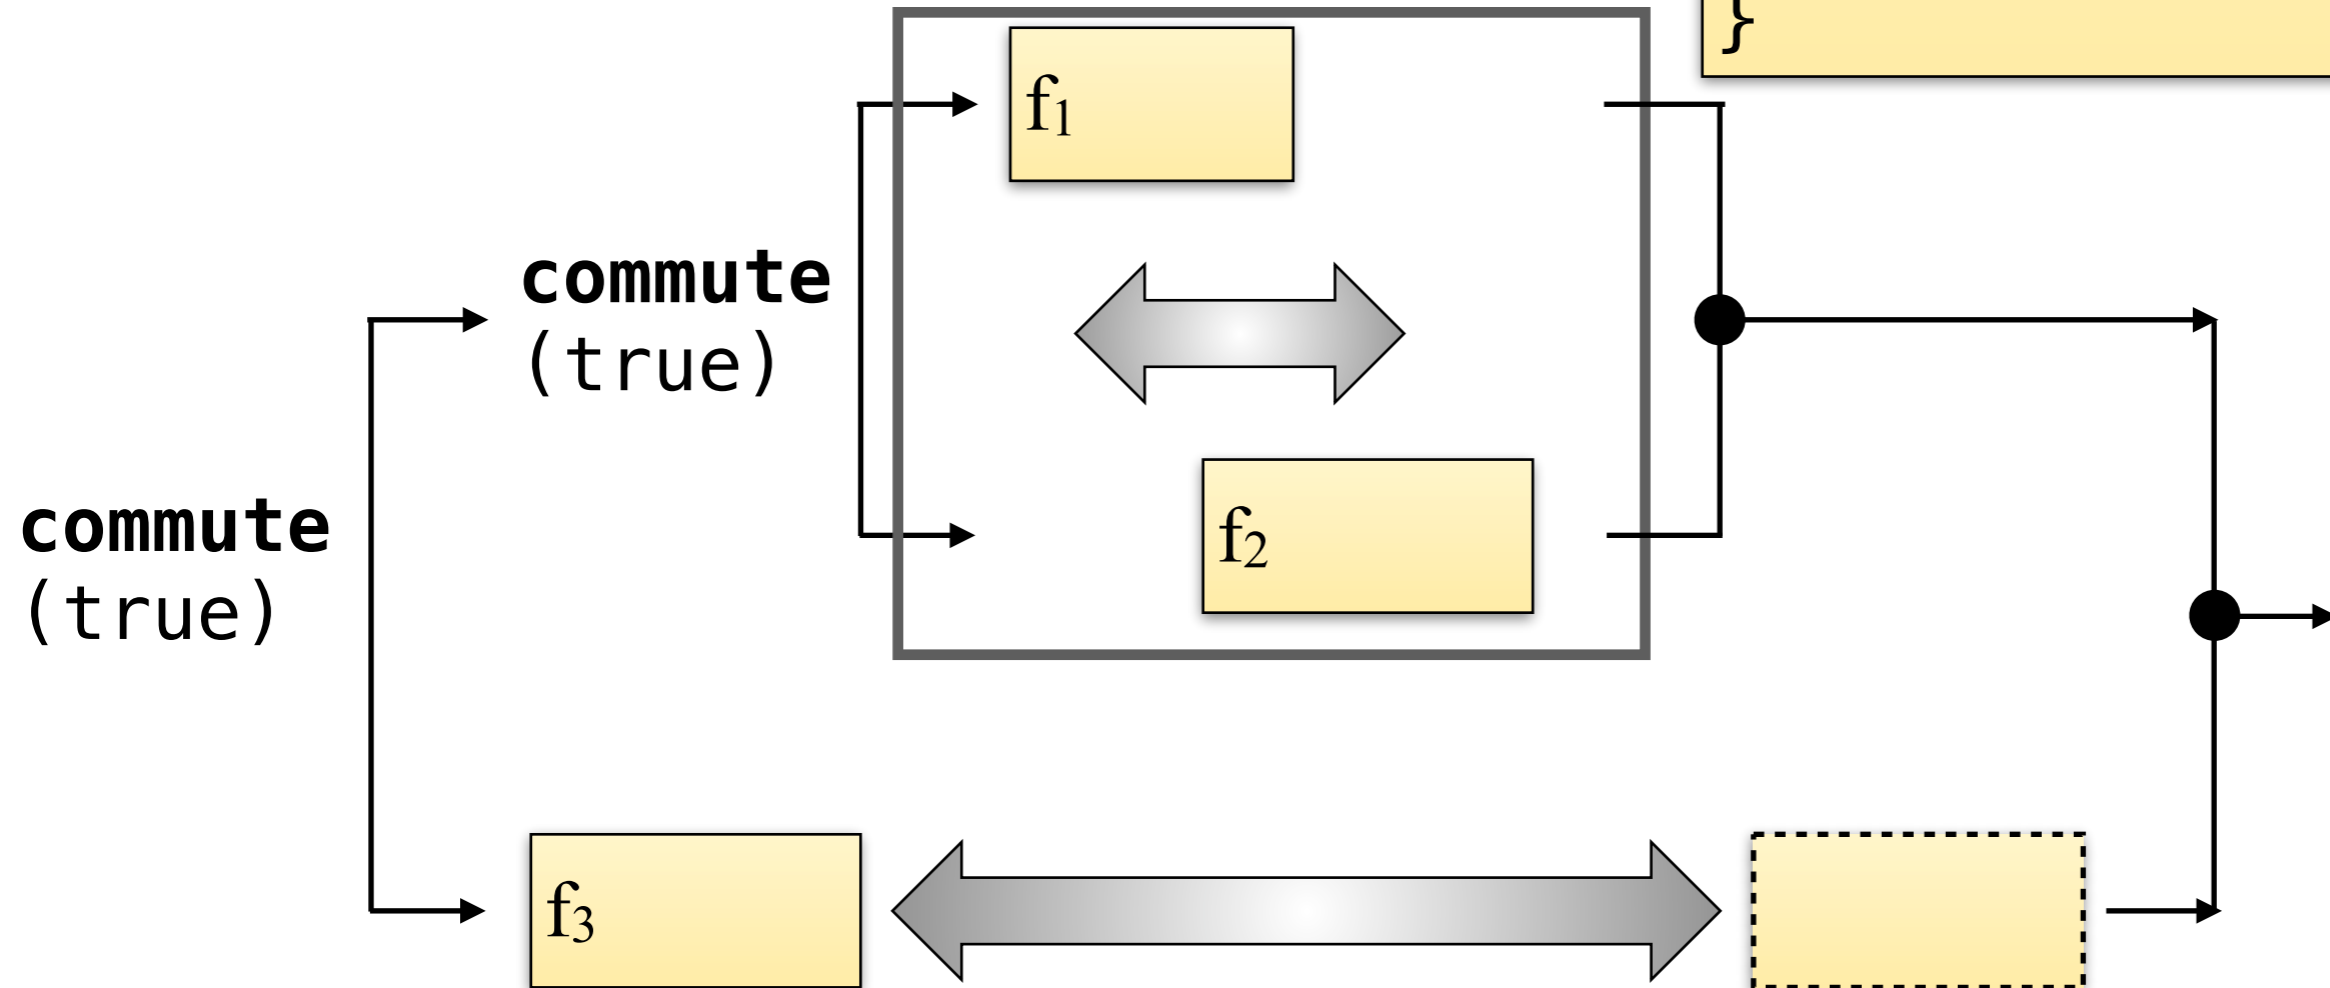
Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

Serializability?

```

y = 0; x = 1;
commute(true) {
  { commute(true)
    f1: { x = 0; }
    f2: { x = x*2; } }
  f3: { if(x >= 2) y = 1; }
}
  
```



$[x=x*2] \cdot [if(x \geq 2) y=1] \cdot [x=0]$ ❌

f_2 f_3 f_1

2. Semantic Implications & Correctness Criteria

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

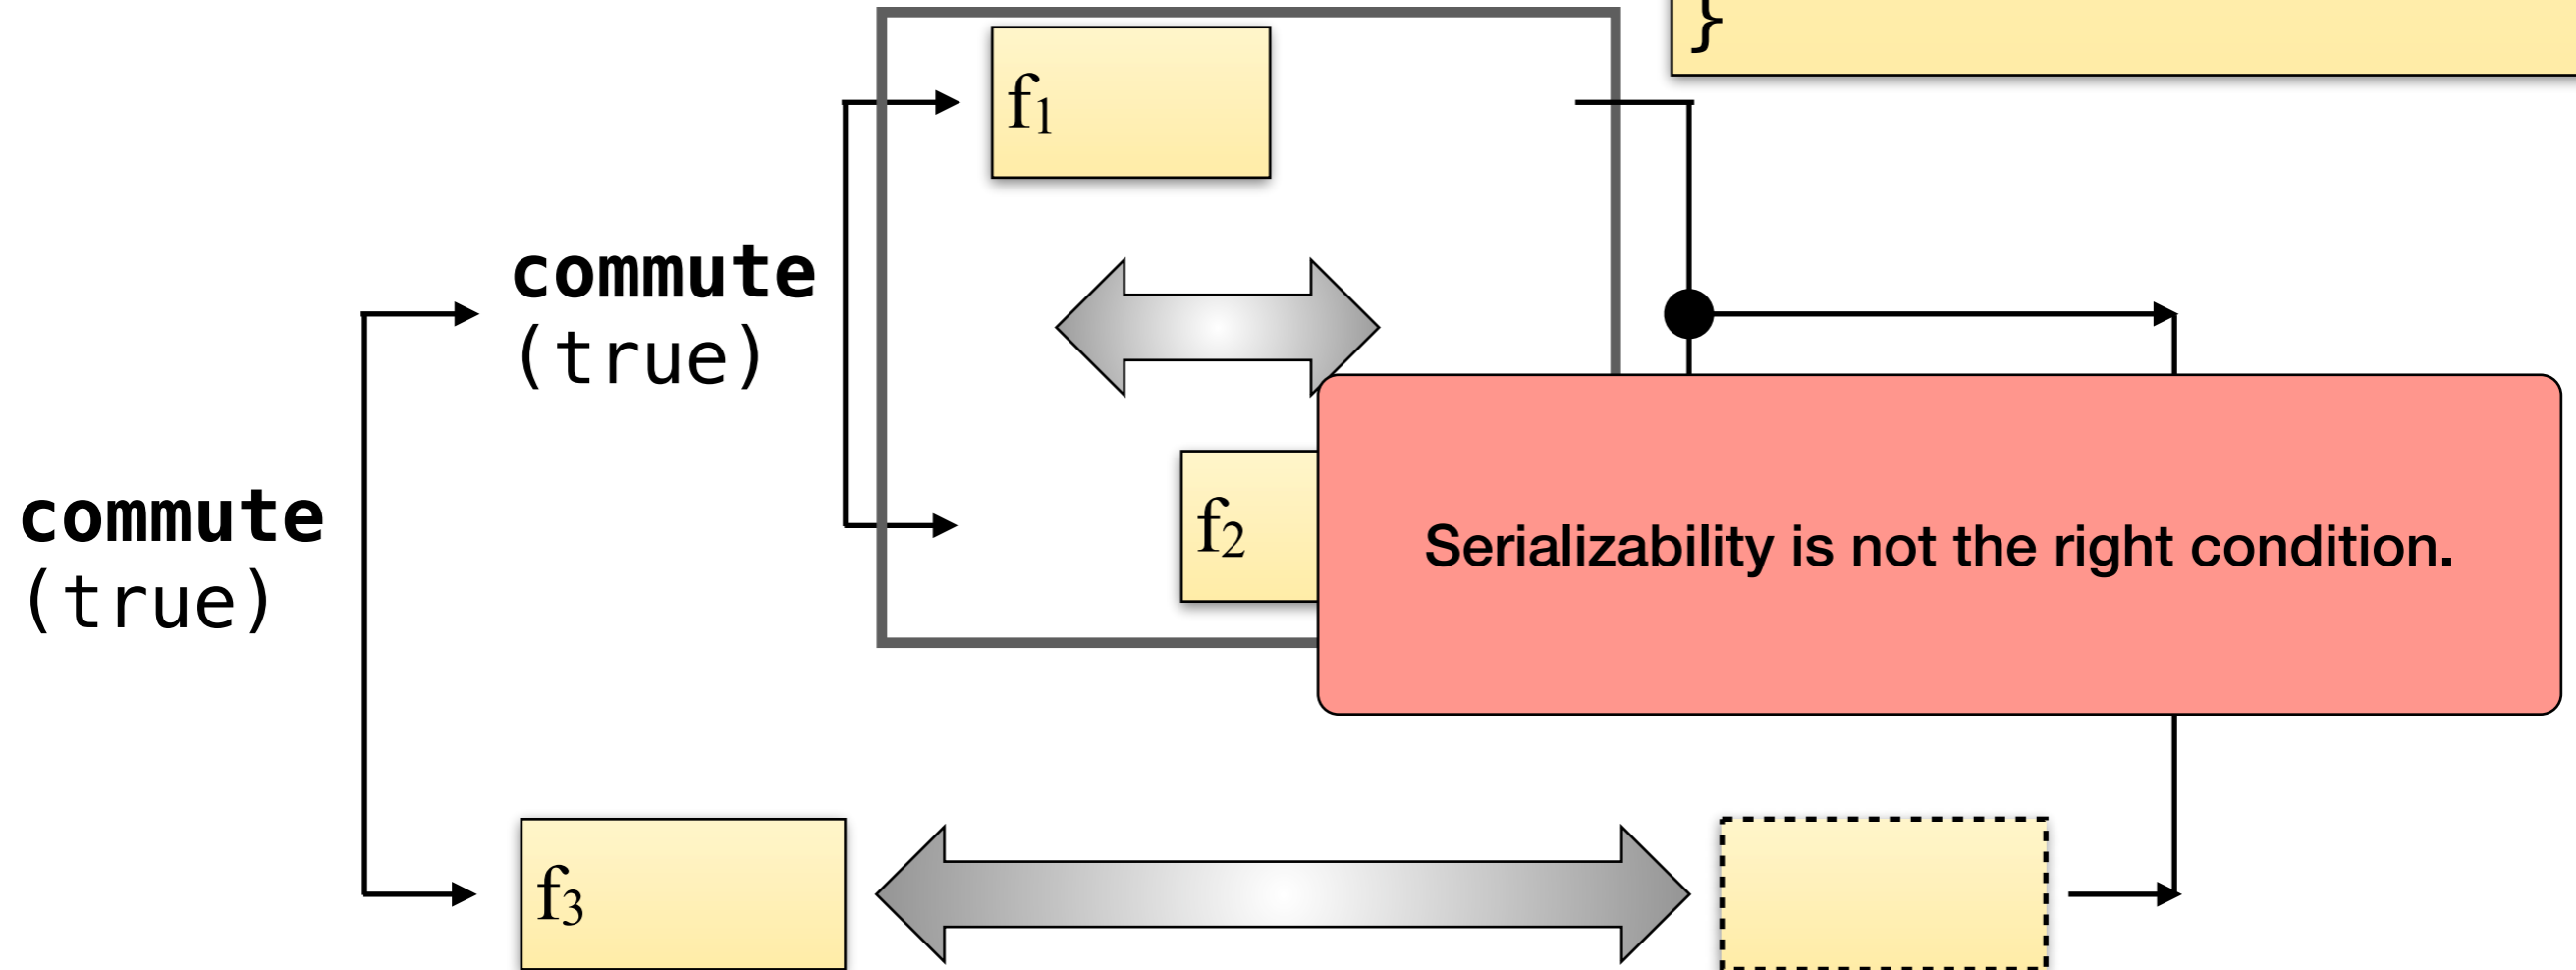
How to ensure equivalence.

Serializability?

```

y = 0; x = 1;
commute(true) {
  { commute(true)
    f1: { x = 0; }
    f2: { x = x*2; } }
  f3: { if(x >= 2) y = 1; }
}

```



$[x=x*2] \cdot [if(x \geq 2) y=1] \cdot [x=0]$ ❌

f_2 f_3 f_1

2. Semantic Implications & Correctness Criteria

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

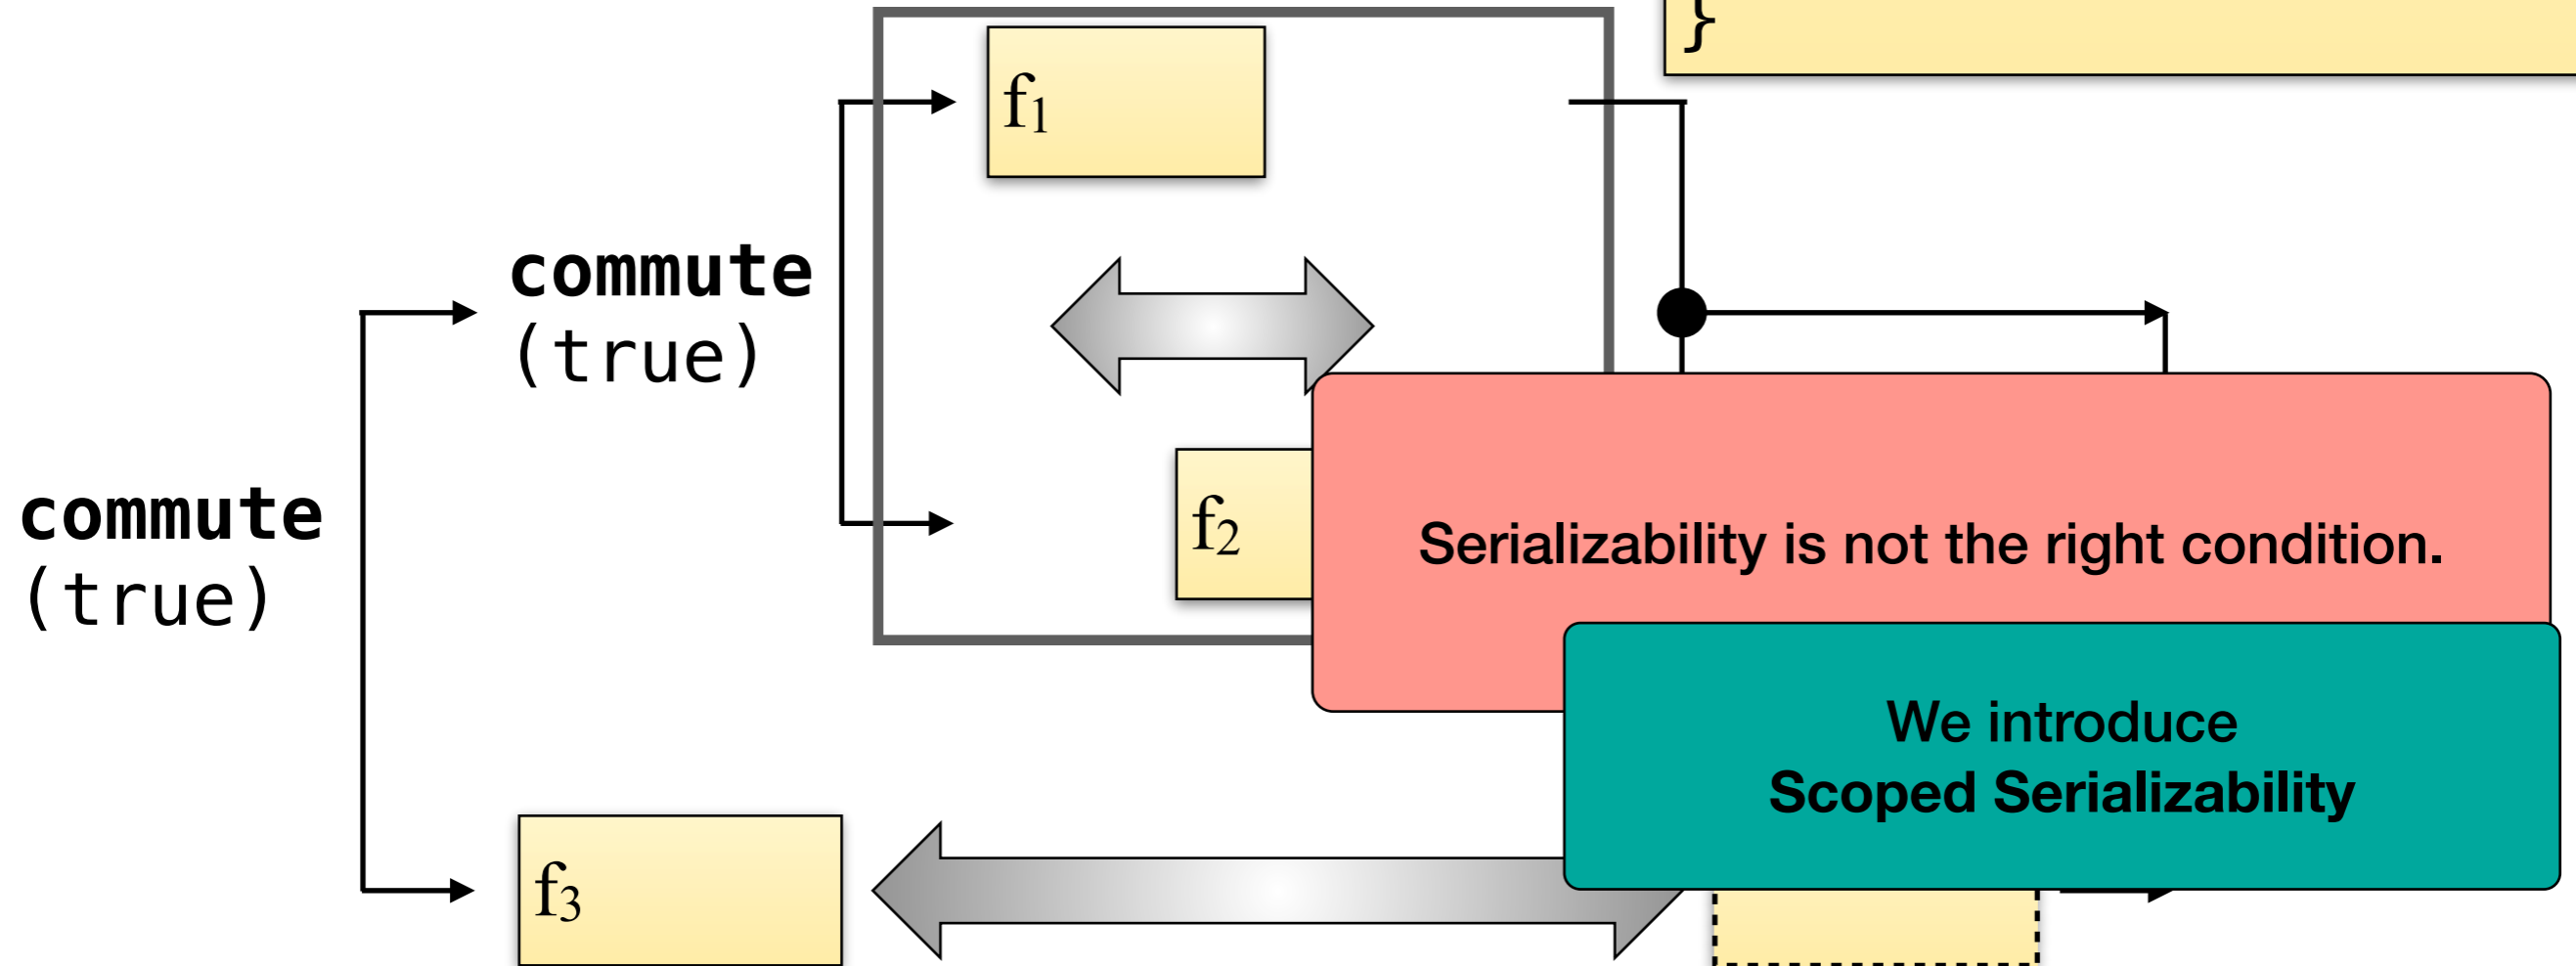
How to ensure equivalence.

Serializability?

```

y = 0; x = 1;
commute(true) {
  { commute(true)
    f1: { x = 0; }
    f2: { x = x*2; } }
  f3: { if(x >= 2) y = 1; }
}

```



$[x=x*2] \cdot [if(x \geq 2) y=1] \cdot [x=0]$ ❌

f_2 f_3 f_1

2. Semantic Implications & Correctness Criteria

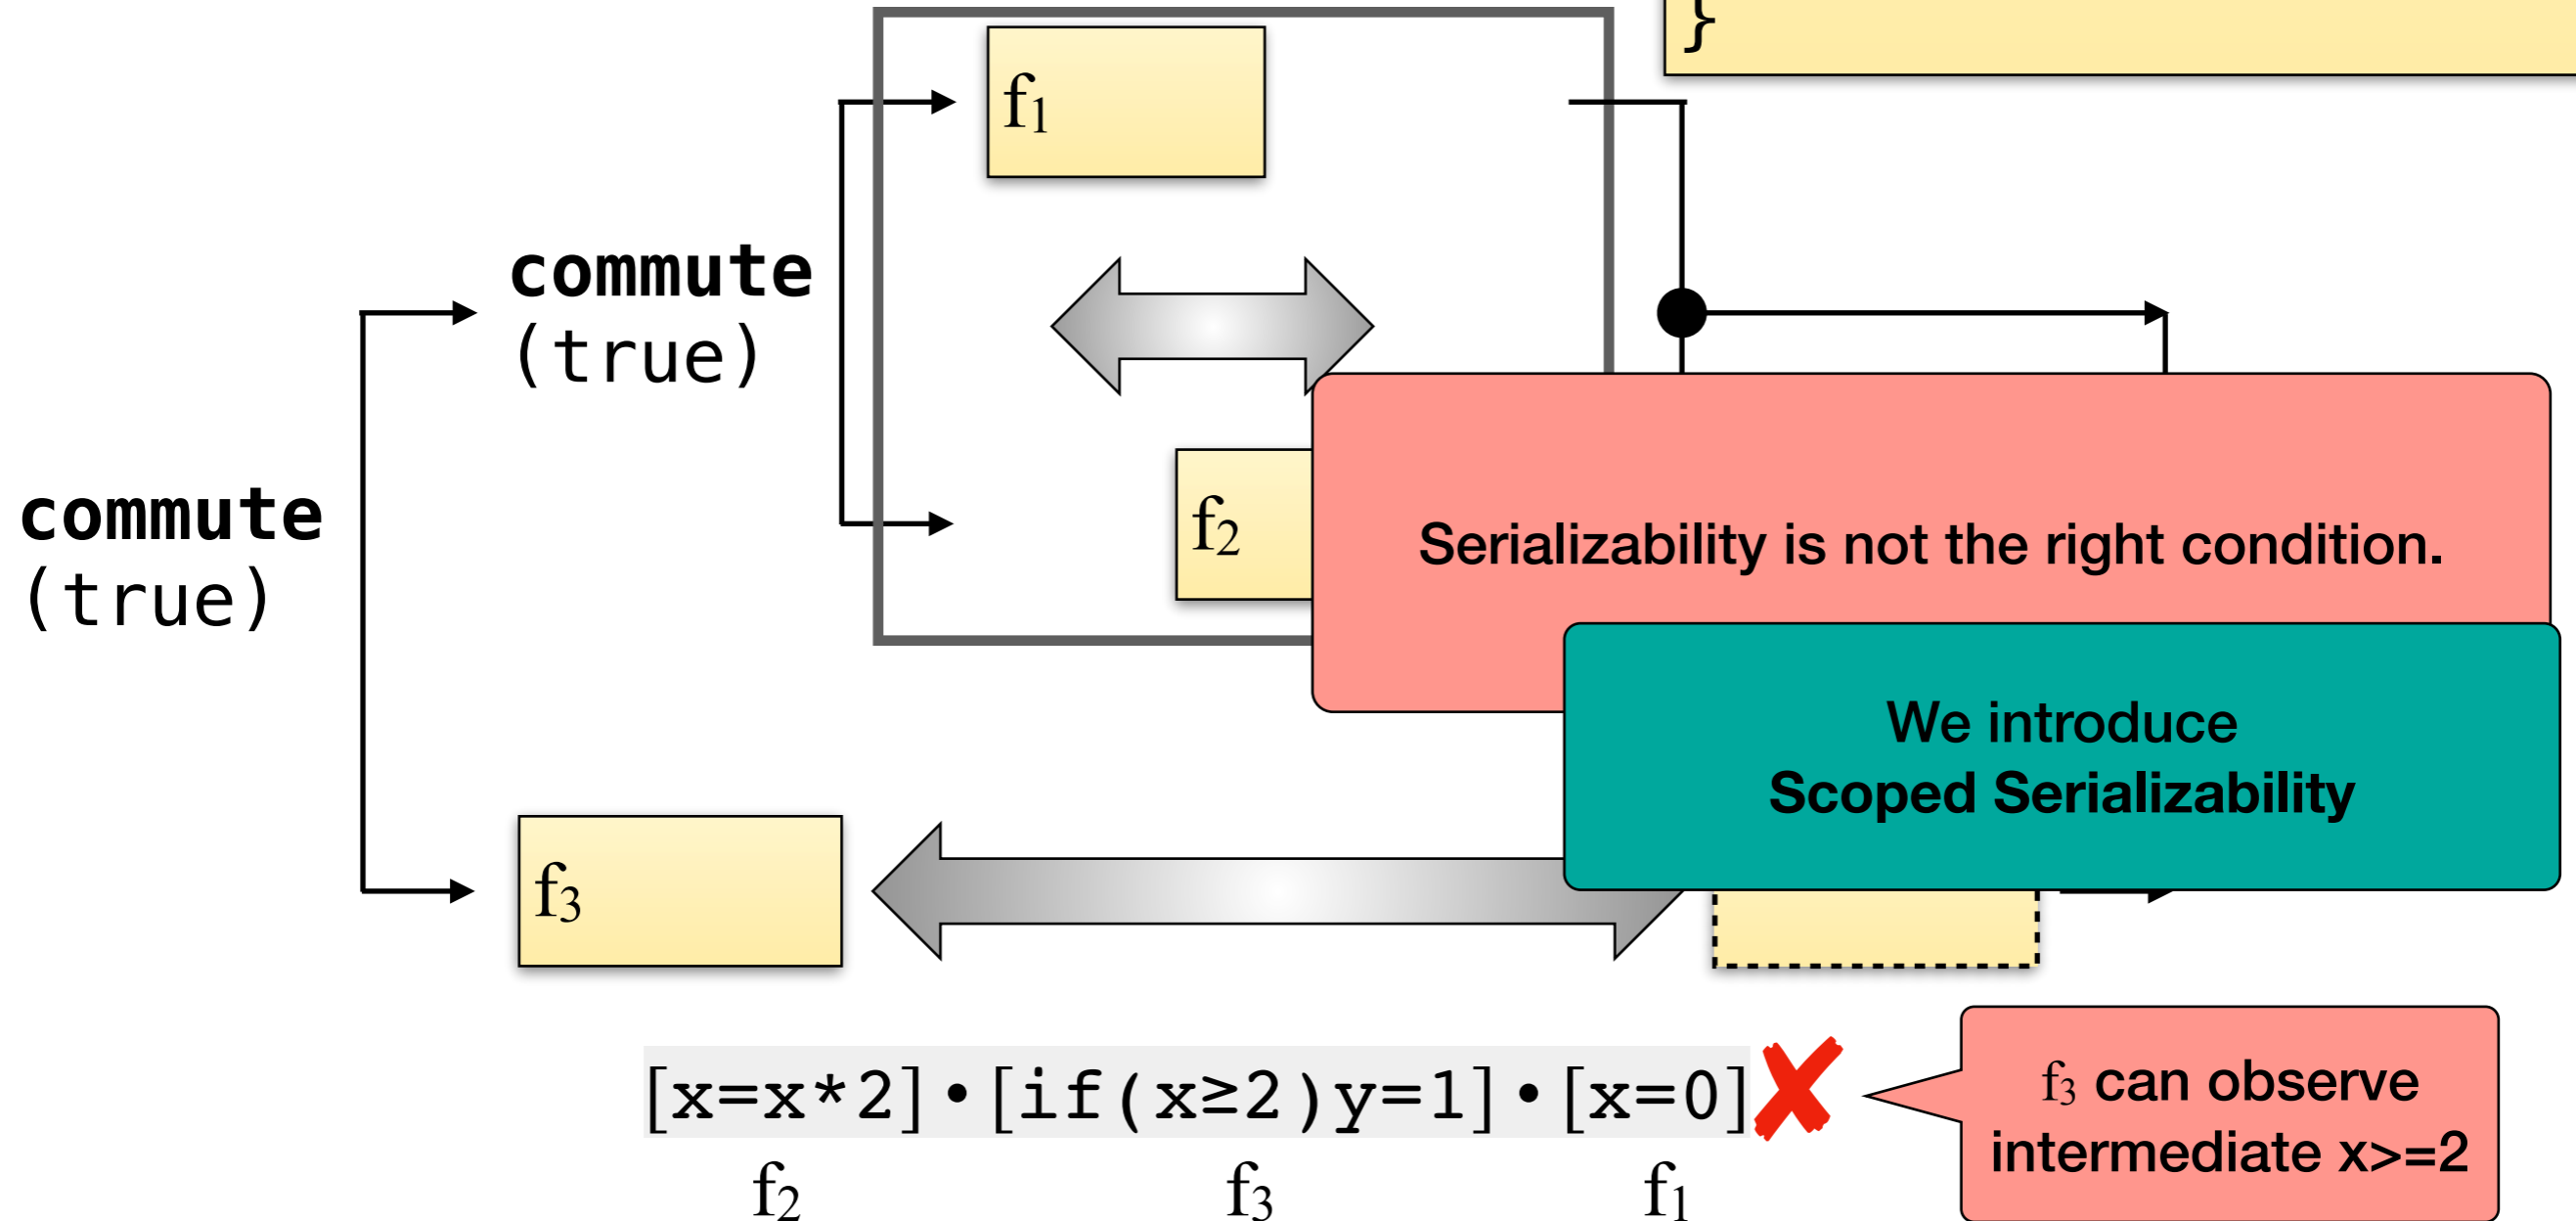
Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

Serializability?

```

y = 0; x = 1;
commute(true) {
  { commute(true)
    f1: { x = 0; }
    f2: { x = x*2; } }
  f3: { if(x >= 2) y = 1; }
}
  
```



Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

5.2 Scoped Serializability

We now define our correctness condition. We begin with a single execution:

Definition 5.2 (Scoped serial execution). Execution ε is scoped serial if:

$$\begin{aligned}
 & \forall p \in \{L_n, R_n \mid n \in \mathbb{N}\}^* : \\
 & ((\forall \ell, \ell' \in \varepsilon : \ell.fr \text{ has prefix } p \cdot L_k \wedge \ell'.fr \text{ has prefix } p \cdot R_k \implies \ell \leq_\varepsilon \ell') \\
 & \vee (\forall \ell, \ell' \in \varepsilon : \ell'.fr \text{ has prefix } p \cdot L_k \wedge \ell.fr \text{ has prefix } p \cdot R_k \implies \ell' \leq_\varepsilon \ell))
 \end{aligned}$$

Above, $\ell.fr$ is the fragment label of ℓ . The key idea here is to identify the *scope* of commute fragments through labels, and then require a serializability condition for the L/R pair of the given scope. Consider, e.g., a single commute block, possibly with children. For an execution to be scoped-serial, we require all of the transitions from one of the fragments to execute prior to all the transitions from its co-fragment (the other statement in the commute). Next, when there are nested commute blocks, the quantification over prefixes requires that we expect this same property to hold locally for all nested commute blocks. Without nesting, we recover the standard notion of serial. We now

Intuition: There exists a reordering of every interleaving into a serial order H in which *pairs of commute blocks are adjacent* in H .

Definition 5.3 (Scoped Serializability). For execution ε such that $\varepsilon \Downarrow_{par} \varepsilon$, ε is scoped serializable, if

Goal: $[[s]]_{nd} = [[s]]_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```


Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**

```
x = calc1(a);  
c = c + (x*x);
```

```
if (c > 0 && y < 0) {  
    c = c - 1;  
    z = calc2(y);  
} else {  
    z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.
- **Not ideal for commute blocks.**

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.
- **Not ideal for commute blocks.**
- **We describe new techniques:**

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.
- **Not ideal for commute blocks.**
- **We describe new techniques:**
 - Achieve shorter windows of locking

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.
- **Not ideal for commute blocks.**
- **We describe new techniques:**
 - Achieve shorter windows of locking
 - Instruction re-ordering

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

- **Synthesize Locks!**
- **Use prior works.** e.g. Flanagan & Qadeer 2003, Cherem et al 2008, Vechev et al 2010, Golan-Gueta et al 2015.
- **Not ideal for commute blocks.**
- **We describe new techniques:**
 - Achieve shorter windows of locking
 - Instruction re-ordering
 - Details in the paper

```
lock x = calc1(a);  
      c = c + (x*x); unlock  
lock if (c > 0 && y < 0) {  
      c = c - 1; unlock  
      z = calc2(y);  
} else {  
      z = calc3(y);  
}
```


Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

Theorem 5.5. If

- every commutativity condition in s is valid,
- and s is scoped-serializable,

Then s is parallelizable, ie, $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

Goal: $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$

How to ensure equivalence.

👉 **How to enforce scoped serializability.**

Theorem 5.5. If

- every commutativity condition in s is valid,
- and s is scoped-serializable,


Then s is parallelizable, ie, $\llbracket s \rrbracket_{nd} = \llbracket s \rrbracket_{par}$



- ✓ Introducing **commute** blocks
- ✓ Semantic Implications & Correctness Criteria *“Scoped Serializability”
and lock synthesis*
- 3 Demo of the Veracity language veracity-lang.org
- 4 Speedup

Part B:

- 5 Symbolic Commutativity Reasoning *TACAS'18, JAR'20, VMCAI'21*
- 6 Inferring/verifying **commute** blocks *Translation to logic*



Veracity

www.veracity-lang.org

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.
- **Support for builtin ADTs.** Arrays, Maps/Dictionaryes.

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.
- **Support for builtin ADTs.** Arrays, Maps/Dictionaries.
- **libcuckoo for parallel maps.** Li *et al*, EuroSys 2014.

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.
- **Support for builtin ADTs.** Arrays, Maps/Dictionaryes.
- **libcuckoo for parallel maps.** Li *et al*, EuroSys 2014.
- **Underlying SMT solvers.** CVC4, CVC5, Z3.

Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.
- **Support for builtin ADTs.** Arrays, Maps/Dictionaries.
- **libcuckoo for parallel maps.** Li *et al*, EuroSys 2014.
- **Underlying SMT solvers.** CVC4, CVC5, Z3.
- **Examples were written in Veracity.** Execute or inference.

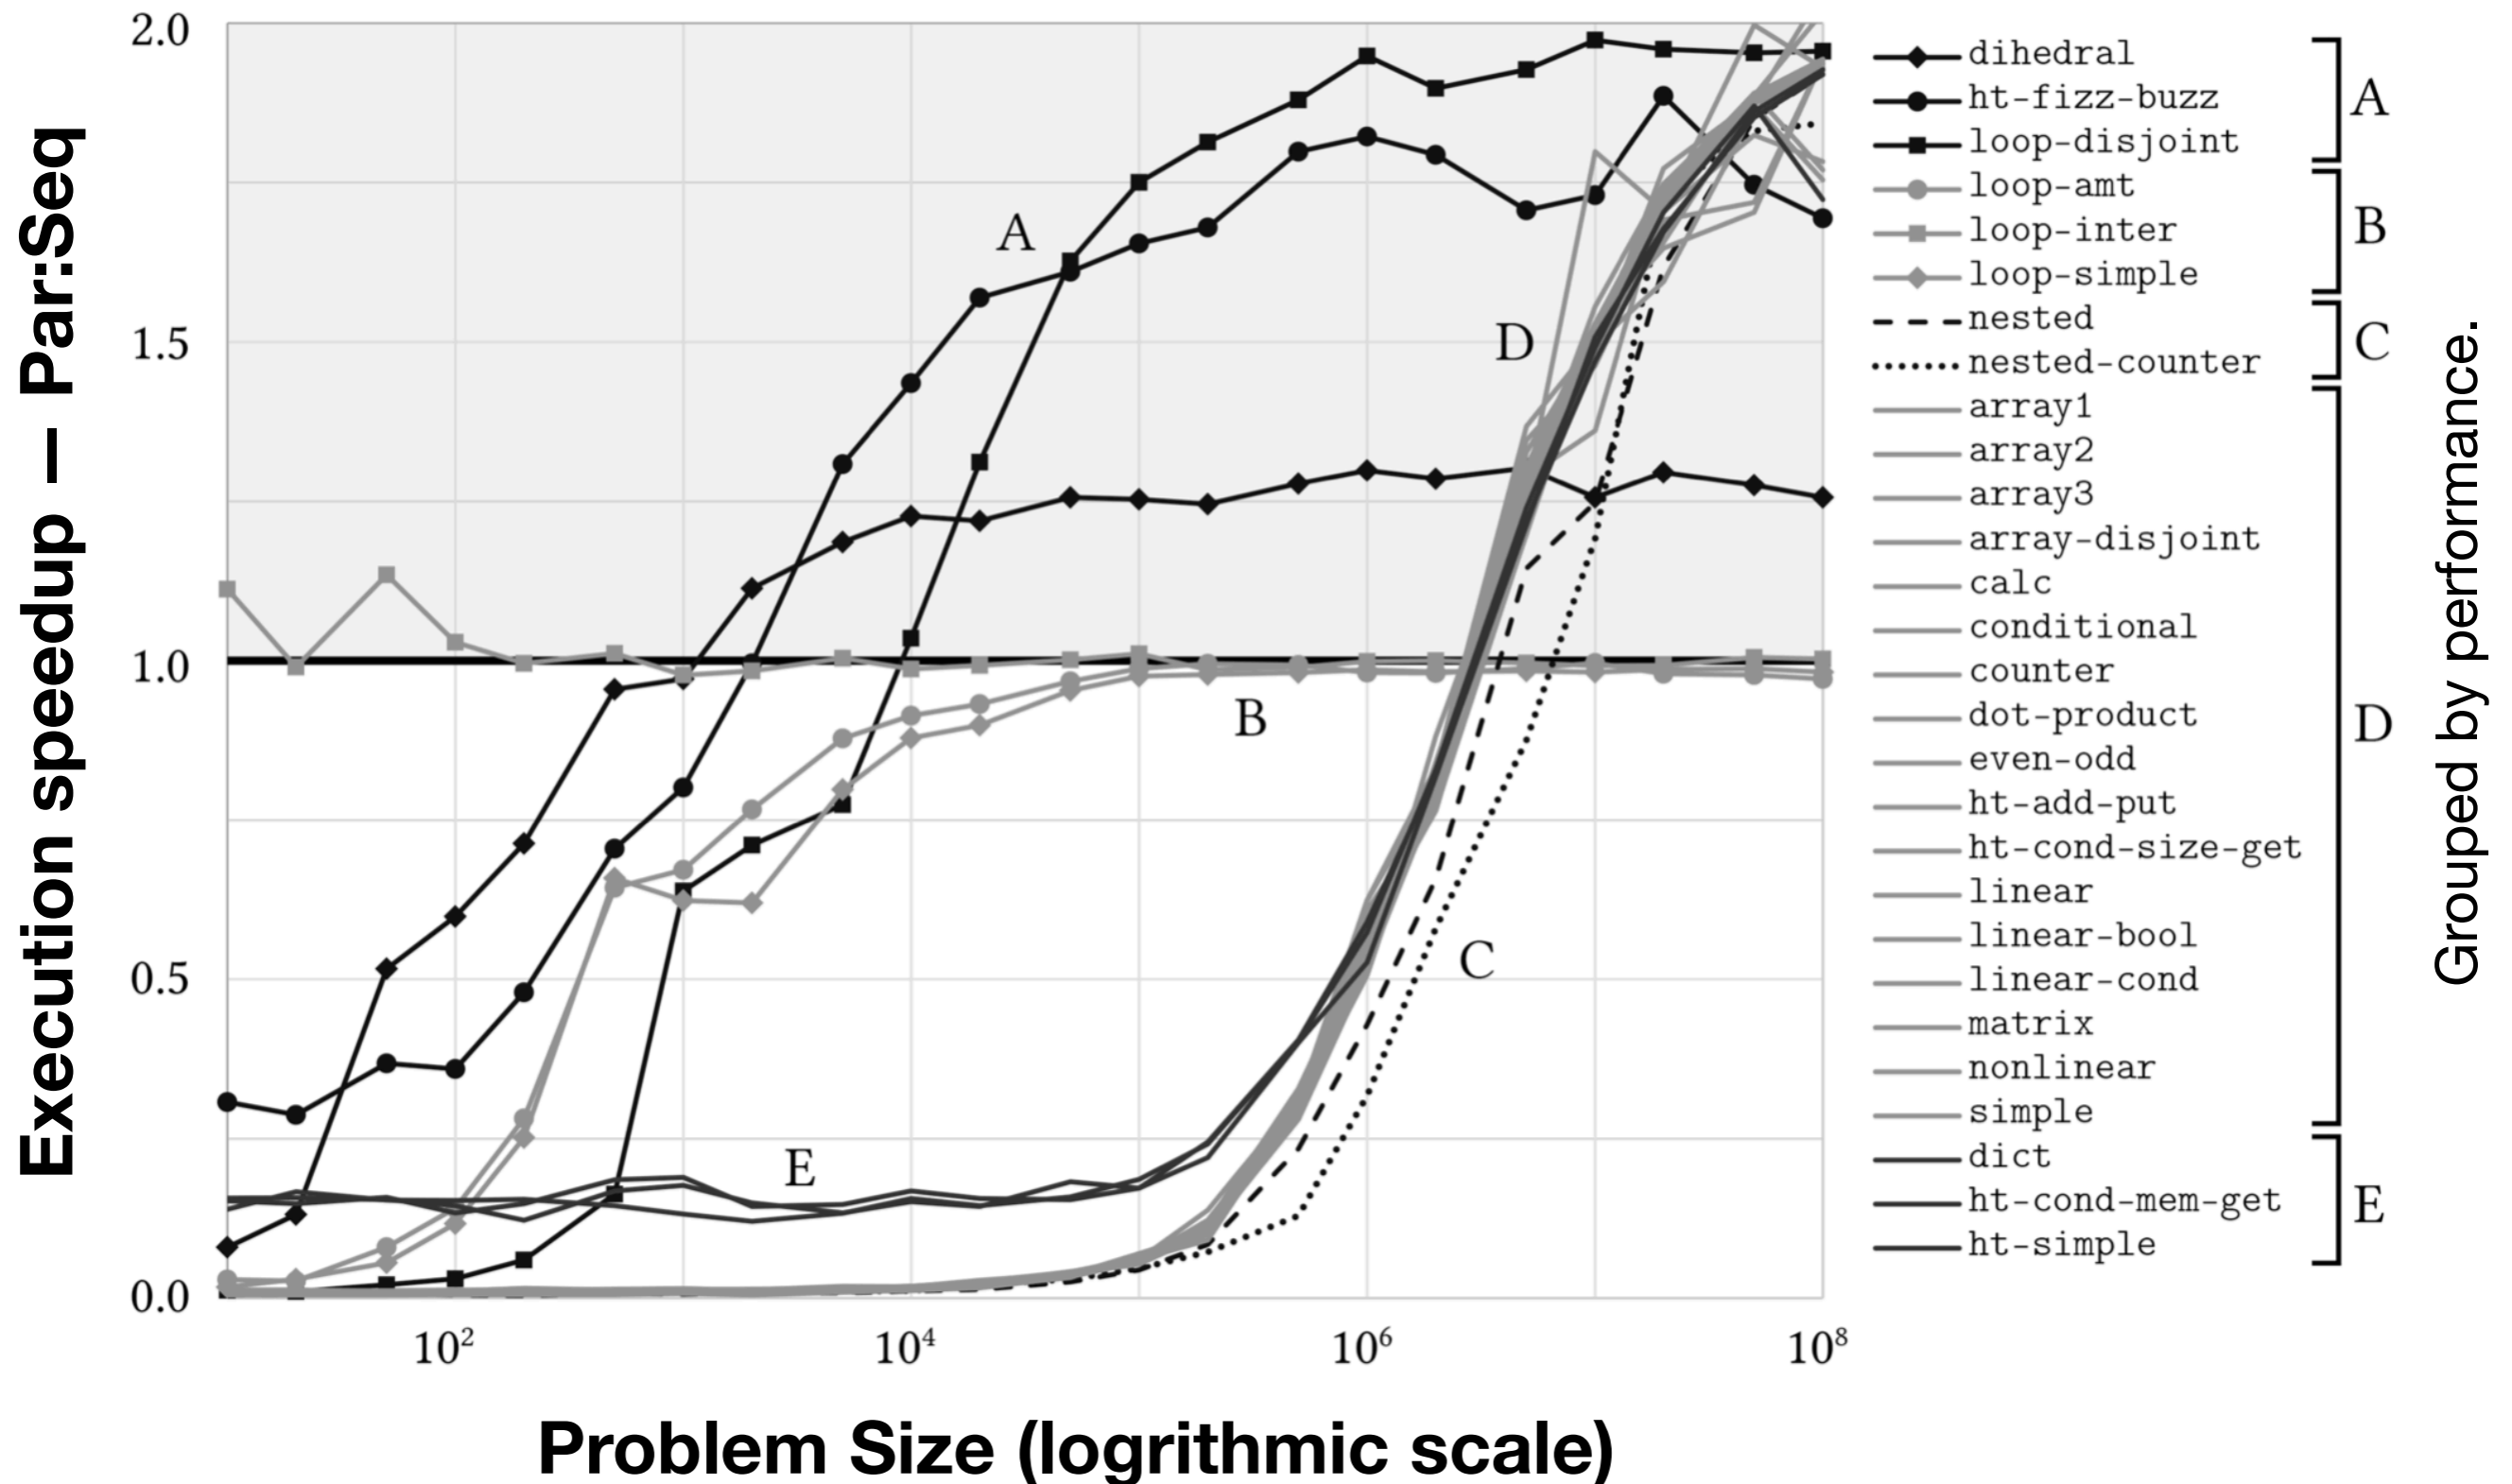
Veracity

www.veracity-lang.org

- **New language, implemented in Multicore OCaml**
- **Interpreter available.** Compiler planned.
- **Support for builtin ADTs.** Arrays, Maps/Dictionaries.
- **libcuckoo for parallel maps.** Li *et al*, EuroSys 2014.
- **Underlying SMT solvers.** CVC4, CVC5, Z3.
- **Examples were written in Veracity.** Execute or inference.

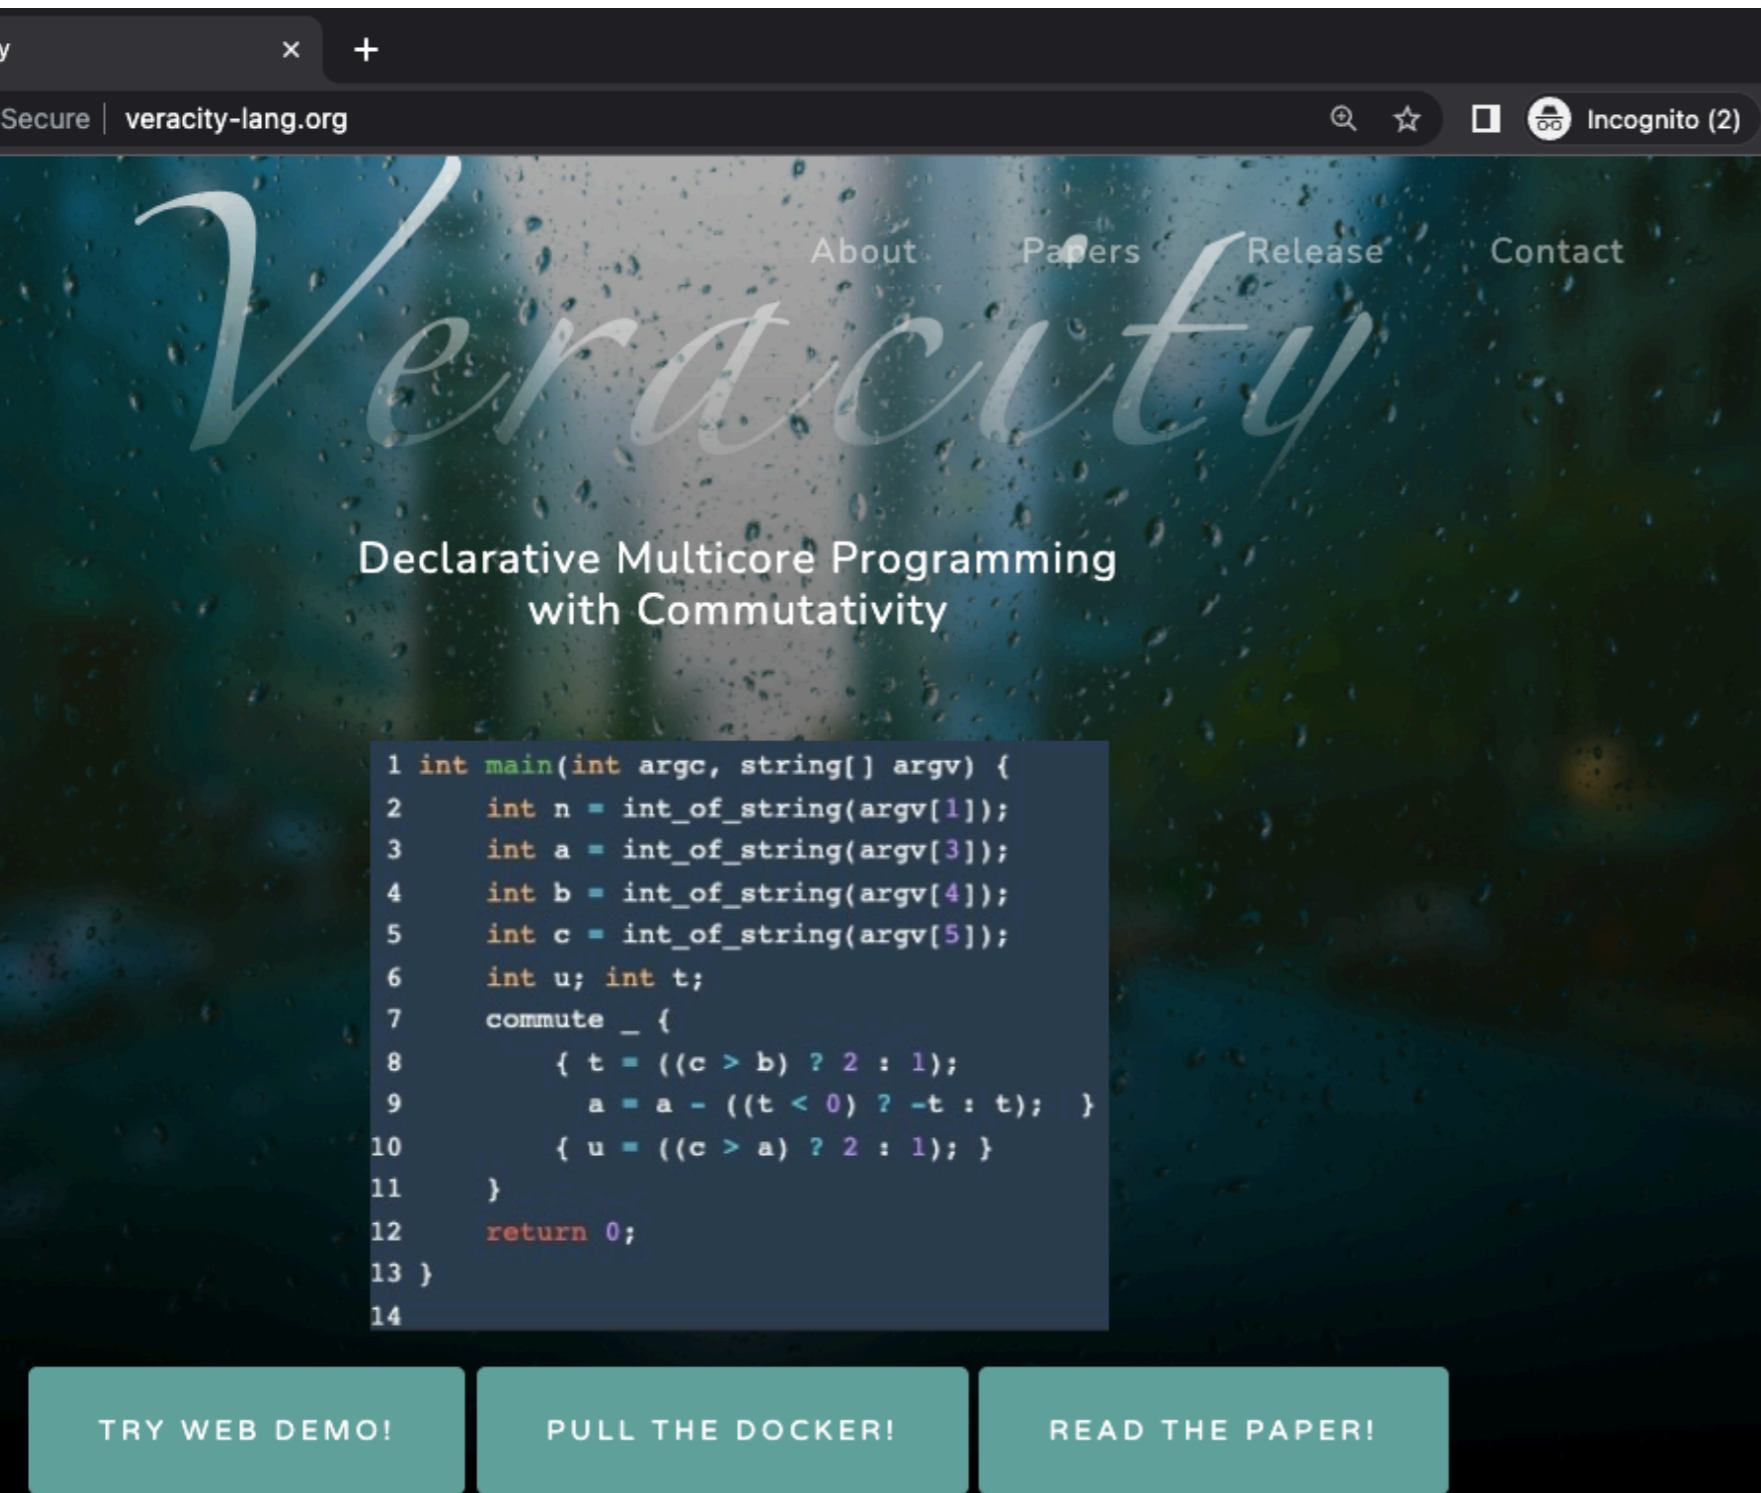
```
ejk@arran:veracity/src$ ./vcy.exe interp ../benchmarks/ht-cond-mem-get.vcy 1 2 3 4 5
Return: 0
ejk@arran:veracity/src$
```

Evaluation: Speedup versus problem size



Programming with Commutativity

www.veracity-lang.org



The screenshot shows the homepage of the Veracity website. The background is a dark, textured image with water droplets. The word "Veracity" is written in a large, light-colored, cursive font. Below it, the text "Declarative Multicore Programming with Commutativity" is displayed. A code block is shown in the lower-left corner, containing C code for a program that demonstrates commutativity. At the bottom of the page, there are three buttons: "TRY WEB DEMO!", "PULL THE DOCKER!", and "READ THE PAPER!".

```
1 int main(int argc, string[] argv) {
2     int n = int_of_string(argv[1]);
3     int a = int_of_string(argv[3]);
4     int b = int_of_string(argv[4]);
5     int c = int_of_string(argv[5]);
6     int u; int t;
7     commute _ {
8         { t = ((c > b) ? 2 : 1);
9           a = a - ((t < 0) ? -t : t); }
10        { u = ((c > a) ? 2 : 1); }
11    }
12    return 0;
13 }
14
```


TRY WEB DEMO! PULL THE DOCKER! READ THE PAPER!

- Express conditional commutativity.
- Parallel speedup with sequential reasoning.
- New correctness condition.
- Infer or verify commute conditions ...

- ✓ Introducing **commute** blocks
- ✓ Semantic Implications & Correctness Criteria *“Scoped Serializability”
and lock synthesis*
- ✓ Demo of the Veracity language veracity-lang.org
- ✓ Speedup

Part B:

- 5 Symbolic Commutativity Reasoning *TACAS'18, JAR'20, VMCAI'21*
- 6 Inferring/verifying **commute** blocks *Translation to logic*



```
commute (  $\varphi$  ) {  
    { x = calc1(a);  
      c = c + (x*x); }  
    { if (c > 0 && y < 0)  
        c = c - 1;  
        z = calc2(y);  
      else  
        z = calc3(y); }  
}
```


What does it mean for `commute` condition φ to be correct?

```

commute (  $\varphi$  ) {
  { x = calc1(a);
    c = c + (x*x); }
  { if (c > 0 && y < 0)
      c = c - 1;
    z = calc2(y);
  else
    z = calc3(y); }
}

```

What does it mean for commute condition φ to be correct?





```

commute (  $\varphi$  ) {
  { x = calc1(a);
    c = c + (x*x); }
  { if (c > 0 && y < 0)
      c = c - 1;
    z = calc2(y);
  else
    z = calc3(y); }
}

```

What does it mean for commute condition φ to be correct?




```

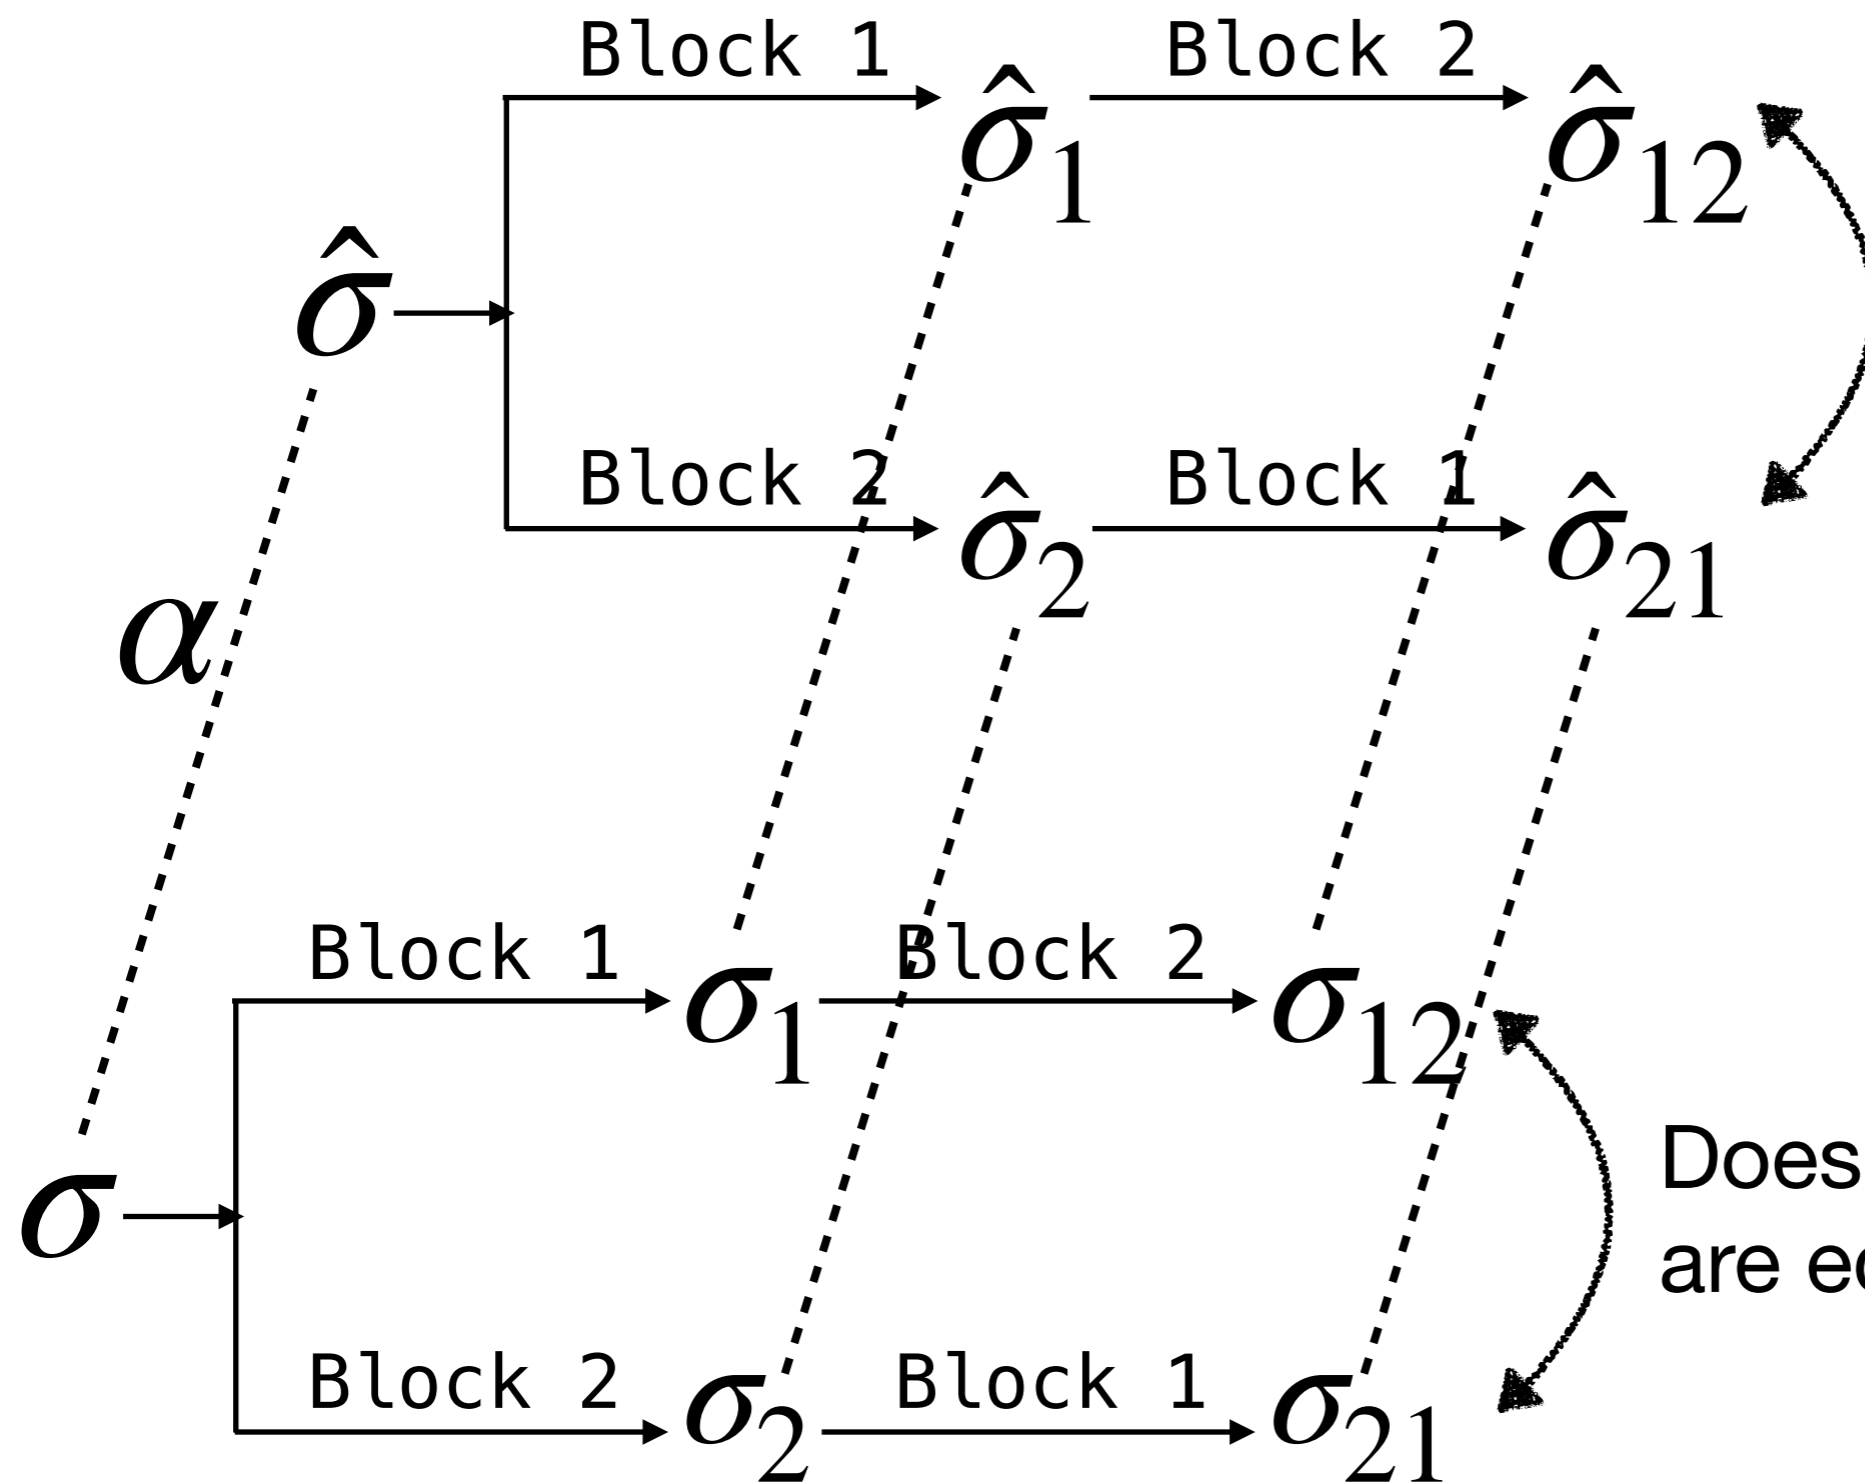
commute (  $\varphi$  ) {
  { x = calc1(a);
    c = c + (x*x); }
  { if (c > 0 && y < 0)
      c = c - 1;
    z = calc2(y);
  else
    z = calc3(y); }
}

```

Need symbolic reasoning

What does it mean for **commute** condition φ to be correct?

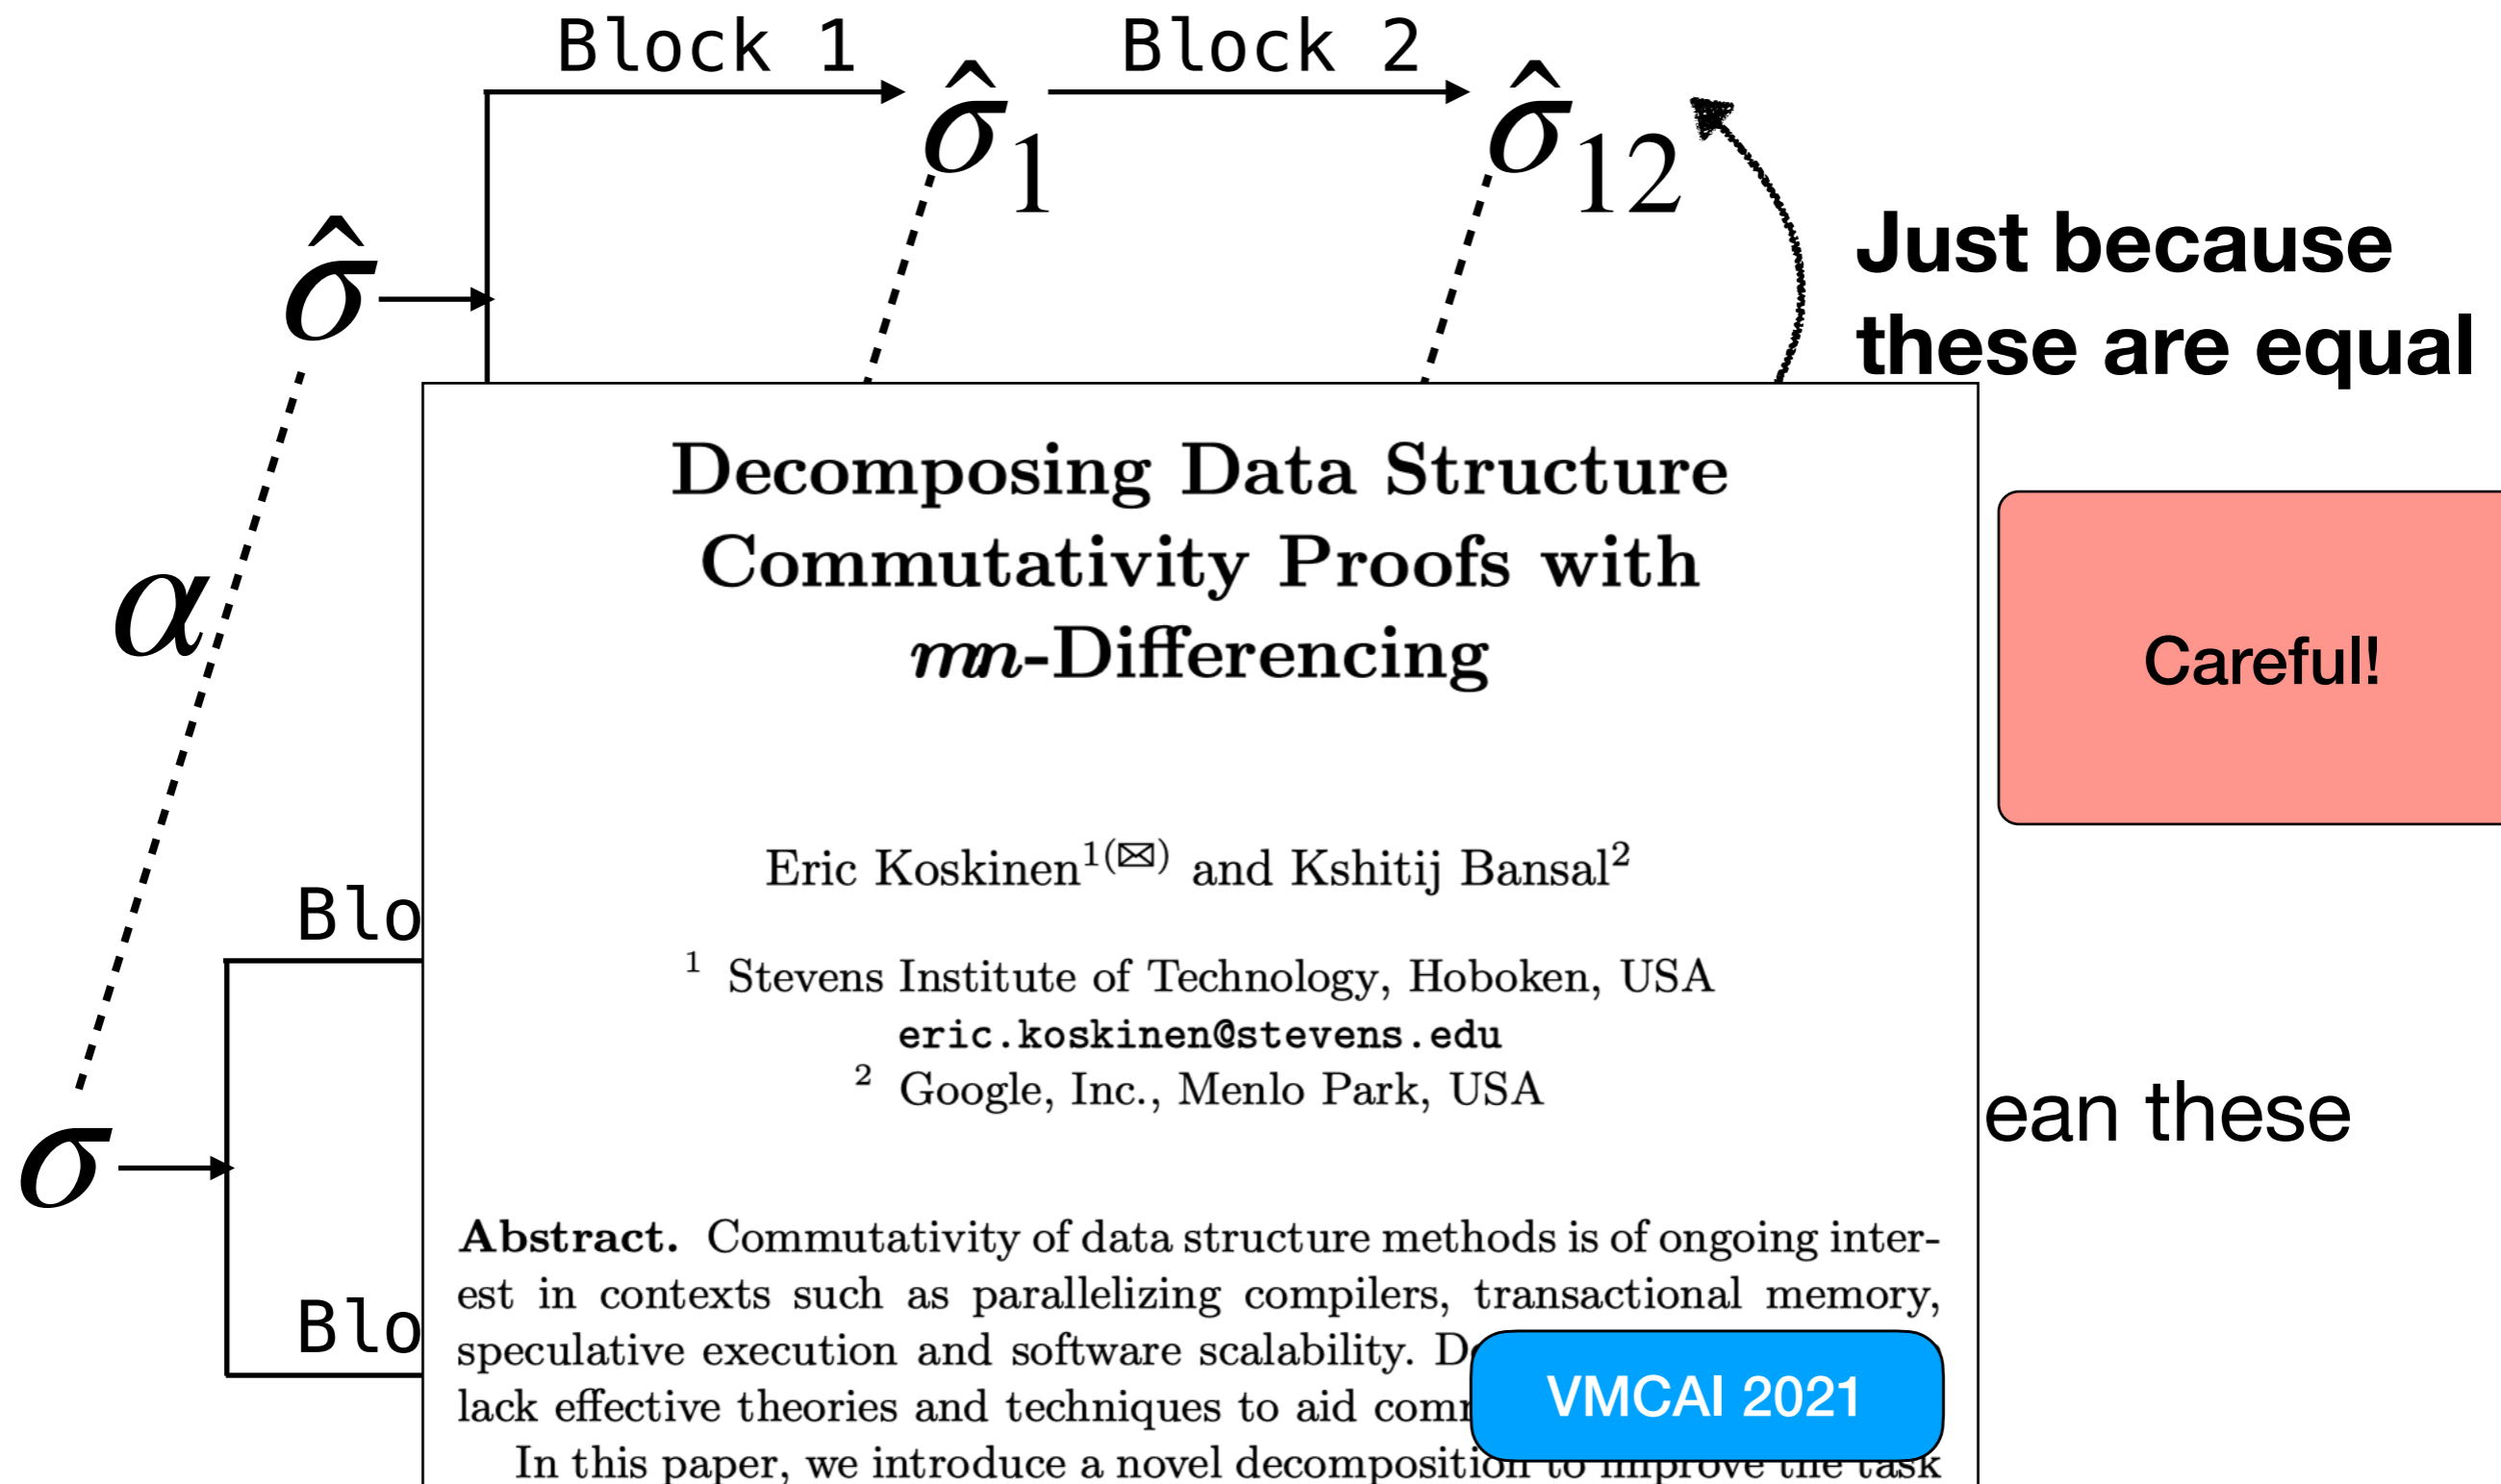




Just because these are equal

Careful!

Doesn't mean these are equal!



Decomposing Data Structure Commutativity Proofs with *mm*-Differencing

Eric Koskinen^{1(✉)} and Kshitij Bansal²

¹ Stevens Institute of Technology, Hoboken, USA
eric.koskinen@stevens.edu

² Google, Inc., Menlo Park, USA

Abstract. Commutativity of data structure methods is of ongoing interest in contexts such as parallelizing compilers, transactional memory, speculative execution and software scalability. Data structure methods lack effective theories and techniques to aid commutativity reasoning. In this paper, we introduce a novel decomposition to improve the task

VMCAI 2021

Just because these are equal

Careful!

can these

Proof Rule for Decomposing Commutativity Reasoning

- (i) : $\{I_\beta\} [s_m]^1(\bar{x}) \mid [s_m]^2(\bar{y}) \{I_\beta\}$
- (ii) : $\{Rch \wedge \varphi_m^n\} \begin{matrix} [s_m]^1(\bar{x}); \\ [s_n]^1(\bar{x}) \end{matrix} \mid \begin{matrix} [s_n]^2(\bar{y}); \\ [s_m]^2(\bar{y}) \end{matrix} \{R_\alpha \wedge C\}$
- (iii) : $(R_\alpha \wedge C) \implies I_\beta$

φ_m^n is a commut. cond. for $m(\bar{x}) \bowtie n(\bar{y})$

Decomposing Data Structure Commutativity Proofs with *m*-Differencing

Eric Koskinen^{1(✉)} and Kshitij Bansal²

¹ Stevens Institute of Technology, Hoboken, USA
eric.koskinen@stevens.edu

² Google, Inc., Menlo Park, USA

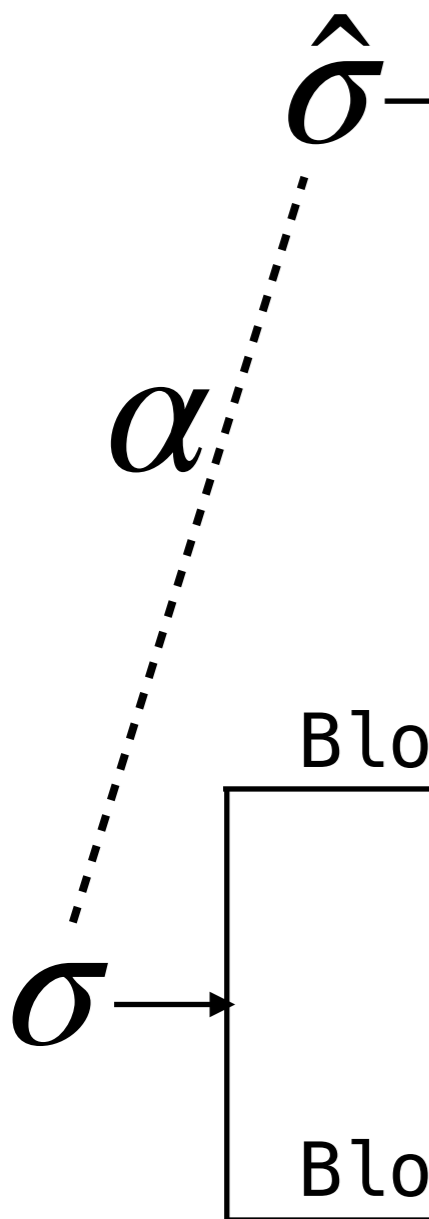
Abstract. Commutativity of data structure methods is of ongoing interest in contexts such as parallelizing compilers, transactional memory, speculative execution and software scalability. Despite the lack of effective theories and techniques to aid commutativity reasoning, these

In this paper, we introduce a novel decomposition to improve the task

Careful!

ean these

VMCAI 2021



ADT Impl.

```
insert() {...}
remove() {...}
contains()...
```

Commutate
Cond.
 φ_m^n

Verification of Semantic Commutativity Conditions and Inverse Operations on Linked Data Structures

Deokhwan Kim Martin C. Rinard

Massachusetts Institute of Technology
{dkim,rinard}@csail.mit.edu

Abstract

We present a new technique for verifying *commutativity conditions*, which are logical formulas that characterize when operations commute. Because our technique reasons with the abstract state of verified linked data structure implementations, it can verify commuting operations that produce semantically equivalent (but not necessarily identical) data structure states in different execution orders. We have used this technique to verify sound and complete commutativity conditions for all pairs of operations on a collection of linked data structure implementations, including data structures that export a set interface (ListSet and HashSet) as well as data structures that export a map interface (AssociationList, HashTable, and ArrayList). This effort involved the specification and verification of 765 commutativity conditions.

Many speculative parallel systems need to undo the effects of speculatively executed operations. *Inverse operations*, which undo these effects, are often more efficient than alternate approaches (such as saving and restoring data structure state). We present a new technique for verifying such inverse operations. We have specified and verified, for all of our linked data structure implementations, an inverse operation for every operation that changes the data structure state.

Together, the commutativity conditions and inverse operations provide a key resource that language designers, developers of program analysis systems, and implementors of software systems can draw on to build languages, program analyses, and systems with strong correctness guarantees.

Categories and Subject Descriptors D.1.3 [Programming Techniques]: Concurrent Programming; D.2.4 [Software Engineering]:

- **Deterministic Parallel Languages:** Including support for commuting operations in deterministic parallel languages increases the expressive power of the language while preserving guaranteed deterministic parallel execution [5, 42].
- **Transaction Monitors:** If a transaction monitor can detect that operations within parallel transactions commute, it can use efficient locking algorithms that allow commuting operations from different transactions to interleave [17, 49]. Because such locking algorithms place fewer constraints on the execution order, they increase the amount of exploitable parallelism.
- **Irregular Parallel Computations:** Exploiting commuting operations has been shown to be critical for obtaining good parallel performance in irregular parallel computations that manipulate linked data structures [28–30]. The reason is essentially the same as for efficient transaction monitors — it enables the use of efficient synchronization algorithms for atomic transactions that execute multiple (potentially commuting) operations on shared objects. For similar reasons, exploiting commuting operations has also been shown to be essential for obtaining good parallel performance for the SPEC benchmarks [7].

Despite the importance of commuting operations, there has been relatively little research in automatically analyzing or verifying the conditions under which operations commute. Indeed, the deterministic parallel language, transaction monitor, and irregular parallel computation systems cited above all rely on the developer to identify commuting operations, with no way to determine whether the operations do, in fact, commute or not. A mistake in identifying commuting operations invalidates both the principles upon which the systems operate and the correctness guarantees that they claim

ADT Impl.

```
insert() {...}
remove() {...}
contains()...
```

Commutate
Cond.
 φ_m^n

Verification of Semantic Commutativity Conditions and Inverse Operations on Linked Data Structures

Deokhwan Kim Martin C. Rinard

Massachusetts Institute of Technology
{dkim,rinard}@csail.mit.edu

Jahob Verification System

Jahob is a verification system for programs written in a subset of Java. Using Jahob, you can verify that your methods satisfy their contracts in all possible executions, as well as that they preserve design constraints.

Jahob is now on github: <https://github.com/epfl-lara/jahob>

Note

- Information below may be outdated
- Java may be outdated. Consider Scala, <http://www.scala-lang.org/>
- To verify Scala, consider tools such as <http://leon.epfl.ch> and its successor

[Some of the data structures verified in Jahob](#)

You may wish to compare

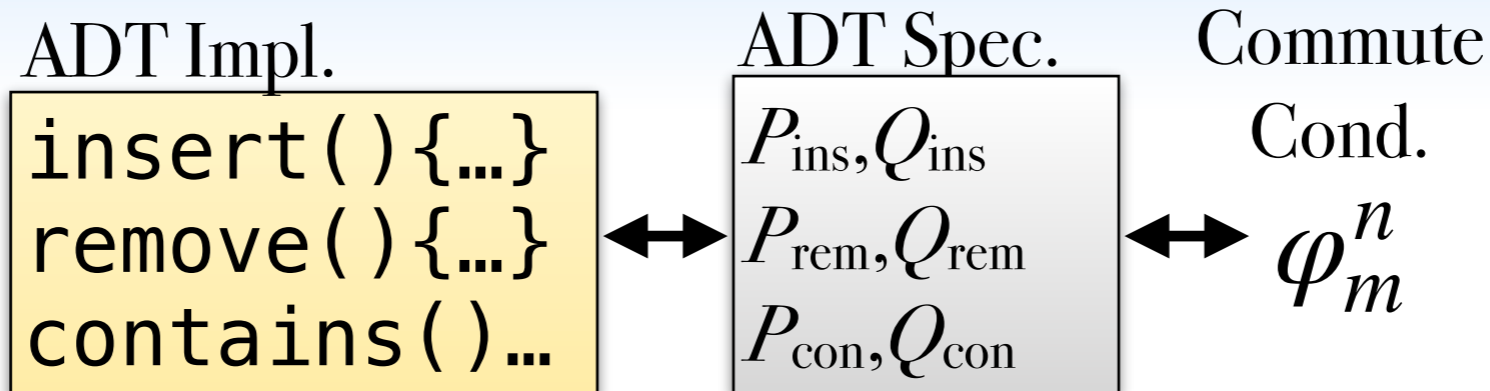
- [ArrayList class JavaDoc](#) to

Parallel Languages: Including support for transactions in deterministic parallel languages increases the expressive power of the language while preserving deterministic parallel execution [5, 42].

Monitors: If a transaction monitor can detect when parallel transactions commute, it can use algorithms that allow commuting operations to interleave [17, 49]. Because such monitors place fewer constraints on the execution of transactions, they increase the amount of exploitable parallelism.

Irregular Computations: Exploiting commuting operations is shown to be critical for obtaining good performance from irregular parallel computations that manipulate linked data structures [28–30]. The reason is essentially the lack of efficient transaction monitors — it enables the use of synchronization algorithms for atomic transactions that allow multiple (potentially commuting) operations to execute in parallel. For similar reasons, exploiting commuting operations has also been shown to be essential for obtaining high performance for the SPEC benchmarks [7].

In the absence of commuting operations, there has been little progress in automatically analyzing or verifying the correctness of programs with operations that do not commute. Indeed, the deterministic semantics of a transaction monitor, and irregular parallel computations cited above all rely on the developer to identify which operations commute, with no way to determine whether they do, commute or not. A mistake in identifying which operations commute invalidates both the principles upon which the monitor is based and the correctness guarantees that they claim.



Verification of Semantic Commutativity Conditions and Inverse Operations on Linked Data Structures

Deokhwan Kim Martin C. Rinard

Massachusetts Institute of Technology
{dkim,rinard}@csail.mit.edu

Jahob Verification System

Jahob is a verification system for programs written in a subset of Java. Using Jahob, you can verify that your methods satisfy their contracts in all possible executions, as well as that they preserve design constraints.

Jahob is now on github: <https://github.com/epfl-lara/jahob>

Note

- Information below may be outdated
- Java may be outdated. Consider Scala, <http://www.scala-lang.org/>
- To verify Scala, consider tools such as <http://leon.epfl.ch> and its successor

[Some of the data structures verified in Jahob](#)

You may wish to compare


- [ArrayList](#) class JavaDoc to

Parallel Languages: Including support for transactions in deterministic parallel languages increases the expressive power of the language while preserving the semantics of deterministic parallel execution [5, 42].

Monitors: If a transaction monitor can detect when parallel transactions commute, it can use algorithms that allow commuting operations to interleave [17, 49]. Because such monitors place fewer constraints on the execution of transactions, they increase the amount of exploitable parallelism.

Irregular Computations: Exploiting commuting operations is shown to be critical for obtaining good performance from irregular parallel computations that manipulate linked data structures [28–30]. The reason is essentially the lack of efficient transaction monitors — it enables the use of synchronization algorithms for atomic transactions that allow multiple (potentially commuting) operations to execute in parallel. For similar reasons, exploiting commuting operations has also been shown to be essential for obtaining good performance for the SPEC benchmarks [7].

In the absence of commuting operations, there has been much research in automatically analyzing or verifying the commutativity of operations. Indeed, the deterministic transaction monitor, and irregular parallel computations cited above all rely on the developer to identify operations that commute or not. A mistake in identifying commuting operations invalidates both the principles upon which the monitor is based and the correctness guarantees that they claim.



Automatic Generation of Precise and Useful Commutativity Conditions

Kshitij Bansal^{1*}, Eric Koskinen^{2†}, and Omer Tripp^{1‡}




	$m(\bar{x})$	$n(\bar{y})$	Simple	Poke	φ_n^m (Poke)
			Qs (time)	Qs (time)	
Counter	decrement \bowtie decrement		3 (0.1)	3 (0.1)	true
	increment \triangleright decrement		10 (0.3)	34 (0.9)	$\neg(0 = c)$
	decrement \triangleright increment		3 (0.1)	3 (0.1)	true
	decrement \bowtie reset		2 (0.1)	2 (0.1)	false
	decrement \bowtie zero		6 (0.1)	26 (0.6)	$\neg(1 = c)$
	increment \bowtie increment		3 (0.1)	3 (0.1)	true
	increment \bowtie reset		2 (0.0)	2 (0.1)	false
	increment \bowtie zero		10 (0.3)	34 (0.8)	$\neg(0 = c)$
	reset \bowtie reset		3 (0.1)	3 (0.1)	true
	reset \bowtie zero		9 (0.2)	30 (0.6)	$0 = c$
zero \bowtie zero		3 (0.1)	3 (0.1)	true	
Acum.	increase \bowtie increase		3 (0.1)	3 (0.1)	true
	increase \bowtie read		13 (0.3)	28 (0.6)	$c + x_1 = c$
	read \bowtie read		3 (0.0)	3 (0.0)	true
Set	add \bowtie add		10 (0.4)	140 (4.4)	$(y_1 = x_1 \wedge y_1 \in S) \vee \neg(y_1 = x_1)$
	add \bowtie contains		10 (0.4)	122 (3.6)	$x_1 \in S \vee (\neg(x_1 \in S) \wedge \neg(y_1 = x_1))$
	add \bowtie getsize		6 (0.2)	31 (0.9)	$x_1 \in S$
	add \bowtie remove		6 (0.2)	66 (2.2)	$\neg(y_1 = x_1)$
	contains \bowtie contains		3 (0.1)	3 (0.1)	true
	contains \bowtie getsize		3 (0.1)	3 (0.1)	true
	contains \bowtie remove		17 (0.5)	160 (4.8)	$S \setminus \{x_1\} = \{y_1\} \vee (\dots \wedge y_1 \in \{x_1\}) \vee$

technology

between data-structure operations including parallelizing are recently, Ethereum smart on automatic generation of ware of any fully automated both sound and effective. by an algorithm that iter- of the commutativity (and methods into an increasingly /when the entire state space t any time to obtain a par- e have generalized our work e completeness. We describe l commutativity conditions, ng refinement and heuristics on.

prototype open-source tool SER- e queries that are dispatched VOIS through two case stud- ons for a range of data struc- or, Coun- called E- y-related- ficient in-

TACAS 2018



Veracity: Declarative Multicore Programming with Commutativity

ADAM CHEN, Stevens Institute of Technology, USA

PARISA FATHOLOLUMI, Stevens Institute of Technology, USA

ERIC KOSKINEN, Stevens Institute of Technology, USA

JARED PINCUS, Stevens Institute of Technology, USA

There is an ongoing effort to provide programming abstractions that ease the burden of exploiting multicore hardware. Many programming abstractions (*e.g.*, concurrent objects, transactional memory, etc.) simplify matters, but still involve intricate engineering. We argue that some difficulty of multicore programming can be meliorated through a declarative programming style in which programmers directly express the independence of fragments of sequential programs.

In our proposed paradigm, programmers write programs in a familiar, sequential manner, with the added ability to explicitly express the conditions under which code fragments sequentially commute. Putting such commutativity conditions into source code offers a new entry point for a compiler to exploit the known connection between commutativity and parallelism. We give a semantics for the programmer's sequential perspective and, under a correctness condition, find that a compiler-transformed parallel execution is equivalent to the sequential semantics. Serializability/linearizability are not the right fit for this condition, so we introduce scoped serializability and show how it can be enforced with lock synthesis techniques.

We next describe a technique for automatically verifying and synthesizing commute conditions via a new reduction from our commute blocks to logical specifications, upon which symbolic commutativity reasoning can be performed. We implemented our work in a new language called Veracity, implemented in Multicore OCaml. We show that commutativity conditions can be automatically generated across a variety of new benchmark programs, confirm the expectation that concurrency speedups can be seen as the computation increases, and apply our work to a small in-memory filesystem and an adaptation of a crowdfund blockchain smart contract.

1 INTRODUCTION

Writing concurrent programs is difficult. Researchers and practitioners, seeking to make life easier,

commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

TACAS'18/JAR'20

commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}


Servois
Synthesis
TACAS'18

φ_m^n

Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

SERVOIS



Automatic Generation of
Precise and Useful Commutativity Conditions


Kshitij Bansal^{1*}, Eric Koskinen^{2†}, and Omer Tripp^{1‡}

¹ Google, Inc.
² Stevens Institute of Technology




Abstract. Reasoning about commutativity between data-structure operations is an important problem with applications including parallelizing compilers, optimistic parallelization and, more recently, Ethereum smart contracts. There have been research results on automatic generation of commutativity conditions, yet we are unaware of any fully automated technique to generate conditions that are both sound and effective. We have designed such a technique, driven by an algorithm that iteratively refines a conservative approximation of the commutativity (and non-commutativity) condition for a pair of methods into an increasingly precise version. The algorithm terminates if/when the entire state space has been considered, and can be aborted at any time to obtain a partial yet sound commutativity condition. We have generalized our work to left-/right-movers [27] and proved relative completeness. We describe aspects of our technique that lead to *useful* commutativity conditions, including how predicates are selected during refinement and heuristics that impact the output shape of the condition. We have implemented our technique in a prototype open-source tool SERVOIS. Our algorithm produces quantifier-free queries that are dispatched to a back-end SMT solver. We evaluate SERVOIS through two case studies: (i) We synthesize commutativity conditions for a range of data structures including Set, HashTable. We consider an Ethereum smart contract that SERVOIS can detect serious bugs. SERVOIS can guide developers to construct re-

TACAS 2018, JAR 2020



$(\varphi_m^n, \overline{\varphi_m^n})$




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

$$\text{valid} \left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$$




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

$$\text{valid} \left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$$

$$\forall \sigma_0, \sigma_1, \sigma_2, x, y, r_m, r_n.$$



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

$$\text{valid} \left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$$

$$\forall \sigma_0, \sigma_1, \sigma_2, x, y, r_m, r_n. \\ \sigma_0 \xrightarrow{m(x)/r_m} \sigma_1 \xrightarrow{n(y)/r_n} \sigma_2 \implies$$

...

commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

valid $\left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$

$$\forall \sigma_0, \sigma_1, \sigma_2, x, y, r_m, r_n. \\ \sigma_0 \xrightarrow{m(x)/r_m} \sigma_1 \xrightarrow{n(y)/r_n} \sigma_2 \implies \\ (\exists \sigma_3. \sigma_0 \xrightarrow{n(y)/r_n} \sigma_3 \xrightarrow{m(x)/r_m} \sigma_2)) \wedge \dots$$

commute ? {
{s1} {s2}}

Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

valid $\left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$

$$\forall \sigma_0, \sigma_1, \sigma_2, x, y, r_m, r_n. \\ \sigma_0 \xrightarrow{m(x)/r_m} \sigma_1 \xrightarrow{n(y)/r_n} \sigma_2 \implies \\ (\exists \sigma_3. \sigma_0 \xrightarrow{n(y)/r_n} \sigma_3 \xrightarrow{m(x)/r_m} \sigma_2)) \wedge \dots$$

Quantifier
alternation

commute ? {
{s1} {s2}}

Translate(commute, s1, s2)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \left\{ \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right. \right\}$$


1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$ for methods $m(\bar{x})/r_m, n(\bar{y})/r_n$

valid $\left\{ \varphi_m^n(\sigma, \bar{x}, \bar{y}) \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$

$$\forall \sigma_0, \sigma_1, \sigma_2, x, y, r_m, r_n. \\ \sigma_0 \xrightarrow{m(x)/r_m} \sigma_1 \xrightarrow{n(y)/r_n} \sigma_2 \implies \\ (\exists \sigma_3. \sigma_0 \xrightarrow{n(y)/r_n} \sigma_3 \xrightarrow{m(x)/r_m} \sigma_2)) \wedge \dots$$

Quantifier
alternation

Avoid introducing quantifiers in
encoding of commutativity



Logical ADT Specification


$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition

Use a form of abstraction-refinement.

Start with candidate commutativity condition H


$$\text{valid} \left\{ H \implies m(\bar{x})/r_m \bowtie n(\bar{y})/r_n \right\}$$




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement




Refine(H)




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement




Refine(H)




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



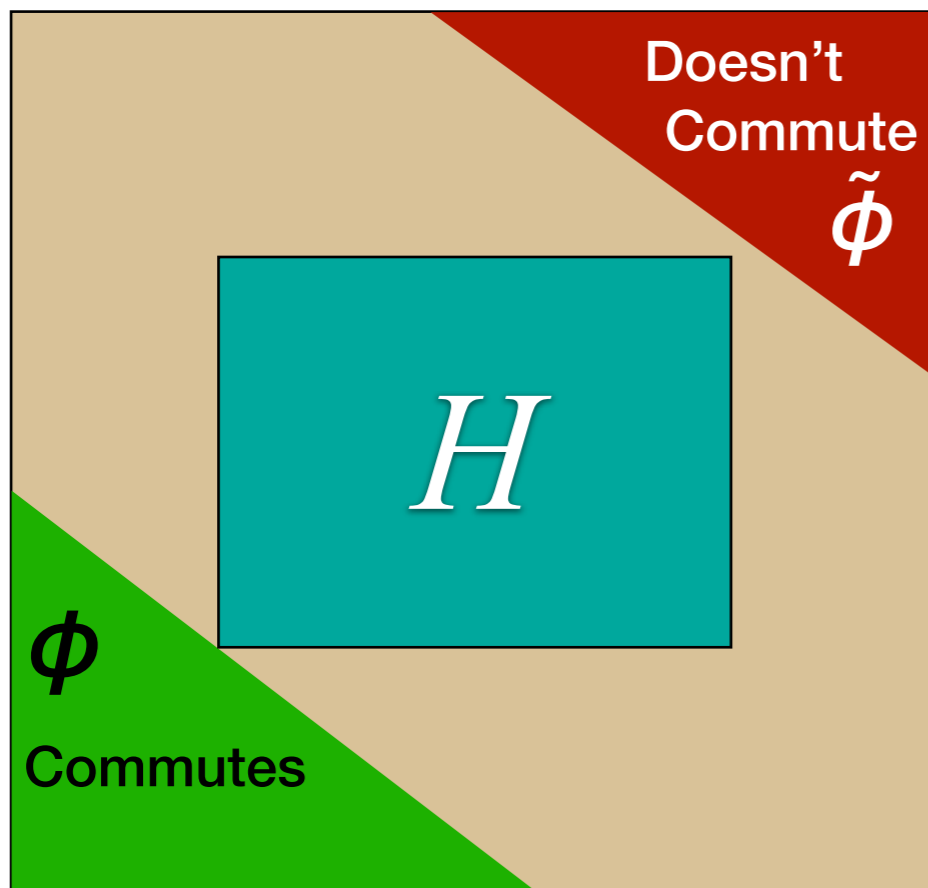
Refine(H)




Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



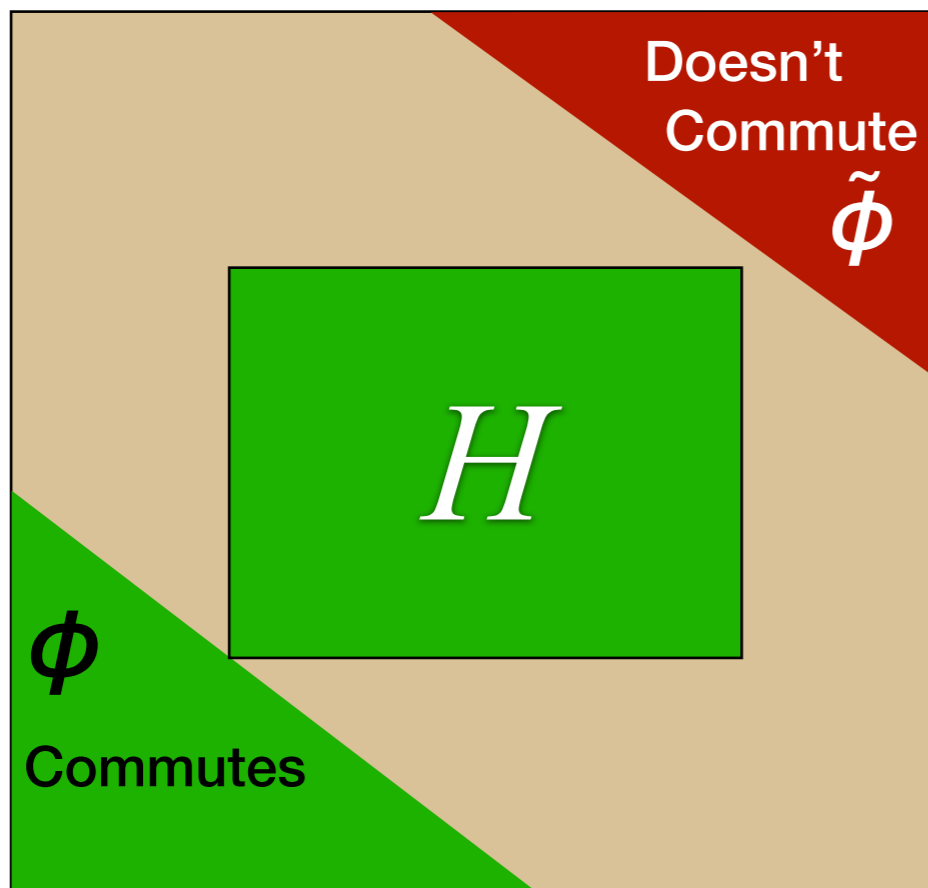
Refine(H)



Logical ADT Specification


$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)

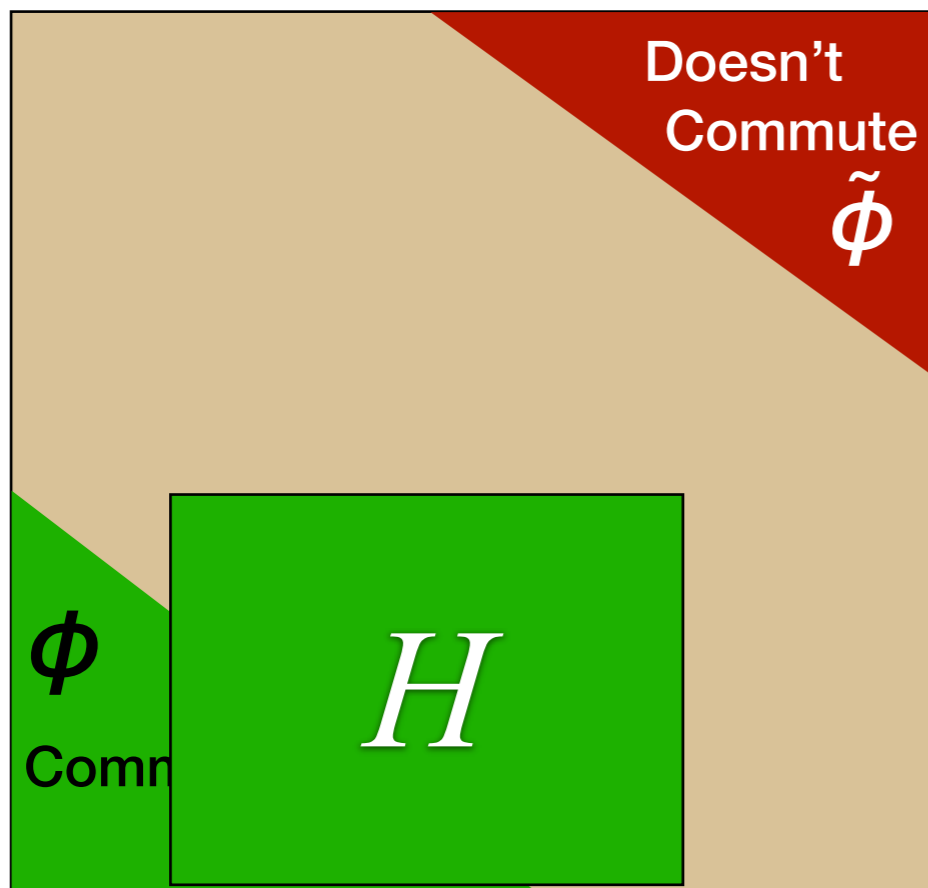
✓ If valid($H \Rightarrow m \bowtie n$): add H to ϕ



Logical ADT Specification


$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)


✓ If valid($H \Rightarrow m \bowtie n$): add H to ϕ



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$


1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)

✓ If valid($H \Rightarrow m \bowtie n$): add H to ϕ


✓ If valid($H \Rightarrow m \not\bowtie n$): add H to $\tilde{\phi}$



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$


1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)

✓ If valid($H \Rightarrow m \bowtie n$): add H to ϕ


✓ If valid($H \Rightarrow m \not\bowtie n$): add H to $\tilde{\phi}$



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$


1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)

✓ If valid($H \Rightarrow m \bowtie n$): add H to ϕ

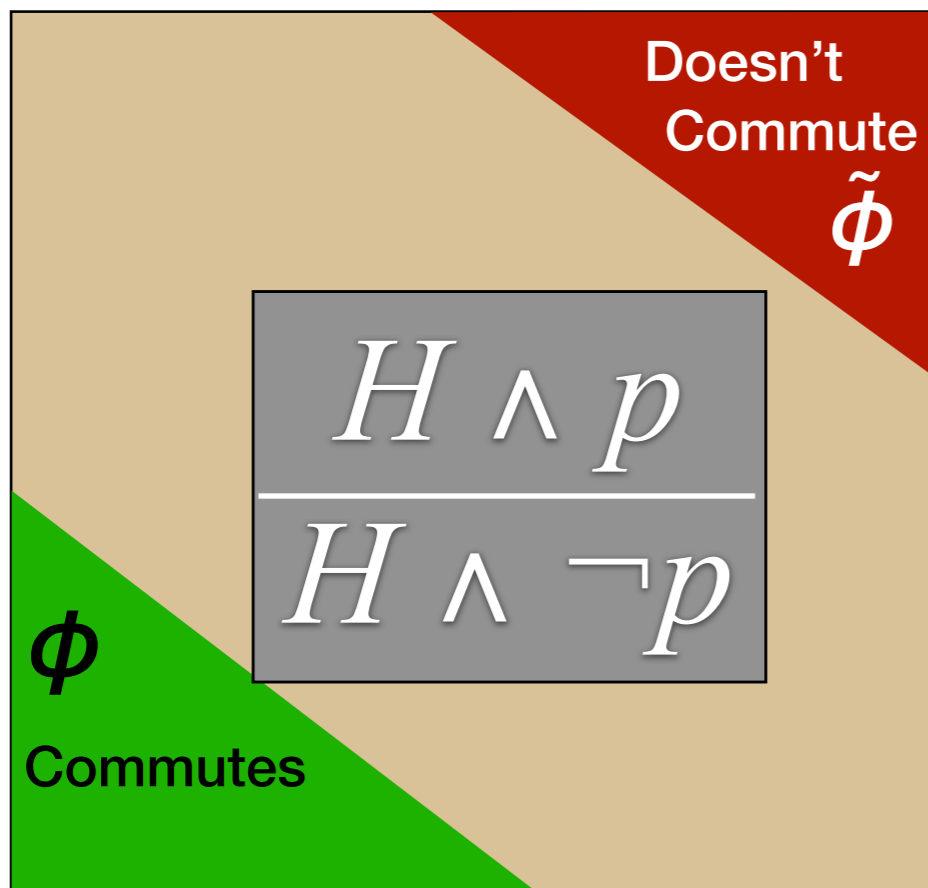
✓ If valid($H \Rightarrow m \not\bowtie n$): add H to $\tilde{\phi}$



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)


✓ If valid($H \Rightarrow m \bowtie n$): **add H to ϕ**

✓ If valid($H \Rightarrow m \not\bowtie n$): **add H to $\tilde{\phi}$**

✓ If neither:

$P = \text{CHOOSE}(\dots)$


$\text{Refine}(H \wedge P); \text{Refine}(H \wedge \neg P)$



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition via abstraction-refinement



Refine(H)

✓ If valid($H \Rightarrow m \bowtie n$): **add H to ϕ**


✓ If valid($H \Rightarrow m \not\bowtie n$): **add H to $\tilde{\phi}$**

✓ If neither:

$P = \text{CHOOSE}(\dots)$

$\text{Refine}(H \wedge P); \text{Refine}(H \wedge \neg P)$

Use counterex's and "poke" heuristic. Generate preds, etc.



Logical ADT Specification


$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

1. Verifying commute condition $\varphi_m^n(\sigma, \bar{x}, \bar{y})$
2. **Synthesize** commute condition

```

1 REFINEnm(H, P) {
2   if valid(H ⇒ m ⋈ n) then
3     φ := φ ∨ H;
4   else if valid(H ⇒ m ⋉ n) then
5     φ̃ := φ̃ ∨ H;
6   else
7     let χc, χnc = counterex. to ⋈ and ⋉ (resp.) in
8     let p = CHOOSE(H, P, χc, χnc) in
9       REFINEnm(H ∧ p, P \ {p});
10      REFINEnm(H ∧ ¬p, P \ {p});
11 }
12 main {
13   φ := false; φ̃ := false;
14   try { REFINEnm(true, P); }
15   catch (InterruptedException e) { skip; }
16   return(φ, φ̃);
17 }

```




Set Abstract Data Type


S

`contains(x) / bool`, which performs a side-effect-free check whether the element x is in S ; and

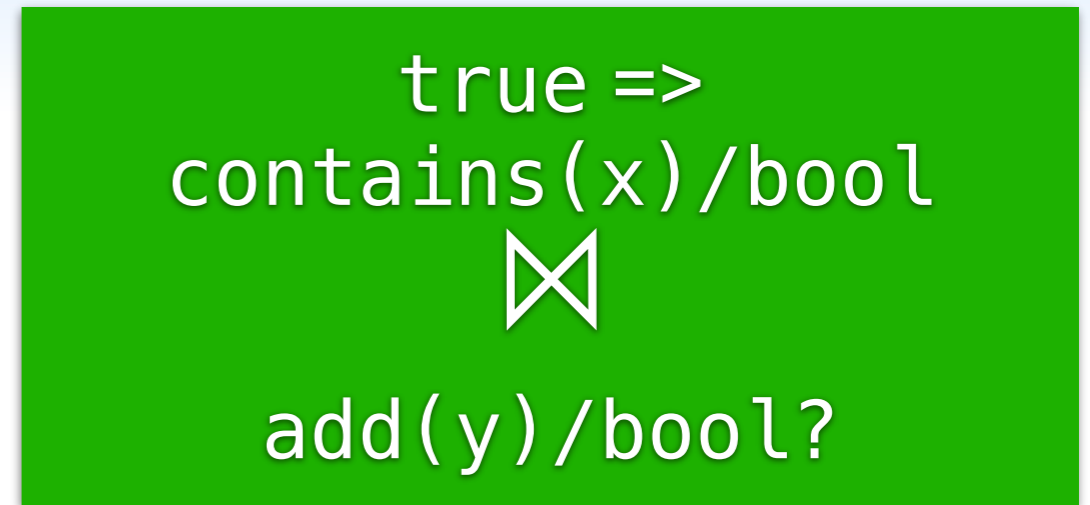


`add(y) / bool`, which adds y to S if it is not already in there and returns `true`, or otherwise returns `false`.



Don't Know



H = true




commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}



φ_m^n



Counterexample (X)

true =>
contains(x)/bool
X
add(y)/bool?


commute ? {
{s1} {s2}

Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n



Counterexample (⊗)

true =>
contains(x)/bool
⊗
add(y)/bool?

true =>
contains(x)/bool
⊗
add(y)/bool?

commute ? {
{s1} {s2}}


Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

Counterexample (~~⊠~~)



Counterexample (~~⊠~~)

true =>
contains(x)/bool
~~⊠~~
add(y)/bool?

true =>
contains(x)/bool
~~⊠~~
add(y)/bool?

commute ? {
{s1} {s2}}

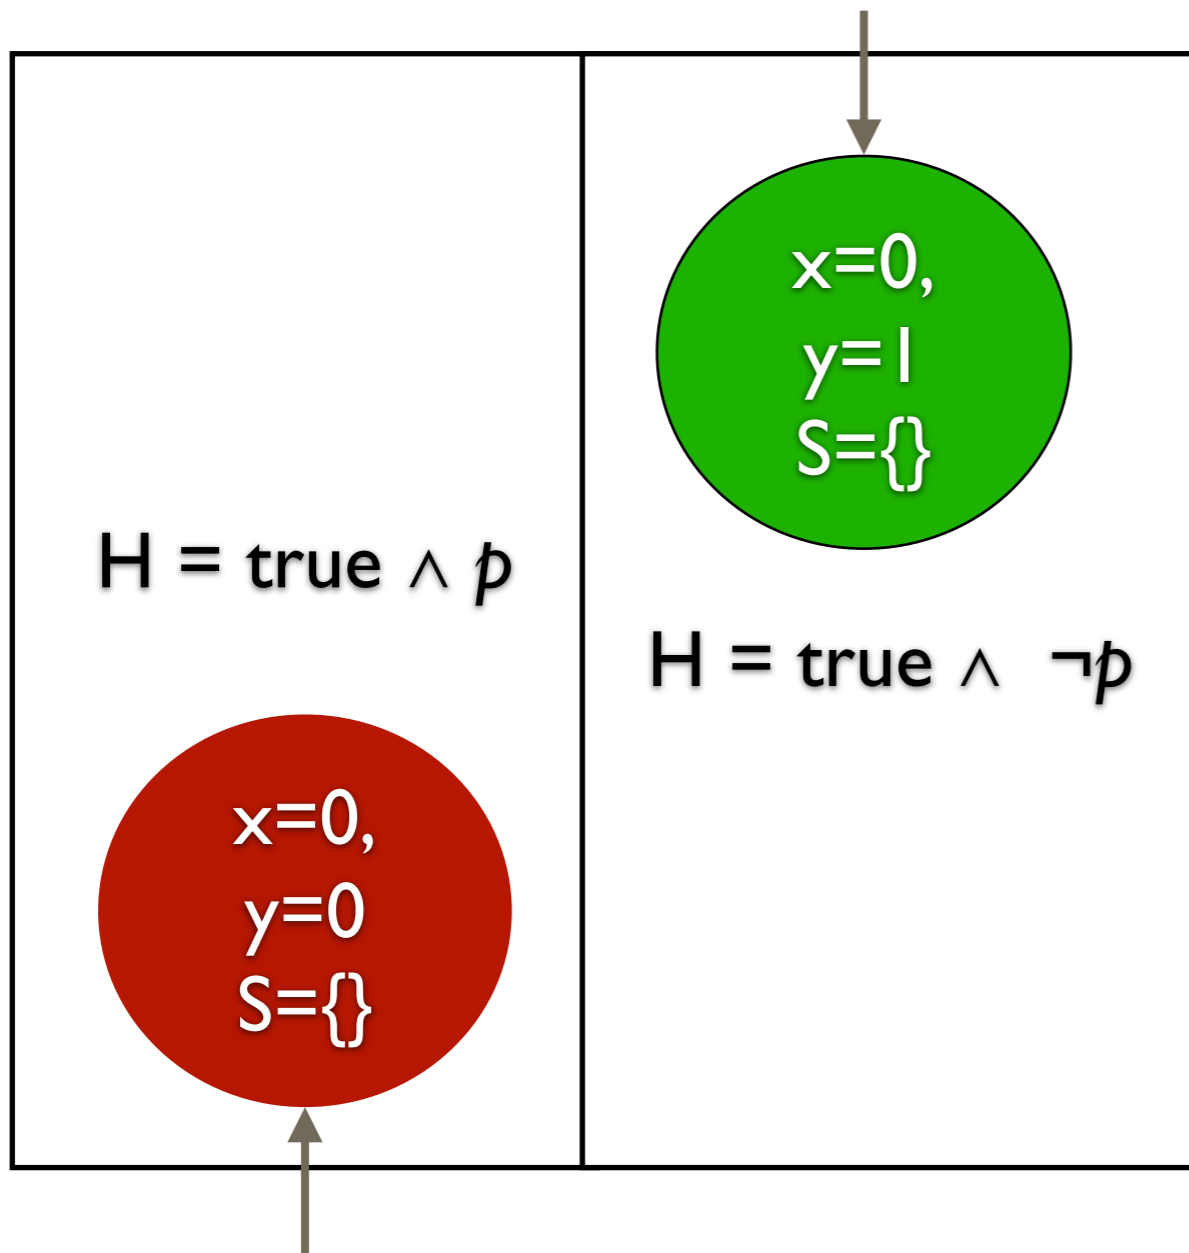
Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n



Counterexample (~~⊠~~)



Counterexample (~~⊠~~)

true =>
contains(x)/bool
~~⊠~~
add(y)/bool?

true =>
contains(x)/bool
~~⊠~~
add(y)/bool?



$$p: (x = y)$$

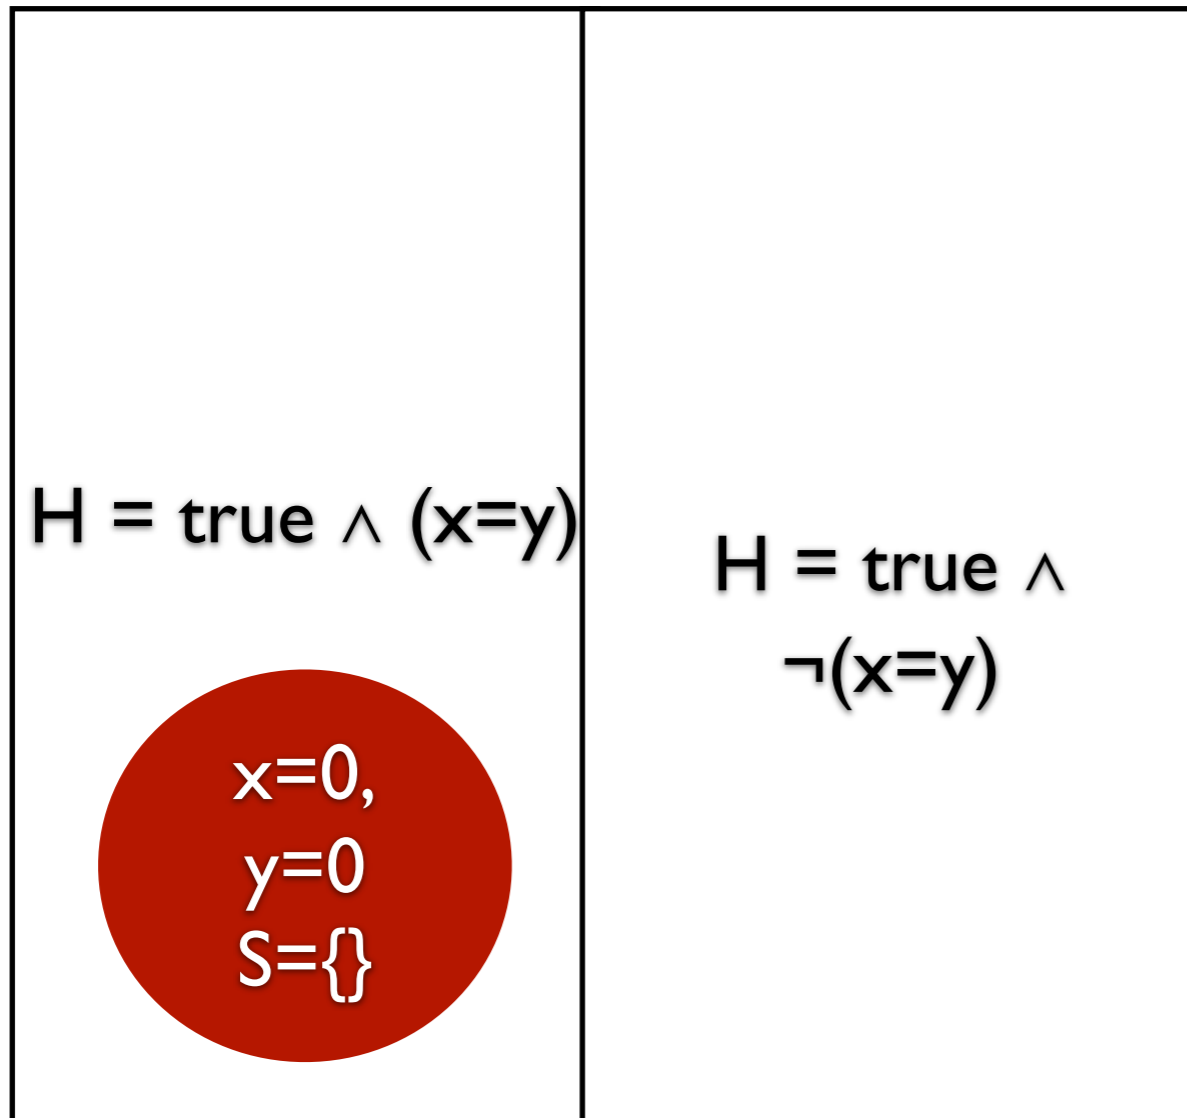
commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n



$p: (x = y)$

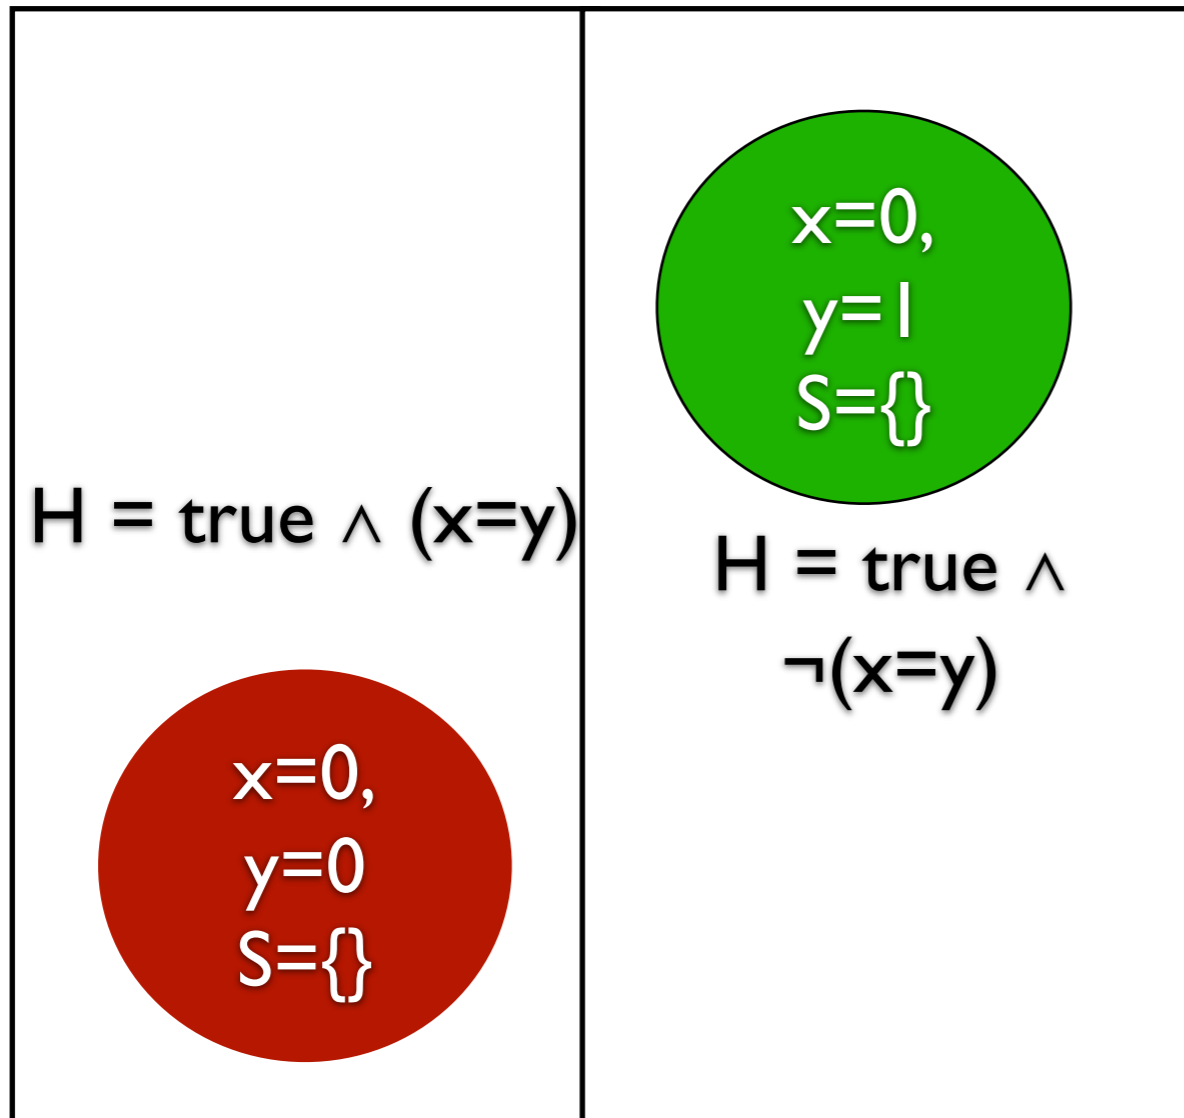
commute ? {
{s1} {s2}}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}



φ_m^n



$p: (x = y)$

commute ? {
{s1} {s2}


Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

$H: \neg(x = y)$



$H \Rightarrow$
 contains(x)/bool
~~×~~
 add(y)/bool?

Yes!

commute ? {
{s1} {s2}

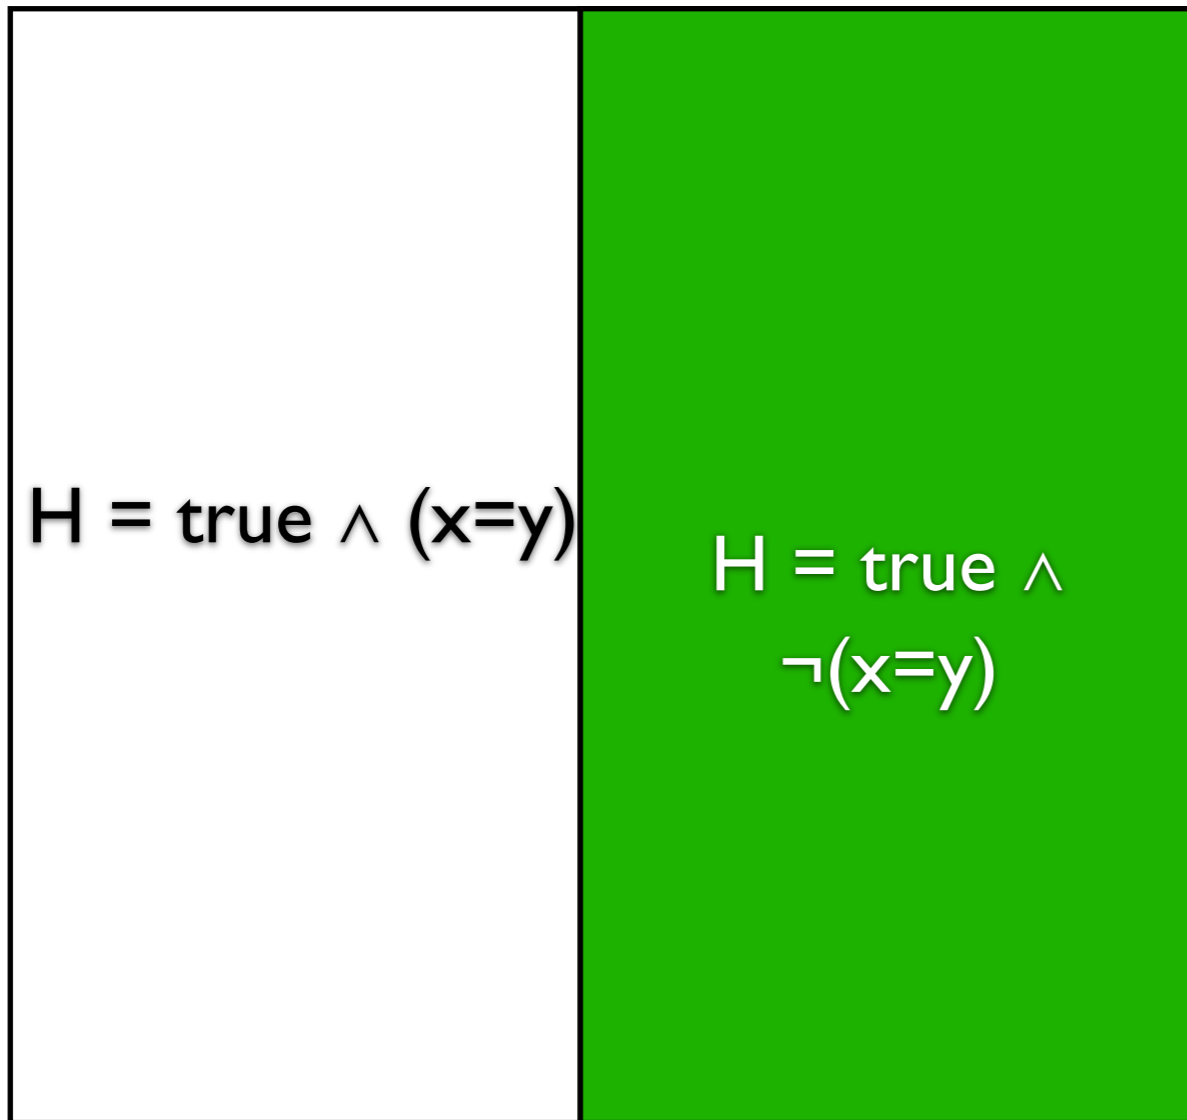
Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}



φ_m^n

H: (x = y)



H =>
contains(x)/bool
~~⊗~~
add(y)/bool?

H =>
contains(x)/bool
~~⊗~~
add(y)/bool?

commute ? {
{s1} {s2}

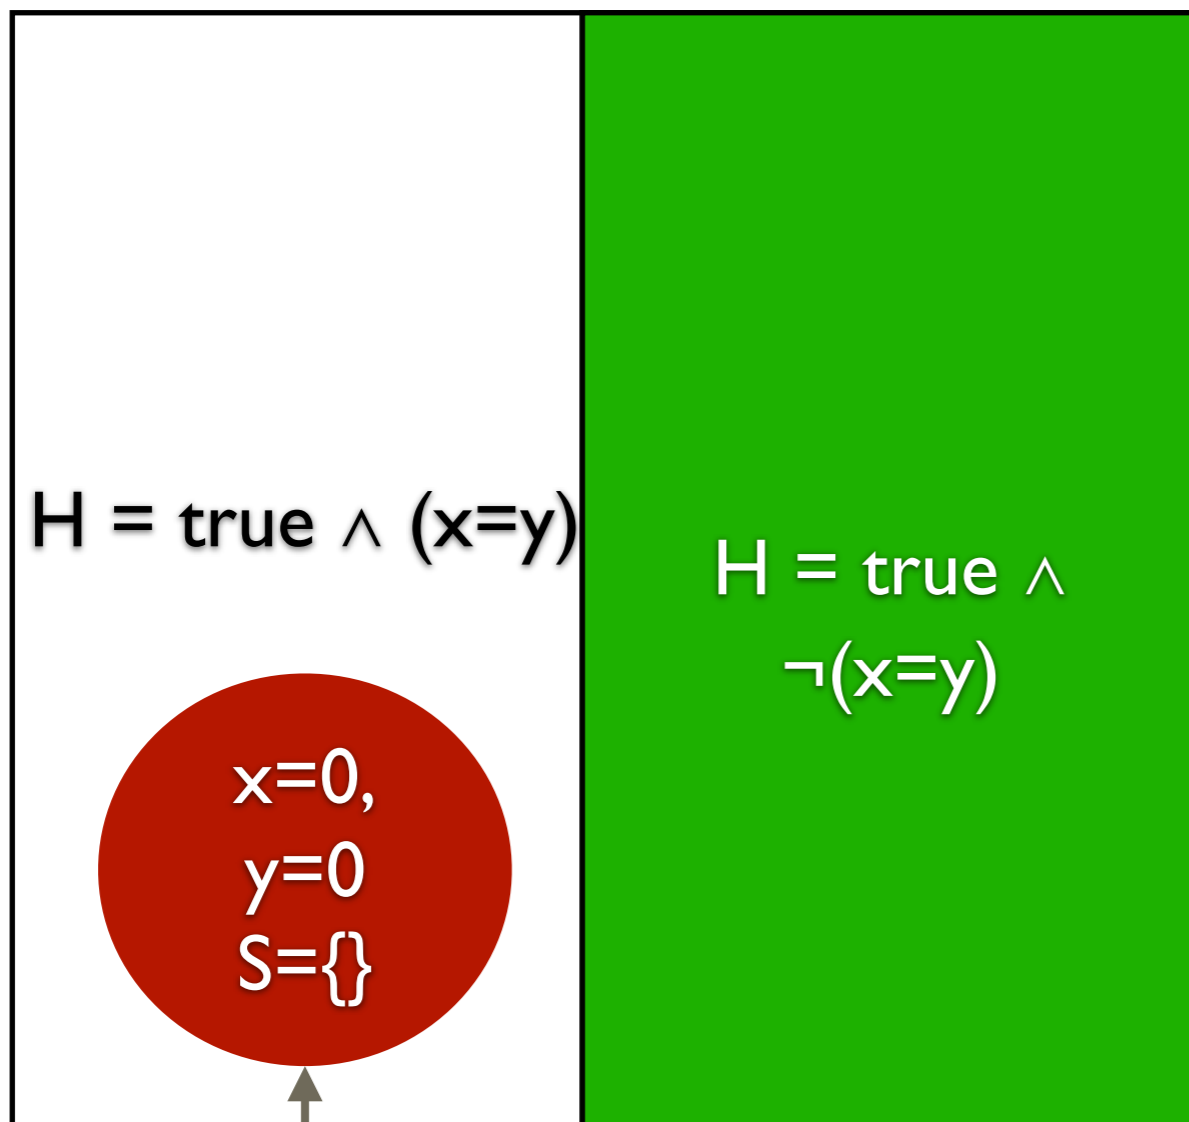
Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

H: (x = y)



Counterexample (\boxtimes)

H =>
contains(x)/bool
 \boxtimes
add(y)/bool?

H =>
contains(x)/bool
 \boxtimes
add(y)/bool?

commute ? {
{s1} {s2}}


Translate(**co**
mmute, s₁, s₂)

P_{ins}, Q_{ins}
P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

Counterexample (~~⊗~~)



Counterexample (~~⊗~~)

H: (x = y)

H =>
contains(x)/bool
~~⊗~~
add(y)/bool?

H =>
contains(x)/bool
~~⊗~~
add(y)/bool?

commute ? {
{s1} {s2}}

Translate(**co**
mmute, s₁, s₂)



P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

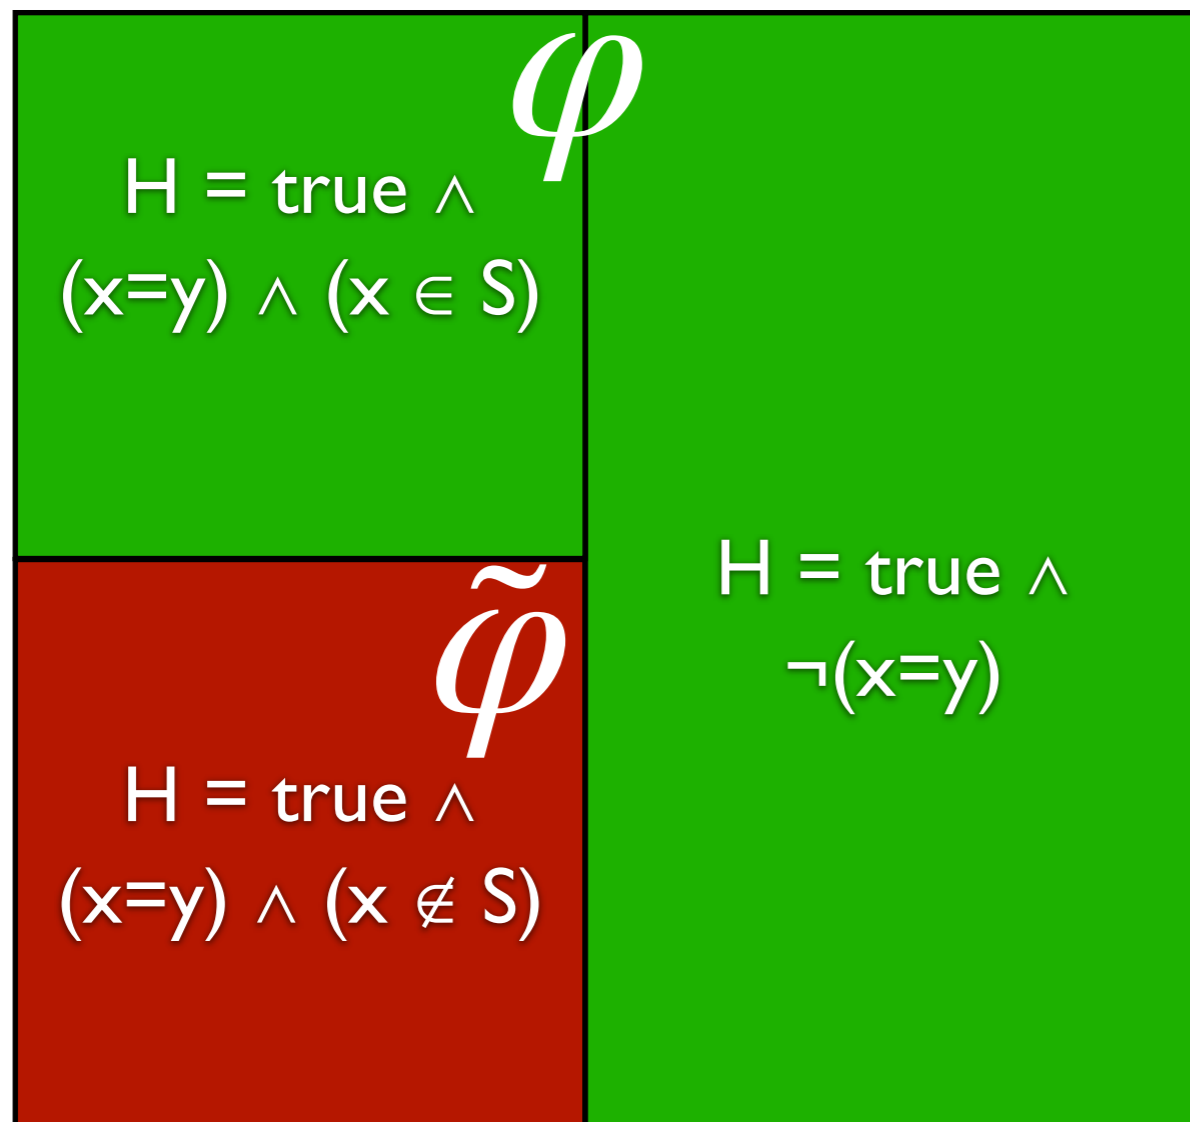



φ_m^n

$H = \text{true} \wedge (x=y) \wedge (x \in S)$	$H = \text{true} \wedge \neg(x=y)$
$H = \text{true} \wedge (x=y) \wedge (x \notin S)$	

$p': (x \in S)$





$$\varphi \equiv x \neq y \vee (x = y \wedge x \in S)$$

commute ? {
{s1} {s2}

Translate(commute, s₁, s₂)

P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n

	$m(\bar{x})$		$n(\bar{y})$	Simple Qs (time)	Poke Qs (time)	φ_n^m (Poke)
Counter	decrement	⊗	decrement	3 (0.1)	3 (0.1)	true
	increment	▷	decrement	10 (0.3)	34 (0.9)	$\neg(0 = c)$
	decrement	▷	increment	3 (0.1)	3 (0.1)	true
	decrement	⊗	reset	2 (0.1)	2 (0.1)	false
	decrement	⊗	zero	6 (0.1)	26 (0.6)	$\neg(1 = c)$
	increment	⊗	increment	3 (0.1)	3 (0.1)	true
	increment	⊗	reset	2 (0.0)	2 (0.1)	false
	increment	⊗	zero	10 (0.3)	34 (0.8)	$\neg(0 = c)$
	reset	⊗	reset	3 (0.1)	3 (0.1)	true
	reset	⊗	zero	9 (0.2)	30 (0.6)	$0 = c$
zero	⊗	zero	3 (0.1)	3 (0.1)	true	
Acum.	increase	⊗	increase	3 (0.1)	3 (0.1)	true
	increase	⊗	read	13 (0.3)	28 (0.6)	$c + x_1 = c$
	read	⊗	read	3 (0.0)	3 (0.0)	true
Set	add	⊗	add	10 (0.4)	140 (4.4)	$(y_1 = x_1 \wedge y_1 \in S) \vee \neg(y_1 = x_1)$
	add	⊗	contains	10 (0.4)	122 (3.6)	$x_1 \in S \vee (\neg(x_1 \in S) \wedge \neg(y_1 = x_1))$
	add	⊗	getsize	6 (0.2)	31 (0.9)	$x_1 \in S$
	add	⊗	remove	6 (0.2)	66 (2.2)	$\neg(y_1 = x_1)$
	contains	⊗	contains	3 (0.1)	3 (0.1)	true
	contains	⊗	getsize	3 (0.1)	3 (0.1)	true
	contains	⊗	remove	17 (0.5)	160 (4.8)	$S \setminus \{x_1\} = \{y_1\} \vee (\dots \wedge y_1 \in \{x_1\}) \vee$...
	getsize	⊗	getsize	3 (0.1)	3 (0.1)	true
getsize	⊗	remove	13 (0.3)	37 (1.0)	$\neg(y_1 \in S)$	

commute ? {
{s1} {s2}}

Translate(commute, s1, s2)


P_{ins}, Q_{ins}
 P_{rm}, Q_{rm}

Servois
Synthesis
TACAS'18

φ_m^n



$m(\bar{x})$		$n(\bar{y})$	Simple Qs (time)	Poke Qs (time)	φ_n^m (Poke)
decrement	⊗	decrement	3 (0.1)	3 (0.1)	true
increment	▷	decrement	10 (0.3)	34 (0.9)	$\neg(0 = c)$
get	⊗	get	3 (0.1)	3 (0.1)	true
get	⊗	haskey	3 (0.1)	3 (0.1)	true
put	▷	get	13 (0.4)	74 (2.3)	$(H[x_1 \leftarrow x_2] = H \wedge y_1 \in keys)$ $\vee (\neg(H[x_1 \leftarrow x_2] = H) \wedge \neg(y_1 = x_1))$
get	▷	put	10 (0.3)	48 (1.5)	$[H[y_1] = y_2] \vee [\neg(H[y_1] = y_2) \wedge \neg(y_1 = x_1)]$
remove	▷	get	3 (0.1)	3 (0.1)	true
get	▷	remove	13 (0.4)	40 (1.2)	$\neg(y_1 = x_1)$
get	⊗	size	3 (0.1)	3 (0.1)	true
haskey	⊗	haskey	3 (0.1)	3 (0.1)	true
haskey	⊗	put	10 (0.3)	52 (1.6)	$[y_1 \in keys] \vee [\neg(y_1 \in keys) \wedge \neg(y_1 = x_1)]$
haskey	⊗	remove	17 (0.5)	44 (1.3)	$[x_1 \in keys \wedge \neg(y_1 = x_1)] \vee [\neg(x_1 \in keys)]$
haskey	⊗	size	3 (0.1)	3 (0.1)	true
put	⊗	put	24 (0.9)	357 (13.5)	$\dots \vee (\neg(H[y_1] = y_2) \wedge \neg(y_1 = x_1))$
put	⊗	remove	6 (0.3)	33 (1.2)	$\neg(y_1 = x_1)$
put	⊗	size	6 (0.2)	23 (0.8)	$x_1 \in keys$
remove	⊗	remove	21 (0.8)	192 (6.9)	$[keys \setminus \{x_1\} = \{y_1\}] \vee [\dots]$
contains	⊗	remove	11 (0.5)	100 (4.0)	$S \setminus \{x_1\} = \{y_1\} \vee (\dots \wedge y_1 \in \{x_1\})$
getsize	⊗	getsize	3 (0.1)	3 (0.1)	true
getsize	⊗	remove	13 (0.3)	37 (1.0)	$\neg(y_1 \in S)$

HashTable



Applications of Commutativity Synthesis

- **Smart Contracts.** Ensure determinism.
- **Concurrent verification.** Partial Order reduction, transactional memory, etc.
- **Testing for interactions between code blocks.**
- **CRDTs.** Distributed computing.
- **Refactoring (and other relational reasoning).**
- **Code synthesis.** Eg, synthesized conditions become specification for synchronization synthesis.
- **Commutate blocks in Veracity!**



SERVOIS

Implemented in Python with CVC4.


Available on GitHub.



Kshitij Bansal
PhD student at NYU
Now at Google



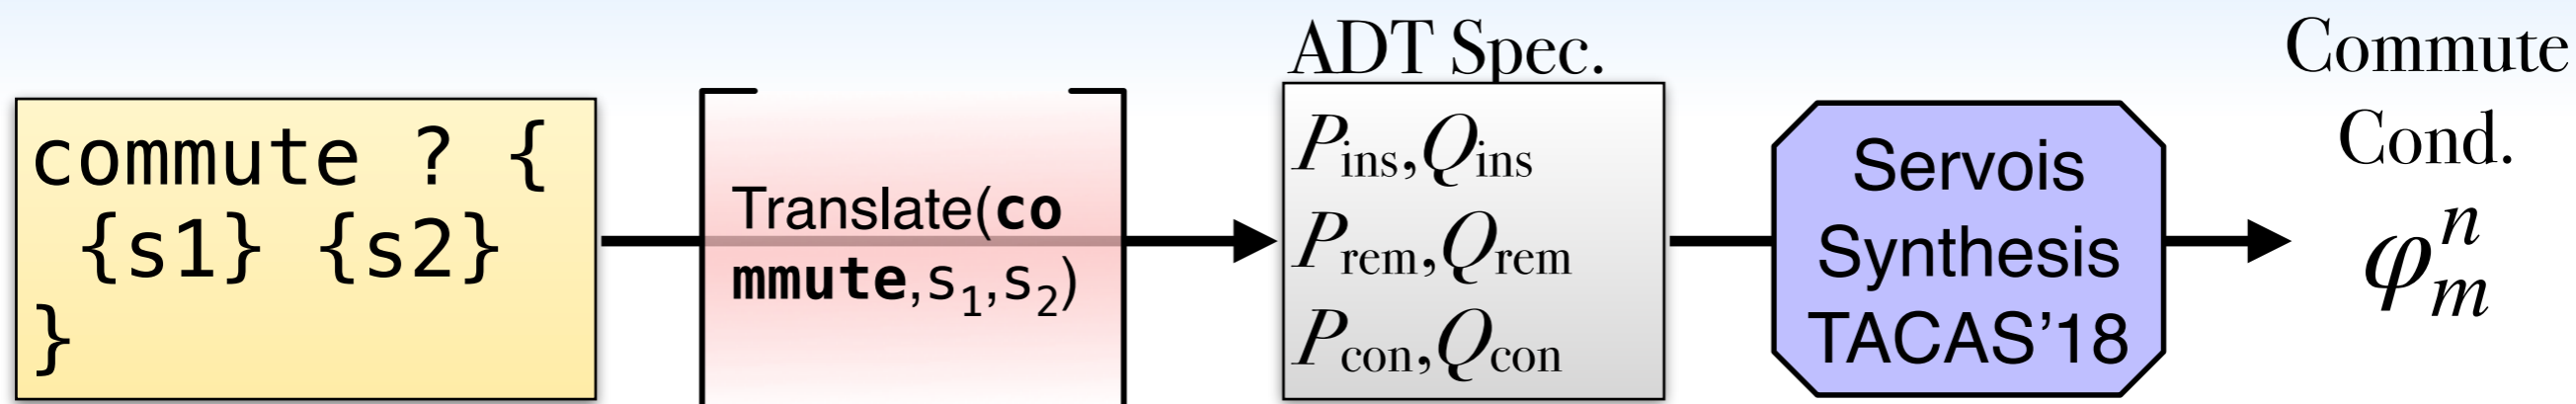

Omer Tripp
PhD student at TAU
Now at Amazon



SERVOIS 2.0

Coming very soon!

- **More solvers!** CVC4, CVC5, Z3, ...
- **More theories!** e.g bitvectors.
- **Faster!** Reimplemented in OCaml from scratch.
- **Better predicate generation.**
- **Better predicate selection.**
- **Command-line or Library API.**



Veracity: Declarative Multicore Programming with Commutativity

ADAM CHEN, Stevens Institute of Technology, USA

PARISA FATHOLOLUMI, Stevens Institute of Technology, USA

ERIC KOSKINEN, Stevens Institute of Technology, USA

JARED PINCUS, Stevens Institute of Technology, USA


There is an ongoing effort to provide programming abstractions that ease the burden of exploiting multicore hardware. Many programming abstractions (*e.g.*, concurrent objects, transactional memory, etc.) simplify matters, but still involve intricate engineering. We argue that some difficulty of multicore programming can be meliorated through a declarative programming style in which programmers directly express the independence of fragments of sequential programs.

In our proposed paradigm, programmers write programs in a familiar, sequential manner, with the added ability to explicitly express the conditions under which code fragments sequentially commute. Putting such commutativity conditions into source code offers a new entry point for a compiler to exploit the known connection between commutativity and parallelism. We give a semantics for the programmer's sequential perspective and, under a correctness condition, find that a compiler-transformed parallel execution is equivalent to the sequential semantics. Serializability/linearizability are not the right fit for this condition, so we introduce scoped serializability and show how it can be enforced with lock synthesis techniques.

We next describe a technique for automatically verifying and synthesizing commute conditions via a new reduction from our commute blocks to logical specifications, upon which symbolic commutativity reasoning can be performed. We implemented our work in a new language called Veracity, implemented in Multicore OCaml. We show that commutativity conditions can be automatically generated across a variety of new benchmark programs, confirm the expectation that concurrency speedups can be seen as the computation increases, and apply our work to a small in-memory filesystem and an adaptation of a crowdfund blockchain smart contract.

1 INTRODUCTION

Writing concurrent programs is difficult. Researchers and practitioners, seeking to make life easier,



Logical ADT Specification


$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

m_1 :

```
x = calc1(a);
c = c + (x*x);
```

m_2 :

```
if (c > 0 && y < 0) {
  c = c - 1;
  z = calc2(y);
} else {
  z = calc3(y);
}
```

Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$


m_1 :

```
x = calc1(a);
c = c + (x*x);
```

But m_1 is code, not a logical spec.

m_2 :

```
if (c > 0 && y < 0) {
  c = c - 1;
  z = calc2(y);
} else {
  z = calc3(y);
}
```



Logical ADT Specification

$$O = \left\{ \begin{array}{ll} \text{state} & : (Var, Type)list; \\ \text{eq} & \subseteq \text{state} \times \text{state}; \end{array} \quad \begin{array}{ll} \text{methods} & : \text{Meth list}; \\ \text{spec} & : \text{Meth} \rightarrow (P, Q) \end{array} \right\}$$

m_1 :

```
x = calc1(a);
c = c + (x*x);
```


But m_1 is code, not a logical spec.

m_2 :

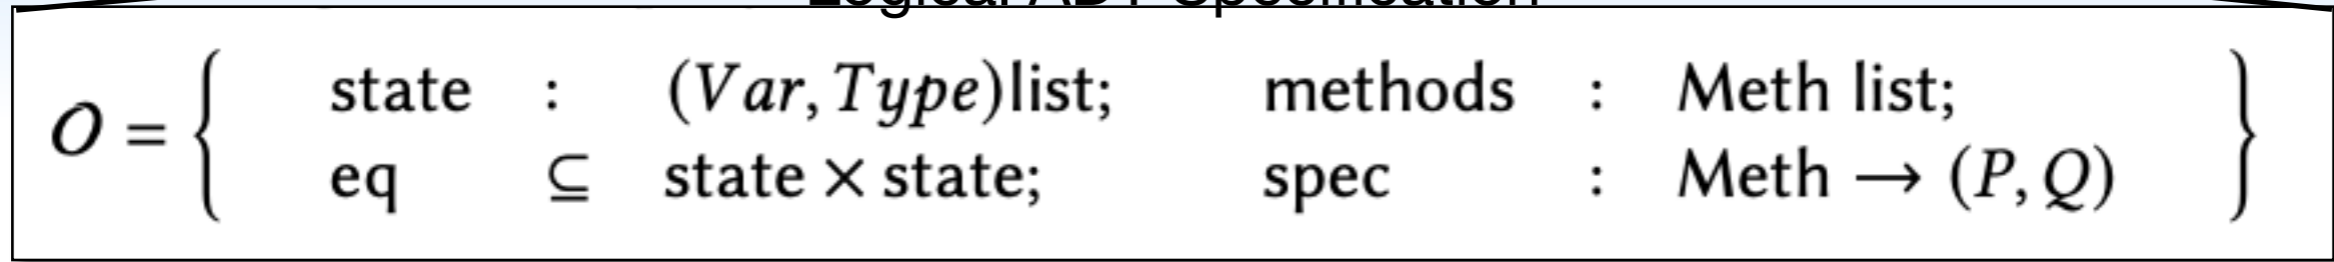
```
if (c > 0 && y < 0) {
  c = c - 1;
  z = calc2(y);
} else {
  z = calc3(y);
}
```

Translate to a logical post-condition.

```
(or (and
  (let ((x_1 a)
        (let ((c_1 (+ c (* x_1 x_1))))
          (and (= c_new c_1) (= x_new x_1))))
    (= a_new a) (= z_new z) (= y_new y)
    (= size_new size))))
```



Logical ADT Specification




Translation

- **Nested commute statements?** Treat them as sequential composition!

$$Tr(\mathbf{commute} \ c \ s_1 \ s_2) = Tr(s_1; s_2)$$

- **Built-in ADTs?**

$$Tr(\mathbf{tbl}[e_1] = e_2) = inlineSpec(HT, \mathbf{tbl}, \dots)$$




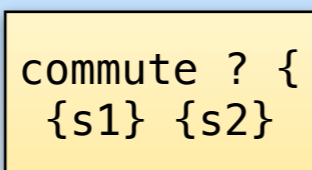




```


int main(int argc, string[] argv) {
    hashtable[int,int] tbl = new hashtable[int,int];
    int n = int_of_string(argv[1]);
    int x = int_of_string(argv[2]);
    int y = int_of_string(argv[3]);
    int z = int_of_string(argv[4]);
    tbl[x] = 42;
    tbl[z] = 42;
    commute _ {
        {
            calc1(n);
            if(ht_mem(tbl, x))
                y = tbl[x];
        }
        { if(ht_mem(tbl, z)) {
            y = tbl[z];
        }
        calc2(n);}
    }
    return 0;
}
  
```

```

ejk@arran:veracity/src$ ./vcy.exe infer ../benchmarks/ht-cond-mem-get.vcy
Inferred condition at ../benchmarks/ht-cond-mem-get.vcy:
[11.2-22.3]: tbl[x] == tbl[z] && !(x == z) || x == z
ejk@arran:veracity/src$
  
```

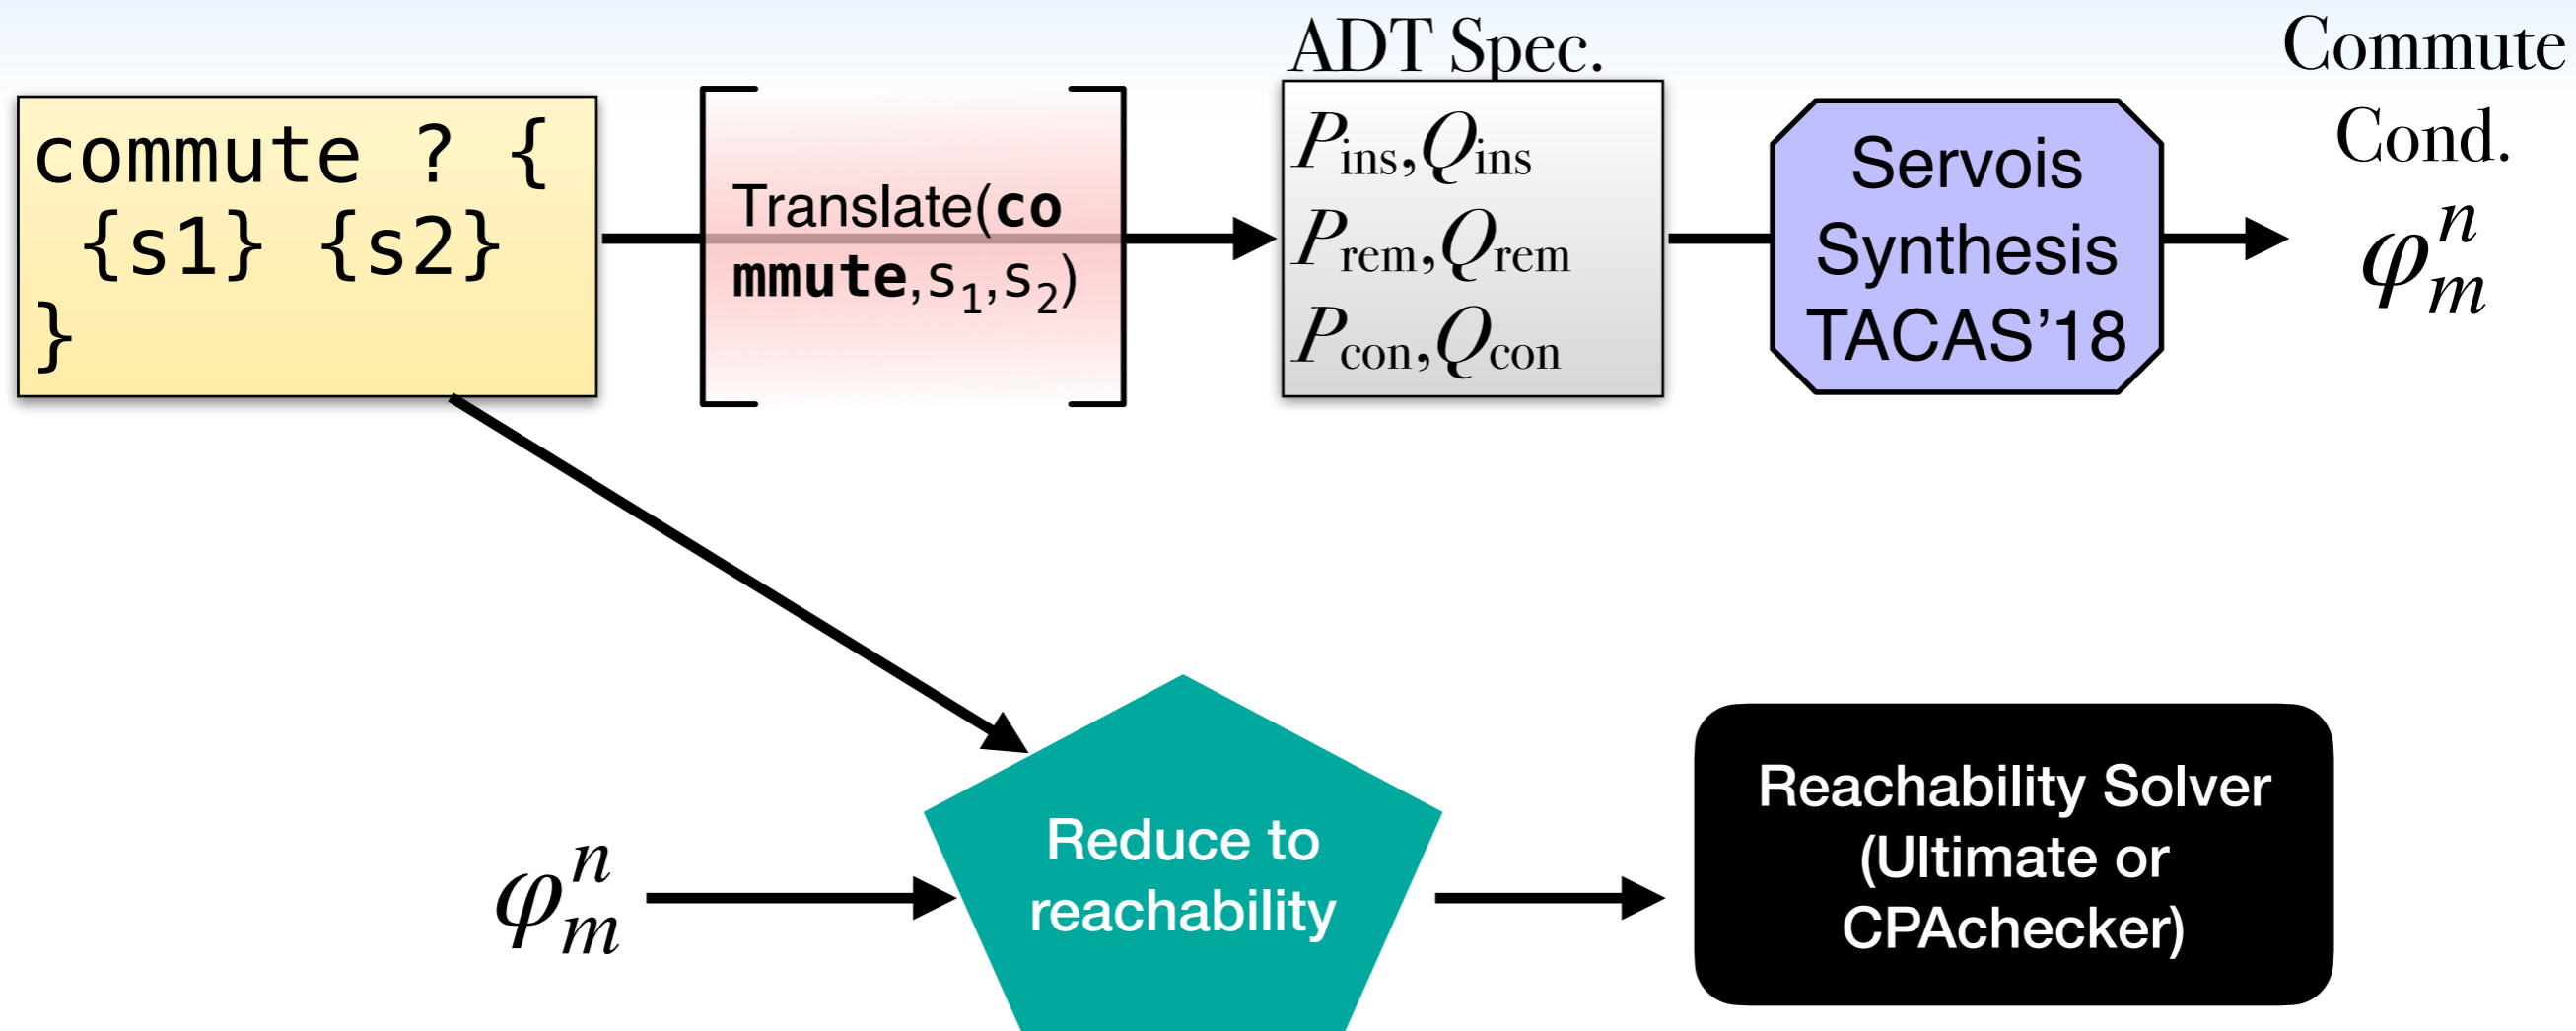


 φ_m^n
Group 1: Automatically Inferred Commute Conditions. All benchmarks, except those below in group (3).

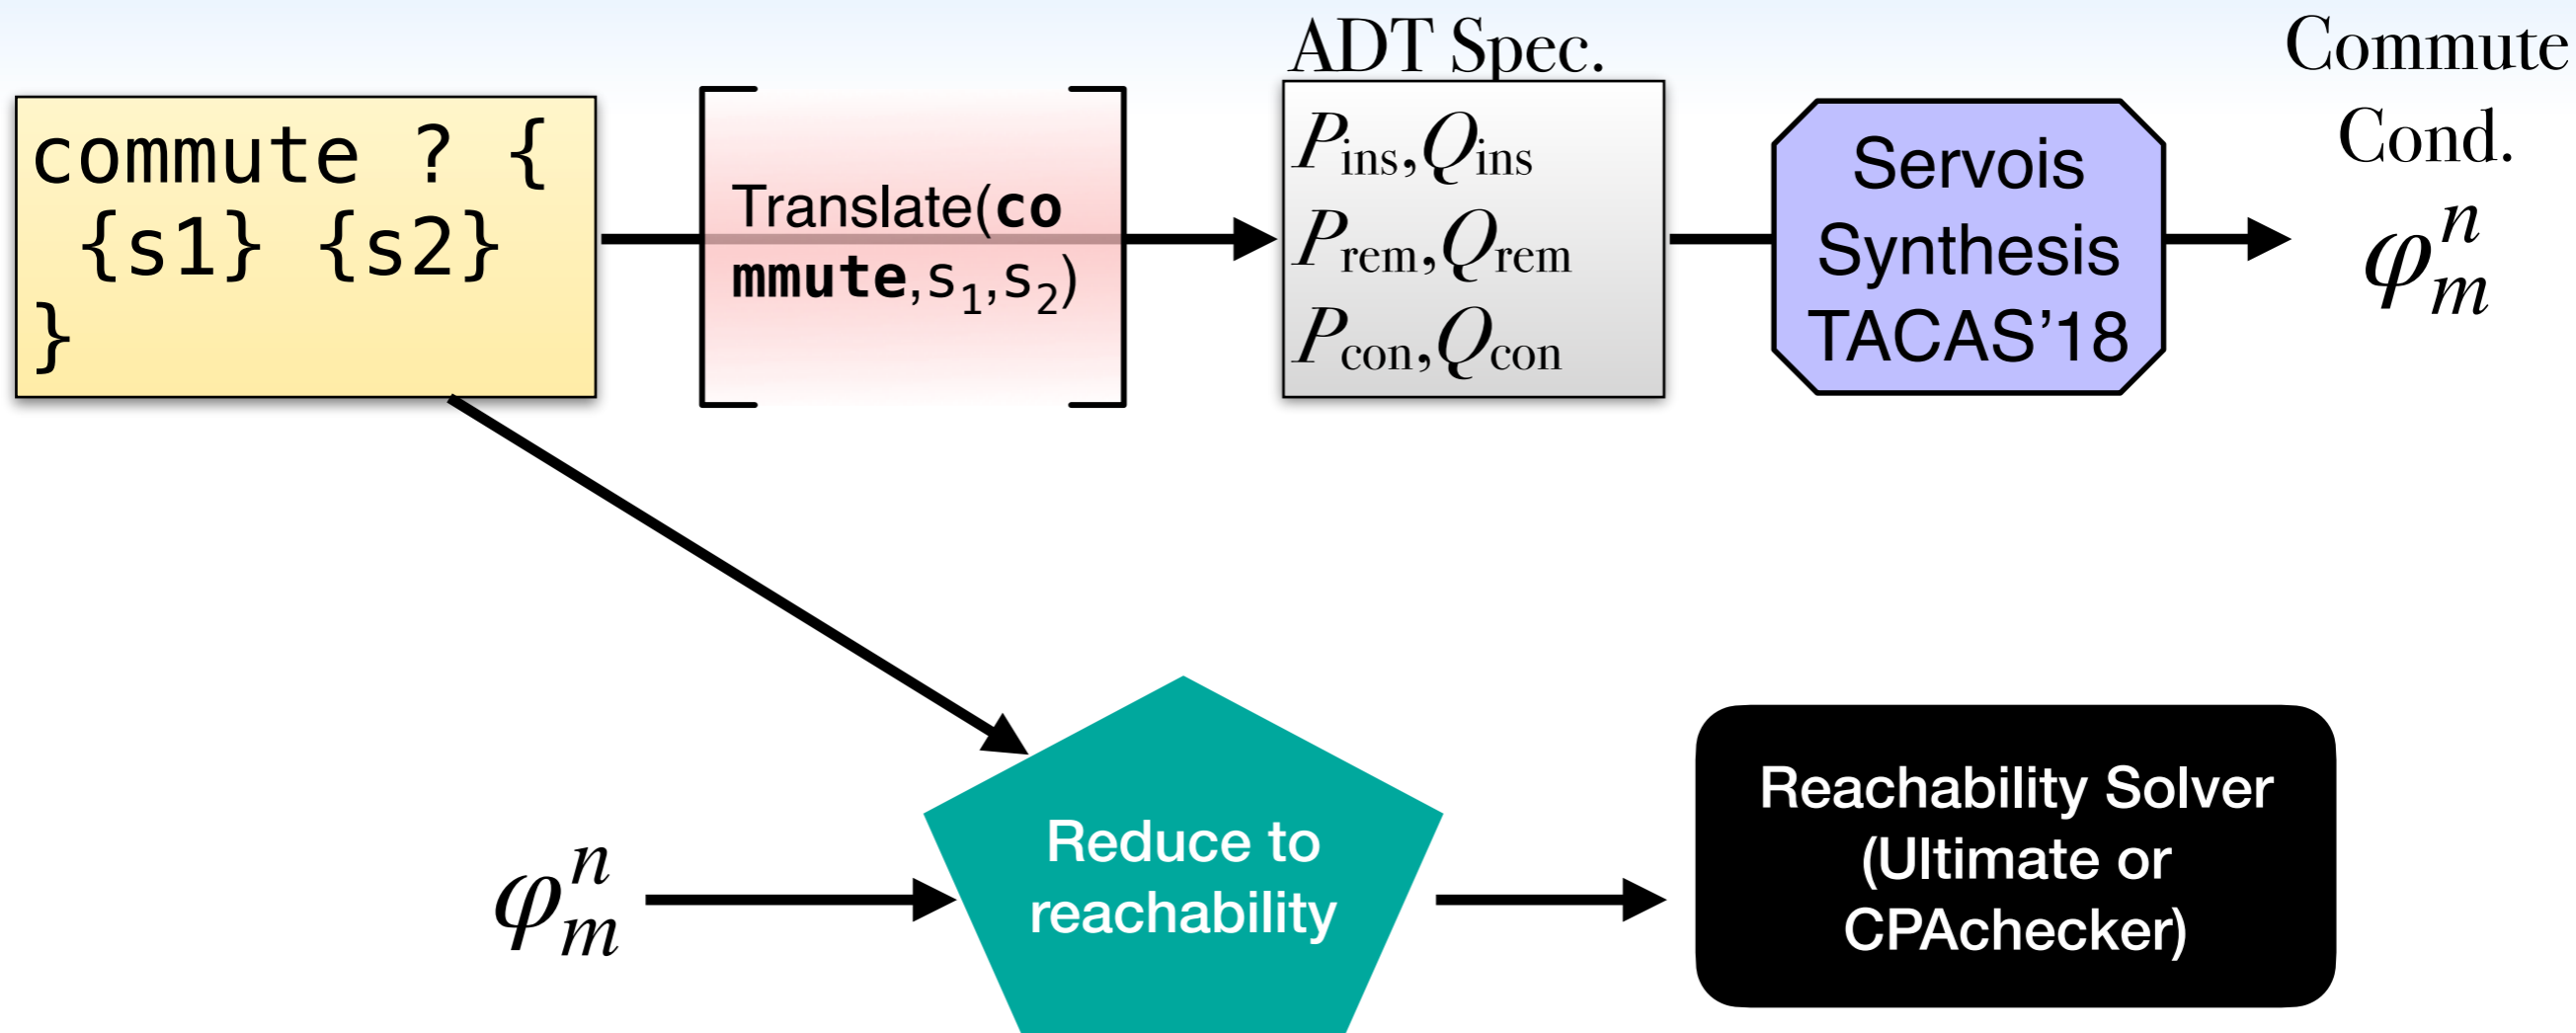

Program	Time (s)	Inferred Conditions
array-disjoint	0.63	$i \neq j \ \&\& \ x \neq y \ \ x == y$
array1	0.75	$1 \neq r[0] \ \&\& \ r[0] + 1 \neq y \ \&\& \ r[0] \leq 1 \ \ r[0] + 1 == y \ \&\& \ r[0] \leq 1$
array2	1.11	$0 > a[0] \ \&\& \ 1 \neq x \ \ 1 == x$
array3	1.13	$d \neq e \ \&\& \ a \neq b \ \ a == b$
calc	 1.13	$1 == y \ \&\& \ 0 \neq y \ \&\& \ 1 > c \ \&\& \ 1 \neq c \ \ \dots \ \ 1 == c$
conditional	0.18	$x > 0$
counter	0.20	$0 \neq c$
dict	3.82	$i \neq r \ \&\& \ c + x \neq y \ \ c + x == y$
dot-product	0.24	true
even-odd	1.18	$x \% 2 == x + y \ \&\& \ 0 \neq y \ \ 0 == y$
ht-add-put	2.24	$tbl[z] == u + 1 \ \&\& \ u + 1 \neq z$
ht-cond-mem-get	1.54	$tbl[x] == tbl[z] \ \&\& \ x \neq z \ \ x == z$
ht-cond-size-get	0.83	$ht_size(tbl) \leq 0 \ \&\& \ 0 \neq z \ \ 0 == z$
ht-simple	30.64	$x + a \neq z \ \&\& \ 3 == tbl[z] \ \&\& \ y \neq z$
linear-bool	3.62	$0 \leq y \ \&\& \ 3 == x \ \&\& \ 2 \neq x \ \&\& \ 1 \neq x \ \&\& \ x > 0 \ \&\& \ 0 \neq x \ \ 0 > y + 3 * x \ \&\& \ 2 == x \ \&\& \ 1 \neq x \ \&\& \ x > 0 \ \&\& \ 0 \neq x$
linear-cond	2.65	$2 \leq y \ \&\& \ 2 \neq y \ \&\& \ 1 \neq y \ \ \dots \ \ 1 == y$
linear	0.25	true
loop-amt	 0.25	$0 == i \ \&\& \ amt == i_pre \ \&\& \ ctr - 1 > i_pre \ \&\& \ i_pre \leq amt \ \&\& \ 0 \neq i_pre \ \&\& \ i_pre \leq ctr \ \&\& \ amt \neq amt_pre \ \&\& \ ctr - 1 > amt_pre \ \&\& \ amt_pre \leq amt \ \&\& \ 0 \neq amt_pre \ \&\& \ amt_pre \leq ctr \ \&\& \ ctr - 1 \neq 1 \ \&\& \ 1 \neq ctr \ \&\& \ 1 \neq amt \ \&\& \ 1 == ctr + amt \ \ \dots \ \ amt == i \ \&\& \ 1 == ctr \ \&\& \ 1 \neq amt \ \&\& \ 1 == ctr + amt$
loop-disjoint	0.02	true
loop-inter	4.63	$0 == x \ \&\& \ 0 \neq y \ \ 0 == y$
loop-simple	0.06	true
matrix	0.71	$0 == y$
nested-counter	6.25	$0 \neq c \ \&\& \ c \neq t \ \ c == t; \ c \neq x \ \&\& \ c \leq x \ \&\& \ 1 \neq x \ \&\& \ t == x \ \ \dots \ $



Group 2: Automatically Verified Commute Conditions. Benchmarks for which inference output was suboptimal.

Program	Time (s)	Verified?	Complete?	Provided Condition
array1	0.02	✓	✓	<code>r[0] <= 0 r[0] == 1 && y == 2</code>
calc	0.07	✓	?	<code>c > 0</code>
counter	0.02	✗	—	<code>true</code>
even-odd	0.04	✓	✓	<code>y % 2 == 0</code>
linear-bool	0.02	✓	✓	<code>y < 0 - 3 * x && x == 2 y >= 0 && x == 3</code>
linear-cond	0.02	✓	✓	<code>y > 0 0 == y && x + 2 == z</code>
loop-amt	0.04	✓	✗	<code>i % 2 == 0 && (ctr > 0 && amt > 0 ctr <= 0 && amt < -ctr)</code>
nested-counter	0.05	✓	✓	First commute block: <code>0 != c && c != t c == t</code>
(cont.)		✓	✗	Second commute block: <code>x == t && (x > c x == c && x > 1)</code>
simple	0.04	✓	✗	<code>c > a</code>





Decomposing Data Structure Commutativity Proofs with *mm*-Differencing

Eric Koskinen¹(✉) and Kshitij Bansal²

¹ Stevens Institute of Technology, Hoboken, NJ, USA
eric.koskinen@stevens.edu

Future

Beyond the interpreter.

Combine with promises/futures?

Beyond N-way commute blocks?

Combine with invariant generation:

```
{ y < 0 } /* Example invariant */  
commute - {  
  { y = y + 3*x; } /* if x is negative, this will reduce y */  
  { if (y<0) { x=2; } else { x=3; } } /* sensitive to whether y went below 0 */  
}
```

Future

Beyond the interpreter.

Combine with promises/futures?

Beyond N-way commute blocks?

Combine with invariant generation:

```
{ y < 0 } /* Example invariant */  
commute - {  
  { y = y + 3*x; } /* if x is negative, this will reduce y */  
  { if (y<0) { x=2; } else { x=3; } } /* sensitive to whether y went below 0 */  
}
```

$$0 > y + 3*x \ \&\& \ 2 == x$$

Interested in Commutativity?

Workshop at PLDI next month!



PLDI
San Diego 2022

Mon 13 - Fri 17 June 2022
San Diego, California, United States

Attending ▾ Program ▾ Tracks ▾ Organization ▾ 🔍 ↻ Sign in Sign up

🏠 [PLDI 2022 \(series\)](#) / [Commute \(series\)](#) /

Commutativity Reasoning and Applications

Commute

This is the first instance of the Commutativity Reasoning & Applications workshop (Commute 2022).



Commutativity Reasoning is becoming increasingly common and appears in many contexts. Commutativity is used in the design of systems, in the design of data structures, in proof methodologies, in parallel execution schemes, etc. The goal of this workshop is to bring together researchers that are working in a variety of areas, with a common need for commutativity, to share ideas and goals. We aim to include researchers who work on commutativity in many contexts: compilers, program logics, automata, concurrency, distributed systems/CRDTs, ML applications, etc.

The workshop will be held on Monday June 13th and Tuesday June 14th.

Call for Papers

The workshop is open to all who are interested and/or working in the area of Commutativity Reasoning and Applications. This includes researchers in

Organizing Committee

-  **Constantin Enea**
Ecole Polytechnique / LIX / CNRS
France
-  **Azadeh Farzan**
University of Toronto
Canada
-  **Eric Koskinen**
Stevens Institute of Technology

Interested in Commutativity?

Come visit!





CYPRESS

CYBERSECURITY
PROGRAMMING LANGUAGES
AND SYSTEMS AT STEVENS





Other recent things I didn't have time for ...

- Automatic **temporal** verification *CAV'11, POPL'11, PLDI'13, LICS'14, LICS'18*
- Automatic **relational** verification *PLDI'17, OOPSLA'19, CAV'21*
- Automatic crash recoverability *POPL'16*
- Verifying binaries *APLAS'21, IEEE S&P'21*
- Transactional implementations *APLAS'19, PODC'17, VMCAI'17, PPOPP'08*
- Semantics of transactions *POPL'10, PLDI'15*

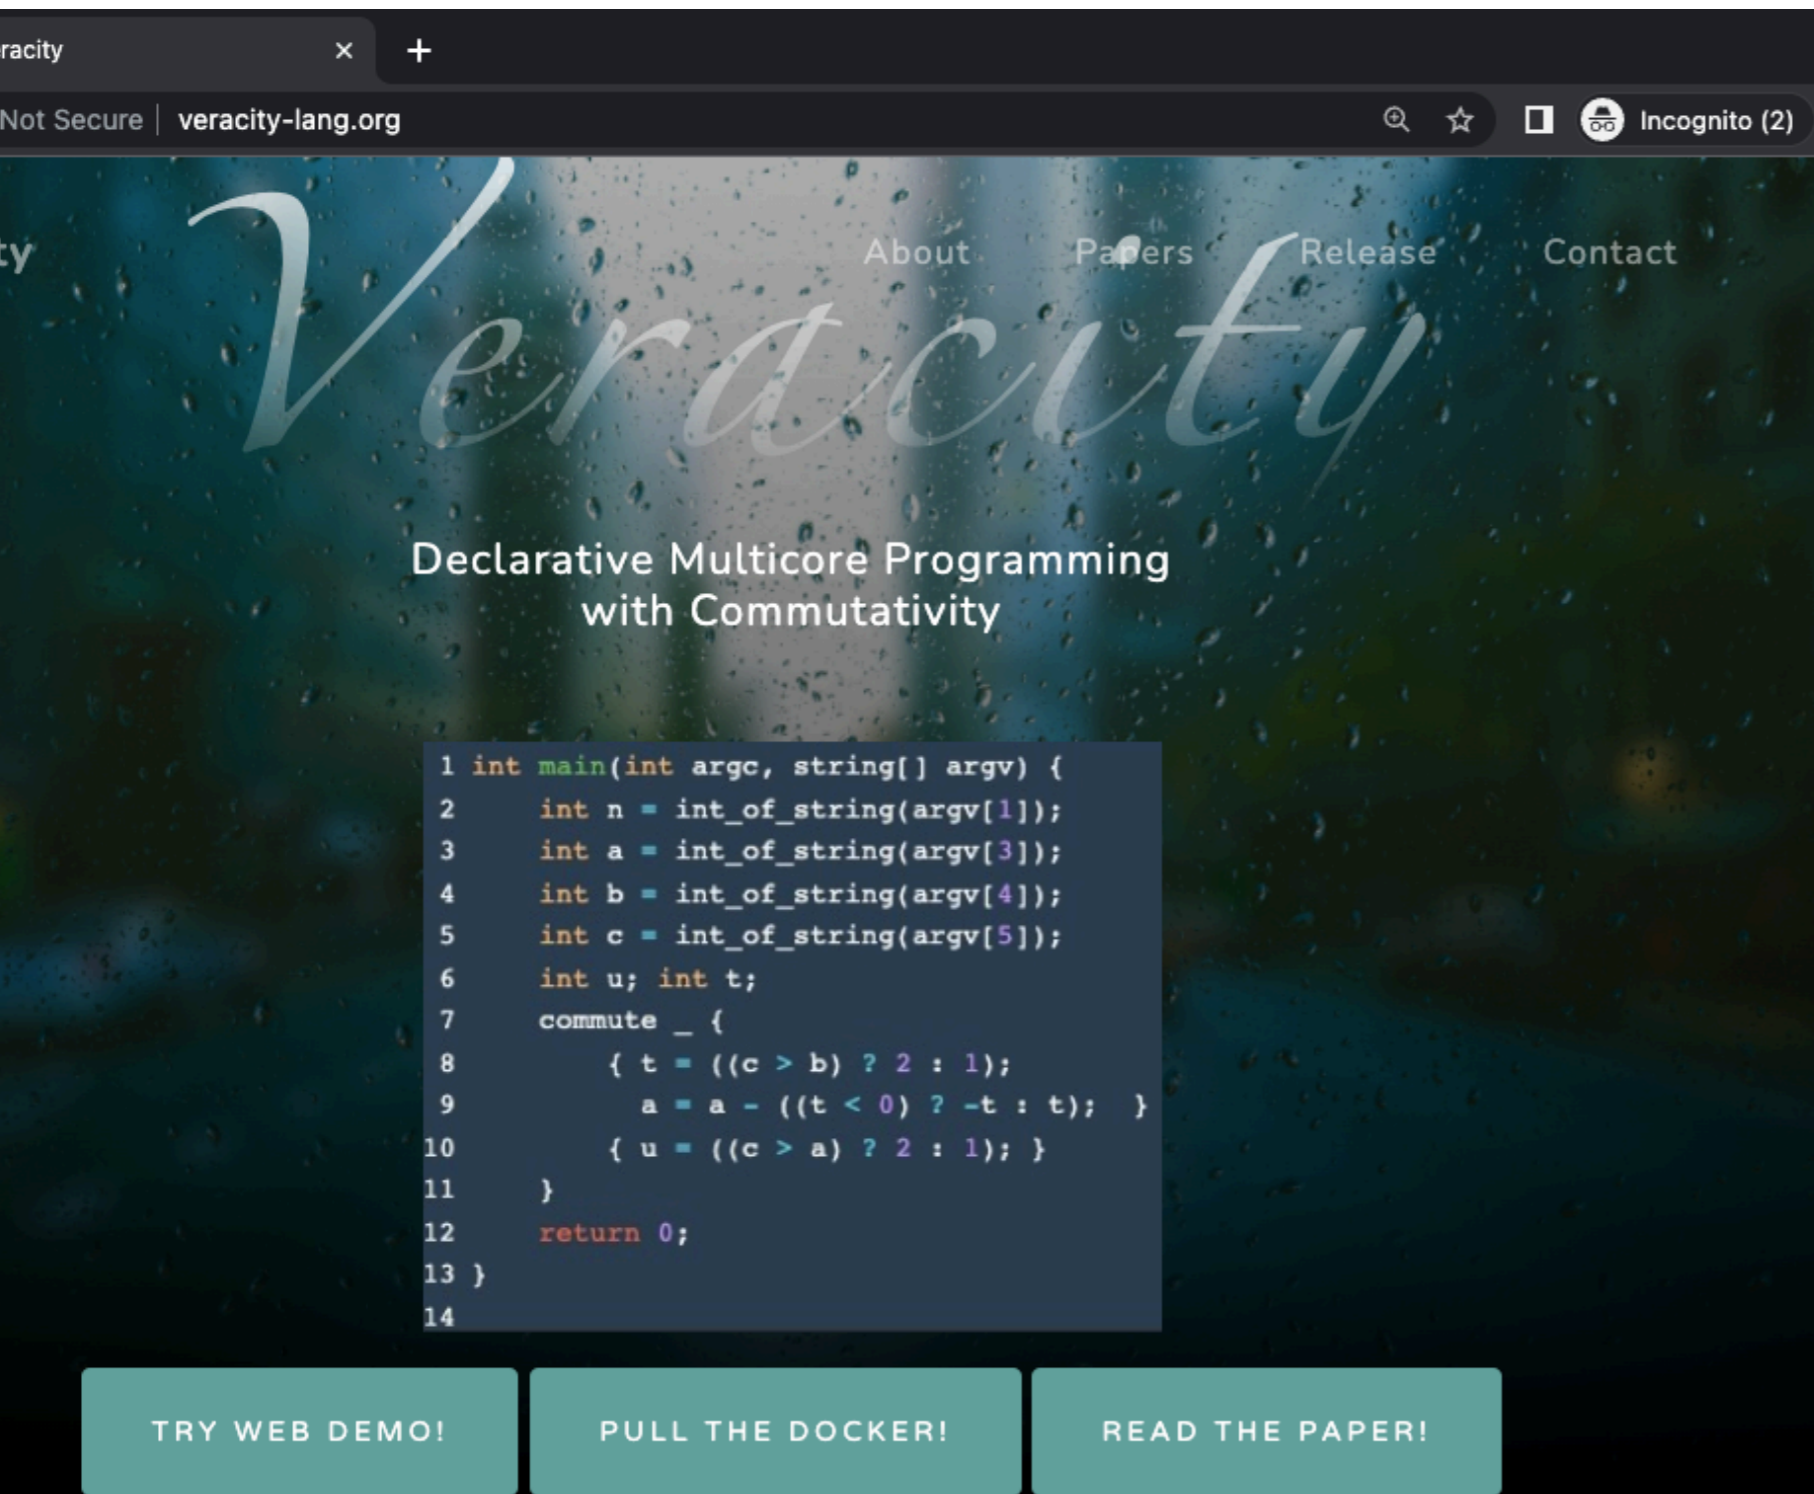


www.erickoskinen.com



Thank you!

Thank you!



Veracity

About Papers Release Contact

Declarative Multicore Programming
with Commutativity

```
1 int main(int argc, string[] argv) {
2     int n = int_of_string(argv[1]);
3     int a = int_of_string(argv[3]);
4     int b = int_of_string(argv[4]);
5     int c = int_of_string(argv[5]);
6     int u; int t;
7     commute _ {
8         { t = ((c > b) ? 2 : 1);
9           a = a - ((t < 0) ? -t : t); }
10        { u = ((c > a) ? 2 : 1); }
11    }
12    return 0;
13 }
14
```

TRY WEB DEMO! PULL THE DOCKER! READ THE PAPER!

www.veracity-lang.org

www.erickoskinen.com