

Reasoning about Eventual Consistency and Replicated Data Types

Alexey Gotsman

IMDEA Software Institute

MSR-IMDEA Collaboration Workshop, 04.04.14

*Joint work with Sebastian Burckhardt (Microsoft Research),
Hongseok Yang (Oxford) and Marek Zawirski (UPMC & INRIA)*

Shared-memory concurrency



Overview **Package** Class Use Tree Deprecated Index Help
PREV PACKAGE NEXT PACKAGE FRAMES NO FRAMES All Classes

Package java.util.concurrent

Utility classes commonly useful in concurrent programming.

See: [Description](#)

Interface Summary

BlockingDeque<E>	A Des
BlockingQueue<E>	A Que
Callable<V>	A tas
CompletionService<V>	A ser
ConcurrentMap<K,V>	A Map
ConcurrentNavigableMap<K,V>	A Con
Delayed	A mi
Executor	An o
ExecutorService	An E
Future<V>	A Fu
RejectedExecutionHandler	A ha
RunnableFuture<V>	A Fu
RunnableScheduledFuture<V>	A Sc
ScheduledExecutorService	An E
ScheduledFuture<V>	A de
ThreadFactory	An o

Class Summary

AbstractExecutorService
ArrayBlockingQueue<E>
ConcurrentHashMap<K,V>
ConcurrentLinkedQueue<E>
ConcurrentSkipListMap<K,V>
ConcurrentSkipListSet<E>

LEARN DOWNLOAD NEWS CONTRIBUTE LOGIN Search -Any-

Intel® Threading Building Blocks (Intel® TBB)

4.2 Update 3

DOWNLOAD NOW 18.42 MB
TBB 4.2 Update 3 for Linux*
All Operating Systems & Previous Releases

A cartoon illustration of a yellow bird perched on a brown branch.

Task Parallel Library (TPL)

.NET Framework 4.5 | [Other Versions](#) | 35 out of 45 rated this helpful - [Rate this topic](#)

The Task Parallel Library (TPL) is a set of public types and APIs in the [System.Threading](#) and [System.Threading.Tasks](#) namespaces. The purpose of the TPL is to make developers more productive by simplifying the process of adding parallelism and concurrency to applications. The TPL scales the degree of concurrency dynamically to most efficiently use all the processors that are available. In addition, the TPL handles the partitioning of the work, the scheduling of threads on the [ThreadPool](#), cancellation support, state management, and other low-level details. By using TPL, you can maximize the performance of your code while focusing on the work that your program is designed to accomplish.

Starting with the .NET Framework 4, the TPL is the preferred way to write multithreaded and parallel code. However, not all code is suitable for parallelization; for example, if a loop performs only a small amount of work on each iteration, or it doesn't run for many iterations, then the overhead of parallelization can cause the code to run more slowly. Furthermore, parallelization like any multithreaded code adds complexity to your program execution. Although the TPL simplifies multithreaded scenarios, we recommend that you have a basic understanding of threading concepts, for example, locks, deadlocks, and race conditions, so that you can use the TPL effectively.

Intro to Intel® TBB parallel_for

Task Parallel Library (TPL)

.NET Framework 4.5 | [Other Versions](#) | 35 out of 45 rated this helpful - [Rate this topic](#)

The Task Parallel Library (TPL) is a set of public types and APIs in the [System.Threading](#) and [System.Threading.Tasks](#) namespaces. The purpose of the TPL is to make developers more productive by simplifying the process of adding parallelism and concurrency to applications. The TPL scales the degree of concurrency dynamically to most efficiently use all the processors that are available. In addition, the TPL handles the partitioning of the work, the scheduling of threads on the [ThreadPool](#), cancellation support, state management, and other low-level details. By using TPL, you can maximize the performance of your code while focusing on the work that your program is designed to accomplish.

Starting with the .NET Framework 4, the TPL is the preferred way to write multithreaded and parallel code. However, not all code is suitable for parallelization; for example, if a loop performs only a small amount of work on each iteration, or it doesn't run for many iterations, then the overhead of parallelization can cause the code to run more slowly. Furthermore, parallelization like any multithreaded code adds complexity to your program execution. Although the TPL simplifies multithreaded scenarios, we recommend that you have a basic understanding of threading concepts, for example, locks, deadlocks, and race conditions, so that you can use the TPL effectively.

Distributed systems



Distributed systems

The image shows a screenshot of the Amazon.co.uk homepage. At the top, the browser address bar displays "Amazon.co.uk: Low Prices in Electronics, Books, Sports Equipment & more" and "www.amazon.co.uk". The Amazon logo is on the left, with navigation links for "Your Amazon.co.uk", "Today's Deals", "Gift Cards", and "Help". On the right, there are links for "January Deals", "Sign in Your Account", "Basket", and "Wish List".

Below the navigation bar, there is a section for Spanish customers: "¿Compras desde España? Shopping from Spain? Visita amazon.es". A horizontal menu lists "Amazon MP3", "Cloud Player", "Kindle", "LOVEFILM", "Appstore for Android", and "Audible".

The main content area features a large "Meet the Kindle Family" promotion. It includes images of Kindle devices and lists three models: "Kindle > £69", "Kindle Paperwhite > from £109", and "Kindle Fire HD > from £159". There is also a "Protect and personalise your Kindle" section with a link to "Kindle Accessories".

On the right side, there are two promotional boxes. The first is for "January Deals" with a "Shop now" link. The second is for a "Two-Hour Flying Lesson" for £99 (was £299), with a "See the deal" link and the "amazonlocal" logo. Below this is a "SHAMBALLA BRACELETS" promotion with a "Shop now" link and an image of a bracelet.

At the bottom, there is a "THE AMAZON CLOTHING STORE" banner titled "Like a Lady". It features images of women in elegant dresses and text: "Shop our sophisticated dress edit from Sugarhill Boutique, Eva Franco, Fornarina and more." There are links for "Shop Dresses" and "Shop All Clothing".

Below the clothing store banner is a section titled "What Other Customers Are Looking At Right Now" with a row of product thumbnails, including a book titled "SAFE HOUSE" and a Kindle device.

Geo-replicated databases



- Every data centre stores a complete replica of data
- Purpose: fault tolerance, minimising latency

Geo-replicated databases



- Every data centre stores a complete replica of data
- Purpose: **fault tolerance**, minimising latency

Geo-replicated databases



- Every data centre stores a complete replica of data
- Purpose: fault tolerance, **minimising latency**

Strong consistency



- Database behaves like a single replica
- Implementation: ensure replicas are in sync → wait until other replicas get updated

Strong consistency




- Database behaves like a single replica
- Implementation: ensure replicas are in sync →
wait until other replicas get updated

Strong consistency




- Problem: **high latency**, can't tolerate network partitions
- CAP theorem: impossible to get all of strong Consistency, Availability, Partition-tolerance

Strong consistency




- Problem: high latency, **can't tolerate network partitions**
- CAP theorem: impossible to get all of strong Consistency, Availability, Partition-tolerance

Strong consistency



- Problem: high latency, can't tolerate network partitions
- **CAP** theorem: impossible to get all of strong **C**onsistency, **A**vailability, **P**artition-tolerance

Strong consistency



- Problem: high latency, can't tolerate network partitions
- **CAP** theorem: impossible to get all of ~~strong~~ ~~Consistency~~, Availability, Partition-tolerance

Weak consistency




- Update your replica now, propagate to others later
- Weak consistency: exhibits anomalies

Weak consistency




- Update your replica now, propagate to others later
- **Weak consistency: exhibits anomalies**

Weak consistency




- Update your replica now, propagate to others later
- **Weak consistency: exhibits anomalies**

Weak consistency



- Update your replica now, propagate to others later
- **Weak consistency: exhibits anomalies**

Weak consistency



- Also an issue with mobile devices: operate when disconnected



Weak consistency

in shared memory

- Processors and programming languages don't provide strong consistency: **weak memory models**
- A multiprocessor is really a distributed system: cache-coherence protocol
- Hot topic now, but first had to define the memory models



The Semantics of x86-CC Multiprocessor Machine Code

Susmit Sarkar¹ Peter Sewell¹ Francesco Zappa Nardelli²

Scott Owens¹ Tom Ridge¹ Thomas Braibant² Magnus O. Myreen¹ Jade Alglave²

¹University of Cambridge ²INRIA

Abstract

Multiprocessors are not as simple as they do not provide the sequential semantics assumed by most work on multiprocessors. They have subtle relaxed memory models described only in an informal way that leads to confusion.

We develop a rigorous semantics for multiprocessor programs against a relaxed memory model. We formalise semantics against actual hardware examples, and give an operational characterisation of our axioms. We prove that they are (in some precise sense) in HOL that their behaviour matches the actual ARM behaviour. We also contrast the x86-CC model with ARM behaviour.

This provides a solid foundation for formal analysis and compilation.

Categories and Subject Descriptors: D.1.3 [Concurrent Programming]; F.3.1 [Specification and Verification]

General Terms: Documentation, Theory, Verification

Keywords: Relaxed Memory Model

1. Introduction

Problem Multiprocessors have been acting on a shared memory since the 1960s, but have suddenly become ubiquitous in the last few years: laptops, desktops, servers, or 16 cores, and the trend is to continue. Meanwhile, the current systems have grown in size. Last 40 years on semantics, the work has almost always been on sequential models. Internally

Permission to make digital or hard copies for personal or classroom use is granted by ACM, provided that the copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the copyright notice. To copy otherwise, to republish, to post on servers, to redistribute to lists, requires prior specific permission and may be charged a fee. Copyright © 2009 ACM.

A Better x86 Memory Model: x86-TSO

Scott Owens Susmit Sarkar Peter Sewell

University of Cambridge

Abstract

Abstract. This paper introduces a new memory model for x86-CC multiprocessors, called x86-TSO. It is a refinement of the x86-CC model, and is designed to be more precise and easier to verify. We discuss the motivation for this model, and how it differs from the current x86-CC model. We also provide a formal semantics for x86-TSO, and show that it is a refinement of the x86-CC model. Both the current x86-CC model and the new x86-TSO model are shown to be valid models of the hardware. We witness the correctness of the new model more precisely.

1. Introduction

Most previous work on multiprocessors assumes sequential consistency. This is a very restrictive model, and does not capture the behaviour of multiprocessors. In particular, it does not capture the behaviour of multiprocessors with relaxed memory models. This paper introduces a new memory model for x86-CC multiprocessors, called x86-TSO. It is a refinement of the x86-CC model, and is designed to be more precise and easier to verify. We discuss the motivation for this model, and how it differs from the current x86-CC model. We also provide a formal semantics for x86-TSO, and show that it is a refinement of the x86-CC model. Both the current x86-CC model and the new x86-TSO model are shown to be valid models of the hardware. We witness the correctness of the new model more precisely.

Abstract

We develop a rigorous semantics for multiprocessor programs, including their hardware behaviour and the behaviour of reasonable fraction sets. The semantics is mechanised in HOL. It is a refinement of the current x86-CC model, and is designed to be more precise and easier to verify.

This should provide a good basis for formal verification of low-level consistent architectures, and, together with the hardware semantics, for the design and compilation of multiprocessor languages.

Categories and Subject Descriptors: D.1.3 [Concurrent Programming]; F.3.1 [Specification and Verification]

General Terms: Documentation, Theory, Verification

Keywords: Relaxed Memory Model, x86-CC, ARM

The Semantics of Power and ARM Multiprocessor Machine Code

Jade Alglave² Anthony Fox¹ Samin Ishtiaq³ Magnus O. Myreen¹
Susmit Sarkar¹ Peter Sewell¹ Francesco Zappa Nardelli²

¹University of Cambridge

Understanding POWER Multiprocessors

Susmit Sarkar¹ Peter Sewell¹ Jade Alglave^{2,3} Luc Maranget³ Derek Williams⁴

¹University of Cambridge ²Oxford University ³INRIA ⁴IBM Austin

Abstract

Exploiting today's multiprocessors requires high-performance and correct concurrent systems code (optimising compilers, language runtimes, OS kernels, etc.), which in turn requires a good understanding of the observable processor behaviour that can be relied on. Unfortunately this critical hardware/software interface is not at all clear for several current multiprocessors.

In this paper we characterise the behaviour of IBM POWER multiprocessors, which have a subtle and highly relaxed memory model (ARM multiprocessors have a very similar architecture in this respect). We have conducted extensive experiments on several generations of processors: POWER G5, 5, 6, and 7. Based on these, on published details of the microarchitectures, and on discussions with IBM staff, we give an abstract-machine semantics that abstracts

many years had aggressive implementations, providing high performance but exposing a very relaxed memory model, one that requires careful use of dependencies and memory barriers to enforce ordering in concurrent code. A priori, one might expect the behaviour of a multiprocessor to be sufficiently well-defined by the vendor architecture documentation, here the Power ISA v2.06 specification [Pow09]. For the sequential behaviour of instructions, that is very often true. For concurrent code, however, the observable behaviour of Power multiprocessors is extremely subtle, as we shall see, and the guarantees given by the vendor specification are not always clear. We therefore set out to discover the actual processor behaviour and to define a rigorous and usable semantics, as a foundation for future system building and research.

The programmer-observable relaxed-memory behaviour of these multiprocessors emerges as a whole system prop-

The Semantics of x86-CC Multiprocessor Machine Code

Susmit Sarkar¹ Peter Sewell¹ Francesco Zappa Nardelli²

Scott Owens¹ Tom Ridge¹ Thomas Braibant² Magnus O. Myreen¹ Jade Alglave²

¹University of Cambridge ²INRIA

Abstract

Multiprocessors are not as simple as they do not provide the sequential semantics assumed by most work on multiprocessors; they have subtle relaxed memory models described only in an informal way that leads to confusion.

We develop a rigorous semantics for multiprocessor programs against a relaxed memory model, and give examples of behaviours that are not captured by the semantics. We characterise our axioms and prove that they are (in some sense) the most liberal in HOL that their behaviour is consistent with ARM behaviour.

This provides a solid foundation for formal analysis, and compilation.

Categories and Subject Descriptors: D.1.3 [Concurrent Programming]; F.3.1 [Specification and Verification]

General Terms: Documentation, Theory, Verification

Keywords: Relaxed Memory Model

1. Introduction

Problem Multiprocessors have been acting on a shared memory since the 1960s, but have suddenly become ubiquitous in the last few years: laptops, desktops, servers, or 16 cores, and the trend is to continue. Meanwhile, the current systems has grown in size. Last 40 years on semantics work has almost always been about sequential execution, and share a sequentially consistent memory model. Internally

Permission to make digital or hard copies for personal or classroom use is granted by ACM, provided that the copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the copyright notice. To copy otherwise, to republish, to post on servers, to redistribute to lists, requires prior specific permission and may be charged a fee. Copyright © 2009 ACM.

A Better x86 Memory Model: x86-TSO

Scott Owens Susmit Sarkar Peter Sewell

University of Cambridge

Abstract

This paper introduces a new memory model for x86 multiprocessors, called x86-TSO. It is a simplification of the x86-TSO model, and is designed to be more amenable to formal verification. We discuss the motivation for this model, and how it differs from the current x86-TSO model. We also discuss the implications of this model for the design of multiprocessors.

1. Introduction

Most previous work on multiprocessors assumes sequential consistency. This is a very restrictive model, and does not capture the behaviour of multiprocessors. In particular, it does not capture the behaviour of multiprocessors that use relaxed memory models. This paper introduces a new memory model for x86 multiprocessors, called x86-TSO. It is a simplification of the x86-TSO model, and is designed to be more amenable to formal verification. We discuss the motivation for this model, and how it differs from the current x86-TSO model. We also discuss the implications of this model for the design of multiprocessors.

Abstract

We develop a rigorous semantics for multiprocessor programs, including their behaviour and the behaviour of reasonable fraction sets. The semantics is mechanised in HOL. This should provide a good basis for formal verification of low-level consistent architectures, and, together with the semantics, for the design and compilation of multiprocessor languages.

This should provide a good basis for formal verification of low-level consistent architectures, and, together with the semantics, for the design and compilation of multiprocessor languages.

Categories and Subject Descriptors: D.1.3 [Concurrent Programming]; F.3.1 [Specification and Verification]

General Terms: Documentation, Theory, Verification

Keywords: Relaxed Memory Model, x86-CC, ARM

The Semantics of Power and ARM Multiprocessor Machine Code

Jade Alglave² Anthony Fox¹ Samin Ishtiaq³ Magnus O. Myreen¹
Susmit Sarkar¹ Peter Sewell¹ Francesco Zappa Nardelli²

¹University of Cambridge

Understanding POWER Multiprocessors

Susmit Sarkar¹ Peter Sewell¹ Jade Alglave^{2,3} Luc Maranget³ Derek Williams⁴

¹University of Cambridge ²Oxford University ³INRIA ⁴IBM Austin

Abstract

Exploiting today's multiprocessors requires high-performance and correct concurrent systems code (optimising compilers, language runtimes, OS kernels, etc.), which in turn requires a good understanding of the observable processor behaviour that can be relied on. Unfortunately this critical hardware/software interface is not at all clear for several current multiprocessors.

In this paper we characterise the behaviour of IBM POWER multiprocessors, which have a subtle and highly relaxed memory model (ARM multiprocessors have a very similar architecture in this respect). We have conducted extensive experiments on several generations of processors: POWER G5, 5, 6, and 7. Based on these, on published details of the microarchitectures, and on discussions with IBM staff, we give an abstract-machine semantics that abstracts

many years had aggressive implementations, providing high performance but exposing a very relaxed memory model, one that requires careful use of dependencies and memory barriers to enforce ordering in concurrent code. A priori, one might expect the behaviour of a multiprocessor to be sufficiently well-defined by the vendor architecture documentation, here the Power ISA v2.06 specification [Pow09]. For the sequential behaviour of instructions, that is very often true. For concurrent code, however, the observable behaviour of Power multiprocessors is extremely subtle, as we shall see, and the guarantees given by the vendor specification are not always clear. We therefore set out to discover the actual processor behaviour and to define a rigorous and usable semantics, as a foundation for future system building and research.

The programmer-observable relaxed-memory behaviour of these multiprocessors emerges as a whole system prop-

If no new updates are made to the database, then replicas will eventually reach a consistent state

practice

DOI:10.1145/1435417.1435432

Building reliable distributed systems at a worldwide scale demands trade-offs between consistency and availability.

BY WERNER VOGELS

Eventually Consistent

AT THE FOUNDATION of Amazon's cloud computing are infrastructure services such as Amazon's S3 (Simple Storage Service), SimpleDB, and EC2 (Elastic Compute Cloud) that provide the resources for constructing Internet-scale computing platforms and a great variety of applications. The requirements placed on these infrastructure services are very strict; they need to score high marks in the areas of security, scalability, availability, performance, and cost-effectiveness, and they need to meet these requirements while serving millions of customers around the globe, continuously.

Under the covers these services are massive distributed systems that operate on a worldwide scale. This scale creates additional challenges, because when a system processes trillions and trillions of requests, events that normally have a low probability of occurrence are now guaranteed to happen and must be accounted for upfront in the design and architecture of the system. Given the worldwide scope of these systems, we use replication techniques ubiquitously to guarantee consistent performance and high availability. Although replication brings us closer to our goals, it cannot achieve them in a perfectly

transparent manner; under a number of conditions the customers of these services will be confronted with the consequences of using replication techniques inside the services.

One of the ways in which this manifests itself is in the type of data consistency that is provided, particularly when many widespread distributed systems provide an *eventual consistency* model in the context of data replication. When designing these large-scale systems at Amazon, we use a set of guiding principles and abstractions related to large-scale data replication and focus on the trade-offs between high availability and data consistency. Here, I present some of the relevant background that has informed our approach to delivering reliable distributed systems that must operate on a global scale. (An earlier version of this article appeared as a posting on the "All Things Distributed" Weblog and was greatly improved with the help of its readers.)

Historical Perspective

In an ideal world there would be only one consistency model: when an update is made all observers would see that update. The first time this surfaced as difficult to achieve was in the database systems of the late 1970s. The best "period piece" on this topic is "Notes on Distributed Databases" by Bruce Lindsay et al.¹ It lays out the fundamental principles for database replication and discusses a number of techniques that deal with achieving consistency. Many of these techniques try to achieve *distribution transparency*—that is, to the user of the system it appears as if there is only one system instead of a number of collaborating systems. Many systems during this time took the approach that it was better to fail the complete system than to break this transparency.²

In the mid-1990s, with the rise of larger Internet systems, these practices were revisited. At that time people began to consider the idea that availability was perhaps the most impor-

If no new updates are made to the database, then replicas will eventually reach a consistent state

But updates never stop!

So what does this tell database clients?

practice

DOI:10.1145/1435417.1435432

Building reliable distributed systems at a worldwide scale demands trade-offs between consistency and availability.

BY WERNER VOGELS

Eventually Consistent

AT THE FOUNDATION of Amazon's cloud computing are infrastructure services such as Amazon's S3 (Simple Storage Service), SimpleDB, and EC2 (Elastic Compute Cloud) that provide the resources for constructing Internet-scale computing platforms and a great variety of applications. The requirements placed on these infrastructure services are very strict; they need to score high marks in the areas of security, scalability, availability, performance, and cost-effectiveness, and they need to meet these requirements while serving millions of customers around the globe, continuously.

Under the covers these services are massive distributed systems that operate on a worldwide scale. This scale creates additional challenges, because when a system processes trillions and trillions of requests, events that normally have a low probability of occurrence are now guaranteed to happen and must be accounted for upfront in the design and architecture of the system. Given the worldwide scope of these systems, we use replication techniques ubiquitously to guarantee consistent performance and high availability. Although replication brings us closer to our goals, it cannot achieve them in a perfectly

transparent manner; under a number of conditions the customers of these services will be confronted with the consequences of using replication techniques inside the services.

One of the ways in which this manifests itself is in the type of data consistency that is provided, particularly when many widespread distributed systems provide an *eventual consistency* model in the context of data replication. When designing these large-scale systems at Amazon, we use a set of guiding principles and abstractions related to large-scale data replication and focus on the trade-offs between high availability and data consistency. Here, I present some of the relevant background that has informed our approach to delivering reliable distributed systems that must operate on a global scale. (An earlier version of this article appeared as a posting on the "All Things Distributed" Weblog and was greatly improved with the help of its readers.)

Historical Perspective

In an ideal world there would be only one consistency model: when an update is made all observers would see that update. The first time this surfaced as difficult to achieve was in the database systems of the late 1970s. The best "period piece" on this topic is "Notes on Distributed Databases" by Bruce Lindsay et al.¹ It lays out the fundamental principles for database replication and discusses a number of techniques that deal with achieving consistency. Many of these techniques try to achieve *distribution transparency*—that is, to the user of the system it appears as if there is only one system instead of a number of collaborating systems. Many systems during this time took the approach that it was better to fail the complete system than to break this transparency.²

In the mid-1990s, with the rise of larger Internet systems, these practices were revisited. At that time people began to consider the idea that availability was perhaps the most impor-

50 shades of eventual consistency

Don't Settle for Eventual: Scalable Causal Consistency for Wide-Area Storage with COPS

Wyatt Lloyd*, Michael J. Freedman*, Michael Kaminsky†, and David G. Andersen‡
*Princeton University, †Intel Labs, ‡Carnegie Mellon University

Consistency-Based Service Level Agreements for Cloud Storage

Douglas B. Terry, Vijayan Prabhakaran, Ramakrishna Kotla,
Mahesh Balakrishnan, Marcos K. Aguilera, Hussam Abu-Libdeh†

Microsoft Research Silicon Valley
University

Eventually Consistent Transactions

Sebastian Burckhardt¹, Daan Leijen¹, Manuel Fähndrich¹, and Mooly Sagiv²

¹ Microsoft Research

² Tel-Aviv University

Conflict-free Replicated Data Types *

Marc Shapiro^{1,5}, Nuno Preguiça^{2,1}, Carlos Baquero³, and Marek Zawirski^{1,4}

¹ INRIA, Paris, France

² CITI Universidade Nova de Lisboa, Portugal

³ Universidade Nova de Lisboa, Portugal

50 shades of eventual consistency

**Don't Settle for Eventual:
Scalable Causal Consistency for Wide-Area Storage with COPS**

Wyatt Lloyd*, Michael J. Freedman*, Michael Kaminsky†, and David G. Andersen‡
*Princeton University, †Intel Labs, ‡Carnegie Mellon University

Pileus →

**Consistency-Based Service Level Agreements
for Cloud Storage**

Douglas B. Terry, Vijayan Prabhakaran, Ramakrishna Kotla,
Mahesh Balakrishnan, Marcos K. Aguilera, Hussam Abu-Libdeh†
Microsoft Research Silicon Valley
University

Eventually Consistent Transactions

Sebastian Burckhardt¹, Daan Leijen¹, Manuel Fähndrich¹, and Mooly Sagiv²

¹ Microsoft Research
² Tel-Aviv University

Conflict-free Replicated Data Types *

Marc Shapiro^{1,5}, Nuno Preguiça^{2,1}, Carlos Baquero³, and Marek Zawirski^{1,4}

¹ INRIA, Paris, France

² CITI Universidade Nova de Lisboa, Portugal
³ CITI Universidade Nova de Lisboa, Portugal
⁴ CITI Universidade Nova de Lisboa, Portugal
⁵ CITI Universidade Nova de Lisboa, Portugal

TouchDevelop

50 shades of eventual consistency

**Don't Settle for Eventual:
Scalable Causal Consistency for Wide-Area Storage with COPS**

Wyatt Lloyd*, Michael J. Freedman*, Michael Kaminsky†, and David G. Andersen‡
*Princeton University, †Intel Labs, ‡Carnegie Mellon University

Pileus →

**Consistency-Based Service Level Agreements
for Cloud Storage**

Douglas B. Terry, Vijayan Prabhakaran, Ramakrishna Kotla,
Mahesh Balakrishnan, Marcos K. Aguilera, Hussam Abu-Libdeh†
Microsoft Research Silicon Valley
University

Eventually Consistent Transactions

Sebastian Burckhardt¹, Daan

TouchDevelop

- + Stronger guarantees and programming interfaces with nontrivial semantics
- Low-level semantics definitions or none at all: hard to reason about database behaviour

Key issues beyond 'eventual'

If updates stop, replicas will eventually reach the same state

1. Which **anomalies** can we see before this?
E.g., does a user always see his own actions?
2. Which state will replicas converge to?
Users can make conflicting updates.
How does the database resolve the **conflicts**?

Set ~ Shopping cart



`set.add(laptop)`

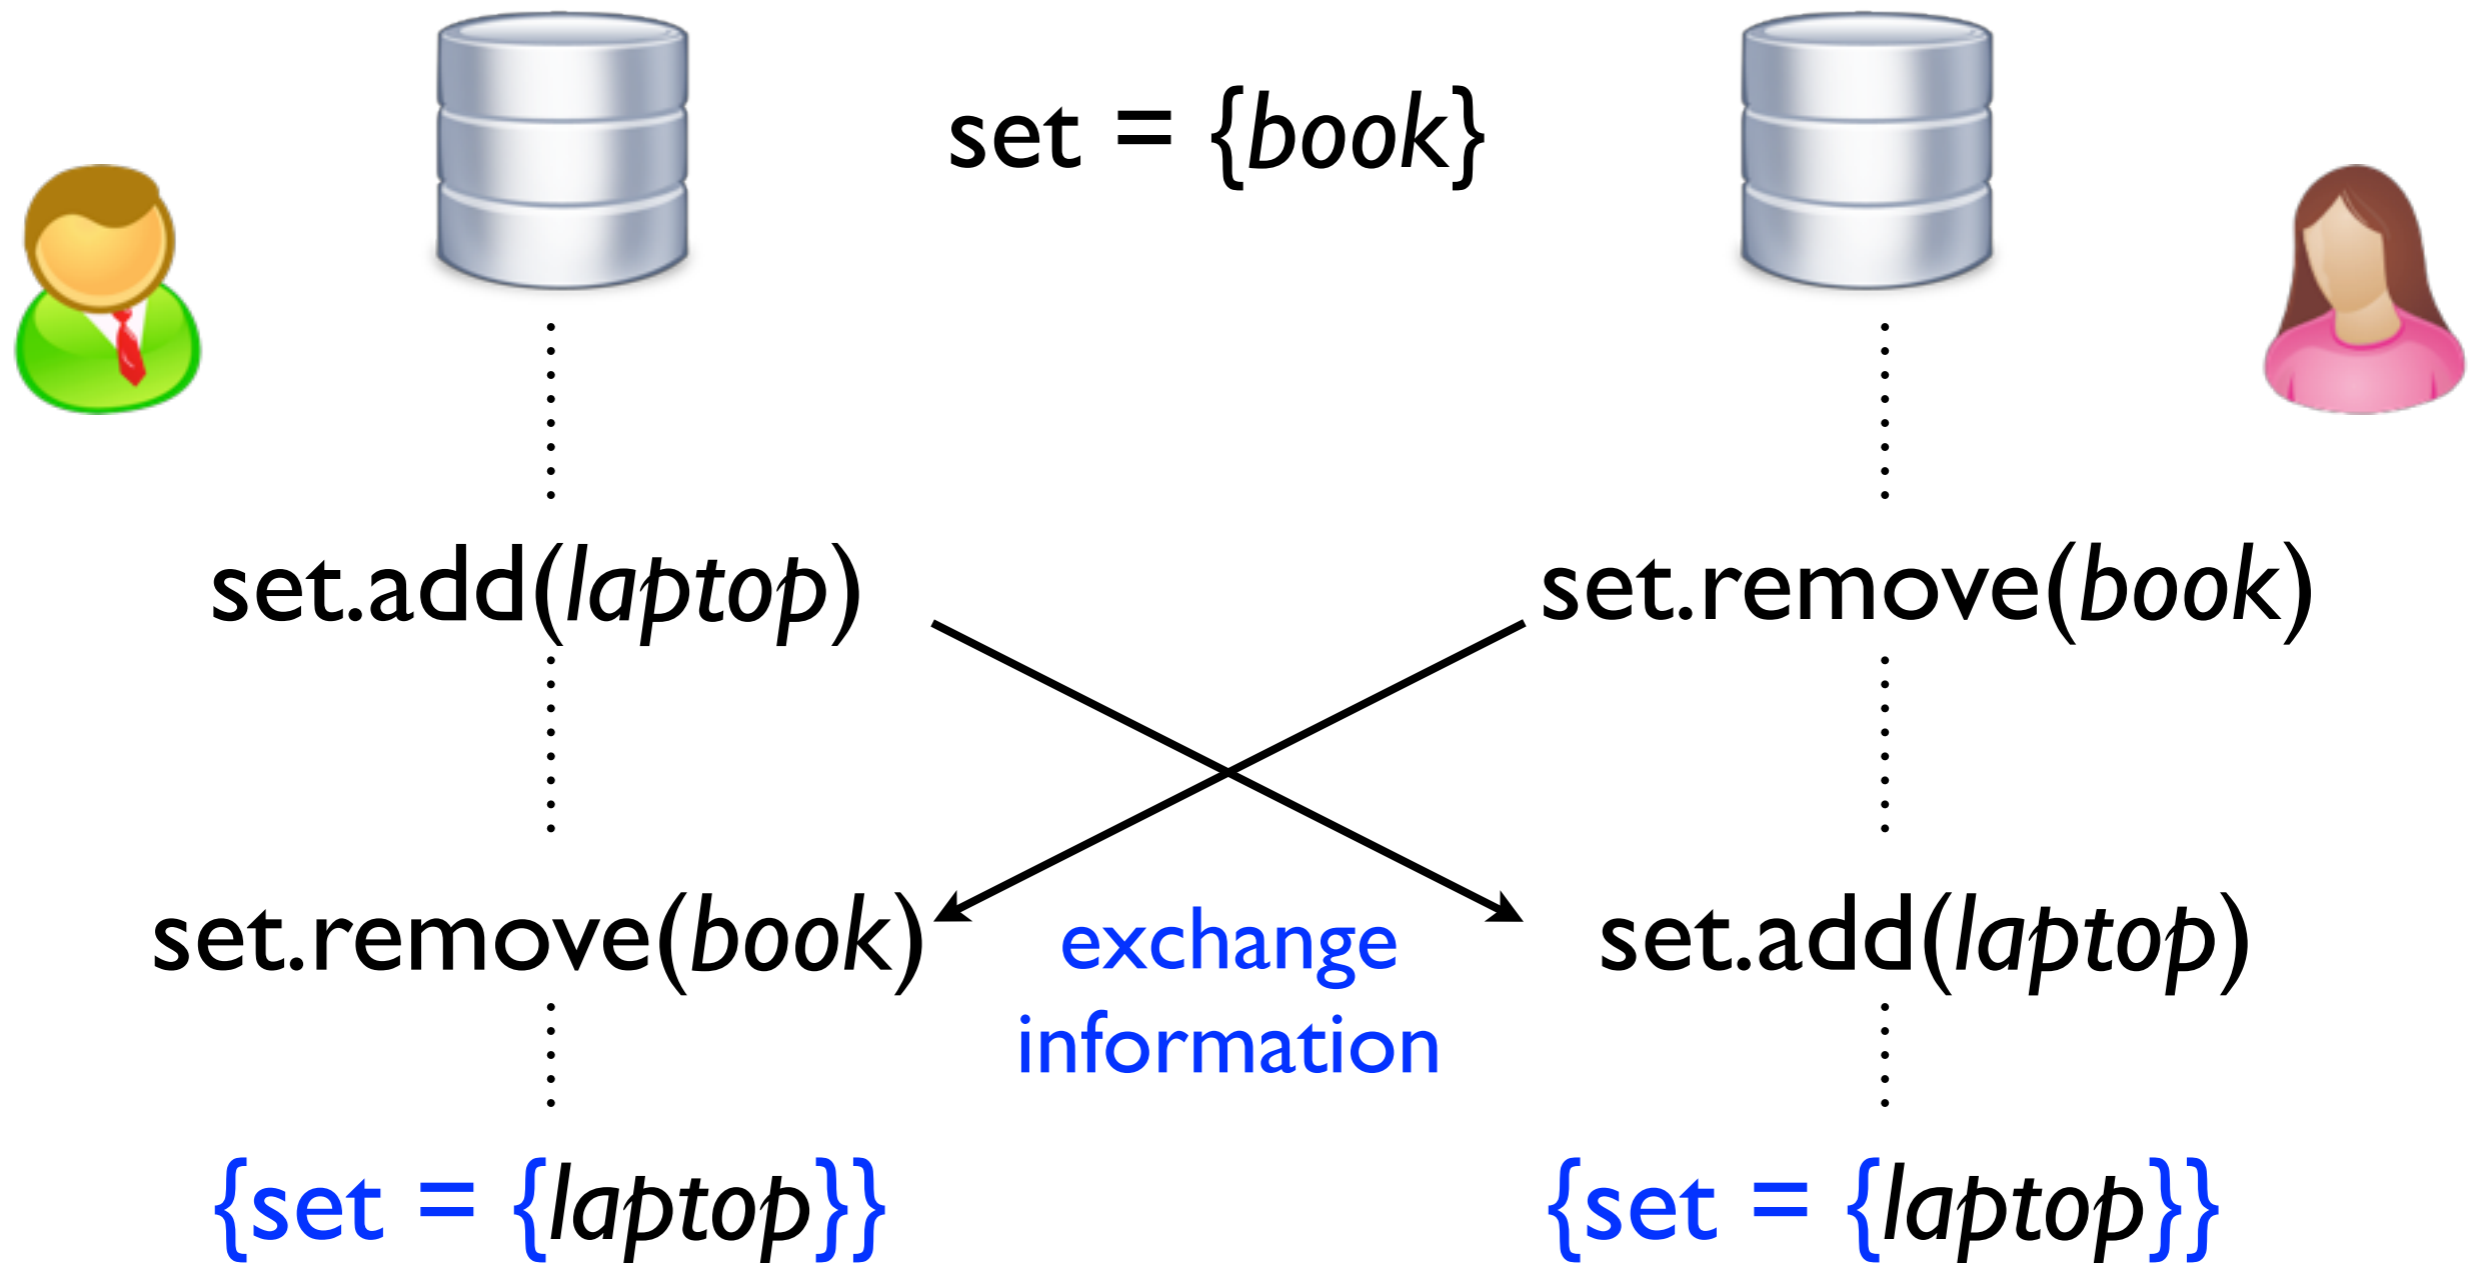
`set = {book}`



`set.remove(book)`




Set ~ Shopping cart



Operations **commute** → eventual consistency OK

Set ~ Shopping cart



`set.add(book)`

`set = {book}`



`set.remove(book)`



Set ~ Shopping cart



`set.add(book)`

`set = {book}`

Conflict!



`set.remove(book)`



Should the remove cancel the concurrent add?

Depends on application requirements

Set ~ Shopping cart



`set.add(book)`

`set = {book}`



`set.remove(book)`



Conflict!

Remove wins:

`set = \emptyset`

Add wins:

`set = {book}`

Last writer wins:

choose based on operation
time-stamps

Set ~ Shopping cart



`set.add(book)`

`set = {book}`



`set.remove(book)`



Conflict!

Remove wins:

`set = ∅`

Add wins:

`set = {book}`

Last writer wins:

choose based on operation
time-stamps

Replicated data types

aka CRDTs, cloud types

Object → Type → Conflict resolution policy

- Many data types: registers, counters, graphs, lists, file systems [Shapiro+ 2011]
- Nontrivial implementations

Replicated data types

aka CRDTs, cloud types

Object → Type → Conflict resolution policy

- Many data types: registers, counters, graphs, lists, file systems [Shapiro+ 2011]
- Nontrivial implementations

So far: implementation is your specification

Long-term goal

Use formal techniques to:

- Define the semantics of eventually consistent databases
- Develop tools for reasoning about their behaviour
- Improve programmability and efficiency

Results [POPL'14]

- **Specification:**
 - ▶ Conflict resolution ~ replicated data types
 - ▶ Anomalies

Results [POPL'14]

- **Specification:**
 - ▶ **Conflict resolution ~ replicated data types**
 - ▶ **Anomalies**
- **Verification:** framework for proving correctness of replicated data type implementations

Results [POPL'14]

- **Specification:**
 - ▶ **Conflict resolution ~ replicated data types**
 - ▶ Anomalies
- **Verification:** framework for proving correctness of replicated data type implementations
- **Optimality:**
 - ▶ Data types maintain metadata for conflict resolution
 - ▶ Method for proving lower bounds on metadata space requirements

Results [POPL'14]

- **Specification:**
 - ▶ **Conflict resolution** ~ replicated data types
 - ▶ Anomalies
- **Verification:** framework for proving correctness of replicated data type implementations
- **Optimality:**
 - ▶ Data types maintain metadata for conflict resolution
 - ▶ Method for proving lower bounds on metadata space requirements
- **Applications** to nontrivial data types

Results [POPL'14]

- **Specification:**
 - ▶ Conflict resolution ~ replicated data types
 - ▶ Anomalies

Results [POPL'14]

- **Specification:**

- ▶ Conflict resolution ~ replicated data types
- ▶ Anomalies



Replicated data
type specifications

Results [POPL'14]

- **Specification:**

- ▶ Conflict resolution ~ replicated data types
- ▶ Anomalies

Consistency
axioms

The diagram consists of two rectangular boxes with blue borders. The left box contains the text 'Consistency axioms'. The right box contains the text 'Replicated data type specifications'. A blue arrow points from the top-right corner of the left box to the top-left corner of the right box. Another blue arrow points from the top-right corner of the right box to the text 'Conflict resolution ~ replicated data types' in the list above.

Replicated data
type specifications

Sequential data type semantics

Strong consistency \rightarrow operations are totally ordered:

`set.add(book)`



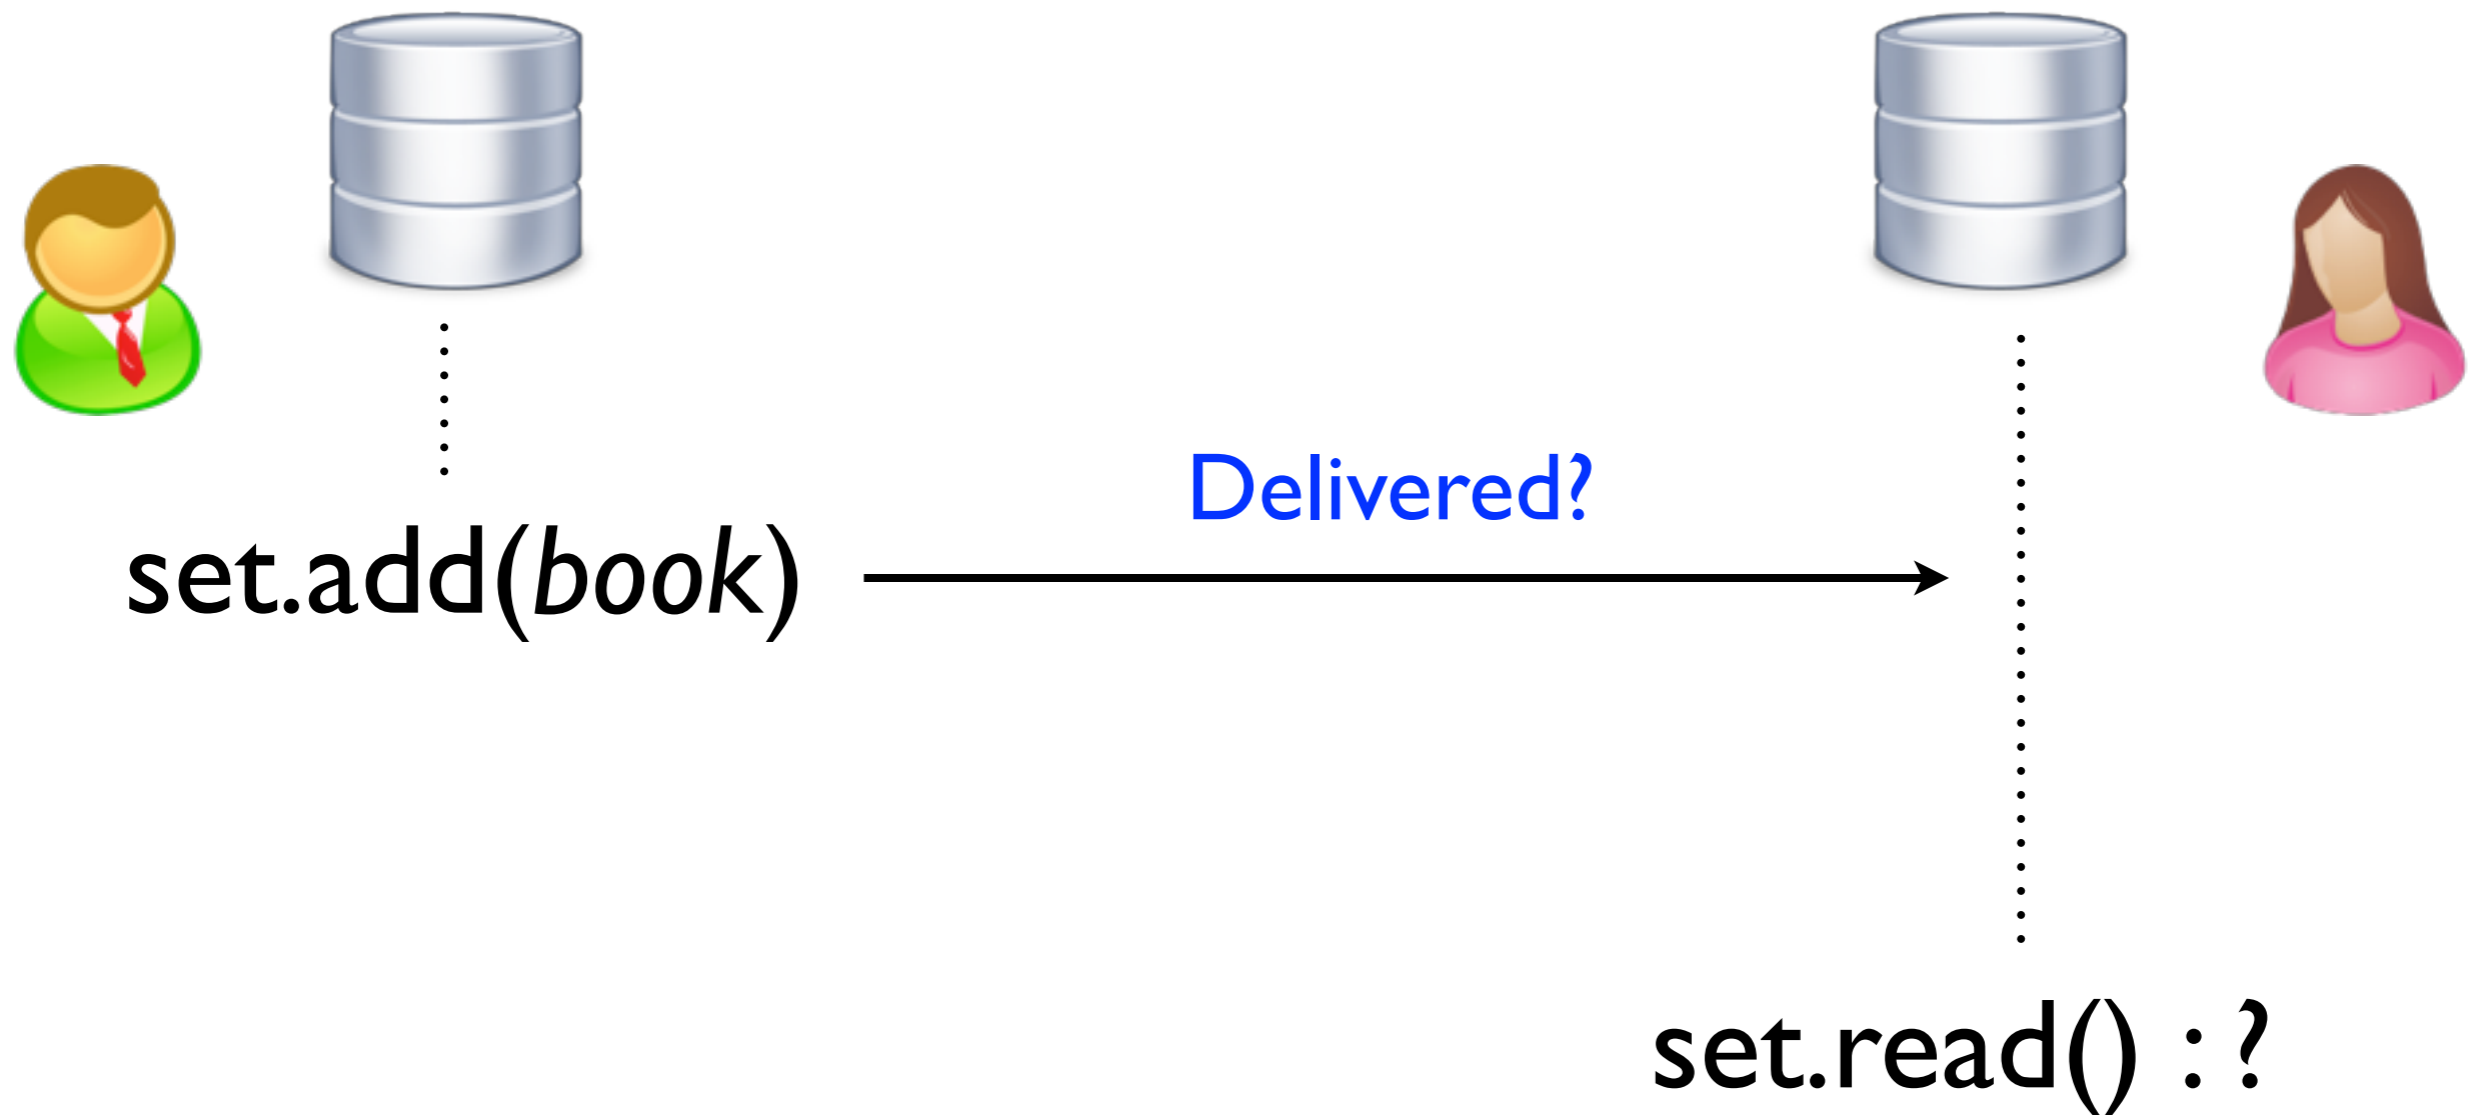
`set.remove(book)`



`set.read() : \emptyset`

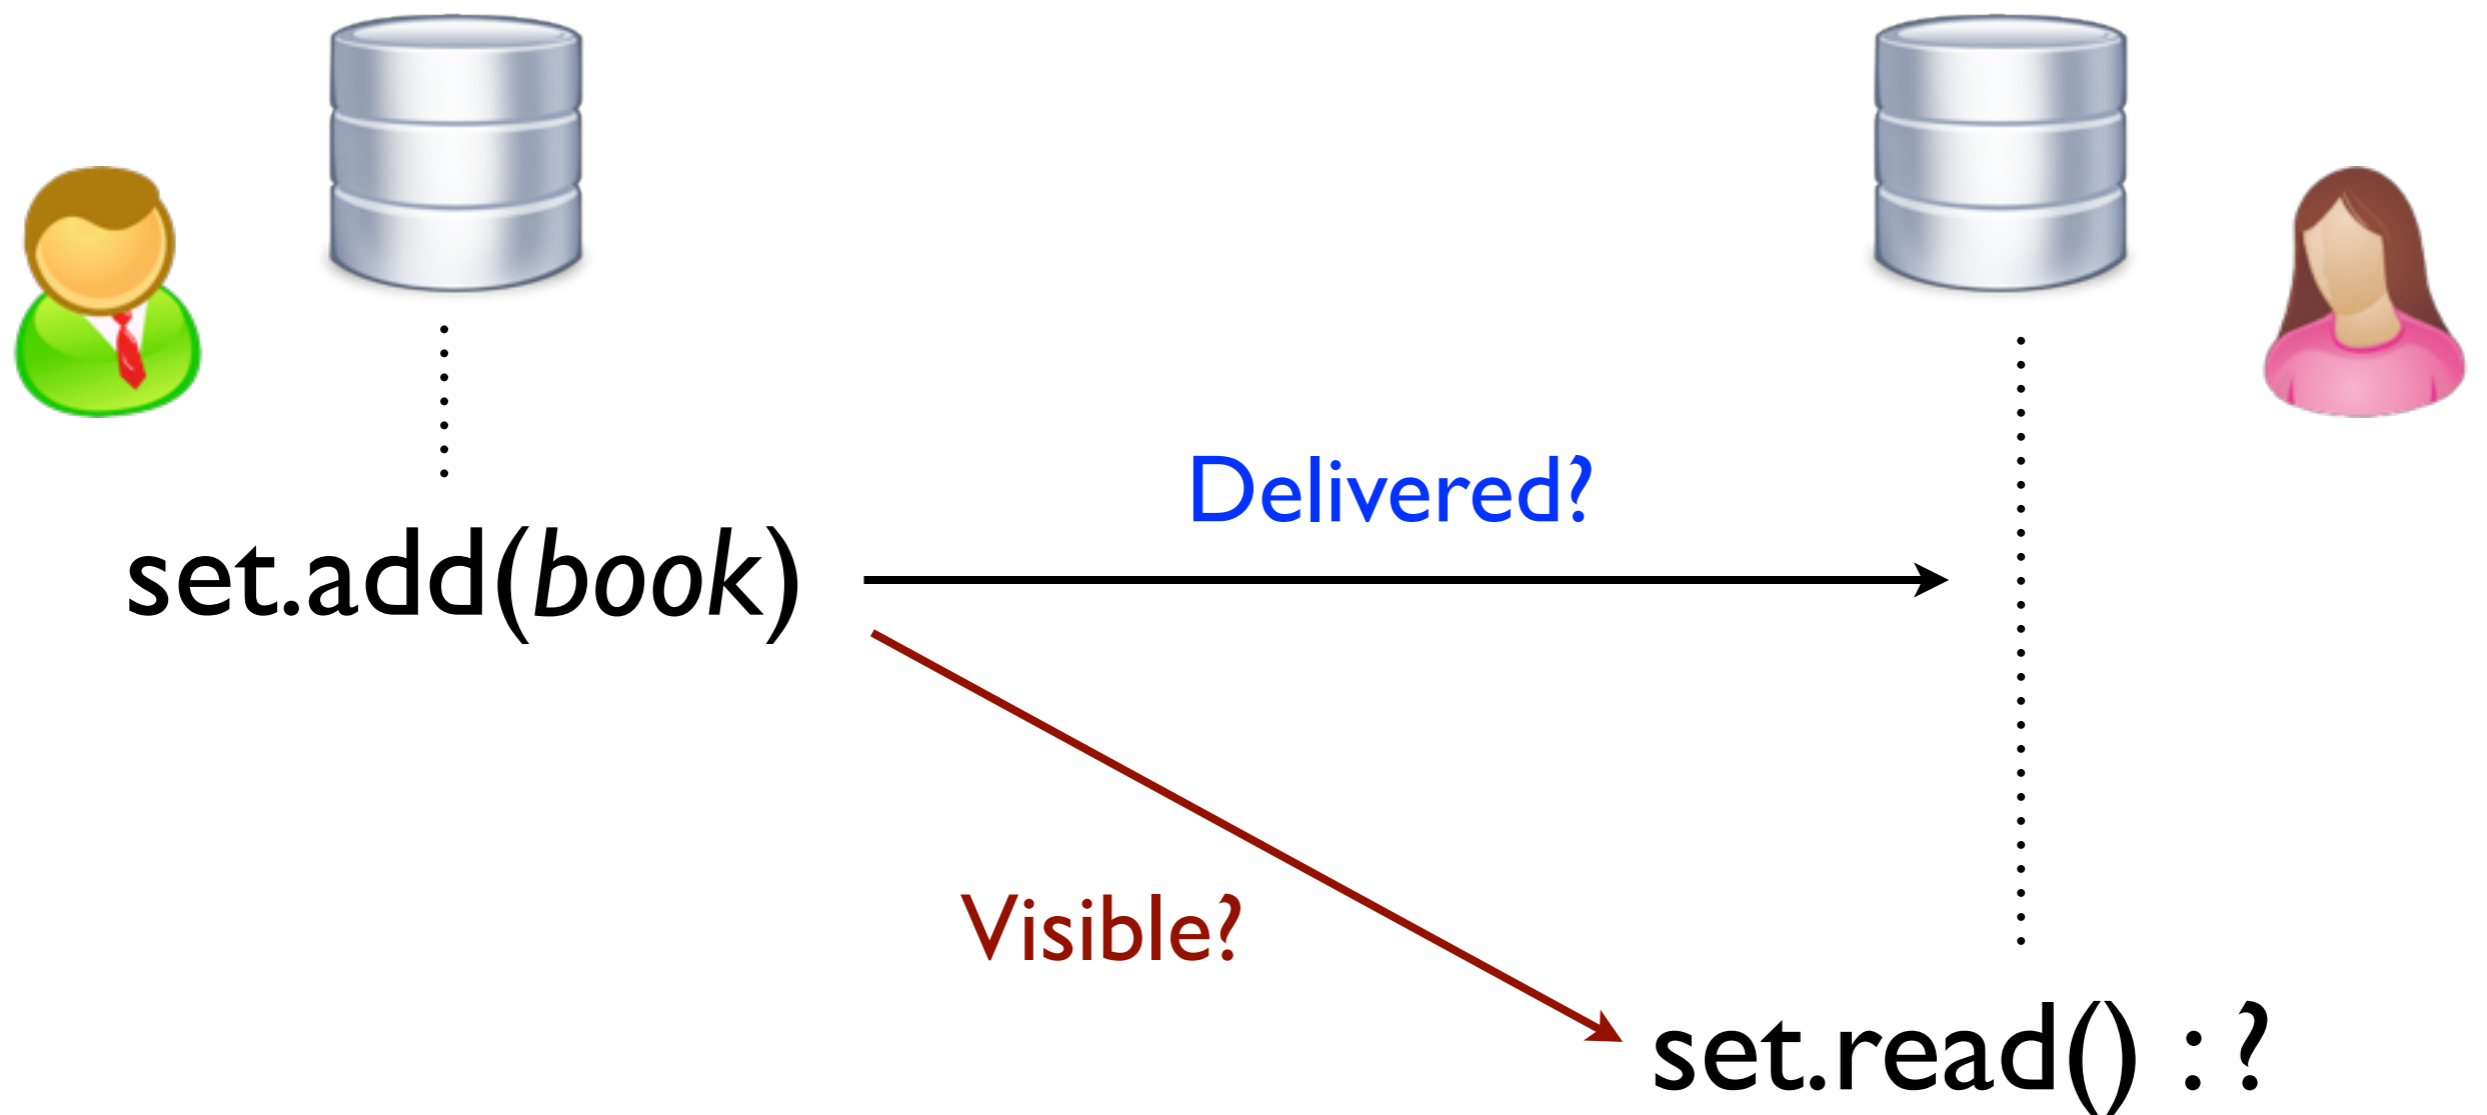
Compute the result by applying operations in sequence

Replicated data type semantics



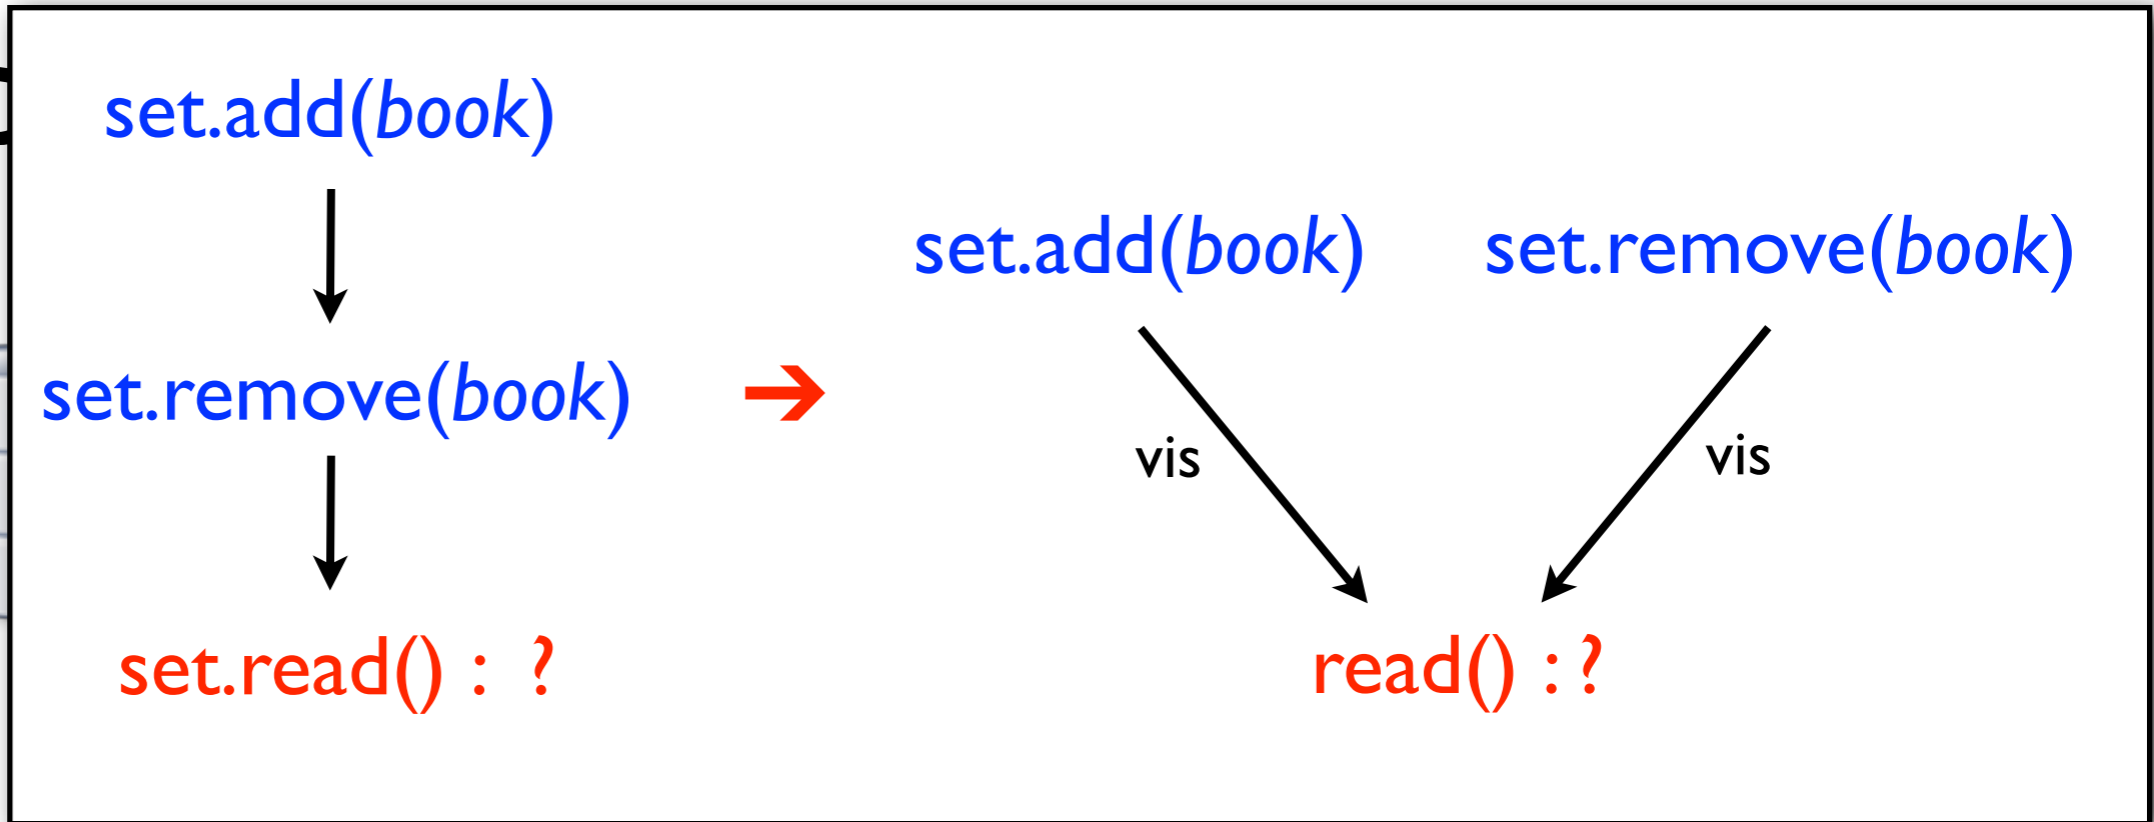
Only updates that have been delivered to the replica performing the operation are important

Replicated data type semantics



Abstract by the **visibility** relation on operations (acyclic, ...)


Rep



set.add(*book*)



Visible?



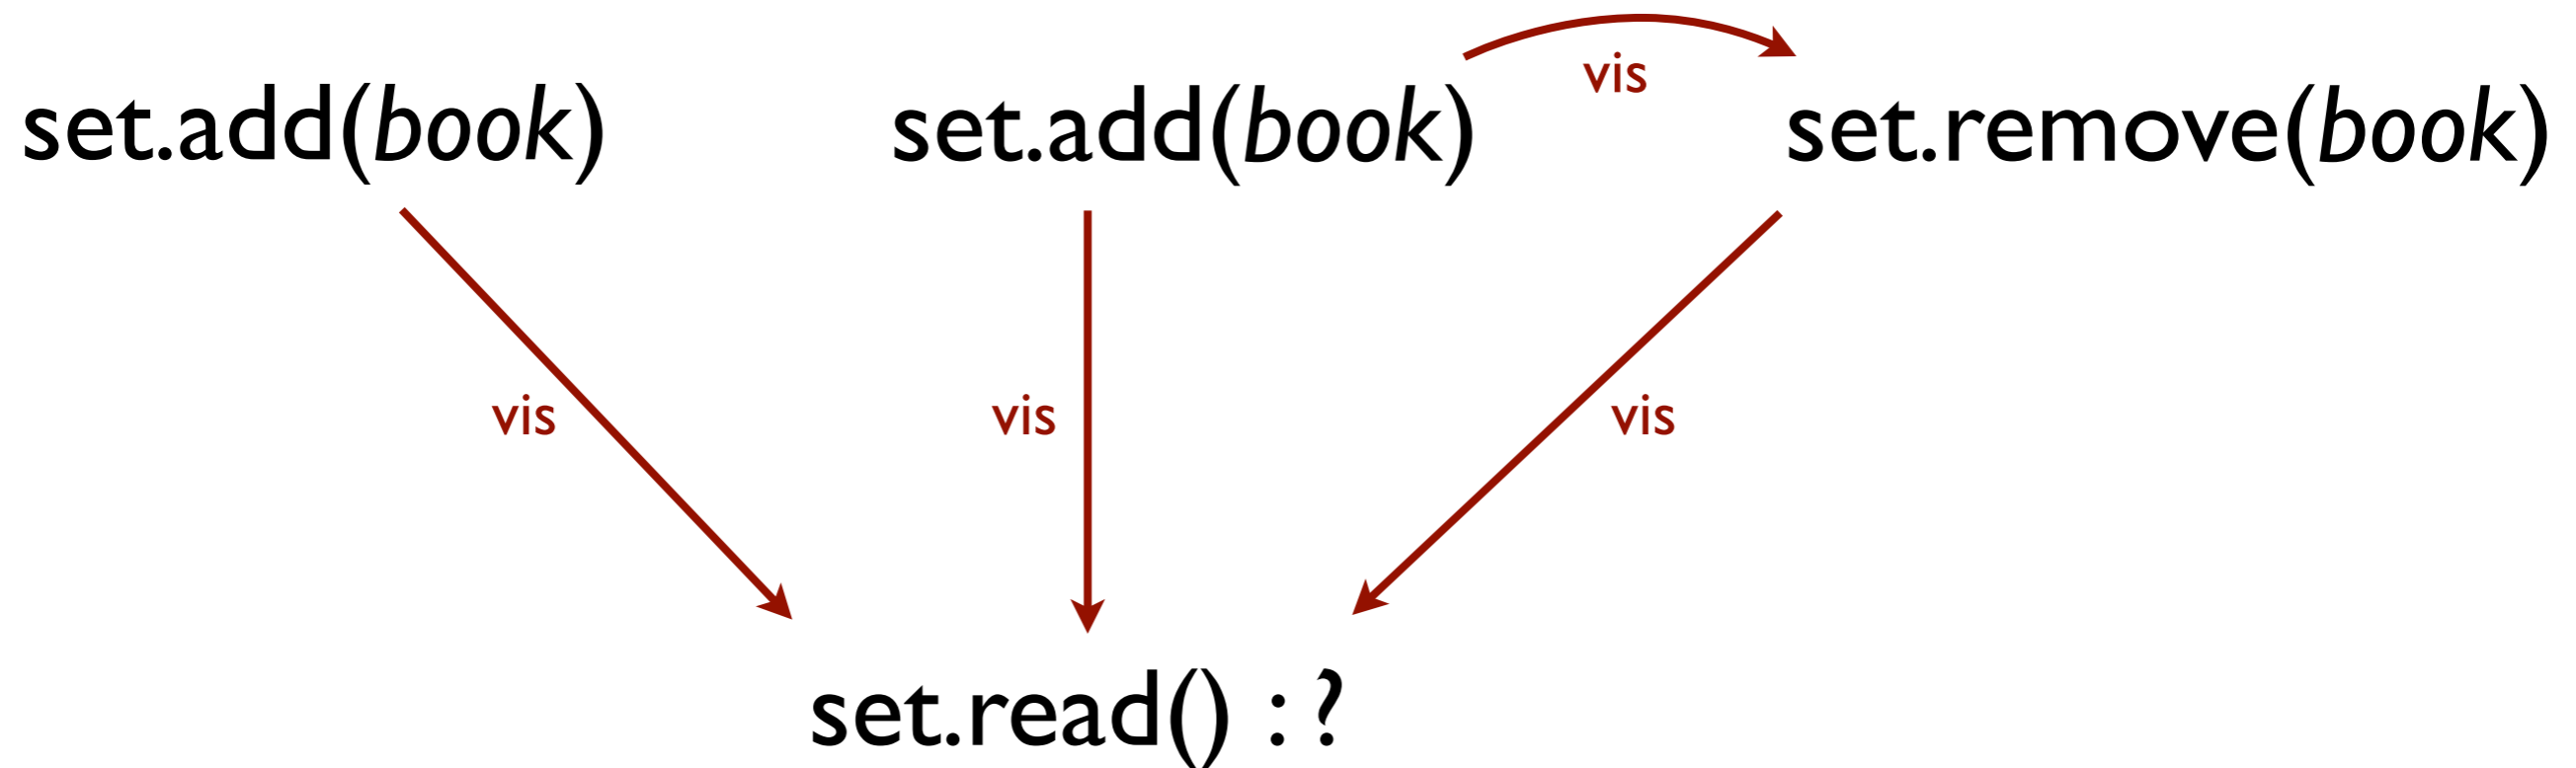
set.read() : ?

Abstract by the **visibility** relation on operations (acyclic, ...)

Replicated data type specification

F: **context**(op) → **return value**(op)

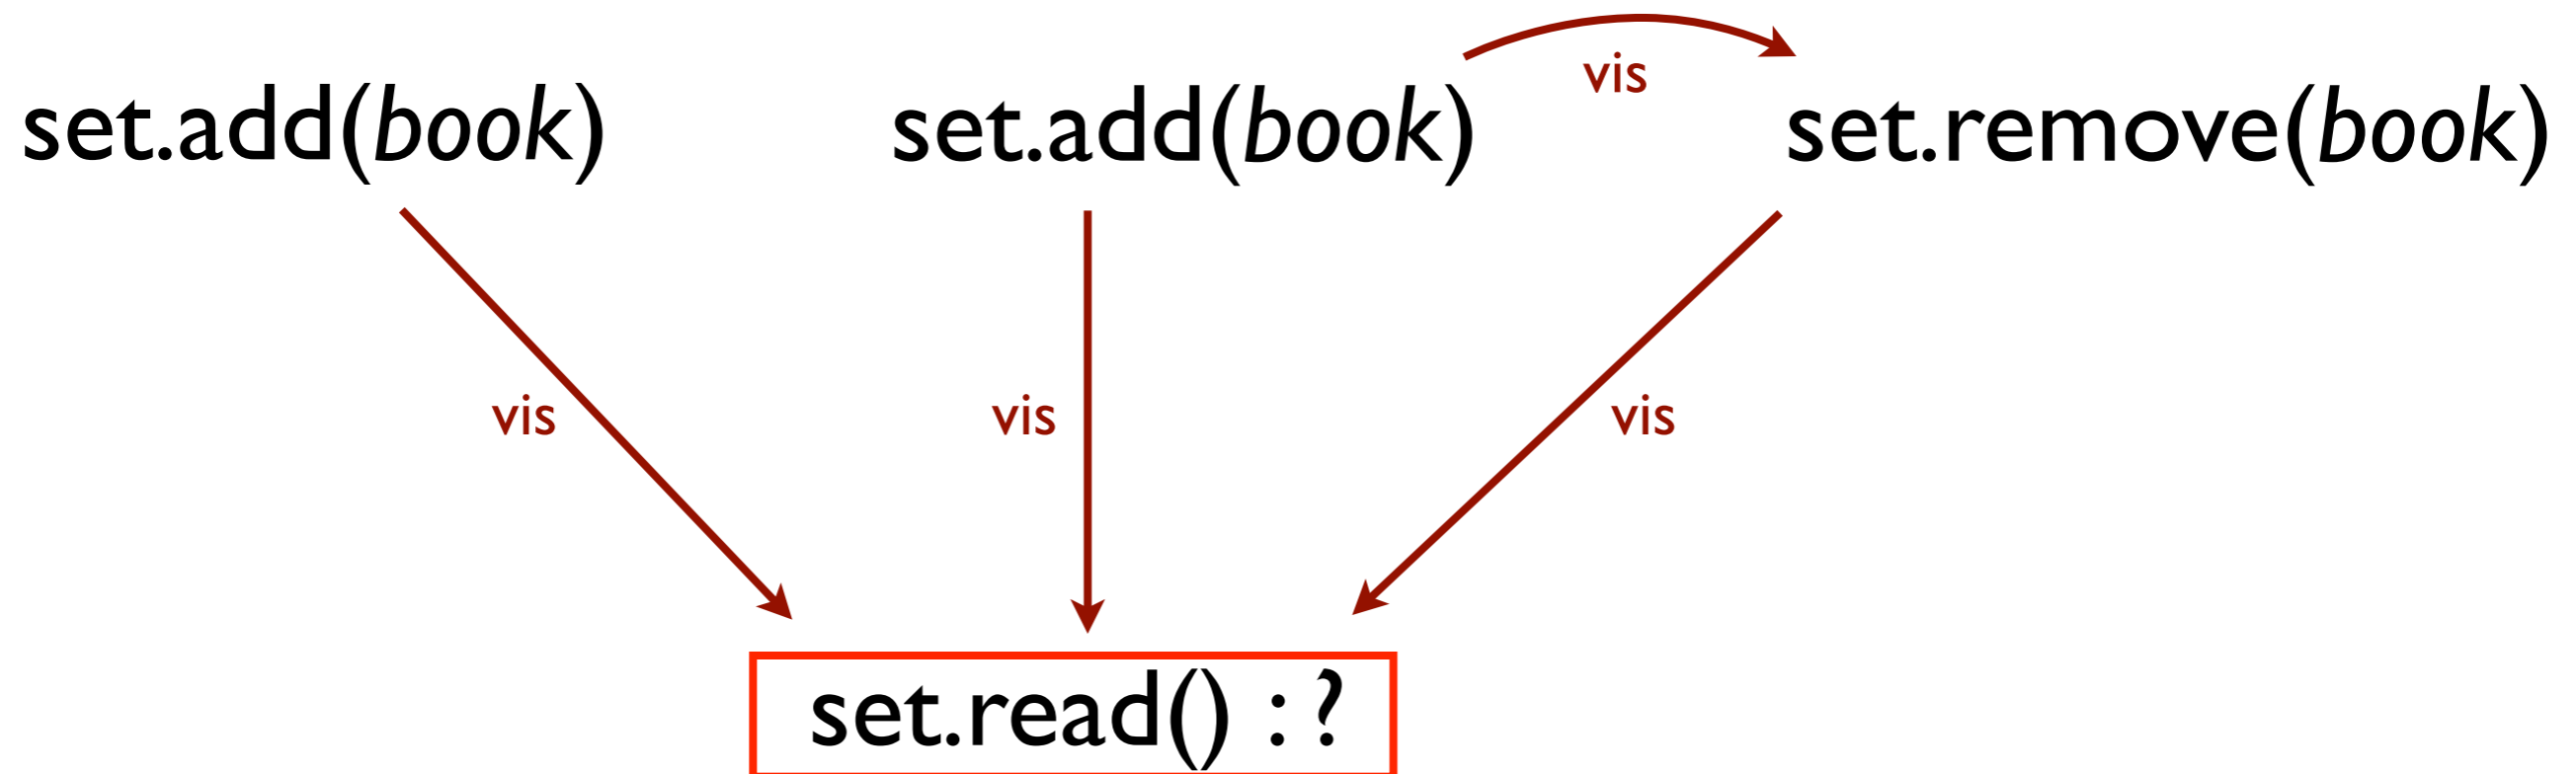
Context: all updates visible to the operation and the visibility relation between them + some other things



Replicated data type specification

F: **context**(op) → **return value**(op)

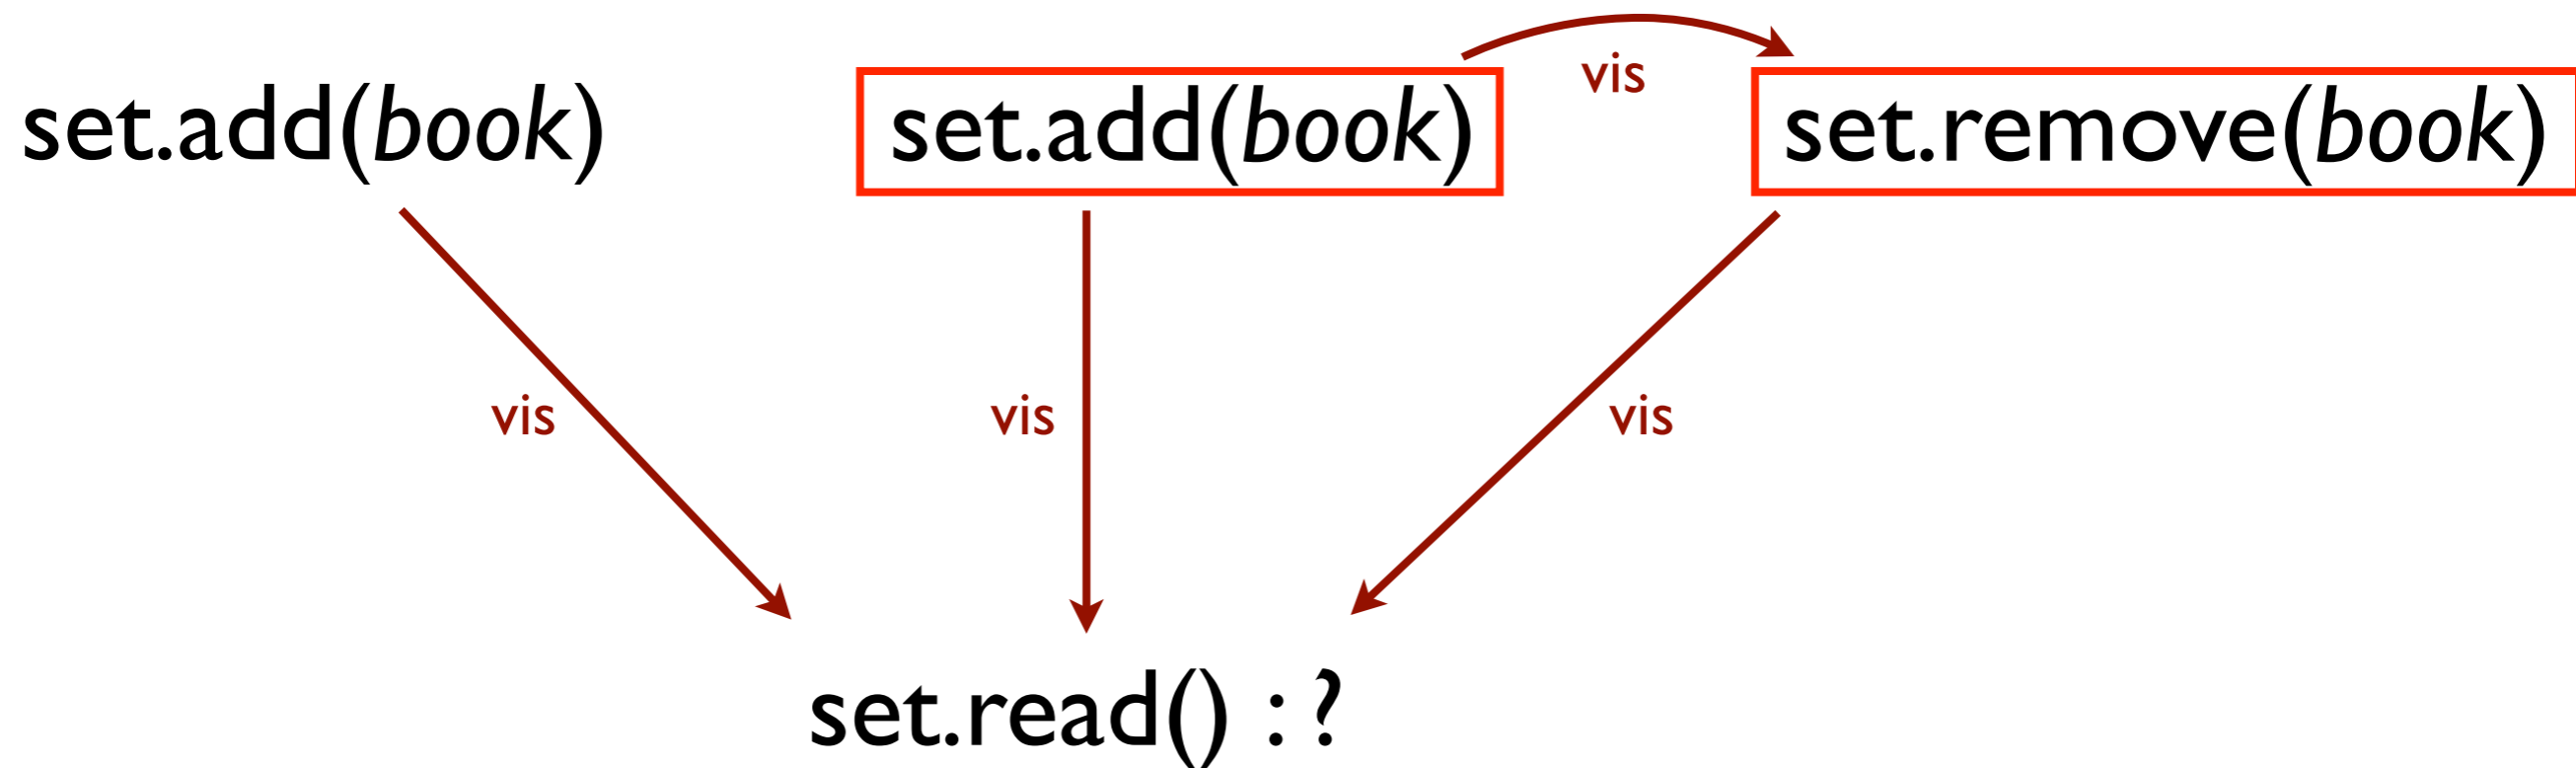
Context: all updates visible to the operation and the visibility relation between them + some other things



Replicated data type specification

F: **context**(op) → **return value**(op)

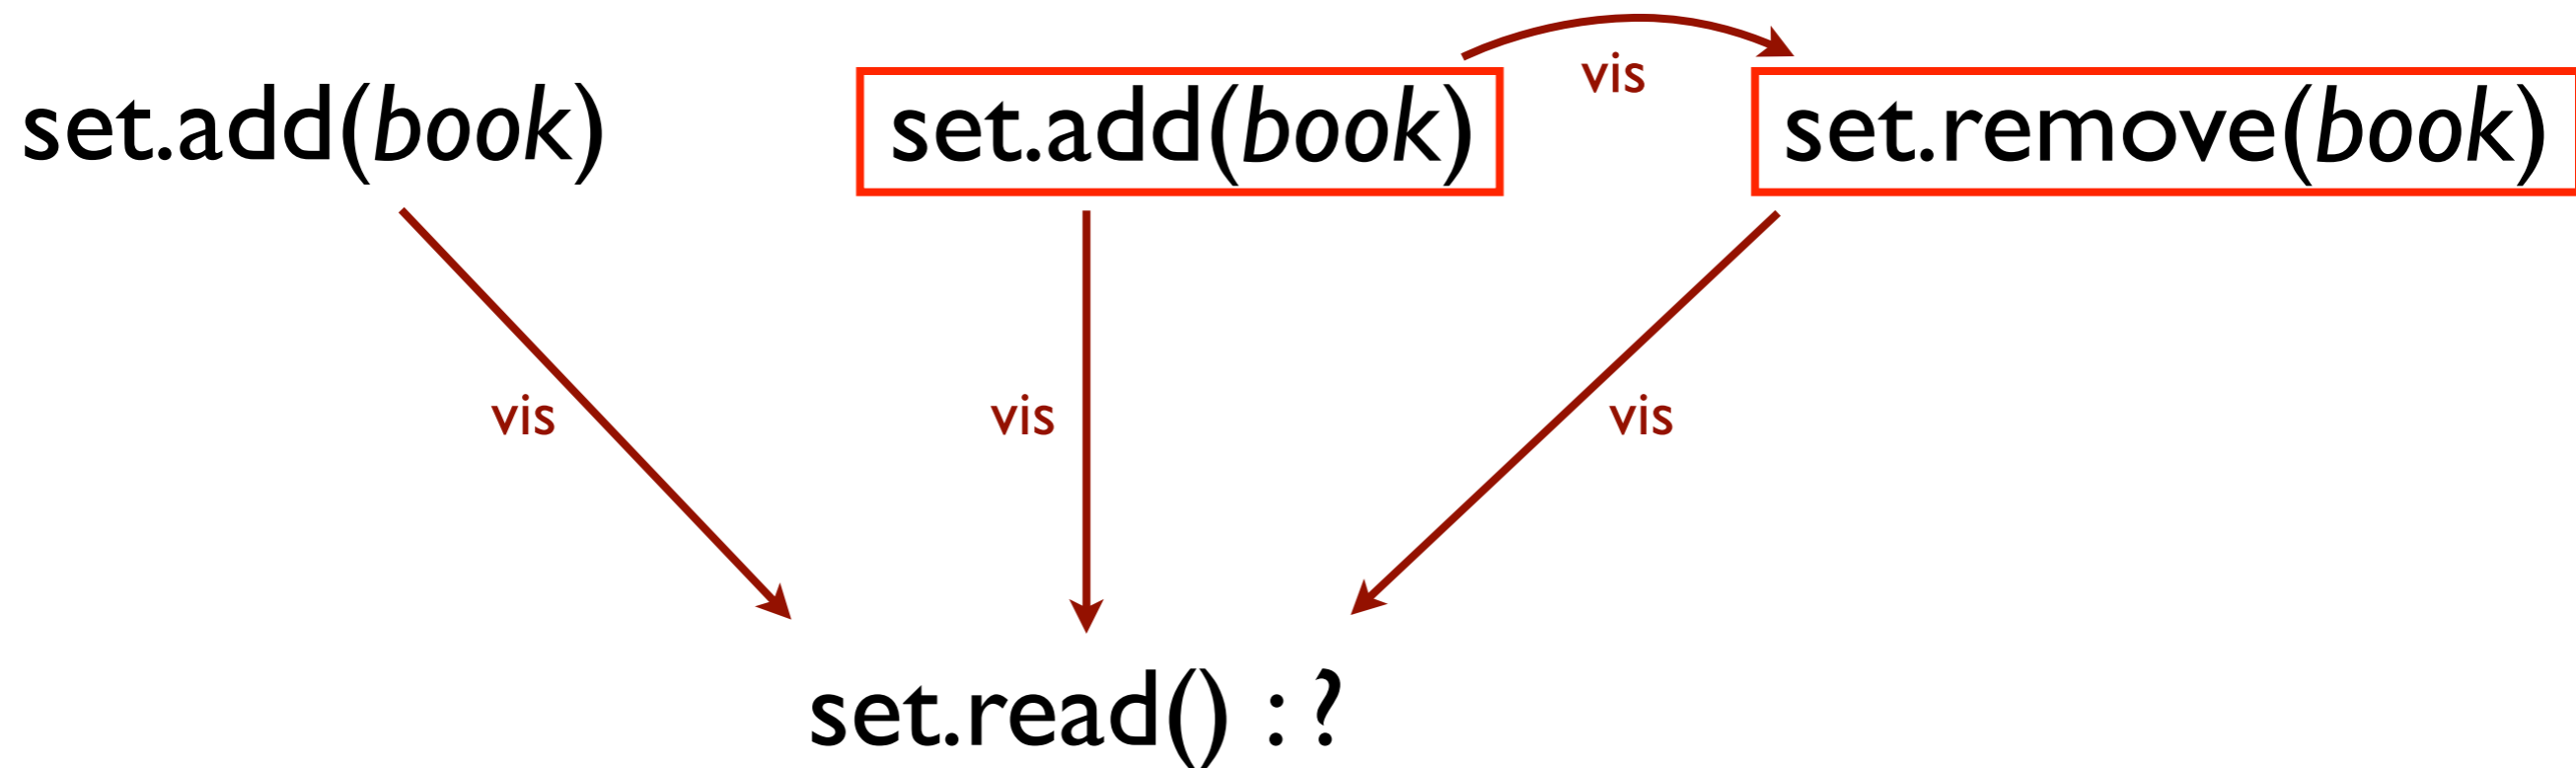
Context: all updates visible to the operation and the visibility relation between them + some other things



Add-wins set

F: **context**(op) → **return value**(op)

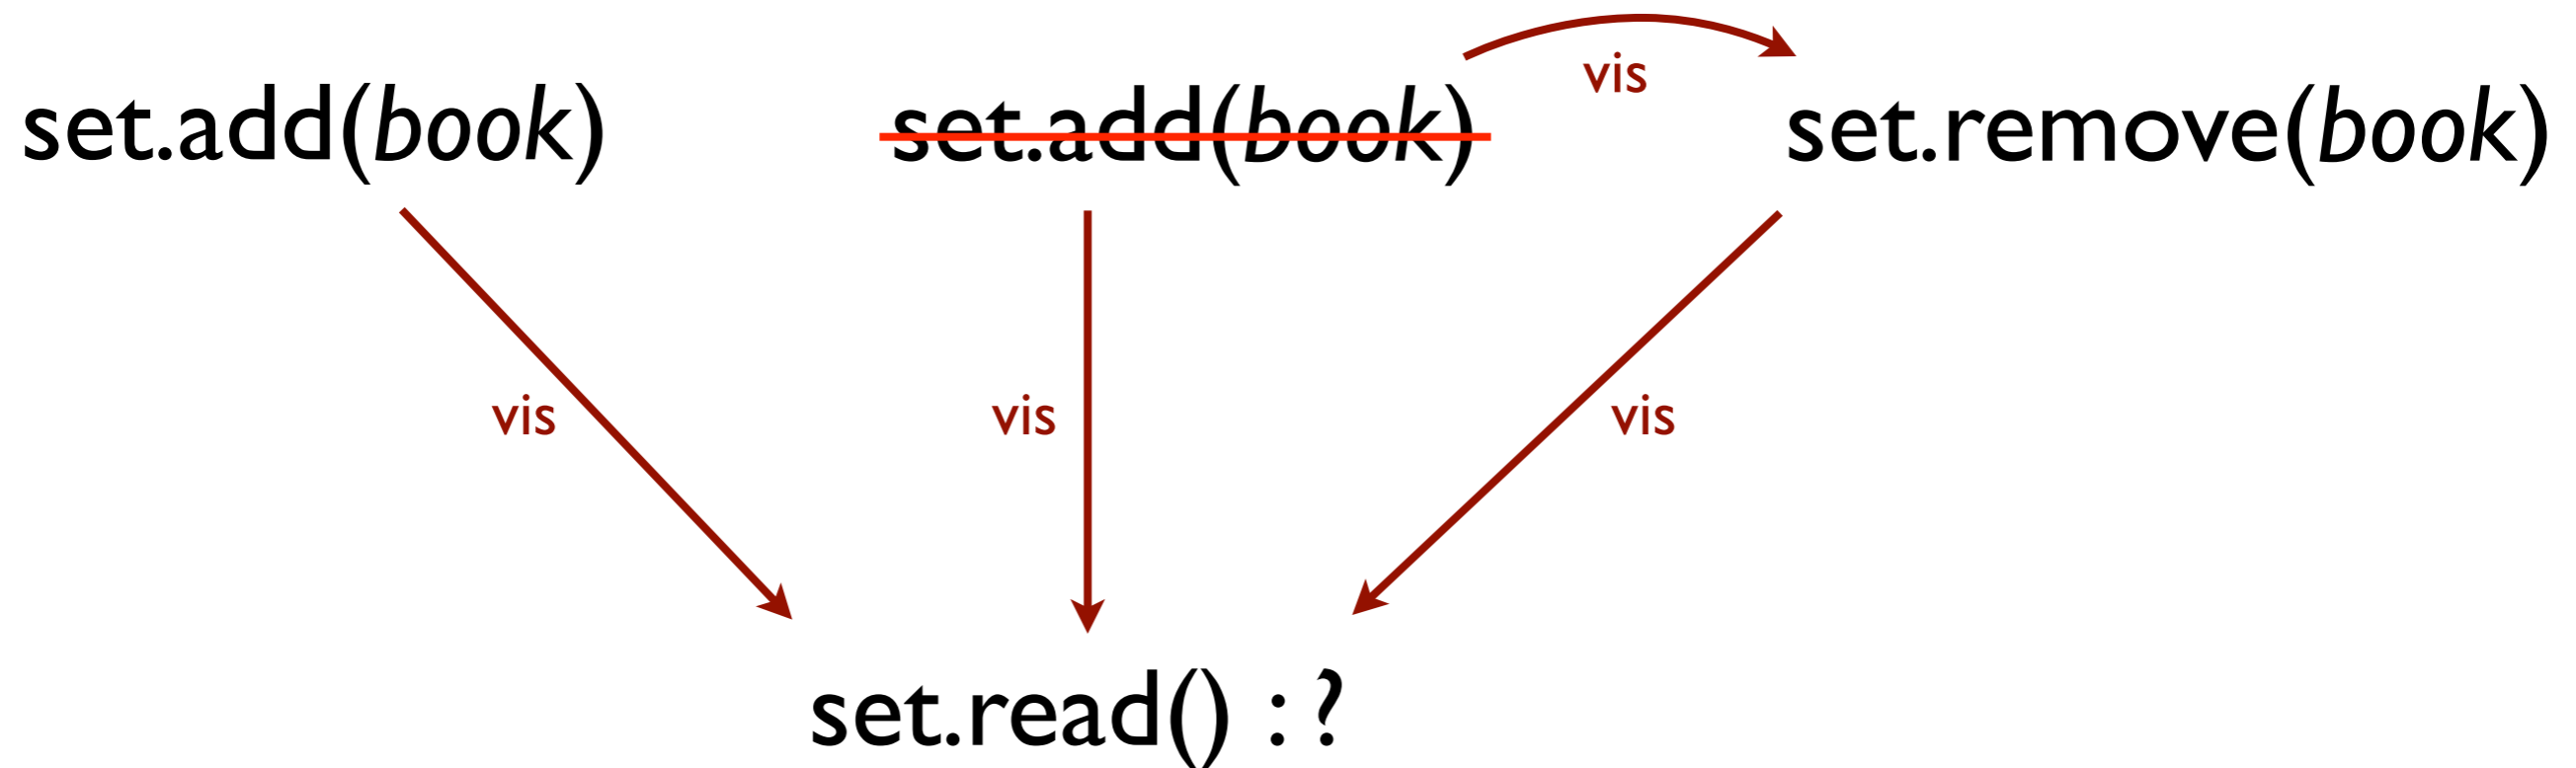
Context: all updates visible to the operation and the visibility relation between them + some other things



Add-wins set

F: **context**(op) → **return value**(op)

Context: all updates visible to the operation and the visibility relation between them + some other things

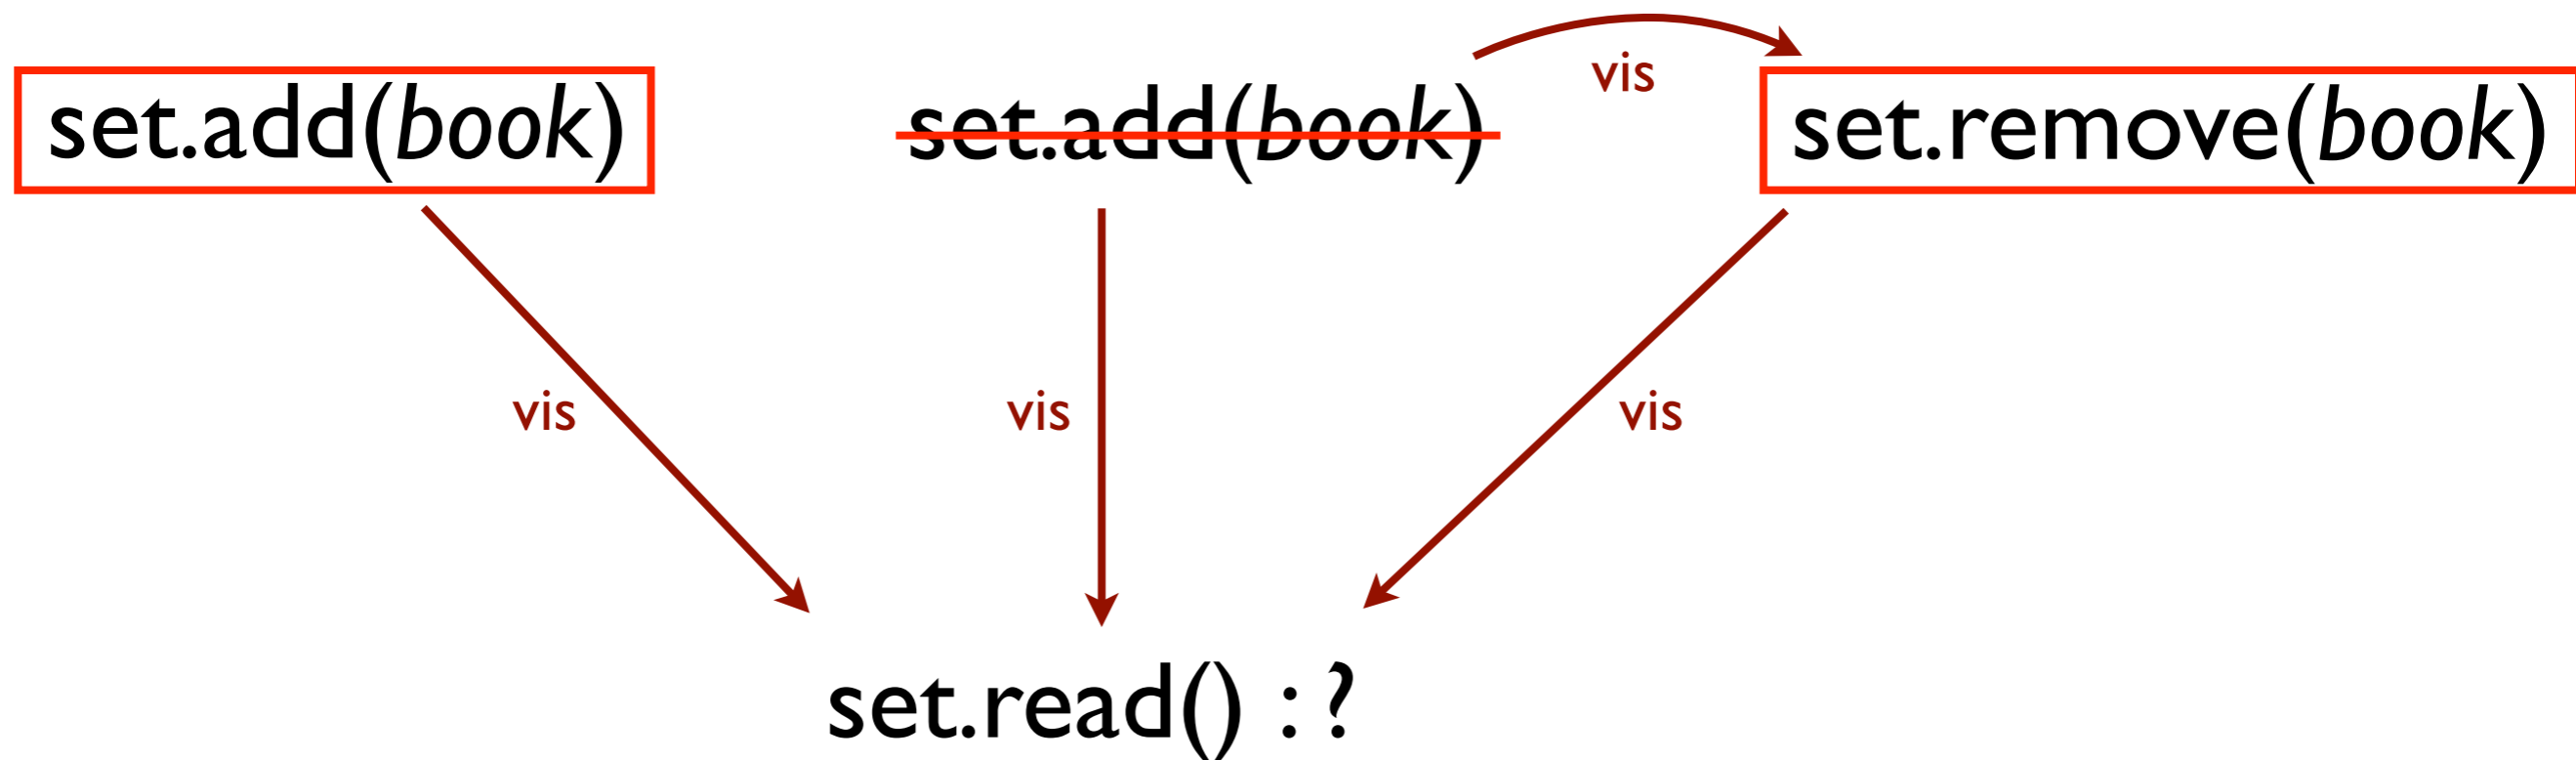


If you saw it, it's not a conflict

Add-wins set

F: **context**(op) → **return value**(op)

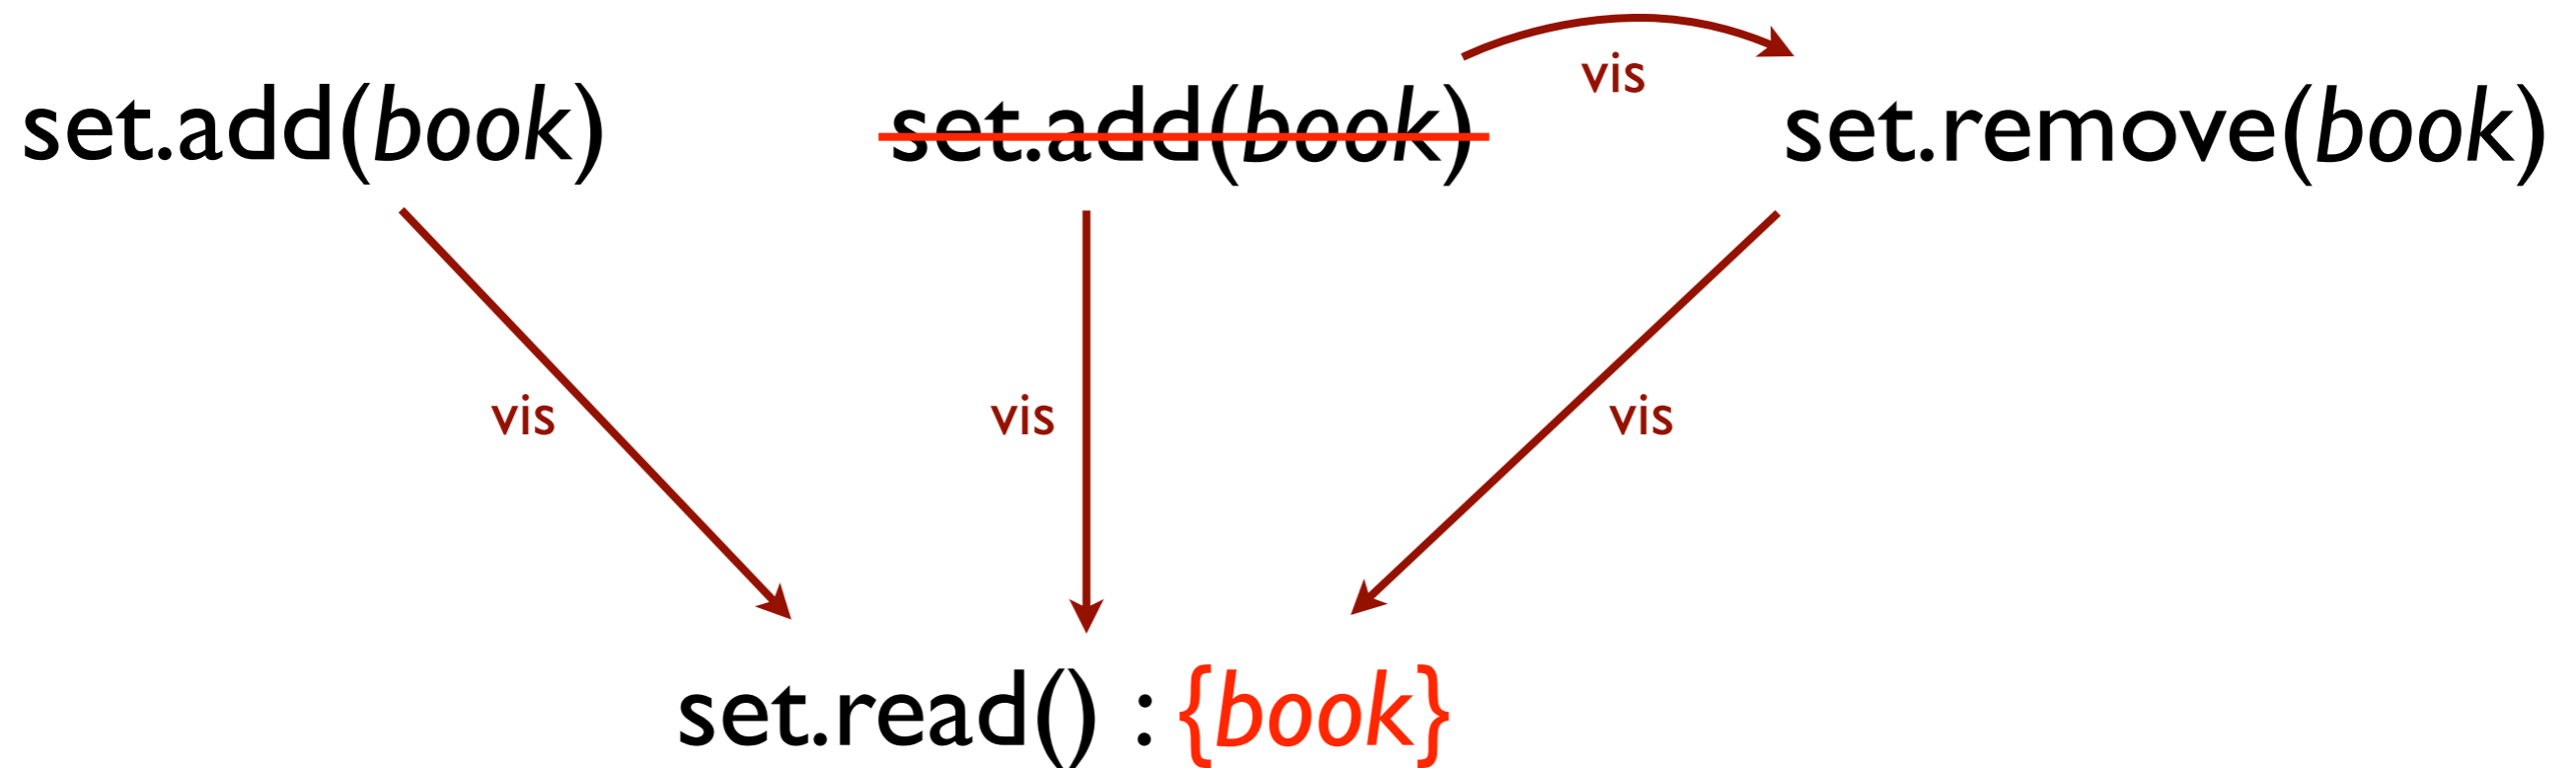
Context: all updates visible to the operation and the visibility relation between them + some other things



Add-wins set

F: **context**(op) → **return value**(op)

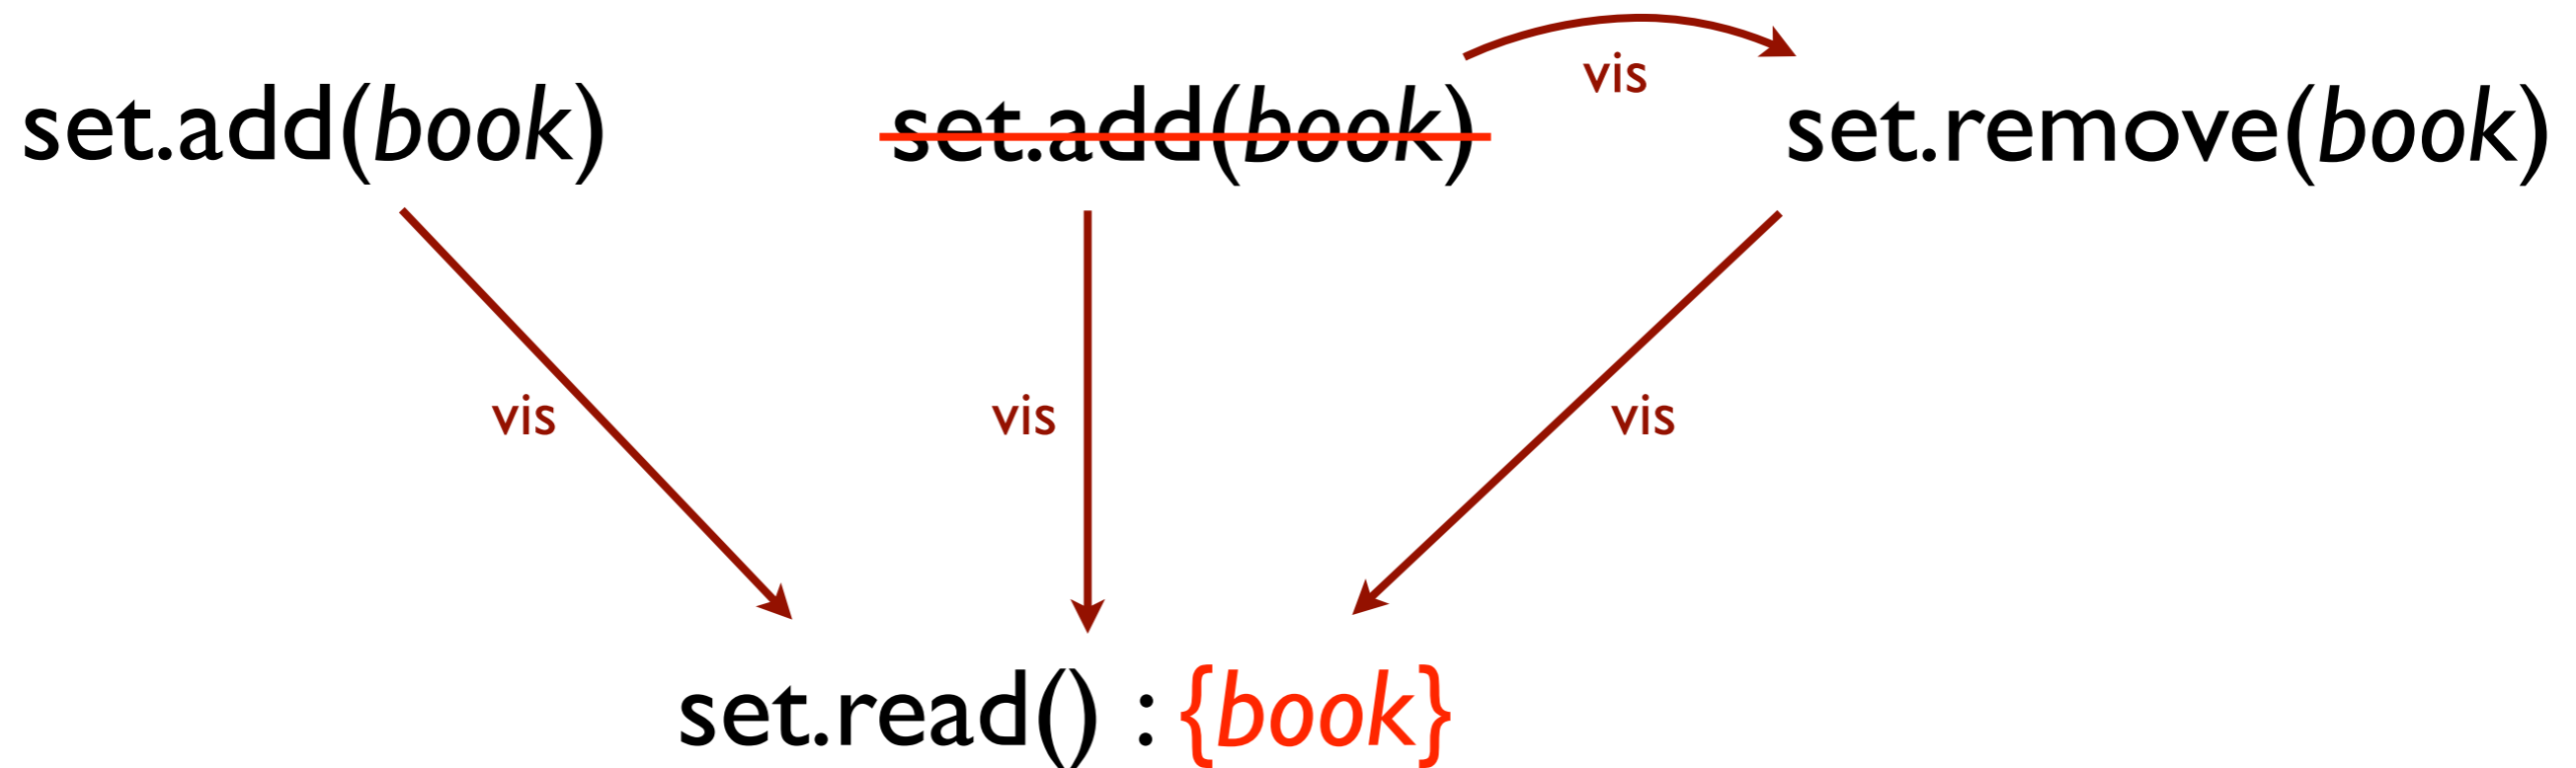
Context: all updates visible to the operation and the visibility relation between them + some other things



Add-wins set

F: **context**(op) → **return value**(op)

Context: all updates visible to the operation and the visibility relation between them + some other things

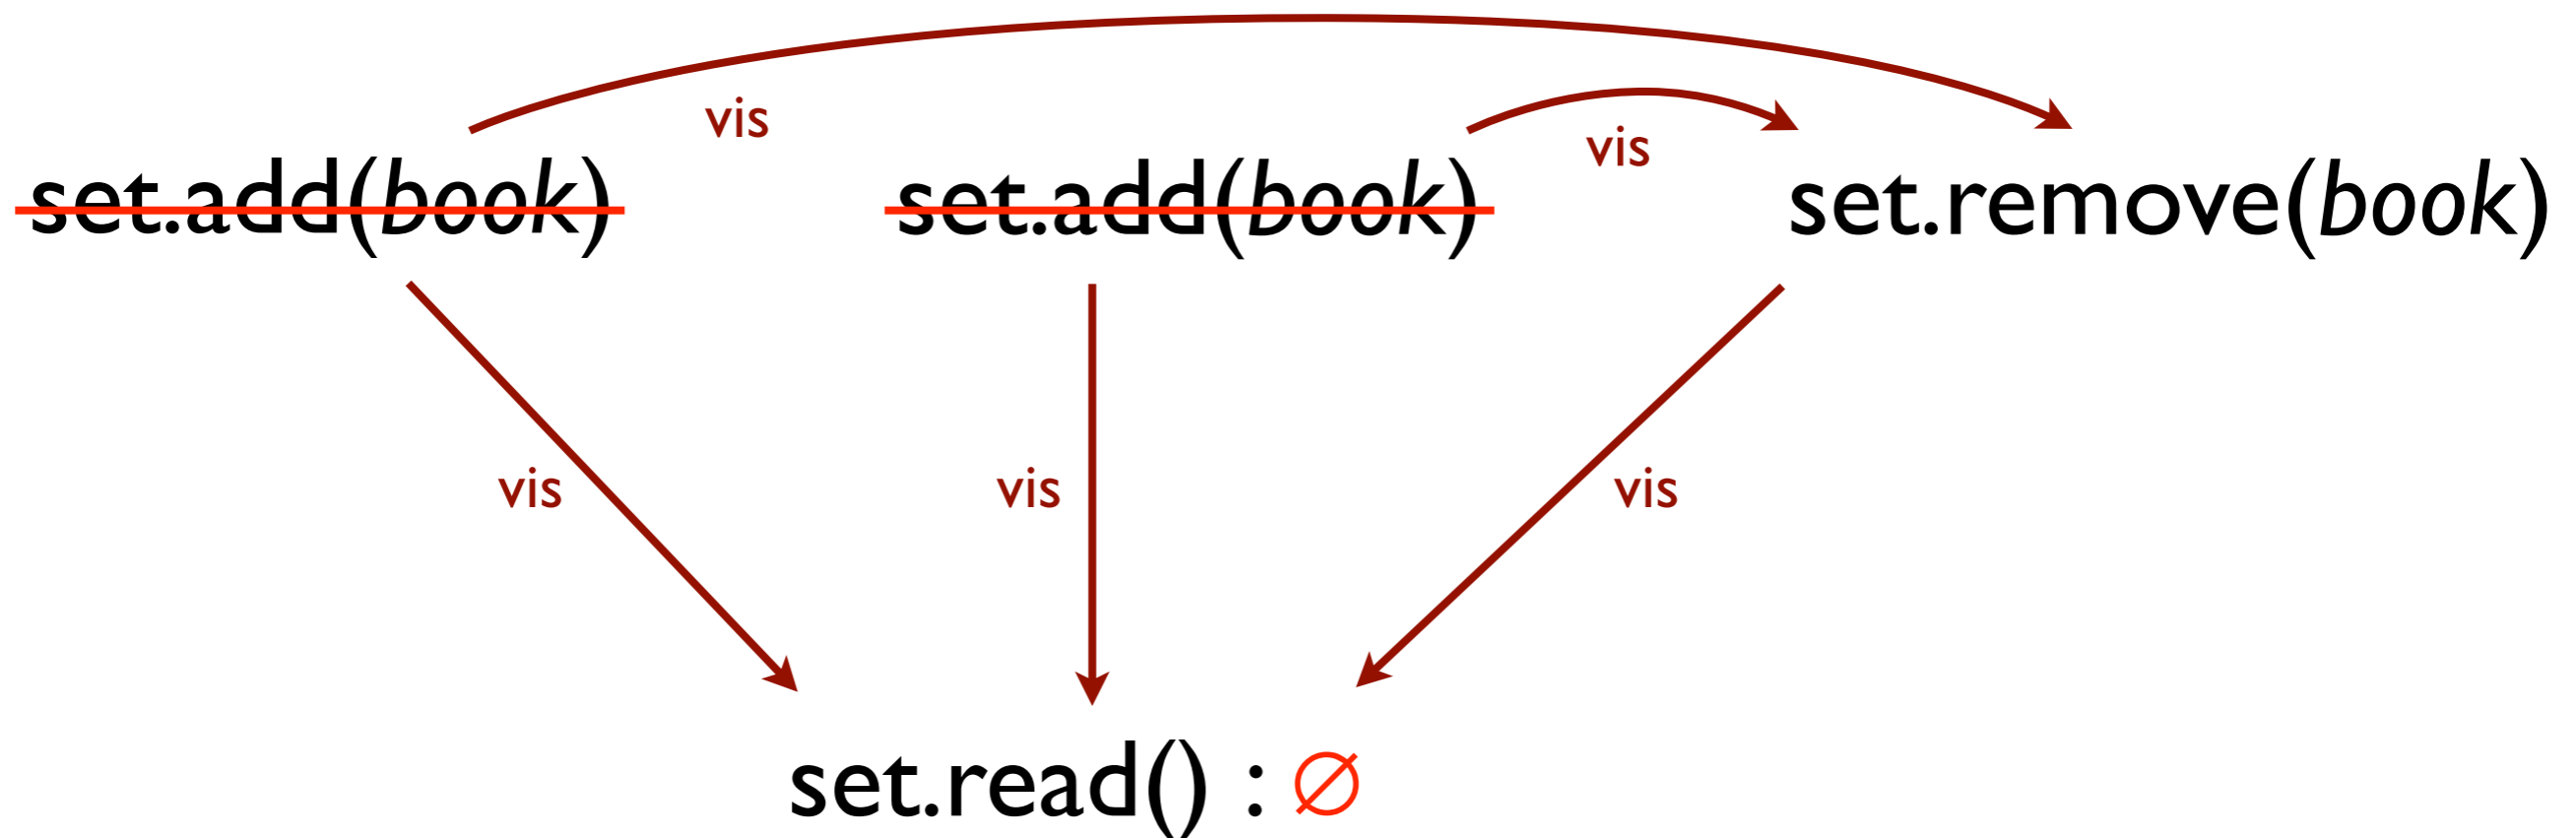


F: cancel all adds seen by a remove

Add-wins set

F: $\text{context}(\text{op}) \rightarrow \text{return value}(\text{op})$

Context: all updates visible to the operation and the visibility relation between them + some other things



F: cancel all adds seen by a remove

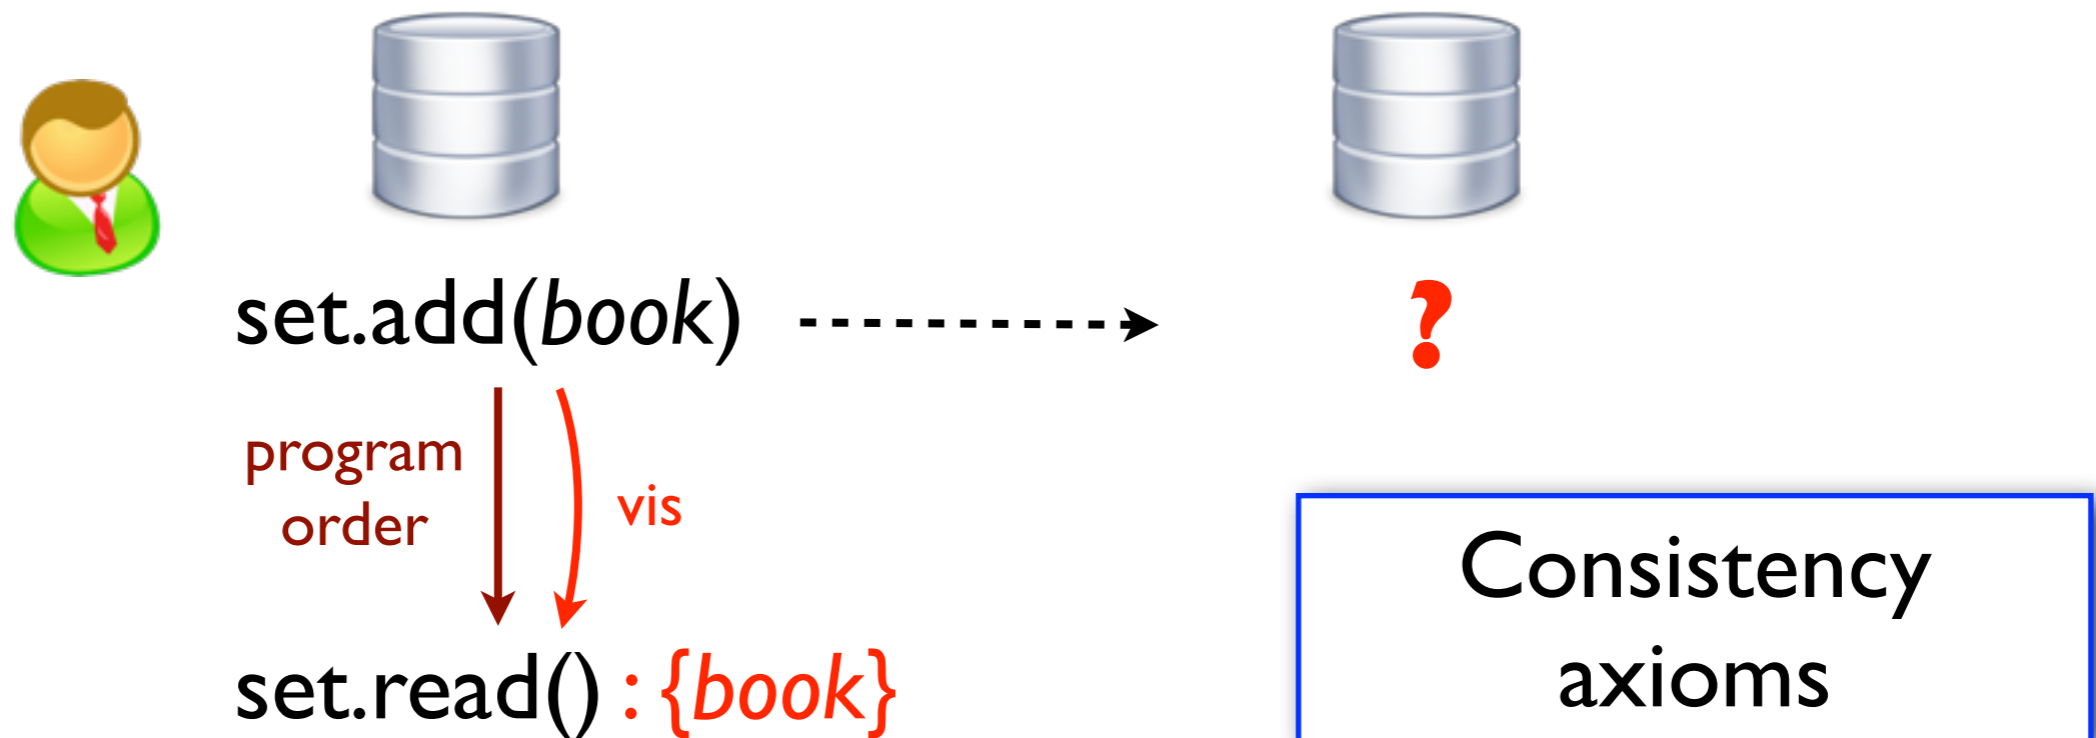
Where does **vis** come from?

Almost arbitrary: little control over when updates are visible to other replicas




Where does **vis** come from?

Almost arbitrary: little control over when updates are visible to other replicas



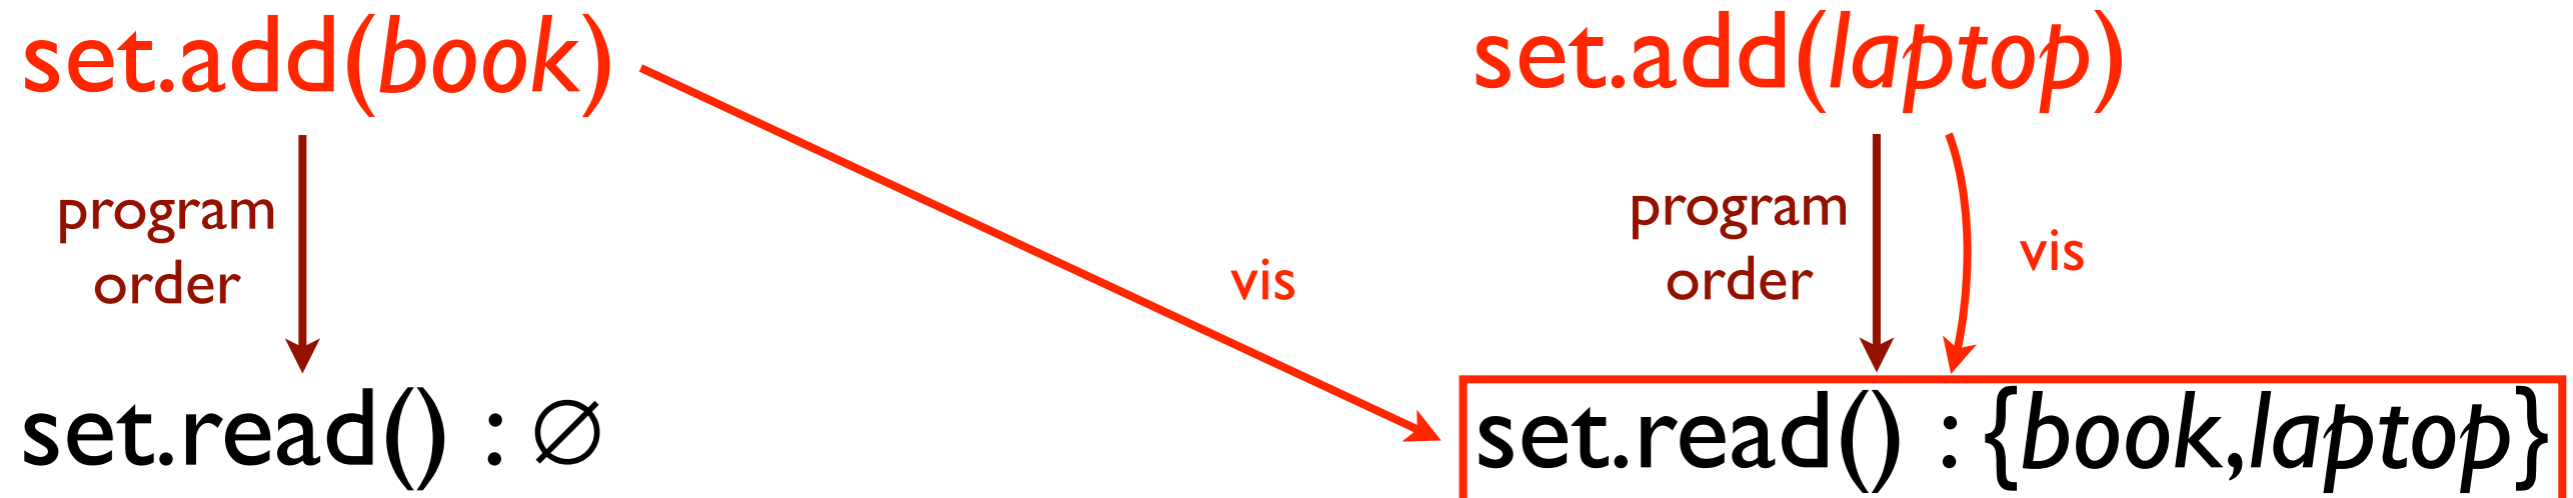
But may guarantee that they don't change unpredictably between operations = anomalies disallowed

Abstract executions: (E, po, vis)



- All operations in a database run, on all objects
- Operations grouped by clients and arranged in program order

Abstract executions: (E, po, vis)




Determines the context of every operation:

$\text{Context}(op)$ = projection onto events visible to op

$\text{return value}(op) = F(\text{Context}(op))$

Abstract executions: (E, po, vis)




Determines the context of every operation:

$\text{Context}(op)$ = projection onto events visible to op


$\text{return value}(op) = F(\text{Context}(op))$

Consistency axioms



- Consistency axioms disallow anomalies by constraining executions
- Read Your Writes: $po \cap \text{same-object} \subseteq vis$
- Principle: strengthen consistency by mandating that more edges be included into vis

Consistency axioms



- Consistency axioms disallow anomalies by constraining executions
- Read Your Writes: $po \cap \text{same-object} \subseteq vis$
- Principle: strengthen consistency by mandating that more edges be included into vis



Basic eventual consistency

Session guarantees

Per-object causal consistency

Causal consistency

Strong consistency

Figure 1. Axioms of eventual consistency

WELL-FORMEDNESS AXIOMS

SOWF: so is the union of transitive, irreflexive and total orders on actions by each session

VISWF: $\forall a, b. a \xrightarrow{\text{vis}} b \implies \text{obj}(a) = \text{obj}(b)$

ARWF: $\forall a, b. a \xrightarrow{\text{ar}} b \implies \text{obj}(a) = \text{obj}(b)$,
ar is transitive and irreflexive, and
 $\text{ar}|_{\text{vis}^{-1}(a)}$ is a total order for all $a \in A$

AUXILIARY RELATIONS

Per-object session order: $\text{soo} = (\text{so} \cap \text{sameobj})$

Per-object causality order: $\text{hbo} = (\text{soo} \cup \text{vis})^+$

Causality order: $\text{hb} = (\text{so} \cup \text{vis})^+$

BASIC EVENTUAL CONSISTENCY AXIOMS

RVAL: $\forall a \in A. \text{rval}(a) = F_{\text{type}(a)}(\text{cone}(a))$

EVENTUAL:

$\forall a \in A. \neg(\exists \text{ infinitely many } b \in A. \text{sameobj}(a, b) \wedge \neg(a \xrightarrow{\text{vis}} b))$

THINAIR: $\text{so} \cup \text{vis}$ is acyclic

SESSION GUARANTEES

RYW (Read Your Writes): $\text{soo} \subseteq \text{vis}$

MR (Monotonic Reads): $(\text{vis}; \text{soo}) \subseteq \text{vis}$

WFRV (Writes Follow Reads in Visibility): $(\text{vis}; \text{soo}^*; \text{vis}) \subseteq \text{vis}$

WFRA (Writes Follow Reads in Arbitration): $(\text{vis}; \text{soo}^*) \subseteq \text{ar}$

MWV (Monotonic Writes in Visibility): $(\text{soo}; \text{vis}) \subseteq \text{vis}$

MWA (Monotonic Writes in Arbitration): $\text{soo} \subseteq \text{ar}$

CAUSALITY AXIOMS

POCV (Per-Object Causal Visibility): $\text{hbo} \subseteq \text{vis}$

POCA (Per-Object Causal Arbitration): $\text{hbo} \subseteq \text{ar}$

COCV (Cross-Object Causal Visibility): $(\text{hb} \cap \text{sameobj}) \subseteq \text{vis}$

COCA (Cross-Object Causal Arbitration): $\text{hb} \cup \text{ar}$ is acyclic



Basic events

- Our specifications similar to weak memory model definitions
- Eventual consistency axioms for registers \approx C/C++ memory model

Session guarantees

Per-object causal consistency

\approx 2011 C/C++ relaxed

Causal consistency

\approx 2011 C/C++ release/acquire

Strong consistency

Per-object session order: $soo = (so \cap sameobj)$

Per-object causality order: $hbo = (soo \cup vis)^+$

Causality order: $hb = (so \cup vis)^+$

BASIC EVENTUAL CONSISTENCY AXIOMS

RVAL: $\forall a \in A. rval(a) = F_{type(a)}(cone(a))$

EVENTUAL:

$\forall a \in A. \neg(\exists \text{ infinitely many } b \in A. sameobj(a, b) \wedge \neg(a \xrightarrow{vis} b))$

THINAIR: $so \cup vis$ is acyclic

SESSION GUARANTEES

RYW (Read Your Writes): $soo \subseteq vis$

MR (Monotonic Reads): $(vis; soo) \subseteq vis$

WFRV (Writes Follow Reads in Visibility): $(vis; soo^*; vis) \subseteq vis$

WFRA (Writes Follow Reads in Arbitration): $(vis; soo^*) \subseteq ar$

MWV (Monotonic Writes in Visibility): $(soo; vis) \subseteq vis$

MWA (Monotonic Writes in Arbitration): $soo \subseteq ar$

CAUSALITY AXIOMS

POCV (Per-Object Causal Visibility): $hbo \subseteq vis$

POCA (Per-Object Causal Arbitration): $hbo \subseteq ar$

COCV (Cross-Object Causal Visibility): $(hb \cap sameobj) \subseteq vis$

COCA (Cross-Object Causal Arbitration): $hb \cup ar$ is acyclic

Specification summary

Conflict resolution policies




Data type spec

Anomalies



Consistency axioms



(E, po, vis)

Specification summary

Conflict resolution policies




Data type spec

Anomalies



Consistency axioms



(E, po, vis)



request₁
response₁
request₂
response₂
...



request₁
response₁
request₂
response₂
...

Specification summary

Conflict resolution policies




Data type spec

Anomalies



Consistency axioms



(E, po, vis)



request₁
response₁
request₂
response₂
...



request₁
response₁
request₂
response₂
...

Events (E, po) allowed iff
 \exists execution (E, po, vis)
satisfying data type specs
and axioms

Specification summary

Conflict resolution policies



Data type spec

Anomalies



Consistency axioms

Quick & dirty proof of correspondence with algorithms
used in systems [TR]



Specification summary

Conflict resolution policies



Data type spec

Anomalies



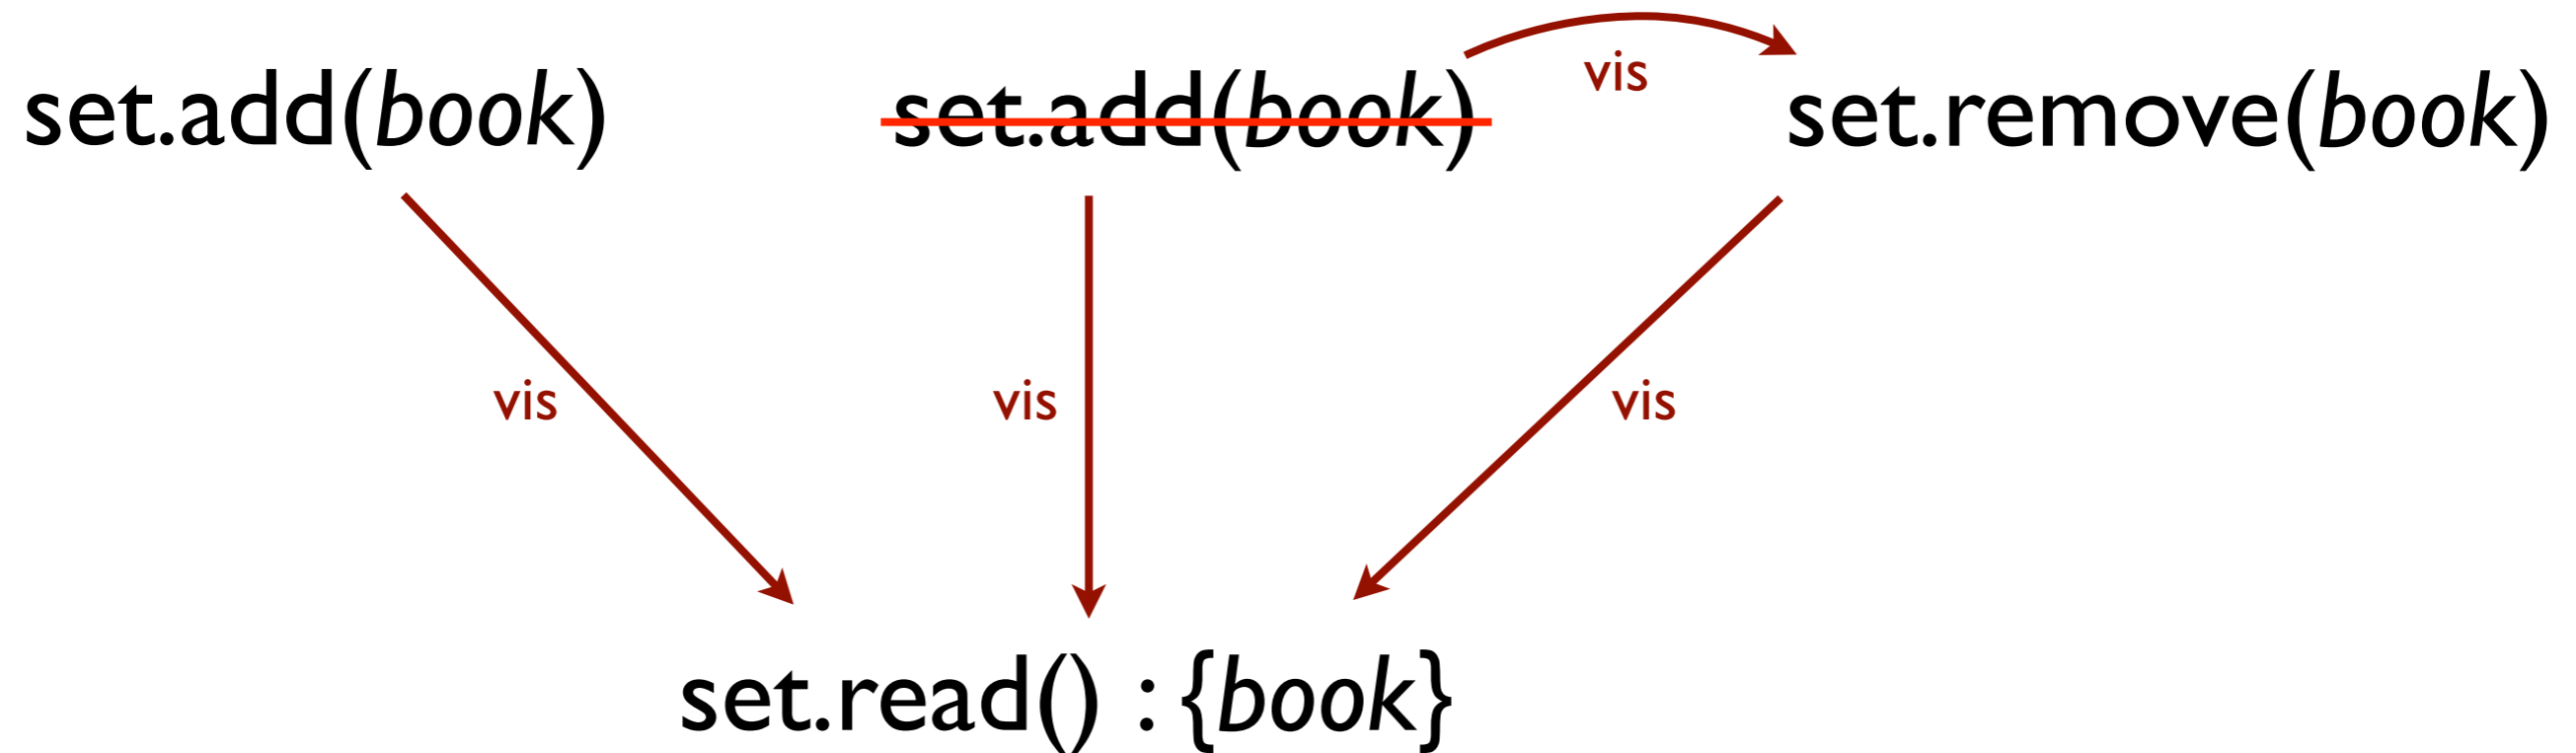
Consistency axioms

Quick & dirty proof of correspondence with algorithms
used in systems [TR]



Verifying data type implementations

Naive add-wins set implementation



Implementation challenge: remove behaves differently wrt different adds of the same element



$S = \{(book, 1)\}$

⋮

`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}$

- Each add creates a new element instance:
(element, unique instance id)



$S = \{(book, 1)\}$

⋮

`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}$

⋮

`set.read() : {book}`

- Each add creates a new element instance: (element, unique instance id)
- Instance ids ignored when reading the set



$S = \{(book, 1)\}$

⋮

`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}$

⋮

`set.read() : {book}`



$S = \{(book, 1)\}$

⋮

`set.remove(book)`

⋮

$S = \emptyset$

- Remove should remove all currently present instances of *book* from *S*



$S = \{(book, 1)\}, T = \emptyset$

⋮

`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

`set.read() : {book}`



$S = \{(book, 1)\}, T = \emptyset$

⋮

`set.remove(book)`

⋮

$S = \emptyset, T = \{(book, 1)\}$

- But maintain the set of **tombstones** T : element instances removed
- Remove moves all instances of *book* in S to T



$S = \{(book, 1)\}, T = \emptyset$

⋮


set.add(book)

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

set.read() : {book}



$S = \{(book, 1)\}, T = \emptyset$

⋮

set.remove(book)

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = ?, T = ?$

State-based implementation:
sends its state snapshot to other replicas



$S = \{(book, 1)\}, T = \emptyset$

⋮


set.add(book)

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

set.read() : {book}



$S = \{(book, 1)\}, T = \emptyset$

⋮

set.remove(book)

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = ?, T = \{(book, 1)\}$

State-based implementation:
sends its state snapshot to other replicas



$S = \{(book, 1)\}, T = \emptyset$

⋮


`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

`set.read() : {book}`



$S = \{(book, 1)\}, T = \emptyset$

⋮

`set.remove(book)`

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = \{\}, T = \{(book, 1)\}$

- Ignore arriving instances that are in T



$S = \{(book, 1)\}, T = \emptyset$

⋮

`set.add(book)`

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

`set.read() : {book}`



$S = \{(book, 1)\}, T = \emptyset$

⋮


`set.remove(book)`

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = \{(book, 2)\}, T = \{(book, 1)\}$



- Ignore arriving instances that are in T
- Add new arriving instances to S



$S = \{(book, 1)\}, T = \emptyset$

⋮


set.add(*book*)

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

set.read() : {*book*}



$S = \{(book, 1)\}, T = \emptyset$

⋮

set.remove(*book*)

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = \{(book, 2)\}, T = \{(book, 1)\}$

⋮

set.read() : {*book*}



$$S = \{(book, 1)\}, T = \emptyset$$

⋮

set.add(book)

⋮

$$S = \{(book, 1), (book, 2)\}, T = \emptyset$$

⋮



$$S = \{(book, 1)\}, T = \emptyset$$

⋮

set.remove(book)

⋮

$$S = \emptyset, T = \{(book, 1)\}$$

⋮

S, T

- State grows linearly with the number of removes
- Realistic implementations represent T compactly: motivation for investigating space optimality
- We prove that space is $\Omega(\log(\text{number of operations}))$



$S = \{(book, 1)\}, T = \emptyset$

⋮


set.add(book)

⋮

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

⋮

set.read() : {book}



$S = \{(book, 1)\}, T = \emptyset$

⋮

set.remove(book)

⋮

$S = \emptyset, T = \{(book, 1)\}$

⋮

$S = \{(book, 2)\}, T = \{(book, 1)\}$

⋮

set.read() : {book}

Impl \models F



~~set.add(book)~~

$S = \{(book, 1)\}, T = \emptyset$

$S = \{(book, 1)\}, T = \emptyset$

set.add(book)

set.remove(book)

$S = \{(book, 1), (book, 2)\}, T = \emptyset$

$S = \emptyset, T = \{(book, 1)\}$

set.read() : {book}

$S = \{(book, 2)\}, T = \{(book, 1)\}$

set.read() : {book}

Impl \models F

vis

vis

vis

S, T

vis

vis

Data type correctness: $\text{Impl} \models F$

- \forall concrete execution of the implementation with any sequence of client operations
- \exists corresponding abstract execution satisfying data type specifications and consistency axioms

Data type correctness: $\text{Impl} \models F$

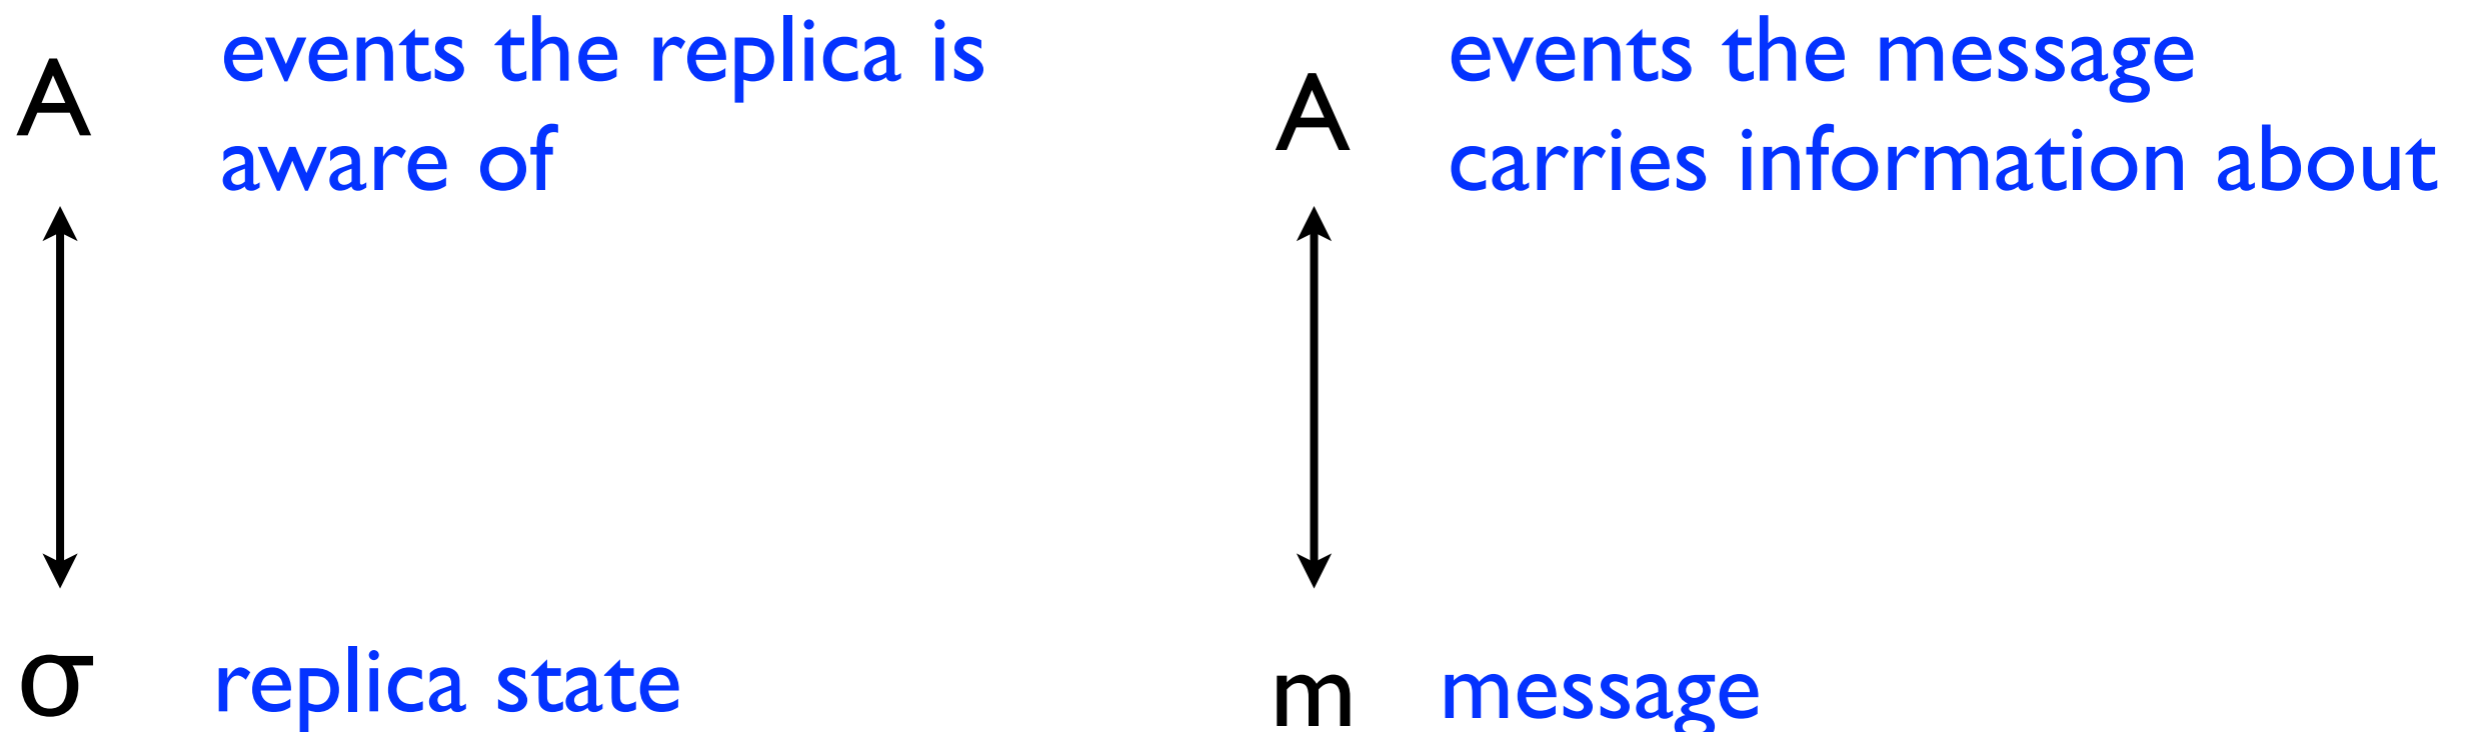
- \forall concrete execution of the implementation with any sequence of client operations
- \exists corresponding abstract execution satisfying **data type specifications** and consistency axioms

Data type correctness: $\text{Impl} \models F$

- \forall concrete execution of the implementation with any sequence of client operations
 - \exists corresponding abstract execution satisfying **data type specifications** and consistency axioms
-
- Requires reasoning about all replicas and interactions between them
 - Want to **modularise** reasoning: construct the abstract execution from separate system configuration components

Replication-aware simulations

- Generalise simulation relations for abstract data types to replicated case
- Replica state or message associated with an abstract execution part describing events that led to it



Simulation for add-wins set

(S,T)

Set S: $\{(book,2), (laptop,3)\}$

Tombstones T: $\{(book,1)\}$



Simulation for add-wins set

(S,T)

A

Set S: $\{(book,2), (laptop,3)\}$

Tombstones T: $\{(book,1)\}$



Simulation for add-wins set

(S,T)

A

Set S: $\{(book,2), (laptop,3)\}$

Tombstones T: $\{(book,1)\}$

\leftrightarrow

$add(book)^1$ $add(laptop)^3$

$add(book)^2$

$(elt, id) \in S \cup T \quad \leftrightarrow \quad add(elt)^{id} \in A$

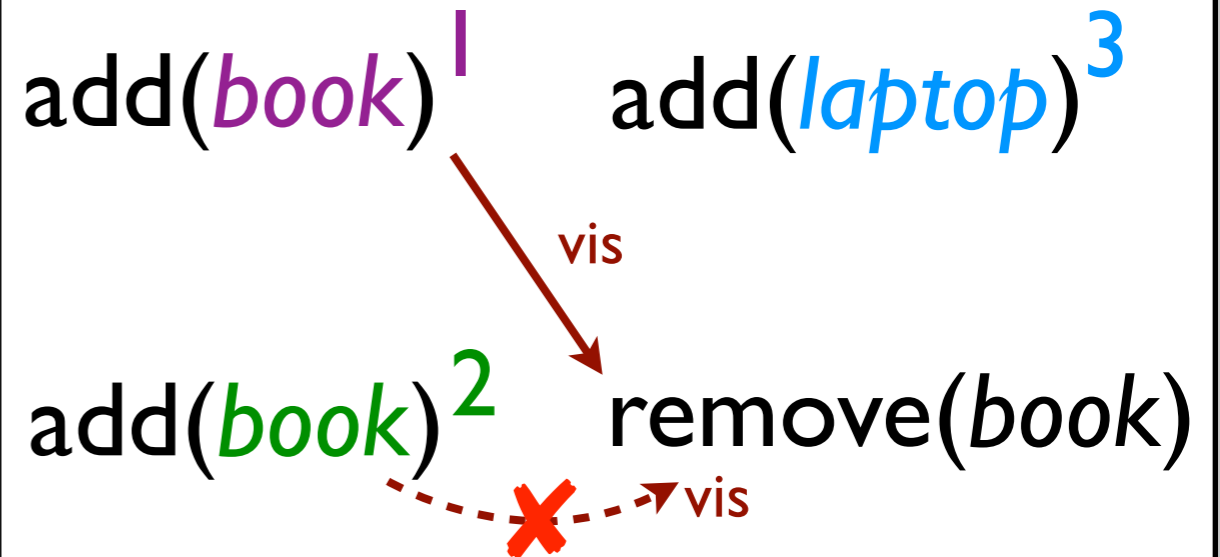
Simulation for add-wins set

(S,T)

Set S: $\{(book,2), (laptop,3)\}$

Tombstones T: $\{(book,1)\}$

A



\leftrightarrow

$(elt, id) \in S \cup T \quad \leftrightarrow \quad add(elt)^{id} \in A$

$(elt, id) \in T \quad \rightarrow \quad remove(elt) \xleftarrow{vis} add(elt)^{id}$

$(elt, id) \in S \quad \rightarrow \quad \neg add(elt)^{id} \xrightarrow{vis} remove(elt)$

Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification

Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification

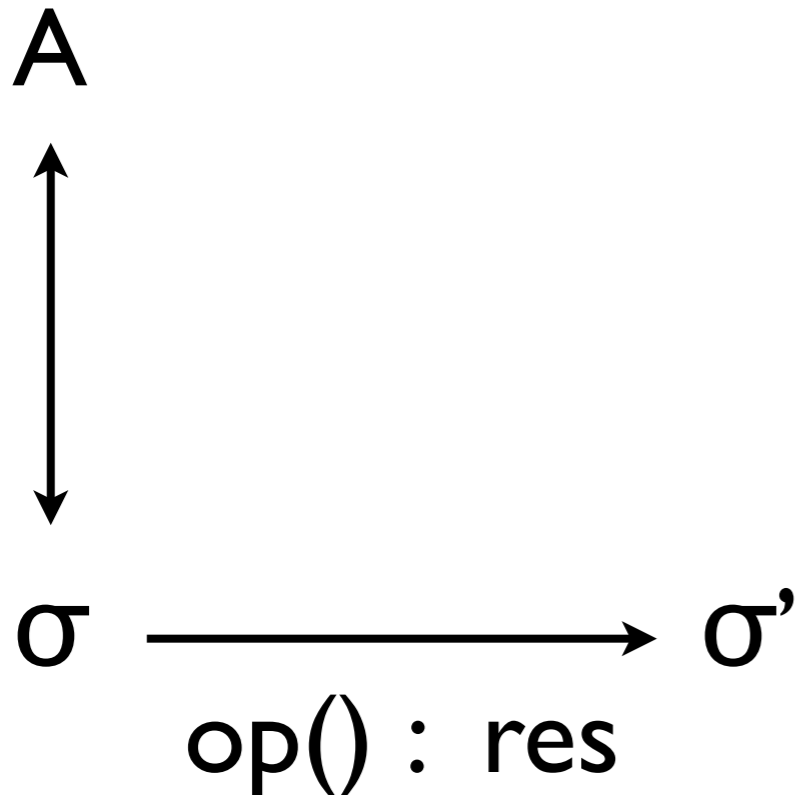
Executing an operation:

$$\sigma \xrightarrow{\text{op}() : \text{res}} \sigma'$$

Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification


Executing an operation:



Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification


Executing an operation:



Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification

Executing an operation:




And check
 $\text{res} = F(\text{Context}_{A'}(\text{op}))$

Proof obligations

- Relations are preserved during a system run
- Relations imply that the abstract execution satisfies the data type specification

Executing an operation:




And check
 $\text{res} = F(\text{Context}_{A'}(\text{op}))$

Modular: considers the state of a single replica


Receiving a message

$(\sigma, m) \xrightarrow{\text{receive}(m)} \sigma'$


Receiving a message



Receiving a message




Receiving a message



Good news: modular - consider the state of a single replica and a message


Bad news: modularity leads to incompleteness - loses required global information

Source of incompleteness



- A and B parts of the same abstract execution \rightarrow can be correlated by some invariants
 - ▶ *Visibility can't contradict on events common to A and B*
 - ▶ *Union of visibility relations in A and B itself a well-formed visibility relation \rightarrow acyclic*
- Simulation relations per-component \rightarrow don't give this

Solution: 2-stage verification



- ▶ Fix a class of data types implementations with similar messaging behaviour

State-based: propagate information by sending full replica state

- ▶ Prove key global invariants non-modularly
- ▶ Unpleasant, but done once for the class

Solution: 2-stage verification

- 1 ▶ Fix a class of data types implementations with similar messaging behaviour
 - State-based: propagate information by sending full replica state*
 - ▶ Prove key global invariants non-modularly
 - ▶ Unpleasant, but done once for the class
- 2 ▶ For any implementation within the class
 - ▶ Verify it modularly using replication-aware simulations while assuming the global invariants

Solution: 2-s

Technical details in the paper

- 1
 - ▶ Fix a class of data types implementations with similar messaging behaviour
 - State-based: propagate information by sending full replica state*
 - ▶ Prove key global invariants non-modularly
 - ▶ Unpleasant, but done once for the class
- 2
 - ▶ For any implementation within the class
 - ▶ Verify it modularly using replication-aware simulations while assuming the global invariants

Summary

- First techniques for reasoning about eventual consistency and replicated data types
 - ▶ Specifying the intended semantics
 - ▶ Verifying replicated data type correctness
- Only the first step
 - ▶ Replicated data types only one system component
 - ▶ More work needed even for them: list data type, used for collaborative editing (Office Online, Google Docs)

Programming languages/verification vs distributed systems

- Put eventually consistent distributed systems onto the PL/verification agenda
- Usual paradigm: developing verification techniques
- But also: helping systems researchers design architectures and programming interfaces
 - ▶ Tricky to figure out semantics & implementation for complex interfaces: multiple consistency levels, transactions

Common ground: weak memory models

- Lot of recent work on weak memory
- Opportunity: apply weak memory technology to distributed systems

Common ground: weak memory models

- Lot of recent work on weak memory
- Opportunity: apply weak memory technology to distributed systems
- Processor and language models have very little known motivation
- Distributed systems are different: implemented algorithms motivate models

