



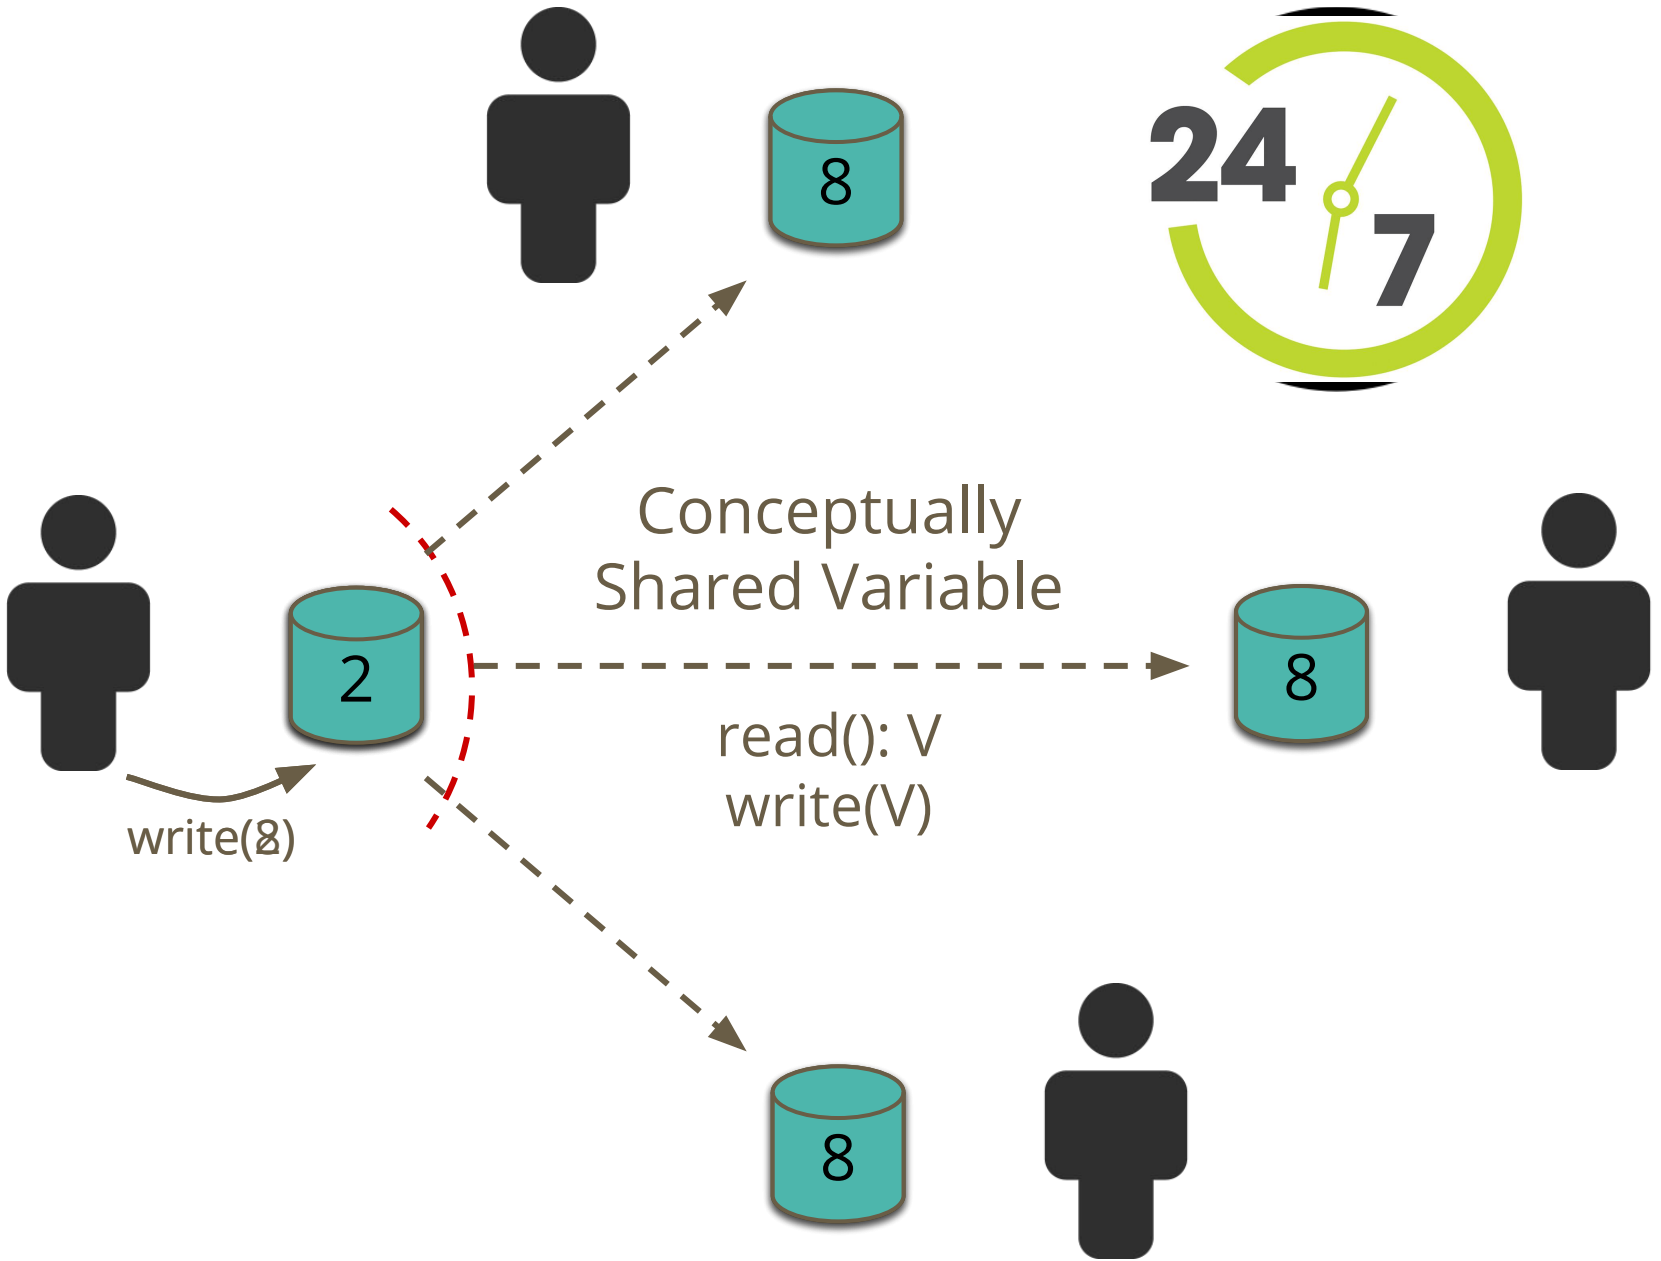
CScript: Supporting Consistent and Available Replication

(Deliverables: D2.1.2, D2.2.1)

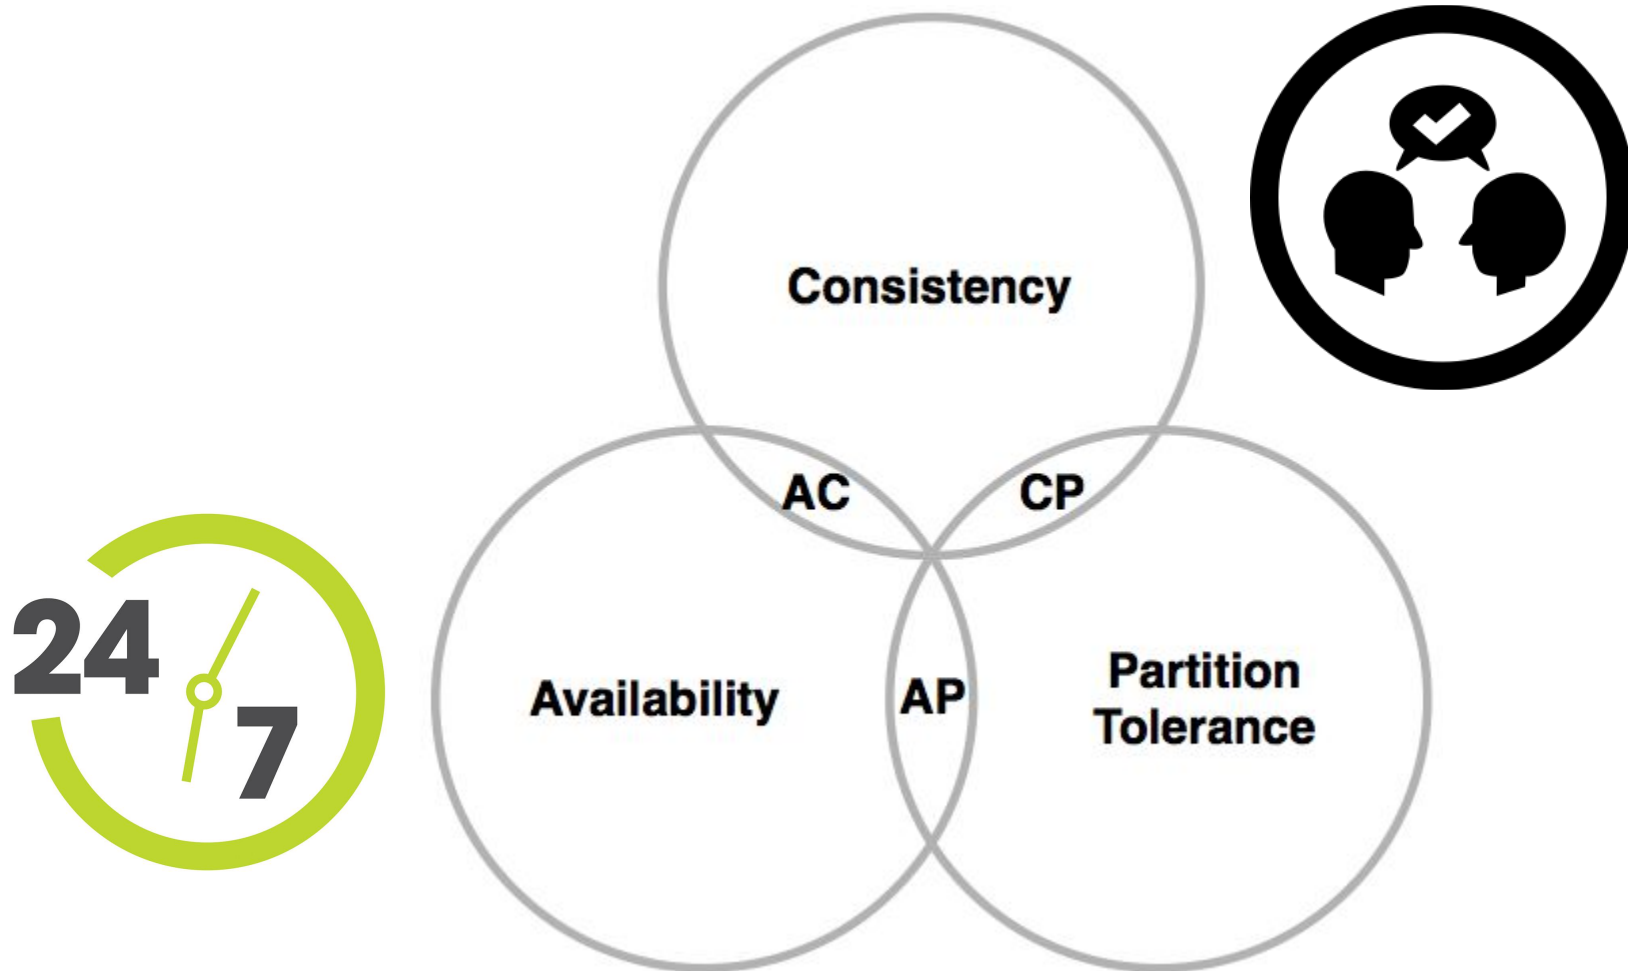
Kevin De Porre
kdeporre@vub.be







CAP Theorem




Brewer, E. A. (2000). Towards Robust Distributed Systems. In *Proceedings of the Nineteenth Annual ACM Symposium on Principles of Distributed Computing* (p. 7). New York, NY, USA. doi: 10.1145/343477.343502

WP 2: End-to-end Eventual Consistency



Problem Statement



Lack of language support

Language	AP	CP
Argus	X	Actions
Distributed Smalltalk	Copy	RMI
Emerald	Move	RMI
AmbientTalk	Isolates	Far References
Geo	Caching	Linearizability & Cache Coherence Protocol
Lasp	CRDTs	X


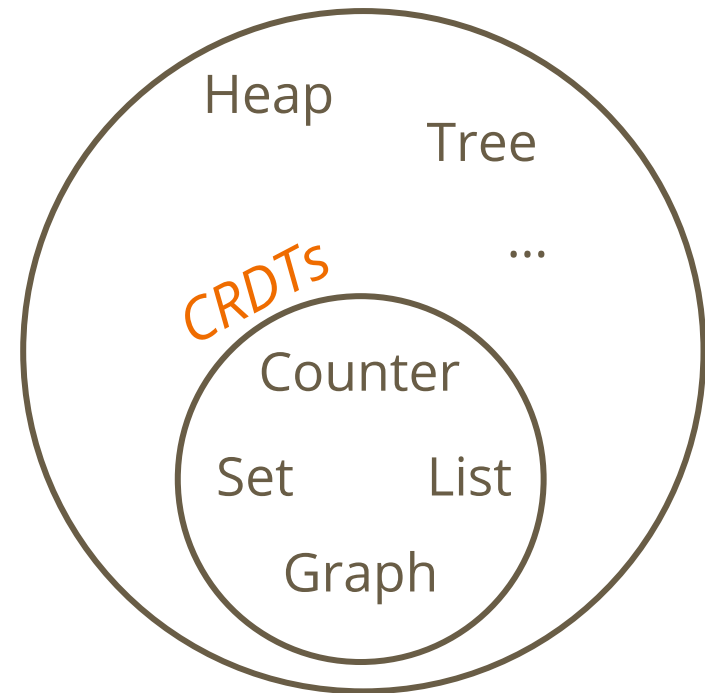



Conflict-free Replicated Data Types (CRDTs)

- Commutativity design
 - Avoids conflicts
 - But.. strong restriction


✗ Limited set of CRDTs

✗ Not generally applicable



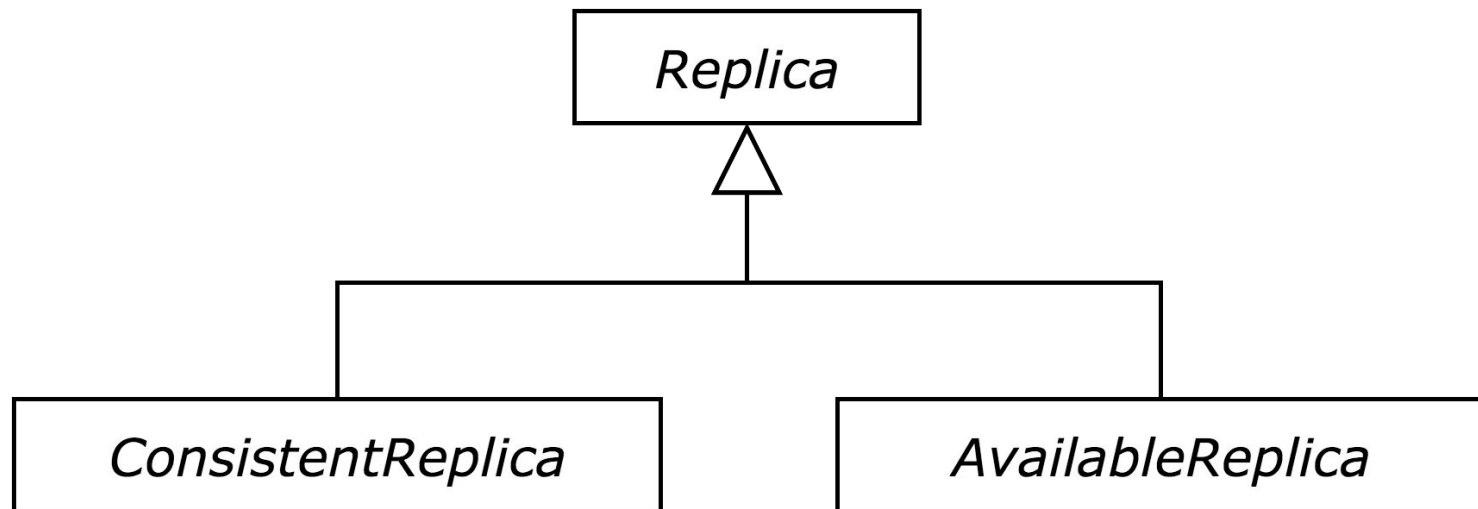


Lack of language support



CScript

- Extends JS
- Replicas
 - Consistent
 - Available





Use Case: Grocery App



Grocery Lists

New

Name...

Author...


Create

Lists

- Shopping List by Kevin +
- Mangos (3/5) +
- Lasagna (0/2) +
- Pizza (0/1) +


Each list item has a shopping bag icon and a close icon (X).





Grocery App

```
service GroceryService {  
  rep list = new GroceryList();  
  rep inventory = new Inventory();  
  
  constructor(name, author)  
    this.name = name;  
    this.author = author;  
}  
  
add(item, qty) {  
  return this.list.add(item, qty);  
}  
  
delete (itemName) {  
  return this.list.delete(itemName);  
}  
  
buy(itemName, qty) { /* ... */ }  
}
```








Grocery App: Inventory

```
class Inventory {
  constructor() {
    this.stock = new Map();
  }

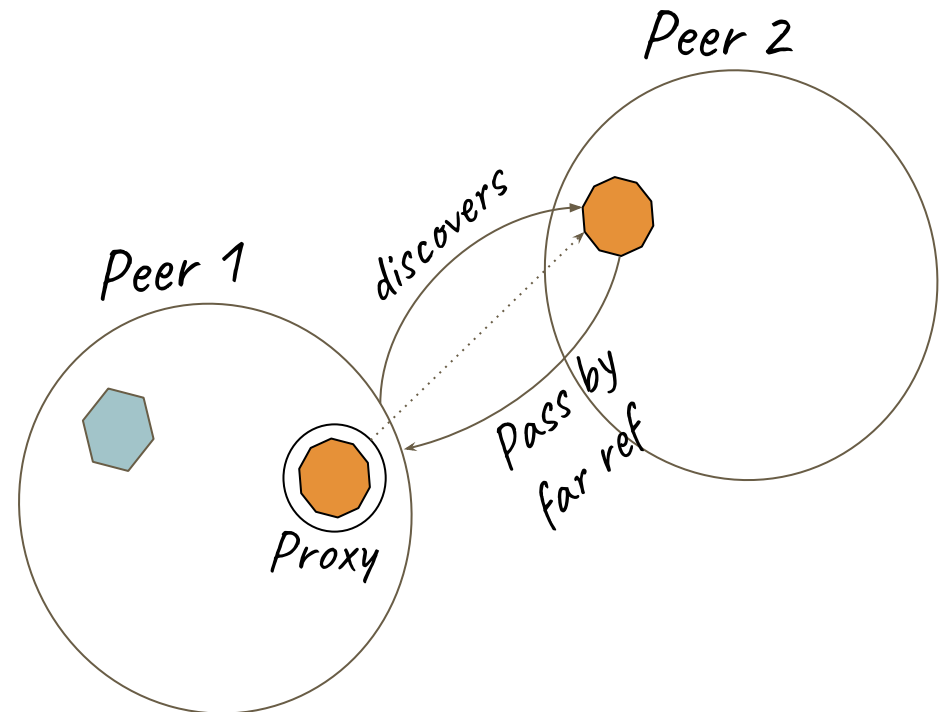
  approve(itemName, stockQuantity, buyingQuantity) {
    if (buyingQuantity <= 0)
      return false;


    const trueStock = this.stock.getOrElse(itemName, 0);
    if (trueStock === stockQuantity) {
      this.stock.set(itemName, trueStock + buyingQuantity);
      return true;
    }
    else {
      return false;
    }
  }
}
```



Consistent Replicas

- Pass by far reference
 - Single copy
 - Serializes operations
- ✓ Strong consistency





Overview



Grocery Lists

New

Name...


Author...

Create

Lists

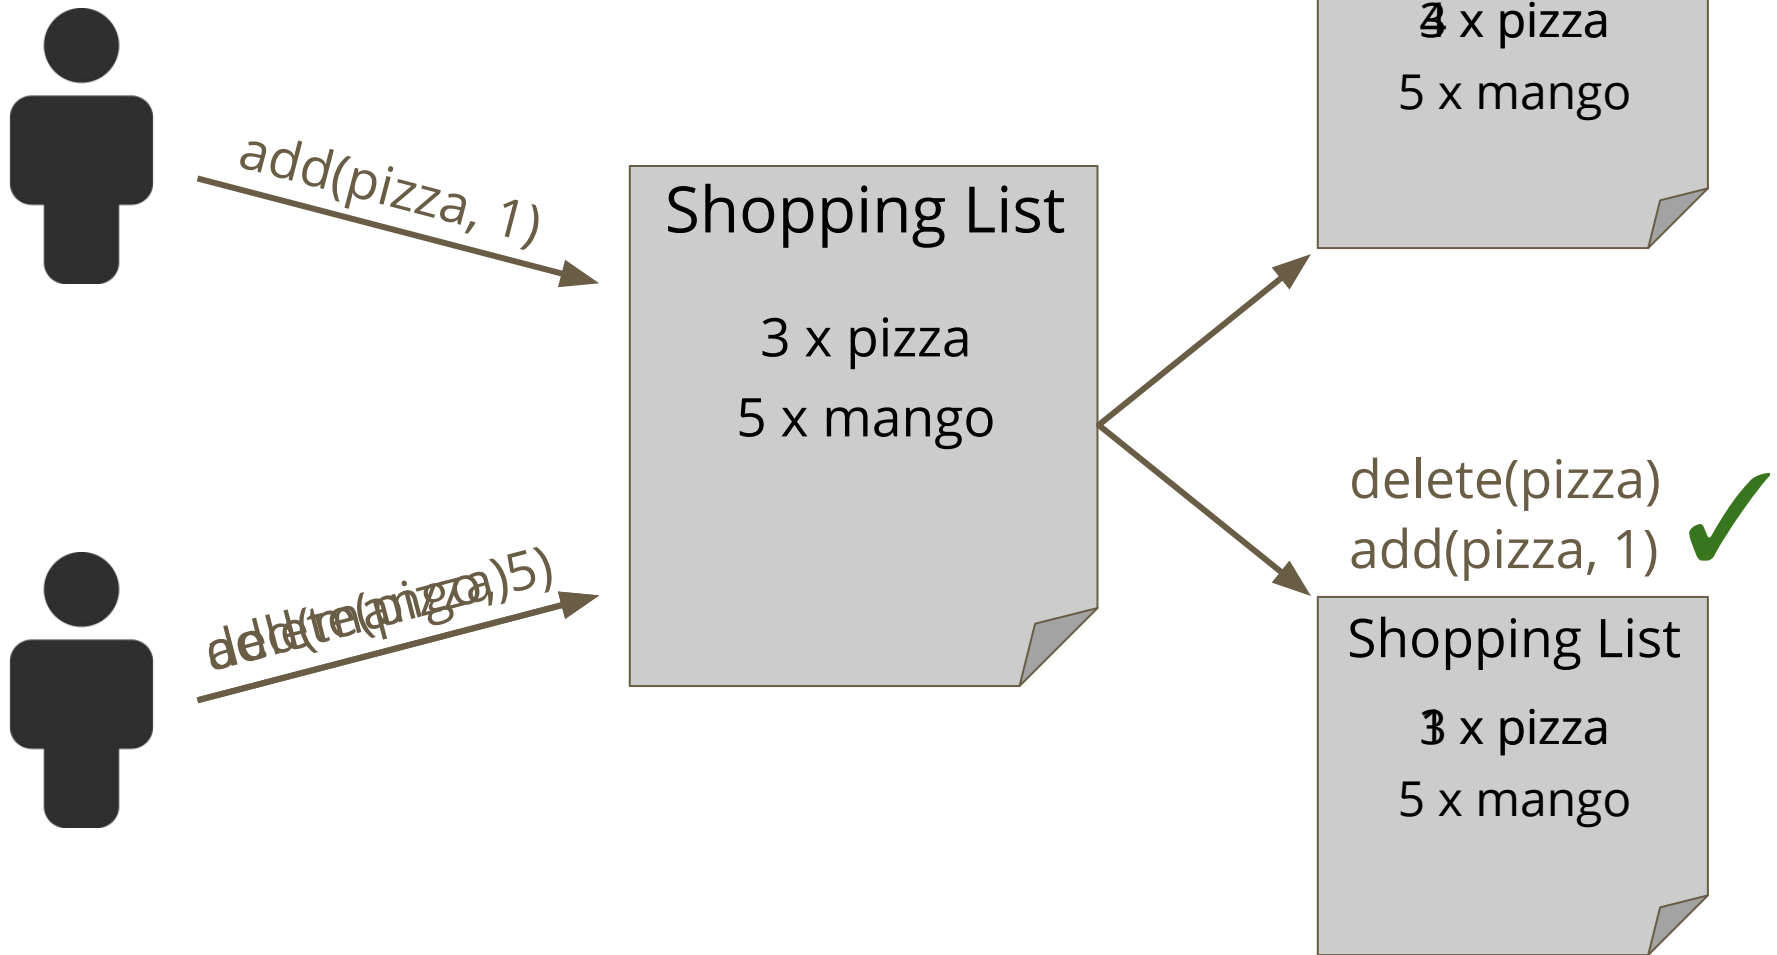
- Shopping List by Kevin +
- Mangos (3/5) +
- Lasagna (0/2) +
- Pizza (0/1) +

Each list item has a shopping bag icon and a close icon (X) to its right.



CRDTs are not generally applicable

NA



Strong Eventually Consistent Replicated Object

- General-purpose replicated data type
 - ✓ Availability
 - ✓ SEC
- No restrictions on operations
- Specify concurrent behavior
 - State validators



Grocery App: List

```
class GroceryList extends SECRO {  
  constructor() {  
    super();  
    this.items = new Map();  
  }  
  
  add(item, qty) {  
    const description =  
      this.items.getOrElse(item.name, {requested: 0, bought: 0});  
    description.requested += qty;  
    this.items.set(item.name, description);  
  }  
}
```



```
post add(oState, state, args, res) {  
  const [item] = args,  
        addedQty = item.requested,  
        resQty = state.items.getOrElse(item.name, 0).requested;  
  return resQty >= addedQty;  
}
```

```
delete(itemName) {  
  this.items.delete(itemName);  
}
```

```
}
```

SECRO: Algorithm

```

post add(oState, state, args, res) {
  /* ... */
  return resQty >= addedQty;
}

```



add(pizza, 1)



delete(pizza)

Shopping List


~~3~~ x pizza

5 x mango

Shopping List

3 x ~~pizza~~

5 x mango




~~[add(pizza, 1), delete(pizza)]~~

[delete(pizza), add(pizza, 1)]

[delete(pizza), add(pizza, 1)]

[add(pizza, 1), delete(pizza)]




Lack of language support






Limited applicability
of available solutions

Evaluation (I)



- Collaborative Text Editor
 - SECROs in CScript
 - JSON CRDT





Code Comparison

- Similar implementations
 - API delegates to list
- SECROs: pre/post conditions
 - More LOC
 - Higher flexibility
- Tree-based text editor
 - ✓ SECROs: re-use AVL tree
 - ✗ JSON CRDT: not possible

<pre>1 insertAfter(charID, char) { 2 const prev = charID === null ? 3 0 : this.indexOf(charID), 4 newChar = new Character(char, generate_id()); 5 this._content.idx(prev).insertAfter(newChar); 6 return newChar; 7 } 8 9 pre insertAfter(state, charID, char) { 10 return charID === null 11 state.indexOf(charID) !== -1; 12 } 13 14 post insertAfter(oState, state, args, newChar) { 15 /* ... */ 16 } 17 18 delete(charID) { 19 var idx = this.indexOf(charID); 20 if (idx !== -1) 21 this._content.idx(idx).delete(); 22 } 23 24 post delete(originalState, state, args, res) { 25 /* ... */ 26 }</pre>	<pre>1 insertAfter(charID, char) { 2 const idx = charID === null ? 3 0 : this.indexOf(charID), 4 newChar = new Character(char, generate_id()); 5 if (idx !== -1) { 6 cjrdt.doc.idx(idx).insertAfter(newChar); 7 return newChar; 8 } 9 } 10 11 12 13 14 15 16 17 18 delete(charID) { 19 const idx = this.indexOf(charID); 20 if (idx !== -1) 21 cjrdt.doc.idx(idx).delete(); 22 } 23 24 25 26</pre>
---	--

Evaluation (II)



- Benchmarks
 - Time complexity
 - Throughput
 - Memory

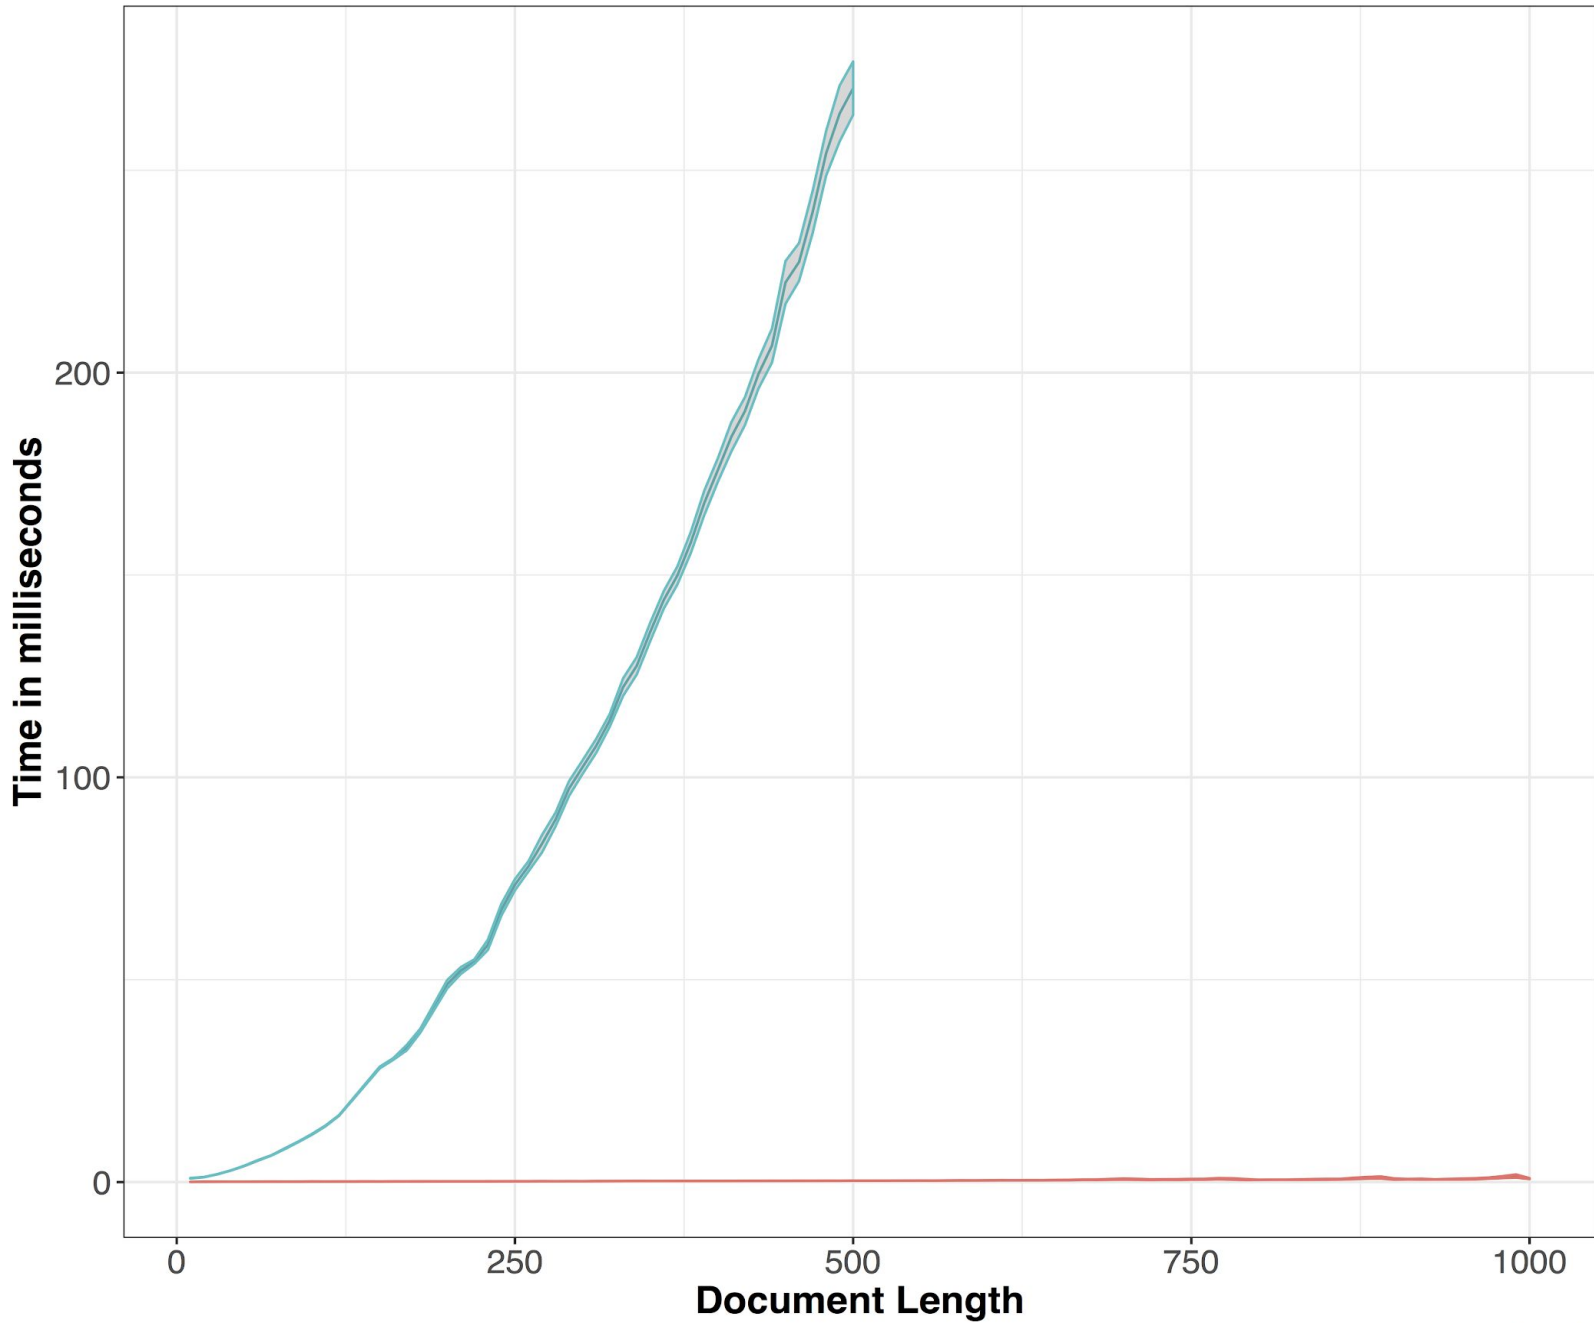
Methodology

- Warm-up rounds
- Disabled JIT compiler
- At least 30 repetitions
- Discard samples affected by GC
- Compute avg and 95% CI

Execution time of the list-based text editors

SECRO vs JSON CRDT


NA



Version
SECRO
JSON CRDT

Bands represent the 95% CI.

```
post insert(...){
  /* ... */
}
```



commit()

Document
abc
v₂



insert(null, "a")

insert("b", "c")

insert("a", "b")

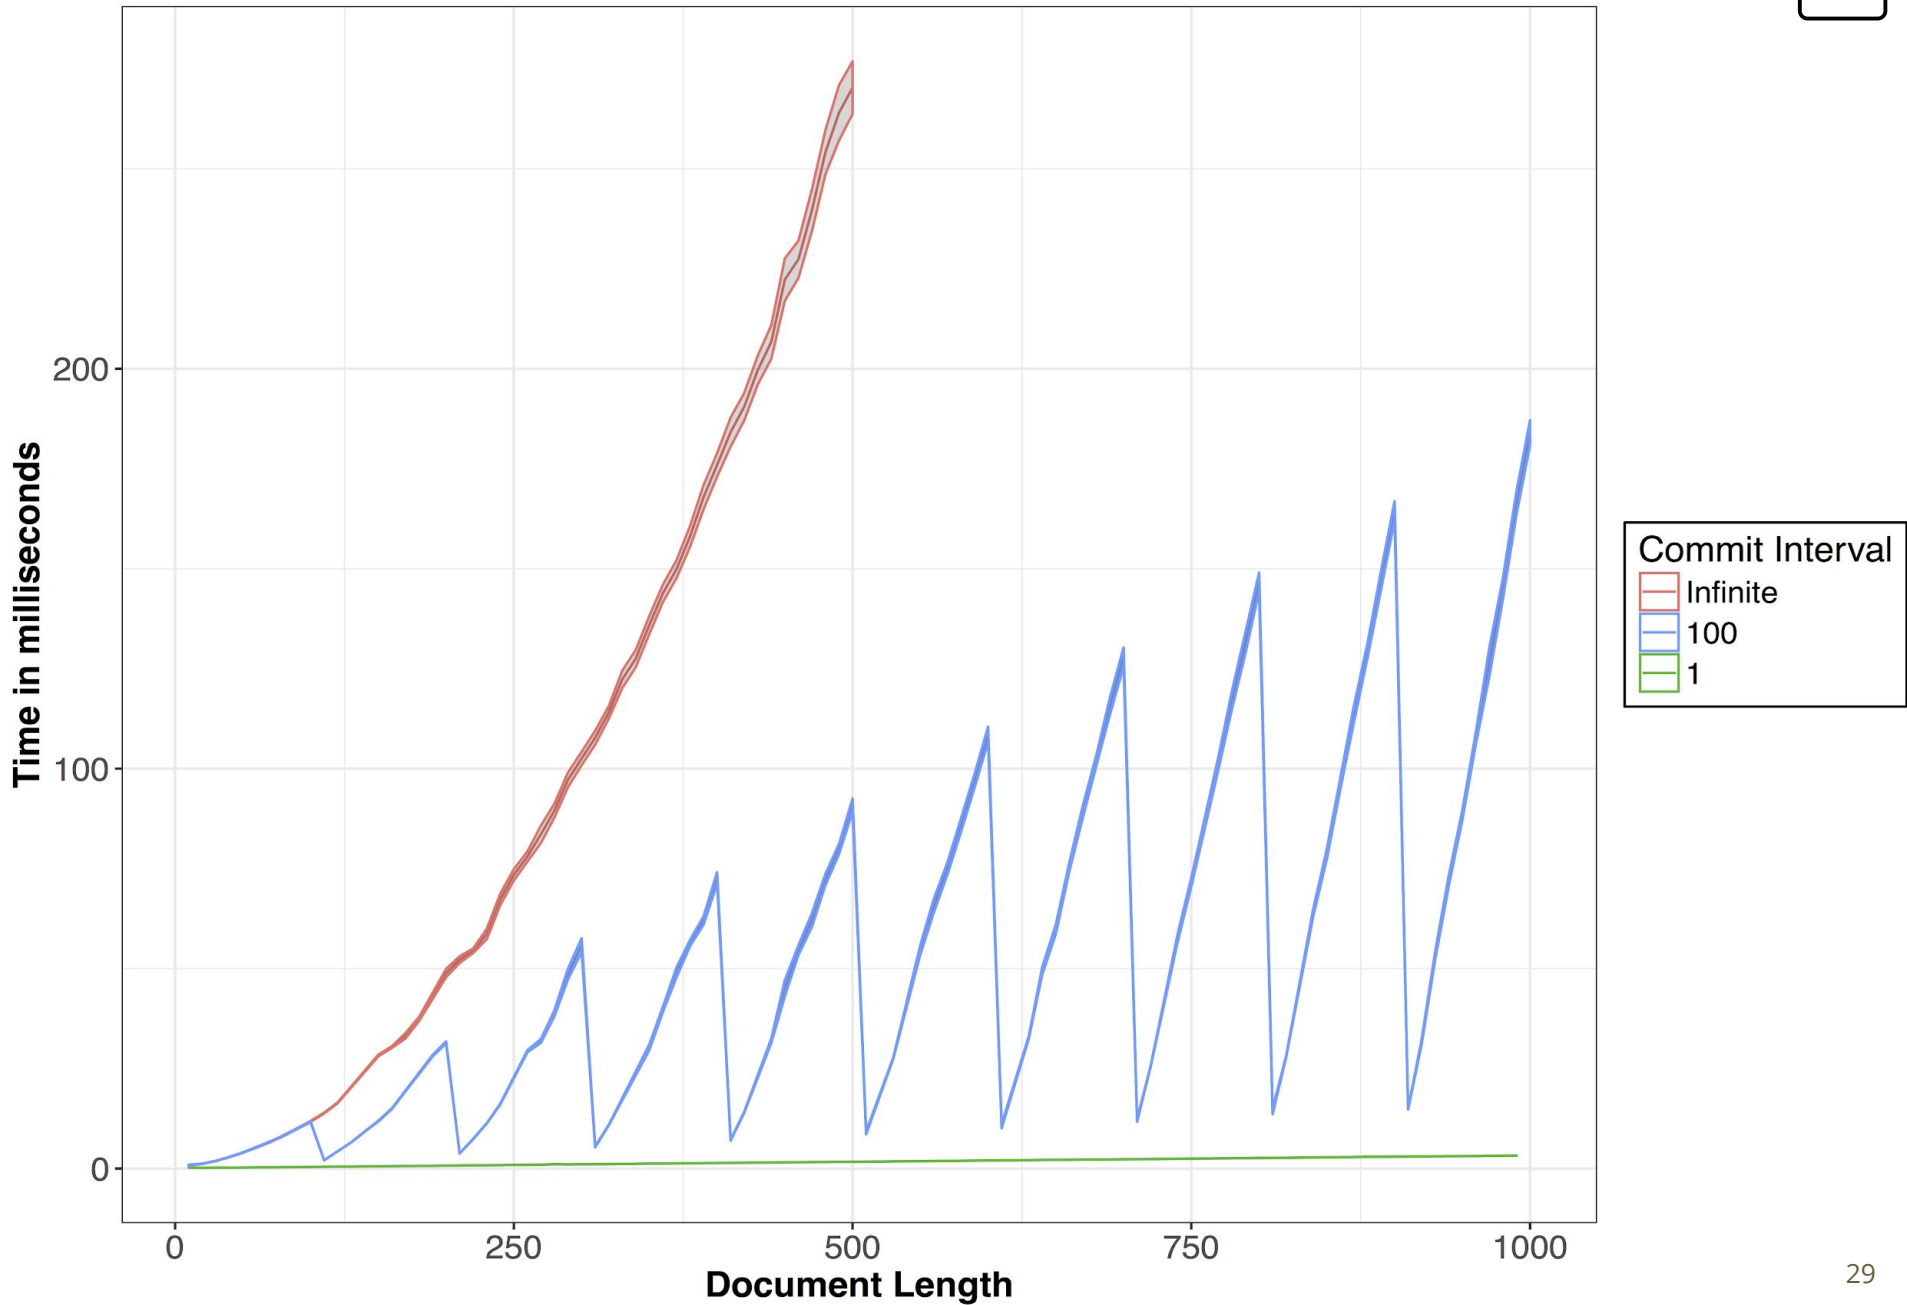


[insert(null, "a"), insert("a", "b"), insert("b", "c")]

Execution time of the text editor

For different commit intervals.

NA




Bands represent the 95% CI.

Throughput Of The Text Editor

Comparison of the naive JSON CRDT and SECRO implementations

NA




Error bars represent the 95% CI.
Commit interval of 100 for the SECRO version.

Memory usage of the text editors

JSON CRDT vs list and tree SECRO implementations

NA

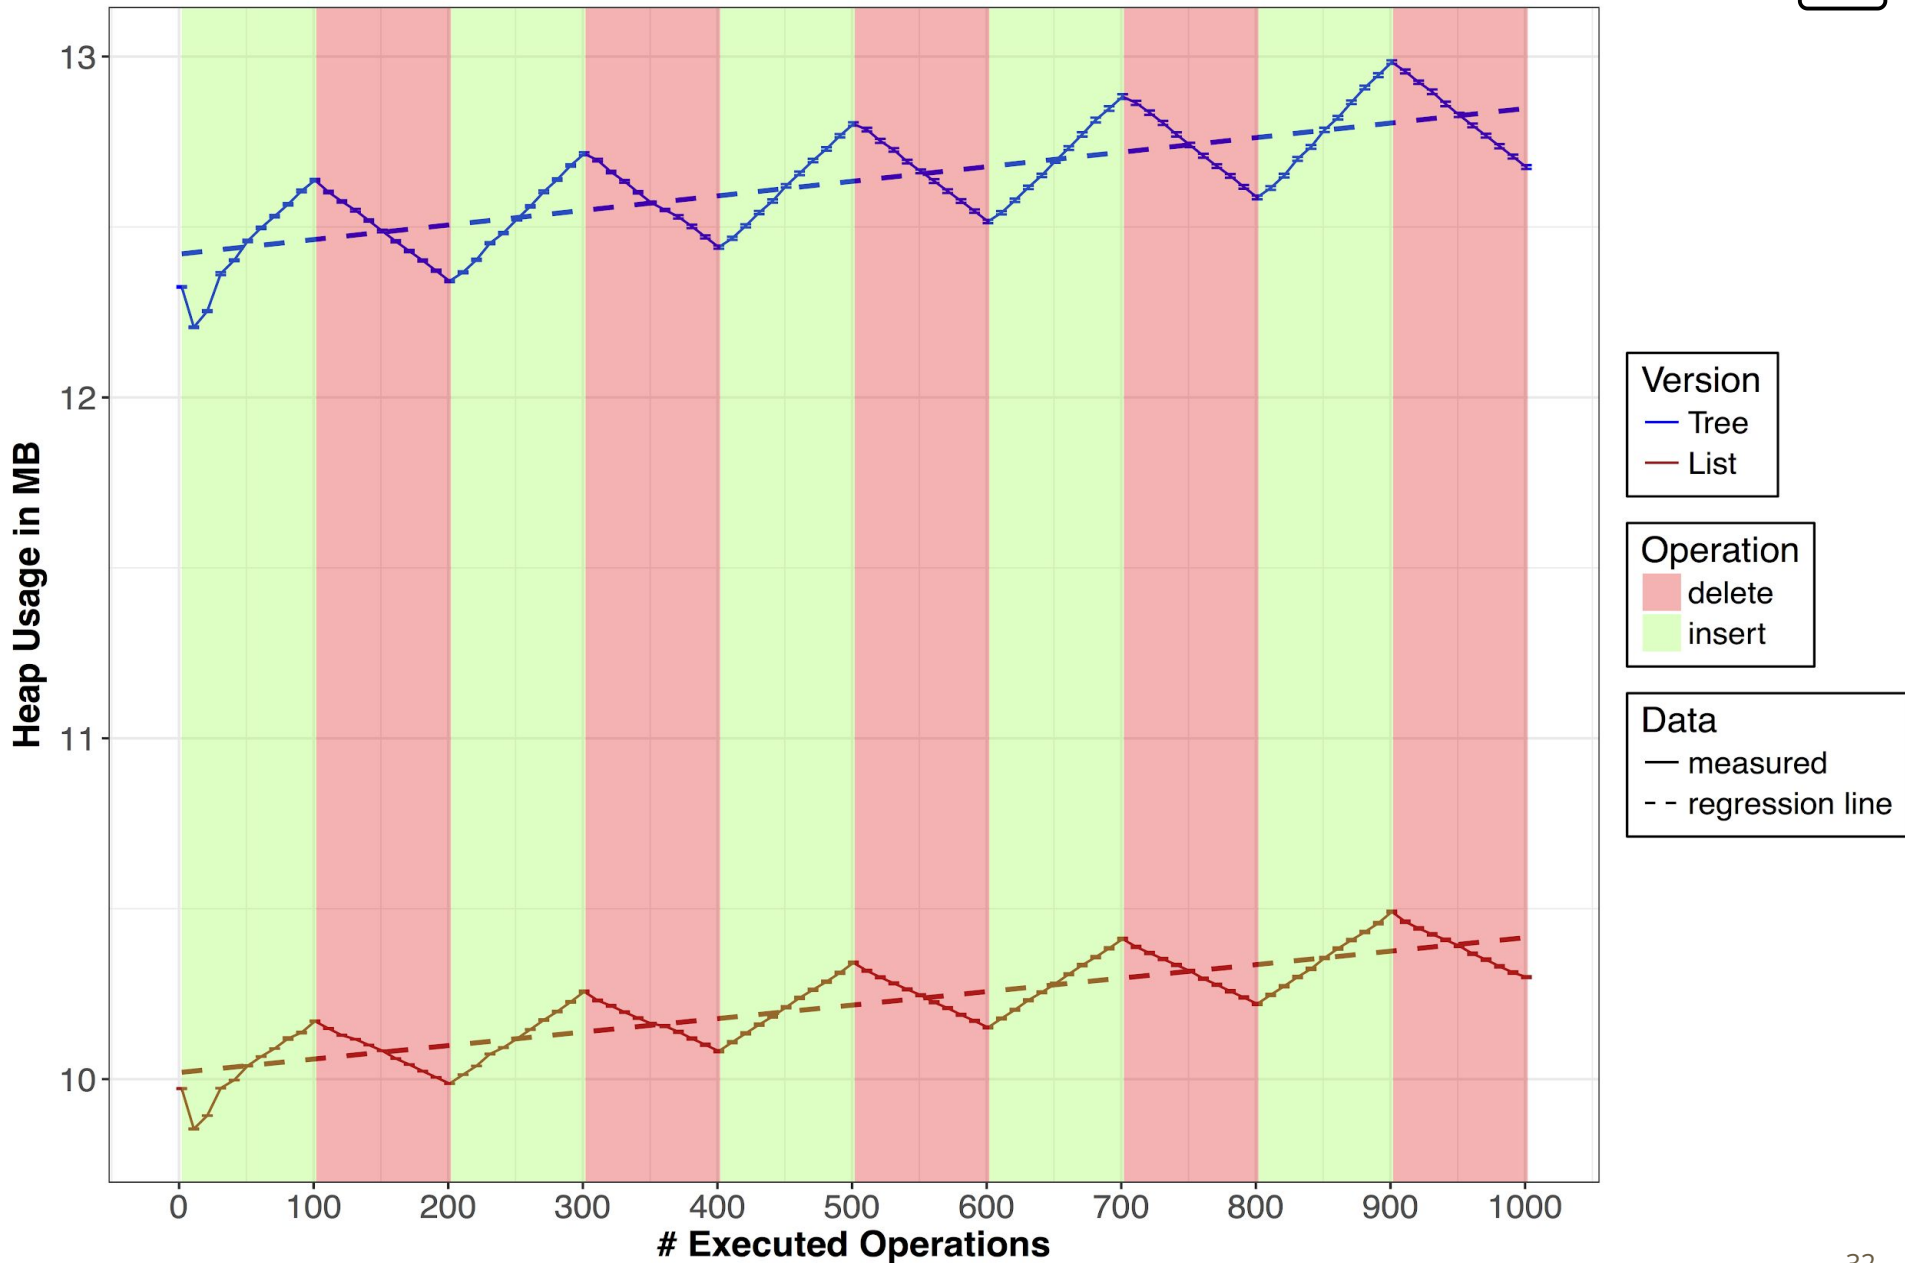


Error bars represent the 95% CI.
Commit and forced garbage collect after every operation.

Memory usage of the text editors


List vs tree SECRO implementations

NA



Error bars represent the 95% CI.
Commit and forced garbage collect after every operation.

Results




- More LOC
- But more flexible



- Uses less memory
 - But slower
- ⇒ `commit`

Conclusion



CScript: replicas




SECROs

- General-purpose
- Available
- SEC




Time to append a character

For the list and tree implementations of the SECRO text editor.




Error bars represent the 95% CI.
Replicas were not committed.

Copying Operation Postcondition Precondition




Copying Operation Postcondition Precondition



Time to append 100 characters

For the list and tree implementations of the SECRO text editor.



Error bars represent the 95% CI.
Replicas were not committed.