

Anticipation of Method Execution in Mixed Consistency Systems

Marco Giunti, Hervé Paulino, António Ravara

NOVA School of Science and Technology, Portugal

ACM/SIGAPP Symposium On Applied Computing
March 30 2023

Weak Consistency

- ▶ Citing Burckhardt (Foundations and Trends in Programming Languages 2014):
- ▶ In globally distributed systems, shared state is never perfect.
- ▶ When communication is neither fast nor reliable, we cannot achieve at the same time
 1. strong consistency
 2. low latency
 3. availability
- ▶ *Weak consistency* to overcome these limitations in *replicated systems*

Commutative concurrent operations

- ▶ We are interested in analysing concurrent operations in replicated systems
- ▶ Standard consistency prerequisite: operations must *commute*
- ▶ Focus on *language level*: Object-Oriented Programming (*OOP*)
- ▶ Rephrasing, our goal is:

identify commutative method calls and gather information on calls that in distributed runtime featuring replicas can be **anticipated**

Example: executing calls concurrently in two Sites

- ▶ **Locally permissible** ops immediately executed
- ▶ **Strongly consistent** ops require coordination among replicated sites

▶ Site 1

```
acc1.deposit(20);  
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

▶ Site 2

```
acc3.getBalance();  
acc1.deposit(20);  
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

Example: executing calls concurrently in two Sites

- ▶ **Locally permissible** ops immediately executed
- ▶ **Strongly consistent** ops require coordination among replicated sites

▶ Site 1

```
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

▶ Site 2

```
acc1.deposit(20);  
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

Example: executing calls in Site 1

- ▶ While `acc2.withdraw(5)` is put under coordination, Site 1 processes the next call
- ▶ Can we locally execute `acc2.deposit(10)`?
- ▶ Let's try

▶ Site 1

```
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

▶ Under coord.

```
acc2.withdraw(5);
```

▶ Site 2

```
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

Example: executing calls in Site 2

- ▶ After a step the system might **fatally diverge**
- ▶ E.g. if balance before deposit was < 5 then Site 1 allows withdrawal while Site 2 disallows it
- ▶ Eventual consistency is broken

▶ Site 1

```
acc3.deposit(20);  
acc2.getBalance();
```

▶ Under coord.

```
acc2.withdraw(5);
```

▶ Site 2

```
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```

Example: executing calls in Site 1

- ▶ Therefore in previous step deposit is put on hold as well
- ▶ Differently, the **next call** can be executed since it operates on a **different account**
- ▶ That is, the call can be **“anticipated”** w.r.t. operations received before

▶ Site 1

```
acc3.deposit(20);  
acc2.getBalance();
```

▶ Under coord.

```
acc2.withdraw(5);  
acc2.deposit(10);
```

▶ Site 2

```
acc2.withdraw(5);  
acc2.deposit(10);  
acc3.deposit(20);  
acc2.getBalance();
```


Related work

- ▶ Logical-based specifications (e.g. Bansal et al., JAR 2020)

```

name: set
preamble: (declare-sort E 0)
state: -name: S -type: (Set E) -name: size -type: Int
methods:
  name: add
  args: -name: v -type: E
  return: - name: result -type: Bool
  requires: true
  ensures:
    (ite (member v S)
      (and (= S_new S)
            (= size_new size)
            (not result))
      (and (= S_new (union S (singleton v)))
            (= size_new (+ size 1)) result))
  
```

Related work

- ▶ Abstract languages (e.g. Houshmand & Lesani, POPL 2019)

```

Class Courseware
let Student = Set <sid:SIId> in
let Course = Set <cid:CId> in
let Enrolment = Set <esid : SIId, ecid : CId> In
 $\Sigma$  = Student  $\times$  Course  $\times$  Enrolment
I =  $\lambda$  <ss,cs,es>.
    refIntegrity(es, esid, ss, sid)  $\wedge$ 
    refIntegrity(es, ecid, cs, cid)
addCourse(c) =  $\lambda$  <ss, cs, es>.
    <T, <ss, cs  $\cup$  c, es>,  $\perp$ >
deleteCourse(c) =  $\lambda$  <ss, cs, es>.
    <T, <ss, cs  $\setminus$  c, es>,  $\perp$ >
...
  
```

Challenge

- ▶ Can we process **source code** at *compile-time* to gather information on calls to **anticipate**?
- ▶ Idea: code analysis decides that deposit cannot anticipate withdraw on the *same instance*

```
class Account {  
  balance : int  
  ...  
  
  def withdraw(amount : int) : Unit {  
    this.balance -= amount  
  }  
  
  def deposit(amount : int) : Unit {  
    this.balance += amount  
  }  
  
  ...  
}
```

Purpose

- ▶ How we will use the **information on anticipation** generated at compile-time?
- ▶ Idea: populate a *table* that can be accessed **fast** by the runtime system
- ▶ Scheduler checks the table in order to anticipate `acc2.deposit(10)` w.r.t. `acc2.withdraw(5)`
- ▶ Anticipation is **forbidden**
- ▶ Instead, `acc3.deposit(20)` can anticipate the previous calls `acc2.withdraw(5)` and `acc2.deposit(10)`

```
/* WITHOUT TABLE */
acc2.withdraw(5);
acc2.deposit(10);
acc3.deposit(20);
acc2.getBalance();
```

```
/* VIRTUAL VIEW WITH TABLE */
acc3.deposit(20);
acc2.withdraw(5); /* ON HOLD */
acc2.deposit(10); /* ON HOLD */
acc2.getBalance();
```

Anticipation of Method Execution in Mixed Consistency Systems

Methodology: language-level consistency requirements

- ▶ We need to provide the consistency semantics
- ▶ Idea: *single keyword* to allow **weak** consistency of *fields*
- ▶ Remaining fields have strong consistency
- ▶ Design choice: field decoration requires less effort and fragmentation than assigning consistency to operations

```
class Account {  
  balance : int weak  
  ...  
  
  def withdraw(amount : int) : Unit {  
    this.balance -= amount  
  }  
  
  def deposit(amount : int) : Unit {  
    this.balance += amount  
  }  
}
```

State invariant

- ▶ We need to provide state invariant to infer the permissibility of ops
- ▶ Idea: field invariants and method preconditions

```
class Account {  
  balance : int weak [this.balance ≥ 0]  
  ...  
  
  def withdraw(amount : int) : Unit [?] {  
    this.balance -= amount  
  }  
  
  def deposit(amount : int) : Unit [amount > 0] {  
    this.balance += amount  
  }  
  
  ...  
}
```

Method's preconditions

- ▶ We cannot rely on the value of the **weak** field balance
- ▶ Withdraw has no preconditions

```
class Account {  
  balance : int weak [this.balance ≥ 0]  
  ...  
  
  def withdraw(amount : int) : Unit {  
    this.balance -= amount  
  }  
  
  def deposit(amount : int) : Unit [amount > 0] {  
    this.balance += amount  
  }  
  
  ...  
}
```

OOP core language

- ▶ We stand on the tradition of languages with formal semantics
- ▶ We consider a variant of Oolong (Castegren&Wrigstad, ACM SAC 2018)

d	$::= x \mid x.f \mid v$	<i>(Invariant values)</i>
c	$::= d_1 \text{ Rel } d_2$	<i>(Invariants)</i>
Md	$::= m(x : t_1) : t_2 [\check{c}] \{e\}$	<i>(Methods)</i>
Fd	$::= f : t \sim \text{weak} [\check{c}]$	<i>(Fields)</i>
e	$::= sv \mid x.f \mid x.f = e \mid x.m(e) \mid$ $\text{let } x = e_1 \text{ in } e_2 \mid \text{new } C \mid (t) e$	<i>(Expressions)</i>
sv	$::= v \mid x \mid sv_1 \text{ Op } sv_2$	<i>(Symbolic Values)</i>

Program reductions

- ▶ Standard Heap (H)/Stack (V) single-thread semantics

$$V(x) = \iota \quad H(\iota) = (\text{Account}, \text{balance} \mapsto n) \quad \text{this}', y' \text{ fresh} \\ e = \text{let } z = \text{this}'.\text{balance} \text{ in } \text{this}'.\text{balance} = z + y'$$

$$\langle H, V, -, -, x.\text{deposit}(10) \rangle \hookrightarrow \langle H, V[\text{this}' \mapsto \iota, y' \mapsto 10], -, -, e \rangle$$

- ▶ What about methods' preconditions?

$$\langle H, V, -, \epsilon, x.\text{deposit}(10) \rangle \hookrightarrow \langle H, V', -, 10 > 0, e \rangle$$

- ▶ A state Σ is composed by the four entries above

$$\langle \Sigma, x.\text{deposit}(10) \rangle \hookrightarrow \langle \Sigma', e \rangle$$

Formal semantics of permissible operations

- ▶ We leverage the Hamsaz model of conflicts (Houshmand & Lesani, POPL 2019)
- ▶ State invariant holds iff $f : t \sim_{\text{weak}} [\tilde{c}] \in \Sigma$ and $c \in \tilde{c}$ implies instantiation of c evals to true
- ▶ A call is *guarded* if post-state satisfies constraint, e.g. $10 > 0$
- ▶ A call is *permissible* in pre-state if pre-state invariant implies post-state invariant
- ▶ A method is *locally permissible (LP)* if all guarded calls are permissible

Method calls under coordination

- ▶ Calls of non-LP methods require **coordination**
- ▶ E.g. state invariant of *balance* requires non-negativity

```
class Account {  
  balance : int weak [this.balance ≥ 0]  
  ...  
}
```

- ▶ The call `x.withdraw(5)` can break state invariant
- ▶ E.g. if $V(x) = \iota$ and $H(\iota) = (\text{Account}, \text{balance} \mapsto 3)$
- ▶ Therefore, *withdraw* is non-LP

Formal semantics of commutative ops

- ▶ Commutative calls defined as expected
- ▶ E.g. $x_1.deposit(v_1)$ and $x_2.withdraw(v_2)$ commute in Σ iff
 1. Sequence $x_1.deposit(v_1); x_2.withdraw(v_2)$ gives rise to Σ_1
 2. Sequence $x_2.withdraw(v_2); x_1.deposit(v_1)$ gives rise to Σ_2
 3. $\Sigma_1 = \Sigma_2$

Call anticipation algorithm

- ▶ Novel notion relying on *weak fields integer generalization*
- ▶ Quantification over all possible integer values

```
let anticipation ( $\Sigma$  : state) (mc1 mc2 : call) : bool =
if commute  $\Sigma$  mc1 mc2 then
  let m1 = methodOf mc1 in let m2 = methodOf mc2 in
  match LP  $\Sigma$  m2  $\forall \forall$  n1, ..., nm. Permissible (WIFG  $\Sigma$  (n1, ..., nm)) mc2 with
  | true -> (* mc1 LP in post-state for all possible weak ints *)
            (* or mc1 preserves permissibility *) ...
  | false -> false
else false
```

- ▶ Limitation: algorithm is non-effective for runtime use
- ▶ Overhead of invoking constraint solver on all post-states unbearable

Aim of static analysis: parametric call anticipation

- ▶ Establishing **call anticipation** as yes or no is too restrictive
- ▶ We need to generate parameters for anticipations

```
def interest(interest : int) : Unit {
    this.balance += this.balance * interest / 100
}
...
x1.method(n);
x2.interest(i);
```

	getBalance	deposit	withdraw	interest
interest	✓	$x_1 \neq x_2 \wedge i \geq -100$	$x_1 \neq x_2 \wedge i \geq -100$	$x_1 \neq x_2 \wedge i \geq -100$

Symbolic memory

- ▶ The static analysis is built on top of *symbolic values*
 $sv ::= v \mid x \mid sv_1 \text{ Op } sv_2$
- ▶ **Symbolic semantics** : transitions with open terms to establish **commutativity**
- ▶ We consider **transitions of methods**, rather than method calls
- ▶ Example: *deposit* and *withdraw* commute?
- ▶ Symbolic heap: *balance* $\mapsto x$, where x is fresh
- ▶ Technique: execute sequences *deposit*; *withdraw* and *withdraw*; *deposit* and produce *heap equations*

Method commutativity, statically

- ▶ Execution *deposit; withdraw*, same “instance”
 1. Symbolic execution of *deposit*, d is formal parameter, leads to heap: $balance \mapsto x + d$
 2. Symb. execution of *withdraw*, w is formal parameter, leads to $balance \mapsto (x + d) - w$
- ▶ Execution *withdraw; deposit*; same “instance”
 1. Symb. execution of *withdraw* leads to $balance \mapsto x - w$
 2. Symbolic execution of *deposit* leads to heap: $balance \mapsto (x - w) + d$
- ▶ Algorithm produces equation $((x + d) - w, (x - w) + d)$
- ▶ Equation is SAT for all values $x \geq 0, d > 0$ (inferred from state invariant)
- ▶ In general, symbolic values in fields and methods parametrize commutativity

Method anticipation, statically

- ▶ Our objective is populate table with parameters that allow call anticipation
- ▶ Use: fast access by runtime system to take decision
- ▶ Methodology: generate a *list of logical conjunctions*
- ▶ We distinguish equality case: e.g. $this_1 = this_2$ and $other_1 \neq other_2$
- ▶ Init state populated with init objects with symbolic integers

```

let ant (h : eqCase) (c1 c2 : classDef) (md1 md2 : methodDef) : prop list =
  let  $\Sigma$  = gen h md1 md2 in
  let eqs, cc = scommute h md1 md2 in
  let eff2 = hasEffect c2 md2 in
  let  $\langle \Sigma'', \_ , \_ \rangle$  = update_s  $\Sigma$  eff2 null in
  eqs :: cc :: sLP  $\Sigma$  md2 ::
    (sLP  $\Sigma''$  md1  $\vee$ 
      $\forall x$  in (weak_inv  $\Sigma$ ).
      ((sP  $\Sigma$  md1 => sP  $\Sigma''$  md1)  $\wedge$ 
       (sNP  $\Sigma$  md1 => sNP  $\Sigma''$  md1))) :: [SFNI h md1 md2]
  
```

Example: anticipating LP method

► Formula:

```
eqs :: cc :: sLP  $\Sigma$  md2 ::
  ( _  $\vee$ 
     $\forall x$  in (weak_inv  $\Sigma$ ).
      ((sP  $\Sigma$  md1 => sP  $\Sigma''$  md1)  $\wedge$ 
       (sNP  $\Sigma$  md1 => sNP  $\Sigma''$  md1))) :: [SFNI h md1 md2]
```

► Consider positive example that relies on right disjunction

```
persons int weak [this.persons mod 2 = 0]
tables   int weak [this.tables  $\geq$  0]
```

```
/* commutative methods */
def addTable() { this.tables += 1 } /* LP */
def addPerson () { this.persons += 1 } /* non-LP */
```

Anticipation of `addTable` w.r.t. `addPerson` ($\text{this}_1 = x = \text{this}_2$)

$$\Sigma = x \mapsto (-, \text{pers} \mapsto p, \text{tbl} \mapsto t)$$

$$\Sigma' = x \mapsto (-, \text{pers} \mapsto p + 1, \text{tbl} \mapsto t)$$

$$\Sigma'' = x \mapsto (-, \text{pers} \mapsto p, \text{tbl} \mapsto t + 1)$$

$$\text{inv}(\Sigma) = p \bmod 2 = 0 \wedge t \geq 0$$

$$\text{inv}(\Sigma') = p + 1 \bmod 2 = 0 \wedge t \geq 0$$

$$\text{inv}(\Sigma'') = p \bmod 2 = 0 \wedge t + 1 \geq 0$$

► Simplified formula:

$$\text{eqs} :: \text{cc} :: \text{sLP } \Sigma \text{ addT} :: \\
\forall p, t. ((\text{sP } \Sigma \text{ addP} \Rightarrow \text{sP } \Sigma'' \text{ addP}) \wedge (\text{sNP } \Sigma \text{ addP} \Rightarrow \text{sNP } \Sigma'' \text{ addP}))$$

► We have the following subformulas

$$\text{eqs} = (p + 1, p + 1), (t + 1, t + 1)$$

$$\text{cc} = \text{true}, \text{true}$$

$$\text{sLP } \Sigma \text{ addT} = \forall p, t. \text{inv}(\Sigma) \Rightarrow \text{inv}(\Sigma'')$$

$$\text{sP } \Sigma \text{ addP} = \text{inv}(\Sigma) \Rightarrow \text{inv}(\Sigma')$$

$$\text{sP } \Sigma'' \text{ addP} = \text{inv}(\Sigma'') \Rightarrow p + 1 \bmod 2 = 0 \wedge t + 1 \geq 0$$

$$\text{snP } \Sigma \text{ addP} = \text{inv}(\Sigma) \Rightarrow \neg \text{inv}(\Sigma')$$

$$\text{snP } \Sigma'' \text{ addP} = \text{inv}(\Sigma'') \Rightarrow \neg(p + 1 \bmod 2 = 0 \wedge t + 1 \geq 0)$$

Anticipation of `addTable` w.r.t. `addPerson` ($\text{this}_1 = x = \text{this}_2$)

$$\Sigma = x \mapsto (-, \text{pers} \mapsto p, \text{tbl} \mapsto t)$$

$$\Sigma' = x \mapsto (-, \text{pers} \mapsto p + 1, \text{tbl} \mapsto t)$$

$$\Sigma'' = x \mapsto (-, \text{pers} \mapsto p, \text{tbl} \mapsto t + 1)$$

$$\text{inv}(\Sigma) = p \bmod 2 = 0 \wedge t \geq 0$$

$$\text{inv}(\Sigma') = p + 1 \bmod 2 = 0 \wedge t \geq 0$$

$$\text{inv}(\Sigma'') = p \bmod 2 = 0 \wedge t + 1 \geq 0$$

► Simplified formula:

$$\text{eqs} :: \text{cc} :: \text{sLP } \Sigma \text{ addT} :: \\ \forall p, t. ((\text{false} \Rightarrow \text{sP } \Sigma'' \text{ addP}) \wedge (\text{true} \Rightarrow \text{sNP } \Sigma' \text{ addP}))$$

► We have the following subformulas

$$\text{eqs} = (p + 1, p + 1), (t + 1, t + 1)$$

$$\text{cc} = \text{true}, \text{true}$$

$$\text{sLP } \Sigma \text{ addT} = \forall p, t. \text{inv}(\Sigma) \Rightarrow \text{inv}(\Sigma'')$$

$$\text{sP } \Sigma \text{ addP} = \text{false}$$

$$\text{sP } \Sigma'' \text{ addP} = _$$

$$\text{sNP } \Sigma \text{ addP} = \text{true}$$

$$\text{sNP } \Sigma'' \text{ addP} = \text{inv}(\Sigma'') \Rightarrow \neg(p + 1 \bmod 2 = 0 \wedge t + 1 \geq 0)$$

Formula is SAT! `addTable` is anticipated Anticipation of Method Execution in Mixed Consistency Systems

Use cases

Results obtained with Java prototype (Thank Rúben Vaz)

Use-case	# weak/ # strong	# invs	# mtds	# non-LP	# pairs	# conflicts
Account	1/1	2	6	3	21	5
Auction	1/1	2	4	0	9	3
Counter	1/0	0	3	0	6	0
Register	3/0	0	2	0	3	0

Use cases

$m_2 \xrightarrow{c} m_1$ indicates that, given sequence $m_1; m_2$, method m_2 can anticipate m_1 when c holds

Use-case	# anticipations
Account	$a \xrightarrow{i \geq -100} g, a \xrightarrow{\text{this}_1 \neq \text{this}_2 \wedge i_2 \geq -100} \{d, t, w, a\},$ $g \rightarrow *, d \rightarrow \{d, g\}, d \xrightarrow{\text{this}_1 \neq \text{this}_2} \{t, w, a\}$
Auction	$b \rightarrow cb, cb \rightarrow *, c \rightarrow c, b \rightarrow b$
Counter	$* \rightarrow *$
Register	$g \rightarrow *, * \rightarrow g, s \rightarrow s,$

Thanks!

- ▶ Evaluation with Java prototype
 1. querying the table < 0.001 ms
 2. several orders of magnitude lower than performing global synchronization
- ▶ Main limitation: consistency only supported for primitive types
- ▶ Directions for future work
 - ▶ Extend the language, e.g. conditionals, loops. . .
 - ▶ Mechanization of proof of soundness, e.g. in the Coq proof assistant