

Hoare-style Specifications as Correctness Conditions for Non-Linearizable Concurrent Objects

Ilya Sergey



joint work with

Aleks Nanevski, Anindya Banerjee, and Germán Andrés Delbianco

Linearizable Concurrent Objects

Linearizability: A Correctness Condition for Concurrent Objects

MAURICE P. HERLIHY and JEANNETTE M. WING
Carnegie Mellon University

A concurrent object is a data object shared by concurrent processes. Linearizability is a correctness condition for concurrent objects that exploits the semantics of abstract data types. It permits a high degree of concurrency, yet it permits programmers to specify and reason about concurrent objects using known techniques from the sequential domain. Linearizability provides the illusion that each operation applied by concurrent processes takes effect instantaneously at some point between its invocation and its response, implying that the meaning of a concurrent object's operations can be given by pre- and post-conditions. This paper defines linearizability, compares it to other correctness conditions, presents and demonstrates a method for proving the correctness of implementations, and shows how to reason about concurrent objects, given they are linearizable.

Non-overlapping calls to methods of a *concurrent* object should appear to take effect in their *sequential* order.

Linearizability is expensive

Laws of Order: Expensive Synchronization in Concurrent Algorithms Cannot be Eliminated

Hagit Attiya

Technion
hagit@cs.technion.il

Rachid Guerraoui

EPFL
rachid.guerraoui@epfl.ch

Danny Hendler

Ben-Gurion University
hendlerd@cs.bgu.ac.il

Petr Kuznetsov

TU Berlin/Deutsche Telekom Labs
pkuznets@acm.org

Maged M. Michael

IBM T. J. Watson Research Center
magedm@us.ibm.com

Martin Vechev

IBM T. J. Watson Research Center
mtvechev@us.ibm.com

An alternative to linearizability?

The advent of multicore processors as the standard computing platform will force major changes in software design.

BY NIR SHAVIT

Data Structures in the Multicore Age

Relaxing the correctness condition would allow one to implement concurrent data structures more efficiently, as they would be free of synchronization bottlenecks.

Alternatives to linearizability

- Quiescent Consistency [[Aspnes-al:JACM94](#)]
- Quasi-Linearizability [[Afek-al:OPODIS10](#)]
- Quantitative Relaxation [[Henzinger-al:POPL13](#)]
- Concurrency-Aware Linearizability [[Hemed-Rinetzky:PODC14](#)]
- Quantitative Quiescent Consistency [[Jagadeesan-Riely:ICALP14](#)]
- Local Linearizability [[Haas-al:arXiv15](#)]
- ...

Challenges of diversity

- *Composing* different conditions (CAL, QC, QQC) in a single program, which uses multiple objects;
- Providing *syntactic* proof methods for establishing all these conditions (akin to *linearization points*);
- Employing these criteria for *client-side reasoning* (*uniformity*).

Hoare-style Specifications as Correctness Conditions for Non-Linearizable Concurrent Objects

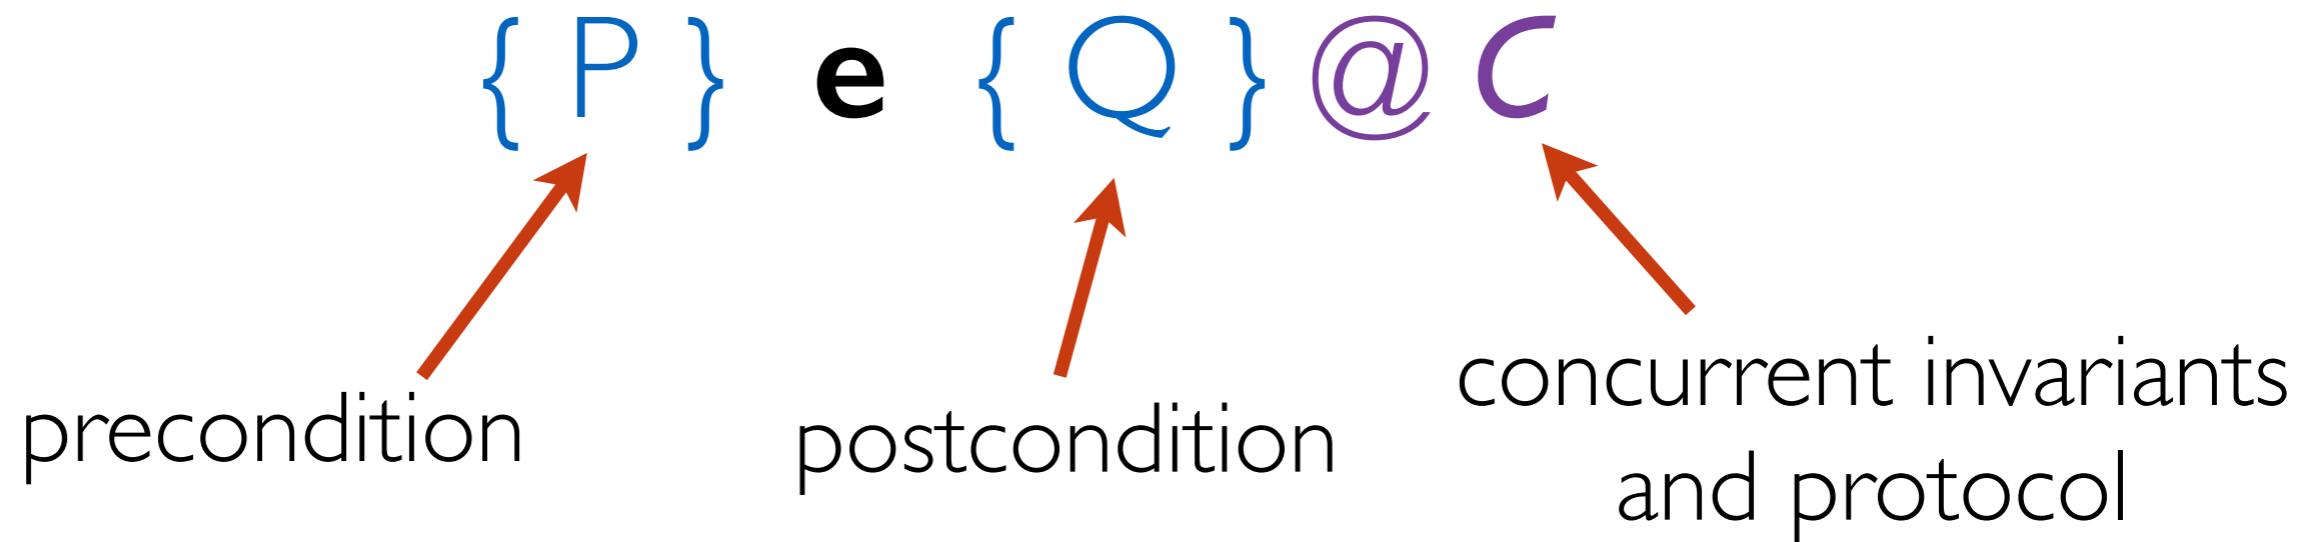
Ilya Sergey



joint work with

Aleks Nanevski, Anindya Banerjee, and Germán Andrés Delbianco

Hoare-style Specifications



If the initial *state* satisfies P , then, after \mathbf{e} terminates, the final *state* satisfies Q (no matter the *interference* manifested by C).

Hoare-style Specifications

$$\{ P \} \mathbf{e} \{ Q \} @ C$$

- *Compositional* — substitution principle;
- *Syntactic proof method* — inference rules;
- *Uniform* — reasoning about objects and their clients in the *same* proof system.

Hoare-style Specifications as CAL, QC, QQC

Concurrency-Aware Linearizability (CAL):

Effects of some concurrent method calls should appear to happen simultaneously.

Quiescent Consistency (QC):

Method calls separated by a period of quiescence should appear to take effect in their order.

This talk

Quantitative Quiescent Consistency (QQC):

The number of out-of-order method results is bounded by the number of interfering threads (with a constant factor).

Simple Counting Network

```
def getAndInc() : nat
```

Simple Counting Network


```
def getAndInc() : nat = {  
  n ← &x;  
  b ← CAS(x, n, n + 1);  
  if b then  
    return n;  
  else getAndInc();  
}
```

high contention location




Simple Counting Network

```
def getAndInc() : nat = {  
  b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```




Sequential Execution (T_1)

```
def getAndInc() : nat = {  
  b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```



Sequential Execution (T_1)

```
def getAndInc() : nat = {  
→ b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```



$$T_1.b_1 = 0$$

Sequential Execution (T_1)

```
def getAndInc() : nat = {  
  b ← flip(bal);  
→ res ← fetchAndAdd2(x + b);  
  return res;  
}
```

bal
1

x
2


x+1
1

$T_1.b_1 = 0$

$T_1.res_1 = 0$

Sequential Execution (T_1)

```
def getAndInc() : nat = {  
→ b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```



$$T_1.b_1 = 0$$

$$T_1.res_1 = 0$$

$$T_1.b_2 = 1$$

Sequential Execution (T_1)

```
def getAndInc() : nat = {  
  b ← flip(bal);  
  → res ← fetchAndAdd2(x + b);  
  return res;  
}
```

bal
0

x
2

x+1
3

$$T_1.b_1 = 0$$

$$T_1.res_1 = 0$$

$$T_1.b_2 = 1$$

$$T_1.res_2 = 1$$

Concurrent Execution (T_1 , T_2)

```
def getAndInc() : nat = {  
  b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```


bal
0

x
0

x+1
1

Concurrent Execution (T_1 , T_2)


```
def getAndInc() : nat = {  
→ b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```



$T_1.b_1 = 0$

Concurrent Execution (T_1, T_2)

```
→ def getAndInc() : nat = {  
→   b ← flip(bal);  
   res ← fetchAndAdd2(x + b);  
   return res;  
}
```



$T_1.b_1 = 0$

$T_2.b_1 = 1$

Concurrent Execution (T_1 , T_2)

```
def getAndInc() : nat = {  
  → b ← flip(bal);  
  → res ← fetchAndAdd2(x + b);  
  return res;  
}
```

bal
0

x
0

x+1
3


$T_1.b_1 = 0$

$T_2.b_1 = 1$

$T_2.res_1 = 1$

Concurrent Execution (T_1 , T_2)

```
→ def getAndInc() : nat = {  
  → b ← flip(bal);  
  res ← fetchAndAdd2(x + b);  
  return res;  
}
```



$$T_1.b_1 = 0$$

$$T_2.b_1 = 1$$

$$T_2.res_1 = 1$$

$$T_2.b_2 = 0$$

Concurrent Execution (T_1 , T_2)

```
def getAndInc() : nat = {  
  → b ← flip(bal);  
  → res ← fetchAndAdd2(x + b);  
  return res;  
}
```

bal
1

x
2

x+1
3

$T_1.b_1 = 0$

$T_2.b_1 = 1$

$T_2.res_1 = 1$

$T_2.b_2 = 0$

$T_2.res_2 = 0$

Correctness Conditions for Counting Network

- ~~**R₀**: calls to `getAndInc()` take effect in their sequential order~~
- **R₁**: different calls return *distinct* results (strong concurrent counter)
- **R₂**: two calls, separated by *period of quiescence*, take effect in their sequential order (QC)
- **R₃**: results of *two calls* in the same thread are out of order by no more than $2 * (\text{number of calls interfering with both})$ (QQC)

Invariants of the Counting Network

“Tokens”

- Every flip of the balancer grants thread a *capability* to add 2 to a counter (x or $x+1$);
- Each of the counters (x and $x+1$) changes *continuously* wrt. *even/odd* values;
- Threads, which gained capabilities but haven't yet incremented, cause one counter to “*run ahead*” of another one, leading to *out-of-order* anomalies.

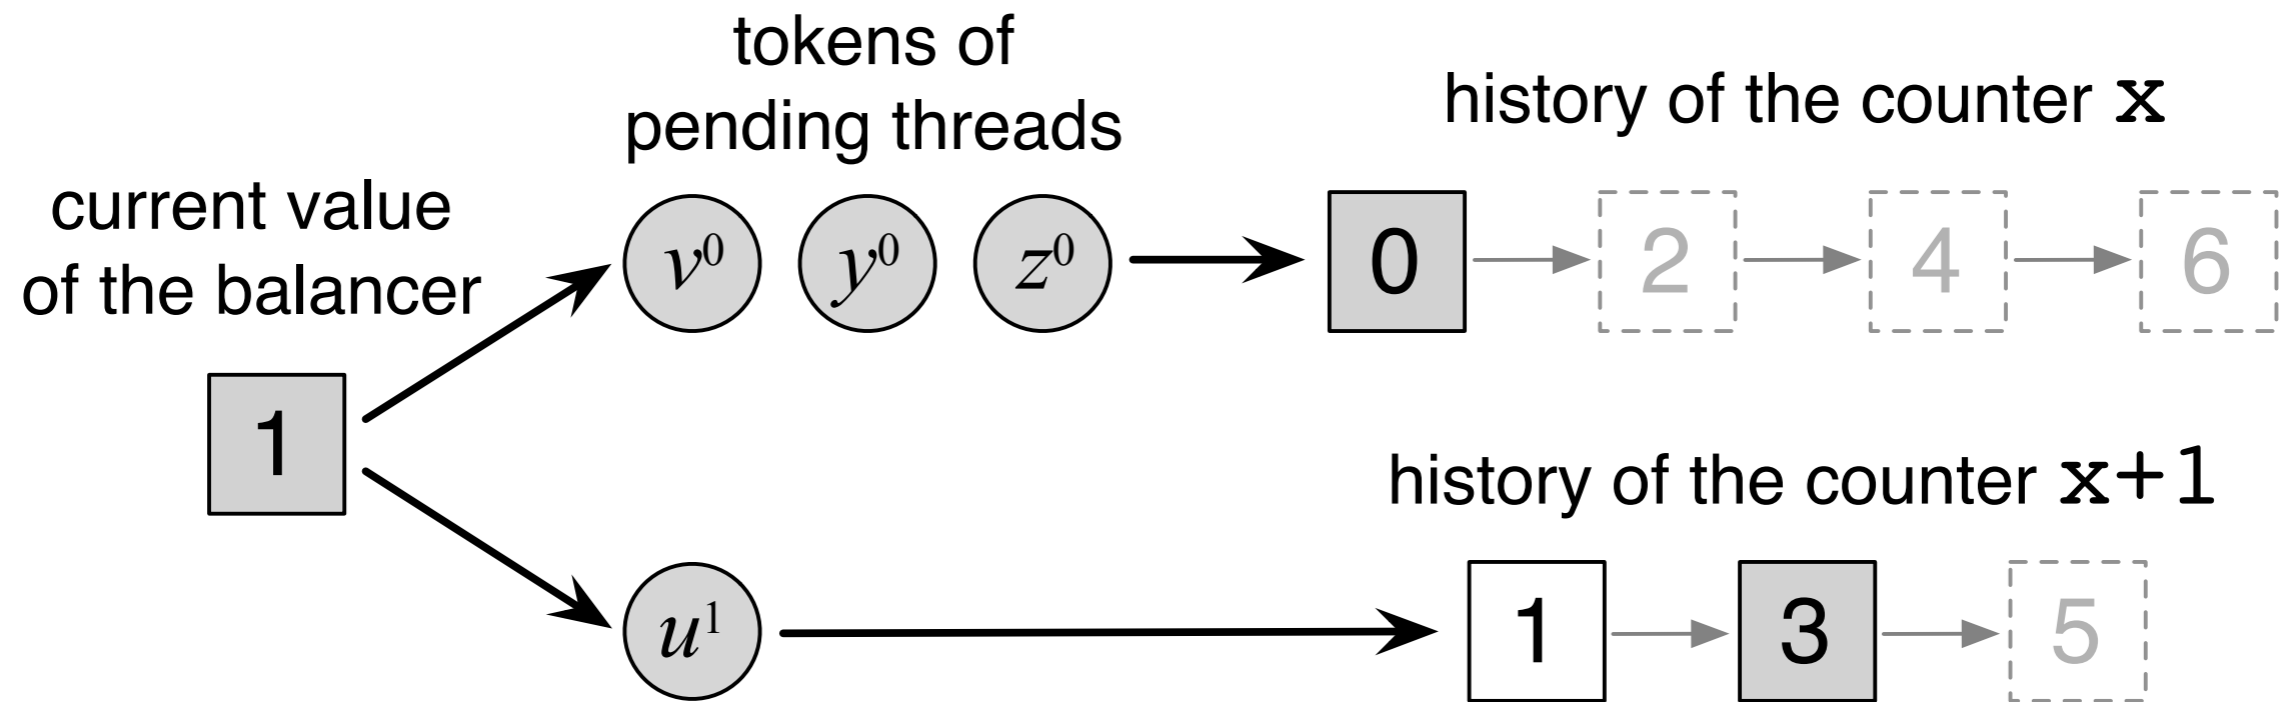
“Histories”

Sergey-al:ESOP15

Real and Auxiliary State

- Hoare-style specs constrain **state**, *auxiliary* or *real*
- **Real state** — *heap* (pointers `ba1`, `x`, `x+1`);
- **Auxiliary state** — any *fictional splittable* resource, represented as a PCM $(S, \oplus, \mathbf{0})$, e.g.,
 - ◆ *Tokens* — disjoint sets;
 - ◆ *Histories* — partial maps with `nat` as domain.

Tokens and Histories of the Network



- New *unique* tokens are emitted upon calling `flip()`;
- Calling `fetchAndAdd2()` *consumes* a token and *adds* an entry to the history.

Interference-capturing histories

$$\eta = \{ \dots, \boxed{t} \mapsto (l, z), \dots \}$$

“timestamp”, a value written to a counter x or $x+1$ (0, 1, 2, etc.)

Interference-capturing histories

$$\eta = \{ \dots, t \mapsto (l, z), \dots \}$$

sets of tokens, held by interfering threads
at the moment the entry has been written

Interference-capturing histories

$$\eta = \{ \dots, t \mapsto (l, \boxed{z}, \dots) \}$$

a token, spent to increment x or $x+1$ from $t-2$ to t

Notation for *Subjective* Histories and Tokens

- χ_s, χ_o — histories, contributed by *self* and *other* threads;
- τ_s, τ_o — tokens, held by *self* and *other* threads;
- η, ι — logical variables for histories and tokens.

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s, \\ \eta_o \subseteq \chi_o, \\ l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\}$$

`getAndInc()`

$$\left\{ \exists l, z, \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z), \\ \eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \\ \text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 | l \cap l_o | \end{array} \right\} @ C$$

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \boxed{\tau_s = \emptyset, \chi_s = \eta_s,} \leftarrow \text{initial tokens} \\ \eta_o \subseteq \chi_o, \\ \text{lo} \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\} \text{and self-history}$$

`getAndInc()`

$$\left\{ \exists l, z, \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z), \\ \eta_o \subseteq \chi_o, \text{lo} \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \\ \text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 | \iota \cap \text{lo} | \end{array} \right\} @ C$$

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s, \\ \eta_o \subseteq \chi_o, \\ l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\}$$

boundary on initial
other-history and tokens



`getAndInc()`

$$\left\{ \exists l, z, \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z), \\ \eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \\ \text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 | l \cap l_o | \end{array} \right\} @ C$$

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s, \\ \eta_o \subseteq \chi_o, \\ l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\}$$

final tokens
and self-history

`getAndInc()`

$\{\exists l, z,$

$$\tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z),$$

$$\eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o),$$

$$\text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 | l \cap l_o | \} @ C$$

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s, \\ \eta_o \subseteq \chi_o, \\ l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\}$$

`getAndInc()`

constraining final
other-history and tokens

$\{\exists l, z,$

$$\tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z),$$

$$\eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o),$$

$$\text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 | l \cap l_o | \} @ C$$

Specification of `getAndInc()`

$$\left\{ \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s, \\ \eta_o \subseteq \chi_o, \\ l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o) \end{array} \right\}$$

`getAndInc()`

$$\left\{ \exists l, z, \begin{array}{l} \tau_s = \emptyset, \chi_s = \eta_s \cup \text{res} + 2 \mapsto (l, z), \\ \eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \\ \text{last}(\eta_s \cup \eta_o) < \text{res} + 2 + 2 \mid l \cap l_o \mid \end{array} \right\} @ C$$

result + 2 is
*greater than any previous value
of the counters, recorded in history
(modulo past n present interference)*

Implications of the derived spec

Trivial from invariants: each result corresponds to a new history entry


- **R₁**: different calls return *distinct* results (strong concurrent counter)
- **R₂**: two calls, separated by *period of quiescence*, take effect in their sequential order (QC)
- **R₃**: results of *two calls* in the same thread are out of order by no more than $2 * (\text{number of calls interfering with both})$ (QQC)

Implications of the derived spec

- **R₁**: different calls return *distinct* results (strong concurrent counter)
- **R₂**: two calls, separated by *period of quiescence*, take effect in their sequential order (QC)
- **R₃**: results of *two calls* in the same thread are out of order by no more than $2 * (\text{number of calls } \textit{interfering with both})$ (QQC)

Exercising Quiescent Consistency

“quiescent moment”



```
(res1, -) ← (getAndInc() || e1);
```

```
(res2, -) ← (getAndInc() || e2);
```

```
return (res1, res2);
```

```
{ ; res1 < res2? }
```

Generic spec for Interference

$$\{ \tau_s = \emptyset, \chi_s = \emptyset, l_0 \subseteq \tau_0 \cup \text{spent}(\chi_0) \}$$

e_i

$$\{ \exists \eta_i, \tau_s = \emptyset, \chi_s = \eta_i, l_0 \subseteq \tau_0 \cup \text{spent}(\chi_0) \} @ C$$

arbitrary contribution to the history



Spec for parallel composition

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \dots \}$

`getAndInc()` || e_i

$\{ \exists l, z, \eta_i, \tau_s = \emptyset, \chi_s = \eta_s \cup \eta_i \cup \text{res} + 2 \mapsto (l, z), \eta_o \subseteq \chi_o,$
 $l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o),$
 $\text{last}(\eta_s \cup \eta_o) < \text{res.1} + 2 + 2 | l \cap l_o | \} @ C$

Spec for parallel composition

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \eta_o \subseteq \chi_o, l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o), \dots \}$

`getAndInc()` || e_i

$\{ \exists l, z, \eta_i, \tau_s = \emptyset, \chi_s = \eta_s \cup \eta_i \cup \text{res}+2 \mapsto (l, z), \eta_o \subseteq \chi_o,$
 $l_o \subseteq \tau_o \cup \text{spent}(\chi_o \setminus \eta_o),$
 $\text{last}(\eta_s \cup \eta_o) < \text{res.1} + 2 + 2 | l \cap l_o | \} @ C$

$$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$$

`(res1, -) ← (getAndInc() || e1);`

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
 where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -$, $\eta_0 = \chi_0$ and $\iota_0 = \tau_0 \}$

`(res2, -) ← (getAndInc() || e2);`

$\{ \exists \eta_1, \eta_2, \iota, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $\iota_0 \subseteq \tau_0 \cup \text{spent}(\chi_0 \setminus \eta_0),$

$\text{last}(\eta'_s \cup \eta_0) < \text{res}_2 + 2 + 2 \mid \iota \cap \iota_0 \}$

return `(res1, res2);`

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 \subseteq \tau_0 \cup \text{spent}(\chi_0 \setminus \eta_0),$
 $\text{last}(\eta'_s \cup \eta_0) < \text{res}_2 + 2 + 2 | \text{in } l_0 | \}$

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 \subseteq \emptyset,$
 $\text{last}(\eta'_s \cup \eta_0) < \text{res}_2 + 2 + 2 | \text{len } l_0 | \}$

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

$\text{last}(\eta'_s \cup \eta_0) < \text{res}_2 + 2 + 2 | \text{len } l_0 | \}$

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

last($\eta'_s \cup \eta_0$) < res₂ + 2 + 2 | ln l₀ | }

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

last($\eta'_s \cup \eta_0$) < res₂ + 2 + 2 | \emptyset |}

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$
where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -, \eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$
last($\eta'_s \cup \eta_0$) < res₂ + 2 }

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

(res₁, -) ← (getAndInc() || e₁);

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$

where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -$, $\eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

(res₂, -) ← (getAndInc() || e₂);

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

last($\eta'_s \cup \eta_0$) < res₂ + 2 }

return (res₁, res₂);

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

$(res_1, -) \leftarrow (getAndInc() \ || \ e_1);$

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$

where $\eta'_s = \eta_s \cup \eta_1 \cup res_1 + 2 \mapsto -$, $\eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

$(res_2, -) \leftarrow (getAndInc() \ || \ e_2);$

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup res_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

$res_1 + 2 < res_2 + 2 \}$

return $(res_1, res_2);$

$\{ \tau_s = \emptyset, \chi_s = \eta_s, \dots \}$

`(res1, -) ← (getAndInc() || e1);`

$\{ \exists \eta_1, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0,$

where $\eta'_s = \eta_s \cup \eta_1 \cup \text{res}_1 + 2 \mapsto -$, $\eta_0 = \chi_0$ and $l_0 = \tau_0 \}$

`(res2, -) ← (getAndInc() || e2);`

$\{ \exists \eta_1, \eta_2, l, \tau_s = \emptyset, \chi_s = \eta'_s, \eta_0 \subseteq \chi_0, \chi_s = \eta'_s \cup \eta_2 \cup \text{res}_2 + 2 \mapsto -$
 $l_0 = \emptyset,$

`res1 < res2`

`return (res1, res2);`

Implications of the derived spec

- **R₁**: different calls return *distinct* results (strong concurrent counter)
- **R₂**: two calls, separated by *period of quiescence*, take effect in their sequential order (QC)
- **R₃**: results of *two calls* in the same thread are out of order by no more than $2 * (\text{number of calls } \textit{interfering with both})$ (QQC)

Summary of the proof pattern

- Express interference that matters via auxiliary state — *tokens*;
- Capture past interference and results in auxiliary *histories*;
- Assume closed world to *bound* interference.

Not discussed today

- Full formal specification of the counting network;
- Formal proofs of QC and QQC properties for the network;
- Discussion on applying the technique for QC-queues;
- Spec and verification of `java.util.concurrent.Exchanger`;
- Verification of an exchanger client in the spirit of *concurrency-aware linearizability* (CAL).



To take away

Hoare-style Specifications for Non-linearizable Concurrent Objects

- *Compositional* — substitution principle;
- *Syntactic proof method* — inference rules;
- *Uniform* — reasoning about objects and their clients in the *same* proof system.

Specification is in the eye of the beholder.

Thanks!