# Linearizability with ownership transfer

Hongseok Yang University of Oxford

Joint work with Alexey Gotsman (IMDEA Software Institute, Spain)

#### Concurrent libraries

- Encapsulate efficient concurrent algorithms.
  - Java: java.util.concurrent
  - C++:Threading Building Blocks
  - C#: System.Collections.Concurrent
- Implement stacks, queues, skip lists, hash tables, etc.
- But it is not easy to understand why they are correct.









#### 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.

Categories and Subject Descriptors: D.1.3 [Programming Techniques]: Concurrent Programming; D.2.1 [Software Engineering]: Requirements/Specifications; D.3.3 [Programming Languages]: Language Constructs—abstract data types, concurrent programming structures, data types and structures; F.1.2 [Computation by Abstract Devices]: Modes of Computation—parallelism; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs—pre- and post-conditions, specification techniques

General Terms: Theory, Verification

Additional Key Words and Phrases: Concurrency, correctness, Larch, linearizability, multiprocessing, serializability, shared memory, specification

#### 1. INTRODUCTION

#### 1.1 Overview

Informally, a concurrent system consists of a collection of sequential processes that communicate through shared typed objects. This model encompasses both message-passing architectures in which the shared objects are message queues,

A preliminary version of this paper appeared in the Proceedings of the 14th ACM Symposium on Principles of Programming Languages, January 1987 [21].

This research was sponsored by IBM and the Defense Advanced Research Projects Agents (DOD), ARPA order 4976 (Amendment 20), under contract F33615-87-C-1499, monitored by the Avionics Laboratory, Air Force Wright Aeronautical Laboratories, Wright-Patterson AFB. Additional sport for J. M. Wing was provided in part by the National Science Foundation under grant CCR-8620027. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

Authors' address: Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3890.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1990 ACM 0164-0925/90/0700-0463 \$01.50

ACM Transactions on Programming Languages and Systems, Vol. 12, No. 3, July 1990, Pages 463-492.





#### A binary relation L ⊑ L' on libraries L and L'.

 Usually, L is an impl. and L' is a spec.

# Problem

#### Shared address space



#### Shared address space



 Linearizability assumes a complete isolation between the library and its client.

#### Shared address space



- Linearizability assumes a complete isolation between the library and its client.
- If client and library share address space, they can corrupt each other.

#### Ownership transfer



- Boundary between data structures owned by the library and its client is not static.
- Ownership = right to access.
- Ownership of cells transferred between library and client.

• Typically stores pointers to data structures.



• Typically stores pointers to data structures.



• Typically stores pointers to data structures.



• Typically stores pointers to data structures.



Can't access B after transferring its ownership!

### Our paper

- I. We generalised linearizability to a realistic setting: shared address space, ownership transfer.
- 2. Our generalisation gives the Abstraction Theorem:
- $L \sqsubseteq L' \Rightarrow Can replace L by L' in a proof of C in C(L).$



# Review of linearizability

#### Histories



(tl, call push(42))(t2, call pop)(tl, ret push)(t2, ret pop(42))(t2, call push(11)) ...

#### Histories





- We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.



- We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.



- We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.



- We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.



- $H \sqsubseteq H'$ 
  - We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.

$$L \sqsubseteq L' \iff \forall H \in \llbracket L \rrbracket, \exists H' \in \llbracket L' \rrbracket, H \sqsubseteq H'$$

$$H t : (t1, call push(42)) (t1, ret push) (t1, call isEmpty) (t1, ret isEmpty(no)) (t2, call push(11)) (t2, ret push) (t1, call push(42)) (t1, call isEmpty) (t2, call push(11)) (t2, ret push) (t1, call push(42)) (t1, call isEmpty) (t2, call push (11)) (t2, ret push) (t1, ret isEmpty(no)) (t2, ret pop(42)) (t2, call push(11)) (t2, ret push)$$

 $H \sqsubseteq H'$ 

- We can permute calls and returns by different threads.
- But non-overlapping method invocations can't be rearranged.

# Our results

# Specified library "L: []"

- Γ defines a light-weight specification {p<sub>m</sub>}m{q<sub>m</sub>} for every method m in L.
- E.g.  $\{n\mapsto _\}push(n)\{emp\}, \{emp\}pop\{res\mapsto _\}.$
- pm determines cells transferred from client to library.
- qm determines cells transferred from library to client.

#### Histories with ownership transfer



- Before: (t, call m(k)), (t, ret m(k)) for k in Integer.
- Now: (t, call m( $\theta$ )),  $\theta \in p_m$ ; (t, ret m( $\theta$ )),  $\theta \in q_m$ .
- Linearizability as before.

### Why were we surprised?

• Important properties of standard linearizability rely on the movability of call and ret actions.



### Why were we surprised?

• Important properties of standard linearizability rely on the movability of call and ret actions.



### Why were we surprised?

• Important properties of standard linearizability rely on the movability of call and ret actions.



But this movability might not hold in the presence of transferred cells.

#### Well-balanced histories

 Histories that have correct accounting of cells owned by library and its client.



#### Well-balanced histories

 Histories that have correct accounting of cells owned by library and its client.



#### Well-balanced histories

 Histories that have correct accounting of cells owned by library and its client.



# Semantics of specified library [[L:[]]







- new *p*: allocates an arbitrary state satisfying *p*.
- delete q: removes the part of state satisfying q.

# Semantics of specified library [[L:[]]



- [L:Γ] consists of histories generated by this most general client.
- Lemma: Only well-behaved histories are generated.



#### • Libraries L and L' have the same specifications $\Gamma$ .

#### Abstraction Theorem



- Client C and libraries L, L' :  $\Gamma$  are safe.
- $L \sqsubseteq L'$ . Then client( $\llbracket C(L) \rrbracket$ )  $\subseteq$  client( $\llbracket C(L') \rrbracket$ )

#### Can replace L by L' in a proof of C:



#### $C(L') \models P \Rightarrow C(L) \models P$



#### Can replace L by L' in a proof of C:



#### $C(L') \models P \Rightarrow C(L) \models P$