Machine-Verifying Concurrent Data Structures

Siddhartha Jayanti

Google Research Dartmouth

Collaborators and Co-authors

[POPL 2024] with Prasad Jayanti, Ugur Yavuz, Lizzie Hernandez [DIST 2021, PODC 2019, PODC 2016] with Robert Tarjan [Google Research] with several, including Laxman Dhulipala, Guy Blelloch

Motivation: Fast Data Structures



Motivation: Fast Concurrent Data Structures

- Fast Data Structures \rightarrow Fast Algorithms

- Fast Concurrent Data Structures → Fast Parallel Algorithms



Goal

Design and Deploy **Fast and Scalable** Multicore (i.e., Concurrent) Data Structures

Google open source graph-mining library

google / graph-mining	Public			A Notifications	rk 15 🖓 Star 514 →	
<> Code • Issues 1 Pi	ull requests 🕟 Actions 🖽 Pro	ojects ① Security i insights	Go to file Code +	About		
	Laxman Dhulipala Fix graph_mining namespaces		a last week 🕲 11 commits	No description, website, or topics provided.		
 docs examples in_memory utils .bazelrc .gitignore BUILD.oss CONTRIBUT 	docs	Boilerplate for new Google open source project	2 years ago	다 Readme 회 Apache-2.0 license		
	in_memory	Fix graph_mining namespaces	last week	 ⊗ Code of conduct ▲ Security policy Activity 514 stars 0 14 watching 15 forks 		
	utils	Remove double-license and update README.md	last week			
	gitignore	Update gitignore and add example from cjcarey and gzuzic's	CLs. last week			
	BUILD.oss CONTRIBUTING.md	V0 of graph-mining repo V0 of graph-mining repo	last week	Report repository		
	LICENSE	Boilerplate for new Google open source project	2 years ago last week	Releases		
	WORKSPACE.bazel	V0 of graph-mining repo	last week	Destroyee		
	i≣ README.md			No packages published		
	The Graph Mir	ning Library 🖉	Languages			
	This project includes some tools by the <u>Google Graph Mining team</u> , namely in-memory clustering. Our tools can be used for solving data mining and machine learning problems that either inherently have a graph structure or can be formalized as graph problems. For more information, see our <u>NeurIPS'20 workshop</u> .			 C++ 93.1% Starlark 6.7% C 0.2% 		
	Among others, this repository contains shared memory parallel clustering algorithms which scale to graphs with tens of billions of edges and are based on the following research papers:					

Challenge

Asynchrony makes concurrent algorithm design notoriously hard

- p^t possible executions of length t
- Uncountably many infinite executions
- ALL possible executions must produce desired behavior

Subtle Race Conditions: some "hard to detect" executions goes wrong



Asynchrony makes concurrent algorithms hard to design and prove correct Subtle race conditions have plagued even mission-critical concurrent code



Mars Pathfinder Rover

Multi-million dollar space project jeopardized



Northeast Blackout of 2003

~45 Million people in the US affected by power cuts which lasted 8 hours



Therac-25

Patients given Radiation doses >100x desired amount; 3 deaths

Goal

Machine-verified proofs of correctness (i.e., linearizability) for concurrent data objects

Correctness = Linearizability

[Herlihy & Wing, 1990]



Time

Each operation *must appear to take effect atomically*, i.e.,

instantaneously at some point between its invocation and completion

Forward Simulation: A Simple Proof Technique [Jonsson, 1991]



Time

- Maintain an atomic reference object
- Induct over arbitrary run
- Identify Linearization Points as they occur
- Show that return values of implemented object match those of the atomic object if operations taking place at linearization points

But what if...



Time

- Is this linearizable?
- Actually, it is linearizable: just with different linearization points
- Lessons:
 - We can't always determine linearization points as they occur
 - Our choice of linearization points of past operations may depend on **future** outcomes

Future-Dependence

Some implementations are inherently future-dependent (e.g., the Herlihy-Wing queue)

So, standard forward simulation is <u>simple</u> but <u>not complete</u> (forward simulation can't be used to prove all data structures)

Is there a *simple* proof technique for linearizability, that is **complete** and can be used to obtain **machine-verified proofs**?

Previous Work

- Linearizability Introduced
- Forward Simulation
- Prophecy Variables
- Aspect-Oriented Proofs
- Backward Simulation
- Commitment Points
- Partial-Order Maintenance
- Prophecy Variables in Separation Logic
- Category Theory Based

Herlihy and Wing 1990 Jonsson 1991 Vafeiadis 2008 Henzinger et al. 2013 Schellhorn et al. 2014 Bouajjani et al. 2017 Khyzha et al. 2017 Jung et al. 2019 Olivera-Vale et al. 2023

Color key: Incomplete, Not Universal, Unmechanized, Not Simple

Our Contributions: Proof Technique

Meta-Configurations Tracking proof technique for Linearizability

- Universal: Objects of **any type** can be handled
- Complete: Any algorithm that is linearizable can be proved so

First such **forward reasoning** technique

Proofs can be Machine-Certified

Our Contributions: Strong Linearizability

We also introduce a related proof technique, called **Singleton Tracking**, which is a universal, sound, and complete forwardreasoning technique to prove Strong Linearizability (a close-cousin of linearizability).

Our Contributions: Machine-verified Proofs

Herlihy-Wing Queue

- Intricate future-dependent linearization structure
- Jayanti Snapshot (single-writer, single-scanner)
 - Efficient algorithm with future-dependent linearization
- Jayanti-Tarjan Union-Find Objects
 - Widely-used across domains & in Google's graph-mining library

<u>Computation</u>	<u>Citation</u>	<u>#cores</u>	<u>Benchmark</u>	<u>speed-up</u>
Spatial Clustering	[Wang et al]	36	Sequential	5 – 33×
Structural Clustering	[Tseng et al]	48	Sequential	5 – 32×
Model Checking	[Bloemen]	64	Sequential	14.16 – 24.35 ×
Connected Components	[Dhulipala et al]	72	Multicore	1.5 – 4.02×

Our Idea

- Given original algorithm A:
 - Design Augmentation A^* that *tracks* the set *M* of *all possible* atomic reference objects, corresponding to all possible linearizations



Meta-Configurations





Meta-Configurations





How we track all meta-configurations

Initially $M := \{m_0\}$, where m_0 has the initial state and all processes are idle

Three Update Rules:

1. When π invokes op(arg)

For each $m \in M$: update $m.\pi$ from (-, -, -) to (op, arg, -)

2. When π returns **r** for op(arg)

Filter out each $m \in M$ if $m.\pi$'s third field is not r

otherwise, update m. π to (–, –, –) since π is now idle

3. When π executes a step

Evolve each $m \in M$ to reflect that any set of pending operations can linearize in any order.

Meta-Configurations Tracking

Initialize

 $M := \{m_0\}$, where m_0 .state is initial state, $\forall \pi \in \Pi$, $m_0.\pi = (-, -, -)$ Update Rules:

- When <u>π invokes op(arg)</u>
 M := {m' | ∃m∈ M:m'.π = (op, arg, -) and otherwise m' = m}
- When <u>π executes a step</u>

 $M := \{m' \mid \exists m \in M, \exists S \subseteq pending(m), \exists \alpha \in perm(S): m' = \alpha(m)\}$

• When π returns **r** for op(arg)

 $M := \{m' \mid \exists m \in M: m.\pi.res = r, m'.\pi = (-, -, -), and o/w m' = m\}$

Meta-Configurations Tracking: Main Theorem

 For an algorithm A, let A* be the same algorithm augmented with a variable M that tracks all meta-configurations.



Initially $A[0,1,2,...] = [\bot, \bot, \bot,...]$ and X = 0

Enqueue(v):

// take next empty slot & place v into it.
s := X.F&Inc()
A[s] := v

Dequeue(): // read X; scan A[0,...,X-1] and return first found element; try again, if no element is found. while (true) len := Xfor (s := 0; s < len; s++) $r := A[s].swap(\perp)$ if $r != \bot$ then return r

Herlihy-Wing Queue: Proof

Want to prove: $M \neq \{\}$ is an invariant of A^*

• Main Intellectual Work: Identify structure of *M*.

• Strengthen to Inductive Invariant

• **Proof by Induction:** over the length of A*

• Encode the proof in TLAPS and let the machine verify!

Herlihy-Wing Queue: Proof

 $I \equiv I_L \wedge I_X \wedge I_Q \wedge I_v \wedge I_i \wedge I_l \wedge I_j \wedge I_x \wedge I_{pc} \wedge I_{\mathcal{M}}$ $\wedge I_{1,5} \wedge I_2 \wedge I_3 \wedge I_4 \wedge I_6 \wedge I_7 \wedge I_8 \wedge I_9 \wedge I_a \wedge I_b$

In the above expression, the various conjuncts on the right hand side are defined below.

• $I_L \equiv \mathcal{M} \neq \emptyset$ • $I_X \equiv X \in \mathbb{N}^+$, $I_Q \equiv \forall k \in \mathbb{N}^+ : Q[k] \in \mathbb{N}^+ \cup \{\bot\}$, $I_v \equiv \forall \pi \in \Pi : v_\pi \in \mathbb{N}^+$, $I_i \equiv \forall \pi \in \Pi : i_\pi \in \mathbb{N}^+$ $I_{l} \equiv \forall \pi \in \Pi : l_{\pi} \in \mathbb{N}^{+}, I_{j} \equiv \forall \pi \in \Pi : j_{\pi} \in \mathbb{N}^{+}, I_{x} \equiv \forall \pi \in \Pi : x_{\pi} \in \mathbb{N}^{+} \cup \{\bot\}, I_{pc} \equiv \forall \pi \in \Pi : pc_{\pi} \in [9]$ • $I_{\mathcal{M}} \equiv \mathcal{M} \subseteq \{(\sigma, f) : \sigma \in \bigcup_{n \in \mathbb{N}} (\mathbb{N}^+)^n, f \in (\{\text{Enqueue}, \text{Dequeue}\} \times (\mathbb{N}^+ \cup \{\bot\}) \times (\mathbb{N}^+ \cup \{ack, \bot\}))^{\Pi}\}$ • $I_{1,5} \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} \in \{1, 5\} \implies f(\pi) = (\bot, \bot, \bot)$ • $I_2 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 2 \implies f(\pi) = (\text{Enqueue}, v_{\pi}, \bot)$ • $I_3 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 3 \implies f(\pi) \in \{\text{ENOUEUE}\} \times \{v_{\pi}\} \times \{ack, \bot\} \land (1 \le i_{\pi} < X) \land (O[i_{\pi}] = \bot)$ $\wedge (\forall \pi' \in \Pi - \{\pi\} : pc_{\pi'} \in \{3,4\} \implies i_{\pi'} \neq i_{\pi})$ • $I_4 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 4 \implies f(\pi) \in \{\text{Enqueue}\} \times \{v_{\pi}\} \times \{ack, \bot\} \land (1 \le i_{\pi} < X)$ $\wedge (\forall \pi' \in \Pi - \{\pi\} : pc_{\pi'} \in \{3, 4\} \implies i_{\pi'} \neq i_{\pi})$ • $I_6 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 6 \implies f(\pi) = (\text{Dequeue}, \bot, \bot)$ • $I_7 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 7 \implies f(\pi) = (\text{Dequeue}, \bot, \bot) \land 1 \le l_{\pi} \le X$ • $I_8 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 8 \implies f(\pi) = (\text{Dequeue}, \bot, \bot) \land 1 \le j_{\pi} < l_{\pi} \le X$ • $I_9 \equiv \forall \pi \in \Pi : \forall (\sigma, f) \in \mathcal{M} : pc_{\pi} = 9 \implies f(\pi) \in \{\text{Dequeue}\} \times \{\bot\} \times (\mathbb{N}^+ \cup \{\bot\})$ • $I_a \equiv \forall k \in \mathbb{N}^+ : k > X - 1 \implies O[k] = \bot$ • $I_b \equiv \forall A \subseteq [X-1] : GoodEngSet(A) \implies (\forall s \in Perm(A) : (JInvSeq(s))$ \implies ($\exists C \in \mathcal{M} : GoodRes(A, C) \land ValuesMatchInds(s, C))))$ • GoodEngSet(A) $\triangleq \forall k \in [X-1]: (Q[k] \neq \bot \implies k \in A) \land ((Q[k] = \bot \land k \in A) \implies \exists \pi' \in \Pi: pc_{\pi'} = 3 \land i_{\pi'} = k)$ • $JInvSeq(s) \triangleq \forall m, n \in [|s|] : (n < m \land s_m < s_n \land Q[s_m] \neq \bot) \implies (\exists \pi' \in \Pi : pc_{\pi'} = 8 \land s_n < l_{\pi'} \land s_m < j_{\pi'})$ • GoodRes(A, C) $\triangleq \forall \pi' \in \Pi : C.f(\pi').res = \begin{cases} ack & \text{if } (pc_{\pi'} = 3 \land i_{\pi'} \in A) \lor pc_{\pi'} = 4 \\ x_{\pi'} & \text{if } pc_{\pi'} = 9 \\ \bot & \text{otherwise.} \end{cases}$ • *ValuesMatchInds*(*s*, *C*) \triangleq *C*. $\sigma = (\alpha_1, \alpha_2, \dots, \alpha_{|s|})$ where $\forall k \in [|s|]$ $\alpha_k = \begin{cases} Q[s_k] & \text{if } Q[s_k] \neq \bot \\ v_{\pi'} & \text{where } \pi' \in \Pi : pc_{\pi'} = 3 \land i_{\pi'} = s_k. \end{cases}$

Fig. 5. Invariant I of $\mathcal{A}(\overline{O})$, where \overline{O} is the implementation of the queue tracker in Figure 4.

Herlihy-Wing Queue Machine Verification in TLAPS

* Proof of full inductive invariant
THEOREM InductiveInvariant == Spec => []Inv
BY InductiveInvariantNL, LinearizabilityFromInvNL, PTL DEF Inv, InvNL

* Proof of linearizability
THEOREM Linearizability == Spec => [](M # {})
BY InductiveInvariant, PTL DEF Inv, Linearizable

THEOREM InductiveInvariantNL == Spec => []InvNL <1> USE AckBotDef DEF InvNL, PosInts <1> SUFFICES /\ (Init => InvNL) /\ (InvNL /\ [Next]_vars => InvNL') **BY PTL DEF Spec** <1>1. Init => InvNL <2> SUFFICES ASSUME Init PROVE InvNL OBVIOUS <2> USE DEF Init <2>1. Type0K BY InitTypeSafety <2>2. Inv_E2 BY DEF Inv E2 <2>3. Inv_E3 BY DEF Inv_E3 <2>4. Inv D2 BY DEF Inv_D2 <2>5. Inv_D3 BY DEF Inv_D3 <2>6. Inv Q BY DEF Inv_Q <2>7. Inv_Main <3> SUFFICES ASSUME NEW A \in SUBSET (1..(X-1)), GoodEngSet(A), NEW seq $\ Perm(A)$, JInvSeq(seq) PROVE \E c \in M : /\ ValuesMatchInds(seq, c.sigma) /\ GoodRes(A, c.fres) BY Zenon DEF Inv_Main <3>1. A = {} BY DEF GoodEngSet <3>2. Cardinality(A) = 0 BY <3>1, Zenon DEF Cardinality <3>3. seq = << >> BY <3>1, <3>2 DEF Perm, JInvSeq <3> 0ED BY <3>1, <3>3 DEF ValuesMatchInds, GoodRes <2> QED BY <2>1, <2>2, <2>3, <2>4, <2>5, <2>6, <2>7 DEF InvNL <1>2. InvNL /\ [Next]_vars => InvNL' <2> SUFFICES ASSUME InvNL /\ [Next]_vars

Conclusion

- Machine-verification of concurrent data structures can help avoid errors in mission critical deployments.
- We introduced Meta-Configurations Tracking: the first universal, sound, and complete, forward-reasoning proof technique for linearizability
 - Easy to express proofs of even complex future-dependent algorithms (e.g., Herlihy-Wing queue, Jayanti's Snapshot)
 - We have verified widely used concurrent objects (e.g., Jayanti-Tarjan Union-Find objects)
 - Proved using TLA+ Proof System
- <u>Ongoing work:</u> We are expanding the technique to variants of Linearizability
- We would love to see these techniques used to machine-certify more important and interesting concurrent algorithms
- We'd love to see support for Meta-Configurations Tracking in verification tools

Links

github.com/google/graph-mining

[POPL 24] Jayanti, Jayanti, Yavuz, Hernandez

Meta-Configuration Notation

