

Sound, Complete, and Tractable Linearizability Monitoring for Concurrent Collections

Michael Emmi

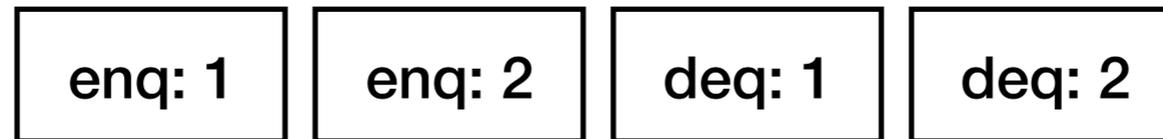
SRI International

Constantin Enea

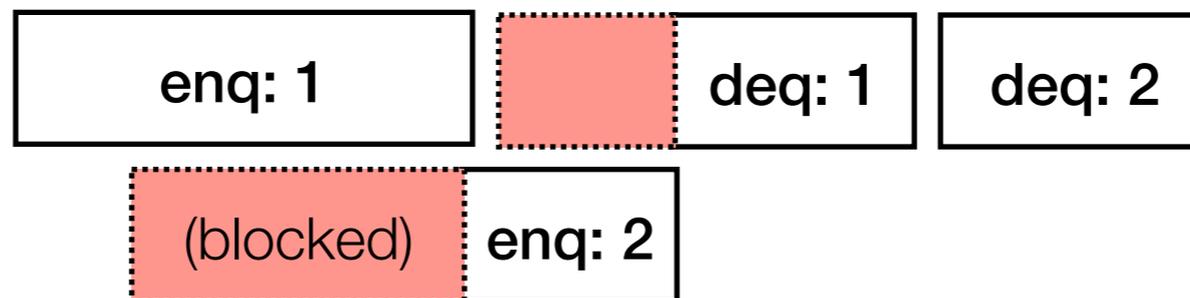
Université Paris 7

Concurrent Objects

Abstract data type (Queue)



Blocking reference implementation



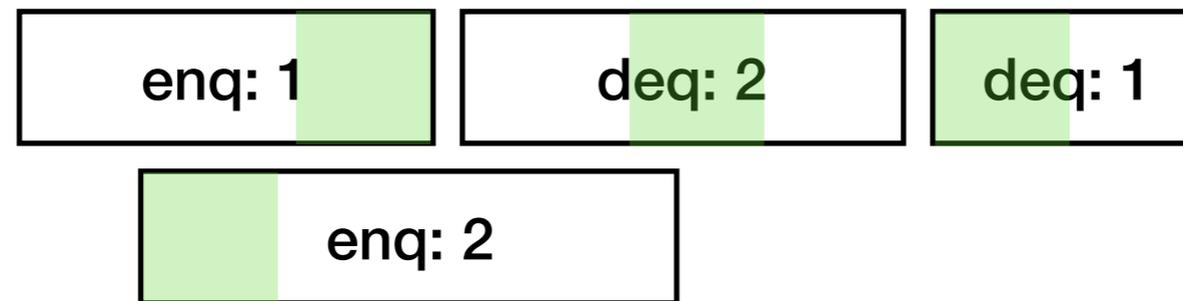
Efficient nonblocking implementation



Linearizability

Effects of each invocation appear to occur instantaneously

Execution history



Linearization admitted by Queue type



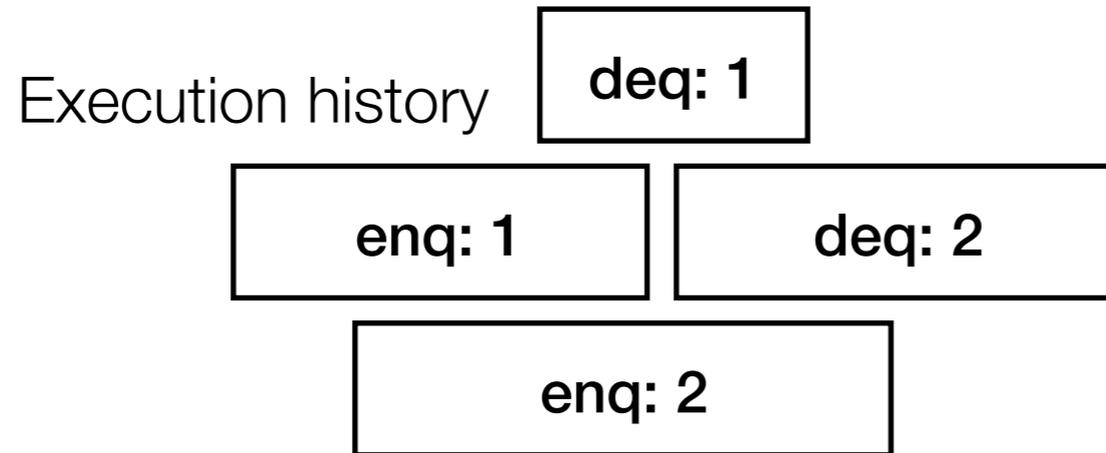
Theorem (Herlihy and Wing, 1990)

Linearizability implies observational refinement

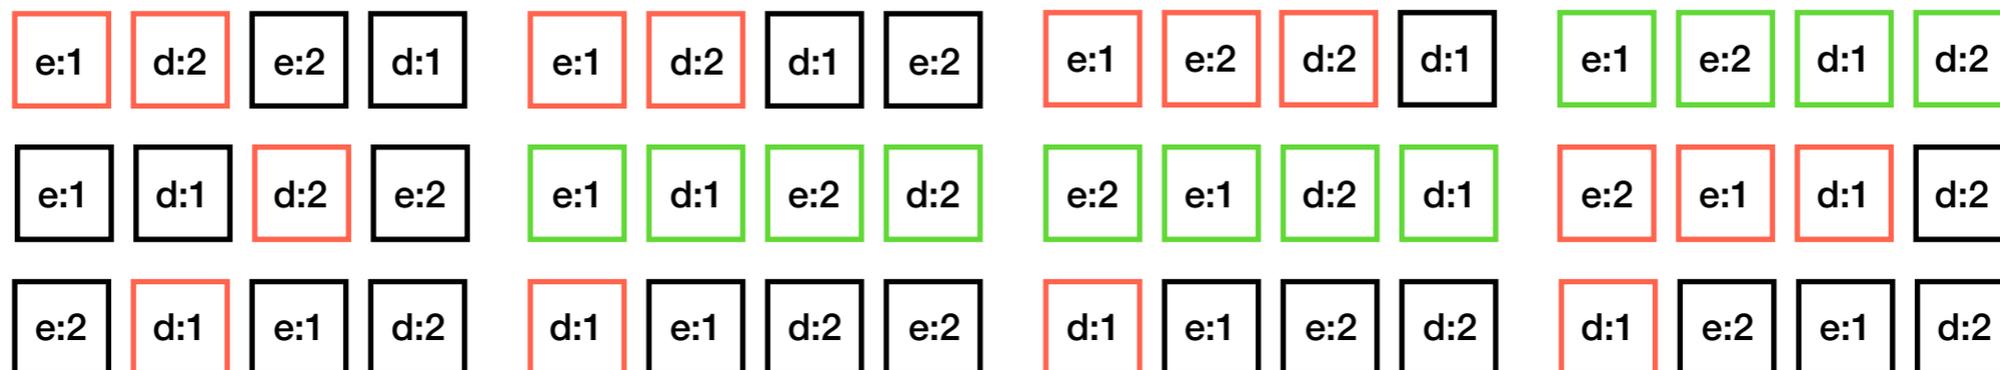
Theorem (Filipovic et al, 2010)

Observational refinement implies linearizability

Hardness



Exponentially many linearizations to consider

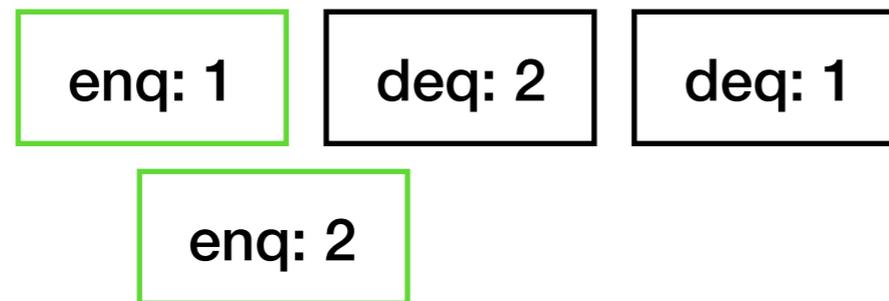


Theorem (Gibbons and Korach, 1997)
Checking linearizability is NP-hard

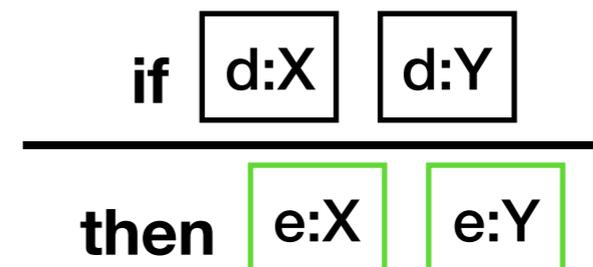
Result

Replace enumeration by monotonic deductive inference

Execution history / partial linearization



Queue rule



(partial) linearization

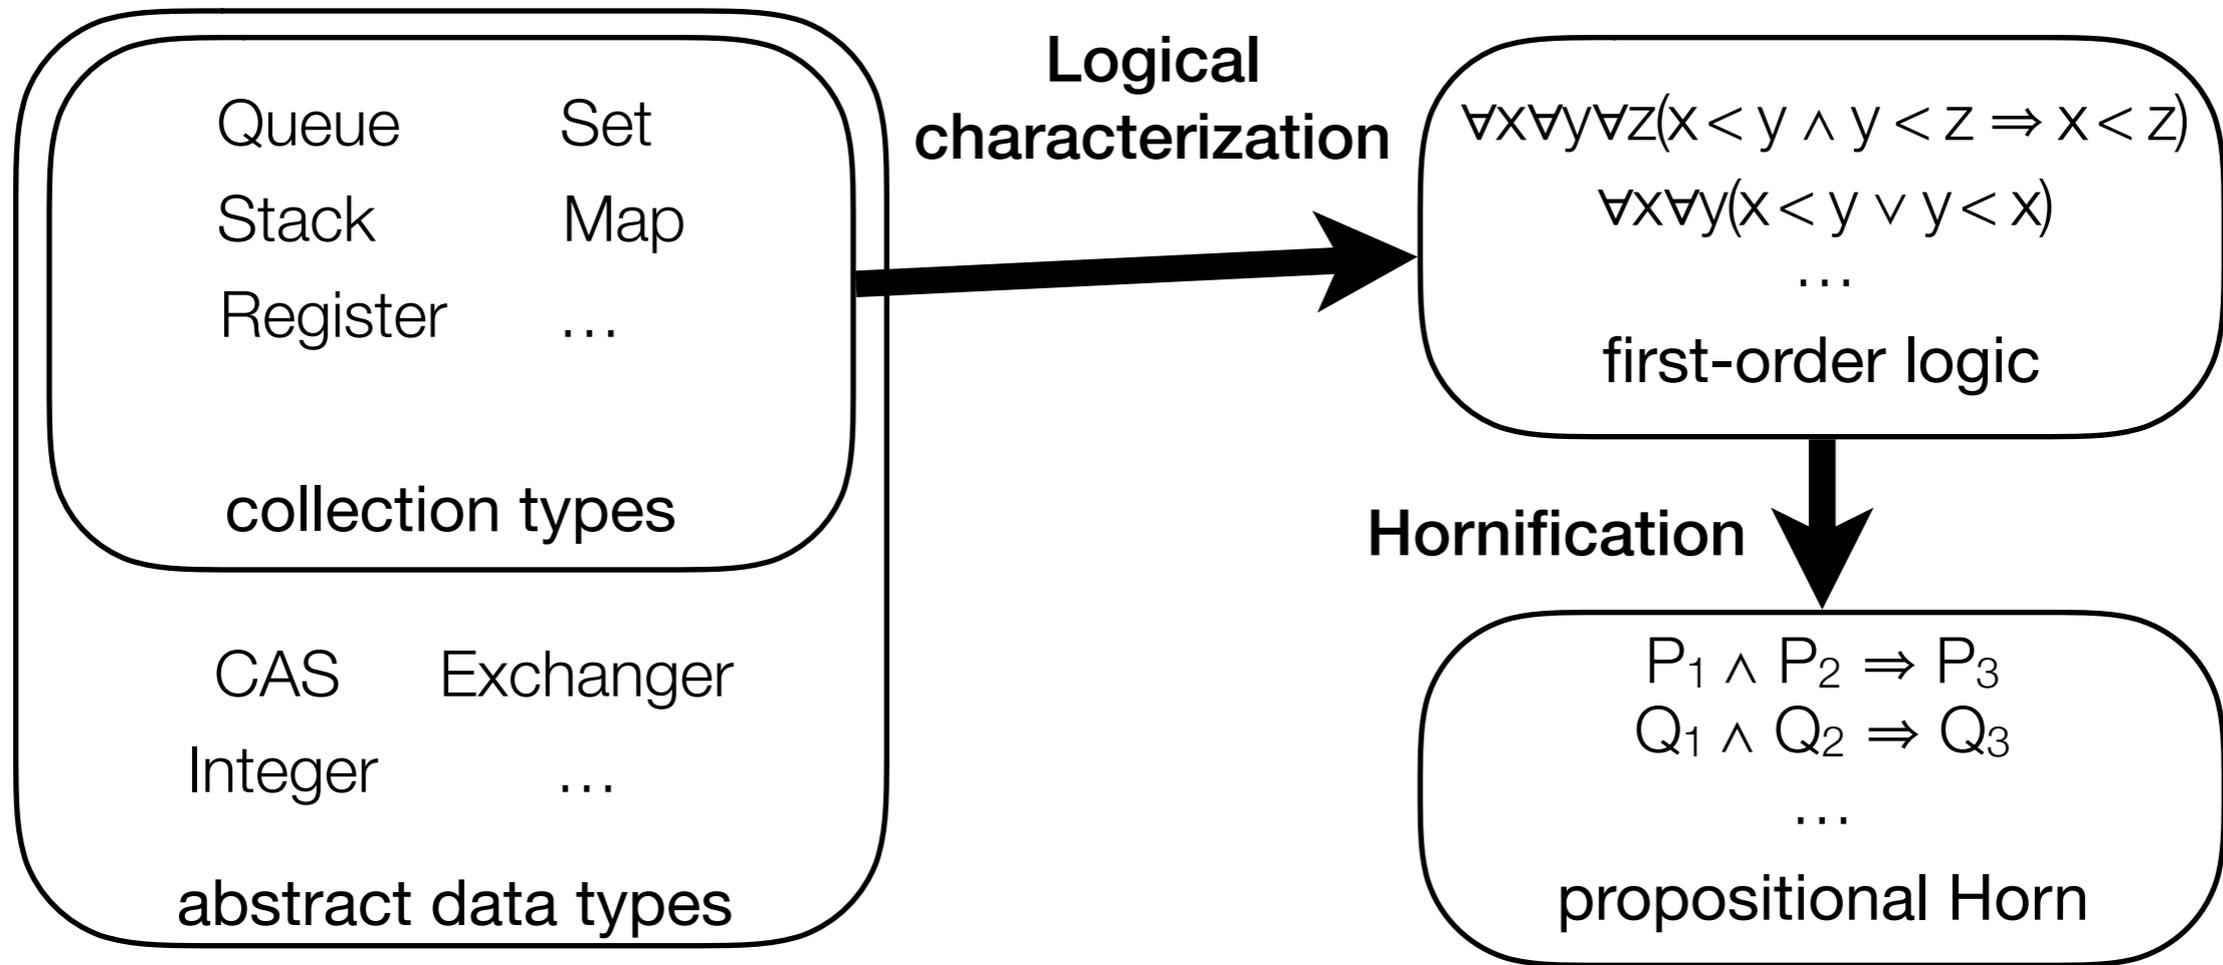


Theorem

Checking linearizability is polynomial-time for *collection types*

Approach

Reduction to logical satisfiability problem

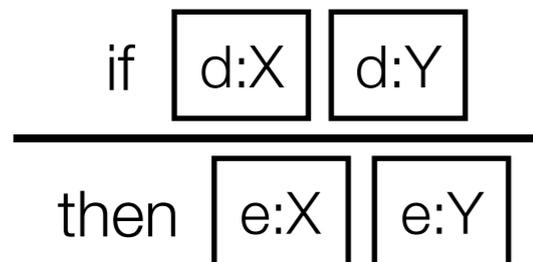


Theorem (Dowling and Gallier, 1984)

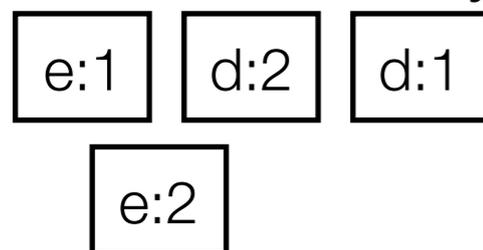
Horn satisfiability is solvable in linear time

Logical Characterization

Queue rules



Execution history



Total order axioms

$$\forall x \forall y \forall z (x < y \wedge y < z \Rightarrow x < z)$$

$$\forall x \forall y (x = y \vee x < y \vee y < x)$$

...

Queue axioms

$$\forall x_1 \forall x_2 \forall y_1 \forall y_2 (\mathbf{match}(x_1, x_2) \wedge \mathbf{match}(y_1, y_2) \wedge x_1 < y_1 \Rightarrow x_2 < y_2)$$

$$\wedge \dots$$

Ground formula

$$e:1 < d:2 \wedge d:2 < d:1 \wedge \mathbf{match}(e:1, d:1) \wedge \mathbf{match}(e:2, d:2)$$

$$\wedge \dots$$

Theorem

Linearizability equivalent to satisfiability
for *unambiguous* histories

ambiguous



unambiguous



Hornification

Original clause

$$A < B \vee B < C \vee \neg(A < B)$$

Translated clauses

$$A < B \vee \neg(\mathbf{C < B}) \vee \neg(A < B)$$

$$\neg(\mathbf{B < A}) \vee B < C \vee \neg(A < B)$$

$$\neg(\mathbf{B < A}) \vee \neg(\mathbf{C < B}) \vee \neg(A < B)$$

$$\neg(\mathbf{B < A}) \vee \neg(\mathbf{C < B}) \vee \mathbf{B < A}$$

Theorem

Translation is equisatisfiable and polynomial-time computable

Collection Types

Parametricity

{3,5} deq:3 {5}

{7,5} deq:7 {5}

Locality

{3,5} deq:3 {5}

{3,5} size:2 {3,5}

Value invariance

{3,5} deq:3 {5}

{3,5} sum:8 {}

Reducibility

{3,5,7} d:3; d:5; d:7 {}



{3,5} d:3; d:5 {}

{3,7} d:3; d:7 {}

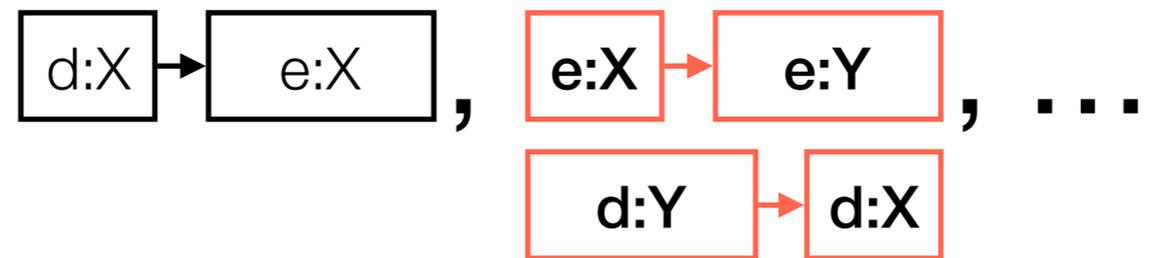
{5,7} d:5; d:7 {}

Theorem

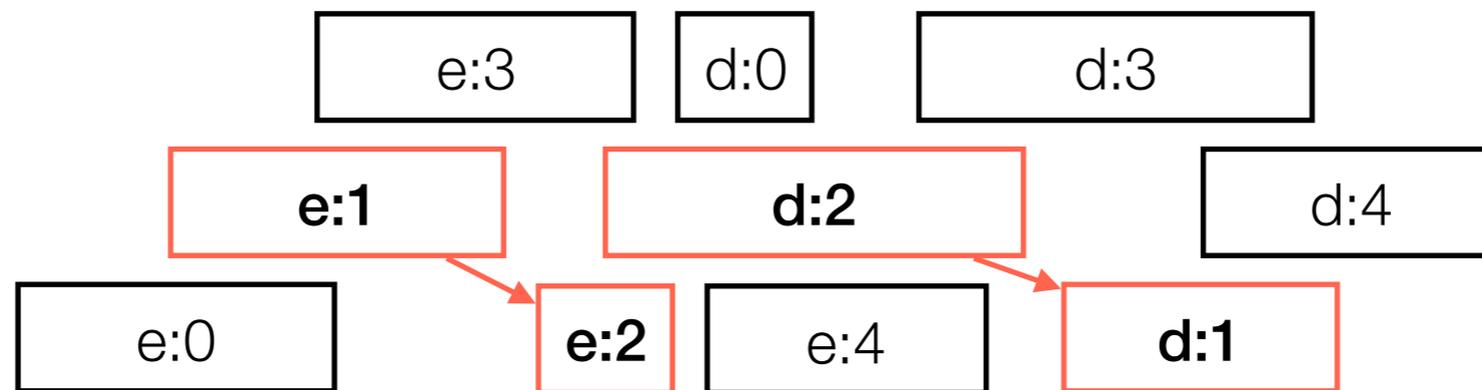
Restriction to unambiguous histories is sound for collections

Bounded Violations

Minimal violations



Minimal violation embedding

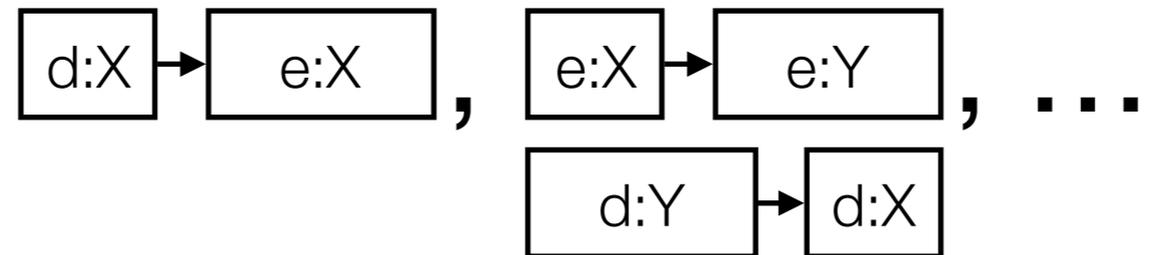


Theorem

Collections have bounded minimal violations

Logical Representation

Minimal violations



Quantified conjunction of negations

$\forall x_1 \forall x_2 \forall y_1 \forall y_2$

$\neg(\mathbf{m}(x_1, x_2) \wedge x_2 < x_1)$

$\wedge \neg(\mathbf{m}(x_1, x_2) \wedge \mathbf{m}(y_1, y_2) \wedge x_1 < y_1 \wedge y_2 < x_2)$

$\wedge \dots$

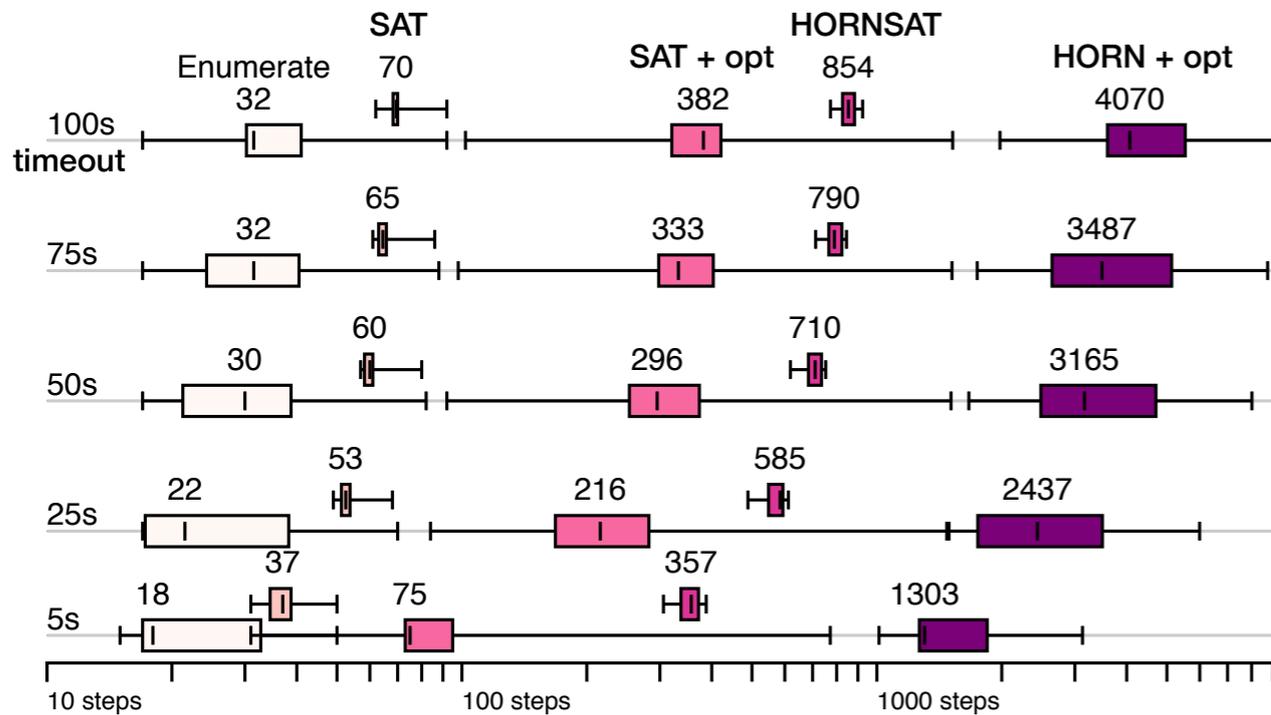
Theorem

Collections have sound logical specifications

Corollary

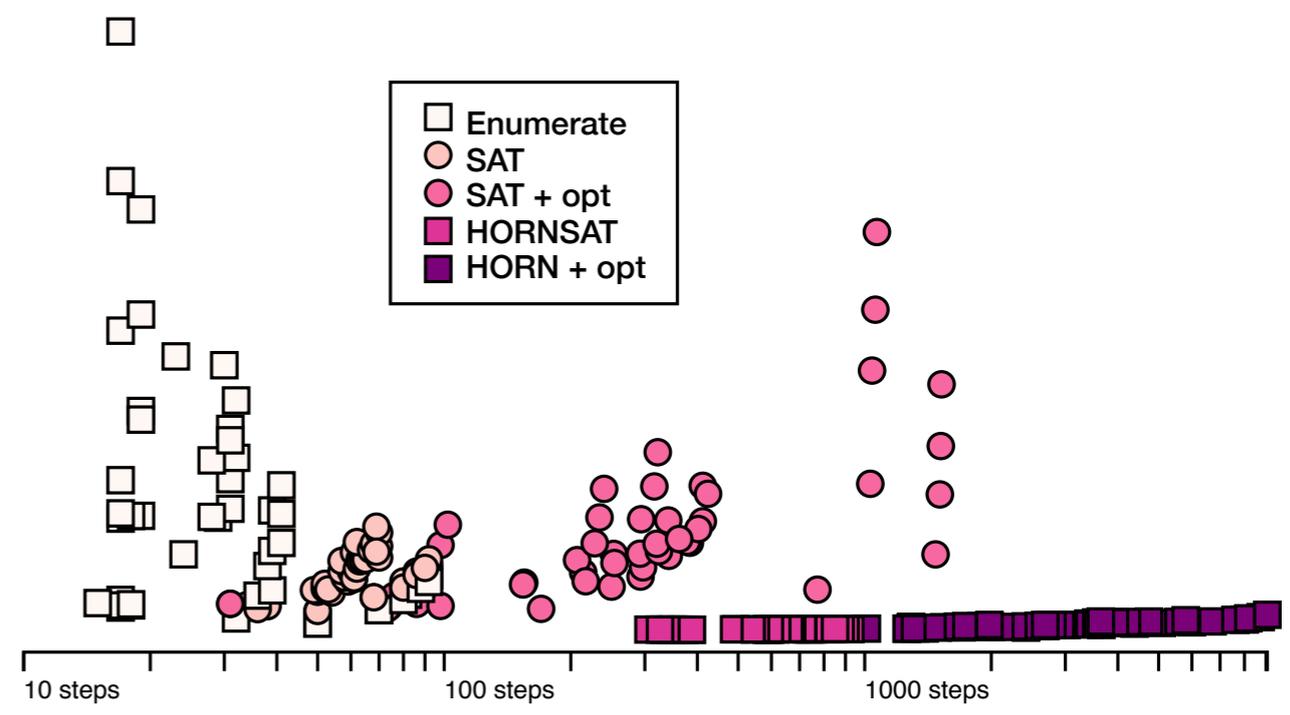
Linearizability is PTIME computable for collections

Empirical Evaluation



Performance

Over 10 histories with 10K steps each



Scalability

Normalized by (collection size)²

Observation

Orders of magnitude speedup

Related Work

Exponential enumeration

Wing and Gong, 1993. *Testing and Verifying Concurrent Objects*

NP Hardness

Gibbons and Korach, 1997. *Testing shared memories*

Asymptotically-equivalent optimizations

Burckhardt et al, 2010. *Line-up: a complete and automatic linearizability checker*

Shacham et al, 2011. *Testing atomicity of composed concurrent operations*

Horn et al, 2015. *Faster linearizability checking via P-compositionality*

Lowe, 2016. *Testing for linearizability*

Tractable approximation

Bouajjani et al, 2015. *Tractable refinement checking for concurrent objects*

Logical characterization

Emmi et al, 2015. *Monitoring refinement via symbolic reasoning*

Logical specification inference

Emmi et al, 2016. *Symbolic abstract data type inference*