

The Anchor Verifier for Blocking and Non-blocking Concurrent Software: Supplementary Appendix

CORMAC FLANAGAN, University of California, Santa Cruz, USA
STEPHEN N. FREUND, Williams College, USA

Verifying the correctness of concurrent software with subtle synchronization is notoriously challenging. We present the ANCHOR verifier, which is based on a new formalism for specifying synchronization disciplines that describes both (1) what memory accesses are permitted, and (2) how each permitted access commutes with concurrent operations of other threads (to facilitate reduction proofs). ANCHOR supports the verification of both lock-based blocking and cas-based non-blocking algorithms. Experiments on a variety of concurrent data structures and algorithms show that ANCHOR significantly reduces the burden of concurrent verification.

CCS Concepts: • **Theory of computation** → **Program verification**; **Program specifications**; • **Software and its engineering** → **Concurrent programming languages**; **Formal software verification**;

Additional Key Words and Phrases: concurrent program verification, reduction, synchronization

ACM Reference Format:

Cormac Flanagan and Stephen N. Freund. 2020. The Anchor Verifier for Blocking and Non-blocking Concurrent Software: Supplementary Appendix. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 156 (November 2020), 34 pages. <https://doi.org/10.1145/3428224>

A SUPPLEMENTARY APPENDIX

A.1 Generic Reduction Theorem

We start by presenting a generic reduction theorem that extends Theorem 1 of [Flanagan and Qadeer 2003]. The key extension is that we now consider two operations to commute even if swapping them could introduce new errors (but never hide old errors), as any errors from a preemptive execution will still also be present in the corresponding cooperative execution obtained via repeated swapping.

Given a set of states $State$, with error states $E \subseteq State$, and transition relations $\rightarrow_i, \rightarrow_j \subseteq State \times State$, we say that \rightarrow_i commutes with \rightarrow_j with respect to E if whenever $s_1 \rightarrow_i s_2 \rightarrow_j s_3$ then either:

- (1) $\exists e \in E. s_1 \rightarrow_j e$, or
- (2) $\exists s_4 \in State, e \in E. s_1 \rightarrow_j s_4 \rightarrow_i e$, or
- (3) $\exists s_4 \in State. s_1 \rightarrow_j s_4 \rightarrow_i s_3$.

Thus, swapping \rightarrow_i and \rightarrow_j transitions either does not affect the final state s_3 or introduces a new error state e .

Authors' addresses: Cormac Flanagan, University of California, Santa Cruz, Santa Cruz, CA, USA; Stephen N. Freund, Williams College, Williamstown, MA, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/11-ART156

<https://doi.org/10.1145/3428224>

For $S \subseteq State$, we define

$$\begin{aligned} S/\rightarrow_i &= (\rightarrow_i) \cap (S \times State) \\ \rightarrow_i \setminus S &= (\rightarrow_i) \cap (State \times S) \end{aligned}$$

Theorem 6 (Generic Reduction). For all i , let R_i, L_i , and E_i be sets of states and \rightarrow_i be a transition relation. Suppose, for all i ,

- (1) R_i, L_i , and E_i are pairwise disjoint.
- (2) $(L_i/\rightarrow_i \setminus R_i)$ is false.

and for all $j \neq i$,

- (3) \rightarrow_i and \rightarrow_j are disjoint.
- (4) $(\rightarrow_i \setminus R_i)$ commutes with \rightarrow_j with respect to E .
- (5) \rightarrow_j commutes with (L_i/\rightarrow_i) with respect to E .
- (6) if $p \rightarrow_i q$, then $R_j(p) \Leftrightarrow R_j(q), L_j(p) \Leftrightarrow L_j(q), E_j(p) \Leftrightarrow E_j(q)$.
- (7) $\forall q \in L_i. q \rightarrow_i^* q' \in N_i$, i.e., q is non-blocking.
- (8) $E_i/\rightarrow_i = \emptyset$.

Let

$$\begin{aligned} N_i &= \neg(R_i \vee L_i) \\ N &= \forall i. N_i \\ E &= \exists i. E_i \\ \rightarrow &= \exists i. \rightarrow_i \\ \twoheadrightarrow &= \exists i. (\forall j \neq i. N_j)/\rightarrow_i \end{aligned}$$

Suppose $p \in N$ with $p \not\rightarrow^* E$. Then

- (1) $p \not\rightarrow^* E$, and
- (2) For all $q \in N$, if $p \rightarrow^* q$ then $p \twoheadrightarrow^* q$.

PROOF. The proof of this theorem follows the structure of the proof of Theorem 1 in [Flanagan and Qadeer 2003], extended to permit the reduced trace to go wrong more often than the original.

We begin by defining several terms.

A *pre-commit sequence* by thread i is a sequence of states starting in N_i where each transition is a right-mover. We summarize the beginning and ending states of a pre-commit sequence as

$$Pre_i = N_i / ((\rightarrow_i \setminus R_i)^*)$$

A *post-commit sequence* by thread i is a pre-commit sequence followed by a committing transition \rightarrow_i , followed by left-movers (L_i/\rightarrow_i) .

$$Post_i = Pre_i; (\rightarrow_i \setminus L_i); (L_i/\rightarrow_i)^*$$

where $;\cdot$ denotes the composition of relations, i.e.:

$$Pre_i; \rightarrow_i = \{ (s, s'') \mid \exists s'. (s, s') \in Pre_i \wedge (s', s'') \in \rightarrow_i \}$$

A *finished sequence* by thread i is a post-commit sequence that ends in N_i :

$$Finish_i = Post_i \setminus N_i$$

We note that $Pre_i, Post_i, Finish_i \subseteq State \times State$. Let

$$\begin{aligned} Pre &= \exists i. Pre_i \\ Post &= \exists i. Post_i \\ Finish &= \exists i. Finish_i \end{aligned}$$

We first prove that

$$\text{If } (p, q) \in \text{Post}^*; \text{Pre}^* \text{ then } q \notin E \quad (1)$$

Let $L(q) = \{ i \mid q \in L_i \}$. Proof is by induction on $|L(q)|$.

- Base Case: Suppose $L(q) = \emptyset$, and assume $q \in E_i$ for some i . Then $(p, q) \in \text{Finish}^*; \text{Pre}_i$. Furthermore, since a Pre_i step cannot end in E_i , $\exists q'. (p, q') \in \text{Finish}^* \wedge q' \in E_i$. Hence, $p \xrightarrow{*} q'$, contradicting the assumption that $p \not\xrightarrow{*} q'$.
- Inductive Case: Assume $(p, q) \in \text{Post}^*; \text{Pre}^*$ and $j \in L(q)$, i.e. $q \in L_j$. By the non-blocking assumption, $q \xrightarrow{*} q' \in N_j$. Thus, we can left-commute these $q \xrightarrow{*} q'$ steps to the end of the last Post_j block to yield one of the following cases:
 - $(p, q') \in \text{Post}^*; \text{Pre}^*$ and $L(q') \subset L(q) \setminus \{j\}$ if no commuting steps yield an error. By induction, $q' \notin E$, and thus $q \notin E$.
 - If a commuting step yields an error state $q'' \in E$, then $L(q'') \subset L(q)$, but by induction this is impossible.

We next prove that

$$\text{If } p \xrightarrow{*} q \text{ then } (p, q) \in \text{Post}^*; \text{Pre}^* \quad (2)$$

The proof is by induction on the length of sequence $p \xrightarrow{*} q$.

- Base Case: If $p = q$, this is trivial.
- Inductive Case: Assume $p \xrightarrow{*} q \xrightarrow{-i} q'$ and $(p, q) \in \text{Post}^*; \text{Pre}^*$. There are several cases to consider:
 - $(q, q') \in N_i \times (N_i \vee L_i)$. Then all Pre blocks right-commute over this $\xrightarrow{-i}$ step, and we can add this $\xrightarrow{-i}$ step as additional Post block, yielding $(p, q') \in \text{Post}^*; \text{Pre}^*$.
 - $(q, q') \in N_i \times R_i$. Then we can add this step as a new Pre block.
 - $q \in L_i$. Then the last i block must be a Post_i block, and $q \xrightarrow{-i} q'$ left commutes to the end of this Post_i block.
 - $q \in R_i$. Then the last i block must be a Pre_i block, and the subsequent fsPre blocks right-commute over $q \xrightarrow{-i} q'$. Thus we can add this step at the end of the Pre_i block.

We now prove the two parts of this theorem. To show part 1, assume $p \xrightarrow{*} q$. By (2), $(p, q) \in \text{Post}^*; \text{Pre}^*$. By (1), $q \notin E$.

To show part 2, assume $q \in N$ and $p \xrightarrow{*} q$. Then $(p, q) \in \text{Post}^*; \text{Pre}^*$ by (2). But since $q \in N$, $(p, q) \in \text{Post}^*$. Hence, each Post_i block must end in N_i , so $(p, q) \in \text{Finish}^*$, and $p \xrightarrow{*} q$. \square

A.2 Proofs for Section 6

We now leverage the generic reduction theorem to prove our central reduction argument for ANCHOR.

First, we prove that \Rightarrow -steps right-commute (or go wrong) in the PRE phase and left-commute (or go wrong) in the POST phase. Proving this property requires a notion like strict stability from Definition 5.2. That notion of stability is satisfied by many synchronization specifications. However, strict stability is sometimes overly restrictive. For example, it prohibits interleaved writes from increasing the access permissions of a thread. For maximal expressiveness, we instead derive a more relaxed notion of stability from the conditions necessary to prove the following commuting lemma for all pairs of possible execution steps.

Definition A.1 (Stability). *Specification* R, W is *stable* if for all l, k, t, H, v, w , and u where $u \neq t$:

A.1	$W_l(t, H, v) \sqsubseteq R \wedge W_k(u, H[l := v], w) \sqsubseteq E \Rightarrow W_k(u, H, w) \in \{W_k(u, H[l := v], w), E\}$
A.2	$W_l(t, H, v) \neq E \wedge W_k(u, H[l := v], w) \not\sqsubseteq L \Rightarrow W_k(u, H, w) \not\sqsubseteq L$
A.3	$k \neq l \wedge W_l(t, H, v) \neq E \wedge W_k(u, H[l := v], w) \sqsubseteq L \Rightarrow W_k(u, H, w) \in \{W_k(u, H[l := v], w), E\}$
C	$W_l(t, H, v) \sqsubseteq E \wedge W_k(u, H, w) \sqsubseteq L \Rightarrow W_l(t, H[k := w], v) \in \{W_l(t, H, v), E\}$
D	$W_l(t, H, v) \sqsubseteq R \wedge W_k(u, H[l := v], w) \sqsubseteq N \Rightarrow W_l(t, H[k := w], v) \in \{W_l(t, H, v), E\}$
E	$W_l(t, H, v) \sqsubseteq N \wedge W_k(u, H[l := v], w) \sqsubseteq L \Rightarrow W_l(t, H[k := w], v) \in \{W_l(t, H, v), E\}$
F	$R_l(t, H) \sqsubseteq R \wedge W_k(u, H, w) \sqsubseteq N \Rightarrow R_l(t, H[k := w]) \in \{R_l(t, H), E\}$
H	$R_l(t, H, v) \sqsubseteq E \wedge W_k(u, H, w) \sqsubseteq L \Rightarrow R_l(t, H[k := w]) \in \{R_l(t, H), E\}$
I	$k \neq l \wedge R_l(t, H) \sqsubseteq N \wedge W_k(u, H, w) \sqsubseteq N \Rightarrow R_l(t, H[k := w]) \in \{R_l(t, H), E\}$
J	$W_l(t, H, v) \sqsubseteq R \wedge R_k(u, H) \sqsubseteq E \Rightarrow R_k(u, H) \in \{R_k(u, H[l := v]), E\}$
K	$W_l(t, H, v) \sqsubseteq N \wedge R_k(u, H[l := v]) \sqsubseteq L \Rightarrow R_k(u, H) \in \{R_k(u, H[l := v]), E\}$
L	$W_l(t, H, v) \sqsubseteq N \wedge R_k(u, H[l := v]) \not\sqsubseteq L \Rightarrow R_k(u, H) \not\sqsubseteq L$
M	$W_l(t, H, v) \sqsubseteq N \wedge W_k(u, H, w_1) \sqsubseteq E \Rightarrow W_k(u, H[k := w_1], w_2) \in \{W_k(u, H[l := v][k := w_1], w_2), E\}$
N	$W_l(t, H, v_1) \sqsubseteq E \wedge W_k(u, H, w) \sqsubseteq L \Rightarrow W_l(t, H[k := w][l := v_1], v_2) \in \{W_l(t, H[l := v_1], v_2), E\}$

Shaded terms are trivially true but included for symmetry among the rules. Rules M and N are only applicable if l or k is manipulated via a cas operation.

Lemma 1 (Commuting). Suppose R, W is a valid and stable synchronization discipline, $u \neq t$, and

$$\Pi_1 \Rightarrow_t \Pi_2 \Rightarrow_u \Pi_3$$

where $\Pi_2 = T_2 \cdot H_2 \cdot C_2 \cdot P_2$, and $P_2(t) = \text{PRE}$ or $P_2(u) = \text{POST}$. Then:

- (1) $\exists \Pi_4. \Pi_1 \Rightarrow_u \Pi_4 \Rightarrow_t \Pi_3$, or
- (2) $\exists \Pi_4, \Pi_5. \Pi_1 \Rightarrow_u \Pi_4 \Rightarrow_t \Pi_5$ and t is wrong in Π_5 , or
- (3) $\exists \Pi_4. \Pi_1 \Rightarrow_u \Pi_4$ and u is wrong in Π_4 .

PROOF. By case analysis on all pairs of possibly conflicting operations. \square

Next, we define the predicates $\mathbb{R}_i, \mathbb{L}_i, \mathbb{E}_i, \mathbb{N}_i, \mathbb{N} \subseteq \Pi$ over an instrumented state $\Pi = T \cdot H \cdot CP$:

$$\begin{aligned} \mathbb{R}_i(\Pi) &= P(i) = \text{PRE} \wedge \neg \mathbb{N}_i(\Pi) \\ \mathbb{L}_i(\Pi) &= P(i) = \text{POST} \wedge \neg \mathbb{N}_i(\Pi) \\ \mathbb{E}_i(\Pi) &= \text{thread } i \text{ is wrong in } \Pi \\ \mathbb{N}_i(\Pi) &= \text{thread } i \text{ is yielding in } \Pi \\ \mathbb{N}(\Pi) &= \forall i. \mathbb{N}_i(\Pi) \end{aligned}$$

These predicates characterize whether thread i is in the pre-commit, or right-mover, part of a reducible code sequence (\mathbb{R}_i); in the post-commit, or left-mover, part (\mathbb{L}_i); gone wrong (\mathbb{E}_i); or is yielding (\mathbb{N}_i).

We can now formalize the *nonblocking* requirement that any transaction that has passed its commit point must be able to terminate, i.e. if $\Pi_0 \Rightarrow^* \Pi'$ and $\mathbb{L}_i(\Pi')$ then $\Pi' \Rightarrow_i^* \Pi''$ such that $\mathbb{N}_i(\Pi'')$.

Restatement of Theorem 3 (Reduction). If $\Pi = T_0 \cdot H_0 \cdot C_0 \cdot P_0$ and Π_0 does not go wrong under \Rightarrow , then:

- (1) Π_0 does not go wrong under \Rightarrow .
- (2) If $\Pi_0 \Rightarrow^* \Pi$ where Π is yielding, then $\Pi_0 \Rightarrow^* \Pi$.

PROOF. By Theorem 6, where the preconditions of that theorem are satisfied as follows. For all $i, j \in \text{ThreadId}$ where $i \neq j$:

- (1) $\mathbb{R}_i, \mathbb{L}_i$, and \mathbb{E}_i are pairwise disjoint by construction.

- (2) $\mathbb{L}_i / \Rightarrow_i \setminus \mathbb{R}_i$ is false. \mathbb{L}_i says that $P(i) = \text{POST}$ and thread i is not yielding. \mathbb{R}_i says that $P(i) = \text{PRE}$. However, the only transition of thread i from POST to PRE is for yield .
- (3) \Rightarrow_i and \Rightarrow_j are disjoint.
- (4) $\Rightarrow_i / \mathbb{R}_i$ commutes with \Rightarrow_j with respect to E_i . By Lemma 1.
- (5) \Rightarrow_j commutes with $\mathbb{L}_i \setminus \Rightarrow_i$ with respect to E_i . By Lemma 1.
- (6) \Rightarrow_i does not change the phase or statements of other threads.
- (7) By the non-blocking assumption.
- (8) By definition of \Rightarrow_i .

□