

AWAMOCHÉ: Proof of Correctness

1 OPERATIONAL STEPS

READ

$$\frac{r \in \text{avail}(G) \cap R_I \setminus R_I^{\text{conf}} \quad w \in G.W_I \quad \neg(\text{spec}(r, w) \vee \text{excl}(r, w))}{\text{block}(r, w) \Rightarrow w = \max_{G.\text{co}}\{w' \in G.W_I\}}$$

$$G \xrightarrow{r@w} \text{AddRead}(G, r, w)$$

READ-SPEC

$$\frac{r \in \text{avail}(G) \cap R_I \setminus R^{\text{conf}} \quad w \in G.W_I \quad \text{spec}(r, w) \vee \text{excl}(r, w)}{\forall r' \in R_I^{\text{spec,excl}}. \text{completed}(G, r') \quad \nexists r' \in G.R_I^{\text{spec,excl}}. G.\text{rf}(r') = w}$$

$$G \xrightarrow{r@w} \text{AddRead}(G, r, w)$$

READ-BOT

$$\frac{r \in \text{avail}(G) \cap R_I \setminus R^{\text{conf}} \quad \forall r' \in R_r. \neg \text{bottom}(G, r)}{\exists r' \in G.R_I^{\text{spec,excl}}. \neg \text{completed}(G, r') \wedge (\text{spec}(r, G.\text{rf}(r')) \vee \text{excl}(r, G.\text{rf}(r')))}$$

$$G \xrightarrow{r@ \perp} \text{AddRead}(G, r, \perp)$$

READ-CONF

$$\frac{r \in \text{avail}(G) \cap R_I^{\text{conf}} \quad w = G.\text{rf}(\max_{G.\text{po}}\{e \in G.R_I^{\text{spec}} \mid \text{tid}(e) = \text{tid}(w)\})}{G \xrightarrow{r@w} \text{AddRead}(G, r, w)}$$

WRITE

$$\frac{w \in \text{avail}(G) \cap W_I \setminus W^{\text{excl,conf}} \quad w_p \in G.W_I}{G \xrightarrow{w@w_p} \text{SetRFs}(\text{AddWrite}(G, w, w_p), G.MBR_I, w)}$$

WRITE-RMW

$$\frac{w \in \text{avail}(G) \cap W^{\text{excl,conf}} \quad w_p = G.\text{rf}(\max_{G.\text{po}}\{e \in G.E \mid \text{tid}(e) = \text{tid}(w)\})}{G \xrightarrow{w@w_p} \text{SetRFs}(\text{AddWrite}(G, w, w_p), G.MBR_I \cup G.R_I^\perp, w)}$$

PLAIN

$$\frac{e \in \text{avail}(G) \setminus R \cup W}{G \xrightarrow{e@-} \text{Add}(G, e)}$$

2 MAXIMAL STEPS

READ

$$\frac{r \in R_I \setminus R^{\text{conf}} \quad G \xrightarrow{r@w} G' \quad w = \max_{G.\text{co}}\{e \in G.W_I\} \vee w = \perp}{G \rightsquigarrow_r G'}$$

WRITE

$$\frac{w \in W_I \setminus W^{\text{excl,conf}} \quad G \xrightarrow{w@w_p} G' \quad w_p = \max_{G.\text{co}}\{e \in G.W_I\}}{G \rightsquigarrow_w G'}$$

REST

$$\frac{e \in R^{\text{conf}} \cup W^{\text{excl,conf}} \cup (E \setminus R \cup W) \quad G \xrightarrow{e@-} G'}{G \xrightarrow{e} G'}$$

3 ALGORITHMIC STEPS

NON-REVISIT STEP

$$\frac{e \in \text{nextCandidates}(G) \quad \neg \text{racy}(G, e) \quad G \xrightarrow{e@p} G' \quad p \neq \perp \quad \text{consistent}(G')}{G \xrightarrow[nr]{e@p} G'}$$

WRITE REVISIT STEP

$$\frac{\begin{array}{l} w \in \text{nextCandidates}(G) \cap W_I \quad \neg \text{racy}(G, w) \quad r \in R_I \setminus R^{\text{block},\perp} \\ [d_1, \dots, d_n] = \text{sort}_{<}(\{e \in G.E \mid r < e \wedge \langle e, w \rangle \notin \text{Add}(G, w, -).\text{porf}\}) \\ G_\emptyset \sqsubseteq G'' \xrightarrow[r]{r} \xrightarrow[d_1]{d_1} \dots \xrightarrow[d_n]{d_n} G \quad G'' \xrightarrow[w@w_p]{w} \xrightarrow[r@w]{r} G' \quad \text{consistent}(G') \end{array}}{\langle G, \langle \rangle \xrightarrow[r]{w@w_p} \langle G', \text{Restrict}(\langle, G''.E \cup \{r\}) \rangle \rangle \dashv\vdash [w]}$$

SPEC REVISIT STEP

$$\frac{\begin{array}{l} r \in \text{nextCandidates}(G) \cap R_I^{\text{spec,excl}} \quad w \in W_I \quad \text{spec}(r, w) \vee \text{excl}(r, w) \\ r' \in G.R_I^{\text{spec,excl}} \quad G.\text{rf}(r') = w \\ [d_1, \dots, d_n] = \text{sort}_{<}(\{e \in G.E \mid r' < e \wedge \langle e, r \rangle \notin \text{Add}(G, r, w).\text{porf}\}) \\ G_\emptyset \sqsubseteq G'' \xrightarrow[r']{r'} \xrightarrow[d_1]{d_1} \dots \xrightarrow[d_n]{d_n} G \quad G'' \xrightarrow[r@w]{r} \xrightarrow[r'@\perp]{r'} G' \quad \text{consistent}(G') \end{array}}{\langle G, \langle \rangle \xrightarrow[r']{r@w} \langle G', \text{Restrict}(\langle, G''.E \cup \{r'\}) \rangle \rangle \dashv\vdash [r]}$$

4 DEFINITIONS

Definition 4.1 (Prefix). Given consistent executions G and G' , we say that G is a *prefix* of G' , and write $G \sqsubseteq G'$, if G can be extended to G' with a series of operational steps.

$$G \sqsubseteq G' \triangleq G \rightarrow^* G'$$

Definition 4.2 (Available prefix). Given a consistent execution G and an event $e \in \text{avail}(G)$, we define the *available prefix* of G w.r.t. e , as the restriction of G to the events in the $G.\text{porf}$ -prefix of the po-last event e' of the thread of e in G , i.e.,

$$\text{availPrefix}(G, e) \triangleq G|_{\text{dom}(G.\text{porf}; [e'])}, \text{ where } e' \triangleq \max_{G.\text{po}} \{e'' \in G.E \mid \text{tid}(e') = \text{tid}(e)\}.$$

Definition 4.3 (Full Execution). Given an execution G , we say that G is full, and write $\text{full}(G)$, if $\text{avail}(G) = \emptyset$.

Definition 4.4 (Pending Read). Given an execution G and a read $r \in G.R$, we say that r is pending in G , and write $\text{pending}(G, r)$, if either $r \in R^{\text{conf}}$, $G.\text{rf}(r) \neq \perp$, and $\neg \text{completed}(G, r)$, or $r \in R^{\text{excl}} \setminus R^{\text{conf}}$, $\text{rng}([r]; G.\text{po}) = \emptyset$, and $\text{val}(G.\text{rf}(r))$ makes the RMW operation succeed.

Definition 4.5 (Well-picked). Given an execution G and an event $e \in \text{avail}(G)$, we say that e is well-picked for G , and write $\text{well-picked}(G, e)$, if

$$(e \in W^I \vee (e \in R^I \wedge \text{spec}(e) \cup \text{excl}(e) \neq \emptyset)) \Rightarrow \nexists e' \in G'.R_I^{\text{spec,excl}}. \text{pending}(G, e')$$

Definition 4.6 (Next Candidates). Given an execution G , we define the set of next candidates

$$\text{nextCandidates}(G) \triangleq \{e \in \text{avail}(G) \mid \text{well-picked}(G, e)\}$$

4.1 Blocking

Definition 4.7. Given an execution G , we write $G.\text{BR}$ for the set of blocked read events.

$$G.\text{BR} \triangleq \{r \in G.R \mid \text{valw}(G.\text{rf}(r)) \in \text{blockv}(r)\}$$

Definition 4.8. Given an execution G , we write $G.MBR$ for the set of coherence-maximal blocked read events.

$$G.MBR \triangleq \{r \in G.BR \mid \text{rng}([r]; G.fr) = \emptyset\}$$

4.2 Confirmation

Definition 4.9. Given an execution G and a read $r \in G.R^{\text{Spec}}$, we say that r is completed in G , and write $\text{completed}(G, r)$, if r is followed by a same-location confirmation write, i.e.,

$$\text{completed}(G, r) \triangleq \exists e \in W_{\text{loc}(r)}^{\text{conf}}. \langle r, e \rangle \in G.po$$

4.3 Lemmas

LEMMA 4.10 (PROGRESS). *Given an execution $G \sqsupseteq G_0$, if $\text{avail}(G) \neq \emptyset$ then $\text{nextCandidates}(G) \neq \emptyset$.*

LEMMA 4.11 (MAXIMAL EXTENSIBILITY). *Given an execution $G \sqsupseteq G_0$ such that $\neg \text{racy}(G)$ and $e \in \text{nextCandidates}(G)$, there is a unique execution G' such that $G \overset{e}{\rightsquigarrow} G'$.*

LEMMA 4.12 (MAXIMAL CONSISTENCY). *Given two execution such that $G_0 \sqsubseteq G \rightsquigarrow G'$, if G is consistent and $\neg \text{racy}(G)$, then G' is consistent.*

LEMMA 4.13 (PREFIX CLOSEDNESS). *Given two executions G and G' s.t. $G \sqsubseteq G \sqsubseteq G'$, if G' is consistent, then G is consistent.*

LEMMA 4.14 (PREFIX EXTENSION). *Given two executions G and G' s.t. $G_0 \sqsubseteq G \sqsubset G'$, and an event $e \in G'.E \setminus G.E$, there is a unique execution G_p s.t. $G \sqsubseteq G_p \sqsubseteq G'$ and $G_p.E \setminus G.E = \text{dom}(G'.\text{porf}; [e]) \setminus G.E$. We denote call this execution the prefix-extension of G until the event e of G' , and denote it as $\text{PREFIXEXTENSION}(G, G', e)$.*

LEMMA 4.15 (NEXT PREFIX). *Given two executions G_s and G_t s.t. $G_0 \sqsubseteq G_s \sqsubset G_t$, and an event $e \in \text{avail}(G_s) \cap G_t.E$, either there exists a step $t = e @ _$ and an execution G' s.t. $G_s \xrightarrow{t} G' \sqsubseteq G_t$, or one of the following hold:*

- $e \in W$ and there is a step $t = e @ _$ and executions G' and G_r s.t. $G_s \xrightarrow{t} G' \sqsubseteq G_r$, $\text{full}(G_r)$, $\text{consistent}(G_r)$, and $\text{racy}(G_r)$.
- $e \in R$ and there is an execution G' s.t. $\text{PREFIXEXTENSION}(G_s, G_t, \text{event}(t)) \xrightarrow{t} \xrightarrow{e@p} G' \sqsubseteq G_t$, where $(t, p) = \text{GETREV}(G_s, G_t, e)$, and there is no execution G'' s.t. $e \in G''.E$ and $G_s \sqsubset G'' \sqsubset G'$.

4.4 Completeness

LEMMA 4.16. *Given a consistent, full execution G_f of P such that $G_f.BR = G_f.MBR$, and every confirmation CAS reads from the same write as the preceding speculative read, it is $G_0 \sqsubseteq G_f$.*

LEMMA 4.17. *Given a consistent, full execution $G_f \sqsupseteq G_0$, a call to $\text{PRODUCTIONSEQUENCE}(\emptyset, G_f)$ will return a production sequence S such that either $\text{APPLY}(S) = G_f$ or $\text{race}(S)$.*

PROOF. We will prove by induction that given a production sequence S such that $\neg \text{race}(S)$ and a full, consistent execution $G_f \sqsupseteq G_0$ s.t. $\text{APPLY}(S) \sqsubseteq G_f$, a call to $\text{PRODUCTIONSEQUENCE}(S, G_f)$ will return a production sequence S' and either $\text{APPLY}(S) = G_f$ and $S' = S$, or $\text{APPLY}(S) \neq G_f$, S' extends S , and $\text{APPLY}(S') \vee \text{race}(S')$. Then the lemma follows immediately by setting $S = \emptyset$.

The induction is on the length of S . The base case is $\text{APPLY}(S) = G_f$ (there is no extension of S), which is trivial.

Algorithm 1 Production sequence

```

1: procedure PRODUCTIONSEQUENCE( $S, G_f$ )
2:   while APPLY( $S$ )  $\neq G_f$  do
3:      $S \leftarrow \text{GETNEXT}(S, G_f)$ 
4:     if race( $S$ ) then return  $S$ 
5:   return  $S$ 

6: procedure GETNEXT( $S_0, G_t$ )
7:    $G_s \leftarrow \text{APPLY}(S_0)$ 
8:    $e \leftarrow \text{next}_p(G_s)$ 
9:   if racy( $G_s, e$ ) then return  $S_0 ++ R(e)$ 
10:  if  $\exists t = (e @ p). G_s \xrightarrow{t} G' \wedge G' \sqsubseteq G_t$  then return  $S_0 ++ \xrightarrow[nr]{t}$ 
11:  if  $e \in W$  then
12:    Let  $G'', \hat{G}$  be executions s.t.  $G_s \xrightarrow{e @ p} G'' \sqsubseteq \hat{G}$ ,  $\text{full}(\hat{G})$ , and  $\text{racy}(\hat{G})$ 
13:    return PRODUCTIONSEQUENCE( $S_0 ++ \xrightarrow[nr]{e @ p}, \hat{G}$ )
14:   $r_0 \leftarrow \text{next}_p(G_s)$ 
15:   $\langle e_0 @ p, p_{r_0} \rangle \leftarrow \text{GETREV}(G_t, r_0)$ 
16:   $G_p \leftarrow \text{PREFIXEXTENSION}(G_s, G_t, e_0)$ 
17:   $A \leftarrow \emptyset$ 
18:   $S \leftarrow S_0$ 
19:  while true do
20:     $G \leftarrow \text{APPLY}(S)$ 
21:     $e \leftarrow \text{next}_p(G)$ 
22:    if availPrefix( $G, e$ )  $\sqsubseteq G_p \wedge e \in G_p.E \cup \{e_0\}$  then
23:      if  $e = e_0$  then
24:        return  $S ++ \xrightarrow[rv \ r_0]{e_0 @ p}$ 
25:      else
26:         $A \leftarrow A ++ e$ 
27:         $G_{MC} \leftarrow \text{MAXCOMPLETION}(G_p, A)$ 
28:         $S \leftarrow \text{GETNEXT}(S, G_{MC})$ 
29:      if race( $S$ ) then return  $S$ 

30: procedure GETREV( $G_s, G_t, r$ )
31:  if  $G_t.\text{rf}(r) \in W^{\text{conf}} \wedge \text{spec}(r, G_t.\text{rf}(G_t.\text{spec}(G_t.\text{rf}(r))))$  then
32:     $r' \leftarrow G_t.\text{spec}(G_t.\text{rf}(r))$ 
33:    return  $\langle r' @ G_t.\text{rf}(r'), \perp \rangle$ 
34:  else
35:
36:     $w_t \leftarrow \max_{G.\text{co}} \{w \in G.W_{1\text{oc}}(r) \mid \langle w, G.\text{rf}(r) \rangle \in G.\text{co}? \wedge \text{block}(r, G.\text{co}|_{\text{imm}}^{-1}(w))\}$ 
37:     $G_p \leftarrow \text{PREFIXEXTENSION}(G_s, G_t, w_t)$ 
38:     $w_p \leftarrow \max_{G_p.\text{co}} \{w \in G_p.W_{1\text{oc}}(w_t)\}$ 
39:    return  $\langle w_t @ w_p, w_t \rangle$ 

```

For the inductive step, we assume a production sequence S such that $\neg \text{race}(S)$, $G_\emptyset \sqsubseteq \text{APPLY}(S) \sqsubset G_f$, $\text{full}(G_f)$, and $\text{consistent}(G_f)$, and prove that $\text{PRODUCTIONSEQUENCE}(S, G_f)$ will return a production sequence S' that extends S such that either $\text{APPLY}(S') = G_f$ or $\text{race}(S')$. The inductive hypothesis states that this implication holds for any production sequence S'' that is longer than S .

To prove the inductive step, we will show that given a production sequence S such that $\neg\text{race}(S)$ and a consistent execution G_t s.t. $G_0 \sqsubseteq G_s \sqsubset G_t$ and $\text{next}_P(G_s) \in G_t.E$, where $G_s \triangleq \text{APPLY}(S)$, a call to $\text{GETNEXT}(S, G_t)$ returns a production sequence S' that extends S such that either $\text{race}(S')$ or $\neg\text{race}(S')$, $G_s \sqsubset \text{APPLY}(S') \sqsubseteq G_t$ and $S' \setminus S$ does not revisit events of G_s . The inductive step then follows immediately: as long as a race is not detected (i.e., $\neg\text{race}(S)$), the calls to $\text{GETNEXT}(S, G_f)$ repeatedly grow S until it leads to G_f . The precondition $G_0 \sqsubseteq \text{APPLY}(S) \sqsubseteq G_f$ holds because for a production sequence $S' = \text{GETNEXT}(S, G_f)$ such that $\neg\text{race}(S')$, it is $\text{APPLY}(S) \sqsubset \text{APPLY}(S') \sqsubseteq G_f$. The precondition $\text{next}_P(G) \in G_f$ holds trivially for any full execution G_f and $G \sqsubseteq G_f$.

We will prove the claim by induction on the size of the arguments at each $\text{GETNEXT}(S, G_t)$ call, with $|\langle S, G_t \rangle| = \langle -|\text{APPLY}(S).E|, |G_t.E| \rangle$. Since G_s and G_t are consistent executions of a program with finite-size executions and $G_s \sqsubset G_t$, this measure is bounded below.

Consider a call $\text{GETNEXT}(S_0, G_t)$, where S_0 is a production sequence such that $\neg\text{race}(S_0)$, G_t is a consistent execution such that $G_0 \sqsubseteq G_s \sqsubset G_t$, and $\text{next}_P(G_s) \in G_t.E$, where $G_s \triangleq \text{APPLY}(S_0)$.

If the test in Line 9 succeeds, then $\text{GETNEXT}(S, G_t)$ will return $S' = S_0 \uparrow R(e)$, which extends S_0 and $\text{race}(S')$. Otherwise, it is not $\neg\text{racy}(G_s, e)$.

If the test in Line 10 succeeds, we only need to show that if $G_s \xrightarrow{e@p} G'$, then $G_s \xrightarrow[nr]{e@p} G'$. If $e \in W$, then $\text{VISIT}_P(G)$ does consider every **co** placement p . If $e \in R$, then the only **rf** choices not considered by the algorithm are the \perp , and writes w such that there is completed same-location speculative read that reads from w (if e is a speculative read). The latter cannot be the case by definition of \rightarrow .

For the former, assume that $G_s \xrightarrow{e@\perp} G'$. Since $e = \text{next}_P(G_s)$, there cannot exist any same-location speculative read r' that is not completed and is not reading \perp . But is also cannot be that there is an r' that reads \perp : it is $G_0 \sqsubseteq G_s$ and thus when r' was added to read \perp in this \rightarrow sequence, there was a same-location speculative r'' read that was reading from a write $w \neq \perp$. Since all speculative reads that do not read \perp are completed in G_s , this r'' must be followed by a confirmation write, which in-place revisited r' . Therefore, all speculative reads in G_s are completed, contradicting that $G_s \xrightarrow{e@\perp} G'$.

If the test in Line 10 does not succeed, from Lemma 4.15 there are two cases. The first case is that $e \in W$ and there is a step $t = e@_$ and executions G' and G_r such that $G_s \xrightarrow{t} G' \sqsubseteq G_r$, $\text{full}(G_r)$, $\text{consistent}(G_r)$, and $\text{racy}(G_r)$. Obviously, \Rightarrow can match this \rightarrow step and thus from the inductive hypothesis on the $\text{PRODUCTIONSEQUENCE}$ calls (the first argument is at least one step longer), we have that Line 13 returns an extension S' of the first argument such that either $\text{APPLY}(S') = \hat{G}$, or $\text{race}(S')$. The former cannot happen, since \hat{G} is racy , and thus a race would have been detected. Therefore, $\text{GETNEXT}(S_0, G_t)$ in Line 13 will return an extension S' of S_0 such that $\text{race}(S')$. The second case is that $e \in R$ and there is an execution G' s.t. $\text{PREFIXEXTENSION}(G_s, G_t, \text{event}(t)) \xrightarrow{t} \xrightarrow{e@p} G_t$, where $\langle t, p \rangle = \text{GETREV}(G_s, G_t, e)$.

From Lemma 4.14, we have that $G_s \sqsubseteq G_p \sqsubseteq G_t$. We will show that in every iteration before the one where the test in Line 23 succeeds, either an extension S' of S_0 is returned such that $\text{race}(S')$ (Line 29), or if S' and A' are the new values of S and A , respectively, at the end of the loop, $G \triangleq \text{APPLY}(S)$, and $G' \triangleq \text{APPLY}(S')$, then $\text{MAXCOMPLETION}(G_p, A') \neq \perp$, $G \sqsubset G' \sqsubseteq \text{MAXCOMPLETION}(G_p, A')$, $A' \sqsubset G'.E$, and no event of G_s is revisited by $S' \setminus S_0$.

In the first iteration, it is $G = G_s$, $\text{next}_P(G) = r_0$, $A = \emptyset$. By construction of G_p it is $G \sqsubseteq G_p$. The test in Line 22 will fail because $r_0 \notin G_p$, and thus it will be $A' = [r_0]$. From $r_0 \in \text{nextCandidates}(G_s)$ and $G_s \sqsubseteq G_p$, we have $r_0 \in \text{nextCandidates}(G_p)$: G_p can only have more enabled events (caused by an in-place revisit). From Lemmas 4.11 and 4.12, there is an execution G_{MC} such that $G \xrightarrow{r_0} G_{MC}$. The first argument in the recursive call of Line 28 is S_0 and the second (G_{MC}) has at least one

fewer event than $G_t(e_0)$. Since $r_0 \in G_{MC}$, we can use the inductive hypothesis and get that for the resulting extension S' of S_0 , it is either $\text{race}(S')$, in which case $\text{next}_P(S_0, G_t)$ returns S' and fulfills its postcondition, or $\neg \text{race}(S')$, $G \sqsubset G' \sqsubseteq G_{MC}$ and no step in $S' \setminus S_0$ revisits an event of G_s . Since $G' \sqsupset G$, $\text{next}_P(G) = r_0$, and no event of G is revisited by $S' \setminus S_0$, it will be $r_0 \in G'.E$ and thus $A' \subseteq G'.E$.

For any next iteration, let S be the initial value of S and S' the value at the end of the iteration. Again, if a race is detected, the an extension S' of S_0 is returned which fulfills the postcondition of $\text{GETNEXT}_{S_0, G_t}$. Otherwise, let $G \triangleq \text{APPLY}(S)$, and $G' \triangleq \text{APPLY}(S')$. From the previous iteration we have $G \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$, $A \subseteq G.E$, and $S \setminus S_0$ does not revisit events of G_s . Additionally, since $G \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$, it is also $\text{next}_P(G) \neq \perp$.

If the test in line Line 22 succeeds, we get that $\text{MAXCOMPLETION}(G_p, A') \neq \perp$ from the previous iteration ($A' = A$). Otherwise, it suffices to show that for the execution $G_{MC} = \text{MAXCOMPLETION}(G_p, A)$, it is $e \in \text{nextCandidates}(G_{MC})$. Assume that $e \notin \text{avail}(G_{MC})$. Since $e \in \text{avail}(G)$ and $G \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$, it is $e \in G_p$: it cannot be that $e \in A$: every time an event e is added to A , in Line 28 we get an execution that includes e and hence it cannot be that $e \in \text{avail}(G'')$ for a later G'' . Since $G \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$ and $e \in \text{avail}(G)$, it is $\text{availPrefix}(G, e) \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$. Together with $e \in G_p$, we get that $\text{availPrefix}(G, e) \sqsubseteq G_p$, which contradicts the fact that the test in line Line 22 fails. Thus, $e \in \text{avail}(G_{MC})$. Assume that $e \notin \text{nextCandidates}(G_{MC})$. Then it must be that $e \in R_l^{\text{spec}}$, and there are is a speculative read $r_b \in R_l^{\text{spec}}$ that is not completed. This contradicts that $e \in \text{nextCandidates}(G)$ since no speculative read event in $G_p.E \setminus G.E$ can be not completed.

We thus have $\text{MAXCOMPLETION}(G_p, A) \neq \perp$ in both cases. Since the first iteration is already executed, we have $G_s \sqsubset G$. Using the inductive hypothesis in Line 28 we get that $G \sqsubset G' \sqsubseteq \text{MAXCOMPLETION}(G_p, A')$ and $S' \setminus S$ does not revisit events of G , and thus of G_s ($G_s \sqsubset G$). To see that the remaining invariant, i.e., $A' \subseteq G'.E$, holds, we consider again two cases, depending on whether the test in Line 22 succeeds. If the test succeeds, we also get that $A' \subseteq G'.E$, since $A' = A \subseteq G.E \subseteq G'.E$ ($G \sqsubset G'$). If the test fails, it suffices we show that $e \in G'.E$. This follows from $\text{next}_P(G) = e$, $G' \sqsupset G$, and no event of G is revisited by $S' \setminus S$.

If a race is not detected, the loop will eventually terminate by the call in Line 24. To see this, observe that at each iteration, G' is a consistent execution of the program that has at least one more event than the previous iteration ($G \sqsubset G'$). From the success of the test in Line 23, $G \sqsubseteq \text{MAXCOMPLETION}(G_p, A)$, and $A \subseteq G.E$, we have that $G = \text{MAXCOMPLETION}(G_p, A)$ and $\text{next}_P(G) = e_0$. From the invariants of the loop, we have that the events of A appear in the sequence in the same order as they were added (once added they are never deleted). The events that would be affected by a revisit step of e_0 to r_0 are the events of $\text{MAXCOMPLETION}(G_p, A)$ that are not in G_p , thus the events of A . Therefore, that revisit step $\xrightarrow[r_v r_0]{e_0 @p}$ will be enabled and for the resulting execution G' we have $G_0 \sqsubseteq G_s \xrightarrow{e_0 @p r_0} \xrightarrow{@p r_0} G' \sqsubseteq G_t$. Finally, the revisit step does not revisit any event of G_s ($r_0 \notin G_s$). \square

4.5 Optimality

LEMMA 4.18. *Given consistent execution G and G' s.t. $G_0 \Rightarrow^* G \xrightarrow[r_v r]{e} G'$, it is $G.\text{rf}(r) \in G'.E$.*

PROOF. If e is a speculative read that revisited r , then in G' it is reading from $G.\text{rf}(r)$ by definition of the revisit step. Otherwise, e is a write event. If e can also be added with a maximal step consistently, then it must be that $\neg \text{racy}(G, e)$, since $G \Rightarrow G'$, which implies that all same-location writes, including $G.\text{rf}(r)$ are in the porf -prefix of e , and thus it would be $G.\text{rf}(r) \in G'.E$. If e

cannot be added with a maximal step, it must be a write that is part of an RMW that reads from a write w that is already read by another completed RMW $\langle r', w' \rangle$. Since G' is consistent, it must be that the revisit step either revisits or deletes r' . In the first case, i.e., $r = r'$, it is obvious that $w = G.\text{rf}(r)$ is in $G'.E$. In the second case, assume that $G.\text{rf}(r) \notin G'.E$. Then, it must be that the write that r is reading from in G was removed by a series of \rightsquigarrow , and thus in the execution \hat{G} that results by removing the events of the deleted set, r is blocking. However, this contradicts that the revisit step takes place. \square

LEMMA 4.19. *Let $\langle G, \langle \rangle \rangle$ be a configuration s.t. $\langle G_0, \emptyset \rangle \xRightarrow[r_v^-]{e @ -} \langle G, \langle \rangle \rangle$ for some event e . Then, there is no algorithmic step after $\langle G, \langle \rangle \rangle$ that revisits or deletes e .*

PROOF. Let S the production sequence of an execution G , i.e., $G = \text{APPLY}(S)$, such that $\neg \text{race}(S)$. We define $rv(S)$ as the set of events in G that revisited an event, and $mrv(S)$ the subset of $rv(S)$ of events that revisited an event that was not revisited or deleted later.

We will first prove by induction that given a production sequence $S \ ++ \ t$, t step cannot revisit or delete an event of $mrv(S)$. The base case is $S = \emptyset$, which is trivial. For the inductive case, we will prove that there is no step t that revisits or deletes an event of $mrv(S)$. Let $G = \text{APPLY}(S)$, $e = \text{next}_P(G)$, and t a step from S that revisits an event r and revisits or deletes an event $b \in mrv(S)$. For the event a that b revisited, it is $a < b$.

The first case is that b is a speculative read event, and thus a is another speculative read event. From the inductive hypothesis, b was not revisited or deleted after it was added and revisited a , and thus it is reading from the write that a was reading before the revisit from b . In execution G , b is completed: it does not read \perp , and for $\text{next}_P(G)$ to pick another speculative read is for b to be followed by the matching confirmation write. Therefore, in execution G , b is completed and a reads from the matching confirmation CAS w of b . Because the revisit step t is enabled, there is an execution \hat{G} s.t. $\hat{G} \rightsquigarrow^* G_b \xrightarrow{b} G'_b \rightsquigarrow^+ G_w \xrightarrow{w} G'_w \rightsquigarrow^* G$. In G'_w , a reads from w , since it does so in G and no \rightsquigarrow would change this **rf** edge. In G_w , a reads \perp since $G_w \xrightarrow{w} G'_w$, and a reads from w in G'_w . Event a reads \perp in G'_b , since it does so in G_w and no \rightsquigarrow can change this **rf** edge. This leads to contradiction, since $G_b \xrightarrow{b} G'_b$ but in G_b there is a same-location speculative read a that reads \perp .

The second case is that b is a write event, and thus a is a read event. Let G_1 and G_2 be the executions before and after the revisit step t' of b to a , i.e., $G_1 \xrightarrow[r_v a]{b} G_2$. From Lemma 4.18, it is $\hat{w} \triangleq G_1.\text{rf}(a) \in G_2.E$, and since G_2 is not racy, it is $\langle \hat{w}, b \rangle \in G_2.\text{porf}$. From the induction hypothesis, b is in G and a has not been revisited. Since G_2 is not racy, all writes W_b of G_2 to $\text{loc}(b)$ are in the **porf?**-prefix of b in G_2 , and thus the writes in W_b are also in the **porf?**-prefix of b in G . Additionally, since the revisit step t' was enabled from G_1 , \hat{w} is the **co**-maximal of the writes in $W_b \setminus \{b\}$ in G_2 . It is easy to see that there are no other writes in the **porf**-prefix of b in G apart from those in W_b , since they would have to revisit an event in the **porf**-prefix of b , and thus delete b . Execution G is also not racy, and thus a reads from a write b' that is **porf?**-after b (a might have been blocked by reading from b , and was in-place revisited later). The revisit step t is enabled in G , and hence there is an execution \hat{G} s.t. $\hat{G} \rightsquigarrow^* G_b \xrightarrow{b} G'_b \rightsquigarrow^* G$. Execution G'_b includes a ; otherwise, a is in the **porf**-prefix of e , and hence the same holds for b since it is $\langle b, a \rangle \in \text{porf?}; [b']; \text{rf}$. Any write **porf**-after b is not in G'_b , and thus $G'_b.\text{rf}(a) = b$. From $G_b \xrightarrow{b} G'_b$, we have that b is blocking by reading from the **co**-maximal write of G_b . It is easy to see that the writes of G_b to $\text{loc}(b)$ are exactly the writes in $W_b \setminus \{b\}$, and a does not block when reading from the **co**-maximal write \hat{w} from this set. Therefore, we reached a contradiction.

Having proved that no event of $mrv(S)$ can be revisited or deleted, we will show that the same holds for $rv(S)$ by proving that every event of $rv(S)$ is either in $mrv(S)$ or in the porf -prefix of an event in $mrv(S)$; in the latter case the event of $rv(S)$ cannot be revisited or deleted because an event of $mrv(S)$ would also have to be revisited or deleted. We will prove our claim by induction on the production sequence S . The base case is $S = \emptyset$, which is trivial.

Let $G = \text{APPLY}(S)$, $e = \text{next}_P(G)$, t a step from S , G' the resulting execution, and b an event of $rv(S)$. Assume that t is a non-revisit step. The first case is $b \in mrv(S)$. Then it follows that it is also $b \in mrv(S \ ++ \ t)$. The second case is $\langle b, e' \rangle \in G.\text{porf}$, for an event $e' \in mrv(S)$. Even if t includes an in-place revisit, it is easy to see that $\langle b, e' \rangle \in G'.\text{porf}$. In both cases $b \in mrv(S \ ++ \ t)$ or b is in the $G'.-prefix of an event in $mrv(S \ ++ \ t)$.$

Assume now that t is a revisit step $\xrightarrow{e}_{rv \ r}$. The first case is $b \in mrv(S)$. We will show that if $b \notin mrv(S \ ++ \ t)$, it is $\langle b, e \rangle \in G'.. To see this, observe that b revisited an event $a \in G.E$ that was either revisited or deleted by t . Since $a < b$ in G , but b was not deleted (events of $mrv(S)$ cannot be revisited or deleted), b is in the porf -prefix of e in G' , and $e \in mrv(S \ ++ \ t)$. The second case is $\langle b, b' \rangle \in G.\text{porf}$, for an event $b' \in mrv(S)$ that revisited $a' \in G$. If $b' \in mrv(S \ ++ \ t)$, it must be that a' was either revisited or deleted by t . For the same reason as in the previous case, b' must be in the porf -prefix of e in G' , and thus it is also that b is in the porf -prefix of e in G' . $\square$$

LEMMA 4.20. *Given an execution G reachable by the algorithm, G has at most one pending read event.*

PROOF. Given a configuration $\langle G, < \rangle$ reachable by the algorithm, we will prove by induction on the length of the production sequence S of $\langle G, < \rangle$ that (1) G has at most one pending read event and it was added or revisited in the exact previous step; and (2) if there are events $r, w, e \in G.E$ such that $\langle r, w \rangle \in G.RMW$, $e \in R$, and $r < e < w$, then e is in the porf prefix of an event of G that revisited in S .

The base case is the initial execution G_0 , which is trivial. For the inductive case, we will prove that if $\langle G_0, \emptyset \rangle \xRightarrow{*} \langle G, < \rangle \xrightarrow{e}_t \langle G', <' \rangle$ and every configuration up to $\langle G, < \rangle$ satisfies properties (1) and (2), then so does $\langle G', <' \rangle$.

We first consider the case where $t = nr$. If G has no pending read, then the added event is not a write that is a part of an RMW, and thus condition (2) holds for $\langle G', <' \rangle$. Condition (1) also holds since G' can have at most one pending read: the event e that is just added. If G has a pending read a , then e is the matching write, by definition of $\text{next}_P(G)$. Thus, execution G' has no pending read and condition (1) holds. Note that there are no await CASes, and thus adding a write cannot introduce any pending reads. It remains to show that if G has a pending read a , condition (2) holds as well. From the inductive hypothesis, a was just added or revisited in G . If a was just added, then condition (2) holds trivially by the inductive hypothesis. If a was just revisited, then any event $<'$ -after a is the porf -prefix of the write that revisited, and condition (2) still holds.

The other case is $t = rv \ a$. We will show that, if G' has a pending read, it must be a . If G has a pending read, then e is the matching RMW. Thus, it suffices to show that the only pending read of G' can be the read a . Suppose that this is not the case. Since there are no await CASes, the only way for G' to have a pending read is by deleting the write of an RMW, but not the read, i.e., there are event $r, w \in G.E$ s.t. $\langle r, w \rangle \in G.RMW$, $r \neq a$, and G' contains r but not w . If r was in the prefix of e , then so would w . Therefore, r is not in the prefix of e , but instead $r < a$. Since w was deleted, it is $a < w$. From condition (2) for G , a is in the prefix of a event that revisited, which contradict that a is revisited in t (Lemma 4.19). Therefore condition (1) holds for G' , since it has at most one pending read. Condition (2) also holds since we showed that if there is a pending read in G' , it must be a , which was the revisited event in t . \square

LEMMA 4.21. *Let S be a production sequence such that $\neg \text{race}(S)$ and $\text{APPLY}(S) = G_t$. Then there is no execution G in S s.t.*

- G results from a $\xRightarrow{e}_{rv\ r}$ step
- $G \setminus \{r\} \sqsubseteq G_t$
- $e \in W \Rightarrow \langle G_t.\text{rf}(r), e \rangle \in G.\text{porf}$
- $e \in R^{\text{spec}} \Rightarrow G_t.\text{rf}(r) = G.\text{rf}(e)$

PROOF. Assume the opposite and let G be the last such execution in S . In both cases, $G_t.\text{rf}(r)$ is in the **porf**-prefix of e , and thus it cannot be deleted in a later step (Lemma 4.19). The only way for G_t to be reached from G is for another read of G to be revisited so that r is deleted. Let G' be the first execution in S after G that does not include r , resulting from a step t that revisits a read r' of an execution \hat{G} by an event e' , i.e., $\hat{G} \xRightarrow{e'}_{rv\ r'} G'$. In G , all events \prec -after r are in the **porf**-prefix of e , thus the only events of G that can be revisited are \prec -before r or r itself. Additionally, no events of G are deleted until \hat{G} . Since t is the first step that deletes r ($r' < r$), there are no revisits to events of G between G and \hat{G} in S : any other revisit by an event a would have to not delete r , and thus a must be **porf**-after r , which would imply that r cannot be deleted later, contradicting the assumption about t deleting r . Let $w_t = G_t.\text{rf}(r')$ and $\hat{w} = \hat{G}.\text{rf}(r')$.

Assume that $e' \in W$. From $G \setminus \{r\} \sqsubseteq G_t$, we have that either $G.\text{rf}(r') = w_t$, or r' reads from a write w' in G and blocks, and all writes **porf**-after w' and before $w_t = G_t.\text{rf}(r')$ in G_t block r' , so that it can be in-placed revisited by all of them and read from w_t in G_t . From Lemma 4.18, it is $\hat{w} \in G'.E$, and since G' is not racy, it is $\langle \hat{w}, e' \rangle \in G'.\text{porf}$.

Let $w' \triangleq G.\text{rf}(r')$. Since no revisit happened from G to \hat{G} to the events of $G.E$, $\hat{G}.\text{rf}(r') = \hat{w}$ is either w' , or a write **porf**-after w' (otherwise \hat{G} is racy). From Lemma 4.18, \hat{w} is in $G'.E$, and thus it is **porf**-before e' (otherwise G' is racy). Therefore, $\langle w', e' \rangle \in G'.\text{porf}$. Assume that $w' \neq w_t$. Since r' was reading from \hat{w} in \hat{G} , $\hat{w} \in G'.E$, and the revisit step was taken from \hat{G} to G' , \hat{w} was **co**-maximal in \hat{G} and did not block r' . Since G' eventually leads to G_t , Lemma 4.19 implies that G_t also includes e' and its **porf**-prefix. If $w_t \in G'$, then $\langle w_t, e' \rangle \in G'.\text{porf}$. Otherwise, w_t is in the **porf**-suffix of e' in G_t , and we have thus reached a contradiction: $G \setminus \{r\} \sqsubseteq G_t$, $G.\text{rf}(r') = w'$, but the writes in the **porf** path from w' to w_t in G_t do not all block r' (\hat{w} does not). Therefore, we have $\langle w_t, e' \rangle \in G'.\text{porf}$.

Assume now that $e' \in R^{\text{spec}}$, which implies $r' \in R^{\text{spec}}$. From $G \setminus \{r\} \sqsubseteq G_t$, we have that $w_t = G.\text{rf}(r')$. Since no events of G are deleted or revisited until \hat{G} , it is $w_t = \hat{w}$. Since r' is revisited from e' , it is $G'.\text{rf}(e') = w_t$, which gives us that $\langle w_t, e' \rangle \in G'.\text{porf}$.

To reach a contradiction, it remains to be shown that $G' \setminus \{r'\} \sqsubseteq G_t$, which would imply that G' is an execution after G in S that also satisfies the properties in question. To see that $G' \setminus \{r'\}$, notice that $G \setminus \{r\} \sqsubseteq G_t$, no event of G is revisited or deleted until \hat{G} , G' does not contain r , and all events in G' that are not in G also exist in G_t , since they are events in the **porf**?-prefix of e' , e' cannot be deleted by a later step in S , and S reaches G_t .

We have now reached a contradiction: G' is an execution after G in S that results from a step $\xRightarrow{e'}_{rv\ r'}$ and satisfies the respective properties. □

LEMMA 4.22. *Let S be a production sequence that reaches a consistent execution G_t from G_0 , and G be an execution in S such that $G \sqsubseteq G_t$. Then there is no algorithm step in S after G that revisits an event of G .*

PROOF. Assume the opposite and take t to be the first step after G in S that revisits a read r of G . Let \hat{G} and G' be the executions before and after t , respectively, i.e., $G \xRightarrow{*}_{rv\ r} \hat{G} \xRightarrow{e@-}_{rv\ r} G'$.

Assume that $e \in W$, and let $\hat{w} = \hat{G}.rf(r)$. From Lemma 4.18, \hat{w} is also in G' , and thus in the **porf**-prefix of e (G' is not racy). Since $G \sqsubseteq G_t$, r is either reading from $w_t = G_t.rf(r)$ in G , or is reading from a **co** maximal write $w_b \neq w_t$ and blocks. In the first case, assume that $w_t \neq \hat{w}$. This implies that r blocks by reading from w_t and was in-place revisited by the **co**-later \hat{w} . From Lemma 4.19 and the fact that S reaches G_t , \hat{w} also exists in G_t , and therefore in G_t r reads from a write w_t that blocks it but it is not **co**-maximal, which is a contradiction. Therefore, we have $w_t = \hat{w}$ in the first case. In the second case, in \hat{G} r is also reading from a **co**-maximal write \hat{w} , (either w_b , or a **co**-later write that in-placed revisited r). Since the revisit step t is enabled, r is not blocking in \hat{G} . Execution G' includes \hat{w} (Lemma 4.18), and is in the **porf**-prefix of e (G' is not racy). Execution G_t can be reached from G' and e must also exist in G_t . Therefore, it must be $\hat{w} = w_t$: it is $G \sqsubseteq G_t$, and \hat{w} is the **porf**-last write that r can be in-place revisited from, when it initially reads from a write w_b in G that blocks r . We thus have $\langle G_t.rf(r), e \rangle \in G'.porf$ in either case.

Assume that $e \in R^{spec}$. Then it must be that $r \in R^{spec}$ and r is either reading \perp or from $G_t.rf(r)$ in G , since it is $G \sqsubseteq G_t$. In the former case, since execution G is reachable from G_0 , and thus $G_0 \sqsubseteq G$, G contains a speculative read r_s in the same location as r that has no matching confirmation CAS, and in G_t , r reads from this confirmation CAS w_t , i.e., $G_t.rf(r) = w_t$. Execution \hat{G} also contains r_s , and since $next_P(\hat{G}) = e \in R^{spec}$, the matching confirmation CAS w_t of r_s is in \hat{G} and r reads from w_t , i.e., $\hat{G}.rf(r) = w_t$. In the latter case, r continues reading from w_t in \hat{G} . In both cases, $\hat{G}.rf(r) = G_t.rf(r)$, and since t revisits r from e , it is $G'.rf(e) = G_t.rf(r)$.

Finally, it is easy to see that $G' \setminus \{r\} \sqsubseteq G_t$: the $G'.porf$ -prefix of e is included in G_t , $G \sqsubseteq G_t$, and no event of G is revisited in S between G and \hat{G} . We have now reached a contradiction due to Lemma 4.21. \square

THEOREM 4.23. *Given a full, consistent execution $G_f \sqsupseteq G_0$, there is at most one production sequence that reaches G_f .*

PROOF. We will show that if there is a production sequence that leads to G_f , then this production sequence can only be $PRODUCTIONSEQUENCE(\emptyset, G_f)$. We will prove by induction (with the same measure as in Lemma 4.17) that given a production sequence S that reaches G_s (i.e., $G_s = APPLY(S)$) and a consistent execution G_t s.t. $G_0 \sqsubseteq G_s \sqsubset G_t$ and $next_P(G_s) \in G_t.E$, any production sequence that extends S and reaches an execution $\hat{G} \sqsupseteq G_t$, (1) does not delete the first event added (i.e., $next_P(G_s)$) and (2) must be a suffix of $GETNEXT(S, G_t)$. Then our result follows by induction on the number of iterations of the loop in $PRODUCTIONSEQUENCE(\emptyset, G_f)$. If at any step $next_P(G_s)$ returns a production sequence such that $race(S)$, there no production sequence that reaches G_f exists (since it cannot be S). Termination of the loop is proven in Lemma 4.17.

Consider a production sequence S_0 and a consistent execution G_t such that $G_0 \sqsubseteq G_s \sqsubset G_t$ and $next_P(G_s) \in G_t.E$, where $G_s = APPLY(S_0)$. Also consider a production sequence \hat{S} that extends S_0 and reaches an execution $\hat{G} \sqsupseteq G_t$. The test in Line 9 cannot succeed, otherwise G_t is racy and thus there is no production sequence that reaches \hat{G} . If the test in Line 10 succeeds, it must be that $\hat{S} \setminus S_0$ starts with the $\xrightarrow{t}{nr}$. Otherwise, to reach an execution $\hat{G} \sqsupseteq G_s$, \hat{S} must revisit an event of G_s . To see this, consider the case where the added event is a read and assume that another step t' is taken and examine two cases, depending on whether $w \triangleq G_t.rf(r)$ is in $G_s.E$. If $w \in G_s.E$, t is the step that adds r to read from w . A step t' would add r to read from a $w' \neq w$ and thus a later step needs to delete r or w' , which means that an event of G_s is revisited. If $w \notin G_s.E$, it cannot be that $r \in R^{spec}$: no read option for r will allow it to read from another write, without a revisit. Hence $r \notin R^{spec}$, t is the step that adds r to read from the **co**-maximal write w_m and block. Any other step t' would either end with r reading from a non **co**-maximal write and either block or not block. In the former case, the execution would be dropped. In the latter case, r cannot be eventually revisited by w since

w_m would be in the execution and thus r would not be reading from a **co**-maximal write. Thus an event of G_s must be revisited in \hat{S} , which leads to a contradiction from Lemma 4.22.

If the test in Line 10 fails, we will show that \hat{S} must reach execution G_p and first follow the steps returned in Line 24. The production sequence \hat{S} will first add the read event r_0 . Since $G_s.\text{rf}(r_0) \notin G_s.E$ (the test in Line 10 fails), the event r_0 will read from another write and must either be deleted and added again, or be revisited. It cannot be that r_0 eventually reads from $G_s.\text{rf}(r_0)$ by a series of in-place revisits, again because the test in Line 10 fails. From Lemma 4.22, it cannot be that r_0 is deleted (it would require a read of G_s being revisited). Thus r_0 will be eventually revisited, which will lead to execution G_p .

We will show that as soon as a configuration $\langle G', <' \rangle$ is reached s.t. $\text{availPrefix}(G', e_0) \sqsubseteq G_p$ and $\text{next}_P(G') = e_0$, \hat{S} will perform a revisit step and reach G_p . Assume that this is not the case: such an execution G' is reached but \hat{S} does not proceed with a revisit of r_0 .

We first consider the case where \hat{S} proceeds with a revisit step. Assume that e_0 is a speculative read, and thus r_0 is a speculative read. If e_0 reads from a write that is **porf**-after the one r_0 is reading, and thus **porf** after r_0 since r_0 is a completed speculative read ($\text{next}_P(\cdot)$ picks another speculative read), then r_0 will keep reading from $G'.\text{rf}(r_0)$ in any later execution in \hat{S} , since it will be in the **porf**-prefix of an event that revisited, which contradicts that \hat{S} reaches an execution $\hat{G} \sqsupseteq G_t$. If e_0 read from a write that is **porf**-before the one r_0 is reading, \hat{S} starts with a revisit of an event of G_s , which contradicts Lemma 4.22. The other case is that e_0 is a write w_0 . If w_0 performs another revisit step, it will never be deleted in a later step (Lemma 4.19). Also, r_0 will not be deleted since no event of G_s will be revisited (Lemma 4.22), and thus \hat{S} cannot reach an execution $\hat{G} \sqsupseteq G_t$. To see this, observe that r_0 does not block in G' (since the revisit to r_0 is enabled), and in the execution after a different revisit step it is not **co**-maximal and therefore its **rf** edge cannot change by an in-place revisit.

The other case is for e_0 to be added with a **nr** step and later to be deleted and eventually added again to revisit e_0 . Let $\xrightarrow[r_v^-]{e' @ -}$ be the first revisit step to events of G' after this **nr** step. If e' is in the **porf**-suffix of e_0 , then e_0 cannot be later deleted, and thus r_0 will never be revisited from e_0 . If e' is not in the **porf**-suffix of e_0 , then r_0 cannot be later revisited by e_0 : the revisit would have to delete e' , which contradicts Lemma 4.19. Note that e' cannot have been deleted in a step before the hypothetical revisit of e_0 , again due to Lemma 4.19.

It remains to show that \hat{S} necessarily reaches the executions in Line 28 (for each iteration), using the corresponding steps returned by the recursive call, before ending with the revisit step in Line 24. We will show that at each iteration, \hat{S} must follow exactly the steps in S and the events in A cannot be deleted by a later step in \hat{S} . Initially this holds since $S = \emptyset$, $A = \emptyset$, and no events of G_s can be revisited (Lemma 4.22) in \hat{S} .

Consider an iteration of the loop. The first case is for the test in Line 22 to succeed. Then A remains the same. Since the events in A cannot be deleted by a later step in \hat{S} , it must be that an execution G' is reached such that $G' \sqsupseteq \text{MAXCOMPLETION}(G_p, A)$; otherwise, the maximality check will fail. From the inductive hypothesis, \hat{S} must match the steps returned by the recursive call. The second case is for the test in Line 22 to fail. Since the events in A (before adding e) cannot be deleted by a later step in \hat{S} , it must again be that an execution G' is reached such that $G' \sqsupseteq \text{MAXCOMPLETION}(G_p, A)$. From the inductive hypothesis, the first event added, i.e., e , cannot be deleted by a later step in \hat{S} . Thus, it must be that an execution G'' is reached such that $G' \sqsupseteq \text{MAXCOMPLETION}(G_p, A + e)$. From the inductive hypothesis, \hat{S} must again match the steps returned by the recursive call.

We have thus shown that at each iteration, \hat{S} matches the steps of S . We note that at no point a race can be detected, because \hat{S} does reach G_t (from the hypothesis), and any \hat{S} matches the steps of

S. The proof is concluded by observing that the first event added after G_s in the unique production sequence that reaches an execution $G' \sqsupseteq G_t$, i.e., r_0 , is not deleted by any step in the production sequence (Lemma 4.22).

□