

Dag-Calculus: A Calculus for Parallel Computation

Technical Appendix

This appendix consists of two main parts: in the first, we prove correct the translations of parallel primitives into the dag calculus, in the second, we prove that the scheduling algorithm implements the rules of the dag calculus. Concretely, Appendix A provides auxiliary definitions and properties for the translations; Appendices B, C and D provide the proofs of translations of fork-join, async-finish and futures, respectively; Appendix E provides the proof of the implementation of dag calculus.

A Auxiliary definitions

Here we recall the syntax and semantics of the dag calculus, which we denote in short dag calculus. These constructs and rules are described in detail in Section 3 of the paper.

Note that while the syntax of the calculus is identical to the one presented in the paper, the formulation of the reduction rules is slightly different: the dag-modification constructs are assumed to always appear in the context of a let-binding or sequential composition, and the reduction reduces them together with the binding/composition. This approach is clearly equivalent to the presentation in the paper, but it significantly reduces the number of otherwise not interesting intermediate steps that would be possible to encounter in the translations, thus shortening the translations and the proofs significantly.

Definition 1. A *computation dag* is a pair (V, E) , where the *vertex-map* V is a finite map from threads to a pair consisting of expressions, and status and *edges* E is a edges

$$\begin{aligned}
 e ::= & v \mid e \oplus e \mid e \otimes e \mid (e, e) \mid \text{fst } e \mid \text{snd } e \mid \text{let } x = e \text{ in } e \mid \\
 & e \mid \text{alloc}() \mid e := e \mid !e \mid \text{if } e \text{ then } e \text{ else } e \mid \\
 & \text{newTd } e \mid \text{release } e \mid \text{newEdge } (e, e) \mid \text{self}() \mid \text{yield}() \\
 v ::= & x \mid \ell \mid \mathbf{t} \mid \mathbf{n} \mid () \mid (v, v) \mid \text{fun } f \text{ x is } e \text{ end} \\
 K ::= & \bullet \mid \text{let } x = K \text{ in } e \mid K := e \mid v := K \mid !K \mid \text{release } K \\
 & \mid \text{newEdge } (K, e) \mid \text{newEdge } (v, K) \mid \dots
 \end{aligned}$$

Figure 1: Abstract syntax for DAG-Calculus.

$$\begin{array}{c}
\frac{}{\sigma, \text{fst}(v_1, v_2) \rightarrow v_1, \sigma} \text{FST} \quad \frac{}{\sigma, \text{snd}(v_1, v_2) \rightarrow v_2, \sigma} \text{SND} \quad \frac{l \notin \text{dom}(\sigma)}{\sigma, \text{alloc} \rightarrow l, \sigma[l \mapsto ()]} \text{ALLOC} \\
\frac{\sigma(l) = v}{\sigma, !l \rightarrow v, \sigma} \text{DEREF} \quad \frac{l \in \text{dom}(\sigma)}{\sigma, (l := v) \rightarrow (), \sigma[l \mapsto v]} \text{ASSIGN} \\
\frac{}{\sigma, ((\text{fun } f \text{ is } e \text{ end}) v) \rightarrow e[f \mapsto \text{fun } f \text{ is } e \text{ end}][x \mapsto v], \sigma} \text{APPLY} \\
\frac{\sigma, e \rightarrow e', \sigma'}{\sigma, K[e] \rightarrow K[e'], \sigma'} \text{CONTEXT} \\
\frac{V(t) = (e_1, X) \quad \sigma_1, e_1 \rightarrow e_2, \sigma_2}{V, E, \sigma_1 \rightarrow V[t \mapsto (e_2, X)], E, \sigma_2} \text{STEP} \quad \frac{V(t) = (v, X) \quad E' = E \setminus \{(t, t') \mid t' \in \text{dom}(V)\}}{V, E, \sigma \rightarrow V[t \mapsto ((), F)], E', \sigma} \text{STOP} \\
\frac{V(t) = (K[\text{let } x = \text{newTd } e \text{ in } e'], X) \quad t' \text{ fresh}}{V, E, \sigma \rightarrow V[t \mapsto (K[e'][x \mapsto t']), X]][t' \mapsto (e, N)], E, \sigma} \text{NEWTD} \quad \frac{V(t) = (e, R) \quad \{t' \mid (t', t) \in E\} = \emptyset}{V, E, \sigma \rightarrow V[t \mapsto (e, X)], E, \sigma} \text{START} \\
\frac{V(t) = (K[\text{release } t'; e], X) \quad V(t') = (e', N)}{V, E, \sigma \rightarrow V[t \mapsto (K[e], X)][t' \mapsto (e', R)], E, \sigma} \text{RELEASE} \\
\frac{V(t) = (K[\text{newEdge } t_1 \ t_2; e], X) \quad t_1, t_2 \in \text{dom}(V) \quad (\text{status}(V(t_2)) \in \{N, R\} \vee t_2 = t) \quad E' = E \uplus (\text{if } \text{status}(V(t_1)) = F \text{ then } \emptyset \text{ else } \{(t_1, t_2)\}) \quad E' \text{ cycle-free}}{V, E, \sigma \rightarrow V[t \mapsto (K[e], X)], E', \sigma} \text{NEWEDGE} \\
\frac{V(t) = (K[\text{self}], X)}{V, E, \sigma \rightarrow V[t \mapsto (K[t], X)], E, \sigma} \text{SELF} \quad \frac{V(t) = (K[\text{yield}; e], X)}{V, E, \sigma \rightarrow V[t \mapsto (K[e], R)], E, \sigma} \text{YIELD}
\end{array}$$

Figure 2: Dynamic semantics for dag calculus. **N** stands for thread status *new*, **R** for *released*, **X** for *executing* and **F** for *finished*. The operation $\text{status}(V(t))$ denotes the status of thread t , i.e. the second component of $V(t)$.

(edges) between threads, more precisely:

$$\begin{aligned}
s & ::= \mathbf{N} \mid \mathbf{R} \mid \mathbf{X} \mid \mathbf{F} \\
V & : \mathcal{T} \rightarrow_{\text{fin}} (e \times s) \\
E & \subseteq \text{dom}(V) \times \text{dom}(V).
\end{aligned}$$

Definition 2. The operational semantics of dag calculus is a relation \rightarrow between states made of a computation dag and a store mapping locations to values. It is given in Figure 2. We take a subset of the transitions **START** and **STOP** as *scheduler* transitions, written \rightarrow_s .

Lemma 1. For any dag (V, E) and any state σ , $V, E, \sigma \not\rightarrow_s^\infty$.

Proof. Give weight to a dag by taking for each vertex weight according to its status:

$$w(s) = \begin{cases} 0 & \text{if } s = \mathbf{N} \text{ or } s = \mathbf{F} \\ 1 & \text{if } s = \mathbf{X} \\ 2 & \text{if } s = \mathbf{R} \end{cases}$$

Take the weight of the dag as the sum of the weights of its vertices, then proceed by induction on the weight of the dag. Observe that each scheduler transition decreases the weight of the node it operates on, and hence the whole dag, which ends the proof. \square

The proofs of translations to the dag calculus presented in the following appendices involve one key difficulty, related to administrative reduction steps, i.e. to the fact that one reduction step in the source language may correspond to several reduction steps in the target language. Most compiler proofs deal with administrative reduction steps using well-known proofs techniques, typically based on simulation diagrams. However, we have found that these techniques were not directly applicable to the parallel semantics of a language such as that the fork-join language considered here.

When the target program takes a reduction step, this step corresponds either to an administrative step, or to a real step from the source program. Consider the latter case. With a sequential semantics, when the target program takes a real step, the target program then typically corresponds exactly to the translation of the source program. However, with a parallel semantics, this might not be the case. For example, since the two branches of a parallel pair may reduce independently, one branch may take a real step while the other branch is in the middle of performing a sequence of administrative steps. Although the target program takes a real step, it is not, at this point, the translation of any source program. Thus, we cannot easily close a simulation diagram.

To address this challenge, we introduce an *instrumented* programming language, similar to the source programming language, except that each parallel construct gets annotated with information about identities associated with its representation in the target language: number of administrative steps already performed, thread identifiers, locations for the results, etc. We then set up a two-layer simulation diagram: the first layer relates the source program with the instrumented program, while the second layer relates the instrumented program with the target program. We are able to reason about both layers independently, using conventional simulation diagrams, and then conclude by relating the source and the target programs.

B Correctness of translations of fork-join

B.1 Syntax and Operational Semantics

The syntax and semantics in this section are precisely the same as in the paper, and repeated here for convenience.

$$\begin{aligned}
 e &::= v \mid \text{fst } v \mid \text{snd } v \mid \text{let } x = e \text{ in } e \mid v \ v \mid \text{forkjoin } (e, e) \\
 v &::= x \mid n \mid (v, v) \mid \text{fun } f \ x \text{ is } e \text{ end} \\
 K &::= \bullet \mid \text{let } x = K \text{ in } e \mid \text{forkjoin } (K, e) \mid \text{forkjoin } (e, K)
 \end{aligned}$$

$$\begin{aligned}
 K[\text{fst } (v_1, v_2)] &\rightarrow K[v_1] \\
 K[\text{snd } (v_1, v_2)] &\rightarrow K[v_2] \\
 K[\text{let } x = v \text{ in } e] &\rightarrow K[e[x \mapsto v]] \\
 K[\text{fun } f \ x \text{ is } e \text{ end } v] &\rightarrow K[e[x \mapsto v, f \mapsto \text{fun } f \ x \text{ is } e \text{ end}]] \\
 K[\text{forkjoin } (v_1, v_2)] &\rightarrow K[(v_1, v_2)]
 \end{aligned}$$

B.2 Instrumented Syntax and Operational Semantics

The general pattern of the instrumentation, both for fork-join parallelism and for the other constructs, is to annotate the parallel constructs, which need to be compiled in the translation to dag calculus, with additional information. The crucial part of these annotations is a number that describes how far the evaluation of the compiled code has progressed. Intuitively, 0 always corresponds to an expression, whose translated equivalent has not yet started execution, and each consecutive step of compiled execution takes the number one step higher. If in the evaluation process new information is obtained, such as identities of relevant threads, locations, etc., these are also included in the annotation (cf. the grammar of a_p , the annotations on parallel compositions, below)

Since in dag calculus the program parts are stored in the vertices of the dag, parts of the program that correspond to separate vertices also need to be annotated. In case that values are passed through the store, like in the compiled version of fork-join, these also need to be present in the annotation, and some additional steps may be possible (cf. the grammar for a_t , the annotations of expressions that correspond to the dag calculus threads, below). In the particular case of fork-join, 0 means that the expression has not yet had a thread allocated to it and released, $(1, l, t)$ — that the expression is running in thread t and will write its result to location l , and $(2, l)$ — that the expression has finished the evaluation, and the result *has been written* to l . The extra step required to perform the write action is why we increment the number in the annotation here.

$$\begin{aligned}
e &::= v \mid \text{fst } v \mid \text{snd } v \mid v \ v \mid \text{let } x = e \text{ in } e \mid (\text{forkjoin } (e^{a_t}, e^{a_t}))^{a_p} \\
v &::= x \mid n \mid (v, v) \mid \text{fun } f \ x \text{ is } e \text{ end} \\
L &::= \bullet \mid \text{let } x = L \text{ in } e \\
K &::= \bullet \mid L^{a_t}[(\text{forkjoin } (K, e^{a_t}))^{a_p}] \mid L^{a_t}[(\text{forkjoin } (e^{a_t}, K))^{a_p}] \\
a_t &::= 0 \mid (1, l, t) \mid (2, l) \\
a_p &::= 0 \mid (1, l) \mid (2, l, l) \mid (3, l, l, t) \mid (n, l, l, t) \text{ where } 4 \leq n \leq 7 \\
&\quad \mid (8, l, l, t) \mid (n, l, l) \text{ where } 9 \leq n \leq 11
\end{aligned}$$

A *surface* expression is one where all annotations have value 0. Note that these are isomorphic to the expressions of the source language. We define well-formed expressions as follows:

$$\text{wf}(e) \equiv \begin{cases} a_1 = 0 \wedge a_2 = 0 \wedge \text{surf}(e_1) \wedge \text{surf}(e_2) \wedge \text{surf}(L) & \text{if } e = L[\text{forkjoin } (e_1^{a_1}, e_2^{a_2})^a] \text{ and } \text{num}(a) \leq 7 \\ \text{loc}(a_1) = l_1 \wedge a_2 = 0 \wedge \text{wf}(e_1) \wedge \text{surf}(e_2) \wedge \text{surf}(L) & \text{if } e = L[\text{forkjoin } (e_1^{a_1}, e_2^{a_2})^{(8, l_1, l_2, t_2)}] \\ \text{loc}(a_1) = l_1 \wedge \text{loc}(a_2) = l_2 \wedge \text{wf}(e_1) \wedge \text{wf}(e_2) \wedge \text{surf}(L) & \text{if } e = L[\text{forkjoin } (e_1^{a_1}, e_2^{a_2})^{(n, l_1, l_2)}] \text{ and } n \in \{9, 10\} \\ a_1 = (2, l_1) \wedge a_2 = (2, l_2) \wedge \text{surf}(e_1) \wedge \text{surf}(e_2) \wedge \text{surf}(L) & \text{if } e = L[\text{forkjoin } (e_1^{a_1}, e_2^{a_2})^{(11, l_1, l_2)}] \\ \text{surf}(e) & \text{otherwise} \end{cases}$$

We extend well-formedness to annotated expressions e^{a_t} by setting $\text{wf}(e^a) \equiv a \neq 0$, and to parallel contexts, by setting $\text{wf}(K) \equiv \forall e. \text{wf}(e) \Rightarrow \text{wf}(K[e])$.

Operational semantics is defined on annotated, well-formed expressions, e^{a_t} . First,

we define some of the transitions between the parallel composition annotations.

$$\begin{aligned}
& 0 \rightsquigarrow_p (1, l_1) \text{ where } l_1 \text{ fresh} & (1, l_1) \rightsquigarrow_p (2, l_1, l_2) \text{ where } l_2 \text{ fresh} \\
(2, l_1, l_2) \rightsquigarrow_p (3, l_1, l_2, t_1) \text{ where } t_1 \text{ fresh} & (3, l_1, l_2, t_1) \rightsquigarrow_p (4, l_1, l_2, t_1, t_2) \text{ where } t_2 \text{ fresh} \\
(n, l_1, l_2, t_1, t_2) \rightsquigarrow_p (n+1, l_1, l_2, t_1, t_2) \text{ for } n \in \{4, 5, 6\} & (9, l_1, l_2) \rightsquigarrow_p (10, l_1, l_2)
\end{aligned}$$

These transitions can be used to make simple administrative reductions on the parallel composition form, as seen in the operational semantics below. Since the transitions through the remaining three steps are not entirely local, they are defined directly in the rules.

$$\begin{array}{c}
\text{fst } (v_1, v_2) \circ \rightarrow v_1 \quad \text{snd } (v_1, v_2) \circ \rightarrow v_2 \quad \text{let } x = v \text{ in } e \circ \rightarrow e[x \mapsto v] \\
\text{fun } f \text{ x is } e \text{ end } v \circ \rightarrow e[x \mapsto v, f \mapsto \text{fun } f \text{ x is } e \text{ end}] \\
\frac{a \rightsquigarrow_p a'}{\text{forkjoin } (e_1^{a_1}, e_2^{a_2})^a \circ \rightarrow \text{forkjoin } (e_1^{a_1}, e_2^{a_2})^{a'}} \\
\text{forkjoin } (e_1^{(0)}, e_2^{(0)})^{(7, l_1, l_2, t_1, t_2)} \circ \rightarrow \text{forkjoin } (e_1^{(1, l_1, t_1)}, e_2^{(0)})^{(8, l_1, l_2, t_2)} \\
\text{forkjoin } (e_1^{a_1}, e_2^{(0)})^{(8, l_1, l_2, t_2)} \circ \rightarrow \text{forkjoin } (e_1^{a_1}, e_2^{(1, l_2, t_2)})^{(9, l_1, l_2)} \\
\text{forkjoin } (v_1^{(2, l_1)}, v_2^{(2, l_2)})^{(10, l_1, l_2)} \circ \rightarrow \text{forkjoin } (v_1^{(2, l_1)}, v_2^{(2, l_2)})^{(11, l_1, l_2)} \\
\text{forkjoin } (v_1^{(2, l_1)}, v_2^{(2, l_2)})^{(11, l_1, l_2)} \circ \rightarrow (v_1, v_2) \\
\frac{\text{wf}(K) \quad e \circ \rightarrow e'}{K[(L[e])^{(1, l, t)}] \rightarrow K[(L[e'])^{(1, l, t)}]} \quad \frac{\text{wf}(K)}{K[v^{(1, l, t)}] \rightarrow K[v^{(2, l)}]}
\end{array}$$

B.3 Connecting the instrumented and source languages

We define a map $\lfloor - \rfloor : \mathcal{E}_I \rightarrow \mathcal{E}_O$ that removes annotations from the instrumented expressions. We use it to connect the instrumented language to the source language. Note that a map that adds an annotation 0 to parallel compositions is the right inverse of $\lfloor - \rfloor$. For *surface* term, it is also the left inverse.

Lemma 2 (Instr-Step). *For any two expressions $e, e' \in \mathcal{E}_I$ and annotations $a, a' \in \mathcal{A}_I$, if $e^a \rightarrow_I e'^{a'}$, then either $\lfloor e \rfloor = \lfloor e' \rfloor$ or $\lfloor e \rfloor \rightarrow_O \lfloor e' \rfloor$.*

Proof. By cases on the reduction. If the reduction is one that changes the annotation on a parallel pair or a thread, the erasure is trivially preserved. In the other cases, we can simply pick the corresponding reduction rule to match the erasure of the right-hand side. \square

Definition 3. An *administrative* reduction in the instrumented language is a reduction $e^a \rightarrow_I e'$ such that $\lfloor e \rfloor = \lfloor e' \rfloor$. We write $e^a \rightarrow_a e'^{a'}$ for administrative steps.

Lemma 3 (Admin-Fin). *For well-formed annotated expressions e^a , there are no infinite sequence of administrative reductions, i.e., if $\text{wf } e^a$ then $e^a \not\rightarrow_a^\infty$.*

Proof. Assign a weight to the expression based on the count of its annotations in evaluation positions. Take $\text{num}(a)$ to be the natural number in the annotation. Then define

$$w(e) = \begin{cases} 15 - \text{num}(a) - \text{num}(a_1) - \text{num}(a_2) + w(e_1) + w(e_2) & \text{if } e = (\text{forkjoin } ((e_1)^{a_1}, (e_2)^{a_2})) \\ w(e_1) & \text{if } e = \text{let } x = e_1 \text{ in } e_2 \\ 0 & \text{otherwise} \end{cases}$$

and extend it to annotated threads as $w(e^a) \equiv 2 - \text{num}(a) + w(e)$. Observe that all administrative reductions decrease the weight of the expression. Thus, by induction on the weight we can conclude that the lemma holds. \square

Lemma 4 (Instrument). *For any expression $e \in \mathcal{E}_l$, value $v \in \mathcal{V}_l$ annotation $a \in \mathcal{A}_l$ and location l , if $e^a \rightarrow_l^* v^{(2,l)}$ then $\llbracket e \rrbracket \rightarrow_O^* \llbracket v \rrbracket$. Moreover, if $e^a \rightarrow_l^\infty$, then $\llbracket e \rrbracket \rightarrow_O^\infty$.*

Proof. First part, by induction on the reduction sequence and Lemma 2.

Similarly, for the second part we proceed by coinduction: since $e \rightarrow_l^\infty$, by Lemmas 3 and 2 there exists e' such that $e \rightarrow_l^* e'$, $\llbracket e \rrbracket \rightarrow_O \llbracket e' \rrbracket$ and $e' \rightarrow_l^\infty$. Hence, $\llbracket e' \rrbracket \rightarrow_O^\infty$ and finally $\llbracket e \rrbracket \rightarrow_O^\infty$. \square

B.4 Connecting instrumented language to dag calculus

Translation of surface expressions into dag calculus First, let us define the translation for surface expressions. We define a map $\llbracket - \rrbracket^S : \mathcal{E}_l \rightarrow \mathcal{E}_{dagcalculus}$, as follows. Note that the map extends in a simple way to sequential evaluation contexts L that are surface contexts. We write the latter one $\llbracket - \rrbracket^L : \mathcal{L}_l \rightarrow \mathcal{K}_{dagcalculus}$.

$$\begin{aligned} \llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket^S &= \text{let } x = \llbracket e_1 \rrbracket^S \text{ in } \llbracket e_2 \rrbracket^S \\ \llbracket \text{fun } f \text{ } x \text{ is } e \text{ end} \rrbracket^S &= \text{fun } f \text{ } x \text{ is } \llbracket e \rrbracket^S \text{ end} \\ \llbracket v_1 \ v_2 \rrbracket^S &= \llbracket v_1 \rrbracket^S \ \llbracket v_2 \rrbracket^S \end{aligned}$$

For the parallel pair, we have

```

1  $\llbracket (\text{forkjoin } (e_1^{(0)}, e_2^{(0)})) \rrbracket^S =$ 
2   let l1 = alloc
3     l2 = alloc
4     t1 = newTd (l1 :=  $\llbracket e_1 \rrbracket^S$ )
5     t2 = newTd (l2 :=  $\llbracket e_2 \rrbracket^S$ )
6     t = self ()
7   in newEdge (t1, t); newEdge (t2, t);
8     release t1; release t2; yield ();
9     (!l1, !l2)
```

Translation of well-formed expressions into dag calculus Now we define the translation for well-formed expressions and annotated expressions. For well-formed, annotated expressions we take an annotated expression and a set of its dependencies:

$$\llbracket e^{(l,t)} \rrbracket(T) = \llbracket e \rrbracket(l, t, T) \quad \llbracket v^{(2,l)} \rrbracket(T) = \emptyset, \emptyset, [l \mapsto v]$$

For well-formed expressions we have $\llbracket - \rrbracket : \mathcal{E}_I \rightarrow \mathcal{L} \times \text{TId} \times \mathcal{P}_{\text{fin}}(\text{TId}) \rightarrow \mathcal{S}$, where

\mathcal{S} denotes the dag states. We have:

$$\begin{aligned}
\llbracket e \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket e \rrbracket^{\mathcal{S}}, \mathbf{R}), \{(t, t') \mid t' \in T\}, [l \mapsto ()] \quad \text{if } \text{surf}(e) \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(1, l_1)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par1}(\llbracket e_1 \rrbracket^{\mathcal{S}}, \llbracket e_2 \rrbracket^{\mathcal{S}}, l_1), \mathbf{X}], \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(2, l_1, l_2)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par2}(\llbracket e_1 \rrbracket^{\mathcal{S}}, \llbracket e_2 \rrbracket^{\mathcal{S}}, l_1, l_2), \mathbf{X}], \\
& \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(3, l_1, l_2, t_1)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par3}(\llbracket e_2 \rrbracket^{\mathcal{S}}, l_1, l_2, t_1), \mathbf{X}], t_1 \mapsto (l_1 := \llbracket e_1 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(4, l_1, l_2, t_1, t_2)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par4}(l_1, l_2, t_1, t_2), \mathbf{X}], t_1 \mapsto (l_1 := \llbracket e_1 \rrbracket^{\mathcal{S}}, \mathbf{N}), t_2 \mapsto (l_2 := \llbracket e_2 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(5, l_1, l_2, t_1, t_2)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par5}(l_1, l_2, t_1, t_2, t), \mathbf{X}], t_1 \mapsto (l_1 := \llbracket e_1 \rrbracket^{\mathcal{S}}, \mathbf{N}), t_2 \mapsto (l_2 := \llbracket e_2 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(6, l_1, l_2, t_1, t_2)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par6}(l_1, l_2, t_1, t_2, t), \mathbf{X}], t_1 \mapsto (l_1 := \llbracket e_1 \rrbracket^{\mathcal{S}}, \mathbf{N}), t_2 \mapsto (l_2 := \llbracket e_2 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\} \cup \{(t_1, t)\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^0, (e_2)^0))^{(7, l_1, l_2, t_1, t_2)}] \rrbracket(l, t, T) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\text{par7}(l_1, l_2, t_1, t_2), \mathbf{X}], t_1 \mapsto (l_1 := \llbracket e_1 \rrbracket^{\mathcal{S}}, \mathbf{N}), t_2 \mapsto (l_2 := \llbracket e_2 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\} \cup \{(t_1, t), (t_2, t)\}, [l \mapsto (), l_1 \mapsto (), l_2 \mapsto ()] \\
\llbracket L[(\text{forkjoin } ((e_1)^{a_1}, (e_2)^0))^{(8, l_1, l_2, t_2)}] \rrbracket(l, t, T) &= \\
& \llbracket (e_1)^{a_1} \rrbracket(l_1, t_1, \{t\}) \uplus ([t \mapsto (l := \llbracket L \rrbracket^L[\text{par8}(l_1, l_2, t_2), \mathbf{X}], t_2 \mapsto (l_2 := \llbracket e_2 \rrbracket^{\mathcal{S}}, \mathbf{N}), \\
& \{(t, t') \mid t' \in T\} \cup \{(t_2, t)\}, [l \mapsto (), l_2 \mapsto ()]) \\
\llbracket L[(\text{forkjoin } ((e_1)^{a_1}, (e_2)^{a_2}))^{(9, l_1, l_2)}] \rrbracket(l, t) &= \\
& \llbracket (e_1)^{a_1} \rrbracket(l_1, t_1, \{t\}) \uplus \llbracket (e_2)^{a_2} \rrbracket(l_2, t_2, \{t\}) \uplus \\
& ([t \mapsto (l := \llbracket L \rrbracket^L[\text{par9}(l_1, l_2), \mathbf{X}], \{(t, t') \mid t' \in T\}, [l \mapsto ()]) \\
\llbracket L[(\text{forkjoin } ((e_1)^{a_1}, (e_2)^{a_2}))^{(10, l_1, l_2)}] \rrbracket(l, t) &= \\
& \llbracket (e_1)^{a_1} \rrbracket(l_1, t_1, \{t\}) \uplus \llbracket (e_2)^{a_2} \rrbracket(l_2, t_2, \{t\}) \uplus \\
& ([t \mapsto (l := \llbracket L \rrbracket^L[(! l_1, ! l_2), \mathbf{R}], \{(t, t') \mid t' \in T\}, [l \mapsto ()]) \\
\llbracket L[(\text{forkjoin } ((e_1)^{a_1}, (e_2)^{a_2}))^{(11, l_1, l_2)}] \rrbracket(l, t) &= \\
& [t \mapsto (l := \llbracket L \rrbracket^L[\llbracket v_1 \rrbracket^{\mathcal{S}}, ! l_2), \mathbf{X}], \{(t, t') \mid t' \in T\}, [l \mapsto (), l_1 \mapsto \llbracket v_1 \rrbracket^{\mathcal{S}}, l_2 \mapsto \llbracket v_2 \rrbracket^{\mathcal{S}}]
\end{aligned}$$

The macrodefinitions are successive partial evaluations of the translation of the parallel composition. These are as follows:

```

1  par1(e1, e2, l1) =
2    let l2 = alloc
3      t1 = newTd (l1 := e1)
4      t2 = newTd (l2 := e2)
5      t = self()
6    in newEdge (t1, t); newEdge (t2, t);
7      release t1; release t2; yield();
8      (!l1, !l2)
9
10 par2(e1, e2, l1, l2) =
11   let t1 = newTd (l1 := e1)
12     t2 = newTd (l2 := e2)
13     t = self()
14   in newEdge (t1, t); newEdge (t2, t);
15     release t1; release t2; yield();
16     (!l1, !l2)
17
18 par3(e2, l1, l2, t1) =
19   let t2 = newTd (l2 := e2)
20     t = self()
21   in newEdge (t1, t); newEdge (t2, t);
22     release t1; release t2; yield();
23     (!l1, !l2)
24
25 par4(l1, l2, t1, t2) =
26   let t = self()
27   in newEdge (t1, t); newEdge (t2, t);
28     release t1; release t2; yield();
29     (!l1, !l2)
30
31 par5(l1, l2, t1, t2, t) =
32   newEdge (t1, t); newEdge (t2, t);
33   release t1; release t2; yield();
34   (!l1, !l2)
35
36 par6(l1, l2, t1, t2, t) =
37   newEdge (t2, t);
38   release t1; release t2; yield();
39   (!l1, !l2)
40
41 par7(l1, l2, t1, t2) =
42   release t1; release t2; yield();
43   (!l1, !l2)
44
45 par8(l1, l2, t2) =
46   release t2; yield(); (!l1, !l2)
47
48 par9(l1, l2) =
49   yield (); (!l1, !l2)

```

Correctness With the translations defined, we can connect the annotated expressions with the dag calculus states (through the definition of *matching* states ∞). Then, Lemma 5 allows us to locate any executing thread as a subterm of our computation. This is of

crucial importance, since this is what restricts the reduction steps that the dag calculus can take in that thread: the code must come from the translation of that subcomputation. Using this property, we can prove Lemma 6, which states that every step on the dag calculus side is either a scheduler step, or can be matched on the side of the instrumented expressions. Since we know the precise subexpression where the reduction happens, we can generally restrict ourselves to ensuring that the administrative reductions on the instrumented expression side match, step for step, the execution of the translation of the parallel pair — which is how the translation was designed in the first place. Finally, Lemma 7 gives us the backwards simulation of dag calculus with the instrumented semantics, and the Correctness Theorem composes it with the instrumentation lemma, Lemma 4. The proof process is much the same for the other translations.

Definition 4. We say that a well-formed annotated instrumented expression e^a such that $\text{wf}(e^a)$ matches a dag calculus state V, E, σ , written $e^a \times V, E, \sigma$ if there exists a dag state V', E', σ' such that $\sigma' \sqsubseteq \sigma$ and $\llbracket e^a \rrbracket(\emptyset) = V', E', \sigma'$, and a finite map of vertices V'' such that $\forall t \in \text{dom}(V''). V''(t) = ((), \mathbb{R})$, and that

$$V' \uplus V'', E', \sigma \rightarrow_s^* V, E, \sigma.$$

Lemma 5. For any e^a, V, E, σ, t if $\text{wf}(e^a), e^a \times V, E, \sigma, t \in \text{dom}(V)$ and $\text{status}(V(t)) = \mathbb{X}$, then either

1. $V(t) = ((), \mathbb{X})$,
2. there exist K, e', l such that $e^a = K[e'^{(1,l,t)}]$,

Lemma 6. For any $e^a, V, V', E, E', \sigma, \sigma'$ if $\text{wf}(e^a), e^a \times V, E, \sigma$ and $V, E, \sigma \rightarrow V', E', \sigma'$ then either $e^a \times V', E', \sigma'$ or there exists $e'^{a'}$ such that $e^a \rightarrow e'^{a'}$ and $e'^{a'} \times V', E', \sigma'$.

Lemma 7. For any $e, a, V, V', E, \sigma, \sigma'$, if $\text{wf}(e^a), e^a \times V, E, \sigma$ and $\forall t \in \text{dom}(V'). V'(t) = ((), \mathbb{F})$, then the following simulation holds:

$$\begin{aligned} V, E, \sigma \rightarrow^* V', \emptyset, \sigma' &\Rightarrow \exists v. e^a \rightarrow^* v^{(2, \text{loc}(a))} \wedge \sigma'(l) = \llbracket v \rrbracket^S \\ V, E, \sigma \rightarrow^\infty &\Rightarrow e^a \rightarrow^\infty \end{aligned}$$

Proof. Follows from Lemmas 6 and 1 by the same inductive and coinductive argument as used in the proof of Lemma 4. \square

Theorem 1 (Correctness). Let t be the identifier of the main thread, $e \in \mathcal{E}_O$ be the source program stored in this thread, and l be a designated location in which to store the final result. Let $e_l \in \mathcal{E}_l$ be e annotated with the 0 annotations on all parallel compositions. For any integer result n , final state σ such that $\sigma(l) = n$, and final set of vertices V , assuming that all threads t' in V are finished (i.e. $\text{status}(V(t')) = \mathbb{F}$), we have:

$$[t \mapsto (l := \llbracket e_l \rrbracket^S, \mathbb{R}), \emptyset, [l \mapsto ()] \rightarrow^* V, \emptyset, \sigma \Rightarrow e \rightarrow_O^* n.$$

Furthermore, divergence in the dag calculus entails divergence in the source language:

$$[t \mapsto (l := \llbracket e_l \rrbracket^S, \mathbb{R}), \emptyset, [l \mapsto ()] \rightarrow^\infty \Rightarrow e \rightarrow_O^\infty.$$

Proof. This theorem follows from the composition of the two simulation diagrams, given by Lemmas 4 and 7. Clearly, $\llbracket e_I \rrbracket = e$. Moreover, since $\text{surf}(e_I)$, we also have $\llbracket e_I^{(1,l,t)} \rrbracket(\emptyset) = [t \mapsto (l := \llbracket e_I \rrbracket^S, \mathbb{R})], \emptyset, [l \mapsto ()]$, which gives us $e_I^{(1,l,t)} \propto [t \mapsto (l := \llbracket e_I \rrbracket^S, \mathbb{R})], \emptyset, [l \mapsto ()]$.

For the termination case, by Lemma 7 we obtain that there exists a value v such that $e_I^{(1,l,t)} \rightarrow_I^* v^{(2,l)}$ and $\sigma'(l) = \llbracket v \rrbracket^S$. However, since $\sigma'(l) = n$, we can conclude that $v = n$. By Lemma 4 we can now conclude that $e \rightarrow_O^* \llbracket n \rrbracket$, which ends the proof, as $\llbracket n \rrbracket = n$.

For the nontermination case, from Lemma 7 conclude that $e_I^{(1,l,t)} \rightarrow_I^\infty$, which, by Lemma 4 implies that $e \rightarrow_O^\infty$. This ends the proof. \square

C Correctness of the translation of async-finish

C.1 Syntax and Operational Semantics

$$\begin{aligned}
e &::= v \mid \text{fst } v \mid \text{snd } v \mid \text{let } x = e \text{ in } e \mid v \ v \\
&\quad \mid \text{alloc} \mid v := v \mid !v \mid \text{async}(e) \mid \text{finish}(S) \\
v &::= x \mid n \mid l \mid (v, v) \mid \text{fun } f \ x \text{ is } e \text{ end} \\
S &\equiv \{e_1, e_2, \dots, e_n\} \\
K &::= L \mid L[\text{finish}(\{K\} \uplus S)] \\
L &::= \bullet \mid \text{let } x = L \text{ in } e
\end{aligned}$$

The input programs are restricted to use the form $\text{finish}(\{e\})$. The operational semantics is given below. It operates on the store $\mathcal{H} = \mathcal{L} \rightarrow_{\text{fin}} \mathcal{V}$

$$\begin{array}{c}
\frac{}{h, \text{fst } (v_1, v_2) \rightarrow v_1, h} \quad \frac{}{h, \text{snd } (v_1, v_2) \rightarrow v_2, h} \quad \frac{}{h, \text{let } x = v \text{ in } e \rightarrow e[x \mapsto v], h} \\
\frac{}{h, \text{fun } f \ x \text{ is } e \text{ end } v \rightarrow e[x \mapsto v, f \mapsto \text{fun } f \ x \text{ is } e \text{ end}], h} \\
\frac{l \notin \text{dom}(h)}{h, \text{alloc} \rightarrow l, h[l \mapsto ()]} \quad \frac{h(l) = v}{h, !l \rightarrow v, h} \quad \frac{l \in \text{dom}(h)}{h, l := v \rightarrow (), h[l \mapsto v]} \\
\frac{}{h, \text{finish}(\{L[\text{async}(e)]\} \uplus S) \rightarrow \text{finish}(\{L[()], e\} \uplus S), h} \\
\frac{\forall e \in S. e = ()}{h, \text{finish}(S) \rightarrow (), h} \quad \frac{h, e \rightarrow e', h'}{h, K[e] \rightarrow K[e'], h'}
\end{array}$$

Note the operational semantics is changed slightly wrt. the paper: the finished async threads are not removed one by one, but rather stay as values associated with the enclosing finish block, which finishes when all the async computations terminate. This is done to facilitate the proof, since the removal of a single finished async would correspond to a scheduler transition (STOP), and it is easier to keep the scheduler transitions separate from the others. It is a simple exercise to show this presentation equivalent to the one presented in the paper: since the removals of finished async threads are nondeterministic,

they can all happen right before the finish block should reduce. The other direction is even simpler.

The following lemma allows us to only consider programs that write their final value to a predetermined location in the state in the following, since these can simulate the other programs.

Lemma 8. *For any program e and location l and result state h , we have the following properties:*

$$[l \mapsto ()], \text{let } x = e \text{ in } l := x \rightarrow^* () , h \Rightarrow \emptyset, e \rightarrow^* h(l), h \setminus [l \mapsto h(l)],$$

$$[l \mapsto ()], \text{let } x = e \text{ in } l := x \rightarrow^\infty \Rightarrow \emptyset, e \rightarrow^\infty .$$

Proof. Simple induction. □

C.2 Instrumented Syntax and Operational Semantics

When there is a chance of confusion between the instrumented and original language, we annotate the constructs of the instrumented language with a subscript I , and the constructs of original language with O . The instrumentations are a_a for async and a_f for finish. Additionally, we annotate roots of expressions with thread identifiers.

$$\begin{aligned} e &::= v \mid \text{fst } v \mid \text{snd } v \mid \text{let } x = e \text{ in } e \mid v v \mid \text{alloc} \mid v := v \mid !v \\ &\quad \mid \text{async}(e)^{a_a} \mid \text{finish}(e)^{a_f} \mid \text{finish}(S)^{a_g} \\ v &::= x \mid n \mid l \mid (v, v) \mid \text{fun } f \text{ x is } e \text{ end} \\ S &\equiv \{e_1^{t_1}, e_2^{t_2}, \dots, e_n^{t_n}\} \\ L &::= \bullet \mid \text{let } x = L \text{ in } e \\ K &::= \bullet \mid (L[\text{finish}(\{K\} \uplus S)^a])^t \\ a_a &::= 0 \mid (1, t) \mid (2, t) \\ a_f &::= 0 \mid 1 \mid (2, t) \mid (3, t) \\ a_g &::= 1 \mid 2 \end{aligned}$$

We define the surface expressions, $\text{surf}(e)$ as ones that only contain annotations of 0 and do not contain subexpressions of the form $\text{finish}(S)$. This notion extends in a trivial way to the sequential contexts L . With this, we can proceed to define the well-formed expressions, $\text{wf}(e)$:

$$\text{wf}(e) = \begin{cases} \text{surf}(e') \wedge \text{surf } L & \text{if } e = L[\text{async}(e')^a] \\ \text{surf}(e') \wedge \text{surf } L & \text{if } e = L[\text{finish}(e')^a] \\ (\forall e', t. (e') \in S \Rightarrow \text{wf}(e')) \wedge \text{surf } L & \text{if } e = L[\text{finish}(S)^a] \\ \text{surf } e & \text{otherwise} \end{cases}$$

This notion simply extends to contexts K .

For operational semantics, let us first define the allowed transitions on annotations, \rightsquigarrow_a for the async annotations, and \rightsquigarrow_f and \rightsquigarrow_g for the finish ones. The essence is that

the associated number can increase by one, and if any new thread identifiers appear, they must be globally fresh. Otherwise, thread ids are preserved.

$$\begin{aligned}
0 &\rightsquigarrow_a (1, t) \text{ where } t \text{ is fresh} & (1, t) &\rightsquigarrow_a (2, t) \\
0 &\rightsquigarrow_f 1 & 1 &\rightsquigarrow_f (2, t) \text{ where } t \text{ is fresh} \\
(2, t) &\rightsquigarrow_f (3, t) & 1 &\rightsquigarrow_g 2
\end{aligned}$$

Next, we can define the operational semantics for the instrumented language. For the most part, it is the same as the source language, with key differences being the reduction for annotation steps and the separate preparatory finish construct.

$$\begin{array}{c}
\frac{}{h, \text{fst}(v_1, v_2) \rightarrow v_1, h} \quad \frac{}{h, \text{snd}(v_1, v_2) \rightarrow v_2, h} \quad \frac{}{h, \text{let } x = v \text{ in } e \rightarrow e[x \mapsto v], h} \\
\frac{}{h, \text{fun } f \text{ x is } e \text{ end } v \rightarrow e[x \mapsto v, f \mapsto \text{fun } f \text{ x is } e \text{ end}], h} \\
\frac{l \notin \text{dom}(h)}{h, \text{alloc} \rightarrow l, h[l \mapsto ()]} \quad \frac{h(l) = v}{h, !l \rightarrow v, h} \quad \frac{l \in \text{dom}(h)}{h, l := v \rightarrow (), h[l \mapsto v]} \\
\frac{t \text{ is fresh}}{h, \text{async}(e)^0 \rightarrow \text{async}(e)^{(1,t)}} \quad \frac{}{h, \text{finish}(\{(L[\text{async}(e)^{(1,t)}]\}' \uplus S\}) \rightarrow \text{finish}(\{(L[\text{async}(e)^{(2,t)}]\}' \uplus S\}), h} \\
\frac{}{h, \text{finish}(\{(L[\text{async}(e)^{(2,t)}]\}' \uplus S\}) \rightarrow \text{finish}(\{(L[()] \uplus S\})', e'\} \uplus S), h} \\
\frac{a \rightsquigarrow_f a'}{h, \text{finish}(e)^a \rightarrow \text{finish}(e)^{a'}, h} \quad \frac{}{h, \text{finish}(e)^{(3,t)} \rightarrow \text{finish}(\{e'\}^1, h)} \\
\frac{}{h, \text{finish}(S)^1 \rightarrow \text{finish}(S)^2, h} \quad \frac{\forall e, t. e' \in S \Rightarrow e = ()}{h, \text{finish}(S)^2 \rightarrow (), h} \quad \frac{h, e \rightarrow e', h'}{h, K[(L[e])'] \rightarrow K[(L[e'] \uplus S)'], h'}
\end{array}$$

Lemma 9. *Evaluation preserves well-formedness, i.e., for any e, e', h, h' , if $\text{wf}(e)$ and $h, e \rightarrow e', h$, then $\text{wf}(e')$.*

C.3 Connecting instrumented language and source language

In order to connect the instrumented language expressions to the source language ones, we introduce an *erasure* function: $[-] : \mathcal{E}_I \rightarrow \mathcal{E}_O$, that removes annotations from terms and maps $\text{finish}(e)^a$ to $\text{finish}(\{[e]\})$. Erasure function extends in the obvious way to heaps and evaluation contexts.

Lemma 10 (Instr-Step). *For any expressions $e, e' \in \mathcal{E}_I$ and heaps $h, h' \in \mathcal{H}_I$, if $h, e \rightarrow_I e', h'$, then either $[e] = [e']$ and $h = h'$, or $[h], [e] \rightarrow_O [e'], [h']$.*

Proof. By induction on the reduction judgment. These split into three forms: the administrative reductions of `async` and `finish`, which give the same erasure for both the redex and the reduct, the reductions of sequential constructs, which can be trivially matched, and the two rules that spawn an `async` thread and terminate the `finish` block — which also can be easily matched. \square

Definition 5. An *administrative* reduction in the instrumented language is a reduction $h, e \rightarrow e', h'$ such that $[e] = [e']$ and $h = h'$. We write $h, e \rightarrow_a e', h'$ for administrative steps.

Lemma 11 (Admin-Fin). *There are no infinite sequences of administrative reductions, i.e., for any $e \in \mathcal{E}_I$ and $h \in \mathcal{H}_I$, if $\text{wf}(e)$ then $h, e \not\rightarrow_a^\infty$.*

Proof. Assign as a weight of the expression based on the number in each of its annotations in evaluation positions. Take $\text{num}(a)$ to denote the number associated with the annotation a . Then define

$$w(e) = \begin{cases} 2 - \text{num}(a) & \text{if } e = L[\text{async}(e')^a] \\ 3 - \text{num}(a) & \text{if } e = L[\text{future}(e')^a] \\ 2 - a + \sum_{e' \in S} w(e) & \text{if } e = L[\text{future}(S)^a] \\ 0 & \text{otherwise} \end{cases}.$$

Observe that each administrative reduction decreases the weight. The lemma then holds by simple induction on the weight. \square

Lemma 12 (Instrument). *For any expression $e \in \mathcal{E}_I$ and heaps $h, h' \in \mathcal{H}_I$ such that $h, e \rightarrow_I^* ()$, $h', [h], [e] \rightarrow_O^* ()$, $[h']$. Moreover, if $h, e \rightarrow_I^\infty$ then $[h], [e] \rightarrow_O^\infty$.*

Proof. First part, by induction on the reduction sequence and Lemma 10.

Similarly, for the second part we proceed by coinduction: since $h, e \rightarrow_I^\infty$, by Lemma 18 and 10 there exist e' and h' such that $h, e \rightarrow_I^* e', h', [h], [e] \rightarrow_O [e'], [h']$ and $h', e' \rightarrow_I^\infty$. Hence, $[h'], [e'] \rightarrow_O^\infty$ and finally $[h], [e] \rightarrow_O^\infty$. \square

C.4 Connecting instrumented language and dag calculus

Translation of surface expressions into dag calculus Note that the translation trivially extends to the evaluation contexts of the instrumented language. We define a translation $\llbracket - \mid - \rrbracket^S : \mathcal{V}_{\text{dagcalculus}} \times \mathcal{E}_I \rightarrow \mathcal{E}_{\text{dagcalculus}}$, where the first argument keeps track of the *nearest enclosing finish clock*. Note that this differs from the form presented in the paper only in the order of the arguments. The bulk of the translation is a standard destination-passing translation:

$$\begin{aligned} \llbracket t \mid \text{let } x = e_1 \text{ in } e_2 \rrbracket &= \text{let } x = \llbracket t \mid e_1 \rrbracket \text{ in } \llbracket t \mid e_2 \rrbracket \\ \llbracket t \mid \text{fun } f \text{ } x \text{ is } e \text{ end} \rrbracket &= \text{fun } f(x, t') \text{ is } \llbracket t' \mid e \rrbracket \text{ end} \\ \llbracket t \mid v_1 \ v_2 \rrbracket &= \llbracket t \mid v_1 \rrbracket (\llbracket t \mid v_2 \rrbracket, t) \end{aligned}$$

The translation for parallel primitives is defined as follows:

```

1  $\llbracket t \mid \text{async}(e)^0 \rrbracket =$ 
2   let  $t' = \text{newTd } (\llbracket t \mid e \rrbracket)$ 
3   in  $\text{newEdge } (t', t); \text{release } t'; ()$ 
4
5  $\llbracket t \mid \text{finish}(e)^0 \rrbracket =$ 
6   let  $t_2 = \text{self } ()$ 
7    $t_1 = \text{newTd } (\llbracket t_2 \mid e \rrbracket)$ 
8   in  $\text{newEdge } (t_1, t_2); \text{release } t_1; \text{yield}(); ()$ 

```

Note that due to its destination-passing flavour, the translation for values does not depend on the thread identifier. Thus, we can also extend this translation to work on heaps: we denote it $\llbracket - \rrbracket^H : \mathcal{H}_I \rightarrow \mathcal{H}_{\text{dagcalculus}}$.

Translation of the well-formed expressions into dag calculus The translation for well-formed, thread-annotated expressions is as follows, defined mutually recursively with the translation for expressions. The latter has the form $\llbracket - \rrbracket : \mathcal{E}_I \rightarrow \text{TId} \times \text{TId}^+ \rightarrow \mathcal{D}$, where \mathcal{D} denotes computation dags of the form (V, E) . The argument that tracks the continuation vertex is \perp outside the topmost finish block. We define a helper function

$$\text{mkEdge}(t, t') = \begin{cases} \emptyset & \text{if } t' = \perp \\ \{(t, t')\} & \text{otherwise} \end{cases}$$

The definition for thread-annotated expressions is just shorthand: the thread-annotation is passed as an argument to the main translation.

$$\llbracket e' \rrbracket(t') = \llbracket e \rrbracket(t, t')$$

$$\begin{aligned} \llbracket e \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid e \rrbracket^S, \mathbf{R}), \text{mkEdge}(t, t')] \quad \text{if } \text{surf}(e) \\ \llbracket L[\text{async}(e)^{(1, t_a)}] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{async1}(t_a, t'), \mathbf{X}], t_a \mapsto (\llbracket t' \mid e \rrbracket^S, \mathbf{N}), \text{mkEdge}(t, t')] \\ \llbracket L[\text{async}(e)^{(2, t_a)}] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{release } t_a, \mathbf{X}], t_a \mapsto (\llbracket t' \mid e \rrbracket^S, \mathbf{N}), \text{mkEdge}(t, t') \cup \text{mkEdge}(t_a, t')] \\ \llbracket L[\text{finish}(e)^1] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{finish1}(\llbracket t \mid e \rrbracket^S, t), \mathbf{X}], \text{mkEdge}(t, t')] \\ \llbracket L[\text{finish}(e)^{(2, t_a)}] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{finish2}(t_a, t), \mathbf{X}], t_a \mapsto (\llbracket t \mid e \rrbracket^S, \mathbf{N}), \text{mkEdge}(t, t')] \\ \llbracket L[\text{finish}(e)^{(3, t_a)}] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{finish3}(t_a, t), \mathbf{X}], t_a \mapsto (\llbracket t \mid e \rrbracket^S, \mathbf{N}), \text{mkEdge}(t, t') \cup \{(t_a, t)\}] \\ \llbracket L[\text{finish}(S)^1] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{yield}(); \text{C}], \mathbf{X}), \text{mkEdge}(t, t')] \uplus \bigsqcup_{e' \in S} \llbracket e' \rrbracket(t) \\ \llbracket L[\text{finish}(S)^2] \rrbracket(t, t') &= \\ & [t \mapsto (\llbracket t' \mid L \rrbracket^L[\text{C}], \mathbf{R}), \text{mkEdge}(t, t')] \uplus \bigsqcup_{e' \in S} \llbracket e' \rrbracket(t) \end{aligned}$$

The macrodefinitions used above are as follows:

```

1  async1(t, t') = newEdge (t, t'); release t; C
2
3  finish1(e, t') =
4    let t = newTd (e)
5    in newEdge (t, t'); release t; yield(); C
6

```

```

7 finish2(t, t') =
8   newEdge (t, t'); release t; yield(); ()
9
10 finish3(t) =
11   release t; yield(); ()

```

Correctness

Definition 6. We say that a configuration of annotated instrumented expression and state e^t, h such that $\text{wf}(e^t)$ and $\text{surf}(h)$ matches a dag calculus state V, E, σ , written $e^t, h \propto V, E, \sigma$ if $\sigma = \llbracket h \rrbracket^H$, and there exists a dag (V', E') such that $\llbracket (C) \mid e^t \rrbracket = V', E'$ and a finite map of vertices V'' such that $\forall t \in \text{dom}(V'')$. $V''(t) = (C), R$, and that

$$V' \uplus V'', E', \sigma \rightarrow_s^* V, E, \sigma.$$

Lemma 13. For any e^t, h, V, E, σ, t if $\text{wf}(e)$, $e^t, h \propto V, E, \sigma$, $t \in \text{dom}(V)$ and $\text{status}(V(t)) = X$, then either $V(t) = (C), X$, or there exist K, e' such that $e^t = K[e']$.

Lemma 14. For any $e^t, h, V, V', E, E', \sigma, \sigma'$ if $\text{wf}(e^t)$, $e^t, h \propto V, E, \sigma$ and $V, E, \sigma \rightarrow V', E', \sigma'$ then either $e^t, h \propto V', E', \sigma'$ or there exist e'^t, h' such that $h, e^t \rightarrow e'^t, h'$ and $e'^t, h' \propto V', E', \sigma'$.

Lemma 15. For any $e, t, h, V, V', E, \sigma, \sigma'$, if $\text{wf}(e^t)$, $e^t, h \propto V, E, \sigma$ and $\forall t' \in \text{dom}(V')$. $V'(t') = (C), F$, then the following simulation holds:

$$\begin{aligned} V, E, \sigma \rightarrow^* V', \emptyset, \sigma' &\Rightarrow \exists v, t', h'. h, e^t \rightarrow^* v', h' \wedge \sigma' = \llbracket h' \rrbracket^H \\ V, E, \sigma \rightarrow^\infty &\Rightarrow h, e^t \rightarrow^\infty \end{aligned}$$

Proof. Follows from Lemmas 14 and 1 by the same inductive and coinductive argument as used in the proof of Lemma 12. \square

Theorem 2 (Correctness). Let t be the identifier of the main thread, e be the source program stored in this thread, l be a designated location in which to store the final result, and e_l be the result of annotating e with 0 on all `finish` and `async` constructs. For any integer result n , final state σ such that $\sigma(l) = n$, and final set of vertices V , assuming that all threads t' in V are finished (i.e. $\text{status}(V(t')) = F$), we have:

$$[t \mapsto (\text{let } x = \llbracket (C) \mid e_l \rrbracket^S \text{ in } l := x, R)], \emptyset, [l \mapsto (C)] \rightarrow^* V, \emptyset, \sigma \Rightarrow \exists h. \emptyset, e \rightarrow_O^* n, h$$

Furthermore, divergence in the dag calculus entails divergence in the source language:

$$[t \mapsto (\text{let } x = \llbracket (C) \mid e_l \rrbracket^S \text{ in } l := x, R)], \emptyset, [l \mapsto (C)] \rightarrow^\infty \Rightarrow \emptyset, e \rightarrow_O^\infty.$$

Proof. This theorem follows from the composition of the two simulation diagrams, given by Lemmas 12 and 15. Clearly, $[e_l] = e$. Moreover, since $\text{surf}(e_l) = e$, we also have $\llbracket (\text{let } x = e_l \text{ in } l := x)^t \rrbracket(\perp) = [t \mapsto (\text{let } x = \llbracket (C) \mid e_l \rrbracket^S \text{ in } l := x, R)], \emptyset$, which gives us $(\text{let } x = e_l \text{ in } l := x)^t, [l \mapsto ()] \propto [t \mapsto (\text{let } x = \llbracket (C) \mid e_l \rrbracket^S \text{ in } l := x, R)], \emptyset, [l \mapsto ()]$.

For the termination case, by Lemma 15 we obtain that there exist v, t' and h such that $[l \mapsto ()], (\text{let } x = e_l \text{ in } l := x)^t \rightarrow_{t'}^* v', h$ and $\sigma' = \llbracket h \rrbracket^H$ (and hence $h(l) = n$).

Since the final command of the program is assignment, we can conclude that $v = ()$. By Lemma 12 we can now conclude that $[l \mapsto ()], \text{let } x = e \text{ in } l := x \rightarrow_O^* (), [h]$. Since $h(l) = n$, by Lemma 8 this gives us that $\emptyset, e \rightarrow_O^* n, [h] \setminus [l \mapsto n]$, which ends the proof.

For the nontermination case, from Lemma 15 conclude that $[l \mapsto ()], (\text{let } x = e_l \text{ in } l := x)^t \rightarrow_l^\infty$, which, by Lemma 12 implies that $[l \mapsto ()], \text{let } x = e \text{ in } l := x \rightarrow_O^\infty$. Finally, by Lemma 8 we obtain $\emptyset, e \rightarrow_O^\infty$ which ends the proof. \square

D Correctness of the translation of futures

D.1 Syntax and Operational Semantics

The syntax and semantics in this section are the same as in the paper and repeated here for convenience.

$$\begin{aligned}
e &::= x \mid n \mid f \mid (v, v) \mid \text{fst } v \mid \text{snd } v \mid \text{let } x = e \text{ in } e \mid v \ v \\
&\quad \mid \text{fun } g \ x \text{ is } e \text{ end} \mid \text{future}(e) \mid \text{force}(v) \\
v &::= x \mid n \mid f \mid (v, v) \mid \text{fun } g \ x \text{ is } e \text{ end} \\
K &::= \bullet \mid \text{let } x = K \text{ in } e
\end{aligned}$$

We take the set of configurations $\mathcal{M} \equiv \mathcal{F} \rightarrow_{\text{fin}} \mathcal{E}$, where \mathcal{F} is the set of future identities and \mathcal{E} — the set of expressions, and write $M \in \mathcal{M}$ for these configurations. We also pick an identifier of the main program, f_0 : thus, the initial state is $[f_0 \mapsto e]$ for some program e . The evaluation reaches a terminal state when all the futures in M have values associated with them.

$$\begin{aligned}
\text{fst } (v_1, v_2) \circ \rightarrow v_1 \quad \text{snd } (v_1, v_2) \circ \rightarrow v_2 \quad \text{let } x = v \text{ in } e \circ \rightarrow e[x \mapsto v] \\
\text{fun } g \ x \text{ is } e \text{ end } v \circ \rightarrow e[x \mapsto v, g \mapsto \text{fun } g \ x \text{ is } e \text{ end}]
\end{aligned}$$

$$\begin{aligned}
&\frac{M(f) = K[e] \quad e \circ \rightarrow e'}{M \rightarrow M[f \mapsto K[e']]} \\
&\frac{M(f) = K[\text{future}(e)] \quad f' \text{ fresh}}{M \rightarrow M[f \mapsto K[f'], f' \mapsto e]} \\
&\frac{M(f) = K[\text{force}(f')] \quad M(f') = v}{M \rightarrow M[f \mapsto K[v]]}
\end{aligned}$$

D.2 Instrumented Syntax and Operational Semantics

The instrumented language is presented below. The input terms can only be annotated with 0. When there is a chance of confusion between the instrumented and original

language, we annotate the constructs of the instrumented language with a subscript I , and the constructs of original language with O . The instrumentations are a_s for future creation, a_f for future access, and a_w for future evaluation.

$$\begin{aligned}
e &::=x \mid n \mid t \mid (v, v) \mid \text{fst } v \mid \text{snd } v \mid \text{let } x = e \text{ in } e \mid v \ v \\
&\quad \mid \text{fun } f \ x \ \text{is } e \ \text{end} \mid \text{future}(e)^{a_s} \mid \text{force}(v)^{a_f} \\
v &::=x \mid n \mid t \mid (v, v) \mid \text{fun } f \ x \ \text{is } e \ \text{end} \\
a_s &::=0 \mid (1, l) \mid (2, l, t) \\
a_f &::=n \text{ where } 0 \leq n \leq 4 \\
a_w &::=0 \mid 1 \\
K &::=\bullet \mid \text{let } x = K \ \text{in } e
\end{aligned}$$

The annotations a_s and a_f can *progress*, written $a \rightsquigarrow_s a'$ (and, respectively $a \rightsquigarrow_f a'$), if the step of a' is one greater than the step of a . All the other parts of a' present in a have to agree, while those absent in a are globally fresh, as follows:

$$\begin{aligned}
0 &\rightsquigarrow_s (1, l) \text{ where } l \text{ is fresh} & (1, l) &\rightsquigarrow_s (2, l, t) \text{ where } t \text{ is fresh} \\
n &\rightsquigarrow_f n + 1 \text{ for } 0 \leq n \leq 3
\end{aligned}$$

Sequential evaluation within a future proceeds as follows:

$$\begin{array}{c}
\text{fst } (v_1, v_2) \circlearrowright v_1 \quad \text{snd } (v_1, v_2) \circlearrowright v_2 \quad \text{let } x = v \ \text{in } e \circlearrowright e[x \mapsto v] \\
\text{fun } f \ x \ \text{is } e \ \text{end } v \circlearrowright e[x \mapsto v, f \mapsto \text{fun } f \ x \ \text{is } e \ \text{end}] \\
\frac{a \rightsquigarrow_f a'}{\text{future}(e)^a \circlearrowright \text{future}(e)^{a'}} \quad \frac{a \rightsquigarrow_s a'}{\text{force}(e)^a \circlearrowright \text{force}(e)^{a'}}
\end{array}$$

The configurations are defined as $\mathcal{M} = \mathcal{F} \rightarrow_{\text{fin}} (\mathcal{L} \times \text{Tid} \times \mathcal{E} \times a_w)$, and well-formed configurations have all the locations and thread id's disjoint. The special starting future, f_0 has a special first thread id, t_0 . The evaluation then proceeds as follows.

$$\begin{array}{c}
\frac{M(f) = (l, t, K[e], 0) \quad K[e] \circlearrowright K[e']}{M \rightarrow M[t \mapsto (l, t, K[e'], 0)]} \\
\frac{M(f) = (l, t, K[\text{future}(e)^{(3, l', t')}], 0) \quad f' \text{ fresh}}{M \rightarrow M[f \mapsto (l, t, K[f'], 0), f' \mapsto (l', t', e, 0)]} \\
\frac{M(f) = (l, t, K[\text{force}(f')^4], 0) \quad M(f') = (l', t', v, 1)}{M \rightarrow M[f \mapsto (l, t, K[v], 0)]} \\
\frac{M(f) = (l, t, v, 0)}{M \rightarrow M[f \mapsto (l, t, v, 1)]}
\end{array}$$

An expression is a *surface* expression, written $\text{surf}(e)$ if all its annotations are 0. Note that these are isomorphic to the terms of the original language. A term is

well-formed, written $\text{wf}(e)$, if it has at most one non-zero annotation, in an evaluation position, i.e.,

$$\text{wf}(e) \equiv \begin{cases} \text{wf}(e_1) \wedge \text{surf}(e_2) & \text{if } e = \text{let } x = e_1 \text{ in } e_2 \\ \text{surf}(e') & \text{if } e = \text{future}(e')^a \\ \text{surf}(v) & \text{if } e = \text{force}(v)^a \\ \text{surf}(e) & \text{otherwise} \end{cases}$$

We write $\text{wf}(M)$ to mean that each of the expressions associated with the futures is well-formed.

Lemma 16. *Evaluation preserves well-formedness, i.e., if $\text{wf}(e)$ and $e \circ \rightarrow e'$ then $\text{wf}(e')$. Similarly, if $\text{wf}(M)$ and $M \rightarrow M'$, then $\text{wf}(M')$.*

Proof. The first part proceeds by simply checking the reductions. The second follows by simple induction on the structure of the expression associated with the future that reduces. \square

D.3 Connecting the instrumented and source languages

We define a map $\lfloor - \rfloor : \mathcal{E}_I \rightarrow \mathcal{E}_O$ that removes annotations from the instrumented expressions. This map extends to the configurations, by also removing the additional information in the futures. We use it to connect the instrumented language to the source language. Note that a map that adds an annotation 0 to `future` and `force` is a right inverse of $\lfloor - \rfloor$. For *surface* term, it is also the left inverse.

Lemma 17 (Instr-Step). *For any configurations $M, M' \in \mathcal{M}_I$, if $M \rightarrow_I M'$, then either $\lfloor M \rfloor = \lfloor M' \rfloor$ or $\lfloor M \rfloor \rightarrow_O \lfloor M' \rfloor$.*

Proof. By cases on the reduction. In cases of spawning a future, forcing a future of finishing a future, trivially true, since $\lfloor K[e] \rfloor = \lfloor K \rfloor \lfloor e \rfloor$. For the same reason, suffices to check the same statement for any $e \circ \rightarrow e'$.

These split into two forms: the administrative reductions of `future` and `force`, which give the same erasure for both the redex and the reduct, and the reductions of sequential constructs, which can be trivially matched. \square

Definition 7. An *administrative* reduction in the instrumented language is a reduction $M \rightarrow M'$ such that $\lfloor M \rfloor = \lfloor M' \rfloor$. We write $M \rightarrow_a M'$ for administrative steps.

Lemma 18 (Admin-Fin). *For well-formed configurations M , there are no infinite sequence of administrative reductions, i.e., for any $M \in \mathcal{M}_I$, if $\text{wf } M$ then $M \not\rightarrow_a^\infty$.*

Proof. Assign as a weight to each future based on the count of their annotations in evaluation positions. Take the $\text{cnt}(a)$ to be the natural number in the annotation. Then define

$$w(e) = \begin{cases} 3 - \text{cnt}(a) & \text{if } e = \text{future}(e')^a \\ 4 - a & \text{if } e = \text{force}(e')^a \\ w(e_1) & \text{if } e = \text{let } x = e_1 \text{ in } e_2 \\ 0 & \text{otherwise} \end{cases},$$

$w(l, t_1, t_2, e, n) = w(e) + 1 - n$, and $w(M) = \sum_{f \in \text{dom}(M)} w(M(f))$. From there, the proof is analogous to the proof of Lemma 1: by simple induction, each administrative reduction decreases the weight, and since the weight of a given configuration is finite, there is no infinite chain of reductions administrative. \square

Lemma 19 (Instrument). *For any configurations M, M' such that M' is final and $M \rightarrow_I^* M'$, $\llbracket M \rrbracket \rightarrow_O^* \llbracket M' \rrbracket$ and $\llbracket M' \rrbracket$ is final. Moreover, if $M \rightarrow_I^\infty$, then $\llbracket M \rrbracket \rightarrow_O^\infty$.*

Proof. First part, by induction on the reduction sequence and Lemma 17. Since final configurations of the instrumented language map to final configurations of the source language, this ends the proof.

Similarly, for the second part we proceed by coinduction: since $M \rightarrow_I^\infty$, by Lemmas 18 and 17 there exists an M' such that $M \rightarrow_I^* M'$, $\llbracket M \rrbracket \rightarrow_O \llbracket M' \rrbracket$ and $M' \rightarrow_I^\infty$. By coinduction, we obtain that $\llbracket M' \rrbracket \rightarrow_O^\infty$, and so $\llbracket M \rrbracket \rightarrow_O^\infty$. \square

D.4 Connecting the instrumented language and dag calculus

Translation of surface expressions into dag calculus Note that the translation trivially extends to the evaluation contexts of the instrumented language. We define a translation $\llbracket - \rrbracket^S : \mathcal{E}_I \rightarrow \mathcal{M}_I \rightarrow \mathcal{E}_{dagcalculus}$.

$$\begin{aligned} \llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket^S(M) &= \text{let } x = \llbracket e_1 \rrbracket^S(M) \text{ in } \llbracket e_2 \rrbracket^S(M) \\ \llbracket \text{fun } f \text{ } x \text{ is } e \text{ end} \rrbracket^S(M) &= \text{fun } f \text{ } x \text{ is } \llbracket e \rrbracket^S(M) \text{ end} \\ \llbracket v_1 \ v_2 \rrbracket^S(M) &= \llbracket v_1 \rrbracket^S(M) \ \llbracket v_2 \rrbracket^S(M) \end{aligned}$$

The translations for the pairs, variables and numbers are analgous. For future id's we use the state M to map them to the thread-location pairs:

$$\llbracket f \rrbracket^S(M) = (t, l) \quad \text{where } M(f) = (l, t, \rightarrow, -).$$

For the other constructs, we have:

```

1  $\llbracket \text{future}(e)^0 \rrbracket^S =$ 
2    $\text{let } l = \text{alloc}$ 
3      $t = \text{newTd } (l := \llbracket e \rrbracket^S)$ 
4    $\text{in release } t; (t, l)$ 
5
6  $\llbracket \text{force}(v)^0 \rrbracket^S =$ 
7    $\text{let } (ft, fl) = \llbracket v \rrbracket^S$ 
8    $\text{in newEdge } (ft, \text{self}()); \text{yield}(); !fl$ 

```

As mentioned, since all expressions in the well-formed evaluation contexts are surface expressions, we can extend this translation to contexts, giving us $\llbracket - \rrbracket^K : \mathcal{K}_I \rightarrow \mathcal{K}_{dagcalculus}$.

Translation of well-formed futures into dag calculus Here, we define a translation $\llbracket - \rrbracket : \mathcal{E}_I \rightarrow \mathcal{L} \times \text{TId} \times \mathcal{M}_I \rightarrow \mathcal{S}$, where \mathcal{S} are the states of dag calculus. In case the expression is a surface expression, we simply use the surface translation and add the dag structure:

$$\llbracket e \rrbracket(l, t, M) = [t \mapsto (l := \llbracket e \rrbracket^S, \mathbb{R}), \emptyset, [l \mapsto ()].$$

This leaves us with non-surface cases: the creation and reading of futures. The translation of these constructs takes into account the partial evaluation of the code that their annotations denote. For the `future` construct, we have as follows:

$$\begin{aligned} \llbracket K[\text{future}(e)^{(1,l)}] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[\text{ft1}(\llbracket e \rrbracket^S, l), \mathbb{R}], \emptyset), [l_f \mapsto (), l \mapsto ()] \\ \llbracket K[\text{future}(e)^{(2,l,t)}] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[\text{ft2}(l, t), \mathbb{R}]; t \mapsto (\llbracket e \rrbracket^S, \mathbb{N}), \emptyset), [l_f \mapsto (), l \mapsto ()] \end{aligned}$$

In the above, the helper macros are defined as partial evaluation states of the surface translation of `future`:

```

1 ft1(e, l) =
2   let t = newTd (l := e)
3   in release t; (t, l)
4
5 ft2(l, t) =
6   release t; (t, l)

```

We tackle the `force` operation in a similar manner, taking throughout the definition $(t, l) = \llbracket v \rrbracket^S(M)$.

$$\begin{aligned} \llbracket K[\text{force}(v)^1] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[\text{fr1}(t, l), \mathbb{X}], \emptyset), [l_f \mapsto ()] \\ \llbracket K[\text{force}(v)^2] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[\text{fr2}(t, l, t_f), \mathbb{X}], \emptyset), [l_f \mapsto ()] \\ \llbracket K[\text{force}(v)^3] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[\text{fr3}(l), \mathbb{X}], \{(t, t_f)\}), [l_f \mapsto ()] \\ \llbracket K[\text{force}(v)^4] \rrbracket(l_f, t_f, M) &= \\ & [t_f \mapsto (l_f := \llbracket K \rrbracket^K[!l, \mathbb{R}], \{(t, t_f)\}), [l_f \mapsto ()] \end{aligned}$$

Similar to before, we use helper macros that define partial evaluation of the surface translation of `force`:

```

1 fr1(ft, fl) =
2   newEdge (ft, self()); yield(); !fl
3
4 fr2(ft, fl, t) =
5   newEdge (ft, t); yield(); !fl
6
7 fr3(fl) =
8   yield(); !fl

```

Finally we define the translation for a future — and then as a translation of the complete state we take a disjoint union of the translation of the individual futures. We have:

$$\begin{aligned} \llbracket (l, t, e, 0) \rrbracket(M) &= \llbracket e \rrbracket(l, t, M) \\ \llbracket (l, t, v, 1) \rrbracket(M) &= \emptyset, \emptyset, [l \mapsto \llbracket v \rrbracket^S(M)] \end{aligned}$$

Definition 8. We say that an instrumented configuration M *matches* a dag Calculus state V, E, σ , written $M \propto V, E, \sigma$ if there exists a dag state V', E' such that $\llbracket M \rrbracket = V', E', \sigma$ and a finite map of threads V'' such that $\forall t \in \text{dom}(V'')$. $V''(t) = (\text{C}), \text{R}$ and that

$$V' \uplus V'', E', \sigma \rightarrow_s^* V, E, \sigma.$$

Lemma 20. For any M, V, E, σ, t if $\text{wf}(M)$, $M \propto V, E, \sigma$, $t \in \text{dom}(V)$ and $\text{status}(V(t)) = \text{X}$, then either $V(t) = (\text{C}), \text{X}$ or there exist f, l and e such that $M(f) = (l, t, e, 0)$.

Lemma 21. For any $M, V, V', E, E', \sigma, \sigma'$ if $\text{wf}(M)$, $M \propto V, E, \sigma$ and $V, E, \sigma \rightarrow V', E', \sigma'$ then either $M \propto V', E', \sigma'$ or there exists an M' such that $M \rightarrow M'$ and $M' \propto V', E', \sigma'$.

Proof. Consider the reduction $V, E, \sigma \rightarrow V', E', \sigma'$. It is either a START reduction, in which case it is an administrative reduction, and $M \propto V', E', \sigma'$, or it happens in some thread $t \in \text{dom}(V)$ such that $\text{status}(V(t)) = \text{X}$. Thus, by Lemma 20 we either have $V(t) = (\text{C}), \text{X}$ or there exists a future f such that $M(f) = (l, t, e, 0)$. In the first case, the only applicable reduction rule is STOP, which is an administrative reduction. This leaves us with the second case.

In this case we have $M(f) = (l, t, e, 0)$ and a reduction in thread t . We can decompose e uniquely into K and e' such that $e = K[e']$ and e' is not a let-binding. If e' is neither a value nor a parallel construct, then it's at $\llbracket e' \rrbracket^S(M)$ that the reduction occurs on the dag calculus side. Since the translation is entirely structural, we can match it by using the corresponding rule. The case if e' is a value is similar: the redex is the innermost let-binding in K , and can be matched by the let-rule. Finally, for parallel primitives, the location of the redex is based on the annotation: however, the translations of the sequences of annotated primitives are specifically designed so that they correspond to one step of reduction on the dag calculus side. Thus, we can locate the redex, obtain any new threads or locations created on the dag calculus side, and use them in the matching reduction. \square

Lemma 22. For any $M, V, V', E, \sigma, \sigma'$ if $\text{wf}(M)$, $M \propto V, E, \sigma$ and $\forall t \in \text{dom}(V')$. $V'(t) = (\text{C}), \text{F}$, then the following simulation holds:

$$\begin{aligned} V, E, \sigma \rightarrow^* V', \emptyset, \sigma' &\Rightarrow \exists M'. M \rightarrow_1^* M' \wedge M' \propto V', \emptyset, \sigma' \\ V, E, \sigma \rightarrow^\infty &\Rightarrow M \rightarrow_1^\infty \end{aligned}$$

Proof. Follows from Lemmas D.4 and 1 by the same inductive and coinductive argument as used in the proof of Lemma 19. \square

Theorem 3 (Correctness). Let t be the identifier of the main thread, e be the source program stored in this thread, and l be a designated location in which e stores its final result. Let e_l to be the result of annotating all instances of future and force in e with 0. For any number n , final state σ such that $\sigma(l) = n$, and final set of vertices V , assuming that all threads t' in V are finished (i.e. $\text{status}(V(t')) = \text{F}$), if

$$[t \mapsto (l := \llbracket e_l \rrbracket^S(\emptyset), \text{R}), \emptyset, [l \mapsto ()] \rightarrow^* V, \emptyset, \sigma$$

then there is a final configuration M such that $[f_0 \mapsto e] \rightarrow^* M$ and $M(f_0) = n$. Furthermore, divergence is preserved:

$$[t \mapsto (l := \llbracket e_l \rrbracket^S(\emptyset), R)], \emptyset, [l \mapsto ()] \rightarrow^\infty \Rightarrow [f_0 \mapsto e] \rightarrow^\infty .$$

Proof. This theorem follows from the composition of the two simulation diagrams, given by Lemmas 19 and 22. Clearly, $\llbracket e_l \rrbracket = e$. Moreover, since $\text{surf}(e_l)$, we also have $\llbracket [f_0 \mapsto (l, t_0, t, e_l, 0)] \rrbracket(\emptyset) = [t \mapsto (l := \llbracket e_l \rrbracket^S(\emptyset), R)], \emptyset, [l \mapsto ()]$, which gives us $[f_0 \mapsto (l, t_0, t, e_l, 0)] \propto [t \mapsto (l := \llbracket e_l \rrbracket^S(\emptyset), R)], \emptyset, [l \mapsto ()]$.

For the termination case, by Lemma 22 we obtain that there exists a configuration M such that $[f_0 \mapsto (l, t_0, t, e_l, 0)] \rightarrow_l^* M$ and $M \propto V', \emptyset, \sigma'$. Inspecting the translation, we find that this requires that the annotations of all the futures in M are equal to 1, and so M is a final configuration. Since this means that $M(f_0) = (l, t_0, t', v, 1)$ for some t' and v , we learn that $\sigma'(l) = \llbracket v \rrbracket^S(M)$, and so, since $\sigma'(l) = n$, that $v = n$. By Lemma 19 we can now conclude that $[f_0 \mapsto e] \rightarrow_O^* \llbracket M \rrbracket$, which ends the proof, since $\llbracket M \rrbracket(f_0) = n$.

For the nontermination case, from Lemma 22 conclude that $[f_0 \mapsto (l, t_0, t, e_l, 0)] \rightarrow_l^\infty$, which, by Lemma 4 implies that $[f_0 \mapsto e] \rightarrow_O^\infty$. This ends the proof. \square

```

1 processor_local queue<vertex*> workQueue // bag of vertices
2 processor_local vertex* current // running vertex
3 processor_local cont* proc_cont // continuation of the scheduler
4
5 void schedulerLoop() // executed by each processor
6   while true // termination details omitted
7     if workQueue.empty()
8       // implementation-dependent load balancing
9       acquireWork() // blocking call
10    current = workQueue.pop()
11    current->releaseHandle = increment(current->in)
12    if(current->cont == null) // initialize the continuation
13      current->cont = new_cont(&enter)
14    swap_cont(proc_cont, current->cont) // execute the vertex
15    if(current->cont == null) // vertex has finished
16      parallelNotify(current->out)
17    else // vertex has yielded
18      decrement(current->releaseHandle)
19
20 void enter() // execute the current vertex, assuming it has never yielded
21   current->run()
22   current->cont = null // mark vertex finished
23   jump_cont(proc_cont)
24
25 // 'newTd e' is short for 'createThread(fun () => e)'
26 vertex* createThread(runMethod)
27   vertex* v = new vertex
28   v->run = runMethod
29   v->in = new_incounter(v)
30   v->out = new_outset()
31   v->releaseHandle = increment(v->in)
32   return v
33
34 void release(vertex* v)
35   decrement(v->releaseHandle)
36
37 void newEdge(vertex* v1, vertex* v2)
38   incounterHandle* h = increment(v2->in)
39   bool success = add(v1->out, h)
40   if not success // vertex v1 has already completed
41     decrement(h) // roll back on edge creation
42
43 void yield()
44   swap_cont(current->cont, proc_cont)
45
46 vertex* self()
47   return current

```

Figure 3: Realization of the scheduler loop and primitive operations. Details of load balancing and termination detection are omitted.

E Correctness of the implementation

Recall the code of the primitive dag operations from the paper, presented in Figure 3.

We continue with the definitions and theorem statement from Section 6 of the paper.

First, we provide the formal statements of the error rules for the semantics.

$$\frac{V(t) = (K[\text{release } t'], X) \quad \text{status}(V(t')) \neq \text{N}}{V, E, \sigma \rightarrow \perp} \text{RELEASE-ERR}$$

$$\frac{V(t) = (K[\text{newEdge } (t_1, t_2)], X) \\ t_1 \notin \text{dom}(V) \vee t_2 \notin \text{dom}(V) \vee \text{status}(V(t_2)) = \text{F} \vee \\ (\text{status}(V(t_2)) = \text{X} \wedge t \neq t_2) \vee (\{(t_1, t_2)\} \cup E) \text{ is cyclic}}{V, E, \sigma \rightarrow \perp} \text{NEWEDGE-ERR}$$

To prove the theorem, we proceed in two steps. In the first step, we present a slightly modified set of rules for the dag calculus, proved equivalent to the original one. This allows for a tighter fit with the implementation, in particular accounting for the fact that outgoing edges may be removed incrementally and not atomically. The second key step is the statement of a global invariant, that binds the dag calculus configurations with corresponding states of the machine. We prove that this invariant is preserved by each evaluation step of the machine, some of those steps corresponding to transitions from the (alternative presentation of the) dag calculus.

In the following paragraphs, we first discuss the necessary refinement of the dag calculus, and then describe the invariants used in the second part of the proof and the points of the program at which dag reductions happen. We elide some of the technical details involved in relating the computations, particularly when it comes to the compilation relation, and refer the interested reader to the appendix.

Refined semantics To refine the semantics, we introduce a reduction relation \rightarrow' . Most of the rules carry over directly from the original semantics; we only refine the `STOP` rule, since the implementation removes edges, one-by-one, potentially in parallel—whereas in the original semantics all the edges are removed in one step. Thus, to obtain a less atomic semantics, we replace this rule with two rules: `STOP'` and `REMEDGE'`. The first of these changes the thread's status to `F` (finished), but does not remove any outgoing edges, while the second one removes a single edge from a vertex with status `F`.

$$\frac{V(t) = (v, X)}{V, E, \sigma \rightarrow' V[t \mapsto (C, F)], E, \sigma} \text{STOP} \quad \frac{V(t) = (C, F)}{V, (E \setminus (t, t')), \sigma \rightarrow' V, E, \sigma} \text{REMEDGE'}$$

With the refined semantics set up, we can prove that it is equivalent to the original one. This means we gain a more fine-grained and less atomic way of talking about programs that is still equivalent to the original formulation.

Lemma 23. *For any dag calculus state V, E, σ , and a terminal state V', \emptyset, σ' such that $\text{status}(V'(t)) = \text{F}$ for any $t \in \text{dom}(V')$, we have*

$$V, E, \sigma \rightarrow^* V', \emptyset, \sigma' \iff V, E, \sigma \rightarrow'^* V', \emptyset, \sigma'.$$

Proof. The left-to-right direction of this lemma holds trivially, since we can encode `STOP` in the refined calculus by using `STOP'` and `REMEDGE'`. For the right-to-left direction, we have a derivation with edge removals scattered after it's marked `F`. The informal argument for why we can replace these with a bulk removal, is that it is always safe to remove an edge earlier, rather than later, as long as the removal happens after the vertex was marked as finished—which exactly matches the semantics of `STOP`. \square

E.1 The invariant

In order to prove the statement of the correctness theorem, we must now relate the states of the dag calculus (with refined semantics), to the machine states. This is done through an invariant of the following shape, where each of the conjuncts represents a specific sub-invariant defined further on, and where H_I and H_O are two finite maps that bind vertex identifiers to sets of incounter handles. These two maps determine the handles stored in the incouters and outsets of each of the vertices.

$$\begin{aligned} \mathcal{I}((V, E, \sigma), (M, S)) &\equiv \text{wf}(M) \wedge \exists H_I, H_O. \\ &\mathcal{I}_V(M, V, S, H_I, H_O) \wedge \mathcal{I}_Q(V, E, S) \wedge \mathcal{I}_S(M, \sigma) \wedge \\ &\mathcal{I}_E(M, V, E, S, H_I, H_O) \wedge \mathcal{I}_R(M, V, S) \end{aligned}$$

Before we turn to discuss specific parts of the invariant, we briefly describe how formalization of the machine state, and of the continuation, incounter and outset data structures.

Description of the machine state This consists of two parts, the memory, M , which maps addresses to values (C-like records, integers, pointers, etc.), and the per-thread state S . The invariant ensures that the memory is well-formed, i.e., we can split it into five disjoint sections, denoted M_V, M_I, M_O, M_K and M_S . These sections represent the parts of storage used to represent, respectively, vertices, incouters, outsets, continuations and the mutable state of the dag calculus. The state of a processor p is modeled as the tuple $(L_p, Q_p, C_p, K_p, \Gamma_p)$, where L_p denotes the line of program currently evaluated by the processor, Q_p denotes its work-queue, C_p denotes the current vertex, K_p denotes the stored continuation `proc_cont`, and Γ_p denotes the environment that maps variables to values. The evaluation proceeds by progressing the L_p counter according to the control-flow, updating the environment Γ_p and other variables as necessary. By convention, we thread the call-stack in the environment, by associating the return address and environment with a `ret` variable in the environment of the callee.

Specification of auxiliary data structures We need to assume specifications of the data structures that the scheduler uses: work queues, continuations, incouters, and outsets. Work queues are simply modeled as sets of vertices. For continuations, we assume a representation predicate $\text{Cont}(K, L, \Gamma)$, which relates a representation of the continuation in memory, K to the pair (L, Γ) of a line number in the program and an environment. The incounter is represented by the predicate $\text{InCounter}(I, H, v)$, which states that the incountered represented in the memory by I has a set of handles H attached to itself and is itself attached to node v (which it will wake when the count is decremented to 0). Finally, the outset is represented by the predicate $\text{OutSet}(O, H)$, which matches the memory representation O with the set of the added elements, H . The functions that operate on these objects use the predicates to specify their effect; we provide these specifications in the technical appendix.

Definition of the sub-invariant I_V The predicate I_V ensures that the dag vertices are properly represented as records with appropriate values, according to the status of the vertex. We call r, i, o, h, k the contents of the fields of a vertex record.

$$\begin{aligned}
I_V(M, V, S, H_I, H_O) \equiv & \\
\forall t, e, s. V(t) = (e, s). \exists r, i, o, h, k. M_V(t) = \{r, i, o, h, k\} \wedge & \\
\text{InHandles}(M_I(i), H_I(t), t) \wedge \text{OutHandles}(M_O(o), H_O(t)) \wedge & \\
s = \text{N} \Rightarrow (\text{Compile}(e, r) \wedge k = \text{null} \wedge h \in H_I(t)) \wedge & \\
s = \text{R} \Rightarrow (\text{if } k = \text{null} \text{ then } \text{Compile}(e, r) & \\
\text{else } \text{EvalCont}(e, M_K(k)) \wedge & \\
(h \in H_I(t) \iff \exists p. L_p \in \{11 \dots 15\} \cup \{18\} \wedge C_p = t) \wedge & \\
s = \text{X} \Rightarrow (\exists p. C_p = t \wedge \text{EvalEnter}(e, L_p, \Gamma_p) \wedge & \\
h \in H_I(t) \wedge (L_p = 21 \Rightarrow k = \text{null})) \wedge & \\
s = \text{F} \Rightarrow k = \text{null} \wedge (\exists p. C_p = t \wedge L_p \in \{15, 16\} \vee H_O(t) = \emptyset) &
\end{aligned}$$

Above, we ensure the correct representation of the incounters and outlets using the `InHandles` and `OutHandles` predicates. The invariant then describes the state of the memory based on what is the status of a given vertex. If the status is `N` (new), the run field, represented by r , should point to a line number where the compiled version of e is stored. The same should be true of a `R` (released) vertex, if its continuation is not yet set up. If the continuation is set up, on the other hand, it is the continuation that represents the computation e . This is achieved through the `EvalCont` predicate, which states that the continuation stored at k corresponds to some program state, and that state is in process of evaluating the `enter` function with computation e . Formally, it is defined as shown below, where `Eval`(e, L, Γ) is a predicate that relates a partially executed dag computation e with a program state (L, Γ) . Note that this predicate is related to `Compile`. (We keep both predicates abstract and reason about them axiomatically.)

$$\begin{aligned}
\text{EvalCont}(e, K) \equiv \exists L, \Gamma. \text{EvalEnter}(e, L, \Gamma) \wedge \text{Cont}(K, L, \Gamma) & \\
\text{EvalEnter}(e, L_p, \Gamma_p) \equiv & \\
(L_p = 21 \wedge \exists L. \text{Compile}(e, L)) \vee & \\
(L_p \in \{22, 23\} \wedge e = v) \vee & \\
(\text{Eval}(e, L_p, \Gamma_p) \wedge \Gamma_p(\text{ret}) = (22, [v \mapsto C_p])) &
\end{aligned}$$

We also use `EvalEnter` if the status of the vertex is `X` (executing); the only difference is that we relate the program state of a processor p , rather than a stored continuation—which matches the intuition about vertices under execution. Finally, if the status of the vertex is `F` (finished), it can only have outgoing edges if the outset is about to be processed.

Definition of the sub-invariant I_Q We now turn to the next invariant, I_Q , which together with I_V contains most of the important properties. I_Q describes the state of the work queues, namely it states that the work queues and the vertices just popped

from them always contain exactly those released nodes that have no incoming edges, discounting the vertices that have just yielded and still need to have the artificial in-edge removed.

$$\begin{aligned}
\mathcal{I}_Q(V, E, S) &\equiv \\
&\bigoplus_{p \in \{1 \dots P\}} (Q_p \uplus \text{if } L_p \in \{11 \dots 14\} \text{ then } \{C_p\} \text{ else } \emptyset) = \\
&\{t \in \text{dom}(V) \mid \text{status}(V(t)) = \text{R} \wedge \neg \text{JustYielded}(t, S) \wedge \\
&\quad \forall t'. (t', t) \notin E\} \\
\text{JustYielded}(t, S) &\equiv \exists p \in \{1 \dots P\}. C_p = t \wedge L_p \in \{15, 18\}
\end{aligned}$$

The next part of the invariant, \mathcal{I}_S , states that the values found in the part of memory that models the store, M_S , match the values stored in σ . We use the same relation of compilation for values as in the theorem statement, `CompileVal`, which ensures that part of the conclusion holds. As with the other compilation relations, we leave this abstract.

$$\mathcal{I}_S(\sigma, M) \equiv \text{dom}(\sigma) = \text{dom}(M) \wedge \forall l \in \text{dom}(\sigma). \text{CompileVal}(\sigma(l), M_S(l))$$

Definition of the sub-invariant \mathcal{I}_E The second-to-last part of the invariant, \mathcal{I}_E , enforces the properties of edges: each edge, as well as each artificial in-edge is uniquely represented by an incouter handle, and the appropriate handles match the representation of incouters and outsets, H_I and H_O . Since this mostly focuses on necessary separation properties, it is less important to the understanding of the relationship. From the technical standpoint, we define this invariant by quantifying over an extension of E that also contains the matching handle for each of the edges and artificial in-edges, and stating the properties with respect to that relation.

$$\begin{aligned}
\mathcal{I}_E(M, V, E, S, H_I, H_O) &\equiv \exists \bar{E} \subseteq \text{Tid}^2 \times \text{Handle}. \\
&(\forall (t_1, t_2, h) \in \bar{E}. (t'_1, t'_2, h') \in \bar{E}. \\
&\quad \text{if } h = h' \text{ then } (t_1 = t'_1 \wedge t_2 = t'_2) \text{ else } (t_1 \neq t'_1 \vee t_2 \neq t'_2)) \wedge \\
&(\forall t, t'. t \neq t' \Rightarrow (t, t') \in E \iff \exists h. (t, t', h) \in \bar{E}) \wedge \\
&(\forall (t_1, t_2, h) \in \bar{E}. t_1 \neq t_2 \Rightarrow h \in H_I(t_2) \wedge h \in H_O(t_1)) \wedge \\
&(\forall (t, t, h) \in \bar{E}. h \in H_I(t) \wedge M_V(t).h = h) \wedge \\
&(\forall h, t. h \in H_O(t) \Rightarrow \exists t' \neq t. (t, t', h) \in \bar{E}) \wedge \\
&(\forall h, t. h \in H_I(t) \Rightarrow (\exists t'. (t', t, h) \in \bar{E} \vee \\
&\quad \exists p. L_p \in \{40 \dots 42\} \wedge h = \Gamma_p(h) \wedge \Gamma_p(v2) = t))
\end{aligned}$$

Definition of the sub-invariant \mathcal{I}_R The final component of the invariant, \mathcal{I}_R , describes additional properties of the per-processor state, i.e., the stored continuation K_p , and the information that lets us distinguish between threads that return to the scheduler because they yielded from the ones whose termination finished. The invariant on K_p states that any processor that evaluates the `enter` function has a stored continuation that represents the `schedulerLoop` function suspended at line 15, while the final lines

state that the status of the thread can be checked by checking if its continuation is set to `null`.

$$\begin{aligned} \mathcal{I}_R(M, V, S) &\equiv \forall p \in \{1 \dots P\}. \\ &(\text{enter} \in \text{CallStack}(p) \Rightarrow \text{Cont}(M_K(K_p), 15, \emptyset)) \wedge \\ &L_p = 15 \Rightarrow (\text{if } M_V(C_p).k = \text{null} \text{ then } \text{status}(V(C_p)) = \text{F} \\ &\quad \text{else } \text{status}(V(C_p)) = \text{R}) \end{aligned}$$

E.2 Structure of the proof

The proof proceeds by three lemmas that can be combined using induction over the evaluation of the machine state. The first, Lemma 24 relates the initial states of the dag calculus and the machine through an invariant. The second, Lemma 25 provides properties of the final state. Finally, Lemma 26 ensures that the invariant is preserved through the evaluation of the machine.

Lemma 24. *Assume the initial state of the theorem:*

- e_0 to be a dag calculus expression,
- t_0 to be a thread identifier,
- l_0 a location for the final result,
- r_0 the code pointer to the compiled code for e_0 , i.e. in the sense that $\text{Compile}(e_0, r_0)$ holds (we describe Compile later),
- $V_0 = [t \mapsto (e_0, \text{R})]$, which describes the initial vertex with body e_0 and released status,
- $E_0 = \emptyset$, which describes the initial set of edges,
- $\sigma_0 = [l_0 \mapsto ()]$, which describes the initial heap,
- $M_0 = [l_0 \mapsto (), t_0 \mapsto \text{InitVertex}(r_0)]$, which describes the initial memory state, with a memory cell at location l_0 , and a representation of the initial vertex with run method r_0 , fresh incounter and outset, and null continuation and releaseHandle (as described by the auxiliary InitVertex operator),
- S_0 , which describes the initial state of the processors, by asserting that they are entering `schedulerLoop` (i.e., $\forall p \in \{1 \dots P\}. L_p = 6$, where L_p denotes the line of processor p), and asserting that all work queues are empty except one that contains exactly t_0 (i.e., $\exists p \in \{1 \dots P\}. Q_p = \{t_0\} \wedge (\forall p' \in \{1 \dots P\}. p \neq p' \Rightarrow Q_{p'} = \emptyset)$, where Q_p denotes the work queue of processor p).

Then, $\mathcal{I}((V_0, E_0, \sigma_0), (M_0, S_0))$ holds.

Proof. By definition of the invariant: the configurations match each other directly. \square

Lemma 25. *For any M, S, V, E, σ such that $\mathcal{I}((V, E, \sigma), (M, S))$, if S is a terminal state (i.e., $L_p = 9 \wedge Q_p = \emptyset$ for any processor p), then $I_S(\sigma, M_S)$ and (V, E, σ) is a terminal configuration (i.e., $\text{status}(V(t)) = \text{F}$ for any $t \in \text{dom}(V)$ and $E = \emptyset$).*

Proof. By definition of the invariant. □

Lemma 26. *For any two machine states (M, S) and (M', S') such that $(M, S) \rightarrow (M', S')$ and for any configuration (V, E, σ) such that $\mathcal{I}((V, E, \sigma), (M, S))$, if $\neg ((V, E, \sigma) \rightarrow^* \perp)$ then there exists a dag configuration (V', E', σ') such that $(V, E, \sigma) \rightarrow^* (V', E', \sigma')$ and $\mathcal{I}((V', E', \sigma'), (M', S'))$.*

Proof. We need to check that every possible reduction in our program preserves the invariant. Reductions within the body of one of the run methods, including ones where STEP should apply, are governed by the Eval predicate. However, this still leaves us with checking all the lines of our code. Among these, most evaluation steps do not require a change of the dag calculus state and the concomitant reduction in the dag calculus; the interesting steps are the ones in which we do need to perform a reduction. We consider these below, and explain in which line of code we perform the matching reduction in the dag calculus.

The high-level justifications of why each of the lines of our code preserves the invariants follow below:

- $L_p = 7$: ensures we only pop from non-empty queues
- $L_p = 9$: by specification of `acquireWork`, ensures queue non-empty
- $L_p = 10$: by specification of `workQueue.pop()`; preserves \mathcal{I}_Q
- $L_p = 11$: by specification of `increment`; note the handle obtained is fresh, so invariant \mathcal{I}_E is preserved
- $L_p = 12$: ensures we only restore non-null continuations; invariant \mathcal{I}_V is preserved
- $L_p = 13$: by specification of `new_cont` and `enter`; invariant \mathcal{I}_V is preserved
- $L_p = 14$: by specification of `swap_cont`; apply the START rule, changing status from R to X; preserves the \mathcal{I}_V invariant
- $L_p = 15$: status has changed from X; checks whether returning vertex has status R or F; exploits \mathcal{I}_R to deduce information
- $L_p = 16$: by specification of `parallelNotify`; ensures that the outset of the current vertex is empty; apply REMEDGE' rule for each outgoing edge of the current vertex, noting that these might execute concurrently to other steps.
- $L_p = 18$: by specification of `decrement`; reestablishes the normal state of a released node (i.e., without an artificial in-edge); note status is already R.
- $L_p = 21$: by the relationship of Compile and Eval predicates: $\text{Compile}(e, r) \Rightarrow \text{Eval}(e, r, \emptyset)$
- $L_p = 22$: establishes that the continuation is null in preparation for switching status to F
- $L_p = 23$: by specification of `jump_cont`; the R-F distinction enforced by the invariant established in the previous line; apply the STOP' rule, changing status

from X to F.

- $L_p = 27$: creates a vertex in M_V ; since the vertex is not yet in V , it does not violate the invariants
- $L_p = 28$: set up the run field; from this point the $\text{Compile}(e, r)$ holds
- $L_p = 29$: set up the incounter; from this point on the InHandles predicate holds for the vertex
- $L_p = 30$: set up the outset; from this point on the OutHandles predicate holds
- $L_p = 31$: set up the release handle by specification of `increment`; note the disjointness preserved since the handle is fresh
- $L_p = 32$: by a property of `Eval`; apply the `NEWTD` rule; invariant \mathcal{I}_V holds for this vertex with status N

- $L_p = 35$: performs a `RELEASE` reduction; the released vertex satisfies the invariant, since its `cont` field still contains null; preserves invariant \mathcal{I}_V , status changes from N to R

- $L_p = 38$: by specification of `increment`; the handle `h` is fresh; at this point enter the special branch of \mathcal{I}_E
- $L_p = 39$: by specification of `add`; performs a `NEWEDGE` transition; if addition successful, the edge was added, if not — the vertex has state F, so transition is allowed
- $L_p = 40$: ensures we do not decrement if addition was successful
- $L_p = 41$: by specification of `decrement`; leaves the special branch of \mathcal{I}_E

- $L_p = 44$: by specification of `swap_cont`; the R-F distinction enforced by `swap_cont` itself; applies the `YIELD` rule, status changes from X to R; preserves \mathcal{I}_V

- $L_p = 47$: by a property of `Eval` and \mathcal{I}_V ; performs a `SELF` reduction

□

Theorem 4 (Correctness). *As stated in the paper.*

Proof. First, we deduce from Lemma 24 that the invariant holds. Then, proceed by induction on the execution of the machine. If the derivation is empty, then the configuration of the machine is terminal, and the conclusion follows easily from Lemma 25.

In the inductive case, we can use Lemma 26 on the first step of the derivation. Thus, we establish that the invariant is preserved. Since the correctness of the program is also preserved by the evaluation, we can use the induction hypothesis to end the proof. □