

Appendix of DimSum: A Decentralized Approach to Multi-language Semantics and Verification

MICHAEL SAMMLER, MPI-SWS, Germany

SIMON SPIES, MPI-SWS, Germany

YOUNGJU SONG, MPI-SWS, Germany

EMANUELE D'OSUALDO, MPI-SWS, Germany

ROBBERT KREBBERS, Radboud University Nijmegen, The Netherlands

DEEPAK GARG, MPI-SWS, Germany

DEREK DREYER, MPI-SWS, Germany

A COMBINATORS

Filter. The filter combinator $M \mid_{\sigma} M' \triangleq (S_{\text{filter}} \times S_M \times S_{M'}, \rightarrow_{\text{filter}}, (\sigma, \sigma_M^0, \sigma_{M'}^0))$ takes in a module $M \in \text{Module}(E_1)$ and a filter $M' \in \text{Module}(\text{FilterEvents}(E_1, E_2))$ and then produces a module with events drawn from E_2 . The states of the filter combinator are given by $S_{\text{filter}} \triangleq \{P, F\} \cup \{P(e) \mid e \in E_1\} \cup \{F(e) \mid e \in E_1\}$ and the transitions are depicted in Fig. 1. The events of the filter module are drawn from the set

$$\text{FilterEvents}(E_1, E_2) \triangleq \{\text{Receive}(e_1) \mid e_1 \in E_1\} \cup \{\text{Emit}(e_2) \mid e_2 \in E_2\} \cup \{\text{Return}(e_1) \mid e_1 \in \text{option}(E_1)\}$$

where $\text{Receive}(e_1)$ is for accepting an incoming event, $\text{Emit}(e_2)$ is for emitting an outgoing event, and $\text{Return}(e_1)$ is for returning control to the inner module. In the last case, the event e_1 is optional to allow us to force the inner module to emit the event e_1 next (*i.e.*, all visible transitions of the inner module except ones emitting event e_1 are blocked).

Linking. The linking operator $M_1 \oplus_X M_2$ is defined on modules $M_1, M_2 \in \text{Module}(E_{?!})$ where $E_{?!}$ is (an event type that is isomorphic to) $E \times \{?, !\}$. The parameter $X = (S, \rightsquigarrow, s^0)$ determines how the events are linked. It consists of a set of linking-internal states S , an initial state $s^0 \in S$, and a relation $\rightsquigarrow \subseteq (D \times S \times E) \times ((D \times S \times E) \cup \{\zeta\})$ describing how events should be translated. Formally, linking can be defined as $M_1 \oplus_X M_2 \triangleq M_1 \times M_2 \mid_{\text{link}_X}$.¹ The module link_X is defined as $\text{link}_X \triangleq (S_{\text{link}} \times S_X, \rightarrow_{\text{link}}, (\text{Wait}, s_X^0))$ where $S_{\text{link}} \triangleq (\{\text{Wait}, \text{Ub}\} \cup \{\text{Emit}(e, \sigma) \mid e \in E_{?!}, \sigma \in S_{\text{link}}\} \cup \{\text{Return}(e) \mid e \in \text{option}(E_{?!})\})$ and $\rightarrow_{\text{link}}$ is defined in Fig. 2.

(Kripke) wrappers. The combinator $\lceil M \rceil_X$ translates a modules with events E_1 to a module with events E_2 . This combinator is parametrized by $X = (S, R, \leftarrow, \rightarrow, s^0, F^0)$ where S is a set of states and s^0 is an initial state. These states were omitted in the main paper for simplicity. They don't give additional expressive power but make writing the wrapper $\lceil \cdot \rceil_{r \Rightarrow a}$ more pleasant. \leftarrow and \rightarrow are relations that describe how the wrapper transforms the incoming and outgoing events. Concretely, \leftarrow describes how to translate an event $e_2 \in E_2$ to an event $e_1 \in E_1$ and \rightarrow describes the translation from $e'_1 \in E_1$ to $e'_2 \in E_2$.

¹The Coq development defines linking via more low-level combinators that we omit from the presentation here. Also the Coq development allows undefined behavior via a Boolean on the right side of \rightsquigarrow instead of a separate ζ result.

Authors' addresses: Michael Sammler, MPI-SWS, Saarland Informatics Campus, Germany, msammler@mpi-sws.org; Simon Spies, MPI-SWS, Saarland Informatics Campus, Germany, spies@mpi-sws.org; Youngju Song, MPI-SWS, Saarland Informatics Campus, Germany, youngju@mpi-sws.org; Emanuele D'Ossualdo, MPI-SWS, Saarland Informatics Campus, Germany, dosualdo@mpi-sws.org; Robbert Krebbers, Radboud University Nijmegen, The Netherlands, mail@robbertkrebbers.nl; Deepak Garg, MPI-SWS, Saarland Informatics Campus, Germany, dg@mpi-sws.org; Derek Dreyer, MPI-SWS, Saarland Informatics Campus, Germany, dreyer@mpi-sws.org.

$$\begin{array}{c}
\text{FILTER-STEP-PROG-NONE} \\
\frac{\sigma = P \vee \sigma = P(e) \quad \sigma_1 \xrightarrow{\tau} \Sigma}{(\sigma, \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(\sigma, \sigma'_1, \sigma_2) \mid \sigma'_1 \in \Sigma\}} \\
\\
\text{FILTER-STEP-PROG-RECV} \\
\frac{\sigma_1 \xrightarrow{e} \Sigma}{(P(e), \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(P, \sigma'_1, \sigma_2) \mid \sigma'_1 \in \Sigma\}} \\
\\
\text{FILTER-STEP-FILTER-RECV} \\
\frac{\sigma_2 \xrightarrow{\text{Receive}(e)} \Sigma}{(F(e), \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(F, \sigma_1, \sigma'_2) \mid \sigma'_2 \in \Sigma\}} \\
\\
\text{FILTER-STEP-FILTER-RETURN} \\
\frac{\sigma_2 \xrightarrow{\text{Return}(e)} \Sigma}{(F, \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(\text{if } e = \text{Some}(e') \text{ then } P(e') \text{ else } P, \sigma_1, \sigma'_2) \mid \sigma'_2 \in \Sigma\}} \\
\\
\text{FILTER-STEP-FILTER-NONE} \\
\frac{\sigma = F \vee \sigma = F(e) \quad \sigma_2 \xrightarrow{\tau} \Sigma}{(\sigma, \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(\sigma, \sigma_1, \sigma'_2) \mid \sigma'_2 \in \Sigma\}} \\
\\
\text{FILTER-STEP-PROG} \\
\frac{\sigma_1 \xrightarrow{e} \Sigma}{(P, \sigma_1, \sigma_2) \xrightarrow{\tau}_{\text{filter}} \{(P(e), \sigma'_1, \sigma_2) \mid \sigma'_1 \in \Sigma\}} \\
\\
\text{FILTER-STEP-FILTER-EMIT} \\
\frac{\sigma_2 \xrightarrow{\text{Emit}(e)} \Sigma}{(F, \sigma_1, \sigma_2) \xrightarrow{e}_{\text{filter}} \{(F, \sigma_1, \sigma'_2) \mid \sigma'_2 \in \Sigma\}}
\end{array}$$

Fig. 1. Definition of $\rightarrow_{\text{filter}}$.

$$\text{to}(d, e) = \begin{cases} \text{Return}(\text{left}(e?, L)) & \text{if } d = L \\ \text{Return}(\text{right}(e?, R)) & \text{else if } d = R \\ \text{Emit}(e!, \text{Return}(\text{None})) & \text{else if } d = E \end{cases}$$

$$\begin{array}{c}
\text{LINK-STEP-WAIT-L} \\
\frac{(L, s, e) \rightsquigarrow (d, s', e')}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{left}(e!, d))}_{\text{link}} \{(\text{to}(d, e'), s')\}} \\
\\
\text{LINK-STEP-WAIT-R} \\
\frac{(R, s, e) \rightsquigarrow (d, s', e')}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{right}(e!, d))}_{\text{link}} \{(\text{to}(d, e'), s')\}} \\
\\
\text{LINK-STEP-WAIT-N} \\
\frac{(E, s, e) \rightsquigarrow (d, s', e')}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{env}(d))}_{\text{link}} \{(\text{Emit}(e'?, \text{to}(d, e')), s')\}} \\
\\
\text{LINK-STEP-EMIT} \\
(\text{Emit}(e, \sigma), s) \xrightarrow{\text{Emit}(e)}_{\text{link}} \{(\sigma, s)\} \\
\\
\text{LINK-STEP-RETURN} \\
(\text{Return}(e), s) \xrightarrow{\text{Return}(e)}_{\text{link}} \{(\text{Wait}, s)\} \\
\\
\text{LINK-STEP-UB} \\
(\text{Ub}, s) \xrightarrow{\tau}_{\text{link}} \emptyset \\
\\
\text{LINK-STEP-WAIT-L-UB} \\
\frac{(L, s, e) \rightsquigarrow \downarrow}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{left}(e!, d))}_{\text{link}} \{(\text{Ub}, s)\}} \\
\\
\text{LINK-STEP-WAIT-R-UB} \\
\frac{(R, s, e) \rightsquigarrow \downarrow}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{right}(e!, d))}_{\text{link}} \{(\text{Ub}, s)\}} \\
\\
\text{LINK-STEP-WAIT-N-UB} \\
\frac{(E, s, e) \rightsquigarrow \downarrow}{(\text{Wait}, s) \xrightarrow{\text{Receive}(\text{env}(d))}_{\text{link}} \{(\text{Ub}, s)\}}
\end{array}$$

Fig. 2. Definition of $\rightarrow_{\text{link}}$.

As mentioned in the paper, these relations are separation logic relations. As such, they are of type $E_1 \times S \times E_2 \times S \rightarrow \text{UPred}(R)$. The second component of X is a resource algebra R [Jung et al. 2018] that determines a separation logic of uniform predicates over the resource algebra $\text{UPred}(R)$. F^0 denotes the initial set of resources owned by the wrapper. We define $[M]_X \triangleq M \Vdash \llbracket \text{state}(s^0, F^0) \rrbracket_s$

$$\text{Instr} \ni \mathbf{c} \triangleq \text{syscall}; \mathbf{c} \mid \text{upd}(\mathbf{x}, \mathbf{r}, \mathbf{v}); \mathbf{c} \mid \text{ldr}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{v}, \mathbf{v}'); \mathbf{c} \mid \text{str}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{v}, \mathbf{v}'); \mathbf{c} \mid \text{jump}$$
Fig. 3. Micro-Instructions of **Asm**

where the filter module is given by the following Spec program:²

$$\begin{aligned} & \text{state}(s_2, F_2) \triangleq_{\text{coind}} \\ & \quad \exists e_2; \text{vis}(\text{Emit}(e_2)); \forall e_1, s_1, F_1; \text{assume}(\text{sat}(F_1 * F_2 * (e_1, s_1) \leftarrow (e_2, s_2))); \text{vis}(\text{Return}(e_1)); \\ & \quad \exists e'_1; \text{vis}(\text{Receive}(e'_1)); \exists e'_2, s'_2, F'_2; \text{assert}(\text{sat}(F_1 * F'_2 * (e'_1, s_1) \rightarrow (e'_2, s'_2))); \text{vis}(\text{Emit}(e'_2)); \\ & \quad \text{state}(s'_2, F'_2) \end{aligned}$$

Intuitively, $\text{state}(s_2, F_2)$ works as follows: Given an initial state s_2 and a proposition describing resource ownership of the translation F_2 , state synchronizes with the environment on an event e_2 . Then it angelically chooses an event e_1 for the inner module, a new state s_1 , ownership of the environment F_1 , and a proof that the ownership of the translation together with the ownership of the environment and the precondition $(e_1, s_1) \leftarrow (e_2, s_2)$ is satisfiable. Then state sends e_1 to the inner module M . Next, it receives an event e'_1 from M and (demonically) chooses an event e'_2 to emit to the environment, a new state s'_2 , new ownership of the translation F'_2 , and a proof that the ownership of the translation together with the ownership of the environment and the postcondition $(e'_1, s_1) \rightarrow (e'_2, s'_2)$ is satisfiable. After emitting e'_2 , the process repeats with state s'_2 and F'_2 .

B MICRO-INSTRUCTIONS OF **Asm**

Inspired by Sammler et al. [2022a], instructions \mathbf{c} in **Asm** are sequences of *micro instructions* (i.e., simple instructions that, when composed together, form an actual instruction), depicted in Fig. 3. The instruction $\text{syscall}; \mathbf{c}$ does a syscall and then executes \mathbf{c} . The instruction $\text{upd}(\mathbf{x}, \mathbf{r}, \mathbf{v}); \mathbf{c}$ updates the register \mathbf{x} according to the map $\mathbf{r} \mapsto \mathbf{v}$ applied to the current register values \mathbf{r} and then executes \mathbf{c} . The instruction $\text{ldr}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{v}, \mathbf{v}'); \mathbf{c}$ takes the value stored in \mathbf{x}_2 , applies the transformation $\mathbf{v} \mapsto \mathbf{v}'$ to it to obtain an address, loads from the memory at that address, stores the result in \mathbf{x}_1 , and then executes \mathbf{c} . The instruction $\text{str}(\mathbf{x}_1, \mathbf{x}_2, \mathbf{v}, \mathbf{v}'); \mathbf{c}$ takes the value stored in \mathbf{x}_1 , applies the transformation $\mathbf{v} \mapsto \mathbf{v}'$ to it to obtain an address, stores in the memory at that address the value in \mathbf{x}_1 , and then executes \mathbf{c} . The instruction jump reads the pc register and then jumps to the address stored there.

The reason for the micro instruction representation is that we can represent a large instruction set by chaining few primitives. For example, the instructions used in **print** are derived as follows:

$$\begin{aligned} \text{ret} & \triangleq \text{upd}(\text{pc}, \mathbf{r}, \mathbf{r}(\mathbf{x}30)); \text{jump} & \text{syscall} & \triangleq \text{syscall}; \text{next} & \text{mov } \mathbf{x}, \mathbf{v} & \triangleq \text{upd}(\mathbf{x}, \mathbf{r}, \mathbf{v}); \text{next} \\ \text{sle } \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 & \triangleq \text{upd}(\mathbf{x}_1, \mathbf{r}, \text{if } \mathbf{r}(\mathbf{x}_2) \leq \mathbf{r}(\mathbf{x}_3) \text{ then } 1 \text{ else } 0); \text{next} \end{aligned}$$

where we abbreviate $\text{next} \triangleq \text{upd}(\text{pc}, \mathbf{r}, \mathbf{r}(\text{pc}) + 1); \text{jump}$.

C SEMANTIC LINKING FOR **Asm**

The full definition of the semantic linking relation \rightsquigarrow for **Asm** can be found in Fig. 4. Compared to the excerpt shown in the paper, it contains two additional cases, **ASM-LINK-SYSCALL** and **ASM-LINK-SYSCALL-RETURN**. The rule **ASM-LINK-SYSCALL** makes sure syscalls are passed on to the environment (and never come from the environment). When a syscall is triggered, we store the current turn d in

²The Coq development defines an equivalent module directly using a step relation, but we give the definition here using Spec for readability.

$$\begin{array}{c}
\text{ASM-LINK-JUMP} \\
\frac{(d' = L \wedge \mathbf{r}(\mathbf{pc}) \in \mathbf{d}_1) \vee (d' = R \wedge \mathbf{r}(\mathbf{pc}) \in \mathbf{d}_2) \vee (d' = E \wedge \mathbf{r}(\mathbf{pc}) \notin \mathbf{d}_1 \cup \mathbf{d}_2) \quad d \neq d'}{(d, \text{None}, \mathbf{Jump}(\mathbf{r}, \mathbf{m})) \rightsquigarrow_{\mathbf{d}_1, \mathbf{d}_2} (d', \text{None}, \mathbf{Jump}(\mathbf{r}, \mathbf{m}))} \\
\text{ASM-LINK-SYSCALL} \\
\frac{d \neq E}{(d, \text{None}, \mathbf{Syscall}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{m})) \rightsquigarrow_{\mathbf{d}_1, \mathbf{d}_2} (E, \text{Some}(d), \mathbf{Syscall}(\mathbf{v}_1, \mathbf{v}_2, \mathbf{m}))} \\
\text{ASM-LINK-SYSCALL-RETURN} \\
\frac{d' \neq E}{(E, \text{Some}(d'), \mathbf{SyscallRet}(\mathbf{v}, \mathbf{m})) \rightsquigarrow_{\mathbf{d}_1, \mathbf{d}_2} (d', \text{None}, \mathbf{SyscallRet}(\mathbf{v}, \mathbf{m}))}
\end{array}$$

Fig. 4. Definition of semantic linking relation \rightsquigarrow for *Asm*.

$$\begin{array}{l}
\text{Library} \ni \mathbf{R} \triangleq (\mathbf{fn} \mathbf{f}(\bar{x}) \triangleq \overline{\text{local } \mathbf{y}[n]}; \mathbf{e}), \mathbf{R} \mid \emptyset \\
\text{Expr} \ni \mathbf{e} \triangleq \mathbf{v} \mid \mathbf{x} \mid \mathbf{e}_1 \oplus \mathbf{e}_2 \mid \text{let } \mathbf{x} := \mathbf{e}_1 \text{ in } \mathbf{e}_2 \mid \text{if } \mathbf{e}_1 \text{ then } \mathbf{e}_2 \text{ else } \mathbf{e}_3 \mid \mathbf{e}_1(\overline{\mathbf{e}_2}) \mid !\mathbf{e} \mid \mathbf{e}_1 \leftarrow \mathbf{e}_2 \\
\text{BinOp} \ni \oplus \triangleq + \mid < \mid == \mid \leq \\
\text{Runtime Expr} \ni \mathbf{E} \triangleq \dots \mid \text{alloc_frame } \overline{(x, n)} \mathbf{E} \mid \text{free_frame } \overline{(\ell, n)} \mathbf{E} \mid \text{Ret}(b, \mathbf{E}) \mid \text{Wait}(b)
\end{array}$$

Fig. 5. Grammar of *Rec*.

the private state of the linking operator. This way, we can make sure that when we return from a `syscall` (`ASM-LINK-SYSCALL-RETURN`), the execution continues with the module that triggered the `syscall`.

D *Rec*

The language *Rec* is a simple, high-level language with arithmetic operations, let bindings, memory operations, conditionals, and (potentially recursive) function calls (depicted in Fig. 5). The libraries \mathbf{R} of *Rec* are lists of function declarations. Each function declaration contains the name of the function \mathbf{f} , the argument names \bar{x} , local variables \bar{y} which are allocated in the memory, and a function body \mathbf{e} . The set of function names $|\mathbf{R}|$ of a library \mathbf{R} is defined as the names of the functions in the list \mathbf{R} .

Module semantics. The semantics of a *Rec* library \mathbf{R} is the module $\llbracket \mathbf{R} \rrbracket_r$. The states of the module are of the form $\sigma = (E, m, \mathbf{R})$ where E is the current *runtime expression* (explained below). We write (\rightarrow_r) for the transition system (shown in Fig. 6) and the initial state is $(\text{Wait}(\text{false}), \emptyset, \mathbf{R})$.

To define the transition relation \rightarrow_r , we extend the static expressions \mathbf{e} to runtime expressions \mathbf{E} , which have operations for allocating and deallocating stack frames as well as two distinguished expressions $\text{Ret}(b, \mathbf{E})$ and $\text{Wait}(b)$. These expressions are used to control when the module emits call and return events: Initially, the module is waiting and willing to accept any incoming call to the functions of the library (see `REC-START`). Once it starts, the function call is wrapped in the $\text{Ret}(b, \cdot)$ expression to ensure an event is emitted after the function finishes executing (see `REC-RET-RETURN`). A call to functions of the library (see `REC-CALL-INTERNAL`), will trigger the allocation of the local variables and, subsequently, the execution of the function body. A call to an external function (see `REC-CALL-EXTERNAL`) will emit a $\text{Call}!(\mathbf{f}, \bar{\mathbf{v}}, \mathbf{m})$ and proceed to the waiting state. The flag for the

$$\begin{array}{c}
\text{REC-BINOP} \\
(v_1 \oplus v_2, m, R) \xrightarrow{\tau}_r \{(v, m, R) \mid \text{eval}_{\oplus}(v_1, v_2, v)\} \\
\\
\text{REC-LOAD} \\
(!v_1, m, R) \xrightarrow{\tau}_r \{(v_2, m, R) \mid \exists \ell. v_1 = \ell \wedge m(\ell) = v_2\} \\
\\
\text{REC-STORE} \\
(v_1 \leftarrow v_2, m, R) \xrightarrow{\tau}_r \{(v_2, m[\ell \mapsto v_2], R) \mid \exists \ell. v_1 = \ell \wedge \text{heap_alive}(m, \ell)\} \\
\\
\text{REC-IF} \\
(\text{if } v \text{ then } e_1 \text{ else } e_2, m, R) \xrightarrow{\tau}_r \{(e, m, R) \mid \exists b. v = b \wedge \text{if } b \text{ then } e = e_1 \text{ else } e = e_2\} \\
\\
\begin{array}{cc}
\text{REC-LET} & \text{REC-VAR} \\
(\text{let } x := v \text{ in } e, m, R) \xrightarrow{\tau}_r \{(e[v/x], m, R)\} & (x, m, R) \xrightarrow{\tau}_r \emptyset
\end{array} \\
\\
\text{REC-ALLOC} \\
\frac{\text{heap_alloc_list}(\bar{n}, \bar{\ell}, m_1, m_2)}{(\text{alloc } \overline{(y, n)} e, m_1, R) \xrightarrow{\tau}_r \{(\text{free_frame } \overline{(\ell, n)} (e[\bar{\ell}/\bar{y}]), m_2, R) \mid \forall m \in \bar{n}. m > 0\}} \\
\\
\text{REC-FREE} \\
(\text{free_frame } \overline{(\ell, n)} v, m_1, R) \xrightarrow{\tau}_r \{(v, m_2, R) \mid \text{heap_free_list}(\overline{(\ell, n)}, m_1, m_2)\} \\
\\
\text{REC-START} \\
\frac{f \in R}{(\text{Wait}(b), m, R) \xrightarrow{\text{Call}?(f, \bar{v}, m')} \{(\text{Ret}(b, f(\bar{v})), m', R)\}} \\
\\
\text{REC-CALL-INTERNAL} \\
\frac{(\text{fn } f(\bar{x}) \triangleq \text{local } y[n]; e) \in R}{(f(\bar{v}), m, R) \xrightarrow{\tau}_r \{(\text{alloc } \overline{(y, n)} (e[\bar{v}/\bar{x}]), m, R) \mid |\bar{x}| = |\bar{v}|\}} \\
\\
\text{REC-CALL-EXTERNAL} \\
\frac{f \notin R}{(f(\bar{v}), m, R) \xrightarrow{\text{Call}!(f, \bar{v}, m)} \{(\text{Wait}(\text{true}), m, R)\}} \quad \text{REC-RET-INCOMING} \\
(\text{Wait}(\text{true}), m, R) \xrightarrow{\text{Return}?(v, m')} \{(v, m', R)\} \\
\\
\text{REC-RET-RETURN} \\
(\text{Ret}(b, v), m, R) \xrightarrow{\text{Return}!(v, m)} \{(\text{Wait}(b), m, R)\} \\
\\
\text{REC-EVAL-CTX} \\
\frac{(E, m, R) \xrightarrow{e}_r \Sigma}{(K[E], m, R) \xrightarrow{e}_r \{(K[E'], m', R') \mid (E', m', R') \in \Sigma\}}
\end{array}$$

Fig. 6. Operational semantics of `Rec`.

waiting becomes true, because the module is now willing to accept a return to the function call that was just issued (see `REC-RET-INCOMING`). The language `Rec` is an evaluation-context based language,

$$\begin{array}{c}
\text{REC-LINK-CALL} \\
\frac{(d' = L \wedge f \in d_1) \vee (d' = R \wedge f \in d_2) \vee (d' = E \wedge f \notin d_1 \cup d_2) \quad d \neq d'}{(d, \bar{d}_s, \text{Call}(f, \bar{v}, m)) \rightsquigarrow_{d_1, d_2} (d', d :: \bar{d}_s, \text{Call}(f, \bar{v}, m))} \\
\\
\text{REC-LINK-RET} \\
\frac{d \neq d'}{(d, d' :: \bar{d}_s, \text{Return}(v, m)) \rightsquigarrow_{d_1, d_2} (d', \bar{d}_s, \text{Return}(v, m))}
\end{array}$$

Fig. 7. Definition of semantic linking relation $\rightsquigarrow_{d_1, d_2}$ for **Rec**.

$$\begin{array}{l}
(e_1, s_1) \rightarrow (e_2, s_2) \triangleq \exists r \ m \ \bar{v}. e_2 = \text{Jump}!(r, m) * \text{inv}(r(\text{sp}), m, \text{mem}(e_1)) * \\
\quad (\exists f \ \bar{v} \ m. e_1 = \text{Call}!(f, \bar{v}, m) * f \notin d * r(x30) \in d * a_f = r(\text{pc}) * \\
\quad \quad s_2 = r :: s_1 * \quad * \quad v \leftrightarrow v \\
\quad \quad \quad v, v \in \bar{v}, \text{take}(|\bar{v}|, r(x0 \dots x8)) \\
\quad \vee \exists v \ m \ r'. e_1 = \text{Return}!(v, m) * r' :: s_2 = s_1 * r(\text{pc}) = r'(x30) * \\
\quad \quad r(x19 \dots x29, \text{sp}) = r'(x19 \dots x29, \text{sp}) * v \leftrightarrow r(x0)) \\
\\
(e_1, s_1) \leftarrow (e_2, s_2) \triangleq \exists r \ m \ \bar{v}. e_2 = \text{Jump}?(r, m) * \text{inv}(r(\text{sp}), m, \text{mem}(e_1)) * \\
\quad (\exists f \ \bar{v} \ m. e_1 = \text{Call}?(f, \bar{v}, m) * f \in d * r(x30) \notin d * a_f = r(\text{pc}) * \\
\quad \quad s_1 = r :: s_2 * \quad * \quad v \leftrightarrow v \\
\quad \quad \quad v, v \in \bar{v}, \text{take}(|\bar{v}|, r(x0 \dots x8)) \\
\quad \vee \exists v \ m \ r'. e_1 = \text{Return}?(v, m) * r' :: s_1 = s_2 * r(\text{pc}) = r'(x30) * \\
\quad \quad r(x19 \dots x29, \text{sp}) = r'(x19 \dots x29, \text{sp}) * v \leftrightarrow r(x0))
\end{array}$$

Fig. 8. Definition of (\leftarrow) and (\rightarrow) for $[\cdot]_{r \Rightarrow a}$.

meaning reductions can happen inside of an arbitrary evaluation context (see **REC-EVAL-CTX**). The definition of the evaluation contexts **K** can be found in the Coq development [Sammler et al. 2022b].

Linking. Syntactically, linking of two **Rec** libraries (i.e., $R_1 \cup_r R_2$) denotes merging the function definitions in R_1 and R_2 . In case of overlapping function names, the function declaration of the left library is chosen. (This choice is arbitrary.) If we semantically link two **Rec** modules (i.e., $M_1 \overset{d_1}{\oplus}_r \overset{d_2}{M_2}$), then we have to synchronize based on the function call and return events. To define the linking $M_1 \overset{d_1}{\oplus}_r \overset{d_2}{M_2}$, we use the combinator $M_1 \oplus_X M_2$. In the case of **Rec**, we pick the relation **R** depicted in Fig. 7. The most interesting difference to **Asm** is that linking in **Rec** has to build up and then wind down a call-stack, which is maintained as the internal state of (\rightsquigarrow) .

E $[\cdot]_{r \Rightarrow a}$ WRAPPER

Before we can give the definition of the wrapper $[\cdot]_{r \Rightarrow a}$, we first need to describe its full form: $[M]_{r \Rightarrow a}^{a, d, d, m}$. In particular, the wrapper is parametrized by a mapping a_{-} from **Rec** function names

to **Asm** addresses, by the instruction address of the **Asm** code **d**, by the function names of the **Rec** code **d**, and by a (fragment of) the initial memory **m**, which can be used for global variables.

To define the wrapper, we pick a suitable flavor of separation logic. Instead of directly presenting the technical details of the resource algebra that we choose for $R_{r \Rightarrow a}$, we instead describe the connectives of the resulting separation logic:

- $p \leftrightarrow v$ states that the **Rec** block id **p** is mapped to **Asm** address **v**. We lift this relation to locations by $\ell \leftrightarrow v_2 \triangleq \exists v_1. \ell.\text{blockid} \leftrightarrow v_1 * v_2 = v_1 + \ell.\text{offset}$ and to values (i.e., $v \leftrightarrow v$) by relating **Rec** integers with the same integer in **Asm** and Boolean values with 0 and 1.
- $v_1 \mapsto_a v_2$ asserts ownership of the address **v**₁ in **Asm** memory **m** and asserts that it contains the value **v**₂. The $v_1 \mapsto_a v_2$ connective is useful for asserting private ownership of **Asm** memory in assembly libraries (e.g., it is used internally by the coroutine library to manage its global state).
- $p \mapsto_r V$ where **V** is a map from offsets to values asserts that the block with id **p** contains exactly **V**. The $p \mapsto_r V$ connective is useful for asserting ownership of locations in the **Rec** memory, e.g., for locations that are not mapped to the **Asm** memory.
- $\text{inv}(v, m, m)$ asserts that **m** and **m** are in an invariant such that all the aforementioned assertions (i.e., $p \leftrightarrow v$, $v_1 \mapsto_a v_2$, and $p \mapsto_r V$) have the meaning described above and **v** points to a valid stack.

This separation logic is used to define the relations (\leftarrow) and (\rightarrow) (depicted in Fig. 8) that are used in the definition of $[\cdot]_{r \Rightarrow a}$. Note that these definitions build on the definition of the Kripke wrapper in Appendix A as they maintain the state *s* for tracking the call stack in addition to the separation logic predicates. We define:

$$[M]_{r \Rightarrow a}^{a, d, d, m} \triangleq [M]_X \quad \text{where} \quad X \triangleq (\text{List}(\mathbf{Registers}), R_{r \Rightarrow a}, \leftarrow, \rightarrow, [], \bigstar_{v_1 \mapsto_a v_2 \in m})$$

F COROUTINE LINKING

Formally, $M_1 \oplus_{\text{coro}} M_2$ is defined using the generic linking operator $M_1 \oplus_X M_2$. Concretely, we define $M_1 \overset{d_1}{\oplus}_{\text{coro}} \overset{d_2, f}{\oplus} M_2 \triangleq M_1 \oplus_{X_{\text{coro}}} M_2$ where

$$X_{\text{coro}} \triangleq ((D \times \text{option}(\mathbf{FnName})), \rightsquigarrow_{\text{coro}}^{d_1, d_2}, (E, \text{Some}(f)))$$

Note that this linking operator is parametrized by a function name **f** of the initial function on the right side of the linking (**stream** in the example). The effect of linking is described by $\rightsquigarrow_{\text{coro}}$ shown in Fig. 9. There are many transitions, but most of them are straightforward. The rule **CORO-LINK-YIELD** encodes the core idea of \oplus_{coro} : If either the left side or the right side performs a call to **yield**, control switches to the other side, and the event is transformed to a **Return?(v, m)** event. There is one special case to consider: When M_1 calls **yield** the first time, there is no **yield** in M_2 from which to return. Instead this first call to **yield** becomes the invocation of a designated start function **f** in M_2 (**stream** in the example), as stated by **CORO-LINK-L-YIELD-INIT**. **CORO-LINK-INIT** handles the initial call from the environment to M_1 . If the environment tries to call a function in M_1 , the behavior is undefined (**CORO-LINK-INIT-UB**). **CORO-LINK-L-RETURN** handles the return from M_1 to the environment. M_2 should never return and thus **CORO-LINK-R-RETURN** states that doing so would lead to undefined behavior. Finally, **CORO-LINK-CALL** and **CORO-LINK-E-RETURN** allow both M_1 and M_2 to call external function (like **print**). However, M_1 and M_2 cannot directly call a function in the other module (without going through **yield**) (**CORO-LINK-CALL-UB**) and the environment may not call them back recursively (**CORO-LINK-CALL-UB**).

$$\begin{array}{c}
\text{CORO-LINK-YIELD} \\
\frac{(d = L \wedge d' = R) \vee (d = R \wedge d' = L)}{(d, (d, \text{None}), \text{Call}(\text{yield}, [v], m)) \rightsquigarrow_{\text{coro}}^{d_1, d_2} (d', (d', \text{None}), \text{Return}(v, m))} \\
\\
\text{CORO-LINK-YIELD-UB} \\
\frac{d = L \vee d = R \quad |\bar{v}| \neq 1}{(d, (d, \text{None}), \text{Call}(\text{yield}, \bar{v}, m)) \rightsquigarrow_{\text{coro}}^{d_1, d_2} \not\downarrow} \\
\\
\text{CORO-LINK-L-YIELD-INIT} \\
(L, (L, \text{Some}(f)), \text{Call}(\text{yield}, [v], m)) \rightsquigarrow_{\text{coro}}^{d_1, d_2} (R, (R, \text{None}), \text{Call}(f, [v], m)) \\
\\
\text{CORO-LINK-L-YIELD-INIT-UB} \\
\frac{|\bar{v}| \neq 1}{(L, (L, \text{Some}(f)), \text{Call}(\text{yield}, \bar{v}, m)) \rightsquigarrow_{\text{coro}}^{d_1, d_2} \not\downarrow} \\
\\
\text{CORO-LINK-INIT} \qquad \qquad \qquad \text{CORO-LINK-INIT-UB} \\
\frac{f \in |M_1|}{(E, (E, f^0), \text{Call}(f, \bar{v}, m)) \rightsquigarrow_{\text{coro}} (L, \text{Call}(f, \bar{v}, m), (L, f^0))} \qquad \frac{f \notin |M_1|}{(E, (E, f^0), \text{Call}(f, \bar{v}, m)) \rightsquigarrow_{\text{coro}} \not\downarrow} \\
\\
\text{CORO-LINK-L-RETURN} \qquad \qquad \qquad \text{CORO-LINK-R-RETURN} \\
(L, (L, f^0), \text{Return}(v, m)) \rightsquigarrow_{\text{coro}} (E, (E, f^0), \text{Return}(v, m)) \qquad (R, R, \text{Return}(v, m)) \rightsquigarrow_{\text{coro}} \not\downarrow \\
\\
\text{CORO-LINK-CALL} \\
\frac{f \neq \text{yield} \quad (d = L \wedge f \notin |M_2|) \vee (d = R \wedge f \notin |M_1|)}{(L, (d, f^0), \text{Call}(f, \bar{v}, m)) \rightsquigarrow_{\text{coro}} (E, (d, f^0), \text{Call}(f, \bar{v}, m))} \\
\\
\text{CORO-LINK-CALL-UB} \\
\frac{f \neq \text{yield} \quad (d = L \wedge f \in |M_2|) \vee (d = R \wedge f \in |M_1|)}{(L, (d, f^0), \text{Call}(f, \bar{v}, m)) \rightsquigarrow_{\text{coro}} \not\downarrow} \\
\\
\text{CORO-LINK-E-RETURN} \\
\frac{(s = L \wedge d = L) \vee (s = R \wedge d = R) \quad e = \text{Return}(_, _)}{(E, (d, f^0), e) \rightsquigarrow_{\text{coro}} (d, (d, f^0), e)} \\
\\
\text{CORO-LINK-E-CALL-UB} \\
\frac{(s = L \wedge d = L) \vee (s = R \wedge d = R) \quad e = \text{Call}(_, _, _)}{(E, (d, f^0), e) \rightsquigarrow_{\text{coro}} \not\downarrow}
\end{array}$$

Fig. 9. Definition of linking relation $\rightsquigarrow_{\text{coro}}^{d_1, d_2}$.

REFERENCES

- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Michael Sammler, Angus Hammond, Rodolphe Lepigre, Brian Campbell, Jean Pichon-Pharabod, Derek Dreyer, Deepak Garg, and Peter Sewell. 2022a. Islaris: verification of machine code against authoritative ISA semantics. In *PLDI*. ACM, 825–840. <https://doi.org/10.1145/3519939.3523434>
- Michael Sammler, Simon Spies, Youngju Song, Emanuele D’Osualdo, Robbert Krebbers, Deepak Garg, and Derek Dreyer. 2022b. DimSum: A Decentralized Approach to Multi-language Semantics and Verification (Coq development). <https://>

[//gitlab.mpi-sws.org/iris/dimsum/](https://gitlab.mpi-sws.org/iris/dimsum/).