

Modular Verification of Liveness Properties of the I/O Behavior of Imperative Programs

Bart Jacobs^[0000-0002-3605-249X]

KU Leuven, Department of Computer Science, imec-DistriNet Research Group
`bart.jacobs@cs.kuleuven.be`

Abstract. One way of verifying systems whose components interact by exchanging messages, such as distributed systems or certain types of concurrent systems, is by defining a protocol that governs the communication between the components and then verifying that each component's input and output (I/O) actions comply with its role in the protocol. In this paper, we propose a separation logic-based approach for specifying and verifying liveness properties of the I/O behavior of such components implemented as imperative programs, such as the property that a server eventually responds to each request. Our approach builds on earlier work for specifying safety properties of the I/O behavior of programs in separation logic by means of abstract nested Hoare triples, and encodes a liveness property verification problem into a termination verification problem by specifying that some appropriately chosen I/O operation (for example, the response to the N 'th request, for some unknown but fixed N) will cause the program to terminate.

1 Introduction

Consider the following program:

```
loop (let msg = recv() in send(msg))
```

This program implements a simple echo server. It sits in an infinite loop. In each iteration, it receives a message and echoes it back out. The problem we address in this paper is: how to specify and verify the safety and liveness properties of the I/O behavior of programs such as this one in Hoare logic [4]? We target Hoare logic so that we obtain a modular verification approach. For the example program specifically, we want to be able to specify the safety property that it only sends messages that it has received, and that it sends a message at most once¹; furthermore, we want to be able to specify the liveness property that it sends each message it has received at least once.

For specifying the safety properties of the I/O behavior, we apply our earlier work [9, 8, 7] on *abstract nested Hoare triples in separation logic* [11]. For specifying liveness properties, we build on earlier work [5] on verification of *basic*

¹ For simplicity, in this paper we assume a program never receives the same message more than once.

liveness. We combine and extend this work to obtain an approach that supports specifying a wide range of liveness properties. Adopting the terminology of Chang *et al.* [2], we show how our approach supports not just *guarantee properties* (of the form $\diamond p$, *eventually p* in temporal logic) (example: the program eventually terminates) and *simple response properties* (of the form $\Box \diamond p$, *always eventually p*) (example: the program always eventually responds to each request), but also *general response properties* (of the form $\bigwedge_i \Box \diamond p_i$) (example: the program always eventually responds to each request and always eventually emits a heartbeat signal), *simple reactivity properties* (of the form $\Box \diamond p \Rightarrow \Box \diamond q$) (example: if the program always eventually receives a request, it always eventually responds to each request), *general reactivity properties* (of the form $\bigwedge_i (\Box \diamond p_i \Rightarrow \Box \diamond q_i)$) (example: the program always eventually emits a heartbeat signal, and if it always eventually receives a request, it always eventually responds to each request) and *persistence properties* (of the form $\diamond \Box p$) (example: the program eventually always sends messages using the new message format).

The structure of this paper is as follows. In §2, we present our approach informally. In §3, we formalize the syntax and semantics of a simple programming language with I/O. In §4, we formalize our logic for liveness. We discuss related work in §5 and offer a conclusion in §6.

2 Our Approach, Informally

In this section, we present our approach informally. In subsequent sections, we formalize the programming language and the logic.

2.1 Safety

For specifying the safety properties of the I/O behavior, we apply our earlier work [9, 8, 7]: we use *abstract nested Hoare triples* in *separation logic* [11]. For example, a specification for the program `send("Hello"); send("world!")` could look as follows:

$$\{P_1 \wedge \text{send}_.(P_1, \text{"Hello"}, P_2) \wedge \text{send}_.(P_2, \text{"world!"}, P_3)\} \\ \text{send}(\text{"Hello"}); \text{send}(\text{"world!"}) \\ \{P_3\}$$

where the specification for `send` is

$$\{P \wedge \text{send}_.(P, m, Q)\} \text{send}(m) \{Q\}$$

and $\text{send}_.(P, m, Q)$ is a *higher-order predicate* that states that any resource that satisfies separation logic predicate P is sufficient to send message m ; doing so consumes such a resource and produces a resource that satisfies separation logic predicate Q . We refer to higher-order predicates such as $\text{send}_.$ as *abstract nested Hoare triples* because their meaning is similar to that of the Hoare triple $\{P\} \text{send}(m) \{Q\}$. The difference between abstract nested Hoare triples and actual nested Hoare triples is that the abstract ones are in fact just user-defined

predicates, so their meaning is entirely defined by the user, rather than having a fixed meaning assigned by the logic. The meaning of predicate `send_`, for example, is defined by the author of the module that implements function `send`.

We can read the specification for the "Hello, world!" example as follows: for any P_1 , P_2 , and P_3 , if a resource (satisfying) P_1 allows me to send "Hello" and obtain a resource P_2 , and P_2 allows me to send "world!" and obtain P_3 , and I have P_1 , then I can run safely and if I terminate, I will end up with P_3 .

In a similar fashion we can specify programs that perform both input and output. For example, the body of the echo loop can be specified as follows:

$$\begin{aligned} & \{P_1 \wedge \text{rcv}_-(P_1, m, P_2) \wedge \text{send}_-(P_2, m, P_3)\} \\ & \text{let msg = rcv() in send(msg)} \\ & \{P_3\} \end{aligned}$$

where the specification for `rcv` is

$$\{P \wedge \text{rcv}_-(P, m, Q)\} \text{rcv}() \{Q \wedge \text{res} = m\}$$

In postconditions, we use `res` to denote the result of a command. The predicate $\text{rcv}_-(P, m, Q)$ means that if you have P , you are allowed to receive. This will consume P and produce Q . Furthermore, the message you will receive is m . It follows that the specification for the single echo iteration states that it is allowed to receive a message, and then send the message it received, and that if it terminates, it shall indeed have performed these actions.

Using these ingredients, we can specify the full echo server, by first coinductively defining two predicates for expressing permission to receive and send an infinite sequence of messages, respectively:

$$\begin{aligned} \text{rcv_stream}(P, m \cdot \mu) &= \exists Q. \text{rcv}_-(P, m, Q) \wedge \text{rcv_stream}(Q, \mu) \\ \text{send_stream}(P, m \cdot \mu) &= \exists Q. \text{send}_-(P, m, Q) \wedge \text{send_stream}(P, \mu) \end{aligned}$$

We use μ to range over infinite sequences of messages², and $m \cdot \mu$ to denote the sequence with head m and tail μ . One possible specification for the safety of the echo server is then:

$$\begin{aligned} & \{P_r * P_s \wedge \text{rcv_stream}(P_r, \mu) \wedge \text{send_stream}(P_s, \mu)\} \\ & \text{loop (let msg = rcv() in send(msg))} \\ & \{\text{False}\} \end{aligned}$$

This is where separation logic's separating conjunction $-*$ comes in; it allows us to express that I have permission P_r to receive, and *separately* I have permission P_s to send. We can apply separation logic's *frame rule*

$$\frac{\{P\} c \{Q\}}{\{P * R\} c \{Q * R\}}$$

² For simplicity, we assume the same message does not appear twice in such a sequence. As a result, "echoing out the message at position i in μ " is equivalent to "responding to the i 'th request".

to the triple

$$\{P_r \wedge \text{recv_stream}(P_r, m \cdot \mu)\} \text{recv}() \{\exists P'_r. P'_r \wedge \text{recv_stream}(P'_r, \mu) \wedge \text{res} = m\}$$

taking *frame* $R = (P_s \wedge \text{send_stream}(P_s, m \cdot \mu))$ to obtain

$$\begin{aligned} & \{P_r * P_s \wedge \text{recv_stream}(P_r, m \cdot \mu) \wedge \text{send_stream}(P_s, m \cdot \mu)\} \\ & \text{recv}() \\ & \{\exists P'_r. P'_r * P_s \wedge \text{recv_stream}(P'_r, \mu) \wedge \text{send_stream}(P_s, \text{res} \cdot \mu)\} \end{aligned}$$

If we only had $P_r \wedge P_s$ instead of $P_r * P_s$, we would not be able to conclude that $\text{recv}()$ preserves P_s .

Note, however, that this specification requires that the server reply to requests in the order in which it receives them. This is overly strict, and forbids concurrent implementations. To allow the server to reply to requests in any order, we replace predicate send_stream by send_all , defined as follows:

$$\text{send_all}(P, m \cdot \mu) = \exists P_1, P_2. (P \Rightarrow P_1 * P_2) \wedge \text{send_}(P_1, m, \text{True}) \wedge \text{send_all}(P_2, \mu)$$

where $P \Rightarrow Q$, called a *view shift* [6], essentially simply means that P implies Q .³

2.2 Basic liveness

By itself, the logic we proposed in earlier work for verification of safety properties of the I/O behavior of programs guarantees that the program does not perform any I/O that is not permitted by the specification, and that *if* the program terminates, it will have performed the I/O prescribed by the specification, but it does not guarantee that the program will indeed perform the prescribed I/O. It might instead go into a silent infinite loop, and not perform any I/O from some point on. Or it might respond to some requests, but allow others to starve.

Of course, we could combine the I/O logic with any logic for verifying termination, such as the one we proposed in earlier work [5]. Interpreted in such a combined logic, the specifications shown above for the "Hello, world!" example and for the single echo loop iteration express full total correctness. However, this approach is not appropriate for the complete echo server: it is not supposed to terminate. (Indeed, the specification for the echo server shown above, interpreted in a total logic, is unimplementable.)

In earlier work [5], we proposed an approach for verifying *basic liveness* of non-terminating programs. Basic liveness means that the program always eventually performs I/O, i.e. that it never completely stops responding. Our approach was to reduce the basic liveness verification problem to a termination verification problem by imagining that the N 'th I/O operation performed by the program,

³ More precisely, $P \Rightarrow Q$ means that a state satisfying P can be transformed into a state satisfying Q by optionally performing an update of the *ghost state*. However, the reader can ignore the concept of ghost state for now.

for some fixed but unknown N , causes abrupt termination of the program. For example, in this approach, we can prove basic liveness of **loop** (**beep**()) by assuming the following specification for **beep**():

$$\{\mathbf{IO}(n)\} \mathbf{beep}() \{0 < n \wedge \mathbf{IO}(n-1)\}$$

where $\mathbf{IO}(n)$ means the $(n+1)$ 'th next I/O operation will terminate the program (so, in particular, $\mathbf{IO}(0)$ means that the very next I/O operation will terminate the program), and then proving

$$\{\mathbf{IO}(n)\} \mathbf{loop} (\mathbf{beep}()) \{\mathbf{False}\}$$

in a total logic. In words, this specification states that, assuming that the $(n+1)$ 'th I/O operation performed by the program terminates the program, the program terminates, and furthermore, the program does not terminate normally (because the postcondition is **False**, so **skip** and **beep**(); **beep**() do not satisfy this specification⁴). Here, the number of I/O operations left before the program is terminated can be used as a loop variant.

2.3 Simple responsiveness

Notice that for the echo server, basic liveness is insufficient as a specification: an implementation could satisfy it simply by receiving requests, but not responding to any of them.

In this paper, we improve upon our earlier work by proposing an approach for specifying and verifying not just basic liveness, but richer, application-specific liveness properties of a program's I/O behavior as well, such as the property that our echo server eventually responds to each request. We again reduce the liveness property verification problem to a termination verification problem by imagining that a particular well-chosen I/O operation causes abrupt termination of the program. For example, to prove that the echo server eventually responds to each request, it is sufficient to prove that it terminates, under the assumption that the response to the k 'th request, for some fixed but unknown k , terminates the program. Since the choice of which I/O operation terminates the program is application-dependent, we do not encode it directly into the I/O primitives' specification, as we did for basic liveness verification. Instead, we integrate it with the program's I/O safety specification, by stating in the program's precondition that the postcondition of the I/O operation that terminates the program is **False**. For example, if we apply this approach to the echo server, we obtain the following specification:

$$\begin{aligned} & \{P_r * P_s \wedge \mathbf{recv_stream}(P_r, \mu) \wedge \mathbf{send_all}'(P_s, k, \mu) \wedge 0 \leq k\} \\ & \mathbf{loop} (\mathbf{let} \text{ msg} = \mathbf{recv}() \mathbf{in} \mathbf{send}(\text{msg})) \\ & \{\mathbf{False}\} \end{aligned}$$

⁴ Actually, **beep**(); **beep**() does satisfy the specification for $n < 2$, but the point is that it does not satisfy the specification $\forall n. \{\mathbf{IO}(n)\} - \{\mathbf{False}\}$.

where $\text{send_all}'(P, k, \mu)$ is defined in exactly the same way that $\text{send_all}(P, \mu)$ was defined above (so it means that resource P gives permission to send all of the messages in μ , in any order), except that the postcondition of sending the message at index k in μ is **False** (so sending the message at index k in μ terminates the program) (if $0 \leq k$)⁵:

$$\begin{aligned} \text{send_all}'(P, k, m \cdot \mu) = \\ \exists P_1, P_2. (P \Rightarrow P_1 * P_2) \wedge \text{send}_.(P_1, m, k \neq 0) \wedge \text{send_all}'(P_2, k - 1, \mu) \end{aligned}$$

where $k \in \mathbb{Z}$. Again, k can be used as a loop variant to prove termination of the loop.

This approach also allows us to verify other echo server implementations, such as one that performs buffering and reordering of requests, against the same specification:

```

{Pr * Ps ∧ recv_stream(Pr, μ) ∧ send_all'(Ps, k, μ)}
loop (
  let msg1 = recv() in
  let msg2 = recv() in
  let msg3 = recv() in
  send(msg3);
  send(msg1);
  send(msg2)
)
{False}

```

Other simple responsiveness properties (i.e. properties of the form $\Box \Diamond p$) can be encoded similarly.

2.4 General responsiveness

Suppose the echo server should eventually respond to each request, and also eventually log each request. We can reduce this property to termination by imagining that *either* the response to the k 'th request, for some k , *or* logging the j 'th request, for some j , terminates the program:

```

{ Pr * Ps * Pl ∧ recv_stream(Pr, μ) ∧
  send_all'(Ps, k, μ) ∧ log_all'(Pl, j, μ) ∧ (0 ≤ k ∨ 0 ≤ j) }
loop (let msg = recv() in send(msg); log(msg))
{False}

```

Since the program only knows that *either* $0 \leq k$ *or* $0 \leq j$, it must *both* log all requests *and* respond to all of them to be sure of termination.

Other general responsiveness properties (i.e. properties of the form $\bigwedge_i \Box \Diamond p_i$) can be encoded similarly.

⁵ Notice that if $k < 0$, $\text{send_all}'(P, k, \mu)$ is equivalent to $\text{send_all}(P, \mu)$.

2.5 Reactivity

Suppose receiving can suffer transient failures. To prove that the echo server always eventually sends, we need to assume that receiving always eventually succeeds. This can be expressed as follows:

```

{Pr * Ps ∧ recv_stream'(Pr, ν, μ) ∧ send_all'(Ps, k, μ)}
loop (let msg = recv() in if msg ≠ ⊥ then send(msg))
{False}

```

where $\text{recv_stream}'(P, \nu, \mu)$ means that receiving message μ_i will fail ν_i times before succeeding, for all $i \in \mathbb{N}$:

$$\begin{aligned} \text{recv_stream}'(P, 0 \cdot \nu, m \cdot \mu) &= \exists Q. \text{recv_}(P, m, Q) \wedge \text{recv_stream}'(Q, \nu, \mu) \\ \text{recv_stream}'(P, (n+1) \cdot \nu, \mu) &= \exists Q. \text{recv_}(P, \perp, Q) \wedge \text{recv_stream}'(Q, n \cdot \nu, \mu) \end{aligned}$$

and $\nu \in \mathbb{N}^\omega$ ranges over infinite sequences of natural numbers.

Other reactivity properties (i.e. properties of the form $\Box \diamond p \Rightarrow \Box \diamond q$) can be encoded similarly.

2.6 General reactivity

Suppose the echo server should always eventually emit a heartbeat signal, even if receiving fails persistently:

```

{ Pr * Ps * Ph ∧ recv_stream'(Pr, ν̃, μ) ∧
  { send_all'(Ps, k, μ) ∧ heartbeats'(Ph, j) ∧ (ν̃ ∈ ℕω ∧ 0 ≤ k ∨ 0 ≤ j) }
loop (let msg = recv() in heartbeat()); if msg ≠ ⊥ then send(msg)
{False}

```

where $\tilde{\nu} \in (\mathbb{N} \cup \{\infty\})^\omega$ and $\text{heartbeats}'$ is defined as follows:

$$\begin{aligned} \text{heartbeats}'(P, j) &= \\ &\exists P_1, P_2. (P \Rightarrow P_1 * P_2) \wedge \text{heartbeat_}(P_1, j \neq 0) \wedge \text{heartbeats}'(P_2, j-1) \end{aligned}$$

Other general reactivity properties (i.e. properties of the form $\bigwedge_i (\Box \diamond p_i \Rightarrow \Box \diamond q_i)$) can be encoded similarly.

2.7 Persistence

Suppose the echo server is allowed to drop a finite number of requests:

```

{ Pr * Ps ∧ recv_stream(Pr, μ) ∧ send_all''(Ps, μ) }
commit(false); commit(false); commit(false);
recv(); let msg = recv() in send(msg); recv();
commit(true);
loop (let msg = recv() in send(msg))
{False}

```

which uses the following auxiliary definition:

$$\begin{aligned} \text{send_all}''(P, m \cdot \mu) = & \\ & (\exists Q_1, Q_2. \text{commit}_-(P, \text{false}, Q_1 * Q_2) \wedge \text{send}_-(Q_1, m, \text{True}) \wedge \text{send_all}''(Q_2, \mu)) \wedge \\ & (\exists Q, k \geq 0. \text{commit}_-(P, \text{true}, Q) \wedge \text{send_all}'(Q, k, m \cdot \mu)) \end{aligned}$$

To terminate, the program has to finish the phase where it is allowed to drop requests (represented by predicate `send_all''`) and enter the phase where it responds to each request (represented by predicate `send_all'`); it needs to signal this by performing the *ghost I/O action* `commit(true)`.⁶

Other persistence properties (i.e. properties of the form $\diamond \square p$) can be encoded similarly.

3 A Programming Language with I/O

We present our approach in the context of a simplified ML-like programming language with support for I/O. Its grammar is as follows:

$$\begin{aligned} v \in \text{Vals}, x \in \text{Vars}, t \in \text{IOPrims} \\ e \in \text{Exprs} ::= v \mid x \\ c \in \text{Cmds} ::= e \mid t(e) \mid \mathbf{if} \ e = e \ \mathbf{then} \ c \ \mathbf{else} \ c \mid \mathbf{let} \ x = c \ \mathbf{in} \ c \mid \mathbf{loop} \ c \end{aligned}$$

We assume a set *Vals* of values, *Vars* of program variables, and *IOPrims* of I/O primitives.⁷ We assume *Vals* contains at least the unit value `()`.

We define $c; c' = \mathbf{let} \ x = c \ \mathbf{in} \ c'$ where x does not appear in c' .

We define the I/O actions $\alpha \in \text{IOActions} ::= t(v, v')$; in $t(v, v')$, we call v the *argument* and v' the *result*. We coinductively define the set of traces $\tau \in \text{Traces} ::= \text{Div} \mid \text{Ret}(v) \mid t(v, v') \cdot \tau$. A trace ending in `Div` is called a *diverging trace*; it represents the behavior where the program eventually neither terminates nor performs I/O. A trace ending in `Ret(v)` is a finite trace; it represents the behavior where the program eventually terminates with result value v .

We define concatenation of traces coinductively as follows:

$$\text{Ret}(v);_v \tau = \tau \quad \text{Div};_v \tau = \text{Div} \quad \frac{\tau;_v \tau' = \tau''}{t(v', v'') \cdot \tau;_v \tau' = t(v', v'') \cdot \tau''}$$

We define the language's semantics by means of a big-step relation $c \Downarrow \tau$, which relates a command c to a trace τ . We define the relation coinductively by

⁶ The `commit` *ghost commands* are inserted into the program text for verification purposes only. Since they do not have any observable effect, any properties proven about the ghost-instrumented program hold also for the original program (also known as the *erased* program).

⁷ In the examples, we assume `send, recv, beep, log, heartbeat, commit` $\in \text{IOPrims}$.

means of the following rules:

$$\begin{array}{c}
v \Downarrow \text{Ret}(v) \quad \frac{c \Downarrow \tau}{\mathbf{if } v = v' \mathbf{ then } c \mathbf{ else } c' \Downarrow \tau} \quad \frac{v \neq v' \quad c' \Downarrow \tau}{\mathbf{if } v = v' \mathbf{ then } c \mathbf{ else } c' \Downarrow \tau} \\
\\
\frac{c \Downarrow \tau \quad c'[v/x] \Downarrow \tau' \quad \tau;_v \tau' = \tau''}{\mathbf{let } x = c \mathbf{ in } c' \Downarrow \tau''} \quad \frac{c; \mathbf{loop } c \Downarrow \tau}{\mathbf{loop } c \Downarrow \tau} \quad t(v) \Downarrow t(v, v') \cdot \text{Ret}(v')
\end{array}$$

Notice that we have $\mathbf{loop } () \Downarrow \text{Div}$. We also have $\mathbf{loop } () \Downarrow \tau$ for any other trace τ as well; this imprecision is harmless, since we will be proving liveness.

3.1 Liveness properties

If we define $c_{\text{echo}} = \mathbf{loop } (\mathbf{let } \text{msg} = \text{recv}() \mathbf{ in } \text{send}(\text{msg}))$, we can state the responsiveness property of the echo server as follows:

$$\forall \tau, \mu. c_{\text{echo}} \Downarrow \tau \wedge \tau|_{\text{recv}} \preceq \mu \Rightarrow \forall m \in \mu. \text{send}(m, ()) \in \tau$$

where we define $\tau|_{\text{recv}} \preceq \mu$ coinductively as follows:

$$\begin{array}{c}
\text{Div}|_{\text{recv}} \preceq \mu \quad \text{Ret}(v)|_{\text{recv}} \preceq \mu \quad \frac{\tau|_{\text{recv}} \preceq \mu}{(\text{recv}(-, m) \cdot \tau)|_{\text{recv}} \preceq m \cdot \mu} \\
\\
\frac{\tau|_{\text{recv}} \preceq \mu}{(\text{send}(-, ()) \cdot \tau)|_{\text{recv}} \preceq \mu}
\end{array}$$

To make the link with temporal logic: this is more or less equivalent to

$$\forall \tau. c_{\text{echo}} \Downarrow \tau \Rightarrow \tau \models \Box \Diamond (\exists m. \text{recv}(-, m)) \wedge \Box (\forall m. \text{recv}(-, m) \Rightarrow \Diamond \text{send}(m, -))$$

(It is equivalent if we ignore the possibility of messages being sent before they are received.)

In the next section, we define a separation logic for modularly verifying properties such as this one, and argue formally that it is indeed adequate for this purpose.

4 A Program Logic for I/O Liveness

4.1 Exit actions and exit traces

We use the notation $\dot{v} \in \text{Vals} \cup \{\top\}$ to denote either a program value or the special *exit value* \top . An *exit action* $\dot{\alpha} = t(v, \dot{v})$ is like an action except that its result may be the exit value. An *exit trace* $\dot{\tau}$ is a finite sequence of exit actions. We say an exit action is *exiting* if its result is the exit value.

The I/O safety logic from our earlier work [8] proves that a program's partial traces are prefixes of a given set of traces. A total logic (such as [5]) proves that a

program's traces end with $\text{Ret}(_)$, or, equivalently, that they start with a partial trace that ends with $\text{Ret}(_)$. By combining and slightly generalizing these, we can obtain a logic whose correctness judgments imply that each of a program's traces starts with one of the partial traces from a given set of *exit traces*, i.e. partial traces that, as soon as an exiting I/O action is performed, cause the program to be considered "terminated". Examples of such *exit sets* are: the set of partial traces T_1 where the program has responded to the first request, the set of partial traces T_2 where the program has responded to the second request, etc. By proving a universally quantified correctness judgment, we can obtain that a program's traces start with a partial trace from T_1 *and* with a partial trace from T_2 , etc., that is, that the program responds to the first request *and* to the second request, etc., allowing us to conclude that the program satisfies the responsiveness property.

In the remainder of this section, we elaborate this approach.

4.2 Petri nets

The I/O safety logic is based on the idea that the program must own particular resources in order to be allowed to perform a given I/O action. When using the logic, there is no need to specify the particular nature of those resources. However, for the sake of proving adequacy of the logic, we do need to introduce a particular ontology of resources. For this purpose, we here use a particular type of *Petri nets*. Indeed, these resources serve very much like *tokens* in a Petri net. If, in a Petri net, there is a token in each *pre-place* of a *transition*, the transition can *fire*, which removes one token from each pre-place and adds one to each *post-place*. This, in turn, can enable a new transition, and so on. If we label the transitions by I/O actions, we obtain that a *marking* $V \in \mathcal{P} \rightarrow \mathbb{N}$ of a Petri net (which maps each place to the number of tokens present at that place) defines a set of traces.

Specifically, we will be labeling transitions by exit actions, so that a marking defines a set of exit traces.

Let \mathcal{P} be a set of places. We use p and q to range over places. We consider Petri nets where the set of transitions N is a subset of the set \mathcal{N} defined as

$$\chi \in \mathcal{N} ::= t(V_{\text{pre}}, v, \dot{v}, V_{\text{post}}) \mid \mathbf{noop}(V_{\text{pre}}, V_{\text{post}})$$

where V_{pre} and V_{post} are multisets of places, called the *pre-places* and *post-places* of the transition, respectively.

A Petri net defines a labeled step relation \rightarrow and a corresponding labeled reachability relation \twoheadrightarrow on markings:

$$\frac{t(V_{\text{pre}}, v, \dot{v}, V_{\text{post}}) \in N}{V \uplus V_{\text{pre}} \xrightarrow{t(v, \dot{v})} V \uplus V_{\text{post}}} \quad \frac{\mathbf{noop}(V_{\text{pre}}, V_{\text{post}}) \in N}{V \uplus V_{\text{pre}} \xrightarrow{\epsilon} V \uplus V_{\text{post}}} \quad V \xrightarrow{\epsilon} V$$

$$\frac{V \xrightarrow{\dot{\tau}} V' \quad V' \xrightarrow{\dot{\tau}'} V''}{V \xrightarrow{\dot{\tau} \cdot \dot{\tau}'} V''}$$

where $V \uplus V' = \lambda p. V(p) + V'(p)$. We define $\text{Traces}_N(V) = \{\dot{\tau} \mid \exists V'. V \xrightarrow{\dot{\tau}} V'\}$.

4.3 Assertions, correctness judgments, view shifts

We define the set of *assertions* semantically as the set of sets of markings. That is, an assertion describes a marking. An assertion $\mathbf{tokens}(V)$ describes a marking that includes V : $\mathbf{tokens}(V) = \{V' \mid \forall p. V'(p) \geq V(p)\}$. The separating conjunction $P * P'$ describes a marking that can be split into one that satisfies P and one that satisfies P' : $P * P' = \{V \uplus V' \mid V \in P \wedge V' \in P'\}$.

We define the meaning of a correctness judgment $\{P\} c \{Q\}$, where precondition P is an assertion, c is a command, and postcondition Q maps a value to an assertion, as follows:

$$\{P\} c \{Q\} \Leftrightarrow \forall V, \tau. V \in P \wedge c \Downarrow \tau \Rightarrow \mathbf{safe}(\tau, \text{Traces}_N(V), Q)$$

where $\mathbf{safe}(\tau, T, Q)$ is defined inductively by the following rules:

$$\frac{\exists V \in Q(v). \text{Traces}_N(V) \subseteq T}{\mathbf{safe}(\text{Ret}(v), T, Q)} \qquad \frac{t(v, \top) \in T}{\mathbf{safe}(t(v, v') \cdot \tau, T, Q)}$$

$$\frac{t(v, v'') \in T \quad v' \neq v''}{\mathbf{safe}(t(v, v') \cdot \tau, T, Q)} \qquad \frac{t(v, v') \in T \quad \mathbf{safe}(\tau, \{\dot{\tau}' \mid t(v, v') \cdot \dot{\tau}' \in T\}, Q)}{\mathbf{safe}(t(v, v') \cdot \tau, T, Q)}$$

The set T can be seen as a *specification* that expresses safety and liveness properties of the program, as well as assumptions about the environment. Notice that infinite or diverging traces are safe only if they contain an input that conflicts with the specification (i.e. the environment assumptions are violated) or perform an action that corresponds to an exiting action of the specification.

Informally, $\{P\} c \{Q\}$ states that for every marking V that satisfies P , assuming that the environment provides only inputs allowed by the traces of V , command c performs only outputs allowed by the traces of V and terminates, either by performing an exiting action, or by returning with a result v and a marking V' such that $V' \in Q(v)$.

We say that P view-shifts to Q , written as $P \Rightarrow Q$, if $\forall V \in P. \exists V' \in Q. \text{Traces}_N(V') \subseteq \text{Traces}_N(V)$. That is, we can replace the current marking by another one that is equivalent or more restrictive in terms of the program outputs it allows.⁸

⁸ Note: reducing the set of traces does not strengthen the assumptions on the environment, because if two traces of a specification make conflicting assumptions about environment behavior, the resulting assumption is the *conjunction* of these, i.e. **False**.

4.4 Proof rules

Given these definitions, the following proof rules are admissible:

$$\frac{\{Q(v)\} v \{Q\} \quad \frac{\{P \wedge v = v'\} c \{Q\} \quad \{P \wedge v \neq v'\} c' \{Q\}}{\{P\} \mathbf{if} v = v' \mathbf{then} c \mathbf{else} c' \{Q\}}}{\{Q(v)\} v \{Q\}}$$

$$\frac{\{P\} c \{Q\} \quad \forall v. \{Q(v)\} c'[v/x] \{R\}}{\{P\} \mathbf{let} x = c \mathbf{in} c' \{R\}} \quad \frac{\forall n. \{P_n\} c \{\exists n' < n. P_{n'}\}}{\{P_m\} \mathbf{loop} c \{\mathbf{False}\}}$$

$$\{\mathbf{tokens}(V_{\text{pre}}) \wedge t(V_{\text{pre}}, v, \dot{v}, V_{\text{post}}) \in N\} t(v) \{\mathbf{tokens}(V_{\text{post}}) \wedge \dot{v} \neq \top \wedge \mathbf{res} = \dot{v}\}$$

$$\frac{P \Rightarrow P' \quad \frac{\{P'\} c \{Q\} \quad Q \Rightarrow Q'}{\{P\} c \{Q'\}} \quad \frac{\{P\} c \{Q\}}{\{P * R\} c \{Q * R\}}}{\frac{\{P\} c \{Q\} \quad Q \Rightarrow Q'}{\{P\} c \{Q'\}} \quad \frac{\{P\} c \{Q\}}{\{P * R\} c \{Q * R\}}}$$

$$\frac{\forall i \in I. \{P_i\} c \{Q\}}{\{\exists i \in I. P_i\} c \{Q\}} \quad \frac{P \Rightarrow Q}{P \Rightarrow Q} \quad \frac{P \Rightarrow Q}{P * R \Rightarrow Q * R}$$

$$\frac{\mathbf{noop}(V_{\text{pre}}, V_{\text{post}}) \in N}{\mathbf{tokens}(V_{\text{pre}}) \Rightarrow \mathbf{tokens}(V_{\text{post}})} \quad \frac{P \Rightarrow P' \quad P' \Rightarrow P''}{P \Rightarrow P''}$$

We sometimes write postconditions as assertions with a free variable `res`.

4.5 Abstract nested Hoare triples notation

We define the abstract nested Hoare triple notation $t_{-}(P, v, v', Q)$ as follows:

$$t_{-}(P, v, v', Q) = \frac{P \Rightarrow \exists V_{\text{pre}}, V_{\text{post}}. \mathbf{tokens}(V_{\text{pre}}) \wedge (t(V_{\text{pre}}, v, v', V_{\text{post}}) \in N \wedge \mathbf{tokens}(V_{\text{post}}) \Rightarrow Q \vee t(V_{\text{pre}}, v, \top, V_{\text{post}}) \in N)}{t_{-}(P, v, v', Q)}$$

It follows that we have $\{P \wedge t_{-}(P, v, v', Q)\} t(v) \{Q \wedge \mathbf{res} = v'\}$.

4.6 Example: simple responsiveness

We show now how we can use our logic to verify the responsiveness property of the echo server, repeated here:

$$\forall \tau, \mu. c_{\text{echo}} \Downarrow \tau \wedge \tau|_{\text{recv}} \preceq \mu \Rightarrow \forall m \in \mu. \mathbf{send}(m, ()) \in \tau$$

First of all, we fix μ . Again, we assume no message appears more than once in μ .

It is sufficient to prove

$$\forall k \geq 0, \tau. c_{\text{echo}} \Downarrow \tau \wedge \tau|_{\text{recv}} \preceq \mu \Rightarrow \mathbf{send}(\mu_k, ()) \in \tau$$

where μ_k denotes the message at index k in μ .

We fix k . We define exit set T as follows:

$$T = \{\dot{\tau} \mid \dot{\tau}|_{\text{recv}} \preceq \mu \wedge (\forall \text{send}(v, \dot{v}) \in \dot{\tau}. v = \mu_k \vee \dot{v} = ())\}$$

where we define $\dot{\tau}|_{\text{recv}} \preceq \mu$ as follows:

$$\epsilon|_{\text{recv}} \preceq \mu \quad \frac{\dot{\tau}|_{\text{recv}} \preceq \mu}{(\text{recv}(_, m) \cdot \dot{\tau})|_{\text{recv}} \preceq m \cdot \mu} \quad \frac{\dot{v} \in \{(), \top\} \quad \dot{\tau}|_{\text{recv}} \preceq \mu}{(\text{send}(_, \dot{v}) \cdot \tau)|_{\text{recv}} \preceq \mu}$$

It is sufficient to prove

$$\forall \tau. c_{\text{echo}} \Downarrow \tau \Rightarrow \text{safe}(\tau, T, \text{False})$$

where μ_k denotes the message at index k in μ .

Indeed, we can prove, by induction on the derivation of $\text{safe}(\tau, T, \text{False})$, that for any τ , if $\text{safe}(\tau, T, \text{False})$ and $\tau|_{\text{recv}} \preceq \mu$, then $\text{send}(\mu_k, ()) \in \tau$.

We construct a Petri net (\mathcal{P}, N) and a marking V such that $\text{Traces}_N(V) \subseteq T$. We define

$$\mathcal{P} = \{\text{recv}_i \mid 0 \leq i\} \cup \{\text{send}_i \mid 0 \leq i\}$$

and

$$N = \{\text{recv}(\{\{\text{recv}_i\}, (), \mu_i, \{\{\text{recv}_{i+1}\}\}) \mid i \geq 0\} \cup \{\text{send}(\{\{\text{send}_i\}, \mu_i, (), \mathbf{0}\} \mid i \geq 0 \wedge i \neq k\} \cup \{\text{send}(\{\{\text{send}_k\}, \mu_k, \top, \mathbf{0}\})\}$$

and

$$V = \lambda p. \begin{cases} 1 & \text{if } p = \text{recv}_0 \\ 1 & \text{if } \exists i. p = \text{send}_i \\ 0 & \text{otherwise} \end{cases}$$

where $\{\{p\}\}$ denotes the singleton multiset with a single token at p , and $\mathbf{0}$ denotes the empty multiset.

It is easy to check that indeed $\text{Traces}_N(V) \subseteq T$, and that $V \in P$ where $P = \exists P_r, P_s. P_r * P_s \wedge \text{recv_stream}(P_r, \mu) \wedge \text{send_all}'(P_s, k, \mu)$.

Then, by the meaning of correctness judgments, we have that $\{P\} c_{\text{echo}} \{\text{False}\}$ implies the goal.

4.7 General responsiveness

The adequacy of the verification approach for general responsiveness properties presented in §2 follows directly from the adequacy for simple responsiveness properties, combined with the observation that $\{P_1 \vee P_2\} c \{Q\}$ implies both $\{P_1\} c \{Q\}$ and $\{P_2\} c \{Q\}$. Therefore, given a successful verification of the program, each constituent simple responsiveness property can then be established separately as above.

4.8 Example: persistence

We wish to prove that c'_{echo} , defined as

$$c'_{\text{echo}} = \text{rcv}(); \mathbf{let} \text{ msg} = \text{rcv}() \mathbf{in} \text{ send}(\text{msg}); \text{rcv}(); \\ \mathbf{loop} (\mathbf{let} \text{ msg} = \text{rcv}() \mathbf{in} \text{ send}(\text{msg}))$$

drops only finitely many messages, i.e. that it eventually always responds. Formally:

$$\forall \mu, \tau. c'_{\text{echo}} \Downarrow \tau \wedge \tau|_{\text{rcv}} \preceq \mu \Rightarrow \exists k_0 \geq 0. \forall k \geq k_0. \text{send}(\mu_k, ()) \in \tau$$

We consider a *ghost-instrumented* version \tilde{c}'_{echo} of c'_{echo} , defined as follows:

$$\tilde{c}'_{\text{echo}} = \text{commit}(\text{false}); \text{commit}(\text{false}); \text{commit}(\text{false}); \\ \text{rcv}(); \mathbf{let} \text{ msg} = \text{rcv}() \mathbf{in} \text{ send}(\text{msg}); \text{rcv}(); \\ \text{commit}(\text{true}); \\ \mathbf{loop} (\mathbf{let} \text{ msg} = \text{rcv}() \mathbf{in} \text{ send}(\text{msg}))$$

Obviously, it is sufficient to prove

$$\forall \mu, \tau. \tilde{c}'_{\text{echo}} \Downarrow \tau \wedge \text{erasure}(\tau)|_{\text{rcv}} \preceq \mu \Rightarrow \exists k_0 \geq 0. \forall k \geq k_0. \text{send}(\mu_k, ()) \in \tau$$

where $\text{erasure}(\tau)$ removes the `commit` actions (replacing an infinite sequence of `commit` actions with `Div`). We define $\#\tau$ as the index of the first `commit(true, -)` action in τ , or 0 if it does not contain such an action. It is sufficient to prove

$$\forall \mu, \tau. \tilde{c}'_{\text{echo}} \Downarrow \tau \wedge \text{erasure}(\tau)|_{\text{rcv}} \preceq \mu \Rightarrow \forall k \geq 0. \text{send}(\mu_{\#\tau+k}, ()) \in \tau$$

We fix μ and k and define exit set T as follows:

$$T = \left\{ \dot{\tau} \mid \begin{array}{l} \text{erasure}(\dot{\tau})|_{\text{rcv}} \preceq \mu \wedge \\ (\forall \text{send}(v, \dot{v}) \in \dot{\tau}. \text{commit}(\text{true}, -) \in \dot{\tau} \wedge v = \mu_{\#\dot{\tau}+k} \vee \dot{v} = ()) \end{array} \right\}$$

It is sufficient to prove

$$\forall \tau. \tilde{c}'_{\text{echo}} \Downarrow \tau \Rightarrow \text{safe}(\tau, T, \text{False})$$

Indeed, one can again show, by induction on the derivation of $\text{safe}(\tau, T, \text{False})$ that for all τ , if $\text{safe}(\tau, T, \text{False})$ and $\text{erasure}(\tau)|_{\text{rcv}} \preceq \mu$, then $\text{send}(\mu_{\#\tau+k}, ()) \in \tau$.

The remainder of this example proceeds as above: we construct a Petri net and a marking that satisfies the precondition we used for verifying the persistence example in §2 and whose traces are included in T . Successful verification then implies the goal.

5 Related work

We are not aware of existing work on Hoare logics for verifying liveness properties of the I/O behavior of programs. To the best of our knowledge, most approaches for verifying liveness properties of the I/O behavior of systems have so far been based on a representation of the system as some kind of a state machine, or a set of interacting processes or state machines, rather than a program. In these approaches, the liveness properties of interest are very often specified using temporal logic [2].

Even when it comes to Hoare logics for liveness properties of other aspects of program execution, there is very little existing work. We are aware of only two lines of work. Boström and Müller [1] verify that blocked threads in a multithreaded program are always eventually unblocked. In ongoing work, D’Osualdo *et al.* [3] verify termination under a fair scheduler of multithreaded programs that involve synchronization based on busy-waiting.

For a discussion of related work on verifying safety of I/O behavior and on verifying program termination, we refer to our earlier work [9, 5].

6 Conclusion

We presented a Hoare logic-based approach for the specification and verification of liveness properties of the I/O behavior of program modules. Our approach is based on the idea of reducing the problem to a termination verification problem and then applying existing approaches for I/O safety verification and termination verification. Our approach can be applied straightforwardly in existing verification tools that support separation logic, higher-order predicates, and termination verification, such as our VeriFast tool [12].

In this paper, we considered a very simple programming language with no dynamically allocated memory, higher-order functions or dynamic method binding, or concurrency. However, we believe the ideas of this paper can be integrated straightforwardly into the separation logic for total correctness of multithreaded object-oriented programs from our earlier work [5] to verify liveness properties of the I/O behavior of a large class of realistic programs.

A limitation of this earlier work, however, is that it does not support *busy waiting*, a common practice in programs for multiprocessor machines. In recent work [10], we propose a logic for verifying termination under fair scheduling of programs where threads busy-wait for other threads to abruptly terminate the program. By combining that logic with the ideas from this paper, one can obtain a logic for verifying responsiveness of a multithreaded server where one thread receives requests and another responds to them: the first thread can be seen as busy-waiting for the second thread to terminate the program by responding to the k ’th request.

We have not addressed the question of completeness: does our approach support all possible liveness properties? Chang *et al.* [2] claim that any temporal logic formula is equivalent to a general reactivity formula. However, their setting is not quite the same as ours, as evidenced by the fact that we need to use

ghost I/O operations to support persistence properties, whereas in their framework persistence properties are subsumed by reactivity properties. A thorough investigation of this question is future work.

References

1. Boström, P., Müller, P.: Modular verification of finite blocking in non-terminating programs. In: Boyland, J.T. (ed.) 29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5-10, 2015, Prague, Czech Republic. LIPIcs, vol. 37, pp. 639–663. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015)
2. Chang, E., Manna, Z., Pnueli, A.: The safety-progress classification. In: Bauer, F.L., Brauer, W., Schwichtenberg, H. (eds.) Logic and Algebra of Specification. pp. 143–202. Springer Berlin Heidelberg, Berlin, Heidelberg (1993)
3. D’Osualdo, E., Farzan, A., Gardner, P., Sutherland, J.: TaDA Live: Compositional reasoning for termination of fine-grained concurrent programs. CoRR **abs/1901.05750** (2019)
4. Hoare, C.A.R.: An axiomatic basis for computer programming. Commun. ACM **12**(10), 576–580 (1969)
5. Jacobs, B., Bosnacki, D., Kuiper, R.: Modular termination verification of single-threaded and multithreaded programs. ACM Trans. Program. Lang. Syst. **40**(3), 12:1–12:59 (2018)
6. Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. J. Funct. Program. **28**, e20 (2018)
7. Penninckx, W., Jacobs, B., Piessens, F.: Sound, modular and compositional verification of the input/output behavior of programs. In: Vitek, J. (ed.) Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9032, pp. 158–182. Springer (2015)
8. Penninckx, W., Timany, A., Jacobs, B.: Abstract I/O specification. CoRR **abs/1901.10541** (2019)
9. Penninckx, W., Timany, A., Jacobs, B.: Specifying I/O using abstract nested Hoare triples in separation logic. In: Proceedings of the 21st Workshop on Formal Techniques for Java-like Programs. FTfJP ’19, Association for Computing Machinery, New York, NY, USA (2019)
10. Reinhard, T., Timany, A., Jacobs, B.: A separation logic to verify termination of busy-waiting for abrupt program exit. In: FTfJP (2020), to appear.
11. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: 17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings. pp. 55–74. IEEE Computer Society (2002)
12. Vogels, F., Jacobs, B., Piessens, F.: Featherweight VeriFast. Logical Methods in Computer Science **11**(3) (2015)