Theorems for Free from Separation Logic Specifications

LARS BIRKEDAL, Aarhus University, Denmark

THOMAS DINSDALE-YOUNG, Concordium, Denmark

ARMAËL GUÉNEAU, Aarhus University, Denmark

GUILHEM JABER, Université de Nantes, France

KASPER SVENDSEN, Uber, Denmark

NIKOS TZEVELEKOS, Queen Mary University of London, United Kingdom

Separation logic specifications with abstract predicates intuitively enforce a discipline that constrains when and how calls may be made between a client and a library. Thus a separation logic specification of a library intuitively enforces a protocol on the trace of interactions between a client and the library. We show how to formalize this intuition and demonstrate how to derive “free theorems” about such interaction traces from abstract separation logic specifications. We present several examples of free theorems. In particular, we prove that a so-called *logically atomic* concurrent separation logic specification of a concurrent module operation implies that the operation is linearizable. All the results presented in this paper have been mechanized and formally proved in the Coq proof assistant using the Iris higher-order concurrent separation logic framework.

CCS Concepts: • **Theory of computation** → **Program specifications; Separation logic.**

Additional Key Words and Phrases: Separation Logic, Program Specifications, Linearizability, Iris

ACM Reference Format:

Lars Birkedal, Thomas Dinsdale-Young, Armaël Guéneau, Guilhem Jaber, Kasper Svendsen, and Nikos Tzevelekos. 2021. Theorems for Free from Separation Logic Specifications. *Proc. ACM Program. Lang.* 5, ICFP, Article 81 (August 2021), 29 pages. <https://doi.org/10.1145/3473586>

1 INTRODUCTION

Separation logic [Ishtiaq and O’Hearn 2001; Reynolds 2002] provides a powerful formalism for specifying an interface between a library and a client in terms of resources. For example, a client may obtain an “opened file” resource by calling an open operation of a file library, which it can then use to access the file by calling a read operation. Crucially, abstract predicates [Biering et al. 2007; Parkinson and Bierman 2005] hide implementation details: a client cannot rely on what an “opened file” resource actually consists of; it can only rely on the functionality for using the resource that the library provides. A client that is verified with respect to the abstract specification will behave correctly with any implementation of the library.

In separation logic, functions are specified with preconditions and postconditions. We can think of the precondition as specifying resources that a client must provide in order to call the function.

Authors’ addresses: Lars Birkedal, birkedal@cs.au.dk, Aarhus University, Denmark; Thomas Dinsdale-Young, tyoung@cs.au.dk, Concordium, Denmark; Armaël Guéneau, armael@cs.au.dk, Aarhus University, Denmark; Guilhem Jaber, guilhem.jaber@univ-nantes.fr, Université de Nantes, France; Kasper Svendsen, ksvendsen@cs.au.dk, Uber, Denmark; Nikos Tzevelekos, nikos.tzevelekos@qmul.ac.uk, Queen Mary University of London, United Kingdom.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/8-ART81

<https://doi.org/10.1145/3473586>

The postcondition then specifies the resources that the client receives when the function returns. For example, a simplified specification for a file library might look like:

$$\{\text{closed}\} \text{open}() \{\text{open}\} \quad \{\text{open}\} \text{close}() \{\text{closed}\} \quad \{\text{open}\} \text{read}() \{\text{open}\}$$

A client acquires an open resource, represented by an abstract predicate, by calling `open`. With this file resource, the client can call `close` and relinquish the resource, or call `read` and retain the resource. The rules of separation logic allow us to prove specifications for clients that use the library correctly, such as:

$$\{\text{closed}\} \text{open}(); \text{read}(); \text{close}() \{\text{closed}\}$$

On the other hand, we cannot prove any useful specifications¹ for programs that use the library incorrectly, like the following: `open(); close(); read()`.

Intuitively, separation logic specifications imply properties about the trace of interactions between a library and a client. For example, the specification for the file library ensures that a file can only be accessed if it has previously been opened and not subsequently closed. This strongly depends on the fact that the specification is abstract: a client has no way to obtain the open resource except by calling the `open` operation. (Indeed, if the client were able to forge the open resource then it could violate the trace property.) In other words, an abstract separation logic specification of a library entails a theorem ‘for free’; that is, a theorem that holds for *any* implementation of the library. While this intuition is broadly used in the separation logic community, it has not previously been formalised.

In this paper we present a formal approach to establishing free theorems from abstract separation logic specifications. In contrast to the well-known free theorems that one can obtain from polymorphic types [Reynolds and Plotkin 1993; Reynolds 1983; Wadler 1989], which are all concerned with equality properties (e.g., that every element of a polymorphic type is *equal to*² the polymorphic identity function), the free theorems we focus on here are all concerned with intensional trace properties describing how a library and a client interacts. Since the properties are intensional and not observable in the standard operational semantics, we establish the free theorems by placing a *wrapper* between a client and a library that generates a trace of the interactions between the client and library. We stress that the wrapper has no bearing on the underlying semantics of the program (when traces are ignored) but simply allows us to formally capture intensional trace properties. Then, supposing that a library implements an abstract separation logic specification, we show that the wrapped library also satisfies this specification and, moreover, maintains the desired trace property as an invariant. For the latter step, we use ghost trace resources, which can be used for instantiating the abstract predicates in the original specification. In the context of a client that uses the library in accordance with the specification, the trace properties are thus guaranteed to hold.

Our approach establishes *temporal* trace properties from separation logic specifications that are *not* inherently temporal, independently of implementation details. While ours is not the first approach incorporating temporal reasoning in program logics, and separation logic in particular, it is the first one to derive trace properties for libraries that have abstract separation logic specifications. Previous approaches [da Rocha Pinto et al. 2014; Fu et al. 2010; Gotsman et al. 2013; Sergey et al. 2015] achieve temporal reasoning by specifying and verifying the underlying libraries using trace-oriented specifications; the novelty of our work is in deriving temporal trace properties “for free” from trace-agnostic specifications.

¹It is always possible to prove a vacuous specification with precondition $\perp$, but a useful specification should at least have a satisfiable precondition. For complete, closed programs we would typically expect the precondition to be `emp` (denoting the empty heap), thus requiring the ownership of no specific resource.

²Be it in a denotational semantics model or contextual equivalence based on operational semantics.

We remark that informal English-language specifications (e.g., the POSIX file system [Gardner et al. 2014]) are often given in terms of trace properties, so the free theorems we obtain also serve to link formal separation logic specifications to more informal intuitive trace properties.

We demonstrate our approach on a variety of examples, which establish that separation logic specifications can imply a variety of free theorems about interaction traces, such as:

- An iterator over a collection should only be used if the collection has not been modified since the iterator was created (§5.2).
- A higher-order function calls its argument exactly once (§5.3).
- A traversal on a stack invokes a given method on each value that has been pushed (but not popped) in the order in which they will be popped (§5.4).

As a more advanced example of a free theorem, we further prove that a *logically atomic* concurrent separation logic specification of a concurrent module operation [da Rocha Pinto et al. 2014; Jung et al. 2015a] implies that the operation is linearizable. To the best of our knowledge, this is the first formal such result. Note that this is indeed a free theorem in the sense it holds for *any* operation satisfying a logically atomic specification, i.e., the proof only relies on the specification and not on the specific operation.

We first outline our methodology on a motivating example in Section 2. We then formally define the programming language that we consider (Section 3), and the program logic that allows us to reason about interaction traces (Section 4). Then, we describe in Section 5 how our approach can be applied on a number of examples, culminating to the formal proof in Section 6 that logical atomicity implies linearizability, as a free theorem.

This work has been entirely mechanized in the Coq proof assistant using the Iris framework. It is publically available at either <https://zenodo.org/record/5054898> (stable artifact) or <https://github.com/logsem/free-theorems-sl> (development repository).

2 MOTIVATING EXAMPLE

As a motivating example, consider a library that provides a stack with push and pop operations. In Separation Logic, these operations can be specified as follows:

$$\begin{aligned} & \forall \alpha, x. \{ \text{stack}(\alpha) * [x \neq ()] \} \text{push}(x) \{ \text{stack}(x :: \alpha) \} \\ & \forall \alpha. \left\{ \text{stack}(\alpha) \right\} \text{pop}() \left\{ r. ([r = () \wedge \alpha = \varepsilon] * \text{stack}(\alpha)) \vee \right. \\ & \quad \left. (\exists \alpha'. [\alpha = r :: \alpha'] * \text{stack}(\alpha')) \right\} \end{aligned}$$

The push operation simply prepends the given value to the stack, which is represented by the stack abstract predicate. The value must be distinct from the unit $()$, which is used to indicate an empty stack. The pop operation returns $()$ if the stack is empty and otherwise removes and returns the head value. (We use $[\cdot]$ to inject pure propositions into the world of Separation Logic assertions.) The library would also include a way to create an initially empty stack, but we elide that here (see the more detailed treatment in §5).

We can also specify the stack in terms of trace properties that are satisfied by interactions between a client and a stack library. For instance, a simple trace property that we might wish to show is this:

Each (non-unit) value returned by an invocation of pop was an argument of a previous invocation of push.

This property is not captured by most models of Separation Logic, in part simply because it is an intensional temporal trace property. We now outline how we can derive this property as a *free theorem*.

In order to capture trace properties, we define a wrapper that instruments the library operations with code to emit trace events annotated with push and pop labels:

$$\text{push}' \triangleq \lambda v. \text{push}(v); \text{emit}(\langle \text{push}, v \rangle) \quad \text{pop}' \triangleq \lambda _. \text{let } r = \text{pop}() \text{ in } \text{emit}(\langle \text{pop}, r \rangle); r$$

We can then define a language of traces that captures our desired invariant:

$$\mathcal{L} = \{t \mid \forall i, v. t[i] = \langle \text{pop}, v \rangle \wedge v \neq () \implies \exists j. j < i \wedge t[j] = \langle \text{push}, v \rangle\}$$

Our separation logic includes two resources that allow us to reason about code that emits trace events: $\text{trace}(t)$ expresses that t is the current trace; and $\text{tracelnv}(I)$ expresses that the trace invariant is I . The proof rule for emit updates the trace resource, while requiring that the new trace satisfies the invariant. Using these, we can define a ‘wrapped’ version of the abstract stack predicate so that the wrapped operations will satisfy the same abstract specification, but also enforce the trace invariant:

$$\text{stack}'(\alpha) \triangleq \text{stack}(\alpha) * \text{tracelnv}(\mathcal{L}) * \exists t. \text{trace}(t) * [t \in \mathcal{L} \wedge \forall x \in \alpha. \exists i. t[i] = \langle \text{push}, x \rangle]$$

Both the wrapped push and pop operations can be verified similarly. The proof of the wrapped push operation proceeds as follows:

$$\begin{aligned} & \{\text{stack}'(\alpha) \wedge v \neq ()\} \\ & \{\text{stack}(\alpha) \wedge v \neq () * \text{tracelnv}(\mathcal{L}) * \exists t. \text{trace}(t) * [t \in \mathcal{L} \wedge \forall x \in \alpha. \exists i. t[i] = \langle \text{push}, x \rangle]\} \\ & \quad \text{push}(v); \\ & \{\text{stack}(v :: \alpha) * \text{tracelnv}(\mathcal{L}) * \exists t. \text{trace}(t) * [t \in \mathcal{L} \wedge \forall x \in \alpha. \exists i. t[i] = \langle \text{push}, x \rangle]\} \\ & \{\text{stack}(v :: \alpha) * \text{tracelnv}(\mathcal{L}) * \exists t. \text{trace}(t) * [(t \cdot \langle \text{push}, v \rangle) \in \mathcal{L}] * \\ & \quad [\forall x \in (v :: \alpha). \exists i. (t \cdot \langle \text{push}, v \rangle)[i] = \langle \text{push}, x \rangle]\} \\ & \quad \text{emit}(\langle \text{push}, v \rangle); \\ & \{\text{stack}(v :: \alpha) * \text{tracelnv}(\mathcal{L}) * \exists t. \text{trace}(t) * [t \in \mathcal{L} \wedge \forall x \in \alpha. \exists i. t[i] = \langle \text{push}, x \rangle]\} \\ & \{\text{stack}'(v :: \alpha)\} \end{aligned}$$

Using an Adequacy theorem for our program logic, which links trace invariants to actual program executions (Theorem 4.1), we can thus conclude that the stack indeed satisfies the desired trace property. It is a free theorem because the proofs for the wrapped operations only depend on the specification for push and pop, not the actual code implementing the operations.

The above example demonstrates our technique on a first-order library (the library cannot make call-backs to the client), but it also applies in a higher-order setting. For instance, consider extending the stack with a `foreach` operation that traverses the stack and calls a client-supplied function on each element in order. In separation logic, this operation can be specified as:

$$\forall \alpha, f, I. \left\{ \text{stack}(\alpha) * I(\epsilon) * \forall \beta, x. \left\{ I(\beta) \right\} f(x) \left\{ I(x :: \beta) \right\} \right\} \text{foreach}(f) \left\{ \text{stack}(\alpha) * I(\text{rev}(\alpha)) \right\}$$

This specification is subtle. The `foreach` operation takes a function f that is specified (using a nested triple) with an invariant $I(\beta)$ whose parameter records the list of values that f has so far been called on. The operation is given the predicate $I(\epsilon)$ initially, and it returns with $I(\text{rev}(\alpha))$, where $\text{rev}(\alpha)$ is the reversal of the list α . Note that while the predicate stack is abstract to the client (the library determines the interpretation), the predicate I is abstract to the library (the client determines the interpretation). This means that for the library to obtain the $I(\text{rev}(\alpha))$ predicate for the postcondition from the $I(\epsilon)$ given in the precondition, it must call f on each element of α in order. Moreover, it cannot make any further calls to f , since separation logic treats the predicate $I(\beta)$ as a resource, which must be given up by the library each time it calls f , and which the library has no means of duplicating since in separation logic $A \implies A * A$ does not hold in general.

Defining a wrapper for higher-order libraries is more complex than in the first-order case, since we wish the trace to capture all interactions between the client and the library, including call-backs

between them. The wrapper must therefore emit events at the call and return of library functions, as well as functions that are passed as arguments to library functions. We can then specify a number of properties that traces generated by client-library interactions will have:

- The trace of push and pop operations obeys the stack discipline.
- Between invocation and return, $\text{foreach}(f)$ calls f on each element of the stack in order, with no further calls.
- The invocations of push, pop and f (the argument of a call of foreach) are atomic — there are no further events between the call and return.

While the first property can be seen as a straightforward consequence of the push and pop specifications, the others are more subtle. In particular, they depend on the foreach specification's parametricity in I , which prevents the library from using its argument in an arbitrary fashion. For particular instantiations of I , (e.g. $I(\beta) = \text{true}$), foreach could call its argument an arbitrary number of times, or even store it and call it from future invocations of pop . However, parametricity ensures that it will behave the same independent of how I is instantiated, and so it cannot do that since I can be instantiated so as to enforce trace properties.

This example illustrates several important aspects of our approach. Firstly, we support higher-order functions, such as foreach . We also deal with expressive trace properties: the language of traces is not context-free, since the stack may be traversed multiple times. Moreover, the connection between the separation logic specification and the trace properties that follow from it is subtle.

In §5 we revisit this example, among others, in detail. Before doing so, we present the formal setting of our approach.

3 PROGRAMMING LANGUAGE

We consider a fairly standard higher-order concurrent imperative language. We define the language in Figure 1: it is an untyped call-by-value λ -calculus with references, products and fork-based concurrency.

Additionally, our language features an emit primitive, which is used to output values as trace events, and a fresh primitive, which generates a fresh event tag and emits an event at the same time. These operations are used to implement the wrappers that instrument the operations of a given library so as to materialize the events of interest. We will formulate properties in terms of the event traces produced during program execution by these two primitives.

The definition of values (Val), expressions (Exp) and *evaluation contexts* (Ectx) is mostly standard. Most notable are the addition of the two primitives emit and fresh , and tags values τ , which are simply represented as strings. Tags are used when defining wrappers so as to label emitted events and associate them with the corresponding function of a library interface. The fresh primitive can be used to emit an event annotated with a fresh tag, which is guaranteed not to already appear in the current trace. In simpler cases, the code of wrappers will simply refer to tag literals (i.e. strings) to decorate events emitted using emit . We will write mytag for a tag represented by the string “mytag”. Importantly, tags can be passed around but cannot be compared, so they have no impact on the control flow of the program.

We equip the language with a small-step operational semantics. The head reduction judgment $e; \sigma \rightarrow_h e'; \sigma'; \vec{e}_f$ describes a single step of computation as can be performed by an individual thread. It expresses that, starting from state σ , the expression e may reduce in one step to expression e' , while changing the state to σ' and forking off a list of threads running expressions $\vec{e}_f$. A state σ corresponds to a pair of a memory heap h (a finite map from locations to values) and the current trace t (a list of values). The threadpool reduction judgment $T; \sigma \longrightarrow T'; \sigma'$ lifts head reduction to lists of concurrently running threads.

$$\begin{aligned}
& \oplus ::= + \mid - \mid \times \mid = \mid < \\
& v \in \text{Val} ::= \lambda x. e \mid \langle v, v \rangle \mid () \mid n \mid b \mid \ell \mid \tau \\
& e \in \text{Exp} ::= v \mid x \mid \lambda x. e \mid e \mid e \oplus e \mid \text{if } e \text{ then } e \text{ else } e \mid \langle e, e \rangle \mid \text{fst } e \mid \text{snd } e \\
& \quad \mid \text{ref } e \mid !e \mid e \leftarrow e \mid \text{fork } e \mid \text{CAS}(e, e, e) \mid \text{emit } e \mid \text{fresh } e \\
& K \in \text{Ectx} ::= (\text{standard left-to-right evaluation contexts}) \\
& n \in \mathbb{N}, b \in \mathbb{B}, \ell \in \text{Loc}, \tau \in \text{Tag} \triangleq \text{String} \\
\\
& h \in \text{Heap} \triangleq \text{Loc} \rightarrow_{\text{fin}} \text{Val} \qquad t \in \text{Trace} \triangleq \text{Val}^* \\
& \sigma \in \text{State} \triangleq \text{Heap} \times \text{Trace} \qquad T \in \text{Tpool} \triangleq \text{Exp}^* \\
& \text{tags}(\varepsilon) \triangleq \varepsilon \qquad \text{tags}(t \cdot \langle \tau, v \rangle) = \text{tags}(t) \cdot \tau \qquad \text{tags}(t \cdot v) = \text{tags}(t) \text{ otherwise} \\
\\
& \boxed{e; \sigma \rightarrow_h e'; \sigma'; \vec{e}_f} \qquad \text{Head reduction (selected rules)} \\
\\
& (\lambda x. e) v; \sigma \rightarrow_h e[v/x]; \sigma; \varepsilon \qquad \text{fork } e; \sigma \rightarrow_h (); \sigma; e \\
\\
& \text{ref } v; (h, t) \rightarrow_h \ell; (h \uplus [l \mapsto v], t); \varepsilon \qquad \frac{\sigma.h(\ell) = v}{! \ell; \sigma \rightarrow_h v; \sigma; \varepsilon} \qquad \frac{h(\ell) = v}{\ell \leftarrow v'; (h, t) \rightarrow_h (); (h[\ell \mapsto v'], t); \varepsilon} \\
\\
& \frac{\sigma.h(\ell) \neq v_1}{\text{CAS}(\ell, v_1, v_2); \sigma \rightarrow_h \text{false}; \sigma; \varepsilon} \qquad \text{CAS}(\ell, v_1, v_2); (h \uplus [\ell \mapsto v_1], t) \rightarrow_h \text{true}; (h \uplus [\ell \mapsto v_2], t); \varepsilon \\
\\
& \text{emit } v; (h, t) \rightarrow_h (); (h, t \cdot v); \varepsilon \qquad \frac{\tau \notin \text{tags}(t)}{\text{fresh } v; (h, t) \rightarrow_h \tau; (h, t \cdot \langle \tau, v \rangle); \varepsilon} \\
\\
& \boxed{T; \sigma \longrightarrow T'; \sigma'} \qquad \text{Threadpool reduction} \\
\\
& \frac{e; \sigma \rightarrow_h e'; \sigma'; \vec{e}_f}{(T_1 \uparrow\uparrow K[e] \uparrow\uparrow T_2); \sigma \longrightarrow (T_1 \uparrow\uparrow K[e'] \uparrow\uparrow T_2 \uparrow\uparrow \vec{e}_f); \sigma'}
\end{aligned}$$

Fig. 1. Syntax and semantics of the language.

Our language includes a compare-and-swap (CAS) primitive that can be used to implement locks or other concurrency primitives. Note that it is not terribly important what the exact set of features of our language is, as long as it is rich enough to allow for interesting implementations of the libraries and their clients for the specifications that we consider. In the rest of the paper, the code that we will directly verify is the code of wrappers, which only rely on fresh and emit, but it is important that we consider a language featureful enough for the whole setup to be realistic.

4 PROGRAM LOGIC

We use a program logic built as an instantiation of the Iris framework. For the purposes of this paper we present a particular instance of this framework for the language introduced in §3 and we refer to this instance simply as Iris.

4.1 Iris Basics

Figure 2 describes the subset of Iris that is relevant for this paper. We write $iProp$ for the universe of Iris propositions. These include the standard connectors of higher-order logic and separation logic, including the separating conjunction $*$ and the magic wand $\multimap$ (also known as separating

$P, Q \in iProp ::=$	True False $\forall x. P$ $\exists x. P$...	higher-order logic
	$P * Q$ $P \multimap Q$ $\lceil \phi \rceil$ $\Box P$	basic separation logic
	$\{P\} e \{v.Q\}$ $\langle P \rangle e \langle v.Q \rangle$	program logic
	$\boxed{P}$ $\boxed{m}^\gamma$	invariants and ghost state
	trace(t) hist(t) traceInv($\mathcal{L}$)	trace-related resources

Fig. 2. The subset of Iris relevant for this paper.

implication). The proposition $\lceil \phi \rceil$ asserts that the pure proposition ϕ holds, where ϕ is a proposition from the meta logic.

The modality $\Box$ is used to assert persistence of assertions. Iris assertions can be divided in two categories: *ephemeral* assertions and *persistent* assertions. Ephemeral assertions describe facts or resources that are available at a given point but might become false or unavailable later. Persistent assertions describe facts that never cease to be true. The assertion $\Box P$ then holds if both P holds and P is persistent. Formally, we say that an assertion P is persistent, written $\text{persistent}(P)$, if $P \dashv\vdash \Box P$, where $\dashv\vdash$ is the logical equivalence of Iris propositions. As the knowledge associated with a persistent assertion can never be invalidated, persistent assertions can be freely duplicated.

Iris allows specifying programs: a Hoare triple $\{P\} e \{v.Q\}$ asserts that, when the precondition P holds, then e is safe to execute, and if e reduces to a value v , then the postcondition Q holds (note that v acts as a binder in Q). Logically atomic triples $\langle P \rangle e \langle v.Q \rangle$ are used to specify concurrent operations, and are explained in more details in Section 6.1.

The assertion $\boxed{P}$ is an Iris invariant: it asserts that P should hold at all times, and is therefore a persistent assertion. Iris also provides a way of tying the evolution of pure data m to the execution of a program, through the use of ghost state, described by propositions $\boxed{m}^\gamma$ (where γ is a logical name identifying a piece of ghost state). We will only make use of invariants and (indirectly) ghost state in Section 6.

4.2 Trace Primitives

Our instantiation of Iris is augmented with three (definable) assertions, $\text{trace}(t)$, $\text{hist}(t)$ and $\text{traceInv}(\mathcal{L})$, for reasoning about the trace of events emitted during execution.

The resource $\text{trace}(t)$ expresses that the trace of events emitted so far is exactly t and asserts exclusive right to emit further trace events. Since $\text{trace}(t)$ asserts exclusive right to emit events, it cannot be duplicated, and allows a single owner to reason precisely about the current trace.

The $\text{trace}(t)$ resource suffices for examples where only a single resource needs to refer to the current trace. To improve expressiveness, we use the $\text{hist}(t)$ resource to reason about prefixes of the current trace. The resource $\text{hist}(t)$ asserts that the trace t is a prefix of the trace of events emitted so far. Note that this property is preserved by emission of new events: if t is a prefix of the current trace then t is also a prefix of any extension of the current trace. This resource is persistent (and thus duplicable), and given the $\text{trace}(t)$ resource we can construct a $\text{hist}(t)$. Moreover, if we have a $\text{trace}(t_1)$ resource and a $\text{hist}(t_2)$ resource, we can conclude that t_2 is a prefix of t_1 . These properties are given formally at the top of Figure 3.

Finally, the $\text{traceInv}(\mathcal{L})$ resource defines a trace invariant that everyone must obey. Here $\mathcal{L}$ is a set of traces and $\text{traceInv}(\mathcal{L})$ asserts that, at every point of the execution, the current trace must belong to $\mathcal{L}$. Since it specifies an invariant, the $\text{traceInv}(\mathcal{L})$ resource is duplicable (see Figure 3). (Internally, traceInv is defined as an Iris invariant of the form $\boxed{\dots}$, from which we prove as derived properties the simpler reasoning rules shown in Figure 3.)

$$\begin{array}{c}
\text{tracelnv}(\mathcal{L}) \vdash \Box \text{tracelnv}(\mathcal{L}) \qquad \text{hist}(t) \vdash \Box \text{hist}(t) \qquad \text{trace}(t) \vdash \text{trace}(t) * \text{hist}(t) \\
\text{trace}(t_1) * \text{hist}(t_2) \vdash \text{trace}(t_1) * [t_2 \leq_{\text{pref}} t_1] \qquad \text{trace}(t) * \text{tracelnv}(\mathcal{L}) \vdash \text{trace}(t) * [t \in \mathcal{L}] \\
\\
\frac{t \cdot v \in \mathcal{L}}{\text{tracelnv}(\mathcal{L}) \vdash \{\text{trace}(t)\} \text{emit } v \{\text{trace}(t \cdot v)\}} \\
\\
\frac{\forall r. r \notin \text{tags}(t) \Rightarrow t \cdot \langle r, v \rangle \in \mathcal{L}}{\text{tracelnv}(\mathcal{L}) \vdash \{\text{trace}(t)\} \text{fresh } v \{r. [r \notin \text{tags}(t)] * \text{trace}(t \cdot \langle r, v \rangle)\}}
\end{array}$$

Fig. 3. Properties of the resources trace, hist and tracelnv, and specifications for emit and fresh.

We can now give Separation Logic specifications for the emit and fresh operations, which are shown in Figure 3. Both specifications require ownership over the trace resource $\text{trace}(t)$, and that the trace is governed by some trace invariant $\mathcal{L}$. Given ownership of the trace resource we can emit an event v which is appended at the end of the trace. Emitting a fresh event v returns a tag r that does not appear in the current trace, and emits the event $\langle r, v \rangle$.

As a side-condition, the updated trace is required to satisfy the invariant $\mathcal{L}$ (in the case of fresh, it is required to satisfy $\mathcal{L}$ for any fresh tag, since we cannot know in advance the tag that will be returned by fresh). Finally, both specifications hand out an updated trace resource in the postcondition, with the emitted event appended at the end.

We emphasize that the predicates and the properties in Figure 3 are all defined and proved formally in Iris, see the accompanying Coq formalization for details.

4.3 Adequacy Theorem

Ultimately, a program logic is only a technical device: the end goal is to establish properties that can be shown to hold with respect to the operational semantics of the language, independently of the details of the logic. This is typically achieved by proving an “adequacy” or “soundness” theorem of the logic. We present one such theorem, Theorem 4.1, tailored to our goal of establishing trace properties for the interaction between a verified library and its client.

THEOREM 4.1 (ADEQUACY). *Given any specification Φ , resource P_0 , initialiser e_{init} , library implementation v_{lib} , memory heap h , client e , and language $\mathcal{L}$, if the following conditions hold:*

- $\Phi : i\text{Prop} \times \text{Val} \rightarrow i\text{Prop}$ and $P_0 : i\text{Prop}$
- $\varepsilon \in \mathcal{L}$
- $\forall \text{lib}, P. \Phi(P, \text{lib}) \vdash \{P\} e \text{ lib } \{\top\}$
- $\vdash \{\top\} e_{\text{init}} \{P_0\}$
- $\vdash \Phi(P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}), v_{\text{lib}})$

then for any T, h', t , provided $(e_{\text{init}}; e v_{\text{lib}}), (h, \varepsilon) \longrightarrow^ T, (h', t)$ then we have $t \in \mathcal{L}$.*

To understand the theorem, first recall that the general pattern we follow is to prove that for any library implementation satisfying an abstract library specification, the wrapped library implementation satisfies the same abstract specification and moreover the traces generated by the wrapping satisfy a given invariant. Client programs verified against the abstract library interface can thus be linked with the wrapped library implementation to conclude that the traces generated by the wrapping satisfy the given invariant. Theorem 4.1 formalises this idea. Here e_{init} is an initialiser operation to initialize the internal state P_0 of the library. The theorem requires that we

provide a library implementation satisfying the abstract specification Φ and generating traces in $\mathcal{L}$. It then guarantees that, for any client e verified against the abstract specification of the library, the trace produced by running the client linked with the library implementation v_{lib} satisfies the trace invariant $\mathcal{L}$ at every step of the execution.

It is worth noting that the assumptions of Theorem 4.1 rules out clients containing calls to `emit` or `fresh`, which could in principle invalidate the trace property. Indeed, following assumption 3 of the theorem, we only consider *verified* clients e that can be proved safe assuming an arbitrary proposition P as pre-condition and the specification of the library $\Phi(P, lib)$. Neither of these propositions can give access to the `trace(·)` and `tracelnv(·)` predicates necessary to verify calls to `emit` and `fresh` (we have to establish Φ in assumption 5 which prevents it from providing these predicates for any P).

To use Theorem 4.1 to derive a trace property from a separation logic specification, the idea is to define a suitable wrapper function `wrap` for the library in question and prove that if a library implementation M satisfies Φ then so does the wrapped version `wrap(M)` and, additionally, the wrapped version generates traces in the $\mathcal{L}$ language:

$$\forall P_0, M. \Phi(P_0, M) \vdash \Phi(P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}), \text{wrap}(M)). \quad (1)$$

One can then apply Theorem 4.1 by taking v_{lib} to be `wrap(M)`. For each of the examples coming up next, we will establish a lemma of the form (1). We stress that each of these results is not just a fact that holds in our logic: each one can be composed with Theorem 4.1 to obtain that the trace property of interest is preserved *at the level of the operational semantics*. (We refer to the Coq formalization for the details of each specific instantiation.) We also emphasize that (1) is shown for any M and thus the trace property only depends on the specification Φ and in that sense we obtain a free theorem.

5 PROVING TRACE PROPERTIES AS FREE THEOREMS

In this section, we demonstrate the approach sketched earlier through a series of increasingly complex examples, starting with the basic file library example from the introduction.

It is worth noting that, even though we work in a concurrent language, the examples presented in this section are sequential in nature: the specifications considered prevent concurrent calls to the underlying library. We will look more closely into the case of concurrent libraries in Section 6.

5.1 File Library

Recall the Separation Logic specification of the file library from the introduction, here written more formally, where P_0 denotes the resources needed to initialize the library.

$$\Phi_{\text{file}}(P_0, (\text{open}, \text{close}, \text{read})) \triangleq \exists \text{open}, \text{closed} : iProp. \Box(P_0 \multimap \text{closed}) \wedge \\ \{\text{closed}\} \text{open}() \{\text{open}\} \wedge \{\text{open}\} \text{close}() \{\text{closed}\} \wedge \{\text{open}\} \text{read}() \{\text{open}\}$$

Note the use of different fonts: *open*, *close* and *read* denote functions (values of the language), while `open` and `closed` denote abstract Separation Logic predicates used in the specification of the functions. Below we will also use the `open`, `close` and `read` tags (recall that they correspond to mere strings) for writing the wrapper code.

We wish to prove that this Separation Logic specification enforces that verified clients can only `close` and `read` when the file is `open`. To capture this property, we define a wrapper function `wrapfile` that instruments an implementation of the file library to emit events annotating calls to the different operations with respective tags `open`, `close` and `read`. Formally, we take a file library to be a triple

consisting of an *open*, a *close* and a *read* function.

$$\begin{aligned} \mathbf{wrap}_{\text{file}} &\triangleq \lambda(\text{open}, \text{close}, \text{read}). \\ &(\lambda_{-}. \text{open}(); \text{emit open}, \lambda_{-}. \text{close}(); \text{emit close}, \lambda_{-}. \text{read}(); \text{emit read}) \end{aligned}$$

We can now formalize the protocol as constraints on the traces generated by linking an instrumented file library with a client. In particular, we consider traces belonging to the language $\mathcal{L}_{\text{file}}$ of all strings $t \in \Sigma^*$ such that t is a valid file trace, $\text{file}_{\text{trace}}(t)$, where $\Sigma = \{\text{open}, \text{close}, \text{read}\}$ and $\text{file}_{\text{trace}}$ is defined as:

$$\begin{aligned} \text{file}_{\text{trace}}(t) &\triangleq \forall n. t[n] = \text{read} \vee t[n] = \text{close} \implies \text{isopen}(t, n) \\ \text{isopen}(t, n) &\triangleq \exists m < n. t[m] = \text{open} \wedge \text{noclose}(t, m, n) \\ \text{noclose}(t, m, n) &\triangleq \forall k. m < k < n \implies t[k] \neq \text{close} \end{aligned}$$

The valid file trace predicate, $\text{file}_{\text{trace}}(t)$, expresses that the trace t only contains read and close events when the file is open. To prove that the specification enforces the intended protocol, we proceed by proving that the wrapping preserves satisfaction of the specification and generates traces in $\mathcal{L}_{\text{file}}$.

LEMMA 5.1. $\forall P_0, \text{ops}. \Phi_{\text{file}}(P_0, \text{ops}) \vdash \Phi_{\text{file}}(P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{file}}), \mathbf{wrap}_{\text{file}}(\text{ops})).$

PROOF SKETCH. We first need to define new wrapped versions of the abstract representation predicates, which relate the abstract resources to the current trace. The idea is that the wrapped open resource, open_w , should express that the current trace t is in $\mathcal{L}_{\text{file}}$ and that the trace t is open. Likewise, the wrapped closed resource, closed_w can simply express that the current trace is in $\mathcal{L}_{\text{file}}$.

$$\begin{aligned} \text{open}_w &\triangleq \text{open} * \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{file}}) * [\text{isopen}(t, |t|)] \\ \text{closed}_w &\triangleq \text{closed} * \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{file}}) \end{aligned}$$

We need to prove $P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{file}}) \multimap \text{closed}_w$ assuming $P_0 \multimap \text{closed}$, which follows trivially from the definition of closed_w . It remains to be proved that the wrapped methods satisfy their specifications.

We give on the right a proof outline for the wrapped open method. Here we use the assumed specification of the underlying open method to verify the call to open, and we used the following property of $\mathcal{L}_{\text{file}}$ and isopen to prove that emitting open would result in a trace in $\mathcal{L}_{\text{file}}$ that was open (recall that $\text{trace}(t)$ combined with $\text{tracelnv}(\mathcal{L}_{\text{file}})$ entails $t \in \mathcal{L}_{\text{file}}$):

$$\begin{aligned} &\{\text{closed}_w\} \\ &\{\text{closed} * \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{file}})\} \\ &\quad \text{open}(); \\ &\{\text{open} * \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{file}})\} \\ &\quad \mathbf{emit\ open}; \\ &\{\text{open} * \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{file}}) * \\ &\quad [\text{isopen}(t, |t|)]\} \\ &\{\text{open}_w\} \end{aligned}$$

$$\forall t \in \mathcal{L}_{\text{file}}. t \cdot \text{open} \in \mathcal{L}_{\text{file}} \wedge \text{isopen}(t \cdot \text{open}, |t \cdot \text{open}|).$$

The proofs for the close and read operations are similar, but rely on the following property to justify the emits:

$$\forall t \in \mathcal{L}_{\text{file}}. \text{isopen}(t, |t|) \implies t \cdot \text{close} \in \mathcal{L}_{\text{file}} \wedge t \cdot \text{read} \in \mathcal{L}_{\text{file}}.$$

By combining Lemma 5.1 with the Adequacy theorem (Theorem 4.1) we have thus shown the free theorem that any verified client can only close and read when the file is open.

5.2 Iterators on Collections

We consider a collections library that provides methods for modifying a collection as well as iterating over it. The combination of these two features means that we have to be wary of *iterator invalidation*: it is typically unsafe to use an iterator after the collection it is iterating over has been modified. Following Krishnaswami et al. [2009], we assume for such a library a cut-down Separation Logic library specification that does not track the contents of the underlying collections, and consists of five operations. The *size* operation is non-destructive and returns the size of the given collection; the *add* and *remove* operations destructively modify the given collection by adding or removing an element; and finally, *iterator* returns a new iterator for the collection, while *next* returns the next element of a given iterator.

$$\begin{aligned}
 \Phi_{\text{coll}}(P_0, (\text{size}, \text{add}, \text{remove}, \text{iterator}, \text{next})) &\triangleq \\
 \exists \text{coll} : \text{Val} \rightarrow i\text{Prop}. \exists \text{iter} : \text{Val} \times \text{Val} \rightarrow i\text{Prop}. \Box(P_0 \multimap \exists c : \text{Val}. \text{coll}(c)) \wedge \\
 \forall c : \text{Val}. \{\text{coll}(c)\} \text{size}() \{x. \text{coll}(c)\} \wedge \\
 \forall c, x : \text{Val}. \{\text{coll}(c)\} \text{add}(x) \{\exists c' : \text{Val}. \text{coll}(c')\} \wedge \\
 \forall c, x : \text{Val}. \{\text{coll}(c)\} \text{remove}(x) \{\exists c' : \text{Val}. \text{coll}(c')\} \wedge \\
 \forall c : \text{Val}. \{\text{coll}(c)\} \text{iterator}() \{r. \text{coll}(c) * \text{iter}(r, c)\} \wedge \\
 \forall c, x : \text{Val}. \{\text{coll}(c) * \text{iter}(x, c)\} \text{next}(x) \{\text{coll}(c) * \text{iter}(x, c)\}
 \end{aligned}$$

A seasoned practitioner of Separation Logic will be able to read from this specification that it successfully prevents the pitfall of iterator invalidation, i.e. it prevents the use of iterators that have been invalidated by a modification to the collection. This is not completely obvious, and comes from the combination of three observations. Firstly, the $\text{coll}(c)$ resource, which represents the ownership over the collection, is indexed by a value c which can be seen as an abstract version number. Secondly, destructive operations such as *add* or *remove* update the version number to a new *existentially quantified* version number, which can only be assumed to be distinct from all the previous versions of the collection. Finally, the $\text{iter}(p, c)$ resource expresses that p is an iterator over a collection which was at version c when the iterator was created; the specification of *next* then requires the ownership of both an iterator and its collection with matching version numbers.

In other words, by simply reading the specification we can deduce the following free theorem:

An iterator over a collection can only be used if the underlying collection has not been destructively modified since the iterator was created.

This is a trace property of the interaction between the collection library and clients. To capture it formally, we define as before a wrapper for the library that produces appropriate trace events. The instrumentation is fairly straightforward and simply emits a suitable event indicating the operation called and the argument and/or return value of the given operation, when relevant:

$$\begin{aligned}
 \text{wrap}_{\text{coll}} &\triangleq \lambda(\text{size}, \text{add}, \text{remove}, \text{iter}, \text{next}). \\
 &\left(\begin{array}{l} \lambda y. \text{let } r = \text{size}(y) \text{ in emit size}; r, \\ \lambda y. \text{add}(y); \text{emit add}, \\ \lambda y. \text{remove}(y); \text{emit remove}, \\ \lambda _. \text{let } r = \text{iter}() \text{ in emit}\langle \text{iterator}, r \rangle; r, \\ \lambda y. \text{let } r = \text{next}(y) \text{ in emit}\langle \text{next}, y \rangle; r \end{array} \right)
 \end{aligned}$$

The traces ignore the arguments to *add* and *remove* and the return values of *size* and *next*, as they are irrelevant for the protocol.

With this instrumentation we can now express the informal protocol as a language of permissible interaction traces between the client and library. We let the trace alphabet be the countable set:

$$\Sigma = \{\text{size}, \text{add}, \text{remove}\} \cup \{\langle \text{next}, \ell \rangle, \langle \text{iterator}, \ell \rangle \mid \ell \in \text{Loc}\}$$

The language $\mathcal{L}_{\text{coll}}$ of safe behaviours contains all strings $t \in \Sigma^*$ such that, for all $0 \leq i < |t|$:

if $t[i] = \langle \text{next}, \ell \rangle$ then there is $j < i$ such that $t[j] = \langle \text{iterator}, \ell \rangle$ and, for all $j < k < i$, $t[k] \notin \{\text{add}, \text{remove}\}$.

That is, every call to the *next* method of an iterator ℓ must be preceded by a call to *iterator*() which returns ℓ . In addition, there should be no modification of the collection between those two events.

We now aim to prove that, for an arbitrary library implementation M that satisfies the collections specification, the wrapped library implementation $\text{wrap}_{\text{coll}}(M)$ also satisfies the collections specification *and* the traces generated by the wrapped implementation are in the language $\mathcal{L}_{\text{coll}}$.

LEMMA 5.2. $\forall P_0, ops. \Phi_{\text{coll}}(P_0, ops) \vdash \Phi_{\text{coll}}(P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{coll}}), \text{wrap}_{\text{coll}}(ops))$.

PROOF SKETCH. We first need to define new wrapped versions of the coll and iter resources that relate the abstract version number to the trace state. The key idea is that the collection parameter of the wrapped coll resource will consist of a pair (c, n) , where the c component is the parameter of the underlying coll resource and n is the index in the trace of the last add or remove event. We want the wrapped coll((c, n)) resource to assert that there are no add or remove events in the current trace after the n -th element. Likewise, the wrapped iter($r, (c, n)$) resource should assert that there is an iterator event for iterator r in the current trace after the n -th element. Hence, if we own both coll((c, n)) and iter($r, (c, n)$) then we know that no add or remove events were emitted since the iterator r was created.

Let coll and iter denote the non-wrapped representation predicates that exist by the $\Phi_{\text{coll}}(P_0, ops)$ assumption and define the wrapped representation predicates coll_w and iter_w as follows:

$$\begin{aligned} \text{coll}_w(x) &\triangleq \exists y : \text{Val}. \exists n : \mathbb{N}. [x = (y, n)] * \text{coll}(y) * \\ &\quad \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] \\ \text{iter}_w(r, x) &\triangleq \exists y : \text{Val}. \exists n : \mathbb{N}. [x = (y, n)] * \text{iter}(y) * \exists t. \text{hist}(t) * [\langle \text{iterator}, r \rangle \in t[n..]] \end{aligned}$$

where we use $t[n..]$ as notation for the subtrace of t starting from the n -th element.

It thus remains to show that the wrapped library satisfies the collections specification. First, we need to prove that we obtain a wrapped collection resource from the initial resources: $P_0 * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * \text{trace}(\varepsilon) \multimap \exists c : \text{Val}. \text{coll}_w(c)$. This follows easily from the $P_0 \multimap \exists c' : \text{Val}. \text{coll}(c')$ assumption by taking the second component of c to be 0.

Next, we have to show that each of the wrapped operations satisfies the corresponding Hoare triple. The *size* method is particularly simple, as any trace $t \in \mathcal{L}_{\text{coll}}$ can trivially be extended with a size event $t \cdot \text{size} \in \mathcal{L}_{\text{coll}}$. We will thus skip the *size* method. The *add* method is more interesting, as we have to update the index into the trace for the last *add* event. Below is a proof outline for the wrapped *add* method applied to argument z .

$$\begin{aligned} &\{\text{coll}_w(c)\} \\ &\{\exists x, n, t. [c = (x, n)] * \text{coll}(x) * \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|]\} \\ &\quad \text{add}(z); \\ &\{\exists x', n, t. \text{coll}(x') * \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|]\} \\ &\quad \text{emit add}; \\ &\{\exists x', n, t. \text{coll}(x') * \text{trace}(t \cdot \text{add}) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|]\} \\ &\{\exists c' : \text{Val}. \text{coll}_w(c')\} \end{aligned}$$

This leaves us with two proof obligations: firstly, we have to show that we are allowed to emit the add event (i.e., that $t \cdot \text{add} \in \mathcal{L}_{\text{coll}}$); and secondly for the last step we have to show that:

$$\forall x', n, t. \left(\text{coll}(x') * \text{trace}(t \cdot \text{add}) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) \right) \multimap \exists c' : \text{Val}. \text{coll}_w(c')$$

This follows easily by taking c' to be $(x', |t \cdot \text{add}|)$, as $t \cdot \text{add}$ contains no add or remove events after the $|t \cdot \text{add}|$ -th element. The proof for *remove* follows the same structure as for *add*.

For the *iterator* method, we emit an iterator event and create a new hist resource to record the trace at the time of the creation of the iterator. On the right, we give a proof outline for the *iterator* method. We are left with two proof obligations: $t \cdot \langle \text{iterator}, r \rangle \in \mathcal{L}_{\text{coll}}$, and:

$$\begin{array}{l} \forall x, n, t. \\ \quad ([c = (x, n)] * \text{coll}(x) * \text{iter}(r, x) * \\ \quad \text{trace}(t \cdot \langle \text{iterator}, r \rangle) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * \\ \quad [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|]) \\ \multimap \text{coll}_w(c) * \text{iter}_w(r, c) \end{array} \quad \begin{array}{l} \{ \text{coll}_w(c) \} \\ \{ \exists x, n, t. [c = (x, n)] * \text{coll}(x) * \text{trace}(t) * \\ \quad \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] \} \\ \quad \text{let } r = \text{iterator}() \text{ in} \\ \{ \exists x, n, t. [c = (x, n)] * \text{coll}(x) * \text{iter}(r, x) * \text{trace}(t) * \\ \quad \text{tracelnv}(\mathcal{L}_{\text{coll}}) * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] \} \\ \quad \text{emit}(\text{iterator}, r); \\ \{ \exists x, n, t. [c = (x, n)] * \text{coll}(x) * \text{iter}(r, x) * \\ \quad \text{trace}(t \cdot \langle \text{iterator}, r \rangle) * \text{tracelnv}(\mathcal{L}_{\text{coll}}) * \\ \quad [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] \} \\ \{ \text{coll}_w(c) * \text{iter}_w(r, c) \} \\ \quad r \\ \{ r. \text{coll}_w(c) * \text{iter}_w(r, c) \} \end{array}$$

To discharge this last proof obligation, we use the rule from Figure 3 that allows us to introduce a history resource $\text{hist}(t \cdot \langle \text{iterator}, r \rangle)$ and since $n \leq |t|$ it follows that $\langle \text{iterator}, r \rangle \in (t \cdot \langle \text{iterator}, r \rangle)[n..]$, as required by $\text{iter}_w(r, c)$.

We are left with next, which is the most interesting case as it requires us to prove the iterator we are trying to use is still valid. We give a proof outline for the next method applied to a value x :

$$\begin{array}{l} \{ \text{coll}_w(c) * \text{iter}_w(x, c) \} \\ \{ \exists y, n, t, t'. [c = (y, n)] * \text{coll}(y) * \text{iter}(x, c) * \text{trace}(t) * \text{hist}(t') * \text{tracelnv}(\mathcal{L}_{\text{coll}}) \\ \quad * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] * [\langle \text{iterator}, x \rangle \in t'[n..]] \} \\ \quad \text{let } r = \text{next}(x); \\ \{ \exists y, n, t, t'. [c = (y, n)] * \text{coll}(y) * \text{iter}(x, c) * \text{trace}(t) * \text{hist}(t') * \text{tracelnv}(\mathcal{L}_{\text{coll}}) \\ \quad * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] * [\langle \text{iterator}, x \rangle \in t'[n..]] \} \\ \quad \text{emit}(\text{next}, x); \\ \{ \exists y, n, t, t'. [c = (y, n)] * \text{coll}(y) * \text{iter}(x, c) * \text{trace}(t \cdot \langle \text{next}, x \rangle) * \text{hist}(t') * \text{tracelnv}(\mathcal{L}_{\text{coll}}) \\ \quad * [\text{add}, \text{remove} \notin t[n..] \wedge n \leq |t|] * [\langle \text{iterator}, x \rangle \in t'[n..]] \} \\ \{ \text{coll}_w(c) * \text{iter}_w(x, c) \} \\ \quad r \\ \{ \text{coll}_w(c) * \text{iter}_w(x, c) \} \end{array}$$

To verify the emit expression, we further have to prove that $t \cdot \langle \text{next}, x \rangle \in \mathcal{L}_{\text{coll}}$, that is, that the iterator x is still valid. This relies on the following key property of the $\mathcal{L}_{\text{coll}}$ language:

$$t \in \mathcal{L}_{\text{coll}} \wedge \text{add}, \text{remove} \notin t[n..] \wedge \langle \text{iterator}, x \rangle \in t[n..] \implies t \cdot \langle \text{next}, x \rangle \in \mathcal{L}_{\text{coll}}$$

To apply this property we use the adequate rule from Figure 3 to conclude from $\text{trace}(t) * \text{hist}(t')$ that $t' \leq_{\text{pref}} t$ and thus that $\langle \text{iterator}, x \rangle \in t'[n..] \implies \langle \text{iterator}, x \rangle \in t[n..]$.

5.3 Well-bracketing Protocols

Libraries that allow clients to acquire, access and release resources often impose a well-bracketing protocol whereby clients are required to acquire resources before accessing and releasing them. The file library in §5.1 was a particularly simple example of such a protocol. In this section we

consider a more advanced and realistic variant thereof for a library with a higher-order function that takes care of acquiring and releasing the underlying resource for clients.

Consider a library with a higher-order method *withRes* for acquiring, accessing and subsequently releasing some resource (e.g. a file) and an operation *op* for accessing the resource. The *withRes* operation takes as argument a function *f* provided by the client for accessing the resource and takes care of acquiring the resource before *f* is called and subsequently releasing it again. For such a library, we typically wish to make sure that (1) clients only access the resource after they have acquired it, and (2) clients do not try to acquire resources they already hold. Such a library can be given a specification in Separation Logic as follows:

$$\begin{aligned} \Phi_{\text{brac}}(P_0, (\text{withRes}, \text{op})) &\triangleq \\ &\exists \text{locked} : iProp. \exists \text{unlocked} : Val \rightarrow iProp. \Box(P_0 \multimap \text{locked}) \wedge \\ &\forall P, Q : iProp. \forall f : Val. \{\text{locked} * P * S(P, Q, \text{unlocked}, f)\} \text{withRes}(f) \{\text{locked} * Q\} \wedge \\ &\forall x, y : Val. \{\text{unlocked}(y)\} \text{op}(x) \{\text{unlocked}(y)\} \\ \\ S(P, Q, \text{unlocked}, f) &\triangleq \forall x, y. \{\text{unlocked}(y) * P\} f(x) \{\text{unlocked}(y) * Q\} \end{aligned}$$

This specification relies on two abstract resources, *unlocked* and *locked*, to capture the well-bracketing aspect of the protocol. One can see that calling *withRes* requires the client to relinquish ownership of the *locked* resource. Since the function provided by the client is only given ownership of the abstract *unlocked* resource, it cannot make reentrant calls to *withRes*. Furthermore, to call *op* requires ownership of the *unlocked* resource, thus ensuring that only the callback provided by the client to *withRes* can call *op*.

But this specification entails another less obvious property. In the specification of *withRes*, the library must behave parametrically with respect to *P* and *Q*, which are provided by the client. Thus, the specification ensures that *withRes* is forced to call the function provided by the client *exactly once*, as the only way it can transform *P* into *Q* is by calling the function provided by the client.

As before, we wish to establish these properties as free theorems obtained from the specification. We define a wrapping function that instruments the library to emit call and return events for all calls where control passes between client and library.

$$\begin{aligned} \text{wrap}_{\text{brac}} &\triangleq \lambda(\text{withRes}, \text{op}). (\lambda f. \text{emit}\langle \text{call}, \text{withRes}, f \rangle; \\ &\quad \text{withRes}(\lambda x. \text{emit}\langle \text{call}, f \rangle; f(x); \text{emit}\langle \text{ret}, f \rangle); \\ &\quad \text{emit}\langle \text{ret}, \text{withRes}, f \rangle, \\ &\quad \lambda x. \text{emit}\langle \text{call}, \text{op} \rangle; \text{op}(x); \text{emit}\langle \text{ret}, \text{op} \rangle) \end{aligned}$$

We now state formally the desired property as a well-bracketing property of the traces generated by the instrumented library. Let us define $\mathcal{L}_{\text{bracfull}} \subseteq Val^*$ the language of traces of the form:

$$s' \cdot \langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s_{\text{op}} \cdot \langle \text{ret}, f \rangle \cdot \langle \text{ret}, \text{withRes}, f \rangle$$

for some $f \in Val$, $s_{\text{op}} \in (\langle \text{call}, \text{op} \rangle \cdot \langle \text{ret}, \text{op} \rangle)^*$ and $s' \in \mathcal{L}_{\text{bracfull}}$. That is, the traces in $\mathcal{L}_{\text{bracfull}}$ are well-bracketed sequences of events formed of subsequences of the form $\langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s_{\text{op}} \cdot \langle \text{ret}, f \rangle \cdot \langle \text{ret}, \text{withRes}, f \rangle$, and subsequences thereof, where s_{op} a sequence of consecutive calls and returns of *op*. Then, we take $\mathcal{L}_{\text{brac}}$, the language of traces we wish to enforce, to be the prefix closure of $\mathcal{L}_{\text{bracfull}}$.

To prove that the specification enforces the trace property we proceed as usual, by proving that for any implementation *M* that satisfies the specification, the wrapped implementation $\text{wrap}_{\text{brac}}(M)$ also satisfies the specification and generates traces in $\mathcal{L}_{\text{brac}}$.

$$\text{LEMMA 5.3. } \forall P_0, \text{ops. } \Phi_{\text{brac}}(P_0, \text{ops}) \vdash \Phi_{\text{brac}}(P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{brac}}, \text{wrap}_{\text{brac}}(\text{ops})).$$

PROOF SKETCH. To prove that the wrapped version satisfies the specification and produces traces in $\mathcal{L}_{\text{brac}}$, we first need to define wrapped versions of the abstract representation predicates. The idea is to let the wrapped locked resource, locked_w , express that the current trace t is in $\mathcal{L}_{\text{brac}}$ and the $\langle \text{call}, \text{withRes}, f \rangle$ and $\langle \text{ret}, \text{withRes}, f \rangle$ events in t are well-balanced and well-bracketed. For the wrapped unlocked resource, $\text{unlocked}_w(x)$, the idea is to use the argument x to track the name f of the last unbalanced $\langle \text{call}, \text{withRes}, f \rangle$ event in t .

To simplify the definitions and subsequent proofs, we first introduce a number of auxiliary resources, T_0, T_1, T_2 and T_3 . T_0 expresses that the current trace is well-balanced. We set $O = (\langle \text{call}, \text{op} \rangle \cdot \langle \text{ret}, \text{op} \rangle)^*$. $T_1(f)$ expresses that the current trace t has the form $s \cdot \langle \text{call}, \text{withRes}, f \rangle$ where s is well-balanced. $T_2(f)$ expresses that the current trace t has the form $s \cdot \langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s'$ where s is well-balanced and $s' \in O$. Finally, $T_3(f)$ expresses that the current trace t has the form $s \cdot \langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s' \cdot \langle \text{ret}, f \rangle$ where s is well-balanced and $s' \in O$.

$$\begin{aligned}
T_0 &= \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) * [t \in \mathcal{L}_{\text{bracfull}}] \\
T_1(f) &= \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) * \exists t' \in \mathcal{L}_{\text{bracfull}}. [t = t' \cdot \langle \text{call}, \text{withRes}, f \rangle] \\
T_2(f) &= \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) * \\
&\quad \exists t' \in \mathcal{L}_{\text{bracfull}}, s_{\text{op}} \in O. [t = t' \cdot \langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s_{\text{op}}] \\
T_3(f) &= \exists t. \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) * \\
&\quad \exists t' \in \mathcal{L}_{\text{bracfull}}, s_{\text{op}} \in O. [t = t' \cdot \langle \text{call}, \text{withRes}, f \rangle \cdot \langle \text{call}, f \rangle \cdot s_{\text{op}} \cdot \langle \text{ret}, f \rangle]
\end{aligned}$$

With these resources, we can now define unlocked_w and locked_w :

$$\begin{aligned}
\text{unlocked}_w(x) &\triangleq \exists y, z : \text{Val}. [x = (y, z)] * \text{unlocked}(y) * T_2(z) \\
\text{locked}_w &\triangleq \text{locked} * T_0
\end{aligned}$$

It follows easily that $P_0 * \text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) \multimap \text{locked}_w$ from $\text{trace}(\varepsilon) * \text{tracelnv}(\mathcal{L}_{\text{brac}}) \multimap T_0$ and the assumption $P_0 \multimap \text{locked}$.

It remains to show that the two instrumented operations satisfy their specifications. We begin by showing that the instrumented *withRes* operation satisfies its specification:

$$\forall P, Q : i\text{Prop}. \forall f : \text{Val}.$$

$$\{\text{locked}_w * P * S(P, Q, \text{unlocked}_w, f)\} \text{fst } (\text{wrap}_{\text{brac}}(\text{withRes}, \text{op}))(f) \{\text{locked}_w * Q\}$$

assuming *withRes* satisfies its specification:

$$\forall P, Q : i\text{Prop}. \forall f : \text{Val}. \{\text{locked} * P * S(P, Q, \text{unlocked}, f)\} \text{withRes}(f) \{\text{locked} * Q\}.$$

We give a proof outline for the instrumented *withRes* operation:

<pre> {locked_w * P * S(P, Q, unlocked_w, f)} {locked * T₀ * P * S(P, Q, unlocked_w, f)} emit⟨call, withRes, f⟩; {locked * T₁(f) * P * S(P, Q, unlocked_w, f)} let g = λx. emit⟨call, f⟩; f(x); emit⟨ret, f⟩ in {locked * T₁(f) * P * S(P * T₁(f), Q * T₃(f), unlocked, g)} withRes(g); {locked * T₃(f) * Q} emit⟨ret, withRes, f⟩ {locked * T₀ * Q} {locked_w * Q} </pre>		<pre> {unlocked(y) * P * T₁(f) * S(P, Q, unlocked_w, f)} emit⟨call, f⟩; {unlocked(y) * P * T₂(f) * S(P, Q, unlocked_w, f)} {unlocked_w((y, f)) * P * S(P, Q, unlocked_w, f)} f(x); {unlocked_w((y, f)) * Q} {unlocked(y) * Q * T₂(f)} emit⟨ret, f⟩ {unlocked(y) * Q * T₃(f)} </pre>
---	--	---

The interesting step, that we detail on the right-hand side of the proof outline, is showing that from the assumed specification of f we can derive the desired specification for the instrumented version of f :

$$\forall P, Q : iProp. \forall f : Val. S(P, Q, \text{unlocked}_w, f) \Rightarrow S(P * T_1(f), Q * T_3(f), \text{unlocked}, \lambda x. \text{emit}(\text{call}, f); f(x); \text{emit}(\text{ret}, f)).$$

Lastly, we need to show that the instrumented op function satisfies its specification. This follows easily from the fact that unlocked_w enforces $T_2(f)$, which is itself preserved when emitting a new $\langle \text{call}, op \rangle$ event. $\square$

5.4 Traversable Stack Example

To further demonstrate that our approach can express and enforce strong trace properties, recall the stack example from §2. We have a stack with a *push* and a *pop* method, and a *foreach* method that takes a function argument and applies the given function to every element of the stack, in order, starting from the top-most element. Here the protocol on the interaction between client and library imposes restrictions on both the client and the library. In particular, we wish to ensure that the function provided by the client cannot call back into the stack-library and potentially modify the underlying stack during the iteration of the stack. We also wish to ensure that the library calls the function provided by the client with every element currently on the stack and in the right order.

A higher-order separation logic specification for such a stack data structure is the following.

$$\Phi(P_{init}, (\text{push}, \text{pop}, \text{foreach})) \triangleq$$

$$\exists \text{stack} : Val^* \rightarrow iProp. \Box(P_{init} \multimap \text{stack}(\varepsilon)) \wedge$$

$$\forall \alpha, x. \{\text{stack}(\alpha) * [x \neq ()]\} \text{push}(x) \{\text{stack}(x :: \alpha)\} \wedge$$

$$\forall \alpha. \{\text{stack}(\alpha)\} \text{pop}() \{r. ([r = () \wedge \alpha = \varepsilon] * \text{stack}(\alpha)) \vee (\exists \alpha'. [r = \alpha' \wedge \text{stack}(\alpha')])\} \wedge$$

$$\forall \alpha, f, I. \{\text{stack}(\alpha) * I(\varepsilon) * \forall \beta, x. \{I(\beta)\} f(x) \{I(x :: \beta)\}\} \text{foreach}(f) \{\text{stack}(\alpha) * I(\text{rev}(\alpha))\}$$

It asserts existence of an abstract stack representation predicate $\text{stack}(\alpha)$ that tracks the exact sequence of elements currently on the stack using the mathematical sequence α . The specification for *push* and *pop* is straightforward: pushing and popping elements pushes or pops elements from this mathematical sequence, with a few special cases for pushing $()$ or popping from an empty stack. The specification for *foreach* is more interesting. It is parametrised by a predicate I , to be chosen by the client. This predicate is indexed by a sequence β and $I(\beta)$ is intended to capture the client's state after the function provided by the client has been called on each element of β , in reverse order. This accounts for the $I(\text{rev}(\beta))$ in the post-condition, where rev is the reverse operator on sequences.

We now wish to show as a free theorem that any implementation satisfying this Separation Logic stack specification will satisfy the informal trace property described above. We can formalize the intended protocol as the language $\mathcal{L}_{\text{stack}}$ defined as the prefix closure of the language of all traces t such that $\text{stk}_{\text{tr}}(t, \varepsilon)$ holds, where, given $\alpha \in Val^*$, we define $\text{stk}_{\text{tr}}(t, \alpha)$ by:

$$\text{stk}_{\text{tr}}(t, \alpha) \triangleq (t = \alpha = \varepsilon) \vee$$

$$(t = t' \cdot \langle \text{call}, \text{push}, x \rangle \cdot \langle \text{ret}, \text{push} \rangle \wedge \alpha = x :: \alpha' \wedge \text{stk}_{\text{tr}}(t', \alpha')) \vee$$

$$(t = t' \cdot \langle \text{call}, \text{pop} \rangle \cdot \langle \text{ret}, \text{pop}, () \rangle \wedge \alpha = \varepsilon \wedge \text{stk}_{\text{tr}}(t', \varepsilon)) \vee$$

$$(t = t' \cdot \langle \text{call}, \text{pop} \rangle \cdot \langle \text{ret}, \text{pop}, x \rangle \wedge \text{stk}_{\text{tr}}(t', x :: \alpha)) \vee$$

$$(t = t' \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t'' \cdot \langle \text{ret}, \text{foreach} \rangle \wedge \text{stk}_{\text{tr}}(t', \alpha) \wedge \text{trav}(t'', \alpha, f))$$

$$\text{trav}(t, \alpha, f) \triangleq (t = \alpha = \varepsilon) \vee (t = \langle \text{call}, f, x \rangle \cdot \langle \text{ret}, f \rangle \cdot t' \wedge \alpha = x :: \alpha' \wedge \text{trav}(t', \alpha', f))$$

To prove the free theorem we then define a suitable library wrapper that tracks all calls to push and pop and all calls to the function argument provided by the client when calling foreach.

$$\begin{aligned} \mathbf{wrap}_{\text{stack}}(\text{push}, \text{pop}, \text{foreach}) \triangleq & \\ (\lambda x. \text{emit}\langle \text{call}, \text{push}, x \rangle; \text{push}(x); \text{emit}\langle \text{ret}, \text{push} \rangle, & \\ \lambda_. \text{emit}\langle \text{call}, \text{pop} \rangle; \text{let } x = \text{pop}() \text{ in } \text{emit}\langle \text{ret}, \text{pop}, x \rangle; x, & \\ \lambda f. \text{emit}\langle \text{call}, \text{foreach}, f \rangle; \text{foreach}(\lambda x. \text{emit}\langle \text{call}, f, x \rangle; f(x); & \\ \text{emit}\langle \text{ret}, f \rangle); \text{emit}\langle \text{ret}, \text{foreach} \rangle) & \end{aligned}$$

LEMMA 5.4. $\forall P_0, ops. \Phi(P_0, ops) \vdash \Phi(P_0 * \text{tracelnv}(\mathcal{L}_{\text{stack}}) * \text{trace}(\varepsilon), \mathbf{wrap}_{\text{stack}}(ops)).$

PROOF SKETCH. We proceed by defining a wrapped version of the stack predicate that asserts that the sequence of elements α matches the contents of the stack as per the current trace t .

$$\text{stack}_w(\alpha) \triangleq \text{stack}(\alpha) * \exists t. [\text{stk}_{\text{tr}}(t, \alpha)] * \text{trace}(t) * \text{tracelnv}(\mathcal{L}_{\text{stack}})$$

Clearly we have that $P_{\text{init}} * \text{tracelnv}(\mathcal{L}_{\text{stack}}) * \text{trace}(\varepsilon) \multimap \text{stack}_w(\varepsilon)$ and $P_{\text{init}} \multimap \text{stack}(\varepsilon)$.

It remains to prove the wrapped library methods satisfy the specification instantiated with the wrapped stack predicate. The proofs for *push* and *pop* are straightforward and have been omitted. For *foreach* we are given a predicate I from the client and need to prove the following triple:

$$\begin{aligned} & \{ \text{stack}_w(\alpha) * I(\varepsilon) * \forall \beta, x. \{ I(\beta) \} f(x) \{ I(x :: \beta) \} \} \\ & \quad (\text{snd}(\text{snd}(\mathbf{wrap}_{\text{stack}}(ops)))(f)) \\ & \{ \text{stack}_w(\alpha) * I(\text{rev}(\alpha)) \} \end{aligned}$$

In the call to the underlying *foreach* method, we can pick a suitably wrapped version of the I predicate, $I_w(\beta)$. The idea is that it should assert $I(\beta)$ and that we have emitted an opening *foreach* call and called the function argument on all the elements of β so far.

$$I_w(\beta) \triangleq I(\beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(\beta), f)] * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2)$$

We need to prove that the wrapped function argument updates the wrapped I predicate appropriately. This follows from:

$$\begin{aligned} & \{ I_w(\beta) \} \\ & \{ I(\beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(\beta), f)] * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2) \} \\ & \quad \mathbf{emit} \langle \text{call}, f, x \rangle; \\ & \{ I(\beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(\beta), f)] * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2 \cdot \langle \text{call}, f, x \rangle) \} \\ & \quad f(x); \\ & \{ I(x :: \beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(\beta), f)] * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2 \cdot \langle \text{call}, f, x \rangle) \} \\ & \quad \mathbf{emit} \langle \text{ret}, f \rangle; \\ & \{ I(x :: \beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(\beta), f)] \\ & \quad * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2 \cdot \langle \text{call}, f, x \rangle \cdot \langle \text{ret}, f \rangle) \} \\ & \{ I(x :: \beta) * \exists t_1, t_2. [\text{stk}_{\text{tr}}(t_1, \alpha) \wedge \text{trav}(t_2, \text{rev}(x :: \beta), f)] * \text{trace}(t_1 \cdot \langle \text{call}, \text{foreach}, f \rangle \cdot t_2) \} \\ & \{ I_w(x :: \beta) \} \end{aligned}$$

The second to last step follows from the following property:

$$\forall \alpha, x, t. \text{trav}(t, \alpha, f) \implies \text{trav}(t \cdot \langle \text{call}, f, x \rangle \cdot \langle \text{ret}, f \rangle, \alpha \cdot x, f)$$

Now the rest of the proof of *foreach* is just an application of the specification of the underlying *foreach* method and the EMIT rule for the emission of the *foreach* call and return events. $\square$

6 RELATING LOGICAL ATOMICITY AND LINEARIZABILITY

The specifications that we have considered so far are sequential in nature. Indeed, even though we are working in a concurrent language, and even though both the implementation of a library and its client can in principle make use of concurrency primitives, the specifications presented earlier enforce that a client may only call the operations of the library in a sequential fashion.

For instance, the specification for the *push* operation of the stack library requires the client to give up exclusive ownership over the stack data structure for the duration of the call, thus preventing concurrent operations to be made.

And indeed, nothing requires an implementation of this stack library to be thread-safe: during the execution of *push*, the internal invariants on the stack data structure might be temporarily broken, in a way that would be observable by concurrent operations. An unverified client attempting to use the library in a concurrent fashion might then observe bogus results or corrupt the data structure.

Proper concurrent data structures are instead carefully designed to ensure that all operations are safe to call concurrently. In particular, they usually ensure that operations appear to take effect atomically, in order to alleviate the need for clients to reason about interleavings of the individual instructions within each operation.

Classically this idea of atomicity is formalized as a property of traces of the interactions between library and client called *linearizability* [Herlihy and Wing 1990]. More recently, in the setting of program logics such as TaDa [da Rocha Pinto et al. 2014] and Iris [Jung et al. 2015a] (see also [Birkedal and Bizjak 2020, Sec. 13]), an alternative notion called *logical atomicity* has been proposed as an alternative to linearizability more amenable to program verification in Hoare logic. Logical atomicity allows expressing the idea of linearizability as an abstract specification for a library, thus allowing clients to reason about operations that appear atomic using the same reasoning principles as physically atomic operations.

In this section, we show how to establish a formal connection between linearizability and logical atomicity. More precisely, we show that any library that can be verified against a logically atomic specification is in fact linearizable.

6.1 Logical Atomicity

In this subsection we recall how logical atomicity is captured in Iris. As mentioned above, a logically atomic specification for an operation is intended to express that the operation behaves *as if it was atomic*. Now, from a logical perspective, what reasoning power do we gain from this property, when reasoning about calls to logically atomic operations? The answer is that logical atomicity *allows us to access invariants* around logically atomic operations. In Iris, invariants are the key mechanism used to manage resources that are shared and updated between concurrent threads according to certain protocols. Invariants (which have to hold at every step of the execution) can only be accessed for the duration of a single atomic execution step, during which the resources held by the invariant become accessible. This is described by the invariant opening rule INV-OPEN:³

$$\frac{\text{INV-OPEN} \quad \{P * R\} \ e \ \{v. Q * R\} \quad e \text{ is atomic}}{\boxed{R} \vdash \{P\} \ e \ \{v. Q\}}$$

What logical atomicity grants us is the power to use a similar invariant opening rule around the call to a logically atomic operation, even though its implementation is not physically atomic (it performs more than one step of computation), as long as it behaves as if it were.

³We omit the technical details related to namespaces and masks throughout this section. They are important for soundness but not for the presentation here, and we refer the reader to the Coq formalization for further details.

Logically atomic triples. Logically atomic operations are specified using logically atomic Hoare triples, $\langle P \rangle e \langle v. Q \rangle$. These triples support an invariant opening rule similar to INV-OPEN:

$$\frac{\text{LOGATOM-INV-OPEN} \quad \langle P * R \rangle e \langle v. Q * R \rangle}{\boxed{R} \vdash \langle P \rangle e \langle v. Q \rangle}$$

Note that in contrast to INV-OPEN, this rule does not require the expression e to be physically atomic. Unlike a normal Hoare triple, which expresses that the resources in precondition P are updated to Q after possibly many execution steps by running e , a logically atomic triple requires that updates occur during a single execution step. As such, a triple $\langle P \rangle e \langle v. Q \rangle$ is perhaps best seen as the specification for the *linearization point* that occurs during the execution of e : the single internal step that atomically updates the observable state of the data structure.

Using and proving logically atomic triples. Logically atomic triples are defined in Iris in terms of a more primitive notion of *atomic update*. We have the following rule:

$$\text{LOGATOM-AU} \quad \langle P \rangle e \langle v. Q \rangle \dashv \vdash \forall \Psi. \{ \text{AU}_{P,Q}(\Psi) \} e \{ \Psi \}.$$

That is, a logically atomic triple is equivalent to a normal triple where we do not directly get access to the precondition P , but instead to a resource $\text{AU}_{P,Q}(\Psi)$ representing the right and obligation to update resources in P into resources in Q in one step. As witnessed by the universally quantified postcondition Ψ , when proving a logically atomic triple, one is indeed required to perform the atomic update (or, in other words, choose a linearization point) as it is the only way of obtaining Ψ .

Atomic updates are not themselves primitives to Iris, but for our purposes we will keep their definition abstract (more details can be found in [Jung et al. 2015b] and in the Coq formalization of Iris). Informally, it should be enough to remember that having ownership over an atomic update $\text{AU}_{P,Q}$ gives us the liberty to choose a linearization point where we get access to resources P and must update them into Q , while being able to access invariants.

6.2 Specification of a Concurrent Library

In this subsection we show what kind of logically atomic specification we consider for a library with a concurrent operation op . To simplify the presentation, we consider a library that just exposes a single operation, which can then be called concurrently by a client, possibly with different arguments. Note that this is in fact equivalent to considering a library exposing several operations: it is always possible for an operation to dispatch to several concurrent sub-operations depending on its argument.

Since we wish to cover a wide range of possible operations and specifications, we consider a fairly abstract specification, where a program value is related to a type S of abstract states and where two mathematical functions f and r ⁴, of type $S \times \text{Val} \rightarrow S$ and $S \times \text{Val} \rightarrow \text{Val}$ respectively, are used to specify what the updated abstract state and return value should be after executing the operation, depending on its argument. Thus we will use a specification of the form:

$$\langle s. \text{PState}(x, s) \rangle op(x, y) \langle z. \text{PState}(x, f(s, y)) * [z = r(s, y)] \rangle.$$

Here x is a pointer to the data structure that the library implements and which is accessed and modified by op ; and y is the argument passed to this specific call to op . PState corresponds to the abstract representation predicate describing the Separation Logic resources associated with x , and corresponding to the logical model s of the data structure (for instance, for a concurrent stack data structure, s would be the sequence of values that it stores). Notice that s is bound in the precondition

⁴ f and r could be easily taken to be relations instead, a generalization that we have implemented in the Coq formalization.

of the triple, instead of being bound on the outside as usual. This reflects the fact that the value of s we are considering will only be known at the linearization point: because the execution of op might take several steps to execute, the value of s might *change* during its execution, and thus cannot be predicted upfront.

Now, it is common for concurrent operations to rely on additional resources which are held in an invariant. To allow for that as well, we parameterize the specification further by a persistent predicate IsP , which we assume to hold. To relate the instances of IsP and $PState$ that correspond to the same instantiation of the library, we use a ghost name γ . To sum up, $IsP(\gamma, x)$ can be shared between different threads, and expresses that the name γ is associated with the structure at address x . $PState(\gamma, s)$ represents the knowledge that the current abstract state of the library is s , and is non-duplicable as it represents ephemeral information that gets updated by calling op . We thus refine the above specification to:

$$IsP(\gamma, x) \vdash \langle s. PState(\gamma, s) \rangle op(x, y) \langle z. PState(\gamma, f(s, y)) * [z = r(s, y)] \rangle.$$

We finally assume that we are provided an initial abstract state $s_{init} \in S$ for the library, and a function op_{init} which, given initial resources P_{init} , initializes the library and returns a new instance.

Combining all these elements, our complete specification for a generic concurrent library is written as follows:

$$\begin{aligned} \Phi(op_{init}, op, (S, s_{init}, f, r), P_{init}, IsP, PState) \triangleq \\ \forall \gamma, x. \text{persistent}(IsP(\gamma, x)) \wedge \\ \{P_{init}\} op_{init}() \{r. \exists \gamma. IsP(\gamma, r) * PState(\gamma, s_{init})\} \wedge \\ \forall \gamma, x, y. IsP(\gamma, x) \vdash \langle s. PState(\gamma, s) \rangle op(x, y) \langle z. PState(\gamma, f(s, y)) \wedge [z = r(s, y)] \rangle. \end{aligned}$$

Example: a concurrent counter library. It can be useful to see a concrete example of the above generic specification. To this end, let us consider a simple counter library, exposing an `init` operation that creates a fresh counter, and an `incr` concurrent operation which increments the integer value held by the counter and returns its previous value. One can for instance implement `incr` by using the compare-and-swap primitive (CAS) to ensure that the increment occurs atomically. The counter library can be then specified as follows, using a logically atomic triple for `incr`:

$$\begin{aligned} \{\text{True}\} \text{init}() \{x. \exists \gamma. IsCnt(\gamma, x) * Cnt(\gamma, 0)\} \wedge \\ \forall \gamma, x. IsCnt(\gamma, x) \vdash \langle n. Cnt(\gamma, n) \rangle \text{incr}(x) \langle m. Cnt(\gamma, n + 1) * [m = n + 1] \rangle, \end{aligned}$$

where Cnt and $IsCnt$ are abstract predicates and $IsCnt$ is persistent. By letting $\text{incr}'(x, ()) \triangleq \text{incr}(x)$, this counter specification can be seen as an instance of our generic specification above: it corresponds to $\Phi(\text{init}, \text{incr}', (\mathbb{N}, 0, \text{succ}_{\mathbb{N}}, \text{succ}_{Val}), \text{True}, IsCnt, Cnt)$, where $\text{succ}_{\mathbb{N}}$ is the successor function on natural numbers, and succ_{Val} is the successor function on program values representing integers.

6.3 Linearizability

Linearizability is the standard correctness condition for (first-order) operations on thread-safe concurrent data structures [Herlihy and Wing 1990]. It expresses that, even if concurrent operations of a library are effectively interleaved at runtime, they behave the same as a sequence of sequential calls to the library. It can be formalized by requiring that between every call and return in a given interaction trace, there must exist a *single physical step* where the observable state of the data structure changes according to the specification of the operation. This point is known as the *linearization point*, and the behavior of the data structure must then match the sequential behavior of the operations when ordered according to the linearization points.

Let us define an alphabet Σ of call and return events and an alphabet Σ_{lin} of call, linearization and return events as follows:

$$\begin{aligned}\Sigma &\triangleq \{\langle \tau, \langle \text{call}, v \rangle \rangle \mid \tau \in \text{Tag}, v \in \text{Val}\} \cup \{\langle \tau, \langle \text{ret}, v, v' \rangle \rangle \mid \tau \in \text{Tag}, v, v' \in \text{Val}\} \\ \Sigma_{\text{lin}} &\triangleq \Sigma \cup \{\langle \tau, \langle \text{lin}, v, v' \rangle \rangle \mid \tau \in \text{Tag}, v, v' \in \text{Val}\}\end{aligned}$$

Here τ is a tag used to relate calls, linearization and return events and v and v' are values representing the argument and return-value, respectively. We define the language $\mathcal{L}_{\text{linfull}}^\tau$ of calls, linearizations and returns tagged with tag τ as follows. We then denote its prefix closure as $\mathcal{L}_{\text{lin}}^\tau$.

$$\mathcal{L}_{\text{linfull}}^\tau ::= \langle \tau, \langle \text{call}, v \rangle \rangle \cdot \langle \tau, \langle \text{lin}, v, v' \rangle \rangle \cdot \langle \tau, \langle \text{ret}, v, v' \rangle \rangle \cdot \mathcal{L}_{\text{linfull}}^\tau \mid \varepsilon$$

As before, we assume that the intended sequential behaviour of the library is given by a tuple $(S, s_{\text{init}}, (f, r))$ consisting of a set of abstract states S along with two functions (f, r) and an initial state $s_{\text{init}} \in S$. The function call $f(s, v)$ models a call to the operation in abstract state s and with argument v ; it returns the new abstract state. The function call $r(s, v)$ gives the return value of the operation when abstract state is s and argument is v . A trace of linearization events satisfies the specification $(S, s_{\text{init}}, (f, r))$ if the return values in the trace match the intended return values given by r . This is captured by the sound predicate, defined below:

$$\text{sound}(\alpha) \triangleq \exists s, \text{soundWith}(\alpha, s_{\text{init}}, s)$$

where

$$\begin{aligned}\text{soundWith}(\varepsilon, s, s') &\triangleq s = s', \\ \text{soundWith}(\alpha \cdot \langle \text{lin}, v, v_r \rangle, s, s'') &\triangleq \exists s', \text{soundWith}(\alpha, s, s') \wedge s'' = f(s', v) \wedge v_r = r(s', v)\end{aligned}$$

For a linearization trace $\beta \in \Sigma_{\text{lin}}^*$, we write $\pi(\beta) \in \Sigma^*$ for the trace consisting of only the `call` and `ret` events of β ; we write $\pi_\tau(\beta) \in \Sigma_{\text{lin}}^*$ for the sub-trace of events with tag τ ; and $\pi_{\text{lin}}(\beta) \in \text{Val}^*$ for the list of linearization events associated to `lin` events in β . We finally define the language $\mathcal{L}_{\text{lin}}$ of linearizable traces as follows:

Definition 6.1 (Linearizability). An interaction trace $\alpha \in \Sigma^*$ is linearizable if there exists a trace $\beta \in \Sigma_{\text{lin}}^*$ such that $\forall \tau. \pi_\tau(\beta) \in \mathcal{L}_{\text{lin}}^\tau$, $\alpha = \pi(\beta)$, and $\text{sound}(\pi_{\text{lin}}(\beta))$.

6.4 Linearizability as a Trace Invariant

To prove that logical atomicity implies linearizability, we follow the usual procedure of introducing a wrapper that emits suitable trace events and proving that the wrapped library satisfies the same logically atomic specification as the original library and additionally a suitable invariant about the trace. The wrapping for the operation op is given by the following function:

$$\text{wrap}(op) \triangleq \lambda(x, y). \text{let } \tau = \text{fresh } \langle \text{call}, y \rangle \text{ in let } r = op(x, y) \text{ in emit } \langle \tau, \langle \text{ret}, y, r \rangle \rangle; r$$

This wrapping function emits suitable call and return events before and after the call to the underlying op operation. We use the `fresh` operation to generate a fresh identifier for every call event to allow us to uniquely associate return events with call events.

The following theorem then states the key correctness property: given any library satisfying a logically atomic specification, the wrapped library satisfies the same specification, and the trace it emits is linearizable.

THEOREM 6.2.

$$\begin{aligned}&\forall op_{\text{init}}, op, spec, P_{\text{init}}, IsP, PState. \\ &\Phi(op_{\text{init}}, op, spec, P_{\text{init}}, IsP, PState) \Rightarrow \\ &\exists IsP_w, PState_w. \Phi(op_{\text{init}}, \text{wrap}(op), spec, P_{\text{init}} * \text{trace}(\varepsilon) * \text{traceInv}(\mathcal{L}_{\text{lin}}), IsP_w, PState_w)\end{aligned}$$

As before, we can finally combine Theorem 6.2 with the Adequacy theorem (Theorem 4.1) and obtain as a free theorem that any library which admits a logically atomic specification produces an interaction trace in $\mathcal{L}_{\text{lin}}$, i.e. is linearizable.

In the rest of this section, we detail the key ideas behind the proof of Theorem 6.2.

6.5 Proving Linearizability

Our proof relies on two key ideas.

First, the library specification that we assume exposes that the operation op contains a linearization point during its execution. We can thus use this point in the execution as the moment where to record a lin event in the “ghost” linearization trace that we maintain as part of our trace invariant. Since the call to op in $\text{wrap}(op)$ occurs between the instructions emitting a call and ret event, we get that the lin event will indeed occur between the corresponding call and ret events.

Second, a key property we need to keep track of is that our use of tags and the fresh primitive indeed identify different calls to the wrapped operation with different tags. More precisely, we need to show that during the execution of $\text{wrap}(op)$, once a fresh tag has been generated, the currently running thread will be the only one emitting events with this tag. Thus, even though other threads can run concurrently, at every point of the execution of $\text{wrap}(op)$ —after the call to fresh , after the call to op , and after the final emit —we know exactly what the linearization trace is *for events of that tag*: it is equal, respectively, to a single call event, a call event followed by a lin event, and finally a call , lin , and ret event.

We now show in more detail how these key ideas materialize in our proof.

Tracking ghost state in Iris. Our proof makes use of somewhat more advanced features of Iris, which we briefly introduce now. In particular, we define certain ghost theories (predicates and rules) on top of Iris’ primitive notion of ghost state, which supports predicates of the form $\llbracket \bar{a} \rrbracket^\gamma$ that asserts the ownership of the resource a at ghost name $\gamma \in \text{Name}$. Here the resource a is an element of some resource algebra which restricts how it might be updated over time.

We use two constructions built on top of the Iris primitives for handling ghost state. How they are defined in terms of $\llbracket \cdot \rrbracket^\gamma$ is not particularly relevant, as long as they satisfy a few rules, which will suffice for carrying out the rest of the proof and which we present below:

The first construction involves a pair of predicates $\text{map}_\bullet$ and $\text{map}_\circ$, used for tracking the state of a finite map M . Given a name γ which identifies a specific instance of these predicates, ownership over $\text{map}_\bullet(\gamma, M)$ asserts that the map of interest is exactly M , while $\text{map}_\circ(\gamma, M')$ asserts that the map of interest is *at least* M' . The predicates satisfy the following rules, which we explain below.

$$\begin{array}{ll}
 \text{GHOST-MAPFRAC-DUP} & \text{GHOST-MAP-SUB} \\
 \text{map}_\circ(\gamma, M) \vdash \text{map}_\circ(\gamma, M) * \text{map}_\circ(\gamma, M) & \text{map}_\bullet(\gamma, M) * \text{map}_\circ(\gamma, M') \vdash \llbracket M' \subseteq M \rrbracket \\
 \\
 \text{GHOST-MAP-UPD} & \\
 \text{map}_\bullet(\gamma, M) * \text{map}_\circ(\gamma, M) \vdash \text{map}_\bullet(\gamma, M \uplus [a \mapsto b]) * \text{map}_\circ(\gamma, M \uplus [a \mapsto b]) &
 \end{array}$$

Given the combined knowledge of both predicates, one can deduce that M' is a subset of M (GHOST-MAP-SUB). In the case where both maps match, then we can update the map by monotonically extending it with a new key-value pair (GHOST-MAP-UPD). Finally, $\text{map}_\circ(\gamma, M')$ is duplicable (GHOST-MAPFRAC-DUP), as it represents a “snapshot” of the state of the map at some point in the past, which remains sound as the map is extended monotonically.

The second construction involves a pair of predicates $\text{val}_\bullet$ and $\text{val}_\circ$, which are used to track the state of a discrete value of an arbitrary type from two different places. Ownership over either $\text{val}_\bullet(\gamma, x)$ or $\text{val}_\circ(\gamma, y)$ tells us that the value of interest is exactly x . The predicates satisfy the

following rules, which we explain below.

GHOST-VAL-EQ

$$\text{val}_\bullet(\gamma, x) * \text{val}_\circ(\gamma, y) \vdash [x = y]$$

GHOST-VAL-UPD

$$\text{val}_\bullet(\gamma, x) * \text{val}_\circ(\gamma, x) \vdash \text{val}_\bullet(\gamma, z) * \text{val}_\circ(\gamma, z)$$

When combining the knowledge of both predicates, we can deduce that the associated values match (GHOST-VAL-EQ) or update it to a new value (GHOST-VAL-UPD). Unlike $\text{map}_\circ$, $\text{val}_\circ$ is not duplicable as it needs to stay synchronized with the matching $\text{val}_\bullet$ predicate.

Wrapped representation predicates and main invariant. Let us assume $op_{init}, op, (S, s_{init}, f, r), IsP$ and $PState$ such that $\Phi(op_{init}, op, (S, s_{init}, f, r), P_{init}, IsP, PState)$ holds. The first step of the proof is then to pick suitable wrapped representation predicates IsP_w and $PState_w$.

$$\begin{aligned} IsP_w &\triangleq \lambda(\gamma, x). \exists \gamma_i, \gamma_s, \gamma_m. IsP(\gamma_i, x) * \text{tracelnv}(\mathcal{L}_{lin}) * \boxed{\exists s. I(\gamma, \gamma_i, \gamma_s, \gamma_m, s)} \\ PState_w &\triangleq \lambda(\gamma, s). \exists \gamma_i, \gamma_s, \gamma_m. PState(\gamma_i, s) * \text{val}_\circ(\gamma, (\gamma_i, \gamma_s, \gamma_m)) * \text{val}_\circ(\gamma_s, s) \end{aligned}$$

Both IsP_w and $PState_w$ take as argument a ghost name γ which identifies a specific instance of the library. Because we need a few extra ghost names to track the various parts of ghost state that we need, we associate to γ three ghost names γ_i, γ_s and γ_m , using a $\text{val}_\bullet, \text{val}_\circ$ pair to ensure that they match between IsP_w and $PState_w$. The $\text{val}_\circ$ component of the pair appears in the definition of $PState_w$, and the $\text{val}_\bullet$ component will appear next, in the definition of our key invariant I , which is itself held by IsP_w .

IsP_w , the persistent part of the representation predicate, holds the corresponding abstract IsP predicate provided by the library specification, the $\text{tracelnv}(\mathcal{L}_{lin})$ predicate enforcing that the trace is indeed linearizable at each step, and an additional Iris invariant I , which we define below. While $\text{tracelnv}(\mathcal{L}_{lin})$ is ultimately the only property that we need to instantiate Theorem 4.1 and obtain that linearizability holds, during the proof we need to maintain additional invariants—held in I —to be able to show that it does hold as a trace invariant.

$PState_w$, which corresponds to knowledge that the library is in a specific state, asserts ownership over the underlying $PState$, and additionally holds halves of a $\text{val}_\bullet, \text{val}_\circ$ pair, associating respectively γ to the auxiliary ghost names used, and associating the state of the library s to ghost name γ_s . The corresponding $\text{val}_\bullet$ halves appear next, in the definition of I .

$$\begin{aligned} I(\gamma, \gamma_i, \gamma_s, \gamma_m, s) &\triangleq \exists \beta \in \Sigma_{lin}^*. \exists M : \text{Tag} \rightarrow_{fin} \text{Name}. [\text{dom}(M) = \text{tags}(\pi(\beta))] * \\ &\quad \text{val}_\bullet(\gamma, (\gamma_i, \gamma_s, \gamma_m)) * \text{val}_\bullet(\gamma_s, s) * \text{map}_\bullet(\gamma_m, M) * \text{map}_\circ(\gamma_m, M) * \\ &\quad \text{trace}(\pi(\beta)) * [\text{sound}(\pi_{lin}(\beta)) \wedge \forall \tau. \pi_\tau(\beta) \in \mathcal{L}_{lin}^\tau] * \\ &\quad \text{tagsState}(M, \beta) \end{aligned}$$

The trace β is the candidate linearization trace, composed of `call`, `lin` and `ret` events, such that its restriction to `call` and `ret` events $\pi(\beta)$ corresponds to the current physical trace.

The finite map M associates a ghost name for each tag currently appearing in the trace. As mentioned before, for each fresh tag that has been allocated, we wish to track the point of the execution where the corresponding wrapper is. M gets us halfway there: for each tag, it associates a ghost name which can then be used to track precisely the state of the wrapper as a ghost resource. This indirection allows M to grow monotonically as new tags are emitted, while updates to the state of a wrapper can be handled separately by updating the corresponding ghost resource.

The invariant I holds a half of a $\text{val}_\bullet, \text{val}_\circ$ pair for both the auxiliary ghost names and the state of the library s , matching the other halves held by the $PState_w$ predicate. It also holds both $\text{map}_\bullet$ and $\text{map}_\circ$ resources for tracking the state of the map M . This serves two purposes: first, it allows updating the map when a new tag is emitted by fresh using rule GHOST-MAP-UPD; second, it allows

handing out afterwards copies of $\text{map}_o(\gamma_m, M \uplus [\tau \mapsto \gamma_\tau])$ (using rule `GHOST-MAPFRAC-DUP`) to serve as witnesses that a tag τ has been emitted and is associated to a ghost name γ_τ .

Invariant I holds the resource trace $(\pi(\beta))$ granting the exclusive permission to update the physical trace. Reasoning about the calls to the `fresh` and `emit` primitive—which update this resource—will thus require opening I and re-establish it afterwards. I moreover enforces that the linearization events are sound, and that β is per-tag linearized, as required to prove linearizability.

Finally, the assertion $\text{tagsState}(M, \beta)$ (which we define now) describes the collection of ghost resources tracking the individual state of wrappers for each tag in M , relating it to the current linearization trace β .

$$\begin{aligned} \text{TagState} &\triangleq \text{AfterCall}(v) \mid \text{AfterLin}(v, v') \mid \text{Done} \\ \text{tagsState}(M, \beta) &\triangleq \\ &\quad \bigstar_{(\tau, \gamma_\tau) \in M} \left(\begin{array}{l} \exists \zeta : \text{TagState}. \text{val}_\bullet(\gamma_\tau, \zeta) * \\ \left[\begin{array}{l} \text{match } \zeta \text{ with} \\ \mid \text{AfterCall}(v) \Rightarrow \pi_\tau(\beta) = \langle \tau, \langle \text{call}, v \rangle \rangle \\ \mid \text{AfterLin}(v, v') \Rightarrow \pi_\tau(\beta) = \langle \tau, \langle \text{call}, v \rangle \rangle \cdot \langle \tau, \langle \text{lin}, v, v' \rangle \rangle \\ \mid \text{Done} \Rightarrow \text{True} \end{array} \right] \end{array} \right) \end{aligned}$$

Values of type *TagState*, here defined as an inductive datatype, describe the state of a thread running an instance of *wrap*(*op*): *AfterCall*(*v*) indicates that its execution is after a call to `fresh`(*call*, *v*) but before the linearization point; *AfterLin*(*v*, *v'*) indicates that it is after a linearization point *<lin, v, v'>* but before the final emit; and *Done* indicates that it is done, i.e. after the final emit. Then, for each pair of a tag τ and a ghost name γ_τ in M , $\text{tagsState}(M, \beta)$ names ζ the state *TagState* of the thread which has first allocated τ , and firstly, ties it with the ghost name γ_τ using the assertion $\text{val}_\bullet(\gamma_\tau, \zeta)$; then secondly, enforces that the value of ζ indeed gives us the expected information about events emitted so far with tag τ in β .

The matching assertion $\text{val}_\bullet(\gamma_\tau, \zeta)$ is not stored inside the invariant I . Instead, it will be kept “on the side” in the proof context, as a token expressing that the current state for tag τ is exactly ζ , for a concrete value of ζ . This token grants the unique, exclusive permission to emit new events with tag τ , updating the state ζ in the process. Consequently, owning such a token during the proof of *wrap*(*op*) guarantees that no other thread is able to emit events with the current tag of interest: this is one of the two key properties that we need to make the proof work.

Outline of the proof. There are three main steps to the proof: one to reason about the call to `fresh`, one for calling to the underlying operation *op*, and one for reasoning about the final emit.

At the beginning of the proof, we are given $\text{IsP}_w(\gamma, x)$, and are asked to prove a logically atomic triple. We first use the rule `LOGATOM-AU` to turn this triple into a normal Hoare triple:

$$\begin{aligned} &\{ \text{IsP}_w(\gamma, x) * \text{AU}_{(\forall s. \text{PState}_w(\gamma, s)), (\text{PState}_w(\gamma, f(s, y)))}(\Psi) \} \\ &\quad \text{let } \tau = \text{fresh } \langle \text{call}, y \rangle \text{ in let } r = \text{op}(x, y) \text{ in emit } \langle \tau, \langle \text{ret}, y, r \rangle \rangle; r \\ &\{ \Psi \} \end{aligned}$$

We can then start and step through the wrapper code:

- Call to `fresh` $\langle \text{call}, y \rangle$.

Because `fresh` is physically atomic (as a primitive of the language), this reasoning step follows the usual pattern of reasoning with invariants in concurrent separation logic: we unfold IsP_w and open the invariant containing I around the instruction. We thus get access to $I(\gamma, \gamma_i, \gamma_s, \gamma_m, s)$ for some $\gamma_i, \gamma_s, \gamma_s$ and s , which in turn gives us access to the trace resource required to call `fresh`. It is easy to check that adding a `call` event with a `fresh` tag preserves the linearizability of the trace.

Afterwards we need to close the invariant. Because we added a new tag to the trace, we need to extend the tagsState resource it contains to account for the new tag. We thus allocate a new pair $\text{val}_\bullet(\gamma_\tau, \text{AfterCall}(y)), \text{val}_\circ(\gamma_\tau, \text{AfterCall}(y))$ for a fresh γ_τ .

After closing the invariant, we obtain the following extra resources as witnesses of the operation, to keep around in our proof context for later:

$$\exists M, \gamma_\tau. \text{map}_\circ(\gamma_m, M) * [M(\tau) = \gamma_\tau] * \text{val}_\circ(\gamma_\tau, \text{AfterCall}(y))$$

- Call to $\text{op}(x, y)$.

Since op is logically atomic, we can again open the invariant containing I . Then, we need to establish the atomic update from the specification of op ($\text{AU}_{(\forall s. \text{PState}(\gamma_i, s)), (\text{PState}(\gamma_i, f(s, y)))}$) using the one we have in the context ($\text{AU}_{(\forall s. \text{PState}_w(\gamma, s)), (\text{PState}_w(\gamma, f(s, y)))}$).

This is the linearization point. Given a way of updating in one step the state of the unwrapped operation $\text{PState}(\gamma_i, s)$, we need to show that we can similarly update the state of the wrapped operation $\text{PState}_w(\gamma, s)$, while preserving our invariants.

Updating $\text{PState}_w(\gamma, s)$ to $\text{PState}_w(\gamma, f(s, y))$ requires updating the $\text{val}_\circ(\gamma_s, s)$ resource it contains. This is easy: we have the corresponding $\text{val}_\bullet$ resource available as part of the opened invariant.

Then, thanks to the $\text{val}_\circ(\gamma_\tau, \text{AfterCall}(y))$ resource that we have in the context, we can deduce that the current trace only contains one call event with tag τ . We can then show that $I(\gamma, \gamma_i, \gamma_s, \gamma_m, s)$ can be reformulated into $I(\gamma, \gamma_i, \gamma_s, \gamma_m, f(s, y))$ by recording a $\langle \text{lin}, y, r(s, y) \rangle$ event in the linearization trace, and that linearizability is preserved.

In the process, the state associated with γ_τ has to be updated, which means that we now hold $\text{val}_\circ(\gamma_\tau, \text{AfterLin}(y, r(s, y)))$ in the context.

- Call to emit $\langle \tau, \langle \text{ret}, y, r \rangle \rangle$.

As with the call to fresh, this is a call to a physically atomic operation, so we can reason as usual by opening then closing the main invariant. Because we hold $\text{val}_\circ(\gamma_\tau, \text{AfterLin}(y, r(s, y)))$, we deduce that the current linearization trace contains a call event followed by a lin event. This is enough to show that emitting a corresponding ret event preserves both the invariant I and the linearizability of the trace, thus completing the proof.

6.6 Related Work

Our key contribution is a technique for deriving free theorems from Separation Logic specifications, by formally relating separation logic specifications with the temporal properties they enforce, through wrapping of abstract resources with assertions about traces. As alluded to in the introduction, the idea of deriving free theorems from specifications is akin to the idea of deriving free theorems from polymorphic types [Reynolds and Plotkin 1993; Reynolds 1983; Wadler 1989]. In both cases, the term “free theorem” is used to highlight that a theorem is obtained from the specification / type; in particular, it holds for all programs satisfying the specification / type. In both cases, a “free theorem” needs proof. In the case of polymorphic types, one often uses a relationally parametric model of types (often formulated using logical relations) and the key point in the proof is then to pick a suitable relation for the quantified type variable. In our case, the key point in the proof is to wrap the assumed library implementation with operations that generate a suitable trace of the interactions between a library and a client and then show that the wrapped library satisfies the specification.

We now discuss some further related work not already discussed earlier in the paper.

Approaches based on type systems. Static analyses based on type systems have been widely used to check resource usage, like our file example, and more generally enforce interaction protocols.

A particularly influential approach has been Tpestates [DeLine and Fähndrich 2004; Strom and Yemini 1986]. These can be seen as specifying trace properties using pre/post-conditions. They have been used to check safety temporal properties of programs, by associating abstract states to objects, then specifying which methods can be called at each state and how they make the state evolve. In [Bierhoff and Aldrich 2007; Bierhoff et al. 2009] they are combined with aliasing information to give a sound and modular analysis of API usage protocols, and to check a specification for the iterator module similar to ours, yet based on a single object. Multi-object properties, like the iterator one, can be captured by an extension of tpestates using tracematches to specify intensional properties [Naeem and Lhoták 2008].

Session types [Honda et al. 1998; Vasconcelos 2012] are another way of enforcing interaction patterns between interacting program components. Contracts have also been studied in the context of session types to check for safety trace-like properties [Gommerstadt et al. 2018]. Properties of interest are then checked by monitors—wrappers that run alongside the code—reminiscent of our wrappers emitting events. The main difference being then that monitors run dynamic checks (without changing the results of the instrumented program), while our wrappers only expose the interactions between library and context by emitting events, the trace properties of interest being proved statically.

Linear types are used in [DeLine and Fähndrich 2001] to develop Vault, a programming language used to design device drivers, where resource management protocols can be specified explicitly using annotations in the source code. An automatic analysis has then been developed in [Igarashi and Kobayashi 2002]. One a more foundational level, core calculus based on linear types have been studied to support language features related to resource management, such as strong updates [Morisett et al. 2005] or the mechanisms for safe memory management in low-level systems code used in the Cyclone language [Fluet et al. 2006].

System F° [Mazurak et al. 2010], an extension of System F with linear types, is shown to be expressive enough to encode as a F° type any protocol representable as a DFA. The authors show a parametricity result for the language, and show in particular the example of a file interface similar to ours. The operational semantics of the language also uses traces to track the flow of values through the program (when they are passed to a function, and when they are used). The authors show that their system can encode tpestate-like properties, by proving (on paper) that any context using a value of a type encoding a protocol (as a DFA) will indeed always use it in a way that satisfies the corresponding DFA. They however do not exhibit a systematic proof method for establishing results similar to our “free theorems”, and the expressivity of their type system is limited compared to the expressive separation logic that we consider in this work.

Type and effect systems have been used in [Skalka et al. 2008] to infer resource usage, represented by an LTS, and combined with model checking to verify trace properties of programs. Such systems have been applied to Featherweight Java in [Skalka 2008], where challenges coming from object orientation like inheritance and dynamic dispatch are tackled.

Higher-order model checking was used to provide a sound and complete resource analysis, using higher-order recursion scheme model checking for a fragment of the μ -calculus [Kobayashi 2009].

The Mezzo programming language [Balabonski et al. 2016; Pottier and Protzenko 2015] uses a type system based on ideas from Separation Logic, and as such is able to express library specifications in a style similar to ours. The work on Mezzo however focuses on the design and expressivity of its type system, and does not study which “free theorems” could be deduced from its types.

Such approaches based on type systems have been typically designed to support automated static verification and thus trade off expressiveness for automation. In this work we have made the opposite trade-off and focused on being able to capture expressive trace properties. As seen, we can specify non-regular properties (§5.3, the language is visibly pushdown [Alur and Madhusudan

2004]), others that rely on tracking an unbounded number of objects (§5.2, where we need to track all valid iterators), and we can even go beyond context-free languages (§5.4, the language requires an order-2 pushdown automaton [Maslov 1976]). In an orthogonal direction, working with a logic with quantification we can specify traces from infinite alphabets of events (§5.2, §5.3, and §5.4).

Approaches based on program logics. We have mentioned in the introduction a number of approaches relying on program logics to establish temporal trace properties. These approaches typically verify trace-oriented specifications for library implementations; in contrast, our work establishes temporal trace properties from separation logic *specifications* that are not inherently temporal.

We additionally mention here the F7 and F* programming languages, which provide another technique for reasoning about trace properties [Bengtson et al. 2011; Swamy et al. 2013]. The technique is primarily aimed at verifying cryptographic primitives, but has also been applied to access control policies about interactions between a client and resources managed through libraries [Borgström et al. 2011]. It is based on extending the base programming language with a primitive for assuming that a given formula holds and an assert primitive that fails if a given formula does not follow from all previously assumed formulas. Access control policies are encoded by inserting appropriate assume and assert statements and proving that no assert can fail. In contrast to our work, the F7 and F* approach does not establish a formal connection between the inserted assume/assert statements and the property enforced on the execution. It is also non-local in that any assume statement can introduce a contradiction and break adequacy.

Finally, we mention the more recent work on using interaction trees [Xia et al. 2019] in combination with the VST framework [Cao et al. 2018] to verify the trace behavior of a networked server [Koh et al. 2019]. In that work, interaction trees are used to specify traces of behaviors for C programs, verified using Separation Logic, where the trace of a given program records external calls made by that program. Similarly to our `tracelnv` predicate (parameterized by a set of traces), the authors rely on an `ITree` Separation Logic predicate (parameterized by an interaction tree) to specify the trace behavior of a given program. The authors also establish a form of linearizability property, as a refinement between interaction trees specifying the network behavior of the web server. The concrete interaction tree specifying a network behavior with possibly interleaved events is shown to refine a simpler interaction tree where network events have been linearized. However, since this connection is made purely at the level of interaction trees, it is only superficially related to our result of Section 6, where we relate linearizability as a trace property and logically atomic triples which do not involve traces.

7 CONCLUSION

We have presented a formal approach for deriving free theorems from Separation Logic specifications. We have focused on free theorems expressing trace properties on the interaction between clients verified against the specified library and the library itself. Since our main goal has been to establish a theoretical foundation relating specifications and trace properties, we focused on expressiveness rather than automation.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments and suggestions. We also thank Jérémy Ledent for answering the third author's questions regarding linearizability. This work was supported in part by a Villum Investigator grant (no. 25804), Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation. Guilhem Jaber was supported by the EPSRC grant EP/L022478/1. Nikos Tzevelekos was supported by the RAEng Research Fellowship 10216/111.

REFERENCES

- Rajeev Alur and P. Madhusudan. 2004. Visibly pushdown languages. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*. 202–211. <https://doi.org/10.1145/1007352.1007390>
- Thibaut Balabonski, François Pottier, and Jonathan Protzenko. 2016. The Design and Formalization of Mezzo, a Permission-Based Programming Language. *ACM Trans. Program. Lang. Syst.* 38, 4, Article 14 (Aug. 2016), 94 pages. <https://doi.org/10.1145/2837022>
- Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. 2011. Refinement Types for Secure Implementations. *ACM Trans. Program. Lang. Syst.* 33, 2 (Feb. 2011), 8:1–8:45. <https://doi.org/10.1145/1890028.1890031>
- Kevin Bierhoff and Jonathan Aldrich. 2007. Modular Typestate Checking of Aliased Objects. In *Proceedings of the 22nd Annual ACM SIGPLAN Conference on Object-oriented Programming Systems and Applications (OOPSLA'07)*. ACM, 301–320. <https://doi.org/10.1145/1297027.1297050>
- Kevin Bierhoff, Nels E Beckman, and Jonathan Aldrich. 2009. Practical API Protocol Checking with Access Permissions. In *Proceedings of the 23rd European Conference on Object-Oriented Programming (ECOOP'09)*. Springer-Verlag, 195–219. https://doi.org/10.1007/978-3-642-03013-0_10
- Bodil Biering, Lars Birkedal, and Noah Torp-Smith. 2007. BI-hyperdoctrines, Higher-order Separation Logic, and Abstraction. *ACM Trans. Program. Lang. Syst.* 29, 5, Article 24 (Aug. 2007). <https://doi.org/10.1145/1275497.1275499>
- Lars Birkedal and Aleš Bizjak. 2020. Lecture Notes on Iris: Higher-Order Concurrent Separation Log. (2020). <http://iris-project.org/tutorial-pdfs/iris-lecture-notes.pdf>
- Johannes Borgström, Andrew D. Gordon, and Riccardo Pucella. 2011. Roles, Stacks, Histories: A Triple for Hoare. *Journal of Functional Programming* 21 (2011), 159–207. Issue 02. <https://doi.org/10.1017/S0956796810000134>
- Qinxiang Cao, Lennart Beringer, Samuel Gruetter, Josiah Dodds, and Andrew W. Appel. 2018. VST-Floyd: A Separation Logic Tool to Verify Correctness of C Programs. *Journal of Automated Reasoning* 61, 1-4 (2018), 367–422. <https://doi.org/10.1007/s10817-018-9457-5>
- Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *Proceedings of ECOOP*. https://doi.org/10.1007/978-3-662-44202-9_9
- Robert DeLine and Manuel Fähndrich. 2001. Enforcing High-level Protocols in Low-level Software. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation (PLDI '01)*. ACM, 59–69. <https://doi.org/10.1145/381694.378811>
- Robert DeLine and Manuel Fähndrich. 2004. Typestates for Objects. In *Object-Oriented Programming: Proceedings of the 18th European Conference on Object-Oriented Programming (ECOOP'04)*. Springer-Verlag, 465. https://doi.org/10.1007/978-3-540-24851-4_21
- Matthew Fluet, Greg Morrisett, and Amal Ahmed. 2006. Linear Regions Are All You Need. In *Programming Languages and Systems*, Peter Sestoft (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 7–21. https://doi.org/10.1007/11693024_2
- Ming Fu, Yong Li, Xinyu Feng, Zhong Shao, and Yu Zhang. 2010. Reasoning about Optimistic Concurrency Using a Program Logic for History. In *In Proceedings of CONCUR*. https://doi.org/10.1007/978-3-642-15375-4_27
- Philippa Gardner, Gian Ntzik, and Adam Wright. 2014. Local Reasoning for the POSIX File System. In *ESOP*. 169–188. https://doi.org/10.1007/978-3-642-54833-8_10
- Hannah Gommerstadt, Limin Jia, and Frank Pfenning. 2018. Session-Typed Concurrent Contracts. In *Programming Languages and Systems*, Amal Ahmed (Ed.). Springer International Publishing, Cham, 771–798. https://doi.org/10.1007/978-3-319-89884-1_27
- Alexey Gotsman, Noam Rinetky, and Hongseok Yang. 2013. Verifying Concurrent Memory Reclamation Algorithms with Grace. In *Proceedings of ESOP*. https://doi.org/10.1007/978-3-642-37036-6_15
- Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: a correctness condition for concurrent objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (July 1990), 463–492. <https://doi.org/10.1145/78969.78972>
- Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. 1998. Language primitives and type discipline for structured communication-based programming. In *Programming Languages and Systems*, Chris Hankin (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 122–138. <https://doi.org/10.1007/BFb0053567>
- Atsushi Igarashi and Naoki Kobayashi. 2002. Resource Usage Analysis. In *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '02)*. ACM, 331–342. <https://doi.org/10.1145/503272.503303>
- Samin S. Ishtiaq and Peter W. O'Hearn. 2001. BI As an Assertion Language for Mutable Data Structures. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'01)*. ACM, 14–26. <https://doi.org/10.1145/360204.375719>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015a. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*. 637–650. <https://doi.org/10.1145/2676726.2676980>

- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015b. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*. 637–650. <https://doi.org/10.1145/2676726.2676980>
- Naoki Kobayashi. 2009. Types and Higher-order Recursion Schemes for Verification of Higher-order Programs. In *Proceedings of the 36th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '09)*. ACM, 416–428. <https://doi.org/10.1145/1480881.1480933>
- Nicolas Koh, Yao Li, Yishuai Li, Li-yao Xia, Lennart Beringer, Wolf Honoré, William Mansky, Benjamin C. Pierce, and Steve Zdancewic. 2019. From C to Interaction Trees: Specifying, Verifying, and Testing a Networked Server. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs (Cascais, Portugal) (CPP 2019)*. Association for Computing Machinery, New York, NY, USA, 234–248. <https://doi.org/10.1145/3293880.3294106>
- Neelakantan R. Krishnaswami, Jonathan Aldrich, Lars Birkedal, Kasper Svendsen, and Alexandre Buisse. 2009. Design Patterns in Separation Logic. In *Proceedings of the 4th International Workshop on Types in Language Design and Implementation (TLDI'09)*. ACM, 105–116. <https://doi.org/10.1145/1481861.1481874>
- A. N. Maslov. 1976. Multilevel stack automata. *Problems Inform. Transmission* 12, 1 (1976), 38–42.
- Karl Mazurak, Jianzhou Zhao, and Steve Zdancewic. 2010. Lightweight Linear Types in System F^{*}. In *Proceedings of the 5th ACM SIGPLAN Workshop on Types in Language Design and Implementation (Madrid, Spain) (TLDI '10)*. Association for Computing Machinery, New York, NY, USA, 77–88. <https://doi.org/10.1145/1708016.1708027>
- Greg Morrisett, Amal Ahmed, and Matthew Fluet. 2005. L3: A Linear Language with Locations. In *Typed Lambda Calculi and Applications*, Pawel Urzyczyn (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 293–307. https://doi.org/10.1007/11417170_22
- Nomair A Naeem and Ondrej Lhoták. 2008. Typestate-like analysis of multiple interacting objects. *OOPSLA'08: Proceedings of Object-Oriented Programming, Systems, Languages and Applications* (2008). <https://doi.org/10.1145/1449764.1449792>
- Matthew Parkinson and Gavin Bierman. 2005. Separation Logic and Abstraction. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'05)*. ACM, 247–258. <https://doi.org/10.1145/1040305.1040326>
- François Pottier and Jonathan Protzenko. 2015. A few lessons from the Mezzo project. In *Summit on Advances in Programming Languages (SNAPL) (Leibniz International Proceedings in Informatics, Vol. 32)*. Asilomar, United States. <https://doi.org/10.4230/LIPICs.SNAPL.2015.221>
- J.C. Reynolds and G.D. Plotkin. 1993. On Functors Expressible in the Polymorphic Typed Lambda Calculus. *Information and Computation* 105, 1 (1993), 1–29. <https://doi.org/10.1006/inco.1993.1037>
- John C. Reynolds. 1983. Types, Abstraction and Parametric Polymorphism. In *Information Processing 83*. 513–523. http://www.cse.chalmers.se/edu/year/2010/course/DAT140_Types/Reynolds_typesabpara.pdf
- John C Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS'02)*. IEEE, 55–74. <https://doi.org/10.1109/LICS.2002.1029817>
- Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2015. Specifying and Verifying Concurrent Algorithms with Histories and Subjectivity. In *Proceedings of ESOP*. https://doi.org/10.1007/978-3-662-46669-8_14
- Christian Skalka. 2008. Types and trace effects for object orientation. *Higher-Order and Symbolic Computation* 21, 3 (2008), 239–282. <https://doi.org/10.1007/s10990-008-9032-6>
- Christian Skalka, Scott Smith, and David Van Horn. 2008. Types and trace effects of higher order programs. *Journal of Functional Programming* 18, 02 (2008), 179–249. <https://doi.org/10.1017/S0956796807006466>
- Robert E Strom and Shaula Yemini. 1986. Typestate: A Programming Language Concept for Enhancing Software Reliability. *IEEE Trans. Softw. Eng.* 12, 1 (Jan. 1986), 157–171. <https://doi.org/10.1109/TSE.1986.6312929>
- Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. 2013. Secure Distributed Programming with Value-Dependent Types. *Journal of Functional Programming* 23, 4 (2013), 402–451. <https://doi.org/10.1017/S0956796813000142>
- Vasco T. Vasconcelos. 2012. Fundamentals of session types. *Information and Computation* 217 (2012), 52–70. <https://doi.org/10.1016/j.ic.2012.05.002>
- Philip Wadler. 1989. Theorems for free! 347–359. <https://doi.org/10.1145/99370.99404>
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2019. Interaction Trees: Representing Recursive and Impure Programs in Coq. *Proc. ACM Program. Lang.* 4, POPL, Article 51 (Dec. 2019), 32 pages. <https://doi.org/10.1145/3371119>