

Views: Compositional Reasoning for Concurrent Programs

Thomas Dinsdale-Young[†], Lars Birkedal^{*}, Philippa Gardner[†], Matthew Parkinson[‡], Hongseok Yang[§]

[†]Imperial College London

^{*}IT University of Copenhagen

[‡]Microsoft Research Cambridge

[§]University of Oxford

Abstract—We present a framework for reasoning compositionally about concurrent programs. At its core is the notion of a *view*: an abstraction of the state that takes account of the possible interference due to other threads. Threads’ views are composable, and an update to the state by one thread must preserve the views of other threads. We prove soundness for our framework, and demonstrate its utility by studying examples. In particular, we show that several variants of concurrent separation logic, type systems, and concurrent abstract predicates can all be seen as instantiations of our framework. Soundness for particular instantiations simply follows from the soundness of our framework.

I. INTRODUCTION

There has been a recent flurry of research on models, logics and type systems for reasoning modularly about programs with dynamically allocated mutable state. There have been type systems extended with linear types [1], [2], [3] and related capability systems [4] that enforce a mixture of local and global properties. In particular, they allow for so-called strong update, where an assignment to a variable can change the type of the variable. On the program logic side there have been a series of logics extending separation logic [5] to reason about various notions of sharing. In the sequential setting, the antiframe rule [6], [7] and higher-order frame rules [8] allow for the preservation of selected global properties. For concurrent languages, SAGL [9], RGSep [10], logics for storable locks [11], [12], [13], deny-guarantee [14] and concurrent abstract predicates [15], [16] have been proposed to reason about shared usages of mutable state. There has even been a combination of separation logic with a type system [17].

We present a framework for reasoning compositionally about concurrent programs, which unifies all these styles of reasoning. Our fundamental idea is that threads have different *views* of the machine. Intuitively, a thread’s view consists of information about the current state of the machine, the *right* of the thread to modify the state as long as the environment’s view is stable (invariant) with respect to such changes, and the thread’s *right* to the stability of its own view with respect to changes being made by the environment. Our framework captures the subtle interplay between a thread’s view and the environment’s view, and allows us to express when views are compatible and hence can be *composed*. This compositional structure of views has close connections with the structure of BI-algebras [18], [19], from which separation logic evolved.

One may worry whether our views framework is more complicated than the earlier concrete models that it unifies. As we shall see, that is not the case. Indeed, because our views framework abstracts and distills the essence of compositional concurrent reasoning, the framework itself is arguably a lot simpler than most of the recent advanced concrete models.

Examples Views enable us to reason about a wide range of type systems and program logics. For example, with a standard type system, a view of a thread might be: a variable x has type τ , the memory is well-typed, and the thread has the right to make any type-preserving changes to the memory (including allocation of new memory cells). The environment must hold a compatible view: it must agree with the typing of x . The environment’s view also permits it to make type-preserving changes. The thread and environment’s views are stable with respect to such changes, since the changes only make the typing information grow monotonically. We are also able to handle non-monotonic type systems, in which updates may change the type of variables, by simply requiring more separation between the thread and the environment’s view of the variables and memory.

Another example is disjoint concurrent separation logic where a view might be: location 10 has value 4, and the thread owns location 10. Since the thread owns location 10, the environment cannot own it. Consequently, the thread has the right to change it, but the environment does not, and cannot depend upon it remaining unchanged. Conversely, the thread has no right to change any locations in the environment’s view, as they must be separate from the thread’s view. The thread and environment’s views are thus stable, and two views can be composed provided they are disjoint. Following the original work on concurrent separation logic, there have been several proposals for reasoning about shared mutable state, for example, separation logic with permissions, SAGL, and deny-guarantee. We handle these examples by simply allowing more sharing in a thread’s view.

Views A *view* of a thread provides a *partial, abstract* description of the concrete state of a machine. It is partial in that it only describes the state accessed by the thread, such as the cell with location 10. It is abstract in that the verifier can use any additional information to help with their reasoning, such as types, ghost state or permissions. Such additional information has no representation in the concrete

state but is a useful fiction for the verifier. Since views describe partial state, an essential property of views is that they are *compositional*. This means that the views of two threads must either be disjoint (e.g., disjoint heaps in separation logic) or must agree on any overlap (e.g., consistent on their types).

The verifier has the flexibility to change the view, using an *angelic preorder*, which does not change the underlying concrete state but does change the abstract state: for example, types can change using a monotonic subtyping relation and permissions can be reassigned across abstract states.

Program Logic We provide a *program logic* for a simple concurrent programming language, parameterised by a set of atomic commands. The views provide the pre- and postconditions of the commands. The logic is parameterised by axioms for the atomic commands and the structural rules include a frame rule and an angelic consequence rule. We present several example program logics using this framework, including weak and strong type theories, several concurrent separation logics, and a combination of separation logic with type theory.

Soundness We provide a *soundness* result, relating our program logic to the underlying operational semantics. The operational semantics is defined on concrete machine states, parameterised by an interpretation of the atomic commands. To relate the program logic to the operational semantics, we require that any abstract view can be *erased* to a set of concrete machine states. This erasure function satisfies two properties: *monotonicity*, declaring that the angelic preorder is erased to the subset relation on sets of machine states; and *atomic soundness*, declaring that the view axiomatisation of atomic commands fits with the interpretation of the atomic commands in the operational semantics. With these properties, we prove a general soundness result, using a technique first introduced by Vafeiadis [20]. To prove soundness for specific type theories and separation logics, it is then enough to display an erasure function satisfying monotonicity and atomic soundness.

Interference Using the constructions presented so far, we can reason about several type theories and concurrent separation logics. These examples have the common pattern that race conditions can be prevented by separation (composition) and abstraction. However, many examples have finer race conditions than this, with a subtle interplay between the effect of the current thread and the effect of the environment. To reason about such behaviour, we need a notion of *interference*: the interference of the environment on a thread’s view, and the interference of the thread on the environment’s view.

We define an *interference operator* R on views, with simple properties capturing the subtle interplay between the thread’s and environment’s view. The key properties are that R provides the minimal *stabilisation* of a view with respect to the environment, and that it is well-behaved with respect to *composition*: that is, interference on the composition of views does no more than the interference on the component views. For Kripke models of type theories, the interference operator corresponds to the future world relation of the Kripke model, and thus it allows all type-preserving changes. With disjoint concurrent separation logic, the interference operation is the

identity since the environment cannot change a thread’s portion of the heap. With complex concurrent separation logics, the interference relation is constructed from the permissions the environment has to change a thread’s view.

The interplay between compositionality and interference, plus stability, leads to the surprising property that the subset of stable views satisfies all the properties of a view model. Thus, interference operators provide a way of constructing views which allow more fine-grained races between threads. Interference operators do not change the underlying theory, and our general soundness result still applies. We thus believe that views capture the fundamental interplay between composition and interference that is inherent in concurrency.

II. PROGRAM LOGIC

We present a program logic for a concurrent programming language, parameterised by atomic commands, whose assertions are views. We give some simple illustrative examples here, and more complicated examples in section IV-A.

Our language is built from standard composite commands, and parameterised by a set of atomic commands.

Parameter 1 (Atomic Commands). *Assume a countable set of (syntactic) atomic commands* Atom , *ranged over by* a .

Definition 1 (Language Syntax). *The set of (syntactic) commands*, Comm , *ranged over by* C , *is defined by the following grammar*: $C ::= a \mid \text{skip} \mid C; C \mid C + C \mid C \parallel C \mid C^*$.

Views form assertions in our program logic. Their properties are axiomatised in our definition of a view model.

Definition 2 (View Model). *A view model consists of a set* View , *ranged over by* p, q, r , *with the following structure*:

- a commutative monoid $(\text{View}, *, u)$,
- a complete join-semilattice $(\text{View}, \models, \bigvee, \perp)$, and it is subject to the following law:
- Distributivity: *the $*$ operator distributes over \bigvee : that is,*

$$p * \bigvee \{q_i\}_{i \in I} = \bigvee \{p * q_i\}_{i \in I}.$$

Lemma 1. *With respect to \models , the $*$ operator is monotone in both of its arguments.*

Remark. The definition of a view model implies additional structure that is not part of the definition. Since views form a complete join-semilattice, they form a complete lattice:

$$\bigwedge \{p_i\}_{i \in I} \stackrel{\text{def}}{=} \bigvee \{q \mid \forall i. q \models p_i\}.$$

Also, the distributivity property implies that the composition $*$ is a residuated monoid, with $q * -$ having right adjoint $q * - : q * r \stackrel{\text{def}}{=} \bigvee \{p \mid q * p \models r\}$. The adjunction $q * p \models r \iff p \models q * r$ is trivial from left to right, and follows from distributivity and monotonicity from right to left. The corresponding distributivity property may not hold for \wedge , and so a view model is not a Heyting algebra (it does not necessarily have an implication, \rightarrow). With implication, a view model is a BI-algebra [18]. Conversely, a BI-algebra is a view model if it is complete as a lattice.

An important class of BI-algebras arises from *separation algebras*, a concept introduced by Calcagno, O’Hearn and Yang [21] to generalise separation logic. We use a variant of separation algebras with multiple units [22], [23] and no cancellativity as a common source of view models.

Definition 3 (Separation Algebra). A separation algebra $(\mathcal{M}, \bullet, I)$ is a partial, commutative monoid with multiple units. Namely, it is a set \mathcal{M} equipped with a partial binary operator $\bullet : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ and designated unit set $I \subseteq \mathcal{M}$ satisfying:

- Commutativity: $m_1 \bullet m_2 = m_2 \bullet m_1$ when either is defined;
- Associativity: $m_1 \bullet (m_2 \bullet m_3) = (m_1 \bullet m_2) \bullet m_3$ when either is defined;
- Existence of Unit: for all $m \in \mathcal{M}$ there exists $i \in I$ such that $i \bullet m = m$; and
- Minimality of Unit: for all $m \in \mathcal{M}$ and $i \in I$, if $i \bullet m$ is defined then $i \bullet m = m$.

Definition 4 (Separation View Model). Each separation algebra $(\mathcal{M}, \bullet, I)$ induces a separation view model $(\mathcal{P}(\mathcal{M}), *, I, \subseteq, \bigcup, \emptyset)$, where

$$p_1 * p_2 \stackrel{\text{def}}{=} \{m_1 \bullet m_2 \mid m_1 \in p_1, m_2 \in p_2\}.$$

Views provide an partial, abstract description of the concrete state of the machine. The angelic preorder enables the verifier to change the instrumentation (e.g. weakening the type of a variable to a supertype, reassigning permissions), as long as the underlying concrete state remains the same.

Definition 5 (Angelic Preorder). Given a view model, View , an angelic preorder $\prec \subseteq \text{View} \times \text{View}$ is a reflexive, transitive relation on View that is subject to the following laws:

- Locality: for all $p, q, r \in \text{View}$, if $p \prec q$ then $p * r \prec q * r$
- \models -closure: $\models \subseteq \prec$
- \bigvee -closure: if $p_i \prec q$ for all $i \in I$ then $\bigvee \{p_i\}_{i \in I} \prec q$.

For any view model, \models is always an angelic preorder (by Lemma 1), and is inherently minimal with respect to all possible angelic preorders on the view model.

We define a *program logic* for our programming language, in which views provide the pre- and postconditions of the commands. We declare axioms for the atomic commands, and rules for composition (frame and disjoint concurrency) and angelic consequence.

Parameter 2 (View Model). Assume a view model

$$(\text{View}, *, u, \models, \bigvee, \perp).$$

Parameter 3 (Angelic Preorder). Assume an angelic preorder on View , \prec .

Parameter 4 (Axiomatisation). Assume a set of axioms $\text{Axiom} \subseteq \text{View} \times \text{Atom} \times \text{View}$.

Definition 6 (Program Logic). The program logic consists of judgements of the form $\{p\} C \{q\}$, where $p, q \in \text{View}$ provide

$$\begin{array}{c} \frac{(p, a, q) \in \text{Axiom}}{\{p\} a \{q\}} \quad \frac{p \prec p' \quad \{p'\} C \{q'\} \quad q' \prec q}{\{p\} C \{q\}} \\ \frac{\{p\} C \{q\}}{\{p * r\} C \{q * r\}} \quad \frac{\forall i \in I. \{p_i\} C \{q\}}{\{\bigvee \{p_i\}_{i \in I}\} C \{q\}} \quad \frac{}{\{p\} \text{skip} \{p\}} \\ \frac{\{p\} C_1 \{r\} \quad \{r\} C_2 \{q\}}{\{p\} C_1; C_2 \{q\}} \quad \frac{\{p\} C_1 \{q\} \quad \{p\} C_2 \{q\}}{\{p\} C_1 + C_2 \{q\}} \\ \frac{\{p_1\} C_1 \{q_1\} \quad \{p_2\} C_2 \{q_2\}}{\{p_1 * p_2\} C_1 \parallel C_2 \{q_1 * q_2\}} \quad \frac{\{p\} C \{p\}}{\{p\} C^* \{p\}} \end{array}$$

Fig. 1. The Proof Rules for our Program Logic.

the precondition and postcondition of command $C \in \text{Comm}$. The proof rules for these judgements are given in Fig. 1.

The intended semantics of $\{p\} C \{q\}$ is that if the program C is run to termination from an initial state that is described by the view p , then the resulting state will be described by the view q . This is a partial correctness interpretation: the judgements say nothing about non-terminating executions.

Most of the proof rules are standard rules from disjoint concurrent separation logic. They include the frame rule, which captures the intuition that a program’s view can be extended with a composable view, and the disjoint concurrency rule, which allows the views of two threads to be composed. The axiom rules are self-explanatory. The rule for the angelic preorder is a modified version of the standard rule of consequence, which it subsumes. It enables the verifier to modify the instrumentation without performing a machine-level update.

A. Examples

We illustrate our program logic with some simple examples. These examples build on a language whose atomic primitives are heap update commands.

Definition 7 (Atomic Heap Commands). Assume a set of variable names Var , ranged over by x and y , and a set of values Val , ranged over by v , of which a subset $\text{Loc} \subseteq \text{Val}$ represents heap addresses, ranged over by l . The syntax of atomic heap commands, Atom_H , is defined by the grammar: $a ::= x := y \mid [x] := v \mid [x] := y \mid x := [y] \mid x := \text{ref } y$.

The first of these commands loads the contents of variable y into variable x . The second and third load, respectively, a literal value v and the contents of y into the heap location whose address is in x . The fourth loads the contents of the heap location whose address is in y into x . The final command allocates a new heap location with initial contents from y and loads its address into x .

1) *Disjoint Concurrent Separation Logic*: Judgements of disjoint concurrent separation logic are, as in the views framework, triples of the form $\{p\} C \{q\}$. Abstractly, the state is treated as a resource, which is divided up by individual variables and heap addresses. Thus, p and q describe resources, which hold information about part of the state. Formally, p and q are views from the separation view model induced by the

separation algebra $(\mathcal{M}_{\text{DCSL}}, \uplus, \{\emptyset\})$, where

$$\mathcal{M}_{\text{DCSL}} \stackrel{\text{def}}{=} (\text{Var} \uplus \text{Loc}) \rightarrow_{\text{fin}} \text{Val}.$$

That is, $\mathcal{M}_{\text{DCSL}}$ is the set of finite partial functions from variables and heap addresses to values, with the partial monoid operation given by the union of partial functions with disjoint domains \uplus , and the unit consisting of only the partial function with the empty domain, \emptyset .

Elements of $\mathcal{M}_{\text{DCSL}}$ declare ownership of the variables and heap addresses that belong to their domains, as well as defining their values. Significantly, they do not declare information about parts of the state which are not owned. Views $p, q \in \mathcal{P}(\mathcal{M}_{\text{DCSL}})$ are sets of these abstract states.

The view $x \Rightarrow v$ denotes the singleton set of the partial function mapping variable x to value v , and $x \Rightarrow _$ denotes the set of all partial functions that only map variable x to a value. Similarly, the views $l \mapsto v$ and $l \mapsto _$ map heap address l to v or any value respectively. The view $\exists v. p(v)$ is the (infinite) join of $p(v)$ for all values of v .

Separation logic does not allow angelic update, and so the angelic preorder is simply \vDash which is itself \subseteq .

The axiomatisation for separation logic is given by the schema:

$$\begin{aligned} & \{x \Rightarrow _ * y \Rightarrow v\} x := y \{x \Rightarrow v * y \Rightarrow v\} \\ & \{x \Rightarrow l * l \mapsto _ \} [x] := v \{x \Rightarrow l * l \mapsto v\} \\ & \{x \Rightarrow l * l \mapsto _ * y \Rightarrow v\} [x] := y \{x \Rightarrow l * l \mapsto v * y \Rightarrow v\} \\ & \{x \Rightarrow l * l \mapsto v * y \Rightarrow _ \} y := [x] \{x \Rightarrow l * l \mapsto v * y \Rightarrow v\} \\ & \{x \Rightarrow _ * y \Rightarrow v\} x := \text{ref } y \{\exists l. x \Rightarrow l * l \mapsto v * y \Rightarrow v\} \end{aligned}$$

2) *Weak-update Type System:* Consider a simple type system for heap update, which types variables and heap cells with the set of types Type , ranged over by τ , and defined by:

$$\tau ::= \text{val} \mid \text{ref } \tau.$$

The type val indicates that a variable or heap cell contains some unspecified value, while the type $\text{ref } \tau$ indicates that it contains the address of a heap cell whose contents is typed as τ . A typing context $\Gamma : \text{Var} \rightarrow \text{Type}$ is a partial function which assigns types to variables. In a weak update type system, the types of variables are fixed, and all assignments must preserve the typing. For the heap update language, we define such a type system by the typing rules:

$$\begin{array}{c} \frac{}{x : \tau, y : \tau \vdash x := y} \quad \frac{}{x : \text{ref val} \vdash [x] := v} \\ \frac{}{x : \text{ref } \tau, y : \tau \vdash [x] := y} \quad \frac{}{x : \text{ref } \tau, y : \tau \vdash y := [x]} \\ \frac{}{x : \text{ref } \tau, y : \tau \vdash x := \text{ref } y} \quad \frac{}{\Gamma \vdash \text{skip}} \\ \frac{\Gamma \vdash C_1 \quad \Gamma \vdash C_2 \quad \text{op} \in \{;, +, \parallel\}}{\Gamma \vdash C_1 \text{ op } C_2} \quad \frac{}{\Gamma \vdash C} \quad \frac{}{\Gamma, \Gamma' \vdash C} \end{array}$$

The intended meaning of a typing judgement $\Gamma \vdash C$ is that, whenever the program C is executed from an initial state in

which the variables can be typed according to Γ , the program does not fault and results in a state in which the variables can still be typed according to Γ .

We can form a separation algebra $(\mathcal{M}_{\text{WTS}}, \cup, \{\emptyset\})$ of typing contexts, where $\mathcal{M}_{\text{WTS}} \stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Type}$, the monoid operator is union (of relations, restricted to when the result is a function) and the unit is the set of the partial function with the empty domain. In this separation algebra, two typing contexts can be combined when they agree on the types of all variables that are common to both: this is the implicit meaning of the “ \cup ” operator in the typing rules above. This induces a view model in which the views are sets of typing contexts.

We fit the type system into the views framework by interpreting the judgement $\Gamma \vdash C$ as $\{\{\Gamma\}\} C \{\{\Gamma\}\}$. By taking the first five rules of the type system as axioms, and \vDash as the angelic partial order, we obtain an instance of the views framework. The remaining rules of the type system are then easily justified by the proof rules of the views program logic. The most interesting of these is the last, the weakening rule, which is an instance of the frame rule, with the frame $\{\Gamma'\}$.

3) *Strong-update Type System:* In the previous example, every command preserves the types of variables: they are weak updates. We now consider a type system in which strong updates, which may change the type of variables, are permitted. In this system, each thread has its own local variables, which allows the types of the variables to be updated, since they are not shared with any other threads. The types of heap locations cannot be updated, however, since multiple threads may have aliases to the same location.

Typing judgements here use the same typing contexts as in the weak-update example. Since type-changing updates are permitted, judgements have input and output typing contexts. The type system is defined by the following typing axioms (the rules are derived directly from the views framework):

$$\begin{array}{c} \frac{}{x : \tau_0, y : \tau \vdash x := y \dashv x : \tau, y : \tau} \\ \frac{}{x : \text{ref val} \vdash [x] := v \dashv x : \text{ref val}} \\ \frac{}{x : \text{ref } \tau, y : \tau \vdash [x] := y \dashv x : \text{ref } \tau, y : \tau} \\ \frac{}{x : \text{ref } \tau, y : \tau_0 \vdash y := [x] \dashv x : \text{ref } \tau, y : \tau} \\ \frac{}{x : \tau_0, y : \tau \vdash x := \text{ref } y \dashv x : \text{ref } \tau, y : \tau} \end{array}$$

The interpretation of a typing judgement $\Gamma \vdash C \dashv \Gamma'$ is that, when C is executed to termination from an initial state satisfying the typing context Γ , it will not fault and will result in a state satisfying the typing context Γ' .

We define the views model by considering an appropriate separation algebra. This time, the separation algebra is $(\mathcal{M}_{\text{STS}}, \uplus, \{\emptyset\})$, where $\mathcal{M}_{\text{STS}} \stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Type}$. Rather than the union, as for weak update, we take the disjoint union of partial functions, as for separation logic. This enforces the intended discipline of ownership: threads may modify the types and values of variables that they own, but have no knowledge or rights to any other variables. Again, this

separation algebra induces a model in which views are sets of typing contexts.

We interpret the type judgement $\Gamma \vdash C \dashv \Gamma'$ as $\{\{\Gamma\}\} C \{\{\Gamma'\}\}$ within the views framework. As before, we take the typing axioms as the axioms of our framework instance. By choosing the angelic preorder \vDash , the standard rules of this type system can be seen simply as instances of the rules from the views framework.

We can add the following subtyping rule to the system:

$$\frac{\Gamma, x : \text{val} \vdash C \dashv \Gamma'}{\Gamma, x : \text{ref } \tau \vdash C \dashv \Gamma'}$$

We can justify this rule by choosing an angelic preorder such that $\{\Gamma, x : \text{ref } \tau\} \prec \{\Gamma, x : \text{val}\}$. A suitable choice is

$$p \prec q \stackrel{\text{def}}{\iff} \forall \Gamma \in p. \exists \Gamma' \in q. \text{dom}(\Gamma) = \text{dom}(\Gamma') \text{ and } \forall x \in \text{dom}(\Gamma). \Gamma'(x) \in \{\text{val}, \Gamma(x)\}$$

which satisfies the requirements of an angelic preorder.

4) *Typed Separation Logic*: Separation logic has been used to reason about programming languages with types, such as Java [24]. However, separation-logic reasoning typically ignores the types, although the work of Tan *et al.* [17] is a notable exception. Since type information is shared and global, and separation-logic reasoning is local, it previously seemed difficult to integrate the two systems. However, with our framework, it is easy. The system we present here only allows weak updates of the heap. We also have an example with strong updates [25].

We can axiomatise the atomic commands by combining the axioms of the strong-update type system and separation logic:

$$\frac{\Gamma_1 \vdash a \dashv \Gamma_2 \quad \{p\} a \{q\}}{\Gamma_1 \vdash \{p\} a \{q\} \dashv \Gamma_2}$$

If the atomic command is allowed by the type system and the separation logic, then it is allowed in the combined system. We can also add axioms, such as the following, which derive from only one of the systems:

$$\frac{\Gamma_1 \vdash x := [y] \dashv \Gamma_2}{\Gamma_1 \vdash \{x \Rightarrow _ \} x := [y] \{x \Rightarrow _ \} \dashv \Gamma_2}$$

$$\frac{\{p\} x := [y] \{q\}}{x : _ \vdash \{p\} x := [y] \{q\} \dashv x : \text{val}}$$

Importantly, anything that is changed must be allowed by both systems, but the ability to access something only needs to be justified in one of the underlying systems.

To model this in the framework, we simply take the product of the two earlier models, $\text{View}_{\text{TSL}} \stackrel{\text{def}}{=} \text{View}_{\text{STL}} \times \text{View}_{\text{DCSL}}$, and lift the operators in the obvious way. We interpret the judgements $\Gamma_1 \vdash \{p\} C \{q\} \dashv \Gamma_2$ as $\{\{\Gamma_1\}, p\} C \{\{\Gamma_2\}, q\}$.

III. OPERATIONAL SEMANTICS AND SOUNDNESS

The operational semantics of our language is parametrised by a model of machine states and an interpretation of the atomic commands as state transformers on this model.

$$\frac{C_1 \xrightarrow{\alpha} C'_1}{C_1; C_2 \xrightarrow{\alpha} C'_1; C_2} \quad \frac{}{C_1 + C_2 \xrightarrow{\text{id}} C_1} \quad \frac{}{C_1 + C_2 \xrightarrow{\text{id}} C_2}$$

$$\frac{C_1 \xrightarrow{\alpha} C'_1}{C_1 \parallel C_2 \xrightarrow{\alpha} C'_1 \parallel C_2} \quad \frac{C_2 \xrightarrow{\alpha} C'_2}{C_1 \parallel C_2 \xrightarrow{\alpha} C_1 \parallel C'_2} \quad \frac{}{a \xrightarrow{[a]} \text{skip}}$$

$$\frac{}{\text{skip}; C_2 \xrightarrow{\text{id}} C_2} \quad \frac{}{\text{skip} \parallel C_2 \xrightarrow{\text{id}} C_2} \quad \frac{}{C_1 \parallel \text{skip} \xrightarrow{\text{id}} C_1}$$

$$\frac{}{C^* \xrightarrow{\text{id}} C; C^*} \quad \frac{}{C^* \xrightarrow{\text{id}} \text{skip}}$$

$$\frac{C_1 \xrightarrow{\alpha} C_2 \quad s_2 \in \alpha(s_1) \quad C_2, s_2 \rightarrow^* C_3, s_3}{C_1, s \rightarrow^* C, s} \quad \frac{}{C_1, s_1 \rightarrow^* C_3, s_3}$$

Fig. 2. The Operational semantics

Parameter 5 (Machine States). Assume a set of machine states \mathcal{S} , ranged over by s .

Definition 8 (Machine Actions). Machine Actions, ranged over by α , are defined as a set of non-deterministic state transformers $\text{Action} \stackrel{\text{def}}{=} \mathcal{S} \rightarrow \mathcal{P}(\mathcal{S})$. We use id to denote the identity action: $\text{id}(s) \stackrel{\text{def}}{=} \{s\}$.

Where necessary, we lift actions to sets of states: for $S \in \mathcal{P}(\mathcal{S})$, $\alpha(S) = \bigcup \{\alpha(s) \mid s \in S\}$. This lifting is monotone.

Parameter 6 (Interpretation of Atomic Commands). Assume a function $\llbracket - \rrbracket : \text{Atom} \rightarrow \text{Action}$ that associates each atomic command with a machine action.

From machine state s , the set of states $\llbracket a \rrbracket(s)$ is the set of possible outcomes of running the atomic command a . If the set is empty, then the command blocks. Here, we consider partial correctness, and so ignore blocking executions.

We define the operational semantics of the language using a labelled transition system. Transitions are between commands, and are labelled by the machine action that is executed in performing the transition. The action id plays a special role, since it represents the effect of computation steps in which the state is not changed.

The labelled transition system splits the control-flow aspect of execution, represented by the transitions between commands, and the state-transforming aspect of execution, represented by the labelling of the transitions. This makes it easy to reinterpret programs over a more abstract state-space, while preserving the control-flow structure, which simplifies the soundness proof of our logic.

Definition 9 (Labelled Transition System and Operational Semantics). The labelled transition relation $- \xrightarrow{-} - : \text{Comm} \times \text{Action} \times \text{Comm}$ and the multi-step operational transition relation $- , - \rightarrow^* - , - : (\text{Comm} \times \mathcal{S}) \times (\text{Comm} \times \mathcal{S})$ are defined by the rules given in Fig. 2.

We prove that our program logic is sound with respect to the operational semantics. To do this, we must relate the views (partial, abstract states) with the machine states (concrete, complete states). This is achieved by an *erasure* function.

Parameter 7 (Erasure). Assume an erasure function

$$\lfloor - \rfloor : \text{View} \rightarrow \mathcal{P}(\mathcal{S})$$

which maps views to sets of machine states, and is a join-semilattice homomorphism: $\lfloor \bigvee \{p_i\}_{i \in I} \rfloor = \bigcup \{\lfloor p_i \rfloor\}_{i \in I}$.

This erasure function must satisfy two properties: *monotonicity*, stating that the erasure of the angelic preorder is set inclusion, and *atomic soundness*, stating that the erasure of axioms corresponds to the interpretation of the atomic commands in the operational semantics.

Property 8 (Monotonicity of Erasure). The erasure function is monotone from (View, \prec) to $(\mathcal{P}(\mathcal{S}), \subseteq)$:

$$\forall p, q \in \text{View}. p \prec q \implies \lfloor p \rfloor \subseteq \lfloor q \rfloor.$$

Remark. We have already noted that \vdash is the minimal angelic preorder. For a given erasure function, there is also a well-defined notion of a maximal angelic preorder, which is the greatest angelic preorder for which monotonicity holds. This preorder, \prec_{\max} , may be defined as follows:

$$p \prec_{\max} q \stackrel{\text{def}}{\iff} \forall r. \lfloor p * r \rfloor \subseteq \lfloor q * r \rfloor.$$

It is easy to see that this satisfies the required properties, and that monotonicity of erasure will hold. It is also clear that if $p \prec q$ for some angelic preorder for which monotonicity holds, it will be the case that $p \prec_{\max} q$. This means that \prec_{\max} really is maximal with respect to angelic preorders.

To define atomic soundness, we first define an action judgement, $\alpha \Vdash \{p\}\{q\}$, which interprets α as updating a thread's view from p to q provided it preserves any environment view.

Definition 10 (Action Judgement).

$$\alpha \Vdash \{p\}\{q\} \stackrel{\text{def}}{\iff} \forall r \in \text{View}. \alpha(\lfloor p * r \rfloor) \subseteq \lfloor q * r \rfloor.$$

Lemma 2 (Basic Locality). For all $p, q, r \in \text{View}$, $\alpha \in \text{Action}$, if $\alpha \Vdash \{p\}\{q\}$ then $\alpha \Vdash \{p * r\}\{q * r\}$.

Lemma 3 (Basic \prec -closure). For all $p, p', q, q' \in \text{View}$, $\alpha \in \text{Action}$, if $p \prec p'$, $\alpha \Vdash \{p'\}\{q'\}$ and $q' \prec q$ then $\alpha \Vdash \{p\}\{q\}$.

Property 9 (Atomic Soundness). For every $(p, a, q) \in \text{Axiom}$, we have $\llbracket a \rrbracket \Vdash \{p\}\{q\}$.

The following lemma gives an alternative characterisation of the action judgement for separation view models, which is typically simpler to check.

Lemma 4. If View is a separation view model induced by $(\mathcal{M}, \bullet, I)$, then

$$\alpha \Vdash \{p\}\{q\} \iff \forall m \in p, m' \in \mathcal{M}. \alpha(\lfloor \{m \bullet m'\} \rfloor) \subseteq \lfloor q * \{m'\} \rfloor.$$

Definition 11 (Semantic Judgement). For command $C \in \text{Comm}$ and views $p, q \in \text{View}$, the semantic judgement $\Vdash \{p\} C \{q\}$ is defined coinductively as follows. $\Vdash \{p\} C \{q\}$ holds if and only if the following two conditions are satisfied:

- 1) for all $\alpha \in \text{Action}$, $C' \in \text{Comm}$, if $C \xrightarrow{\alpha} C'$, then there exists $p' \in \text{View}$ such that $\alpha \Vdash \{p\}\{p'\}$ and $\Vdash \{p'\} C' \{q\}$;

- 2) if $C \equiv \text{skip}$ then $p \prec q$.

For each of the proof rules there is a corresponding lemma establishing that it holds for the semantic judgement. Lemma 2 is used in the proofs for the frame and concurrency rules, while Lemma 3 is used for the angelic consequence rule. Together, these lemmas establish the following result.

Lemma 5. If $\{p\} C \{q\}$ is derivable in the program logic, then $\Vdash \{p\} C \{q\}$.

To establish soundness, we link the semantic judgement to the multi-step operational semantics.

Lemma 6. If $\Vdash \{p\} C \{q\}$, then for all $s \in \lfloor p \rfloor$ and $s' \in \mathcal{S}$, if $(C, s) \rightarrow^* (\text{skip}, s')$ then $s' \in \lfloor q \rfloor$.

This lemma follows from the definitions. The proof makes use of the fact that $*$ has a unit, and is the only place in the paper that we explicitly use this fact.

Theorem 1 (Soundness). Assume that $\{p\} C \{q\}$ is derivable in the program logic. Then, for all $s \in \lfloor p \rfloor$ and $s' \in \mathcal{S}$, if $(C, s) \rightarrow^* (\text{skip}, s')$ then $s' \in \lfloor q \rfloor$.

These results have been machine checked with Coq, and the proof scripts are available [25].

A. Examples

We return to the examples that we introduced in II-A to consider their soundness. First, we provide an operational interpretation for the heap-update commands.

Definition 12 (Heap States). Machine states for the heap-update commands are partial functions from variables and locations to values. There is also an exceptional faulting state, denoted \downarrow , which represents the result of an invalid memory access. Formally, $\mathcal{S}_H \stackrel{\text{def}}{=} ((\text{Var} \uplus \text{Loc}) \rightarrow_{\text{fin}} \text{Val}) \uplus \{\downarrow\}$.

Definition 13 (Heap Command Semantics). The semantics of the atomic heap-update commands are given by:

$$\begin{aligned} \llbracket x := y \rrbracket(s) &\stackrel{\text{def}}{=} \begin{cases} \{s[x \mapsto s(y)]\} & \text{if } y \in \text{dom}(s) \\ \{\downarrow\} & \text{otherwise} \end{cases} \\ \llbracket [x] := v \rrbracket(s) &\stackrel{\text{def}}{=} \begin{cases} \{s[s(x) \mapsto v]\} & \text{if } x, s(x) \in \text{dom}(s) \\ \{\downarrow\} & \text{otherwise} \end{cases} \\ \llbracket [x] := y \rrbracket(s) &\stackrel{\text{def}}{=} \begin{cases} \{s[s(x) \mapsto s(y)]\} & \text{if } x, y, s(x) \in \text{dom}(s) \\ \{\downarrow\} & \text{otherwise} \end{cases} \\ \llbracket y := [x] \rrbracket(s) &\stackrel{\text{def}}{=} \begin{cases} \{s[y \mapsto s(s(x))]\} & \text{if } x, s(x) \in \text{dom}(s) \\ \{\downarrow\} & \text{otherwise} \end{cases} \\ \llbracket x := \text{ref } y \rrbracket(s) &\stackrel{\text{def}}{=} \begin{cases} \{s[x \mapsto l, l \mapsto s(y)] \mid l \in \text{Loc} \setminus \text{dom}(s)\} & \text{if } y \in \text{dom}(s) \\ \{\downarrow\} & \text{otherwise.} \end{cases} \end{aligned}$$

Here we extend dom to \mathcal{S}_H by taking $\text{dom}(\downarrow) = \emptyset$.

1) *Disjoint Concurrent Separation Logic*: Since separation-logic views are sets of partial functions from variables and locations to values, they can be seen as sets of heap states. Thus, we can define a simple notion of erasure.

Definition 14 (DCSL Erasure). *The DCSL erasure function $\lfloor - \rfloor_{\text{DCSL}} : \text{View}_{\mathcal{M}_{\text{DCSL}}} \rightarrow \mathcal{P}(\mathcal{S})$ is defined by $\lfloor p \rfloor_{\text{DCSL}} \stackrel{\text{def}}{=} p$.*

It is trivial to see that $\lfloor - \rfloor_{\text{DCSL}}$ is a join-semilattice homomorphism, and that it satisfies Property 8 (Monotonicity of Erasure). All that remains to be shown is Property 9 (Atomic Soundness). As a representative example, we consider the heap-lookup axiom:

$$\{x \Rightarrow l * l \mapsto v * y \Rightarrow _ \} y := [x] \{x \Rightarrow l * l \mapsto v * y \Rightarrow v\}.$$

By appeal to Lemma 4, it is sufficient to show that, for arbitrary $m_r \in \mathcal{M}_{\text{DCSL}}$ and $w \in \text{Val}$,

$$\llbracket y := [x] \rrbracket (\{ [x \mapsto l, l \mapsto v, y \mapsto w] \bullet m_r \}) \subseteq \{ [x \mapsto l, l \mapsto v, y \mapsto v] \bullet m_r \}.$$

If x , y or l is in $\text{dom}(m_r)$ then the composition yields \emptyset and so the inclusion holds trivially. Otherwise, it follows from semantics of command.

Remark. Although mathematically they may be defined the same way, intuitively a separation-logic state and a heap state represent different things. In a separation-logic state, the partiality of the function means that the rest of the state is unknown and not accessible by the thread. In a heap state, the partiality of the function means that the undefined parts are unallocated.

Erasure can be seen as completing the state by treating the unknown regions as unallocated. We could define erasure differently, by completing the state in *every possible way*. For example $\lfloor x \Rightarrow 5 \rfloor$ would be the set of all states in which x has value 5. By changing the notion of erasure like this, we obtain a slightly different notion of soundness.

Note that the erasure function does not cover the machine state space: in particular, there is no view p with $\frac{1}{2} \in \lfloor p \rfloor$. This means that separation-logic triples do not permit memory faults to occur, which is part of the standard interpretation.

2) *Weak-update Type System*: We interpret typing contexts as the set of states which are well-typed with respect to the context. Consequently, in order to define erasure, we must define a notion of typing for states.

Definition 15 (State Typing). *The state typing judgement $\Gamma; \Theta \vdash s$, where $\Gamma : \text{Var} \rightarrow \text{Type}$, $s \in \mathcal{S}_{\text{H}}$ and $\Theta : \text{Loc} \rightarrow \text{Type}$ ranges over heap typing contexts, is defined as follows:*

$$\Gamma; \Theta \vdash s \stackrel{\text{def}}{\iff} \forall x \in \text{dom}(\Gamma). \Theta \vdash s(x) : \Gamma(x) \wedge \forall l \in \text{dom}(\Theta). \Theta \vdash s(l) : \Theta(l)$$

where $\Theta \vdash v : \tau \stackrel{\text{def}}{\iff} \tau = \text{val} \vee \tau = \text{ref}(\Theta(v))$.

The state typing essentially ensures that every typed variable and location has a value consistent with its type. Specifically, this means that references must refer to addresses that have the

appropriate type. Note that it would not be possible to have x and y referencing the same location in the typing context $x : \text{ref val}, y : \text{ref ref val}$. This is necessary, since otherwise an update to the location via x could invalidate the type of y .

Definition 16 (Weak Type Erasure). *The WTS erasure function $\lfloor - \rfloor_{\text{WTS}} : \text{View}_{\mathcal{M}_{\text{WTS}}} \rightarrow \mathcal{P}(\mathcal{S}_{\text{H}})$ is defined by*

$$\lfloor \Gamma \rfloor_{\text{WTS}} \stackrel{\text{def}}{=} \{s \in \mathcal{S}_{\text{H}} \mid \exists \Theta. \exists \Gamma \in \Gamma. \Gamma; \Theta \vdash s\}$$

It is again clear that this definition of erasure is a join-semilattice homomorphism, and that it satisfies Property 8 (Monotonicity of Erasure). It remains to show Property 9 (Atomic Soundness). This is straightforward; for further details, consult [25].

3) *Strong-update Type System*: For the strong-update type system, we use an analogous erasure to that of Definition 16 for weak update. It is straightforward to prove monotonicity and atomic soundness. For details, see [25].

4) *Typed separation logic*: The soundness of this system follows primarily from the two systems on which it is built. We define the erasure as the intersection of the erasures of the underlying systems.

Definition 17 (TSL erasure). *The TSL erasure function $\lfloor - \rfloor_{\text{TSL}} : \text{View}_{\text{TSL}} \rightarrow \mathcal{P}(\mathcal{S}_{\text{H}})$ is defined by*

$$\lfloor \Gamma, p \rfloor_{\text{TSL}} \stackrel{\text{def}}{=} \lfloor \Gamma \rfloor_{\text{STS}} \cap \lfloor p * \text{true} \rfloor_{\text{DCSL}}.$$

In the separation logic component, we allow the state to be larger than specified by the formula to account for any state that the type system might be additionally describing. The type system's erasure is already closed under larger states.

The soundness of the combined axiom follows directly from the soundness of the underlying two axioms, and the angelic preorder satisfies monotonicity as the two underlying orders satisfy monotonicity. The new axioms' soundness proofs are straightforward and given in the extended version [25].

IV. VIEWS AND INTERFERENCE

Views model partial, abstracted information about machine states that are immune to interference from other threads. However, typically we start with a view model that does not consider possible interference. We then account for interference by using an *interference operator* that captures how a given view may evolve under interference.

Definition 18 (Interference Operator). *Given view model View, an interference operator $R : \text{View} \rightarrow \text{View}$ is function on View that satisfies the following properties:*

$$R(p * q) \models R(p) * R(q) \tag{1}$$

$$R(u) = u \tag{2}$$

$$R\left(\bigvee \{q_i\}_{i \in I}\right) = \bigvee \{R(q_i)\}_{i \in I} \tag{3}$$

$$p \models R(p) \tag{4}$$

$$R(R(p)) \models R(p) \tag{5}$$

Property (3) implies that R is monotone:

$$p \vDash q \implies R(p) \vDash R(q) \quad (6)$$

Monotonicity (6), inflation (4) and idempotency (5) correspond to R being a closure operator.

A view p is *stable* if $R(p) = p$. Given an unstable view p , $R(p)$ is the *stabilisation* of p . In fact, $R(p)$ is the least (with respect to \vDash) stable view that is greater than p . (Minimality can be seen as a consequence of properties (6) and (5).)

The separation property of interference operators (1) has the motto “more resource means less interference”: the interference acting on the combined resource $p * q$ does no more than the interference acting separately on p and q . With (4) and (5), it implies that the composition of stable views is also stable:

$$\begin{aligned} R(p) * R(q) &\vDash R(R(p) * R(q)) \\ &\vDash R(R(p)) * R(R(q)) \vDash R(p) * R(q) \\ \therefore R(p) * R(q) &= R(R(p) * R(q)) \end{aligned}$$

Consequently, the subset of stable views themselves form a view model with the same composition, unit and join.

Lemma 7. *If R is an interference operator on View then $R(\text{View}) = \{R(p) \mid p \in \text{View}\}$ is a view model.*

We can also show that stable views are closed under meets. Since meet is a greatest lower bound on View, it must also give a greatest lower bound on $R(\text{View})$. As a corollary, if View is a Heyting algebra then $R(\text{View})$ will also be a Heyting algebra.

Interference operators provide an approach to constructing view models, which may be used in our proof system. By restricting the proof system to the stable views, we alter the set of possible axioms: we forbid axioms with unstable pre- and postconditions; and we add new axioms that only preserve stable frames—such axioms could fail to preserve some unstable frame and be unsound in the original proof system.

When a view model is generated from a separation algebra, we can characterise interference using interference relations.

Definition 19 (Interference Relation). *An interference relation $R \subseteq \mathcal{M} \times \mathcal{M}$ is a preorder which satisfies the following decomposition property: for all $m_1, m_2, m, m' \in \mathcal{M}$ with $m = m_1 \bullet m_2$ and $m R m'$, there exist $m'_1, m'_2 \in \mathcal{M}$ with $m_1 R m'_1$, $m_2 R m'_2$ and $m' = m'_1 \bullet m'_2$.*

By lifting the interference relation to sets, we can obtain an interference operator. In fact, the following lemma establishes that there is a one-to-one correspondence between interference relations and interference operators.

Lemma 8. *Let $\text{View}_{\mathcal{M}}$ be the view model induced by separation algebra $(\mathcal{M}, \bullet, I)$. Then there is a bijective correspondence, $\hat{\cdot}$, between interference relations on \mathcal{M} and interference operators on $\text{View}_{\mathcal{M}}$ with*

$$m' \in \hat{R}(p) \iff \exists m \in p. m R m' .$$

In this setting, the action judgement (Definition 10) has an equivalent formulation that is typically simpler to check.

Lemma 9.

$$\alpha \Vdash \{p\}\{q\} \iff \forall m \in \mathcal{M}. \alpha (\lfloor p * \{m\} \rfloor) \subseteq \lfloor q * \hat{R}(\{m\}) \rfloor .$$

A. Examples

We give examples of views with interference. We first explore the running examples of the paper. Although our example view models have been constructed without interference, it is instructive to consider how interference could be used in their definition. Interference allows us to separate instantaneous knowledge about the program state (which is subject to change as a result of the actions of other threads) and knowledge about how the state can evolve as a result of other threads’ actions. The knowledge about the state can be seen as unstructured and global, whereas the instantaneous knowledge about the evolution of the state is a separable resource that is local to each thread. We also study the fine-grained reasoning of concurrent abstract predicates [15]. We would not know how to construct this example without using interference.

1) *Separation Logic with Ownership*: Rather than the view model introduced in II-A, we construct a view model for disjoint concurrent separation logic by instrumenting machine states (excluding \downarrow) with an ownership mask, which provides explicit permissions for stating which variables and addresses are “owned”. Our set of instrumented states is:

$$\mathcal{M}_{\text{MSL}} \stackrel{\text{def}}{=} (\text{Var} \uplus \text{Loc}) \rightarrow_{\text{fin}} (\text{Val} \times \{0, 1\}) .$$

Given instrumented state $m \in \mathcal{M}_{\text{MSL}}$, for each variable x (or address l), the first component of $m(x)$ (or $m(l)$) is its actual value in the machine, while the second component indicates whether or not the variable (or location) is owned; if $m(x)$ is undefined, the variable is not in the state at all; if $m(l)$ is undefined then l is not allocated. Composition is defined by:

$$\begin{aligned} m_1 \bullet m_2 &= m \stackrel{\text{def}}{\iff} \text{dom}(m_1) = \text{dom}(m_2) = \text{dom}(m) \\ &\quad \wedge \forall k \in \text{dom}(m). m_1(k) \downarrow_1 = m_2(k) \downarrow_1 = m(k) \downarrow_1 \\ &\quad \wedge m_1(k) \downarrow_2 + m_2(k) \downarrow_2 = m(k) \downarrow_2 \end{aligned}$$

Composition requires that the state components of the composed states are the same as that of the composite, and that their ownership masks sum to give the mask of the composite. This ensures that each variable and location is uniquely owned.

This composition is well-defined, associative and commutative. To complete the separation algebra $(\mathcal{M}_{\text{MSL}}, \bullet, I)$, it remains to give the unit:

$$I = \{m \in \mathcal{M}_{\text{MSL}} \mid \forall k \in \text{dom}(m). m(k) \downarrow_2 = 0\} .$$

It is also easy to see that this is indeed the unit.

If we were to construct a view model based on this separation algebra, the commands we could reason about would be very limited: they could not alter the (machine) state at all. To see this, suppose we are able to derive the triple $\{\{m\}\} C \{q\}$. Then, by the frame rule, we could also derive the triple $\{\{m \bullet m_i\}\} C \{q * \{m_i\}\}$, where $m_i \in I$ is the unit element for m . By the definition of composition, this is equivalent to

$$\{\{m\}\} C \left\{ \left\{ m' \in q \mid \text{dom}(m) = \text{dom}(m') \wedge \forall k \in \text{dom}(m). m(k) \downarrow_1 = m'(k) \downarrow_1 \right\} \right\}$$

Thus, C cannot change the values of any variables or locations.

The problem here is that programs are required to preserve all frames, and therefore all values. However, the intention behind ownership is that only variables and locations that are owned are guaranteed to be preserved by other threads. Thus, instead of preserving all frames, we wish only to preserve all *stable* frames for a suitable notion of stability. Such a notion can be obtained by defining an interference relation:

$$m R m' \stackrel{\text{def}}{\iff} \forall k \in \text{dom}(m). m(k) \downarrow_1 > 0 \implies m'(k) = m(k)$$

Intuitively, this interference relation expresses that the environment can do anything that does not alter the variables and locations owned by the thread. It is not difficult to see that the interference relation satisfies the decomposition property: any change that does not alter the variables or locations owned by either thread does not alter the variables or locations owned by each thread individually.

If we now consider the view model induced by the separation algebra under this interference relation, $R(\text{View}_{\text{MSL}})$, we obtain a notion of view that is specific about variables and locations that are owned, but can say nothing at all about variables and locations that are not owned. Now, threads are at liberty to mutate variables and heap locations they own, and allocate locations that are not owned by other threads, since these operations preserve stable frames. (Note that composition plays an important role here, since it enforces that the environment cannot also own variables and locations that belong to the thread.)

We define an operation $\iota : R(\text{View}_{\text{MSL}}) \rightarrow \text{View}_{\text{DCSL}}$ by:

$$\iota(p) \stackrel{\text{def}}{=} \left\{ m \in \mathcal{M}_{\text{DCSL}} \mid \begin{array}{l} \exists m' \in p. \forall k, v. \\ m(k) = v \iff m'(k) = (v, 1) \end{array} \right\}$$

This operation in fact defines an isomorphism between the two view models, so we can really think of the separation logic model as being constructed in this way.

It is a small step to define a separation logic with fractional permissions [26] in a similar fashion. Instead of a bit mask from the set $\{0, 1\}$, variables and locations are associated with fractional permissions from the interval $[0, 1]$. Composition and interference can be defined in the same way as above. A thread may depend upon the environment preserving the value of anything for which it holds a non-zero permission. Conversely, in order to mutate a variable or location, a thread must ensure the environment cannot have a non-zero permission: the thread must have permission 1 itself.

2) *Type Systems*: We construct a model for the weak-update type system, in a similar fashion to above, by instrumenting machine states with type information. The set of instrumented states can be constructed by: $\mathcal{M}_{\text{WUTS}} \stackrel{\text{def}}{=} \{(s, \Gamma, \Theta) \mid \Gamma; \Theta \vdash s\}$. We turn this into a separation algebra $(\mathcal{M}_{\text{WUTS}}, \bullet, \mathcal{M}_{\text{WUTS}})$ by defining

$$(s, \Gamma_1, \Theta_1) \bullet (s, \Gamma_2, \Theta_2) \stackrel{\text{def}}{=} (s, \Gamma_1 \cup \Gamma_2, \Theta_1 \cup \Theta_2)$$

with \bullet undefined when the states are different or the result is not well-typed.

Again, if we construct a view model from this separation algebra, operations will not be allowed to mutate the state. However, we can define an interference relation, R , on $\mathcal{M}_{\text{WUTS}}$ to be the least relation such that $(s, \Gamma, \Theta) R (s', \Gamma, \Theta')$ for all $s, s', \Gamma, \Theta, \Theta'$. This leads to a definition of views in which only the typing of variables is fixed. We interpret typing contexts as views in the following way: $[\Gamma] \stackrel{\text{def}}{=} \{(s, \Gamma, \Theta) : (s, \Gamma, \Theta) \in \mathcal{M}_{\text{WUTS}}\}$. This interpretation gives rise to an isomorphism between this view model and the one considered in II-A.

The strong-update type system can be modelled by augmenting the instrumentation with an ownership mask on variables (though not the heap). Composition is defined to ensure disjoint ownership, as in the separation logic example. Interference then permits any changes that preserve the types of owned variables.

3) *Concurrent Abstract Predicates*: Concurrent Abstract Predicates (CAP) [15] extends separation logic with shared regions that can be mutated by multiple threads through atomic operations. Each shared region is associated with some interference environment which stipulates how it may be mutated, by defining a collection of actions. Threads can mutate the shared state only by performing actions for which they have a capability resource, which may be exclusive or non-exclusive. An example of a CAP assertion is the following:

$$\exists r. [\text{LOCK}]_{\pi}^r * \boxed{(x \mapsto 0 * [\text{UNLOCK}]_1^r) \vee x \mapsto 1}_{I(r)}^r \quad (7)$$

This assertion represents a shared mutex in the heap at address x . The boxed part of the assertion describes the shared region: either the mutex is available (x has value 0) or it is locked (x has value 1). If it is available, the exclusive capability $[\text{UNLOCK}]_1^r$ belongs to the region. The thread itself has a non-exclusive capability $[\text{LOCK}]_{\pi}^r$, which will permit it to lock the mutex. The actions corresponding to the capabilities are defined by the interference environment $I(r)$:

$$I(r) \stackrel{\text{def}}{=} \left(\begin{array}{ll} \text{LOCK} & : x \mapsto 0 * [\text{UNLOCK}]_1^r \rightsquigarrow x \mapsto 1, \\ \text{UNLOCK} & : x \mapsto 1 \rightsquigarrow x \mapsto 0 * [\text{UNLOCK}]_1^r \end{array} \right)$$

The interference environment stipulates that a thread holding a LOCK capability may (atomically) update x from 0 to 1 in the shared region, simultaneously removing the capability $[\text{UNLOCK}]_1^r$; a thread holding an UNLOCK capability may update x from 1 to 0 while returning $[\text{UNLOCK}]_1^r$ to the shared region. By locking the mutex, a thread acquires the exclusive capability to unlock it. This guarantees that only one thread can have locked the mutex at a time, and only that thread can unlock it. Having locked the mutex (through an atomic compare-and-swap operation), a thread's view may be described by the following assertion:

$$\exists r. [\text{UNLOCK}]_1^r * [\text{LOCK}]_{\pi}^r * \boxed{x \mapsto 1}_{I(r)}^r \quad (8)$$

Whereas in (7) the thread does not know whether the mutex is locked or not, in (8) it knows that it is locked. It can only make such an assertion because it holds the exclusive capability $[\text{UNLOCK}]_1^r$; without that, the assertion would not

be stable. (In separation logic, capabilities are always directly associated with the information they pertain to, as in $l \mapsto 7$, whereas in CAP they can be separated.)

CAP assertions are modelled by sets of states instrumented with ownership, capabilities and region information. To use CAP, we must work with stable assertions — views — in order to account for possible interactions between threads. These stable assertions are defined by closure under an interference relation, which allows the environment to update regions in any way for which it can have a suitable capability.

In our previous examples, interference was not essential. Since the interference relations were equivalences, views were sets of equivalence classes, and it was easy to define canonical representations of these classes. CAP, on the other hand, is difficult to model without considering interference explicitly. One reason for this might be that interference is directional, and not simply an equivalence. Consequently, expressing entailments as inclusions between closed sets is likely to be the simplest way of constructing a suitable view model. Another reason is that it is convenient to construct views from assertions that are not themselves stable, hence it is convenient to have a model in which we can represent both stable and unstable assertions.

V. RELATED WORK

Our composition operator provides a logical notion of separation, which, as we have demonstrated by examples, need not be realized by physical separation in the concrete machine. This idea of fictional separation has been used in recent work on separation logics for concurrent languages [15], [9], [10], [14]. Similar ideas are also useful in a purely sequential setting to enable modular reasoning about abstract data structures implemented using physical sharing, but for which a logical notion of separation can be defined [27], [28], [29].

The soundness of Pottier’s capability system [30] is based on an axiom that is similar to our definition of interference relation, and the soundness proof of concurrent abstract predicates [15] also uses an equivalent lemma. Our framework does not have an explicit notion of guarantee, so many of the other properties required in both Pottier’s work and concurrent abstract predicates are not required. Feng’s LRG [31] also provides conditions such that the stable predicates can be composed. The condition requires fences to delimit the scope of interference, which we do not require. LRG is the only combination of rely-guarantee with separation logic we have not encoded into our framework, and remains future work.

VI. CONCLUSIONS

We have introduced views as a general framework in which a wide variety of compositional reasoning approaches can be constructed, understood and proved sound. We find it surprising and revealing that diverse approaches such as separation logic and type theory can be understood in an elegant, unifying setting. We have shown a selection of program logics that can be seen as applications of our framework. Many others are possible; details of some can be found in the technical report [25].

Ten years ago, Reynolds asked “whether the dividing line between types and assertions can be erased” [32]. We have shown there is no dividing line: they are just different kinds of *view*.

REFERENCES

- [1] G. Morrisett, D. Walker, K. Cray, and N. Glew, “From system F to typed assembly language,” *TOPLAS*, vol. 21, no. 3, pp. 527–568, 1999.
- [2] A. Ahmed, M. Fluett, and G. Morrisett, “L³: A linear language with locations,” *Fundam. Inform.*, vol. 77, no. 4, pp. 397–449, 2007.
- [3] F. Smith, D. Walker, and J. G. Morrisett, “Alias types,” in *ESOP*, 2000.
- [4] A. Charguéraud and F. Pottier, “Functional translation of a calculus of capabilities,” in *ICFP*, 2008, pp. 213–224.
- [5] P. W. O’Hearn, J. C. Reynolds, and H. Yang, “Local reasoning about programs that alter data structures,” in *CSL*, 2001, pp. 1–19.
- [6] J. Schwinghammer, H. Yang, L. Birkedal, F. Pottier, and B. Reus, “A semantic foundation for hidden state,” in *FOSSACS*, 2010, pp. 2–16.
- [7] J. Schwinghammer, L. Birkedal, and K. Støvring, “A step-indexed Kripke model of hidden state via recursive properties on recursively defined metric spaces,” in *FOSSACS*, 2011.
- [8] L. Birkedal, N. Torp-Smith, and H. Yang, “Semantics of separation-logic typing and higher-order frame rules for Algol-like languages,” *Logical Methods in Computer Science*, vol. 2, no. 5:1, 2006.
- [9] X. Feng, R. Ferreira, and Z. Shao, “On the relationship between concurrent separation logic and assume-guarantee reasoning,” in *ESOP*, 2007, pp. 173–188.
- [10] V. Vafeiadis and M. J. Parkinson, “A marriage of rely/guarantee and separation logic,” in *CONCUR*, 2007, pp. 256–271.
- [11] A. Gotsman, J. Berdine, B. Cook, N. Rinetzky, and M. Sagiv, “Local reasoning for storable locks and threads,” in *APLAS*, 2007.
- [12] A. Hobor, “Oracle semantics,” Ph.D. dissertation, Princeton University, Department of Computer Science, Princeton, NJ, October 2008.
- [13] A. Buisse, L. Birkedal, and K. Støvring, “A step-indexed Kripke model of separation logic for storable locks,” in *MFPS*, 2011.
- [14] M. Dodds, X. Feng, M. J. Parkinson, and V. Vafeiadis, “Deny-guarantee reasoning,” in *ESOP*, 2009, pp. 363–377.
- [15] T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis, “Concurrent abstract predicates,” in *ECOOP*, 2010.
- [16] M. Dodds, S. Jagannathan, and M. J. Parkinson, “Modular reasoning for deterministic parallelism,” in *POPL*, 2011.
- [17] G. Tan, Z. Shao, X. Feng, and H. Cai, “Weak updates and separation logic,” in *APLAS*. Springer-Verlag, 2009, pp. 178–193.
- [18] P. W. O’Hearn and D. J. Pym, “The logic of bunched implications,” *Bulletin of Symbolic Logic*, vol. 5, no. 2, pp. 215–244, 1999.
- [19] D. J. Pym, *The Semantics and Proof Theory of the Logic of Bunched Implications*, ser. Applied Logic Series. Springer, 2002, vol. 26.
- [20] V. Vafeiadis, “Concurrent separation logic and operational semantics,” in *MFPS*, 2011.
- [21] C. Calcagno, P. W. O’Hearn, and H. Yang, “Local action and abstract separation logic,” in *LICS*. IEEE Computer Society, 2007, pp. 366–378.
- [22] C. Calcagno, P. Gardner, and U. Zarfaty, “Local reasoning about data update,” *ENTCS*, vol. 172, pp. 133–175, 2007.
- [23] R. Dockins, A. Hobor, and A. W. Appel, “A fresh look at separation algebras and share accounting,” in *APLAS*, 2009.
- [24] M. J. Parkinson, “Local reasoning for java,” Ph.D. dissertation, University of Cambridge, November 2005.
- [25] T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang, “Additional material,” <http://sites.google.com/site/viewmodel/>, 2012.
- [26] R. Bornat, C. Calcagno, P. O’Hearn, and M. Parkinson, “Permission accounting in separation logic,” in *POPL’05*, 2005.
- [27] T. Dinsdale-Young, P. Gardner, and M. Wheelhouse, “Abstraction and refinement for local reasoning,” in *VSTTE*, 2010.
- [28] N. Krishnaswami, L. Birkedal, and J. Aldrich, “Verifying event-driven programs using ramified frame properties,” in *TLDI*, 2010.
- [29] J. B. Jensen and L. Birkedal, “Fictional separation logic,” in *Proceedings of ESOP*, 2012, to Appear.
- [30] F. Pottier, “Syntactic soundness proof of a type-and-capability system with hidden state,” INRIA, Tech. Rep., 2011, (available from the author).
- [31] X. Feng, “Local rely-guarantee reasoning,” in *POPL*, 2009.
- [32] J. C. Reynolds, “Separation logic: A logic for shared mutable data structures,” in *LICS*, 2002.