



On Algebraic Abstractions for Concurrent Separation Logics

FRANTIŠEK FARKA, IMDEA Software Institute, Spain

ALEKSANDAR NANEVSKI, IMDEA Software Institute, Spain

ANINDYA BANERJEE, IMDEA Software Institute, Spain

GERMÁN ANDRÉS DELBIANCO, Nomadic Labs, France

IGNACIO FÁBREGAS, Universidad Complutense de Madrid, Spain

Concurrent separation logic is distinguished by transfer of state ownership upon parallel composition and framing. The algebraic structure that underpins ownership transfer is that of partial commutative monoids (PCMs). Extant research considers ownership transfer primarily from the logical perspective while comparatively less attention is drawn to the algebraic considerations. This paper provides an algebraic formalization of ownership transfer in concurrent separation logic by means of structure-preserving partial functions (i.e., morphisms) between PCMs, and an associated notion of separating relations. Morphisms of structures are a standard concept in algebra and category theory, but haven't seen ubiquitous use in separation logic before. Separating relations are binary relations that generalize disjointness and characterize the inputs on which morphisms preserve structure. The two abstractions facilitate verification by enabling concise ways of writing specs, by providing abstract views of threads' states that are preserved under ownership transfer, and by enabling user-level construction of new PCMs out of existing ones.

CCS Concepts: • **Theory of computation** → **Separation logic; Hoare logic; Type theory**; • **Software and its engineering** → *Formal software verification*; • **Computing methodologies** → *Concurrent algorithms*.

Additional Key Words and Phrases: Program Logics for Concurrency, Hoare/Separation Logics, Coq

ACM Reference Format:

František Farka, Aleksandar Nanevski, Anindya Banerjee, Germán Andrés Delbianco, and Ignacio Fábregas. 2021. On Algebraic Abstractions for Concurrent Separation Logics. *Proc. ACM Program. Lang.* 5, POPL, Article 5 (January 2021), 32 pages. <https://doi.org/10.1145/3434286>

1 INTRODUCTION

The algebraic foundations of separation logic are rooted in the discovery that the structure of *partial commutative monoids* (PCMs) underpins the semantics of the key inference rules of framing and parallel composition [Calcagno et al. 2007; Dinsdale-Young et al. 2013; Pym et al. 2004]. The PCMs do so by mathematically representing the essential notions of state ownership and ownership transfer, while abstracting the details of the concrete memory models used by the programs.

In a nutshell, a PCM is a structure $(A, \bullet, \mathbb{1})$ on a carrier set A , equipped with a (partial) binary operation \bullet (pronounced “join”), which is commutative, associative, and has $\mathbb{1}$ as the unit. The elements of the carrier A model the *private* state of individual threads, and \bullet models how the private states of two children threads combine into the state of their parent. The operation \bullet is

Authors' addresses: František Farka, IMDEA Software Institute, Spain, frantisek.farka@imdea.org; Aleksandar Nanevski, IMDEA Software Institute, Spain, aleks.nanevski@imdea.org; Anindya Banerjee, IMDEA Software Institute, Spain, anindya.banerjee@imdea.org; Germán Andrés Delbianco, Nomadic Labs, France, german@nomadic-labs.com; Ignacio Fábregas, Universidad Complutense de Madrid, Spain, fabregas@fdi.ucm.es.



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

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

2475-1421/2021/1-ART5

<https://doi.org/10.1145/3434286>

commutative and associative because the order of threads in a thread pool is irrelevant for the computation. The operation \bullet is partial to signify that some state combinations are impossible. For example, if $x \bullet y$ is undefined, then x and y can't be the private states of two different concurrent threads, *simultaneously*. The unit element represents the empty private state.

The canonical PCM in separation logic is that of heaps, which are finite maps from pointers (positive natural numbers) to values. The \bullet is the *disjoint* union of heaps. It is undefined if the operand heaps have a pointer in common, thus modeling that the private heaps of two concurrent threads can't share pointers. The unit is the heap with no pointers allocated. When a parent forks two children threads, then its private heap is divided disjointly among the children. Upon joining, the private, disjoint heaps of the children are unioned to derive the heap of the parent. This transfer of heap ownership between parent and children threads is the defining pattern of separation logic.

While PCMs were originally used to explain the *semantics* of separation logic, more recent separation logics [Appel et al. 2014; Jensen and Birkedal 2012; Jung et al. 2015; Ley-Wild and Nanevski 2013; Nanevski et al. 2014] take a step further and employ PCMs in *program specifications* (henceforth: specs). In these logics the user may introduce various PCMs to model custom notions of ghost state relevant to the verification problem. Examples include PCMs of permissions [Bornat et al. 2005], and PCMs of histories [Sergey et al. 2015b] for representing temporal (i.e., execution order) properties in the style of linearizability and other consistency criteria [Delbianco et al. 2017; Nanevski et al. 2019; Sergey et al. 2016]. Having arbitrary PCMs also facilitates the verification of graph algorithms [Sergey et al. 2015a], which has been notoriously difficult in heap-only separation logics. These approaches therefore usefully combine the algebra of PCMs with logical reasoning about state ownership and transfer.

In this paper, we take the PCM-based approach to specification significantly further by introducing a theory of structure-preserving functions (*morphisms*), and structure-preserving relations (*separating relations*) on PCMs. Morphisms are *partial*, as they preserve the PCM structure only on *some* inputs. Separating relations are binary relations that describe the inputs on which a morphism is structure-preserving, and abstractly generalize heap disjointness.

The above development has two relevant consequences for separation logic. First, it immediately provides powerful user-level support for constructing new PCMs out of existing ones. To see why such construction is desirable, consider that to specify both spatial and temporal properties of programs, the user may want to combine the PCMs of heaps and histories into their Cartesian product, itself also a PCM. But a standard use of morphisms in abstract algebra and category theory is precisely in the definition of algebraic constructions, where morphisms relate a construction to its components, e.g., how a Cartesian product is associated with projection and pairing morphisms. We illustrate this aspect of our contribution by introducing the algebraic construction of a *sub-PCM*, and showing how it applies to verification in separation logic.

Second, the two concepts (morphisms and separating relations) provide ways to *abstract* from the concrete thread states; morphisms can functionally *compute* novel abstractions from a state, whereas separating relations relate the states of a thread and its concurrent environment. Being structure-preserving means that both respect the ownership transfer of separation logic, as we shall see. Together, the two concepts thus present: a novel foundation for separation logic that facilitates systematic introduction of algebraic concepts into specs; and a way to mathematically model the essentials of a verification problem while abstracting from details of program state.

1.1 Morphisms as Ownership-Preserving Abstractions

Glossing over the partiality of PCMs, to which we return in Section 1.2, the standard algebraic definition says that a morphism from the monoid $(A, \bullet_A, \mathbb{1}_A)$ to the monoid $(B, \bullet_B, \mathbb{1}_B)$ is a function

$\phi : A \rightarrow B$ that preserves the monoidal structure:

$$\phi(\mathbb{1}_A) = \mathbb{1}_B \quad (1)$$

$$\phi(x \bullet_A y) = \phi(x) \bullet_B \phi(y) \quad (2)$$

We previously described \bullet as a way to combine private states of two children threads into the state of the parent. The above equations then characterize ϕ as computing a view—an abstraction—of a thread’s private state, while preserving the thread-private nature of the view.

To illustrate, consider how PCMs may model a mutually exclusive lock that threads race to acquire. We first require the PCM O that formalizes lock ownership. O has the carrier $\{\text{own}, \overline{\text{own}}\}$, where own (resp. $\overline{\text{own}}$) signifies that the thread owns (resp. doesn’t own) the lock. The \bullet computes the lock ownership of the parent thread from those of the children by the following table, where $\text{own} \bullet \text{own}$ is undefined as two threads can’t own the lock simultaneously, and $\overline{\text{own}}$ is the unit.

x	y	$x \bullet y$
own	own	undefined
own	$\overline{\text{own}}$	own
$\overline{\text{own}}$	own	own
$\overline{\text{own}}$	$\overline{\text{own}}$	$\overline{\text{own}}$

The table says that the lock is transferred from child to parent upon joining, analogously to how the heap of a child is transferred to the parent, as discussed before. If neither child owns the lock, then the parent doesn’t own the lock either.

A concrete implementation of the lock will typically require threads to internally store much more private lock-related state than merely an element of O . The extra state may be used for synchronization purposes, or it may be ghost state required to formulate the logical invariants of the locking algorithm, as often necessary for verification. Let this private state be modeled by a PCM X . The concrete definition of X may differ between lock implementations and proofs, but each should exhibit a function $\alpha : X \rightarrow O$ that computes the lock ownership status αx of a thread from the thread’s private state $x \in X$.

Moreover, α must be structure-preserving, and in particular must satisfy equation (2). To see what goes wrong if α doesn’t, suppose there’re states x and y such that, e.g., $\alpha x = \alpha y = \overline{\text{own}}$, but $\alpha(x \bullet y) = \text{own}$, to consider but one bad combination of values for α (the other bad combinations are similarly absurd). Then we have two children threads that don’t own the lock, but their parent is granted the lock upon joining, out of thin air. Such α violates the transfer of lock ownership between children and parent threads, and thus doesn’t model locking.

1.2 Partiality and Separating Relations as Abstraction of Disjointness

Taking into account that \bullet_A may be undefined on some inputs, it’s clear that equation (2) can’t hold as stated, but must be prefixed by some condition on x and y . At the very least, such condition should entail that $x \bullet_A y$ is defined, so that ϕ has an input value to which to apply, and on which ϕ itself is defined. More generally, we associate ϕ with a binary relation \perp_ϕ that captures when ϕ distributes over \bullet , via the updated axiom

$$\text{if } x \perp_\phi y \text{ then } x \bullet_A y \text{ and } \phi(x \bullet_A y) \text{ are defined and } \phi(x \bullet_A y) = \phi(x) \bullet_B \phi(y) \quad (3)$$

The relation \perp_ϕ will be a *separating relation*, thus satisfying a number of properties that we outline in Section 3.2. One of the properties that $x \perp_\phi y$ entails is that $x \bullet_A y$ is defined, or, equivalently, that x and y are *separate* (denoted $x \perp y$). Clearly, this notion generalizes disjointness of heaps and applies it to arbitrary PCMs. Then a separating relation \perp_ϕ represents a morphism-specific notion of separateness that strengthens the one inherited from the underlying PCM.

Because separateness determines when states of two threads combine into a parent state, separating relations essentially provide a custom notion of when two PCM elements can be considered as states of concurrent threads, and thus also when a PCM element can be transferred from one thread to another. A related important use of separating relations is in the construction of sub-PCMs of the PCM A , whereby \bullet_A is restricted to the inputs admitted by the separating relation. These uses are illustrated in Sections 2.7 and 3.2.

We also show in Section 3.2 that morphisms and separating relations are closed under basic algebraic constructions. For example, morphism kernels and equalizers are separating relations; restricting a morphism by a separating relation produces a new morphism, etc. Thus, separating relations are a natural algebraic structure to describe the inputs on which a partial PCM morphism is structure-preserving (and defined).

1.3 Morphisms and Separating Relations in the Abstract

We further consider how morphisms and separating relations interact to support framing (or more generally, parallel composition) *in the abstract*. In other words, if we have a spec involving morphisms and separating relations whose exact definitions we want to hide, what properties must be exposed to make it possible to frame the spec? In Section 4, we argue that what must be exposed is that the morphisms and the separating relations respectively satisfy the novel property of *invertibility*, in addition to being structure-preserving functions and relations. Framing in the abstract has been considered in related work on concurrent abstract predicates (CAP) [Dinsdale-Young et al. 2010]. The novelty of our approach is the use of morphisms (i.e., functions) rather than predicates (i.e., relations). When possible, functions are preferred to relations, as results of functions needn't be named; hence one can avoid existential quantification (e.g., consider function vs. relation composition). Section 5 discusses the relationship to concurrent abstract predicates.

1.4 Use of Morphisms in Specs

We show that morphisms allow the user to compute, directly in specs, PCM values out of the state, without requiring almost any other logical connectives familiar from separation logic. Thus, for the most part, our specs won't use separating conjunction¹ or separating implication, or the numerous recent additions to separation logic of *assertions* in the form of modalities and custom notions of implication [Bizjak et al. 2019; Dinsdale-Young et al. 2010; Jung et al. 2018, 2015] and quantification [da Rocha Pinto et al. 2014]. Instead, we rely only on standard constructs from higher-order logic to make and combine statements about morphism values and separating relations. However, ours is still a separation logic as we're concerned with PCMs and ownership transfer.

As morphisms are just a special class of functions, they are particularly well-suited to a formalization as a *shallow embedding* in a system based on type theory such as Coq. We have thus mechanized all the results from the paper by building on the recent formulation of separation logic in Coq by Nanevski et al. [2019]. Morphisms and separating relations integrate very naturally into this ambient theory, and don't require any particular automation by tactics in order to be used effectively. The resulting mechanization is available as a separate artefact [Farka et al. 2020a].

2 PCM ABSTRACTIONS BY EXAMPLE

2.1 Ticket Lock

To illustrate the issue at hand consider a simple synchronization primitive, a ticket lock [Lamport 1974; Mellor-Crummey and Scott 1991]. Ticket lock consists of two shared pointers, the ticket dispenser tdr , and the display dsp . The thread that wishes to acquire the lock first increments tdr

¹Though we'll define a similar notion for use in proof outlines.

by the `inc_and_fetch` primitive.² The thread then loops until the pointer `dsp` matches the value read from `tdr`. The thread unlocks by incrementing the value of `dsp`.

$$\begin{aligned} \text{lock} &\hat{=} x \leftarrow \text{inc_and_fetch}(tdr); \\ &\quad \text{do } y \leftarrow !dsp \text{ until } x = y \\ \text{unlock} &\hat{=} \text{inc_and_fetch}(dsp) \end{aligned}$$

Intuitively, the ticket lock's workflow resembles the ticket queue management system that guides customers to a counter in a bakery [Lampert 1974]. Incrementing `tdr` corresponds to taking a new ticket from the ticket dispenser, thus fixing a customer's position in the queue. Looping corresponds to awaiting the ticket's turn. Incrementing `dsp` signals, on the display, the next customer's turn. The initial value of `tdr` is 0; thus, the first ticket drawn is 1. The initial value of `dsp` is 1; thus, the first thread that draws 1 can immediately be served. The sequel continues this analogy.

Specs. The specs of the two ticket lock programs should say that `lock` acquires exclusive ownership of the ticket lock, and `unlock` releases it. We denote that by the following type ascriptions.³

$$\text{lock} : \{\lambda s. \alpha_s s = \overline{\text{own}}\} \{\lambda s. \alpha_s s = \text{own}\} @ \text{TL} \quad (4)$$

$$\text{unlock} : \{\lambda s. \alpha_s s = \text{own}\} \{\lambda s. \alpha_s s = \overline{\text{own}}\} @ \text{TL} \quad (5)$$

Unlike most separation logics, we make the binding of the state s in the assertions explicit by means of λ , as customary in higher-order logic. In the above specs, α is a morphism from the underlying PCM of the state s , whose exact definition we want to keep abstract, to the PCM O from Section 1.1. Several questions arise. Although a client can reason with the specs, they appear too abstract: how can the specs be established in the first place? After all, on inspection of the implementations of `lock` and `unlock` above, it isn't obvious how morphism α is even involved. It turns out that we will require concrete specs of the implementations and then hide implementation-level details to define α and obtain the abstract specs (4) and (5). But then how do morphisms and separating relations interact with the *concrete* specs? How do they work with framing of the concrete specs? How do they work under abstraction? The sequel answers these questions after first introducing the basics of our type-theoretic approach.

2.2 Hoare Types, States and Specifications

Hoare Types. A Hoare type [Nanevski et al. 2019, 2006] is a dependently typed state and concurrency (and divergence) monad, indexed with a spec in the style of separation logic. Concretely, in the judgment $e : \{P\}\{Q\} @ V$, P and Q , both predicates over state s , are respectively the pre- and postcondition of a program e , in the sense of partial correctness. V is a *resource*, i.e., a state transition system describing the atomic state changes that e is permitted. Two programs can be safely composed, sequentially or in parallel, only if they are typed by the same resource. The resource thus serves as a bound on the interference that concurrent threads can perform on each other's executions, enabling a form of *rely-guarantee* reasoning [Jones 1983]. As `lock` and `unlock` share the resource `TL` (to be defined soon), they can be composed.

States. In our ambient type theory [Nanevski et al. 2019], states are *subjective* [Ley-Wild and Nanevski 2013]. That is, each state s is a pair (s_s, s_o) , where s_s and s_o are referred to as *self* and *other* components, respectively. The s_s component describes the private state of a thread, whereas

²*Increment-and-fetch* is a generic RMW operation [Herlihy and Shavit 2008] that atomically increments the value stored at `tdr` and returns the incremented value. Similar primitives exist in many systems, e.g., `--atomic_add_fetch` of gcc.

³For simplicity, we don't consider lock invariants that describe the heap that the lock protects. Attaching such invariants is an orthogonal issue to the topic of this paper and has been discussed in [Nanevski et al. 2019].




Fig. 1. States of concurrent threads. *Self* components are in light shade, *other* components are in dark. Adapted from [Nanevski et al. 2019].

s_o describes the combined state of all the other threads, that is, the concurrent environment.⁴ Thus, *self* and *other* components model, *at the level of state*, the same dichotomy modeled by the rely and guarantee *transitions* of rely-guarantee reasoning. The value of s_o may be used in specs, but a program can't alter it. Both s_s and s_o are elements of one and the same PCM. When we access state components by a morphism ϕ , we attach the subscript to the morphism and write $\phi_s s$ and $\phi_o s$ instead of ϕs_s and ϕs_o respectively. We write $\hat{\phi} s$ for $\phi_s s \bullet \phi_o s$. We also implicitly assume that in every state s , the components s_s and s_o are *separate*; that is, $s_s \bullet s_o$ is defined in the PCM of the resource. Using the notation from Section 1.2, this is denoted as $s_s \perp s_o$.

Figure 1 illustrates the interaction among the state components of concurrent threads. Consider three threads, θ_1 , θ_2 , and θ_3 , running concurrently, and without any additional threads. Their respective states must have the forms $s_1 = (a_1, a_2 \bullet a_3)$, $s_2 = (a_2, a_3 \bullet a_1)$, and $s_3 = (a_3, a_1 \bullet a_2)$, because any of the two threads combined represent the concurrent environment of the third. Thus, the join of the *self*'s of any two threads must equal the *other* of the third. If θ is the parent thread of θ_1 and θ_2 , then its state is $s = (a_1 \bullet a_2, a_3)$, since θ is the combination of θ_1 and θ_2 , and has θ_3 as its environment. In particular, the join of the *self* and *other* components is invariant across all the threads. Figure 1 illustrates these relations. Moreover, we abbreviate the relationship of the state s of the parent thread θ and the states s_1 and s_2 of children threads θ_1 and θ_2 by $s = s_1 \star s_2$.

Morphisms and Separating Relations. In the types of lock and unlock, α computes the lock ownership information from s_s . It's therefore apparent that the types capture what's desired: that the lock program starts not owning the lock (precondition $\alpha_s s = \overline{\text{own}}$), and acquires the lock upon termination (postcondition $\alpha_s s = \text{own}$), and conversely for unlock. We'll see examples of other morphisms and separating relations shortly, when we discuss the internal definition of the state.

2.3 Internal State of the Ticket Lock

Recall that our goal is to define morphism α and reach the abstract specs (4) and (5) via concrete specs of the implementations of lock and unlock. To that end, we next design the ghost state of the ticket lock so that we can express the internal logical invariants needed for the typing derivations of the implementations of lock and unlock. Later, the morphism α will abstract these internals to an element of O . We use the following PCM U for the internals.

$$U = \mathbb{N}^+ \rightarrow_{\text{fin}} L \quad \text{where } L = \{\text{wait, serve, used}\} \quad (6)$$

Here, $\mathbb{N}^+ \rightarrow_{\text{fin}} L$ is the type of finite (partial) maps from positive natural numbers, representing tickets. Given a ticket t , the value of the map at t is one of the *labels* in the set L , denoting the status of the ticket according to the ticket lock workflow from Section 2.1: wait means that t has been drawn from the dispenser and the thread holding t is waiting to be called on the display; serve means that t has been called on the display and the thread has begun its turn holding the

⁴States in [Nanevski et al. 2019] also contain the third component s_j describing shared state, but we won't need it here.

lock; and used means t 's turn at the counter has finished, and the thread holding t has unlocked by signaling $t + 1$'s turn on the display. Notice that we don't throw away tickets, but just change their status in the map to reflect their progress through the bakery. The map thus serves as a form of history of the bakery. Similarly to heaps, the type $\mathbb{N}^+ \rightarrow_{\text{fin}} L$ is a PCM under the operation of disjoint union \cup of maps, which is undefined if the two operands share a ticket. The unit is the empty (i.e., nowhere defined) map \emptyset . We take \mathbb{N}^+ as the domain instead of \mathbb{N} in order to exclude the ticket 0, as the latter can't be drawn from the dispenser.

Given a ticket map $x \in U$ that represents the history of tickets in the bakery, we can compute out of x the ticket called on the display by the following definition, where we assume that \max of the empty set of natural numbers is by default the value 0.

$$\psi x = \max\{t \in \text{dom}(x) \mid x t = \text{used}\} + 1$$

Indeed, according to the workflow of ticket locks, upon finishing its turn with the lock a thread holding the ticket t sets the display to $t + 1$ to call the next thread in the queue. Thus, the value of the display, which at that point denotes the currently called ticket, is one larger than the largest used ticket in x . The function $\psi : U \rightarrow \mathbb{N}^+$ is our first example of a morphism, where we endow \mathbb{N}^+ with the PCM structure $(\mathbb{N}^+, \max, 1)$. Indeed, it's easy to see that for any two disjoint ticket maps $x \perp y$, we have

$$\psi(x \bullet y) = \max(\psi x, \psi y)$$

Moreover, $\psi \emptyset = 1$, and 1, being the smallest element of \mathbb{N}^+ , is the unit w.r.t. \max .

Morphism Notation for Ticket Locks. Let us name the identity morphism on U as σ . Giving a special name to the identity morphism will provide for uniform notation in our specs, where we apply σ , ψ and other morphisms to compute various values from states. In particular, when applying morphisms σ, ψ to state $s = (s_s, s_o)$, and according to the morphism notation from the previous section, we use the following expressions to denote various ticket maps and values.

- $\sigma_s s$ denotes the *self* map of tickets. These are the tickets, and their status, that the thread under consideration (henceforth “we” or “us”) has drawn from the dispenser.
- $\sigma_o s$ denotes the *other* ticket map. These are the tickets, and their status, that every other thread but “us” (henceforth “others”) has drawn from the dispenser.
- $\psi_s s$ denotes the *self* value of the last *called* ticket. This is the ticket that “we” have called by incrementing the display upon finishing our last turn at the counter, to call the next thread in the ticket queue.
- $\psi_o s$ is the ticket last called by “others”, when they finished their turns at the counter.

The combined ticket map $\hat{\sigma} s = \sigma_s s \bullet \sigma_o s$ and the value $\hat{\psi} s = \psi_s s \bullet \psi_o s = \max(\psi_s s, \psi_o s)$ have further important meanings. As tickets are drawn in order, we can compute the current value of the ticket dispenser pointer tdr as $\max(\text{dom}(\hat{\sigma} s)) + 1$. Similarly, we can compute the value of the display pointer dsp as $\hat{\psi} s$. Therefore, our specs needn't explicitly store the values of tdr and dsp , or any other shared state. In *specs*, any shared state can generally be computed out of *self* and *other* ghost components that suitably track the history of the updates to that shared state, just like $\hat{\sigma}$ and $\hat{\psi}$ compute the values of tdr and dsp out of the *self* and *other* ticket maps.⁵

2.4 Concrete Specs, Ghost Code, and Proof Outlines

With the internal state defined, we can next establish the following types for the implementations, in Section 2.1, of lock and unlock. The types are *concrete*, because they specify lock and unlock

⁵Of course, one needs to relate the ghost to concrete program state, shared or private, but that's beyond our scope here. We refer to [Nanevski et al. 2019] for more details on how this relationship is made in the ambient theory.

in terms of components of the underlying PCM U using morphisms σ and ψ , thus exposing the internal state of ticket lock. In the example, we denote by $t \mapsto l$ the singleton map that assigns label l to a ticket t , and is undefined elsewhere.

$$\text{lock} : \{\lambda s. \sigma_s s = \emptyset\} \{\lambda s. \sigma_s s = (\hat{\psi} s) \mapsto \text{serve}\} @ \text{TL}$$

$$\text{unlock} : [t]. \{\lambda s. \sigma_s s = t \mapsto \text{serve} \wedge t = \hat{\psi} s\} \{\lambda s. \sigma_s s = t \mapsto \text{used}\} @ \text{TL}$$

The spec for lock says that initially the ghost ticket map is empty. Thus, as customary in separation logic, it can be framed to any ticket map. Upon termination, we hold the ticket being displayed and this ticket is labeled as serve in our map ($\sigma_s s = (\hat{\psi} s) \mapsto \text{serve}$). Notice that the value $\hat{\psi} s$ in the postcondition is *stable* under interference, as other threads can't change the display because we hold the lock when we're being served. In particular, they can't change $\psi_o s$ which is a factor in the computation of $\hat{\psi} s = \psi_s s \bullet \psi_o s$.

The spec for unlock says that we hold the displayed ticket t ($\sigma_s s = t \mapsto \text{serve}$ and $t = \hat{\psi} s$). Upon termination, we still hold t , but it's now labeled as used, to indicate we finished our turn. The Hoare type for unlock explicitly binds the variable t , denoted by $[t]$, to snapshot the initial value of the display and to allow its use both in the precondition and the postcondition. The scope of t extends through the precondition and postcondition to the right of the binding $[t]$. We couldn't have ascribed to unlock the postcondition $\lambda s. \sigma_s s = \hat{\psi} s \mapsto \text{used}$ because the value $\hat{\psi} s$ in the postcondition isn't stable. Indeed, after we unlock, other threads can get their turn at the counter and increment the display. Thus, we use t to explicitly bind the stable value that the display has when unlock is invoked, and we hold the lock.

We emphasize how morphisms in the above specs combine in the standard mathematical fashion to compute various required values. For example, we apply ψ to s_s and s_o to obtain $\psi_s s$ and $\psi_o s$, and then combine the two into the expression $\hat{\psi} s = \psi_s s \bullet \psi_o s$, to define $\sigma_s s$. But for this to be possible, we had to make the binding of the state s *explicit* in the assertions of the Hoare triple, so that ψ and σ could be applied to the different projections of the same state s . Had we kept s *implicit*, as customary in separation logic assertions, expressing the above specs would have required somewhat more logical machinery. This convenience afforded by morphisms and explicit states extends to proof outlines, and to the definitions of resource transitions (see below) which relate two states, the input and output states, that are usefully differentiated by the explicit naming.

Transitions and State Space of the Resource TL. Before we can derive the types for lock and unlock, we need to annotate the programs with *ghost code*, i.e., code that manipulates the ghost state expressed in terms of σ and ψ . In our ambient type theory, the ghost code is formed by transitions of the resource (i.e., the state-transition system) of the specs; in the current example, transitions of the resource TL. In Figure 2 we show the three transitions that define TL: `taketx_tr`, `lock_tr`, and `unlock_tr`. Each is a relation over the initial state s and final state s' , and defines one of the three basic changes that ticket lock programs can perform over the state. We denote by $f[x \mapsto a]$ the function obtained by changing the value of function f at point x to the value a .

In the transition `taketx_tr`, the smallest undrawn ticket in the state s , $\max(\text{dom}(\hat{\sigma} s)) + 1$, is added into the *self* component $\sigma_s s'$ and labeled `wait`. Thus `taketx_tr` models a thread drawing a fresh ticket. In the transition `lock_tr`, the value $t = \hat{\psi} s$ is the ticket on display. This transition updates the ticket map σ at t from `wait` to `serve` to model that the thread noticed its ticket called on the display, and took its turn at the counter. The `unlock_tr` transition checks that the ticket t being displayed is owned by the thread and is being served: so $(\sigma_s s)(t) = \text{serve}$. The transition updates the status of t to `used` to model finishing the turn. Note that from the definition of ψ (Section 2.3), one can immediately compute that $\psi_s s' = t + 1$ because t is the largest used ticket in s' .

$$\begin{aligned}
\text{taketx_tr } s \ s' &\hat{=} \sigma_s s' = (t \Rightarrow \text{wait}) \cup \sigma_s s && \text{where } t = \max(\text{dom}(\hat{\sigma} s)) + 1 \\
\text{lock_tr } s \ s' &\hat{=} (\sigma_s s)(t) = \text{wait} \wedge \sigma_s s' = (\sigma_s s)[t \mapsto \text{serve}] && \text{where } t = \hat{\psi} s \\
\text{unlock_tr } s \ s' &\hat{=} (\sigma_s s)(t) = \text{serve} \wedge \sigma_s s' = (\sigma_s s)[t \mapsto \text{used}] && \text{where } t = \hat{\psi} s \\
s \in \Sigma_{\text{TL}} &\hat{=} \text{ordered}(\hat{\sigma} s) \wedge \text{no_gaps}(\hat{\sigma} s) \\
\text{where} &&& \\
\text{ordered}(x) &\hat{=} \{t \mid x(t) = \text{used}\} < \{t \mid x(t) = \text{serve}\} \wedge && \\
&\quad \{t \mid x(t) = \text{serve}\} < \{t \mid x(t) = \text{wait}\} \wedge && \\
&\quad \{t \mid x(t) = \text{used}\} < \{t \mid x(t) = \text{wait}\} && \\
\text{no_gaps}(x) &\hat{=} \forall t \in \mathbb{N}^+. t + 1 \in \text{dom}(x) \Rightarrow t \in \text{dom}(x)
\end{aligned}$$




Fig. 2. The state transition system TL. For sets S, T , the notation $S < T$ means $\forall s \in S, t \in T. s < t$.

We emphasize that resource transitions aren't arbitrary relations on states. Rather, as customary in separation logic [Calcagno et al. 2007], they must satisfy the important property of *locality*. The latter constrains the behavior of a transition under ownership transfer, and is necessary for the soundness of the rules of frame and parallel composition. The precise definition of locality in the subjective setting is given by the ambient type theory [Nanevski et al. 2019, Definition 3.5]. Here, we just mention that the three transitions of TL are all local, which we proved in the Coq code. Because locality constrains ownership transfer, these proofs essentially rely on the functions σ and ψ being morphisms. In Section 2.5 we illustrate how morphisms behave under ownership transfer and specifically under the rule of frame.

TL, being a state transition system, requires a state space in addition to transitions. The state space Σ_{TL} , given in Figure 2, is a subset of $U \times U$ that the transitions preserve. Thus, the state space imposes natural properties of ticket locks that: (1) tickets go through the bakery in order, i.e. all used tickets are smaller than serve tickets, which in turn are smaller than all wait tickets, as defined by the predicate `ordered` in Figure 2; and (2) tickets are drawn consecutively from the dispenser and none are skipped, as defined by the predicate `no_gaps` in Figure 2.

Ghost Code Annotation. We elide the discussion on how to formally factor transitions into the ghost code, and refer to the ambient type theory [Nanevski et al. 2019] for details. Instead, we decorate lock and unlock below to *informally* illustrate when the various transitions are invoked to change the ghost components of the state.

$$\begin{aligned}
\text{lock} &\hat{=} x \leftarrow \langle \text{inc_and_fetch}(tdr); \text{taketx_tr} \rangle; \\
&\quad \text{do } y \leftarrow !dsp \text{ until } x = y; \\
&\quad \langle \text{lock_tr} \rangle \\
\text{unlock} &\hat{=} \langle \text{inc_and_fetch}(dsp); \text{unlock_tr} \rangle
\end{aligned}$$

In the above code, angle brackets $\langle _ \rangle$ signify that the code they enclose executes atomically, that is without interference from other threads. In the first case, the value returned from the agglomeration of actual with ghost code is the value returned by the actual code itself. For example, the lock program executes `taketx_tr` atomically with the call to `inc_and_fetch(tdr)`, to bind to x the incremented value of tdr , and set the status of x in the ghost state to wait. When the condition $x = y$ is satisfied, since y is assigned the value of dsp , the ticket x is called on the display. The lock program then executes `lock_tr` as its final command to set the status of ticket x to serve. This

models taking the turn at the counter and completes the acquisition of the lock. Similarly, unlock executes `unlock_tr` to record in the ghost state that the display is incremented upon unlocking.

Proof Outlines. We next present the proof outline for lock and discuss its key points.

1. $\{\sigma_s s = \emptyset\}$
2. $x \leftarrow \langle \text{inc_and_fetch}(tdr); \text{taketx_tr} \rangle$
3. $\{\sigma_s s = x \Rightarrow \text{wait} \wedge \hat{\psi} s \leq x\}$
4. **do** $y \leftarrow !dsp$
5. $\{\sigma_s s = x \Rightarrow \text{wait} \wedge y \leq \hat{\psi} s \leq x\}$
6. **until** $x = y;$
7. $\{\sigma_s s = x \Rightarrow \text{wait} \wedge y = \hat{\psi} s = x\}$
8. $\langle \text{lock_tr} \rangle;$
9. $\{\sigma_s s = \hat{\psi} s \Rightarrow \text{serve}\}$

Line 1 is the precondition for lock. Line 3 shows that after the execution of `taketx_tr`, the drawn ticket x is the (only) ticket in $\sigma_s s$. Moreover, x is computed by `inc_and_fetch`, and hence is one larger than the last ticket drawn. More precisely, x is bound to $\max(\text{dom}(\hat{\sigma} s)) + 1$, for s taken at line 2. Now, from the definition of ψ , it must be that $\hat{\psi} s \leq x$ at the state s taken at line 2. Indeed, ψ computes the largest used ticket, and x equals the largest ticket, used or not. The property $\hat{\psi} s \leq x$ propagates to line 3 and beyond because it's stable under interference. Other threads can execute the transitions of TL over their own states to increase the display (by increasing $\psi_o s$ and thus also $\hat{\psi} s$), but can't increase the display beyond x . In [Farka et al. 2020b, Appendix B] we formally establish this stability property. For the do-until loop (lines 4-6), the loop invariant is on line 5: it conjoins the property that y is smaller than the displayed ticket ($y \leq \hat{\psi} s$). This property holds in the loop because line 4 stores the display value into y , after which the display may be further incremented by other threads. Line 7 marks the exit from the loop, thus the loop invariant holds together with the condition $x = y$ for exiting the loop. This immediately gives that $x = \hat{\psi} s$, which is a precondition for `lock_tr`. Finally, line 9 directly follows from line 7 by the definition of `lock_tr`.

2.5 Framing and Morphisms

The above spec for lock is in the small footprint style, where the spec's precondition uses \emptyset for $\sigma_s s$. A natural question is how this spec—which employs morphisms—can be lifted to large footprints. In other words, how do we employ the frame rule by using, as a frame, an arbitrary ticket map k for $\sigma_s s$ in the precondition? Framing is a standard operation in separation logic, but works somewhat differently in the setting with *self* and *other* variables, and in the presence of morphisms.

1. $\{\sigma_s s = k\}$
2. $\{\exists s_1 s_2. s = s_1 \star s_2 \wedge \sigma_s s_1 = \emptyset \wedge \sigma_s s_2 = k\}$
3. $\{((\lambda s. \sigma_s s = \emptyset) * (\lambda s. \sigma_s s = k))(s)\}$
4. lock
5. $\{((\lambda s. \sigma_s s = \hat{\psi} s \Rightarrow \text{serve}) * (\lambda s. \sigma_s s = k))(s)\}$
6. $\{\exists s_1 s_2. s = s_1 \star s_2 \wedge \sigma_s s_1 = \hat{\psi} s_1 \Rightarrow \text{serve} \wedge \sigma_s s_2 = k\}$
7. $\{\sigma_s s = (\hat{\psi} s \Rightarrow \text{serve}) \cup k\}$

On line 1, we start with $\sigma_s s = k$; thus $s = (k, s_o)$. Line 2 expands line 1 into a form suitable for applying the frame rule. It posits that s can be split into states s_1 and s_2 such as $s = s_1 \star s_2$. It's easy to see that this holds: we can represent $s_s = \emptyset \bullet k$, thus pick $s_1 = (\emptyset, k \bullet s_o)$ and $s_2 = (k, \emptyset \bullet s_o)$ (see Figure 1). Line 3 represents line 2 using separating conjunction, which isn't a primitive of our

logic, but is defined in the ambient theory in the customary way, modulo the use of subjective state splitting (Figure 1):

$$P * Q = \lambda s. \exists s_1 s_2. s = s_1 \star s_2 \wedge P s_1 \wedge Q s_2$$

Line 5 applies the frame rule to the intermediate spec for lock and the frame $(\lambda s. \sigma_s s = k)$ (given in **color** above). Line 6 unfolds the definition of separating conjunction, and line 7 collapses line 6, relying on the following two critical points.

The first critical point is that $\hat{\psi} s = \hat{\psi} s_1$. Indeed,

$$\begin{aligned} \hat{\psi} s &= \psi_s s \bullet \psi_o s && \text{(by definition of } \hat{\psi} \text{)} \\ &= (\psi_s s_1 \bullet \psi_s s_2) \bullet \psi_o s && \text{(because } s_s = s_{1_s} \bullet s_{2_s} \text{ by Figure 1, and } \psi \text{ distributes over } \bullet \text{)} \\ &= \psi_s s_1 \bullet (\psi_s s_2 \bullet \psi_o s) && \text{(by associativity of } \bullet \text{)} \\ &= \psi_s s_1 \bullet \psi_o s_1 && \text{(because } s_{1_o} = s_{2_s} \bullet s_o \text{ by Figure 1, and } \psi \text{ distributes over } \bullet \text{)} \\ &= \hat{\psi} s_1 && \text{(by definition of } \hat{\psi} \text{)} \end{aligned}$$

Notice that the proof of the property $\hat{\psi} s = \hat{\psi} s_1$ *doesn't rely on the definition of* ψ , but only on ψ being a morphism (with a trivial separating relation). Thus, the above is a general property of morphisms that follows because the join of *self* and *other* components are *invariant for parent and children states*. In this particular proof of lock, it allows replacing $\hat{\psi} s_1$ in line 6 with $\hat{\psi} s$ in line 7.

The second critical point is that $\sigma_s s = \sigma_s s_1 \cup \sigma_s s_2$. This holds because σ is a morphism, and $s = s_1 \star s_2$ implies that $s_s = s_{1_s} \bullet s_{2_s}$ by Figure 1, so σ can distribute over \bullet . Thus $\sigma_s s = (\hat{\psi} s \Rightarrow \text{serve}) \cup k$. Again, in this argument we didn't rely on the definition of σ .

2.6 Morphisms as Functional Abstractions

We next proceed to transform the concrete specs of lock and unlock into specs using a morphism $\alpha : U \rightarrow O$ to more abstractly express lock ownership. We define α as follows.

$$\alpha x \hat{=} \begin{cases} \text{own} & \text{if } t \Rightarrow \text{serve} \in \sigma x \text{ for some } t \\ \overline{\text{own}} & \text{otherwise} \end{cases} \quad (7)$$

As before, in the definition of α , one should think of x as the *self* component of a thread. Then the definition says that the thread owns the lock iff it holds a ticket labeled *serve* in the *self* set of tickets (σx) .

Structure Preservation and Partiality of α . Just like the morphism properties of σ and ψ were important for the internal specs to behave correctly under framing, so any spec using α requires α to be a morphism. And indeed, α satisfies the equation (3) from Section 1.2. In particular, $\alpha(x \bullet y)$ is defined and $\alpha(x \bullet y) = \alpha x \bullet \alpha y$ but only under the condition that x and y don't both contain a ticket labeled *serve*. In the latter case $\alpha x = \alpha y = \text{own}$ so their join is undefined.

A formal way to say this is that α is associated with the following separating relation, where $\#_{\text{serve}}(a)$ equals the number of *serve* tickets in the ticket map a .

$$x \perp_\alpha y \hat{=} \#_{\text{serve}}(\sigma x) + \#_{\text{serve}}(\sigma y) \leq 1 \wedge x \perp y \quad (8)$$

The definition directly captures that together x and y contain at most one served ticket. We shall see in Section 3.2 that \perp_α is indeed a separating relation, and moreover (Example 3.11) that $\#_{\text{serve}}(-)$ itself is a morphism, composed out of map filter and map counter functions, both of which are morphisms.

For now it suffices to observe that if we want to use α in a Hoare triple, then, at the very least, we must also attach the property $s_s \perp_\alpha s_o$ to the pre- and postcondition.⁶ Otherwise we won't be able to derive the framed Hoare triples generically, i.e., by using only the property that α is a morphism, without relying on α 's definition. Framing essentially relies on a morphism distributing over \bullet , as we've previously seen for σ and ψ , and the distribution of α is conditional upon $s_s \perp_\alpha s_o$.

Deriving Abstract Specs. We thus continue to establish the following abstract, but still intermediate, types of lock and unlock via α .

$$\begin{aligned} \text{lock} &: \{\lambda s. \alpha_s s = \overline{\text{own}} \wedge s_s \perp_\alpha s_o\} \{\lambda s. \alpha_s s = \text{own} \wedge s_s \perp_\alpha s_o\} @ \text{TL} \\ \text{unlock} &: \{\lambda s. \alpha_s s = \text{own} \wedge s_s \perp_\alpha s_o\} \{\lambda s. \alpha_s s = \overline{\text{own}} \wedge s_s \perp_\alpha s_o\} @ \text{TL} \end{aligned}$$

The derivations follow straightforwardly from the large footprint specs for lock and unlock. Below we just present the proof outline for lock; the one for unlock is in [Farka et al. 2020b, Appendix A].

1. $\{s_s \perp_\alpha s_o\}$
2. $\{\sigma_s s = k \wedge s_s \perp_\alpha s_o\}$
3. lock
4. $\{\sigma_s s = (\hat{\psi} s \Rightarrow \text{serve}) \cup k \wedge s_s \perp_\alpha s_o\}$
5. $\{\alpha_s s = \text{own} \wedge s_s \perp_\alpha s_o\}$

Line 1 weakens the desired precondition by eliding that $\alpha_s s = \overline{\text{own}}$, as this property isn't actually required by the proof. Indeed, if lock is invoked by a thread that already holds the lock, i.e., where $\alpha_s s = \text{own}$, the (partial correctness) Hoare triple for lock holds trivially because lock diverges. Line 2 snapshots $\sigma_s s$ into k , and gives the large footprint precondition for lock conjoined with $s_s \perp_\alpha s_o$. The latter property is an *invariant* of the resource TL. In other words, it isn't only stable under interference of other threads, but also it's preserved by the actions of our own thread, as we show in [Farka et al. 2020b, Appendix B]. In particular, $s_s \perp_\alpha s_o$ can strengthen the precondition and weaken the postcondition of any well-typed program that has TL as its resource type. In the ambient type theory [Nanevski et al. 2019] this is formally captured by a variant of the standard Hoare logic rule of consequence that admits the conjunction of resource invariants to the assertions in a Hoare triple. Therefore, line 4 follows from the large footprint spec for lock and the above invariance property. Finally, line 5 follows immediately by the definition of α .

2.7 Sub-PCM

Construction Overview. To obtain the ultimately desired compact specs (4) and (5) our algebraic approach provides the sub-PCM construction. The construction mods out the PCM U by \perp_α , to obtain a sub-PCM U_{TL} . Two ticket maps $x, y \in U_{\text{TL}}$ are considered disjoint only if $x \perp_\alpha y$, i.e., if x and y have at most one serve ticket in total.

$$U_{\text{TL}} \hat{=} U / \perp_\alpha$$

In U_{TL} , the \bullet operation restricts that of U so that $x \bullet_{U_{\text{TL}}} y$ equals $x \bullet_U y$ if $x \perp_\alpha y$, and is undefined otherwise. Consequently, $x \perp_{U_{\text{TL}}} y$ iff $x \perp_\alpha y$. Therefore, the relation \perp_α is the default notion of separateness in U_{TL} . It's thus assumed of every state, and doesn't need to be explicitly listed in any assertion.

It's essential for the sub-PCM construction that the condition by which we mod out be a separating relation, otherwise $\bullet_{U_{\text{TL}}}$ won't be commutative, associative, and admit a unit. But once we know that the condition is a separating relation, there is a generic proof (Section 3.4) that the

⁶We'll see in Section 4 that we'll also require α to be an *invertible* morphism, but that property is tied to α and needn't appear in Hoare triples.

construction results in a PCM. We also note that morphisms play a role in relating a PCM U and a sub-PCM U/R , for a separating relation R . As customary in algebraic definitions of substructures, the sub-PCM construction comes with two morphisms: *injection* $\iota : U/R \rightarrow U$ and *retraction* $\rho : U \rightarrow U/R$ that allow us to transfer values and reasoning between U and U/R . We explain the properties of ι and ρ in Section 3.4. In our case, the sub-PCM U_{TL} comes with the attendant injection $\iota_{\text{TL}} : U_{\text{TL}} \rightarrow U$, and retraction $\rho_{\text{TL}} : U \rightarrow U_{\text{TL}}$. The injection is a morphism with the separating relation $\perp_{U_{\text{TL}}}$, and the retraction is a morphism with separating relation \perp_{α} .

Use in Specifications. Once we obtain the restricted PCM U_{TL} , we can proceed to construct a *sub-resource* TL' which restricts the PCM U of TL to U_{TL} . The formal discussion of resources is given in [Nanevski et al. 2019]. Here, we just mention that TL' *simulates* TL , intuitively, because each transition of TL preserves \perp_{α} . The latter is easy to check: if in a state s the map $\hat{\sigma}$ has at most one used ticket, then so does a state s' obtained by executing one of the transitions of TL in s .

The ambient type theory provides an inference rule by which one can compositionally change the resource of a program from TL to TL' ,⁷ while precomposing the morphisms in the specs with the injection ι_{TL} . Thus, we can transform the previous specs using α into the ones given below where $\alpha' = \alpha \circ \iota_{\text{TL}}$. The condition $s_{\text{s}} \perp_{\alpha} s_0$ transforms into $s_{\text{s}} \perp_{U_{\text{TL}}} s_0$ and can thus be elided. This yields the specs we set out to obtain, modulo the renaming of α and TL into α' and TL' .

$$\begin{aligned} \text{lock} &: \{\lambda s. \alpha'_s(s) = \overline{\text{own}}\} \{\lambda s. \alpha'_s(s) = \text{own}\} @ \text{TL}' \\ \text{unlock} &: \{\lambda s. \alpha'_s(s) = \text{own}\} \{\lambda s. \alpha'_s(s) = \overline{\text{own}}\} @ \text{TL}' \end{aligned}$$

We emphasize that the simple Hoare specs are not the only benefit of the sub-PCM construction. By constructing U_{TL} , we not only restricted the states of TL , but we did so in a way that promoted \perp_{α} into the new default notion of separateness. Thus, we can now reason about \perp_{α} using the support that the ambient type theory provides for separateness in the form of lemma libraries and decision procedures, and which wouldn't have applied if \perp_{α} is simply listed as a conjunct in the assertions, and tracked as just another hypothesis in the proof state. We shall see in Section 3 that the ordered property can also be viewed as a separating relation and thus moved from the state space Σ_{TL} into the PCM by a sub-PCM construction. On the other hand, `no_gaps` doesn't admit such a move. We demonstrate in Section 3 that `no_gaps` isn't a separating relation; it doesn't generalize a disjointness relation between states of two threads, but rather represents a global property of the aggregated state of all threads.

It's also worth mentioning that we could have obtained the above specs in several alternative ways. For example, we could have started our example immediately by using U_{TL} instead of U . Correspondingly, instead of σ and ψ , we would have used $\sigma' = \sigma \circ \iota_{\text{TL}}$ and $\psi' = \psi \circ \iota_{\text{TL}}$ in our specs and proof outlines. The whole development that we carried out in this section then retraces easily. This shows that the approach is flexible enough to achieve the same specs and proofs by different order and arrangement.

We could also have chosen a different internal representation altogether. For example, we could have stored the tickets not into a map, but into three disjoint sets: one set for wait, one for serve, and one for used tickets, with the restriction that the set for serve tickets has at most one element. The algorithm would then shuffle tickets between sets to track the progress of the ticket through the bakery. A PCM implementing this alternative representation would be *isomorphic* to U_{TL} . But, to be able to formally speak of PCM isomorphism, one first has to have a notion of PCM *morphism*, as it applies to separation logic. Developing such a notion, along with the associated notion of separating relation, is the contribution of this paper.

⁷Or to any resource simulating TL .

3 PCM ABSTRACTIONS FORMALLY

3.1 Making Partiality Explicit

In previous sections, our discussion of partiality has been implicit, as we merely postulated that \bullet and various PCM morphisms are allowed to be undefined on some inputs. In practical formalization, however, it's useful to make partiality explicit by enriching the carriers with a new element that a function returns whenever it's supposed to be undefined. This is a common practice in theories of partial functions, e.g., domain theory [Abramsky and Jung 1995], and in symbolic execution in separation logic [Berdine et al. 2005]. We make a similar enrichment here as well.

Definition 3.1. A *topped partial commutative monoid* is a 5-tuple $(A, \bullet, \mathbb{1}, \top, D)$ such that \bullet is a *total* commutative and associative operation on A , with $\mathbb{1}$ as the unit. The element $\top \in A$ is the canonical *undefined* element, and $D \subseteq A$ is the set of *defined* elements. The two satisfy the following properties.

- (1) $\top \notin D$
- (2) $\mathbb{1} \in D$
- (3) if $x \bullet y \in D$ then $x, y \in D$
- (4) $x \bullet \top = \top \bullet x = \top$

We say that a topped PCM is *normal*, if $A = D \cup \{\top\}$, i.e., \top is the only undefined element.

The definition introduces the element \top which functions are supposed to return to signal undefinedness. For technical reasons that we explain below, we allow topped PCMs that have multiple undefined elements, but \top is a distinct one among them, and in particular, (1) $\top \notin D$. The above properties further say that (2) $\mathbb{1}$ is defined, and that (3) a join with an undefined element must be undefined. More strongly, (4) a join with \top must equal \top , i.e., \top is the *absorbing* element of A (also known as the *zero*). We continue to say that x and y are *separate*, written $x \perp y$, if $x \bullet y$ is defined, i.e., $x \bullet y \in D$. It's easy to see that $x \in A$ is defined iff it's separate from $\mathbb{1}$. Thus we write $x \perp \mathbb{1}$ instead of $x \in D$ to say that x is defined. As all the PCMs we consider in this paper are topped, we dispense with the adjective.

Example 3.2. The PCM $\mathbb{N}^+ \rightarrow_{\text{fin}} L$, which we used to represent the internal state of ticket locks is an instance of a more general PCM of finite maps. Finite maps $A \rightarrow_{\text{fin}} B$ form a topped normal PCM: take $(A \rightarrow_{\text{fin}} B) \cup \{\top\}$ as the carrier, $D = A \rightarrow_{\text{fin}} B$ for the defined elements, the empty map \emptyset as the unit, and the join defined as

$$f \bullet g \hat{=} \begin{cases} f \cup g & \text{if } f, g \neq \top \text{ and } f, g \text{ are maps with disjoint domains} \\ \top & \text{otherwise} \end{cases}$$

PROPOSITION 3.3. *Given (topped) PCMs A, B , the Cartesian product $A \times B$ is a (topped) PCM with $\bullet, \mathbb{1}$ and \top defined pointwise: $(a_1, b_1) \bullet (a_2, b_2) \hat{=} (a_1 \bullet_A a_2, b_1 \bullet_B b_2)$, $\mathbb{1} \hat{=} (\mathbb{1}_A, \mathbb{1}_B)$ and $\top \hat{=} (\top_A, \top_B)$, and the set of defined elements $D = D_A \times D_B$.*

The above proposition shows that $V = A \times B$ is a topped PCM whenever A and B are, but V isn't necessarily normal. Indeed, V contains elements of the form (a, \top_B) and (\top_A, b) , where $a \in D_A$ and $b \in D_B$. These elements can't be defined (hence, neither is in D_V), but they're all distinct from $\top_V = (\top_A, \top_B)$. The common way to avoid the proliferation of undefined elements in theories of partiality is to consider *smash products* instead of Cartesian products. In this paper, we instead allow PCMs such as Cartesian products that aren't normal. We also allow morphisms on them, e.g., the projections $\pi_1 : V \rightarrow A$ and $\pi_2 : V \rightarrow B$. When required, we rely on the sub-PCM construction (to be defined shortly) to normalize a PCM by removing the undefined elements other than \top .

3.2 Separating Relations

We next define *separating relations*, a strengthening of disjointness \perp of the underlying PCM. Having in mind that our specs apply separating relations to *self* and *other* components of a state, one can thus view separating relations as determining when two PCM values can be used to model the state of two concurrent threads.

Definition 3.4. Relation R on the carrier of the PCM A , is a *separating relation* if it satisfies the following laws which make R a structure-preserving relation on A .

- (1) (*definedness*) if $x R y$ then $x R \mathbb{1}$
- (2) (*strengthening*) if $x R y$ then $x \perp y$
- (3) (*unit*) $\mathbb{1} R \mathbb{1}$
- (4) (*symmetry*) $x R y$ iff $y R x$
- (5) (*associativity*) if $x R y$ and $(x \bullet y) R z$ then $x R (y \bullet z)$ and $y R z$

The law (1) restricts the separating relation R to defined elements only, as only a defined element should represent the state of a thread. Law (2) says that R strengthens the separating relation of the underlying PCM. Law (3) says that empty state is a valid state for any two threads, and law (4) says that the order in which threads appear in the relation is irrelevant.

The associativity law (5) describes when we can transfer ownership of state between two threads. Let's assume that we have two concurrent threads θ_1 and θ_2 . Correspondingly, their states are related by R . Let z be the state of θ_2 , and let θ_1 be a parent of two other concurrent threads with states x and y , respectively. Thus $x R y$ and $(x \bullet y) R z$. The law says that we can transfer y from θ_1 to θ_2 , which essentially corresponds to re-associating the child of θ_1 owning y to θ_2 . Intuitively, this is possible because the ordering and grouping of the threads in a thread pool is irrelevant.

Notice that from $x R y$ and $(x \bullet y) R z$, by symmetry of R and commutativity of \bullet , we get $y R x$ and $(y \bullet x) R z$, which by associativity implies $y R (x \bullet z)$ and $x R z$ as well. Thus, it's convenient to introduce the following notation for the antecedent of the associativity law:

$$x R y R z \hat{=} x R y \wedge (x \bullet y) R z$$

to say that x , y and z represent states of three concurrent threads, which are pairwise separate, and each is separate from the join of the other two.

PROPOSITION 3.5. *Let U be a PCM, and R a separating relation on U . Then $x R y$ implies $(x \bullet y) R \mathbb{1}$.*

PROOF. From $x R y$ we derive $\mathbb{1} R x$ by the definedness and symmetry laws for R , and $(\mathbb{1} \bullet x) R y$ because $\mathbb{1}$ is the unit. Then by associativity $\mathbb{1} R (x \bullet y)$, and by symmetry $(x \bullet y) R \mathbb{1}$. \square

The proof of Proposition 3.5 uses associativity, and we can explain the proposition using threads similarly to how we explained associativity. The proposition says: if x, y are valid states of two concurrent threads, then joining them produces a parent whose state $(x \bullet y)$ is valid as well.

Basic Examples of Separating Relations. The smallest separating relation of a PCM A is induced by A 's unit. We denote it $\perp_{\mathbb{1}_A}$, and define it by

$$x \perp_{\mathbb{1}_A} y \hat{=} x = \mathbb{1}_A \wedge y = \mathbb{1}_A$$

The relation clearly satisfies the required laws. Similarly, the PCM A itself induces the *trivial* separating relation \perp (or \perp_A when we want to make A explicit), defined as

$$x \perp_A y \hat{=} (x \bullet y) \text{ is defined}$$

This is the largest separating relation on A , since any larger relation violates the strengthening property (2). The intersection of two separating relations is also a separating relation. The *join*

relation J on A^2 defined as $(a_1, a_2) J (b_1, b_2)$ iff $(a_1 \bullet a_2) \perp (b_1 \bullet b_2)$ iff $(a_1 \bullet a_2 \bullet b_1 \bullet b_2) \perp \mathbb{1}_A$ is a separating relation.

Separating Relation \perp_α . Our formalization of ticket locks in Section 2 uses the invariant $x \perp_\alpha y$ on the PCM U , defined in (8) to state that at most one of the maps σx and σy holds the serve ticket. The property ensures that α is a morphism.

It's easy to see that \perp_α is a separating relation. The conditions (1-4) of Definition 3.4 are immediate. To show associativity, assume that $(x \bullet y) \perp_\alpha z$ (we don't need $x \perp_\alpha y$). Because $\#_{\text{serve}}$ distributes over \bullet (to be shown in a more general form in Example 3.11), the assumption gives us

$$(\#_{\text{serve}}(\sigma x) + \#_{\text{serve}}(\sigma y)) + \#_{\text{serve}}(\sigma z) \leq 1$$

But then $\#_{\text{serve}}(\sigma x) + (\#_{\text{serve}}(\sigma y) + \#_{\text{serve}}(\sigma z)) \leq 1$ and $\#_{\text{serve}}(\sigma y) + \#_{\text{serve}}(\sigma z) \leq 1$, i.e., $x \perp_\alpha (y \bullet z)$ and $y \perp_\alpha z$. We see that associativity in this example says that three threads may group in any way while preserving \perp_α because at most one of them can hold the serve ticket.

Separating Relation ordered. In Section 2.4, we defined the state space Σ_{TL} of the resource TL using the predicate ordered to capture that used tickets are smaller than serve ticket, which in turn is smaller than wait tickets. While ordered is defined as a predicate over a single PCM element $x \in U$, it easily lifts to a binary relation as follows:

$$x \omega y \hat{=} \text{ordered}(x \bullet y) \wedge x \perp y$$

It's easy to see that ω is a separating relation; again, the key property is associativity: $(x \omega y)$ and $(x \bullet y) \omega z$ imply $(y \omega z)$ and $x \omega (y \bullet z)$. By definition of ω , we must show: ordered $(x \bullet y)$ and ordered $(x \bullet y \bullet z)$ together imply ordered $(y \bullet z)$ and ordered $(x \bullet y \bullet z)$. This holds because if a map is ordered, then trivially, every subset of that map is ordered as well. The conjunct $x \perp y$ ensures the strengthening property (2) of Definition 3.4. Thus, our construction of TL' could have moved ordered from the state space Σ_{TL} into the definition of the sub-PCM U_{TL} .

Non-Example of Separating Relation: no_gaps. The state space Σ_{TL} also used the predicate no_gaps to capture that the tickets are drawn consecutively starting from ticket 1. Similarly to ordered, no_gaps can be made into a binary relation:

$$x v y \hat{=} \text{no_gaps}(x \bullet y) \wedge x \perp y$$

In contrast to ω , however, the relation v isn't associative, and hence isn't a separating relation. For example, let x, y and z be ticket maps with domains $\text{dom}(x) = \{2\}$, $\text{dom}(y) = \{1\}$, and $\text{dom}(z) = \{3\}$, respectively. Then $x v y$ and $(x \bullet y) v z$ hold because $\text{dom}(x \bullet y) = \{1, 2\}$ and $\text{dom}(x \bullet y \bullet z) = \{1, 2, 3\}$ contain consecutive tickets starting from ticket 1. But clearly $y \not v z$ because $\text{dom}(y \bullet z) = \{1, 3\}$ has a gap, missing the ticket 2 (see Figure 3).

In this sense, no_gaps is a *global* property. It holds of the collective state of all threads taken together, but, unlike separating relations, doesn't relate any two individual threads. In particular, no_gaps can't be moved from Σ_{TL} into U_{TL} . In other words, PCMs and separating relations encode local properties of thread states, while resource state spaces encode global ones.

3.3 Morphisms

Definition 3.6. A morphism $\phi : A \rightarrow B$ with a separating relation \perp_ϕ is a function ϕ from A to B such that.

- (1) (*preservation of unit*) $\phi \mathbb{1}_A = \mathbb{1}_B$
- (2) (*preservation of undefinedness*) $\phi \top_A = \top_B$
- (3) (*distributivity*) if $x \perp_\phi y$ then $\phi x \perp \phi y$, and $\phi(x \bullet y) = \phi x \bullet \phi y$

We say that ϕ is a *total* PCM morphism if \perp_ϕ equals the trivial separating relation \perp_A .




Fig. 3. Binary relation over predicate `no_gaps`. Adjacent elements satisfy the relation.

Some basic examples of morphisms include the identity function $\iota_A : A \rightarrow A$, which is a total morphism on A . So is the always-unit function $- \mapsto \mathbb{1}_B : A \rightarrow B$, as are the projections out of the product PCM. We also have the morphism $\text{join}_A : A^2 \rightarrow A$ defined as $\text{join}_A(a, b) = a \bullet b$, which is a morphism under the separating relation J (Section 3.2).

Morphisms and separating relations support a number of common algebraic constructions.

Definition 3.7. Let α, β be PCM morphisms. The *composition* $\alpha \circ \beta$, *tensor product* $\alpha \otimes \beta$, and *arrow product* $\alpha \times \beta$ are defined as below. All three are morphisms, under the given separating relations.

$$\begin{aligned}
 (\alpha \circ \beta) x &\hat{=} \alpha(\beta x) && \text{with} && x \perp_{\alpha \circ \beta} y &\hat{=} x \perp_{\beta} y \wedge \beta x \perp_{\alpha} \beta y \\
 (\alpha \otimes \beta) x &\hat{=} (\alpha x, \beta x) && \text{with} && x \perp_{\alpha \otimes \beta} y &\hat{=} x \perp_{\alpha} y \wedge x \perp_{\beta} y \\
 (\alpha \times \beta)(x_1, x_2) &\hat{=} (\alpha x_1, \beta x_2) && \text{with} && (x_1, x_2) \perp_{\alpha \times \beta} (y_1, y_2) &\hat{=} x_1 \perp_{\alpha} y_1 \wedge x_2 \perp_{\beta} y_2
 \end{aligned}$$

We can also define kernels, equalizers and restrictions of PCM morphisms, as customary in various algebraic theories. We don't apply these constructions in the ticket lock example but comment below why they are useful. Importantly, our theory is closed under these constructions, as equalizers and kernels of morphisms are separating relations, and a restriction of a morphism by a separating relation is a morphism. This shows that separating relations and morphisms are natural notions to consider together. Moreover:

THEOREM 3.8. *Morphism composition is associative, with the identity morphism as unit. Thus, the structure with (topped) PCMs as objects and PCM morphisms as arrows, forms a category.*

Definition 3.9. Let α and β be PCM morphisms. The kernel $\ker \alpha$ and equalizer $\text{eql } \alpha \beta$ are defined below. Both are separating relations.

- $x(\ker \alpha) y \hat{=} x \perp_{\alpha} y \wedge \alpha x = \alpha y = \mathbb{1}$, and
- $x(\text{eql } \alpha \beta) y \hat{=} x \perp_{\alpha} y \wedge x \perp_{\beta} y \wedge \alpha x = \beta x \wedge \alpha y = \beta y$.

Equalizers are useful whenever one wants to equate components of a PCM. For example, it's common in practice to have PCMs A and B , and to consider pairs $(a, b) \in A \times B$, but only under the restriction that $\phi(a) = \psi(b)$ for some morphisms $\phi : A \rightarrow X$ and $\psi : B \rightarrow X$. The morphisms ϕ and ψ would typically be projections, thus forcing that A and B are “stitched” along the projected components. The desired pairs are described by the equalizer $\text{eql}(\phi \times \iota_B)(\iota_A \times \psi)$. Kernels are a special case of equalizers, when one of the morphisms is the always-unit function.

Definition 3.10. A *restriction* of a morphism α with separating relation R is defined below. It's a morphism under the given separating relation.

$$(\alpha/R) x \hat{=} \begin{cases} \alpha x & \text{if } x R \mathbb{1} \\ \top & \text{otherwise} \end{cases} \quad \text{with} \quad x \perp_{\alpha/R} y \hat{=} x \perp_{\alpha} y \wedge x R y$$

Returning to our leading example of ticket lock, we can identify several other examples of morphisms that we used.

Section 2.6 mentioned that $\#_{\text{serve}}(-)$ is a morphism. This morphism, however, can be decomposed into two simpler morphisms.

Example 3.11. The function $\text{filter}_l : (A \multimap_{\text{fin}} B) \rightarrow (A \multimap_{\text{fin}} B)$ over a finite map selects only the entries that map to the label l .

$$\text{filter}_l m \hat{=} \{z \Rightarrow l \mid mz = l\} \quad \text{with } m \perp_{\text{filter}_l} n \hat{=} m \perp n$$

The filter function is a total morphism. Similarly, the counting function $\# : (A \multimap_{\text{fin}} B) \rightarrow \mathbb{N}$ computing the size of the domain of a finite map is a total morphism. Then, we can define $\#_{\text{serve}}(x) = (\# \circ \text{filter}_{\text{serve}})(x)$ and since morphisms compose, it's a morphism as well.

3.4 Sub-PCMs

In Section 2.7, we restricted the PCM U with a separating relation \perp_α . Formally, this construction is developed as a sub-object of a PCM, a *sub-PCM*. As customary in abstract algebra, we present the construction through two *morphisms* on the objects; *injection* of a sub-PCM into a PCM and a *retraction* from a PCM into its sub-PCM.

Definition 3.12. A PCM A is a *sub-PCM* of a PCM B if there exists a total PCM morphism $\iota : A \rightarrow B$ (an injection) and a morphism $\rho : B \rightarrow A$ (a retraction), such that:

- (1) $\rho(\iota x) = x$,
- (2) if $u \perp_\rho \mathbb{1}$ then $\iota(\rho u) = u$
- (3) if $\rho u \perp_A \rho v$ then $u \perp_\rho v$
- (4) if $\iota x \perp_B \mathbb{1}$ then $x \perp_A \mathbb{1}$

Property (1) says that ι is injective, i.e., if we coerce $x \in A$ into ιx , we can recover x back by ρ , since no other element of A maps by ι into ιx . This is a common property in sub-object constructions, and essentially says that ι embeds A into a subset of B . The dual property (2) allows the same for the elements of B that are defined according to \perp_ρ . Hence, A is in 1-1 correspondence with that subset of B . Property (3) extends the correspondence between A and the subset of B to their respective notions of disjointness. That is, \perp_A , when considered on images under ρ , implies (and hence, by properties of morphisms *equals*) \perp_ρ . Property (4) says that ι preserves the undefined elements, so that embedding A into B doesn't accidentally turn an undefined element into a defined one. A similar property of ρ is a consequence of (3). Finally, ι is total in order to embed the whole of A into B . A partial ι would embed only a subset of A into B , but that can be modeled by considering a total morphism from that subset into B .

As a simple example, we note that A is a sub-PCM of itself with the identity injection and retraction and trivial separating relations.

Definition 3.12 says what it means to be a sub-PCM abstractly, in terms of morphisms and separating relations. We next proceed to give a concrete construction that mods out a PCM U by a separateness relation R to obtain a PCM U/R , that is a sub-PCM of U according to Definition 3.12. It is this construction that we used in Section 2 to obtain the PCM U_{TL} out of U . The construction starts by defining the carrier set U/R , and the unit and \bullet as follows.

$$U/R \hat{=} \{z : U \mid zR \mathbb{1}_U\} \cup \{\top\} \quad \mathbb{1}_{U/R} \hat{=} \mathbb{1}_U \quad x \bullet_{U/R} y \hat{=} \begin{cases} x \bullet_U y & \text{if } xRy \\ \top & \text{otherwise} \end{cases}$$

Elements of the sub-PCM are the elements of U that are defined wrt. R , and the unit and \bullet are obtained by lifting the operations of U . Notice that the operations are well-defined. In particular, $\mathbb{1}_{U/R}$ is in the carrier set U/R , since $\mathbb{1}_U R \mathbb{1}_U$ by the properties of separating relations. Also, if $x, y \in U/R$, then $x \bullet_{U/R} y \in U/R$. This is proved by case analysis on whether x and y are defined

or \top . The interesting case is when they're defined and $x R y$. Then by Proposition 3.5, $(x \bullet y) R \mathbb{1}$, so the conclusion follows immediately.

LEMMA 3.13. *The definitions of $\mathbb{1}_{U/R}$ and $\bullet_{U/R}$ satisfy the following properties:*

- (1) $\bullet_{U/R}$ is commutative, i.e., $x \bullet_{U/R} y = y \bullet_{U/R} x$,
- (2) $\bullet_{U/R}$ is associative, i.e., $(x \bullet_{U/R} y) \bullet_{U/R} z = x \bullet_{U/R} (y \bullet_{U/R} z)$,
- (3) $\mathbb{1}_{U/R}$ is the unit for $\bullet_{U/R}$, i.e., $\mathbb{1}_{U/R} \bullet_{U/R} x = x \bullet_{U/R} \mathbb{1}_{U/R} = x$,
- (4) $\top \notin \{z : U \mid z R \mathbb{1}\}$,
- (5) $\mathbb{1} \in \{z : U \mid z R \mathbb{1}\}$, and
- (6) if $x \bullet_{U/R} y \in \{z : U \mid z R \mathbb{1}\}$ then $x, y \in \{z : U \mid z R \mathbb{1}\}$.

PROOF. (Sketch.) By easy analysis of the cases in the definition of U/R and $\bullet_{U/R}$. The proof essentially requires all the separating relation properties of R . For example, the commutativity property (1) relies on the symmetry of R , the associativity property (2) relies on the associativity of R , and the unit property (3) relies on the unit law of separating relations. Thus, separating relation laws are directly obtained as a requirement for proving this lemma. \square

Thus, by the above lemma, we have a PCM:

$$(U/R, \bullet_{U/R}, \mathbb{1}_{U/R}, \top, \{z : U \mid z R \mathbb{1}\})$$

It remains to show that this PCM is a sub-PCM in the sense of Definition 3.12. To that purpose, we define the two necessary morphisms:

$$\iota x \hat{=} x \text{ with } \perp_\iota \hat{=} \perp_{U/R} \quad \rho u \hat{=} \begin{cases} u & \text{if } u R \mathbb{1} \\ \top & \text{otherwise} \end{cases} \text{ with } \perp_\rho \hat{=} R$$

Note that these functions are indeed morphisms. That ι and ρ preserve unit and \top is trivial to show, and so is that $\iota(x \bullet y) = \iota x \bullet \iota y$ when $x \perp_\iota y$. It remains to show that $\rho(u \bullet v)$ is defined and $\rho(u \bullet v) = \rho u \bullet \rho v$, if $u R v$. To see this, assume $u R v$ and observe that from the law of defined elements of separating relation, this implies $u R \mathbb{1}$ and $v R \mathbb{1}$. Thus $\rho u = u$ and $\rho v = v$ and so $\rho u \bullet \rho v = u \bullet v$. By Proposition 3.5, we also have $(u \bullet v) R \mathbb{1}$; thus $(u \bullet v)$ is defined and $\rho(u \bullet v)$ equals $u \bullet v$, concluding that ρ is a morphism. Now it's also easy to see that the injection ι is total (by definition, since it has the trivial separating relation), and that ι and ρ satisfy the requirements of Definition 3.12. Therefore U/R is a sub-PCM of U .

We conclude this section by noticing that U/R is a normal PCM, since \top is its only undefined element. Thus, we can use the sub-PCM construction to normalize PCMs, when desirable. Given a non-normal PCM A , the PCM A/\perp_A is normal and contains all the defined elements of A .

3.5 Histories, Morphisms, and Separating Relations

This section illustrates how PCM morphisms and separating relations apply to reasoning about data structures specified via time-stamped histories. Histories are a common and general abstraction in concurrency, used, for example, in the formulation of consistency criteria such as linearizability [Herlihy and Wing 1990]. Here, we specifically focus on their application to locking.

An abstract locking history of a thread is a finite map from timestamps represented by positive natural numbers to set $\text{Op} = \{L, U\}$, i.e. $\text{Hist} = \mathbb{N}^+ \rightarrow_{\text{fin}} \text{Op}$. If a thread's history has the value L at timestamp t , that signifies that the thread has *locked* at time t . Similarly, if the value is U then the thread has *unlocked* at time t . If the history of a thread is undefined at t , then the thread was idle at that moment, and some other thread may have locked or unlocked at time t . We overload the notation from Section 2.4 and write $t \Rightarrow L$ (resp. $t \Rightarrow U$) for a singleton history containing only

the timestamp t with the locking (resp. unlocking) operation. Histories form a PCM under disjoint union, with the nowhere defined map (i.e., empty history) as unit.

Using PCM Hist In a Resource. Let us assume that we have defined an internal state of some (unspecified) locking algorithm and a corresponding resource with a state space and transitions, similar to how we defined the resource TL in Section 2. Then we will typically have a morphism, which we name τ here, that projects the history component of the underlying state; that is $\tau_s s$ is the history of “our” thread in state s , and $\tau_o s$ is the history of all “other” threads combined.

Moreover, we will also typically use the resource state space to specify global properties of histories, similar to `no_gaps` from Figure 2. For example, at the very least, we will require that the global history $\hat{\tau} s$ alternates the operations L and U, i.e., $\hat{\tau} s$ has the form $\hat{\tau} s = 1 \Rightarrow L \bullet 2 \Rightarrow U \bullet 3 \Rightarrow L \bullet \dots$. The alternation property specifies the key relationship between locking and unlocking, but doesn’t form a separating relation itself. However, as we show presently, there’s an important separating relation \perp_ω over locking histories that gives rise to a morphism $\omega : \text{Hist} \rightarrow O$ for computing lock ownership out of a thread’s history.

Histories of Exclusive Locking. Note that histories, even with the alternation property imposed, don’t exclude the possibility that one thread may take the lock, which is then released by another thread. In our subjective setting, we may represent the situation as follows

$$\tau_s(s)(t) = L \quad \tau_o(s)(t+1) = U$$

The equations say that we have locked at time t , but another thread has unlocked at $t+1$. Modeling such behavior is desirable because there exist locking algorithms that admit it. For example, even simple spin locks physically allow that the locking and unlocking threads are different. Readers-writers locks [Courtois et al. 1971], which can be built over spin locks, allow an initial reader thread to acquire a lock and a possibly different reader thread to release it. In a setting where a lock can be released by any thread, one can’t really speak about lock ownership. Thus, structures that admit such behavior and that can utilize the general histories above, typically require richer internal ghost state in order to specify the desired locking discipline. For example, readers-writes locks require enrichment with permissions [Bornat et al. 2005], which we forego here.

Nevertheless, even without enrichment, we can already illustrate how to impose on locking histories a more restricted behavior, whereby the thread that unlocks must be the one that currently holds the lock. Such “mutually exclusive” histories form a sub-PCM of general locking histories, and thus the property of mutual exclusion can be captured as a separating relation. Analogous to the ticket lock example, we then construct the morphism ω that computes lock ownership.

Let us first define the separating relation:

$$\begin{aligned} x \perp_\omega y \hat{=} & (\forall t. x(t) = L \Rightarrow \max(\text{dom}(x \bullet y)) \leq t \vee x(t+1) = U) \wedge \\ & (\forall t. y(t) = L \Rightarrow \max(\text{dom}(y \bullet x)) \leq t \vee y(t+1) = U) \wedge x \perp y \end{aligned} \quad (9)$$

Intuitively, the relation states that whenever the thread with history x locked at time t then the thread with history y couldn’t have proceeded. On the other hand, the thread with history x could have proceeded by unlocking at the immediate time $t+1$. Similarly to the separating relations in the previous sections, the relation symmetrically applies to the history y as well, and requires that the join of x and y be valid, i.e., that the histories of two threads don’t share timestamps.

LEMMA 3.14. *The relation \perp_ω is a separating relation.*

PROOF. The proof shows that properties of Definition 3.4 hold:

- (1) Definedness: Assume that $x \perp_{\omega} y$, and show that $x \perp_{\omega} \mathbb{1}$. Indeed, consider t such that $x(t) = L$. By case analysis on t , either $t = \max(\text{dom}(x)) = \max(\text{dom}(x \bullet \mathbb{1}))$, or $t \leq \max(\text{dom}(s)) = \max(\text{dom}(x \bullet \mathbb{1}))$. In either case, trivially $x \perp_{\omega} \mathbb{1}$.
- (2) Strengthening: Follows trivially from definition of \perp_{ω} .
- (3) Unit property: follows from symmetry (4) and definedness (1).
- (4) Symmetry: Immediate from symmetry of conjunction and symmetry of $x \perp y$.
- (5) Associativity: Assume $(x \bullet y) \perp_{\omega} z$ and $x \perp_{\omega} y$ to show $y \perp_{\omega} z$ and $x \perp_{\omega} (y \bullet z)$. Let t be a timestamp such that $y(t) = L$ (the cases when $x(t) = L$ or $z(t) = L$ are similar). Then from $x \perp_{\omega} y$ we get that $\max(\text{dom}(x \bullet y)) \leq t$ (and more specifically $\max(\text{dom}(x)) < t$), or $y(t+1) = U$. In the first case, it must be $(x \bullet y)(t) = L$. Thus from $(x \bullet y) \perp_{\omega} z$, we infer that either $\max(\text{dom}(z)) < t$, and thus $y \perp_{\omega} z$ and $x \perp_{\omega} (y \bullet z)$, or $(x \bullet y)(t+1) = U$, which implies that $y(t+1) = U$, which we consider as part of the second case. In the second case, i.e., when $y(t+1) = U$, the property $y \perp_{\omega} z$ is immediate. On the other hand, we also have $(y \bullet z)(t) = L$ and $(y \bullet z)(t+1) = U$, thus $x \perp_{\omega} (y \bullet z)$ holds as well.

□

Finally, we can define the morphism $\omega : \text{Hist} \rightarrow O$.

$$\omega h \hat{=} \begin{cases} \top & \text{if } h = \top \\ \text{own} & \text{if } \max(\text{dom}(h)) > 0 \text{ and } h(\max(\text{dom}(h))) = L \\ \overline{\text{own}} & \text{otherwise} \end{cases} \quad (10)$$

LEMMA 3.15. *The map ω is a morphism with separating relation \perp_{ω} .*

PROOF. The properties of Definition 3.6 hold as follows:

- (1) Map ω clearly preserves unit since $\max(\text{dom}(\mathbb{1}_{\text{Hist}})) = 0$; thus $\omega(\mathbb{1}_{\text{Hist}}) = \overline{\text{own}} = \mathbb{1}_O$.
- (2) Undefinedness is preserved trivially.
- (3) To show distributivity, assume that $x \perp_{\omega} y$ and let $t = \max(\text{dom}(x \bullet y))$. We consider only the interesting case when $t > 0$, and w.l.o.g., $t \in x$, and $x(t) = L$. Then by definition of ω , $\omega x = \text{own}$. But it must also be that $\omega y = \overline{\text{own}}$, for if otherwise, then by $x \perp_{\omega} y$, the history y must have an unlocking entry at time $t+1$ and thus contains a timestamp beyond $t = \max(\text{dom}(x \bullet y))$. Therefore $\omega(x \bullet y) = \text{own} = \omega x \bullet \omega y$.

□

4 INVERTIBLE MORPHISMS AND SEPARATING RELATIONS

4.1 Invertibility of Morphisms

As we have seen in the previous sections, the key property of a morphism $\phi : A \rightarrow B$ is ϕ distributes over \bullet . In other words, if the *argument* of ϕ splits into $a_1 \bullet a_2$, then the result splits as well, that is:

$$\phi(a_1 \bullet a_2) = \phi a_1 \bullet \phi a_2$$

under a suitable condition on a_1 and a_2 expressed as a separating relation $a_1 \perp_{\phi} a_2$.

In verification practice, however, we often have to show the converse: that if the *result* of ϕ is defined and splits into $b_1 \bullet b_2$, then the argument must split as well, that is:

$$a \perp_{\phi} \mathbb{1} \wedge \phi a = b_1 \bullet b_2 \Rightarrow \exists a_1 a_2. a = a_1 \bullet a_2 \wedge a_1 \perp_{\phi} a_2 \wedge \phi a_1 = b_1 \wedge \phi a_2 = b_2 \quad (11)$$

We call this property *invertibility*, because it can be seen as imposing a form of distributivity on the *inverse image* $\phi^{-1} : \mathcal{P}(B) \rightarrow \mathcal{P}(A)$, where we take only inverses that are separate from $\mathbb{1}$ by \perp_{ϕ} , i.e. $\phi^{-1}(X) = \{x \in A \mid \phi x \in B \wedge x \perp_{\phi} \mathbb{1}\}$. Indeed, property (11) can be restated compactly as

$$\phi^{-1}\{b_1 \bullet b_2\} \subseteq (\phi^{-1}\{b_1\}) \bullet_{\perp_{\phi}} (\phi^{-1}\{b_2\}) \quad (12)$$

where $\bullet_{\perp\phi}$ is a special case of the more general operation \bullet_R that *lifts* a separating relation R on A to an operation on sets $X_1, X_2 \in \mathcal{P}(A)$ as follows.

$$X_1 \bullet_R X_2 = \{a_1 \bullet a_2 \mid a_1 \in X_1 \wedge a_2 \in X_2 \wedge a_1 R a_2\}$$

Invertibility of morphisms appears naturally in separation logic when reasoning by framing or parallel composition. For example, imagine a program e with the following spec, similar to our abstract spec for lock.

$$e : \{\lambda s. \phi_s s = b_1\} \{\lambda s. \phi_s s = b'_1\}$$

Here ϕ is a total morphism (i.e., ϕ has a trivial separating relation), and we want to frame it by $(\lambda s. \phi_s s = b_2)$. The direct application of the frame rule, unfolding the definition of \star that we introduced in Section 2, derives

$$e : \{\lambda s. \exists s_1 s_2. s = s_1 \star s_2 \wedge \phi_s s_1 = b_1 \wedge \phi_s s_2 = b_2\} \\ \{\lambda s. \exists s_1 s_2. s = s_1 \star s_2 \wedge \phi_s s_1 = b'_1 \wedge \phi_s s_2 = b_2\}$$

Of course, we would like to strengthen the precondition and weaken the postcondition of this spec into the more compact and ultimately desirable form

$$e : \{\lambda s. \phi_s s = b_1 \bullet b_2\} \{\lambda s. \phi_s s = b'_1 \bullet b_2\}$$

Here's where invertibility comes in. It's easy to see that the postcondition readily weakens into the desired form just by using that ϕ is a (total) morphism, and the fact that $s_s = s_{1s} \bullet s_{2s}$. However, the precondition doesn't strengthen immediately. We need to show

$$\phi_s s = b_1 \bullet b_2 \implies \exists s_1 s_2. s = s_1 \star s_2 \wedge \phi_s s_1 = b_1 \wedge \phi_s s_2 = b_2$$

but this doesn't follow from distributivity of ϕ . It does follow, however, if ϕ is invertible. To see this, assume that $s_s = a$. Then $\phi_s s = b_1 \bullet b_2$ transforms into $\phi a = b_1 \bullet b_2$. From the assumption that $s_s \perp s_o$ and the properties of separating relations, we get $s_s \perp \mathbb{1}$ and thus $a \perp \mathbb{1}$ as well. Then invertibility of ϕ gives us a_1 and a_2 such that $a = a_1 \bullet a_2$, $a_1 \perp_\phi a_2$ (which equals $a_1 \perp a_2$ because ϕ is a total morphism), $\phi a_1 = b_1$, and $\phi a_2 = b_2$. Choosing $s_1 = (a_1, a_2 \bullet s_o)$ and $s_2 = (a_2, a_1 \bullet s_o)$ gives us $s = s_1 \star s_2$ such that $\phi_s s_1 = \phi a_1 = b_1$ and $\phi_s s_2 = \phi a_2 = b_2$. This strengthens the precondition as desired.

4.2 Invertibility of Separating Relations

Similar style of reasoning applies if ϕ isn't total, but has a non-trivial separating relation \perp_ϕ . It turns out, however, that then we need to impose an additional condition of \perp_ϕ , thus giving rise to a notion of *invertible* separating relations also. To see what this condition should be, imagine that we have a program e with the following spec, similar to our intermediate abstract spec for lock.

$$e : \{\lambda s. \phi_s s = b_1 \wedge s_s \perp_\phi s_o\} \{\lambda s. \phi_s s = b'_1 \wedge s_s \perp_\phi s_o\}$$

Because ϕ is not total, we include the conjunct $s_s \perp_\phi s_o$ into the spec to ensure that ϕ distributes when framed. We now want to frame with $(\lambda s. \phi_s s = b_2 \wedge s_s \perp_\phi s_o)$. Similarly to the previous Section 4.1, unfolding the definition of \star derives us the following spec:

$$e : \{\lambda s. \exists s_1 s_2. s = s_1 \star s_2 \wedge \phi_s s_1 = b_1 \wedge s_{1s} \perp_\phi s_{1o} \wedge \phi_s s_2 = b_2 \wedge s_{2s} \perp_\phi s_{2o}\} \\ \{\lambda s. \exists s_1 s_2. s = s_1 \star s_2 \wedge \phi_s s_1 = b'_1 \wedge s_{1s} \perp_\phi s_{1o} \wedge \phi_s s_2 = b_2 \wedge s_{2s} \perp_\phi s_{2o}\} \quad (13)$$

However, we ultimately desire to obtain a compact spec in the following form:

$$e : \{\lambda s. \phi_s s = b_1 \bullet b_2 \wedge s_s \perp_\phi s_o\} \{\lambda s. \phi_s s = b'_1 \bullet b_2 \wedge s_s \perp_\phi s_o\} \quad (14)$$

As before, we need to prove two implications to weaken (13) to (14).

$$\phi_s s = b_1 \bullet b_2 \wedge s_s \perp_\phi s_o \Rightarrow \quad (15)$$

$$\begin{aligned} & \exists s_1 s_2. s = s_1 * s_2 \wedge \phi_s s_1 = b_1 \wedge s_{1s} \perp_\phi s_{1o} \wedge \phi_s s_2 = b_2 \wedge s_{2s} \perp_\phi s_{2o} \\ & s = s_1 * s_2 \wedge \phi_s s_1 = b'_1 \wedge s_{1s} \perp_\phi s_{1o} \wedge \phi_s s_2 = b_2 \wedge s_{2s} \perp_\phi s_{2o} \Rightarrow \quad (16) \\ & \phi_s s = b'_1 \bullet b_2 \wedge s_s \perp_\phi s_o \end{aligned}$$

Or alternatively, if we replace the state variables by pairs of *self* and *other* components, e.g., $s = (a, a')$, $s_1 = (a_1, a_2 \bullet a')$, $s_2 = (a_2, a_1 \bullet a')$, we obtain after some simplification:

$$\phi a = b_1 \bullet b_2 \wedge a \perp_\phi a' \Rightarrow \quad (17)$$

$$\begin{aligned} & \exists a_1 a_2. a = a_1 \bullet a_2 \wedge \phi a_1 = b_1 \wedge a_1 \perp_\phi (a_2 \bullet a') \wedge \phi a_2 = b_2 \wedge a_2 \perp_\phi (a_1 \bullet a') \\ & \phi a_1 = b'_1 \wedge a_1 \perp_\phi (a_2 \bullet a') \wedge \phi a_2 = b_2 \wedge a_2 \perp_\phi (a_1 \bullet a') \Rightarrow \quad (18) \\ & \phi (a_1 \bullet a_2) = b'_1 \bullet b_2 \wedge (a_1 \bullet a_2) \perp_\phi a' \end{aligned}$$

If we assume that ϕ is invertible, then from $a \perp_\phi a'$, we get $a \perp_\phi \mathbb{1}$ by the definedness property of separating relations, and then (17) immediately follows by associativity of separating relations. However, to obtain the first conjunct in the conclusion of (18), we require that $a_1 \perp_\phi a_2$, so that we can distribute ϕ over $a_1 \bullet a_2$ and then use that $\phi a_1 = b_1$ and $\phi a_2 = b_2$. To obtain the second conjunct in (18), we need to reassociate a_1, a_2 and a' , which can be done if $a_1 \perp_\phi a_2 \perp_\phi a'$. Thus, we obtain the required condition that makes it possible to derive (14).

$$a_1 \perp_\phi (a_2 \bullet a') \wedge a_2 \perp_\phi (a_1 \bullet a') \Rightarrow a_1 \perp_\phi a_2 \perp_\phi a'$$

To establish this implication it suffices to show that either $a_1 \perp_\phi a'$ or $a_2 \perp_\phi a'$ as the consequent $a_1 \perp_\phi a_2 \perp_\phi a'$ then follows from associativity of separating relations.

4.3 Duality of Invertibility of Morphisms and Separating Relations

We note an interesting duality in the interplay of ϕ and \perp_ϕ in the above framing process. When strengthening the precondition, it's the invertibility of ϕ that provides the split of a into $a_1 \bullet a_2$ such that $a_1 \perp_\phi a_2$, which is then used to reassociate \perp_ϕ . When weakening the postcondition, the situation is dual. We start with a already split into $a = a_1 \bullet a_2$, but it's the invertibility of \perp_ϕ that ensures the split is such that ϕ can distribute over it. Thus, in the precondition, ϕ helps \perp_ϕ and in the postcondition \perp_ϕ helps ϕ .

Thus, to summarize, we have the following definitions of invertibility for separating relations and morphisms that enable framing in the abstract of above form, i.e., without relying on the definitions of morphism or its separating relation.

Definition 4.1. A separating relation R on the PCM A is *invertible* if for all a_1, a_2, a' such that $a_1 R (a_2 \bullet a')$ and $a_2 R (a_1 \bullet a')$, it must also be $a_1 R a_2 R a'$. Moreover, it suffices to prove $a_1 R a'$ or $a_2 R a'$, as $a_1 R a_2 R a'$ follows by associativity.

Definition 4.2. A morphism $\phi : A \rightarrow B$ is *invertible* if \perp_ϕ is an invertible separating relation and for all $a \in A$ such that $a \perp_\phi \mathbb{1}$, and $b_1, b_2 \in B$ where $\phi a = b_1 \bullet b_2$, there exist $a_1, a_2 \in A$, such that $a = a_1 \bullet a_2$, $a_1 \perp_\phi a_2$, $\phi a_1 = b_1$ and $\phi a_2 = b_2$.

We now demonstrate the invertibility of various constructions we introduced earlier. First, separating relations of total morphisms are always invertible.

PROPOSITION 4.3. *Let A be a PCM. The trivial separating relation \perp_A is invertible.*

PROOF. Let $a_1 \perp_A (a_2 \bullet a')$. Recall that the trivial separating relation is given by $x \perp_A y \hat{=} (x \bullet y)$ is defined. Hence we obtain that $a_1 \bullet (a_2 \bullet a')$ is defined, and, using commutativity and associativity of join \bullet , we have that $(a_1 \bullet a') \bullet a_2$ is defined. Thus, by law (3) of Definition 3.1, also $a_1 \bullet a'$ is defined whence $a_1 \perp_A a'$. Similarly for $a_2 \perp_A a'$. \square

Similarly, other basic constructions on separating relation preserve invertibility. So do the construction on morphisms. Recall the composition of morphisms, tensor and arrow product in Definition 3.7.

PROPOSITION 4.4. *Let α, β be invertible morphisms. Then $\alpha \circ \beta$ and $\alpha \times \beta$ are invertible morphisms.*

PROOF. We just show the case for \circ as the one for \times is simple. Let $\alpha : C \rightarrow A, \beta : B \rightarrow C$ be invertible morphisms. First, we show that $\perp_{\alpha \circ \beta}$, the induced separating relation of $\alpha \circ \beta$, is invertible. Consider a_1, a_2, a' , such that $a_1 \perp_{\alpha \circ \beta} (a_2 \bullet a')$, and $a_2 \perp_{\alpha \circ \beta} (a_1 \bullet a')$. We need to show that $a_1 \perp_{\alpha \circ \beta} a_2 \perp_{\alpha \circ \beta} a'$. From Definition 3.7 of composition, we obtain $a_1 \perp_\beta (a_2 \bullet a') \wedge \beta a_1 \perp_\alpha \beta (a_2 \bullet a')$, and $a_2 \perp_\beta (a_1 \bullet a') \wedge \beta a_2 \perp_\alpha \beta (a_1 \bullet a')$. Since β is an invertible morphism also its separating relation \perp_β is invertible. We use invertibility of \perp_β and the first conjunct to obtain that $a_1 \perp_\beta a_2 \perp_\beta a'$. Now considering the second conjuncts, we get by distributivity of β , $\beta a_1 \perp_\alpha (\beta (a_2) \bullet \beta (a'))$ and $\beta a_2 \perp_\alpha (\beta (a_1) \bullet \beta (a'))$. Because α is invertible, so is \perp_α . We therefore obtain $\beta a_1 \perp_\alpha \beta a_2 \perp_\alpha \beta a'$. Thus $\perp_{\alpha \circ \beta}$ is invertible.

Second, we show that $\alpha \circ \beta$ is an invertible morphism. Assume $b \in B$ and $a_1, a_2 \in A$ such that $(\alpha \circ \beta) b = a_1 \bullet a_2$. Also assume that $b \perp_{\alpha \circ \beta} \mathbb{1}$; that is, $b \perp_\beta \mathbb{1}$ and $\beta b \perp_\alpha \mathbb{1}$. Using invertibility of α on $\beta b \in C$, we obtain $c_1, c_2 \in C$, such that $\beta b = c_1 \bullet c_2, c_1 \perp_\alpha c_2, \alpha c_1 = a_1$, and $\alpha c_2 = a_2$. Using invertibility of β on b we further obtain $b_1, b_2 \in B$, such that $b = b_1 \bullet b_2, b_1 \perp_\beta b_2, \beta b_1 = c_1$, and $\beta b_2 = c_2$. Consequently $\beta b_1 \perp_\alpha \beta b_2$. Hence, using $b_1 \perp_\beta b_2$, we obtain $b_1 \perp_{\alpha \circ \beta} b_2$. Finally, $(\alpha \circ \beta) b_1 = \alpha (\beta b_1) = \alpha (c_1) = a_1$ and similarly for a_2 . Therefore, we have $b_1, b_2 \in B$ such that $b = b_1 \bullet b_2, b_1 \perp_{\alpha \circ \beta} b_2, (\alpha \circ \beta) b_1 = a_1$ and $(\alpha \circ \beta) b_2 = a_2$. Hence morphism $\alpha \circ \beta$ is invertible. \square

Notice that $\alpha \otimes \beta$ is an example of a morphism that isn't necessarily invertible, even if α and β are. By definition, $(\alpha \otimes \beta) x = (\alpha x, \beta x)$. Thus, if we're given $(\alpha \otimes \beta) x = (y, z)$, we can induce one split of x by α and y , and another by β and z . However, there's no reason to expect that these splits are equal, which is required for $\alpha \otimes \beta$ to be invertible.

We also introduced the notions of kernel and equalizer, which are separating relations. These illustrate constructions that turn invertible morphisms into invertible separating relations.

PROPOSITION 4.5. *Let α, β be morphisms with invertible separating relations. Then $\ker \alpha$ is an invertible separating relation, while $\text{eql } \alpha \beta$ is so if the range PCM of α and β is cancellative.⁸*

PROOF. We show the proof for equalizers. A kernel is a special case of an equalizer when one of the morphisms is the always-unit one, which circumvents the need for cancellativity. Consider x, y, z , such that $x (\text{eql } \alpha \beta) (y \bullet z)$ and $y (\text{eql } \alpha \beta) (x \bullet z)$. By Definition 4.1, it suffices to show $y (\text{eql } \alpha \beta) z$; that is $y \perp_\alpha z \wedge y \perp_\beta z \wedge \alpha y = \beta y \wedge \alpha z = \beta z$. From the assumptions, we get $x \perp_\alpha (y \bullet z) \wedge x \perp_\beta (y \bullet z) \wedge \alpha x = \beta x \wedge \alpha (y \bullet z) = \beta (y \bullet z)$, and $y \perp_\alpha (x \bullet z) \wedge y \perp_\beta (x \bullet z) \wedge \alpha y = \beta y \wedge \alpha (x \bullet z) = \beta (x \bullet z)$. Since \perp_α and \perp_β are both invertible separating relations, this obtains $y \perp_\alpha z \wedge y \perp_\beta z$. Thus, we can distribute α and β over $y \bullet z$ to derive: $\alpha y \bullet \alpha z = \beta y \bullet \beta z$. Since we already have $\alpha y = \beta y$, we apply cancellativity to derive $\alpha z = \beta z$ and conclude the proof. \square

⁸A PCM is cancellative if $a \bullet b = a \bullet c$ implies $b = c$, whenever $a \perp b$ and $a \perp c$.

Sub-PCM. Section 2.7 demonstrates how to use the sub-PCM construction to provide a compact spec. First, we start with a spec like the following:

$$e : \{\lambda s. \alpha_s(s) = b_1 \wedge s_s \perp_\alpha s_o\} \{\lambda s. \alpha_s(s) = b'_1 \wedge s_s \perp_\alpha s_o\}$$

Using the sub-PCM construction, we can write the spec compactly as follows:

$$e : \{\lambda s. (\alpha \circ \iota)_s s = b_1\} \{\lambda s. (\alpha \circ \iota)_s s = b'_1\}$$

Note that, implicitly, we also have that $s_s \perp s_o$. We show the following theorem that states that invertibility is preserved by such construction:

THEOREM 4.6. *Let $\alpha : A \rightarrow B$ be an invertible morphism and let $\iota : A/\perp_\alpha \rightarrow A$ be a sub-PCM injection. Then $\alpha \circ \iota : A/\perp_\alpha \rightarrow B$ is invertible.*

PROOF. Recall the sub-PCM construction in Section 3.4 and use \perp_α as the separating relation for the construction of sub-PCM. Then $\iota x \hat{=} x$ and both ι and $\alpha \circ \iota$ are total morphisms, with the separating relation \perp_{A/\perp_α} . This separating relation is a restriction of \perp_α to the set $A/\perp_\alpha = \{a \in A \mid a \perp_\alpha \mathbb{1}_A\}$. The separating relation is also trivial and thus invertible, by Proposition 4.3.

Now we proceed with the proof of invertibility itself. Assume that we are given $a \in A/\perp_\alpha$ and $b_1, b_2 \in B$ such that $(\alpha \circ \iota) a = \alpha(\iota a) = \alpha a = b_1 \bullet b_2$ and $a \perp_{A/\perp_\alpha} \mathbb{1}$. The second conjunct implies $a \perp_\alpha \mathbb{1}_A$. Now, because α is invertible, there exist $a_1, a_2 \in A$ such that $a = a_1 \bullet_A a_2$, $a_1 \perp_\alpha a_2$, $\alpha a_1 = b_1$, and $\alpha a_2 = b_2$. But, because $a_1 \perp_\alpha a_2$ it follows that $a_1, a_2 \in A/\perp_\alpha$, $a = a_1 \bullet_{A/\perp_\alpha} a_2$, and $a_1 \perp_{A/\perp_\alpha} a_2$. Since also $(\alpha \circ \iota) a_i = \alpha a_i = b_i$, the morphism $\alpha \circ \iota$ is invertible. \square

We can now show that the morphisms and separating relations used in our abstract specs (both the intermediate and final one) of ticket lock are invertible.

LEMMA 4.7. *The separating relation \perp_α from (8) is invertible.*

PROOF. Assume that there are a_1, a_2, a' such that $a_1 \perp_\alpha (a_2 \bullet a')$ and $a_2 \perp_\alpha (a_1 \bullet a')$. Using the definition of \perp_α , we obtain $\#_{\text{serve}}(\sigma a_1) + \#_{\text{serve}}(\sigma(a_2 \bullet a')) \leq 1 \wedge a_1 \perp (a_2 \bullet a')$ and $\#_{\text{serve}}(\sigma a_2) + \#_{\text{serve}}(\sigma(a_1 \bullet a')) \leq 1 \wedge a_2 \perp (a_1 \bullet a')$. This gives us, using the second conjuncts and commutativity and associativity of \bullet as in the proof of Proposition 4.3, that $a_1 \perp a_2 \perp a'$. Further, since $\#_{\text{serve}}(-)$ and σ are morphisms, we obtain, using either of the first conjuncts, $\#_{\text{serve}}(\sigma a_1) + \#_{\text{serve}}(\sigma a_2) + \#_{\text{serve}}(\sigma a') \leq 1$. Therefore $\#_{\text{serve}}(\sigma a_1) + \#_{\text{serve}}(\sigma a') \leq 1$ and we conclude that $a_1 \perp_\alpha a'$. The rest follows from associativity of separating relations. \square

LEMMA 4.8. *The morphism α from (7) is invertible.*

PROOF. The separating relation \perp_α is invertible by Lemma 4.7. Now, assume that there are $a \in U$ and $b_1, b_2 \in O$ such that $\alpha a = b_1 \bullet b_2$ and $a \perp_\alpha \mathbb{1}$. We must show there exist $a_1, a_2 \in U$, such that $a = a_1 \bullet a_2$, $a_1 \perp_\alpha a_2$, $\alpha a_1 = b_1$, and $\alpha a_2 = b_2$. Proceed by case analysis on $b_1 \bullet b_2$.

Case: $b_1 \bullet b_2 = \text{own}$. W.l.o.g. $b_1 = \text{own}$ and $b_2 = \overline{\text{own}}$. Choose $a_1 = a$, $a_2 = \mathbb{1}$. Then trivially $a = a_1 \bullet a_2$ and by assumption $a_1 \perp_\alpha a_2$. Also, $\alpha a_1 = \alpha a = \text{own} = b_1$ and $\alpha a_2 = \alpha \mathbb{1} = \mathbb{1}_O = \overline{\text{own}} = b_2$.

Case: $b_1 \bullet b_2 = \overline{\text{own}}$. Then $b_1 = b_2 = \overline{\text{own}}$. Choose $a_1 = a$, $a_2 = \mathbb{1}$. Again trivially $a = a_1 \bullet a_2$ and by assumption $a_1 \perp_\alpha a_2$. Also, $\alpha a_1 = \alpha a = \overline{\text{own}} = b_1$ and $\alpha a_2 = \alpha \mathbb{1} = \mathbb{1}_O = \overline{\text{own}} = b_2$. \square

COROLLARY 4.9. *The morphism $\alpha' = \alpha \bullet_{\text{TL}}$ from Section 2.7 is invertible.*

The same holds also for the morphism and the separating relation we discussed in Section 3.5:

LEMMA 4.10. *The separating relation \perp_ω from (9) is invertible.*

PROOF. Assuming $a_1 \perp_\omega (a_2 \bullet a')$ and $a_2 \perp_\omega (a_1 \bullet a')$, by associativity of \perp_ω , it suffices to establish $a_1 \perp_\omega a'$. In Definition 4.1, we only consider the clause whereby $a'(t) = L$ implies $\max(\text{dom}(a_1 \bullet a')) \leq t$ or $a'(t+1) = U$. From $a'(t) = L$, it follows that $(a_2 \bullet a')(t) = L$. Therefore, $a_1 \perp_\omega (a_2 \bullet a')$ derives that $\max(\text{dom}(a_1 \bullet (a_2 \bullet a'))) \leq t$ or $(a_2 \bullet a')(t+1) = U$. In the first case, it must also be $\max(\text{dom}(a_1 \bullet a')) \leq \max(\text{dom}(a_1 \bullet (a_2 \bullet a'))) \leq t$, which completes the proof. In the second case, it can be either $a'(t+1) = U$ or $a_2(t+1) = U$. The first case also completes the proof. The second case contradicts the assumption $a_2 \perp_\omega (a_1 \bullet a')$, and is thus impossible. \square

LEMMA 4.11. *The morphism ω from (10) is invertible.*

PROOF. The separating relation \perp_ω is invertible by Lemma 4.10. The rest of the proof follows similarly as in the case of Lemma 4.8. \square

4.4 Invertibility and Separating Conjunction

We close this section with two lemmas that show how invertible morphisms and separating relations interact with separating conjunction. We'll elaborate more on these properties in Section 5 on the related work.

LEMMA 4.12. *Let S be an invertible separating relation, and let $R = \lambda s. (s_s S s_o)$. Then R is duplicable, i.e., $R \Leftrightarrow R * R$.*

PROOF. For the \Rightarrow direction, let's assume that $s = (a, a')$ and $R s$; that is $a S a'$. Consider states $s_1 = s = (a, a')$ and $s_2 = (\mathbb{1}, a \bullet a')$. By definition, $s = s_1 \star s_2$. For s_1 , we do have $s_{1_s} S s_{1_o}$. Indeed, the latter by definition equals $a S a'$, and thus holds by assumption. For s_2 , we do have $s_{2_s} S s_{2_o}$. Indeed, the latter by definition equals $\mathbb{1} S (a \bullet a')$, which holds by Proposition 3.5. But then $R s_1$ and $R s_2$, and thus $(R * R)(s)$.

For the \Leftarrow direction, let's assume $s = s_1 \star s_2$ where $s = (a_1 \bullet a_2, a')$, $s_1 = (a_1, a_2 \bullet a')$ and $s_2 = (a_2, a_1 \bullet a')$, such that $R s_1$ and $R s_2$. That is, for s_1 : $a_1 S (a_2 \bullet a')$. And for s_2 : $a_2 S (a_1 \bullet a')$. By invertibility of S then $a_1 S a_2 S a'$, and thus by associativity $(a_1 \bullet a_2) S a'$, i.e. $R s$. \square

LEMMA 4.13. *Let ϕ be an invertible morphism, and let $F(b) = \lambda s. (\phi_s s = b \wedge s_s \perp_\phi s_o)$. Then $F(b_1 \bullet b_2) \Leftrightarrow F b_1 * F b_2$.*

PROOF. For the \Rightarrow direction, let $s = (a, a')$ and $F(b_1 \bullet b_2)(s)$; that is $\phi a = b_1 \bullet b_2$ and $a \perp_\phi a'$. By defined elements property of \perp_ϕ , it must be $a \perp_\phi \mathbb{1}$. Then by invertibility of ϕ , there exist a_1 and a_2 , such that $a = a_1 \bullet a_2$, $a_1 \perp_\phi a_2$, $\phi a_1 = b_1$ and $\phi a_2 = b_2$. From $(a_1 \bullet a_2) \perp_\phi a'$ and $a_1 \perp_\phi a_2$, by associativity of separating relations, we get $a_1 \perp_\phi (a_2 \bullet a')$ and $a_2 \perp_\phi (a_1 \bullet a')$. Combined with $\phi a_1 = b_1$ and $\phi a_2 = b_2$, we get $F(b_1)(s_1)$ and $F(b_2)(s_2)$, where $s_1 = (a_1, a_2 \bullet a')$ and $s_2 = (a_2, a_1 \bullet a')$. Because also $s = s_1 \star s_2$, we get $(F(b_1) * F(b_2))(s)$.

For the \Leftarrow direction, let $s = s_1 \star s_2$ where $s = (a_1 \bullet a_2, a')$, $s_1 = (a_1, a_2 \bullet a')$ and $s_2 = (a_2, a_1 \bullet a')$, such that $F(b_1)(s_1)$ and $F(b_2)(s_2)$. That is, $\phi a_1 = b_1$ and $a_1 \perp_\phi (a_2 \bullet a')$ and $\phi a_2 = b_2$ and $a_2 \perp_\phi (a_1 \bullet a')$. By invertibility of \perp_ϕ , then $a_1 \perp_\phi a_2 \perp_\phi a'$, and by associativity $(a_1 \bullet a_2) \perp_\phi a'$. By distributivity of ϕ , also $\phi(a_1 \bullet a_2) = \phi a_1 \bullet \phi a_2 = b_1 \bullet b_2$. In other words, $F(b_1 \bullet b_2)(s)$. \square

5 RELATED WORK

PCMs in Separation Logics. PCMs arise as the structure underpinning the semantics of (concurrent) separation logic: the PCMs of heaps capture the dynamics of ownership transfer which is quintessential to separation logics. Initially, cancellative PCMs, also known as *separation algebras* [Calcagno et al. 2007] were used to provide abstract semantic treatment of separation logic. Later, Cao et al. [2017] unified different semantics of separation logics using *ordered* separation

algebras to account for *affine* aspects of various memory models; that is, to model whether deallocation is explicitly allowed to the user, or is carried out implicitly by garbage collection. Several program logics continue this trend, adding further properties to PCMs to give semantics to (higher-order) ghost state [da Rocha Pinto et al. 2014; Dinsdale-Young et al. 2013, 2010; Gotsman et al. 2007; Hobor et al. 2010; Jung et al. 2018, 2015; Krishnaswami et al. 2012; Svendsen and Birkedal 2014; Svendsen et al. 2013; Swamy et al. 2020; Turon et al. 2014]. In this paper we don't consider higher-order state and focus on the algebraic treatment of PCMs without additional properties, as these aren't required by our ambient logic, which admits explicit deallocation. We expect that in the future morphisms and separating relations can be developed for these enriched PCMs.

Recently, several program logics, most notably those that are built on top of the Iris framework [Bizjak et al. 2019; Hinrichsen et al. 2020; Jung et al. 2018, 2020, 2015], the SteelCore framework [Swamy et al. 2020], VST [Appel et al. 2014], and also FCSL [Ley-Wild and Nanevski 2013; Nanevski et al. 2014; Sergey et al. 2015a,b, 2016], have allowed PCMs to be declared at the user level, and sometimes even constructed by means of a predetermined set of combinators.

However, none of these logics have considered morphisms over PCMs, as we do here. Instead, when the state space of a program has to be restricted by some property, that is usually done by conjoining the property to the state space of the underlying state transition system. In contrast, with PCM morphisms, we can restrict the PCM itself, thus promoting the property into a new notion of separateness. The move makes it possible to provide clients with the PCM most suitable to their needs. The new PCM may also be subjected to mathematical theories and their mechanizations that are parametric in the PCM, such as, for example, our theory of invertibility, to facilitate the reuse of mechanized proofs.

Morphisms are a standard component in the study of structures in algebra and category theory. They provide the user with the most general and systematic way to define her own PCM combinators and, as we illustrated, are also useful in specs. Morphisms generally are also essential in the definitions of functors and natural transformations which we plan to consider in the PCM setting in the future. In contrast to our morphism-based specifications, most of the related program logics follow the specification style originating from the work on Concurrent Abstract Predicates (CAP) [Dinsdale-Young et al. 2010], to which we compare below.

We aren't aware of any other work that considers separating relations as a standalone concept. That said, the key separating relation property of associativity (property 5 in Definition 3.4) has been considered before [Jacobs 2018; Krebbers 2015], though as a property of the disjointness relation \perp of the underlying PCM. In our setting, the latter is just one possible separating relation, associated with total PCM morphisms.

Comparison with concurrent abstract predicates (CAP). The intermediate abstract specs for ticket locks we developed in Section 2.6 are similar to the lock specs from CAP [Dinsdale-Young et al. 2010]. We show the CAP specs below, ignoring lock invariants (see Footnote 3 in Section 2), adapted to our type-based notation with explicit binding of the state s in the assertions, and using \wedge instead of $*$.

$$\begin{aligned} \text{lock}_{CAP} &: \{\lambda s. \text{is_lock } s \wedge \text{unlocked } s\} \{\lambda s. \text{is_lock } s \wedge \text{locked } s\} \\ \text{unlock}_{CAP} &: \{\lambda s. \text{is_lock } s \wedge \text{locked } s\} \{\lambda s. \text{is_lock } s \wedge \text{unlocked } s\} \end{aligned}$$

Here is_lock , locked , and unlocked are separation logic assertions (hence, predicates over s). The predicate is_lock captures the internal conditions required of s to represent a lock, and locked and unlocked capture that the lock is taken and free, respectively.⁹ The definitions of the predicates

⁹In [Dinsdale-Young et al. 2010], the unlocked predicate is replaced by separation logic emp , and thus elided. We include it here explicitly to exemplify the similarity with our specs.

are hidden from the clients, but for the specs to be usable wrt. framing and parallel composition, one must export a number of their properties, such as (a) $\text{locked} * \text{locked} \Rightarrow \perp$, and (b) is_lock is duplicable, i.e. $\text{is_lock} \Leftrightarrow \text{is_lock} * \text{is_lock}$.

We could turn our specs of lock and unlock into the same format by setting, for example:

$$\text{is_lock } s \hat{=} s_s \perp_{\alpha} s_o \quad \text{unlocked } s \hat{=} \alpha_s s = \overline{\text{own}} \wedge s_s \perp_{\alpha} s_o \quad \text{locked } s \hat{=} \alpha_s s = \text{own} \wedge s_s \perp_{\alpha} s_o$$

and because \perp_{α} is an invertible separating relation and α an invertible total morphism, by Lemmas 4.12 and 4.13, the equations (a) and (b) above hold. In this sense, we see our algebraic formulation as explaining why the two different kinds of abstract predicates appear in CAP: the duplicable predicates are a lifting of invertible separating relations as in Lemma 4.12, and the non-duplicable ones are a lifting of invertible morphisms as in Lemma 4.13.

Of course, morphisms and separating relations have uses where abstract predicates simply don't apply. Examples are the algebraic constructions that we introduced in Section 3, or the sub-PCM construction which we used to obtain the ultimately simplest abstract specs in Section 2.7. Furthermore, by being functions, morphisms can compute values out of the state,¹⁰ and thus lead to convenient specs and proofs in a formalization based on type theory. We thus propose that abstract specs be given directly in terms of morphisms and separating relations, instead of using their coercion into abstract predicates.

Comparison with the ambient type theory. This paper builds on previous work by Nanevski et al. [2019] which provides a type-theoretic formulation of concurrent separation logic. Nanevski et al. consider an algebraic treatment of *state transition systems of resources*, introduces notions of *resource morphisms* and *simulations*. While that paper focuses on the logic of Hoare triples, in the present paper we focus on the logic of *assertions* and the associated algebraic constructions.

The goal of Nanevski et al. is to provide a systematic way of coercing a program from one resource type to another, as long as the target resource simulates the source one. The system provides an inference rule in the style of Hoare's rule of invariance, to reason about the coerced programs. We utilized this rule implicitly in Section 2.7 to coerce lock from a resource with PCM U to one with PCM U_{TL} . A program is coerced from resource V to resource W by means of a resource morphism, which modifies the behavior of the program on the ghost state. Programmatically, the coercion may be seen as *re-instrumenting* a program with a ghost code specific to W , a posteriori to the proof of the program against the initial ghost instrumentation specific to V , and using the resources as a type-style interface. The same mechanism of resource morphisms provides a scoped way to allocate a new resource into the private state of another resource. Resource morphisms are similar in spirit to the refinement mappings of Abadi and Lamport [1991], and enable a form of refinement-style reasoning within separation logic.

PCM morphisms versus homomorphisms in effect algebras. *Effectus theory* [Cho et al. 2015] is a fairly new field of category theory whose aim is to describe quantum computation and its logic, hence generalizing probabilistic and Boolean logic. The mathematical backbone of effectus theory is effect algebras, which essentially are PCMs with an orthosupplement, i.e. a total unary negation operation. An effectus is a category with finite coproducts and final object that satisfies three technical properties: 1) a form of partial pairing for compatible partial maps; 2) disjointness of coprojections; and 3) joint monicity of partial projections. In effectuses, predicates are total maps of the form $X \rightarrow Y + 1$ which, as usual, are equivalent to partial maps of the form $X \rightarrow Y$. In particular, given an effectus B , the category $\text{Par}(B)$ of partial maps over B is enriched over the category of PCMs.

¹⁰For example, how we used $\hat{\psi}$ in the concrete specs for ticket lock to compute the displayed ticket.

Interestingly, the notion of homomorphism for effect algebras [Cho 2015; Cho et al. 2015; Jacobs 2018] is similar to our notions of PCM morphisms. Indeed, the similarities between our Definition 3.4 and Definition 3.6 with [Cho et al. 2015, Definition 12] are clear. One difference, however, is that their definition only considers PCM morphisms with trivial separating relation (what we call total morphisms), whereas our morphisms can have more general separating relations, and are thus properly partial. The origins of PCM morphisms, as described in our paper, lie in separation logic and we have explored their applications to verification of concurrent programs. Nevertheless the close relation to effectuses encourages us to explore future applications of our work to recent extensions of separation and Hoare logic such as quantum relational Hoare logic [Unruh 2019a,b], relational proofs of quantum programs [Barthe et al. 2020b], and probabilistic separation (and other program) logics [Barthe et al. 2020a; Batz et al. 2019; Sato et al. 2019; Tassarotti and Harper 2019].

6 CONCLUSION AND FUTURE WORK

Morphisms are a standard notion in algebra and category theory, where algebraic structures give rise to structure-preserving functions, i.e. morphisms, between them. We adapt the notion of morphisms to the structure of PCMs, thereby extending standard algebraic and categorical approaches to concurrent separation logics.

The mathematics behind this adaptation gives rise to separating relations, which delineate the domain where a function is structure preserving and thus a morphism. We introduce invertibility as a property of morphisms and separating relations that allows working with morphisms under abstraction. Our exposition of PCMs and their morphisms is natural; we recover the standard algebraic constructions (e.g. that of a sub-object, a sub-PCM), show that the constructions preserve structure (e.g., composition of morphisms is a morphism, equalizer of morphisms is a separating relation, etc.), and show that invertibility is preserved under composition and products of morphisms. Morphisms are useful in specs to compute values out of the state; structure preservation ensures that morphisms are well behaved under ownership transfer.

In the future, we will build on the scaffolding provided by PCM morphisms, along with resource morphisms and simulations [Nanevski et al. 2019], to obtain an algebraic theory of *linearizable* resources. Such a formalism will unite logical, categorical, and type-theoretic foundations [Harper 2011], while supporting the verification of a wide range of realistic concurrent programs.

ACKNOWLEDGMENTS

We thank Gordon Stewart and Joe Tassarotti for their comments on various drafts of the paper. We thank the anonymous reviewers from the POPL'21 PC and AEC for their feedback. This research was partially supported by the Spanish MICINN projects BOSCO (PGC2018-102210-B-I00) and ProCode-UCM (PID2019-108528RB-C22), the European Research Council project Mathador (ERC2016-COG-724464) and the US National Science Foundation (NSF). Any opinions, findings, and conclusions or recommendations expressed in the material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- Martin Abadi and Leslie Lamport. 1991. The existence of refinement mappings. *Theoretical Computer Science (TCS)* 82, 2 (1991), 253–284. [https://doi.org/10.1016/0304-3975\(91\)90224-P](https://doi.org/10.1016/0304-3975(91)90224-P)
- Samson Abramsky and Achim Jung. 1995. *Domain Theory*. Oxford University Press, Inc., USA, 1–168.
- Andrew W. Appel, Robert Dockins, Aquinas Hobor, Lennart Beringer, Josiah Dodds, Gordon Stewart, Sandrine Blazy, and Xavier Leroy. 2014. *Program Logics for Certified Compilers*. Cambridge University Press. <https://doi.org/10.1017/CBO9781107256552>

- Gilles Barthe, Justin Hsu, and Kevin Liao. 2020a. A probabilistic separation logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 55:1–55:30. <https://doi.org/10.1145/3371123>
- Gilles Barthe, Justin Hsu, Mingsheng Ying, Nengkun Yu, and Li Zhou. 2020b. Relational proofs for quantum programs. *Proc. ACM Program. Lang.* 4, POPL (2020), 21:1–21:29. <https://doi.org/10.1145/3371089>
- Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Thomas Noll. 2019. Quantitative separation logic: a logic for reasoning about probabilistic pointer programs. *Proc. ACM Program. Lang.* 3, POPL (2019), 34:1–34:29. <https://doi.org/10.1145/3290347>
- Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2005. Symbolic Execution with Separation Logic. In *Programming Languages and Systems, Third Asian Symposium, APLAS 2005, Tsukuba, Japan, November 2-5, 2005, Proceedings*. 52–68. https://doi.org/10.1007/11575467_5
- Aleš Bizjak, Daniel Gratzer, Robbert Krebbers, and Lars Birkedal. 2019. Iron: managing obligations in higher-order concurrent separation logic. *Proc. ACM Program. Lang.* 3, POPL, 65:1–65:30. <https://doi.org/10.1145/3290378>
- Richard Bornat, Cristiano Calcagno, Peter W. O’Hearn, and Matthew J. Parkinson. 2005. Permission accounting in separation logic. In *ACM Symposium on Principles of Programming Languages (POPL)*. 259–270. <https://doi.org/10.1145/1040305.1040327>
- Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. 2007. Local Action and Abstract Separation Logic. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007, 10-12 July 2007, Wrocław, Poland, Proceedings*. IEEE Computer Society, 366–378. <https://doi.org/10.1109/LICS.2007.30>
- Qinxiang Cao, Santiago Cuellar, and Andrew W. Appel. 2017. Bringing Order to the Separation Logic Jungle. In *Programming Languages and Systems - 15th Asian Symposium, APLAS 2017, Suzhou, China, November 27-29, 2017, Proceedings*. 190–211. https://doi.org/10.1007/978-3-319-71237-6_10
- Kenta Cho. 2015. Total and Partial Computation in Categorical Quantum Foundations. In *Proceedings 12th International Workshop on Quantum Physics and Logic, QPL 2015, Oxford, UK, July 15-17, 2015 (EPTCS, Vol. 195)*. 116–135. <https://doi.org/10.4204/EPTCS.195.9>
- Kenta Cho, Bart Jacobs, Bas Westerbaan, and Abraham Westerbaan. 2015. An Introduction to Effectus Theory. *CoRR* abs/1512.05813 (2015). arXiv:1512.05813 <http://arxiv.org/abs/1512.05813>
- P. J. Courtois, F. Heymans, and D. L. Parnas. 1971. Concurrent control with “readers” and “writers”. *Commun. ACM* 14, 10 (1971), 667–668. <https://doi.org/10.1145/362759.362813>
- Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A logic for time and data abstraction. In *European Conference on Object-Oriented Programming (ECOOP)*. 207–231. https://doi.org/10.1007/978-3-662-44202-9_9
- Germán Andrés Delbianco, Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2017. Concurrent data structures linked in time. In *European Conference on Object-Oriented Programming (ECOOP)*. 8:1–8:30. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.8>
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: compositional reasoning for concurrent programs. In *ACM Symposium on Principles of Programming Languages (POPL)*. 287–300. <https://doi.org/10.1145/2429069.2429104>
- Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. 2010. Concurrent abstract predicates. In *European Conference on Object-Oriented Programming (ECOOP)*. 504–528. https://doi.org/10.1007/978-3-642-14107-2_24
- František Farka, Aleksandar Nanevski, Anindya Banerjee, Germán Andrés Delbianco, and Ignacio Fábregas. 2020a. On Algebraic Abstractions for Concurrent Separation Logics (artefact). (October 2020). <https://doi.org/10.5281/zenodo.4069513>
- František Farka, Aleksandar Nanevski, Anindya Banerjee, Germán Andrés Delbianco, and Ignacio Fábregas. 2020b. On Algebraic Abstractions for Concurrent Separation Logics (extended version). *CoRR* (October 2020). <https://arxiv.org/abs/2010.12686>
- Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. 2007. Local Reasoning for Storable Locks and Threads. In *Programming Languages and Systems, 5th Asian Symposium, APLAS 2007, Singapore, November 29-December 1, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4807)*, Zhong Shao (Ed.). Springer, 19–37. https://doi.org/10.1007/978-3-540-76637-7_3
- Robert Harper. 2011. The Holy Trinity. <https://existentialtype.wordpress.com/2011/03/27/the-holy-trinity/>.
- Maurice Herlihy and Nir Shavit. 2008. *The art of multiprocessor programming*. M. Kaufmann. <https://doi.org/10.1108/03684920810907904>
- Maurice Herlihy and Jeannette M. Wing. 1990. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 12, 3 (1990), 463–492. <https://doi.org/10.1145/78969.78972>
- Jonas Kastberg Hinrichsen, Jesper Bengtson, and Robbert Krebbers. 2020. Actris: session-type based reasoning in separation logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 6:1–6:30. <https://doi.org/10.1145/3371074>

- Aquinas Hobor, Robert Dockins, and Andrew W. Appel. 2010. A theory of indirection via approximation. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*. 171–184. <https://doi.org/10.1145/1706299.1706322>
- Bart Jacobs. 2018. From probability monads to commutative effectuses. *J. Log. Algebr. Meth. Program.* 94 (2018), 200–237. <https://doi.org/10.1016/j.jlamp.2016.11.006>
- Jonas Braband Jensen and Lars Birkedal. 2012. Fictional Separation Logic. In *ESOP*.
- Cliff B. Jones. 1983. Tentative Steps Toward a Development Method for Interfering Programs. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 5, 4 (1983). <https://doi.org/10.1145/69575.69577>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming (JFP)* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. 2020. The future is ours: prophecy variables in separation logic. *Proc. ACM Program. Lang.* 4, POPL (2020), 45:1–45:32. <https://doi.org/10.1145/3371113>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: monoids and invariants as an orthogonal basis for concurrent reasoning. In *ACM Symposium on Principles of Programming Languages (POPL)*. 637–650. <https://doi.org/10.1145/2676726.2676980>
- Robbert Krebbers. 2015. *The C standard formalized in Coq*. Phd Thesis. Radboud University Nijmegen. <https://hdl.handle.net/2066/147182>
- Neelakantan R. Krishnaswami, Aaron Turon, Derek Dreyer, and Deepak Garg. 2012. Superficially substructural types. In *ICFP*.
- Leslie Lamport. 1974. A New Solution of Dijkstra’s Concurrent Programming Problem. *Commun. ACM* 17, 8 (1974), 453–455. <https://doi.org/10.1145/361082.361093>
- Ruy Ley-Wild and Aleksandar Nanevski. 2013. Subjective auxiliary state for coarse-grained concurrency. In *ACM Symposium on Principles of Programming Languages (POPL)*. 561–574. <https://doi.org/10.1145/2429069.2429134>
- John M. Mellor-Crummey and Michael L. Scott. 1991. Algorithms for Scalable Synchronization on Shared-memory Multiprocessors. *ACM Transactions on Computer Systems (TOCS)* 9, 1 (1991), 21–65. <https://doi.org/10.1145/103727.103729>
- Aleksandar Nanevski, Anindya Banerjee, Germán Andrés Delbianco, and Ignacio Fábregas. 2019. Specifying concurrent programs in separation logic: morphisms and simulations. *PACMPL* 3, OOPSLA (2019), 161:1–161:30. <https://doi.org/10.1145/3360587>
- Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating state transition systems for fine-grained concurrent resources. In *European Symposium on Programming (ESOP)*. 290–310. https://doi.org/10.1007/978-3-642-54833-8_16
- Aleksandar Nanevski, Greg Morrisett, and Lars Birkedal. 2006. Polymorphism and separation in Hoare Type Theory. In *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming (ICFP 2006)*. <https://doi.org/10.1145/1159803.1159812>
- David J. Pym, Peter W. O’Hearn, and Hongseok Yang. 2004. Possible worlds and resources: the semantics of BI. *Theor. Comput. Sci.* 315, 1 (2004), 257–305. <https://doi.org/10.1016/j.tcs.2003.11.020>
- Tetsuya Sato, Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Justin Hsu. 2019. Formal verification of higher-order probabilistic programs: reasoning about approximation, convergence, Bayesian inference, and optimization. *Proc. ACM Program. Lang.* 3, POPL (2019), 38:1–38:30. <https://doi.org/10.1145/3290351>
- Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2015a. Mechanized Verification of Fine-grained Concurrent Programs. In *ACM Conference on Programming Languages Design and Implementation (PLDI)*. <https://doi.org/10.1145/2737924.2737964>
- Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2015b. Specifying and verifying concurrent algorithms with histories and subjectivity. In *European Symposium on Programming (ESOP)*. 333–358. https://doi.org/10.1007/978-3-662-46669-8_14
- Ilya Sergey, Aleksandar Nanevski, Anindya Banerjee, and Germán Andrés Delbianco. 2016. Hoare-style specifications as correctness conditions for non-linearizable concurrent objects. In *ACM Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. 92–110. <https://doi.org/10.1145/3022671.2983999>
- Kasper Svendsen and Lars Birkedal. 2014. Impredicative concurrent abstract predicates. In *European Symposium on Programming (ESOP)*. 149–168. https://doi.org/10.1007/978-3-642-54833-8_9
- Kasper Svendsen, Lars Birkedal, and Matthew J. Parkinson. 2013. Modular reasoning about separation of concurrent data structures. In *European Symposium on Programming (ESOP)*. 169–188. https://doi.org/10.1007/978-3-642-37036-6_11
- Nikhil Swamy, Aseem Rastogi, Aymeric Fromherz, Denis Merigoux, Danel Ahman, and Guido Martínez. 2020. SteelCore: an extensible concurrent separation logic for effectful dependently typed programs. *Proc. ACM Program. Lang.* 4, ICFP (2020), 121:1–121:30. <https://doi.org/10.1145/3409003>

- Joseph Tassarotti and Robert Harper. 2019. A separation logic for concurrent randomized programs. *Proc. ACM Program. Lang.* 3, POPL (2019), 64:1–64:30. <https://doi.org/10.1145/3290377>
- Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. 2014. GPS: navigating weak memory with ghosts, protocols, and separation. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014*, Andrew P. Black and Todd D. Millstein (Eds.). ACM, 691–707. <https://doi.org/10.1145/2660193.2660243>
- Dominique Unruh. 2019a. Quantum Hoare Logic with Ghost Variables. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*. IEEE, 1–13. <https://doi.org/10.1109/LICS.2019.8785779>
- Dominique Unruh. 2019b. Quantum relational Hoare logic. *Proc. ACM Program. Lang.* 3, POPL (2019), 33:1–33:31. <https://doi.org/10.1145/3290346>