arXiv:2411.12857v1 [cs.PL] 19 Nov 2024

An Abstract Domain for Heap Commutativity

Jared Pincus[1][0000−0001−6708−5262] and Eric Koskinen[2][0000−0001−7363−634X]

[1] Boston University, Boston MA 02215, USA pincus@bu.edu
[2] Stevens Institute of Technology, Hoboken NJ 07030, USA
eric.koskinen@stevens.edu

Abstract. Commutativity of program code (the equivalence of two code fragments composed in alternate orders) is of ongoing interest in many settings such as program verification, scalable concurrency, and security analysis. While some recent works have explored static analysis for code commutativity, few have specifically catered to heap-manipulating programs. We introduce an abstract domain in which commutativity synthesis or verification techniques can safely be performed on abstract mathematical models and, from those results, one can directly obtain commutativity conditions for concrete heap programs. This approach offloads challenges of concrete heap reasoning into the simpler abstract space. We show this reasoning supports framing and composition, and conclude with commutativity analysis of programs operating on example heap data structures. Our work has been mechanized in Coq and is available in the supplement.

Keywords: Commutativity · Abstract Interpretation · Observational Equivalence · Separation Logic

1 Introduction

Commutativity describes circumstances in which two programs executed in sequence yield the same result regardless of their execution order. Fundamentally, commutativity offers a kind of independence that enables parallelization or simplifies program analysis by collapsing equivalent cases. As such, this property has long been of interest in many areas including databases [34], parallelizing compilers [31], scalable [8] and distributed systems [32], and smart contracts [12,28]. Commutativity is also used in verification, including efforts to simplify proofs of parameterized programs [18], and concurrent program properties [17] such as termination [24], dynamic race detection [19], and information flow [13].

Efforts in *commutativity analysis* seek to infer or verify *commutativity conditions*—the initial conditions (if any) under which code fragments commute. For example, Rinard and Kim [21] verify commutativity conditions by relating a data structure to its pre/post-conditions, then showing commutativity of that abstraction. Later, methods were developed to automatically synthesize commutativity conditions for specifications and for imperative program fragments [2,3,7,6], which can be used to automatically parallelize programs [31,29,6].

Objectives. Considering its many applications, it would be appealing to extend commutativity analysis to the heap memory model. However, few works thus far have targeted heap-modifying programs. Pîrlea *et al.* [28] describe a commutativity analysis with some heap treatment, but with the limited scope of heap disjointness/ownership. Eilers *et al.* [13] employ, but do not verify, commutativity to reason about information flow security in concurrent separation logic. This paper thus aims to analyze commutativity of heap programs.

Challenges & Contributions. Commutativity reasoning can be challenging for heap-manipulating programs, but is more straightforward for functions over simple mathematical objects. With that in mind, this paper presents an abstract domain which is amenable to existing commutativity condition verification and synthesis techniques [2,3,7], and allows conditions found in said abstract space to immediately yield concrete separation-logic-style conditions for heap programs. For instance, given a simple mathematical list representation of a stack data structure, we can easily find that the functions *push* and *pop* operating on these lists commute when the first (i.e. top) list element equals the value being pushed. We may then thread down and apply this condition to concrete `push` and `pop` methods. For this threading to be sound, we design the domain to handle various complications of heap programs, including observational equivalence of heap data structures, nondeterminism of address allocation, and improperly allocated heaps and resultant program failure.

In summary, the contributions of this paper include:

(Sec. 4) We introduce a simple abstract domain that is amenable to abstract commutativity reasoning, which can soundly entail concrete commutativity. The domain is designed to lift heap-allocated structures to mathematical objects (e.g. value lists to encode a stack) via separation-logic-style predicates, while cleanly handling ambiguous or improperly allocated heaps. We then describe the semantics of abstract programs, and establish a soundness relationship between abstract and concrete programs in the style of abstract interpretation [10], which enables the reduction of concrete to abstract commutativity.

(Sec. 5) We define concrete and abstract commutativity, and present our main soundness theorem: that a valid commutativity condition for two concrete programs may be derived immediately from commutativity reasoning performed (much more easily) on their corresponding abstract programs. These derived concrete conditions are also *framed*, i.e. if the programs commute under a given initial heap layout, they also commute in any larger layout.

(Sec. 6) We define a way to compose abstractions, which mirrors the separating conjunction on concrete heaps, and present compositional commutativity properties. Composition aids in settings with multiple allocated data structures, and working with arguments/return values.

(Sec. 7) We share four complete examples: a non-negative counter, two-place unordered set, linked-list stack, and the composition of a counter and stack.

(Supplement) Our work is implemented in Coq, on top of Separation Logic Foundations [5]. All theorems, lemmas and properties have been proved in Coq (except for the example programs written in the concrete semantics of §3.3).

Limitations. In this paper we establish the theory necessary for deriving heap commutativity from abstract commutativity reasoning. We leave implementation of this theory, including automation and synthesis, future work.

2 Overview

We summarize our work through the example of a stack, implemented as a heap-allocated linked list, with typical `push` and `pop` methods (`pop` fails if the stack is empty). For simplicity, method arguments or return values are stored in method-specific pre-allocated heap cells (e.g. a).

```
push^a_ℓ ≔ let p = !ℓ in
  let v = !a in
  let x = (v, p) in
  let q = ref x
  in ℓ ← q

pop^a_ℓ ≔ let p = !ℓ in
  if p = null then fail else
  let x = !p in
  let v = fst x in let q = snd x
  in ℓ ← q ; free p ; a ← v
```

Inconvenience of concrete heap commutativity. We would like to know under what precondition P, if any, do `push` and `pop` *commute*. Verifying whether `assume`(P);`push`;`pop` is equivalent to `assume`(P);`pop`;`push` is not amenable to typical relational reasoning techniques [35], because the swapped orders of `push` and `pop` (or, in general, f and g) do not yield meaningful alignment points [25,1]. So, the problem amounts to performing four forward symbolic executions per method pair, meaning quadratically many executions as the number of methods grows. There are, however, efficient CEGAR/CEGIS methods for verifying and even synthesizing a condition P of logical method pre/postconditions [2,3,7]. We now discuss a simple abstraction of the Stack methods as mathematical functions, how our theory relates the concrete heap to this abstraction, and how this enables *concrete* commutativity conditions to be immediately derived from the results of simpler (automatable) abstract commutativity reasoning.

Abstract commutativity. Let us, for the moment, set aside the concrete heap implementation, and consider a simple mathematical representation of a stack as a list of values, along with functions over the stack with an input/output "cell":

$$push(s, v) \coloneqq (v :: s, v) \qquad pop([], \_) \coloneqq \bot \qquad pop(v :: s, \_) \coloneqq (s, v)$$

Here, *push* appends the input value to the stack, and *pop* removes and returns the topmost value, but fails (yields $\bot$) on an empty stack. Reasoning about commutativity of these mathematical objects is convenient, and synthesizing commutativity conditions can now be automated [2,3,7]. For this example, given an initial stack s, $push(u)$ and $push(v)$ commute when $u = v$; $push(u)$ and $pop()$ when $s = u :: \_$; and $pop()$ and $pop()$ when $s = x :: x :: \_$ for some x.

The question we ask in this paper, is how we might abstract the concrete programs so that, from these abstract commutativity conditions, we can immediately obtain concrete conditions. We show that this is indeed possible; for example, in the case of $\texttt{push}^a_\ell$ and $\texttt{push}^b_\ell$ we obtain the (separation logic) precondition: $\exists u.\ \mathsf{stk}^+(\_) * a \mapsto u * b \mapsto u$, where $\mathsf{stk}^+(\_)$ asserts that the heap contains

a properly allocated stack (we formally define stk^+ below). We now summarize how to achieve this using the Stack example.

Step 1: Projecting heaps to abstract Stacks. To soundly lift concrete heap structures into an abstract domain, we must accommodate various complications of the heap, such as improperly allocated heaps, observationally equivalent layouts, and allocation nondeterminism. We achieve this by choosing a set of abstract values X (such as value lists to represent stacks), and a projection function π which maps concrete heaps into X. Whether a heap is "well-structured" depends on the choice of π. For the Stack, π encodes that the only valid heaps for the abstraction are those containing a correctly allocated linked stack. We say that such heaps are *within the purview* of the abstraction at hand, meaning they will be amenable to abstract commutativity reasoning.

Of course, π must also account for heaps which are *not* properly allocated to be representable by a value in X. To this end, every abstract domain includes a special element ✗ ("cross"), to which π maps all inappropriately structured, or "out-of-purview", heaps. Commutativity w.r.t. such heaps cannot be reasoned about meaningfully in the abstract space.

For our Stack example, the abstract values X are lists of non-null concrete values, and the projection function π maps a heap to such a list by checking (via separation logic heap predicates [26]) whether it contains a valid linked list, and extracting the contents of said list. Precisely:

$$
\begin{array}{ll}
\pi(h) & \coloneqq s \text{ if } \mathsf{stk}^+(s,h) \text{ for } s \in \mathrm{list}(\mathrm{value}) \text{ else } ✗, \text{ where} \\
\\
\mathsf{stk}(s) & \coloneqq \exists p \in (\mathrm{address} + \mathsf{null}).\ \ell \mapsto p * \psi(s,p) \\
\psi(v :: s, p) & \coloneqq \exists q \in (\mathrm{address} + \mathsf{null}).\ p \neq \mathsf{null} * p \mapsto (v,q) * \psi(s,q) \\
\psi([], p) & \coloneqq p = \mathsf{null}
\end{array}
$$

By existentially quantifying the intermediate addresses of the linked stack and omitting them from the extracted abstract value, we render stacks observationally equivalent if and only if their contained values match.

Step 2: Abstract domain for commutativity. The next question is how this projection should induce an abstract domain, while soundly respecting concrete programs, and enabling commutativity reasoning. The key is to "wrap" the abstract values X into a domain that also includes elements which track if abstract commutativity still entails concrete commutativity. Namely, we add the abstract value ✗ for out-of-purview heaps, ✓ ("check") for an ambiguous collection of in-purview heaps, $\bot$ as a failure state, and $\top$ for a collection of heaps both in- and out-of-purview. Instrumenting π appropriately induces a Galois connection [11] between this abstract domain and the concrete heap domain.

For abstract commutativity to soundly entail concrete commutativity, we establish a soundness relation between concrete and abstract programs, in the style of abstract interpretation [10]. It demands that $push^a_\ell$ overapproximate `push`, and that $push^a_\ell$ fail when `push` fails on an in-purview heap.

Summary. From our user-provided Stack abstraction, and typical abstract interpretation techniques, we can immediately derive sound heap commutativity conditions from simple mathematical conditions:

Concrete Programs	Stack Functions	Abstract Condition	Derived Concrete Condition
$\texttt{push}^a_\ell\ \texttt{push}^b_\ell$	$push(u), push(v)$	$u = v$	$\exists u.\ \mathsf{stk}^+(\_) * a \mapsto u * b \mapsto u$
$\texttt{push}^a_\ell\ \texttt{pop}^b_\ell$	$push(u), pop()$	$s = u :: \_$	$\exists u.\ \mathsf{stk}^+(u :: \_) * a \mapsto u * b \mapsto \_$
$\texttt{pop}^a_\ell\ \texttt{pop}^b_\ell$	$pop(), pop()$	$s = x :: x :: \_$	$\mathsf{stk}^+(\_ :: \_ :: \_) * a \mapsto \_ * b \mapsto \_$

In the remainder of this paper we formalize the above discussion, share properties of commutativity, and define composition on abstract domains. Then in §7 we elaborate on the Stack example (adding a `peek` operation) along with other examples. All definitions, lemmas, theorems, and examples have been mechanized in Coq (available in the supplement), except for those concerning the example concrete semantics in §3.3.

3 Preliminaries

Here we outline conventions and notations for heaps and heap predicates, as well as our choice of concrete heap semantics.

3.1 Heaps and Heap Predicates

A *heap* is a finite partial map from addresses ($\mathbb{L}$) to concrete values ($\mathbb{V}$). Denote the set of all heaps with $\mathbb{H}$, the set of non-null values as $\mathbb{V}^* = \mathbb{V} \setminus \{\mathsf{null}\}$, and the finite address domain of heap h with $\mathsf{dom}(h)$.

Concrete values may be any typical datatypes (integers, booleans, pairs, etc.) which are generally fixed in size (e.g. an unbounded list would not typically be stored as one concrete value). Addresses themselves can be treated as concrete values. Assume unlimited addresses are available, namely $\forall h.\, \mathsf{dom}(h) \subsetneq \mathbb{L}$.

Heaps h_1 and h_2 are *disjoint* ($h_1 \perp h_2$) iff $\mathsf{dom}(h_1) \perp \mathsf{dom}(h_2)$. The *union* of disjoint heaps ($h_1 + h_2$) yields the combined partial map of h_1 and h_2. Dually, a heap may be partitioned into smaller disjoint heaps. h_1 is a *subheap* of h_2 ($h_1 \subseteq h_2$) iff $h_1 + k = h_2$ for some $k \perp h_1$. Heap equality is defined extensionally.

A *heap predicate* characterizes the addresses and values of heaps. Given heap predicate Ψ on $\mathbb{H}$, $\Psi(h)$ denotes that h satisfies Ψ, and $\mathbb{H}(\Psi)$ denotes the set of all Ψ-satisfying heaps. Sometimes we parameterize a heap predicate on some domain V. In this case, $\Psi(v, h)$ denotes that h satisfies Ψ with $v \in V$.

Aside from heap predicates containing pure logical expressions and quantifiers, we will describe heap contents directly with two standard predicates from separation logic [30,26]. The *points-to* predicate $p \mapsto v$, for $p \in \mathbb{L}$ and $v \in \mathbb{V}$, holds for h iff $\mathsf{dom}(h) = \{p\}$ and $h(p) = v$. The *separating conjunction* operator $\Psi * \Phi$, for predicates Ψ and Φ, holds for h iff $h = h_1 + h_2$ for some h_1 and h_2 s.t. $h_1 \perp h_2$ and $\Psi(h_1)$ and $\Phi(h_2)$. $*$ is commutative and associative.

We will also want to characterize when a *portion* of a heap satisfies a particular predicate. For this, we define heap predicate *extension*:

Definition 1. *For predicate Ψ, h satisfies* extension Ψ^+ *iff a subheap of h satisfies Ψ. Namely, $\Psi^+(h) \equiv (\exists \Phi.\, \Psi * \Phi)(h)$ and $\Psi^+(v, h) \equiv (\exists \Phi.\, \Psi(v, \cdot) * \Phi)(h)$.*

3.2 Concrete Semantics

We choose to represent concrete programs as total functions over heaps, so that our semantics are highly general, and thus broadly applicable; we admit semantics which are nondeterministic, but must be constructive and must terminate. Nondeterminism mainly arises during address allocation, though this work is agnostic to the source of nondeterminism.

Definition 2. *A* (concrete) program *is a total function* $f : \mathbb{H} \to \wp(\mathbb{H})$*, subject to the following properties:*

1. *On any initial heap* h*,* f *terminates or fails in finite time.*
2. *If* f *successfully terminates on* h*, then* $f(h)$ *yields the set of all possible final heaps in which* f *may terminate. If* f *fails on* h*, then* $f(h) = \varnothing$*.*
3. f *reads any input arguments from, and writes any outputs to, fixed locations on the heap. We explore this in §7.*
4. f acts locally*. Namely, if* $f(h) \neq \varnothing$*, then for any* $h' \perp h$*:* $f(h + h') \neq \varnothing$ *and* $\forall k \in f(h + h').\ h' \subseteq k$*.*

Local action enforces that f does not modify data outside its footprint, as a weak form of the typical frame rule of separation logic [30]. For any particular semantics that satisfies framing by construction, the local action constraint should immediately hold for all concrete programs. When working with compound abstractions, we will define a stronger framing property (Def. 14).

For two programs to commute, they must terminate without failure in the first place. To that end, we classify the initial heaps upon which f can execute. *Sufficient* heaps contain the necessary addresses, which map to appropriate values, for f to run. In contrast, heaps in the *footprint* of f contain appropriate addresses, but may not contain appropriate values; thus a heap in f's footprint on which f fails cannot be "corrected" with further allocations. Precisely:

$$\mathsf{suff}(f) \triangleq \{h \mid f(h) \neq \varnothing\}$$
$$\mathsf{foot}(f) \triangleq \mathsf{suff}(f) \cup \{h \mid \forall h' \perp h.\, h + h' \notin \mathsf{suff}(f)\}$$

Our *concrete domain*, as the counterpart to our eventual abstract domain, will be $\mathscr{C} \triangleq \wp(\mathbb{H})$. *Concrete transformers* over this domain map from sets of heaps to sets of heaps. We derive these transformers directly from programs:

Definition 3. *The* concrete transformer *constructed from* f*, denoted* $\bar{f}$*, is defined as* $\bar{f}(C) = \bigcup_{h \in C} f(h)$ *if* $C \subseteq \mathsf{suff}(f)$*, else* $\varnothing$*.*

$\bar{f}$ yields the set of all possible output heaps given a set of inputs, but fails entirely if f fails on any individual input. Relatedly, we define the composition/sequence of two programs so that the second program fails entirely if it fails on any one output of the first program; namely, $f; g \triangleq \bar{g} \circ f$.

3.3 An Example Concrete Language

We share here a simple language called Himp, which featuring allocation (`ref`), deallocation, writing (←), reading (`!`), and branching. We will use Himp in our example programs throughout the paper, as an instance of the general semantics above. (N.B. Himp has not been formalized in Coq. However, all of our results regarding the general semantics have been proved in Coq.)

$$\begin{aligned} \mathit{stmt} &::= \mathit{id} \leftarrow \mathit{val} \mid \texttt{free}\ \mathit{id} \mid \mathit{stmt}\ \texttt{;}\ \mathit{stmt} \mid \texttt{fail} \mid \texttt{skip} \\ &\quad \mid \texttt{let}\ \mathit{id} \coloneqq \mathit{exp}\ \texttt{in}\ \mathit{stmt} \mid \texttt{if}\ \mathit{val}\ \texttt{then}\ \mathit{stmt}\ \texttt{else}\ \mathit{stmt} \\ \mathit{exp} &::= \texttt{ref}\ \mathit{val} \mid \texttt{!}(\mathit{id} \mid \mathbb{L}) \mid \mathit{val} + \mathit{val} \mid \mathit{val} = \mathit{val} \mid \cdots \\ \mathit{val} &::= \mathit{id} \mid \mathbb{V} \end{aligned}$$

Himp has a nondeterministic allocation procedure: for any initial heap h, `ref` yields an address from a set $\mathsf{fresh}(h)$ where $\mathsf{fresh}(h) \subset \mathbb{L}$ and $\mathsf{fresh}(h) \perp \mathsf{dom}(h)$. All other terms are deterministic. We elide the remaining details of Himp's semantics, which are typical for an imperative heap language. When working with Himp programs throughout the paper, we implicitly lift them to the concrete programs described by Def. 2.

4 Abstract Domain

In this section we formalize our abstract domain and abstract programs operating within this domain, and establish a soundness relationship between concrete and abstract programs. This will let us perform commutativity reasoning easily with abstract objects, and obtain concrete commutativity results for free.

4.1 Constructing the Domain

We design our abstract domain to lift heap-allocated structures to mathematical objects. Each instance of an abstraction consists of a set of *abstract values* (e.g. integers, sets, sequences) and a mapping from heaps into said values. A heap with an appropriately allocated structure to be mapped to an abstract value is said to be *"within the purview"* of a given abstraction. On the other hand, poorly structured heaps are *"outside the purview"* of the abstraction, and get mapped to a special abstract value, ✗ ("cross").

When constructing an abstraction, we impose a *locality constraint* which demands that information extracted from heaps be finite in scope, so that any heap *containing* a valid structure is in-purview. Some results of this work actually do not depend on this constraint. However, it is a valuable property to strive for, and indeed it should hold for any typical data structure abstraction.

The abstract domain is defined formally as follows:

Definition 4. *Build an abstraction* $\mathsf{A} = \langle X, \pi \rangle$ *from a set of* abstract values $X = \{x_1, x_2, \dots\}$*, and a* projection function $\pi : \mathbb{H} \to X + ✗$ *subject to the* locality

constraint *that for any disjoint h and h', $\pi(h) \in X \Rightarrow \pi(h+h') = \pi(h)$. The full domain is $\mathscr{A}_\mathsf{A} \triangleq X+\bot+\top+\text{✗}+\text{✓}$, with partial ordering $\leq_\mathsf{A}$, where $\bot \leq_\mathsf{A} \text{✗} \leq_\mathsf{A} \top$ and $\forall x \in X.\ \bot \leq_\mathsf{A} x \leq_\mathsf{A} \text{✓} \leq_\mathsf{A} \top$ and $\forall x_1, x_2 \in X.\ x_1 \leq_\mathsf{A} x_2 \Leftrightarrow x_1 = x_2$.*

Denote all in-purview heaps of A with $\mathsf{Prv}(\mathsf{A}) \triangleq \{h \mid \pi(h) \in X\}$, and denote its abstract values and projection function with X_A and π_A.

$\mathscr{A}_\mathsf{A}$ is a complete lattice, and is thus equipped with a typical join ($\sqcup$) operation. Recalling from §3.2 that our concrete domain is $\mathscr{C} = \wp(\mathbb{H})$, we connect the concrete and abstract domains as follows:

Definition 5. *Define* abstraction *and* concretization *functions $\alpha_\mathsf{A} : \mathscr{C} \to \mathscr{A}_\mathsf{A}$ and $\gamma_\mathsf{A} : \mathscr{A}_\mathsf{A} \to \mathscr{C}$, where $\alpha_\mathsf{A}(C) = \bigsqcup_{h \in C} \pi(h)$ and $\gamma_\mathsf{A}(x) = \{h \in \mathbb{H} \mid \pi(h) \leq_\mathsf{A} x\}$.*

α_A and γ_A are each monotone, and together they form a Galois connection [11] which induces meanings for values in $\mathscr{A}_\mathsf{A}$: $\bot$ is the failure state, $x_i \in X$ abstracts sets of in-purview heaps which all map to x_i, ✓ abstracts any set of in-purview heaps, ✗ out-of-purview heaps, and $\top$ any heaps.

We now share three basic examples of abstract domains. In subsequent sections we will reason about concrete and abstract programs which operate on the defined structures. In §7 we explore each example in full.

Example 1 (Non-Negative Counter).

A "hello world" data structure in commutativity is the non-negative counter (NNC)—an integer which may be incremented, decremented (but not below 0), and read. We abstract an NNC located at address p as $\mathsf{Ctr}_p \coloneqq \langle \mathbb{N}, \pi \rangle$, where $\pi(h) \coloneqq n$ if $(p \mapsto n)^+(h)$ for $n \in \mathbb{N}$, else ✗. The *extension* (Def. 1) on the predicate $(p \mapsto n)$ ensures that Ctr_p satisfies locality, allowing the domain to capture any heap *containing* a counter.

Example 2 (Two-Set).

Consider an unordered set containing at most two elements of type $T \subseteq \mathbb{V}^*$. These elements are stored at fixed addresses p and q, with null otherwise stored as a placeholder. One might construct X and π for this structure in a few equivalent ways, the key property being that concrete two-sets containing the same elements should map to the same abstract value, regardless of their order. In the following approach, the abstract values are simply mathematical sets of size at most two. We define $\mathsf{Set}^T_{p,q} \coloneqq \langle \{S \in \wp(T) \mid |S| \leq 2\}, \pi \rangle$ with:

$$\mathsf{set}_{p,q}\big((u,v),\cdot\big) \coloneqq p \mapsto u * q \mapsto v$$

$$\pi(h) \coloneqq \begin{cases} \{u,v\} & \mathsf{set}^+_{p,q}\big((u,v),h\big) \wedge u \neq v \text{ for } u,v \in T \\ \{u\} & \mathsf{set}^+_{p,q}\big((u,\mathsf{null}),h\big) \vee \mathsf{set}^+_{p,q}\big((\mathsf{null},u),h\big) \text{ for } u \in T \\ \varnothing & \mathsf{set}^+_{p,q}\big((\mathsf{null},\mathsf{null}),h\big) \\ \text{✗} & \text{else} \end{cases}$$

Note how we use $\mathsf{set}^+_{p,q}$ throughout π to satisfy locality. One technicality is ensuring that π is well-defined, namely that $\pi(h)$ yields one unambiguous abstract value for each h. In Coq we prove this π is indeed well-defined.

Example 3 (Linked Stack).

Reiterating our overview example (§2), consider a stack implemented on the heap as a linked list and accessed at fixed address ℓ. Constructing new cells when pushing to the stack involves nondeterministic address allocation; two stacks containing the same values but different intermediate addresses should be considered observationally equivalent. To achieve this, our recursive construction of π will existentially quantify intermediate addresses, and not lift them explicitly into the abstract domain. We define $\mathsf{Stk}_\ell \coloneqq \langle \mathrm{list}(\mathbb{V}^*), \pi \rangle$ with

$$\begin{array}{ll} \pi(h) & \coloneqq s \text{ if } \mathsf{stk}^+(s,h) \text{ for } s \in \mathrm{list}(\mathbb{V}^*) \text{ else } \text{✗}, \text{ where} \\ \mathsf{stk}(s) & \coloneqq \exists p \in \mathbb{L} + \mathsf{null}.\ \ell \mapsto p * \psi(s,p) \\ \psi(v :: s, p) & \coloneqq \exists q \in \mathbb{L} + \mathsf{null}.\ p \neq \mathsf{null} * p \mapsto (v,q) * \psi(s,q) \\ \psi([],p) & \coloneqq p = \mathsf{null} \end{array}$$

Again, we must show that each h satisfies $\mathsf{stk}^+(h,s)$ for at most one $s \in \mathrm{list}(\mathbb{V}^*)$.

4.2 Isomorphism of Abstract Domains

A notion of isomorphism between abstract domains will also prove useful:

Definition 6. *A bijection $\varphi : X_\mathsf{A} \to X_\mathsf{B}$ induces an isomorphism between A and B, denoted $\mathsf{A} \cong_\varphi \mathsf{B}$, if $\forall h.\ \varphi(\pi_\mathsf{A}(h)) = \pi_\mathsf{B}(h)$. A and B are* isomorphic ($\mathsf{A} \cong \mathsf{B}$) *if such a bijection exists.*

This definition is stronger than necessary for some of our results; future work may warrant distinguishing isomorphism from a weaker equivalence which relaxes the bijectivity requirement of φ. See the Appendix for an example where we construct a two-set abstraction which is isomorphic to that of Ex. 2.

4.3 Abstract Semantics and Soundness

We design our abstract semantics so that concrete heap behavior (valid or invalid) can be threaded up to the abstract domain with enough precision that, later, abstract commutativity will soundly entail concrete commutativity (§5).

An *abstract program* in the context of abstraction $\mathsf{A} = \langle X, \pi \rangle$ is any function from X to $\mathscr{A}$, from which we derive an *abstract transformer* over $\mathscr{A}$ (recall from Def. 4 that $\mathscr{A} = X + \bot + \top + \text{✗} + \text{✓}$). While these transformers could be constructed in various ways, for our purposes the following will suffice:

Definition 7. *$\widehat{m}$ denotes the abstract transformer derived from $m : X \to \mathscr{A}$, where $\widehat{m}(x) = \bot$ for $x = \bot$, $m(x)$ for $x \in X$, and $\top$ otherwise.*

To relate the behavior of an abstract program m defined in A to a concrete program f, we establish a notion of *soundness* (recall α from Def. 5):

Definition 8. *m* soundly abstracts *f in A, denoted $m \rightsquigarrow_\mathsf{A} f$, if*

1. $\forall C, x.\ \alpha_\mathsf{A}(C) \leq_\mathsf{A} x \implies \alpha_\mathsf{A}\left(\bar{f}(C)\right) \leq_\mathsf{A} \widehat{m}(x)$, *and*

2. $\forall C.\ \alpha_{\mathsf{A}}(C) \in X \ \wedge\ \alpha_{\mathsf{A}}\left(\bar{f}(C)\right) = \bot \implies \widehat{m}(\alpha_{\mathsf{A}}(C)) = \bot.$

The first condition is typical for abstract interpretation, demanding that m *over-approximate* f. This lets us be imprecise with the definition of abstract programs if desirable (e.g. to avoid complex behaviors), the trade-off being that concrete commutativity conditions eventually derived may yield false negatives (i.e. not be the weakest precondition). The second soundness condition demands that m fail (yield $\bot$) whenever f fails on an in-purview input. Because the failure sink state is at the bottom of the abstract lattice (Def. 4), without this condition there could be scenarios where two abstract programs successfully commute, but their corresponding concrete programs fail to execute.

Throughout our examples in §7 we will implicitly use the fact that program soundness is preserved by isomorphism. Namely, if $m \rightsquigarrow_{\mathsf{A}} f$ and $\mathsf{A} \cong_{\varphi} \mathsf{B}$, then $\varphi \circ m \circ \varphi^{-1} \rightsquigarrow_{\mathsf{B}} f$.

We now return to the example of the non-negative counter and its abstraction Ctr_p (see Ex. 1), to define the corresponding concrete and abstract programs for increment and decrement. We cannot define *read* in Ctr_p, because *read* must write its output to an additional location on the heap. We similarly must wait to revisit our two-set and linked stack examples, whose methods all have inputs and/or outputs. We develop the machinery needed for inputs and outputs in §6, and reason fully about all three examples in §7.

Example 4 (Non-Negative Counter).

Consider an NNC allocated at address p. We begin with concrete implementations of increment and decrement, written in the Himp language (§3.3), and implicitly lift them to concrete programs.

$$\begin{array}{ll} \texttt{incr}_p \coloneqq \texttt{let}\ c = \texttt{!}p\ \texttt{in} & \texttt{decr}_p \coloneqq \texttt{let}\ c = \texttt{!}p\ \texttt{in}\ \texttt{let}\ i = c - 1\ \texttt{in} \\ \quad \texttt{let}\ i = c + 1\ \texttt{in}\ p \leftarrow i & \quad \texttt{if}\ i < 0\ \texttt{then skip else}\ p \leftarrow i \end{array}$$

Next, defining abstract increment and decrement functions in Ctr_p is intuitive:

$$\mathit{incr}_p(n) \coloneqq n + 1 \qquad \mathit{decr}_p(n) \coloneqq \max(0, n - 1)$$

We can show that $\mathit{incr}_p \rightsquigarrow_{\mathsf{Ctr}_p} \texttt{incr}_p$ and $\mathit{decr}_p \rightsquigarrow_{\mathsf{Ctr}_p} \texttt{decr}_p$ with typical symbolic execution techniques. Had we defined `decr` to fail when the counter is 0, rather than skip, we would have to define $\mathit{decr}_p(0) = \bot$ to maintain soundness.

5 Sound Commutativity

With our concrete and abstract semantics established, we now define commutativity in the concrete and abstract spaces. We then relate the two with a soundness theorem that reduces concrete commutativity reasoning to simpler reasoning about abstract programs.

5.1 Defining Commutativity

We say concrete programs f and g *commute*, under a notion of observational equivalence on a subset of all heaps, when $f;g$ and $g;f$ both terminate successfully and yield observationally equivalent outcomes. We denote an obs. eq. relation with $[\Psi, \sim]$, for a predicate Ψ on $\mathbb{H}$ and an eq. relation $\sim$ on $\mathbb{H}(\Psi^+)$. Then concrete commutativity is defined precisely as:

Definition 9. *For* commutativity condition P *on* $\mathbb{H}$*,* f *and* g commute under P w.r.t. $[\Psi, \sim]$*, denoted* $f \bowtie^P_{[\Psi,\sim]} g$*, if* $\forall h \in \mathbb{H}(P).\ (f;g)(h), (g;f)(h) \neq \varnothing$ *and* $(f;g)(h) \cup (g;f)(h) \subseteq [h']_\sim$ *for some* $h' \in \mathbb{H}(\Psi^+)$*.*

The eq. relation we use for concrete commutativity is often very similar to the relation implicitly induced by an abstraction. We can make this explicit by deriving a concrete relation in terms of an abstraction, a technique we will use extensively in the remainder of the paper.

Definition 10. *For abstraction* A*, define concrete equivalence* $[\mathsf{Prv}(\mathsf{A}), \sim_\mathsf{A}]$ *with* $h \sim_\mathsf{A} h' \coloneqq \pi_\mathsf{A}(h) = \pi_\mathsf{A}(h')$*. Define the shorthand* $f \bowtie^P_\mathsf{A} g$ *for* $f \bowtie^P_{[\mathsf{Prv}(\mathsf{A}),\sim_\mathsf{A}]} g$*.*

Naturally, a user is responsible for choosing an observational equivalence relation (or in light of Def. 10, an abstraction) which is sufficiently descriptive for their purposes. Properties of concrete commutativity include:

$$\frac{\mathbb{H}(P') \subseteq \mathbb{H}(P) \qquad f \bowtie^P_{[\Psi,\sim]} g}{f \bowtie^{P'}_{[\Psi,\sim]} g} \qquad \frac{\mathsf{A} \cong \mathsf{B} \qquad g \bowtie^P_\mathsf{A} f}{g \bowtie^P_\mathsf{B} f} \qquad \frac{f \bowtie^P_{[\Psi,\sim]} g}{g \bowtie^P_{[\Psi,\sim]} f} \qquad \frac{\forall h, h'.\ h \sim h' \implies h \sim' h' \qquad f \bowtie^P_{[\Psi,\sim]} g}{f \bowtie^P_{[\Psi',\sim']} g}$$

Next we define *abstract commutativity*, in a manner so that we can achieve a useful soundness guarantee in relation to concrete commutativity.

Definition 11. *Abstract programs* m *and* n *defined in* A commute under *predicate* Q *on* X_A*, denoted* $m \mathrel{\hat{\bowtie}}^Q_\mathsf{A} n$*, if* $\forall x \in X_\mathsf{A}.\ Q(x) \Rightarrow \widehat{n}(m(x)) = \widehat{m}(n(x)) \in X_\mathsf{A}$*.*

Note how we only accept outcomes within the purview of A. We reject $\bot$, as non-failure is a prerequisite for commutativity. We also reject ✗, ✓, and $\top$, as such outcomes are not precise enough to guarantee meaningful commutativity.

5.2 Sound Commutativity Theorem

To relate abstract and concrete commutativity, we must first relate the relevant abstraction and observational equivalence relation:

Definition 12. A captures $[\Psi, \sim]$ *iff* $\mathbb{H}(\Psi^+) \subseteq \mathsf{Prv}(\mathsf{A})$ *and*
$\forall h, h' \in \mathbb{H}(\Psi).\ \pi_\mathsf{A}(h) = \pi_\mathsf{A}(h') \in X_\mathsf{A} \implies h \sim h'$*.*

This definition of *capture* is preserved by abstraction isomorphism. Additionally, A always captures $[\mathsf{Prv}(\mathsf{A}), \sim_\mathsf{A}]$, a fact we will frequently use implicitly.

We now present the central result of our work, the *sound commutativity theorem*, which offloads the verification (via symbolic execution) of concrete commutativity conditions to much simpler reasoning about abstract programs:

Theorem 1. *If* A *captures* $[\Psi, \sim]$, $m \rightsquigarrow_{\mathsf{A}} f$, $n \rightsquigarrow_{\mathsf{A}} g$, *and* $m \mathrel{\hat{\bowtie}}^{Q}_{\mathsf{A}} n$, *then* $f \bowtie^{P^+}_{[\Psi,\sim]} g$, *where* $P(h) \equiv \pi_{\mathsf{A}}(h) \in X_{\mathsf{A}} \wedge Q(\pi_{\mathsf{A}}(h))$.

Sound commutativity takes a valid abstract commutativity condition, and transforms it automatically into a concrete condition. For instance, suppose π_{A} takes the fairly common form $\pi_{\mathsf{A}}(h) = x$ if $\Phi(x, h)$ for $x \in X_{\mathsf{A}}$, else $\boldsymbol{\times}$, for some heap predicate Φ. In this case, the concrete condition we derive from abstract condition Q is $\exists x \in X.\ \Phi^+(x, h) * Q(x)$. Indeed, any condition we derive with Thm. 1 is *extended* (Def. 1). This provides *framing* for free—it guarantees if f and g commute in some heap context, then they also commute in any larger context.

With sound commutativity established, we can summarize the overall pattern of reasoning: (i) construct an abstract domain, (ii) define abstract programs for each concrete program and prove soundness between them using symbolic execution, (iii) verify or synthesize commutativity conditions for abstract program pairs, and (iv) immediately derive concrete conditions.

Example 5 (Non-Negative Counter).

Returning to the NNC programs defined in Ex. 4, we can perform abstract commutativity reasoning easily. For instance, to verify that $incr_p \mathrel{\hat{\bowtie}}^{Q}_{\mathsf{Ctr}_p} decr_p$ for $Q(n) \equiv n > 0$, we simply evaluate:

$$\widehat{decr}_p(incr_p(n)) = \widehat{decr}_p(n+1) = n \qquad \widehat{incr}_p(decr_p(n)) = \widehat{incr}_p(n-1) = n$$

With proofs of soundness (Def. 8) of *incr* and *decr* w.r.t. `incr` and `decr`, we can immediately derive via Thm. 1 that $f \bowtie^{P^+}_{\mathsf{Ctr}_p} g$ where $P \equiv \exists n \in \mathbb{N}^+.\ p \mapsto n$. Note how we use the equivalence imposed by Ctr_p as our concrete observational equivalence relation (Def. 10).

We might *synthesize* an abstract condition interactively (we leave automation to future work), by trying to prove that *incr* and *decr* commute under $\top$, and seeing where the proof gets stuck. In this case, we will get stuck showing that $1 = 0$ under an initial abstract value of 0. By constraining that $n > 0$, the proof can be completed; thus $n > 0$ is a valid precondition.

6 Abstract Domain Composition

In this section, we expand upon our established tools for abstracting data structures and performing commutativity analysis, by defining composition of abstract domains. This enables compositional reasoning in settings with multiple data structures, and provides machinery needed to reason about concrete programs with heap-allocated inputs and outputs. Various compositional operators on abstract domains have been used in abstract interpretation, such as Cartesian product, reduced product, and reduced cardinal power [9]. While an existing operator is likely suitable, we elect to design a simple one from scratch.

A composition of abstractions should yield a new abstraction, and should respect the disjointness of the concrete structures which are individually abstracted. We achieve this with the *abstract conjunction* operator, denoted $\mathsf{A} * \mathsf{B}$,

which is defined to mirror the concrete separating conjunction. The purview of $\mathsf{A} * \mathsf{B}$ will contain heaps which can be partitioned into joint halves that are respectively in the purviews of A and B.

Definition 13. $\mathsf{A} * \mathsf{B} \triangleq \langle X_{\mathsf{A}} \times X_{\mathsf{B}}, \pi \rangle$, *where* $\pi(h) = (\pi_{\mathsf{A}}(h_1), \pi_{\mathsf{B}}(h_2))$ *if* $h = h_1 + h_2$ *for some* $h_1 \perp h_2$ *s.t.* $h_1 \in \mathsf{Prv}(\mathsf{A})$ *and* $h_2 \in \mathsf{Prv}(\mathsf{B})$, *else* ✗.

A technicality of Def. 13 is ensuring that $\pi(h)$ maps to a *unique* (a, b) pair (or to ✗) for every h. We show in Coq that this uniqueness does hold, due to the locality constraint of Def. 4. Additionally, when A and B satisfy locality, so too does $\mathsf{A} * \mathsf{B}$. Abstract conjunction is compatible with, and is commutative and associative up to, abstraction isomorphism; we use these facts implicitly throughout §7. Much like the heap separating conjunction, $\mathsf{A} * \mathsf{B}$ may be *unsatisfiable*, i.e. $\pi_{\mathsf{A}*\mathsf{B}}$ maps all heaps to ✗. Commutativity reasoning w.r.t. an unsatisfiable domain is still valid, albeit meaningless, as there are no in-purview heaps. Applying the concrete observational equivalence construction of Def. 10 to composition, we get $[\mathsf{Prv}(\mathsf{A} * \mathsf{B}), \sim_{\mathsf{A}*\mathsf{B}}]$ where $h \sim_{\mathsf{A}*\mathsf{B}} h'$ if $h = h_a + h_b$ and $h' = h'_a + h'_b$ s.t. $h_a \sim_{\mathsf{A}} h'_a$ and $h_b \sim_{\mathsf{B}} h'_b$.

6.1 Concrete Program Capture

Effective compositional reasoning will often require that behavior of the concrete programs of interest are "captured" by the abstract domains in use. For instance, `incr` has well-defined behavior w.r.t. Ctr_p, but not Stk_ℓ. If we prove properties about `incr` w.r.t. Ctr_p, we should be able to derive how `incr` behaves w.r.t. $\mathsf{Ctr}_p * \mathsf{Stk}_\ell$ (namely it leaves the stack untouched). However, if we reason about `incr` w.r.t. Stk_ℓ, we will learn nothing about how `incr` behaves w.r.t. $\mathsf{Stk}_\ell * \mathsf{Ctr}_p$.

To this end, we say that an abstraction A *captures* program f if (1) heaps in the purview of A are in the footprint of f, (2) f maps in-purview heaps to in-purview heaps, and (3) f maintains its specific in-purview mappings when any disjoint heaps are tacked onto the initial heap (thus strengthening *local action* in an abstraction-specific manner). Precisely:

Definition 14. A captures *concrete program* f *if:*

1. $\mathsf{Prv}(\mathsf{A}) \subseteq \mathsf{foot}(f)$, *and*
2. $\forall h \in \mathsf{Prv}(\mathsf{A}) \cap \mathsf{suff}(f).\ \alpha_{\mathsf{A}}(f(h)) \in X_{\mathsf{A}}$, *and*
3. $\forall h \in \mathsf{Prv}(\mathsf{A}) \cap \mathsf{foot}(f), h' \perp h.\ \alpha_{\mathsf{A}}\big(\{k - h' \mid k \in f(h + h')\}\big) = \alpha_{\mathsf{A}}(f(h))$, *where* $k - h'$ *is the subheap of* k *with domain* $\mathsf{dom}(k) \setminus \mathsf{dom}(h')$.

Proving capture can be less burdensome than it appears—if we have shown that $m \rightsquigarrow_{\mathsf{A}} f$ for m s.t. $\mathsf{image}(m) \subseteq X_{\mathsf{A}} + \perp$, then conditions (1) and (2) immediately hold for f. We also expect that (3) should hold automatically for all abstractions and all Himp programs (and other typical heap languages). Capture is preserved by abstraction isomorphism, and if A captures f then $\mathsf{A} * \mathsf{B}$ captures f.

Capture yields a common form of concrete commutativity: *commutativity from noninterference*. If two programs operate on disjoint structures, then they automatically commute whenever they individually terminate without failure:

Theorem 2. *If* A *and* B *capture* f *and* g *respectively, then* $f \bowtie^{P^+}_{\mathsf{A}*\mathsf{B}} g$ *where* $P(h) \equiv h \in \mathsf{suff}(f) \cap \mathsf{suff}(g) \cap \mathsf{Prv}(\mathsf{A} * \mathsf{B})$.

Finally, we will often discuss program capture and soundness in the same breath, so we provide a convenient shorthand:

Definition 15. $m \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f$ *denotes that* $m \rightsquigarrow_{\mathsf{A}} f$ *and* A *captures* f.

6.2 Compositional Commutativity

In service of sound commutativity reasoning in compound abstract domains, we define a way to compose abstract programs:

Definition 16. *Given* m *and* n *defined in* A *and* B *respectively, their* conjunction *is* $m * n$ *in* $\mathsf{A} * \mathsf{B}$*, where for* $a \in X_{\mathsf{A}}$ *and* $b \in X_{\mathsf{B}}$, $(m * n)(a, b) = \bot$ *if* $m(a) = \bot \vee n(b) = \bot$, $(m(a), n(b))$ *if* $m(a) \in X_{\mathsf{A}} \wedge n(b) \in X_{\mathsf{B}}$, *else* $m(a) \sqcup n(b)$.

The *compound program soundness theorem* then states that the conjunction of two abstract programs soundly abstracts the *sequence* of two concrete programs (in either order), so long as these concrete programs operate on disjoint structures (via *capture*). We also derive an *abstract frame rule* of sorts.

Theorem 3. *If* $m \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f$ *and* $n \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{B}} g$, *then* $m*n \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}*\mathsf{B}} f;g$ *and* $m*n \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}*\mathsf{B}} g;f$.

Corollary 1. *If* $m \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f$, *then* $m * \mathrm{id} \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}*\mathsf{B}} f$.

With composition of abstract domains and their programs established, we can demonstrate new *compositional* commutativity soundness properties, beginning with *compound abstract commutativity*:

Lemma 1. *If* $m \hat{\bowtie}^{Q}_{\mathsf{A}} n$ *and* $m' \hat{\bowtie}^{Q'}_{\mathsf{A}'} n'$, *then* $m * m' \hat{\bowtie}^{P}_{\mathsf{A}*\mathsf{A}'} n * n'$ *where* $P(a, b) \equiv Q(a) \wedge Q'(b)$.

From this, along with sound commutativity (Thm. 1) and compound program soundness (Thm. 3), we yield *compound sound commutativity*, which allows us to compose individual commutativity results about pairs of programs operating on disjoint structures:

Theorem 4. *If* $m \hat{\bowtie}^{Q}_{\mathsf{A}} m'$, $n \hat{\bowtie}^{Q'}_{\mathsf{B}} n'$, $m \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f$, $m' \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f'$, $n \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{B}} g$, *and* $n' \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{B}} g'$, *then* $f;g \bowtie^{P^+}_{\mathsf{A}*\mathsf{B}} f';g'$ *where* $P(h) \equiv \pi_{\mathsf{A}*\mathsf{B}}(h) = (a, b) \in X_{\mathsf{A}*\mathsf{B}} \wedge Q(a) \wedge Q'(b)$.

Note how this result still holds even if $\mathsf{A}*\mathsf{B}$ is unsatisfiable, because the concrete commutativity condition requires that the initial heap maps to a valid abstract value. From this theorem we derive *framed sound commutativity*:

Corollary 2. *If* $m \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} f$ *and* $n \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} g$, *and* $m \hat{\bowtie}^{Q}_{\mathsf{A}} n$, *then* $f \bowtie^{P^+}_{\mathsf{A}*\mathsf{B}} g$, *where* $P(h) \equiv h \in \mathsf{Prv}(\mathsf{A} * \mathsf{B}) \wedge Q(\mathrm{fst}(\pi_{\mathsf{A}*\mathsf{B}}(h)))$.

Let us contrast this result with our prior sound commutativity result (Thm. 1), which also includes a form of framing. Thm. 1 says that, *within* the parameters of a chosen abstraction and concrete observational eq. relation, the derived concrete commutativity condition is framed. On the other hand, Cor. 2 guarantees—under the stronger premise of *capture*—that the programs commute under a framed concrete condition within A, as well as within the purview and observational equivalence imposed by any compound abstraction *containing* A.

7 Examples in Full

In this section we reiterate and expand on our working examples: a non-negative counter, a two-set, a linked-list stack, and the composition of a stack and counter. For each example, we (i) define concrete programs; (ii) construct an abstract domain and abstract programs; (iii) perform abstract commutativity reasoning; (iv) prove soundness between the concrete and abstract programs; and (v) derive concrete commutativity conditions. As (v) follows directly from (iii) and (iv), we sometimes omit it for brevity.

Throughout these examples, we will need to abstract the arguments and return values of methods. To that end we define the "address domain" Addr, which simply captures the value stored at a given heap address. For address p and values $V \subseteq \mathbb{V}$ considered valid, construct $\mathsf{Addr}_p^V \coloneqq \langle V, \pi \rangle$, where $\pi(h) \coloneqq v$ if $(p \mapsto v)^+(h)$ for $v \in V$, else ✗.

We will also use some shorthands. We may omit abstract program compositions rendered trivial by (Cor. 1), e.g. if m is defined in A and we are reasoning in $\mathsf{A} * \mathsf{B}$, we may write m rather than $m * id$. Furthermore, we will leverage the associativity and commutativity of abstract composition up to isomorphism to reason about abstract programs which have been defined in a similar but "out-of-order" compound abstraction. We use the addresses on which an abstract program is parameterized as unique identifiers for how the input and output tuples of said program should be implicitly permuted.

7.1 Non-negative Counter

Consider a non-negative counter structure allocated at p, featuring the following concrete operations written in Himp:

```
incr_p :=                 decr_p := let c = !p in      read_p^r :=
  let c = !p in             let i = c − 1 in             let c = !p in
  let i = c + 1 in          if i < 0 then skip           r ← c
  p ← i                     else p ← i
```

Note how `read` writes its output to a location r. Define $\mathsf{Ctr}_p \coloneqq \langle \mathbb{N}, \pi \rangle$, where $\pi(h) \coloneqq n$ if $(p \mapsto n)^+(h)$ for $n \in \mathbb{N}$, else ✗. Next construct the abstract *incr* and *decr* programs in Ctr_p, and *read* in $\mathsf{A} = \mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}$ to accommodate its output:

$$incr_p(n) \coloneqq n + 1 \qquad decr_p(n) \coloneqq \max(0, n - 1) \qquad read_p^r(n, \_) \coloneqq (n, n)$$

We verify commutativity conditions for *incr* and *decr* w.r.t. an initial value $n \in X_{\mathsf{Ctr}_p}$. By evaluating the abstract programs (as in Ex. 5) we find that $incr_p$ and $incr_p$ commute under $\top$, $incr_p$ and $decr_p$ under $n > 0$, and $decr_p$ and $decr_p$ under $\top$. We cannot compare *incr* or *decr* with *read* directly, as they are defined in different domains. So we consider $incr_p * id$ and $decr_p * id$ in $\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}$, while preserving soundness w.r.t. the concrete programs via Cor. 1. We can verify that, for an initial value $(n, \_) \in X_{\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}}$, $incr_p * id$ and $read_p^r$ never commute, and $decr_p * id$ and $read_p^r$ commute under $n = 0$.

Lastly, we pair *read* with itself, recognizing that the two *read* executions should write their outputs to different locations (this idea returns in the subsequent examples). So with addresses r and s, we consider $read_p^r * id$ and $read_p^s * id$ defined in $\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}} * \mathsf{Addr}_s^{\mathbb{N}}$, and find that they commute under $\top$.

To derive concrete commutativity from these abstract results, we first find with symbolic execution the soundness and capture facts (recall shorthand 15) that $incr_p \overset{\mathsf{CAP}}{\rightsquigarrow}_{\mathsf{Ctr}_p} \texttt{incr}_p$, $decr_p \overset{\mathsf{CAP}}{\rightsquigarrow}_{\mathsf{Ctr}_p} \texttt{decr}_p$, and $read_p^r \overset{\mathsf{CAP}}{\rightsquigarrow}_{\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}} \texttt{read}_p^r$. Now we may immediately derive *framed* concrete commutativity conditions via Thm. 1. For each program pair, we characterize $f \bowtie_{\mathsf{A}}^{P^+} g$ (recall shorthand 10):

f	g	A	P
$\texttt{incr}_p$	$\texttt{incr}_p$	Ctr_p	$\exists n \in \mathbb{N}.\ p \mapsto n$
$\texttt{incr}_p$	$\texttt{decr}_p$	Ctr_p	$\exists n \in \mathbb{N}^+.\ p \mapsto n$
$\texttt{decr}_p$	$\texttt{decr}_p$	Ctr_p	$\exists n \in \mathbb{N}.\ p \mapsto n$
$\texttt{incr}_p$	$\texttt{read}_p^r$	$\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}$	$\bot$
$\texttt{decr}_p$	$\texttt{read}_p^r$	$\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}}$	$\exists x \in \mathbb{N}.\ p \mapsto 0 * r \mapsto x$
$\texttt{read}_p^r$	$\texttt{read}_p^s$	$\mathsf{Ctr}_p * \mathsf{Addr}_r^{\mathbb{N}} * \mathsf{Addr}_s^{\mathbb{N}}$	$\exists n, x, y \in \mathbb{N}.\ p \mapsto n * r \mapsto x * s \mapsto y$

7.2 Two-Set

Consider an unordered set containing at most two elements of type $T \subseteq \mathbb{V}^*$. We define concrete add, remove, and member-test methods, each parameterized on the set locations p and q (we omit p and q subscripts for brevity), as well as input address a and/or output address r.

```
add_a ≔ let u = !p in
  let v = !q in
  let x = !a in
  if u = null then
    p ← x else
  if v = null then
    q ← x else
  skip
```

```
rem_a ≔ let u = !p in
  let v = !q in
  let x = !a in
  if u = x then
    p ← null else
  if v = x then
    q ← null else
  skip
```

```
mem_a^r ≔ let u = !p in
  let v = !q in
  let x = !a in
  if u = x then
    r ← true else
  if v = x then
    r ← true else
  r ← false
```

Define abstraction $\mathsf{Set}^T_{p,q} \coloneqq \langle \{S \in \wp(T) \mid |S| \leq 2\}, \pi \rangle$ with:

$$\mathsf{set}_{p,q}\big((u,v),\cdot\big) \coloneqq p \mapsto u * q \mapsto v$$

$$\pi(h) \coloneqq \begin{cases} \{u,v\} & \mathsf{set}^+_{p,q}\big((u,v),h\big) \wedge u \neq v \text{ for } u,v \in T \\ \{u\} & \mathsf{set}^+_{p,q}\big((u,\mathsf{null}),h\big) \vee \mathsf{set}^+_{p,q}\big((\mathsf{null},u),h\big) \text{ for } u \in T \\ \varnothing & \mathsf{set}^+_{p,q}\big((\mathsf{null},\mathsf{null}),h\big) \\ \text{✗} & \text{else} \end{cases}$$

Next we construct abstract programs *add* and *rem* in $\mathsf{B} = \mathsf{Set}^T_{p,q} * \mathsf{Addr}^T_a$, and *mem* in $\mathsf{C} = \mathsf{Set}^T_{p,q} * \mathsf{Addr}^T_a * \mathsf{Addr}^{\mathbb{B}}_r$:

$$\begin{aligned} \mathit{add}_a(S,v) &\coloneqq (S',v) \text{ where } S' = S \cup \{v\} \text{ if } |S \cup \{v\}| \leq 2 \text{ else } S \\ \mathit{rem}_a(S,v) &\coloneqq (S \setminus \{v\}, v) \\ \mathit{mem}^r_a(S,v,\_) &\coloneqq \big(S, v, \mathsf{true} \text{ if } v \in S \text{ else } \mathsf{false}\big) \end{aligned}$$

We analyze commutativity w.r.t. $\mathsf{A} = \mathsf{Set}^T_{p,q} * \mathsf{Addr}^T_a * \mathsf{Addr}^T_r * \mathsf{Addr}^T_b * \mathsf{Addr}^T_s$ (though only the *mem*/*mem* pair utilizes all four Addr domains).

Programs	Condition on $(S,u,\_,v,\_) \in X_{\mathsf{A}}$
add_a add_b	$S = \varnothing \vee \lvert S\rvert = 2 \vee u \in S \vee v \in S \vee u = v$
add_a rem_b	$\lvert S\rvert < 2 \wedge u \neq v \vee$
	$\lvert S\rvert = 2 \wedge (u, v \notin S \vee u \in S \wedge u \neq v)$
rem_a rem_b	$\top$
mem^r_a add_b	$u \neq v \vee v \in S \vee \lvert S\rvert = 2$
mem^r_a rem_b	$u \neq v \vee v \notin S$
mem^r_a mem^s_b	$\top$

After proving that $\mathit{add}_a \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{B}} \texttt{add}_a$, $\mathit{rem}_a \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{B}} \texttt{rem}_a$, and $\mathit{mem}^r_a \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{C}} \texttt{mem}^r_a$, we again can derive concrete commutativity conditions for `add`, `rem`, and `mem`. In particular, we use the observational equivalence induced by A as our concrete eq. relation, which ensures that concrete two-sets with the same values in different orders are considered equivalent. We omit for brevity the concrete conditions, which each amount to "the program pair commutes w.r.t. $[\mathsf{Prv}(\mathsf{A}), \sim_{\mathsf{A}}]$ on h if $\pi(h)$ satisfies the abstract condition."

7.3 Linked Stack

We now expand upon the example from the overview (§2) of a stack of non-null values, implemented as a linked list accessible at address ℓ. We consider its `push`, `pop`, and `peek` methods, where `pop` fails on an empty stack. To account for the lack of an `is_empty` operation, `peek` outputs null on an empty stack.

$\texttt{push}^a_\ell \coloneqq \texttt{let}\ p = \texttt{!}\ell\ \texttt{in}$
$\quad \texttt{let}\ v = \texttt{!}a\ \texttt{in}$
$\quad \texttt{let}\ x = (v, p)\ \texttt{in}$
$\quad \texttt{let}\ q = \texttt{ref}\ x\ \texttt{in}$
$\quad \ell \leftarrow q$

$\texttt{pop}^a_\ell \coloneqq \texttt{let}\ p = \texttt{!}\ell\ \texttt{in}$
$\quad \texttt{if}\ p = \mathsf{null}\ \texttt{then}$
$\quad\quad \texttt{fail else}$
$\quad \texttt{let}\ x = \texttt{!}p\ \texttt{in}$
$\quad \texttt{let}\ v = \texttt{fst}\ x\ \texttt{in}$
$\quad \texttt{let}\ q = \texttt{snd}\ x\ \texttt{in}$
$\quad \ell \leftarrow q\ \texttt{;}\ \texttt{free}\ p\ \texttt{;}$
$\quad a \leftarrow v$

$\texttt{peek}^a_\ell \coloneqq \texttt{let}\ p = \texttt{!}\ell\ \texttt{in}$
$\quad \texttt{if}\ p = \mathsf{null}\ \texttt{then}$
$\quad\quad a \leftarrow \mathsf{null}\ \texttt{else}$
$\quad \texttt{let}\ x = \texttt{!}p\ \texttt{in}$
$\quad \texttt{let}\ v = \texttt{fst}\ x\ \texttt{in}$
$\quad a \leftarrow v$

Once again, we abstraction define $\mathsf{Stk}_\ell \coloneqq \langle \mathrm{list}(\mathbb{V}^*), \pi \rangle$ where

$$\pi(h) \coloneqq s \text{ if } \mathsf{stk}^+(s,h) \text{ for } s \in \mathrm{list}(\mathbb{V}^*) \text{ else } \text{✗}$$

$$\mathsf{stk}(s) \coloneqq \exists p \in \mathbb{L} + \mathsf{null}.\ \ell \mapsto p * \psi(s,p)$$
$$\psi(v :: s, p) \coloneqq \exists q \in \mathbb{L} + \mathsf{null}.\ p \neq \mathsf{null} * p \mapsto (v,q) * \psi(s,q)$$
$$\psi([], p) \coloneqq p = \mathsf{null}$$

Define abstract *push*/*pop* in $\mathsf{B} \coloneqq \mathsf{Stk}_\ell * \mathsf{Addr}_a^{\mathbb{V}^*}$, and *peek* in $\mathsf{C} \coloneqq \mathsf{Stk}_\ell * \mathsf{Addr}_a^{\mathbb{V}}$:

$$\mathit{push}^a_\ell(s,v) \coloneqq (v :: s, v) \qquad \mathit{pop}^a_\ell([], \_) \coloneqq \bot \qquad \mathit{peek}^a_\ell(v :: s, \_) \coloneqq (v :: s, v)$$
$$\mathit{pop}^a_\ell(v :: s, \_) \coloneqq (s, v) \qquad \mathit{peek}^a_\ell([], \_) \coloneqq ([], \mathsf{null})$$

Next we verify commutativity conditions for each abstract program pair w.r.t. $\mathsf{A} = \mathsf{Stk}_\ell * \mathsf{Addr}_a^U * \mathsf{Addr}_b^V$, where U and V depend on the pair.

Programs		U	V	Cond. on $(s,u,v) \in X_\mathsf{A}$
push^a_ℓ	push^b_ℓ	$\mathbb{V}^*$	$\mathbb{V}^*$	$u = v$
push^a_ℓ	pop^b_ℓ	$\mathbb{V}^*$	$\mathbb{V}^*$	$s = u :: \_$
push^a_ℓ	peek^b_ℓ	$\mathbb{V}^*$	$\mathbb{V}$	$s = u :: \_$
pop^a_ℓ	pop^b_ℓ	$\mathbb{V}^*$	$\mathbb{V}^*$	$s = \_ :: \_ :: \_$
pop^a_ℓ	peek^b_ℓ	$\mathbb{V}^*$	$\mathbb{V}$	$\exists x.\ s = x :: x :: \_$
peek^a_ℓ	peek^b_ℓ	$\mathbb{V}$	$\mathbb{V}$	$\top$

We then find through symbolic execution that $\mathit{push}^a_\ell \overset{\mathsf{CAP}}{\rightsquigarrow}_\mathsf{B} \texttt{push}^a_\ell$, $\mathit{pop}^a_\ell \overset{\mathsf{CAP}}{\rightsquigarrow}_\mathsf{B} \texttt{pop}^a_\ell$, and $\mathit{peek}^a_\ell \overset{\mathsf{CAP}}{\rightsquigarrow}_\mathsf{C} \texttt{peek}^a_\ell$. Finally, we freely derive commutativity conditions for each concrete pair w.r.t. $[\mathsf{Prv}(\mathsf{A}), \sim_\mathsf{A}]$, which we again omit for brevity.

7.4 Composition of Stack and Counter

Suppose we augment our stack data structure with a counter to track the stack's size. The concrete operations for this structure would be

$$\texttt{cpush}^a_{\ell,p} \coloneqq \texttt{push}^a_\ell\ \texttt{;}\ \texttt{incr}_p \qquad \texttt{cpeek}^a_\ell \coloneqq \texttt{peek}^a_\ell$$
$$\texttt{cpop}^a_{\ell,p} \coloneqq \texttt{pop}^a_\ell\ \texttt{;}\ \texttt{decr}_p \qquad \texttt{csize}^r_p \coloneqq \texttt{read}^r_p$$

We can similarly define abstract programs with composition. Namely, the following are defined in $\mathsf{Stk}_\ell * \mathsf{Ctr}_p$, with Addr domains joined on as appropriate:

$$\mathit{cpush}^a_{\ell,p} \coloneqq \mathit{push}^a_\ell * \mathit{incr}_p \qquad \mathit{cpeek}^a_\ell \coloneqq \mathit{peek}^a_\ell * \mathit{id}$$
$$\mathit{cpop}^a_{\ell,p} \coloneqq \mathit{pop}^a_\ell * \mathit{decr}_p \qquad \mathit{csize}^r_p \coloneqq \mathit{id} * \mathit{read}^r_p$$

We analyze commutativity of these abstract programs in $\mathsf{A} = \mathsf{Stk}_\ell * \mathsf{Ctr}_p * \mathsf{Addr}_a^U * \mathsf{Addr}_b^V$. For each pair, we leverage Lem. 1 to derive a condition for free from our prior analyses (ℓ and p subscripts are omitted):

Programs		U	V	Cond. on $(s, n, u, v) \in X_{\mathsf{A}}$
$cpush^a$	$cpush^b$	$\mathbb{V}^*$	$\mathbb{V}^*$	$u = v$
$cpush^a$	$cpop^b$	$\mathbb{V}^*$	$\mathbb{V}^*$	$s = u :: \_ \wedge n > 0$
$cpush^a$	$cpeek^b$	$\mathbb{V}^*$	$\mathbb{V}$	$s = u :: \_$
$cpush^a$	$csize^b$	$\mathbb{V}^*$	$\mathbb{N}$	$\bot$
$cpop^a$	$cpop^b$	$\mathbb{V}^*$	$\mathbb{V}^*$	$s = \_ :: \_ :: \_$
$cpop^a$	$cpeek^b$	$\mathbb{V}^*$	$\mathbb{V}$	$\exists x.\ s = x :: x :: \_$
$cpop^a$	$csize^b$	$\mathbb{V}^*$	$\mathbb{N}$	$s = \_ :: \_ \wedge n = 0$
$cpeek^a$	$cpeek^b$	$\mathbb{V}$	$\mathbb{V}$	$\top$
$cpeek^a$	$csize^b$	$\mathbb{V}$	$\mathbb{N}$	$\top$
$csize^a$	$csize^b$	$\mathbb{N}$	$\mathbb{N}$	$\top$

For our requisite soundness and capture facts, we use compound soundness (Thm. 3) to derive for free that $cpush^a_{\ell,p} \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} \texttt{cpush}^a_{\ell,p}$, $cpop^a_{\ell,p} \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} \texttt{cpop}^a_{\ell,p}$, $cpeek^a_\ell \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} \texttt{cpeek}^a_\ell$, and $csize^a_p \overset{\text{CAP}}{\rightsquigarrow}_{\mathsf{A}} \texttt{csize}^a_p$. With this and our abstract commutativity conditions, we may freely derive concrete conditions for `cpush`, `cpop`, `cpeek`, and `csize`. One interesting case is the *cpop*/*csize* condition, which demands a nonempty stack with size 0. This cannot occur with a correctly initialized stack and counter, meaning in practice `cpop` and `csize` never commute.

8 Related Work

To the best of our knowledge there are no existing works on commutativity reasoning specifically geared toward heap-based programs, aside from preliminary work by Pincus [27] who only considers deterministic programs. We covered some other works in §1, and here discuss some of those and others in more detail.

Commutativity without the heap. Some prior works focused on verifying and even inferring commutativity properties, though not with a heap memory model. These include aforementioned work such as Kim and Rinard [21] who verified commutativity properties in two steps: verifying an implementation satisfies its ADT specs, and then verifying commutativity of the ADT spec. Further in that direction, Bansal *et al.* [2,3] showed that commutativity conditions of ADT specs could be synthesized, and introduced the tool Servois, which was later improved as Servois2 [7]. Chen *et al.* [6] adapted these approaches to perform commutativity analysis on (non-heap) imperative programs. Finally, Koskinen and Bansal [22] also verified commutativity, but not of heap-based programs.

Commutativity with the heap. Pîrlea *et al.* [28] describe the CoSplit tool, which performs a commutativity analysis on a smart contract language. Their analysis determines commutativity of heap operations, but only when commutativity can be determined based on heap ownership. Thus, their approach could

not handle the examples in this paper including the Counter, which involves commutative updates to the same heap cell. Eilers *et al.* [13] leverage commutativity modulo user-specified abstractions to prove information flow security in the space of relational concurrent separation logic.

Exploiting commutativity for verification. Commutativity is a widely used abstraction in many verification tools, which use commutativity specifications as user-provided inputs. In the context of concurrent programs, QED [14] and later CIVL [23] both use commutativity (more specifically, left/right movers) to build atomic sections. The Anchor [20] verifier also uses commutativity for a more automated approach of verifying concurrent programs. Commutativity is also used for proofs of parameterized programs [18], operational-style proofs of concurrent objects [15] and termination of concurrent programs [24]. Abstract commutativity relations have also been used in reductions for verification [17]. A summary of the use of commutativity in verification was given by [16].

Abstract domains for the heap. Other works have focused on the intersection of abstract interpretation and heap logics, although they do not specifically target commutativity, and generally seek to abstract the semantics of separation logic rather than reasoning about particular allocated structures. Sims [33] constructs a detailed abstract domain for representing separation logic heap predicates. Calcagno *et al.* [4] introduce separation algebras, which generalize the semantics of separation logic.

9 Conclusion

We have demonstrated how commutativity analysis on concrete heap programs can be reduced to much simpler reasoning about mathematical objects in an abstract space. We have designed our abstract domain to account for allocation nondeterminism, concrete observational equivalence, and improperly allocated heap structures. We formalized this domain, laid out the abstract semantics and program soundness relation, established a sound commutativity theorem, and introduced composition of abstractions for multi-structure program settings. We then worked through several examples illustrating the convenience of analyzing abstract program commutativity and deriving concrete results. Our work has been implemented in Coq and is available in the supplement.

In future work we plan to pursue automation by implementing an abstract interpreter, and exploring how it can be combined with existing commutativity synthesis techniques in the abstract domain [2,7]. Furthermore, we will explore how our model of data structure abstraction could be applicable to various kinds of heap-style reasoning other than commutativity. We are also interested in augmenting our abstract domain to feature an abstract value subset lattice, which would permit nondeterministic abstract programs and reasoning about bounded divergence commutativity.

Acknowledgement. We thank Marco Gaboardi, David Naumann, VFC, and the anonymous reviewers for their feedback on earlier versions of this draft. Both authors

were partially supported by NSF award #2008633. Koskinen was partially supported by NSF award #2315363. Pincus was partially supported by NSF award #1801564.

Disclosure of Interests. todo

References

1. Antonopoulos, T., Koskinen, E., Le, T.C., Nagasamudram, R., Naumann, D.A., Ngo, M.: An algebra of alignment for relational verification. Proc. ACM Program. Lang. **7**(POPL), 573–603 (2023). https://doi.org/10.1145/3571213
2. Bansal, K., Koskinen, E., Tripp, O.: Automatic generation of precise and useful commutativity conditions (extended version). CoRR **abs/1802.08748** (2018), http://arxiv.org/abs/1802.08748
3. Bansal, K., Koskinen, E., Tripp, O.: Synthesizing precise and useful commutativity conditions. J. Autom. Reason. **64**(7), 1333–1359 (2020). https://doi.org/10.1007/S10817-020-09573-W
4. Calcagno, C., O'Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: 22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007). pp. 366–378 (2007). https://doi.org/10.1109/LICS.2007.30
5. Charguéraud, A.: Separation Logic Foundations, Software Foundations, vol. 6. Electronic textbook (2023), http://softwarefoundations.cis.upenn.edu, version 2.0
6. Chen, A., Fathololumi, P., Koskinen, E., Pincus, J.: Veracity: Declarative multicore programming with commutativity. Proc. ACM Program. Lang. **6**(OOPSLA2) (oct 2022). https://doi.org/10.1145/3563349
7. Chen, A., Fathololumi, P., Nicola, M., Pincus, J., Brennan, T., Koskinen, E.: Better predicates and heuristics for improved commutativity synthesis. In: André, É., Sun, J. (eds.) Automated Technology for Verification and Analysis - 21st International Symposium, ATVA 2023, Singapore, October 24-27, 2023, Proceedings, Part II. Lecture Notes in Computer Science, vol. 14216, pp. 93–113. Springer (2023). https://doi.org/10.1007/978-3-031-45332-8_5
8. Clements, A.T., Kaashoek, M.F., Zeldovich, N., Morris, R.T., Kohler, E.: The scalable commutativity rule: Designing scalable software for multicore processors. In: Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles. pp. 1–17. SOSP '13, Association for Computing Machinery, New York, NY, USA (2013). https://doi.org/10.1145/2517349.2522712
9. Cortesi, A., Costantini, G., Ferrara, P.: A survey on product operators in abstract interpretation. Electronic Proceedings in Theoretical Computer Science **129**, 325–336 (sep 2013). https://doi.org/10.4204/eptcs.129.19
10. Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages. pp. 238–252. POPL '77, Association for Computing Machinery, New York, NY, USA (1977). https://doi.org/10.1145/512950.512973
11. Cousot, P., Cousot, R.: Abstract interpretation and application to logic programs. The Journal of Logic Programming **13**(2), 103–179 (1992). https://doi.org/10.1016/0743-1066(92)90030-7
12. Dickerson, T., Gazzillo, P., Herlihy, M., Koskinen, E.: Adding concurrency to smart contracts. In: Proceedings of the ACM Symposium on Principles of Distributed Computing. pp. 303–312. PODC '17, ACM, New York, NY, USA (2017). https://doi.org/10.1145/3087801.3087835

13. Eilers, M., Dardinier, T., Müller, P.: Commcsl: Proving information flow security for concurrent programs using abstract commutativity. Proc. ACM Program. Lang. **7**(PLDI), 1682–1707 (2023). https://doi.org/10.1145/3591289
14. Elmas, T., Qadeer, S., Tasiran, S.: A calculus of atomic actions. ACM SIGPLAN Notices **44**(1), 2–15 (2009)
15. Enea, C., Koskinen, E.: Scenario-based proofs for concurrent objects. Proc. ACM Program. Lang. (to appear) (OOPSLA2) (2024)
16. Farzan, A.: Commutativity in automated verification. In: LICS. pp. 1–7 (2023). https://doi.org/10.1109/LICS56636.2023.10175734
17. Farzan, A., Klumpp, D., Podelski, A.: Stratified commutativity in verification algorithms for concurrent programs. Proc. ACM Program. Lang. **7**(POPL), 1426–1453 (2023). https://doi.org/10.1145/3571242
18. Farzan, A., Klumpp, D., Podelski, A.: Commutativity simplifies proofs of parameterized programs. Proc. ACM Program. Lang. (POPL) (2024)
19. Farzan, A., Mathur, U.: Coarser equivalences for causal concurrency. Proc. ACM Program. Lang. **8**(POPL), 911–941 (2024). https://doi.org/10.1145/3632873
20. Flanagan, C., Freund, S.N.: The anchor verifier for blocking and non-blocking concurrent software. Proceedings of the ACM on Programming Languages **4**(OOPSLA), 1–29 (2020)
21. Kim, D., Rinard, M.C.: Verification of semantic commutativity conditions and inverse operations on linked data structures. In: Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 528–541. PLDI '11, Association for Computing Machinery, New York, NY, USA (2011). https://doi.org/10.1145/1993498.1993561
22. Koskinen, E., Bansal, K.: Decomposing data structure commutativity proofs with mn-differencing. In: International Conference on Verification, Model Checking, and Abstract Interpretation. pp. 81–103. Springer (2021)
23. Kragl, B., Qadeer, S.: The CIVL verifier. In: 2021 Formal Methods in Computer Aided Design (FMCAD). pp. 143–152. IEEE (2021)
24. Lette, D., Farzan, A.: Commutativity for concurrent program termination proofs. In: Enea, C., Lal, A. (eds.) Computer Aided Verification - 35th International Conference, CAV 2023, Paris, France, July 17-22, 2023, Proceedings, Part I. Lecture Notes in Computer Science, vol. 13964, pp. 109–131. Springer (2023). https://doi.org/10.1007/978-3-031-37706-8_6
25. Nagasamudram, R., Naumann, D.A.: Alignment completeness for relational hoare logics. In: 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021. pp. 1–13. IEEE (2021). https://doi.org/10.1109/LICS52264.2021.9470690
26. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: Proceedings of the 32nd ACM SIGPLAN-SIGACT symposium on Principles of programming languages. pp. 247–258 (2005)
27. Pincus, J.: Commutativity Reasoning for the Heap. Master's thesis, Stevens Institute of Technology (2022), https://www.proquest.com/docview/2681771819
28. Pîrlea, G., Kumar, A., Sergey, I.: Practical Smart Contract Sharding with Ownership and Commutativity Analysis, pp. 1327–1341. Association for Computing Machinery, New York, NY, USA (2021), https://doi.org/10.1145/3453483.3454112
29. Prabhu, P., Ghosh, S., Zhang, Y., Johnson, N.P., August, D.I.: Commutative set: A language extension for implicit parallel programming. In: Proceedings of the 32nd ACM SIGPLAN conference on Programming language design and implementation. pp. 1–11 (2011). https://doi.org/10.1145/1993316.1993500

30. Reynolds, J.: Separation logic: a logic for shared mutable data structures. In: Proceedings 17th Annual IEEE Symposium on Logic in Computer Science. pp. 55–74 (2002). https://doi.org/10.1109/LICS.2002.1029817
31. Rinard, M.C., Diniz, P.C.: Semantic foundations of commutativity analysis. In: Euro-Par '96 Parallel Processing, Second International Euro-Par Conference, Lyon, France, August 26-29, 1996, Proceedings, Volume I. pp. 414–423 (1996). https://doi.org/10.1007/3-540-61626-8_55
32. Shapiro, M., Preguiça, N., Baquero, C., Zawirski, M.: A comprehensive study of convergent and commutative replicated data types. Ph.D. thesis, Inria–Centre Paris-Rocquencourt; INRIA (2011)
33. Sims, E.J.: An abstract domain for separation logic formulae. In: Proceedings of the 1st International Workshop on Emerging Applications of Abstract Interpretation (EAAI06). pp. 133–148. ENTCS, Vienna, Austria (2006)
34. Weihl, W.E.: Data-dependent concurrency control and recovery (extended abstract). In: Proceedings of the second annual ACM symposium on Principles of distributed computing (PODC'83). pp. 63–75. ACM Press, New York, NY, USA (1983). https://doi.org/10.1145/800221.806710
35. Yang, H.: Relational separation logic. Theor. Comput. Sci. **375**(1-3), 308–334 (2007). https://doi.org/10.1016/J.TCS.2006.12.036

Appendix

To share a specific example of isomorphism, as well as an alternative way to construct abstract domains, we return to our two-set example (Ex. 2), again containing values of type $T \subseteq \mathbb{V}^*$ at addresses p and q. Defining π in that example involved explicitly tracking different versions of the structure, taking care that the concrete order of values did not affect the abstract mapping.

Here, we take a different, perhaps more intuitive approach. We will explicitly construct an equivalence relation on heaps, and derive X and π from there. Here is said equivalence relation, using a heap predicate similar to what Ex. 2 used:

$$h \sim h' \coloneqq \exists u, v.\ \wedge \begin{cases} \mathsf{setq}^+_{p,q}\big((u,v),h\big) \\ \mathsf{setq}^+_{p,q}\big((u,v),h'\big) \vee \mathsf{setq}^+_{p,q}\big((v,u),h'\big) \end{cases}$$

$$\mathsf{setq}_{p,q}\big((u,v),\cdot\big) \coloneqq p \mapsto u * q \mapsto v * (u, v \in T + \mathsf{null}) * (u = v \implies u = \mathsf{null})$$

Our abstract values will be the equivalence classes of heaps containing two-sets, and the projection function will extend the canonical projection. Specifically, $\mathsf{Setq}^T_{p,q} = \langle \mathbb{H}(\mathsf{setq}^+_{p,q})/\!\sim, \pi \rangle$, where $\pi(h) = [h]_\sim$ if $\mathsf{setq}^+_{p,q}(h)$, else ✗.

Indeed, we show in Coq that $\mathsf{Set}^T_{p,q} \cong \mathsf{Setq}^T_{p,q}$. Comparing the two constructions, we see how this eq. relation approach implicitly and unambiguously captures each possible set cardinality, as well as concrete order irrelevance. On the other hand, explicit representations of abstract values are often desirable for human comprehension (e.g. sets of size 0–2), but equivalence classes are more opaque in their meaning. This complicates the construction and analysis of abstract programs. For this reason, we only use $\mathsf{Set}^T_{p,q}$ when working with two-set programs in §7.2.