

Mode Crossing

BENJAMIN PETERS, MPI-SWS, Germany

JULES JACOBS, Jane Street, USA

DIANA KALINICHENKO, Jane Street, USA

LIAM STEVENSON, Jane Street, USA

ASPEN SMITH, Jane Street, USA

DEREK DREYER, MPI-SWS, Germany

RICHARD A. EISENBERG, Jane Street, USA

OxCaml extends the OCaml type system with support for safe low-level systems programming via *modes*. For example, OxCaml’s modal *portability* and *contention* axes ensure that concurrent OxCaml programs have no data races. In practice, however, modal tracking can reject programs that are obviously safe, such as when the data shared across threads is immutable. To remedy this problem, we introduce *mode crossing*—the ability to automatically strengthen modes (e.g., from **nonportable** to **portable**) for values of certain types. Mode crossing significantly reduces the annotation burden associated with modal types. To support mode crossing in the presence of abstract type specifications, we further introduce a new type system feature, *modal kinds*.

We present a type-theoretic account of modal kinds, interpret them as monotone functions on a lattice of modes, and extend this interpretation to recursive and abstract types. We verify soundness of the modal kind system in Rocq on top of Iris. We design an inference procedure that reduces kind checking and subsumption to constraints solved by a dedicated lattice solver. Our design is implemented in the OxCaml compiler and deployed in a large industrial codebase, demonstrating practical usability.

CCS Concepts: • **Computing methodologies** → **Concurrent programming languages**; • **Theory of computation** → **Type theory**.

Additional Key Words and Phrases: Modal types, modal kinds, OxCaml, Iris, Rocq

ACM Reference Format:

Benjamin Peters, Jules Jacobs, Diana Kalinichenko, Liam Stevenson, Aspen Smith, Derek Dreyer, and Richard A. Eisenberg. 2026. Mode Crossing. *Proc. ACM Program. Lang.* 10, ICFP, Article 150 (August 2026), 54 pages.

1 Introduction

OxCaml is a backward-compatible extension of OCaml for safe concurrent and systems programming [12, 13, 18]. It aims to provide strong safety guarantees (notably data-race freedom and safe stack allocation) while preserving OCaml’s ergonomics and ecosystem compatibility. As the name OxCaml (“Oxidized Caml”) suggests, the language is inspired by Rust and its support for safe, low-level, performance-oriented programming, but OxCaml makes some very different design choices. First, whereas Rust supports manual memory management and ensures safety by prohibiting aliased mutable state by default, OxCaml (building on OCaml) is garbage-collected and safely allows arbitrarily aliased mutable state by default. Second, whereas Rust starts from an

Authors’ Contact Information: Benjamin Peters, MPI-SWS, Saarland Informatics Campus, Germany, bpeters@mpi-sws.org; Jules Jacobs, Jane Street, New York, USA, julesjacobs@gmail.com; Diana Kalinichenko, Jane Street, New York, USA, dkalinichenko@janestreet.com; Liam Stevenson, Jane Street, New York, USA, lstevenson@janestreet.com; Aspen Smith, Jane Street, New York, USA, aspsmith@janestreet.com; Derek Dreyer, MPI-SWS, Saarland Informatics Campus, Germany, dreyer@mpi-sws.org; Richard A. Eisenberg, Jane Street, New York, USA, reisenberg@janestreet.com.



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

© 2026 Copyright held by the owner/author(s).

ACM 2475-1421/2026/8-ART150

<https://doi.org/>

ownership-based type system (with rigid control over aliasing) and makes it more flexible through the use of *lifetimes* and *borrowing*, OCaml starts from a very flexible language and supports more careful control over aliasing through the use of *modal types*.

This paper addresses a practical usability problem in the previously presented OCaml design. We describe a language extension—already deployed in the OCaml compiler—that supports what we call *mode crossing* and *modal kinds*. To illustrate the problem, consider a simple OCaml fork-join example: a recursive divide-and-conquer summation that splits an immutable input array in half, processes both halves in parallel, and maps each integer with f .

```

50 let rec sum_map (f : int → int) @ portable (xs : int iarray) =
51   match Iarray.length xs with
52   | 0 → 0
53   | 1 → f xs.(0)
54   | _ → let s_left, s_right =
55         fork_join2
56         (fun () → sum_map f (Iarray.first_half_slice xs))
57         (fun () → sum_map f (Iarray.second_half_slice xs))
58         in
59     s_left + s_right

```

OCaml ensures race freedom by requiring callbacks passed to `fork_join2` to be @ `portable`:¹

```

60 val fork_join2 : (unit → 'a) @ portable → (unit → 'b) @ portable → 'a * 'b

```

Functions are portable if they can be executed in parallel without introducing data races: they neither access unsynchronized shared mutable state themselves nor call functions that do. (In this example, we assume that top-level functions, such as `fork_join2` and the functions in `Iarray`, are themselves portable.) Accordingly, `sum_map` requires $f : (int \rightarrow int) @ portable$ in order to allow f to be called in child threads. Under that caller-side precondition, the remaining data flow involves only immutable arrays and integers, so `sum_map` is race-free.

The problem. Georges et al. [12] presented a semantic foundation for OCaml’s modal type system. Under their typing rules, because xs in the above program flows to other threads via `fork_join2`, the program type-checks only if the xs argument of `sum_map` is also @ `portable`:

```

61 let rec sum_map (f : (int → int) @ portable) (xs : int iarray @ portable) = ...

```

redundant requirement in baseline OCaml

One would similarly have to establish portability-annotated types for the auxiliary functions called by `sum_map`:

```

62 val first_half_slice, second_half_slice : int iarray @ portable → int iarray @ portable

```

For immutable `int iarray`, however, these additional annotations impose a redundant requirement because *every* value of type `int iarray` is portable! Unfortunately, the modal type system in Georges et al. [12] does not exploit that type-specific property, as it treats modes and types *independently*: its mode-checking rules must be uniformly sound for *all* types and are therefore pessimistic, requiring both f and xs to be portable.

In our experience with a large industrial OCaml codebase at Jane Street (see §5), this is a major usability issue: redundant annotations pollute type signatures and burden callers. For example, in practice, although all `int iarray` values are @ `portable`, it is surprisingly non-trivial to convince

¹OCaml has multiple modal axes, including portability, contention, affinity, uniqueness, locality, etc. We use portability as the running example axis in this introduction and generalize later to include contention also.

the type-checker of this fact. Doing so requires burdensome annotations, nested modalities [12], and mode polymorphism across the library ecosystem used to produce those values. It also makes it much harder for code that uses modes to interoperate with code that does not.

Mode crossing and modal kinds. In this paper, we show how to solve this usability problem in OxCaml’s modal type system using what we call *mode crossing*. Mode crossing lets the type checker strengthen the mode of a value based on its type, thereby eliminating redundant annotations such as `@ portable` on `int iarray` while still soundly tracking mode-sensitive types. For instance, `int iarray` values can always be treated as `portable`, but `int → int` values generally cannot; in this case, we say that the type `int iarray` “crosses” portability. We track the mode-crossing properties of types using *modal kinds*.² For example, we would say that `int iarray` has modal kind `value mod portable`, whereas `int → int` has only modal kind `value`.³

To reiterate, mode crossing and modal kinds *reduce the annotation burden of a modal type system* by moving annotations where they belong. Instead of annotating every function argument and result (including mode-polymorphism annotations on type-polymorphic functions), we place the burden on abstract type specifications. For concrete types, we infer the most accurate modal kind. This reduction is crucial in practice: without mode crossing, adopting OxCaml in a large industrial codebase like Jane Street’s would have been infeasible.

Although the basic idea of mode crossing is simple, its implementation and metatheory in terms of modal kinds is surprisingly subtle. In particular, we identify the following design highlights:

With-bounds (§2.5) Modal kinds of parameterized types should be computed compositionally: the kind of `int iarray` should follow from the kinds of `int` and `iarray`, which requires a kind language for type constructors. We express constructor dependence via `with`-bounds (e.g., `'a iarray : value mod portable with 'a`) and model type constructors semantically as functions between modal kinds.

Recursive types (§2.6) Recursive type constructors, especially those containing non-uniform recursion, cannot be handled by naive unfolding, as that approach can fail to terminate. We instead give recursive kinds a principled lattice-theoretic least fixpoint semantics to support sound and terminating checking.

Modalities (§2.7) Modalities create axis-specific dependence: a type constructor may be insensitive to an argument on one axis while still depending on it on another. We capture this by allowing modality-qualified with-bounds such as `with 'a @@ portable`, so kinding tracks how modality application transforms crossing behavior.

Abstract types (§2.8) Kinding must remain sound across module abstraction boundaries, where concrete representations are hidden. We therefore track dependence on abstract types and abstract constructors (e.g., `with t` and `with 'a t`), and use this information both for type checking and for sub-kinding obligations that arise in signature inclusion and substitution.

Lattice solver (§2.9) The compiler must solve kind inclusion checks, including kinds that are fixpoints of recursive equations, under hypotheses induced by abstract types with with-bounds. We reduce this to a decision problem in lattice functions and solve it with a dedicated solver for practical kind inference and sub-kinding checks.

²An alternative approach would be to track such properties in the way that Rust tracks modal properties of types (e.g., portability), using so-called *marker traits* like `Send` and `Sync`. We employ kinds, in part so that we can build on OxCaml’s existing kind infrastructure for unboxed types (OxCaml does not support traits or type classes). Furthermore, Rust’s marker traits do not address a number of the design challenges described in this section (see §6 for details).

³The actual kinds are more complex due to the presence of other modal axes.

Contributions. We make the following contributions:

Design (§2) We motivate *mode crossing* and introduce *modal kinds*, giving a detailed account of modal kinds for base types, modalities, recursive types, and abstract types, together with sub-kinding and inclusion principles.

Soundness (§3) We present SMOki (a System of Modal Kinds), a core modal type system with an explicit kind calculus that formalizes kinding constraints and their interaction with typing. We prove memory safety and data-race freedom for SMOki (Thm. 6, in §3.6) using a logical relation mechanized in Rocq with Iris.

Inference (§4) We describe an inference procedure that reduces kind checking and sub-kinding to lattice constraints solved by a dedicated solver, and incorporate it into the OCaml compiler.

Implementation The ideas in this paper are all implemented in open-source and are used in production to compile Jane Street’s codebase.

We conclude our paper with some discussion of how mode crossing has manifested in practice (§5) and with a review of related work (§6).

Non-goals and limitations. Our core type system SMOki does not handle all features of OCaml. In particular, it does not formalize all OCaml modal axes, although it does handle all axes considered by Georges et al. [12]. It also does not directly model a full-blown module system, although it supports typechecking in the presence of abstract type specifications. Lastly, SMOki does not handle GADTs, for which the construction of a semantic model in Iris is an independently challenging problem [25]. In contrast, our OCaml implementation handles all modal axes of OCaml, supports the full-blown OCaml module system, and deals with GADTs conservatively (inferring a sound modal kind but not necessarily a principal one).

2 Key Ideas

This section introduces *modal kinds*: kinds that describe which modal axes a type can safely ignore (*mode crossing*). We first recap portability and contention, then show by example that plain mode checking is too conservative for immutable data. Next, we interpret kinds as lattice-valued functions, use that view for recursion and modalities, and end with the normal-form view used by the implementation.

2.1 Background: The Portability and Contention Modes

The central idea behind OCaml’s data-race freedom guarantee concerns control of mutable state. Unless we use synchronization primitives (which are outside the scope of this paper), we must never mutate state using a reference that has passed from one thread to another. We thus assign all values that pass from one thread to another the **contended** mode, and we prevent reading or writing mutable fields from **contended** values. This might seem sufficient, except that a closure could *capture* an **uncontended** reference, then be transferred to another thread and executed there, defeating this protection. We thus say that a function capturing **uncontended** references is **nonportable**; a function capturing only **contended** references is **portable**. Functions that are **portable** are allowed to be passed to other threads; **nonportable** functions may not be. (A **portable** function must additionally capture only other **portable** functions.) For example, the function `fun () → r := !r + 1` (for some captured `r : int ref @ uncontended`) is **nonportable**, because it needs **uncontended** access to `r` to mutate it. On the other hand, `fun r → r := !r + 1` is fully **portable**: it captures only **portable** operations, including addition and reference reads and writes. (Intuitively, the notion of portability is similar to that of Rust’s `Send`, in that it determines whether values can move between threads or not.)




Fig. 1. The portability and contention modes (§2.1). *Sub-modding* follows solid arrows, and *sub-kinding* (§2.5) follows dashed arrows. Bolded modes are default, and underlined modes have a modality (§2.7) associated with them.

```

208 module Stack : sig @@ portable
209   type 'a t
210   val empty : 'a t
211   val push : 'a → 'a t → 'a t
212   val pop : 'a t → ('a * 'a t) option
213 end = struct
214   type 'a t = | Nil | Cons of 'a * 'a t
215   let empty = Nil
216   let push x st = Cons (x, st)
217   let pop st = match st with ...
218 end
219 val print_int_stack : (int Stack.t → unit)
220   @@ portable

```

Fig. 2. An API for an immutable stack.

```

int : value mod portable contended
'a list : value mod portable contended with 'a
'a ref : value mod portable with 'a
'a * 'b : value mod portable contended
         with 'a with 'b
'a → 'b : value mod contended
'a @@ m : value mod portable contended
         with 'a @@ m

```

Fig. 3. Some examples of types and their modal kinds.

We have now met two *modal axes*: **portable/ nonportable** on the *portability* axis, and **contended/ uncontended** on the *contention* axis. These two axes are summarized in Fig. 1. Each axis also has a third element: a **shared** value allows read-only access to mutable fields, and a **shareable** function may capture only read-only state (**shared** or **contended**) and read-only functions (**shareable** or **portable**).

All values in OCaml are assigned a mode from each of the modal axes. Accordingly, the function type $\text{arg_ty} @ \text{arg_mode} \rightarrow \text{res_ty} @ \text{res_mode}$ specifies the expected input and output modes of the function, separate from the mode of the function value itself. Omitted modal axes are set to their **default** value, chosen for backward compatibility with OCaml.

OCaml also supports *sub-modding*. Each axis arranges its modes along a lattice, and we can always consider a value with a stronger (lower) mode to instead have a weaker (higher) mode. For example, if we have an $\text{int} \rightarrow \text{string}$ function at **portable**, we can use that in a context expecting a **nonportable** value, because **portable** < **nonportable**. With sub-modding, we can refine our capture rules slightly: when a **portable** function captures an **uncontended** value, that value can be downgraded (that is, upcast) to **contended** before capture rather than becoming inaccessible.

Finally, modes in OCaml are by default *deep*. That is, if a record is at mode m , then projecting a field from that record produces a value of mode m as well.

2.2 Motivation for Mode Crossing

Portability and contention are expressive enough to type-check interesting programs, especially in conjunction with concurrency APIs using other OCaml modes [12]. However, limitations

246 appear quickly in real-world use. Consider an interface for immutable stacks (Fig. 2) and a function
 247 `print_int_stack` defined separately from the module. The `@@ portable` annotation on the signa-
 248 ture states that all values in the module are **portable**. Similarly, the `@@ portable` annotation on
 249 `print_int_stack` states that it, too, is **portable**.

250 Intuitively, the implementation’s immutability guarantee makes it safe to use an `int Stack.t`
 251 from another thread without risking data races. OCaml, however, rejects the following program:

```
252 let st = Stack.(empty |> push 4 |> push 2) in
253   fork_join2 (fun _ → print_int_stack st) (fun _ → ...)
254           ^^
```

255 **Error: The value st is nonportable but is expected to be portable**

256 We can see why this is so by examining the types and modes. The `Stack` module signature does
 257 not actually expose the fact that `int Stack.t` is safe to access concurrently. Specifically, `st` is
 258 **nonportable**: the return mode of `Stack.push` is just the default mode **nonportable uncontended**.
 259 (Note that the function `Stack.push` itself is **portable**, but not its inputs and output.) Hence, `st`
 260 cannot be captured by the **portable** closure passed to `fork_join2`.

261 How do we generalize the `Stack` module while exposing the thread-safety guarantees that come
 262 from immutability? After all, an `int Stack.t` has no functions or mutability—there is no way it
 263 could cause trouble when used in different threads.

265 2.3 Mode Crossing and Kinds

266 We exploit the insight that some types are naturally agnostic to some modal axes. We say that a
 267 type *crosses* (or *mode-crosses*) an axis when all of its values have no interaction with that axis. In
 268 our case, `int Stack.t` crosses both portability and contention. This holds because `int` crosses both
 269 axes and because `Stack.t` preserves this property. We can encode this information in the stack
 270 signature in Fig. 2 by adding a kind annotation to the second line:

```
272 type 'a t : value mod portable contended with 'a
```

273 This specifies the *kind* of abstract `'a t`, restricting admissible instantiations while giving users of
 274 `'a t` additional capabilities (the ability to mode-cross). Such a kind has three components:

- 275 (1) The layout **value**. In OCaml, layouts support unboxed representations [4, 6], inspired by
 276 [5] and [7]. Layouts are orthogonal to modal kinds and beyond this paper’s scope.
- 277 (2) The modal bound **mod portable contended** states that values cross portability and con-
 278 tention. (Why do we specify modes instead of specifying whole axes? Read on to §2.7.)
- 279 (3) The list of with-bounds **with 'a** states that the type can only cross axes that are also crossed
 280 by `'a`. With-bounds can mention arbitrary types.

281 A kind annotation on an abstract type describes a bound on admissible instantiations of that type.
 282 Any instantiation must have a kind at least as strong as this bound. The default kind annotation, if
 283 none is given, is just **value**. It expresses that the type does not cross any modal axes. Adding the
 284 above kind (**value mod portable contended with 'a**) to the signature of the `Stack` module allows
 285 OCaml to compile the example from before without any other changes.

287 2.4 Inferring Modal Kinds for Concrete Type Definitions

288 OCaml automatically infers modal kinds of concrete types. Computing these kinds is sometimes
 289 challenging. In this section, we describe the rules for inferring kinds for some basic types and
 290 present an example where inferring the kind requires more care. We list examples of kinds in Fig. 3.

291 The kind of a record with only immutable fields is easy to infer:
 292
 293
 294

```

295 type t = { a : string ; b : int }
296 (* Inferred kind: value mod portable contended with string with int,
297    equivalent to value mod portable contended *)

```

298 The inferred kind includes two with-bounds, because `string` and `int` are contained in `t`. In many
 299 cases, the with-bounds of a type’s kind are simple: they list all of the (potential) components of
 300 a type. However, users need not reason about with-bounds to understand `t`, because this kind
 301 simplifies to **value mod portable contended**: both `string` and `int` cross portability and contention.

302 What about the kind of a record with a mutable field?

```

303 type 'a ref = { mutable contents : 'a } (* Inferred kind: value mod portable with 'a *)
304

```

305 A record with a mutable field can never cross contention, because contention is used to track
 306 whether it can be accessed/mutated or not (but it may cross portability, if `'a` does so). The modal
 307 bound on its kind is thus **value mod portable**, and because it contains an `'a`, we need to add a
 308 with-bound **with 'a**.

309 Functions, on the other hand, carry *computation* and no value components: their kind has no
 310 with-bounds. Functions never cross portability (because that’s how we track their thread-safety),
 311 but they do cross contention. So, their overall kind is always **value mod contended**.

312 We have not yet talked about the kinds of composite types like `'a ref ref`. The syntax of
 313 kinds does not make it obvious how to compose and simplify them. These remaining gaps in our
 314 understanding of kinds become more pronounced when we consider recursive types, like this one:

```

315 type 'a nested_refs =
316   | Nil of 'a
317   | Cons of 'a ref nested_refs
318

```

319 This recursion is *non-uniform*: each recursive step adds another `ref` to the argument. Computing
 320 its kind requires reasoning about unbounded nesting.

322 2.5 Key Idea I: Kinds as Functions

323 Having seen several examples of kinds in action, we now formalize their meaning. So far we have
 324 explained kinds informally as a way to express the mode-crossing behavior of types. Our first
 325 key idea is to interpret the kinds of type constructors as functions that map the modes crossed by
 326 their type arguments to the modes crossed by their results. This yields a principled way to handle
 327 recursive types and later check *sub-kinding*.

328 Base types like `int` or `string ref` have *modal base kinds*, which specify their mode-crossing
 329 behavior along each axis—for example, **mod portable contended**. Modal base kinds do not have
 330 with-bounds. The kinds of type constructors are more interesting. For instance, recall the kind of
 331 `'a ref` from §2.4:

```

332 type 'a ref : value mod portable with 'a
333

```

334 The kind of `'a ref` depends only on the kind of `'a`, and not on its concrete instantiation. This means
 335 that we can interpret the kinds of type constructors as functions that take the modal base kind of
 336 `'a` and compute the modal base kind of `'a ref`. What should this function be, given the kind of `'a`?
 337 Specifically, what is the meaning of the with-bound **with 'a**?

338 A with-bound **with t** *trims* the possible mode-crossing behaviors; It can only reduce the number
 339 of axes that are crossed. We intend to set this up as a join in a lattice, but we first need a more precise
 340 representation of modal base kinds. We define modal base kinds as elements of the set of modes,
 341 but with a different (partial) order on this set compared to the sub-moding order: We flip the order
 342 on the contention axis, as shown in Fig. 1. We call this the *sub-kinding* order. (Overall, a “larger”
 343

kind comprises more types and permits less mode-crossing behavior.) The bottom modes with respect to sub-kinding (**portable** and **contended**) represent that a type can cross the respective axis, like modal bound annotations **mod portable contended** (§2.3). Conversely, top modes (**nonportable** and **uncontended**) represent that the respective axis is not crossed. The reason for this choice is that mode crossing is intimately connected to modalities (§2.7), and their behavior determines which mode sits at \perp .

Let us get back to the kind of 'a ref. The modal-bound **value mod portable** (represented by (**portable, uncontended**)) states the “floor” of mode-crossing behavior, that is, that 'a ref could potentially cross portability, but never crosses contention. However, its mode-crossing behavior also depends on 'a: the with-bound **with** 'a potentially *removes* per-axis mode-crossing behavior, based on whether 'a crosses **portable** or not. So **with** 'a means taking the least upper bound of the rest of the kind with the kind of 'a. This is precisely the join operation on the sub-kinding lattice. Given the base kind of 'a, call it $\hat{\alpha}$, we get the following so-called *kind function* for 'a ref:

$$f(\hat{\alpha}) = (\mathbf{portable, uncontended}) \sqcup \hat{\alpha}$$

2.6 Key Idea II: Fixpoints of Kind Functions for Recursive Types

Recall the 'a nested_refs type from §2.4. It is not obvious what its mode-crossing behavior is. We use kind functions as a tool to analyze the mode-crossing behavior of such complicated recursive types: we set up a *recursive equation* that must hold for its kind function. Using the type’s definition and the kind function of 'a ref we just computed, we can set up the following equation on the kind function f of 'a nested_refs:

$$f(\hat{\alpha}) = (\mathbf{portable, contended}) \sqcup \hat{\alpha} \sqcup f((\mathbf{portable, uncontended}) \sqcup \hat{\alpha})$$

Note that we use composition of f with the kind function of 'a ref to interpret the with-bound **with** 'a ref nested_refs that mentions a composed type. There are many monotone solutions for f that satisfy the equation. But since we are working in a finite and hence complete lattice, we know that there always exists a least solution (a least fixpoint) to such equations.⁴ This least solution corresponds to the strongest kind function we can choose. In this case, the solution is:

$$f(\hat{\alpha}) = (\mathbf{portable, uncontended}) \sqcup \hat{\alpha}$$

It corresponds to the kind **value mod portable with** 'a.

2.7 Key Idea III: Invariant Modalities and Mode Crossing

We have discussed how several of OxCaml’s features (records, mutable state, and recursive types) interact with modal kinds. So far, we have ignored one important feature of OxCaml: modalities. Modalities are used to adjust the underlying mode of selected record fields. Mode crossing and modalities are closely related: A type crosses a modal axis if and only if it is *invariant* under a certain modality. In this section, we connect mode crossing and modalities, use that connection to justify our sub-kinding lattice from the previous section, and then describe the mode-crossing behavior of modalities themselves.

Consider the following record:

```
type 'a port = { x : 'a @@ portable }
```

Its only field is equipped with a *modality annotation* @@ **portable**, expressing that the value in the x field is always **portable**, regardless of the mode of the record. (Note that the @@ **portable** annotation itself is not a type former, just an annotation on a record field.) We refer to 'a port as

⁴To be pedantic: We are forming a least fixpoint on the lattice of functions from modal base kinds to modal base kinds, which is complete if the sub-kinding lattice itself is.

393 the **portable** *modality* because it wraps its only field with @@ **portable**. OCaml also supports the
 394 **contended** modality, among others.

395 We can observe the following connection between the **portable** modality and mode crossing:
 396 A type t crosses portability iff it is *invariant* under the **portable** modality, i.e., if t and t port are
 397 morally “equivalent”. The proof is simple: if t crosses portability, we can define trivial coercions
 398 between t and t port; conversely, these coercions suffice to justify mode crossing. Analogously,
 399 a type crosses contention iff it is invariant under the **contended** modality. Indeed, the modes in a
 400 modal bound **mod portable contended** do not refer to the portability and contention axes, but to
 401 the **portable** and **contended** modalities. They express that a type is invariant under the **portable**
 402 and **contended** modalities, and so crosses portability and contention. Hence we want **portable** and
 403 **contended** to be the bottom values in the sub-kinding order.

404 This is why we flipped the order of the contention axis in the sub-kinding lattice relative to the
 405 sub-modging lattice in §2.5. Take another look at Fig. 1: modalities behave differently on portability
 406 versus contention. The **portable** and **contended** modalities correspond to the bottom and top modes
 407 of their axes with respect to sub-modging, respectively. A **nonportable** or an **uncontended** modality
 408 would be unsound: they could be used to store **nonportable** (**uncontended**) data in a **portable**
 409 (**contended**) value and would allow it to be smuggled to another thread.

410 OCaml also supports **shareable** and **shared** modalities that correspond to *middle modes* on their
 411 respective axes. Reading a record field annotated with the **shared** modality yields its value at **shared**
 412 mode, except when the surrounding record is at **contended**: then the field value is only accessible
 413 at **contended** mode (cf. [12, §4.3]). We say that a type crosses **shared** if it is invariant under the
 414 **shared** modality. In this case, values of that type can move freely between the **uncontended** and
 415 **shared** modes, but not between those and **contended**. The **shareable** modality behaves similarly: it
 416 collapses the **shareable** and **nonportable** modes.

417 We have discussed the interactions between mode crossing and modalities so far, but not the
 418 mode-crossing behavior of modalities themselves. The type $'a$ port always crosses portability, and
 419 only crosses contention if $'a$ also does so. Unfortunately, we cannot express this mode-crossing
 420 behavior using the syntax of modal kinds we have introduced so far: when interpreting with-bounds
 421 as joins on the sub-kinding lattice, we cannot “selectively apply” the mode-crossing behavior on
 422 some axes. From now on, we allow modality annotations to show up in with-bounds to express
 423 this behavior. The inferred kind of $'a$ port is:

424 `type 'a port : value mod portable contended with 'a @@ portable`

425 Semantically, we can read **with** as a join (combining restrictions) and @@ as a meet (weakening
 426 restrictions):

$$427 \quad f(\hat{\alpha}) = (\text{portable}, \text{contended}) \sqcup (\hat{\alpha} \sqcap (\text{portable}, \text{uncontended}))$$

428 We interpret @@ **portable** as taking the meet with $(\text{portable}, \top) = (\text{portable}, \text{uncontended})$ as to
 429 only affect the portability axis.

432 2.8 Key Idea IV: Abstract Kinds for Abstract Types

433 In our kind-function account of OCaml’s kind system, we have so far handled only types whose
 434 definitions are visible. But as we have seen in our introductory example of kinds in §2.3, kinds also
 435 interact with OCaml’s module system. Using modules and functors, we can mention so-called
 436 *abstract types* in with-bounds. An abstract type is a type declaration without a definition, either
 437 because it has not been instantiated or because its definition has been hidden.

438 Consider the following example of a map functor with mode crossing annotations:

439
 440
 441

```

442 module type OrderedType = sig
443   type t
444   val compare : (t → t → int) @@ portable
445 end
446 module Map (Ord : OrderedType) : sig @@ portable
447   type key = Ord.t
448   type 'a t : value mod portable contended with 'a with Ord.t
449   ...
450 end = struct ... end

```

Given a type with an associated compare function, the Map functor generates an implementation of an immutable map. The compare function is required to be **portable** so that it can be called concurrently from multiple threads. The functor exposes a type constructor `'a t` whose mode-crossing behavior is constrained by `'a` and `Ord.t`. Naturally, we want `'a t` to have as much mode-crossing behavior as `Ord.t` does. However, this means that the with-bound `with Ord.t` depends not on the declared kind of `Ord.t`—the default kind **value**—but on the as-yet-unknown kind of the *instantiation* of `Ord.t`. We need to treat `with Ord.t` as referring to an *abstract kind* that is instantiated along with `Ord`.

Checking sub-kinding in the presence of abstract kinds is challenging. Consider this:

```

459 type 'a t (* some abstract type in our context *)
460 module Foo : sig
461   type 'a t2 : value mod portable contended with 'a t
462 end = struct
463   type 'a t2 = { x : 'a t t }
464 end

```

The compiler has to determine whether **value mod portable contended with 'a t t** is a sub-kind of **value mod portable contended with 'a t**. It is. The example type-checks successfully. But why? To answer this question, we need a more precise description of how kinds can behave.

2.9 Key Idea V: Kinds as Coefficients for Inference

We now shift from semantics to implementation: we discuss how to represent kind functions by coefficients and reduce sub-kinding to coefficient constraints. Previously, we claimed that **value mod portable contended with 'a t t** is a sub-kind of **value mod portable contended with 'a t**. How do we show this? Let us use f to refer to the abstract kind function of `'a t`. We need to check:

$$\forall \hat{\alpha}. (\text{portable}, \text{contended}) \sqcup f(f(\hat{\alpha})) \leq (\text{portable}, \text{contended}) \sqcup f(\hat{\alpha})$$

Kind functions admit a normal form: for unary f , there are constants c_0, c_1 such that⁵

$$f(\hat{\alpha}) = c_0 \sqcup (c_1 \sqcap \hat{\alpha}).$$

This fact follows from the structure of kind functions: They are compositions of functions that consist of constants, arbitrary joins, meets with constants, and least fixpoints (of recursive types). All of these operations preserve this structure, because in the sub-kinding lattice, meets and joins are *distributive*. Finally, by some basic computation on lattice terms, we can justify the claim:⁶

$$f(f(\hat{\alpha})) = c_0 \sqcup (c_1 \sqcap (c_0 \sqcup (c_1 \sqcap \hat{\alpha}))) = c_0 \sqcup (c_1 \sqcap c_0) \sqcup (c_1 \sqcap c_1 \sqcap \hat{\alpha}) = c_0 \sqcup (c_1 \sqcap \hat{\alpha}) = f(\hat{\alpha})$$

For n -ary constructors, the normal form generalizes to:

$$f(\hat{\alpha}_1, \dots, \hat{\alpha}_n) = c_0 \sqcup (c_1 \sqcap \hat{\alpha}_1) \sqcup \dots \sqcup (c_n \sqcap \hat{\alpha}_n)$$

⁵Our actual normal form theorem(s) require additional properties c_0 and c_1 to ensure uniqueness.

⁶We actually show something stronger: the with-bounds `with 'a t` and `with 'a t t` are equivalent for all unary type constructors `'a t`.

Note, however, that these constants may themselves depend on the kinds of other types. Specifically, these c_i can be expressed using constants, joins, and (arbitrary) meets of coefficients of the kind functions of other ambient types (either concrete or abstract).

Generally, a sub-kinding query asks whether one kind function f is less than another g for all inputs. How would we go about checking this using coefficients? For simplicity, let us consider unary functions f and g again: let f be represented by coefficients c_0, c_1 , and g be represented by coefficients d_0, d_1 . We want to verify the following property:

$$\forall \hat{\alpha}. c_0 \sqcup (c_1 \sqcap \hat{\alpha}) \leq d_0 \sqcup (d_1 \sqcap \hat{\alpha})$$

How do we check this? Using $x \leq y \iff x \sqcup y = y$ we can rewrite it as:

$$\forall \hat{\alpha}. (c_0 \sqcup d_0) \sqcup ((c_1 \sqcup d_1) \sqcap \hat{\alpha}) = d_0 \sqcup (d_1 \sqcap \hat{\alpha})$$

We compute a normal form of both sides, and check for syntactic equality of the coefficient formulas!

This is promising news for our sub-kinding solver (§4): it can manipulate kind functions in terms of coefficients. It represents the kinds of concrete types using *equalities* $c_i = \dots$ on their coefficients computed from type structure, and represents the kinds of abstract types using *inequalities* $c_i \leq \dots$ on their coefficients computed from their modal kind bounds. The right-hand sides of these (in-)equalities may mention constants and arbitrary meets and joins of other coefficients. Similarly, the solver reduces sub-kinding queries themselves to inequalities on coefficients only.

To solve such inequalities on coefficients in the presence of constraints on coefficients, the sub-kinding solver iteratively *inlines* more and more constraints into the inequalities to be checked, while simultaneously tying recursive knots that show up on the way. It continues until no constraints are left, at which point subsumption is easy to check. This procedure is explained by example in §4.

3 Type System

This section presents *SMoKi*, a pen-and-paper type system for the language from §2. We describe the kind calculus in §3.2, recall OxCaml-style modal type systems in §3.3, and instantiate the calculus to model *SMoKi* in §3.4. In §3.5 we justify module reasoning via substitution, and in §3.6 we sketch soundness. We summarize definitions from this section in Fig. 4.

All orderings and lattice operations in this section are wrt. sub-kinding, not sub-moding.

3.1 Preliminaries

We recall some standard definitions and facts about finite lattices. A finite lattice L is a set with a partial order \leq and operations meet \sqcap and join \sqcup . The element $\ell_1 \sqcap \ell_2$ is the greatest element ℓ_3 such that $\ell_3 \leq \ell_1$ and $\ell_3 \leq \ell_2$. Join is dual to meet, working up the lattice instead of down. Our lattices are always *bounded*, i.e. they have distinguished top \top and bottom \perp elements. On finite lattices, all monotone functions have least fixpoints.

Recall the definition of a Bi-Heyting algebra [e.g., 2, 8]:

Definition 1. A *Bi-Heyting algebra* $(L, \sqcup, \sqcap, \top, \perp, \leq, /, \backslash)$ is a bounded lattice equipped with additional binary operations $/$ and \backslash called *implication* and *subtraction*, respectively, such that

$$a \leq b / c \iff a \sqcap b \leq c \qquad a \backslash b \leq c \iff a \leq b \sqcup c.$$

Simple examples of Bi-Heyting algebras are lattices based on *bounded linear orders*. In that case, the implication and subtraction operations behave as follows:

$$a / b := \begin{cases} \top & \text{if } a \leq b \\ b & \text{if } a > b \end{cases} \qquad a \backslash b := \begin{cases} \perp & \text{if } a \leq b \\ a & \text{if } a > b \end{cases}$$

Products of Bi-Heyting algebras are again Bi-Heyting algebras under pointwise lifting of all operations. Hence, the lattices of sub-moding and of sub-kinding (Fig. 1) are Bi-Heyting algebras: they are products of bounded linear orders. Bi-Heyting algebras are distributive lattices, *i.e.*, meets and joins distribute in both directions. We use this fact implicitly throughout this text.

3.2 Kind Calculus

We now describe the SMOki components specific to kinds: the type context Δ , a syntactic notion of kind functions, and the sub-kinding judgment $\Delta \vDash \phi \leq \psi$. All three components are agnostic to the set of types, the underlying type system, the choice of modes, and the sub-kinding lattice. For the purpose of this section, fix an *arbitrary finite sub-kinding lattice* M . Our components are agnostic to what kinds are used for: we defer the mode-crossing rule that connects kinds to typing to §3.4.

Type contexts Δ are lists of abstract types. Abstract types are either type variables α , which are always unrestricted, or abstract type constructors t with an associated arity and kind. All type variables and abstract type constructors in the context are associated with an implicit abstract base kind or abstract kind function, denoted by a hat like $\hat{\alpha}$ or \hat{t} , respectively. Type constructors appear in type contexts as elements $t : \bar{\alpha}. \phi$, where the *kind term* ϕ describes a kind bound on t given the argument kinds $\bar{\alpha}$. The kind term ϕ may also mention the kinds of abstract types in its context Δ .

Kind terms provide a syntactic description of kind functions. A kind term may contain base kind constants m , joins of two kind terms, and meets of a kind term with a base kind (matching the form of kind functions in §2). Finally, kind terms can form (non-uniform) fixpoints on the underlying lattice, modeled using a parameterized μ construct. Such fixpoints always exist because all kind terms correspond to monotone functions. Note that kind terms have no way of referring to non-abstract types like `int`: instead, we describe how to translate types to their kinds in §3.4.

Finally, we discuss the sub-kinding judgment $\Delta \vDash \phi \leq \psi$. Intuitively, it should mean the following: the base kind determined by ϕ should be at most that determined by ψ , under any instantiation of abstract base kinds $\hat{\alpha}$ (for type variables) and abstract kind functions \hat{t} (for abstract type constructors) from Δ .

Actually defining this judgment is a little tricky. We define it *semantically* (hence the \vDash): we first quantify over all “valid” abstract base kinds and kind functions determined by Δ , then evaluate $\phi \leq \psi$ in the underlying lattice.

First, we define a semantic interpretation from kind terms to base kinds with respect to a *substitution* that maps variables $\hat{\alpha}$ and \hat{t} to base kinds or kind functions, respectively:

Definition 2. A *substitution* is a function $\rho : (\text{TyCon} \cup \text{TyVar}) \rightarrow \text{Mon}(M^*, M)$, where $A \rightarrow B$ denotes the set of partial functions and $\text{Mon}(A, B)$ denotes the set of monotone functions from A to B . (Substitutions don’t fix the arity of abstract kind function symbols \hat{t} to simplify the notation.) Given a substitution ρ , we define the *evaluation* $\llbracket \phi \rrbracket_\rho$ of a kind term to a base kind as follows:

$$\begin{aligned} \llbracket \hat{\alpha} \rrbracket_\rho &:= \rho(\alpha)() & \llbracket \phi \sqcup \psi \rrbracket_\rho &:= \llbracket \phi \rrbracket_\rho \sqcup \llbracket \psi \rrbracket_\rho \\ \llbracket \hat{t} \bar{\phi} \rrbracket_\rho &:= \rho(t)(\overline{\llbracket \bar{\phi} \rrbracket_\rho}) & \llbracket \phi \sqcap m \rrbracket_\rho &:= \llbracket \phi \rrbracket_\rho \sqcap m \\ \llbracket m \rrbracket_\rho &:= m & \llbracket (\mu \hat{t} \bar{\alpha}. \phi) \bar{\psi} \rrbracket_\rho &:= (\mu(\lambda f. \lambda \bar{m}. \llbracket \phi \rrbracket_{\rho, t \mapsto f, \bar{\alpha} \mapsto \bar{m}})) \overline{\llbracket \bar{\psi} \rrbracket_\rho} \end{aligned}$$

Note that we use a semantic least fixpoint μF over the lattice of functions $\text{Mon}(M^n, M)$ from n -tuples of base kinds to base kinds to interpret a syntactic fixpoint.

Now we define when a substitution ρ is valid for a type context Δ , by requiring that the inequalities of kinds determined by Δ hold, and use it to define sub-kinding.

589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637

$$\begin{aligned}
& \alpha, \beta \in \text{TyVar} := \dots \quad t \in \text{TyCon} := \dots \\
& p \in \text{Portability} := \mathbf{nonportable} \mid \mathbf{shareable} \mid \mathbf{portable} \\
& c \in \text{Contention} := \mathbf{uncontended} \mid \mathbf{shared} \mid \mathbf{contended} \\
& w \in W := \text{Portability} \quad n \in N := \text{Contention} \quad m \in M := W \times N \\
& \phi, \psi, \chi \in \text{KindTerm} ::= \hat{\alpha} \mid \hat{t} \bar{\phi} \mid m \mid \phi \sqcup \psi \mid \phi \sqcap m \mid (\mu \hat{t} \bar{\alpha}. \phi) \bar{\psi} \\
& A, B, C \in \text{Type} ::= \alpha \mid t \bar{A} \mid A @ m_1 \rightarrow B @ m_2 \mid \square^m A \mid (\mu t \bar{\alpha}. A) \bar{B} \mid \mathbb{1} \mid \mathbb{B} \mid \mathbb{Z} \mid A \times B \mid A + B \mid \\
& \quad \text{ref } A \mid \text{atomic } A \mid \dots \\
& \Delta ::= \emptyset \mid \Delta, \alpha \mid \Delta, t : \bar{\alpha}. \phi \\
& \Gamma ::= \emptyset \mid \Gamma, x : - \mid \Gamma, x : A @ m \\
& \quad \text{Substitution (selected rules)} \\
& (t \bar{A})[t \mapsto \bar{\alpha}. C] := C[\bar{\alpha} \mapsto \overline{A[t \mapsto \bar{\alpha}. C]}] \\
& (t' \bar{A})[t \mapsto \bar{\alpha}. C] := t' \overline{A[t \mapsto \bar{\alpha}. C]} \quad (t \neq t') \\
& ((\mu t' \bar{\beta}. A) \bar{B})[t \mapsto \bar{\alpha}. C] := (\mu t' \bar{\beta}. A[t \mapsto \bar{\alpha}. C]) \overline{B[t \mapsto \bar{\alpha}. C]} \\
& \quad \text{The "kind of" operation } \text{kind}(A) \\
& \text{kind}(\alpha) \quad := \hat{\alpha} \qquad \text{kind}(A_1 \times A_2) \quad := \text{kind}(A_1) \sqcup \text{kind}(A_2) \\
& \text{kind}(t \bar{A}) \quad := \hat{t} \text{kind}(A) \qquad \text{kind}(A_1 + A_2) \quad := \text{kind}(A_1) \sqcup \text{kind}(A_2) \\
& \text{kind}(\mathbb{1}) \quad := (\mathbf{portable}, \mathbf{contended}) \quad \text{kind}(A @ m_1 \rightarrow B @ m_2) := (\mathbf{nonportable}, \mathbf{contended}) \\
& \text{kind}(\mathbb{B}) \quad := (\mathbf{portable}, \mathbf{contended}) \quad \text{kind}(\text{ref } A) \quad := (\mathbf{portable}, \mathbf{uncontended}) \sqcup \text{kind}(A) \\
& \text{kind}(\mathbb{Z}) \quad := (\mathbf{portable}, \mathbf{contended}) \quad \text{kind}(\text{atomic } A) \quad := (\mathbf{portable}, \mathbf{contended}) \\
& \text{kind}(\square^m A) := \text{kind}(A) \sqcap m \quad \text{kind}((\mu t \bar{\alpha}. A) \bar{C}) \quad := (\mu \hat{t} \bar{\alpha}. \text{kind}(A)) \overline{\text{kind}(C)} \\
& \boxed{\Delta; \Gamma \vdash e : A @ m} \\
& \text{CROSS} \quad \frac{\Delta; \Gamma \vdash e : A @ (w, n) \quad \Delta \vDash \text{kind}(A) \leq (w', n')}{\Delta; \Gamma \vdash e : A @ (w' \sqcap w, n' / n)} \qquad \text{WEAKEN} \quad \frac{\Delta; \Gamma \vdash e : A @ (w, n) \quad w \leq w' \quad n' \leq n}{\Delta; \Gamma \vdash e : A @ (w', n')} \\
& \text{LAM} \quad \frac{\Delta; \mathbb{w}_w(\Gamma), x : A @ m_1 \vdash e : B @ m_2}{\Delta; \Gamma \vdash \lambda x. e : (A @ m_1 \rightarrow B @ m_2) @ (w, n)} \qquad \text{APP} \quad \frac{\Delta; \Gamma \vdash e_1 : (A @ m_1 \rightarrow B @ m_2) @ m_3 \quad \Delta; \Gamma \vdash e_2 : A @ m_1}{\Delta; \Gamma \vdash e_1 e_2 : B @ m_2} \\
& \text{BOX} \quad \frac{\Delta; \Gamma \vdash e : A @ m_1 \sqcap m_2}{\Delta; \Gamma \vdash \text{box } e : \square^{m_2} A @ m_1} \qquad \text{UNBOX} \quad \frac{\Delta; \Gamma \vdash e : \square^{m_2} A @ m_1}{\Delta; \Gamma \vdash \text{unbox } e : A @ m_1 \sqcap m_2} \\
& \text{ROLL} \quad \frac{\forall i. \Delta \vdash B_i \text{ ok} \quad \Delta; \Gamma \vdash e : A[\bar{\alpha} \mapsto \bar{B}, t \mapsto \bar{\alpha}. (\mu t \bar{\alpha}. A) \bar{\alpha}] @ m}{\Delta; \Gamma \vdash \text{roll } e : (\mu t \bar{\alpha}. A) \bar{B} @ m} \qquad \text{UNROLL} \quad \frac{\Delta; \Gamma \vdash e : (\mu t \bar{\alpha}. A) \bar{B} @ m}{\Delta; \Gamma \vdash \text{unroll } e : A[\bar{\alpha} \mapsto \bar{B}, t \mapsto \bar{\alpha}. (\mu t \bar{\alpha}. A) \bar{\alpha}] @ m}
\end{aligned}$$

Fig. 4. Selected SMOki rules. The remaining substitution rules are standard. Parallel substitution of type variables into types $A[\bar{\alpha} \mapsto \bar{B}]$ and substitutions on lattice terms are defined analogously. The judgement $\Delta \vdash A \text{ ok}$ is defined as usual to check whether type constructors are applied correctly. All lattice operations are with respect to *sub-kindings*.

Definition 3. We write $\rho \in \llbracket \Delta \rrbracket$ if

- (1) $\text{dom } \rho = \text{dom } \Delta$.
- (2) for all $(t : \bar{\alpha}. \phi) \in \Delta$, for all \bar{m} , $\rho \ t \ \bar{m} \leq \llbracket \phi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\rho$.
- (3) ρ is valid: for all $(t : \bar{\alpha}. \phi) \in \Delta$, there exists χ such that for all \bar{m} , $\rho \ t \ \bar{m} = \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\emptyset$.

Definition 4. We say that $\Delta \vDash \phi \leq \psi$ if for every $\rho \in \llbracket \Delta \rrbracket$ we have $\llbracket \phi \rrbracket_\rho \leq \llbracket \psi \rrbracket_\rho$.

There is a subtle validity condition (3) in Def. 3. For sub-kinding, intuitively, we need to check $\llbracket \phi \rrbracket_\rho \leq \llbracket \psi \rrbracket_\rho$ only for ρ where all functions $\rho \ t$ could have come from a valid kind term, *i.e.*, are “realizable” kind functions. However, some functions in $\text{Mon}(M^*, M)$ don’t correspond to any such realizable kind. (Specifically, only those with a normal form as presented in §2.9 do.) The condition ensures that we only need to check $\llbracket \phi \rrbracket_\rho \leq \llbracket \psi \rrbracket_\rho$ for realistic substitutions ρ . It does so by requiring $\rho \ t$ be representable by a kind term χ wrt. the empty substitution, which can be shown equivalent to the normal form condition from §2.9.

3.3 Modal Type Systems

We now instantiate §3.2 to a concrete model of OxCamL. Key rules are in Fig. 4; the supplementary material gives the remaining definitions [23]. DRFCamL background appears in [12, §4].

We distinguish between the lattices of comonadic modes $w \in W := \text{Portability}$ and monadic modes $n \in N := \text{Contention}$. Comonadic modes are those that track properties of functions, and monadic modes are those that track properties of references. To accommodate more modal axes, the W and N lattices can also be instantiated with (the products of) multiple modal axis lattices. Modes themselves are drawn from $m \in M := W \times N$. We order these sets as the sub-kinding lattice, rather than by sub-moding as in previous work. The choice of ordering affects, for instance, the WEAKEN rule: It allows *weakening* from mode (w, n) to (w', n') when $w \leq w'$ and $n' \leq n$ (note the flipped direction for the monadic component).

Our typing judgment $\Delta; \Gamma \vdash e : A @ m$ can mention modes in three places: as part of the output (alongside the type), within types themselves, and in variable bindings in the term context Γ . Binders in Γ are either associated with a type and a mode, or *inaccessible*, denoted with $x : -$.

Inaccessible binders model the capture behavior of closures. Consider the LAM rule for lambda expressions $\lambda x. e$. It has type $A @ m_1 \rightarrow B @ m_2$ if, given $x : A @ m_1$, the body e has type B at mode m_2 . The mode of the closure itself, (w, n) , also affects typing, because it determines what the closure can capture from its environment, and at what mode it can access it. This depends only on the comonadic component (so portability). The LAM rule expresses this constraint by applying a *lock* \mathfrak{L}_w to the term context Γ before the binder $x : A @ m_1$. Locks are operations on contexts that change the modes of binders in the context or render them inaccessible, depending on w .

$$\begin{aligned} \mathfrak{L}_w(\emptyset) &:= \emptyset \\ \mathfrak{L}_w(\Gamma, x : -) &:= \mathfrak{L}_w(\Gamma), x : - \\ \mathfrak{L}_{w_1}(\Gamma, x : A @ (n_2, w_2)) &:= \begin{cases} \mathfrak{L}_{w_1}(\Gamma), x : A @ (n_2, w_2) & \text{if } w_1 = \text{nonportable} \\ \mathfrak{L}_{w_1}(\Gamma), x : A @ (n_2, \text{contended}) & \text{if } w_1 = w_2 = \text{portable} \\ \mathfrak{L}_{w_1}(\Gamma), x : - & \text{otherwise} \end{cases} \end{aligned}$$

Since locks are an operation (and not syntactically part of contexts), the variable rule of our system is simple: it just looks up a variable in the context. (Keeping inaccessible binders in the context instead of just removing them ensures that locks and variable shadowing interact well together.)

3.4 Instantiation to OCaml

We now instantiate the kind system from §3.2. The sub-kinding lattice M is the sub-kinding lattice over the set of modes from §2.5. We present OCaml-specific types and their kinds in Fig. 4, including integers, products, sums, references, and iso-recursive types, among others.

The kind of a type A is computed by a function $\text{kind}(A)$ from types to kind terms in Fig. 4. The key typing rule in SMOki is CROSS. Its antecedent $\Delta \models \text{kind}(A) \leq (w', n')$ requires that the kind of A , has at least the mode-crossing behavior (w', n') . Recall from §2.7 that for a type with mode-crossing behavior (w', n') , all modes *above* (w', n') in sub-kinding collapse; they are equivalent and can be crossed between. We need to determine the “best” mode a value currently at mode (w, n) can cross to, that is, the least mode *with respect to sub-moding*.

On comonadic axes, sub-kinding and sub-moding coincide: for each axis, if the crossing mode w' is less than the current mode w , we want to cross to w' , and otherwise stay at w . This is the meet $w \sqcap w'$. The situation for monadic axes is different: for each axis, if the crossing mode n' is less than or equal to the current mode n , we want to cross to the *top mode* on that axis. This top mode corresponds to the bottom (“strongest”) mode with respect to sub-moding. Otherwise, if $n' > n$, we want to stay at mode n . We actually have an operation that does this: it is the implication $/$ from a Bi-Heyting algebra! Hence, the monadic mode to cross to is n' / n .

For instance, if we have a $n = \text{contended}$ value that crosses $n' = \text{contended}$, then it should cross to $\text{contended} / \text{contended} = \text{uncontended}$. (Recall the definition of $/$ in bounded linear orders from §3.1.) If it only crossed $n' = \text{shared}$, then it would cross to $\text{shared} / \text{contended} = \text{contended}$.

The modality type \square^m is used to model record fields annotated with a modality (§2.7). It is annotated with an arbitrary mode m , and its typing rules take a meet (on the sub-kinding lattice) of the ambient mode with m . We can recover all modalities considered in §2.7 by choosing m as follows, for instance for @@ **portable**, @@ **contended**, and @@ **shared**, respectively:

(**portable, uncontended**) (**nonportable, contended**) (**nonportable, shared**)

Our formalization supports non-uniform recursive type constructors. To avoid syntactic overhead, we desugar mutually recursive types to non-mutually recursive types using *Bekić’s theorem* [1]. The kind of a recursive type is defined as a fixpoint of its underlying type function, as expected.

3.5 Modules

SMOki does not explicitly formalize a module system: it only formalizes abstract types, and so disregards many aspects of a full module system. Instead, the examples of modules in this paper should be viewed through the lens of abstract types, as formalized by typing contexts Δ .

For example, an ambient opaque module signature

```

722 module Foo : sig
723   type t : value mod portable contended
724   val check : t → bool
725 end
726

```

can be read as SMOki context assumptions

$$\Delta = t : (\text{portable, contended}) \quad \Gamma = \text{check} : t \rightarrow \text{bool}.$$

A client may use the fact that `Foo.t` crosses portability and contention, but it cannot inspect the representation of `Foo.t`. Correctness of module instantiation, then, corresponds to a form of substitution. For instance, a concrete implementation may define `Foo.t` as `int`:

```

733 module Foo : sig ... end = struct
734   type t = int
735

```

```

736   let check x = x = 0
737   end

```

The *module inclusion check* requires the usual sub-kinding obligation:

$$\text{kind}(\mathbb{Z}) \leq (\text{portable}, \text{contended}).$$

After this check, replacing the abstract type by its implementation should preserve typing. This is the role of the substitution theorem:

Theorem 5 (Substitution). *The following rules are admissible, where all objects must be closed in their respective contexts:*

$$\frac{\Delta, \alpha; \Gamma \vdash e : A @ m}{\Delta[\hat{\alpha} \mapsto \text{kind}(B)]; \Gamma[\alpha \mapsto B] \vdash e : A[\alpha \mapsto B] @ m}$$

$$\frac{\Delta, t : \bar{\alpha}. \phi; \Gamma \vdash e : A @ m \quad \Delta, \bar{\alpha} \models \text{kind}(B) \leq \phi}{\Delta[\hat{t} \mapsto \bar{\alpha}. \text{kind}(B)]; \Gamma[t \mapsto \bar{\alpha}. B] \vdash e : A[t \mapsto \bar{\alpha}. B] @ m}$$

The first rule substitutes a concrete type for an abstract variable ($\alpha \mapsto B$, $\hat{\alpha} \mapsto \text{kind}(B)$). The second does the analogous higher-order substitution for abstract constructors. Both follow from similar substitution lemmas for $\Delta \models \phi \leq \psi$. Although we leave a formal treatment of OxCaml modules to future work, we believe the above substitution theorem is the key sanity check needed to ensure soundness of modes in the presence of modules.

3.6 Correctness

We build on the DRFCaml semantic model in Rocq with Iris [16, 17] to prove the soundness of SMOki. The mechanized proof is part of the supplementary materials [23]. Here we restate the soundness theorem and outline the changes required to support mode crossing.

DRFCaml defines an operational semantics for a language with an explicit separation between heap and stack, concurrency, and data race detection à la RustBelt [15]. Its logical relation supports the portability, contention, locality, affinity, and uniqueness modal axes. Although we present SMOki only with the portability and contention axes, this is only a projection used for exposition: our soundness proof handles all of DRFCaml's axes. The complete rule set appears in the supplementary material [23]. When instantiated with all modal axes, SMOki strictly extends DRFCaml. We extend the DRFCaml logical relation to prove the following theorem in Rocq:

Theorem 6. *SMOki is sound; that is, if a program type-checks in the empty context $\emptyset; \emptyset \vdash e : A @ m$, then no execution of e crashes or has a data race.*

The mechanized soundness theorem is in the style of RustBelt [15], where contexts (specifically Δ) are not modeled explicitly. Instead, RustBelt models such contexts as assumptions in the meta-level context. Similarly, we do not explicitly formalize the function that computes the kind of a type $\text{kind}(A)$; instead, we define a predicate $\text{crosses}(\tau, m)$ stating that a semantic type τ crosses modal base kind m and prove properties of this predicate for all type constructors.

Conceptually, the $\text{crosses}(\tau, m)$ predicate is defined as follows:

$$\text{crosses}(\tau, m) := \forall m' : M, v : \text{val}. \tau(m', v) \Leftrightarrow \tau(m \sqcap m', v)$$

In DRFCaml's semantic model, semantic types τ are predicates not only over values, but also over the mode m of the value. (They are also parameterized over a *world* Δ that we omit here for simplicity.) The $\text{crosses}(\tau, m)$ predicate states that, if a type τ has modal base kind m , then for its




Fig. 5. Reduction architecture of inference.

inhabitants, all modes above m collapse. Indeed, if we know $\text{crosses}(\tau, m)$ and have two modes $m_1 \geq m$ and $m_2 \geq m$, then for all v :

$$\tau(m_1, v) \xleftrightarrow{\text{cross}} \tau(m \sqcap m_1, v) \xleftrightarrow{\text{absorb}} \tau(m, v) \xleftrightarrow{\text{absorb}} \tau(m \sqcap m_2, v) \xleftrightarrow{\text{cross}} \tau(m_2, v)$$

For many types, such as products and functions, establishing their mode-crossing behavior in the semantic model is straightforward. Proving the correct mode-crossing rules for mutable references and recursive types required some innovation; see the supplementary material for details [23].

Interestingly, the proof that the mode-crossing behavior of a recursive type corresponds to a least fixpoint is almost trivial: it follows directly by induction on the fixpoint used to define the semantic model of recursive types.

4 Inference

This section describes the part of the compiler that infers and checks the modal kinds of type expressions and decides sub-kinding obligations. It does not cover ordinary type or mode inference; those phases use the resulting mode-crossing information, but are separate from the kind inference problem studied here. The supplementary material gives the formal reduction and solver details [23]. We cover non-recursive coefficients, abstract-kind constraints, recursive fixpoint equations, and solver architecture.

Inference is coefficient-driven: sub-kinding obligations and recursive kind equations reduce to coefficient constraints. These coefficients are then treated symbolically for normalization and solving. See Fig. 5 for an overview.

4.1 Computing the Kind of a Type Expression

This subsection sets up the representation used by the rest of the algorithm: the modal kind of every type expression is interpreted as a lattice expression and then normalized to coefficient form. The later subsections all use this same representation of modal kinds.

Recall the kinds in Fig. 3. In general, the “standard form” of the kind of a type constructor t is:

```
('a, 'b) t : value mod t.0 with 'a @@ t.1 with 'b @@ t.2
```

where the “kind coefficients” $t.0$, $t.1$, $t.2$ range over the sub-kinding lattice:

- $t.0$ is the base contribution of $(\text{'a}, \text{'b}) t$. It records which modes are crossed by t itself, independent of its arguments. For example, `list` crosses `portability` and `contention`, while `ref` crosses only `portability`.
- $t.1$, $t.2$ describe how arguments `'a`, `'b` affect the mode-crossing behavior of $(\text{'a}, \text{'b}) t$. For example, `'a → 'b` ignores its arguments entirely, while `'a @@ m` weakens `'a`'s restrictions by m .

834 This yields the following kind coefficients for the base type constructors, where \perp is **portable contended**
 835 and \top includes no modes:

```

836         int.0 =  $\perp$ 
837
838         list.0 =  $\perp$            list.1 =  $\top$ 
839         ref.0 = portable       ref.1 =  $\top$ 
840         (*).0 =  $\perp$            (*).1 =  $\top$            (*).2 =  $\top$ 
841         ( $\rightarrow$ ).0 = contended   ( $\rightarrow$ ).1 =  $\perp$        ( $\rightarrow$ ).2 =  $\perp$ 
842
843         (@@ m).0 =  $\perp$            (@@ m).1 =  $m$ 

```

844 Here we have used that a plain **with** 'a bound is equivalent to **with** 'a @@ \top , and **with** 'a @@ \perp is
 845 equivalent to no with-bound at all. Semantically, read a surface-syntax kind as a lattice expression:
 846 interpret **with** as join (combine restrictions) and @@ as meet (weaken restrictions):

```

847         ('a, 'b) t : value mod t.0 with 'a @@ t.1 with 'b @@ t.2
848
849                 means

```

```

850         kind(('a, 'b) t) = t.0  $\sqcup$  (t.1  $\sqcap$  kind('a))  $\sqcup$  (t.2  $\sqcap$  kind('b))
851

```

852 Computing the kind of a compound type expression is straightforward: apply the rule above
 853 repeatedly. For example, for (int list) ref, we have:

```

854         kind((int list) ref) = ref.0  $\sqcup$  (ref.1  $\sqcap$  kind(int list))
855                               = ref.0  $\sqcup$  (ref.1  $\sqcap$  (list.0  $\sqcup$  (list.1  $\sqcap$  int.0)))
856                               = portable  $\sqcup$  ( $\top$   $\sqcap$  ( $\perp$   $\sqcup$  ( $\top$   $\sqcap$   $\perp$ )))
857                               = portable
858

```

859 This gives a systematic way to compute the kind of any type expression.

861 4.2 Non-recursive Type Definitions

862 For non-recursive definitions, inference computes the kind of the right-hand side, normalizes it to
 863 coefficient form, and stores the resulting coefficients for later uses of the constructor.

864 For a non-recursive type definition

```

865         type t = (right-hand side)
866

```

867 we store the kind of t by setting t.0 = kind(right-hand side). For type constructor definitions:

```

868         type ('a, 'b) t = (right-hand side mentioning 'a and 'b)
869

```

870 we compute the right-hand-side kind as a *lattice expression*. Because the right-hand side can mention
 871 'a and 'b, the resulting lattice expression need not be a ground lattice element; it may contain
 872 kind('a) and kind('b) as subexpressions. We simplify it to coefficient form

```

873         kind(right-hand side mentioning 'a and 'b) = A  $\sqcup$  (B  $\sqcap$  kind('a))  $\sqcup$  (C  $\sqcap$  kind('b))
874

```

875 then extract t.0 := A, t.1 := B, t.2 := C. (Note how this coefficient form never has a term containing
 876 kind('a) \sqcap kind('b) because no type has this kind.) For example:

```

877         type ('a, 'b) t = ('a * ('b @@ portable)) ref
878

```

879 Here we have

```

880         kind(('a * ('b @@ portable)) ref) = portable  $\sqcup$  ( $\top$   $\sqcap$  kind('a))  $\sqcup$  (portable  $\sqcap$  kind('b))
881
882

```

and extract the kind coefficients for t from this formula:

$$t.\theta := \text{portable}, \quad t.1 := \top, \quad t.2 := \text{portable}$$

Later uses of t reuse these coefficients directly.

4.3 Abstract Types

Abstract type declarations introduce the need for a solver for sub-kinding checks. The reason is that unknown coefficients from abstract type declarations and with-bound annotations produce hypotheses in the context; the algorithm must eliminate those hypotheses and reduce each check to comparison of normalized lattice polynomials.

With abstract types in scope, coefficients need not be ground constants; they become lattice polynomials over abstract kind symbols. For an abstract unary type constructor u , we use unknown coefficients $u.\theta$, $u.1$ and write

$$\text{kind}('a u) = u.\theta \sqcup (u.1 \sqcap \text{kind}('a)).$$

If present, a kind annotation constrains these unknowns. First consider no annotation:

```

899 type elt
900 type 'a u
901 type 'a t = ('a u * elt) ref

```

Normalizing the right-hand side gives

$$\text{kind} (('a u * \text{elt}) \text{ ref}) = \text{portable} \sqcup u.\theta \sqcup \text{elt}.\theta \sqcup (u.1 \sqcap \text{kind}('a)),$$

so we extract

$$t.\theta := \text{portable} \sqcup u.\theta \sqcup \text{elt}.\theta, \quad t.1 := u.1.$$

As before, later uses of t consult these coefficients directly; the difference is that coefficients are now symbolic *lattice polynomials*.

Sub-kind checks with abstract types. A sub-kinding check $\Delta \vDash \phi \leq \psi$ quantifies over all abstract-type instantiations in Δ . We reduce it to equality via

$$\phi \leq \psi \iff \phi \sqcup \psi = \psi$$

then decide equality by reducing both sides to *minimal normal form*.

Minimal normal form. To introduce minimal normal form, consider checking that the kind of $'a u u$ equals the kind of $'a u$:

$$\begin{aligned} \text{kind}('a u) &= u.\theta \sqcup (u.1 \sqcap \text{kind}('a)) \\ \text{kind}('a u u) &= u.\theta \sqcup (u.1 \sqcap \text{kind}('a u)) \\ &= u.\theta \sqcup (u.1 \sqcap u.\theta) \sqcup (u.1 \sqcap \text{kind}('a)) \\ &= u.\theta \sqcup (u.1 \sqcap \text{kind}('a)) \end{aligned}$$

These kinds normalize to the same formula, so they are equal. Here is how we normalize in general:

- First, we distribute meets over joins, use idempotence, and combine like terms to obtain a formula in *polynomial form*, which looks like this for two abstract kind symbols x and y :

$$p(x, y) = c_0 \sqcup (c_1 \sqcap x) \sqcup (c_2 \sqcap y) \sqcup (c_3 \sqcap x \sqcap y)$$

where c_0, c_1, c_2, c_3 are constants and x, y are abstract kind symbols. (“polynomial” by analogy with ordinary polynomials over numbers, where $+$ is join and \cdot is meet)

- Second, we reduce the polynomial to minimal form by subtracting coefficients of all smaller terms from coefficients of larger terms, which gives:

$$p(x, y) = c_0 \sqcup ((c_1 \setminus c_0) \sqcap x) \sqcup ((c_2 \setminus c_0) \sqcap y) \sqcup ((c_3 \setminus (c_0 \sqcup c_1 \sqcup c_2)) \sqcap x \sqcap y)$$

where \setminus is the co-Heyting subtraction operation on the base lattice.

The second step is essential: distribution alone does not yield a canonical form. For example:

$$x \sqcup (y \sqcap x) = x$$

$$(\text{portable} \sqcap x) \sqcup (y \sqcap x) = (\text{portable} \sqcap x) \sqcup (\text{contended} \sqcap y \sqcap x)$$

That is, we choose the minimal value for coefficients of the polynomial. The next example illustrates this concretely.

Example. Consider the following type definitions:

```

945 type 'a x
946 type 'a y
947 type 'a f = 'a x y * int x
948 type 'a g = 'a y x * int y

```

After computing kind coefficients, we obtain:

$$\begin{array}{l}
 951 \quad f.0 = \underbrace{y.0 \sqcup (y.1 \sqcap x.0)}_{\text{from 'a x y}} \sqcup \underbrace{x.0}_{\text{from int x}} \\
 952 \\
 953 \quad f.1 = \underbrace{x.1 \sqcap y.1}_{\text{from 'a x y}} \\
 954 \\
 955 \\
 956 \quad g.0 = \underbrace{x.0 \sqcup (x.1 \sqcap y.0)}_{\text{from 'a y x}} \sqcup \underbrace{y.0}_{\text{from int y}} \\
 957 \\
 958 \quad g.1 = \underbrace{x.1 \sqcap y.1}_{\text{from 'a y x}}
 \end{array}$$

The linear coefficients already match $f.1 = g.1 = x.1 \sqcap y.1$. Absorption ($a \sqcup (b \sqcap a) = a$) simplifies only the base coefficients, yielding:

$$f.0 = x.0 \sqcup y.0 \qquad g.0 = x.0 \sqcup y.0$$

Thus, these two types have the same kind. Note that if we instead had `type 'a f = 'a x y` and `type 'a g = 'a y x`, we would get different kinds:

$$\begin{array}{l}
 962 \quad f.0 = y.0 \sqcup (y.1 \sqcap x.0) \qquad g.0 = x.0 \sqcup (x.1 \sqcap y.0) \\
 963 \\
 964 \quad f.1 = x.1 \sqcap y.1 \qquad g.1 = x.1 \sqcap y.1
 \end{array}$$

$f.0$ and $g.0$ are genuinely different; unlike above, there are no absorbed terms. This example used absorption in an ad-hoc way; the following lemma gives the general normalization rule.

LEMMA 7 (MINIMAL NORMAL FORM). *For every polynomial $p(x_1, \dots, x_n)$, define*

$$\text{MNF}(p) = \bigsqcup_{S \subseteq \{1, \dots, n\}} \left(p(\chi_S) \setminus \bigsqcup_{T \subset S} p(\chi_T) \right) \sqcap \prod_{i \in S} x_i.$$

Here χ_S is the valuation with $\chi_S(x_i) = \top$ if $i \in S$, and $\chi_S(x_i) = \perp$ otherwise (with $\prod_{i \in \emptyset} x_i = \top$). A valuation ρ maps each variable x_i to an element of the base lattice. Then $\forall \rho. \text{MNF}(p)(\rho) = p(\rho)$, and for all polynomials p, q :

$$\text{MNF}(p) \text{ and } \text{MNF}(q) \text{ are syntactically equal} \iff \forall \rho. p(\rho) = q(\rho).$$

The supplementary material contains the proof [23]. The subset-based formulation of MNF is exponential in arity; in practice we use LDD-based normalization and solver operations, which are described in the supplementary material [23].

With-bounds and abstract annotations. With-bounds that appear in kind annotations of abstract types are compiled into solver hypotheses. For

```
type ('a1, ..., 'an) u : value mod m0 with 'a1 @@ m_1 ... with 'an @@ m_n
```

we keep $u.\emptyset, u.1, \dots, u.n$ symbolic and register

$$\begin{aligned} \forall a_1, \dots, a_n. u.\emptyset \sqcup (u.1 \sqcap a_1) \sqcup \dots \sqcup (u.n \sqcap a_n) \\ \leq m_0 \sqcup (m_1 \sqcap a_1) \sqcup \dots \sqcup (m_n \sqcap a_n). \end{aligned}$$

This is reduced using the following lemma.

LEMMA 8 (QUANTIFIER ELIMINATION FOR WITH-BOUNDS). *The quantified inequality*

$$\forall a_1, \dots, a_n. u.\emptyset \sqcup (u.1 \sqcap a_1) \sqcup \dots \sqcup (u.n \sqcap a_n) \leq m_0 \sqcup (m_1 \sqcap a_1) \sqcup \dots \sqcup (m_n \sqcap a_n)$$

is equivalent to the following set of quantifier-free inequalities on the $u.i$ symbols:

$$u.\emptyset \leq m_0 \quad \text{and} \quad u.i \leq m_i \sqcup m_0 \quad (1 \leq i \leq n).$$

LEMMA 9 (BOUNDED-QUANTIFIER SUBSTITUTION). *For predicate P on lattice elements and lattice-polynomial bound $b(x)$,*

$$\forall x. (x \leq b(x) \Rightarrow P(x)) \iff \forall x. P(x \sqcap b(x)).$$

Together, these lemmas reduce a judgment $\Delta \vDash \phi \leq \psi$ to $\vDash \phi' \leq \psi'$. First, quantified inequalities on lattice functions in Δ are reduced to quantifier-free inequalities on coefficient symbols. Second, bounded implications are rewritten by substitution. The resulting simple judgment is then checked by testing whether $\text{MNF}(\phi' \sqcup \psi') \stackrel{?}{=} \text{MNF}(\psi')$.

Example: Map functor signature inclusion. We reuse the map functor from §2.8:

```
module type OrderedType = sig
  type t
  val compare : t -> t -> int
end
module Map (Ord : OrderedType) = struct
  type 'a t : value mod portable contended with 'a with Ord.t = ...
end
```

Now build maps on key pairs by currying two map applications:

```
module PairMap (A : OrderedType) (B : OrderedType) : sig
  type 'v t : value mod portable contended with 'v with A.t with B.t
end = struct
  module MB = Map(B) module MA = Map(A)
  type 'v t = ('v MB.t) MA.t
end
```

We must check that the implementation kind is a sub-kind of the signature kind. In the functor body, $A.t$ and $B.t$ are abstract, with symbols $A.t.\emptyset, B.t.\emptyset$. The signature requires:

$$\text{kind}('v t)_{\text{sig}} = \text{kind}('v) \sqcup A.t.\emptyset \sqcup B.t.\emptyset.$$

The implementation defines $'v t = ('v MB.t) MA.t$, so:

$$\text{kind}('v t)_{\text{impl}} = MA.t.\emptyset \sqcup (MA.t.1 \sqcap (MB.t.\emptyset \sqcup (MB.t.1 \sqcap \text{kind}('v))))$$

By the definition of the Map functor, we have:

$$MA.t.\emptyset = A.t.\emptyset, \quad MA.t.1 = T, \quad MB.t.\emptyset = B.t.\emptyset, \quad MB.t.1 = T.$$

1030 So the implementation kind simplifies to $\text{kind}('v \ t)_{\text{impl}} = A.t.\emptyset \sqcup B.t.\emptyset \sqcup \text{kind}('v)$ This is precisely
 1031 as required by the signature.

1032 **Map functor inclusion with an abstract signature.** Now consider the same construction as
 1033 above, but where Map is abstracted behind a signature:
 1034

```
1035 module Map (Ord : OrderedType) : sig
1036   type 'a t : value mod portable contended with 'a with Ord.t
1037   end = struct ... end
1038 module PairMapViaSig (A : OrderedType) (B : OrderedType) : sig
1039   type 'v t : value mod portable contended with 'v with A.t with B.t
1040   end = struct
1041     module MB = Map(B) module MA = Map(A)
1042     type 'v t = ('v MB.t) MA.t
1043   end
```

1044 This is identical to the previous PairMap example, except that Map is abstract, so we only know its
 1045 behavior through its signature. For $MA = \text{Map}(A)$ and $MB = \text{Map}(B)$, this yields quantified hypotheses:

$$1046 \quad \forall x. MA.t.\emptyset \sqcup (MA.t.1 \sqcap x) \leq A.t.\emptyset \sqcup x \quad \forall x. MB.t.\emptyset \sqcup (MB.t.1 \sqcap x) \leq B.t.\emptyset \sqcup x.$$

1047 Applying Lemma 8 to each yields coefficient bounds:

$$1048 \quad MA.t.\emptyset \leq A.t.\emptyset, \quad MA.t.1 \leq T, \quad MB.t.\emptyset \leq B.t.\emptyset, \quad MB.t.1 \leq T.$$

1050 The implementation and signature kinds are

$$1051 \quad \text{kind}('v \ t)_{\text{impl}} = MA.t.\emptyset \sqcup (MA.t.1 \sqcap (MB.t.\emptyset \sqcup (MB.t.1 \sqcap \text{kind}('v))))),$$

$$1052 \quad \text{kind}('v \ t)_{\text{sig}} = \text{kind}('v) \sqcup A.t.\emptyset \sqcup B.t.\emptyset.$$

1054 So inclusion reduces to the following universally quantified implication (omitting trivial T -bounds):

$$1055 \quad (MA.t.\emptyset \leq A.t.\emptyset) \wedge (MB.t.\emptyset \leq B.t.\emptyset) \implies \text{kind}('v \ t)_{\text{impl}} \leq \text{kind}('v \ t)_{\text{sig}}$$

1057 where the free coefficient symbols (and $\text{kind}('v)$) are universally quantified. Using Lemma 9 under
 1058 this universal closure, we substitute bounded coefficients in the conclusion, reducing to

$$1059 \quad (MA.t.\emptyset \sqcap A.t.\emptyset) \sqcup (MA.t.1 \sqcap ((MB.t.\emptyset \sqcap B.t.\emptyset) \sqcup (MB.t.1 \sqcap \text{kind}('v)))) \leq$$

$$1060 \quad \text{kind}('v) \sqcup A.t.\emptyset \sqcup B.t.\emptyset,$$

1062 which is then discharged by MNF equality.

1064 4.4 Recursive Types

1065 Recursive definitions require a least-fixpoint calculation. The solver computes the right-hand side in
 1066 coefficient form with recursive occurrences left symbolic, reduces the resulting fixpoint to equations
 1067 on coefficients, and eliminates those equations one coefficient at a time.

1068 For a recursive type constructor

```
1069   type ('a1, ..., 'an) t = (right-hand side mentioning 'a1, ..., 'an and t)
```

1071 we write its unknown kind in coefficient form as

$$1072 \quad \text{kind}(('a1, \dots, 'an) \ t) = t.\emptyset \sqcup (t.1 \sqcap a_1) \sqcup \dots \sqcup (t.n \sqcap a_n)$$

1074 Here a_i abbreviates the kind of the i -th type parameter. Unlike the non-recursive case, the right-
 1075 hand side may mention t itself. So, after computing $\text{kind}(\text{right-hand side})$ compositionally with t
 1076 treated abstractly, every recursive occurrence $\text{kind}(t(\phi_1, \dots, \phi_n))$ is expanded as

$$1077 \quad t.\emptyset \sqcup (t.1 \sqcap \text{kind}(\phi_1)) \sqcup \dots \sqcup (t.n \sqcap \text{kind}(\phi_n)).$$

1078

1079 After substitution and normalization, we collect the base part and the a_i -parts; $R_0(t.\theta, \dots, t.n)$ is
 1080 the resulting base coefficient, and $R_i(t.\theta, \dots, t.n)$ is the coefficient of a_i . Thus

1081 $\text{kind}(\text{right-hand side}) = R_0(t.\theta, \dots, t.n) \sqcup (R_1(t.\theta, \dots, t.n) \sqcap a_1) \sqcup \dots \sqcup (R_n(t.\theta, \dots, t.n) \sqcap a_n)$
 1082

1083 To make this equal to the left-hand side, we need to solve the system of equations:

$$1084 \quad t.\theta = R_0(t.\theta, \dots, t.n) \quad t.1 = R_1(t.\theta, \dots, t.n) \quad \dots \quad t.n = R_n(t.\theta, \dots, t.n)$$

1086 In fact, we want the *least* solution to this system of equations.

1087
 1088 **LEMMA 10 (COEFFICIENTWISE REDUCTION OF RECURSIVE KINDS).** *For a recursive constructor t with*
 1089 *right-hand side matching coefficient form R_i , the lattice function fixpoint is equivalent to finding the*
 1090 *least solution of the following system of coefficient equations:*

$$1091 \quad t.i = R_i(t.\theta, \dots, t.n) \quad (0 \leq i \leq n).$$

1093 The supplementary material validates this coefficientwise reduction [23].

1094 To find the least solution, we use that the $R_i(t.\theta, \dots, t.n)$ are *lattice polynomials* in the unknowns
 1095 $t.\theta, \dots, t.n$, meaning expressions involving only meets, joins, constants, and unknowns.⁷ This
 1096 means that as a function of $t.\theta$, R_0 can be written as:

$$1097 \quad R_0(t.\theta, t.1, \dots, t.n) = A(t.1, \dots, t.n) \sqcup (B(t.1, \dots, t.n) \sqcap t.\theta)$$

1098 for some A and B involving only $t.1, \dots, t.n$. Now note that $A(t.1, \dots, t.n)$ is a fixpoint of R_0 with
 1099 respect to $t.\theta$, because:

$$1100 \quad R_0(A(t.1, \dots, t.n), t.1, \dots, t.n) = A(t.1, \dots, t.n) \sqcup (B(t.1, \dots, t.n) \sqcap A(t.1, \dots, t.n))$$

$$1101 \quad = A(t.1, \dots, t.n)$$

1102 This is also the *least* fixpoint: if $A'(t.1, \dots, t.n)$ is another fixpoint, then:

$$1103 \quad A'(t.1, \dots, t.n) = R_0(A'(t.1, \dots, t.n), t.1, \dots, t.n)$$

$$1104 \quad = A(t.1, \dots, t.n) \sqcup (B(t.1, \dots, t.n) \sqcap A'(t.1, \dots, t.n))$$

$$1105 \quad \geq A(t.1, \dots, t.n)$$

1106 Because these equations are monotone over a complete lattice, Bekić-style decomposition applies:
 1107 having found the least fixpoint of the first equation with respect to $t.\theta$ in terms of the other
 1108 unknowns, we substitute the solution into the remaining equations and repeat this process for the
 1109 other unknowns, thus finding the least solution to the system.
 1110

1111 **LEMMA 11 (ONE-COEFFICIENT ELIMINATION).** *The least fixpoint of $x = A \sqcup (B \sqcap x)$ is $x = A$.*

1112 **Example: simple recursive type.** Consider the type definition:

1113 `type rlist = Nil | Cons of int ref * rlist`

1114 There is one coefficient, `rlist.θ`. Computing the right-hand side kind gives

$$1115 \quad \text{kind}(\text{int ref * rlist}) = \text{portable} \sqcup \text{rlist.}\theta,$$

1116 so the fixpoint equation is `rlist.θ = portable` \sqcup `rlist.θ`. Its least solution is `rlist.θ = portable`.

1117
 1118 ⁷The constants here may be formulas involving abstract kind symbols in scope, but they are constant w.r.t. the $t.i$ is.

1128 **Example: non-uniform parameterized recursion.** Consider the following non-uniform recur-
 1129 sive type from §2.6:

1130 `type 'a nref = Nil of 'a | Cons of 'a ref nref`

1131 The recursive occurrence is non-uniform because it appears at 'a ref. Write

1132 $\text{kind}('a \text{ nref}) = \text{nref}.0 \sqcup (\text{nref}.1 \sqcap \text{kind}('a)).$

1133 Computing the right-hand side kind while keeping nref abstract yields

1134 $\text{kind}('a) \sqcup \text{kind}('a \text{ ref nref}) = \text{nref}.0 \sqcup (\text{nref}.1 \sqcap \text{portable}) \sqcup ((\top \sqcup \text{nref}.1) \sqcap \text{kind}('a)).$

1135 Matching coefficients gives the system

1136 $\text{nref}.0 = \text{nref}.0 \sqcup (\text{nref}.1 \sqcap \text{portable}) \quad \text{nref}.1 = \top \sqcup \text{nref}.1.$

1137 The least solution is $\text{nref}.1 = \top$ and $\text{nref}.0 = \text{portable}$, so 'a nref gets kind **value mod portable**
 1138 with 'a.

1142 4.5 Implementation

1143 The preceding subsections reduce modal-kind inference and sub-kinding to operations on abstract
 1144 lattice polynomials. To implement this efficiently, our implementation uses a new LDD (lattice
 1145 decision diagram) datatype, a variation on BDD-style decision diagrams for lattice formulas rather
 1146 than Boolean formulas [3, 21]. This shared DAG representation supports canonicalization and
 1147 memoization of the normalization and solving steps used by inference. Details of the reduction,
 1148 correctness-oriented formulation, and concrete algorithmic pseudocode for the LDD operations are
 1149 given in the supplementary material [23].

1150 **From OCaml inference to SMOki sub-kinding.** We do not relate this section's algorithm to
 1151 SMOki's semantic sub-kinding judgement $\Delta \models \phi \leq \psi$ directly. This is intentional: this algorithm
 1152 implicitly operates on a representation of types closer to the internal one in OCaml's compiler, as
 1153 opposed to the idealized version in §3. We can, however, assemble the results from this section into
 1154 a sound and complete decision procedure for $\Delta \models \phi \leq \psi$, as described in Appendix B.

1157 5 Mode Crossing in Practice

1158 Mode crossing and modal kinds are essential for practical use of OCaml in Jane Street's codebase
 1159 of 80 million lines of code. Uptake has been fast: as of February 2026, the codebase had roughly 6,000
 1160 modal kind annotations, though the feature has been available (in a preliminary, less principled
 1161 implementation) only since early 2025. These annotations were not added to decorate the code; they
 1162 were added only when needed to fix concrete type errors. Furthermore, our colleagues urgently
 1163 wanted a more expressive kind system than the initial versions, holding up important (closed-source)
 1164 features until the type system was ready.

1165 Here we explore examples inspired by real code where these annotations are necessary.

1167 5.1 Regular Expressions

1168 Consider a regular-expression library with abstract compiled-regex type `t`. A simplified API is:

1169 `type t : value mod portable contended`
 1170 `val create : string → t Or_error.t`
 1171 `val find_all : t → string → string list Or_error.t`
 1172 `val matches : t → string → bool`

1173 Compiled regexes are immutable, so cross-thread use is safe; mode crossing exposes this without
 1174 API noise. Without mode crossing, the signatures would need explicit mode decorations:

1176

```

1177 val create : string → t Or_error.t @ portable
1178 val find_all : t @ contended → string → string list Or_error.t
1179 val matches : t @ contended → string → bool

```

1180 Without these decorations, mode errors would arise when capturing a compiled regular expression
 1181 in a **portable** function. Many other functions in this API and many similarly shaped libraries
 1182 (defining an abstract type with no functions or mutable state) would also need annotations, along
 1183 with any polymorphic functions that might be instantiated with these types. The annotation burden
 1184 would simply be impractical in a large codebase. Instead, mode crossing allows us to avoid this
 1185 work, enabling parallel programming with a lower annotation burden.

1187 5.2 Finite Maps

1188 Recall the functorial Map example from §2.8. There, we aimed to ensure that immutable maps over
 1189 immutable values should cross portability, and we achieved that by requiring the compare function
 1190 used by the map (carried in OrderedType) to be **portable**. This is appropriate for maps meant to cross
 1191 portability, but when we interact with legacy OCaml code, comparison functions might only be
 1192 available at **nonportable** mode, in which case the corresponding maps should *not* cross portability.
 1193 One way to accommodate such maps would be to require them to be constructed using a different
 1194 Map functor, but that would result in completely different types of **portable** and **nonportable** maps.
 1195 Here, we discuss an alternative, unifying interface for maps, deployed at Jane Street, by which the
 1196 distinction between **portable** and **nonportable** maps can be encoded via a *phantom type* parameter,
 1197 thus providing an interesting showcase for how modal kinds enable the expression of novel APIs.

1198 The key idea of this API is to piggyback on a solution to a different problem, namely how to
 1199 differentiate between maps that share the same key type but use different comparison functions. It is
 1200 important that such maps *not* be treated as interoperable since the internal representation invariant
 1201 of the map depends crucially on the comparison function—*e.g.*, merging maps built from different
 1202 comparison functions would result in a broken invariant. Rossberg et al. [24] suggest a technique
 1203 for how to enforce this dependency soundly and systematically by introducing phantom type
 1204 declarations as static representatives of all value declarations, and then using a variant of OCaml’s
 1205 applicative functor types to effectively parameterize the map type by the static representative of
 1206 its corresponding comparison function. Our API achieves the same effect in a less systematic but
 1207 more lightweight manner, via explicit type parameterization rather than applicative functors.

1208 Figure 6 shows an excerpt of the API. First, observe that we index the type of maps not only by
 1209 the types of keys and values that they operate on (as usual) but also by a *comparator witness* type
 1210 'cmp: This 'cmp type is a *phantom type* parameter: it is not associated with runtime data, but merely
 1211 serves as a static representative of the comparison function on which the map type depends. (We
 1212 will return to the with-bounds on this type declaration in a moment.)

1213 Comparator witness phantom types, in turn, are generated by the make functions in the Comparator
 1214 module (Figure 6). These functions return first-class modules [10], so that each returned module
 1215 generates a fresh abstract comparator_witness type. Because this type is abstract, clients are forced
 1216 to distinguish the witnesses from different comparators.

1217 Having explained how we encode the dependency of map types on their comparison functions,
 1218 let us now return to the original problem: how can we ensure that map types using **portable**
 1219 comparison functions cross portability whereas map types using **nonportable** functions do not? .

1220 The trick is to employ *two* make functions, one for **portable** comparators and another for
 1221 **nonportable** comparators. In the former case the freshly generated abstract comparator_witness
 1222 type is declared to cross portability, whereas in the latter case it is not. Then, thanks to the with-
 1223 bounds on Comparator.t and Map.t, the mode-crossing behavior of comparator_witness determines
 1224

1225

```

1226 module Map : sig
1227   type ('key, 'val, 'cmp) t : value mod portable contended
1228     with 'key with 'val with ('key, 'cmp) Comparator.t
1229 end
1230 module Comparator : sig
1231   type ('a, 'cmp) t : value mod portable contended with 'cmp @@ contended
1232   module type S = sig
1233     type element      type comparator_witness
1234     val comparator : (element, comparator_witness) t
1235   end
1236   val make : ('a → 'a → int) → (module S with type element = 'a)
1237   module type S_portable = sig
1238     type element      type comparator_witness : value mod portable
1239     val comparator : (element, comparator_witness) t
1240   end
1241   val make_portable : ('a → 'a → int) @ portable → (module S_portable with type element = 'a)
1242 end

```

Fig. 6. Excerpt of the phantom type-based Map and Comparator APIs.

that of ('key, comparator_witness) Comparator.t and ('key, 'val, comparator_witness) Map.t, ensuring that the map type crosses portability only if the comparison function did.

The implementation of Comparator.t itself is simple:

```

1247 type ('a, 'cmp) t : value mod portable contended with 'cmp @@ contended =
1248   { compare : 'a → 'a → int } [@@unsafe_allow_any_mode_crossing]
1249

```

The payload is just the comparison function. However, as is common with phantom type APIs, the type checker cannot verify the associated invariants: it has no way of knowing that (thanks to module-level data abstraction) the function is guaranteed to be **portable** if the phantom 'cmp parameter crosses portability. Hence, we use the unsafe annotation [@@unsafe_allow_any_mode_crossing] to work around the type system. (See Georges et al. [12] for another example of a modal, phantom type-based API, whose implementation also relies on unsafe code.) In summary, this example shows a more advanced application for modal kinds in a real API; without them we would need either modality-parameterized types or duplicated portable/nonportable variants of the map type.

5.3 Cost of Mode Crossing

The examples in this paper illustrate why mode crossing is useful, but the compiler must now compute crossing information while type checking. The cost of computing modal kinds for type definitions is small, but the Cross typing rule (Fig. 4) is potentially expensive: the compiler repeatedly queries a type's mode-crossing behavior when propagating modes through values, patterns, applications, field accesses, subtyping, and signature inclusion.

We tested the cost of mode crossing in our implementation on plain OCaml files from the standard library and the compiler implementation. They do not need modes to type check, but the OCaml compiler still runs the full mode-crossing machinery on them because mode information is used in later optimization passes. To reduce measurement noise, we added a compiler flag to repeat mode-crossing checks multiple times, and normalized the numbers afterwards. The estimates may slightly understate ordinary compilation cost because repeating the same checks can improve cache and branch-predictor behavior. Even so, on the files we tested, the approximate cost of mode crossing stays below 1% of total native compile time, and about 0–6% of the type-checking phase. The full measurements are shown in Appendix A.

6 Related Work

Modal type systems for safe concurrency in OCaml. OCaml’s mode system, formalized in DRFCaml, establishes memory safety and data-race freedom by enforcing safety restrictions through modes [12, 13, 18]. Our work extends this line of work: we do not change the underlying safety story, but add *mode crossing* and a kind system that relaxes certain modes based on value types. (Specifically, SMOki is an extension of DRFCaml.) This addresses a common source of false negatives: client code need not satisfy modal constraints when the type’s representation and operations make those axes semantically irrelevant. From a technical standpoint, SMOki builds on the DRFCaml logical relation by extending it with mode crossing, iso-recursive types, and the **shareable** mode.

Marker traits and protocols (Rust/Swift). The closest conceptual analogues to modal kinds are *marker traits* such as Rust’s Send and Sync and Swift’s Sendable, classifying types as safe for transfer or sharing across threads [28, 29]. Both languages structurally analyze types to automatically determine whether they satisfy these traits. For instance, Rust can use coinduction to infer whether a (uniform) recursive type implements Sync [27]. We also infer type properties, but our setting is more complex because of abstract types (from ML-style modules) and non-uniform recursive type constructors. In particular, the combination of recursion and abstraction forces us to reason about least fixpoints in the presence of unknown (abstract) kind functions.

Many languages also have “deriving” mechanisms (Haskell’s deriving, Rust’s derive macros, OCaml’s ppx_compare) to derive properties or implementations from type structure. These mechanisms often handle recursive types and, in the case of Haskell, even non-uniform recursive types. However, these mechanisms cannot express interactions with abstract type constructors like we can using with-bounds. Simultaneously, our use of kind functions over lattices helps us solve these more complicated constraints.

Qualified types and type classes. Type qualifiers are a classic framework for tracking type properties, often with constraint-based inference and extensible checking frameworks [9, 22]. Modes and modal kinds resemble constraint-based polymorphism (*qualified types*) and type classes: they record obligations that must hold when instantiating type arguments. Our setting differs: our constraints range over a lattice of mode-crossing behaviors and are discharged by lattice-specific entailment rather than instance resolution, and they interact well with non-uniform recursion and module abstraction, as described above. More broadly, similar phenomena appear in nullable and non-nullable types, universe hierarchies (e.g., Prop vs. Type), and linear/non-linear modalities (e.g., session-typed languages such as GV/Par), though with different semantics than mode crossing.

Boolean and lattice-valued constraint solving for type inference. Our inference procedure reduces kind checking and subsumption to constraints over a distributive lattice, represented as lattice polynomials and solved via a dedicated solver (§4). This connects to prior work relating effect inference to Boolean constraints and Boolean unification [19], and more broadly to type parameters over lattices in Boolean-kinded type systems [30]. Both extend HM-style type systems with algebraic type arguments from *Boolean lattices*. They rely on results in Boolean unification to compute *most general unifiers* on *Boolean formulas*. Our work goes in a different direction: We consider lattice-based properties of types, and not lattice *arguments* to types. Additionally, it is unclear to us whether the unification approach generalizes to Bi-Heyting lattices.

Decision diagrams. Our LDD representation (§4) is analogous to Boolean decision diagrams (BDDs/ZDDs) [3, 21]: a shared canonical DAG, but for lattice polynomials. Our contribution is a solver architecture specialized to modal-kind algebra: co-Heyting-normalized lattice polynomials plus a disciplined fragment of monotone linear functions and fixpoints, sufficient in practice.

Table 1. Mode-crossing cost on selected OCaml files. TC abbreviates type checking.

File	Total (s)	TC (s)	TC/total	Cross (ms)	Cross/TC	Cross/total
<i>Compiler files</i>						
utils/misc.ml	0.735	0.175	23.85%	2.7	1.53%	0.36%
parsing/ast_mapper.ml	0.625	0.120	19.22%	3.4	2.83%	0.54%
lambda/lambda.ml	0.765	0.159	20.72%	5.3	3.34%	0.69%
lambda/simplif.ml	0.619	0.097	15.74%	3.4	3.51%	0.55%
lambda/translcore.ml	1.293	0.169	13.09%	7.3	4.34%	0.57%
lambda/matching.ml	1.628	0.278	17.10%	11.5	4.13%	0.71%
driver/compile_common.ml	0.161	0.042	25.86%	0.3	0.83%	0.21%
typing/typecore.ml	4.531	0.645	14.23%	35.4	5.50%	0.78%
typing/ctype.ml	2.456	0.466	18.97%	23.2	4.99%	0.95%
<i>Standard library files</i>						
stdlib/list.ml	0.197	0.056	28.63%	0.7	1.24%	0.35%
stdlib/array.ml	0.213	0.074	34.78%	0.2	0.21%	0.07%
stdlib/map.ml	0.321	0.088	27.45%	2.2	2.44%	0.67%
stdlib/hashtbl.ml	0.253	0.064	25.44%	1.6	2.49%	0.63%
stdlib/format.ml	0.467	0.088	18.84%	1.8	2.04%	0.38%
stdlib/buffer.ml	0.174	0.055	31.58%	0.4	0.75%	0.24%
stdlib/string.ml	0.135	0.046	33.78%	0.4	0.98%	0.33%

APPENDIX

A Mode-Crossing Benchmark Table

See Table 1 for our measurements on the cost of mode crossing as described in §5.3. We measured the compilation time of different OCaml standard library and compiler files, how much of compilation time was spent type checking, and how much of type checking was spent checking for mode crossing. These measurements are taken over three runs.

To better assess the effect of mode crossing on compilation time, we added a compiler flag that repeats mode-crossing checks N times. We measure total compilation time with `ocamlopt.opt -c`, type checking time with `ocamlopt.opt -i`, and estimate the time for one mode-crossing pass as $(T_{10} - T_1)/9$, where T_N is type checking time with `-mode-crossing-repetitions N`.

On the compiler files, mode crossing takes about 1–6% of overall type checking time; on the standard-library files, it is about 0–3%. In this setup, type checking itself accounts for roughly 13–35% of total native compilation time.

B A Decision Procedure for SMoKi Sub-Kinding

The inference algorithm of §4 does not directly decide the semantic sub-kinding judgment $\Delta \vDash \phi \leq \psi$ of Def. 4, for the reasons discussed in §4.5. Here we show how its results compose into a sound and complete decision procedure for $\Delta \vDash \phi \leq \psi$.

To be precise: we prove that the *overall procedure* underlying §4 is sound and complete for $\Delta \vDash \phi \leq \psi$. This procedure represents each kind by coefficients, reduces the constraints of Δ to coefficient bounds, inlines them, and decides the residual inequality by minimal-normal-form comparison. For some subroutines, though, we establish only that the subproblem is *decidable* rather than verifying the more efficient algorithm §4 uses in its place. The recursive case is the one

1373 such spot: Lemma 12 computes second-order fixpoints by iteration over the finite function lattice,
 1374 whereas §4.4 solves the same fixpoint more efficiently via a coefficient system.

1375 Unfolding Def. 4, the judgment asks whether

$$1376 \quad \forall \rho \in \llbracket \Delta \rrbracket. \llbracket \phi \rrbracket_\rho \leq \llbracket \psi \rrbracket_\rho \quad (\star)$$

1377
 1378 The procedure represents every abstract kind function by coefficients, rewrites both the constraints
 1379 that Δ places on those coefficients and the inequality itself into symbolic form, and discharges the
 1380 result with the minimal-normal-form test of §4.3. Throughout, we treat each type variable α in Δ
 1381 as a nullary abstract constructor t with trivial bound \top .

1382 The reduction rests on a notion of *join-linearity* for lattice functions. Call a function $g : M^k \rightarrow M$
 1383 *join-linear* if it can be written in *coefficient form* $g(m_1, \dots, m_k) = g_0 \sqcup (g_1 \sqcap m_1) \sqcup \dots \sqcup (g_k \sqcap m_k)$
 1384 for some coefficients g_0, \dots, g_k . Recall from §4.3 that a *lattice polynomial* in some variables is a
 1385 function built from those variables and constants by arbitrary meets and joins. We write c for mode
 1386 constants (or, equivalently, constant polynomials without variables), and Φ or Ψ for arbitrary lattice
 1387 polynomials.

1388 LEMMA 12 (COEFFICIENT FORM OF KIND TERMS). *Let χ be a kind term, let $\bar{\alpha}$ be distinguished type*
 1389 *variables, and let ρ be a substitution (in the sense of Def. 2) that assigns a join-linear function with*
 1390 *constant coefficients \bar{c} to every other free variable in χ :*

$$1392 \quad (\rho t)(\bar{m}) = c_0^t \sqcup \bigsqcup_i (c_i^t \sqcap m_i)$$

1394 Then $\bar{m} \mapsto \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\rho$ is join-linear with coefficient form

$$1396 \quad \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\rho = \Phi_0 \sqcup \bigsqcup_i (\Phi_i \sqcap m_i),$$

1398 and each of its coefficients Φ_i can be written as a lattice polynomial in \bar{c} .

1400 PROOF. By induction on χ . Most cases preserve join-linearity and the required coefficient form
 1401 by distributivity. The only case that needs more than distributivity is the second-order μ -fixpoint

$$1403 \quad \chi = (\mu t \bar{\beta}. \phi) \bar{\psi}.$$

1404 Note that we are working in a *finite* lattice, and so that $M^n \rightarrow M$ is also finite. Hence, we can
 1405 compute the second-order fixpoint

$$1407 \quad \bar{m} \mapsto \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\rho = \mu(\lambda f. \lambda \bar{m}'. \llbracket \phi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_{\rho, t \mapsto f, \bar{\beta} \mapsto \bar{m}'}) \overline{\llbracket \psi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_\rho}$$

1409 by iterating from $\bar{m} \mapsto \perp$ until it stabilizes, which takes at most $|M^n \rightarrow M|$ iterations. Each step of
 1410 the iteration preserves join-linearity and the required coefficient form by induction. \square

1411 For instance, for unary abstract t, u with coefficient forms $(\rho t)(\hat{\alpha}) = g_0^t \sqcup (g_1^t \sqcap \hat{\alpha})$ and $(\rho u)(\hat{\alpha}) =$
 1412 $g_0^u \sqcup (g_1^u \sqcap \hat{\alpha})$, the kind term $\hat{u}(\hat{t} \hat{\alpha})$ has coefficient form

$$1414 \quad g_0^u \sqcup (g_1^u \sqcap (g_0^t \sqcup (g_1^t \sqcap \hat{\alpha}))) \rightsquigarrow g_0^u \sqcup (g_1^u \sqcap g_0^t) \sqcup ((g_1^u \sqcap g_1^t) \sqcap \hat{\alpha})$$

1415 with coefficients $g_0^u \sqcup (g_1^u \sqcap g_0^t)$ and $g_1^u \sqcap g_1^t$.

1416 We use Lemma 12 twice, once to bring an arbitrary kind function into coefficient form, and here
 1417 in degenerate form to represent entire substitutions in terms of constant coefficients:

1419 LEMMA 13 (COEFFICIENT REPRESENTATION OF VALID SUBSTITUTIONS). *For every $\rho \in \llbracket \Delta \rrbracket$ and*
 1420 *every $t \in \text{dom } \Delta$, the function ρt is join-linear, with constant coefficients c_0^t, \dots, c_k^t .*

1421

PROOF. By clause (3) of Def. 3, $(\rho t)(\bar{m}) = \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_{\emptyset}$ for a kind term χ with only free variables $\bar{\alpha}$. We apply Lemma 12 to χ with the empty substitution \emptyset , yielding coefficients Φ_i^t such that

$$(\rho t)(\bar{m}) = \llbracket \chi[\bar{\alpha} \mapsto \bar{m}] \rrbracket_{\emptyset} = \Phi_0^t \sqcup \bigsqcup_i (\Phi_i^t \sqcap m_i).$$

Note that the Φ_i^t are constant coefficients c_i^t because we applied Lemma 12 to the empty substitution, and hence the Φ_i^t are polynomials in no variables. \square

With these two lemmas, the reduction proceeds in four steps.

Represent substitutions by coefficients By Lemma 13, each $\rho \in \llbracket \Delta \rrbracket$ is represented by a finite family of constants $\bar{c} = (c_i^t)_{t,i}$:

$$(\rho t)(\bar{m}) = c_0^t \sqcup \bigsqcup_i (c_i^t \sqcap m_i) \quad (1)$$

Quantifying over $\rho \in \llbracket \Delta \rrbracket$ thus becomes quantifying over \bar{c} , subject to the constraints of the next step.

To make sure that the procedure described here is complete, we also need to argue that any such family of constants gives rise to a substitution $\rho \in \llbracket \Delta \rrbracket$. Clause (1) of Def. 3 is trivial, clause (2) is taken care of by the next step, and clause (3) follows because ρt can be represented as a kind term χ such that $(\rho t)(\bar{m}) = \llbracket \chi \rrbracket_{\emptyset}$ by following the structure in (1).

Reduce the constraints on ρ to constraints on \bar{c} For each abstract t with declared bound χ_t , clause (2) of Def. 3 requires

$$\forall \bar{m}. (\rho t)(\bar{m}) \leq \llbracket \chi_t[\bar{\alpha} \mapsto \bar{m}] \rrbracket_{\rho}$$

The left side has a coefficient form by the previous step, and the right side has a coefficient form by Lemma 12, with coefficients that are lattice polynomials Ψ_i^t in \bar{c} :

$$\forall \bar{m}. c_0^t \sqcup \bigsqcup_i (c_i^t \sqcap m_i) \leq \Psi_0^t \sqcup \bigsqcup_i (\Psi_i^t \sqcap m_i)$$

Comparing the two reshaped sides of the equation, Lemma 8 reduces the quantified inequality to finitely many bounds of the form $c_0^t \leq \Psi_0^t$ or $c_i^t \leq \Psi_0^t \sqcup \Psi_i^t$ (for $i \geq 1$). Applying Lemma 12 to ϕ and ψ likewise turns the conclusion $\llbracket \phi \rrbracket_{\rho} \leq \llbracket \psi \rrbracket_{\rho}$ into a lattice-polynomial inequality $\Phi \leq \Psi$ in \bar{c} . So putting it all together, (\star) is equivalent to

$$\forall \bar{c}. \left(\bigwedge_{t,i} c_i^t \leq \Phi_i^t \right) \implies \Phi \leq \Psi$$

for finitely many polynomials Φ_i^t derived from the Ψ_i^t .

Inline the bounds We inline the coefficient bounds iteratively: each application of Lemma 9 discharges one hypothesis $c_i^t \leq \Phi_i^t$ by substituting $c_i^t \mapsto c_i^t \sqcap \Phi_i^t$ until none remain. (Note that Lemma 9 holds even if Φ_i^t mentions c_i^t .) This leaves an unconditional inequality $\forall \bar{c}. \Phi' \leq \Psi'$ between lattice polynomials equivalent to (\star) .

Decide The residual inequality $\forall \bar{c}. \Phi' \leq \Psi'$ is a symbolic query that can be decided by the minimal-normal-form comparison of §4.3.

By composing these different results, we obtain:

Theorem 14 (Decision procedure for SMoKi sub-kinding). *For every well-formed SMoKi context Δ and kind terms ϕ, ψ , the procedure above accepts $\phi \leq \psi$ under Δ exactly when $\Delta \models \phi \leq \psi$ holds.*

References

- [1] Hans Bekić. 1984. Definable Operations in General Algebras, and the Theory of Automata and Flowcharts. In *Programming Languages and Their Definition*. Lecture Notes in Computer Science, Vol. 177. Springer, 30–55. doi:10.1007/BFb0048939
- [2] T. S. Blyth. 2005. *Lattices and Ordered Algebraic Structures*. Springer.
- [3] Randal E. Bryant. 1986. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Trans. Comput.* C-35, 8 (1986). doi:10.1109/TC.1986.1676819
- [4] Stephen Dolan. 2019. Unboxed Types for OCaml. Jane Street Tech Talk. <https://www.janestreet.com/tech-talks/unboxed-types-for-ocaml/>
- [5] Paul Downen, Zena M. Ariola, Simon Peyton Jones, and Richard A. Eisenberg. 2020. Kinds are calling conventions. In *ICFP*. doi:10.1145/3408986
- [6] Richard A. Eisenberg, Stephen Dolan, and Leo White. 2022. Unboxed types for OCaml. In *ML Workshop*. <https://icfp22.sigplan.org/details/mlfamilyworkshop-2022-papers/13/Unboxed-types-for-OCaml>
- [7] Richard A. Eisenberg and Simon Peyton Jones. 2017. Levity polymorphism. In *PLDI*. doi:10.1145/3062341.3062357
- [8] Leo Esakia, Guram Bezhanishvili, Wesley H. Holliday, and Anton Evseev. 2019. *Heyting Algebras: Duality Theory* (1st ed.). Springer.
- [9] Jeffrey S. Foster, Manuel Fähndrich, and Alexander Aiken. 1999. A Theory of Type Qualifiers. In *PLDI*. doi:10.1145/301618.301665
- [10] Alain Frisch and Jacques Garrigue. 2010. First-Class Modules and Composable Signatures in Objective Caml 3.12. In *ACM SIGPLAN Workshop on ML*. Baltimore, MD.
- [11] Lennard Gäher, Michael Sammler, Ralf Jung, Robbert Krebbers, and Derek Dreyer. 2024. RefinedRust: A Type System for High-Assurance Verification of Rust Programs. *PACMPL* 8, PLDI (2024). doi:10.1145/3656422
- [12] Aina Linn Georges, Benjamin Peters, Laila Elbeheiry, Leo White, Stephen Dolan, Richard A. Eisenberg, Chris Casinghino, François Pottier, and Derek Dreyer. 2025. Data Race Freedom à la Mode. *PACMPL* 9, POPL (2025). doi:10.1145/3704859
- [13] Jane Street. 2026. OxCaml. Project website. <https://oxcaml.org/> (accessed February 2, 2026).
- [14] Ralf Jung. 2020. *Understanding and evolving the Rust programming language*. Ph.D. Dissertation. Saarland University. doi:10.22028/D291-31946
- [15] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2017. RustBelt: securing the foundations of the Rust programming language. *Proc. ACM Program. Lang.* 2, POPL, Article 66 (2017), 34 pages. doi:10.1145/3158154
- [16] Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018). <https://doi.org/10.1017/S0956796818000151>
- [17] 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 *POPL*. <https://doi.org/10.1145/2676726.2676980>
- [18] Anton Lorenzen, Leo White, Stephen Dolan, Richard A. Eisenberg, and Sam Lindley. 2024. Oxidizing OCaml with Modal Memory Management. *PACMPL* 8, ICFP (2024). doi:10.1145/3674642
- [19] Magnus Madsen and Jaco van de Pol. 2020. Polymorphic types and effects with Boolean unification. *PACMPL* 4, OOPSLA (2020). doi:10.1145/3428222
- [20] Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2019. Time Credits and Time Receipts in Iris. In *ESOP*. doi:10.1007/978-3-030-17184-1_1
- [21] Shin-ichi Minato. 1993. Zero-Suppressed BDDs for Set Manipulation in Combinatorial Problems. In *DAC*. doi:10.1145/157485.164890
- [22] Matthew M. Papi, Mahmood Ali, Telmo Luis Correa Jr., Jeff H. Perkins, and Michael D. Ernst. 2008. Practical pluggable types for Java. In *ISSTA*. doi:10.1145/1390630.1390656
- [23] Benjamin Peters, Jules Jacobs, Diana Kalinichenko, Liam Stevenson, Aspen Smith, Derek Dreyer, and Richard A. Eisenberg. 2026. Supplementary Materials for Mode Crossing. Zenodo. doi:10.5281/zenodo.20272263
- [24] Andreas Rossberg, Claudio Russo, and Derek Dreyer. 2014. F-ing modules. *Journal of Functional Programming* 24, 5 (2014), 529–607. doi:10.1017/S0956796814000264
- [25] Filip Sieczkowski, Sergei Stepanenko, Jonathan Sterling, and Lars Birkedal. 2024. The Essence of Generalized Algebraic Data Types. *Proc. ACM Program. Lang.* 8, POPL, Article 24 (Jan. 2024), 29 pages. doi:10.1145/3632866
- [26] Thomas Somers, Jonas Kastberg Hinrichsen, Lennard Gäher, and Robbert Krebbers. 2026. Building Blocks for Step-Indexed Program Logics. In *CPP*. doi:10.1145/3779031.3779095
- [27] The Rust Project Developers. 2026. Coinduction. Rust Compiler Development Guide. <https://rustc-dev-guide.rust-lang.org/solve/coinduction.html>
- [28] The Rust Project Developers. 2026. Send and Sync. The Rustonomicon. <https://doc.rust-lang.org/nomicon/send-and-sync.html>

- 1520 [29] The Swift Project Authors. 2026. The Swift Programming Language: Concurrency. Swift.org Documentation.
1521 <https://docs.swift.org/swift-book/documentation/the-swift-programming-language/concurrency/>
- 1522 [30] Joseph A. Zullo. 2026. Let Generalization, Polymorphic Recursion, and Variable Minimization in Boolean-Kinded Type
1523 Systems. *PACMPL* 10, POPL (2026). doi:10.1145/3776644
- 1524
- 1525
- 1526
- 1527
- 1528
- 1529
- 1530
- 1531
- 1532
- 1533
- 1534
- 1535
- 1536
- 1537
- 1538
- 1539
- 1540
- 1541
- 1542
- 1543
- 1544
- 1545
- 1546
- 1547
- 1548
- 1549
- 1550
- 1551
- 1552
- 1553
- 1554
- 1555
- 1556
- 1557
- 1558
- 1559
- 1560
- 1561
- 1562
- 1563
- 1564
- 1565
- 1566
- 1567
- 1568

C Inference Solver Details

Inference reduces kind checking and sub-kinding to constraints on a lattice of mode crossings, solved by a dedicated solver. We present the solver as four layers over a shared base lattice.

Layer 0 (Appendix C.1) The base lattice

Layer 1 (Appendix C.2) Lattice polynomials and operations

Layer 2 (Appendix C.3) Fixpoint constraints

Layer 3 (Appendix C.4) Linear lattice functions

Layer 4 (Appendix C.5) From kinds to constraints

Connection to simplified LDD operations. This section gives the layer-by-layer reduction used by inference. The concrete Layer 1 data structure and operations are documented in Appendix D, which also covers the operations reused by Layer 2 after inlining solved variables.

C.1 Layer 0: The base lattice

We work over a distributive lattice $(L, \leq, \sqcup, \sqcap)$. A *lattice* is a poset in which every pair of elements has a join $a \sqcup b$ (their least upper bound) and a meet $a \sqcap b$ (their greatest lower bound). The lattice is *distributive* when joins and meets distribute over each other:

$$a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c) \quad \text{and} \quad a \sqcup (b \sqcap c) = (a \sqcup b) \sqcap (a \sqcup c).$$

In a finite lattice, the top and bottom elements exist and are given by $\top = \bigsqcup L$ and $\perp = \bigsqcap L$.

Bi-Heyting structure. A lattice is *Heyting* if for all $a, b \in L$ the implication a / b exists and is characterized by the adjunction

$$x \leq (a / b) \quad \text{iff} \quad x \sqcap a \leq b.$$

Equivalently, a / b is the greatest c such that $a \sqcap c \leq b$. Dually, a lattice is *co-Heyting* if for all $a, b \in L$ the subtraction $a \setminus b$ exists and satisfies

$$(a \setminus b) \leq x \quad \text{iff} \quad a \leq b \sqcup x,$$

i.e. $a \setminus b$ is the least c such that $a \leq b \sqcup c$. A *bi-Heyting lattice* is both Heyting and co-Heyting.

Finite distributive lattices are bi-Heyting. Let L be finite and distributive. Define the implication by the finite join

$$a / b = \bigsqcup \{x \in L \mid x \sqcap a \leq b\}.$$

Finiteness ensures the join exists. By distributivity,

$$(a / b) \sqcap a = \bigsqcup_{x \sqcap a \leq b} (x \sqcap a) \leq b,$$

so a / b itself satisfies the defining inequality, and it is the greatest such element because every x in the set is below the join. Similarly, define the subtraction by the finite meet

$$a \setminus b = \bigsqcap \{x \in L \mid a \leq b \sqcup x\}.$$

Using the dual distributive law, we get $a \leq b \sqcup (a \setminus b)$, and minimality follows because the meet is below every x in the defining set. Hence every finite distributive lattice is bi-Heyting.

Conversely. Every Heyting lattice is distributive (and dually for co-Heyting). Indeed, let $d = (a \sqcap b) \sqcup (a \sqcap c)$. The inequality $d \leq a \sqcap (b \sqcup c)$ holds in any lattice. For the other direction, the adjunction gives $b \leq a / d$ and $c \leq a / d$, hence $b \sqcup c \leq a / d$, so $a \sqcap (b \sqcup c) \leq d$. Thus $a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c)$.

Examples. Finite chains. The linear order $L_n = \{0, 1, \dots, n\}$ is a distributive lattice with order $0 \leq 1 \leq \dots \leq n$, join $\sqcup = \max$, and meet $\sqcap = \min$. The implication and subtraction have a simple closed form:

$$a/b = \begin{cases} n & \text{if } a \leq b \\ b & \text{otherwise} \end{cases} \quad \text{and} \quad a \setminus b = \begin{cases} 0 & \text{if } a \leq b \\ a & \text{otherwise.} \end{cases}$$

Intuitively, if $a \leq b$ then $a/b = \top$ because *any* x satisfies $x \sqcap a \leq b$; otherwise the largest x with $x \sqcap a \leq b$ is exactly b . Dually, if $a \leq b$ then $a \setminus b = \perp$, and otherwise the least x with $a \leq b \sqcup x$ is a . For instance in L_3 , $2/1 = 1$ while $1/2 = 3$, and $2 \setminus 1 = 2$ while $1 \setminus 2 = 0$.

Products. If L_1 and L_2 are distributive lattices, then so is $L_1 \times L_2$ with componentwise order, join, and meet; thus $(a_1, a_2) \leq (b_1, b_2)$ iff $a_i \leq b_i$ for each component. Top and bottom are (\top_1, \top_2) and (\perp_1, \perp_2) , and when L_1, L_2 are Heyting/co-Heyting the operations are also componentwise: $(a_1, a_2)/(b_1, b_2) = (a_1/b_1, a_2/b_2)$ and $(a_1, a_2) \setminus (b_1, b_2) = (a_1 \setminus b_1, a_2 \setminus b_2)$. For example in $L_3 \times L_3$, $(2, 1)/(1, 0) = (1, 0)$ and $(2, 1) \setminus (1, 0) = (2, 1)$, while $(1, 0)/(2, 1) = (3, 3)$.

Running concrete lattice. Concrete product-lattice examples below use $L = L_3 \times L_3$.

C.2 Layer 1: Lattice polynomials and operations

We now introduce lattice polynomials, written in algebraic notation to mirror familiar numeric polynomials.

Algebraic notation. To streamline formulas, we will write $a + b$ for the join $a \sqcup b$ and ab (or $a \cdot b$) for the meet $a \sqcap b$. We also write 0 for \perp and 1 for \top , so $a + 0 = a$ and $a1 = a$. With this notation, join and meet are commutative, associative, and idempotent, and meet distributes over join (and vice versa):

$$\begin{aligned} a + b &= b + a & ab &= ba, \\ a + (b + c) &= (a + b) + c & a(bc) &= (ab)c, \\ a + a &= a & aa &= a, \\ a(b + c) &= ab + ac \end{aligned}$$

The dual distributive law

$$a + bc = (a + b)(a + c)$$

also holds in distributive lattices, but we will not use it explicitly. The order is recovered by $a \leq b$ iff $a + b = b$ (equivalently $ab = a$).

Polynomials. We work with polynomials over such a base lattice L , by analogy to ordinary polynomials over numbers, using $+$ for join and multiplication for meet. Polynomials are generated by

$$p, q ::= c \mid x \mid p + q \mid pq$$

for constants $c \in L$ in the base lattice and variables x . An example of such a polynomial is:

$$p(x, y) = (1, 0) + (2, 3)x + (3, 1)y + (1, 2)xy$$

with coefficients $(1, 0)$, $(2, 3)$, $(3, 1)$, and $(1, 2)$ in $L = L_3 \times L_3$.

Normal forms. By repeatedly distributing meet over join and then using commutativity and associativity to flatten, any term from the grammar can be expanded into the form above: a finite sum of monomials of the form $cx_1 \cdots x_k$ (with $c \in L$) where the variables are distinct. Duplicate variables are removed using idempotence, explicitly $x \cdot x = x$, and duplicate monomials (same variable set) can be combined by joining their coefficients. *This expansion is not a normal form:* for instance, $x + xy$ and x denote the same polynomial (since $xy \leq x$, so $x + xy = x$), yet they are

1667 syntactically distinct. Moreover, even after removing absorbed monomials, coefficients are not
 1668 canonical: $(1, 0)x + (0, 1)xy$ equals $(1, 0)x + (1, 1)xy$ because $(1, 1)xy = (1, 0)xy + (0, 1)xy$ and the
 1669 extra $(1, 0)xy$ is absorbed by $(1, 0)x$. However, the $(1, 1)xy$ term cannot be removed entirely. The
 1670 following develops two normal forms that are canonical and correct.

1671 **Maximal normal form.** We define the *maximal normal form* for a multivariate polynomial
 1672 $p(x_1, \dots, x_n)$ as

$$1673 \sum_{S \subseteq \{1, \dots, n\}} p(\chi_S) \prod_{i \in S} x_i,$$

1674 where χ_S assigns $x_i = 1$ for $i \in S$ and $x_i = 0$ otherwise. Here sums and products are join and meet,
 1675 with the empty sum 0 and empty product 1. For the bivariate case this specializes to

$$1676 p(0, 0) + p(1, 0)x + p(0, 1)y + p(1, 1)xy.$$

1677 This is a normal form for two reasons:

1678 (A) *Correctness.* Without loss of generality write $p(x, y)$ in multilinear form as $c_{00} + c_{10}x + c_{01}y +$
 1679 $c_{11}xy$ using distributivity, commutativity, associativity, and idempotence. Then

$$1680 \begin{aligned} 1681 p(0, 0) &= c_{00} & p(1, 0) &= c_{00} + c_{10} \\ 1682 p(0, 1) &= c_{00} + c_{01} & p(1, 1) &= c_{00} + c_{10} + c_{01} + c_{11}, \end{aligned}$$

1683 so substituting into the normal form and expanding with distributivity yields terms like $c_{00}x$
 1684 and $c_{10}xy$, which are absorbed by c_{00} and $c_{10}x$ respectively, leaving $c_{00} + c_{10}x + c_{01}y + c_{11}xy$.

1685 (B) *Canonicity.* Semantically equivalent polynomials agree on $(0, 0)$, $(1, 0)$, $(0, 1)$, and $(1, 1)$, so
 1686 they produce the same normal form.

1687 *Examples.*

1688 (a) *Bivariate formula.* The definition specializes to

$$1689 p(0, 0) + p(1, 0)x + p(0, 1)y + p(1, 1)xy.$$

1690 (b) *Worked example.* For $p(x, y) = (1, 0)x + (0, 1)y$ we have $p(0, 0) = (0, 0)$, $p(1, 0) = (1, 0)$,
 1691 $p(0, 1) = (0, 1)$, and $p(1, 1) = (1, 1)$, so the maximal normal form is $(1, 0)x + (0, 1)y + (1, 1)xy$.

1692 **Minimal normal form.** The *minimal normal form* uses the same basis of monomials, but
 1693 replaces each coefficient $p(\chi_S)$ by a co-Heyting subtraction of all smaller terms:

$$1694 \sum_{S \subseteq \{1, \dots, n\}} \left(p(\chi_S) \setminus \sum_{S' \subset S} p(\chi_{S'}) \right) \prod_{i \in S} x_i.$$

1695 This uses the co-Heyting subtraction \setminus on the base lattice. *Examples.*

1696 (a) *Bivariate formula.* The definition specializes to

$$1697 \begin{aligned} 1698 p(0, 0) &+ (p(1, 0) \setminus p(0, 0))x + (p(0, 1) \setminus p(0, 0))y \\ 1699 &+ (p(1, 1) \setminus (p(1, 0) + p(0, 1) + p(0, 0)))xy. \end{aligned}$$

1700 (b) *Coefficient reduced.* For $p(x, y) = (1, 0)x + (1, 1)xy$, the xy coefficient reduces to $(1, 1) \setminus (1, 0) =$
 1701 $(0, 1)$, giving minimal normal form $(1, 0)x + (0, 1)xy$.

1702 (c) *Term removed.* For $p(x, y) = (1, 0)x + (1, 0)xy$, we have $p(1, 1) = (1, 0)$ and $(1, 0) \setminus (1, 0) =$
 1703 $(0, 0)$, so the xy term disappears and the minimal normal form is just $(1, 0)x$.

1704 (d) *Term removed (mixed).* For $p(x, y) = (0, 1)x + (1, 0)y + (1, 1)xy$, we have $p(1, 0) = (0, 1)$,
 1705 $p(0, 1) = (1, 0)$, and $p(1, 1) = (1, 1)$, so the xy coefficient reduces to $(1, 1) \setminus ((0, 1) + (1, 0)) =$
 1706 $(0, 0)$, and the minimal normal form is $(0, 1)x + (1, 0)y$.

1707

(e) *Coefficient reduced (mixed)*. In $L_3 \times L_3$, for $p(x, y) = (2, 0)x + (0, 3)y + (3, 3)xy$, we have $p(1, 0) = (2, 0)$, $p(0, 1) = (0, 3)$, and $p(1, 1) = (3, 3)$, so the xy coefficient reduces to $(3, 3) \setminus ((2, 0) + (0, 3)) = (3, 0)$, giving minimal normal form $(2, 0)x + (0, 3)y + (3, 0)xy$.

(f) *No xy term*. For $p(x, y) = (1, 0)x + (0, 1)y$ we have $p(0, 0) = (0, 0)$, $p(1, 0) = (1, 0)$, $p(0, 1) = (0, 1)$, $p(1, 1) = (1, 1)$, so the minimal normal form is $(1, 0)x + (0, 1)y$.

Pointwise order. For polynomials p, q , we write $p \leq q$ iff $p(\rho) \leq q(\rho)$ for every valuation ρ of the variables into L .

LEMMA 15 (ONE-SIDED MAXIMAL NORMAL FORM SUFFICES). *Let*

$$p = \sum_{S \subseteq \{1, \dots, n\}} c_S \prod_{i \in S} x_i \quad q = \sum_{S \subseteq \{1, \dots, n\}} d_S \prod_{i \in S} x_i,$$

where missing monomials have coefficient 0. Assume only that q is in maximal normal form, i.e. $d_S = q(\chi_S)$ for all $S \subseteq \{1, \dots, n\}$. Then $p \leq q$ iff $c_S \leq d_S$ for all S .

PROOF. For any S , evaluation at χ_S gives

$$p(\chi_S) = \sum_{T \subseteq S} c_T,$$

since $x_T(\chi_S) = 1$ iff $T \subseteq S$ and 0 otherwise. Hence $c_S \leq p(\chi_S)$.

(\Leftarrow) Assume $c_S \leq d_S$ for all S . For any valuation ρ , monotonicity of meet gives $c_S \prod_{i \in S} \rho(x_i) \leq d_S \prod_{i \in S} \rho(x_i)$, and taking the join over all S preserves inequality, so $p(\rho) \leq q(\rho)$. Thus $p \leq q$.

(\Rightarrow) Assume $p \leq q$. Fix $S \subseteq \{1, \dots, n\}$ and show $c_S \leq d_S$.

1. From the corner-evaluation formula above, $p(\chi_S) = \sum_{T \subseteq S} c_T$, so in particular $c_S \leq p(\chi_S)$ because c_S is one of the join-summands.
2. By the pointwise assumption $p \leq q$, we have $p(\chi_S) \leq q(\chi_S)$.
3. By maximality of q , $q(\chi_S) = d_S$.

Chaining these yields $c_S \leq d_S$, and since S was arbitrary we get coefficientwise inequality for all S . \square

If q is not maximal, coefficientwise comparison can fail: $x + xy \leq x$ but the right-hand side has $d_{\{1,2\}} = 0$. Also note that coefficientwise comparison on minimal normal forms is unsound: $p(x, y) = xy$ and $q(x, y) = x + y$ are both minimal and satisfy $p \leq q$, but the xy coefficient is 1 for p and 0 for q .

Operations and interface. The first solver layer exposes a small set of operations on polynomials:

- *Introduce constants and rigid variables.* Constants $c \in L$ and rigid variables $r \in R$ can be used as polynomials directly.
- *Combine by addition and multiplication.* Polynomials are closed under $p + q$ and pq .
- *Create a fresh flexible variable.* We may introduce a fresh flexible variable $x \in X$ to stand for an unknown subterm.
- *Order testing for ground polynomials.* When p and q contain no unsolved flexible variables, we check $p \leq q$ by computing the residual $p \setminus q$ and testing whether its rounded-up value is \perp .
- *Co-Heyting subtraction.* We can compute $p \setminus q$ directly in the DAG representation; this is the core normalizing operation.
- *Rounding to a lattice element.* Given a ground polynomial p , we can compute its least upper approximation in L , written $\lceil p \rceil$, by setting all rigid variables to 1.

1765 Here $p \setminus q$ is co-Heyting subtraction in the lattice of positive lattice polynomials represented by
 1766 the DAGs: it is the least positive polynomial r such that $p \leq q + r$. It is not pointwise base-lattice
 1767 subtraction. Thus

$$1768 \quad p \leq q \iff p \setminus q = 0 \iff \lceil p \setminus q \rceil = \perp.$$

1769 The last equivalence uses monotonicity: a positive polynomial is 0 exactly when its value with all
 1770 rigid variables set to 1 is \perp . Operationally, the reference implementation checks

$$1771 \quad p \leq q \iff \lceil p \setminus q \rceil = \perp$$

1772
 1773 by computing $\text{diff} = \text{sub_subsets}(\text{inline_solved_vars } p) (\text{inline_solved_vars } q)$ and
 1774 then $\text{witness} = \text{round_up } \text{diff}$; the check succeeds exactly when witness is bottom. These
 1775 operations form the interface of the first layer; later layers will encode kind inference and subsump-
 1776 tion constraints into this language. Fixpoint assignments are handled in Layer 2. The simplified
 1777 reference implementation of these operations (including `canonicalize`, `sub_subsets`, `round_up`,
 1778 and `leq_with_reason`) is documented in Appendix D.

1779 C.2.1 Efficient representation in minimal normal form.

1781 **Representation.** The solver represents polynomials in a form reminiscent of binary decision
 1782 diagrams (BDDs). We pick an order on the variables $v \in R \cup X$, and represent polynomials as a tree
 1783 (or DAG) on these variables:

$$1784 \quad p, q ::= c \mid p + (vq) \quad \text{where } c \in L \text{ and } v \in R \cup X$$


1785 i.e.,

$$1786 \quad \text{type poly} = \text{Leaf of } L \mid \text{Node of } (R \cup X) \times \text{poly} \times \text{poly}$$

1788 Leaves of the DAG are constants $c \in L$, and a node for variable v has two children: p and q , and
 1789 represents the polynomial $p + vq$. We draw the p -edge as a dashed edge and the q -edge as a solid
 1790 edge labeled by the variable.

1791 Each root-to-leaf path contributes a product term: when a path follows the q -edge out of a
 1792 variable node, we include that variable in the product, and then multiply the result with the leaf
 1793 constant. The polynomial is the sum of all these path terms.

1794 **Example.** For the base lattice $L = L_3 \times L_3$ and variable order $v_1 < v_2 < v_3$, consider the following
 1795 DAG with a shared subtree:



1808 This DAG denotes the polynomial

$$1809 \quad t + v_1(t + v_2(2, 3)) \quad \text{where } t = (0, 0) + v_3(1, 2).$$

1810 For the DAG above, the path terms are:

$$1811 \quad (0, 0), \quad (v_3(1, 2)), \quad (v_1(0, 0)), \quad (v_1v_3(1, 2)), \quad (v_1v_2(2, 3)).$$

1813

1814 **DAG normal form.** The solver maintains these polynomial DAGs in a normal form that corre-
1815 sponds to minimal normal form with respect to the chosen variable order:

1816 **Variable Order:** When walking down the DAG, variables always appear according to the chosen
1817 variable order: the variables of children are always later in the order than the variable of
1818 the parent. Note that different paths down the DAG may visit different sets of variables.

1819 **Split Irredundancy:** The q -child is never 0. If it is, the split is redundant and we can remove the
1820 node.

1821 **Normalization:** The q -child is irredundant with the p -child according to co-Heyting subtraction:
1822 $q \setminus p = q$.

1823 **Normalization example.** In the running product lattice $L = L_3 \times L_3$ with componentwise order,
1824 consider the two terms

$$(1, 0)v_1 \quad \text{and} \quad (1, 2)v_1v_2.$$

1827 Since $v_1v_2 \leq v_1$, the second term is only relevant on the subregion where the first already applies,
1828 so normalization subtracts the overlap by replacing the coefficient $(1, 2)$ with

$$(1, 2) \setminus (1, 0) = (0, 2),$$

1830 yielding the normalized polynomial

$$(1, 0)v_1 + (0, 2)v_1v_2.$$

1833 If the coefficients are equal, e.g. $(1, 0)v_1$ and $(1, 0)v_1v_2$, then subtraction yields $(1, 0) \setminus (1, 0) = 0$ and
1834 the more specific term disappears; in the DAG this corresponds to a node whose q -child is 0, so the
1835 split is removed by split irredundancy.

1836 **Canonicity.** For a fixed variable order, this normal form is canonical: if two polynomials in
1837 normal form are equal for all valuations of the variables, then they are syntactically identical (and
1838 conversely).

1839 **Operations on DAGs.** Multiplication, addition, and co-Heyting subtraction are defined by
1840 recursion on the tree structure, followed by normalization to restore the invariants above.

1843 C.3 Layer 2: Fixpoint constraints

1844 This layer adds fixpoint assignments on top of lattice polynomials.

1846 *Interface.*

- 1847 • *Assign a flexible variable by a fixpoint.* We allow asserting the value of a flexible variable by
1848 $x := p$ where p is a polynomial. Because p may itself contain x , we have two variants with
1849 different fixpoint equations:

$$1850 \quad x :=_{\text{lfp}} p \quad \text{and} \quad x :=_{\text{gfp}} p$$

1852 which take the least or greatest fixpoint of the induced monotone function.

1853 Operationally, solving is done in two phases. First we force solved substitutions in the right-hand
1854 side using `inline_solved_vars`; this avoids reasoning about stale references to already-solved
1855 flexible variables. Then we install the solved definition of the assigned variable. In the simplified
1856 reference solver, this assignment step is:

$$1857 \quad x :=_{\text{lfp}} p : x \mapsto p[x := 0] \quad \text{and} \quad x :=_{\text{gfp}} p : x \mapsto p[x := 1],$$

1858 implemented by `assign_bot` and `assign_top` after inlining. This one-step rule is valid because p
1859 is a positive lattice polynomial. Viewed as a polynomial in x , it has the form $A + Bx$ for polynomials
1860 A and B that do not mention x . Hence $\mu x. p(x) = A + p[0/x]$ and $\nu x. p(x) = A + B + p[1/x]$.

1863 For systems of variables, the solver applies this rule by sequential elimination in the usual Bekic
 1864 style. This is the same solver interface used by the optimized implementation; the appendix gives
 1865 pseudocode and discusses pending GFP queues and forcing points (Appendix D).

1866 C.4 Layer 3: Linear lattice functions

1868 Layer 3 lifts Layer 1 from lattice elements to join-linear functions between products of lattices. We
 1869 use the form

$$f(x_1, \dots, x_n) = c_0 + c_1x_1 + \dots + c_nx_n,$$

1870 where each coefficient c_i is itself a Layer 1 polynomial (possibly over rigid variables). Coefficients are
 1872 extracted from corner evaluations. With e_i the valuation that maps $x_i \mapsto 1$ and all other arguments
 1873 to 0, we have:

$$\begin{aligned} 1875 \quad c_0 &= f(0, 0, \dots, 0), \\ 1876 \quad c_i &= f(e_i) \setminus c_0 \quad (1 \leq i \leq n). \end{aligned}$$

1877 Equivalently, one may use the redundant maximal coefficients

$$\begin{aligned} 1879 \quad d_0 &= f(0, 0, \dots, 0), \\ 1880 \quad d_i &= f(e_i), \end{aligned}$$

1881 with $d_i = c_0 + c_i$. We use this redundant view when comparing functions: after rewriting the
 1882 right-hand side into maximal normal form, Layer 1's coefficientwise comparison lemma applies
 1883 directly.

1884 Fixpoints at this layer reduce to Layer 2 by representing function coefficients as Layer 1 polyno-
 1885 mials and solving their assignment constraints. The next subsection presents the construction for
 1886 one ternary function symbol; the implementation uses the same construction for arbitrary arities,
 1887 with one coefficient per basis point. Other abstract type constructors in scope are treated as fixed
 1888 coefficient parameters while solving the distinguished recursive function.

1890 C.4.1 Coefficient-based fixpoints for linear lattice terms.

1891 C.4.2 *Setting.* Let $(L, \leq, \sqcup, \sqcap, \perp, \top)$ be a *bounded distributive lattice*. We order function spaces
 1892 pointwise: for $g, h : L^3 \rightarrow L$, $g \leq h$ iff $\forall (x, y, z) \in L^3$, $g(x, y, z) \leq h(x, y, z)$.

1893 **Definition 16** (Linear functions). A function $f : L^3 \rightarrow L$ is *linear* (in the sense used in this paper)
 1894 if there exist coefficients $c_0, c_1, c_2, c_3 \in L$ such that

$$1897 \quad f(x, y, z) = c_0 \sqcup (c_1 \sqcap x) \sqcup (c_2 \sqcap y) \sqcup (c_3 \sqcap z). \quad (2)$$

1898 Write $\text{Lin}_3(L)$ for the set of all such functions.

1899 C.4.3 *Canonical coefficient representation.* A naive coefficient tuple (c_0, c_1, c_2, c_3) for (2) is not
 1900 unique. For fixpoint computation it is convenient to work with a canonical representation.

1901 Define the *basis points*

$$1902 \quad p_0 = (\perp, \perp, \perp), \quad p_1 = (\top, \perp, \perp), \quad p_2 = (\perp, \top, \perp), \quad p_3 = (\perp, \perp, \top).$$

1903 **Definition 17** (Coefficient domain, abstraction, concretization). Define

$$1904 \quad \text{Coef}_3(L) = \{(a_0, a_1, a_2, a_3) \in L^4 \mid a_0 \leq a_1, a_0 \leq a_2, a_0 \leq a_3\},$$

1905 ordered componentwise.

1906 Define $\gamma : \text{Coef}_3(L) \rightarrow \text{Lin}_3(L)$ by

$$1907 \quad \gamma(a_0, a_1, a_2, a_3)(x, y, z) = a_0 \sqcup (a_1 \sqcap x) \sqcup (a_2 \sqcap y) \sqcup (a_3 \sqcap z). \quad (3)$$

1911

Define $\delta : \text{Lin}_3(L) \rightarrow \text{Coef}_3(L)$ by evaluation on basis points:

$$\delta(f) = (f(p_0), f(p_1), f(p_2), f(p_3)). \quad (4)$$

LEMMA 18 (NORMALIZATION AND UNIQUENESS VIA BASIS POINTS). *For every $f \in \text{Lin}_3(L)$ we have $\delta(f) \in \text{Coef}_3(L)$ and*

$$\gamma(\delta(f)) = f.$$

Conversely, for every $a \in \text{Coef}_3(L)$,

$$\delta(\gamma(a)) = a.$$

Hence γ and δ are mutually inverse order-isomorphisms between $\text{Coef}_3(L)$ and $\text{Lin}_3(L)$.

PROOF. Let f be linear, so f has the form (2). Evaluating at p_0 gives $f(p_0) = c_0$. Evaluating at p_1 gives $f(p_1) = c_0 \sqcup (c_1 \sqcap \top) = c_0 \sqcup c_1$. Similarly $f(p_2) = c_0 \sqcup c_2$ and $f(p_3) = c_0 \sqcup c_3$. Thus $\delta(f) = (c_0, c_0 \sqcup c_1, c_0 \sqcup c_2, c_0 \sqcup c_3) \in \text{Coef}_3(L)$.

Now compute $\gamma(\delta(f))$:

$$\gamma(\delta(f))(x, y, z) = c_0 \sqcup ((c_0 \sqcup c_1) \sqcap x) \sqcup ((c_0 \sqcup c_2) \sqcap y) \sqcup ((c_0 \sqcup c_3) \sqcap z).$$

Using distributivity, $(c_0 \sqcup c_1) \sqcap x = (c_0 \sqcap x) \sqcup (c_1 \sqcap x)$. After distributing similarly for y, z and using absorption $c_0 \sqcup (c_0 \sqcap x) = c_0$ (and likewise for y, z), we obtain $\gamma(\delta(f))(x, y, z) = c_0 \sqcup (c_1 \sqcap x) \sqcup (c_2 \sqcap y) \sqcup (c_3 \sqcap z) = f(x, y, z)$.

Conversely, let $a = (a_0, a_1, a_2, a_3) \in \text{Coef}_3(L)$. Then $\gamma(a)(p_0) = a_0$. Moreover $\gamma(a)(p_1) = a_0 \sqcup (a_1 \sqcap \top) = a_0 \sqcup a_1 = a_1$ because $a_0 \leq a_1$. Similarly $\gamma(a)(p_2) = a_2$ and $\gamma(a)(p_3) = a_3$. Hence $\delta(\gamma(a)) = a$.

Monotonicity of γ and δ under componentwise/pointwise orders is immediate, and mutual inverses yield an order-isomorphism. \square

C.4.4 *A grammar of lattice expressions with a function symbol.* Fix a constant set of lattice elements $K \subseteq L$ (the set of allowed constants). We consider expressions generated by

$$e ::= k \mid x \mid y \mid z \mid f(e_1, e_2, e_3) \mid (e_1 \sqcup e_2) \mid (e \sqcap k), \quad (k \in K).$$

Given a function $g : L^3 \rightarrow L$, the standard denotation $\llbracket e \rrbracket_g : L^3 \rightarrow L$ is defined as usual:

$$\begin{aligned} \llbracket k \rrbracket_g(x, y, z) &= k, \\ \llbracket x \rrbracket_g(x, y, z) &= x, \quad \llbracket y \rrbracket_g(x, y, z) = y, \quad \llbracket z \rrbracket_g(x, y, z) = z, \\ \llbracket e_1 \sqcup e_2 \rrbracket_g &= \llbracket e_1 \rrbracket_g \sqcup \llbracket e_2 \rrbracket_g, \quad \llbracket e \sqcap k \rrbracket_g = \llbracket e \rrbracket_g \sqcap k, \\ \llbracket f(e_1, e_2, e_3) \rrbracket_g(x, y, z) &= g(\llbracket e_1 \rrbracket_g(x, y, z), \llbracket e_2 \rrbracket_g(x, y, z), \llbracket e_3 \rrbracket_g(x, y, z)). \end{aligned}$$

C.4.5 *Coefficient semantics (four-point evaluation).* The key observation is that, when g is linear, every expression denotes a linear function, and it suffices to track the four basis-point values.

Definition 19 (Coefficient semantics). For $a \in \text{Coef}_3(L)$, define $\llbracket e \rrbracket_a^\# \in \text{Coef}_3(L)$ by structural recursion:

Variables.

$$\llbracket x \rrbracket_a^\# = (\perp, \top, \perp, \perp), \quad \llbracket y \rrbracket_a^\# = (\perp, \perp, \top, \perp), \quad \llbracket z \rrbracket_a^\# = (\perp, \perp, \perp, \top).$$

Constants.

$$\llbracket k \rrbracket_a^\# = (k, k, k, k).$$

Join and meet-with-constant (componentwise). If $u = (u_0, u_1, u_2, u_3)$ and $v = (v_0, v_1, v_2, v_3)$, define $u \sqcup^4 v = (u_0 \sqcup v_0, \dots, u_3 \sqcup v_3)$ and $u \sqcap^4 k = (u_0 \sqcap k, \dots, u_3 \sqcap k)$. Then

$$\llbracket e_1 \sqcup e_2 \rrbracket_a^\# = \llbracket e_1 \rrbracket_a^\# \sqcup^4 \llbracket e_2 \rrbracket_a^\#, \quad \llbracket e \sqcap k \rrbracket_a^\# = \llbracket e \rrbracket_a^\# \sqcap^4 k.$$

1961 *Application.* Let $b^i = \llbracket e_i \rrbracket_a^\# = (b_0^i, b_1^i, b_2^i, b_3^i)$ for $i \in \{1, 2, 3\}$. Write $a = (a_0, a_1, a_2, a_3)$. Define
 1962 $\text{app}(a, b^1, b^2, b^3) \in L^4$ componentwise by

$$1963 \quad (\text{app}(a, b^1, b^2, b^3))_j = a_0 \sqcup (a_1 \sqcap b_j^1) \sqcup (a_2 \sqcap b_j^2) \sqcup (a_3 \sqcap b_j^3), \quad j \in \{0, 1, 2, 3\}. \quad (5)$$

1964 Set

$$1965 \quad \llbracket f(e_1, e_2, e_3) \rrbracket_a^\# = \text{app}(a, b^1, b^2, b^3).$$

1966 LEMMA 20 (SOUNDNESS OF COEFFICIENT SEMANTICS). *For all $a \in \text{Coef}_3(L)$ and all expressions e ,*

$$1968 \quad \gamma(\llbracket e \rrbracket_a^\#) = \llbracket e \rrbracket_{\gamma(a)}. \quad (6)$$

1969 Equivalently, $\llbracket e \rrbracket_a^\# = \delta(\llbracket e \rrbracket_{\gamma(a)})$.

1970 PROOF. By structural induction on e .

1971 Case $e = k$. Immediate from the definitions: $\gamma(k, k, k, k)$ is the constant function with value k .

1972 Case $e \in \{x, y, z\}$. Immediate from (3) and the definitions in Definition 19.

1973 Case $e = e_1 \sqcup e_2$. By induction hypothesis, $\gamma(\llbracket e_i \rrbracket_a^\#) = \llbracket e_i \rrbracket_{\gamma(a)}$ for $i = 1, 2$. Since γ is defined by
 1974 joins/meets and \sqcup^4 is componentwise, $\gamma(u \sqcup^4 v) = \gamma(u) \sqcup \gamma(v)$ pointwise. Hence

$$1975 \quad \gamma(\llbracket e \rrbracket_a^\#) = \gamma(\llbracket e_1 \rrbracket_a^\# \sqcup^4 \llbracket e_2 \rrbracket_a^\#) = \gamma(\llbracket e_1 \rrbracket_a^\#) \sqcup \gamma(\llbracket e_2 \rrbracket_a^\#) = \llbracket e_1 \rrbracket_{\gamma(a)} \sqcup \llbracket e_2 \rrbracket_{\gamma(a)} = \llbracket e \rrbracket_{\gamma(a)}.$$

1976 Case $e = e' \sqcap k$. Similarly, componentwise meets satisfy $\gamma(u \sqcap^4 k) = \gamma(u) \sqcap k$ pointwise, using
 1977 distributivity of \sqcap over \sqcup in L .

1978 Case $e = f(e_1, e_2, e_3)$. Let $b^i = \llbracket e_i \rrbracket_a^\#$ and by induction hypothesis $\gamma(b^i) = \llbracket e_i \rrbracket_{\gamma(a)}$. Let $g = \gamma(a)$.
 1979 Then for any $(x, y, z) \in L^3$,

$$1980 \quad \llbracket e \rrbracket_g(x, y, z) = g(\gamma(b^1)(x, y, z), \gamma(b^2)(x, y, z), \gamma(b^3)(x, y, z)).$$

1981 Expanding $g = \gamma(a)$ using (3) and expanding each $\gamma(b^i)$ again using (3), we obtain an expression
 1982 built from \sqcup, \sqcap . Distributivity allows pushing each meet by a_1, a_2, a_3 through the joins coming from
 1983 the arguments. After regrouping the constant part and the x -, y -, and z -parts, the resulting function
 1984 is exactly $\gamma(c_0, c_1, c_2, c_3)$ where $(c_0, c_1, c_2, c_3) = \text{app}(a, b^1, b^2, b^3)$ as defined in (5). Thus

$$1985 \quad \llbracket e \rrbracket_{\gamma(a)} = \gamma(\llbracket e \rrbracket_a^\#),$$

1986 as required. □

1987 COROLLARY 21 (CLOSURE). *If $a \in \text{Coef}_3(L)$ then $\gamma(\llbracket e \rrbracket_a^\#)$ is linear, i.e., $\llbracket e \rrbracket_{\gamma(a)} \in \text{Lin}_3(L)$. In
 1988 particular, every expression in the grammar denotes a linear function when f is linear.*

1989 C.4.6 Coefficient transformer and fixpoint transfer. Fix an expression E in the grammar and define
 1990 the functional $F : \text{Lin}_3(L) \rightarrow \text{Lin}_3(L)$ by

$$1991 \quad F(g) = \llbracket E \rrbracket_g.$$

1992 (When needed for fixpoints, we assume F is monotone; this holds for the above grammar because
 1993 all constructors are monotone.)

1994 Define the induced coefficient transformer

$$1995 \quad \widehat{F} : \text{Coef}_3(L) \rightarrow \text{Coef}_3(L), \quad \widehat{F}(a) = \llbracket E \rrbracket_a^\#. \quad (7)$$

1996 LEMMA 22 (CONJUGACY). *For all $a \in \text{Coef}_3(L)$,*

$$1997 \quad F(\gamma(a)) = \gamma(\widehat{F}(a)).$$

1998 Equivalently, $\widehat{F} = \delta \circ F \circ \gamma$.

1999 PROOF. Immediate from Lemma 20 with $e = E$. □

2000

Theorem 23 (Fixpoint computation over coefficients). *Assume L is a complete lattice (so $\text{Coef}_3(L)$ and $\text{Lin}_3(L)$ are complete under the respective orders) and F is monotone. Then both least fixpoints exist and satisfy*

$$\mu F = \gamma(\widehat{\mu F}).$$

PROOF. By Lemma 18, $\gamma : \text{Coef}_3(L) \rightarrow \text{Lin}_3(L)$ is an order-isomorphism with inverse δ . By Lemma 22, $F \circ \gamma = \gamma \circ \widehat{F}$, i.e. F and \widehat{F} are conjugate through the isomorphism. Order-isomorphisms preserve least fixpoints of monotone maps (apply γ to the Knaster–Tarski characterization of μ as the least pre-fixpoint, or equivalently transport the complete lattice structure). Hence $\mu F = \gamma(\widehat{\mu F})$. \square

C.4.7 Operational computation (Kleene iteration). If L has finite height (or more generally if Kleene iteration stabilizes), $\widehat{\mu F}$ can be computed by iterating in $\text{Coef}_3(L)$:

$$a^{(0)} = (\perp, \perp, \perp, \perp), \quad a^{(n+1)} = \widehat{F}(a^{(n)}).$$

Upon stabilization $a^{(n+1)} = a^{(n)} = a^*$, we return $f_{\text{fix}} = \gamma(a^*)$. This is the generic mathematical computation for the coefficient transformer. The implementation uses the specialized Layer 2 polynomial solver described above, which eliminates coefficients one at a time instead of iterating the whole function transformer.

C.4.8 Remark: n -ary generalization. All definitions extend to arity n by using $n+1$ basis points $p_0 = (\perp, \dots, \perp)$ and p_i with \top in coordinate i and \perp elsewhere. Then $\text{Coef}_n(L) \subseteq L^{n+1}$ consists of tuples (a_0, \dots, a_n) with $a_0 \leq a_i$, and $\gamma(a)(x_1, \dots, x_n) = a_0 \sqcup \bigsqcup_{i=1}^n (a_i \sqcap x_i)$. The coefficient semantics evaluates expressions on these basis points and the fixpoint transfer theorem remains unchanged.

C.5 Layer 4: From kinds to constraints

We now reduce judgments from §3 to Layer 3 constraints. Given a type context Δ , we introduce one unknown join-linear function F_t for each abstract type constructor $t \in \text{dom}(\Delta)$. Lattice terms are translated compositionally:

$$\llbracket \hat{\alpha} \rrbracket = \alpha,$$

$$\llbracket m \rrbracket = m,$$

$$\llbracket \phi \sqcup \psi \rrbracket = \llbracket \phi \rrbracket + \llbracket \psi \rrbracket,$$

$$\llbracket \phi \sqcap m \rrbracket = \llbracket \phi \rrbracket m,$$

$$\llbracket \hat{t} \bar{\phi} \rrbracket = F_t(\llbracket \bar{\phi} \rrbracket),$$

$$\llbracket (\mu \hat{t} \bar{\alpha}. \phi) \bar{\psi} \rrbracket = (\mu X. \lambda \bar{a}. \llbracket \phi \rrbracket [\hat{t} \mapsto X, \hat{\alpha}_i \mapsto a_i]) (\llbracket \bar{\psi} \rrbracket).$$

Here X is a fresh unknown function symbol replacing \hat{t} in the body. The fixpoint is taken at the function level and is applied to $\llbracket \bar{\psi} \rrbracket$ only after solving the recursive function; the substitutions $\hat{\alpha}_i \mapsto a_i$ are simultaneous. For each declaration $(t : \bar{\alpha}. \phi_t) \in \Delta$, we add the hypothesis constraint

$$F_t(\bar{\alpha}) \leq \llbracket \phi_t \rrbracket.$$

To decide $\Delta \models \phi \leq \psi$, we ask the solver whether

$$\llbracket \phi \rrbracket \leq \llbracket \psi \rrbracket$$

is entailed under these hypotheses. Before calling the Layer 1 order test, these hypothesis constraints are eliminated by bounded substitution: a hypothesis $x \leq b(x)$ is discharged by replacing x with $x \sqcap b(x)$ in the conclusion. After this step, the final Layer 1 query is unconditional. This is the constraint check used by the CROSS rule (Fig. 4) and therefore by kind checking and subsumption in the type system.

(Layer 4 kinds) \implies (Layer 3 function constraints)
 \implies (Layer 2 fixpoint assignments)
 \implies (Layer 1 polynomial DAG operations).

Detailed pseudocode for the solver operations used by this reduction is given in Appendix D.

D Simplified LDD Operations

These supplementary materials include the complete OCaml compiler used in the paper, including both an optimized LDD implementation and a pedagogical simplified LDD implementation used as a reference model. This section documents the simplified version, whose purpose is to expose the mathematical structure directly, without low-level representation tricks.

Placement in the inference pipeline. This appendix section is the concrete operational companion to Layer 1 and parts of Layer 2 in Appendix C. In particular, `sub_subsets` is the Layer 1 co-Heyting subtraction operation on DAGs, and Layer 2 fixpoint assignments reuse these normalized DAG operations.

Design goal. The simplified implementation keeps exactly the semantic invariants used by the optimized solver, but removes engineering optimizations that obscure the core algorithm. In particular, it does not use unboxed leaves, cached down-edges, pointer-equality short-cuts, or mutation-based fast paths. As a result, each operation is written as direct structural recursion over the LDD syntax tree.

Representation and invariant. An LDD node has the form

$$\text{Node}(v, \ell, h) \quad \text{representing} \quad \ell \sqcup (v \sqcap h).$$

Variables are totally ordered and each path follows increasing variable ids. The key normalization invariant is:

$$h = h \setminus \ell.$$

Intuitively, the high branch stores only the contribution that is not already present in the low branch. Normalization is enforced by a dedicated constructor `node v ~lo ~hi`, which computes `canonicalize(hi, lo)` before building the node.

Core operations. The algebraic operations `join`, `meet`, and `sub_subsets` follow one common pattern: they compare the top variables of the two inputs and recurse according to the variable order. When roots coincide, recursion proceeds on corresponding subtrees; when one root is smaller, recursion descends only on that side. This is the same merge discipline used in ordered BDD-style structures, and it preserves variable ordering by construction.

The only non-trivial operation is canonicalization (co-Heyting subtraction). Its recursive cases mirror the ordered merge above and maintain $h = h \setminus \ell$ at each rebuilt node. In the simplified code this procedure is intentionally explicit rather than factored through specialized fast paths.

```
(* Node(v, lo, hi) denotes lo  $\sqcup$  (v  $\sqcap$  hi). *)
```

```
let rec canonicalize hi lo =
  match hi, lo with
  | _, _ when hi == lo  $\rightarrow$  bot
  | _, _ when is_bot lo  $\rightarrow$  hi
  | Leaf a, Leaf b  $\rightarrow$  Leaf (a \ b)
```

```

2108 | Leaf a, Node _ → Leaf (a \ down0 lo)
2109 | Node _, Leaf b → map_leaves (fun a → a \ b) hi
2110 | Node (v,l1,h1), Node (w,l2,h2) →
2111   let order = compare v w in
2112   if order = 0 then
2113     let lo' = canonicalize l1 l2 in
2114     let hi' = canonicalize (canonicalize h1 h2) l2 in
2115     node_raw v lo' hi'
2116   else if order < 0 then
2117     node_raw v (canonicalize l1 lo) (canonicalize h1 lo)
2118   else
2119     canonicalize hi l2
2120
2121 let node v ~lo ~hi = node_raw v lo (canonicalize hi lo)
2122 let sub_subsets a b = canonicalize a b

```

Here `canonicalize` is the core co-Heyting subtraction routine on DAGs: it computes $hi \setminus lo$ while respecting variable order and preserving the normalization invariant. `node` is the public smart constructor; it is the only constructor that enforces normalization. `sub_subsets` is the exported subtraction operation, defined directly by canonicalization.

Why the order = 0 case is subtle. Suppose both roots have the same variable:

$$tree_1 = \text{Node}(v, l_1, h_1), \quad tree_2 = \text{Node}(v, l_2, h_2).$$

Then we must compute

$$(l_1 \sqcup (v \sqcap h_1)) \setminus (l_2 \sqcup (v \sqcap h_2)).$$

The key point is that l_2 is active in *both* worlds ($v = \perp$ and $v = \top$), while h_2 is active only when $v = \top$. Therefore:

$$\begin{aligned} lo' &= l_1 \setminus l_2, \\ hi1' &= h_1 \setminus h_2, \\ hi' &= hi1' \setminus l_2 = (h_1 \setminus h_2) \setminus l_2. \end{aligned}$$

So the result is `Node(v, lo', hi')`. The second subtraction by l_2 is exactly the non-obvious step: without it, the high branch would incorrectly retain mass already covered by the low branch of b .

```

2142 let rec join a b =
2143   match a, b with
2144   | Leaf x, Leaf y → Leaf (x \sqcup y)
2145   | Leaf x, _ → join_with_leaf x b
2146   | _, Leaf y → join_with_leaf y a
2147   | Node (v,l1,h1), Node (w,l2,h2) →
2148     let order = compare v w in
2149     if order = 0 then
2150       node_raw v
2151         (join l1 l2)
2152         (join (canonicalize h1 l2) (canonicalize h2 l1))
2153     else if order < 0 then
2154       node_raw v (join l1 b) (canonicalize h1 b)
2155     else
2156       node_raw w (join a l2) (canonicalize h2 a)

```

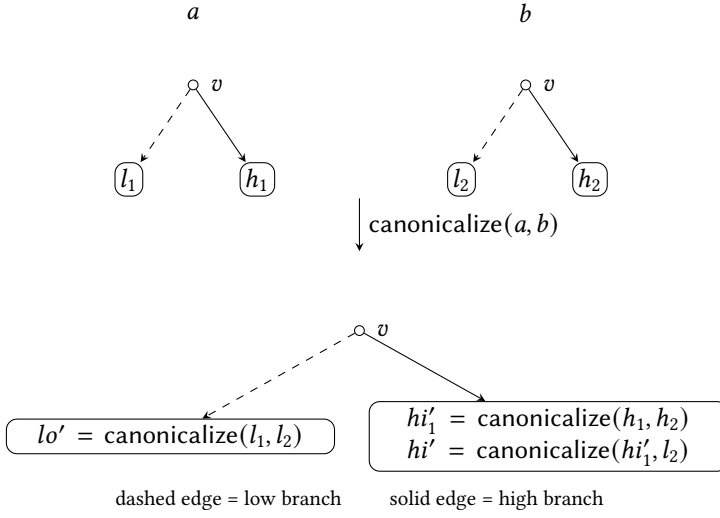


Fig. 7. Equal-root case in canonicalize: both inputs split on the same variable v , so subtraction is computed by separate low/high residuals, with the high residual requiring an additional subtraction by l_2 .

```

and join_with_leaf c n =
  match n with
  | Leaf d → Leaf (c ⊔ d)
  | Node (v, lo, hi) →
      node_raw v (join_with_leaf c lo) (canonicalize hi (Leaf c))

let rec meet a b =
  match a, b with
  | Leaf x, _ → meet_with_leaf x b
  | _, Leaf y → meet_with_leaf y a
  | Node (v, l1, h1), Node (w, l2, h2) →
      let order = compare v w in
      if order = 0 then
        let lo = meet l1 l2 in
        let hi = meet (join h1 l1) (join h2 l2) in
        node v ~lo ~hi
      else if order < 0 then
        node v ~lo:(meet l1 b) ~hi:(meet h1 b)
      else
        node w ~lo:(meet a l2) ~hi:(meet a h2)

and meet_with_leaf c n =
  match n with
  | Leaf d → Leaf (c ⊓ d)
  | Node (v, lo, hi) →
      node v ~lo:(meet_with_leaf c lo) ~hi:(meet_with_leaf c hi)

```

join merges two normalized DAGs according to top-variable order and re-normalizes only where overlap can arise. join_with_leaf is the specialized branch for adding a constant to a structured

diagram; it pushes the constant down the low spine and subtracts it from the high branch. meet is dual in structure and combines same-variable nodes by intersecting the induced upper sets. meet_with_leaf is the base case for meeting with a constant.

Naive baseline and why this code is faster. A fully naive presentation would define join and meet by first constructing an unnormalized candidate and then always calling the smart constructor node, which runs canonicalization on the high branch. For example, with equal roots $a = \text{Node}(v, l_1, h_1)$, $b = \text{Node}(v, l_2, h_2)$, a naive join could be written as

$$\text{node}(v, \text{join}(l_1, l_2), \text{join}(\text{join}(l_1, h_1), \text{join}(l_2, h_2))),$$

and then rely on node to subtract the low part from the high part. The implementation above is equivalent, but faster, because it performs only the subtraction work that is actually needed and in several branches uses node_raw (no extra canonicalization pass).

The key equalities rely on canonical inputs:

$$h_1 = h_1 \setminus l_1, \quad h_2 = h_2 \setminus l_2.$$

These invariants are used in the following non-obvious places.

- (1) **Equal-root join (order = 0).** The naive high branch is

$$hi_{\text{naive}} = (l_1 \sqcup h_1) \sqcup (l_2 \sqcup h_2).$$

After subtracting $lo' = \text{join}(l_1, l_2)$, we get

$$hi_{\text{naive}} \setminus lo' = (h_1 \setminus l_2) \sqcup (h_2 \setminus l_1),$$

exactly $\text{join}(\text{canonicalize}(h_1, l_2), \text{canonicalize}(h_2, l_1))$. This is why the code can directly build with node_raw: the residual is already low-disjoint.

- (2) **Unequal-root join (order < 0 /> 0).** In the order < 0 branch, the naive high would be $\text{join}(h_1, b)$, then normalized against $\text{join}(l_1, b)$. Also, because inputs are variable-ordered and $v < w$, b cannot contain a test on v ; this is why recursion may keep b unchanged while descending under a 's v -node. Using $h_1 \setminus l_1 = h_1$, this reduces to $h_1 \setminus b$, i.e. $\text{canonicalize}(h_1, b)$, so the code avoids constructing and re-normalizing $\text{join}(h_1, b)$. The order > 0 case is symmetric.
- (3) **join_with_leaf.** Naively, one would normalize $\text{join}(c, hi)$ against $\text{join}(c, lo)$. By canonicity of (lo, hi) , this simplifies to $hi \setminus c$, so the code uses $\text{canonicalize}(hi, \text{Leaf}(c))$ directly.
- (4) **meet.** The implementation already follows the denotational split, but still depends on canonical representation: for equal roots, the high branch must use activated values $(l_i \sqcup h_i)$, not just h_i , because $\text{Node}(v, l, h)$ denotes $l \sqcup (v \sqcap h)$. For unequal roots, ordering of canonical DAGs ensures that if $v < w$, then b contains no test on v , so descending only on a 's side is sound.

If inputs were not canonical, these identities would fail and the node_raw-based fast paths in join would no longer be correct.

Why the Node-Node case of join is subtle. Let

$$a = \text{Node}(v, l_1, h_1), \quad b = \text{Node}(w, l_2, h_2).$$

The join recursion splits into three structurally different cases.

- (1) $v = w$ (**order = 0**). Both diagrams branch on the same variable, so low branches combine directly:

$$lo' = \text{join}(l_1, l_2).$$

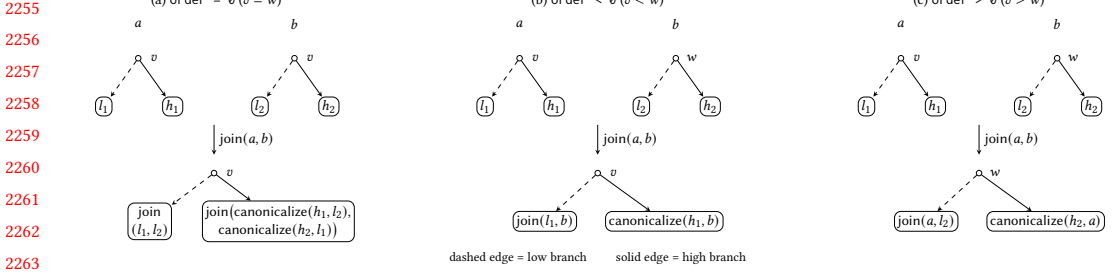


Fig. 8. Case split for join on two non-leaf roots. The three cases are determined by the top-variable order and differ in where normalization (canonicalize) is required to maintain $hi = hi \setminus lo$.

The high branch is subtler: each side's high contribution must be normalized against the other side's low contribution before being joined:

$$hi' = \text{join}(\text{canonicalize}(h_1, l_2), \text{canonicalize}(h_2, l_1)).$$

This is what preserves the invariant $hi = hi \setminus lo$.

- (2) $v < w$ (**order** < 0). Variable v must remain above w , so we keep v at the root and push all of b under both branches:

$$lo' = \text{join}(l_1, b), \quad hi' = \text{canonicalize}(h_1, b).$$

- (3) $v > w$ (**order** > 0). This is the symmetric case:

$$lo' = \text{join}(a, l_2), \quad hi' = \text{canonicalize}(h_2, a),$$

and the output root is w .

Why the Node-Node case of meet is subtle. Again let

$$a = \text{Node}(v, l_1, h_1), \quad b = \text{Node}(w, l_2, h_2).$$

The recursion has the same three-way split by variable order, but the $\text{order} = 0$ branch differs fundamentally from join. When $v = w$, expanding the denotation $\text{Node}(v, l, h) = l \sqcup (v \sqcap h)$ gives

$$\begin{aligned} a \sqcap b &= (l_1 \sqcup (v \sqcap h_1)) \sqcap (l_2 \sqcup (v \sqcap h_2)) \\ &= (l_1 \sqcap l_2) \sqcup (v \sqcap ((l_1 \sqcup h_1) \sqcap (l_2 \sqcup h_2))). \end{aligned}$$

So the high branch must combine the *activated* values $(l_i \sqcup h_i)$, not just h_i .

- (1) $v = w$ (**order** $= 0$). The low branch is pointwise:

$$lo' = \text{meet}(l_1, l_2).$$

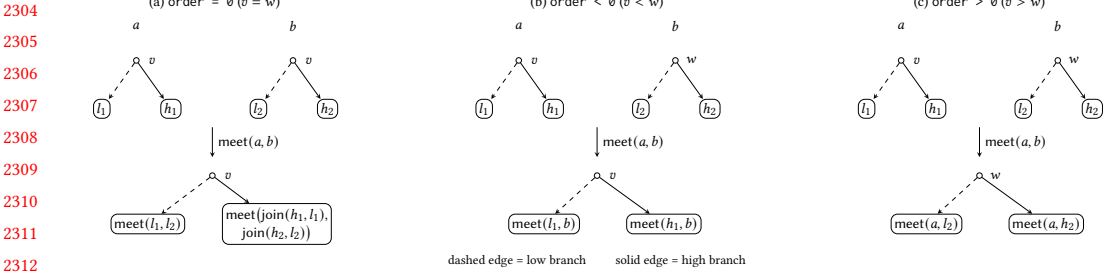
For the high branch, each side can contribute either via its low branch or via its high branch once v is true, so we first form the activated values $\text{join}(h_1, l_1)$ and $\text{join}(h_2, l_2)$, then meet them:

$$hi' = \text{meet}(\text{join}(h_1, l_1), \text{join}(h_2, l_2)).$$

The result is rebuilt as $\text{node}(v, lo', hi')$, which re-normalizes.

- (2) $v < w$ (**order** < 0). Keep v at the root and meet both branches of a with all of b :

$$lo' = \text{meet}(l_1, b), \quad hi' = \text{meet}(h_1, b).$$



2314 Fig. 9. Case split for meet on two non-leaf roots. In the $order = \emptyset$ case, the high branch meets $join(h_1, l_1)$
 2315 with $join(h_2, l_2)$ before re-normalization via node.

2316 (3) $v > w$ (**order** $> \emptyset$). Symmetrically:

$$2317 lo' = meet(a, l_2), \quad hi' = meet(a, h_2),$$

2318 and the output root is w .

```

2322 let rec assign_bot ~var n =
2323   match n with
2324   | Leaf _ → n
2325   | Node (v, lo, hi) →
2326     let order = compare var v in
2327     if order < 0 then n
2328     else if order = 0 then lo
2329     else node v ~lo:(assign_bot ~var lo) ~hi:(assign_bot ~var hi)
2330
2331 let rec assign_top ~var n =
2332   match n with
2333   | Leaf _ → n
2334   | Node (v, lo, hi) →
2335     let order = compare var v in
2336     if order < 0 then n
2337     else if order = 0 then join lo hi
2338     else node v ~lo:(assign_top ~var lo) ~hi:(assign_top ~var hi)
2339
2340 let rec inline_solved_vars n =
2341   match n with
2342   | Leaf _ → n
2343   | Node (v, lo, hi) →
2344     match v.state with
2345     | Rigid _ → n
2346     | Solved d →
2347       let lo' = inline_solved_vars lo in
2348       let hi' = inline_solved_vars hi in
2349       let d' = inline_solved_vars d in
2350       v.state <- Solved d';
2351       join lo' (meet hi' d')
2352   | Unsolved →
2353     node v ~lo:(inline_solved_vars lo) ~hi:(inline_solved_vars hi)
  
```

2353 `assign_bot` computes substitution $var := \perp$ by removing the high branch at the matching variable.
 2354 `assign_top` computes substitution $var := \top$ by joining low and high branches at the matching variable.
 2355 `inline_solved_vars` eliminates solved flexible variables by replacing each solved occurrence
 2356 with its definition, while leaving rigid variables untouched.

2357

2358

2359

2360

2361

2362

2363

2364

2365

2366

2367

2368

2369

2370

2371

2372

2373

2374

2375

2376

2377

2378

2379

2380

2381

2382

2383

2384

2385

2386

2387

2388

2389

2390

2391

2392

2393

2394

2395

2396

2397

2398

2399

2400

2401

Fixpoint variables. The simplified implementation keeps the same solver-level interface as the optimized code: `solve_lfp`, `solve_gfp`, pending GFP queues, and solved-variable inlining. The simplified module is therefore a reference model for the optimized module, covering full solver behavior rather than only basic lattice operations.

A key interface point is that solved substitutions are *lazy* with respect to existing trees. Calling `solve_lfp` or `solve_gfp` updates variable state (`var.state <- Solved ...`) but does not traverse and rewrite every LDD currently in scope. Hence multiple trees may remain “in flight” after solving, still containing references to variables that are now marked solved. The operation `inline_solved_vars` is the forcing step that pushes those substitutions into a given tree. Operationally, after solving, any tree that is reused should be forced with `inline_solved_vars`, unless the called operation already does so internally.

```

let solve_lfp var rhs =
  match var.state with
  | Unsolved →
    let rhs' = inline_solved_vars rhs in
    var.state <- Solved (assign_bot ~var rhs')
  | _ → error "invalid variable state"

let solve_gfp var rhs =
  match var.state with
  | Unsolved →
    let rhs' = inline_solved_vars rhs in
    var.state <- Solved (assign_top ~var rhs')
  | _ → error "invalid variable state"

let enqueue_gfp var rhs = pending := (var, rhs) :: !pending
let solve_pending () =
  let work = !pending in
  pending := [];
  iter (fun (v,rhs) → solve_gfp v rhs) work

```

`solve_lfp` and `solve_gfp` implement one-variable least and greatest fixpoint solving, respectively, after forcing solved substitutions in the right-hand side. `enqueue_gfp` records deferred GFP constraints, and `solve_pending` snapshots and clears the queue, then discharges the snapshot in sequence. Neither `solve_lfp` nor `solve_gfp` rewrites arbitrary external trees; they only install the solved definition on the variable.

```

let rec round_up n =
  match inline_solved_vars n with
  | Leaf c → c
  | Node (_,lo,hi) → round_up lo ⊔ round_up hi

let leq_with_reason a b =

```

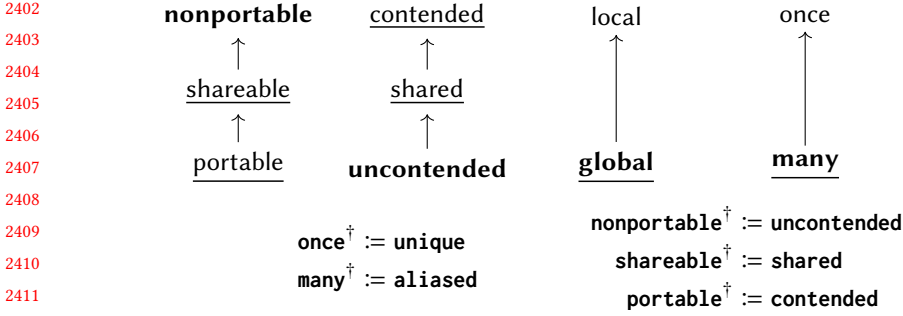


Fig. 10. The portability, contention, locality, affinity, and uniqueness modal axes, and their duality operation. Portability, locality, and affinity are comonadic, and contention and uniqueness are monadic.

```

2417 let diff = sub_subsets (inline_solved_vars a) (inline_solved_vars b) in
2418 let witness = round_up diff in
2419 non_bot_axes witness

```

round_up evaluates a ground diagram by interpreting each variable as \top , i.e., by joining all leaves reachable in the normalized DAG. leq_with_reason checks $a \leq b$ via $a \setminus b$, then extracts non-bottom lattice axes as a concrete witness when the inequality fails. Both routines explicitly force solved substitutions, so they are safe entry points when trees are still in flight.

Relationship to the optimized implementation. The code above implements the same API as the optimized implementation. The optimized module has better node representation, cached summaries, and aggressive short-circuiting. To validate this correspondence, we tested (1) concrete examples (canonicalization, assignments, LFP/GFP solving), (2) randomized algebraic/property tests, and (3) cross-checks between optimized and simplified modules.

E SMOKi Type System Definitions

The complete SMOKi with all supported modal axes (locality, affinity, uniqueness, portability, and contention) is described in Figures 10 and 12 to 16. These rules additionally use the locality, affinity, and uniqueness modes. Specifically, the typing rules make heavy use of the context joining operation $+$ to encode substructural features. These figures are heavily based on Georges et al. [12].

F Important Changes to the Logical Relation

Proving the correct mode-crossing rules for mutable references and recursive types (§3.4) required interesting changes to the model of DRFCaml [12]. We briefly describe these changes in this section.

DRFCaml's type system distinguishes between two reference types $\text{ref}_p A$ using a portability-mode annotation. This annotation represents the portability mode of the reference's contents: Only a $\text{ref}_{\text{portable}} A$ reference can be allocated at **portable** mode, and a $\text{ref}_{\text{nonportable}} A$ would always be at **nonportable** mode. This distinction was necessary precisely because DRFCaml's model of references did not support mode crossing. In OCaml, by contrast, an `int ref` allocated at **nonportable** mode can cross portability and become **portable**. We improved DRFCaml's model of references to support this mode-crossing behavior. We achieve this by additionally requiring a reference at **portable** (**shareable**) mode to carry a witness that its inner type crosses portability (crosses **shareable**).

$$\begin{array}{c}
2500 \quad \boxed{\Delta \vdash A \text{ ok}} \\
2501 \\
2502 \quad \frac{\Delta \vdash A \text{ ok} \quad \Delta \vdash B \text{ ok}}{\Delta \vdash A \times B \text{ ok}} \quad \frac{\Delta \vdash A \text{ ok} \quad \Delta \vdash B \text{ ok}}{\Delta \vdash A + B \text{ ok}} \quad \frac{\Delta \vdash A \text{ ok}}{\Delta \vdash \Box^{m'} A \text{ ok}} \quad \frac{(\alpha : m) \in \Delta}{\Delta \vdash \alpha \text{ ok}} \\
2503 \\
2504 \quad \frac{\forall i. \Delta \vdash A_i \text{ ok} \quad (t : \bar{a}. \varphi) \in \Delta}{\Delta \vdash t \bar{A} \text{ ok}} \quad \frac{\Delta, t : \bar{a}. \top, \bar{\alpha} : \bar{\top} \vdash A \text{ ok} \quad \forall i. \Delta \vdash B_i \text{ ok}}{\Delta \vdash (\mu t \bar{a}. A) \bar{B} \text{ ok}} \\
2505 \\
2506 \\
2507
\end{array}$$

Fig. 13. Auxiliary SMOki judgements.

2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538
2539
2540
2541
2542
2543
2544
2545
2546
2547
2548

Something similar was already observed in RustBelt, where delayed sharing is presented as a solution [14]. It is unclear how delayed sharing would be adapted to the DRFCaml logical relation, or whether it generalizes to all of the laws we require. Instead, we develop a technique based on later credits and additive time receipts for later credit generation [20] and on RefinedRust's *logical step modality* [11, 26] to work around these limitations.

$$\begin{array}{c}
2549 \\
2550 \\
2551 \\
2552 \\
2553 \\
2554 \\
2555 \\
2556 \\
2557 \\
2558 \\
2559 \\
2560 \\
2561 \\
2562 \\
2563 \\
2564 \\
2565 \\
2566 \\
2567 \\
2568 \\
2569 \\
2570 \\
2571 \\
2572 \\
2573 \\
2574 \\
2575 \\
2576 \\
2577 \\
2578 \\
2579 \\
2580 \\
2581 \\
2582 \\
2583 \\
2584 \\
2585 \\
2586 \\
2587 \\
2588 \\
2589 \\
2590 \\
2591 \\
2592 \\
2593 \\
2594 \\
2595 \\
2596 \\
2597
\end{array}$$

$$\begin{array}{c}
\frac{}{\Delta; \Gamma \vdash () : \mathbb{1} @ m} \text{UNIT} \quad \frac{}{\Delta; \Gamma \vdash b : \mathbb{B} @ m} \text{BOOL} \quad \frac{}{\Delta; \Gamma \vdash z : \mathbb{Z} @ m} \text{INT} \quad \frac{}{\Delta; \Gamma, x : A @ m, \Gamma' \vdash x : A @ m} \text{VAR} \\
\frac{\Delta; \mathbf{\clubsuit}_w(\Gamma), x : A @ m_1 \vdash e : B @ m_2}{\Delta; \Gamma \vdash \lambda^{w.l} x, e : (A @ m_1 \rightarrow B @ m_2) @ (w, n)} \text{LAM} \\
\frac{w.o = \text{many} \quad \Delta; \mathbf{\clubsuit}_w(\Gamma), x : A @ m_1, f : (A @ m_1 \rightarrow B @ m_2) @ (w, n) \vdash e : B @ m_2}{\Delta; \Gamma \vdash \lambda^{w.l} f x, e : (A @ m_1 \rightarrow B @ m_2) @ (w, n)} \text{REC} \\
\frac{\Delta; \Gamma_1 \vdash e_1 : (A @ m_1 \rightarrow B @ m_2) @ m_3 \quad \Delta; \Gamma_2 \vdash e_2 : A @ m_1}{\Delta; \Gamma_1 + \Gamma_2 \vdash e_1 e_2 : B @ m_2} \text{APP} \\
\frac{\Delta; \Gamma_1 \vdash e_1 : \mathbb{1} @ m_1 \quad \Delta; \Gamma_2 \vdash e_2 : A @ m_2}{\Delta; \Gamma_1 + \Gamma_2 \vdash (e_1; e_2) : A @ m_2} \text{SEQ} \quad \frac{\Delta; \Gamma_1 \vdash e_1 : A @ m_1 \quad \Delta; \Gamma_2, x : A @ m_1 + e_2 : B @ m_2}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{let } x := e_1 \text{ in } e_2 : B @ m_2} \text{LET} \\
\frac{\Delta; \Gamma_1 \vdash e_1 : A @ (p', l, \text{many}, c', \text{unique}) \quad \Delta; \Gamma_2, x : A @ (p', \text{local}, \text{many}, c', \text{aliased}) \vdash e_2 : B @ (p, \text{global}, o, c, u) \quad \Delta; \Gamma_3, x : A @ (p', l, \text{many}, c', \text{unique}), y : B @ (p, \text{global}, o, c, u) \vdash e_3 : C @ m}{\Delta; \Gamma_1 + \Gamma_2 + \Gamma_3 \vdash \text{borrow } x := e_1 \text{ for } y := e_2 \text{ in } e_3 : C @ m} \text{BORROW} \\
\frac{\Delta; \mathbf{\clubsuit}(\text{global, once, portable})(\Gamma) \vdash e : A @ m_1}{\Delta; \Gamma \vdash \text{fork}(e) : \mathbb{1} @ m_2} \text{FORK} \quad \frac{\Delta; \Gamma_1 \vdash e_1 : A @ m \quad \Delta; \Gamma_2 \vdash e_2 : B @ m}{\Delta; \Gamma_1 + \Gamma_2 \vdash (e_1, e_2) : A \times B @ m} \text{PAIR} \\
\frac{\Delta; \Gamma_1 \vdash e_1 : A \times B @ m_1 \quad \Delta; \Gamma_2, y : B @ m_1, x : A @ m_1 \vdash e_2 : C @ m_2}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{unpair } e_1 \text{ as } (x, y) \text{ in } e_2 : C @ m_2} \text{UNPAIR} \\
\frac{\Delta; \Gamma \vdash e : A @ m}{\Delta; \Gamma \vdash \text{inl}(e) : A + B @ m} \text{INL} \quad \frac{\Delta; \Gamma \vdash e : B @ m}{\Delta; \Gamma \vdash \text{inr}(e) : A + B @ m} \text{INR} \\
\frac{\Delta; \Gamma_1 \vdash e : A + B @ m_1 \quad \Delta; \Gamma_2, x : A @ m_1 \vdash e_1 : C @ m_2 \quad \Delta; \Gamma_2, x : B @ m_1 \vdash e_2 : C @ m_2}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{case } e \{ \text{inl } x \rightarrow e_1; \text{inr } y \rightarrow e_2 \} : C @ m_2} \text{CASE} \quad \frac{\Delta; \Gamma_1 \vdash e : \mathbb{B} @ m_1 \quad \Delta; \Gamma_2 \vdash e_1 : A @ m_2 \quad \Delta; \Gamma_2 \vdash e_2 : A @ m_2}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{if } e \text{ then } e_1 \text{ else } e_2 : A @ m_2} \text{IF} \\
\frac{\Delta; \Gamma \vdash e : A @ (p, \text{global}, o, c, u)}{\Delta; \Gamma \vdash \text{region}(e) : A @ (p, \text{global}, o, c, u)} \text{REGION} \quad \frac{\Delta; \Gamma_1 \vdash e_1 : A @ m \quad \Delta; \Gamma_2 \vdash e_2 : B @ m \quad \text{binopTyped}(\oplus, A, B, C)}{\Delta; \Gamma_1 + \Gamma_2 \vdash e_1 \oplus e_2 : C @ m} \text{BINOP} \\
\frac{\Delta; \Gamma \vdash e : A @ m \quad \text{unopTyped}(\oplus, A, B)}{\Delta; \Gamma \vdash \oplus(e) : B @ m} \text{UNOP} \quad \frac{\text{BOX} \quad \Delta; \Gamma \vdash e : A @ m_1 \sqcap m_2}{\Delta; \Gamma \vdash \text{box } e : \square^{m_2} A @ m_1} \quad \frac{\text{UNBOX} \quad \Gamma; \Delta \vdash e : \square^{m_2} A @ m_1}{\Gamma; \Delta \vdash \text{unbox } e : A @ m_1 \sqcap m_2} \\
\frac{\Delta; \Gamma \vdash e : A @ (w, n) \quad w \leq w' \quad n' \leq n}{\Delta; \Gamma \vdash e : A @ (w', n')} \text{SUB} \quad \frac{\Gamma_1 \geq \Gamma_2 \quad \Delta; \Gamma_1 \vdash e : A @ m}{\Delta; \Gamma_2 \vdash e : A @ m} \text{SUB-CTX} \quad \frac{\Delta; \Gamma \vdash e : A @ (w, n) \quad \Delta \vDash \text{kind}(A) \leq (w', n')}{\Delta; \Gamma \vdash e : A @ (w' \sqcap w, n' / n)} \text{CROSS}
\end{array}$$

Fig. 14. SMOki typing rules.

$$\begin{array}{c}
2598 \\
2599 \\
2600 \\
2601 \\
2602 \\
2603
\end{array}
\frac{\forall i. \Delta \vdash B_i \text{ ok} \quad \Delta; \Gamma \vdash e : A[\overline{\alpha \mapsto B}, t \mapsto \bar{\alpha}. (\mu t \bar{\alpha}. A) \bar{\alpha}] @ m}{\Delta; \Gamma \vdash \text{roll } e : (\mu t \bar{\alpha}. A) \bar{B} @ m} \text{ROLL} \quad \frac{\Delta; \Gamma \vdash e : (\mu t \bar{\alpha}. A) \bar{B} @ m}{\Delta; \Gamma \vdash \text{unroll } e : A[\overline{\alpha \mapsto B}, t \mapsto \bar{\alpha}. (\mu t \bar{\alpha}. A) \bar{\alpha}] @ m} \text{UNROLL}$$

Fig. 15. SMOki typing rules for recursive types.

$$\begin{array}{c}
2604 \\
2605 \\
2606 \\
2607 \\
2608 \\
2609 \\
2610 \\
2611 \\
2612 \\
2613 \\
2614 \\
2615 \\
2616 \\
2617 \\
2618 \\
2619 \\
2620 \\
2621 \\
2622 \\
2623 \\
2624 \\
2625 \\
2626 \\
2627 \\
2628 \\
2629 \\
2630 \\
2631 \\
2632 \\
2633 \\
2634 \\
2635 \\
2636 \\
2637 \\
2638 \\
2639 \\
2640 \\
2641 \\
2642 \\
2643 \\
2644 \\
2645 \\
2646
\end{array}
\frac{\Delta; \Gamma \vdash e : A @ (\text{nonportable}, l, \text{many}, \text{uncontended}, \text{aliased})}{\Delta; \Gamma \vdash \text{alloc}^l(e) : \text{ref}(A) @ (\text{nonportable}, l, \text{many}, \text{uncontended}, \text{unique})} \text{NAALLOC}$$

$$\frac{\Delta; \Gamma \vdash e : \text{ref}(A) @ (p, l, o, c, u) \quad c \neq \text{contended}}{\Delta; \Gamma \vdash !^{\text{NA}} e : A @ (p, l, \text{many}, c, \text{aliased})} \text{NALOAD}$$

$$\frac{\Delta; \Gamma_1 \vdash e_1 : \text{ref}(A) @ (p, l, o, \text{uncontended}, u) \quad \Delta; \Gamma_2 \vdash e_2 : A @ (p, \text{global}, \text{many}, \text{uncontended}, \text{aliased})}{\Delta; \Gamma_1 + \Gamma_2 \vdash e_1 \leftarrow^{\text{NA}} e_2 : \mathbb{1} @ m} \text{NASTORE}$$

$$\frac{\Delta; \Gamma \vdash e : A @ m \quad m.l = l}{\Delta; \Gamma \vdash \text{alloc}^l(e) : \text{box}(A) @ m} \text{IMALLOC} \quad \frac{\Delta; \Gamma \vdash e : \text{box}(A) @ m}{\Delta; \Gamma \vdash !^{\text{NA}} e : A @ m} \text{IMLOAD}$$

$$\frac{\Delta; \Gamma \vdash e : A @ (\text{portable}, l, \text{many}, \text{contended}, \text{unique})}{\Delta; \Gamma \vdash \text{alloc}^l(e) : \text{atomic}(A) @ (\text{portable}, l, \text{many}, \text{contended}, \text{unique})} \text{ATALLOC}$$

$$\frac{\Delta; \Gamma \vdash e : \text{atomic}(A) @ m}{\Delta; \Gamma \vdash !^{\text{AT}} e : A @ (p, l, o, \text{contended}, \text{aliased})} \text{ATLOAD}$$

$$\frac{\Delta; \Gamma_1 \vdash e_1 : \text{atomic}(A) @ (\text{portable}, l, o, \text{contended}, u) \quad \Delta; \Gamma_2 \vdash e_2 : A @ (\text{portable}, \text{global}, \text{many}, \text{contended}, \text{aliased})}{\Delta; \Gamma_1 + \Gamma_2 \vdash e_1 \leftarrow^{\text{AT}} e_2 : \mathbb{1} @ m} \text{ATSTORE}$$

$$\frac{\Delta; \Gamma_1 \vdash e_1 : \text{atomic}(A) @ (\text{portable}, l, o, \text{contended}, u) \quad \Delta; \Gamma_2 \vdash e_2 : A @ m \quad \Delta; \Gamma_3 \vdash e_3 : A @ (\text{portable}, \text{global}, \text{many}, \text{contended}, \text{aliased}) \quad \text{typeCmpSafe}(A)}{\Delta; \Gamma_1 + \Gamma_2 + \Gamma_3 \vdash \text{cmpXchg}(e_1, e_2, e_3) : A \times B @ (\text{portable}, l, \text{many}, \text{contended}, \text{aliased})} \text{CMPXCHG}$$

$$\frac{\Delta; \Gamma_1 \vdash e_1 : \text{atomic}(A) @ (\text{portable}, l, o, \text{contended}, u) \quad \Delta; \Gamma_2 \vdash e_2 : A @ (\text{portable}, \text{global}, \text{many}, \text{contended}, \text{aliased})}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{xchg}(e_1, e_2) : A @ (\text{portable}, l, \text{many}, \text{contended}, \text{aliased})} \text{XCHG}$$

$$\frac{\Delta; \Gamma_1 \vdash e_1 : \text{atomic}(\mathbb{Z}) @ (\text{portable}, l, o, \text{contended}, u) \quad \Delta; \Gamma_2 \vdash e_2 : \mathbb{Z} @ m_1}{\Delta; \Gamma_1 + \Gamma_2 \vdash \text{faa}(e_1, e_2) : \mathbb{Z} @ m_2} \text{FAA}$$

Fig. 16. SMOki typing rules for state.