



TypeDis: A Type System for Disentanglement

ALEXANDRE MOINE, New York University, USA
STEPHANIE BALZER, Carnegie Mellon University, USA
ALEX XU, Carnegie Mellon University, USA
SAM WESTRICK, New York University, USA

Disentanglement is a runtime property of parallel programs guaranteeing that parallel tasks remain oblivious to each other's allocations. As demonstrated in the MaPLE compiler and run-time system, disentanglement can be exploited for fast automatic memory management, especially task-local garbage collection with no synchronization between parallel tasks. However, as a low-level property, disentanglement can be difficult to reason about for programmers. The only means of statically verifying disentanglement so far has been DisLog, an Iris-fueled variant of separation logic, mechanized in the Rocq proof assistant. DisLog is a fully-featured program logic, allowing for proof of functional correctness as well as verification of disentanglement. Yet its employment requires significant expertise and per-program proof effort.

This paper explores the route of automatic verification via a type system, ensuring that any well-typed program is disentangled and lifting the burden of carrying out manual proofs from the programmer. It contributes TypeDis, a type system inspired by region types, where each type is annotated with a timestamp, identifying the task that allocated it. TypeDis supports iso-recursive types as well as polymorphism over both types and timestamps. Crucially, timestamps are allowed to change during type-checking, at join points as well as via a form of subtyping, dubbed *subtiming*. The paper illustrates TypeDis and its features on a range of examples. The soundness of TypeDis and the examples are mechanized in the Rocq proof assistant, using an improved version of DisLog, dubbed DisLog2.

CCS Concepts: • Software and its engineering → Parallel programming languages; • Theory of computation → Type theory; Separation logic.

Additional Key Words and Phrases: disentanglement, parallelism, type system, separation logic

ACM Reference Format:

Alexandre Moine, Stephanie Balzer, Alex Xu, and Sam Westrick. 2026. TypeDis: A Type System for Disentanglement. *Proc. ACM Program. Lang.* 10, POPL, Article 13 (January 2026), 30 pages. <https://doi.org/10.1145/3776655>

1 Introduction

A recent line of work has identified a key memory property of parallel programs called *disentanglement* [Acar et al. 2015; Arora et al. 2024, 2021, 2023; Guatto et al. 2018; Moine et al. 2024; Raghunathan et al. 2016; Westrick et al. 2022, 2020]. Roughly speaking, disentanglement is the property that concurrent tasks remain oblivious to each other's memory allocations. As demonstrated by the MaPLE compiler [Acar et al. 2020], this property makes it possible to perform task-local memory management (allocations and garbage collection) independently, in parallel, without any synchronization between concurrent tasks. MaPLE in particular features a provably efficient memory management system for a dialect of Parallel ML—a parallel functional programming

Authors' Contact Information: Alexandre Moine, alexandre.moine@nyu.edu, New York University, New York, USA; Stephanie Balzer, balzers@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, USA; Alex Xu, alexxu@andrew.cmu.edu, Carnegie Mellon University, Pittsburgh, USA; Sam Westrick, shw8119@nyu.edu, New York University, New York, USA.



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

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

ACM 2475-1421/2026/1-ART13

<https://doi.org/10.1145/3776655>

```

let r = ref ""
let rec write_max x =
  let current = !r in
  if x <= current || compare_and_swap r current x
  then () else write_max x
let entangled =
  par(fun () -> write_max (Int.to_string 1234),
       fun () -> write_max (Int.to_string 5678))

```

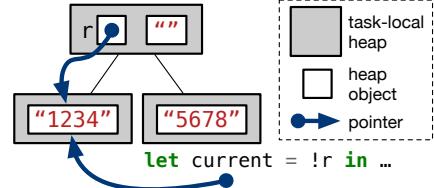


Fig. 1. Entanglement example

language—and offers competitive performance in practice relative to low-level parallel code written in languages such as C/C++ [Arora et al. 2023].

This line of work aims to gain control over the synchronization costs of parallel garbage collection by taking advantage of structured fork-join parallelism. The idea is to synchronize the garbage collector only at application-level forks and joins, thereby making these synchronization costs predictable at the source level, and avoiding the need for any global synchronization of the garbage collector. At each fork and join, the runtime system performs $O(1)$ work to maintain a dynamic tree of heaps which mirrors the parent/child relationships between tasks. Each task thus has its own task-local heap, in which it allocates memory objects and may perform garbage collection independently, in parallel. The independence of these task-local garbage collections hinges upon *disentanglement*, which can be defined as a “no cross-pointers” invariant. Specifically, disentanglement allows for *up-pointers* from descendant heaps to ancestors, as well as *down-pointers* from ancestors to descendants, but disallows *cross-pointers* between concurrent tasks (siblings, cousins, etc.). The existence of cross-pointers is called *entanglement*. When two tasks become entangled with a cross-pointer, neither task can perform garbage collection without synchronizing with the other. These additional synchronizations lead to significant performance degradations [Arora et al. 2023], and in this sense, entanglement is a performance hazard.

Entanglement arises from a particular communication pattern, where one task allocates a local heap object and then another task (executing concurrently, relative to the first task) acquires a pointer to the object. An example is shown in Figure 1. The example uses the fork-join primitive `par(f1, f2)` to execute two functions in parallel; these two function calls correspond to two child tasks. In the example, the two child tasks perform `write_max` concurrently, both attempting to update `r` to point to a locally allocated string. After one task finishes, the other task reads the updated `r` and acquires a cross-pointer, which constitutes entanglement.

As shown below, this example could be rewritten to be disentangled by moving the allocations of the strings “up” into the parent task (making all pointers involved up-pointers).

```

... (* same definitions of r and write_max, as in Figure 1 *) ...
let disentangled =
  let a = Int.to_string 1234 in let b = Int.to_string 5678 in
  par(fun () -> write_max a, fun () -> write_max b)

```

Preventing entanglement. One way to rule out entanglement is to disallow side effects entirely. Indeed, the original study of disentanglement emerged out of an interest in improving the performance of parallel functional programming techniques, which naturally have a high rate of allocation and whose scalability and efficiency is largely determined by the performance of automatic memory management. In this setting, disentanglement is guaranteed by construction due to a lack of side effects. But the full power of disentanglement lies in its expressivity beyond purely functional programming—in particular, disentanglement allows for judicious utilization of side effects such as in-place updates and irregular and/or data-dependent access patterns in shared memory. These side effects are crucial for efficiency in state-of-the-art implementations of parallel algorithms, such as

those in the PBBS Benchmark Suite [Abdi et al. 2024; Anderson et al. 2022; Shun et al. 2012], which have been found to be naturally disentangled [Westrick et al. 2022].

In these more general settings, where there are numerous opportunities to efficiently utilize side effects, it is easy for a programmer to accidentally entangle concurrent tasks. Ideally, it would be evident at the source level where entanglement may or may not occur. However, in practice, this is not the case. To reason about entanglement, the programmer effectively has to know the memory allocation and access patterns of the entire program. This makes it especially difficult to reason about higher-order functions, because the memory effects of a function taken as argument are unknown. Other high-level programming features also complicate the matter, such as parametric polymorphism which allows for code to be specialized for both “boxed” (heap-allocated) and “unboxed” types, potentially resulting in entanglement in one case but not the other. These details can be formally considered using the program logic DisLog [Moine et al. 2024], but verifying disentanglement using DisLog requires significant expertise and effort, even in small examples.

An interesting question therefore is whether it is possible to guarantee disentanglement statically through a type system. This would have the advantage of being mostly automatic, requiring (ideally) only a modest amount of type annotation. Most importantly, a type system would raise the level of abstraction at which the programmer can reason about disentanglement, clarifying how the property interacts with high-level abstractions such as parametric polymorphism, higher-order functions, algebraic datatypes, and other desirable features.

A type system for disentanglement. In this paper, we present **TypeDis**, the first static type system for disentanglement. We intend for TypeDis to be the type system for a high-level ML-like language with structured fork-join parallelism, in-place atomic operations on shared memory, and disentangled parallel garbage collection. The language features a single parallel construct, written $\text{par}(f_1, f_2)$, which calls $f_1()$ and $f_2()$ in parallel, waits for both to complete, and returns their results as a pair. Here, we think of the execution of the two function calls as two child tasks, which themselves might execute $\text{par}(\dots)$ recursively, creating a dynamic tree (parent-child) relationship between tasks.

TypeDis identifies tasks with timestamp variables δ , and annotates every value computed during execution with the timestamp of the task that allocated that value. This is tracked explicitly in the type of the value. For example, $s : \text{string}@\delta$ indicates that the value s is a string that was allocated by a task δ . The type system implicitly maintains a partial order over timestamps, written $\delta' \preccurlyeq \delta$, intuitively corresponding to the tree relationship between tasks. Crucially, TypeDis guarantees an invariant that we call the **up-pointer invariant**: for every task running at timestamp δ , every value accessed by this task must have a timestamp $\delta' \preccurlyeq \delta$, i.e., the value must have been allocated “before” the current timestamp. In other words, the key insight in this paper is to *restrict all memory references to point backwards in time*, which is checked statically. This restriction is a deep invariant over values: every data structure will only contain values allocated at the same timestamp or a preceding timestamp. As a result, all loads in the language are guaranteed to be safe for disentanglement.

The up-pointer invariant statically rules out one feature of disentanglement: down-pointers. This restriction is mild, however, because down-pointers are fairly rare. Quantitatively, there has been at least one relevant study: in their work on entanglement detection, [Westrick et al. 2022] observe in multiple benchmarks that down-pointers do not arise at all, and more broadly they measure that the number of objects containing down-pointers is small. The creation of a down-pointer requires a combination of dynamic allocation and pointer indirection, each of which is typically avoided in parallel performance-sensitive code to reduce memory pressure and improve cache efficiency. In this paper, we have found the up-pointer invariant to be sufficiently expressive to encode a number of interesting examples (§6), fully typed within TypeDis, and therefore guaranteed disentangled.

The up-pointer invariant is especially well-suited for immutable data (which naturally adheres to the invariant), as well as parallel batch processing of pre-allocated data. The up-pointer invariant also allows for structure sharing, even in the presence of mutable state.

Maintaining the up-pointer invariant. To maintain the up-pointer invariant in the presence of mutable state, TypeDis places a restriction on writes (in-place updates), requiring that the timestamp of the written value precedes the timestamp of the reference pointing to it. This restriction is implemented in the type system with a form of subtyping, dubbed **subtiming**, which affects only the timestamps of values within their types. The idea is to allow for any value to be (conservatively) restamped with a *newer* timestamp. Subtiming makes it possible to express the restriction on writes as a simple unification over the type of the contents of a mutable reference or array.

Restamping with an *older* timestamp would be unsound in TypeDis, as it would allow for a child's (heap-allocated) value to be written into a parent's container, potentially making that value accessible to a concurrent sibling. This is prevented throughout the type system, except in one place: at the join point of `par`. At this point, the two sub-tasks have completed and their parent inherits the values they allocated. To allow the parent task to access these values, TypeDis restamps the result of `par` with the timestamp of the parent. We dub this operation **backtiming**.

TypeDis features first-class function types $(\alpha \rightarrow^{\delta} \beta)$, annotated by a timestamp variable δ , indicating which task the function may be called by. Timestamp variables can be universally quantified, effectively allowing for **timestamp polymorphism**. For example, pure functions that have no side-effects are type-able as $(\forall \delta. \alpha \rightarrow^{\delta} \beta)$, indicating that the function may be safely called by any task. TypeDis also allows for **constrained timestamp polymorphism**. For example, a function of type $(\forall \delta' \preceq \delta. \text{string}@{\delta'} \rightarrow^{\delta} ())$ only accepts as argument strings timestamped at some δ' that precede the timestamp δ of the calling task. Typically, such constraints arise from the use of closures, especially those that close over mutable state.

Soundness. The soundness of TypeDis is verified in the Rocq prover (the new name of the Coq proof assistant) on top of the Iris higher-order concurrent separation logic framework [Jung et al. 2018b]. We use the approach of *semantic typing* [Constable et al. 1986; Martin-Löf 1982; Timany et al. 2024], and define a logical relation targeting a variation of DisLog [Moine et al. 2024], from which we reuse the technical parts. As illustrated by RustBelt [Jung et al. 2018a], semantic typing facilitates manual verification of programs that are correct (e.g. disentangled), but ill-typed, by carrying out a logical relation inhabitation proof using the program logic—overcoming incompleteness inherent to any type system. For example, in the case of TypeDis, this allows the user to verify part of the code that use down-pointers.

We note that, similar to many other type systems (such as those in OCaml and Haskell), TypeDis relies on dynamic checks in the operational semantics to enforce memory safety for out-of-bounds (OOB) array accesses. The formal statement of soundness (§5.1) therefore explicitly distinguishes between three kinds of program states: those that have terminated, those that can step, and those that are stuck due to OOB. The soundness theorem states that all executions of programs typed within TypeDis always remain disentangled throughout execution.

Contributions. Our contributions include:

- TypeDis, the first static type system for disentanglement. It includes the notion of a timestamp to track which object is accessible by which task. TypeDis offers (iso-)recursive types as well as polymorphism over types and over timestamps. Moreover, TypeDis supports polymorphic recursion over timestamps, and offers a relaxation of the value restriction.
- Two mechanisms to update a timestamp annotation: via subtiming, a form of subtyping, and specifically at join points via the new operation of backtiming.

- A new model for disentanglement with cyclic computation graphs. We prove this model equivalent to the standard one and explain why it is more amenable to verification.
- A soundness proof of TypeDis mechanized in the Rocq prover using the Iris framework. We use semantic typing [Timany et al. 2024] and DisLog₂, an improved version of DisLog.
- A range of case studies, including building and iterating over an immutable tree in parallel, as well the challenging example of deduplication via concurrent hashing.

2 Key Ideas

In this section, we cover the key ideas of our work. We start by recalling the definition of disentanglement (§2.1). We then present the main idea of TypeDis: adding task identifiers, specifically *timestamp variables*, to types (§2.2). Based on examples, we then illustrate two core principles of TypeDis, allowing for updating timestamps within types: backtiming (§2.3) and subtiming (§2.4).

2.1 Preliminaries

Nested fork-join parallelism and task trees. We consider programs written in terms of a single parallel primitive: $\text{par}(f_1, f_2)$, which creates two new child tasks $f_1()$ and $f_2()$ to execute in parallel, waits for both of the child tasks to complete, and then returns the results of the two calls as an immutable pair. Creating the two child tasks is called a *fork*, and waiting for the two children to complete is called a *join*. The behavior of the par primitive guarantees that every fork has a corresponding join. Any task may (recursively) fork and join, facilitating *nested* parallelism and giving rise to a dynamic tree during execution called the *task tree*. The nodes of the task tree correspond to (parent) tasks that are waiting for their children to join, and the leaves of the task tree correspond to tasks which may actively take a step. Whenever two sibling tasks join, the children are removed from the tree and the parent resumes as a leaf task. The task tree therefore dynamically grows and shrinks as tasks fork and join. In this paper, we will use the letter t to denote tasks (leaves of the task tree), and will equivalently refer to these as *timestamps*.

Computation graphs. The evolution of the task tree over time can be recorded as a *computation graph*, where vertices correspond to tasks and edges correspond to scheduling dependencies. The computation graph records not just the current task tree, but also the history of tasks that have joined. When a task t forks into two children t_1 and t_2 , two edges (t, t_1) and (t, t_2) are added to the graph; later when t_1 and t_2 join, two edges (t_1, t) and (t_2, t) are added to the graph. We say that t precedes t' in graph G and write $G \vdash t \preccurlyeq t'$, when there exists a sequence of edges from t to t' . Note that \preccurlyeq is reflexive. Two tasks are *concurrent* when neither precedes the other.

Cyclic versus standard computation graphs. We contribute a new definition of computation graphs, which we call the *cyclic approach*, that differs slightly from the standard presentation used in prior work [Acar et al. 2016; Moine et al. 2024; Westrick et al. 2020]. The standard approach is to use a fresh task identifier at each join point, effectively renaming the resumed parent task. In the cyclic approach, we instead use the same task identifier after the join point. Figure 2 illustrates the difference between the two approaches. It presents two computation graphs representing the same computation: Figure 2a shows the standard approach, and Figure 2b shows the (new) cyclic approach. The distinction occurs when two tasks join. In Figure 2a tasks t_3 and t_4 join and form a new task t'_2 whereas in Figure 2b the two tasks join by going back to task t_2 . This distinction occurs again when t'_2 (resp. t_2) join to form t'_0 (resp. t_0).

The cyclic approach considerably reduces the need to manipulate timestamps, not only in our proofs (for example the soundness proof of backtiming), but also in the design of the type system itself as well as in the underlying program logic (§5.3). We prove the two approaches equivalent for the purpose of verifying disentanglement [Moine et al. 2025a]. Intuitively, the two approaches

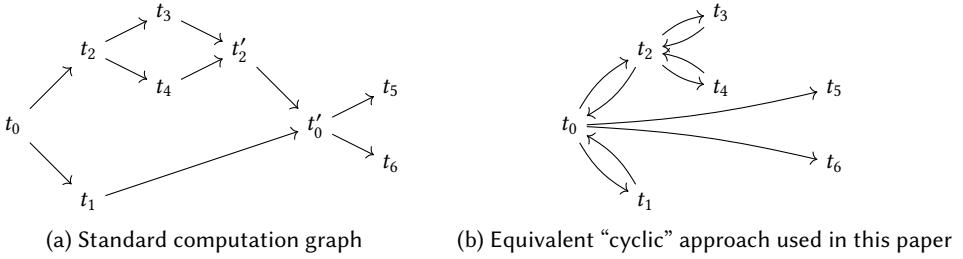


Fig. 2. Comparison of two computation graphs equivalent for disentanglement

are equivalent because we never need to check reachability between two tasks that have both completed. We have formally proven this equivalence with a simulation theorem: every reduction in the semantics with the standard approach implies the existence of a reduction with the same scheduling reaching the same expression in the semantics with the cyclic approach, and vice-versa. Moreover, if one state is disentangled in one semantics, so it is in the other.

Roots. At any moment, every task has a set of task-local **roots** which are the memory locations directly mentioned within a subexpression of that task. For example, the expression ‘let $x = (\ell_1, \ell_2)$ in $\text{fst}(x)$ ’ has roots $\{\ell_1, \ell_2\}$, where (formally) ℓ_1 and ℓ_2 are locations within the memory store. Note that the roots of a task change over time: for example, the above expression eventually steps to ℓ_1 at which point it only has one root, $\{\ell_1\}$. The set of roots can grow due to allocations and loads from memory.

Disentanglement. Disentanglement restricts the set of possible task-local roots. A **program state is disentangled if each root of a task has been allocated by some preceding task**. More precisely, a program state with a computation graph G is disentangled if, for a root ℓ of a task t , ℓ has been allocated by a task t' such that $G \vdash t' \preccurlyeq t$, that is, such that t' precedes t in G . Following the computation graph definition, preceding tasks include t itself, parent tasks, but also children tasks that have terminated. The formal definition of disentanglement appears in Section 3.3.

TypeDis, the type system we present, verifies that a program is disentangled, that is, every reachable program state is disentangled.

2.2 TypeDis 101: Timestamps in Types

In order to keep track of which task allocated which location, TypeDis incorporates timestamps in types. More precisely, every heap-allocated (“boxed”) type is annotated by a *timestamp variable*, written δ , which can be understood as the timestamp of the task that allocated the underlying location. For example, a reference allocated by task δ on an (unboxed) integer has type $\text{ref}(\text{int})@\delta$.

Timestamp polymorphism. Functions in TypeDis are annotated by a timestamp variable, restricting which task they may run on. Such a variable can be universally quantified, allowing for functions to be run by different tasks. For example, consider the function `fun` $x \rightarrow \text{newref}(x)$ which allocates a new mutable reference containing an integer x . This function can be given the type $\forall \delta. \text{int} \rightarrow^\delta \text{ref}(\text{int})@\delta$. The superscript δ on the arrow indicates that the function must run on a task at timestamp δ , and the result type $\text{ref}(\text{int})@\delta$ indicates that the resulting reference will be allocated at the same timestamp δ . By universally quantifying δ , the function is permitted to run on any task, with the type system tracking that the resulting reference will be allocated at the same timestamp as the caller.

```
newref :  $\forall \alpha. \forall \delta. \alpha \rightarrow^{\delta} \text{ref}(\alpha)@\delta$ 
get :  $\forall \alpha. \forall \delta \delta'. \text{ref}(\alpha)@\delta' \rightarrow^{\delta} \alpha$ 
set :  $\forall \alpha. \forall \delta \delta'. \text{ref}(\alpha)@\delta' \rightarrow \alpha \rightarrow^{\delta} ()$ 
```

Fig. 3. Example: typing reference primitives

```
let r = newref "hello"    r : ref(string@ $\delta_0$ )@ $\delta_0$ 
let w = "world"          w : string@ $\delta_0$ 
let f () = set r w      f : ( $\forall \delta. () \rightarrow^{\delta} ()$ )@ $\delta_0$ 
let g i =                g : (int  $\rightarrow^{\delta_0} ()$ )@ $\delta_0$ 
  set r (Int.to_string i)
```

Fig. 4. Example: typing closures

```
type tree@ $\delta$  = (int + (tree@ $\delta$   $\times$  tree@ $\delta$ )@ $\delta$ )@ $\delta$ 
let leaf x = inj1 x
let node x y = inj2 (x, y)
let rec build n x =
  if n <= 0 then leaf x else
  let n' = n - 1 in
  let (l, r) = par (fun () -> build n' x) (fun () -> build n' (x + pow2 n')) in
  node l r
```

Fig. 5. Example: building a tree in parallel

Type polymorphism. TypeDis allows type variables α to be universally quantified. Using type polymorphism, we can now give the function `fun x -> newref(x)` the more general type $\forall \alpha. \forall \delta. \alpha \rightarrow^{\delta} \text{ref}(\alpha)@\delta$, indicating that it is polymorphic in the type α of the contents of the mutable reference. Corresponding get and set primitives for mutable references are then typed as shown in Figure 3, all of which are polymorphic in the type variable α . For functions with multiple arguments, such as `set`, we adopt the notational convention to only specify the timestamp variable on the last arrow.

The up-pointer invariant. In Figure 3, the type of `get` is given as $\forall \alpha. \forall \delta \delta'. \text{ref}(\alpha)@\delta' \rightarrow^{\delta} \alpha$. Note that this type is parameterized over both a caller time δ as well as a (potentially different) timestamp δ' associated with the input reference. Intuitively, this type specifies that `get` is safe to call at any moment, by any task, with any reference given as argument. The design of TypeDis in general guarantees that all loads from memory, both mutable and immutable, are always safe. Specifically, this is guaranteed by enforcing an invariant that we call the **up-pointer invariant**: all data structures in the language may only contain values allocated at the same timestamp or a preceding timestamp. For example, given two non-equal timestamps δ_1 and δ_2 where $\delta_1 \prec \delta_2$, the type `ref(ref(int)@ δ_1)@ δ_2` is valid, but `ref(ref(int)@ δ_2)@ δ_1` is not.

Closures. In TypeDis, functions are first-class values and may be passed as arguments to other functions, or stored in data structures, etc. Function values are implemented as heap-allocated closures [Appel 1992; Landin 1964], and must be given a timestamp indicating when they were allocated. For example, consider the definition of function f in Figure 4, which closes over a mutable reference r and an immutable string w , both allocated at timestamps δ_0 which (in this example) is the timestamp of the current task. We can give f the type $(\forall \delta. () \rightarrow^{\delta} ())@\delta_0$, indicating that f itself was allocated at timestamp δ_0 . Additionally, the type of f specifies that it may be freely called at any timestamp; this is safe for disentanglement because f preserves the up-pointer invariant, regardless of *when* it will be called. Contrast this with the definition of function g , which (when called) allocates a new string and writes this string into the reference r . If g were called at some timestamp δ_1 where $\delta_0 \prec \delta_1$, then this would violate the up-pointer invariant for r . The function g does however admit the type $(\text{int} \rightarrow^{\delta_0} ())@\delta_0$, indicating that g may be safely called only by tasks at time δ_0 (the same timestamp as the reference r).

2.3 Backtimimg the Result of a par

As explained earlier (§2.1), we consider in this paper the parallel primitive `par(...)`, which executes two closures in parallel and returns their result as an immutable pair. The `par` primitive can be

```

let rec selectmap p f t =
  match t with
  | inj1 x -> if p x then leaf (f x) else t
  | inj2 (l,r) ->
    let (nl,nr) = par (fun () -> selectmap p f l) (fun () -> selectmap p f r) in
    if nl == l && nr == r then t else node nl nr

```

$\text{selectmap} : (\forall \delta_p \delta_f \delta_t. (\forall \delta'. \text{int} \rightarrow^{\delta'} \text{bool}) @ \delta_p \rightarrow (\forall \delta'. \text{int} \rightarrow^{\delta'} \text{int}) @ \delta_f \rightarrow \text{tree} @ \delta_t \rightarrow^{\delta} \text{tree} @ \delta) @ \delta$

Fig. 6. Example: the selectmap function

used to build data structures in parallel. Consider the code presented in [Figure 5](#). The recursive type $\text{tree} @ \delta = (\text{int} + (\text{tree} @ \delta \times \text{tree} @ \delta) @ \delta) @ \delta$ describes a binary tree with integer leaves. It consists of an immutable sum of either an integer (a leaf) or a product of two subtrees (a node). All the parts of a tree are specified in the type to have been allocated at the same timestamp δ . A leaf is built with the first injection, and a node with the second injection. The function $\text{build } n \times$ builds in parallel a binary tree of depth n , with leaves labeled from x to $x + 2^n - 1$ in left-to-right order.

TypeDis type-checks build with the type $\forall \delta. \text{int} \rightarrow \text{int} \rightarrow^{\delta} \text{tree} @ \delta$. The reader may be surprised: we announced that the type $\text{tree} @ \delta$ has all of its parts allocated at the same timestamp δ , but we are showing a function that builds a tree in parallel, hence with some parts allocated by different tasks at different timestamps. What's the trick?

The key observation is that *we can pretend that the objects allocated by a completed sub-task were instead allocated by its parent*. Indeed, disentanglement prevents sharing of data allocated in parallel, but as soon as the parallel phase has ended, there is no restriction anymore!

In TypeDis, the `par` primitive implements **backtiming**, meaning that it replaces the timestamp of the child task by the timestamp of the parent task in the return type of the closures executed in parallel. Indeed, the `par` primitive admits the following, specialized for `build`, type:

$$\forall \delta \delta_l \delta_r. (\forall \delta'. () \rightarrow^{\delta'} \text{tree} @ \delta') @ \delta_l \rightarrow (\forall \delta'. () \rightarrow^{\delta'} \text{tree} @ \delta') @ \delta_r \rightarrow^{\delta} (\text{tree} @ \delta \times \text{tree} @ \delta) @ \delta$$

This type for `par` does exactly what we need: it returns the result of the two closures in a pair as if they were called at time δ . Backtiming is a powerful feature: it reduces parallelism to almost an implementation detail. Indeed, the type of `build` does not reveal its internal use of parallelism.

2.4 Making Something New out of Something Old with Subtimimg

A common practice (especially in functional programming) is data structural *sharing*, where components of an old structure are reused inside part of a new structure. In the context of TypeDis, data structural sharing is interesting in that it mixes data of potentially different timestamps within the same structure. Here we consider one such example and describe a key feature of TypeDis which enables such “mixing” of timestamps.

[Figure 6](#) presents the `selectmap p f t` function, which selectively applies the function `f` to the leaves of the tree `t`, following a predicate `p` on integers. The `selectmap` function traverses the tree in parallel and crucially preserves sharing as much as possible. Specifically, when none of the leaves of the tree satisfy the predicate, the function returns the original input tree as-is, instead of building another identical tree. To type this function in TypeDis, it may not be immediately clear what the timestamp of the resulting tree should be: `selectmap` might directly return the argument passed as argument (potentially coming from an older task), or it might return a new tree. TypeDis type-checks `selectmap` with the type

$$\forall \delta \delta_p \delta_f \delta_t. (\forall \delta'. \text{int} \rightarrow^{\delta'} \text{bool}) @ \delta_p \rightarrow (\forall \delta'. \text{int} \rightarrow^{\delta'} \text{int}) @ \delta_f \rightarrow \text{tree} @ \delta_t \rightarrow^{\delta} \text{tree} @ \delta$$

This type universally quantifies over δ (the timestamp at which `selectmap` will run), δ_p and δ_f (the timestamps of the two closure arguments), and δ_t (the timestamp of the tree argument). Crucially, the result is of type $\text{tree} @ \delta$, as if the whole result tree was allocated by δ . What's the trick?

Values	$v, w ::= () \mid b \in \{\text{true}, \text{false}\} \mid i \in \mathbb{Z} \mid \ell \in \mathcal{L} \mid \text{vfold } v$	
Blocks	$r ::= \vec{w} \mid (v, v) \mid \text{inj}_{i \in \{0,1\}} v \mid \hat{\mu} f. \lambda \vec{x}. e$	
Primitives	$\bowtie ::= + \mid - \mid \times \mid \div \mid \text{mod} \mid == \mid < \mid \leq \mid > \mid \geq \mid \vee \mid \wedge$	
Expressions	$e ::= v \mid x \in \mathcal{V} \mid \text{let } x = e \text{ in } e \mid \text{if } e \text{ then } e \text{ else } e \mid e \bowtie e$ $\mid \mu f. \lambda \vec{x}. e \mid e \vec{e}$ $\mid (e, e) \mid \text{proj}_{i \in \{1,2\}} e$ $\mid \text{inj}_{i \in \{1,2\}} e \mid \text{match } e \text{ with } \text{inj}_1 x \Rightarrow e \mid \text{inj}_2 x \Rightarrow e \text{ end}$ $\mid \text{alloc } e \mid e.[e] \mid e.[e] \leftarrow e \mid \text{length } e$ $\mid \text{fold } e \mid \text{unfold } e$ $\mid \text{par}(e, e) \mid e \parallel e \mid \text{CAS } e \ e \ e \ e$	<i>closures</i> <i>pairs</i> <i>sums</i> <i>arrays</i> <i>iso-recursive types</i> <i>parallelism and concurrency</i>
Contexts	$K ::= \text{let } x = \square \text{ in } e \mid \text{if } \square \text{ then } e \text{ else } e \mid \text{alloc } \square \ e \mid \text{alloc } v \ \square \mid \text{length } \square$ $\mid \square.[e] \mid v.[\square] \mid \square.[e] \leftarrow e \mid v.[\square] \leftarrow e \mid v.[v] \leftarrow \square$ $\mid \square \bowtie e \mid v \bowtie \square \mid \square \vec{e} \mid v(\vec{v} + \square + \vec{e}) \mid \text{fold } \square$ $\mid \text{unfold } \square \mid (\square, e) \mid (v, \square) \mid \text{proj}_i \square \mid \text{inj}_i \square$ $\mid \text{match } \square \text{ with } \text{inj}_1 x \Rightarrow e \mid \text{inj}_2 x \Rightarrow e \text{ end} \mid \text{par}(\square, e) \mid \text{par}(v, \square)$ $\mid \text{CAS } \square \ e \ e \ e \mid \text{CAS } v \ \square \ e \ e \mid \text{CAS } v \ v \ \square \ e \mid \text{CAS } v \ v \ v \ \square$	

Fig. 7. Syntax of DisLang₂. Constructs in blue are runtime-level.

TypeDis supports **subtiming**, that is, a way of “advancing” timestamps within a type, following the precedence. The rules of subtiming are as follows. For a mutable type (e.g. an array or a reference), subtiming is *shallow*: the outermost timestamp can be updated, but not the inner timestamps; this is due to well-known variance issues [Pierce 2002, §15]. For an immutable type (e.g. products and sums), subtiming is *deep*: any timestamp within the type can be advanced, as long as the up-pointer invariant is preserved.

For `selectmap`, we need to use deep subtiming on the recursive immutable type $\text{tree}@\delta_t$, in order to update it to $\text{tree}@\delta$. How can we be sure that δ_t , the timestamp of the tree, precedes δ , the timestamp at which we call `selectmap`? We unveil a key invariant of TypeDis: every timestamp of every memory location in scope precedes the “current” timestamp, that is, the timestamp of the task executing the function. In our case the current timestamp is precisely δ . We hence deduce that δ_t precedes δ , allowing us to use subtiming to “restamp” the value $t : \text{tree}@\delta_t$ as $t : \text{tree}@\delta$.

To allow the user to express additional knowledge about the dependencies between timestamps, TypeDis annotates universal timestamp quantification with a set of *constraints*, which are supposed to hold while typing the function body, and are verified at call sites. For example, the following function `let` $\text{par}' f g = \text{ignore} (\text{par } f g)$ that executes two closures f and g from unit to unit in parallel and ignores the result can be given the type:

$$\forall \delta \delta_1 \delta_2. (\forall \delta' \delta \preceq \delta'. () \rightarrow^{\delta'} ()) @ \delta_1 \rightarrow (\forall \delta' \delta \preceq \delta'. () \rightarrow^{\delta'} ()) @ \delta_2 \rightarrow^{\delta} ()$$

This type says that, if par' gets called at timestamp δ with arguments f and g , then f and g can assume that they will be called at timestamp δ' such that $\delta \preceq \delta'$. These constraints are discussed in Section 4, and the fully general type of `par` is presented in Section 4.6.

3 Syntax and Semantics

The formal language we study, dubbed DisLang₂, can be understood as an extension of DisLang, the language studied by Moine et al. [2024]. DisLang₂ adds support for immutable pairs and sums, iso-recursive types, and directly offers the `par` primitive for fork-join parallelism. We present the syntax of DisLang₂ (§3.1), its semantics (§3.2), and the formal definition of disentanglement (§3.3).

3.1 Syntax

The syntax of DisLang_2 appears in Figure 7. The constructs in blue are forbidden in the source program and occur only at runtime.

A value $v \in \mathcal{V}$ can be the unit value $()$, a boolean $b \in \{\text{true}, \text{false}\}$, an idealized integer $i \in \mathbb{Z}$, a *memory location* $\ell \in \mathcal{L}$, where \mathcal{L} is an infinite set of locations, or a folded value $\text{vfold } v$, witnessing our use of iso-recursive types [Pierce 2002, §20].

A *block* describes the contents of a heap cell, amounting to either an array of values, written \vec{w} , an immutable pair (v, v) , the first injection $\text{inj}_1 v$ or the second injection $\text{inj}_2 v$ of an immutable sum, or a λ -*abstraction* $\mu f. \lambda \vec{x}. e$. Lambdas can close over free variables, compilers of functional languages usually implement them as *closures* [Appel 1992; Landin 1964]. A closure is a heap-allocated object carrying a code pointer as well as an environment, recording the values of the free variable. Thus, acquiring a closure can create entanglement. Moreover, because functions and tuples are heap allocated, currying and uncurrying—that is, converting a function taking multiple arguments to a function taking a tuple of arguments and vice-versa—does not come for free. Hence, we chose to support a version of the language where every function takes possibly multiple arguments. Closure allocation is written $\mu f. \lambda \vec{x}. e$. This notation binds a recursive name f , argument names \vec{x} in the expression e . A function call is written $e \vec{e}$.

In DisLang_2 , fork-join parallelism is available via the parallel primitive $\text{par}(e_1, e_2)$, which reduces e_1 and e_2 to closures, calls them in parallel, and returns their result as an immutable pair. This parallel computation is represented by the *active* parallel pair $e_1 \parallel e_2$, appearing only at runtime. DisLang_2 supports a compare-and-swap instruction $\text{CAS } e \ e \ e \ e$, which targets an array, and is parameterized by 4 arguments: the location of the array, the index in the array, the old value and the new value. A (sequential) evaluation context K describes a term with a hole, written \square . The syntax of evaluation contexts dictates a left-to-right *call-by-value* evaluation strategy. Note that evaluation contexts K in this presentation are sequential. Specifically, we intentionally excluded active parallel pairs $(- \parallel -)$ from the grammar of K . The evaluation strategy for active parallel pairs allows for interleaving of small steps, which is handled separately by a “scheduler reduction” relation in the operational semantics (§3.2).

3.2 Operational Semantics

Head reduction relation. A *head configuration* $\sigma \setminus \alpha \setminus e$ is composed of a store σ , an allocation map α , and an expression e . The store σ represents the heap and consists of a finite map of locations to blocks. The allocation map α is a finite map of locations to timestamps, recording the timestamps at which locations were allocated. Figure 8 presents parts of the definition of the head reduction relation between two head configurations $G, t \vdash \sigma \setminus \alpha \setminus e \longrightarrow \sigma' \setminus \alpha' \setminus e'$ occurring at the (local) task of timestamp t in the (global) computation graph G . A head configuration consists of the expression e being evaluated, the store σ , and an allocation map α . Figure 8 omits rules for the length array primitive as well as the atomic compare-and-swap on arrays.

We write $\sigma(\ell)$ to denote the block stored at the location ℓ in the store σ . We write $[\ell := r]\sigma$ for the insertion of block r at location ℓ in σ . Note that only arrays can be updated; closures, pairs and sums are immutable. We write $\vec{w}(i)$ to refer to the index i of an array \vec{w} . We write $[i := v]\vec{w}$ for an update to an array, and we similarly write $[\ell := t]\alpha$ for an insertion in the allocation map. We write v^n for an array of length n , where each element of the array is initialized with the value v .

HEADALLOC allocates an array, extending the store and the allocation map. **HEADLOAD** acquires the value v from an index of an array. **HEADSTORE**, **HEADLETVAL**, **HEADIFTRUE** and **HEADIFFALSE** are standard. **HEADCLOSURE** allocates a closure and **HEADCALL** calls a closure. **HEADCALLPRIM** calls a primitive, whose result is computed at the meta-level by the $\xrightarrow{\text{pure}}$ relation. **HEADPAIR** and

$\begin{array}{c} \text{HEADALLOC} \\ \dfrac{0 < n \quad \ell \notin \text{dom}(\sigma) \quad \ell \notin \text{dom}(\alpha)}{G, t \vdash \sigma \setminus \alpha \setminus \text{alloc } n \ v \longrightarrow [\ell := v^n] \sigma \setminus [\ell := t] \alpha \setminus \ell} \end{array}$	$\begin{array}{c} \text{HEADLOAD} \\ \dfrac{\sigma(\ell) = \vec{w} \quad 0 \leq i < \vec{w} \quad \vec{w}(i) = v}{G, t \vdash \sigma \setminus \alpha \setminus \ell.[i] \longrightarrow \sigma \setminus \alpha \setminus v} \end{array}$
$\begin{array}{c} \text{HEADSTORE} \\ \dfrac{\sigma(\ell) = \vec{w} \quad 0 \leq i < \vec{w} }{G, t \vdash \sigma \setminus \alpha \setminus \ell.[i] \leftarrow v \longrightarrow [\ell := [i := v] \vec{w}] \sigma \setminus \alpha \setminus ()} \end{array}$	$\begin{array}{c} \text{HEADLETVAL} \\ \dfrac{G, t \vdash \sigma \setminus \alpha \setminus \text{let } x = v \text{ in } e \longrightarrow \sigma \setminus \alpha \setminus [v/x]e}{G, t \vdash \sigma \setminus \alpha \setminus \ell.[i] \leftarrow v \longrightarrow [\ell := [i := v] \vec{w}] \sigma \setminus \alpha \setminus ()} \end{array}$
$\begin{array}{c} \text{HEADIFTRUE} \\ \dfrac{G, t \vdash \sigma \setminus \alpha \setminus \text{if true then } e_1 \text{ else } e_2 \longrightarrow \sigma \setminus \alpha \setminus e_1}{G, t \vdash \sigma \setminus \alpha \setminus \ell.[i] \leftarrow v \longrightarrow [\ell := [i := v] \vec{w}] \sigma \setminus \alpha \setminus ()} \end{array}$	$\begin{array}{c} \text{HEADIFFALSE} \\ \dfrac{G, t \vdash \sigma \setminus \alpha \setminus \text{if false then } e_1 \text{ else } e_2 \longrightarrow \sigma \setminus \alpha \setminus e_2}{G, t \vdash \sigma \setminus \alpha \setminus \ell.[i] \leftarrow v \longrightarrow [\ell := [i := v] \vec{w}] \sigma \setminus \alpha \setminus ()} \end{array}$
$\begin{array}{c} \text{HEADCLOSURE} \\ \dfrac{\ell \notin \text{dom}(\sigma) \quad \ell \notin \text{dom}(\alpha)}{G, t \vdash \sigma \setminus \alpha \setminus \hat{\mu}f. \lambda \vec{x}. e \longrightarrow [\ell := \hat{\mu}f. \lambda \vec{x}. e] \sigma \setminus [\ell := t] \alpha \setminus \ell} \end{array}$	$\begin{array}{c} \text{HEADCALL} \\ \dfrac{\sigma(\ell) = \hat{\mu}f. \lambda \vec{x}. e \quad \vec{x} = \vec{w} }{G, t \vdash \sigma \setminus \alpha \setminus \ell \vec{w} \longrightarrow \sigma \setminus \alpha \setminus [\ell / f] [\vec{w} / \vec{x}] e} \end{array}$
$\begin{array}{c} \text{HEADCALLPRIM} \\ \dfrac{v_1 \bowtie v_2 \xrightarrow{\text{pure}} v}{G, t \vdash \sigma \setminus \alpha \setminus v_1 \bowtie v_2 \longrightarrow \sigma \setminus \alpha \setminus v} \end{array}$	$\begin{array}{c} \text{HEADPAIR} \\ \dfrac{\ell \notin \text{dom}(\sigma) \quad \ell \notin \text{dom}(\alpha)}{G, t \vdash \sigma \setminus \alpha \setminus (v_1, v_2) \longrightarrow [\ell := (v_1, v_2)] \sigma \setminus [\ell := t] \alpha \setminus \ell} \end{array}$
$\begin{array}{c} \text{HEADPROJ} \\ \dfrac{\sigma(\ell) = (v_1, v_2)}{G, t \vdash \sigma \setminus \alpha \setminus \text{proj}_i \ell \longrightarrow \sigma \setminus \alpha \setminus v_i} \end{array}$	$\begin{array}{c} \text{HEADINJ} \\ \dfrac{\ell \notin \text{dom}(\sigma) \quad \ell \notin \text{dom}(\alpha)}{G, t \vdash \sigma \setminus \alpha \setminus \text{inj}_i v \longrightarrow [\ell := \text{inj}_i v] \sigma \setminus [\ell := t] \alpha \setminus \ell} \end{array}$
$\begin{array}{c} \text{HEADCASE} \\ \dfrac{}{G, t \vdash \sigma \setminus \alpha \setminus (\text{match } \ell \text{ with } \text{inj}_1 x_1 \Rightarrow e_1 \mid \text{inj}_2 x_2 \Rightarrow e_2 \text{ end}) \longrightarrow \sigma \setminus \alpha \setminus [v / x_i] e_i} \end{array}$	$\begin{array}{c} \text{HEADFOLD} \\ \dfrac{G, t \vdash \sigma \setminus \alpha \setminus \text{fold } v \longrightarrow \sigma \setminus \alpha \setminus \text{vfold } v}{G, t \vdash \sigma \setminus \alpha \setminus \text{unfold } (\text{vfold } v) \longrightarrow \sigma \setminus \alpha \setminus v} \end{array}$
$\begin{array}{c} \text{HEADUNFOLD} \\ \dfrac{G, t \vdash \sigma \setminus \alpha \setminus \text{unfold } (\text{vfold } v) \longrightarrow \sigma \setminus \alpha \setminus v}{G, t \vdash \sigma \setminus \alpha \setminus \text{unfold } (\text{vfold } v) \longrightarrow \sigma \setminus \alpha \setminus v} \end{array}$	

Fig. 8. Head reduction (selected rules)

HEADPROJ allocate and project immutable pairs, respectively. **HEADINJ** and **HEADCASE** allocate and case over immutable sums, respectively. **HEADFOLD** and **HEADUNFOLD** handle iso-recursive types in a standard way.

Scheduler reduction relation. In order to keep track of the timestamp of each task and whether the task is activated or suspended, we follow Westrick et al. [2020] and enrich the semantics with an auxiliary structure called a *task tree*, written T , of the following formal grammar: $T \triangleq t \in \mathcal{T} \mid T \otimes_t T$. A leaf t indicates an active task denoted by its timestamp. A node $T_1 \otimes_t T_2$ represents a suspended task t that has forked two parallel computations, recursively described by the task trees T_1 and T_2 .

Figure 9 presents the scheduling reduction relation $\sigma / \alpha / G / T / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G' / T' / e'$ as either a head step, a fork, or a join. In this reduction relation, σ is a store, α an allocation map, G a computation graph, T a task tree, and e an expression. **SCHEDHEAD** reduction describes a head reduction. **SCHEDFORK** reduction describes a fork: the task tree consists of a leaf t and the expression $\text{par}(v_1, v_2)$, where both v_1 and v_2 are closures to be executed in parallel. The reduction generates two fresh timestamps t_1 and t_2 , adds the corresponding edges to the computation graph, and updates the task tree to comprise the node with two leaves $t_1 \otimes_t t_2$. The reduction then updates the expression to the active parallel pair $v_1 [()] \parallel v_2 [()]$, reflecting the parallel call of the two closures v_1 and v_2 , each one called with a single argument, the unit value $()$. **SCHEDJOIN** reduction describes a join and differs from prior semantics for disentanglement [Moine et al. 2024; Westrick et al. 2022] because it reuses a timestamp (§2.1). The task tree is at a node t with two leaves $t_1 \otimes_t t_2$, and both leaves reached a value. The reduction adds edges (t_1, t) and (t_2, t) to the computation

$$\begin{array}{c}
\text{SCHEDHEAD} \quad \text{SCHEDFORK} \\
\frac{G, t \vdash \sigma \setminus \alpha \setminus e \longrightarrow \sigma' \setminus \alpha' \setminus e'}{\sigma / \alpha / G / t / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G / t / e'} \quad \frac{t_1, t_2 \notin \text{vertices}(G)}{G' = G \cup \{(t, t_1), (t, t_2)\} \quad e' = v_1 [() \parallel v_2 [()]} \\
\text{SCHEDJOIN} \quad \text{STEPBIND} \\
\frac{\ell \notin \text{dom}(\sigma) \quad \ell \notin \text{dom}(\alpha) \quad G' = G \cup \{(t_1, t), (t_2, t)\}}{\sigma / \alpha / G / t_1 \otimes_t t_2 / v_1 \parallel v_2 \xrightarrow{\text{sched}} [\ell := (v_1, v_2)] \sigma / [\ell := t] \alpha / G' / t / \ell} \\
\text{STEPSCHEDE} \quad \text{STEPBIND} \\
\frac{\sigma / \alpha / G / T / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G' / T' / e'}{(\sigma, \alpha, G) / T / e \xrightarrow{\text{step}} (\sigma', \alpha', G') / T' / e'} \quad \frac{S / T / e \xrightarrow{\text{step}} S' / T' / e'}{S / T / K[e] \xrightarrow{\text{step}} S' / T' / K[e']} \\
\text{STEPPARL} \quad \text{STEPPARR} \\
\frac{S / T_1 / e_1 \xrightarrow{\text{step}} S' / T'_1 / e'_1}{S / T_1 \otimes_t T_2 / e_1 \parallel e_2 \xrightarrow{\text{step}} S' / T'_1 \otimes_t T_2 / e'_1 \parallel e_2} \quad \frac{S / T_2 / e_2 \xrightarrow{\text{step}} S' / T'_2 / e'_2}{S / T_1 \otimes_t T_2 / e_1 \parallel e_2 \xrightarrow{\text{step}} S' / T_1 \otimes_t T'_2 / e_1 \parallel e'_2}
\end{array}$$

Fig. 9. Reduction under a context and parallelism

$$\begin{array}{c}
\text{DELEAF} \quad \text{DEPAR} \\
\frac{\forall \ell. \ell \in \text{roots}(e) \implies G \vdash \alpha(\ell) \preccurlyeq t}{\text{Disentangled } (_, \alpha, G) / t / e} \quad \frac{\text{Disentangled } S / T_1 / e_1 \quad \text{Disentangled } S / T_2 / e_2}{\text{Disentangled } S / T_1 \otimes_t T_2 / e_1 \parallel e_2} \\
\text{DEBIND} \\
\frac{\begin{array}{c} S = (_, \alpha, G) \quad \text{Disentangled } S / T_1 \otimes_t T_2 / e \\ \forall \ell. \ell \in \text{roots}(K) \implies \forall t'. t' \in \text{leaves}(T_1) \cup \text{leaves}(T_2) \implies G \vdash \alpha(\ell) \preccurlyeq t' \end{array}}{\text{Disentangled } S / T_1 \otimes_t T_2 / K[e]}
\end{array}$$

Fig. 10. Definition of Disentanglement

graph, and allocates a memory cell to store the result of the (active) parallel pair. It then updates the task tree to the leaf t .

Parallelism and reduction under a context. The lower part of Figure 9 presents the main reduction relation $S / T / e \xrightarrow{\text{step}} S' / T' / e'$, which describes a scheduling reduction inside the whole parallel program [Moine et al. 2024]. A configuration $S / T / e$ consists of the program state S , the task tree T , and an expression e . This expression e can consist of multiple tasks, governed by the nesting of active parallel pairs ($e_1 \parallel e_2$). The corresponding timestamps of these tasks are given by the accompanying task tree T . A state S consists of the tuple (σ, α, G) , denoting a store σ , an allocation map α , and a computation graph G . **STEPSCHEDE** reduction describes a scheduling step. The other reductions describe *where* the scheduling reduction takes place in the whole parallel program. **STEPBIND** reduction describes a reduction under an evaluation context. **STEPPARL** and **STEPPARR** reductions are non-deterministic: if a node of the task tree is encountered facing an active parallel pair, the left side or the right side can reduce.

3.3 Definition of Disentanglement

The property $\text{Disentangled } S / T / e$ asserts that, given a program state S and a task tree T , the expression e is disentangled—that is, the roots of each task in e were allocated by preceding tasks. Figure 10 gives the inductive definition of $\text{Disentangled } S / T / e$. If the program state has an allocation map α and a computation graph G , and if the task tree is a leaf t , **DELEAF** requires

Timestamp variables δ
Type variables α
Logical graphs $\Delta \triangleq \Delta, \delta \preceq \delta \mid \emptyset$
Kinds $\kappa \triangleq \star \mid \mathbf{x} \Rightarrow \kappa$
Unboxed types $\tau \triangleq () \mid \mathbf{bool} \mid \mathbf{int}$
Boxed types $\sigma \triangleq \mathbf{array}(\rho) \mid (\rho \times \rho) \mid (\rho + \rho) \mid \forall \vec{\delta} \Delta. \vec{\rho} \rightarrow^{\delta} \rho$
Types $\rho \triangleq \tau \mid \lambda \delta. \rho \mid \rho \delta \mid \forall \alpha :: \kappa. \rho \mid \mu \alpha. \sigma @ \delta \mid \alpha \mid \sigma @ \delta$
Environments $\Gamma \triangleq x : \rho, \Gamma \mid \alpha :: \kappa, \Gamma \mid \emptyset$

Fig. 11. Syntax of types

for every location ℓ in $\text{roots}(e)$, that is, the set of locations syntactically occurring in e , that the location ℓ has been allocated by a task $\alpha(t)$ preceding t in G . If the task tree is a node $T_1 \otimes_t T_2$, there are two cases. In the first case, if the expression is an active parallel pair, **DEPAR** requires that the two sub-expressions are disentangled. Otherwise, the expression must be of the form $K[e]$, and then **DEBIND** requires that e itself is disentangled and that for every location ℓ occurring in the evaluation context K , the location ℓ has been allocated before every leaf ℓ' of T_1 and T_2 .

Difference with Previous Semantics for Disentanglement. Inspired by Westrick et al. [2022], we equip DisLang₂ with a mostly standard semantics, instrumented with a computation graph and an allocation map. We then distinguish disentangled states using the Disentangled property, resembling the “rootsde” invariant proposed by Westrick et al. [2022]. The novelty of our approach resides in the instrumentation with the, more amenable to verification, cyclic computation graph (§ 2.1).

DisLog [Moine et al. 2024] chooses a slightly different formalization in which the semantics gets stuck if entanglements is detected. Each time a task acquires a location from the heap, the semantics performs a check to verify that the location was allocated by a preceding task. Intuitively, this check ensures by construction that a program’s evaluation reaches only states satisfying the Disentangled property. Conversely, guaranteeing the Disentangled property at every step ensures that a disentanglement check cannot fail.

4 Type System

In this section, we describe TypeDis in depth. First, we present the formal syntax of types (§ 4.1) as well as the typing judgment (§ 4.2). We then comment on typing rules for mutable heap blocks (§ 4.3), which enforce disentanglement. Next, we present the rules for creating and calling closures (§ 4.4), which are crucial for understanding our approach for typing the par primitive (§ 4.5). We then focus on advanced features of TypeDis: general recursive types and type polymorphism (§ 4.6). We conclude by presenting subtyping (§ 4.7).

4.1 Syntax of Types

To reason statically about the runtime notions of timestamps t and computation graphs G (§ 3.2), we introduce their corresponding static notions: *timestamp variables* δ and *logical graphs* Δ , respectively. A logical graph Δ is a set of pairs $\delta_1 \preceq \delta_2$, asserting that the timestamp δ_1 *precedes* the timestamp δ_2 , that is, everything allocated by the task at δ_1 is safe to acquire for the task at δ_2 . Figure 11 summarizes these notions together with the syntax of types.

A powerful feature of our type system is its support for *timestamp polymorphism*, facilitated through *higher-order types*. This higher-order feature is instrumental in typing the par primitive (§ 4.5), and thus supporting the cyclic approach detailed in § 2.1. Because our system is higher-order, we introduce *kinds*, written κ , which capture the number of timestamps a type expects as

arguments. The ground kind, written \star , indicates that the type does not take a timestamp argument. The successor kind, written $\mathbf{x} \Rightarrow \kappa$, indicates that the type expects $\kappa + 1$ timestamp arguments.

A base type τ describes an *unboxed* value, that is, a value that is not allocated on the heap. Base types include the unit type, Booleans, and integers.

The syntax of types ρ is mutually inductive with the syntax of boxed types σ . A type ρ is either a base type τ , a type taking a timestamp argument $\lambda\delta.\rho$, an application of a type to a timestamp $\rho\delta$, a universal quantification of a type variable with some kind $\forall\alpha :: \kappa.\rho$, a recursive type $\mu\alpha.\sigma@\delta$, a type variable α , or a boxed type annotated with a timestamp $\sigma@\delta$. When the timestamp δ does not matter, we write $\sigma@_$. A boxed type σ is either an array $\text{array}(\rho)$, an immutable pair $(\rho \times \rho)$, an immutable sum $(\rho + \rho)$, or a function $\forall\delta\Delta.\vec{\rho} \rightarrow^\delta \rho$.

Types support α -equivalence for both type and timestamp variables, as well as β -reduction.

4.2 The Typing Judgment

A typing environment Γ is a map from free program variables to types, and from free type variables to kinds. The general form of the typing judgment of TypeDis is:

$$\Delta \mid \Gamma \vdash e : \rho \triangleright \delta$$

where Δ is a logical graph, Γ a typing environment, e the expression being type-checked at type ρ and at current timestamp δ .

Selected rules of the type system appear in Figure 12. The rules adopt Barendregt’s convention [Barendregt 1984], assuming bound variables to be distinct from already existing free variables in scope. The reader might notice that several rules (for example, **T-ABS** or **T-TABS**) require the user to manually decide where to apply these rules and with which arguments. We leave to future work the design of syntactic features together with a type inference mechanism for simplifying this process. Various rules are standard: **T-VAR** type-checks variables and **T-UNIT**, **T-INT**, and **T-BOOL** type-check base types. The structural rules **T-LET** and **T-IF** are also standard, and type-check let bindings and if statements, respectively. In the remainder, we discuss the rules that deserve special attention with regard to disentanglement.

4.3 Typing Rules for Heap Blocks

Heap blocks must be handled with care to guarantee disentanglement: every time the program acquires a location—that is, the address of a heap block—we must ensure that this location has been allocated by a preceding task. Otherwise, this newly created root would break the disentanglement invariant (§3.3). Because load operations are so common in programming languages, we chose to enforce the following invariant on the typing judgment $\Delta \mid \Gamma \vdash e : \rho \triangleright \delta$: every location that can be acquired from Γ was allocated before the current timestamp δ (§2.4). Hence, load operations (from immutable blocks and from mutable blocks) do not have any timestamp check.

Operations on immutable blocks are type-checked by **T-PAIR** and **T-PROJ**, for pairs, and by **T-INJ** and **T-CASE**, for sums. In particular, **T-PAIR** and **T-INJ** reflect that pair creation and injection *allocate heap blocks*, hence, the resulting type is annotated with $@\delta$, denoting the allocating timestamp.

Operations on mutable blocks are type-checked by **T-ARRAY**, **T-STORE**, and **T-LOAD**.

4.4 Abstractions and Timestamp Polymorphism

A function can be seen as a *delayed computation*. In our case, this notion of “delay” plays an interesting role: a function can run on a task distinct from the one that allocated it. Hence, functions in TypeDis have three non-standard features related to timestamps, roughly describing the status of the computation graph when the function will run. First, a function takes timestamp parameters,

$\begin{array}{c} \text{T-VAR} \\ \frac{\Gamma(x) = \rho}{\Delta \mid \Gamma \vdash x : \rho \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-UNIT} \\ \frac{}{\Delta \mid \Gamma \vdash () : \text{unit} \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-INT} \\ \frac{}{\Delta \mid \Gamma \vdash i : \text{int} \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-BOOL} \\ \frac{}{\Delta \mid \Gamma \vdash b : \text{bool} \triangleright \delta} \end{array}$
$\begin{array}{c} \text{T-LET} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \rho' \triangleright \delta \quad \Delta \mid x : \rho', \Gamma \vdash e_2 : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash \text{let } x = e_1 \text{ in } e_2 : \rho \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-IF} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \text{bool} \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : \rho \triangleright \delta \quad \Delta \mid \Gamma \vdash e_3 : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \rho \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-PAIR} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \rho_1 \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : \rho_2 \triangleright \delta}{\Delta \mid \Gamma \vdash (e_1, e_2) : (\rho_1 \times \rho_2) @ \delta \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-PROJ} \\ \frac{\Delta \mid \Gamma \vdash e : (\rho_1 \times \rho_2) @ _ \triangleright \delta}{\Delta \mid \Gamma \vdash \text{proj}_i e : \rho_i \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-INJ} \\ \frac{\Delta \mid \Gamma \vdash e : \rho_i \triangleright \delta}{\Delta \mid \Gamma \vdash \text{inj}_i e : (\rho_1 + \rho_2) @ \delta \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-CASE} \\ \frac{\Delta \mid \Gamma \vdash e : (\rho_1 + \rho_2) @ _ \triangleright \delta \quad \Delta \mid x_1 : \rho_1, \Gamma \vdash e_1 : \rho \triangleright \delta \quad \Delta \mid x_2 : \rho_2, \Gamma \vdash e_2 : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash \text{match } e \text{ with } \text{inj}_1 x_1 \Rightarrow e_1 \mid \text{inj}_2 x_2 \Rightarrow e_2 \text{ end} : \rho \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-ARRAY} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \text{int} \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash \text{alloc } e_1 e_2 : \text{array}(\rho) @ \delta \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-STORE} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \text{array}(\rho) @ _ \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : \text{int} \triangleright \delta \quad \Delta \mid \Gamma \vdash e_3 : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash e_1.[e_2] \leftarrow e_3 : () \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-LOAD} \\ \frac{\Delta \mid \Gamma \vdash e_1 : \text{array}(\rho) @ _ \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : \text{int} \triangleright \delta}{\Delta \mid \Gamma \vdash e_1.[e_2] : \rho \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-ABS} \\ \frac{\Delta, \Delta_1, \delta \preccurlyeq \delta_f \mid f : (\forall \vec{\delta}_1 \Delta_1. \vec{\rho}_1 \rightarrow^{\delta_f} \rho_2) @ \delta, (\vec{x} : \vec{\rho}_1), \Gamma \vdash e : \rho_2 \triangleright \delta_f}{\Delta \mid \Gamma \vdash \mu f. \lambda \vec{x}. e : (\forall \vec{\delta}_1 \Delta_1. \vec{\rho}_1 \rightarrow^{\delta_f} \rho_2) @ \delta \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-APP} \\ \frac{\delta = [\vec{\delta}_1 / \vec{\delta}_1] \delta_f \quad \vec{\rho}'_1 = [\vec{\delta}'_1 / \vec{\delta}_1] \vec{\rho}_1 \quad \rho'_2 = [\vec{\delta}'_1 / \vec{\delta}_1] \rho_2 \quad \Delta'_1 = [\vec{\delta}'_1 / \vec{\delta}_1] \Delta_1 \quad \Delta \mid \Gamma \vdash e : (\forall \vec{\delta}_1 \Delta_1. \vec{\rho}_1 \rightarrow^{\delta_f} \rho_2) @ _ \triangleright \delta \quad \Delta \vdash \Delta'_1 \quad \Delta \mid \Gamma \vdash e' : \vec{\rho}'_1 \triangleright \delta}{\Delta \mid \Gamma \vdash e e' : \rho'_2 \triangleright \delta} \end{array}$			
$\begin{array}{c} \text{T-PAR} \\ \frac{\Delta \mid \Gamma \vdash e_1 : (\forall \delta' \delta \preccurlyeq \delta'. () \rightarrow^{\delta'} \varphi_1 \delta') @ _ \triangleright \delta \quad \Delta \mid \Gamma \vdash e_2 : (\forall \delta' \delta \preccurlyeq \delta'. () \rightarrow^{\delta'} \varphi_2 \delta') @ _ \triangleright \delta}{\Delta \mid \Gamma \vdash \text{par}(e_1, e_2) : (\varphi_1 \delta \times \varphi_2 \delta) @ \delta \triangleright \delta} \end{array}$			
$\begin{array}{c} \text{T-FOLD} \\ \frac{\Gamma \vdash \mu \alpha. \sigma @ \delta :: \star \quad \Delta \mid \Gamma \vdash e : ([\mu \alpha. \sigma @ \delta / \alpha] \sigma) @ \delta \triangleright \delta}{\Delta \mid \Gamma \vdash \text{fold } e : \mu \alpha. \sigma @ \delta \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-UNFOLD} \\ \frac{\Gamma \vdash \mu \alpha. \sigma @ \delta :: \star \quad \Delta \mid \Gamma \vdash e : \mu \alpha. \sigma @ \delta \triangleright \delta}{\Delta \mid \Gamma \vdash \text{unfold } e : ([\mu \alpha. \sigma @ \delta / \alpha] \sigma) @ \delta \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-TABS} \\ \frac{\Delta \mid \alpha :: \kappa, \Gamma \vdash e : \rho \triangleright \delta \quad \text{veryPure } e}{\Delta \mid \Gamma \vdash e : \forall \alpha :: \kappa. \rho \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-TAPP} \\ \frac{\Gamma \vdash \rho' :: \kappa \quad \Delta \mid \Gamma \vdash e : \forall \alpha :: \kappa. \rho \triangleright \delta}{\Delta \mid \Gamma \vdash e : [\rho' / \alpha] \rho \triangleright \delta} \end{array}$		
$\begin{array}{c} \text{T-GETROOT} \\ \frac{\Gamma(x) = \text{array}(\sigma') @ \delta' \vee \Gamma(x) = \mu \alpha. \sigma' @ \delta' \quad \Delta, \delta' \preccurlyeq \delta \mid \Gamma \vdash e : \rho \triangleright \delta}{\Delta \mid \Gamma \vdash e : \rho \triangleright \delta} \end{array}$	$\begin{array}{c} \text{T-SUBTIMING} \\ \frac{\Delta \mid \Gamma \vdash e : \rho \triangleright \delta \quad \Delta \vdash \rho \subseteq_{\delta} \rho'}{\Delta \mid \Gamma \vdash e : \rho' \triangleright \delta} \end{array}$		

Fig. 12. The type system (selected rules)

which are universally quantified. Second, a function takes a constraint over these timestamps, as a logical graph. Third, a function is annotated with a timestamp representing the task it will run on.

Let us focus on the abstraction rule **T-ABS**. This rule type-checks a function definition of the form $\mu f. \lambda \vec{x}. e$, and requires the user to provide timestamp parameters $\vec{\delta}_1$, logical graph Δ_1 , and a running timestamp δ_f . The current timestamp is δ and the type associated to the function is $(\forall \vec{\delta}_1 \Delta_1. \vec{\rho}_1 \rightarrow^{\delta_f} \rho_2) @ \delta$. This type asserts that, if (i) there is some instantiation of $\vec{\delta}_1$ satisfying Δ_1 , (ii) there are some arguments of type $\vec{\rho}_1$, and (iii) the timestamp of the calling task is δ_f , then the function will produce a result of type ρ_2 . This type also reminds us that a function is a heap-allocated object, and is hence annotated with the task that allocated it, here δ . The premise of **T-ABS** changes the current timestamp to be δ_f , the timestamp of the invoking task, and requires the body e to be of type ρ_2 . **T-ABS** is in fact the sole rule of the system “changing” the current timestamp while type-checking. The logical graph is augmented with Δ_1 plus the knowledge that δ precedes δ_f , conveying the fact that a function can only be called at a subsequent timestamp. The environment Γ is extended with the parameters $(\vec{x} : \vec{\rho}_1)$ as well as the recursive name f . Note that timestamp parameters $\vec{\delta}_1$ and logical graph Δ_1 are before the arguments \vec{x} . This means that the body e will be able to recursively call f with different timestamp arguments (potentially including a different δ_f), for example after it forked. We already saw such an example for the build and selectmap functions in [Sections 2.3](#) and [2.4](#).

Let us now focus on **T-APP**, type-checking a function application. The conclusion type-checks the expression $e \vec{e}'$ to be of type ρ'_2 at the current timestamp δ . The premise of **T-APP** requires e to be a function of type $\forall \vec{\delta}_1 \Delta_1. \vec{\rho}_1 \rightarrow^{\delta_f} \rho_2$, allocated by some irrelevant task. The premise then substitutes in all the relevant parts the user-supplied timestamps $\vec{\delta}'_1$ in place of $\vec{\delta}_1$. Hence, the result type ρ'_2 is equal to $[\vec{\delta}'_1 / \vec{\delta}_1] \rho_2$. In particular, the premise $\delta = [\vec{\delta}'_1 / \vec{\delta}_1] \delta_f$ requires that the running timestamp δ_f to be equal to δ , the current timestamp. The premise also requires the logical graph Δ'_1 to be a subgraph of the logical graph Δ , written $\Delta \vdash \Delta'_1$, meaning that every pair of vertices reachable in Δ'_1 must also be reachable in Δ . This property is formally defined in the Appendix [[Moine et al. 2025b](#)]. Finally, the premise requires the arguments \vec{e}' to be of the correct type $\vec{\rho}'_1$.

4.5 The Par Rule

The typing rule for the par primitive is at the core of TypeDis. **T-PAR** type-checks $\text{par}(e_1, e_2)$ at current timestamp δ . Recall ([§3.2](#)) that the results of e_1 and e_2 must be closures; these closures are then called in parallel and their results are returned as an immutable pair. To preserve disentanglement, the two closures must not communicate allocations they make with each other. Hence, the premise of **T-PAR** requires the two expressions e_1 and e_2 to be of type $\forall \delta' \delta \preceq \delta'. () \rightarrow^{\delta'} \dots$, signaling that they must be closures that are expected to run on a task δ' , universally quantified, and subsequent to δ . Because of this universal quantification over the running timestamp δ' and because the rules allocating blocks (**T-ARRAY**, **T-PAIR**, **T-PROJ** and **T-ABS**) always tag the value they allocate with the running timestamp, the tasks will not be able to communicate allocations they make.

After these two closure calls terminate, and their underlying tasks join, the parent task gains access to everything the two children allocated. In fact, from the point of view of disentanglement, we can even pretend that the parent task itself allocated these locations! **T-PAR** does more than pretending and **backtimes** the return types of the two closures, by substituting the running timestamp of the children δ' by the running timestamp of the parent δ . Indeed, the return types of the closures, $\varphi_1 \delta'$ for e_1 and $\varphi_2 \delta'$ for e_2 , signal that these two closures will return some type, parametrized by the running timestamp δ' . This formulation allows the rule to type-check the original $\text{par}(e_1, e_2)$ as $(\varphi_1 \delta \times \varphi_2 \delta) @ \delta$, that is, a pair of the two types returned by the closures, *but where the running timestamp of the child δ' was replaced by the running timestamp of the parent δ* .

4.6 Recursive Types and Type Polymorphism

Recursive types. TypeDis supports iso-recursive types [Pierce 2002, §20.2]. In TypeDis, a recursive type takes the form $\mu\alpha. \sigma@\delta$, binding the recursive name α in the boxed type σ which must have been allocated at δ . This syntax ensures that types are well-formed, and forbids meaningless types $\mu\alpha. \alpha$ as well as useless types $\mu\alpha. \mu\beta. \rho$. **T-FOLD** and **T-UNFOLD** allow for going from $\mu\alpha. \sigma@\delta$ to $([\mu\alpha. \sigma@\delta/\alpha]\sigma)@\delta$ and vice-versa. Note that this approach requires that the recursive occurrences of α are all allocated at the same timestamp; all the nodes of the recursive data structures must have been allocated at the same timestamp. This may seem restrictive, but subtyping will relax this requirement (§4.7).

Let us give an example. The type of lists allocated at timestamp δ containing integers is:

$$\mu\alpha. ((()) + (\text{int} \times \alpha)@\delta)@\delta$$

This type describes that a list of integers is either the unit value (describing the nil case), or the pair of an integer and a list of integers (describing the cons case).

Type polymorphism. TypeDis supports type polymorphism, through type abstraction **T-TABS** and type application **T-TAPP**. Whereas the former is standard, the latter has an unusual premise `veryPure e`, our variant of the value restriction.

The *value restriction* [Wright 1995] is a simple syntactic restriction guaranteeing soundness of polymorphism in the presence of mutable state—a combination that is well known to be unsound if unrestricted. In particular, the value restriction permits only *values* to be polymorphic. However, DisLang₂ has an unusual aspect: functions are *not* values, they are allocated on the heap (§3.1). Hence, the value restriction is not applicable as-is, yet it is crucial to allow universal type quantification in front of functions. We contribute a variant of the value restriction, that allows type quantification in front of any *pure expression that does not call a function, project a pair, case over a sum, or fork new tasks*. This includes function allocation, pair allocation, sums injection, as well as other control-flow constructs. This syntactic check is ensured by the predicate `veryPure e` that appears as a premise of the type abstraction rule **T-TABS**. The predicate `veryPure e` is defined in the Appendix [Moine et al. 2025b]. It can be seen as an alternative to the solution proposed by [de Vilhena \[2022\]](#), in which every arrow has a *purity attribute*, indicating if the function interacts with the store. Contrary to [de Vilhena](#)'s solution, we support some benign interactions with the store: the allocation of immutable data structures.

TypeDis supports higher-kind type polymorphism. For example, reminding of the typing rule **T-PAR**, one could present `par` as a higher-order function of the following type

$$\text{par} : \forall(\varphi_1 :: \mathbf{x} \Rightarrow \star)(\varphi_2 :: \mathbf{x} \Rightarrow \star).$$

$$\forall\delta_1\delta_2. (\forall\delta' \delta \preccurlyeq \delta'. () \rightarrow^{\delta'} \varphi_1 \delta')@\delta_1 \rightarrow (\forall\delta' \delta \preccurlyeq \delta'. () \rightarrow^{\delta'} \varphi_2 \delta')@\delta_2 \rightarrow^{\delta} (\varphi_1 \delta \times \varphi_2 \delta)@\delta$$

Taking $\varphi_1 = \varphi_2 = \lambda\delta. \text{tree}@\delta$ and doing β -reduction matches the type presented in Section 2.3.

4.7 Subtyping

As presented so far, backtiming—that is, substituting the timestamp of a child task by the one of its parent task at the join point—is the only way of changing a timestamp inside a type (§4.5). We propose here another mechanism that we dub *subtyping*. As the name suggests, subtyping is a form of subtyping [Pierce 2002, §15] for timestamps.

At a high-level, subtyping allows for “advancing” a timestamp within a type, as long as this update makes sense. This notion of “advancing” relates to the notion of *precedence*, describing reachability between two timestamps. We write $\Delta \vdash \delta_1 \preccurlyeq \delta_2$ to describe that δ_1 can reach δ_2 in Δ (the Appendix [Moine et al. 2025b]). Equipped with this reachability predicate, we make a first attempt

$$\begin{array}{c}
\begin{array}{c}
\text{S-REFL} \quad \Delta \vdash \rho \subseteq_{\delta} \rho \\
\text{S-REFLAT} \quad \Delta \vdash \sigma \subseteq_{\delta} \sigma
\end{array}
\qquad
\begin{array}{c}
\text{S-TABS} \quad \frac{\Delta \vdash \rho_1 \subseteq_{\delta} \rho_2}{\Delta \vdash \forall \alpha :: \kappa. \rho_1 \subseteq_{\delta} \forall \alpha :: \kappa. \rho_2} \\
\text{S-SUM} \quad \frac{\Delta \vdash \rho l_1 \subseteq_{\delta} \rho l_2 \quad \Delta \vdash \rho r_1 \subseteq_{\delta} \rho r_2}{\Delta \vdash (\rho l_1 + \rho r_1) \subseteq_{\delta} (\rho l_2 + \rho r_2)}
\end{array}
\qquad
\begin{array}{c}
\text{S-PAIR} \quad \frac{\Delta \vdash \rho l_1 \subseteq_{\delta} \rho l_2 \quad \Delta \vdash \rho r_1 \subseteq_{\delta} \rho r_2}{\Delta \vdash (\rho l_1 \times \rho r_1) \subseteq_{\delta} (\rho l_2 \times \rho r_2)} \\
\text{S-AT} \quad \frac{\Delta \vdash \delta_1 \preccurlyeq \delta_2 \quad (\delta_1 \neq \delta_2 \implies \Delta \vdash \delta_2 \preccurlyeq \delta)}{\Delta \vdash \sigma_1 \subseteq_{\delta_2} \sigma_2} \\
\text{S-REC} \quad \frac{\Delta \vdash \delta_1 \preccurlyeq \delta_2 \quad (\delta_1 \neq \delta_2 \implies \Delta \vdash \delta_2 \preccurlyeq \delta) \quad \Delta \vdash \sigma_1 \subseteq_{\delta_2} \sigma_2 \quad \Delta \mid \alpha \mapsto \delta_2 \vdash_{\delta_2} \sigma_2}{\Delta \vdash \mu\alpha. \sigma_1 @ \delta_1 \subseteq_{\delta} \mu\alpha. \sigma_2 @ \delta_2}
\end{array}
\qquad
\begin{array}{c}
\text{S-ABS} \quad \frac{\Delta' = \Delta \cup \Delta_2 \quad \Delta' \vdash \rho \vec{s}_2 \subseteq_{\delta_f} \rho \vec{s}_1 \quad \Delta' \vdash \rho_1 \subseteq_{\delta_f} \rho_2}{\Delta \vdash \forall \vec{s} \Delta_1. \rho \vec{s}_1 \rightarrow^{\delta_f} \rho_1 \subseteq_{\delta} \forall \vec{s} \Delta_2. \rho \vec{s}_2 \rightarrow^{\delta_f} \rho_2} \\
\text{S-INST} \quad \frac{\delta_2 = [\delta_y / \delta_x] \delta_1 \quad \rho \vec{s}_2 = [\delta_y / \delta_x] \rho \vec{s}_1 \quad \Delta_2 = [\delta_y / \delta_x] \Delta_1}{\Delta \vdash \forall (\delta s_l + [\delta_x] + \delta s_r) \Delta_1. \rho \vec{s}_1 \rightarrow^{\delta_1} \rho_1 \subseteq_{\delta} \forall (\delta s_l + \delta s_r) \Delta_2. \rho \vec{s}_2 \rightarrow^{\delta_2} \rho_2}
\end{array}
\end{array}$$

Fig. 13. The subtyping judgment

at capturing the idea of subtyping as follows:

$$\text{SPECIALIZED-SUBTIMING} \quad \frac{\Delta \mid \Gamma \vdash e : \sigma @ \delta_1 \triangleright \delta \quad \Delta \vdash \delta_1 \preccurlyeq \delta_2 \quad \Delta \vdash \delta_2 \preccurlyeq \delta}{\Delta \mid \Gamma \vdash e : \sigma @ \delta_2 \triangleright \delta}$$

SPECIALIZED-SUBTIMING asserts that an expression of type $\sigma @ \delta_1$ can be viewed as an expression of type $\sigma @ \delta_2$ as long as δ_1 precedes δ_2 and δ_2 is not ahead of time, that is δ_2 precedes the current timestamp δ . Indeed, TypeDis enforces that, if $\Delta \mid \Gamma \vdash e : \sigma @ \delta' \triangleright \delta$ holds, then δ' precedes δ .

While **SPECIALIZED-SUBTIMING** is admissible in TypeDis, it is not general enough, as it only considers the timestamp at the root of a type. This motivates rule **T-SUBTIMING** in Figure 12, which relies on the subtyping judgment $\Delta \vdash \rho \subseteq_{\delta} \rho'$, given in Figure 13, and acts as a subsumption rule. Intuitively, the judgment $\Delta \vdash \rho \subseteq_{\delta} \rho'$ captures the fact the timestamps in ρ precede the timestamps in ρ' under logical graph Δ , knowing that every timestamp occurring in ρ' must precede δ . The definition of the judgment now allows changing the timestamps inside immutable types. Because of variance issues (see [Pierce 2002, §15.5]), however, subtyping for mutable types is only shallow: a timestamp can be changed only at the root of an array type.

The subtyping judgment $\Delta \vdash \rho \subseteq_{\delta} \rho'$ assumes that types are in β -normal form. **S-REFL** and **S-REFLAT** assert that the subtyping judgment is reflexive. **S-TABS** asserts that subtyping goes below type quantifiers (which are irrelevant here, the subtyping judgment tolerates open terms). **S-PAIR** and **S-SUM** reflect that subtyping for immutable types is deep.

S-AT illustrates the case presented in **SPECIALIZED-SUBTIMING**. This rule asserts that, with logical graph Δ and maximum allowed timestamp δ , the boxed type $\sigma_1 @ \delta_1$ is a subtype of $\sigma_2 @ \delta_2$ if three conditions are met. First, δ_1 must precede δ_2 . Second, if subtyping is applied here, that is, if $\delta_1 \neq \delta_2$, then δ_2 must precede δ , the maximum timestamp allowed. Third, σ_1 must recursively be a subtype of σ_2 , with maximum timestamp allowed δ_2 . Indeed, recall that TypeDis allows only for up-pointers: every timestamp in σ_2 must precede δ_2 .

S-REC allows subtyping for recursive types $\mu\alpha. \sigma_1 @ \delta_1$ and $\mu\alpha. \sigma_2 @ \delta_2$. The first three premises (in the left-to-right, top-to-bottom order) are the same as for **S-AT**. The fourth premise $\Delta \mid \alpha \mapsto \delta_2 \vdash_{\delta_2} \sigma_2$ requires explanations. This predicate, dubbed the “valid variable” judgment and formally defined in the Appendix [Moine et al. 2025b], ensures two properties. First, that α does not appear in an array type (because subtyping is not allowed at this position) or in an arrow or another recursive type (for simplicity). Second, that if α appears under a timestamp δ , then δ_2 must precede δ .

OOB-ALLOC $i < 0$	OOB-LOAD $\sigma(\ell) = \vec{v}$ $i < 0 \vee i \geq \vec{v} $	OOB-STORE $\sigma(\ell) = \vec{v}$ $i < 0 \vee i \geq \vec{v} $	OOB-CAS $\sigma(\ell) = \vec{v}$ $i < 0 \vee i \geq \vec{v} $
OOB σ (alloc i v)	OOB σ ($\ell.[i]$)	OOB σ ($\ell.[i] \leftarrow w$)	OOB σ (CAS $\ell i w_1 w_2$)
RED-SCHED $S/T/e \xrightarrow{\text{sched}} S'/T'/e'$	RED-OOB $\sigma(\ell) = \vec{v}$	RED-CTX AllRedOrOOB $S/T/e$	RED-CTX AllRedOrOOB $S/T/K[e]$
AllRedOrOOB $S/T/e$	AllRedOrOOB $(\sigma, \alpha, G) / t / e$		
RED-PAR $(e_1 \notin \mathcal{V} \vee e_2 \notin \mathcal{V})$ $(e_1 \notin \mathcal{V} \implies \text{AllRedOrOOB } S/T_1/e_1)$ $(e_2 \notin \mathcal{V} \implies \text{AllRedOrOOB } S/T_2/e_2)$	SAFE-FINAL Safe $S/t/v$	SAFE-NONFINAL AllRedOrOOB $S/T/e$	SAFE-NONFINAL Safe $S/T/e$
AllRedOrOOB $S/T_1 \otimes_t T_2 / e_1 \parallel e_2$			

Fig. 14. The OOB, AllRedOrOOB and Safe predicates

S-ABS allows subtyping for function types $\forall \vec{\delta} \Delta_1$. $\vec{\rho} \vec{s}_1 \rightarrow^{\delta_f} \rho_1$ and $\forall \vec{\delta} \Delta_2$. $\vec{\rho} \vec{s}_2 \rightarrow^{\delta_f} \rho_2$. The quantified timestamps $\vec{\delta}$ s and the calling timestamp δ_f must be the same. The extended logical graph Δ' , equal to $\Delta \cup \Delta_2$, must subsume Δ_1 . Moreover, the arguments $\vec{\rho} \vec{s}_2$ must subtype $\vec{\rho} \vec{s}_1$ (note the polarity inversion). The return type ρ_1 must subtype ρ_2 .

S-INST allows for specializing a universally-quantified timestamp and has a more standard subtyping flavor. In this rule, the quantified timestamp δ_x is being instantiated with δ_y (similarly to the instantiation occurring in **T-APP**).

Before using subtyping, information about precedence may be needed. TypeDis guarantees a strong invariant: every timestamp occurring in the typing environment comes before the current timestamp. Such an invariant is illustrated by **T-GETROOT**, which allows adding to the logical graph Δ an edge (δ', δ) , where δ' is a timestamp in the environment and δ the current timestamp.

5 Soundness

In this section, we state the soundness of TypeDis and give an intuition for its proof, which takes the form of a logical relation in Iris and is mechanized in Rocq [Moine et al. 2025a]. We first enunciate the soundness theorem (§ 5.1). We then recall the concepts of Iris we need (§ 5.2) and present DisLog₂ (§ 5.3), the verification logic we use. We then devote our attention to the formal proof, by presenting the high-level ideas of the logical relation (§ 5.4) we developed and its fundamental theorem (§ 5.5). We conclude by assembling all the building blocks we presented and sketching the soundness proof of TypeDis (§ 5.6).

5.1 Soundness Statement of TypeDis

Our soundness statement adapts Milner [1978]’ slogan “well-typed programs cannot go wrong” by proving that the reduction of a well-typed program reaches only configurations that are *safe* and *disentangled*.

We already formally defined the concept of disentanglement (§ 3.3). What about safety? Intuitively, a configuration is safe if *all tasks can take a step* or, conversely, *no task is stuck*. However, this property is too strong for our type system due to reasons unrelated to disentanglement. Being purposefully designed for disentanglement, our type system is not capable of verifying arbitrary functional correctness conditions. In particular, while the semantics of DisLang₂ ensures that accesses to arrays by load and store operations are within bounds and thus cannot cause a task to get stuck, our type system does not enforce that. This restriction comes at the advantage of freeing

programmers from carrying out correctness proofs themselves, which are carried out by the type-checker instead. Intuitively, we say that a configuration is safe if it is final, or each task can either take a step or encounters a load or a store operation out-of-bounds. We formalize these properties in Figure 14. The property $\text{OOB } \sigma \ e$ asserts that the expression e faces an out-of-bounds operation: that is, an allocation, a load, a store, or a CAS outside the bounds. The property $\text{AllRedOrOOB } S / T / e$ asserts that, within the configuration of the program state S , the task tree T and the expression e , every task of the task tree can either take a step or faces an out-of-bounds operation. RED-SCHED asserts that the configuration can take a scheduling step (that is, either a head step, a fork, or a join). RED-OOB asserts that the configuration is at a leaf and faces an out-of-bounds operation. RED-CTX asserts that an expression under evaluation is reducible if this very expression is reducible. RED-PAR asserts that an active parallel pair $e_1 \parallel e_2$ is reducible if at least one of its components e_i is not a value and any e_i that is not a value is reducible. (If both expressions are values, a join is possible). The property $\text{Safe } S / T / e$ asserts that the configuration $S / T / e$ is either final (SAFE-FINAL), that is, the task tree is at a leaf and the expression is a value, or that every task of the task tree can either take a step or faces an out-of-bounds operation (SAFE-NONFINAL).

An expression e is *always safe and disentangled* if $(\emptyset, \emptyset, \{(t_0, t_0)\}) / t / e \xrightarrow{\text{step}}^* S' / T' / e'$ implies that $\text{Safe } S' / T' / e'$ and $\text{Disentangled } S' / T' / e'$ hold, for some initial timestamp t_0 .

THEOREM 5.1 (SOUNDNESS OF TYPEDIS). *If $\emptyset \mid \emptyset \vdash e : \rho \triangleright \delta$ then e is always safe and disentangled.*

PROOF. We prove this theorem using a logical relation [Timany et al. 2024], which makes use of DisLog_2 , a variation of DisLog [Moine et al. 2024]. We present the proof sketch in Section 5.6. \square

5.2 Iris Primer

We set up our proofs in Iris [Jung et al. 2018b], and recall here the base notations. Iris’ assertions are of type $iProp$. We write Φ for an assertion, $\lceil P \rceil$ for an assertion of the meta-logic (that is, Rocq), $\Phi_1 * \Phi_2$ for a separating conjunction, and $\Phi_1 -* \Phi_2$ for a separating implication. We write a postcondition—that is, a predicate over values—using Ψ .

One of the most important features of Iris are *invariants*. An invariant assertion Φ , written $\boxed{\Phi}$, holds true in-between every computation step. (Formally, invariants are annotated with so-called masks [Jung et al. 2018b, §2.2], we omit them for brevity.) Invariants, as well as other logical resources in Iris, are implemented using *ghost state*. We write $\Phi_1 \Rightarrow \Phi_2$ to denote a *ghost update*—that is, an update of the ghost state between Φ_1 and Φ_2 .

Iris features a variety of modalities. In this work we use two of them extensively. First, the *persistence* modality, written $\Box \Phi$, asserts that the assertion Φ is persistent, meaning in particular that $\Box \Phi$ is duplicable. Second, the *later* modality, written $\triangleright \Phi$, asserts that Φ holds “one step of computation later”.

We write $\ell \mapsto \vec{v}$ to denote that ℓ points-to an array with contents \vec{v} . We write $\ell \mapsto_{\Box} r$, with a *discarded fraction* [Vindum and Birkedal 2021], to denote that ℓ points-to an immutable block r (that is, either a closure, an immutable pair, or an immutable sum). This latter assertion is persistent.

5.3 Taking Advantage of the Cyclic Approach with DisLog_2

Moine et al. [2024] contributed DisLog , the first program logic for verifying disentanglement. DisLog depends on the very definition of disentanglement, and uses the standard approach presented in Section 2.1: when two tasks join, they form a new task with a fresh timestamp. This choice impacts the logic: the *weakest precondition* (WP) modality of DisLog takes the form $\text{wp} \langle t, e \rangle \{ \lambda t' v. \Phi \}$ and asserts that the expression e running on *current timestamp* t is disentangled, and if the evaluation of e terminates, it does so on the *end timestamp* t' , with final value v and satisfying the assertion Φ . In particular, t and t' may not be the same, for example if e contains a call to `par`.

$\begin{array}{c} \text{D-LOAD} \\ \frac{\lceil 0 \leq i < \vec{w} \wedge \vec{w}(i) = v \rceil \quad \ell \mapsto_p \vec{w} \quad v \odot t}{\text{wp} \langle t, \ell.[i] \rangle \{ \lambda v'. \lceil v' = v \rceil * \ell \mapsto_p \vec{w} \}} \end{array}$	$\begin{array}{c} \text{D-LOADOOB} \\ \frac{\lceil i < 0 \vee i \geq \vec{w} \rceil \quad \ell \mapsto_p \vec{w}}{\text{wp} \langle t, \ell.[i] \rangle \{ \lambda _. \perp \}} \end{array}$
$\begin{array}{c} \text{D-PAR} \\ \frac{\forall t_1 t_2. \quad t \preceq t_1 * t \preceq t_2 \Rightarrow \exists \Psi_1 \Psi_2. \quad \text{wp} \langle t_1, \ell_1 [() \rangle \{ \Psi_1 \} * \text{wp} \langle t_2, \ell_2 [() \rangle \{ \Psi_2 \} * \\ (\forall v_1 v_2 \ell. \quad \Psi_1 v_1 * \Psi_2 v_2 * t_1 \preceq t * t_2 \preceq t * \ell \mapsto (v_1, v_2) * \Psi \ell)}{\text{wp} \langle t, \text{par}(\ell_1, \ell_2) \rangle \{ \Psi \}} \end{array}$	$\begin{array}{c} \text{D-CLOCKMONO} \\ \frac{v \odot t_1 \quad t_1 \preceq t_2}{v \odot t_2} \end{array}$

Fig. 15. Selected rules of DisLog₂

To take advantage of the cyclic approach for disentanglement (§2.1), we had to develop a new version of DisLog, yielding the logic DisLog₂. DisLog₂ allows reusing the timestamp of the forking task for the child tasks upon join. As a result, the current timestamp and end timestamp of an expression always coincide, allowing us to simplify the WP of DisLog by simply removing the end timestamp parameter of the postcondition. Formally, the WP of DisLog₂ then takes the form

$$\text{wp} \langle t, e \rangle \{ \lambda v. \Phi \}$$

and asserts that the expression e running on timestamp t is disentangled, and if the evaluation of e terminates, it does so with final value v and satisfying the assertion Φ . In contrast to DisLog, DisLog₂ tolerate out-of-bounds accesses to cater to the TypeDis type system which only enforces disentanglement. (In practice, DisLog₂ is parameterized by a boolean flag which can be used to enable or disable inboundedness proof obligations; when such obligations are enabled, DisLog₂ has the same expressive power as DisLog.) Otherwise, DisLog₂ adapts all the ideas of DisLog. In particular, the logic features two persistent assertions related to timestamps. First, the clock assertion $\ell \odot t$ asserts that location ℓ was allocated by a task that precedes t . Similarly, $v \odot t$ has the same meaning, if v is a location ℓ , or otherwise denotes $\lceil \text{True} \rceil$. Second, the precedence assertion $t_1 \preceq t_2$ asserts that task t_1 precedes task t_2 in the underlying computation graph. The precedence assertion forms a pre-order: it is reflexive and transitive. Crucially, the clock assertion is *monotonic with respect to the precedence pre-order* [Moine et al. 2024].

In the remainder of the paper, we write $t_1 \approx t_2$ to denote that t_1 and t_2 are *equivalent*, that is, both $t_1 \preceq t_2$ and $t_2 \preceq t_1$ hold.

Selected rules of DisLog₂. Figure 15 presents four key rules of DisLog₂. The premise of these rules are implicitly separated by a separating conjunction $*$.

D-LOAD, targeting a load operation on the array ℓ at offset i on task t , ensures disentanglement. Indeed, the rule requires that ℓ points-to the array \vec{w} and that the offset i in \vec{w} corresponds to the value v . It also requires the assertion $v \odot t$, witnessing that if v is a location, then this location must have been allocated by a preceding task. **D-LOADOOB** is unusual for a program logic and reflects that we purposefully allow for OOB accesses in verified programs, because our type system does. Because an OOB access results in a crash, the postcondition of the WP is $\lceil \text{False} \rceil$, allowing the user to conclude anything. **D-PAR** is at the heart of DisLog₂ and allows verifying a parallel call to two closures ℓ_1 and ℓ_2 at timestamp t . The premise universally quantifies over t_1 and t_2 , the two timestamps of the forked tasks, that are both preceded by t . Then, the user must provide two postconditions, Ψ_1 and Ψ_2 for the two tasks, and verify that the closure call $\ell_1 [()$ (resp. $\ell_2 [()$) is safe at timestamp t_1 (resp. t_2) with postcondition Ψ_1 (resp. Ψ_2). The second line of the premise requires the user to prove that, after the two tasks terminated and joined, the initial postcondition $\Psi \ell$ must hold, for some location ℓ pointing to the pair (v_1, v_2) where v_1 is the final result of t_1 and v_2 of t_2 . **D-CLOCKMONO** formalizes monotonicity of the clock assertion.

The adequacy theorem of DisLog₂. The adequacy theorem of DisLog₂ asserts that if e can be verified using the program logic, then e is always safe and disentangled.

THEOREM 5.2 (ADEQUACY OF DISLOG₂). *If $wp\langle t, e \rangle \{\Psi\}$ holds then e is always safe and disentangled.*

PROOF. Similar to the adequacy proof of DisLog; see our mechanization [Moine et al. 2025a]. \square

5.4 A Logical Relation

The very heart of the soundness proof of TypeDis is a *logical relation*, set up in Iris using DisLog₂. *Logical relations* [Girard 1972; Pitts and Stark 1998; Plotkin 1973; Statman 1985; Tait 1967] are a technique that allows one to prescribe properties of valid programs in terms of their *behavior*, as opposed to solely their static properties. We adopt the *semantic approach* [Constable et al. 1986; Martin-Löf 1982; Timany et al. 2024], which admits terms that are not necessarily (syntactically) well-typed to be an inhabitant of the logical relation and has been successfully deployed in the RustBelt project [Jung et al. 2018a], for example. Our logical relation is presented in the Appendix [Moine et al. 2025b]; we comment next on the high-level ideas.

As usual, our (unary) logical relation gives the *interpretation* of a type ρ with kind \star as a predicate on values. Values satisfying the predicate are said to inhabit the relation. Because our types have higher kinds, our logical relation includes predicates on timestamps. In particular, the interpretation of a type ρ with kind κ is a function taking κ timestamp arguments (where \star indicates zero timestamps and $x \Rightarrow \kappa$ indicates $\kappa + 1$ timestamps) and producing a predicate over values.

The presented relations involve two sorts of closing substitutions for variables occurring in types. First, a *timestamp substitution*, written h , which is a finite map from timestamp variables δ to concrete timestamps t . Second, a *type substitution*, written m , which is a finite map from type variables to tuples of a kind κ and a tuple of two functions depending on κ . The first function takes κ timestamps and produces a predicate over values; it represents the semantic interpretation of the type by which the variable will be instantiated. The second function takes κ timestamps and produces a timestamp; its result corresponds to the root timestamp of the type by which the variable will be instantiated. The interpretation of a type guarantees that the type only contains *up-pointers* (§2.2), that is, the interpretation of $\sigma @ \delta$ ensures that if δ' appears in σ , then δ' precedes δ . To enforce this invariant, our approach makes use of transitivity: the interpretation of $\sigma @ \delta$ ensures that, for each outermost ρ encountered in σ , the root timestamp of ρ —conceptually, the outermost timestamp in ρ —precedes δ . Because this invariant is enforced at each stage of the type interpretation, and because precedence is transitive, we guarantee that there are only up-pointers. The Appendix [Moine et al. 2025b] presents a function that computes the root timestamp of a type and defines the assertion $\text{root } \rho \preceq_m^h \delta$, asserting that the root timestamp of ρ comes before δ with the mappings h and m .

The main relation is the *type relation* $[\rho]_m^h \kappa$. It produces a predicate waiting for κ timestamps, a value v , and captures that v is of type ρ , given the timestamp mapping h and type mapping m .

Apart from timestamps, the seasoned reader of logical relations in Iris will not be surprised by our approach, as it follows the standard recipe [Timany et al. 2024]: a recursive type is interpreted using a guarded fixpoint, universal type quantification is interpreted as a universal quantification in the logic, an array is interpreted using an invariant, and an arrow using WP. Moreover, every predicate is designed such that it is persistent.

5.5 Interpretation of Typing Judgments

We now focus on the interpretation of the TypeDis typing judgment, paving our way to state the fundamental theorem of the logical relation. Figure 16 gives its interpretation, appealing to the WP of

$$\begin{aligned}
\llbracket \Delta \mid \Gamma \vdash e : \rho \triangleright \delta \rrbracket &\triangleq \square \forall h m u. \Gamma \text{ dom } \Gamma = \text{dom } u \neg * \\
&\quad \Gamma \forall \alpha \kappa \Psi r. m(\alpha) = (\kappa, (\Psi, r)) \implies \text{proper}_\kappa \Psi \wedge \text{regular}_\kappa r \neg * \\
&\quad \star_{(t_1, t_2) \in \Delta} h(t_1) \preceq h(t_2) \neg * \\
&\quad \star_{(x, \rho) \in \Gamma, (x, v) \in u} (\text{root } \rho \preceq_m^h \delta * \llbracket \rho \rrbracket_m^h \star v) \neg * \\
&\quad \forall t. t \approx h(\delta) \neg * \text{ wp } \langle t, [u/]e \rangle \{ \lambda v. \text{root } \rho \preceq_m^h \delta * \llbracket \rho \rrbracket_m^h \star v \}
\end{aligned}$$

Fig. 16. The interpretation of typing judgments

DisLog₂. For the judgment with logical graph Δ , type environment Γ , and expression e with type ρ at timestamp δ , the interpretation starts by quantifying over three closing substitutions: the timestamp mapping h , the type mapping m , as well as a *variable mapping* u , a map from term variables to values. The variable mapping must have the same domain as the environment Γ . The type mapping m is restricted such that type variables are given only a proper interpretation (via the $\text{proper}_\kappa \Psi$ property) and a regular root function (via the $\text{regular}_\kappa r$ property). The property $\text{proper}_\kappa \Psi$ captures that any timestamp parameter of Ψ can be replaced by an equivalent one. The property $\text{regular}_\kappa r$ captures that the function r either ignores all its arguments or returns one of them. These two properties are needed in order to prove the correctness of **T-PAR**.

Then, the interpretation requires that Δ is a valid logical graph, that is, each edge between two timestamp variables in Δ corresponds to an edge between their mapping. The interpretation also requires that, for every variable x that has type ρ in Γ and is associated to value v in u , the root timestamp of v precedes the interpretation of δ and v inhabits the interpretation of ρ . Next, the definition quantifies over a timestamp t , equivalent to the interpretation of δ , and asserts WP at timestamp t of the expression e in which variables are substituted by values following the variable mapping u . The postcondition asserts that the root timestamp of v precedes the interpretation of δ and that the returned value v inhabits the interpretation of type ρ .

Having the interpretation of typing judgment defined, we can state the *fundamental theorem*, ensuring that syntactically well-typed terms (§4) inhabit the logical relation.

THEOREM 5.3 (FUNDAMENTAL). *If $\Delta \mid \Gamma \vdash e : \rho \triangleright \delta$ holds, then $\llbracket \Delta \mid \Gamma \vdash e : \rho \triangleright \delta \rrbracket$ holds too.*

PROOF. By induction over the typing derivation; see our mechanization [Moine et al. 2025a]. \square

5.6 Putting Pieces Together: The Soundness Proof of TypeDis

We can finally unveil the proof of the soundness **Theorem 5.1** of TypeDis, which we formally establish in Rocq. Let us suppose that $\emptyset \mid \emptyset \vdash e : \rho \triangleright \delta$ holds. Making use of the fundamental **Theorem 5.3**, we deduce that $\llbracket \emptyset \mid \emptyset \vdash e : \rho \triangleright \delta \rrbracket$ holds too. Unfolding the definition (Figure 16), instantiating the timestamp mapping h with the singleton map $[\delta := t_0]$ —for some initial timestamp t_0 —and the type mapping m and the variable mapping u with empty maps, and simplifying trivial premises concerning these mappings, we are left with the statement $\forall t. t \approx t_0 \neg * \text{ wp } \langle t, e \rangle \{ \lambda v. \text{root } \rho \preceq_m^h \delta * \llbracket \rho \rrbracket_{\emptyset}^{[\delta := t_0]} \star v \}$. Instantiating t with t_0 , we deduce that $\text{wp } \langle t_0, e \rangle \{ \lambda v. \text{root } \rho \preceq_m^h \delta * \llbracket \rho \rrbracket_{\emptyset}^{[\delta := t_0]} \star v \}$ holds. We finally use the adequacy **Theorem 5.2** of DisLog₂ and deduce that e is always safe and disentangled.

6 Case Studies

We evaluate the usefulness of TypeDis by type-checking several case studies in Rocq using the rules presented in Section 4.

We verify the examples presented in the “Key Ideas” Section 2. These examples illustrate: simple mechanics of the type system (§2.2), backtiming (§2.3), and subtiming (§2.4). In particular, the last two examples, `build` and `selectmap`, illustrate the use of TypeDis with higher-order functions and

```

parfor  $\triangleq$   $\mu f. \lambda[a; b; k].$ 
  if  $b \leq a$  then () else if  $(b - a) == 1$  then  $k [a]$ 
  else let  $mid = a + ((b - a)/2)$  in  $\text{par}(f [a; mid; k], f [mid; b; k])$ 
parfor :  $\forall \delta \delta_k. \text{int} \rightarrow \text{int} \rightarrow (\forall \delta' \delta \preceq \delta'. \text{int} \rightarrow^{\delta'} ()) @ \delta_f \rightarrow^{\delta} ()$ 

```

Fig. 17. Implementation and type of the parfor primitive

a recursive immutable type (a binary tree with integer leaves). For more details than the intuitions we already gave for these examples, we refer the reader to our formalization [Moine et al. 2025a].

Our largest case study consists of the typing of a parallel deduplication algorithm via concurrent hashing. This example is a case study of DisLog [Moine et al. 2024, §6.3]. Deduplication consists of removing duplicates from an array—something that can be done efficiently in a parallel, disentangled setting [Westrick 2022]. The algorithm relies on a folklore [VerifyThis 2022] concurrent, lock-free, fixed-capacity hash set using *open addressing* and *linear probing* to handle collisions [Knuth 1998]. The main deduplication function allocates a new hash set, inserts in parallel every element into the hash set using a parallel for loop, and finally returns the elements of the set. We first comment on the parallel for loop (§6.1) and then on the main deduplication algorithm (§6.2).

6.1 The Parallel For Loop

Our implementation of the parallel for loop appears in the upper part of Figure 17 and is a direct translation of MaPLe standard library’s implementation [Acar et al. 2020]. The function parfor takes three arguments: a lower bound a , a higher bound b , and a closure k to execute at each index between these bounds. The function parfor is defined recursively: it returns immediately if $b \leq a$, executes the closure $k [a]$ if $b - a = 1$, and otherwise calls itself recursively *in parallel*, splitting the range in two. The type we give to parfor appears in the lower part of Figure 17, and is as one could expect. Indeed, the type quantifies over two timestamps δ , at which parfor will be called, and δ_k , the (irrelevant) timestamp of the closure. The function then requires two integers, and a closure that will be called at some timestamp δ' that succeeds δ .

The type-checking of parfor is non-trivial because it involves *polymorphic recursion*. Indeed, parfor’s type universally quantifies over the calling timestamp, but calls itself recursively after a *par*—that is, at another (subsequent) timestamp. TypeDis supports natively such a pattern thanks to **T-ABS**. Interestingly, polymorphic recursion introduce a need for subtyping. Indeed, while type-checking parfor’s body at current timestamp δ , the closure k has type $(\forall \delta' \delta \preceq \delta'. \text{int} \rightarrow^{\delta'} ())$. However, the recursive call happens after a *par*, hence at a new current timestamp δ_1 such that $\delta \preceq \delta_1$. But in order to type-check the recursive call, the user has to give to k the type $(\forall \delta' \delta_1 \preceq \delta'. \text{int} \rightarrow^{\delta'} ())$ —notice the difference between the precedence information on δ' . This is a typical use of subtyping, and because $\delta \preceq \delta_1$, we conclude using **S-ABS**.

6.2 Internals of the Deduplication Case Study

Let us now focus on the code for our deduplication algorithm, which appears in the upper part of Figure 18. This code assumes a maximum size C for the underlying hash set. The function dedup takes three arguments: a hashing function h , a dummy element d in order to populate the result array, and the array to deduplicate ℓ . The function first allocates the hash set a and then calls in parallel the add function for every index in ℓ . The function add consists of a CAS loop, that tries to insert the element in the first available slot. Finally, dedup filters the remaining dummy elements d using an omitted function filter_compact. Because it involves no fork or join, the function add admits a simple polymorphic type, shown in the lower part of Figure 18, quantifying over the

```

add  ≡  λ[h; a; d; x].
let put = μf.λ[i].
  if (CAS a i d x ∨ a.[i]==x) then ()
  else f [(i + 1) mod C] in
  put [h [x] mod C]

dedup  ≡  λ[h; d; ℓ].
let a = alloc C d in
let k = λ[i]. add [h; a; d; ℓ.[i]] in
parfor [0; length ℓ; k] ;
filter_compact [a; d]

add : ∀α :: ★. ∀δ δ₁ δ₂. (α →δ int)@δ₁ → array(α)@δ₂ → α → α →δ ()
dedup : ∀α :: ★. ∀δ δ₁ δ₂. (forall δ' δ < δ'. α →δ' int)@δ₁ → α → array(α)@δ₂ →δ array(α)@δ

```

Fig. 18. Case study: deduplication of an array by concurrent hashing

type α of the elements of the array to deduplicate, the calling timestamp δ and two timestamps δ_1 and δ_2 . The first argument is a closure of a hashing function on α that will be called at timestamp δ . The second argument is the hash set, a $\text{array}(\alpha)$. The third and fourth arguments are of type α and correspond to the dummy element and the element to insert, respectively.

Using this type for `add`, we are able to type-check `dedup` with the type shown in the lower part of Figure 18. This type quantifies again over the type α of the elements of the array to deduplicate, and then quantifies over δ , the calling timestamp, and δ_1 and δ_2 , the (irrelevant) allocation timestamps of the first and third argument, respectively. The first argument is a closure of a hashing function on α , that will be called at subsequent tasks δ' . The second argument is a dummy element. The third argument is the array to deduplicate. Again, type-checking `dedup` requires subtyping: the `add` function expects a hashing function at its calling timestamp δ , whereas the supplied h is more general, because it is polymorphic with respect to its calling timestamp. We use subtyping (**S-ABS** and **S-INST**) to convert the latter into the former.

7 Related Work

Disentanglement. The specific property we consider in this paper is based on the definition by Westrick et al. [2020] which was later formalized by Moine et al. [2024]. Most of the existing work on disentanglement considers structured fork-join parallel code, as we do in this paper. More recently, Arora et al. [2024] showed that disentanglement is applicable in a more general setting involving parallel futures, and specifically prove deadlock-freedom in this setting. We plan to investigate whether TypeDis could be extended to support futures.

Verification of Disentanglement. Two approaches to check for and/or verify disentanglement have been proposed prior to TypeDis. First, as currently implemented in the MaPLE compiler, the programmer can rely on a runtime *entanglement detector* [Westrick et al. 2022]. This approach is similar in principle to dynamic race detection [Flanagan and Freund 2009]. In the case of entanglement, dynamic detection has been shown to have low overhead, making it suitable for automatic run-time management of entanglement [Arora et al. 2023]. However, run-time detection cannot guarantee disentanglement due to the inherent non-determinism of entanglement, which typically arises due to race conditions and may or may not occur in individual executions. The second approach, as developed by Moine et al. [2024], is full-blown static verification of disentanglement using a separation logic called DisLog, proven sound in Rocq. This approach can be used to statically verify disentanglement for a wide variety of programs—for example, even for non-deterministic programs that utilize intricate lock-free data structures in shared memory. However, static verification with DisLog is difficult, requiring significant effort even to verify small examples.

Region-based Systems. TypeDis associates timestamp variables with values in their types. Immediately, we note similarities with region-based type and effect systems [Grossman et al. 2002; Tofte et al. 2004; Tofte and Talpin 1997] which have also recently received attention in supporting parallelism [Elsman and Henriksen 2023]. The timestamps in our setting are somewhat analogous to regions, with parent-child relationships between timestamps and the up-pointer invariant of TypeDis bearing resemblance to the stack discipline of region-based memory management systems. However, there are a number of key differences. In region-based systems, allocations may occur within any region, and all values within a region are all deallocated at the same moment; one challenge in such systems is statically predicting or conservatively bounding the lifetime of every value. In contrast, in TypeDis, allocations only ever occur at the “current” timestamp, and timestamps tell you nothing about deallocation—every value in our approach is dynamically garbage collected. Each timestamp in TypeDis is associated with a task within a nested fork-join task structure, and values with the same timestamp are all allocated by the same task (or one of its subtasks).

Possible Worlds Type Systems. Our type system falls into what can broadly be categorized as a *possible worlds* type system. These type systems augment the typing judgment with world modalities (in our case timestamp variables δ) that occur as syntactic objects in propositions (a.k.a. types), and typing is then carried out relative to an accessibility relation (in our case the logical graph Δ). While our work is the first to contribute a possible worlds type system for disentanglement, world modalities have been successfully used for other purposes. In the context of fork-join parallelism, Muller et al. [2017] employed world modalities to track priorities of tasks and guarantee absence of priority inversions, ensuring responsiveness and interactivity. While Muller et al. [2017] also require their priorities to be partially ordered, as we require timestamps to be partially ordered, their priorities are fixed, whereas ours are not. In the context of message-passing concurrency, world modalities have been employed to verify deadlock-freedom [Balzer et al. 2019], domain accessibility [Caires et al. 2019], and information flow control [Derakhshan et al. 2021, 2024]. This line of work not only differs in underlying computation model, considering a process calculus, but also adopts linear typing to control data races and non-determinism. While disentanglement does not forbid races, adopting some form of linear typing may be an interesting avenue for future work, to admit even more disentangled programs as well-typed, e.g. those with down-pointers.

Information Flow Control Type Systems. Information flow type systems [Sabelfeld and Myers 2003; Smith and Volpano 1998; Volpano et al. 1996] can also be viewed as representatives of possible worlds type systems, where modalities capture confidentiality (or integrity) and pc labels, and the accessibility relation is a lattice. Typically, modalities can change by typing. For example, when type-checking the branches of an if statement the pc label is increased to the join of its current value and the confidentiality label of the branching condition. A similar phenomenon happens in TypeDis upon type-checking a fork, where the sibling threads are type-checked at a later timestamp. Besides these similarities in techniques employed, the fundamental invariants preserved by type-checking are different. In our setting it is the “no cross-pointers invariant”, whereas it is noninterference for IFC type systems. As a result, the metatheory employed also differs: whereas we use a unary logical relation, noninterference demands a binary logical relation. Such a binary logical relation for termination-insensitive noninterference in the context of a sequential, higher-order language with higher-order store, for example, has been contributed by Gregersen et al. [2021]. The authors develop an IFC type system, in the spirit of Flow Caml [Pottier and Simonet 2003; Simonet 2003], with label polymorphism, akin to our timestamp polymorphism. Like our work, the authors use the semantic typing approach supported by the Iris separation logic framework. Similarly, the authors support subtyping on labels, allowing a label to be raised in accordance with the lattice, akin to our subtyping, in accordance with the precedence relation.

Type Systems for Parallelism and Concurrency. There has been significant work on developing static techniques, especially type systems, to guarantee correctness and safety properties (such as race-freedom, deadlock-freedom, determinism, etc.) for parallel and concurrent programs. For example, the idea of ownership [Clarke et al. 1998; Dietl and Müller 2005; Müller 2002; Noble et al. 1998] has been exploited to rule out races and deadlocks among threads [Boyapati et al. 2002, 2003; Boyapati and Rinard 2001]. Ownership is also enforced by linear type systems [Wadler 1990], which rule out races by construction and have been successfully employed in message-passing concurrency [Caires et al. 2019; Wadler 2012]. The approach has then been popularized by Rust [Klabnik and Nichols 2023], in particular, focusing on statically restricting aliasing and mutability [Jung et al. 2018a], which in Rust takes the form of ownership and borrowing as well as reference-counted mutexes for maximal flexibility. Recently flexible mode-based systems have been explored, as present in the work on DRFCaml [Georges et al. 2025], which exploits modes (extending Lorenzen et al. [2024]) to distinguish values that can and cannot be safely shared between threads. Other systems leverage region-based techniques to restrict concurrent threads, ensuring safe disjoint access to the heap with minimal annotations [Milano et al. 2022], or leveraging explicit annotations to limit the set of permissible effects on shared parts of the heap [Bocchino Jr. et al. 2009].

Much of these related works focus on the hazards of concurrency: data races, race conditions, non-determinism, and similar issues. Disentanglement (and by extension, TypeDis) focuses on an equally important but different issue, namely, the *performance* of parallel programs. TypeDis in particular is designed to allow for unrestricted sharing of immutable data (as demonstrated in Section 2.4) mixed with disentangled sharing of mutable data (for example, in Section 6). This support for data sharing is motivated by the implementation of efficient parallel algorithms, many of which rely upon access to shared memory with irregular and/or data-dependent access patterns, which are difficult to statically analyze for safety. For example, Abdi et al. [2024] find that many standard implementations of parallel algorithms are rejected by the Rust type system, yet these same implementations have been shown to be disentangled [Westrick et al. 2022]. We consider one such implementation as a case study in Section 6 and confirm that it is typeable under TypeDis.

8 Conclusion and Future Work

Disentanglement is an important property of parallel programs, which can in particular serve for improving performance. This paper introduces TypeDis, a static type system that proves disentanglement. TypeDis annotates types with timestamps, recording for each object the task that allocated it. Moreover, TypeDis supports iso-recursive types, as well as type and timestamp polymorphism. TypeDis allows restamping the timestamps in types using a particular form of subtyping we dub *subtiming*.

This paper focuses on *type-checking*, that is, given a program annotated with types, checking if these types are valid. We are currently working on a prototype type-checker, written in OCaml. An immediate direction for future work is *type inference*, that is, generating a valid type for a program. For future work, we plan to use the framework of Odersky et al. [1999], which adapts Hindley-Milner to a system with constrained universal quantification. We believe subtiming and backtiming will be inferrable. One challenging case will be mixing polymorphic recursion with `par`, which might require annotations in order to remain decidable (this is a known problem in region-based type systems [Tofte and Birkedal 1998]).

Acknowledgments

We thank Umut A. Acar for sharing his insights during early design discussions and helping us shape the context for this work. We also thank Kashish Raimalani for reviewing an initial draft, and we thank the anonymous reviewers for their helpful comments.

References

Javad Abdi, Gilead Posluns, Guozheng Zhang, Boxuan Wang, and Mark C. Jeffrey. 2024. When Is Parallelism Fearless and Zero-Cost with Rust?. In *Proceedings of the 36th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2024, Nantes, France, June 17-21, 2024*, Kunal Agrawal and Erez Petrank (Eds.). ACM, 27–40. [doi:10.1145/3626183.3659966](https://doi.org/10.1145/3626183.3659966)

Umut A. Acar, Jatin Arora, Matthew Fluet, Ram Raghunathan, Sam Westrick, and Rohan Yadav. 2020. MPL: A high-performance compiler for Parallel ML. <https://github.com/MLLang/mpf>

Umut A. Acar, Guy E. Blelloch, Matthew Fluet, Stefan K. Muller, and Ram Raghunathan. 2015. Coupling Memory and Computation for Locality Management. In *1st Summit on Advances in Programming Languages, SNAPL 2015, May 3-6, 2015, Asilomar, California, USA (LIPIcs, Vol. 32)*, Thomas Ball, Rastislav Bodík, Shriram Krishnamurthi, Benjamin S. Lerner, and Greg Morrisett (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 1–14. [doi:10.4230/LIPICS.SNAPL.2015.1](https://doi.org/10.4230/LIPICS.SNAPL.2015.1)

Umut A. Acar, Arthur Charguéraud, Mike Rainey, and Filip Sieczkowski. 2016. Dag-calculus: a calculus for parallel computation. In *International Conference on Functional Programming (ICFP)*. 18–32. <https://doi.org/10.1145/2951913.2951946>

Daniel Anderson, Guy E. Blelloch, Laxman Dhulipala, Magdalen Dobson, and Yihan Sun. 2022. The problem-based benchmark suite (PBBS), V2. In *PPoPP '22: 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, Seoul, Republic of Korea, April 2 - 6, 2022*, Jaejin Lee, Kunal Agrawal, and Michael F. Spear (Eds.). ACM, 445–447. [doi:10.1145/3503221.3508422](https://doi.org/10.1145/3503221.3508422)

Andrew W. Appel. 1992. *Compiling with Continuations*. Cambridge University Press. <http://www.cambridge.org/9780521033144>

Jatin Arora, Stefan K. Muller, and Umut A. Acar. 2024. Disentanglement with Futures, State, and Interaction. *Proc. ACM Program. Lang.* 8, POPL (2024), 1569–1599. [doi:10.1145/3632895](https://doi.org/10.1145/3632895)

Jatin Arora, Sam Westrick, and Umut A. Acar. 2021. Provably space-efficient parallel functional programming. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–33. [doi:10.1145/3434299](https://doi.org/10.1145/3434299)

Jatin Arora, Sam Westrick, and Umut A. Acar. 2023. Efficient Parallel Functional Programming with Effects. *Proc. ACM Program. Lang.* 7, PLDI (2023), 1558–1583. [doi:10.1145/3591284](https://doi.org/10.1145/3591284)

Stephanie Balzer, Bernardo Toninho, and Frank Pfenning. 2019. Manifest Deadlock-Freedom for Shared Session Types. In *28th European Symposium on Programming (ESOP) (Lecture Notes in Computer Science, Vol. 11423)*. Springer, 611–639. [doi:10.1007/978-3-030-17184-1_22](https://doi.org/10.1007/978-3-030-17184-1_22)

Henk P. Barendregt. 1984. *The Lambda Calculus, Its Syntax and Semantics*. Elsevier. http://www.elsevier.com/wps/find/bookdescription.cws_home/501727/description

Robert L. Bocchino Jr., Vikram S. Adve, Danny Dig, Sarita V. Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patrick Simmons, Hyojin Sung, and Mohsen Vakilian. 2009. A type and effect system for deterministic parallel Java. In *Proceedings of the 24th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2009, October 25-29, 2009, Orlando, Florida, USA*, Shail Arora and Gary T. Leavens (Eds.). ACM, 97–116. [doi:10.1145/1640089.1640097](https://doi.org/10.1145/1640089.1640097)

Chandrasekhar Boyapati, Robert Lee, and Martin C. Rinard. 2002. Ownership types for safe programming: preventing data races and deadlocks. In *ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*. ACM, 211–230. [doi:10.1145/582419.582440](https://doi.org/10.1145/582419.582440)

Chandrasekhar Boyapati, Barbara Liskov, and Liuba Shrira. 2003. Ownership types for object encapsulation. In *30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. ACM, 213–223. [doi:10.1145/604131.604156](https://doi.org/10.1145/604131.604156)

Chandrasekhar Boyapati and Martin C. Rinard. 2001. A Parameterized Type System for Race-Free Java Programs. In *ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*. ACM, 56–69. [doi:10.1145/504282.504287](https://doi.org/10.1145/504282.504287)

Luís Caires, Jorge A. Pérez, Frank Pfenning, and Bernardo Toninho. 2019. Domain-Aware Session Types. In *30th International Conference on Concurrency Theory (CONCUR) (LIPIcs, Vol. 140)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 39:1–39:17. [doi:10.4230/LIPICS.CONCUR.2019.39](https://doi.org/10.4230/LIPICS.CONCUR.2019.39)

David G. Clarke, John Potter, and James Noble. 1998. Ownership Types for Flexible Alias Protection. In *ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*. ACM, 48–64. [doi:10.1145/286936.286947](https://doi.org/10.1145/286936.286947)

Robert L. Constable, Stuart F. Allen, Mark Bromley, Rance Cleaveland, J. F. Cremer, Robert Harper, Douglas J. Howe, Todd B. Knoblock, Nax Paul Mendler, Prakash Panangaden, James T. Sasaki, and Scott F. Smith. 1986. *Implementing Mathematics with the Nuprl Proof Development System*. Prentice Hall. <http://dl.acm.org/citation.cfm?id=10510>

Paulo Emilio de Vilhena. 2022. *Proof of Programs with Effect Handlers*. Theses. Université Paris Cité. <https://inria.hal.science/tel-03891381>

Farzaneh Derakhshan, Stephanie Balzer, and Limin Jia. 2021. Session Logical Relations for Noninterference. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society, 1–14. [doi:10.1109/LICS52264.2021.9470654](https://doi.org/10.1109/LICS52264.2021.9470654)

Farzaneh Derakhshan, Stephanie Balzer, and Yue Yao. 2024. Regrading Policies for Flexible Information Flow Control in Session-Typed Concurrency. In *38th European Conference on Object-Oriented Programming (ECOOP) (LIPIcs, Vol. 313)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:29. [doi:10.4230/LIPIcs.ECOOP.2024.11](https://doi.org/10.4230/LIPIcs.ECOOP.2024.11)

Werner Dietl and Peter Müller. 2005. Universes: Lightweight Ownership for JML. *Journal of Object Technology* 4, 8 (2005), 5–32. [doi:10.5381/JOT.2005.4.8.A1](https://doi.org/10.5381/JOT.2005.4.8.A1)

Martin Elsman and Troels Henriksen. 2023. Parallelism in a Region Inference Context. *Proc. ACM Program. Lang.* 7, PLDI (2023), 884–906. [doi:10.1145/3591256](https://doi.org/10.1145/3591256)

Cormac Flanagan and Stephen N. Freund. 2009. FastTrack: efficient and precise dynamic race detection. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, 121–133. [doi:10.1145/1542476.1542490](https://doi.org/10.1145/1542476.1542490)

Aïna 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. *Proc. ACM Program. Lang.* 9, POPL (2025), 656–686. [doi:10.1145/3704859](https://doi.org/10.1145/3704859)

Jean-Yves Girard. 1972. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse d'État. Université Paris 7. <https://girard.perso.math.cnrs.fr/These.pdf>

Simon Oddershede Gregersen, Johan Bay, Amin Timany, and Lars Birkedal. 2021. Mechanized logical relations for termination-insensitive noninterference. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–29. [doi:10.1145/3434291](https://doi.org/10.1145/3434291)

Dan Grossman, J. Gregory Morrisett, Trevor Jim, Michael W. Hicks, Yanling Wang, and James Cheney. 2002. Region-Based Memory Management in Cyclone. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany, June 17–19, 2002*, Jens Knoop and Laurie J. Hendren (Eds.). ACM, 282–293. [doi:10.1145/512529.512563](https://doi.org/10.1145/512529.512563)

Adrien Guatto, Sam Westrick, Ram Raghunathan, Umut A. Acar, and Matthew Fluet. 2018. Hierarchical memory management for mutable state. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPoPP 2018, Vienna, Austria, February 24–28, 2018*, Andreas Krall and Thomas R. Gross (Eds.). ACM, 81–93. [doi:10.1145/3178487.3178494](https://doi.org/10.1145/3178487.3178494)

Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. 2018a. RustBelt: Securing the Foundations of the Rust Programming Language. *Proceedings of the ACM on Programming Languages* 2, POPL (2018), 66:1–66:34. <https://people.mpi-sws.org/~dreyer/papers/rustbelt/paper.pdf>

Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Aleš Bizjak, Lars Birkedal, and Derek Dreyer. 2018b. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming* 28 (2018), e20. <https://people.mpi-sws.org/~dreyer/papers/iris-ground-up/paper.pdf>

Steve Klabnik and Carol Nichols. 2023. *The Rust programming language*. No Starch Press.

Donald E. Knuth. 1998. *The Art of Computer Programming, Volume 3: (2nd Ed.) Sorting and Searching*. Addison Wesley Longman Publishing Co., Inc., USA.

Peter J. Landin. 1964. The Mechanical Evaluation of Expressions. *Computer Journal* 6, 4 (Jan. 1964), 308–320.

Anton Lorenzen, Leo White, Stephen Dolan, Richard A. Eisenberg, and Sam Lindley. 2024. Oxidizing OCaml with Modal Memory Management. *Proc. ACM Program. Lang.* 8, ICFP (2024), 485–514. [doi:10.1145/3674642](https://doi.org/10.1145/3674642)

Per Martin-Löf. 1982. Constructive Mathematics and Computer Programming. In *Logic, Methodology and Philosophy of Science VI. Studies in Logic and the Foundations of Mathematics*, Vol. 104. Elsevier, 153–175. [doi:10.1016/S0049-237X\(09\)70189-2](https://doi.org/10.1016/S0049-237X(09)70189-2)

Mae Milano, Joshua Turcotti, and Andrew C. Myers. 2022. A flexible type system for fearless concurrency. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 458–473. [doi:10.1145/3519939.3523443](https://doi.org/10.1145/3519939.3523443)

Robin Milner. 1978. A Theory of Type Polymorphism in Programming. *J. Comput. System Sci.* 17, 3 (Dec. 1978), 348–375. <http://citeserx.ist.psu.edu/viewdoc/summary?doi=10.1.1.67.5276>

Alexandre Moine, Stephanie Balzer, Alex Xu, and Sam Westrick. 2025a. TypeDis: A Type System for Disentanglement (Artifact). [doi:10.5281/zenodo.1733635](https://doi.org/10.5281/zenodo.1733635)

Alexandre Moine, Stephanie Balzer, Alex Xu, and Sam Westrick. 2025b. TypeDis: A Type System for Disentanglement (Extended Version). (Nov. 2025). <https://arxiv.org/abs/2511.23358>

Alexandre Moine, Sam Westrick, and Stephanie Balzer. 2024. DisLog: A Separation Logic for Disentanglement. *Proc. ACM Program. Lang.* 8, POPL, Article 11 (Jan. 2024), 30 pages. [doi:10.1145/3632853](https://doi.org/10.1145/3632853)

Peter Müller. 2002. *Modular Specification and Verification of Object-Oriented Programs*. Lecture Notes in Computer Science, Vol. 2262. Springer. [doi:10.1007/3-540-45651-1](https://doi.org/10.1007/3-540-45651-1)

Stefan K. Müller, Umut A. Acar, and Robert Harper. 2017. Responsive parallel computation: bridging competitive and cooperative threading. In *38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, 677–692. [doi:10.1145/3062341.3062370](https://doi.org/10.1145/3062341.3062370)

James Noble, Jan Vitek, and John Potter. 1998. Flexible Alias Protection. In *12th European Conference on Object-Oriented Programming (ECOOP) (Lecture Notes in Computer Science, Vol. 1445)*. Springer, 158–185. [doi:10.1007/BF0054091](https://doi.org/10.1007/BF0054091)

Martin Odersky, Martin Sulzmann, and Martin Wehr. 1999. Type Inference with Constrained Types. *Theory and Practice of Object Systems* 5, 1 (1999), 35–55. <https://infoscience.epfl.ch/handle/20.500.14299/221715>

Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press.

Andrew M. Pitts and Ian Stark. 1998. Operational Reasoning for Functions with Local State. *Higher Order Operational Techniques in Semantics (HOOTS)* (1998), 227–273.

Gordon D. Plotkin. 1973. *Lambda-definability and logical relations*. Technical Report. University of Edinburgh.

François Pottier and Vincent Simonet. 2003. Information flow inference for ML. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 25, 1 (2003), 117–158. [doi:10.1145/596980.596983](https://doi.org/10.1145/596980.596983)

Ram Raghunathan, Stefan K. Muller, Umut A. Acar, and Guy E. Blelloch. 2016. Hierarchical memory management for parallel programs. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18–22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 392–406. [doi:10.1145/2951913.2951935](https://doi.org/10.1145/2951913.2951935)

Andrei Sabelfeld and Andrew C. Myers. 2003. Language-Based Information-Flow Security. *IEEE J. Sel. Areas Commun.* 21, 1 (2003), 5–19.

Julian Shun, Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, Aapo Kyrola, Harsha Vardhan Simhadri, and Kanat Tangwongsan. 2012. Brief announcement: the problem based benchmark suite. In *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25–27, 2012*, Guy E. Blelloch and Maurice Herlihy (Eds.). ACM, 68–70. [doi:10.1145/2312005.2312018](https://doi.org/10.1145/2312005.2312018)

Vincent Simonet. 2003. Flow Caml in a Nutshell. In *1st APPSEM-II Workshop*, Graham Hutton (Ed.).

Geoffrey Smith and Dennis M. Volpano. 1998. Secure Information Flow in a Multi-Threaded Imperative Language. In *POPL*. ACM, 355–364.

Richard Statman. 1985. Logical Relations and the Typed λ -calculus. *Information and Control* 65, 2/3 (1985), 85–97. [doi:10.1016/S0019-9958\(85\)80001-2](https://doi.org/10.1016/S0019-9958(85)80001-2)

William W. Tait. 1967. Intensional Interpretations of Functionals of Finite Type I. *The Journal of Symbolic Logic* 32, 2 (1967), 198–212. <http://www.jstor.org/stable/2271658>

Amin Timany, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2024. A Logical Approach to Type Soundness. *J. ACM* 71, 6, Article 40 (Nov. 2024), 75 pages. [doi:10.1145/3676954](https://doi.org/10.1145/3676954)

Mads Tofte and Lars Birkedal. 1998. A Region Inference Algorithm. *ACM Trans. Program. Lang. Syst.* 20, 4 (1998), 724–767. [doi:10.1145/291891.291894](https://doi.org/10.1145/291891.291894)

Mads Tofte, Lars Birkedal, Martin Elsman, and Niels Hallenberg. 2004. A Retrospective on Region-Based Memory Management. *Higher-Order and Symbolic Computation* 17, 3 (Sept. 2004), 245–265. <https://doi.org/10.1023/B:LISP.0000029446.78563.a4>

Mads Tofte and Jean-Pierre Talpin. 1997. Region-based memory management. *Information and Computation* 132, 2 (1997), 109–176. <http://www.irisa.fr/prive/talpin/papers/ic97.pdf>

VerifyThis. 2022. Challenge 3 - The World's Simplest Lock-Free Hash Set. <https://ethz.ch/content/dam/ethz/special-interest/infk/chair-program-method/pm/documents/Verify%20This/Challenges2022/verifyThis2022-challenge3.pdf>

Simon Friis Vindum and Lars Birkedal. 2021. Contextual refinement of the Michael-Scott queue. In *Certified Programs and Proofs (CPP)*. 76–90. <https://cs.au.dk/~birke/papers/2021-ms-queue-final.pdf>

Dennis M. Volpano, Cynthia E. Irvine, and Geoffrey Smith. 1996. A Sound Type System for Secure Flow Analysis. *J. Comput. Secur.* 4, 2/3 (1996), 167–188.

Philip Wadler. 1990. Linear Types Can Change the World!. In *IFIP Working Group 2.2, 2.3 on Programming Concepts and Methods*. North-Holland, 561.

Philip Wadler. 2012. Propositions as Sessions. In *ACM SIGPLAN International Conference on Functional Programming (ICFP)*. ACM, 273–286. [doi:10.1145/2364527.2364568](https://doi.org/10.1145/2364527.2364568)

Sam Westrick. 2022. *Efficient and Scalable Parallel Functional Programming through Disentanglement*. Ph.D. Dissertation. Department of Computer Science, Carnegie Mellon University.

Sam Westrick, Jatin Arora, and Umut A. Acar. 2022. Entanglement Detection with Near-Zero Cost. *Proc. ACM Program. Lang.* 6, ICFP, Article 115 (aug 2022), 32 pages. [doi:10.1145/3547646](https://doi.org/10.1145/3547646)

Sam Westrick, Rohan Yadav, Matthew Fluet, and Umut A. Acar. 2020. Disentanglement in Nested-Parallel Programs. *Proc. ACM Program. Lang.* 4, POPL, Article 47 (jan 2020), 32 pages. [doi:10.1145/3371115](https://doi.org/10.1145/3371115)

Andrew K. Wright. 1995. Simple Imperative Polymorphism. *Lisp and Symbolic Computation* 8, 4 (Dec. 1995), 343–356. <http://www.cs.rice.edu/CS/PLT/Publications/Scheme/lasc95-w.ps.gz>

Received 2025-07-10; accepted 2025-11-06