DisLog: A Separation Logic for Disentanglement

ALEXANDRE MOINE, Inria, France
SAM WESTRICK, Carnegie Mellon University, USA
STEPHANIE BALZER, Carnegie Mellon University, USA

Disentanglement is a run-time property of parallel programs that facilitates task-local reasoning about the memory footprint of parallel tasks. In particular, it ensures that a task does not access any memory locations allocated by another concurrently executing task. Disentanglement can be exploited, for example, to implement a high-performance parallel memory manager, such as in the MPL (MaPLe) compiler for Parallel ML. Prior research on disentanglement has focused on the design of optimizations, either trusting the programmer to provide a disentangled program or relying on runtime instrumentation for detecting and managing entanglement. This paper provides the first static approach to verify that a program is disentangled: it contributes DisLog, a concurrent separation logic for disentanglement. DisLog enriches concurrent separation logic with the notions necessary for reasoning about the fork-join structure of parallel programs, allowing the verification that memory accesses are effectively disentangled. A large class of programs, including race-free programs, exhibit memory access patterns that are disentangled "by construction". To reason about these patterns, the paper distills from DisLog an almost standard concurrent separation logic, called DisLog+. In this high-level logic, no specific reasoning about memory accesses is needed: functional correctness proofs entail disentanglement. The paper illustrates the use of DisLog and DisLog+ on a range of case studies, including two different implementations of parallel deduplication via concurrent hashing. All our results are mechanized in the Coq proof assistant using Iris.

CCS Concepts: • **Software and its engineering** → **Parallel programming languages**; • **Theory of computation** → **Separation logic**; **Program verification**.

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

1 INTRODUCTION

Recent work has shown that parallel functional programming can deliver the same efficiency and scalability as imperative and procedural approaches. The key to this line of work is a memory property known as *disentanglement* [Arora et al. 2021, 2023; Guatto et al. 2018; Raghunathan et al. 2016; Westrick et al. 2022, 2020], which restricts parallel tasks to access data that was allocated "before" the task executed. This restriction enables tasks to allocate and garbage-collect memory locally and independently—that is, without synchronizing with other parallel tasks. Utilizing disentanglement, Arora et al. [2023] developed a provably efficient memory manager for functional programs which also provides full support for effects. All of this work is implemented in MPL

Authors' addresses: Alexandre Moine, alexandre.moine@inria.fr, Inria, Paris, France; Sam Westrick, swestric@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, USA; Stephanie Balzer, balzers@cs.cmu.edu, Carnegie Mellon University, Pittsburgh, USA.

("maple"), an open-source[1] compiler for Parallel ML. In practice, MPL has been shown to be fast, scalable, and competitive with lower-level and imperative language implementations.

This line of work relies on disentanglement to ensure efficiency and scalability, and leaves it up to the programmer to reason about disentanglement and its performance impact. For purely functional code, reasoning about disentanglement is not an issue: purely functional programs are guaranteed to be disentangled, by construction, due to the lack of mutation (in-place updates). Programmers can also use pure libraries that are implemented under-the-hood with in-place updates for efficiency. For example, common parallel operations (such as map, reduce, scan, etc.) can be implemented efficiently with mutable arrays, hidden behind a pure interface, and can then be used to write disentangled code. However, in this example, the developers of high-performance libraries still need to reason carefully about disentanglement. More generally, whenever in-place updates and other low-level optimizations are necessary for efficiency, disentanglement needs to be taken into account.

In the context of in-place updates and other memory effects, reasoning about disentanglement is subtle. Programmers may wish to use concurrent data structures (for example, lock-free hash tables) to improve efficiency. Such data structures can be made disentangled [Westrick 2022], but reasoning about their correctness is challenging, even for experts. If disentanglement is violated, there can be significant consequences for performance, in terms of increased time and space usage [Arora et al. 2023]. In this sense, disentanglement can be considered a "safety" condition for performance-oriented code.

Therefore, we shift our attention to static verification of disentanglement. Our goal is to support reasoning about both high-level and low-level code, including atomic in-place updates and concurrent data structures, which can require identifying intricate invariants. In this setting, concurrent separation logic [Brookes 2007; O'Hearn 2007] and its modern variants [Jung et al. 2018; Nanevski et al. 2014] have proven to be successful vehicles for verifying safety and correctness properties of programs in the presence of challenging concurrent features. An intriguing question is whether or not separation logic can be used to prove disentanglement, which we address in this paper.

To verify disentanglement statically, we develop **DisLog**, the first program logic for proving disentanglement, and formally prove its soundness. At a high level, DisLog is a concurrent separation logic built on Iris [Jung et al. 2018] endowed with assertions which describe dependencies between parallel tasks and permissions to make disentangled loads from the heap. This approach makes the logic powerful enough to verify disentanglement even in complex and subtle situations, such as programs with lock-free data structures and algorithms using atomic in-place reads and writes.

Going further, on top of DisLog, we develop **DisLog+**, a standard concurrent separation logic which hides the details of disentanglement, allowing for standard separation logic proofs while also getting proofs of disentanglement for free. DisLog+ is applicable for a wide variety of programs, including purely functional programs, race-free programs, and even programs that have "benign" memory races (for example, write-write races). Importantly, DisLog and DisLog+ work seamlessly together, allowing for DisLog+ proofs to drop into the more powerful DisLog where needed (for example, for verifying non-pure segments of mostly pure programs), and otherwise stay at a high level of abstraction.

To evaluate DisLog and DisLog+, we consider several case studies, including key parallel primitives, as well as sophisticated parallel algorithms involving concurrent data structures. In all cases, we prove that the programs are disentangled. Our experience has shown that, using the logics developed in this paper, the effort of proving disentanglement is typically small. Furthermore, when a formal proof of functional correctness is desired, using DisLog+ often yields a proof of disentanglement for free.

---

[1]https://github.com/mpllang/mpl

Our contributions include:

- DisLog, the first program logic to verify that a program is disentangled (§4). It employs the notion of **timestamps** to reason about the nested fork-join parallelism of a program and introduces a novel **clock** assertion to prove that memory accesses do not cause entanglement.
- DisLog+, a high-level logic built on top of DisLog that shields the user from timestamp management (§ 5). As a result, race-free programs can be verified in DisLog+ with the standard reasoning rules of concurrent separation logic.
- Two mechanisms allowing to reason about benign races in DisLog+, including fractional write-only assertions for write-write races (§5.4), and a set of rules for read-write races on pre-allocated data (§5.5).
- A range of case studies (§6), including multiple parallel primitives, parallel lookup in a lazy collection, and two examples of deduplication via concurrent hashing.
- A formalization in the Coq proof assistant using Iris [Jung et al. 2018]. All our results are mechanized (§7) in Coq, including: the two program logics, their soundness theorems, and the case studies [Moine et al. 2023b].

2 KEY IDEAS

2.1 Background

Nested Fork-Join Parallelism. We consider programs written using a single parallel primitive: the parallel tuple $e_1 || e_2$. It executes e_1 and e_2 in parallel, and returns their results as a pair. Here, two **tasks** are spawned to execute expressions e_1 and e_2 in parallel. Parallel tuples may be arbitrarily nested. For example, the expression e_1 might itself execute another parallel tuple. This leads to a dynamic nesting structure of tasks called the **task tree**, where the leaves are tasks that may take steps in parallel. In an operational semantics, the task tree is maintained by two distinguished reductions: a **fork**, where two tasks are spawned to execute e_1 and e_2 in parallel, and a **join**, when the tasks complete and return their results as a pair. Parallel tasks can be understood as concurrent threads *with a structure*: if they terminate, the two tasks forked by a parallel pair ultimately join.

This style of programming is known as *nested fork-join parallelism*, or sometimes *nested task parallelism*. The arbitrary nesting of parallel tasks allows programmers to write parallel recursive divide-and-conquer style algorithms. For example, a "parallel for-loop" can be implemented by splitting the index range in half and then recursively executing the two halves in parallel (§6.2).

Acquiring locations. During execution, each task may allocate locations in memory and perform memory effects such as reads and writes on these locations, including atomic compare-and-swaps (CAS). The reads in particular are important for the disentanglement property we consider. In our operational semantics (§3.3), a read occurs in three distinct cases: (1) a memory load inside an array reads the indexed value, (2) a closure call reads the environment of the closure, and (3) a CAS reads the scrutinized value. When a task performs a read, if the result of that read is a location, we say that the task **acquires** the resulting location.

Disentanglement. Disentanglement limits communication between concurrent tasks by restricting which locations may be acquired: each task may always acquire its own allocations, and additionally, each task may acquire any location allocated "before" the task began. The notion of "before" relates to the dependencies induced by forks and joins. A forking task comes before the two tasks it forks, and conversely, two joining tasks come before their join point. If a task ever acquires a location allocated by some other task that is executing concurrently, this constitutes *entanglement*. **The logics developed in this paper allow proving that a program is disentangled**, i.e., that in every possible execution of the program, entanglement will never occur.

```
 1  fun scratch () =
 2    let
 3      val shared = newScratchpad()
 4      val {tryLock, releaseLock} =
 5        newLock()
 6      fun myTask() =
 7        if tryLock() then
 8          (doWork shared;
 9           clearScratchpad shared;
10           releaseLock())
11        else
12          let val x = newScratchpad()
13          in doWork x
14    in
15      // two calls in parallel (could be
16      // generalized to many calls if desired)
17      (myTask() ∥ myTask())
```

```
18  fun newLock () =
19    let val r = ref false
20    in { tryLock = fn () ⇒ CAS(r, false, true),
21         releaseLock = fn () ⇒ r := false }
22
23  type elem = ...
24  val defaultElem : elem = ...
25
26  type scratchpad = elem array
27  fun newScratchpad () =
28    Array.allocate (N, defaultElem)
29  fun clearScratchpad scratch =
30    for i from 0 to N − 1 do
31      scratch[i] := defaultElem
32
33  fun doWork scratchpad =
34    // ... read and write to scratchpad
```

(a) The scratch function calls doWork multiple times in parallel, and provides a suitable scratchpad for each call.

(b) Auxiliary code for scratch, including locks, scratchpads, and a function doWork that requires a scratchpad as temporary space.

Fig. 1. The scratch example used to illustrate our approach.

Atomic operations and determinacy races. A key feature of disentanglement is that it allows for non-deterministic interleaving of atomic in-place operations such as atomic loads, stores, and CASes. Such operations are commonly used under-the-hood in the implementation of high-performance libraries (for example, in the implementation of a high-throughput lock-free hash table). Proving disentanglement in this setting requires reasoning carefully about atomic operations that can acquire locations.

In other words, one of the challenges is to prove that entanglement is impossible whenever there is a *determinacy race* [Feng and Leiserson 1999]. Concretely, a **determinacy race** occurs whenever two atomic in-place operations are performed concurrently at the same location, and at least one of the operations modifies the location. For example, an atomic load could race with an atomic store, or an atomic load could race with a CAS, or two CASes could race with each other, etc. As the name suggests, determinacy races can lead to non-deterministic execution, which we allow.

In this paper, we permit only atomic (i.e., properly synchronized) in-place operations, and therefore avoid all *data races* [Adve 2010; Boehm 2011; Dolan et al. 2018] by construction. For simplicity, we assume a sequentially consistent memory model. Because of the lack of data races, throughout the paper, we use the term **race** to refer only to determinacy races.

2.2 Running Example

To illustrate the ideas in the paper, we use a running example, called scratch, shown in Fig. 1. This function is non-deterministic due to a determinacy race, yet is disentangled. At a high level, scratch calls the function doWork two times in parallel. Each call to doWork uses an array (called a "scratchpad") as temporary space. Note that it would be safe to allocate a fresh scratchpad for every call to doWork. The goal of the example is to optimize performance by reducing the number of scratchpads that are allocated. (The example only calls doWork twice, but this could be generalized to any number of calls in parallel, which would make the optimization more significant.)

To reduce the number of allocated scratchpads, scratch implements a simple strategy. First, a shared scratchpad is allocated together with a lock to protect it (Fig. 1, lines 3–5). Then, before
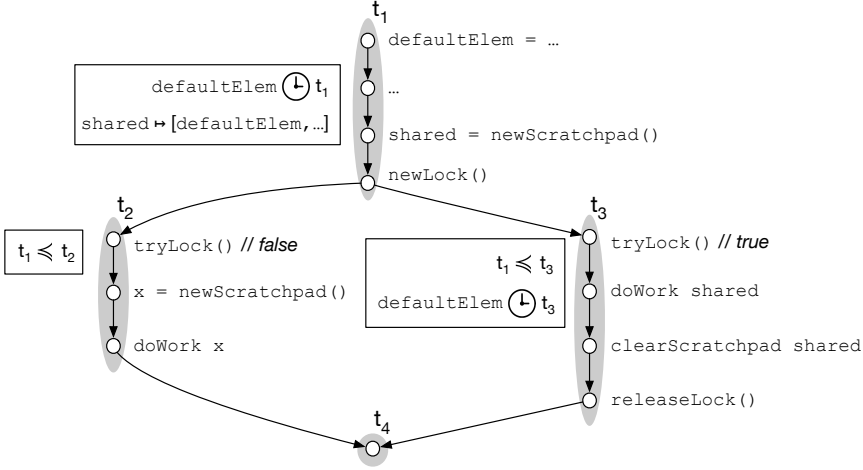
Fig. 2. One possible execution of the example of Fig. 1

each call to doWork, scratch will attempt to claim access to the shared scratchpad by calling tryLock (line 7). If this succeeds, then doWork may use the shared scratchpad (line 8); otherwise, scratch falls back on allocating a fresh scratchpad (line 12). Whenever a call to doWork is finished using the shared scratchpad, the shared scratchpad is cleared (line 9), before finally releasing the lock (line 10). In this way, scratch reduces the number of allocations by reusing the shared scratchpad as much as possible. In particular, if scratch is executed using only a single processor, then every call to doWork will be able to use the shared scratchpad, and no additional scratchpads will be allocated.

The auxiliary code for the example is shown in Fig. 1b, which defines scratchpads and locks. The details of doWork are not important as long as it performs reads and writes on the scratchpad. Locks are implemented by a pair of closures with a mutable boolean, indicating whether the lock has been locked. The closure tryLock is implemented using an atomic compare-and-swap (CAS) which returns **true** if the CAS succeeds, and **false** otherwise.

Disentanglement in scratch. Proving that the scratch example is disentangled is subtle. In particular, doWork may read or write to the scratchpad, and we need to show that reading from the scratchpad will never return a value allocated by a concurrent task. Thankfully, the scratch function guarantees a strong precondition: when doWork begins, the argument scratchpad will contain only the value defaultElem, which is allocated before every call to doWork, and therefore is safe with respect to disentanglement.

The precondition on doWork is easily satisfied on line 13, because the scratchpad is freshly allocated. Showing that the precondition is also satisfied on line 8 is more subtle, because there is an invariant on the shared scratchpad which is determined by the state of the lock. Informally, the invariant is: "while the lock is not held, for every i, shared[i] = defaultElem." This invariant is re-established by calling clearScratchpad (line 9) before releasing the lock. Note that removing this call to clearScratchpad may lead to an entangled state. Indeed, after a call to doWork, the scratchpad may contain locally-allocated data, which is hence available for the other task.

2.3 Disentanglement: Timestamps, Reads, and How to Reason about Them

Partial orders on tasks through timestamps. In Sec. 3, we present a semantics that facilitates reasoning about disentanglement. To this end, we enrich the semantics with the notion of a *timestamp*, a unique identifier for each parallel task. We then assign every heap-allocated location

a timestamp, marking *when* (i.e., by which task) the location was allocated, and restrict every task to only depend on locations allocated at timestamps that come *before* their timestamp. Timestamps form a partial order which respects the dependencies induced by forks and joins. When a task forks, the semantics generates two timestamps (one for each task) that are preceded by the timestamp of the forking task. Conversely, when two tasks join, the semantics generates a timestamp that is preceded by both timestamps of the two joining tasks. Forks and joins are the only operations extending the partial order of timestamps. Tasks otherwise step independently.

Fig. 2 visualizes one possible execution of the running example. Each shaded oval represents a task, and is labeled by its timestamp t_i. The framed content is not relevant yet. Execution begins on a task t_1, which allocates the scratchpad and the lock. Then, a fork occurs, generating two tasks t_2 and t_3. The task t_2 fails to win the lock, whereas the task t_3 succeeds. After completing, the two tasks join, forming a new task t_4.

A program logic for timestamp orders. In Sec. 4, we develop a program logic, DisLog, that incorporates timestamps. Every expression in DisLog is associated with a *current* timestamp and an *end* timestamp. The program logic uses a *weakest precondition* (WP) modality which takes the form

$$\mathsf{wp} \langle t,\, e \rangle \, \{\lambda\, t'\, v.\ \Phi\}$$

asserting that the expression e is currently evaluated by a task at timestamp t, that e is disentangled and can reduce, and if its reduction terminates, then it does so at end timestamp t', yielding a value v such that Φ holds. The partial order of timestamps is encoded into the assertions of the logic, via the precedence assertion $t \leqslant t'$. The precedence assertion is *persistent* and hence duplicable at will. This assertion describes the parallel structure of the program being verified.

Examples of the precedence assertion $t \leqslant t'$ appear in the framed boxes of Fig. 2. To reason about the task t_2, the user gets an assertion $t_1 \leqslant t_2$. Dually, the user gets an assertion $t_1 \leqslant t_3$ to reason about the task t_3. At the join point t_4, the user gets the assertions $t_2 \leqslant t_4$ and $t_3 \leqslant t_4$. Making use of the fact the $\leqslant$ is a pre-order, the user can use transitivity and deduce for example that $t_1 \leqslant t_4$.

Preserving disentanglement. Acquiring a memory location ℓ from the heap puts disentanglement at risk. Disentanglement is preserved if and only if ℓ was allocated at some timestamp preceding the timestamp t of the acquiring task—in that case, we say that ℓ was allocated before t. To represent such a requirement, we introduce the clock assertion, written $\ell \odot t$, which precisely asserts that the ℓ was allocated before t. This assertion appears for example in the precondition of DisLog's Load rule (§4.3). To illustrate this rule, we show a specialized instantiation allowing to load the element at index 0 in the shared scratchpad, here named ℓ. The premises are implicitly separated by the separating conjunction $*$.

$$\textsc{SpecializedLoad}\ \frac{\mathsf{shared} \mapsto [\ell;\dots] \qquad \ell \odot t}{\mathsf{wp}\,\langle t,\, \mathsf{shared}[0]\rangle\,\{\lambda\,t'\,v.\ \ulcorner t' = t \wedge v = \ell \urcorner\ *\ \mathsf{shared} \mapsto [\ell;\dots]\,\}}$$

The SpecializedLoad rule first requires, as in standard separation logic, ownership of the shared location via a points-to assertion. Crucially, this rule also requires that the loaded value ℓ was allocated before the current timestamp t via the $\ell \odot t$ assertion. In the postcondition of the WP, the rule asserts that the end timestamp t' is equal to the previous timestamp t, the returned value v is ℓ, and the user still has the points-to ownership.

The clock assertion is persistent, giving the user great flexibility. Moreover, it is *monotonic with respect to the precedence pre-order*. Hence, if the user knows that the location ℓ was allocated before t and that t precedes t' they can then deduce that ℓ was allocated before t'. This mechanism is illustrated in Fig. 2. Indeed, the user can produce an assertion $\mathsf{defaultElem} \odot t_1$ upon the allocation of $\mathsf{defaultElem}$. Then, while reasoning about t_3 the user can use the assertion $t_1 \leqslant t_3$ to generate

a new clock assertion defaultElem $\odot$ t_3. As explained next, our high-level logic DisLog+ takes full advantage of the monotonicity of the clock assertion.

2.4 Going High-Level: Simple Programs Should Have Simple Proofs

Readers familiar with proofs of realistic programs may be worried by timestamps polluting the logic, and the additional proof burden imposed on a common rule such as Load. In practice, many programs "don't poke the bear" and subtle reasoning about timestamps should not be needed. For example, Westrick et al. [2020] show that race-free programs are always disentangled, as they prevent the communication of concurrently-allocated locations. Verifying such programs should be as cheap as a standard separation logic proof.

In Sec. 5, we present DisLog+, a high-level separation logic where timestamps, clocks, and precedence are confined to very few occurrences. DisLog+ allows reasoning on race-free programs with the standard reasoning rules of concurrent separation logic. The sole difference is a restriction on ghost state, effectively preventing races (§5.3). What is the secret of the DisLog+ logic? The key observation we make on race-free programs is that (1) the content of a freshly allocated location is always safe to acquire for the allocating task and (2) "being safe to acquire" is a *monotonic* property: if a location is safe to acquire for a given task, it is safe to acquire for every subsequent task. As long as a program does not write carelessly to a shared location and break monotonicity, locations are always safe to acquire and no reasoning about timestamps is needed.

Technically, we define assertions of DisLog+ as monotonic predicates of DisLog over an ambient timestamp, the latter being implicitly threaded through during the proof. Points-to assertions of DisLog+ store not only ownership information but also the proof that all pointed-to locations are safe to acquire at the ambient timestamp. Hence, DisLog+ provides a standard Load reasoning rule. We stress that DisLog+ is a light abstraction over DisLog. At any moment during the proof, the user of DisLog+ can fall back to DisLog for fine timestamp-related reasoning.

Our definition of DisLog+ is directly inspired by separation logics for weak-memory models: iGPS [Kaiser et al. 2017], iRC11 [Dang et al. 2020] and Cosmo [Mével et al. 2020]. These high-level logics are defined in terms of low-level logics by implicitly threading through a monotonic *view* of the memory.

Beyond race freedom. It turns out that DisLog+ is not confined to reasoning about race-free programs, but additionally provides two new sets of high-level reasoning rules to accommodate the most elementary disentangled races. The first one consists of *fractional write-only assertions* (§5.4) allowing the user to reason about write-write races within DisLog+. As a write-write race does not acquire any location, such a race is always disentangled. The second one consists of a set of rules unveiling just enough timestamps to reason about races on "obviously safe" data (§5.5). These data include data that was allocated before the beginning of the parallel phase, and *unboxed* data—that is, data that is not allocated in the heap.

The language we model supports the atomic operation compare-and-swap (CAS). A CAS is an entanglement hazard. Indeed, a CAS reads the scrutinized value, which must be safe to acquire. In the scratch running example (Fig. 1b), we use CAS to implement a spin-lock. Here, we exploit unboxed data to allow parallel tasks to safely communicate via a race on shared reference r. A "race on unboxed data" should ring a bell: it perfectly fits in the realm of DisLog+ and its extensions. We show in Sec. 6 how to reason about our locks and scratch entirely within the high-level DisLog+.

3 LANGUAGE AND SEMANTICS

Our language, DisLang, is an imperative lambda-calculus with fork-join parallelism. We equip Dis-Lang with a small-step, substitution-based, call-by-value semantics, guaranteeing disentanglement.

Values	$\mathcal{V}$	$v, w ::= ()$	$b \in \{\text{true}, \text{false}\}$	$i \in \mathbb{Z}$	$\ell \in \mathcal{L}$	$\hat{\mu}f.\lambda\vec{x}.e$ where $fv(e) \subseteq (\{f\} \cup \vec{x})$		

Values $\mathcal{V}$ $v, w ::= () \mid b \in \{\text{true}, \text{false}\} \mid i \in \mathbb{Z} \mid \ell \in \mathcal{L} \mid \hat{\mu}f.\lambda\vec{x}.e$ where $fv(e) \subseteq (\{f\} \cup \vec{x})$

Blocks $r ::= \vec{w} \mid \mu f.\lambda\vec{x}.e$

Primitives $\bowtie ::= + \mid - \mid \times \mid \div \mid \text{mod} \mid {==} \mid < \mid \leq \mid > \mid \geq \mid \vee \mid \wedge$

Expressions $e ::= v$ *value* $\text{alloc } e \; e$ *array allocation*

           x *variable* $e[e]$ *array load*

           $\text{let } x = e \text{ in } e$ *sequencing* $e[e] \leftarrow e$ *array store*

           $\text{if } e \text{ then } e \text{ else } e$ *conditional* $\text{length } e$ *array length*

           $e\,\vec{e}$ *call* $e \,||\, e$ *parallel tuple*

           $e \bowtie e$ *primitive operation* $\text{CAS } e\,e\,e\,e$ *compare-and-swap*

           $\mu f.\lambda\vec{x}.e$ *closure allocation*

Contexts $K ::= \text{let } x = \square \text{ in } e \mid \text{if } \square \text{ then } e \text{ else } e \mid \text{alloc } \square\, e \mid \text{alloc } v\, \square \mid \text{length } e$

               $\mid \square[e] \mid v[\square] \mid \square[e] \leftarrow e \mid v[\square] \leftarrow e \mid v[v] \leftarrow \square$

               $\mid \square \bowtie e \mid v \bowtie \square \mid \square\,\vec{t} \mid v\,(\vec{v} + \square + \vec{t})$

               $\mid \text{CAS } \square\, t\, t\, t \mid \text{CAS } v\, \square\, t\, t \mid \text{CAS } v\, v\, \square\, t \mid \text{CAS } v\, v\, v\, \square$

Fig. 3. Syntax of DisLang

3.1 Syntax

The syntax of DisLang appears in Fig. 3. 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 *top-level function* $\hat{\mu}f.\lambda\vec{x}.e$. A top-level function is closed in the sense that the only variables available in the function body e are the function's name f and the formal parameters $\vec{x}$.

A *block* describes the contents of a heap cell, amounting to either an array of values, written $\vec{w}$, or a *λ-abstraction* $\mu f.\lambda\vec{x}.e$. Lambdas, as opposed to top-level functions $\hat{\mu}f.\lambda\vec{x}.e$, are not values. Instead, they are compiled to heap-allocated *closures* [Appel 1992; Landin 1964]. Hence, acquiring a lambda can create entanglement. Top-level functions can be seen as closures that are pre-allocated outside the heap, which thus cannot create entanglement. In DisLang, fork-join parallelism is available via the parallel tuple $e_1 \,||\, e_2$, representing the expressions e_1 and e_2 to be computed in parallel. DisLang 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. An 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.

3.2 Computation Graphs and Disentanglement

The dynamics of DisLang, presented in the next section, makes use of a *computation graph*, capturing the nested fork-join parallel structure of a program. A computation graph is a directed acyclic graph where vertices, or *tasks*, represent sequential computations, and edges represent the dependencies between them [Acar et al. 2016]. We label each task with a unique *timestamp* t, from an infinite set $\mathcal{T}$. When a task t_0 forks two fresh tasks t_1 and t_2, the computation graph is extended with edges (t_0, t_1) and (t_0, t_2). Conversely, when two completed tasks t_1 and t_2 join to form a fresh task t_3 the computation graph is extended with edges (t_1, t_3) and (t_2, t_3). As discussed earlier, an example computation graph for the scratch example (§2) is shown in Fig. 2.

In a computation graph G, we say that t *precedes* t' and write $t \preccurlyeq_G t'$ when there exists a sequence of edges in G from t to t'. In particular, we say that two tasks are *concurrent* when neither precedes the other. Entanglement occurs when a task acquires a location that was allocated by a concurrent task. Recall our running example (§2, Fig. 1): a particular implementation of doWork can store a locally-allocated location in the shared scratchpad. If it were possible for two concurrent tasks to both win the lock, without proper cleaning of the scratchpad, this locally-allocated location could be acquired by the concurrent task, generating entanglement.

HEADALLOC
$$\frac{0 \le n \qquad \ell \notin \mathrm{dom}(\sigma) \qquad \ell \notin \mathrm{dom}(\alpha)}{G, t \vdash \sigma \backslash \alpha \backslash \mathsf{alloc}\, n\, v \longrightarrow [\ell := v^n]\sigma \backslash [\ell := t]\alpha \backslash \ell}$$

HEADCALLPRIM
$$\frac{v_1 \bowtie v_2 \xrightarrow{\mathrm{pure}} v}{G, t \vdash \sigma \backslash \alpha \backslash v_1 \bowtie v_2 \longrightarrow \sigma \backslash \alpha \backslash v}$$

HEADCLOSURE
$$\frac{\ell \notin \mathrm{dom}(\sigma) \qquad \ell \notin \mathrm{dom}(\alpha)}{G, t \vdash \sigma \backslash \alpha \backslash \mu f. \lambda \vec{x}. e \longrightarrow [\ell := \mu f. \lambda \vec{x}. e]\sigma \backslash [\ell := t]\alpha \backslash \ell}$$

HEADLENGTH
$$\frac{\sigma(\ell) = \vec{w}}{G, t \vdash \sigma \backslash \alpha \backslash \mathsf{length}\, \ell \longrightarrow \sigma \backslash \alpha \backslash |\vec{w}|}$$

HEADSTORE
$$\frac{\sigma(\ell) = \vec{w} \qquad 0 \le i < |\vec{w}|}{G, t \vdash \sigma \backslash \alpha \backslash \ell[i] \leftarrow v \longrightarrow [\ell := [i := v]\vec{w}]\sigma \backslash \alpha \backslash ()}$$

HEADLETVAL
$$G, t \vdash \sigma \backslash \alpha \backslash \mathsf{let}\, x = v \,\mathsf{in}\, e \longrightarrow \sigma \backslash \alpha \backslash [v/x]e$$

HEADIFTRUE
$$G, t \vdash \sigma \backslash \alpha \backslash \mathsf{if}\,\mathsf{true}\,\mathsf{then}\, e_1 \,\mathsf{else}\, e_2 \longrightarrow \sigma \backslash \alpha \backslash e_1$$

HEADIFFALSE
$$G, t \vdash \sigma \backslash \alpha \backslash \mathsf{if}\,\mathsf{false}\,\mathsf{then}\, e_1 \,\mathsf{else}\, e_2 \longrightarrow \sigma \backslash \alpha \backslash e_2$$

HEADLOAD
$$\frac{\sigma(\ell) = \vec{w} \qquad 0 \le i < |\vec{w}| \qquad \vec{w}(i) = v \qquad \boxed{(v \in \mathcal{L} \implies \alpha(v) \le_G t)}}{G, t \vdash \sigma \backslash \alpha \backslash \ell[i] \longrightarrow \sigma \backslash \alpha \backslash v}$$

HEADCALL
$$\frac{(v = \hat{\mu} f. \lambda \vec{x}. e) \vee (v \in \mathcal{L} \wedge \sigma(v) = \mu f. \lambda \vec{x}. e) \qquad |\vec{x}| = |\vec{w}| \qquad \boxed{(\forall \ell.\ \ell \in locs(e) \implies \alpha(\ell) \le_G t)}}{G, t \vdash \sigma \backslash \alpha \backslash v\, \vec{w} \longrightarrow \sigma \backslash \alpha \backslash [v/f][\vec{w}/\vec{x}]e}$$

HEADCASSUCC
$$\frac{\sigma(\ell) = \vec{w} \qquad 0 \le i < |\vec{w}| \qquad \vec{w}(i) = v_0 \qquad v_0 = v \qquad \boxed{(v_0 \in \mathcal{L} \implies \alpha(v_0) \le_G t)}}{G, t \vdash \sigma \backslash \alpha \backslash \mathsf{CAS}\, \ell\, i\, v\, v' \longrightarrow [\ell := [i := v']\vec{w}]\sigma \backslash \alpha \backslash \mathsf{true}}$$

HEADCASFAIL
$$\frac{\sigma(\ell) = \vec{w} \qquad 0 \le i < |\vec{w}| \qquad \vec{w}(i) = v_0 \qquad v_0 \ne v \qquad \boxed{(v_0 \in \mathcal{L} \implies \alpha(v_0) \le_G t)}}{G, t \vdash \sigma \backslash \alpha \backslash \mathsf{CAS}\, \ell\, i\, v\, v' \longrightarrow \sigma \backslash \alpha \backslash \mathsf{false}}$$

Fig. 4. Head reduction. The disentanglement proof obligation is highlighted.

3.3 Operational Semantics

Head Reduction. Fig. 4 defines the head reduction relation $G, t \vdash \sigma \backslash \alpha \backslash e \longrightarrow \sigma' \backslash \alpha' \backslash e'$ between two *head configurations* $\sigma \backslash \alpha \backslash e$ and $\sigma' \backslash \alpha' \backslash e'$, where G is the (global) computation graph and t the timestamp of the (local) task at which the reduction takes place. A head configuration consists of the expression e being evaluated, the *store* σ, and an *allocation map* α. A store σ is a finite map of locations to blocks, representing the heap, and an allocation map α is a finite map of locations to timestamps, recording the timestamps at which locations were allocated.

We write $\sigma(\ell)$ to denote the block stored at the location ℓ in the store σ. To insert a block into the store or update the store, we write $[\ell := r]\sigma$. Note that only arrays can be updated; closures are immutable. To refer to the index i of an array $\vec{w}$, we write $\vec{w}(i)$, and to update an array, we write $[i := v]\vec{w}$. 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.

The HEADALLOC and HEADCLOSURE reductions allocate heap blocks, arrays and closures, respectively, extending the store with the desired block and the allocation map with the current timestamp. The HEADCALLPRIM reduction encompasses a reduction $\xrightarrow{\mathrm{pure}}$ to compute a primitive operation. The HEADSTORE reduction updates the field of an array, and the HEADLENGTH reduction returns the length of an array. The HEADLETVAL reduction substitutes a variable by its value. The HEADIFTRUE and HEADIFFALSE reductions reduce an if-then-else construction where the conditional is evaluated.

Entanglement may only occur when a task acquires a location. Locations are acquired during the reductions HEADLOAD, HEADCALL, HEADCASSUCC and HEADCASFAIL. A load reads the indexed value, a call the environment of the closure, and a CAS the scrutinized value. The HEADCALL

SCHEDHEAD
$$\frac{G, t \vdash \sigma \backslash \alpha \backslash e \longrightarrow \sigma' \backslash \alpha' \backslash e'}{\sigma / \alpha / G / t / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G / t / e'}$$

SCHEDFORK
$$\frac{t_1, t_2 \notin \text{vertices}(G) \qquad G' = G \cup \{(t_0, t_1), (t_0, t_2)\}}{\sigma / \alpha / G / t_0 / e_1 \| e_2 \xrightarrow{\text{sched}} \sigma / \alpha / G' / t_1 \otimes t_2 / e_1 \| e_2}$$

SCHEDJOIN
$$\frac{\ell \notin \text{dom}(\sigma) \qquad \ell \notin \text{dom}(\alpha) \qquad t_3 \notin \text{vertices}(G) \qquad G' = G \cup \{(t_1, t_3), (t_2, t_3)\}}{\sigma / \alpha / G / t_1 \otimes t_2 / v_1 \| v_2 \xrightarrow{\text{sched}} [\ell := [v_1; v_2]]\sigma / [\ell := t_3]\alpha / G' / t_3 / \ell}$$

STEPSCHED
$$\frac{\sigma / \alpha / G / T / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G' / T' / e'}{(\sigma, \alpha, G) / T / e \xrightarrow{\text{step}} (\sigma', \alpha', G') / T' / e'}$$

STEPBIND
$$\frac{S / T / e \xrightarrow{\text{step}} S' / T' / e'}{S / T / K[e] \xrightarrow{\text{step}} S' / T' / K[e']}$$

STEPPARL
$$\frac{S / T_1 / e_1 \xrightarrow{\text{step}} S' / T_1' / e_1'}{S / T_1 \otimes T_2 / e_1 \| e_2 \xrightarrow{\text{step}} S' / T_1' \otimes T_2 / e_1' \| e_2}$$

STEPPARR
$$\frac{S / T_2 / e_2 \xrightarrow{\text{step}} S' / T_2' / e_2'}{S / T_1 \otimes T_2 / e_1 \| e_2 \xrightarrow{\text{step}} S' / T_1 \otimes T_2' / e_1 \| e_2'}$$

Fig. 5. Reduction under a context and parallelism

reduction distinguishes between invoking a top-level function and a closure. Calling a closure loads the values of its environment, which may contain locations. As we use a substitution-based semantics, these locations are the location literals occurring in the function body e, which are computed by the $locs(e)$ function. To prevent entanglement, all the mentioned rules include the same kind of precondition (highlighted in Fig. 4): if ℓ is acquired, then its allocation timestamp $\alpha(\ell)$ must precede the timestamp t of the task at which the reduction takes place. These preconditions amount to a *proof obligation* during the verification of a program. Verified programs will satisfy the obligation and will thus never get stuck. As we will see in Sec. 4.4 and in Sec. 5.3, soundness of both of our logics entail the invariant that the physical program state are always disentangled.

Parallelism and Reduction under a Context. To keep track of the currently active and suspended tasks of an executing parallel program, 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: $T \triangleq t \in \mathcal{T} \mid T \otimes T$. A leaf represents an active task and is denoted by its timestamp t. A node $T_1 \otimes T_2$ represents a suspended task that has forked two parallel computations, recursively described by the task trees T_1 and T_2.

Taking advantage of task trees, we define the semantics of parallel reductions and reductions under a context in Fig. 5. We define a scheduling reduction $\sigma / \alpha / G / T / e \xrightarrow{\text{sched}} \sigma' / \alpha' / G' / T' / e'$ 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. The SCHEDHEAD reduction describes a head reduction. The SCHEDFORK reduction describes a fork: the task tree is at a leaf t_0 and faces a parallel tuple. 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 the node with two leaves $t_1 \otimes t_2$. The SCHEDJOIN reduction describes a join: the task tree is at a node with two leaves $t_1 \otimes t_2$, and both leaves reached a value. The reduction generates a fresh timestamp t_3, updates the computation graph, and allocates a memory cell to store the result of the parallel tuple. It then updates the task tree to the leaf t_3.

The main reduction relation $S / T / e \xrightarrow{\text{step}} S' / T' / e'$ describes a scheduling reduction inside the whole parallel program. A tuple $S / T / e$ consists of the program state S, the task tree T, and an expression e. A state S consists of the tuple (σ, α, G), denoting a store σ, an allocation map α, and a computation graph G. The STEPSCHED reduction describes a scheduling step. The other reductions

describe *where* the scheduling reduction takes place. The STEPBIND reduction describes a reduction under an evaluation context. The STEPPARL and STEPPARR reductions unveil the non-determinism of the parallel reduction. If a node of the task tree is encountered facing a parallel tuple, the left side or the right side can reduce.

4 DISLOG, A PROGRAM LOGIC FOR DISENTANGLEMENT

In this section, we present the details of DisLog. We first give an Iris primer and explain our notations (§ 4.1). Then, we showcase how timestamps appear in the program logic (§ 4.2), and present other reasoning rules (§ 4.3). Finally, we discuss the soundness theorem of DisLog (§ 4.4).

4.1 Assertions and Weakest Preconditions

We build DisLog on top of Iris [Jung et al. 2018], adopting Iris' syntax. In particular, we write Φ for an Iris assertion (of type *iProp*), $\Phi * \Phi'$ for a separating conjunction, and $\Phi \twoheadrightarrow \Phi'$ for a separating implication. If U is a proposition of the meta logic, we call U *pure* and write $\ulcorner U \urcorner$. We write $\Phi \dashv\vdash \Phi'$ for the equivalence of assertions.

Our program logic features a weakest precondition (WP) modality which takes the form:

$$\text{wp} \langle t, e \rangle \{\lambda t' v. \Phi\}$$

This modality adapts a standard Iris' WP to the semantics of DisLang, and in particular, enriches it to account for timestamps. In the above assertion, t is the timestamp of the task which symbolically executes the expression e. We call this timestamp the *current timestamp* of the expression. A postcondition takes the form $\lambda t' v. \Phi$ where the variables t' and v are bound in Φ. The variable v denotes the resulting value and the variable t' the *end* timestamp, the timestamp of the returning task. We write $\text{wp} \langle t, e \rangle \{\lambda t' \ell. \Phi\}$, where the variable ℓ denotes a location, as a syntactic sugar for $\text{wp} \langle t, e \rangle \{\lambda t' v. \exists \ell. \ulcorner v = \ell \urcorner * \Phi\}$. We similarly do so for booleans b and integers i. If we want to abstract over the details of the postcondition, we write Ψ instead of $\lambda t' v. \Phi$.

Our WP is subject to the standard structural rules of separation logic. DisLog supports in particular the FRAME rule that we present below, as a warm-up to our notations. We write reasoning rules as inference rules, where premises are separated by the separating conjunction $*$ and entail the conclusion. In particular, if the conclusion is a WP, premises amount to preconditions.

$$\text{FRAME} \frac{\Phi_0 \qquad \text{wp} \langle t, e \rangle \{\lambda t' v. \Phi_1\}}{\text{wp} \langle t, e \rangle \{\lambda t' v. \Phi_0 * \Phi_1\}}$$

Separation logic triples can be obtained with the standard definition $\{\Phi\} \langle t, e \rangle \{\Psi\} \triangleq \Box(\Phi \twoheadrightarrow \text{wp} \langle t, e \rangle \{\Psi\})$, where $\Box$ stands for the persistence modality of Iris. The persistence modality characterize *persistent* assertions (an assertion Φ is persistent when $\Phi \dashv\vdash \Box \Phi$). Once a persistent assertion holds, it holds forever. In particular, persistent assertions are duplicable.

Iris features *ghost state*, which is hence available in DisLog. We write $\Phi \Rrightarrow \Phi'$ for a *ghost update* (or *fancy update*) that updates the ghost state. We omit the so-called *masks* for the sake of readability. Thanks to the ghost state, DisLog supports Iris *invariants* [Jung et al. 2018, §2.2], with a standard interface. Our WP allows the user to assume (or *open*) an invariant before reasoning about an *atomic* expression and generates an obligation to restore (or *close*) the invariant in the postcondition. An atomic expression is an expression that can reduce to a value in a single head step of computation. We syntactically characterize such assertions with the Atomic e pure predicate.

DisLog makes use of fractional [Bornat et al. 2005; Boyland 2003] and discardable [Vindum and Birkedal 2021] points-to assertions of the form $\ell \mapsto_p \vec{w}$, where p denotes either a positive fraction less than or equal to 1, or a discarded fraction written $\Box$. The latter makes the points-to assertion

$$\frac{\text{CLOCKMONO}}{\ell \oplus t \quad t \preccurlyeq t'}{\ell \oplus t'} \qquad \frac{\text{PRECREFL}}{t \preccurlyeq t} \qquad \frac{\text{PRECTRANS}}{t \preccurlyeq t' \quad t' \preccurlyeq t''}{t \preccurlyeq t''} \qquad \frac{\text{MEMENTOPRE}}{\ulcorner \ell \in locs(e) \urcorner \quad \ell \oplus t \mathbin{-\!*} \text{wp}\, \langle t, e \rangle\, \{\Psi\}}{\text{wp}\, \langle t, e \rangle\, \{\Psi\}}$$

$$\frac{\text{MEMENTOPOST}}{\text{wp}\, \langle t, e \rangle\, \{\lambda\, t'\, v.\, v \oplus t' \mathbin{-\!*} \Psi\, t'v\}}{\text{wp}\, \langle t, e \rangle\, \{\Psi\}} \qquad \frac{\text{TEMPUSFUGIT}}{\text{wp}\, \langle t, e \rangle\, \{\lambda\, t'\, v.\, t \preccurlyeq t' \mathbin{-\!*} \Psi\, t'v\}}{\text{wp}\, \langle t, e \rangle\, \{\Psi\}} \qquad \frac{\text{TEMPUSATOMIC}}{\ulcorner \text{Atomic}\, e \urcorner \quad \text{wp}\, \langle t, e \rangle\, \{\lambda\, \_\, v.\, \Psi\, t\, v\}}{\text{wp}\, \langle t, e \rangle\, \{\Psi\}}$$

Fig. 6. Reasoning rules for clocks and precedence assertions

persistent. When $p = 1$ we write $\ell \mapsto \vec{w}$. Points-to assertions of DisLog do not carry information about timestamps: this role is devoted to two new assertions described in the next Section.

4.2 Timestamps Management

A central aspect of our disentanglement logic is the management of timestamps. To this end, DisLog features two new assertions.

- The *clock* assertion $\ell \oplus t$, indicating that the location ℓ was allocated before the timestamp t in the underlying computation graph.
- The *precedence* assertion $t \preccurlyeq t'$, witnessing that the timestamp t precedes the timestamp t' in the underlying computation graph.

Both assertions are persistent and work hand-in-hand. Given the assertion $\ell \oplus t$, a task at timestamp t can safely acquire the location ℓ. Moreover, given both the assertions $\ell \oplus t$ and $t \preccurlyeq t'$, a task at timestamp t' can safely acquire the location ℓ as well. A benefit of phrasing a location's allocation timestamp relative to another timestamp, rather than absolute, is that the user never needs to know precisely at which timestamp a location was allocated: disentanglement is ensured as soon as the acquired location was allocated by a preceding task. Similarly, the user never needs to know the whole computation graph: precedence information suffices for proving disentanglement. We overload the clock assertion to arbitrary values v and introduce assertions of the form $v \oplus t$. If v is a location ℓ, then this assertion is defined as $\ell \oplus t$. Otherwise, it is defined as $\ulcorner True \urcorner$. We overload again this assertion to a collection of values, and write $\vec{w} \oplus t$ for the iterated conjunction $\mathord{\scalebox{1.3}{$*$}}_{(v \in \vec{w})} (v \oplus t)$.

Fig. 6 summarizes the rules governing the clock and the precedence assertions. The CLOCKMONO rule illustrates the monotonicity of the clock predicate with respect to the precedence pre-order: if the location ℓ was allocated before t and t precedes t', then it is safe to conclude that ℓ was allocated before t'. We emphasize that the precedence assertion forms a pre-order: this assertion is reflexive (PRECREFL) and transitive (PRECTRANS). The MEMENTOPRE and MEMENTOPOST rules are the only rules generating a clock predicate. The MEMENTOPRE rule asserts that if the location ℓ occurs in the expression e at current timestamp t, then the user can gain a witness $\ell \oplus t$ that ℓ was allocated at a timestamp preceding t. The MEMENTOPOST rule asserts that the value returned by a task was allocated before the end timestamp of this task.

The TEMPUSFUGIT rule distills the semantics of DisLang: it is safe to suppose that the current timestamp precedes the end timestamp. The TEMPUSATOMIC rule asserts that the current timestamp and the end timestamp of an atomic expression are the same. The TEMPUSATOMIC rule is more precise than needed: the clock predicate and the precedence predicate are both monotonic with respect to the precedence pre-order, via the CLOCKMONO rule and the PRECTRANS rule, respectively. However, the TEMPUSATOMIC rule relieves the user from the burden of always applying the CLOCKMONO and PRECTRANS rules by hand when the timestamp is effectively preserved.

VALUE
$$\frac{\Psi\, t\, v}{\text{wp}\,\langle t, v\rangle\,\{\Psi\}}$$

ALLOC
$$\frac{\ulcorner 0 \le n \urcorner}{\text{wp}\,\langle t, \text{alloc}\ n\ v\rangle\,\{\lambda\_\ell.\ \ell \mapsto v^n\}}$$

LOAD
$$\frac{\ulcorner 0 \le i < |\vec{w}|\ \wedge\ \vec{w}(i) = v \urcorner \qquad \ell \mapsto_p \vec{w} \qquad v \circleddash t}{\text{wp}\,\langle t, \ell[i]\rangle\,\{\lambda\_v'.\ \ulcorner v' = v \urcorner\ *\ \ell \mapsto_p \vec{w}\}}$$

CLOSURE
$$\text{wp}\,\langle t, \mu f.\lambda\vec{x}.e\rangle\,\{\lambda\_\ell.\ \text{Func}\ \ell\ f\ \vec{x}\ e\}$$

TOPLEVEL
$$\frac{\ulcorner v = \hat{\mu}f.\lambda\vec{x}.e \urcorner}{\text{Func}\ v\ f\ \vec{x}\ e}$$

LENGTH
$$\frac{\ell \mapsto_p \vec{w}}{\text{wp}\,\langle t, \text{length}\ \ell\rangle\,\{\lambda\_i.\ \ulcorner i = |\vec{w}| \urcorner\ *\ \ell \mapsto_p \vec{w}\}}$$

CALLPRIM
$$\frac{\ulcorner v_1 \bowtie v_2 \xrightarrow{\text{pure}} v \urcorner}{\text{wp}\,\langle t, v_1 \bowtie v_2\rangle\,\{\lambda\_v'.\ \ulcorner v' = v \urcorner\}}$$

BIND
$$\frac{\text{wp}\,\langle t, e\rangle\,\{\lambda\,t'\,v.\ \text{wp}\,\langle t', K[v]\rangle\,\{\Psi\}\}}{\text{wp}\,\langle t, K[e]\rangle\,\{\Psi\}}$$

LETVAL
$$\frac{\text{wp}\,\langle t, [v/x]e\rangle\,\{\Psi\}}{\text{wp}\,\langle t, \text{let}\ x = v\ \text{in}\ e\rangle\,\{\Psi\}}$$

CALL
$$\frac{\ulcorner |\vec{x}| = |\vec{w}| \urcorner \qquad \text{Func}\ v\ f\ \vec{x}\ e \qquad \text{wp}\,\langle t, [v/f][\vec{w}/\vec{x}]e\rangle\,\{\Psi\}}{\text{wp}\,\langle t, v\ \vec{w}\rangle\,\{\Psi\}}$$

IFTRUE
$$\frac{\text{wp}\,\langle t, e_1\rangle\,\{\Psi\}}{\text{wp}\,\langle t, \text{if true then } e_1 \text{ else } e_2\rangle\,\{\Psi\}}$$

IFFALSE
$$\frac{\text{wp}\,\langle t, e_2\rangle\,\{\Psi\}}{\text{wp}\,\langle t, \text{if false then } e_1 \text{ else } e_2\rangle\,\{\Psi\}}$$

CASSUCC
$$\frac{\ulcorner 0 \le i < |\vec{w}|\ \wedge\ \vec{w}(i) = v_0\ \wedge\ v_0 = v \urcorner \qquad \ell \mapsto \vec{w} \qquad v_0 \circleddash t}{\text{wp}\,\langle t, \text{CAS}\ \ell\ i\ v\ v'\rangle\,\{\lambda\_b.\ \ulcorner b = \text{true} \urcorner\ *\ \ell \mapsto [i := v']\vec{w}\}}$$

CASFAIL
$$\frac{\ulcorner 0 \le i < |\vec{w}|\ \wedge\ \vec{w}(i) = v_0\ \wedge\ v_0 \ne v \urcorner \qquad \ell \mapsto_p \vec{w} \qquad v_0 \circleddash t}{\text{wp}\,\langle t, \text{CAS}\ \ell\ i\ v\ v'\rangle\,\{\lambda\_b.\ \ulcorner b = \text{false} \urcorner\ *\ \ell \mapsto_p \vec{w}\}}$$

STORE
$$\frac{\ulcorner 0 \le i < |\vec{w}| \urcorner \qquad \ell \mapsto \vec{w}}{\text{wp}\,\langle t, \ell[i] \leftarrow v\rangle\,\{\lambda\_\_.\ \ell \mapsto [i := v]\vec{w}\}}$$

PAR
$$\frac{\forall t_1\, t_2.\ t \preccurlyeq t_1\ *\ t \preccurlyeq t_2 \Rrightarrow \exists \Psi_1\, \Psi_2.\ \text{wp}\,\langle t_1, e_1\rangle\,\{\Psi_1\}\ *\ \text{wp}\,\langle t_2, e_2\rangle\,\{\Psi_2\}\ *}{\text{wp}\,\langle t, e_1 \,||\, e_2\rangle\,\{\Psi\}}$$
$$\big(\forall t_1'\, v_1\, t_2'\, v_2\, t'\, \ell.\ \Psi_1\, t_1'\, v_1\ *\ \Psi_2\, t_2'\, v_2\ *\ t_1' \preccurlyeq t'\ *\ t_2' \preccurlyeq t'\ *\ \ell \mapsto [v_1; v_2] \rightarrowtail \Psi\, t'\, \ell\big)$$

Fig. 7. Syntax-directed rules of DisLog

4.3 Reasoning Rules for Expressions

Fig. 7 gives the syntax-directed reasoning rules of DisLog (we hide "later" modalities for brevity).
The rules ALLOC, LENGTH, CALLPRIM, LETVAL, IFTRUE, IFFALSE, and STORE are standard, apart from
their mention of timestamps. In particular, the ALLOC rule does not generate a clock assertion. If
desired, such an assertion can be obtained by applying the MEMENTOPOST rule.

The VALUE rule asserts that if the symbolic evaluation of an expression ended at timestamp t,
yielding a value v, then the postcondition $\Psi\, t\, v$ should hold. The LOAD rule extends the standard
separation logic rule to prevent entanglement. Indeed, the $v \circleddash t$ assertion in the precondition
witnesses that if v is a location, then it must have been allocated before the current timestamp t.
The CASSUCC and CASFAIL rules are similarly extended: they prevent entanglement by requiring
that if the scrutinized value is a location, then it was allocated before the current timestamp.

The CLOSURE and TOPLEVEL rules produce an assertion Func $v\, f\, \vec{x}\, e$ certifying that calling v as a
function will not cause entanglement. Obtaining this assertion for closures may be surprising at
first, but is warranted by the following facts: *(i)* all the timestamps of locations captured by the
closure are guaranteed to precede the closure's allocation timestamp t (by rule MEMENTOPRE), and
(ii) closures are immutable objects and, as such, cannot themselves create entanglement [Westrick
et al. 2022]. Phrased differently, the locations of the environment are allocated before the closure
itself, and thanks to immutability, this fact never changes. The Func predicate is persistent. The
CALL rule allows calling a function, given the Func predicate. Proving that the environment was
allocated before the current timestamp amounts to proving that the closure's location itself was

allocated before the current timestamp, which is true since the closure's location is already part of the expression (§4.4).

The BIND rule gives meaning to the notion of the "current timestamp" of an expression. Operationally, the evaluation of a term $K[e]$ at timestamp t reduces the sub-expression e until it reaches a value v and an end timestamp t'. Then, the whole term $K[v]$ starts reducing at the new timestamp t'. The BIND rule paraphrases this operational behavior. The rule asserts that the user first has to reason about the sub-expression e at the same current timestamp. The user has then to reason about the filled term $K[v]$ at a current timestamp t', under the precondition that the sub-expression e reduced to a value v at end timestamp t'.

Reasoning About a Parallel Tuple. A pivotal rule of DisLog is PAR. Let's first derive a naive version PARWEAK of this rule below before focusing on the ultimate rule given in Fig. 7.

$$\frac{\forall t_1\, t_2. \qquad t \leqslant t_1 \twoheadrightarrow \mathsf{wp}\, \langle t_1,\, e_1\rangle\, \{\Psi_1\} \qquad t \leqslant t_2 \twoheadrightarrow \mathsf{wp}\, \langle t_2,\, e_2\rangle\, \{\Psi_2\}}{\mathsf{wp}\, \langle t,\, e_1 \,||\, e_2\rangle \left\{\lambda\, t'\, \ell. \begin{array}{c} \exists t_1'\, v_1\, t_2'\, v_2.\ \Psi_1\, t_1'\, v_1 * \Psi_2\, t_2'\, v_2 \\ t_1' \leqslant t' * t_2' \leqslant t' * \ell \mapsto [v_1; v_2] \end{array}\right\}}\ \text{PARWEAK}$$

This rule allows reasoning about a parallel tuple $e_1 \,||\, e_2$ at current timestamp t. The premise universally quantifies over two fresh timestamps t_1 and t_2, which are used for e_1 and e_2, respectively. We focus on the left-hand side of the tuple as the right-hand side is handled similarly. The user should verify e_1 with the postcondition Ψ_1, under the hypothesis $t \leqslant t_1$ witnessing that t precedes t_1. This information allows the user to safely acquire any location that was safe to acquire from t. Indeed, if the user has an assertion $\ell \odot t$, they can use the CLOCKMONO rule to obtain an assertion $\ell \odot t_1$.

We emphasize that the above rule ensures that the two fresh timestamps t_1 and t_2 are *unrelated*. Hence, an assertion $\ell \odot t_1$ *cannot* be converted to an assertion $\ell \odot t_2$. This would indeed be unsafe, as a location allocated by the left task could be acquired by the right one, creating entanglement.

After the join point, the postcondition of the PARWEAK rule asserts that e_1 reduced to a value v_1 at end timestamp t_1', and that e_2 reduced to a value v_2 at end timestamp t_2'. The postcondition also produces witnesses that t_1' and t_2' precede the (new) current timestamp t'. Thanks to these two assertions any locations allocated by either of the two tasks are now accessible by any task at timestamp t'' such that $t' \leqslant t''$. Finally, the postcondition asserts that the parallel tuple itself reduced to a location ℓ, pointing to the two resulting values v_1 and v_2.

Unfortunately, the PARWEAK rule is tedious to use in practice. It fails to support a common pattern underlying our proof rules, which would allow the postconditions Ψ_1 and Ψ_2 of the newly forked tasks to depend on the tasks' timestamps t_1 and t_2. This dependence is rendered impossible by the universal quantification of the timestamps t_1 and t_2. Our final rule PAR presented in Fig. 7 facilitates the wished-for pattern. The premise of the PAR rule quantifies universally over the two timestamps, and *after* the quantification, allows the user to choose two existentially quantified postconditions Ψ_1 and Ψ_2 that can depend on the two timestamps. Moreover, the user is free to choose the postconditions after a potential ghost update; for example, to allocate an invariant that depends on both t_1 and t_2. The user should then verify the two parts of the parallel tuple with their respective timestamps, as in the PARWEAK rule. Finally, the user has to show that the resulting values and timestamps entail the postcondition Ψ. This is formally expressed in the second line of the precondition of the rule.

4.4 Soundness

Finally, we devote our attention to stating and proving soundness of DisLog. For our disentanglement logic to be sound it has to hold that the reduction of a program verified using the rules of DisLog leads to a disentangled program state. Our semantics is phrased in terms of a transition system that

REDPAR

REDSCHED

$\dfrac{S\,/\,T\,/\,e \xrightarrow{\text{sched}} S'\,/\,T'\,/\,e'}{\text{red}\,S\,T\,e}$ REDCTX $\dfrac{\text{red}\,S\,T\,e}{\text{red}\,S\,T\,(K[e])}$ $\dfrac{(e_1 \notin \mathcal{V} \;\vee\; e_2 \notin \mathcal{V})}{(e_1 \notin \mathcal{V} \implies \text{red}\,S\,T_1\,e_1) \qquad (e_2 \notin \mathcal{V} \implies \text{red}\,S\,T_2\,e_2)}$
$\text{red}\,S\,(T_1 \otimes T_2)\,(e_1 \,||\, e_2)$

Fig. 8. Reducibility of a configuration

$$\text{wp}\,\langle t,\,e\rangle\,\{\Psi\} \quad\triangleq\quad \text{wpg}\,\langle t,\,e\rangle\,\{\Psi\}$$

$$\text{wpg}\,\langle T,\,e\rangle\,\{\Psi\} \quad\triangleq$$
$$\left(\ulcorner e \in \mathcal{V} \urcorner \;*\; (\forall S.\ \text{interp}\,S\,T\,e \implies \ulcorner T \in \mathcal{T}\urcorner \;*\; \text{interp}\,S\,T\,e \;*\; \Psi\,T\,e) \right)$$
$$\vee \Big(\ulcorner e \notin \mathcal{V}\urcorner \;*\; (\forall S.\ \text{interp}\,S\,T\,e \implies$$
$$\ulcorner \text{red}\,S\,T\,e\urcorner \;*\; \forall S'\,T'\,e'.\ \ulcorner S\,/\,T\,/\,e \xrightarrow{\text{step}} S'\,/\,T'\,/\,e'\urcorner \implies \rhd \implies \text{interp}\,S'\,T'\,e' \;*\; \text{wpg}\,\langle T',\,e'\rangle\,\{\Psi\}) \Big)$$

Fig. 9. Definition of the weakest precondition modalities

gets stuck if entanglement is encountered, ensured by the highlighted premises for head reductions in Fig. 4. Soundness of our logic thus must entail that verified programs cannot get stuck.

Since we use a small-step semantics in a parallel world, the definition of "not getting stuck" needs careful wording. In particular, it is not enough to say that "the configuration can take a step". Indeed, one step of DisLang corresponds to a step of *one* task, whereas we want to ensure that *every task* can take a proper step. The adequacy theorem of the Iris WP makes use of the notion of *reducibility* to capture the fact that a thread can take a proper step. The judgment red $S\,T\,e$ presented in Fig. 8 adapt this notion, ensuring that every task of the task tree can take a step. The REDSCHED rule asserts that a configuration that can make a scheduling step (either a head step, a fork, or a join) is reducible. The REDCTX rule asserts that the reducibility of a configuration facing an expression under an evaluation context amounts to the reducibility of this very expression. The REDPAR rule asserts that a configuration facing a node of the task tree and a parallel tuple is reducible if at least one side of the pair is not a value (otherwise, a join should be possible), and each side that is not a value is reducible.

An expression e is *safe* if $(\emptyset, \emptyset, \emptyset)\,/\,t\,/\,e \xrightarrow{\text{step}}{}^{*} S'\,/\,T'\,/\,e'$ implies that either the configuration $S'\,/\,T'\,/\,e'$ is reducible, or that e' is a value and T' a single leaf. Our soundness theorem asserts that if an expression e can be verified using DisLog, then it is safe.

THEOREM 4.1 (SOUNDNESS OF DISLOG). *If wp* $\langle t,\,e\rangle\,\{\lambda \_\_.\ \ulcorner True\urcorner\}$ *holds, then e is safe.*

As the semantics of DisLang cannot progress when entanglement is detected, the soundness theorem asserts that e cannot reach an entangled state. The formal proof Theorem 4.1 can be found in our Coq formalization [Moine et al. 2023b]. We detail below the main definitions and invariants.

Definition of the Weakest Precondition. The formal definition of the wp $\langle t,\,e\rangle\,\{\Psi\}$ assertion can be found in Fig. 9. The key to our approach is to define the wp modality with respect to a more general—hidden from the user—weakest precondition modality that we refer to as wpg modality. The wpg modality is parameterized not with a single timestamp, but a whole task tree: we found this generalization crucial for the various proofs to succeed. Nevertheless, reasoning always takes place at the leaves. Hence, we can hide the details of the task tree from the user.

The definition of the wpg modality appears also in Fig. 9 and follows the traditional Iris recipe [Jung et al. 2018, §6]. As usual, the WP is defined as a guarded fixpoint, and makes use of a *state interpretation predicate* (or *central invariant*), written interp, which relates the ghost state and the physical state. The definition of wpg cases on whether the expression is a value or not. If the expression is a value, we can access the state interpretation, and deduce that the task

RDeLeaf
$$\frac{\forall \ell.\ \ell \in locs(e) \implies \alpha(\ell) \preccurlyeq_G t}{\text{rootsde } \alpha\, G\, t\, e}$$

RDePar
$$\frac{\text{rootsde } \alpha\, G\, T_1\, e_1 \qquad \text{rootsde } \alpha\, G\, T_2\, e_2}{\text{rootsde } \alpha\, G\, (T_1 \otimes T_2)\, (e_1 \,||\, e_2)}$$

RDeCtx
$$\frac{\text{rootsde } \alpha\, G\, T\, e \qquad \forall \ell\, t.\ \ell \in locs(K)\ \wedge\ t \in \mathsf{leaves}(T) \implies \alpha(\ell) \preccurlyeq_G t}{\text{rootsde } \alpha\, G\, T\, (K[e])}$$

$$\mathsf{interp}\,(\sigma, \alpha, G)\, T\, e \quad\triangleq\quad \begin{array}{l} \ulcorner \mathsf{dom}(\sigma) = \mathsf{dom}(\alpha)\ \wedge\ \text{rootsde } \alpha\, G\, T\, e \urcorner\ * \\ \overline{\lfloor \bullet G \rfloor}^{\gamma}\ *\ \mathsf{Heap}\,\sigma\ *\ \text{\Large$*$}_{(\ell,t)\in\alpha}\,(\mathsf{meta}\,\ell\,t) \end{array}$$

$$\mathsf{edge}\,t\,t' \triangleq \overline{\lfloor \circ\{(t,t')\} \rfloor}^{\gamma} \qquad t \preccurlyeq t' \triangleq \mathsf{rtc}\ \mathsf{edge}\,t\,t' \qquad \ell \odot t \triangleq \exists t_0.\ \mathsf{meta}\,\ell\,t_0\ *\ t_0 \preccurlyeq t$$

$$\mathsf{Func}\,v\,f\,\vec{x}\,e \triangleq \ulcorner v = \hat{\mu} f.\,\lambda \vec{x}.\,e \urcorner \vee$$
$$\left(\ulcorner v \in \mathcal{L} \urcorner\ *\ v \mapsto_\square (\mu f.\,\lambda \vec{x}.\,e)\ *\ \exists t.\ \mathsf{meta}\,v\,t\ *\ \text{\Large$*$}_{\ell' \in locs(e)}\,(\ell' \odot t) \right)$$

Fig. 10. Definition of the state interpretation predicate and of base assertions

tree consists of a single leaf and the postcondition. Otherwise, if the expression is not a value, the wpg modality asserts that the configuration is reducible, and that for any possible step, the state interpretation must continue to hold, as well as the WP of the reduced program. Apart from its mention of timestamps, our WP distinguishes itself from the standard Iris WP by making the state interpretation available in the value case, and making use of our custom red judgment.

The State Interpretation Predicate. Our WP maintains a state interpretation predicate between each reduction step, which is defined in Fig. 10. We review its definition next.

We first focus on the roots disentanglement judgment rootsde $\alpha\, G\, T\, e$. This judgment asserts each task of the expression e only uses locations allocated before its associated timestamp. This judgment allows stating the MementoPre rule. If the task tree consists of a single leaf t, the RDeLeaf rule requires that the locations of e were allocated before t. The RDePar rule requires that both sides of a parallel tuple satisfy the rootsde judgment. In the case of an evaluation context, RDeCtx requires that the judgment holds for the expression under the context, and that the locations occurring in the evaluation context itself are allocated before all the leaves of the task tree.

The state interpretation predicate also gives meaning to the ghost state from the physical state. We first briefly explain the construction of ghost state in Iris abstractly, before detailing the part of interp that concerns ghost state. In Iris, ghost state is defined in terms of so-called *cameras* (CMRA) which can be thought as "step-indexed partial commutative monoids" [Jung et al. 2018], detailing a resource algebra. Iris provides predefined notions of resource algebras. For example, the resource algebra $Auth(M)$ describes the *authoritative resource algebra* over the resources M. This resource algebra gives access to $\bullet a$, the *authoritative* ownership of a, and $\circ b$, the *fragmentary* ownership of b. Together, these two assertions entail that there exists an element c of the algebra such that $a = b \cdot c$. The resource algebra $Set(M)$ describes the *set resource algebra*, where the composition of resources is described by set union. For our state interpretation predicate, we define a ghost cell γ which we equip with the resource algebra $Auth(Set(\mathcal{T} \times \mathcal{T}))$. The ghost cell γ stores the computation graph and gives meaning to the precedence assertion.

Iris moreover provides a generic construction to define points-to assertions via the gen_heap library [Iris Development Team 2023]. This library defines a certain piece of ghost state, defines an assertion Heap σ that ties a store σ to this ghost state, and defines the points-to assertion $\ell \mapsto_p \vec{w}$ in terms of this ghost state. Moreover, the gen_heap library allows associating persistent information

to locations via a mechanism of *meta* assertions. In our case, we associate to each location ℓ the timestamp t of the task that allocated it, and write meta $\ell\, t$. The main property of this assertion is that, from the knowledge meta $\ell\, t$ and meta $\ell\, t'$, we can deduce that $t = t'$.

We are now able to review the details of the definition of our state interpretation predicate shown in the lower part of Fig. 10. First, it asserts that the domain of the store is the same as the allocation map, and the roots disentanglement judgment. The state interpretation also asserts the ghost authoritative ownership of the computation graph $\lceil \bullet G \rceil^\gamma$ and the ownership of the store via the Heap σ assertion. Moreover, the state interpretation asserts, for every mapping from a location ℓ to a timestamp t in the allocation map α, that the persistent knowledge meta $\ell\, t$ was set.

We define an edge between t and t' as a ghost fragmentary ownership of the singleton $\lceil \circ\{(t, t')\} \rceil^\gamma$. The conjunction of $\lceil \bullet G \rceil^\gamma$ and $\lceil \circ\{(t, t')\} \rceil^\gamma$ allows to deduce that $(t, t') \in G$. We define the precedence assertion $t \preccurlyeq t'$ as the reflexive-transitive-closure (rtc) over the edge predicate. The clock assertion $\ell \odot t$ is defined as a paraphrase of its informal definition. The location ℓ was allocated before timestamp t if there exists a timestamp t_0 such that ℓ was allocated at t_0, and t_0 precedes t.

The representation predicate of a λ-abstraction Func $v\, f\, \vec{x}\, e$ is a disjunction: either v is a top-level function, or a heap-allocated closure. In that case, we use a discarded fraction for the points-to assertion, as the closure is immutable. The predicate also asserts the existence of a timestamp t at which the closure was allocated, and that every location of its environment (the locations occurring in e) was allocated before t. We make use of this knowledge to verify the CALL rule of Fig. 7. The closure's location is allocated before the current timestamp (thanks to the rootsde judgment), but since the locations of the environment were allocated before the allocation time of the closure itself, they are also allocated before the current timestamp, and hence safe to acquire.

5 A HIGH-LEVEL LOGIC: DISLOG+

In this section, we introduce DisLog+, an almost standard concurrent separation logic allowing proof of disentanglement for a large class of programs. DisLog+ is defined in terms of DisLog using monotonicity arguments, a technique that first appeared in program logics targeting weak-memory models [Dang et al. 2020; Kaiser et al. 2017; Mével et al. 2020].

5.1 Don't Poke the Bear

Disentanglement is preserved by restricting reads: when a task acquires a location, the programmer must ensure that this location was allocated by a preceding task. However, numerous programs "don't poke the bear", that is, are disentangled because they do not comprise reads of shared, and hence potentially hazardous, locations.

Determinacy-race-free programs are an example of such cautious programs. A *determinacy race* [Feng and Leiserson 1999] occurs when two concurrent tasks access the same location atomically, and at least one of these accesses is a write. As noticed by Westrick et al. [2020], such race-free programs are trivially disentangled: shared locations cannot be written to from different tasks, which prevent the communication of freshly allocated data between tasks. Moreover, they noticed that there exist some races that are also trivially disentangled. Races that fall into this category are: *(i)* write-write races, because a write does not acquire a location, and *(ii)* read-write races on data that was allocated before the beginning of the parallel phase, because tasks are allowed to communicate previously-allocated data.

What is the common denominator of all these cautious programs? Rather than categorically restricting reads, they demand a more nuanced consideration of writes. More precisely, these programs ensure that when a task writes a value to a location, this value is safe to read for any task that can access the said location. Race-free programs prevent concurrent reads, because a task

$$vProp \triangleq \mathcal{T} \xrightarrow{mon} iProp \qquad\qquad P \vdash_{vProp} P' \triangleq \forall t.\, P\,t \vdash_{iProp} P'\,t$$

$$P * P' \triangleq \lambda t.\, P\,t * P'\,t \qquad\qquad\qquad \lceil \Phi \rceil \triangleq \lambda\_.\, \Phi$$

$$P \twoheadrightarrow P' \triangleq \lambda t.\, \forall t'.\, t \leqslant t' \twoheadrightarrow P\,t' \twoheadrightarrow P'\,t' \qquad\qquad P@t \triangleq P\,t$$

$$\forall x.\, P \triangleq \lambda t.\, \forall x.\, P\,t \qquad\qquad\qquad \ell \odot now \triangleq \lambda t.\, \ell \odot t$$

$$\exists x.\, P \triangleq \lambda t.\, \exists x.\, P\,t \qquad\qquad \ell \mapsto_p \vec{w} \triangleq \lceil \ell \mapsto_p \vec{w} \rceil * \vec{w} \odot now$$

$$\mathsf{wpm}\; e\; \{Q\} \triangleq \lambda t.\, \forall t'.\, t \leqslant t' \twoheadrightarrow \mathsf{wp}\; \langle t', e \rangle\; \{\lambda t''\, v.\, (Q\,v)@t''\}$$

Fig. 11. DisLog+ separation logic and assertions

must have unique ownership of a location to write to it. Write-write races do not restrict writes, as there is no concurrent read, and read-writes races are permitted as long as the written value was allocated before the beginning of the parallel phase.

For this large class of programs that don't poke the bear, we provide an almost traditional separation logic called DisLog+. By "traditional", we mean that the weakest precondition of DisLog+ takes a form which does not mention timestamps (§5.2). Moreover, its syntax-directed reasoning rules do not mention clock nor precedence assertions: they are the standard reasoning rules of concurrent separation logic (§5.3). By "almost", we stress that DisLog+ restricts the use of ghost state to prevent races, but is otherwise a standard separation logic. DisLog+ is hence ideally suited to reason about race-free programs. To cater to the benign races identified above, we extend DisLog+ with *(i)* write-only assertions to reason about write-write races (§5.4) and *(ii)* three rules to reason about read-write races on previously allocated data that we refer to as the objectivity lemmas (§5.5).

5.2 Monotonicity to the Rescue

Our development of DisLog+ was triggered by two observations about race-free programs: *(i)* race-free programs ensure that, when a task accesses a location, any value referenced by the location is safe for the task to acquire, and that *(ii)* this property is *monotonic* with respect to the precedence pre-order. Indeed, if all the pointed-to values are safe to read for a given task, then these values are also safe to read for any of the task's descendants in the computation graph.

We define a new separation logic, in which every assertion is parameterized by a timestamp, called the *ambient* timestamp, and is monotonic with respect to the precedence pre-order. Fig. 11 presents the formal definitions of these assertions, written P and of type $vProp$.

The user can always *project* a $vProp$ assertion P to a particular timestamp t in $iProp$ via the construction $P@t$. Conversely, the lifting construction $\lceil \Phi \rceil$ allows to lift an $iProp$ assertion Φ into $vProp$. Hence, the whole ghost-state theory of $iProp$ is available in $vProp$. When the context allows it, we write Φ instead of $\lceil \Phi \rceil$ for the lifting of an $iProp$ assertion into $vProp$. The entailment of $vProp$ (written $\vdash_{vProp}$) is defined using the entailment of $iProp$ (written $\vdash_{iProp}$). The definition ensures that an entailment $P \vdash_{vProp} P'$ is valid if and only if, for any timestamp t the projection $P@t$ of the premise entails the projection $P'@t$ of the conclusion.

Fig. 11 also defines the assertions relative to DisLog+. The $\ell \odot now$ assertion asserts that ℓ was allocated before the ambient timestamp. This is a persistent assertion, whose monotonicity is ensured by the CLOCKMONO rule. Again, we overload this assertion to arbitrary values and collection of values. The key idea of DisLog+ is its definition of the points-to assertion. The assertion $\ell \mapsto_p \vec{w}$ is defined as the conjunction of the points-to assertion in $iProp$, written $\lceil \ell \mapsto_p v \rceil$, as well as the knowledge that every value pointed-to by the location was allocated before the ambient timestamp, using the $\vec{w} \odot now$ assertion. Hence, the points-to assertion of $vProp$ asserts that every load for this location is safe for this particular task, and any subsequent task!

The WP of DisLog+ takes the form $\mathsf{wpm}\; e\; \{\lambda v.\, P\}$, where the variable v denotes the resulting value of e and is bound in P. To abstract over the details of the postcondition, we write Q instead

ALLOC+
$$\text{wpm } (\text{alloc } n \, v) \, \{\lambda \ell. \, \ell \mapsto v^n\}$$

STORE+
$$\frac{\ulcorner 0 \leq i < |\vec{w}| \urcorner \qquad \ell \mapsto \vec{w}}{\text{wpm } (\ell[i] \leftarrow v) \, \{\lambda \_. \, \ell \mapsto [i := v]\vec{w}\}}$$

LOAD+
$$\frac{\ulcorner 0 \leq i < |\vec{w}| \, \wedge \, \vec{w}(i) = v \urcorner \qquad \ell \mapsto_p \vec{w}}{\text{wpm } (\ell[i]) \, \{\lambda v'. \, \ulcorner v' = v \urcorner \, * \, \ell \mapsto_p \vec{w}\}}$$

PAR+
$$\frac{\text{wpm } e_1 \, \{Q_1\} \qquad \text{wpm } e_2 \, \{Q_2\}}{\text{wpm } (e_1 \, || \, e_2) \, \{\lambda \ell. \, \exists v_1 \, v_2. \, \ell \mapsto [v_1; v_2] \, * \, Q_1 \, v_1 \, * \, Q_2 \, v_2\}}$$

Fig. 12. Selected rules of DisLog+

of $\lambda v. \, P$. The assertion wpm $e \, \{Q\}$ asserts that e is safe to execute at any timestamp succeeding the ambient one, and that if e reaches a value v, then $Q \, v$ holds at the end timestamp (or at any subsequent timestamp, since Q is monotonic).

The user can freely go between DisLog+ and DisLog using the following CONVERSION rule.

$$\left(P \vdash_{vProp} \text{wpm } e \, \{Q\}\right) \iff \left(\forall t. \, P@t \vdash_{iProp} \text{wp } \langle t, e \rangle \, \{\lambda t' \, v. \, (Q \, v)@t'\}\right) \quad \text{(CONVERSION)}$$

This rule needs a careful reading. It asserts that the precondition P entails wpm $e \, \{Q\}$ in $vProp$ if and only if (in the meta-logic), for any timestamp, the projection of the precondition at this timestamp entails the WP of DisLog with the postcondition projected at the end timestamp. Notice that the CONVERSION rule is an equivalence. While it is not surprising that a specification in DisLog+ is valid in DisLog (the former being more restrictive than the latter), the converse is also true: the user can use the full power of DisLog rules to verify a DisLog+ interface.

The soundness of DisLog+ is a direct corollary of the soundness of DisLog (Theorem 4.1), thanks to the CONVERSION rule.

THEOREM 5.1 (SOUNDNESS OF DISLOG+). *If wpm $e \, \{\lambda \_. \, \ulcorner True \urcorner\}$ holds, then e is safe.*

5.3 Reasoning Rules of DisLog+

We showcase the most important reasoning rules of DisLog+ in Fig. 12, which are similar to the reasoning rules of the original concurrent separation logic with fractional permissions [Bornat et al. 2005]. These rules are expressed at the $vProp$ level, where the horizontal bar stands for $vProp$ entailment. In particular, the points-to assertion occurring in the rules is the one defined in Fig. 11, guaranteeing that any load will be safe. Nevertheless, all the rules of Fig. 12 are derived from the rules of DisLog (§4.3) using the CONVERSION rule.

The rules we present in Fig. 12 prevent races. Indeed, the only way to allow a race is by sharing a points-to assertion between tasks, which is only possible via invariants [Jung et al. 2018, §2.2]. Because invariants can only be installed for assertions of type $iProp$, but points-to assertions in DisLog+ are of type $vProp$, DisLog+ rules out races by construction. We alluded to this restricted use of Iris ghost state by referring to DisLog+ as an "almost" standard separation logic (§5.1).

The ALLOC+ rule produces a valid points-to assertion, that is, both the ownership information and the proof that the default value is safe to read at the ambient timestamp. As the default value v occurs in the expression alloc $i \, v$, if v is a location, then it was already acquired, and hence already safe. This is reminiscent of the MEMENTOPRE rule (Fig. 6). The STORE+ rule is also standard and preserves the fact that any subsequent load will be safe. The stored value v occurs in the expression $\ell[i] \leftarrow v$, and is hence safe to read. The LOAD+ rule is the standard rule of separation logic, as it internally rests on the fact that all the values pointed-to by the location are safe to read. Finally, the PAR+ rule heavily makes use of the monotonicity of $vProp$ assertions. Indeed, the two postconditions Q_1 and Q_2 are valid for the end timestamp of the two forked tasks. Hence, they are also valid for the end timestamp of the parallel tuple, that succeeds them.

$$\text{WOSTART} \atop \ell \mapsto [v] \;\Rightarrow\; \exists \delta.\ \mathrm{orig}^\delta\, v * \ell \Mapsto^\delta_1 \emptyset \qquad \text{WOFRAC} \atop \ell \Mapsto^\delta_{(p_1+p_2)} (X_1 \cup X_2) \;\dashv\vdash\; \ell \Mapsto^\delta_{p_1} X_1 * \ell \Mapsto^\delta_{p_2} X_2$$

$$\text{WOSTORE} \atop \dfrac{\ell \Mapsto^\delta_p X}{\mathrm{wpm}\,(\ell[0] \leftarrow v)\,\{\lambda\_.\,\ell \Mapsto^\delta_p \{v\}\}} \qquad \text{WOCANCEL} \atop \mathrm{orig}^\delta\, v * \ell \Mapsto^\delta_1 \emptyset \;\Rightarrow\; \ell \mapsto [v] \qquad \text{WOEND} \atop \dfrac{X \neq \emptyset}{\ell \Mapsto^\delta_1 X \;\Rightarrow\; \exists v.\ \ulcorner v \in X \urcorner * \ell \mapsto [v]}$$

Fig. 13. Fractional write-only points-to assertions

5.4 Write-Write Races are Disentangled: Fractional Write-Only Assertions

A write-write race occurs when two or more tasks race to write to a shared location, but neither of them (or any other task) ever reads from the said location. Write-write races are always disentangled, as a write does not acquire any location. However, from the point of view of functional correctness, write-write races are subtle: once the tasks join and the program reads the shared location, all the outcomes of the race should be taken into account.

To verify such races in more standard Iris settings, the user typically installs an invariant containing the points-to assertion, quantifies existentially over the pointed-to value, and constrains it using ghost state. This existential quantification allows the user to change the pointed-to value while preserving the invariant. In DisLog+, invariants are restricted to *iProp* assertions, and thus the user cannot store a (*vProp*) points-to assertion inside an invariant. To allow the verification of write-write races in DisLog+ without the need of an invariant, we introduce the notion of a fractional *write-only assertion* presented in Fig. 13.

A write-only assertion takes the form $\ell \Mapsto^\delta_p X$, where δ is a list of ghost names, p a positive fraction less or equal to 1, and X a set of possible values. When $p = 1$, the set X contains all the values possibly written to ℓ. The write-only assertion comes with a companion assertion $\mathrm{orig}^\delta\, v$, which is persistent and describes the *original* value of the points-to assertion. The WOSTART rule consumes a points-to assertion and produces an orig assertion and an empty write-only assertion. The WOFRAC rule asserts that the write-only assertion is fractional: the user can always arbitrarily split and join it. The WOSTORE rule allows executing a store operation, overwriting the set of possible values. This rule only requires a fraction of the write-only assertion: a concurrent task could have another fraction and race to write.

Fig. 13 also includes two rules for getting back a points-to assertion from a write-only assertion. In both cases, the full fraction 1 must be given back. The WOCANCEL rule can be used if no write occurred, as witnessed by the empty write-only assertion. The rule produces the original points-to assertion. If at least one write has occurred, the WOEND rule can be used. The rule produces a points-to assertion and a proof that the pointed-to value is in the set of possible values.

The definition of write-only assertions appears in our Coq formalization [Moine et al. 2023b]. This definition makes use of standard ghost state and in particular a *cancellable invariant* [Jung et al. 2018, §7.1] in which a points-to assertion of DisLog is stored. Notably, the assertion $\ell \Mapsto^\delta_p X$ carries not only information on the contents of the cancellable invariant, but also a witness $X \oplus now$ that the set of possible values was allocated before the ambient timestamp. Hence, after an application of the WOEND rule, we can reconstruct back a *vProp* points-to by canceling the invariant and exhibiting a witness that the pointed-to value was allocated before the ambient timestamp.

While write-only assertions fit well in the context of *vProp*, we stress that they are not specific to it. Similar definitions can be proposed for regular points-to assertions at the *iProp* level, dropping assertions related to timestamps. In our discussion of write-only assertions, we focus on *references*, which in DisLang are represented by arrays of size 1. To target arbitrary arrays, we assume that our approach can be generalized to a more detailed interface with a per-index write-only points-to assertion. This generalization should be purely mechanical.

GETCLOCK
$$\frac{\ell \mapsto_p \vec{w}}{\ell \mapsto_p \vec{w} \;*\; \vec{w} \odot now}$$

MEMENTOPRE+
$$\frac{\ulcorner \ell \in locs(e) \urcorner \qquad \ell \odot now \;*\; \text{wpm}\, e\, \{Q\}}{\text{wpm}\, e\, \{Q\}}$$

SPLITSUBJOBJ
$$P \dashv\vdash \exists t.\, \uparrow t \;*\; \lceil P@t \rceil$$

OBJECTIVIZE
$$\frac{\ell \mapsto_p \vec{w} \qquad \lceil locs(\vec{w}) \odot t \rceil}{\lceil (\ell \mapsto_p \vec{w})@t \rceil}$$

Fig. 14. Objectivity lemmas

5.5 Many Read-Write Races are Disentangled: Objectivity Lemmas

Read-write races can create entanglement: a task could communicate a memory location it allocated to a concurrent task. Verifying that such races are safe often requires the full expressiveness of Dis-Log. However, some read-write races are trivially disentangled: if written values are unboxed (that is, not allocated in the heap), or were allocated before the beginning of the race. This section explains how to reason about such races within DisLog+.

In the extreme case where a location points to a block of unboxed values, read-write races are tolerated without additional work. Recall (Fig. 11) that $\ell \mapsto_p \vec{w} \triangleq \lceil \ell \mapsto_p \vec{w} \rceil \;*\; \vec{w} \odot now$. If $\vec{w}$ contains only unboxed values, the assertion $\vec{w} \odot now$ holds trivially true, and we thus have that $\ell \mapsto_p \vec{w} \dashv\vdash \lceil \ell \mapsto_p \vec{w} \rceil$. Therefore, the points-to assertion of ℓ can be stored directly inside an invariant, and read-write races on ℓ can be verified, as long writes concern unboxed values.

In a more general case, Fig. 14 presents an interface of "objectivity lemmas" to reason about trivially disentangled read-write races. Objectivity lemmas offer two mechanisms. First, they allow witnessing that a set of locations were allocated before a program point (rules GETCLOCK and MEMENTOPRE+). Second, they allow sharing a points-to assertion through an invariant, and allow updating this points-to assertion as long as new values are unboxed or were allocated before the installation of the invariant (rules SPLITSUBJOBJ and OBJECTIVIZE).

To witness that a set of locations were allocated before a given program point, the user can use two rules and combine their result using the equivalence $(A_1 \cup A_2) \odot now \dashv\vdash A_1 \odot now * A_2 \odot now$. First, the GETCLOCK rule allows extracting from a points-to assertion that every pointed-by value was allocated before the ambient timestamp. Second, the MEMENTOPRE+ rule asserts that every location occurring in the current expression was allocated before the ambient timestamp.

The SPLITSUBJOBJ rule allows in particular sharing a points-to assertion through an invariant, while fixing the set of possibilities for newly written values. We present the general form of the rule, adapted from Cosmo [Mével et al. 2020, §4.1]. The SPLITSUBJOBJ rule asserts that owning a *vProp* assertion P is equivalent to owning its *iProp* projection to some timestamp $\lceil P@t \rceil$ (the *objective* part of P), and the information that this timestamp precedes the ambient timestamp, written $\uparrow t$ (the *subjective* part of P). The assertion $\uparrow t$ is persistent and defined as $\lambda t'.\, t \le t'$. The objective part $P@t$ is an *iProp* assertion and can hence be shared through an invariant. The subjective part $\uparrow t$ cannot be shared through an invariant, but it can be given to subsequent tasks and used to convert back the objective part into the original *vProp* assertion, by using again the SPLITSUBJOBJ rule.

We now focus on a contrived example to illustrate how the objectivity lemmas are intended to be used, and why an additional rule, OBJECTIVIZE, is needed. Sec. 6.1 and 6.3 provide more realistic examples, based on the same idea. Our contrived example is:

$$\text{let } \ell = \text{alloc } 1\,()\text{ in } (\ell[0] \,||\, \ell[0] \leftarrow \ell)$$

This example allocates a reference ℓ (that is, an array of size 1), initialized to (). It then forks two tasks, one reading from ℓ and the other storing ℓ inside itself. This is a trivially disentangled read-write race: the only written value is allocated before the beginning of the race.

To verify this example, we first use the ALLOC+ rule (Fig. 12) and obtain the assertion $\ell \mapsto [()]$. Then, we use the MEMENTOPRE+ rule to generate a witness $\ell \odot now$ that ℓ was allocated before the execution of the par. Then, we use the SPLITSUBJOBJ rule on the assertion $(\ell \mapsto [()] * \ell \odot now)$. We obtain the assertion $\uparrow t * \lceil (\ell \mapsto [()] * \ell \odot now)@t \rceil$. Unfolding definitions, the objective part is equivalent to $(\ell \mapsto [()])@t * \ell \odot t$. We then allocate an invariant containing the assertion:

$$\exists v.\ (\ell \mapsto [v])@t\ *\ \lceil v = ()\ \vee\ v = \ell \rceil$$

Next, we apply the PAR+ rule (Fig. 12), and give each task a copy of the invariant, of the assertion $\uparrow t$ and of the assertion $\ell \odot t$. For the two tasks, the proof is similar. First, we open the invariant. Then, use the SPLITSUBJOBJ rule to convert back the assertion $(\ell \mapsto [v])@t$ into $\ell \mapsto [v]$, instantiating the existential with t. We then execute the load or the store.

Finally, comes the time to close the invariant. We cannot use the SPLITSUBJOBJ rule again because the existential quantification of this rule would generate a fresh timestamp t', distinct from t. Thankfully, the OBJECTIVIZE rule comes to save the day. This rule allows us to generate an assertion $\lceil (\ell \mapsto \vec{w})@t \rceil$ as long as we can provide the assertion $\lceil locs(\vec{w}) \odot t \rceil$, where $locs(\vec{w})$ denotes the set of locations of the list of values $\vec{w}$. We apply the OBJECTIVIZE rule, instantiating $\vec{w}$ by $[v]$, do a case analysis on v and end the proof, given that we have at hand the assertion $\ell \odot t$. Notice that, instead of writing ℓ into itself, we could have written any unboxed value, or any location that was allocated before ℓ itself.

6 EVALUATION

We showcase DisLog+ and DisLog via a range of case studies. We first focus on the scratch example of Sec. 2, and prove it correct in DisLog+ (§6.1). Then, we illustrate how to reason about a write-write race using write-only assertions (§6.2) based on a parallel lookup in a lazy collection. We conclude with a case study on concurrent hashing and deduplication (§6.3 and §6.3).

Deduplication refers to the process of removing duplicates from a collection. This task can be efficiently done in parallel using concurrent hashing: each task tries to insert elements into a shared concurrent hash set, which by construction does not store duplicates. If the collection is fully allocated beforehand, we use a folklore hash set [VerifyThis 2022], and verify both the hash set interface and the duplication itself entirely within DisLog+ (§6.3), thanks to the objectivity lemmas. In the case of a lazy collection, where elements may not be already allocated prior to the parallel phase, the previous approach cannot be directly used: naïvely applying the previous deduplication algorithm would result in entanglement, and so a different deduplication algorithm is needed. We address this issue by first partially removing duplicates in parallel with a more subtle hash set, then getting rid of the remaining duplicates by calling the previous deduplication function. Interestingly, the proof requires the full power of DisLog (§6.4).

In the case studies, we write a non-recursive function as $\lambda \vec{x}.e$, which is a sugar for $\mu\_.\lambda \vec{x}.e$. where _ denotes an anonymous binding. We add a hat and write $\hat{\lambda}$ to distinguish top-level functions. We write $e_1\ ;\ e_2$ for a sequence, which is encoded as let $\_ = e_1$ in e_2.

6.1 The scratch Example

We first give an interface to a spin-lock. The verified code is a direct translation of the spin-lock presented earlier (§2.2), which is implemented as a pair of closures sharing a reference to a boolean named r, initialized to false. The first closure attempts a lock by doing a CAS on r from false to true. The second closure releases the lock by setting the reference r to false.

Our specification of locks in DisLog+ appears in Fig. 15 and is very similar to the standard specification of locks in high-order separation logic [Gotsman et al. 2007; Svendsen and Birkedal 2014]. In DisLog+, a lock must be restricted to protect an *iProp* assertion Φ, as a lock protects

$$\frac{\ulcorner \Phi \urcorner}{\mathsf{wpm}\,(\mathsf{new\_lock}\,[])\,\{\lambda\,\ell.\,\exists l\,u.\,\ell \mapsto [l;u] * \ulcorner \mathsf{lock}\,l\,u\,\Phi \urcorner\}}$$

$$\frac{\ulcorner \mathsf{lock}\,l\,u\,\Phi \urcorner}{\mathsf{wpm}\,(l\,[])\,\{\lambda\,b.\,\ulcorner b = \mathsf{true} \urcorner \twoheadrightarrow (\ulcorner \mathsf{locked}\,l \urcorner * \ulcorner \Phi \urcorner)\}} \qquad \frac{\ulcorner \mathsf{lock}\,l\,u\,\Phi \urcorner \quad \ulcorner \mathsf{locked}\,l \urcorner \quad \ulcorner \Phi \urcorner}{\mathsf{wpm}\,(u\,[])\,\{\lambda\,\_.\,\ulcorner True \urcorner\}}$$

Fig. 15. Case study: specification of a spin-lock

an invariant. The precondition of the specification of new_lock consumes Φ. The postcondition produces a location ℓ pointing to a pair of closures (l, u), for locking and unlocking, respectively, as well as an *iProp* assertion lock $l\,u\,\Phi$, which is persistent and asserts that the l and u describe a valid lock. The specification of a call to the closure l requires a valid lock, and returns a boolean. If this boolean is true, then the lock was successfully locked: the user gains the protected assertion Φ as well an exclusive token locked l, witnessing that the lock is now locked. This token is required to call the closure u, as well as Φ, which has to be given back when the user wants to unlock the lock.

We conduct the proofs of the interface of Fig. 15 entirely within DisLog+. Indeed, we are in an extreme case: the shared reference points-to a boolean, which is unboxed. Hence, the points-to of DisLog+ is equivalent to the points-to of DisLog (§5.5). Thus, we are able to define the lock $l\,u\,\Phi$ assertion using an invariant that stores directly the points-to assertion of the shared reference.

We make use of the interface of locks we just presented, as well as the objectivity lemmas, and verify the following interface for the scratch example.

$$\frac{\mathbin{\text{\Large$\ast$}}_{i \in \{0,1\}} \quad \forall \ell.\, \ell \mapsto \mathit{defaultElem}^{\mathsf{N}} \twoheadrightarrow \mathsf{wpm}\,(\mathsf{doWork}\,[\ell])\,\{\lambda\,\_.\,Q_i \;\ast\; \exists \vec{w}.\, \ell \mapsto \vec{w}\}}{\mathsf{wpm}\,(\mathsf{scratch}\,[])\,\{\lambda\,\_.\,Q_0 \;\ast\; Q_1\}}$$

Recall that the scratch example allocates a *shared* reference, which is protected by a lock while two parallel tasks attempt to access it. After allocating the *shared* reference, we have a points-to assertion of the form *shared* $\mapsto \mathit{defaultElem}^{\mathsf{N}}$. We use the GETCLOCK rule to extract an assertion $\mathit{defaultElem} \odot \mathit{now}$. Then, we use the SPLITSUBJOBJ rule and obtain the assertion $\uparrow t$ as well as the assertion that will be protected by the lock $(\mathit{shared} \mapsto \mathit{defaultElem}^{\mathsf{N}})@t$. Each task gets an assertion $\uparrow t * \mathit{defaultElem} \odot t$. Then, if a task wins the lock, we reconstruct a points-to assertion with the SPLITSUBJOBJ rule, call the doWork function, call the cleanScratchPad function, and restore the invariant protected by the lock using the OBJECTIVIZE rule.

6.2 Parallel Lookup in a Lazy Collection

The left part of Fig. 16 presents the code of a *parallel loop* parfor $[a; b; h]$, calling the function h for each index between a and b. The presented code is a direct translation of the implementation used in the standard library of MPL [MPL Development Team 2022]. The specification of parfor appears below and should be unsurprising. The precondition requires, for every index i between a and b, that $h\,[i]$ is safe and satisfies a postcondition $Q\,i$. The postcondition of parfor $[a; b; h]$ produces the iterated conjunction of the postconditions.

The right part of Fig. 16 presents the code of the lookup $[k; n]$ function that searches for a non-unit value in the lazy collection k up to index n. To do so, the function uses a reference r, and a closure h that takes an index i, produces the i-th index of the lazy collection, and writes it in r if it is non-unit. The closure h is then called in parallel for every index between 0 and n. This is a typical example of a write-write race: each call on h may write in r, but never read from it.

The specification of lookup $[k; n]$ appears at the bottom of Fig. 16. Its precondition requires that k is valid lazy collection: between indices 0 and n, k produces a value satisfying a predicate K. The postcondition of the specification produces a value v and asserts the existence of $\vec{w}$, the collection

$$parfor \triangleq \hat{\mu}\, f.\lambda[a; b; h].$$
$$\quad\text{if } (b - a) == 0 \text{ then } ()$$
$$\quad\text{else if } (b - a) == 1 \text{ then } h\,[a]$$
$$\quad\text{else let } mid = a + ((b - a)/2) \text{ in}$$
$$\quad\quad (f\,[a; mid; h]) \,||\, (f\,[mid; b; h])$$

$$lookup \triangleq \hat{\lambda}[k; n].$$
$$\quad\text{let } r = \text{alloc } 1\,() \text{ in}$$
$$\quad\text{let } h = \lambda[i].\ \text{let } x = k\,[i] \text{ in}$$
$$\quad\quad\text{if } x == () \text{ then } () \text{ else } r[0] \leftarrow x \text{ in}$$
$$\quad parfor\,[0; n; h]\ ;\ r[0]$$

$$\frac{\bigast_{i\in[a;b)}\ \mathsf{wpm}\,(h\,[i])\,\{\lambda\_.\,Q\,i\}}{\mathsf{wpm}\,(parfor\,[a; b; h])\,\{\lambda\_.\ \bigast_{i\in[a;b)}\,(Q\,i)\}} \qquad \frac{\bigast_{i\in[0;n)}\ \mathsf{wpm}\,(k\,[i])\,\{K\}}{\mathsf{wpm}\,(lookup\,[k; n])\,\{\lambda v.\ \exists \vec{w}.\ \ulcorner \mathsf{found}\ n\ \vec{w}\ v \urcorner\ *\ \bigast_{v\in\vec{w}}\,(K\,v)\}}$$

Fig. 16. Case study: parallel lookup in a lazy collection

itself. The $\mathsf{found}\ n\ \vec{w}\ v$ judgment asserts that $\vec{w}$ has size n and that either v is not the unit value and occurs in $\vec{w}$, or v is the unit value and every value in $\vec{w}$ is the unit value. It is defined as:

$$\mathsf{found}\ n\ \vec{w}\ v \triangleq |\vec{w}| = n\ \wedge\ (v \neq ()\ \wedge\ v \in \vec{w})\ \vee\ (v = ()\ \wedge\ \forall w.\ w \in \vec{w} \implies w = ())$$

The proof of our specification makes use of a write-only assertion (§5.4). Indeed, just after the allocation of r, we convert its points-to assertion into a write-only assertion $r \mapsto_1^\delta \emptyset$, and split it into n fractions. Then, we call the parfor specification and instantiates Q with:

$$\lambda\_.\ \exists v.\ K\,v\ *\ ((\ulcorner v \neq ()\urcorner\ *\ r \mapsto_{1/n}^\delta \{v\})\ \vee\ (\ulcorner v = ()\urcorner\ *\ r \mapsto_{1/n}^\delta \emptyset))$$

This postcondition asserts that the lazy collection produced a value and that if it is not the unit value, then it was written into r, else nothing was written to r. After applying the specification of parfor, we gather all the fractions of the write-only assertion. We then do a case analysis on whether a non-unit value is in the lazy collection k, and convert the write-only assertion back to a normal points-to assertion accordingly.

6.3 Deduplication via Concurrent Hashing

In the next two sections, we suppose a user-chosen capacity C, which bounds the number of elements within hash sets. We also suppose a hash function from values to integers.

We present a folklore [VerifyThis 2022] concurrent, lock-free, fixed-capacity hash set using *open addressing* and *linear probing* to handle collision [Knuth 1998]. The code of our hash set appears in the upper part of Fig. 17. The hash set consists of an array of size C. When created, the array is filled with a dummy element d, which cannot be inserted in the hash set as it denotes an empty slot. Inserting an element is done by the add $[s; d; x]$ function, where s is the hash set, d the dummy element and x the element being inserted. The function calls the *put* auxiliary closure, which tries to insert x at a given index, originally the hash of x, using a CAS. If this index is already taken by a distinct value, potentially due to a collision, the *put* closure tries the next index. (We do not resize: the *put* function loops if the table is full.) The function elems $[s; d]$ function returns the elements of s: that is, all the elements distinct from the dummy element d. Occurrences of d are removed via a call to a dedicated filter_compact function, which filters the array, and returns a compacted array where d does not appear anymore. For brevity here, we omit the implementation of filter_compact, as it can be implemented entirely race-free and therefore disentangled. We instead focus on the nuance of disentanglement for concurrent operations on the hash set.

The hash set can be used to insert values concurrently and in parallel. However, in order to preserve disentanglement, the user should only insert values that were allocated before the beginning of the parallel phase [Westrick 2022]. Indeed, the *put* auxiliary function does a CAS operation on an *a priori* arbitrary index, which may have been filled by a concurrent task. Our interface hence restricts insertions to a set of values that were allocated before the hash set itself.

$$\text{init} \triangleq \hat{\lambda}[d].\ \text{alloc}\ C\ d \qquad\qquad \text{elems} \triangleq \hat{\lambda}[s;d].\ \text{filter\_compact}\ [s;d]$$

$$\text{add} \triangleq \hat{\lambda}[s;d;x]. \qquad\qquad\quad \text{dedup} \triangleq \hat{\lambda}[d;\ell].$$
$$\quad \text{let}\ put\ =\ \mu f.\lambda[i]. \qquad\qquad\qquad \text{let}\ s\ =\ \text{init}\ [d]\ \text{in}$$
$$\qquad \text{if}\ (\text{CAS}\ \ell\ i\ d\ x\ \vee\ \ell[i] == x)\ \text{then}\ () \qquad \text{let}\ h\ =\ \lambda[i].\ \text{add}\ [s;d;\ell[i]]\ \text{in}$$
$$\qquad \text{else}\ f\ [(i+1)\ \text{mod}\ C]\ \text{in} \qquad\qquad \text{parfor}\ [0;\text{length}\ \ell;h]\ ;$$
$$\quad put\ [\text{hash}\ [x]\ \text{mod}\ C] \qquad\qquad\qquad \text{elems}\ [s;d]$$

$$\frac{\ulcorner d \notin A \urcorner \qquad A \oplus now}{\text{wpm}\ (\text{init}\ [d])\ \{\lambda s.\ \text{hset}\ s\ d\ A\ \emptyset\ 1\}} \qquad \frac{\ulcorner x \in A \urcorner \qquad \text{hset}\ s\ d\ A\ X\ p}{\text{wpm}\ (\text{add}\ [s;d;x])\ \{\lambda\_.\ \text{hset}\ s\ d\ A\ (X \cup \{x\})\ p\}}$$

$$\frac{\text{hset}\ s\ d\ A\ X\ 1}{\text{wpm}\ (\text{elems}\ [s;d])\ \{\lambda\ell.\ \exists\vec{w}.\ \ell \mapsto \vec{w}\ *\ \ulcorner \text{deduped}\ \vec{w}\ X \urcorner\}}$$

$$\frac{\ulcorner d \notin \vec{v} \urcorner \qquad \ell \mapsto_p \vec{v}}{\text{wpm}\ (\text{dedup}\ [d;\ell])\ \{\lambda\ell'.\ \exists\vec{w}.\ \ell' \mapsto \vec{w}\ *\ \ell \mapsto_p \vec{v}\ *\ \ulcorner \text{deduped}\ \vec{w}\ \vec{v} \urcorner\}}$$

Fig. 17. Case study: deduplication of an array by concurrent hashing

The representation predicate of a hash set s is written hset $s\ d\ A\ X\ p$, where d is the dummy element, A a set of values that were witnessed as allocated before the hash set, X a set of values that were inserted, and p a fraction in $(0;1]$. This predicate can be split and joined, allowing for parallel use.

$$\text{hset}\ s\ d\ A\ (X_1 \cup X_2)\ (p_1 + p_2)\quad \dashv\vdash\quad \text{hset}\ s\ d\ A\ X_1\ p_1\ *\ \text{hset}\ s\ d\ A\ X_2\ p_2$$

Such a predicate is created by the init $[d]$ function. Its precondition requires a witness that a set A of values were allocated before the current timestamp, via the $A \oplus now$ assertion. Such an assertion can be obtained via the GetClock and SplitSubjObj rules. The precondition also requires that the dummy element d is not an element of A. The postcondition produces a valid empty hash set with fraction 1. The specification of add $[s;d;x]$ requires a valid hash set with an arbitrary fraction, and that the element being inserted is in the authorized set of values. The specification of elems $[s;d]$ consumes a hash set with fraction 1 with content X and produces an array ℓ with content $\vec{w}$. The deduped $\vec{w}\ X$ assertion asserts that $\vec{w}$ contains no duplicate and has the same elements as X:

$$\text{deduped}\ \vec{w}\ X\quad \triangleq\quad \text{NoDup}\ \vec{w}\ \wedge\ (\forall v.\ v \in \vec{w} \iff v \in X)$$

The proofs of the hash set interface of Fig. 17 rest on the objectivity lemmas. Indeed, parallel insertion may imply read-write races, but only on data allocated before the parallel phase. Intuitively, the hset predicate involves a cancellable invariant, storing an assertion $(s \mapsto \vec{w})@t$, as well as an assertion $A \oplus t$, where t comes from an application of the SplitSubjObj rule. The invariant also records that $locs(\vec{w}) \subseteq A$, allowing the use of the Objectivize rule during the proofs. In addition to this cancellable invariant, the hset predicate also involves the (persistent) assertion $\uparrow t$, allowing to retrieve and update the points-to assertion of s using the SplitSubjObj rule.

Fig. 17 also presents the code and specification of the dedup $[d;\ell]$ function. This function deduplicates the array ℓ using our concurrent hash set. The function first creates a hash set s. Then, the function allocates a closure which, given an index i, inserts the element $\ell[i]$ inside s. Next, the function calls the closure in parallel for every index of the array. Finally, the function returns the elements of the hash set. The precondition requires that ℓ points to an array $\vec{v}$ and the existence of a dummy element d that is not in the array. The postcondition returns a fresh location ℓ' pointing to an array $\vec{w}$ that is a deduplicated version of $\vec{v}$. Making use of our hash set, the proof is straightforward: each task gets an assertion hset $s\ d\ \vec{v}\ \emptyset\ (1/|\vec{v}|)$, enabling them to

$$
\begin{aligned}
\text{init} \;&\triangleq\; \hat{\lambda}[d]. \\
&\quad \text{pair } [\text{alloc } C \, d; \text{alloc } C \, (-1)] \\[4pt]
\text{add} \;&\triangleq\; \hat{\lambda}[s; d; x; u]. \\
&\quad \text{let } put \;=\; \mu f. \lambda[i]. \\
&\qquad \text{if tryput } [s; d; x; u; i] \text{ then } () \\
&\qquad \text{else } f \, [(i+1) \bmod C] \text{ in} \\
&\quad put \, [\text{hash } [x] \bmod C]
\end{aligned}
\qquad
\begin{aligned}
\text{tryput} \;&\triangleq\; \hat{\lambda}[s; d; x; u; i]. \\
&\quad \text{let } l_1 \;=\; s[0] \text{ in let } l_2 \;=\; s[1] \text{ in} \\
&\quad \text{if } l_2[i] == u \vee \text{CAS } l_2 \, i \, (-1) \, u \text{ then} \\
&\qquad \text{let } e \;=\; l_1[i] \text{ in} \\
&\qquad \text{if } e == d \text{ then } l_1[i] \leftarrow x \,;\, \text{true} \\
&\qquad \text{else } e == x \\
&\quad \text{else false}
\end{aligned}
$$

$$
\frac{\ulcorner n \neq 0 \urcorner}{\text{wp } \langle t, \text{ init } [d] \rangle \, \{\lambda t' \, s. \; \text{\Large\textasteriskcentered}_{u \in [0;n)} \, (\text{lhset } s \, d \, u \, \emptyset \, t')\}}
\qquad
\frac{\ulcorner x \neq d \urcorner \qquad \text{lhset } s \, d \, u \, X \, t}{\text{wp } \langle t, \text{ add } [s; d; x; u] \rangle \, \{\lambda t' \, \_. \; \text{lhset } s \, d \, u \, (X \cup \{x\}) \, t'\}}
$$

$$
\frac{\text{\Large\textasteriskcentered}_{i \in [0;n)} \, \text{wpm } (k \, [i]) \, \{\lambda v. \, \ulcorner v \neq d \urcorner \, * \, K \, v\}}{\text{wpm } (\text{dedup\_lazy } [d; k; n]) \, \{\lambda \ell. \, \exists \vec{v} \, \vec{w}. \, \ulcorner |\vec{v}| = n \wedge \text{deduped } \vec{w} \, \vec{v} \urcorner \, * \, \ell \mapsto \vec{w} \, * \, \text{\Large\textasteriskcentered}_{v \in \vec{v}} \, (K \, v)\}}
$$

Fig. 18. Case study: deduplication of a lazy collection by concurrent hashing

insert their element. At the end, the fractions of the hset predicate are joined, and the specification of elems concludes the proof.

6.4 Deduplication of a Lazy Collection via Concurrent Hashing

We cannot reuse the hash set of the previous section to deduplicate a lazy collection in parallel: its elements might not be allocated before the parallel phase. To address this issue, we implement and verify a more subtle hash set that can store elements allocated by concurrent tasks, while having a small number of duplicates. After the parallel phase, we use the previous dedup function to get rid of the remaining duplicates.

The hash set consists of a pair of arrays of the same size C, and takes inspiration from the *lock striping* technique [Herlihy and Shavit 2012]. The first array is similar to the hash set of the previous section and contains the inserted elements. The second array stores *task identifiers*, represented as (unboxed) integers. Each task is given a distinct identifier. Intuitively, before loading or writing an index of the first array, a task must ensure with a CAS that its identifier is written inside the second array at the same index. Otherwise, the task is not allowed to load or write the desired index in the first array. The races on the second array are disentangled: it stores only unboxed integers. This design has two consequences. First, duplicates may occur in the hash set: two tasks could insert the same element at distinct indexes. The number of tasks bounds the number of duplicates. Second, disentanglement of loads on the first array relies on a subtle invariant: if a task has written its identifier at an index of the second array, this task uniquely owns the same index of the first array.

Fig. 18 presents the implementation of our new hash set. The init $[d]$ function initializes the first array with the dummy element d and the second array with a dummy identifier -1. The add $[s; d; x; u]$ function inserts the element x inside the hash set s with dummy element d and the task identifier u. This function differs from the add function of the previous section in two points. First, it is parameterized by the identifier of the task inserting the element. Second, instead of inserting an element at index i using a direct CAS, it uses the auxiliary function tryput.

The tryput function first loads the first array l_1 and the second array l_2 of the hash set. Then, the function tests if it has already written its identifier at index i inside l_2. If yes, or if a CAS succeeds in writing the identifier u, the task has unique ownership of the index i in l_1. The task can then load the content of $l_1[i]$, and tries to insert the desired element.

To prove this data structure correct, the full power of DisLog is needed. The lower part of Fig. 18 shows the specifications of init and add, where n is the number of tasks that will use the hash set. They involve a representation predicate lhset $s\ d\ u\ X\ t$, which asserts that the hash set s with dummy element d can be used with the task identifier u, contains elements in X and is valid at timestamp t. The specification of init generates n pieces of the lhset predicate, one for each task. The add function must be called with the correct identifier. The key idea is that the lhset representation predicate is monotonic with respect to the precedence pre-order. We derive similar specifications in DisLog+ using the CONVERSION rule, confining the timestamp-related reasoning.

Given a number of tasks P, our deduplication function dedup_lazy $[d; k; n]$ allocates a hash set and generates P parallel tasks using parfor. For each element of its designated range, each task sequentially generates the element on-the-fly before trying to insert it inside the lazy collection. When all tasks end, the dedup_lazy function extracts the elements of the hash set, and remove the remaining duplicates using the previous dedup function.

The last part of Fig. 18 presents the specification of our deduplication function dedup_lazy $[d; k; n]$. The precondition requires that the lazy collection k is safe between index 0 and n, that it returns a value that is not the dummy element, and that satisfies a given postcondition K. A call to dedup_lazy returns a location ℓ and guarantees the existence of $\vec{v}$ and $\vec{w}$ such that $\vec{v}$ contains n elements and $\vec{w}$ is a deduplicated version of $\vec{w}$. The postcondition then asserts that the returned location ℓ points to $\vec{w}$, and that for every value v in $\vec{v}$, the assertion $K\ v$ holds.

7 MECHANIZATION

All our results are mechanized in the Coq proof assistant [Moine et al. 2023b] using Iris [Jung et al. 2018] and its dedicated Proof Mode [Krebbers et al. 2018]. Rounding and excluding comments, the definition of the language takes 1200LOC, the proofs of the two logics and their soundness theorems, 4600LOC, and the verification of case studies 3700LOC. We provide tactics to be used while reasoning with the two logics, and automation to DisLog+ thanks to the Diaframe library [Mulder et al. 2022].

8 RELATED WORK

Disentanglement. There has been a variety of work on disentanglement [Arora et al. 2021, 2023; Guatto et al. 2018; Raghunathan et al. 2016; Westrick 2022; Westrick et al. 2022, 2020]. This work focuses on dynamic techniques that exploit disentanglement for improved efficiency, especially for parallel memory management. In particular, Arora et al. [2021] developed a provably efficient memory manager for functional programs based on disentanglement, and Arora et al. [2023] extended this approach to support unrestricted effects by accounting for the cost of entanglement. These works rely on disentanglement for efficiency and scalability, and leave the task of reasoning about disentanglement to the programmer. The first formal definition for disentanglement was given by Westrick et al. [2020] using traces of memory operations, and Westrick et al. [2022] developed a semantics which detects entanglement during execution. Our semantics for disentanglement is similar in the sense that it becomes stuck when entanglement occurs. In this context, the logics developed in this paper statically verify that execution never becomes stuck.

Linearity and Concurrency. Our "plain vanilla" DisLog+ (i.e., without fractional write-only assertions and objectivity lemmas) is related to reasoning approaches establishing race freedom by a linear treatment of resources. These approaches comprise type systems for the π-calculus [Igarashi and Kobayashi 2001, 2004] as well as session type systems [Balzer and Pfenning 2017; Caires et al. 2016; Jacobs et al. 2022; Lindley and Morris 2015; Toninho et al. 2013]. The latter are based on a Curry-Howard correspondence established between linear logic and the session-typed

π-calculus [Caires and Pfenning 2010; Wadler 2012]. Most closely related to our work in terms of employed techniques is the work by Jacobs et al. [2022], which mechanizes safety of a session-typed language in Coq, where safety encompasses freedom of memory leaks and deadlocks. The authors introduce the notion of a connectivity graph, which is acyclic by construction due to linearity, and use concurrent separation logic to prove acyclicity-preservation graph transformations. Our work, in contrast, is not confined to a linear setting.

Separation Logics. Multiple Iris-based concurrent separation logic were developed. Among them, logics targeting weak-memory models inspired DisLog+, namely iGPS [Kaiser et al. 2017], iRC11 [Dang et al. 2020] and Cosmo [Mével et al. 2020]. Indeed, they all build a high-level logic on top of a low-level logic using monotonicity arguments. In their case, assertions are monotonic predicates over the view of the memory: assertions remain valid even after observing additional memory events. In their case, the view ordering of the memory is a pure assertion. We generalize their approach to a pre-order within *iProp*. Moine et al. [2023a] present a separation logic to reason about heap space for a sequential language with garbage collection. Their language is similar to the sequential subset of DisLang. In particular, they also make the difference between top-level functions and heap-allocated closures. Disentanglement is closely related to garbage collection: disentanglement ensures in particular that locations occurring in the program are always safe to read for a task-local garbage collector: this is reminiscent of the free variable rule [Felleisen and Hieb 1992]. Outside the Iris world, Fu et al. [2010] present a concurrent separation logic with temporal reasoning. Contrary to them, our notion of time only relates to the pre-order induced by the fork-join structure of the program.

9   CONCLUSION AND FUTURE WORK

Disentanglement is an important property for parallel performance, but prior work leaves the challenge of reasoning about disentanglement to the programmer. We address this challenge by presenting DisLog, the first program logic to formally verify that a program is disentangled. Additionally, we present DisLog+, which allows for mostly standard separation logic proofs and offers proofs of disentanglement "for free" for many programs. Using these logics, we prove disentanglement for a number of examples, including several lock-free data structures. Our experience with DisLog and DisLog+ is that the effort required to prove disentanglement is often small and can be confined to the daring parts of the program. In future work, we plan to develop a type system to automatically infer disentanglement where possible. We hope that a semantic type soundness approach making use of DisLog could be used to prove such a system sound.

DATA AVAILABILITY STATEMENT

The DisLog and DisLog+ logics, their soundness proofs and all our case studies are mechanized (§7). This mechanization is recorded in an artifact available on Zenodo [Moine et al. 2023b].

ACKNOWLEDGMENTS

REFERENCES

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

Sarita V. Adve. 2010. Data races are evil with no exceptions: technical perspective. *Commun. ACM* 53, 11 (2010), 84.

Andrew W. Appel. 1992. *Compiling with Continuations*. Cambridge University Press. http://www.cambridge.org/9780521033114

Jatin Arora, Sam Westrick, and Umut A. Acar. 2021. Provably Space-Efficient Parallel Functional Programming. *Proc. ACM Program. Lang.* 5, POPL, Article 18 (jan 2021), 33 pages. 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, Article 170 (jun 2023), 26 pages. https://doi.org/10.1145/3591284

Stephanie Balzer and Frank Pfenning. 2017. Manifest Sharing with Session Types. *Proceedings of the ACM on Programming Languages* 1, ICFP (2017), 37:1–37:29. https://doi.org/10.1145/3110281

Hans-Juergen Boehm. 2011. How to Miscompile Programs with "Benign" Data Races. In *3rd USENIX Workshop on Hot Topics in Parallelism, HotPar'11, Berkeley, CA, USA, May 26-27, 2011*.

Richard Bornat, Cristiano Calcagno, Peter O'Hearn, and Matthew Parkinson. 2005. Permission accounting in separation logic. In *Principles of Programming Languages (POPL)*. 259–270. http://www.cs.ucl.ac.uk/staff/p.ohearn/papers/permissions_paper.pdf

John Boyland. 2003. Checking Interference with Fractional Permissions. In *Static Analysis Symposium (SAS) (Lecture Notes in Computer Science, Vol. 2694)*. Springer, 55–72. http://www.cs.uwm.edu/~boyland/papers/permissions.pdf

Stephen Brookes. 2007. A semantics for concurrent separation logic. *Theoretical Computer Science* 375, 1–3 (2007), 227–270. https://doi.org/10.1016/j.tcs.2006.12.034

Luís Caires and Frank Pfenning. 2010. Session Types as Intuitionistic Linear Propositions. In *21th International Conference onf Concurrency Theory (CONCUR) (Lecture Notes in Computer Science, Vol. 6269)*. Springer, 222–236. https://doi.org/10.1007/978-3-642-15375-4_16

Luís Caires, Frank Pfenning, and Bernardo Toninho. 2016. Linear Logic Propositions as Session Types. *Mathematical Structures in Computer Science* 26, 3 (2016), 367–423. https://doi.org/10.1017/S0960129514000218

Hoang-Hai Dang, Jacques-Henri Jourdan, Jan-Oliver Kaiser, and Derek Dreyer. 2020. RustBelt meets relaxed memory. *Proceedings of the ACM on Programming Languages* 4, POPL (2020), 34:1–34:29. https://hal.inria.fr/hal-02351793/

Stephen Dolan, K. C. Sivaramakrishnan, and Anil Madhavapeddy. 2018. Bounding data races in space and time. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*. 242–255.

Matthias Felleisen and Robert Hieb. 1992. The Revised Report on the Syntactic Theories of Sequential Control and State. *Theoretical Computer Science* 103, 2 (1992), 235–271. https://www2.ccs.neu.edu/racket/pubs/tcs92-fh.pdf

Mingdong Feng and Charles E. Leiserson. 1999. Efficient Detection of Determinacy Races in Cilk Programs. *Theory of Computing Systems* 32, 3 (1999), 301–326. https://doi.org/10.1007/s002240000120

Ming Fu, Yong Li, Xinyu Feng, Zhong Shao, and Yu Zhang. 2010. Reasoning about Optimistic Concurrency Using a Program Logic for History. In *CONCUR 2010 - Concurrency Theory*, Paul Gastin and François Laroussinie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 388–402.

Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. 2007. Local Reasoning for Storable Locks and Threads. In *Asian Symposium on Programming Languages and Systems (APLAS) (Lecture Notes in Computer Science, Vol. 4807)*. Springer, 19–37. http://dx.doi.org/10.1007/978-3-540-76637-7_3

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. https://doi.org/10.1145/3178487.3178494

Maurice Herlihy and Nir Shavit. 2012. *The Art of Multiprocessor Programming, Revised Reprint* (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Atsushi Igarashi and Naoki Kobayashi. 2001. A Generic Type System for the Pi-calculus. In *8th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 128–141. https://doi.org/10.1145/360204.360215

Atsushi Igarashi and Naoki Kobayashi. 2004. A Generic Type System for the Pi-calculus. *Theoretical Computer Science* 311, 1-3 (2004), 121–163. https://doi.org/10.1016/S0304-3975(03)00325-6

Iris Development Team. 2023. `iris.base_logic.lib.gen_heap`. https://plv.mpi-sws.org/coqdoc/iris/iris.base_logic.lib.gen_heap.html.

Jules Jacobs, Stephanie Balzer, and Robbert Krebbers. 2022. Connectivity Graphs: a Method for Proving Deadlock Freedom Based on Separation Logic. *Proceedings of the ACM on Programming Languages* 6, POPL (2022), 1–33. https://doi.org/10.1145/3498662

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

Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris. In *European Conference on Object-Oriented Programming (ECOOP)*. 17:1–17:29. https://people.mpi-sws.org/~dreyer/papers/iris-weak/paper.pdf

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

Robbert Krebbers, Jacques-Henri Jourdan, Ralf Jung, Joseph Tassarotti, Jan-Oliver Kaiser, Amin Timany, Arthur Charguéraud, and Derek Dreyer. 2018. MoSeL: A General, Extensible Modal Framework for Interactive Proofs in Separation Logic. *Proc. ACM Program. Lang.* 2, ICFP, Article 77 (jul 2018), 30 pages. https://doi.org/10.1145/3236772

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

Sam Lindley and J. Garrett Morris. 2015. A Semantics for Propositions as Sessions. In *24th European Symposium on Programming (ESOP) (Lecture Notes in Computer Science, Vol. 9032)*. Springer, 560–584. https://doi.org/10.1007/978-3-662-46669-8_23

Alexandre Moine, Arthur Charguéraud, and François Pottier. 2023a. A High-Level Separation Logic for Heap Space under Garbage Collection. *Proc. ACM Program. Lang.* 7, POPL, Article 25 (jan 2023), 30 pages. https://doi.org/10.1145/3571218

Alexandre Moine, Sam Westrick, and Stephanie Balzer. 2023b. *DisLog: A Separation Logic for Disentanglement - Artifact.* https://doi.org/10.5281/zenodo.8414566 Last version available at: https://gitlab.inria.fr/amoine/dislog.

MPL Development Team. 2022. The MaPLe (MPL) compiler v0.3. https://github.com/MPLLang/mpl.

Ike Mulder, Robbert Krebbers, and Herman Geuvers. 2022. Diaframe: Automated Verification of Fine-Grained Concurrent Programs in Iris. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (San Diego, CA, USA) *(PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 809–824. https://doi.org/10.1145/3519939.3523432

Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2020. Cosmo: A Concurrent Separation Logic for Multicore OCaml. *Proceedings of the ACM on Programming Languages* 4, ICFP, Article 96 (June 2020), 29 pages. http://cambium.inria.fr/~fpottier/publis/mevel-jourdan-pottier-cosmo-2020.pdf

Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating State Transition Systems for Fine-Grained Concurrent Resources. In *Programming Languages and Systems*, Zhong Shao (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 290–310.

Peter W. O'Hearn. 2007. Resources, Concurrency and Local Reasoning. *Theoretical Computer Science* 375, 1–3 (May 2007), 271–307. http://www.cs.ucl.ac.uk/staff/p.ohearn/papers/concurrency.pdf

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. https://doi.org/10.1145/2951913.2951935

Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *European Symposium on Programming (ESOP) (Lecture Notes in Computer Science, Vol. 8410)*. Springer, 149–168. http://cs.au.dk/~birke/papers/icap-conf.pdf

Bernardo Toninho, Luís Caires, and Frank Pfenning. 2013. Higher-Order Processes, Functions, and Sessions: A Monadic Integration. In *22nd European Symposium on Programming (ESOP) (Lecture Notes in Computer Science, Vol. 7792)*. Springer, 350–369. https://doi.org/10.1007/978-3-642-37036-6_20

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

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

Sam Westrick. 2022. *Efficient and Scalable Parallel Functional Programming Through Disentanglement.* Ph. D. Dissertation. Carnegie Mellon University. https://www.cs.cmu.edu/~swestric/22/thesis.pdf

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. 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. https://doi.org/10.1145/3371115