EMMA NARDINO, ENS Lyon, Univ Lyon, UCBL, CNRS, Inria, LIP, France LUDOVIC HENRIO, CNRS, Univ Lyon, ENS Lyon, UCBL, Inria, LIP, France GABRIEL RADANNE, Inria, Univ Lyon, ENS Lyon, UCBL, CNRS, LIP, France YANNICK ZAKOWSKI, Inria, Univ Lyon, ENS Lyon, UCBL, CNRS, LIP, France

This article extends tail-call optimisation by applying it to asynchronous calls. We first introduce *Tail-Modulo-Await*, a novel code transformation for asynchronous tail recursive functions that prevents the creation of unnecessary tasks. We then show how to combine Tail-Modulo-Await with the existing Tail-Modulo-Cons optimisation; we obtain an optimisation able to turn a recursive function with multiple tail calls under constructors into a parallel version of the function, also optimised in space.

We formalise both optimisations over representative calculi, and prove them correct through backward simulations. Finally, we provide a proof-of-concept implementation as an OCaml syntax extension and evaluate it experimentally, showing our approach optimises both memory and execution time.

ACM Reference Format:

1 Introduction

The Tail Call Optimisation has been a staple of functional programming languages since its introduction in 1977 in Steele Jr. [17]'s seminal paper "LAMBDA: The Ultimate GOTO". This optimisation hinges on the observation that, when a function call is in *tail position*, i.e., a return position in the program, it can be compiled as a simple JUMP instruction, thus dramatically improving the efficiency of procedure calls. This transformation, which also saves space in the function stack, made its way in numerous languages and compilers, whether enabled by default in languages such as OCaml, Haskell, or Scala; or as an on-demand optimisation in compilers such as LLVM and GCC.

The notion of tail position was initially quite restrictive: only calls in exact return position were considered. Tail-Call *modulo Constructor* [3] extends the class of accepted program with single tail-positions under a data constructor. Thanks to this notion, most functions over lists are now considered tail-recursive, and thus liberated from stack constraints.

However, while lists remain a data structure of choice for functional programmers, shouldn't trees deserve the same care? More broadly, can we hope for a similar optimisation in presence of *multiple* calls in tail position? Surely, this makes little sense in a sequential context: how could we launch several tail-calls without any synchronisation? However, this makes perfect sense in an *asynchronous* context, and more specifically in the presence of futures [5],¹ which are entities representing the result of an ongoing computation. In fact, a notion similar to asynchronous tail calls already exists for OO method calls in asynchronous languages, dubbed forward [8].

In this article, we further extend the notion of tail-position to be "modulo Await". Similarly to how tail-calls in sequential contexts are optimised to use constant stack space, we show how to optimise tail-calls in asynchronous context in order to use no spurious scheduler space. Furthermore, when

¹Coincidentally, both works were published in 1977, and both originated from the LISP community!

Authors' Contact Information: Emma Nardino, emma.nardino@ens-lyon.fr, ENS Lyon, Univ Lyon, UCBL, CNRS, Inria, LIP, Lyon, France; Ludovic Henrio, ludovic.henrio@cnrs.fr, CNRS, Univ Lyon, ENS Lyon, UCBL, Inria, LIP, Lyon, France; Gabriel Radanne, gabriel.radanne@inria.fr, Inria, Univ Lyon, ENS Lyon, UCBL, CNRS, LIP, Lyon, France; Yannick Zakowski, yannick.zakowski@inria.fr, Inria, Univ Lyon, ENS Lyon, UCBL, CNRS, LIP, Lyon, France.

2025. ACM XXXX-XXXX/2025/4-ART https://doi.org/10.1145/nnnnnnnnnnn

```
1 type 'a tree = Leaf | Node of 'a * 'a tree * 'a tree
2
3 val map : ('a -> 'b) -> 'a tree -> 'b tree future
4 let%async rec map f t = match t with
5 | Leaf -> Leaf
6 | Node (x , tl, tr) -> Node (f x, await (map f tl), await (map f tr))
Fig. 1. A parallel and asynchronous map on trees.
```

combined with Tail-Modulo-Cons (tmc), we obtain a Tail-Modulo-Cons-Await optimisation (tmca) able to optimise functions with multiple recursive calls under a constructor. We demonstrate this feature over trees.

A tail-recursive map on trees

Fig. 1 showcases an implementation of an asynchronous map function on trees in an OCaml dialect with asynchronous functions.² Before going over details, let us state our selling point: using our code transformation, this function is fully parallel (all calls to f are executed concurrently), uses a constant memory space per thread, and doesn't use any superfluous scheduler space.

We provide the signature for map on Line 3: it takes a function f, a tree t, and returns the tree where f has been applied to every element of t. In our code, asynchronous-related operations are indicated in **blue** and code internal to our optimisation in **peach**. Concretely, we consider a parametric tree type, defined on Line 1, where each Node is binary and contains a value. As is common for asynchronous functions, map returns more precisely a *future* for this tree. By tagging the function as **async** on Line 4, calls to this function create a future and launch the associated computation in a parallel sub-task. The result can then be retrieved using the **await** primitive: this is done for both recursive calls on Line 6. Finally, we assemble the value into a Node and return it, implicitly in a future.

In the absence of any optimisation, the behaviour of this function would be, on each Node constructor, to launch the task for one branch, wait for it to finish, launch the task for the other branch, wait for it to finish, then allocate the new constructor and return it. The stack would grow linearly in the height of the tree; furthermore, the execution would be overall sequential, while paying additional cost for synchronisation. The asynchronous programmer would however not write the code above and would favour a more parallel version by replacing Line 6 with:

let tl' = map f tl in let tr' = map f tr in Node (f x, await tl', await tr')
This version parallelizes the treatment of the left and right branches (this introduces task interleaving
as discussed below). Its main drawback is that it can have as many "active" tasks as there are nodes
in the tree; indeed all the finished tasks which have at least one unfinished task below it in the tree
structure are still blocked in an await statement.

Our approach optimises the two preceding programs by both parallelizing subtasks, and terminating all tasks that are only awaiting others. To achieve this, we translate the function to *Destiny Passing Style*, an asynchronous interpretation of Destination Passing Style [16]. An idealized version of the transformed function map_dps is shown on Fig. 2a; its execution is illustrated on Fig. 2b.

At the top level, a single future fut is tied to the filling of a memory cell, a destiny d (Line 11). From there, a helper function map_dps is called. The function map_dps takes as an additional argument a *destiny* d (Line 1) which represents the memory location where the result of the asynchronous computation should be written to. When returning a value, such as Leaf on Line 2, we set the destiny d with d < Leaf. In the recursive case, we perform the allocation ahead of time—similarly to what

2

²Our dialect is implemented with an OCaml syntax extension.

[,] Vol. 1, No. 1, Article . Publication date: April 2025.



Fig. 2. Translation of map on trees from Fig. 1 using Tail-Modulo-Cons-Await

happens in tmc- but put empty holes in lieu of its arguments. The current destiny can therefore be immediately set to v, the freshly allocated node. Both recursive calls are then forked in their own thread, one for each child: we must pass them their respective destiny, i.e., the corresponding hole in v they are responsible for filling-we address fields of constructors by position. The calls to fork dps handling recursive calls each fork a new task running the function map dps with the arguments passed as parameters; they create no stack for executing map dps and discard the value returned by the function so that the call uses a minimal space. Finally, we compute the value f x and set it in v. When all parallel tasks have finished, there are no more holes in the structure, the future is marked as resolved, and the value can be retrieved. Fig. 2b shows an intermediate state of the execution, the left side shows the status in memory of the tree being computed, the right side shows the tasks involved in the computation. Scheduler space is represented in blue. Futures and destiny both play a synchronisation role and serve as pointers to actual data, and are thus striped. The original future fut is yet unresolved (X). Tasks map_dps d f t, map_dps d.2 f t2, and map dps d.2.1 f t21 are finished, the corresponding destinies have been filled and the tasks have been garbage-collected. Task map dps d.2.2 f t22 has just performed the allocation of a Node (with holes). Task map dps d.1 f t1 is still ongoing. This illustrates that only the tasks actually computing are still active on the scheduler side, contrarily to the versions that do not use tmca.

This code transformation automatically reveals parallelism present in the source program. In fact, it doesn't fully respect the sequential semantics. Indeed on Line 6 in the code above, the transformed code can *interleave* subtasks in map f tl and map f tr, while the sequential semantics would enforce an execution order, without interleaving. Thus, this automatic parallelisation comes at the risk of introducing undesired behaviors. We mitigate this risk in two ways.

First, we only modify code *at the level of a constructor in tail position*. This means the programmer can easily enforce sequentiality via simple let-bindings. For instance, we could replace Line 6 by:

let x' = f x in Node (x', await (map f tl), await (map f tr))

In this case, it would wait for f to return before proceeding to the two subtrees in parallel (thus bounding the parallelism by the width of the tree). This aligns with good practice in languages like OCaml, where users who care about order of operation should state it explicitly via let-bindings.³

Second, we introduce two variants for each constructor: *parallel* constructors which behave as above, and *sequential* ones which don't introduce parallelism and only allow one tail position. With these variants, we show that our transformation preserves the semantics.

³Naturally, this might come with some cost in scheduler- or stack- space, since the only way to optimise these aspects is through parallelism and thus interleaving of tasks.

Contribution and Plan

In this paper, we introduce the following contributions.

- In Section 2, we introduce a novel code transformation for asynchronous tail-recursive function dubbed Tail-Modulo-Await. This code transformation generalises and automatises the treatment of forward present in the literature on asynchronous computing.
- In Section 3, we combine Tail-Modulo-Await and Tail-Modulo-Cons to derive a transformation allowing for having multiple tail calls launched concurrently. This code transformation can, depending on the users' choices, maximise parallelism, or strictly preserve sequential semantics.
- We formalise both on minimal calculi and prove them sound via backward simulations.
- We provide a proof-of-concept implementation as an OCaml syntax extension and evaluate it experimentally in Section 4.

Tail-Modulo-Await 2

In this section, we set aside constructors and focus on an optimisation dubbed Tail-Modulo-Await (tma)—we shall extend this optimisation to handle our introductory example later, in Section 3. After illustrating informally the optimisation on an example, we introduce two core calculus, a source and a target, over which we formalise the transformation.

Motivating example 2.1

Fig. 3 presents a toy program illustrating tma. The check_urls function iterates recursively over a list of urls in order to check whether they actually point to some resource. Such checks require to communicate over the network, which takes time: to avoid halting the whole system, we hence perform them via a call to an asynchronous operation check_url: url -> bool future. This operation, provided by our favourite Web Client, is wrapped in an

```
val check urls : url list -> bool future
2 let%async rec check urls files =
   match files with
   [] -> true
    url :: rest ->
     if await (check url url) then
       await (check urls rest)
     else false
```

```
Fig. 3. Checking all URLs, with async/await
```

await on Line 6. But a similar issue arises on Line 7 in the recursive call! The check urls function itself must therefore be declared asynchronous, and recursive calls must be wrapped in an await.

Recursion and async/await seem to play well together: we have elegantly turned our recursive function into a non-blocking one. But one may wonder what the space behaviour of such a function is? Indeed, the synchronous equivalent (where async/awaits are removed) is clearly tail-recursive, and hence its recursive calls can operate in constant stack. Hopefully, the same is true of this asynchronous variant. But in this case, check urls's memory behaviour deserves closer inspection, as illustrated on Fig. 4. Strictly speaking, the function does not consume stack-space either: it delegates to async, where each call to await is in tail position. However, as illustrated in Section 1, in an asynchronous context we must worry about scheduler-space as well! Each recursive call creates



Fig. 4. Memory behaviour of check_urls



Fig. 5. Checking all URLs with the Tail-Modulo-Await transformation

a future, which needs to be awaited for, forming a chain the size of the list 10 taken as argument. As the chain grows, the memory gets cluttered with un-active futures and tasks.

As per its name, Tail-Modulo-Await optimises such functions that are almost tail recursive, but whose tail calls happen under await. Consider the effect of the optimisation in Fig. 5: the stack remains of constant size, only one recursive task is ever alive, and the chain of futures is reduced to length one. To achieve this, the recursive computation check_urls_dps is crucially not asynchronous any more. Furthermore, it takes a destiny d as argument, a *write-once* memory cell in which we shall set the boolean resulting from its computation. The function check_urls is then defined as a top level wrapper that creates a pair of a destiny to be passed to the helper, and of a future that will get resolved by this destiny: such a pair is referred to as a *promise* in the literature. To ensure we are non-blocking, the helper is forked in a separate thread, before returning the future. Naturally, calls to await which are not in tail position remain, such as the one to check_url on Line 5.

One might wonder what happens if a terminal position contains a call to a function which is not in Destiny-Passing-Style ? For instance, if we were to unroll the check_urls function once, we would obtain the following extra-clause in our pattern matching:

| [url] -> await (check_url url)

The check_url function, which is provided by an external library, doesn't necessarily have a DPS version. Nevertheless, we would like to avoid creating an intermediary task whose sole purpose is to wait before filling the destiny. This exactly corresponds to forward [8] from the literature. forward (d, fut) registers that when future fut is resolved, the result should be used to set destiny d. This can be implemented efficiently in the scheduler to avoid an extra task. Furthermore, it works for any future, not just asynchronous function application. This efficiently bridges the gap between the DPS world and the rest of the asynchronous world.

2.2 Calculus

We now present our calculus, a first order imperative and asynchronous λ -calculus which formalises this first transformation, Tail-Modulo-Await, without any constructor. Fig. 6 shows at once the syntax of both the source language and the target one, using the following code colour: *Light blue* elements are source-specific, *peach* elements are target-specific, and Grey elements are runtime syntax. dynamic semantics of the language.

Shared core. Both languages includes standard control flow and support for references, with operations for creation (newref(e)), reads (!e), and updates (e := e). As motivated earlier, this first calculus contains no data type, static values are reduced to basic types. Its functional side is first order: we assume all functions are λ -lifted in a table of function definitions *Fun*. In both languages,

(expressions) $Expr \ni e ::= v \mid x \in Var \mid let x = e in e'$ (functions) $::= f \mapsto \lambda x. e$ fd |(f e) $f \mapsto \lambda_{async} x. e$ | if e then e_t else e_f $f \mapsto \lambda_{dps} \delta, x. e$ $|\operatorname{newref}(e)| ! e | e := e$ $::= \overline{fd}$ Fun await(e) $(f #_t (e', e))$ $Val \ni v ::= b \in \mathbb{B} \mid c \in Const$ (values) forward(e, e') $\ell \in Loc \mid fut \in \mathbb{F}$ refine(e, e') $d \in Dest$

Fig. 6. Syntax – *Light blue* elements only exist in the source language and *peach* elements only in the target language. Grey elements denote runtime syntax.

functions are tagged with their calling behaviour when declared: sequential by default, or with an asynchronous behaviour. When a task t is created, a future *fut* is associated to it, and any task can get the result of the computation by performing await(fut). If t is not finished yet, we say *fut* is *unresolved*, in which case await(fut) is blocking; otherwise, it retrieves the computed value. Futures are runtime values.

Source language. The only syntax for function calls is (f e). The behaviour of the call directly depends on the tag associated to f in its declaration: synchronous by default, unless tagged by λ_{async} . In the latter case, the behaviour is asynchronous: each call spawns both a task and a fresh future, whose resolution is tied to the task.

Target language. Here, there are no such coupling between a future and a task responsible for fulfilling it. Instead, each future is mapped to exactly one *destiny*: a value acting as a write-once memory location which can be used to fill the future. A pair of a read-only future and a write-once reference is often referred to as a *promise* in asynchronous programming languages. These destinies are not bound to a particular task, and can be passed by argument to subsequent calls to asynchronous functions. To this end, asynchronous functions are declared in destiny-passing style (dps), i.e., of the shape ($\lambda_{dps}\delta$, *x. e*), where δ is the variable used to pass the destiny.

There are two distinct ways to call a dps function. The first one, as (f e), is similar to the source language: it creates a task, a promise (read a pair *fut*, *d*, where *d* is the destiny resolving future *fut*) and maps the destiny to *d*. This corresponds to automatically introducing the wrapper function in Fig. 5. The second one is specific to tail calls to asynchronous functions: we explicitly represent such calls in the syntax as $(f \#_t (d, e))$. Instead of creating a fresh promise, such tail calls *delegate* the resolution of the promise to the newly created task, by passing its destiny *d* as the new task's destiny. This behaviour corresponds to **fork_dps** calls in Fig. 5.

Unlike futures, the resolution of destinies is not associated to the end of a task but must be performed explicitly, via the statement refine(d, e). In the absence of constructor, this corresponds to d < e in our informal syntax.⁴ Finally, forward(fut, d) binds the resolution of a promise to the resolution of another one: the destiny *d* is automatically filled when the promise *fut* is resolved.

2.3 Semantics

We now equip both languages with an operational, small-step, semantics. We assume a table of function definitions *Fun* implicitly parametrising the semantics.

Shared core. At runtime, values are extended with locations and futures (Fig. 6). Both languages share a thread-local core, whose reduction is defined over *local* configurations (σ , e) carrying a store and an expression. The local reduction $\cdot \rightarrow \cdot$ is provided on Fig. 8. This reduction relation is

⁴As the name hints, refine takes on more responsibilities when constructors are introduced, in the next section

(store)
$$\sigma \qquad ::= \left[\overline{\ell \mapsto v} \right]$$

(evaluation contexts)
$$Ectx \ni C ::= \Box \mid \text{let } x = C \text{ in } e \mid \text{if } C \text{ then } e_t \text{ else } e_f \mid (f C) \mid newref(C) \mid C := e \mid \ell := C \mid await(C) \mid (f \#_t C) \mid forward(C, e') \mid refine(e, C)$$

Fig. 7. Runtime syntax for Stores
$$\sigma$$
 and Contexts C

$$\frac{\sigma, e \to \sigma', e'}{\sigma, C[e] \to \sigma', C[e']} \qquad \frac{\ell \notin \operatorname{dom}(\sigma)}{\sigma, \operatorname{newref}(v) \to \sigma[\ell \mapsto v], \ell} \qquad \frac{(\ell \mapsto v) \in \sigma}{\sigma, !\ell \to \sigma, v}$$
$$\frac{(f \mapsto \lambda x. e) \in \operatorname{Fun}}{\sigma, (\ell \coloneqq v) \to \sigma[\ell \mapsto v], ()} \qquad \frac{(f \mapsto \lambda x. e) \in \operatorname{Fun}}{\sigma, (f v) \to \sigma, e[x \mapsto v]}$$

Fig. 8. Sequential semantics (rules for if and let are omitted) $-\sigma$, $e \rightarrow \sigma'$, e'

(unordered list of tasks) $Task ::= (main \mapsto e) \sqcup \overline{(fut \mapsto e)}$ (source configurations) $c ::= \sigma, Task$

(a) Runtime configurations

Fig. 9. Source Configuration

 $\sigma_s = \emptyset$ Task_s = (main $\mapsto e$) (b) Initial configurations

$$\begin{split} & \overbrace{\sigma, rask}^{\text{STEP}} \underbrace{\sigma, e \to \sigma', e'}_{\sigma, Task \ \sqcup \ (fut \mapsto e) \to \sigma', Task \ \sqcup \ (fut \mapsto e')} \\ & \overbrace{\sigma, Task \ \sqcup \ (fut \mapsto e) \to \sigma', Task \ \sqcup \ (fut \mapsto e')} \\ & \underbrace{(fut' \mapsto v) \in Task}_{\sigma, Task \ \sqcup \ (fut \mapsto C[\texttt{await}(fut')]) \to \sigma, Task \ \sqcup \ (fut \mapsto C[v])} \\ & \overbrace{\sigma, Task \ \sqcup \ (fut \mapsto C[(fut')]) \to \sigma, Task \ \sqcup \ (fut \mapsto C[v])} \\ & \overbrace{\sigma, Task \ \sqcup \ (fut \mapsto C[(fv)]) \to \sigma, Task \ \sqcup \ (fut' \mapsto e[x \mapsto v]) \ \sqcup \ (fut \mapsto C[fut'])} \\ & \overbrace{Fig. 10. \ Global \ source \ reduction \ -\sigma, Task \ \to \ \sigma', Task'} \\ \end{split}$$

completely standard: evaluation contexts are defined on Fig. 7, and the four other rules specify the operations over the store, and the beta reduction for sequential function calls.

Source language. In the source, task creation and synchronisation is entirely managed through futures: runtime configurations (Fig. 9) pair a store with a set of *tasks*, each bound to a different *future* (with the exception of the main thread). Initial configurations have an empty store and a single *main* task, while final configurations have all tasks evaluated down to a value.

The concurrent reductions are given on Fig. 10. The STEP rule allows any task to progress according to the local semantics. Rule AWAIT-RESOLVED evaluates an await(fut') statement: if the computation of fut' is reduced down to a value, this value is substituted. Finally, the ASYNC-CALL rule handles task creation by creating a fresh future and a new entry in the set of tasks.

Target language. Runtime configurations for the target language are described on Fig. 11. Contrary to the source, *any* thread can fulfil a future: *Fut* hence keeps track of their status. We write [$fut \mapsto \mathfrak{U} d$]

(maps of futures)	Fut	$::= \mathbb{F} \xrightarrow{\text{fin}} Dest + Val$
(Tasks)	Task	$::= (main \mapsto e) \sqcup \overline{(tid \mapsto e)}$
(target configurations)	С	$::= \sigma$, Task, Fut

 $\sigma_t = \emptyset$ $Task_t = (main \mapsto e)$ $Fut_t = \emptyset$ (b) Initial configuration

(a) Runtime configuration

Fig. 11. Target Configuration

STEP
$\sigma, e \rightarrow \sigma', e'$
$\overline{\sigma, Task} \ \sqcup \ (tid \mapsto e), Fut \rightarrow \sigma', Task \ \sqcup \ (tid \mapsto e'), Fut$
AWAIT-TARGET
$Fut [fut] = \Re v$
$\sigma, Task \sqcup (tid \mapsto C[\texttt{await}(fut)]), Fut \to \sigma, Task \sqcup (tid \mapsto C[v]), Fut$
CHAIN
$Fut [fut] = \mathfrak{U} d \qquad Fut [fut'] = \mathfrak{R} v$
$\overline{\sigma, Task} \sqcup (tid \mapsto \mathbf{forward}(fut', d)), Fut \to \sigma, Task, Fut[fut \mapsto \Re v]$
DPS-CALL
$(f \mapsto \lambda_{dps}\delta, x. e) \in Fun$ fut', d, tid' fresh
$\sigma, Task \sqcup (tid \mapsto C[(f v)]), Fut \rightarrow$
$\sigma, Task \sqcup (tid \mapsto C[fut']) \sqcup (tid' \mapsto e[\delta \mapsto d][x \mapsto v]), Fut[fut' \mapsto \mathfrak{U}d]$
TAIL-CALL
$(f \mapsto \lambda_{dps}\delta, x.e) \in Fun$ $tid' \notin dom(Task)$
$\overline{\sigma, Task} \sqcup (tid \mapsto (f \#_t (d, v))), Fut \to \sigma, Task \sqcup (tid' \mapsto e[x \mapsto v][\delta \mapsto d]), Fut$
FUTURE-FILL
$Fut [fut] = \mathfrak{U} d$
$\sigma, Task \sqcup (tid \mapsto refine(d, v)), Fut \xrightarrow{\varepsilon} \sigma, Task, Fut [fut \mapsto \Re v]$
Fig. 12. Global target reduction $-\sigma$, <i>Task</i> , <i>Fut</i> $\rightarrow \sigma'$, <i>Task'</i> , <i>Fut'</i>

when *fut* is yet unresolved, and tied to the destiny *d*; and [*fut* $\mapsto \Re v$] when *fut* is resolved with value *v*. The target tasks are not indexed by futures, but rather by anonymous identifiers, written *tid*, that cannot be awaited. Initial configurations are similar to the source language, with an empty map of futures. A configuration is *final* if its main thread is reduced to a value, it is the only thread left, and every future in *Fut* is resolved.

Reduction is defined on 12. Local reductions can still be lifted to any task (STEP). AWAIT-TARGET handles futures synchronisation: **await**(*fut*) can only be evaluated if *fut* is resolved, i.e., if *Fut* [*fut*] = $\Re v$ for some value v.

CHAIN evaluates a **forward**(*fut'*, *d*) statement when *fut'* is resolved, and in turn resolves the future linked to destiny *d*. Its behaviour is seemingly similar to the one of **await**: it blocks until *fut'* is resolved. However **forward** can only occur in tail position, as captured by the absence of an evaluation context in CHAIN. Thus tasks of the form (*tid* \mapsto **forward**(*fut*, *d*)) are kept asleep

(tail contexts) $Tctx_n \ni C_n^{tl} ::= \Box : Tctx_1 | \operatorname{let} x = e \operatorname{in} C_n^{tl} : Tctx_n | \operatorname{if} e \operatorname{then} C_n^{tl} \operatorname{else} C^{tl}_m : Tctx_{n+m}$ $\Box_1[e] := e (\operatorname{let} x = e \operatorname{in} C_n^{tl})[e_1 | \cdots | e_n] := \operatorname{let} x = e \operatorname{in} (C_n^{tl}[e_1 | \cdots | e_n])$ (if $e \operatorname{then} C_n^{tl} \operatorname{else} C_m^{tl})[e_1 | \cdots | e_{n+m}] := \operatorname{if} e \operatorname{then} (C_n^{tl}[e_1 | \cdots | e_n]) \operatorname{else} (C_m^{tl}[e_{n+1} | \cdots | e_{n+m}])$

(a) Tail contexts: definition and filling operation

$$\mathcal{D}\llbracket C_n^{tl} \llbracket e_1 & | \cdots & | e_n \rrbracket \rrbracket_{\delta} \coloneqq C_n^{tl} \llbracket \mathcal{D}\llbracket e_1 \rrbracket_{\delta} & | \cdots & | \mathcal{D}\llbracket e_n \rrbracket_{\delta} \rrbracket_{\delta}$$

$$Fun_s \Rightarrow Fun_t$$

$$f \mapsto \lambda x. e \rightsquigarrow f \mapsto \lambda x. e$$

$$\mathcal{D}\llbracket await(e) \rrbracket_{\delta} \coloneqq \mathcal{D}_ a \llbracket e \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e \rrbracket_{\delta} \coloneqq refine(\delta, e) \quad (otherwise)$$

$$\mathcal{D}\llbracket e \rrbracket_{\delta} \coloneqq refine(\delta, e) \quad (otherwise)$$

$$\mathcal{D}_ a \llbracket C_n^{tl} \llbracket e_1 & | \cdots & | e_n \rrbracket \rrbracket_{\delta} \coloneqq C_n^{tl} \llbracket \mathcal{D}_ a \llbracket e_1 \rrbracket_{\delta} & | \cdots & | \mathcal{D}_ a \llbracket e_n \rrbracket_{\delta} \rrbracket$$

$$\mathcal{D}\llbracket \emptyset, T_{main} \rrbracket = \emptyset, T_{main}, \emptyset$$

$$(c) \text{ Transformation over initial configurations}$$

$$\mathcal{D}\llbracket e \rrbracket_{\delta} \coloneqq refine(\delta, e) \quad (otherwise)$$

$$\mathcal{D}_ a \llbracket C_n^{tl} \llbracket e_1 & | \cdots & | e_n \rrbracket \rrbracket_{\delta} \coloneqq C_n^{tl} \llbracket \mathcal{D}_ a \llbracket e_1 \rrbracket_{\delta} & | \cdots & | \mathcal{D}_ a \llbracket e_n \rrbracket_{\delta} \rrbracket$$

$$\mathcal{D}_ a \llbracket e \rrbracket_{\delta} \coloneqq f = f \text{ orward}(e, \delta) \quad (otherwise)$$

$$\mathcal{D}_ a \llbracket e \rrbracket_{\delta} \coloneqq f \text{ orward}(e, \delta) \quad (otherwise)$$

Fig. 13. The Tail-Modulo-Await transformation $-\mathcal{D}[\![\cdot]\!]$

in the configuration until *fut* is resolved, but do not block pending computations. An optimised implementation could store such **forward**ed destinies directly on the future *fut* itself.

DPS-CALL and TAIL-CALL handle asynchronous function calls, as described previously. DPS-CALL creates a fresh promise (thus a new mapping [$fut \mapsto \mathfrak{U} d$] in Fut), and a new task resolving said promise by executing the called function. TAIL-CALL handles specifically calls in tail position: it passes the current destiny to the called function and hosts it in a new task, but furthermore kills the current task as we know there is nothing left to computer afterwards. Finally, FUTURE-FILL handles future resolution, killing the current task as well, as **refine** statements only exist in tail position.

2.4 The Tail-Modulo-Await Transformation

We can finally define the Tail-Modulo-Await transformation in Fig. 13. It is applied to all asynchronous functions $\lambda_{async} \cdot$. in *Fun* (Fig. 13b). The rest of the program, and in particular the main thread, is not affected (Fig. 13c).

Its core is a recursive function over expressions $\mathcal{D}[\![\cdot]\!]_{\delta}$ where δ is a variable name bound to a destiny (Fig. 13d). It relies on the notion of *tail-context* (Fig. 13a) which captures in an expression the set of positions at which a call would be tail.

The function $\mathcal{D}[\![\cdot]\!]$ itself proceeds by pattern-matching, recursing through the program towards tail positions, leaving the rest unperturbed. If it reaches an $\mathsf{await}(e)$, it switches mode, now recursing e via $\mathcal{D}_{\mathbf{a}}[\![\cdot]\!]$, until it reaches either an asynchronous function call to transform it into a tail-call, or any other expression to transform it into a **forward**. Finally, for other expressions e, $\mathcal{D}[\![\cdot]\!]$ emits **refine** (δ, e) to fill the current destiny.

Correctness of Tail-Modulo-Await 2.5

We prove that the tma transformation does not introduce new behaviors: if a source program p_s is transformed into a target program p_t , then any execution of p_t can be matched by an execution of p_s , leading to similar final states, and identical stores (i.e., p_t is a behavioural refinement of p_s). As is standard, we prove it by exhibiting a backward simulation. The simulation admits stuttering steps, capturing "administrative" steps occurring in the target language but matched by no step in the source. These steps exactly correspond to FUTURE-FILL, whose transitions are identified by an " ε ". We hence need to ensure that a transformed program cannot take infinitely many ε steps in a row.

Definition 2.1 (backward simulation). We say that \succeq is a backward simulation if for any c_s, c_t, c'_t such that $c_s \geq c_t$, we have:

if ct → ct, then cs ≥ ct
if ct → ct, there exists cs such that cs → cs, and cs ≥ ct

Which allows us to state the correctness theorem for tma

THEOREM 2.2. There exists a relation \succeq such that:

- (1) \succ is a backward simulation
- (2) forall c_s initial source configuration, $c_s \geq \mathcal{D}[\![c_s]\!]$
- (3) if $c_s \succeq c_t$, and c_t is a final configuration, then there is a final source configuration c'_s such that $c_s \rightarrow^* c'_s$, and c'_s and c_t have the same final main value.

In the remaining of this section, we define the relation \succ and sketch the simulation proof.

2.5.1 Definition of \succeq . We define \succeq on Fig. 14c in three successive layers. First at the level of expressions, the relation \sim_d is parametrised by a destiny d. It relates a source expression e to its syntactic transformation, in which the free variable δ is substituted by d (EXPR-TRANSFO). This alignment can however be broken: the target may reduce to expressions like refine(d, await(e)), with synchronous function calls in tail-positions (EXPR-REFINE).

At the level of tasks, \simeq (Fig. 14b) relates sets of threads which collectively capture the same computation. The relation is of the form $T_s \simeq T_t$, $\int fut_0 \mapsto \Re v | \mathfrak{U} d |$ where $T_{s,t}$ denote sets of threads computing the same future. On the source side, a task computing an asynchronous function and its subsequent recursive calls is related on the target side to the corresponding pair of thread and future/destiny mapping. One of the main purpose of tma is to prevent chains of hanging tasks all awaiting the following, i.e., groups of tasks with the shape $(fut_0 \mapsto await(fut_1)) \sqcup \cdots (fut_{n-1} \mapsto await(fut_n))$ await (fut_n)). We hence introduce the notation Aw [fut_n .. fut_n] for such chains, and capture this intuition in TASK-COMP. We must furthermore keep track of resolved futures on either sides (TASK-VAL), and bookkeep intermediary states featuring a resolved future yet to be garbage collected (TASK-AWAIT-VAL).

Finally, Fig. 14c composes task-level relations into a simulation relation between configurations. The stores $(\sigma_{s,t})$ and the main thread (main $\mapsto e$) are in exact agreement. The other threads on each side can be partitioned into *n* pools $(T_{s,t}^{(i)})$, one for each future in the target, such that they are pairwise related by \simeq . Intuitively, a new pool is created when an asynchronous function is called from a non-tail position (by rule ASYNC-CALL/DPS-CALL), and are never removed from the configuration. Finally, we maintain two invariants: (6) that no two pools can use the same destiny name, and (7) that the futures resulting from tail-calls in a pool are *internal*, i.e., they cannot appear in the store or any other thread.

2.5.2 *Proof sketch.* The proof of backward simulation relies on the observation that (1) a step in a thread belonging to a pool leaves all other pools unchanged, and (2) it is simulated in the source by EXPR-TRANSFO

EXPR-REFINE

 $e \sim_d refine(d, e)$

$$e \sim_d \mathcal{D}[\![e]\!]_{\delta}[\delta \mapsto d]$$

(a) Expression-level relation

$$\frac{e \sim_{d} e'}{Aw[fut_{0}..fut_{n}] \sqcup (fut_{n} \mapsto e) \simeq (tid \mapsto e'), [fut_{0} \mapsto \mathfrak{U} d]} \qquad \frac{TASK-VAL}{(fut_{i} \mapsto v)^{0 \le i \le n} \simeq \emptyset, [fut_{0} \mapsto \mathfrak{R} v]} \\
\frac{TASK-AWAIT-VAL}{Aw[fut_{0}..fut_{n}] \sqcup (fut_{n} \mapsto v) \simeq \emptyset, [fut_{0} \mapsto \mathfrak{R} v]} \\
(b) Task-level relation
$$\begin{pmatrix}
(1) & \sigma_{s} = \sigma_{t} \\
(2) & Task_{s} = (main \mapsto e) \sqcup T_{s}^{(1)} \sqcup \cdots T_{s}^{(n)} \\
(3) & Task_{t} = (main \mapsto e) \sqcup T_{t}^{(1)} \sqcup \cdots T_{t}^{(n)} \\
(4) & Fut_{t} = [\overline{fut_{i} \mapsto \overline{\tau}_{t}^{1 \le i \le n}}] \quad (\overline{\tau}_{i} = \mathfrak{U} d_{i} | \mathfrak{R} v_{i}) \\
(5) & \forall i, T_{s}^{(i)} \simeq T_{t}^{(i)}, [fut_{i} \mapsto \overline{\tau}_{i}] \\
(6) & \forall i \ne j, d_{i} \ne d_{j} \quad (if defined) \\
(7) & \forall i, fut \in dom(T_{s}^{(i)}) \setminus \{fut_{i}\}, \begin{cases} fut \notin Im(\sigma) \\ \forall i \ne T_{s}^{(j)} \end{cases}$$$$

(c) Configuration-level relation



steps taken in the corresponding pool. We then distinguish two broad cases: a step in main must be exactly mirrored, while a step in pool $T_t^{(i)}$ must be matched in $T_s^{(i)}$ and restore (5) for index *i*. We refer the reader to Appendix C for the full case analysis, and detail here two informative cases.

Stepping in main. Since the target main is syntactically identical to the source one, the only steps it can take are local STEPS, DPS-CALL, and AWAIT-TARGET.

To illustrate thread creation, consider a DPS-CALL step: it is matched by an ASYNC-CALL in the source, using the same future name. The stores are omitted from the configurations.

$$Task_{s} \sqcup (\min \mapsto C[(f v)]) \succcurlyeq Task_{t} \sqcup (\min \mapsto C[(f v)]), Fut_{t}$$

$$\downarrow^{\text{DPS-CALL}} \downarrow (f \mapsto \lambda_{async} x. e) \in Fun_{s}$$

$$Task_{s} \sqcup (\min \mapsto C[fut']) \rightarrowtail (fut' \mapsto e[x \mapsto v]) \succcurlyeq Task_{t} \sqcup (\min \mapsto C[fut']) \sqcup (tid \mapsto \mathcal{D}[e]_{\delta}[\delta \mapsto d][x \mapsto v]),$$

$$Fut_{t} \sqcup [fut' \mapsto \mathfrak{U} d]$$

The main threads (highlighted in *red*) are still identical, and the newly created tasks (in *blue*) verify clause (5) by TASK-COMP. *d* and *fut*' are fresh so (6) and (7) are satisfied.



Fig. 15. A parametric map on trees

Stepping in a parallel thread. Given a target step in pool $T_t^{(i)}$, clause (5) constraints the shape of the corresponding pool $T_s^{(i)}$. Considering a TAIL-CALL step as illustration:

 $\begin{array}{ccc} Task_{s} \sqcup \mathsf{Aw}[fut_{0}..fut_{n}] \\ \sqcup (fut_{n} \mapsto \mathsf{await}((f \ v))) \\ & & & \\ & & \\ & & \\ \hline Task_{s} \sqcup \mathsf{Aw}[fut_{0}..fut_{n+1}] \\ \sqcup (fut_{n+1} \mapsto e[x \mapsto v]) \end{array} & \succcurlyeq & Task_{t} \sqcup (tid \mapsto (f \ \#_{t} (d, v))), \\ & & Fut_{t} \sqcup [fut_{0} \mapsto \mathfrak{U} d] \\ & & \\ & & \\ & & \\ & & \\ Task_{t} \sqcup (tid' \mapsto \mathcal{D}[e]]_{\delta}[\delta \mapsto d][x \mapsto v]), \\ & Fut_{t} \sqcup [fut_{0} \mapsto \mathfrak{U} d] \end{array}$

The task-level invariant \simeq ensures via TASK-COMP that we can take a matching ASYNC-CALL in the source. Since fut_{n+1} is fresh, (7) is preserved.

Initial and final states. Initial configurations are trivially related. If $c_s \geq c_t$, with c_t a final target configuration, then the main thread is reduced to a value v, and the same is true in c_s by clauses (2) and (3). Furthermore, every future is resolved, so each $T_s^{(i)}$ is either of the form $\overline{(fut_i \mapsto v_i)}$ by TASK-VAL, and thus satisfy finality; or $\operatorname{Aw}[fut_0...fut_n] \sqcup (fut_n \mapsto v)$, in which n AWAIT-RESOLVED steps can be taken to reach a final configuration.

3 Tail-Modulo-Cons-Await

Before adding constructors to the calculus, let us revisit some subtle semantic considerations, using the example of a tail-recursive map on trees introduced in Section 1.

As discussed, we consider two possible semantics for constructors in the presence of async/await. Constructors tagged as sequential ensure their arguments are evaluated in sequence, while those tagged *parallel* maximise parallelism. These different behaviours are illustrated on Fig. 15, which features the map on trees from Fig. 1, but where the final Node constructor is parametrised by a tag ρ . We capture the meaning of this tag by translating map into two programs semantically equivalent to the original one, depending on the value of ρ , but whose order of evaluation is made explicit through let bindings. Beware, this is *not* the result Tail-Modulo-Cons-Await transformation, which aims to pre-allocate the constructor and avoid intermediate task creations, as showcased in Section 1.

For ρ = seq, the *sequential* semantics showcased in Fig. 15b prevents any interleaving, and fixes an order of execution (here, left-to-right). In this case, only one **await** statement can be considered in tail position, the rest of the body must be awaited on to launch the final call. This semantics is similar to the Tail-Modulo-Cons transformation formalized by Allain et al. [3].

For $\rho = ||$, evaluation can maximise parallelism, as showcased in Fig. 15c. The semantics is equivalent to first launching both sub-tasks (here, map f tl and map f tr), and then the sequential code (f x), hence removing any synchronisation between the tasks.

Fig. 16. Syntax extension for constructors – Extends Fig. 6 - Light blue elements only exist in the source language and *peach* elements only in the target language. Grey elements denote runtime syntax.

(store) σ ::= ... (eval contexts) $Ectx \ni C$::= ··· | [t, C, e] | [t, v, C] | [t, \blacksquare , C] | e.(C) | C.(v) | e.(e) $\leftarrow C$ | e.(C) $\leftarrow v$ | C.(v) $\leftarrow v$ | refine(e, C)

Fig. 17. Runtime syntax for constructors – Extends Fig. 7

$$\operatorname{sta}(C_n^{tl}[\operatorname{await}(e_1) | \cdots | \operatorname{await}(e_n)]) ::= C_n^{tl}[e_1 | \cdots | e_n]$$

Fig. 18. Utility function to strip awaits in tail positions – sta(e)

In both cases, tmca pre-allocates the Node constructor and constructs the structure in memory piece-by-piece, as shown in Section 1. But crucially in the case of the parallel semantics, the intermediary structure during execution might have several holes. The top-level future is resolved when there are no more holes in the structure.

In the remainder of this section, we extend the calculus introduced in Section 2 in order to formalise the Tail-Modulo-Cons-Await code transformation, and prove its correctness.

3.1 Calculus

To model data-constructors, we add binary constructors to the language as tagged mutable memory blocks { t, e_1 , e_2 }, with notations and semantics inspired by Allain et al.'s formalization of Tail-Modulo-Cons [3], but with the additional *sequential* (t_{seq}) and *parallel* ($t_{||}$) discussed above. As in [3], we have a separate, runtime-only block construction [t, e_1 , e_2], in which evaluation order is always left-to-right, and which performs allocation when all fields have finished evaluating.

Memory locations are still runtime-only, and there is no general pointer arithmetic. Once allocated, constructor fields are accessible using ℓ . (*i*), where *i* is an offset in {0, 1, 2}. Destinies are now actual memory locations, in the store.

The purpose of the **refine** instruction is extended: refine(d, v) takes a value v which may have some *holes*—syntactic markers (written \blacksquare) to denote positions in a value that we want to compute later—, plugs said value at location d, and returns the locations of said *holes* in v. The positions marked are initialized with value \checkmark (read *Undef*). For simplicity, we keep arguments of **refine** as *shallow* constructors (no sub-constructors) with at most two arguments. We extend slightly this construct for our implementation in Section 4.1.

3.2 Semantics

Both the source and target runtime configurations are the same as for the Tail-Modulo-Await transformation. The same is true for initial and final configurations.

Shared core. We add support for constructors with contexts in Fig. 17, utility functions in Fig. 18 and reduction rules in Fig. 19.

$\begin{array}{l} \text{CONSTR-READ} \\ \sigma \left[\ell + i \right] = v \end{array}$	$\begin{array}{c} \text{constr-write} \\ \ell+i \in \operatorname{dom}(\sigma) \end{array}$	$\begin{array}{l} \text{CONSTR-ALLOG} \\ \forall i \in Inc. \end{array}$	c <i>lex</i> , <i>ℓ</i> + <i>i</i> ∉ dom($σ$)
$\sigma, \ell.(i) \rightarrow \sigma, v$	$\sigma, \ell . (i) \leftarrow v \rightarrow \sigma \left[\ell + i \mapsto v \right]$], () σ , [t, v ₁ , v ₂]	$\rightarrow \sigma \left[\ell \mapsto t, v_1, v_2 \right], \ell$
$\begin{array}{l} \text{CONSTR-FULL-PAR} \\ \text{sta}(e_1) = \end{array}$	e_1' sta $(e_2) = e_2'$	$\begin{array}{l} \text{CONSTR-PAR-LEFT} \\ \texttt{sta}(e_1) = e_1' \end{array}$	$sta(e_2)$ undefined
$\sigma, \{ t_{ }, e_1, e_2 \} \rightarrow \sigma,$	$let x_1 = e'_1 in$ $let x_2 = e'_2 in$ $[t_{ }, await(x_1), await(x_2)]$	$\sigma, \{ t_{ }, e_1, e_2 \} \rightarrow$	$\begin{array}{l} \operatorname{let} x_1 = e_1' \operatorname{in} \\ \sigma, \operatorname{let} x_2 = e_2 \operatorname{in} \\ \left[t_{ }, \operatorname{await}(x_1), x_2 \right] \end{array}$
$\begin{array}{l} \text{CONSTR-PAR-RIGHT} \\ \texttt{sta}(e_2) = e_2' \end{array}$	$sta(e_1)$ undefined	CONSTR-SEQ-LEFT sta $(e_1) = e'_1$	$sta(e_2)$ undefined
$\sigma, \{ t_{ }, e_1, e_2 \} \rightarrow$	let $x_2 = e'_2$ in σ , let $x_1 = e_1$ in $[t_{ }, x_1, await(x_2)]$	$\sigma, \{ t_{\text{seq}}, e_1, e_2 \} \to \sigma,$	let $x_2 = e_2$ in let $x_1 = e'_1$ in [t_{seq} , await (x_1) , x_2]
	CONSTR-BASE other	cases	
	$\overline{\sigma, \{t, e_1, e_2\}} \rightarrow \sigma, let x_1 = e_1$	$1 \text{ inlet } x_2 = e_2 \text{ in } [t, x_1, x_2]$	x ₂]

Fig. 19. Sequential semantics for constructors $-\sigma$, $e \rightarrow \sigma'$, e' – Extends Fig. 8

To specify their reduction, we need to capture expressions of shape $C_n^{tl}[\operatorname{await}(e_1) | \cdots | \operatorname{await}(e_n)]$ where C_n^{tl} is a tail context (see Fig. 13a). In such expressions, all execution path lead to an await. Remarkably, such tail contexts "commute" with awaits, which is crucial to formalize which positions in a constructor are indeed task creations. To capture this, we define a partial function over expressions: strip-tail-awaits (shortened to sta), defined in Fig. 18, which removes all tail awaits. Intuitively, this removal separates synchronisation from computation: consider $e' = \operatorname{sta}(e) = C_n^{tl}[e_1 | \cdots | e_n]$, then e is semantically equivalent to first binding fut to e', launching a parallel task, and then evaluating await(fut). Conversely, if sta(e) is undefined, expression e must be treated sequentially.

We can define reduction for constructors on Fig. 19. Stores are unchanged and evaluation contexts are easily extended in Fig. 17. Read (CONSTR-READ) and write (CONSTR-WRITE) accesses to constructors are defined by memory accesses. The final constructor allocation (CONSTR-ALLOC) is immediate by store access. Recall that { t, e_1 , e_2 } denotes syntactic constructors, which can have sub-expressions, while [t, x_1 , x_2] is runtime syntax for a constructor allocation and can not have sub-expressions. The main challenge is in defining the evaluation of sub-expressions in a constructor depending on its tag. While the rules are numerous, they simply enumerate all possible node tags and positions for holes (one or two holes, left or right). We use let-bindings to represent evaluation order: for *parallel* constructors $t_{||}$, we first evaluate the "parallel" fields (i.e., such that sta is defined), then the sequential fields (similarly to tmc), and then perform the **await** for the parallel fields just before allocation; conversely for *sequential* constructors t_{seq} : non-parallel fields are computed first, thus preserving evaluation order.

Source language. The concurrent reduction of the source semantics is unchanged.

Target language. Our target semantics crucially relies on the fact that we can *resolve* a future only if it has no holes. To model this, we assume a magic function holes : $\ell \to \mathbb{N}$ (see Section 4 for an efficient implementation) which returns the number of locations set to \checkmark in the tree rooted at ℓ .

Sequential Semantics. We add rules for **refine** in the sequential semantics. In a transformed program, the second argument of a refine can be one of four cases, each described by one of the rules in Fig. 20: either a block with a parallel tag and two holes (REFINE-PAR), a block with a sequential

REFINE-PAR $\sigma [d] = \checkmark$ $\ell, \ell + 1, \ell + 2 \notin \operatorname{dom}(\sigma)$	REFINE-SEQ-RIGHT $\sigma [d] = \ell \ell, \ell+1, \ell+2 \notin \operatorname{dom}(\sigma)$
$\sigma' = \sigma \left[d \mapsto \ell, \ell \mapsto t_{ }, \ell + 1, \ell + 2 \mapsto \checkmark \right]$	$\sigma' = \sigma \left[d \mapsto \ell, \ell \mapsto t_{\text{seq}}, \ell + 1 \mapsto v_1, \ell + 2 \mapsto \checkmark \right]$
$\overline{\sigma, \texttt{refine}(d, [t_{ }, \blacksquare, \blacksquare])} \rightarrow \sigma', (\ell+1, \ell+2)$	$\sigma, \texttt{refine}(d, [t_{seq}, v_1, \blacksquare]) \rightarrow \sigma', (\ell+2)$
$ \begin{array}{l} \text{REFINE-SEQ-LEFT} \\ \sigma \left[\ d \ \right] = \nearrow \ell, \ell+1, \ell+2 \notin \text{dom}(\sigma) \\ \sigma' = \sigma \left[\ d \mapsto \ell, \ell \mapsto t_{\text{seq}}, \ell+1 \mapsto \checkmark, \ell+2 \mapsto \end{array} $	v_2 FILL $\sigma [d] = \checkmark$
$\sigma, \texttt{refine}(d, [t_{\texttt{seq}}, \blacksquare, v_2]) \rightarrow \sigma', (\ell+1)$	$\sigma, \texttt{refine}(d, v) \stackrel{\varepsilon}{\to} \sigma \left[d \mapsto v \right], ()$
Fig. 20. Sequential	semantics for refine
CHAIN-UNRESOLVED $Fut [fut] = \mathfrak{U} d'$	$C \neq \Box$ tid' fresh
$\frac{c}{\sigma, Task \sqcup (tid \mapsto C[forward(fut, d)]), Fut \stackrel{\varepsilon}{\to} \sigma, Ta.}$	$sk \sqcup (tid \mapsto C[()]) \sqcup (tid' \mapsto \mathbf{forward}(fut, d)), Fut$
CHAIN-RESOLVED	$[tt] = \Re v$
$\sigma, Task \sqcup (tid \mapsto \mathbb{C}[\texttt{forward}(fut, d)]), Fu$	$t \to \sigma [d \mapsto v], Task \sqcup (tid \mapsto C[())), Fut$
dps-call $(f \mapsto \lambda_{dps} \delta, x. e) \in F$	<i>'un fut', d, tid</i> ' fresh
$\sigma, Task \sqcup (tid \mapsto \sigma [d \mapsto \swarrow], Task \sqcup (tid \mapsto C[fut']) \sqcup (tid \mapsto C[fu$	$ c[(f v)]), Fut \to id' \mapsto e[\delta \mapsto d][x \mapsto v]), Fut [fut' \mapsto \mathfrak{U} d] $
Tail-Call $(f \mapsto \lambda_{dps} \delta, x. e)$	\in Fun tid' fresh
$ \begin{array}{ccc} \sigma, & & \sigma, \\ Task \sqcup (tid \mapsto \mathbb{C}[(f \#_t (d, v))]), & \to & Task \sqcup \\ Fut & & Fut \end{array} $	$ (tid \mapsto C[()]) \sqcup (tid' \mapsto e[x \mapsto v][\delta \mapsto d]), $
FUTURE-RESOLVE Fut [fut] = $\mathfrak{U} d$ holes(d) = 0 $\sigma [d] = a$) PRUNE-TASKS
$\sigma, \mathit{Task}, \mathit{Fut} \xrightarrow{\varepsilon} \sigma \setminus d, \mathit{Task}, \mathit{Fut} [\mathit{fut} \mapsto \Re v]$	$\overline{\sigma, Task} \sqcup (tid \mapsto ()), Fut \to \sigma, Task, Fut$

Fig. 21. Global Target reduction $-\sigma$, *Task*, *Fut* $\rightarrow \sigma'$, *Task'*, *Fut'* - Extends Fig. 12. Greyed out rules are similar to tma, with the modified parts highlighted in *violet*.

tag and one hole (REFINE-SEQ-LEFT or REFINE-SEQ-RIGHT), or a value (FILL). Except for that last case, **refine** performs the allocation when every field of the block is either a value or a hole—setting the location to \checkmark in that case—, sets the destiny argument to the corresponding value, and returns all locations initialized at \checkmark .

Concurrent Semantics. To define the target concurrent semantics with constructors, we extend the one defined for Tail-Modulo-Await, but with additional contexts. The language constructs that previously appeared exclusively in tail-position (and therefore could only be reduced under empty evaluation context), can now appear under context: before a sequence, when evaluating asynchronous fields of transformed constructor allocations. Therefore, we add contexts in the reduction rules for these constructs (namely **forward** and $(\cdot \#_t \cdot))$ —for clarity, we write the rule in grey and highlight the new additions in *violet*. Furthermore, these rules cannot kill the current

task any more, and must instead return (). Instead, the new rule PRUNE-TASKS kills the threads which "have nothing left to do". Similarly, since **forward** can now appear under non-empty context, it shouldn't block its thread evaluation if the future to forward is still unresolved. Rule CHAIN-UNRESOLVED delegates the forward to a dedicated task, until it can be evaluated once the future is resolved.

Filling a destiny does not necessarily resolve a future immediately. The rule FUTURE-RESOLVE sets a future to \Re in *Fut* when it has no holes left; it also removes its associated destiny from the store. This last removal is not necessary, but simplifies proofs by keeping stores consistent between source and target configurations (these locations do not exist in the source).

3.3 Transformation

Finally, Fig. 22 describes the Tail-Modulo-Cons-Await transformation, extending Tail-Modulo-Await with constructors. It uses the sta function (see Fig. 18) to determine which positions end in awaits or not, and re-order the expressions accordingly. The bulk of the constructor is always replaced by an allocation and a **refine** to obtain the appropriate destinies. Note that, in the t_{seq} case, the transformation is not recursively called on sub-expressions as there is only one tail-position.

$\mathcal{D}\llbracket e rbracket_{\delta}$::= .		
$\mathcal{D}[\![\{ t_{ },e_1,e_2\}]\!]_\delta \coloneqq$	$\begin{split} & \texttt{let}\delta_1,\delta_2=\texttt{refine}(\delta,[t_{ },\blacksquare,\blacksquare])\texttt{in}\\ & \mathcal{D}[\![e_2]\!]_{\delta_2};\mathcal{D}[\![e_1]\!]_{\delta_1} \end{split}$	(when $sta(e_2)$ defined, and $sta(e_1)$ undefined)
$\mathcal{D}[\![\{t_{ },e_1,e_2\}]\!]_\delta \coloneqq =$	$\begin{split} & \texttt{let}\delta_1,\delta_2=\texttt{refine}(\delta,[t_{ },\blacksquare,\blacksquare])\texttt{in}\\ & \mathcal{D}[\![e_1]\!]_{\delta_1};\mathcal{D}[\![e_2]\!]_{\delta_2} \end{split}$	(otherwise)
$\mathcal{D}\llbracket \{ t_{\text{seq}}, e_1, e_2 \} \rrbracket_{\delta} ::=$	$\begin{split} & \texttt{let}\delta_1 = \texttt{refine}(\delta,[t_{\texttt{seq}},\blacksquare,e_2])\texttt{in}\\ & \mathcal{D}[\![e_1]\!]_{\delta_1} \end{split}$	(when $sta(e_1)$ defined)
$\mathcal{D}\llbracket \{ t_{\text{seq}}, e_1, e_2 \} \rrbracket_{\delta} ::=$	$\begin{split} & \texttt{let}\delta_2 = \texttt{refine}(\delta,[t_{\texttt{seq}},e_1,\blacksquare])\texttt{in}\\ & \mathcal{D}[\![e_2]\!]_{\delta_2} \end{split}$	(when $sta(e_1)$ undefined)

Fig. 22. The Tail-Modulo-Cons-Await transformation $-\mathcal{D}[\![\cdot]\!]$ – Extends Fig. 13

3.4 Correctness of tmca

We extend the proof of correctness introduced in Section 2.5 for tma. The simulation relation essentially extends the previous one. However, the amount of administrative details we need to track of gets hairy: we therefore refer the interested reader to the Appendix C for details, and favour here a graphical representation to convey the intuition of the proof.

The management of the stuttering aspect of the backward simulation is slightly more involved than in the tma case: the FILL reduction step can be stuttering in some contexts, but not all. The identification of ε steps is therefore now contextual, rather than purely hardcoded in the semantics.

3.4.1 Definition of \succeq . Figs. 23 and 24 introduce graphical representations for respectively source and target configurations in relation: we take this informal graphical depiction as definition for \succeq . As for tma, the main thread is identical. However, because we allocate constructors in the store at different points in execution, and also because destinies are now actual memory locations, the stores cannot be synchronised: nonetheless, the store for a source configuration is always included in the target store (σ), which we single out on the left of the representation.

The remaining threads of the configurations are still divided in matching pairs of pools, but with additional structure, represented as boxes. Each box is associated to exactly one future. When, in





 $T_i \underset{fut_i}{\sim} \approx_{\ell+i} \sigma_i, T'_i, i \in \{1, 2\}$

Fig. 25. The last case of the relation for parallel constructors

the target, this future fut_i maps to an unresolved destiny d_i , it owns the share of store σ_i , which holds the tree of constructors "in construction" on which the resolution of d_i depends. As in tma,

the source and target threads pools are related via a relation $T_s^{(i)} = \frac{\sigma_i}{fut_i} \approx \frac{\sigma_i}{d_i} \frac{\sigma_i}{T_t^{(i)}}$, annotated with the future and destiny being resolved by the box. Otherwise, if the future is resolved (case fut_j), the box carries the remaining set of lingering anonymous tasks on the target side, and resolved

futures on the source side.

Recall the tree-like memory layout at runtime presented in Fig. 2b. This structure, internal to a "box", therefore needs to be captured in \approx . It therefore reflects two aspects of the current state of the structure in memory. First, at each node, it recursively ensures that the subtrees are related. Furthermore, it must keep track of the status of the allocation of a given node as one of three successive states: whether none, one, or two subtasks have been launched yet.⁵ By identifying this last case, we can identify the operation in the target filling the final hole, and against which the source program can perform the corresponding allocations, resynchronising both configurations.

Fig. 25 illustrates the case of a *parallel* constructor with both fields asynchronous, in the situation where both subtasks have already been launched. On the source side, we have $\alpha_i = \text{await}(fut_i)$ if fut_i is not resolved yet, and $\alpha_i = v_i$ where $\sigma_i(l+i) = v_i$ otherwise.

3.4.2 *Proof sketch.* First, observe that similarly to the tma case, the \approx relation allows us to frame out part of the configurations when proving the backward simulation. We can therefore strengthen the statement of simulation we establish, specifying the shape of the configurations we reach.

Consider
$$c_s = \frac{\sigma_s}{Task_s \sqcup T_s} \approx \frac{\sigma_t \sqcup s_t}{Fut_t} = c_t$$
, s.t. $T_s \int_{fut} a_s s_t$, T_t , then:

⁵In the case of sequential constructors, at most one subtask is launched.



Fig. 26. Sub-configurations before a synchronisation step. Idle denotes a possibly empty list of threads all reduced to ()

LEMMA 3.1. If we have

$$c_t \rightarrow \begin{array}{c} \sigma'_t \sqcup s'_t, \\ Task_t \sqcup T'_t (\sqcup T_t^{\text{new}})^? \\ Fut_t (\sqcup F_t^{\text{new}})^? \end{array}$$

then there exist T'_s, σ'_s (and maybe T^{new}_s) such that

$$c_s \rightarrow^* \quad \frac{\sigma'_s}{Task_s} \sqcup T'_s (\sqcup T^{\text{new}}_s)^? \quad and \ T'_s \quad _{fut} \approx_d s'_t, T'_t$$

Moreover, we can characterise more finely T'_s in the case where we finish computing the tasks associated with *d*. Indeed, if in c_t we have holes(d) = 1, and if the step under consideration in the target fills this last hole in the subtree of *d*, then $T'_s = (fut \mapsto v) \sqcup (fut_i \mapsto v_i)$, with $v = s'_t(d)$. This situation arises each time a destiny is introduced, but most importantly at the end of the evaluation of constructors. In the target, as soon as all values are filled, we are done with evaluating the constructor, and its future can be resolved and awaited. Therefore, the last step filling a destiny needs to be matched, in the source, by all allocations and administrative steps required to get to a configuration where the block can be allocated, and finally performing said allocation.

This lemma relies in particular on the fact that for related configuration, for any destiny *d* such that holes(d) = 0 associated to related pools $T_s |_{fut} \approx_d \sigma_t T_t$, we have that *fut* is resolved and $T_s \supseteq (fut \mapsto v)$ where $v = \sigma_t(d)$.

Backward simulation: illustrative case. Finally, we illustrate the case from Fig. 25 in the proof of lemma 3.1 conducted by induction on the hypothesis $T_s _{fut} \approx_d s_t, T_t$.

Fig. 26 depicts the situation, where the target is at the bottom. The step it takes fills the last location under destiny *d*. We know that holes(d) = 1, with said hole necessarily in either the left or right sub-tree. Let's assume it is in the right sub-tree, rooted at $\ell + 2$, with the other one completely resolved. By invariant, we have $\alpha_1 = v_1$, where $v_1 = \sigma_t [\ell + 1]$ and $T_2 \int_{fut_2} \approx_{\ell+2} \sigma_2, T'_2$, with $holes(\ell + 2) = 1$. By induction hypothesis, the step can be matched in the source, until fut_2 is resolved, to value v_2 with v_2 the value at $\ell + 2$ in the target configuration, represented on Fig. 26 by the first sequence of transition at the top of the figure.

At this point, the $await(fut_2)$ remaining in the constructor can be resolved, and the block allocated. As this block is now allocated in the store on the source side, it is already accounted for in the "synchronised" logical view of the target store, and the corresponding mappings are removed from the σ_i sub-parts.

We finally reach a "value" state of the \approx relation: all source threads are resolved futures, all target thread idle (read reduced to () or already pruned), with a single mapping from the current destiny to its value in the sub-part of the store; which satisfies the lemma, and the invariant.

4 Practical considerations

In this section, we showcase an implementation and evaluation of our code transformation. We first describe a refined version of Tail-Modulo-Cons-Await which avoids creating unnecessary destinies. We then describe our prototype OCaml implementation, and evaluate it in practice.

4.1 Optimized Code Transformation

The tmca code transformation proposed in Section 3.3 is quite inefficient in the presence of nested constructors. This inefficiency, already noted by Allain et al. [3], is best demonstrated on an example. Let us consider the source code below on the left, which builds a value using constructors $A_{||}$ and $B_{||}$ (both in a parallel fashion). It will result in the target code on the right.

$$\{A_{||}, \{B_{||}, \mathsf{await}(e_1), e_2\}, \mathsf{await}(e_3)\} \Rightarrow \begin{bmatrix} \mathsf{let}\,\delta_l, \delta_r = \mathsf{refine}(\delta, [A_{||}, \blacksquare, \blacksquare]) \, \mathsf{in} \\ \mathsf{forward}(e_3, \delta_r); \\ \mathsf{let}\,\delta_{ll}, \delta_{lr} = \mathsf{refine}(\delta_l, [B_{||}, \blacksquare, \blacksquare]) \, \mathsf{in} \\ \mathsf{forward}(e_1, \delta_{lr}); \, \mathsf{refine}(\delta_{ll}, e_2) \end{bmatrix}$$

While semantically fine, the destiny δ_l is unnecessary: it is immediately filled by the allocated block $[B_{||}, \blacksquare, \blacksquare]$. As we show in next section, creating a destiny in practice is a fairly costly operation: it requires cooperation from the scheduler, and, in a language like OCaml, induces a write-barrier when filled. It would be better to set the value directly. As such, and following a similar idea from Allain et al. [3], we emit the following code:

let
$$\delta_{ll}, \delta_{lr}, \delta_{r} = \text{refine}(\delta, [A_{||}, [B_{||}, \blacksquare, \blacksquare], \blacksquare]) \text{ in } forward(e_3, \delta_r); forward(e_1, \delta_{lr}); \text{refine}(\delta_{ll}, e_2)$$

This new code builds a single more complex "value with holes", and has only one call to **refine**. It nonetheless preserves the execution order of sub-expressions e_1 , e_2 and e_3 induced by the naive transformation.

The detailed transformation is shown in Fig. 27. Only the base case of the $\mathcal{D}[\![\cdot]\!]_{\delta}$ changes, by calling a collecting function, $\mathcal{D}_{cons}[\![\cdot]\!]$, which accumulates two things: a list of expressions to be evaluated (either synchronously or asynchronously), and a "value with holes", made of constructors and constants. This transformation ensures that only one call to **refine** is done per tail-position. Except for this optimisation, our prototype implementation closely follows the formalisation we presented in the previous sections.

4.2 Proof-of-concept Implementation

We provide a proof-of-concept implementation in OCaml. The implementation is available in the joint anonymous material (and will be available as free software online for publication). The syntactic transformation is implemented as a PPX syntax extension on OCaml's syntax. This is not a battle-ready implementation, which should be implemented directly inside a compiler, but is sufficient to test our transformation on real-world examples. The runtime is implemented as a thin layer atop an existing OCaml concurrency library.⁶ An overview of the runtime API is provided on Figs. 28 and 30, and the actual translation of map on trees on Fig. 29.

⁶Most existing library would satisfy our needs, we use Picos (https://github.com/ocaml-multicore/picos) for its clear semantics and to allow testing with different schedulers

$$\mathcal{D}\llbracket \mathcal{C}_{n}^{tl}\llbracket e_{1} | \cdots | e_{n}
bracket \rrbracket_{\delta} \coloneqq \mathcal{C}_{n}^{tl}\llbracket \mathcal{D}\llbracket e_{1}
bracket_{\delta} | \cdots | \mathcal{D}\llbracket e_{n}
bracket_{\delta}
bracket$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathcal{D}\llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathcal{D}\llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathcal{D}\llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathcal{D} \llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathcal{D} \llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D} \llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathbb{D} \llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D} \llbracket e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D} \blacksquare e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathbb{D} \blacksquare e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathbb{D} \blacksquare e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \coloneqq \mathbb{D} \blacksquare e_{n} \rrbracket_{\delta}$$

$$\mathcal{D}\llbracket e_{n} \rrbracket_{\delta} \vDash \mathbb{D} \blacksquare e_{n} \blacksquare e_{n$$

Fig. 27. Optimized Tail-Modulo-Cons-Await transformation for nested constructors $-\mathcal{D}[\![\cdot]\!]$ – Extends Fig. 22

```
1 module Loc : sig
                                                       1 (** Launch a new task. *)
2 (** Location inside an OCaml memory block,
                                                       2 val fork : (unit -> unit) -> unit
   composed of a value and an offset. *)
3
                                                       3
4 type 'h t = { pointer : Obj.t ; offset : int }
                                                       4 module Destiny : sig
                                                       5 (** A Destiny with a hole of type ['a] *)
6 (** An heterogeneous list of locations *)
                                                       6 type 'a t = {
7 type 'a list =
                                                          counter : int Atomic.t; (* Number of holes *)
                                                       7
                                                         resolve : unit -> unit; (* Top-level resolve *)
8 | [] : unit list
                                                       8
    | (::) : 'h t * 'h2 list -> ('h * 'h2) list
                                                       9
                                                          set : 'a -> unit; (* Set the current destiny *)
9
10
                                                      10 }
11 (** A dummy value to insert into blocks. *)
                                                      12 (** An heterogeneous list of destinies *)
12 val dummy : 'a
                                                      13 type 'a list =
14 (** [Loc.mk o i] is the location in value [o]
                                                      14 | [] : unit list
                                                      15 | (::) : 'h t * 'h2 list -> ('h * 'h2) list
      at offset [i]. [o] must b a pointer to a
15
16
      block. *)
                                                      16
17 val mk : 'a -> int -> 'b t
                                                      17 (** [mk ()] returns a synchronised destiny and
18 end
                                                      18
                                                            future. The future will be fulfilled only
                                                      19
                                                            when all holes in the destiny are filled. *)
          Fig. 28. Runtime library for locations
                                                      20 val mk : unit -> 'a t * 'a Future.t
                                                      21
1 let rec map$dps d4 f t = match t with
                                                      22 (** [refine d v hs] fills the destiny [d] with a
   | Leaf -> Destiny.refine d4 Leaf []
                                                            value [v] with holes [hs]. It returns a list
2
                                                      23
    | Node (v, tl, tr) ->
3
                                                      24
                                                            of destinies corresponding to each holes.
      let v5 = Node(Loc.dummy,Loc.dummy,Loc.dummy) in25
 4
                                                            [d] should not be used again. *)
                                                      26 val refine : 'b t -> 'b -> 'a Loc.list -> 'a list
      let [d6; d7; d8] =
5
        Destiny.refine d4 v5
6
                                                      27
7
          [Loc.mk v5 0; Loc.mk v5 1; Loc.mk v5 2]
                                                      28 (** [forward c d] will fill the destiny [d] with
                                                            the value of the asynchronous computation [c]
8
      in
                                                      29
      fork (fun () -> map$dps d8 f tr);
                                                            when it is finished. *)
9
                                                      30
                                                      31 val forward : 'a Future.t -> 'a t -> unit
10
      fork (fun () -> map$dps d7 f tl);
      Destiny.refine d6 (f v) []
                                                      32 end
                                                               Fig. 30. Runtime library for destinies
13 let map f l =
14 let d, c = Destiny.mk () in
```

Fig. 29. Real translation of parallel map from Fig. 1

, Vol. 1, No. 1, Article . Publication date: April 2025.

fork (fun () -> map\$dps d f l);

15 16

21

The main difference between our implementation and the formalism presented so far is that hole creation, memory locations and destinies were previously fused into the near-magical refine operation. These operations are now split to allow finer control. To pre-allocate a constructor, we first create a value from a constructor, with holes implemented by a dummy value Loc.dummy. We then take locations for each of its holes, using the Loc.mk function, and gather them in an heterogeneous Loc.list. We then use Destiny.refine with the explicit list of locations to obtain the list of destinies (as a Destiny.list). This is demonstrated in Line 4 to 7 in Fig. 29.

The implementation of Loc relies on details of the OCaml memory model and some unsafe internal API (Obj). Destiny, in turn, only uses regular OCaml with concurrency. Destinies are a triplet composed of a set to a location, a reference to a shared atomic counter representing the current number of holes in the full structure, and a resolve function which marks the top-level future as resolved. The functions **forward** and **refine** decrement and increment the counter when holes are added or removed from the structure. The key trick here is that, as soon as the atomic counter hits zero, there are no more holes in the structure and the counter can never rise up again. We can then resolve the top-level future immediately.

To implement DPS calls (Line 9, 10 and 15 in Fig. 29), we use a regular **fork** function which takes a callback containing the actual call. The call is syntactically a tail-call, ensuring that OCaml's compiler uses the desired call convention. Finally, for the purpose of this implementation we consider all constructors to be parallel: indeed, as described in Section 3, sequential constructors can be emulated by let-bindings. The rest of the implementation follows our formalism closely.

4.3 Experimental Evaluation

We now evaluate our full code transformation experimentally. Our evaluation setup is a Thinkpad T470 equipped with i5-7200U with 4 cores and 11GiB of memory. All the implementations evaluated are provided in Appendix B. We explore the following scientific questions:

What is the effect of Tail-Modulo-Cons-Await on memory consumption? Evaluating stack and scheduler space precisely is quite difficult. We settle for a lesser goal: to evaluate the total live memory during the execution of a program. More specifically, we consider the execution of several variants of List.map on a list of 50000000 elements. The implementations (see Fig. 36 in Appendix B) are: a naive non-tail-rec List.map, a version using Tail-Modulo-Cons, an asynchronous version with a sequential function parameter, an asynchronous version with a sequential constructor, and an asynchronous version with a parallel constructor. For more details, see Appendix A. The overall memory footprint of the overall trace is shown in Fig. 31. There is only one run, since memory footprint is deterministic. The List.map runs out of stack before allocating anything, and crashes. All the other functions complete their execution while consuming the exact same amount of memory, which corresponds exactly to an allocation of the output list. Neither stack nor scheduler space is used in this execution. The synchronous version using Tail-Modulo-Cons runs faster than the asynchronous version, which is expected since it doesn't involve a scheduler. This confirms that Tail-Modulo-Cons-Await behaves as advertised and avoids superfluous stack and scheduler space.

What is the effect of Tail-Modulo-Cons-Await on performances and parallelism? To evaluate the effect of Tail-Modulo-Cons-Await on time, we run several variants of map on randomly generated trees of growing size. Figure 32 shows the *time per elements*, i.e., the total time divided by the number of elements, for Tree.map f. We consider two distinct f functions which are applied at each Node: a very fast function on the left (multiply by 2) and a slow function (sleep a random amount between 0 and 50ms). The evaluated map implementations (see Fig. 37 in Appendix B) are as follows: map_tmc is a synchronous sequential version which uses Tail-Modulo-Cons (with only one recursive tail call); all the others are asynchronous, with and without Tail-Modulo-Cons-Await: map_seq uses a



Order of execution:

- Creation of the list of 50000000 elements
- map_naive: Not tail-recursive
- map_tmc: With Tail-Modulo-Cons
- map_async_sync: Synchronous function f
- map async seq: Sequential constructor ::
- map async par: Parallel constructor ::

Fig. 31. Memory measurement of several version of List.map. Global Memory footprint vs. Time elapsed since start of execution in seconds.



Fig. 32. Time measurement for several versions of Tree.map. Lower is better.

sequential semantics Node (as in Fig. 15b); map_par_full uses a full parallel semantics (as in Fig. 15c); map_par_width also uses a parallel semantics, but applies f before any recursive calls, thus bounding the parallelism by width (see Section 1). The function compiled with Tail-Modulo-Cons-Await are marked with _tmca and are represented using dashed lines. The scheduler is allowed to spawn tasks on all four available processors. All time measurements are obtained with repeated runs until a correlation of $r^2 > 0.8$ is obtained.

First, we can observe that the parallel semantics is, indeed, parallel! Regardless of the function f used, the parallel versions are exactly 4 times faster than the sequential versions.

Second, let us look at raw performance. With the fast f function, we can observe that Tail-Modulo-Cons-Await consistently improves performances, especially on large trees, although never to the point of equating the performance of the synchronous version. The cost of the slow function completely dominates any of the costs introduced by the implementation of map, and the cohort are separated into two groups: the sequential and the parallel versions.

5 Related Work

5.1 Tail recursion

The Tail-Modulo-Cons transformation is as old as functional programming, seeing birth in the LISP community in the early 70th. During this decade, two academic publications describe the

transformation, first Risch for the REMREC system [15], then Friedman and Wise [11] over a pure LISP with cons compiled down to a machine language.

After having remained folklore for half a century, the transformation has seen revived interest in the community through essentially two independent pieces of work. Leijen and Lorenzen have implemented tmc for Koka, a strongly typed language with effect types and handlers. In doing so, they have rephrased the transformation in an equational style, generalised it to work with an abstract notion of *contexts*, and adapted it to the linear typing discipline followed by Koka. Independently, Allain et al. [3] have also implemented tmc, but in the OCaml compiler. In doing so, they have formalised the optimisation over a core calculus from which we took inspiration for designing our own, and have furthermore mechanised its proof of soundness using a variant of Simuliris [12], a simulation technique based on separation logic.

5.2 Asynchronous Programming

Future and Promises. Futures are limited in the sense that they must be fulfilled by the task that was associated with their creation. This limitation gives guarantees on the fact that futures are eventually fulfilled. On the contrary, promises associate their fulfilment with a handler given to the programmer. Hence, they do not suffer from this limitation, but it cannot be ensured that a promise is resolved exactly once [1]. The fact that promises allow for the optimisation of scheduler space compared to futures was already the key idea of the **forward** construct [8]. In this work, **forward** promises are only used internally for optimisation purposes, as we do in the work presented in this paper. Indeed, the forward optimisation is only valid when applied to an asynchronous call in tail position. The present article can therefore be seen as an extension of [8] with more automation, added parallelism, and more general application setting.

Dataflow synchronisation on futures [9] is a paradigm that makes it impossible to observe chains of futures: chains are automatically traversed upon each synchronisation (synchronisation is insensible to successive asynchronous tail calls). Chappe et al. [6] showed that with dataflow futures, it is safe to optimise tail-calls with promises, and that return has the same semantics as **forward** in this setting. This is another setting that makes it possible to optimise scheduler space for futures (also by replacing futures by promises) and in which the tmaa approach should be applicable to make the approach more general. Unfortunately dataflow synchronisation on futures implies a type system where 'a future unifies with ('a future) future. Such typing rules, where monadic flatten is an axiom, are particularly difficult to integrate in ML-like languages such as OCaml.

Other Asynchronous Paradigms. Futures are massively used in actors and active object languages [2, 7]: this family of languages thus offers a privileged application domain for all optimisations of asynchronous calls, notably the one we propose. The fact that, by nature, all calls to an actor create a future multiplies the number of futures, often with tail-calls, and hence makes future optimisation crucial in these languages. Indeed **forward** was created in an active object language, and the present work would clearly benefit languages with prominent Actors libraries [2, 13, 14], including the recently designed OCaml active object library [4].

Our parallel constructors could also be compared to the parT construct of the Encore language [10] or the par construct of Haskell [18, 19]. Both combinators generate parallelism and allow the coordination of parallel tasks. Rather, our vision here is to create parallelism automatically when using data constructors. We believe the integration of data production and parallelism fits well with the notion of futures and promises, but approaches based on parallel combinator features richer synchronisation patterns compared to our work. We might study if it would be possible to integrate a few richer coordination patterns in our context, which could for example enable some form of pipelining when performing two asynchronous maps in a row (even over complex data structures).

References

- [1] Erika Ábrahám, Immo Grabe, Andreas Grüner, and Martin Steffen. 2009. Behavioral interface description of an object-oriented language with futures and promises. *The Journal of Logic and Algebraic Programming* 78, 7 (2009), 491 518. doi:10.1016/j.jlap.2009.01.001 The 19th Nordic Workshop on Programming Theory (NWPT 2007).
- [2] Gul Agha. 1986. Actors: a model of concurrent computation in distributed systems. MIT Press.
- [3] Clément Allain, Frédéric Bour, Basile Clément, François Pottier, and Gabriel Scherer. 2025. Tail Modulo Cons, OCaml, and Relational Separation Logic. Proceedings of the ACM on Programming Languages 9, POPL (Jan. 2025), 2337–2363. doi:10.1145/3704915
- [4] Martin Andrieux, Ludovic Henrio, and Gabriel Radanne. 2024. Active Objects based on Algebraic Effects. In Active Object Languages: Current Research Trends. Lecture Notes in Computer Science, Vol. 14360. 3–36. doi:10.1007/978-3-031-51060-1_1
- [5] Henry. G. Baker Jr. and Carl Hewitt. 1977. The Incremental Garbage Collection of Processes. In Proc. Symp. on Artificial Intelligence and Programming Languages. New York, NY, USA, 55–59.
- [6] Nicolas Chappe, Ludovic Henrio, Amaury Maillé, Matthieu Moy, and Hadrien Renaud. 2021. An Optimised Flow for Futures: From Theory to Practice. CoRR abs/2107.07298 (2021). arXiv:2107.07298 https://arxiv.org/abs/2107.07298
- [7] Frank de Boer, Vlad Serbanescu, Reiner Hähnle, Ludovic Henrio, Justine Rochas, Crystal Chang Din, Einar Broch Johnsen, Marjan Sirjani, Ehsan Khamespanah, Kiko Fernandez-Reyes, and Albert Mingkun Yang. 2017. A Survey of Active Object Languages. ACM Comput. Surv. 50, 5, Article 76 (Oct. 2017), 39 pages. doi:10.1145/3122848
- [8] Kiko Fernandez-Reyes, Dave Clarke, Elias Castegren, and Huu-Phuc Vo. 2018. Forward to a Promising Future. In Conference proceedings COORDINATION 2018.
- [9] Kiko Fernandez-Reyes, Dave Clarke, Ludovic Henrio, Einar Broch Johnsen, and Tobias Wrigstad. 2019. Godot: All the Benefits of Implicit and Explicit Futures. In 33rd European Conference on Object-Oriented Programming (ECOOP 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 134), Alastair F. Donaldson (Ed.). Schloss Dagstuhl– Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2:1–2:28. doi:10.4230/LIPIcs.ECOOP.2019.2 Distinguished artefact.
- [10] Kiko Fernandez-Reyes, Dave Clarke, and Daniel S. McCain. 2016. ParT: An Asynchronous Parallel Abstraction for Speculative Pipeline Computations. In Proc. 18th IFIP WG 6.1 Intl. Conf. on Coordination Models and Languages (COORDINATION 2016), Alberto Lluch-Lafuente and José Proença (Eds.). Vol. 9686. 101–120.
- [11] Daniel P Friedman and David S Wise. 1975. Unwinding stylized recursions into iterations. Comput. Sci. Dep., Indiana University, Bloomington, IN, Tech. Rep 19 (1975).
- [12] Lennard G\u00e4her, Michael Sammler, Simon Spies, Ralf Jung, Hoang-Hai Dang, Robbert Krebbers, Jeehoon Kang, and Derek Dreyer. 2022. Simuliris: a separation logic framework for verifying concurrent program optimizations. *Proc.* ACM Program. Lang. 6, POPL, Article 28 (Jan. 2022), 31 pages. doi:10.1145/3498689
- [13] Philipp Haller and Martin Odersky. 2009. Scala Actors: Unifying thread-based and event-based programming. *Theoretical Computer Science* 410, 2 (2009), 202 220. doi:10.1016/j.tcs.2008.09.019 Distributed Computing Techniques.
- [14] Einar Broch Johnsen and Olaf Owe. 2007. An Asynchronous Communication Model for Distributed Concurrent Objects. Softw. Syst. Model. 6, 1 (2007), 39–58. doi:10.1007/s10270-006-0011-2
- [15] Tore Risch. 1973. REMREC A program for Automatic Recursion Removal. https://api.semanticscholar.org/CorpusID: 60315470
- [16] Amir Shaikhha, Andrew W. Fitzgibbon, Simon Peyton Jones, and Dimitrios Vytiniotis. 2017. Destination-passing style for efficient memory management. In *Proceedings of the 6th ACM SIGPLAN International Workshop on Functional High-Performance Computing, FHPC@ICFP 2017, Oxford, UK, September 7, 2017, Phil Trinder and Cosmin E. Oancea* (Eds.). ACM, 12–23. doi:10.1145/3122948.3122949
- [17] Guy L. Steele Jr. 1977. Debunking the "expensive procedure call" myth or, procedure call implementations considered harmful or, LAMBDA: The Ultimate GOTO. In *Proceedings of the 1977 annual conference, ACM '77, Seattle, Washington,* USA, October 16-19, 1977, James S. Ketchel, Harvey Z. Kriloff, H. Blair Burner, Patricia E. Crockett, Robert G. Herriot, George B. Houston, and Cathy S. Kitto (Eds.). ACM, 153–162. doi:10.1145/800179.810196
- [18] Philip W. Trinder, Kevin Hammond, James S. Mattson Jr., Andrew S. Partridge, and Simon L. Peyton Jones. 1996. GUM: A Portable Parallel Implementation of Haskell. In *Proceedings of the ACM SIGPLAN'96 Conference on Programming Language Design and Implementation (PLDI), Philadephia, Pennsylvania, USA, May 21-24, 1996*, Charles N. Fischer (Ed.). ACM, 79–88. doi:10.1145/231379.231392
- [19] Philip W. Trinder, Kevin Hammond, Hans-Wolfgang Loidl, and Simon L. Peyton Jones. 1998. Algorithms + Strategy = Parallelism. J. Funct. Program. 8, 1 (1998), 23–60. doi:10.1017/S0956796897002967

A Yet Another example: a medley of map

We consider several asynchronous versions of the function map_async over lists, mapping a synchronous or asynchronous function over a list, returning an asynchronous computation itself. We first consider in Fig. 33 a version where only the recursive call to map_async is asynchronous, with the mapped function synchronous (map_async_sync : ('a -> 'b) -> 'a list -> 'b list future). Using Tail-Modulo-Cons-Await, we can transform it into the code in Fig. 33b.

```
1let%async rec map async sync f l =
                                                       1let rec map async sync dps d f l =
                                                          match l with
2 match l with
                                                       2
  | Nil -> Nil
                                                           | Nil -> d < Nil
3
                                                       3
  | Cons (x, xs) ->
                                                           | Cons (x, xs) ->
4
                                                       4
5
     Cons (f x, await (map async sync f xs))
                                                       5
                                                             let v = Cons (f x, \blacksquare) in
                                                       6
                                                             d < v;
                                                             fork_dps map async sync dps v.1 f xs
               (a) Source implementation
                                                        (b) Translation using tmca
                                                        The wrapping code for the initial destiny is identical to Fig. 5
```

Fig. 33. An asynchronous, sequential, map on lists with async/await, mapping a sequential function

When both calls are asynchronous (map_async : ('a -> 'b future) -> 'a list -> 'b list future), the function admits two semantics: *sequential* where an application of f must terminate before the next, or *parallel* which launches all tasks at once. These two semantics are differentiated in our formalisation by the annotation over constructor tags. Let us first consider the sequential version in Fig. 34. The Tail-Modulo-Cons-Awaittransformation affects this version in the same way as ths version in Fig. 33.

```
1let%async rec map async seq f l =
                                                     1let rec map async seq dps d f l =
2 match l with
                                                     2 match l with
   | Nil -> Nil
                                                         | Nil -> d < Nil
                                                     3
3
  | Cons (x, xs) ->
                                                         | Cons (x, xs) ->
4
                                                     4
     let y = await (f x) in
                                                           let y = await (f x) in
5
                                                     5
     Cons (y, await (map async seq f xs))
                                                           let v = Cons (y, ■) in
6
                                                     6
                                                     7
                                                           d ◀ v;
                                                           fork_dps map_async_seq_dps v.1 f xs
                                                     8
               (a) Source implementation
                                                       (b) Translation using tmca
                                                       The wrapping code for the initial destiny is identical to Fig. 5
```

Fig. 34. An asynchronous, sequential, map on lists with async/await

Finally, the parallel version is illustrated in Fig. 35. The transformed code in Fig. 35b uses **forward** for the call to f, rather than **fork_dps**, as f might be an external asynchronous function without a dps version.

```
1let%async rec map_async_par f l =
                                                 1let rec map_async_par_dps d f l =
2 match l with
                                                 2 match l with
   | Nil -> Nil
                                                    | Nil -> d < Nil
3
                                                 3
   | Cons (x, xs) ->
                                                    | Cons (x, xs) ->
4
                                                 4
     Cons (await (f x),
                                                      let v = Cons (■, ■) in
5
                                                 5
           await (map_async_par f xs))
                                                      d ◀ v;
6
                                                 6
                                                      let fut = f x in
                                                 7
                                                 8
                                                      forward v.0 fut;
              (a) Source implementation
                                                      fork_dps map_async_par_dps v.1 f xs
                                                 9
```

(b) Translation using tmca The wrapping code for the initial destiny is identical to Fig. 5

Fig. 35. An asynchronous, parallel, map on lists with async/await

26

B Supporting code for the Experimental Evaluation

```
1 let rec map naive f = function
2
   | [] -> []
   | h :: t -> f h :: map naive f t
3
5 let[@tail mod cons] rec map tmc f l =
6 match l with
7
   | [] -> []
8
   | x :: xs ->
     let y = f x in
9
10
    y :: map_tmc f xs
12 let%async rec map_async_sync f l =
13 match l with
14
   | [] -> []
   | x :: xs ->
15
16
     f x :: await (map_async_sync f xs)
18 let%async rec map_async_seq f l =
19 match l with
20
   | [] -> []
21
   | x :: xs ->
    let y = await (f x) in
22
    y :: await (map_async_seq f xs)
23
24
25 let%async rec map async par f l =
26 match l with
    | [] -> []
27
   | x :: xs ->
28
29
     await (f x) :: await (map_async_par f xs)
```

```
Fig. 36. Implementation of map on lists
```

```
1 type 'a tree = N of 'a * 'a tree * 'a tree | L
3 let[@tail mod cons] rec map tmc f t = match t with
   | L -> L
4
    | N (v , tl , tr ) ->
 5
      let y = f v in
 6
      let tl' = map_tmc f tl in
7
8
      N (y, tl', map tmc f tr)
10 let rec map_seq f t =
11
   async @@ fun () -> match t with
    | L -> L
12
   | N (v , tl , tr ) ->
13
14
      let y = f v in
     let tl' = await (map_seq f tl) in
15
     let tr' = await (map_seq f tr) in
16
     N (y, tl', tr')
17
18
19 let rec map_par_width f t =
   async @@ fun () -> match t with
20
    | L -> L
21
   | N (v , tl , tr ) ->
22
      let y = f v in
23
      let task tl = map par width f tl in
24
      let task_tr = map_par_width f tr in
25
      N (y, await task tl, await task tr)
26
28 let rec map_par_full f t =
29 async @@ fun () -> match t with
    | L -> L
30
    | N (v , tl , tr ) ->
31
      let task tl = map par full f tl in
32
      let task_tr = map_par_full f tr in
33
34
      N (f v, await task_tl, await task_tr)
35
36 let%async rec map seq tmca f t = match t with
    | L -> L
37
    | N (v , tl , tr ) ->
38
      let y = f v in
39
      let tl' = await (map_seq_tmca f tl) in
40
      N (y, tl', await (map_seq_tmca f tr))
41
42
43 let%async rec map par width tmca f t = match t with
    | L -> L
44
    | N (v , tl , tr ) ->
45
46
      let y = f v in
47
      N (y,
         await (map_par_width_tmca f tl),
48
         await (map par width tmca f tr))
49
50
51 let%async rec map_par_full_tmca f t = match t with
52
   | L -> L
53
    | N (v , tl , tr ) ->
      N (fv,
54
         await (map_par_full_tmca f tl),
55
         await (map_par_full_tmca f tr))
56
         Fig. 37. Implementation of map on trees
```

C Backward-simulation proofs

C.1 Tail-Modulo-Await backward simulation proof

Cases for the main thread. :

Since the main thread is unaffected by the transformation, its term is only comprised of source syntax. Thus, the only steps that can be taken are local steps, DPS-CALL and AWAIT-TARGET

STEP any local step is immediately closed if they don't affect the store. If they affect the store, then they affect it in the same way in source and target, and since stores are equal, we close immediately.

DPS-CALL:

 $\begin{array}{ll} Task_{s} \sqcup (\text{main} \mapsto C[(f v)]) & \cong_{D} & Task_{t} \sqcup (\text{main} \mapsto C[(f v)]), Fut_{t} \\ & \downarrow_{(f \mapsto \lambda_{async} x. e) \in Fun} \\ Task_{s} \sqcup (\text{main} \mapsto C[fut']) \\ & \sqcup (fut' \mapsto e[x \mapsto v]) & \cong_{D \sqcup \{d\}} & \Piask_{t} \sqcup (\text{main} \mapsto C[fut']) \\ & \sqcup (tid \mapsto \mathcal{D}[\![e]]_{\delta}[\delta \mapsto d][x \mapsto v]), \\ Fut_{t} \sqcup [fut' \mapsto \mathfrak{U}[d] \end{array}$

Main threads in red still equal, blue closed by TASK-COMP AWAIT-TARGET: The challenge is:

$$Task_{s} \sqcup (\text{main} \mapsto C[\text{await}(fut)]) \cong_{D} Task_{t} \sqcup (\text{main} \mapsto C[\text{await}(fut)]), Fut_{t}$$

$$\downarrow^{\text{AWAIT-TARGET}}_{Fut_{t}[fut]=\Re v}$$

$$Task_{t} \sqcup (\text{main} \mapsto C[v]), Fut_{t}$$

We need to have, on the LHS, ($fut \mapsto v$), which we will get by exploiting $Fut_t [fut] = \Re v$ in the RHS. There are two cases: TASK-VAL:

$Task_s \sqcup (main \mapsto C[await(fut)])$	\cong_D	$Task_t \sqcup (main \mapsto C[await(fut)]), Fut_t$
AWAIT-RESOLVED		$\sqrt{\frac{AWAIT-TARGET}{Fut_{l} [fut] = \Re v}}$
$Task_s \sqcup (main \mapsto C[v])$	\cong_D	$Task_t \sqcup (main \mapsto C[v]), Fut_t$

TASK-AWAIT-VAL: We have some $T_s^{(j)} = \operatorname{Aw}[fut, fut_1..fut_n] \sqcup (fut_n \mapsto v) \subseteq Task_s$. We close by n + 1 applications of AWAIT-RESOLVED.

$Task_s \sqcup (main \mapsto C[await(fut)])$	\cong_D	$Task_t \sqcup (main \mapsto C[await(fut)]), Fut_t$
AWAIT-RESOLVED $i=n.0$		AWAIT-TARGET $Fut_{l} [fut] = \Re v$
$Task_s \sqcup (main \mapsto C[v])$	\cong_D	$Task_t \sqcup (main \mapsto C[v]), Fut_t$

We are left with $T_s^{(j)'} = (fut \mapsto v) \sqcup \overline{(fut_1 \mapsto v)}^{1 \le i \le n}$, which is closed using TASK-VAL. Which concludes the cases for the main thread.

28

other groupings. In the following, we thus assume that we are in one of the $T_t^{(i)}$. We omit writing the await chains when they are not affected.

Note that we have the following property: For any C_t , e such that $e_t = C_t[e]$, there exists a context C_s such that $e_s = C[e]$.

Moreover, for any e', $C_s[e'] \sim_d C_t[e']$

Case STEP. All sub-cases follow directly from the properties presented above, and the arguments mentioned for the main thread.

Case AWAIT-RESOLVED. For the reasons explained above:

$$\begin{aligned} Task_{s} \sqcup (fut \mapsto C_{s}[\texttt{await}(fut_{r})]) &\cong_{D \sqcup \{d\}} & Task_{t} \sqcup (tid \mapsto C_{t}[\texttt{await}(fut_{r})]), Fut_{t} \\ & \downarrow_{(fut_{r} \mapsto v) \in Task_{s}(*)}^{\text{AWAIT-RESOLVED}} & \downarrow_{Fut_{t}}^{\text{AWAIT-RESOLVED}} \\ & Task_{s} \sqcup (fut_{0} \mapsto C_{s}[v]) &\cong_{D \sqcup \{d\}} & Task_{t} \sqcup (tid \mapsto C_{t}[v]), Fut_{t} \end{aligned}$$

(*) by exploiting $Fut_t [fut_r] = \Re v$, in the case TASK-VAL. See the proof in the case of the main thread for the case of TASK-AWAIT-VAL

Case CHAIN.

Using the same arguments as above to obtain (*).

Closed with TASK-VAL or TASK-AWAIT-VAL depending on the number of awaits in LHS (i.e. depending on n)

Case DPS-CALL.

$$\begin{aligned} Task_{s} \sqcup (fut \mapsto C_{s}[(f v)]) &\cong_{D \sqcup \{d\}} & Task_{t} \sqcup (tid \mapsto C_{t}[(f v)]), Fut_{t} \sqcup [fut \mapsto \mathfrak{U} d] \\ & \downarrow_{(f \mapsto \lambda_{async} x. e) \in Fun}^{Async-CALL} & \downarrow_{(f \mapsto \lambda_{dps} \delta, x. \mathcal{D}[e]_{\delta}) \in Fun'}^{DPS-CALL} \\ Task_{s} \sqcup (fut \mapsto C_{s}[fut']) & \cong_{D \sqcup \{d,d'\}} & Task_{t} \sqcup (tid \mapsto C_{t}[fut']) \\ & \sqcup (fut' \mapsto e[x \mapsto v]) & \cong_{D \sqcup \{d,d'\}} & \Pi(tid' \mapsto \mathcal{D}[e]_{\delta}[\delta \mapsto d'][x \mapsto v]), \\ Fut_{t} \sqcup [fut \mapsto \mathfrak{U} d] \sqcup [fut' \mapsto \mathfrak{U} d'] \end{aligned}$$

Chain of awaits omitted. Both painted subsets closed with TASK-COMP.

Case TAIL-CALL.

$$\begin{array}{ccc} Task_{s} \sqcup \operatorname{Aw}[fut_{0}..fut_{n}] \\ \sqcup (fut_{n} \mapsto \operatorname{await}((fv))) \\ & & \downarrow^{\operatorname{ASYNC-CALL}} \\ Task_{s} \sqcup \operatorname{Aw}[fut_{0}..fut_{n+1}] \\ \sqcup (fut_{n+1} \mapsto e[x \mapsto v]) \end{array} \xrightarrow{\cong_{D \sqcup \{d\}}} \begin{array}{c} Task_{t} \sqcup (tid \mapsto (f \#_{t}(d, v))), \\ Fut_{t} \sqcup [fut_{0} \mapsto \mathfrak{U} d] \\ & & \downarrow^{\operatorname{TAIL-CALL}} \\ Task_{t} \sqcup (tid' \mapsto \mathcal{D}\llbracket e \rrbracket_{\delta}[\delta \mapsto d][x \mapsto v]), \\ Fut_{t} \sqcup [fut_{0} \mapsto \mathfrak{U} d] \end{array}$$

With fut_{n+1} fresh and awaited nowhere else. Closed with TASK-COMP.

Case FUTURE-FILL (ε).

$$\begin{array}{l} Task_{s} \sqcup \operatorname{Aw}[fut_{0}..fut_{n}] \\ \sqcup (fut_{n} \mapsto v) \end{array} \cong_{D \sqcup \{d\}} & Task_{t} \sqcup (tid \mapsto \operatorname{refine}(d, v)), Fut_{t} \sqcup \left[fut_{0} \mapsto \mathfrak{U} d\right] \\ \cong_{D} & \downarrow_{\varepsilon}^{| \operatorname{FUTURE-FILL}} \\ & \operatorname{Task}_{t}, Fut_{t} \sqcup \left[fut_{0} \mapsto \mathfrak{R} v\right] \end{array}$$

Closed using TASK-AWAIT-VAL.

There can be at most as many FUTURE-FILL steps as there are unresolved futures in Fut_t , of which there are finitely many at any given point; which ensures that we cannot stutter indefinitely.

Which concludes the proof.

C.2 Tail-Modulo-Cons-Await backward simulation proof

C.2.1 Simulation relation. We first define the full backward simulation relation, and all necessary sub-relations.

EXPR-TRANSFO	EXPR-FILL-AWAIT	EXPR-AWAIT
$\overline{e \sim_d \mathcal{D}[\![e]\!]_{\delta}[\delta \mapsto d]}$	$e \sim_d \texttt{refine}(d, e)$	$e \overset{\text{await}}{\sim} {}_{d} \mathcal{D}_{a}\llbracket e \rrbracket_{\delta} [\delta \mapsto d]$

Fig. 38. Expr-level relations

$$\frac{\text{THREAD-DEST}}{(fut \mapsto e)_{fut} \simeq_d [d \mapsto \measuredangle], (tid \mapsto e')}$$

Fig. 39. Thread-level relation

```
\frac{T_{s \ fut_{n}} \simeq_{d} \sigma_{t}, T_{t}}{Aw[fut_{0}..fut_{n}] \ \sqcup \ T_{s \ fut_{0}} \approx_{d}^{\{fut_{1..n}\}} \sigma_{t}, T_{t}} \qquad \frac{THREAD-VAL}{(fut \mapsto v) \ \sqcup \ \overline{(fut_{i} \mapsto v)} \ _{fut} \approx_{d}^{\overline{fut_{i}}} \ [d \mapsto v], \overline{(tid \mapsto ())}^{+}}
```



$$\begin{array}{l} \overbrace{\alpha = \mathsf{await}(fut) \quad T_s \quad fut}^{\mathsf{FIELD}-\mathsf{THREAD}-\mathsf{AWAIT}} \\ \overbrace{\alpha : T_s \quad \mathcal{R}_d^I \quad \sigma_t, \ T_t}^{d} \quad d \in \operatorname{dom}(\sigma_t) \\ \hline \{\alpha\} : T_s \quad \mathcal{R}_d^I \quad \sigma_t, \ T_t \end{array} \qquad \begin{array}{l} \overbrace{\alpha = \mathsf{await}(fut) \quad T_t = \mathsf{forward}(d, \ fut)}^{\mathsf{FIELD}-\mathsf{RESOLVED}} \\ \hline \{v\} : T_s \quad \mathcal{R}_d^{\overline{fut_i}} \ [\ d \mapsto v \], \ \emptyset \end{array} \\ \hline \begin{array}{l} \overbrace{\alpha = \mathsf{await}(fut) \quad T_t = \mathsf{forward}(d, \ fut)}^{\mathsf{FIELD}-\mathsf{RESOLVED}} \\ \hline \{v\} : T_s \quad \mathcal{R}_d^{\overline{fut_i}} \ [\ d \mapsto v \], \ \emptyset \end{array} \end{array}$$

Fig. 41. For keeping sub-tasks for constructors together

 $\begin{array}{c} \begin{array}{c} \begin{array}{c} \begin{array}{c} e_{1} \stackrel{\operatorname{await}}{\longrightarrow} d_{1} e_{1}' & e_{2} \stackrel{\operatorname{await}}{\longrightarrow} d_{2} e_{2}' \\ \hline e_{1} x_{1} = e_{1} \operatorname{in} & \left[d \mapsto \ell, \ell \mapsto t_{||}, / , / \right], \\ (fut \mapsto \ \operatorname{let} x_{2} = e_{2} \operatorname{in} & \left[d \mapsto \ell, \ell \mapsto t_{||}, / , / \right], \\ (fut \mapsto \ \operatorname{let} x_{2} = e_{2} \operatorname{in} & \left[t_{||}, \operatorname{await}(x_{1}), \operatorname{await}(x_{2}) \right] \end{array} \right) fut \stackrel{\sim}{\rightarrow} d & \left(tid \mapsto \ \begin{array}{c} \operatorname{let}(d_{1}, d_{2}) = (\ell + 1, \ell + 2) \operatorname{in} \\ e_{1}' : e_{2}' \end{array} \right) \\ \end{array} \\ \end{array} \\ \begin{array}{c} \begin{array}{c} \begin{array}{c} \operatorname{PAR-BOTH-1} & \left[e_{1} \stackrel{\operatorname{await}}{\longrightarrow} \ell_{\ell+1} e_{1}' & e_{2} \stackrel{\operatorname{await}}{\longrightarrow} \ell_{\ell+2} e_{2}' \\ \hline & \left[\operatorname{let} x_{1} = e_{1} \operatorname{in} \\ (fut \mapsto \ \operatorname{let} x_{2} = e_{2} \operatorname{in} \\ & \left[t_{||}, \operatorname{await}(x_{1}), \operatorname{await}(x_{2}) \right] \end{array} \right) fut \stackrel{\sim}{\rightarrow} d & \left[d \mapsto \ell, \ell \mapsto t_{||}, / , / \right], (tid \mapsto e_{1}'; e_{2}') \end{array} \end{array} \\ \end{array} \\ \begin{array}{c} \begin{array}{c} PAR-BOTH-2 \\ & \bigvee \left\{ \begin{array}{c} \left\{ \alpha_{1} \right\} : T_{1} \int_{\ell=1}^{\mathrm{fwd}} \mathcal{R}_{\ell+1} \sigma_{1}, T_{1}' \\ & \left\{ \alpha_{1} \right\} : T_{1} \mathcal{R}_{\ell+1}^{I} \sigma_{1}, T_{1}' \\ & \left[t_{||}, \alpha_{1}, \operatorname{await}(x_{2}) \right] \end{array} \right) \sqcup T_{1} \int_{fut} \stackrel{\sim}{\rightarrow} d & \left[d \mapsto \ell, \ell \mapsto t_{||}, \ell + 2 \mapsto / \right] \sqcup \sigma_{1}, (tid \mapsto ((\cdot);)^{?} e_{2}') \sqcup T_{1}' \end{array} \\ \\ \begin{array}{c} \begin{array}{c} \operatorname{PAR-BOTH-2} \\ & \bigvee \left\{ \begin{array}{c} \left\{ \alpha_{1} \right\} : T_{1} \int_{\ell=1}^{\mathrm{fwd}} \mathcal{R}_{\ell+1} \sigma_{1}, T_{1}' \\ & \left[t_{1} \mapsto \alpha_{1}, \operatorname{await}(x_{2}) \right] \end{array} \right) \sqcup T_{1} \int_{fut} \stackrel{\sim}{\rightarrow} d & \left[d \mapsto \ell, \ell \mapsto t_{||}, \ell + 2 \mapsto / \right] \sqcup \sigma_{1}, (tid \mapsto ((\cdot);)^{?} e_{2}') \sqcup T_{1}' \end{array} \\ \\ \begin{array}{c} \operatorname{PAR-BOTH-3} \\ & \bigvee \left\{ \begin{array}{c} \left\{ \alpha_{1} \right\} : T_{1} \mathcal{R}_{\ell+1}^{I} \sigma_{1}, T_{1}' \\ & \left\{ \alpha_{2} \right\} : T_{2} \mathcal{R}_{\ell+2}^{L} \sigma_{2}, T & T = (tid \mapsto (\cdot))^{?} & \operatorname{holes}(d) \geq 1 \end{array} \\ \\ \begin{array}{c} \operatorname{PAR-BOTH-3} \\ & \bigvee \left\{ \begin{array}{c} \left\{ \alpha_{1} \right\} : T_{1} \mathcal{R}_{\ell+1}^{L} \sigma_{1}, T_{1}' \\ & \left\{ \alpha_{2} \right\} : T_{2} \mathcal{R}_{\ell+2}^{L} \sigma_{2}, T & T & T_{1}' \sqcup \sigma_{2}, T \sqcup T_{1}' \sqcup T_{2}' \end{array} \right\} \end{array} \right\} \end{array}$

Fig. 42. paratag, both awaits

, Vol. 1, No. 1, Article . Publication date: April 2025.

par-left-0

$$\begin{array}{c|c} e_1 \stackrel{\text{await}}{\sim} e_1' & e_2 \sim_{d_2} e_2' \\ \hline e_1 x_1 = e_1 \text{ in } & \left[d \mapsto \ell, \ell \mapsto t_{||}, /, / \right], \\ (fut \mapsto \begin{array}{c} \operatorname{let} x_2 = e_2 \text{ in } \\ [t_{||}, \operatorname{await}(x_1), x_2 \end{array}) & fut^{\sim} d \end{array} & \begin{array}{c} \operatorname{let} (d_1, d_2) = (\ell + 1, \ell + 2) \text{ in } \\ e_1'; e_2' \end{array} \end{array}$$

PAR-LEFT-1

$$\begin{array}{c} e_1 \stackrel{\text{await}}{\sim} e_1 e_1 & e_2 \\ \hline e_1 \stackrel{\text{await}}{\sim} e_1 e_1 & e_2 \\ \hline e_1 & e_1 & e_1 \\ (fut \mapsto e_1 e_1 & e_1 \\ e_2 & e_2 & e_2 \\ \hline e_1 & e_1 & e_1 \\ e_1 & e_1 \\ e_2 & e_2 \\ e_2 & e_2 \\ e_1 & e_1 \\ e_1 & e_1 \\ e_1 & e_1 \\ e_2 & e_2 \\ e_2 & e_2 \\ e_1 & e_1 \\ e_1 & e_1 \\ e_1 & e_1 \\ e_2 & e_2 \\ e_2 & e_2 \\ e_1 & e_1 \\ e_$$

par-left-2

$$\vee \begin{cases} \{\alpha_1\} : T_1 \stackrel{\text{fwd}}{\to} \mathcal{R}_{\ell+1} \sigma_1, T_1' \\ \{\alpha_1\} : T_1 \mathcal{R}_{\ell+1}^I \sigma_1, T_1' \end{cases} \qquad e_2 \sim_{\ell+2} e_2'$$

 $(fut \mapsto \begin{array}{c} \operatorname{let} x_2 = e_2 \operatorname{in} \\ [t_{||}, \alpha_1, x_2] \end{array}) \sqcup T_1 \quad fut \cong \begin{array}{c} \overset{I}{d} \left[d \mapsto \ell, \ell \mapsto t_{||}, \ell_2 \mapsto \swarrow \right] \sqcup \sigma_1, (tid \mapsto (();)^2 e_2') \sqcup T_1' \end{array}$

$$\frac{\operatorname{PAR-LEFT-3}}{\left(fut \mapsto [t_{||}, \alpha_{1}, v_{2}]\right) \sqcup T_{1} \sqcup T_{2} \operatorname{fut}^{\simeq} d\left[d \mapsto \ell, \ell \mapsto t_{||}, \ell + 2 \mapsto v_{2}\right] \sqcup \sigma_{1}, T \sqcup T_{1}} \qquad T = (tid \mapsto ())^{?} \qquad \operatorname{holes}(d) \geq 1$$



 $\frac{e_{1} \sim_{d_{1}} e'_{1} \qquad e_{2} \sim_{d_{2}} e'_{2}}{e_{1} \quad \text{let } x_{1} = e_{1} \text{ in } \qquad \left[d \mapsto \ell, \ell \mapsto t_{||}, /, /\right], \\ (fut \mapsto \begin{array}{c} \operatorname{let } x_{2} = e_{2} \text{ in } \\ [t_{||}, x_{1}, x_{2}] \end{array}) \quad fut^{\simeq} d \quad (tid \mapsto \begin{array}{c} \operatorname{let } (d_{1}, d_{2}) = (\ell + 1, \ell + 2) \text{ in } \\ e'_{1}; e'_{2} \end{array}) \\ \\
\frac{PAR-NONE-1}{\left[t_{||}, x_{1}, x_{2}\right]} \quad e_{1} \sim_{\ell+1} e'_{1} \qquad e_{2} \sim_{\ell+2} e'_{2} \\ \hline (fut \mapsto \begin{array}{c} \operatorname{let } x_{1} = e_{1} \text{ in } \\ (fut \mapsto \begin{array}{c} \operatorname{let } x_{2} = e_{2} \text{ in } \\ [t_{||}, x_{1}, x_{2}\end{array}) \quad fut^{\simeq} d \quad \left[d \mapsto \ell, \ell \mapsto t_{||}, /, /\right], (tid \mapsto e'_{1}; e'_{2}) \\ \hline \\ PAR-NONE-2 \\ \hline e_{2} \sim_{\ell+2} e'_{2} \\ \hline (fut \mapsto \begin{array}{c} \operatorname{let } x_{2} = e_{2} \text{ in } \\ [t_{||}, v_{1}, x_{2}\end{array}) \quad fut^{\simeq} d \quad \left[d \mapsto \ell, \ell \mapsto t_{||}, v_{1}, /\right], (tid \mapsto ((1);)^{2} e'_{2}) \\ \hline \end{array}$

Fig. 44. partag, no awaits

seq-none-0

$$\begin{array}{c} \overbrace{(fut \mapsto [t_{seq}, v_1, x_2]]}^{\sim} & \underbrace{e_2 \sim_{d_2} e'_2} \\ \hline (fut \mapsto [t_{seq}, v_1, x_2]] & [d \mapsto \ell, \ell \mapsto t_{seq}, v_1, \swarrow], \\ \overbrace{t_{seq}, v_1, x_2}] & fut \stackrel{\sim}{\rightarrow} d & (tid \mapsto [e'_2] \\ \hline e'_2 & (tid \mapsto e'_2) \\ \hline e'_2 & (t$$

seq-none-1

$$\begin{array}{c} e_2 \sim_{\ell+2} e'_2 \\ \hline (fut \mapsto \begin{array}{c} \mathsf{let} x_2 = e_2 \mathsf{in} \\ [t_{\mathsf{seq}}, v_1, x_2] \end{array}) fut^{\simeq}_d \left[d \mapsto \ell, \ell \mapsto t_{\mathsf{seq}}, v_1, \swarrow \right], (tid \mapsto e'_2) \end{array}$$



seq-left-0

$$\frac{e_1 \overset{\text{await}}{\sim} d_1 e'_1}{(fut \mapsto \underset{[t_{\text{seq}}, await(x_1), v_2]}{\text{let } x_1 = e_1 \text{ in } [t_{\text{seq}}, d_1 e'_1]} \begin{bmatrix} d \mapsto \ell, \ell \mapsto t_{\text{seq}}, \neq, v_2 \end{bmatrix}, \\ (fut \mapsto \underset{[t_{\text{seq}}, await(x_1), v_2]}{\text{let } d_1 e'_1} \end{bmatrix}_{fut} \overset{\text{let } d_1 = \ell + 1 \text{ in } e'_1}{(tid \mapsto \underset{e'_1}{e'_1} e'_1)}$$

SEQ-LEFT-1

$$\begin{array}{c} \underbrace{e_1 \stackrel{\text{await}}{\sim} e_{1}}_{\ell+1} e'_1 \\ \hline \\ (fut \mapsto \begin{array}{c} \texttt{let } x_1 = e_1 \texttt{ in } \\ \texttt{[} t_{\texttt{seq}}, \texttt{await}(x_1), v_2 \texttt{]} \end{array})_{fut} \stackrel{}{\simeq} \underbrace{\left[\begin{array}{c} d \mapsto \ell, \ell \mapsto t_{\texttt{seq}}, \swarrow, v_2 \end{array}\right],}_{(tid \mapsto e'_1)} \end{array}$$

.

seq-left-2

_

$$\frac{\bigvee \left\{ \begin{array}{l} \left\{ \alpha_{1} \right\} : T_{1} \ ^{\text{twd}}\mathcal{R}_{\ell+1} \sigma_{1}, T_{1}^{\prime} \\ \left\{ \alpha_{1} \right\} : T_{1} \mathcal{R}_{\ell+1}^{I} \sigma_{1}, T_{1}^{\prime} \end{array} \right. T = (tid \mapsto ())^{?} \quad \text{holes}(d) \ge 1}{\left(fut \mapsto \ [t_{\text{seq}}, \alpha_{1}, v_{2}] \) \sqcup T_{1} \ _{fut}^{\simeq}_{d} \ \begin{array}{l} \left[d \mapsto \ell, \ell \mapsto t_{\text{seq}}, \ell+2 \mapsto v_{2} \right] \sqcup \sigma_{1}, \end{array} \right.}$$

Fig. 46. seqtag, await on the left (sym for right)

$$\frac{T_{s}}{T_{s}} \stackrel{\simeq I}{\cong} \stackrel{s_{t}, T_{t}}{s_{t}, T_{t}, [fut \mapsto \mathfrak{U} d]} \qquad \frac{\text{BOX-VAL}}{(fut \mapsto v) \sqcup (fut_{i} \mapsto v)^{i}} \stackrel{\simeq \overline{fut_{i}}^{i}}{\cong \overline{fut_{i}}^{i}} [d \mapsto v], \emptyset, [fut \mapsto \mathfrak{U} d]$$

$$\frac{\text{BOX-RESOLVED}}{(fut \mapsto v) \sqcup (fut_{i} \mapsto v)^{i}} \stackrel{\simeq \overline{fut_{i}}^{i}}{\cong \overline{fut_{i}}^{i}} \emptyset, T, [fut \mapsto \mathfrak{R} v]$$

Fig. 47. Top-level-future-level rel (=boxes)

Config-level. The simulation relation between source and target configurations:

$$\sigma_{s}, Task_{s} \succeq \sigma_{t}, Task_{t}, Fut_{t} \Leftrightarrow \begin{cases} (1) \quad Task_{s} = (\min \mapsto e) \sqcup T_{s}^{(1)} \sqcup \cdots T_{s}^{(n)} \\ (2) \quad Task_{t} = (\min \mapsto e) \sqcup T_{t}^{(1)} \sqcup \cdots T_{t}^{(n)} \\ (3) \quad \sigma_{t} = \sigma_{s} \sqcup \sigma_{1} \sqcup \cdots \sigma_{n} \\ (4) \quad Fut_{t} = \left[\overline{fut_{i} \mapsto \mathcal{F}_{i}}^{1 \le i \le n} \right] \quad (\mathcal{F}_{i} = \mathfrak{U} d_{i} \mid \mathfrak{R} v_{i}) \\ (5) \quad \forall i, \ T_{s}^{(i)} \cong^{I_{i}} \sigma_{i}, T_{t}^{(i)}, \left[fut_{i} \mapsto \mathcal{F}_{i} \right] \end{cases}$$

with each I_i internal sets of futures in the sense given in Section 2.5. Like in tma, steps are "local" to one thread pool.

C.2.2 Proof sketch. We follow the same sketch as in tma:

- (1) isolate in which set of threads the step operates: either the main, or one of the parallel pools
- (2) case analysis on \cong , then on the step.
- (3) in the parallel pools, in case BOX-THREAD, use lemma 3.1

Steps in the main thread. The same arguments as for tma are still valid here, with one caveat: the stores are not fully synchronised. However, only the synchronised parts are accessible by the main thread.

Steps in parallel thread pools. We are thus in a $T_t^{(i)}$. We write the configs as

$$\sigma_s \sqcup \sigma_{others} \sqcup \sigma_i$$

$$\sigma_s, Task_s \sqcup T_s^{(i)} \succeq Task_t \sqcup T_t^{(i)},$$

$$Fut_t \sqcup F_i$$

where $T_i \cong^I \sigma_i, T_i, F_i$

We also have the same property as for single-hole: if $e_s \sim_d e_t$, then for any C_t , e such that $e_t = C_t[e]$, there exists a context C_s such that $e_s = C[e]$.

Moreover, for any e', $C_s[e'] \sim_d C_t[e']$

BOX-VAL. Only step that can be taken is FUTURE-RESOLVE (ε). We close immediately with BOX-RESOLVED. Cannot have infinitely many FUTURE-RESOLVE steps in a row: each consumes an unresolved future, which there are finitely many at each step.

BOX-RESOLVED. Only possible step is PRUNE-TASKS (ε) (only if $T \neq \emptyset$), and we close immediately.

BOX-THREAD. We give here most elements of the proof of lemma 3.1, from which the proof for this case follows immediately

THREAD-DEST. This case handles almost every computation, except constructors being built

- **local steps** All local steps except the ones for **refine** are treated the same as for the single-hole optimisation, using the contexts property, and the same arguments as in the main thread. For the store argument, only refine can read or write to unsynchronised parts of the store.
- **refine** (\neq FILL) A refine step necessarily happens in a context of the form **let** $\delta^+ = \text{refine}(d, \Box)$ in *e*. There are several cases depending on the shape of the second argument of **refine** (sequential or parallel tag, which and how many fields are in the domain of sta). They are all matched with one of the CONSTR-PAR/SEQ rules, and we enter one of the chains described earlier, in a state PAR/SEQ-<_>-0.
- FILL We close with THREAD-VAL after *n* applications of AWAIT-RESOLVED.

$$\begin{array}{ccc} \mathsf{Aw}[\mathit{fut}_0..\mathit{fut}_n] \sqcup (\mathit{fut}_n \mapsto v) & {}_{\mathit{fut}_0} \approx_d & [\mathit{d} \mapsto \swarrow], (\mathit{tid} \mapsto \mathsf{refine}(\mathit{d}, v)) \\ & & & & \\ & & & \downarrow^{\text{AWAIT-RESOLVED}} & & & & \downarrow^{\text{FILL}} \\ & & & & & \downarrow^{\text{FILL}} \\ & & & & & \downarrow^{\text{fut}_0} \approx_d & [\mathit{d} \mapsto v], (\mathit{tid} \mapsto ()) \end{array}$$

Even though we might answer the challenge with 0 steps on the source side, we cannot take infinitely many FILL steps in a row: each FILL step consumes a binding [$d \mapsto \checkmark$] from the store. The number of these bindings in the store is thus a strictly decreasing measure for possibly- ε steps. It is moreover always finite, thus we cannot stutter indefinitely. This also verifies the second part of the lemma.

CHAIN-RESOLVED We have $e \sim_d e' = C[\mathbf{forward}(fut', d)]$. The only possible case for C is \Box . Thus, with $Fut_t [fut'] = \Re v$, we get by BOX-RESOLVED that $(fut' \mapsto v) \in Task_s$

$$\begin{split} \mathsf{Aw}[\mathit{fut}_0..\mathit{fut}_n] \sqcup (\mathit{fut}_n \mapsto \mathsf{await}(\mathit{fut}')) & {}_{\mathit{fut}_0} \approx_d & [d \mapsto \nearrow], (\mathit{tid} \mapsto \mathsf{forward}(\mathit{fut}', d)) \\ \downarrow^{\mathrm{AWAIT-RESOLVED}} & \downarrow^{\mathrm{CHAIN-RESOLVED}} \\ (\mathit{fut}_0 \mapsto v) \sqcup \overline{(\mathit{fut}_i \mapsto v)}^{i=1..n} & {}_{\mathit{fut}_0} \approx_d & [d \mapsto v], (\mathit{tid} \mapsto ()) \end{split}$$

Closing again with THREAD-VAL, same as fill, with possibly n + 1 awaits.

This also verifies the second part of the lemma.

DPS-CALL Very similar to the case for the single-hole optimisation. Here, we create a new "box":



Red part is the same, blue by BOX-THREAD, THREAD-DEST TAIL-CALL Exactly the same as for the single-hole optimisation. CHAIN-UNRESOLVED, FUTURE-RESOLVE, PRUNE-TASKS None of these rules applicable in this case.

Continuing the cases of \approx :

THREAD-VAL. Two rule can be applied: either PRUNE-TASKS, or FUTURE-RESOLVE.

- PRUNE-TASKS stutters and we either stay in the same state, or if if it was the last task to prune, go to BOX-VAL.
- FUTURE-RESOLVE matched with *n* awaits, then closed with BOX-RESOLVED

The sequence PAR-BOTH-X. We give details here fro one of the chains handling constructors: the case of parallel-tagged constructors, with both fields "asynchronous" (read: in the domain of sta). The arguments for the other sequences are the same.

PAR-BOTH-0 Only case is reducing the let, stuttering, going to PAR-BOTH-1

- PAR-BOTH-1 Any step under context in e'_1 will be matched in e_1 , stay in same state. Last steps can be either TAIL-CALL or one of the chain ones, and need to be matched with the corresponding source step, then the substitution (e_1 is now just a future), and we go to PAR-BOTH-2
- PAR-BOTH-2 Any step will be either in the thread *tid*, or in T'_1 . Steps in T'_1 keep us in the same state, we go to PAR-BOTH-3 only with the last step in e'_2 .
 - in *tid* First step will be getting rid of the $(();)^?$, which is epsilon. Then, any subsequent steps in e'_2 will be matched in e_2 in the same way as e'_1 in the previous case. The last step will either create a T'_2 , or directly resolve to a value, in the same way.
 - in T'_1 if its under FIELD-CHAIN, only possible step is CHAIN-RESOLVED, ok. Otherwise, induction hypothesis.

After the last step (which by ind ends with $(fut_1 \mapsto v_1)$ on the LHS), we must resolve the **await**(*fut*₁), and we close this part with FIELD-RESOLVED.

PAR-BOTH-3 Same as the previous one, but now steps might also be in T'_2 . They are however for the most part treated in the same way.

The last step in this state is already detailed in Section 3.4