

ParT: An Asynchronous Parallel Abstraction for Speculative Pipeline Computations

Kiko Fernandez-Reyes^(✉), Dave Clarke, and Daniel S. McCain

Department of Information Technology, Uppsala University, Uppsala, Sweden
kiko.fernandez@it.uu.se

Abstract. The ubiquity of multicore computers has forced programming language designers to rethink how languages express parallelism and concurrency. This has resulted in new language constructs and new combinations or revisions of existing constructs. In this line, we extended the programming languages Encore (actor-based), and Clojure (functional) with an asynchronous parallel abstraction called ParT, a data structure that can dually be seen as a collection of asynchronous values (integrating with futures) or a handle to a parallel computation, plus a collection of combinators for manipulating the data structure. The combinators can express parallel pipelines and speculative parallelism. This paper presents a typed calculus capturing the essence of ParT, abstracting away from details of the Encore and Clojure programming languages. The calculus includes tasks, futures, and combinators similar to those of Orc but implemented in a non-blocking fashion. Furthermore, the calculus strongly mimics how ParT is implemented, and it can serve as the basis for adaptation of ParT into different languages and for further extensions.

1 Introduction

The ubiquity of multicore computers has forced programming language designers to rethink how languages express parallelism and concurrency. This has resulted in new language constructs that, for instance, increase the degree of asynchrony while exploiting parallelism. A promising direction is programming languages with constructs for tasks and actors, such as Clojure and Scala [8, 16], due to the lightweight overhead of spawning parallel computations. These languages offer coarse-grained parallelism at the task and actor level, where futures act as synchronisation points. However, these languages are lacking in high-level coordination constructs over these asynchronous computations. For instance, it is not easy to express dependence on the first result returned via a bunch of futures and to safely terminate the computations associated with the other futures. The task of terminating speculative parallelism is quite delicate, as the futures

Partly funded by the EU project FP7-612985 UPSCALE: From Inherent Concurrency to Massive Parallelism through Type-based Optimisations.

may have attached parallel computations that depend on other futures, creating complex dependency patterns that need to be tracked down and terminated.

To address this need, this paper presents the design and implementation of ParT, a non-blocking abstraction that asynchronously exploits futures and enables the developer to build complex, data parallel coordination workflows using high-level constructs. These high-level constructs are derived from the combinators of the orchestration language Orc [11, 12]. ParT is formally expressed in terms of a calculus that, rather than being at a high level of abstraction, strongly mimics how this asynchronous abstraction is implemented and is general enough to be applied to programming languages with notions of futures.

The contributions of the paper are as follows: the design of an asynchronous parallel data abstraction to coordinate complex workflows, including pipeline and speculative parallelism, *and* a typed, non-blocking calculus modelling this abstraction, which integrates futures, tasks and Orc-like combinators, supports the separation of the realisation of parallelism (via tasks) from its specification, and offers a novel approach to terminating speculative parallelism.

2 Overview

To set the scene for this paper, we begin with a brief overview to asynchronous computations with futures and provide an informal description of the ParT abstraction and its combinators. A SAT solver example is used as an illustration.

In languages with notions of tasks and active objects [2, 8, 16], asynchronous computations are created by spawning tasks or calling methods on active objects. These computations can exploit parallelism by decoupling the execution of the caller and the callee [7]. The result of a spawn or method call is immediately a future, a container that will eventually hold the result of the asynchronous computation. A future that has received a value is said to be *fulfilled*. Operations on futures may be blocking, such as getting the result from a future, or may be asynchronous, such as attaching a callback to a future. This second operation, called *future chaining* and represented by $f \rightsquigarrow \text{callback}$, immediately returns a new future, which will contain the result of applying the callback function *callback* to the contents of the original future after it has been fulfilled. A future can also be thought of as a handle to an asynchronous computation that can be extended via future chaining or even terminated. This is an useful perspective that we will further develop in this work. In languages with notions of actors, such as Clojure and Encore [2], asynchrony is the rule and blocking on futures suffers a large performance penalty. But creating complex coordination patterns based on a collection of asynchronous computations without blocking threads (to maintain the throughput of the system) is no easy task.

To address this need, we have designed an abstraction, called ParT, which can be thought of as a handle to an ongoing parallel computation, allowing the parallel computation to be manipulated, extended, and terminated. A ParT is a functional data structure, represented by type *Par t*, that can be empty $\{\}$:: *Par t*, contain a single expression $\{-\}$:: $t \rightarrow \text{Par } t$, or futures attached to

computations producing values, using $(-)^{\circ} :: Fut\ t \rightarrow Par\ t$, or computations producing ParTs, embedded using $(-)^{\dagger} :: Fut\ (Par\ t) \rightarrow Par\ t$. Multiple ParTs can be combined using the par constructor, $\parallel :: Par\ t \rightarrow Par\ t \rightarrow Par\ t$. This constructor does not necessarily create new parallel threads of control, as this would likely have a negative impact on performance, but rather specifies that parallelism is available. The scheduler in the ParT implementation can choose to spawn new tasks as it sees fit — this is modelled in the calculus as a single rule that nondeterministically spawns a task from a par (rule RED-SCHEDULE).

The combinators can express complex coordination patterns and operate on them in a non-blocking manner, and safely terminate speculative parallelism even in the presence of complex workflows. These combinators will be illustrated using an example, then explained in more detail.

Illustrative Example. Consider a portfolio-based SAT solver (Fig. 1), which creates numerous strategies, of which each finds an assignment of variables to Boolean values for a given proposition, runs them in parallel, and accepts the first solution found. Each strategy tries to find a solution by selecting a variable and creating two instances of the formula, one where the variable is assigned true, the other where it is assigned false (called splitting) — strategies differ in the order they select variables for splitting. These new instances can potentially be solved in parallel.

The example starts in function `process` (line 20) which receives an array of strategies and the formula to solve. Strategies do not interact with each other and can be lifted to a ParT, creating a parallel pipeline (line 21) using the `each`

```

1  def fut2Par(f: Fut(Maybe a)): Par a
2    (f ~> \m: Maybe a ->
3      match m with Nothing => {}; Just val => {val})†
4
5  def evaluateFormula(form: Formula, a: Assignment): (Maybe bool, Assignment)
6    ...
7
8  def sat(st: Strategy, fml: Formula, a: Assignment): Par Assignment
9    let variable = st.getVariable(fml, a)
10   a1 = a.extendAssignment(variable, true)
11   a2 = a.extendAssignment(variable, false)
12   in
13     ({evaluateFormula(fml, a1)} || {evaluateFormula(fml, a2)}) >>=
14     \result: (Maybe bool, Assignment) ->
15       match result with
16         (Nothing, ar) => sat(st, fml, ar);
17         (Just true, ar) => {ar};
18         (Just false, ar) => {};
19
20  def process(sts: [Strategy], fml: Formula): Par Assignment
21    fut2Par << (each(sts) >>= \s: Strategy ->
22      (async sat(s, fml, new Assignment()))†)

```

Fig. 1. A SAT solver in Encore.

and `bind` (\gg) combinators. As soon as one strategy finds an assignment, the remaining computations are terminated via the `prune` (\ll) combinator.

For each strategy, a call to the `sat` function (line 8) is made in parallel using a call to `async`, which in this case returns a value of type `Fut (Par Assignment)`. Function `sat` takes three arguments: a strategy, a formula and an assignment object containing the current mapping from variables to values. This function uses the strategy object to determine which variable to split next, extends the assignment with new valuations (lines 9–11), recursively solves the formula (by again calling `sat`), and returns an assignment object if successful. The evaluation of the formula, `evaluateFormula` returns, firstly, an optional Boolean to indicate whether evaluation has completed, and if it has completed, whether the formula is satisfiable, and secondly, the current (partial) variable assignment. The two calls to `evaluateFormula` are grouped into a new `ParT` collection (using `||`) and, with the use of the \gg combinator, a new asynchronous pipeline is created to either further evaluate the formula by calling `sat`, to return the assignment in the case that a formula is satisfiable as a singleton `ParT`, or `{}` when the assignment does not satisfy the formula (lines 14–18).

Finally, returning back to `process`, the `prune` combinator (\ll) (line 21) is used to select the first result returned by the recursive calls to `sat`, if there is one. This result is converted from an option type to an empty or singleton `ParT` collection (again asynchronously), which can then be used in a larger parallel operation, if so desired. The `prune` combinator will begin poisoning and safely terminating the no longer needed parallel computations, which in this case will be an ongoing parallel pipeline of calls to `sat` and `evaluateFormula`.

ParT Combinators. The combinators are now described in detail. The combinators manipulate `ParT` collections and were derived from `Orc` [11, 12], although in our setting, they are typed and redefined to be completely asynchronous, never blocking the thread. Primitive combinators express coordination patterns such as pipeline and speculative parallelism, and more complex patterns can be expressed based on these primitives.

Pipeline parallelism is expressed in `ParT` with the sequence and bind combinators. The sequence combinator, $\gg :: Par\ t \rightarrow (t \rightarrow t') \rightarrow Par\ t'$, takes a `ParT` collection and applies the function to each element in the collection, potentially in parallel, returning a new `ParT` collection. The bind combinator (derived from other combinators) $\ggg :: Par\ t \rightarrow (t \rightarrow Par\ t') \rightarrow Par\ t'$ is similar to the sequence combinator, except that the function returns a `ParT` collection and the resulting nested `ParT` collection is flattened. (*Par* is a monad!¹) In the presence of futures inside a `ParT` collection, these combinators use the future chaining operation to create independent and asynchronous pipelines of work.

Speculative parallelism is realised by the `peek` combinator, `peek :: Par\ t \rightarrow Fut (Maybe t)`, which sets up a speculative computation, asynchronously waits for a single result to be produced, and then safely terminates the speculative work. To terminate speculative work the `ParT` abstraction poison these specu-

¹ The monad operations on *Par* are essentially the same as for lists but parallelised.

lative computations, which may have long parallel pipelines to which the poison spreads recursively, producing a pandemic infection among futures, tasks and pipelines of computations. Afterwards, poisoned computations that are no longer needed can safely be terminated. Metaphorically, this is analogous to a tracing garbage collector.

The value produced by `peek` is a future to an option type. The option type is used to capture whether the parallel collection was empty or not. The empty collection `{}` results in `Nothing`, and a non-empty collection results in a `Just v`, where v is the first value produced. The conversion to option type is required because ParTs cannot be tested for emptiness without blocking. The `peek` combinator is an internal combinator, i.e., it is not available to the developer and is used by the prune \ll combinator (explained below).

Built on top of `peek` is the prune combinator, $\ll :: (Fut (Maybe t) \rightarrow Par\ t') \rightarrow Par\ t \rightarrow Par\ t'$, which applies a function in parallel to the future produced by `peek`, and returns a parallel computation.

Powerful combinators can be derived from the ones mentioned above. An example of a derived combinator, which is a primitive in Orc, is the otherwise combinator, $\times :: Par\ t \rightarrow Par\ t \rightarrow Par\ t$ (derivation is shown in Sect. 3.1). Expression $e_1 \times e_2$ results in e_1 unless it is an empty ParT, in which case it results in e_2 .

Other ParT combinators are available. For instance, `each` $:: [t] \rightarrow Par\ t$ and `extract` $:: Par\ t \rightarrow [t]$ convert between sequential (arrays) and ParTs. The latter potentially requires a lot of synchronisation, as all the values in the collection need to be realised. Both have been omitted from the formalism, because neither presents any real technical challenge — the key properties of the formalism, namely, deadlock-freedom, type preservation and task safety (Sect. 3.5), still hold with these extensions in place.

3 A Typed ParT Calculus

This section presents the operational semantics and type system of a task-based language containing the ParT abstraction. The formal model is roughly based on the Encore formal semantics [2, 5], with many irrelevant details omitted.

3.1 Syntax

The core language (Fig. 2) contains expressions e and values v . Values include constants c , variables, futures f , lambda abstractions, and ParT collections of values. Expressions include values v , function application $(e\ e)$, task creation, future chaining, and parallel combinators. Tasks are created via the `async` expression, which returns a future. The parallel combinators are those covered in Sect. 2 (`||`, `>>`, `peek` and `<<`), plus some derived combinators, together with the low-level combinator `join` that flattens nested ParT collections. Recall that `peek` is used under-the-hood in the implementation of `<<`. Status π controls how `peek` behaves: when π is \odot and the result in `peek` is an empty ParT collection,

the value is discarded and not written to the corresponding future. This status helps to ensure that precisely one speculative computation writes into the future and that a speculative computation fails to produce a value only when all relevant tasks fail to produce a value.

$$\begin{aligned}
e &::= v \mid e \mid e \mid \text{async } e \mid e \rightsquigarrow e \mid \{e\} \mid e \parallel e \\
&\quad \mid e \gg e \mid e \ll e \mid e^\circ \mid e^\dagger \mid \text{join } e \mid \text{peek}^\pi e \\
v &::= c \mid f \mid x \mid \lambda x.e \mid \{ \} \mid \{v\} \mid f^\circ \mid f^\dagger \mid v \parallel v \\
\pi &::= _ \mid \emptyset
\end{aligned}$$

Fig. 2. Syntax of the language.

ParT collections are monoids, meaning that the composition operation $e \parallel e$ is associative and has $\{ \}$ as its unit. As such, ParT collections are sequences, though no operations such as getting the first element are available to access them sequentially. As an alternative, adding in commutativity of \parallel would give multiset semantics to the ParT collections — the operational semantics is otherwise unchanged. Two for one!

A number of the constructs are defined by translation into other constructs.

$$\begin{aligned}
\text{let } x = e \text{ in } e' &\hat{=} (\lambda x.e') e \\
e_1 \times e_2 &\hat{=} \text{let } x = e_1 \text{ in} \\
&\quad (\lambda y.(y \rightsquigarrow (\lambda z.\text{match } z \text{ with Nothing} \rightarrow e_2; _ \rightarrow x)))^\dagger \ll x \\
e_1 \gg e_2 &\hat{=} \text{join } (e_1 \gg e_2) \\
\text{maybe2par} &\hat{=} \lambda x.\text{match } x \text{ with Nothing} \rightarrow \{ \}; \text{Just } y \rightarrow \{y\}
\end{aligned}$$

The encoding of **let** is standard. In $e_1 \times e_2$, pruning \ll is used to test the emptiness of e_1 . If it is not empty, the result of e_1 is returned, otherwise the result is e_2 . The definition of \gg is a standard definition of monadic bind in terms of **map** (\gg) and **join**. We assume for convenience a **Maybe** type and pattern matching on it.

3.2 Configurations

Running programs are represented by configurations (Fig. 3). Configurations can refer to the global system or a partial view of the system. A global configuration $\{ \text{config} \}$ captures the complete global state, e.g., $\{ (\text{fut}_f) (\text{task}_f e) \}$ shows a global system containing a single task running expression e . Local configurations, written as config , show a partial view of the state of the program. These are multisets of tasks, futures, poison and future chains. The empty configuration is represented by ϵ . Future configurations, (fut_f) and $(\text{fut}_f v)$, represent unfulfilled and fulfilled futures, respectively. Poison is the configuration $(\text{poison } f)$ that will eventually terminate tasks and chains writing to future f and their dependencies. A running task $(\text{task}_f^\alpha e)$ has a body e and will write its result to

future f . The chain configuration $(\text{chain}_f^\alpha g e)$ depends on future g that, when fulfilled, will then run expression e on the value stored in g , and write its value into future f . Concatenation of configurations, $\text{config config}'$, is associative and commutative with the empty configuration ϵ as its unit (Fig. 12).

$$\begin{aligned} g\text{config} &::= \{\text{config}\} \\ \text{config} &::= \epsilon \mid (\text{fut}_f) \mid (\text{fut}_f v) \mid (\text{poison } f) \mid (\text{task}_f^\alpha e) \mid (\text{chain}_f^\alpha g e) \mid \text{config config}' \\ \alpha &::= _ \mid \dagger \end{aligned}$$

Fig. 3. Runtime configurations.

Tasks and chains have a flag α that indicates the poisoned state of the computation. Whitespace ‘ $_$ ’ indicates that the computation has not been poisoned, and \dagger indicates that the computation has been poisoned and can be safely terminated, if it is not needed (see Rule RED-TERMINATE of Fig. 10).

The initial configuration to evaluate expression e is $\{(\text{task}_f e) (\text{fut}_f)\}$, where the value written into future f is the result of the expression.

3.3 Reduction Rules

The operational semantics is based on a small-step, reduction-context based rules for evaluation within tasks, and parallel reduction rules for evaluation across configurations. Evaluation is captured by expression-level evaluation context E containing a hole \bullet that marks where the next step of the reduction will occur (Fig. 4). Plugging an expression e into an evaluation context E , denoted $E[e]$, represents both the subexpression to be evaluated next and the result of reducing that subexpression in context, in the standard fashion [21].

$$\begin{aligned} E &::= \bullet \mid E e \mid v E \mid E \rightsquigarrow e \mid v \rightsquigarrow E \mid \{E\} \mid E \parallel e \mid v \parallel E \mid E \gg e \mid v \gg E \\ &\mid E \ll e \mid E^\circ \mid E^\dagger \mid \text{join } E \mid \text{peek}^\pi E \end{aligned}$$

Fig. 4. Expression-level evaluation contexts.

Reduction of configurations is denoted $\text{config} \rightarrow \text{config}'$, which states that config reduces in a single step to config' .

Core Expressions. The core reduction rules (Fig. 5) for functions, tasks and futures are well-known or derived from earlier work [5]. Together, the rules RED-CHAIN and RED-CHAINV describe how future chaining works, initially attaching a closure to a future (via the chain configuration), then evaluating the closure in a new task after the future has been fulfilled.

$$\begin{array}{c}
\frac{}{(\text{task}_g^\alpha E[(\lambda x.e) v]) \rightarrow (\text{task}_g^\alpha E[e[v/x]])} \quad \frac{}{(\text{task}_g^\alpha E[\text{async } e]) \rightarrow (\text{fut}_f) (\text{task}_f^\alpha e) (\text{task}_g^\alpha E[f])} \\
\text{(RED-}\beta\text{-RED)} \qquad \qquad \qquad \text{(RED-ASYNC)} \\
\text{fresh } f \\
\frac{}{(\text{task}_f^\alpha v) (\text{fut}_f) \rightarrow (\text{fut}_f v)} \quad \frac{}{(\text{task}_g^\alpha E[f \rightsquigarrow v]) \rightarrow (\text{fut}_h) (\text{chain}_h^\alpha f v) (\text{task}_g^\alpha E[h])} \\
\text{(RED-FUTV)} \qquad \qquad \qquad \text{(RED-CHAIN)} \\
\text{fresh } h \\
\frac{}{(\text{chain}_g^\alpha f e) (\text{fut}_f v) \rightarrow (\text{task}_g^\alpha (e v)) (\text{fut}_f v)} \\
\text{(RED-CHAINV)}
\end{array}$$

Fig. 5. Core reduction rules.

$$\begin{array}{c}
\frac{}{(\text{task}_g^\alpha E[\{\} \gg v]) \rightarrow (\text{task}_g^\alpha E[\{\}])} \quad \frac{}{(\text{task}_g^\alpha E[\{v\} \gg v']) \rightarrow (\text{task}_g^\alpha E[\{v' v\}])} \\
\text{(RED-SEQS)} \qquad \qquad \qquad \text{(RED-SEQV)} \\
\frac{}{(\text{task}_g^\alpha E[f^\circ \gg v]) \rightarrow (\text{task}_g^\alpha E[(f \rightsquigarrow v)^\circ])} \quad \frac{}{(\text{task}_g^\alpha E[f^\dagger \gg v]) \rightarrow (\text{task}_g^\alpha E[(f \rightsquigarrow (\lambda x.x \gg v))^\dagger])} \\
\text{(RED-SEQF)} \qquad \qquad \qquad \text{(RED-SEQFP)} \\
\frac{}{(\text{task}_g^\alpha E[(v_1 \parallel v_2) \gg v]) \rightarrow (\text{task}_g^\alpha E[(v_1 \gg v) \parallel (v_2 \gg v)])} \\
\text{(RED-SEQP)}
\end{array}$$

Fig. 6. Reduction rules for the sequence \gg combinator.

Sequencing. The sequencing combinator \gg creates pipeline parallelism. Its semantics are defined inductively on the structure of ParT collections (Fig. 6). The second argument must be a function (tested in function application, but guaranteed by the type system). In RED-SEQS, sequencing an empty ParT results in another empty ParT. A ParT with a value applies the function immediately (RED-SEQV). A lifted future is asynchronously accessed by chaining the function onto it (RED-SEQF). Rule RED-SEQP recursively applies $\gg v$ to the two sub-collections. A future whose content is a ParT collection chains a recursive call to $\gg v$ onto the future and lifts the result back into a ParT collection (RED-SEQFP).

Join. The `join` combinator flattens nested ParT collections of type *Par* (*Par t*) (Fig. 7). Empty collections flatten to empty collections (RED-JOINS). Rule RED-JOINV extracts the singleton value from a collection. A lifted future that contains a ParT (type *Fut (Par t)*) is simply lifted to a ParT collection (RED-JOINF). In RED-JOINFP, a future containing a nested ParT collection (type *Fut (Par (Par t))*), chains a call to `join` to flatten the inner structure. Rule RED-JOINP applies the `join` combinator recursively to the values in the ParT collection.

Prune and Peek. Pruning is the most complicated part of the calculus, though most of the work is done using the `peek` combinator (Fig. 8). Firstly, rule RED-PRUNE spawns a new task that will peek the collection v' , and passes this new task's future to the function v . The essence of the `peek` rules is to set up a

$$\begin{array}{c}
 \text{(RED-JOINS)} \qquad \qquad \qquad \text{(RED-JOINV)} \\
 (\text{task}_g^\alpha E[\text{join } \{\}]) \rightarrow (\text{task}_g^\alpha E[\{\}]) \qquad (\text{task}_g^\alpha E[\text{join } \{v\}]) \rightarrow (\text{task}_g^\alpha E[v]) \\
 \\
 \text{(RED-JOINF)} \qquad \qquad \qquad \text{(RED-JOINFP)} \\
 (\text{task}_g^\alpha E[\text{join } f^\circ]) \rightarrow (\text{task}_g^\alpha E[f^\dagger]) \qquad (\text{task}_g^\alpha E[\text{join } f^\dagger]) \rightarrow (\text{task}_g^\alpha E[(f \rightsquigarrow (\lambda x. \text{join } x))^\dagger]) \\
 \\
 \text{(RED-JOINP)} \\
 (\text{task}_g^\alpha E[\text{join } (v_1 \parallel v_2)]) \rightarrow (\text{task}_g^\alpha E[(\text{join } v_1) \parallel (\text{join } v_2)])
 \end{array}$$

Fig. 7. Reduction rules for the join combinator.

$$\begin{array}{c}
 \text{(RED-PRUNE)} \\
 \frac{\text{fresh } f}{(\text{task}_g^\alpha E[v \ll v']) \rightarrow (\text{fut}_f) (\text{task}_f^\alpha (\text{peek } v')) (\text{task}_g^\alpha E[v f])} \\
 \\
 \text{(RED-PEEK}^\circ\text{S)} \qquad \qquad \qquad \text{(RED-PEEK}^\circ\text{S)} \\
 (\text{task}_g^\alpha E[\text{peek}^\circ \{\}]) \rightarrow \epsilon \qquad (\text{task}_g^\alpha E[\text{peek } \{\}]) (\text{fut}_g) \rightarrow (\text{fut}_g \text{ Nothing}) (\text{poison } g) \\
 \\
 \text{(RED-PEEK}^\circ\text{V)} \\
 (\text{task}_g^\alpha E[\text{peek}^\circ (\{v\} \parallel v')]) (\text{fut}_g) \rightarrow (\text{fut}_g (\text{Just } v)) (\text{poison } g) \bigcup_{h \in \text{deps}(v')} (\text{poison } h) \\
 \\
 \text{(RED-PEEK}^\circ\text{F)} \\
 (\text{task}_g^\alpha E[\text{peek}^\circ (f^\circ \parallel v)]) \rightarrow (\text{chain}_g^\alpha f (\lambda x. \text{peek}^\circ \{x\})) (\text{task}_g^\alpha (\text{peek}^\circ v)) \\
 \\
 \text{(RED-PEEK}^\circ\text{FP)} \\
 \frac{\text{fresh } h}{(\text{task}_g^\alpha E[\text{peek}^\circ (f^\dagger \parallel v)]) \rightarrow} \\
 (\text{chain}_g^\alpha f (\lambda x. \text{peek}^\circ (x \parallel (h \rightsquigarrow \text{maybe2par})^\dagger))) \\
 (\text{fut}_h) (\text{task}_h^\alpha (\text{peek } v)) (\text{chain}_g^\alpha h (\lambda x. \text{peek}^\circ (\text{maybe2par } x)))
 \end{array}$$

Fig. 8. Reduction rules for pruning. For singleton collections are handled via equality $v = v \parallel \{\}$.

bunch of computations that compete to write into a single future, with the strict requirement that **Nothing** is written only when all competing tasks cannot produce a value—that is, when the ParT being peeked is empty. This is challenging due to the lifted future ParTs (type *Fut (Par t)*) within a collection, because such a future may be empty, but this fact cannot easily be seen in a non-blocking way. Another challenge is to avoid introducing sequential dependencies between entities that can potentially run in parallel, to avoid, for instance, a non-terminating computation blocking one that will produce a result.

A task that produces a ParT containing a value (rule RED-PEEK[∘]V) writes the value, wrapped in an option type, into the future and poisons all computations writing into that future, recursively poisoning direct dependencies. The [∘] status on **peek** prevents certain **peek** invocations from writing a final empty result, as in rule RED-PEEK[∘]S. Contrast with RED-PEEK[∘]S, in which a task resulting in an empty ParT writes **Nothing** into the future — in this case it is guaranteed that no other **peek** exists writing to the future.

$$\begin{array}{c}
\text{(RED-SCHEDULE)} \\
\text{fresh } f \\
\hline
(\mathbf{task}_g^\alpha E[e_1 \parallel e_2]) \rightarrow (\mathbf{task}_g^\alpha E[e_1 \parallel f^\dagger]) (\mathbf{fut}_f) (\mathbf{task}_f^\alpha e_2) \\
\hline
\end{array}$$

Fig. 9. Spawning of tasks inside a ParT.

A lifted future f is guaranteed to produce a result, though it may not produce it in a timely fashion. This case is handled (rule RED-PEEKF) by chaining a function onto it that will ultimately write into future g when the value is produced, if it wins the race. Otherwise, the result of peeking into v is written into g , unless the value produced is $\{\}$ (which is controlled by \emptyset).

A lifted future to a ParT is not necessarily guaranteed to produce a result, and neither is any ParT that runs in parallel with it. Thus, extra care needs to be taken to ensure that **Nothing** is written if and only if both are actually empty. This is handled in rule RED-PEEKFP. Firstly, a function is chained onto the lifted future to get access to the eventual ParT collection. This is combined with future h that is used to peek into v via a new task.

In all cases, computations propagate the poison state α to new configurations.

Scheduling. Rule RED-SCHEDULE (Fig. 9) models the non-deterministic scheduling of parallelism within a task, converting some of the parallelism latent in a ParT collection into a new task. Apart from this rule, expressions within tasks are evaluated sequentially.

$$\begin{array}{c}
\text{(RED-POISON)} \\
(\mathbf{poison } f) (PC_f e) \rightarrow (\mathbf{poison } f) (PC_f^\dagger e) \cup (\mathbf{poison } g) \\
\qquad\qquad\qquad g \in \text{deps}((PC_f e)) \\
\text{(RED-TERMINATE)} \\
\frac{\neg(\mathbf{config} \vdash \text{needed}(f))}{\{(PC_f^\dagger e) \mathbf{config}\} \rightarrow \{\mathbf{config}\}} \\
\hline
\end{array}$$

Fig. 10. Poisoning reduction rules.

Poisoning and Termination. The rules for poisoning and termination (Fig. 10) are based on a *poisoned carrier configuration* defined as $(PC_f^\alpha e) ::= (\mathbf{task}_f^\alpha e) \mid (\mathbf{chain}_f^\alpha g e)$; these rules rely on the definition of when a future is needed (Definition 2), which in turn is defined in terms of the futures on which a task depends to produce a value (Definition 1).

Definition 1. *The dependencies of an expression e , $\text{deps}(e)$, is the set of the futures upon which the computation of e depends in order to produce a value:*

$$\begin{aligned}
\text{deps}(f) &= \{f\} \\
\text{deps}(c) &= \text{deps}(\{\}) = \text{deps}(x) = \emptyset
\end{aligned}$$

$$\begin{aligned}
 \text{deps}(\{e\}) &= \text{deps}(\lambda x. e) = \text{deps}(\text{async } e) = \text{deps}(e^\circ) = \text{deps}(e^\dagger) = \\
 &\quad \text{deps}(\text{peek}^\pi e) = \text{deps}(\text{join } e) = \text{deps}(e) \\
 \text{deps}(e \parallel e') &= \text{deps}(e \ e') = \text{deps}(e \gg e') = \text{deps}(e \ggg e') = \\
 &\quad \text{deps}(e \times e') = \text{deps}(e \rightsquigarrow e') = \text{deps}(e \ll e') = \text{deps}(e) \cup \text{deps}(e') \\
 \text{deps}(\text{task}_f^\alpha e) &= \text{deps}(e) \\
 \text{deps}(\text{chain}_f^\alpha g e) &= \{g\} \cup \text{deps}(e).
 \end{aligned}$$

Definition 2. A future f is needed in configuration config , denoted $\text{config} \vdash \text{needed}(f)$, whenever some other element of the configuration depends on it:

$$\text{config} \vdash \text{needed}(f) \text{ iff } (PC_g^\alpha e) \in \text{config} \wedge f \in \text{deps}((PC_g^\alpha e)) \wedge (\text{fut}_f) \in \text{config}.$$

Configurations go through a two step process before being terminated. In the first step (rule RED-POISON) the poisoning of future f poisons any task or chain writing to f , marks it with poison , and the poison is transmitted to the direct dependencies of the expression e in the task or chain. In the second step (RED-TERMINATE), a poisoned configuration is terminated when there is no other configuration relying on its result — that is, a poisoned task or chain is terminated if there is no expression around to keep it alive. This rule is global, referring to the entire configuration. Termination can be implemented using tracing garbage collection, though in the semantics a more global specification of dependency is used.

An example (Fig. 11) illustrates how poisoning and termination work to prevent a task that is still needed from being terminated. Initially, there is a bunch of tasks (squares) and futures (circles) (Fig. 11A), where one of the tasks completes and writes a value to future f . This causes all of the other tasks writing to f to be poisoned, via Rule RED-PEEKV (Fig. 11B). After application of rule RED-POISON, the dependent tasks and futures are recursively poisoned (Fig. 11C). Finally, the application of rule RED-TERMINATE terminates tasks that are not needed (Fig. 11D). Task e_1 is not terminated, as future g is required by the task computing $e \ g$.

Configurations. The concatenation operation on configurations is commutative and associative and has the empty configuration as its unit (Fig. 12). We assume that these equivalences, along with the monoid axioms for \parallel , can be applied at any time during reduction.

The reduction rules for configurations (Fig. 13) have the individual configuration reduction rules at their heart, along with standard rules for parallel evaluation of non-conflicting sub-configurations, as is standard in rewriting logic [14].

3.4 Type System

The type system (Fig. 14) assigns the following types to terms:

$$\tau ::= K \mid \text{Fut } \tau \mid \text{Par } \tau \mid \text{Maybe } \tau \mid \tau \rightarrow \tau$$




Fig. 11. Safely poisoning and terminating a configuration. The letter in the top right corner indicates the order. Tasks are represented by squares, contain a body and have an arrow to the future they write to. Futures (circles) have dotted arrows to tasks that use them. Grey represents poisoned configurations. Terminated configurations are removed (Color figure online).

$$\epsilon \text{ config} \equiv \text{config} \quad \epsilon \equiv \text{config} \quad \text{config} \text{ config}' \equiv \text{config}' \text{ config}$$

$$\text{config} (\text{config}' \text{ config}'') \equiv (\text{config} \text{ config}') \text{ config}'' \quad \frac{\text{config} \equiv \text{config}'}{\{\text{config}\} \equiv \{\text{config}'\}}$$

Fig. 12. Configuration equivalence modulo associativity and commutativity.

$$\frac{\text{config} \rightarrow \text{config}'}{\text{config} \text{ config}'' \rightarrow \text{config}' \text{ config}''} \quad \frac{\text{config}_0 \rightarrow \text{config}'_0 \quad \text{config}_1 \rightarrow \text{config}'_1}{\text{config}_0 \text{ config}_1 \rightarrow \text{config}'_0 \text{ config}'_1}$$

$$\frac{\text{config}_0 \text{ config}'' \rightarrow \text{config}'_0 \text{ config}'' \quad \text{config}_1 \text{ config}'' \rightarrow \text{config}'_1 \text{ config}''}{\text{config}_0 \text{ config}_1 \text{ config}'' \rightarrow \text{config}'_0 \text{ config}'_1 \text{ config}''}$$

$$\frac{\text{config} \rightarrow \text{config}'}{\{\text{config}\} \rightarrow \{\text{config}'\}}$$

Fig. 13. Configuration reduction rules

where K represents the basic types, Fut τ is the type of a future containing a value of type τ , Par τ is the type of a ParT collection of type τ , $Maybe$ τ represents an option type, and $\tau \rightarrow \tau$ represents function types. We also let ρ range over types.

(TS-CONST) $\frac{c \text{ is a constant of type } \tau}{\Gamma \vdash_{\rho} c : \tau}$	(TS-FUT) $\frac{f : Fut \ \tau \in \Gamma}{\Gamma \vdash_{\rho} f : Fut \ \tau}$	(TS-X) $\frac{x : \tau \in \Gamma}{\Gamma \vdash_{\rho} x : \tau}$	(TS-APP) $\frac{\Gamma \vdash_{\rho} e_1 : \tau' \rightarrow \tau \quad \Gamma \vdash_{\rho} e_2 : \tau'}{\Gamma \vdash_{\rho} e_1 e_2 : \tau}$
(TS-FUN) $\frac{\Gamma, x : \tau \vdash_{\rho} e : \tau'}{\Gamma \vdash_{\rho} \lambda x. e : \tau \rightarrow \tau'}$	(TS-ASYNC) $\frac{\Gamma \vdash_{\rho} e : \tau}{\Gamma \vdash_{\rho} \text{async } e : Fut \ \tau}$	(TS-CHAIN) $\frac{\Gamma \vdash_{\rho} e_1 : Fut \ \tau' \quad \Gamma \vdash_{\rho} e_2 : \tau' \rightarrow \tau}{\Gamma \vdash_{\rho} e_1 \rightsquigarrow e_2 : Fut \ \tau}$	
(TS-EMPTYPAR) $\frac{}{\Gamma \vdash_{\rho} \{\} : Par \ \tau}$	(TS-SINGLETONPAR) $\frac{\Gamma \vdash_{\rho} e : \tau}{\Gamma \vdash_{\rho} \{e\} : Par \ \tau}$	(TS-LIFTF) $\frac{\Gamma \vdash_{\rho} e : Fut \ \tau}{\Gamma \vdash_{\rho} e^{\circ} : Par \ \tau}$	(TS-LIFTFP) $\frac{\Gamma \vdash_{\rho} e : Fut \ (Par \ \tau)}{\Gamma \vdash_{\rho} e^{\dagger} : Par \ \tau}$
(TS-PAR) $\frac{\Gamma \vdash_{\rho} e_1 : Par \ \tau \quad \Gamma \vdash_{\rho} e_2 : Par \ \tau}{\Gamma \vdash_{\rho} e_1 \parallel e_2 : Par \ \tau}$		(TS-SEQUENCE) $\frac{\Gamma \vdash_{\rho} e_1 : Par \ \tau' \quad \Gamma \vdash_{\rho} e_2 : \tau' \rightarrow \tau}{\Gamma \vdash_{\rho} e_1 \gg e_2 : Par \ \tau}$	
(TS-JOIN) $\frac{\Gamma \vdash_{\rho} e : Par \ (Par \ \tau)}{\Gamma \vdash_{\rho} \text{join } e : Par \ \tau}$	(TS-OTHERWISE) $\frac{\Gamma \vdash_{\rho} e_1 : Par \ \tau \quad \Gamma \vdash_{\rho} e_2 : Par \ \tau}{\Gamma \vdash_{\rho} e_1 \times e_2 : Par \ \tau}$	(TS-PEEK) $\frac{\Gamma \vdash_{Maybe \ \rho} e : Par \ \rho}{\Gamma \vdash_{Maybe \ \rho} \text{peek}^{\pi} e : \tau}$	
(TS-PRUNE) $\frac{\Gamma \vdash_{\rho} e_1 : Fut \ (Maybe \ \tau) \rightarrow Par \ \tau' \quad \Gamma \vdash_{\rho} e_2 : Par \ \tau}{\Gamma \vdash_{\rho} e_1 \ll e_2 : Par \ \tau'}$		(TS-BIND) $\frac{\Gamma \vdash_{\rho} e_1 : Par \ \tau' \quad \Gamma \vdash_{\rho} e_2 : \tau' \rightarrow Par \ \tau}{\Gamma \vdash_{\rho} e_1 \ggg e_2 : Par \ \tau}$	

Fig. 14. Expression Typing.

The key judgement in the type system is $\Gamma \vdash_{\rho} e : \tau$ which asserts that, in typing context Γ , the expression e is a well-formed term with type τ , where ρ is the expected return type of the task in which this expression appears — ρ is required to type `peek`. The typing context contains the types of both free variables and futures.

Rule TS-ASYNC gives the type for task creation and rule TS-CHAIN shows how to operate on such values — future chaining has the type of map for the `Fut` constructor. Rules TS-EMPTYPAR, TS-SINGLETONPAR, TS-LIFTF, TS-LIFTFP, and TS-PAR give the typings for constructing ParT collections. Rule TS-SEQUENCE implies that sequencing has the type of map for the `Par` constructor. TS-BIND and TS-JOIN give \ggg and `join` the types of the monadic bind and join operators for the `Par` constructor, respectively. Rule TS-PRUNE captures the communication between the two parameters via the future passed as an argument to the first parameter — the future will contain the first value of the second parameter if there is one, captured by the `Maybe` type. Rule TS-PEEK captures the conversion of the singleton or empty argument of `peek` from `Par` ρ to `Maybe` ρ , the expected result type of the surrounding task. Because `peek` terminates the task and does not return locally, its return type can be any type.

Well-formed configurations (Fig. 15) are expressed by the judgement $\Gamma \vdash \text{config} \text{ ok}$, where Γ is the assumptions about the future types in `config`. Rules T-TASK and T-CHAIN propagate the eventual expected result type on the turn-style \vdash when typing the enclosed expression. Rule T-CONFIG depends upon the following definition, a function that collects all futures defined in a configuration:

$\frac{}{\Gamma \vdash (\mathbf{fut}_f) \text{ ok}} \quad \text{(T-FUT)}$	$\frac{f : \mathbf{Fut} \tau \in \Gamma \quad \Gamma \vdash_\tau v : \tau}{\Gamma \vdash (\mathbf{fut}_f v) \text{ ok}} \quad \text{(T-FUTV)}$	$\frac{f : \mathbf{Fut} \tau \in \Gamma}{\Gamma \vdash (\mathbf{poison} f) \text{ ok}} \quad \text{(T-POISON)}$
$\frac{f : \mathbf{Fut} \tau \in \Gamma \quad \Gamma \vdash_\tau e : \tau}{\Gamma \vdash (\mathbf{task}_f^\alpha e) \text{ ok}} \quad \text{(T-TASK)}$	$\frac{f_1 : \mathbf{Fut} \tau_1 \in \Gamma \quad f_2 : \mathbf{Fut} \tau_2 \in \Gamma \quad \Gamma \vdash_{\tau_2} e : \tau_1 \rightarrow \tau_2}{\Gamma \vdash (\mathbf{chain}_{f_2}^\alpha f_1 e) \text{ ok}} \quad \text{(T-CHAIN)}$	
$\frac{\Gamma \vdash \mathit{config}_1 \text{ ok} \quad \Gamma \vdash \mathit{config}_2 \text{ ok} \quad \mathit{futset}(\mathit{config}_1) \cap \mathit{futset}(\mathit{config}_2) = \emptyset}{\Gamma \vdash \mathit{config}_1 \mathit{config}_2 \text{ ok}} \quad \text{(T-CONFIG)}$		
$\frac{\Gamma \vdash \mathit{config} \text{ ok} \quad \mathit{dom}(\Gamma) = \mathit{futset}(\mathit{config}) \quad \mathbf{TASKSAFE}(\mathit{config}) \quad \mathbf{ACYCLICDEP}(\mathit{config})}{\Gamma \vdash \{\mathit{config}\} \text{ ok}} \quad \text{(T-GCONFIG)}$		

Fig. 15. Configuration Typing.

Definition 3. Define $\mathit{futset}(\mathit{config})$ as:

$$\begin{aligned} \mathit{futset}((\mathbf{fut}_f)) &= \mathit{futset}((\mathbf{fut}_f v)) = \{f\} \\ \mathit{futset}((\mathit{config}_1 \mathit{config}_2)) &= \mathit{futset}(\mathit{config}_1) \cup \mathit{futset}(\mathit{config}_2) \\ \mathit{futset}(-) &= \emptyset. \end{aligned}$$

Rule T-GCONFIG defines the well-formedness of global configurations, judgement $\Gamma \vdash \{\mathit{config}\} \text{ ok}$. This rule depends on a number of definitions that capture properties on futures and tasks and on the dependency between futures. The invariance of these properties is ultimately used to proof type soundness and other safety properties of the system.

Definition 4. Define the following functions for collecting the different kinds of tasks and chains of a configuration:

$$\begin{aligned} \mathit{regular}_f(\mathit{config}) &= \{(\mathbf{task}_f e) \in \mathit{config} \mid e \neq \mathbf{peek}^\pi e'\} \\ &\quad \cup \{(\mathbf{chain}_f g e) \in \mathit{config} \mid e \neq \lambda _ . \mathbf{peek}^\pi e'\} \\ \mathit{peeker}_f(\mathit{config}) &= \{(\mathbf{task}_f (\mathbf{peek} e)) \in \mathit{config}\} \\ &\quad \cup \{(\mathbf{chain}_f g (\lambda _ . \mathbf{peek} e)) \in \mathit{config}\} \\ \mathit{peeker}_f^\circ(\mathit{config}) &= \{(\mathbf{task}_f (\mathbf{peek}^\circ e)) \in \mathit{config}\} \\ &\quad \cup \{((\mathbf{chain}_f g (\lambda _ . \mathbf{peek}^\circ e)) \in \mathit{config})\} \end{aligned}$$

Tasks with no \mathbf{peek} expression are called *regular tasks*, while *peeker tasks* have the \mathbf{peek} expression — there are both \circ - and non- \circ -peeker tasks. These functions can be used to partition the tasks and chains in a configuration into these three kinds of tasks and chains. These definitions consider \mathbf{peek} expressions only at the top level of a task, although the syntax allows them to be anywhere. Based on the reduction rules, one can prove that \mathbf{peek} only appears at the top level of a task or chain, so no task or chain is excluded by these definitions.

Definition 5. Define predicate $\text{TASKSAFE}(config)$ as follows:

$$\begin{aligned}
 & \text{TASKSAFE}(config) \text{ iff for all } f \in \text{futset}(config) \\
 & \quad |regular_f(config) \cup peeker_f(config)| \leq 1 \\
 & \wedge (\text{fut}_f) \in config \wedge (\text{poison } f) \notin config \Rightarrow \\
 & \quad |regular_f(config) \cup peeker_f(config)| = 1 \\
 & \wedge |regular_f(config)| = 1 \Rightarrow peeker_f(config) \cup peeker_f^\circ(config) = \emptyset \\
 & \wedge (\text{task}_f^\alpha(\text{peek } \{ \})) \in config \Rightarrow (\text{fut}_f) \in config \wedge peeker_f^\circ(config) = \emptyset
 \end{aligned}$$

Predicate $\text{TASKSAFE}(config)$ (Definition 5) describes the structure of the configuration $config$. It states that:

- there is at most one regular or non- \circ -peeker task per future;
- if a future has not yet been fulfilled and it is not poisoned, then there exists exactly one regular task or non- \circ -peeker task that fulfils it;
- regular tasks and peeker tasks do not write to the same futures; and
- if a peeker task that is about to fulfil a future with **Nothing**, then the future is unfulfilled and no \circ -peeker task fulfilling the same future exists.

The following definition establishes dependencies between futures. Predicate $config \vdash f \triangleleft g$ holds for all future g whose eventual value could influence the result stored in future f .

Definition 6. Define the predicate $config \vdash f \triangleleft g$ as the least transitive relation satisfying the following rules:

$$\begin{array}{c}
 \frac{(\text{task}_f^\alpha e) \in config \quad g \in \text{deps}(e)}{config \vdash f \triangleleft g} \quad \frac{(\text{chain}_f^\alpha h e) \in config \quad g \in \text{deps}(e) \cup \{h\}}{config \vdash f \triangleleft g} \\
 \\
 \frac{(\text{fut}_f v) \in config \quad g \in \text{deps}(v)}{config \vdash f \triangleleft g}
 \end{array}$$

Definition 7. Predicate $\text{ACYCLICDEP}(config)$ holds iff relation \triangleleft is acyclic, where \triangleleft is defined for $config$ in Definition 6.

Rule T-GCONFIG for well-formed global configurations requires that precisely the futures that appear in the typing environment Γ appear in the configuration, that the configuration is well-formed, and that it satisfies the properties TASKSAFE and ACYCLICDEP . By including these properties as a part of the well-formedness rule for global configurations, type preservation (Lemma 1) makes these invariants. These invariants on the structure of tasks and the dependency relation together ensure that well-typed configurations are deadlock-free, as we explore next.

3.5 Formal Properties

The calculus is sound and deadlock-free. These results extend previous work [15] to address the pruning combinator.

Lemma 1 (Type Preservation). *If $\Gamma \vdash \{\text{config}\} \text{ok}$ and $\{\text{config}\} \rightarrow \{\text{config}'\}$, then there exists a Γ' such that $\Gamma' \supseteq \Gamma$ and $\Gamma' \vdash \{\text{config}'\} \text{ok}$.*

Proof. By induction on derivation $\{\text{config}\} \rightarrow \{\text{config}'\}$. In particular, the invariance of ACYCLICDEP is shown by considering the changes to the dependencies caused by each reduction rule. The only place where new dependencies are introduced is when new futures are created. Adding a future to the dependency relation cannot introduce cycles. \square

The following lemma states that the notion of *needed*, which determines whether or not to garbage collect a poisoned task or chain, is anti-monotonic, meaning that after a future is no longer needed according to the definitions, it does not subsequently become needed.

Lemma 2 (Safe Task Kill). *If $\Gamma \vdash \{\text{config}\} \text{ok}$ and $\{\text{config}\} \rightarrow \{\text{config}'\}$, then $\neg(\text{config} \vdash \text{needed}(f))$ implies $\neg(\text{config}' \vdash \text{needed}(f))$.*

Proof. A future is initially created in a configuration where it is needed. If ever a future disappears from $\text{deps}(e)$, it can never reappear. \square

This lemma rules out the situation where a task is poisoned and garbage collected, but is subsequently needed. For instance, the application of rule RED-TERMINATE in Fig. 11C kills tasks e_2, e_3, e_5 and e_6 (shown in Fig. 11D). If the future into which these tasks were going to write is needed afterwards, there would be a deadlock as a new task could chain on that future but never be fulfilled.

Definition 8 (Terminal Configuration). *A global configuration $\{\text{config}\}$ is terminal iff every element of config has one of the following shapes: (fut_f) , $(\text{fut}_f v)$ or $(\text{poison } f)$.*

Lemma 3 (Deadlock-Freedom/Progress). *If $\Gamma \vdash \{\text{config}\} \text{ok}$, then config is a terminal configuration, or there exists a config' such that $\{\text{config}\} \rightarrow \{\text{config}'\}$.*

Proof. By induction on a derivation of $\{\text{config}\} \rightarrow \{\text{config}'\}$, relying on the invariance of ACYCLICDEP and Lemma 2. \square

Deadlock-freedom guarantees that some reduction rule can be applied to a well-typed, non terminal, global configuration — this is essentially the progress property required to prove type safety. It implies further that there are no local deadlocks, such as a deadlocked configuration like $(\text{chain}_f g e) (\text{chain}_g f e')$. Such a configuration fails to satisfy the ACYCLICDEP invariant, thus cannot exist. If mutable state is added to the calculus, deadlock-freedom is lost.

Implementations. There are two prototypes of the ParT abstraction. In the first prototype,² ParT has been written as an extension to the Encore compiler (written in Haskell) and runtime (written in C) but, it can be implemented in well-established languages with notions of tasks and futures. This prototype integrates futures produced by tasks and active objects with the ParT abstraction. The other prototype has been written in Clojure,³ which is not statically typed. Both prototypes follow the semantics to guide the implementation. In practice, this means that the semantic rules are written in such a way that they can be easily mimicked in a library or in a language runtime.

4 Related Work

Our combinators have been adapted from those of the Orc [11, 12] programming language. In ParT, these combinators are completely asynchronous and are integrated with futures. ParTs are first class citizens and can be nested *Par* (*Par t*), neither of which is possible in Orc, which sits on top of the expression being coordinated and a flat collection of values.

Meseguer et al. [1] used rewriting logic semantics and Maude to provide a distributed implementation of Orc. Their focus on the semantic model allows them to model check Orc programs. In this paper, our semantics is more fine-grained, and guides the implementation in a multicore setting.

ParT uses a monad to encapsulate asynchronous computations, which is not a new idea [3, 13, 20]. For instance, $F\#$ expresses asynchronous workflows using a continuation monad [20] but cannot create more parallelism within the monad, making the model better suited for event-based programming. In contrast, our approach can spawn parallel computations and include them within ParTs.

Other work implements Orc combinators in terms of a monad within the pure functional language Haskell [3, 13]. One of these approaches [3] relies on threads and channels and implements the *prune* \ll combinator using sequential composition, losing potential parallelism. The other approach [13] uses Haskell threads and continuations to model parallel computations and re-designs the *prune* \ll combinator in terms of a *cut* combinator that sparks off parallel computations, waits until there is a value available and terminates, in bulk, the remaining computations. In contrast, the ParT abstraction relies on more lightweight tasks instead of threads, has fully asynchronous combinators, which maintain the throughput of the system, and terminates speculative work by recursively poisoning dependencies and terminating computations that are not needed.

An approach to increase parallelism is to create parallel versions of existing collections. For instance, Haskell [10] adds parallel operations to its collections, and the Scala parallel collections [18] adds new methods to their collection, `par` and `seq`, that return a parallel and a sequential version of the collection. However these approaches cannot coordinate complex workflows, which is possible with the ParT abstraction.

² Encore ParT prototype: <http://52.50.101.143/kompile/encore/>.

³ Clojure ParT prototype: <https://github.com/kikofernandez/ParT>.

Recent approaches to creating pipeline parallelism are the Flowpool [19] and FlumeJava [4] abstractions. In the former, functions are attached to Flowpool and, with the `foreach` combinator, the attached functions are applied to items asynchronously added to the Flowpool thereby creating parallel pipelines of computations. The latter, FlumeJava, is a library extending the MapReduce framework; it provides high-level constructs to create efficient data-parallel pipelines of MapReduce jobs, via an optimisation phase. The ParT abstraction can create data-parallel pipelines with the sequence \gg and bind $\gg=$ combinators (at the moment there is no optimisation phase) and further can terminate speculative work.

Existing approaches to safely terminating speculative parallelism [6, 9, 17] did not integrate well with the ParT abstraction. For instance, the Cilk programming language provides the `abort` keyword to terminate all speculative work generated by a procedure [6]. The termination does not happen immediately, instead, computations are marked as non-runnable; already running computations would get marked as non-runnable but do not stop execution until their work is finished. In other approaches, the developer specifies termination checkpoints at which a task may be terminated [9, 17]. This solves the previous problem and improves responsiveness but, adds an extra overhead (for the checking) and puts the responsibility on the developer, who specifies the location of the checkpoints. In our design, the developer does not need to specify these checkpoints and speculative work is terminated as soon as there are no dependencies. No other approach considers that the results of tasks may be needed elsewhere.

5 Conclusion and Future Work

This paper presented the ParT asynchronous, parallel collection abstraction, and a collection of combinators that operate over it. ParT was formalised as a typed calculus of tasks, futures and Orc-like combinators. A primary characteristic of the calculus is that it captures the non-blocking implementation of the combinators, including an algorithm for pruning that tracks down dependencies and is safe with respect to shared futures. The ParT abstraction has prototypes in the Encore (statically typed) and Clojure (dynamically typed) programming languages.

Currently, the calculus does not support side-effects. These are challenging to deal with, due to potential race conditions and terminated computations leaving objects in an inconsistent state. We expect that Encore’s capability type system [2] can be used to avoid data races, and a run-time, transactional mechanism can deal with the inconsistent state. At the start of the paper we mentioned that ParT was integrated into an actor-based language, but the formalism included no actors. This work abstracted away the actors, replacing them by tasks and futures—message sends in the Encore programming language return results via futures—which were crucial for tying together the asynchronous computations underlying a ParT. Actors can easily be re-added as soon as the issues of shared mutable state have been addressed. The distribution aspect of actors has not yet

been considered in Encore or in the ParT abstraction. This would be an interesting topic for future work. Beyond these extensions, we also plan to extend the range of combinators supporting the ParT abstraction.

References

1. AlTurki, M., Meseguer, J.: Dist-Orc: A rewriting-based distributed implementation of Orc with formal analysis. In: Ólveczky, P.C., (ed.) Proceedings First International Workshop on Rewriting Techniques for Real-Time Systems, RTRTS 2010, 6–9 April, 2010, vol. 36 of EPTCS, pp. 26–45. Longyearbyen, Norway (2010)
2. Brandauer, S., et al.: Parallel objects for multicores: a glimpse at the parallel language encore. In: Bernardo, M., Johnsen, E.B. (eds.) Formal Methods for Multicore Programming. LNCS, pp. 1–56. Springer, Switzerland (2015)
3. Campos, M.D., Barbosa, L.S.: BarbosaImplementation of an orchestration language as a haskell domain specific language. *Electr. Notes Theor. Comput. Sci.* **255**, 45–64 (2009)
4. Chambers, C., Raniwala, A., Perry, F., Adams, S., Henry, R.R., Bradshaw, R., Weizenbaum, N.: Flumejava: Easy, efficient data-parallel pipelines. In: Proceedings of the 31st ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2010, pp. 363–375, New York, NY, USA. ACM (2010)
5. Clarke, D., Wrigstad, T.: Vats: a safe, reactive storage abstraction. In: Ábrahám, E., Bonsangue, M., Johnsen, E.B. (eds.) Theory and Practice of Formal Methods. LNCS, vol. 9660, pp. 140–154. Springer, Switzerland (2016)
6. Frigo, M., Leiserson, C.E., Randall, K.H.: The implementation of the Cilk-5 multithreaded language. In: Davidson, J.W., Cooper, K.D., Berman, A.M., (eds.) Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation (PLDI), pp. 212–223. ACM, Montreal, 17–19 June 1998
7. Robert Jr., R.H.: Multilisp: A language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.* **7**(4), 501–538 (1985)
8. Hickey, R.: The clojure programming language. In: Brichau, J. (ed.) Proceedings of the 2008 Symposium on Dynamic Languages, DLS 2008, p. 1. ACM, Paphos, Cyprus (2008)
9. Imam, S., Sarkar, V.: The Eureka programming model for speculative task parallelism. In: Boyland, J.T. (ed.) 29th European Conference on Object-Oriented Programming, ECOOP 2015, vol. 37 of LIPIcs, pp. 421–444, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Prague, Czech Republic, 5–10 July 2015
10. Peyton Jones, S.: Harnessing the multicores: nested data parallelism in haskell. In: Ramalingam, G. (ed.) APLAS 2008. LNCS, vol. 5356, p. 138. Springer, Heidelberg (2008)
11. Kitchin, D.E., Cook, W.R., Misra, J.: A language for task orchestration and its semantic properties. In: Baier, C., Hermanns, H. (eds.) CONCUR 2006. LNCS, vol. 4137, pp. 477–491. Springer, Heidelberg (2006)
12. Kitchin, D., Quark, A., Cook, W., Misra, J.: The Orc programming language. In: Lee, D., Lopes, A., Poetzsch-Heffter, A. (eds.) FMOODS 2009. LNCS, vol. 5522, pp. 1–25. Springer, Heidelberg (2009)
13. Launchbury, J., Elliott, T.: Concurrent orchestration in haskell. In: Gibbons, J., (ed.) Proceedings of the 3rd ACM SIGPLAN Symposium on Haskell, Haskell 2010, pp. 79–90. ACM, Baltimore, MD, USA, 30 September 2010

14. Martí-Oliet, N., Meseguer, J.: Rewriting logic: roadmap and bibliography. *Theor. Comput. Sci.* **285**(2), 121–154 (2002)
15. McCain, D.: Parallel combinators for the Encore programming language. Master's thesis, Uppsala University (2016)
16. Odersky, M., Spoon, L., Venners, B.: *Programming in Scala: A Comprehensive Step-by-step Guide*, 1st edn. Artima Incorporation, USA (2008)
17. Peierls, T., Goetz, B., Bloch, J., Bowbeer, J., Lea, D., Holmes, D.: *Java Concurrency in Practice*. Addison-Wesley Professional, Boston (2005)
18. Prokopec, A., Bagwell, P., Rompf, T., Odersky, M.: A generic parallel collection framework. In: Jeannot, E., Namyst, R., Roman, J. (eds.) *Euro-Par 2011, Part II*. LNCS, vol. 6853, pp. 136–147. Springer, Heidelberg (2011)
19. Prokopec, A., Miller, H., Schlatter, T., Haller, P., Odersky, M.: FlowPools: a lock-free deterministic concurrent dataflow abstraction. In: Kasahara, H., Kimura, K. (eds.) *LCPC 2012*. LNCS, vol. 7760, pp. 158–173. Springer, Heidelberg (2013)
20. Syme, D., Petricek, T., Lomov, D.: The F# asynchronous programming model. In: Rocha, R., Launchbury, J. (eds.) *PADL 2011*. LNCS, vol. 6539, pp. 175–189. Springer, Heidelberg (2011)
21. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Inf. Comput.* **115**(1), 38–94 (1994)