Forward to a Promising Future

Kiko Fernandez-Reyes, Dave Clarke, Elias Castegren, Huu-Phuc Vo

► To cite this version:

Kiko Fernandez-Reyes, Dave Clarke, Elias Castegren, Huu-Phuc Vo. Forward to a Promising Future. 20th International Conference on Coordination Languages and Models (COORDINATION), Jun 2018, Madrid, Spain. pp.162-180, 10.1007/978-3-319-92408-3_7 . hal-01821490

HAL Id: hal-01821490

<https://hal.inria.fr/hal-01821490>

Submitted on 22 Jun 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Distributed under a Creative Commons Attribution| 4.0 International License

Forward to a Promising Future

Kiko Fernandez-Reyes^{<https://orcid.org/0000-0001-8654-118X>}, Dave
Clarke^{<https://orcid.org/0000-0002-1970-6607>}, Elias
Castegren^{<https://orcid.org/0000-0003-4918-6582>}, and Huu-Phuc Vo*

Dept. of Information Technology
Uppsala University
Uppsala, Sweden

Abstract. In many actor-based programming models,¹ asynchronous method calls communicate their results using futures, where the fulfilment occurs under-the-hood. Promises play a similar role to futures, except that they must be explicitly created and explicitly fulfilled; this makes promises more flexible than futures, though promises lack fulfilment guarantees: they can be fulfilled once, multiple times or not at all. Unfortunately, futures are too rigid to exploit many available concurrent and parallel patterns. For instance, many computations block on a future to get its result only to return that result immediately (to fulfil their own future). To make futures more flexible, we explore a construct, *forward*, that delegates the responsibility for fulfilling the current implicit future to another computation. *Forward* reduces synchronisation and gives futures promise-like capabilities. This paper presents a formalisation of the *forward* construct, defined in a high-level source language, and a compilation strategy from the high-level language to a low-level, promised-based target language. The translation is shown to preserve semantics. Based on this foundation, we describe the implementation of *forward* in the parallel, actor-based language Encore,² which compiles to C.

1 Introduction

Futures extend the actor programming model to express call-return synchronisation of message sends [1]. Each actor is single-threaded, but different actors execute concurrently. Communication between actors happens via asynchronous method calls (messages), which immediately return a future; futures are placeholders for the eventual result of these asynchronous method calls. An actor processes one message at a time and each message has associated a future that

* We are grateful to Joachim Parrow and Johannes Borgström for their comments regarding the bisimulation relation. We also thank the anonymous referees for their useful comments. The underlying research was funded by the Swedish VR project: SCADA.

¹ This paper focuses on futures. From this perspective we consider the actor-, task-, and active object-based models as synonymous.

² <https://github.com/parapluu/encore>

will be fulfilled with the returned value of the method. Futures are first-class values, and operations on them may be blocking, such as getting the result out of the future (`get`), or asynchronous, such as attaching a callback to a future. This last operation, known as future chaining ($f \overset{x}{\rightsquigarrow} e$), attaches a closure $\lambda x.e$ to the future f and immediately returns a new future that will contain the result of applying the closure to the value eventually stored in future f .

Consider the following code (in the actor-based language Encore [2]) that implements the broker delegation pattern: the developer's intention is to connect clients (the callers of the `Broker` actor) to a pool of actors that will actually process a job (lines 6–7):

```

1  active class Broker
2    val workers: Buffered[Worker]
3    var current: uint
4
5    def run(job: Job): int
6      val worker = this.workers[++this.current % workers.size()]
7      val future : Fut[int] = worker!start(job)
8      return get(future)
9    end
10 end

```

The problem with this code is that the connection to the `Broker` cannot be completed immediately without blocking the `Broker`'s thread of execution: returning the result of the worker running the computation requires that the `Broker` blocks until the future is fulfilled (line 8). This implementation makes the `Broker` the bottleneck of the application.

One obvious way to avoid this bottleneck is by returning the future, instead of blocking on it, as in the following code:

```

1  def run(job: Job): Fut[int]
2    val worker = this.workers[++this.current % workers.size()]
3    return worker!start(job)
4  end

```

This solution removes the blocking from `Broker`, but returns a future, which results in the client receiving a future containing a future `Fut (Fut int)`, cluttering client code and making the typing more complex.

Another way to avoid the bottleneck is to not block but yield the current thread until the future is fulfilled. This can be done using the `await` command [3, 2], which frees up the `Broker` to do other work:³

```

1  def run(job: Job): int
2    val worker = this.workers[++this.current % workers.size()]
3    val future = worker!start(job)
4    await(future)
5    return get(future)
6  end

```

This solution frees up the `Broker`, but can result in a lot of memory being consumed to hold the waiting instances of calls `Broker.run()`.

³ The essential difference between `get` and `await` is that `get` blocks an actor, whereas `await` blocks only the current method invocation and frees up the actor.

Another alternative is to use promises [4]. A promise can be passed around and fulfilled explicitly at the point where the corresponding result is known. Passing a promise around is akin to passing the responsibility to provide a particular result, thereby fulfilling the promise.

```

1  def run(job: Job, promise: Promise[int]): unit
2    val worker = this.workers[++this.current % workers.size()]
3    worker!start(job, promise)
4  end
5
6  class Worker
7    def start(job: Job, promise: Promise[int]) : unit
8      // actually do job
9      promise.fulfil(result)
10   end
11 end

```

Promises are problematic because they diverge from the commonplace call-return control flow, there is no explicit requirement to actually fulfil a promise, and care is required to avoid fulfilling multiple times. This latter issue, fulfilling a promise multiple times, can be solved by a substructural type system, which guarantees a single writer to the promise [5, 6]. Substructural type systems are more complex and not mainstream, which rules out adoption in languages such as Java and C#. Our solution relies on futures and is suitable for mainstream languages.

The main difference between promises and futures are that developers *explicitly* create and fulfil promises, whereas futures are *implicitly* created and fulfilled. Promises are thus more flexible at the expense of any fulfilment guarantees.

This paper explores a construct called *forward* that retains the guarantees of using futures, while allowing some degree of delegation of responsibility to fulfil a future, as in promises. This construct was first proposed a while ago [7], but only recently has been implemented in the language Encore [2].

With *forward*, the `run` of `Broker` method now becomes:

```

1  def run(job: Job): int
2    val worker = this.workers[++this.current % workers.size()]
3    forward(worker!start(job))
4  end

```

Forward delegates the fulfilment of the future that `run` will put its result in, to the call `worker!start(job)`. Using *forward* frees up the `Broker` object, as `run` completes immediately, though the future is fulfilled only when `worker!start(job)` produces a result.

The paper makes the following contributions:

- a formalisation and soundness proof of the *forward* construct in a concise, high-level language (Section 2);
- a formalisation of a low-level, promise-based language (Section 3),
- a translation from the high-level language to the low-level language, a proof of program equivalence, between the high-level language and its translation to the low-level language (Section 4); and
- microbenchmarks that compare the `get-and-return` and `await-and-get` pattern versus the `forward` construct (Section 5).

2 A Core Calculus of Futures and Forward

This section presents a core calculus that includes tasks, futures and operations on them, and forward. The calculus consists of two levels: expressions and configurations. Expressions correspond to programs and what tasks evaluate. Configurations capture the run-time configuration; they are collections of tasks (running expressions), futures, and chains. This calculus is much more concise than the previous formalisation of forward [7].

The syntax of the core calculus is as follows:

$$\begin{aligned} e &::= v \mid e e \mid \text{async } e \mid e \overset{x}{\rightsquigarrow} e \mid \text{if } e \text{ then } e \text{ else } e \mid \text{forward } e \mid \text{get } e \\ v &::= c \mid f \mid x \mid \lambda x.e \end{aligned}$$

Expressions include values (v), function application ($e e$), spawning asynchronous computations ($\text{async } e$), future chaining ($e \overset{x}{\rightsquigarrow} e'$), which attaches $\lambda x.e'$ onto a future to run as soon as the future produced by e is fulfilled, *if-then-else* expressions, **forward**, and **get**, which extracts the value from a future. Values are constants (c), futures (f), variables (x) and lambda abstractions ($\lambda x.e$). The calculus has neither actors nor message sends/method calls. For our purposes, tasks play the role of actors and spawning asynchronous computations is analogous to message sends.

Configurations, *config*, give a partial view on the system and are (non-empty) multisets of tasks, futures and chains. They have the following syntax:

$$\text{config} ::= (\text{fut}_f) \mid (\text{fut}_f v) \mid (\text{task}_f e) \mid (\text{chain}_f f e) \mid \text{config config}$$

Future configurations are (fut_f) and $(\text{fut}_f v)$, representing an unfulfilled future f and a fulfilled future f with value v . Configuration $(\text{task}_f e)$ is a task running expression e that will write the result of e in future f .⁴ Configuration $(\text{chain}_f g e)$ denotes a computation that waits until future g is fulfilled, applies expression e to the value stored in g in a new task whose result will be stored in future f .

The initial configuration for program e is $(\text{task}_f e) (\text{fut}_f)$, where the result of e will be written into future f at the end result of the program's execution.

2.1 Operational Semantics

The operational semantics use a small-step semantics with reduction-based, contextual rules for evaluation within tasks. Evaluation contexts E contains a hole $\bullet$ that denotes where the next reduction step happens [8]:

$$E ::= \bullet \mid E e \mid v E \mid E \overset{x}{\rightsquigarrow} e \mid \text{forward } E \mid \text{get } E \mid \text{if } E \text{ then } e \text{ else } e$$

⁴ A reviewer suggested that (fut_f) , $(\text{fut}_f v)$, and $(\text{task}_f e)$ could be combined into a single configuration component. We have considered this conflation in the past. While it would reduce the complexity of the calculus, it would also make compilation into the target calculus and the proofs of correctness more complex.

(RED-IF-TRUE)	(RED-β)
$(task_f E[\text{if true then } e \text{ else } e']) \rightarrow (task_f E[e])$	$(task_f E[\lambda x.e v]) \rightarrow (task_f E[e[v/x]])$
(RED-IF-FALSE)	(RED-FWD-FUT)
$(task_f E[\text{if false then } e \text{ else } e']) \rightarrow (task_f E[e'])$	$(task_f E[\text{forward } h]) \rightarrow (chain_f h \lambda x.x)$
(RED-CHAIN-RUN)	(RED-GET)
$(chain_g f e) (fut_f v) \rightarrow (task_g (e v)) (fut_f v)$	$(task_f E[\text{get } h]) (fut_h v) \rightarrow (task_f E[v]) (fut_h v)$
(RED-FUT-FULFIL)	(RED-ASYNC)
$(task_f v) (fut_f v) \rightarrow (fut_f v)$	$\frac{fresh\ f}{(task_g E[\text{async } e]) \rightarrow (fut_f) (task_f e) (task_g E[f])}$
	(RED-CHAIN-CREATE)
	$\frac{fresh\ g}{(task_f E[h \rightsquigarrow e]) \rightarrow (fut_g) (chain_g h \lambda x.e) (task_f E[g])}$

Fig. 1: Reduction Rules. f, g, h range over futures.

$\frac{config \rightarrow config''}{config\ config' \rightarrow config''\ config'}$	$\frac{config \equiv config' \quad config' \rightarrow config'' \quad config'' \equiv config'''}{config \rightarrow config'''}$
---	---

Fig. 2: Configuration evaluation rules. Equivalence $\equiv$ (omitted) captures the fact that configurations are a multiset of basic configurations.

The evaluation rules are given in Fig. 1. The evaluation of *if-then-else* expressions and functions applications proceed in the standard fashion (RED-IF-TRUE, RED-IF-FALSE, and RED-β). The *async* construct spawns a new task to execute the given expression, and creates a new future to store its result (RED-ASYNC). When the spawned task finishes its execution, it places the value in the designated future (RED-FUT-FULFIL). To obtain the contents of a future, the blocking construct *get* stops the execution of the task until the future is fulfilled (RED-GET). Chaining an expression on a future results immediately in a new future that will eventually contain the result of evaluating the expression, and a chain configuration storing the expression is connected with the original future (RED-CHAIN-CREATE). When the future is fulfilled, any chain configurations become task configurations and start evaluating the stored expression on the value stored in the future (RED-CHAIN-RUN). Forward applies to a future where the result of the future computation will be the result of the current computation, stored in the future associated with the current task. Forwarding to future h throws away the remainder of the body of the current task and chains the identity function on the future, the effect of which is to copy the eventual result stored in h into the current future (RED-FWD-FUT).

The configuration evaluation rules (Fig. 2) describe how configurations make progress, which is either by some subconfiguration making progress, or by rewriting a configuration to one that will make progress using the equations of multi-sets.

Example and Optimisations The following example illustrates some aspects of the calculus.

$$\begin{aligned}
& (task_f E[async (forward h)]) (fut_h 42) \\
& \xrightarrow{\text{RED-ASYNC}} (task_f E[g]) (fut_h 42) (fut_g) (task_g forward h) \\
& \xrightarrow{\text{RED-FWD-FUT}} (task_f E[g]) (fut_h 42) (fut_g) (chain_g h \lambda x.x) \\
& \xrightarrow{\text{RED-CHAIN-RUN}} (task_f E[g]) (fut_h 42) (fut_g) (task_g (\lambda x.x) 42) \\
& \xrightarrow{\text{RED-}\beta} (task_f E[g]) (fut_h 42) (fut_g) (task_g 42) \\
& \xrightarrow{\text{RED-FUT-FULFIL}} (task_f E[g]) (fut_h 42) (fut_g 42)
\end{aligned}$$

Firstly, a new task is spawned with the use of **async**. This task forwards the responsibility to fulfil its future to (the task fulfilling) future h , i.e. future g gets fulfilled with the value contained in future h .

Two special cases of **forward** can be given more direct reduction sequences, which correspond to optimisations performed in the Encore compiler. The first case corresponds to forwarding directly to another method call, which is the primary use case for **forward**, namely, forwarding to another method **forward(e!m())**. The optimised reduction rule is

$$(task_f E[forward (async e)]) \rightarrow (task_f e)$$

For comparison, the standard reduction sequence⁵ is

$$\begin{aligned}
& (task_f E[forward (async e)]) \rightarrow (task_f E[forward g]) (task_g e) (fut_g) \\
& \rightarrow (chain_f g \lambda x.x) (task_g e) (fut_g) \rightarrow^* (chain_f g \lambda x.x) (task_g v) (fut_g) \\
& \rightarrow (chain_f g \lambda x.x) (fut_g v) \rightarrow (task_f (\lambda x.x) v) (fut_g v) \rightarrow (task_f v) (fut_g v)
\end{aligned}$$

This can be seen as equivalent to the reduction sequence

$$(task_f E[forward (async e)]) \rightarrow (task_f e) \rightarrow^* (task_f v)$$

because the future g will no longer be accessible.

Similarly, forwarding a future chain can be reduced directly to a chain configuration:

$$(task_f E[forward (h \rightsquigarrow^x e)]) \rightarrow (chain_f h \lambda x.e)$$

In both cases, forward can be seen as making a call-with-current-future.

⁵ $\rightarrow^*$ is the reflexive, transitive closure of the reduction relation $\rightarrow$.

(T-CONSTANT)	(T-FUTURE)	(T-VARIABLE)	(T-ABSTRACTION)
$\frac{\text{c is a constant of type } \tau}{\Gamma \vdash_{\rho} c : \tau}$	$\frac{f : Fut \tau \in \Gamma}{\Gamma \vdash_{\rho} f : Fut \tau}$	$\frac{x : \tau \in \Gamma}{\Gamma \vdash_{\rho} x : \tau}$	$\frac{\Gamma, x : \tau \vdash_{\bullet} e : \tau'}{\Gamma \vdash_{\rho} \lambda x. e : \tau \rightarrow \tau'}$
(T-APPLICATION)	(T-IF-THEN-ELSE)	(T-GET)	
$\frac{\Gamma \vdash_{\rho} e_1 : \tau \rightarrow \tau' \quad \Gamma \vdash_{\rho} e_2 : \tau}{\Gamma \vdash_{\rho} e_1 e_2 : \tau'}$	$\frac{\Gamma \vdash_{\rho} e : bool \quad \Gamma \vdash_{\rho} e' : \tau \quad \Gamma \vdash_{\rho} e'' : \tau}{\Gamma \vdash_{\rho} \text{if } e \text{ then } e' \text{ else } e'' : \tau}$	$\frac{\Gamma \vdash_{\rho} e : Fut \tau}{\Gamma \vdash_{\rho} \text{get } e : \tau}$	
(T-ASYNC)	(T-CHAIN)	(T-FORWARD)	
$\frac{\Gamma \vdash_{\tau} e : \tau}{\Gamma \vdash_{\rho} \text{async } e : Fut \tau}$	$\frac{\Gamma \vdash_{\rho} e : Fut \tau \quad \Gamma, x : \tau \vdash_{\tau'} e' : \tau'}{\Gamma \vdash_{\rho} e \overset{x}{\rightsquigarrow} e' : Fut \tau'}$	$\frac{\Gamma \vdash_{\rho} e : Fut \rho \quad \rho \neq \bullet}{\Gamma \vdash_{\rho} \text{forward } e : \tau}$	

Fig. 3: Typing Rules

2.2 Static Semantics

The type system has basic types, K , and future types:

$$\tau ::= K \mid Fut \tau$$

The typing rules (Fig. 3) define the judgement $\Gamma \vdash_{\rho} e : \tau$, which states that in the typing environment Γ , which gives the types of futures and free variables, expression e has type τ , where ρ is the *expected task type*, the result type of the task in which the expression appears. ρ ranges over both types τ and symbol $\bullet$ which is not a type. $\bullet$ is used to prevent the use of **forward** in contexts where the expected task type is not clear, specifically within closures, as a closure can be passed between tasks and run in a context different from their defining contexts. The types of constants are assumed to be provided (Rule T-CONSTANT). Variables and futures types are defined in the typing environment (Rules T-VARIABLE and T-FUTURE). Function application and abstraction have the standard typing rules (Rules T-APPLICATION and T-ABSTRACTION), except that within the body of a closure the expected task type is not known. When **async** is applied to an expression e , a new task is created and the expected task type changes to the type of the expression. The result type of the **async** call is a future type of the expression's type (Rule T-ASYNC). Chaining is essentially mapping for the *Fut* type constructor, and rule T-CHAIN reflects this fact. In addition, because chaining ultimately creates a new task to run the expression, the expected task type ρ changes to the return type of the expression. Getting the value from a future of some type results in a value of that type (Rule T-GET). Forwarding requires the argument to **forward** to be a future of the same type as the expected task type (Rule T-FORWARD). As **forward** does not return locally, the result type is arbitrary.

Well-formed configurations, $\Gamma \vdash \text{config } ok$, are typed against environment, Γ , that gives the types of futures (Fig. 4). The type rules depend on the following definitions.

$\frac{\text{(FUT)}}{f \in \text{dom}(\Gamma)} \quad \Gamma \vdash (\text{fut}_f) \text{ ok}$	$\frac{\text{(F-FUT)}}{f : \text{Fut } \tau \in \Gamma \quad \Gamma \vdash_\bullet v : \tau} \quad \Gamma \vdash (\text{fut}_f v) \text{ ok}$	$\frac{\text{(TASK)}}{f : \text{Fut } \tau \in \Gamma \quad \Gamma \vdash_\tau e : \tau} \quad \Gamma \vdash (\text{task}_f e) \text{ ok}$
$\frac{\text{(CHAIN)}}{f : \text{Fut } \tau \in \Gamma \quad g : \text{Fut } \tau' \in \Gamma \quad \Gamma \vdash_\tau e : \tau' \rightarrow \tau} \quad \Gamma \vdash (\text{chain}_f g e) \text{ ok}$	$\frac{\text{(CONFIG)} \quad \Gamma \vdash \text{config}_1 \text{ ok} \quad \Gamma \vdash \text{config}_2 \text{ ok} \quad \text{defs}(\text{config}_1) \cap \text{defs}(\text{config}_2) = \emptyset \quad \text{writers}(\text{config}_1) \cap \text{writers}(\text{config}_2) = \emptyset}{\Gamma \vdash \text{config}_1 \text{ config}_2 \text{ ok}}$	

Fig. 4: Configuration typing

Definition 1. The function $\text{defs}(\text{config})$ extracts the set of futures present in a configuration config .

$$\begin{aligned} \text{defs}((\text{fut}_f)) &= \text{defs}((\text{fut}_f v)) = \{f\} \\ \text{defs}((\text{config}_1 \text{ config}_2)) &= \text{defs}(\text{config}_1) \cup \text{defs}(\text{config}_2) \\ \text{defs}(_) &= \emptyset \end{aligned}$$

Definition 2. The function $\text{writers}(\text{config})$ extracts the set of writers to futures in configuration config .

$$\begin{aligned} \text{writers}((\text{chain}_f g e)) &= \text{writers}((\text{task}_f e)) = \{f\} \\ \text{writers}(\text{config}_1 \text{ config}_2) &= \text{writers}(\text{config}_1) \cup \text{writers}(\text{config}_2) \\ \text{writers}(_) &= \emptyset \end{aligned}$$

Rules FUT and F-FUT define well-formed future configurations. Rules TASK and CHAIN define well-formed task and future chaining configurations and set the expected task types. Rule CONFIG defines how to build larger configurations from smaller ones. Each future may be defined at most once and there is at most one writer to each future.

The rules for well-formed configurations apply to partial configurations. Complete configurations can be typed by adding extra conditions to ensure that all futures in Γ have a future configuration, there is a one-to-one correspondence between tasks/chains and unfulfilled futures, and dependencies between tasks are acyclic. These definitions have been omitted and are similar to those found in our earlier work [9].

Formal Properties The proof of soundness of the type system follows standard techniques [8]. The proof of progress requires that there is no deadlock, which follows as there is no cyclic dependency between tasks [9].

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

Proof. By induction on the derivation of $\text{config} \rightarrow \text{config}'$. $\square$

Definition 3 (Terminal Configuration). A complete configuration config is terminal iff every element of the configuration has the shape: $(\text{fut}_f v)$.

Lemma 2 (Progress). For a complete configuration config , 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 the acyclicity of task dependencies. $\square$

3 A Promising Implementation Calculus

The implementation of *forward* in the Encore programming language is via compilation into C, linking with Pony's actor-based run-time [10]. At this level, Encore's futures are treated like promises in that they are passed around to the place where the result of a method call is known in order to be fulfilled. To model this implementation approach, we introduce a low-level target calculus based on tasks and promises. This section presents the formalised target calculus, and the next section presents the compilation strategy from the source to the target language.

The syntax of the target language is as follows:

$$\begin{aligned} e &::= v \mid e \ e \mid \text{Task}(e, e) \mid \text{stop} \mid e; e \mid \text{Prom} \mid \text{fulfil}(e, e) \mid \text{get } e \\ &\quad \mid \text{Chain}(e, e, e) \mid \text{if } e \text{ then } e \text{ else } e \\ v &::= c \mid f \mid x \mid \lambda x. e \mid () \end{aligned}$$

Expressions consist of values, function application ($e \ e$), sequential composition of expressions ($e; e$), the spawning and stopping of tasks ($\text{Task}(e, e)$ and stop), the creation, fulfilment, reading, and chaining of promises (Prom , $\text{fulfil}(e, e)$, $\text{get } e$, and $\text{Chain}(e, e, e)$) and the standard *if-then-else* expression. Values are constants, futures, variables, abstractions and unit $()$. The main differences with the source language are that tasks have to be explicitly stopped, which captures non-local exit, and promises must be explicitly created and fulfilled.

3.1 Operational Semantics

The semantics of the target calculus is analogous to the source calculus. The evaluation contexts are:

$$\begin{aligned} E &::= \bullet \mid E \ e \mid v \ E \mid E; e \mid \text{get } E \mid \text{fulfil}(E, e) \mid \text{fulfil}(v, E) \\ &\quad \mid \text{Task}(E, e) \mid \text{Chain}(e, E, e) \mid \text{Chain}(E, v, e) \mid \text{Chain}(v, v, E) \\ &\quad \mid \text{if } E \text{ then } e \text{ else } e \end{aligned}$$

(RI-IF-TRUE)	(RI-ERROR)
$(\mathbf{task} \ E[\mathbf{if} \ \mathit{true} \ \mathbf{then} \ e \ \mathbf{else} \ e']) \rightarrow (\mathbf{task} \ E[e])$	$(\mathit{prm}_f \ v) (\mathbf{task} \ E[\mathbf{fulfil}(f, v')]) \rightarrow \mathbf{ERROR}$
(RI-IF-FALSE)	(RI-PROMISE)
$(\mathbf{task} \ E[\mathbf{if} \ \mathit{false} \ \mathbf{then} \ e \ \mathbf{else} \ e']) \rightarrow (\mathbf{task} \ E[e'])$	$\frac{\mathit{fresh} \ f}{(\mathbf{task} \ E[\mathbf{Prom}]) \rightarrow (\mathit{prm}_f) (\mathbf{task} \ E[f])}$
(RI-STATEMENT)	(RI-CHAIN)
$(\mathbf{task} \ E[v; e]) \rightarrow (\mathbf{task} \ E[e])$	$(\mathbf{task} \ E[\mathbf{Chain}(f, g, (\lambda x.e))]) \rightarrow (\mathbf{chain} \ g \ e[f/x]) (\mathbf{task} \ E[f])$
(RI- β)	(RI-FULFIL)
$(\mathbf{task} \ E[(\lambda x.e) \ v]) \rightarrow (\mathbf{task} \ E[e[v/x]])$	$(\mathit{prm}_f) (\mathbf{task} \ E[\mathbf{fulfil}(f, v)]) \rightarrow (\mathit{prm}_f \ v) (\mathbf{task} \ E[()])$
(RI-STOP)	(RI-TASK)
$(\mathbf{task} \ E[\mathbf{stop}]) \rightarrow \epsilon$	$(\mathbf{task} \ E[\mathbf{Task}(f, (\lambda x.e))]) \rightarrow (\mathbf{task} \ E[f]) (\mathbf{task} \ e[f/x])$
(RI-CONFIG-CHAIN)	(RI-GET)
$(\mathbf{chain} \ g \ e) (\mathit{prm}_g \ v) \rightarrow (\mathbf{task} \ (e \ v)) (\mathit{prm}_g \ v)$	$(\mathbf{task} \ E[\mathbf{get} \ h]) (\mathit{prm}_h \ v) \rightarrow (\mathbf{task} \ E[v]) (\mathit{prm}_h \ v)$

Fig. 5: Target reduction rules

Configurations are multisets of promises, tasks, and chains:

$$\mathit{config} ::= \epsilon \mid (\mathit{prm}_f) \mid (\mathit{prm}_f \ v) \mid (\mathbf{task} \ e) \mid (\mathbf{chain} \ f \ e) \mid \mathit{config} \ \mathit{config}$$

The empty configuration is represented by ϵ , an unfulfilled promise is written as (prm_f) and a fulfilled promise holding value v is written as $(\mathit{prm}_f \ v)$.

Tasks and chains work in the same way as in the source language, except that they work now on promises (Fig. 5). Promises are handled much more explicitly than futures are, and need to be passed around like regular values. The creation of a task needs a promise and a function to run; the spawned task runs the function, has access to the passed promise and leaves the promise reference in the spawning task (RI-TASK). Stopping a task just finishes the task (RI-STOP). The construct **Prom** creates an empty promise (RI-PROMISE). Fulfilling a promise results in the value being stored if the promise was empty (RI-FULFIL), or an error otherwise (RI-ERROR). Promises are chained in a similar fashion to futures: the construct **Chain**(f, g, e) immediately passes the promise f to expression e — the intention being that f will hold the eventual result; the chain then waits on promise g , and passes the value it receives into expression $(e \ f)$ (RI-CHAIN and RI-CONFIG-CHAIN). The target language borrows the configuration evaluation rules from the source language (Fig. 2).

Example For illustration purposes we translate the example from the high-level language, $(\mathit{fut}_f) (\mathit{task}_f \ E[\mathbf{forward}(\mathbf{async} \ e)])$ shown in Section 2, and show the reduction steps of the low-level language:

$$\begin{aligned}
& (prm_f) \text{ (task } E[\text{Chain}(f, \text{Task}(\text{Prom}, (\lambda d'. \text{fulfil}(d', e); \text{stop})), \lambda d'. \lambda x. \text{fulfil}(d', x); \text{stop}); \text{stop})] \\
& \longrightarrow (prm_f) (prm_g) \text{ (task } E[\text{Chain}(f, \text{Task}(g, (\lambda d'. \text{fulfil}(d', e); \text{stop})), \lambda d'. \lambda x. \text{fulfil}(d', x); \text{stop}); \text{stop})] \\
& \longrightarrow (prm_f) (prm_g) \text{ (task } E[\text{Chain}(f, g, \lambda d'. \lambda x. \text{fulfil}(d', x); \text{stop}); \text{stop})] \text{ (task fulfil}(g, e); \text{stop}) \\
& \longrightarrow (prm_f) (prm_g) \text{ (task } E[f; \text{stop}] \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task fulfil}(g, e); \text{stop}) \\
& \longrightarrow (prm_f) (prm_g) \text{ (task } E[\text{stop}] \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task fulfil}(g, e); \text{stop}) \\
& \longrightarrow (prm_f) (prm_g) \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task fulfil}(g, e); \text{stop}) \\
& \longrightarrow^* (prm_f) (prm_g) \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task fulfil}(g, v); \text{stop}) \\
& \longrightarrow (prm_f) (prm_g v) \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task (); \text{stop}) \\
& \longrightarrow (prm_f) (prm_g v) \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) (task stop) \\
& \longrightarrow (prm_f) (prm_g v) \text{ (chain } g \text{ (}\lambda x. \text{fulfil}(f, x); \text{stop})) \\
& \longrightarrow (prm_f) (prm_g v) \text{ (task (}\lambda x. \text{fulfil}(f, x); \text{stop}) v) \\
& \longrightarrow (prm_f) (prm_g v) \text{ (task fulfil}(f, v); \text{stop}) \\
& \longrightarrow (prm_f v) (prm_g v) \text{ (task (); \text{stop})} \\
& \longrightarrow (prm_f v) (prm_g v) \text{ (task stop)} \\
& \longrightarrow (prm_f v) (prm_g v)
\end{aligned}$$

We show how the compilation strategy proceeds in Section 4.

3.2 Static Semantics

The type system has basic types, K , and promise types defined below:

$$\tau ::= K \mid \text{Prom } \tau$$

The type rules define the judgment $\Gamma \vdash e : \tau$ which states that, in the environment Γ , which records the types of promises and free variables, expression e has type τ . The rules for constants, promises, and variables, *if-then-else*, abstraction and function application are analogous to the source calculus, except no expected task type is recorded. The unit value has type **unit** (TI-UNIT); the **stop** expression finishes a task and has any type (TI-STOP). The creation of a promise has type $\text{Prom } \tau$ (TI-PROMISE-NEW); the fulfilment of a promise $\text{fulfil}(e, e')$ has type **unit** and requires the first parameter to be a promise and the second to be an expression that matches the type of the promise (TI-FULFIL). To spawn a task (**Task**(e, e)), the first argument of the task must be a promise and the second a function that takes a promise having the same type as the first argument (TI-TASK); promises can be chained on with functions that run if the promise is fulfilled: **Chain**(e, e', e'') has type $\text{Prom } \tau$ and e and e' are promises and e'' is an abstraction that takes arguments of the first and second promise types. Both task and chain constructors return the promise that is passed to them, for convenience in the compilation scheme.

Soundness of the type system is proven using standard techniques.

(TI-CONSTANT) $\frac{c \text{ is a constant of type } \tau}{\Gamma \vdash c : \tau}$	(TI-PROMISE) $\frac{f : \text{Prom } \tau \in \Gamma}{\Gamma \vdash f : \text{Prom } \tau}$	(TI-VARIABLE) $\frac{x : \tau \in \Gamma}{\Gamma \vdash x : \tau}$	(TI-UNIT) $\frac{}{\Gamma \vdash () : \mathbf{unit}}$
(TI-STOP) $\frac{}{\Gamma \vdash \mathbf{stop} : \tau}$	(TI-PROMISE-NEW) $\frac{}{\Gamma \vdash \mathbf{Prom} : \text{Prom } \tau}$	(TI-IF) $\frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash e' : \tau \quad \Gamma \vdash e'' : \tau}{\Gamma \vdash \mathbf{if } e \mathbf{ then } e' \mathbf{ else } e'' : \tau}$	
(TI-STATEMENT) $\frac{\Gamma \vdash e_1 : \tau' \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1; e_2 : \tau}$	(TI-ABSTRACTION) $\frac{\Gamma, x : \tau \vdash e : \tau'}{\Gamma \vdash \lambda x. e : \tau \rightarrow \tau'}$	(TI-APP) $\frac{\Gamma \vdash e : \tau' \rightarrow \tau \quad \Gamma \vdash e' : \tau'}{\Gamma \vdash e e' : \tau}$	
(TI-FULFIL) $\frac{\Gamma \vdash e : \text{Prom } \tau \quad \Gamma \vdash e' : \tau}{\Gamma \vdash \mathbf{fulfil}(e, e') : \mathbf{unit}}$		(TI-TASK) $\frac{\Gamma \vdash e : \text{Prom } \tau \quad \Gamma \vdash e' : \text{Prom } \tau \rightarrow \tau'}{\Gamma \vdash \mathbf{Task}(e, e') : \text{Prom } \tau}$	
(TI-GET) $\frac{\Gamma \vdash e : \text{Prom } \tau}{\Gamma \vdash \mathbf{get } e : \tau}$	(TI-CHAIN) $\frac{\Gamma \vdash e : \text{Prom } \tau \quad \Gamma \vdash e' : \text{Prom } \tau' \quad \Gamma \vdash e'' : \text{Prom } \tau \rightarrow \tau' \rightarrow \tau''}{\Gamma \vdash \mathbf{Chain}(e, e', e'') : \text{Prom } \tau}$		

(PROM) $\frac{f \in \text{dom}(\Gamma)}{\Gamma \vdash (\text{prm}_f) \text{ ok}}$	(F-PROM) $\frac{f : \text{Prom } \tau \in \Gamma}{\Gamma \vdash (\text{prm}_f \ v) \text{ ok}}$	(CHAIN-TARGET) $\frac{\Gamma \vdash f : \text{Prom } \tau \quad \Gamma \vdash e : \tau \rightarrow \tau''}{\Gamma \vdash (\mathbf{chain } f \ e) \text{ ok}}$
(TASK-TARGET) $\frac{\Gamma \vdash e : \tau}{\Gamma \vdash (\mathbf{task } e) \text{ ok}}$		(CONFIG-TARGET) $\frac{\Gamma \vdash \text{config}_1 \text{ ok} \quad \Gamma \vdash \text{config}_2 \text{ ok}}{\Gamma \vdash \text{config}_1 \text{ config}_2 \text{ ok}}$

Fig. 6: Typing rules for expressions and configurations in the target language

4 Compilation: From Futures and Forward to Promises

This section presents the compilation function from the source to the target language and outlines a proof that it preserves semantics. The compilation strategy is defined inductively (Fig. 7); the compilation of expressions, denoted $\mathcal{C}[e]_{\text{destiny}}$, takes an expression e and a meta-variable **destiny** which holds the promise that the current task should fulfil, and produces an expression in the target language.

Futures are translated to promises, and most other expressions are translated homomorphically. The constructs where something interesting happens are **async**, **forward** and future chaining; these constructs adopt a common pattern implemented using a two parameter lambda abstraction: the first parameter, variable **destiny'**, is the promise to be fulfilled and the second parameter is the value that fulfils the promise. The best illustration of how **forward** behaves differently from a regular asynchronous call is the difference in the rules

$\mathcal{C}\llbracket e \rrbracket_{\text{destiny}}$	Compilation Strategy
$\begin{aligned} \mathcal{C}\llbracket f \rrbracket_{\text{destiny}} &= f & \mathcal{C}\llbracket x \rrbracket_{\text{destiny}} &= x & \mathcal{C}\llbracket c \rrbracket_{\text{destiny}} &= c \\ \mathcal{C}\llbracket \lambda x. e \rrbracket_{\text{destiny}} &= \lambda x. \mathcal{C}\llbracket e \rrbracket_{\text{destiny}} \\ \mathcal{C}\llbracket e_1 \ e_2 \rrbracket_{\text{destiny}} &= \mathcal{C}\llbracket e_1 \rrbracket_{\text{destiny}} \ \mathcal{C}\llbracket e_2 \rrbracket_{\text{destiny}} \\ \mathcal{C}\llbracket \text{get } e \rrbracket_{\text{destiny}} &= \text{get } \mathcal{C}\llbracket e \rrbracket_{\text{destiny}} \\ \mathcal{C}\llbracket \text{async } e \rrbracket_{\text{destiny}} &= \text{Task}(\text{Prom}, (\lambda \text{destiny}' . \text{fulfil}(\text{destiny}', \mathcal{C}\llbracket e \rrbracket_{\text{destiny}'}); \text{stop})) \\ \mathcal{C}\llbracket \text{forward } e \rrbracket_{\text{destiny}} &= \text{Chain}(\text{destiny}, \mathcal{C}\llbracket e \rrbracket_{\text{destiny}}, \lambda \text{destiny}' . \lambda x. \text{fulfil}(\text{destiny}', x); \text{stop}); \text{stop} \\ \mathcal{C}\llbracket e \xrightarrow{x} e' \rrbracket_{\text{destiny}} &= \text{Chain}(\text{Prom}, \mathcal{C}\llbracket e \rrbracket_{\text{destiny}}, (\lambda \text{destiny}' . \lambda x. \text{fulfil}(\text{destiny}', \mathcal{C}\llbracket e' \rrbracket_{\text{destiny}'}); \text{stop})) \end{aligned}$	
$\mathcal{C}\llbracket e \rrbracket_{\text{destiny}}$	Optimised Compilation Strategy
$\begin{aligned} \mathcal{C}\llbracket \text{forward}(\text{async}(e)) \rrbracket_{\text{destiny}} &= \\ &\quad \text{Task}(\text{destiny}, (\lambda \text{destiny}' . \text{fulfil}(\text{destiny}', \mathcal{C}\llbracket e \rrbracket_{\text{destiny}'}); \text{stop})); \text{stop} \\ \mathcal{C}\llbracket \text{forward } (e \xrightarrow{x} e') \rrbracket_{\text{destiny}} &= \\ &\quad \text{Chain}(\text{destiny}, \mathcal{C}\llbracket e \rrbracket_{\text{destiny}}, \lambda \text{destiny}' . \lambda x. \text{fulfil}(\text{destiny}', \mathcal{C}\llbracket e' \rrbracket_{\text{destiny}'}); \text{stop}); \text{stop} \end{aligned}$	
$\mathcal{T}\llbracket \text{config} \rrbracket$	Configuration Compilation Strategy
$\begin{aligned} \mathcal{T}\llbracket (\text{fut}_f) \rrbracket &= (\text{prm}_f) & \mathcal{T}\llbracket (\text{task}_f \ e) \rrbracket &= (\text{task } \text{fulfil}(f, \mathcal{C}\llbracket e \rrbracket_f); \text{stop}) \\ \mathcal{T}\llbracket (\text{fut}_f \ v) \rrbracket &= (\text{prm}_f \ \mathcal{C}\llbracket v \rrbracket_f) & \mathcal{T}\llbracket \text{config } \text{config}' \rrbracket &= \mathcal{T}\llbracket \text{config} \rrbracket \ \mathcal{T}\llbracket \text{config}' \rrbracket \\ & & \mathcal{T}\llbracket (\text{chain}_f \ g \ e) \rrbracket &= (\text{chain } g \ (\lambda x. \text{fulfil}(f, \mathcal{C}\llbracket e \rrbracket_f \ x); \text{stop})) \\ & & \text{where } x \text{ is fresh} & \\ & & \mathcal{T}\llbracket (\text{task}_f \ e) \rrbracket &= (\text{task } (\text{fulfil}(f, \mathcal{C}\llbracket e \rrbracket_f); \text{stop})) \end{aligned}$	
$\mathcal{C}\llbracket \tau \rrbracket$	Type translation
$\begin{aligned} \mathcal{C}\llbracket \text{ok} \rrbracket &= \text{ok} & \mathcal{C}\llbracket K \rrbracket &= K \\ \mathcal{C}\llbracket \text{Fut } \tau \rrbracket &= \text{Prom } \mathcal{C}\llbracket \tau \rrbracket & \mathcal{C}\llbracket \tau \rightarrow \tau' \rrbracket &= \mathcal{C}\llbracket \tau \rrbracket \rightarrow \mathcal{C}\llbracket \tau' \rrbracket \end{aligned}$	
$\mathcal{T}\llbracket \Gamma \vdash f : \tau \rrbracket$	Environment Translation
$\begin{aligned} \mathcal{T}\llbracket \Gamma \vdash_{\rho} \text{config} \rrbracket &= \mathcal{C}\llbracket \Gamma \rrbracket \vdash \mathcal{T}\llbracket \text{config} \rrbracket & \mathcal{C}\llbracket \emptyset \rrbracket &= \epsilon \\ \mathcal{C}\llbracket \Gamma, f : \text{Fut } \tau \rrbracket &= \mathcal{C}\llbracket \Gamma \rrbracket, \mathcal{C}\llbracket f : \text{Fut } \tau \rrbracket & \mathcal{C}\llbracket x : \tau \rrbracket &= x : \mathcal{C}\llbracket \tau \rrbracket \\ \mathcal{C}\llbracket \Gamma, x : \tau \rrbracket &= \mathcal{C}\llbracket \Gamma \rrbracket, \mathcal{C}\llbracket x : \tau \rrbracket & \mathcal{C}\llbracket f : \text{Fut } \tau \rrbracket &= f : \mathcal{C}\llbracket \text{Fut } \tau \rrbracket \end{aligned}$	

Fig. 7: Compilation strategy of terms, configurations, types and typing rules

for `async e` and the optimised rule for `forward (async e)`. The translation of `async e` creates a new promise to store e 's result value, whereas the translation of `forward (async e)` reuses the promise from the context, namely the one passed in via the `destiny` variable.

The compilation of configurations, denoted $\mathcal{T}[\![config]\!]$, translates configurations from the source language to the target language. For example, the compilation of the source configuration $(fut_f) (task_f \text{ forward } (async e))$ compiles into:

$$\begin{aligned} \mathcal{T}[\![(fut_f) (task_f \text{ forward } (async e))]\!] &= \\ \mathcal{T}[\![(fut_f)]\!] \mathcal{T}[\![(task_f \text{ forward } (async e))]\!] &= \\ (prm_f) (\text{task fulfil}(f, \mathcal{C}[\![\text{forward } (async e)]\!]_f)) & \end{aligned}$$

The optimised compilation of $\mathcal{C}[\![\text{forward } (async e)]\!]_f$ is:

$$(prm_f) (\text{task } E[\text{Task}(f, (\lambda d'. \text{fulfil}(d', \mathcal{C}[\![e]\!]_{d'}); \text{stop})); \text{stop}])$$

For comparison, the base compilation gives:

$$(prm_f) (\text{task } E[\text{Chain}(f, \text{Task}(\text{Prom}, (\lambda d'. \text{fulfil}(d', \mathcal{C}[\![e]\!]_{d'}); \text{stop})), \lambda d'. \lambda x. \text{fulfil}(d', x); \text{stop}); \text{stop}])$$

Types and typing rules are compiled inductively (Fig. 7). The following lemmas guarantee that the compilation strategy does not produce broken target code and state the correctness of the translation.

4.1 Correctness

The correctness of the translation is proven in a number of steps.

The first step involves converting the reduction rules to a labelled transition system where communication via futures is made explicit. This involves splitting several rules involving multiple primitive configurations on the left-hand side to involve single configurations, and labelling the values going into and out of futures. For example, $(task_f v) (fut_f) \rightarrow (fut_f v)$ is replaced by the two rules:

$$(task_f v) \xrightarrow{\overline{f \downarrow v}} \epsilon \quad (fut_f) \xrightarrow{f \downarrow v} (fut_f v)$$

The other rules introduced are:

$$\begin{aligned} (fut_f v) \xrightarrow{f \uparrow v} (fut_f v) \quad (task_f E[\text{get } h]) \xrightarrow{\overline{h \uparrow v}} (task_f E[v]) \\ (chain_g f e) \xrightarrow{\overline{f \uparrow v}} (task_g e[v/x]) \end{aligned}$$

Label $f \downarrow v$ captures a value being written to a future, and label $f \uparrow v$ captures a value being read from a future, both from the future's perspective. Labels $\overline{f \downarrow v}$ and $\overline{f \uparrow v}$ are the duals from the perspective of the remainder of the configuration. The remainder of the rules are labelled with τ to indicate that no observable behaviour occurs. The same pattern is applied to the target language.

It is important to note that the values in the labels of the source language are the compiled values, while the values in the labels of the target language remain the same.⁶ This is needed so that labelled values such as lambda abstraction match during the bisimulation game.

The composition rules are adapted to propagate or match labels in the standard way. For instance, the rule for matching labels in parallel configurations is:

$$\frac{\text{config} \xrightarrow{l} \text{config}'' \quad \text{config}' \xrightarrow{\bar{l}} \text{config}'''}{\text{config config}' \xrightarrow{\tau} \text{config}'' \text{config}'''}$$

The following theorems capture correctness of the translation.

Theorem 1. *If $\Gamma \vdash \text{config ok}$, then $\mathcal{C}[\Gamma] \vdash \mathcal{T}[\llbracket \text{config} \rrbracket] \text{ok}$.*

Theorem 2. *If $\Gamma \vdash \text{config ok}$, then $\text{config} \sim \mathcal{T}[\llbracket \text{config} \rrbracket]$.*

The first theorem states that translating well-typed configurations results in well-typed configurations. The second theorem states that any well-typed configuration in the source language is bisimilar to its translation. The precise notion of bisimilarity used is bisimilarity up-to expansion [11]. This notion of bisimilarity compresses the administrative, unobservable transitions introduced by the translation.

The proof involves taking each derivation rule in the adapted semantics for the source calculus (described above) and showing that each source configuration is bisimilar to its translation. This is straightforward for the base cases, because tasks are deterministic in both source and target languages, and at most two unobservable transitions are introduced by the translation. To handle the parallel composition of configurations, bisimulation is shown to be compositional, meaning that if $\text{config} \sim \mathcal{T}[\llbracket \text{config} \rrbracket]$ and $\text{config}' \sim \mathcal{T}[\llbracket \text{config}' \rrbracket]$, then $\text{config config}' \sim \mathcal{T}[\llbracket \text{config config}' \rrbracket]$; now by definition $\mathcal{T}[\llbracket \text{config config}' \rrbracket] = \mathcal{T}[\llbracket \text{config} \rrbracket] \mathcal{T}[\llbracket \text{config}' \rrbracket]$, hence $\text{config config}' \sim \mathcal{T}[\llbracket \text{config} \rrbracket] \mathcal{T}[\llbracket \text{config}' \rrbracket]$.

5 Experiments

We benchmarked the implementation of **forward** by comparing it against the blocking pattern **get-and-return** and an implementation that uses the **await-and-get** (both described in Section 1). The micro-benchmark used is a variant of the broker pattern with 4 workers, compiled with aggressive optimisations (**-O3**). We report the average time (wall clock) and memory consumption of 5 runs of this micro-benchmark under different workloads (Fig. 8). The processing of each message sent involves complex nested loops with quadratic complexity (in the Workload value) written in such a way to avoid the compiler optimising them away — the higher the workload, the higher the probability that the **Broker** actor blocks or awaits in the non-**forward** implementations.

⁶ We have omitted the notation from the translation to keep it simple to read

Performance (in seconds)			
Workload	Get	Await+Get	Forward
100	0.03	0.03	0.00
500	0.47	0.25	0.02
1000	1.85	0.94	0.06
3000	16.55	8.29	0.39
5000	45.77	23.01	1.03
7500	103.43	51.62	2.26
10000	183.04	91.86	4.02

Memory consumption (in kilobytes)			
Workload	Get	Await+Get	Forward
100	12697	49446	7334
500	12292	49676	6608
1000	12451	49927	6832
3000	12222	49070	7793
5000	12427	48584	7269
7500	12337	48016	7853
10000	12484	48316	8475

Fig. 8: Elapsed time (left) and memory consumed (right) by the Broker microbenchmark (the lower the better).

The performance results (Fig. 8) show that the **forward** version is always faster than the **get-and-return** and **await-and-get** version. In the first case, this is expected as blocking prevents the **Broker** actor from processing messages, while the **forward** version does not block. In the second case, we also expected the **forward** version to be faster than the **await-and-get**: this is due to the overhead of the context switching operation performed on each **await** statement.

The **forward** version consumes the least amount of memory, while the **await-and-get** version consumes the most (Fig. 8). This is expected: **forward** creates one fewer future per message sent than the other two versions; the **await-and-get** version has around 5 times more overhead than the **forward** implementation, as it needs to save the context (stack) whenever a future cannot immediately be fulfilled.

Threats to validity The experiments use a microbenchmark, which provides useful information but is not as comprehensive as a case study would be.

6 Related work

Baker discovered futures in 1977 [12]; later Liskov introduced promises to Argus [4]. Around the same time, Halstead introduced implicit futures in Multisp [13]. Implicit futures do not appear as a first-class construct in the programming language at either the term or type level, as they do in our work.

The **forward** construct was introduced in earlier work [7], in the formalisation of an extension to the active object-based language Creol [14]. The main differences with our work are: our core calculus is much smaller, based on tasks rather than active objects; our calculus includes closures, which complicate the type system, and future chaining; we defined a compilation strategy for **forward**, and benchmark its implementation.

Caromel *et al.* [15] formalise an active object language that transparently handles futures, prove determinism of the language using concepts similar to weak bisimulation, and provide an implementation [16]. In contrast, our work uses a task-based formalism built on top of the lambda calculus and uses fu-

tures explicitly. It is not clear whether **forward** can be used in conjunction with transparent futures.

Proving semantics preservation of whole programs is not a new idea [17–22]. We highlight the work from Lochbihler, who added a new phase to the verified, machine-checked Jinja compiler [23] that proves that the translation from multi-threaded Java programs to Java bytecode is semantics preserving, using a delay bisimulation. In contrast, our work uses an on-paper proof using weak bisimilarity up-to expansion, proving that the compilation strategy preserves the semantics of the high-level language.

Ábrahám et al [5] present an extension of the Creol language with promises. The type system uses linear types to track the use of the write capability (fulfilment) of promises to ensure that they are fulfilled precisely once. In contrast to the present work, their type system is significantly more complex, and no **forward** operation is present. Curiously, Encore supports linear types, though lacks promises and hence does not use linear types to keep promises under control.

Niehren et al [6] present a lambda calculus extended with futures (which are really promises). Their calculus explores the expressiveness of programming with promises, by using them to express channels, semaphores, and ports. They also present a linear type system that ensures that promises are assigned only once.

7 Conclusion

One key difference between futures, futures with forward and promises is that the responsibility to fulfil a future cannot be delegated. The **forward** construct allows such delegation, although only of the implicit future receiving the result of some method call, while promises allow arbitrary delegation of responsibility. This paper presented a formal calculus capturing the **forward** construct, which retains the static fulfilment guarantees of futures. A translation of the source calculus into a target calculus based on promises was provided and proven to be semantics preserving. This translation models how **forward** is implemented in the Encore compiler. Microbenchmarks demonstrated that **forward** improves performance in terms of speed and memory overhead compared to two alternative implementations in the Encore language.

References

1. 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. A survey of active object languages. *ACM Comput. Surv.*, 50(5):76:1–76:39, October 2017.
2. Stephan Brandauer, Elias Castegren, Dave Clarke, Kiko Fernandez-Reyes, Einar Broch Johnsen, Ka I Pun, Silvia Lizeth Tapia Tarifa, Tobias Wrigstad, and Albert Mingkun Yang. Parallel objects for multicores: A glimpse at the parallel language Encore. In Marco Bernardo and Einar Broch Johnsen, editors, *Formal*

Methods for Multicore Programming - 15th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2015, Bertinoro, Italy, June 15-19, 2015, Advanced Lectures, volume 9104 of *Lecture Notes in Computer Science*, pages 1–56. Springer, 2015.

3. Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In Bernhard K. Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *Formal Methods for Components and Objects - 9th International Symposium, FMCO 2010, Graz, Austria, November 29 - December 1, 2010. Revised Papers*, volume 6957 of *Lecture Notes in Computer Science*, pages 142–164. Springer, 2010.
4. Barbara Liskov and Liuba Shrira. Promises: Linguistic support for efficient asynchronous procedure calls in distributed systems. In Richard L. Wexelblat, editor, *Proceedings of the ACM SIGPLAN’88 Conference on Programming Language Design and Implementation (PLDI), Atlanta, Georgia, USA, June 22-24, 1988*, pages 260–267. ACM, 1988.
5. Erika Ábrahám, Immo Grabe, Andreas Grüner, and Martin Steffen. Behavioral interface description of an object-oriented language with futures and promises. *J. Log. Algebr. Program.*, 78(7):491–518, 2009.
6. Joachim Niehren, Jan Schwinghammer, and Gert Smolka. A concurrent lambda calculus with futures. *Theor. Comput. Sci.*, 364(3):338–356, 2006.
7. Dave Clarke, Einar Broch Johnsen, and Olaf Owe. Concurrent objects à la carte. In Dennis Dams, Ulrich Hannemann, and Martin Steffen, editors, *Concurrency, Compositionality, and Correctness, Essays in Honor of Willem-Paul de Roever*, volume 5930 of *Lecture Notes in Computer Science*, pages 185–206. Springer, 2010.
8. Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Inf. Comput.*, 115(1):38–94, 1994.
9. Kiko Fernandez-Reyes, Dave Clarke, and Daniel S. McCain. ParT: An asynchronous parallel abstraction for speculative pipeline computations. In Alberto Lluch-Lafuente and José Proença, editors, *Coordination Models and Languages - 18th IFIP WG 6.1 International Conference, COORDINATION 2016, Held as Part of the 11th International Federated Conference on Distributed Computing Techniques, DisCoTec 2016, Heraklion, Crete, Greece, June 6-9, 2016, Proceedings*, volume 9686 of *Lecture Notes in Computer Science*, pages 101–120. Springer, 2016.
10. Sylvan Clebsch and Sophia Drossopoulou. Fully concurrent garbage collection of actors on many-core machines. In Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes, editors, *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, pages 553–570. ACM, 2013.
11. Damien Pous and Davide Sangiorgi. Enhancements of the bisimulation proof method. In Davide Sangiorgi and Jan Rutten, editors, *Advanced Topics in Bisimulation and Coinduction*. Cambridge University Press, 2012.
12. Henry G. Baker and Carl Hewitt. The incremental garbage collection of processes. *SIGART Newsletter*, 64:55–59, 1977.
13. Robert H. Halstead Jr. Multilisp: A language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.*, 7(4):501–538, 1985.
14. Einar Broch Johnsen, Olaf Owe, and Ingrid Chieh Yu. Creol: A type-safe object-oriented model for distributed concurrent systems. *Theor. Comput. Sci.*, 365(1-2):23–66, 2006.

15. Denis Caromel, Ludovic Henrio, and Bernard P. Serpette. Asynchronous sequential processes. *Inf. Comput.*, 207(4):459–495, 2009.
16. Denis Caromel, Christian Delbe, Alexandre Di Costanzo, and Mario Leyton. ProActive: an integrated platform for programming and running applications on grids and P2P systems. *Computational Methods in Science and Technology*, 12:issue 1, 2006.
17. Xavier Leroy. A formally verified compiler back-end. *J. Autom. Reasoning*, 43(4):363–446, 2009.
18. Xavier Leroy. Formal certification of a compiler back-end or: programming a compiler with a proof assistant. In J. Gregory Morrisett and Simon L. Peyton Jones, editors, *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2006, Charleston, South Carolina, USA, January 11-13, 2006*, pages 42–54. ACM, 2006.
19. Andreas Lochbihler. Verifying a compiler for Java threads. In Andrew D. Gordon, editor, *Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings*, volume 6012 of *Lecture Notes in Computer Science*, pages 427–447. Springer, 2010.
20. Adam Chlipala. A certified type-preserving compiler from lambda calculus to assembly language. In Jeanne Ferrante and Kathryn S. McKinley, editors, *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*, pages 54–65. ACM, 2007.
21. Mitchell Wand. Compiler correctness for parallel languages. In John Williams, editor, *Proceedings of the seventh international conference on Functional programming languages and computer architecture, FPCA 1995, La Jolla, California, USA, June 25-28, 1995*, pages 120–134. ACM, 1995.
22. Xinxin Liu and David Walker. Confluence of processes and systems of objects. In Peter D. Mosses, Mogens Nielsen, and Michael I. Schwartzbach, editors, *TAPSOFT’95: Theory and Practice of Software Development, 6th International Joint Conference CAAP/FASE, Aarhus, Denmark, May 22-26, 1995, Proceedings*, volume 915 of *Lecture Notes in Computer Science*, pages 217–231. Springer, 1995.
23. Gerwin Klein and Tobias Nipkow. A machine-checked model for a Java-like language, virtual machine, and compiler. *ACM Trans. Program. Lang. Syst.*, 28(4):619–695, 2006.