

Deterministic Parallelism

Vikraman Choudhury

April, 2018

1 Deterministic Parallelism

LVish as described in Kuper et al. [2014] is a deterministic parallel programming library which provides the advantages of parallel speedup while avoiding the nondeterministic bugs that plague concurrent code. This is made possible by the use of *LVars* which provides a shared memory construct with well-defined semantics. However, the implementation of *LVish* relies on dynamic enforcement of such guarantees, which makes it difficult to formalize the static and dynamic semantics of the language.

We formalize a toy language with a single *IVar* (instead of *LVars*) which is deterministic by construction. The language is deeply embedded in Agda, and admits an interpretation to the host language. The formalization follows the style used in the *Dependently Typed Metaprogramming* notes by McBride [2013].

1.1 Syntax

The language is a simply typed lambda calculus, with a base type, products and exponentials.

```
data ★ : Set where
  ! : ★
  _⊗_ : ★ → ★ → ★
  _⇒_ : ★ → ★ → ★
```

Contexts are defined as snoc-lists.

```
data Cx (A : Set) : Set where
  ε : Cx A
  _┌_ : Cx A → A → Cx A
```

We setup a De Bruijn index over contexts to track bound identifiers.

```
data _∈_ (τ : ★) : Cx ★ → Set where
  here : ∀ {Γ} → τ ∈ Γ , τ
  there : ∀ {Γ σ} → τ ∈ Γ → τ ∈ Γ , σ
```

A store, parameterized over a type, is either empty or full, containing something of that type.

```
data St (A : Set) : Set where
  ⟨-⟩ : St A
  ⟨_⟩ : (a : A) → St A
```

Next, we setup a typing judgment of the following shape $\Sigma \triangleleft \Gamma \vdash \tau \triangleright \Xi$, where Σ and Ξ are stores, Γ is the context, and τ the type. The intended meaning is that, the derivation of τ in the context Γ changes the store state from Σ to Ξ . The typing relation is defined as a logical relation indexed on the type alongwith the pre and post stores. Due to a technical limitation of the \LaTeX rendering mechanism in Agda, we have to define the term `Tm` but it is an alias for the same operator.

```
data Tm (Γ : Cx ★) : St ★ → ★ → St ★ → Set where
```

The typing rules are defined in a syntax-directed manner. If a type is bound in the context, the derivation is immediate, with no change to the store.

```
var : ∀ {τ Σ}
      → τ ∈ Γ
      -----
      → Σ △left Γ ⊢ τ ▷ Σ
```

A bound type in the extended context is abstracted by a function type, and a function type can be applied, with no change to the store.

```
lam : ∀ {τ σ Σ}
       → Σ △left Γ , σ ⊢ τ ▷ Σ
       -----
       → Σ △left Γ ⊢ σ ⇒ τ ▷ Σ

app : ∀ {τ σ Σ}
       → Σ △left Γ ⊢ σ ⇒ τ ▷ Σ
```

$$\begin{array}{l}
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \triangleright \Sigma \\
\hline
\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Sigma
\end{array}$$

Products are introduced and eliminated as usual, and do not change the store.

$$\begin{array}{l}
\text{prd} : \forall \{\tau \sigma \Sigma\} \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \triangleright \Sigma \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Sigma \\
\hline
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \otimes \tau \triangleright \Sigma
\end{array}$$

$$\begin{array}{l}
\text{fst} : \forall \{\tau \sigma \Sigma\} \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \otimes \tau \triangleright \Sigma \\
\hline
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \triangleright \Sigma
\end{array}$$

$$\begin{array}{l}
\text{snd} : \forall \{\tau \sigma \Sigma\} \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \otimes \tau \triangleright \Sigma \\
\hline
\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Sigma
\end{array}$$

A get operation can only be done on a full store. This is enforced by a full store for the start state. The second argument can access the read value in the extended context.

$$\begin{array}{l}
\text{get} : \forall \{\tau \sigma \Sigma\} \\
\rightarrow \langle \sigma \rangle \triangleleft \Gamma, \sigma \vdash \tau \triangleright \Sigma \\
\hline
\rightarrow \langle \sigma \rangle \triangleleft \Gamma \vdash \tau \triangleright \Sigma
\end{array}$$

A put operation can only be done on an empty store. This is enforced by an empty store for the start state. The first argument of put updates the state, and the updated state is available to the second argument.

$$\begin{array}{l}
\text{put} : \forall \{\tau \sigma \Sigma\} \\
\rightarrow \langle - \rangle \triangleleft \Gamma \vdash \sigma \triangleright \langle - \rangle \\
\rightarrow \langle \sigma \rangle \triangleleft \Gamma \vdash \tau \triangleright \Sigma \\
\hline
\rightarrow \langle - \rangle \triangleleft \Gamma \vdash \tau \triangleright \Sigma
\end{array}$$

For a fork operation, two threads can run concurrently only if they make compatible updates to the store. There are two symmetric cases, one thread does not update the store, while the other thread does, so these threads can be successfully joined. This is expressed by a single rule, and the product of the result of the two computations is returned.

$$\begin{array}{l}
\text{fork} : \forall \{ \tau \ \sigma \ \Sigma \ \Xi \} \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \triangleright \langle - \rangle \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Xi \\
\hline
\rightarrow \Sigma \triangleleft \Gamma \vdash \sigma \otimes \tau \triangleright \Xi
\end{array}$$

Although not strictly necessary for the interpreter, we can implement substitution for this calculus for sanity-checking. First we define the friendly fish operator for context extension with a list.

$$\begin{array}{l}
_ \langle \rangle _ : \forall \{ A \} \rightarrow \text{Cx } A \rightarrow \text{List } A \rightarrow \text{Cx } A \\
\Gamma \langle \rangle \langle [] \rangle = \Gamma \\
\Gamma \langle \rangle \langle (a :: as) \rangle = (\Gamma , a) \langle \rangle \langle as \rangle
\end{array}$$

Substitutions are morphisms of contexts.

$$\begin{array}{l}
\text{Sub} : \text{Cx } \star \rightarrow \text{Cx } \star \rightarrow \text{Set} \\
\text{Sub } \Gamma \Delta = \forall \{ \tau \ \Sigma \ \Xi \} \rightarrow \tau \in \Gamma \rightarrow \Sigma \triangleleft \Delta \vdash \tau \triangleright \Xi \\
\\
\text{Shub} : \text{Cx } \star \rightarrow \text{Cx } \star \rightarrow \text{Set} \\
\text{Shub } \Gamma \Delta = \forall \Pi \rightarrow \text{Sub } (\Gamma \langle \rangle \langle \Pi \rangle) (\Delta \langle \rangle \langle \Pi \rangle)
\end{array}$$

Substitution on terms is type-preserving and acts functorially.

$$\begin{array}{l}
\text{subst} : \forall \{ \Gamma \ \Delta \ \tau \ \Sigma \ \Xi \} \rightarrow \text{Shub } \Gamma \Delta \\
\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Xi \rightarrow \Sigma \triangleleft \Delta \vdash \tau \triangleright \Xi \\
\text{subst } \theta (\text{var } i) = \theta [] i \\
\text{subst } \theta (\text{lam } t) = \text{lam } (\text{subst } (\lambda \rho \rightarrow \theta (_ :: \rho)) t) \\
\text{subst } \theta (\text{app } f s) = \text{app } (\text{subst } \theta f) (\text{subst } \theta s) \\
\text{subst } \theta (\text{prd } s t) = \text{prd } (\text{subst } \theta s) (\text{subst } \theta t) \\
\text{subst } \theta (\text{fst } t) = \text{fst } (\text{subst } \theta t) \\
\text{subst } \theta (\text{snd } t) = \text{snd } (\text{subst } \theta t) \\
\text{subst } \theta (\text{get } f) = \text{get } (\text{subst } (\lambda \rho \rightarrow \theta (_ :: \rho)) f)
\end{array}$$

$\text{subst } \theta (\text{put } p \ t) = \text{put } (\text{subst } \theta \ p) (\text{subst } \theta \ t)$
 $\text{subst } \theta (\text{fork } s \ t) = \text{fork } (\text{subst } \theta \ s) (\text{subst } \theta \ t)$

We consider a few examples of programs in this calculus that successfully type-check. We begin by fixing types τ and σ .

module `Examples` ($\tau \ \sigma : \star$) where

e_1 reads a value from the store and applies it to the function in the context.

$e_1 : \langle \sigma \rangle \triangleleft (\varepsilon, \sigma \Rightarrow \tau \vdash \tau \triangleright \langle \sigma \rangle$
 $e_1 = \text{get } (\text{app } (\text{var } (\text{there here})) (\text{var here}))$

e_2 computes a value by function application and writes it to the store, and the return value is read back.

$e_2 : \langle - \rangle \triangleleft (\varepsilon, \sigma \Rightarrow \tau), \sigma \vdash \tau \triangleright \langle \tau \rangle$
 $e_2 = \text{put } (\text{app } (\text{var } (\text{there here})) (\text{var here}))$
 $\quad (\text{get } (\text{var here}))$

e_3 runs two threads which compute different values while only one of them updates the store.

$e_3 : \langle - \rangle \triangleleft (\varepsilon, \sigma \Rightarrow \tau), \sigma \vdash \tau \otimes \sigma \triangleright \langle \tau \rangle$
 $e_3 = \text{fork } (\text{app } (\text{var } (\text{there here})) (\text{var here}))$
 $\quad (\text{put } (\text{app } (\text{var } (\text{there here})) (\text{var here}))$
 $\quad \quad (\text{var here}))$

1.2 Semantics

Now, we can define a semantics for each syntactic construct, by interpreting them back to the host language. The base type is interpreted to the type of natural numbers, and the product and exponential are interpreted to the corresponding types in Agda.

$\llbracket _ \rrbracket \star : \star \rightarrow \text{Set}$
 $\llbracket \iota \rrbracket \star = \mathbb{N}$
 $\llbracket \sigma \otimes \tau \rrbracket \star = \llbracket \sigma \rrbracket \star \times \llbracket \tau \rrbracket \star$
 $\llbracket \sigma \Rightarrow \tau \rrbracket \star = \llbracket \sigma \rrbracket \star \rightarrow \llbracket \tau \rrbracket \star$

A context of types can be interpreted back, by abstracting over the interpretation of types.

$$\begin{aligned} \llbracket _ \rrbracket_{\text{Cx}} &: \{A : \text{Set}\} \rightarrow \text{Cx } A \rightarrow (A \rightarrow \text{Set}) \rightarrow \text{Set} \\ \llbracket \varepsilon \rrbracket_{\text{Cx}} \text{ v} &= \top \\ \llbracket \Gamma, \sigma \rrbracket_{\text{Cx}} \text{ v} &= \llbracket \Gamma \rrbracket_{\text{Cx}} \text{ v} \times \text{v } \sigma \end{aligned}$$

De Bruijn indices are interpreted by doing a recursive lookup in the context.

$$\begin{aligned} \llbracket _ \rrbracket_{\in} &: \forall \{\Gamma \tau \text{ v}\} \rightarrow \tau \in \Gamma \rightarrow \llbracket \Gamma \rrbracket_{\text{Cx}} \text{ v} \rightarrow \text{v } \tau \\ \llbracket \text{here} \rrbracket_{\in} (\gamma, t) &= t \\ \llbracket \text{there } i \rrbracket_{\in} (\gamma, s) &= \llbracket i \rrbracket_{\in} \gamma \end{aligned}$$

Stores are interpreted similarly using the interpretation of the type.

$$\begin{aligned} \llbracket _ \rrbracket_{\text{St}} &: \{A : \text{Set}\} \rightarrow \text{St } A \rightarrow (A \rightarrow \text{Set}) \rightarrow \text{Set} \\ \llbracket \langle - \rangle \rrbracket_{\text{St}} \text{ v} &= \top \\ \llbracket \langle \sigma \rangle \rrbracket_{\text{St}} \text{ v} &= \text{v } \sigma \end{aligned}$$

Finally, we give semantics to the typing judgment. We interpret the states using a state monad, where the start state is Σ and the end state is Ξ .

$$\begin{aligned} \llbracket _ \rrbracket_{\vdash} &: \forall \{\Gamma \tau \Sigma \Xi\} \\ &\rightarrow \Sigma \triangleleft \Gamma \vdash \tau \triangleright \Xi \\ &\rightarrow \llbracket \Gamma \rrbracket_{\text{Cx}} \llbracket _ \rrbracket_{\star} \\ &\rightarrow \llbracket \Sigma \rrbracket_{\text{St}} \llbracket _ \rrbracket_{\star} \rightarrow \llbracket \Xi \rrbracket_{\text{St}} \llbracket _ \rrbracket_{\star} \times \llbracket \tau \rrbracket_{\star} \\ \llbracket \text{var } i \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, \llbracket i \rrbracket_{\in} \gamma \\ \llbracket \text{lam } t \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, \lambda s \rightarrow \text{proj}_2 (\llbracket t \rrbracket_{\vdash} (\gamma, s) \text{ st}) \\ \llbracket \text{app } f s \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, \text{proj}_2 (\llbracket f \rrbracket_{\vdash} \gamma \text{ st}) (\text{proj}_2 (\llbracket s \rrbracket_{\vdash} \gamma \text{ st})) \\ \llbracket \text{prd } s t \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, (\text{proj}_2 (\llbracket s \rrbracket_{\vdash} \gamma \text{ st}), \text{proj}_2 (\llbracket t \rrbracket_{\vdash} \gamma \text{ st})) \\ \llbracket \text{fst } t \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, \text{proj}_1 (\text{proj}_2 (\llbracket t \rrbracket_{\vdash} \gamma \text{ st})) \\ \llbracket \text{snd } t \rrbracket_{\vdash} \gamma \text{ st} &= \text{st}, \text{proj}_2 (\text{proj}_2 (\llbracket t \rrbracket_{\vdash} \gamma \text{ st})) \\ \llbracket \text{get } f \rrbracket_{\vdash} \gamma \text{ st} &= \llbracket f \rrbracket_{\vdash} (\gamma, \text{st}) \text{ st} \\ \llbracket \text{put } p t \rrbracket_{\vdash} \gamma \text{ st} &= \llbracket t \rrbracket_{\vdash} \gamma (\text{proj}_2 (\llbracket p \rrbracket_{\vdash} \gamma \text{ st})) \\ \llbracket \text{fork } s t \rrbracket_{\vdash} \gamma \text{ st} &= \text{let } (\text{st}', \text{vt}) = \llbracket t \rrbracket_{\vdash} \gamma \text{ st} \\ &\quad (_, \text{vs}) = \llbracket s \rrbracket_{\vdash} \gamma \text{ st} \\ &\quad \text{in } \text{st}', (\text{vs}, \text{vt}) \end{aligned}$$

Since we have a provably terminating, total function which interprets each term in the calculus, we know that the evaluation is deterministic! We can try to run our examples with the interpreter and get expected results.

```
open Examples ι ι
```

```
e1 ↦ : [[ e1 ]] ⊢ (tt , (λ n → 3 * n)) 19 ≡ (19 , 57)  
e1 ↦ = refl
```

```
e2 ↦ : [[ e2 ]] ⊢ ((tt , (λ n → 3 * n)) , 19) tt ≡ (57 , 57)  
e2 ↦ = refl
```

```
e3 ↦ : [[ e3 ]] ⊢ ((tt , (λ n → 3 * n)) , 19) tt ≡ (57 , 57 , 19)  
e3 ↦ = refl
```

1.3 Conclusion

Using the decidable typechecking algorithm in Agda, we're able to typecheck programs in this language, which implies that we can extract a typechecking algorithm for a real implementation of this language. Alternatively, this could be implemented by deep embedding into Haskell using its recent dependently typed programming features as described in Weirich et al. [2017], or using refinement types in Liquid Haskell from Vazou et al. [2017]. It should also be possible to extend this technique to implement the full LVars calculus.

References

Lindsey Kuper, Aaron Turon, Neelakantan R Krishnaswami, and Ryan R Newton. Freeze after writing: Quasi-deterministic parallel programming with lvars. *ACM SIGPLAN Notices*, 49(1):257–270, 2014.

Conor McBride. Dependently typed metaprogramming (in agda). 2013.

Stephanie Weirich, Antoine Voizard, Pedro Henrique Avezedo de Amorim, and Richard A Eisenberg. A specification for dependent types in haskell. *Proceedings of the ACM on Programming Languages*, 1(ICFP):31, 2017.

Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G. Scott, Ryan R. Newton, Philip Wadler, and Ranjit Jhala. Refinement reflection: Complete verification with smt. *Proc. ACM Program. Lang.*, 2(POPL):53:1–53:31, December 2017. ISSN 2475-1421. doi: 10.1145/3158141. URL <http://doi.acm.org/10.1145/3158141>.