

The Function Passing Model: Types, Proofs, and Semantics

Philipp Haller, Normen Müller, Heather Miller

May 2016

1 Overview

We formalize our programming model in the context of a typed lambda calculus with records. Figure 1 shows the abstract syntax of our core language. Besides standard terms, the language includes terms related to (a) spores, (b) silos, and (c) futures. The `spore` term creates a new spore. It contains a list of variable definitions, the spore header, and a closure which may only refer to its parameter and variables in the spore header. The `spawn` term creates a new host capable of hosting silos. The `populate` term initializes a new silo on a given host with a given data value. The `map`, `flatMap`, and `persist` terms create lineages of silo transformations represented as silo references. The `send` term forces the materialization of the silo corresponding to its argument silo reference; `send` returns a future which is asynchronously completed with the silo’s value. The `await` term waits for the completion of its argument future and returns the future’s value. Locations ι are used to refer to futures and hosts, both of which can be created dynamically using the above terms.

Values in our language are as expected: besides abstractions and record values they include spore values, locations, and silo references. Locations and silo references are not part of the “surface syntax” of our language; they are only introduced by evaluation (see Section 1.1). Silo reference values are values of a simple datatype with constructors *Mat*, *Mapped*, *FMapped*, and *Persist*. The constructors include all information required for *materializing* a silo with the result of applying the described transformations. Therefore, a silo reference value is also called the *lineage* of its corresponding silo. We defer a detailed explanation of the transformations described by a lineage to the following Section 1.1.

In addition to standard function and record types, the language has types for spores, hosts, silo references, and futures. A spore type $T \Rightarrow T' \{ \text{type } C = \bar{T} \}$ includes the types \bar{T} of the variables declared in the header of the spore.

1.1 Operational Semantics

In the following we present a small-step operational semantics of the introduced core language. The semantics is clearly stratified into a deterministic layer and a non-deterministic (concurrent) layer. Importantly, this means our programming model can benefit from existing reasoning techniques for sequential programs. Program transformations that are correct for sequential programs are also cor-

$t ::=$ x $ (x : T) \Rightarrow t$ $ t t$ $ \{\overline{l} = t\}$ $ t.l$ $ \mathbf{spore} \{ \overline{x : T = t} ; (x : T) \Rightarrow t \}$ $ \mathbf{spawn}(t)$ $ \mathbf{populate}(t, t)$ $ \mathbf{map}(t, t)$ $ \mathbf{flatMap}(t, t)$ $ \mathbf{persist}(t)$ $ \mathbf{send}(t)$ $ \mathbf{await}(t)$ $ \iota$ $ r$	<i>terms:</i> variable abstraction application record construction selection spore spawn host populate silo map flatMap persist send await future location silo reference
$v ::=$ $(x : T) \Rightarrow t$ $ \{\overline{l} = v\}$ $ p$ $ \iota$ $ r$	<i>values:</i> abstraction value record value spore value location silo reference
$p ::= \mathbf{spore} \{ \overline{x : T = v} ; (x : T) \Rightarrow t \}$	
$r ::=$ $\mathbf{Mat}(\omega)$ $ \mathbf{Mapped}(\omega, r, p)$ $ \mathbf{FMapped}(\omega, r, p)$ $ \mathbf{Persist}(\omega, r, v)$	<i>silo reference values:</i> materialized lineage with map lineage with flatMap lineage with persist
$\omega ::= (h, i) \quad \text{where } i \in \mathbb{N}$	decentralized identifier
$T ::=$ $T \Rightarrow T$ $ \{\overline{l} : T\}$ $ T \Rightarrow T \{ \text{type } \mathcal{C} = \overline{T} \}$ $ \mathbf{Host}$ $ \mathbf{SiloRef}[T]$ $ \mathbf{Future}[T]$	<i>types:</i> abstraction type record type spore type host type silo reference type future type

Figure 1: Abstract syntax of core language.

rect for distributed programs. Our programming model shares this property with some existing approaches [?].

The semantics is based on three reduction relations for (a) sequential reduction of terms, (b) deterministic reduction of hosts, and (c) non-deterministic reduction of sets of hosts. The reduction relations use the definition of evalua-

$E ::=$		<i>evaluation contexts:</i>
	[]	hole
	$E t$	application (fun)
	$v E$	application (arg)
	$\{\overline{l = v}; l_i = E; \overline{l' = t}\}$	record
	$E.l$	selection
	$\text{spore } \{ \overline{x : T = v}; x_i : T_i = E; \overline{x' : T = t}; (x : T) \Rightarrow t \}$	spore
	$\text{spawn}(E)$	spawn
	$\text{populate}(E, t)$	populate (host)
	$\text{populate}(v, E)$	populate (spore)
	$\text{map}(E, t)$	map (ref)
	$\text{map}(v, E)$	map (fun)
	$\text{flatMap}(E, t)$	flatMap (ref)
	$\text{flatMap}(v, E)$	flatMap (fun)
	$\text{persist}(E)$	persist
	$\text{send}(E)$	send
	$\text{await}(E)$	await

Figure 2: Evaluation context.

tion contexts shown in Figure 2. Evaluation contexts capture the notion of the “next subterm to be evaluated.” Following a standard approach [?], we write $E[t]$ for the term obtained by replacing the hole in evaluation context E with term t .

Figure 3 shows the rules for sequential reduction. The sequential reduction relation has the form $E[t] \mid \mu \rightarrow^h E[t'] \mid \mu'$ with stores μ and μ' . Stores are required for the dynamic allocation of futures and hosts. A store μ is a partial function mapping locations ι to values v . The annotation with host h is used for creating *decentralized identifiers* $\omega = (h, i)$ for silo references. Rules R-APPABS and R-PROJRCB are completely standard. Analogous to rule R-APPABS, rule R-APPSPORE describes the application of a spore value to an argument value. Rule R-AWAIT reduces $\text{await}(\iota)$ to v if future ι is already completed with v in μ .

Rules R-MAP, R-FMAP, R-PERSIST and R-UNPERSIST describe the creation of lineages. Rules R-MAP and R-FMAP create silo reference values using the constructors *Mapped* and *FMapped*, respectively. The new silo reference has a fresh identifier (h, i) which uniquely identifies the corresponding (logical) silo. In each case, the spore value p is stored in the new silo reference; this enables a materialization of the silo identified by (h, i) using parent silo reference r and spore p . Rules R-PERSIST and R-UNPERSIST create silo reference values using the *Persist* constructor. *Persist* contains a function enabling host h to persist $(\cdot \cup \cdot)$ or unpersist $(\cdot \setminus \cdot)$ silo r , respectively.

The deterministic reduction relation has the form $(E[t], \mu, Q, S)^h \longrightarrow (E[t'], \mu', Q', S')^h$ where Q is a *message queue* and S is a *silo store*. Figure 4 shows the definition of message queues. A message queue Q may contain three kinds of messages. A message of the form $\text{Req}_s(h, r, \omega)$ requests the value of silo r to be sent to host h for materialization of identifier ω . A message of the form $\text{Res}_s(\omega, v, P)$ represents the corresponding response, containing the identifier ω to be materialized, value v , and persist set P (the set of hosts which have persisted the

$$\begin{array}{c}
\text{R-APPABS} \\
\frac{}{E[(x : T) \Rightarrow t] v' \mid \mu \rightarrow^h E[[x \mapsto v']t] \mid \mu} \\
\\
\text{R-PROJRCD} \\
\frac{}{E[\{l_i = v_i^{i \in 1..n}\}.l_j] \mid \mu \rightarrow^h E[v_j] \mid \mu} \\
\\
\text{R-APPSPORE} \\
\frac{}{E[(\text{spore } \{ x : T = v ; (x : T) \Rightarrow t \}) v'] \mid \mu \rightarrow^h E[\overline{[x \mapsto v]}[x \mapsto v']t] \mid \mu} \\
\\
\begin{array}{cc}
\text{R-AWAIT} & \text{R-MAP} \\
\frac{\mu(\iota) = \text{Some}(v)}{E[\text{await}(\iota)] \mid \mu \rightarrow^h E[v] \mid \mu} & \frac{r' = \text{Mapped}((h, i), r, p) \quad i \text{ fresh}}{E[\text{map}(r, p)] \mid \mu \rightarrow^h E[r'] \mid \mu'} \\
\\
\text{R-FMAP} & \text{R-PERSIST} \\
\frac{r' = \text{FMapped}((h, i), r, p) \quad i \text{ fresh}}{E[\text{flatMap}(r, p)] \mid \mu \rightarrow^h E[r'] \mid \mu'} & \frac{r' = \text{Persist}((h, i), r, \cdot \cup \cdot) \quad i \text{ fresh}}{E[\text{persist}(r)] \mid \mu \rightarrow^h E[r'] \mid \mu'} \\
\\
\text{R-UNPERSIST} \\
\frac{r' = \text{Persist}((h, i), r, \cdot \setminus \cdot) \quad i \text{ fresh}}{E[\text{unpersist}(r)] \mid \mu \rightarrow^h E[r'] \mid \mu'}
\end{array}
\end{array}$$

Figure 3: Sequential reduction.

$$\begin{array}{ll}
Q ::= & \text{message queue values:} \\
\epsilon & \text{empty queue} \\
| \text{Req}_s(h, r, \omega) :: Q & \text{request (silo)} \\
| \text{Res}_s(\omega, v, P) :: Q & \text{response (silo)} \\
| \text{Req}_\iota(\iota, \omega) :: Q & \text{request (future)}
\end{array}$$

Figure 4: Message queues.

silo identified by ω). A message of the form $\text{Req}_\iota(\iota, \omega)$ requests future ι to be completed with the value of silo ω . A *silo store* S is a partial function mapping identifiers ω to values of the form $(\text{Val}(v), P)$ or $(\text{Fwd}(r), P)$ where P is a set of hosts which have persisted the silo (the persist set). The former represents a materialized silo with value v . The latter represents a *proxy* forwarding to the silo specified by lineage r .

The deterministic reduction rules use helper functions *host*, *id*, *parent*, and *consume*, which are defined as follows:

Definition 1.1 (Host). *The host of a silo reference.*

$$\text{host}(r) := \begin{cases} h & \text{if } r = \text{Mat}((h, i)) \\ \text{host}(r') & \text{if } r = \text{Mapped}(-, r', -) \\ \text{host}(r') & \text{if } r = \text{FMapped}(-, r', -) \\ \text{host}(r') & \text{if } r = \text{Persist}(-, r', -) \end{cases}$$

Definition 1.2 (Silo reference identifier). *The identifier of a silo reference.*

$$\text{id}(r) := \begin{cases} \omega & \text{if } r = \text{Mat}(\omega) \\ \omega & \text{if } r = \text{Mapped}(\omega, r', -) \\ \omega & \text{if } r = \text{FMapped}(\omega, r', -) \\ \omega & \text{if } r = \text{Persist}(\omega, r', -) \end{cases}$$

Definition 1.3 (Silo reference parent). *The parent of a silo reference.*

$$\text{parent}(r) := \begin{cases} \text{None} & \text{if } r = \text{Mat}(-) \\ \text{Some}(r') & \text{if } r = \text{Mapped}(-, r', -) \\ \text{Some}(r') & \text{if } r = \text{FMapped}(-, r', -) \\ \text{Some}(r') & \text{if } r = \text{Persist}(-, r', -) \end{cases}$$

Definition 1.4 (Consume silo). *Consume silo ω with persist set P in silo store S .*

$$\text{consume}(\omega, P, S) := \begin{cases} S - \omega & \text{if } P = \emptyset \\ S & \text{otherwise} \end{cases}$$

We discuss the deterministic reduction rules in two steps. First, we discuss the rules shown in Figure 5. Rule R-SEQ reduces $\text{host}(E[t], \mu, Q, S)^h$ in case $E[t]$ reduces in μ . Rule R-SEND1LOCAL reduces $\text{send}(r)$ to a completed future ι if the corresponding silo is already materialized in silo store S . Rule R-SEND2LOCAL covers the case where the requested silo is not yet materialized. In this case, two request messages are added to the queue: a first message $\text{Req}_s(h, r, \text{id}(r))$ requesting the materialization of silo $\text{id}(r)$, and a second message requesting the value of silo $\text{id}(r)$ for completing future ι . Rule R-REQF1 processes a message $\text{Req}_\iota(\iota, \omega)$ by completing future ι with the value of the materialized silo ω . Rule R-REQF2 delays such a request in case silo ω is not yet materialized by moving the request to the back of the queue.

Figure 6 shows the remaining deterministic reduction rules. Rule R-RES processes a message $\text{Res}_s(\omega, v, P)$ by materializing silo ω with value v , yielding silo store S' . Rules R-REQ1LOCAL and R-REQ2LOCAL process a message $\text{Req}_s(h, r, \omega)$ where silo store S forwards $\text{id}(r)$ to another silo $\text{id}(r')$. Rules R-REQMAPLOCAL and R-REQFMAPLOCAL evaluate a silo reference containing *Mapped* or *FMapped*, respectively, in case the parent silo reference is materialized. In both cases, spore value p , stored in r , is applied to the value of the parent silo. In case of R-REQMAPLOCAL, the silo store is updated with the materialization result v' . In case of R-REQFMAPLOCAL, the silo store

$$\begin{array}{c}
\text{R-SEQ} \\
\frac{E[t] \mid \mu \rightarrow^h E[t'] \mid \mu'}{(E[t], \mu, Q, S)^h \rightarrow (E[t'], \mu', Q, S)^h} \\
\\
\text{R-SEND1LOCAL} \\
\frac{\text{host}(r) = h \quad S(\text{id}(r)) = (\text{Val}(v), P) \quad \iota \text{ fresh} \quad \mu' = [\iota \mapsto \text{Some}(v)]\mu}{(E[\text{send}(r)], \mu, Q, S)^h \rightarrow (E[\iota], \mu', Q, S)^h} \\
\\
\text{R-SEND2LOCAL} \\
\frac{\text{host}(r) = h \quad \text{id}(r) \notin \text{dom}(S) \quad \iota \text{ fresh} \quad \mu' = [\iota \mapsto \text{None}]\mu}{(E[\text{send}(r)], \mu, Q, S)^h \rightarrow (E[\iota], \mu', Q \cdot \text{Req}_s(h, r, \text{id}(r)) \cdot \text{Req}_\iota(\iota, \text{id}(r)), S)^h} \\
\\
\text{R-REQF1} \\
\frac{Q = \text{Req}_\iota(\iota, \omega) :: Q' \quad S(\omega) = (\text{Val}(v), P) \\ S' = \text{consume}(\omega, P, S) \quad \mu' = [\iota \mapsto \text{Some}(v)]\mu}{(E[\text{await}(\iota')], \mu, Q, S)^h \rightarrow (E[\text{await}(\iota')], \mu', Q', S')^h} \\
\\
\text{R-REQF2} \\
\frac{Q = \text{Req}_\iota(\iota, \omega) :: Q' \quad \omega \notin \text{dom}(S)}{(E[\text{await}(\iota')], \mu, Q, S)^h \rightarrow (E[\text{await}(\iota')], \mu, Q' \cdot \text{Req}_\iota(\iota, \omega), S)^h}
\end{array}$$

Figure 5: Deterministic reduction (future).

is updated with a forwarding reference to r'' , the result of the spore application. Finally, the parent silo $\text{id}(r')$ is consumed (removed from silo store S'') in case the persist set P is empty, which means that $\text{id}(r')$ was not persisted. Rule R-REQPERSISTLOCAL materializes silo ω' under a persist set P' which is obtained by modifying the persist set P of parent silo $\text{id}(r')$ according to the operator \star stored in r . Rule R-REQPARENTLOCAL enqueues a materialization request $\text{Req}_s(h, r', \text{id}(r'))$ in case the parent $\text{id}(r')$ of a requested silo $\text{id}(r)$ is not materialized yet.

Figure 7 shows the non-deterministic reduction rules. The non-deterministic reduction relation has the form $H \rightarrow H'$ where H and H' are sets of hosts of the form $(t, \mu, Q, S)^h$. Rule R-SCHEDULE reduces a host chosen non-deterministically from the set of hosts. Rule R-SPAWN creates a new host whose initial term is given by the application of the provided spore to the unit value $\{\}$. The new host has an empty store, an empty queue, and an empty silo store. Rule R-POPULATE materializes a silo with a fresh identifier ω on host h' using value v . Rules R-REQ1-3 and R-SEND are analogous to the corresponding deterministic reduction rules. The main difference is that messages are exchanged between different hosts in the case of non-deterministic reduction.

1.2 Type Assignment

Type assignment is based on a judgment of the form $\Gamma; \Sigma; \Delta \vdash t : T$ which assigns term t type T . Γ is a standard type environment; Σ is a standard store typing; Δ is a *silo store typing* which is new. Δ maps identifiers ω to types,

$$\begin{array}{c}
\text{R-RES} \\
\frac{Q = \text{Res}_s(\omega, v, P) :: Q' \quad S' = [\omega \mapsto (\text{Val}(v), P)]S}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q', S')^h} \\
\\
\text{R-REQ1LOCAL} \\
\frac{Q = \text{Req}_s(h, r, \omega) :: Q' \quad S(\text{id}(r)) = (\text{Fwd}(r'), P) \quad S(\text{id}(r')) = (\text{Val}(v), P')}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Res}_s(\omega, v, P), S')^h} \\
\\
\text{R-REQ2LOCAL} \\
\frac{Q = \text{Req}_s(h, r, \omega) :: Q' \quad S(\text{id}(r)) = (\text{Fwd}(r'), P) \quad \text{id}(r') \notin \text{dom}(S)}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Req}_s(h, r', \omega), S')^h} \\
\\
\text{R-REQMAPLOCAL} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q' \quad r = \text{Mapped}(\omega', r', p) \quad S(\text{id}(r')) = (\text{Val}(v), P) \\ v' = p(v) \quad S' = [\omega' \mapsto (\text{Val}(v'), \emptyset)]S \quad S'' = \text{consume}(\text{id}(r'), P, S')}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Req}_s(h', r, \omega), S'')^h} \\
\\
\text{R-REQFMAPLOCAL} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q' \quad r = \text{FMapped}(\omega', r', p) \quad S(\text{id}(r')) = (\text{Val}(v), P) \\ r'' = p(v) \quad S' = [\omega' \mapsto (\text{Fwd}(r''), \emptyset)]S \quad S'' = \text{consume}(\text{id}(r'), P, S')}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Req}_s(h', r'', \omega), S'')^h} \\
\\
\text{R-REQPERSISTLOCAL} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q' \quad r = \text{Persist}(\omega', r', \star) \quad \omega' = (h'', i) \quad S(\text{id}(r')) = (\text{Val}(v), P) \\ P' = P \star \{h''\} \quad S' = [\omega' \mapsto (\text{Val}(v), P')]S \quad S'' = \text{consume}(\text{id}(r'), P, S')}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Res}_s(\omega, v, P'), S'')^h} \\
\\
\text{R-REQPARENTLOCAL} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q' \quad \text{Some}(r') = \text{parent}(r) \quad \text{id}(r') \notin \text{dom}(S)}{(E[\text{await}(\iota)], \mu, Q, S)^h \longrightarrow (E[\text{await}(\iota)], \mu, Q' \cdot \text{Req}_s(h, r', \text{id}(r')) \cdot \text{Req}_s(h', r, \omega), S')^h}
\end{array}$$

Figure 6: Deterministic reduction (silos).

$$\begin{array}{c}
\text{R-SCHEDULE} \\
\frac{(t, \mu, Q, S)^h \rightarrow (t', \mu', Q', S')^h}{\{(t, \mu, Q, S)^h\} \cup H \rightarrow \{(t', \mu', Q', S')^h\} \cup H} \\
\\
\text{R-SPAWN} \\
\frac{h' \text{ fresh} \quad \iota \text{ fresh} \quad \mu' = [\iota \mapsto h']\mu}{\{(E[\text{spawn}(\text{spore } \{x : T = v ; (x : T) \Rightarrow t\})], \mu, Q, S)^h\} \cup H \rightarrow \{(E[\iota], \mu', Q, S)^h, (\text{spore } \{x : T = v ; (x : T) \Rightarrow t\}) \{\}, \epsilon, \epsilon, \epsilon\}^h\} \cup H} \\
\\
\text{R-POPULATE} \\
\frac{\mu(\iota) = h' \quad S'' = [\omega \mapsto (\text{Val}(v), \emptyset)]S' \quad \omega = (h', i) \quad i \text{ fresh}}{\{(E[\text{populate}(\iota, v)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H \rightarrow \{(E[\text{Mat}(\omega)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H} \\
\\
\text{R-REQ1} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q'' \quad S(\text{id}(r)) = (\text{Val}(v), P) \quad m = \text{Res}_s(\omega, v, P)}{\{(E[\text{await}(\iota)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H \rightarrow \{(E[\text{await}(\iota)], \mu, Q'', S)^h, (t', \mu', Q' \cdot m, S')^{h'}\} \cup H} \\
\\
\text{R-REQ2} \\
\frac{Q = \text{Req}_s(h', r, \omega) :: Q'' \quad S(\text{id}(r)) = (\text{Fwd}(r'), P) \\ S(\text{id}(r')) = (\text{Val}(v), P') \quad m = \text{Res}_s(\omega, v, P)}{\{(E[\text{await}(\iota)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H \rightarrow \{(E[\text{await}(\iota)], \mu, Q'', S)^h, (t', \mu', Q' \cdot m, S')^{h'}\} \cup H} \\
\\
\text{R-REQ3} \\
\frac{Q = \text{Req}_s(h'', r, \omega) :: Q'' \quad S(\text{id}(r)) = (\text{Fwd}(r'), P) \\ \text{id}(r') \notin \text{dom}(S) \quad h' = \text{host}(r') \quad m = \text{Req}_s(h'', r', \omega)}{\{(E[\text{await}(\iota)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H \rightarrow \{(E[\text{await}(\iota)], \mu, Q'', S)^h, (t', \mu', Q' \cdot m, S')^{h'}\} \cup H} \\
\\
\text{R-SEND} \\
\frac{\text{host}(r) = h' \quad h' \neq h \quad m = \text{Req}_s(h, r, \text{id}(r)) \quad \iota \text{ fresh} \quad \mu'' = [\iota \mapsto \text{None}]\mu}{\{(E[\text{send}(r)], \mu, Q, S)^h, (t', \mu', Q', S')^{h'}\} \cup H \rightarrow \{(E[\iota], \mu'', Q, S)^h, (t', \mu', Q' \cdot m, S')^{h'}\} \cup H}
\end{array}$$

Figure 7: Non-deterministic reduction.

$$\begin{array}{c}
\text{T-VAR} \\
\frac{x : T \in \Gamma}{\Gamma; \Sigma; \Delta \vdash x : T} \\
\\
\text{T-LOC} \\
\frac{\Sigma(l) : T}{\Gamma; \Sigma; \Delta \vdash l : T} \\
\\
\text{T-ABS} \\
\frac{\Gamma, x : T; \Sigma; \Delta \vdash t : T'}{\Gamma; \Sigma; \Delta \vdash ((x : T) \Rightarrow t) : T \Rightarrow T'} \\
\\
\text{T-APP} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : T \Rightarrow T' \quad \Gamma; \Sigma; \Delta \vdash t' : T}{\Gamma; \Sigma; \Delta \vdash (t t') : T'} \\
\\
\text{T-RECORD} \\
\frac{\Gamma; \Sigma; \Delta \vdash \bar{t} : \bar{T}}{\Gamma; \Sigma; \Delta \vdash \{\bar{l} = t\} : \{\bar{l} : \bar{T}\}} \\
\\
\text{T-SELECT} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \{\bar{l} : \bar{T}\}}{\Gamma; \Sigma; \Delta \vdash t.l_i : T_i} \\
\\
\text{T-SPORE} \\
\frac{\Gamma; \Sigma; \Delta \vdash \bar{t} : \bar{T} \quad \overline{x : T}, x : T; \emptyset; \Delta \vdash t : T' \quad \forall T_i \in \bar{T}. \text{serializable}(T_i)}{\Gamma; \Sigma; \Delta \vdash (\text{spore } \{ \overline{x : T = t}; (x : T) \Rightarrow t \}) : T \Rightarrow T' \{ \text{type } C = \bar{T} \}} \\
\\
\text{T-APPSPORE} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : T \Rightarrow T' \{ \text{type } C = \bar{T} \} \quad \Gamma; \Sigma; \Delta \vdash t' : T}{\Gamma; \Sigma; \Delta \vdash (t t') : T'} \\
\\
\text{T-SPAWN} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : (\{\} \Rightarrow T \{ \text{type } C = \bar{T} \})}{\Gamma; \Sigma; \Delta \vdash \text{spawn}(t) : \text{Host}} \\
\\
\text{T-POPULATE} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{Host} \quad \Gamma; \Sigma; \Delta \vdash t' : T \quad \text{serializable}(T)}{\Gamma; \Sigma; \Delta \vdash \text{populate}(t, t') : \text{SiloRef}[T]} \\
\\
\text{T-MAP} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{SiloRef}[T] \quad \Gamma; \Sigma; \Delta \vdash t' : (T \Rightarrow T' \{ \text{type } C = \bar{T} \})}{\Gamma; \Sigma; \Delta \vdash \text{map}(t, t') : \text{SiloRef}[T']} \\
\\
\text{T-FMAP} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{SiloRef}[T] \quad \Gamma; \Sigma; \Delta \vdash t' : (T \Rightarrow \text{SiloRef}[T'] \{ \text{type } C = \bar{T} \})}{\Gamma; \Sigma; \Delta \vdash \text{flatMap}(t, t') : \text{SiloRef}[T']} \\
\\
\text{T-PERSIST} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{SiloRef}[T]}{\Gamma; \Sigma; \Delta \vdash \text{persist}(t) : \text{SiloRef}[T]} \\
\\
\text{T-SEND} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{SiloRef}[T]}{\Gamma; \Sigma; \Delta \vdash \text{send}(t) : \text{Future}[T]} \\
\\
\text{T-SILOREF} \\
\frac{\Delta(\text{id}(r)) = T \quad \Delta \vdash r}{\Gamma; \Sigma; \Delta \vdash r : \text{SiloRef}[T]} \\
\\
\text{T-AWAIT} \\
\frac{\Gamma; \Sigma; \Delta \vdash t : \text{Future}[T]}{\Gamma; \Sigma; \Delta \vdash \text{await}(t) : T}
\end{array}$$

Figure 8: Type assignment.

$$\begin{array}{c}
\text{S-RECORD} \\
\frac{\forall T_i \in \bar{T}. \text{serializable}(T_i)}{\text{serializable}(\{\bar{l} : T\})} \\
\\
\text{S-SPORE} \\
\frac{\forall T_i \in \bar{T}. \text{serializable}(T_i)}{\text{serializable}(T \Rightarrow T' \{ \text{type } \mathcal{C} = \bar{T} \})} \\
\\
\text{S-SILOREF} \\
\text{serializable}(\text{SiloRef}[T])
\end{array}$$

Figure 9: Serializable types.

thereby providing a typing for silo stores S . Figure 8 shows the rules for type assignment. Rules T-VAR, T-LOC, T-ABS, T-APP, T-RECORD, and T-SELECT are unchanged compared to a standard typed lambda calculus with records [?].

Rule T-SPORE assigns a type to spore literals. Importantly, the body of the spore’s closure, t , must be well-typed in a type environment containing only the closure parameter x and the variables \bar{x} in the spore’s header, as well as an empty store typing. Furthermore, the types of captured variables must be serializable. The predicate *serializable* is defined in Figure 9. These constraints ensure that spore values are always independent of the environment and store of the creating host. This independence is expressed by the following theorem:

Theorem 1.1. (Serializable Values) *If $\Gamma; \Sigma; \Delta \vdash v : T$ and $\text{serializable}(T)$ then $\emptyset; \emptyset; \Delta \vdash v : T$.*

Proof. By induction on the derivation of $\Gamma; \Sigma; \Delta \vdash v : T$. See Appendix ??.

Rule T-APPSPORE is analogous to rule T-APP. Rule T-SPAWN requires argument t to be a spore with domain type unit; the result has type `Host`. Rule T-POPULATE leverages the *serializable* predicate to ensure the value of the silo to be populated is independent of its source context. Rules T-MAP, T-FMAP, and T-PERSIST are straightforward; note that `map` and `flatMap` are polymorphic in the types of the captured variables of their spore argument types. Rules T-SEND and T-AWAIT are entirely unsurprising. Rule T-SILOREF is the only rule that uses the silo store typing Δ . Analogous to rule T-LOC, the type of silo $id(r)$ is looked up in Δ . Furthermore, T-SILOREF requires r to be well-formed in Δ , written $\Delta \vdash r$ (see below).

1.3 Well-Formed Configurations

Figure 10 shows the rules for well-formed configurations. These rules are essential for establishing subject reduction (see Section 2). Rules WF-STORE1 and WF-STORE2 are standard. Rules WF-REF1-2 require the types given by the silo store typing Δ to be consistent with the corresponding type of spore p . Rule WF-REF3 requires the type of silo ω to be equal to the type of its parent silo $id(r)$ in silo store typing Δ . Rule WF-REF4 requires Δ to be defined for the identifier of a materialized silo. Finally, rules WF-REF1-3 require parent silo references to be well-formed. Rules WF-SILOSTORE1-3 require a well-formed silo store to be consistent with silo store typing Δ . Rules WF-Q1-4 specify well-formedness of message queues in Δ and Σ . Rules WF-HOSTCONFIG, WF-HOST1, and WF-HOST2 combine the previous rules in the expected way.

$$\begin{array}{c}
\text{WF-STORE1} \\
\frac{}{\emptyset \vdash \emptyset} \\
\\
\text{WF-REF1} \\
\frac{\Delta(\omega) = T \quad \Delta(\text{id}(r)) = T' \quad \exists \Gamma, \Sigma. \Gamma; \Sigma; \Delta \vdash p : T' \Rightarrow T \quad \{\dots\} \quad \Delta \vdash r}{\Delta \vdash \text{Mapped}(\omega, r, p)} \\
\\
\text{WF-STORE2} \\
\frac{\Sigma \vdash \mu}{[\iota \mapsto T] \Sigma \vdash [\iota \mapsto v] \mu} \\
\\
\text{WF-REF2} \\
\frac{\Delta(\omega) = T \quad \Delta(\text{id}(r)) = T' \quad \exists \Gamma, \Sigma. \Gamma; \Sigma; \Delta \vdash p : T' \Rightarrow \text{SiloRef}[T] \quad \{\dots\} \quad \Delta \vdash r}{\Delta \vdash \text{FMapped}(\omega, r, p)} \\
\\
\text{WF-REF3} \\
\frac{\Delta(\omega) = T \quad \Delta(\text{id}(r)) = T \quad \Delta \vdash r}{\Delta \vdash \text{Persist}(\omega, r, \star)} \\
\\
\text{WF-REF4} \\
\frac{\omega \in \text{dom}(\Delta)}{\Delta \vdash \text{Mat}(\omega)} \\
\\
\text{WF-SILOSTORE1} \\
\Delta \vdash \emptyset \\
\\
\text{WF-SILOSTORE2} \\
\frac{\Delta(\omega) = T \quad \emptyset; \emptyset; \Delta \vdash v : T \quad \Delta \vdash S}{\Delta \vdash [\omega \mapsto (\text{Val}(v), P)]S} \\
\\
\text{WF-SILOSTORE3} \\
\frac{\Delta(\text{id}(r)) = \Delta(\omega) \quad \Delta \vdash r \quad \Delta \vdash S}{\Delta \vdash [\omega \mapsto (\text{Fwd}(r), P)]S} \\
\\
\text{WF-Q1} \\
\Delta; \Sigma \vdash \epsilon \\
\\
\text{WF-Q2} \\
\frac{\Delta(\omega) = T \quad \Sigma(\iota) = \text{Future}[T] \quad \Delta; \Sigma \vdash Q}{\Delta; \Sigma \vdash \text{Req}_\iota(\iota, \omega) :: Q} \\
\\
\text{WF-Q3} \\
\frac{\Delta(\omega) = T \quad \emptyset; \emptyset; \Delta \vdash v : T \quad \Delta; \Sigma \vdash Q}{\Delta; \Sigma \vdash \text{Res}_s(\omega, v, P) :: Q} \\
\\
\text{WF-Q4} \\
\frac{\Delta(\text{id}(r)) = \Delta(\omega) \quad \Delta \vdash r \quad \Delta; \Sigma \vdash Q}{\Delta; \Sigma \vdash \text{Req}_s(h, r, \omega) :: Q} \\
\\
\text{WF-HOSTCONFIG} \\
\frac{\Sigma \vdash \mu \quad \Delta \vdash S \quad \Delta; \Sigma \vdash Q \quad \Gamma; \Sigma; \Delta \vdash t : T}{\Delta; \Sigma \vdash (t, \mu, Q, S)^h} \\
\\
\text{WF-HOST1} \\
\Delta \vdash \emptyset \\
\\
\text{WF-HOST2} \\
\frac{\exists \Sigma. \Delta; \Sigma \vdash (t, \mu, Q, S)^h \quad \Delta \vdash H}{\Delta \vdash \{(t, \mu, Q, S)^h\} \cup H}
\end{array}$$

Figure 10: Well-formedness.

2 Subject Reduction

This section establishes a subject reduction theorem for the presented core language. The complete proof is provided in the appendix; here, we restrict ourselves to summarizing the main results.

Lemma 2.1. (Substitution) *If $\Gamma, x : T'; \Sigma; \Delta \vdash t : T$ and $\Gamma; \Sigma; \Delta \vdash v : T'$ then $\Gamma; \Sigma; \Delta \vdash [x \mapsto v]t : T$.*

Proof. By induction on the derivation of $\Gamma, x : T'; \Sigma; \Delta \vdash t : T$. □

Lemma 2.2. (Queue Concatenation) *If $\Delta; \Sigma \vdash Q$ and $\Delta; \Sigma \vdash Q'$ then $\Delta; \Sigma \vdash Q :: Q'$.*

Proof. By induction on the length of Q . See Appendix ???. □

Theorem 2.1. (Subject Reduction)

1. *If $\Gamma; \Sigma; \Delta \vdash t : T$, $\Sigma \vdash \mu$, and $t \mid \mu \rightarrow^h t' \mid \mu'$ then $\Gamma; \Sigma'; \Delta' \vdash t' : T$, and $\Sigma' \vdash \mu'$ for some $\Sigma' \supseteq \Sigma$ and $\Delta' \supseteq \Delta$.*
2. *If $\Delta; \Sigma \vdash (t, \mu, Q, S)^h$ and $(t, \mu, Q, S)^h \rightarrow (t', \mu', Q', S')^h$ then $\Delta'; \Sigma' \vdash (t', \mu', Q', S')^h$ for some $\Delta' \supseteq \Delta$ and $\Sigma' \supseteq \Sigma$.*
3. *If $\Delta \vdash H$ and $H \rightarrow H'$ then $\Delta' \vdash H'$ for some $\Delta' \supseteq \Delta$.*

Proof. Part 1: by induction on the derivation of $t \mid \mu \rightarrow^h t' \mid \mu'$. Part 2: by induction on the derivation of $(t, \mu, Q, S)^h \rightarrow (t', \mu', Q', S')^h$. Part 3: by induction on the derivation of $H \rightarrow H'$. See Appendix ?? for the complete proof. □