

# Distributed Systems As Programs With Placement Types

Luna Phipps-Costin, Arjun Guha

## Background

Distributed systems like Roblox experiences are often written as a number of programs running on different nodes. These programs interact by passing messages or interacting with a replicated database. However, systems like these are difficult to program and prone to error. Programs with erroneous message passing might result in deadlock, inconsistent state, or race conditions, and all interaction with a replicated database must be correct with respect to the consistency guarantees of the database.

## Introduction

Programs in our language instead describe an entire (distributed) computation. Programmers use type annotations to describe on which nodes computation and effects are performed. Using these annotations, the implementation generates the message-passing necessary to run the program across multiple nodes, and the type system ensures that the protocol is correct. Further, the type system is amenable to best-effort constraint-based type inference, which we plan to leverage to allow programs to elide type annotations while maintaining performance.

## A Note

This poster represents in-progress research. Results may be invalidated and the language presented has multiple problems yet to be worked out.

## Type Inference

If we wrap all expressions in `send`, every simply well-typed program has a set of placement type assignments that type-checks. Further, this type inference task is undecidable (if only due to the use of System F). These two facts suggest a best-effort constraint-based type inference using (possibly programmer-provided) heuristics to guide typing.

## Syntax

$e ::= \text{true } p$	$  \text{ send}_p e$	$u ::= \text{bool}$
$  \text{ false } p$	$  \Lambda P. e$	$  t \rightarrow t$
$  \lambda_p x : t. e$	$  e \text{ } p$	$  \mu \alpha. u$
$  e \text{ } e$	$  \text{ fold}_{\mu \alpha. t} e$	$  \alpha$
$  \text{ ref}_p e$	$  \text{ unfold}_{\mu \alpha. t} e$	$  \text{ ref } t$
$  !e$	$  \text{ pack } \{p, e\} \text{ as } \exists P. t$	$  \{l_1 : u_1, \dots\}$
$  e := e$	$  \text{ unpack } \{P, x\} = v \text{ in } e$	$p ::= \text{node } n \mid P$
$  \{l = e, \dots\}_p$		$t ::= p \text{ } u$
$  e. l$		$  \forall P. t$
		$  \exists P. t$

The standard base of our language is System F plus records, mutation, and recursive types.

## Types

$\frac{\Delta, p \vdash u \text{ ok}}{\Delta, p \vdash p \text{ } u \text{ ok}}$	$\frac{\Gamma \vdash \text{true } p : p \text{ } \text{bool}}{\Gamma, x : t \vdash e : s \quad \{p\} = \{p   p \text{ } u \in \text{fv}(e)\}}$	$\frac{\Gamma \vdash e : q \text{ } u}{\Gamma \vdash \text{send}_p e : p \text{ } u}$
$\frac{\Delta, P \vdash t \text{ ok}}{\Delta \vdash \forall P. t \text{ ok}}$	$\frac{\lambda_p x : t. e : p \text{ } (t \rightarrow s)}{\Gamma \vdash e_1 : q \text{ } (p \text{ } u \rightarrow t) \quad \Gamma \vdash e_2 : p \text{ } u}$	$\frac{\Delta \cup \{P\}, \Gamma \vdash e : t}{\Delta, \Gamma \vdash \Lambda P. e : \forall P. t}$
$\frac{\Delta, P \vdash t \text{ ok}}{\Delta \vdash \exists P. t \text{ ok}}$	$\frac{\Gamma \vdash e_1 \text{ } e_2 : t}{\Gamma \vdash e_1 \text{ } e_2 : t}$	$\frac{\Delta \vdash p \text{ ok} \quad \Delta, \Gamma \vdash e : \forall P. t}{\Delta, \Gamma \vdash e \text{ } p : t[P / p]}$
$\frac{\Delta \vdash \text{bool ok}}{\Delta \vdash \text{bool ok}}$	$\frac{\Gamma \vdash e : p \text{ } u}{\Gamma \vdash \text{ref}_p e : p \text{ } \text{ref } u}$	$\frac{\Gamma \vdash e : \mu \alpha. t}{\Gamma \vdash \text{unfold } e : t[\alpha / \mu \alpha. t]}$
$\frac{\Delta \vdash t_1 \text{ ok} \quad \Delta \vdash t_2 \text{ ok}}{\Delta \vdash t_1 \rightarrow t_2 \text{ ok}}$	$\frac{\Gamma \vdash e : t}{\Gamma \vdash \text{ref}_p e : p \text{ } \text{ref } t}$	$\frac{\Gamma \vdash e : t[\alpha / \mu \alpha. t]}{\Gamma \vdash \text{fold } e : \mu \alpha. t}$
$\frac{\Delta \vdash t \text{ ok}}{\Delta \vdash \text{ref } t \text{ ok}}$	$\frac{\Gamma \vdash e : p \text{ } \text{ref } t}{\Gamma \vdash !e : t}$	$\frac{\Delta, \Gamma \vdash e : t[P / p] \quad \Delta \vdash \exists P. t \text{ ok}}{\Delta, \Gamma \vdash \text{pack } \{P, e\} \text{ as } \exists P. t : \exists P. t}$
Valid types	$\frac{\Gamma \vdash e_1 : p \text{ } \text{ref } t \quad \Gamma \vdash e_2 : t}{\Gamma \vdash e_1 := e_2 : p \text{ } \text{ref } t}$	$\frac{\Delta, \Gamma \vdash e_1 : \exists P. t_1 \quad \Delta \cup \{P\}, \Gamma, x : t \vdash e_2 : t_2 \quad \Delta \vdash t_2 \text{ ok}}{\Delta, \Gamma \vdash \text{unpack } \{P, x\} = e_1 \text{ in } e_2 : t_2}$
	$\frac{\Gamma \vdash e : p \text{ } u, \dots}{\Gamma \vdash \{l = e, \dots\}_p : p \text{ } \{l : u, \dots\}}$	
	$\frac{\Gamma \vdash e : p \text{ } \{l : u, \dots\}}{\Gamma \vdash e. l : p \text{ } u}$	

We extend the conventional type system by:

- Pairing each type with a placement annotation,
- supporting placement polymorphism à la System F, and
- supporting placement existentials à la ML modules.

Note that records are homogenously placed, that referenced values can never be sent, and that lambda must live where its free variables live.

## Semantics

$s; \lambda_p x : t. e \rightarrow_p s; p \lambda x : t. e$	$\frac{s; H[r] \rightarrow_q s'; H[r]; p \neq q}{p; s; H[r] \xrightarrow{\text{msg}(p, q, H[r])} q; s; H[r]}$	
$s; (p \text{ } \lambda x. e) \text{ } v \rightarrow_p s; e[x / v]$		
$s; \Lambda P. e \rightarrow_{\text{?}} s; e$		
$s; \{l = p \text{ } c, \dots\}_p \rightarrow_p s; p \text{ } \{l = c, \dots\}$		Using the steps, the above describes message passing.
$s; (p \text{ } \{l = c, \dots\}). l \rightarrow_p s; p \text{ } c$		
$s; \text{ref}_q(p \text{ } c) \rightarrow_p s, l \mapsto p \text{ } c; \text{send}_q(p \text{ } l) \text{ } (l \text{ fresh})$		
$s, l \mapsto v; !(p \text{ } l) \rightarrow_p s, l \mapsto v; v$		It is obvious from our formalization that each step occurs on exactly one node.
$s; (p \text{ } l) := v \rightarrow_p s, l \mapsto v; p \text{ } l$		
$s; \text{send}_q(p \text{ } c) \xrightarrow{\text{msg}(p, q, c)}_p s; q \text{ } c$		
$s; \text{fold}(p \text{ } c) \rightarrow_p s; p \text{ } c$		
$s; \text{unfold}(p \text{ } c) \rightarrow_p s; p \text{ } c$		

We also intend to prove that a node only ever dereferences locations it allocated.

We elide type well-formedness checking in most typing rules; they are standard.

## Next Steps

Our next task is to implement the dynamic semantics in our prototype. Using our implementation, we will be able to write example programs to gain confidence in the usefulness of the current design. Our first benchmarks will be an arbitrary-node producer-consumer queue and a simple client-server benchmark from the literature. Once we are confident in the explicitly typed language, we will begin work on our type inference system. In the longer term, much further work could be done on the basis of Mycelium, such as support for concurrency, automatic parallelization, and client speculation.

## Related Work

The PLT Scheme Web Server (Krishnamurthi et al 2007) supports sending continuations from the server to the client and vice versa, modeling Web navigation as invoking continuations. Swift (Chong et al 2007) usse Information-Flow Control (based on JIF) to describe programs, which then are split between client and server using a heuristic based on the min-flow problem. Pyxis (Cheung et al 2012) allows parts of programs to be offloaded to the database server with no annotations to improve performance. It also uses a program graph to optimize the split. Links (Hutchison et al 2007) is a language with server and client annotations on functions, allowing programs to be written that span both. Fission (Guha et al 2017) distributes the program dynamically.

Our language is type-based, distributes on the expression level, and supports functions which capture their environment, while the above work supports at most one of these three features. Fabric (Liu et al 2009) represents a scalable, multi-node, mutually distrustful programming system for distributed systems. While Fabric emphasizes concurrency and state synchronization, our language prioritizes linear, deterministic protocols proven by the type system.