

Introduction

The Message Passing Interface (MPI) [1] is the standard for high-performance computing (HPC) applications, ensuring correct communication in massively parallel systems. However, since MPI is often implemented in imperative languages, programs are inherently hard to reason about. We present λ_{MPI} , a new Concurrent Lambda Calculus to formalize the semantics of MPI.

System Overview

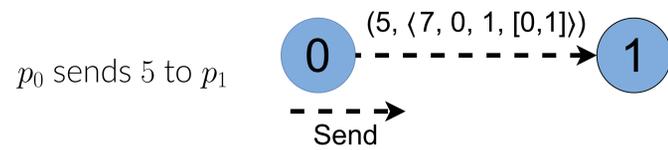
System Level:

- Program ρ is a set number of named processes p_n , each running an expression e_n in parallel.
- State σ is a set of FIFO queues between each ordered pair of nodes $(p_n p_m)$.
- A configuration $\langle \rho, \sigma \rangle$ captures full system.

Local Level:

- Processes know only their expression.
- Labelled transition system (LTS) forms operational semantics.

Point to Point Example



$p_0 \triangleright \text{send } 5 \text{ tag } 7 \text{ to } 1 \text{ of } \mathbb{G}[0, 1]^{i_1} \xrightarrow{5 \rightsquigarrow \langle 7, 0, 1, [0, 1]^{i_1} \rangle} ()$

$p_1 \triangleright \text{recv } x \text{ in } x \text{ select } \langle \text{tag } 7, \text{from } 0 \rangle \text{ of } \mathbb{G}[0, 1]^{i_2} \xrightarrow{\langle 7, 0, 1, 5 \rangle \rightsquigarrow \mathbb{G}[0, 1]^{i_1}} 5$

Envelopes: $\langle \text{tag}, \text{source}, \text{dest}, \text{group} \rangle$

- All information of where a message need to go.
- Components appear in expressions, but full envelopes are only found in labels.

λ_{MPI} supports **blocking** and **non-blocking sends** and **receives**.

Collective Communication Example

$p_6 \triangleright$
gather 1 from
[scatter 1 of $[0; 0; 0; 0]^i$
as x in $(x + (\text{rank } \mathbb{G}^i))$
to $\mathbb{G}[6, 3, 0, 2]^i$]^j
and $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{[[0]^i; [0]^i; [0]^i; [0]^i] \rightsquigarrow_{\text{sc}} \mathbb{G}[6, 3, 0, 2]^i}$

gather 1 from
[$0 + (\text{rank } \mathbb{G}^i)$]^j
and $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{\text{rank}(\mathbb{G}^i) \rightsquigarrow 6}$

gather 1 from $[6]^j$
and $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{[[3]^i; [0]^i; [2]^i; [2]^i] \rightsquigarrow_{\text{ga}} \mathbb{G}[6, 3, 0, 2]^i}$

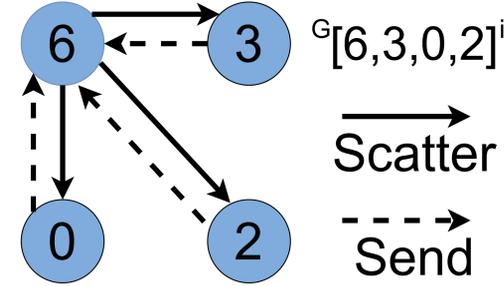
$[6; 3; 0; 2]^i$

$p_3 \triangleright$
send
[recv x in
 $(x + (\text{rank } \mathbb{G}^i))$
select $\langle \text{tag } *, \text{from } 0 \rangle$
of $\mathbb{G}[6, 3, 0, 2]^i$]^j
tag 0 to 0 of $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{\langle 0, 0, 1, \mathbb{G}[6, 3, 0, 2]^i \rangle \rightsquigarrow [0]^j}$

send $[(0 + (\text{rank } \mathbb{G}^i))]^j$
tag 0 to 0 of $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{\text{rank}(\mathbb{G}^i) \rightsquigarrow 3}$

send $[3]^j$
tag 0 to 0 of $\mathbb{G}[6, 3, 0, 2]^i$
 $\xrightarrow{[3]^j \rightsquigarrow_3 \langle 0, 1, 0, \mathbb{G}[6, 3, 0, 2]^i \rangle}$

()



All Processes:

- Receive the value 0 from the root.
- Compute their rank in $\mathbb{G}[6, 3, 0, 2]^i$.
- Returns their rank added to the 0.

Collective Communication:

λ_{MPI} supports MPI **Broadcast**, **Scatter**, and **Gather** as primitives as well as **Reduce** and **Scan**.

Selection of Syntax and Stepping Rules

Expressions $e ::= x \mid \tilde{n} \mid () \mid \text{error}(i) \mid * \mid \langle \text{tag } e_1, \text{from } e_2 \rangle \mid \text{send } e \text{ tag } n \text{ to } p \text{ of } g \mid \text{recv } x \text{ in } e_1 \text{ select } e_2 \text{ of } g$
 $\mid \text{scatter } n \text{ of } e_1 \text{ as } x \text{ in } e_2 \text{ to } g \mid \text{gather } n \text{ from } e \text{ and } g \mid [e_0; \dots; e_n]^i \mid \mathbb{G}[e_0, \dots, e_n]^i \mid \text{rank } g$
 Values $v ::= () \mid \tilde{n} \mid \text{error}(i) \mid * \mid \langle \text{tag } v_1, \text{from } v_2 \rangle \mid [v_1; \dots; v_n]^i \mid \mathbb{G}^i \mid \emptyset^i \mid \mathbb{G}[n]^i \mid \mathbb{G}[n_0, \dots, n_m]^i$
 Labels $l ::= \epsilon \mid v \rightsquigarrow \langle t, p, q, g \rangle \mid \langle t, p, q, g \rangle \rightsquigarrow v \mid v \rightsquigarrow_{\text{br}} g \mid v \rightsquigarrow_{\text{sc}} g \mid g \rightsquigarrow_{\text{ga}} v \mid \text{rank}(g) \rightsquigarrow n$
 Programs $\rho ::= p_1 \triangleright e_1 \parallel \dots \parallel p_n \triangleright e_n$

Send $\frac{}{\text{send } v \text{ tag } t \text{ to } q \text{ of } g \xrightarrow{v \rightsquigarrow \langle t, \text{rank } g, q, g \rangle} ()}$ Recv $\frac{}{\text{recv } x \text{ in } e \text{ select } \langle \text{tag } t_1, \text{from } p_1 \rangle \text{ of } g \xrightarrow{\langle t_2, p_2, q, g \rangle \rightsquigarrow v} e[x \mapsto v]}$

Scatter $\frac{(n_2 + 1) = n_1(n_3 + 1) \quad (\forall b. 1 \leq b \leq n_3 \wedge a_b = [v_{bm_1}; \dots; v_{(b+1)n_1-1}]^c, \text{ where } c \text{ is fresh})}{\text{scatter } n_1 \text{ of } [v_0; \dots; v_{n_2}]^{i_1} \text{ as } x \text{ in } e \text{ to } \mathbb{G}[p_0, \dots, p_{n_3}]^{i_2} \xrightarrow{[a_1; \dots; a_{n_3}]^{i_3} \rightsquigarrow_{\text{sc}} \mathbb{G}[p_0, \dots, p_{n_3}]^{i_2}} e[x \mapsto [v_1; \dots; v_{(n_1-1)}]^{i_4}]}$

$\rho = \mathbb{G}[x_0, \dots, x_n]^i \quad 0 \leq p, q \leq n \quad \rho(p_{x_p}) = e_{x_p} \quad \rho(p_{x_q}) = e_{x_q} \quad e_{x_p} \xrightarrow{v \rightsquigarrow \langle t_1, p, q, g \rangle} e'_{x_p}$
 $\frac{e_{x_q} \xrightarrow{\langle t_2, p_2, q, g \rangle \rightsquigarrow v} e'_{x_q} \quad \text{true} = \text{tmatch}(t_1, t_2) \quad \text{true} = \text{rmatch}(p, p_2) \quad \sigma(\text{peek } p_{x_p} p_{x_q}) = \perp}{\text{PSend} \quad \langle \rho, \sigma \rangle \rightarrow_{\rho} \langle \rho[p_{x_p} \mapsto e'_{x_p}, p_{x_q} \mapsto e'_{x_q}], \sigma \rangle}$

$\rho(p_{x_0}) = e_{x_0} \quad e_{x_0} \xrightarrow{[v_1; \dots; v_n]^{i_2} \rightsquigarrow_{\text{sc}} g} e'_{x_0} \quad (\forall m. 1 \leq m \leq n \wedge \rho(p_{x_m}) = e_{x_m} \wedge e_{x_m} \xrightarrow{\langle t, 0, m, g \rangle \rightsquigarrow v_m} e'_{x_m} \wedge \text{peek } p_{x_0} p_{x_m} = \perp)$
 $\frac{g = \mathbb{G}[x_0, x_1, \dots, x_n]^{i_1}}{\text{PSc} \quad \langle \rho, \sigma \rangle \rightarrow_{\rho} \langle \rho[p_{x_m} \mapsto e'_{x_m} \mid \forall m. 0 \leq m \leq n], \sigma \rangle}$

Group Operations

MPI Groups:

- Logical collections of processes.
- Each process p has a rank 0 to $(n - 1)$.
- Rank 0 is the **root** of a group.
- Written as lists of **global ranks**, index of the rank being the **local rank**.

Default groups: n processes and $0 \leq m < n$

- Global Group: $\mathbb{G}^i = \mathbb{G}[0; \dots; (n - 1)]^i$
- Empty Group: $\emptyset^i = \mathbb{G}[]^i$
- Singleton Group: $\mathbb{G}[m]^i$

With the default groups λ_{MPI} can construct all valid groups, such as $\mathbb{G}[6; 3; 0; 2]^i$, where global rank 6 has local rank 0.

Group Constructors:

- **gunion** $\mathbb{G}[6; 3; 0; 2]^{i_1} \mathbb{G}[1; 2; 5; 3]^{i_2} \xrightarrow{\epsilon} \mathbb{G}[6; 3; 0; 2; 1; 5]^{i_3}$
- **gsplit** $(\text{fun } f \ x := x \ \text{mod } 2) \mathbb{G}[6; 3; 0; 2]^{i_1} \xrightarrow{\epsilon} [\mathbb{G}[6; 0]^{i_3}; \mathbb{G}[3; 2]^{i_4}]^{i_2}$
- **gshuffle** $[3; 0; 2; 1]^{i_1} \mathbb{G}[6, 3, 0, 2]^{i_2} \xrightarrow{\epsilon} \mathbb{G}[2, 6, 0, 3]^{i_3}$

λ_{MPI} also supports other MPI group constructors, such as **intersection** and **difference**, using these three defined operations.

Future Work

- Formalize λ_{MPI} in the Rocq proof assistant.
- Use λ_{MPI} as the target language in a choreographic [2] HPC language based on MPI.
- Explore program analysis and verification using this choreographic language.

References

- 1 Message Passing Interface Forum. 2025. MPI: A Message-Passing Interface Standard Version 5.0. <https://www.mpi-forum.org/docs/mpi-5.0/mpi50-report.pdf>
- 2 Fabrizio Montesi. 2022. Introduction to Choreographies. Cambridge University Press. <https://doi.org/10.1017/9781108981491>