

SPMD IR: Unifying SPMD and Multi-value IR Showcased for Static Verification of Collectives

Semih Burak, I. Ivanov, J. Domke, M. Müller

EuroMPI 2024
25 - 27.09.24

Motivation

- HPC systems get more complex and **heterogeneous**; the **software** and hardware

Motivation

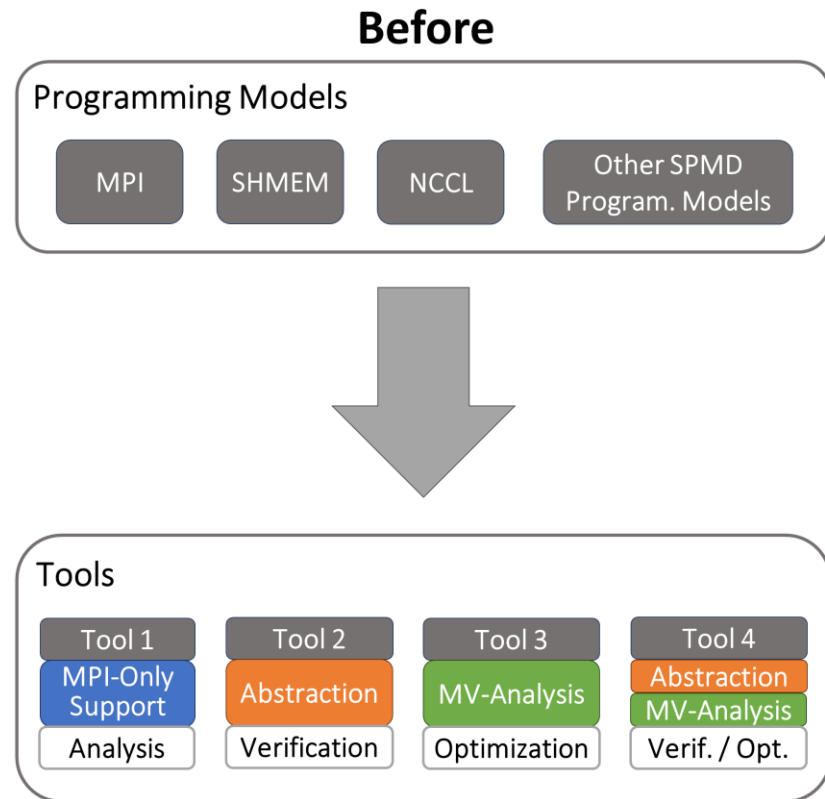
- HPC systems get more complex and **heterogeneous**; the **software** and hardware
- For **effective utilization** of HPC clusters,
 - inter-node comm. and SPMD programming models such as MPI are **inevitable**

Motivation

- HPC systems get more complex and **heterogeneous**; the **software** and hardware
- For **effective utilization** of HPC clusters,
 - inter-node comm. and SPMD programming models such as MPI are **inevitable**
- There exist different **programming models** with distinct specialization but with **common features**

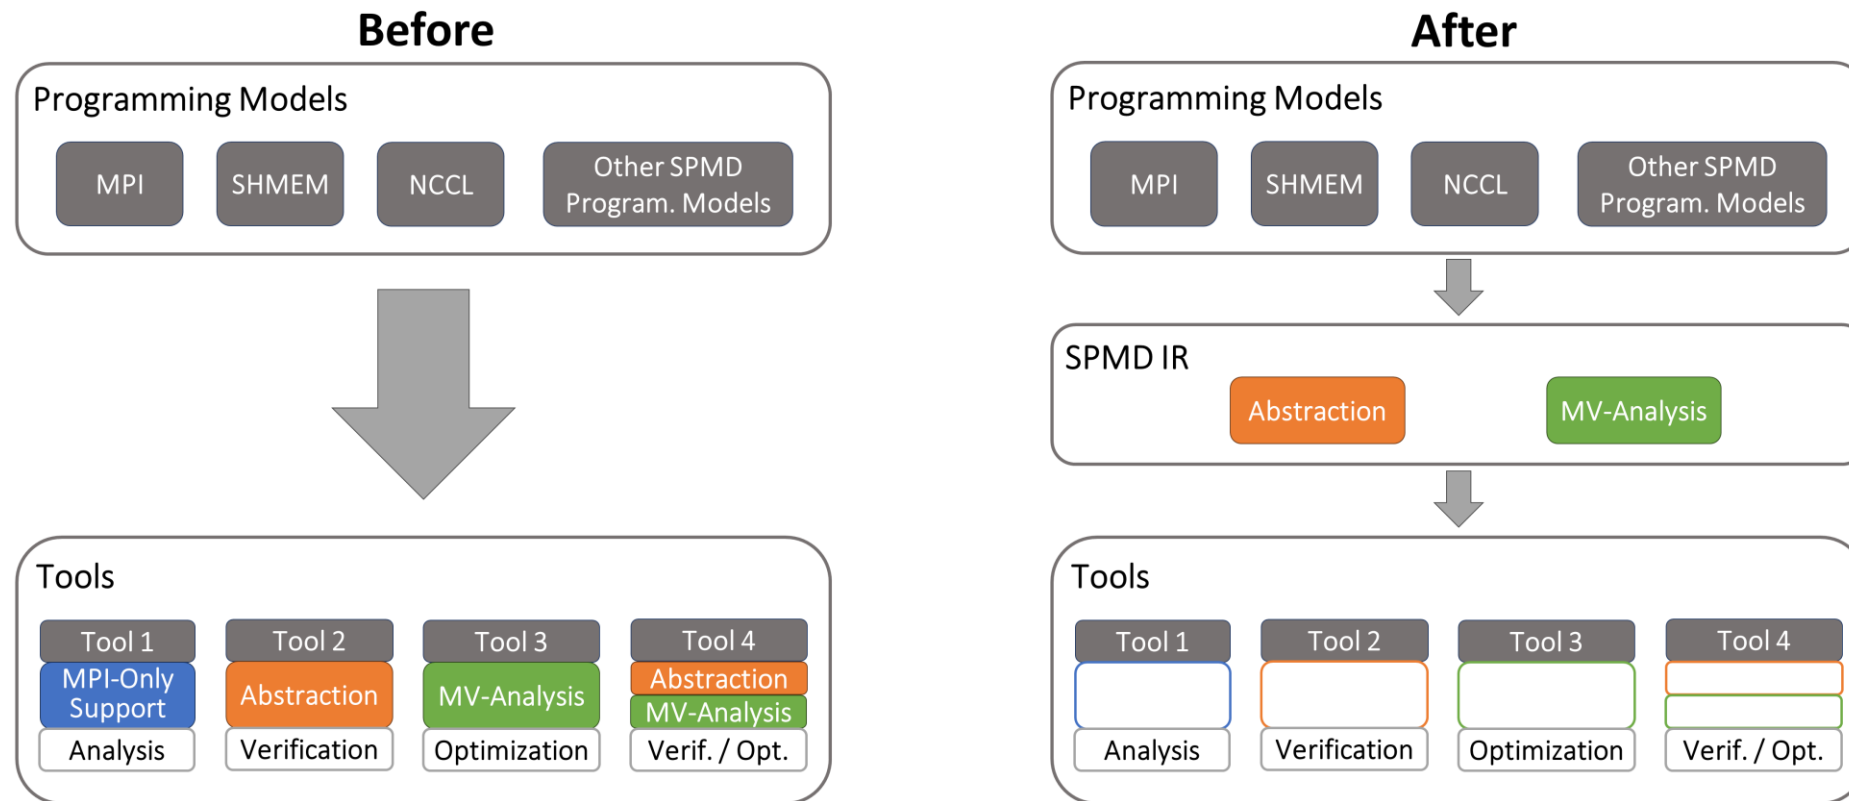
Motivation

- HPC systems get more complex and **heterogeneous**; the **software** and hardware
- For **effective utilization** of HPC clusters,
 - inter-node comm. and SPMD programming models such as MPI are **inevitable**
- There exist different **programming models** with distinct specialization but with **common features**



Motivation

- HPC systems get more complex and **heterogeneous**; the **software** and hardware
- For **effective utilization** of HPC clusters,
 - inter-node comm. and SPMD programming models such as MPI are **inevitable**
- There exist different **programming models** with distinct specialization but with **common features**



Outline

- Background
 - MLIR
 - Multi-Values

Outline

- Background
 - MLIR
 - Multi-Values
- SPMD IR
 - Extended Multi-Value Analysis
 - Scope
 - Example
 - Workflow

Outline

- Background
 - MLIR
 - Multi-Values
- SPMD IR
 - Extended Multi-Value Analysis
 - Scope
 - Example
 - Workflow
- Evaluation
 - Static Verification of Collectives
 - Comparison to PARCOACH

Outline

- Background
 - MLIR
 - Multi-Values
- SPMD IR
 - Extended Multi-Value Analysis
 - Scope
 - Example
 - Workflow
- Evaluation
 - Static Verification of Collectives
 - Comparison to PARCOACH
- Discussion

Background: MLIR (Multi-Level Intermediate Representation)

[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021

Background: MLIR (Multi-Level Intermediate Representation)

[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021
- The purpose of IRs in compilers is generally to facilitate optimizations
 - before lowering to low-level or machine code

Background: MLIR (Multi-Level Intermediate Representation)

[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021
- The purpose of IRs in compilers is generally to facilitate optimizations
 - before lowering to low-level or machine code
- Compared to LLVMIR, MLIR
 - is more extensible, reusable
 - can mix low-level and high-level abstractions

Background: MLIR (Multi-Level Intermediate Representation)

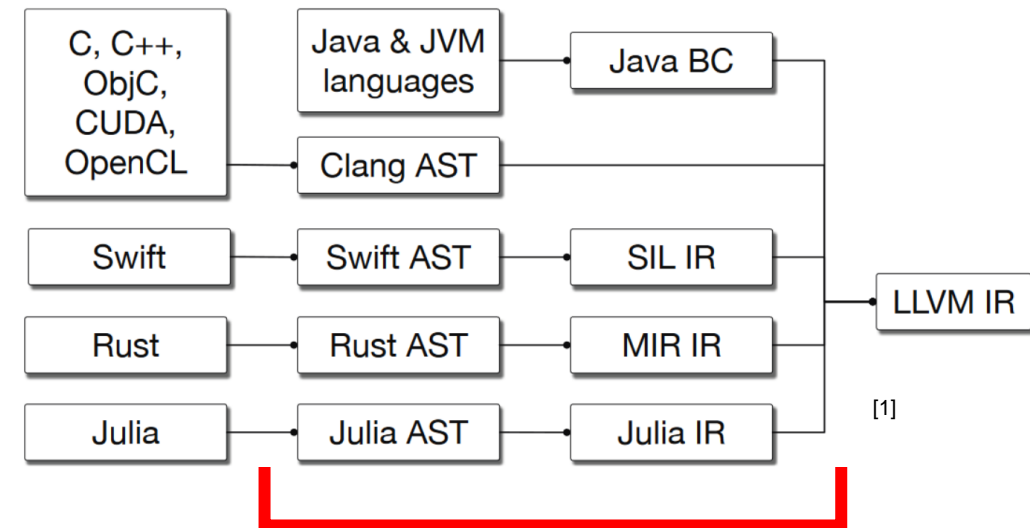
[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021
- The purpose of IRs in compilers is generally to facilitate optimizations
 - before lowering to low-level or machine code
- Compared to LLVMIR, MLIR
 - is more extensible, reusable
 - can mix low-level and high-level abstractions
 - can encapsulate domain-specific semantics by dialects
 - a collection of custom operations, types, and attributes

Background: MLIR (Multi-Level Intermediate Representation)

[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021
- The purpose of IRs in compilers is generally to facilitate optimizations
 - before lowering to low-level or machine code
- Compared to LLVMIR, MLIR
 - is more extensible, reusable
 - can mix low-level and high-level abstractions
 - can encapsulate domain-specific semantics by dialects
 - a collection of custom operations, types, and attributes

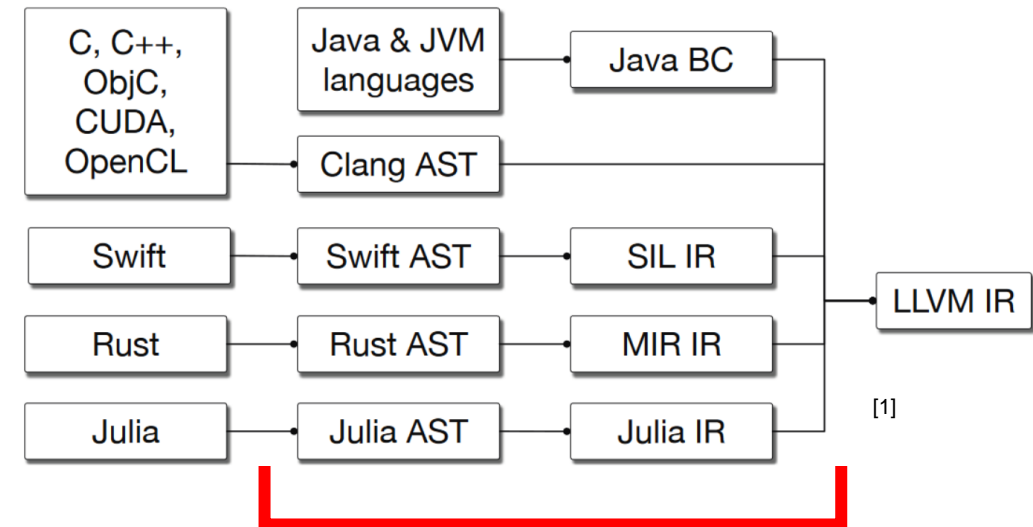


MLIR also helps here, replacing custom high-level IRs

Background: MLIR (Multi-Level Intermediate Representation)

[1] Lattner et al: MLIR: Scaling Compiler Infrastructure for Domain Specific Computation. CGO '21

- IR and compiler framework building upon and part of LLVM, published 2021
- The purpose of IRs in compilers is generally to facilitate optimizations
 - before lowering to low-level or machine code
- Compared to LLVMIR, MLIR
 - is more extensible, reusable
 - can mix low-level and high-level abstractions
 - can encapsulate domain-specific semantics by dialects
 - a collection of custom operations, types, and attributes



- Idea:
 - Express SPMD-related information in its own dialect and a set of attributes
 - Next to other special-purpose dialects, e.g., for shared-memory parallelism

MLIR also helps here, replacing custom high-level IRs

Background: MLIR (2)

```
1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14 func.func @main(...) -> i32 {  
15     ... // define constants and other omitted operations  
16     %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17     ... // fill "array" (ref) with values  
18     func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19     ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

Background: MLIR (2)

```
1 func.func @multPosByX(%arg_ref: memref<?xf64>, %arg_size: index,  
    %arg_c_x: f64) -> () {  
2   %c0_index = arith.constant() {value = 0 : index} () -> index  
3   ... // define other constants  
4  
5  
6  
7  
8  
9  
10  
11  
12 }  
13  
14 func.func @main(...) -> i32 {  
15   ... // define constants and other omitted operations  
16   %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17   ... // fill "array" (ref) with values  
18   func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19   ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

Background: MLIR (2)

```
1 func.func @multPosByX(%arg_ref: memref<?xf64>, %arg_size: index,  
    %arg_c_x: f64) -> () {  
2   %c0_index = arith.constant() {value = 0 : index} () -> index  
3   ... // define other constants  
4   scf.for %idx = %c0_index to %arg_size step %c1_index  
    : (index, index, index, i32) -> () {  
5  
6  
7     scf.if(%cmpRes) : (i1) -> () {  
8  
9  
10    }  
11  }  
12 }  
13  
14 func.func @main(...) -> i32 {  
15   ... // define constants and other omitted operations  
16   %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17   ... // fill "array" (ref) with values  
18   func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19   ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

Background: MLIR (2)

```
1 func.func @multPosByX(%arg_ref: memref<?xf64>, %arg_size: index,  
    %arg_c_x: f64) -> () {  
2   %c0_index = arith.constant() {value = 0 : index} () -> index  
3   ... // define other constants  
4   scf.for %idx = %c0_index to %arg_size step %c1_index  
    : (index, index, index, i32) -> () {  
5     %value = memref.load(%arg_ref, %idx) : (memref<?xf64>, index) -> f64  
6     %cmpRes = arith.cmpf (%value, %c0_f64) {...olt..} : (f64, f64) -> i1  
7     scf.if(%cmpRes) : (i1) -> () {  
8  
9  
10    }  
11  }  
12 }  
13  
14 func.func @main(...) -> i32 {  
15   ... // define constants and other omitted operations  
16   %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17   ... // fill "array" (ref) with values  
18   func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19   ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

Background: MLIR (2)

```
1 func.func @multPosByX(%arg_ref: memref<?xf64>, %arg_size: index,  
    %arg_c_x: f64) -> () {  
2   %c0_index = arith.constant() {value = 0 : index} () -> index  
3   ... // define other constants  
4   scf.for %idx = %c0_index to %arg_size step %c1_index  
    : (index, index, index, i32) -> () {  
5     %value = memref.load(%arg_ref, %idx) : (memref<?xf64>, index) -> f64  
6     %cmpRes = arith.cmpf (%value, %c0_f64) {...olt..} : (f64, f64) -> i1  
7     scf.if(%cmpRes) : (i1) -> () {  
8       %newValue = arith.mulf(%value, %arg_c_x) : (f64, f64) -> f64  
9       memref.store(%newValue, %arg_ref, %idx):(f64,memref<?xf64>,index) ->()  
10    }  
11  }  
12 }  
13  
14 func.func @main(...) -> i32 {  
15   ... // define constants and other omitted operations  
16   %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17   ... // fill "array" (ref) with values  
18   func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19   ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

Background: MLIR (2)

```
1 func.func @multPosByX(%arg_ref: memref<?xf64>, %arg_size: index,  
    %arg_c_x: f64) -> () {  
2   %c0_index = arith.constant() {value = 0 : index} () -> index  
3   ... // define other constants  
4   scf.for %idx = %c0_index to %arg_size step %c1_index  
    : (index, index, index, i32) -> () {  
5     %value = memref.load(%arg_ref, %idx) : (memref<?xf64>, index) -> f64  
6     %cmpRes = arith.cmpf (%value, %c0_f64) {...olt..} : (f64, f64) -> i1  
7     scf.if(%cmpRes) : (i1) -> () {  
8       %newValue = arith.mulf(%value, %arg_c_x) : (f64, f64) -> f64  
9       memref.store(%newValue, %arg_ref, %idx):(f64,memref<?xf64>,index) ->()  
10    }  
11  }  
12 }  
13  
14 func.func @main(...) -> i32 {  
15   ... // define constants and other omitted operations  
16   %ref = memref.alloca(%size) : (index) -> memref<?xf64>  
17   ... // fill "array" (ref) with values  
18   func.call @multPosByX(%ref, %size, %c_x):(memref<?xf64>, index, f64) -> ()  
19   ...  
20 }
```

- Red highlights the dialects
- Orange the values
- Blue the types

- As each dialect can seamlessly interact with other dialects
 - MLIR provides a toolset to develop an expressive IR composed of multiple levels and semantics

Background: Multi-Values

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4
5
6
7
8
9
10
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4  if (rank == 0) { // MVed conditional
5
6
7
8  }
9
10
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4  if (rank == 0) { // MVed conditional
5      a = 42; // "a" -> MV
6      b = 38; // MV
7      ...
8  }
9
10
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4  if (rank == 0) { // MVed conditional
5      a = 42; // "a" -> MV
6      b = 38; // MV
7      ...
8  }
9  b = 64; // "b" -> SV
10 int c = a + b; // MV
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4  if (rank == 0) { // MVed conditional
5      a = 42; // "a" -> MV
6      b = 38; // MV
7      ...
8  }
9  b = 64; // "b" -> SV
10 int c = a + b; // MV
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise
- **MV-seed** operations yield MVs **per definition**

Background: Multi-Values

```
1  int a = 13; // SV
2  int rank, b;
3  MPI_Comm_rank(MPI_COMM_WORLD, &rank); // MV-Seed Op: "rank" -> MV
4  if (rank == 0) { // MVed conditional
5      a = 42; // "a" -> MV
6      b = 38; // MV
7      ...
8  }
9  b = 64; // "b" -> SV
10 int c = a + b; // MV
```

- A variable or value is called **multi-value (MV)**
 - if its runtime-value **may differ** for a subset of processes
 - and called **single-value (SV)** otherwise
- **MV-seed** operations yield MVs **per definition**
- **New MVs** can result from **operations using MVs**
 - or through **control-flow operations**

SPMD IR: Extended Multi-Value Analysis

[2] Tiotto et al.: Experiences Building an MLIR-Based SYCL Compiler. CGO '24

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR

SPMD IR: Extended Multi-Value Analysis

[2] Tiotto et al.: Experiences Building an MLIR-Based SYCL Compiler. CGO '24

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR
- Add [execution kind](#) and [executing processes](#) analysis

SPMD IR: Extended Multi-Value Analysis

[2] Tiotto et al.: Experiences Building an MLIR-Based SYCL Compiler. CGO '24

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR
- Add [execution kind](#) and [executing processes](#) analysis
- Express the results of this [extended MV-analysis](#) in the IR by [attributes](#)

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR
- Add [execution kind](#) and [executing processes](#) analysis
- Express the results of this [extended MV-analysis](#) in the IR by [attributes](#)
 - `isMultiValued`: True or False
 - Tagging [conditionals](#) as depending on an MV

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR
- Add [execution kind](#) and [executing processes](#) analysis
- Express the results of this [extended MV-analysis](#) in the IR by [attributes](#)

- [isMultiValued](#): True or False
 - Tagging [conditionals](#) as depending on an MV

- [executionKind](#): All, AllBut, Static, Dynamic, One, AllButOne
 - Specify for each operation the information given at [compile time](#)

- Adapt [existing work](#)² for MV-analysis for SPMD programs in MLIR
- Add [execution kind](#) and [executing processes](#) analysis
- Express the results of this [extended MV-analysis](#) in the IR by [attributes](#)

- [isMultiValued](#): True or False
 - Tagging [conditionals](#) as depending on an MV

- [executionKind](#): All, AllBut, Static, Dynamic, One, AllButOne
 - Specify for each operation the information given at [compile time](#)

- [executed\(Not\)By](#): List of process IDs (not) executing the operation
 - Given for Static and AllBut cases

SPMD IR: Extended Multi-Value Analysis

```
1 int a = b + 5; // Executed by 'All'
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank); // Executed by 'All'
3 if (rank == 0 || rank == 1) { //isMVed
4     MPI_Barrier(...); // Executed by [0, 1] ('Static')
5 }
6 else {
7     MPI_Barrier(...); // Executed by 'AllBut' [0, 1]
8 }
```

- **isMultiValued**: True or False
 - Tagging **conditionals** as depending on an MV
- **executionKind**: All, AllBut, Static, Dynamic, One, AllButOne
 - Specify for each operation the information given at **compile time**
- **executed(Not)By**: List of process IDs (not) executing the operation
 - Given for Static and AllBut cases

SPMD IR: Extended Multi-Value Analysis

```
1  if (rank == i) { // isMVed and "i" being SV but not known at compile time
2      MPI_Barrier(...); // Executed by 'One'
3  }
4  else {
5      MPI_Barrier(...); // Executed by 'AllButOne'
6  }
```

- **isMultiValued**: True or False
 - Tagging **conditionals** as depending on an MV
- **executionKind**: All, AllBut, Static, Dynamic, One, AllButOne
 - Specify for each operation the information given at **compile time**
- **executed(Not)By**: List of process IDs (not) executing the operation
 - Given for Static and AllBut cases

SPMD IR: Scope

Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

SPMD IR: Scope

Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management		
Data Management Collective Comm.		
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)		

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management Collective Comm.		
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)		

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management Collective Comm.	MPI, SHMEM, NCCL	malloc, realloc, free
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)		

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofComm, getRankInComm, getDeviceInComm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)		

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	send, recv

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commsplit, commdestroy, commsplitstrided, commworld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reducescatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**
 - **custom** dialect (spmd) types for the remaining, e.g.:

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**
 - **custom** dialect (spmd) types for the remaining, e.g.:
 - a communicator of type **spmd comm**

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**
 - **custom** dialect (spmd) types for the remaining, e.g.:
 - a communicator of type **spmd comm**
 - an error/success value of type **spmd error**

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofComm, getRankInComm, getDeviceInComm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**
 - **custom** dialect (spmd) types for the remaining, e.g.:
 - a communicator of type **spmd comm**
 - an error/success value of type **spmd error**

- For memory references, **IR** specifies **side-effects**, e.g., read-only

SPMD IR: Scope

Concept	Supported by	SPMD IR Operations
Process Management	MPI, SHMEM, NCCL	init, finalize, getsizeofcomm, getrankincomm, getdeviceincomm
Communicator Management	MPI, SHMEM, NCCL	commSplit, commDestroy, commSplitStrided, commWorld
Data Management	MPI, SHMEM, NCCL	malloc, realloc, free
Collective Comm.	MPI, SHMEM, NCCL	bcast, reduce, allreduce, scatter, reduceScatter, gather, allgather, alltoall, scan, exscan, barrier
Point-To-Point Comm.	MPI, NCCL	send, recv
Non-Blocking Semantics (P2P and Collectives)	MPI, NCCL	wait{All,Some,Any}, test{All,Some,Any}

- Idea:
 - Having a **minimal** set of operations, types, and attributes that can cover those features for **unification** purposes

- Attributes for **blocking** or **buffered** communication
- Types:
 - **integers** for process and device **IDs**, or **tags**
 - any **memref** type for communicated **data**
 - **custom** dialect (spmd) types for the remaining, e.g.:
 - a communicator of type **spmd comm**
 - an error/success value of type **spmd error**

- For memory references, **IR** specifies **side-effects**, e.g., read-only
- RMA is work in progress

SPMD IR: Example

```
1  int rank;  
2  MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3  if (rank == 0) {  
4      MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5  }
```

SPMD IR: Example

```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

SPMD IR: Example

```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

```
1 %rank, %error1 = spmd.getRankInComm(%comm)  
  
2 %cmpRes = arith.cmpi eq (%rank, %c0)  
3 scf.if (%cmpRes) {  
4     %error2 = spmd.bcast (%comm, %sendBuf, %recvBuf, %count, %i32Type, %c0)  
  
5 }
```


SPMD IR: Example

```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

```
1 %rank, %error1 = spmd.getRankInComm(%comm)  
    {spmd.usedModel=0, spmd.execKind="All"}  
2 %cmpRes = arith.cmpi eq (%rank, %c0)  
3 scf.if (%cmpRes) {  
4     %error2 = spmd.bcast (%comm, %sendBuf, %recvBuf, %count, %i32Type, %c0)  
  
5 }
```

SPMD IR: Example

```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

```
1 %rank, %error1 = spmd.getRankInComm(%comm)  
    {spmd.usedModel=0, spmd.execKind="All"}  
2 %cmpRes = arith.cmpi eq (%rank, %c0) {spmd.execKind="All"}  
3 scf.if (%cmpRes) {  
4     %error2 = spmd.bcast (%comm, %sendBuf, %recvBuf, %count, %i32Type, %c0)  
  
5 }
```

SPMD IR: Example

```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

```
1 %rank, %error1 = spmd.getRankInComm(%comm)  
    {spmd.usedModel=0, spmd.execKind="All"}  
2 %cmpRes = arith.cmpi eq (%rank, %c0) {spmd.execKind="All"}  
3 scf.if (%cmpRes) {  
4     %error2 = spmd.bcast (%comm, %sendBuf, %recvBuf, %count, %i32Type, %c0)  
        {spmd.usedModel=0, spmd.isBlocking=true, spmd.executedBy=[0],  
        spmd.execKind="Static"}  
5 }
```

SPMD IR: Example

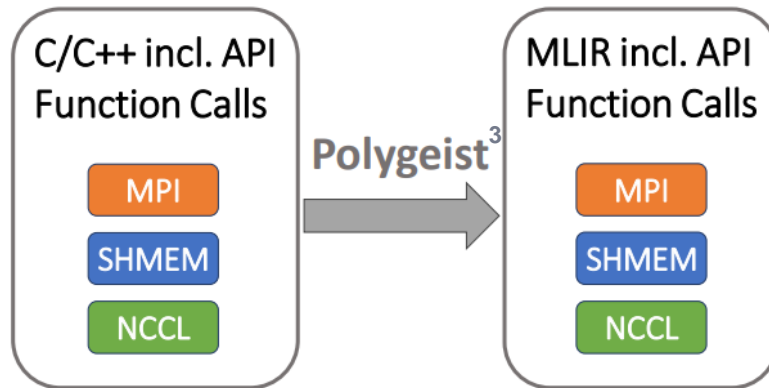
```
1 int rank;  
2 MPI_Comm_rank(MPI_COMM_WORLD, &rank);  
3 if (rank == 0) {  
4     MPI_Bcast(buf, 1, MPI_INT, 0, MPI_COMM_WORLD);  
5 }
```

```
1 int rank = shmem_my_pe();  
2 if (rank == 0) {  
3     shmem_int_broadcast(SHMEM_TEAM_WORLD, recvBuf, sendBuf, 1, 0);  
4 }
```

```
1 %rank, %error1 = spmd.getRankInComm(%comm)  
    {spmd.usedModel=0, spmd.execKind="All"}  
2 %cmpRes = arith.cmpi eq (%rank, %c0) {spmd.execKind="All"}  
3 scf.if (%cmpRes) {  
4     %error2 = spmd.bcast (%comm, %sendBuf, %recvBuf, %count, %i32Type, %c0)  
        {spmd.usedModel=0, spmd.isBlocking=true, spmd.executedBy=[0],  
        spmd.execKind="Static"}  
5 } {spmd.execKind="All", spmd.isMV=true}
```

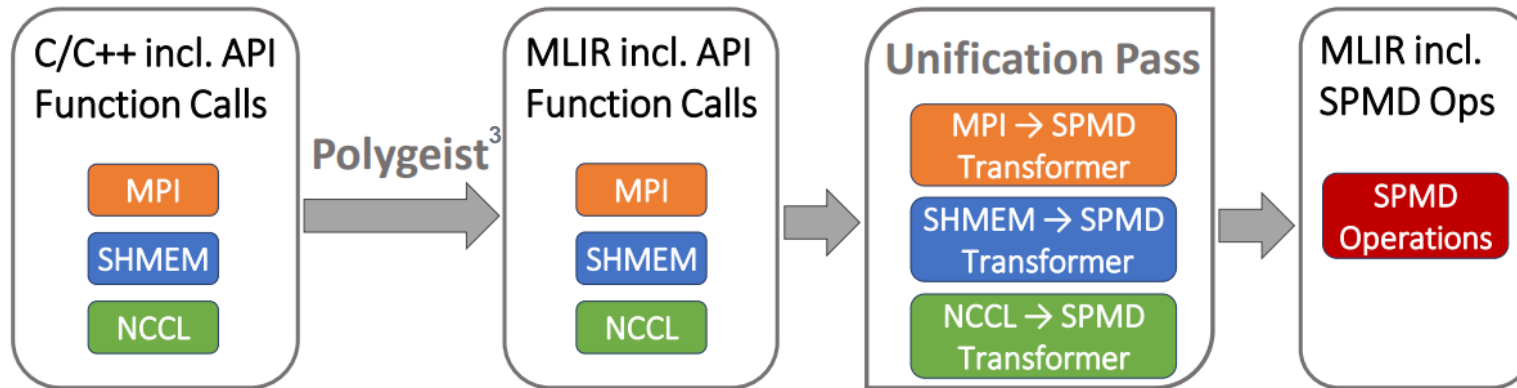
SPMD IR: Workflow

[3] Moses et al.: Polygeist: Raising C to Polyhedral MLIR. PACT '21



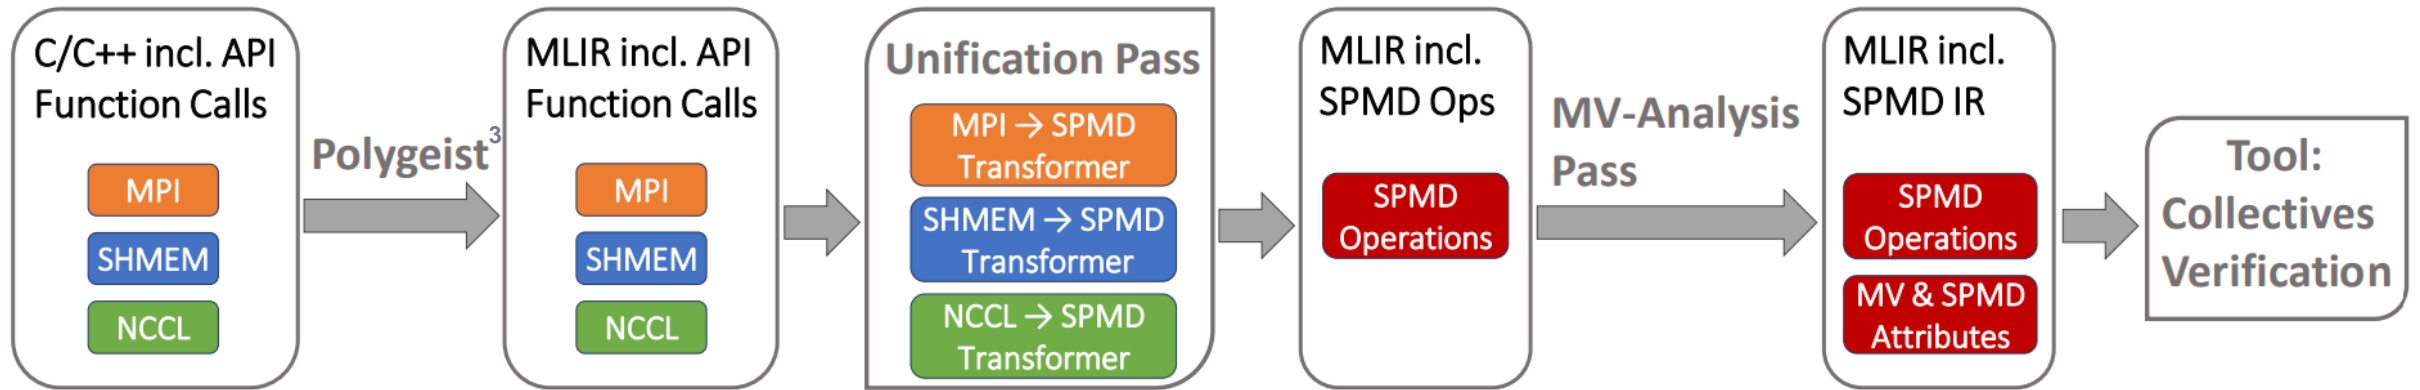
SPMD IR: Workflow

[3] Moses et al.: Polygeist: Raising C to Polyhedral MLIR. PACT '21



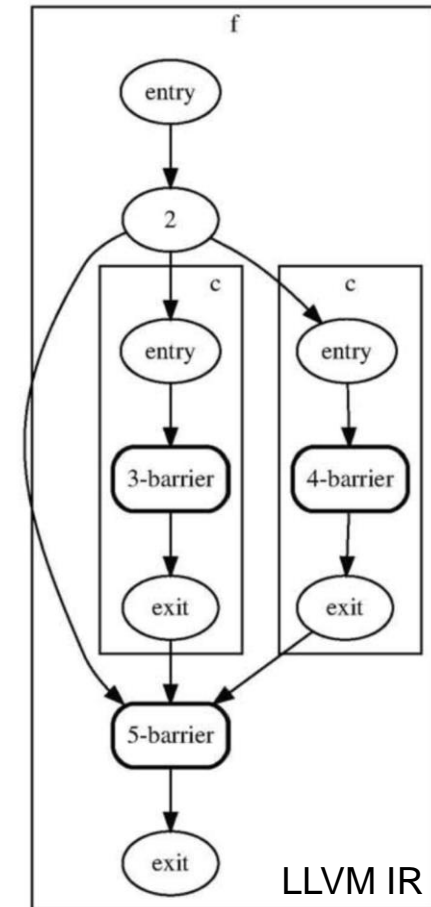
SPMD IR: Workflow

[3] Moses et al.: Polygeist: Raising C to Polyhedral MLIR. PACT '21

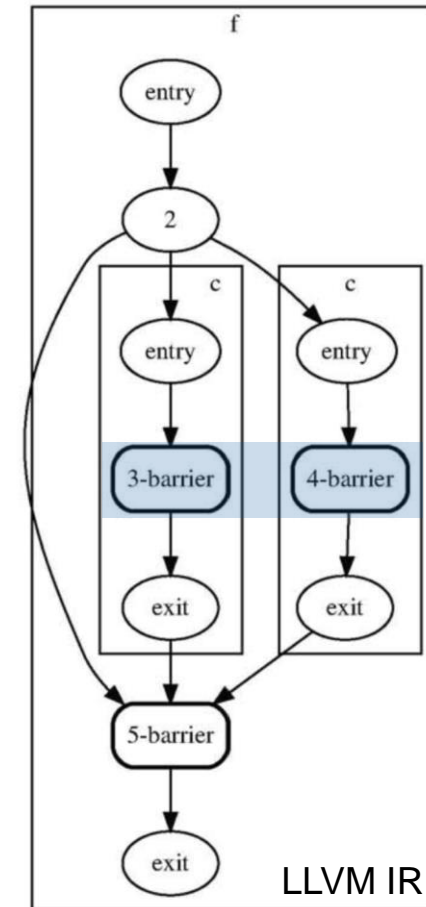


- **Use case:** static verification of collective communiation
 - The same **kind** of **collectives** has to be called **in order** by **all processes** of the **same comm.**

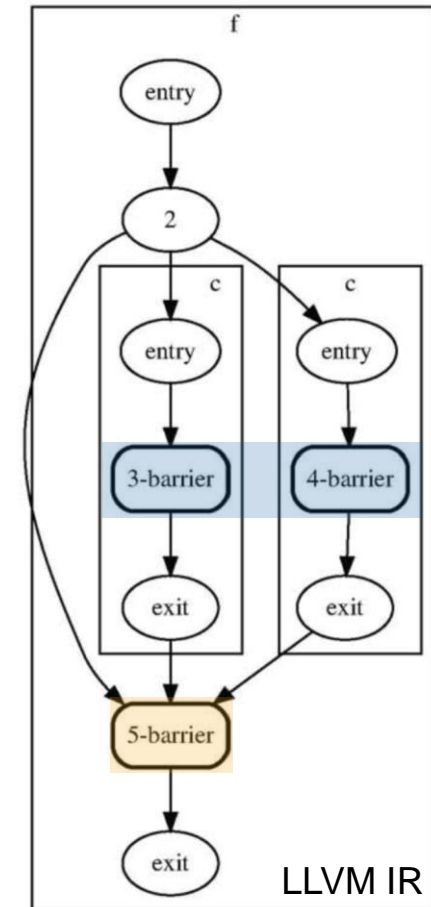
- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.




- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.



- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.




- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.




Evaluation

- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.



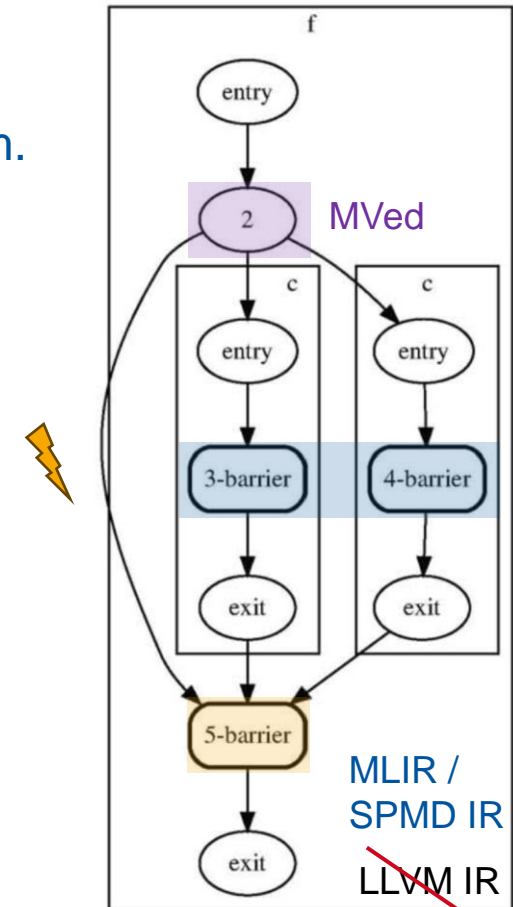
Evaluation

- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.
- Port the approach of PARCOACH⁴ to the SPMD IR with two extensions




Evaluation

- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.
- Port the approach of PARCOACH⁴ to the SPMD IR with two extensions



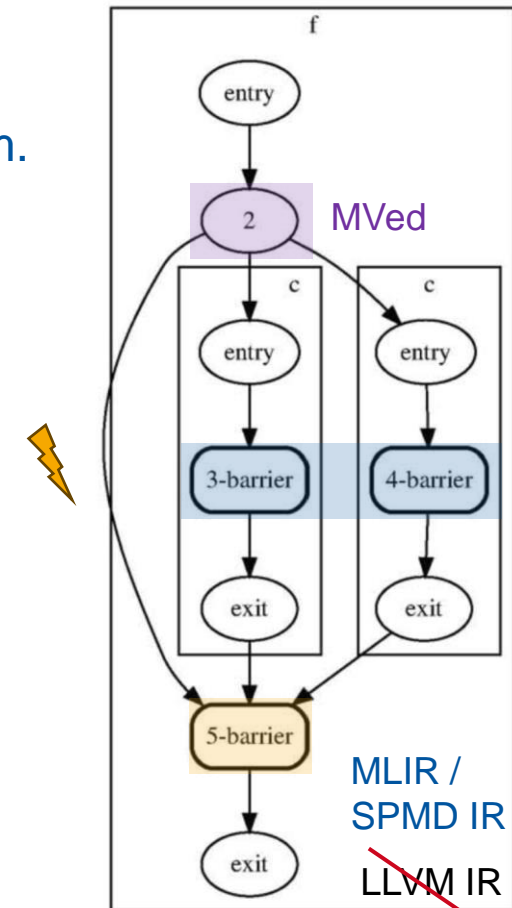
Evaluation

- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.
- Port the approach of PARCOACH⁴ to the SPMD IR with two extensions
 - Considering execution counts of collectives induced by loops




Evaluation

- Use case: static verification of collective communication
 - The same kind of collectives has to be called in order by all processes of the same comm.
- Port the approach of PARCOACH⁴ to the SPMD IR with two extensions
 - Considering execution counts of collectives induced by loops
 - Static cases where the total number and executing processes of operations are known



Evaluation


- **Use case:** static verification of collective communication
 - The same **kind of collectives** has to be called **in order** by **all processes** of the **same comm.**
- **Port** the approach of PARCOACH⁴ to the **SPMD IR** with **two extensions**
 - Considering **execution counts** of collectives induced by **loops**
 - **Static cases** where the **total number** and **executing processes** of operations are **known**
- **Assessment** based on PARCOACH's **micro-benchmark suite** and **custom test cases**
 - Port original **MPI** codes to **SHMEM** and **NCCL** and provide **hybrid** cases



Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴



Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19


Micro-Benchmark Suite⁴



Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19


Micro-Benchmark Suite⁴



Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴



S \triangleq SPMD IR Verification
P \triangleq PARCOACH

True Positive
True Negative
False Positive
False Negative
Not Supported

```
1 if (ncclRank % 2) {  
2   ncclAllGather (...);  
3 }  
4 else {  
5   ncclGroupStart ();  
6   ncclAllGather (...);  
7   ncclGroupEnd ();  
8 }
```

Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19


Micro-Benchmark Suite⁴



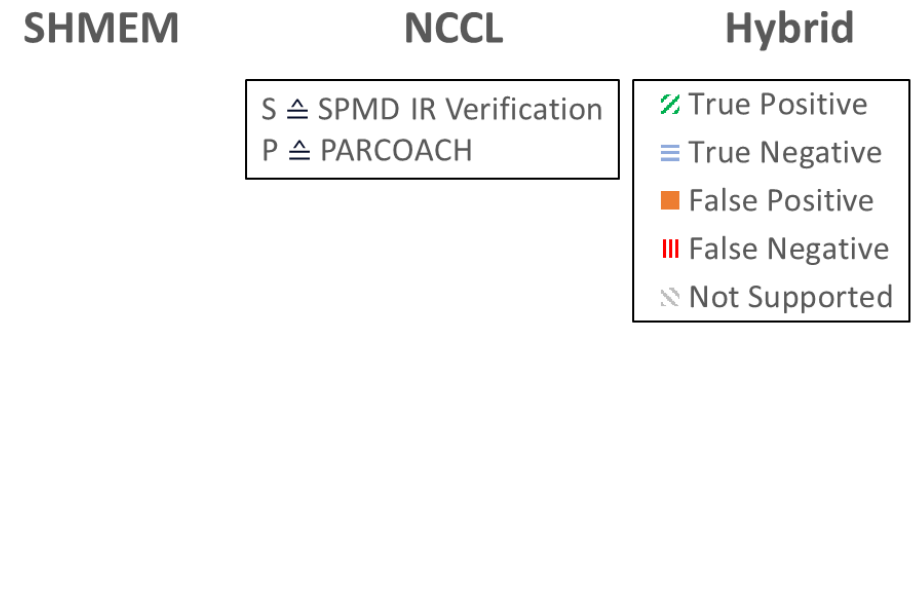
Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴




Custom Test Cases



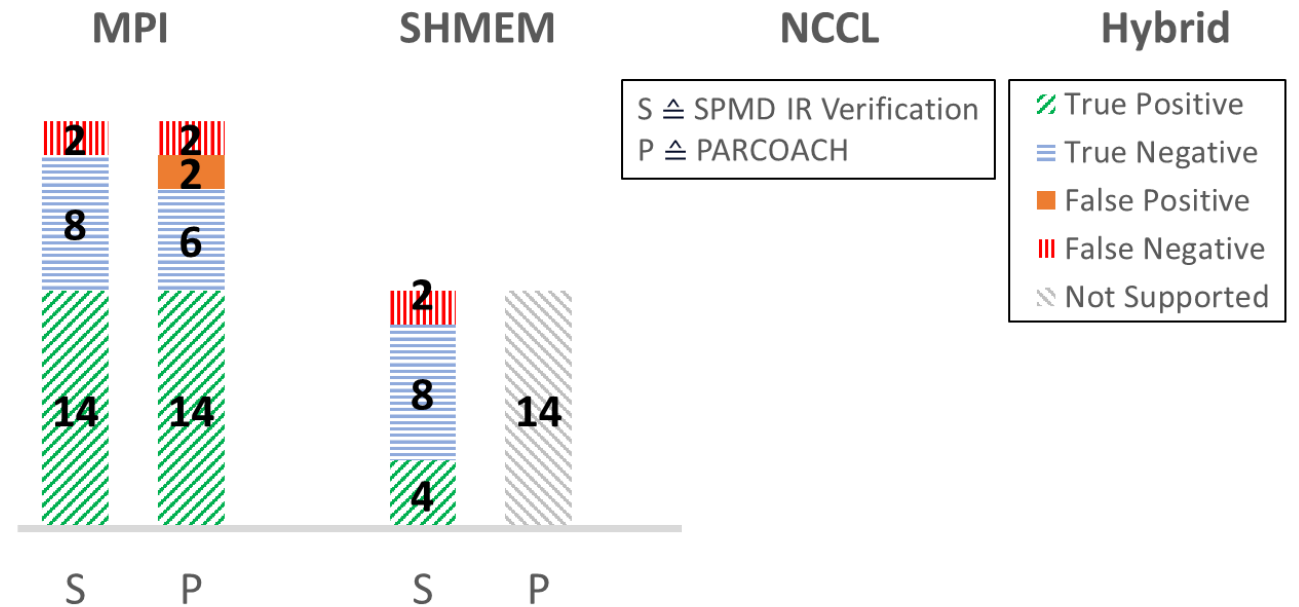
Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴




Custom Test Cases



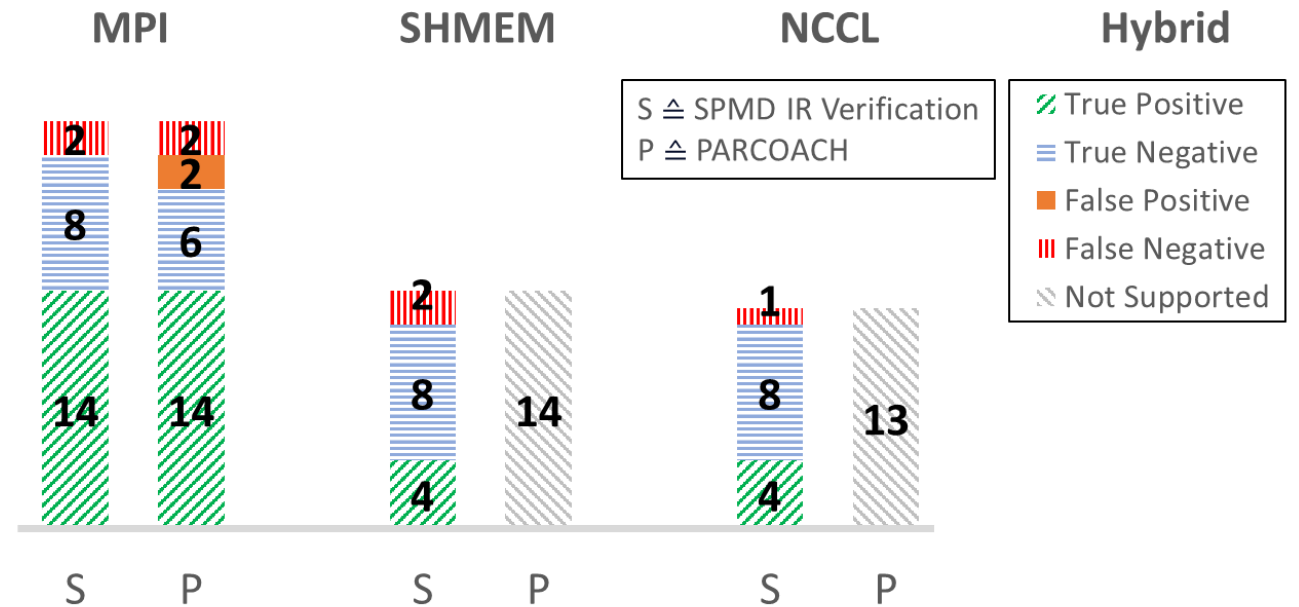
Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴




Custom Test Cases




Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19

Micro-Benchmark Suite⁴




Custom Test Cases



Evaluation: Results


[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19



- SPMD IR
 - covers **two** test cases **correctly**, where PARCOACH **fails**

Evaluation: Results


[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19



- SPMD IR
 - covers **two** test cases **correctly**, where PARCOACH **fails**
 - has one **false positive** for one **NCCL** port

Evaluation: Results

[4] MPI Codes from Huchant et al:
Multi-Valued Expression Analysis for
Collective Checking. Euro-Par '19



- SPMD IR
 - covers **two** test cases **correctly**, where PARCOACH **fails**
 - has one **false positive** for one **NCCL** port
 - besides MPI, **also supports** SHMEM, NCCL, and their **hybrid** combinations

Evaluation: Hybrid Example

```
1 int myPE = shmem_my_pe();
2 MPI_Comm_split(MPI_COMM_WORLD, 0, myPE, &shmem_comm);
3 MPI_Comm_rank(shmem_comm, &myRank);
4 assert(myPe == myRank); // Ensured by Line 2 and 3
5
6
7
8
9
10
11
12
```

Evaluation: Hybrid Example

```
1  int myPE = shmem_my_pe();
2  MPI_Comm_split(MPI_COMM_WORLD, 0, myPE, &shmem_comm);
3  MPI_Comm_rank(shmem_comm, &myRank);
4  assert(myPe == myRank); // Ensured by Line 2 and 3
5  if (myRank == 0) {
6      shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
7      MPI_Bcast(..., shmem_comm);
8  }
9  else {
10     MPI_Bcast(..., shmem_comm);
11     shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
12 }
```


Evaluation: Hybrid Example

```
1 int myPE = shmem_my_pe();
2 MPI_Comm_split(MPI_COMM_WORLD, 0, myPE, &shmem_comm);
3 MPI_Comm_rank(shmem_comm, &myRank);
4 assert(myPe == myRank); // Ensured by Line 2 and 3
5 if (myRank == 0) {
6     shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
7     MPI_Bcast(..., shmem_comm);
8 }
9 else {
10    MPI_Bcast(..., shmem_comm);
11    shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
12 }
```

- Understanding both program. models separately is **not sufficient**
- A **tool** needs to understand also the **interaction** / their combination

Evaluation: Hybrid Example

```
1 int myPE = shmem_my_pe();
2 MPI_Comm_split(MPI_COMM_WORLD, 0, myPE, &shmem_comm);
3 MPI_Comm_rank(shmem_comm, &myRank);
4 assert(myPe == myRank); // Ensured by Line 2 and 3
5 if (myRank == 0) {
6     shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
7     MPI_Bcast(..., shmem_comm);
8 }
9 else {
10    MPI_Bcast(..., shmem_comm);
11    shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
12 }
```

- Understanding both program. models separately is **not sufficient**
- A **tool** needs to understand also the **interaction** / their combination
- The conversion to and representation in the unified **SPMD IR** addresses this issue

Evaluation: Hybrid Example

```
1 int myPE = shmem_my_pe();
2 MPI_Comm_split(MPI_COMM_WORLD, 0, myPE, &shmem_comm);
3 MPI_Comm_rank(shmem_comm, &myRank);
4 assert(myPe == myRank); // Ensured by Line 2 and 3
5 if (myRank == 0) {
6     shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
7     MPI_Bcast(..., shmem_comm);
8 }
9 else {
10    MPI_Bcast(..., shmem_comm);
11    shmem_int_sum_reduce(SHMEM_TEAM_WORLD, ...);
12 }
```

```
1 %commWorld = spmd.commWorld() {spmd.usedModel=0}
2 %rank = spmd.getRankInComm(%commWorld) {spmd.usedModel=0, ...}
3 %cmpRes = arith.cmpi eq (%rank, %c0)
4 scf.if (%cmpRes) {
5     spmd.allreduce (%commWorld, ...) {spmd.usedModel=1, ...}
6     spmd.bcast (%commWorld, ...) {spmd.usedModel=0, ...}
7 } else {
8     spmd.bcast (%commWorld, ...) {spmd.usedModel=0, ...}
9     spmd.allreduce (%commWorld, ...) {spmd.usedModel=1, ...}
10 }{spmd.isMV=true, ...}
```

Discussion

- MPI-centric design
 - MPI is an **extensive** standard and programming model

Discussion

- MPI-centric design
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming

Discussion

- MPI-centric design
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly**, the **API calls** could be converted to **SPMD IR operations**

Discussion

- MPI-centric design
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly**, the **API calls** could be converted to **SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information

Discussion

- MPI-centric design
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly**, the **API calls** could be converted to **SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information
 - **Unifying PGAS** memory management calls **with non-PGAS** ones

Discussion

- **MPI-centric design**
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly, the API calls could be converted to SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information
 - **Unifying PGAS** memory management calls **with non-PGAS** ones
- **Loss of information** due to abstraction
 - E.g., multiple calls could be mapped to a single operation

Discussion

- **MPI-centric design**
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly, the API calls could be converted to SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information
 - **Unifying PGAS** memory management calls **with non-PGAS** ones
- **Loss of information due to abstraction**
 - E.g., multiple calls could be mapped to a single operation
 - Similar case led to one **false positive** in the evaluation for a NCCL port

Discussion

- **MPI-centric design**
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly, the API calls could be converted to SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information
 - **Unifying PGAS** memory management calls **with non-PGAS** ones
- **Loss of information due to abstraction**
 - E.g., multiple calls could be mapped to a single operation
 - Similar case led to one **false positive** in the evaluation for a NCCL port
 - **Could be approached** by introducing an ID that distinguishes originally separate calls
 - **but** might be **contrary** to a unifying IR


Discussion

- **MPI-centric design**
 - MPI is an **extensive** standard and programming model
 - Large **community** support
 - De facto **standard** for distributed-memory programming
- **Mostly, the API calls could be converted to SPMD IR operations**
 - **Non-blocking** communication **differences** between MPI and NCCL
 - Possibly **non-static** incorporation of information
 - **Unifying PGAS** memory management calls **with non-PGAS** ones
- **Loss of information due to abstraction**
 - E.g., multiple calls could be mapped to a single operation
 - Similar case led to one **false positive** in the evaluation for a NCCL port
 - **Could be approached** by introducing an ID that distinguishes originally separate calls
 - **but** might be **contrary** to a unifying IR
- **So far, executing kind and processes analyses limited to a few patterns for static cases**

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization



Concept

Process Management

Communicator Management


Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs



Concept

Process Management

Communicator Management


Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties



Concept

Process Management

Communicator Management


Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)



Concept

Process Management

Communicator Management


Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI




Concept
Process Management
Communicator Management
Data Management Collective Comm.
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)

Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI
- **Extended** micro-benchmark suite and provide **ports**



Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)


Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI
- **Extended** micro-benchmark suite and provide **ports**

Future Work:

- **Extend** the **execution kind** and **executing processes** analysis, explore **new use cases**



Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)


Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI
- **Extended** micro-benchmark suite and provide **ports**

Future Work:

- **Extend** the **execution kind** and **executing processes** analysis, explore **new use cases**
- **Apply** the SPMD IR to **other program. models** such as NVSHMEM, GASPI, and UPC++



Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)


Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI
- **Extended** micro-benchmark suite and provide **ports**

Future Work:

- **Extend** the execution kind and executing processes analysis, explore **new use cases**
- **Apply** the SPMD IR to other program. models such as NVSHMEM, GASPI, and UPC++
- **Integrate** other SPMD features / paradigms such as **RMA** or **OpenMP**



Concept

Process Management

Communicator Management

Data Management
Collective Comm.

Point-To-Point Comm.
Non-Blocking Semantics
(P2P and Collectives)


Conclusion

Summary:

- Established unifying and abstracting **SPMD IR**
 - Allowing programming-model-**independent** analysis and optimization
 - For **SPMD-like** programming models, incl. distributed-memory, PGAS, and GPU programs
- Introduced a **multi-value IR** and extended analysis for important SPMD properties
- Evaluated on **collectives verification**, **outperforming** PARCOACH (programming model support, hybrid codes, two extensions)
 - **First** collectives verification for NCCL and SHMEM, and their hybrid combinations with MPI
- **Extended** micro-benchmark suite and provide **ports**

Future Work:

- **Extend** the execution kind and executing processes analysis, explore **new use cases**
- **Apply** the SPMD IR to other program. models such as NVSHMEM, GASPI, and UPC++
- **Integrate** other SPMD features / paradigms such as **RMA** or **OpenMP**
- **Assess** approach and workflow on more realistic codes such as **proxy apps**



Concept
Process Management
Communicator Management
Data Management Collective Comm.
Point-To-Point Comm. Non-Blocking Semantics (P2P and Collectives)
