

# Marawacc: A Framework for Heterogeneous Computing in Java

Juan Fumero, Michel Steuer, Christophe Dubach



The University of Edinburgh

UK Many-Core Developer Conference 2016



Motivation

Marawacc-API

Runtime Code  
Generation

Runtime  
Management

Results

Conclusion

Marawacc

## Heterogeneous Hardware



Motivation

Marawacc-API

Runtime Code  
Generation

Runtime  
Management

Results

Conclusion

## Hardware

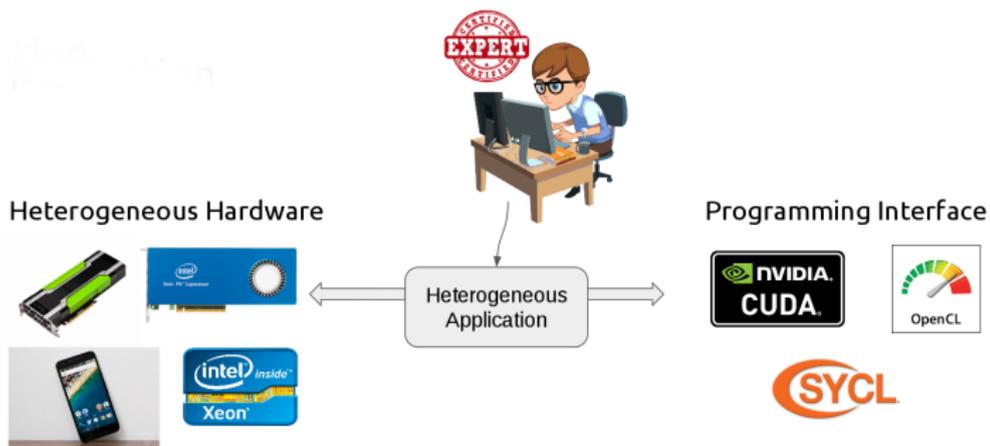
### Heterogeneous Hardware



### Programming Interface



# Motivation



Motivation

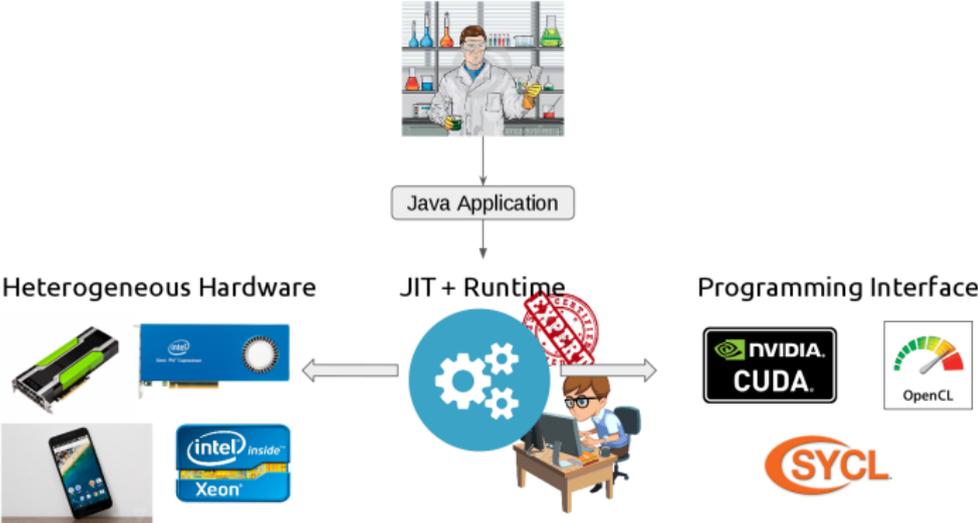
Marawacc-API

Runtime Code  
GenerationRuntime  
Management

Results

Conclusion

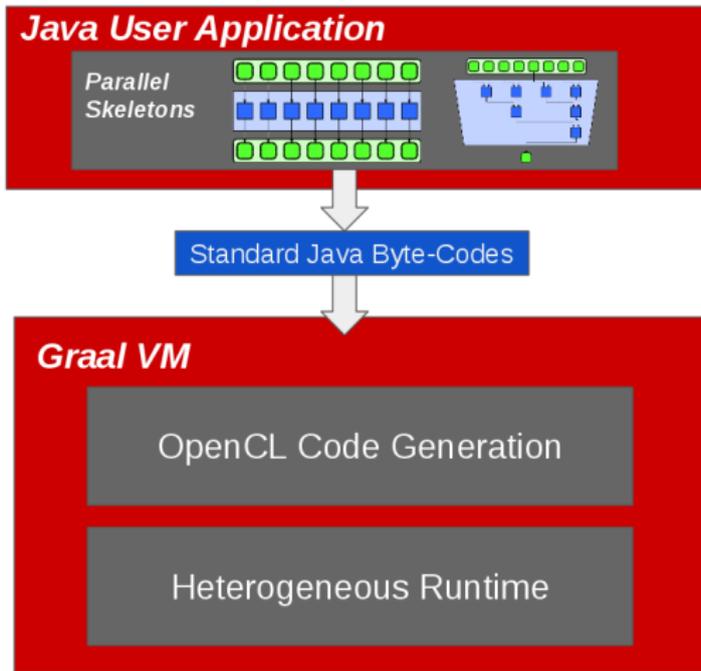
# Motivation



- Motivation
- Marawacc-API
- Runtime Code Generation
- Runtime Management
- Results
- Conclusion

# Marawacc: our approach

Three levels of abstraction



Motivation

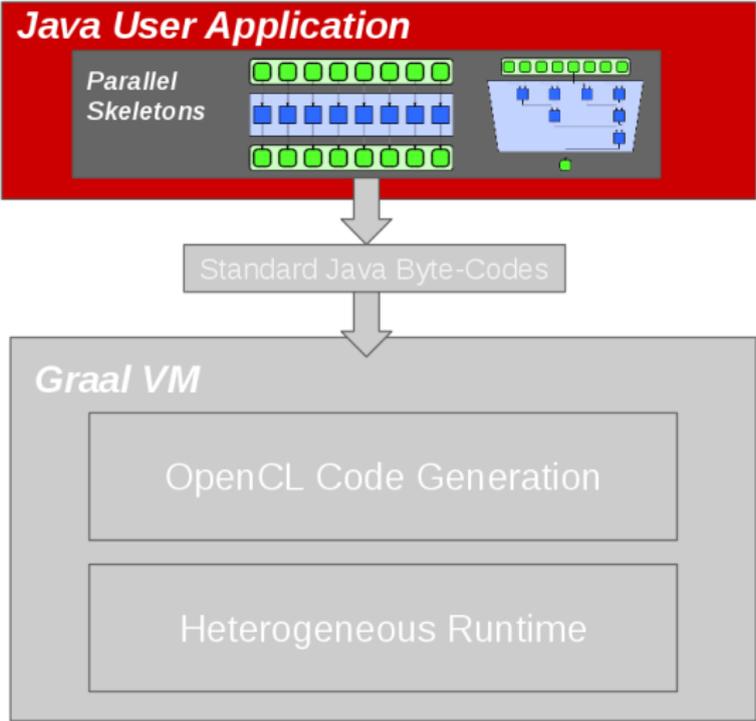
Marawacc-API

Runtime Code  
Generation

Runtime  
Management

Results

Conclusion



- Motivation
- Marawacc-API
- Runtime Code Generation
- Runtime Management
- Results
- Conclusion

# Example: Saxpy in Java

[Motivation](#)[Marawacc-API](#)[Runtime Code  
Generation](#)[Runtime  
Management](#)[Results](#)[Conclusion](#)

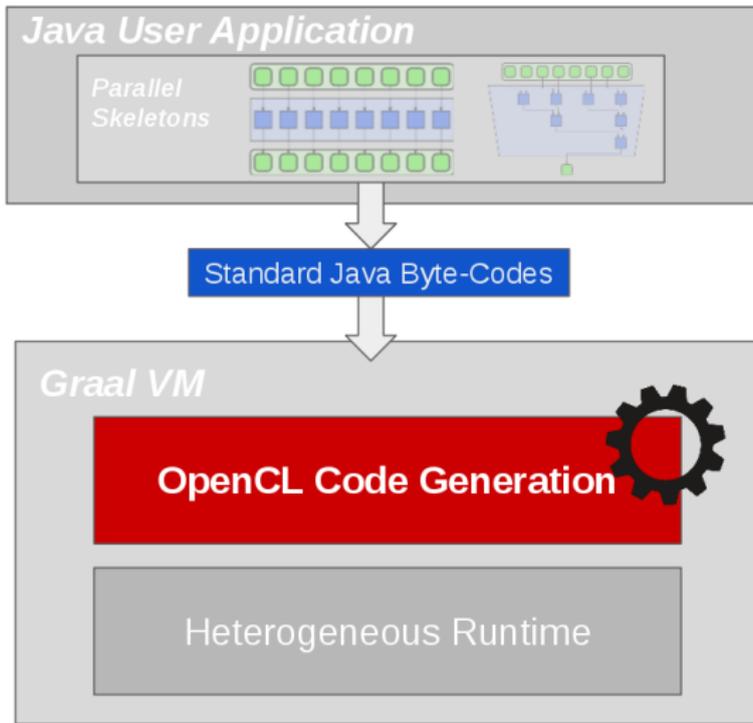
```
1 float [] v1 = new float [size];
2 float [] v2 = new float [size];
3 float [] result = new float [size];
4
5 for (int i = 0; i < size; i++) {
6     result[i] = alpha * v1[i] + v2[i];
7 }
```

# Example: Saxpy in Java

[Motivation](#)[Marawacc-API](#)[Runtime Code  
Generation](#)[Runtime  
Management](#)[Results](#)[Conclusion](#)

```
1 Float [] v1 = new Float[size];
2 Float [] v2 = new Float[size];
3
4 ArrayFunc<Tuple2<Float , Float>, Float> f;
5 f = new MapFunction<>(t -> alpha * t._1() + t._2());
6
7 Float [] result = f.zip(v1, v2).apply();
```

# Runtime Code Generation



Motivation

Marawacc-API

Runtime Code Generation

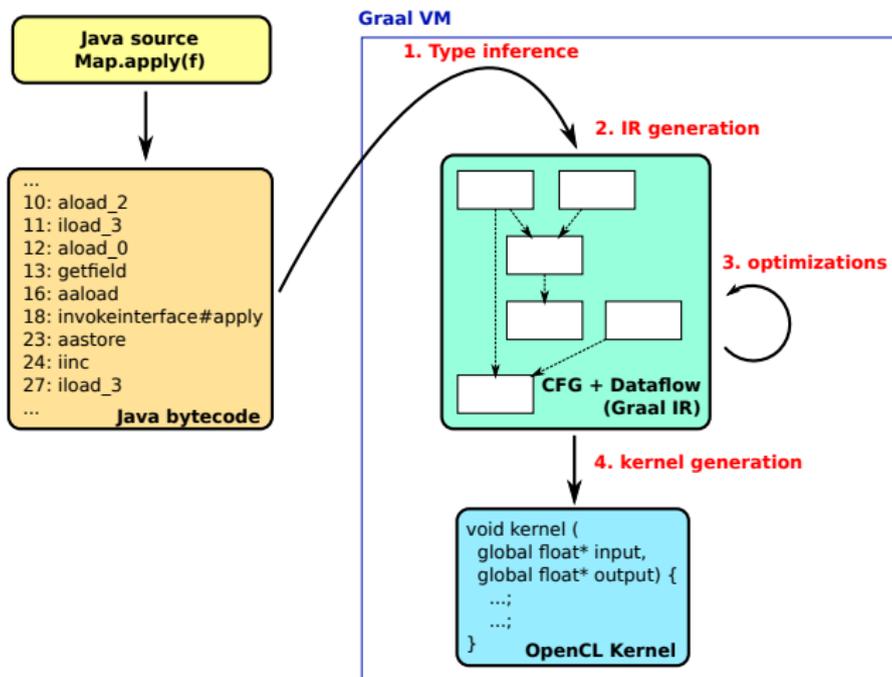
Runtime Management

Results

Conclusion

# Runtime Code Generation

## Workflow



Motivation

Marawacc-API

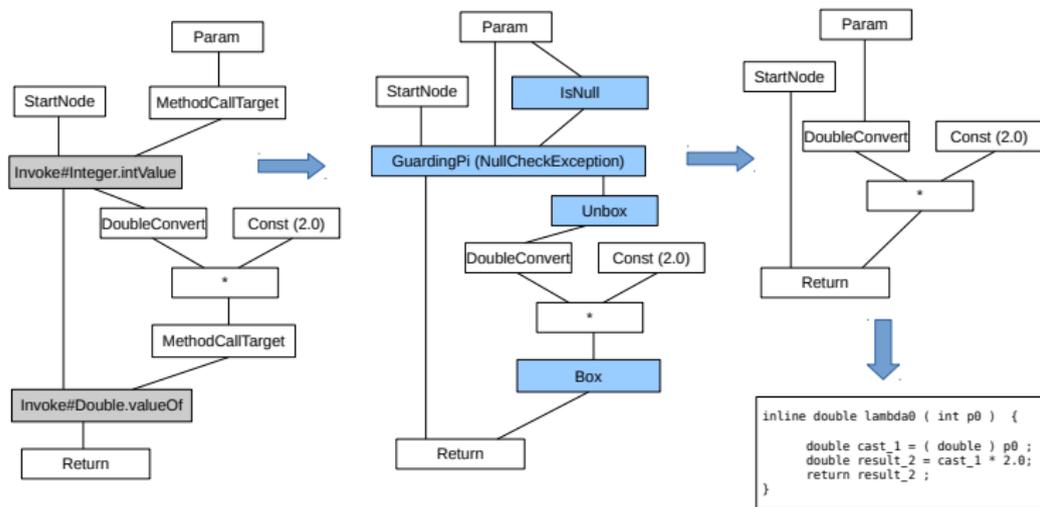
Runtime Code  
GenerationRuntime  
Management

Results

Conclusion

# Runtime Code Generation

MapFunction < Integer, Double > (x -> x \* 2.0)



Motivation

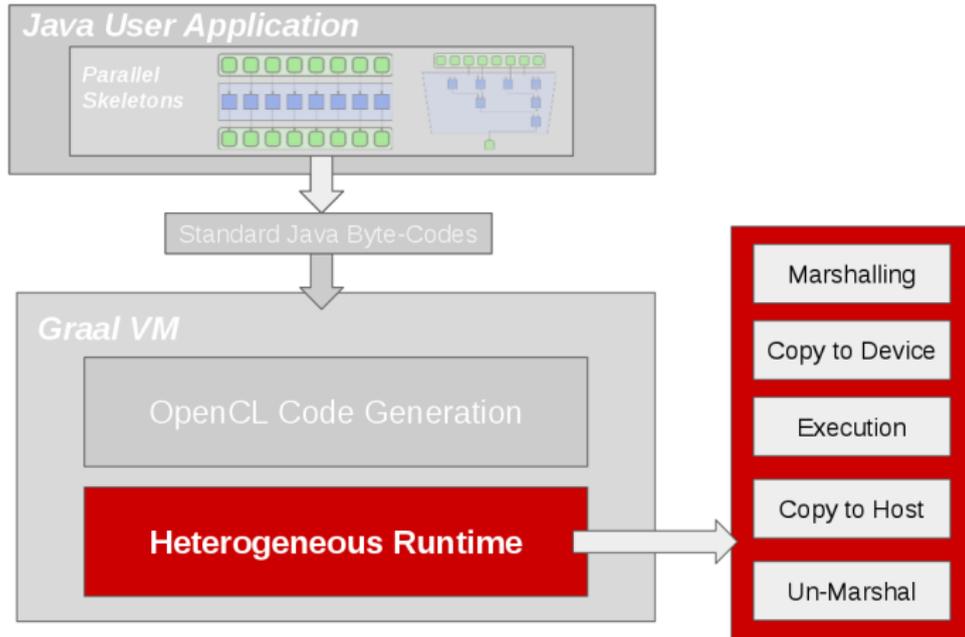
Marawacc-API

Runtime Code  
GenerationRuntime  
Management

Results

Conclusion

# Marawacc: Runtime Management

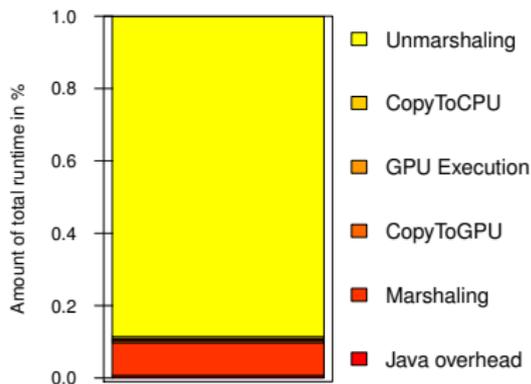


- Motivation
- Marawacc-API
- Runtime Code Generation
- Runtime Management
- Results
- Conclusion

# Where the time is spent?

Black-scholes benchmark.

$Float[] \implies Tuple2 < Float, Float > []$



- ▶ Un/marshal data takes up to 90% of the time
- ▶ Computation step should be dominant

This is not acceptable. Can we do better?

Motivation

Marawacc-API

Runtime Code Generation

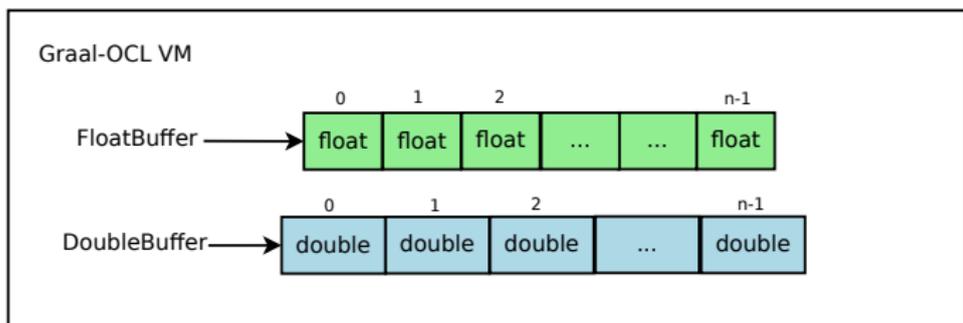
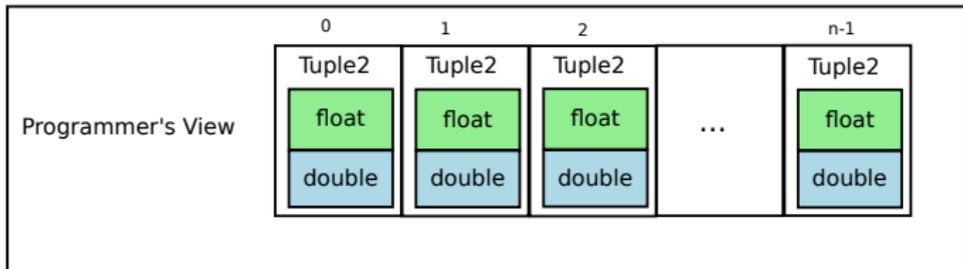
Runtime Management

Results

Conclusion

# Custom Array Type: PArray

PArray<Tuple2<Float,Double>>



With this layout, un/marshal operations are not necessary

Motivation

Marawacc-API

Runtime Code  
GenerationRuntime  
Management

Results

Conclusion

# Sapy example

[Motivation](#)[Marawacc-API](#)[Runtime Code  
Generation](#)[Runtime  
Management](#)[Results](#)[Conclusion](#)

```
1 Float [] v1 = new Float[size];
2 Double [] v2 = new Double[size];
3
4 ArrayFunc<Tuple2<Float , Double>, Double> f;
5 f = new MapFunction<>(t -> alpha * t._1() + t._2());
6
7 Float [] result = f.zip(v1, v2).apply();
```

# Saxpy with our Custom PArrays

[Motivation](#)[Marawacc-API](#)[Runtime Code  
Generation](#)[Runtime  
Management](#)[Results](#)[Conclusion](#)

```
1 Float[] v1 = new Float[size];
2 Double[] v2 = new Double[size];
3 PArray input= new PArray(v1, v2);
4
5 ArrayFunc<Tuple2<Float, Double>, Double> f;
6 f = new MapFunction<>(t -> alpha * t._1() + t._2());
7 PArray<Double> output = f.apply(input);
```

[Motivation](#)[Marawacc-API](#)[Runtime Code  
Generation](#)[Runtime  
Management](#)[Results](#)[Conclusion](#)

# Results

# OpenCL GPU Execution

## AMD R9 and NVIDIA GeForce GTX Titan

Motivation

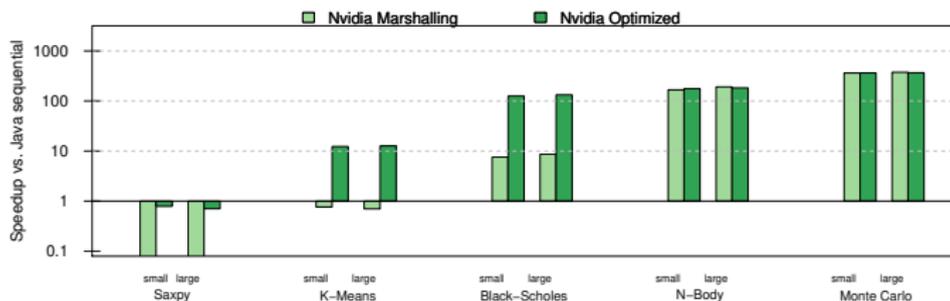
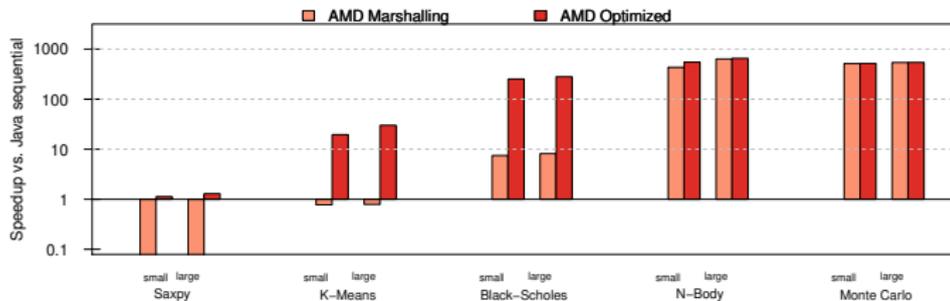
Marawacc-API

Runtime Code Generation

Runtime Management

Results

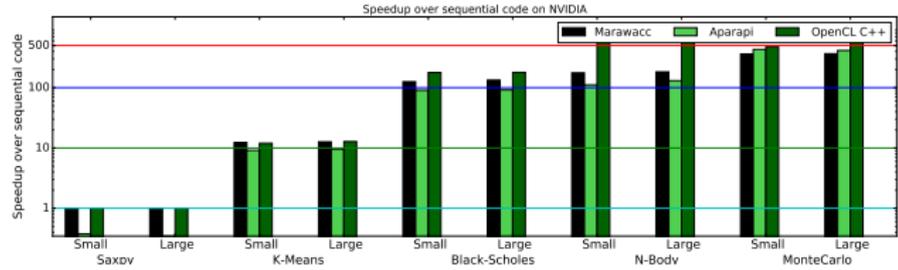
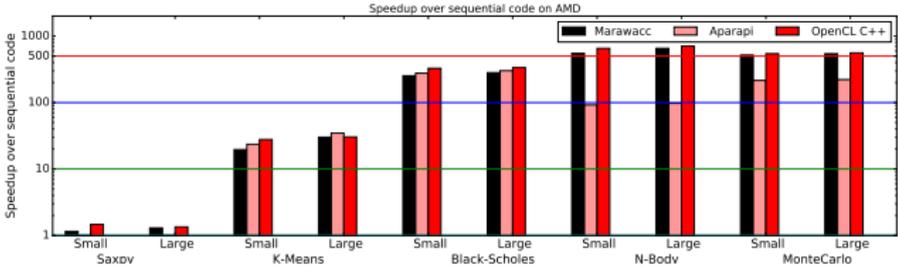
Conclusion



# Comparison with OpenCL C++

AMD R9 and NVIDIA GeForce GTX Titan

- Motivation
- Marawacc-API
- Runtime Code Generation
- Runtime Management
- Results
- Conclusion



# .zip(Conclusions).map(Future)

## Present

- ▶ We have presented Marawacc framework for programming GPUs from Java
- ▶ Custom array type to reduce overheads when transforming the data
- ▶ Runtime system to run heterogeneous applications within Java

## Future

- ▶ Code generation for multiple devices
- ▶ Runtime scheduling (Where is the best place to run the code?)

Motivation

Marawacc-API

Runtime Code  
Generation

Runtime  
Management

Results

Conclusion

# Thanks so much for your attention

Marawacc



Juan Fumero

<juan.fumero@ed.ac.uk>

Motivation

Marawacc-API

Runtime Code  
Generation

Runtime  
Management

Results

Conclusion

# OpenCL code generated

```

1  double lambda0(float p0) {
2      double cast_1 = (double) p0;
3      double result_2 = cast_1 * 2.0;
4      return result_2;
5  }
6  kernel void lambdaComputationKernel (
7      global float * p0,
8      global int *p0_index_data ,
9      global double *p1,
10     global int *p1_index_data) {
11     int p0_dim_1 = 0; int p1_dim_1 = 0;
12     int gs = get_global_size(0);
13     int loop_1 = get_global_id(0);
14     for ( ; ; loop_1 += gs) {
15         int p0_len_dim_1 = p0_index_data[p0_dim_1];
16         bool cond_2 = loop_1 < p0_len_dim_1;
17         if (cond_2) {
18             float auxVar0 = p0[loop_1];
19             double res = lambda0(auxVar0);
20             p1[p1_index_data[p1_dim_1] + loop_1]
21                 = res;
22         } else { break; }
23     }
24 }

```

Motivation

Marawacc-API

Runtime Code  
GenerationRuntime  
Management

Results

Conclusion