



PHI

A Modern C++ Library for Parallel Pattern Composition

Santiago Veigas Ramírez

 sveigas@inf.uc3m.es

Daniel Martínez Davies
J. Daniel García Sánchez

TABLE OF CONTENTS

01 Introduction

02 Background

- C++ 20 Ranges
- Parallel Patterns
- Syntax

03 Design

- Interface
- Range Adaptors & RACOs
- Parallel Patterns

04 Evaluation

05 Conclusion

- Programming languages evolve, so does the way to express code.
- Pattern-based programming improves readability and maintainability.
- Ranges enable more expressive, composable, and safer code.
- Transition from iterators to ranges.
- Parallel code can benefit from similar abstractions.

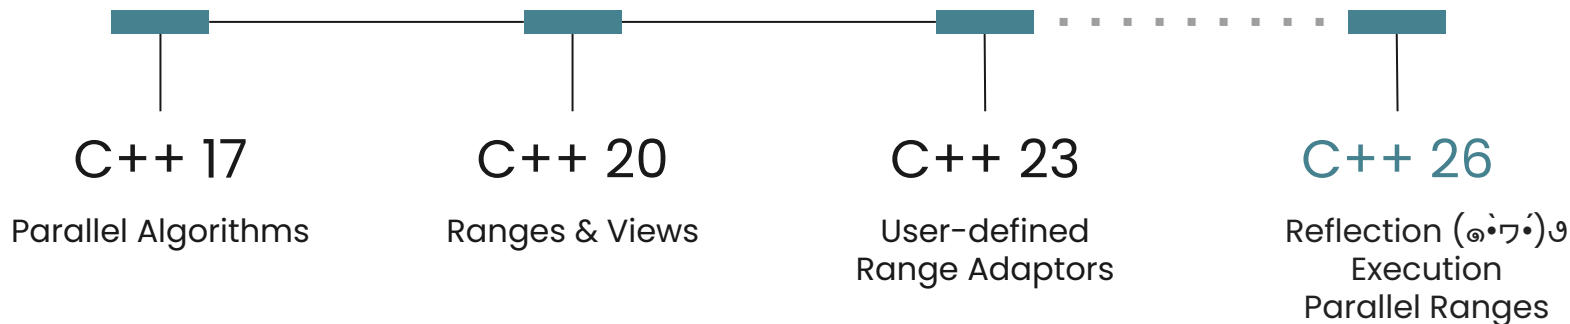


TABLE OF CONTENTS

01 Introduction

02 Background

- C++ 20 Ranges
- Parallel Patterns
- Syntax

03 Design

- Interface
- Range Adaptors & RACOs
- Parallel Patterns

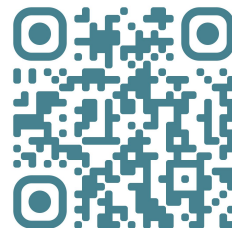
04 Evaluation

05 Conclusion

C++ 20 Ranges | Background

- High-level abstraction that supersedes iterators.
 - Introduces the concepts of ranges & views.
- Replaces the use of iterator-pairs with a pipe-like syntax.
- Generally lazily evaluated.

```
0 // View Composition
1 auto && squares = std::views::iota(1)                // Infinite Range [1 ... N]
2 | std::views::transform([] (int x) { return x * x; }) // Square elements
3 | std::views::filter([] (int x) { return x > 100; })  // Keep elements > 100
4 | std::views::take(10)                               // Keep 10 elements
5 ;
6 // View Evaluation (Print Elements)
7 for (auto element: squares) std::print("{} ", element);
8 // Another Evaluation (Print Average)
9 std::print("[{}]", std::ranges::fold_left(squares, 0.0, std::plus{}) / 10.0);
```



<https://godbolt.org/z/ehvIEfsze>

- Programming patterns are widely recognized as a best practice.
- They define a clear mapping between input and output data, using a transformation function to express the computation to be done in parallel
- Parallel control patterns (data parallel patterns)
 - Map, reduce, stencil
- Parallel data management patterns
 - Pipeline, pack/unpack, scatter/gather

- Commonly found in OpenMP.
- Primarily based on pragma directives.
- Requires minimal code changes.
- Limited type support.
- Obscure error messages.

Macros

```
0 void openmp(float a[], float b[], int n) {  
1   int i, j;  
2   #pragma omp parallel shared(a,b,n)  
3   {  
4     #pragma omp for schedule(dynamic,1) private (i,j) nowait  
5     for (i = 1; i < n; i++)  
6         for (j = 0; j < i; j++)  
7             b[j + n*i] = (a[j + n*i] + a[j + n*(i-1)]) / 2.0;  
8   }  
9 }
```

- No method chaining.
- Often stateless and reusable.
- Could be composable but not "fluent".
- Most common among parallel frameworks and libraries.

Free Functions

```
0 // Intel TBB
1 tbb::parallel_for(0, static_cast<int>(v1.size()), [&](int i) {
2     v1[i] = i * i;
3 });
```

```
0 // C++17 Parallel Algorithms
1 std::for_each(std::execution::parallel_policy{,
2     v1.begin(), v1.end(),
3     [] (int &x) { x = x*2; }
4 );
```

```
0 // GrPPI
1 grppi::map(grppi:: parallel_execution_native{, v1, v1,
2     [](int &x) { x = x * 2; }
3 );
```

- Declarative, functional style.
- Chainable operations.
- Returns *this or new object for chaining.
- Similar to range-like pipes.

Fluent / Chaining Style

```
0 // RxCPP
1 auto values = rxcpp::observable<>::range(1, 10)
2   .map([](int v) { return v * 2; })
3   .filter([](int v) { return v % 2 == 0; })
4   .reduce(0, [](int acc, int v) { return acc + v; });
```

- Heterogenous frameworks
 - SYCL
 - CUDA
 - OpenCL
- Work-items and work-groups
- Manual index calculation
- Direct memory access

Kernel

```
0 // SYCL 2020
1 q.submit([&](sycl::handler& h) {
2     auto in = in_buf.get_access<sycl::access::mode::read>(h);
3     auto out = out_buf.get_access<sycl::access::mode::write>(h);
4     h.parallel_for(sycl::range<1>{N}, [=](sycl::id<1> i) {
5         out[i] = in[i] * in[i];
6     });
7 });
```

```
0 // CUDA
1 __global__ void square_kernel(int* input, int* output, int N) {
2     int i = blockIdx.x * blockDim.x + threadIdx.x;
3     if (i < N)
4         output[i] = input[i] * input[i];
5 }
```

TABLE OF CONTENTS

01 Introduction

02 Background

- C++ 20 Ranges
- Parallel Patterns
- Syntax

03 Design

- Interface
- Range Adaptors
- Parallel Patterns

04 Evaluation

05 Conclusion



- PHI builds on the foundations and concepts of GrPPI.
- PHI's core abstractions:
 - Range adaptors.
 - Data-parallel patterns.
 - Execution backends.
- Pattern composition & execution should be decoupled.
- Parallel Patterns are modeled after Range Adaptors.
- Avoid explicitly composed patterns like map-reduce.
- Type constraints.



<https://github.com/arcosuc3m/grppi>

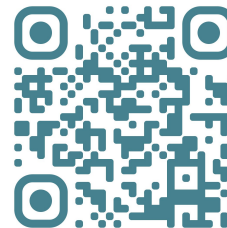
- Range Adaptors are Customization Point Objects (CPOs) that transform a range into another range-like object.
- They create deeply nested types.

```
0 // Using pipe syntax
1 auto && view = range | adaptor_1 | adaptor_2;
2 // Equivalent without RACOs
3 auto && view = adaptor_2(adaptor_1(range));
```

```
0 // std::ranges::views::__adaptor::_Partial<std::views::_Transform, lambda(int)>
1 auto && T = std::views::transform([] (int x) {return x;});
2
3 // std::ranges::transform_view<std::ranges::iota_view<int, std::unreachable_sentinel_t>, lambda(int)>
4 auto && I = std::views::iota(1) | T;
```

- RACOs are unary function objects that enables the piping mechanism.
 - When applied to a range: the output is a view.
 - When applied to a RACo: the output is a closure object.

```
0 struct identity_fn : std::ranges::range_adaptor_closure<identity_fn> {  
1     template<std::ranges::viewable_range Rng>  
2     constexpr auto operator()(Rng && range) const {  
3         return std::forward<Rng>(range);  
4     }  
5 };  
6 inline constexpr auto identity = identity_fn{};  
7  
8 // Usage Example  
9 int main() {  
10     for (auto x : std::vector{1,2,3,4} | identity)  
11         std::print("{} ", x); // Prints 1 2 3 4  
12 }
```



<https://godbolt.org/z/835GhMaE8>

- Modeled after Unary functions.
- N-ary interface substituted by zip views over data.
- Unary Transformer: **U res = op(x)**
- Transformer function must be pure.

```
0 // Illustrative Interface - Unary Map Transformer
1 auto transformer = [ ] (T element) → U
2     { return /* Transformation Op */ ; };
3
4 // Usage Example
5 auto pipeline = input_range<T>
6     | phi::map(transformer)
7     | phi::map([ ] (U element) { return 42 * element; })
8     ;
```

- Modeled after a combiner function & identity value.
- Combiner: **$T\ res = cmb(T\ x, U\ y)$** .
- Combiner function should be pure & associative.
- Special type of adaptor \rightarrow "Terminal Operation / Range"

```
0 // Illustrative Interface - Reduce Combiner
1 auto combiner = [ ] (T element_1, U element_2)  $\rightarrow$  T
2     { return /* Combination Op */; };
3
4 // Usage Example
5 auto pipeline = input_range<T>
6     | phi::map([ ] (T e)  $\rightarrow$  T { return e * e; } )
7     | phi::reduce([ ] (T e1, U e2) { return e1 + e2; } , 0) // Identity Value
8 // | phi::reduce(std::plus{}, 0)                          // Equivalent Expression
9 ;
```

- Performance impact & implementation challenges
 - Irregularly structured data in memory.
 - Data that cannot be partitioned into parallel units.
- C++ Proposal P3179 authors' identify random access iterators as a common requirement across many existing parallel execution models.
- Ranges must be bounded in size.
 - `std::ranges::iota_view` produces unbounded sequences.

TABLE OF CONTENTS

01 Introduction

02 Background

- C++ 20 Ranges
- Parallel Patterns
- Syntax

03 Design

- Interface
- Range Adaptors & RACOs
- Parallel Patterns

04 Evaluation

05 Conclusion



- Embarrassingly parallel escape-time algorithm.
- Well-suited for map-pattern decomposition.
- The algorithm has these stages:
 - Initialization of a grid of data points.
 - Mapping each grid index to a corresponding coordinate in the complex plane.
 - Running the escape-time iteration for each complex point.

Mandelbrot | Evaluation

Traditional Algorithm

- All stages of the computation are embedded directly in the inner loop.
- Executed eagerly at each iteration
- Algorithm and implementation details are coupled.
- More complex algorithms could introduce errors with indexed array accesses.

```
0 // Precompute the change in real and imaginary components (x_step, y_step)
1 double const x_step = (plane.x_max - plane.x_min) / image_width;
2 double const y_step = (plane.y_max - plane.y_min) / image_height;
3
4 // Initialize the matrix (std::vector<unsigned>) from 0 to N
5 std::ranges::iota(matrix, 0);
6
7 // Execution loop
8 for (unsigned const n : matrix) {
9     // Convert 1D index to 2D index (i, j)
10    unsigned x = n % matrix.width;
11    unsigned y = n / matrix.width;
12
13    // Map index to complex plane coordinates (x, y)
14    double real = plane.x_min + (x * x_step);
15    double imag = plane.y_min + (y * y_step);
16
17    // Compute escape time for the complex point
18    std::complex<double> c(real, imag);
19    std::complex<double> z(0.0, 0.0);
20    unsigned iterations = 0;
21    while (std::norm(z) ≤ escape_value * 2 && iterations < max_iterations) {
22        z = std::pow(z, 2) + c;
23        ++iterations;
24    }
25    output[n] = iterations; // Assume that output is correctly sized
26 }
```

```

0 // Precompute the change in real and imaginary components (x_step, y_step)
1 // x_min, x_max, y_min, and y_max represent a window in the complex plane
2 double const x_step{(plane.x_max - plane.x_min) / image_width};
3 double const y_step{(plane.y_max - plane.y_min) / image_height};
4
5 // Initialize the matrix (std::vector<unsigned>) from 0 to N
6 std::ranges::iota(matrix, 0);
7
8 // Transform iota value (n) to an idx value (i, j) → 1D index to 2D index
9 auto iota_to_idx = [&] (unsigned const value) → coordinates<unsigned> {
10     return {x = value % matrix.width, .y = value / matrix.width};
11 };
12
13 // Transform the idx pair to a coordinate pair belonging to the Mandelbrot set
14 auto idx_to_xy = [=] (coordinates<unsigned> const & ij)
15     → coordinates<double> {
16     return {x = x_min + (ij.x * x_step), .y = y_min + (ij.y * y_step)};
17 };
18
19 // Compute escape time for the complex point
20 auto xy_to_escape = [=] (coordinates<double> const & xy) → unsigned {
21     unsigned iterations{0};
22     std::complex const c{xy.x, xy.y};
23     std::complex<double> z{0, 0};
24     while (std::norm(z) ≤ escape_value * 2 && iterations < max_iterations) {
25         z = pow(z, 2) + c;
26         ++iterations;
27     }
28     return iterations;
29 };

```

Mandelbrot | Evaluation

Lambda Algorithmic Skeleton

- Loop has been replaced with lambdas representing data transformations.
 - A single lambda would've sufficed.
- Less intuitive than the iterative version, but easier to compose, reuse, and abstract.
- None of the lambdas are eagerly executed.

- Each lambda can now be composed into declarative pipelines.
- Lambdas can be re-used with minimal assumptions about the underlying data range.
- The pipeline can be executed iteratively as a range-loop; or passed to an offloading backend.

```
0 // Pattern composition
1 auto && pipeline =
2     matrix
3     | phi::map(iota_to_idx)
4     | phi::map(idx_to_xy)
5     | phi::map(xy_to_escape)
6     ;
7
8 // Pipeline execution
9 phi::execute(pipeline, output);
```

Mandelbrot | Evaluation

Machine Configuration

CPU (X86_64) AMD Ryzen 9 5950X

Cores 16 (2 threads per core)

Base - Max Frequency 2200 MHz - 5083 MHz

L1 Cache 512 KiB + 512 KiB

L2 Cache 8 MiB

L3 Cache 64 MiB (2x)

Compilation Configuration

Compiler GCC 14.2.0

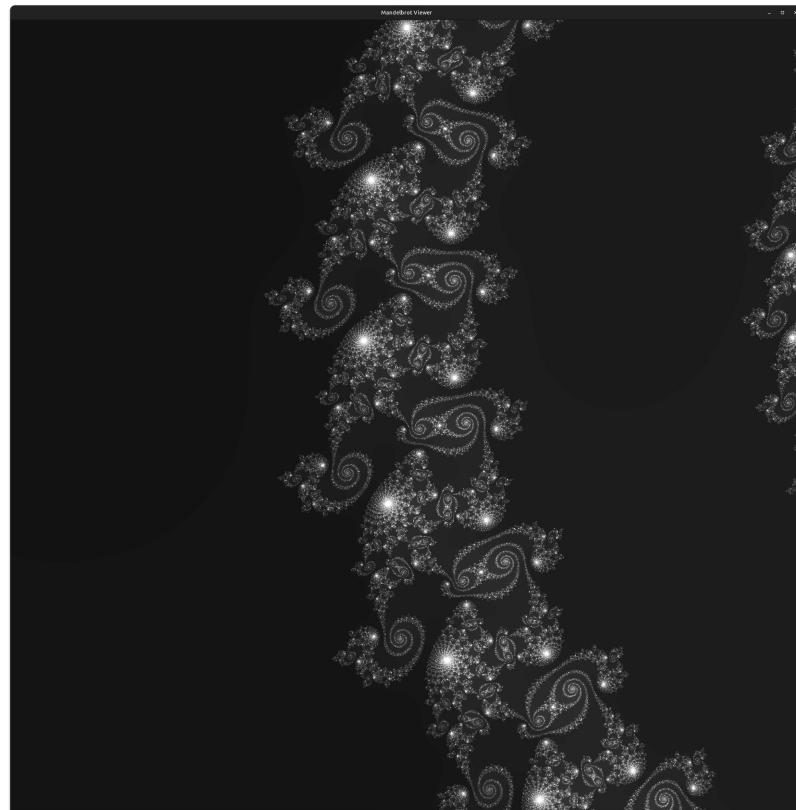
Flags -O3 -DNDEBUG

Mandelbrot Configuration

Position -0.761574 -0.0847596 3125

Iterations 10,000

Escape Value 2.0



Mandelbrot | Evaluation

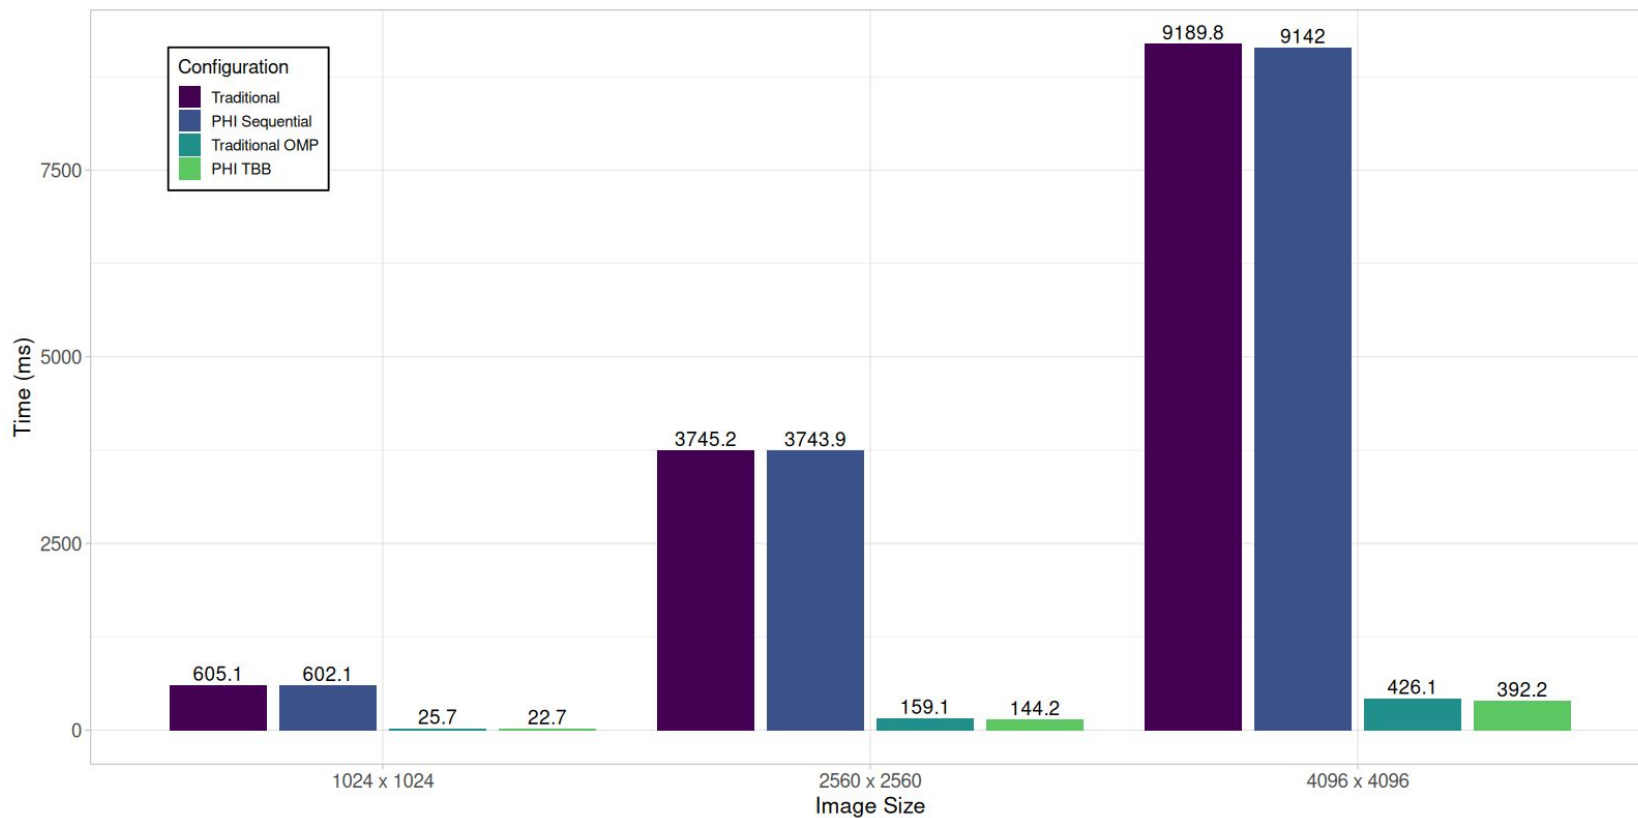


TABLE OF CONTENTS

01 Introduction

02 Background

- C++ 20 Ranges
- Parallel Patterns
- Syntax

03 Design

- Interface
- Range Adaptors & RACOs
- Parallel Patterns

04 Evaluation

05 Conclusion



- PHI leverages more idiomatic C++ range-like interfaces.
- Range adaptors and closure objects are the foundation of function composition for lazy evaluation.
- PHI is in early stages of development. Improvements need to be made.
 - More pattern adaptors: ***stencil - scan - filter - zip***.
 - Support for in-place data. Requires careful study.
 - Improved support for backend integration, and more backends.
 - Deeply nested types are difficult to work with. Study a better representation.
 - Non-linear chains of transformation.
 - Additional and more complex use-cases.



PHI

A Modern C++ Library for Parallel Pattern Composition

Santiago Veigas Ramírez

 sveigas@inf.uc3m.es

Daniel Martínez Davies
J. Daniel García Sánchez