



Towards Scriptable C++ Refactorings with Coccinelle



Michele MARTONE^{*}, Julia LAWALL[†]

^{*}michele.martone@lrz.de, Leibniz Supercomputing Centre, Garching near Munich, Germany

[†]julia.lawall@inria.fr, Inria, Paris, France

Context: Problems with large scientific HPC C/C++ Codebases

In our experience and in our environment, the problems of **large** HPC codebases are not so unlike those of non-HPC codebases:

- Keeping up to date to dependencies (e.g. libraries or parallelism APIs)
- Uncluttering old constructs (e.g. once new standards appear)
- Making code future proof (e.g. amenable to parallelism)
- Refactoring for performance (e.g. data layout change)
- Ensuring code properties (e.g. thread safety or against *code smells*)

A difference is perhaps in longer lifespans and smaller teams. HPC development tends to happen in bursts, e.g., when an idea is being developed and usually associated to a publication. Typically, *future proofing* and correctness efforts are tackled by the original author, who is not a computer scientist. At times, access to staff with specific expertise in Software Engineering or HPC is granted as a service.

Managing the Unfeasible

Our collaboration (see [C3PO]) was initiated from the need to obtain a programmable code refactoring that not only could be performed in little time (a matter of seconds), but could also be maintained by the code owners, who are astrophysicists, not C++ or parsing experts.

Manipulate Expressions, Globally

In [C3PO] we had to change all expressions involving a few *array of structures* accesses into corresponding *structure of arrays* accesses. This change impacted nearly the entire code (tens of thousands of *change sites*), and yielded an estimated 2 to 5× speedup on representative runs. We deem *global expression manipulation* to be **the most important application** of our approach.

We hope that the general C++ community can find uses of our otherwise HPC-oriented collaboration.

Low-Cost Performance Experiments

Rules can be instructed to act on selected quantities only, thus enabling refactorings that are *partial*. This opens the door to low-effort performance experiments, including also #pragmas, loops, and function manipulations like in the following subsections.

Modern C++ Transformation Example

Introduction of C++23's multi-index transformations is now possible (e.g. for `std::mdspan`, see [MDSPAN]). The following *semantic patch* with two rules:

```
1 #spatch --c++=23
2 @tomultiindex@
3 symbol a;
4 expression i,j,k;
5 @@
6 - a[i][j][k]
7 + a[i, j, k]
8
9 @@
10 symbol b;
11 @@
12 - b[...]
13 + b[0]
```

leads to the following code patch (context deliberately expanded):

```
1 @@ -1,8 +1,8 @@
2 int main()
3 {
4     int a[1][1][1];
5     int b[1][1][1];
6     int i=0, j=0, k=0;
7 - a[i][j][k]++;
8 - b[i][j][k]++;
9 + a[i, j, k]++;
10 + b[0][j][k]++;
11 }
```

Notice that lines 5–6 of the semantic patch (rule `@tomultiindex@`) modify expressions of **arbitrarily complicated** statements.

Current Work

We are expanding COCCINELLE's C++ support:

- constructors, destructors
- template declarations and instantiation
- namespaces, misc keywords
- lambda functions
- variadic operators

Unsupported syntax in source code leads to skipping transformations. Once we have enough syntax covered, we will develop use cases.

Background Work

- Joint expertises:
 - Performance-oriented refactoring of large HPC codes (Martone)
 - Formal methods for program rewriting (Lawall)
- Successful development of a *programmable, replayable* refactoring for a **200 kLoC** C codebase (see [C3PO])
- Using *semantic patching* technology of COCCINELLE

What this Poster is About

- Preliminary work for a submitted Franco-German project proposal
- In form of *code rewriting rules* and resulting *diff* output
- Refactorings recently enabled in COCCINELLE:
 - HPC-affine introduction of `mdspan` (modern C++)
 - A sample rule to replace loops with *STL algorithms*
 - Another sample rule to enforce a programming guideline

No Raw Loops Example

Coccinelle can match several constructs, comprehensive of control flow. The following example replaces a loop on elements with the use of an *STL algorithm*, which is good practice [NRL].

```
1 #spatch --c++=17
2 @@ @@
3 #include <iostream>
4 + #include <algorithm>
5 + #include <functional>
6
7 @@
8 type T;
9 constant k;
10 identifier elem, result, arrid;
11 @@
12 - bool result = false;
13     ...
14 - for ( T &elem : arrid )
15 -   if ( \( elem == k \ | k == elem \ ) )
16 -   {
17 -     ...
18 -     result = true;
19 -     break;
20 -   }
21 + const bool result =
22 +   (find(begin(arrid), end(arrid), k) !=
23 +    end(arrid));
```

```
1 @@ -1,20 +1,15 @@
2 #include <vector>
3 #include <iostream>
4 + #include <algorithm>
5 + #include <functional>
6 int main()
7 {
8     using namespace std;
9     vector v = {1,2,3};
10 - bool has_zero = false;
11
12     v[2] = 0;
13
14 - for ( int &a : v )
15 -   if ( 0 == a )
16 -   {
17 -     cout << "doing things\n";
18
19 -     has_zero = true;
20 -     break;
21 -   }
22 + const bool has_zero =
23 +   (find(begin(v), end(v), 0) != v.end());
24     cout << has_zero << endl;
25 }
```

Coccinelle

About:

- Design started in 2004 (open source in 2008)
- Targeted *"collateral evolutions"* (recurring changes motivated by changes in API interfaces), in **LINUX** device driver code (C)
- Contributed to around 9000 commits in the **LINUX** kernel
- Available in various **LINUX** distros

Distinguishing features:

- Self-contained project
- Unique *diff-like* (`+ ./- /- ..`) patch specification language

What *novel* uses do we foresee:

- Large-scale refactorings
- Data-layout changes
- Advanced expression manipulations
- Modern C/C++ standards and GPU-specific language extensions

Enforcing Coding Guidelines Example

Matching programming guidelines, whether project-specific or idiomatic. Here an example to recognize F.17 from the guidelines collection of Stroustrup and Sutter [F17].

```
1 #spatch --c++
2 @r1@
3 type T;
4 identifier f;
5 parameter list pl;
6 @@
7
8 T f(pl) {}
9
10 @r2@
11 typedef A, B;
12 type heavy_type = {A, B};
13 type r1.T;
14 identifier r1.f;
15 symbol i;
16 @@
17
18 // Note: heavy copy!
19 T f (
20     ..., heavy_type i, ...
21 ) { ... }
```


```
1 @@ -1,13 +1,16 @@
2 #include <array>
3 struct A { std::array<int,999> a; }; // heavy
4 struct B { std::array<int,999> a; }; // heavy
5 struct C { std::array<int, 16> a; }; // light
6 void if1(int i) {}
7 void if2(int &arg) {}
8 void if3(const int &arg) {}
9 // Note: heavy copy!
10 void af1(A i) {}
11 // Note: heavy copy!
12 void bf1(B i) {}
13 // Note: heavy copy!
14 void bf2(const B i) {}
15 void bf3(const B &i) {}
16 void cf1(C i) {}
17 int main() { }
```

Key Points

- Best if used on codebases with code conventions in place
- Transformations preserve most spacing and comments
- Specificity vs generality is at the user's discretion
- This is *preliminary* work
- C++ codes are very diverse
- **We're curious to hear about your refactoring patterns**

Acknowledgements

- This work was partially funded by *SiVeGCS*
- Work visits in this collaboration have been supported by the BayFrance'23 scheme (<https://www.bayern-france.org/>), financed by the *Bavarian Ministry of State for Education, Culture, Science and Art (StMBW)* and the *French Ministry of Europe and Foreign Affairs (MEAE)*



References

- [C3PO] M. Martone, J. Lawall; "Refactoring for Performance with Semantic Patching: Case Study with Recipes"; Proceedings of the "Compiler-assisted Correctness Checking and Performance Optimization for HPC" Workshop at ISC'21 (**preprint**: <https://hal.inria.fr/hal-03266521>; **doi**: https://link.springer.com/chapter/10.1007/978-3-030-90539-2_15)
- [NRL] Sean Parent. *GoingNative 2013 C++ Seasoning* at <https://www.youtube.com/watch?v=W2tW0dZgXHA>
- [F17] Bjarne Stroustrup and Herb Sutter. *F.17: For "in-out" parameters, pass by reference to non-const* at <https://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines>
- [MDSPAN] VV.AA. <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2022/p0009r18.html>