

Migrating large codebases to C++ Modules

Y Takahashi¹, O Shadura² and V Vassilev³

¹University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8654, Japan

²University of Nebraska-Lincoln, 1400 R St, Lincoln, NE 68588, USA

³Princeton University, Princeton, New Jersey 08544, USA

E-mail: yukatkh@is.s.u-tokyo.ac.jp

Abstract. ROOT has several features which interact with libraries and require implicit header inclusion. This can be triggered by reading or writing data on disk, or user actions at the prompt. Often, the headers are immutable, and reparsing is redundant. C++ Modules are designed to minimize the reparsing of the same header content by providing an efficient on-disk representation of C++ Code. ROOT has released a C++ Modules-aware technology preview, which intends to become the default for the ROOT 6.20 release.

In this paper, we will summarize our experience with migrating C++ Modules to LHC experiment's software codebases, particularly for CMS software (CMSSW). We outline the challenges with adopting C++ Modules for CMSSW, including the integration of C++ Modules support in the CMS build system and we will evaluate the CMSSW performance benefits.

1. Introduction

In High Energy Physics (HEP), experiments such as CMS, produce a large amount of data each second, and we expect to have even more data generated during the Hi-Lumi LHC [1] era. Thus software for HEP is always striving to archive better performance, which could improve the users data analysis as well the cost optimisation for the HEP resources used in the data centers.

ROOT [2] is a core software used in HEP not only for data analysis but also as a backend for LHC experiment's software such as CMSSW. Thus the performance improvement in ROOT will result in also improving performance of the experiment's software, which requires more resources than ROOT itself.

C++ Modules [3] has been supported in ROOT since version 6.16, in order to improve its performance. The cost of header re-parsing can be negligible in small to medium size codebases, but can be critical in larger codebases. In this case compile-time scalability will be affected, even though it will not affect the programs at runtime. However, ROOT is different – its C++ interpreter Cling [10] processes code at program execution time and avoidance of redundant content re-parsing yields better runtime performance.

2. Background

2.1. C++ Modules in ROOT's Dictionaries

C++ Modules represent the compiler internal state serialized on disk and deserialized on demand to avoid repetitions of parsing operations, as we have described in [4]. There are different implementations of the concept in Clang, GCC, and MSVC. ROOT uses Clang through its

interpreter Cling and adopts its C++ Modules implementation, which is one of the most sophisticated on the market as of today.

Many operations in ROOT require loading a shared library and its corresponding header files. For example, when we serialize or deserialize a ROOT file, ROOT needs to know which library contains how to stream its object. This meta information is widely known as ROOT "dictionary" information. The dictionary is generated by a platform-independent tool, *rootcling*, which processes ROOT-aware code and produces a C++ source code file with a prefix *G_*. The file is later compiled and linked into the library it describes.

Generally, there are two ways to load a library: implicitly or explicitly. Both models require the library description to be available before or no later than library loading time. That is, all header files describing the library should be processed by the interpreter blindly without knowing if they will be used. This is a severe performance bottleneck which ROOT addresses with several technologies: *ROOTMAPs*, *RDICTs*, and a *PCH*. The ROOTMAP file represents a lightweight version of the library descriptor containing forward declarations and their mapping to the corresponding library. The RDICT file represents a cache of a subset of *TClass* objects which will be created when the library is loaded. The PCH file contains almost all of the headers in ROOT in an optimized form which is loaded at startup to avoid header re-parsing.

The ROOTMAPs and RDICTs are home-grown fixes to the inherent PCH problem – it is impossible to extend a PCH without fully recompiling it. C++ Modules are the industry solution to this problem – they represent composable PCH files called *PCMs*.

A PCM is generated by *rootcling* using a special file containing the mapping between a header file and the module to be created, called a *modulemap* file. It describes how a collection of existing headers corresponds to the logical structure of a module. It is important to note that transitive includes not present in a modulemap file are persisted multiple times. That is, if *A.h* includes *C.h*, *B.h* includes *C.h* and the modulemap maps *A.h* to module *Alpha* and *B.h* to module *Beta*, the content of *C.h* will be duplicated in both Modules. Such duplication can reduce performance (increasing parsing time) dramatically and should be avoided at almost any cost. This important detail suggests that "modularization" should be done bottom-up, namely, starting from external dependencies onward.

2.2. C++ Modules in CMSSW

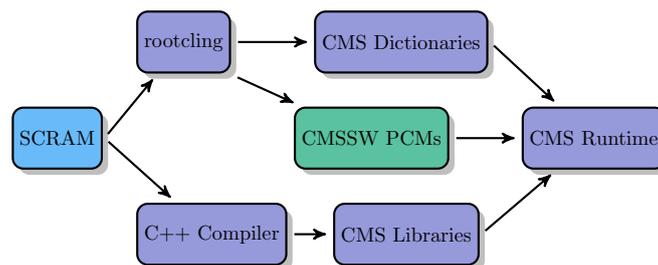


Figure 1: Dependency Graph of C++ Modules in CMSSW

CMS [5] is one of the largest experiments in LHC and performance improvement of its software is crucial for the coming HL-LHC. CMS experiment develops its software stack called CMSSW (CMS SoftWare) which uses ROOT as its backend.

CMSSW utilizes their own build system, SCRAM [6]. SCRAM is a configuration management tool, a distribution system, a build system, and a resource manager, with local resources and applications managed transparently.

As shown in Figure 1, the SCRAM build system resolves library dependencies and executes *genreflex* and C++ compiler such as *gcc*. *Genreflex* is a wrapper of *rootcling*, which supports

legacy interface for experiments using ROOT. Rootcling generates CMS dictionaries such as rootmaps and RDICTs, as it was in ROOT. Simultaneously, C++ Compiler compiles CMS C++ files and generates CMS shared object libraries and CMSSW executable. Generated CMS libraries and dictionaries are being loaded at CMS Runtime.

We added C++ Modules to the existing CMS system, shown in Figure 1. Rootcling generates CMSSW PCMs using a modulemap generated by SCRAM, and they will be used at CMS Runtime to save its parsing overhead.

3. Migrating CMSSW to C++ Modules

Modularizing CMSSW requires two steps: generating a modulemap file and enabling C++ Modules in *genreflex(rootcling)*. The generation of modulemap file happens at configuration time where we produce a file describing each library as a module and each library header file as a submodule. This can be subtle in some build systems as it requires tracking of header files which usually is not required. SCRAM already supports this and it was trivial to synthesize that description. Enabling C++ Modules in rootcling is done by adding an extra *-cxxmodule* flag. We started doing that gradually library-by-library to avoid unnecessary complications.

The mixed mode support ensures an incremental migration path to C++ Modules while having a stable system throughout the migration period. For instance, the dictionaries of CMSSW may use the old dictionary system while the dictionaries of ROOT use the new technology. Ultimately, the mixed mode will not manifest into performance improvements because the C++ Modules technology intends to address dictionaries outside of the ROOT PCH, namely, in third-party software.

This section describes the steps taken in CMSSW in order to migrate to C++ Modules.

3.1. Header Sanitizing

In most cases, a module corresponds to a single dictionary or a library. Each module enumerates every header file in a submodule. Each submodule needs to be able to compile in isolation. Illustratively, a separate compiler instance is run on each header file. This assumes that every header file should be able to compile on its own. *Standalone* header files include what they use and are resilient to configuration macros.

CMSSW codebase had a lot of header files inaccuracies thoroughly described in the GitHub C++ Modules Meta Issue [7]. The issues can be classified into the following categories:

- Incomplete headers – header files which do not include what they use. They are easy to fix because the C++ Modules system usually is able to suggest which are the omitted header files;
- Broken headers (unnecessary headers) – header files which were never compiled. Those header files were never included in any translation unit but were ”part” of a library. They are easy to deprecate and remove;
- Cyclic headers – header files which include graph contains a cycle. For example, header *A* includes header *B* which includes *A*. Header *A* is in module *Alpha* and header *B* in *Beta*. Usually this is a signal for a layering violation – concepts from one library depend on another an vice versa. In many cases, using forward declarations of such entities resolves the problem. In some cases more sophisticated engineering techniques are necessary, such as, refactoring, moving the two dependent headers together or splitting the common dependent logic into its own library and module. In rare cases, if the mutual dependence of headers is by design they have to be in a single submodule.
- Macro headers – header files which contain predominantly macro definitions which can be expanded differently in each translation unit. For example, *<assert.h>* which conditionally

defines the *assert* if *NDEBUG* is not defined. Macro headers are meant to be always textually included and they should be marked as *textual headers* in the modulemap file.

- Token generating headers – header files which enable preprocessor metaprogramming should be excluded from the modulemap file.

After sanitizing the header files, a toolchain to protect regressions can be introduced. In CMSSW every header in a pull request is checked for the aforementioned issues. This is done by precompiling each header on its own. In future, we will deploy the include graph sanitization tool which we used to detect include cycles [8].

3.2. Modularizing External Dependencies

In order to fully see the expected performance benefits, we should modularize external dependencies as well. An external dependency should be modularized if a module transitively includes it. For example if header *A* from module *Alpha* includes directly or indirectly *<vector>*, the external library, *libstdc++*, containing that file should be modularized.

It can be challenging to modularize external dependencies. On one hand, it may be difficult to know what its modulemap file should contain, on the other hand, external dependencies can be located in non-writable locations on the file system. We have prepared a set of modulemap files for all external dependencies of CMSSW [9]. The second challenge is solved by a file which instructs rootcling to pretend that the modulemap file is located at the non-writable folder in the file system. The *virtual file system overlay file (VFSOF)* is programmatically synthesized by Cling.

3.3. Automatic Generation of Modulemap Files and Virtual File System Overlay File

As was mentioned in the section 2.1, one of the important features to be adjusted is generation modulemaps. Well-behaved header files are trivial to enumerate in the modulemap file, as illustrated in Listing 1. The modulemap file syntax allows easy automation and most of its content can be generated by a build system. In CMSSW, modulemap autogeneration was simple due to its build system and header files structure. The codebase has the library interface headers files in a separate folder. We automatically generate the modulemap by iterating through those interface headers. It is expected that the modulemap is generated at configuration time but it must be generated no later then the first invocation of *rootcling*.

```
module DataFormatsTrackerCommon_xr {
  module "TrackerTopology" {header "DataFormats/TrackerCommon/interface/TrackerTopology.h" export *}
  module "TrackerDetSide" {header "DataFormats/TrackerCommon/interface/TrackerDetSide.h" export *}
  module "ClusterSummary" {header "DataFormats/TrackerCommon/interface/ClusterSummary.h" export *}

  export *
}
```

Listing 1: An example of a C++ Module definition in the CMSSW modulemap file.

The process of modularization subsumes external dependencies such as *libc*, *libstdc++* and *libxml*. Usually, external libraries do not ship with modulemap files. Listing 1 demonstrates that the header location is relative to the include path in modulemaps. However, system headers are located in non-writable areas and modulemaps are impossible to be placed in order to access the headers with relative paths, unless super user privileges are granted.

Figure 2 shows the organization of the C++ Modules infrastructure together with the module map system. Blue rectangles describe the physical directories, and purple rectangles describes the physical files. Black thick arrows denote physical files and dashed lines show the virtual file system relationship.

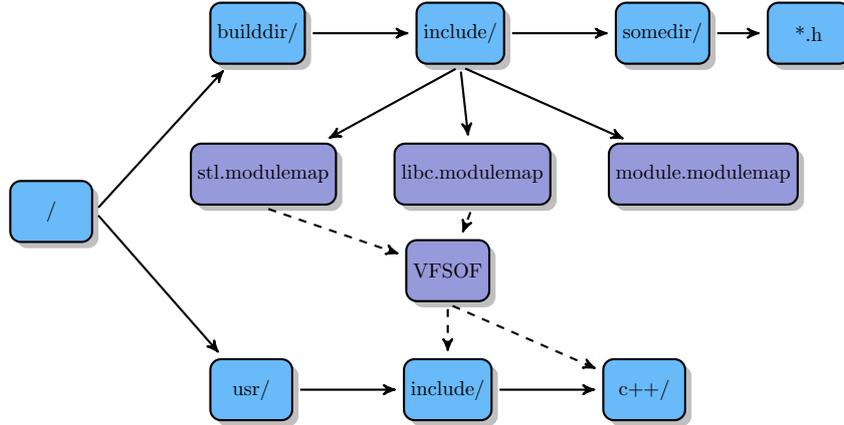


Figure 2: Modulemap directory physical vs virtual structure.

```

{ 'version': 0,
  'roots': [
    { 'name': '/usr/include/c++/', 'type': 'directory',
      'contents': [
        { 'name': 'module.modulemap', 'type': 'file',
          'external-contents': '/builddir/include/stl.modulemap' }]],
    { 'name': '/usr/include/', 'type': 'directory',
      'contents': [
        { 'name': 'module.modulemap', 'type': 'file',
          'external-contents': '/builddir/include/libc.modulemap'
        }]]}
  ]}]

```

Listing 2: An example of a VFSOF for libc.modulemap and stl.modulemap.

Modulemap files for libc and stl (libstdc++) need to be present in `/usr/include` and `/usr/include/c++` respectively. Listing 2 illustrates how the VFSOF instructs the infrastructure to consider `libc.modulemap` and `stl.modulemap` as if they were physically present in the expected folders. This approach allows us to "mount" modulemap files anywhere on the system enabling full software stack modularization.

The pre-configured VFSOF did not match the deployment process of CMSSW due to its static nature. The PCMs became non-relocatable from build directory as the path was hardcoded to PCMs. A more dynamic, on-memory representation of VFSOF was introduced to solve the problem. As it is a flexible virtual file on memory, paths do not need to be hardcoded at the configuration time, and thus it can determine header paths at runtime after the binary was distributed.

4. Preliminary Performance Results

The mixed run mode in ROOT allows us to make partial performance studies to assess the impact of this technology. The performance study was conducted on the CMSSW build server, where all the CMS software is continuously integrated. As CMSSW is large and there exist huge variations of possible physics workflows, we measured realistic CMS workflow tests which were expected to represent the actual code that physicists would run for their analysis.

Figure 3 shows the benchmark of CPU time and RSS memory usage. On every plot, *ROOT Master* is the case where ROOT and CMSSW built without C++ Modules, and serves as a baseline. *CMS PCMs* is the configuration where ROOT (96 pcms) and CMSSW (25 pcms), which are both built with C++ Modules. Although CMSSW has over 200 libraries, we have modularized only 25 out of them at this point of the study.

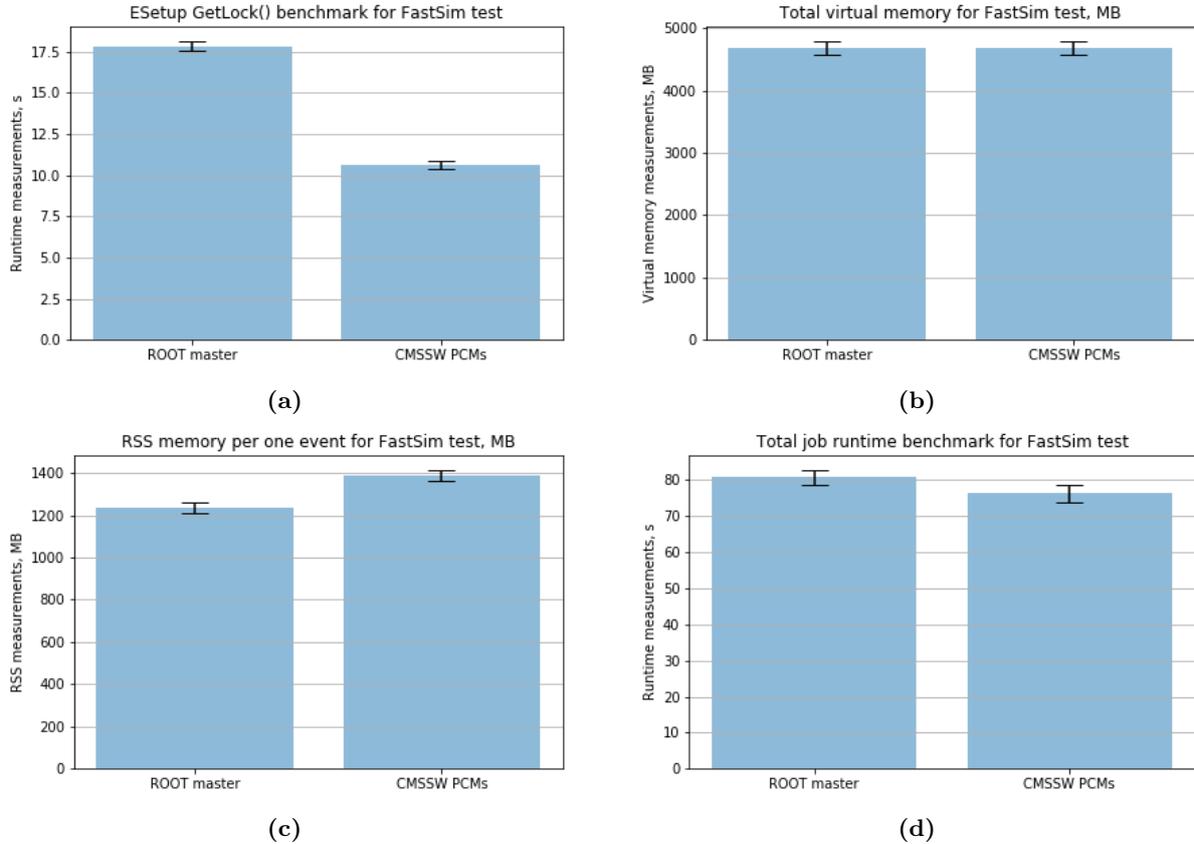


Figure 3: Performance results: (a) shows EventSetup GetLock() run time (RT) measurements for the fast simulation test. (b) shows virtual memory measurements for the fast simulation test. (c) shows RSS measurements for the fast simulation test. (d) shows the total job measurements for the the fast simulation test.

Figures 3c and 3d show the improvements for run time measurements, particularly for the event preparation phase (ESetup GetLock()). Figure 3d proves the absence of performance degradation, which was the intent of the study at this stage of CMSSW modularization. Figures 3b, 3c with memory measurements don't show the expected improvements. We miss memory performance improvements mainly due to two reasons: the incomplete modularization (only 15 % of CMSSW libraries were modularized) and the preloading of all modules at startup time in sparse workflows, which is a reason why CMS PCMs use more RSS memory (see Figure 3d). Error bars are caused by using different set of the events for CMSSW tests and possible fluctuations due to inconsistent CPU load of computing nodes used for tests.

5. Conclusion

We have shown the current advancements in the modularization of the CMSSW codebase. Despite the partial migration, we already observed some improvements. We are working towards modularization of the rest of the CMSSW components and external dependencies. We have implemented tools to aid this process and we are working on a further RSS memory reduction by avoiding preloading of all modules at startup time. It requires global module content indexing and loading the corresponding modules on demand. We expect more sophisticated benchmarking to be done, improving the coverage of many more workflows.

6. Acknowledgments

This work has been supported by an Intel Parallel Computing Center grant, by U.S. National Science Foundation grants PHY-1450377, ACI-1450323 and PHY-1624356, and by the U.S. Department of Energy, Office of Science.

The authors are thankful to Shahzad Malik Muzaffar and Mircho Rodozov from the CMSSW development team, CERN/EP-SFT and the ROOT team.

References

- [1] The HL-LHC Project. 2019. The HL-LHC Project — High Luminosity Large Hadron Collider. [ONLINE] Available at: <http://hilumilhc.web.cern.ch/>. [Accessed 15 May 2019].
- [2] R. Brun, F. Rademakers. ROOT - An Object Oriented Data Analysis Framework. 1997 *Nucl. Inst. & Meth. in Phys. Res. A* **389**, Proceedings AIHENP'96 Workshop.
- [3] Vassil Vassilev, Optimizing ROOT's Performance Using C++ Modules, 2017 *Journal of Physics: Conf. Ser.*, **898**, 072023
- [4] Yuka Takahashi *et. al.*, 2018 Optimizing Frameworks Performance Using C++ Modules Aware ROOT, *CoRR*, arXiv:1812.03992 [cs.PL]
- [5] CMS Collaboration, CMS: The computing project. 2005 *Technical design report*, CERN-LHCC-2005-023.
- [6] GitHub. 2019. GitHub - cms-sw/SCRAM: Software Configuration And Management - CMS internal build tool. [ONLINE] Available at: <https://github.com/cms-sw/SCRAM>. [Accessed 15 May 2019].
- [7] GitHub. 2019. C++ Modules support (based on Clang) - Issue 15248 -cms-sw/cmssw GitHub. [ONLINE] Available at: <https://github.com/cms-sw/cmssw/issues/15248>. [Accessed 15 May 2019].
- [8] GitHub. 2019. Circle Break - Teemperor/circle-break. [ONLINE] Available at: <https://github.com/Teemperor/circle-break>. [Accessed 15 May 2019].
- [9] GitHub. 2019. GitHub - Teemperor/ClangAutoModules: Automatically mounts clang modules for your system libraries (and more). [ONLINE] Available at: <https://github.com/Teemperor/ClangAutoModules>. [Accessed 15 May 2019].
- [10] V.Vassilev *et. al.*, Cling – The New Interactive Interpreter for ROOT 6, 2012, *Journal of Physics, Conf. Ser.*, **396**, 052071