Interactive C++ for Data Science





VASSIL VASSILEV 20 21 October 24-29





Speaker



V.Vassilev – Interactive C++ for Data Science

27-Oct-2021

Vassil Vassilev, Research Software Engineer, Princeton/CERN



Outline

- Introduction
 - Key insights of Interactive C++
 - Interpreting C++. Tools and technology
- Interactive, interpretative C++ for Data Science
 - Eval-style programming, C++ in notebooks, CUDA C++
- Beyond just interpreting C++
 - Template instantiation on demand, Language Interop on Demand
- Interpreter/Compiler as a service
 - Extensions, Automatic differentiation on the fly, Lifelong Optimization
- Conclusion

Execution of statements, Execution Results, Entity redefinitions, Error Recovery, Code Undo



Acknowledgement & Disclaimer

- This talk includes technologies developed by various individuals and organizations in the area of interpretative C++ since 1998
- This talk is about work conducted by me but also the work of dozens colleagues and contributors from many domains in science and industry. In the slides I have tried to mention individuals and organizations where possible.
- Any characterizations, mischaracterizations, emphasis and errors are solely mine and do not necessarily represent the views of other individuals or organizations.



The current work is partially supported by National Science Foundation under Grant OAC-1931408. Any opinions, findings, conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.







Introduction

Despite its power, C++ is often seen as difficult to learn, difficult to teach and inconsistent with rapid application development/exploratory programming.

Exploration and prototyping is slowed down by the long edit-compile-run cycles during development.

Is C++ a "compiled" language?
Is C++ an "experts" language?



Choice of translation

The language translation should be an implementation choice.

Language design could affect the choice of translation and the choice of translation can affect language design.

V.Vassilev – Interactive C++ for Data Science

27-Oct-2021

What would it enable if we start interpreting C++?



Interactive C++

11:16:58-vvassilev~/Downloads/cling-opengl-demo\$

The invisible compile-run cycle aids interactive use offering a different programming experience and enhanced productivity. It becomes trivial to orient a shape, choose size and color or compare to previous settings.

27-Oct-2021

V.Vassilev – Interactive C++ for Data Science

cling-opengl-demo — -bash — 137×35

Video Credits: A. Penev





Exploratory Programming

- An effective way to gain understanding of code/algorithms
- users
- Significantly reduces the edit-compile-run cycle

Interactive probing of data and interfaces makes complex libraries accessible to



Exploratory Programming

- shorten the compile-run cycle but generally have a noticeable cost
- may also put syntax sugar to improve convenience and terseness
- compilation technology

Languages which enable exploratory programming tend to have interpreters to

• Language designers who acknowledge the use case of exploratory programming

Performance is mitigated nowadays by just-in-time (JIT) or ahead-of-time (AOT)



Interactive C++. Key Insights

- Incremental Compilation
- Handling errors
 - Syntactic
 - Semantic
- Execution of statements
- Displaying execution results
- Entity redefinition

[cling] #include <vector>
[cling] std::vector<int> v = {1,2,3,4,5};

[cling] std.sort(v.begin(), v.end()); input_line_1:1:1: error: unexpected namespace name 'std': expected expression std.sort(v.begin(), v.end()); ^

[cling] std::sort(v.begin(), v.end()); [cling] v // No semicolon (std::vector<int> &) { 1, 2, 3, 4, 5 }

[cling] std::string v = "Hello World"
(std::string &) "Hello World"





Interactive C++. Key Insights

 Eval-style programming – enables embedding in frameworks

```
[cling]$ #include <cling/Interpreter/Value.h>
[cling]$ #include <cling/Interpreter/Interpreter.h>
[cling]$ int i = 1;
[cling]$ cling::Value V;
[cling]$ gCling->evaluate("++i", V);
[cling]$ i
(int) 2
[cling]$ V
(cling::Value &) boxes [(int) 2]
[cling]$ ++i
(int) 3
[cling]$ V
(cling::Value &) boxes [(int) 2]
```





Interpreters for C++

- beginners to learn mathematics, computing, numerical analysis (numeric methods), and programming in C/C++. Last release 2017
- CINT a command line C/C++ interpreter that was originally developed by Last release 2013
- physics to replace CINT in ROOT. Last release 2021

• Ch – a proprietary cross-platform C and C++ interpreter and scripting language environment, originally designed by Harry H. Cheng as a scripting language for

Masaharu Goto and later included in the ROOT data analysis software package.

• Cling – an LLVM-based C/C++ interpreter developed in the field of high-energy



Interpreting C++. Cling



27-Oct-2021



Interpreting C++. Cling

- Cling has originally developed in the field of high energy physics to enable developers.
- and virtual inheritance.
- makes the language more welcoming for use.
- Cling compiles C++ code incrementally and relies on JIT compilation
- Cling enables exploratory programming for C++

interactivity, dynamic interoperability and rapid prototyping capabilities to C++

Cling supports the full C++ feature set including the use of templates, lambdas,

Cling adds a small set of extensions in C++ to allow interactive exploration and



Interactive C++ in Data Science

Data Science

- theoretical and computational science.
- roughly 11.5 million new jobs).

 Data science is a "concept to unify statistics, data analysis, informatics, and their related methods" in order to "understand and analyze actual phenomena" with data. Considered by some as the "fourth paradigm" of science next to empirical,

Per LinkedIn, there has been a 650% increase in data science jobs since 2012.

• The U.S. Bureau of Labor Statistics sees strong growth in the data science field and predicts the number of jobs will increase by about 28% through 2026 (that is



Xeus-Cling. C++ in Notebooks

- Xeus-Cling is a Cling-based Jupyter Notebook kernel
- If you use C++ in Jupyter you use xeus-cling
- Developed and maintained by the private company QuantStack





Xeus-Cling. C++ in Notebooks



Rich mime type rendering in Jupyter

27-Oct-2021

	()	File Edit View Run Kernel Tabs Settings Help		
	Files	Untitled.ipynb ● B + % □ □ ▶ ■ C Code ∨		
	Running	In [1]: ?std::vector Create account	Search	
	Commands	Page Discussion C++ Containers library std::vector std::Vector	View	Edit
	ell Tools (Defined in header <vector> template< class T, class Allocator = std::allocator<t> class vector; (1)</t></vector>	=	
Hub ∂Binder Code ∽	Tabs C	<pre>namespace pmr { template <class t=""> using vector = std::vector<t, std::pmr::polymorphic_allocator<t="">>; } 1) std::vector is a sequence container that encapsulates dynamic size arrays.</t,></class></pre> (2)	(since C++17)	
		In []:		
$+10\sqrt{2}x^9 + 720x^2 + 1680x^4$	4 + 840	^{x⁶ + 90x⁸ + x¹⁰ Direct access to docu}	ument	ati
		▼ Mode: Edit ⊗ Ln 1, Col 1 symengine.ipynb		







Xeus-Cling. C++ in Notebooks

()	File Edit Viev	v Run Kernel Tabs Settings Help			
Files	xwidgets.ipyn	t● □ □ ▶ ■ C Code ∽ C	++		
Running	Numerical widgets				
nands		Defining a Slider Widget			
Comr	In [1]:	<pre>#include "xwidgets/xslider.hpp"</pre>			
	In [*]:	xw::slider <double> slider;</double>			
bs Cell Tools	In []:	slider			
		slider.value = 20;			
	In []:	slider.value()			
Та	In []:	<pre>// changine some more properties slider.max = 40; slider.style().handle_color = "blue"; slider.orientation = "vertical"; slider.description = "A slider";</pre>			
	In []:	<pre>#include "xcpp/xdisplay.hpp"</pre>			
		<pre>using xcpp::display;</pre>			

Xwidgets – User-defined controls

S. Corlay, Quantstack, *Deep dive into the Xeus-based Cling kernel for Jupyter*, May 2021, compiler-research.org

27-Oct-2021

V.Vassilev – Interactive C++ for Data Science



Xleaflet – Interactive Geo Information System



Interactive CUDA C++

- Interactive CUDA backend
- Developed and maintained by HZDR

[cling]	#incl
[cling]	#incl
[cling]	#prag
// set	parame
[cling]	glo
[cling]	? i
[cling]	? i
[cling]	?
[cling]	? }
[cling]	
[cling]	// la
[cling]	init<
[cling]	init<
[cling]	
[cling]	cubla
dim, &a	lpha,
[cling]	cubla
[cling]	cudaG
(cudaEr	ror_t)

```
.ude <iostream>
ude <cublas v2.h>
ma cling(load "libcublas.so")
eters, allocate memory ...
bal void init(float *matrix, int size) {
nt x = blockIdx.x * blockDim.x + threadIdx.x;
f (x < size)
matrix[x] = x;
```

unching a function direct in the global space <<blocks, threads>>>(d A, dim*dim); <<blocks, threads>>>(d B, dim*dim);

```
sSgemm(handle, CUBLAS OP N, CUBLAS OP N, dim, dim,
d A, dim, d B, dim, &beta, d C, dim);
sGetVector(dim*dim, sizeof(h_C[0]), d_C, 1, h_C, 1);
GetLastError()
```

```
(cudaError::cudaSuccess) : (unsigned int) 0
```





Interactive CUDA C++

		JupyterLab - Mozilla Firefox	- 0	8
C Jupyter	rLab	\times +		
$\overleftarrow{\leftarrow}$ \rightarrow	C' û	t 🗊 🛈 localhost:8888/lab … 🖾 🏠	lii\ 🗊	Ξ
📁 File	Edit	View Run Kernel Tabs Settings Help		
8	+ %	Code ∽	C++14-CUDA	0
•	[1:	<pre>global void compute(int * data, int width){ int x = blockIdx.x * blockDim.x + threadIdx.x; int y = blockIdx.y * blockDim.y + threadIdx.y; int id = y * width + x; }</pre>		
	[]:	<pre>for(int i = 0; i < 8; ++i){ compute<<<1,dim3(4,4,1)>>>(d_data, width); cudaMemcpy(h_data, d_data, m_size, cudaMemcpyDeviceToHost display_data(h_data, i+1); }</pre>);	
0 5 2	0 C	++14-CUDA Idle Mode: Edit ⊘ Ln 6, Col 5 simple	e_cuda_example.ir	oynb

S. Ehrig, HZDR, Cling's CUDA Backend: Interactive GPU development with CUDA C++, Mar 2021, compiler-research.org

27-Oct-2021



ROOT – Scientific Data Analysis

- data facilitating more than 1000 scientific publications
- to use it as a reflection information service for data serialization



 Developed and maintained by the field of high-energy physics and organizations such as CERN, FNAL, GSI, University of Nebraska, UC San Diego, Princeton

• Last 5 years the ROOT and Cling technology are used to store around 1EB physics

• The ROOT data analysis package embeds Cling to enable interactive C++ but also



Dynamic Scopes. Runtime Lookup

gCling->EvaluateT</*ret type*/void>("ntuple->GetTitle()", /*context*/); Eile Edit View Options Iools

[root] ntuple->GetTitle() error: use of undeclared identifier 'ntuple' [root] TFile::Open("tutorials/hsimple.root"); ntuple->GetTitle() (const char *) "Demo ntuple" [root] gFile->ls(); *TFile*** tutorials/hsimple.root Demo ROOT file with histograms tutorials/hsimple.root Demo ROOT file with histograms *TFile** This is the px distribution : 0 at: 0x7fadbb84e39 OBJ: TH1F hpx OBJ: TNtuple ntuple Demo ntuple : 0 at: 0x7fadbb93a890 KEY: TH1F hpx;1 This is the px distribution [...] KEY: TNtuple ntuple;1 Demo ntuple [root] hpx->Draw()

Eval-style programming enables Cling to be embedded in frameworks.

27-Oct-2021





Impact of Interactive C++ in Physics



[1] K. Morishima et al, Discovery of a big void in Khufu's Pyramid by observation of cosmic-ray muons, Nature, 2017 [2] Abbott et al, Observation of gravitational waves from a binary black hole merger. Physical review letters, 2016 [3] CMS Collab, Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. Physics Letters B, 2012

27-Oct-2021

[3]

Scientific breakthroughs such as the discovery of the big void in the Khufu's Pyramid, the gravitational waves and the Higgs boson heavily rely on the ROOT software package





Key Insights

- efficient incremental execution of C++
- Cling is used in several high-performance systems to provide reflection and introspection information
- regions can be annotated with specific optimization levels
- to defer for the target platform

Cling is not just a Repl, it is an embeddable and extensible execution system for

• Cling can produce efficient code for performance-critical tasks where hot-spot

Cling allows us to decide how much we want to compile statically and how much



Beyond Just Interpreting C++

Automatic Language InterOp Classification of Prior Work

- Static binding generators SWIG, SIP, Boost.Python, CxxWrap.jl the other language knows how to call. Advanced C++ features such as reference counting are hard to support.
- Dynamic / Automatic bindings cppyy, Cxx.jl

How to define to whether two languages are interoperable?

CaaS Monthly, Oct, 2021

Interactive C++: A Language InterOp Layer, V. Vassilev

V. Vassilev, Princeton, *Interactive C++: A Language InterOp Layer*, Oct 2021, compiler-research.org

V.Vassilev – Interactive C++ for Data Science

They rely on limited parsers and limited type introspection to generate wrapper functions which

They operate on demand and generate only what's necessary when necessary. They rely on a compiler available at program runtime: cling and the Julia JIT respectively. They are efficient and can support about advanced C++ features such as template instantiation. Harder to implement.



3

Automatic Language InterOp. Python

cppyy: Yet another Python – C++ binder?!

- Yes, but it has its niche: *bindings are runtime*
 - Python is all runtime, so runtime is more natural
 - C++-side runtime-ness is provided by Cling
- Very complete feature-set (not just "C with classes")
- Good performance on CPython; great with PyPy*
- pip: <u>https://pypi.org/project/cppyy/</u>
- conda: <u>https://anaconda.org/conda-forge/cppyy</u>
- git: https://github.com/wlav/cppyy
- docs: https://cppyy.readthedocs.io/en/latest/

For HEP users: *cppyy in ROOT is* an old fork. It won't run all the examples here, doesn't work with PyPy, and has worse performance.

(★) PyPy support lags CPython

- 2 -

W. Lavrijsen, LBL, <u>cppyy</u>, Sep 2021, compiler-research.org

Cppyy connects the Cling interpreter to the Python interpreter to co-operate.









Automatic Language InterOp. Python

Performance Compared to Static Approaches

No fundamental CPU performance difference

Note carefully that *everything* in Python is runtime: compile-time just means that the bindings *recipe* is compiled, not the actual bindings themselves!

- But heavy Cling/LLVM dependency:
 - ~25MB download cost; ~100MB memory overhead

- 24 -

Complex installation (and worse build)

ENERGY

W. Lavrijsen, LBL, <u>cppyy</u>, Sep 2021, compiler-research.org

.....

The approach does not require the project maintainer to bother providing static bindings

Basic Performance Test: overload

ΤοοΙ	Execution time (ms/call)*
C++ (Cling w/ -O2; out-of-line)	1.8E-6
срруу / руру-с	0.50
cppyy / CPython	1.25
swig (builtin)	1.29
swig (default)	4.23
pybind11	6.97

- 27 -

 \Rightarrow C++ overload is resolved at compile time, not based on dynamic type

⇒ Largest overhead: Python instance type checking (avoidable, but clumsy)

⇒ There is no obvious benefit to "static" over runtime bindings

(*) lower is better





Automatic Language InterOp. D

sil-cling: Architecture



A. Militaru, Symmetry Investments, <u>Calling C++ libraries from a D-written DSL: A cling/cppyy-based approach</u>, Feb 2021, compiler-research.org

Sil-Cling connects D to C++ via Cppyy at runtime

sil-cling: Interface with cppyy

binds with cppyy through the direct inclusion of the latter's C header using dpp

1 //cling.dpp

3 #include "capi.h" // cppyy's C header 5 // D code ↓ 6 import std.string : fromStringz, toStringz; 7 8 string resolveName(string cppItemName) 9 { import core.stdc.stdlib : free; 10 11 // Calling cppyy_resolve_name ↓ char* chars = cppyy_resolve_name(cppItemName.toStringz); 12 13 string result = chars.fromStringz.idup; 14 free (chars); 15 return result; 16 }

→ ~ dub run dpp -- cling.dpp --keep-d-files -c







Automatic Language InterOp. Julia

Features - Interactive C++ REPL

	Documentation: https://docs.julialang.org		
	Type "?" for help, "]?" for Pkg help.		
	Version 1.1.0 (2019–01–21) Official https://julialang.org/ release		
[julia> using Cxx			
[C++ > // Press '<' to act	ivate C++ mode		
[C++ > #include <iostream> true</iostream>			
[C++ > std::cout << "Welcome to Cxx.jl" << std::endl; Welcome to Cxx.jl			
<pre>[C++ > std::string cxx = "C++";</pre>			
<pre>[julia> println("Combine Julia and ", String(icxx"cxx;")) Combine Julia and C++</pre>			
julia> 🛛			

Cxx.jl, similarly to Cppyy uses incremental compilation (not based on Cling).

27-Oct-2021

V.Vassilev – Interactive C++ for Data Science



K. Fischer, Julia Computing, <u>A brief history of Cxx.jl</u>, Aug 2021, compiler-research.org



Compiler As A Service

Compiler As A Service

The design of Cling, just like Clang, allows it to be used as a library. In the next example we show how to incorporate libCling in a C++ program. The program can by compiled by your favorite compiler. Cling can be used on-demand, as a service, to compile, modify, describe or extend C++ but also use CUDA.



Compiler As A Service

```
// caas-demo.cpp
// g++ ... caas-demo.cpp; ./caas-demo
int main(int argc, const char* const* argv) {
  cling::Interpreter interp(argc, argv, LLVMDIR);
```

```
callInterpretedFn(interp);
return 0;
```

27-Oct-2021

```
/// Call an interpreted function using its symbol address.
void callInterpretedFn(cling::Interpreter& interp) {
  // Declare a function to the interpreter. Make it extern "C"
  // to remove mangling from the game.
  interp.declare("#pragma cling optimize(1)"
      extern \"C\" int cube(int x) { return x * x * x; }");
  void* addr = interp.getAddressOfGlobal("cube");
  using func t = int(int);
  func t* pFunc = cling::utils::VoidToFunctionPtr<func t*>(addr);
  std::cout << "7 * 7 * 7 = " << pFunc(7) << '\n';</pre>
```







Compiler As A Service. Extensions

int main(int argc, const char* const* argv) { std::vector<const char*> argvExt(argv, argv+argc); argvExt.push back("-fplugin=etc/cling/plugins/lib/clad.so"); cling::Interpreter interp(argvExt.size(), &argvExt[0], LLVMDIR); gimme pow2dx(interp); return 0;



Compiler As A Service. Extensions

#include <...>

// Derivatives as a service.

void gimme pow2dx(cling::Interpreter &interp) { // Definitions of declarations injected also into cling. interp.declare("double pow2(double x) { return x*x; }"); interp.declare("#include <clad/Differentiator/Differentiator.h>"); interp.declare("#pragma cling optimize(2)"); interp.declare("auto dfdx = clad::differentiate(pow2, 0);");

cling::Value res; // Will hold the evaluation result. interp.process("dfdx.getFunctionPtr();", &res);

```
using func t = double(double);
func t* pFunc = res.getAs<func t*>();
printf("dfdx at 1 = \frac{1}{n}, pFunc(1));
```

interp.process("dfdx.getCode();", &res); printf("dfdx code: %s\n %s\n", res.getAs<const char*>());

```
vvassilev@vv-nuc ~/.../builddir $ ./caas-demo
dfdx at 1 = 2.000000
dfdx code: double pow2_darg0(double x) {
   double _d_x = 1;
   return _d_x * x + x * _d_x;
vvassilev@vv-nuc ~/.../builddir $
```





Lifelong Optimization



27-Oct-2021



Incremental Compilation in LLVM

- possible to use incremental C++ in Clang
- in 120 LOC

In the context of the compiler-as-a-service project we are gradually making it

• We have put the infrastructure to required to instantiate and execute a template



Instantiating a C++ template in C

```
// gcc ... template instantiate_demo.C
#include "InterpreterUtils.h"
int main(int argc, char **argv) {
 Clang Parse("void* operator new( SIZE TYPE ,
void* p);"
      "extern \"C\" int printf(const char*,...);"
      "class A {};"
      "\n #include <typeinfo> \n"
     "struct B {"
     " template<typename T>"
     " void callme(T) {"
      " printf(\" Instantiated with [%s] \\n \",
typeid(T).name());"
      ··· } ··
      "};");
  const char * InstArgs = "A*";
  Decl t T = Clang LookupName("A");
 Decl t TemplatedClass = Clang LookupName("B");
```

```
vvassilev@vv-nuc ~/.../cpptemplate $ LD_LIBRARY_PATH="." ./template_instantiate_demo.out
Instantiated with [P1A]
Instantiated with [P8MyClass1]
vvassilev@vv-nuc ~/.../cpptemplate $
```

```
// ...
 // Instantiate B::callme with the given types
  Delc t Inst
     = Clang InstantiateTemplate(TemplatedClass,
"callme", InstArgs);
 // Get the symbol to call
  typedef void (*fn def) (void*);
  fn def callme fn ptr
    = (fn def) Clang GetFunctionAddress(Inst);
 // Create object of type T
 void* NewT = Clang CreateObject(T);
  callme fn ptr(NewT);
  return 0;
```

vvassilev@vv-nuc ~/.../cpptemplate \$ LD_LIBRARY_PATH="." ./template_instantiate_demo.out "class MyClass1{};" "MyClass1" "MyClass1*"







Conclusion

- Is C++ a "compiled" language?
- Is C++ an "experts" language?
- What would it enable if we start interpreting C++?



Thank You!

Selected References

- •https://blog.llvm.org/posts/2020-11-30-interactive-cpp-with-cling/
- •https://blog.llvm.org/posts/2020-12-21-interactive-cpp-for-data-science/
- https://blog.llvm.org/posts/2021-03-25-cling-beyond-just-interpreting-cpp/
- https://Compiler-Research.org
- https://root.cern
- https://root.cern/cling



