

A short course on object-oriented numerics

Sylwester Arabas

Faculty of Physics, University of Warsaw

local organisers:

Juan A. Añel, Orlando García

EPHysLab, Universidade de Vigo, Ourense

day 2 (June 4th 2014)

plan for today: getting acquainted with Blitz++

Veldhuizen & Jernigan 1997

Will C++ be faster than Fortran?

Todd L. Veldhuizen and M. Ed Jernigan

Department of Systems Design Engineering, University of Waterloo,
Waterloo Ontario N2L 3G1 Canada
tveldhui@monet.uwaterloo.ca

Abstract. After years of being dismissed as too slow for scientific computing, C++ has caught up with Fortran and appears ready to give it stiff competition. We survey the reasons for the historically poor performance of C++ (pairwise expression evaluation, the abstraction penalty, aliasing ambiguities) and explain how these problems have been resolved. C++ can be faster than Fortran for some applications, due to template techniques (such as expression templates and template metaprograms) which permit optimizations beyond the ability of current Fortran compilers.

documentation (0.9)

[http://blitz.sf.net/
resources/blitz-0.9.pdf](http://blitz.sf.net/resources/blitz-0.9.pdf)

packages (0.10)

Debian/Ubuntu: `libblitz-dev`
Fedora: `blitz-devel`

Example 3: Blitz++ API basics

```
1 #include <blitz/array.h>
2
3 int main()
4 {
5     blitz::Array<int, 2> a(4, 4);
6     // dimensionality.^ ^^^-size
7
8     // basic loop-free arithmetics
9     a = 0; // before - undefined values!
10    a += 2;
11    std::cerr << a << std::endl;
12
13 { // array-index arithmetics
14     using namespace blitz::tensor;
15     a = (i == j); // 1 on diagonal, else 0
16 }
17    std::cerr << a << std::endl;
18
19 { // element addressing
20     a(0, 0) = 2;
21
22     blitz::Range i(1, 2), j(1, 2);
23     a(i, j) = 8;
24
25     blitz::RectDomain<2> ij({i, j});
26     a(ij) -= 4;
27 }
28    std::cerr << a << std::endl;
29
30 { // reductions
31     using namespace blitz::tensor;
32     blitz::Array<int, 1> s(sum(a, j));
33     std::cerr << s << std::endl;
34 }
35 }
```

```
$ clang++ -std=c++11 -lblitz -Ofast \
           -march=native ex2.cpp
$ ./a.out
(0,3) x (0,3)
[ 2 2 2 2
  2 2 2 2
  2 2 2 2
  2 2 2 2 ]

(0,3) x (0,3)
[ 1 0 0 0
  0 1 0 0
  0 0 1 0
  0 0 0 1 ]

(0,3) x (0,3)
[ 2 0 0 0
  0 4 4 0
  0 4 4 0
  0 0 0 1 ]

(0,3)
[ 2 8 8 1 ]
```

Example 3: Blitz++ API basics (debug mode)

what if we change “`i(1,2)`” into `i(1,22)`

- ▶ compiles fine
- ▶ may run fine!
- ▶ if lucky, one may get a run-time error (segfault or alike)

Blitz++ debug mode: recompile with `-DBZ_DEBUG`

(to compile with clang++, see <http://sf.net/p/blitz/bugs/55/>)

```
$ g++ -std=c++11 -lblitz -DBZ_DEBUG ex2.cpp  
$ ./a.out
```

...

```
Bad array slice: Range(1, 22, 1). Array is Range(0, 3)  
a.out: /usr/include/blitz/array/slicing.cc:340:  
      void blitz::Array<P_numtype, N_rank>::slice(int, blitz::Range)  
      [with P_numtype = int; int N_rank = 2]: Assertion '0' failed.  
Aborted
```

Example 4: Blitz++ API / element-wise fun. evaluation

```
1 #include <blitz/array.h>
2 #include "fun.hpp"
3
4 int main()
5 {
6     // custom array ranges
7     blitz::Range
8     x(-3, 3), y(-3, 3);
9     blitz::Array<double, 2>
10    a(x, y), b(x, y);
11
12    // output settings
13    std::cout.precision(2);
14    std::cout.setf(std::ios::fixed);
15
16    // built-int fun. evaluation
17    {
18        using namespace blitz::tensor;
19        a = exp(-(pow(i,2) + pow(j,2)));
20    }
21    std::cout << a << std::endl;
22
23    // user-defined fun. evaluation
24    b = fun({.a = -1, .b = 1})(a);
25    std::cout << b << std::endl;
26 }
```

```
$ clang++ -std=c++11 -lblitz ex3.cpp
$ ./a.out
(-3,3) x (-3,3)
[ 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
  0.00 0.00 0.01 0.02 0.01 0.00 0.00
  0.00 0.01 0.14 0.37 0.14 0.01 0.00
  0.00 0.02 0.37 1.00 0.37 0.02 0.00
  0.00 0.01 0.14 0.37 0.14 0.01 0.00
  0.00 0.00 0.01 0.02 0.01 0.00 0.00
  0.00 0.00 0.00 0.00 0.00 0.00 0.00 ]
(-3,3) x (-3,3)
[ 1.00 1.00 1.00 1.00 1.00 1.00 1.00
  1.00 1.00 0.99 0.98 0.99 1.00 1.00
  1.00 0.99 0.86 0.63 0.86 0.99 1.00
  1.00 0.98 0.63 0.00 0.63 0.98 1.00
  1.00 0.99 0.86 0.63 0.86 0.99 1.00
  1.00 1.00 0.99 0.98 0.99 1.00 1.00
  1.00 1.00 1.00 1.00 1.00 1.00 1.00 ]
```

Example 4: Blitz++ API / element-wise fun. evaluation

```
1 #include <blitz/array.h>
2 #include "fun.hpp"
```

fun.hpp

```
1 #pragma once
2
3 #include <blitz/array.h>
4
5 struct fun
6 {
7     const double a, b;
8
9     double operator()(double x) const
10    {
11        return a * x + b;
12    }
13
14 // make it work with Blitz++ types
15 BZ_DECLARE_FUNCTOR(fun)
16 };
```

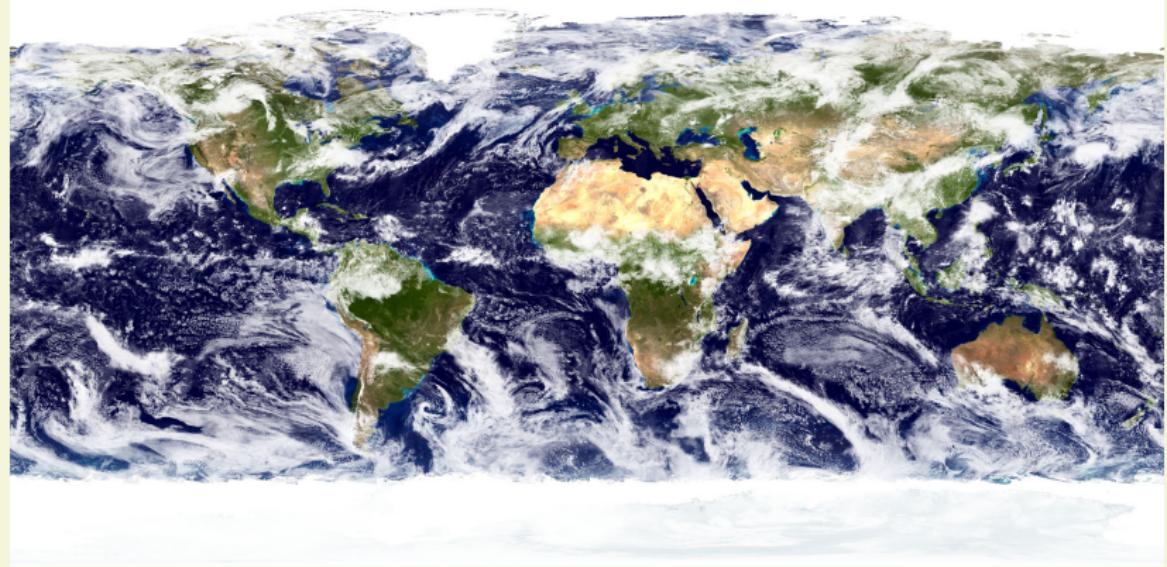
```
23 // user-defined fun. evaluation
24 b = fun({.a = -1, .b = 1})(a);
25 std::cout << b << std::endl;
26 }
```

```
$ clang++ -std=c++11 -lblitz ex3.cpp
$ ./a.out
(-3,3) x (-3,3)
[ 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
  0.00 0.00 0.01 0.02 0.01 0.00 0.00
  0.00 0.01 0.14 0.37 0.14 0.01 0.00
  0.00 0.02 0.37 1.00 0.37 0.02 0.00
  0.00 0.01 0.14 0.37 0.14 0.01 0.00
  0.00 0.00 0.01 0.02 0.01 0.00 0.00
  0.00 0.00 0.00 0.00 0.00 0.00 0.00 ]

(-3,3) x (-3,3)
[ 1.00 1.00 1.00 1.00 1.00 1.00 1.00
  1.00 1.00 0.99 0.98 0.99 1.00 1.00
  1.00 0.99 0.86 0.63 0.86 0.99 1.00
  1.00 0.98 0.63 0.00 0.63 0.98 1.00
  1.00 0.99 0.86 0.63 0.86 0.99 1.00
  1.00 1.00 0.99 0.98 0.99 1.00 1.00
  1.00 1.00 1.00 1.00 1.00 1.00 1.00 ]
```

Example 5: a draft of an advection-equation solver

$$\partial_t \psi + \nabla \cdot (\vec{u} \psi) = 0$$



http://earthobservatory.nasa.gov/Features/BlueMarble/BlueMarble_2002.php

Example 5: a draft of an advection-equation solver

$$\partial_t \psi + \nabla \cdot (\vec{u} \psi) = 0$$

donor-cell algorithm for constant $\vec{u} > 0$

$$C = \left\{ u_x \frac{\Delta t}{\Delta x}, \quad u_y \frac{\Delta t}{\Delta y} \right\}$$

$$\begin{aligned}\psi_{i,j}^{n+1} = \psi_{i,j}^n &- \left(C_x \psi_{i,j}^n - C_x \psi_{i-1,j}^n \right) \\ &- \left(C_y \psi_{i,j}^n - C_y \psi_{i,j-1}^n \right)\end{aligned}$$

(note: basic donor-cell has significant numerical diffusion)

Example 5: a draft of an advection-equation solver

```
1 #include <blitz/array.h>
2 #include <array>
3
4 int main()
5 {
6     // some compile-time constants
7     enum { n_dims = 2, n_tlev = 2, x=0, y=1, n=0 };
8
9     // ``run-time'' parameters
10    int nx = 3, ny = 3, nt = 4;
11    std::array<double, n_dims> C = {.5, .5};
12
13    // objects representing ``psi'', ``i'' & ``j''
14    std::array<blitz::Array<double, n_dims>, n_tlev> psi;
15    blitz::Range i(1, nx), j(1, ny);
16
17    // mem alloc + zeroing (C++11 for loop syntax)
18    for (auto arr : {&psi[n], &psi[n+1]}) {
19        arr->resize(nx+2, ny+2);
20        *arr = 0;
21    }
22
23 }
```

Example 5: a draft of an advection-equation solver

```
24 // output settings
25 std::cout.setf(std::ios::fixed);
26 std::cout.precision(2);
27
28 // initial condition
29 psi[n](i, j)(0, 0) = 1;
30 std::cout << psi[n](i,j) << std::endl;
31
32 // integration loop
33 for (int t = 0; t < nt; ++t)
34 {
35     // donor-cell advection (constant C>0)
36     psi[n+1](i,j) = psi[n](i,j)
37         - (C[x] * psi[n](i, j) - C[x] * psi[n](i-1, j))
38         - (C[y] * psi[n](i, j) - C[y] * psi[n](i, j-1));
39
40     // output
41     std::cout << psi[n+1](i,j) << std::endl;
42
43     // ``old'' psi becomes ``new'' psi (no copying)
44     blitz::cycleArrays(psi[n], psi[n+1]);
45 }
46 }
```

Example 5: a draft of an advection-equation solver

```
$ ./a.out
(0,2) x (0,2)
[ 1.00 0.00 0.00
  0.00 0.00 0.00
  0.00 0.00 0.00 ]

(0,2) x (0,2)
[ 0.00 0.50 0.00
  0.50 0.00 0.00
  0.00 0.00 0.00 ]

(0,2) x (0,2)
[ 0.00 0.00 0.25
  0.00 0.50 0.00
  0.25 0.00 0.00 ]

(0,2) x (0,2)
[ 0.00 0.00 0.00
  0.00 0.00 0.38
  0.00 0.38 0.00 ]

(0,2) x (0,2)
[ 0.00 0.00 0.00
  0.00 0.00 0.00
  0.00 0.00 0.38 ]
```

blackboard

$$\begin{aligned}\psi_{i,j}^{n+1} = \psi_{i,j}^n & - \left(C_x \psi_{i,j}^n - C_x \psi_{i-1,j}^n \right) \\ & - \left(C_y \psi_{i,j}^n - C_y \psi_{i,j-1}^n \right)\end{aligned}$$

Blitz++

```
psi[n+1](i,j) = psi[n](i,j)
  - (C[x] * psi[n](i, j) -
      C[x] * psi[n](i-1, j))
  - (C[y] * psi[n](i, j) -
      C[y] * psi[n](i, j-1));
```

- ▶ loop-free code
- ▶ execution within a single loop
- ▶ no temporary objects

Summary

Blitz++

- ▶ C++ multi-dimensional array containers
suitable for high-performance computing

Summary

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics

Summary

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics
- ▶ works on array-expressions rather than arrays
 - ~~> no temporary objects, single-loop traversals
(by design, in Fortran it depends on compiler optimisations)

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics
- ▶ works on array-expressions rather than arrays
 - ~~> no temporary objects, single-loop traversals
 - (by design, in Fortran it depends on compiler optimisations)
- ▶ OOP interface (incl. multi-dim. addressing with single object)
 - ~~> no analogues in Fortran

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics
- ▶ works on array-expressions rather than arrays
 - ~~> no temporary objects, single-loop traversals
(by design, in Fortran it depends on compiler optimisations)
- ▶ OOP interface (incl. multi-dim. addressing with single object)
 - ~~> no analogues in Fortran
- ▶ C-pointer API for interfacing C, Fortran, NumPy, ...

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics
- ▶ works on array-expressions rather than arrays
 - ~~> no temporary objects, single-loop traversals
 - (by design, in Fortran it depends on compiler optimisations)
- ▶ OOP interface (incl. multi-dim. addressing with single object)
 - ~~> no analogues in Fortran
- ▶ C-pointer API for interfacing C, Fortran, NumPy, ...
- ▶ numerous alternatives:
 - ▶ C++: NT2, Armadillo, ...
 - ▶ Python: NumPyPy, numexpr, ...
 - ▶ Fortran: built-in arrays (not OO)

Blitz++

- ▶ C++ multi-dimensional array containers suitable for high-performance computing
- ▶ loop-free arithmetics
- ▶ works on array-expressions rather than arrays
 - ~~> no temporary objects, single-loop traversals
 - (by design, in Fortran it depends on compiler optimisations)
- ▶ OOP interface (incl. multi-dim. addressing with single object)
 - ~~> no analogues in Fortran
- ▶ C-pointer API for interfacing C, Fortran, NumPy, ...
- ▶ numerous alternatives:
 - ▶ C++: NT2, Armadillo, ...
 - ▶ Python: NumPyPy, numexpr, ...
 - ▶ Fortran: built-in arrays (not OO)
- ▶ implemented using template meta-programming
 - ~~> much easier to use with C++11

That's it for today.
Thanks for your attention!