

An Embedded Language for Vector Operations in OpenCL

Brian Brennan

Texas Tech University

June 5, 2012

Goal

Generate FEM code to run on various architectures:

Goal

Generate FEM code to run on various architectures:

$$\int_{\Omega} \nabla u \cdot \nabla v dx = \int_{\Omega} f v dx$$

Goal

Generate FEM code to run on various architectures:

$$\int_{\Omega} \nabla u \cdot \nabla v dx = \int_{\Omega} f v dx$$



Architectures



Architectures



- Multicore CPU

Architectures



- Multicore CPU
- NVIDIA GPU

Architectures



- Multicore CPU
- NVIDIA GPU
- AMD GPU

Architectures



- Multicore CPU
- NVIDIA GPU
- AMD GPU
- Intel MIC

Motivation

Why are we interested in GPU computing?

Motivation

Why are we interested in GPU computing?

Pros:

- Very Fast
- Inexpensive
- Many Simple Cores

Motivation

Why are we interested in GPU computing?

Pros:

- Very Fast
- Inexpensive
- Many Simple Cores

Cons:

- Difficult to Program
- Architecture Specific
- Communication with CPU is Costly

Where to Begin?



Where to Begin?



Back to the basics!

Where to Begin?



Back to the basics!

$$y = a * (b - c) - d**2$$

PyOpenCL

What does PyOpenCL have to offer?

PyOpenCL

What does PyOpenCL have to offer?

PyOpenCL = Python + OpenCL

PyOpenCL

What does PyOpenCL have to offer?

PyOpenCL = Python + OpenCL

Python

- Easy to program
- Essential packages such as NumPy, SciPy, PyTrilinos, and many more

PyOpenCL

What does PyOpenCL have to offer?

PyOpenCL = Python + OpenCL

Python

- Easy to program
- Essential packages such as NumPy, SciPy, PyTrilinos, and many more

OpenCL

- General purpose parallel programming
- Vendor-neutral

PyOpenCL Example

```
1 a = numpy.random.rand(50000).astype(numpy.float32)
2 b = numpy.random.rand(50000).astype(numpy.float32)
3
4 ctx = cl.create_some_context()
5 queue = cl.CommandQueue(ctx)
6 a_dev = cl_array.to_device(queue, a)
7 b_dev = cl_array.to_device(queue, b)
8 dest_dev = cl_array.empty_like(a_dev)
9
10 prg = cl.Program(ctx, """
11     __kernel void sum(__global const float *a,
12                      __global const float *b, __global float *c)
13     {
14         int gid = get_global_id(0);
15         c[gid] = a[gid] + b[gid];
16     }
17     """ ). build ()
18
19 prg.sum(queue, a.shape, None, a_dev.data, b_dev.data, dest_dev.data)
```

arraysum.py

Code Generation

We want all the ease of NumPy with the efficiency of PyOpenCL

Code Generation

We want all the ease of NumPy with the efficiency of PyOpenCL

$$y = a + b - 5.0*c - d**2$$



```
1 prg = cl.Program(ctx, """  
2     __kernel void sum(__global const float *a,  
3     __global const float *b, __global const float *c  
4     __global const float *d, __global float *y)  
5     {  
6         int gid = get_global_id(0);  
7         y[gid] = a[gid] + b[gid] - 5.0*c[gid] - d[gid]*d[gid];  
8     }  
9     """).build()
```

kernel.py

Main

```
1 if __name__ == "__main__":
2     ctx = cl.create_some_context()
3     queue = cl.CommandQueue(ctx)
4
5     # declare random numpy arrays here #
6
7     y_dev = cl_array.to_device(queue, a)
8     a_dev = cl_array.to_device(queue, a)
9     b_dev = cl_array.to_device(queue, b)
10    c_dev = cl_array.to_device(queue, c)
11    D_dev = cl_array.to_device(queue, c)
12    Y = Vec( "y_dev" )
13    A = Vec( "a_dev" )
14    B = Vec( "b_dev" )
15    C = Vec( "c_dev" )
16    D = Vec( "d_dev" )
17
18    kernel = assignVector(Y, A + B - 5.0*C - D*D)
19
20    prg = cl.Program( ctx, kernel ).build()
21    prg.op(queue, a.shape, None, y_dev.data, a_dev.data, b_dev.data,
           c_dev.data, d_dev.data)
```

mymain.py

Loop Kernel Output

```
2  __kernel void op( __global float *y_dev,  
4  __global const float *a_dev, __global const float *b_dev,  
6  __global const float *c_dev, __global const float *d_dev )  
    {  
        int gid = get_global_id(0);  
        y[gid] = a[gid] + b[gid] - 5.0*c[gid] - d[gid]*d[gid];  
    }
```

mykernel.py

GPU vs. CPU Results







Figure: Ratio of CPU run times / GPU run times for 1, 2, 3, 4 basic vector operations.

What's
Next?

Loopy

$$\int_{\Omega} \nabla u \cdot \nabla v dx = \int_{\Omega} f v dx$$



```
1 for c in range(num_cells):  
2     for i in range(num_bf):  
3         for j in range(num_bf):  
4             K_loc[c, i, j] = 0.0  
5             for k in range(num_qp):  
6                 K_loc[c, i, j] += jacs_det[c] * qwts[k] \  
7                     * ( (jacs_inv[c, 0, 0] * bgrads[i, k, 0] + jacs_inv[c, 1, 0] *  
8                         bgrads[i, k, 1]) \  
9                         * (jacs_inv[c, 0, 0] * bgrads[j, k, 0] + jacs_inv[c, 1, 0] * bgrads[  
10                            j, k, 1]) \  
11                            + (jacs_inv[c, 0, 1] * bgrads[i, k, 0] + jacs_inv[c, 1, 1] * bgrads[  
12                               i, k, 1]) \  
13                               * (jacs_inv[c, 0, 1] * bgrads[j, k, 0] + jacs_inv[c, 1, 1] * bgrads[
```

poissonloop.py

```
1 for c in range(num_cells):
2     for i in range(num_bf):
3         for j in range(num_bf):
4             K_loc[c, i, j] = 0.0
5             for k in range(num_qp):
6                 K_loc[c, i, j] += jacs_det[c] * qwts[k] \
7                     * ( (jacs_inv[c, 0, 0] * bgrads[i, k, 0] + jacs_inv[c, 1, 0] *
8                         bgrads[i, k, 1] ) \
9                     * (jacs_inv[c, 0, 0] * bgrads[j, k, 0] + jacs_inv[c, 1, 0] * bgrads[
10                        j, k, 1] ) \
11                     + (jacs_inv[c, 0, 1] * bgrads[i, k, 0] + jacs_inv[c, 1, 1] * bgrads[
12                        i, k, 1] ) \
13                     * (jacs_inv[c, 0, 1] * bgrads[j, k, 0] + jacs_inv[c, 1, 1] * bgrads[
14                        j, k, 1] ) )
```

poissonloop.py



In a Galaxy Far Far Away

Expand to large scale software packages:

In a Galaxy Far Far Away

Expand to large scale software packages:



In a Galaxy Far Far Away

Expand to large scale software packages:



In a Galaxy Far Far Away

Expand to large scale software packages:



Trilinos: Sundance Example

```
1 // create symbolic objects for test and unknown functions
  Expr v = new TestFunction(new Lagrange(2));
3 Expr u = new UnknownFunction(new Lagrange(2));

5 // create symbolic differential operators
  Expr dx = new Derivative(0,1);
7 Expr dy = new Derivative(1,1);
  Expr grad = List(dx, dy);

9 // Write symbolic weak equation and Neumann and Robin BCs
11 Expr poisson = Integral(-(grad*v)*(grad*u)-f*v, new GaussQuadrature(2))
    + Integral(top, v/3.0) + Integral(right, v*(rightBCExpr - u));
13

15 // Write essential BCs:
  // Bottom:  $u=x^2$ 
17 EssentialBC bc = EssentialBC(bottom, v*(u - 0.5*x*x), new
    GaussQuadrature(4));

19 // Assemble everything into a problem object, with a specification that
  // Petra be used as the low-level linear algebra representation
21 StaticLinearProblem prob(mesh, poisson, bc, v, u, petra);
```

heat2d.cpp

Questions



Thank you!

- Dr. Andreas Klöeckner

Thank you!

- Dr. Andreas Klöeckner
- Dr. Robert Kirby

Thank you!

- Dr. Andreas Klöeckner
- Dr. Robert Kirby
- Simula