

Programming Cuda and OpenCL

A Case Study Using Modern C++ Libraries

Frameworks

- Cuda
 - NVIDIA
 - Large set of libraries
 - Compute kernels compiled to PTX (low level)
- OpenCL
 - Cross platform
 - API - Boilerplate code
 - Compute kernels compiled to C-like sources (higher level)

Libraries

- (C)MTL4 (The Matrix Template Library)
 - Linear algebra library
 - DSL embedded in c++
 - High level, compile time transformations
 - Cuda
- VexCL (Vector Expression Template Library)
 - Convenient vector and matrix
 - OpenCL
 - Reduce boilerplate code
- ViennaCL (The Vienna Computing Library)
 - Linear Algebra
 - Cuda and OpenCL (only OpenCL in article)
- Thrust
 - Resembles c++ STL
 - Reference point

Ordinary differential equation

- Derivatives with respect to only one variable $\frac{dx}{dt} = \dot{x} = f(x, t), \quad x(0) = x_0.$
- With PDE, surface change over time, ODE particle moving through time
- Eulers method:

$$x(t_0) \quad x(t + \Delta t) \quad x(t + \Delta t) = x(t) + \Delta t f(x(t), t).$$

Odeint

- C++ library for solving ODE's numerically $\frac{dx}{dt} = \dot{x} = f(x, t), \quad x(0) = x_0.$
- Use odeint solving capabilities with gpgpu libraries
- State type, algebra, operation.

$$\dot{x} = -\sigma(x - y), \quad \dot{y} = Rx - y - xz, \quad \dot{z} = -bz + xy.$$

Odeint – Stepper (runge-kutta)

1. `typename State`

The state type.

2. `typename Value = double`

The value type.

3. `typename Deriv = State`

The type representing the time derivative of the state.

4. `typename Time = Value`

The time representing the independent variable - the time.

5. `typename Algebra = typename algebra_dispatcher< State >::algebra_type`

The algebra type.

6. `typename Operations = typename operations_dispatcher< State >::operations_type`

The operations type.

7. `typename Resizer = initially_resizer`

The resizer policy type.

```
template<typename State, typename Value = double, typename Deriv = State,  
        typename Time = Value,  
        typename Algebra = typename algebra_dispatcher< State >::algebra_type,  
        typename Operations = typename operations_dispatcher< State >::operations_type,  
        typename Resizer = initially_resizer>
```

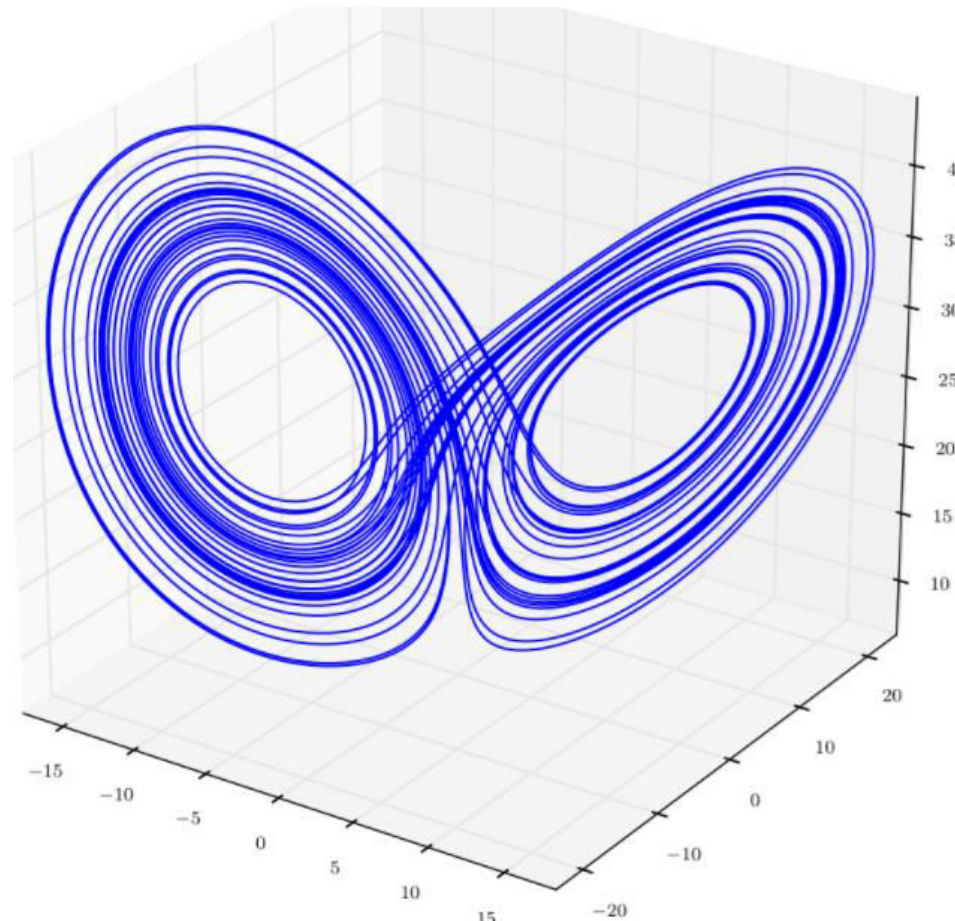
Odeint - integrate

```
size_t integrate_const(Stepper stepper, System system, State & start_state,  
                      Time start_time, Time end_time, Time dt,  
                      Observer observer);
```

```
integrate_const( stepper_type() , lorenz , x , value_type(0.0) , t_max , dt );
```

Lorenz system

$$\dot{x} = -\sigma(x - y), \quad \dot{y} = Rx - y - xz, \quad \dot{z} = -bz + xy.$$



Lorenz - Thrust

```
typedef thrust::device_vector<double> state_type;
```

```
struct lorenz_system {  
    size_t N;  
    const state_type &R;  
    lorenz_system(size_t n, const state_type &r) : N(n), R(r) { }  
    void operator()(const state_type &x, state_type &dxdt, double t) const;  
};
```

```
void lorenz_system::operator()(const state_type &x, state_type &dxdt,  
                              double t) const  
{  
    thrust::for_each(  
        thrust::make_zip_iterator( thrust::make_tuple(  
            R.begin(),  
            x.begin(), x.begin() + N, x.begin() + 2 * N,  
            dxdt.begin(), dxdt.begin() + N, dxdt.begin() + 2 * N ) ),  
        thrust::make_zip_iterator( thrust::make_tuple(  
            R.end(),  
            x.begin() + N, x.begin() + 2 * N, x.end(),  
            dxdt.begin() + N, dxdt.begin() + 2 * N, dxdt.end() ) ),  
        lorenz_functor() );  
}
```

```
struct lorenz_functor {  
    template< class T >  
    __host__ __device__ void operator()( T t ) const {  
        double R = thrust::get<0>(t);  
        double x = thrust::get<1>(t);  
        double y = thrust::get<2>(t);  
        double z = thrust::get<3>(t);  
        thrust::get<4>(t) = sigma * ( y - x );  
        thrust::get<5>(t) = R * x - y - x * z;  
        thrust::get<6>(t) = -b * z + x * y ;  
    }  
};
```

Lorenz - CMTL4

```
1  typedef mtl::dense_vector<double>  vector_type;
2  typedef mtl::multi_vector<vector_type> state_type;
3
4  struct lorenz_system {
5      const vector_type &R;
6      explicit lorenz_system(const vector_type &R) : R(R) { }
7
8      void operator()(const state_type& x, state_type& dxdt, double t) {
9          dxdt.at(0) = sigma * (x.at(1) - x.at(0));
10         dxdt.at(1) = R * x.at(0) - x.at(1) - x.at(0) * x.at(2);
11         dxdt.at(2) = x.at(0) * x.at(1) - b * x.at(2);
12     }
13 };
```

$$\begin{aligned} (\text{ lazy(dxdt.at(0)) } &= \text{ sigma * (x.at(1) - x.at(0)) }) \parallel \\ (\text{ lazy(dxdt.at(1)) } &= \text{ R * x.at(0) - x.at(1) - x.at(0) * x.at(2) }) \parallel \\ (\text{ lazy(dxdt.at(2)) } &= \text{ x.at(0) * x.at(1) - b * x.at(2) }); \end{aligned}$$

150 % overhead with a 3-component vector with 4K entries compared to one vector of size 12K

Lorenz - VexCL

```
1  typedef vex::multivector<double, 3> state_type;
2
3  struct lorenz_system {
4      const vex::vector<double> &R;
5      lorenz_system(const vex::vector<double> &r) : R(r) {}
6
7      void operator()(const state_type &x, state_type &dxdt, double t) const {
8          dxdt(0) = sigma * (x(1) - x(0));
9          dxdt(1) = R * x(0) - x(1) - x(0) * x(2);
10         dxdt(2) = x(0) * x(1) - b * x(2);
11     }
12 };
```

```
dxdt = std::tie(
    sigma * (x(1) - x(0)),
    R * x(0) - x(1) - x(0) * x(2),
    x(0) * x(1) - b * x(2)
);
```

1 Kernel call instead of 3 -> 25% performance gain
(Large systems)

ViennaCL

```
1  typedef fusion::vector<
2      viennacl::vector<double>, viennacl::vector<double>, viennacl::vector<double>
3      > state_type;
4
5  struct lorenz_system {
6      const viennacl::vector<double> &R;
7      lorenz_system(const viennacl::vector<double> &r) : R(r) {}
8
9      void operator()(const state_type &x, state_type &dxdt, double t) const {
10         typedef viennacl::generator::vector<value_type> vec;
11
12         const auto &X = fusion::at_c<0>(x);
13         const auto &Y = fusion::at_c<1>(x);
14         const auto &Z = fusion::at_c<2>(x);
15
16         auto &dX = fusion::at_c<0>(dxdt);
17         auto &dY = fusion::at_c<1>(dxdt);
18         auto &dZ = fusion::at_c<2>(dxdt);
19
20         viennacl::generator::custom_operation op;
21         op.add( vec(dX) = sigma * (vec(Y) - vec(X)) );
22         op.add( vec(dY) = element_prod(vec(R), vec(X)) - vec(Y)
23                 - element_prod(vec(X), vec(Z)) );
24         op.add( vec(dZ) = element_prod(vec(X), vec(Y)) - b * vec(Z) );
25         op.execute()
26     }
27 };
```

Kernel is created once and buffered

Results

	Lorenz attractor	
	Time	T-put
NVIDIA		
Thrust	242.78	105 (71%)
CMTL4	237.91	108 (73%)
VexCL	246.58	104 (70%)
ViennaCL	259.85	99 (66%)
AMD Radeon		
VexCL	149.49	171 (65%)
ViennaCL	148.69	172 (65%)
Intel Core		
Thrust	2 336.14	11 (43%)
VexCL (AMD)	2 329.00	11 (43%)
VexCL (Intel)	2 372.70	11 (42%)
ViennaCL (AMD)	2 322.78	11 (43%)
ViennaCL (Intel)	2 322.39	11 (43%)

