

Solving the Kuramoto Oscillator Model on Random Graphs

Jeffrey Kelling,
Géza Ódor, Sibylle Gemming

12th July 2019

hzdr



HELMHOLTZ
ZENTRUM DRESDEN
ROSSENDORF

Where am I from?

HZDR

HELMHOLTZ
ZENTRUM DRESDEN
ROSSENDORF

- outside of Dresden, Germany



Jürgen-M. Schuler <http://dresden-luftfoto.de>

about me:

- member of computational science group
- background in statistical and theoretical solid state physics

HZDR

1 Introduction

2 Implementation

3 Performance

4 Conclusion

1 Introduction

2 Implementation

3 Performance

4 Conclusion

The Kuramoto Model

- describes a network of coupled oscillators
- system of ordinary differential equations (ODEs)

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \neq j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

⇒ integration to study time-evolution

Using things that already exist

- `boost::numeric::odeint` odeint.com
 - template library of ODE solvers
 - `boost::numeric` supports various vector backends for accelerators:
e.g. `Thrust` (CUDA), `VexCL` (CUDA/OpenCL)

Using things that already exist

- `boost::numeric::odeint` odeint.com
 - template library of ODE solvers
 - `boost::numeric` supports various vector backends for accelerators:
e.g. `Thrust` (CUDA), `VexCL` (CUDA/OpenCL)
- `VexCL`
 - library for offloading vector expressions via CUDA or OpenCL
 - direct support for custom kernels

Using things that already exist



- `boost::numeric::odeint` odeint.com
 - template library of ODE solvers
 - `boost::numeric` supports various vector backends for accelerators:
e.g. `Thrust` (CUDA), `VexCL` (CUDA/OpenCL)
 - `VexCL`
 - library for offloading vector expressions via CUDA or OpenCL
 - direct support for custom kernels

 - we use 4th order Runge-Kutta form `odeint`
- ⇒ computing derivatives remains and is the most time-consuming part

+ offloading vector expressions, which is what `boost::compute` relies on

```
1  std::vector<double> host(N, 2);  
2  vex::vector<double> device(context, host);  
3  
4  device *= device;  
5  
6  vex::copy(device, host);
```

+ offloading vector expressions, which is what `boost::compute` relies on

```
1  std::vector<double> host(N, 2);  
2  vex::vector<double> device(context, host);  
3  
4  device *= device;  
5  
6  vex::copy(device, host);
```

– pseudo single-source: kernel compilation at runtime

– no custom function templates

⇒ have to use custom kernel and inject string to get “template”

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \neq j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- parallel implementations depend on network topology

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \neq j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- parallel implementations depend on network topology
- fully connected graph:
 - N^2 -problem, vectorizable


$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \neq j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- parallel implementations depend on network topology
- fully connected graph:
 - N^2 -problem, vectorizable
- regular lattice / band matrix:
 - stencil integration

Shape of the Network II

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$


- sparse, random graph



Shape of the Network II

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$


- sparse, random graph
 - requires explicit storage network topology
i.e. sparse representation, neighbor lists
 - random neighbor sums



Shape of the Network II

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- sparse, random graph
 - requires explicit storage network topology
i.e. sparse representation, neighbor lists
 - random neighbor sums
- ⇒ **techniques for SIMT vectorization by tuned operation and memory ordering**




1 Introduction

2 Implementation

3 Performance


4 Conclusion

Recap: GPU Architecture



- Single-Instruction-Multiple-Thread (SIMT) workers in lock-step
- vector memory transactions (> 64 byte)

Recap: GPU Architecture



- Single-Instruction-Multiple-Thread (SIMT) workers in lock-step
- vector memory transactions (> 64 byte)
- actually, the same goes for CPU (SIMD + Cache-lines) GPUs just have wider vectors and more simultaneous multi threading (SMT)

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- vectorizing over oscillators j
 - sum over k too short on average ($\lesssim 51$),
too little parallelism
 - avoid need for reduction

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data local to j s is continuous

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data local to j s is continuous
- data in naive neighbor lists would lead to scattered memory access

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data **local to j** is continuous
- data in naive **neighbor lists** would lead to scattered memory access
- data of **remote site k** is at random positions

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data **local to j** is continuous
 - data in naive **neighbor lists** would lead to scattered memory access
 - data of **remote site k** is at random positions
- + no branches, vectorizable expression

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data **local to j** is continuous
 - data in naive **neighbor lists** would lead to scattered memory access
 - data of **remote site k** is at random positions
- + no branches, vectorizable expression
- no predictable data reuse within thread block:
shared memory of not useful, but caches may be
- low computational density, mostly streaming data

Vectorization II: Memory Locality

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- data **local to j** is continuous
 - data in naive **neighbor lists** would lead to scattered memory access
 - data of **remote site k** is at random positions
- + no branches, vectorizable expression
- no predictable data reuse within thread block:
shared memory of not useful, but caches may be
- low computational density, mostly streaming data
- ⇒ maximize memory locality of reads
- ⇒ minimize load imbalances

$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

Memory Layout

#nodes


:

... with n^{th} links;

array of n^{th} links

prefix sum

:



$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

Memory Layout

#nodes


:

... with n^{th} links;

array of n^{th} links


prefix sum

:



$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

Memory Layout




$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

Memory Layout

#nodes	... with n^{th} links;	prefix sum
\vdots	array of n^{th} links	\vdots
1	5 th links	24
2	4 th links	22
5	3 rd links	17
7	2 nd links	10
10	1 st links	0
$j = 0 \quad 1 \quad 2 \quad 3 \quad 4 \quad 5 \quad 6 \quad 7 \quad 8 \quad 9$		



Memory Layout




1 Introduction

2 Implementation


3 Performance

4 Conclusion

long-tailed human brain connectome vs. random graph



Benchmarks



$$\frac{\partial \phi_j(t)}{\partial t} = \omega_j + \sum_{k \text{ NN of } j} \lambda_{jk} \cdot \sin [\phi_k(t) - \phi_j(t)]$$

- profile on tesla P100

- global load efficiency: $\sim 47\%$
 - saturating gross load bandwidth to $\sim 70\%$
- data requests dominant stall reason $\sim 50\%$

⇒ remains memory-latency bound, due to random accesses to neighbors

Conclusion

1 Introduction

2 Implementation

3 Performance

4 Conclusion

- efficient implementation for integration on random graphs
~ 20× improved throughput over single CPU socket.
- easily adaptable to other models: we use it for 2nd order Kuramoto, too

- efficient implementation for integration on random graphs
~ 20× improved throughput over single CPU socket.
- easily adaptable to other models: we use it for 2nd order Kuramoto, too

- *handle randomness on GPU by sorting data to maximise the likelihood of efficient memory access and load balance*

Thank You.