

CUDA-Based Implementation of GSLIB: The Geostatistical Software Library

Daniel Baeza
dabaeza@alges.cl

Oscar Peredo
operedo@alges.cl



The Mining Process

Exploration



Evaluation



Planning



Operation





GSLIB: The Geostatistical Software Library

GSLIB: The Geostatistical Software Library

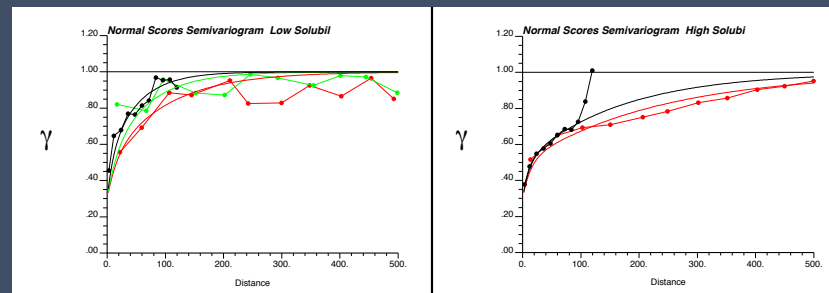
- GSLIB is a software package composed by a set of utilities and applications related with geostatistics
- Full implemented in Fortran 77/90
- Run in OSX, Linux and Windows
- Widely used for academics, researchers, engineers




GSLIB: Geostatistical Software Library and user's guide (1998)
Deutsch, Clayton V, Journé, André G

Variogram calculation with GSLIB

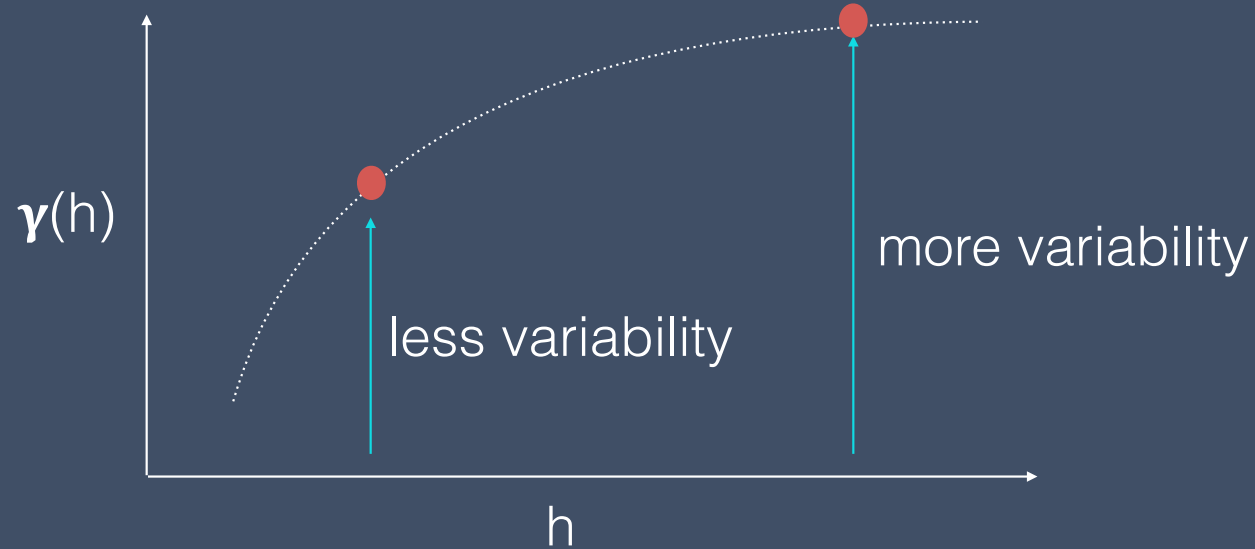
- **gamv** is the GSLIB variogram calculation method
- It's a fundamental tool in geostatistics
- Allow to quantify the spatial variability of a variable
- Used in geostatistical estimation and simulation
- High computational cost




Variogram calculation



$$2\gamma(h) = \frac{1}{N(h)} \sum [z(u) - z(u+h)]^2$$





Variogram computation

Sequential & Parallel Implementations

Sequential implementation

Input:

- (\mathbf{VR}, Ω) : sample data values \mathbf{VR} (m columns) defined in a 3D domain of coordinates Ω
- $nvar$: number of variables ($nvar \leq m$)
- $nlag$: number of lags
- h : lag separation distance
- $ndir$: number of directions
- $\mathbf{h}_1, \dots, \mathbf{h}_{ndir}$: directions
- $\tau_1, \dots, \tau_{ndir}$: geometrical tolerance parameters
- $nvarg$: number of variograms
- $(ivtype_1, ivtail_1, ivhead_1), \dots, (ivtype_{nvarg}, ivtail_{nvarg}, ivhead_{nvarg})$: variogram types

Setup parameters
&
Load data

```

1 Read input parameter file;
2 Read sample data values file;
3  $\beta \leftarrow \text{zeros}(nvar \times nlag \times ndir \times nvarg)$ ;
4 for  $i \in \{1, \dots, |\Omega|\}$  do
5     for  $j \in \{i, \dots, |\Omega|\}$  do
6         for  $id \in \{1, \dots, ndir\}$  do
7             for  $iv \in \{1, \dots, nvarg\}$  do
8                 for  $il \in \{1, \dots, nlag\}$  do
9                      $p_i = (x_i, y_i, z_i) \in \Omega$ ;
10                     $p_j = (x_j, y_j, z_j) \in \Omega$ ;
11                    if  $(p_i, p_j)$  satisfy tolerances  $\tau_{id}$  and  $\|p_i - p_j\| \approx \mathbf{h}_{id} \times il \times h$  then
12                        Save  $(\mathbf{VR}_{i,ivhead_{iv}}$  or  $\mathbf{VR}_{i,ivtail_{iv}})$  and  $(\mathbf{VR}_{j,ivhead_{iv}}$  or  $\mathbf{VR}_{j,ivtail_{iv}})$  according to variogram  $ivtype_{iv}$  into  $\beta$ ;
13  $\gamma \leftarrow$  build variogram using statistics  $\beta$  Write  $\gamma$  in the output file

```

Main computation: Loop over pairs of points

Read computation results

Output: Output file with γ values

Sequential implementation

```
4 for  $i \in \{1, \dots, |\Omega|\}$  do
5   for  $j \in \{i, \dots, |\Omega|\}$  do
6     for  $id \in \{1, \dots, ndir\}$  do
7       for  $iv \in \{1, \dots, nvarg\}$  do
8         for  $il \in \{1, \dots, nlag\}$  do
9            $p_i = (x_i, y_i, z_i) \in \Omega$ ;
10           $p_j = (x_j, y_j, z_j) \in \Omega$ ;
11          if  $(p_i, p_j)$  satisfy tolerances  $\tau_{id}$  and  $\|p_i - p_j\| \approx \mathbf{h}_{id} \times il \times h$  then
12            Save  $(\mathbf{VR}_{i,ivhead_{iv}}$  or  $\mathbf{VR}_{i,ivtail_{iv}})$  and  $(\mathbf{VR}_{j,ivhead_{iv}}$  or  $\mathbf{VR}_{j,ivtail_{iv}})$  according to variogram  $ivtype_{iv}$  into  $\beta$ ;
```

Sequential implementation

```
4 for  $i \in \{1, \dots, |\Omega|\}$  do
5   for  $j \in \{i, \dots, |\Omega|\}$  do
6     for  $id \in \{1, \dots, ndir\}$  do
7       for  $iv \in \{1, \dots, nvarg\}$  do
8         for  $il \in \{1, \dots, nlag\}$  do
9            $p_i = (x_i, y_i, z_i) \in \Omega$ ;
10           $p_j = (x_j, y_j, z_j) \in \Omega$ ;
11          if  $(p_i, p_j)$  satisfy tolerances  $\tau_{id}$  and  $\|p_i - p_j\| \approx \mathbf{h}_{id} \times il \times h$  then
12            Save  $(\mathbf{VR}_{i,ivhead_{iv}}$  or  $\mathbf{VR}_{i,ivtail_{iv}})$  and  $(\mathbf{VR}_{j,ivhead_{iv}}$  or  $\mathbf{VR}_{j,ivtail_{iv}})$  according to variogram  $ivtype_{iv}$  into  $\beta$ ;
```


Sequential implementation

```
4 for  $i \in \{1, \dots, |\Omega|\}$  do
5   for  $j \in \{i, \dots, |\Omega|\}$  do
6     for  $id \in \{1, \dots, ndir\}$  do
7       for  $iv \in \{1, \dots, nvarg\}$  do
8         for  $il \in \{1, \dots, nlag\}$  do
9            $p_i = (x_i, y_i, z_i) \in \Omega;$ 
10           $p_j = (x_j, y_j, z_j) \in \Omega;$ 
11          if  $(p_i, p_j)$  satisfy tolerances  $\tau_{id}$  and  $\|p_i - p_j\| \approx \mathbf{h}_{id} \times il \times h$  then
12            Save  $(\mathbf{VR}_{i,ivhead_{iv}}$  or  $\mathbf{VR}_{i,ivtail_{iv}})$  and  $(\mathbf{VR}_{j,ivhead_{iv}}$  or  $\mathbf{VR}_{j,ivtail_{iv}})$  according to variogram  $ivtype_{iv}$  into  $\beta$ ;
```

COMPUTE_STATISTICS(p_i, p_j)

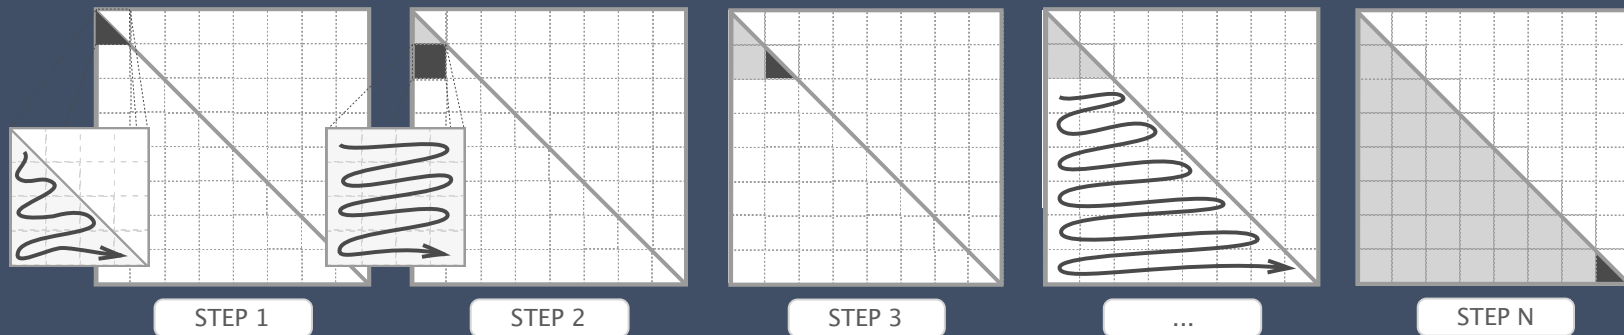
Sequential implementation

```

4 for  $i \in \{1, \dots, |\Omega|\}$  do
5   for  $j \in \{i, \dots, |\Omega|\}$  do
6     for  $id \in \{1, \dots, ndir\}$  do
7       for  $iv \in \{1, \dots, nvar\}$  do
8         for  $il \in \{1, \dots, nlag\}$  do
9            $p_i = (x_i, y_i, z_i) \in \Omega;$ 
10           $p_j = (x_j, y_j, z_j) \in \Omega;$ 
11          if  $(p_i, p_j)$  satisfy tolerances  $\tau_{id}$  and  $\|p_i - p_j\| \approx \mathbf{h}_{id} \times il \times h$  then
12            Save  $(\mathbf{VR}_{i,ivhead_{iv}}$  or  $\mathbf{VR}_{i,ivtail_{iv}})$  and  $(\mathbf{VR}_{j,ivhead_{iv}}$  or  $\mathbf{VR}_{j,ivtail_{iv}})$  according to variogram  $ivtype_{iv}$  into  $\beta$ ;

```

COMPUTE_STATISTICS(p_i, p_j)



Parallel implementation

GPU Kernel

```
1  $id_x = \text{blockId.x} * \text{blockDim.x} + \text{threadId.x}$           /* x threads coord in the GPU grid */
2  $id_y = \text{blockId.y} * \text{blockDim.y} + \text{threadId.y}$           /* y threads coord in the GPU grid */
3 Set and Initialize shared memory in the block
4 __syncthreads()
5  $iter_x = id_x$ 
6  $iter_y = id_y$ 
7 while ( $iter_x$  &  $iter_y \in \text{ChunkPoints}$ )          /* Chunk points belongs thread (x,y) */
8 do
9      $j = iter_x + \frac{|\Omega|}{2}$ 
10     $i = iter_y$ 
11    COMPUTE_STATISTICS( $p_i, p_j$ ) ←
12    store result in shared memory via atomic functions
13    if ( $iter_x > iter_y$ ) then
14         $i = iter_y$ 
15         $j = iter_x$ 
16        COMPUTE_STATISTICS( $p_i, p_j$ ) ←
17        store result in shared memory via atomic functions
18    else
19        if ( $iter_x == iter_y$ ) then
20             $i = iter_y$ 
21             $j = iter_y$ 
22            COMPUTE_STATISTICS( $p_i, p_j$ ) ←
23            store result in shared memory via atomic functions
24         $i = iter_x + \frac{|\Omega|}{2}$ 
25         $j = iter_y + \frac{|\Omega|}{2}$ 
26        COMPUTE_STATISTICS( $p_i, p_j$ ) ←
27        store result in shared memory via atomic functions
28    UP_DATE( $iter_x, iter_y$ )
29 __syncthreads()
30 save Statistics values that are in shared memory into globa memory via atomic functions
```

Output: Array β

Each thread compute
only a couple of correlation values

Parallel implementation

```
1  $id_x$  = blockIdx.x*blockDim.x + threadIdx.x          /* x threads coord in the GPU grid */
2  $id_y$  = blockIdx.y*blockDim.y + threadIdx.y          /* y threads coord in the GPU grid */
3 Set and Initialize shared memory in the block
4 __syncthreads()
5  $iter_x = id_x$ 
6  $iter_y = id_y$ 
7 while ( $iter_x$  &  $iter_y \in ChunkPoints$ )                /* Chunk points belongs thread (x,y) */
8 do
9      $j = iter_x + \frac{|\Omega|}{2}$ 
10     $i = iter_y$ 
11    COMPUTE_STATISTICS( $p_i, p_j$ )
12    store result in shared memory via atomic functions
13    if ( $iter_x > iter_y$ ) then
14         $i = iter_y$ 
15         $j = iter_x$ 
16        COMPUTE_STATISTICS( $p_i, p_j$ )
17        store result in shared memory via atomic functions
18    else
19        if ( $iter_x == iter_y$ ) then
20             $i = iter_y$ 
21             $j = iter_y$ 
22            COMPUTE_STATISTICS( $p_i, p_j$ )
23            store result in shared memory via atomic functions
24         $i = iter_x + \frac{|\Omega|}{2}$ 
25         $j = iter_y + \frac{|\Omega|}{2}$ 
26        COMPUTE_STATISTICS( $p_i, p_j$ )
27        store result in shared memory via atomic functions
28    UP_DATE( $iter_x, iter_y$ )
29 __syncthreads()
30 save Statistics values that are in shared memory into global memory via atomic functions
```

Output: Array β


Parallel implementation

```

1  $id_x = blockIdx.x * blockDim.x + threadIdx.x$           /* x threads coord in the GPU grid */
2  $id_y = blockIdx.y * blockDim.y + threadIdx.y$           /* y threads coord in the GPU grid */
3 Set and Initialize shared memory in the block
4 __syncthreads()
5  $iter_x = id_x$ 
6  $iter_y = id_y$ 
7 while ( $iter_x$  &  $iter_y \in ChunkPoints$ )                /* Chunk points belongs thread (x,y) */
8 do
9      $j = iter_x + \frac{|\Omega|}{2}$ 
10     $i = iter_y$ 
11    COMPUTE_STATISTICS( $p_i, p_j$ )
12    store result in shared memory via atomic functions
13    if ( $iter_x > iter_y$ ) then
14         $i = iter_y$ 
15         $j = iter_x$ 
16        COMPUTE_STATISTICS( $p_i, p_j$ )
17        store result in shared memory via atomic functions
18    else
19        if ( $iter_x == iter_y$ ) then
20             $i = iter_y$ 
21             $j = iter_x$ 
22            COMPUTE_STATISTICS( $p_i, p_j$ )
23            store result in shared memory via atomic functions
24         $i = iter_x + \frac{|\Omega|}{2}$ 
25         $j = iter_y + \frac{|\Omega|}{2}$ 
26        COMPUTE_STATISTICS( $p_i, p_j$ )
27        store result in shared memory via atomic functions
28    UP_DATE( $iter_x, iter_y$ )
29 __syncthreads()
30 save Statistics values that are in shared memory into globa memory via atomic functions

Output: Array  $\beta$ 

```



STEP 1


Parallel implementation

```

1   $id_x = blockIdx.x * blockDim.x + threadIdx.x$           /* x threads coord in the GPU grid */
2   $id_y = blockIdx.y * blockDim.y + threadIdx.y$           /* y threads coord in the GPU grid */
3  Set and Initialize shared memory in the block
4  __syncthreads()
5   $iter_x = id_x$ 
6   $iter_y = id_y$ 
7  while ( $iter_x$  &  $iter_y \in ChunkPoints$ )                /* Chunk points belongs thread (x,y) */
8  do
9       $j = iter_x + \frac{|\Omega|}{2}$ 
10      $i = iter_y$ 
11     COMPUTE_STATISTICS( $p_i, p_j$ )
12     store result in shared memory via atomic functions
13     if ( $iter_x > iter_y$ ) then
14          $i = iter_y$ 
15          $j = iter_x$ 
16         COMPUTE_STATISTICS( $p_i, p_j$ )
17         store result in shared memory via atomic functions
18     else
19         if ( $iter_x == iter_y$ ) then
20              $i = iter_y$ 
21              $j = iter_x$ 
22             COMPUTE_STATISTICS( $p_i, p_j$ )
23             store result in shared memory via atomic functions
24          $i = iter_x + \frac{|\Omega|}{2}$ 
25          $j = iter_y + \frac{|\Omega|}{2}$ 
26         COMPUTE_STATISTICS( $p_i, p_j$ )
27         store result in shared memory via atomic functions
28     UP_DATE( $iter_x, iter_y$ )
29 __syncthreads()
30 save Statistics values that are in shared memory into globa memory via atomic functions

Output: Array  $\beta$ 

```



STEP 1

Parallel implementation


```

1  $id_x = \text{blockId.x} * \text{blockDim.x} + \text{threadId.x}$           /* x threads coord in the GPU grid */
2  $id_y = \text{blockId.y} * \text{blockDim.y} + \text{threadId.y}$           /* y threads coord in the GPU grid */
3 Set and Initialize shared memory in the block
4 __syncthreads()
5  $iter_x = id_x$ 
6  $iter_y = id_y$ 


7 while ( $iter_x$  &  $iter_y \in \text{ChunkPoints}$ )          /* Chunk points belongs thread (x,y) */
8 do
9      $j = iter_x + \frac{|\Omega|}{2}$ 
10     $i = iter_y$ 
11    COMPUTE_STATISTICS( $p_i, p_j$ )
12    store result in shared memory via atomic functions
13    if ( $iter_x > iter_y$ ) then
14         $i = iter_y$ 
15         $j = iter_x$ 
16        COMPUTE_STATISTICS( $p_i, p_j$ )
17        store result in shared memory via atomic functions
18    else
19        if ( $iter_x == iter_y$ ) then
20             $i = iter_y$ 
21             $j = iter_x$ 
22            COMPUTE_STATISTICS( $p_i, p_j$ )
23            store result in shared memory via atomic functions
24         $i = iter_x + \frac{|\Omega|}{2}$ 
25         $j = iter_y + \frac{|\Omega|}{2}$ 
26        COMPUTE_STATISTICS( $p_i, p_j$ )
27        store result in shared memory via atomic functions
28    UP_DATE( $iter_x, iter_y$ )
29 __syncthreads()
30 save Statistics values that are in shared memory into global memory via atomic functions

Output: Array  $\beta$ 

```



STEP 1



STEP 2

Parallel implementation


```

1  idx = blockIdx.x*blockDim.x + threadIdx.x          /* x threads coord in the GPU grid */
2  idy = blockIdx.y*blockDim.y + threadIdx.y          /* y threads coord in the GPU grid */
3  Set and Initialize shared memory in the block
4  __syncthreads()
5  iterx = idx
6  itery = idy


7  while (iterx & itery ∈ ChunkPoints)                /* Chunk points belongs thread (x,y) */
8  do
9      j = iterx +  $\frac{|\Omega|}{2}$ 
10     i = itery
11     COMPUTE_STATISTICS(pi, pj)
12     store result in shared memory via atomic functions
13     if (iterx > itery) then
14         i = itery
15         j = iterx
16         COMPUTE_STATISTICS(pi, pj)
17         store result in shared memory via atomic functions
18     else
19         if (iterx == itery) then
20             i = itery
21             j = itery
22             COMPUTE_STATISTICS(pi, pj)
23             store result in shared memory via atomic functions
24         i = iterx +  $\frac{|\Omega|}{2}$ 
25         j = itery +  $\frac{|\Omega|}{2}$ 
26         COMPUTE_STATISTICS(pi, pj)
27         store result in shared memory via atomic functions
28     UP_DATE(iterx, itery)
29 __syncthreads()
30 save Statistics values that are in shared memory into globa memory via atomic functions

Output: Array  $\beta$ 


```



STEP 1




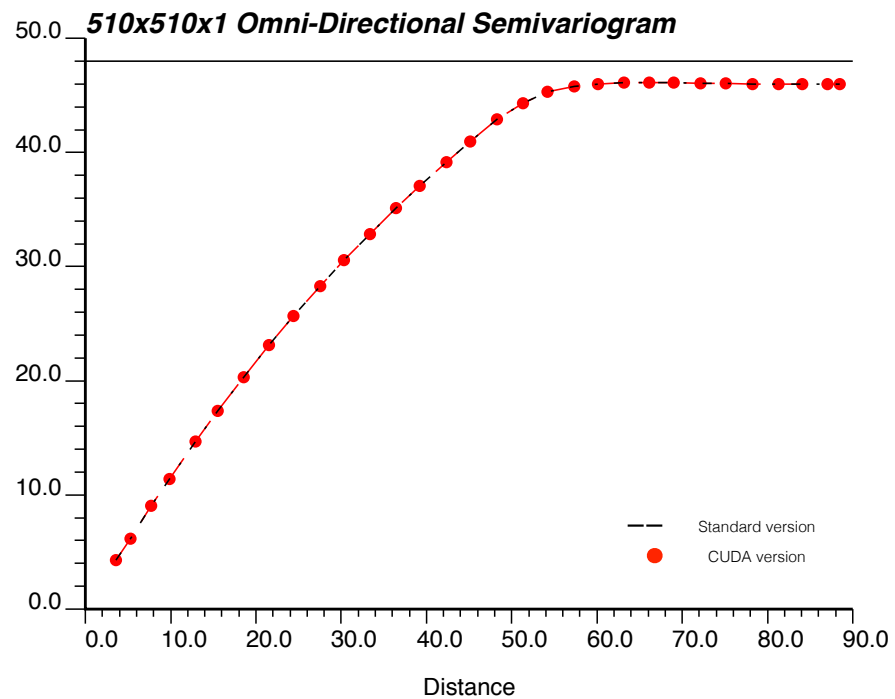
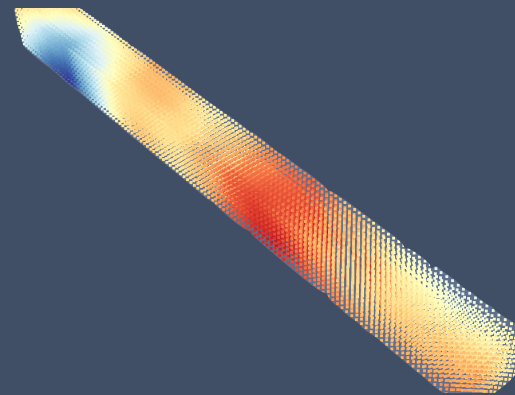

STEP 2



STEP 3




1,443.03	6,817.11	3,878.89	6,066.89	2,011.11
4,887.11	2,860.78	4,537.89	4,384.91	3,459.31
1,880.78	1,680.53	210.78	1,869.93	9,899.11
1,680.54	4,248.81	3,130.91	3,690.36	2,904.03
6,889.89	1,558.04	6,890.43	4,555.89	1,239.03
1,563.04	5,118.81	4,835.95	3,890.36	1,902.31
6,993.83	3,808.88	2,884.99	3,890.36	7,890.36
908.33	4,585.14	3,930.22	3,890.36	8,700.78
4,685.96	8,118.38	623.09	3,210.78	1,781.97
3,918.88	2,884.54	3,478.31	8,302.04	8,822.11
734.69	8,121.84	459.93	2,981.98	9,090.28
3,929.02	5,810.18	989.09	5,108.03	1,883.03
5,890.93	3,985.47	8,392.71	9,920.88	3,107.00
1,781.07	4,818.78	9,930.77	5,091.99	898.31
4,318.73	4,885.18	3,091.99	5,000.21	9,934.93
4,885.93	571.01	946.18	8,398.91	3,470.37
671.09	8,410.81	3,110.91	3,890.00	2,430.37
3,490.31	3,184.08	3,630.90	1,881.93	6,940.37
2,984.03	6,904.51	7,890.83	811.46	1,781.27
5,744.69	4,340.44	8,557.97	8,437.04	6,890.93
4,230.00	5,189.84	9,738.95	955.48	8,912.88
1,252.34	1,811.11	7,093.09	1,890.99	1,755.37
499.11	7,410.18	879.93	6,441.38	2,981.28
7,890.93	1,671.47	3,989.08	3,881.03	3,106.17
1,781.27	5,810.18	9,279.03	3,450.38	4,814.39
6,890.93	8,322.81	3,909.88	3,310.41	1,989.93
8,912.88	1,581.33	3,772.21	8,867.80	1,989.93
1,755.37	1,181.66	8,100.80	2,310.87	1,989.93
2,981.28	4,105.39	2,319.83	2,138.26	1,989.93
3,106.17	4,814.39	8,771.93	6,071.09	1,989.93
		640.70	920.44	1,989.93
		1,989.93		1,989.93

γ  γ 

Dell Precision T7600


- Intel Xeon E5
- 3.3 GHz CPU Frequency
- 10 MB cache
- 16GB RAM (1,6 MHz)
- 1 TB HDD
- Linux









Tesla c2075

- 448 CUDA cores
- 1.15 GHz Frequency of CUDA cores
- 6GB RAM
- 114 GB/sec Memory bandwidth






SpeedUp





Current and future work

- Finish the CUDA version of indicator simulation and gaussian simulation of GSLIB
- Release the first GSLIB-CUDA version with these three methods



Thanks! Questions?