Parallel Space-Time Kernel Density Estimation on a GPU

Tahmid Efaz, Berea College 2019 // tahmidefaz.com // Advisor: Dr. Erik Saule, UNC Charlotte

1 Introduction

Sometimes, the dataset in a Geographical Information System (GIS) comprises of events concentrated in space and time. Visualization of such data may involve representing the latitude and longitude on the x-axis and y-axis against time in a three-dimensional graph.
Space-Time Kernel Density Estimation (STKDE) is a method of calculating the first step of the visualization pipeline in a GIS. The 3-d space in a GIS consists of cubes called voxels and the events occurred in time are represented as points. Using STKDE, we can estimate the density of voxels based on the points around it. The voxel-based algorithm of the STKDE involves iterating through all the voxels in the 3-d space and estimating their density.

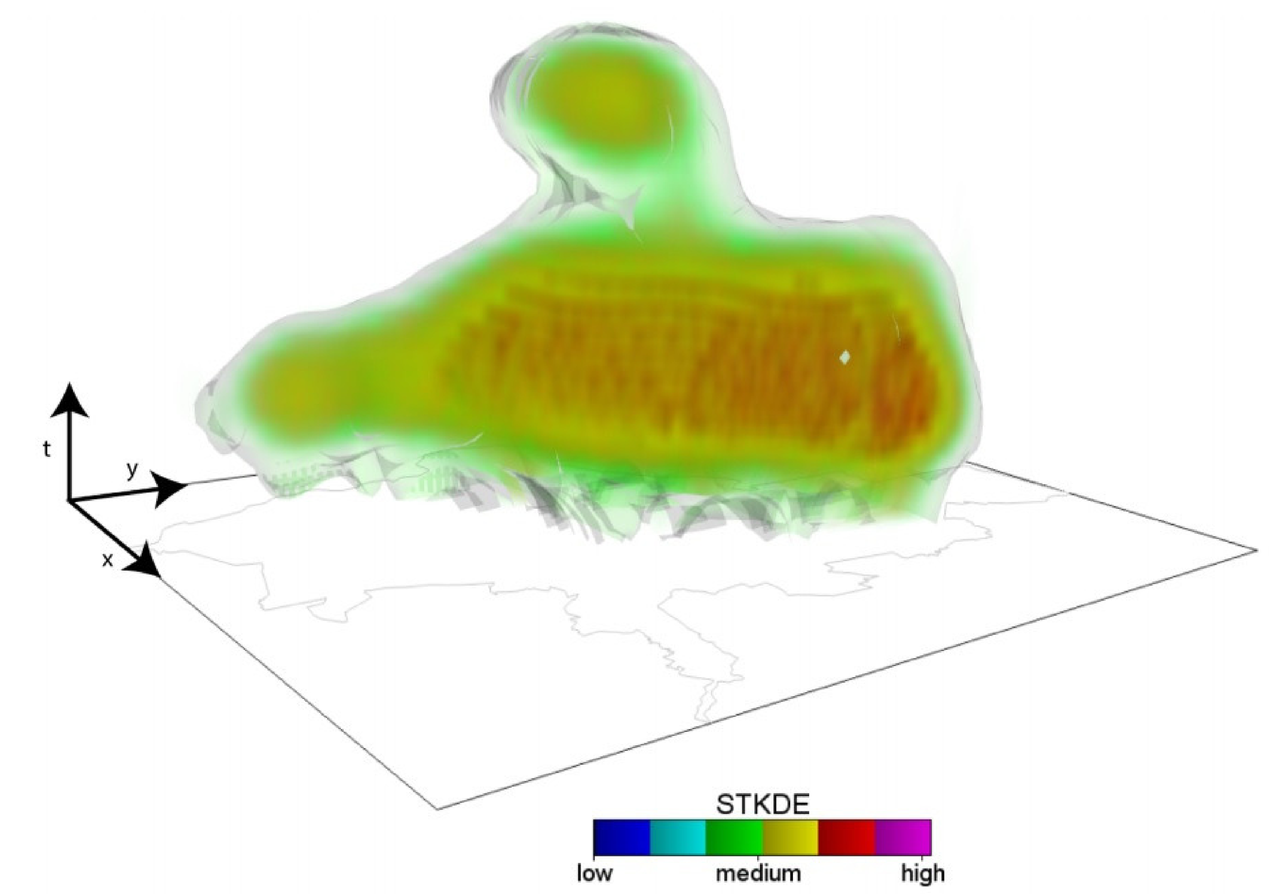

Figure: Example of data visualization in GIS using the STKDE

2 The Issue

Voxel based implementation of the STKDE is very computationally expensive. A CPU implementation of the algorithm can have runtime in hours depending on the number of observations (points).

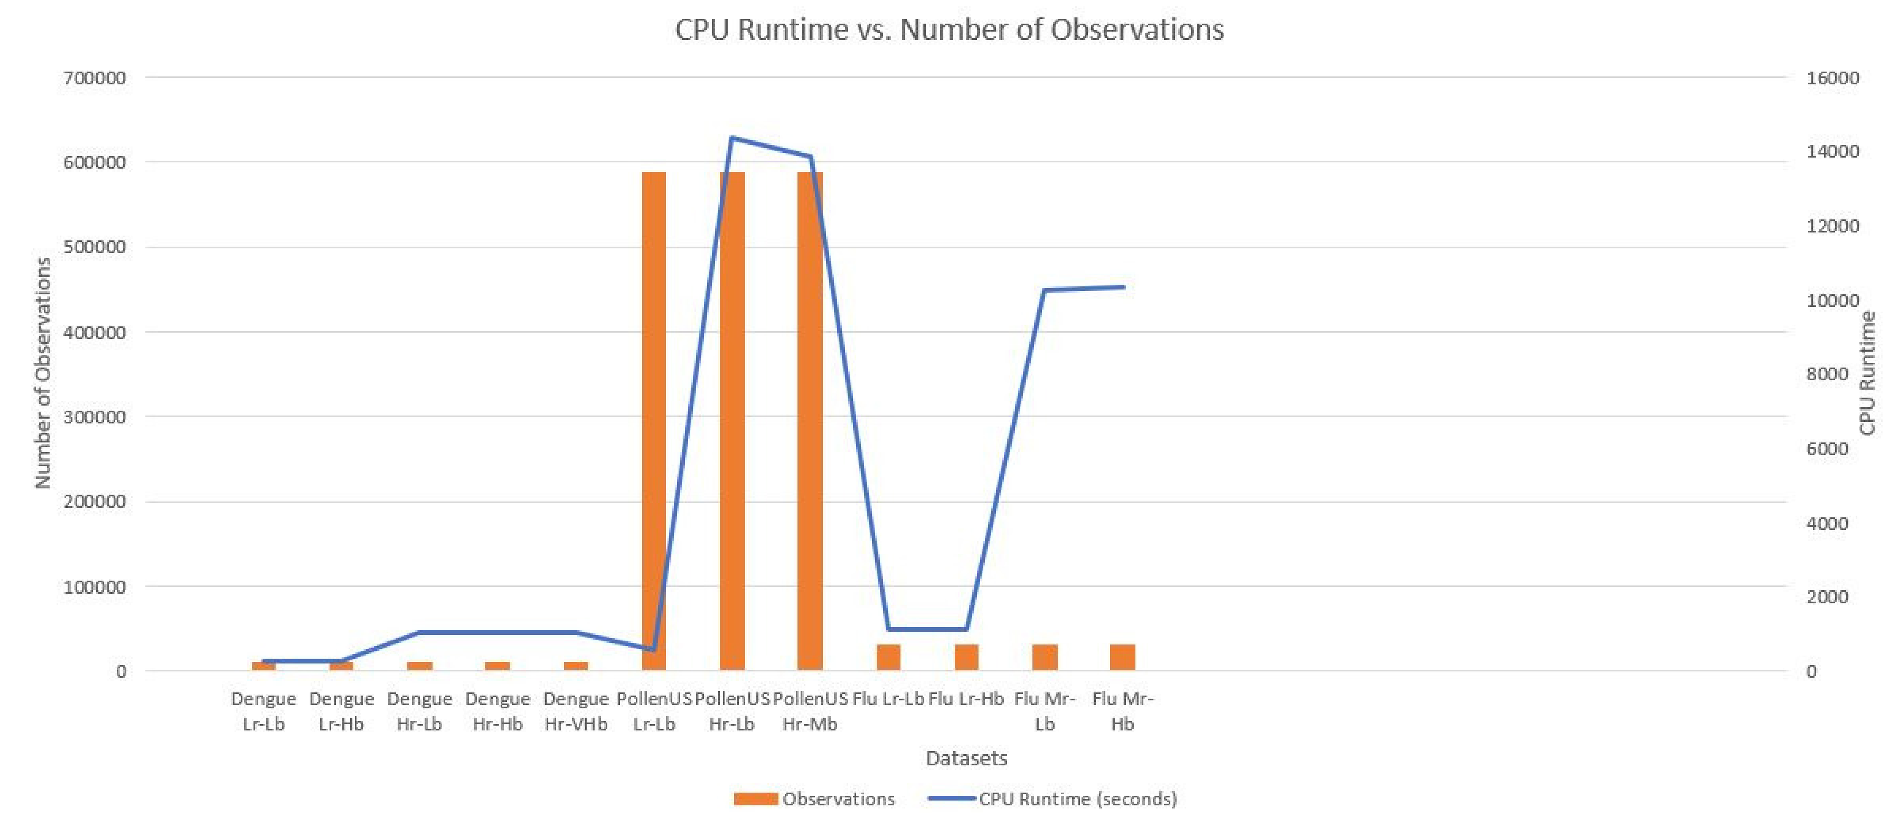

Question:
Can a GPU based implementation of the voxel-based algorithm decrease the runtime?

3 CPU vs. GPU

GPU consists of thousands of cores whereas the CPU only have multiple cores. This allows for immense parallel computations on the GPU.

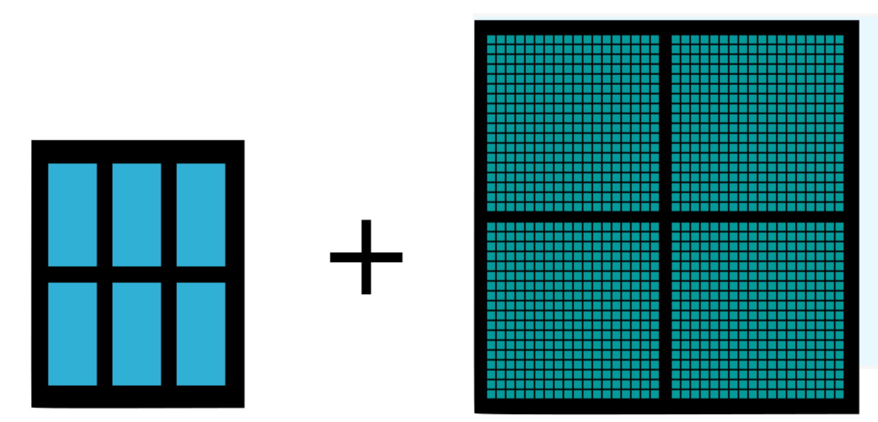

4 Method

- Port the voxel-based implementation of the single-core CPU code to GPU using the NVIDIA CUDA framework.
- Compare the CPU runtime on different datasets with the GPU runtime.

5 Result

Dataset	CPU Runtime (seconds)	GPU Runtime (seconds)	Speedup GPU over CPU	Observations	Domain size in voxels G_x X G_y X G_t	Size (MB)	Bandwidth in voxels H_s	H_t
Dengue Lr-Lb	259.98	4.98	52.17	11056	148x194x728	79	3	1
Dengue Lr-Hb	258.80	5.14	50.39	11056	148x194x728	79	25	1
Dengue Hr-Lb	1024.46	18.06	56.72	11056	294x386x728	315	2	1
Dengue Hr-Hb	1022.13	17.95	56.95	11056	294x386x728	315	50	1
Dengue Hr-VHb	1023.14	17.96	56.96	11056	294x386x728	315	50	14
PollenUS Lr-Lb	559.48	8.94	62.6	588189	131x61x84	2	2	3
PollenUS Hr-Lb	14363.57	192.84	74.48	588189	651x301x84	62	10	3
PollenUS Hr-Mb	13863.23	192.16	72.14	588189	651x301x84	62	25	7
Flu Lr-Lb	1105.80	19.05	58.03	31478	117x308x851	117	1	1
Flu Lr-Hb	1107.98	18.93	58.54	31478	117x308x851	117	2	3
Flu Mr-Lb	10270.69	176.60	58.16	31478	233x615x1985	1085	2	3
Flu Mr-Hb	10345.78	174.91	59.15	31478	233x615x1985	1085	4	7

We saw immense speed up of the runtime on the GPU. GPU code was 50 to 70 times faster depending on the data set.

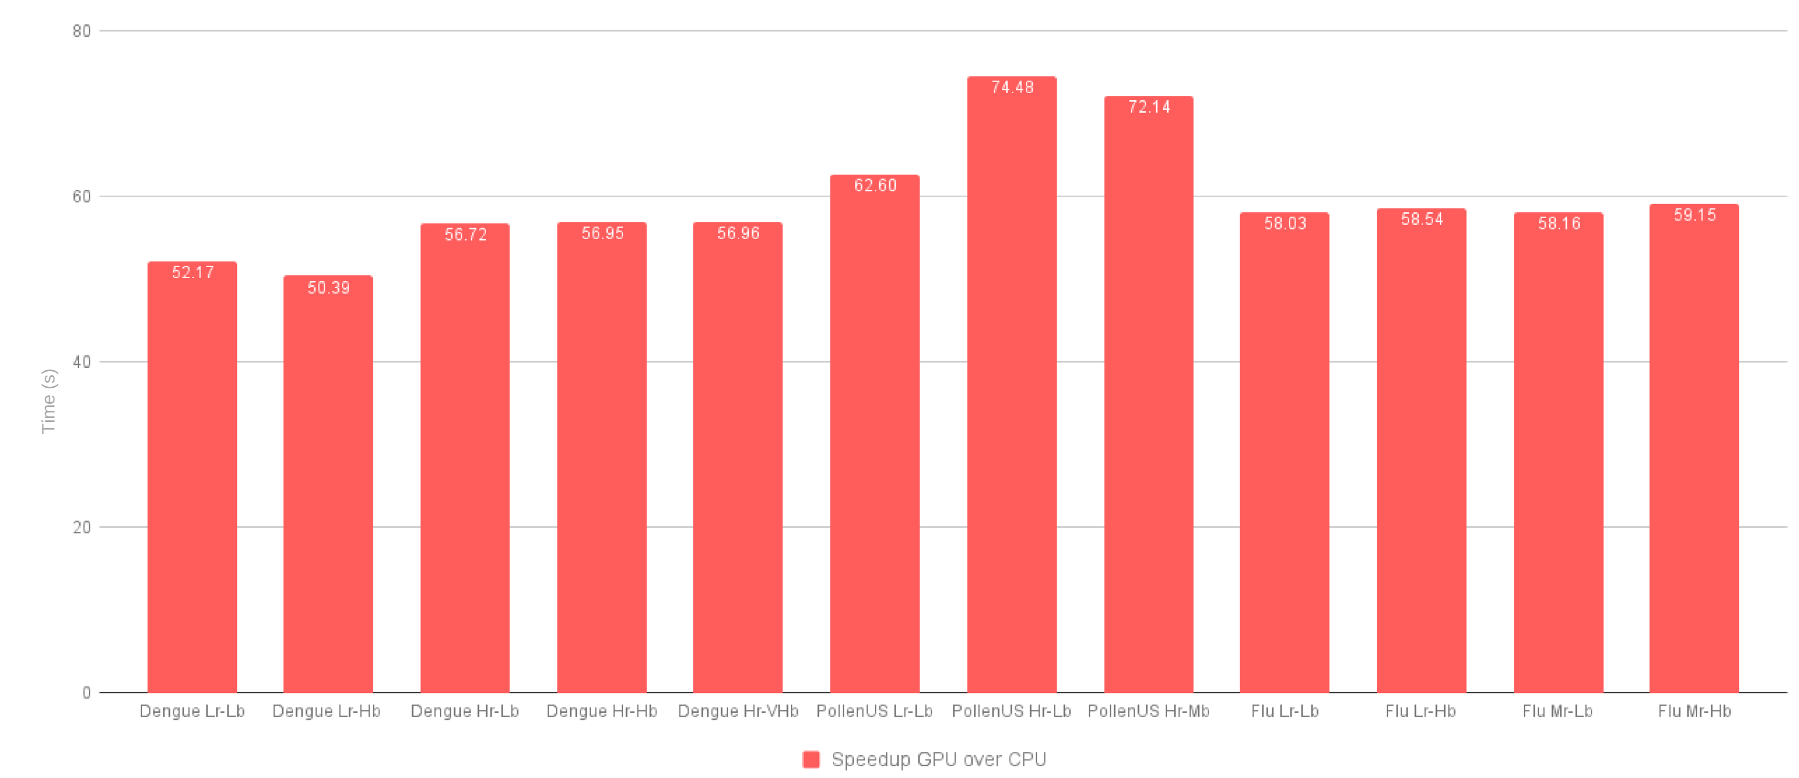

Runtime comparison on the PollenUS Hr-Lb data:

Acknowledgements

- Paper: "Parallel Space-Time Kernel Density Estimation" by {Erik Saule, Dinesh Panchananam, Alexander Hohl, Wenwu Tang, Eric Delmelle} UNC Charlotte
- NVIDIA CUDA framework
- UNC Charlotte University Research Computing

UNC CHARLOTTE

BEREA COLLEGE
GOD HAS MADE OF ONE BLOOD ALL PEOPLES OF THE EARTH