

Accelerating the Gabor Transform with a GPU for SAR Image Compression

Conner McInnes, Shadi Alawneh

cmcinn@oakland.edu, shadialawneh@oakland.edu

Department of Electrical and Computer Engineering

School of Engineering and Computer Science, Oakland University

Introduction

The objective of this research was to illustrate the potential gains using CUDA could have over other languages for an implementation of the Gabor transform used for synthetic aperture radar (SAR) image compression. The Gabor transform is a very versatile transform often utilized in artificial intelligence (AI) algorithms and in signal analysis. However, its versatility does not end there, it also can be featured in a compression algorithm for images. In particular a technique exists that utilizes the Gabor Transform that makes use of Fast Fourier transforms (FFT's) to compress the size of SAR images. This performs an optimized lossy compression which reduces the size of the image while retaining enough information for it to still be useful. Issues arise when the time taken to calculate and compress these images is accounted for. Even though they offer the ability to take up less space through minimal lossy compression, performing this compression can be taxing on a central processing unit (CPU). Numerous libraries exist in a multitude of programming languages that can assist in performing this transform, but one in particular stands out due to its ability to compute the transform extremely quickly. CUDA's cuFFT library is based upon the open source library called Fastest Fourier Transform in the West (FFTW) available in C++, but unlike FFTW, cuFFT is able to leverage the power of the graphics processing unit (GPU) to compute FFT's. One issue with cuFFT is that the time that is required to transfer information between the CPU and GPU is very significant. This causes a paradox where even though cuFFT is able to compute FFT's much quicker than FFTW, the data transfer time will bottleneck its performance. As a result, this leaves the use of CUDA's GPU accelerated library as a potential solution that requires benchmarking to optimize its use.

Design

Two programs would be necessary for a proper benchmark. One was created in C++ using the FFTW library, the other in CUDA using cuFFT. The programs were made to be as structurally similar as possible. They both utilized the same C++ code when not calling functions from their respective libraries. Multiple images were used at a number of different resolutions to directly compare how resolution affected performance. The GPU used to perform these runs was an Nvidia Geforce RTX 2080 Ti running at 2070 MHz on the clock, and 7800 MHz on the memory, both of which were boost locked. The CPU used was an AMD Ryzen 3900x running at an all core 4.2 GHz overclock at 1.275 Volts. Each program was ran five times alternating between each run using the chrono library to collect the time elapsed values. The average of the five runs was taken for each category. Figures 1 – 4 depict the test images that were Fast Fourier transformed for this benchmark. Tables 1 – 3 illustrate the data gathered on the sampled image, resolution, and runtime.

Results



Figure 1: Lab



Figure 2: Pantheon



Figure 3: Tent



Figure 4: Castle

Image	Resolution	cuFFT Time(sec)	FFTW Time(sec)	% Faster
Lab	1920 x 1280	0.000036	0.131609	369587%
Pantheon	640 x 426	0.000035	0.019579	55841%
Pantheon	1280 x 853	0.000026	0.080414	306824%
Pantheon	1920 x 1280	0.000035	0.131356	377360%
Tent	640 x 427	0.000032	0.013402	42045%
Tent	1280 x 853	0.000020	0.080337	405643%
Tent	1920 x 1280	0.000039	0.132188	340591%
Tent	3000 x 2000	0.000033	0.170687	520286%
Castle	640 x 439	0.000022	0.021397	98958%
Castle	1280 x 879	0.000019	0.061001	317615%
Castle	1920 x 1318	0.000022	0.153738	711650%
Castle	5705 x 3917	0.000085	N / A	N / A

Table 1: Execution of Plan

Image	Resolution	cuFFT Time(sec)	FFTW Time(sec)
Lab	1920 x 1280	0.151330	0.134483
Pantheon	640 x 426	0.152355	0.021616
Pantheon	1280 x 853	0.156869	0.082607
Pantheon	1920 x 1280	0.150899	0.134224
Tent	640 x 427	0.152653	0.015591
Tent	1280 x 853	0.158958	0.082551
Tent	1920 x 1280	0.150712	0.135054
Tent	3000 x 2000	0.151648	0.173599
Castle	640 x 439	0.152623	0.023422
Castle	1280 x 879	0.156284	0.063296
Castle	1920 x 1318	0.153273	0.157001
Castle	5705 x 3917	0.181563	N / A

Table 2: Creation of Plan

Image	Resolution	cuFFT Time(sec)	FFTW Time(sec)
Lab	1920 x 1280	0.424639	0.176053
Pantheon	640 x 426	0.324040	0.025484
Pantheon	1280 x 853	0.359570	0.100223
Pantheon	1920 x 1280	0.420902	0.175412
Tent	640 x 427	0.323964	0.020099
Tent	1280 x 853	0.364874	0.099208
Tent	1920 x 1280	0.408482	0.175953
Tent	3000 x 2000	0.648835	0.290701
Castle	640 x 439	0.303318	0.028166
Castle	1280 x 879	0.356702	0.080740
Castle	1920 x 1318	0.425278	0.202571
Castle	5705 x 3917	1.480622	N / A

Table 3: FFT Function

An important note is that for the runs that used the 5705 x 3917 image, the C++ program was unable to compile due to the limits on the sequential memory available to Visual Studio. The programs will need to be redesigned to use a linked list to access the more available non-sequential RAM to allow that test. Nevertheless, it is easily able to be deduced from Table 1 that the cuFFT program is much faster at performing the computations required of the FFT. These results show the extreme promise that taking a GPU accelerated approach to FFT's can have for the Gabor transform in speeding up the computational time. However, Table 2 and 3 highlight the bottleneck that data transfer has on the GPU. If methods can be devised that can allow the GPU to access and allocate information from the CPU quicker, or when very large images are supplied, CUDA could become extremely valuable in speeding up the computational time of the Gabor transform.