Mapping and Characterizing Global-Scale Human Settlements Using HPC

Dilip R. Patlolla
patlolladr@ornl.gov

Harini Sridharan
sridharanh@ornl.gov

Vincent C. Paquit
paquitvc@ornl.gov

Jeanette E. Weaver
weaverje@ornl.gov

Mark A. Tuttle
tuttlema@ornl.gov

Anil M. Cheriyadat*
cheriyadatam@ornl.gov

ABSTRACT

Spatiotemporal analytics derived from large-scale analysis of satellite imagery is a key component in the models used for global population mapping, risk analysis, and critical infrastructure assessment. Making informed national-level decisions based on image derived analytics require processing petabytes of high resolution satellite image data. The computational process often involves extracting, representing, and classifying visual data patterns to detect different types of man-made structures [1]. Computing numerical feature descriptors from raw image data is an expensive but key step in the computational process. The scale of the computational process and the need to repeat the process at time stipulated intervals can easily saturate conventional CPUs. As a feasible solution, we have built a computational system based on GPGPU architecture for deriving big image data analytics [7]. In this study, we demonstrate an efficient GPU implementation for multi-scale feature extraction and exploit multi-GPU architecture for large-scale image analysis for characterizing human settlements.

General Terms

Algorithms, Performance, Design, Experimentation

Keywords

Algorithms, GPU , Parallel Programming, Feature Extraction, Settlement, Detection

1. COMPUTATIONAL PROCESS

Our first step involves automatically detecting settlement (buildings, houses, etc.) in sub-meter resolution satellite images. Given the variations in the appearance and density of settlements this problem requires intelligent representation of pixel data to uniquely characterize the pixels belonging to settlement regions. We begin by normalizing the input image data to have zero mean and unit standard deviation. The color image is then converted to a gray band. To carefully design the data access and execution strategies that can fully exploit GPU's parallel computational capabilities we divide the image into non-overlapping square pixel-blocks. For each of the pixel-block we compute feature descriptors at several scales based on edge orientations and co-occurrence of pixel intensity patterns [7]. From the edge orientations measurements, we computed histogram of orientation (HOG) and from the co-occurrence of pixel patterns we derived statistics from gray level co-occurrence matrix (GLCM). To capture the characteristics of spatial context around each pixel-block we computed these features at different scales by considering different sized pixel windows around each block. The multi-scale feature descriptors computed from the image data are classified into settlement or non-settlement set based on a linear support vector machine (SVM) classifier.

The second step of the process involves further classifying the settlement regions to fine-grained categories such as residential, industrial, commercial, or mobile home parks. In our current system we further scan the settlement regions to identify mobile home park locations. Mobile home parks present a challenging class of scenes for automated detection from satellite images given their large within class variations (mobile home parks in California and Louisiana might look different in terms of appearance and spatial arrangements) and small between class variances (mobile home park units might look more like residential neighborhoods). For this step we detect lines from the image by connecting pixels having similar orientations. We then examine line patterns and generate numerical feature descriptors to represent line patterns for an image patch. We examine overlapping image patches and identify patches containing mobile home parks by applying a pre-learned classifier. The line extraction and line pattern representation are computationally intensive steps which are parallelized in the GPU processing allowing us to detect lines in parallel as instances of multiple threads and compute feature representation for each image in parallel manner.

Today, commodity GPUs provides hundreds of cores at their disposal [8] offering ample opportunities to accelerate algorithms by exploiting the parallel processing power. Since the introduction of GPUs, the general trend in increasing

*Corresponding Author.
Authors Affiliation:
Oak Ridge National Laboratory

Supercomputing 13, 'November 17-22, 2013, Denver, USA

the computational power is to exploit parallelism by executing single instruction set on multiple data (SIMD) on many processors. Recently, computer vision algorithms exploiting SIMD architecture have shown to provide speedups of several orders compared to their CPU counterparts [2, 4, 5, 6]. The data parallel nature of the feature descriptor generation and the pixel processing opens up a wide range of opportunities for the acceleration and scaling of geospatial processing algorithms. We have used the Compute Unified Device Architecture (CUDA) to harness the tremendous computational power and memory bandwidth of the GPU. The CUDA programming model conveniently allowed the processing of multiple sub-tasks in parallel over the pixel blocks at several scales.

It is also important to make careful choices in the data layout to fully exploit the parallel computational capabilities provided by the GPU processors. In our implementation the image raster is read using the GDAL (Geospatial Data Abstraction Library) [4] is initially transferred to the global memory. A subset of the image data, required by each thread block, is loaded into a local memory called shared memory and all the threads within the block can access this memory at high speed in a highly parallel manner. The constant memory is read-only and provides faster and parallel data access for CUDA kernel execution. In our implementation this memory is used to store the pre-computed SVM classification parameters. As all the threads in a thread block access the same SVM parameters in Constant memory, the reads are parallelized resulting in a single read.

2. RESULTS

Our computational system was tested on a benchmarking dataset consisting of 0.6-meter spatial resolution satellite imagery representing urban cities, rural villages, and sparse structures from different parts of the world. The system yielded an accuracy of 85

To further evaluate the performance of the GPU-based settlement characterization for mobile homes, a benchmark dataset comprising of mobile home park scenes from 10 different cities in the US was compiled. The database comprised of 610 positive samples and 7,813 negative samples of size 300mx300m collected from USDAâĂŹs National Agriculture Imagery Program (NAIP) imagery having a spatial resolution of one meter. The average overall accuracy is 96.80

3. ACKNOWLEDGEMENTS

Authors would like to thank the support and close supervision of Budhendra Bhaduri. This manuscript has been authored by employees of Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, Tennessee 37831-6285, managed by UT-Battelle, LLC for the U. S. Department of Energy under contract no. DEAC05-00OR22725.

4. REFERENCES

[1] A. Cheriyadat, E. Bright, D. Potere, and B. Bhaduri. Mapping of settlements in high-resolution satellite imagery using high performance computing. *GeoJournal*, 69(1-2):119–129, 2007.

[2] N. Cornelis and L. V. Gool. Fast scale invariant feature detection and matching on programmable graphics hardware. In *IEEE Computer Vision and Pattern Recognition, CVPR* , June 2008.

[3] GDAL. Geospatial data acquisition library http://www.gdal.org.

[4] A. Plaza, S. Bernabe, P. Marpu and J. Benediktsson. A new parallel tool for classification of remotely sensed imagery. *Computers & Geosciences*, 2012.

[5] M. J. Thurley and V. Danell. Fast morphological image processing open-source extensions for GPU processing with CUDA. *IEEE Journal of Selected Topics in Signal Processing,*, PP(99):1, 2012.

[6] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek. Empowering visual categorization with the GPU. *IEEE Transactions on Multimedia*, 13(1):60–70, 2011.

[7] D. Patlolla, J. Weaver, E. Bright, and A. Cheriyadat. Accelerating Satellite Image Based Large-Scale Settlement Detection with GPU. *ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data*, 43–51, 2012.