

# Improving performance of geospatial data processing using OpenMP

Natalija Stojanović<sup>1</sup>, Dragan Stojanović<sup>2</sup>

**Abstract** – In this paper, high-performance processing of geospatial data on multicore computer architecture is considered. We use OpenMP (Open Multi-Processing) to implement parallel processing of common Geographic Information Systems (GIS) algorithms on multi-core processors. This approach is used to accelerate application for map-matching computation over large spatial datasets consisting of moving points and road network segments. Also, the algorithm for viewshed analysis is parallelized and performed on a multicore system. Experimental evaluation validates our approach and shows feasibility of high-performance computing in GIS.

**Keywords** – high-performance processing, GIS, multicore processors, OpenMP, map-matching, viewshed analysis

## I. INTRODUCTION

Many today's data and computing intensive applications require more processing power than even before. Computer games, database searching, Web search engines, financial and economic forecasting, climate modeling, medical imaging, etc. are application domains from a broad range of applications that demand accelerating and performance improvements [1]. One of the approaches for achieving performance improvements, which has been used in last decades, is to decrease the size of components used to build computers. The technology progress has made possibility to put billions of transistors on a single chip. As the size of transistors has decreased their speed also has increased and therefore the speed of overall integrated circuit has to be increased. But, increasing of the transistor speed, as a major approach of computer performance advance, has limited due to heat-dissipation. As a consequence of heat dissipation integrated circuit could become unreliable [2]. These limitations have led to alternative strategies for creating more powerful computers, and have made the use of parallelism to become one of the crucial solutions for accelerating applications. For example, chip manufactures have started to produce processors with several computing units, called cores, on one chip, that have independent control and have access to the same, shared memory. Therefore, using multicore processor has made each desktop or a laptop computer a commodity parallel system. It is important that application software must be able to make effective use of parallelism that is present in available hardware resources. Software

developers can't expect that the increasing of computer power can be automatically used by their application programs. This is due to the fact that there is not appropriate automatic transformer from a sequential program to a parallel program that runs efficiently on the new architectures. Research in the area of parallelizing compilers has shown that in many situations it is not possible to extract enough parallelism from sequential programs. Instead, a software developer has to transform its software to run efficiently on new architectures.

Therefore, OpenMP (Open-Multiprocessing) was developed to enable creation of parallel programs for shared-memory multiprocessor platforms. OpenMP represents a set of compiler directives, library routines, and environment variables that provide programmer to tell the compiler which instructions to execute in parallel. Also, OpenMP provides programmer to define how to distribute them among the threads that will run the code.

Many modern software applications, which are developed to model real world, process a large amount of data and thus cause long execution times on today's computers. One such application domain that could significantly benefit from parallel processing is Geographic Information System (GIS) that includes processing and analytics of massive geospatial data (Big Geospatial Data). In this paper, we use OpenMP parallel implementation for a map matching, a process of integration of raw vehicle trajectory data to underlying road network providing richer semantics in registered movement. Also, the algorithm for viewshed analysis is parallelized and accelerated on a multicore system. The performances of our implementations have been evaluated. Experimental results are examined and discussed.

## II. HIGH-PERFORMANCE COMPUTING IN GIS

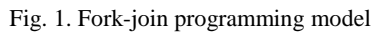
There are many ways to parallelize applications in order to improve their performance. Just as there are several different classes of parallel hardware, so there are different models of parallel programming. Therefore, OpenMP (Open-Multiprocessing) was developed to enable development of parallel programs for shared-memory multiprocessor platforms. Also, there are other high performance parallel/distributed techniques and methods. For example, MPI (Message Passing Interface) programming library use multiple computers connected by high-speed LAN for distributed computing, while OpenCL (Open Computing Language) and CUDA (Compute Unified Device Architecture) enable general-purpose applications to access massively parallel Graphics Processing Units (GPU) for non-graphical computing [2].

OpenMP is not a new programming language. Rather it is a notation that can be added to a sequential program in C, C++

<sup>1</sup> Natalija Stojanović is with the Faculty of Electronic Engineering, University of Nis, Aleksandra Medvedeva 14, 18000 Niš, Serbia  
E-mail: natalija.stojanovic@elfak.ni.ac.rs

<sup>2</sup> Dragan Stojanović is with the Faculty of Electronic Engineering, University of Nis, Aleksandra Medvedeva 14, 18000 Niš, Serbia  
E-mail: dragan.stojanovic@elfak.ni.ac.rs

OpenMP supports the so-called fork-join programming model. This approach assumes that the program begins execution as a single thread, just like an ordinary sequential program. The thread that executes this piece of code is called initial (main) thread. Every time this thread encounters OpenMP parallel construct during program execution, it creates a set of threads. It then becomes a parent thread and cooperates with other threads of the program execution. (Fig. 1). At the end of parallel construct only the initial (parental, main) thread continues execution, interrupting the execution of others.



In our previous paper [8] we evaluate two parallel/distributed architectures and programming models: MPI (Message Passing Interface) over network of workstations (NoW) and CUDA (Compute Unified Device

### III. GEOSPATIAL DATA PROCESSING USING OPENMP

Determining the shortest distance between a point and a segment depends on a value of scalar dot between vectors NM and MP, i.e. scalar dot between vectors MN and NP and is calculated using following expression:

where  $A = y_1 - y_2$ ,  $B = x_2 - x_1$  and  $C = x_1 y_2 - x_2 y_1$ .

Viewshed analysis is one of the common algorithms in terrain spatial analysis and use of terrain models. A viewshed

is the area(s) of the land/terrain surface that is visible from one or more viewpoints. Viewshed analysis is used in a variety of applications such as locating radio and TV transmitters and cellular communication base stations for maximum coverage, site selection for the forest lookout stations and determining the areas of the resort that would be visible from the new restaurant to determine its scenic qualities.

The location and height of the viewpoint can change the size and shape of a viewshed. Viewing parameters can also include the viewing angle, the search distance and even the tree height. The process of deriving viewsheds is called viewshed or visibility analysis. A viewshed analysis requires two input datasets. The first dataset represents a point layer containing one or more defined viewpoints. The second input dataset represents the land surface and can be defined as a DEM (Digital Elevation Model) - a raster spatial data containing elevation at spatial grid points, or a TIN (Triangulated Irregular Network).

The viewshed analysis is performed following a series of steps [9]. First, the location of the viewpoint is connected by a line of sight or a ray to all specified location in the terrain. Second, a set of intermediate points is derived along each sightline. These intermediate points are determined from the intersection between the sightline and the grid lines of the elevation raster (DEM). Finally the computing algorithm examines the elevations of the intermediate points, looking for locations that are higher and thus visible from the viewpoint or not (Fig. 3.) Higher points along this sightline obscure the lower points that are behind them. The above procedure is repeated for each grid point in the elevation raster as a target, or only for specified locations of interest. In the parallel OpenMP implementation the viewshed computing for the whole set of grid points is evenly distributed among several threads (e.g. 2, 4, 6, 8, 10, etc.) and thus can be performed on separate cores in a multi core architecture. Using OpenMP the effective use of a multicore architecture is achieved.

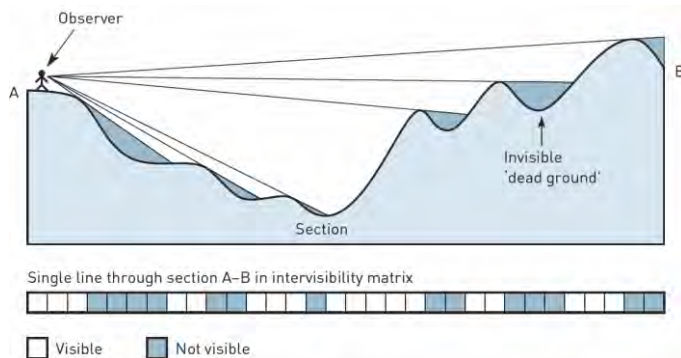


Fig. 3. Ray tracing for visibility analysis [9]

Through repeated 'ray tracing' a viewshed map is built. This is a Boolean raster map that classifies raster cells into the visible and invisible groups and indicates which areas of terrain are visible or not from the location of interest. This map is more effectively interpreted and understood by its visualization over the terrain surface (Fig. 4.).

Various algorithms have been developed for computing viewsheds and appropriate tools are available in several contemporary commercial and open-source GIS software. Such tools provide basic capabilities for setting up the parameters for viewshed analysis but usually do not provide choices for algorithms and information of adopted algorithm.




Fig. 4. Results of viewshed analyses for single point

Both map matching and viewshed analysis algorithms are time consuming and computationally intensive operations that are performed by applying relatively simple geometric operations over massive vector or raster geospatial datasets. As such these operations are well suited for speeding up by employing parallel processing techniques in the computation. In the next section we will examine application of OpenMP on common GIS algorithms, map matching and viewshed analysis to show improvement in performance and feasibility of effective use of multicore architecture for high performance computing in GIS.

#### IV. EXPERIMENTAL EVALUATION

In order to estimate proposed parallel solutions we used the speedup as measure. Speedup ( $S_p$ ) is used to compare the execution time of observed computation on  $p$  processors ( $T_p$ ) to the execution time of sequential version ( $T_1$ ) i.e.

$$S_p = \frac{T_1}{T_p}$$

As the values of speedup values are indication of processor's performance, the maximum obtained speedup can be used to locate the best number of threads to be used for corresponding application on observed architecture. In our experiments for both applications we used Intel Core 2 Duo T5870 2GHz CPU, 4GB of RAM, and Intel i7-2670QM 2,2GHz, 4GB of RAM multicore architecture. Intel Core 2 Duo has two cores, while Intel Core i7 has 4 cores with implemented hyper-threading (HT) technology. With hyper-threading technology, the instructions from two threads are interleaved in the processor pipeline. Circuits of the processor which store architectural state of processor are duplicated, but the main execution resources are shared and not duplicated. For the development of parallel application on multicore computer, Visual Studio 2008 with included support for OpenMP is used.

In the case of map-matching application two datasets are used and stored in corresponding files. The first dataset represents the set of segments (each polyline consists of

straight line segments) where each record contains segment identifier and (x, y) coordinates of the segment endpoints. The segment dataset contains 7035 records. The second dataset contains the set of moving points where each record contains point identifier, sequence number, the timestamp, (x, y) coordinate of point and the current speed. The point dataset consists of 329273 records. In the case of viewshed application DEM with 1201x1201 points is used.

Performance of map-matching application as well as viewshed application, obtained for different number of threads on the Intel Core 2 Duo and Intel Core i7 are shown in Table I and Table II. In both tables the execution time (in seconds) and speedup for corresponding thread count are shown.

**TABLE I**  
**EXPERIMENTAL RESULTS FOR MAP-MATCHING**  
algorithm

Core 2 Duo			i7		
threads	T <sub>p</sub>	S <sub>p</sub>	threads	T <sub>p</sub>	S <sub>p</sub>
2	181,006	1,824	2	71,585	1,954
4	186,489	1,771	4	42,661	3,278
6	191,441	1,725	6	31,016	4,509
8	222,351	1,485	8	25,823	5,416

**TABLE II**  
**EXPERIMENTAL RESULTS FOR VIEWSHED**  
ALGORITHM

Core 2 Duo			i7		
threads	T <sub>p</sub>	S <sub>p</sub>	threads	T <sub>p</sub>	S <sub>p</sub>
2	221,185	1,991	2	149,575	1,826
4	250,743	1,756	4	84,276	3,240
6	282,675	1,558	6	73,443	3,719
8	245,047	1,797	8	65,632	4,162

For Core 2 Duo the best performance results are when using two threads, while performance decreases for both algorithms when engaging more than two threads. For the Core i7 processor, hyper-threading technology enhances the speedup for applications with two to eight threads. After that, it stagnates. The maximum speedup in the case of map-matching is obtained for eight threads and its value is 5,416. For the viewshed algorithm the maximum obtained speedup is for eight threads and its value is 4,162. The speedup values more than 4 are obtained due to HT technology.

From Tables I and II it can be concluded that the obtained speedup makes this OpenMP implementations a candidate for sequential acceleration solutions. As can be concluded from the experimental evaluation, the use of OpenMP implementation for map-matching algorithm and viewshed

analysis is a major challenge, and as a result gives satisfactory performance improvements.

## V. CONCLUSION

With advances in remote sensing, geosensor networks and pervasive positioning, the amount of geospatial data that needs to be processed and analyzed has exploded in recent years. It leads to a rising interest in using high performance parallel and distributed techniques for large scale geospatial data processing and analysis. This paper shows that the application of OpenMP parallel processing technique to effectively employ multicore computer architecture could improve the performance in executing common computation and data intensive GIS algorithms, such as map matching and viewshed analysis, over large scale raster and vector geospatial data.

The future research will consider the usage of other HPC techniques and platforms in GIS, such as cloud computing (MapReduce/Hadoop) and the adaptation of various GIS algorithms to cloud infrastructure.

## VI. ACKNOWLEDGMENTS

Research presented in this paper is funded by Ministry of education, science and technological development, Republic of Serbia as part of the project 'Environmental Protection and Climate Change Monitoring and Adaptation', Nr. III-43007.

## REFERENCES

- [1] S.Patel, Wen-mei W. Hwu, "Accelerator Architectures", IEEE Micro, Vol. 28,4(2008), 4-12.
- [2] P. Pacheco, "An Introduction to Parallel Programming", Morgan Kaufman, 2011.
- [3] B. Chapman, G. Jost, and R. Van Der Pas. "Using OpenMP: portable shared memory parallel programming" MIT Press, 2008.
- [4] A. Clematis, M. J. Mineter, and R. Marciano, "High performance computing with geographical data", *Special issue Parallel Computing*, pp. 1275-1279, 2003.
- [5] S. Shekhar, "High Performance Computing with Spatial Datasets", Invited talk, International Workshop on High Performance and Distributed GIS, San Jose, California, 2010.
- [6] S. Akhter, K. Aida and Y. Chemin, "GRASS GIS on High Performance Computing with MPI, OpenMP and Ninf-G Programming Framework", Proceeding of ISPRS 2010, Japan.
- [7] J. Zhang, "Towards Personal High-Performance Geospatial Computing (HPC-G): Perspectives and a Case Study", International Workshop on High Performance and Distributed GIS, pp. 3-10, San Jose, California, 2010.
- [8] N. Stojanovic, D. Stojanovic, "High-performance computing in GIS: techniques and applications", International Journal of Reasoning-based Intelligent Systems - IJRIS, accepted for publication - in press, 2013  
<http://www.inderscience.com/info/ingeneral/forthcoming.php?code=ijris>
- [9] I. Heywood, S. Cornelius and S. Carver, "An Introduction to Geographical Information Systems, 4th Edition", Prentice Hall, 2012.