A Parallel Sweep Line Algorithm for Visibility Computation

Chaulio R. Ferreira[1], Marcus V. A. Andrade[1],

Salles V. G. Magalhes[1], W. R. Franklin[2], Guilherme C. Pena[1]

[1]Departamento de Informática – Universidade Federal de Viçosa (UFV)

Campus da UFV – 36.570-000 – Viçosa – MG – Brazil

[2]Rensselaer Polytechnic Institute – Troy – NY – USA

```
{chaulio.ferreira,marcus,salles,guilherme.pena}@ufv.br,
                   wrf@ecse.rpi.edu
```

Abstract. *This paper describes a new parallel raster terrain visibility (or viewshed) algorithm, based on the sweep-line model of [Van Kreveld 1996]. Computing the terrain visible from a given observer is required for many GIS applications, with applications ranging from radio tower siting to aesthetics. Processing the newly available higher resolution terrain data requires faster architectures and algorithms. Since the main improvements on modern processors come from multi-core architectures, parallel programming provides a promising means for developing faster algorithms. Our algorithm uses the economical and widely available shared memory model with OpenMP. Experimentally, our parallel speedup is almost linear. On 16 parallel processors, our algorithm is up to 12 times faster than the serial implementation.*

1. Introduction

An important group of Geographical Information Science (GIS) applications on terrains concerns visibility, i.e., determining the set of points on the terrain that are visible from some particular observer, which is usually located at some height above the terrain. This set of points is known as *viewshed* [Franklin and Ray 1994] and its applications range from visual nuisance abatement to radio transmitter siting and surveillance, such as minimizing the number of cellular phone towers required to cover a region [Ben-Moshe et al. 2007], optimizing the number and position of guards to cover a region [Magalhães et al. 2011], analysing the influences on property prices in an urban environment [Lake et al. 1998] and optimizing path planning [Lee and Stucky 1998]. Other applications are presented in [Champion and Lavery 2002].

Since visibility computation is quite compute-intensive, the recent increase in the volume of high resolution terrestrial data brings a need for faster platforms and algorithms. Considering that some factors (such as processor sizes, transmission speeds

and economic limitations) create practical limits and difficulties for building faster serial computers, the parallel computing paradigm has become a promising alternative for such computing-intensive applications [Barney et al. 2010]. Also, parallel architectures have recently become widely available at low costs. Thus, they have been applied in many domains of engineering and scientific computing, allowing researchers to solve bigger problems in feasible amounts of time.

In this paper, we present a new parallel algorithm for computing the viewshed of a given observer on a terrain. Our parallel algorithm is based on the (serial) sweep line algorithm firstly proposed by [Van Kreveld 1996], which is described in Section 2.3.3. Comparing to the original algorithm, our new algorithm achieved speedup of up to 12 times using 16 parallel cores, and up to 3.9 times using four parallel cores.

2. Related Work

2.1. Terrain representation

In what follows, our region of interest is small compared to the radius of the earth, thus, for this discussion the earth can be considered to be flat.

A *terrain* τ is a $2\frac{1}{2}$ dimensional surface where any vertical line intersects τ in at most one point. The terrain is usually represented approximately either by a *Triangulated Irregular Network (TIN)* or a *Raster Digital Elevation Model (DEM)* [Li et al. 2005]. A TIN is a partition of the surface into planar triangles, i.e., a piecewise linear triangular spline, where the elevation of a point p is a bilinear interpolation of the elevations of the vertices of the triangle containing the projection of p. On the other hand, a DEM is simply a matrix storing the elevations of regularly spaced positions or posts, where the spacing may be either a constant number of meters or a constant angle in latitude and longitude. In this paper, we will use the DEM representation because of its simpler data structure, ease of analysis, and ability to represent discontinuities (cliffs) more naturally. Finally, there is a huge amount of data available as DEMs.

2.2. The viewshed problem

An *observer* is a point in space from where other points (the *targets*) will be visualized. Both the observer and the targets can be at given heights above τ, respectively indicated by h_o and h_t. We often assume that the observer can see only targets that are closer than the *radius of interest*, ρ. We say that all cells whose distance from O is at most ρ form the *region of interest* of O. A target T is visible from O if and only if the distance of T from O is, at most, ρ and the straight line, the *line of sight*, from O to T is always strictly above τ; see Figure 1.
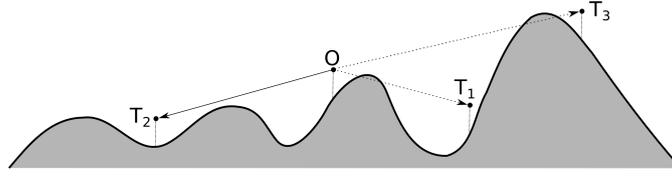
Figure 1. Targets' visibility: T_1 **and** T_3 **are not visible but** T_2 **is.**

The *viewshed* of O is the set of all terrain points vertically below targets that can be seen by O; formally,

$$viewshed(O) = \{p \in \tau \mid \text{the target above } p \text{ is visible from } O\}$$

with ρ implicit. The viewshed representation is a square $(2\rho + 1) \times (2\rho + 1)$ bitmap with the observer at the center.

Theoretically, determining whether a target T is visible from O requires verifying all points in the line of sight connecting O to T. But since τ is represented with a finite resolution, only points close to the rasterized line segment connecting the projections of O and T onto the horizontal plane will be verified. Which points those might be, is one difference between competing algorithms, as the ones we will describe in Section 2.3. The visibility depends on the line segment rasterization method used, see Figure 2, and how the elevation is interpolated on those cells where the segment does not intersect the cell center.
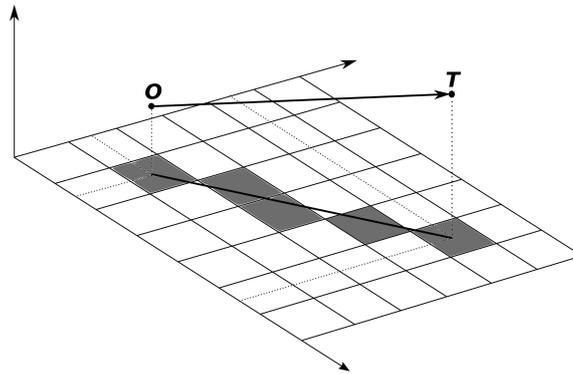
Figure 2. The rasterization of the line of sight projection.

The visibility of a target above a cell c_t can be determined by checking the slope of the line connecting O and T and the cells' elevation on the rasterized segment. More precisely, suppose the segment is composed of cells $c_0, c_1, \cdots, c_t$ where c_0 and c_t correspond to the projections of O and T respectively. Let α_i be the slope of the line connecting O to c_i, that is,

$$\alpha_i = \frac{\zeta(c_i) - (\zeta(c_0) + h_o)}{dist(c_0, c_i)} \tag{1}$$

where $\zeta(c_0)$ and $\zeta(c_i)$ are, respectively, the elevation of cells c_0 and c_i and $dist(c_0, c_i)$ is the 'distance' (in number of cells) between these two cells. The target on c_t is visible if and only if the slope $\frac{\zeta(c_t)+h_t-(\zeta(c_0)+h_o)}{dist(c_0,c_t)}$ is greater than α_i for all $0 < i < t$. If yes, the corresponding cell in the viewshed matrix is set to 1; otherwise, to 0.

2.3. Viewshed algorithms

Different terrain representations call for different algorithms. A TIN can be processed by the algorithms proposed by [Cole and Sharir 1989] and [De Floriani and Magillo 2003]. For a DEM, we can point out [Van Kreveld 1996] and RFVS [Franklin and Ray 1994], two very efficient algorithms. Another option for processing DEMs is the well-known R3 algorithm [Shapira 1990]. Although this one is not as efficient as the other two, it has higher accuracy and may be suitable for small datasets.

These three algorithms differ from each other not only on their efficiency, but also on the visibility models adopted. For instance, R3 and Van Kreveld's algorithms use a center-of-cell to center-of-cell visibility, that is, a cell c is visible if and only if the ray connecting the observer (in the center of its cell) to the center of c does not intersect a cell blocking c. On the other hand, RFVS uses a less restrictive approach where a cell c may be considered visible if its center is not visible but another part of c is.

Therefore, the viewsheds obtained by these methods may be different. Without knowing the application and having a model for the terrain's elevation between the known points, it is impossible to say which one is better. Some applications may prefer a viewshed biased in one direction or the other, while others may want to minimize error computed under some formal terrain model. For instance, since Van Kreveld's algorithm presents a great tradeoff between efficiency and accuracy [Fishman et al. 2009], it may be indicated for applications that require a high degree of accuracy. On the other hand, if efficiency is more important than accuracy, the RFVS algorithm could be preferred.

Considering that each one of these algorithms might be suitable for different applications, we will describe them briefly in the next sections.

2.3.1. R3 algorithm

The R3 algorithm provides a straightforward method of determining the viewshed of a given observer O with a radius of interest ρ. Although it is considered to have great accuracy [Franklin et al. 1994], this algorithm runs in $\Theta(n^{\frac{3}{2}})$, where $n = \Theta(\rho^2)$. It works as follows: for each cell c inside the observer's region of interest, it uses the digital differential analyzer (DDA) [Maćiorov 1964] to determine which cells the line of sight (from O to the center of c) intersects. Then, the visibility of c is determined by calculating the

slope of all cells intersected by this line of sight, as described in Section 2.2. In this process, many rules to interpolate the elevation between adjacent posts may be used, such as average, linear, or nearest neighbour interpolations.

2.3.2. RFVS algorithm

RFVS [Franklin and Ray 1994] is a fast approximation algorithm that runs in $\Theta(n)$. It computes the terrain cells' visibility along rays (line segments) connecting the observer (in the center of a cell) to the center of all cells in the boundary of a square of side $2\rho + 1$ centered at the observer (see Figure 3(a)). In each column, it tests the line of sight against the closest cell. Although a square was chosen for implementation simplicity, other shapes such as a circle would also work.

RFVS creates a ray connecting the observer to a cell on the boundary of this square, and then rotates it counter-clockwise around the observer to follow along the boundary cells (see Figure 3(a)). The visibility of each ray's cells is determined by walking along the segment, which is rasterized following [Bresenham 1965]. Suppose the segment is composed of cells $c_0, c_1, \cdots, c_k$ where c_0 is the observer's cell and c_k is a cell in the square boundary. Let α_i be the slope of the line connecting the observer to c_i determined according to Equation (1) in Section 2.2. Let μ be the highest slope seen so far when processing c_i, i.e., $\mu = \max\{\alpha_1, \alpha_2, \cdots, \alpha_{i-1}\}$. The target above c_i is visible if and only if the slope $(\zeta(c_i) + h_t - (\zeta(c_0) + h_o))/dist(c_0, c_i)$ is greater than μ. If yes, the corresponding cell in the viewshed matrix is set to 1; otherwise, to 0. Also, if $\alpha_i > \mu$ then μ is updated to α_i. We say that a cell c_i blocks the visibility of the target above c_j if c_i belongs to the segment $\overline{c_0 c_j}$ and α_i is greater or equal to the slope of the line connecting the observer to the target above c_j.
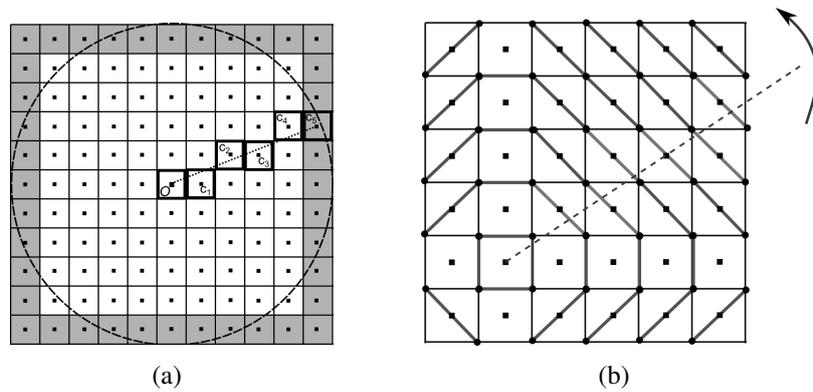
2.3.3. Van Kreveld's algorithm

Van Kreveld's algorithm [Van Kreveld 1996] is another fast viewshed algorithm. According to [Zhao et al. 2013], it has accuracy equivalent to the R3 algorithm, while running in $\Theta(n \, log \, n)$. Its basic idea is to rotate a sweep line around the observer and compute the visibility of each cell when the sweep line passes over its center (see Figure 3(b)). For that, it maintains a balanced binary tree (the *agenda*) that stores the slope of all cells currently being intersected by the sweep line, keyed by their distance from the observer. When this sweep line passes over the center of a cell c, the *agenda* is searched to check c's visibility. More specifically, this algorithm works as follows:

For each cell, it defines three types of events: *enter*, *center*, and *exit* events to indi-

cate, respectively, when the sweep line starts to intersect a cell, passes over the cell center and stops to intersect a cell. The algorithm creates a list E containing these three types of events for all cells inside the region of interest. The events are then sorted according to their azimuth angle.

To compute the viewshed, the algorithm sweeps the list E and for each event it decides what to do depending on the type of the event:

- If it is an *enter* event, the cell is inserted into the *agenda*.
- If it is an *center* event of cell c, the *agenda* is searched to check if it contains any cell that lies closer to the observer than c and has slope greater or equal to the slope of the line of sight to c; if yes, then c is not visible, otherwise it is.
- If it is an *exit* event, the cell is removed from the *agenda*.

(a)   (b)

Figure 3. Viewshed algorithms: (a) RFVS; (b) Van Kreveld - adapted from [Fishman et al. 2009].

2.3.4. Parallel viewshed algorithms

Parallel computing has become a mainstream of scientific computing and recently some parallel algorithms for viewshed computation have been proposed. [Zhao et al. 2013] proposed a parallel implementation of the R3 algorithm using Graphics Processing Units (GPUs). The RFVS algorithm was also adapted for parallel processing on GPUs by [Osterman 2012]. [Chao et al. 2011] proposed a different approach for parallel viewshed computation using a GPU, where the algorithm runs entirely within the GPUs visualization pipeline used to render 3D terrains. [Zhao et al. 2013] also discuss other parallel approaches.

However, we have not found any previous work proposing a parallel implementation of Van Kreveld's algorithm. In fact, [Zhao et al. 2013] stated that "a high degree of sequential dependencies in Van Kreveld's algorithm makes it less suitable to exploit

parallelism". In Section 3 we show how we have overcome this difficulty and describe our parallel implementation of Van Kreveld's sweep line algorithm.

2.4. Parallel Programming models

There are several parallel programming models, such as distributed memory/message passing, shared memory, hybrid models, among others [Barney et al. 2010]. In this work, we used the shared memory model, where the main program creates a certain number of tasks (*threads*) that can be scheduled and carried out by the operating system concurrently. Each thread has local data, but the main program and all threads share a common address space, which can be read from and written to asynchronously. In order to control the concurrent access to shared resources, some mechanisms such as locks and semaphores may be used. An advantage of this model is that there is no need to specify explicitly the communication between threads, simplifying the development of parallel applications.

For the algorithm implementation, we used OpenMP (*Open Multi-Processing*) [Dagum and Menon 1998], a portable parallel programming API designed for shared memory architectures. It is available for C++ and Fortran programming languages and consists of a set of compiler directives that can be added to serial programs to influence their run-time behaviour, making them parallel.

3. Our parallel sweep line algorithm

As described in Section 2.3.3, Van Kreveld's algorithm needs information about the cells intersected by the sweep line. It maintains these information by processing the *enter* and *exit* events to keep the *agenda* up to date as the sweep line rotates. Therefore, processing a *center* event is dependent upon all earlier *enter* and *exit* events.

In order to design a parallel implementation of this algorithm, this dependency had to be eliminated. We did that by subdividing the observer's region of interest in S sectors around the observer, O (see Figure 4(a), where $S = 8$). Our idea is to process each one of these sectors independently using Van Kreveld's sweep line algorithm, such that it can be done in parallel.

More specifically, consider sector s defined by the interval $[\alpha, \beta)$, where α and β are azimuth angles. Let a and b be the line segments connecting O to the perimeter of its region of interest, with azimuth angles α and β, respectively (see Figure 4(a)). To process s, the algorithm creates rays connecting O to all cells on the perimeter of the region of interest that are between (or intersected by) a and b (see Figure 4(b)). These rays are rasterized using the DDA method [Maćiorov 1964] and the events related to the intersected cells are inserted into s's own list of events, E_s. Since the grid cells are convex,

this process inserts into E_s the events for all cells inside s or intersected by a or b. The inserted cells are shown in Figure 4(b).
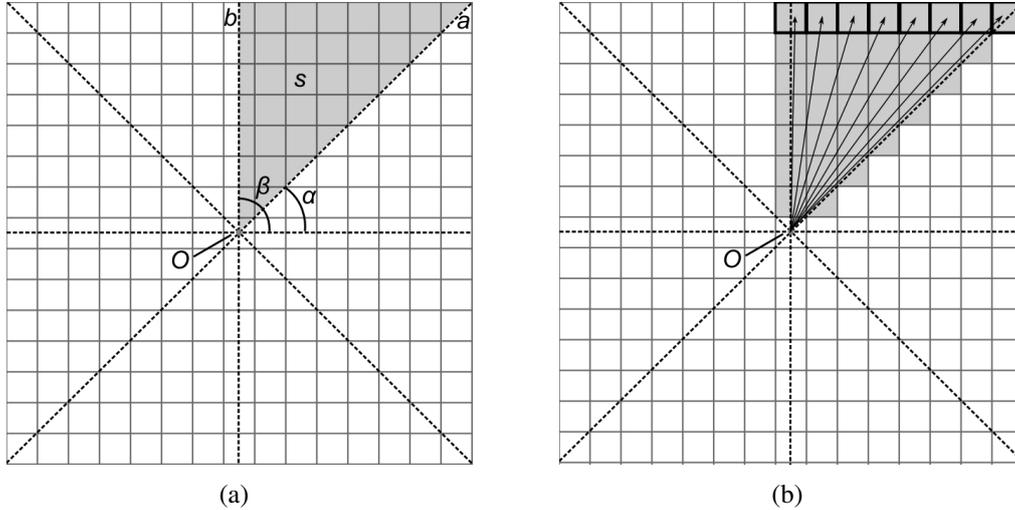
Figure 4. Sector definition: (a) The subdivision of the region of interest and the sector s, defined by the interval $[\alpha, \beta]$; (b) The cells in the perimeter of the region of interest, the rays used to determine which cells are intersected by s and the cells inserted into E_s (shaded cells).

Then, the algorithm sorts E_s by the events' azimuth angles and sweeps it in the same manner as Van Kreveld's algorithm. Note that, because we have distributed the events into different lists and each list contains all events that are relevant to its sector, each sector may be processed independently, each one with its own *agenda*. This allows a straightforward parallelization of such processing. Also, note that the events of a cell may be included in more than one sector's event list and therefore some cells may be processed twice. But that is not a problem, since this will happen only to a few cells, and it will not affect the resulting viewshed.

It is also important to note that our algorithm might be faster than the original one even with non-parallel architectures. For instance, we achieved up to 20% speedup using only one processor (see Section 4). This happens because both implementations have to sort their lists of events and, while the original (serial) algorithm sorts a list of size n, our algorithm sorts S lists of size about $\frac{n}{S}$. Since sorting can be done in $\Theta(n \, log n)$, the latter one is faster. In practice, we empirically concluded that, for a computer with N cores, using $S > N$ achieved better results than using $S = N$. This will be further discussed in Section 4, as long with our experimental results.

4. Experimental results

We implemented our algorithm in C++ using OpenMP. We also implemented the original (serial) Van Kreveld's algorithm in C++. Both algorithms were compiled with g++ 4.6.4

and optimization level -O3. Our experimental platform was a Dual Intel Xeon E5-2687 3.1GHz 8 core. The operational system was Ubuntu 12.04 LTS, Linux 3.5 Kernel.

The tests were done using six different terrains from SRTM datasets and, in all experiments, the observer was sited in the center of the terrain, with $h_O = 100$ meters and $h_T = 0$. The radius of interest, ρ, was set to be large enough to cover the whole terrain.

Another important parameter for our program is the number of sectors S into which the region of interest will be subdivided. Changing the number of sectors may significantly modigy the algorithm performance. Empirically, we determined that good results are achieved when the region is subdivided such that each sector contained about 40 cells from the perimeter of the region of interest, so we adopted that strategy. Other strategies for choosing the number of sectors should be further investigated and it could be an interesting topic for future work.

To evaluate our algorithm performance, we compared it to the original (serial) algorithm. We ran several experiments limiting the number of parallel threads to the following values: 16, 8, 4, 2 and 1. The results are given in Table 1 and plotted in Figure 5(a), where the times are given in seconds and refer just to the time needed to compute the viewshed. That is, we excluded the time taken to load the terrain data and to write the computed viewshed into disk, since it was insignificant (less than 1% of the total time in all cases). Also, the time represents the average time for five different runs of the same experiment.
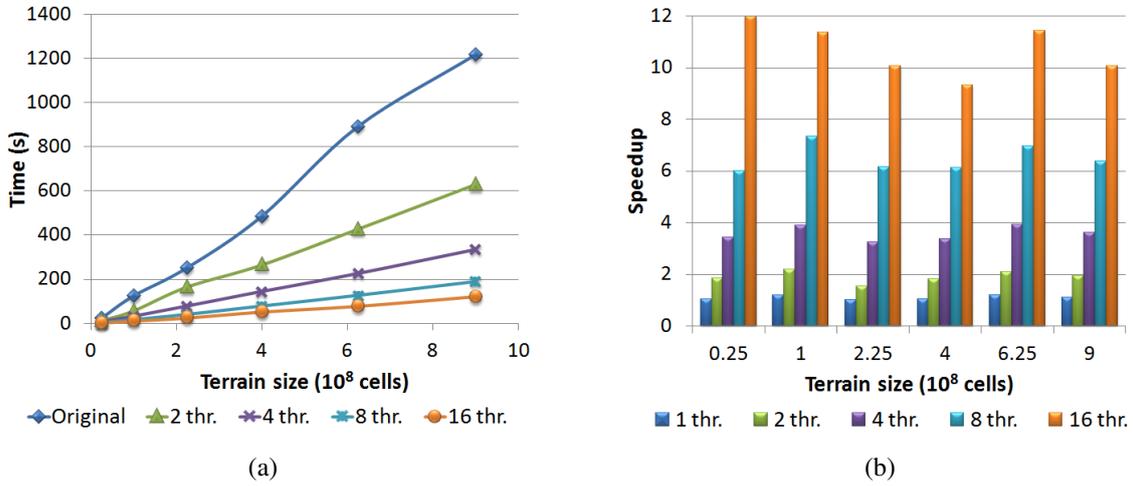
Table 1. Running times for the serial algorithm and the parallel algorithm with different number of threads.

Terrain size		Serial Alg.	Parallel Alg. Number of threads				
# cells	GiB		16	8	4	2	1
$5\,000^2$	0.09	24	2	4	7	13	23
$10\,000^2$	0.37	125	11	17	32	57	105
$15\,000^2$	0.83	252	25	41	78	165	246
$20\,000^2$	1.49	485	52	79	144	265	464
$25\,000^2$	2.33	891	78	128	226	427	740
$30\,000^2$	3.35	1\,216	121	191	335	629	1\,100

We calculated our algorithm speedup compared to the original algorithm. They are presented in Table 2 and plotted in Figure 5(b). Our algorithm has shown very good performance, achieving up to 12 times speedup, when running 16 concurrent threads. It is also important to notice that with only four threads we achieved a speedup of 3.9 times for two terrains and more than 3 times for all other terrains. Considering that processors

Table 2. Speedups achieved by our parallel algorithm, with different number of threads.

Terrain size		Parallel Alg. Number of threads				
# cells	GiB	16	8	4	2	1
$5\,000^2$	0.09	12.00	6.00	3.43	1.85	1.04
$10\,000^2$	0.37	11.36	7.35	3.91	2.19	1.19
$15\,000^2$	0.83	10.08	6.15	3.23	1.53	1.02
$20\,000^2$	1.49	9.33	6.14	3.37	1.83	1.05
$25\,000^2$	2.33	11.42	6.96	3.94	2.09	1.20
$30\,000^2$	3.35	10.05	6.37	3.63	1.93	1.11

(a)

(b)

Figure 5. (a) Running times for the serial algorithm and the parallel algorithm with different number of threads; (b) Speedups achieved by our parallel algorithm, with different number of threads.

with four cores have become usual and relatively cheap nowadays, these improvements may be useful for real users with regular computers. Finally, as discussed in Section 3, the experiments with only one thread show that our strategy can be faster than the original program even with serial architectures.

5. Conclusions and future work

We proposed a new parallel sweep line algorithm for viewshed computation, based on an adaptation of Van Kreveld's algorithm. Compared to the original (serial) algorithm, we achieved speedup of up to 12 times with 16 concurrent threads, and up to 3.9 times using four threads. Even with a single thread, our algorithm was better than the original one, running up to 20% faster.

Compared to other parallel viewshed algorithms, ours seems to be the only to use

Van Kreveld's model, which presents a great tradeoff between efficiency and accuracy [Fishman et al. 2009]. Also, most of them use other parallel models, such as distributed memory/message passing and general purpose GPU programming. On the other hand, ours uses the shared memory model, which is simpler, requires cheaper architectures and is supported by most current computers.

As future work, we can point out the development of other strategies for defining S, the number of sectors into which the region of interest is subdivided. We also intent to develop another adaptation of Van Kreveld's model using GPU programming. Since GPU architectures are much more complex, this will not be a straightforward adaptation.

Acknowledgements

References

Barney, B. et al. (2010). Introduction to parallel computing. *Lawrence Livermore National Laboratory*, 6(13):10.

Ben-Moshe, B., Ben-Shimol, Y., and Y. Ben-Yehezkel, A. Dvir, M. S. (2007). Automated antenna positioning algorithms for wireless fixed-access networks. *Journal of Heuristics*, 13(3):243–263.

Bresenham, J. (1965). An incremental algorithm for digital plotting. *IBM Systems Journal*, 4(1):25–30.

Champion, D. C. and Lavery, J. E. (2002). Line of sight in natural terrain determined by L_1-spline and conventional methods. In *23rd Army Science Conference*, Orlando, Florida.

Chao, F., Chongjun, Y., Zhuo, C., Xiaojing, Y., and Hantao, G. (2011). Parallel algorithm for viewshed analysis on a modern gpu. *International Journal of Digital Earth*, 4(6):471–486.

Cole, R. and Sharir, M. (1989). Visibility problems for polyhedral terrains. *J. Symb. Comput.*, 7(1):11–30.

Dagum, L. and Menon, R. (1998). Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55.

De Floriani, L. and Magillo, P. (2003). Algorithms for visibility computation on terrains: a survey. *Environment and Planning B: Planning and Design*, 30(5):709–728.

Fishman, J., Haverkort, H. J., and Toma, L. (2009). Improved visibility computation on massive grid terrains. In Wolfson, O., Agrawal, D., and Lu, C.-T., editors, *GIS*, pages 121–130. ACM.

Franklin, W. R. and Ray, C. (1994). Higher isn't necessarily better: Visibility algorithms and experiments. In Waugh, T. C. and Healey, R. G., editors, *Advances in GIS Research: Sixth International Symposium on Spatial Data Handling*, pages 751–770, Edinburgh. Taylor & Francis.

Franklin, W. R., Ray, C. K., Randolph, P. W., Clark, L., Ray, K., and Mehta, P. S. (1994). Geometric algorithms for siting of air defense missile batteries.

Lake, I. R., Lovett, A. A., Bateman, I. J., and Langford, I. H. (1998). Modelling environmental influences on property prices in an urban environment. *Computers, Environment and Urban Systems*, 22(2):121–136.

Lee, J. and Stucky, D. (1998). On applying viewshed analysis for determining least-cost paths on digital elevation models. *International Journal of Geographical Information Science*, 12(8):891–905.

Li, Z., Zhu, Q., and Gold, C. (2005). *Digital Terrain Modeling — principles and methodology*. CRC Press.

Maćiorov, F. (1964). *Electronic digital integrating computers: digital differential analyzers*. Iliffe Books (London and New York).

Magalhães, S. V. G., Andrade, M. V. A., and Franklin, W. R. (2011). Multiple observer siting in huge terrains stored in external memory. *International Journal of Computer Information Systems and Industrial Management (IJCISIM)*, 3:143–149.

Osterman, A. (2012). Implementation of the r. cuda. los module in the open source grass gis by using parallel computation on the nvidia cuda graphic cards. *ELEKTROTEHNISKI VESTNIK*, 79(1-2):19–24.

Shapira, A. (1990). Visibility and terrain labeling. *Master's thesis, Rensselaer Polytechnic Institute*.

Van Kreveld, M. (1996). Variations on sweep algorithms: efficient computation of extended viewsheds and class intervals. In *In Proceedings of the Symposium on Spatial Data Handling*, pages 15–27.

Zhao, Y., Padmanabhan, A., and Wang, S. (2013). A parallel computing approach to viewshed analysis of large terrain data using graphics processing units. *International Journal of Geographical Information Science*, 27(2):363–384.