

Natural Neighbor Interpolation Based Grid DEM Construction Using a GPU *

Alex Beutel, Thomas Mølhave, Pankaj K. Agarwal

Department of Computer Science

Duke University

{alex.beutel,thomasm,pankaj}@cs.duke.edu

ABSTRACT

With modern LiDAR technology the amount of topographic data, in the form of massive point clouds, has increased dramatically. One of the most fundamental GIS tasks is to construct a grid digital elevation model (DEM) from these 3D point clouds. In this paper we present a simple yet very fast algorithm for constructing a grid DEM from massive point clouds using natural neighbor interpolation (NNI). We use a graphics processing unit (GPU) to significantly speed up the computation. To handle the large data sets and to deal with graphics hardware limitations clever blocking schemes are used to partition the point cloud. For example, using standard desktop computers and graphics hardware, we construct a high-resolution grid with 150 million cells from two billion points in less than thirty-seven minutes. This is about one-tenth of the time required for the same computer to perform a standard linear interpolation, which produces a much less smooth surface.

Categories and Subject Descriptors: D.2 [Software]: Software Engineering; F.2.2 [Nonnumerical Algorithms and Problems]: Geometrical problems and computations; H.2.8 [Database Management]: Database Applications—*Data Mining, Image Databases, Spatial Databases and GIS*

General Terms: Performance, Algorithms

Keywords: LIDAR, Massive data, GIS, Natural Neighbor Interpolation

1 Introduction

The revolution in sensing and mapping technologies is providing an unprecedented opportunity to characterize and understand the earth's surface and dynamics. For instance, modern airborne LiDAR technology can map the earth's surface at a 15-20cm horizontal resolution, and the future generation of LiDAR scanners are expected to generate high-resolution maps of other planets; see Figure 1(a). It is essential for many applications to exploit the high-resolution data sets that are available. An example of this can be

*Supported by NSF under grants CNS-05-40347, CCF-06 -35000, IIS-07-13498, and CCF-09-40671, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, by an NIH grant 1P50-GM-08183-01, and by a grant from the U.S.–Israel Binational Science Foundation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GIS 2010 November 2-5, 2010, San Jose, CA, USA.

Copyright 2010 ACM 978-1-4503-0031-5 ...\$10.00.

seen in a simple flood mapping application. Figure 1(b,c) shows the result of the flood risk mapping for the island of Mandø in the Wadden-Sea off the west-coast of Denmark. The island has an approximately five meter tall perimeter dike which protects it from the sea. Because of the small width of the perimeter dike, this feature is not present in low- or mid-resolution grids. Thus, when flood maps are constructed for a water level of 2 meters, it looks as if most of the island will be underwater; see Figure 1(b) for an example using a 90m grid (the SRTM grid available from NASA [12]). The same computation performed on a 2m-resolution grid, shown in Figure 1(c), correctly finds that the dikes, now present in the terrain model, block the water from entering the lower-lying areas inside the perimeter.

Capitalizing on opportunities made feasible by high resolution data sets and transforming this massive amount of topographic data into useful information for vastly different types of users requires solving several challenging algorithmic problems. For example, in order to fully explore topographic data, one must often first extract a terrain from the scattered set of points generated by the LiDAR equipment. Most GIS applications do not work directly on the point cloud S gathered by a LiDAR scanner, but instead operate on a *digital elevation model (DEM)* of the terrain surface. Thus one of the most important of these problems is to generate a DEM from S .

Because of its simplicity and efficiency, the most widely used DEM is a uniform two-dimensional grid in which an elevation value is stored at each cell. However, point clouds are not acquired on a uniform grid but can be seen as a set S of n (arbitrary) points in \mathbb{R}^2 with an associated elevation function $h : S \rightarrow \mathbb{R}$. Thus, to construct a grid DEM, h has to be extended via interpolation to a uniform grid $G \subset \mathbb{R}^2$ at the desired resolution. Numerous interpolation methods, ranging from sophisticated but computationally expensive methods to simpler and efficient methods, for grid DEMs have been developed; see [19] for a review. Regularized splines with tension (RST) is a well-known method, which is sophisticated but computationally expensive due to its use of non-trivial polynomials [1, 20]. RST and similar highly sophisticated interpolation methods are especially good when the input data is sparse and lots of interpolation has to be performed. On the other hand, constructing a triangulation on input points and linearly interpolating the elevation of grid points across the triangles is one of the simplest interpolation methods. It, however, does not produce a smooth surface, especially when the data is relatively sparse. The resulting grid DEM can appear jagged both when viewed directly and also in derived products, such as contour maps.

In this paper we use the well-known *natural neighbor interpolation* strategy [25]. Based on the Voronoi diagram of S , it produces a smooth surface. Although NNI is more efficient than RST and other similar interpolation methods, its traditional implementations




Figure 1. (a) Grid DEM constructed from LiDAR data over a region in Afghanistan (data source: Army Research Office), trees are clearly visible. (b) A flood risk mapping of the island of Mandø in Denmark, using the 90m, and (c) the same using 2m grid; both figures are screenshots of a custom map application built on Google Maps.

are significantly slower than linear interpolation and are therefore not widely used for very large data sets.

Over the last decade modern PCs have started to become equipped with advanced and increasingly powerful graphics processing units (GPUs). Although originally designed for rapidly transforming 3D geometric scenes into pixels on the image plane (screen) and extensively used in video games, they can be regarded as massively parallel vector processors suitable for general purpose computing. Known as general purpose GPUs (GPGPUs), their tremendous computational power and memory bandwidth make them attractive for applications far beyond the original goal of rendering complex 3D scenes, and they have been used for a wide range of applications, e.g., geometric computing [3], robotic collision detection [13], database systems [14], fluid dynamics [18], and solving sparse linear systems [8, 7]. As GPUs have become more flexible and programmable (e.g. NVIDIA’s CUDA [22] library), their applicability has also increased tremendously; see [23] for a recent survey. In the context of grid DEM construction, Fan *et al.* [11] have described a GPU based algorithm for natural neighbor interpolation (NNI), which is considerably faster than a CPU based algorithm.

Our results. In this paper we present a simple yet very fast GPU based algorithm for constructing a grid DEM from large LiDAR point clouds using a variant of the natural neighbor interpolation method. LiDAR scanners provide dense (high resolution) point cloud of elevation data at most locations, but there are gaps, usually at large bodies of water or human-made objects that have been removed from the point cloud in a preprocessing step. In such cases we wish to label the corresponding “gap” cells in the grid DEM with “nodata” instead of interpolating elevation based on points that are far away. We introduce the notion of *region of influence* for each input point, similar to the one used in α -shapes [10]. For a grid point, we use only those points to compute its elevation whose regions of influence contain the query point. See Section 4 for details. Although our algorithm is similar to that of Fan *et al.* [11], there are three main differences:

- (i) Our algorithms handles gaps differently, as described above.
- (ii) Exploiting the fact that we are interpolating elevations at grid points, it uses a clever “blocking” scheme to expedite the computation considerably. In contrast to the algorithm in [11], which performs NNI interpolation at ≤ 32 points (or ≤ 128 points depending on hardware properties of the GPU

card) in one step, it can answer more than 10^6 NNI queries at grid points in one step.

- (iii) It exploits CUDA to substantially improve its efficiency by performing the majority of the computation on the GPU, thereby minimizing the communication between main and GPU memory.

These techniques lead to an extremely fast algorithm for computing a grid DEM. For example, our algorithm computes a grid DEM covering a 600km^2 region at 2m resolution (i.e., ≈ 150 million grid points) from two billion input points in less than thirty-seven minutes on a 3GHz Intel Core 2 Duo processor with a NVIDIA GeForce GTX 470 graphics card. Our CPU based linear-interpolation algorithm takes more than five and a half hours on the same PC. Not only is this a significant speedup, NNI interpolation also produces a smoother grid DEM than the linear-interpolation method. The more sophisticated RST-based algorithm takes about thirty-four hours on the same data set, even after throwing away a fraction of the points for efficiency, and the output between NNI and RST-based interpolations is nearly indistinguishable. Another advantage of our algorithm over linear or RST interpolation is that it can be trivially parallelized, so it could be implemented easily on a GPU cluster.

The paper is organized as follows. Section 2 provides a brief overview of the GPU model of computation, and Section 3 describes a GPU based algorithm for computing the Voronoi diagram of a set of points. Section 4 describes a slight variant of the Fan *et al.* [11] algorithm for computing natural neighbor interpolation, and Section 5 describes the new algorithm for computing NNI interpolation on a grid. Finally, Section 6 contains implementation details and information about the speed and quality experiments we have performed.

2 GPU Model of Computation

Primarily designed to achieve high performance for interactive graphics applications, modern programmable GPUs consist of a large number of processors (e.g., up to 480 for the newest NVIDIA GeForce 4-series) with a high memory bandwidth (177.4 GB/sec for NVIDIA’s GeForce GTX 480) and achieve higher floating-point throughput than the CPUs. This high throughput has led to a tremendous effort for developing GPU based numerical algorithms; see the recent survey by Owens *et al.* [23] and the references therein.




Figure 2. (a) Voronoi diagram $\text{Vor}(S)$ of a set S of points. (b) Pixelized Voronoi diagram $\text{PVor}(S)$. (c) Truncated pixelized Voronoi diagram $\text{TPVor}(S)$ of S with $r = 5$, k was set high enough for the k -gons to be indistinguishable from cones.

The graphics computation in all GPUs follows a similar pipeline, called the *graphics pipeline*, which draws a three dimensional scene, composed of many objects, onto a two dimensional image plane Π of pixels as seen from a specified viewpoint o . Because of their simplicity and flexibility, these objects are almost always triangles. For each pixel $\pi = (x, y)$ where x, y is a global coordinate, the GPU finds all objects $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\}$ which ray $o\pi$ intersects. To maintain high throughput, each stage of the computation is implemented in hardware and computation on different parts of Π is performed in parallel. See e.g. [17, 21] for details of NVIDIA’s GeForce 6 series. Here we note that, a GPU contains several two dimensional arrays of pixels called buffers (or texture memory). We mention the two most basic ones, which we will use:

- The *depth buffer* \mathbb{D} stores the distance to the nearest object from o for each pixel π . Given that p_j is the point of intersection for ray $o\pi$ and object ω_j , the GPU calculates

$$\mathbb{D}[\pi] = \min_{1 \leq j \leq n} \|op_j\|.$$

The modern GPUs provide the flexibility of performing various *semigroup* operations on op_j ’s instead of simply computing the minimum and also of performing them in a conditional manner. For example, \mathbb{D} can be in the *read-only* mode.

- The *color buffer* \mathbb{C} stores the color of the scene as viewed from o . If for each object $\omega_j \in \Omega$ we have a color χ_j , we define a blending function that computes the color at each pixel:

$$\mathbb{C}[\pi] = \sum_{1 \leq j \leq n} \alpha_j \chi_j,$$

where $\alpha_j \in [0, 1]$ is the blending parameter of ω_j . Typically, α_j is based on depth buffer so that \mathbb{C} stores the color of the foremost object. Again, the modern GPUs provide several other binary operations on the colors. We will like our blending function to compute the bitwise-OR of the colors. This can be done by setting $\alpha_j = 1$ for $1 \leq j \leq n$, as long as χ_j are bitwise-disjoint.

During graphical computations, the color and depth buffers reside in memory on the graphics card. Objects can be drawn onto these buffers with specific APIs such as OpenGL[24] or Microsoft DirectX [6]. However, in some cases we will want to use the values in these buffers for computation on the CPU. For this, we have

to read the buffer back to the computer’s main memory. Unfortunately, since this involves transferring large amounts of data over the relatively slow bus systems, read backs are very slow.

For using the GPU’s parallel processing capabilities, the popular graphics card manufacturer NVIDIA has created the CUDA parallel computing architecture. CUDA makes it easy to divide a task into many threads, where threads can work in parallel but also when necessary share memory and work together on procedures that aren’t trivially parallelizable. Additionally advantageous is that CUDA operations are performed directly on the graphics card and can efficiently access buffers, which reside in nearby GPU memory. As an example, CUDA can split into one thread for each pixel of a buffer and read from each pixel simultaneously without reading the buffer back to main memory. Inversely, modern GPUs allow multiple threads to write to the same memory with atomic functions that provide thread synchronization and serialization.

3 Pixelized Voronoi Diagram

Let $S = \{p_1, \dots, p_n\}$ be a set of n points in \mathbb{R}^2 . For each point $p \in S$, its *Voronoi cell*, denoted by $\text{Vor}_S(p)$, is defined as

$$\text{Vor}_S(p) = \{x \in \mathbb{R}^2 \mid \|xp\| \leq \|xq\| \forall q \in S\},$$

where $\|\cdot\|$ is the Euclidean distance, i.e., a point $x \in \text{Vor}_S(p)$ if p is the point in S closest to x . The *Voronoi diagram* of S , $\text{Vor}(S)$, is the planar subdivision induced by the Voronoi cells of points in S . See Figure 2 (a).

Hoff *et al.*[16] have described a GPU based algorithm for computing the Voronoi diagram of a set of points. Since we use a slightly different algorithm, we describe the algorithm for the sake of clarity and completeness. An image plane Π consisting of $N \times N$ pixels can be regarded as the square $[0, N - 1] \times [0, N - 1]$ in \mathbb{R}^2 . Any square $R \subseteq \mathbb{R}^2$ can be mapped to Π using an affine transformation. Given the set S and a square R , we are interested in computing a *pixelized (discretized) Voronoi diagram* of S within R , which we define below. We assume that R is mapped to the image plane Π . Each pixel of Π corresponds to a (tiny) square of area $\rho^2 = \text{Area}(R)/N^2$. We refer to ρ as the *resolution* of Π . For a pixel $\pi \in \Pi$, let $\varphi(\pi, S)$ be the point in S whose Voronoi cell contains π . Since π is a (tiny) square region, it may intersect multiple Voronoi cells, in which case $\varphi(\pi, S)$ is assigned to one of the points using standard methods. For a point $p \in S$, we define the




Figure 3. Voronoi diagram as the lower envelope of a set of cones. The outer cells of a Voronoi diagram are infinite, but in this figure their sizes are limited because the cones are truncated.

pixelized Voronoi cell of p to be

$$\text{PVor}_S(p) = \{\pi \mid \varphi(\pi, S) = p\},$$

i.e., the set of pixels that lie in $\text{Vor}_S(p)$; see Figure 2 (b). The quantity $\rho^2 |\text{PVor}_S(p)|$ approximates the area of $\text{Vor}_S(p)$ within R . The approximation error depends on ρ . For a fixed R , the error decreases as we increase N , namely,

$$\lim_{N \rightarrow \infty} \rho^2 |\text{PVor}_S(p)| = \text{Area}(\text{Vor}_S(p)).$$

For our purpose, we assume that $\text{PVor}(S)$ is stored as follows. If $\varphi(\pi, S) = p_i$, then the color buffer $\mathbb{C}[\pi] = i$, i.e., we view each cell of the color buffer as a single word (concatenation of R, G, B, A components) that stores the index of the point; $\mathbb{D}[\pi]$ stores the value of $\|\pi\varphi(\pi, S)\|$, the distance from the center of π to its nearest neighbor in S .

The problem of computing $\text{PVor}(S)$ can be formulated as that of rendering a 3D scene. For a point $p_i \in S$, let $f_i : \mathbb{R}^2 \rightarrow \mathbb{R}$ be defined as $f_i(x) = \|xp_i\|$. The lower envelope f of $\{f_1, \dots, f_n\}$ is defined to be

$$f(x) = \min_{1 \leq i \leq n} f_i(x),$$

which is the distance from x to its nearest neighbor. $\text{Vor}(S)$ is the projection of the graph of f on the xy -plane. Let C be the circular cone $C : z = \sqrt{x^2 + y^2}$ in \mathbb{R}^3 . Then the graph of each f_i is a circular cone $C_i = C + p_i$, with p_i as its apex. See Figure 3. Let $\mathcal{C} = \{C_1, \dots, C_n\}$. A point $x \in \text{Vor}_S(p_i)$ if $f(x)$ is realized by the function f_i at x , i.e., the line oriented in the $+z$ direction hits C_i first. In other words, $\varphi(\pi, S) = p_i$ if C_i is the cone seen at pixel π when the set C is viewed from $z = -\infty$. If we set the color of C_i to i , then the color and depth buffers store the desired information.

It is not easy to render a circular cone using a GPU, so we approximate a circular disk by a regular k -gon and approximate the circular cone by using this k -gon as its base (see Figure 4). The resulting polygonal cone C^\diamond is composed of k triangles. We replace C_i by $C_i^\diamond = C^\diamond + p_i$. The error in tessellation induced by this approximation can be controlled by choosing the value of k appropriately.




Figure 4. Approximating disk by a k -gon, and the corresponding polyhedral cone.

Finally, we note that we want to limit the *region of influence* for the points. We do this by using a truncated Voronoi diagram. We define the *radius of influence* r of each point in S and the notion of a *truncated pixelized Voronoi cell*:

$$\text{TPVor}_S(p) = \{\pi \mid \varphi(\pi) = p \wedge \|p\pi\| < r\}.$$

See Figure 2 (c). Thus, a pixel π that is farther than r away from all points of S does not belong to the Voronoi cell of any point. Let D_r denote the disk of radius r centered at origin. We can assume that $S \subset R + D_r$, as no point outside this region will contain any pixel of Π in its Voronoi cell. This truncation is realized by limiting the height of the cones C_i^\diamond ; with a slight abuse of notation we use C_i^\diamond to denote the truncated cone as well. For each $p_i \in S$, we set the color of each triangle of C_i^\diamond to i and pass them to the graphics pipeline with $z = -\infty$ as the viewpoint. \mathbb{C} and \mathbb{D} together contain $\text{TPVor}(S)$. We refer to this algorithm as $\text{GPUVORONOI}(S)$. As mentioned above, there might be pixels that are not touched by $\text{GPUVORONOI}(S)$. We assume that \mathbb{C} is initialized with a value that allows us to distinguish these pixels from the pixels that are part of the truncated diagram, e.g., we set their color to 0.

4 Natural Neighbor Interpolation

In this section we first formally define natural neighbor interpolation (NNI), then describe a GPU algorithm for answering NNI queries, which is a small variant of the algorithm by Fan *et al.* [11]. A height function $h : S \rightarrow \mathbb{R}$ can be extended to entire \mathbb{R}^2 using natural neighbor interpolation. In particular, for a point $x \in \mathbb{R}^2$,

$$h(x) = \sum_{p \in S} w_p(x) h(p),$$

where $w_p(x)$ is the fractional area of $\text{Vor}_{S \cup \{x\}}(x)$ that belongs to $\text{Vor}_S(p)$ (Figure 5), i.e.,

$$w_p(x) = \frac{\text{Area}(\text{Vor}_S(p) \cap \text{Vor}_{S \cup \{x\}}(x))}{\text{Area}(\text{Vor}_{S \cup \{x\}}(x))}.$$

Since we use truncated pixelized Voronoi diagrams, we redefine the height function as

$$h(x) = \sum_{p \in S} \bar{w}_p(x) \cdot h(p) \quad (1)$$




Figure 5. Natural neighbor interpolation. Shaded cell is $\text{Vor}_{S \cup \{q\}}(q)$, and each color denotes the area stolen from each cell of $\text{Vor}(S)$.




Figure 6. (a) $\text{Vor}(S)$ of a set S . (b) $\text{Vor}(S \cup \{q_1\})$ and $\text{Vor}(S \cup \{q_2\})$ for two query points q_1 and q_2 . (c) $\text{PVor}(S \cup \{q_1\})$ and $\text{PVor}(S \cup \{q_2\})$. The colors correspond to the bitwise-OR colors of the query point. Query point q_1 and q_2 have colors 01 and 10 respectively, the pixels in $\text{PVor}_{S \cup \{q_1\}}(q_1) \cap \text{PVor}_{S \cup \{q_2\}}(q_2)$ thus get the color $01 \vee 10 = 11$, where \vee is bitwise or.

Algorithm 1 BUFFERANALYSIS($\mathcal{C}^1, \mathcal{C}^2$)

for all $\pi \in \mathcal{C}^2$ do
 if $\mathcal{C}^2[\pi] \neq 0$ and $\mathcal{C}^1[\pi] \neq 0$ then
 $N_q \leftarrow N_q + \mathcal{C}^1[\pi]$
 $D_q \leftarrow D_q + 1$
return N_q/D_q

where

$$\bar{w}_p(x) = \frac{|\text{TPVor}_S(p) \cap \text{TPVor}_{S \cup \{x\}}(x)|}{|\text{TPVor}_{S \cup \{x\}}(x)|}. \quad (2)$$

Answering an NNI query. For a point $x \in \mathbb{R}^2$, let $D_r(x) = D_r + x$ denote the disk of radius r centered at x . Let q be a query point such that $D_r(q) \subseteq R$. The algorithm for computing $h(q)$ works in two phases. The first phase calls GPU Voronoi (S) with the following twist: the color of each triangle of the cone C_i^\diamond is set to $h(p_i)$ (instead of i). After the first phase $\mathcal{C}[\pi]$ stores $h(p_i)$ for all pixels $\pi \in \text{TPVor}_S(p_i)$. We read back the color buffer; let \mathcal{C}^1 denote the resulting two-dimensional array. We then clear the color buffer. The depth buffer \mathbb{D} is not touched, i.e., $\mathbb{D}[\pi]$ continues to store $\|\pi\varphi(\pi, S)\|$, the distance from the center of π to $\varphi(\pi, S)$. In the second phase, we set \mathbb{D} to read-only mode so that it is not overwritten and draw a polygonal cone $qC^\diamond = C^\diamond + q$ with q as the apex. Adding qC^\diamond is the same as computing $\text{TPVor}_{S \cup \{q\}}(q)$. However, the color buffer was cleared before the second phase and thus has non-zero entries¹ (corresponding to the color of qC^\diamond) only for $\text{TPVor}_{S \cup \{q\}}(q)$. Let \mathcal{C}^2 denote the color buffer contents after the second phase, and \mathcal{C}^2 the array resulting from reading back \mathcal{C}^2 into memory. The value of $h(q)$ can be computed by adding the values of $\mathcal{C}^1[\pi]$ for all π for which $\mathcal{C}^2[\pi] \neq 0$ and finally dividing the sum by the number of non-zero values in \mathcal{C}^2 (this is the denominator of (2)). We refer to this final step of the algorithm as BUFFERANALYSIS; Algorithm 1 gives the pseudo code.

There are four main sources of difference between our method and the traditional natural neighbor interpolation:

- (i) The tessellation error caused by using k -gons instead of cones.
- (ii) The discretization error.

¹We assume without loss of generality that all points of S have a positive height.

(iii) The limited precision of the depth buffer \mathbb{D} (which can cause problems at the boundaries between two Voronoi cells).

(iv) The radius of influence r .

Batching the queries. The above algorithm is very inefficient if we want to compute heights at many points, especially since reading back a buffer is a slow step. Fan *et al.* [11] have shown that by exploiting the power of modern GPUs, many queries can be batched and answered in one pass. More precisely, if each pixel of the color buffer \mathcal{C} has w bits, then $m \leq w$ queries can be batched in one pass by encoding the colors cleverly in \mathcal{C} : One bit of $\mathcal{C}[\pi]$ is assigned for each query point. Let q_1, q_2, \dots, q_m be the m query points, and let $qC_i^\diamond = C^\diamond + q_i$ be the cone corresponding to the query point q_i . The color of all triangles in qC_i^\diamond is set to 2^i . This ensures that colors of the m query points are bitwise-disjoint.

The first phase of the algorithm is the same as before. Let \mathcal{C}^1 be the same as above. In the second phase, we again set the depth buffer to the read-only mode. We draw $qC_1^\diamond, qC_2^\diamond, \dots, qC_m^\diamond$ but the color buffer now operates as follows. Suppose the graphics pipeline is rendering a triangle t of color χ_t . If the depth of t at a pixel π is larger than $\mathbb{D}[\pi]$, then nothing happens. Otherwise $\mathcal{C}[\pi]$ is set to $\mathcal{C}[\pi] \leftarrow \mathcal{C}[\pi] \vee \chi_t$, where \vee is the bitwise OR operation. Recall that $\mathbb{D}[\pi]$ is not updated, as it is in the read-only mode and $\mathbb{D}[\pi]$ stores $\|\pi\varphi(\pi)\|$. After the second phase, the i 'th bit of $\mathcal{C}[\pi]$ is 1 if $\pi \in \text{TPVor}_{S \cup \{q_i\}}(q_i)$ see Figure 6(c).

We read back the color buffer; let \mathcal{C}^2 denote the contents of the buffer after the second phase. We compute $h(q_i)$ by summing the values of $\mathcal{C}^1[\pi]$ for all pixels π for which the i 'th bit of $\mathcal{C}^2[\pi]$ is 1 and then dividing the sum by the number of such pixels. Algorithm 2 gives the pseudo code for doing this step efficiently; N_q and D_q are arrays of length m ; for a bit-vector v , v_i denotes its i 'th bit; and H is an array of length m , where $H[i] = h(q_i)$.

5 NNI on Grids

We now describe a faster algorithm for answering NNI queries when the query points lie on an $M \times M$ grid \mathbb{Q} . The algorithm can easily be extended to handle rectangular grids. For convenience we will use $\mathbb{Q}[i, j]$ to denote the (i, j) 'th query point of \mathbb{Q} for $0 \leq i, j < M$. Let $s\rho$ denote the size of each grid cell in \mathbb{Q} , for some positive integer s . We refer to s as the *scaling parameter*, and for simplicity we assume that s is odd and that $\rho(s-1)/2 < 2r$.

Algorithm 2 BUFFERANALYSIS($\mathcal{C}^1, \mathcal{C}^2, m$)

```

for  $i \leftarrow 1$  to  $m$  do
   $N_q[i] \leftarrow 0$ 
   $D_q[i] \leftarrow 0$ 
for all  $\pi \in \mathcal{C}^2$  do
  if  $\mathcal{C}^2[\pi] \neq 0 \wedge \mathcal{C}^1[\pi] \neq 0$  then
     $v = \mathcal{C}^2[\pi]$ 
    for  $i \leftarrow 1$  to  $m$  do
      if  $v_i = 1$  then  $\{\pi \in \text{TPVor}_{S \cup \{q_i\}}\}$ 
         $N_q[i] \leftarrow N_q[i] + \mathcal{C}^1[\pi]$ 
         $D_q[i] \leftarrow D_q[i] + 1$ 
for  $i \leftarrow 1$  to  $m$  do
   $H[i] \leftarrow N_q[i]/D_q[i]$ 
return  $H$ 

```

\mathbb{Q} can be mapped to the image plane Π so that each grid point of \mathbb{Q} lies at the center of an $s \times s$ array of pixels of Π . Since the radius of influence is r , we can assume that all points of S lie within distance $2r$ from \mathbb{Q} ; see below for an explanation. Let $\sigma = [-2r, +2r]^2$ and $\mathcal{Q} = \mathbb{Q} + \sigma$, so $S \subset \mathcal{Q}$. For now we assume that $N \geq sM + 4r/\rho$, i.e., $M \leq (N - 4r/\rho)/s$. Later we will describe how to handle larger query grids. Our assumption ensures that \mathcal{Q} can be mapped to Π with ρ being the resolution of each pixel. Let $\alpha = (s - 1)/2 + 2r/\rho$. We map the bottom left corner of \mathcal{Q} to that of Π , so the query point $\mathbb{Q}[i, j]$ maps to the pixel $(s \cdot i + \alpha, s \cdot j + \alpha)$; see Figure 7.

Let $B = \lfloor \sqrt{w} \rfloor$, where w , as above, is the number of bits in the color buffer. For simplicity, we assume that B is a divisor of M . We partition \mathbb{Q} into $(M/B)^2$ query blocks, each of size $B \times B$, with the $(i, j)^{\text{th}}$ block, for $0 \leq i, j < M/B$, being

$$\mathbb{Q}_{i,j} = \mathbb{Q}[iB, (i+1)B - 1][jB, (j+1)B - 1].$$

See Figure 8 (a). A query point $q \in \mathbb{Q}$ can be represented by a pair (\mathbf{a}, \mathbf{t}) , where $\mathbf{a} \in [0, M/B - 1]^2$ is the index of the query block that contains q , $\mathbf{t} \in [0, B - 1]^2$ is the offset of q within that query block.

We process all $B^2 \leq w$ NNI queries in each block in one pass, using the algorithm described in Section 4. Processing all queries in \mathbb{Q} thus requires $(M/B)^2$ passes, each involving the expensive operation of reading back the GPU memory to the main memory.

One-pass algorithm. By exploiting the grid structure of query points and the fact that a query q is affected by only those points p for which $D_r(p)$ and $D_r(q)$ intersect (i.e., $\|pq\| \leq 2r$), we show that we can answer all queries in one pass provided that

(A1) $M \leq (N - 4r/\rho)/s$, and

(A2) $r \leq s\rho B/2$.




Figure 7. Embedding \mathbb{Q} on Π ; $s = 3$ and $r = 2\rho$.

The second assumption is reasonable for high resolution LiDAR data sets because the height of a point can be interpolated from the nearby sampled points. For a graphics card with $w = 32$ and with $s = 5$ we will assume that $r < 5$ meter for an output grid with a resolution of $s\rho = 2$ meter. In other words, the result of a query is not affected by an input point that is more than 20 meters away. With high-resolution LiDAR data sets this is not a bad assumption. Such gaps usually appear when buildings and other features are removed using a classification algorithm, or at lakes and similar features.

Here is the key idea that enables us to answer all queries in one pass assuming that (A1) and (A2) hold. We call two points $p, q \in S$ independent if $\text{TPVor}_S(p)$ and $\text{TPVor}_S(q)$ are disjoint. If $\|pq\| > 2r$, then, by definition of $\text{TPVor}(S)$, p and q are independent. Let q_1, \dots, q_u be a set of query points such that $\|q_i q_j\| > 2r$ for any pair $i \neq j$. Then $q_i C_i^\diamond$ and $q_j C_j^\diamond$ are disjoint. We can potentially use the same color for all of these query points because no pixel will be rendered by two such cones in the second pass of the NNI query-answering algorithm. The difficulty with using the same color for all q_i 's is that the color of $q_i C_i^\diamond$ no longer encodes the value of i , so the algorithm does not know which query point colored a given pixel. However, q_i 's being independent implies that there is at most one query point for each pixel π that could color π , namely, the query point closest to π and it lies within distance r from π .

In our case, the query points lie on a grid \mathbb{Q} and we assume that $r < s\rho B/2$. Therefore for any two query blocks \mathbb{Q}_a and \mathbb{Q}_b and an offset \mathbf{t} , the query points (\mathbf{a}, \mathbf{t}) and (\mathbf{b}, \mathbf{t}) are independent, as the distance between query points with the same offset in two adjacent blocks is $s\rho B$. In other words, for any $\mathbf{t} \in [0, B - 1]^2$, all points in

$$\mathbb{Q}_{|\mathbf{t}} = \{(\mathbf{a}, \mathbf{t}) \mid \mathbf{a} \in [0, M/B - 1]^2\},$$

the set of all query points with offset \mathbf{t} , are independent; see Figure 8 (b). We can therefore assign the same color, say, χ , to all points in $\mathbb{Q}_{|\mathbf{t}}$. If a pixel π is colored χ , we can determine in $O(1)$ time which point in $\mathbb{Q}_{|\mathbf{t}}$ colored π . Hence, we proceed as follows.

For $\mathbf{t} = (t_1, t_2) \in [0, B - 1]^2$, we set $\chi(\mathbf{t}) = 2^{t_1 B + t_2}$ and assign the color $\chi(q)$ to all triangles of the query cone $C^\diamond + q$ for $q \in \mathbb{Q}_{|\mathbf{t}}$. Let

$$\mathcal{C} = \{q C_{ij}^\diamond = C^\diamond + \mathbb{Q}[i, j] \mid 0 \leq i, j < M\}$$

be the set of all query cones.

The first pass of the algorithm is the same as in Section 4, i.e., we compute $\text{TPVor}(S)$. In the second pass, we render all cones in




Figure 8. (a) Splitting \mathbb{Q} into query blocks of size $B = 4$ query points. (b) All the $\mathbb{Q}_{[1,1]}$ query points are independent since their areas of influence (depicted by circles in the figure) are disjoint.

Algorithm 3 BUFFERANALYSIS($\mathbb{C}^1, \mathbb{C}^2, \mathbb{Q}$)

```

for  $i \leftarrow 0$  to  $M - 1$  do
  for  $j \leftarrow 0$  to  $M - 1$  do
     $N_q[i, j] \leftarrow 0, D_q[i, j] \leftarrow 0$ 
  for all  $\pi \in \mathbb{C}^2$  do
    if  $\mathbb{C}^2[\pi] \neq 0$  and  $\mathbb{C}^1[\pi] \neq 0$  then
       $v = \mathbb{C}^2[\pi]$ 
      for  $\ell \leftarrow 0$  to  $w - 1$  do
        if  $v_\ell = 1$  then
           $(i, j) = \text{NN}(\pi, \ell)$ 
           $N_q[i, j] \leftarrow N_q[i, j] + \mathbb{C}^1[\pi]$ 
           $D_q[i, j] \leftarrow D_q[i, j] + 1$ 
      for  $i \leftarrow 0$  to  $M - 1$  do
        for  $j \leftarrow 0$  to  $M - 1$  do
           $H[i, j] \leftarrow N_q[i, j] / D_q[i, j]$ 
return  $H$ 

```

\mathbb{C} one by one while keeping the depth buffer in the read-only mode. Next, we read the color buffer back to main memory, and let \mathbb{C}^2 denote its content. We process each pixel π as follows: If the ℓ^{th} bit of π is 1, i.e., it has been rendered by a query point with the offset $\mathbf{l} = (\lfloor \ell/B \rfloor, \ell \bmod B)$, then we compute in $O(1)$ time the query point $\mathbb{Q}[i, j]$ that rendered π , i.e., the nearest point to π in the set \mathbb{Q}_1 . Let $\text{NN}(\pi, \ell)$ denote this procedure. We update the height of $\mathbb{Q}[i, j]$ appropriately. Algorithm 3 gives the pseudo-code of this step.

Handling larger grids. The preceding algorithm assumed that \mathbb{Q} was small enough that the entire \mathbb{Q} could be mapped to Π . This is not always realistic since the value of N is limited by the graphics hardware. For example, $N \leq 2^{14} = 16384$ on modern graphics cards such as the NVIDIA GeForce GTX 470. With a scaling parameter $s = 5$, $M \leq \lfloor 2^{14}/5 \rfloor = 3276$, implying that we can process $3276^2 \approx 10^7$ grid query points in a single pass. Recall that each pass consists of two rendering phases and the subsequent buffer analysis. For $s\rho = 2$ meter, this corresponds to computing a grid DEM for a region of roughly $70 \times 70 \text{ km}^2$ in area. However, we often want to generate grid DEMs that are considerably larger, in which case we proceed as follows.

Let $\mu = (N - 4r/\rho)/s$, the largest grid of query points we can handle in one pass is $\mu \times \mu$. Thus, if $M > \mu$ we partition \mathbb{Q} into $\mu \times \mu$ sub-grids; see Figure 9(a) and process these sub-grids individually. Let $m = M/\mu$, then m^2 is the number of sub-grids. We define

$$\mathbb{Q}^{i,j} = \bigcup_{(l,k) \in [0, \mu-1]^2} \mathbb{Q}[l + i\mu, k + j\mu]$$

to be the (i, j) 'th sub-grid, for $i, j < m$. For sub-grid \mathbb{Q}^l we let $\mathbb{Q}^l = \mathbb{Q}^l + \sigma$. We interpolate each \mathbb{Q}^l on the GPU independently, using the algorithm described above. Thus we need to find the set $S^l = \mathbb{Q}^l \cap S$ of input points relevant for the queries in \mathbb{Q}^l . Note that the S^l 's are not disjoint; see Figure 9(b). Let \mathcal{M} and \mathcal{B} be the amount of points that fit in main memory and in a disk block respectively. For simplicity we assume that \mathcal{B} divides \mathcal{M} .

If $|S^l| < \mathcal{M}$, then extracting the set S^l is not hard — we can construct a two-dimensional table to store S and extract S^l efficiently for each sub-grid \mathbb{Q}^l . This is, however, more challenging and expensive when, as is typically the case, S is too large to fit in the main memory. For example, the Denmark data set we have consists of 25 billion points. In this case, we preprocess S as follows.

We can keep up to \mathcal{M}/\mathcal{B} streams of points by storing a block of \mathcal{B} elements for each stream in memory, and the rest on disk. If




Figure 9. (a) When \mathbb{Q} is too big to handle in one pass, we split it up into sub-grids of size μ . To filter the points of S to the right place we may need several passes where the grids are split into \mathcal{M}/\mathcal{B} partitions of size n by n . (b) This figure shows $\mathbb{Q}^{1,1} = \mathbb{Q}^{1,1} + \sigma$ (contained in the dotted square) and $\mathbb{P}^{1,1} = \mathbb{P}^{1,1} + \sigma$ (contained in the dashed square).

$m^2 \leq \mathcal{M}/\mathcal{B}$, then we construct a stream for each bin \mathbb{Q}^l . We then distribute each point $p \in S$ to the bins for the sub-grids $\{\mathbb{Q}^l | p \in \mathbb{Q}^l\}$. If $m^2 > \mathcal{M}/\mathcal{B}$ we cannot hold enough streams in memory and instead use a recursive procedure. We partition the m^2 sub-grids into \mathcal{M}/\mathcal{B} square partitions of dimension $n = m/\sqrt{\mathcal{M}/\mathcal{B}}$.

$$\mathbb{P}^{i,j} = \bigcup_{(l,m) \in [0, n-1]^2} \mathbb{Q}^{l+in, k+jn},$$

for $0 \leq i, j < \sqrt{\mathcal{M}/\mathcal{B}}$ (see Figure 9(b)). Since there are \mathcal{M}/\mathcal{B} partitions, we can construct a stream for each of them. We then distribute each point $p \in S$ into the stream representing partition \mathbb{P}^x if $p \in \mathbb{P}^x = \mathbb{P}^x + \sigma$, where $\sigma = [-2r, +2r]^2$. Following this distribution step we recurse on each partition individually. The depth of the recursion is $O(\log_{\sqrt{\mathcal{M}/\mathcal{B}}} m) = O(\log_{\mathcal{M}/\mathcal{B}} M/\mu)$.

6 Implementation and Experimentation

Here we describe implementation details that contribute to the efficiency and quality of our algorithm. We subsequently offer empirical results for tests of the algorithm's speed and quality on real-world terrains.

Platform. We ran our experiments on an Intel Core2 Duo CPU E6850 at 3.00GHz with 4GB of internal memory. We used Ubuntu 10.4 and two 1TB SATA disk drives in a RAID0 configuration. Additionally, the machine contained a NVIDIA GeForce GTX 470 graphics card running CUDA 3.0. This card has 1.2 gigabytes of memory, 448 CUDA cores, and 14 multiprocessors.

The algorithm was implemented in C++ using OpenGL to interact with the graphics card. The \mathbb{C}^1 , \mathbb{C}^2 , and \mathbb{D} buffers were implemented using OpenGL's frame-buffer and render-buffer objects. Additionally, we used a display list to render the cones \mathbb{C}^\diamond . As described, our algorithm uses the same radius of influence r for all \mathbb{C}^\diamond , however, for flexibility our implementation supports using one radius for input points r_s and another one for queries r_q . Adjusting these values separately allows us to optimize r_q based on hardware bounds imposed by w , while adjusting r_s based on properties of our data sets.

6.1 Reducing communication complexity

The computational efficiency is one of two important factors in the real-world performance of our algorithm. We have taken great care to minimize the other major component, communication cost, as

	Afghanistan	DKPART	Fort Leonard Wood
Size of input (10^6)	186	1038	2180
Size of output (10^6)	9.5	213	151
NNI with CUDA	163	1238	2190
Binning Time	67	558	1030
Interpolation Time	96	680	1160
NNI without CUDA	1252	14323	11164
Binning Time	91	569	1036
Interpolation Time	1161	13754	10128
Linear	962	7377	20307
RST	5698	66729	122305

Table 1. Time comparison of competing interpolation algorithms (times in seconds)

well. As mentioned in the introduction, the cost of transferring buffers between GPU and main memory is substantial, but the cost of transferring data between the hard drive and the main memory is also substantial, especially for large data sets that do not fit in main memory.

GPU to CPU communication. As described in Section 4, the colors buffers \mathbb{C}^1 , \mathbb{C}^2 are read back into memory resulting in \mathbb{C}^1 and \mathbb{C}^2 . They are then used by the BUFFERANALYSIS algorithm, the final step in the interpolation. However, each of these buffers contains an $s \times s$ square of pixels for each query point of \mathbb{Q} , which means that each buffer is a factor of s^2 larger than \mathbb{Q} . Thus, we are transferring far more data between GPU and CPU memory than would be required if we could just transfer the final interpolated values (i.e. the buffer H from BUFFERANALYSIS).

Therefore, we have used CUDA to implement BUFFERANALYSIS directly on the graphics card. As discussed previously, CUDA can directly access the color buffers from GPU memory and likewise can keep two dimensional arrays of its own for N_q , D_q , and H . In performing this final summation step in CUDA, each pixel in the color buffer is accessed in parallel, and for each bit that is set to one, the appropriate values in N_q and D_q are incremented. Because this may cause the same memory location to be written to by multiple threads simultaneously, we use a CUDA function for atomic addition, providing serialization and synchronization of the many threads. Thus, we only perform one read from GPU memory to main memory, with only one 32-bit word being transferred for each query point. This has a drastic effect on running times of our algorithm, which will be demonstrated in our experiments.

Reducing disk-transfers. As described in Section 5 we preprocess S by binning the data into sub-grids of size $\mu \times \mu$, before feeding them to the GPU one by one. We have used the efficient disk-based stream abstractions provided by the Templated Portable I/O Environment (TPIE) [5] library to implement the recursive algorithm that performs this binning.

Additionally, the final output from the interpolation routine is a stream of points $(i, j, h(i, j))$ which typically need to be converted into a row-major raster grid sorted on i and j . For this we use the external-memory sorting algorithm [4] from TPIE, which is asymptotically optimal with respect to disk-memory transfers. The time taken by the final step is not included in the results presented in this section since this step is the same for all of the algorithms presented and it is not the bottleneck in the running time.

6.2 Performance

In analyzing our algorithm, we compared our results against results from applying a regularized spline with tension (RST) and linear interpolation based on the global Delaunay triangulation of the data.

The RST implementation is from Danner *et al.*'s previous work on TerraSTREAM[9]. We also tested our NNI implementation against the linear interpolation, in which we compute the Delaunay triangulation of the entire data set, as presented in [2], followed by interpolating the value of each query point. Both of these algorithms are sequential.

Data sets. For our tests, we ran our different interpolation schemes on three main data sets. The first was LiDAR data that covers most of Denmark supplied by COWI A/S. The entire point cloud is 1.5 terabytes in size with 26 billion data points. For a number of our experiments we used a portion of this data set, which we will refer to as DKPART. DKPART contains 1 billion data points over a 10 kilometer by 90 kilometer region and is 27 gigabytes on disk. This gives a point density of 0.9 points per square meter on average.

Because Denmark is relatively flat, we also ran tests for both speed and quality on a point cloud of a mountainous region in the Paktika province of Afghanistan (data courtesy of ARO). This data set is 3.5 gigabytes on disk and contains 186 million data points over an approximately 4000 m^2 region. This is approximately 6.5 points/ m^2 on average. Because the Afghanistan data set comes from a mountainous region, the data is useful for comparing how different algorithms handle steep slopes and ridges.

The third dataset covers an approximately 600 km^2 region around Fort Leonard Wood in Missouri with a dense point cloud consisting of about 2.2 billion points (data courtesy of ARO) and takes up 57 gigabytes on disk. That is about 3.6 points/ m^2 on average.

The Afghanistan point cloud has not been substantially filtered and contains many non-ground points (such as points on vegetation). For the quality tests we used a subset of the points without many of the non-ground points². This data set, which we refer to as Afghanistan-1, has a point density of 0.26 points per square meter. To test the significance of the density of the points, we removed a portion of the points, keeping only one out of sixteen data points at random. This produces a point density of 0.016 points/ m^2 . We denote this data set as Afghanistan-2.

Parameter choices. Within our algorithm there are numerous parameters that can be adjusted to shift speed and quality trade-offs. The algorithm's precision can be modified through its scaling parameter s , the number of faces k of the C° 's, and the cone radii r_s and r_q . For our tests we set $k = 6$ and $r_s = 20$ meters. We set r_q based on ρ and w , such that $r_q = s\rho B/2$. In our tests the word size w of the color buffers was set to 32 bits (though this can easily be increased to 128 bits on modern graphics cards), and thus $B = 5$. For the scaling parameter, we tested the algorithm with varying values, but in experiments below we used $s = 5$. These parameters were chosen to offer a sufficiently high quality of output without unnecessarily slowing the implementation.

²This was done by only using the "last return"-points for each pulse.

	Without CUDA		With CUDA	
Grid Resolution (m.)	0.8	2	0.8	2
GPU VORONOI (S)	411	73	76	74
Read \mathbb{C}^1	814	116	N/A	N/A
Draw Query Cones	51	5.84	39	6.96
Read \mathbb{C}^2	875	135	N/A	N/A
BUFFER ANALYSIS	102	9.57	183	0.46
Write points	4.01	0.92	4.2	0.8
Total running time	2289	371	337	105

Table 2. CUDA’s effect on NNI algorithm timing, running on the Afghanistan data set (times in seconds)

Efficiency. We ran the various interpolation algorithms on our Afghanistan, Fort Leonard Wood and DKPART data sets. For these general tests, we used a grid resolution of 2 meters. As shown in Table 1, our NNI implementations run significantly faster than the linear interpolation and the RST based interpolation. The NNI algorithm takes only 17% of the time of the linear interpolation for both the Afghanistan and DKPART data sets, and only 11% of the time on the Fort Leonard Wood data set. In comparison with the RST algorithm, the NNI implementation takes only 2.8% of the time on the Afghanistan data set and approximately 1.8% of the time on both the DKPART and Fort Leonard Wood data sets.

Worth noting is the breakdown of the time spent binning the data and the time spent performing the interpolation. On all three data sets when running the CUDA-based implementation, the binning takes a little under half of the running time of the algorithm. While necessary in some cases, many data sets are already stored in tiles such that binning could be skipped. In contrast, the RST method spends most of the time on interpolation, while the linear-interpolation algorithm spends most of the time on constructing the TIN.

The time results also show the significant advantage provided by CUDA, with the CUDA based implementation taking 10%-20% of the time of the non-CUDA based implementation.

It is worthwhile to compare the bottlenecks in the two different NNI implementations. The experiments were again done on the Afghanistan data set. As shown in Table 2, without using CUDA the bottleneck in the implementation is reading \mathbb{C} to main memory. At both grid resolutions, reading \mathbb{C} takes approximately 70% of the total running time of the interpolation. And, since a smaller grid resolution requires more tiles, there are more read backs on the 0.8 meter resolution and thus the time of the algorithm increases significantly.

With our CUDA implementation, we remove the high cost of GPU memory reads since we are reading back a fraction of the data, but the BUFFER ANALYSIS step is highly dependent on the grid resolution. With a 2 meter resolution, the bottleneck is clearly merely drawing the cones for the Voronoi diagrams, and BUFFER ANALYSIS takes negligible time. With a higher grid resolution of 0.8 meters, the expensive atomic add operation is performed many more times in the BUFFER ANALYSIS step and thus it takes far more time. However, even with the bottleneck shifting to the CUDA step, the algorithm is still much faster with CUDA than without it.

Quality of output. It is, of course, necessary to compare not just the speed but also the quality of these different interpolation schemes. To do this we compared the smoothness of the contour maps from the NNI and linear interpolation over the Afghanistan data set. The entire Afghanistan data set is very dense, with 6.5 points per square meter. Even at a 0.8 meter grid resolution, very few differences in contour maps were seen between the two inter-

polation methods. This was also obstructed by the many trees in the region which cause frequent changes in elevation in the grid. Therefore, to perform these quality tests we used the Afghanistan-1 and Afghanistan-2 data sets.

We ran both the linear interpolation and NNI on the data, with a scaling parameter of 5 and a grid resolution of 2 meters. Finally, we use the GRASS GIS system [15] to compute the contour lines from the interpolated data with a 1 meter increment between contour levels.

We present in Figure 10 sample images from each pair of contour maps. Comparing the maps, it can be clearly seen that at certain points, especially around curves, the linear interpolation from the triangulation produces jagged results while the NNI maintains its smoothness. It is clear that as the input data becomes sparser in Afghanistan-2, the linear interpolation becomes increasingly jagged, while the NNI output remains smooth.

Acknowledgements

We thank Andrew Danner and Peter Hachenberger for contributions to the RST and linear interpolation code that we used. We thank COWI A/S, ARO and NASA for access to data.

References

- [1] P. K. Agarwal, L. Arge, and A. Danner. From point cloud to grid DEM: A scalable approach. In A. Riedl, W. Kainz, and G. Elmes, editors, *Progress in Spatial Data Handling. 12th International Symposium on Spatial Data Handling*, pages 771–788. Springer-Verlag, 2006.
- [2] P. K. Agarwal, L. Arge, and K. Yi. I/O-efficient construction of constrained Delaunay triangulations. In *Proc. European Symposium on Algorithms*, pages 355–366. Springer-Verlag, 2005.
- [3] P. K. Agarwal, S. Krishnan, N. H. Mustafa, and S. Venkatasubramanian. Streaming geometric optimization using graphics hardware. In *In Proc. 11th European Symposium on Algorithms*, pages 544–555. Springer-Verlag, 2003.
- [4] L. Arge. External-memory algorithms with applications in geographic information systems. In M. van Kreveld, J. Nievergelt, T. Roos, and P. Widmayer, editors, *Algorithmic Foundations of GIS*, pages 213–254. LNCS 1340. Springer-Verlag, 1997.
- [5] L. Arge, R. Barve, D. Hutchinson, O. Procopiu, L. Toma, D. E. Vengroff, and R. Wickeremesinghe. *TPIE User Manual and Reference (edition 0.9.01b)*. Duke University, 1999. The manual and software distribution are available on the web at <http://www.cs.duke.edu/TPIE/>.
- [6] D. Blythe. The Direct3D 10 system. *ACM Transactions on Graphics*, 25(3):724–734. ACM, 2006.
- [7] J. Bolz, I. Farmer, E. Grinspun, and P. Schröder. Sparse matrix solvers on the gpu: conjugate gradients and multigrid. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, pages 917–924. ACM, 2003.
- [8] L. Buatois, G. Caumon, and B. Levy. Concurrent number cruncher: a gpu implementation of a general sparse linear solver. *International Journal of Parallel, Emergent and Distributed Systems*, 24(3):205–223. Taylor & Francis, Inc, 2009.
- [9] A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge, and H. Mitasova. TerraStream: from elevation data to watershed hierarchies. In *GIS '07: Proc. of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems*. ACM, 2007.




Figure 10. Contour map comparison. In both figures, the grid used for the red (resp. black) contours were generated from the Afghanistan data using linear (resp. natural neighbor) interpolation.

- [10] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Transactions on Graphics*, 13(1):43–72. ACM, 1994.
- [11] Q. Fan, A. Efrat, V. Koltun, S. Krishnan, and S. Venkatasubramanian. Hardware-assisted natural neighbor interpolation. In *In: Proc. 7th Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 111–120. SIAM, 2005.
- [12] T. G. Farr, P. A. Rosen, E. Caro, R. Crippen, R. Duren, S. Hensley, M. Kobrick, M. Paller, E. Rodriguez, L. Roth, D. Seal, S. Shaffer, J. Shimada, J. Umland, M. Werner, M. Oskin, D. Burbank, and D. Alsdorf. The shuttle radar topography mission. *Reviews of Geophysics*, 45, 5. AGU, 2007.
- [13] N. K. Govindaraju, M. C. Lin, and D. Manocha. Fast and reliable collision detection using graphics processors. In *SCG '05: Proc. of the Twenty-First Annual Symposium on Computational Geometry*, pages 384–385. ACM, 2005.
- [14] N. K. Govindaraju, N. Raghuvanshi, and D. Manocha. Fast and approximate stream mining of quantiles and frequencies using graphics processors. In *SIGMOD '05: Proc. of the 2005 ACM SIGMOD International Conference on Management of Data*, pages 611–622. ACM, 2005. ACM.
- [15] GRASS Development Team. GRASS GIS homepage. <http://www.baylor.edu/grass/>.
- [16] K. E. Hoff, III, J. Keyser, M. Lin, D. Manocha, and T. Culver. Fast computation of generalized voronoi diagrams using graphics hardware. In *SIGGRAPH '99: Proc. of the 26th Annual Conference On Computer Graphics and Interactive Techniques*, pages 277–286. ACM/Addison-Wesley Publishing Co, 1999.
- [17] E. Kilgariff and R. Fernando. *The GeForce 6 series GPU architecture*, pages 471–491. Addison Wesley Professional, 2005.
- [18] W. Li, Z. Fan, X. Wei, and A. Kaufman. Gpu-based flow simulation with complex boundaries. In M. Pharr and R. Fernando, editors, *GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation (Gpu Gems)*, pages 747–764. Addison-Wesley Professional, 2005.
- [19] L. Mitasa and H. Mitasova. Spatial interpolation. In P. Longley, M. F. Goodchild, D. J. Maguire, and D. W. Rhind, editors, *Geographical Information Systems: Principles, Techniques, Management and Applications*, volume 1, pages 481–492. Wiley, 1999.
- [20] H. Mitasova and L. Mitasa. Interpolation by regularized spline with tension: I. theory and implementation. *Mathematical Geology*, 25:641–655. Springer-Verlag, 1993.
- [21] J. Montrym and H. Moreton. The GeForce 6800. *Micro, IEEE*, 25(2):41 – 51. IEEE, 2005.
- [22] NVIDIA. CUDA Homepage. <http://nvidia.com/cuda>, 2010. 3.0.
- [23] J. D. Owens, D. Luebke, N. Govindaraju, M. Harris, J. Krüger, A. Lefohn, and T. J. Purcell. A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum*, 26(1):80–113. Wiley, 2007.
- [24] D. Shreiner, M. Woo, J. Neider, and T. Davis. *OpenGL(R) Programming Guide: The Official Guide to Learning OpenGL(R), Version 2.1 (5th Edition) (OpenGL)*. Addison-Wesley Professional, 2005.
- [25] R. Sibson. A brief description of natural neighbour interpolation. In V. Barnett, editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester, 1981.