

# UC Irvine

## ICS Technical Reports

### Title

LOD-based clustering techniques for optimizing large-scale terrain storage and visualization

### Permalink

<https://escholarship.org/uc/item/6h45c9t9>

### Authors

Bao, Xiaohong  
Pajarola, Renato

### Publication Date

2002

Peer reviewed

# ICS

## TECHNICAL REPORT

### **LOD-based Clustering Techniques for Optimizing Large-scale Terrain Storage and Visualization**

*Xiaohong Bao and Renato Pajarola*

UCI-ICS Technical Report No. 02-16  
Department of Information & Computer Science  
University of California, Irvine

June 2002

**Notice: This Material  
may be protected  
by Copyright Law  
(Title 17 U.S.C.)**

Information and Computer Science  
University of California, Irvine



# LOD-based Clustering Techniques for Optimizing Large-scale Terrain Storage and Visualization

Xiaohong Bao and Renato Pajarola  
Department of Information and Computer Science  
University of California Irvine, CA 92697-3425  
{xbao,pajarola}@ics.uci.edu

## Abstract

Large grid-digital terrain data sets used in scientific visualization, GIS and training & simulation applications are far too complex to be rendered at interactive frame rates as a whole, and easily exceed available physical main memory capacity. Therefore, to avoid excessive paging in virtual memory, the terrain data must be maintained on disk and dynamically loaded into main memory as required by the rendering algorithm. Furthermore, the elevation data must be organized in a multiresolution triangulation framework to allow efficient rendering at different levels-of-detail. In this paper, we propose novel clustering algorithms and data structures to map multiresolution terrain data to external memory such that dynamic loading (paging) of elevation data at varying level-of-detail is very efficient and minimizes the number of page faults (I/O).

## 1 Introduction

Interactive visualization of very large scale terrain data in scientific visualization, GIS or simulation & training applications is a hard problem. Because the grid digital terrain elevation models are not only too large to be rendered in real-time but also exceed physical main memory capacity, traditional in-memory multi-resolution triangulation and rendering techniques do not provide a sufficient solution. Relying only on virtual memory and the operating system's paging mechanism does not sufficiently take into account spatial as well as level-of-detail (LOD) relations in a multi-resolution triangulation framework. For rendering large terrain from out-of-core, the multi-resolution data structure and rendering algorithm themselves must provide an efficient paging of LOD data from disk.

In [8] several different multidimensional access methods and spatial indexing structures are discussed. However, those methods are not very suitable for terrain data. Vertices in terrain data cannot be physically clustered only based on their spatial location but must also be clustered based on the LOD and the triangle mesh connectivity. For similar reasons, the clustering technique presented in [5] that minimizes the external path length for disk based data structures is not directly applicable to multi-resolution terrain data. Our approach is to incorporate the LOD (geometric approximation error) of vertices into the process of page mapping, integrating three constraints: *space constraint*, *vertex dependency constraint* and the *LOD constraint*.

As main contributions of this paper we propose novel clustering algorithms and data structures to map terrain data to secondary memory (disk blocks) such that dynamic loading (paging) of elevation data is very efficient and minimizes the number of page faults (I/O). The multi-

resolution triangulation framework we use is a hierarchical quadtree based terrain triangulation method.

The paper is organized as follows: related work is discussed in Section 2, in Section 3 our clustering techniques are introduced, the physical data file storage structure is given in Section 4, in Section 5 we present experimental results, and Section 6 concludes the paper.

## 2 Related Work

### 2.1 Multi-resolution triangulation

Many mesh simplification and multi-resolution triangulation methods have been developed over the last decade. We refer to the literature for overviews on general mesh simplification and multi-resolution modeling [9, 4, 14] and only point out closely related multi-resolution terrain triangulation methods here.

For grid-digital terrain data sets hierarchical quadtree [12, 16] or bin-tree [6, 3, 17, 7] based multi-resolution triangulation methods have shown to be exceptionally efficient. These methods all define the same class of triangulations on a regular rectilinear grid of points but differ in algorithms, data structures and error metrics used to efficiently generate and render an adaptively triangulated terrain surface. We adopt the notion of a quadtree based triangulation as in [16] that uses the dependency relation presented in [12] to guarantee crack-free conforming LOD triangulations. Figure 1 shows the basic recursive triangulation of the quadtree hierarchy.




Figure 1: Two-stage subdivision of the hierarchical quadtree triangulation shown in a), b) and c). An adaptive conforming triangulation example shown in d).

Figure 2 shows the dependency relations that have to be satisfied when generating an adaptive triangulation. If a vertex is selected for a particular LOD triangulation then also the vertices pointed to by the dependency graph have to be selected to guarantee a conforming triangulation. Note that these dependencies always propagate from finer to coarse grid resolutions.

The object space error metric  $\delta$  used in this paper is the vertical distance between corresponding vertices in the original and the simplified terrain surface. For view-dependent triangulation and rendering an




Figure 2: Parent-child dependency shown in a) and c), neighboring vertex dependency shown in b) and d).

image-space error metric  $\rho$  is used which is a perspective projection of  $\delta$ .

## 2.2 Visualization from external memory

While efficient out-of-core algorithms and data structures have been studied in the literature [2, 1] and applied to some extent in scientific visualization (see special issue journal [15]), only little work has been done for multi-resolution terrain rendering from external memory. In [16] a multi-resolution terrain quadtree is maintained in an object-oriented database for efficient out-of-core access, however, the physical clustering is mostly left to the database system. In [10] the terrain data is partitioned and clustered into square tiles without taking any LOD information into account. Furthermore, this approach does not cluster data into fixed size disk blocks. In [13] the multi-resolution terrain triangulation hierarchy is linearized into an array and a memory-mapped file mechanism (supported by the operating system) is used to provide efficient out-of-core access. The main drawbacks are that the terrain data is only clustered on disk with respect to the linearization of the triangulation hierarchy and that the storage cost is comparatively high.

## 3 Clustering Algorithms

In interactive terrain visualization, whether or not a vertex or triangle is rendered not only depends on its spatial location with respect to the current view frustum (*space constraint*) but also is determined by its approximation error and a given image- or object-space error threshold (*LOD constraint*). Considering the adaptive triangle mesh refinement, it is very likely that not all vertices within a query window (i.e. the view frustum) are required for rendering. For view-dependent terrain triangulation and rendering only vertices with a projected image-space error  $\rho$  larger than a user defined threshold are needed. To achieve a conforming triangulation additional vertices may have to be selected (*triangulation dependency constraint*).

If we want to reduce terrain access time from external memory we must avoid loading too many vertices which are not used for rendering. To satisfy the space constraint terrain data can be organized in a spatial index structure. However, within the spatial region of interest vertices with projected image-space error  $\rho$  below the given threshold should not be loaded from disk. Previously this was only addressed by ordering terrain data by level within the hierarchy. Unfortunately the error of a vertex cannot be expressed by its level information precisely. Furthermore,  $\rho$  is view-dependent and cannot be used for clustering vertices. However,  $\rho$  is strongly related to the object space approximation error  $\delta$  which can be precomputed. The LOD constraint can thus be satisfied partially by grouping vertices on disk with respect to  $\delta$ .

As shown in Figure 3, suppose each page can hold 21 quadtree nodes. We can cluster vertices in set A, a small balanced sub-quadtree, into one page (tiling strategy). Thus the error variation among vertices is ignored. Another way is to cluster sets B, C, D, E into one page of the same size. This scheme not only aims at spatial clustering but also

limits the error variation within one page. Accessing all vertices with error greater than 70, the first approach loads 8 pages (page of A and 7 more pages just below A) while the second method only loads one page (containing the relevant sets B, C, D, E).




Figure 3: Clustering of quadtree nodes with associated object space error  $\delta$ , and page capacity of 21 nodes. Set A contains 21 nodes representing a complete three-level subtree of the quadtree hierarchy. Sets B, C, D and E together represent a more flexible clustering of 20 nodes.

Based on the above observations we propose clustering techniques for quadtree based multiresolution terrain triangulation incorporating spatial location,  $\delta$  values and dependency relation to minimize disk access times and page faults.

### 3.1 Basic clustering technique ( $C_{basic}$ )

Let us first define some notations that we will use in the remainder of the paper.

**Node** is a quadtree node storing five vertices (*Center, West, South, East and North*) as shown in Figure 4.

**Evenness  $E$**  is a measure of variation of object space error  $\delta$  among a set of nodes or vertices.

**Primitive element (PE)** is a set of nodes with similar object space errors, satisfying some evenness threshold. Each PE is a small complete quadtree, and for storage efficiency we define the minimal PE to be a two-level quadtree consisting of five nodes.

**Linker** is the pointer from a parent PE's leaf node to the child PE's root node (see also Figure 6). *Intra-linker* connect two PEs within the same page, and *inter-linker* connect PEs between different pages.

**Page height** is the total number of pages to be visited from the root to load a set of nodes or vertices.




Figure 4: A, B, C, and D are quadtree nodes. Each node is associated with 5 vertices: *Center, West, South, East and North*

Intuitively, given an evenness threshold  $\epsilon$  we can partition the complete terrain quadtree into a set of PEs and map this set to secondary memory pages. However, this mapping cannot be done arbitrarily. The idea of this mapping is to minimize the page height as well as the total page number for the entire data set. The following issues have to be considered:

- (i). What is the optimal size of a PE? We have defined the lower bound above, and the upper bound is clearly the size of an external memory page.
- (ii). How can the page height be minimized for loading a set of vertices in a given region at a given LOD?
- (iii). How can the terrain quadtree be mapped to a minimum number of pages?

Studies have shown that optimizing (ii) and (iii) is NP-complete even all PEs have the same size. It can be reduced to the bin-packing problem easily. The detailed proof can be found in [5].

To address point (i) we grow PEs according to an evenness threshold  $\epsilon$ . Starting with a minimal PE we grow level by level until the evenness  $E$  of the leaf nodes of the current PE exceeds  $\epsilon$  or the size of the PE exceeds the current page size. Based on the actual traversal of the multiresolution terrain quadtree we introduce the following heuristic for point (ii). If a node is visited and selected for triangulation its children and neighbors will have the highest possibility to be selected too. Therefore, we should merge PEs which are close in space and LOD into the same page. To address point (iii), if the PE stored in a page cannot be grown anymore and the remaining free page memory is larger than a minimal PE, we start a new PE in the same page. This strategy guarantees that the wasted memory per page is not larger than the size of a minimal PE (a negligible fraction for reasonable page sizes). The pseudo-code of the algorithm is given in Figure 5.

Each time before generating a new PE, the current disk page is checked whether it has enough room for a new PE or not. If not, the current page is finished and a new empty page becomes the current page to be filled. A new minimal PE is generated starting at an unprocessed node from a priority queue, or if it is empty, from the set of unprocessed quadtree nodes. Siblings of this PE root node are pushed into the priority queue  $Q_{page}$ . The current PE is grown level by level until the current page is full, the leaf level of the quadtree is reached, or the evenness threshold  $\epsilon$  is exceeded. Upon completion of a PE its children nodes are pushed into  $Q_{page}$ .

The priority queue  $Q_{page}$  stores the potential root nodes for new PEs within the same page.  $Q_{page}$  is ordered based on the nodes' object space errors in relation to the current page's first PE. This strategy assures that PEs in the same page are not only close in space but have also little difference in object space error. Upon completion of a page this priority queue is emptied.

Furthermore, to preserve memory locality for efficient retrieval it is desirable that consecutive pages on disk store vertices and nodes closely related in space and LOD. We implement this by top-down depth-first traversal of the terrain quadtree to get the root node of the first PE of each page. Another advantage of the depth-first traversal is that we do not need to hold the entire quadtree in main memory. Instead, only the nodes located along the same path from the root node to the current page are maintained in main memory, making this a real out-of-core pre-process.

Finally, let us define the evenness  $E$  used above for PE generation. We define absolute  $E_{abs}$  and relative evenness  $E_{rlt}$  as follows:

$$E_{abs} = \frac{(\delta_{max} - \delta_{min})}{\delta_{max}} \quad (1)$$

$$E_{rlt} = (\delta_{node} - \delta_{child}) - (\delta_{root} - \delta_{node}) \quad (2)$$

For a PE,  $\delta_{max}$  and  $\delta_{min}$  are the max/min object space errors of all unprocessed nodes just below of the PE.  $\delta_{node}$  is the object space error of one node below the current PE, and  $\delta_{child}$  is the min object space

```

Algorithm Generate-Primitive-Element(quad-tree t)
{
  if(there is no room in the current page for a new primitive element)
    return (the current page is full);
  Node n = Get-Primitive-Element-Root(t) ; //start a new PE.
  if(n==NULL)
    return (the whole quad-tree is finished);
  Construct a minimum PE c rooted at n;
  store n's unprocessed siblings into the priority queue q
  as the potential roots for the next new PEs.
  while(check-next-level(t, c))
    //the current primitive element is growing
    expand the PE c to the next level;
  store c's children into the priority queue q
  as the potential roots for the next new PEs;
  return c ;
}

Algorithm Get-Primitive-Element-Root(quad-tree t)
{
  if(the priority queue q is not empty)
    return (dequeue(q)) ;
  If there is(are) node(s) in t which is not processed
    return the first node searched by Depth-first-search;
  else
    return NULL;
}

Algorithm check-next-level(quad-tree t, PE c)
{
  if (the current level of c is leaf level)
    return false ; //stop the current PE.
  if(there is no room in the current page for the next level)
    return false ; //stop the current PE.
  if(the next level is leaf level)//load leaf level anyway.
    return true ;
  if(the evenness of the next level of c meets the requirement)
    return true ; //grow to the next level.
  return false ; //stop the current PE
}

```

Figure 5: Pseudo-code for generating Primitive Elements

error of the four children of this node.  $\delta_{root}$  is the root node object space error of the current PE.

$E_{abs}$  measures the evenness among the nodes of the next level below the current PE. If  $E_{abs}$  is equal or less than a given threshold  $\epsilon$  then the current PE is grown by one more level. The relative evenness  $E_{rlt}$  measures whether a node in the next level is closer to its children or closer to the root of the current PE in terms of object space error. If  $E_{rlt}$  of all nodes just below the current PE is equal or less than zero then the current PE is grown by one more level.

### 3.2 Optimized clustering technique ( $C_{opt}$ )

We observed from experiments as shown in Figure 7 that around 25% or more of the vertices in any frame are selected not based on their LOD but only due to triangulation dependency constraints (see Figure 4). The *parent-child dependency* enforces that when a vertex from a child node is selected also vertices from the parent node are selected, and is accounted for in the above clustering technique. However, the *neighboring dependency* that propagates vertex selection to sibling nodes is

not incorporated. We introduce a subtle change in the clustering strategy here to improve clustering by grouping sibling nodes.

As shown in Figure 6 a minimal PE now includes four sibling nodes, each being the root of a small two-level quadtree. Growing of PEs with this optimized clustering method  $C_{opt}$  is similar to the basic method  $C_{basic}$  discussed above, all child nodes of the current leaf nodes of the PE are added level by level.




Figure 6: Logical clustering structure of similar size PEs according to strategies  $C_{basic}$  and  $C_{opt}$ . In  $C_{basic}$  a PE is a small complete quadtree with  $n$  levels with each leaf node having 4 linkers pointing to the respective child PEs. In  $C_{opt}$  a PE contains 4 sub-quadtrees of  $n - 1$  levels of sibling nodes with each leaf node having only 1 linker to the PE containing its 4 children.

There are two advantages to the  $C_{opt}$  clustering strategy. First, it reduces the number of linkers between PEs to 1/4, thus reducing storage cost significantly. Second, in 50% of the cases the neighboring dependency relation points to within the same PE, thus reducing I/O cost for loading PEs to resolve triangulation dependencies. Our experimental results in Section 5 show the effect of these two advantages and the resulting performance increase of this subtle change from  $C_{basic}$  to  $C_{opt}$ .




Figure 7: Shows the percentage of vertices which are loaded due to their screen projection error  $\rho$  exceeding the threshold  $\tau = 0.015$  over a test-run of 12000 frames. Thus about 20%, and up to 60% in the worst case, are loaded due to triangulation dependencies only.

## 4 Storage structure

As shown in the Figure 6, each PE contains one ( $C_{basic}$  method) or four ( $C_{opt}$  strategy) small complete quadtrees. Each PE stores information on its spatial location (its SW corner), number of quadtree levels it covers, and boundary type (touches none, the East, South, or both boundaries). Furthermore, a PE stores the maximum object space error and a bounding box of all its descendant vertices. Conservative bounding boxes of nodes are dynamically computed at rendering time based on their position in the quadtree within a PE (only geographical extent,

elevation range is kept constant for all nodes within the same PE due to the small variation).

Each node in the leaf level of a PE contains 4 linker elements for  $C_{basic}$ , respectively 1 for  $C_{opt}$ . The leaf nodes of the entire terrain quadtree do not store any linker pointers. For I/O efficiency, the object space error and a bounding box of the child PE are both stored with the *inter-linker* element in the parent PE. Therefore, pages are never loaded in vain only to check for LOD or visibility. On the other hand, the *intra-linkers* only store a pointer to the child PE because accessing the child PE within the same page does not cause any additional I/O cost.

The nodes of each PE are mapped to the disk page ordered by level. Considering redundancy, only three vertices (*Center*, *West* and *North*) of each node are actually stored with the node itself (see also Figure 4). The other vertices *East* and *South* are stored in sibling nodes. Only at the boundary of the terrain data set we actually store all vertices with a node to avoid degenerate nodes along the boundary. Since selection of any non-center vertex causes both adjacent nodes to be loaded due to dependency constraints as shown in Figures 2 and 4, this storage layout will eliminate any vertex storage redundancy and does not increase I/O cost. For example, in Figure 4 if the *East* vertex  $V$  of node A is to be loaded then also the *Center* vertex of node B is required and thus node B is loaded which actually stores the initially requested vertex  $V$  (as its *West* vertex).

Each vertex only stores the terrain elevation value and its object space error  $\delta$ .

The entire data file consists of a file header and a set of disk pages as shown in Figure 8. The file header contains the following information: the page size, the total number of quadtree levels of the entire terrain data set, the resolution of the elevation values, the geographical range of the terrain data, the elevation values of the bounding box corners, the information about the top-level quadtree root node, and the four pointers to the child PEs of the root node for  $C_{basic}$  (or one pointer for  $C_{opt}$ ). Each page stores the number of PEs as the only meta information.




Figure 8: Logical storage layout of data file. A Node contains at least the vertices *West*, *Center* and *North* (and also *South*, *East* for boundary nodes). PEs contain 4 ( $C_{basic}$ ) or 1 ( $C_{opt}$ ) linker elements.

## 5 Results

The main objectives of the conducted experiments are to show the effectiveness of our clustering techniques and to determine the influence of the adjustable parameters. For this we implemented a simple terrain viewer that simulates a series of user operations such as changing the LOD parameter or viewing area. The viewer operates on our external memory terrain data structure and dynamically loads elevation data from disk as required by the LOD or viewing parameters. While not optimized for speed, the triangulation and rendering algorithm incorporates simple view-dependent vertex selection as outlined in the following section.

For comparison, we also implemented the quadtree-based indexing scheme and file mapping of [13] (with embedded *white quadtree*). Ad-

ditionally, we implemented a hierarchical tile partitioning with tiles of  $(2^k + 1) \times (2^k + 1)$  vertices and  $(2^k)^2$  child nodes similar to the VGIS approach [11].

The following three performance aspects are studied:

- (i). File sizes of entire data set.
- (ii). Number of page faults (i.e. how many pages are accessed and retrieved, total I/O cost).
- (iii). Loading and processing times of pages (relative I/O cost with respect to total rendering cost).

## 5.1 Experimental environment

The hardware configuration is a one-processor 1.5GHz Pentium IV PC with the Microsoft Windows 2000 operating system and 1.5GB of main memory (actually used main memory for elevation and triangulation data was 20MB).

The source terrain data used for our statistical experiments is  $(2^{13} + 1) \times (2^{13} + 1)$  (67M vertices) ranging from latitude  $35^\circ$ , longitude  $-120^\circ$  to latitude  $41^\circ 49' 37''$ , longitude  $-113^\circ 10' 20''$  at 250 meters horizontal and 1 meter vertical resolution. (From [http://edc.usgs.gov/glis/hyper/guide/1\\_dgr\\_demfig/nj11.html](http://edc.usgs.gov/glis/hyper/guide/1_dgr_demfig/nj11.html)).

Elevation data and object space errors  $\delta$  are 4-byte values, and the bounding box occupies 8 bytes. In our implementation of a tiled terrain partitioning and the simulation of the approach presented in [13] each vertex occupies 16 bytes (compared to 20 bytes in [13]).

The basic algorithm we use for terrain mesh generation and rendering is the recursive top-down vertex selection algorithm from [16]. Our implementation is not optimized for rendering speed but rather focused on efficient I/O and dynamic terrain loading from external memory. The viewer application incorporates view-dependent vertex selection that includes view-frustum culling according to the bounding box information of quadtree nodes, and a screen projection error threshold  $\tau$ . From the geometric approximation error  $\delta$ , the image-space error is computed by

$$\rho = \frac{\delta}{\|p - e\|_2}, \quad (3)$$

with  $p$  being the vertex position and the viewpoint position  $e$ . For back-tracking we use  $\rho = \frac{\delta}{\|p - e\|_2 - r}$  at *Center* vertices of nodes with bounding sphere radius  $r$ .

We implemented our approaches with clustering strategies  $C_{basic}$  and  $C_{opt}$  as discussed in Section 3. They are represented by PE1 and PE2 respectively in the below figures. With *LP* we denote the embedded quadtree layout and indexing proposed in [13], and *Tile* denotes a regular hierarchical terrain partitioning. All approaches are implemented in the same framework using the same main memory rendering algorithm and only differ in the indexing and physical data layout on disk. We believe this provides a fair comparison of page faults and other I/O measures such as data locality or vertex utility rate.

For page loading experiments we limited the main memory usage of our test application to 20MB for storing elevation data, and pages are freed according to a least-recently-used caching strategy. If not specified otherwise the viewpoint is at 20000 meters altitude with a FOV of  $45^\circ$ . For fly-through tests the flight path consisted of four straight line segments between points  $(38^\circ 24' 48'', -116^\circ 35' 11'')$ ,  $(38^\circ 24' 48'', -119^\circ 55' 11'')$ ,  $(35^\circ 4' 48'', -119^\circ 55' 11'')$  and  $(35^\circ 4' 48'', -116^\circ 35' 11'')$  in latitude, longitude coordinates. The fly-through test rendered 12000 frames in total.

Page size(bytes)	LP	Tile	$C_{basic}$	$C_{opt}$
1024	1.33GB	N/A	0.711GB	0.518GB
2048	1.33GB	2.00GB	0.662GB	0.5GB
5120	1.33GB	1.25GB	0.606GB	0.576GB

Table 1: File sizes of different clustering approaches for varying page sizes. Our clustering techniques  $C_{basic}$  and  $C_{opt}$  used an absolute evenness threshold of  $\epsilon = 0.2$ .

## 5.2 Experimental results

### 5.2.1 Data file size

In Table 1 we compare the file sizes resulting from different partitioning and clustering techniques with respect to different page sizes. For the hierarchical *Tile* approach the file size greatly depends on the page size because the tile size is limited by the disk page size, and because the occupancy rate is related to the combination of page and tile size. A high occupancy rate of 90.3% can only be achieved with a large page size of 5K that stores a  $17 \times 17$  tile per page. The file size of the *LP* approach is independent of the page size and only depends on the size of vertex elements (16 bytes) and the occupancy rate. In comparison, the approach in [10] requires 50 bytes per vertex (the 70 bytes reported in the paper include some additional information not used in here).

The file sizes of our clustering techniques and storage layout depend to some (minor) extent on the actual page size. That is because of variations in the sizes of PEs and the number of linkers. Furthermore, since the wasted memory per page is at most the size of a minimal PE the occupancy rate is very high even for small pages and increases with page size.

From Table 1 we can observe that the file sizes of our approach are significantly smaller than the other two approaches despite some added extra information such as linker elements. Besides the high occupancy rate and otherwise compact representation of PE the main reason for the space savings is that bounding box information is only recorded for each PE rather than for every single vertex. As mentioned in Section 4 bounding boxes of nodes are dynamically computed from the PE bounding box only at run-time for efficient back-tracking of the top-down vertex selection algorithm. Note that the error introduced by this conservative bounding box measure is very little as shown below.

First, for the same screen projection error threshold  $\tau$  more triangles are generated by our techniques because of the conservative bounding box computation. Second, the ratio of extra triangles is no more than 2.5%. Thus the significant space savings of 50% or more come only at the expense of a very slight increase of number of rendered triangles per frame.

The strategy of storing the bounding box information only with specific elements or nodes cannot be applied to the *LP* [13] method, in particular because the memory mapped file organization does not allow array elements of varying size.

### 5.2.2 Static I/O

Figure 9 shows the number of page faults for a static view frustum under different screen projection error thresholds  $\tau$ . This test measures how many pages have to be loaded from external memory to generate an adaptive triangulation satisfying the image-space tolerance  $\tau$ . Our  $C_{opt}$  clustering strategy achieves best results with fewer page faults for any threshold  $\tau$  and saves almost 50% page faults when  $\tau$  is small, compared to *LP*. The *LP* approach achieves similar results as  $C_{basic}$  for larger threshold values  $\tau > 0.05$ . But  $C_{basic}$  performs much better than *LP* for small tolerances  $\tau < 0.05$ .

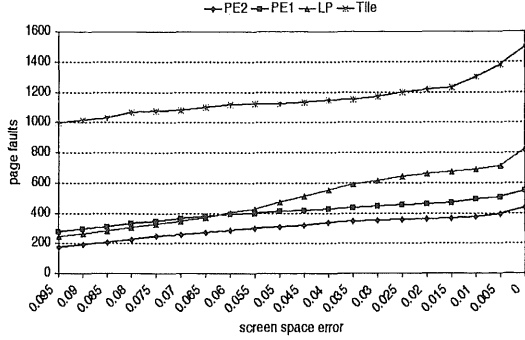


Figure 9: Page faults for different screen projection error thresholds  $\tau$  to render the fixed view centered at ( $lat : 38^{\circ} 15' 48''$ ,  $lon : -116^{\circ} 35' 11''$ ) with FOV of  $45^{\circ}$ .

### 5.2.3 Dynamic I/O

In this section we report tests of dynamic access to the terrain data based on a fly-through sequence as outlined in Section 5.1 consisting of 12000 rendered frames. If not specified otherwise, the absolute evenness threshold is set to  $\varepsilon = 0.2$  and the page size is 2KB.

Figure 10 shows the cumulative number of page faults for each of the 12000 frames of a fly-through test. We can observe that also for continuously changing viewpoints our clustering techniques show constantly better page fault performance and  $C_{opt}$  has about half as many page faults as LP.




Figure 10: Accumulated number of page faults for a fly-through test with 12000 frames and  $\tau = 0.015$ .

Figure 11 shows the cumulative page faults for different image-space error tolerances  $\tau$ . It can be seen that the relative advantages of our clustering techniques are largely maintained for varying thresholds, only for large values of  $\tau$  the LP approach reaches the same or better efficiency as  $C_{basic}$ .

In Figure 12 the cumulative page faults are related to different page sizes. We can observe that the page size does not affect the relative page fault performance among the different techniques. Our optimized clustering technique  $C_{opt}$  has constantly fewer page faults than all other approaches irrespective of which parameters are changed. The hierarchical Tile approach is always the worst.

In Figure 13 we measured the loading times of pages from external memory for the different clustering techniques for a fly-through test. Assuming that fixed sized pages require constant loading time it is to be expected and shown in Figure 13 that the cumulative loading time statistic is strongly related to the cumulative page faults of Figure 10.




Figure 11: Accumulated number of page faults for different screen projection thresholds  $\tau$ .

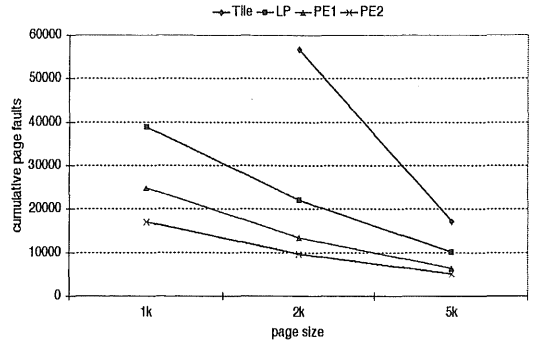


Figure 12: Accumulated number of page faults for different page sizes and  $\tau = 0.015$ .

We can observe that  $C_{opt}$  and also  $C_{basic}$  are both more than two times faster than the LP technique. Compared to the page faults shown in Figure 10, this improved I/O efficiency is achieved by a good physical memory locality of the elevation data on disk. From a different viewpoint we can say that  $C_{basic}$  loads 13274 pages within 10 seconds, and  $C_{opt}$  loads 9677 pages within 7.9 seconds. In contrast, the LP method only loads 7200 pages in 10 seconds or 5473 pages in 7.9 seconds. The loading times per page for  $C_{opt}$ ,  $C_{basic}$ , LP and Tile are 0.76ms, 0.82ms, 1.13ms and 0.63ms respectively.




Figure 13: Accumulated loading times for a fly-through test with 12000 frames and  $\tau = 0.015$ .

Figure 14 demonstrates the cumulative time of data retrieving (including data loading and processing) for the different clustering approaches. Here the data processing in run time consists of two sub-






Figure 14: Cumulative loading and data processing times for a fly-through test with 12000 frames and  $\tau = 0.015$ .

processes: deriving bounding sphere information from bounding box information and computing global coordinate  $(x, y, z)$  for each vertex. By comparing Figure 14 & Figure 13, we can see that the total *retrieving* time of using  $C_{basic}$  or  $C_{opt}$  is even less than the pure *loading* time of using LP or Tile. In LP, part of run-time processing can be avoided using the data structure proposed in [13]. However each vertex will occupy more space, thus increasing page faults.

Figure 15 shows the dependency of cumulative loading times of a fly-through test on the evenness parameter  $\epsilon$ . The corresponding cumulative page faults are about 13200 and 9700 for  $C_{basic}$  and  $C_{opt}$  respectively. While the page faults stay similar for different  $\epsilon$  the loading times change noticeably. From this we can conclude that the evenness parameter has some impact on the resulting physical data locality on external memory (affecting latency and seek times). Best results have been achieved using the absolute evenness threshold  $\epsilon = 0.2$ .




Figure 15: Cumulative loading times of a fly-through with  $\tau = 0.015$  for different absolute evenness thresholds  $\epsilon = 0.1, 0.2, 0.3, 0.5$  and using the relative evenness (rlt) for  $C_{opt}$  and  $C_{basic}$  clustering strategies.

In Figure 16 we recorded the number of loaded vertices during the fly-through test and compare this to the number of actually selected and rendered vertices. Utility rate of total loading vertices by using  $C_{opt}$  or  $C_{basic}$  is higher than that of others when  $\tau < 0.04$ . After that, the utility rate of LP is slightly better than that of  $C_{opt}$  or  $C_{basic}$ .

## 6 Conclusion

In this paper we present novel clustering algorithms and data structures to support interactive large scale terrain visualization from out-of-core.

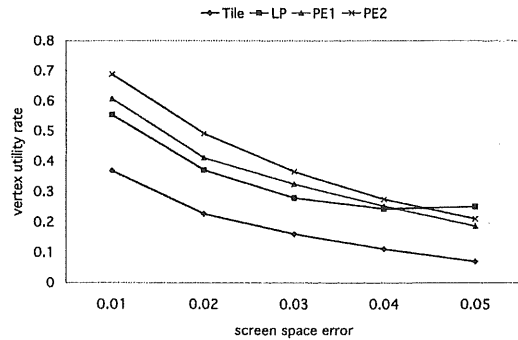


Figure 16: Utility rate of loaded vertices during a fly-through test for different screen projection error thresholds  $\tau$ .

Frame Number	#1	#650	#1000	#1321	#2000
$C_{opt}$	435	1588	2133	3064	4528
$C_{basic}$	645	2089	2782	3944	6205
LP	950	2736	3730	5279	8053

Table 2: Cumulative page faults of different clustering techniques at different frames from the same-route-fly-over sequences, see pictures in Figure 18

Our clustering techniques allow efficient paging of terrain data at different LOD from disk with minimal number of page faults since our clustering techniques are more closely related to space and LOD constrained access patterns. From our experiments we can draw the following conclusions:

1.  $C_{opt}$  is the best clustering technique among the tested methods under all conditions.
2. Our clustering techniques make the data file have a better property of data locality.
3. In most cases, by using  $C_{opt}$  or  $C_{basic}$  approach, the total page faults (I/O) can be reduced to less than 20% (compared to the hierarchical tile structure) or less than 50% (compared to the LP algorithm). And the total data loading time can be reduced to less than 30%.
4. Because  $C_{opt}$  and  $C_{basic}$  approaches can reduce the total page faults so greatly, more benefits will be gained when the large scale terrain visualization applications are distributed.

To our knowledge this is the first work to address the clustering of elevation data into disk blocks according to spatial location as well as LOD information.

## Acknowledgments

This work was partially supported by NSF grant CCR-0119053. We would also like to recognize and thank the US Geological Survey from where all the digital elevation models originate that are used in this paper.

## References

- [1] J. Abello and J. S. Vitter. *External Memory Algorithms*. American Mathematical Society, Providence, R.I., 1999.




Figure 17: Static view at the same location with different screen space errors. when  $\tau = 0.015$ , the page faults of using  $C_{opt}$ ,  $C_{basic}$ , LP and Tile are 1639, 2157, 2501, and 5880 respectively; when  $\tau = 0.035$ , the page faults of using  $C_{opt}$ ,  $C_{basic}$ , LP and Tile are 958, 1406, 1214, and 4072 respectively.




Figure 18: Frames from fixed-route fly-over sequences of using  $C_{opt}$ ,  $C_{basic}$  and LP respectively.  $\tau = 0.015$  and the viewpoint height is 20000m. See Table 2 for the page faults of each of them.

[2] L. Arge. *Efficient External-Memory Data Structures and Applications*. PhD thesis, Department of Computer Science, University of Aarhus (Denmark), 1996.

[3] L. Balmelli, S. Ayer, and M. Vetterli. Efficient algorithms for embedded rendering of terrain models. In *Proceedings IEEE Int. Conf. on Image Processing ICIP 98*, pages 914–918, 1998.

[4] P. Cignoni, C. Montani, and R. Scopigno. A comparison of mesh simplification algorithms. *Computers & Graphics*, 22(1):37–54, 1998.

[5] A. Diwan, S. Rane, S. Seshadri, and S. Sudarshan. Clustering techniques for minimizing external path length. In *Proceedings of the 22th VLDB Conf.*, pages 342–353, 1996.

[6] M. Duchaineau, M. Wolinsky, D. E. Sigeti, M. C. Miller, C. Aldrich, and M. B. Mineev-Weinstein. Roaming terrain: Real-time optimally adapting meshes. In *Proceedings IEEE Visualization 97*, pages 81–88, 1997.

[7] W. Evans, D. Kirkpatrick, and G. Townsend. Right-triangulated irregular networks. *Algorithmica*, 30(2):264–286, March 2001.

[8] V. Gaede and O. Günther. Multidimensional access methods. *ACM Computing Surveys*, 30(2):170–231, 1998.

[9] P. S. Heckbert and M. Garland. Survey of polygonal surface simplification algorithms. SIGGRAPH 97 Course Notes 25, 1997.

[10] H. Hoppe. Smooth view-dependent level-of-detail control and its application to terrain rendering. In *Proceedings IEEE Visualization 98*, pages 35–42. Computer Society Press, 1998.

[11] P. Lindstrom, D. Koller, W. Ribarsky, L. Hodges, and N. Faust. An integrated global GIS and visual simulation system. Technical Report GVU Technical Report 97-0, Georgia Tech Research Institute, 1997. <http://www.gvu.gatech.edu/gvu/virtual/VGIS/>.

[12] P. Lindstrom, D. Koller, W. Ribarsky, L. F. Hodges, N. Faust, and G. A. Turner. Real-time, continuous level of detail rendering of height fields. In *Proceedings SIGGRAPH 96*, pages 109–118. ACM SIGGRAPH, 1996.

[13] P. Lindstrom and V. Pascucci. Visualization of large terrains made easy. In *Proceedings IEEE Visualization 2001*, pages 363–370. Computer Society Press, 2001.

[14] D. P. Luebke. A developer’s survey of polygonal simplification algorithms. *IEEE Computer Graphics & Applications*, 21(3):24–35, May/June 2001.

[15] K.-L. Ma. Large-scale data visualization. *IEEE Computer Graphics & Applications*, 21(4):22–23, July-August 2001.

[16] R. Pajarola. Large scale terrain visualization using the restricted quadtree triangulation. In *Proceedings IEEE Visualization 98*, pages 19–26, 515, 1998.

[17] L. Velho and J. Gomes. Variable resolution 4-k meshes: Concepts and applications. *Computer Graphics Forum*, 19(4):195–214, 2000.