

Interactive rendering of massive terrains on PC clusters

V. Gouranton¹, S. Madougou^{1,2}, E. Melin¹ and C. Nortet¹

¹Laboratoire d'Informatique Fondamentale d'Orléans, LIFO Orléans, France
²BRGM, France

Abstract

We describe a parallel framework for interactive smooth rendering of massive terrains. We define a parallelization scheme for level of detail algorithms in cluster-based environments. The scheme relies on modern PC clusters capabilities to address the scalability issue of level of detail algorithms. To achieve this, we propose an efficient tile-based data partitioning method that allows both reducing load imbalance and solving the well-known border problem. At runtime level of detail computations are performed in parallel on cluster nodes. A hierarchical view frustum culling combined to a compression mechanism harnessing the frame-to-frame coherence are used to drastically reduce the inter-tasks communication overhead. We take into account level of detail algorithms visual quality issue by providing geomorphing and texturing supports. We are able to interactively and smoothly render terrains composed of hundreds of millions to billions of polygons on a cluster of 8 PCs.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Parallel simplification, Virtual Reality, Real-time Rendering, Level of Detail Algorithms

1. Introduction

These last two decades, scientific visualization becomes an attractive way for gaining deep understanding in scientific data. Interactive rendering of these data gives researchers the ability to do faster visual data analysis. Due to strides in data acquisition technologies (high-resolution 3D scanners, cameras, scientific simulation), these data can reach gigantic sizes. For instance, USGS elevation data for Washington state at a horizontal resolution of ten meter and a vertical resolution of ten centimeter represent 1,4 billion elevation values [was]. Visualizing this kind of data in an interactive way is not a trivial task [HDJ04]. As the data obviously do not fit in the system memory, direct rendering cannot be used even with high-end graphics dedicated infrastructures such as SGI machines [MBDM97]. These architectures generally provide memory and compute capabilities that can satisfy many visualization needs. However, their restricted scalability and extensibility makes them not appropriate for the rapid increase of today's datasets. An alternative way is to use level of detail (LOD) algorithms. Given a model to be visualized, their goal is to decrease the size of the data sent through the graphics pipeline by discarding the unnecessary detail while preserving an ac-

ceptable visual quality. Although they induce some CPU load overhead, LOD algorithms give the ability to balance between interactivity and detail for large models visualization. First LOD algorithms require the data to fit entirely in the system memory [LKR*96, Hop96, DWS*97]. This naturally bounds their usage for huge datasets. In order to bypass this bottleneck, many desktop-based LOD algorithms resort to out-of-core mechanisms. This often leads to considerable CPU load and memory footprint overheads. To ensure interactivity by hiding the latency due to the external memory access, they must implement complex prefetching mechanisms. However, any extra time spent in LOD algorithm inexorably results into frame rate degradation, thus into interaction disruption.

For interactive rendering of huge datasets, both LOD techniques and high-performance computing are required. Moreover, the anatomy of super-computers is quickly and deeply changing. Clusters of commodity components are becoming the leading choice architecture. They are now usual in the supercomputer top 500 [top]. They are scalable and modular with a high performance-price ratio. These architectures are proved efficient for classical (not interactive) intensive computations. In scientific visualization field, several work

[ARZ03, GJR03] showed the effectiveness of using PC cluster as a means of complex computations (simulation) while the rendering is performed in a multi-display environment.

By using such an environment for LOD algorithms, the major issues related to massive terrain rendering can be addressed. Indeed, as the number of nodes is theoretically unbounded, using parallelism makes complex computations achievable without significant overhead. In the same way, by distributing the terrain data on cluster nodes system memory, the size of manageable terrains depends only on the number of nodes. However, as for external solutions, one must face the communications overhead induced by bus and network traffics. Currently, efficient inter-connecting networks exist for clusters. Their throughput now outperforms that of today's hard discs. Furthermore, in the particular case of graphical applications, in addition to the LOD algorithm which itself is a data reduction process, ways for further reducing the data transfer can be implemented by harnessing intrinsic graphical properties.

LOD algorithms usually require building a hierarchical structure that encompasses all levels of detail. As this preprocessing may be very time-consuming, the first parallelization attempts have been used to tackle this issue [ESV99, DLR00]. Other approaches really parallelize LOD computations [BW00, GLMM04]. In [BW00] a master/slaves paradigm is proposed. In this scheme, slaves perform the LOD computations and the resulting data are gathered on master for rendering. The master node becomes a bottleneck for the scalability of this approach. In a previous work [GLMM04], we present a cluster-based parallelization framework for height fields that does not suffer from this bottleneck. However, this work presents some weaknesses. No load balancing support is provided. It does not propose a satisfactory solution for the partition boundary simplification issue nor it benefits from frame-to-frame coherence. Moreover, there is no support for visual quality improvement such as geomorphing or texturing.

This paper describes a comprehensive parallelization framework built on the model described in [GLMM04]. Critical issues such as the border problem and load balancing are solved. Essential optimizations that benefit from graphical properties such as culling and frame-to-frame coherence are proposed to drastically reduce the communication overhead. In addition, techniques such as geomorphing and texturing, are used to improve the visual quality of the rendering. The paper is organized as follows: in section 2 we briefly recall our previous work. Section 3 describes the parallelization scheme in details. The way communication overhead is minimized is explained in 4. Geomorphing and texturing are described in 5. Algorithm performance is given through benchmarks in section 6. We conclude and plan for future work in section 7.

2. Background

In [GLMM04], we present a cluster-based parallel simplification framework for height fields (fig. 1). This framework partitions cluster nodes into visualization nodes and LOD nodes. The former do only rendering whereas the latter do only computations. At a preprocessing stage, initial height field data are partitioned into n partitions where n is the number of LOD nodes. During the execution, before each step of LOD computations, computations nodes receive view parameters (view point and frusta) from visualization nodes. Then, they independently perform the LOD algorithm on their partition. Resulting data are culled against visualization nodes frusta before being sent for rendering. The way

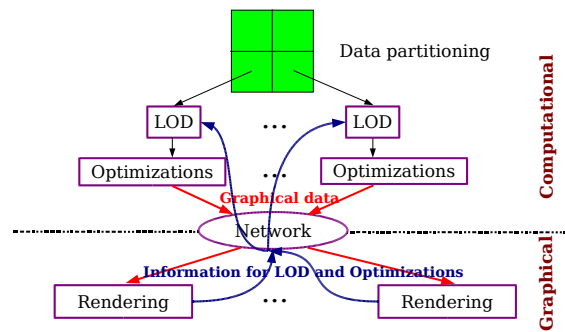


Figure 1: A parallelization model for LOD.

the initial data are partitioned does not allow to reach a satisfying load balancing. As partition is done on a block-per-node basis, load balancing is strongly related to the user's interaction. Moreover, this work does not address the border problem because LOD computations are performed totally independently inside tiles whereas some overlapping data are needed. Indeed, [GLMM04] proposes a parallelization of Roettger's algorithm [RHSS98]. This algorithm dynamically constructs a hierarchical structure, called quadtree, that represents the most relevant data with respect to the current view point and the local terrain topology. The depth of the quadtree at which appears a node is called its LOD value. The subdivision process is controlled by a decision variable f evaluated at each quadtree node (for the sake of clarity, we will call it block). It depends on the block's distance from the current viewpoint as well as the local terrain topology. The local terrain topology is measured as the terrain roughness around the block (called its $d2$ value). Thus, the decision variable of a given block also depends on the adjacent blocks. This information is used to constrain the LOD difference between adjacent blocks to never exceed one. We call this constraint the *fundamental constraint*. As long as this constraint is fulfilled, T-junctions, that cause cracks,

are avoided by skipping drawing central vertices on shared edges between blocks of different LOD values (see figure 2). While using a sequential algorithm, we can easily satisfy this constraint as we hold the entire neighborhood information about each block. Roughness values are then computed in a consistent manner. However, in a parallel framework, adjacent blocks may belong to different partitions. For these blocks, there is a lack of complete neighborhood information. Thus, independently computing $d2$ values inside the partitions, as done in [GLMM04], leads to inconsistencies in border blocks $d2$ values and thereafter to cracks.




Figure 2: Lacking complete neighborhood information leads to cracks between adjacent blocks of different partitions.

3. The Parallelization Scheme

As usual in parallel computing, our algorithm is made up of two phases: a data partitioning phase and a computations phase. Data partitioning is done statically in a preprocessing stage whereas computations step occurs at each frame and may lead to inter-tasks communications. The three key points in writing efficient parallel algorithms, namely inter-tasks communications, load balancing and global coherence, are taken into account. Inter-tasks communication is the topic of the next section. The way data are partitioned takes care of load imbalance and it is discussed in subsection 3.2. The global coherence issue, known as the border problem in the field of LOD, is described and addressed next.

3.1. Addressing the Border Problem

Spatial decomposition is not a novel idea in the LOD field. Many out-of-core algorithms [CRMS02, CGG*03, Hop98, Ulr02, ESC00, YSGM04] partition the initial data into tiles (some authors call them clusters, segments or chunks) to speed up external data access. A known issue related to this partitioning is the border problem. This problem occurs because LOD algorithms need to know about the neighborhood of a particular simplex in order to determine whether it has to be simplified, refined or left unchanged. However, when data are partitioned, this information lacks for border simplices inside the partitions. Several authors have been faced this issue. For instance,

in [Hop98], Hoppe describes one of the first out-of-core LOD algorithms which partitions initial terrain data into rectangular blocks. After simplification, adjacent blocks are stitched together in a hierarchical fashion to form larger blocks. To ensure a conforming stitching, blocks boundaries are left unchanged. While going up in the hierarchy, blocks boundaries are coarsen but at the price of an additional processing. Other authors simply hide the visual artifacts due to the border problem either by using vertical skirts to fill the cracks around the chunks [Ulr02] or by using texturing [BDH00].

In fact, in order to solve the border problem, some information about neighboring simplices are needed for each border simplex [CRMS02]. There are several ways to do this. One way is to fetch the missing information at runtime. This could be costly, especially in a parallel computing context. So, we adopt an alternative way which consists in associating additional neighborhood information into partitions data at the preprocessing stage and then doing some computations to recover the missing information. As stated in section 2, in the LOD algorithm we used, when the fundamental constraint is fulfilled, the only information needed by a block about its neighbors is whether they are refined or not. This information is obtained by computing the decision variable f of those blocks. Thus, we must fulfill two conditions in order to guarantee a solution to the border problem:

1. satisfying the fundamental constraint across partitions boundaries
2. knowing foreign neighbor blocks subdivision state for each border block

To satisfy the first constraint, the $d2$ values must be computed as described in [RHSS98], but on the entire dataset and before the partitioning stage occurs. To satisfy the second condition for a given border block, we must be able to evaluate f for each adjacent block belonging to another tile. Except the $d2$ value, all parameters needed to evaluate the decision variable for adjacent blocks can be computed from the border block. So, in addition to its own $d2$ values, a tile must store the $d2$ values of neighboring tile blocks. These contiguous blocks follow a binary tree pattern, and as neighboring tiles number is at most four, the number of additional $d2$ values to record is $4 * \sum_{i=0}^{n-1} 2^i = 4 * (2^n - 1)$. This represents a few percent of the initial tile data (3% for $n=7$ for instance).

3.2. Data Partitioning

Data distribution is a critical stage in parallel computing. The parallel algorithm efficiency heavily depends on it. The global task and data must be partitioned into subtasks and subsets that are ideally independent and balanced. Apart from embarrassingly parallel algorithms, these goals are difficult to reach. In the LOD field, partitions size must be carefully chosen: they must not be too small nor too large. In




Figure 3: Binary tree of neighboring d_2 values at the border of the tile right edge.

one side, if they are too large, the navigation can introduce load imbalance. On the other side, cutting them too smaller can limit the drastic simplification ability of the algorithm as the number of partitions determines its minimum achievable resolution. Let's assume our initial terrain is a square of side $(2^n + 1)$. On one hand, if terrain is not square, we extend its side to the power of 2 immediately greater than the actual side, the added data being skipped at drawing phase. Next, terrain data are partitioned into square *tiles* of side $(2^m + 1)$, where $m \leq n$. Such tiles perfectly fit to Roettger's algorithm and can be processed on different machines. The number m is chosen such that load imbalance is minimized and the minimal resolution is acceptable. To achieve this, tiles are assigned to the computation nodes using a round-robin pattern. Thus, whatever the navigation, each computations node has some work to do. Moreover, sending the simplified data to the visualization nodes is not done on a per-tile basis. Instead, each computations node data are gathered together before being sent, thus optimizing the network bandwidth use.

4. Minimizing communication overhead

Although they have evolved, networks still remain a bottleneck for data transfer-intensive applications on PC clusters. Therefore, sending all data at each frame from LOD nodes to visualization nodes would bound the model scalability. Instead, we propose a scheme which relies on culling and frame-to-frame coherence to considerably reduce the data transfer overhead by avoiding sending non visible and non modified data.

4.1. Hierarchical frustum culling

Large terrain visualization highly benefits from frustum culling since usually only a small part of the whole model is visible at a time. Systems for interactive rendering of large datasets often comprise many optimization techniques. Hierarchical view frustum culling is one of the most interesting ones. Combining it with LOD allows to maintain a constant frame rate. LOD is more effective for large views of the 3D object whereas frustum culling reaches its full potential at close ranges. In addition, in our scheme, all distribution features take advantage of this combination: faster LOD com-

putations with early geometry pruning and lighter communication, both leading to an effective rendering. In order to perform the culling, a hierarchy of bounding boxes is generated at the preprocessing time. As for d_2 values, a bounding box is computed for each possible quadtree block. At runtime, the view frustum culling is carried out in two stages : *coarse-grained* by not processing every tile fully outside the viewing frustum, and then *fine-grained* inside the LOD process when traversing each tile quadtree. Storing a complete hierarchy of bounding boxes over each tile quadtree would represent up to 66% of the original data. Actually a reasonable tradeoff is to limit the hierarchy to one or two level(s) below its maximum depth, the memory overhead is henceforth of 16% or 4% only whereas the culling efficiency is still maintained.

4.2. Exploiting Temporal Coherence

Despite of LOD which continuously modifies the data visibility state, we observe that, the frame-to-frame coherence property still holds inside tiles. So, between two frames, large parts of the rendered data do not change even with interaction. On the basis of this observation, we implemented a mechanism allowing modified data detection. It should be noted that this stage takes place after the culling step. Therefore, as it is applied to already reduced data, it does not induce a significant CPU overload. Furthermore, the introduced memory overhead is compensated by a substantial reduction of communication overhead. This mechanism is very similar to Unix commands *diff* and *patch*. We recall that after LOD algorithm is performed, data are sent to visualization nodes as triangle fans. For each tile, we record a copy of the list of fans being rendered which is the result of the previous LOD computations. At the end of current computations, a new list of fans is generated. A *diff* operation occurs between this list and the saved one. The delta is then sent to the visualization nodes which perform a *patch* operation between the incoming data and the previously rendered data. This operation results in updating the rendered data with respect to the current view parameters.

5. Rendering Quality

Thanks to advances in graphics technologies, graphics users become more demanding in interactivity and high visual quality even for huge models. LOD-specific visual issues are not tolerated anymore. Indeed, LOD obviously produces lower detail rendering in simplified regions. Moreover, as the view point moves, detail is suddenly added or removed. This leads to visual artifacts known as vertex popping. There are several ways to solve these issues. Terrain texturing can be used to maintain an acceptable detail even in simplified areas. Vertex popping is usually solved by using colors blending or by implementing a geometrical morphing technique (geomorphing) [Hop97]. However, geomorphing is more widely used than blending. Instead of going from one level

of detail to another in one step, geomorphing technique consists to smoothly morph vertex positions between two consecutive levels of detail. This smooth transition is gained at the cost of an additional processing. So, it naturally benefits from our parallel framework as the induced CPU workload is distributed on all computations nodes.

5.1. Geomorphing

Since its introduction by Fergusson et al. [FEK01], many LOD algorithms have incorporated the morphing technique. Different approaches are used for its implementation. Distance-based approaches use transition zones to blend the geometry between levels of detail [COL96, Paj98]. Time-based strategies use a fixed time interval as interpolating parameter [Hop97, DWS*97]. These approaches may induce important computational overhead because finding the appropriate interpolating parameter is not a trivial task [Hop97]. Another approach consists in using the error metric used in the algorithm to parameterize the morphing process [RHSS98, CE01]. What is common to all approaches is that they all may introduce cracks because of the need to early know the interpolation final positions before the subdivision that creates them occurs.

In [RHSS98], the morphing parameter (called the blending factor b) is computed using the decision variable f (the metric error). However, instead of computing the blending factor as specified in that paper, we use the formula $b = \frac{(1-f)}{(0.5+0.5*f_{cmin}-f)}$ [ter] (f_{cmin} being the minimum of the child blocks f values) that takes into account children states. This is essential in order to prevent a block having a child while still morphing. Moreover it allows more linearity and more smoothness to the morphing process.

Our parallel framework perfectly fits to this technique (see figure 4). However, due to the border problem, the cracks occurrence is even more important in our model than the sequential ones. But as the blending factor depends on f and that the global consistence issue regarding this parameter is solved with the border problem, our implementation is free of geomorphing cracks.




Figure 4: New vertices at each level of refinement. Only these vertices have to be geomorphed, while corner ones stay at the same 3D position.

5.2. Texturing

Texture data represent another important category of terrain data. In many applications, terrain data are associated with another layer of data such as a large satellite image, generally in the form of a texture handled by the graphics card [DBH00, BDH00].

To fit the texturing to our tiling framework, we also adopt a tiled texturing approach [Ulr02]. At preprocessing stage, we compute a texture tile for each terrain tile using lightmapping and multi-texturing. These additional tiles are managed by the visualization nodes at rendering time.

Using lightmaps leads us to a static lighting based on the original terrain model. We found this technique more suitable to our framework than dynamic lighting that could require more graphical resources and lead to disturbing effects as we do not provide yet morphing mechanism for lighting attributes. In contrast, our approach avoids this drawback and it requires only one or two bytes per texel which can be compressed using advanced features of modern graphics cards. An alternative way would be to distribute lighting information on computations nodes, packed with height data. But lighting precision will follow the LOD simplification scheme, which leads to unexpected appearance in most simplified areas.

Texturing may also be used to manage missing values by exploiting alpha texturing. Missing data are present in many areas such as GIS. In this framework, we use them to fulfill the tile size constraint by setting to 'missing' all values beyond the real boundary of the tile.

Finally, large scenery textures may be difficult to handle as we do not provide distribution mechanism for texture data. However, since our framework frees visualization nodes from unnecessary geometry storing and processing, the whole graphical resources are available for addressing a maximal video memory, using AGP memory at full extent or even running out-of-core techniques such as clip-mapping [TMJ98].

6. Performance

We perform our tests on a linux cluster of 8 PCs, each equipped with 1 Pentium 4 processor, 1 GB of RAM and 1 NVIDIA GeForce FX 5900 graphics card with 256 MB of video memory. The nodes are connected by a Gigabit Ethernet network over TCP/IP. The same screen resolution of 1024x768 is used in all configurations. We use the publicly available USGS Washington state dataset [was] for all tests. The flythrough over the terrain is fixed in order to keep the interaction almost identical in all cases. As our framework is a tile-based system, we can load a variable number of tiles in order to change our dataset size. We choose a tile size of 257x257. Changing the granularity of the tiles affects load balancing and the minimal global resolution. This

resolution is set to a value that always ensure a good approximation of the model. The parallel environment used is the MPI (see www-unix.mcs.anl.gov/mmpi/mpich/) communication library and we use Net Juggler [AGL*02], as parallel rendering middleware. LOD computations and the rendering are synchronized at frame level by Net Juggler. In the remainder of the section, we note $m - n$ a setup composed of m visualization nodes and n computations nodes.

Using the described configuration, an extended version of the whole USGS dataset representing 3.2 Gigabytes of data was successfully rendered in a 1 - 7 setup at an average rate of 28 frames per seconde (FPS). Each of the 7 computations nodes was loaded with 3600 tiles of 256x256 points, which equals to a complete map of 1.65 billions 16 bits height values. Then, we perform some tests with small datasets. The tests consist in rendering a fixed-size model using different setups. This model is made up of 3600 tiles. In the first setup, we use only one computations node, so that it processes the whole dataset. The second setup uses 2 computations nodes, each one dealing with 1800 tiles. The others setups follow the same pattern. Results being drawn in the dotted curve of the figure 5, we can observe that after a small speed-up of about 1.5, the speed-up stays almost constant. This is not surprisingly for a parallel approach. When data are too small, the parallelization overhead is not compensated. However, when the data size exceeds some threshold, the framework performs at its full efficiency.

The purpose of our work is to increase the size of datasets which can be rendered using level of detail and commodity clusters. So, our first test consists in checking the algorithm data scalability. Our experiments showed that when dealing with geometry only, one needs 3600 tiles to saturate one node's main memory (1 GB). So, each computations node participating to the test is loaded with $L_1 = 3600$ tiles. We then use the frame rate for the 1 - 1 setup as a reference and give it the ratio 1. For the 1 - 2 setup, the global terrain size is $L_2 = 2 \times L_1 = 7200$ tiles. Thus, as the data size is two times that of the 1 - 1 setup, the frame rate ratio is two times the observed frame rate divided by the frame rate of the reference setup. More generally, for a $m - n$ setup, the global terrain size is $L_n = n \times L_1$ and the speed-up is obtained by the formula: $frame\ rate\ speed-up \times \frac{L_n}{L_1}$. The results are summarized in the figure 5. We observe a super-linear speed-up very close to the optimal linear speed-up. It clearly emphasizes the data scalability of our approach. The super-linearity is due to a reduction in the LOD computations time. Indeed, thanks to the view-dependent nature of the LOD algorithm, whatever the number of computations nodes and their load, the intensive work is always located around the nearest areas to the view point. Moreover, we note that we do not necessarily enlarge these areas by adding more tiles, so as our algorithm is parallel and load-balanced, the LOD computations time does not grow in the same proportion as the dataset size.




Figure 5: Data scalability test.

Tests on the data compression are very conclusive while not activating geomorphing. 80% to 90% of the data remain unchanged on the visualization nodes. However, when activating geomorphing this ratio drops to 50-60% which is not surprisingly as geomorphing introduces more changing vertices (see figure 6). Using our framework, we distribute these additional computations on the cluster nodes. As shown in figure 6, the more they are, the lighter the cost due to geomorphing.




Figure 6: Implementing geomorphing introduces important computation overhead. Observe the distance between the 2 top curves. These curves represent the global computation time when geomorphing is ON (curve on the bottom) and OFF (curve on the top).

As stated in section 5, texturing improves the visual quality. But as the goal of these tests is to clearly establish the pure scalability of our parallel framework in terms of dataset size, we do not activate it (see figure 7). But texturing induces no substantial overhead on the rendering, so activating it do not alter the given results.




Figure 7: Our parallel framework running on USGS Washington State dataset.

7. Conclusion

We presented a parallel framework for interactive smooth rendering of massive terrains. The framework based on the model presented in [GLMM04], gives several improvements on this previous work. The main contributions concern addressing unsolved issues in [GLMM04] by solving the border problem and reducing load imbalance, adding more optimizations to further reduce the communication overhead by providing a hierarchical view frustum culling and a data compression mechanisms. Moreover, geomorphing and texturing features are given to improve the rendering quality. As shown by our benchmarks, this allows our framework to support smooth interactive rendering of very huge terrains.

However, texture data bound our parallel scheme because they are centralized on the visualization nodes. A solution will be to distribute those data among computation nodes which process them and then transfer the resulting color values to the rendering nodes. Furthermore, terrain data may be described as 'triangle soup', our framework is then not usable. In order to extend its applicability to other types of dataset such as triangulated irregular networks (TIN) or time-varying data, we aim at using scattered data approximation techniques [LWS97] to adapt such models to our framework. In addition, to speed up the local computations, the high stream processing capabilities of recent graphics processors can be exploited.

Acknowledgments

We would like to acknowledge the BRGM (French Geological survey) and Région Centre for supporting part of this work. We also would like to acknowledge Sylvain Jubertie for setting up the environment for our tests.

References

- [AGL*02] ALLARD J., GOURANTON V., LECOINTRE L., MELIN E., RAFFIN B.: Net juggler: Running vr juggler with multiple displays on a commodity component cluster. *IEEE VR'02* (2002). 6
- [ARZ03] ALLARD J., RAFFIN B., ZARA F.: Coupling parallel simulation and multi-display visualization on a pc cluster. In *Euro-par 2003* (Klagenfurt, Austria, August 2003). 2
- [BDH00] BAUMANN K., DOLLNER J., HINRICHS K.: Integrated multiresolution geometry and texture models for terrain visualization. *Proceedings joint EuroGraphics-IEEE TCVG 2000* (May 2000). 3, 5
- [BW00] BRODSKY D., WATSON B.: Model simplification through refinement. *Graphics Interface'00* (2000). 2
- [CE01] CLINE D., EGBERT P.: Terrain decimation through quadtree morphing. *IEEE Transaction on Visualization and Computer Graphics 2001* 7, 1 (January 2001), 62–69. 5
- [CGG*03] CIGNONI P., GANOVELLI F., GOBETTI E., MARTON F., PONCHIO F., SCOPIGNO R.: Bdam: Batched dynamic adaptive meshes for high performance terrain visualization. *Proceedings EG2003* (September 2003), 505–514. 3
- [COL96] COHEN-OR D., LEVANONI Y.: Temporal continuity of levels of detail in delaunay triangulated terrain. *IEEE Visualization '96* (October 1996), 37–42. 5
- [CRMS02] CIGNONI P., ROCCHINI C., MONTANI C., SCOPIGNO R.: External memory management and simplification of huge meshes. *IEEE Transaction on Visualization and Computer Graphics* (2002). 3
- [DBH00] DOLLNER J., BAUMANN K., HINRICHS K.: Texturing techniques for terrain visualization. *Proceedings IEEE Visualization'00* (2000), 227–234. 5
- [DLR00] DEHNE F., LANGIS C., ROTH G.: Mesh simplification in parallel. *ICA3PP'00* (2000), 281–290. 2
- [DWS*97] DUCHAINEAU M., WOLINSKY M., SIGETI D., MILLE M., ALDRICH C., MINEEV-WEINSTEIN M. B.: Roaming terrain: Real-time optimally adapting meshes. *IEEE Visualization* (1997), 81–88. 1, 5
- [ESC00] EL-SANA J., CHIANG Y.-J.: External memory view-dependent simplification. *Computer Graphics Forum* 3, 19 (August 2000), 139–150. 3
- [ESV99] EL-SANA J., VARSHNEY A.: Parallel processing for view-dependent polygonal virtual environments. *Proceedings SIGGRAPH'99* (1999). 2
- [FEK01] FERGUSON R., ECONOMY R., KELLY A. AND RAMOS P.: Continuous terrain level of detail for visual simulation. *ACM Symposium on Interactive 3D Graphics* (March 2001), 111–120. 5

- [GJR03] GAUGNE R., JUBERTIE S., ROBERT S.: Distributed multigrid algorithms for interactive scientific simulations on clusters. *ICAT 2003, Japan* (2003). 2
- [GLMM04] GOURANTON V., LIMET S., MADOUGOU S., MELIN E.: a scalable cluster-based parallel simplification framework for height fields. *EuroGraphics/ACM SIGGRAPH, Proceedings Parallel Graphics and Visualization'04* (June 2004), 59–65. 2, 3, 7
- [HDJ04] HWA L., DUCHAINEAU M., JOY K.: Adaptive 4-8 texture hierarchies. *IEEE Visualization '04* (October 2004), 219–226. 1
- [Hop96] HOPPE H.: Progressive meshes. In *proceedings SIGGRAPH'96* (1996), 99–108. 1
- [Hop97] HOPPE H.: View-dependent refinement of progressive meshes. *Computer Graphics (In proceedings SIGGRAPH'97)* (August 1997), 189–198. 4, 5
- [Hop98] HOPPE H.: Smooth view-dependent level-of-detail control and its application to terrain rendering. *IEEE Visualization '98 31* (October 1998), 35–42. 3
- [LKR*96] LINDSTROM P., KOLLER D., RIBARSKY W., HODGES L. F., FAUST N., TURNER G.: Real-time, continuous level of detail rendering of height fields. *Computer Graphics, Proceedings SIGGRAPH'96* (1996), 109–118. 1
- [LWS97] LEE S., WOLBERG G., SHIN Y.: Scattered data interpolation with multilevel b-splines. *IEEE Transaction on Visualization and Computer Graphics* 3, 3 (July-September 1997). 7
- [MBDM97] MONTRYM J., BAUM D., DIGNAM D., MIGDAL C.: InfiniteReality: A Real-Time Graphics System. In *Computer Graphics (SIGGRAPH 97)* (August 1997), ACM Press, pp. 293–303. 1
- [Paj98] PAJAROLA R.: Large scale terrain visualization using the restricted quadtree triangulation. *IEEE Visualization (Proc. IEEE Visualization '98)* (1998), 19–26. 5
- [RHSS98] ROTTGER S., HEIDRICH W., SLUSSALLEK P., SEIDEL H.-P.: Real-time generation of continuous levels of detail for height fields. *Proceedings of the 6th International Conference in Central Europe on Computer Graphics and Visualization* (February 1998), 315–322. 2, 3, 5
- [ter] <http://home.planet.nl/~monstrous/>. 5
- [TMJ98] TANNER C., MIGDAL C., JONES M.: The clipmap: A virtual mipmap. *SIGGRAPH'98 proceedings* (1998), 151–158. 5
- [top] <http://www.top500.org>. 1
- [Ulr02] ULRICH T.: Rendering massive terrains using chunked level of detail control. *SIGGRAPH Course Notes* (2002). 3, 5
- [was] <http://rocky.ess.washington.edu/data/raster/tenmeter/onebytwo10/index.html>. 1, 5
- [YSGM04] YOON S., SALOMON B., GAYLE R., MANOCHA D.: Quick-vdr: Interactive view-dependent rendering of massive models. *IEEE Visualization 2004* (October 2004), 131–138. 3