

距離画像統合の並列化

佐川立昌 西野 恒 Mark D. Wheeler 池内克史
東京大学 東京大学 Cyra Technologies, Inc. 東京大学

概要

本論文では距離画像を用いて物体の表面形状モデルを生成する体積的統合アルゴリズムについて説明する。既に我々が提案したアルゴリズムでは、体積的陰関数表現を用いて複数の距離画像を統合し、マーチングキューブスアルゴリズムを用いて表面形状のメッシュモデルを生成した。ここで、その過程における符号付距離を計算において、計算とメモリ使用の効率化をはかる手法とPCクラスタを用いた並列計算手法を提案する。本手法により主記憶に配置されるデータ量を削減し、最近傍点探索を効率的に行なう。また計算の並列度を高めることができ、計算時間を小さくすることができる。

並列計算手法では、以下のCPUとメモリ使用の効率性において相補的な2つのアルゴリズムを提案する。1. 距離データの分散配置。2. 部分octreeの並列探索。これらをシステムの仕様にしたがって調節して用いることにより、PCクラスタを用いて効率的なモデル生成が可能になる。

Parallel Processing of Range Data Merging

Ryusuke Sagawa Ko Nishino Mark D. Wheeler Katsushi Ikeuchi
Univ. of Tokyo Univ. of Tokyo Cyra Technologies, Inc. Univ. of Tokyo

Abstract

In this paper, we describe a volumetric view-merging algorithm that generates a consensus surface of an object from its range images. Our original method merges a set of range images into a volumetric implicit-surface representation, which is converted to a surface mesh using a variant of the marching-cubes algorithm. We propose a method which increases the computation and memory efficiency of computing signed distances and the method of parallel computing using a PC cluster. Since our method permits a reduction in the data allocated in memory, the closest point is searched efficiently; this allows us to increase the number of the parallel traversals and to reduce the computation time.

We describe the following two algorithms which are complementary in terms of the efficiency of CPUs and memory usage: 1. Distributed allocation of range data. 2. Parallel Traversal of partial octree. By adjusting them according to the system specification, we can build the model efficiently by a PC cluster.

1 Introduction

We have been developing techniques to automatically create virtual reality models through observation of real objects; we refer to these techniques as modeling-from-reality (MFR). In order to explore unforeseen technical difficulties and to further extend our MFR techniques by solving these difficulties, we have begun a project to model Japanese cultural heritage objects through these MFR techniques[1].

Some of Japanese cultural heritage objects are large, but their shapes may be intricate. Thus, the models of these objects' shapes must contain huge amounts of data. In our previous experiments in modeling small, indoor objects, we did not have to consider the computation and memory requirements to build these models. However, building a model of a huge amount of data necessitates our taking these requirements into account. In this paper, we describe our proposed method for modeling the shape of huge, possibly intricate objects.

After scanning the shape of an object by using a range sensor and then aligning all range images into the same coordinate system, our original method[2] converts a set of range images into a volumetric implicit-surface representation. It then obtains a surface mesh using a variant of the marching-cubes algorithm[3]. Unlike previous techniques[4, 5, 6] based on implicit-surface representations, our method estimates the signed distance to the object surface by finding a consensus of locally coherent observations of the surface.

The most costly part of the computation of our method is finding the consensus surface to compute the signed distance. To increase the computation and memory efficiency, we propose a method which reduces the amount of data to be searched, around which point the signed distance is computed.

We utilize octrees to represent volumetric implicit surfaces for effectively reducing the computation and memory requirements of the volumetric representation without sacrificing the accuracy of the resulting surface.

To further ease this size problem, we have developed parallel software that runs on a PC cluster to handle the huge amount of data. The parallel soft-




Figure 1: Zero-crossing interpolation from the grid sampling of an implicit surface

ware consists of the following two components: 1. Distributed allocation of range data. 2. Parallel Traversal of partial octree.

In the following sections, Section 2 describes our original merging algorithm. Section 3 explains the method to increase the computation and memory efficiency. In Section 4, the parallel merging algorithm is shown. Finally, the merged result is shown in Section 5.

2 Data Merging

2.1 Volumetric Modeling and Marching Cubes

Recently, the marching-cubes algorithm[3] has propelled volumetric modeling beyond the confines of "blocky" occupancy grids. Instead of storing a binary value in each voxel to indicate whether the voxel is empty or full, the marching-cubes algorithm requires that the data in the volume grid are samples of an implicit surface. In each voxel, we store the signed distance, $f(\mathbf{x})$, from the center point of the voxel, \mathbf{x} , to the closest point on the object's surface. The sign indicates whether the point is outside, $f(\mathbf{x}) > 0$, or inside, $f(\mathbf{x}) < 0$, the object's surface, while $f(\mathbf{x}) = 0$ indicates that \mathbf{x} lies on the surface of the object(See Figure 1).

The marching-cubes algorithm constructs a surface




Figure 2: Marching Cubes: An implicit surface is approximated of by triangles. \circ : voxels of outside surface. \bullet : voxels of inside surface.

mesh by “marching” around the cubes while following the zero crossings of the implicit surface $f(\mathbf{x}) = 0$. The resulting surfaces are relatively smooth and their accuracy can be greater than the resolution of the volume grid due to sub-voxel interpolation (See Figure 2).

Now we focus on a more easily solved problem: How do we compute $f(\mathbf{x})$? The real problem underlying our simple question is that we do not have a surface; instead, we have many surfaces. Some elements of those surfaces do not belong to the object of interest but rather are artifacts of the image acquisition process or background surfaces. In the next subsection, we present an algorithm that answers the question and does so reliably in spite of the presence of noisy and extraneous surfaces in our data.

2.2 Consensus Surface Algorithm

In this section, we present our answer to the question of how to compute the signed distance function $f(\mathbf{x})$ for arbitrary points \mathbf{x} when given N triangulated surface patches from various views of the object surface. We call our algorithm the consensus-surface algorithm.

We can break down the computation of $f(\mathbf{x})$ into two steps:

- Compute the magnitude: compute the distance, $|f(\mathbf{x})|$, to the nearest object surface from \mathbf{x}
- Compute the sign: determine whether the point is inside or outside of the object




Figure 3: Naive algorithm: An example of inferring the incorrect sign of a voxel’s value, $f(\mathbf{x})$, due to a single noisy triangle.

The previous naive algorithm finds the nearest triangle from all views and uses the distance to that triangle as the magnitude $|f(\mathbf{x})|$. If the normal of the closest surface point is directed toward \mathbf{x} , then \mathbf{x} must be outside the object surface. In Figure 3, the point chosen as the closest point from \mathbf{x} does not belong to the real surface. Thus, the algorithm incorrectly thinks \mathbf{x} is inside the surface based on the normal information from the closest point.

Our solution to these problems is to estimate the surface locally by averaging the observations of the same surface. The trick is to specify a method for identifying and collecting all observations of the same surface.

Nearby observations are compared using their location and surface normal. If the location and normal are within a predefined error tolerance (determined empirically), we can consider them to be observations of the same surface. Given a point on one of the observed triangle surfaces, we can search that region of 3D space for other nearby observations from other views which are potentially observations of the same surface. These searches are efficiently implemented using k-d trees[7].

The consensus-surface algorithm examines the closest point in each image’s triangle set. If there are sufficient surfaces of other triangle sets which are regarded as the same surfaces of the each closest point, the closest point is a consensus surface. The algo-




Figure 4: Consensus surface algorithm: The signed distance is chosen from circled consensus surfaces.

rithm which determines whether two surface observations are sufficiently close in terms of location and normal direction is as follows:

$$\text{SameSurface}(\langle \mathbf{p}_0, \mathbf{n}_0 \rangle, \langle \mathbf{p}_1, \mathbf{n}_1 \rangle) = \begin{cases} \text{True} & (\|\mathbf{p}_0 - \mathbf{p}_1\| \leq \delta_d) \wedge (\mathbf{n}_0 \cdot \mathbf{n}_1 \geq \cos \theta_n) \\ \text{False} & \text{otherwise} \end{cases} \quad (1)$$

where δ_d is the maximum allowed distance and θ_n is the maximum allowed difference in normal directions.

For example, consensus surfaces are circled in Figure 4. The algorithm chooses the closest one of them as the signed distance. In this case, it is correctly determined that \mathbf{x} is the outside surface and \mathbf{x}' is the inside surface.

2.3 Adaptive Resolution by Octree Representation

Volumetric modeling involves a tradeoff between accuracy and efficiency. The octree representation[8] balances these problems while keeping the algorithm implementation simple. Instead of iterating over all elements of the voxel grid, we can apply a recursive algorithm on an octree that samples the volume more finely only when near the surface of the object (See Figure 5).

To interpolate the zero crossings properly, we will need the implicit distance for the voxel containing the surface (the zero crossing) and all voxels neighboring this voxel; these voxels must all be represented at the




Figure 5: The adaptive resolution is high around the surface and low elsewhere

finest level of precision. This constraint means that, if we have a surface at one corner of an octant, the longest possible distance to the center of a neighboring octant is one and one-half diagonals of the voxel cube, which is a distance of $\frac{3\sqrt{3}}{2}$ cube units.

Given the current octant, we can compute the signed distance. If the magnitude of the signed distance, $|f(\mathbf{x})|$, is larger than $\frac{3\sqrt{3}}{2}$ of the octant width, then it is not possible for the surface to lie in the current or neighboring octant. If the surface is not in the current or neighboring octant, we do not care to further subdivide the current octant.

3 Increase the computation and memory efficiency

If the size of mesh data to be merged is huge, it is difficult to allocate all of those to memory. Also, the computation time of the signed distance cannot be ignored. We propose the following method to increase the computation and memory efficiency by reducing the data allocated in the memory.

When the algorithm traverses a part of the octree, the data searched for finding the closest surface is only the local area around the voxel. The data of the other area are never used for computing signed distances while traversing the sub-octree. Moreover, a closest surface is effectively searched using a k-d tree.




Figure 6: Load only the mesh data within the dotted rectangle into memory

However, it is inefficient when the k-d tree contains unnecessary data.

As described in Section 2.3, an octant is subdivided when its signed distance is less than $\frac{3\sqrt{3}}{2}$ cube units. Thus, the data farther than $\frac{3\sqrt{3}}{2}$ cube units is not necessary for finding the closest point of the voxel.

To load the necessary data into memory, we must read all of the data files. Since the overhead of reading files for the every finest octant is too costly, we read the data files for an ancestor octant. Where the width of an ancestor octant is W_0 and the width of the finest octant is W , the area of the mesh data to be loaded is inside the rectangle of a dotted line in Figure 6.

4 Parallel Computing of Signed Distances

In this section, we describe the algorithm for parallel computing of signed distances. There are two motivations for parallel computing signed distances. We now propose the parallel computing method for each motivation:

1. Handle range data of huge size: We distribute the allocation of range data to multiple PCs.
2. Fast merging: We divide the octree to sub-octrees and assign traversal of a sub-octree to each CPU.




Figure 7: Parallel computation of signed-distances

4.1 Distributed Allocation of Range Data

Calculating a signed distance from a point requires consideration of all range data with respect to this point. When the number of the measurement increases, more data should be considered. It becomes difficult to allocate all the range data in a single processor.

We distribute that range data to multiple PCs and compute signed distances in a parallel manner. For example, in Figure 7, Data 1,2,3 are allocated to PC 1,2,3, respectively. Signed distances from the point, x , to Data 1 are computed by PC1. In the same manner, signed distances to Data 2 are computed by PC2, and so on. Since finding the closest point of a mesh data is independent of the others, we can compute signed distances in a parallel manner.

However, the computation times are different among CPUs; After finding the closest points of all data, we have to choose the smallest magnitude of the signed distances. To synchronize, the CPUs have to wait until the remaining CPUs finish computing the signed distances.

4.2 Parallel Traversal of an Octree

Dividing an octree into partial trees enables us to traverse the partial trees. We assign the partial space of an octree to each CPU and traverse partial trees in a parallel manner (See Figure 8). Since the traver-




Figure 8: Assignment of partial space of the octree to each CPU and parallel traversal of partial trees

sals of partial trees are independent of one another, a traversal does not have to synchronize with others, and the computation time can be reduced according to the number of CPUs.

By the method described in Section 3, the area of range data which each process owns is only inside the voxel and its peripheral area. Thus, each process owns only the range data of the local area which it takes charge of in a traversal of a partial octree.

However, each machine must cache range data files in memory to read them efficiently and repeatedly. Since a PC cluster cannot share data as a shared-memory machine can, range data files have to be allocated redundantly; therefore, memory efficiency grows worse as this parallel traversal method is used.

4.3 Combination of Parallel Methods

The above two methods are complementary in terms of the efficiency of CPUs and memory usage. In practice, they should be adjusted according to the system specification by combining those two methods with an appropriate condition. Two methods can be combined by allocating range data distributed in each parallel traversal.

The maximum number of traversals is determined by the system memory size. Thus, the combination strategy maximizes the number of traversals to deal with the memory. If the system has more CPUs than the parallel traversals, each traversal uses multiple CPUs by the method of distributed allocation.

5 Merged Result

We have implemented these algorithms, and constructed one integrated digital Great Buddha. For this project, we have built a PC cluster that consists of eight PCs of dual PentiumIII 800MHz processors with 1GB memory for each PC. The machines are connected by 100BASE-TX Ethernet. Figure 9 shows the obtained geometric model of the Great Buddha; the model contains 3 million points and 5.5 million triangles.

We tested the merging program by changing parameters of the number of traversals and machines of each traversal. Merged data consists of 12 files; of those files, the average contains about 300 thousand points and 600 thousand triangles. The total size is about 150M bytes. The result is shown in Table 1. The computation time is almost inversely proportional to the number of parallel traversals; and the required memory of each machine is inversely proportional to the number of machines in each traversals.

According to the combination strategy, the signed distances are computed by 16 parallel traversals that are allocated to each PC to minimize the computation time. If each PC has only 256MB memory, the signed distances are should be computed by 8 parallel traversals that are allocated and distributed to 2 PCs.

Without reducing the data allocated in the memory, the maximum number of the traversals is four because of the system memory size. It takes 59 hours to build the model where it is computed by 4 traversals that are allocated and distributed to 4 PCs. It has been proven that the method of reducing the data allocated in the memory increases the computation and memory efficiency.

6 Conclusion

In this paper, we proposed a method which increases the computation and memory efficiency of computing signed distances, along with a method for parallel computing using a PC cluster. First, since we reduce the data allocated in the memory, the closest point is searched efficiently. Thus, we can increase




Figure 9: Merging result of Kamakura Great Buddha

Table 1: Results of different parameters of the number of traversals and machines of each traversal.

Number of traversals	Number of machines in each traversal	Average required memory of each machine	Computation Time
1	1	200MB	468 min.
8	1	200MB	58 min.
1	8	20MB	256 min.
8	2	200MB	44 min.
16	1	250MB	23 min.

the number of the parallel traversals and reduce the computation time.

In addition, we described two algorithms which are complementary in terms of the efficiency of CPUs and memory usage. By adjusting them according to the system specifications, we can build the model efficiently by using a PC cluster.

Now we can build models of huge size. In the future, we plan to scan more Japanese cultural heritage objects and build fine models with photometric attributes.

References

- [1] Daisuke Miyazaki, Takeshi Ooishi, Taku Nishikawa, Ryusuke Sagawa, Ko Nishino, Takashi Tomomatsu, Yutaka Takase, and Katsushi Ikeuchi. The great buddha project: Modelling cultural heritage through observation. In *Proceedings of 6th International Conference on Virtual Systems and MultiMedia*, pp. 138–145, Gifu, 2000.
- [2] M.D. Wheeler, Y. Sato, and K. Ikeuchi. Consensus surfaces for modeling 3d objects from multiple range images. In *Proc. International Conference on Computer Vision*, January 1998.
- [3] W. Lorensen and H. Cline. Marching cubes: a high resolution 3d surface construction algorithm. In *Proc. SIGGRAPH'87*, pp. 163–170. ACM, 1987.
- [4] H. Hoppe, T. DeRose, T. Duchamp, J.A. McDonald, and W. Stuetzle. Surface reconstruction from

- unorganized points. In *Proc. SIGGRAPH'92*, pp. 71-78. ACM, 1992.
- [5] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proc. SIGGRAPH'96*, pp. 303-312. ACM, 1996.
- [6] A. Hilton, A.J. Stoddart, J. Illingworth, and T. Windeatt. Reliable surface reconstruction from multiple range images. In *Proceedings of European Conference on Computer Vision*, pp. 117-126, Springer-Verlag, 1996.
- [7] Jerome H. Friedman, Jon Bentley, and Raphael Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software*, Vol. 3, No. 3, pp. 209-226, 1977.
- [8] D. J. R. Meagher. *The octree encoding method for efficient solid modeling*. PhD thesis, Rensselaer Polytechnic Institute, 1980.