

EFFICIENT GEOMETRIC ALGORITHMS FOR CAD

Wm. Randolph Franklin, Narayanaswami Chandrasekhar, and Mohan Kankanhalli

Electrical, Computer, and Systems Engineering Dept., and Computer Science Department, 6026 J.E.C., Rensselaer Polytechnic Institute, Troy, NY, 12180, USA. (518) 276-6077, *Internet*: WRF@ECSE.RPI.EDU, *Bitnet*: WRFRANKL@RPITSMTS, *Fax*: (518) 276-6003, *Telex*: 6716050 RPI TROU.

Varol Akman

Dept. of Computer Engineering and Information Sciences, Bilkent University, PO Box 8, 06572 Maltepe, Ankara, TURKEY. *Bitnet*: AKMAN@TRBILUN.

Peter YF Wu

IBM Thomas J. Watson Research Center, PO Box 704, Yorktown Heights, NY, 10598, USA. *Internet*: PWU@IBM.COM

Efficient parallelizable algorithms that process large geometric data sets with millions of edges are possible using techniques such as uniform grids and local topology data structures. Possible operations include edge intersection detection, and calculation of mass properties of boolean combinations of objects. These techniques are also useful in the low level compute-bound steps in other geometric areas. The uniform grid is a regular grid overlaying the data whose resolution depends on the number and length of the data's edges. Hierarchical data structures such as quadtrees are not necessary for these operations. The grid resolution can be varied a factor of three either way from the optimum without increasing execution time over 65%. The local data structures represent the object as a set of fractional neighborhoods of the vertices. The edges and faces are stored only implicitly. Properties such as total edge length, face area, and volume can be calculated in one pass through the data. Both of these techniques can be parallelized on SIMD machines as vector reduction operations. High speed on large databases is a major advantage of these techniques, and has been demonstrated experimentally. For example two million edges were tested to find all seven million intersections in three minutes on a Sun 4/280.

1. INTRODUCTION

Geometric problems in separate fields, such as interference testing in robotics, VLSI circuit extraction, hidden surface removal in computer graphics, and map overlay in cartography have in common low level operations, such as testing a point for inclusion in a polygon, and finding which of a large set of small edges intersect. These operations then integrate upward to applications such as boolean operations on thousands of polyhedra, visible surface calculations for large scenes, determining mass properties of the union of many rectangles without determining the union polygon itself, and determining the area of the intersection of several polygons each

defined as the union of many rectangles.

This paper presents algorithms to solve some of these problems using the techniques of uniform grids, local topological data structures, and multiple precision rational fractions in Prolog. An efficient algorithm for finding the intersections among large numbers of edges is extended to finding the area and perimeter of the union of many rectangles. These use the uniform grid and the local data structures. These mass properties can be found more quickly than finding the resulting polygon itself. A map overlay algorithm using rationals in Prolog is also presented. Applications to constructive solid geometry tree evaluation are mentioned.

High speed on large databases is a major advantage of these techniques, and has been demonstrated experimentally. We have processed a complete integrated circuit design of 1,819,064 edges to find all 6,941,110 intersections in 178 seconds on a Sun 4/280. When implemented on a 16 processor Sequent, which has separate processors with a common bus and memory, the program ran 10 times as fast with 15 processors as with one when processing the complete USGS Chikamauga DLG (Digital Line Graph) took only 37 CPU seconds. Finding mass properties of the union of many rectangles was also implemented. It processed a union of 100,000 rectangles in 199 seconds.

This paper will first present uniform grids, compare them to other possibilities, and describe their use to intersect edges. Then it will describe local topological data structures. Next it will give applications combining these two ideas, such as finding the area of the union of many polygons, and overlaying two maps in cartography. Finally it will discuss the advantages of these concepts.

2. UNIFORM GRIDS

We use *uniform grids* to determine spatial adjacency and intersection [1,2,10,14]. Here a uniform, regular, grid, with resolution determined by the number and average length of the edges, is laid over the data. For each object, the cells that it passes through are noted. Then the data structure is inverted and the contents of each cell are processed sequentially. This algorithm is useful for operations such as detecting intersections among large numbers of small edges, boolean combinations of polyhedra, and hidden surface calculations of complex scenes.

2.1. Comparison to Quadrees

The obvious, but wrong, objection to uniform grids is that they cannot handle irregular scenes, and that hierarchical methods such as quadrees must be used to give finer cells where the data is denser. There are three answers to this.

1. *Theoretical:* If the data are distributed independently and identically then some cells will have more objects than others, with n_i , the number of objects in cell #i following a Poisson distribution, $p(n_i = k) = \lambda^k e^{-\lambda} / k!$, where λ is the average number of objects per cell. Then the time to process cell #i by comparing all objects in it against each other will be n_i^2 , and its mean will be λ^2 . That is, for the Poisson distribution, the mean of the square is the square of the mean. In this case, the expected time, $\bar{T} = \Theta(N+K)$, where N is the number of input objects, and K the number of output intersections. More details of the analysis are given in [1,10].

If the input objects are distributed unevenly, but the unevenness is "bounded", then the above result is still true. It is sufficient that the densest cell be a constant times the average density as the data sets grow to infinity. Similar restrictions are imposed on data in other fields, such as Lipschitz conditions in partial differential equations, and bounded variation in analysis.

Real data can be worse than the randomness assumptions indicate. For example, the probability of two random edges having the same endpoint is zero, but this happens. However, such correlations are of only local importance, and become relatively more insignificant as $N \rightarrow \infty$.

2. *Experimental:* We have tested uniform grids with several algorithms on data in assorted applications, and have never observed a serious problem. These tests have included apparently quite bad input data, such as a haloed line algorithm on a wire-frame database of 11,000 edges [6]. The data contained faceted cylinders composed of dozens of faces, whose many, close, parallel lines did not intersect each other, but did pass through the same cells several times. This caused some cells to have many more edges than the average. Also, when one cylinder's projection crossed another, all the edges of one intersected all the edges of the other. Nevertheless, calculating the haloed lines took only about 10 minutes on a Prime 500 computer.

Experimentally the actual grid size is not critical, within large variations. Factors of three in grid resolution off from the optimum tend to increase execution time by under 50%. This represents a range of almost 100 in number of cells during which the execution time varies less than 50%. The optimum is so broad because the total time is the sum of two parts: putting objects into cells, which runs faster if there are fewer cells, and processing the cells, which runs slower. This insensitivity of time to grid size explains why the uniform grid can handle regions of varying data density.

3. *Quadtree Performance:* Quadtrees are useful for defining irregular regions of space in image processing and (as octrees) defining irregular objects, but they cannot handle very irregular geometric scenes. Consider a scene with N parallel edges spaced $1/N^2$ apart. A quadtree fine enough to determine that there are no intersections would require more than N^2 time to construct, so the even naive method would be faster.

The problem with hierarchical methods in general is that if $N=1,000,000$ and the quadtree is fine enough that on average each object is in a separate low level cell, then the tree, even if perfectly balanced, is $\log_4 1000000 = 10$ levels deep. In practice, 15 would be more likely. Thus accessing any object requires following ten to fifteen pointers. Pointers to random locations in memory, the usual case, defeat any hardware caching mechanism. The tree can be sorted and compressed, but this is slow and prevents dynamic updating.

The difference between flat uniform data structures and hierarchical adaptive data structures is similar to the difference between relational databases and hierarchical ones. Forcing a database to satisfy the Codd Normal Forms may increase the size, gives increased regularity in return.

Parallelizability is the final advantage of a uniform data structure. The cells that an object passes through can be determined in parallel with each object. After the data set is inverted, each cell can be processed in parallel. In contrast, constructing a tree

is inherently a more sequential operation, and accessing all elements by first traversing the root may cause collisions there (depending on the model of parallel machine) unless the top few nodes of the tree are replicated.

The problems with using trees on parallel machines have been observed by Hillis [12, 13]: "One case where our serial intuition misled us was our expectation that parallel machines would dictate the use of binary trees. It turns that linear linked lists serve almost as well, since they can be inverted to balanced binary trees whenever necessary and are far more convenient for other purposes."

Nevertheless, a limited hierarchical data structure may prove useful in the future if the available storage is hierarchical. On a parallel machine such as a Connection Machine or Hypercube, it may be reasonable to partition the data in subsets small enough to fit onto one processor. Each processor would then apply a uniform grid. This two level data structure would increase computation costs, since it would be harder to parallelize the initial partition step, but it would lower communication costs, which are often more important.

2.2. Comparison to Functional Hierarchies

Although the data are often originally organized into a functional hierarchy, such as described in standards like PHIGS, we ignore this and flatten the directed acyclic graph into a set of objects before processing. This is because two objects' spatial coincidence, especially in projection, often bears little resemblance to their functional relation, or even to their spatial relation in E^3 .

2.3. Comparison to Theoretically Optimal Methods

Chazelle and Edelsbrunner's line intersection algorithm [3], has a worst case time of $T = \Theta(N \log N + K)$. However, the algorithm is much more complicated to implement and appears to be quite difficult to parallelize. There are simpler "scan line" algorithms with $T = \Theta((N+K) \log N)$. However they also appear inherently sequential. One could start P separate scanlines in parallel, but that would incur an overhead in splitting the edges between the processors. The scan line methods also take up to $T = \Theta(N^2 \log N)$, that is, worse than the naive method, for very bad data.

Both of these theoretically optimal methods have a $\log N$ factor in the time because they work with a theoretical machine model which is weaker than real computers. Thus they assume that sorting takes $\Theta(N \log N)$ expected and maximum time, which is true only for pair-comparison models of machines. If we assume that in constant (not \log) time we can write to any one of N words of memory, then sorting i.i.d. keys takes $\bar{T} = \Theta(N)$ time with the address calculation sort, [15]. A radix sort, similar to a bucket sort, is also much faster than the theory on a pair-comparison machine would suggest. The procedure is as follows.

- a) Consider the key to be composed of 8-bit bytes. Read the data once to determine the frequency distribution of values of each byte of the keys.
- b) Read the data and write each record to one of 256 lists depending on the value of lowest order byte of the key. The 256 lists are really a partition of a work space of N records, where the amount of space each list will need was known from the frequency distribution found in step (a).

- c) Read the 256 lists in order and distribute the data into a new set of 256 lists according to the second lowest byte of the keys.
- d) Repeat for each byte of the key.

For B byte keys, this requires reading each record $B+1$ times and writing it B times, with $B = \log_{256} N$ usually. In contrast, the pair comparison model quicksort reads each record about $1.4 \log_2 N$ times.

2.4. Application to Edge Intersections

Finding all the intersections among a large set of small edges was the first application of uniform grids. The algorithm, presented in [1, 2, 10, 14], has expected time equal to the sum of the sizes of the input and output. The sophistication of the algorithm rests in its simplicity, as well as in the appropriate implementation of the abstract data structure for the cell-edge information. The various possibilities examined include these:

1. With G = the number of grid cells on a side, and M = the maximum number of edges per cell, allocate a $G \times G \times M$ array for the edges. This is fast but too big, even with virtual memory.
2. Allocate a $G \times G$ array of pointers for G^2 linked lists of edges. This is smaller but accesses storage randomly. However, that has not been a problem in practice. No sorting is needed here. We used this and the next method alternately.
3. Write a list of (cell, edge) ordered pairs and then sort it. This accesses storage sequentially, and avoids thrashing, but requires a sort, which adds a (small) log factor to the time. Note that if a cell number takes the same space as a pointer, then this method requires the same storage as the previous one. Note that an empty cell requires no storage at all here.
4. Let E = expected number of edges per cell. Allocate $G \times G \times E$ array for edges and use another method for overflow. This is the smallest but most complicated method.

For parallel intersection detection, there are several possibilities. These are designed for the Sequent's model of a common memory with interlocks and separate processors [17].

1. Let P = number of processors, and multiply any sequential data structure above by P . This is too big.
2. Use one (cell, edge) list, or one $G \times G$ array of linked lists, for all the processors and interlock the append operation. This assumes that a relatively small fraction of the time is spent storing the new (cell, edge) information, which may not be true.
3. As above, but each time a processor requires an element to store edges in, then allocate it several elements to reduce the number of interlocks and possible collisions. At the end, compress out the unused elements.




Figure 1: WaRP1 Integrated Circuit

The largest example used so far is a complete design of WaRP1, a VLSI chip by Jim Guilford, figure 1. It has 1,819,064 edges, and we found all 6,941,110 intersections in 178 seconds on a Sun 4/280 with 32 MB of real memory. The program took advantage of the large flat address space by storing all the intermediate and final data in 110 MB of virtual memory. The virtual memory was not thrashed. The algorithm was implemented as a C program using the Sun-supplied compiler, which doesn't optimize. Using commercial quality compilers and assembler modules for inner loops would presumably reduce this time considerably. This version did take advantage of the fact that all the edges were horizontal or vertical, which gave us an estimated factor of three in speed, and simplified the code.

We also implemented a version for general (non-rectilinear) edges, and tested it on the complete USGS (United State Geological Survey) DLG (Digital Line Graph) sampler tape, of Chikamauga, Tennessee, figure 2. This database contains all surface features, such as roads, railroads, rivers, streams, power lines, and pipelines, but not the contour lines. Finding all 144,666 intersections among the 116,896 edges took only 37 CPU seconds. The unevenness of the data can be seen and its irregularity quantified: for the edge lengths, $\mu = 0.00231$, while $\sigma = 0.0081$.

The uniform grid structure is remarkably insensitive to the actual grid spacing. Figure 2 shows the effect of processing the Chikamauga data and varying the grid resolution over a range of from 50×50 to 2000×2000 . Over a factor of eight in grid resolution, from 125 to 1000, the time is less than 65% worse than the minimum. From 175 to 800, the time is less than 30% worse than the minimum, which occurs at 325. The insensitivity is due to the time to insert the edges into cells running slower when




Figure 2: Chikamauga DLG

there are more cells, while the time to process the cells runs faster (within broad limits) because the cells are emptier. Contrary to our expectations, at the optimal grid resolution, most of the time is spent in processing the cells.

When implemented on a 16 processor Sequent Balance 21000, which has separate processors with a common bus and memory, the program ran 10 times as fast with 15 processors in use as with one processor. Finding all 81,373 intersections in a 62,045 edge database (half the Chikamauga data) took only 28 seconds elapsed time. Unfortunately, the Balance with 16 National Semiconductor 32000 processors is an older machine which is slower than the Sun 4. Figure 4 shows the relative parallel efficiency, that is $(\text{time on one processor}) / P / (\text{time on } P \text{ processors})$. It is 66% for $P = 15$, which is quite good.

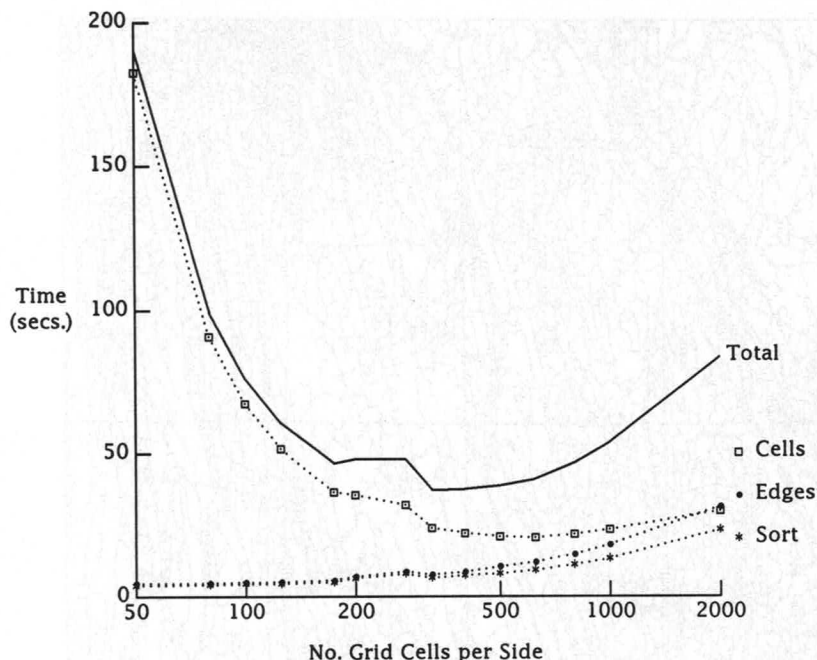


Figure 3: Effect of Grid Size on Edge Intersection Time

3. LOCAL TOPOLOGY DATA STRUCTURES

New formulae for polyhedron edge length, surface area, volume, and point inclusion testing are possible using only local topological information. The input data is the set of vertex, edge, and face coincidences. Neither any global topology, nor even the edges, faces, nor complete vertex neighborhoods are required. Since the formulae merely sum a function of each set element, they are suitable for a vector processor or a SIMD machine. The computation is a vector reduction operation. This contrasts with the usual data structure for representing a polyhedron, which stores the complete topology consisting of components, shells of faces, loops of edges, etc. Special attention is needed for the non-manifold cases [18-20]. That data structure is hierarchical, with numerous pointers.

Another conceivable data structure would be to represent the polyhedron as a result of unions, intersections, and complements of halfplanes defined by the faces. However Guibas has shown that this is not always possible in E^3 , although it is possible in E^2 . One might also try to represent a polyhedron as the disjoint union of tetrahedra with vertices chosen from the polyhedron vertices. This is also impossible, as cited in O'Rourke [16]. The counterexample is a triangular prism with the top rotated 30° relative to the bottom. Some of the following formulae have been independently discovered and used in a solid modeler by Gursoz and Prinz [11]; they call this data




Figure 4: Parallel Efficiency for Intersecting 62,045 Edges on Sequent

structure a *cusp*.

Assume the following notation and definitions.

O is the coordinate origin.

P is a vertex. When it is necessary to refer to individual components, we will use $P = (P_x, P_y, P_z)$.

T is a unit tangent vector from P along an edge.

N is a unit vector normal to T , in the plane of a face adjacent to both P and the edge, and pointing towards the face's interior.

B is a unit vector normal to both T and N , which makes it normal to the face plane, and pointing towards the interior of the polyhedron.

L is the total length of all the faces' perimeters, with each edge counted once for each face that it is adjacent to.

A is the total area of all the faces.

V is the volume of the polyhedron.

σ is the sign function, $\sigma(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases}$

\hat{P} is a unit vector parallel to P .

Consider the polyhedron vertex, and its neighborhood, shown in figure 5.




Figure 5: Polyhedron Vertex - Edge - Face Adjacency

The polyhedron is represented as the set $\{(P, T, N, B)\}$, where there is one element for each occurrence of an adjacency of a vertex, edge, and face.

Given this, the following theorems are true; the proofs are in [8].

$$L = - \sum P \cdot T$$

$$A = \frac{1}{2} \sum P \cdot T \cdot P \cdot N$$

$$V = - \frac{1}{6} \sum P \cdot T \cdot P \cdot N \cdot P \cdot B$$

$$\text{Let } s = \sum \sigma(P_x T_x) \sigma(P_x N_x + P_y T_y) \sigma(P \cdot B)$$

$$\text{Then if } O \text{ is } \begin{cases} \text{outside the polyhedron} \\ \text{inside the polyhedron} \\ \text{on the boundary} \end{cases} \text{ then } s = \begin{cases} 0 \\ -8 \\ 0, -4, \text{ or } -8 \end{cases}$$

If O is on the boundary, then its value is the average of the value for points in the neighborhood above and those below. There is no loss of generality in testing O since any arbitrary point may be transformed to O . Although there are cases, such as display and filleting, where the complete topology is required, these formulae show that there are also many cases where it is unnecessary, and that then the calculations are both faster and parallelizable.

Other local topological methods are also possible. The *vertex neighborhood* method [7], defines the polygon as a set of the neighborhoods of the vertices, that is as $\{(P, \hat{A}, \hat{B})\}$, where P is a vertex's position, and \hat{A} and \hat{B} are unit vectors along the two adjacent edges. \hat{A} and \hat{B} are oriented so that sweeping from \hat{A} to \hat{B} traverses inside the polygon. For each neighborhood, we can define a weight function assigning a value to each point in the plane depending on its relation to the neighborhood. The sum of these weight functions is a characteristic function for the polygon, which can

be used for point inclusion testing and mass property calculation. This also extends to E^3 .

Alternatively, one may split each edge into two semi-infinite rays [4]. Again, point inclusion testing and mass property calculation are a parallelizable vector reduction operation on the set of data, regardless of the global topology.

4. COMBINING SEVERAL TECHNIQUES

4.1. Mass Properties of the Union of Many Polygons

Given a set of many small polygons in the plane, we want to find the covered area and the perimeter of the union. This is useful in VLSI design for circuit extraction from the mask geometry of a chip. The techniques used for this problem are the uniform grid to find edge intersections and the vertex neighborhood technique. With the notation:

N = number of input edges

K = number of output edges,

a summary of the method is as follows.

1. *Each output vertex is either an input vertex or an edge intersection.* Find all edge intersections.
2. *An output vertex of the union is not in any input polygon.* Test each point returned above for inclusion in any input polygon. For this we have an efficient point inclusion testing algorithm which takes expected constant time per point tested.
3. For each surviving point, determine its local neighborhood.
4. Use the vertex neighborhood algorithms to determine area, perimeter or any other mass property.

Note that we never need to determine the number of components, or even the edges of the output. Nevertheless, we could link the vertices if necessary in $\bar{T} = \Theta((N+K)\log(N+K))$. Also, the algorithm is obviously parallelizable.

The algorithm, described in more detail in [9], was implemented in C and run on a Sun 4/280 machine. The algorithm was timed for several database fragments of the chip. The timings obtained for computing area and perimeter are given in Table 1.

4.2. Mass Properties of CSG Trees

The same concepts can be extended to determining the mass properties of constructive solid geometry trees. The procedure in E^2 would be this.

1. Determine all the intersections of any edges anywhere in the tree. The time is $\bar{T} = \Theta(N+K)$ for N input edges and K output intersections.
2. Determine the local neighborhood of each intersection. Based on this delete any intersections that are not vertices of the output object. If the depth of the tree

Database	Number of Rectangles	Number of Intersections of Rectangle edges	Grid Size	Total Time (sec.)
VLSI Data - xfacea.mag	109	960	150	0.41
VLSI Data - xfacell.mag	490	4343	65	0.53
VLSI Data - part of WaRP1	2500	25232	400	5.32
VLSI Data - part of WaRP1	25000	262058	750	34.28
VLSI Data - part of WaRP1	50000	538776	1000	82.76
VLSI Data - part of WaRP1	100000	1457593	1500	199.36

Table 1: Timing for Measure of Unions of Rectangles

is D , then using a uniform grid to test one intersection will take about $T = \Theta(\log D)$ for a total of $T = \Theta(K \log D)$.

3. Make one pass through the remaining set of vertices to calculate the desired mass property, such as area or total edge length.

The above method would fail if many objects, which did not actually generate vertices of the result, intersected. This must be tested by analyzing actual CSG trees, but we do not expect a problem since as the number of object grows, their size shrinks.

4.3. Map Overlay with Multiple Precision Rationals in Prolog

The problem of arithmetic roundoff errors that occur when the intersection of two lines is not exactly representable is well-known. Cartography simplifies the problem compared to general CAD databases since cartography uses only straight lines in E^2 . Therefore the intersections are exactly representable in rational numbers. Although the numerators and denominators may be multiple precision, this is a problem only if the computation tree for the operation is deep. After considering various representational issues [5], we combined this with a test of the usefulness of Prolog as an implementation tool by implementing a multiple precision rational number map overlay package in Prolog [22]. As would be expected, the program ran quite slowly, although this would be improved by the Prolog hardware boards now in existence. On the other hand there were no roundoff errors.

5. SUMMARY

By synthesizing techniques such as uniform grids and local topological data structures, it is possible to derive geometric algorithms which can process large data sets with the following properties.

- *Simple data structures*, that is, sets, instead of trees. This obviates the need to tree balance, which is messy, but essential for optimal searching. Hierarchical methods such as quadtrees are not necessary.

- *Ease of implementation*, because of the above.
- *Parallelizability*, because the operations are mapped over sets of elements. The algorithms can be implemented even on a SIMD or a vector machine. Scheduling the cooperating processes becomes easier because of the simple set-based data structures. This results in lesser overhead, and hence gives better speed-ups.
- *Speed of execution* on all the real data seen so far, including regions of widely differing density.
- 4. *Dynamizability*, since the data structures could be extended to allow us to add and delete objects on-line, and quickly know the changes in intersections and mass properties.

We are now extending these concepts to boolean operations and hidden surface calculations on many polyhedra, fast evaluation of mass properties of the regions of overlaid maps, and to evaluation of CSG trees.

ACKNOWLEDGEMENTS

Partial support for this work was provided by the Directorate for Computer and Information Science and Engineering of the National Science Foundation under Institutional Infrastructure Grant No. CDA-8805910. It was also supported by NSF under PYI grant CCR-8351942. It was implemented on equipment provided by Sun Microsystems at a discount in the Computer Science Department and Rensselaer Design Research Center at RPI.

REFERENCES

1. Akman, Varol and Wm. Randolph Franklin, "Geometric Computing and the Uniform Grid Data Structure," *Computer Aided Design*, 1989. (to appear)
2. Chandrasekhar, Narayanaswami and Manoj Seshan, "The Efficiency of the Uniform Grid for Computing Intersections," M.S. thesis, Electrical, Computer, and Systems Engineering Dept., Rensselaer Polytechnic Institute, December 1987.
3. Chazelle, Bernard and Herbert Edelsbrunner, "An Optimal Algorithm for Intersecting Line Segments in the Plane," *Foundations of Computer Science - 29th Annual Symposium*, White Plains, October 1988.
4. Franklin, Wm. Randolph, "Rays - New Representation for Polygons and Polyhedra," *Computer Graphics and Image Processing*, vol. 22, pp. 327-338, 1983.
5. Franklin, Wm. Randolph, "Problems with Raster Graphics Algorithms," *Proceedings of the Workshop on Data Structures for Raster Graphics*, Springer-Verlag, Steensel, The Netherlands, June 24-28, 1985.
6. Franklin, Wm. Randolph and Varol Akman, *A Simple and Efficient Haloed Line Algorithm for Hidden Line Elimination*, Univ. of Utrecht, CS Dept., Utrecht, October 1985. report number RUU-CS-85-28

7. Franklin, Wm. Randolph, "Polygon Properties Calculated from the Vertex Neighborhoods," *Proceedings of the Third Annual Symposium on Computational Geometry*, pp. 110-118, Waterloo, Canada, 8-10 June 1987.
8. Franklin, Wm. Randolph, *Vertex Based Polyhedron Formulae*, Rensselaer Polytechnic Institute, November 1988. (submitted for publication)
9. Franklin, Wm. Randolph and Mohan Kankanhalli, *Measure and Perimeter of the Union of a Set of Iso-Rectangles*, Rensselaer Polytechnic Institute, November 1988. (submitted for publication)
10. Franklin, Wm. Randolph, Narayanaswami Chandrasekhar, Mohan Kankanhalli, Manoj Seshan, and Varol Akman, "Efficiency of Uniform Grids for Intersection Detection on Serial and Parallel Machines," *Computer Graphics International*, Geneva, May 1988.
11. Gursoz, E. Levent and F.B. Prinz, Carnegie-Mellon University, Rensselaerville, NY, September 1988. presentation at RPI-IFIP-NSF Workshop on Geometric Modeling
12. Hillis, W. Daniel, *The Connection Machine*, MIT Press, Cambridge, MA, 1985.
13. Hillis, W. Daniel and Guy J. Steele, "Data Parallel Algorithms," *Communications of the ACM*, vol. 29, no. 12, pp. 1170-1183, December 1986.
14. Kankanhalli, Mohan, "The Uniform Grid Technique for Fast Line Intersection on Parallel Machines," *M.S. Thesis*, Electrical, Computer & Systems Eng. Dept., Rensselaer Polytechnic Institute, Troy, NY, April 1988.
15. Knuth, D.E., *The Art of Computer Programming, Volume 3: Sorting and Searching*, Addison-Wesley, 1973.
16. O'Rourke, Joseph, *Art Gallery Theorems and Algorithms*, Oxford University Press, 1987.
17. Sequent Computer Systems Inc., *Balance Technical Summary*, 1986.
18. Weiler, Kevin, "Edge Based Data Structures for Solid Modeling in Curved-Surface Environments.," *IEEE Computer Graphics and Applications*, vol. 5, no. 1, pp. 21 - 40, IEEE, January 1985.
19. Weiler, Kevin, "Topological Structures for Geometric Modeling.," *Rensselaer Polytechnic Institute PhD Dissertation*, Troy, NY, August 1986.
20. Weiler, Kevin, "Non-Manifold Geometric Boundary Modeling," *SIGGRAPH '87 Advanced Solid Modeling Tutorial*, July 1987.
21. Wu, Peter Y.F., *Polygon Overlay in Prolog*, ECSE Dept., Rensselaer Polytechnic Institute, Ph.D. Thesis, Troy, NY, August 1987.
22. Wu, Peter Y.F. and Wm. Randolph Franklin, "A Logic Programming Approach to Cartographic Map Overlay," *International Computer Science Conference*, Hong Kong, December 1988.