# Progressive Conversion from B-rep to BSP for Streaming Geometric Modeling

Chandrajit Bajaj<sup>1</sup>, Alberto Paoluzzi<sup>2</sup> and Giorgio Scorzelli<sup>3</sup>

<sup>1</sup>ICES and Computer Sciences Dept., Univ. of Texas at Austin, <u>bajaj@cs.utexas.edu</u> <sup>2</sup>Dip. Informatica e Automazione, Univ. "Roma Tre", <u>paoluzzi@dia.uniroma3.it</u> <sup>3</sup>Dip. Informatica e Automazione, Univ. "Roma Tre", <u>scorzelli@dia.uniroma3.it</u>

### ABSTRACT

We introduce a novel *progressive* approach to generate a Binary Space Partition (*BSP*) and convex cell decomposition from any input triangles boundary representation (B-rep), by utilizing a very fast computation of the surface inertia. A solid model is so generated at progressive levels of detail. This approach relies on a simple variation of standard *BSP* trees, allowing for labeling cells as *in*, *out* and *fuzzy*, and permits a complete representation of a solid mesh as the Hasse diagram of a cell complex. Our new algorithm is embedded in a streamlined computational framework using four types of dataflow processes, that continuously *produce*, *transform*, *combine* or *consume* subsets of cells depending on the number or their input/output streams. A varied collection of geometric modeling techniques are integrated in this streaming framework, including polygonal, spline, solid and heterogeneous modeling with boundary and decompositive representations, Boolean set operations, Cartesian products and adaptive refinement. The real-time B-rep to *BSP* streaming results we report in this paper are a large step towards the unification of rapid conceptual and detailed shape design methodologies.

**Keywords:** Geometric modeling and programming, Representation conversion, Solid modeling. Binary Space Partitions.



Fig. 1. (a) BSP generated by boundary polygons and normals; (b) solid helicoid; (c) progressive intersection.

## **1. INTRODUCTION**

We introduce here a novel progressive conversion from triangle boundary representations (B-rep) to solid Binary Space Partition (BSP) decompositions. The algorithm may be used for importing geometric models into a dataflow (streaming) pipeline to be used for shape modeling of complicated geometries. Streaming pipelines are commonly used

in graphics engines; conversely, their use for geometric modeling is relatively new. We use balanced BSP trees as progressive representations of shapes, giving multiple levels of detail. The algorithm we introduce in this paper generates finer and finer geometric approximations progressively, and closely integrates to streamlined rendering for instant visualization. In our framework, we already used specialized routines to generate progressive representations of primitive objects, such as spheres and cylinders, and introduce here a novel algorithm to import B-rep models from external data stores, and embed them into complex generative expressions, in order to build progressive BSP-trees to be refined on demand. Also, view-dependent data may be propagated upstream our geometric modeling pipeline so that the computation is focused on detailing only the model features currently of interest for the shape designer.

Our novel conversion algorithm from B-reps to solid decompositions with convex cells takes as input a boundary triangulation and does not need any representation of the boundary topology. It produces a natively balanced BSP-tree with only O(n) preprocessing. The preprocessing consists in the computation of the Euler tensor of each input triangle. The Euler tensor is strictly related to the inertia tensor, and takes into account the contribution of the triangle to the spatial distribution of surface points. The tensor matrices are then attached like textures to triangles, and are only recomputed when some triangles are split. A best-fitting parallelepiped is so generated for the input surface using the eigendecomposition of the Euler tensor, and represented as the intersection of six boundary planes using a standard BSP. The interior cell is further split by the principal plane normal to the direction of maximal elongation. Two local best-fitting parallelepipeds are then computed for the two half-surfaces, and their intersection with the two sub-cells is used to add detail to the solid approximation already generated. Some cells of the decomposition are labeled as either external or internal to the boundary and are not further split. The splitting parallelepipeds, until only few input triangles remain in a cell. In the very last step only the boundary triangles are used to add the maximum detail.

The new algorithm integrates seamlessly with our streaming approach to geometric modeling [20]. It provides a native generation of balanced BSP trees with only linear preprocessing of the imported triangulations. As it is well-known, a balanced BSP tree provides an efficient spatial index allowing for fast point location, collision detection, and distance field computations.. We demonstrate the practical validity of this approach on rather complicated geometric models constructed with our prototype system. Also, our approach provides a natural convex decomposition [1] of large-scale models, resulting in a representation that is highly portable and scalable, and can be effectively used on PCs and HPC architectures of different kinds. The approach is aimed at supporting generative geometric modeling, starting from either primitive or imported shapes, that may come from external data stores. Such atomic shapes may be either transformed with affine or projective transformations, or aggregated within hierarchical assemblies, or combined by several operators, including Boolean set operations [18], Cartesian product and Minkowski sums of point-sets. This approach is *embarrassingly parallel* (see [20]), from limited to no communication between processes generating the portions of the model on the nodes of a PC cluster, where communication overhead is negligible and the speedup is nearly linear with the number of processors. After the initial broadcast of the generating expression (the executable code) to all nodes, each job is simply specified by a clipping BSP-tree. Each clipping BSP is simply given exactly by one of the paths of the tip of our progressive BSP-tree. Furthermore, the computation can be easily organized as a queue of independent jobs, that may also be generated hierarchically, and submitted to some computational infrastructure by using high-throughput batch tools for clusters or grids.

### 2. STREAMING GEOMETRIC MODELING

Our streaming approach results in a fine-grain streamlined parallelism where suitable geometric data structures flow between specialized processes, with the resulting shape produced by progressive refinements of a first approximation of the result. The data tokens flowing between different computations are a couple of pointers to the twin representations of the mesh, i.e. (a) a BSP-node (actually either a linear hyperplane or a leaf label) and (b) its associated *d*-cell in the Hasse diagram of the current mesh. For example, when computing a Boolean operation between large-scale objects, the result with Naylor's algorithm [14] is only generated at the end of the entire computation, and may require an intolerably long time. With our approach, a continuously refined estimate of the Boolean result is available from the very beginning. If the output appears unsatisfactory, the task can be instantly terminated. The streaming computation can be also terminated using a local threshold for the geometric approximation error, and possibly depending on the viewer's gaze.



Fig. 2. A chess model: (a) union of a cylinder, a cone, a scaled sphere, two cones, a scaled sphere and a sphere, rendered flat at maximum resolution; (b) exploded geometry; (c) differences from parallelepipeds

### 2.1 Progressive BSP data structures

In our approach two main data structures are used: **BSP-trees** for progressive generation of geometry, and doubly linked list representation of **Hasse diagrams** for the storage of already generated geometry.

A short summary of the main properties of BSP trees [14] is included for completeness. In particular, each node of a BSP tree: (a) is associated to a convex cell; (b) if non-leaf, then contains a hyperplane splitting its cell; (c) if leaf, then contains a label that characterizes its cell as either IN, OUT or FUZZY; (d) is defined as the set intersection of the halfspaces associated to the (unique) path from the node to the root; (e) equals the set union of cells associated to the subtree rooted within it (see Figure 3). With respect to standard BSP trees, we only add a FUZZY label to nodes that can be further split progressively, so that each *frontier* (time-dependent space partition) is always subdivided into solid (IN), empty (OUT) and yet undecided (FUZZY) cells.



Fig. 3. Frontier evolution as progressive refinement of the space partition.

The search for maximum runtime efficiency was the main reason for using the B-rep based on the Hasse diagram [24], which allows for explicit and complete storage of both geometry and topology of the model. In order theory, a Hasse diagram is a graph, whose nodes are the members of a finite partially ordered set S, and where there is an arc from x to y iff: (a) x < y and (b) there is no z such that x < z < y. In this case, we say y covers x, or y is an immediate successor of x. Hasse diagrams can be used to give a complete representation of the structure of d-polytopes and d-meshes, with respect to the inclusion relation between k-faces,  $0 \le k \le d$ . The very basic operation is the *split* of a fuzzy *d*-cell with a hyperplane, with efficient update of the Hasse diagram, performed using the fast algorithm [2]. The inverse operation is called *join*, and updates the representation by substituting two split *d*-cells (and their split *k*-faces) with their convex hulls. These operations are geometrically robust, i.e. able to withstand the effects of numerical errors. Furthermore, they are very fast, having to deal with hundreds of thousands or millions of cells of varying dimensions in practical cases.



Computer-Aided Design & Applications, Vol. x, No. y, 200z, pp ww-uu

#### Fig. 4. Splitting of a *d*-cell *c* with a hyperplane *h*, and corresponding Hasse diagram.

Notice that we only handle collections of piecewise-linear bounded and convex sets, i.e. complexes of polytopal cells. Note also that the approach is *dimension-independent*, in the sense that both the data structures and the operations would work with solid (codimension 0) pointsets of arbitrary dimension *d*.

#### 2.2 Streaming framework

In our streaming framework, geometric objects are always produced by the (progressive) evaluation of a generating expression, which is compiled into a dataflow pipeline made from four kind of processes, denoted respectively as producers, transformers, combiners and consumers, according to the number of their input/output streams. A producer, or *builder*, is a process with no input, and one output stream. It continuously generates progressive polyhedral approximations, at finer and finer levels of detail. A special effort was devoted to supply our streaming technology with a rich set of producers, providing solid primitives, polynomial and rational splines, subdivision surfaces and polygon models. A transformer is a process with one input stream and one output stream. Our transformers either apply affine or projective transformations or produce the complement or extrusion of their input. A combiner is a process with more than one input stream and a single output stream. Such a process combines several progressive BSP streams and return the stream of the result. In each combiner process, an efficient variable-grain merge of splitting planes is used to combine the operation arguments. In particular, a cursor pointer is set to the current input splitting node. This pointer is moved to the following input tree every g splitting nodes, where g stands for granularity. A consumer is a process with one input stream and no output streams, that is used either to compute suitable model properties or to analyse or visualize the progressively generated model. Depending on the computed properties, a consumer process may decide which cells of the input stream should be expanded, i.e. further detailed, or collapsed, (joined) in a single cell, and hence provides a mechanism for adaptivity and control with feedback.

The conversion algorithm described in this paper implements an important **builder** component in our streaming architecture. In particular, it provides a mechanism to *progressively import* into the streaming pipeline models extracted from an *external store* and, in particular, supplied using the weakest B-rep solid representation, i.e. given as a bunch of raw triangles. The only constraint on the input triangles is that they must be coherently oriented. In other words all triangles must be either implicitly or explicitly associated with the correct external normal to the surface they belong to. Such a constraint is normally satisfied by most data structures, and in particular by the data files produced by laser scanners of 3D solids.

#### **3. Progressive CONVERSION ALGORITHM**

The new algorithm generates a solid decomposition into convex cells by producing a balanced BSP-tree, and by using only a coherently-oriented decomposition of the object boundary into raw triangles, and without any need for a representation of the boundary topology.

The idea is very simple. A fast preprocessing, executed in linear time with the number of input triangles, computes for each triangle a 4x4 matrix that codifies numerically its mechanical behavior. The sum of all such matrices gives a good representation of the spatial distribution of the surface points. Such a matrix may be used used to generate a best-fitting ellipsoid, centered in the center of mass of the surface, that is mechanically equivalent to the triangulated surface.

The first solid approximation of the surface is given by the minimal best-fitting parallelepiped, parallel to the principal axes of the ellipsoid and strictly confining the surface. Such a solid is represented as a standard BSP-tree, as intersection of six linear subspaces tangent to the boundary surface. This enclosing solid is then split by a plane passing for the center of mass and normal to the direction of maximal elongation.

The set of input triangles is in turn split into two subsets contained in the two half-spaces of this plane, and the two subsets are associated to the below and above sub-trees of the BSP root. The confinement of each surface subset into a narrower and properly rotated best-fitting parallelepiped and its splitting into the principal direction is recursively repeated for each sub-tree.

As it is well-known, the leaves of a BSP-tree give a partition of the embedding space into convex cells. Some cells are labeled as *empty* or *full*, i.e. as either external or internal to the boundary surface. Such cells are not further split. Other cells cross the boundary, and are provisionally classified as *fuzzy*.

The shape confinement using the local best-fitting parallelepiped is then repeated, either for every fuzzy cell or only for those cells where better detail is needed. Each fuzzy cell is detailed by its intersection with the best-fitting parallelepiped. This intersection of the current cell always produces some empty or full subcell, togheter with a smaller

fuzzy cell. Finally this one is split along the center of mass, and the refinement is recursively repeated, until the boundary planes are used to get the final approximation.

### 3.1 B-rep Streams

It is fairly straightforward to set up a producer process importing a B-rep, externally defined by some standard polygonal format, e.g. either a wavefront and java3D obj file, into an input stream for our geometric pipeline. The boundary representation given by polygons and normals must be *coherently oriented*. A filtering of the input file to cope with nonplanar polygons and other geometric inaccuracies may be needed for generally archived geometric models used primarily in computer graphics [23]. The output stream of coherently-oriented triangles, is then converted into our twin progressive-BSP trees by the algorithmic steps described below.

### 3.2 B-rep to BSP Algorithm Outline

A primary technique of our method is the computation of the inertia of triangle subsets by contraction of the precomputed inertia of each triangle, and the eigendecomposition of the inertia of triangle subsets to bound their shape optimally and recursively.

Please notice that, in the *d*-dimensional case, the shape confinement is made by 2 extremal tangent hyperplanes for each of the *d* eigenvectors of the Euler matrix. The intersection of the corresponding 2d hyperspaces produces the best-fitting (hyper)parallelepiped of the boundary subset contained in the current cell. In 3D, there are 6=2x3 such planes.

#### Initialization

(a) The affinely extended Euler tensors [9] of each input triangle are first computed (in linear time).

(b) The entire set of input triangles is associated with the BSP root.

(c) The entire  $E^3$  space (that is convex) is associated to the root.

(d) The label of the root is set to FUZZY.

#### **Recursive** case

(a) The current FUZZY cell is split by at most 6 orthogonal hyperplanes that are normal to the eigenvectors of the matrix representation of the Euler tensor of the current triangle subset.

(b) Such planes are computed via the minimum and the maximum value of the linear function  $w = a \cdot v$  evaluated on the vertices v of the current triangle subset, where a is the current eigenvector.

(c) For each eigenvector, at most three convex cells are produced by two max-min parallel hyperplanes, that are either  $\{OUT, FUZZY, OUT\}$  or  $\{OUT, FUZZY, IN\}$ .

(d) Each FUZZY cell is further split by the principal hyper-plane associated to the maximum eigenvector.

(e) A smaller subset of triangles is associated to each split cell via containment test of their vertices.

(f) Triangles crossing a splitting plane are split, and the (sub)triangles are associated to node subtrees.

#### Basic case

The recursive inertia-based division stops when the current cell only contains a small number of boundary triangles. A final cell splitting is executed using the planes of the boundary triangles.



Figure 5: Progressive construction of the BSP in 2D starting from edge boundary elements. OUT cells are white, IN cells are black, and FUZZY (i.e. yet undecided) cells are grey. The ellipsoids of inertia of edge subsets are also drawn.

#### 3.3 Fast Inertia Computation

Mass and inertia of compact point-sets (curves, surfaces, volumes) are defined as integrals of certain scalar fields  $f: E^3 \rightarrow R$  over the point-set. If S  $\subset$  E<sup>3</sup> is a surface, then its mass M, first moments  $M_x$ ,  $M_y$ ,  $M_z$ , second moments  $M_{xx}$ ,  $M_{yy}$ ,  $M_{zz}$ , and products of inertia  $M_{yz}$ ,  $M_{xy}$  are defined as

$$\int_{S} f(x, y, z) \, dS = \int_{S} x^{\alpha} y^{\beta} z^{\gamma} dS = I_{S}^{\alpha \beta \gamma}$$

where the scalar field f(x,y,z) is respectively equal to 1 for mass; to x, y and z for first moments; to  $x^2$ ,  $y^2$  and  $z^2$  for second moments; and to yz, xz and xy for products of inertia. The centroid or center of mass  $g=(g_x, g_y, g_z)$  is defined by the ratios

$$g_x = \frac{M_x}{M}, \quad g_y = \frac{M_y}{M}, \quad g_z = \frac{M_z}{M}$$

of the first moments to the mass. For *homogeneous* point-sets, where the density is constant, the centroid depends only on the geometry. In conclusion, an integral over an entire surface S, open or closed, as well as over every subset  $S' \subseteq S$ , is a summation of integrals over the triangles  $\tau \square S'$ : please add epsilon

$$I_{S'}^{\alpha\beta\gamma} = \int_{S'} x^{\alpha} y^{\beta} z^{\gamma} dS = \sum_{\tau \in S'} I_{\tau}^{\alpha\beta\gamma}.$$

The integrals described above for a point-set S can be arranged into a 4x4 matrix, that represents by components the affinely extended Euler tensor [9] of S;

$$I_{S} = \begin{pmatrix} M_{xx} M_{xy} M_{xz} M_{x} \\ M_{yx} M_{yy} M_{yz} M_{y} \\ M_{zx} M_{zy} M_{zz} M_{z} \\ M_{x} M_{y} M_{z} M_{z} \end{pmatrix} = \begin{pmatrix} I^{200} I^{110} I^{101} I^{100} \\ I^{110} I^{020} I^{011} I^{010} \\ I^{101} I^{011} I^{002} I^{001} \\ I^{100} I^{010} I^{001} I^{000} \end{pmatrix}$$

A reference frame with origin in the surface centroid and with the first three eigenvectors of the tensor as its fundamental basis is called the surface's *principal frame*. In the principal frame, products of inertia and first moments are zero and the tensor  $I_S$  is diagonal. Since integrals are additive with respect to the integration domain, the inertia tensor of S is

the sum of tensors of triangles in S. A very fast procedure for computation of the Euler tensor of triangulated surfaces is given in [9]. There is a ten-fold time improvement of such an approach over explicit integration methods [6]. It is about a thousand fold if implemented on a modern GPU [15]. E.g. the Euler tensor of the cow model was computed in 0.17 seconds, on a triangulation with 5804 triangles, using an Intel Centrino 2.00Ghz with 1.047.784 KB of RAM and MS Visual C++ 6.0. The average computation on several models is greater than 35,000 triangles/sec.

### **3.4 Algebraic Details**

With some abuse of notation, let denote the generic input triangle as  $\tau \in S$ . In the following, without lose of generality, we denote as  $S_C \subset S$  the subset of input triangles contained in the current FUZZY cell *C* of the progressive BSP-tree. Suppose that the Euler matrices of triangles in  $S_C$  have been already computed.  $S_C$  may be either closed, i.e. may be the boundary of a solid, or an open surface. Let  $M_C$  denote the symmetric and positive definite (by definition) affinely extended Euler tensor of  $S_C$ , or better, the coordinate representation (i.e. the matrix) of such tensors in world coordinates. So we have:

$$M_C = \sum_{\tau \in S_C} M_{\tau}$$

Since the matrix  $M_C$  is symmetric, and positive definite, its eigenvalues are all reals and positive. The ratios of elements of the 4-th row to the element  $m_{44}$  give the affine coordinates of the centroid  $g_C$  of triangles contained in C. The first three normalized eigenvectors give an orthonormal basis of  $E^3$  where the inertia matrix is diagonalized.



Fig. 6. BSP solid models from B-rep triangulations, and exploded views of the cell complexes.

Consider the linear function  $w: E^3 \rightarrow R$  such that  $w(p) = p \cdot a$ , where a is the eigenvector of M corresponding to the minimum eigenvalue. It is possible to show, that a linear function take maximum and minimum value over extremal (vertex) points of a polyhedral point-set [7,14]. Hence, find a pair of minima and maxima points of w(S) over the set V(S) of vertices of S. Lets call them  $v_{min}$  and  $v_{max}$ , respectively. Clearly this task is O(n), if n=|V(S)|.

Split the cell C in at most three parts using the two parallel hyperplanes

 $p \cdot a - w(v_{min}) = 0$  and  $p \cdot a - w(v_{max}) = 0$ .

The cell will be split in only two parts if one of the hyperplanes contains a boundary facet of C. No less than two parts may arise, one of which certainly to be labeled as IN or OUT, by construction. If no splitting is possible, the entire cell C is certainly empty (OUT) or solid (IN), and the progression in the current cell will stop.

If at least one cell splitting has occurred, and *hence* one of sub-cells is FUZZY, let continue splitting the generated FUZZY sub-cell by using the max and min hyperplanes parallel to the remaining eigenvectors previously computed. The result is a confinement of the geometric data into a smaller (mimimal) rotated parallelepiped, parallel to the principal frame of the current triangles. Finally, split such FUZZY cells with the hyperplane orthogonal to the maximum eigenvector and passing for the center of mass:

 $p \cdot a - w(g(\tau)) = 0$ 

## 3.5. Eigendecomposition of the Euler Matrix

The upper-diagonal 3x3 submatrix of the affine inertia matrix has several nice properties. It is small, dense, symmetric and positive-define, so that we can safely use the simplest algorithms for computation of eigenvectors, i.e. the *direct* and *inverse power* methods [10]. First, only the minimum and maximum eigenvector must be computed iteratively, since the third one may be derived by their vector product. Both methods are iterative, and work by repeatedly applying either the matrix M or the inverse  $M^{-1}$  to an initial trial vector, say (1,1,1), to successively yield an approximation of the maximum (minimum) eigenvector. This process may be accelerated by applying a number of squaring operations to the initial matrix M. For example, after 5 compound applications of the squaring operator to M, the resulting matrix is  $M^2^5 = M^{32}$ . Hence, a single multiplication of the trial vector by it will give the same result than 32 iterations of the multiplication by M.



Fig. 7. Transformation from B-rep triangles to BSP: (a) progressive generation of model from triangles; (b) exploded view of the produced BSP; (c) a close view of the head.

### 5. CONCLUSIONS

We have presented an optimal conversion from triangled B-Rep to a BSP solid representation. The algorithm has a O(n) preprocessing, used to compute the inertia tensor of each triangle, and a  $O(n \log n)$  construction phase of a balanced (progressive) BSP. Furthermore, the algorithm is well suited to GPU implementations [15] and to streaming rendering and combination with other operations, and in particular to perform Boolean set operations on-the-fly [18]. The prototype system is presently implemented as a multithreaded library written in C and named XGE, for eXtreme Geometric Environment. An integration with the PLASM design language just started, with the aim of compiling the user functional environment into a distributed dataflow, able to exploit the available resources on demand. All the main techniques of geometric and solid modeling are well supported by this technology. The next step will concentrate on: (a) porting the library on the *cell processor* architecture, that appears to perfectly fit our streaming approach, and can be implemented using the Stanford's Brooke streaming extension of the C language (see [4,5] and [8]); and (b) a close integration of solid and physical modeling, with the goal of supporting progressive simulations and adaptive, simulation-driven refinements of the generated mesh.



Fig. 8. Progressive transformation from B-rep triangles to BSP: adaptive refinement



Fig. 9. Progressive transformation from B-rep triangles to BSP: (a) progressive solid models of the Max Plank head; (b) exploded views of the produced BSP.

# **6. REFERENCES**

- [1] Bajaj, C. L., and Dey, T. Convex Decompositions and Robustness, *SIAM J. on Computing*, 21, 2, (1992), 339-364.
- [2] Bajaj, C. L., and Pascucci, V. Splitting a complex of convex polytopes in any dimension. In SCG '96: Proc. of the *12th Symposium on Computational Geometry* (1996), ACM Press, pp. 88–97.
- [3] Bajaj, C. L., Pascucci, V., and Zhuang, G. Progressive compressive and transmission of arbitrary triangular meshes. In VIS '99: Proc. of *Visualization '99* (Los Alamitos, CA, USA, 1999), IEEE Computer Society Press, pp. 307–316.
- [4] Buck, I., Foley, T., Horn, D., Sugerman, J., Fatahalian, K., Houston, M., and Hanrahan, P. Brook for GPUs: stream computing on graphics hardware. *ACM Trans. Graph.* 23, 3 (2004), 777–786.
- [5] Buck, I., and Purcell, T. A toolkit for computation on GPUs. In *GPU Gems: Programming Techniques, Tips and Tricks for Real-Time Graphics*, F. Randima, Ed. Addison-Wesley, 2004.

- [6] Cattani, C., and Paoluzzi, A. Boundary integration over linear polyhedra. *Comput. Aided Des.* 22, 2 (1990), 130–135.
- [7] Chvatal, and Vasek. Linear Programming. W. H. Freeman and Company, New York, 1983.
- [8] Dally, W. J., Labonte, F., Das, A., Hanrahan, P., Ahn, J.-H., Gummaraju, J., Erez, M., Jayasena, N., Buck, I., Knight, T. J., and Kapasi, U. J. *Merrimac: Supercomputing with streams*. In SC '03: Proceedings of the 2003 ACM/IEEE conference on Supercomputing (Washington, DC, USA, 2003), IEEE Computer Society, p. 35.
- [9] DiCarlo, A. and Paoluzzi, A. Fast computation of inertia properties through affinely extended Euler tensor, (Submitted paper) 2005.
- [10] Gerald, C.F., and Wheatley, P.O. Applied Numerical Analysis. Addison-Wesley, NY, 1998.
- [11] Lario, R., Pajarola, R., and Tirado, F. Hyperblock-quadtin: Hyper-block quadtree based triangulated irregular networks. In Proc. *IASTED Visualization, Imaging and Image Processing* (2003).
- [12] Milicchio, F., Paoluzzi, A., and Bertoli, C. A visual approach to geometric programming, *Computer-Aided Design & Applications*, Vol. 2, Nos. 1-4, *CAD*'05, 2005, pp 411-420.
- [13] Murty, K. G. Linear Programming. John Wiley & Sons, New York, NY, 1983.
- [14] Naylor, B. F. Binary space partitioning trees as an alternative representation of polytopes. *Computer Aided Design* 22, 4 (1990), 250–252.
- [15] Owens, J. D., Luebke, D., Govindaraju, N., Harris, M., Krger, J., Lefohn, A. E., and Purcell, T. J. A survey of general-purpose computation on graphics hardware. In *Eurographics 2005, State of the Art Reports* (Aug. 2005), pp. 21–51.
- Pan, Z., Tao, Z., Cheng, C., and Shi, J. A new BSP tree framework incorporating dynamic lod models. In VRST '00: Proc. of the ACM symposium on Virtual reality software and technology (2000), ACM Press, pp. 134–141.
- [17] Paoluzzi, A. Geometric Programming for Computer Aided Design. John Wiley & Sons, Chichester, UK, 2003.
- [18] Paoluzzi, A., Pascucci, V., and Scorzelli, G. Progressive dimension-independent Boolean operations. In Proceeding of the 9-th ACM Symposium on Solid Modeling and Applications (2004), G. Elber, N. Patrikalakis, and P. Brunet, Eds., pp. 203–212.
- [19] Pascucci, V., and Bajaj, C. L. Time critical isosurface refinement and smoothing. In VVS '00: Proceedings of the 2000 *IEEE Symposium on Volume Visualization* (New York, NY, USA, 2000), ACM Press, pp. 33–42.
- [20] Scorzelli, G., Paoluzzi, A., and Pascucci, V. Parallel Solid Modeling Using BSP Dataflow. (2006). To appear on *Int. Journal of Computational Geometry & Applications.*
- [21] Stephens, R. A survey of stream processing. Acta Informatica 34, 7 (1997), 491–541.
- [22] Wiley, C., A. T. Campbell, I., Szygenda, S., Fussell, D., and Hudson, F. Multiresolution BSP trees applied to terrain, transparency, and general objects. In Proc. of *Graphics interface '97* (Toronto, Ont., Canada, 1997), Canadian Information Processing Society, pp. 88–96.
- [23] Tao-Ju, Robust Repair of Polygonal Models, Proceedings of ACM SIGGRAPH, 2004 ACM Transactions on Graphics, 23(3):888-895.
- [24] Eric W. Weisstein. "Hasse Diagram." From <u>MathWorld</u>--A Wolfram Web Resource. Url of the specific page: <u>http://mathworld.wolfram.com/HasseDiagram.html</u>