

DISTRIBUTED AND COLLABORATIVE SYNTHETIC ENVIRONMENTS

Chandrajit L. Bajaj and Fausto Bernardini
Department of Computer Science
Purdue University
West Lafayette, IN 47907
{bajaj, fxb}@cs.purdue.edu

Introduction and Overview

Fast graphics workstations and increased computing power, together with improved interface technologies, have created new and diverse possibilities for developing and interacting with synthetic environments^{1,2,3,4,5,6}. A synthetic environment system is generally characterized by the following components:

- Input/output devices, that constitute the interface between the human senses and the synthetic environment generated by the computer. Several degrees of immersion are possible, ranging from simple stereoscopic view of an image on a CRT display to a total immersion in which a head-mounted display, sound and haptic devices (force and torque feedback, tactile stimuli) are used.
- A computation system, running a real-time simulation of the environment. This can sometimes be subdivided in several subsystems, for example a simulation module running on a supercomputer coupled with a scene creation and a rendering module running on a graphics workstation.

To achieve an acceptable level of realism, the display subsystem must generate at least 10 frames per second of high quality graphics. For complex environments, this goal is still far from being reached. Nonetheless, synthetic environments have already been applied successfully in such diverse fields as: Operations in hazardous or remote environments, through telepresence; molecular modeling; flight simulations; battlefield simulation; architectural walk-throughs; surgical planning; collaborative design; education and training; entertainment, and many others.

A basic need of a synthetic environment system is that of giving to the user a plausible reproduction of the visual aspect of the objects he is interacting with. It is well known that populating a synthetic world with objects, be these spaceships in a science fiction movie, prototypes in a manufacturing design or replicas of existing real-world objects in an architectural walk-through, can be extremely time consuming. Not only it must be possible for the users to move in the environment, updating in real time his or her view of it. To convey the impression of

an immersive, active presence in an environment, a real-time simulation of the physical behavior of the various components of the environment is necessary. Moreover, in many applications the users should be allowed to interrogate objects about their associated properties, and to interact with them, for example modifying some of their properties to support *what-if* scenarios.

The goal of our Shastra research project is that of providing a substrate of geometric data structures and algorithms which allow the distributed construction and modification of the environment, efficient querying of objects attributes, collaborative interaction with the environment, fast computation of collision detection and visibility information for efficient dynamic simulation and real-time scene display. In particular, we address the following issues:

- A geometric framework for modeling and visualizing synthetic environments and interacting with them. We highlight the functions required for the *geometric engine* of a synthetic environment system.
- A *distribution* and *collaboration* substrate that supports construction, modification and interaction with synthetic environments on networked desktop machines.

Geometric Engine

A critical subsystem in all synthetic environment systems is the *geometric engine*, or the software module responsible for creating a realistic view of the simulated world, and for allowing the user to interact with it. In a typical scenario, a system requires the display of several objects, that move and must behave realistically. A user wanders in the environment, constantly changing his or her point of view. Other users (or actors) may be sharing the same environment, and a suitable representation of them could be required. The users interact with objects in the environment, for example touching their surface to pick and query an object, grabbing and moving them, or simply colliding with a wall of a room. Simulation of even the simplest form of dynamic and interaction requires collision detection and contact analysis. Fast display of complex environments requires efficient visibility computation. Querying and interacting with objects requires rapid point location and local shape control. In short, the geometric engine of a general-purpose synthetic environment system should provide efficient data structures and algorithms for:

- Shape representation
- Dynamic object insertion and deletion
- Object animation (motion, non-rigid transformation)
- Object location and closest point queries
- Collision detection and contact analysis
- Visibility ordering and culling
- View-volume clipping
- Multi-resolution representation
- Real-time high-quality rendering
- Other operations required by specific applications (e.g., set operation and interactive shape control for collaborative design)

We want to explore the use of data structures suitable to support the operations listed above in a complex, constantly changing environment. Several types of data structures that partially satisfy these needs have been proposed in the past. However, they have mainly been used in Computer Aided Design (CAD) systems or other special-purpose applications, and not as an infrastructure for a complex, dynamic and general-purpose environment. Often, CAD systems use boundary representations (Breps) to describe the geometry of the object being modeled. Breps are well suited to the implementation of the most common modeling operations. However, when this representation is to be used in a general-purpose synthetic environment, it has to be supplemented with additional structures such as octrees or bounding volumes hierarchies to achieve the required efficiency.

Several variants of the *octree* data structure have been proposed, both as a superimposed search index on existing representations of geometry, or as the main representation scheme in the system⁷. The simplest variant (and usually the one requiring the most space) is the *region octree*. In a region octree, a cubic domain is split at each node in eight equal sub-cubes. The decomposition continues recursively for each node that requires a further spatial refinement. Leaves of the octree represent empty or solid regions of space. The major drawbacks in the use of a region octree for representing the geometry are the fact that it is an approximation (unless the object is constituted by mutually orthogonal planar faces only), and its size. Another octree variant is the *PM octree* (PM stands for Polygonal Map), each leaf corresponds to a single vertex, edge or face. The only exception is that a leaf may contain more than one edge (face) if all edges (faces) are incident on the same vertex. A PM octree usually requires much less storage than a region octree^{8,9}, and it permits an exact representation of polyhedral models. PM octrees support a large suite of modeling operations, and several methods are known to convert CSG or Breps models to PM octrees and vice-versa⁷.

A *K-d-tree*¹⁰ (where K is the dimension of the domain space, in our case we will deal with 3-d-trees) is a binary tree in which internal nodes partition the space by a cut hyperplane defined by a value in one of the K dimensions, and external nodes, or *buckets*, store the points in the resulting hyper-rectangles of the partition. They allow $O(\log n)$ insertion, deletion, and point query operation. The semidynamic variant introduced in¹¹ allows constant expected time deletion, undeletion, nearest neighbor searching and fixed-radius near-neighbor searching. When used to model the geometry of an object, they suffer of some of the same disadvantages of region octrees in that they can only approximate the geometry of general objects. However, they have been successfully used in precomputing viewpoint-independent visibility and illumination information for large models of buildings, to speed up a successive interactive walk-through phase^{12,13,14}. Changes in the environment would require a new processing of the visibility information.

Binary Space Partitioning Trees (BSPT)¹⁵ are a generalization of K -d-trees. In a BSPT, the cutting planes associated to the internal nodes are not constrained to be orthogonal. Constructing a BSPT representation of one or more polyhedral objects involves encoding the polygonal faces into a binary tree of cutting planes. Notice that affine transformations on the encoded object or groups of objects do not change the tree structure, but requires only the application of a transformation to the plane equations. The spatial ordering encoded in the tree can be exploited for the speedup of intersection and visibility computation. Moreover, when a BSPT is properly constructed, it offers a *multi-resolution* representation of the object: As one descends a path in the tree, the bounded

region decreases its size monotonically. The effectiveness of this technique in practice has been proved in several applications.

Unstructured tetrahedral grids are extensively used in finite element analysis, and triangulations are pervasive in computational geometry. The use of simplicial complexes as a general approach to representing geometric shape has been advocated by several authors^{16,17}. We are currently investigating the use of *3D Regular* (a *weighted* variant of Delaunay) *triangulations*, coupled with efficient point location data structures, to provide the functionalities needed in a synthetic environment. A representation based on Hierarchical Simplicial Complexes has been recently proposed^{18,19}. (see also an extension to a hierarchy of cell complexes²⁰). In our scheme, the top-level triangulation subdivides the space in tetrahedral cells (with the associated search structure). Each cell contains an object (or a group of objects), and these are described by other triangulations. The scheme can be recursively used to achieve the level of detail in the decomposition needed by the application. Using a hierarchy of triangulations has several advantages. Regular triangulations are acyclic with respect to visibility ordering²¹. This means that, given a viewpoint, it is possible to order the cells of the complex in a sequence such that if A obstructs B, then B precedes A in the sequence. This property is preserved in a hierarchy defined as above. Moreover, a hierarchy of simplicial complexes naturally defines a multiresolution representation of the geometry, and allows fast point location and other types of queries when associated with appropriate data structures.

A 3D Regular triangulation can be built incrementally, via point insertion and topological flipping. A history DAG can be used in the incremental algorithm to allow the efficient location of which tetrahedron in the current triangulation the point to be inserted lies in, as well as to help in other point location queries required by the application. Moreover, a *Power* (weighted Voronoi) *diagram* (eventually of order k), can be easily computed to allow fast answers to closest point queries. We are considering the use of several variations of this approach. A possible version of the data structure consists in using a Delaunay triangulation that conforms to objects boundaries. When a new object is inserted in the environment, the triangulation is updated to accommodate all its faces. Methods to construct conforming Delaunay triangulations in 2D and 3D by inserting extra points are known^{22,23,24,25,26}. However, while a polynomial bound on the number of extra points is known for the two-dimensional case ($O(m^2n)$, for m edges and n vertices²⁷), whether such a bound exists for the 3D case is still an open problem. Paper²⁸ reports statistics from experiments with the 3D algorithm that show a reasonable behavior on models used in real applications. This scheme is restricted to polyhedral objects (curved surface objects can be approximated by small planar faces).

In a different scheme, the triangulation is used as the domain for implicit piecewise-algebraic surfaces^{29,30,31,32,33}. In each tetrahedron containing a piece of some object's boundary, a polynomial function $f(x,y,z)$ of degree n (where n is usually small, often $n \leq 3$) is defined, and the piece of surface is implicitly given by $f(x,y,z) = 0$. The surface patches can be made to join with some degree of continuity (for example C^1 or C^2), by using interpolants of appropriate degree. For $n = 1$, in particular, the surface pieces are planar polygons, and the representation resembles the PM octree, where the cubic cells are replaced by tetrahedral ones. This approach has the advantage of allowing a compact representation for curved objects. Moreover, it makes possible to use a more "relaxed" condition on the conforming triangulation. This needs not to conform to the object boundary, but simply contain a set of tetrahedra suitable to define the piecewise surface. An

implicit piecewise-algebraic surface is suitable for interactive local control. In a 3D synthetic environment, one could very naturally make use of 3D widgets to allow an intuitive control on the surface shape. For example handles could be attached to the surface that can be pulled or pushed to deform it.

Networked Distribution and Collaboration

We advocate the approach of integrating a collection of function-specific tools into a distributed and extensible environment, where tools can easily use facilities provided by other tools^{34, 35, 36}.

Isolation of functionality makes the environment modular, and makes tools easy to develop and maintain. Distribution lets us benefit from the cumulative computation power of workstation clusters. Tool-level cooperation allows us to exploit the commonality that is inherent to many scientific manipulation systems. An enabling infrastructure of communication and interaction tools, display and visualization facilities, symbolic processing substrates, and simulation and animation tools saves avoidable re-implementation of existing functionality, and speeds up the application development.

The Shastra environment consists of multiple interacting tools. Some tools implement scientific design and manipulation functionality (the Shastra Toolkits). Other tools are responsible for managing the collaborative environment (Kernels and Session Managers). Yet others offer specific services for communication and animation (Service Applications). Tools register with the environment at startup, providing information about the kind of services that they offer (Directory), and how and where they can be contacted for those services (Location). The environment supports mechanisms to create remote instances of applications and to connect to them in client-server or peer-peer mode (Distribution). In addition, it provides facilities for different types of multi-user interaction ranging from master-slave blackboarding (Turn Taking) to synchronous multiple-user interaction (Collaboration). It implements functionality for starting and terminating collaborative sessions, and for joining or leaving them. It also supports dynamic messaging between different tools. Tools are thus built on top of the abstract Shastra layer, which is depicted in Figure 1.

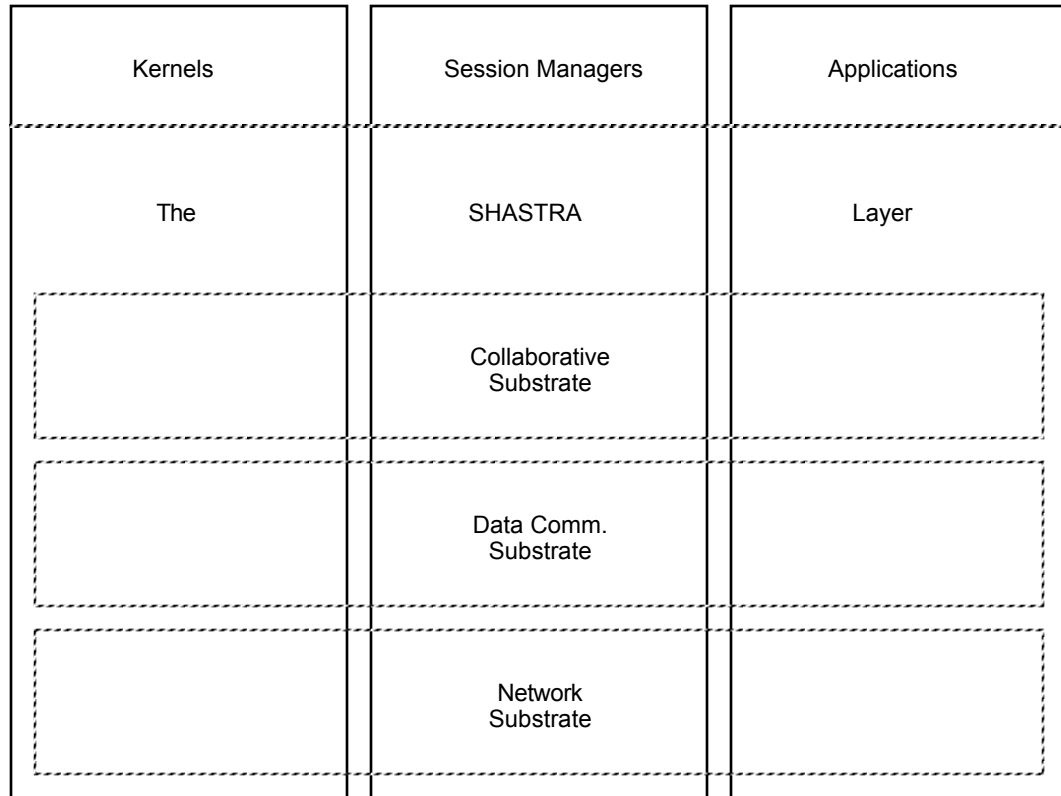


Figure 1 *The Shastra Layer: A connection, communication and collaboration management substrate. Shastra tools inter-operate using facilities provided by this layer.*

The SHASTRA collaborative scientific environment provides mechanisms to support a variety of multi-user interactions spanning the range from demonstrations and walk-throughs, to synchronous multi-user collaboration. In addition, it facilitates synchronous and asynchronous exchange of multimedia information which is useful to successfully communicate at the time of design, and to share the results of scientific tasks, and often necessary to actually solve problems. The infrastructure provides facilities to distribute the input of low computation tasks - to obtain the parallelism benefit of distribution, and the output of compute intensive tasks - to emphasize sharing of resources among applications. It provides a convenient abstraction to the application developer, shielding him from lower level details, while providing him with a rich substrate of high level mechanisms to tackle progressively larger problems.





Figure 2 (left) Distributed Volume Visualization from a large data set. **(right)** Shared Visualization. A group of researchers uses Sha-Poly to share volume visualization images of a head with cutaways (top center and right) and a cadaver (center).

A teleconferencing approach to modeling and analysis of empirical data is presented in³⁷, where the authors hypothesize about a collaborative scientific visualization environment. A discussion of an interactive visualization environment for 3D imaging is presented in³⁸, where the authors adapt the electron microscope to perform as a computer peripheral. An environment like Shastra makes it convenient to build collaborative visualization and manipulation facilities, that support resource sharing in a distributed setting.