

HEAD:

Understanding Scene Graphs

DEK

Using graph-based data structures to organize and manage scene contents

Bio

Aaron is chairman of Mantis Development, and teaches computer graphics and Internet/Web application development at Boston College. He can be contacted at aaron@mantiscorp.com.

Byline

Aaron E. Walsh

**NOTE TO AUTHORS:
PLEASE CHECK ALL URLs AND
SPELLINGS OF NAMES CAREFULLY.
THANK YOU.**

Pullquotes

Before scene-graph programming models, we usually represented scene data and behavior procedurally

Scene-graph nodes represent objects in a scene

2.1 Scene graphs are data structures used
 - to organize and manage the contents
 - of hierarchically oriented scene data.
 - Traditionally considered a high-level
 2.5 data management facility for 3D content,
 - scene graphs are becoming popular as
 - general-purpose mechanisms for manag-
 - ing a variety of media types. MPEG-4, for
 - instance, uses the Virtual Reality Model-
 2.10 ing Language (VRML) scene-graph pro-
 - gramming model for multimedia scene
 - composition, regardless of whether 3D
 - data is part of such content. In this arti-
 - cle, I'll examine what scene graphs are,
 2.15 what problems they address, and scene-
 - graph programming models supported by
 - VRML, Extensible 3D (X3D), MPEG-4, and
 - Java 3D.

2.20 Scene Composition and Management

- Scene graphs address problems that gen-
 - erally arise in scene composition and
 - management. Popularized by SGI Open
 - Inventor (the basis for VRML), scene-
 2.25 graph programming shields you from the
 - gory details of rendering, letting you fo-
 - cus on what to render rather than how to
 - render it.

- As Figure 1 illustrates, scene graphs of-
 2.30 fer a high-level alternative to low-level
 - graphics rendering APIs such as OpenGL
 - and Direct3D. In turn, they provide an ab-
 - straction layer to graphics subsystems re-
 - sponsible for processing, eventually pre-
 2.35 senting scene data to users via monitors,
 - stereoscopic goggles/glasses, projectors,
 - and the like.

- Before scene-graph programming mod-
 - els, we usually represented scene data and
 2.40 behavior procedurally. Consequently, code
 - that defined the scene was often inter-
 - spersed with code that defined the pro-
 - cedures that operated on it. The result was
 - complex and inflexible code that was dif-
 2.45 ficult to create, modify, and maintain—
 - problems that scene graphs help resolve.

- By separating the scene from the op-
 - erations performed on it, the scene-graph
 - programming model establishes a clean
 2.50 boundary between scene representation
 - and rendering. Thus, scenes can be com-
 - posed and maintained independent of rou-
 - tines that operate on them. In addition to
 - making things easier, this lets you create
 2.55 sophisticated content using visual author-
 - ing tools without regard for how work is
 - processed.

- Listing One is VRML code for a scene
 - consisting of a sphere that, when touched,
 2.60 appears yellow. As you can see, the ob-
 - jects and their behavior are represented
 - at a high level. You don't know (or care)
 - how the sphere is rendered—just that it
 - is. Nor do you know or care about how
 - the input device is handled by the un-
 - derlying run-time system to support the
 2.67 "touch" behavior. Ditto for the light.

2.68 At the scene level, you concern your-
 - self with what's in the scene and any as-
 - sociated behavior or interaction among
 - objects therein. Underlying implementa-
 - tion and rendering details are abstracted
 - out of the scene-graph programming mod-
 - el. In this case, you can assume that your
 2.75 VRML browser plug-in handles low-level
 - concerns.

- Nodes and Arcs

- As Figure 2 depicts, scene graphs consist
 2.80 of nodes (that represent objects in a scene)
 - connected by arcs (edges that define re-
 - lationships between nodes). Together,
 - nodes and arcs produce a graph structure
 - that organizes a collection of objects hi-
 2.85 erarchically, according to their spatial po-
 - sition in a scene.

- With the exception of the topmost root
 - node (which defines the entry point into
 - the scene graph), every node in a scene
 2.90 has a parent. Nodes containing other
 - nodes are parent nodes, while the nodes
 - they contain are the child nodes (children)
 - of their parent. Nodes that can contain
 - children are grouping nodes; those that
 - cannot are leaf nodes. Subgraph structures
 2.95 in Figure 2 let a specific grouping of nodes
 - exist as a discrete and independently ad-
 - dressed unit of data within the main scene-
 - graph structure. Operations on the scene
 2.100 can be performed on all nodes in the
 - graph, or they may be restricted to a par-
 - ticular subgraph (scenes can therefore be
 - composed of individual nodes as well as
 - entire subgraphs that may be attached or
 2.105 detached as needed).

- Scene graphs in Figure 2 resemble tree
 - data structures when depicted visually.
 - Not surprisingly, trees are often used for
 - scene-graph programming. The directed
 2.110 acyclic graph (DAG) data structure (also
 - known as an "oriented acyclic graph") is
 - commonly used because it supports node
 - sharing at a high level in the graph (nodes
 - in a DAG can have more than one par-
 2.115 ent) although typically at the expense of
 - additional code complexity and memory
 - consumption. In a DAG, all nodes in the
 - graph have a directed parent-child rela-
 - tionship in which no cycles are allowed—
 2.120 nodes cannot be their own parent.

- Graph Traversal

- Scene-graph nodes represent objects in a
 - scene. Scene graphs used for 3D content,
 2.125 for instance, usually support nodes that
 - represent 3D geometric primitives (prede-
 - fined boxes, cones, spheres, and so forth),
 - arbitrarily complex polygonal shapes, lights,
 - materials, audio, and more. On the other
 2.130 hand, scene-graph programming models
 - for other forms of media might support
 - nodes for audio/video content, timing and
 - synchronization, layers, media control, spe-
 2.134 cial effects, and other functionality for com-

3.1 posing multimedia.

- Scene-graph programming models support a variety of operations through traversals of the graph data structure that typically begin with the root node (root nodes are usually the entry point for scene rendering traversals). Graph traversals are required for a number of operations, including rendering activities related to transformations, clipping and culling (preventing objects that fall outside of the user's view from being rendered), lighting, and interaction operations such as collision detection and picking.

3.15 Nodes affected by a given operation are visited during a corresponding traversal. Upon visitation, a node's internal state may be set or altered (if supported) so that it reflects the state of the operation at that point in time. Rendering traversals occur almost constantly with interactive and animated graphics because the state of affairs change as often as the user's viewpoint, necessitating continual scene-graph queries and updates in response to an ever-changing perspective. To increase performance, effect caching can be used so that commonly applied operations use cached results when possible.

3.30 **Virtual Reality Modeling Language (VRML)**

- VRML is an international Standard for 3D computer graphics developed by the Web3D Consortium (formerly the VRML Consortium) and standardized by ISO/IEC. The complete specification for ISO/IEC 14772-1:1997 (VRML97) is available at <http://web3d.org/>.

3.40 An Internet and web-enabled outgrowth of Open Inventor technology developed by SGI (<http://www.sgi.com/>), VRML standardizes a DAG-based scene-graph programming model for describing interactive 3D objects and 3D worlds. Also intended to be a universal interchange format for integrated 3D graphics and multimedia, the VRML Standard defines nodes that can generally be categorized as:

- 3.50 • Geometry nodes that define the shape or form of an object.
- • Geometric property nodes used to define certain aspects of geometry nodes.
- 3.55 • Appearance nodes.
- • Grouping nodes that define a coordinate space for children nodes they may contain.
- • Light-source nodes that illuminate objects in the scene.
- 3.60 • Sensor nodes that react to environmental or user activity.
- • Interpolator nodes that define a piecewise-linear function for animation purposes.
- • Time-dependent nodes that activate and deactivate themselves at specified times.
- 3.67 • Bindable children nodes that are unique

3.68 because only one of each type can be bound, or affect the user's experience, at any instant in time.

- Every VRML node has an associated type name that defines the formal name for the node—*Box*, *Fog*, *Shape*, and so forth. Each node may contain zero or more fields that define how nodes differ from other nodes of the same type (field values are stored in the VRML file along with the nodes and encode the state of the virtual world) in addition to a set of events, if any, that the node can send or receive. When a node receives an event, it reacts accordingly by changing its state, which might trigger additional events. Nodes can change the state of objects in the scene by sending events. A node's implementation defines how it reacts to events, when it may generate and send events, and any visual or auditory appearance it might have in the scene.

- VRML supports a *Script* node that facilitates dynamic behaviors written in programming languages such as ECMAScript, JavaScript, and Java. *Script* nodes are typically used to signify a change in the scene or some form of user action, receive events from other nodes, encapsulate program modules that perform computations, or effect change elsewhere in the scene by sending events. External programmatic control over the VRML scene graph is possible via the External Authoring Interface (EAI). Currently awaiting final ISO standardization as Part 2 of the VRML97 Standard, EAI is a model and binding for the interface between VRML worlds and external environments.

- All in all, the VRML Standard defines semantics for 54 built-in nodes that implementers, such as VRML browser plug-ins, are obligated to provide. In addition, VRML's PROTO and EXTERNPROTO statements (short for "prototype" and "external prototype") offer extension mechanisms for creating custom nodes and behaviors beyond those defined by the Standard.

- VRML is a text-based language for which a variety of authoring and viewer applications and freely available browser plug-ins exist, making it popular for exploring scene-graph programming fundamentals. The file *human.wrl* (available electronically; see "Resource Center," page 5), for instance, defines the 3D humanoid in Figure 3, which is composed of primitive sphere and cylinder shapes. To view and examine this scene, open *human.wrl* file in your web browser (after installing a VRML plug-in such as Contact, <http://blaxxun.com/>, or Cortona, <http://www.parallelgraphics.com/>).

- The scene graph in *human.wrl* relies heavily on *Transform*, a grouping node

4.1 that contains one or more children. Each
 - *Transform* node has its own coordinate
 - system to position the children it contains
 - relative to the node's parent coordinate
 4.5 system (*Transform* children are typically
 - *Shape* nodes, *Group* nodes, and other
 - *Transform* nodes). The *Transform* node
 - supports transformation operations relat-
 - ed to position, scale, and size that are ap-
 4.10 plied to each of the node's children. To
 - help identify the children of each *Trans-*
 - *form* used in *human.wrl*, I've placed al-
 - phabetical comments (#a, #b, #c, and so
 - on) at the beginning/ending braces of each
 4.15 *children* field.

- As with Listing One, the nodes that
 - compose *human.wrl* are named using
 - VRML's DEF mechanism. After a node
 - name has been defined with DEF (short
 4.20 for "define"), it can then be referenced
 - elsewhere in the scene. Listing One shows
 - how USE is combined with ROUTE to fa-
 - cilitate event routing; *human.wrl* illustrates
 - how specific node instances can be reused
 4.25 via the USE statement. With Figure 3, the
 - arm segments defined for the left side of
 - the body are reused on the right. Like-
 - wise, the skin appearance DEF'ed for the
 - body is USE'd for the skull.

4.30 In addition to enabling node sharing and
 - reuse within the scene, DEF is handy for
 - sharing VRML models with other program-
 - ming environments. *Human.wrl* takes care
 - to DEF a number of nodes based on the
 4.35 naming conventions established by the
 - Web3D Consortium's Humanoid Animation
 - Working Group (H-Anim; <http://hanim.org/>). As a result, the *Human_body*,
 - *Human_r_shoulder*, *Human_r_elbow*, and
 4.40 *Human_skullbase* nodes are accessible
 - to applications that support H-Anim se-
 - mantics for these and other human-like
 - structures. VRML Viewer (VView), [http://](http://coreweb3d.com/)
 - coreweb3d.com/, does this.

4.45 Nodes are discrete building blocks used
 - to assemble arbitrarily complex scenes. If
 - you need lower-level application and plug-
 - in plumbing, check OpenVRML ([http://](http://openvml.org/)
 - openvml.org/) and FreeWRL ([http://www](http://www.crc.ca/FreeWRL/)
 4.50 [.crc.ca/FreeWRL/](http://www.crc.ca/FreeWRL/)). Both are open-source
 - implementations that add VRML support
 - to projects.

- OpenVRML and FreeWRL are open-
 - source VRML implementations hosted by
 4.55 SourceForge (<http://sourceforge.net/>). X3D
 - is the official successor to VRML that
 - promises to significantly reduce develop-
 - ment requirements while advancing state-
 - of-the-art for 3D on and off the Web.

4.60 **Extensible 3D (X3D)**

- Extensible 3D (X3D; <http://web3d.org/x3d/>)
 - enables interactive web- and broadcast-
 - based 3D content to be integrated with
 - multimedia while specifically address-
 - ing limitations and issues with the now-
 4.67 obsolete VRML Standard. X3D adds features

4.68 and capabilities beyond VRML including
 - advanced APIs, additional data-encoding
 - formats, stricter conformance, and a com-
 - ponentized architecture that enables a
 - modular approach to supporting the Stan-
 - dard (as opposed to VRML's monolithic
 - approach).

4.75 X3D is intended for use on a variety of
 - devices and application areas—engi-
 - neering and scientific visualization, multi-
 - media presentations, entertainment and
 - education, web-page enhancement, and
 4.80 shared multiuser environments. As with
 - VRML, X3D is designed as a universal in-
 - terchange format for integrated 3D graph-
 - ics and multimedia. But because X3D sup-
 - ports multiple encodings—including XML
 4.85 encoding—it should surpass VRML as a
 - 3D interchange format.

- X3D was designed as a content devel-
 - opment and deployment solution for a va-
 - riety of systems—number-crunching sci-
 4.90 entific workstations, desktop/laptop
 - computers, set-top boxes, PDAs, tablets,
 - web-enabled cell phones, and devices that
 - don't have the processing power required
 - by VRML. X3D also enables the integra-
 4.95 tion of high-performance 3D facilities into
 - broadcast and embedded devices, and is
 - the cornerstone of MPEG-4's baseline 3D
 - capabilities.

- X3D's componentized architecture en-
 - ables lightweight client players and plug-
 4.100 ins that support add-on components. X3D
 - eliminates VRML's all-or-nothing complexity
 - by breaking functionality into discrete com-
 - ponents loaded at run time. An X3D com-
 4.105 ponent is a set of related functions con-
 - sisting of various objects and services, and
 - is typically a collection of nodes, although
 - a component may also include encodings,
 - API services, or other X3D features.

4.110 The X3D Standard specifies a number
 - of components including a Core compo-
 - nent that defines the base functionality for
 - the X3D run-time system, abstract base-
 - node type, field types, event model, and
 - routing. The Core component provides the
 4.115 minimum functionality required by all X3D-
 - compliant implementations, and may be
 - supported at a variety of levels for imple-
 - mentations conformant to the X3D archi-
 - tecture, object model, and event model.

4.120 The X3D Standard defines components
 - such as Time (nodes that provide the time-
 - based functionality), Aggregation and
 - Transformation (organizing and grouping
 4.125 nodes that support hierarchy in the scene
 - graph), Geometry (visible geometry
 - nodes), Geometric Properties (nodes that
 - specify the basic properties of geometry
 - nodes), Appearance (nodes that describe
 4.130 the appearance properties of geometry
 - and the scene environment), Lighting
 - (nodes that illuminate objects in the
 - scene), and many other feature suites in-
 4.134 cluding Navigation, Interpolation, Text,

5.1 Sound, Pointing Device Sensor, Environ-
 - mental Sensor, Texturing, Prototyping, and
 - Scripting components.

- A number of proposed components are
 5.5 under consideration including nodes need-
 - ed for geometry using Non-Uniform B-
 - Splines (NURBS), for applying multiple
 - textures to geometry using multipass or
 - multistage rendering, to relate X3D worlds
 5.10 to real world locations, humanoid anima-
 - tion nodes (H-Anim), Distributed Interac-
 - tive Simulation (DIS) IEEE 1284 commu-
 - nications nodes, and more. Since it is
 - extensible, you can create your own compo-
 5.15 nents when X3D's predefined compo-
 - nents aren't sufficient.

- X3D also specifies a suite of imple-
 - mentation profiles for a range of applica-
 - tions, including an Interchange Profile for
 5.20 content exchange between authoring and
 - publishing systems; an Interactive Profile
 - that supports delivery of lightweight in-
 - teractive animations; an Extensibility Pro-
 - file that enables add-on components; and
 5.25 a VRML97 Profile that ensures interoper-
 - ability between X3D and VRML97 legacy
 - content.

- By letting scenes be constructed using
 - the Extensible Markup Language (XML),
 5.30 X3D scene graphs can be exposed via
 - markup. This lets you weave 3D content
 - into web pages and XML documents like
 - that of Scalable Vector Graphics (SVG), Syn-
 - chronized Multimedia Integration Language
 5.35 (SMIL), and other XML vocabularies.

- The file `mountains3.x3d.txt` (available
 - electronically) is an X3D scene encoded
 - in XML. In this case, the scene consists of
 - a *NavigationInfo* node that specifies phys-
 5.40 ical characteristics of the viewer's avatar
 - and viewing model, and a *Background*
 - node that specifies ground and sky tex-
 - tures, which create a panoramic backdrop
 - for the scene. Because this scene is ex-
 5.45 pressed in XML, the nodes that make up
 - this scene graph are exposed through the
 - Document Object Model (DOM), and the
 - scene graph itself may be transformed into
 - other formats as needed. In this way, XML-
 5.50 encoded X3D content is a convenient
 - mechanism by which 3D content can be
 - delivered to devices that don't yet support
 - X3D. Figure 4, for instance, shows the
 - X3D scene in `human.wrl` displayed in a
 5.55 VRML-enabled web browser. Here the
 - XML file was transformed into VRML97
 - format, letting the scene be viewed using
 - any VRML product. When the benefits of
 - XML aren't required, an alternate data-
 5.60 encoding format (say, X3D's binary and
 - VRML97 UTF-8 encodings) can be used.

- MPEG-4

- Developed by the Moving Picture Ex-
 - perts Group (MPEG; <http://web3dmedia.com/web3d-mpeg/> and <http://mpeg.telecomitalia.com/>), MPEG-4 is an

5.68 ISO/IEC Standard for delivering multi-
 - media content to any platform over any
 - network. As a global media toolkit for
 - developing multimedia applications
 - based on any combination of still im-
 - agery, audio, video, 2D, and 3D content,
 - MPEG-4 builds on VRML97 while em-
 5.75 bracing X3D. MPEG-4 uses the VRML
 - scene graph for composition purposes,
 - and introduces new nodes and features
 - not supported by the VRML Standard. In
 - addition, MPEG has adopted the X3D In-
 5.80 teractive Profile as its baseline 3D pro-
 - file for MPEG-4, thereby enabling 3D
 - content that can play across MPEG-4 and
 - X3D devices.

- Recall from my article "The MPEG-4
 5.85 Java API & MPEGlets" (*DDJ*, April 2002)
 - that MPEG-4 revolves around the concept
 - of discrete media objects composed into
 - scenes. As such, it builds on scene-graph
 - programming concepts popularized by
 5.90 VRML. MPEG-4 also introduces features
 - not supported by VRML—streaming, bi-
 - nary compression, content synchroniza-
 - tion, face/body animation, layers, intel-
 - lectual property management/protection,
 5.95 and enhanced audio/video/2D.

- MPEG-4's Binary Format for Scenes
 - (BIFS) composes and dynamically alters
 - MPEG-4 scenes. BIFS describes the spatio-
 - temporal composition of objects in a scene
 5.100 and provides this data to the presentation
 - layer of the MPEG-4 terminal. The BIFS-
 - Command protocol supports commands
 - for adding/removing scene objects and
 - changing object properties in a scene. In
 5.105 addition, the BIFS-Anim protocol offers
 - sophisticated object animation capabilities
 - by allowing animation commands to be
 - streamed directly to scene-graph nodes.

- As a binary format, BIFS content is typ-
 5.110 ically 10 to 15 times smaller in size than
 - VRML content stored in plain-text format,
 - and in some cases up to 30 times small-
 - er. (VRML can also be compressed with
 - GZip, although GZip's Lempel-Zip LZ77
 5.115 compression isn't as efficient as binary
 - compression, resulting in files around eight
 - times smaller than the uncompressed
 - VRML file.)

- In its uncompressed state, BIFS content
 5.120 resembles VRML, although nonVRML
 - nodes are often present in the BIFS scene
 - graph. Listing Two, for instance, contains
 - a snippet of the MPEG-4 uncompressed
 - (raw text) ClockLet scene presented in my
 5.125 April article. If you're familiar with VRML,
 - you'll recognize several 2D nodes not de-
 - fined by the VRML Standard. *Back-*
 - *ground2D*, *Transform2D*, and *Material2D*
 - are a few of the new nodes introduced
 5.130 by BIFS, which currently supports over
 - 100 nodes.

- In addition to new nodes, VRML pro-
 - grammers will notice the absence of the
 5.134 `#VRML V2.0 utf8` comment in the first line

6.1 of every VRML97 file. (“Utf8” comments
 - identify version and UTF-8 encoding in-
 - formation.) In MPEG-4, information like
 - this is conveyed in object descriptors (OD).
 6.5 Similar in concept to URLs, MPEG-4 ODs
 - identify and describe elementary streams
 - and associate these streams with corre-
 - sponding audio/visual scene data.

- As Figure 5 illustrates, a media object’s
 6.10 OD identifies all streams associated with
 - that object. In turn, each stream is char-
 - acterized by a set of descriptors that cap-
 - ture configuration information that can be
 - used; for instance, to determine what re-
 6.15 sources the decoder requires or the pre-
 - cision of encoded timing information.
 - Stream descriptors can also convey Qual-
 - ity of Service (QoS) hints for optimal trans-
 - mission.

6.20 MPEG-4 scene descriptions are coded
 - independently from streams related to
 - primitive media objects, during which
 - identification of various parameters be-
 - longing to the scene description are given
 6.25 special attention. In particular, care is
 - taken to differentiate parameters that im-
 - prove object coding efficiency (such as
 - video coding motion vectors) from those
 - that are used as modifiers of an object
 6.30 (such as parameters that specify the po-
 - sition of the object in the scene) so that
 - the latter may be modified without actu-
 - ally requiring decoding of the media ob-
 - jects. By placing parameters that modify
 6.35 objects into the scene description instead
 - of intermingling them with primitive me-
 - dia objects, MPEG-4 lets media be un-
 - bound from its associated behavior.

- In addition to BIFS, MPEG-4 supports
 6.40 a textual representation called Extensible
 - MPEG-4 Textual format (XMT). As an
 - XML-based textual format, XMT enhances
 - MPEG-4 content interchange while pro-
 - viding a mechanism for interoperability
 6.45 with X3D, SMIL, SVG, and other forms of
 - XML-based media.

Java 3D

- Java 3D is a collection of Java classes that
 6.50 define a high-level API for interactive 3D
 - development. As an optional package
 - (standard extension) to the base Java tech-
 - nology, Java 3D lets you construct
 - platform-independent applets/applications
 6.55 with interactive 3D graphics and sound
 - capabilities.

- Java 3D is part of Sun’s Java Media APIs
 - multimedia extensions (<http://java.sun.com/products/java-media/>). Java 3D pro-
 6.60 grams are created using classes in the
 - `javax.media.j3d`, `javax.vecmath`, and `com`
 - `.sun.j3d` packages. Java 3D’s primary
 - functionality is provided by the `javax.me-`
 - `dia.j3d` package (the core Java 3D class-
 - es), which contains more than 100 3D-
 - graphics-related classes. Alternatively, the
 6.67 `javax.vecmath` package contains a collec-

6.68 tion of vector and matrix math classes used
 - by the core Java 3D classes and Java 3D
 - programs. A variety of convenience and
 - utility classes (content loaders, scene-graph
 - assembly classes, and geometry conven-
 - ience classes) are in `com.sun.j3d`.

- Unlike scene-graph programming mod-
 6.75 els, Java 3D doesn’t define a specific 3D
 - file format. Instead, it supports run-time
 - loaders that let Java 3D programs support
 - a range of 3D file formats. Loaders cur-
 - rently exist for VRML, X3D, Wavefront
 6.80 (OBJ), AutoCAD Drawing Interchange File
 - (DXF), Caligari trueSpace (COB), Light-
 - wave Scene Format (LSF), Lightwave Ob-
 - ject Format (LOF), 3D-Studio (3DS), and
 - more. You can also create custom loaders.

6.85 Java 3D uses a DAG-based scene-graph
 - programming model similar to VRML,
 - X3D, and MPEG-4. Java 3D scene graphs
 - are more difficult to construct, however,
 - owing to the inherent complexity of Java.
 6.90 For each Java 3D scene object, transform,
 - or behavior, you must create a new ob-
 - ject instance using corresponding Java 3D
 - classes, set the fields of the instance, and
 - add it to the scene. Figure 6 shows sym-
 6.95 bols visually representing aspects of Java
 - 3D scenes in scene-graph diagrams like
 - those in Figures 7 and 8.

- Although complex, Java 3D’s program-
 - matic approach is expressive: All of the
 6.100 code necessary to represent a scene can
 - be placed in a central structure, over which
 - you have direct control. Altering Java 3D
 - node attributes and values is achieved by
 - invoking instance methods and setting
 6.105 fields.

- The Java 3D term “virtual universe” is
 - analogous to scene or world and describes
 - a 3D space populated with objects. As Fig-
 - ure 7 illustrates, Java 3D scene graphs are
 6.110 rooted to a *Locale* object, which itself is
 - attached to a *VirtualUniverse* object. Vir-
 - tual universes represent the largest possi-
 - ble unit of aggregation in Java 3D, and as
 - such can be thought of as databases. The
 6.115 *Locale* object specifies a high-resolution
 - coordinate anchor for objects in a scene;
 - objects attached to a *Locale* are positioned
 - in the scene relative to that *Locale*’s high-
 - resolution coordinates, specified using
 6.120 floating-point values.

- Together, *VirtualUniverse* and *Locale*
 - objects comprise scene-graph superstruc-
 - tures. Virtual universes can be extremely
 - large and can accommodate more than
 6.125 one *Locale* object. A single *VirtualUni-*
 - *verse* object, therefore, can act as the data-
 - base for multiple scene graphs (each *Lo-*
 - *cale* object is the parent of a unique scene
 - graph).

6.130 The Java 3D renderer is responsible for
 - traversing a Java 3D scene graph and dis-
 - playing its visible geometry in an on-
 - screen window (an applet canvas or ap-
 6.134 plication frame). In addition to drawing

7.1 visible geometry, the Java 3D renderer is
 - responsible for processing user input.
 - Unlike modeling languages such as
 - VRML, rendering APIs such as Java 3D
 7.5 typically give you complete control over
 - the rendering process and often provide
 - control over exactly when items are ren-
 - dered to screen. Java 3D supports three
 - rendering modes—immediate, retained,
 7.10 and compiled retained—which corre-
 - spond to the level of control you have
 - over the rendering process and the
 - amount of liberty Java 3D has to optimize
 - rendering. Each successive rendering
 7.15 mode gives Java 3D more freedom for op-
 - timizing program execution.
 - Java 3D lets you create customized be-
 - haviors for objects that populate a virtual
 - universe. Behaviors embed program log-
 7.20 ic into a scene graph and can be thought
 - of as the capacity of an object to change
 - in response to input or stimulus. *Behav-*
 - *ior* nodes, or objects, can be added to or
 - removed from a scene graph as needed.
 7.25 Every *Behavior* object contains a sched-
 - uling region that defines a spatial volume
 - used to enable the scheduling of the node.
 - The file HelloUniverse.java (available elec-
 - tronically) shows how a simple rotation
 7.30 *Behavior* can be applied to a cube shape
 - in Java 3D. In this case, the rotation *Be-*
 - *havior* makes the cube spin on the y-axis.
 - Java 3D supports a unique view mod-
 - el that separates the virtual world from
 7.35 the physical world users reside in. Al-
 - though more complicated than view mod-
 - els typically employed by other 3D APIs,
 - Java 3D’s approach lets programs operate
 - seamlessly across a range of viewing de-
 7.40 vices: A Java 3D program works just as
 - well when viewed on a monitor as when
 - viewed through stereoscopic video gog-
 - gles. The *ViewPlatform* object represents
 - the user’s viewpoint in the virtual world
 7.45 while the *View* object and its associated
 - components represent the physical (Fig-
 - ure 7). Java 3D provides a bridge between
 - the virtual and physical environment by
 - constructing a one-to-one mapping from
 7.50 one space to another, letting activity in
 - one space affect the other.
 - The Java 3D program HelloUniverse.java
 - (available electronically) is a slightly mod-
 - ified version of Sun’s HelloUniverse pro-
 7.55 gram. Figure 7 is a corresponding scene-
 - graph diagram. The content branch of
 - HelloUniverse consists of a *Transform-*
 - *Group* node that contains a *ColorCube*
 - shape node. A rotation *Behavior* node ani-
 7.60 mates this shape by changing the trans-
 - formation on the cube’s *TransformGroup*.
 - The content branch of this scene graph
 - is on the left side of Figure 7, while the
 - right side illustrates aspects related to view-
 - ing the scene. The SimpleUniverse con-
 - venience utility manages the view branch
 7.67 so that you don’t have to handle these de-

7.68 tails unless you want that level of control.

-
 - **Acknowledgments**
 - Thanks to my *Java 3D Jump-Start* coau-
 - thor Doug Gehringer of Sun Microsystems,
 - and Mikael Bourges-Sevenier, my coau-
 - thor for *Core Web3D* and *MPEG-4 Jump-*
 7.75 *Start*. Thanks also to Steve Wood, Michelle
 - Kim, Jeff Boston of IBM, and Alex
 - MacAulay of Envivio.

7.80 **DDJ**
(Listings begin on page xx.)

7.85
 -
 -
 -
 -
 7.90
 -
 -
 -
 -
 7.95
 -
 -
 -
 -
 7.100
 -
 -
 -
 -
 7.105
 -
 -
 -
 -
 7.110
 -
 -
 -
 -
 7.115
 -
 -
 -
 -
 7.120
 -
 -
 -
 -
 7.125
 -
 -
 -
 -
 7.130
 -
 -
 -
 -
 7.134

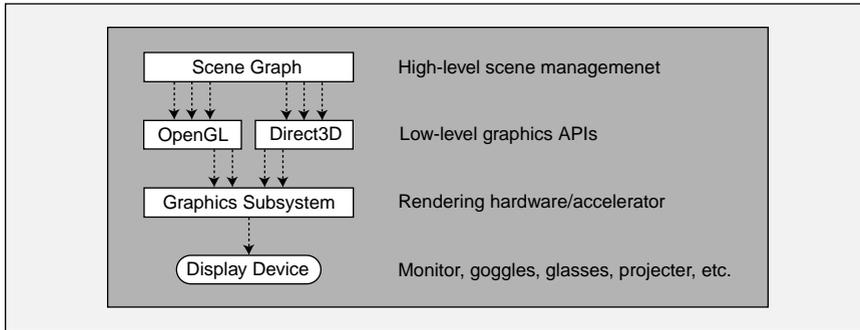


Figure 1: Scene-graph programming models shield you from underlying graphics APIs, and graphics rendering and display devices.

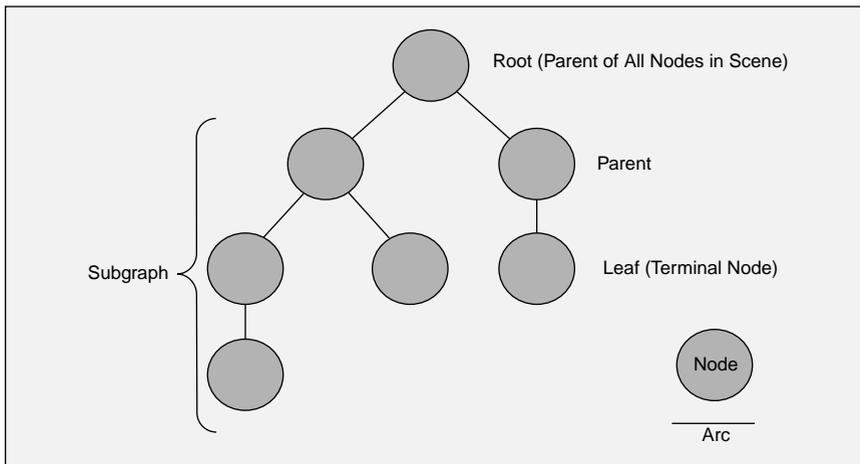


Figure 2: Scene graphs consist of nodes connected by arcs that define node relationships.

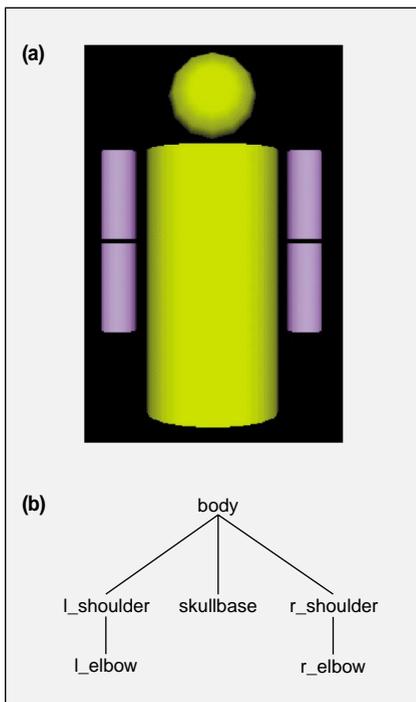


Figure 3: (a) VRML humanoid; (b) Corresponding scene-graph diagram based on human.url.

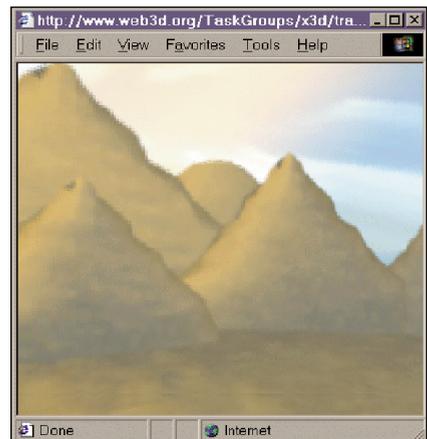


Figure 4: VRML scene generated from the X3D code in human.url. Universal Media images courtesy of Gerardo Quintieri (<http://web3dmedia.com/UniversalMedia/>).

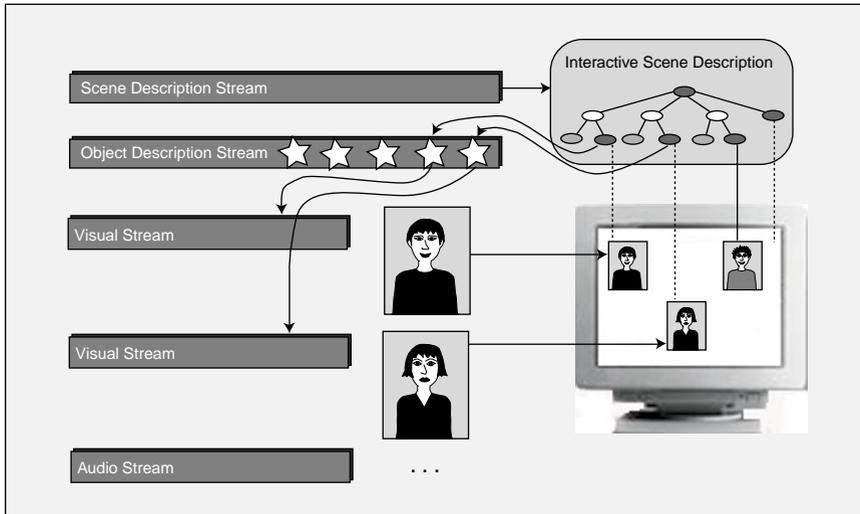


Figure 5: MPEG-4 media streams are composed at the terminal according to a scene description.

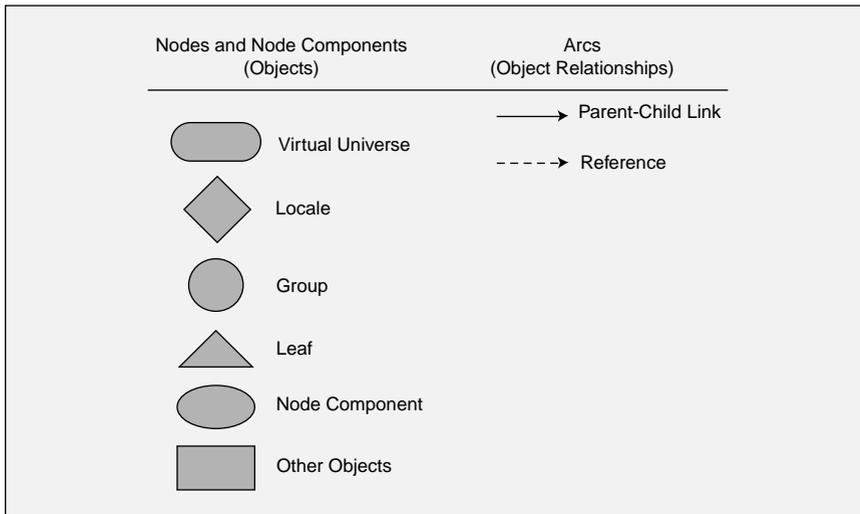


Figure 6: Symbols commonly used to visually depict Java 3D scene graphs.

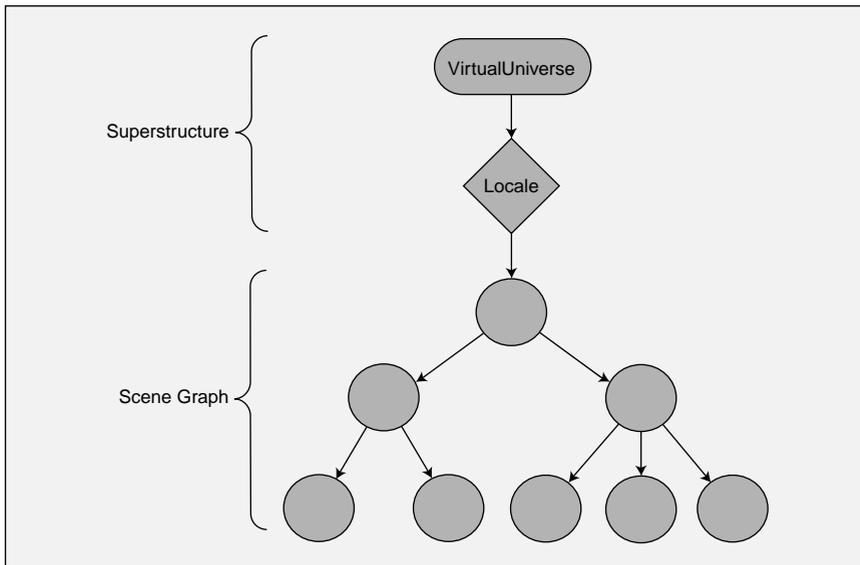


Figure 7: Java 3D scene-graph nodes are rooted to a Locale object that is in turn rooted to a Virtual Universe object.

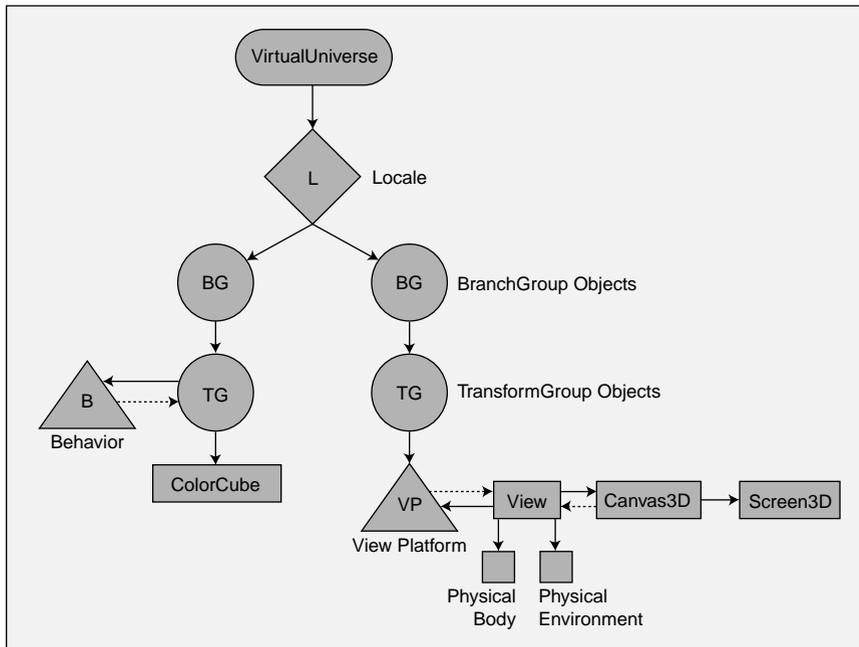


Figure 8: HelloUniverse scene-graph diagram (right branch provided by the SimpleUniverse utility).

Listing One

```
#VRML V2.0 utf8
Group {
  children [
    Shape (
      geometry Sphere {}
      appearance Appearance {material Material{}}
    )
    DEF TOUCH TouchSensor { } # define sensor
    DEF LIGHT DirectionalLight { # define light
      color 1 1 0 # R G B
      on FALSE # start with light off
    }
  ]
  ROUTE TOUCH.isOver TO LIGHT.set_on
}
```

Listing Two

```
Group {
  children [
    Background2D {
      backColor 0.4 0.4 0.4
      url []
    }
    Transform2D (
      children [
        Transform2D (
          children [
            DEF ID0 Shape {
              appearance Appearance {
                material Material2D {
                  emissiveColor 0.6 0.6 0.6
                  filled TRUE
                  transparency 0.0
                }
              geometry Rectangle {size 20.0 20.0}
            }
          ]
          center 0.0 0.0
          rotationAngle 0.0
          scale 1.0 1.0
          scaleOrientation 0.0
          translation 0.0 204.0
        )
      ]
    )
  ]
}
```

DDJ