Emerging Web3D Web Standards and Technologies

Aaron E. Walsh
Mantis Development Corps | Boston College
Boston, MA USA 02115 | Chestnut Hill, MA 02467
voice: +1-617-536-3279 | +1-617-552-3900
e-mail: aaron@mantiscorp.com | walshag@bc.edu
World Wide Web: http://mantiscorp.com/people/aew
Web3D Web: urn:web3Dweb:people:walsh:aaron:e

Abstract

The Web3D Web is an emerging distributed digital media platform designed specifically for three-dimensional (3D) content and other forms of rich, interactive media. Built using traditional Internet and World Wide Web technologies, the Web3D Web enables innovative 3D applications such as interactive movies; stereoscopic cinema; immersive multiplayer games; distributed virtual reality; synthetic environments; three-dimensional multi-user chat; telepresence and telesurgery; simulators; immersive distance learning; and other forms of high-impact content. The Web3D Web supports proprietary and open digital media data types and formats to allow rich media applications to be constructed from a wide range of technologies, including Flash, Shockwave, Virtual Reality Modeling Language (VRML), Extensible 3D (X3D), MPEG-4, and Java 3D.

Keywords: Web3D, Web, Web3D Web, 3D, three-dimensional, Extensible 3D, X3D, MPEG, MPEG-4, Java, Java 3D, Xj3D, humanoid, humanoid animation, H-Anim, computer graphics, virtual reality, VR, Virtual Reality Modeling Language, VRML, multi-user chat, simulation, telesurgery, games, media, universal media, media grid, grid.

1 Introduction

The World Wide Web began life in the late 1980s as a relatively simple hypertext system and has matured over the years to become a vital distributed multimedia system that spans the globe. As the Web itself has grown, so too has the richness and complexity of the digital media that it supports. Today many Web browsers natively support fundamental forms of media such as digital audio (e.g., WAV and MIDI), bitmap images (e.g., GIF, JPEG, and PNG) and even vector graphics (e.g., SVG and VML).

More sophisticated data formats and proprietary media types are rarely supported directly by the browser, however, and typically require a corresponding "player" that comes in the form of a plug-in, Active X control, applet, or helper application that can decode such content and allow the end user to interact with it directly within the Web browser window. Because Web browsers can be easily extended in this way an astonishing array of new digital media types have been introduced over the years that would not be viable if browser vendors were required to support such content natively.

Although the vast majority of visual media available today on the traditional Web is two-dimensional, having only the dimensions of height and width, 3D content that has the added dimension of *depth* is becoming increasingly popular. 3D for the Web, or "Web3D", is nothing new; it has simply taken a long time to become a force to recon with. When development of the Virtual Reality Modeling Language (VRML) began in 1994 to establish a three-dimensional alternative to the two-dimensional HyperText Markup Language (HTML) it ignited the imaginations of Web developers and technologists around the work, sparking a bonfire of hype that the promising new technology could not live up to [1].

The dream of a Web3D Web was dazzling, tantalizing, and out of reach. At that time the vast majority of end users did not have the computer power and network bandwidth demanded by the most compelling 3D content. To complicate matters VRML authoring tools of the day were in short supply and quite primitive as compared to the 3D authoring tools available now. Consequently the average Web user was hard pressed to find high quality VRML throughout the 1990s, relegating the technology to a niche market at best.

In 1997 VRML became the first International Standard for 3D on the Web [2] and inspired a number of proprietary Web3D technologies in the years since. Despite these early technological advances the global, universally accessible Web3D Web remained a dream until now.

2 The Web3D Web

After nearly a decade of hard work by thousands of dedicated individuals and scores of software and hardware companies the Web3D Web is finally possible. Barriers that prevented the Web3D Web from becoming a reality in the past have been steadily eliminated over the years: today's personal computers are fast enough to handle the types of rich 3D content that brought a previous generation of PCs to a grinding halt; broadband is commonplace in the office and is rapidly invading our homes; and 3D content authoring tools are powerful, plentiful, and affordable.

Over the years these once significant issues have literally become non-issues, while at the same time a core suite of 3D technology standards have emerged that together form the backbone of the Web3D Web. Thanks to these advances today's content authors can now create compelling and immersive 3D experiences that the average Web user can actually enjoy, marking a watershed moment in the history of Web3D [1].

2.1 Text is a Second-class Citizen

The Web3D Web is accessible through the traditional Web at http://Web3DWeb.com. Unlike the traditional Web, however, text is a second class citizen on the Web3D Web [3]. Multimedia dominates the Web3D Web, with interactive 3D content leading the way. Text is used sparingly and is relegated to lowly roles such as: displaying game scores; presenting "how-to" documentation that explains how to interact with or use a particular piece of content; text input in 3D chat rooms that is automatically converted into to interactive character gestures and audible speech using text-to-speech engines (Figure 1).

Figure 1: Text is a second-class citizen on the Web3D Web. In this example users take the form of animated 3D characters to chat in an immersive, three-dimensional world. Users who don't have microphones can use their keyboard to type text that is automatically converted into gestures that their character exhibits (such as waving, nodding, smiling, dancing, and so forth) and audible speech that everyone participating in the chat can hear.

2.2 Architecture

The Web3D Web is based on a distributed, multi-tiered network architecture in which application functionality is physically partitioned across various tiers of the network architecture (i.e., client, server, and business logic tiers).

The underlying architecture of the Web3D Web is transparent to end user who interacts with rich media content using a standard Web browser (such as Internet Explorer or Netscape Navigator) that is enhanced with a special Web3D browser as Figure 2 illustrates [4]. As this figure indicates, the Web3D Web delivers multimedia content in a variety of proprietary and open media types and file formats for which corresponding players are required. Users that aren't already equipped with the proper media players are lead through a simple and automated "one click" installation process that ensures the entire stack of player technology is properly installed.

The end user software stack typically consists of a Java-based Web3D Web browser and a variety of industry-standard media players. These media players are broadly classified according to their capabilities, and fall into one of two categories: *Adrenaline* players support “fast-twitch” interactive experiences such as games and virtual reality worlds, whereas *Hollywood* players deliver more traditional media such as movies and music.

Users who have especially powerful computers and adequate bandwidth may optionally become nodes on the Web3D Web network. End users who wish to become so-called “power users” must first pass a battery of automated tests designed to ensure that new nodes measure up to Web3D Web standards. These tests stress computer processing power, bandwidth availability and reliability, and security. Upon passing the power user tests an additional stack of software is installed on the user’s system that transforms their Web3D Web experience. Unlike standard end users who are subject to all of the delays and limitations of the traditional Web, power users act as peer servers on the network (see Figure 3) and as such have instant access to the most sophisticated content available on the Web3D Web [5].

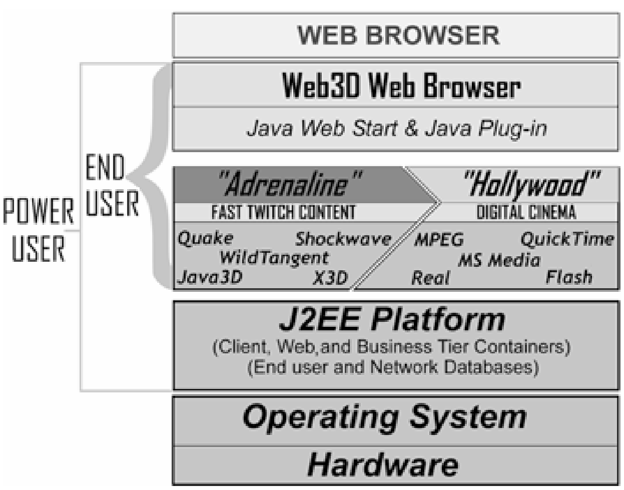

Figure 2: Web3D Web software stacks are provided for standard "end users" and more sophisticated "power users". In both cases a custom Web3D Web browser is combined with a standard suite of media players. Power users must pass a battery of performance and security tests before they can install the optional Java 2 Enterprise Edition (J2EE) platform stack that provides full support for Media Grid services. Because power users are peer nodes on the Web3D Web network they are able to receive the highest quality content at the fastest possible speeds.

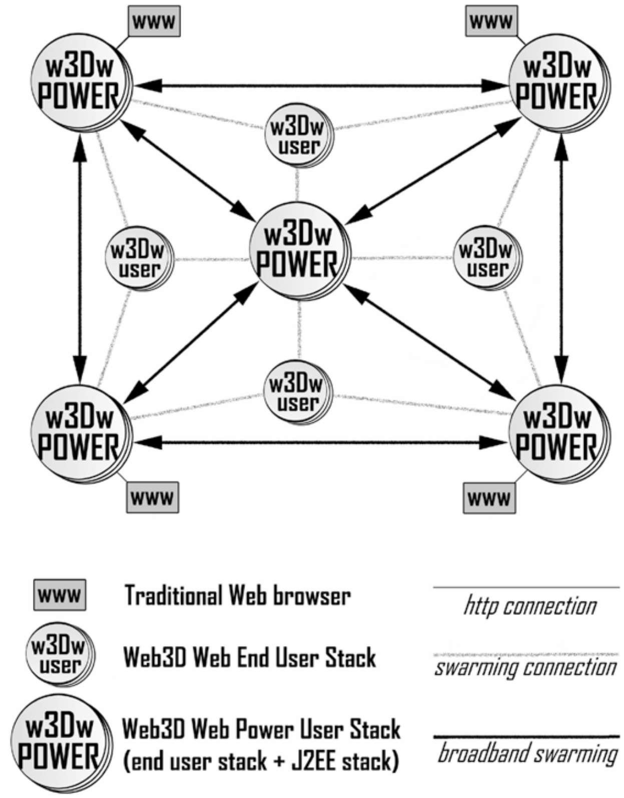

Figure 3: End users interact with the Web3D Web network using a traditional Web browser or a high performance Web3D Web browser, while power users also act as network nodes (see Figure 2). The Web3D Web browser makes it possible for digital media content to be delivered through "swarming" network connections. Swarming enables a single media file to be delivered simultaneously from numerous nodes working in parallel as illustrated by this figure.

2.3 Media Grid Enabled

Although portions of the Web3D Web are available using a standard Web browser the most sophisticated and compelling content requires special power user software that takes advantage of Media Grid services. Currently under development, the Media Grid is a decentralized network infrastructure and software development platform based on distributed computational grid technology that provides high-performance media delivery and media processing services to enable a new generation of digital media applications [6].

The Media Grid combines standard Web services with traditional distributed computing practices to permit parallel processing across decentralized networks comprised of heterogeneous devices. Built using open Internet and Web standards, the Media Grid allows massively scalable, secure, and stable digital media networks to be spontaneously created in an ad-hoc fashion or assembled from specific devices that can be administrated much like a traditional managed network. The Media Grid supports Quality of Service (QoS), broadcast capabilities, and distributed parallel processing and rendering features to create a unique software development platform designed specifically for networked applications that produce or consume massive quantities of media.

The Web3D Web will ultimately use the Media Grid to provide power users with the highest quality content possible. Development of the Media Grid will take place through the forthcoming Media Grid Center hosted by Boston College. Scheduled to open in 2004, the Media Grid Center is an interdisciplinary, cross-cutting technology and academic center through which academic and industry leaders are invited to collaborate in the development of this emerging platform [7].

3 Core Web3D Web Standards

The underlying architecture of the Web3D Web revolves around a core of 3D standards that include X3D, MPEG-4, Java 3D, H-Anim, and Universal Media. These standards make it possible to create highly portable and interoperable 3D content that is accessible to a wide spectrum of computing platforms and operating systems.

3.1 Extensible 3D (X3D)

X3D is the official successor to VRML and has been the primary focus of the Web3D Consortium's development efforts over the past five years [8]. According to the X3D specification "*X3D improves upon VRML with new features, advanced application programmer interfaces, additional data encoding formats, stricter conformance, and a componentized architecture that allows for a modular approach to supporting the standard*" [9]. At the time of this writing the X3D specification was in the final stages of ISO standardization and as such will emerge as the native 3D software standard for the Web3D Web.

X3D was designed to enable new opportunities in the creation and deployment of state-of-the-art 3D graphics on lightweight Web clients and the integration of high-performance 3D into broadcast and embedded devices. To this end X3D addresses a number of long standing issues with VRML while pushing the envelope for 3D both on and off of the traditional Web. In particular, X3D is designed to enable “3D anywhere” by employing an advanced component-based architecture that can scale across a wide range of devices and platforms. Whereas VRML is a monolithic all-or-nothing standard, X3D supports the concept of components and profiles that player and tool vendors will find infinitely more flexible and customizable by comparison.

3.1.1 X3D Components

An X3D component is a set of related functionality that consists of various X3D objects and services. Although a component is usually a collection of X3D nodes it can also include encodings, API services, or other features. X3D revolves around a Core component that defines the base functionality and the set of capabilities (such as architecture, basic abstract node types and field types) for the X3D run-time system.

In addition to the Core component the X3D standard also defines components for a variety of fundamental capabilities such as geometry, appearance, time, lighting, sound, navigation, scripting, texturing and more. X3D also defines components for more advanced features and capabilities such as Non-uniform Rational B-Splines (NURBS), geographic and geospatial functionality, Humanoid Animation (H-Anim), Distributed Interactive Simulation (DIS) and so forth. Vendors can implement any of the pre-defined standard components or define their own for private use.

3.1.2 X3D Profiles

An X3D profile is a subset of the X3D specification that meets a specific market need. More specifically, an X3D profile is a named collection of functionality and corresponding requirements that a product vendor implements to support X3D content that conforms to that profile. Profiles are further defined as a set of components and corresponding support levels, as well as the minimum support criteria for all of the objects contained within that set. The X3D specification specifies the following six profiles:

- **Core** -- defines the absolute minimal file definitions required by X3D; supports minimally defined scenes by explicitly specifying the component and levels required; enables a broad range of implementations by eliminating some of the complexity of a comprehensive X3D implementation.
- **Full** -- supports the complete set of X3D features and, as such, enables the most comprehensive and powerful content possible while remaining conformant to the specification. All other profiles are effectively sub-sets of the Full profile.
- **Immersive** -- enables immersive virtual worlds with complete navigational and environmental sensor control; analogous to the base profile defined by the VRML97 standard.
- **Interactive** -- enables lightweight players that support rich graphics and interactivity; may be implemented in low-footprint engines (such as applets and small browser plug-ins) that require limited navigation and environmental sensor control.
- **Interchange** -- enables exchange of geometry and animations between authoring systems; may be implemented in low-footprint engines; addresses limitations of software-only rendering engines that are not capable of supporting the full X3D lighting model.
- **MPEG-4 interactive** -- defines a base point of interoperability with the MPEG-4 standard; may be implemented by lightweight engines that require rich graphics and interactivity or even by low-footprint engines that require limited navigation and environmental sensor control.

In addition to these pre-defined profiles vendors may also define their own if desired. Custom profiles can be kept proprietary or they may be registered with the Web3D Consortium to make them available to the general public.

3.1.3 Open, Interoperable 3D

Like VRML, X3D is an open, royalty-free standard with a corresponding open source implementation provided free of charge by the Web3D Consortium. The Moving Picture Experts Group (MPEG) has already accepted X3D for the baseline 3D capabilities of the MPEG-4 streaming media standard, while the Web3D Consortium continues to work closely with the World Wide Web Consortium (W3C) and other standards organizations in anticipation of further adoption of X3D across the industry.

Listing 1 contains a simple X3D scene encoded in XML, the rendered results of which are shown in Figure 4. X3D content that is encoded in XML can be dynamically transformed into a variety of other 3D formats on the fly, enabling X3D to act as a master 3D format from which other formats may be derived on demand. The simple X3D scene shown in Listing 1, for example, could be transformed on the client or server into VRML, Shockwave, MPEG-4, Java 3D, or nearly any other 3D format. In this way a single body of 3D content can reach the largest possible audience.

Listing 1: X3D Scene Encoded in XML

```
<?xml version="1.0"
      encoding="UTF-8"?>
<!DOCTYPE X3D PUBLIC
 "ISO//Web3D//DTD X3D 3.0//EN"
 "http://www.web3d.org/
  specifications/x3d-3.0.dtd">
<X3D>
<Scene>
<WorldInfo
 title="Primitive Shapes"
 info='"X3D encoded in XML"'/>
<Transform DEF='dad_GROUND'>
 <Group DEF='GROUND'>
  <Transform DEF='dad_Box1'
   translation='-0.68 -0.04 0.14'>
   <Shape DEF='Box1'>
    <Appearance>
     <Material DEF='Red_mat'
      ambientIntensity='0.200'
      shininess='0.200'
```

```
      diffuseColor='1.0 0.0 0.0'/>
    </Appearance>
    <Box size='1.5 1.5 1.5'/>
   </Shape>
  </Transform>
  <Transform DEF='dad_Sphere1'
   translation='-3.26 0.039 0.02'>
   <Shape DEF='Sphere1'>
    <Appearance>
     <Material USE='Red_mat'/>
    </Appearance>
    <Sphere radius='1.000'/>
   </Shape>
  </Transform>
  <Transform DEF='dad_Cone1'
   translation='1.84 0.09 -0.16'>
   <Shape DEF='Cone1'>
    <Appearance>
     <Material USE='Red_mat'/>
    </Appearance>
    <Cone
     height='2.000'
     bottomRadius='1.000'/>
   </Shape>
  </Transform>
  <Transform DEF='dad_Cylinder1'
   translation='4.3 0.0 0.098'>
   <Shape DEF='Cylinder1'>
    <Appearance>
     <Material USE='Red_mat'/>
    </Appearance>
    <Cylinder
     height='1.750'
     radius='1.000'/>
   </Shape>
  </Transform>
  <Transform DEF='dad_Background1'
   translation='-4.917 1.392 0.0'>
   <Background DEF='Background1'
    skyAngle=''
    skyColor='1.0 1.0 1.0'
    groundAngle=''
    groundColor='0.0 0.0 0.0'/>
  </Transform>
 </Group>
</Transform>
</Scene>
</X3D>
```

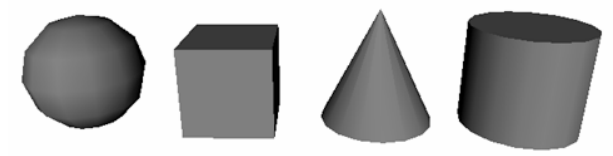

Figure 4: These primitive 3D shapes are rendered from the X3D code shown in Listing 1. X3D content that is encoded in XML can be dynamically transformed into a variety of other 3D formats on the fly, enabling X3D to act as a master 3D format from which other formats may be derived on demand.

3.2 MPEG-4

MPEG-4 is an International Standard for streaming and broadcast multimedia and was developed by the same Moving Picture Experts Group (MPEG) [10] that created the successful MPEG-1 and MPEG-2 digital audio/video standards. MPEG-4 is a comprehensive solution for encoding and delivering many different forms of digital media over a wide assortment of networks and computing platforms and as such can be viewed as a “global media solution” that supports audio, video, still images, 2D, and 3D content. These various forms of MPEG-4 content can be delivered separately or they can be combined to produce more compelling multimedia experiences [1].

MPEG-4's scene description capabilities are based on the VRML97 standard. It's possible to convert VRML content directly into the MPEG-4 format, which is how many MPEG-4 vendors currently create content to demonstrate the 3D capabilities of their products. MPEG-4 also introduces a variety of features not supported by the VRML standard, including: streaming; binary compression; advanced audio capabilities; support for 2D; integration of 2D and 3D; timing model; animation protocols that enable modification of the scene in time; and efficient compression using the Binary Format for Scenes (BIFS) layer of the MPEG-4 standard [11].

Since late 1999 MPEG and the Web3D Consortium have collaborated to develop a cross-standard profile for 3D content based on X3D. At the time of this writing the MPEG-4 Interactive profile (also known as "X3D Interactive Profile for MPEG-4") was in the final stages of standardization and promises to enable fully compatible, interoperable 3D content that can be seamlessly exchanged between MPEG-4 and X3D players and products [12].

In addition the base 3D capabilities currently defined by the MPEG-4 standard a number of new and innovative features are being developed through MPEG's Animation Framework Extension (AFX) activity. AFX adds a number of advanced 2D and 3D capabilities to MPEG-4 that make the standard suitable for high performance multimedia applications such as interactive 3D games. MPEG is also developing multi-user capabilities for MPEG-4 that will enable shared, collaborative 3D worlds systems such as the Web3D chat system shown in Figure 1.

3.3 Java 3D

Java 3D is an optional package that extends the popular Java programming language with 3D capabilities. Developed by Sun Microsystems, the Java 3D application programming interface (API) gives programmers the ability to write comprehensive, platform independent programs for the desktop and the Web. The Java 3D API is analogous to a library of function calls or procedures that developers use to write human-readable source code that is compiled into platform-neutral bytecode and executed by a Java Virtual Machine (JVM) on the end user's computer.

Java 3D is full-featured 3D graphics API that employs a scene-graph programming model similar in nature to that of VRML, X3D and MPEG-4. A Java 3D application program describes a 3D scene in terms of a scene graph structure that the Java 3D runtime is responsible for managing and rendering to the display [13].

Java 3D supports a *loader* mechanism that enables application programs to import data from a file. Loaders are available for a wide variety of 3D file formats, including VRML and X3D. Java 3D developers can also write custom loaders as needed. Because a single Java 3D application program can use any number of loaders it's possible to construct Java 3D applications that consume 3D content created by any number of authoring tools. This, in turn, makes Java 3D agnostic when it comes to file formats since application programs can access any 3D file format for which a corresponding loader exists.

The Web3D Consortium's Xj3D project leverages Java 3D to create an open source toolkit for VRML and X3D [14]. The Xj3D toolkit is used by developers to access VRML and X3D content from their own programs, and it can also be used to create complete 3D applications from the ground up. Because the Xj3D toolkit is available free of charge and is squarely focused on specification-compliant rendering of VRML and X3D you can use it to shave hundreds, if not thousands, of hours off your own Java-based 3D software development efforts.

3.4 Humanoid Animation (H-Anim)

The Web3D Consortium's Humanoid Animation Working Group [15] was originally chartered to define interchangeable humanoids and animations in VRML without extensions, and has subsequently defined the standard H-Anim component for X3D. A primary goal of the effort was to allow humanoids and animations to be authored independently. H-Anim currently supports animated limb movements, facial expressions and lip synchronization with audio sources.

Humanoid forms such as the human body, primates, video game monsters, and similar human-like structures, consist of a number of segments (such as the forearm, hand and foot) that are connected to each other by joints (such as the elbow, wrist and ankle). Because humanoid body segments are typically defined as a polygonal mesh applications that animate such humanoids must somehow alter the locations of the vertices in these meshes. Humanoid animation applications may also require information about grouping vertices in order to perform various mesh deforming operations. The H-Anim specifications define a standard way in which applications can access the joints and alter joint angles for the purposes of animation. In addition, H-Anim specifications define a way to retrieve related information such as joint limits and segment masses.

H-Anim files contain Joint nodes organized into a humanoid hierarchy. Listing 2 contains a partial hierarchy based on the current H-Anim 1.1 specification (segment names appear alongside the joints they’re attached to), the full hierarchy of which is graphically illustrated at [16]. Listing 3 shows a corresponding fragment of X3D code taken from one of the many animated humanoid characters that populate the Web3D Web.

H-Anim Joint nodes can be nested as Listing 3 illustrates, and they can also contain a Segment node that describes the body part associated with that joint. Segments can contain Site nodes that define locations relative to the segment. Site nodes, in turn, can be used to attach clothing and jewelry to the humanoid and also act as end-effectors for inverse kinematics (IK) applications. In addition, sites can also define eyepoints and viewpoint locations. Segment nodes can also contain Displacer nodes that describe which segment vertices correspond to a specific feature or to a configuration of vertices.

Finally, each H-Anim file can contain a single Humanoid node to store documentation about the humanoid, such as author and copyright information, as well as references to all the Joint, Segment and Site nodes in the file. In this sense the Humanoid node serves as a "wrapper" for the humanoid while also providing a top-level Transform for positioning the humanoid in its environment.

Listing 2: Partial H-Anim 1.1 hierarchy

```
r_sternoclavicular : r_clavicle
 r_acromioclavicular : r_scapula
  r_shoulder : r_upperarm
   r_elbow : r_forearm
    r_wrist : r_hand
     r_thumb1 : r_thumb_metacarpal
      r_thumb2 : r_thumb_proximal
       r_thumb3 : r_thumb_distal
     r_index0 :r_index_metacarpal
     r_index1 : r_index_proximal
      r_index2 : r_index_middle
       r_index3 : r_index_distal
     r_middle0 :r_middle_metacarpal
      r_middle1 : r_middle_proximal
       r_middle2 : r_middle_middle
        r_middle3 : r_middle_distal
     r_ring0 : r_ring_metacarpal
      r_ring1 : r_ring_proximal
       r_ring2 : r_ring_middle
        r_ring3 : r_ring_distal
     r_pinky0 : r_pinky_metacarpal
      r_pinky1 : r_pinky_proximal
       r_pinky2 : r_pinky_middle
        r_pinky3 : r_pinky_distal
```

Listing 3: X3D H-Anim Fragment

```
<Group
 DEF='childof_hanim_r_shoulder'>
 <Joint DEF='hanim_r_elbow'
  center='-0.19 1.13 -0.06'
  name="r_elbow">
  <Group
   DEF='childof_hanim_r_elbow'>
   <Joint DEF='hanim_r_wrist'
    center='-0.19 0.86 -0.05'
    name="r_wrist">
    <Group
     DEF='childof_hanim_r_wrist'>
     <Joint DEF='hanim_r_thumb1'
      center='-0.18 0.85 -0.04'
      name="r_thumb1">
```

3.5 Universal Media

Universal Media increases the realism of Web3D content and decreases network downloads by defining a small, cross-platform library of locally resident media elements (textures, sounds and 3D objects) and a Uniform Resource Name (URN) mechanism by which content creators can incorporate these media elements into their worlds. Universal Media allows content authors to create media-rich worlds that can be immediately loaded over even the slowest dial-up modem Internet connections; content created using Universal Media is accessible, on average, 20 to 50 times faster then it would be otherwise.

The Web3D Consortium's Universal Media Working Group [17] has already released Textures 1.0, a library of images in JPEG and PNG format for use in texture mapping 3D objects and as panoramic backgrounds. With X3D nearing ISO standardization the group will soon release 3D objects in X3D format, allowing content authors to populate their worlds and scenes with pre-made geometric shapes and forms such as humans, animals, trees and plants, vehicles, and so forth. Additionally, the group will release sound libraries that compliment the texture and 3D object libraries. All Universal Media content libraries are cross-platform, browser independent, and freely available.

Universal Media elements need to be downloaded once (by the first application that uses them). From that point on the media element is installed into a shared and universally accessible cache located on the user's hard drive or other local storage. Because Universal Media is cached in a shared area that is accessible to all applications content authors can use Universal Media just as they do locally resident resources; there is only a one-time download hit for an element, no matter how many different applications use it subsequently.

In this respect, Universal Media can be thought of as a locally resident suite of 3D media primitives, or media building blocks, that enable Web3D content developers to construct realistic worlds without the penalty of network transfers typically associated with media-rich content. And, since the media elements are of extremely high quality, Web3D content developers can build compelling worlds in considerably less time and at an overall lower cost than if they were to create such media themselves.

Conceptually speaking, Universal Media is to Web3D scenes and worlds what fonts are to plain text documents: Authors can greatly enhance their work with very little effort by using professionally developed elements designed to increase the appeal of their documents. Similar to fonts, Universal Media resides locally on end user systems, eliminating the need for media to be distributed along with the 3D world. As a result, authors can incorporate any number of textures, sounds and 3D objects into their worlds without increasing the time it takes the world to download -- only the world definition itself (such as VRML .wrl files or X3D .x3d files) is sent over the wire, while media is retrieved from the user's local hard drive.

This unique caching feature of Universal Media is accomplished by referencing media elements using a Uniform Resource Name (URN), a sibling to the better-known URL. Unlike URLs that identify resources by location, URNs are globally unique and persistent name identifiers. Applications that support Universal Media URNs search the shared cache before resorting to a network download. In cases where a download is actually required applications that support Universal Media cache the item accordingly so that all other applications can fetch it from local storage instead of requiring redundant network traffic.

4 Future Developments

The Web3D Web is in its infancy. The core standards and technologies that enable the Web3D Web have recently emerged or are still progressing through formal standardization processes. Content authors and developers can now take advantage of these technologies using a number of tools provided at [3] to create sophisticated Web3D applications. Authors and developers are encouraged to submit their works to the Web3D Web for peer review, the best of which will be made available to the global community through the traditional Web at http://Web3DWeb.com.

Although traditional Web browser access to the Web3D Web will open to the general public in early 2004 the native system reserved for power users remains under active development. The native system features special software to access Media Grid services, development of which will continue through the Media Grid Center. Academia and industry are invited to join the Media Grid Center to participate in the design and development of the Media Grid and grid-enabled applications such as the Web3D Web [7].

References

[1] Walsh, A. and Bourges-Sevenier, M., *Core Web3D*, Prentice Hall, 2000, ISBN 0-13-085728-9.

[2] ISO/IEC 14772-1:1997, The Virtual Reality Modeling Language (VRML), 1997, http://web3d.org/.

[3] The Web3D Web on the traditional World Wide Web, http://Web3DWeb.com/.

[4] Walsh, A., *J2EE 1.4 Essential*, John Wiley & Sons, 2003, ISBN 0-7645-2615-4.

[5] Web3D Web architecture and design documents, http://web3Dweb.com/architecture/.

[6] Media Grid Web site, http://MediaGrid.org/.

[7] Media Grid Center overview, http://MediaGrid.org/center/overview.pdf.

[8] Web3D Consortium Web site, http://web3d.org/.

[9] ISO/IEC FCD 19775-7:200x, Extensible 3D (X3D) International Draft Standards, 2003, http://web3d.org/.

[10] MPEG Web site, http://www.chiariglione.org/mpeg/.

[11] Walsh, A. and Bourges-Sevenier, M., *MPEG-4 Jump-Start*, Prentice Hall, 2001, ISBN 0-13-060036-9.

[12] Web3D-MPEG Working Group Web site and specifications, http://Web3DMedia.com/web3d-mpeg/.

[13] Walsh, A. and Gehringer, D., *Java 3D API Jump-Start*, Prentice Hall, 2001, ISBN 0-13-034076-6.

[14] Xj3D Web site, http://xj3d.org/.

[15] Humanoid Animation (H-Anim) Working Group Web site and specifications, http://h-anim.org/.

[16] H-Anim 1.1 skeletal sites diagram, http://h-anim.org/h-anim1_1f.gif.

[17] Universal Media Working Group Web site and specifications, http://Web3DMedia.com/UniversalMedia/.