The Media Grid: A Public Utility for Digital Media

Aaron E. Walsh
Boston College | Grid Institute
Chestnut Hill, MA 02467 | Boston, MA 02111
voice: +1-617-552-3900 | +1-617-777-0452
digital: walshag@bc.edu | http://gridinstitute.com/people/aew

Abstract

The Media Grid is a public utility that provides digital media delivery, storage and processing (compute) services for a new generation of networked applications. Built using Internet and Web standards, the Media Grid combines Quality of Service (QoS) and broadcast features with distributed parallel processing capabilities. Together these features create a unique software development platform designed specifically for networked applications that produce and consume massive quantities of digital media. As an open and extensible software development and delivery platform the Media Grid is designed to enable a wide range of applications not possible with the traditional Internet and World Wide Web. Applications enabled by the Media Grid include: Immersive Education; on-demand digital cinema and interactive movies; distributed film and movie rendering; truly immersive multiplayer games and virtual reality; real-time visualization of complex data (weather, medical, engineering, and so forth); telepresence and telemedicine (remote surgery, medical imaging, drug design, etc.); vehicle and aircraft design and simulation; and similar high-performance media applications.

1 What is a “grid”?

The term “grid” is commonly used to refer to various forms of distributed computing in which the computers, or nodes, of the system are heterogeneous and geographically dispersed. In other words, a grid can be comprised of different types of computers that may be physically located at different places (either at different places within an enterprise or organization, in different cities, or even different countries). Because the term “grid” has no formal meaning or definition, and no single definition, it’s often used to describe a wide range of distributed computer systems, including clusters, render farms, and even peer-to-peer networks.

2 Solving the Global Media Problem

The MediaGrid.org international standards group [1] is leading the design and development of the Media Grid because it is necessary, and nothing like it exists: A forthcoming generation of high-performance, media-oriented software applications and products simply cannot be built using today’s public digital networks (specifically the

Internet, World Wide Web, and public peer-to-peer networks) because of inherent technical limitations that also impact existing applications.

The long awaited “digital media revolution” has yet to fully materialize because of the combined negative effect of several interrelated factors, including: inability of client-server architectures to scale under heavy loads; slow and unreliable access to resources that reside on public networks; competing and incompatible technology standards and file formats; overly complex software development infrastructures; inadequate and incompatible digital media rights protection mechanisms; rampant illegal file sharing; incompatible electronic commerce, transaction, and auditing systems; slow adoption of Quality of Service (QoS) standards; insufficient computational power of end user devices; wide variation of display capabilities especially in terms of screen size, resolution, and bit depth; and lack of support for novel user interaction modes and GUIs.

Consequently, no public network infrastructure currently exists that provides open, secure, scalable, and reliable access to a full range of digital media services such as: instant, on-demand access to high quality content (graphics, animation, video, audio, 3D models, motion data, etc.); distributed storage mechanisms; media manipulation and presentation capabilities (e.g., dynamic scaling and resizing of content to best accommodate the end user’s display device, visualization and simulation, image sharpening and enhancement, scene rendering, titling, special effects, transformations, compositing, format conversion and encoding); advanced user interaction modes; GUIs for rich and interactive media; and access to vast computational resources such as virtualized storage and distributed parallel processing.

Simply put, the Media Grid harnesses global computing power to unleash the untapped potential of people and organizations around the world.

3 MediaGrid.org International Standards Organization

Media Grid standards are developed through the MediaGrid.org open standards organization, whose charter is to “*codify, publish and maintain Media Grid technical specifications, best practices and reference implementations to ensure fair and uniform access to the Media Grid throughout the world*”. To this end Media Grid standards, technologies, best practices and related initiatives are developed by an international collaboration of universities, colleges, research institutes, consortia and companies that includes:

- Boston College
- City of Boston
- Sun Microsystems
- New Media Consortium (NMC)
- Sweden’s Royal Institute of Technology
- Japan’s University of Aizu
- The Israeli Association of Grid Technologies (IGT)

- Singapore's Institute of High Performance Computing (IHPC)
- The James Burke Institute
- Amherst College
- Columbia University
- Media Machines
- Synthespian Studios
- National Aeronautics and Space Administration (NASA)

Visit MediaGrid.org for a complete list of members, to join, or for more information.

4 How the Media Grid Works

The Media Grid is an open digital media network infrastructure and software development platform that provides content delivery, storage and processing (compute) services for use by a wide range of networked applications. The Media Grid is powered by *service providers* (such as rendering farms, clusters, high-performance computer systems, computational grids, and similar systems) that furnish on-demand services to Media Grid *users*. User service requests are received by the Media Grid network over the public Internet and routed to appropriate service providers as illustrated by Figure 1.

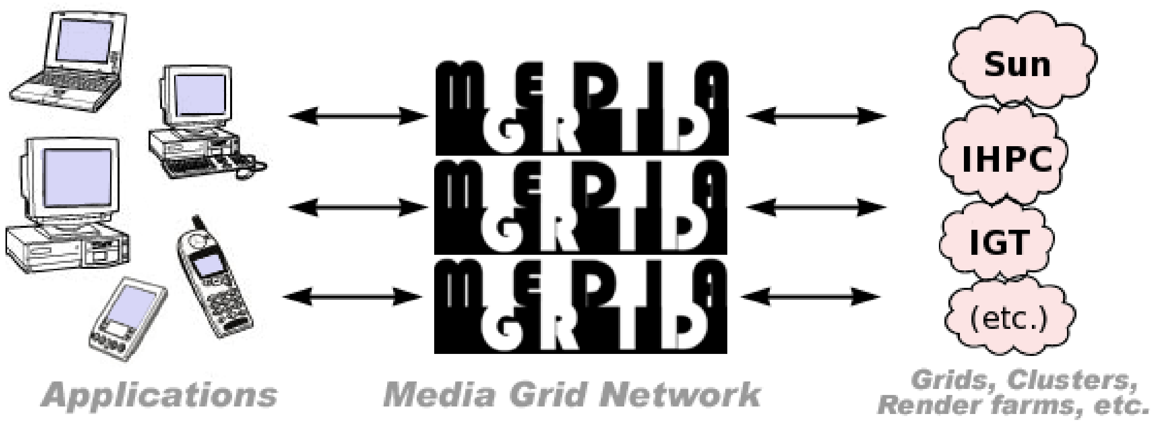

Figure 1: *The Media Grid is a global "grid of grids" that provides users with open, uniform, and simplified access to otherwise complex and proprietary rendering farms, clusters, grids, and other high-performance computing systems. Users (left) make service requests to the Media Grid (middle) which, in turn, routes the job to appropriate service providers (left).*

4.1 Benefit to Users

- **Standardized Pricing**: Transparent, unified pricing model ensures that users pay a low and standardized price for storage, delivery and processing services regardless of which vendors actually provide such services.
- **No Vendor Lock-in**: Open and unified APIs and service calls provide access to a variety of service providers without involving proprietary code. Media Grid applications are vendor neutral: storage, delivery, and processing services are

provided by a diverse suite of service providers that users don't need to know, or even care, about.
- **Simplified Access**: Simplified APIs and service call mechanisms provide user applications with open access to otherwise complex or proprietary storage, delivery, and utility computing infrastructures.

4.2 Benefit to Service Providers

- Revenue generated from services provided to the Media Grid.
- Expanded customer base without requiring sales or marketing efforts.
- Simplified public access to potentially complex or proprietary utility computing infrastructures (open, standardized access to proprietary systems).
- Tax deductions and publicity for national public service campaigns.
- Grants to pay for human and computational resources allocated to the Media Grid.

4.3 A Global "Grid of Grids"

The Media Grid does not replace or circumvent existing grids, clusters, rendering farms and similar compute services—it provides open, uniform and simplified access to them. As with the World Wide Web, which shields users and developers from the complexity of the Internet, the Media Grid provides a unified view to otherwise complex and potentially closed or proprietary systems. The Web simplifies Internet development and provides a standard browser interface for text-oriented information and basic digital media content. Similarly, the Media Grid makes it easy for developers to access services provided by utility computing vendors, rendering farms, high-performance computing systems, clusters, grids, and other service providers. In other words, the Media Grid is a "grid of grids" designed specifically to be a public utility.

By making digital media services available through standardized and unified Application Programming Interfaces (APIs), Uniform Resource Identifiers (URIs, such as URLs and URNs), Grid services, and Web services the Media Grid provides an open public utility that benefits end users, application developers and service providers.

4.4 The "Spine" and "Gateway"

The Media Grid "spine" consists of the public Internet backbone and service providers. The spine, in turn, enables any application connected to the Internet to access Media Grid services [Figure 2]. The Grid Gateway specification defined by the Grid Gateway Technology Group (GGTG) [2] details how service providers connect to, and interact with, the Media Grid network. Figure 3 illustrates at a high level the basic concept of service request processing and routing that is facilitated by a geographically dispersed network of computers that each run Grid Gateway software ("gateway nodes" on the Media Grid network).

The Grid Gateway system is similar, in concept, to the Common Gateway Interface (CGI) mechanism defined for the World Wide Web. Whereas the Web's CGI mechanism is a standard for interfacing external applications with Web (HTTP) servers, the Grid Gateway specification defines a standard for interfacing Media Grid clients and middleware with back-end grids, clusters, render farms, scientific workstations, and similar high-performance computing systems. By defining a uniform gateway interface between client-side applications and back-end service provider systems the Media Grid can be extended to support any form of content storage, delivery, or computing system.

Gateway nodes on the Media Grid network are responsible for:

1. Receiving service requests from client applications
2. Selecting an appropriate service provider(s) based on service request parameters and settings
3. Converting the open service request format into corresponding native format(s)
4. Routing the appropriately formatted service request to the service provider(s)
5. Recording the transaction (job metering, accounting, and auditing)
6. Receiving the results from the service provider
7. Routing the results to the client application

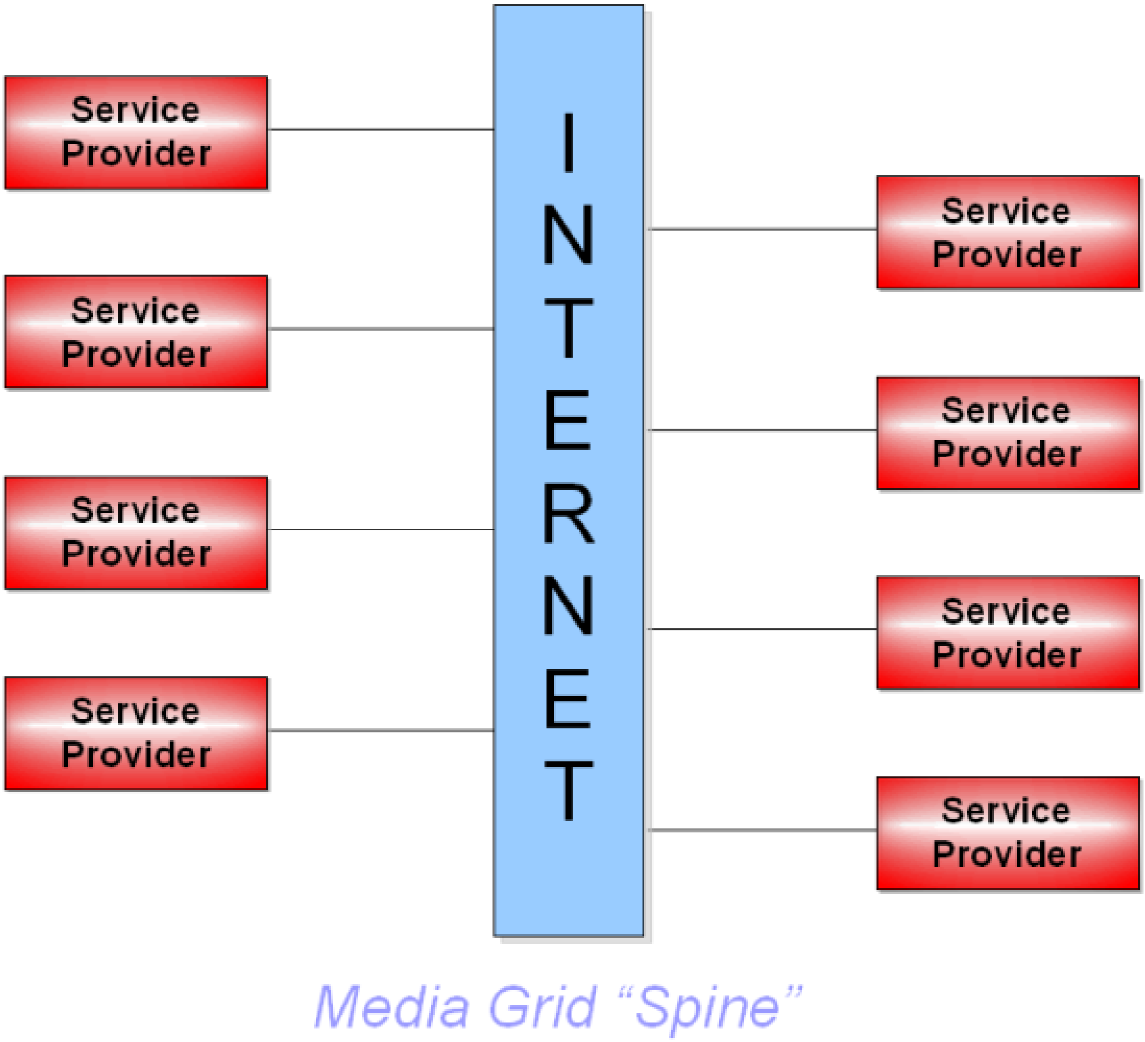

Figure 2*: The Media Grid "spine" is comprised of service provides (render farms, storage service providers, grids, clusters, etc.) connected to the public Internet backbone. Media Grid gateway nodes handle incoming service requests and route them to the appropriate service provider for processing.*

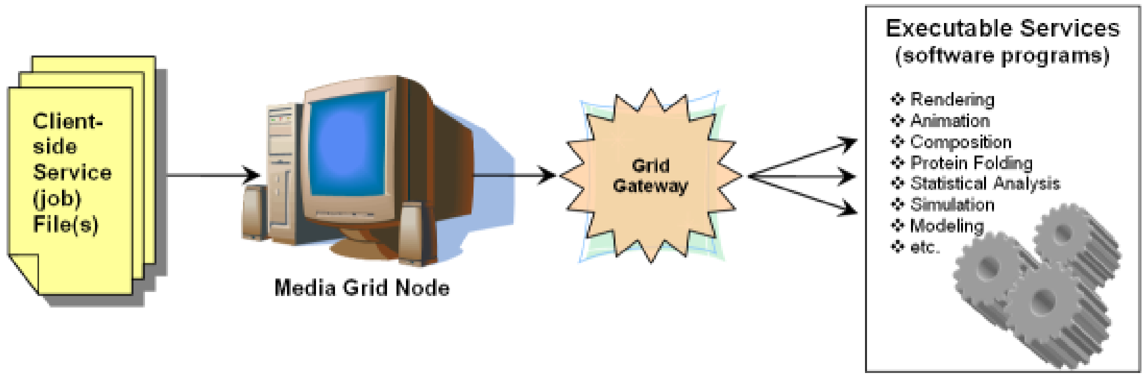

High-level flow of client-side Media Grid service requests to server-side executables that carry out (execute) the requests. Return flow not shown.

Figure 3: *Users (clients) access the Media Grid through gateway nodes that transform and route incoming job requests to back-end service providers. Users make incoming service requests using open (standardized) APIs and service calls. Gateway nodes translate these open service requests into proprietary service calls required by service providers.*

5 Secure, Pre-defined Services

The Media Grid provides on-demand access to pre-defined services. To promote secure, stable, reliable and trusted services the Grid Gateway mechanism is prohibited from accepting or processing service requests that contain arbitrary code or executable code that doesn't correspond directly to services that have been defined (specified) by the MediaGrid.org standards organization.

The Media Grid supports three general types of services:

1. Delivery
2. Storage
3. Processing (compute)

These three types of services can be considered top-level categories for which any number of sub-types (sub-categories) may be defined. All services provided by the Media Grid ultimately derive from top-level, general-purpose base types that define the minimal information necessary for service requests and responses. Service types are used to categorize and group services according to functionality as depicted by Figure 4.

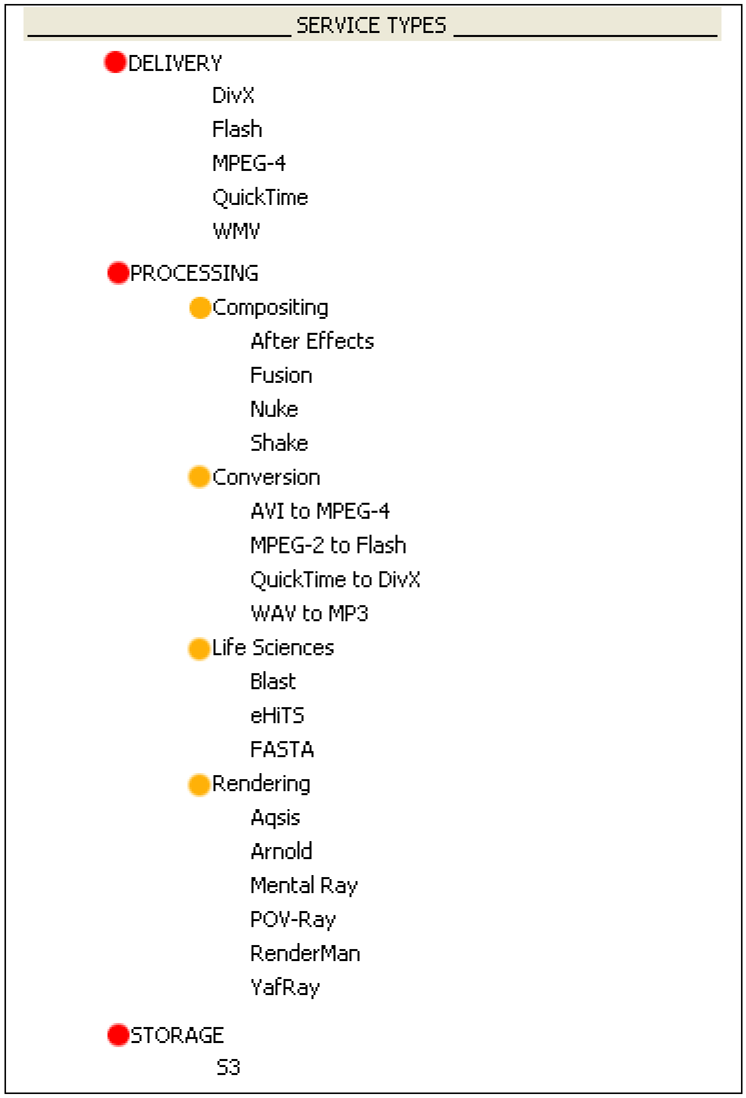

Figure 4: *The Media Grid supports three fundamental types of services (storage, delivery, and compute services) from which any number of specific digital media services can be supported. The Media Grid's RenderMan rendering service, for example, is ultimately classified as a "compute" (processing) service because it takes as input raw 3D data models and outputs fully rendered images and videos (i.e., the input data is processed to create digital media output files in the form of images and video).*

6 Examples: Rendering Services

For users located in the United States the Media Grid's rendering services are today performed by two different service providers that together furnish approximately 2,000 computers (nodes) dedicated exclusively to high quality image and video rendering. Additional rendering service providers are routinely being added to the Media Grid, enabling the system to accommodate growing usage demands. Rendering service provider's nodes are dedicated exclusively to image and video rendering; storage and content delivery services are handled by entirely different nodes.

6.1 How rendering works

In the field of computer graphics "rendering" refers to the process of generating visual output (such as an image, series of images, or video) by processing model data input. Listing 1, for example, is the raw model data that corresponds to the image shown in Figure 5. In this example the model is in the Virtual Reality Modeling Language (VRML) [3] format, and is simple enough to be rendered in real time by the user's own computer. More sophisticated models, such as those used to construct Hollywood movies and special effects, are far too complex to be rendered in real time and are typically rendered offline (not in real time) by so-called "render farms" that consist of numerous high performance computers working in parallel.

Listing 1*: VRML model of a red 3D sphere and blue 3D cube.*

```
#VRML V2.0 utf8
Transform {
  children [
    NavigationInfo { headlight FALSE }

    DirectionalLight {        # First child
        direction 0 0 -1      # Light illuminating the scene
    }

    Transform {               # Second child - a red sphere
      translation 3 0 1
      children [
        Shape {
          geometry Sphere { radius 2.3 } # define the sphere
          appearance Appearance {
            material Material { diffuseColor 1 0 0 }   # Red
         }
        }
      ]
    }

    Transform {               # Third child - a blue box
      translation -2.4 .2 1
      rotation     0 1 1  .9
      children [
        Shape {
          geometry Box {} # define the box
          appearance Appearance {
            material Material { diffuseColor 0 0 1 }  # Blue
         }
        }
      ]
    }

  ] # end of children for world
} # end of this VRML model
```

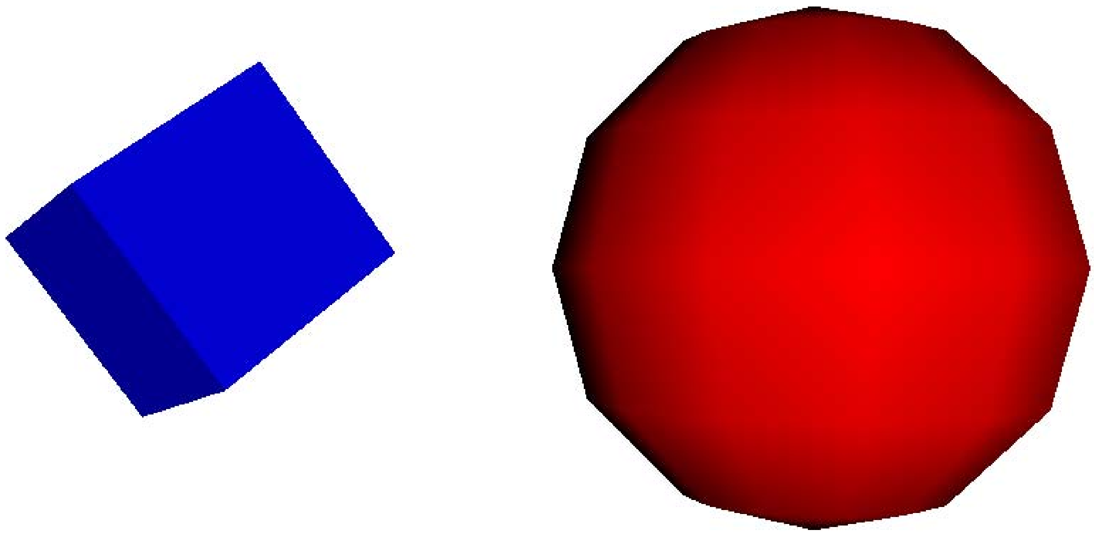

***Figure 5**: This image was created by processing ("rendering") the VRML model data shown in Listing 1. This model is simple enough to be rendered in real time on the end user's computer. In contrast, models used for Hollywood movies and television commercials are extremely complex and can take several days (if not weeks) for a single computer to render. Consequently, commercial-grade rendering is not conducted in real time and is typically offloaded to grids or clusters of high-performance computers that work in parallel (aka "render farms").*

6.2 Concept Testing Results: Benchmark Rendering Tests

Figure 6 shows the first test image ever rendered on the Media Grid. The image consists of a red 3D block and transparent 3D sphere on a flat surface. This particular test scene was fully rendered in 1.9 seconds using approximately thirty of the least powerful computers on the Media Grid (meaning only a very small amount of the Media Grid's rendering capacity was used for this scene). By comparison, the same model took 67.8 seconds to render on the end user's laptop computer. A one hour (60 minute) video or movie consisting of image frames of the same complexity would therefore be expected to take approximately 2034 hours (nearly 3 months) on the end user's computer, but only 57 hours (2.3 days) on the Media Grid.

***Table 1**: End user (client) vs. Media Grid test rendering results*

	End user (client)	**Media Grid**
Time to render 1 image	67.8 SECONDS	1.9 SECONDS
Time to render 1 hour video (est)	2034 HOURS (3 MONTHS)	57 HOURS (2.3 DAYS)
CPU (Central Processing Unit)	1 Pentium 4 M @ 1.86GHz	~ 30 Pentium III @ 3GHz
RAM	1.00 Gigabytes (GB)	~ 60 Gigabytes (GB)
GPU (Graphics Processing Unit)	ATI Radeon X300	*none*

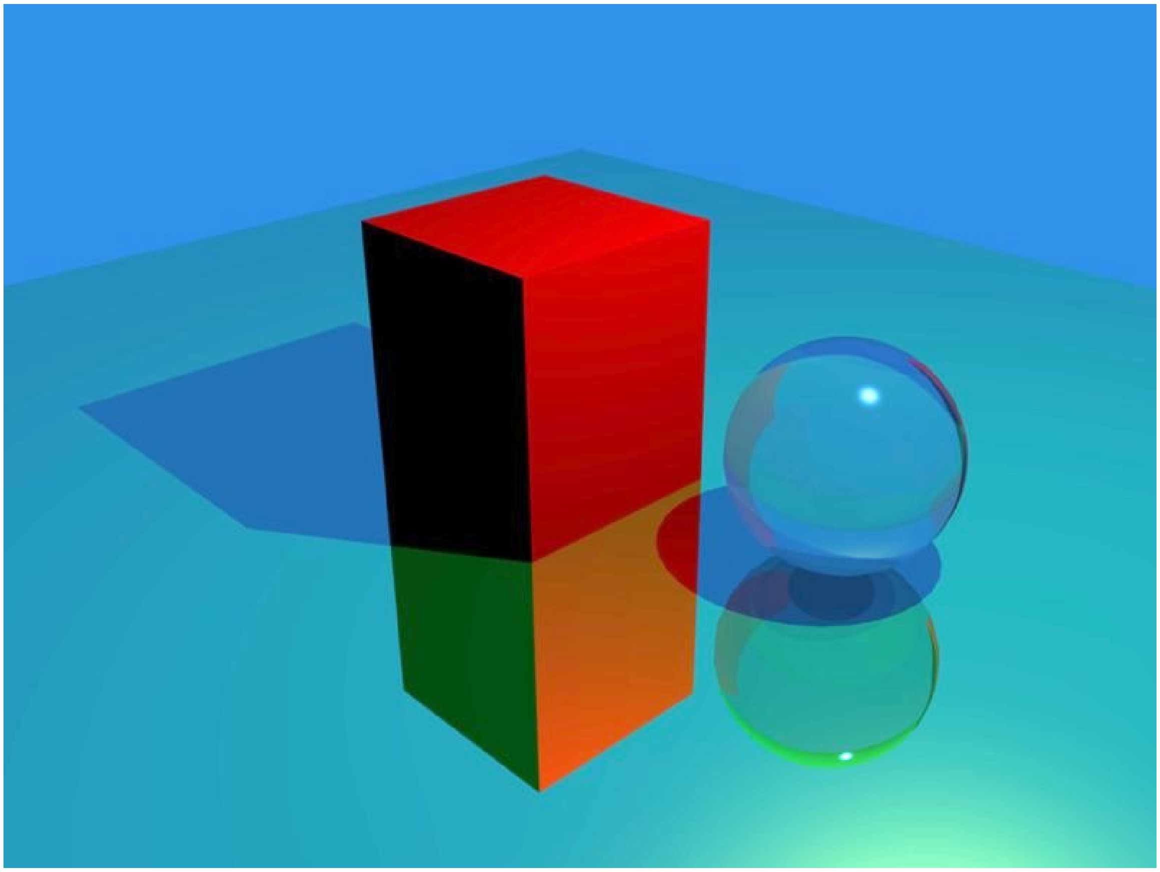

Figure 6. *The first test image rendered on the Media Grid. Image rendered in JPEG format image at 800x600 (800 pixels wide, 600 pixels high) at a resolution of 24 bits per pixel. This image took 67.8 seconds for the client's computer to render, but only 1.9 seconds on the Media Grid. One hour of similar imagery would take the client's computer approximately 3 months to render, but only 2 days on the Media Grid.*

6.3 Real World Results: Hollywood and Television Rendering

Figure 7 shows one of the first commercial (non-test) images rendered on the Media Grid. This image is a single frame from a television commercial created by Synthespian Studios [4] for Sun-Maid [5]. Synthespian Studios is a Hollywood visual effects company that specializes in digital effects for film, video and television. Synthespian Studios used the Media Grid to create the Sun-Maid "Grapes and Sunshine" campaign. The Sun-Maid girl, famous for her red bonnet and for holding a tray of freshly picked grapes, received a digital make-over for her 90th anniversary. Working closely with Sun-Maid, Synthespian Studios carefully guided the character's make-over as they designed, developed, and produced an ad campaign comprised of two pilot television commercials, several print ads and visual imagery that today appear in a variety of marketing and public relations applications including Sun-Maid's website.

On average each individual image of the Sun Maid project took between 300 and 420 minutes (between 5 and 7 hours) to render on the computers at Synthespian Studios, compared to between 10 and 30 minutes on the Media Grid. For this project thousands of images were rendered on the Media Grid, which in turn saved Synthespian Studios over 10,000 hours of computer time. In addition to saving a significant amount of rendering time the Media Grid also "freed up" the client's computers completely; Synthespian Studios simply offloaded rendering to the Media Grid and as a result their own computers were not monopolized by time consuming rendering tasks.

Figure 7*. The first commercial (non-test) image rendering on the Media Grid. This early stage concept art was a first step toward giving the classic "Sun Maid Girl" a digital make-over. Images for this project were rendered in Maya IFF format at 1080x729 (1080 pixels wide, 729 pixels high) at a resolution of 24 bits per pixel. Image copyright © Sun-Maid Growers of California. All rights reserved.*

7 Media Grid Initiatives

The MediaGrid.org standards group is actively applying Media Grid technologies to specific problem spaces, such as distance learning, digital libraries, and the impact of digital media on culture and society. The Media Grid's Immersive Education initiative

[6], for example, is an international collaboration of universities, colleges, research institutes, consortia and companies that are working together to define and develop open standards, best practices, platforms, and communities of support for virtual reality and game-based learning and training systems.

7.1 Immersive Education Initiative

Immersive Education is an award-winning learning platform that combines interactive 3D graphics, commercial game and simulation technology, virtual reality, voice chat (Voice over IP/VoIP), Web cameras (webcams) and rich digital media with collaborative online course environments and classrooms. Immersive Education gives participants a sense of "being there" even when attending a class or training session in person isn't possible, practical, or desirable, which in turn provides educators and students with the ability to connect and communicate in a way that greatly enhances the learning experience. Originally available only to university students, the next generation of Immersive Education is focused on a broad spectrum of academic and non-academic users (higher education, K-12 [kindergarten through high school], and corporate training).

Immersive Education is a merit-based initiative of the MediaGrid.org open standards group. Membership is open to the public but restricted to organizations and individuals that have experience using virtual reality, virtual worlds or game-based learning technology. To join, or for more information, visit MediaGrid.org

Unlike traditional computer-based learning systems, Immersive Education is designed to immerse and engage students in the same way that today's best video games grab and keep the attention of players. Immersive Education supports self-directed learning as well as collaborative group-based learning environments that can be delivered over the Internet or using fixed-media such as CD-ROM and DVD. Shorter mini-games and interactive lessons can be injected into larger bodies of course material to further heighten and enrich the Immersive Education experience.

Since 2004 students at Boston College have had the opportunity to participate in courses conducted entirely within Immersive Education, which has the potential to fundamentally reshape education by providing on-demand learning and simulation technology that can engage and instruct at a level far beyond that of the typical in-person or online course. Following an award-winning 2 year pilot, Boston College (in cooperation with the Grid Institute, Media Grid, Burke Institute for Innovation in Education, Media Machines, City of Boston, Sun Microsystems and other organizations) is now preparing to make Immersive Education available as a community resource for the benefit of educators, students, and researchers.

Figure 8. *The next generation (3rd generation) Immersive Education platform supports high resolution avatars and virtual learning environments.*

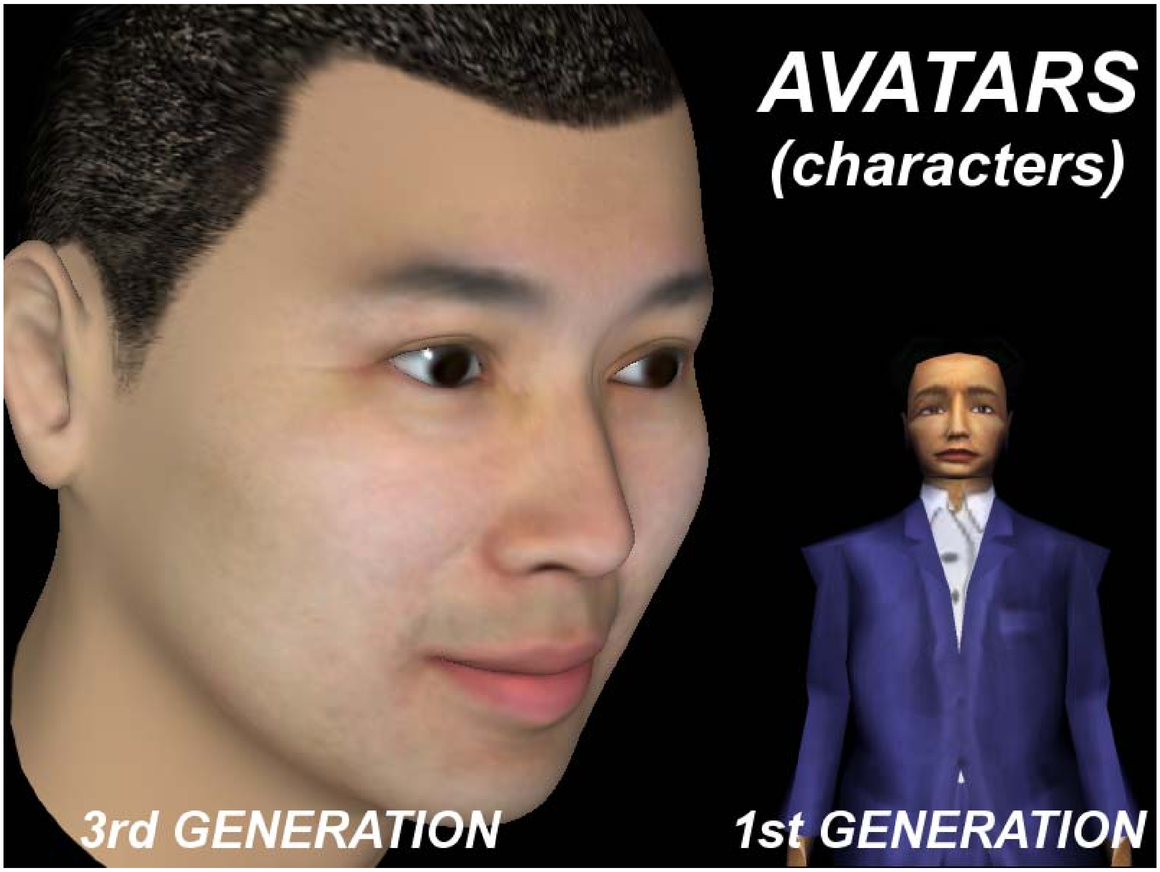

Figure 9. *The first generation of Immersive Education supported relatively low quality avatars (right), whereas the next generation supports high quality avatars (left) that are realistic and lifelike by comparison.*

References

[1] MediaGrid.org international standards organization Web site, http://MediaGrid.org/.
[2] Grid Gateway Technology Group (GGTG) Web site, http://MediaGrid.org/groups/.
[3] ISO/IEC 14772-1:1997, The Virtual Reality Modeling Language (VRML), 1997, http://web3d.org/.
[4] Synthespian Studios Web site, http://synthespians.net/.
[5] Sun-Maid Web site, http://www.sunmaid.com/.
[6] Immersive Education initiative Web site, http://ImmersiveEducation.org/.