VRML: Prelude and Future

Organizer:
Don Brutzman, Naval Postgraduate School
Panelists:
Mark Pesce, Author, VRML List Moderator
Gavin Bell, Silicon Graphics Inc., VRML 2.0 Co-Architect
Andries van Dam, Brown University
Salim AbiEzzi, Microsoft ActiveX Animation

Don Brutzman
`brutzman@nps.navy.mil`

The meteoric development of the Virtual Reality Modeling Language (VRML) is one of the most important standards developments in graphics and internetworking. An informal, intense, open, and collaborative design process has worked. VRML version one defined a concise and workable scene description language. Version two refined this language to incorporate behavioral animation mechanisms compatible with the current World Wide Web. Networked interoperable interactive 3D graphics are now feasible for everyone's computer. We examine lessons learned over the past two years, where VRML is going, and how VRML is triggering fundamental changes in the economics, mindset, and membership of the graphics community.

Although networking has been considered "different" than computer graphics, network considerations are integral to large-scale interactive 3D graphics. Graphics and networks are now two interlocking halves of a greater whole: distributed virtual environments. New capabilities, new applications, and new ideas abound in this rich intersection. Our ultimate goal is to use networked interactive 3D graphics to take full advantage of all computation, content, and people resources available on the Internet.

Realizing the lofty ambitions of VRML has required hands-on attention to myriad technical and people challenges. How do you specify a 3D scene both concisely and compatibly, given a plethora of other formats? How are working (i.e. successful) relationships built among individuals, academia, and companies of all sizes? What is a "networked behavior?" How do we sustainably capture both the specification and the standards process? Does VRML = graphics + the Web + networking + behaviors + everything in the world? In other words, where (or does) VRML stop? Which steps are next?

Graphics, networking, and interoperability breakthroughs repeatedly remove bottlenecks and provide new opportunities. A pattern appears as we attempt to scale up in capability and capacity without limit: every old bottleneck broken reveals another. Understanding bottlenecks, corresponding solutions, and potential upper bounds to growth permits us to develop effective networked graphics. Technically and socially, SIGGRAPH has crossed a threshold in capability. As we overcome current bottlenecks, "effectively networked graphics" will simply mean "applications."

Mark Pesce
`mpesce@netcom.com`

Perhaps the most singular aspects of the VRML "movement" are its out-of-control nature and its persistent strength in the face of well-organized opposition. From the very beginning–when Tony Parisi and I began to share our work with Tim Berners-Lee and others–we practiced a politic of inclusion, keeping the door open to relationships which could be leveraged into successes for VRML. This speaks more of a process of social engineering than software engineering, and articulates the heart of the difference between VRML and any of its potential competitors (ActiveVRML, OpenFlight, WIRL, etc.). Because VRML has remained open in deed as well as word, because anyone can become a member of the community and contribute, VRML has garnered the support of communities across the graphics and networking industries. In fact, despite the persistent lobbying of many large companies—including Netscape, Silicon Graphics, and Microsoft—the consensus process which brought us both VRML 1.0 and VRML 2.0 has proved resistant to tampering by press release or marketing hype. That's one of the real lessons of VRML: the social fabric of the VRML community is the real key to its success. It's my belief that this lesson has wide application outside the limited domain of VRML; wherever virtual communities are to spring up and flourish, the same conditions must apply.

Gavin Bell
`gavin@sgi.com`

VRML 2.0 is an incredibly rich file format for creating interactive 3Dmultimedia experiences that are distributed across the Internet. It isalso a solid foundation for solutions to the next hard problem: multi-userworlds.

VRML is a success because it hasn't tried to solve all of the problems of computer graphics, simulation, and networking all at once. When creating something to meet the needs of a very large group of people, it is difficult to balance the limitless number of features requested against the limited amount of design and implementation time available. Rough agreement on both constraints and goals is the key to getting anything accomplished.

Its design is both solid and practical, in part because the Internet gives system designers an invaluable tool: direct feedback from knowledgeable users. We had to convince users that our proposal would solve their problems or they would take their business elsewhere (i.e. vote for another proposal). Giving concrete answers to the stream of "can I do this..." questions ensured that we were solving relevant problems and constantly testing the design.

What's next? Tackling the multi-user problem will first require agreement on exactly which problem should be solved: multi-user "chat" is a mucheasier problem than general multi-user collaboration in a shared virtual world. Solving the more general problem will require additions to both VRML and the infrastructure of the World Wide Web.

Andries van Dam
`avd@cs.brown.edu`

One of the cliches about standards is that standards bodies produce camels, horses designed by committee. (This cynical comment ignores the clear utility of camels in the desert environment.) Whether they are de facto industry standards promulgated by a leading company or industrial consortium, or de jure official standards promulgated by standards bodies such as the IEEE, the ANSI, and ISO, standards reflect both the strengths and the weaknesses of a technopolitical consensus process that favors compromise.

Recently new standards processes specific to the Internet have emerged, aiming for both greater speed and greater democratic input, e.g. the Internet Engineering Task Force (IETF). The design and adoption of VRML 1.0 and more recently of VRML 2.0 via the VRML list and the ad hoc, self-selected VRML Architecture Group (VAG) is a prime example of an even more rapid process. Both the process and the specification should be of great interest to the graphics community. In particular, I believe that this successful process should force ANSI/ISO to redesign its heavy-weight, overly lengthy standards process. They should now consider a light-weight, fast-track process to review both VRML 2.0 and its legitimate competitors for a net-based multimedia standard. This standard should support various visions of cyberspace: multiple participants, distributed virtual environments that contain autonomous objects whose behavioral interactions with participants and each other must be simulated in real time, etc.

I hope that the pressure of various companies to simply ratify the evolving VRML 2.0 spec (or its competitors) will be resisted, and that an open, technically sound, extensible standard will be designed to last us well into the next century. VRML 2.0 should certainly be considered as a baseline for the future standard. Needless to say, if the standards process is not ultra-fast (12-18 months), VRML 2.0 WILL become the de facto standard, with all the advantages and disadvantages that implies. Then the official standard, if it is not VRML 2.0 (or any other solution) will only displace VRML 2.0 if is demonstrably superior, using the Darwinian Web-virus competition model that many Web enthusiasts believe is the dominant force for change on the Web. In summary, I believe that it is important to reexamine the good work the VRML community has done and to take into account other models, and to do that with all deliberate speed.

Salim AbiEzzi
salimabi@microsoft.com

Graphics and multimedia for the Internet are needed for both interactive illustrations embedded in Web pages, and for distributed immersive shared spaces. On the one hand, the majority of Internet users can benefit greatly from added life to Web pages through interactive illustrations, which could be used for advertising, information, and artistic and entertainment purposes. On the other hand, advanced users are interested in shared spaces and their potential for spectacular applications in group entertainment, collaborative engineering, and other far-reaching experiential applications of the Internet.

The VRML community has been focused on 3D spaces and models primarily; for example, the present VRML 2.0 is not designed for sprite animation, hot spots, and synthetic audio, which could be building blocks as essential for interactive illustrations as 3D might be. Furthermore, in order for large 3D spaces to become viable for the majority of Internet users, much faster 3D texturing, higher-bandwidth networking, and more reliable low-latency communication are required It is going to be quite some time before these capabilities become common aspects of the infrastructure.

As VRML 2.0 is considered for standardization by formal organizations, it is important to carefully consider broader Internet needs in the area of graphics and multimedia. I believe that either VRML 2.0 needs to be extended to better support the broader needs, especially in the area of illustrations, or the situation calls for other, possibly complementary, standards.

ActiveX Animation (formerly ActiveVRML) is a product out of Microsoft. It includes a modeling language and a run-time environment that primarily targets Web-interactive illustrations. It provides a novel approach for modeling rich behaviors, media integration, interaction, and events. ActiveX Animation facilitates multimedia coordination and rich media composition. It is suitable for sprite animation, just as it is for 3D and the interplay between the two. I will contrast it with VRML 2.0 and highlight its value for Web-interactive illustrations.

This panel statement can be found online at:
http://www.stl.nps.navy.mil/~brutzman/vrml/siggraph96panel.html