

Real Time Graphics' Interview

Don Brutzman, Associate Professor, Naval Postgraduate School

Don Brutzman graduated from the Naval Academy in 1978 and served as a naval officer. He went to the Naval Postgraduate School in Monterey, CA in 1990 to study Computer Science. He received his Ph.D. there in 1995, and has since remained on the faculty. He is now an Associate Professor, and his current duties include technical director for Network Virtual Environments and 3-D. He works in the Modeling, Virtual Environments, and Simulation (MOVES) Institute where his current research and development efforts include chairing the Extensible 3D (X3D) graphics specification, implementing the Distributed Interactive Simulation (DIS) protocol in the DIS-Java-VRML working group, and leads the Extensible Modeling & Simulation Framework (XMSF) effort. He also currently serves as a member of the Board of Directors of the Web3D Consortium, the subject of the present interview.

Real Time Graphics Editor Roy Latham conducted the interview on December 10th.

Real Time Graphics [RTG]: What are the objectives of the Web3D Consortium?

Don Brutzman [DB]: The Web3D Consortium is a nonprofit organization dedicated to solving the problem of open 3-D standards for the Web. Typically, we're promoting 3-D interoperability. We think that there's a huge untapped market out there for 3-D based on a lot of fragmented piecemeal implementation of 3-D that are not really available for most people.

What we do is provide an open forum or a private forum as appropriate for companies, for standards organizations, for universities, for individuals, for all players to come together and be able to work on requirements, goals, and end results that achieve this broad goal to get 3-D graphics out on the Web.

RTG: I'm sure there's more than one player, but what are you finding to be the most interested parties?

DB: That's an interesting question. Our member list continues to grow. We have several classes of members. Charter members are organizations that pay extra to help support Web3D activities and to stress their commitment to 3-D graphics on the Web. These include 3Dlabs, Sun Microsystems, and SGDL, a company that specializes in parametric surface representations technology. Military charter members include STRICOM and the Naval Postgraduate School. Recently Intel joined as a charter member, and that has brought in many other companies with interests in CAD.

"The Web3D Consortium is a nonprofit organization dedicated to solving the problem of open 3-D standards for the Web."

We have several dozen members that are backed almost exclusively by corporations. The academic members typically have fewer funds, but often provide more time to support the work. We have over 100 professional and student members, individuals who wish to declare their affiliation and support for this effort.

RTG: When you read down your list of charter members of companies, it sounds like those are people who would like to sell the technology. In other words, they hope that it will become useful so that they can sell software or hardware or something. The two government sponsors sound like people who are interested in using it for training?

DB: Yes, training I think is the biggest motivator for the military, because there is so much complicated gear, and when it comes to understanding how the world works and how things move around, 3-D really helps a lot. I think with the military you also see an untapped potential of what they call "command and control," including the ability to visualize the battle space.

RTG: Situational awareness?

DB: Correct, and now we're not just talking about who's where, what are the lines on the map, where are the good guys, and where are the bad guys. We also visualize how the sensors react with the real world and how to deconflict: keeping people from stepping on each other's toes or getting in trouble. Although this might seem a bit specialized, I think since 9/11 more and more people are aware of how challenging and how important this capability is.

Certainly NPS has a special role there. Sometimes we wear our university hats in teaching and learning, but we also wear our military hats saying what is needed by our forces out there defending the free world today.

So, I think if you keep going down the Web3D member list and looking at interests, it's not just one company. I have trouble thinking of a single member who would not benefit financially from 3-D graphics being much more accessible than it is today.

The technology is not an end in itself. It's an enabler for sophisticated capabilities, and it enables the ability to use resources better and to offer new services.

RTG: It's interesting - the early effort, at least in my perception, was aimed at games and advertising and that's notably not a driving force presently, is that true?

DB: I think that's a good statement. Games have proven to have definite life cycles of their own; they're always tightly tied to proprietary technologies. The idea of advertising on the Web became

prominent during the Internet boom with a lot of people expecting that it would generate major revenue streams. That hasn't happened with or without 3-D graphics. Advertising is an existing market or channel, it's certainly important, but it's nowhere near the dotcom projections that were going on.

It's interesting that when we've done market studies of where the needs of 3-D are, a lot of people go into it looking for a single killer application or two, hoping to find targets to hang their hats on. It turns out that our list is always about 8 or 10 or 12 different areas, all of which are served by 3-D, all of which have overlapping and occasionally unique requirements. This is why Web3D I think has a strong rationale to be there and continues to do important things.

It will prove effective as the third-generation X3D open standards supports the next-generation applications, large-scale simulations, and a good mapping to the Web. We'll continue to see the importance of having this open forum.

RTG: At least in my perception, there was a short list of problems that kept VRML from taking off and I'd like to get your take on those problems. One seeming problem was the very high cost of building 3-D content. It ended up taking a team of five people, and you had to have sound and design and so forth and so on. Is that still a limiting factor in the use of 3-D on the Web?

DB: I think it's a very important issue, and it's a little bit of a chicken or egg situation, so let me explain. If you look at a company that has some specialty 3-D technology, perhaps it's a game company, or perhaps it's some other focus. Certainly a part of their strategy every time is to write authoring tools so they can design levels, or so that people can design their house or something. The need for effective authoring is well recognized. It's time consuming and that has been one of our barriers to achieving widespread, easily authored 3-D content.

It takes money to put these tools together and usually that level of money requires a clear business plan before somebody can say, "Ah - I'm going to do it."

RTG: But assembling the tools may be the easy part of the job. It's the labor-intensive process of actually using the tools, isn't it?

DB: If we enumerated the challenges of this 3-D chicken or egg, the flip side is summarized by the same statement. It's hard; it's expensive to author and re-author 3-D. So I think it's fundamentally important that we have an open standard so that we can continue to have this pretty diverse ecology, if you will, of varied authoring tools. Each one can do what they're good at, yet each could have a way to map their products to the Web with interchange standards.

Specific to X3D at least, we did add authoring tools to our deliverables. It was not on our original list, but we have one now. We have an open source, simple editing tool called X3D-Edit. It's not a three-space authoring environment, but can be used for native authoring of 3-D content. Actually that's turned out

Authoring tools benefit from open standards by allowing diverse approaches to the tools while providing a common way to present the output.

pretty well because I think it's a clear exemplar for tool builders on how X3D works.

So if companies want to build a 3-D tool they can. Frankly, 3-D authoring is a tricky business. Authoring is not an automatic skill, and certainly 3-D authoring requires specialized tools. It's clearly important that we have a simple software exemplar out there, and we think that interchange is important for all of these different tools. We still have hope that the product they make can be re-used, so that maybe we can reduce the cost of re-authoring.

RTG: The flip-side of that is that if someone spends a tremendous amount of money on content, they want to protect it and

keep it from being used by someone else. Has that problem been addressed as to how you lock content?

DB: It's been quantified and we actually have a partial solution available. Let's back up a little - how do people usually do this? Do they usually try to put some encryption algorithm on it or do they simply obfuscate it to make it hard to de-bug or decipher. Security via obscurity is not very strong.

What we have in X3D is a call for requirements for binary format that recognizes that there are multiple aspects to the problem. You want to be able to encrypt content so that people can selectively restrict others from touching it. You also want to be able to watermark it or authenticate it so that you can undeniably find in the content who wrote it, and what the rights and what the permissions are for it.

I should note that there is a metadata node in VRML 97 called "World Info." It is very simplistic and easily subverted by anyone who wanted to get around it.

Usually the issues of 3-D authentication and encryption are usually intimately tied together with the idea of compression, both in binary terms for thrashing out white space and redundant bits. But also in terms of geometric data, we can simplify triangles, quantize colors, make tables of normals and so forth for both compression and encryption.

When you add this up, the shortest list which we could make has nine different capabilities - including different ways of compressing, authenticating, and encrypting.

What we've done is posed it as an interdependent set of requirements. The principals involved in the X3D spec development all agree that they can fulfill the requirements in combination, and that they all probably cannot be provided independently.

Finally, we've intentionally deferred reopening this work until we've locked down and put the final lid on the X3D specification. So this gives you a little look ahead at what I think will happen. We'll have the full spec done in March.

Continued on page 8

Interview: Don Brutzman

Continued from page 7

RTG: I think most people are familiar with the Virtual Reality Modeling Language (VRML). How is X3D distinguished from VRML 97?

DB: There are a lot of requirements for X3D. First among them is backward compatibility with existing VRML 97 contents. This is motivated by that same issue that it's being difficult and expensive to author 3-D content, so we don't want to lose a single polygon or a single scene. We also had some go-forward requirements for compatibility with other modern Web languages, which means XML, the Extensible Mark-up Language. XML is worth looking into because it is very cool and because it is very powerful.

At first XML is often confused a little bit. People tend to think of XML as a language in it's own right. That's the name of it, so that's a reasonable expectation, but actually it's a language for writing other languages. So what we've seen is HTML, the Web page language grow up into XHTML which is the XML form of the same language, but stricter, more rigorous, and it has more rules about how to do things. There are other examples of XML defined languages like Scalable Vector Graphics (SVG). There's Math ML, the mathematics markup language. There's SMIL, the Synchronized Multimedia Interface Language. There's a series of XML-based toolkit or functional languages such as Extensible Stylesheet Languages for Transformation, and XPath/XPointer for being able to point to things in documents.

People should go strolling over to the Worldwide Web Consortium at w3.org. They're like our big brother of the Web3D consortium. They have over 450 members, they have a huge amount of work going on, they have very few numbers of people questioning, they're not full of angst at what they're doing - they actually realize that this is all best business case, best enterprise strategy, even though that's individual strategy for how to be proactive. It's popping back up to our requirements for X3D, the

first one was backwards compatibility, the second one is go-forward compatibility with the Web, and the third is extensibility.

We want an individual or a company to be able to add new extensions for 3-D interchange, with or without getting permission. We have, at my count, five different mechanisms to extend the X3D language.

OK, now some final points on how X3D differs.

There have been lots of advances in graphics hardware and rendering in the last five years, so we have incorporated multitexture, 2-D primitives, keyboard, NURB surfaces, geospatial nodes, humanoid animation, and the Distributed Interactive Simulation protocol. All of these projects have matured in various Web3D working groups. So it's a great time to be looking at X3D as an interchange format.

The requirements for X3D are backwards compatibility with VRML, the ability to go forward with the Web, and extensibility.

We've kept the file-format status quo by formally including both the XML encoding used in the Web languages and keeping the VRML encoding, i.e. all the squiggly and square brackets that look familiar to VRML 97 authors. So, all in all, a lot of good differences. I hope I answered the question for you.

RTG: Yes, that's fine. One of the problems with VRML was that basically the browsers didn't work. They were all supposed to read VRML and in fact, the ones that made it out on the net didn't work very universally. You had to write for different VRML browsers. Is there something that can be done inherently to address this problem?

DB: A very important point. We tackle that in a couple of ways. First is that the

specification is tighter. We try to get rid of phrases like behavior is undefined and get more explicit about what browsers should and should not do. Second, we have built, over the last several years, a very large conformance suite. It's up around 750 scenes right now. It was originally authored by a team at the National Institute for Standards and Technology (NIST), and it exercises most of the nooks and crannies of the X3D specification.

We have test scenes in either encoding, XML X3D or classic VRML, and we give away CDs with the test suite twice a year as part of the X3D Software Development Kit (sdk.web3d.org).

NIST went the whole distance with testing, providing images and movies for reference on what the proper view looks like with proper rendering and proper behavior. So, it's a huge resource. Even so, I think this is not enough. This conformance suite is sort of like "hey, everybody eat your spinach, it's good for you." Companies don't always have the inclination or the resources to perform every test.

We also found that companies, despite goodwill and good intentions, couldn't always execute with each other when there appeared to be interoperability issues. We found a way to fix all that: an open source implementation, which is one of the two implementations required for X3D completion. The name of our open source is called Xj3D, available at xj3d.org and also as part of the Xj3D software development kit.

The code base was originally donated by Sun, so it's based on a lot of plain-vanilla Java for the scene graph along with Java 3D for rendering. Nevertheless, it's grown hugely since that first contribution. The principals among many equals are the guys in Yumetech, Alan Hudson and Justin Couch, plus Rick Goldberg of Aniviza, maintaining the code-base on a daily basis.

They have probably a dozen members that are regular contributors. The source and the executables are available for free under an LGPL open-source license. A few members have been able to fund some new projects. For example,

we have a Java-GL bindings group - your readers may be aware that three or four different implementations have existed over the years - that essentially let people write OpenGL code straight in Java.

Those bindings have all been ad hoc and by external groups. This current effort is looking at all of those prior projects and is sanctioned by Sun. We expect to produce cross platform OpenGL bindings for Java and do that as a separable part of this XJ3D open source. So eventually we'll even make the Java 3D rendering optional.

We're pretty happy with Java 3D, our NPS group is seeing excellent performance of 30 to 60 scenes per second - on very large scenes right now. Still, you can always do better and for small footprint devices, printing, GL is clearly the way to go. I would expect that work to get past the design stages and be in the useful stage probably in the middle of next year.

RTG: So the Java bindings are basically on the browser side of the equation?

DB: Correct, although the source code would be available for any program.

RTG: If you have an open source browser that works, doesn't that kind of defeat any attempt to protect content? I mean, you can always use the browser to read the database and then write it out again. The way the game people protect content is they have a closed source browser. So whatever it is they're doing to get from the database format to the screen, a third party doesn't know.

DB: An important distinction is that we don't attempt to try to force everybody in the world to use open source - quite the opposite. Two mechanisms can be used right now - on the content side, meaning on the authoring side to protect the XML version of X3D, you can use the World Wide Web Consortium recommendation for XML encryption, and independently or simultaneously also use W3C XML authentication to digitally sign your content.

So, you can prevent a browser from opening content, or you can force it to acknowledge that the authentication key is there and content is bought — or belongs to someone. Further, given the existence of the open source, we reduce

the barrier to someone else implementing that, perhaps in a closed or commercial way. Commercial is fine - if people are making money off interoperable 3-D content, we think that's a good thing. We think it's an even better thing when an author can either share or protect content as they wish, not being gated by the particular constraints of one patent holder or another.

If somebody wanted to write a carefully locked browser, it would be a fairly straightforward matter to either look at the open source and see how it's done for your own product or just adopt the open source and make your own specialty browser.

RTG: So you could, for example, bury encryption code somewhere in the source, compile it, and make your own custom proprietary piece?

Hardware is not an obstacle. It's been three years since one could buy a computer without 3-D graphics acceleration.

DB: Yes, it is very easy to do, and licensing is not mandatory. So there's no moral or legal requirement to give your changes back to the open source library.

RTG: We always have a spectrum of users out there with minimal systems and some with the latest and greatest graphics and we always wish that everyone had the latest and greatest. What kind of an obstacle is that posing to the Web3D initiative?

DB: You'll get a variety of opinions perhaps, but I think it's no obstacle at all. It's been three years since you have been able to buy a PC without acceleration, so you have to work pretty hard today just to find such a computer. Second: Moore's Law, great as it is, is surpassed by progress in 3-D chips. Those chips tend to progress 1.3 to 1.5 times Moore's Law - it's a huge growth rate with new capabilities going in all the time.

I consider the hardware problem essentially solved and getting better every day.

The real problem, the real challenge facing the whole industry is content interoperability. The X3D interchange specification is not trying to compete with any other API or any other technology. We're pretty comfortable with letting people do what they do well. What we hope to add to the mix is how to easily show somebody else your work, out on the World Wide Web. Clearly the race is not going to the company that somehow figured out how to get their proprietary plug-in on every workstation and every PC, because if that was going to happen, it would have happened already.

The issue is how do we get interchange, and how do we pass new content around. We're pretty happy with the level of sophistication in the new X3D specification - we've got most people's wish lists covered in terms of broad graphics capabilities. More than that: the extensibility story is so strong, the extension capabilities here are so excellent that people are not going to have to wait another five years for the next great leap. Rather they can start using components and start adding in additions right now. The LatticeXVL parametric surfaces are our commercial exemplar how to do that.

Another good test of extensibility, actually, is when we go in reverse. We have multiple profiles that range from fat to very small, similar to the component levels going on in XHTML and in SVG. We have a Full Profile that covers from soup to nuts, an Immersive Profile similar to VRML 97 capabilities, and we also have a lightweight Interactive and Interchange profiles which are under 30 nodes shared with the MPEG4 specification. We've successfully figured out what's the minimum subset of nodes needed to get renderable content all within a small plug-in or a small device. It's been about two years now that we've had that. Several Web3D companies have shown that good-looking plug-ins on the order of 50K and 100K are possible.

Continued on page 10

Interview: Don Brutzman

Continued from page 9

It seems like a great entry-level mechanism for how to install a browser: first you load the light content - and the light plug-ins, then while people start doing cool work - it's pulling down the rest of the immersive browser component behind it. We think the door is now open for that approach. We also think the door is open for people to start getting serious about 3-D browsers on hand-helds. ParallelGraphic's Cortona has one already. So small is also good. When 3-D space is just a big, mushy changing mass of nodes and capabilities, it's very hard to optimize.

Now that we have this X3D interchange standard, either with an interactive or immersive or full-up profile, the hardware designers can say, "Ah, I know which profile I want to draw fastest;" and the software folks can say, "Ah, I know how to prioritize my software development, and how to get the most bang for the buck first." And then finally, the usually forgotten group is the authors. Authors can now say, "Ah, I can now author my content so that it is initially a very lightweight scene, it will work on a small device or a big device without any impediment, and meanwhile I can figure out how to transition to my way-bigger scene which is where I really want to go."

We've been listening to everybody, and I think we've done an excellent job in combining all the requirements. Now that the browser updates are shipping, and now that the spec is shipping, the browser companies can finally start locking down their implementations over the next half year or so.

RTG: Who's putting money into writing browsers at this point?

DB: There is some money helping to underwrite the open source browser. That remains as wide open as ever, though money's not a pre-requisite to play. There's a new company, Tony Parisi's Media Machines. Tony's first VRML browser shipped with Explorer a number of years ago, and he's got an excellent product. He's also one of the

principals on our X3D spec team as we bang out necessary work, as we decide on what are the key technical requirements that X3D is going to support, or else that we're going to defer.

Let's see - who else? Parallel Graphics has demonstrated their staying power in this area, and they continue to update and improve their browser. That's Connell Gallagher headquartered in Dublin, Ireland with the majority of their very big team back in Moscow. There is also Paul Diefenbach of OpenWorlds. He's made some critical contributions regarding advanced techniques and also serving sometimes as the commercial test case on what it takes for a browser company to be able to integrate X3D into an API toolset.

*Metadata is data
which describes data,
in other words it is a
format description.*

Somewhere on the edges, I hopefully think still a player, is blaxxun (www.blaxxun.com) and their published source base. Maybe a restructuring of blaxxun is a possibility or a spinoff of some of their principal programmers.

With that said, I think we should also note that we never expected anybody to make money on a browser - quite the opposite. We think that browsers are going to be readily available commodities. The corner we're trying to turn is when that's true. When those browsers look the same and act the same, then we all get to be where we really want to be: much cooler applications and scenes available for people on the Web. And that's when things like added-value services, great 3-D authoring, and integration with the whole plethora of XML-based information-technology (IT) tools - that's when things really start to take off for 3-D.

So, if anybody still thinks they might get rich on a browser - gee whiz, please

look beyond that! Browsers are just the door we need to crack open for everybody so that 3-D can shift from being a niche technology with perhaps a hundred incompatible tiny niches, and instead become a full first-class citizen in the Web, a part of commercial and open enterprise strategies for all the big things going on today.

RTG: What's the CAD3D group and how does that relate to the overall effort?

DB: The CAD3D group is new and very cool. Intel approached the Web3D Consortium early in the year, started a dialog and just a few months ago started a group. It's grown very quickly and has taken off. The goal of this group is to expose CAD, computer-aided design files to the Web. The members that Intel has attracted in this effort are diverse. There are of course CAD companies who've joined and there's a liaison organization called OpenHSF which can handle about 30 different CAD formats with an 80% common-denominator format (www.OpenHSF.org).

There's also a number of big-user companies such as Dassault, Boeing, and large manufacturers. One of these companies reported that their corporation's intellectual assets include 70 terabytes of CAD data of all sorts of flavors and sizes. They would like to be able to expose some of that to the Web.

The CAD3D group is meeting face-to-face every other month. It differs from the X3D group in that it's a "Web3D members-only" group. The CAD3D participants have decided they all want serious commitment from all participants.

We've just completed the requirements stage. We spent two days last month up in Portland, Oregon integrating multiple submissions of what the requirements are for a CAD3D binding. So we meet again in January and again in March, probably in France, as part of the 2003 Web3D Symposium. (www.web3d.org/s2003).

Our timetable is to complete this work about a year from now. Intel has really set things up for rapid progress. You can read the press release issued with Web3D this past summer - everything

in there and is right on target. Intel is further dedicating some resources to this - the participating group in Intel expects to produce a free player - as part of this effort. I would not only be able to render such CAD-to-Web content, but also presumably serve as an exemplar for other folks implementing this.

RTG: I believe there are companies like Engineering Animation that do all this over the Web now. You can buy their products, you can read CAD formats, you can interact with CAD models, and you can do anything you want to do. Is this basically replicating the existing technology in a way so that it's cheaper? In other words, a way for ten companies to contribute and own a product that competes in this market with the proprietary efforts?

DB: It may be repeating it, or just formalizing it. Technically and economically, many people have been down this road before. It's hard to do it all yourself and it doesn't make a lot of sense when a solution only works on a small subset of content.

RTG: I think what you're saying is 'yes.' There's no real new technology here, it's just that if you buy it from somebody, it's really expensive, whereas if you team a whole bunch of companies together, then they can afford it and share it.

DB: Right, and that also helps establish a market wherein it's long term, it's sustainable. It's a very competent group, I doubt there are any technical surprises. What is happening, we're trying to optimize across a really big space. I like the way Intel is leading and guiding this group.

RTG: I guess Intel's notion is that if more people have CAD interchange, and if it's cheaper, more people will have it, and if more people have it, then they'll go and buy more computers to support it. Is that the logic there?

DB: Yes. Web3D wants to show people cool, compelling, and economic reasons why they should use 3-D. This lines up pretty well with a company like Intel that wants to show them why they should buy a new computer.

RTG: What is the MOVES Institute and what's it doing?

DB: The MOVES Institute is an academic and research group inside the Naval Postgraduate School. MOVES stands for Modeling, Virtual Environments, and Simulation. We have about 30 faculty, 40 students, and our own degree program. What we did was to keep about half of the computer science program, particularly things with graphics and physically based modeling and human interaction. Then we integrated about half of the operations research and analytic mathematics curriculum, so that our students could do both interactive simulations, such as cool real time games, and also analytically correct, statistically precise types of abstract simulation.

*A binding is a way
to incorporate the
features of one
computer language
into another.*

It's very successful. We're now able to work on completely interdisciplinary projects and graduate students who are extremely talented and not gated by fences.

Fundamentally there are two sets of simulations here - all the things rotating around graphics and gaming and so forth, and all the efforts spinning out of more mathematical representations of the world.

Our military, our allies, need people with both sets of skills, so we're producing a pretty amazing cadre of students out there - who can serve as program managers and serve as leaders to get work done and not be fooled by the particular problem du jour. Our biggest successes to date are the U.S. Army games, with over one million registered users at AmericasArmy.com.

RTG: What did you think of I/ITSEC?

DB: I/ITSEC (Interservice/Industry Training, Simulation, and Education Conference) was tremendous, a big show of around 10,000 people in Orlando, Florida. I/ITSEC is like SIG-

GRAPH in terms of cool tech and applications. Of course, it also has a military slant. We are showing how all types of networked Web-based simulation can connect to the military and while they still utilize simulators and domes and other devices plugged in. That even-broader interoperability is the goal of the XMSF project.

I/ITSEC also has a portion dedicated to Advanced Distributed Learning (ADL), which is the notion that you can take distance learning content and mark it up with consistent metadata so that a variety of different learning management systems might run it for students.

MOVES is becoming involved in ADL more and more. We've produced the SIGGRAPH On-line work, which is about 200 hours of SIGGRAPH video with presentations content. We're also trying to show ADL using X3D content, meaning that all of the strengths of XML for metadata and for multi-lingual internationalization can also play well in the ADL XML metadata space.

So, it's a great mix at I/ITSEC that I recommend heartily to companies that want to make money with the Government. It's also good for folks who are just interested in seeing the state of the art in Orlando, the state of the art of modeling and simulation in our country. The funniest quote overheard in our booth was by an attendee looking at our on-demand creation of X3D terrain from DTED databases via XML. He said, "Great - I'm so tired of paying for a new version of the same old dirt every two years!" So interoperability remains compelling. We'll keep showing the military how to use X3D and XML there, and will be glad when the rest of the 3-D industry joins in.

Contact Don Brutzman:

Code UW/Br, Root Hall 200
Naval Postgraduate School
Monterey, California 93943-5000
Tel: 831-656-2149, Fax: -3679
e-mail: brutzman@nps.navy.mil
<http://web.nps.navy.mil/~brutzman>

Contact the Web3D Consortium, Inc.:

c/o Anders Jepsen
P.O. Box 2272
Orinda, CA 94563
Tel: 925-254-3079, Fax: -5740
e-mail: AJViking@aol.com
http://www.web3d.org/fs_contact.htm