A Transatlantic Research and Development Environment

Michael R. Macedonia
Fraunhofer Center for Research in Computer Graphics

Stefan Noll
Fraunhofer Institute for Computer Graphics

The Internet has undone what Columbus discovered 500 years ago: It has made the earth flat once again.
—Daniel Lynch quoted in The Unpredictable Certainty[1]

Why do we need global virtual environments? Consider: Our world is entering an age where our current understanding of telecommunications and graphics computing is constantly challenged. For instance, the universal advancement of network and graphics technology, new business models, and global infrastructure development are transforming the solitary, platform-centric 3D computing model. With the availability of global information highways, 3D graphical intercontinental collaboration will become a part of our daily work routine. Already, world class auto makers are attempting to reduce car development time to two-year cycles, enlisting global engineering teams. However, this process requires new tools such as shared 3D CAD and distributed product data management systems.[2]

Distributed 3D virtual environments can help researchers conduct experiments globally with remotely located participants. Here, we discuss challenges and opportunities for the shared work environment.

Our organizations, the Fraunhofer Center for Research in Computer Graphics (CRCG) in the United States and the Fraunhofer Institute for Computer Graphics (IGD) in Germany, are looking ahead into this new age by establishing a transcontinental computer graphics research effort (see Figure 1) and a proposed G-7 testbed. We are studying how 3D computer graphics and virtual environments can aid global collaborative work.

We have focused our research efforts on determining how computer networks can transform the distributed workplace into a shared environment, allowing real-time interaction among people and processes without regard to their location. Our research is part of a long-term project—the Transatlantic Research and Development Environment (Trade)—that lets us

- identify key research and development topics for global VEs,
- define and select experiments that use the benefits of 3D environments and high-speed digital networks,
- establish an intercontinental distributed testbed for gaining operational experience and management practice,
- deliver input for defining a standard set of services for common and open use,
- obtain measurement and performance results and overall system evaluations, and
- develop global VEs for research and development, industrial, and commercial applications.

Research

Much needs to be studied. In the 1995 report, "Virtual Reality: Scientific and Technological Challenges,"[3] the National Research Council surveyed key challenges in VR technology and research. Years of effort on distributed interactive 3D graphics applications have clearly identified networking issues as the critical bottlenecks preventing the creation of large-scale VEs as ubiquitously as home-page space on the Web. Those years have also produced important component technologies such as VRML and Java. We believe the time is right to combine the network pieces with the graphics pieces. Trade provides a platform for that development.

Infrastructure issues

A key problem for networked 3D communication that the NRC report and Trade address is that the Internet, in its current state, does not adequately support interactive 3D applications. A common misperception in the graphics community is that bandwidth remains a major impediment to using VEs. However, network technology and infrastructure have progressed with the Internet's growth and the increased competition in the telecommunications industry. This has led to substantial decreases in telecommunication costs in the United

 0272-1716/97/$10.00 © 1997 IEEE

States—a trend expected in Europe in 1998 with the arrival of deregulation.

Digital local-loop technologies will soon permit universal high-speed data access to the home and small business. The available choices include ISDN, Asymmetric Digital Subscriber Line (ADSL), and cable modems. ISDN is widely deployed in Europe and the US. ADSL, which is being used in several testbeds, transmits more than 6 Mbps to a subscriber and as much as 640 Kbps upstream.

Local area networks have also advanced with Asynchronous Transfer Mode (ATM) networks and, soon, Gigabit Ethernet. ATM looks alluring because it might eliminate the distinction between wide and local area networks, thereby providing high-speed, low-latency connectivity from desktops across the world.

1 International network of computer graphics.

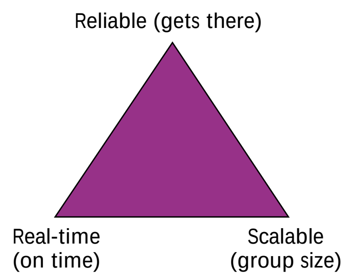

2 VE issues.

International infrastructure

International network providers such as Global One, TeleGlobe Canada, and Deutsche Telekom (DT) are investing in transatlantic ATM and ISDN network infrastructures. Transatlantic communications costs are expected to rapidly decrease because of

- The emerging competition between the three major international blocks led by AT&T, British Telecom, and Global One.
- Enormous increases in available terrestrial links and bandwidth. Currently 300,000 km of undersea fiber optic cable exist. By 1997 more than 100 countries will be connected by an additional 100,000 km of cable. The TAT-12 cable between the US and Europe will have a 10 gigabit-per-second capacity (handling 120,000 phone calls) that can be doubled in five years.[5]
- The ability to use shared links, which support multimedia communication thanks to ATM and IP technology in contrast to dedicated leased lines.

Network maturity

Much work remains before digital networks can provide the services necessary for global VEs. ATM, and to a lesser extent ISDN, are immature technologies with respect to equipment compatibility, application-programming interfaces, and operations. For example, ATM-switched virtual circuits (SVCs) that allow users to "dial up" a connection to a remote location are not widely supported or available. Therefore, links must often be prearranged administratively with the service provider each time a new connection is established. Differences in the ISDN standard for Europe and the US are also not trivial, though they can be overcome. On the other hand, international ATM standards are generally uniform.

Latency

Aside from the bandwidth infrastructure to support global VEs, latency persists as another important issue. Latency represents the dark side of constructing global environments (see Figure 2). For a distributed environment to emulate the real world, it must operate in real time in terms of human perception. In addition, real-time audio and video communication services must be provided.

Overcoming latency becomes a major challenge in systems that use WANs because of delays induced by long paths, switches, and routers. However, communications reliability often forces a compromise between bandwidth and latency. Reliability in this case means that systems can logically assume that data sent is always received correctly. Some networks do not guarantee reliable communication. ATM and UDP/IP offer best-effort service. Cells or packets are dropped "silently" when a switch or router encounters congestion.

To guarantee delivery, the underlying network architecture must use acknowledgment-and-error recovery schemes that at times introduce large variable delays. Protocols like TCP are required to ensure reliable service. However, transport protocols use congestion control mechanisms that do not suit real-time traffic because they throttle back the packet rate if congestion

Hands Across the Atlantic

R. Bowen Loftin
University of Houston

Historically, NASA has trained astronaut teams by bringing them to the Johnson Space Center (JSC) in Houston, Texas to undergo generic training followed by mission-specific training. This latter training occurs after a crew has been selected for a mission and often begins as much as two years before launch.

While some Space Shuttle flights have included an astronaut from a foreign country, the International Space Station will be consistently crewed by teams containing astronauts from two or more of the partner nations. As the task of training crews for the Space Station comes closer to realization, the cost, in terms of both travel and physical "wear and tear," appears ever larger.

NASA, the University of Houston, and the Fraunhofer Institute for Computer Graphics (IGD) have been exploring VR technology as a partial answer to the problem of helping international crews prepare for their missions. The initial effort in determining a shared VE's feasibility for team training took the form of a simple extra vehicular activity (EVA) simulation that engaged two astronauts on opposite sides of the Atlantic ocean.

A Astronaut Ulf Merbold.

On September 20, 1995, Astronaut Bernard Harris (physically located in Houston at the Johnson Space Center) entered a VE with Astronaut Ulf Merbold (physically located at the Fraunhofer Institute for Computer Graphics in Darmstadt, Germany). Their shared environment contained models of the Space Shuttle payload bay and the Hubble Space Telescope (HST) (see Figure A). The two astronauts spent more than thirty minutes performing the major activities associated with the changeout of the HST's Solar Array Drive Electronics (SADE). Their work included the real-time hand-off of replacing the original SADE. When the task ended, the two astronauts shook hands and waved goodbye.

This demonstration was accomplished by generating duplicate graphical environments at each site (using Silicon Graphics Onyx Reality Engine workstations) and exchanging state change data (for example, the movement of one astronaut's hand or the translation of a suited astronaut to a new site) via a commercial Integrated Services Digital Network (ISDN) line connecting the two sites. My colleagues and I used NASA-developed software to render the graphics and Comm++ library, an IGD-developed communications software program, for exchanging state data. Voice communication occurred over a

is detected and induce delay through acknowledgments.

The high ratio of propagation time to cell transmission time demands new methods of handling network congestion. Propagation time for data is roughly 10 ms per time zone and therefore about 60 ms between Providence, Rhode Island and Darmstadt, Germany. Because of the bandwidth-delay product and the windowing scheme used in TCP/IP, a computer in Europe sends 20 megabits of data on an OC-3 link before a host in the US can send a negative acknowledgment (NAK) message informing the computer to stop and resend. A loss of one 53-byte ATM cell carrying IP traffic results in losing an IP packet typically carrying 1,000 bytes. Moreover, reliable transport is not appropriate for interactive VEs, which often can recover from a lost packet more gracefully than from late arrivals.

Methods do exist for ameliorating latency effects, which we are exploring with Trade. Network latency can be reduced somewhat by using dedicated links (or virtual ones using protocols like the Reservation Protocol[6] or ATM), improvements in router and switching technologies, and faster interfaces and computers. Dead-reckoning techniques not only reduce communications loads on the network but also smooth perceived delays as a result of predictive modeling by the local host.[7] However, lag time can never be totally eliminated for environments where the VE remains widely distributed. Therefore, we use techniques such as synthetic fixtures that provide clues to operators in limited domains about that environment.

Testbeds

Several US and European testbeds are conducting feasibility studies on using multimedia services over high-speed networks for a variety of application areas: conferencing, education, engineering, publishing, medical consultation, and tourism (see the sidebars on astronaut training and medical ultrasound). Recent examples include the Magic network, the Berliner Communication Network (Berkom), and the European ATM Pilot. The Navy Research Lab is also investigating using ATM with distributed interactive simulation and other VEs.

The Trade project uses an ATM network between Providence, Rhode Island, and Darmstadt, Germany.

commercial telephone line.

Both astronauts judged the performance acceptable, with very stable round-trip latencies of less than 200 ms. Only 20 percent of the available bandwidth was required (total bandwidth was 128 kilobits per second). This suggests that inexpensive commercial communications channels could be used for this type of shared VE and that at least three sites could be supported with the typical ISDN bandwidth. After the demonstration, both astronauts strongly supported this training method in cooperative task performance and for procedures and timeline development.

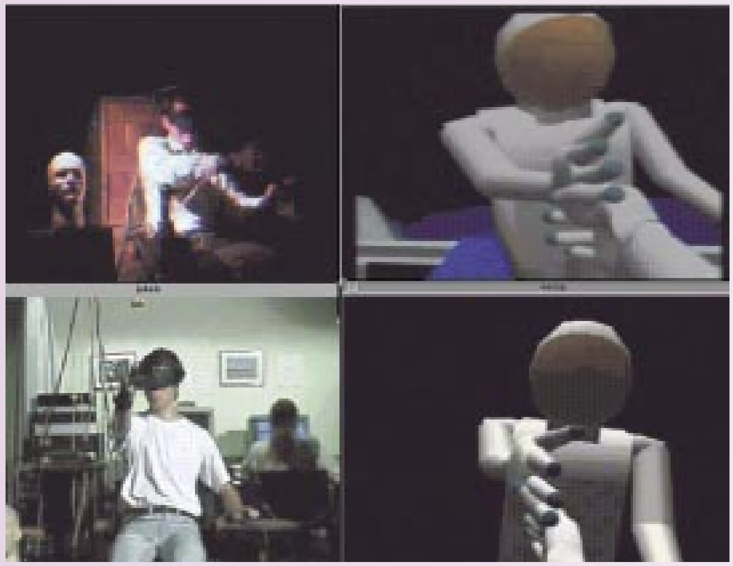

B Views of astronauts in New Orleans (above) and in Darmstadt shaking hands.

Since the successful 1995 demonstration, we have

- examined a variety of communication options in addition to ISDN,
- integrated additional sites into the shared environment, and
- developed a better understanding of human interaction in these shared environments.

During Siggraph 96, the University of Houston and the Fraunhofer IGD demonstrated shared environments within the conference's Digital Bayou (see Figure B) by using ATM protocols over a heterogeneous communications network. ATM can provide communications with minimal delay for collaborative work in VEs. The Fraunhofer CRCG provided resources from the Trade project and its expertise with global ATM.

Given the needs of the current Shuttle Mir and the International Space Station programs, astronauts strongly support routinely sharing VEs between JSC and other overseas sites for many training needs. Clearly, this type of team training via communication networks will become commonplace and eventually expand into orbital settings.

For more information, contact R. Bowen Loftin at bowen@gothamcity.jsc.nasa.gov, www.vetl.uh.edu; Robert Cross, Fraunhofer CRCG Providence, rcross@crcg.edu, www.crcg.edu; or Rolf Ziegler, Fraunhofer IGD Darmstadt, ziegler@igd.fhg.de, www.igd.fhg.de.

Operationally, we establish a permanent virtual channel (PVC) to Darmstadt. CRCG connects to Global One's Reston, Virginia laboratory via Sprint's national ATM network. From there, our traffic transmits via the transatlantic facilities of the Multimedia Applications on Intercontinental Highways (MAY) consortium (see Figure 3). The segment terminates in Berlin. ATM cells are then forwarded to IGD in Darmstadt via Deutsche Telekom's experimental ATM network. The MAY group includes broadband network providers such as DeTeBerkom, TeleGlobe Canada, and Global One as well as users such as Boeing and Daimler-Benz.

We are currently using IP protocols over ATM. This approach permits us to test and compare several applications that had already been developed for IP networks. Moreover, pure ATM does not support multicast communication except when using proprietary schemes, which proves a major deficiency.[8,9]

Future work

Don Brutzman at the Naval Postgraduate School (NPS) stated that networked interactive 3D graphics

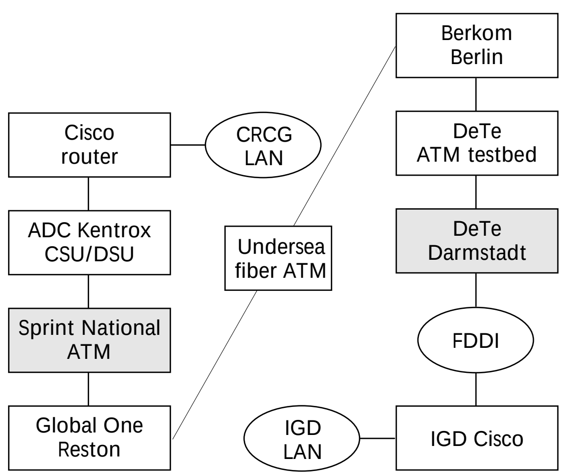

3 Trade network setup.

rendering proves an essential human interface technology for dealing with information overload.[4] A VE network is an inevitable development as access to practically all technologies and disciplines converges via

TeleInViVo: Telemedicine on the Desktop and on the Road

John Coleman, Andrei Savchenko, and Ammo Goettsch
Fraunhofer Center for Research in Computer Graphics

Fraunhofer CRCG has developed TeleInViVo, a medical visualization software program that permits two geographically separated radiologists to view patient data simultaneously and to consult with each other in real time. It is an extension of the Interactive Visualizer of Volumes (InViVo), Fraunhofer IGD's family of highly portable software that has helped bring computer-based visualization power to the desktop (see Figure C).

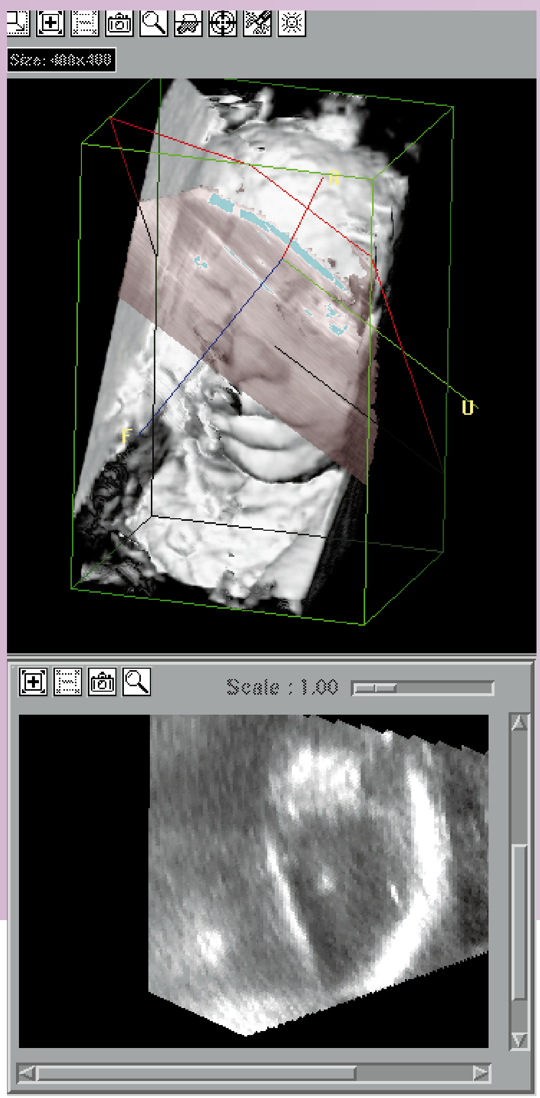

C **3D scan data.**

In 1996, the Battelle Pacific Northwest Laboratories deployed for DARPA the first version of TeleInVivo to Bosnia. They did so as part of the MustPAC (Medical Ultrasound, Three-dimensional and Portable with Advanced Communications) 3D Ultrasound Diagnostic System. We also tested TeleInViVo over the transatlantic ATM link between CRCG Providence and IGD Darmstadt as part of the Trade project.

Architecture

TeleInVivo extends InViVo's software architecture by using a distributed, shared state system. When changes to the local state are made, the scope of these changes is analyzed. Messages communicating state changes are sent over the network. Analogous changes are then made in the state in the remote system. If the TeleInViVo system detects that the distributed views cannot be kept synchronized by these incremental updates, the entire state transfers from one participant to the other, reestablishing consistent views.

Scanning 3D ultrasound data

Fraunhofer CRCG has been studying ways to use TeleInVivo in new domains and application

the Internet. We wish to enable the initial infrastructure for implementing distributed VEs as predicted by "Gibsonian" cyberspace, that is, the futuristic vision of unconstrained VR prevalent in popular literature. An open architecture will use standards-based software and repeatable, exemplary implementations, all taking advantage of the sustained exponential growth of internetworked global information.

For example, the Trade research group, in collaboration with the Naval Postgraduate School, has been studying constructing a Virtual Reality Transfer Protocol (VRTP). VRTP is to be the applications layer protocol used for communicating state information among the various participants in internetworked VE. We plan to examine the entire gamut of issues involved in developing such a protocol, including rapid reconfigurability, tools for rapid specification, scalable network and software architectures that use VRTP, and other associated issues. Application areas for VRTP include distributed training simulation and telemedicine.

Conclusion

As you can see in the sidebars on telemedicine and astronaut training, 3D graphics for global work have already arrived, but many challenges remain. Applications for this research range from simple distributed multimedia visualization of scientific data to scenarios. One such area, which proves increasingly important for this kind of technology, is in 3D ultrasonography. Because of its noninvasive nature, relatively low cost, and portability, ultrasound is quickly becoming the modality of choice when considering difficult-to-serve or expensive-to-serve areas. The MustPAC system was located at the field hospital in Tuzla. A similar workstation was located at Madigan Army Medical Center in the state of Washington (near Seattle), where radiology specialists consulted in real time with their medical counterparts on location in Bosnia. They used TCP/IP over shared military communications channels to make the connection.

The Madigan site was configured with a six degree-of-freedom Immersion Probe input device. Attached to the probe's "stylus" is a mock-up of an ultrasound transducer. This supports remote diagnosis by letting the remote expert arbitrarily scan the data rather than the patient, using the Immersion Probe as though it were an ultrasound sensor. The concept of scanning the data rather than the patient is a relatively new paradigm and significantly aids in remote collaboration.

Collaborative diagnosis

As Figure D shows, TeleInViVo supports shared viewing tools to help facilitate discussion among the collaborators**.** These tools include a telepointer, which has become standard in collaborative applications. Additionally, domain-specific pointing methods are supported by extending proven viewing tools like InViVo's real-time magnifying glass. TeleInViVo allows one user to operate the magnifying glass on the local screen and have a duplicate glass moving synchronously on the remote end. This lets the user call attention to specific details in the data by transmitting part of the user's "way of looking at the data." Because the data being visualized is 3D, common visualization tasks involve frequent changes in the angle from which users view the data. TeleInViVo lets users interactively change the view on both the local and remote machine, while keeping the views consistent. In this manner, both users can move the data around and telepoint on top of the 3D views.

Visualization settings not controlled by direct manipulation are also shared between the users. TeleInViVo's menus manipulate a shared set of settings in this case. By controlling the view on the remote end as well as the current filtering settings, TeleInViVo can be used to teach or assist with a diagnosis session. For example, a specialist could assist a physician from a remote site by selecting an appropriate view and isolating critical detail in the data.

For more information on TeleInViVo, contact John Coleman, Fraunhofer CRCG, jcoleman@crcg.edu or Georgios Sakas, Fraunhofer IGD Darmstadt, gsakas@igd.fhg.de. For the MustPAC system, contact Rik Littlefield, Battelle Pacific Northwest Laboratories, rj_littlefield@pnl.gov.

D Consultation session with TeleInViVo.

distributed VEs for teleconferences and simulators. However, the computer graphics community must work with developers and users to establish progressive forms of telecommunication. Global telecommunication by businesses is currently limited by the cost of transmitting 2D image data and high-quality audio signals. However, we believe that the greatest impediment to 3D communication is that potential users do not understand the requirements and benefits of this advanced technology. We established the Trade project because advancing worldwide computer networks and VE technologies such as 3D graphics, sound, and interaction helps create new methods for global communication and collaboration. ■

Acknowledgments

We wish to acknowledge the support of Hewlett-Packard, Europe; Hans H. Jung and José Luis Encarnação of the Fraunhofer Society; and the members of the MAY research consortium.

References

1. NII 2000 Steering Committee, *The Unpredictable Certainty: Information Infrastructure Through 2000*, Computer Science and Telecommunications Board, Commission on Physical Sciences, Mathematics, and Applications, National Research Council, National Academy Press, Washington, D.C., 1996 (http://www.nap.edu/readingroom/books/unpredictable/).
2. J. Rix, S. Haas, and J. Teixeira, eds., *Virtual Prototyping: Virtual Environments and the Product Design Process*, Chapman Hall, New York, 1995.
3. N. Durlach and A. Mavor, eds., "Committee on Virtual Reality Research and Development," *Virtual Reality: Scientific and Technological Challenges,* National Research Council, National Academy of Sciences Press, Washington, D.C., 1994, pp. 362-372.
4. D. Brutzman, "Graphics Internetworking: Bottlenecks and Breakthroughs," *Digital Illusions*, C. Dodsworth, ed., Addison-Wesley, Reading, Mass., 1996, http://www.stl.nps.navy.mil/~brutzman/vrml/breakthroughs.html.
5. P. Trischitta and C. William, "Global Undersea Communication Networks," *IEEE Communications*, Vol. 34, No. 2, Feb. 1996, pp. 20-21.
6. R. Braden, L. Zhang, and S. Berson, "Resource Reservation Protocol (RSVP)—Version 1 Functional Specification," July 18, 1996, http://www.ietf.org/ids.by.wg/rsvp.html.
7. S.K. Singhal, "Using a Position History-Based Protocol for Distributed Object Visualization," *Designing Real-Time Graphics for Entertainment* (Course Notes for Siggraph 94 Course #14), ACM Press, N.Y., July 1994, pp. 10.1-28.
8. M. Macedonia et al., "Exploiting Reality with Multicast Groups," *IEEE Computer Graphics and Applications*, Vol. 15, No. 5, Sept. 1995, pp. 38-45.
9. M. Macedonia and D. Brutzman, "Mbone Provides Audio and Video Across the Internet," *Computer*, Apr. 1994, Vol. 27, No. 4, pp. 30-36.

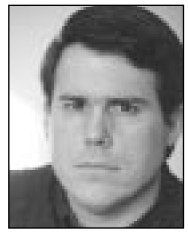

Michael R. Macedonia *is vice president in charge of developing global work environments for commerce, government, and education at the nonprofit Fraunhofer Center for Research in Computer Graphics, in Providence, Rhode Island. He received a BS from the US Military Academy in 1979, an MS in telecommunications from the University of Pittsburgh in 1989, and a PhD in computer science from the Naval Postgraduate School in 1995. He is a member of the IEEE Computer Society, IEEE Communications Society, ACM Siggraph and Sigcomm, Society for Computer Simulation, and Society of Manufacturing Engineers.*

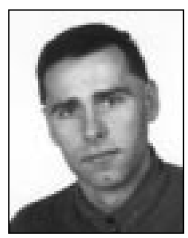

Stefan Noll *is head of the Communication and Cooperation Department at the Fraunhofer Institute for Computer Graphics in Darmstadt, Germany. He received his PhD in Computer Science at Technical University, Darmstadt. Noll's current research interests include multimedia telecommunication and user interfaces in distributed environments.*

Contact Macedonia at Fraunhofer Center for Research in Computer Graphics, 167 Angell Street, Providence, RI, 02906, e-mail macedon@crcg.edu, Web site http://www.crcg.edu.