

SYSTEM & NETWORK

Simon Penny
"Traces"




Fig0. diagramatic representations of the Traces system.

Introduction: What is special about "Traces"?

'Traces' is a project for networked CAVEs (immersive VR spaces). But it is very different in its goals and its nature from any other CAVE or VR project (to the knowledge of the author). The root of the project is a long standing concern over the disembodiment quality of the VR experience, which stands in stark contrast to the rhetoric around VR, which argues that the experience allows the user to interact in a bodily way with digital worlds. As I first argued in my essay *Virtual Reality as the end of the Enlightenment project* (in *Culture on the Brink* Eds Druckrey and Bender, Bay books 1994), conventional HMD (Head Mounted Display) VR dissects the body, privileging visibility to the exclusion of bodily senses. The body is reduced to a single Cartesian point, the body is checked at the door.

When I first used a CAVE, I was fascinated with the visceral sensation of collisions with virtual objects. I realised that part of the disembodiment quality of HMD VR was because when you look down, your body is not there! In a CAVE, you see your body colliding with virtual objects. Because of my interest in the problem of embodiment, the CAVE became an attractive site to work in. But clearly only half the problem was solved, the user could experience virtual object in a more bodily way, but the user was still reduced to a point from the perspective of the machine. Thus it became necessary to build an input system which described the entirety of the users body. After substantial research, we built a multi-camera vision system which constructs a real-time body model of the user. (see Figs 1,2,3,4,5 and of video) This body model is currently of a low spatial resolution, but of a high temporal resolution, the user experiences no 'latency' or lag, between their movements and the virtual structures created.

In Traces, virtually all sound and visual experiences are generated in real time based on the users behavior. Unlike other VR projects, I have no interest here in illusionistic texture mapped

models, the illusion of infinite virtual space or building 'virtual worlds'. All attention is focused on the ongoing bodily behavior of the user.

Traces will be a telematic, networked experience. But creating an illusion of close proximity (like in the work of Paul Sermon) is not the goal. Rather, there is an emphasis on the highly technologically mediated nature of the communication. The users never see each other, only the highly mediated results of each others behavior. The user interacts with gossamer spatial traces which exhibit the dynamics and volumes of bodies, but are translucent and ephemeral.




Fig1. High resolution body model derived from vision system.

Theoretical, Aesthetic and Technical Overview

The focus of Traces is real-time spatial/bodily interaction between distant participants via real-time 3D image (and sound) traces. The work will function as both a telematic performance environment and a public interactive experience. My previous work 'Fugitive' (ZKM 1997) uses infra-red machine vision to extract data about the bodily dynamics of the user (see appended videotape). In 'Traces', each CAVE will use multi-camera machine vision to build real-time body models of participants. These body-models will then be used to generate abstracted graphical bodily traces in the other CAVEs. If, for instance, three CAVEs are networked, then a participant in CAVE A will interact in real time with image traces of participants in CAVEs B and C. These traces will not necessarily be accurate sculptural representations but will be used to drive complex algorithmic processes which will give rise to changing 3D graphical traces which indicate the presence, gesture and movement of the remote participants. Hence a person may be represented as a moving ghostlike transparent and wispy trace (see Traces videotape).

In almost all VR and 3D immersive systems, the users' body is reduced to a single point, a viewpoint, and/or the end of a 'pointer'. This effectively erases the body from the computational system. This erasure is thoroughly consistent with the main streams of Computer Science research. My goal is to build a system with which the user can communicate kinesthetically, where the system come closer to the native sensibilities of the human, rather than the human being required to adopt a system of abstracted and conventionalised signals (buttons, mouse clicks, command line interface...) in order to input data to the system.

In Traces, the interactive experience is not premised on hypertextual paradigms such as interactive story telling. It is based on bodily, temporal, kinesthetic sensibilities. Traces most closely resembles a combination of sculpture-making and dancing. The paths of bodies will persist in the space, each body 'painting' as a three-dimensional brush. Decay time and graphical behavior of these 'traces' will be algorithmically controlled with interactive input. Not a 'virtual body' in the sense of the 'holodeck', but fleeting, hazy, evanescent, ghostly, chimerical representations like time lapse photographic images and the patterns of oil on water, which make explicit the layers of digital mediation. Interactions will take the form of real-time collaborative sculpting with light, created through dancing with telematic partners.




Fig2. Example of images from the four cameras. The body model in Figs 3 and 4 are derived from this image.

This form of remote interaction does not map onto conventional categories such as TV or Internet, it expands the concept of the 'Stage' into the telematic. The important theoretical foci, are:

- use of machine vision techniques for rich, but unintrusive, full body tracking and behavior analysis.

- development of an intuitively meaningful and persuasive kinesthetic interface without need of texts, icons or user training.
- development of the desire-driven 'autopedagogic interface'.
- to enrich and complexify notions of telepresence and immersion.

Technical achievements of the project will include:

- use of 3D machine vision as sensor input to a CAVE
- use of exclusively real-time graphics and real-time spatialised sound in a CAVE
- real-time networking of CAVEs

Traces is concerned with time and history, in two ways: with cultural history and with embodiment and bodily history. One might say that the twentieth century has been the century of the representation of movement. The research of Muybridge and Marey resulted in new pictorial conventions for the representation of motion (time lapse and multiple exposure photography) which were quickly exploited by artists, notable examples being Duchamp's 'Nude Descending a Staircase' and Giacomo Balla's 'Dog on a Leash'. Time and motion were similarly celebrated in the video-feedback works of Nam Jun Paik, The Vasulkas and others in the early days of video art. Most recently the algorithmic 'brush' in computer graphics performs a similar function. In 'Traces', the user's body acts as a three dimensional brush.



Fig3. Body model derived from images in fig2. (see also videotape)

Day

3D machine vision

The task of the vision system is to supply a volumetric body model to the graphics and sound systems, complete with data about the acceleration of body parts. It is critical that the entire system function in real time. For persuasive kinesthetic feedback, spatial resolution is less important than temporal resolution. The vision system and modeling code must be highly optimised to produce model updates with zero latency. The vision system uses an array of low resolution monochrome cameras (see Fig 5), with visible light filters (to filter out light from the projectors (beamers)). The space will be lit with multiple near IR light sources, to minimise shadows.

The vision system was developed by Andre Bernhardt and myself between November 1998 and January 1999, with financial support from the Robotics Institute and the Studio for Creative Enquiry, both of CMU. It was demonstrated on January 14, 1999. Using novel and highly efficient implementations of vision algorithms and techniques, the system runs at 15 frames per second (effectively, 'real-time'), and builds a new body volume per frame. The system uses four cameras and a Pentium PC. Body model data will be piped from the vision system to both local and remote Onyx's (the host computers for each CAVE).

In addition to providing the spatial description of the body, the system also produces measurements of the velocities of five extremities of the body (head, hands and feet). These are called the 'speedpoints'. they have lines attached which indicate velocities (see Fig6.)

The body model is represented for development purposes as a collection of 5cm cubes which we call 'Legoman' (seen in the attached still images and video). The cube size represents the maximum resolution possible in real time on the current hardware. The basic body model data is a point cloud, one for each cube. In the video, the bodymodel is depicted variously as Legoman, as pointcloud and as a variant of Legoman, Octahedron-man.




fig4. rear view of body model in Fig3. (see also videotape)




Fig 5. Miniature video camera mounted in corner of mock-up CAVE.

User Experience

The experience of the user in Traces will become increasingly complex with time, as per my theory of the 'autopedagogic interface'. At the outset, the 'traces' will be a passive residue of the user's behavior, but as time passes, succeeding traces will begin to 'behave' in increasingly complex ways, increasingly independent of the user. Likewise, the data from the body model which is utilised by the system will become more complex. At the outset it will use only volume data (i.e.: which voxels are occupied per frame or timestep). Slowly, the acceleration of various body parts (and simulated momentum) will be used, along with an elementary simulated physics.

The user enters the CAVE. A small cubic virtual room with an hemispherical dome is rendered exactly onto the walls of the CAVE in simple shading. This enclosure immediately negates the illusion of infinite space found in many virtual worlds. It is a cozy space. As the user moves, volumes appear in the space which are representations of the space passed through by the body parts of the user. These volumes have a life span, they become more transparent over a period of about a minute, then disappear. After a minute or two, another set of volumes appear in the space. These are volumes created by the remote user, appearing in real-time. Users can create volumes in response to each other.

After a while, this volume-building behavior changes. The user can create self-contained floating volumes, blobs or particles. Simultaneously, a small four-paned window appears in the back wall of the CAVE. Beyond this window the user can see another room exactly like the one she is in. In it is a very abstract moving anthropomorphic volume. This is the body model of the remote user, perhaps also with some blobs. The wall between the spaces dissolves, leaving the two users in a longer room. The users find that their blobs can be thrown and have a trajectory. They bounce off the walls. Slowly the blobs or particles begin to have their own dynamic, graphical and acoustic behavior. The particles can traverse the whole space, while the users are confined to their original part. The particles interact with each other with some sort of gravitational behavior. The space becomes populated with swarms of autonomous chattering things. Slowly it becomes more difficult to generate particles, and they all die off. Projection intensity dims and the experience is over. The anticipated time length per user is 10 minutes.




Fig6. Legoman with speedpoint balls and velocity vector lines (see also videotape)

3D Graphics

At the outset the 'traces' are reminiscent of time-lapse photographic images of Marey, translucent volumes which fade over time. When the volumes produced by remote users are present in the space, those volumes will be identified by a simple color cast: more pink, more violet, and lower color intensity. Gradually, these traces begin to possess simple autonomous behaviors. The traces might rise like smoke, or drift on a virtual breeze. In a third phase, body movements give rise to Artificial life-like entities, whose complex behavior may persist over time like an infestation of fireflies in the CAVE. Particles are originated by the user in the manner of flicking water from a wet arm: the acceleration of the hand determines the velocity of the particles. But that velocity would eventually subside due to (simulated) friction, to be replaced by the 'natural' behavior of the particles, which swarm, orbit each other, or collide to give rise to graphical, dynamic and acoustic phenomena.

At this point (early April 1999) the graphics development includes: experimentation with various representations of the body model; display of 'ribbons' flowing from up to five body extremity points; tests for the full body volume 'traces' generated through a 3D cellular automaton process, which persist over time with increasing transparency; and development of 'speedpoints', a system to track the velocity of five body extremity points in order to allow the 'throwing' of virtual objects. The supplied videotape contains short sequences of several types of imagery derived from the behavior of the body model. These clips should be understood only as code in development, not as examples of the final imagery.




Fig7. Realtime legoman representation of a user in the CAVE mockup. (see also videotape)

Spatial Sound

In the earlier parts of the project, where the user is building volumes, sound will emanate from the parts of the body from which volumes are being produced, i.e., those parts which are moving. Various approaches to the generation of sound are being researched, including mapping a physical model of a human larynx onto the space, in which the parameters of the model are given changing values generated by the history of the body model.

In the second part, where particles are generated, each particle will produce sound. Sound generation may be derived from a software implementation of the Sympathetic Sentience 2 system (Penny and Schulte 95-96). In this system, multiple simple units emergently generate constantly changing rhythmic patterns which are passed around the group. In Traces, each user will produce their own SS2 system, the total number of particles will be constantly changing due to creation by the user and by die-off, so the complexity of the sounds will change too.

Full eight-channel spatialisation of sound in CAVE environments is not a well established or even thoroughly researched field. We are planning a custom and self contained PC based eight-channel sound system which can be networked into various CAVE systems. The code will attach dynamic sound behaviors to moving virtual objects in such a way that their spatial behavior is accessible to the sound generating systems. The source sound must be varied as it interacts with other aspects of the scenario, and perceptually placed in the appropriate location, utilising both spatial positioning and treatment such as clipping of the dynamic range, and/or reverb.