ARCO and immersive environments, Part 2: Oil industry experience with immersive environments

TRACY J. STARK, GEOFFREY A. DORN, and MARY J. COLE, ARCO, Exploration and Technology Operations, Plano, Texas, U.S.

The first part of this paper (*TLE*, May 2000) reviewed the development of visualization technology at ARCO over the last 11 years. We started with desktop displays using perspective and lighting; then stereo and head-tracking were added to increase functionality and realism. Development of large spatially immersive environments started in 1996, and our first CAVE-like system was installed in 1997. Subsequent software research and development have focused on developing immersive drilling planning, volume-rendering, and attribute interpretation capabilities. ARCO's hardware R&D has focused on developing an environment which includes the best aspects of CAVE-like immersive systems and semi-immersive theater systems (e.g., visionariums) in a single display device. This culminated in ARCO's rapidly reconfigurable Immersive Visualization Environment in 1999.

Our first three years of using these immersive systems have taught us a number of lessons about the design, operation, and application of this environment. These lessons and ideas, discussed below, include the concept of visualization as a learned skill and the importance of level of realism to the effectiveness of the visualization in conveying information to the observer. Elements that contribute to this realism include "smart" data reduction to improve performance without sacrificing realism, and the importance of minimizing or controlling image distortion. The sense of realism is also significantly affected by the use of stereo display and the degree of immersion achieved through the engagement of peripheral vision. Other key aspects of effective immersive visualization include intuitive display interaction and human-machine interface. Finally, we include some thoughts on CAVE etiquette. We believe that consideration of all of these issues is crucial to development of successful visualization environments.

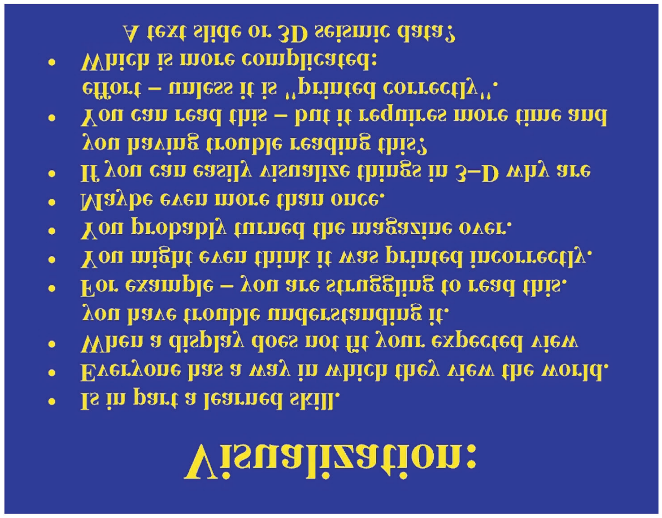

Figure 1. Visualization is a learned skill.

Visualization is a learned skill. Most people, on seeing Figure 1, will think that the editor or printer made a mistake. They have not. You should have trouble reading the figure. At first it appears to be simply printed upside down. The truth is not that simple. If the figure is rotated 180° about an axis perpendicular to the page (which turns the page upside down), it still "does not look right." However, a second 180° rotation about an axis oriented vertically on the page through the image (such that you are looking at the text through the page), will allow the text to be seen in its proper position. You can also read the text using a mirror, something Leonardo da Vinci mastered more than 450 years ago. Figure 11, at the end of this article, has the text in its "proper" orientation. Notice, however, that a simple 90° rotation makes it hard to read again.

The point is that it was easy to see that Figure 1 was a word slide and that something was wrong. In a way, Figure 1 represents the 80% solution; you can get the big picture but not the details. However, reading the text in Figure 1 is quite difficult for most, until the letters are positioned "correctly" as in Figure 11. When reading, we know what the letters are "supposed" to look like. Perhaps that is why orientation is so important. Maybe when we don't know exactly what we are looking for (e.g., the shape of the reservoir), then the orientation of the display does not really matter. Maybe, but probably not. When you expect to see things in a certain way, they are almost invisible when they are in the wrong orientation. If you don't have the ability to orient your data properly, some things will remain invisible.

Many interpreters have experienced this effect with their own data. After working for a while on a 3-D survey, they become accustomed to seeing their data in a certain way. For example, they may always display in-lines from left to right and in reverse polarity. If the same data are displayed with normal polarity and from right to left, many have trouble recognizing the image. It doesn't look right. They have built a mental framework in which the data and interpretation are placed. Things are supposed to look a certain way. The ability to move the data and modify their display characteristics is important. Standardizing on a familiar set of display characteristics frees us to examine the data's detail.

Level of realism. Each time you add another "level of realism" to your display, the faster and easier you will be able to understand what you are looking at (Figure 2). Each time you increase the "level of realism" you are effectively increasing the information content of the display. By using display techniques that mimic the world our visual system evolved to interpret, we can add information content without increasing the user's cognitive load. The set of images that

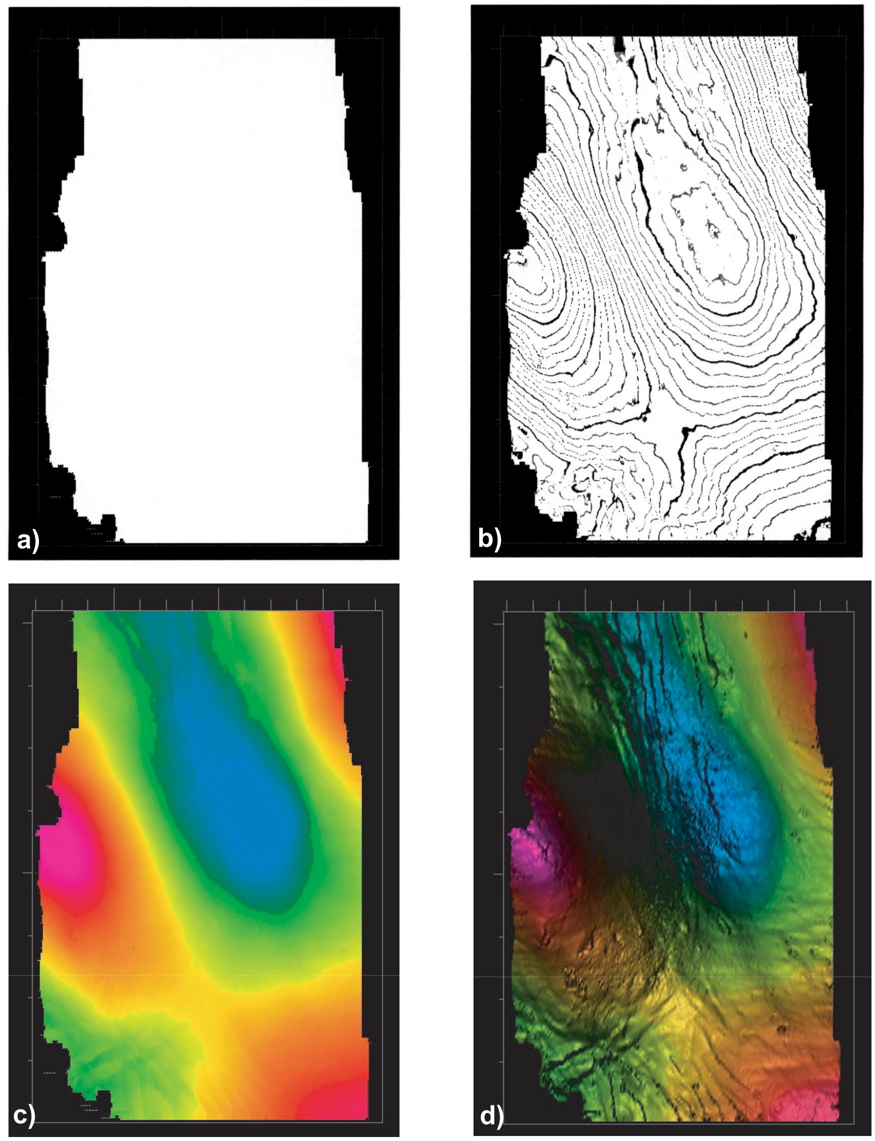

Figure 2. Top down view of interpreted horizon. Increasing "level of realism" (from a to d) improves ability to convey information.

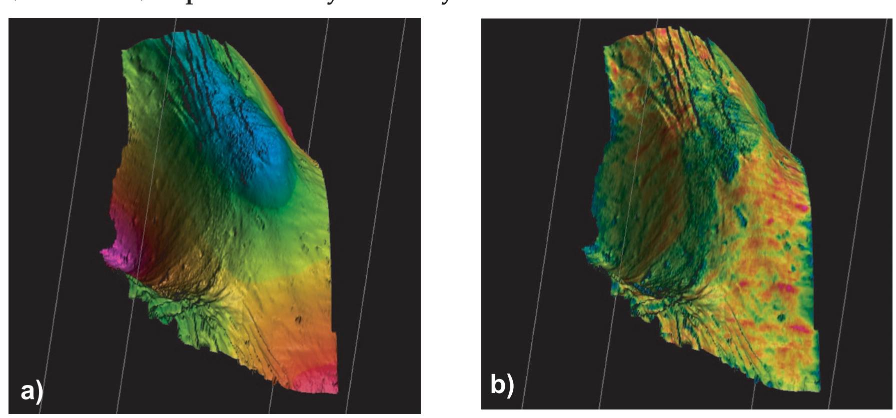

Figure 3. Perspective view of data in Figure 2. Color in a) represent travel time, color in b) represent porosity.

makes up Figure 2 contains the same data, all from the same display orientation. We are above, looking down, on 3-D perspective rendering of a horizon as interpreted from a 3-D data set. The only difference between the panels in Figure 2 is "the level of realism" used in displaying the data. In Figure 2a all you can determine is where the horizon exists (white) and does not exist (black) within the bounds of the survey. Figure 2b contains slightly more information. Here the black lines represent time contours. The closer the lines, the steeper the dip. Offsets in the lines indicate faults. There is not enough information in Figure 2b to determine if the closed contour toward the top of the image is an anticline or a syncline. In Figure 2c the elevation has been color-coded; however, there is no legend explaining what the colors mean, so you don't know if reds represent highs or lows. You don't know if the colors vary smoothly from blues through yellow to magenta, or if there are abrupt changes. In this case the colors vary smoothly; blues are highs and magentas lows. With a color bar, this image allows you to develop a better mental image of the structure than Figure 2a or Figure 2b. Of the images shown so far, this one conveys the most information.

In Figure 2d, a light source casts shadows and highlights on the complex surface. This image conveys significantly more information than any of the three previously discussed. The shadows allow us to see small faults. A legend is still needed to determine overall elevation, but relative changes are easy to see in this one figure.

Perspective displays and lighting/shading are display techniques that mimic the natural world. Color-coding and contour lines are data abstractions. We can interpret these, but this interpretation is a learned response. Different audiences will have different visual vocabularies. When using display techniques to map data into the visual world, rather than mimicking the effects of the visual world, it is important to choose abstractions that will be meaningful to the audience. Carefully chosen display techniques allow us to see both anomalies and the background.

For Figures 3a and 3b, we have moved our viewpoint from on top looking down to oblique. In Figure 3a color still represents depth; in Figure 3b color represents a seismic attribute, such as porosity. In both figures faults are clearly visible. We can see that the chosen seismic attribute does not conform to structure.

In order to completely understand the structure, a single view is not adequate; several are needed, along with the ability to move from one display to another. Other visual cues also are helpful, such as depth shading, motion parallax, and stereo displays. As we move from Figure 2a to Figure 3b and then to moving stereo images, we gain more insight about a particular data set per unit time. This ability to gain additional insight per unit time was a key motivator for building our two gen-

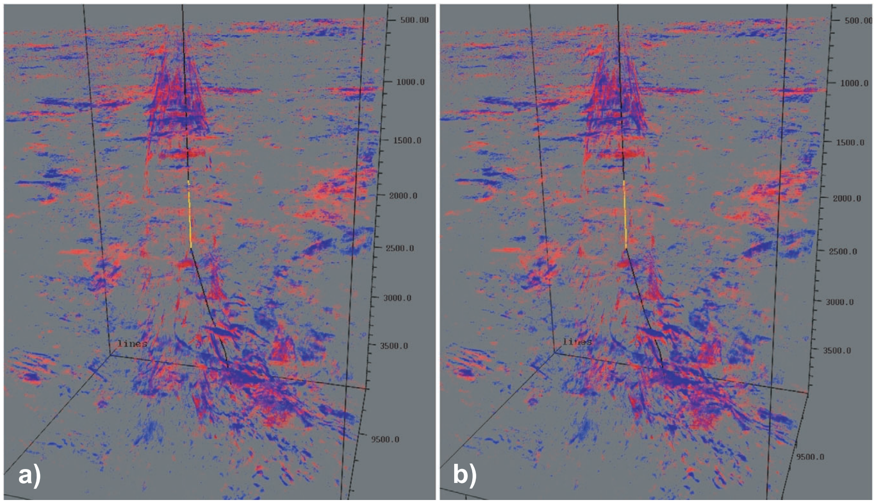

Figure 4. Stereo pair to show relationship of well path to the salt dome and amplitude anomalies.

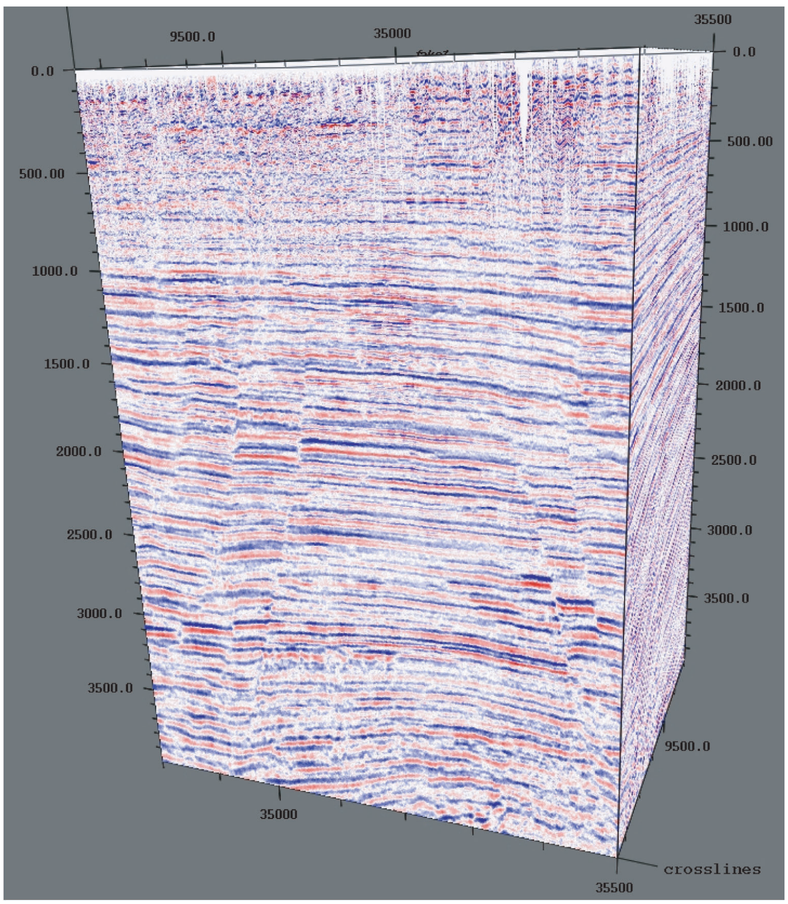

Figure 5. Opaque seismic cube used to generate stereo pair in Figure 4.

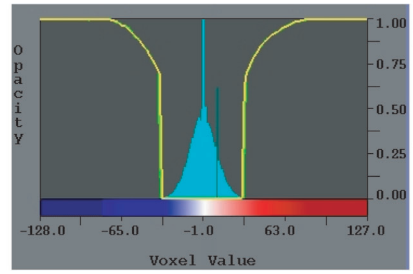

Figure 6. Opacity function used to convert figure 5 to figure 4.

erations of immersive environments. Natural motion is yet another degree of realism that improves our ability to quickly understand the data volume. We will return to this later.

Importance of stereo. Geologists, cartographers, and photo analysts have long used stereo pairs of aerial photographs to map data on the surface of the earth. A single properly generated stereo image can convey more information than a short movie loop or animation. A stereo movie or (better yet) head-tracked stereo images provides even more information than a single stereo image. In an article such as this, it is hard to convey the importance of either motion or stereo. Figure 4 contains a stereo image pair and is an attempt to show the advantage of stereo. When the images are viewed in stereo, it is much easier to determine the relative position of the bright spot and the yellow marker on the well path.

Stereo is a strong cue to depth. However, it is not effective for all users. Nearly 10% of the population has a reduced capacity to interpret stereoscopic displays (stereo blindness). Producing a stereo display that is correct and comfortable requires attention to a number of details, such as refresh rate, projection method, user position, and user eye separation. If these details are not properly addressed, the display may appear distorted or cause user discomfort, ranging from headaches to nausea and "simulator sickness." A correctly calibrated stereo image is an invaluable tool.

Smart data reduction. "Smart data reduction" can significantly improve your ability to understand your data. To do this you need to know what is "signal" for your objective. Everything else is noise—so remove it. As an example, we will use an image of an opaque cube of seismic data (Figure 5). If we are interested in looking for bright spots, we could apply an opacity curve to the data (e.g., Figure 6) to remove all data points that do not fit our bright-spot criteria. The results of applying this opacity curve were shown in Figure 4. Now several bright spots are readily apparent. However, using a single image (e.g., either the left or right image in Figure 4), it is hard to determine exactly where in the cube the bright spots are located. To properly perceive the relative positions, you might use several images or a movie loop. Even better would be to use a stereo pair (e.g., the pair in Figure 4) or a stereo view with animation. (For another example of "smart data reduction" see Interpreter's Corner in the July 1996 *TLE*).

Image distortion. Image distortion will be an issue in building any immersive environment. Using technology available today, an environ-

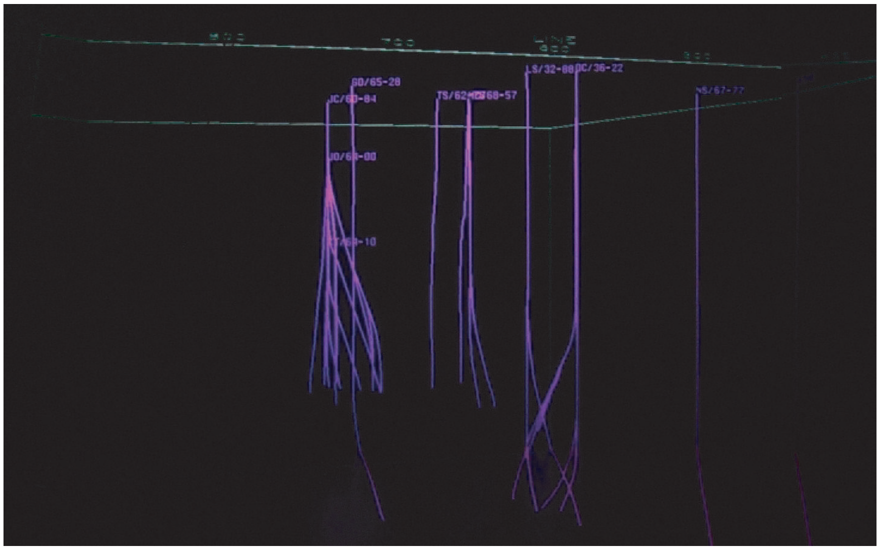

Figure 7. Well paths as seen from "sweet spot" outside of our second generation IVE.

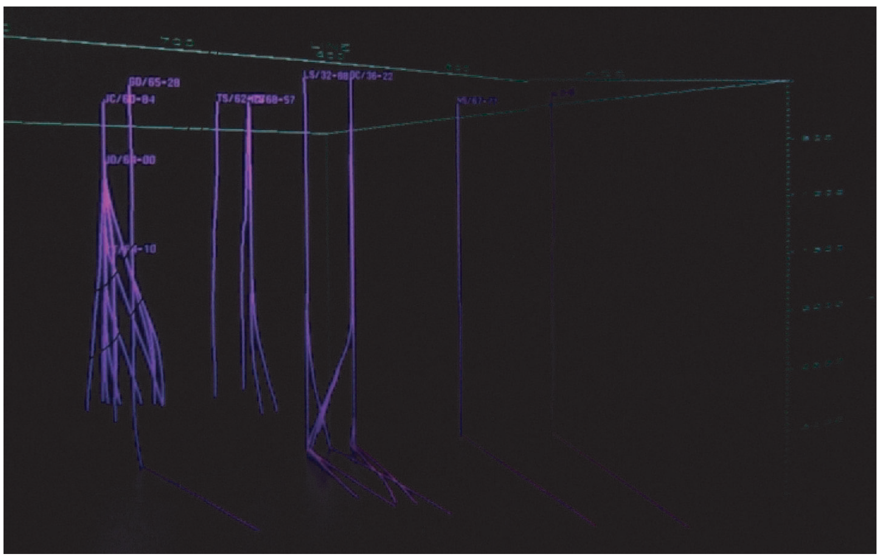

Figure 8. Well paths as seen from about 10 feet away from the "sweet spot" to exaggerate the amount of image distortion seen by untracked individuals.

ment is designed for one optimum viewing position at any one time. Images are generated for a pair of eyes positioned at this viewing position. As eyes move from this optimum eye position, images become distorted. The manner and amount of distortion varies depending on the screen geometry and the viewing location relative to the optimum eye position. It is important to minimize this distortion but because it is still there, other viewers should not be deceived in believing what they are looking at has the correct geometry.

Large screen immersive environments fall into two classes, those with segmented flat screens and those with curved screens. On each flat screen, straight lines will always be straight regardless of the viewing location. However, as one moves away from the optimum viewing position, the lines will kink or bend at the seams of the segmented screens. Curved screens may be a section of a cylinder, a toroid, a sphere, or some combination of these shapes. On curved screens, straight lines appear curved as the viewing location moves away from the optimum viewing position. The curvature increases the farther the viewing location moves away from the optimum position and increases with degree of screen curvature. Because there is no major image discontinuity, it is hard to ascertain how distorted an image you are seeing, and it is therefore easier to make a false interpretation.

In both of ARCO's immersive visualization environments, we used flat, rear-projected screens rather than curved, front-projected screens. A flat screen provides superior immersion for the primary viewer and subjects other viewers to less image distortion. Because our primary user is head-tracked, the optimum viewing position is continually changed, meaning he or she always gets a distortion-free image. Because image distortions for viewers away from the optimum position are concentrated on the seams, they can be detected and are less likely to lead to any false conclusions.

Figures 7 and 8 exaggerate the distortion seen in the closed configuration of our second-generation IVE. The only difference between these figures is the position of the camera—not the images on the screens. In Figure 7 the camera shows the image seen by the head-tracked person (except that this is a mono image instead of stereo). In other words, the camera was at the optimum viewing position, or "sweet spot." Kinks or bends observed in the well paths are due to the well paths themselves, not the camera's (or viewer's) position relative to the generated image. Where the wall screens meet the floor is not readily apparent in this image. The camera was moved a distance away from the sweet spot for Figure 8. This represents what you would see if you were ten or more feet from the head-tracked person. Now the location of the floor and walls is apparent by the "kinks" or bends in the wells. We used a much larger camera separation between Figures 7 and 8 than is typical in a normal work session in order to emphasize the nature of the image distortion. A group working together in a head-tracked immersive environment tends to stay together, thus minimizing the geometric distortion.

We view this localized distortion as an advantage of segmented screen environments. Normally, more than one person is in the environment at a time and, therefore most people are seeing a distorted view of the data. All they need to do is look at a corner or a floor/wall seam to estimate the degree of geometric distortion. With a curved screen system, however, no such visual cue tells observers not at the optimum viewing position that they are seeing a distorted image. The distortion in a curved screen system is distributed over the image. In applications where geometrically accurate perception is important, this is a critical issue.

Immersion. Immersive visualization is very effective for knowledge transfer. You can convince your partners in a few hours that you have the better well location or the right interpretation of the prospect, provided you have access

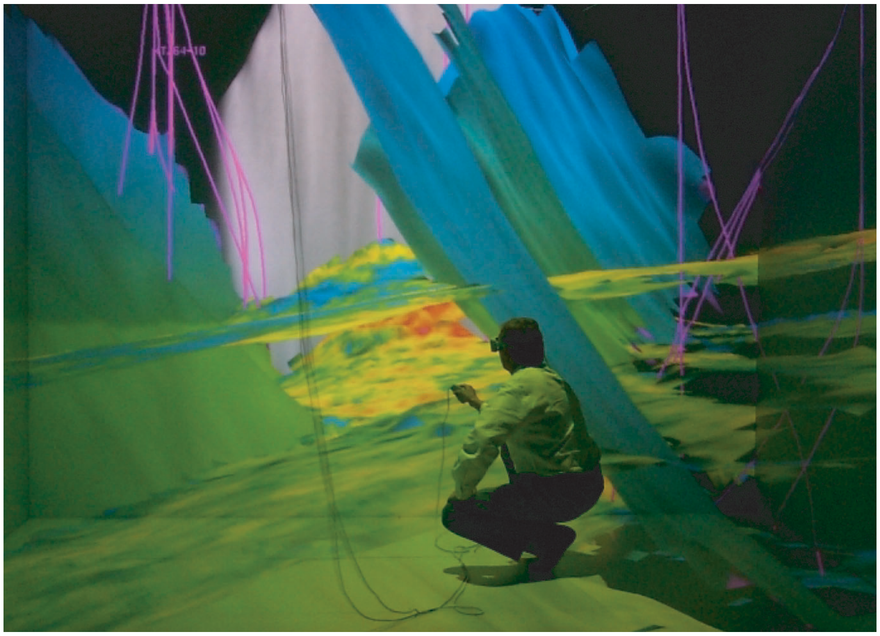

Figure 9. Example of using natural movements to investigate an interpreted 3-D volume.

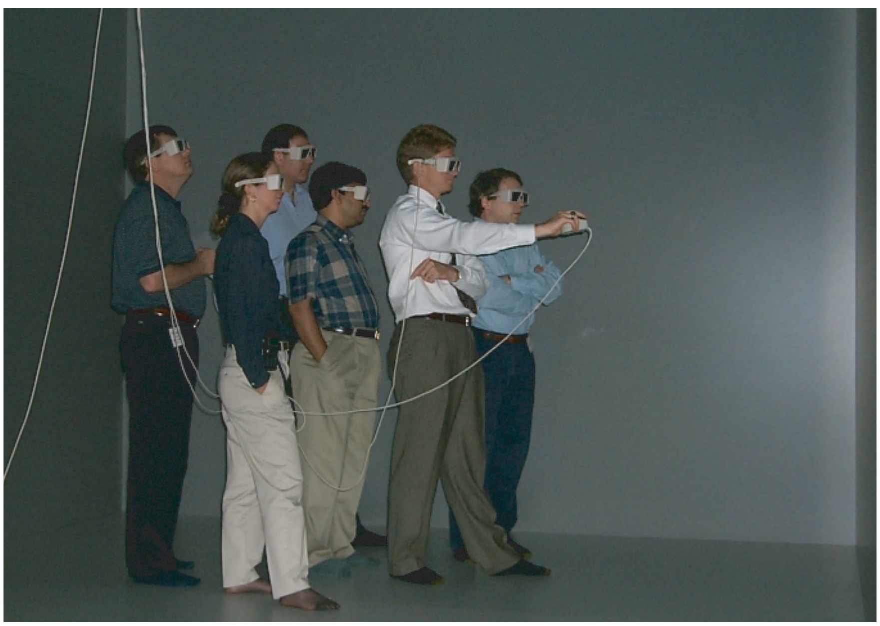

Figure 10. Example of "good schooling" in an IVE.

to the right visualization tools. Otherwise, you need to plan on days of arguing over paper maps. It is the access to these tools that allows you to choose a better well location than your partners (or competitors!).

To obtain a high degree of immersion, a visualization environment requires at least: an image that fills the field of view, off-axis head-tracked stereo images, human navigation, and the ability to reach out and touch instead of just sit back and point. Given limited computer resources, trade-offs must be made among these features. Finding an effective trade-off of these "features" guided our work in building both immersive visualization environments.

First, true immersion is achieved when the entire field of view of the observer is filled by the image—both the field of focus and the peripheral vision. Because the user is moving around, we need to be sure that as much as possible, peripheral vision remains filled with data from the virtual environment. In our case, this means having 10-foot-high screens to keep the edge of the screens out of the user's field of view. Also, the mirror for the floor projection is angled such that when the users look up, they cannot see their reflection. You will get more information out of a large-screen system (flat or curved screen) than out of a desktop system simply because it fills the entire field of view.

Second, displayed images must be in stereo. Smooth motion of the data relative to the viewer can be used to give an illusion of immersion. This works for mono images due to the differential parallax of the various objects in the scene. Once the motion stops, the relative distance of objects from the viewer is lost, and the illusion of immersion is gone. This does not happen with stereo images. With a single stereo image, the user can change their gaze and study various parts of the image or scene to determine relative positions of different objects or geologic features. Therefore, a single stereo image can provide much more information than a small movie loop and may represent less computational or display load on the computer. By combining the ability to track the viewer's position with stereo imaging, you get more out of your physical visualization environment because it allows you to have both an "inside looking out" view and an "outside looking in" view.

Third, we want to use normal human motion to change our perspective of the data. This implies that the location of the user's head is tracked. Every time the user's head moves, a new set of images is generated so he or she sees what should be seen from the new location. This is human-controlled parallax instead of a precomputed parallax provided by a movie loop. If we are walking around a fault or a well, we want it to appear stationary as we move; there should be a one-to-one relationship between the user's motion and the motion of the virtual environment. Of course, we also need some interactions to move the data or ourselves volume rapidly from one place to another. In order to most effectively use human motion to navigate, we must rear-project images onto the screens so we will not cast a shadow on the part of the data we are interested in observing.

This requires a fast computer, or a simple image, to keep the frame rate high enough so that the user does not see a time lag between head movement and generation of a new image. One advantage of a large screen (versus a head-mounted display) is that a larger amount of lag between moving and image update is acceptable. There is a trade-off between image complexity or details and motion. It is possible to move the data in low resolution and then freeze the location temporarily while the image is put into a higher resolution. Alternatively, a small region of interest can be put in high resolution with minimal affect on the overall frame rate.

Fourth, the environment is a working environment, which means we need to be able to interact with the data. We need to modify such things as our interpretation and well paths. To do this effectively requires the use of a 3-D stereo cursor on stereo images that remain stationary relative to the user. It is very hard to pick or interpret a moving surface or data cube.

Display interaction. The more intuitive the method of controlling the visualization, the more easily the interpreter can manipulate the data to extract meaning. One major advantage of head-tracked immersive environments is that they allow interaction with the virtual environment that is very similar to daily interaction with the real 3-D world. Unfortunately, not all interaction in the immersive environment is this straightforward. At some level, most human-machine interfacings require training. As an example, two widely used industry visualization packages, VoxelGeo and GeoViz, have different navigation conventions. In one you move the mouse away from you to zoom (as if you were flying through the data) and in the other you pull the mouse toward you (as if you were holding the object). Users quickly adapt to one method or the other. Both methods can be argued to be somewhat intuitive but nonetheless, both require user familiarity for effective use.

Cave etiquette. In working in an immersive environment, the manner in which the collaborators interact greatly affects the results and value of the interaction. Interaction between an individual and the group needs to change slightly when working in an immersive environment. Numerous sessions in our immersive facilities and a few sessions in immersive facilities at the University of Houston, Norsk Hydro, and GMD (National German Research Center) have led us to form a few "rules" of etiquette.

Don't point with your finger. Whenever you want to discuss a particular feature of interest to you with others in the group, don't point at it with your finger or any other part of your body. You will have the group looking at the wrong object. Use a laser pointer. Objects such as a well will appear to be at a particular location in the room but the apparent location is slightly different for everyone in the room due to the way stereo images are generated. Two images are generated, one for each eye. A laser pointer will allow you to highlight one of those images on the CAVE wall and everyone should be able to unambiguously determine what you are interested in. This applies to both the "navigator" and the observers.

Move slowly and smoothly—but do move. This particularly applies to the navigator. As mentioned earlier, in a head-tracked immersive environment, computer images are generated for the navigator. When the navigator moves, the images change to match the new viewing location. If the navigator moves quickly, the images will (hopefully!) change quickly. As a navigator, moving your head quickly between two points is a good way (from stereo motion parallax) to determine relative distance between objects. It is also a potentially good way to make your observers cybersick—so use it carefully. The navigator should also move around the room to get different perspectives on the data. He or she should not stand in one place. For example, notice in Figure 9 how Don Dean of ARCO has knelt down below the fault so he can look between two horizons to investigate the relationship of the bright spot on the lower horizon and the interpreted salt interface. If the computer has a slow frame rate, the navigator should move slowly to avoid getting ahead of the image calculations.

Watch your frame rate. Normally, it is thought that the higher the frame rate the better. This is not always true. It depends on what is being done in the environment. In general, stereo frame rates need to exceed about 5-7 frames per second. When the rate drops below this, it is hard for the navigator to move slowly enough to prevent the images from "jumping." If very high frame rates are obtainable (>30 fps), it is better to add more data to the scene and slow it down some to ease up on the cybersickness potential of the untracked observers.

Practice good schooling. Figure 10 shows the proper way a group should interact in a CAVE. Note they are close together and facing the same direction. They will move around the room slowly as a group, acting like a calm school of fish. An observer's head should be as close as possible to the navigator's head to minimize image distortion. If the navigator turns his/her head, everyone in the room should do the same. The computer tracks the location of both eyes, so for

Visualization:

- Is in part a learned skill.
- Everyone has a way in which they view the world.
- When a display does not fit your expected view you have trouble understanding it.
- For example – you are struggling to read this.
- You might even think it was printed incorrectly.
- You probably turned the magazine over.
- Maybe even more than once.
- If you can easily visualize things in 3–D why are you having trouble reading this?
- You can read this – but it requires more time and effort – unless it is "printed correctly".
- Which is more complicated:
  A text slide or 3D seismic data?

Figure 11. Visualization is a learned skill.

proper stereo viewing, an observer's head needs to be at the same orientation as the navigator's.

Don't face each other. This is particularly important for the navigator to obey. The navigator should *never* turn around and face the observers! This is not an obvious rule of etiquette and will probably not register with you until you have had a navigator do it to you. Stereo images are based on the orientation of the tracked stereo glasses. When these glasses are facing the observers, their left eye will see the image designed for their right and vice versa—which is not very comfortable.

Share your glasses. The best experience in an immersive environment is when you wear the tracked glasses—being the navigator. Everyone should take the opportunity to navigate through the data. In fact, swapping the tracked glasses along with its associated wires should be treated as a ritual.

Don't touch the walls. The walls in immersive environments are rear-projection screens, and image quality degrades every time someone adds fingerprints to them. Take care of your projection surfaces.

Wear clean socks. The floor is also a projection surface. In order to protect it, users are usually requested to remove their shoes or to cover their shoes with protective slippers. Although floors are now generally made of more durable surface material (e.g., formica) than the floors of five years ago (painted wood), they still need to be clean for good image quality.

Closing remarks. Look again at Figure 11, and ask yourself how easily you can orient, manipulate, and display your data so you can easily "read" it? Can you interact with the data in the ways you have trained yourself to? Can you be retrained? Do your displays have a sufficient level of realism so that you quickly understand them? Have you used stereo to better understand the 3-D nature of your data? Have you removed the parts of the data currently not of interest, thus effectively increasing your signal-to-noise ratio? From where you are looking at the data, how distorted is it? When working in an immersive environment, are you practicing good etiquette?

Answering these questions properly will allow you to get the most from your data with the least effort.

There is a variety of different spatially immersive or semi-immersive environments. The reconfigurable flat-screen system allows the user to move around and explore the data from both inside out and outside in. This requires rear-projected screens and the ability to head-track a user. It satisfies, in a single system, the needs of fully immersive work sessions with small teams and larger semi-immersive presentation sessions. Nontracked front-projected curved screen systems are better suited for the more traditional outside-in data viewing and analysis, providing primarily a large "desktop." Although they support collaboration among large groups, they do not support fully immersive work sessions.

At this point, many companies have experienced the value of visualization through increased efficiency, shortened cycle time, and improved communication. An attendant increase in integration and collaboration has improved the accuracy and completeness of the end product. Visualization can produce tremendous value in a wide variety of applications and can compress months of effort into weeks. Large-scale immersive environments offer the greatest benefit. LE

Acknowledgments: We thank ARCO for allowing us to publish this paper. Many members of ARCO management provided critical support for the visualization program in the late 1990s. These include David S. Campbell, Barry Davis, Dodd DeCamp, Jim O'Connell, and Nancy Sisler.

Corresponding author: tstark3@ibm.net