Integrating True 3-D Digital Video Effects Processing with Real-Time Computer Graphics

By Mark Podesla and David E. Acker

A system has been developed that allows for real-time modeling of true 3-D shapes with live video mapping in an economical digital video effects (DVE) environment. By dividing a video image into a variable number of polygonal patches and manipulating each polygonal patch individually, almost any geometric shape can be modeled and video-mapped in real time. The technique is unlike any conventional video effects processor and is in certain ways more closely related to the computer graphics world of three-dimensional (3-D) modeling and animation, but without the time and cost constraints.

For several years it has been customary to refer to DVE machines with rotation and perspective as 3-D systems. However, with one or two exceptions, these machines only create the illusion of 3-D by distorting flat pictures, limiting their realism to certain angles. A proprietary polygonal patch-based processing technique that models and manipulates true 3-D objects in real time and then performs live video mapping onto these shapes has been developed. This article describes that technique and the Microtime products that employ it — the IMPACT Series line of true 3-D DVE devices.

Traditional Digital Video Effects Processing

Before exploring the polygonal patch process, it would be helpful to review the fundamentals of traditional DVE devices and build upon those basic concepts. In Fig. 1, the traditional DVE processor takes a normal rectangular input picture (ABCD) and converts it into a nonrectangular output picture (A'B'C'D'). The amount and nature of the manipulation is a function of how the picture is being controlled. Size and position changes alone will result in the output remaining rectangular, but changes in rotation will cause some areas of the output picture to become enlarged because they are closer to the viewer and some areas to become compressed because they are further away from the viewer. In this example, the picture itself has been foreshortened to give the appearance that it has been laid back in 3-D space. The area of the source picture around A is manipulated to A', B to B', C to C', D to D', and so on, for all areas of the picture. The edge D'C' is closer to the viewer than edge A'B', and is therefore larger.

When all or part of the picture is either expanded or compressed from its normal size, filtering and interpolation must be performed to strike the optimum balance between resolution and aliasing artifacts. Ideally, the compromise should be a continuously variable selection throughout the picture, depending on the actual amount of expansion or compression at each point in the output video image.

Presented at the 135th SMPTE Technical Conference (paper no. 135-84) in Los Angeles on November 2, 1993. Mark Podesla (who read the paper) is with Microtime, Bloomfield, CT 06002. David E. Acker, former Microtime president, is an industry consultant.

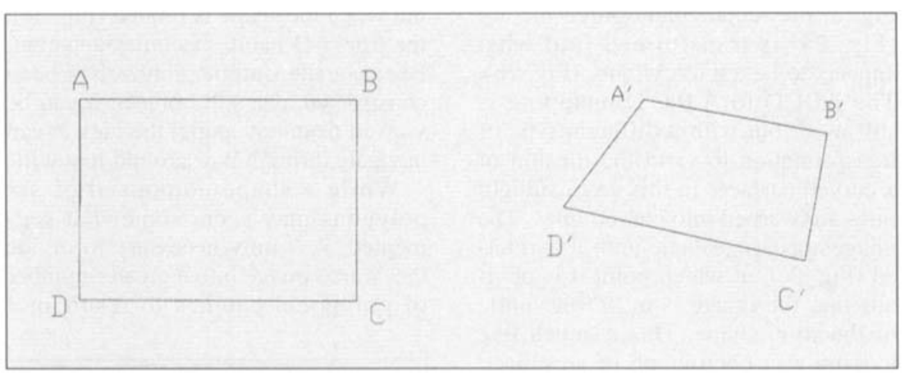

Figure 1. Traditional DVE processing.

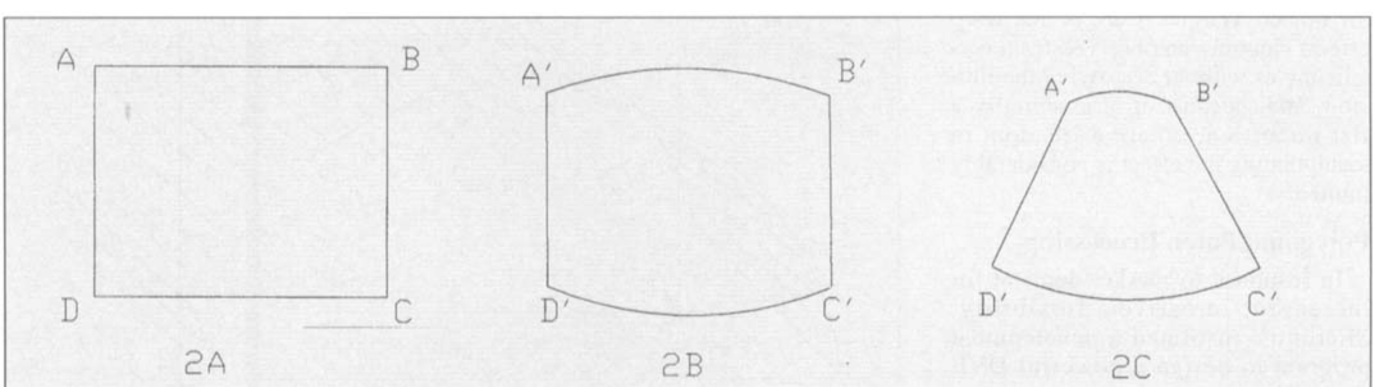

Figure 2. Traditional DVE processing with warp.

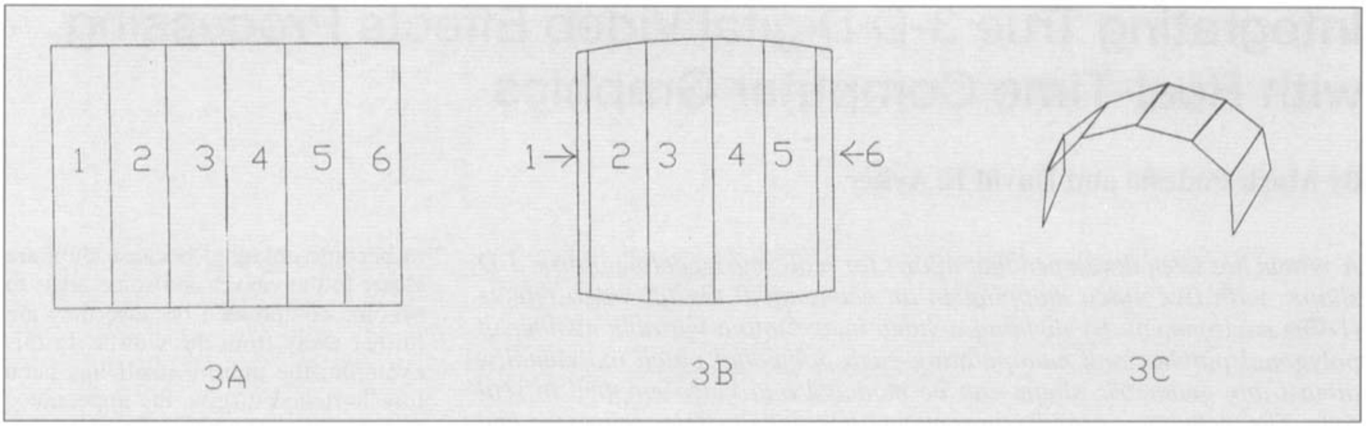

Figure 3. True 3-D polygonal patch processing.

Calculation of the manipulation parameters and the appropriate filtering and interpolation must be computed at up to pixel clock rates to allow real-time operation on live video. This is achieved by the extensive use of specialized digital signal processing hardware. The Ampex Corp. was an early pioneer in this technology with their Ampex Digital Optics (ADO) product line.

Warp Effects

The next step in the evolution of traditional digital video effects was the addition of the so-called "warp." In Fig. 2, the rectangular source picture (Fig. 2a) is transformed into what appears to be a half cylinder (Fig. 2b). The ABCD to A'B'C'D' mapping is still used, but with a different type of transformation to yield the illusion of a curved surface. In this case, straight lines are warped into curved lines. The image appears realistic until it is rotated (Fig. 2c), at which point it is obvious that the image is a 2-D facsimile of the solid shape. This is much like looking at a photograph of an object; turning the photo over does not allow the viewer to look at the side or rear of the object. The net result is that warp effects can only be observed from certain angles without destroying the illusion, and, because it is essentially a flat distortion, creative freedom in manipulating the effect is considerably limited.

Polygonal Patch Processing

In response to market demand for increased creative flexibility, Microtime instituted a development program to design a powerful DVE device that would model and manipulate true 3-D objects and then map video onto those shapes, all in real time and at a cost approaching the conventional flat 3-D/warp machines. Such a system would exploit recent developments in microprocessor technology and essentially integrate many of the characteristics of computer-graphics modeling with real-time video effects processing.

Returning to the half cylinder example, assume that the input picture is divided into a small number — six, for instance — of polygons (Fig. 3a). The result of the initial transformation (Fig. 3b) is roughly the same in appearance as the "warp" approach, but when the shape is rotated (Fig. 3c), the true 3-D nature becomes apparent. Because the output picture has been constructed as a solid object, it can be viewed from any angle; the viewer can navigate through it or around it at will.

While a shape composed of six polygons may seem somewhat segmented, it is only necessary to divide the source image into a greater number of polygonal patches to result in a smoother or more continuous output shape. Figure 4 shows the increased shape resolution that occurs when using 20 or so polygonal patches. Generally, when a greater number of patches is used, the object surface will be smoother, the shape control will be finer, and the shape resolution will be better.

Basic Principles of Polygonal Patch Processing

The mathematics and hardware needed to fly a single rectangular picture in 3-D space have been understood for several years. Figure 3 reveals that the six polygonal patches outlined are rectangular on input and nonrectangular on output. For each individual polygonal patch, the input-to-output transformation method is identical to that of Fig. 1. Polygonal patch 1 is transformed into polygonal patch 1', polygonal patch 2 to 2', polygonal patch 3 to 3', etc. Each patch is manipulated with slightly different size, position, and rotation para-

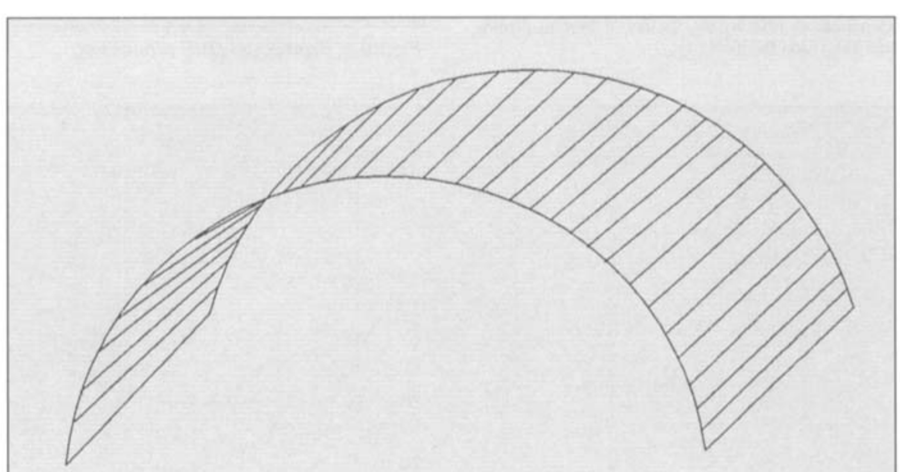
Figure 4. Adding more polygonal patches.

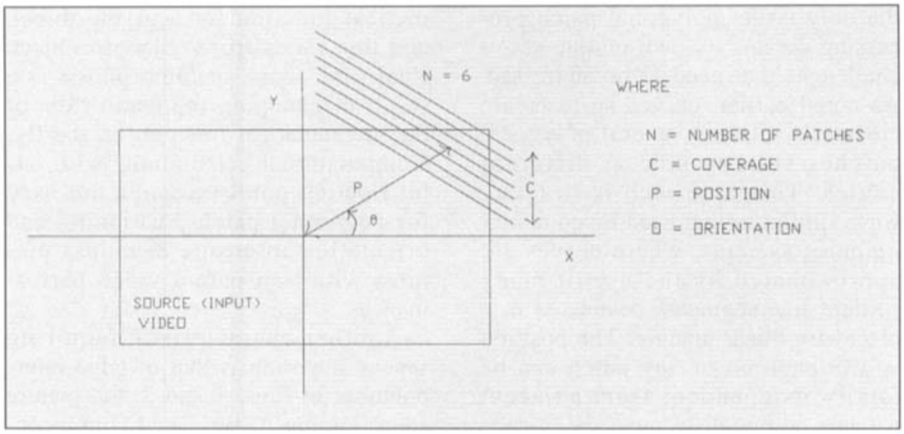

Figure 5. Four basic polygonal patch control parameters.

meters through its own individual transformation matrix to create the solid object. The manipulation parameters vary in an ordered fashion from patch to patch to ensure that all pieces remain connected to each other and create the desired true 3-D shape. Therefore, in this example, polygonal patch processing only requires that the machine be able to process six times in parallel that which the conventional flat DVE machine does: fly rectangular pictures in 3-D space.

By varying four basic polygonal patch operating parameters, numerous true 3-D shapes (in addition to cylinders) can be created; examples include page turns, spirals, waves, and Venetian blinds. These control parameters are number (of polygonal patches), orientation, coverage, and position (Fig. 5). All four are "keyframe-able" in that they can be changed on a keyframe-by-keyframe basis over the course of the effect.

The number parameter simply specifies how many polygonal patches are used to compose the desired image. Figures 3 and 4 demonstrate how changing the number of polygonal patches affects the shape resolution. It should be noted that a larger number of polygonal patches is not always desirable; for example, in the case of a cube or beveled picture frame, a smaller number of polygonal patches is actually more visually pleasing because the angular nature of the shape is accented.

The effect of the orientation parameter on polygonal patches is best demonstrated by a page turn (Fig. 6). Changing this parameter creates the effect of flipping up one of the various corners of a page. (The coverage, position, and number variables are, in this case, set to a specific value and held constant.)

When position is varied (with the other three parameters fixed) the video is pulled over until it is upside down and sliding off the screen (Fig. 7). Coverage controls the area over which the polygonal patches are dispersed (Fig. 8).

If applied to nonparallel, angular, and/or disjointed polygonal patches, many numerous other shapes including (but not limited to) cones, cubes, beveled picture frames, and shattering tiles can be created. The conventional DVE manipulative primitives of rotation, position, and size are all available to change the viewer's vantage point in addition to the polygonal patch-based parameters. Any combination of parameters can be selected to change simultaneously from keyframe to keyframe.

Project Goals

After a number of market surveys, preliminary tests, and computer simulations, the Microtime design team began work on a new-generation DVE device incorporating the polygonal patch architecture. Designated by the name IMPACT, the design goals were as follows:

- Real-time, on-line, true 3-D shape manipulation with live video mapping (in contrast to off-line rendering approaches).
- Curved, angular, and surface-of-revolution true 3-D objects.
- Traditional DVE effects (including light sourcing).
- "Read-side" (target pixel-to-source pixel) processing for optimal video quality and expansion effects.
- Upgradable software to accommodate the latest fashionable effects.
- Intuitive operator interface (i.e., no special programming knowledge required).
- Three video inputs to emulate multichannel operation.

It was concluded that a large number of useful shapes, including realistic page turns; cubes; beveled picture frames; intersecting video planes; exploding tiles; and a wide range of push, pull, and wipe transitions could be created with as little as eight polygonal patches. With as many as 32 patches, a wide range of smooth surface effects such as cylinders and waves were possible. As a result, the design goals were refined to create a family of upgradeable machines — the IMPACT Series — that supported between 8 and 32 polygonal patches. Video performance and basic features are identical for each version, and machines can be field-upgraded to add more polygonal patches. The only differences between versions are the library of true 3-D shapes and the shape control flexibility.

The major challenge was to make the 32-polygon version a marketable,

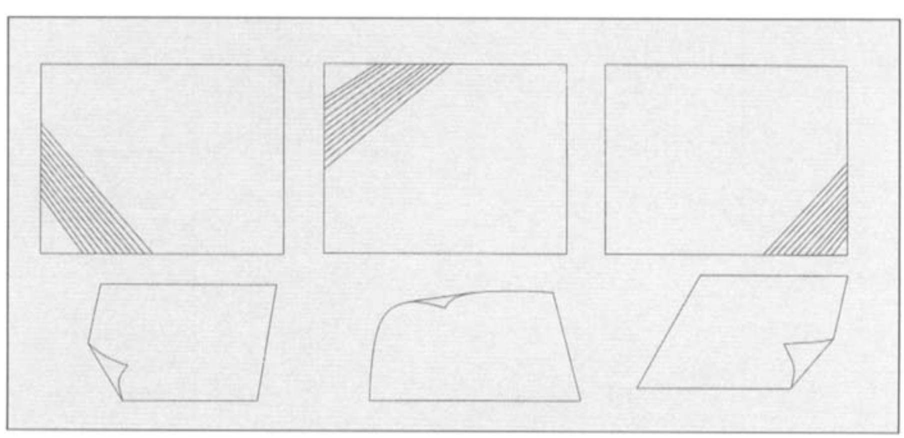

Figure 6. Changing polygonal patch orientation.

economical package. With real-time video, new fields of information arrive 60 times/second (50 in PAL) and new pixel information arrives every 74 nsec, thus fixing the time available for computation. The polygonal patch processing method requires that up to 32 times more calculations be performed in these fixed time slots than in a conventional DVE machine. The Intel i860 microprocessor — a chip ideally suited to graphics number-crunching tasks — was able to handle the computational load for 32 polygons when run at 33 MHz, provided that it was supported by a few other peripheral processors for housekeeping, data communications, operator display, etc.

Computational horsepower was not the only issue; polygonal patch processing creates its own unique set of challenges that need to be addressed. As noted earlier, curved surfaces are created by placing several polygonal patches side by side at different angles. This approach is in many ways similar to that used by computer graphics systems, where curves are approximated by the use of many straight line segments connected in a piecewise linear manner. The position and orientation of any patch can be totally independent from adjacent patches, although in most cases each remains closely related to its neighbors to avoid undesirable gaps or overlaps at the seams. The character of the relationship between the polygonal patches is defined by the mathematical function for a given object, and this varies from object to object. Real-time shape metamorphosis is a result of changing the parameters of the mathematical function on the fly. Computational calculations with 64-bit floating point precision are used for polygonal patch positioning and orientation to ensure seamless pictures with transparent video performance.

Another characteristic requiring special attention is that of edge interpolation. In Figs. 1 and 2, the picture edges joining A, B, C, and D in a conventional DVE machine always remain linked relative to each other, and there are always four of them. Therefore, the normal interior interpolation that is used within the active

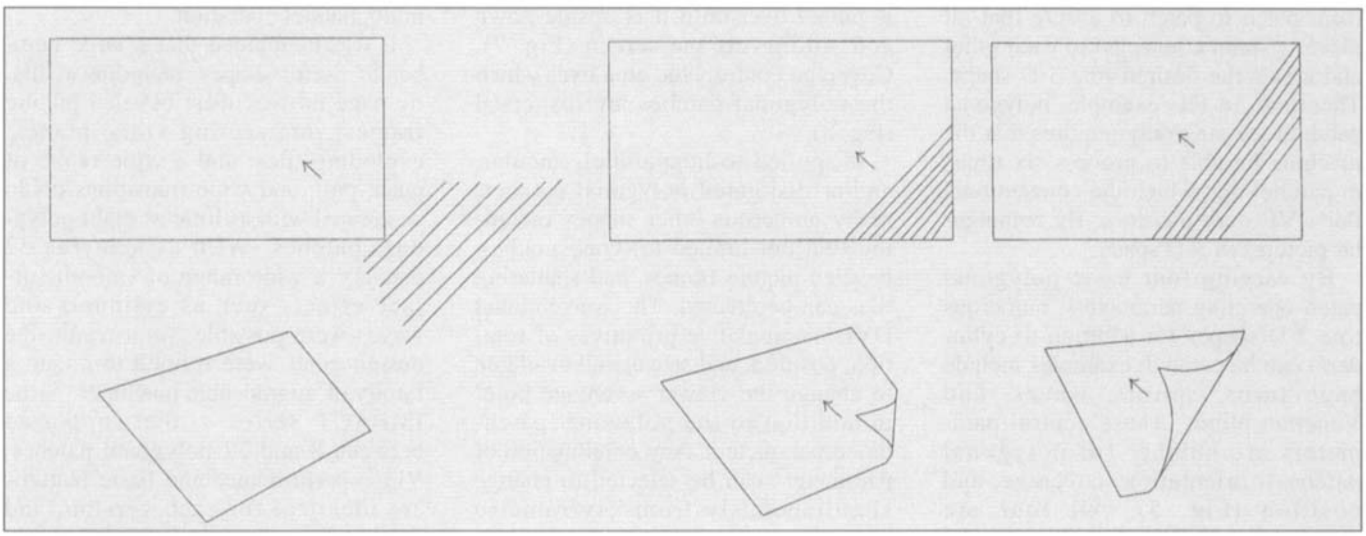

Figure 7. Changing polygonal patch position.

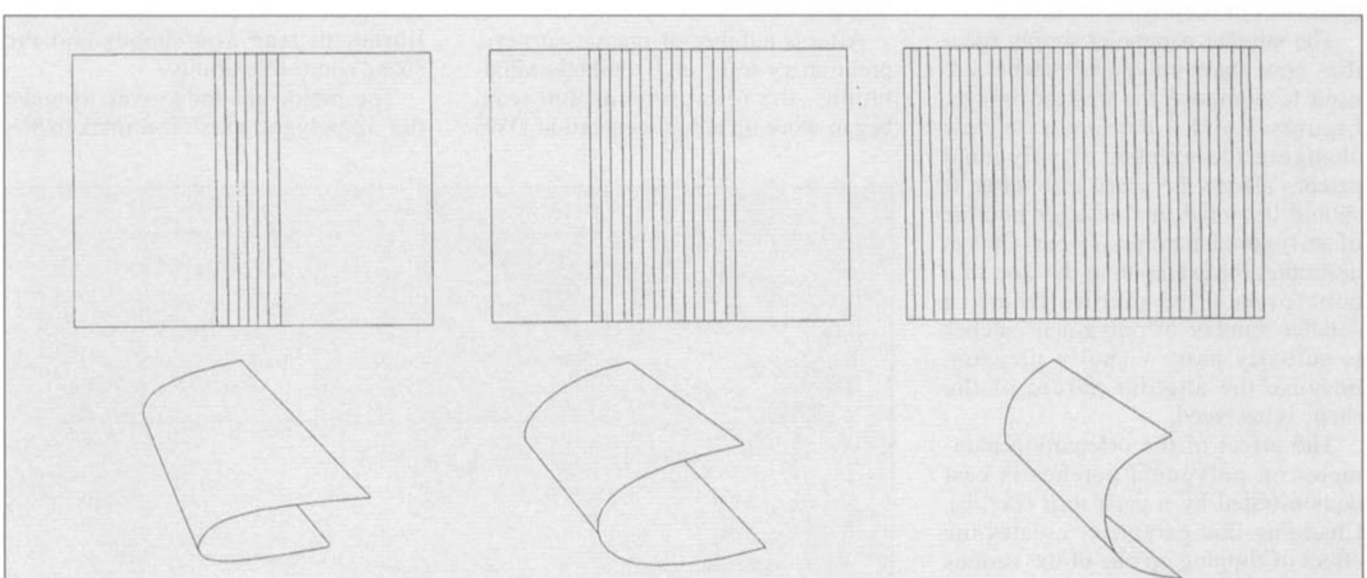

Figure 8. Changing polygonal patch coverage.

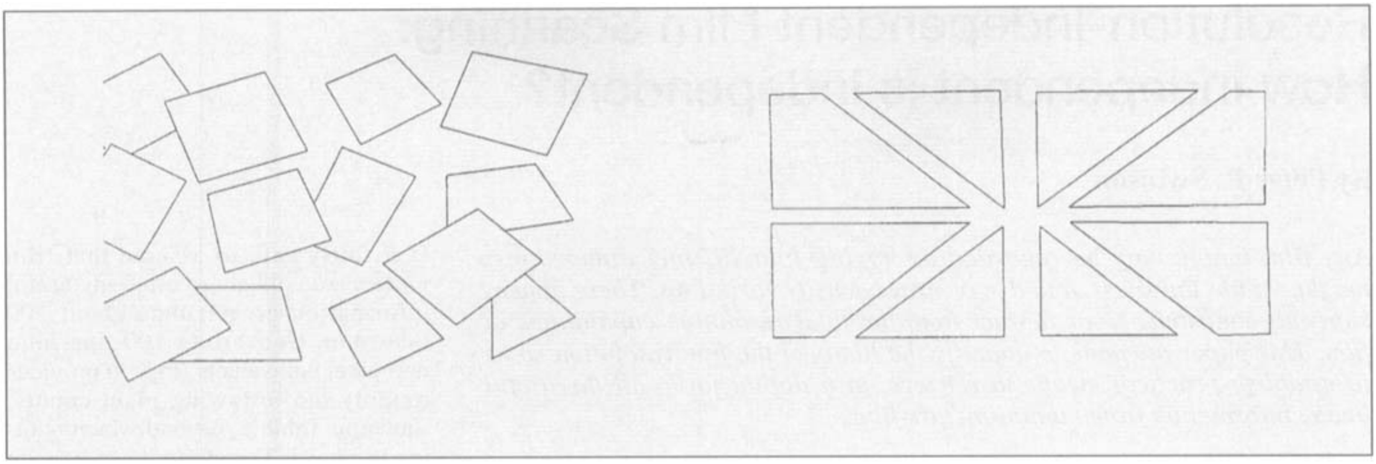

Figure 9. Multiple edges.

picture area can operate from picture edge to picture edge. This minimizes the effect of "jaggies" on the four edges of the output picture. With polygonal patch processing, the situation is a little more complex in that new edges can appear from nowhere as the shape breaks apart. If the picture is broken into tiles or shards (Fig. 9), the total number of edges increases with the number of picture pieces and can vary continuously. Since shape metamorphosis can be controlled in real time under operator control, or during an effects sequence, the number of edges and their locations are dynamic variables which require special handling.

A second set of edge interpolators (in addition to the normal video image content interpolators) is dedicated to eliminating the "jaggies" or "stairstepping" on however many edges exist at any point in time. These interpolators operate on both interior "fold-over" edges and outside perimeter edges.

With any true 3-D object, the polygonal patches on the back are sometimes partially or fully obscured by those on the front. The amount hidden varies from patch to patch and from field to field as the object moves; a back patch becomes a front patch as soon as the object is rotated on its axis. Figure 9 illustrates the unique video prioritization with an object that appears like floor tiles. Many polygonal patches are stacked behind each other in this particular orientation and must be prioritized for every other possible orientation of the object. Ensuring that the correct parts of each polygonal patch are seen by the viewer is analogous to removing hidden surfaces in a computer graphics software program, except that a DVE system must perform this operation in real time. A unique depth priority algorithm was developed to resolve this issue.

Multichannel Effects

To complement its polygonal patch architecture, the IMPACT product line employs some unique video processing characteristics. Up to three video inputs are simultaneously available for mapping onto any true 3-D object. Each video input can be separately bordered, frozen, or adjusted in video levels and have such effects as posterize, solarize, and strobe added. This feature allows IMPACT to create complex effects in a single pass on a single channel, whereas multiple passes on tape or multiple effects channels might otherwise be required.

IMPACT always keeps track of the front side and back side of each object, and each of the polygonal patches used to subdivide the input can be mapped with a selected front-side video source and a different back-side video source. Using a page turn as an example, this means that two different video sources can be simultaneously displayed, one on the front side and another on the back side. As the page turn progresses, more and more polygonal patches become the back side instead of the front, revealing more and more of the second source. If the video on the back of the page contains text, it can also be inverted horizontally to ensure that it is "right-readable" at all times.

Perfectly aligned cubes and boxes are available at the touch of a button, eliminating the tedious entry of data or trial and error methods. Each of the three sources is cropped and mapped to one pair of six polygonal patches (each opposite face has the same source), and then each polygonal patch is positioned and oriented by the system to create a perfect cube. Three of the six faces are visible at all times. Modifying the X, Y, or Z size turns the cube into an unlimited number of asymmetrical boxes, and the faces can be flown apart with all six faces visible — a move that would require six channels of a conventional effects machine or six passes on tape.

Each input can also be mapped onto one of three planes, which can then be positioned and rotated either independently or globally as a group. If the planes intersect each other, the system automatically takes care of hidden surface removal.

Conclusion

Polygonal patch processing yields three significant benefits when compared to conventional DVE systems:

- The "look" of true 3-D images
- The immediacy of on-line shape selection and control
- The ability to emulate multichannel effects in a single-channel environment

This technique, and the resulting IMPACT system, significantly enhances "creative" flexibility with its revolutionary image manipulating power and numerous price/performance advantages. Polygonal patch processing opens up a new dimension in video special effects with its practical method of mapping live video onto variable true 3-D objects, providing a look that integrates the power and versatility of computer graphics modeling with real-time DVE processing.