Motion Quality in Simulator Imagery: Some Effects of Resolution

Julie Mapes Lindholm	Angélique A. Scharine	Byron J. Pierce
Lockheed Martin Systems Management	Army Research Laboratory	Air Force Research Laboratory
Mesa, AZ	Aberdeen Proving Ground, MD	Mesa, AZ
Julie.Lindholm@mesa.afmc.af.mil	Angelique.Scharine@us.army.mil	Byron.Pierce@mesa.afmc.af.mil

ABSTRACT

The level of detail in flight-simulator imagery depends upon the resolutions of the database, the image generator (IG), and the display. In response to the need for greater detail, resources are being devoted to increasing the spatial resolution of each of these system components. Next-generation flight-simulator visual systems will thus be capable of representing smaller environmental features and of representing a given feature at a greater distance. However, flight-simulator imagery is more than a sequence of static, spatial images. During simulated flight, the system creates a three-dimensional (two dimensions of space, one of time) space-time image that approximates the *continuous* changes in the spatial image that would result if the pilot were to actually fly through the synthetic environment. The quality of such space-time images depends upon the temporal as well as the spatial characteristics of the IG and the display.

Here we (a) discuss how database and IG resolutions and simulated-flight speed and altitude affect the temporal frequencies in an image and thus the extent of temporal aliasing likely in simulator imagery, (b) describe effects of a display system's spatial and temporal resolution on the spatiotemporal-frequency spectrum of a display image, (c) summarize characteristics of the human visual system relevant to spatiotemporal-frequency and motion perception, and (d) report preliminary results of a research project in which we are examining the effects of image resolution on perceived-motion quality during simulated flight.

ABOUT THE AUTHORS

Julie Mapes Lindholm is a senior scientist with Lockheed Martin Systems Management at the Air Force Research Laboratory in Mesa, Arizona. Her research focuses on mechanisms of form and motion perception, particularly as they pertain to image generation and display technologies. She received the B.A. degree in psychology from the University of Redlands and the M.S. and Ph.D. degrees in experimental psychology from the University of Wisconsin-Madison.

Angélique A. Scharine is a research psychologist for the Army Research Laboratory/Human Research and Engineering Directorate in Aberdeen, MD. Her research interests are auditory localization and visual perception. She recently completed a Ph.D. in cognitive psychology at Arizona State University. She received degrees in educational psychology (M.S.) and math education (B.S.) from the University of Utah.

Byron J. Pierce is a senior research psychologist employed by the Air Force Research Laboratory, in Mesa, Arizona, where he is team lead for the visual research and development program. Dr. Pierce is responsible for the planning, management, and conduct of visual R&D activities for flight simulation and training. He was awarded a Ph.D. in experimental psychology from Arizona State University, an M.S. in engineering psychology from the University of Illinois, and a B.S. in psychology from Arizona State University.

Motion Quality in Simulator Imagery: Some Effects of Resolution

Julie Mapes Lindholm
Lockheed Martin Systems Management
Mesa, AZ
Julie.Lindholm@mesa.afmc.af.mil

Angélique A. Scharine
Army Research Laboratory
Aberdeen Proving Ground, MD
Angelique.Scharine@us.army.mil

Byron J. Pierce
Air Force Research Laboratory
Mesa, AZ
Byron.Pierce@mesa.afmc.af.mil

INTRODUCTION

In a flight simulator, the spatial resolutions of the database, the image generator (IG), and the display system all affect the sharpness and detail of the display imagery. The effects of a given system-component's resolution depend—in a variety of asymmetric ways—upon the resolutions of the other components. For example, an IG's digital representation of a given environmental location cannot contain more detail than is contained in the database, and the simulated distance at which a given detail can be represented is limited by the resolutions of the IG and the display.

In response to the need for greater detail in simulator imagery, resources are being devoted to increasing the resolutions of all the components in next-generation visual systems: Satellite images with one-meter, or better, resolution are now available for some regions, and ultra-high resolution (~5000 pixel x 4000 pixel) displays and IGs are under development. Flight simulator imagery is, however, more than a sequence of static spatial images. During simulated flight, the visual system creates a three-dimensional (two dimensions in space, one in time) space-time image that approximates the *continuous* changes in the spatial image that would result if the pilot were to actually fly through the synthetic environment. The quality of a space-time image depends upon the temporal as well as the spatial characteristics of a flight simulator's visual system.

Here we describe how (a) the content and resolution of the database, (b) the pixel density and antialiasing filters of the IG, and (c) the speed and altitude of the simulated flight jointly determine the temporal frequencies in the nominal space-time image internal to the IG. (We limit our discussion to the specific case of constant-altitude, constant-velocity, level flight over flat terrain). We also describe how the spatial and temporal resolutions of a display system affect the spatiotemporal-frequency spectrum of the display image. Finally, we provide brief descriptions of the spatial and temporal characteristics of the human visual system and an initial investigation of the effects of IG spatial resolution on the perceived extent of incoherent motion during simulated low-level flight.

COMPONENTS OF SPATIAL RESOLUTION

Database

A flight-simulator database defines the synthetic environment through which a pilot may fly. Ground and object surfaces are represented as polygon meshes. The chromoluminance pattern of a given polygonal surface is typically defined by a two-dimensional digital image called a "texture." A texture pattern can be created in a variety of ways. With current technology, it is often based on a photograph of some object or region of the world.

The individual elements (digital values) of a texture are called "texels." The resolution of a texture is defined by the texel "size"—that is, by the horizontal and vertical distance, in *environmental units* (e.g., meters), between adjacent elements in the texture pattern. In the past, a typical geospecific terrain texture had a resolution of 10 – 20 m. However, terrain textures with submeter resolution are becoming increasingly available. Moreover, with current IGs, the detail in a relatively low-resolution texture can be increased by modulating its values with those of a geotypical microtexture of arbitrarily high resolution.

Figure 1 depicts a 9.6-m and a 0.6-m resolution texture for a 153.6-m x 153.6-m terrain square. (To aid comparison, bilinear resampling was used to increase the size of the low-resolution image to that of the high-resolution image.) Note that the 0.6-m texture is sharper than, and contains features not represented in, the 9.6-m texture. Although a texture with a given resolution (e.g., 1 m) is not necessarily characterized by features that small, it is capable of representing such features.

Image Generator

In an IG, the synthetic environment is projected onto a "view plane" internal to the computer. It is useful to think of this projection as a continuous space-time image that the IG samples, in time and space, to create a sequence of two-dimensional digital images.

(a) (b)

Figure 1. Effects of texture resolution on the representation of a 153.6-m by 153.6-m terrain square. The image in Figure 1a is based on a 16 x 16, 9.6 m/texel pattern; the image in Figure 1b is based on a 256 x 256, 0.6 m/texel pattern. Bilinear resampling was used to equate the sizes of the two images.

The temporal sampling frequency is the rate at which new digital images are computed (i.e., the image update rate). The horizontal and vertical sampling frequencies are defined by the size of the pixel array. For a given field of view, the spatial sampling frequencies determine the visual angle subtended by a pixel.

A continuous image, internal to an IG, is likely to be characterized by extremely (often infinitely) high spatial frequencies. If this image is sampled directly, at the infinitesimal points corresponding to the pixel array, any two-dimensional spatial frequency component with a horizontal or vertical frequency that exceeds half the corresponding sampling frequency will be aliased—that is, it will appear in the display image as a component with a horizontal or vertical frequency less than half that sampling rate.[1] To minimize spatial aliasing, IGs implement one or more antialiasing techniques. All such techniques serve to low-pass filter the continuous image, in *image units* (c/pixel), before it is sampled to create the pixel values. Although the filters differ in quality and none is ideal, here we make the simplifying assumption that an IG's antialiasing filters successfully remove all frequencies greater than half the pixel sampling rate and pass with minimal attenuation all frequencies less than half the sampling rate.

The projected size and shape of an environmental feature depends upon the distance and angle from which it is viewed. In a computer-generated image, then, the relationship between spatial frequency in c/m and spatial frequency in c/pixel, and thus the texture resolution that can be represented, depends upon the simulated distance and altitude. This dependency is illustrated in Figures 2a and 2b, which depict 40° vertical fields of view (with 30° below the viewpoint and the up-axis of the image perpendicular to the ground surface) of a synthetic environment consisting of a blue sky and a flat, textured terrain. The simulated altitude is 300 m (~1000 ft) in Figure 2a and 30 m (~100 ft) in Figure 2b. Each half of the terrain was textured by "tiling" it with copies of a single texture pattern depicting a 614-m by 614-m region of the earth. The resolutions of the textures mapped to the left and right halves were 9.6 m and 0.6 m, respectively. To facilitate comparison, the high-resolution texture was "flipped" horizontally before it was mapped.

[1] For this statement to be consistently true, the display system must pass all spatial frequencies up to half the sampling rate.

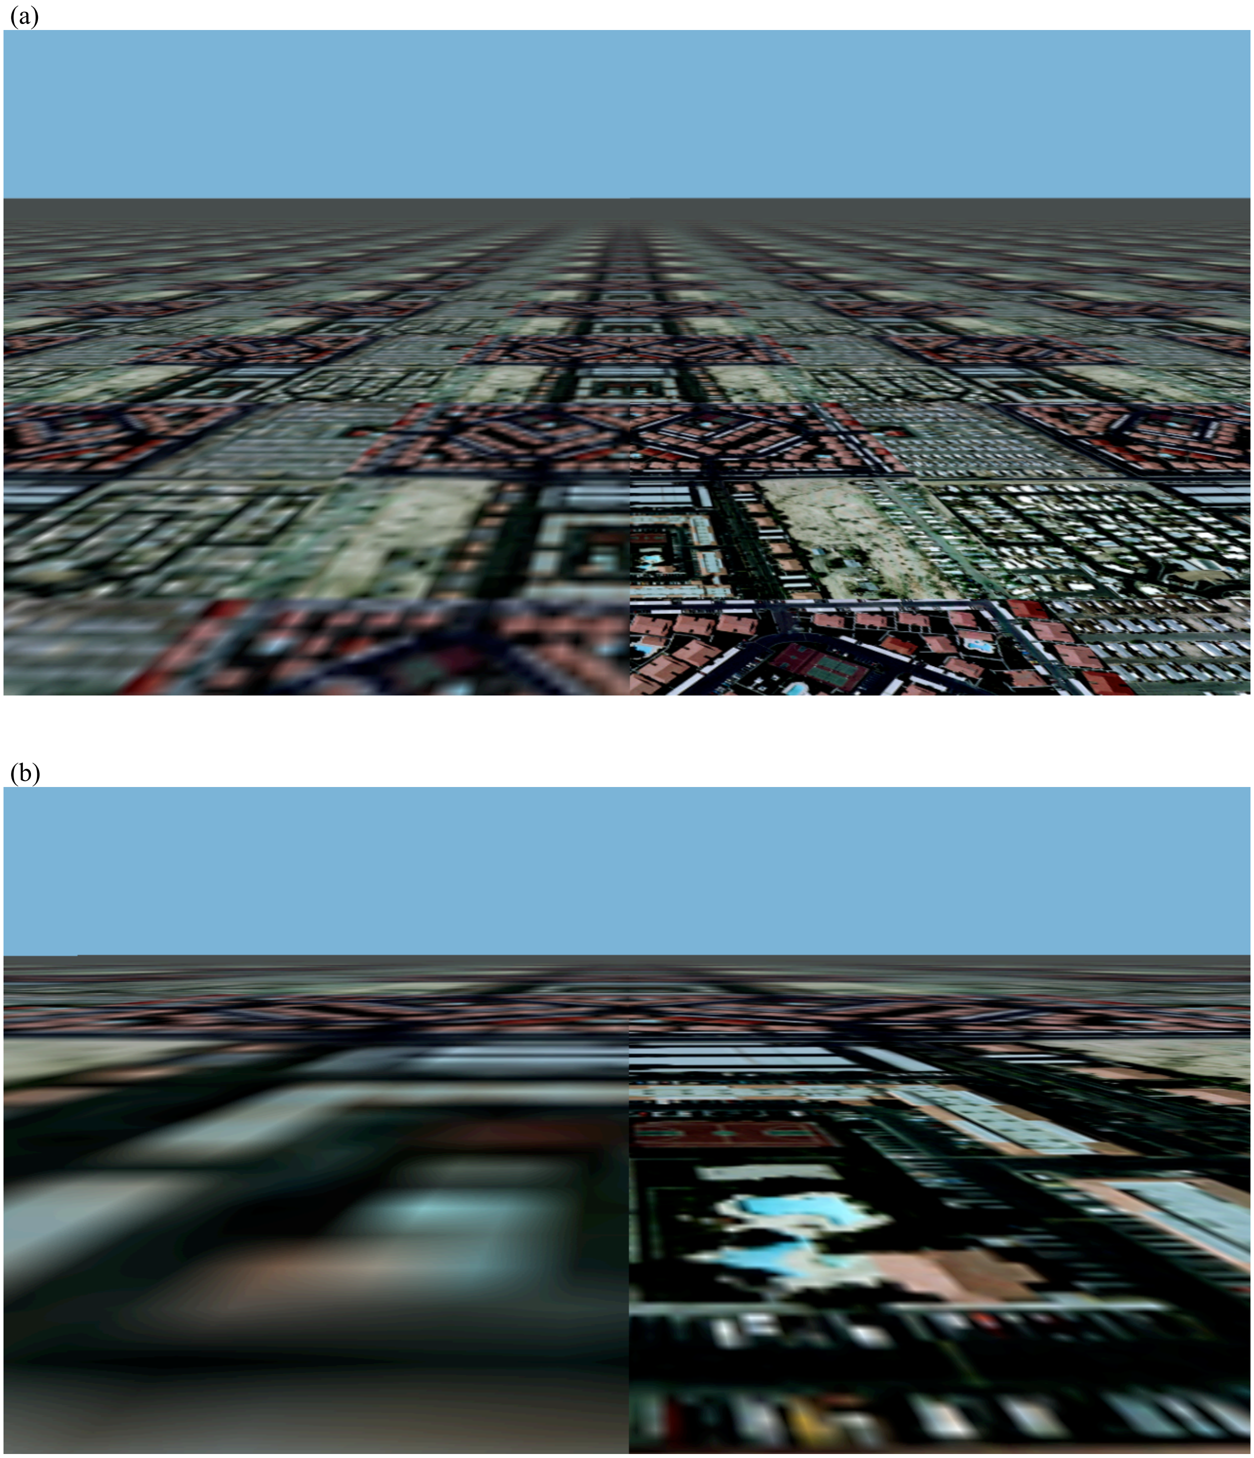

Figure 2. Effects of texture resolution, altitude, and distance on terrain representation. The synthetic environment was created by repeatedly mapping 2 textures of different resolution: 9.6 m on the left and 0.6 m on the right. (The two textures were derived from the same image. For illustrative purposes, the high-resolution texture was "flipped" horizontally before it was mapped.) Figures 2a and 2b depict 40°vertical fields of view, with 30° below the viewpoint, from simulated altitudes of 300 m and 30 m, respectively.

Note that, for both altitudes, the nearest visible terrain is sharper and contains more features when the texture resolution is higher. Similarly, for both altitudes, as the simulated distance along the ground increases, the effects of texture resolution decrease: On the distant terrain, the high- and low-resolution textures appear identical. Note also that the proportion of the image showing a texture-resolution effect is greater for an altitude of 30 m than for an altitude of 300 m.[2]

An antialiasing filter's cutoff frequency is determined on a pixel-by-pixel basis. The cutoff frequency (in c/m) for a given pixel depends upon the number of meters, in the projection on the view plane, subtended by that pixel. The larger the number of meters, the lower the cutoff frequency and, thus, the lower the texture resolution that is represented. For a given view of a synthetic environment, then, the higher the pixel density, the higher the texture resolution (and thus the greater the environmental detail) that can be represented at a given distance and the greater the distance at which a given texture resolution can be represented. Thus, if the images shown in Figure 2 had been generated by an ultra-high resolution system, the effects of texture resolution would extend over a larger proportion of the terrain.

Display System

The display system of a flight simulator creates a visible space-time image from a sequence of computer-generated, digital images. Ideally, the display of the spatial image for a given sampled point in time would be an accurate "reconstruction" of the nominal, antialiased image projected onto the view plane internal to the IG. However, for this to be the case, the display system would have to (a) pass, without attenuation, all spatial frequencies up to half the IG's sampling frequency and (b) remove all higher frequencies introduced by the sampling process (Holst, 1998). The nature and extent of deviations from this ideal depend upon the display technology. However, most display technologies do severely attenuate the high frequencies introduced by spatial sampling. The higher frequencies in the original antialiased image are also likely to suffer substantial attenuation. Various measures of display resolution (e.g., the highest frequency, in c/pixel, that can be produced with a specified amplitude) provide information relevant to the amount of attenuation of frequencies up to half the IG's sampling frequency. In general, if the resolution of a display system is low relative to that of the IG, the finer details represented in the digital image may not be visible in the display image. It is important to note, however, that attenuation increases with spatial frequency in *c/pixel*. As we have illustrated, the relationship between image units and environmental units depends upon the simulated viewing distance and angle. Relatively small environmental details in *near* terrain or objects may be well represented, even on a display that severely attenuates the highest frequencies in a digital image.

TEMPORAL CHARACTERISTICS OF COMPUTER-GENERATED IMAGERY

Continuous Space-Time Image

A continuous, time-varying image, internal to an IG, is sampled in time as well as in space. As with any sampled dimension, in order to prevent temporal aliasing, the temporal sampling rate must be greater than twice the highest temporal frequency in the image.

The Fourier transform of a three-dimensional, space-time image consists of three-dimensional, spatiotemporal-frequency components. The temporal frequency of a given component equals the product of its spatial frequency and its speed. Thus, if an IG uses spatial antialiasing filters to limit the maximum spatial frequency in the continuous space-time image, it will indirectly limit the temporal frequencies in that image. Spatial antialiasing filters do *not*, however, ensure that the temporal frequencies are all less than half the temporal sampling rate.

Temporal Aliasing

With current simulator technology, no direct attempt is made to prevent temporal aliasing—that is, the image is not subjected to a temporal, low-pass, presampling filter. Temporal aliasing is therefore possible. If a spatiotemporal-frequency component is temporally aliased, it will have the original spatial frequency but an erroneous, lower temporal frequency and thus an erroneous velocity.

Given that the temporal frequencies in a continuous image are a function of its spatial frequencies, the *spatially-antialiased*, continuous images of ultra-high resolution systems will typically contain higher temporal frequencies than corresponding images of

[2] Although neither the differences among nor the imperfections of various antialiasing techniques are considered here, the images presented in Figure 2 were affected by the resolution and filters of the IG used to generate them. They were also affected by the resolutions and filters used to resample them for the present document.

current systems. Estimates of the maximum temporal frequencies likely for typical flight scenarios can be used to assess both the extent of temporal aliasing likely with a standard 60-Hz update rate as well as the update rate necessary to eliminate temporal aliasing.

Measures of the temporal-frequency composition of a continuous space-time image could be based on spatial frequencies and velocities (in c/degree and degrees/s or in c/pixel and pixels/s) within spatially-local parts of the image. However, the computations would be complex. It is far simpler to use environmental units (i.e., spatial frequency in c/m and flight velocity in m/s). Consider, for example, the space-time image created during constant-altitude, constant-velocity flight over a flat surface textured with a single spatial frequency (e.g., 0.1 cycle/m). Because of the perspective projection, the image of the ground plane will contain a broad band of spatial frequencies. However, every spatial frequency (> 0.0 c/pixel) will have the same temporal frequency, as will every point in the image where the luminance varies over time—that is, every point in the image where the projected frequency is passed by the spatial antialiasing filter. This temporal frequency will equal the dot product of the two-dimensional spatial-frequency vector and the two-dimensional velocity vector, in environmental units.

For a given update rate, the simulated ground speed determines the maximum spatial frequency (in the direction of motion) that will *not* be temporally aliased—that is, the maximum spatial frequency for which the temporal frequency is less than half the update rate. Figure 3 shows this function for an update rate of 60 Hz. Note that only frequencies higher than about 0.5 c/m will be aliased at a speed of 120 knots, while frequencies as low as about 0.1 c/m will be aliased at speeds greater than 600 knots. If a specific spatial frequency is temporally aliased at a specific speed, then all higher spatial frequencies will also be aliased at that speed.

When the spatial-frequency bandwidth of a texture is limited by the resolution of the texture and not by the ground pattern it represents, an increase in texture resolution may result in temporal aliasing at lower speeds, and, for a given speed, an increase in the number of spatial-frequency components that are temporally aliased. However, an increase in the bandwidth of the texture will not necessarily have these effects: Building a database with a given texture resolution does not guarantee that that resolution will be present in the spatially-antialiased image: IG resolution, field of view, altitude, and distance-along-the-ground all affect the cutoff frequency of the

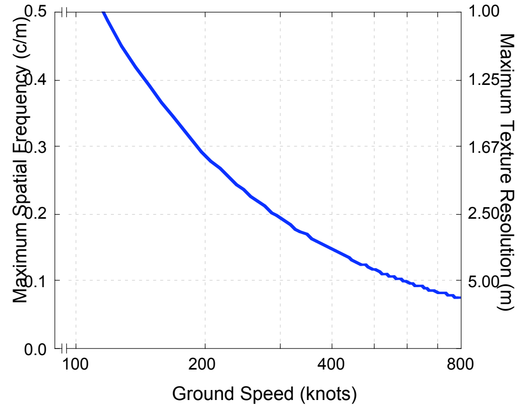

Figure 3. Maximum spatial frequency (in the direction of motion) that will *not* be temporally aliased, as a function of simulated speed (image update rate = 60 Hz). The corresponding maximum texture resolution is given on the right axis.

antialiasing filter (in environmental units) and thus the resolution of the texture that is sampled by the IG. For a given field of view and texture resolution, then, the proportion of a flat ground plane in which temporal aliasing could occur is a function of IG resolution and simulated altitude—in addition to ground speed (see Figure 4). For a particular altitude, as IG resolution increases, the proportion of the ground plane in which temporal aliasing could occur will increase. Correspondingly, an increase in IG resolution will

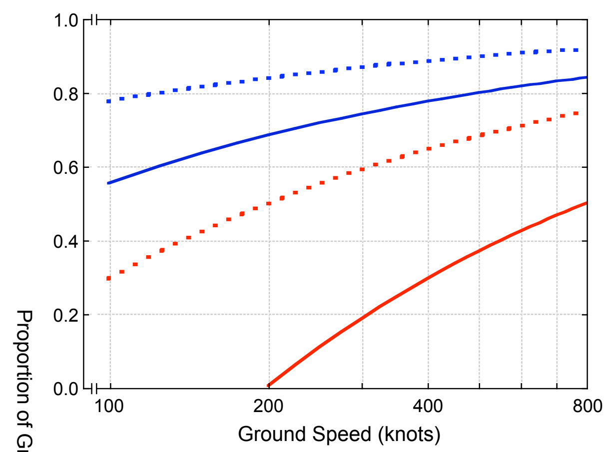

Figure 4. Proportion of ground-plane image (subtending the lower half of a 60° vertical field of view) in which some temporal aliasing could occur, as a function of simulated speed for altitudes of 500 m (red) and 50 m (blue) and IG vertical resolutions of 1024 (solid) and 4096 (dotted).

result in an increase in the altitude for which some aliasing is likely.

Temporal Characteristics of Display Systems

As we noted earlier in this paper, the display system of a flight simulator "reconstructs" a visible space-time image from a sequence of computer-generated, digital images. The fidelity of this reconstruction is affected by the temporal as well as the spatial characteristics of the IG and the display. Ideally, the IG's image-update rate would be high enough to prevent temporal aliasing, and ideally the display system would pass, without attenuation, all spatiotemporal frequency components in the original image and would remove all sampling-induced components with higher temporal frequencies.

Current flight simulators fail to meet this ideal in a variety of ways. As we have indicated, IG technology does not prevent temporal aliasing. Once temporal aliasing occurs, there is no way for a display system to accurately reconstruct the original, nominal image. Moreover, the band of temporal frequencies passed by a display system depends upon the display technology. Some display systems pass, with minimal attenuation, temporal frequencies many times greater than half the update rate (e.g., CRTs), whereas others attenuate many temporal frequencies less than half the update rate (e.g., LCDs). If a display system passes temporal frequencies higher than half the update rate, it will pass all the aliased spatiotemporal frequency components. It will probably also pass all of the spatiotemporal frequency components in the original image as well as many high temporal frequency components introduced by the temporal sampling process. If a display system attenuates temporal frequencies less than half the update rate, it will attenuate the high temporal frequency components introduced by the temporal sampling process, but it is also likely to attenuate some of the spatiotemporal frequencies in the original image. The extent to which it attenuates aliased components will depend upon the temporal frequencies of those components.

PERCEPTION OF ALIASED AND OTHER SPATIOTEMPORAL FREQUENCIES INTRODUCED BY TEMPORAL SAMPLING

During observation of a display image, an image is formed on the retina of each eye. If the observer's direction of gaze is constant during the observation period, the temporal frequencies in the retinal images will match those in the display image. This correspondence does not hold, however, during movement of the observer's eyes or head. In fact, during certain kinds of eye movements (e.g., smooth pursuit of an object in the display image), a temporally aliased component in the display image may not be aliased in the retinal image—and vice versa. Eye movements may thus have a large effect on the perception of aliased as well as other spatiotemporal frequencies introduced by temporal sampling. In particular, when a display image contains both the original spatiotemporal frequency components and temporal-sampling induced components, pursuit eye movements may determine which components are seen.

The human visual system is sensitive to only a limited band of spatiotemporal frequencies. The maximum spatial frequency falls between 30 and 60 c/degree (Wilson et al., 1990); the maximum temporal frequency falls between 50 and 70 Hz (Watson, 1986). Sensitivity is very low to relatively high spatial frequencies with relatively high temporal frequencies (and to very low spatial frequencies with very low temporal frequencies). Presumably, any spatiotemporal-frequency component falling within the passband of the human visual system may have a perceptual effect. Thus, if a 60 Hz display system passes temporal-sampling induced components with temporal frequencies greater than 30 Hz, those with low *spatial* frequencies (in c/deg) are likely to be visible. Aliased frequencies are also likely to be visible: Although their spatial frequencies (in c/deg) may be relatively high over much of the image, their temporal frequencies may be quite low.

If a single moving spatial sinusoid is temporally aliased, it usually appears to move in accord with its aliased temporal frequency rather than with the temporal frequency in the original image—that is, it appears to move at a lower speed and most often[3] in the wrong direction. If *all* of the spatiotemporal frequencies in a more complex pattern are aliased, the pattern usually appears to move in accord with the erroneous velocity of the *lowest* spatial frequency—as in the well known "wagon wheel" effect. However, moving textures in flight simulator imagery will usually be characterized by many of the spatiotemporal

[3] If perception is consistent with the aliased component, the perceived direction of motion will depend upon the relationship between the temporal sampling rate and the temporal frequency in the original image. For example, if the sampling rate is greater than the temporal frequency but less than two times the temporal frequency, as is often the case, the spatial sinusoid will appear to move in the wrong direction, whereas if the sampling rate is slightly higher than two times the temporal frequency, it will appear to move in the right direction (but at an erroneous speed).

frequency components that were in the original continuous image as well as many introduced by temporal sampling. Relatively little is known about the perceptual effects of sampling-induced components within complex, broadband images that also contain many veridical components.

We have previously used IGs with standard resolutions (e.g., 1280 x 1024) to examine perceptual effects of temporal aliasing during low-level flight over flat terrain textured with complex generic or geospecific patterns. In our view, the motion was of poor quality in image regions characterized by large numbers of both nonaliased and aliased component—appearing incoherent or "choppy," rather than smooth and continuous.

EXPERIMENTAL ASSESSMENT OF EFFECTS OF RESOLUTION ON MOTION QUALITY

We have initiated a research program to assess the effects of IG, texture, and display resolution on perceived-motion quality during simulated low-level flight. In support of this effort, we are developing a software package with the capability of generating coordinated images on multiple IGs and of setting IG-pixel-sampling frequency (effective resolution) and graphics, database, and simulated-flight parameters on a trial-by-trial, IG-specific basis.

In an initial experiment, we examined factors affecting the *spatial extent* of perceived incoherent or choppy motion during simulated flight over flat, textured terrain.

Apparatus

Images were generated by three PCs equipped with Nvidia GeForce4 Ti4600 graphics cards. A fourth PC was programmed to synchronize the IGs and to control the experiment.

The images were displayed on three CRT monitors, each of which was set to a pixel resolution of 1280 x 1024 and a color resolution of 32 bits. The brightness and contrast settings of the monitors and the gamma settings of the drivers were adjusted to match CRT output levels and to linearize the relationships between digital image value and display luminance.

The three monitors were aligned horizontally to create a sectioned FOV of about 65º (H) x 17º (V). Although the monitor positions were fixed, the observer could control the vertical position of the viewing frustum (and thus the visible region of an image with a greater vertical FOV).

IG resolution was varied by sampling the continuous image, internal to an IG, at either 1280 x 1024 or 320 x 256 locations. The IGs' viewing frusta, which did not vary with sampling frequency, were such that samples were separated by a visual angle of about 1 arcmin in the higher resolution images and of about 4 arcmin in the lower resolution images. To create an appropriately-sized display image, the lower resolution digital images were "stretched," using bilinear filtering, to fill the full 1280 x 1024 display-buffer memory.

"Database" Characteristics

The characteristics of most geospecific textures vary appreciably across space. Such variation introduces unwanted noise into experimental results. Generic textures, on the other hand, are often unnaturally uniform. To create a texture with constrained, randomly distributed variation in feature size and shape, we followed the following procedure: First, we assigned a random value between 0 and 255 to each of the elements in a 512 x 512 array. Second, we introduced rectangular and diagonal regions (formed by two 4 x 4-texel squares) of uniform value. The locations and values of these larger features, which were allowed to overlap, were determined by a random selection of 0.8% of the original texels. Figure 5 illustrates a 64 x 64-texel excerpt of the resulting texture pattern. (To eliminate artifacts that would be introduced by any CRT color-convergence problems, in this experiment only the green component of the texture pattern was used.)

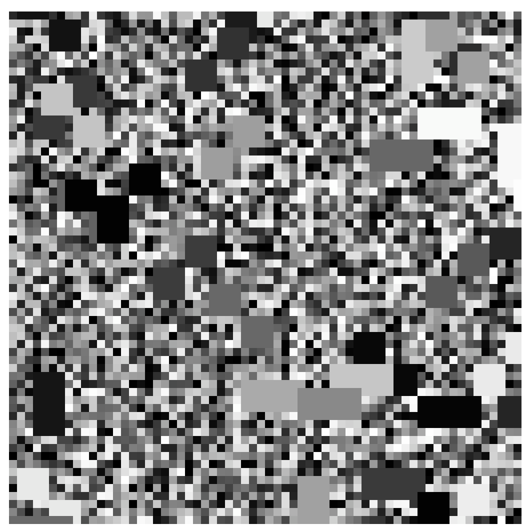

Figure 5. 64 x 64-Texel Excerpt of Texture Pattern

A flat ground plane was textured by tiling the 512 x 512 texture, with 1 texel/m^2. To limit spatial aliasing, we used standard texture mipmapping, with bilinear-magnification and trilinear-minification filters. (A mipmap is a set of prefiltered versions of a texture pattern, with each successive version having half the resolution of the preceding version. The size of a pixel relative to the projected size of a texel determines the mipmap level[s] applied at that pixel. With a bilinear-magnification filter, the value of a pixel is derived from 4 texel values within the highest resolution version of the texture. With a trilinear minification filter, the value of a pixel is derived from 4 texel values within each of 2 mipmap levels.)

When mapped with 1 texel/m, the base level of the texture (Mipmap Level 0) was characterized by spatial frequencies up to 0.5 c/m (2 m/c); Mipmap Level 1 by frequencies up to 0.25 c/m; and Mipmap Level 2 by frequencies up to 0.125 c/m. Thus, there would be some temporal aliasing of the base level texture for any speed greater than 1 m/update, some for Mipmap Level 1 for any speed greater than 2 m/update; and some for Mipmap Level 2 for any speed greater than 4 m/update.

Procedure

During each experimental session, observers viewed a random order of 16 trials representing all possible combinations of four bi-level variables: The IG resolution was either 4 arcmin or 1 arcmin per pixel; the simulated altitude was either 100 or 200 ft (30.5 or 61 m); the simulated straight and level speed was either 2 or 4 m/update (233 or 466 knots); and the initial vertical position of the viewing frustum was either “high” or “low” (second or fourth quarter of a nominal frustum with 27º above and 35º below the eye point). Observers used a joystick to “slide” the viewing frustum up and down until a horizontal line in the center of the middle CRT separated motion that appeared smooth and continuous from motion that appeared choppy or incoherent.

Results

We analyzed both the location of the line, relative to the horizon, and the meter:pixel (i.e., texel:pixel) ratio at that location. As predicted, line height (perceived spatial extent of choppy motion) was significantly greater for (a) an IG resolution of 1 arcmin than for an IG resolution of 4 arcmin, (b) a speed of 4m/update than for a speed of 2 m/update, and (c) an altitude of 100 ft than for an altitude of 200 ft. Moreover, as predicted, the meter:pixel ratio was significantly greater for the higher of the two velocities. However, the meter:pixel ratio was also significantly greater for the lower-resolution IG. With 4 arcmin/pixel, the IG’s texture antialiasing algorithms limited the proportion of the image in which the motion appeared choppy: When the speed was 2 m/update, the line was positioned halfway up the blend region of Mipmap Levels 0 and 1; when the speed was 4 m/update, the line was positioned halfway up the blend region of Mipmap Levels 1 and 2. In contrast, with an IG resolution of 1 arcmin/pixel, the line was positioned well within Mipmap Level 0 for the lower speed and less than halfway up the blend region of Mipmap Levels 0 and 1 for the higher speed. Thus, for the higher resolution IG, the perceived extent of choppy motion did not correspond to the extent of temporal aliasing in the image. These results suggest that the spatial and temporal properties of the human visual system play a primary role in limiting the perception of temporal aliasing when the resolution of an IG is as high as 1 pixel/arcmin.

Conclusions

The spatial resolutions of the databases, IGs, and displays of next-generation flight simulators are expected to be significantly higher than those of current systems. As a consequence, the nominal, space-time images sampled by an IG may have higher temporal frequencies than those that characterize current systems. If the image-update rate is not greater than twice the highest temporal frequency, temporal aliasing will occur.

Most IGs have an image-update rate of about 60 Hz. While this update rate is adequate to prevent temporal aliasing for current texture and IG resolutions, at most flight speeds and altitudes, temporal aliasing is sometimes present at very low altitudes and high speeds. If a 60-Hz update rate is used with an ultra-high resolution system, temporal aliasing will occur at higher simulated altitudes and lower simulated speeds. During low-level flight, a large proportion of the image may be characterized by aliasing, and a large proportion of the components within parts of the image may be aliased. If aliasing is perceptually salient, it could seriously degrade perceived image quality and affect a variety of perceptual tasks.

We have initiated a research program to assess the effects of IG, display, and texture resolution on motion perception and perception-dependent tasks during low-level-flight simulation. The results of an initial experiment suggest that spatial and temporal properties of the human visual system may limit the perceived spatial extent of temporal aliasing in images generated by IGs with spatial resolutions as high as 1 arcmin/pixel. In future research, we plan to use a variety of procedures, tasks, and high-resolution-display technologies in an attempt to determine the database, IG, and display characteristics necessary for high quality, high resolution imagery.

REFERENCES

Holst, G. C. (1998). *Sampling, aliasing, and data fidelity for electronic imaging systems, communications, and data acquisition.* Winter Park, FL: JCD Publishing & Bellingham, WA: SPIE Optical Engineering Press.

Watson, A. B. (1986). Temporal sensitivity. In K. R. Boff, L. Kaufman, & J. P. Thomas, (Eds.), *Handbook of Perception and Human Performance, Vol. 1., Sensory Processes and Perception* (Chapter 6). New York: Wiley.

Wilson, H. R., Levi, D., Maffei, L., Rovamo, J. & DeValois, R. (1990). The perception of form: Retina to striate cortex. In L. Spillman & J. S. Werner (Eds.) *Visual Perception: The Neurophysiological Foundations* (pp. 231 – 272). New York: Academic.