Fourier Synthesis of Ocean Scenes

Gary A. Mastin, Peter A. Watterberg,* and John F. Mareda

Sandia National Laboratories

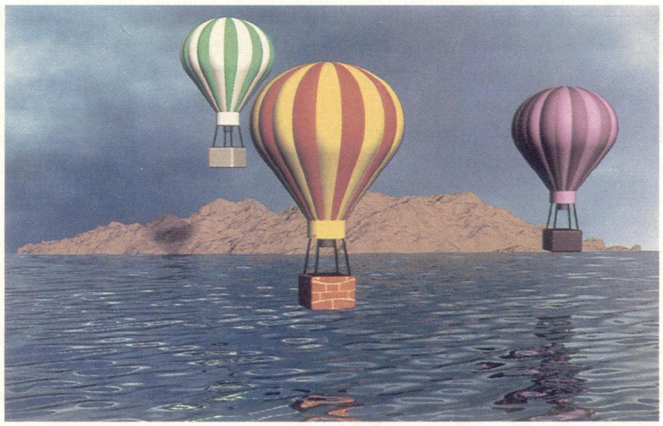

Considerable interest in scene synthesis has focused on fractal techniques such as 1/f noise filtering. However, these techniques produce poor synthetic ocean scenes. A superior technique is presented here that uses an empirical sea spectrum model to represent a fully developed sea in nature. A computer program has been written whereby a user supplies a wind velocity, and a two-dimensional Fourier-domain filter is created whose shape is described by the modified Pierson-Moskowitz spectrum. A white-noise image can then be filtered and rendered with a ray-tracing algorithm to produce a realistic ocean scene. Wave motion may be invoked by manipulating the phase of the Fourier-transformed white-noise image. Several phase-manipulating techniques have been implemented.

Scene synthesis has become very popular in commercial movie production, flight simulators, architectural planning, and the graphic presentation of mathematical entities. In the quest for realism, the art of scene synthesis has progressed from line drawings to shaded polygon tilings to fractal surfaces. Fractal terrains[1-3] represent a major step in our ability to synthesize natural-looking landscapes. Natural scenes, however, consist of more than landforms. Water plays an important role.

Past attempts at rendering water have been hampered by overly simplistic views of water. Max[4] and Schachter[5] had some success at modeling ocean waves as the sum of a limited number of sinusoids, possibly having added noise. Ocean waves, however, are not simple sinusoids or the superposition of simple sinusoids. Neither are they fractal surfaces having the same fractal dimension in all directions. They are complex waveforms born of the momentum transfer from wind. Ocean waves develop changing profiles due to fetch (the distance along open water over which the wind blows), air-sea temperature difference, variations in surface roughness, and alterations in the vertical wind profile. Furthermore, they transfer energy between

*Peter Watterberg is currently employed at the Savannah River Laboratory, which is run by Dupont for the Department of Energy.

0272-1716/87/0300-0016$01.00 © 1987 IEEE

spectral frequencies in a nonlinear fashion, making most linear models rough approximations, at best. Seas should not be looked upon as deterministic entities, rather as statistical processes. In this article, seas are viewed as filtered white noise.

Voss[2] has presented a method of generating fractal scenes by filtering white noise with a 1/f noise (Power law) filter. The rationale is that fractal geometries, with their "self-similarity" over a wide range of scales, can be characterized by Power law filtering which introduces correlations over a wide range of scales. Stated another way, fractal geometries look similar at different scales despite small differences in detail. The same is true for Power law filtered random-number signals.

Voss' landscape pictures in Mandelbrot's book,[1] and the mountain lake scene in Figure 1 were generated by filtering the magnitude of the fast Fourier transform (FFT) with a 1/f filter designed to give the desired fractal dimension, then inverse transforming and rendering by a ray-tracing algorithm. The frequency response of the filter is shown in Figure 2. Note the dominance of the

Figure 1. Mountain lake scene was created from a 1/f filtering of a random intensity image, and rendered by a ray-tracing program (written by Watterberg). The mountains, clouds, and water texture were made from the same white-noise image.

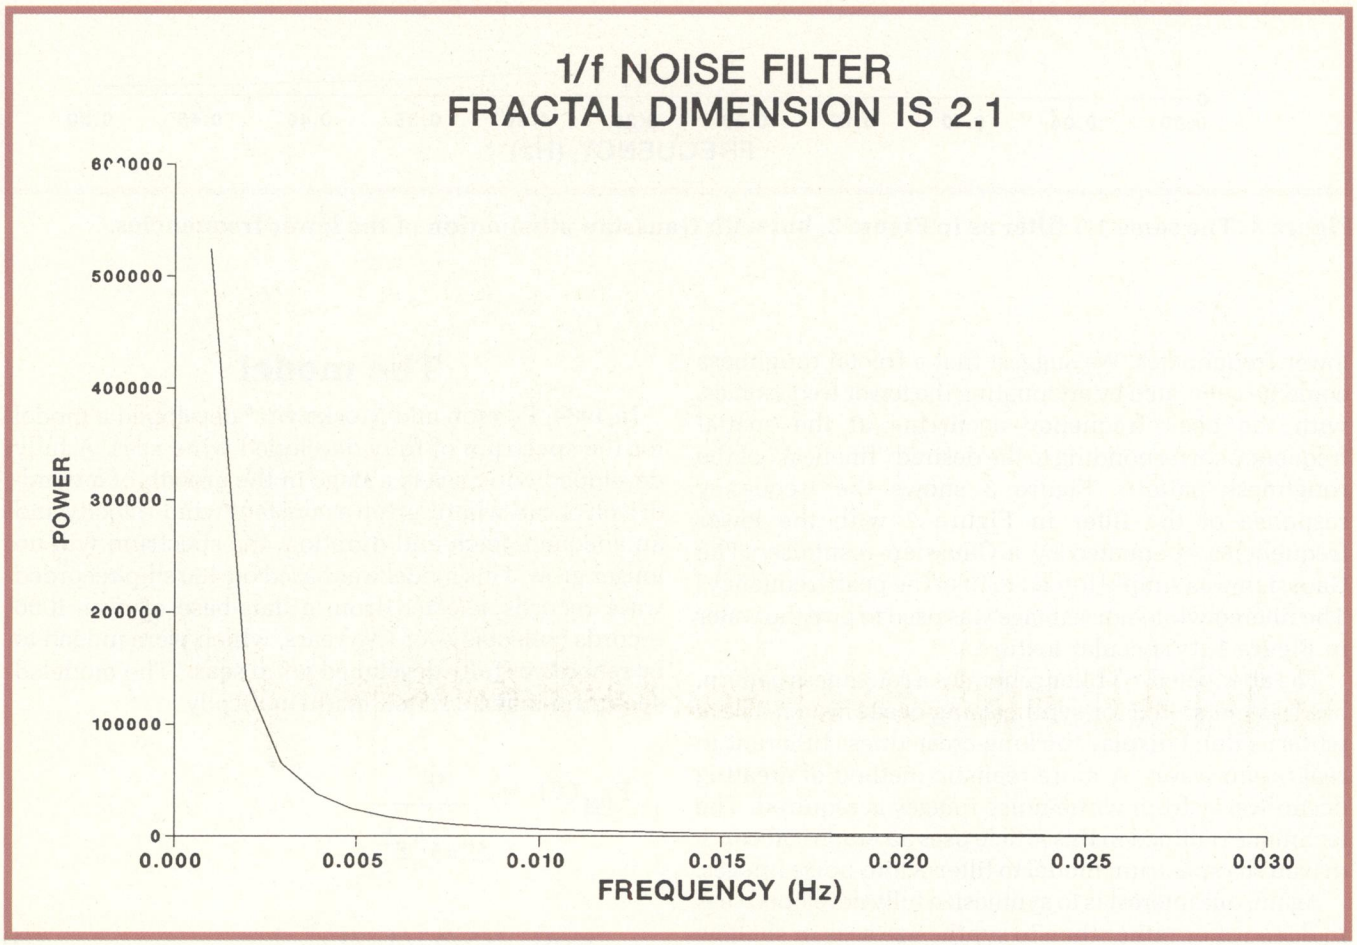

Figure 2. The frequency response of the 1/f filter used in generating the mountains in Figure 1. The lower frequencies correspond to the basic shape of the mountains while the higher frequencies give them roughness.

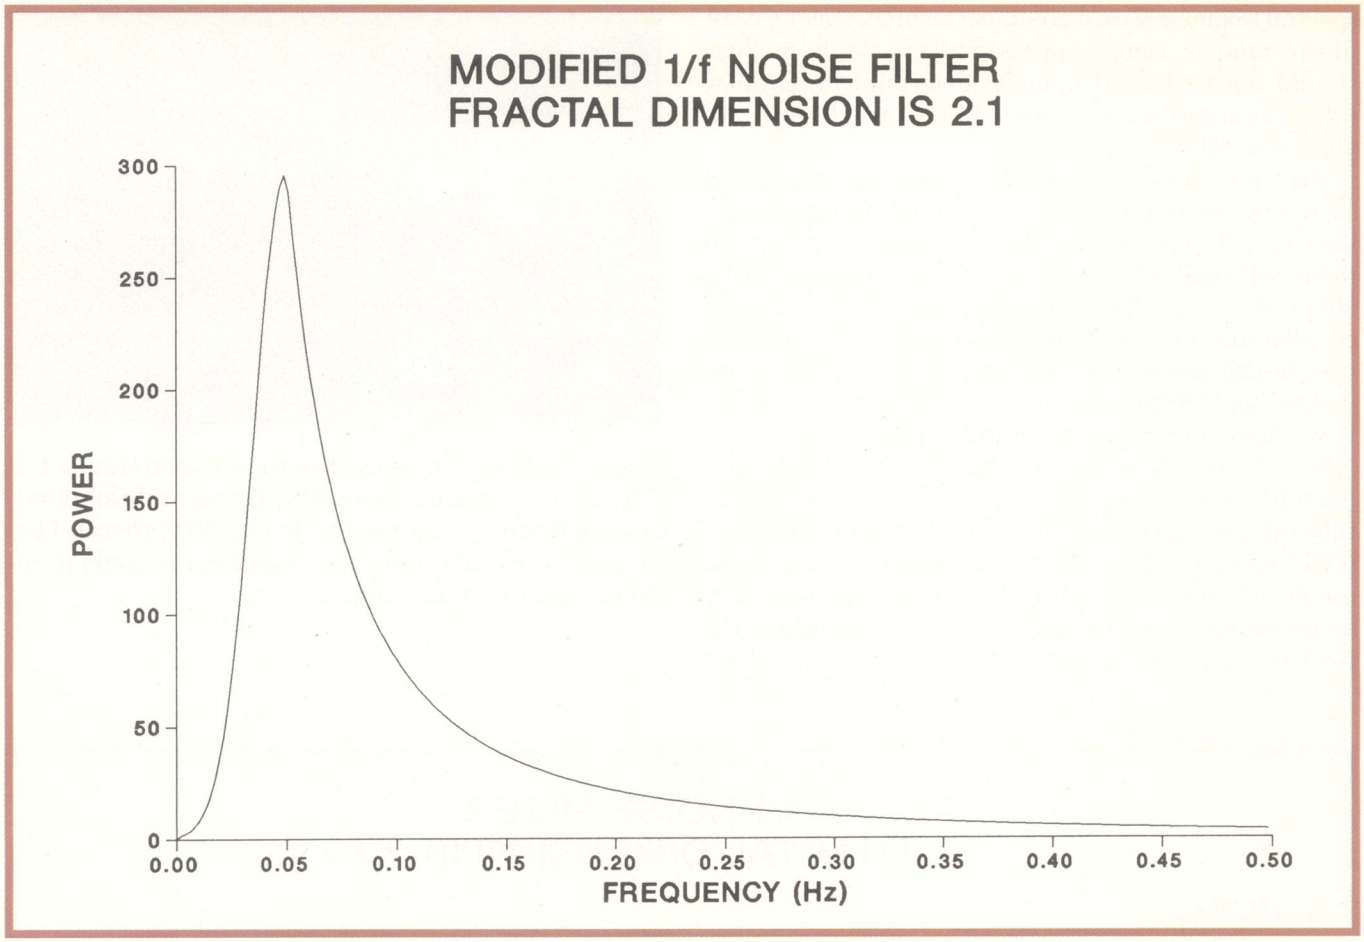

Figure 3. The same 1/f filter as in Figure 2, but with Gaussian attenuation of the lower frequencies.

lower frequencies. We suggest that a fractal roughness could be generated by attenuating the lower frequencies, with the peak frequency occurring at the spatial frequency corresponding to the desired "fineness" of the roughness pattern. Figure 3 shows the frequency response of the filter in Figure 2 with the lower frequencies attenuated by a Gaussian response. (The Gaussian was simply forced to fit at the peak frequency.) The filtered white-noise image was used to give the water in Figure 1 its specular texture.

The attenuated 1/f filter generates a roughness pattern, but it is not suited for synthesizing ocean water. These patterns don't display the long-crestedness inherent in real ocean waves. A more realistic method of creating ocean waves from white-noise images is required. The technique outlined in this article uses an empirical wind-driven sea spectrum model to filter white-noise images.

Again, our interest is to synthesize fully developed seas in deep water rather than to synthesize seas in shallow water where effects like wave refraction and wave breaking exist. Shallow water coastal scene synthesis was covered by Peachey,[6] and also by Fournier and Reeves.[7]

The model

In 1964, Pierson and Moskowitz[8] developed a model for the spectrum of fully developed wind seas. A fully developed wind sea is a stage in the growth of a wind-driven ocean where, given a constant wind velocity and an adequate fetch and duration, the spectrum will no longer grow. This model was based on 460 ship-recorded wave records selected from a data-base of over 1000 records collected over five years, which were judged to be records of fully developed wind seas.[9] The modeled spectrum may be stated mathematically,

$$F_{PM}(f) = \frac{\alpha g^2}{(2\pi)^4 f^5} \exp\left[-(5/4)(f_m/f)^4 \right] \quad (1)$$

where $F_{PM}(f)$ = downwind power spectrum, f = frequency in Hz, f_m = peak frequency in Hz, α = Phillips

constant = 0.0081, and g = gravitational constant. The peak frequency f_m is directly related to the wind speed at a height of 10 m above the sea surface, u_{10}, by the expression,

$$f_m = 0.13\ g\ /\ u_{10}\ . \qquad (2)$$

The significance of this model is that the one-dimensional (in the direction of the wind) spectrum of a fully developed wind sea may be calculated from Equation 1 using the wind speed, u_{10}.

The assumptions inherent in Equations 1 and 2 should be stated. First, we are assuming fetch-limited wave growth in deep water with a constant wind field. Fetch-limited wave growth is divided into three time stages: an initial growth stage, a transitional stage, and a late stage where growth is strongly reduced. The latter stage applies here and is referred to as a fully developed wind-sea state. The upper limit of a fully developed sea is wave breaking, a condition not covered here.

Second, aerodynamic drag does not remain constant, but changes with wind speed and sea state.[10,11] A constant drag coefficient of $C_{10} = 1.8\times10^{-3}$ assumed here is valid for a wind speed u_{10} = 15 m/s. It is approximately valid for wind speeds between and 10 and 20 m/s.

The Pierson-Moskowitz spectrum, with minor modifications,[12] is still considered to be a valid model for fully developed seas. Hasselmann et al.[13] suggested a two-dimensional spectrum,

$$F(f,\theta) = F_{PM}(f)\ D(f,\theta), \qquad (3)$$

based on the Pierson-Moskowitz model, where $D(f,\theta)$ is a directional spreading factor that weights the spectrum at angles θ from the downwind direction. The spreading factor is defined by the relations,

$$D(f,\theta) = N_p^{-1} \cos^{2p}(\theta/2) \qquad (4)$$

where

$$p = 9.77\ (f/f_m)^{\mu}$$

$$\mu = \begin{cases} 4.06 & , \quad f < f_m \\ -2.34 & , \quad f > f_m \end{cases} \qquad (5)$$

and the normalization constant

$$N_p = 2^{1-2p}\pi\ \Gamma(2p+1)/\Gamma^2(p+1) \qquad (6)$$

is defined such that $\int_{-\pi}^{+\pi} D\,d\theta = 1$.

The frequency associated with the peak power is f_m.

The spreading function suggested by Hasselmann et al.[13] differs from the more traditional $\cos^2\theta$ distribution in two ways. First, it creates a narrower profile near the peak frequency in the downwind direction of the spectrum. This modification is in keeping with additional experimental evidence. Secondly, it attenuates the peak frequency and forms a bimodal spectrum shape for angles nearing 90 degrees from downwind. The effect is to suppress the long-crested peak frequency components that would run parallel to the wind direction, while retaining nonpeak frequency components.

A one-dimensional profile of the modified Pierson-Moskowitz spectrum computed for a wind speed of 15 m/s is shown in Figure 4. This profile was extracted in the downwind direction which, in this case, was 45 degrees. Note that the frequency axis has been normalized by f_m and is therefore nondimensional. This is done to provide a reference spectrum often used to compare a variety of empirical data. The frequency scale independence should not be misunderstood as a betrayal of the physical parameters, which are assumed in the model and are scale sensitive.

Probably the most significant reason for performing frequency normalization when comparing empirical data is fetch. The fetches generating empirical data, whether from the field or in laboratory experiments, are "long," but "long" is a relative term. Komen et al.[12] note that the high frequency part of a fully developed sea spectrum is nearly fetch independent. This is not true for the lower frequencies. The peak frequency varies inversely with fetch, however, within the wind speed limits of 10 to 20 m/s, the shape of the fully developed wind-sea spectrum does not vary appreciably. Normalizing by the peak frequency becomes a standardization technique for generating a reference spectrum.

Implementation

Synthetic ocean images may be generated from white-noise images in a very straightforward manner using the forward FFT, a filter based on Equation 3, and the inverse FFT. These steps were implemented on a VAX 11/750 minicomputer. The FFT used here is a mass store two-dimensional FFT provided in the IEEE signal-processing package.[14]

The white-noise image is generated by adding uniformly distributed noise, having intensities between −127 and 127, to a constant intensity image of gray level 128. This results in random gray shades between 0 and 255, inclusive. Figure 5 is an example. A two-dimensional forward FFT is performed on the image to generate a magnitude and a phase image. The result of the FFT is a complex number file. The magnitude and phase images are created by the operations,

$$MAG = SQRT(\ REAL^2 + IMAGINARY^2\ ) \qquad (7)$$

and

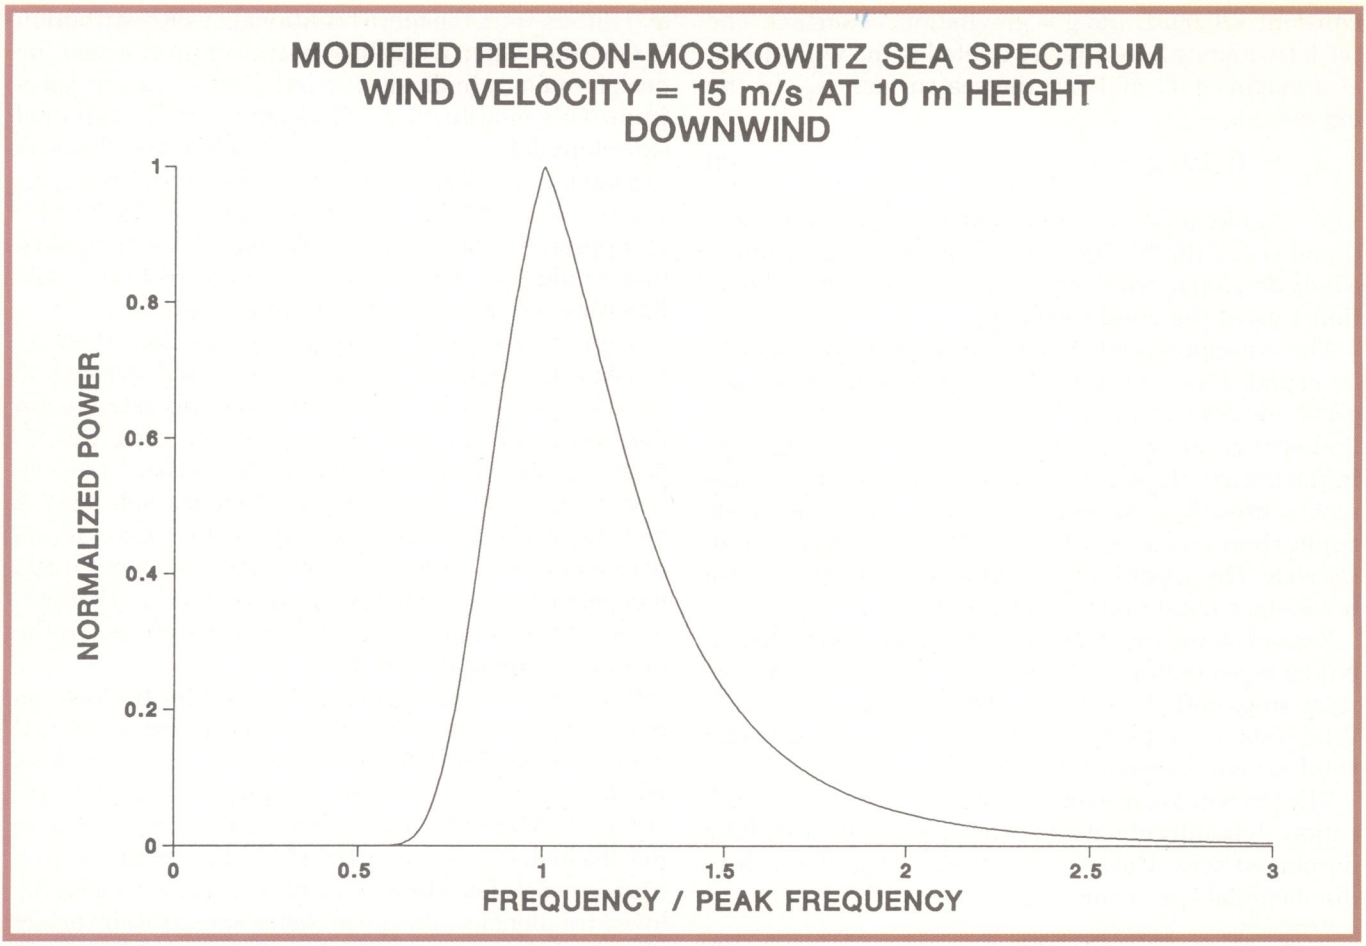

Figure 4. The modified Pierson-Moskowitz spectrum computed for a wind speed of 15 m/s.

Figure 5. A white-noise image with intensities uniformly distributed between 0 and 255, inclusive.

```
PHASE = Arctan ( IMAGINARY / REAL ),
```
(8)

respectively. Only the magnitude image will be filtered. An interesting property of the FFT of a white-noise image is that the magnitude is a white-noise image. All frequencies are present, as shown in Figure 6.

The next step is to filter the FFT magnitude. This is done in the Fourier domain by multiplying the magnitude by $D(f,\theta)$. The filter is shown in Figure 7 for a wind speed of 15 m/s at 45 degrees. The long-crestedness of the resulting image is emphasized by the passing of frequencies around f_m at 45 and 225 degrees (white) and the attenuation of the same frequencies at 135 and 315 degrees (black). The filter is normalized in power to provide a gain of one. It is normalized in frequency and therefore independent of scale.

The filtered FFT magnitude and the original phase are then inverse transformed to create a synthesized wave image. The result of processing the white-noise image in Figure 5 is shown in Figure 8. People who have flown over seas should appreciate the realism of this image.

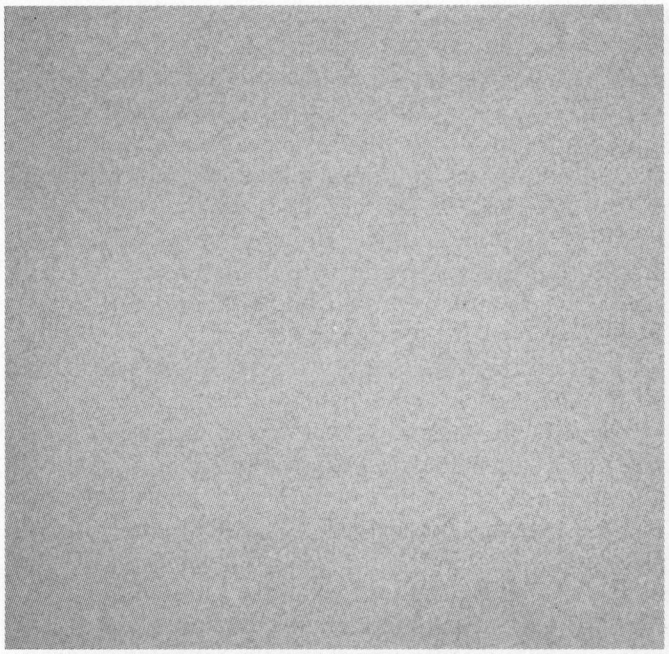

Figure 6. The FFT magnitude of the image in Figure 5.

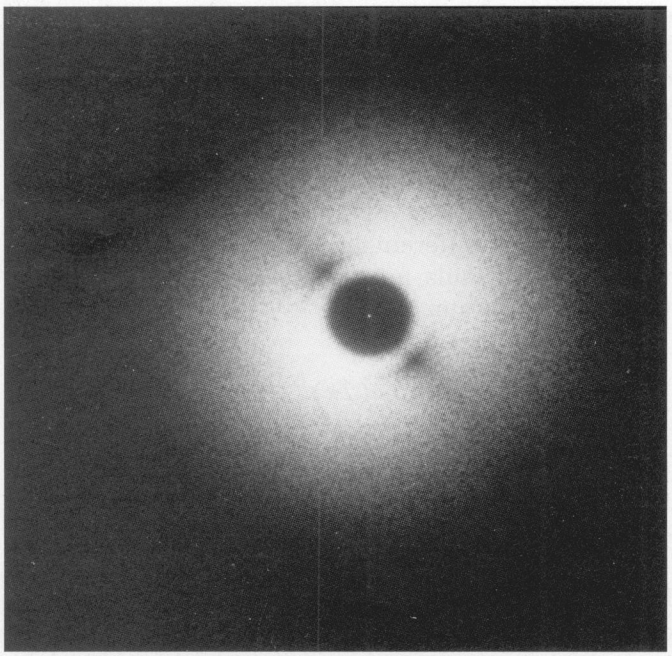

Figure 7. The Fourier-domain filter based on the modified Pierson-Moskowitz spectrum model for u_{10} = 15 m/s at 45 degrees.

Ocean scene animation

A sequence of images portraying wave motion can be animated by manipulating the FFT phase. The simplest manipulation is based on the translation property of the Fourier transform which is[15]

$$\mathcal{F}\{f(x-x_0)\} = \exp(-j\omega x_0)F(\omega) \quad . \qquad (9)$$

This implies that a one-dimensional signal may be translated in the spatial domain by multiplying the Fourier transform of the signal by a linear phase component. For each movie frame, x_0 would be incremented, thereby changing the slope of the linear phase and translating the signal at a constant rate. This is a very unrealistic manipulation since each frame is a translated carbon copy of the preceding one. Realistic wind-driven seas gradually change their structure with time without altering their fundamental periodicity. Two more sophisticated phase manipulation techniques are presented here: One is ad hoc, and the other is based on the deep water model for gravity wave propagation.

The ad hoc technique will be presented first because of its recent use in the SIGGRAPH '86 movie "Hot Air," (produced by J. Mareda, available from ACM). Since changing the derivative (slope) of the linear phase by a constant amount per frame produces translated carbon copy frames, why not shift frequency components around the peak frequency by a constant different from the surrounding frequencies? We defined dominant frequency components as those frequencies bounded by

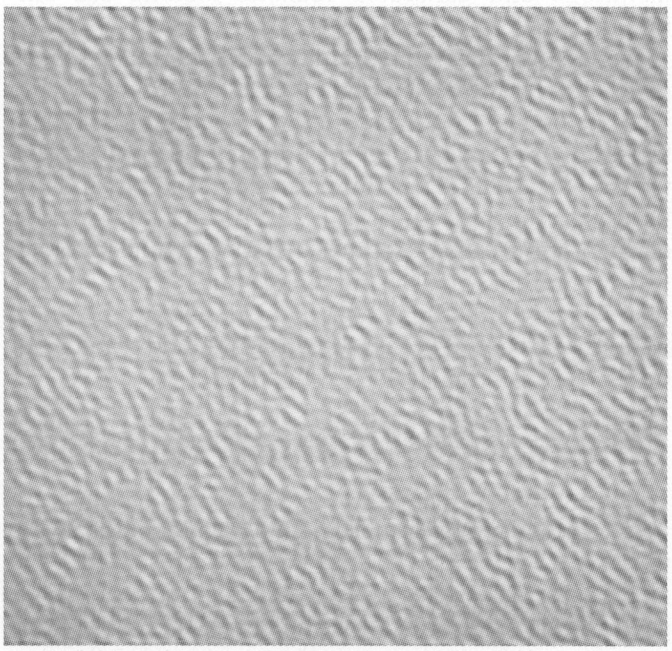

Figure 8. The synthesized ocean image created by processing the white-noise image in Figure 5.

two-thirds of the energy of the peak frequency f_m. The dominant components were shifted independently of the surrounding components. The choice of how much to shift the dominant components relative to the surrounding frequencies was then decided by trial and

error. In "Hot Air," the surrounding frequencies were moved 0.003 radians per frame, while the dominant frequency components were shifted 0.05 radians per frame. Keep in mind that all shifts must be less than π radians to avoid aliasing.

The second phase manipulation technique takes advantage of our idea of shifting different frequency components by different amounts based on a model. A well-known formula for the phase velocity of gravity waves in deep water is[16]

$$c^2 = g/k = (g\lambda)/(2\pi) \quad (10)$$

where c = phase velocity, g = acceleration due to gravity, λ = wavelength, and k = wave number (spatial frequency).

The phase velocity c is dependent upon the wavelength (or spatial frequency), which becomes the mechanism for shifting the FFT phase. By assigning a scale to the image — i.e. how many m/pixel (looking straight down on the scene, disregarding perspective) — we can convert c from m/s to radians/s. For video sequences, we have 30 frames/s, so we can apply the appropriate phase shift c, in radians/frame, to the FFT phase and thereby properly model wave motion. We applied this technique and generated realistic looking sea animations.

Processing times were monitored for the single scenes generated by the ad hoc technique. The times for the wave dispersion model animation are comparable. The times for the forward FFT, filtering, and inverse FFT were on the order of five minutes, a few seconds, and five minutes, respectively, for an ELXSI 6400 (10-processor configuration). The FFT times include the time for creating magnitude, phase and image files and with the forward FFT, creating a log-encoded magnitude image (which was useful for viewing intermediate results, but unnecessary in an actual application).

The ray-tracing image rendition

The synthetic ocean waves are presented here in images produced with experimental ray-tracing software that we developed. The data are represented as a 512×512 array of height values. The technique for intersecting a ray with ocean wave height field is the same as the one used for intersecting the terrain, since they are essentially the same data structure.

To intersect a ray with a height field, the ray is intersected with a series of bounding boxes (or extents) that are successively finer approximations to the surface. While similar to the method presented by Kajiya,[17] it differs in that these extents are rectangular in cross section while Kajiya's are triangular. Since all the sides of the boxes are parallel to the x-, y-, and z-axes, the intersections are trivial and fast. Each box is divided into four smaller boxes for successive refinement. This is essentially a quadtree structure where each node contains the maximum height value for all its subtrees.

To test a ray against the quadtree, start by pushing the root node (a box enclosing the entire height field) on a consideration stack. Iterate as follows: Pop the top node on the consideration stack and test the ray against it. If the ray is not entirely above the box, from one to three of its offspring (determining which ones is a two-dimensional problem) are pushed on the stack (farthest first, since we want to consider it last).

If the stack becomes empty, the ray doesn't intersect the surface. If a leaf is reached, the ray must actually be intersected with the surface. Rather than intersect the ray with a bilinear patch defined by four neighboring points, the approximation presented by Coquillart and Gangnet[18] is used here. Normals are obtained from the bilinear patches and interpolated between patches to remove visual discontinuities at patch boundaries.

For the lighting model, the inherent blue color in the water contributes 30 percent and the reflected environment 70 percent. Incorporation of Fresnel's law into the reflectance model would probably enhance image quality. The images are adaptively antialiased with 16 extra rays per pixel in the neighborhood of large color discontinuities. Figures 1 and 9 were rendered at a resolution of 1217×811 pixels, antialiased, on a VAX 8600. Both images required approximately one hour of CPU time.

Summary

A technique for synthesizing wind-driven seas in computer-generated imagery was presented, based on an empirical model for wind-driven sea spectra. When rendered with a ray-tracing program, these synthesized scenes produce remarkably realistic images. The technique has been extended to generate animated wave motion sequences by processing the Fourier phase. Several schemes for modifying the Fourier phase were tried, and the scheme of shifting the dominant wave frequencies independent of shifting the entire wave ensemble proved to be useful in generating animated sequences for video presentation.

As a final point, the simplicity of synthesizing complex ocean data from a few parameters is especially attractive. One simply supplies a uniformly distributed random noise generating program with a random noise seed, a wind direction, and a wind speed. Animated sequences require only the additional specification of one or two phase shifting parameters. ■

Acknowledgments

We would like to thank D.C. Ghiglia for his many helpful hints and his advice for performing the signal processing. We would also like to thank two unknown reviewers for their constructive criticism and suggestion to apply the wave dispersion model. This work was performed at Sandia National Laboratories, Albuquerque, NM, a prime contractor to the US Department of Energy, under contract DE-AC04-76DP00789.

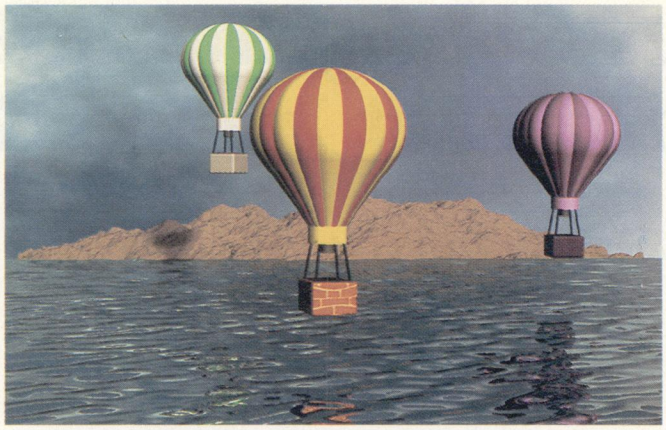

Figure 9. The ray-traced rendition of the Fourier-synthesized image in Figure 8. Several features have been added to make the scene more appealing.

References

1. B.B. Mandelbrot, *The Fractal Geometry of Nature*, W.H. Freeman, New York, 1983.
2. R.F. Voss, "Fourier Synthesis of Gaussian Fractals: 1f Noises, Landscapes, and Flakes," in *State of the Art in Image Synthesis*, Tutorial No. 10, SIGGRAPH 1983, ACM, New York, 1983.
3. A. Fournier, D. Fussell, and L. Carpenter, "Computer Rendering of Stochastic Models," *Comm. ACM*, June 1982, pp. 371-384.
4. N.L. Max, "Vectorized Procedural Models for Natural Terrain: Waves and Islands in the Sunset," *Computer Graphics* (Proc. SIGGRAPH 81), Aug. 1981, pp. 317-324.
5. B. Schachter, "Long Crested Wave Models for Gaussian Fields," *Computer Graphics and Image Processing*, Feb. 1980, pp. 187-201.
6. D.R. Peachey, "Modeling Waves and Surf," *Computer Graphics* (Proc. SIGGRAPH 86), Aug. 1986, pp. 65-74.
7. A. Fournier and W.T. Reeves, "A Simple Model of Ocean Waves," *Computer Graphics* (Proc. SIGGRAPH 86), Aug. 1986, pp. 75-84.
8. W.J. Pierson and L. Moskowitz, "A Proposed Spectral Form for Fully Developed Wind Seas Based on the Similarity Theory of S.A. Kilaigorodskii," *J. Geophysical Research*, Dec. 1964, pp. 5181-5190.
9. L. Moskowitz, "Estimates of the Power Spectrums for Fully Developed Seas for Wind Speeds of 20 to 40 Knots," *J. Geophysical Research*, Dec. 1964, pp. 5161-5179.
10. S.A. Hsu, "A Mechanism for the Increase of Wind Stress (Drag) Coefficient with Wind Speed over Water Surfaces: A Parametric Model," *J. Physical Oceanography*, Jan. 1986, pp. 144-150.
11. N.E. Huang et al., "A Study of the Relationship Among Wind Speed, Sea State, and the Drag Coefficient for a Developing Wave Field," *J. Geophysical Research*, June, 1986, pp. 7733-7742.
12. G.J. Komen, S. Hasselmann, and K. Hasselmann, "On the Existence of a Fully Developed Wind-Sea Spectrum," *J. Physical Oceanography*, Aug. 1984, pp. 1271-1285.
13. D.E. Hasselmann, M. Dunckel, and J.A. Ewing, "Directional Wave Spectra Observed during JONSWAP 1973," *J. Physical Oceanography*, Aug. 1980, pp. 1264-1280.
14. Digital Signal Processing Committee, IEEE Acoustics, Speech and Signal Processing Soc., ed., *Programs for Digital Signal Processing*, IEEE Press, New York, 1979.
15. A. Antoniou, *Digital Filters: Analysis and Design*, McGraw-Hill, New York, 1979, p. 139.
16. A. Sommerfeld, *Mechanics of Deformable Bodies, Lectures on Theoretical Physics*, Vol. II, Academic Press, New York, 1950.
17. J.T. Kajiya, "New Techniques for Ray Tracing Procedurally Defined Objects," *Computer Graphics* (Proc. SIGGRAPH 83), July 1983, pp. 265-274.
18. S. Coquillart and M. Gangnet, "Shaded Display of Digital Maps," *IEEE CG&A*, July 1984, pp. 35-42.

Biographies

Gary Mastin's biographical information may be found at the beginning of this issue in "About the Cover."

Peter Watterberg is a senior staff analyst at the Savannah River Laboratory, which is run by Dupont for the Department of Energy. He supports various graphics applications from finite element analysis to data plotting and presentation graphics. He was previously a member of the technical staff at Sandia National Laboratories where he served in a similar capacity. At every opportunity he spends time generating surreal computer images.

Watterberg received his BS degree in mathematics from the University of Illinois in 1977.

John Mareda is a staff member in the Applied Computer Graphics Division at Sandia National Laboratories in Albuquerque, New Mexico. His primary responsibility is providing system support for graphics software on all of Sandia's mainframe computer systems. He also has implemented a video animation system for use at the laboratories. He uses it in his spare time to create computer generated animations such as those described in this article.

Mareda received a BSEET from Missouri Institute of Technology in 1976 and an MSEE from the University of New Mexico in 1979.

The authors can be contacted through the Digital Image Processing Facility at Sandia National Laboratories, PO Box 5800, Organization 2643, Albuquerque, NM 87185.