FRACTAL MODULATION EFFECTS

Gianpaolo Evangelista

Sound Technology, Digital Media Division
Department of Science and Technology (ITN)
Linköping Institute of Technology (LiTH), Norrköping, Sweden
giaev@itn.liu.se

ABSTRACT

Fractal modulation is obtained by forming a power weighted superposition of scaled and modulated versions of the signal. The resulting signal is self-similar with fractal characteristics. In this paper we explore fractal modulation as a powerful method to generate rich signals, useful both for the synthesis of complex sounds, like the sounds from natural events or ecological sounds, or as control functions of audio effects. The wavelet transform can be used as an efficient tool in order to generate a subset of fractal modulated signals that are power homogeneous. Any signal used as a seed for fractal modulation is transformed into a multiscale sound by means of a tree-structured multirate filter bank. Moreover, by superimposing a structured modulation scheme one can generate pseudo-periodic sounds whose partials have fractal behavior.

1. INTRODUCTION

Introduced in the context of communication, fractal modulation allows for redundant transmission of the same signal at different scales, useful for robust detection in a fading channel scenario or for information hiding, since the characteristics of the fractal modulated signal have a tendency to be more noise-like than the original signal. In this paper we explore fractal modulation as a powerful method to generate rich signals, useful both for the synthesis of complex sounds, like the sounds from natural events or ecological sounds, or as control functions of audio effects.

The fractal modulation scheme is somehow inspired by the construction of the self-similar function known as the Weierstrass cosine

$$w(t) = \sum_{n=0}^{+\infty} \gamma^n \cos 2\pi a^n t \tag{1}$$

in which infinite scaled versions of the same cosine function are weighted by a power of γ and added together, where $0 < \gamma < 1 < a$ and $a\gamma \geq 1$. The result is a nowhere differentiable sound whose partials are further and further spaced away, with their frequencies exponentially growing. The Weierstrass cosine is approximately self-similar in the sense that

$$w(at) = \frac{1}{\gamma}\left(w(t) - \cos 2\pi t\right) \tag{2}$$

so that, time-scaling the function by a factor a obtains the same function $w(t)$, except for a smooth $\cos 2\pi t$ additional term and an amplitude scaling factor γ. The graph of the Weierstrass cosine, shown in Figure 1, is a fractal with box counting dimension $D = 2 + \frac{\log \gamma}{\log a}$. Notice that γ controls the fractal dimension: D approaches 2 when γ approaches 1 and decreases with γ.

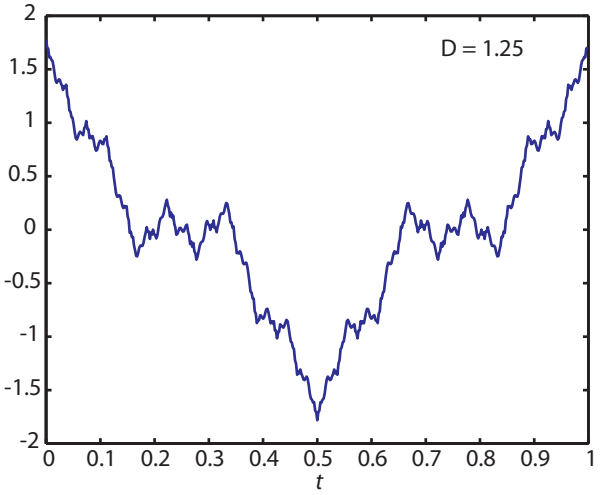

Figure 1: *Weierstrass cosine function with box counting dimension* $D = 1.25$.

2. FRACTAL MODULATION

Similar to the construction of the Weierstrass cosine, fractal modulation employs scaled versions of a signal in order to build a self-similar signal. Moreover, the scaled versions of the signal are modulated to proper band, which is implicit in scaling the cosine functions. The original signal can be considered as a "seed" for the generation of the fractal.

At least in a formal sense, one can construct self-similar signals by adding together a countable number of scaled and modulated versions of the seed. We assume that $a > 1$ and, usually, $1 < a \leq 2$. Suppose that our seed signal $x(t)$ is bandlimited to $[-(a-1)\pi, +(a-1)\pi]$, i.e., that its Fourier transform $X(\omega)$ is zero for ω outside this interval. We will use the notation $x(t) \in BL_{[A,B]}$ to specify bandlimited signals in the interval $[A, B]$.

By the Fourier scaling theorem, the scaled signal

$$x_n(t) = \frac{1}{a^n} x\left(\frac{t}{a^n}\right) \tag{3}$$

has Fourier transform

$$X_n(\omega) = X\left(a^n \omega\right) \tag{4}$$

and, therefore,

$$x_n(t) \in BL_{\left[-\frac{(a-1)\pi}{a^n}, +\frac{(a-1)\pi}{a^n}\right]}.$$

The peculiar modulation scheme we are going to employ is shown in Figure 2. The scaled signal is split into two subbands, one covering negative frequencies and the other one covering positive frequencies:

$$\begin{aligned} X_-(\omega) &= X(\omega)\chi_{]-(a-1)\pi,0]}(\omega) \\ X_+(\omega) &= X(\omega)\chi_{]0,+(a-1)\pi]}(\omega) \end{aligned} \tag{5}$$

where $\chi_{]A,B]}(\omega)$ is the characteristic function of the semiclosed interval $]A, B]$, i.e.

$$\chi_{]A,B]}(\omega) = \begin{cases} 1 & A < \omega \leqslant B \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

The negative frequency band $X_-(\omega)$ is modulated to positive frequencies while the positive frequency band $X_+(\omega)$ is modulated to negative frequencies. In this way one obtains, for each scaled and modulated version of the signal, a real signal whose effective bandwidth is the same as the baseband signal. The type of modulation is somehow arbitrary; the reason for this given choice will become apparent in the next Section. The choice of modulating frequencies

$$\omega_n^{\mp} = \pm\frac{\pi}{a^{n-1}} \tag{7}$$

guarantees that, ideally, the scaled and modulated signals do not overlap in the frequency domain. In fact, the scaled and modulated signals have disjoint Fourier transforms for any n:

$$\begin{aligned} X_-(a^n\omega - a\pi) &= X(a^n\omega - a\pi)\chi_{]+\frac{\pi}{a^n},+\frac{\pi}{a^{n-1}}]}(\omega) \\ X_+(a^n\omega + a\pi) &= X(a^n\omega + a\pi)\chi_{]-\frac{\pi}{a^{n-1}},-\frac{\pi}{a^n}]}(\omega) \end{aligned} \tag{8}$$

In the frequency domain, the self-similar signal $s(t)$ is obtained as a superposition of the signals in (8). For $\omega > 0$ we can define the Fourier transform of the self-similar signal as follows:

$$S(\omega) = \sum_{n=-\infty}^{+\infty} \beta^n X\left(a^n\omega - a\pi\right)\chi_{]\frac{\pi}{a^n},\frac{\pi}{a^{n-1}}]}(\omega), \tag{9}$$

where $\beta > 1$ is an arbitrary parameter. Actually β can be expressed in terms of a as follows:

$$\beta = a^\gamma; \;\; \gamma = \log_a \beta \tag{10}$$

For $\omega < 0$ the Fourier transform of the self-similar signal can be obtained from the relationship

$$S(-\omega) = S^\star(\omega), \tag{11}$$

where the symbol $\star$ denotes complex conjugation. Notice that since the terms of the sum in (9) do not overlap, for any given frequency there is only one nonzero term, hence the sum always converges, with a possible singularity at zero frequency if the original signal has a DC component.

The signal $s(t)$ obtained by taking the inverse Fourier transform of (9) is indeed self-similar with respect to scaling by powers of a. In fact, taking (10) into account and substituting $\frac{\omega}{a^K}$ for ω in (9), by performing a simple index change in the summation one obtains:

$$S\left(\frac{\omega}{a^K}\right) = a^{K\gamma}S\left(\omega\right). \tag{12}$$

for any $K \in \mathbf{Z}$. Therefore, in the time domain we have:

$$s\left(a^K t\right) = a^{K(\gamma-1)}s\left(t\right), \tag{13}$$

which is the required self-similarity property.

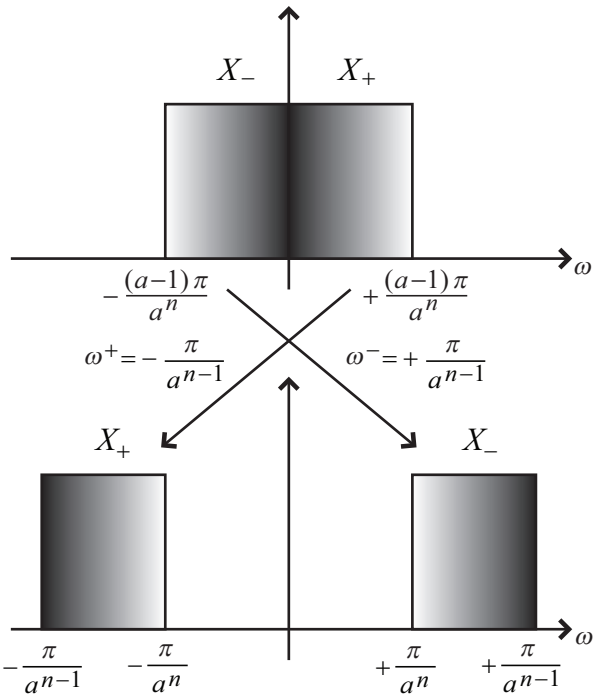

Figure 2: *Modulation scheme employed for the scaled signals.*

3. WAVELET BASED FRACTAL MODULATION

The frequency band subdivision of the scaling-modulation scheme to form self-similar signals illustrated in the previous Section is reminiscent of the structure of wavelet series. Indeed a wavelet based method for fractal modulation was introduced by Wornell and Oppenheim [1]. The simplest form of wavelet series [2, 3] is the expansion

$$s(t) = \sum_{n,m\in\mathbf{Z}} d_n(m)\psi_{n,m}(t) \tag{14}$$

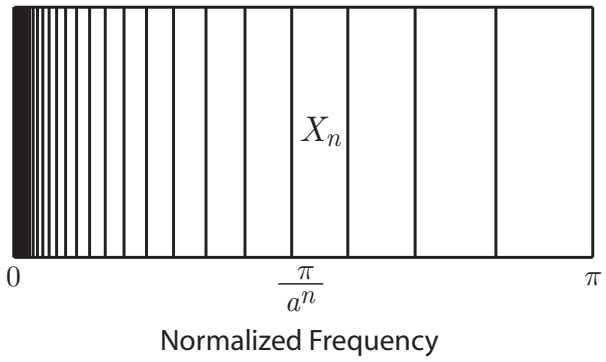

Figure 3: *Frequency band subdivision of the scaled and modulated versions of the signal.*

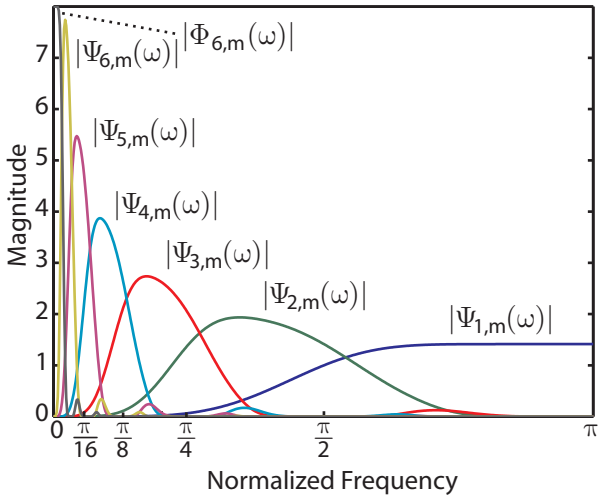

Figure 4: *Magnitude Fourier transform of dyadic wavelets.*

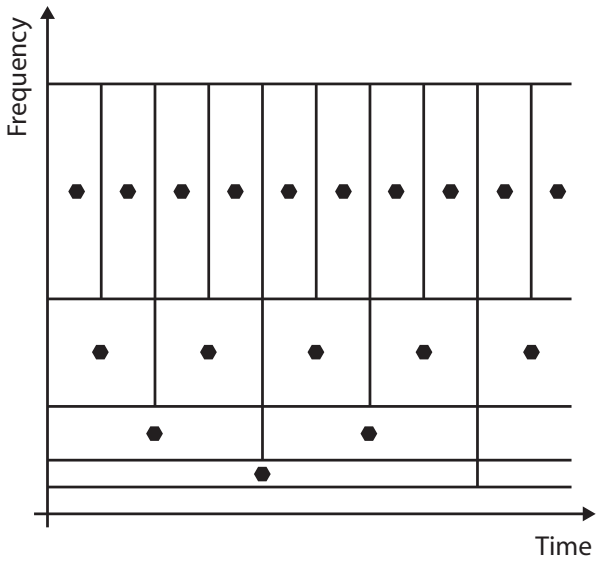

Figure 5: *Time-Frequency cells corresponding to dyadic wavelets.*

of a signal $s(t)$ in terms of a set of functions obtained by time-shifting and scaling a unique function, e.g., $\psi_{0,0}(t)$:

$$\psi_{n,m}(t) = \frac{1}{\sqrt{2^n}}\psi_{0,0}\left(\frac{t}{2^n} - m\right); \quad n, m \in \mathbf{Z}, \tag{15}$$

where the scaling factor is $a = 2$. As shown in Figure 4, the dyadic wavelets have bandwidth one octave, with bandwidth decreasing with the scale index n.

Theoretically the wavelets form the tessellation of the time-frequency plane shown in Figure 5, with higher frequency resolution at large scales and lower frequency resolution at small scales. In practice the cells overlap since the time-frequency uncertainty of the wavelets is larger than the time-frequency sample spacing.

For the orthogonal wavelet bases, the expansion coefficients are obtained by computing the scalar product of the signal with the basis functions $\psi_{n,m}(t)$:

$$d_n(m) = \int_{-\infty}^{+\infty} s(t)\psi^*_{n,m}(t)dt \tag{16}$$

Notice that, by performing the variable change $t = 2^n\tau$ in the integral in (16) and observing (15), one obtains:

$$d_n(m) = \sqrt{2^n}\int_{-\infty}^{+\infty} s\left(2^n\tau\right)\psi^*_{0,m}(\tau)d\tau \tag{17}$$

Suppose that the signal $s\left(t\right)$ is dyadically self-similar, i.e., that (13) is satisfied with $a = 2$, then:

$$d_n(m) = 2^{n\left(\gamma-\frac{1}{2}\right)}d_0(m), \tag{18}$$

i.e., the expansion coefficients of a dyadic self-similar signal on a dyadic wavelet basis are obtained by scaling the amplitude of one and the same coefficient sequence $d_0(m)$. Vice versa, if the amplitude scaled versions

$$2^{n\left(\gamma-\frac{1}{2}\right)}x(m)$$

of a unique discrete-time "seed" signal $x(m)$ are employed as wavelet coefficients, the self-similar signal

$$s(t) = \sum_{m\in\mathbf{Z}} x(m)\sum_{n\in\mathbf{Z}} 2^{n(\gamma-1)}\psi_{0,0}\left(\frac{t}{2^n} - m\right) \tag{19}$$

is obtained. In fact, one can immediately verify that $s(t)$ defined by (19) satisfies (13) with $a = 2$. Indeed, the functions

$$\xi_m(t) = \sum_{n\in\mathbf{Z}} 2^{n(\gamma-1)}\psi_{0,0}\left(\frac{t}{2^n} - m\right) = \sum_{n\in\mathbf{Z}} 2^{n\left(\gamma-\frac{1}{2}\right)}\psi_{n,m}\left(t\right)$$

in the linear combination (19) are self-similar [1].

Intuitively, the wavelet way to fractal modulation is very efficient since, rather than producing scaled and modulated versions of a signal, one pre-scales the known basis functions, an operation that is implicit in the organization of wavelet bases. Moreover, the wavelet series synthesis of a signal can be achieved by means of an iterated two-channel filter bank, where, in each stage, the inputs are upsampled by a factor 2 and a lowpass filter $H(z)$ and a highpass filter $G(z)$ forming a quadrature mirror filter pair [4] are applied. The diagram in Figure 6 serves as a fast generator of self-similar signals.

The computational complexity of the wavelet based fractal modulation scheme depends on the algorithm employed to compute the wavelet transform. However, as we will see, the overall cost is linear with the number of generated samples.

One method to compute fractal modulation is to pre-compute the wavelets at each scale and multiply each of them by the samples of the seed signal. The output is formed by properly time-shifting and adding the pre-multiplied wavelet components. We will refer to this algorithm as the overlap-add method. In this case the complexity can be estimated at N operations per output sample, where N is the number of scales. In fact, the length of the scale n wavelet can be roughly estimated at $2^{n-1}L$ samples obtained as a cascade of upsampling and filtering operators where the order of the filters $L-1$ is an odd integer for orthogonal wavelets.

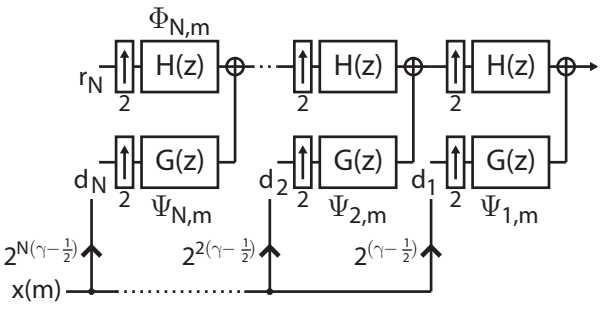

Figure 6: *Wavelet scheme for fractal modulation: an iterated quadrature mirror synthesis filter bank is fed by the same sequence $x(m)$ at any branch. The residue coefficients $r_N(m)$ are arbitrary.*

Thus, in order to generate an output signal of $2^{N-1}L$ samples one needs to multiply a single largest scale wavelet by one sample of the seed sequence; the two next to largest scale wavelets are each multiplied by one samples of the seed sequence and so forth. In general, computation at scale n requires multiplication of 2^{N-n} coefficients by the length $2^{n-1}L$ wavelets. This requires an order of $2^{N-1}L$ operations independently of scale. Therefore the total cost is $2^{N-1}LN$ operations in order to produce $2^{N-1}L$ output samples, which requires a rate of N operations per output sample.

A computational cost of the same order of magnitude is obtained when the scheme in Figure 6 is directly implemented. In this case one can show that the required rate of operations is proportional to the filter length L and does not depend on the number of scales N. It must be pointed out that, for sufficiently large filter orders, FFT based algorithms for the implementation of the filters can be employed to further reduce the overall computational cost.

Due to the finite number of scales obtained by a finite depth filter bank, the synthesized signal is only approximately self-similar. The residue coefficients $r_N(m)$, also known as the coefficients of the scaling function, are arbitrary. If a sufficient number of scales, corresponding to a sufficient number of stages of the filter bank, are employed then the scaling residue represents a very low-frequency signal and it can be disregarded in the synthesis of audio signals. However, in the next Section we will propose the use of self-similar signals also as low-frequency controlling functions of audio effects or as modulators of partials, in which case the scaling residue can be associated either to an amplitude scaled version of the seed signal or to an envelope function.

An important property of the wavelet based self-similar signal generator is found in the frequency domain:

$$S(\omega) = \sum_{n \in \mathbf{Z}} 2^{n\gamma} X\left(e^{j2^n\omega}\right) \Psi_{0,0}\left(2^n\omega\right), \qquad (20)$$

where $X\left(e^{j\omega}\right)$ denotes the DTFT of the seed signal $x(m)$. Equation (20) and Figure 7 show that the frequency spectrum of $x(m)$ is (periodically) contracted by a factor 2^n and directly weights the essential octave band of the corresponding wavelet at scale n, contributing to the frequency spectrum of the self-similar signal for the interval $\left[\frac{\pi}{2^n}, \frac{\pi}{2^{n-1}}\right]$ and its negative frequency mirror.

The type of crossed modulation illustrated in Section 2 emerges in wavelet based fractal modulation: the scaled negative frequency sideband of the seed signal contributes to the spectrum of the self-similar signal in the positive frequencies. Indeed, it is possible to

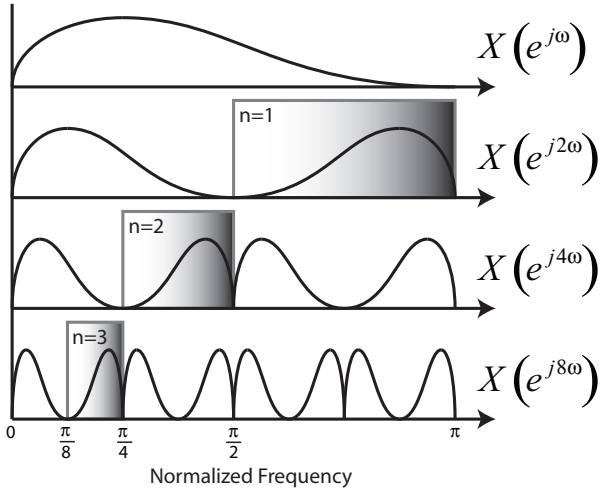

Figure 7: *Frequency spectrum weighting of the wavelet bands (rectangles) by the DTFT of the seed signal.*

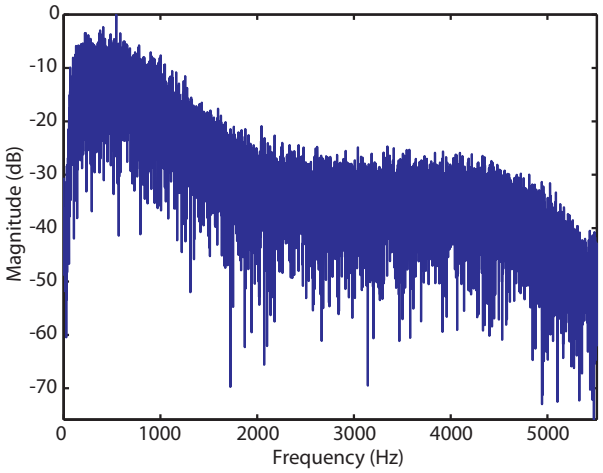

Figure 8: *Frequency spectrum of ocean wave impact sound.*

show that the *sinc* wavelets

$$\psi_{n,m}(t) = \frac{1}{\sqrt{2^n}} \cos\left(\frac{3\pi}{2}\left(\frac{t}{2^n} - m\right)\right) \operatorname{sinc}\left(\frac{1}{2}\left(\frac{t}{2^n} - m\right)\right),$$

which are based on an ideal rectangular frequency subdivision, form an orthogonal set of wavelets. However, the time localization of the ideal wavelets is very poor and a wide class of regular orthogonal wavelets exist which have a better time-frequency behavior [3].

4. AUDIO EFFECTS

Interesting and natural sound textures can be produced by means of fractal modulation. As we showed in the previous Section, the wavelet way to fractal modulation has a low computational cost, which is linear in the number of output samples. The sounds and effects obtained are so different from those produced by conventional, e.g., AM and FM, modulation techniques that an attempt to

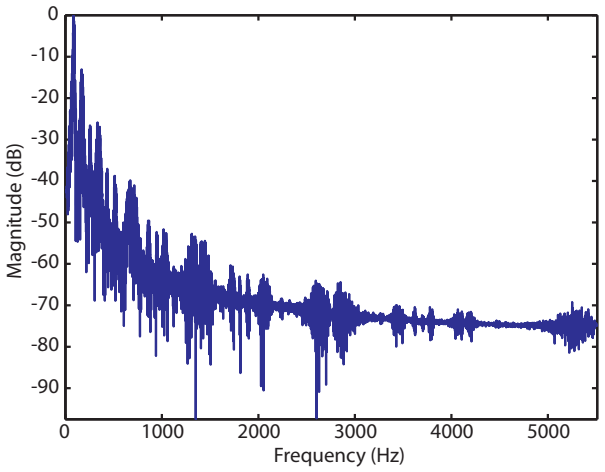

Figure 9: *Frequency spectrum of fractal modulation of the ocean wave impact sound, with $\gamma = 3$ and $N = 7$ scales.*

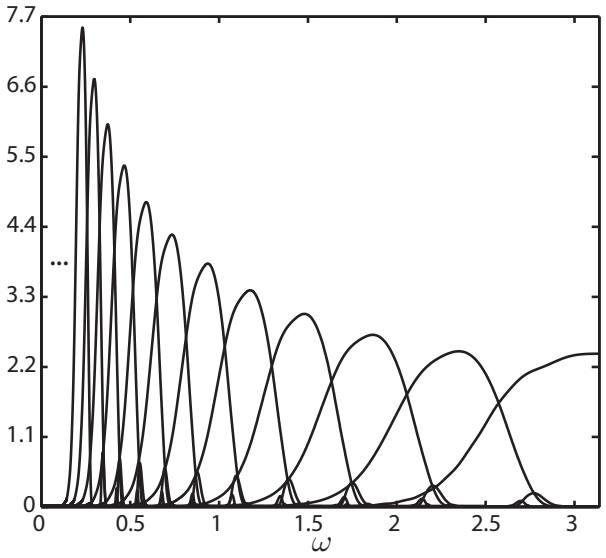

Figure 10: *Magnitude Fourier transform of frequency warped wavelets with 1/3 octave frequency resolution.*

compare them with fractal modulation is totally meaningless both from a computational and qualitative point of view.

4.1. Dense Textures

In the example of Figure 9 we applied the wavelet based fractal modulation scheme to the impact sound of an ocean wave whose frequency spectrum is reported in Figure 8. For the computation we used $N = 7$ dyadic scales with Daubechies wavelets of degree 11 [3]. The frequency spectrum of the transformed sound shows a highly structured organization, together with an underlying $1/f$ decay. The modulated sound has a quite apocalyptic flavor, with multivoiced inserts originating from the small scale components and with rhythmic and reverberating patterns originating from the large scale components.

The parameter γ controls the overall spectral decay: large values ($\gamma \gg 1$) move the spectral centroid to low frequencies, while small values of ($\gamma \approx 1$) tend to produce brighter sounds with flatter spectra, closer to white noise. The parameter γ is indeed inversely related to brightness, as it can be seen from equation (20). For large values of γ a much higher weight is associated to the large scale (low frequency) wavelets while for $\gamma < 1$ the small scale (higher frequency) wavelets take the lead.

Sound examples using the ocean wave sound as seed can be found at the URL: http://www.itn.liu.se/~giaev/soundexamples.html, where fractal modulation is applied for various values of the γ parameter, showing a corresponding shifting of the spectral center of mass of the sound.

4.2. Sparse Textures

The ocean wave sound has the structure of several fine scale events with an overall swooshing evolution. This type of structure is interesting for the seed to generate complex sounds emulating a large number of concurrent sources. On the other extreme, the application of fractal modulation to sounds with sparser textures, such as the cracking noise of fire, produces an augmented rhythmic texture that lies halfway from the original sound to the sound of dripping water, as it can be appreciated from the corresponding example available at the URL http://www.itn.liu.se/~giaev/soundexamples.html.

In sparse textures it is important or even critical to control the number of scales associated with fractal modulation. Limiting the number of scales achieves sounds where the original rhythmic pattern can still be distinguished, while a multiscale structure is superimposed. At the given URL we show patterns generated using $N = 7$ and $N = 3$ scales, which have a different flavor in terms of sharpness of attack and reverberation.

4.3. Extensions

A limitation of the wavelet method for fractal modulation lies in the fact that only dyadic self-similarity with $a = 2$ can be implemented. As a result, the frequency spectrum is organized by octave bands, which can be too coarse for spectrum modeling. However, orthogonal wavelets having rational scale factor a have been devised [5] based on rational sampling rate filter banks. Alternately, arbitrary scale wavelets, shown in Figure 10, have been defined based on iterated frequency warping schemes [6, 7, 8], which are suitable for richer fractal modulation.

Frequency warping can also be used in order to modify the perceived pitch, if any, of the sounds generated by fractal modulation. In an example that can be found at our URL, frequency warping is employed in order to play a short melody with the fractal modulated ocean wave, which renders the acoustics of a stadium-like crowd singing in an unorganized or spontaneous choir. An efficient approximated frequency warping algorithm can be found in [9], which allows for real-time computation.

4.4. Dynamic Fractal Modulation

The fractal modulation parameter γ can be dynamically and interactively changed by suitably scaling the seed signal. Since the sampling rate is different for each wavelet component, suitably downsampled amplitude envelopes must be considered. However,

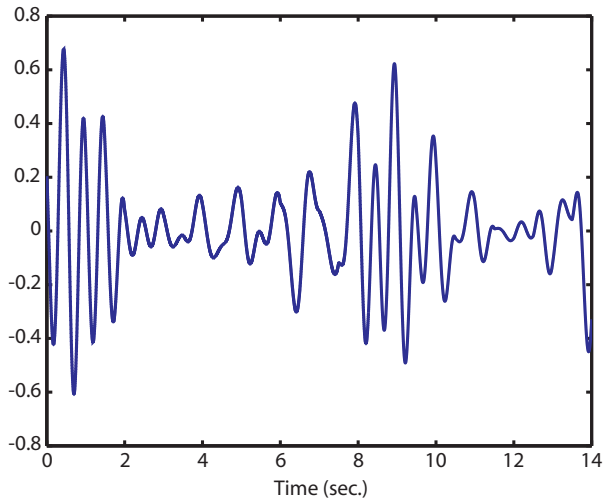

Figure 11: *Pseudo-sinusoid obtained by low-frequency cosine modulated fractal modulation.*

as the low-frequency components are generated by coefficient sequences at a very low sampling rate, excessive downsampling can result in audible distortion of the time envelopes. A more costly solution for higher fidelity time modulation of the γ parameter can be realized by amplitude scaling the synthesized wavelet components at the same sampling rate as the output signal. In the overlap-add method this can be achieved by the pre-computed shifted wavelets by the proper time-varying factor $2^n\gamma(k)$, where $\gamma(k)$ is the time-varying fractal modulation parameter. Interesting brightness transition effects can be realized in his fashion.

4.5. Fractal Modulated Low Frequency Oscillator

Modulated fractal noise has been employed in the Fractal Additive Synthesis technique [10, 11, 12] in order to model the $1/f$ behavior of the harmonic partials in the sounds of natural instruments. There, the seed signals are given by independent colored noise sources for each wavelet scale. Parallel wavelet synthesis structures are cascaded with a cosine modulated filter bank (MDCT) in order to synthesize the partials, where each wavelet section controls a single sideband of a partial. The same scheme but with deterministic signals replacing the colored noise sources can be employed in order to generate self-similar pseudo-periodic sounds.

Using only the lower frequency fraction of the MDCT, one can generate low-frequency pseudo-sinusoids with fractal fluctuations. These signals with a pseudo-random behavior can prove interesting as control functions of audio effects such as phaser, flanger, chorus, vibrato, e.g., by replacing the traditional LFO with a fractal LFO, whose oscillations are shown in Figure 11, where the main sinusoidal behavior, controlled by the scaling residue sequence $r_N(m)$ of Figure 6, is not included in the signal in order to better display the fluctuations. The introduction of a time-varying fractal modulation parameter allows us to dynamically control the depth of the effect, ranging from pure sinusoidal to extremely erratic behavior.

5. CONCLUSIONS

In this paper we explored the use of fractal modulation for the synthesis of complex sounds and of control functions to be employed in audio effects. The technique is very promising and efficient and it allows us to control complex behavior with a very small amount of parameters. Textured sounds with multiscale organization are typically generated by applying fractal modulation, where the small scale similar components produce voiced patterns and the large scale similar components produce rhythmic patterns.

While further experimentation is needed in order to assess the aesthetical value of the proposed modulation scheme in musical composition, simple examples serving the scope of appreciating the new sonority can be found at our URL.

6. REFERENCES

[1] G. W. Wornell and A. V. Oppenheim, "Wavelet-based representations for a class of self-similar signals with applications to fractal modulation," *IEEE Trans. Information Theory*, vol. 38, no. 2, pp. 785–800, Mar. 1992.

[2] S. Mallat, *A Wavelet Tour of Signal Processing*. Cambridge, MA: Academic Press, 1997.

[3] I. Daubechies, *Ten Lectures on Wavelets*, ser. CBMS-NSF Reg. Conf. Series in Applied Math. Philadelphia: SIAM, 1992, vol. 61.

[4] P. Vaidyanathan, *Multirate Systems and Filter Banks*. Upper Saddle River NJ: Prentice Hall Signal Processing Series, 1993.

[5] T. Blu, "A new design algorithm for two-band orthonormal rational filter banks and orthonormal rational wavelets," *IEEE Trans. Sig. Proc.*, vol. 46, no. 6, pp. 1494–1503, June 1998.

[6] G. Evangelista and S. Cavaliere, "Frequency warped filter banks and wavelet transform: A discrete-time approach via Laguerre expansions," *IEEE Trans. Sig. Proc.*, vol. 46, no. 10, pp. 2638–2650, Oct. 1998.

[7] ——, "Discrete frequency warped wavelets: Theory and applications," *IEEE Trans. Sig. Proc.*, vol. 46, no. 4, pp. 874–885, Apr. 1998, special issue on Theory and Applications of Filter Banks and Wavelets.

[8] G. Evangelista, "Dyadic warped wavelets," *Adv. in Imaging and Electron Physics*, vol. 117, pp. 73–171, Apr. 2001.

[9] G. Evangelista, I. Testa, and S. Cavaliere, "The role of frequency warping in physical models of sound sources," in *Proc. Forum Acusticum*, Budapest, Hungary, Aug. 2005, pp. 715–720.

[10] P. Polotti and G. Evangelista, "Analysis and synthesis of pseudo-periodic 1/f-like noise by means of wavelets with applications to digital audio," *EURASIP J. on Applied Sig. Proc.*, vol. 2001, no. 1, pp. 1–14, Mar. 2001.

[11] ——, "Fractal additive synthesis by means of harmonic-band wavelets," *Computer Music J.*, vol. 25, no. 3, pp. 22–37, 2001.

[12] P. Polotti, "Fractal additive synthesis," Ph.D. dissertation, EPFL, Lausanne, Switzerland, 2003. [Online]. Available: http://library.epfl.ch/theses/?nr=2711